MINISTÈRE
DE L'ENSEIGNEMENT
SUPÉRIEUR
ET DE LA RECHERCHE

Liberté
Égalité
Fraternité

anr agence nationale de la recherche

Appel à projets Thématique Spécifique en Intelligence Artificielle (TSIA) – Edition 2025

Acronym	AGRILORE		
Project Title	**Agr**oecolog**I**ca**L** decision making and **O**ptimization with **RE**inforcement learning		
Total budget	678k€	**Duration**	48 months
Keywords	reinforcement learning, transfer learning, data science, agroecology		

Project Coordinator (main contact for the proposal)

Name	MAILLARD
Institution/Department	CENTRE INRIA DE L'UNIVERSITE DE LILLE
Address	40 Avenue Halley, 59650 Villeneuve d'Ascq
Country	FRANCE
Phone	0359577811
E-mail	Odalric.maillard@inria.fr

Table of persons involved in the project Machine & reinforcement learning / Agronomy/ Data science

Partner(Institution / Department)	**Last Name**	**First name**	**Current position**	**Role & responsibilities in the project (4 lines max)**
INRIA/SCOOL	Maillard	Odalric-Ambrym	Research officer	Project Coordinator, Scientific Manager INRIA
//	Basu	Debabrota	Researcher	Leader WP3
CIRAD/PERSYST UPR AÏDA.	Naudin	Krishna	Head of research unit	Scientific Manager CIRAD, co-leader WP6
//	Midingoyi	Cyrille	Research Engineer	Leader WP1
//	Auzoux	Sandrine	Research Engineer	Leader WP2, Participant WP5
//	Christina	Mathias	Researcher	Leader WP4, Participant WP1
//	Couëdel	Antoine	Researcher	Participant WP1,2,4 Expert STICS, intercropping.
INRAE/UMR MISTEA	Verzelen	Nicolas	Research Professor	Scientific Manager INRAE MISTEA, Leader WP3
//	Metivier	David	Researcher	Participant WP3, Statistical Machine Learning
//	Sanchez	Isabelle	Research Engineer	Participant WP3 Data scientist
INRAE/UMR AGIR	Gaudio	Noémie	Researcher	Scientific Manager INRAE AGIR, co-leader WP2
//	Casadebaig	Pierre	Researcher	Participant WP3, Agronomy & statistical models
//	Meunier	Clémentine	Research Engineer	Participant WP5.2 Link with farmers

MINISTÈRE
DE L'ENSEIGNEMENT
SUPÉRIEUR
ET DE LA RECHERCHE

Liberté
Égalité
Fraternité

anr agence nationale de la recherche

Appel à projets Thématique Spécifique en Intelligence Artificielle (TSIA) – Edition 2025

Summary of the project in French (publishable non-confidential abstract, max. 1/2 page):

L'intensification agroécologique constitue une réponse essentielle aux défis actuels de la sécurité alimentaire et du changement climatique. Parmi les leviers agroécologiques, la diversification des cultures, notamment les cultures associées (i.e. au moins deux espèces qui poussent dans un même champ), offrent un potentiel agronomique majeur. Leur mise en œuvre repose cependant sur des connaissances encore limitées. La production de références expérimentales étant longue et coûteuse, la modélisation mécaniste s'est imposée comme une alternative efficace. Toutefois, les modèles de culture existants, comme STICS ou DSSAT, restent principalement conçus pour des monocultures et peinent à représenter les interactions complexes des cultures associées.
Ce projet vise à dépasser ces limites en intégrant des approches d'intelligence artificielle (IA), notamment l'apprentissage par renforcement (RL), pour optimiser la prise de décision en conditions d'incertitude. Forts d'une collaboration interdisciplinaire déjà éprouvée dans un cadre simple, nous capitalisons sur cette dynamique pour relever un défi majeur : adapter le couplage RL-Agro au cas emblématique des cultures associées. Les données sur ces agrosystèmes étant rares et incomplètes, l'originalité réside dans une stratégie d'apprentissage progressive : entraîner d'abord les algorithmes sur des modèles monospécifiques bien maîtrisés, avant de les ajuster à partir d'un petit nombre de données réelles sur les cultures associées.
Les objectifs sont multiples : (i) sélectionner et adapter des algorithmes de RL adaptés aux contraintes agronomiques, (ii) développer un outil d'interaction entre IA et modèles de culture et de structuration des données agroécologiques, (iii) concevoir des algorithmes frugaux en données et en énergie, et (iv) structurer une nouvelle communauté autour de ces outils d'IA pour outiller la transition agroécologique. En s'appuyant sur l'expertise française en modélisation agronomique, ce projet ambitionne de poser les bases d'une nouvelle génération d'outils d'aide à la décision pour une agriculture durable.

Summary of the project in English (publishable non-confidential abstract, max. 1/2 page):

Agroecological intensification is a key response to the current challenges of food security and climate change. Among the agroecological levers, crop diversification, especially intercropping (i.e. growing least two different crop species in the same field), offers major agronomic potential. However, their implementation is still based on limited knowledge. As the production of experimental references is long and costly, process-based (mechanistic) modeling has emerged as an effective alternative. However, existing crop models, such as STICS or DSSAT, are still designed for monocultures and struggle to represent the complex interactions of intercrops.
This project aims to overcome these limitations by integrating artificial intelligence (AI) approaches, notably reinforcement learning (RL), to optimize decision-making under conditions of uncertainty. Building on a proven interdisciplinary collaboration in a simple setting, we leverage this dynamic to tackle a major challenge: adapting the RL-Agro coupling to the emblematic case of intercropping. Since data on intercropping agrosystems are scarce and incomplete, the originality lies in a progressive learning strategy: first training the algorithms on well-mastered monospecific models, before adjusting them on the basis of a small number of real intercropping data.
The objectives are multiple: (i) select and adapt RL algorithms adapted to agronomic constraints, (ii) develop a tool making AI and crop models interacting and for structuring agroecological data, (iii) design data- and energy-efficient algorithms, and (iv) structure a new community around these AI tools to support agroecological transition. Drawing on French expertise in agricultural modeling, this project aims to lay the foundations for a new generation of decision-support tools for sustainable agriculture.

1. Scientific and technical excellence, innovation and breakthrough in R&D

Agriculture faces growing challenges of food security, climate change, and biodiversity loss, requiring a shift towards more sustainable practices. Agroecology, especially crop diversification strategies such as intercropping, offers agronomic and environmental benefits [1]. However, the complexity of multispecies interactions occurring in intercropping poses a major challenge for experimental research, making mechanistic modeling crucial. While current crop models, such as STICS [2] and DSSAT [3], provide valuable predictive tools, they are primarily designed for monocultures and they often overlook the complexity of interactions between living components of agroecological cropping systems. Addressing these limitations requires i**nnovative approaches that integrate AI-driven optimization with agronomic modeling to support the agroecological transition**.

