Learning Crop Management by Reinforcement: gym-DSSAT

Romain Gautron [1,2] Emilio J. Padrón [3] Philippe Preux [4] Julien Bigot [5] Odalric-Ambrym Maillard [4] Gerrit Hoogenboom [6] Julien Teigny [4]

[1]AIDA, Université de Montpellier, CIRAD, France [2]CGIAR Platform for Big Data in Agriculture, Colombia [3]Universidade da Coruña–Computer Architecture Group & CITIC, Spain

[4]Scool: Inria, Université de Lille, CNRS, France [5]Université Paris-Saclay, UVSQ, CNRS, CEA, Maison de la Simulation, France [6]University of Florida Gainesville, USA

`gym-DSSAT`

A `gym` environment for realistic crop management tasks, that is easy to use for training Reinforcement Learning (RL) agent with the **Decision Support System for Agrotechnology Transfer** (`DSSAT`) crop simulator coupled with the `WGEN` **stochastic weather generator**.

DSSAT: **state-of-the-art** *Fortran* mechanistic crop growth simulator.

gym: **standardized** *Python* API to connect a RL agent with a simulator of its environment.

⇒ `gym-DSSAT` is backed by the PDI library which allows loose coupling interaction between Fortran and Python code.

- gym-DSSAT is an on-going effort: 1st version end of 2021.
- DSSAT offers a vast amount of possible simulations. gym-DSSAT currently handles some of them.
- Beyond DSSAT, our approach may be used to turn other C/C++/Fortran monolithic mechanistic models into RL envs.

Crop management problems in `gym-DSSAT`

1. **Fertilization problem**: the agent can apply every day a certain quantity of nitrogen (Table 1). Crops are rainfed, and no irrigation is applied during the growing season. We crafted the default fertilization **return function** as:

$$r(t) = \underbrace{\texttt{trnu}(t,t+1)}_{\text{plant nitrogen uptake (kg/ha)}} - \underbrace{0.5}_{\text{penalty factor}} \times \underbrace{\texttt{anfer}(t)}_{\text{fertilizer quantity (kg/ha)}} \quad (1)$$

2. **Irrigation problem**: the agent can provide every day a certain amount of water to irrigate, as indicated in Table 1.
3. Mixed **fertilization and irrigation problem**: combines both the aforementioned decision problems, i.e. the agent can fertilize and/or irrigate every day.

Action	Description	Range
fertilization	nitrogen amount (kg/ha)	[0,200]
irrigation	water amount (L/m^2)	[0,50]

Table 1. Daily actions available in `gym-DSSAT`

DayAfterPlanting	Quantity (kg N/ha)
40	27
45	35
80	54

Table 2. Expert fertilization policy

Custom scenario definition

- The **observation space** can be <u>easily modified</u> by editing a YAML config. file.
- The **return functions** can also be <u>easily modified</u> by editing a standalone Python file.

Features:

→ Soil conditions and weather (simulated or measured) are available off-the-shelf based on **hundreds of example of real-world measures**.

→ `gym-DSSAT` allows built-in climate change (e.g. for atmospheric CO_2, temperature) for **non-stationary** crop management problems.

A use case: learning an efficient maize fertilization

An episode spans one growing season, i.e. a **finite** number of time steps. The objective function is defined as: $\sum_{t=0}^{\text{harvest}} r(t)$.

Table 3 provides the observation space. We consider three policies:

- The ***null*** policy never fertilizes. As there is always nitrogen in soil before cultivation [3], the reference experiment, or *control*, is the null policy.
- The ***expert*** policy is the one published in the original maize field experiment [1] and defined in Table 2.
- The ***PPO*** policy learned by the Proximal Policy Optimization [4] RL algorithm, as implemented in `Stable-Baselines3 1.4.0` [2] with default hyperparameters as baseline. We trained PPO for 10^6 iterations.

