AI for Agricultural Sciences

Opportunities, Developments, and Challenges

Debabrota Basu

Équipe Scool, Inria, Université de Lille, CNRS, France

Bihar Agricultural University, India, July 2023

The Trajectory

1. Artificial Intelligence: A Brief Introduction
2. AI meets Agroecology
3. Reinforcement Learning (RL) for Agricultural and Ecological Sciences
    3.1 RL: A Brief Overview
    3.2 gym-DSSAT
    3.3 Farmgym: an environment for learning to farm
    3.4 ForestDefender
4. Future Roadmap: Challenges and Opportunities

Intelligence: Human and Artificial

Intelligence is the ability of

Intelligence: Human and Artificial

Intelligence is the ability of

1. accumulating information,
2. processing these information in form of general lessons
3. learning these lessons about the world
4. using them for solving problems and decision making.

Intelligence: Human and Artificial

Intelligence is the ability of

1. accumulating information,
2. processing these information in form of general lessons
3. learning these lessons about the world
4. using them for solving problems and decision making.

AI is also an artefact of its experience like humans.

AI: How does it work?

“AI is the person who goes to a restaurant and asks for the food that the majority likes."

AI: A Marriage of Statistics, Computing, and Optimisation

AI operates by

1. accumulating data
   $\longrightarrow$ **IoT** and **Computation**
   (Surveys, Sensors, Edge computing, Federated learning etc.)

2. processing data and learning patterns about the problem
   $\longrightarrow$ **Statistics**, **Optimisation**, and **Computation**
   (Linear models, Kernels, Deep neural networks, etc.)

3. using the patterns for solving problems and decision making
   $\longrightarrow$ **Statistics** and **Optimisation**

AI meets Agroecology

The future of food production in the face of population growth and climate change

The Silver Lining

The Future of Farming: Artificial Intelligence and Agriculture

08.JAN.2020 7:23 PM . 5 MIN READ

While artificial intelligence (AI) seemed until recently to be science fiction, countless corporations across the globe are now researching ways to implement this technology in everyday life. AI works by processing large quantities of data, interpreting patterns in that data, and then translating these interpretations into actions that resemble those of a human being. Scientists have used it to develop self-driving cars and chess-playing computers, but the technology has expanded into another domain: agriculture. AI has the potential to spur more efficient methods of farming in order to combat global warming, but only with expanded regulation of its development.

Global Warming and Agriculture: A Vicious Cycle

Sydney Young

Sydney is the former Director of Interviews and Perspectives at the HIR. She is interested in health, human rights, and social justice issues.

RECENT POSTS

The Vision of Academia and Governments (1/2)

AGROECOLOGY COCOA

Using artificial intelligence to model complex agroforestry systems

SCIENCE AT WORK • 20 May 2019

An innovative project has just been launched at CIRAD to model cocoa agroforestry systems using artificial intelligence. This ambitious programme, called Deep2DPE, will explore three research fronts: the mathematics of modelling, artificial intelligence (neural networks and machine learning) and agronomy to optimise agroforestry systems by modelling them. This third objective should, for example, enhance understanding of how competition for light between species affects cocoa production.

WORLD ECONOMIC FORUM

Join us Sign in

IMPACT AGRICULTURE, FOOD AND BEVERAGE

AI for agriculture: How Indian farmers are harnessing emerging technologies to sustainably increase productivity

The Vision of Academia and Government (2/2)

Press release

Agriculture and Digital Technology: a white paper by Inria and INRAE to establish the foundations for responsible digital agriculture

Date: 01 Mar. 2022

Home > News and events
> Agriculture and Digital Technology: a white paper by Inria and INRAE to establish the foundations for responsible digital agriculture

The Industrial Momentum

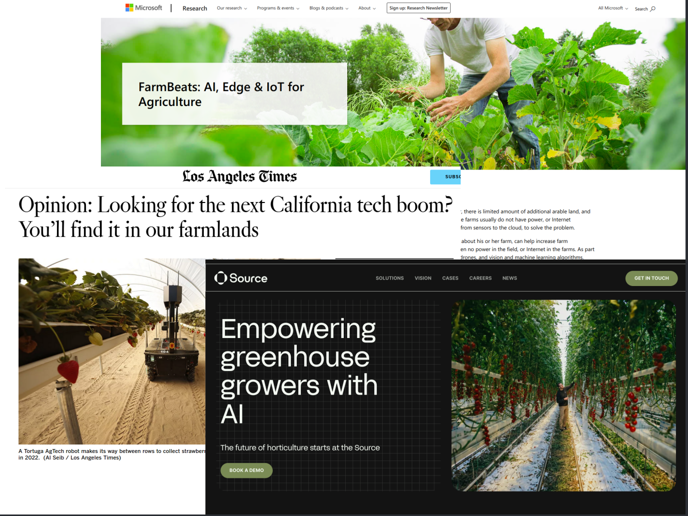

Reinforcement Learning (RL)

