

Article

ACO-Path: ACO-Based Informative Path Planning with Gaussian Processes for Water Monitoring with a Fleet of ASVs

Micaela Jara Ten Kathen ^{1,2,3} , Natalia Benitez ², Mario Arzamendia ⁴  and Daniel Gutiérrez Reina ^{5,*} 

¹ Center for the Advancement of Technology in Society and Industry, Taiwan-Paraguay Polytechnic University, Av. Sebastian Gaboto, Asunción 001018, Paraguay; mjara@uppt.edu.py

² Facultad de Ciencias y Tecnología, Universidad Católica “Nuestra Señora de la Asunción”, Supercarretera, s/n, Hernandarias 100519, Alto Paraná, Paraguay; natalia.benitez@uc.edu.py

³ Centro de Innovación Social y Gestión Territorial, Fundación Parque Tecnológico Itaipú—Paraguay, Supercarretera, Hernandarias 007220, Alto Paraná, Paraguay

⁴ Facultad de Ingeniería, Universidad Nacional de Asunción, San Lorenzo 2160, Central, Paraguay; marzamendia@ing.una.py

⁵ Department of Electronic Engineering, Technical School of Engineering of Seville, C. Americo Vespucio, 41092 Seville, Spain

* Correspondence: dgutierrezreina@us.es

Abstract

Autonomous surface vehicles can support water-quality monitoring, but they require planners that place measurements where they most improve the environmental estimate under mission constraints. This paper proposes ACO-Path, an informative path planner that couples Ant Colony Optimization -Ant System- with online Gaussian Process mapping. During the mission, the Gaussian Process updates a mean or contamination map and a variance or uncertainty map, from which dynamic action zones are derived and used to guide an explicit explore then exploit policy. The method is evaluated in a simulated water resource monitoring scenario inspired by Lake Ypacaraí, considering three exploration distances and two heuristic weights. In a comparison against five baseline planners, ACO-Path achieves the lowest hotspot error, $Error_{peak} = 0.19896 \pm 0.39400$, while remaining competitive in global reconstruction, $MSE_{map} = 0.00144 \pm 0.00348$, $R^2 = 0.96066 \pm 0.09861$. In addition, a turning analysis based on the absolute heading change between consecutive segments $|\Delta\alpha|$ shows that ACO-Path produces smoother trajectories, with fewer sharp turns $|\Delta\alpha| \geq 45^\circ$ than counterpart baselines under the same mission constraints.

Keywords: ant colony optimization; autonomous surface vehicles; Gaussian process; informative path planning; water monitoring



Academic Editors: Ulises Orozco-Rosas and Oscar Montiel Ross

Received: 30 December 2025

Revised: 31 January 2026

Accepted: 2 February 2026

Published: 4 February 2026

Correction Statement: This article has been republished with a minor change. The change does not affect the scientific content of the article and further details are available within the backmatter of the website version of this article.

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The decline in water quality across oceans, rivers, and lakes stems from pressures acting at multiple scales: diffuse nutrient inputs from agriculture, urban and industrial effluents, emerging contaminants, and climate forcing that intensifies hydrometeorological extremes [1,2]. In this context, global assessments that integrate land use, hydrology, and water quality show that clean-water scarcity triples when pollution is included in availability models, reshaping management priorities across thousands of sub-basins [3].

Among emerging contaminants, microplastics have shifted from a surface issue to a full water-column problem: a synthesis of 1885 vertical profiles (2014–2024) reports concentrations from $\approx 10^{-4}$ to 10^4 particles per m^3 and vertical patterns driven by particle size, density, and mixing dynamics evidence that surface only sampling is insufficient [4].

As a result, there are direct effects on organisms and ecosystems: many marine species mistake these particles for food, leading to contamination risks and adverse impacts that propagate through food webs.