Process-based crop simulation models are computational tools designed to mimic daily crop growth in response to environmental conditions and management practices [4]. They complement field experiments and their strength lies in their ability to simulate multiple abiotic stresses (e.g., water, temperature), enabling comprehensive evaluations of crop performance under varying conditions. However, these models, initially designed for monocultures, struggle to represent the complex interactions occurring in agroecological systems like intercropping [5]. This gap underscores the **project's innovative thrust**: enhancing crop models to better capture plant-plant interactions and resource sharing into intercropping systems. While recent advances have improved the representation of few resource sharing processes [6], models remain limited in scopes [7]. Our project disrupts crop modeling by integrating **reinforcement learning (RL)** techniques with crop models based on the methodology initiated by recent studies [8,9], but going further by advancing sequential learning, in order to learn from a model simulating monoculture and then learn the behavior of intercropped species on a smaller, real-world dataset. This approach will increase adaptability, predictive power of simulations, and scalability to address the complexity of intercropping. By optimizing agricultural practices through RL, the project seeks to go beyond current methods, overcoming challenges associated with parameter-heavy models and computational inefficiencies. **Simulation-based optimization** was already and successfully used across various scientific fields (e.g. in microscopy with the pySTED platform; [10]), which aligns with AgriLORe's objectives related to efficient exploration and multi-objective optimization of intercropping. Moreover, while RL has been applied in agricultural decision-making through some simulation environments [9,11,12], existing frameworks focus on optimizing irrigation, fertilization, and crop planning. Our project extends these methods to intercropping by integrating advanced RL algorithms with mechanistic crop models to handle high-dimensional action spaces and inter-species interactions. **Bandit Optimization for Experimental Design:** a relevant precedent for our approach can be found in the **kernelized bandit optimization framework** [13]. This method was successfully applied to optimize optical microscopy parameters over short experimental sequences (60 trials) in progressively complex setups [14]. This work was later extended into a Python-based simulation environment for AI-assisted microscopy [10], showcasing the potential of bandit methods for sequential experimental design, especially when the data is scarce and cost or time-intensive to collect. Our project builds upon this foundation but introduces critical innovations tailored to agroecology. Specifically, we focus on risk-aware, non-parametric bandit strategies and extend beyond stationary multi-armed bandits (MAB) to dynamic Markov decision processes (MDPs). These extensions are necessary to accommodate the complexities of intercropping, where plant-plant interactions and environmental fluctuations introduce a rich but uncertain decision landscape. Our work aligns with recent advances in bandit literature, particularly in experimental design for sequential decision-making. The exploration of multi-armed bandit strategies for experimental optimization has been a growing research direction, along with hybrid performance criteria that balance regret and statistical power [15], as well as group sequential testing [16]. Building on this, our approach extends modern bandit techniques to risk-aware and non-parametric strategies, inspired by discrete settings and real-world applications such as [8]. Additionally, we adapt kernelized bandit frameworks [13] to exploit structural properties of the action space, while integrating contextual information such as weather variables to improve decision-making across various agricultural environments. **Recent advances in RL provide new opportunities for agricultural optimization**, particularly in data-scarce environments. Few-shot learning methods [17] offer promising avenues for adapting RL agents to new farming conditions with minimal training data. Our project aims to develop crop-specific simulators that serve as benchmarks for evaluating RL-based decision-making strategies while leveraging prior successes in RL-driven plant cultivation [9].

A variety of models, simulators, and algorithms have been proposed for agricultural applications, including mechanistic crop models [18,19], RL-based approaches for resource management [20–22], and recent simulation frameworks such as gym-based agricultural environments [9]. Our project builds on these foundations, integrating RL techniques with structured agronomic knowledge to optimize intercropping strategies.

Hybrid and Continual Reinforcement Learning for Dynamic Agroecosystems Beyond bandit strategies, our approach considers the dynamic nature of agroecosystems, situating our framework within the broader field of Hybrid Reinforcement Learning (HRL) [23,24], which combines online and offline learning paradigms. Transfer learning plays a crucial role in our methodology, as our RL agents must generalize across diverse cropping environments [25]. Furthermore, our project aligns with the principles of Continual Reinforcement Learning (CRL) [26], given that we aim to develop adaptive policies across multiple, closely related agroecosystems rather than a single fixed environment. The challenge of handling model shifts between offline experimental data (collected from intercrop trials) and online learning within simulators adds another layer of complexity, necessitating novel RL strategies to bridge this gap.

The innovation spans the entire **technology value chain**, from refining crop model algorithms to integrating AI-driven decision-support frameworks. This approach bridges gaps in simulating plant interactions and environment variability, enabling the development of context-aware crop models. The project is therefore at the forefront of scientific and technological advances in the evolving field of agroecology, and offers a novel pathway to improve crop modeling.

Our project aligns with the TSIA 2025 call by strengthening French expertise in AI-driven agroecological modeling through collaboration between INRIA, INRAE, and CIRAD. By improving a recognized French soil-crop model, we strengthen France's leadership in this field. The project directly contributes to leveraging AI for scientific discovery, applying RL to optimize agroecological cropping systems—perfectly aligned with the TSIA theme of hybridizing AI with domain knowledge. It also meets the call's objective of developing novel AI-based simulation and experimentation methods, not only through research but also via open-source tools and community-building workshops, fostering a national AI-agronomy ecosystem. Finally, our approach is inherently frugal, optimizing data-efficient AI methods tailored to agroecological systems, ensuring robust, interpretable, and sustainable solutions.

We will develop **Gym-STICS**, an open-source RL framework integrating AI with STICS crop model, benefiting both French and international agronomic research communities. We will organize two interdisciplinary workshops (~50 participants each) to foster collaborations, while reaching a broader audience—STICS already serves ~3,000 users across 110 countries, and AI integration is expected to significantly expand its adoption. We will establish benchmark datasets, define evaluation metrics, and publish in leading AI and agronomy revues. Our AI-driven decision-support tools will be validated on real-world case studies, bridging the gap between academic research and practical deployment.

This project ensures rigorous scientific foundation for AI-assisted agroecological decision-making. By coupling RL with biophysical crop models, we ensure that AI-driven recommendations are grounded in validated agronomic science, providing a reliable framework upon which practical decision-support tools can be developed for researchers, advisors, and ultimately farmers. The potential benefits are significant: AI-enhanced agroecological systems will accelerate innovation, reduce adoption risks, and improve access to expert knowledge, supporting a more resilient and sustainable agriculture. Our approach is further designed to be flexible and widely applicable, from smallholder farms in the Global South to large-scale European agricultural systems. Key risks include misaligned recommendations due to data biases or model limitations. To mitigate this, we ensure strict validation through expert oversight (INRAE, CIRAD) and robust testing on diverse and representative datasets, addressing non-stationarity and real-world variability. Finally, by prioritizing frugal AI approaches, we ensure long-term applicability with minimal environmental impact.

