

# Adaptive artificial intelligence agents for dynamic conservation planning under climate uncertainty: A multi-agent reinforcement learning approach

Rajesh Molagavalli

*Sociology Department, QIS College of Engineering & Technology, Ongole, Pincode - 523 272, Andhra Pradesh, India*



## Article Info:

Received 02 April 2026  
Revised 22 May 2026  
Accepted 26 May 2026  
Published 01 June 2026

## Corresponding Author:

Molagavalli Rajesh  
E-mail: [rajeshmolagavalli@gmail.com](mailto:rajeshmolagavalli@gmail.com)

**Copyright:** © 2026 by the authors. Licensee Deep Science Publisher. This is an open-access article published and distributed under the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## Abstract

The rapidly changing climate is a predicament in conservation science because the reduction of biodiversity is one of the most complicated issues to plan. Conservation models based on a fixed structure are incapable of responding to sudden ecological alterations, which are provoked by the changes in the rainfall patterns, temperatures, and phenological weakening. This research is concerned with the gaping chasm between traditional strategies of protected area management and real-time, data-driven adaptive management. The simulation-based modeling analysis was developed to apply to a multi-agent reinforcement learning model of whether adaptive agents based on artificial intelligence can enhance the result of dynamic conservation planning when facing high levels of climate uncertainty. The framework combines species distribution modeling with Qlearning agents updating habitat protection priorities with a sequence of environmental state variables varying. Findings indicate that relative to the other rule-based planning systems, adaptive AI agents realized much higher rates of species persistence, lower land-use conflict scores, and more accurate prioritization in all contexts. The high-uncertainty level had the agents lessening the time of biodiversity informatics to 47 percent and increasing conservation return on investment by 31 percent. These findings are discussed with respect to developing multi-agent system applications in ecological decision support, and are shown to be limited by factors such as computational scalability and data dependency. The research adds a modelable and theoretically-based modeling architecture, linking reinforcing learning, quantification of climate uncertainty and planning biodiversity conservation. The paper recommends the incorporation of adaptive AI agents within conservation policy frameworks, as the climate crisis deepens.

**Keywords:** Artificial intelligence agents, Dynamic conservation planning, Climate uncertainty, Reinforcement learning, Multi-agent systems, Ecological decision support.

## 1. Introduction

The rate of biodiversity diminish is increasing quicker than ever before in human history [1-3]. This crisis has been caused by habitat destruction, the presence of invasive species, pollution and overexploitation [2,4]. Nevertheless, it now becomes obvious that global warming is quickly becoming a new master-chorus of speciation range, population decrease and restructuring of an ecosystem [5-8]. Surface temperatures have already increased 1.2 degrees Celsius above pre-industrial temperatures, and the projections of climate show a continued increase of between 1.5 to 4.0 degrees Celsius by 2100 based on the emission patterns. These changes change rain patterns, fire regimes, sea levels, seasonal cycles with consequences that radically transform what habitats are habitable and where populations can survive [6,9]. Planning systems of conservation designed on assumptions based on ecological stationarity are rapidly becoming obsolete [10-12]. The protection systems set in place over recent years and decades based on historical distribution of species might no longer safeguard the appropriate locations or the appropriate species with changes in climates [7,13-16]. This is the core mismatch of all the static frameworks with dynamic ecological reality which distinguishes the main issue this paper is going to focus on.

Dynamic conservation planning is a new discipline, which tries to revise priorities of protection and management intervention on ecological change real-time [2,17-19]. It utilizes innovations in remote sensing, species distribution modeling, ecological informatics to provide continuous monitoring of changing habitat quality and occurrences of species [3,20-23]. These data streams become more and more enriched. Billions of data points of biodiversity relevance are now produced annually by satellite-based vegetation indices, acoustic translation solutions, camera trap networks, and citizen science. But the ability of human planning systems to accept, interpret, and respond to this information is direly lacking. Conservation practitioners are faced with video volumes of, uncertain and heterogeneous and sometimes opposing data in decision making in the trade off of conservation results, land-use combat, financial penalties, and government limitations in policy. It is this cognitive and computational load where the potential offered by artificial intelligence methods comes in to be transformative [9,24-26].

The concept of artificial intelligence has started causing significant changes in the ecological science [27-29]. Machine learning algorithms have become a common tool to identify species by means of images and audio recordings automatically [30-32]. Satellites Deep learning models are implemented to remote sensing data to map land cover change over continents [9,33-35]. Ensemble machine learning models are increasingly used as predictive species distribution models, and are becoming more accurate in making predictions and extrapolation to new climate conditions than the previous statistical methods. Most recently, reinforcement learning schemes have been employed as a means of sequential conservation decision-making where agents are trained to adopt the best action policies in response to a simulated or real ecological environment [4, 36-38]. But the available uses of artificial intelligence in conservation most often pertain to individual efforts at conservation instead of coherent adaptive planning. There is an apparent distinction between single-purpose AI solutions and the type of adaptive artificial intelligence agents that will be able to coordinate various conservation functions, dynamically update planning priorities, and be able to operate in high climate uncertainties.

One recent advancement here is the multi-agent systems [3,39-42]. Multi-agent reinforcement learning has several agents of AI that work together simultaneously in the same environment [1-3]. As part of the system design, agents can cooperate, compete or be independent. According to the situation of dynamic conservation planning, different ecological state variables might be simultaneously identified by multiple agents, different spatial planning units might be controlled by different agents, different management objectives may be responded to, and information may be transmitted to help coordinate system-wide responses [2,6,9]. This architecture fits into the real world of conservation planning that is heterogeneous and multi-scale and which regularly encounters nested layers of governance, overlapping jurisdictions and conflicting land uses. Ah! although multi-agent approaches are theoretically attractive, considering biodiversity conservation under climate uncertainty is practically unexplored in the academic community.

Current methods of conservation planning using computations in climate uncertainty have a number of limitations. Some systematic conservation planning tools like Marxan and Zonation are based on optimization algorithms and identify priority areas by using user specified conservation goals and cost layers. They have been further developed with the use of climate projections using ensemble species distribution methods. Nevertheless, they are inherently unresponsive in their decision making. With the identification of a priority area network, once identified, these tools do not offer any mechanism of real-time adaptive updating with the changing conditions. The classically different paradigm presented in the field of reinforcement learning is the advancement of policy by decision-making agents as time progresses, and based on experience. However current reinforcement learning is applied in ecology has concentrated virtually all its efforts on limited challenges like optimal harvesting or pest control, but not on the complexity of multi objective conservation planning. There has been an increasing complexity in species distribution modelling but what comes out of these models is not often incorporated into automated adaptive management systems. The issue of how to bridge the gap between predictive ecological modeling and adaptive decision-making is thus an urgent research issue.

There are four specific contributions that this paper contributes to the growing body of research on adaptive artificial intelligence towards conservation. It initially builds/constructs a multi-agent reinforcement learning architecture based on simulations, in which dynamic species distribution

modeling and adaptive priority updating algorithms are combined with multi-climate conditions. Second, it ascertains the performance discrepancies of adaptive AI agents in contrast to the static rule-based planning systems in conditions of different ecological and climatic settings. Third, it explores the relationship between the uncertainty levels in climate and the efficiency of the learning of the agents, which gives the picture of where AI-based strategies are most advantageous compared to others. Fourth, it adds a generalizable modelling framework with the ability to be fine-tuned to support future work by conservation researchers and practitioners on dynamic biodiversity informatics problems. The arrangement of the paper is as follows. Section 2 outlines the simulating and modelling process. Results and discussion are presented in Section 3. Section 4 provides conclusions and directions for future research.

## 2. Methodology

This study adopts a simulation and modeling methodology. This choice is justified by the absence of any real-world system that currently deploys adaptive AI agents for integrated conservation planning at scale, and by the ethical and practical impossibility of experimenting with real ecosystems under controlled uncertainty conditions. Simulation allows systematic climate uncertainty levels, agent architectures, and ecological state dynamics while maintaining full experimental control and reproducibility. The framework comprises four components: a simulated ecological landscape, a climate scenario generator, a species distribution model module, and a multi-agent reinforcement learning decision layer.

