Food Security

Automated Farming

Assignments

Keep working on your projects!

Discussion question: What is the most challenging aspect of your project right now?

Due Friday 5pm

Climate change in the news

Climate change in the news

U.S. ›

U.S. Postal Service starts nationwide electric vehicle fleet, buying 9,250 EVs and thousands of charging stations

BY LI COHEN
UPDATED ON: MARCH 1, 2023 / 10:21 AM / CBS NEWS

The vehicles purchased are Ford E-Transit Battery Electric Vehicles (BEVs), which according to USPS are "100 percent electric." It's part of the agency's plans, announced in December, to make 75% of its newly acquired vehicles, known as Next Generation Delivery Vehicles, over the next five years electric. **After 2026, NGDV purchases will be 100% electric, the agency said.**

Three suppliers were awarded contracts for more than **14,000 charging stations**, as well, USPS said, to kick off its Electric Vehicle Supply Equipment (EVSE) inventory.

A contract has also been awarded for the agency to acquire **9,250 commercial-off-the-shelf internal combustion engine vehicles** "to fill the urgent need for vehicles."

The postal service's drive toward clean energy vehicles is only a recent development that came after it received significant backlash over plans to replace its fleet with 90% gas-powered vehicles.

Summary

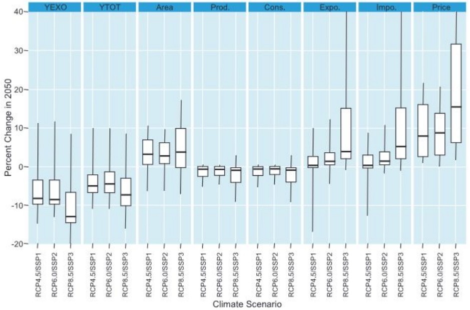

Figure ES-4. Climate-change effects on agricultural commodities in 2050 under different SSPs and RCPs. *The more pessimistic "high concentration/low international cooperation" scenario (RCP8.5/SSP3) shows much larger and more variable climate-change effects for the five commodities (coarse grains, rice, wheat, oilseeds and sugar), than the "medium concentration/middle of the road" (RCP6.0/SSP2) and "low concentration/sustainable development" (RCP4.5/SSP1) scenarios. All are compared to baseline of SSPs with no climate change. Results are from three GCMs and five economic models, aggregated across thirteen regions (n = 75). YEXO = yield effect of climate change without technical or economic adaptation, YTOT = realized yields after adaptation, AREA = agricultural area in production, PROD = total production, CONS = consumption, Expo = exports, IMPO = imports, PRICE = prices.*

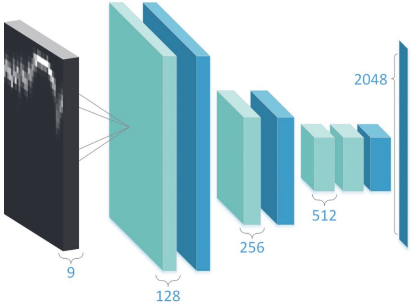

(c) The CNN structure

Farm labor

Farm labor can be dangerous and is made more dangerous by increasing temperatures

Some farm tasks can be automated

But some tasks are too delicate to automate

Changing climate will harm outdoor crops

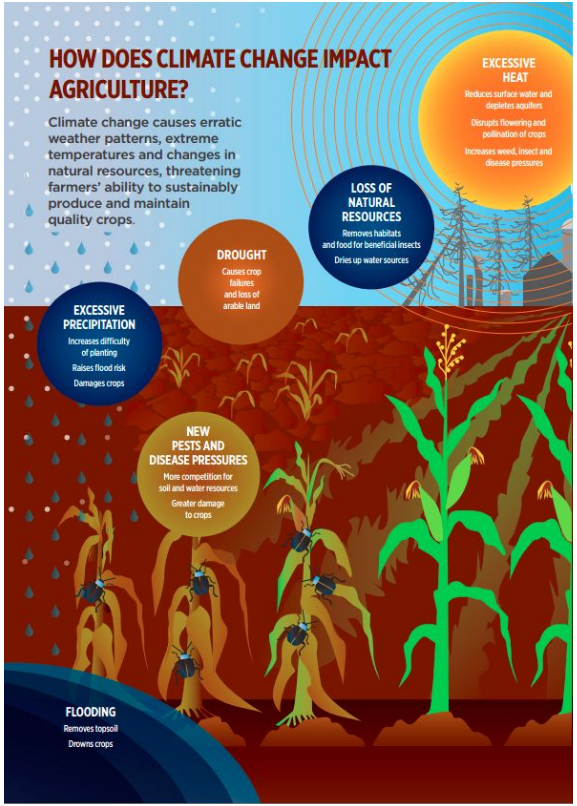

Vertical Farms and Precision Agriculture

Indoor farms allow for climate and pest control

Precision agriculture uses fine-grained data to tailor climate etc to plant needs

Paper Deep Dive

Robust Model-based Reinforcement Learning for Autonomous Greenhouse Control

Wanpeng Zhang ZHANGWP19@MAILS.TSINGHUA.EDU.CN
Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

Xiaoyan Cao HOLMESCAOXY@GMAIL.COM
Department of Informatics, Xiamen University, Xiamen, China

Yao Yao Y-YAO19@MAILS.TSINGHUA.EDU.CN
Zhicheng An AZC19@MAILS.TSINGHUA.EDU.CN
Xi Xiao XIAOX@SZ.TSINGHUA.EDU.CN
Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

Dijun Luo DIJUNLUO@TENCENT.COM
Tencent AI Lab, Shenzhen, China

Background and Motivation

“The traditional agricultural production mode is highly dependent on the weather, which has been unable to meet the increasing demand for fresh and healthy food from global population growth. ”

“Modern high-tech greenhouses are equipped with standard sensors and actuators (such as heating, lighting, CO_2 dosing, irrigation, etc.) to empower precision agriculture. To improve crop yield and quality, ***managers regularly regulate a suitable environment*** for crop growth by overseeing the greenhouse climate and crop growth state.”

