# THE CHIPS TO SYSTEMS CONFERENCE

61

SHAPING THE NEXT GENERATION OF ELECTRONICS

**JUNE 23-27, 2024**

MOSCONE WEST CENTER
SAN FRANCISCO, CA, USA

THE CHIPS TO SYSTEMS CONFERENCE

SHAPING THE NEXT GENERATION OF ELECTRONICS

JUNE 23-27, 2024

MOSCONE WEST CENTER
SAN FRANCISCO, CA, USA

# Go Beyond Black-box Policies: Rethinking the Design of Learning Agent for Interpretable and Verifiable HVAC Control

Ryan (Zhiyu) An, Xianzhong Ding, Wan Du

University of California, Merced

# HVAC control in smart buildings

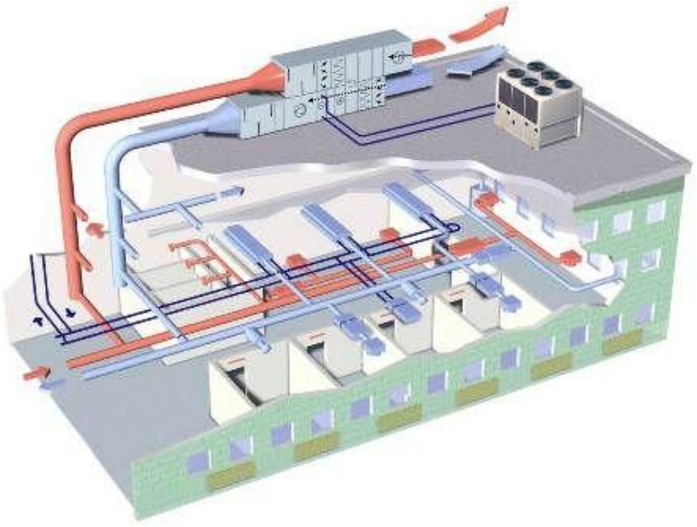

Heating, Ventilation, and
Air Conditioning
(HVAC) system

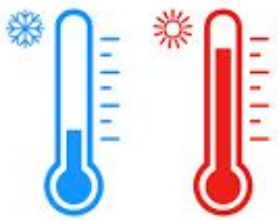

Objective #1:

Thermal Comfort

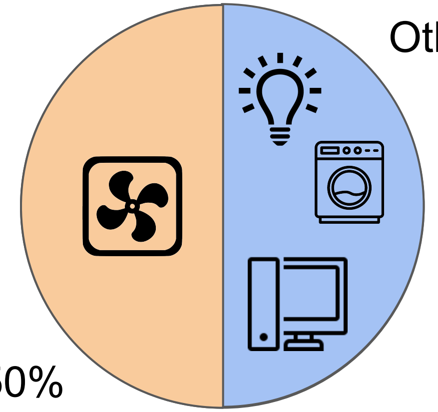


Objective #2:

Power Consumption

Higher

Comfort time/kWh

# HVAC control in smart buildings

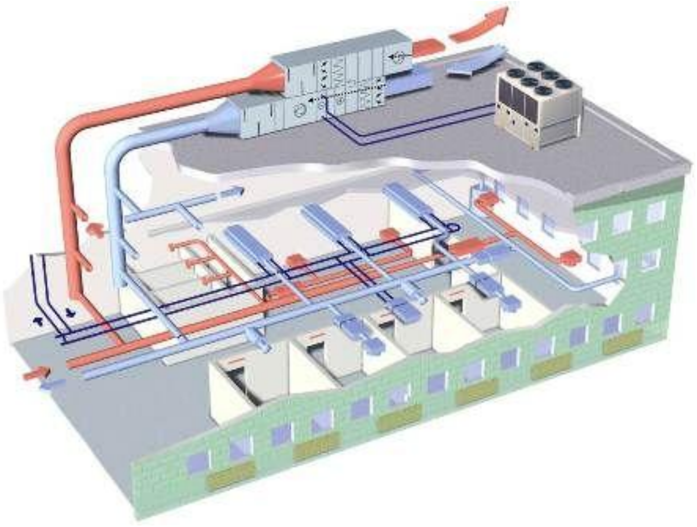

Heating, Ventilation, and
Air Conditioning
(HVAC) system

Rule-based decision-making

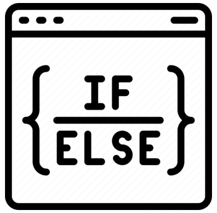


Energy Efficiency

baseline

# HVAC control in smart buildings

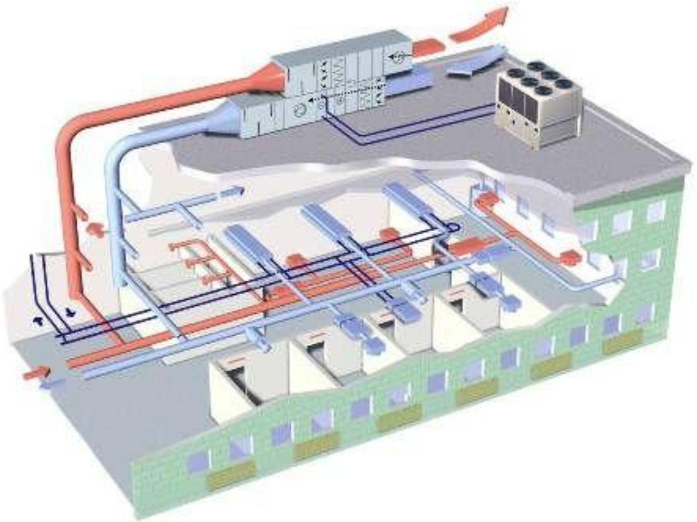

Heating, Ventilation, and
Air Conditioning
(HVAC) system

Rule-based decision-making

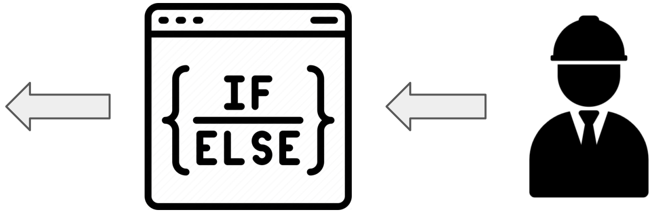


Data-driven predictive control agent

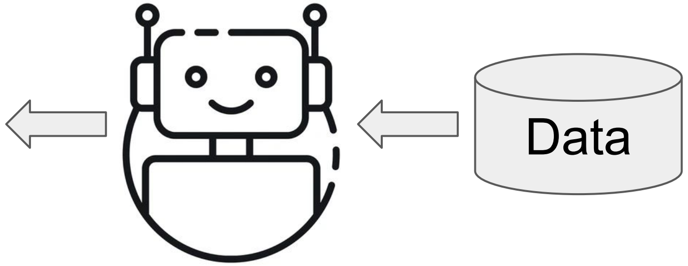


Energy Efficiency

baseline

10% less energy

# HVAC control in smart buildings

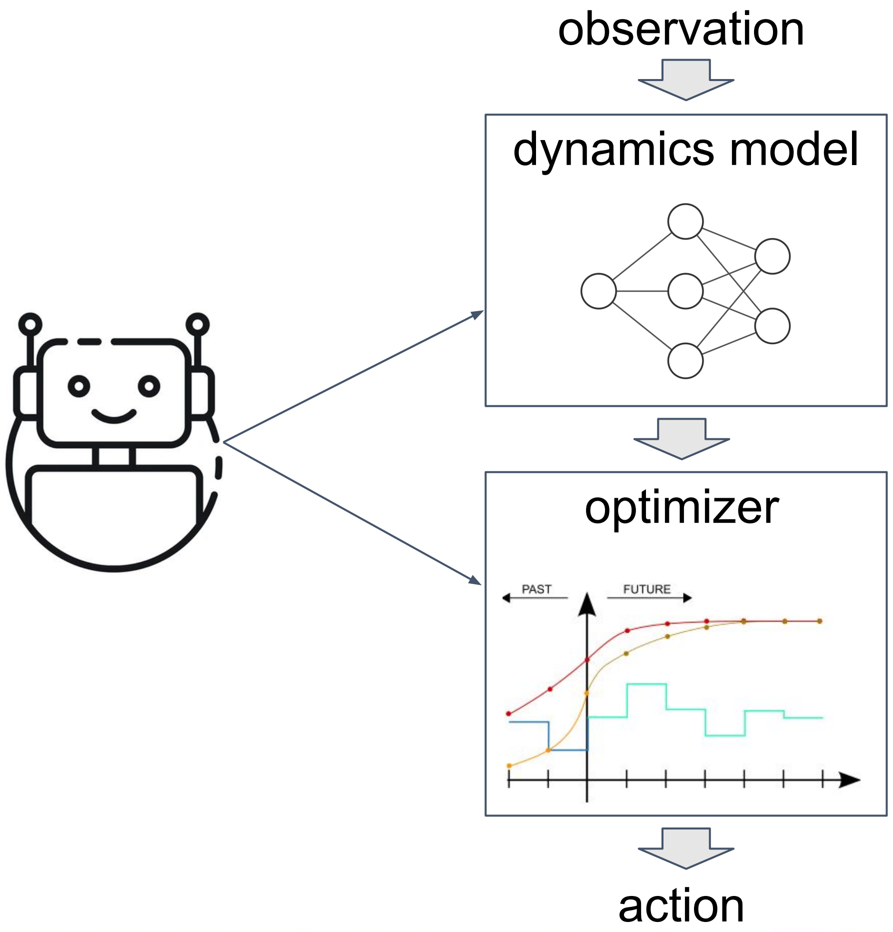

# The trustworthiness challenge of control agents

We want to guarantee that the 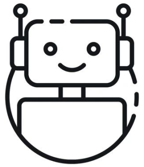 will always make decisions that are $\underbrace{\text{“safe”}}_{\text{interpretable \& verifiable}}$

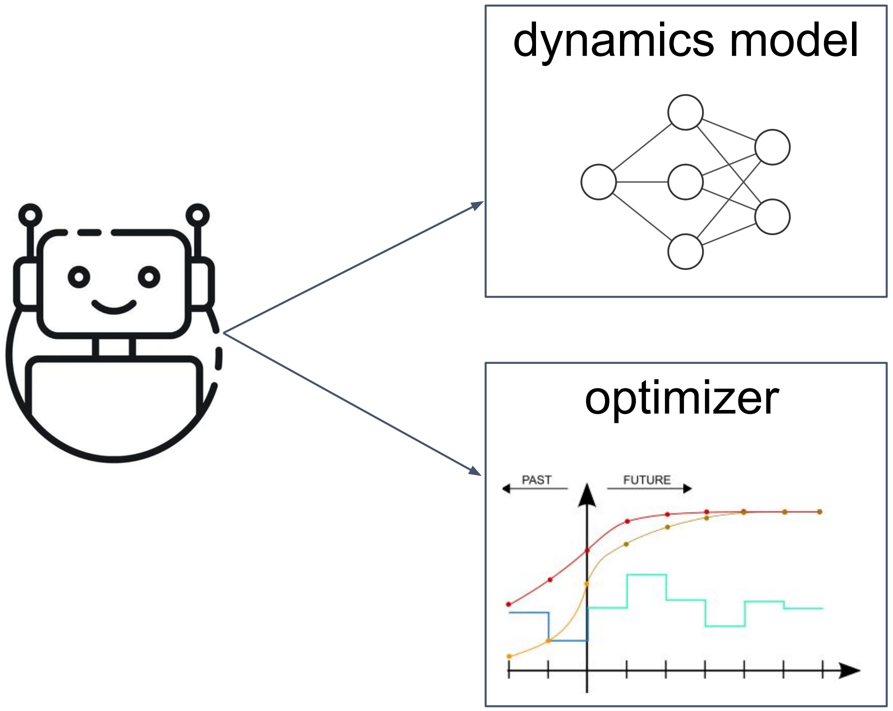

# The trustworthiness challenge of control agents

We want to guarantee that the 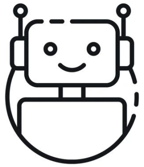 will always make decisions that are “safe”

interpretable & verifiable

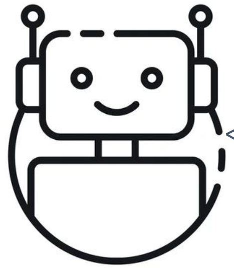

dynamics model

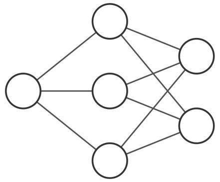

Not interpretable

(incomprehensible decision process)

optimizer

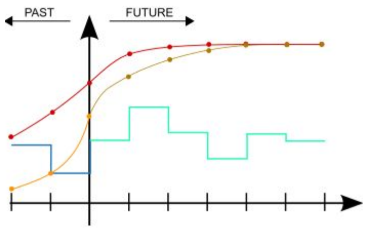


Stochastic

(chance of failure is always > 0)

Rethink the design of the agent!