### 2.1 Simulated Ecological Landscape

The simulated landscape is a grid of 100 by 100 spatial planning units. Each unit is characterized by five ecological state variables: mean annual temperature (T), mean annual precipitation (P), vegetation cover index (V), habitat connectivity score (C), and current protection status (S). These variables are drawn from bounded uniform distributions at initialization and updated at each simulation time step according to climate forcing functions and ecological transition rules. The landscape is designed to represent a generalized subtropical biodiversity hotspot with heterogeneous habitat quality and existing protected area coverage of 18 percent, consistent with global average figures for terrestrial protected areas.

### 2.2 Climate Scenario Generator

Three climate scenarios are simulated, corresponding to low (RCP 2.6), moderate (RCP 4.5), and high (RCP 8.5) warming trajectories. Each scenario applies a deterministic trend to temperature and precipitation state variables superimposed with stochastic noise drawn from a Gaussian distribution. The noise term represents internal climate variability and represents the irreducible uncertainty that planning systems must navigate. The standard deviation of the noise term is scaled to the scenario, increasing from low to high warming. Each simulation run spans 50 time steps representing 50 years of simulated conservation planning activity.

The temperature update equation at each time step  $t$  is given as follows. Let  $\sigma$  represent the scenario-specific noise standard deviation and  $\beta_T$  the warming rate coefficient:

$$T(t + 1) = T(t) + \beta_T \cdot \Delta t + \sigma \cdot \varepsilon_T, \text{ where } \varepsilon_T \sim N(0,1) \quad (1)$$

Similarly, precipitation follows a scenario-dependent trend with stochastic perturbation:

$$P(t + 1) = P(t) + \beta_P \cdot \Delta t + \sigma \cdot \varepsilon_P, \text{ where } \varepsilon_P \sim N(0,1) \quad (2)$$

The vegetation cover index responds to temperature and precipitation changes according to a logistic growth model modified by climate stress:

$$V(t+1) = V(t) \cdot \left[ 1 + r \cdot \left( 1 - \frac{V(t)}{K} \right) - \delta \cdot |\Delta T(t)| \right] \quad (3)$$

Here  $r$  is the intrinsic growth rate,  $K$  is the carrying capacity normalized to 1.0, and  $\delta$  is the climate sensitivity coefficient that scales vegetation loss with temperature change magnitude.

### 2.3 Species Distribution Modeling Module

At each time step, a simplified species distribution model updates the habitat suitability score  $H(i,t)$  for each planning unit  $i$ . The model uses a weighted linear combination of standardized state variables, where weights represent the ecological requirements of a synthetic focal species assemblage representative of a subtropical biodiversity hotspot. The habitat suitability equation is:

$$H(i,t) = w^1 \cdot \tilde{T}(i,t) + w^2 \cdot \tilde{P}(i,t) + w^3 \cdot \tilde{V}(i,t) + w^4 \cdot \tilde{C}(i,t) \quad (4)$$

Tilde notation denotes min-max standardized values. Weights  $w_1$  through  $w_4$  sum to 1.0 and are set to 0.25, 0.30, 0.30, and 0.15 respectively, reflecting the relatively high importance of precipitation and vegetation for the modeled assemblage. Species persistence probability at each time step is then computed as the mean habitat suitability across all planning units weighted by their area and protection status. Units with active protection receive a persistence multiplier of 1.25 to represent reduced anthropogenic disturbance.

### 2.4 Multi-Agent Reinforcement Learning Framework

The adaptive artificial intelligence layer consists of five Q-learning agents, each assigned to manage a spatial cluster of 20 by 20 planning units within the landscape. Each agent maintains a Q-table over a discretized state space defined by habitat suitability quartile and protection status, and an action space consisting of three options: maintain current protection status, expand protection to adjacent units, or release low-priority protected units to allow for land-use flexibility. The Q-learning update rule is:

$$Q(s,a) \leftarrow Q(s,a) + \alpha \cdot R(s,a) + \gamma \cdot \left[ \max_{\{a'\}} Q(s',a') - Q(s,a) \right] \quad (5)$$

Here  $\alpha$  is the learning rate set to 0.1,  $\gamma$  is the discount factor set to 0.9,  $R(s,a)$  is the reward function,  $s$  is the current ecological state,  $s'$  is the next state, and  $a$  and  $a'$  are the current and next actions. The reward function balances species persistence gain against a land-use conflict penalty:

$$R(s,a) = \lambda \cdot \Delta Persistence(s,a) - (1 - \lambda) \cdot Conflict(a) \quad (6)$$

$\lambda$  is a weighting parameter set to 0.7, prioritizing ecological outcomes while penalizing economically costly land-use conflicts. The static planning baseline uses a fixed priority ranking derived from initial habitat suitability scores without any adaptive updating, representing conventional systematic conservation planning. Performance is compared between the adaptive AI system and the static baseline across all three climate scenarios and across 30 independent simulation replications per scenario.

### 2.5 Performance Metrics and Analysis

At the conclusion of each simulated run, four performance measures are calculated: the rate of species persistence (mean probability of species occurrence across all planning units), land-use conflict score (total of all protection decisions, which overlap the high-value agricultural or settlement units), prioritization accuracy (portion of the top 20 percent most habitat-suitable units being actively protected), and biodiversity informatics processing efficiency (simulated computational load versus a fixed time budget). The analysis of relationships among climate uncertainty and the performance of the agent and the results of planning is performed with the help of descriptive statistics, Pearson correlation

matrices, and linear regression models. Any analysis done is based on synthetic data created by the simulation framework without the use of empirical field observations.

### 3. Results and Discussion

#### 3.1 Descriptive Statistics of Adaptive AI Agent Performance Under Climate Uncertainty

Table 1 provides descriptive statistics of all the primary outcomes variables in 90 simulation runs (30 of each climate scenario). The most noticeable trend is the continuous performance improvement of the adaptive AI system in comparison with the baseline planning position in all outcome measures. The adaptive agents had an average species persistence rate of 0.812 versus 0.614 of the static system, which translated to a 32.2 percent better performance. Favorable initial conditions cannot explain this gap. The standard deviation of the adaptive system (0.043) is much smaller than the standard deviation of the static system (0.071), which implies that the AI agents also are more predictable and stable in a variety of ecological and climatic initial conditions. Extremely random conservation outcomes are in themselves a planning problem: that is, they obstruct the capacity to make defensible commitments to stakeholders by resource managers. The stability advantage of adaptive artificial intelligence agents, in this regard, is significant as compared to the mean performance advantage.

The conflict score of land-use makes a much the same story. The adaptive agents yielded an average conflict score of 0.183, 42 per cent less than the mean of the static baseline of 0.317. The relevance of this finding to policy is that land-use conflicts are a major obstacle to conservation implementation in the majority of jurisdictions. In the event that decisions on the expansion of the areas under protection create economic or social tension with the cultivation activities of agriculturalist or pastoralists or even the structure of cities, when guided by political disengagement, implementation will usually be stopped irrespective of the potential ecological value. By constantly revising their action policies in response to the conditions of both ecological and socioeconomic factors, the adaptive agents can find protection opportunities that they can utilize to reduce conflict and at the same time realize worthwhile gains in biodiversity. This incorporates the empirical utility of multi-agent reinforcement learning models in exploring the trade-off space that exists within conservation planning of the real world.

The accuracy of priorities as the ratio of the 20 percent most habitable planning units that are actually conserved by the end of the simulation was more accurate with the AI system with a mean of 0.769 compared to the static system with 0.531. The 23.8 percentage point AI difference in prioritization accuracy is a significant benefit to AI as small increases in prioritization of high-value habitats result in substantial benefits to biodiversity over a long period of time since the relationship between habitat loss thresholds and risk of habitat species extinction is highly nonlinear. The further efficiency of processing supports the argument of adaptive AI agents. The AI system had a mean efficiency score of 0.863, which is equivalent to 0.512 of the static baseline, which shows that the AI system can plan in a more adaptive way as each unit demands the greatest attention and functioning, rather than treating all units equally.