Background and Motivation

“Skilled managers capable of autonomous greenhouse control are scarce. Furthermore, even a seasoned manager is not able to monitor and manage too many greenhouses simultaneously.”

“The grower needs to balance production and resource consumption during a 3-5 months period, which implies a tremendous decision-making space. The complexity of decision-making has led to growers **only giving coarse-grained control strategies**, which do not make full use of the rich greenhouse states information.”

Background and Motivation

“Skilled managers capable of autonomous greenhouse control are scarce. Furthermore, even a seasoned manager is not able to monitor and manage too many greenhouses simultaneously.”

“The grower needs to balance production and resource consumption during a 3-5 months period, which implies a tremendous decision-making space. The complexity of decision-making has led to growers **only giving coarse-grained control strategies**, which do not make full use of the rich greenhouse states information.”

Goal: automate the decision making process for greenhouse control

Brainstorm

What kind of data would you want to have to be able to approach this problem?

What kind of methods would you apply?

How would you measure success?

What difficulties might you face?

Data

Temperature, humidity, CO2, and “photosynthetically active radiation” measured from a greenhouse with a ventilation controller, CO2 producer, Fertigation (irrigation) controller.

Goal: grow the most tomatoes for the least cost

(a) (b) (c) (d) (e) (f)

Figure 2: Equipments in our greenhouse. (a)temperature sensor. (b)CO_2 sensor. (c)PAR sensor. (d)ventilation controller. (e)CO_2 producer. (f)fertigation controller.

Framing the problem

Framing the problem

When using reinforcement learning to solve a control problem, the aim is to develop a **policy** that controls an **agent** to maximize **reward**.

The policy is a function that takes in a **state observation** and produces an **action**.

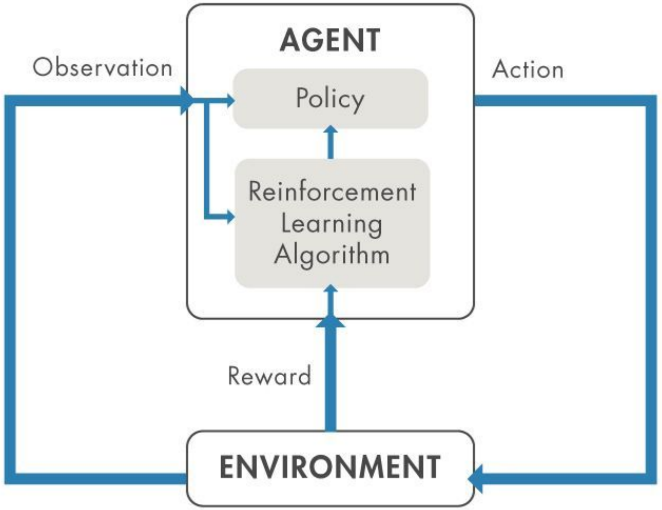

https://www.mathworks.com/company/newsletters/articles/reinforcement-learning-a-brief-guide.html

Framing the problem

When using reinforcement learning to solve a control problem, the aim is to develop a **policy** that controls an **agent** to maximize **reward**.

The policy is a function that takes in a **state observation** and produces an **action**.

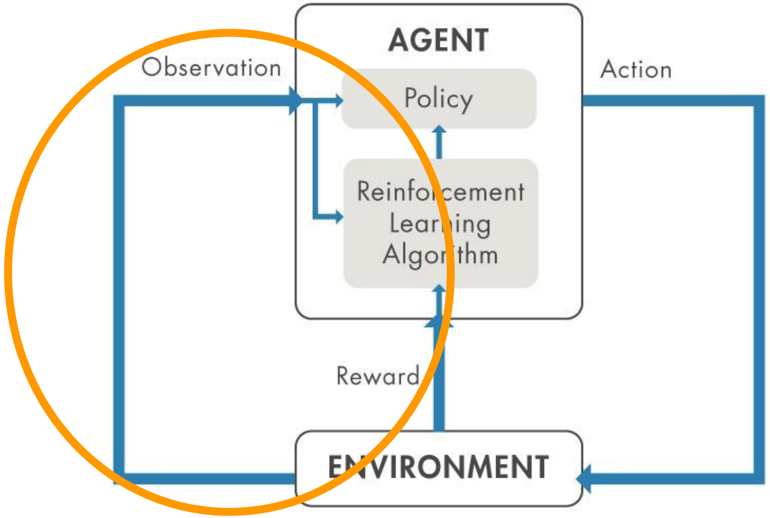

https://www.mathworks.com/company/newsletters/articles/reinforcement-learning-a-brief-guide.html

RL example

Agent = person

State = coordinates, walls, reward

Reward = star

Actions = up,down,left,right

Policy = what action to take when in a certain state

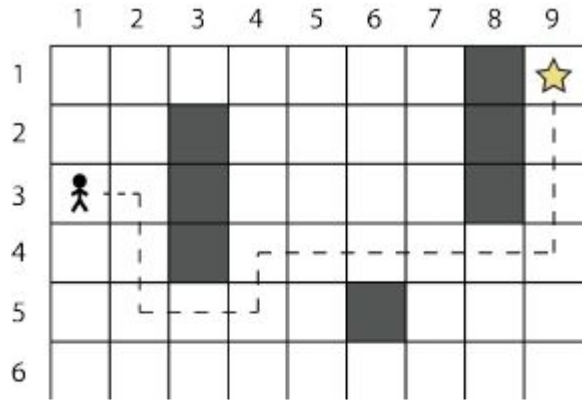

Reinforcement learning is hard because of...