The sub-field of AI that learns from interactions

and optimises a long-term goal over a sequence of interactions

Decision Making under Incomplete Information

1	Fever
2	Pressure
...	...
D	Respiration

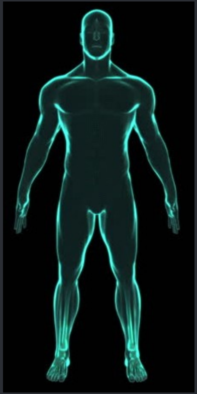

Medical State

s_t

Decision Making under Incomplete Information

with Multiple States: Markov Decision Process (MDP)

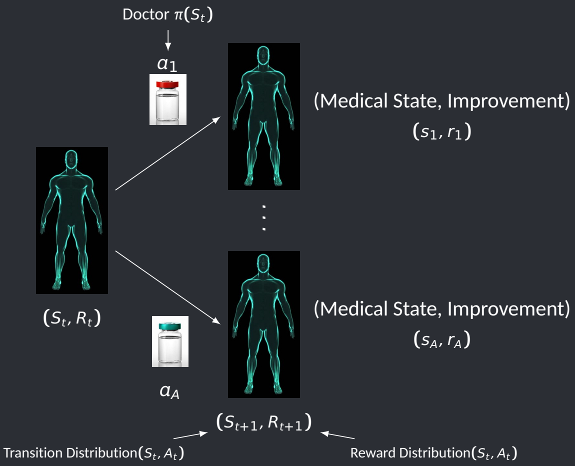

Decision Making under Incomplete Information

with Multiple States: Markov Decision Process (MDP)

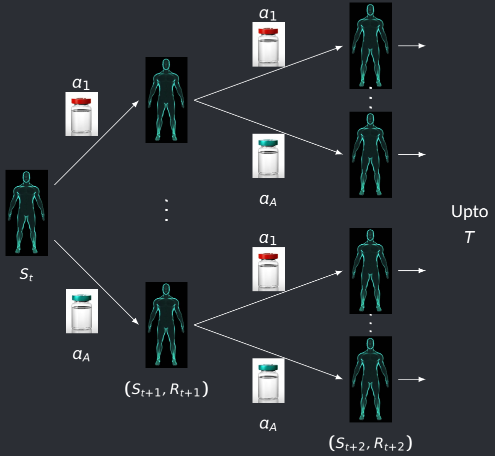

Goal: Maximise the total improvement of the patient $\sum_{t=1}^{T} R_t$

Reinforcement Learning (RL) [Sutton and Barto, 2018]

A Reinforcement Learning (RL) algorithm learns from interaction while optimising a long-term goal.

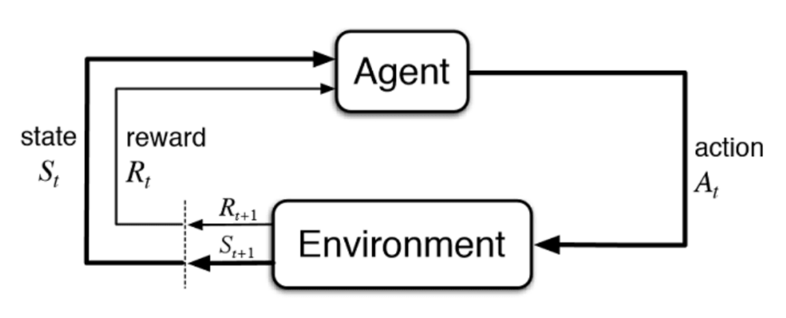

Goal: **Maximise the expected cumulative sum of rewards for agent π, which is the value of following a policy for a given environmental dynamics.**

gym-DSSAT

Using precise simulators to improve and adapt agricultural decision making

[Gautron et al., 2023]

What is gym-DSSAT?

DSSAT: State-of-the-art crop simulator.

gym: Standardized API to connect a reinforcement learning agent with a simulator of its environment.