# Rethinking the design of learning agent

We want to guarantee that the 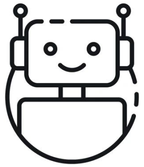 will always make decisions that are “safe”

interpretable & verifiable

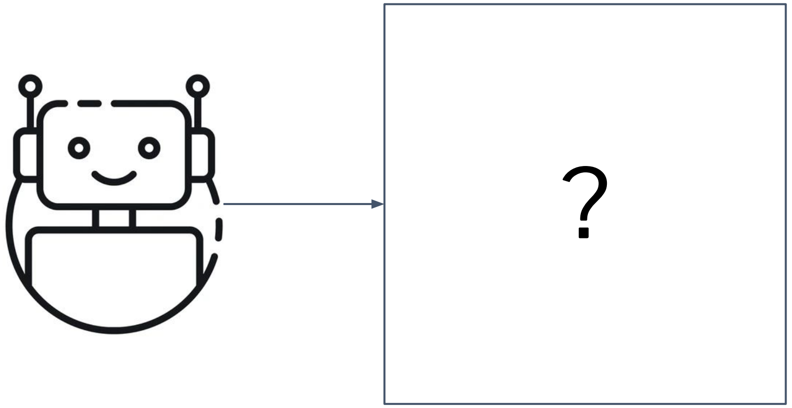


Interpretable

(human-readable decision process)

Deterministic

(enables safety verification)

# Rethinking the design of learning agent

We want to guarantee that the 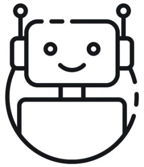 will always make decisions that are “safe”

interpretable & verifiable

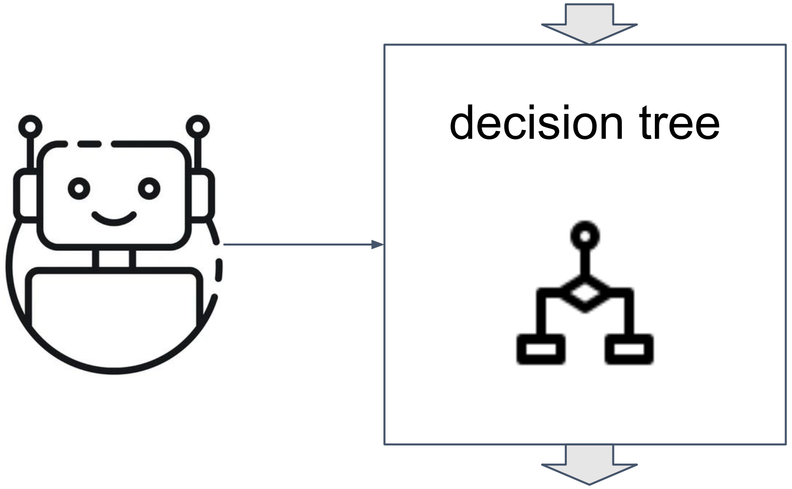


Interpretable

(human-readable decision process)

Deterministic

(enables safety verification)

# Learning the decision tree

New problem: how to generate an decision tree capable of high-efficiency decisions from data?