Variable	Definition
`istage`	`DSSAT` maize growing stage (categorical)
`vstage`	vegetative growth stage (number of leaves)
`topwt`	above the ground crop biomass (kg/ha)
`grnwt`	grain weight dry matter (kg/ha)
`swfac`	index of plant water stress (unitless)
`nstres`	index of plant nitrogen stress (unitless)
`xlai`	leaf area index (m^2 leaf/m^2 soil)
`dtt`	growing degree days (°C.day)
`dap`	days after planting (day)
`cumsumfert`	cumulative nitrogen fertilization (kg N/ha)
`rain`	rainfall for the current day (L/m^2/day)
`ep`	actual plant transpiration rate (L/m^2/day)

Table 3. Default observation space for the fertilization task

Conclusions from experimental results

- An **untuned** PPO was able to learn sustainable fertilization and irrigation policies.
  ⇒ RL has a great potential to learn **sustainable crop management practices** using `gym-DSSAT`.
- `gym-DSSAT` allows exploration of many other agricultural decision problems, e.g. multi-year crop management with crop rotations.

If you want to collaborate for developing `gym-DSSAT`, contact us at:

gym-dssat@inria.fr

References

[1] JM Bennett, LSM Mutti, PSC Rao, and JW Jones. Interactive effects of nitrogen and water stresses on biomass accumulation, nitrogen uptake, and seed yield of maize. *Field Crops Research*, 19(4):297–311, 1989.

[2] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. Stable baselines. `https://github.com/hill-a/stable-baselines`, 2018.

[3] Thomas F Morris, T Scott Murrell, Douglas B Beegle, James J Camberato, Richard B Ferguson, John Grove, Quirine Ketterings, Peter M Kyveryga, Carrie AM Laboski, Joshua M McGrath, et al. Strengths and limitations of nitrogen rate recommendations for corn and opportunities for improvement. *Agronomy Journal*, 110(1):1, 2018.

[4] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

Experimental results

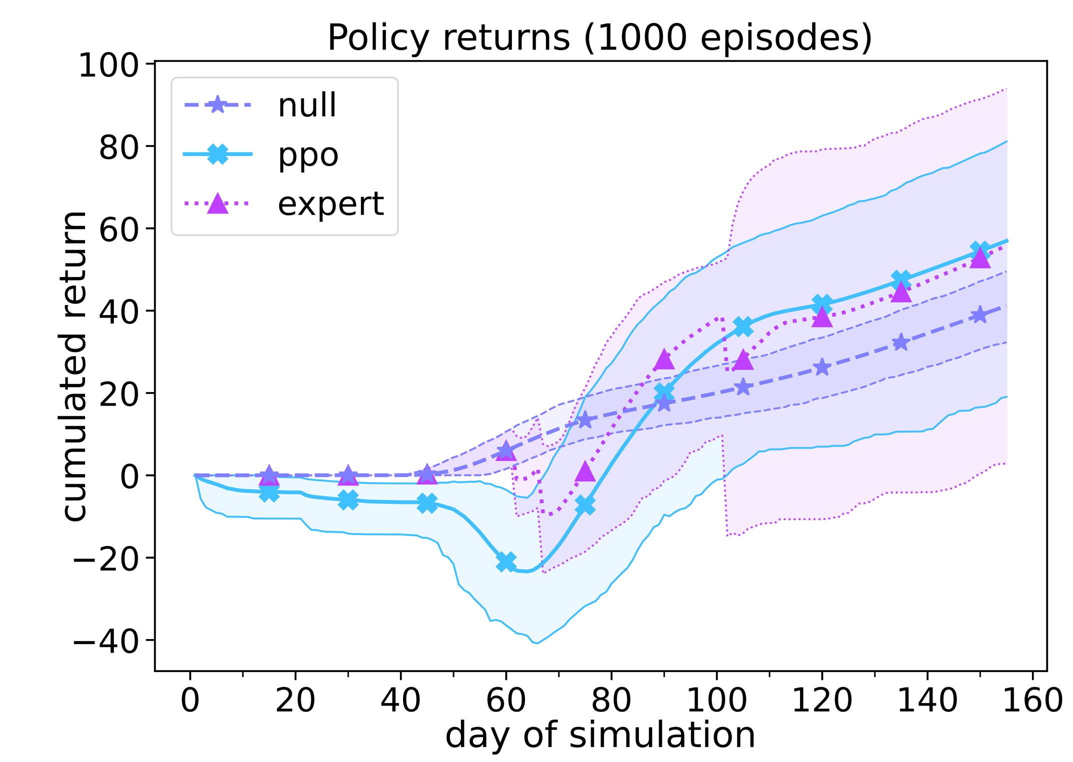

Figure 1. Mean cumulated return of each of the 3 policies against the day of simulation. Shaded area displays the $[0.05, 0.95]$ quantile range for each policy.