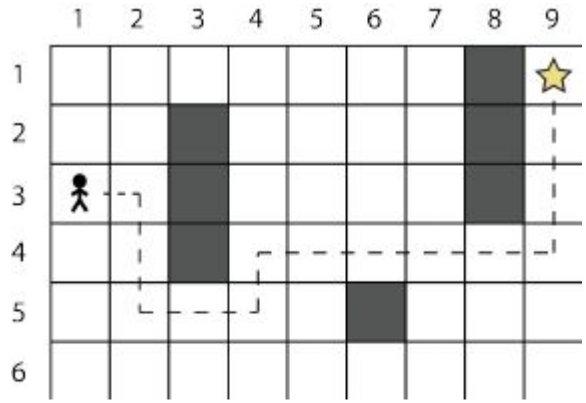

Reinforcement learning is hard because of...

- Sparse feedback
- Long-term dependencies
- Complex observations
- Complex action spaces
- Uncertainty in the world
- Data needs

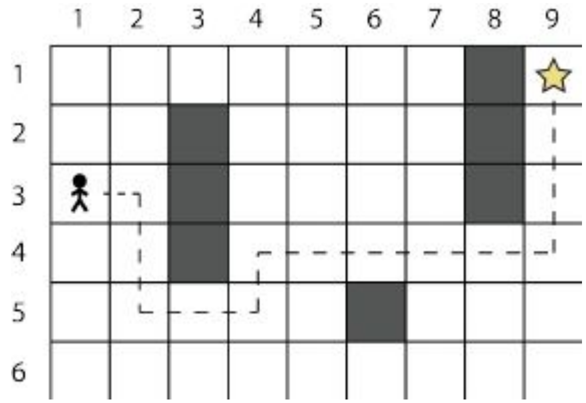

RL concepts

Markov Decision Process (describes the situation in which RL can take place)