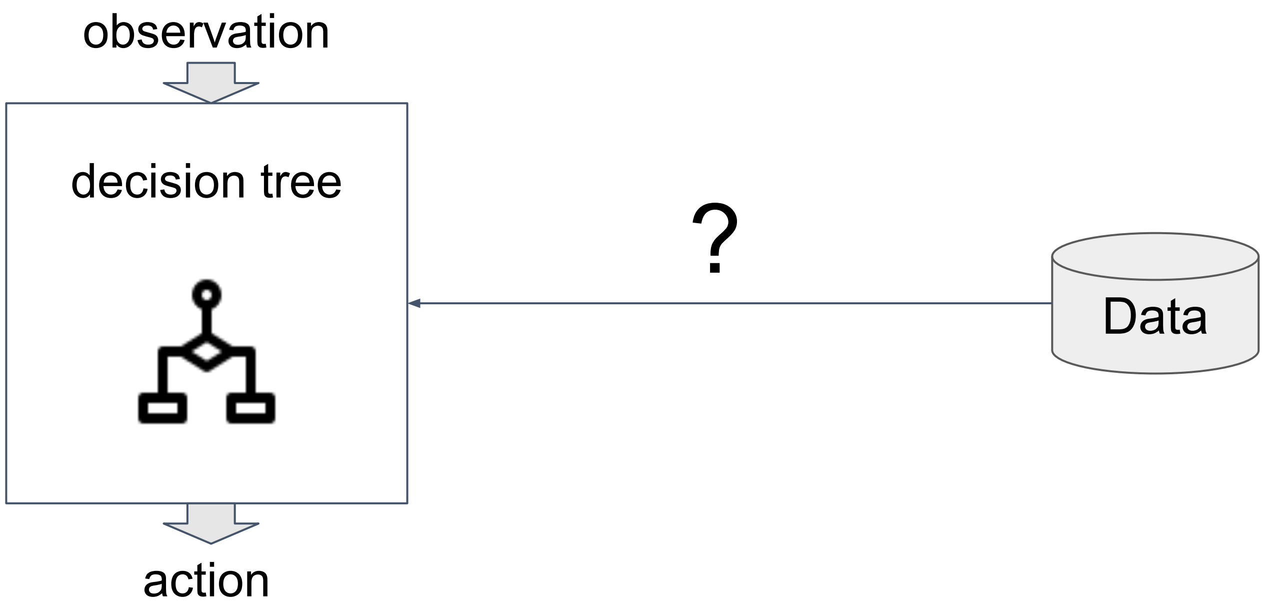

# Learning the decision tree

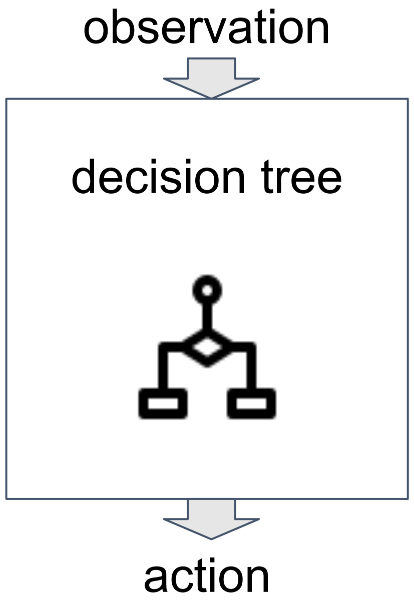


Use the original agent as the optimizer

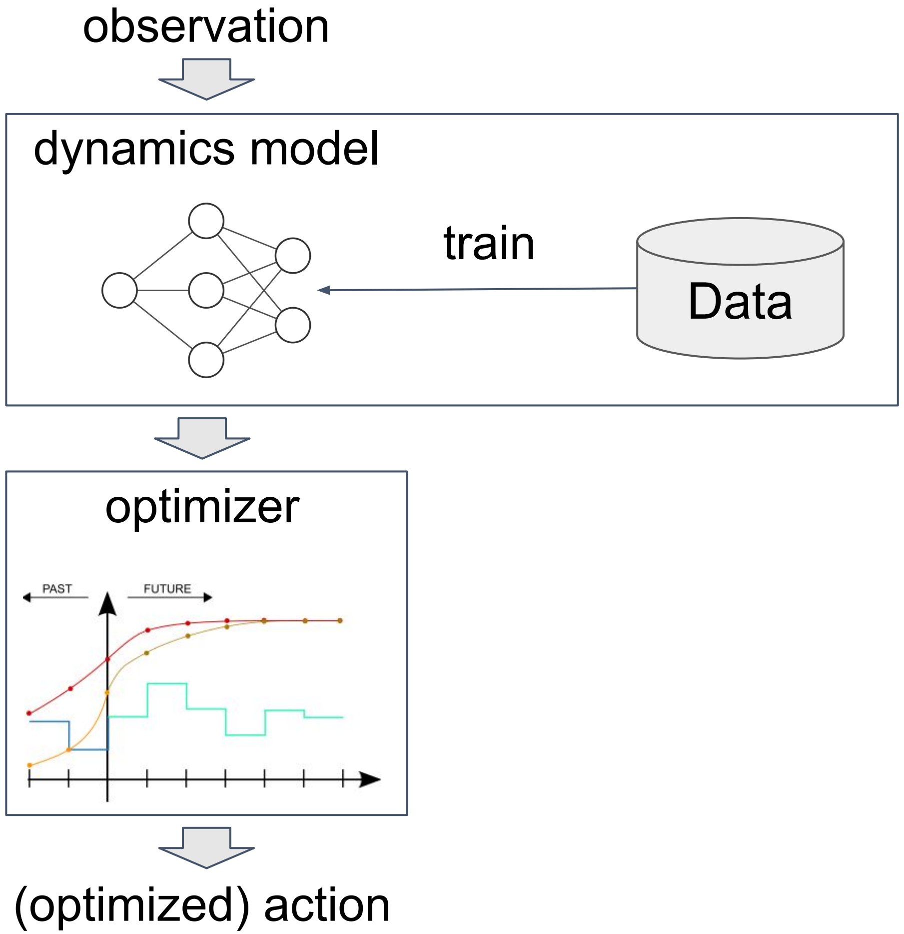

# Learning the decision tree

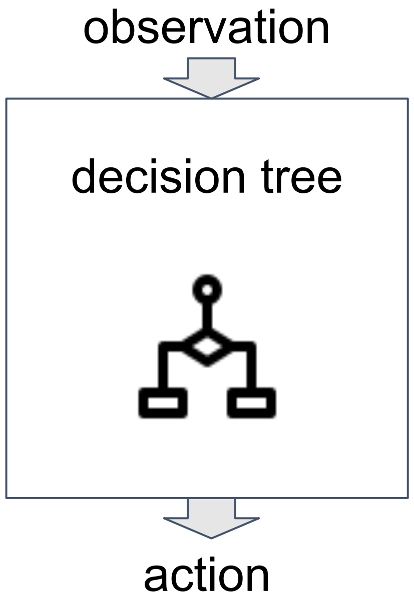


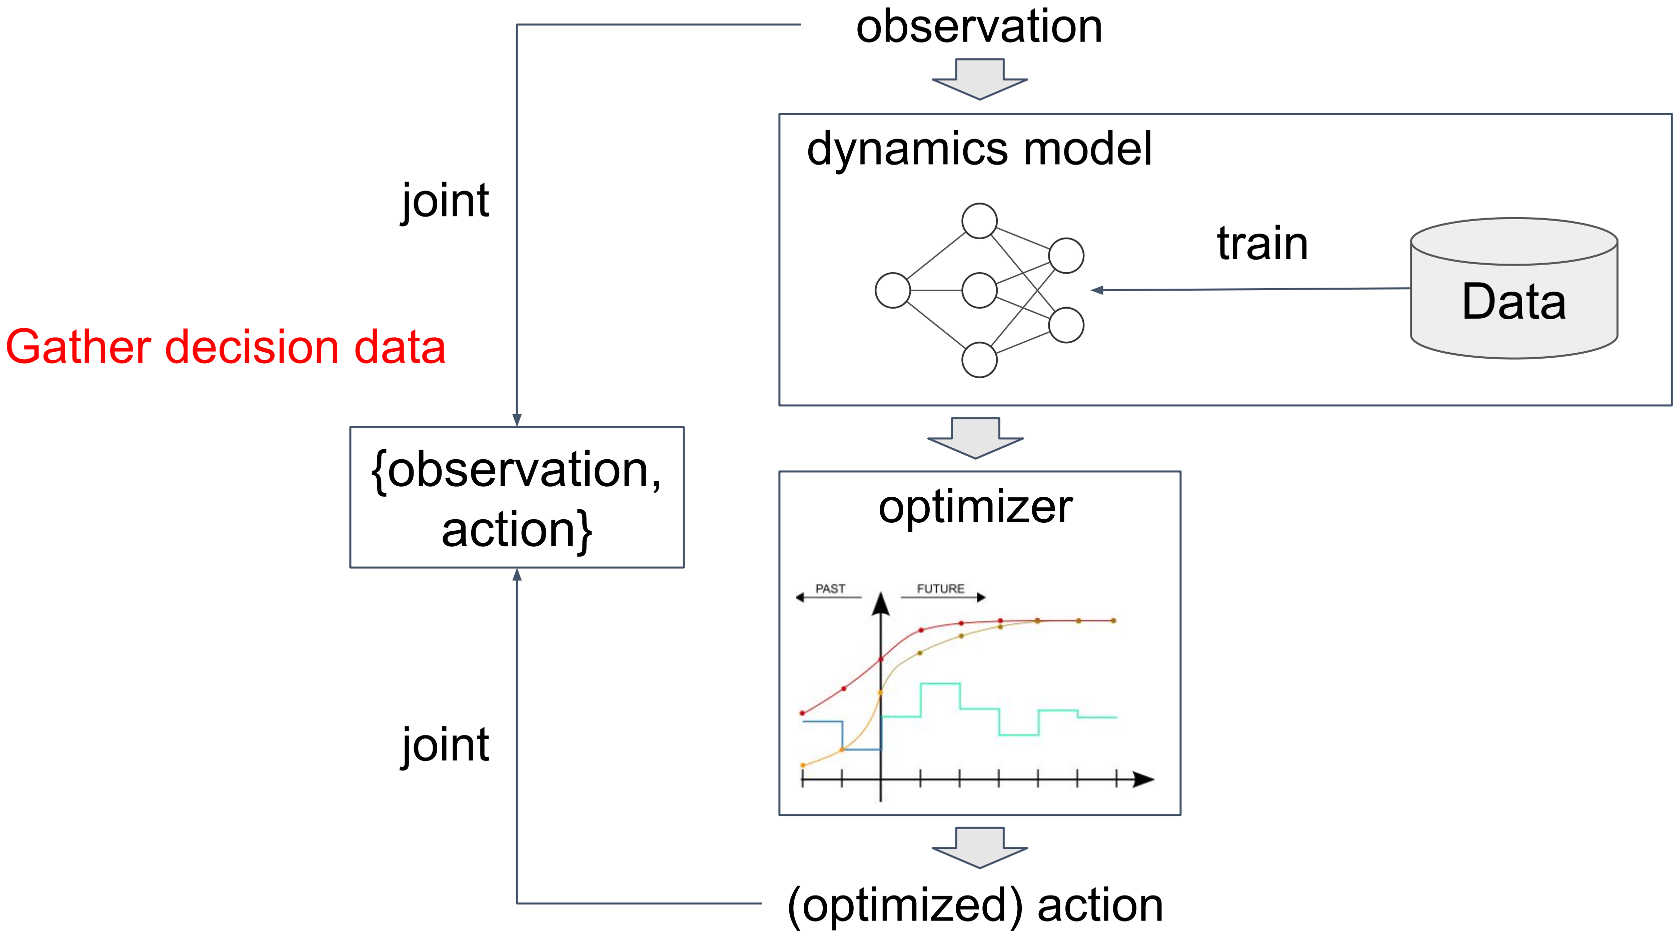

# Learning the decision tree

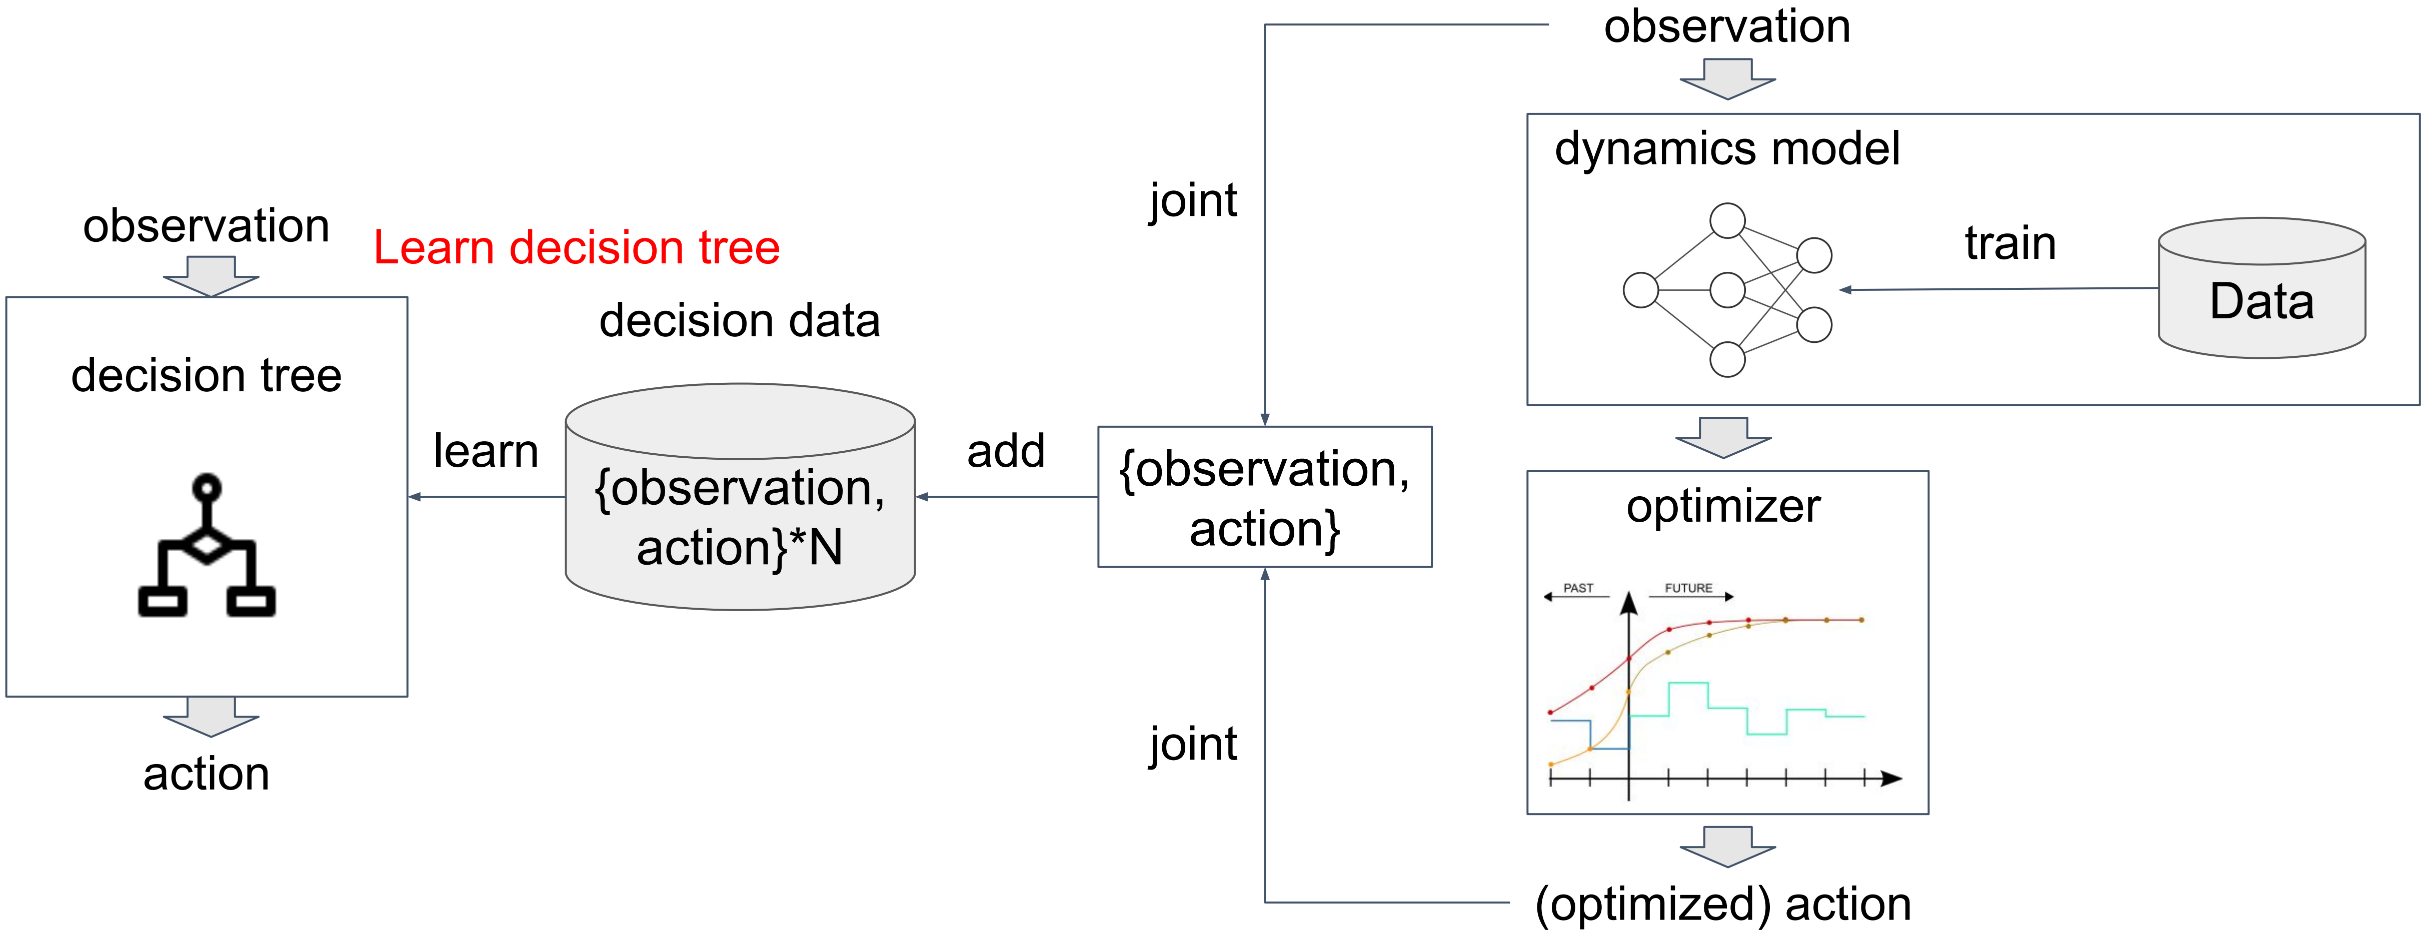

# Learning the decision tree

exhaust all combinations of observation can take prohibitively long time

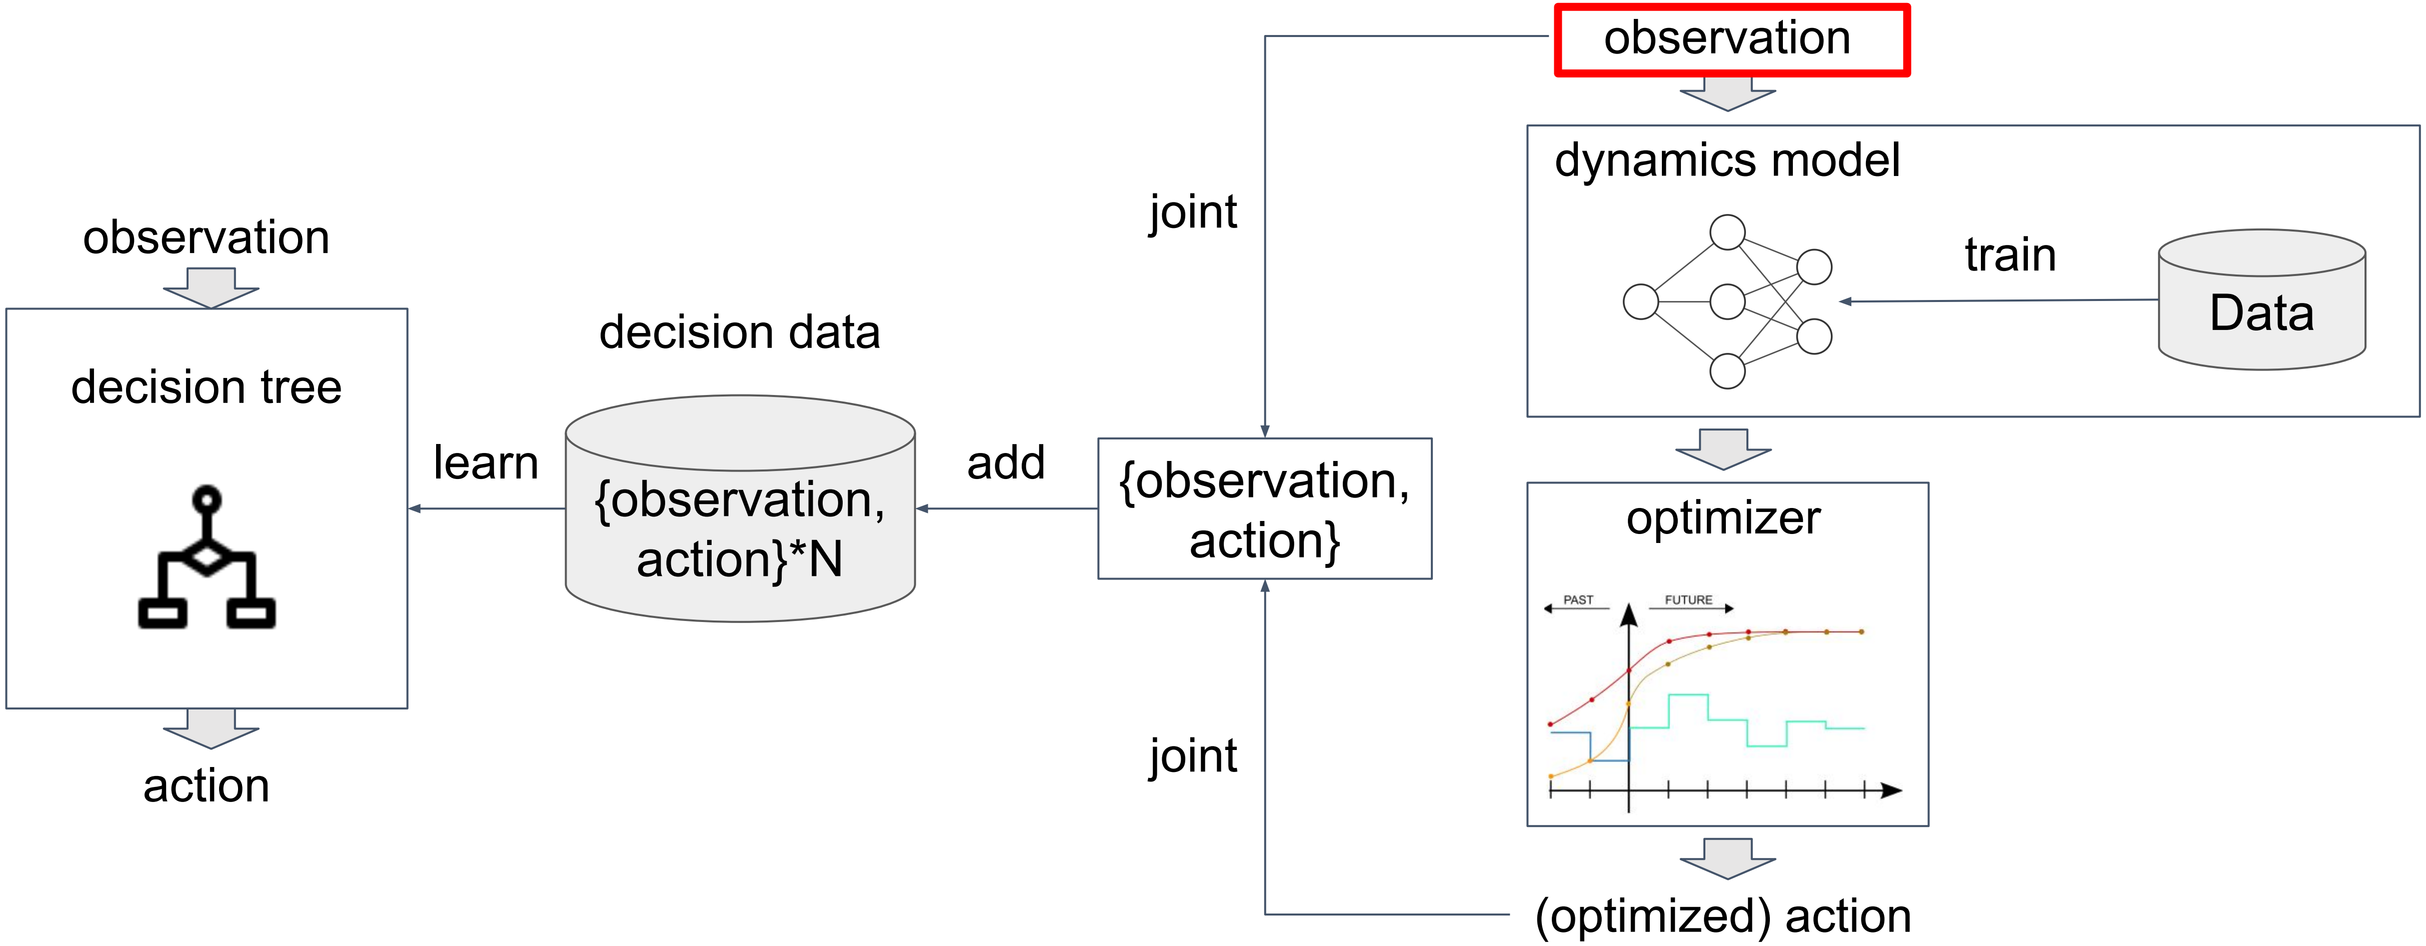

# Learning the decision tree

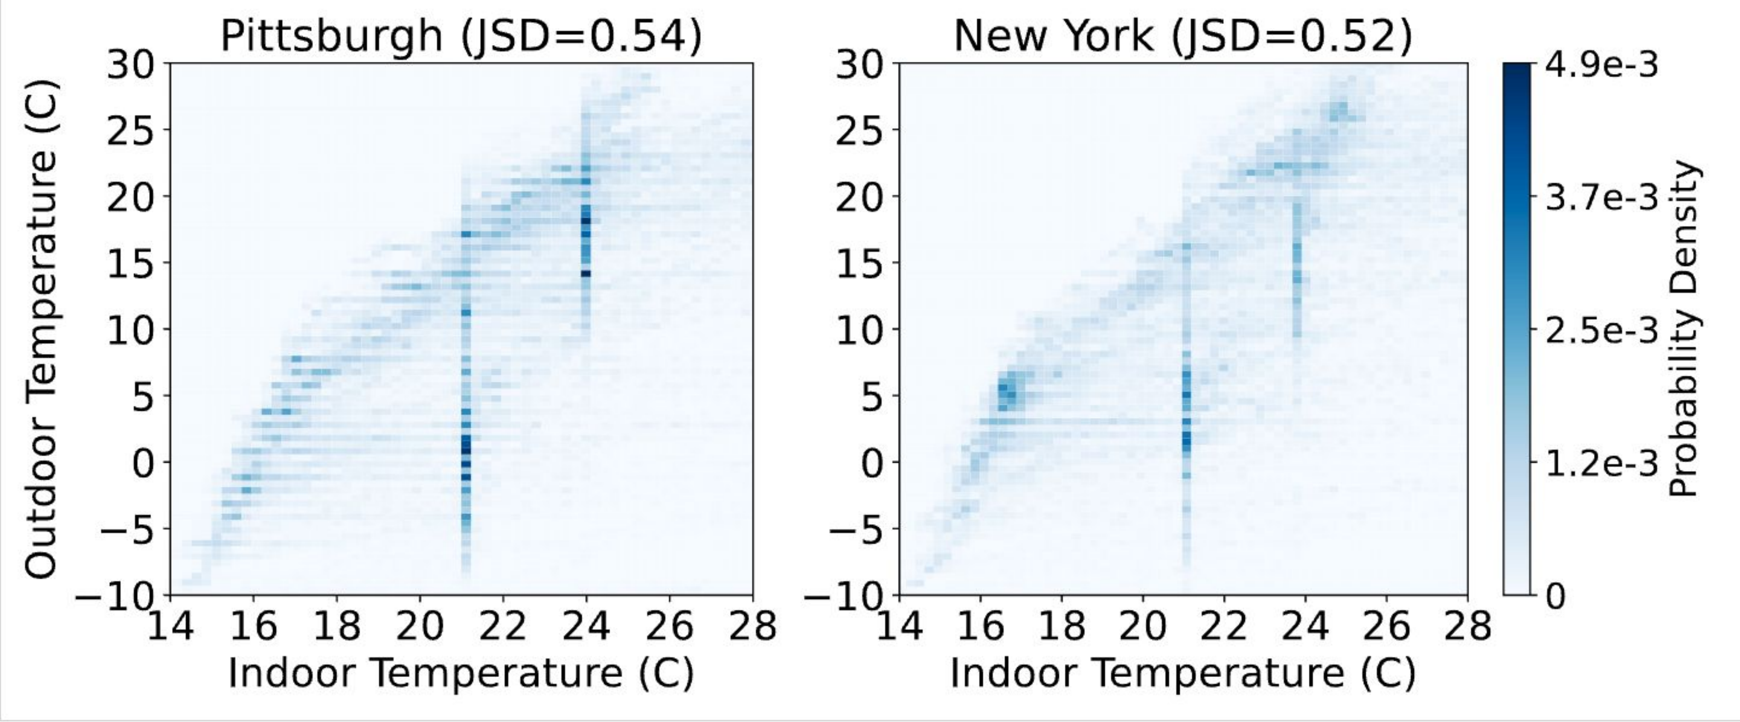


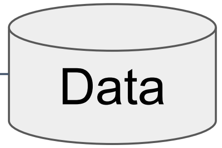

# Learning the decision tree

Insight: real-world HVAC data is heavily biased

➔ No need to exhaust all combinations; just follow the existing data distribution!