- The very first version of gym-DSSAT dates back end of 2021. It is an on-going effort.
- For the RL community, gym-DSSAT is a great, original environment including a wide range of fundamental research problems still loosely investigated, potentially impactful w.r.t. **sustainable development**, based on a state-of-the-art crop simulator.
- For agronomists, a tool to investigate how reinforcement learning may be used to improve crop management.

Decision Support System for Agrotech Transfer [Jones et al., 2003]

- Developed for more than 30 years now, U. Florida, Gainsville (`https://dssat.net`).
- Mechanistic model of crop.
- Simulates very accurately the growth of a plant based on the properties of the soil, the cultivar, the weather conditions, initial soil conditions (residue from previous year), ...
  + the actions made in the field: irrigation, fertilization, tillage, ... on a daily basis.
- Simulates a unit of surface very finely: interactions between the soil properties with roots then growth of the plant (PDE integration over time).

gym-DSSAT features

- gym-DSSAT simulates a crop season.
- Thousands different soils from all over the world.
- 42 potential crops (wheat, maize, rice, chickpea, ...): **only maize for RL problems so far**.
- Various cultivars for each crop.

  Soils and cultivars have been calibrated by agronomists using extensive, multi-year real field trials.
- Weather: Recorded weather or weather generator (from all over the world).

- Observation = Collection of measurements amenable to a real farmer.

  Observed features are defined in a config yaml file.
- Objective: Can be customized, combining various performance indicators.

  The return function is defined in an easy to customize python file.

Out-of-the-box gym-DSSAT

- 3 built-in problems: Based on a maize field experiment [Morris *et al.*, 1982]

  How to manage irrigation or fertilization to maximize the yield of a certain cultivar of maize in a certain soil in certain weather conditions?

  Let's focus on the fertilization problem.

- We look for a policy: ∀ day: (day, amount of fertilizer)

  which is efficient and effective:

  trades-off yield vs. pollution and and cost.

- The daily return is defined by:

  $r(day) = \text{plant N uptake}(day, day+1) - 0.5 \times \text{fertilizer quantity}(day)$

- The goal is to maximize $\sum_{day=0}^{\text{harvest}} r(day)$.

A few preliminary results (1/3)

We compare:

1. A null policy which does not fertilize,
2. An expert policy used in the original 1982 field experiment,
3. A policy learned by RL (basic untuned PPO).

Policies 1. and 2. are fixed and deterministic.
Only the weather is stochastic.

Protocol:

- Null and expert policies are evaluated on 10^3 seasons.
- RL: Trained on 10^6 simulated seasons, then evaluated on 10^3 other seasons.

A few preliminary results (2/3)

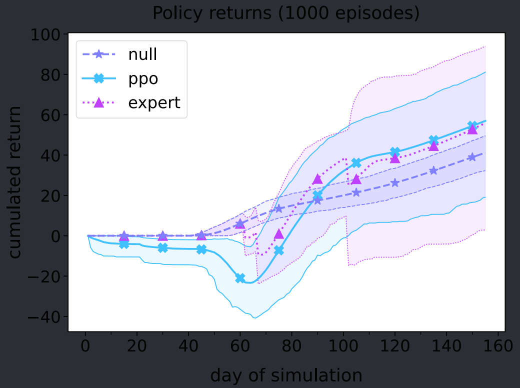

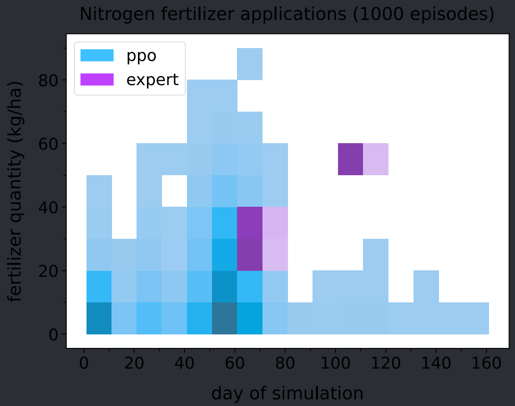

A few preliminary results (3/3)

	null	expert	PPO
grain yield (kg/ha)	1141.1 (344.0)	**3686.5** (1841.0)	3463.1 (1628.4)
massic nitrogen in grains (%)	1.1 (0.1)	**1.7** (0.2)	1.5 (0.3)
total fertilization (kg/ha)	**0** (0)	115.8 (5.2)	82.8 (15.2)
application number	**0** (0)	3.0 (0.1)	5.7 (1.6)
nitrogen use efficiency (kg/kg)	n.a.	22.0 (14.1)	**28.3** (16.7)
nitrate leaching (kg/ha)	**15.9** (7.7)	18.0 (12.0)	18.3 (11.6)

This table contains the mean (std dev) measured on 10^3 evaluation seasons.