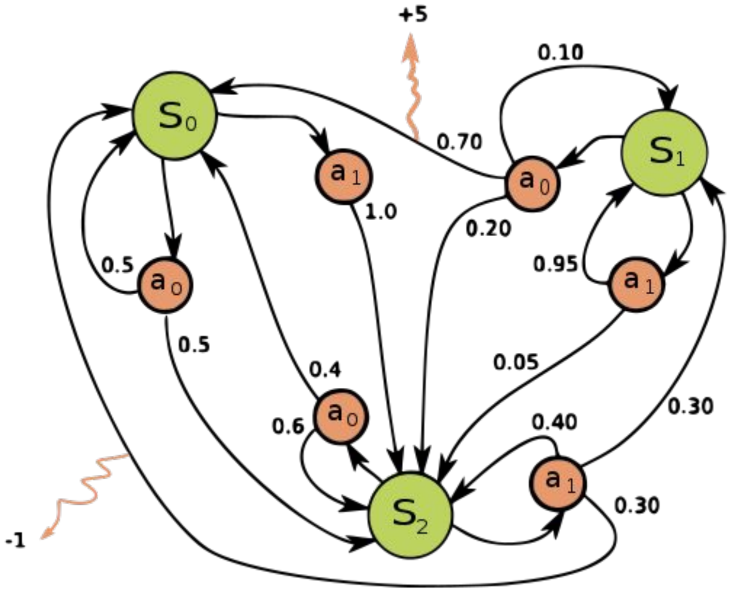

If in a state S_x, each action will lead to State_y, with some probability

RL Concepts

Value functions

RL concepts

Value functions

Policies can be defined as taking the action with the highest value

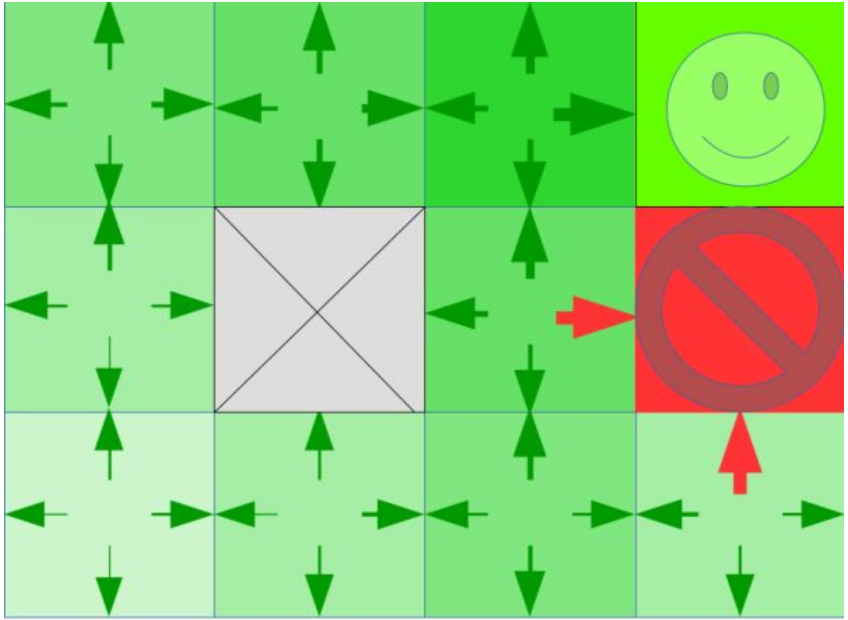

RL concepts

Model-free vs Model-based Agents

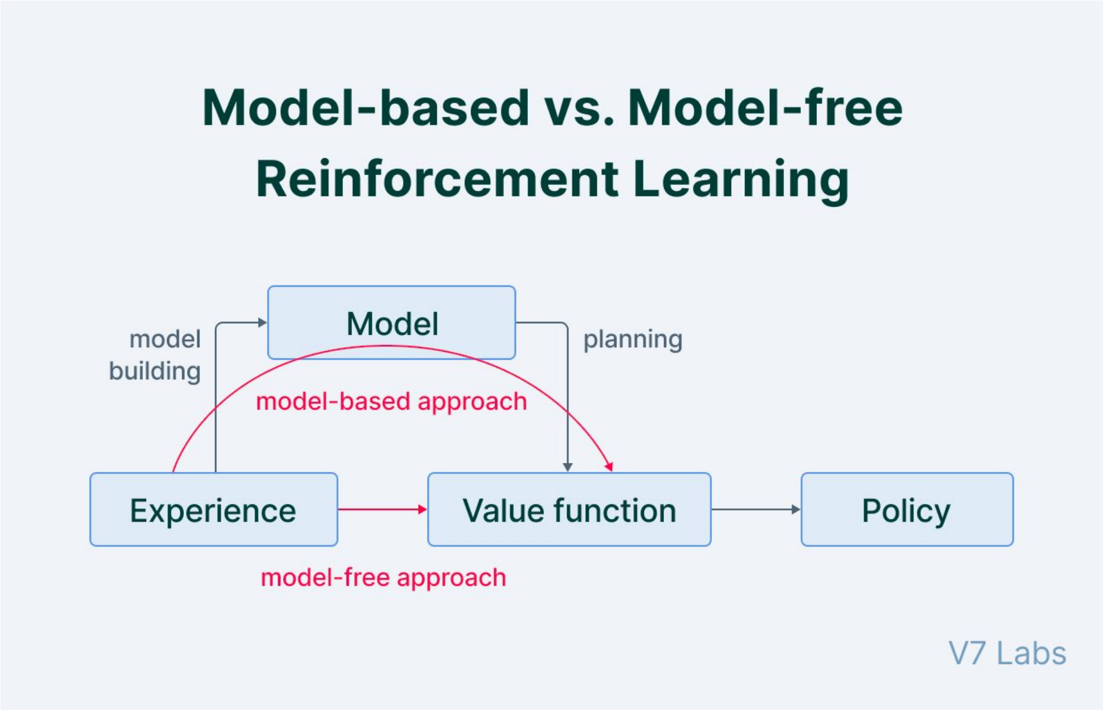

RL concepts

Learning algorithms: trying to find the best value function/policy

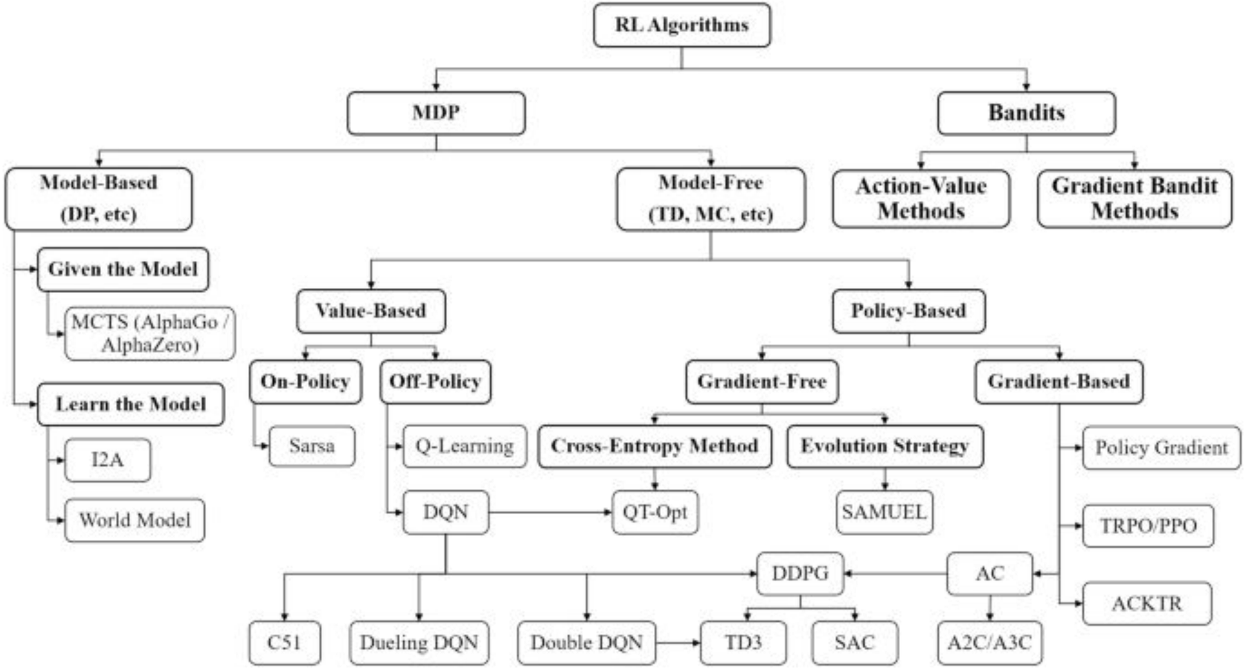

Monte Carlo methods

Model-free Monte Carlo

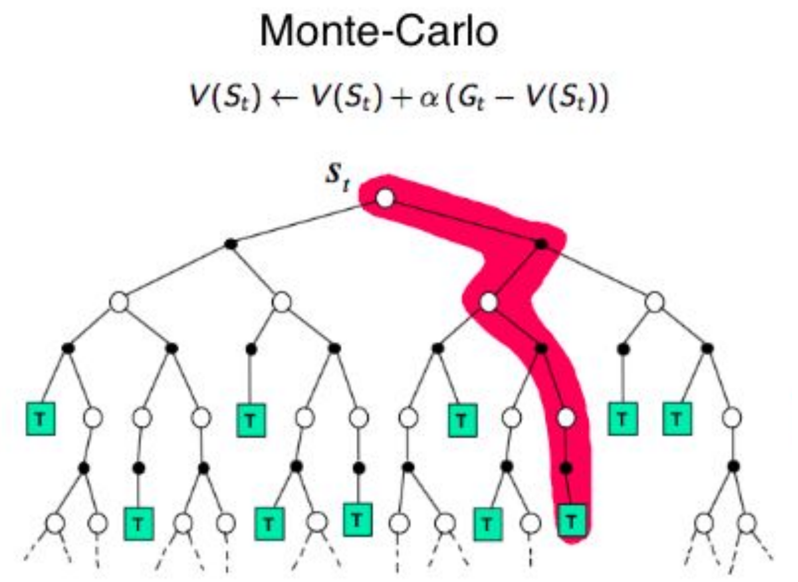

Try random moves and see how much reward it gets you in the end.

Value function of a state is the average reward received after passing through that state.

Monte Carlo methods

Model-free Monte Carlo

Monte-Carlo

$$V(S_t) \leftarrow V(S_t) + \alpha\,(G_t - V(S_t))$$