Another driver of degradation is acidification, which reflects a sustained chemical alteration of the water. The ocean absorbs CO_2 which, once dissolved, forms carbonic acid, lowers pH, reduces carbonate ions (CO_3^{2-}), and decreases the saturation state of calcium carbonate (CaCO_3), making calcification more difficult [5]. The outcome is not a single “toxic” substance but a more corrosive habitat for corals, mollusks, and calcifying plankton, with documented effects on growth and larval survival [6] as well as increased shell fragility [7], with knock-on consequences for food webs and fisheries [8].

In parallel, harmful algal blooms (HABs) have intensified since the 1980s, with a 44% rise when comparing the 2000s to the 2010s across multiple regions, driven by nutrient surpluses linked to urbanization, wastewater, agricultural expansion, and stratification [9]. This trend calls for early warning systems and continuous observing networks. Documented cases include cyanobacterial blooms in Lake Taihu (China), with recent analyses of environmental drivers [10], recurrent events in Lake Erie (USA) associated with cyanotoxin risks [11], and episodes in Lake Okeechobee (USA) linked to nutrient management and releases to urban estuaries [12]. In South America, Lake Ypacaraí (Paraguay) shows eutrophication and cyanobacterial proliferation under watershed pressures, with environmental and social impacts [2]. Overall, these examples underscore the urgent need for efficient, continuous water-quality monitoring capable of informing timely environmental management and policy decisions.

Given this urgency, the 2030 Agenda of the United Nations prioritizes clean water and sanitation (SDG 6), promoting comprehensive and consistent monitoring methods to support sustainable management practices [13]. However, traditional monitoring approaches, such as fixed laboratories and manual sampling, have persistent limitations: high operating costs, limited spatial coverage, staff exposure to hazards, and inflexibility [14].

To address these limits, autonomous surface vehicles (ASVs) have emerged as a robust option in recent years [15,16]. Equipped with specialized sensors, ASVs enable real-time data collection and adaptive environmental monitoring while reducing human exposure. Within the ASV system, Informative Path Planning (IPP) is a key component: it generates trajectories dynamically from collected data to maximize information under energy, time, and sensor constraints [17].

In recent IPP literature, deep reinforcement learning (DRL) has been used to learn adaptive sensing and monitoring policies that can be executed online after an offline training phase, which is useful under partial observability. Among recent applications, ref. [18] proposes an adaptive IPP approach that combines tree search with an offline-trained neural network to predict informative sensing actions. In simulation, the method matches benchmark performance while reducing computing time, and it is additionally validated with real surface-temperature data. In a related direction, ref. [19] presents a framework in which UAVs autonomously acquire training images to retrain semantic segmentation models, reporting performance gains and reduced labeling effort when compared against local planning baselines.

More generally, DRL-based IPP has also been studied with action-space design mechanisms that keep online decision-making tractable, such as constructing a local graph online to restrict available actions while supporting replanning [20]. A broader synthesis of learning-based adaptive IPP methods, including common design choices and open issues such as scenario coverage and transfer validation, is provided in [21].

For aquatic monitoring, the work [22] combines local Gaussian Processes (GPs) with a DRL policy that conditions its decisions on the posterior mean and variance through an

information-gain reward. Safety and coordination are handled through a consensus-based heuristic, and the reported results show lower estimation errors compared to alternative monitoring approaches. A related line of work, ref. [23], formulates continuous water-quality patrolling as a partially observable Markov game for Lake Ypacaraí and adopts a multi-agent deep Q-learning setup with a shared policy for homogeneous ASVs. In that formulation, exploration and intensification are treated as separate phases, and a transition variable controls the shift toward prioritizing highly polluted areas while maintaining revisit behavior.

Swarm-intelligence planners have also shown promise. AquaFeL-PSO [16] combines Particle Swarm Optimization (PSO), federated learning, and Gaussian Processes (GPs) as a surrogate model in two phases: exploration to estimate an initial model and exploitation that partitions the search space into action zones to refine hotspots. In a case study, the planner achieved a 300% improvement in water-quality modeling and about 4000% in peak detection. In [24], the authors present a hybrid HGWO–PSO planner that blends Gray Wolf Optimization (GWO) for exploration with PSO for exploitation in static environments. Validation across four obstacle-rich scenarios for a mobile robot shows shorter paths than PSO, GWO, and other heuristics algorithms.