In short: an untuned PPO learns a very good policy that balances the different criteria.

We obtain the same sort of results on the irrigation task.

The future of gym-DSSAT

Lots of things to do:

- Experiment with the existing out-of-the-box gym-DSSAT.

  E.g. studying the return function; fine tune RL algorithms able to generalize to various soil/weather/economical conditions; study the impact of global warming on learned policies.

- Study the multi-objective aspect of the problem.
- Extend the set of actions to those defined in DSSAT.
- Extending gym-DSSAT to the 41 other crops available in DSSAT.
- Extension to the management of more than 1 field.
- Keep up with DSSAT upgrades.

https://gitlab.inria.fr/rgautron/gym_dssat_pdi

gym-dssat@inria.fr

Check out https://arxiv.org/abs/2207.03270 for further details.

Farmgym

A gamified environment to understand dynamics of a farm and to learn to farm

[Maillard et al., 2023]

Farmgym: a gamified farming environment

- Farmgym is **stochastic** with a dependent ecosystem.
- The learner observe one entity at a time but for a cost.

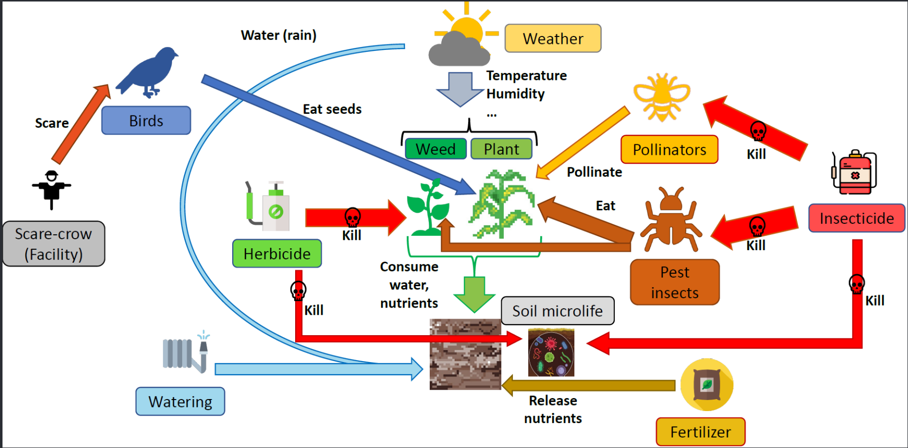

A farm in Farmgym

- **Farm:** a set of *Fields*
- **Field:** a 2d-array representing a *Plot* of land with set of *Entities*.
- **Entities:** Plants, Weeds, Soil, Pollinators, Pests

⟶

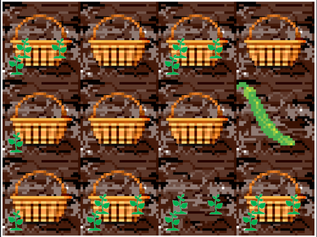

Entities: The Farm's Inhabitants with a Stochastic Dynamics

- Plant: sowing, harvesting, removing a plant,
- Weeds: removing weeds,
- Soil: watering the soil,
- Fertilizer: fertilizing,
- Cides: scattering herbicide, scattering insecticide,
- Facility: adding a scarecrow.....

Each entity has a **stochastic dynamics** modelled by a parametric **Generalised Linear Model (GLM)**

You can add your own entities and their dynamics, Farmgym is modular!

Actions and Scores

Actions can be divided into two categories:

- Observation: *Observe size of plant*
- Intervention: *Scatter fertilizer*

Each action has a **cost** and to trade-off with the **reward** (yield). Farmgym is **multi-objective**

- Soil-health
- Price of resources (manpower, fertilizer, pesticides...)

A new family of RL problems $\longrightarrow$ **Actively-Observed MDPs**

Coupled Dynamics

Effect of watering on yield with/without **pollinators** for **beans** vs **corn**.