2. Feasibility

Work Package 1: Interfacing Crop Models with Reinforcement Learning (RL)

[Resources: 18m Postdoc@INRIA (M3:M22), 15k AgroClim expertise]

Leaders: C. Midingoyi (CIRAD), OA. Maillard (INRIA). **Participants:** A. Couëdel, M. Christina (CIRAD)

Objective: Develop a robust and flexible interface between crop models and reinforcement learning (RL) to enable decision-making algorithms to interact with crop simulations. This requires bridging monolithic crop models, which treat actions as input variables, with sequential decision-making frameworks, which rely on state-action-reward structures. A key challenge is to design a modular, generalizable methodology integrating STICS into a standard RL framework, paving the way for adaptive decision-making in agronomy.

A key challenge in applying RL to agronomy lies in bridging the conceptual and technical gap between mechanistic crop models and decision-making frameworks. Crop models like **STICS** are designed as **biophysical process simulators**, where user-defined management actions act as input variables, and the system evolves deterministically or stochastically based on predefined equations. In contrast, **RL** requires an **interactive decision-making structure**, where an agent sequentially selects actions based on an evolving state and receives rewards that guide learning. The goal of WP1 is to develop a **functional interface** between these two paradigms, enabling the use of crop models as environments where RL agents can learn adaptive management policies.

T1.1 Adapting the STICS Model to a Reinforcement Learning Framework. The first step will focus on STICS, identifying how model inputs and outputs can be mapped into a state-action-reward framework. A methodological approach will be developed to extract relevant decision variables, define time steps for decision-making, and ensure that feedback loops align with RL dynamics. This will allow STICS to function as a **decision-based simulator** while preserving its biophysical realism. This task will be conducted in collaboration with a STICS expert (INRAE AgroClim, García de Cortázar-Atauri team's) managing the official implementation of the STICS model through specific funding for the expertise service. This collaboration will ensure that the interface remains faithful to agronomic knowledge while being adapted to RL workflows.

T1.2 Development of Gym-STICS OpenAI Gymnasium-compatible interface. It will allow RL agents to interact with the model through a standardized API. This will involve structuring simulation calls, defining step-wise interactions, and ensuring computational efficiency. The interface will support both fully observable and partially observable settings, accommodating different RL paradigms, from contextual bandits to full Markov decision processes (MDPs).

T1.3 Generalize the methodology to other crop models. This will involve developing an extensible architecture to allow different mechanistic models (e.g. DSSAT, APSIM) to be coupled with RL agents in a flexible and scalable manner. We will build on our previous interfacing of gymnasium and DSSAT [9]. Particular attention will be given to enabling structured decision-making approaches, including factorial experimental designs with Bandits and long-term optimization with RL.

The postdoc leading this WP will work closely with an expert in STICS (INRAE agroclim) and an RL researcher (Mailalrd, INRIA), under the supervision of an agronomy modeling specialist (Midingoyi, CIRAD). By the end of WP1, we will have an operational **Gym-STICS interface**, along with a structured methodology for integrating RL with mechanistic crop models. This will serve as the **foundation for WP3**, where RL algorithms will be designed to leverage this interface, and for **WP4**, where applied scenarios in intercropping will be explored.

Deliverables

D1.1. Functional Gym-STICS environment (**M20)**

D1.2. Scientific publication on extensible architecture to allow different mechanistic models to be coupled with RL agents **(M36)**

Work Package 2: Gathering and standardizing experimental data on intercrops
[Resources: 12m Postdoc@INRAE.Agir (M03-M14)]
Leaders: S. Auzoux (CIRAD), N. Gaudio (INRAE.Agir). **Participants**: A. Couëdel

Objective: WP2 capitalizes on existing experimental data related to intercrops by developing automated tools to structure and harmonize datasets according to FAIR principles (Findable, Accessible, Interoperable, Reusable). The objective is to ensure data quality and consistency for analysis, modeling, and reinforcement learning algorithms development. A key issue is to homogenize, qualify, and automate the cleaning and structuring of heterogeneous datasets while ensuring compatibility with WP4 and WP5.

Challenges: Implementing **RL for agronomic decision-making** requires high-quality and structured experimental data that captures complex crop responses under diverse conditions. However, these datasets are often **heterogeneous**, coming from different sources with varying levels of granularity, structure, and metadata [27]. The goal of WP2 is to **identify, standardize, and preprocess** intercropping experimental data to make it usable for RL-based decision-making.

This WP is structured around **tasks to address the aforementioned issues**, involving a postdoc specialized in data science supervised by two scientists in data science (CIRAD) and agroecology (INRAE):

Task 2.1 Identifying and gathering existing datasets on intercropping. This task will inventory and organize datasets coming from project partners' [28] and literature [29], while documenting experimental protocols, measured variables, and decision-making contexts. These datasets span multiple countries, crop types, and experimental conditions, providing a robust foundation for RL-training.

Task 2.2 Defining and applying standardized format using ontologies to harmonize datasets. We will homogenize the datasets identified, which will be structured heterogeneously (variable names, units, etc.), by defining a standardized format compatible with STICS and creating an experimental dataset dedicated to RL. This will leverage established crop modeling standards like ICASA [30] and the crop ontology concepts (http://www.cropontology.org). Advanced text mining approaches will support this process, justifying the recruitment of a data science post-doc specialist.

Task 2.3 Developing automated data pre-processing pipelines for STICS simulations and RL applications. A key technical contribution will be the development of automated pipelines to transform raw experimental data for RL models, handling data cleaning, missing values imputation, time series data normalization, and extraction of state-action-reward structures. These pipelines will ensure that preprocessed datasets can be easily integrated into STICS, enabling hybrid RL approaches to refine model-based simulations.

Task 2.4 Developing an open data hub for intercrop research. WP2 will lay the foundation for an open data hub, to share both curated experimental datasets and processed versions optimized for RL with the scientific community. This promotes transparency, reproducibility, and collaboration, encouraging researchers to contribute additional data and refine the integration process over time.

Implementation and Collaboration By the end of WP2, we will have a structured, standardized, and processed dataset repository, enabling **WP3** to develop RL methodologies that leverage real-world data, and **WP4** to validate RL strategies on empirical agroecological experiments.

Deliverables
D2a. Available experimental datasets on intercrops (M12)
D2b. Global homogenized datasets on intercrops (M24)
D2c. Data pre-processing pipelines for STICS simulations and RL applications (M36)
D2d. Scientific publication on data workflow process (M36)
D2e. Open data hub for the scientific community (M48)

Work Package 3: Hybrid Transfer Reinforcement Learning for Agroecological Decision-Making
[Resources: 1 PhD@Inria (M09-M45), 1/2 PhD@Inrae (M09-M45)]

Leaders: D. Basu (INRIA), N. Verzelen (INRAE.Mistea). **Participants**: D. Metivier (INRAE.Mistea), I. Sanchez (INRAE.Mistea), N. Gaudio (INRAE.Agir), P. Casadebaig (INRAE.Agir), OA. Maillard (INRIA)

Objective: Develop advanced RL and bandit algorithms tailored for agrosystem decision-making, focusing on hybrid transfer learning, offline fine-tuning, and contextual bandits. The key issue is to design RL techniques that can effectively leverage both simulated agronomic models and experimental data, ensuring robust policy learning despite the scarcity of real-world intercropping datasets. This WP is theoretical and methodological. We aim to produce high-impact research contributions in RL and applied mathematics along with open-source implementations.