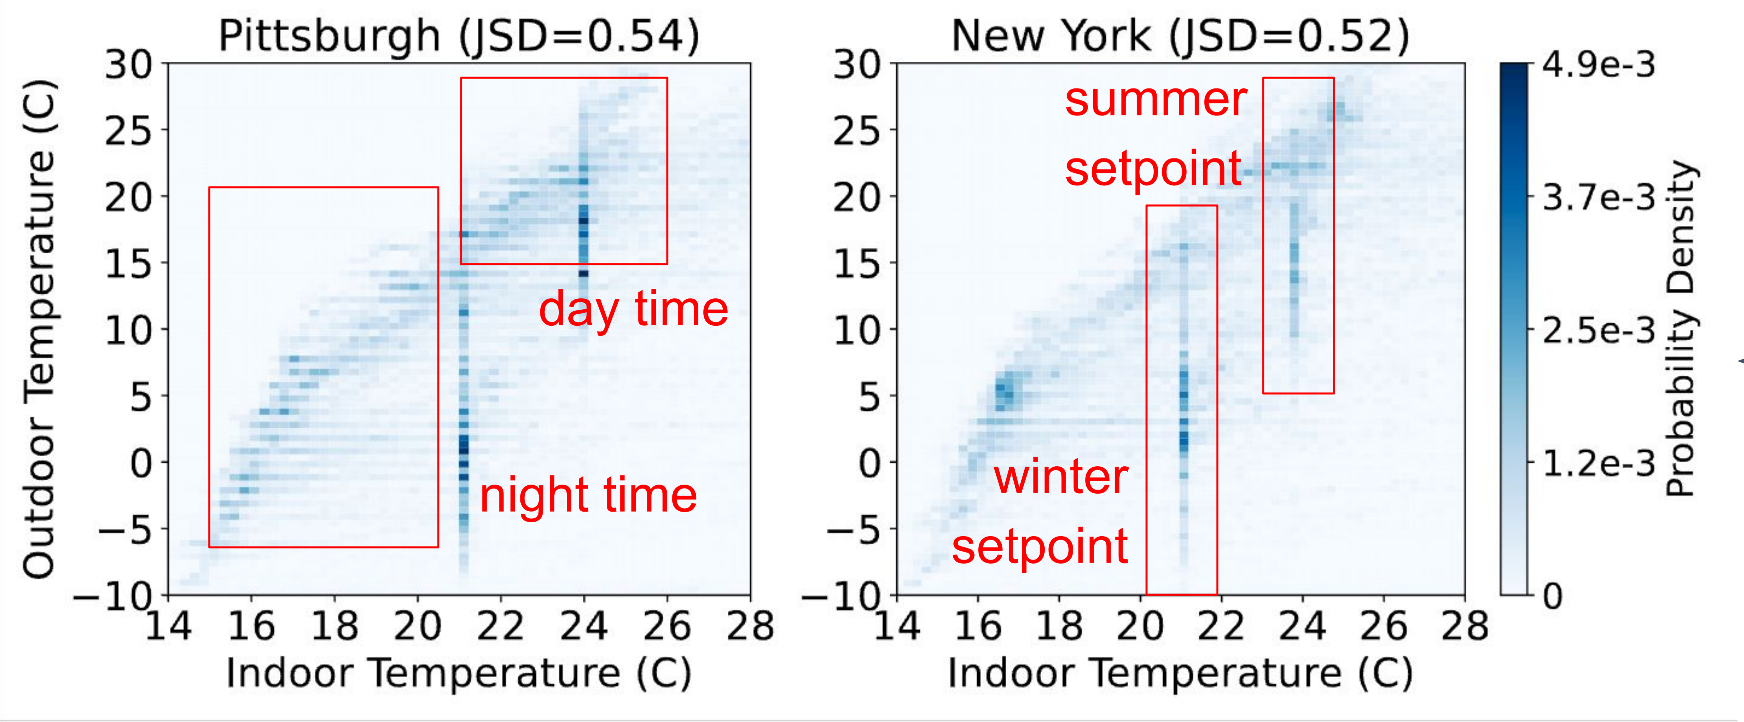


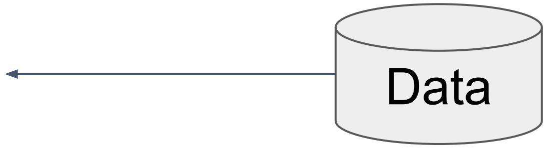

# Learning the decision tree

We use the following noise-addition function as the sampling strategy:

Sampling distribution

$$\widehat{p(x)} = X + \mathcal{N}\left(0, \text{noise\_level} \times \sqrt{\frac{\sum(x_i - \overline{x})^2}{|X|}}\right)$$

data distribution

# Learning the decision tree

We use the following noise-addition function as the sampling strategy:

Sampling distribution $\longrightarrow \widehat{p(x)} = X + \mathcal{N}\left(0, \boxed{\text{noise\_level}} \times \sqrt{\frac{\sum (x_i - \overline{x})^2}{|X|}}\right)$

data distribution

Noise_level is empirically determined to <u>maximize entropy</u> AND <u>minimize distance to the original data</u>:

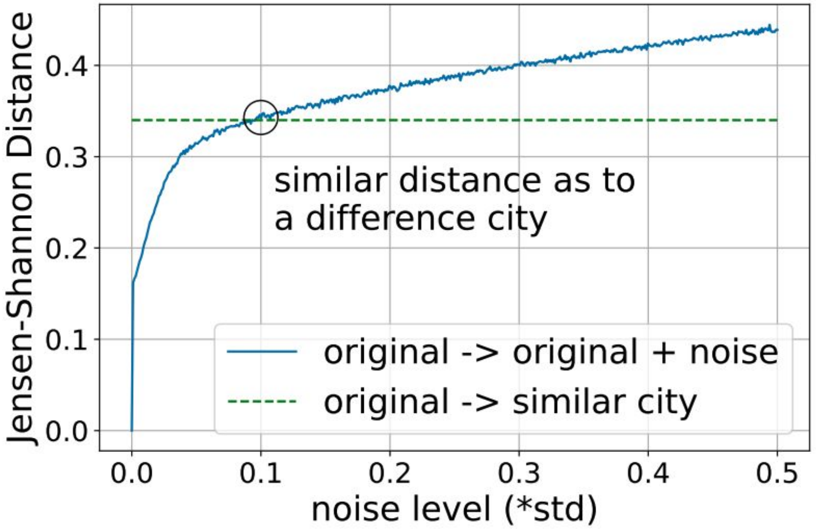


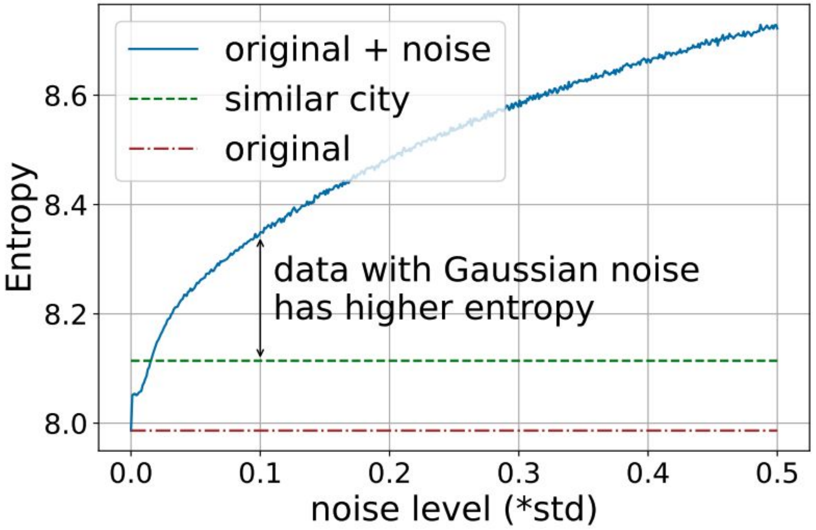

# Learning the decision tree

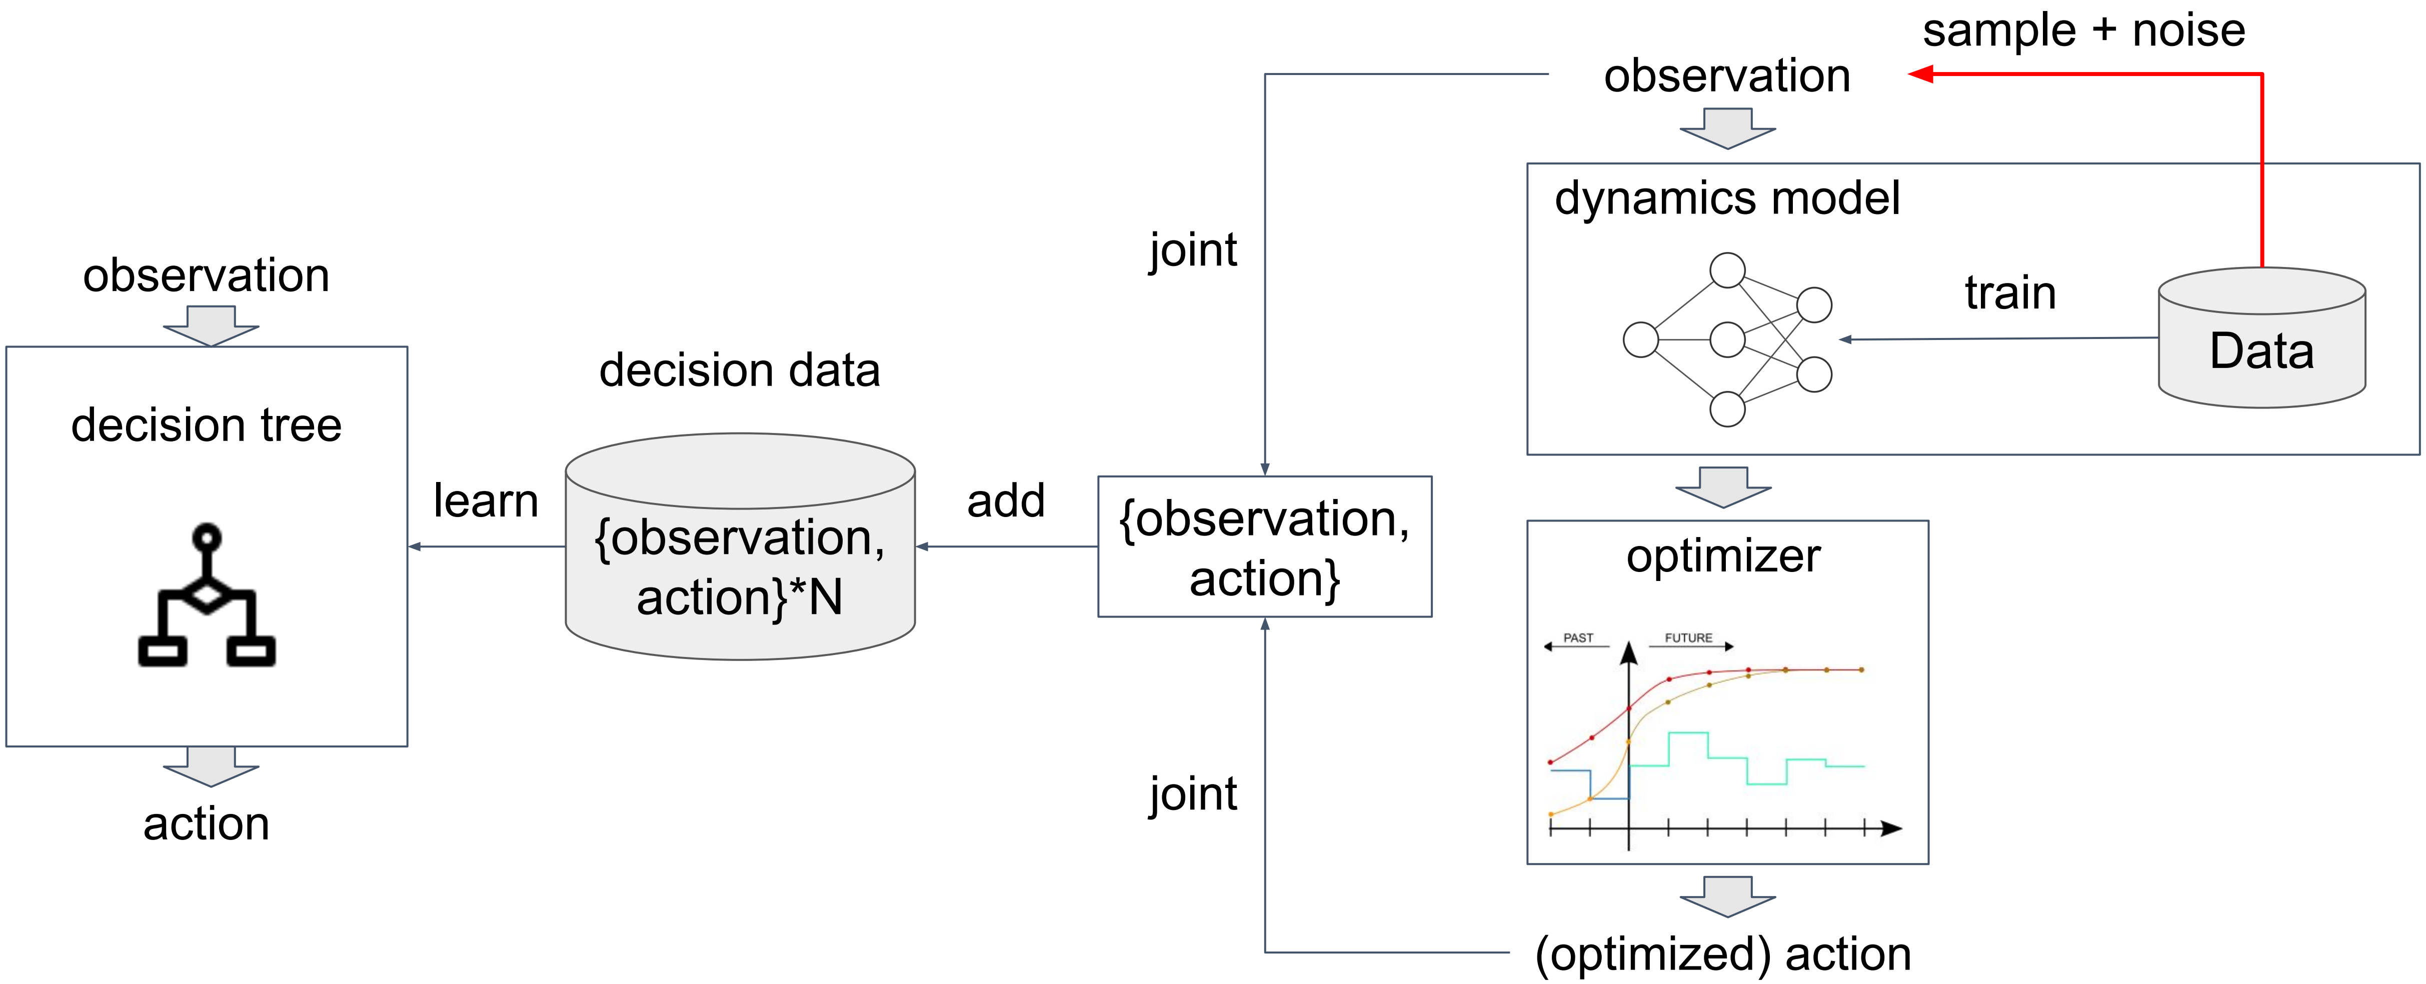

# Verifying the decision tree

How to verify the safety of a decision tree policy?

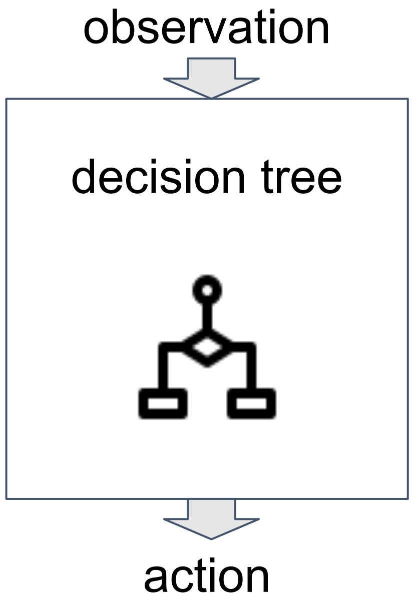

# Verifying the decision tree

How to verify the safety of a decision tree policy?

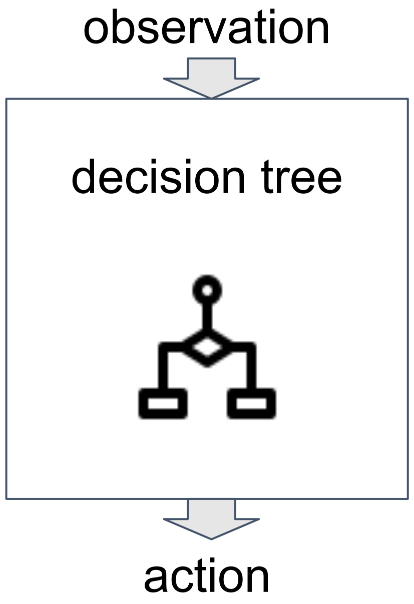


1. Incorporate human heuristic logical propositions between input and output

"If [input x] is in [certain range], then the action should be in [certain range]"

# Verifying the decision tree

How to verify the safety of a decision tree policy?

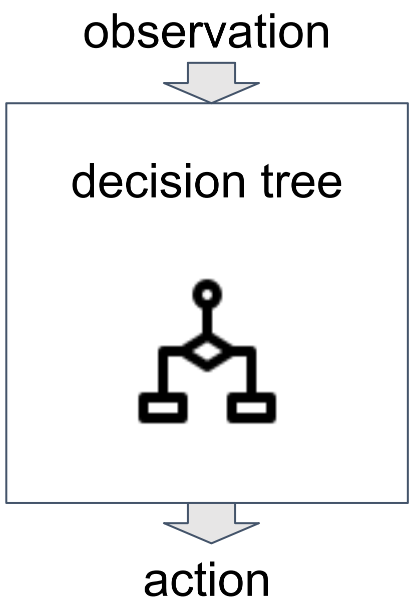


1. Incorporate human heuristic logical propositions between input and output

“If [input x] is in [certain range], then the action should be in [certain range]”

We consider two scenarios:

If the room is occupied and the temperature is cooler than the comfort range

➔ setpoint > current temperature (turn on heating)

If the room is occupied and the temperature is hotter than the comfort range

➔ setpoint < current temperature (turn on cooling)

# Verifying the decision tree

In decision tree, each decision node makes a comparison with one input parameter.
This split the input space with a hyperplane.