Within Ant Colony Optimization (ACO)-based path planners, several recent contributions stand out. Ref. [25] introduces MACOGA, which uses ACO to propose feasible routes and an enhanced Genetic Algorithm (GA) to refine them, incorporating pheromones and heuristic factors with adaptive probabilistic crossover/mutation and a reconnection step to ensure length and smoothness. Across six grid models of different sizes and complexity, MACOGA attains shorter paths and planning times, with 100% success in relatively complex settings. In turn, ref. [26] proposes ADL-ACO (two layers) for dynamic planning and ADWA for real-time obstacle avoidance: the first speeds convergence and improves global search via adaptive parameter tuning, and the second refines length, number of turns, safety, and smoothness using segmented B-splines. In an industrial context, ref. [27] integrates ACO with Bayesian Optimization (BO) to tune critical parameters and prioritize high-demand regions on CAD-derived maps, improving coverage and reducing waste. For multi-vehicle adaptive ocean sampling, ref. [28] presents V-ACO, which couples Voronoi partitioning (with tournament selection) and a modified-heuristic ACO to generate collision-free trajectories under mission time, inter-vehicle spacing, and obstacle constraints. Simulations maximize data collection in high-interest regions, and field tests confirm practical feasibility.

In summary, IPP with autonomous vehicles has advanced in quantifying uncertainty and adopting adaptive replanning. Even so, most ACO-based planners remain focused on global/local path planning and do not integrate an informative map derived from water-quality models to guide sampling. Consequently, current ACO approaches remain disconnected from predictive models that use uncertainty information to guide exploration and contamination data to drive exploitation.

To address this gap, this work proposes the ACO-Path, an IPP method that combines ACO with a GP as an online surrogate model of the WQPs and maintains an exploration–exploitation logic through dynamic action zones. The GP is updated online with ASV measurements and provides two outputs: the posterior mean (contamination map) and the posterior variance (model uncertainty). Building on the notion of action zones introduced in AquaFeL-PSO [16], ACO-Path defines operational regions from the GP models to steer where ASVs are more likely to transit: high-variance zones promote exploration by densifying sampling where knowledge is weak, while high-mean zones promote exploitation by characterizing hotspots.

This paper is presented as a proof of concept for using ACO as the backbone of an IPP strategy in a multi-ASV setting. In this design, the GP updates the action zones and, through them, the set of admissible candidate locations, while ACO optimizes the route within that set. To keep the contribution focused, the ACO component is kept simple by adopting the classical Ant System as the baseline ACO variant [29,30]. More recent ACO variants could be incorporated within the same framework as a follow-up step, but they are not the objective of the present study. The case study is based on the Lake Ypacaraí scenario.

In contrast to DRL pipelines that require extensive training data and careful validation to ensure transfer across scenarios, the present work keeps the planner training-free and uses the GP only to update the action zones and, consequently, the set of candidate locations that ASVs can visit.

The main contributions of this work are as follows:

- A new informative path planner that integrates Ant Colony Optimization (Ant System) with an online Gaussian Process (GP) updated throughout the mission, so that waypoint selection is driven by both the estimated contamination (GP mean) and the associated uncertainty (GP variance). The resulting trajectories are smoother, which can reduce the burden on the low-level controller and also lower energy demand during execution.
- A GP-driven action zone construction that translates the mean and variance maps into candidate sampling regions and implements an explore then exploit policy: early decisions prioritize uncertainty reduction, and later decisions focus on informative (potential hotspot) areas using a combined criterion.