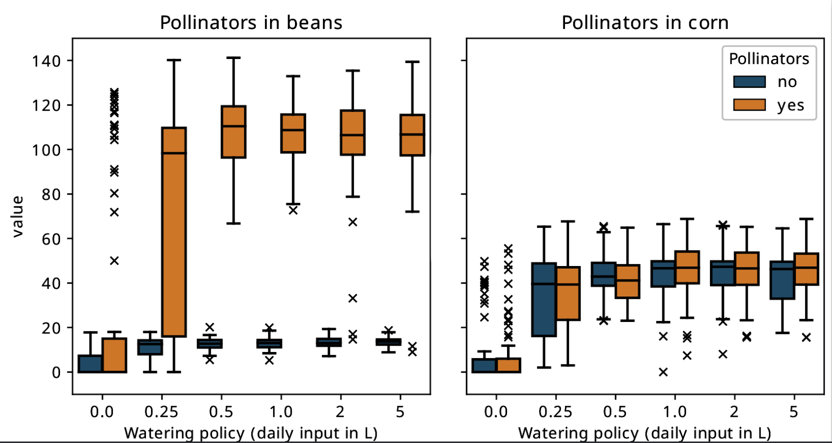

Learning using RL agents

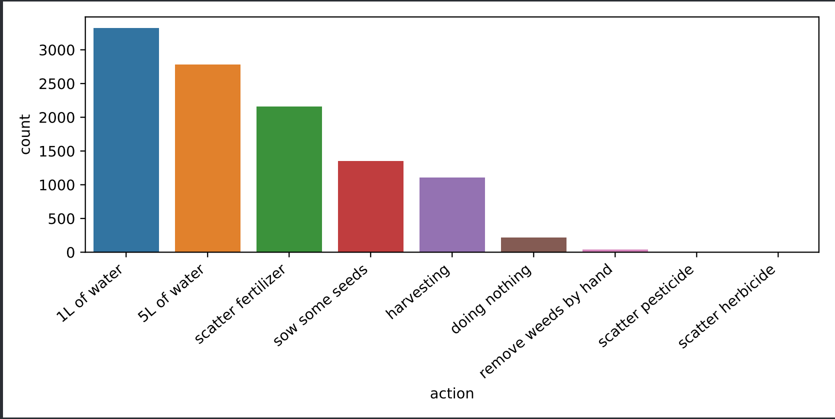

Actions of a trained PPO over 128 evaluation episodes shows RL agent learns to take certain actions and differentiate with others.

Shortcomings of Deep RL algorithms

Episode reward for PPO is around 100 after 400k steps while **Expert agent** achieves more than 400 episode reward.

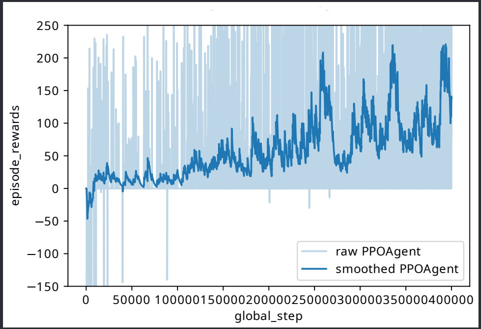

Usual Deep RL algorithms are not used to: take time into account, handle stochastic behavior, do an action only once (harvest, sow...)

Interactive demo

You can be the agent.

Colab notebook for interactive play: `https://tinyurl.com/farmgym-demo`

Farmgym repository: `https://github.com/farm-gym/farm-gym`

ForestDefender

Modelling impacts of diversity and extreme events on forest growth

Forest: Dynamics, Actions and Rewards

Dynamics: Model for trees interacting with their neighbours

- A tree grows at a faster rate when its small and eventually its growth saturate
- Rate of tree (height) growth is higher if the neigbhouring trees have closer or lower height

Actions:

- The agent can either let a tree grow or cut it and directly plant a new one
- Multiple trees can be cut at once
- An harvest action takes place before the growth step and the storms

Rewards:

- While harvesting, reward is proportional to square of the size of the harvested trees