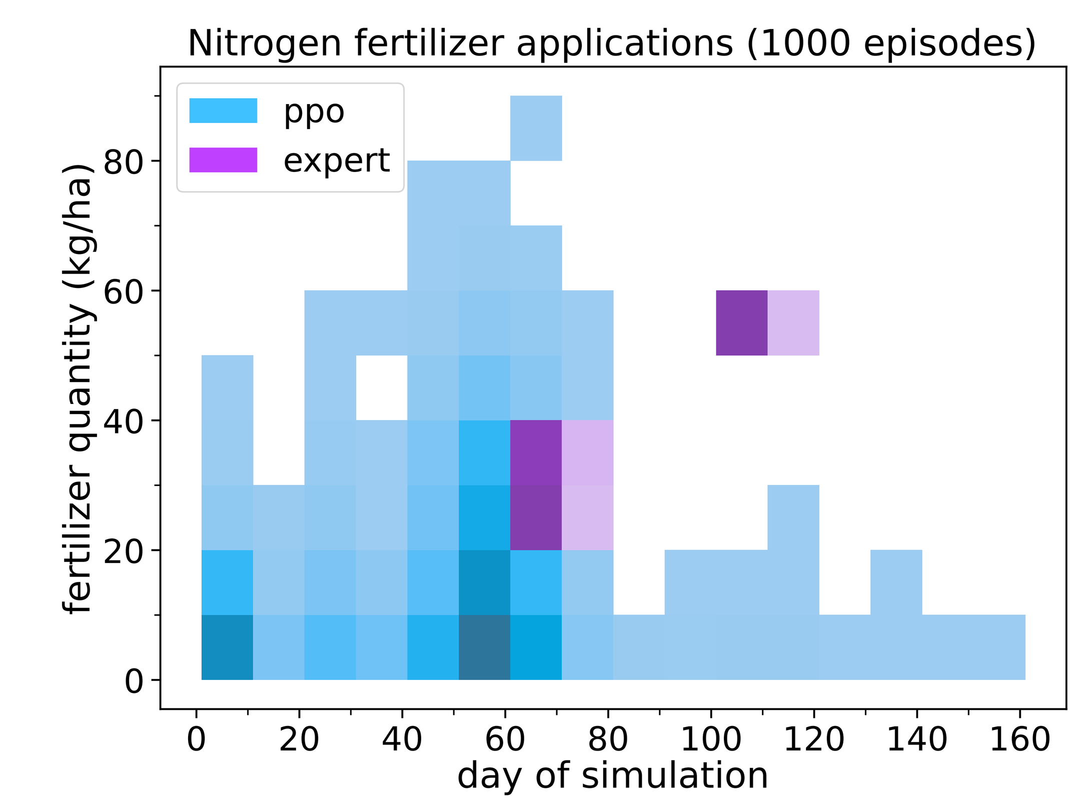

Figure 2. 2D histogram of fertilizer applications (the darker the more frequent).

	null	expert	PPO
grain yield	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	**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	n.a.	22.0 (14.1)	**28.3** (16.7)
nitrate leaching	**15.9** (7.7)	18.0 (12.0)	18.3 (11.6)

Table 4. Mean (st. dev.) of performances computed over 1000 episodes. **Bold** numbers indicate the best performing policy. See Table 5 for interpretation.

Variable	Definition	Comment
`grnwt`	grain yield (kg/ha)	quantitative objective to be maximized
`pcngrn`	nitrogen content in grains (%)	qualitative objective to be maximized
`cumsumfert`	total fertilization (kg/ha)	cost to be minimized
–	application number	cost to be minimized
–	nitrogen use efficiency (kg/kg)	agronomic criteria to be maximized
`cleach`	nitrate leaching (kg/ha)	loss/pollution to be minimized

Table 5. Performance indicators for fertilization policies. '-' means the variable is not provided by default but it can be derived.