This paper is organized as follows. Section 2 formulates the monitoring problem and summarizes the operational assumptions. Section 3 presents the methodological foundations of the proposed IPP approach. Section 4 introduces ACO-Path, the proposed IPP. Section 5 describes the simulation setup, ground-truth generation, evaluation metrics, parameter settings, and experimental results. Section 6 discusses the main findings and their implications. Finally, Section 7 concludes the paper and outlines directions for future work.

2. Problem Formulation

The goal of the monitoring mission is to recover a reliable spatial model of a WQP over the water body. Each ASV p in the fleet carries a WQP sensor S . In this study, the task is single-objective, so $S = 1$, and each observation is a scalar measurement m_k of the target variable. For simulation, a ground truth of the WQP $y(\mathbf{x})$ is defined over planar coordinates $\mathbf{x} = (x, y)$. During the mission, the ASVs collect M measurements, stored as values $\mathcal{M}$ with their sampling locations $\mathcal{Q}$. Given data $(\mathcal{Q}, \mathcal{M})$, a regression model produces an estimate model of the WQP of Lake Ypacaraí $\hat{y}(\mathbf{x})$:

$$\hat{y}(\mathbf{x}) \approx y(\mathbf{x}). \quad (1)$$

The optimization objective is to minimize the discrepancy between the estimated model $\hat{y}$ and the ground truth y under a mission-length budget. Using mean squared error (MSE) over a set of evaluation points $\{\mathbf{x}_i\}_{i=1}^N$, and constraining the average path length per vehicle by a maximum distance $d_{\max}$, the problem is

$$\begin{aligned} \min \quad & f(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N (y(\mathbf{x}_i) - \hat{y}(\mathbf{x}_i))^2 \\ \text{s.t.} \quad & \text{dist}_p \leq d_{\max}, \end{aligned} \quad (2)$$

where dist_p is the path length traveled by ASV p .

Assumptions

Assumption 1. The lake surface is discretized as an occupancy grid $\mathcal{R} \in \{0,1\}^{r \times n}$, where cell $\mathcal{R}_{i,j} = 1$ denotes navigable water (black grids) and $\mathcal{R}_{i,j} = 0$ denotes land, restricted areas, or obstacles (white grids). ASVs may only visit cells with value 1. The grid resolution is denoted by Δ (Figure 1).




Figure 1. Occupancy grid for Lake Ypacaraí; black cells are navigable.

Assumption 2. ASV motions are synchronized. During execution, each vehicle relies on an on-board local path planner to avoid obstacles and prevent inter-ASV collisions, applying short-range corrections while tracking the nominal path provided by the IPP. Initial positions are placed at piers or clear shoreline areas to ensure immediate access to navigable water and to start the mission with safe separation.

Assumption 3. The vehicle mission is constrained by both mobility and autonomy. The ASV speed is bounded by a maximum value $v_{\max}$, and the mission length is enforced through a maximum traveled distance per vehicle, $dist_{\max}$.

Assumption 4. Vehicles communicate reliably with a central server (no dropouts are considered). The coordination is centralized.

Assumption 5. The physico-chemical variables (pH, temperature, conductivity, turbidity) are assumed to vary smoothly enough for GP interpolation. Sensors are calibrated. The measurements are treated as noise-free at this stage. ASVs sample in real time and send data to the server, which updates the GP map as new data arrive.

3. Methodology

This section introduces the methodological components used to build the proposed IPP: the Ant Colony Optimization algorithm, the Gaussian Process model, and the action zone construction adopted from AquaFeL-PSO.

3.1. Ant Colony Optimization

ACO is a metaheuristic inspired by the foraging behavior of ants, in which a population of simple agents builds candidate routes on a graph while exchanging information through virtual “pheromone” trails [30]. In path planning form, nodes represent feasible waypoints and edges represent admissible motions between cells. Artificial ants traverse the graph to construct collision-free paths while the colony biases future moves toward promising edges [31].