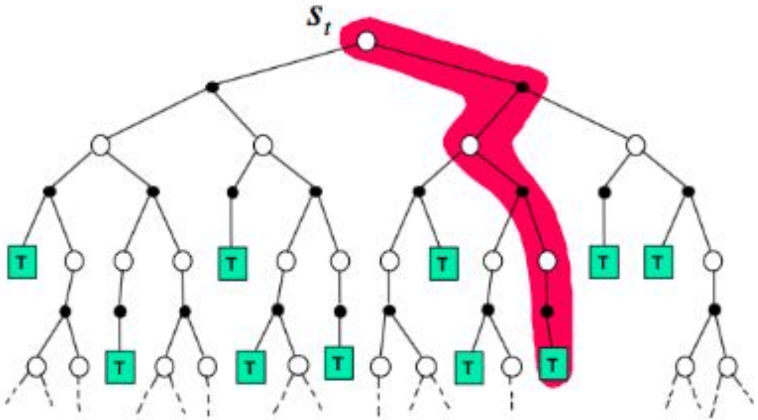

Try random moves and see how much reward it gets you in the end.

Value function of a state is the average reward received after passing through that state.

Important concept: Discounted Reward

Let $V^{\pi_\theta,\mathcal{P}}(s)$ denotes the expected return or expectation of cumulative rewards starting from initial state s, i.e., the expected sum of discounted rewards following policy $\pi_\theta(a|s)$ and state transition function $\mathcal{P}(s,a)$:

$$V^{\pi_\theta,\mathcal{P}}(s) = \mathop{\mathbb{E}}_{\{a_0,s_1,\ldots\}\sim\pi_\theta,\mathcal{P}}\left[\sum_{t=0}^{\infty}\gamma^t r(s_t,a_t) \mid s_0 = s\right] \qquad (1)$$

Monte Carlo methods

Model-based Monte Carlo

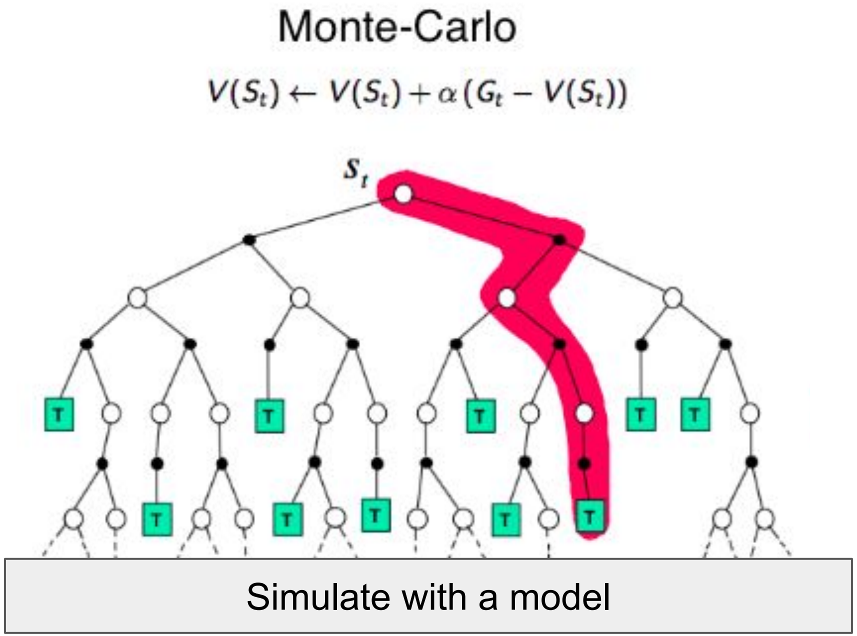

Models can be used to speed up a tree search.