Simulating the Forest

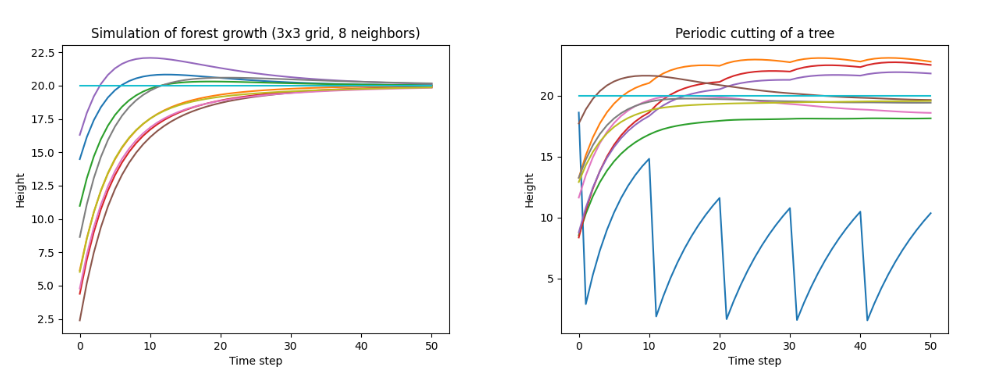

Figure: Evolution of tree height, Left: No action, Right: Periodical harvest of one tree (after its neighbors are fully grown, the regrowth become slower)

A Small Forest

1.png

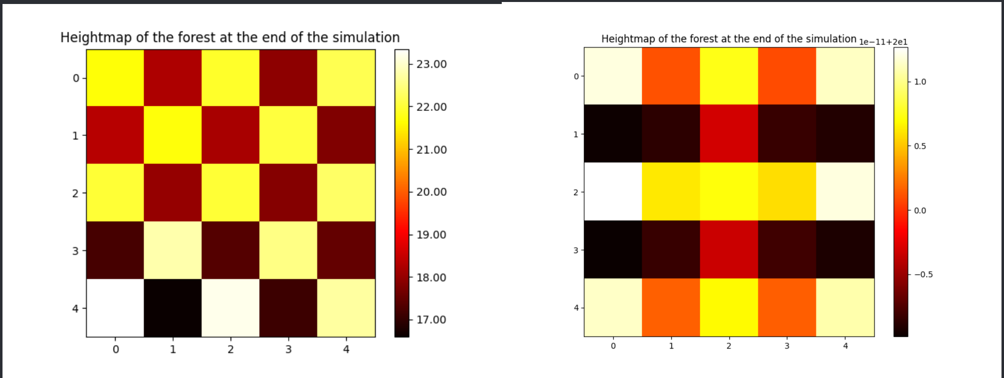

A Large Forest

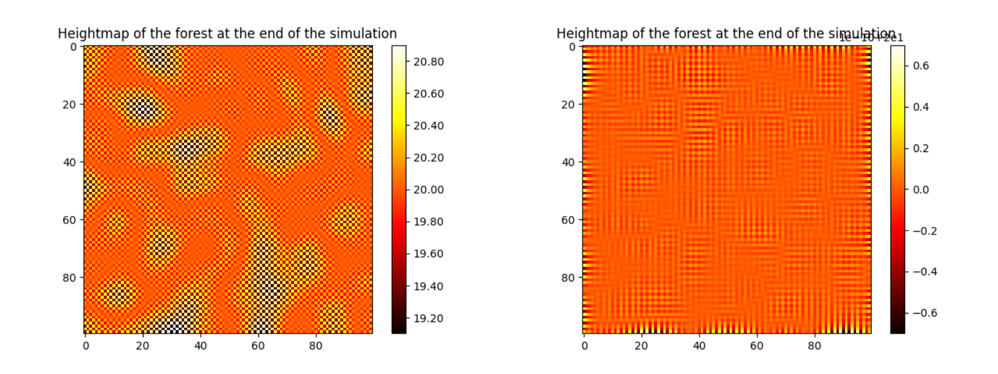

Figure: Left: Interaction with 4 nearest neighbors, Right: Interaction with 8 nearest neighbors