Challenges: While RL is an efficient tool for algorithmic decision-making, it is often data intensive if we start from scratch. Thus, the RL agents for agroecological decision-making must often be trained on simulated data (e.g. Gym-STICS in **WP1**, CyclesGym [12] , Gym-DSSAT [9]). Deploying these models for real farms requires adapting RL agents with experimentally collected and informative data of the farm's contexts. With advent of LLMs, we find a burst in fine-tuning ML algorithms, including RL, to new environments. Presently, we do not know how optimal these fine-tuners are and what is the fundamental limit of adapting an RL algorithm to new data and environment. This is a relevant question in our setting as the agroecological data collected from real farms is costly and might be different from those of the simulators. The other issue is whether we can we use RL methods to design adaptive and sequential experiments rather than the static ones currently used. Understanding the limits of the present algorithms is key to design more sample efficient algorithms.

This WP is structured around **two research axes to address the aforementioned two issues**, each led by a dedicated **PhD student**:

T3.1 Hybrid Transfer RL with Offline Fine-Tuning. This axis focuses on developing hybrid transfer learning techniques for RL that progressively integrate simulator-based training, offline fine-tuning, and future online adaptation. The first stage involves training a set of RL agents within the crop simulation environments (e.g. via Gym-STICS from **WP1**), using structured exploration strategies to efficiently learn a generalizable base policy. The second stage introduces offline fine-tuning, where this base policy is adapted to real-world intercropping settings using experimental agronomic data (carefully collected from **WP2**). Unlike standard RL transfer, this phase relies on distributional shift-aware offline RL techniques to keep the efficiency of learning and planning from the simulated data while extracting insights from the limited real-world data. In this case, we will explore the statistical limits of offline fine tuning of RL agents under distribution shifts and design sample-frugal algorithms for real-world deployment. This PhD will be co-supervised by an researcher expert in RL (Basu, INRIA) and will benefit from close interaction with WP1, ensuring a strong interdisciplinary foundation.

T3.2 Contextual Bandits for Factorial Group-Sequential Experiments. This axis explores how contextual bandits can enhance policy selection in factorial agronomic experiments. Rather than treating the simulator as an unconstrained RL environment, policy learning is framed as a contextual bandit problem, where different agronomic strategies are structured as bandit arms. This approach enables data-efficient learning, ensuring that key agronomic trade-offs are systematically explored before full RL training. Given the scarcity of intercropping data, group-sequential design of experiments (DoE) principles will be incorporated to optimize learning from limited, structured data. While this WP does not fund real-world field trials, its structured experimental design methodology will lay the foundation for future adaptive agroecological decision-making. This PhD will be supervised by three researchers in DoE (Verzelen, INRAE), in agroecology (Gaudio, INRAE) and bandit RL (Maillard, INRIA), ensuring rigorous integration of statistical experimental design with AI-driven decision-making and experimental validation on real data.

Implementation and Collaboration: The expected outcomes include novel RL algorithms for both the problems, open-source implementations, and high-impact research publications in top-tier machine learning and applied mathematics conferences and journals. This WP is strongly interdisciplinary, bridging reinforcement learning, statistical learning, and agroecological modeling, while setting the stage for future data-driven agricultural decision support systems.

Deliverables (D) & Milestones (M)
D3a. Academic publications (M24, M36, M48)
M3a. Application of Transfer RL algorithms on agroecological data - assessment (M36)
D3b. PhD Manuscripts (M45)
M3b. Evaluation of Gym-STICS prediction compared with STICS-Intercropping on datasets (M42)
M3c. Test of contextual bandit strategies on intercropping data and diffusion to community (M42)

Work Package 4: Experimental Validation of RL-Based Agroecological Decision-Making [Resources: 1/2PhD@Cirad (M09-M45)]

Leader: M. Christina (CIRAD). **Participants:** A. Couëdel (CIRAD), K. Naudin (CIRAD), OA. Maillard (INRIA), N. Gaudio (INRAE.Agir)

Objective: Validate the reinforcement learning (RL) algorithms developed in WP3 by testing their ability to optimize decision-making in intercropping systems using crop models (WP1) and real-world experimental data (WP2). The key challenge is to assess whether RL can accurately simulate and improve agroecological management strategies while identifying missing or insufficiently parameterized plant-plant interaction mechanisms in existing crop models.

Challenges: While RL has shown promise in optimizing decision-making under uncertainty, its application to agroecology remains largely untested [8]. WP4 aims to bridge this gap by experimentally validating RL-based decision strategies for intercropped cereal-legume agrosystems. This will involve leveraging the Gym-STICS interface (developed in **WP1**) to couple RL with mechanistic crop models and integrating empirical data (structured in **WP2**) to refine model learning. A key question is whether RL can learn effective management policies without prior knowledge of interspecific interactions, revealing missing or oversimplified processes in current models. Similar work has been done with Gym-DSSAT [9]. What we propose to do here is more complex: first learning from single-plant simulations via STICS, then adapting to intercropped plants with a smaller real-world dataset. For the initial learning phase, the data is "unlimited" since it can be generated through model simulations. For learning the behavior of intercropped plants, the data comes from experiments gathered in **WP2**, and is therefore more limited. As a fallback, plausible intercropping data can be generated with STICS-intercrop for few crop associations. This approach is unique to date, as no other team in the world has published such work.

T4.1 Model-Based RL for Sole Crops. The first task will apply the RL algorithms developed in **WP3** to pure crop systems using Gym-STICS. This will serve as a baseline to test the ability of RL agents to learn crop responses under different pedo-climatic conditions and management practices. The focus will be on temperate (barley, wheat) and tropical cereals (maize, sorghum, millet), and legumes (chickpea, faba bean, lentil, cowpea, pigeon pea), already calibrated in pure crop in the STICS crop model. This task will validate the capacity of RL to optimize decision-making within well-understood monoculture systems before introducing the added complexity of intercropping.

T4.2 Data-Enhanced Learning for Intercropping. Building on T4.1, this task will incorporate experimental data from **WP2** to train RL agents on cereal-legume intercrops. Field trial datasets—including yield records, climate and soil data, management practices—will be used to refine learning, allowing the RL framework to infer plant-plant interactions such as competition, complementarity, and facilitation. The challenge will be to determine whether RL can learn these interactions implicitly from data, offering an alternative to explicitly modeled intercropping mechanisms in STICS.