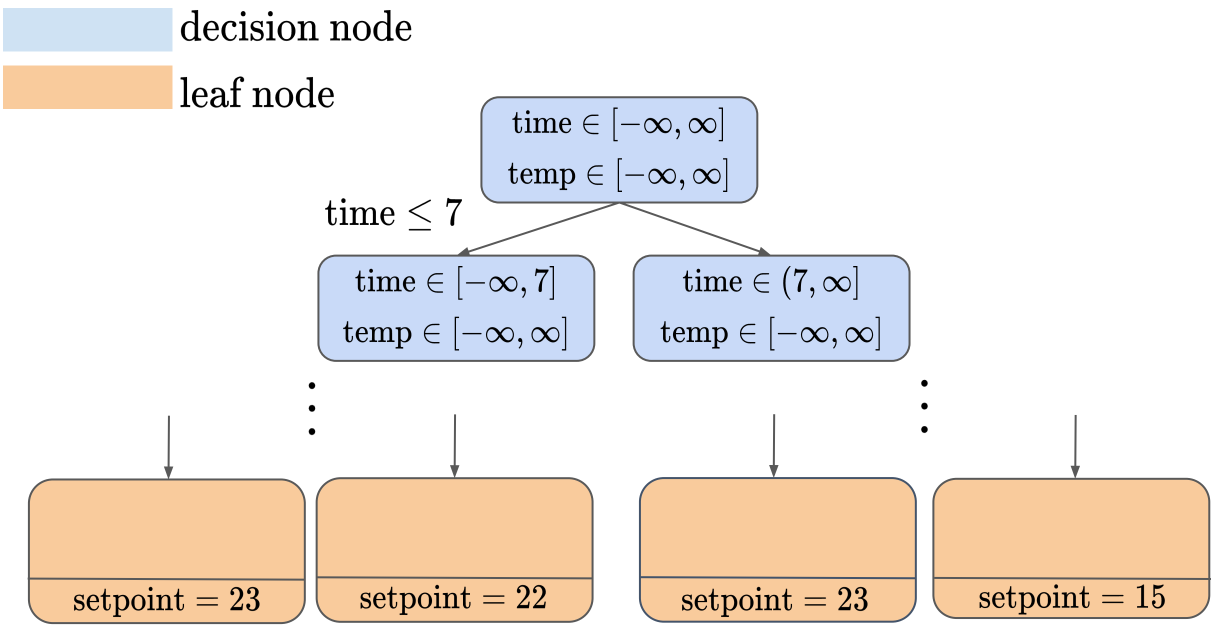


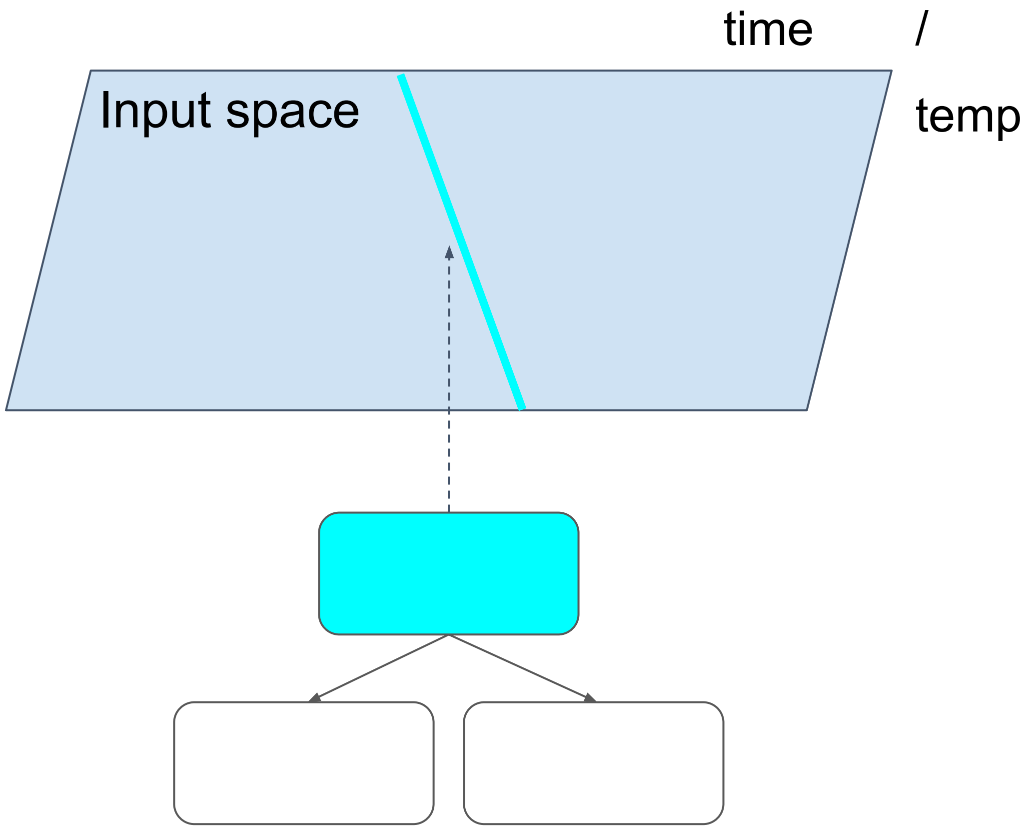

# Verifying the decision tree

Each leaf node corresponds to a region in the input space;
We can find this region for each leaf node

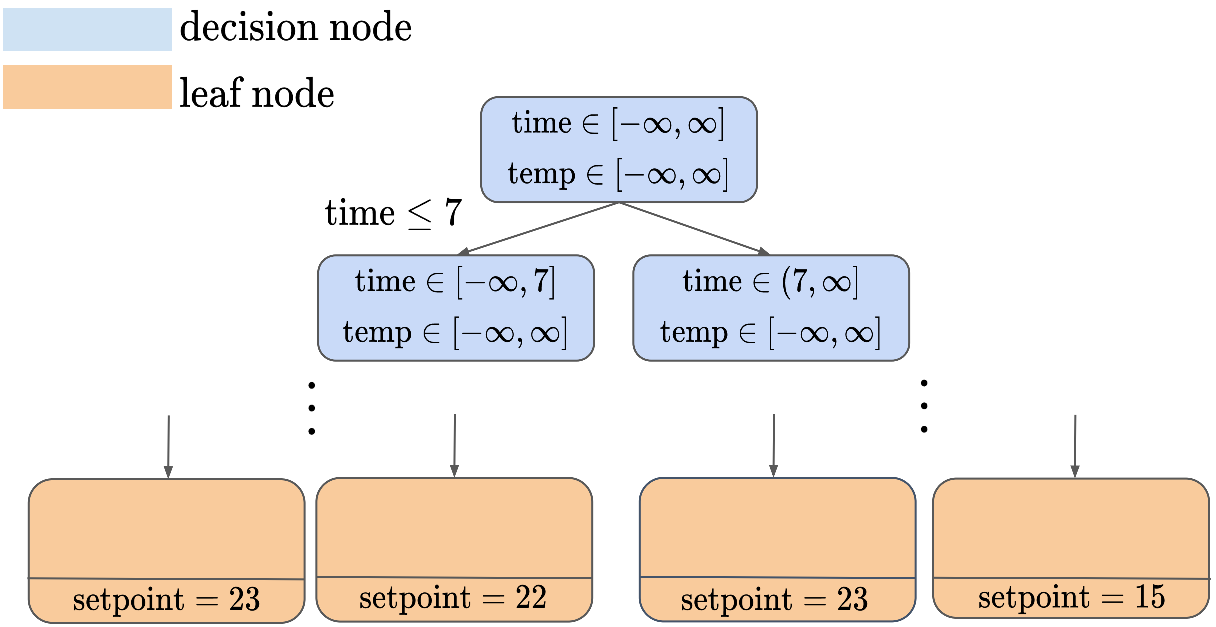


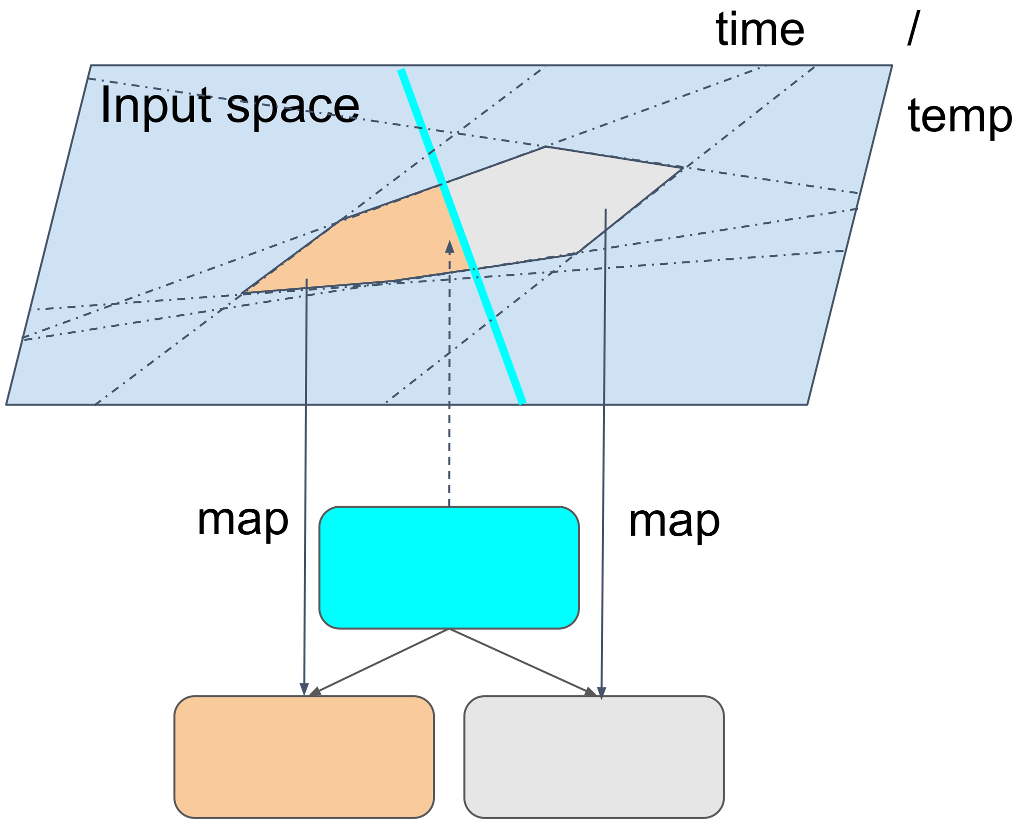

# Verifying the decision tree

Each leaf node corresponds to a region in the input space;
We can find this region for each leaf node

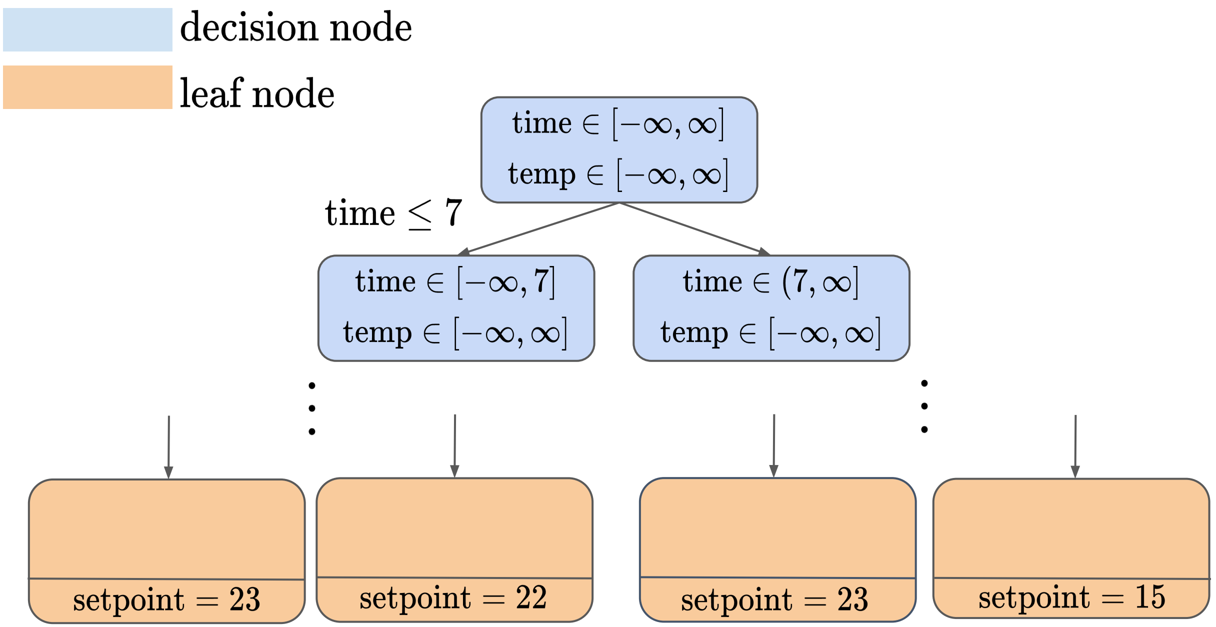


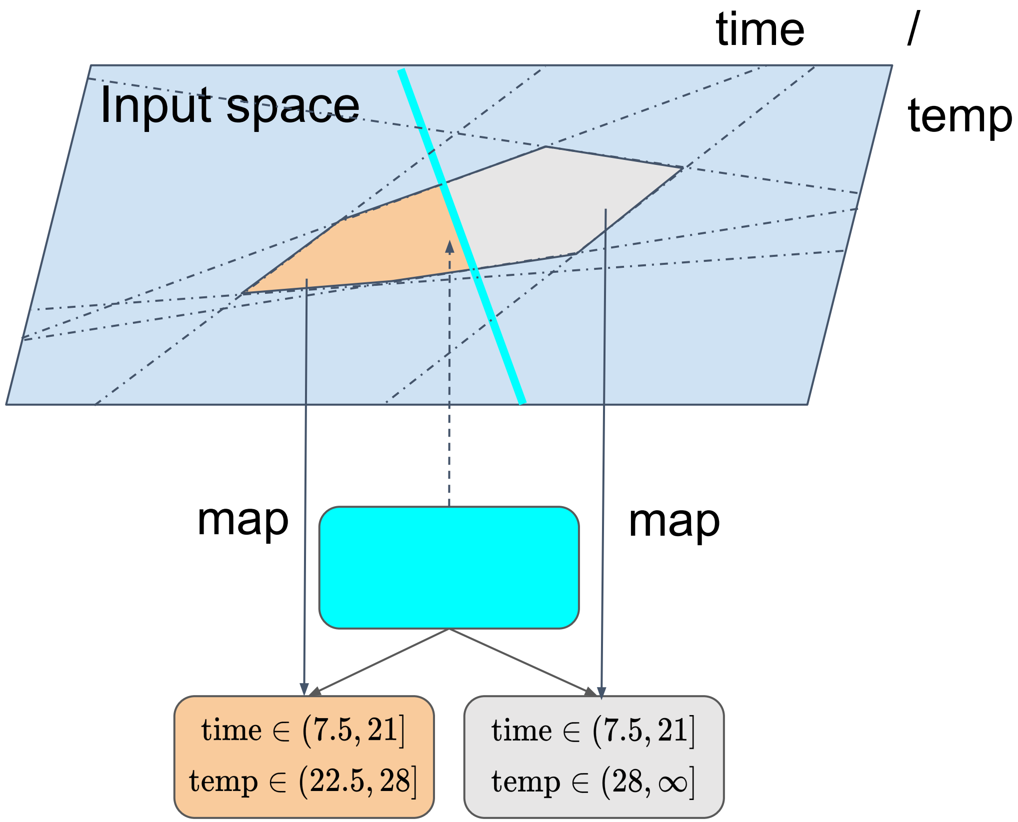

# Verifying the decision tree

Each leaf node corresponds to a region in the input space;