Learned Harvesting Policy

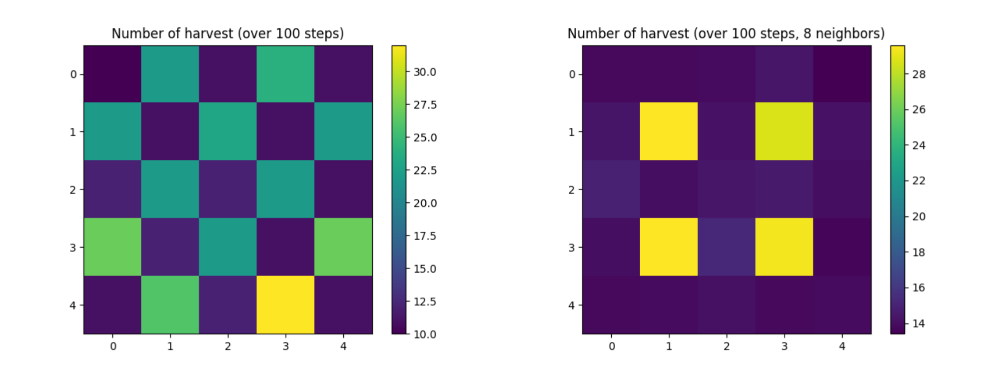

Policy learned using PPO. Left: Interaction with 4 nearest neighbors, Right: Interaction with 8 nearest neighbors

A Model of Storm

- At each time step, the forest experiences a storm with fixed probability p_{storm}
- In case of a storm, each tree is removed independently with a probability p_t^i that depends on the height of its neighbours
- The higher the neighbors, the less likely a tree is to be destroyed
- Destructive power of the storm D_t

Threshold Policy under Storm Risk

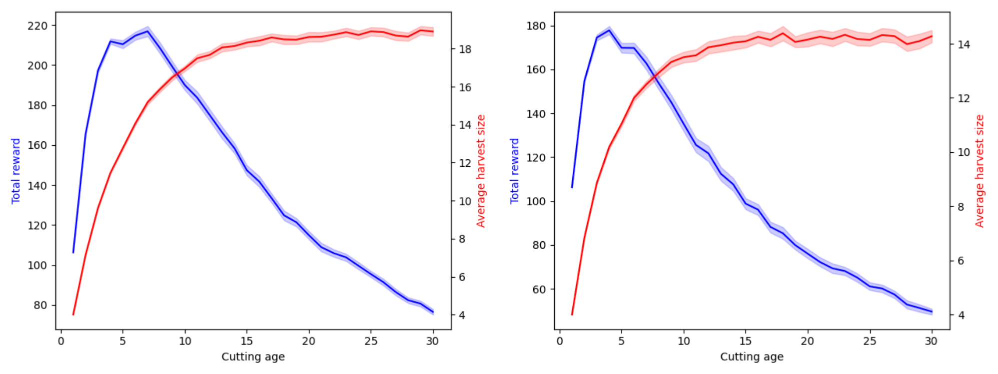

Left: Periodic harvest with storm probability 0.05, Right: Periodic harvest with storm probability 0.2

Learned policy under storm risk

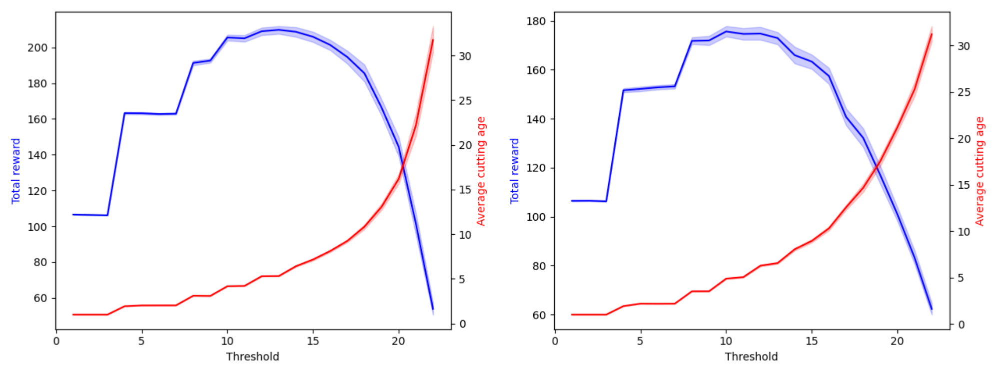

Left: Learned policy without storm, Right: Learned policy with storm probability 0.05