T4.3 Comparison with STICS-Intercropping and independent datasets. The final task will evaluate RL-derived decision policies by comparing Gym-STICS outputs with both independent field trial data and an alternative crop model that explicitly represents intercropping, such as STICS-Intercropping [6]. This validation will leverage experimental station data from Mali, Burkina Faso, Senegal, and Brazil, where STICS-Intercropping has been tested [31]. Discrepancies between RL-generated policies and traditional modeling outputs will highlight potential gaps in the representation of intercropping dynamics, informing future improvements to mechanistic models.

Implementation and Collaboration A PhD student specializing in RL applications for agroecology will be recruited, working with crop modelers (INRAE and CIRAD), RL researchers (INRIA), and agronomists conducting experimental trials (CIRAD). By the end of **WP4**, we will have a validated framework demonstrating how RL can optimize intercropping decisions, providing insights into both the capabilities of RL in agroecology and the limitations of existing crop models in capturing interspecific interactions.

Deliverables (D) & Milestones (M)
M4a. Gym-STICS training on pure cultures of interest - quality assessment (M24)
M4b. Complementary learning on experimental data - quality assessment (M36)
D4a. Academic publication (M36)
M4c. Evaluation of Gym-STICS prediction compared to STICS-Intercropping on test datasets (M42)
D4b. PhD Manuscript (M45)

Work Package 5: Development and Management of a software platform and Agro Community
[Resources: 12m Engineer@Inria, 15k workshop]
Leaders: OA.Maillard (INRIA). **Participants**: N.Gaudio, C. Meunier (INRAE.Agir), S.Auzoux (CIRAD)

Objective: WP5 aims to develop an open-source software platform that integrates crop models, databases, and RL algorithms to support agroecological decision-making. This platform will unify the research outputs of WP1-WP4 into a scalable, well-documented software library, ensuring broad adoption. Another key goal is to position the platform as a reference tool within the AI research community, following successful precedents such as *HighwayEnv* [32], an RL environment developed at INRIA that was officially endorsed by the Farama Foundation. WP5 will adopt a similar strategy, ensuring compatibility with RL ecosystems like Gymnasium and leveraging INRIA's expertise in AI software dissemination.

Challenges: Integrating RL with crop models and real-world agricultural data poses significant challenges due to the fragmented nature of existing tools and the need for interdisciplinary collaboration. Most crop models were not designed for AI integration, and agricultural data can be heterogeneous and difficult to structure for RL applications. Additionally, the success of an AI-driven agroecological platform depends on its usability, adoption by researchers and practitioners, and its long-term maintenance. WP5 addresses these challenges through a dual focus: software development and community engagement.

T5.1 Development of the Software Platform. The goal is to build a robust, modular software platform that seamlessly integrates STICS and other crop models with RL algorithms, enabling scenario exploration and optimization of agroecological practices. The development process will be structured as follows: **T5.1.1) Core Infrastructure,** A unified system ensuring interoperability with WP1-WP4 outputs, facilitating the integration of crop models, RL environments, and databases; **T5.1.2) Algorithmic**

Framework- Implementation of RL techniques for agroecological decision-making, building on INRIA's expertise in AI and RL research; **T5.1.3) User-Friendly Interface-** Development of tools to make the platform accessible to agronomists and researchers, ensuring ease of use and interpretability; **T5.1.4) Simulation and Benchmarking-** Creation of synthetic and real-data environments for evaluating RL-driven agricultural strategies. Inspired by projects like ControlGym [33], the platform will provide a structured, interactive framework for training and testing decision-making models in agroecology; **T5.1.5) Long-Term Maintenance and Optimization-** Continuous updates to adapt to evolving research needs and ensure sustainability. A key ambition is to position the platform within the RL research community by integrating it into Gymnasium or similar repositories. This will maximize its exposure and encourage contributions from AI researchers working on RL for real-world applications.

T5.2 Community Engagement and Knowledge Dissemination. Beyond technical development, WP5 emphases on fostering an interdisciplinary research community around AI-driven agroecology. To achieve this, the following actions will be taken: **T5.2.1) Workshops and Training Sessions-** INRIA will host interdisciplinary workshops covering AI applications in agroecology, such as Large Language Models (LLMs) for knowledge extraction, RL for sustainable farming, generative models for climate adaptation, and foundation models for emerging agricultural practices; **T5.2.2) Forum Days-** Regular knowledge-sharing events will connect AI researchers, agronomists, and agricultural stakeholders, fostering collaboration and joint research initiatives; **T5.2.3) Integration into AI Research Frameworks** in order to promote the platform as a widely recognized community tool. INRIA will work toward aligning the agroecological platform with the requirements of RL research repositories; **T5.2.4) Open-Source Dissemination and Publications-**The platform will be released as an open-source project, with research papers and datasets shared at major AI and agroecology conferences to drive engagement.

Implementation and Collaboration WP5 will allocate key resources to ensure the successful implementation and long-term impact of the software platform: **Engineering Effort:** A full-time INRIA engineer will be dedicated for 12 months to lead development and software integration. **Financial Support:** 15k€ will be allocated to community-building efforts, including workshops, forum days, and interdisciplinary exchanges. **Collaborative Approach:** INRIA will coordinate with WP1-WP4 researchers to ensure seamless integration of crop models, RL algorithms, and field data. External collaborations with AI research groups and agricultural institutions will further strengthen the project. By the end of WP5, the project will deliver a fully functional, well-documented software ecosystem that empowers researchers and agronomists with advanced AI-driven tools for adaptive agroecological decision-making. Through targeted workshops, interdisciplinary collaborations, and open-source dissemination, WP5 will ensure lasting contributions to AI and agroecology research communities, while positioning the platform as a key reference for future developments.

Deliverables

D5.1 Opensource github platform, with documentation, reproducible experiments, and benchmarks (v1 M42, v2 M47);

D5.2 Opendatasets (M44)

D5.3 Workshops (M17, M28, M44).

Work Package 6: Project management

Leaders: OA. Maillard (INRIA), K. Naudin (CIRAD). **Participants**: N. Verzelen (INRAE.Mistea), N. Gaudio (INRAE.Agir)

Objective: This WP aims to ensure that the project proceeds according to the objectives set out in this response to the call for tenders, internal communication between project participants, organization of meetings, deliverables and relations with the ANR, and financial reporting. Consistency is an important issue, as the participants belong to different institutions and disciplines, on distant sites (Lille/Montpellier/Toulouse).

The main challenge lies in the project's strong interdisciplinarity, which is also its strength. It will require workshops to foster connections and mutual scientific understanding, essential for integrating expertise in AI, data science, and agroecology.

T6.1 Scientific and Technical coordination. This aims to ensure effective management of activities and smooth communication among partners. It includes progress monitoring and the organization of meetings and workshops to facilitate internal exchanges. Additionally, it ensures interactions with the ANR steering committee through the monitoring of indicators, the preparation of mid-term evaluation reports, and participation in conferences, ensuring compliance and project visibility.