Framing the problem

"We collect 22 kinds of observation variables, constituting a 275-dimensional observation space S, and 6 kinds of control variables, constituting a 52-dimensional action space"

Table 1: Observation Space (Above) and Action Sapce (Below)

Name	Min	Max	Dim
temperature setpoint	13	32	24
CO_2 setpoint	400	1000	24
light-on time	0	24	1
light-off time	0	24	1
irrigation start time	0	24	1
irrigation stop time	0	24	1
outside solar radiation	0	2000	24
outside temperature	-30	50	24
outside humidity	0	100	24
wind speed	0	25	24
virtual sky temperature	-20	20	24
greenhouse air temperature	-30	100	24
greenhouse air humidity	0	100	24
greenhouse air CO_2 concentration	400	1000	24
light intensity just above crop	0	2000	24
cumulative amount of irrigation per day	0	10	1
cumulative amount of drain per day	0	10	1
leaf area index	0	10	1
current number of growing fruits	0	1000	1
cumulative harvest in terms of fruit fresh weight	0	100	1
cumulative harvest in terms of fruit dry weight	0	100	1
planting days	0	365	1
temperature setpoint	13	32	24
CO_2 setpoint	400	1000	24
light on time	0	24	1
light off time	0	24	1
irrigation start time	0	24	1
irrigation stop time	0	24	1

Framing the problem

We use the Netprofit (USD/m2) as the target reward for training.

Netprofit = Gains − Costs,

where Gains are obtained through yields and price, and Costs include resource consumption (electricity, heat, CO_2, and water) and crop maintenance costs.

Training problem

RL agents can make a lot of mistakes while learning and need many iterations

Training problem

To solve the lack of data, train a "greenhouse simulator" based on existing greenhouse data.

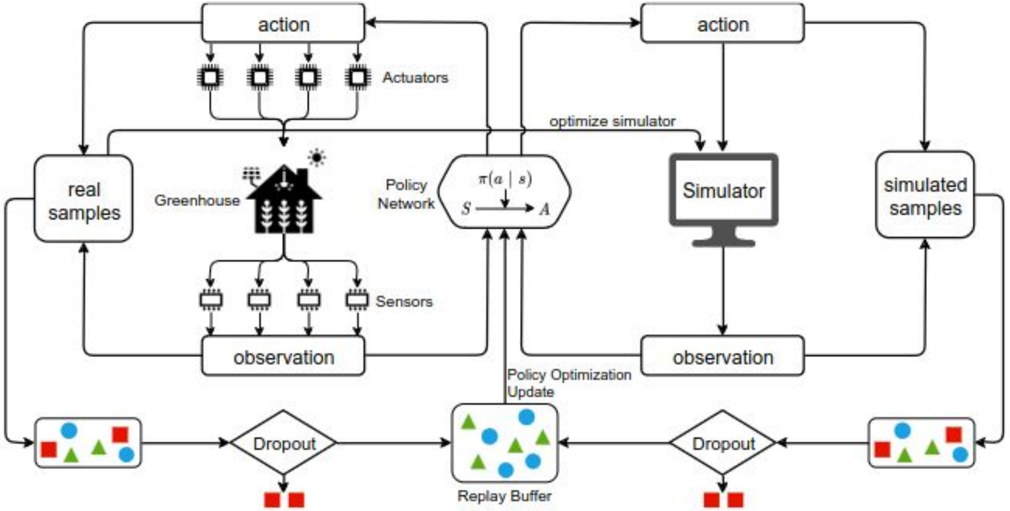

Training problem

To solve the lack of data, train a “greenhouse simulator” based on existing greenhouse data.

To correct for the fact that the simulator won’t be perfectly accurate, use an ensemble

model. We learn a collection of fine-tuned simulator models $\mathcal{M} \doteq \{\mathcal{M}_{\phi_1}, \mathcal{M}_{\phi_2}, \ldots, \mathcal{M}_{\phi_N}\}$. We use parametric notation $\mathcal{M}_\phi, \phi \in \Phi$ to specifically denote the model trained by neural network, where Φ is the parameter space of models. Each member of the collection $\mathcal{M}$ is a probabilistic neural network whose outputs $\mu_{\phi_i}, \sigma_{\phi_i}$ parametrize a Guassian distribution:

$$s' = \mathcal{M}_{\phi_i}(s, a) \sim \mathcal{N}(\mu_{\phi_i}(s, a), \sigma_{\phi_i}(s, a)) \tag{4}$$

Training problem

To solve the lack of data, train a “greenhouse simulator” based on existing greenhouse data.

To correct for the fact that the simulator won’t be perfectly accurate, use an ensemble

model. We learn a collection of fine-tuned simulator models $\mathcal{M} \doteq \{\mathcal{M}_{\phi_1}, \mathcal{M}_{\phi_2}, \ldots, \mathcal{M}_{\phi_N}\}$. We use parametric notation $\mathcal{M}_{\phi}, \phi \in \Phi$ to specifically denote the model trained by neural network, where Φ is the parameter space of models. Each member of the collection $\mathcal{M}$ is a probabilistic neural network whose outputs $\mu_{\phi_i}, \sigma_{\phi_i}$ parametrize a Guassian distribution:

$$s' = \mathcal{M}_{\phi_i}(s, a) \sim \mathcal{N}(\mu_{\phi_i}(s, a), \sigma_{\phi_i}(s, a)) \tag{4}$$

“Individual probabilistic models capture **aleatoric uncertainty** or the noise due to the inherent stochasticity. The bootstrapping ensemble procedure can capture **epistemic uncertainty** or uncertainty in the model parameters aroused from insufficient training data.”