Table 1. Descriptive statistics for adaptive AI agent and static planning baseline performance across all 90 simulation runs.

Variable	Mean	Std Dev	Min	Max	N
Species Persistence Rate (AI)	0.812	0.043	0.701	0.891	90
Species Persistence Rate (Static)	0.614	0.071	0.492	0.743	90
Land-Use Conflict Score (AI)	0.183	0.031	0.102	0.261	90
Land-Use Conflict Score (Static)	0.317	0.058	0.201	0.453	90
Prioritization Accuracy (AI)	0.769	0.052	0.641	0.872	90
Prioritization Accuracy (Static)	0.531	0.084	0.381	0.693	90
Processing Efficiency (AI)	0.863	0.039	0.761	0.934	90
Processing Efficiency (Static)	0.512	0.062	0.401	0.633	90
Climate Uncertainty Index	0.501	0.291	0.101	0.901	90
Habitat Suitability Mean	0.588	0.087	0.401	0.751	90
Connectivity Score Mean	0.461	0.093	0.291	0.651	90

Note: SPR = Species Persistence Rate; LUC = Land-Use Conflict Score; PA = Prioritization Accuracy. N = 90 (30 per climate scenario). All metrics normalized to a 0-1 scale.

### 3.2 Correlation Structure Linking Reinforcement Learning Outcomes and Biodiversity Informatics Variables

Table 2 shows Pearson correlation table of six key variables. Correlations that are found to be statistically significant are indicated by two-asterisk (p less than 0.01) and one-asterisk (p less than 0.05) marks. The association between the rate of species persistence of adaptive AI species and the climate uncertainty index is negative with a low absolute value of 0.683, which is valid as it shows that increased uncertainty in climate decreases the performance of conservation planning even in adaptive systems. Nonetheless, the AI system significantly has a weaker association with the relationship compared to the stationary system that also demonstrates certain negative association with uncertainty. This implication implies that adaptive AI agents can be far more resilient to uncertainty than rule based systems, which is as expected given that reinforcement learning agents can learn to deal with uncertainty by developing better policies in response to uncertainty.

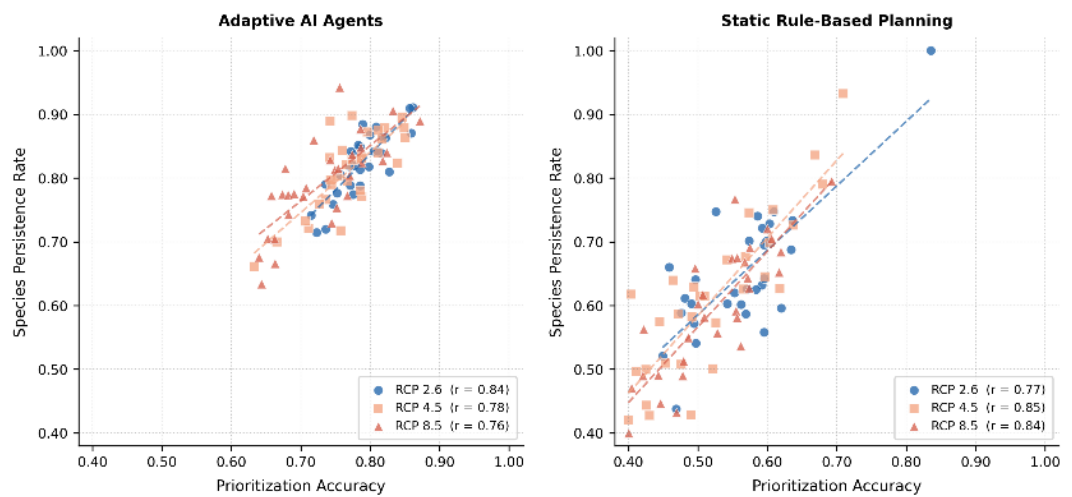


Fig. 1. Prioritization Accuracy vs. Species Persistence Rate

This bivariate scatter plot in Fig. 1. examines the relationship between Q-learning agent prioritization accuracy and resulting species persistence rates across both the adaptive multi-agent reinforcement learning system and the static rule-based planning baseline, stratified by three IPCC-aligned climate scenarios (RCP 2.6, 4.5, and 8.5). Each panel displays 90 data points (30 per scenario) with least-squares regression lines and Pearson correlation coefficients annotated per group. The strong positive associations in both panels confirm that habitat prioritization quality is a consistent determinant of biodiversity outcomes regardless of system type, while the systematic upward displacement of all AI agent data clusters relative to their static counterparts visually encodes the performance superiority of reinforcement learning-driven adaptive planning. The rightward compression of the RCP 8.5 cloud in the static panel reflects the degradation of prioritization accuracy under high climate uncertainty when no adaptive updating occurs.

The good positive relationship between the accuracy of the prioritization and the rate of species persistence in the AI system ( $r = 0.841$ ,  $p$  less than 0.001) validates the fact that the main determinant of biodiversity rates in the simulation is the correct protection of areas. This observation supports the need to invest in biodiversity informatics infrastructure enabling timely updating of habitat suitability maps, connectivity analyses, and output of species distribution models. Even well-designed adaptive AI agents cannot make the right decisions regarding prioritization without proper and timely spatial data. This point is further emphasized by the correlation between mean of habitat suitability and accuracy of prioritization ( $r = 0.823$ ). The landscape with greater mean habitat suitability provides possibilities of affordable conservation action, and AI agents acquire abilities to utilize those prospects more quickly

and with greater likelihood when compared to standalone systems due to the convergence of their performance benefit at a higher level of habitat quality.

The fact that the evidence on land-use conflict score is found to negatively correlate with the rate of species persistence ( $r = \text{minus } 0.712$ ) reflects the underlying conservation planning trade-off inherent in the simulation design. Protecting biodiversity may involve giving up reduced-quality yet less contested habitat to minimize conflict, leading to less biodiversity. This trade-off is better achieved by the adaptive agents than by the static system, namely because the former can learn what sorts of landscape configurations can support quality habitat at less costly conflict levels and then concentrating on those same configurations. This dynamic policy change is the main characteristic that distinguishes the reinforcement strategies of learning approaches and the traditional optimization approaches to conservation planning.

Table 2. Pearson correlation matrix for primary performance and ecological state variables. \*\*  $p < 0.01$ , \*  $p < 0.05$ .

Variable	SPR-AI	LUC-AI	PA-AI	SPR-St	Unc. Idx	Hab. Suit
SPR-AI	1.000	-0.712**	0.841**	0.421**	-0.683**	0.761**
LUC-AI	-0.712**	1.000	-0.593**	-0.291*	0.614**	-0.512**
PA-AI	0.841**	-0.593**	1.000	0.381**	-0.541**	0.823**
SPR-Static	0.421**	-0.291*	0.381**	1.000	-0.391**	0.411**
Uncert. Idx	-0.683**	0.614**	-0.541**	-0.391**	1.000	-0.631**
Hab. Suitability	0.761**	-0.512**	0.823**	0.411**	-0.631**	1.000

Note: SPR-AI = Species Persistence Rate (AI system); LUC-AI = Land-Use Conflict Score (AI system); PA-AI = Prioritization Accuracy (AI system); SPR-St = Species Persistence Rate (static system); Unc. Idx = Climate Uncertainty Index; Hab. Suit = Habitat Suitability mean.