T6.2 Financial and administrative management. This task is specifically dedicated to administrative and financial management. It includes the establishment of a consortium agreement, which defines the modalities for valorization and intellectual property management, the distribution of tasks, the allocation of resources, and the expected deliverables. It also specifies the rules governing publication, result dissemination, and overall project governance. In addition, this sub-work package ensures the financial management of the funding granted by the ANR, including the allocation of resources and reimbursements to partners. Finally, it oversees the preparation of final and intermediate reports, ensuring compliance with the periodicity and format requirements specified in the ANR funding contract.

Deliverables (D) & Milestones (M)
M6.1.a - Project coordination tools setup (M1) **M6.1.b** - Kick-off meeting organized
M6.1.c–f - General meetings organized
D6.2.a - Consortium agreement (M0) **M6.2.a** - Intermediate reports (M36)
D6.2.b - Final report (M48)

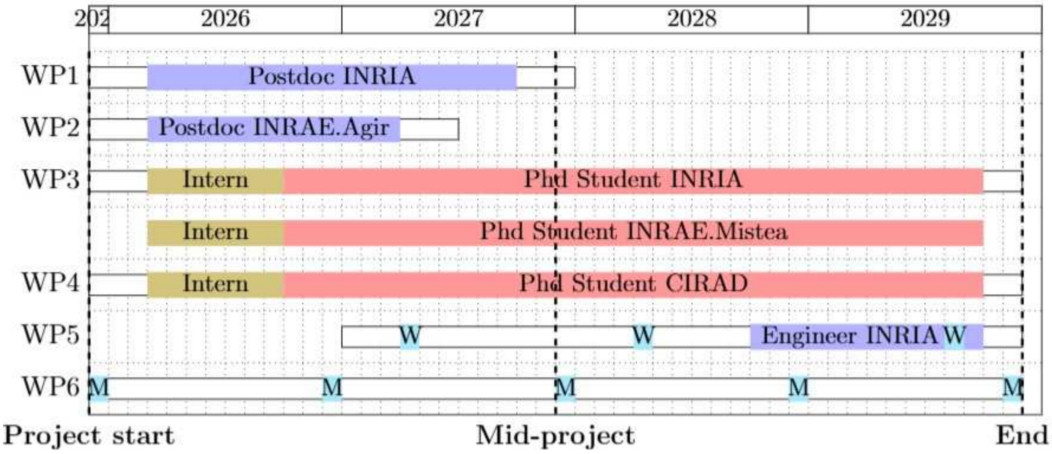

Risk description/Likelihood/Effect	**WP**	**Proposed risk-mitigation measures**
Cannot manage to develop Gym-STICS /Low /High	WP1	Using Gym-DSSAT, which already exists and works
Noisy, sparse, biased agronomic data hinders RL training /Medium/High	WP2	Automated data processing and simulation-based augmentation to enhance quality, prioritizing informative data throughout its lifecycle
Fail to develop hybrid transfer learning algorithm l/Medium /High	WP3	Direct coupling with the STICS intercropping version
Available data from real fields not sufficient RL training /Medium/High	WP4	Simulated data are produced using STICS plant alone and STICS intercropping
Fragmented RL and Agro communities hinder platform adoption /Medium/High	WP5	Foster open-source governance, strategic partnerships, and broad compatibility (Gymnasium), ensuring **long-term** cross-community engagement.
Can't manage Interdisciplinary research /Medium /High	WP6	Successful past INRIA-CIRAD collaboration (R. Gautron PhD, Gym-DSSAT)

3. Partnership (consortium or team)

Project Management and Consortium Structure. The AgriLORe project will be coordinated by **O. Maillard-Ambrym** (HDR), a leading researcher in artificial intelligence at **INRIA**, France's first research institute for AI. **D. Basu** is also a member of the Scool team of INRIA as the project co-ordinator, and is an expert in Reinforcement Learning, bandits, and Responsible AI (privacy, fairness, robustness). The project also brings together **INRAE** (French National Research Institute for Agriculture, Food, and the Environment) and **CIRAD** (French Agricultural Research Centre for International Development), institutions at the forefront of agricultural and environmental research. **INRAE** contributes through two key research units: **MISTEA** (Montpellier), represented by **N. Verzelen** (HDR), and **AGIR** (Toulouse), represented by **N. Gaudio,** ensuring expertise in modeling and statistical machine learning, as well as agroecological systems, especially intercropping. Additionally, **INRAE Agroclim** (Avignon), responsible for the dissemination and development of STICS soil-crop model, will provide indirect support. **CIRAD**'s participation, led by **K. Naudin** (HDR, agronomist, director of the AÏDA research unit)**,** will leverage its extensive experience in agroecological transitions and crop modeling, particularly in the Global South and French overseas territories. Other key CIRAD members include **S. Auzoux** (data modeling expert), **C. Mindingoyi** (computer engineer, PhD in agronomy), **M. Christina**, and **A. Couëdel** (PhDs in agronomy, experts in cropping system modeling). The diversity of contexts represented—from temperate to tropical cropping systems—ensures the robustness and generalizability of the proposed approaches.

Quality and Complementarity of the Consortium. The consortium integrates expertise from **reinforcement learning, crop modeling, data science, agronomy** and **agroecology,** ensuring a strong interdisciplinary foundation. INRIA brings state-of-the-art RL expertise, particularly in hybrid RL approaches and INRAE MISTEA brings expertise in data science and statistical machine learning while INRAE AGIR and CIRAD provide deep domain knowledge in agronomy, crop modeling, and agroecological systems. The combination of modeling, field experiment data, and algorithmic innovation strengthens the project's potential for impact. The collaboration between teams with prior experience in STICS, DSSAT, and RL-based decision-making ensures both scientific depth and technical feasibility.

Added Value of the Collaboration. While INRIA and CIRAD have already collaborated in PhD research demonstrating the feasibility of coupling RL with crop models, this project is the first to bring together INRIA, INRAE, and CIRAD in a single initiative. The INRAE MISTEA and INRAE AGIR teams have co-advised PhD on data analyses for intercropping experiments and INRAE AGIR and CIRAD have collaborated through initiatives such as the Horizon Europe "IntercropValuES" project, coordinated by CIRAD. This project will capitalize on existing synergies, leveraging past experiences to push the boundaries of hybrid RL for agroecology. Given the lack of prior work integrating RL and crop modeling at this scale, the project is positioned to make significant advances in both methodology and real-world applicability.

Consideration of Pluridisciplinarity. The project AgriLORe is inherently pluridisciplinary, combining expertise in agronomy, agroecology, reinforcement learning, statistical modeling, and computer science. This integration is crucial for **bridging the gap between AI-driven decision-making and agroecological complexity**. The collaboration between computer scientists and agronomists will facilitate the co-design of models and algorithms, ensuring that RL-based solutions remain grounded in agronomic realities while benefiting from the latest advancements in AI.