Training problem #2

Especially when applying to real world environments, want to not just have large average reward, but also avoid the possibility of catastrophic outcomes

Training problem #2

Especially when applying to real world environments, want to not just have large average reward, but also avoid the possibility of catastrophic outcomes

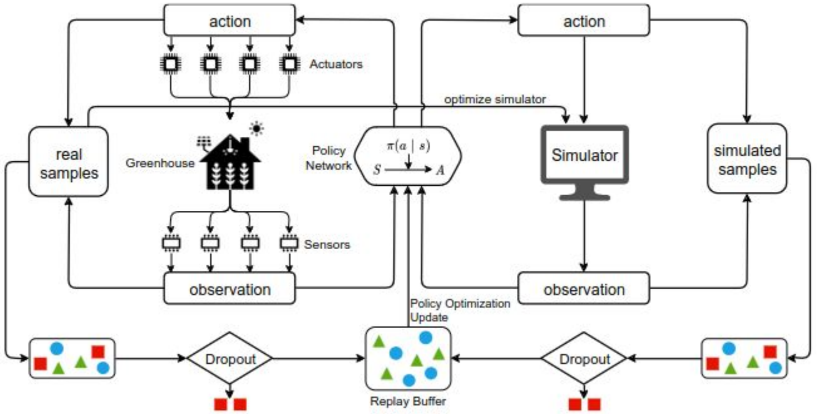

Inspired by previous works Tamar et al. (2015); Rajeswaran et al. (2016) which optimize conditional value at risk (CVaR) to explicitly seek a robust policy, we add a sample dropout module to the RL algorithm, which selectively discards a portion of samples with excessive reward, to focus more on worst-case samples, improving the adaptability of the planting policy in extreme situations, and solve the safety challenges of RL in real-world deployment, aiming to further enhance the safety of the automated planting policy.

Evaluation

Would take too long to test it out on real greenhouse, so use a hand-built computationally expensive greenhouse simulator.

Evaluate on expected return and net profits

Compare to model-free method, learned directly from observed state-action pairs

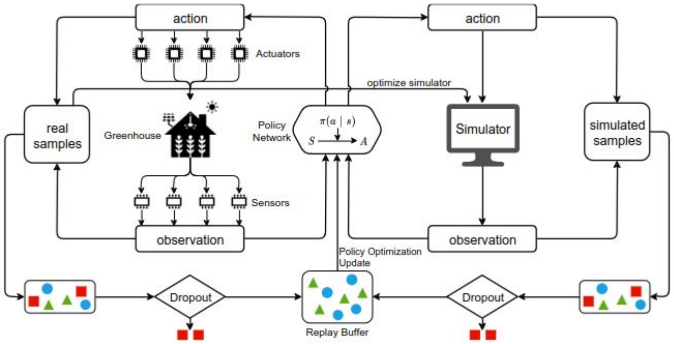

Results

We train two versions of our method on the greenhouse simulator, one with sample dropout ($p = 0.8$) and one without sample dropout ($p = 1$) and compare to a model-free method.

Results

We train two versions of our method on the greenhouse simulator, one with sample dropout (p = 0.8) and one without sample dropout (p = 1) and compare to a model-free method.

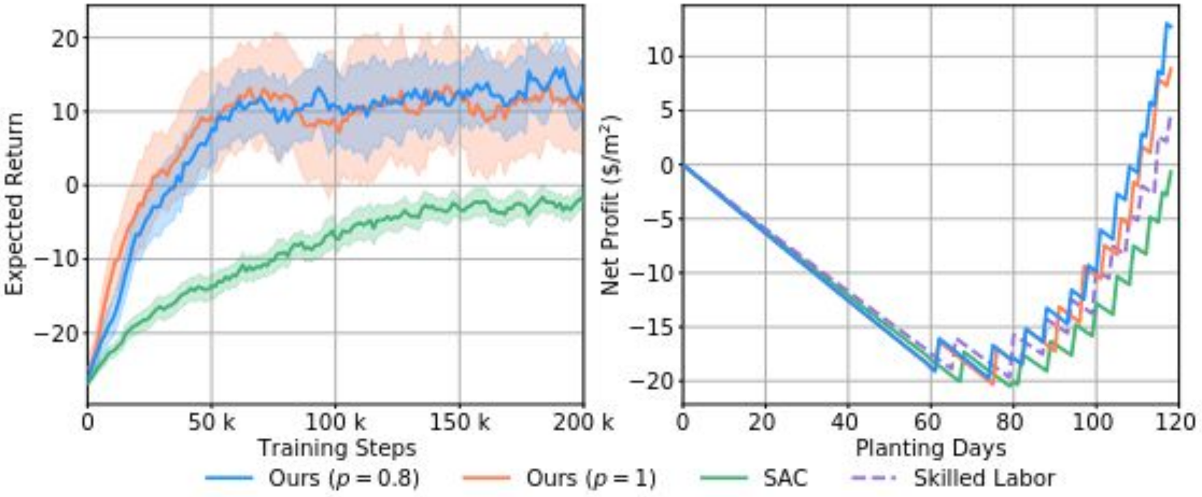

"As for the SAC algorithm, it performs worse than our algorithm, which is caused by the low sample efficiency of the model-free method, making it difficult to learn enough information with limited samples."