Ant System (AS) is the first algorithm in the ACO family, introduced by Dorigo, Maniezzo, and Colomi in the early 1990s [31]. In this work, the Ant System (AS) formulation is used as the baseline ACO variant. AS provides a standard and widely studied update rule, allowing the evaluation to focus on the proposed IPP mechanism and its coupling with the GP-based field model.

Each ant k builds a path step by step. When the ant is at node i , it chooses the next node j from the admissible neighbor set $N_k(i)$ according to the probabilistic rule [29]:

$$p_{ij}^{(k)} = \frac{(\tau_{ij})^\alpha (\eta_{ij})^\beta}{\sum_{h \in N_k(i)} (\tau_{ih})^\alpha (\eta_{ih})^\beta} \quad \text{for } j \in N_k(i). \quad (3)$$

Here, τ_{ij} is the pheromone level on edge (i, j) ; η_{ij} is a heuristic value (e.g., $\eta_{ij} = 1/d_{ij}$ for distance-minimizing objectives); α controls the influence of pheromone and β the influence of the heuristic. Setting $\alpha = 0$ yields a purely heuristic (greedy-biased) policy, whereas $\beta = 0$ produces a purely pheromone driven policy that risks premature stagnation if one pattern is reinforced too early [31].

After all ants complete their tours in an iteration, pheromone is updated in two stages. First, a uniform evaporation removes a fraction $\rho \in (0, 1]$ of every trail, which helps forget stale information and encourages exploration [29]:

$$\tau_{ij} \leftarrow (1 - \rho) \tau_{ij}. \quad (4)$$

Then, each ant k travels along its constructed path S_k stored in its local memory L_k and deposits pheromone on the used edges:

$$\tau_{ij} \leftarrow \tau_{ij} + \Delta\tau_{ij}^{(k)}, \quad \forall (i, j) \in S_k, \quad (5)$$

where the deposited amount $\Delta\tau_{ij}^{(k)}$ is a function of the solution quality $C(S_k)$:

$$\Delta\tau_{ij}^{(k)} = f(C(S_k)). \quad (6)$$

In the classical AS for minimization, a common choice is $f(C(S_k)) = Q/C(S_k)$ with scale $Q > 0$, so that lower-cost tours receive stronger reinforcement [29].

In the present implementation, the tour cost is defined as the total travel distance along the sequence of selected centers,

$$C(S_k) = \sum_{(i,j) \in S_k} d_{ij}, \quad (7)$$

and the pheromone update therefore follows the standard minimization form $\Delta\tau_{ij}^k = Q/C(S_k)$, reinforcing shorter tours.

3.2. Gaussian Process

A GP provides a Bayesian, nonparametric prior over functions, defined by a mean $\mu(x)$ and a covariance (kernel) $k(x, x')$ [32,33]. In environmental mapping without a clear trend, it is standard to set a zero prior mean so that the kernel carries the modeling capacity and inference remains numerically stable [34]. Formally,

$$y = f(x) \sim \mathcal{GP}(\mu(x), k(x, x')), \quad x \in \mathcal{X} \subset \mathbb{R}^d, \quad (8)$$

for any finite set of inputs, the function values are jointly normal with mean from μ and covariance from k [32].

In this work, the GP represents a WQP model over the surface. The posterior mean is the current estimate of the parameter of interest and the posterior variance quantifies the uncertainty of the model, both updated as new measurements arrive.

The kernel encodes smoothness and correlation length. For water-quality fields, this work adopts the Radial Basis Function (RBF) kernel [35],

$$k_{\text{RBF}}(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\ell^2}\right) \quad (9)$$

where the length scale ℓ controls how fast correlations decay [33].