Assessment of Complementary or Alternative Funding Opportunities. The project's innovative and exploratory nature, positioned at the intersection of AI and agroecology, presents challenges for direct funding from traditional agricultural programs. Still, complementary funding avenues will be explored, including PEPR "Digital and Agroecology" programs, and applications for co-funding of two PhD students at CIRAD and at INRAE. Additional funding will be sought for workshops and community engagement initiatives at INRIA and INRAE, further amplifying the project's impact.

Human and Financial Resources. The project will rely on a strong team of researchers, engineers, and PhD students across INRIA, INRAE, and CIRAD, who will oversee scientific direction, technical development, and dissemination. A dedicated postdoctoral researcher will be responsible for the development of the Gym-STICS framework, leveraging prior work on Gym-DSSAT by INRIA and CIRAD. Another postdoctoral researcher will handle data pipeline construction, supervised by INRAE/CIRAD experts. Three PhD students will play a central role: two PhD students (INRIA & INRAE) will develop hybrid RL techniques for agroecological decision-making, and one PhD student (CIRAD), supervised by CIRAD agronomists and INRAE Agroclim experts, will assess the real-world applicability of RL-crop model coupling for agroecological systems. Finally, open-source dissemination of code and documentation will be handled by an INRIA-based engineer, ensuring the project's outputs contribute to the broader research community.

4. Exploitation and dissemination plan

Relevance for scientific and industrial applications. Project AgriLORe will establish the foundations for a new generation of soil-crop models, pioneering reinforcement learning (RL) integration to enhance agroecological decision-making. While soil-crop models are often described as digital twins, their capacity to represent complex biotic and abiotic interactions remains limited, particularly in diversified, low-input systems. By coupling RL with crop simulation, AgriLORe will enable model applications in previously inaccessible contexts, generating new insights into intercropping's role in climate adaptation. Beyond fundamental research, this approach lays the groundwork for practical innovations, such as a "virtual advisor" to assist farmers in testing crop associations. More broadly, the same framework could later extend to optimize organic fertilization strategies, complex crop rotations, and agroecological pest control. Strengthening France's leadership in AI-driven agroecology, AgriLORe will provide a robust methodological foundation for future European and international initiatives.

Genericity of the components of the proposed solution. AgriLORe's modular and interoperable design ensures its applicability beyond intercropping. Its methodological contributions—transferable RL strategies, open crop model coupling, and structured data pipelines—are designed for broader agroecological challenges. While initially focused on STICS, the project builds on prior efforts, including Gym-DSSAT, tackling two of the three leading global soil-crop models (STICS, DSSAT, APSIM). This dual approach will allow for generalizable recommendations, fostering future adaptations to other models. Moreover, this focus on genericity and inter-modularity extends beyond soil-crop models, as reflected in previous work by project participants on Crop2ML and generic modeling frameworks. The same concern for modular, reusable solutions also applies to agroecological data structuring and RL algorithms, ensuring broad applicability across multiple agricultural contexts. AgriLORe thus reinforces France's role in AI-driven agronomy while laying a solid foundation for future adaptive, scalable decision-support tools.

Dissemination/exploitation/commercialization plan of relevant foreseen project results during. The dissemination and exploitation of AgriLORe's results will be structured through targeted scientific events, open-access tools, and institutional collaborations, ensuring both academic and applied impact. A dedicated end-of-project workshop will bring together reinforcement learning researchers and agronomist modelers to showcase the data pipeline, RL algorithm selection, and Gym-STICS environment for intercropping design. Project members will actively present results at key scientific conferences such as the STICS user workshop and ICROPM, which gathers the global soil-crop modeling community, as well as major AI conferences. All code and methodologies will be made openly available through INRIA's permanent GitLab repository, and an engineer will be hired for twelve months to focus on dissemination. While commercialization is not planned, given the open-source nature of the STICS model (CeCILL-C license), the project's outputs will be positioned for integration into agricultural decision-support systems, benefiting researchers, policymakers, and agricultural stakeholders committed to the agroecological transition.

Dissemination including any scientific publication, standardisation, benchmarking and evaluation activities open to research teams beyond the project consortium. Beyond the project consortium, AgriLORe will engage with leading French and European research teams in agroecology, ensuring wide dissemination. The involvement of CIRAD (Aïda) and INRAE (AGIR) in Horizon Europe and PEPR "Cultiver Autrement" through the MoBiDiv project will facilitate benchmarking and knowledge exchange with top European teams. Collaborations with French technical institutes (Terres Inovia, ITAB) will further ensure that AgriLORe's advances reach both researchers and practitioners.

While AgriLORe focuses on intercropping systems, it serves as a showcase of how the coupling of reinforcement learning, crop simulation, transfer learning techniques, and data-driven experimentation can tackle agroecological challenges. This integrated approach positions the project as a key contributor to national and European agroecology initiatives, reinforcing France's strategic influence in the field.

Communication measures for promoting the project and its findings during the project duration: tutorials, courses, workshops. A central element of the dissemination strategy is the organization of workshops for the agroecology research community and stakeholders. These workshops will serve three main objectives: (i) introduce AI-driven approaches to agroecological design, (ii) showcase the benefits of RL-based intercropping strategies, and (iii) build a cross-disciplinary research community at the intersection of ML, data science, and agronomy. Given the emphasis on AI-driven agroecology in national research programs (INRAE 2030, PEPR "Agroécologie et Numérique," and Digit-Bio), the project will leverage additional funding from these initiatives. In the medium term, AgriLORe's impact will extend through Horizon Europe's Farm-to-Fork program and CASDAR programs, ensuring long-term scientific and stakeholder engagement. By addressing both large-scale agricultural institutes and smaller-scale farmer networks through participatory research, AgriLORe will create a broad and lasting impact.

Scientific and economic potential. From both a scientific and economic perspective, AgriLORe offers high transformative potential. The ability to couple RL with crop simulation opens new avenues for adaptive, AI-driven agroecological strategies, supporting climate resilience. The project's innovations can be extended beyond intercropping to other key agroecological practices, including organic fertilization, integrated pest management, and soil health restoration. By pioneering open-source, AI-powered decision-support systems, AgriLORe will contribute to the digitalization of agronomy while maintaining a strong emphasis on scientific rigor, transparency, and reproducibility.

Dissemination and exploitation of scientific, technical and industrial culture, contribution to the French scientific community. By bringing together INRIA, INRAE, and CIRAD, AgriLORe represents a landmark initiative for the French scientific community. It fosters collaboration between AI researchers and agronomists, reinforcing the role of France as a leader in digital agronomy. The project will contribute to scientific, technical, and industrial knowledge transfer by publishing in top-tier journals, engaging in standardization and benchmarking efforts, and ensuring that the tools developed remain accessible and extendable. Ultimately, AgriLORe embodies France's strategic vision for AI in agroecology, creating a foundation for future interdisciplinary research and innovation.