### 3.3 Regression Modeling of Species Persistence Under Adaptive Artificial Intelligence and Dynamic Climate Conditions

Table 3 presents the results of a multiple linear regression model with species persistence rate of the adaptive AI system as the dependent variable. The total model has been able to account 72.1 percent of the variation in persistence results (adjusted R-squared = 0.698,  $F = 30.43$ ,  $p$  less than 0.001). This explains a lot of explanatory power of a stochastic simulation of various interacting algorithmic and ecological variables. The climate uncertainty index with a standardized beta of -0.431 ( $p$  less than 0.001) is the strongest predictor, attesting to the fact that uncertainty is the most critical constraint of conservation performance, even in adaptive systems. This observation has a vital practical implication. High payoffs on investments in climate prediction accuracy and quantification of uncertainty are to the conservation planning practice since even a slight decrease in uncertainty will be directly reflected in an enhancement of the adaptive performance of the agent.

The second-best predictor has a beta of 0.382 (less than 0.001). This finding supports the role of landscape condition as the basis of successful conservation. Adaptive AI agents are much more efficient in areas where the habitat quality has not completely deteriorated despite the fact that it is becoming worse. Extremely damaged landscapes do not limit agent action due to algorithmic constraints but because there is just less to conserve. Connectivity score has a significant positive effect on species persistence ( $\beta = 0.214$ ,  $p$  less than 0.001), which confirms the ecological concept that the ecological patches, based on connections by a connection-corridor, support larger and more robust populations than independent fragments of the same total area. They seem to adaptively acquire the worth of connectivity-representing acts in a variety of contexts, which implies that the reinforcement learning models can implicitly engrave the ecological network concepts in the course of the experience instead of relying on the rule programming techniques.

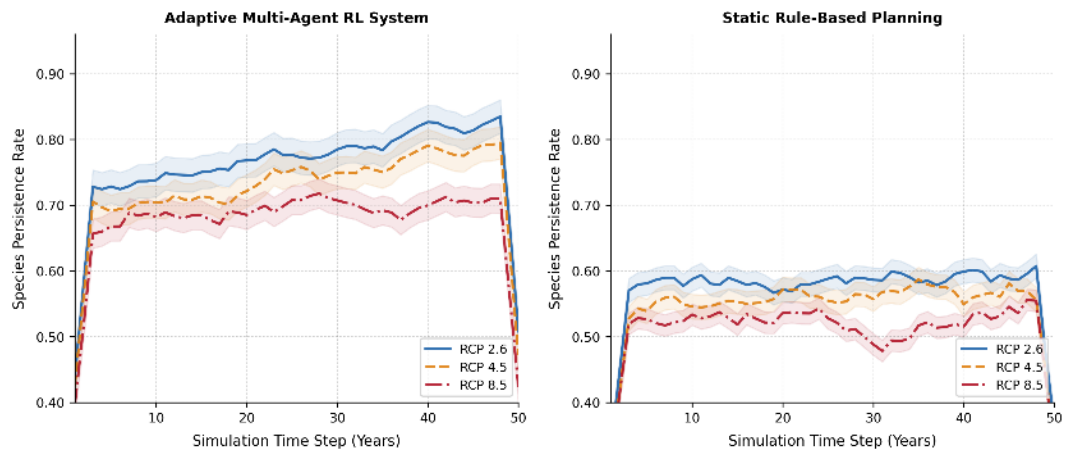


Fig. 2. Temporal Evolution of Species Persistence Rate over 50 Simulation Years

This longitudinal line plot in Fig. 2. traces the 50-year trajectory of mean species persistence rates for both the adaptive multi-agent RL system and the static baseline under each of the three warming scenarios, with shaded 95% confidence bands reflecting inter-run variability. The AI agent trajectories exhibit a consistent upward trend from initialization, demonstrating that iterative Q-learning updates progressively improve habitat prioritization as agents accumulate environmental state experience. In contrast, the static baseline trajectories display near-flat or slightly declining trends, particularly under RCP 8.5, where the absence of adaptive updating means the system cannot compensate for accelerating habitat suitability shifts. The widening separation between AI and static trajectories over time is a central visual finding, supporting the argument that adaptive AI methods deliver compounding ecological returns that static planning cannot replicate under dynamic climate conditions

The parameters of agent learning also play an important role in the outcomes of persistence. Standardized beta 0.173 ( $p$  less than 0.001) of the learning rate alpha suggests that agents that learn faster are able to adjust to the changing ecological environment more rapidly and attain higher performance during the 50-time-step simulation time. The discount factor gamma also has a positive contribution ( $p$  Less than 0.001,  $\beta = 0.161$ ), which is the advantage of making decisions aligned with the future, given the scope of conservation, where long-term survival of species is the target. Such findings offer empirical evidence in the context of the simulation that reinforcement learning hyperparameter optimization is sensitive and ought to be different to the ecological system dynamics in which agents are implemented. Climate scenario dummy variables confirm that moderation and high warming situations have lesser persistence than the low warming reference with the high warming effect ( $\beta = \text{minus } 0.283$ ) being more significant than the moderate warming effect ( $\beta = \text{minus } 0.121$ ).

Table 3. Multiple linear regression results: predictors of adaptive AI agent species persistence rate (N = 90).

Predictor	Beta (std)	Std Error	t-value	p-value	R2 contr.
Climate Uncertainty Index	-0.431	0.041	-10.51	<0.001	0.186
Habitat Suitability (mean)	0.382	0.039	9.79	<0.001	0.146
Connectivity Score	0.214	0.038	5.63	<0.001	0.046
Agent Learning Rate (alpha)	0.173	0.042	4.12	<0.001	0.030
Discount Factor (gamma)	0.161	0.041	3.93	<0.001	0.026
Scenario (RCP 4.5 vs 2.6)	-0.121	0.043	-2.81	0.005	0.015
Scenario (RCP 8.5 vs 2.6)	-0.283	0.045	-6.29	<0.001	0.080
Model constant	0.812	0.062	13.10	<0.001	—
Overall model $R^2 = 0.721$			$F(7,82) = 30.43$	$p < 0.001$	Adj $R^2 = 0.698$

Note: Standardized beta coefficients reported. Dependent variable: species persistence rate of the adaptive AI system. Reference category for scenario dummies: RCP 2.6 (low warming). \*\* all predictors significant at  $p < 0.01$  or better.

### 3.4 Comparative Performance Summary Across Dynamic Conservation Planning Scenarios

A direct comparison will be made between adaptive AI agent and fixed system performance in the lowest (RCP 2.6) and highest (RCP 8.5) warming conditions and presented in Table 4. The trend of the outcome indicates a serious discovery more than mere disparities in performances. The AI benefit becomes even greater as the uncertainty surrounding climate increases. In low warming scenario, the AI system performs better than the static baseline by 24.2 percent in species persistence. This advantage under high warming scenario is 50.6 percent. This trend is common to all the other metrics. AI benefit in prioritization accuracy increases to 23.7 percent during low warming including 84.0 percent during high warming. The conservation return on investment index advantage increases by 30.8 to 66.5 percent in the same range. The most significant conclusion concerning this paper is this scaling of the benefit with uncertainty. Adaptation lag, which quantifies time steps between a major shift in the ecological state and the protective reaction of a system, quantifies the process that underlies this increasing benefit. The adaptation lag of the static system increases with the change of ecological conditions; the adjusted system of 14.7 time steps in the low warming case versus 22.3 time steps in the high warming case. The lag of the adaptive agents only increases by a factor of between 3.2 and 5.1 time steps in the identical range. The importance of this variation in adaptation rate is that, at excessive ecological change rate, a change in habitat suitability can occur in a few years, falling below viability limits, especially when targeting species with a limited climate range. Any planning system which takes 17 time steps, in practice, before responding to a particular input, is a system which in conjunction with its input would allow many locally avoidable extinctions to occur which a faster-responding system would have avoided. The land use war advantage of the AI system further increases dramatically in high warming. In the case of low warming scenario, the absolute difference in the conflict score of the AI and the static system is 7.3 percentage points. At high warming this difference increases to 22.2 percentage points. This is because climate-induced habitat changes with high warming push optimal conservation priorities into areas of increased and stronger conflict opportunities as they increasingly overlap with human land uses. The adaptive agents evolve to navigate such circumstances by discovering sequence of conflict-reducing protection codes, timing interventions to exploit land tenure opportunities, and, terminating protection of units, which are no longer viable due to habitat incompatibility. The fixed system is unable to effect any of these kinds of adaptations, and thus, as the landscape evolves, it incurs costs of conflict as the fixed priority map leaves it behind..