Let $Q = [q_1, \dots, q_N]$ be sampling locations and $m = [m_1, \dots, m_N]^\top$ be the associated measurements. In the noise-free interpolation regime used here, the GP posterior at test input X_* is

$$\mu_* = K_*^T (K + \sigma_o^2)^{-1} y(\mathbf{x}) \quad (10a)$$

$$\sigma_* = K_{**} - K_*^T (K + \sigma_o^2)^{-1} K_* \quad (10b)$$

with $K = k(X, X)$, $K_* = k(X, X_*)$, and $K_{**} = k(X_*, X_*)$ [33]. These expressions provide the posterior mean map used as the current WQP field estimate, and the variance map, which quantifies model uncertainty, for exploration. The σ_o represents the expected level of measurement noise. Incorporating this term helps the GP handle water-quality readings more realistically, as the model accounts for uncertainty in observations during inference [36].

3.3. Action Zone—AquaFeL-PSO

The AquaFeL-PSO [16] is an IPP for fleets of ASVs that combines a multimodal PSO with a GP surrogate of the WQPs. The mission runs in two stages: first, an exploration phase spreads the fleet to gather measurements and fit an initial WQP map; then, an exploitation phase focuses sampling where that map indicates the highest expected values, hotspots.

In the AquaFeL-PSO, an action zone is a compact region of interest defined after the exploration phase. Using the exploration map, the planner flags areas whose estimated WQP levels exceed operational thresholds and centers each zone on local peaks of the estimated contamination map [16]. The procedure is as follows:

1. From the exploration map, grid cells with estimated WQP in the warning band ($34\% \leq \text{level} \leq 65\%$ of the phase maximum value) or in the risk level ($\geq 66\%$) are eligible.

2. The highest peak defines the center of the first zone. A circular sampling area is set so that one ASV (or a small sub-fleet) can cover it effectively. The radius is

$$\text{rad} = \frac{\text{length}}{n_{\text{ASVs}}}, \quad (11)$$

where length is the shortest lakeshore side and n_{ASVs} is the number of vehicles.

3. The zones do not overlap. After placing one zone, its covered cells are removed and the next highest remaining peak seeds the next zone.
4. The number of zones does not exceed the fleet size.

4. ACO-Path: Proposed Informative Path Planner

ACO-Path is a novel IPP designed for autonomous environmental monitoring of water resources using fleets of ASVs. The planner addresses a gap in existing ACO-based path planning approaches: while traditional ACO methods excel at finding collision-free geometric paths, they typically do not integrate informative mapping derived from predictive models of the environment.

The key innovation of ACO-Path lies in its integration of three complementary techniques: (1) ACO, specifically the AS variant for distributed, adaptive path construction with pheromone; (2) GP regression as a surrogate model that provides both a contamination map (mean $\mu(\mathbf{x})$) and an uncertainty map (variance $\sigma(\mathbf{x})$); (3) Dynamic Action Zones, inspired by AquaFeL-PSO [16] but adapted for ACO, that segment the monitoring area into regions of interest based on real-time GP outputs.

By coupling these methods, ACO-Path achieves an exploration–exploitation balance: the GP variance guides exploration (sampling where the model is uncertain), while the GP mean guides exploitation (characterizing contamination hotspots).

4.1. Dynamic Action Zones

Action zones are operational regions defined from the GP maps where ASVs should focus their efforts. Unlike AquaFeL-PSO [16], which generates zones once after an exploration phase, ACO-Path regenerates zones dynamically at every planning cycle to adapt to evolving information.

The incorporation of action zones in ACO-Path serves a dual purpose: (1) to provide a structured candidate set of waypoints for ACO decision-making, and (2) to prevent the ACO algorithm from converging prematurely to local optima.

In classical ACO path planning, artificial ants explore the entire graph uniformly, which can lead to two undesirable behaviors. First, ants may become trapped in locally attractive regions (e.g., nearby nodes with high pheromone concentration) without exploring distant areas that could yield higher information gain [37]. Second, the search space grows combinatorially with the number of nodes, making exhaustive exploration computationally prohibitive for large environments such as Lake Ypacaraí [38].

Action zones address these limitations by dynamically filtering the search space. Instead of evaluating all navigable cells as potential targets, ACO constructs paths among the centers of action zones, a much smaller set of high-value candidate waypoints derived from the GP model.