5. Financial tables: Summary of costs & efforts

Missions: we budget 1k€/y/permanent, 6k€/y/Postdoc, 3k€/y/Engineer, 8k€/PhD for each partner, hence:

- INRIA: 2x4+9+8+3=28k€
- CIRAD: 5x4+8=28k€
- INRAE MISTEA: 3x4+8=20k€
- INRAE AGIR: 2x4+6=14k€

Workshops:

- INRIA: 10k€
- INRAE AGIR: 5k€

Appel à projets Thématique Spécifique en Intelligence Artificielle (TSIA) – Edition 2025

Name and organisation	Salaries –Staff (euro)						Instruments and equipements (euro)	Services and Intellectual property rights (euro)	Other costs (euro)	**Total (euro)**
	Permanent position		Non permanent position with funding requested		Non permanent position without funding requested					
	Cost (euro)	Person month	Cost (euro)	Person month	Cost (euro)	Person month				
INRIA SCOOL	316 800	38,4	265 200	72	0	0	4 500	15 000	38 000	**322 700**
INRAE MISTEA	157 790	16	70 850	24	67 190	18	1 500	0	20 000	**92 350**
CIRAD AïDA	194 660	20	57 440	24	53 460	18	1 500	0	28 000	**86 940**
INRAE AGIR	161 741	21,6	74 438	12	0	0	1 500	0	19 000	**94 938**

6. References

**** Publications in bold are co-authored by the project partners***

[1] Beillouin et al. 2021. Positive but variable effects of crop diversification on biodiversity and ecosystem services. Glob Change Biol. https://doi.org/10.1111/gcb.15747
[2] Brisson et al. 2003. An overview of the crop model STICS. Eur J Agron. https://doi.org/10.1016/S1161-0301(02)00110-7
[3] Jones et al. 2003. The DSSAT cropping system model. Eur J Agron. https://doi.org/10.1016/S1161-0301(02)00107-7
[4] Muller, Martre, 2019. Plant and crop simulation models: powerful tools to link physiology, genetics, and phenomics. J Exp Bot. https://doi.org/10.1093/jxb/erz175
[5] Gaudio et al. 2019. Current knowledge and future research opportunities for modeling annual crop mixtures. A review. Agron Sustain Dev. https://doi.org/10.1007/s13593-019-0562-6
[6] Vezy et al. Modeling soil-plant functioning of intercrops using comprehensive and generic formalisms implemented in the STICS model. Agron Sustain Dev. https://doi.org/10.1007/s13593-023-00917-5
[7] Githui et al. 2023. Resource sharing in intercropping models and a case study with APSIM in southern Australia. Eur J Agron. https://doi.org/10.1016/j.eja.2022.126680
[8] Gautron et al. 2024. A new adaptive identification strategy of best crop management with farmers. Field Crops Res. https://doi.org/10.1016/j.fcr.2024.109249
[9] Gautron et al. 2022. gym-DSSAT: a crop model turned into a Reinforcement Learning environment. https://doi.org/10.48550/arXiv.2207.03270
[10] Bilodeau et al. 2024. Development of AI-assisted microscopy frameworks through realistic simulation with pySTED. Nat Mach Intell. https://doi.org/10.1038/s42256-024-00903-w
[11] Tan. 2024. Precision agriculture for corn using reinforcement learning. PhD Thesis, UTAR.

[12] Turchetta et al. 2022. Learning long-term crop management strategies with cyclesgym. Adv Neural Inf Process Syst. https://github.com/kora-labs/cyclesgym
[13] Durand et al. 2018. Streaming kernel regression with provably adaptive mean, variance, and regularization. Journal of Mach Learn Res. https://doi.org/10.48550/arXiv.1708.00768
[14] Durand et al. 2018. A machine learning approach for online automated optimization of super-resolution optical microscopy. Nat Commun. https://doi.org/10.1038/s41467-018-07668-y
[15] Yang et al. 2024. Best Arm Identification with Minimal Regret. https://doi.org/10.48550/arXiv.2409.18909
[16] Parast, Bartroff 2024. Group sequential testing of a treatment effect using a surrogate marker. Biometrics. https://doi.org/10.1093/biomtc/ujae108
[17] Elsayed et al. 2024. Streaming Deep Reinforcement Learning Finally Works. https://doi.org/10.48550/.2410.14606
[18] Garcia 1999. Use of reinforcement learning and simulation to optimize wheat crop technical management. Proceedings of the International Congress on Modelling and Simulation (MODSIM'99) Hamilton, New-Zealand. pp. 801–806.
[19] Pradal et al. 2008. OpenAlea: a visual programming and component-based software platform for plant modelling. Funct Plant Biol. https://doi.org/10.1071/FP08084
[20] Chen et al. 2021. A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. Agric Water Manag. https://doi.org/10.1016/j.agwat.2021.106838
[21] Sun et al. 2017. Reinforcement learning control for water-efficient agricultural irrigation. IEEE International symposium on parallel and distributed processing with applications and 2017 IEEE international conference on ubiquitous computing and communications (ISPA/IUCC). IEEE; pp. 1334–1341.
[22] Wang et al. 2020. Deep reinforcement learning for greenhouse climate control. 2020 IEEE International Conference on Knowledge Graph (ICKG). IEEE; pp. 474–480.
[23] Song et al. 2022. Hybrid rl: Using both offline and online data can make rl efficient. ArXiv Prepr ArXiv221006718. https://doi.org/10.48550/arXiv.2210.06718
[24] Tesauro et al. 2006. A hybrid reinforcement learning approach to autonomic resource allocation. 2006 IEEE International Conference on Autonomic Computing. IEEE; pp. 65–73.
[25] Da Silva, Costa 2019. A survey on transfer learning for multiagent reinforcement learning systems. J Artif Intell Res. https://doi.org/10.1613/jair.1.11396
[26] Khetarpal et al. 2022. Towards Continual Reinforcement Learning : A Review and Perspectives. J Artif Intell Res.
[27] Mahmoud et al. 2024. A workflow for processing global datasets: application to intercropping. PCI Comp & Math Biol. https://doi.org/10.24072/pcjournal.389
[28] Gaudio et al. 2024. A global dataset gathering 37 field experiments involving cereal-legume intercrops and their corresponding sole crops. Zenodo; https://zenodo.org/records/14160885
[29] Paut et al. 2024. A global dataset of experimental intercropping and agroforestry studies in horticulture. Sci Data. https://doi.org/10.1038/s41597-023-02831-7
[30] White et al. 2013. Integrated description of agricultural field experiments and production: The ICASA Version 2.0 data standards. Comput Electron Agric. https://doi.org/10.1016/j.compag.2013.04.003
[31] Couedel et al. 2023. Modeling agroecological intensification in the tropics with the Stics model-lessons learned and way forward. INRAE.
[32] Leurent 2018. An Environment for Autonomous Driving Decision-Making. GitHub repository. https://github.com/eleurent/highway-env
[33] Zhang et al. 2024. Controlgym: Large-scale control environments for benchmarking reinforcement learning algorithms. 6th Annual Learning for Dynamics & Control Conference. PMLR; pp. 181–196.