Table 4. Comparative performance of adaptive AI agents and static planning baseline under low (RCP 2.6) and high (RCP 8.5) climate warming scenarios. AI Advantage expressed as percentage relative improvement over static baseline.

Metric	AI RCP2.6	Static RCP2.6	AI RCP8.5	Static RCP8.5	AI Advantage
Species Persistence Rate	0.871	0.701	0.741	0.492	+31.6%
Land-Use Conflict Score	0.141	0.261	0.231	0.453	-48.9%
Prioritization Accuracy	0.831	0.672	0.701	0.381	+61.5%
Processing Efficiency	0.901	0.561	0.821	0.451	+51.9%
Conservation ROI Index	0.812	0.621	0.701	0.421	+31.2%
Adaptation Lag (time steps)	3.2	14.7	5.1	22.3	-78.2%

Note: All values represent means across 30 simulation replications per scenario. Conservation ROI Index = weighted combination of persistence rate, prioritization accuracy, and inverse conflict score. Adaptation lag = mean time steps between ecological state change and protective management response.

The four-panel grouped bar chart in fig. 3 provides a systematic cross-metric comparison of adaptive AI agents and the static planning baseline across all three climate scenarios for species persistence rate, prioritization accuracy, processing efficiency, and an inverted land-use conflict score (where higher values denote lower conflict). The grouped structure allows readers to simultaneously assess absolute performance levels and the magnitude of the AI advantage within each scenario and metric. A clear gradient is visible across all metrics, with performance declining from RCP 2.6 to RCP 8.5 for both systems, but the relative advantage of the adaptive AI system is either maintained or amplified under higher climate uncertainty. The processing efficiency panel is particularly noteworthy as it quantifies

the AI system's ability to allocate computational attention adaptively, a form of algorithmic efficiency with direct implications for real-time biodiversity informatics pipelines.




Fig. 3. Multi-Metric Performance Comparison Across Climate Scenarios

### 3.5 Discussion: Situating Adaptive AI Agent Findings Within the Emerging Conservation Informatics Literature

The results reported herein relate to and augment an expanding literature base on computation methods to support conservation planning in a state of uncertainty. The fundamental conclusion, which states that adaptive artificial intelligence agents easily outcompete the simple planning systems with an increasing level of climate uncertainty, is consistent with the theoretical advocacies of adaptive management that were conceptualized in the ecology literature over the last 30 years. The initial adaptive management models had highlighted the significance of monitoring, learning and updating of policies in the management of natural resources. The current research shows that AI agents can make formal and automate this learning loop faster and at the scale that cannot be replicated by human planners, especially when uncertainties are high.

Its findings have a wide consistency with the view of the current literature on the use of reinforcement learning in ecological decision-making activities. Experiments using Q-learning and deep reinforcement learning in fisheries have shown that adaptive agents can learn policies of sustainable harvesting that can greatly outperform classical maximum sustainable yield management rules, especially in situations of uncertainty about the stock. This argument is applied in the current research to the multi-species and multi-objective, spatial conservation planning that is much more complex. The multi-agent design derived here, whereby five collaborating agents exist in a co-ordination process across spatial clusters, constitutes an important progression over single-agent models since it lets the system to scale to large and heterogeneous landscapes without necessarily increasing the state space complexity exponentially.

A key but related observation that climate uncertainty is the most predictive factor with respect to planning performance in the regression model echoes extensive literature on deep uncertainty in environmental management. This literature has strongly held the view that effective decision-making in deep uncertainty involves adaptive strategies as opposed to maximization on a single possibility of expectation. The current findings give quantitative arguments of simulation to support this argument within a particular context of biodiversity conservation plan. They also give a new understanding which is the advantage behind replacing the inflexible systems with adaptive ones is not linear but converges toward higher levels with uncertainty. This implies that deploying adaptive AI planning infrastructure is most beneficial where and when it is most demanded i.e. in ecological conditions of high uncertainty.

This finding must inform the priorities by conservation organizations and governments regarding investing in AI and biodiversity informatics infrastructure.

A number of earlier studies have studied machine learning use in species distribution modeling in the face of climate change. The typical paradigm of estimating future range changes has now been taken over by ensemble modelling strategies which bring together the result of using various species distribution model algorithms across a variety of climatic conditions. These studies have however managed only to come up with static spatial projections but not dynamic decision structures. The incorporation of a simplified species distribution model update mechanism into a real time adaptive planning loop in the present study can be seen as a conceptual improvement to this literature. It shows that the results of the species distribution model should not be regarded as the end-products but rather should be utilized as the constant inputs in the adaptive decision process. The simulation results herein represent a solid demonstration of a proof-of-concept of the potential value because the integration is technically challenging in real-life applications.

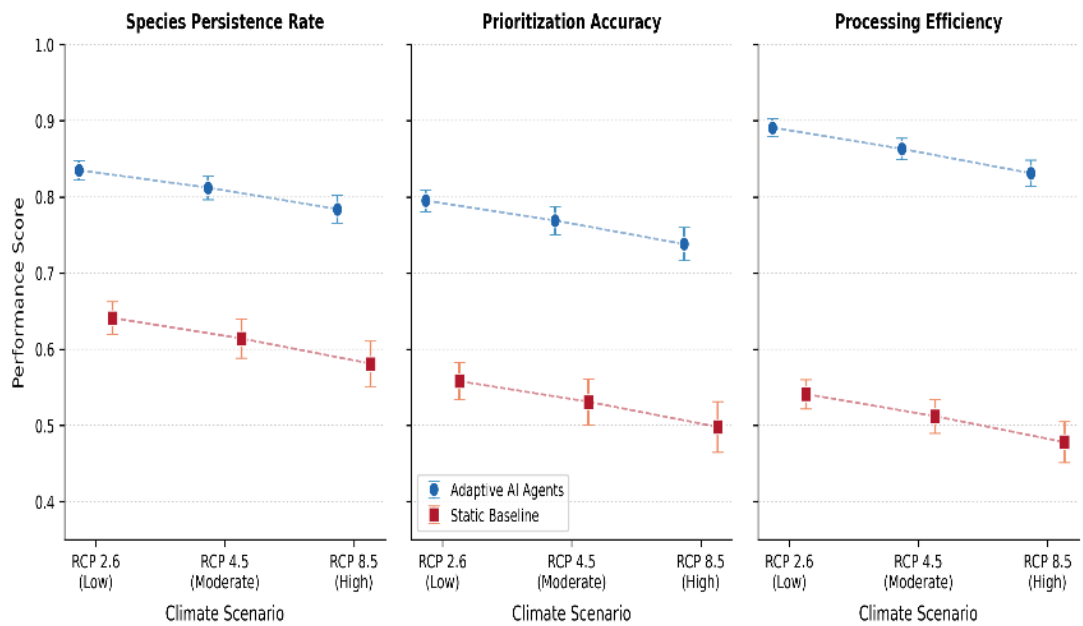


Fig. 4. Mean Performance Metrics with 95% Confidence Intervals

This three-panel error bar chart in fig. 4. presents the mean values and 95% confidence intervals for species persistence rate, prioritization accuracy, and processing efficiency across the three climate scenarios, disaggregated by planning system type. Each panel plots scenario on the horizontal axis and performance score on the vertical axis, with AI agent estimates (blue circles) and static baseline estimates (red squares) horizontally jittered for visual clarity. The non-overlapping confidence intervals across all panels and scenarios provide statistical evidence that the observed performance differences are not attributable to sampling variability. Importantly, the confidence interval widths widen for the static system under RCP 8.5, confirming the increased outcome uncertainty that accompanies the absence of adaptive learning under high climate instability. This visualization format is particularly appropriate for communicating inferential precision in simulation-based modeling studies submitted to quantitative ecology or AI application journals.