A Synergistic Future: Challenges and Opportunities

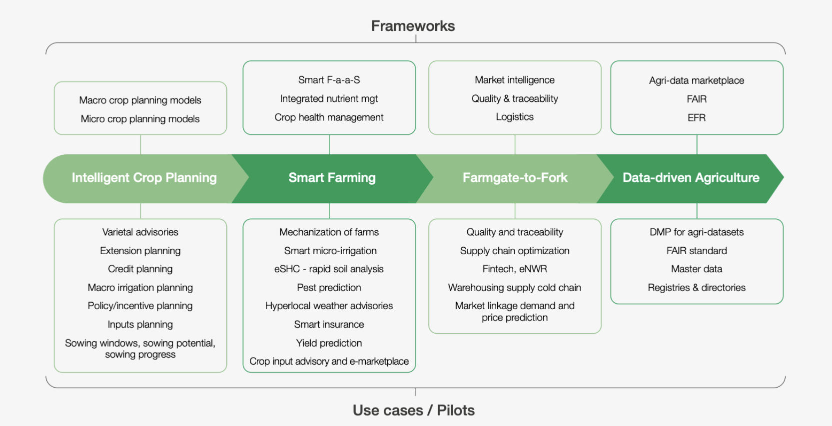

Welcome to this journey of developing and deploying AI for Agricultural and Ecological Sciences!

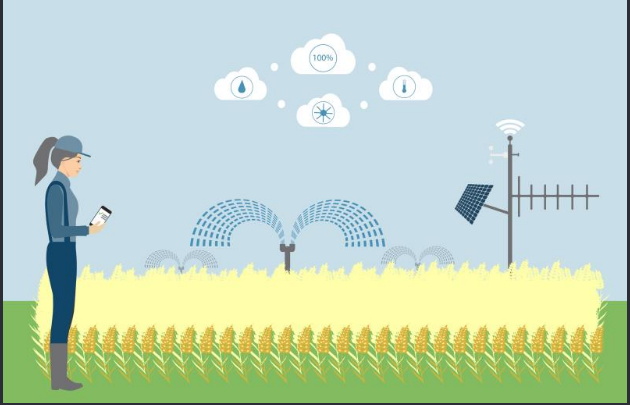

Thanks to the members and collaborators in Scool, who have been central to develop these works.

References

[Gautron et al., 2023] Gautron, R., Padrón, E. J., Preux, P., Bigot, J., Maillard, O.-A., Hoogenboom, G., and Teigny, J. (2023).
Learning crop management by reinforcement: gym-dssat.
In *AIAFS 2023-2nd AAAI Workshop on AI for Agriculture and Food Systems*.

[Jones et al., 2003] Jones, J. W., Hoogenboom, G., Porter, C. H., Boote, K. J., Batchelor, W. D., Hunt, L., Wilkens, P. W., Singh, U., Gijsman, A. J., and Ritchie, J. T. (2003).
The dssat cropping system model.
European journal of agronomy, 18(3-4):235–265.

[Maillard et al., 2023] Maillard, O.-A., Mathieu, T., and Basu, D. (2023).
Farm-gym: A modular reinforcement learning platform for stochastic agronomic games.
In *2nd AAAI Workshop on AI for Agriculture and Food Systems*.

[Sutton and Barto, 2018] Sutton, R. S. and Barto, A. G. (2018).
Reinforcement learning: An introduction.
MIT press.

RL agent and gym

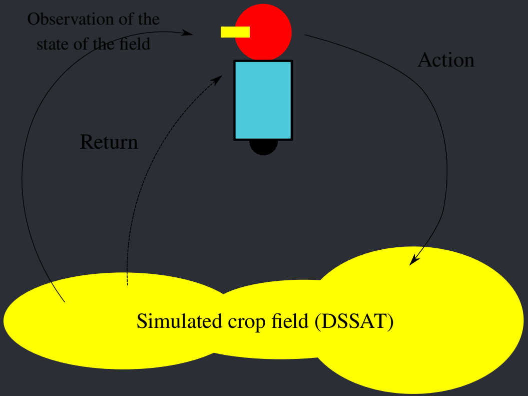

- Create the environment env
- `env.reset ()`
- Iterate on: one day interaction:
  - Choose the next action to perform
  - ```
    next_observation,
        return, finished?,
           d = env.step (action)
    ```
  - update the learning agent

In gym-DSSAT, the agent really gets an observation, not a full state.

Currently, actions consist in applying a certain amount of fertilizer and a certain amount of irrigation per day (per unit of surface).
    ```

Under the hood

- DSSAT is a large software program written in Fortran (300 klocs, 450 files).
- DSSAT reads a set of configuration files, runs the simulation accordingly, and outputs result files.
  No notion of the interaction loop required by RL agents.
- Today, almost all RLers use Python as a scripting language and know nothing about Fortran.
- ⇒ A python/Fortran connection is necessary.

  Based on the PDI library (`https://pdi.dev/master`).
  PDI makes the interaction between python and the information processed in DSSAT much easier to configure.