We can find this region for each leaf node

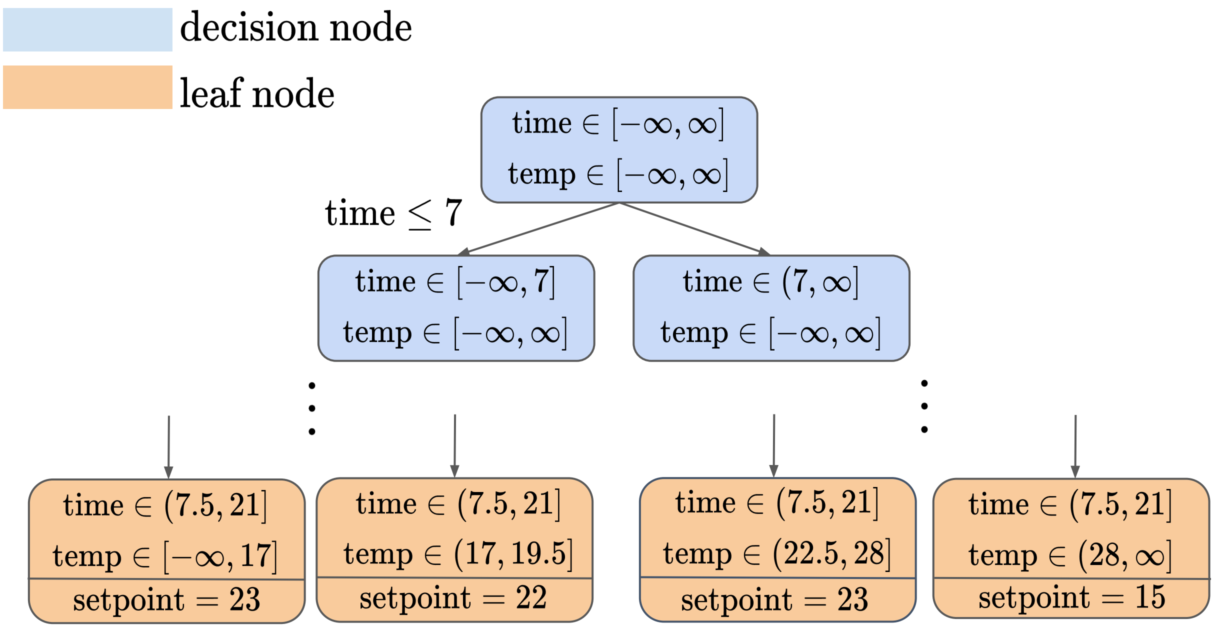


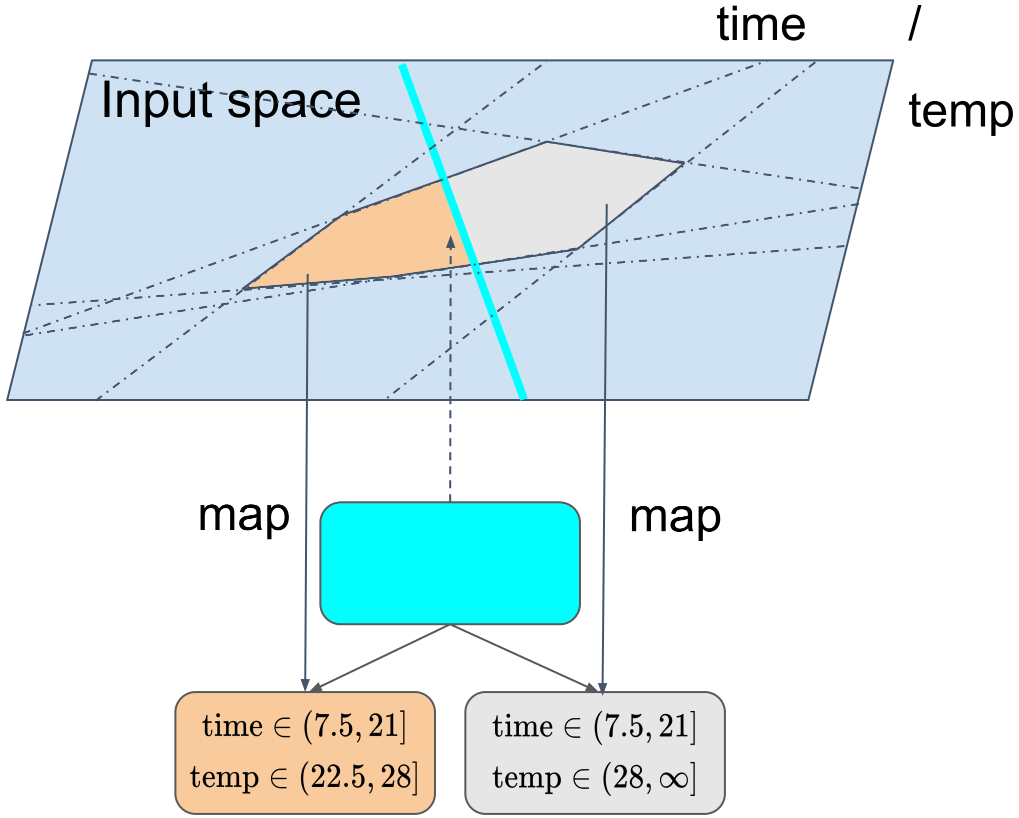

# Verifying the decision tree

Then assert the input-output logic propositions for each leaf node

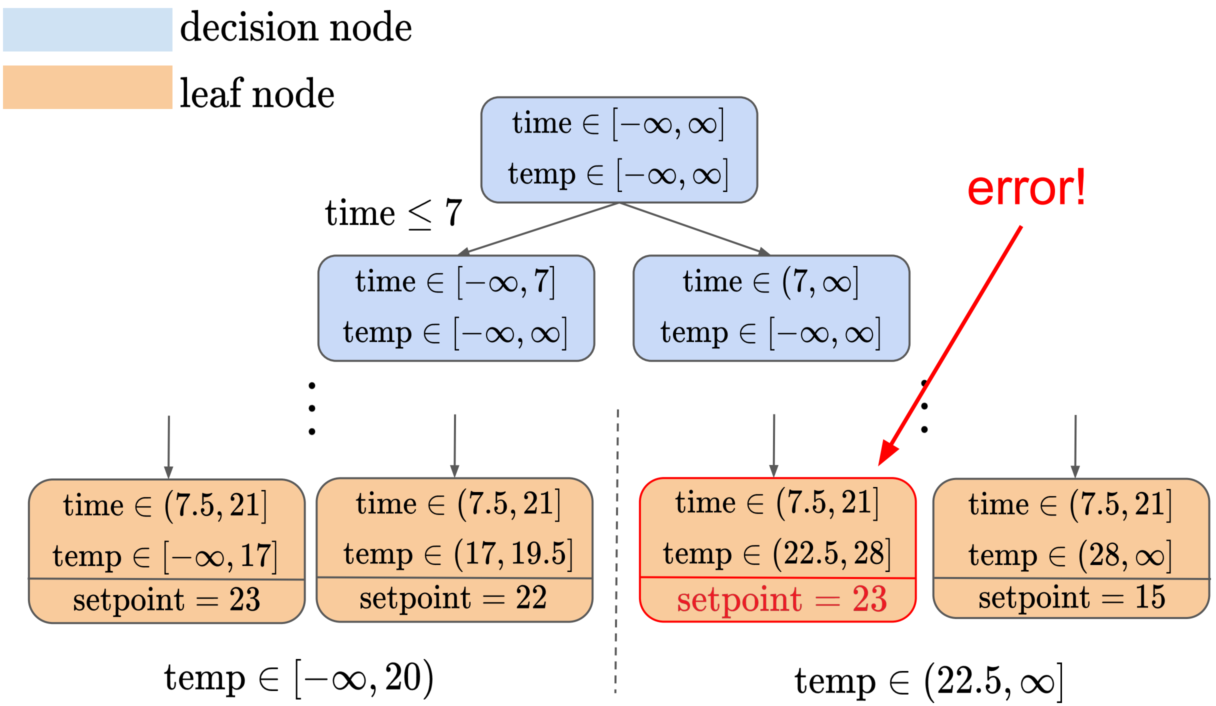


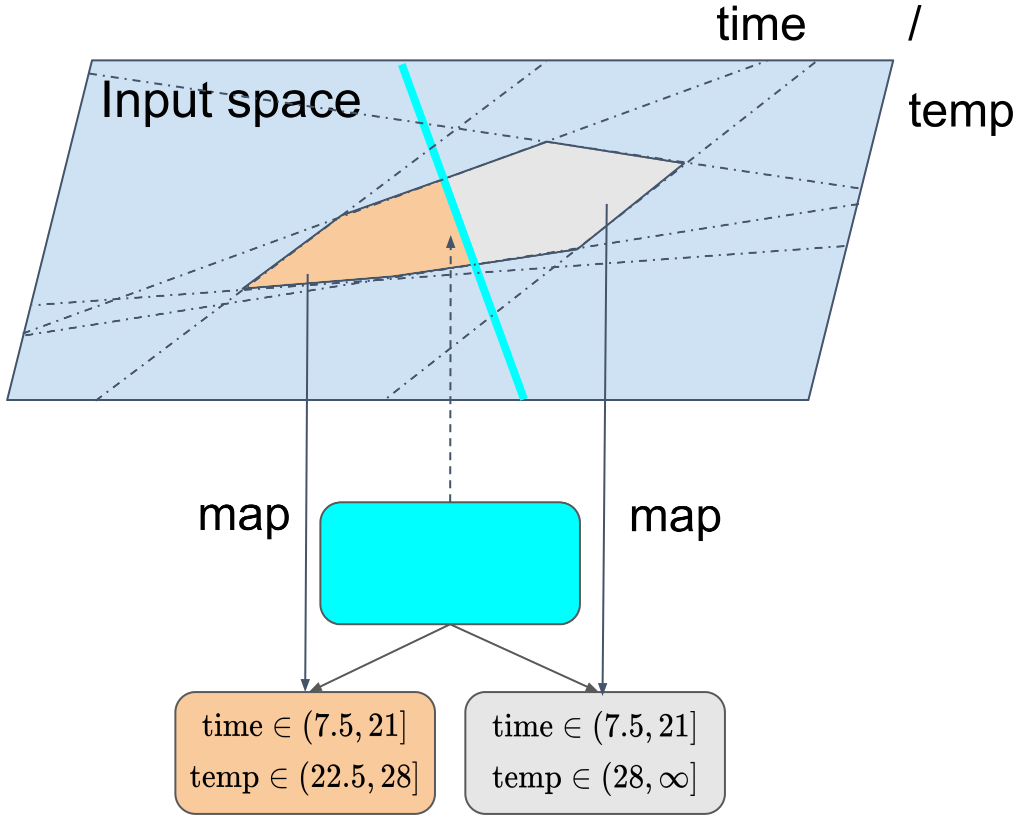

# Verifying the decision tree

We also use probabilistic verification in addition to formal proposition verifications with Monte Carlo method