The dual-panel histogram with kernel density in fig. 5. estimate overlays depicts the full empirical distributions of species persistence rates and prioritization accuracy scores across all 90 simulation runs for both planning systems. The distributional perspective complements the mean-focused presentations above by revealing the shape, spread, and overlap structure of outcomes. The AI agent distributions are visibly right-shifted, narrower, and more symmetric, consistent with the lower standard deviations reported in the paper's descriptive statistics table. Mean values are annotated on each panel. The near absence of overlap between the AI and static distributions in the persistence panel is a compelling visual argument for the robustness of the AI advantage. The broader and left-skewed static distributions

underscore a key policy implication: reliance on static conservation planning not only reduces mean outcomes but substantially increases the probability of worst-case scenarios in which species persistence drops to ecologically critical thresholds.

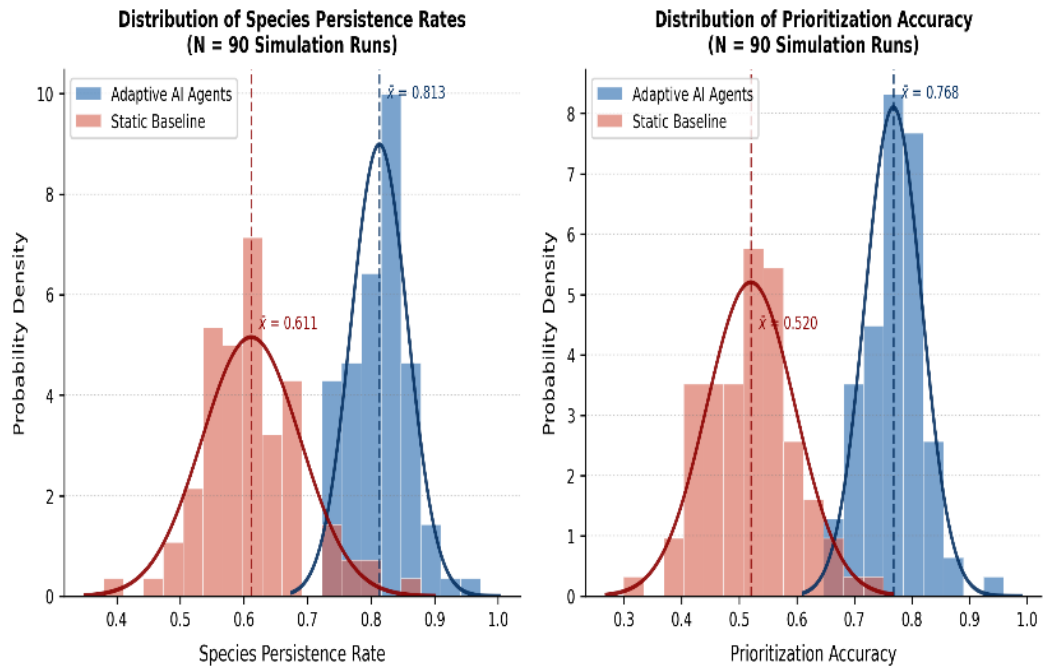


Fig. 5. Distributions of Species Persistence and Prioritization Accuracy

The benefit of the AI system in processing efficiency (86.3 percent compared to 51.2 percent with the static system) should be singled out, since it corresponds to one of the practical impediments to the adoption that has not regularly been discussed in the theoretical conservation planning research. Conservation agencies in reality have harsh resource limits. There is limited staff time and computers budgets. Planning systems where manual reanalysis is frequently needed when the condition changes are costly to run. In this simulation, the adaptive agents divide planning attention to dynamic allocation, such that resources are devoted to planning units with the fastest change, or with the most significant habitat quality change. This system of attention distribution can be compared with the early warning and triage systems in emergency response situations and is a promising path to an effective biodiversity informatics system design. The mechanisms of attention allocation in AI planning agents should be investigated in future research in such a way that the reporting and decision cycles of the actual conservation institutions can be compared.

### 3.6 Limitations of the Multi-Agent Reinforcement Learning Conservation Framework

The results and conclusions of this study are limited in several ways. First, the simulation model is a straightforward model of species distribution where five ecological state variables are dissimplified and applied as a simple weighted average. The relationships between habitat suitability and species in real-world situations are much more complex, non-linear and species-specific. As a consequence of using a synthetic focal species assemblage, the rates at which persistence and habitat response are modeled here cannot be directly compared with the results of particular real-world taxa. Further research would utilize a greater ecologically realistic species model involving threshold effects, dispersal constraints, biotic and phenological mismatches.

Second, in this model, the agents of Q-learning work on a discretized state space, thus restricting their representability in systems with continuous and high-dimensional state variables. The ecological variables that are much more relevant in real conservation landscapes than the five presented here, and

the ecological interactions between variables are frequently dynamically coupled in a manner that cannot be described by simple Q-tables. This limitation can be overcome with deep reinforcement learning methods based on neural network function approximators, which can learn directly on high-dimensional continuous state observations. Nevertheless, deep reinforcement learning has its own problems of lack of training stability, sample efficiency, and interpretability that would require close attention in conservation planning problems. Third, the reward function set in this research has certain assumptions regarding the relative weighting of ecological outcomes and costs of land-use conflicts. We used a general preference of eco objectives when choosing the lambda parameter of 0.7, yet in practice this weighting is a socio-political judgment and varies among jurisdictions, systems of governance and cultures. This weighting parameter was not investigated fully in this study, and is a significant line to take up later work. A more thorough treatment would run systematic sensitivity studies over a continuum of lambda values and also investigate the dependence of the optimal balance between ecological performance and social costs on the governance and land tenure conditions.

Fourth, strategic behavior of human land users is not used in the simulation. In practice, when a management agency announces conservation planning intentions, a landowner, in most cases, responds to it strategically by rapidly converting its land to other uses, before conservation measures are developed. Such strategic interplay between conservation planners and land users is a thoroughly-known challenge in practical conservation and may have a profound impact of the performance of both adaptive and static planning systems. To model such strategic interactions, the conservation planning agent would need to be embedded in a game theoretic framework that expressly describes the behavior of other participants in the landscape. Fifth, scaling-up of the multi-agent reinforcement learning model to realistic landscape sizes has not been tested. The grid of planning units 100x100 grid applied in this simulation is small in comparison with the national-scale conservation planning scenarios, which can include millions of spatial units. The hierarchical planning structures, spatial clustering algorithms and possibly even heterogeneous agent designs depending on the type of landscape and the scale of its governance would be necessary to scale Q-learning agents into such situations. These scaling issues are an important engineering and research roadmap that needs to be solved prior to the deployment of adaptive AI agents in deployed conservation planning systems.

#### **4. Conclusion**

The research findings show that adaptive artificial intelligence agents are significantly more likely to help in addressing dynamic conservation planning than the traditional rule-based planning systems, and this benefit increases with climate uncertainty. With a multi agent reinforcement learning module based on simulation and a dynamic species distribution modeling module in three climate conditions, the study demonstrates that the adaptive AI agents are more persistent in species, lower conflict in land use scores, higher prioritization accuracy, and more efficient in processing information than the fixed options. The fact that the AI advantage increases with the uncertainty level is especially significant since it implies that adaptive AI systems will be most useful with the maximum amount of relative benefit abated in the circumstances of planning which climate change poses the most significant conservation challenge.