Results

Sample dropout makes the model more robust to a wider range of conditions

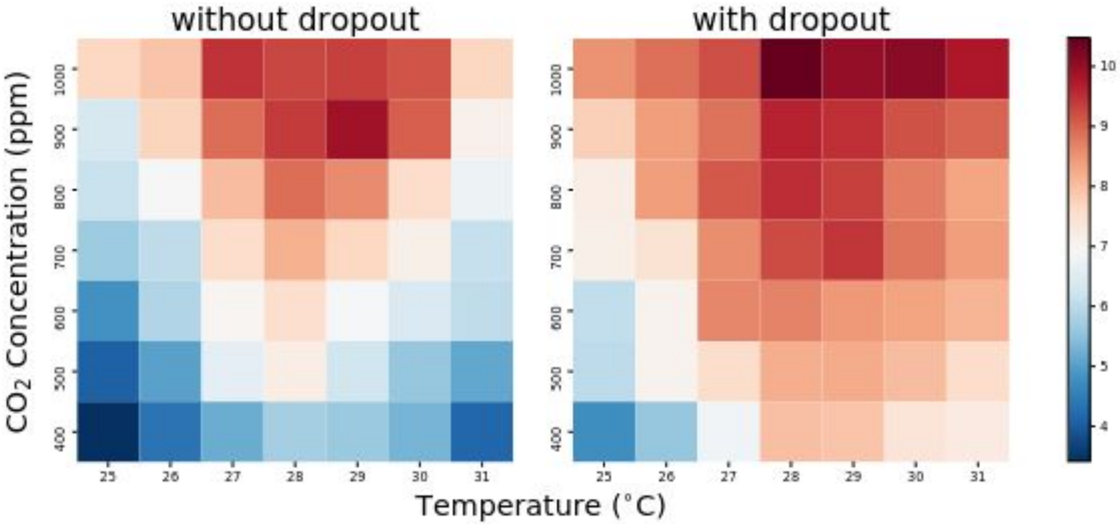

Results

Too much dropout is bad

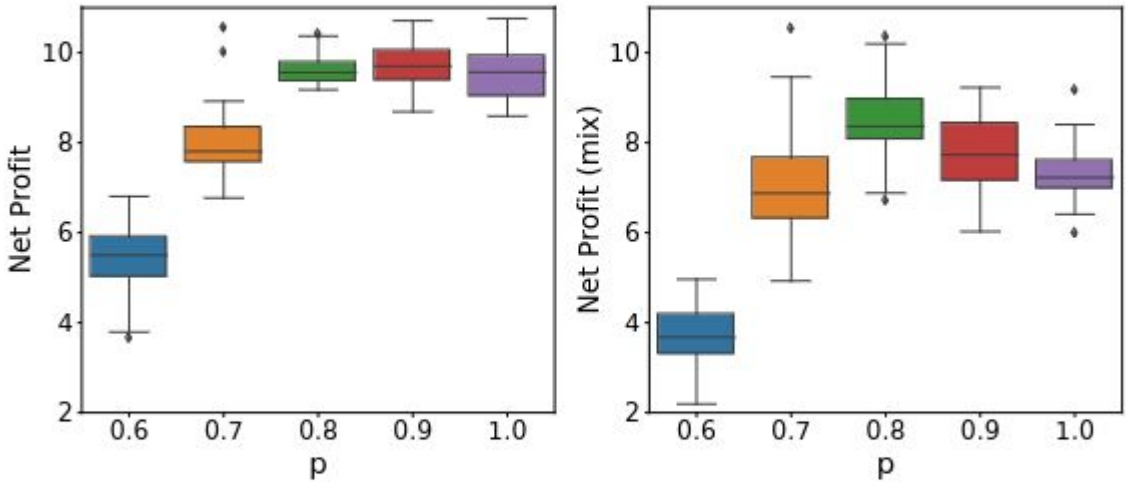

The effect of adjusting parameter p: The left box plot shows the $Netprofit$ that the algorithms can achieve with different parameters in the standard environment; The right box plot shows the average values of $Netprofit$ that the algorithm can achieve with different parameters in the four disturbed environments (Tempera-

Further resources

More on robot farm equipment: https://youtu.be/bpa1iiJmR3Q

https://deeplizard.com/learn/video/nyjbcRQ-uQ8 Online lessons on reinforcement learning

Summary

National Farm Safety and Health Week September 18-24, 2022
Top 5 Farm-Related Injuries
1 Overturning tractors and heavy machinery
2 Falls
3 Toxic chemical exposure to pesticides
4 Suffocation
5 Heat stress
$8.3 Billion
Annual costs of occupational injuries in agriculture**
100 ag workers
daily suffer a lost-work-time injury*
410
farmworkers who died in a work-related injury In 2019
70,000
suffered heat-related injuries during last 25 years*
20x
U.S. crop workers more likely to die of heat-related illness*
EFI provides training and a place for workers to have a voice in the farming operation. Worker-manager collaborative teams play a vital role in creating safer and healthier workplaces.
Learn more at equitablefood.org.
Sources: *National Institute for Occupational Safety and Health (NIOSH); **OSHA
efi
RESPONSIBLY GROWN FARMWORKER ASSURED

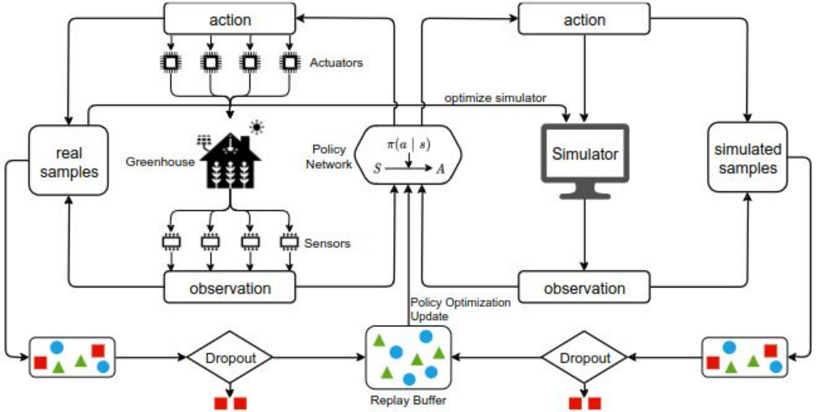

action
Actuators
real samples
Greenhouse
Sensors
observation
Policy Network
$\pi(a \mid s)$
S
A
optimize simulator
action
Simulator
simulated samples
observation
Policy Optimization Update
Dropout
Replay Buffer
Dropout