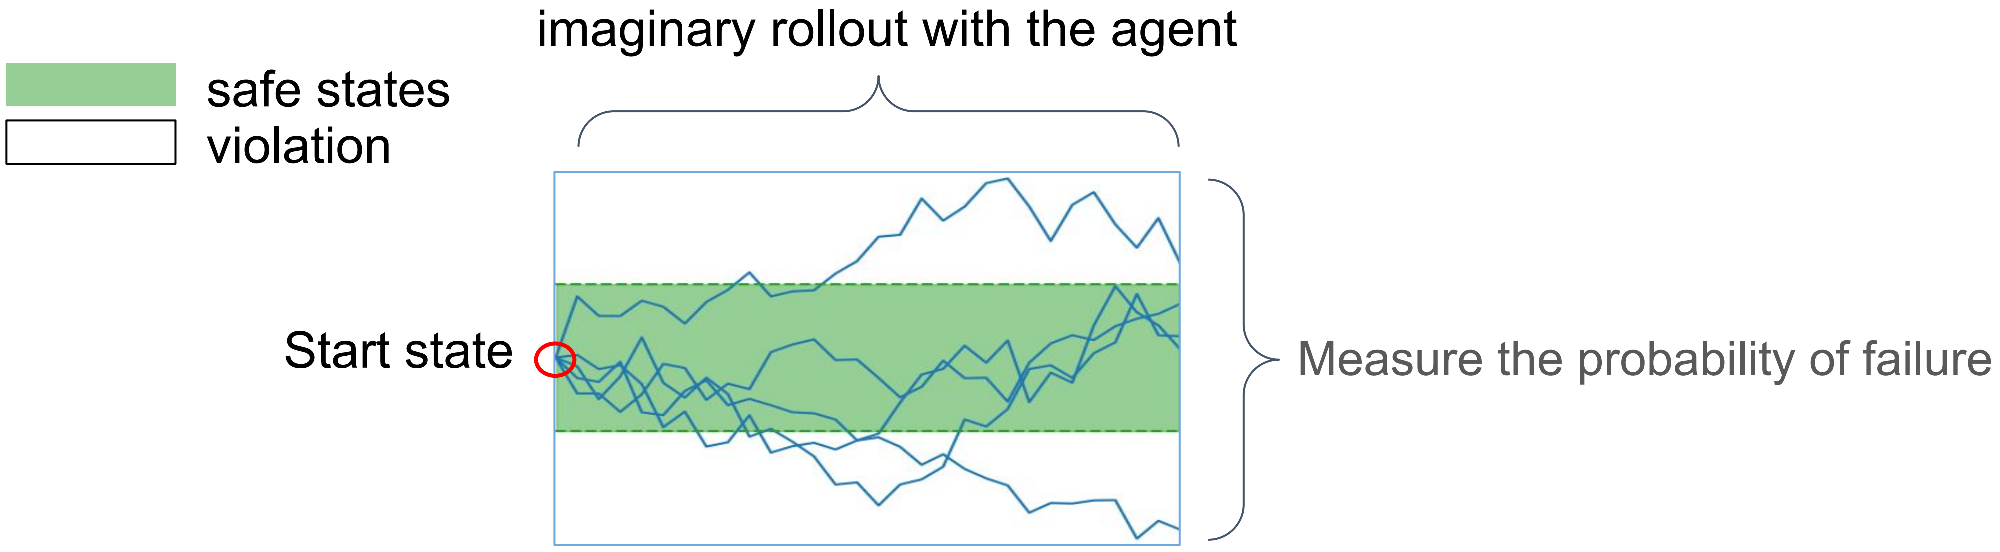

# Summary of our approach

learning the decision tree agent

verifying the agent

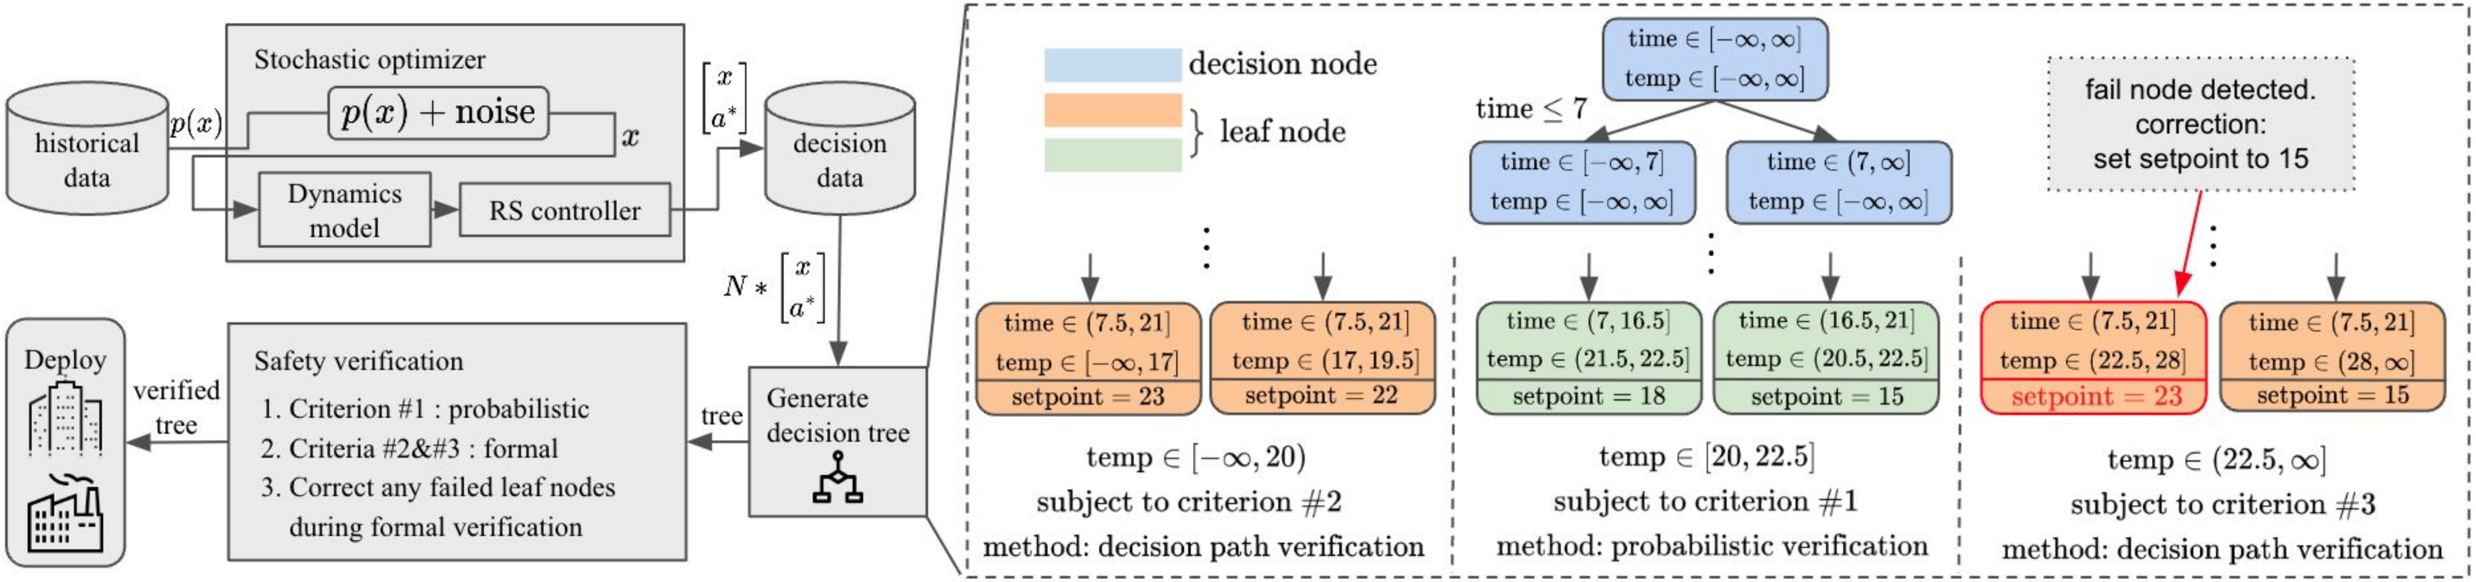

# Experiment Results

- ➔ Simulation experiments with EnergyPlus
  - ◆ Two cities: Pittsburgh, Tucson. Real weather profile.
- ➔ Research questions to answer:
  - ◆ Does our method provide superior energy efficiency?
  - ◆ Does it converge fast enough?
  - ◆ Does the tree size explode to unmanageable scale?
  - ◆ Does it run fast on edge devices?

# Experiment Results

Our method results in superior energy efficiency compared with previous state-of-the-arts

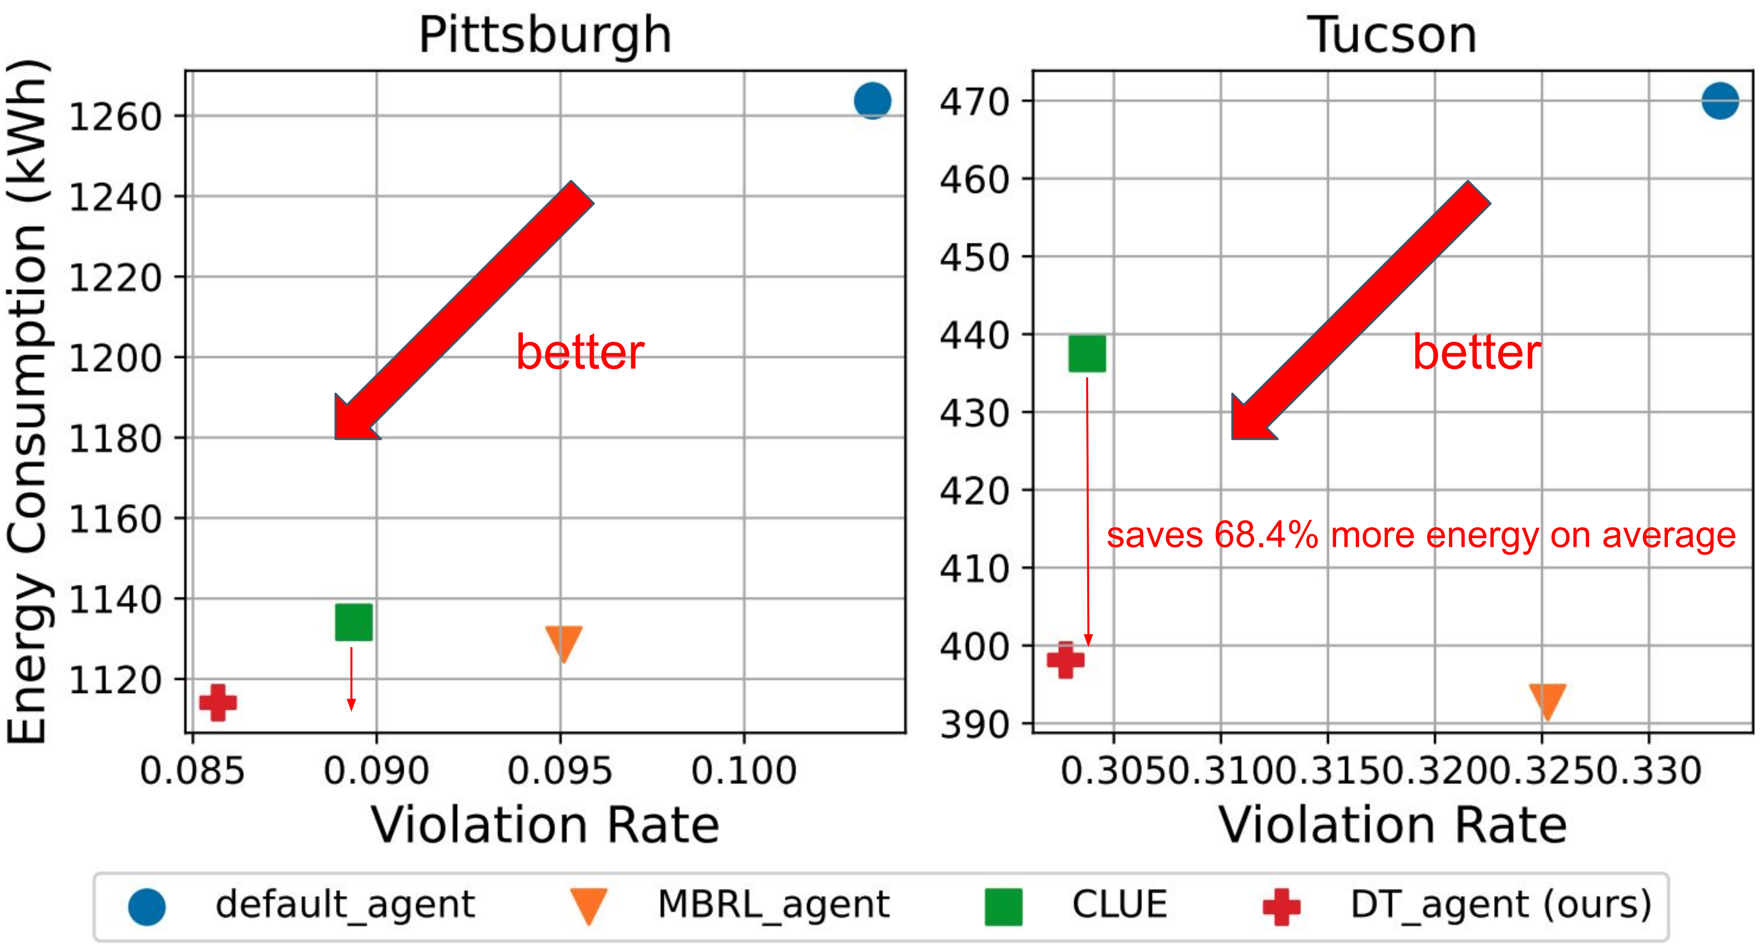

# Experiment Results

Our method converges with small amount of data and decision tree nodes.

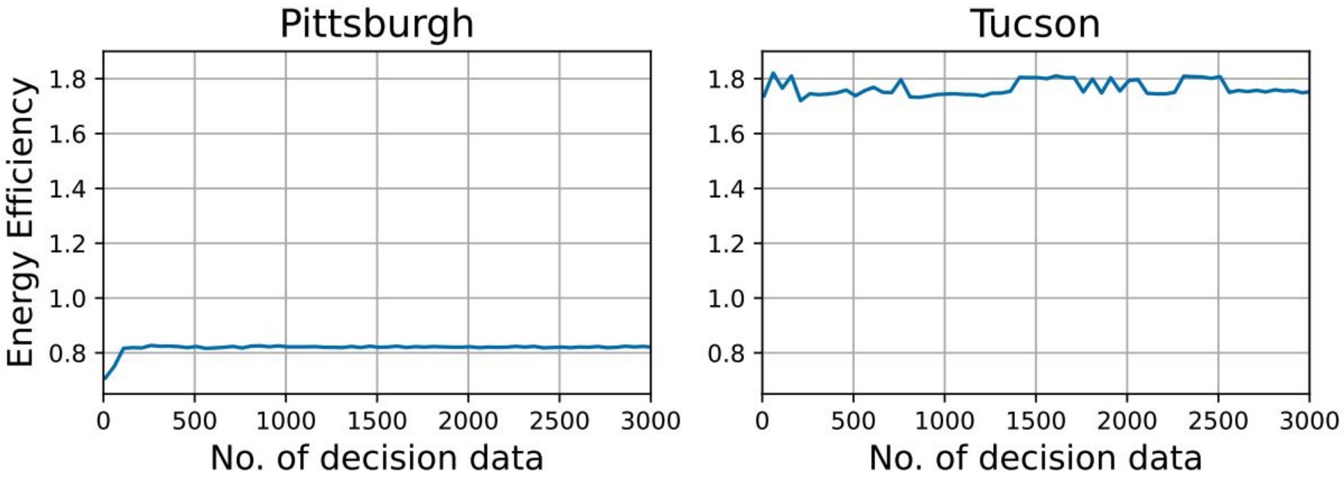


Fast convergence

| | Pittsburgh | Tucson |
|---|---|---|
| Total No. of nodes | 1199 | 3291 |
| No. of leaf nodes (unique path) | 599 | 1646 |
| Safe probability estimated by crit. #1 | 94.6% | 95.1% |
| No. of nodes corrected by crit. #2 | 0 | 0 |
| No. of nodes corrected by crit. #3 | 0 | 88 |

Manageable tree size

# Experiment Results

Our method uses 1127x less time during online inference. Suitable for all edge devices!

| | default [12] | MBRL [9] | CLUE [1] | DT (ours) |
|---|---|---|---|---|
| average (ms) | 0.0 | 212.87 | 326.30 | 0.1888 |
| std (ms) | 0.0 | 266.89 | 102.30 | 0.4423 |

# Thank you

Ryan (Zhiyu) An

zan7@ucmerced.edu

Code available at https://github.com/ryeii/Veri_HVAC