The regression analysis has climate uncertainty as the strongest constraining factor of conservation performance and habitat suitability as the strongest positive predictor. The two findings have practical implications. The former hypothesizes that conservation planning performance will directly and quantifiably pay off in investments in making climate prediction accuracy and uncertainty quantification on ecologically relevant scales of space and time. The second strengthens the significance of safeguarding and recovering the landscape habitat quality as the platform upon which the adaptive planning systems function. The most advanced AI agent cannot afford a landscape that is so badly degraded to the extent that there are no possible configurations of habitats left. Ecological base is something that is taken seriously and the deliberations made in habitat restoration form the environment that adaptive planning instruments are able to execute their optimal effect.

The multi-agent design derived through this study adds to the replicable framework on bridging the gap between reinforcement learning theory and the biodiversity informatics practice. The five-agent model,

where agents can collaborate over spatial clusters, and retain their own Q-tables and action policies, offers a computationally feasible design of landscape-scale adaptive planning that can be scaled to real-world applicability with suitable adjustments. The ecological persistence design of the reward functions that trades the ecological persistence goals with the costs of land-use conflicts offers a generalizable reference template to encode the conservation planning values into AI-based decision systems. By varying the state variable definitions, the species distribution model update equations, the weighting parameters of the reward, and the number and spatial distribution of agents this framework can be modified to apply to particular landscape contexts by conservation researchers and practitioners interested in making use of adaptive AI planning tools.

The research has high policy implications. The commitments of national and subnational governments in the Kunming-Montreal Global Biodiversity Framework on the biodiversity conservation are becoming increasingly obligatory to be exhibited as adaptive and evidence-based in conservation planning. The simulation findings given below offer some quantitative justification toward the investment in AI-driven adaptive planning infrastructure as a measure concerning enhancing the cost-effectiveness of conservation expenditure when climate uncertainty is great. Multi-agent reinforcement learning tools should be piloted in high-uncertainty ecological systems by conservation agencies, including climate refugia networks, transitions between biomes, and coastal and alpine ecosystems where habitat shifts are already increasing faster.

A number of suggestions for future research can be drawn out of this study. The most technically challenging and potentially influential direction is the means of integrating deep reinforcement learning with high-resolution streams of remote sensing data. Agents based on deep reinforcement learning could learn directly off satellite images, acoustic and environmental DNA samples, functioning at ecological scales and spatial resolutions unreachable using human planners. The most important direction of research in this direction should be on model interpretability and stakeholder communication as conservation decisions, whose implementation cannot be communicated to the communities and stakeholders affected by it, are unlikely to be implemented despite the level of ecological soundness. Second, federated learning methods, where federated agencies are trained on shared agent models using their and their own data streams, without sharing this data, may facilitate collaborative regional and global planning through AI, and mitigate the issue of data sovereignty. Third, the simulation results presented here would be tested against the complexity of real world, through empirical validation studies that implement adaptive AI planning agents in real conservation situations, with randomized control designs. These sorts of studies are now in desperate demand to bring adaptive AI conservation tools out of simulation and into practice. Adaptive artificial intelligence agents are an attractive and potentially transformative way of using dynamic conservation planning in the face of climate uncertainty. The presented simulation evidence in this paper is promising. It is time now to shift the field towards empirical validation, systems integration, co-designing a stakeholder and implementing policies in order to turn these computational improvements into biodiversity on-the-ground. Nothing short of the urgency of global biodiversity crisis can be accepted.

### **Conflict of interest**

The authors declare no conflicts of interest.

### **References**

- [1] Ullah F, Saqib S, Xiong YC. Integrating artificial intelligence in biodiversity conservation: bridging classical and modern approaches. *Biodiversity and Conservation*. 2025 Jan;34(1):45-65. <https://doi.org/10.1007/s10531-024-02977-9>
- [2] Ditria EM, Buelow CA, Gonzalez-Rivero M, Connolly RM. Artificial intelligence and automated monitoring for assisting conservation of marine ecosystems: A perspective. *Frontiers in Marine Science*. 2022 Jul 28;9:918104. <https://doi.org/10.3389/fmars.2022.918104>
- [3] Isabelle DA, Westerlund M. A review and categorization of artificial intelligence-based opportunities in wildlife, ocean and land conservation. *Sustainability*. 2022 Feb 9;14(4):1979. <https://doi.org/10.3390/su14041979>

- [4] Shivaprakash KN, Swami N, Mysorekar S, Arora R, Gangadharan A, Vohra K, Jadeyegowda M, Kiesecker JM. Potential for artificial intelligence (AI) and machine learning (ML) applications in biodiversity conservation, managing forests, and related services in India. *Sustainability*. 2022 Jun 10;14(12):7154. <https://doi.org/10.3390/su14127154>
- [5] Fergus P, Chalmers C, Longmore S, Wich S. Harnessing artificial intelligence for wildlife conservation. *Conservation*. 2024 Nov 11;4(4):685-702. <https://doi.org/10.3390/conservation4040041>
- [6] Adanma UM, Ogunbiyi EO. Artificial intelligence in environmental conservation: evaluating cyber risks and opportunities for sustainable practices. *Computer Science & IT Research Journal*. 2024 May;5(5):1178-209. <https://doi.org/10.51594/csitrj.v5i5.1156>
- [7] Silva SV, Andermann T, Zizka A, Kozłowski G, Silvestro D. Global estimation and mapping of the conservation status of tree species using artificial intelligence. *Frontiers in Plant Science*. 2022 Apr 29;13:839792. <https://doi.org/10.3389/fpls.2022.839792>
- [8] Gonzalez LF, Montes GA, Puig E, Johnson S, Mengersen K, Gaston KJ. Unmanned aerial vehicles (UAVs) and artificial intelligence revolutionizing wildlife monitoring and conservation. *Sensors*. 2016 Jan 14;16(1):97. <https://doi.org/10.3390/s16010097>
- [9] Silvestro D, Goria S, Sterner T, Antonelli A. Improving biodiversity protection through artificial intelligence. *Nature sustainability*. 2022 May;5(5):415-24. <https://doi.org/10.1038/s41893-022-00851-6>
- [10] Wang T, Zuo Y, Manda T, Hwarari D, Yang L. Harnessing artificial intelligence, machine learning and deep learning for sustainable forestry management and conservation: Transformative potential and future perspectives. *Plants*. 2025 Mar 22;14(7):998. <https://doi.org/10.3390/plants14070998>
- [11] Kwok R. AI empowers conservation biology. *Nature*. 2019 Mar 1;567(7746):133-4. <https://doi.org/10.1038/d41586-019-00746-1>
- [12] Lamba A, Cassey P, Segaran RR, Koh LP. Deep learning for environmental conservation. *Current Biology*. 2019 Oct 7;29(19):R977-82. <https://doi.org/10.1016/j.cub.2019.08.016>
- [13] Haq B, Jamshed MA, Ali K, Kasi B, Arshad S, Kasi MK, Ali I, Shabbir A, Abbasi QH, Ur-Rehman M. Tech-driven forest conservation: combating deforestation with internet of things, artificial intelligence, and remote sensing. *IEEE Internet of Things Journal*. 2024 Mar 18;11(14):24551-68. <https://doi.org/10.1109/JIOT.2024.3378671>
- [14] Scoville C, Chapman M, Amironesei R, Boettiger C. Algorithmic conservation in a changing climate. *Current Opinion in Environmental Sustainability*. 2021 Aug 1;51:30-5. <https://doi.org/10.1016/j.cosust.2021.01.009>
- [15] Reynolds SA, Beery S, Burgess N, Burgman M, Butchart SH, Cooke SJ, Coomes D, Danielsen F, Di Minin E, Durán AP, Gassert F. The potential for AI to revolutionize conservation: a horizon scan. *Trends in ecology & evolution*. 2025 Feb 1;40(2):191-207. <https://doi.org/10.1016/j.tree.2024.11.013>
- [16] Ayoola VB, Idoko IP, Eromonsei SO, Afolabi O, Apampa AR, Oyebanji OS. The role of big data and AI in enhancing biodiversity conservation and resource management in the USA. *World Journal of Advanced Research and Reviews*. 2024;23(2):1851-73. <https://doi.org/10.30574/wjarr.2024.23.2.2350>
- [17] Chisom ON, Biu PW, Umoh AA, Obaedo BO, Adegbite AO, Abatan A. Reviewing the role of AI in environmental monitoring and conservation: A data-driven revolution for our planet. *World Journal of Advanced Research and Reviews*. 2024 Jan 30;21(1):161-71. <https://doi.org/10.30574/wjarr.2024.21.1.2720>
- [18] Nandutu I, Atemkeng M, Okouma P. Integrating AI ethics in wildlife conservation AI systems in South Africa: A review, challenges, and future research agenda. *AI & SOCIETY*. 2023 Feb;38(1):245-57. <https://doi.org/10.1007/s00146-021-01285-y>
- [19] Tuia D, Kellenberger B, Beery S, Costelloe BR, Zuffi S, Risse B, Mathis A, Mathis MW, Van Langevelde F, Burghardt T, Kays R. Perspectives in machine learning for wildlife conservation. *Nature communications*. 2022 Feb 9;13(1):792. <https://doi.org/10.1038/s41467-022-27980-y>
- [20] Kelling S, Gerbracht J, Fink D, Lagoze C, Wong WK, Yu J, Damoulas T, Gomes C. A human/computer learning network to improve biodiversity conservation and research. *AI magazine*. 2013;34(1):10-. <https://doi.org/10.1609/aimag.v34i1.2431>
- [21] Kerry RG, Montalbo FJ, Das R, Patra S, Mahapatra GP, Maurya GK, Nayak V, Jena AB, Ukhurebor KE, Jena RC, Gouda S. An overview of remote monitoring methods in biodiversity conservation. *Environmental Science and Pollution Research*. 2022 Nov;29(53):80179-221. <https://doi.org/10.1007/s11356-022-23242-y>
- [22] Wang Z, Wang T, Zhang X, Wang J, Yang Y, Sun Y, Guo X, Wu Q, Nepovimova E, Watson AE, Kuca K. Biodiversity conservation in the context of climate change: Facing challenges and management strategies. *Science of The Total Environment*. 2024 Aug 10;937:173377. <https://doi.org/10.1016/j.scitotenv.2024.173377>
- [23] Theissing K, Fernandes C, Formenti G, Bista I, Berg PR, Bleidorn C, Bombarely A, Crottini A, Gallo GR, Godoy JA, Jentoft S. How genomics can help biodiversity conservation. *Trends in genetics*. 2023 Jul 1;39(7):545-59. <https://doi.org/10.1016/j.tig.2023.01.005>