ACO-Path defines two types of zones based on GP outputs:

1. **Exploration Zones ($\mathcal{A}_{\text{expl}}$).** Regions with high GP variance $\sigma(\mathbf{x})$ where model uncertainty is elevated. The goal is to reduce uncertainty by densifying sampling in unexplored areas.

- *Criterion:* Cells with $\sigma(\mathbf{x}) \geq 0.66 \cdot \sigma_{\max}$ define high-priority exploration regions, and cells with $0.34 \leq \sigma(\mathbf{x}) < 0.66 \cdot \sigma_{\max}$ define medium-priority regions, where $\sigma_{\max} = \max[\sigma(\mathbf{x})]$ is the current maximum variance.
 - *Centers:* Local maxima of $\sigma(\mathbf{x})$.
2. **Exploitation Zones ($\mathcal{A}_{\text{expt}}$).** Regions with high GP mean $\mu(\mathbf{x})$ where contamination is likely. The goal is to characterize hotspots with detailed sampling.
- *Criterion:* Cells with $\mu(\mathbf{x}) \geq 0.66 \cdot \mu_{\max}$ define high-priority exploitation regions (risk level), and cells with $0.34 \leq \mu(\mathbf{x}) < 0.66 \cdot \mu_{\max}$ define medium-priority regions (warning level), where $\mu_{\max} = \max[\mu(\mathbf{x})]$ is the current maximum mean.
 - *Centers:* Local maxima of $\mu(\mathbf{x})$.

It is worth noting that the boundaries of each priority region are defined following the approach described in AquaFeL-PSO [16].

Figure 2 summarizes how action zones are generated and then used by ACO-Path. First, the exploration stage fits a GP over the water resource (Figure 2a), yielding a mean map μ (contamination estimate) and a variance map σ (model uncertainty). Next, the planner thresholds these maps to flag regions of interest and instantiates non-overlapping circular action zones sized for the fleet (Figure 2b). Finally, the center of each action zone is extracted (Figure 2c) and provided to the ACO module at time t .




Figure 2. Cont.




Figure 2. From GP mapping to action zones and ACO inputs. (a) GP outputs: estimated WQP map or contamination map (μ —left) and variance or uncertainty map (σ —right). (b) Action zones (A_{expt} —left and A_{expl} —right) generated from the maps of (a). (c) Centers of the action zones used by ACO at time t as candidate waypoints.

4.2. Path Planner

Figure 3 outlines the workflow of the proposed path planner. The main steps are as follows:




Figure 3. Flowchart of ACO-Path.

(1) Initialization

Set the IPP parameters (number of ants and iterations, α , β , ρ for pheromone dynamics, among others).

(2) Sensing and model update

At their current waypoints, ASVs acquire WQP measurements and update the GP, with the following results: the mean $\mu(\mathbf{x})$ (contamination) and the variance $\sigma(\mathbf{x})$ (uncertainty) maps.

(3) Action zone generation

After updating the GP, action zones are recomputed on both maps. Let $\mathcal{C}_\mu$ be the set of centers from the mean map and $\mathcal{C}_\sigma$ be the set from the uncertainty map. The candidate set $\mathcal{C}$ is defined as

$$\mathcal{C} = \begin{cases} \mathcal{C}_\sigma \text{ (from } \mathcal{A}_{\text{expl}}), & \text{if } \text{dist}_{\text{total}} < \text{dist}_{\text{exploration}}, \\ \mathcal{C}_\sigma \text{ (from } \mathcal{A}_{\text{expl}}) \cup \mathcal{C}_\mu \text{ (from } \mathcal{A}_{\text{expt}}), & \text{otherwise.} \end{cases}$$

Early in the mission, data are scarce, and mean peaks tend to lie near the current positions of the ASVs, which can trigger premature exploitation of a single region. Temporarily restricting candidates to $\mathcal{C}_\sigma$ encourages coverage of poorly sampled areas. Once $\text{dist}_{\text{exploration}}$ is exceeded, both sets are used within the ACO to balance exploration and exploitation.