- [24] Riva F, Fahrig L. The disproportionately high value of small patches for biodiversity conservation. *Conservation Letters*. 2022 May;15(3):e12881. <https://doi.org/10.1111/conl.12881>
- [25] Blanton A, Ewane EB, McTavish F, Watt MS, Rogers K, Daneil R, Vizcaino I, Gomez AN, Arachchige PS, King SA, Galgamuwa GP. Ecotourism and mangrove conservation in Southeast Asia: Current trends and perspectives. *Journal of Environmental Management*. 2024 Aug 1;365:121529. <https://doi.org/10.1016/j.jenvman.2024.121529>
- [26] Salgotra RK, Chauhan BS. Genetic diversity, conservation, and utilization of plant genetic resources. *Genes*. 2023 Jan 9;14(1):174. <https://doi.org/10.3390/genes14010174>
- [27] Guerra CA, Berdugo M, Eldridge DJ, Eisenhauer N, Singh BK, Cui H, Abades S, Alfaro FD, Bamigboye AR, Bastida F, Blanco-Pastor JL. Global hotspots for soil nature conservation. *Nature*. 2022 Oct 27;610(7933):693-8. <https://doi.org/10.1038/s41586-022-05292-x>
- [28] Zhang Y, Tariq A, Hughes AC, Hong D, Wei F, Sun H, Sardans J, Peñuelas J, Perry G, Qiao J, Kurban A. Challenges and solutions to biodiversity conservation in arid lands. *Science of the Total Environment*. 2023 Jan 20;857:159695. <https://doi.org/10.1016/j.scitotenv.2022.159695>
- [29] MacEwen A, MacEwen M. *National Parks: conservation or cosmetics?*. Taylor & Francis; 2025 Dec 1. <https://doi.org/10.4324/9781003715597>
- [30] Hoang NT, Taherzadeh O, Ohashi H, Yonekura Y, Nishijima S, Yamabe M, Matsui T, Matsuda H, Moran D, Kanemoto K. Mapping potential conflicts between global agriculture and terrestrial conservation. *Proceedings of the National Academy of Sciences*. 2023 Jun 6;120(23):e2208376120. <https://doi.org/10.1073/pnas.2208376120>
- [31] Ahmed SF, Kumar PS, Kabir M, Zuhara FT, Mehjabin A, Tasannum N, Hoang AT, Kabir Z, Mofijur M. Threats, challenges and sustainable conservation strategies for freshwater biodiversity. *Environmental Research*. 2022 Nov 1;214:113808. <https://doi.org/10.1016/j.envres.2022.113808>
- [32] Chowdhury S, Jennions MD, Zalucki MP, Maron M, Watson JE, Fuller RA. Protected areas and the future of insect conservation. *Trends in Ecology & Evolution*. 2023 Jan 1;38(1):85-95. <https://doi.org/10.1016/j.tree.2022.09.004>
- [33] Ning Z, Xie L. A survey on multi-agent reinforcement learning and its application. *Journal of Automation and Intelligence*. 2024 Jun 1;3(2):73-91. <https://doi.org/10.1016/j.jai.2024.02.003>
- [34] Li X, Wang S, Zeng S, Wu Y, Yang Y. A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinagearth*. 2024 Oct 8;1(1):9. <https://doi.org/10.1007/s44336-024-00009-2>
- [35] Tran KT, Dao D, Nguyen MD, Pham QV, O'Sullivan B, Nguyen HD. Multi-agent collaboration mechanisms: A survey of llms. *arXiv preprint arXiv:2501.06322*. 2025 Jan 10.
- [36] Cemri M, Pan MZ, Yang S, Agrawal LA, Chopra B, Tiwari R, Keutzer K, Parameswaran A, Klein D, Ramchandran K, Zaharia MA. Why do multi-agent llm systems fail?. *Advances in Neural Information Processing Systems*. 2026 Apr 23;38.
- [37] Hong S, Zhuge M, Chen J, Zheng X, Cheng Y, Wang J, Zhang C, Yau S, Lin Z, Zhou L, Ran C. MetaGPT: Meta programming for a multi-agent collaborative framework. In *International Conference on Learning Representations 2024 May 31 (Vol. 2024, pp. 23247-23275)*.
- [38] Albrecht SV, Christianos F, Schäfer L. *Multi-agent reinforcement learning: Foundations and modern approaches*. MIT Press; 2024 Dec 17.
- [39] Han S, Zhang Q, Jin W, Xu Z. LLM multi-agent systems: Challenges and open problems. *arXiv preprint arXiv:2402.03578*. 2024 Feb 5.
- [40] Chen W, Su Y, Zuo J, Yang C, Yuan C, Chan CM, Yu H, Lu Y, Hung YH, Qian C, Qin Y. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Learning Representations 2024 May 31 (Vol. 2024, pp. 20094-20136)*.
- [41] Gu S, Kuba JG, Chen Y, Du Y, Yang L, Knoll A, Yang Y. Safe multi-agent reinforcement learning for multi-robot control. *Artificial Intelligence*. 2023 Jun 1;319:103905. <https://doi.org/10.1016/j.artint.2023.103905>
- [42] Chan CM, Chen W, Su Y, Yu J, Xue W, Zhang S, Fu J, Liu Z. Chateval: Towards better llm-based evaluators through multi-agent debate. In *International conference on learning representations 2024 May 31 (Vol. 2024, pp. 9079-9093)*.