(4) ACO routing over action zones

Given the admissible set of action zone centers $\mathcal{C}$, the routing problem is defined on the graph induced by these centers (nodes), with edge weights given by the travel distance between centers. Ants then construct candidate routes by sampling transitions according to pheromone information (τ) and a distance-based heuristic (η). After all ants finish, pheromones evaporate and are reinforced according to route quality.

The best-ranked solution yields an assignment of each ASV to a target center and an ordered sequence of centers to be visited. In ACO-Path, however, this plan is executed in a receding-horizon manner. Only the first center of the selected route is used as the next waypoint for the first vehicle. This is illustrated in Figure 4, where the first waypoint of the best route is highlighted with a red box. The assigned center is then removed from the candidate set and the ACO step is run again for the next vehicle to avoid selecting the same target. After each ASV reaches its assigned center and collects new measurements, the GP is updated, action zones and their centers are recomputed, and the ACO routing step is executed again to select the next waypoint under the updated belief.




Figure 4. ACO routing over GP-driven action-zone centers.

(5) Measurement acquisition

During the mission, ASVs move toward their assigned zones and acquire new measurements only when the traveled distance between consecutive samples $dist_{meas}$ reaches the adaptive sampling distance l [35] (Equation (12)).

$$l = \lambda \ell^t, \quad (12)$$

where λ is a scaling factor that controls how strict the sampling spacing is, and ℓ^t is the current GP length scale at time t . This choice provides enough measurements for modeling while avoiding oversampling and unnecessary GP updates.

(6) Stopping criterion

Consistent with Assumption 3, the monitoring mission terminates once the average per-vehicle path length $dist_{total}$ reaches $dist_{max} = 20$ km.

5. Experiments

This section describes the experimental setup, the ground-truth generation procedure, the evaluation metrics, and the parameter settings for both the planner and the simulation environment. All methods are evaluated under the same fleet configuration and mission distance to ensure a fair comparison.

5.1. Setup

The proposed IPP (https://github.com/Natitesis/ACO_Nati.git, accessed on 1 February 2026) was implemented in Python 3.10 using NumPy 1.19.5, Pandas 1.2.0, Matplotlib 3.3.3, scikit-learn 0.23.2, and SciPy 1.5.4. All simulations ran on a workstation with an 11th-gen Intel Core i7@2.80 GHz, 16 GB RAM, and a 64-bit OS.

5.2. Case Study

Lake Ypacaraí was adopted as case study due to its documented eutrophication and contamination pressures from sewage and agricultural runoff [2]. This water resource is the largest lake of Paraguay (about 60 km²), fed mainly by Pirayú (SE) and Yukyry (NW) streams, with the Salado river as its natural outlet [39].

The simulated search area follows the discretization in Figure 1. Since in situ WQP fields are unavailable, the ground-truth contamination map is synthesized with a multimodal Shekel function:

$$f_{\text{Shekel}}(\mathbf{x}) = \sum_{i=1}^M \frac{1}{c_i + \sum_{j=1}^L (x_j - a_{ij})^2}, \quad (13)$$

where $\mathbf{A} \in \mathbb{R}^{M \times L}$ sets peak locations and $\mathbf{C} \in \mathbb{R}^{M \times 1}$ controls peak prominence. For evaluation, 10 ground-truth maps were generated with 2–4 peaks over (x, y) . Peak locations and the entries of $\mathbf{C}$ were drawn at random, producing multimodal contamination fields with varying hotspot prominence. An example ground truth is shown in Figure 5. In these tests, no cross-WQP correlation is modeled.

5.3. Evaluation Metrics

The analysis reports the following: (i) mean squared error over the lake grid,

$$\text{MSE}_{\text{map}}(y(\mathbf{x}), \hat{y}(\mathbf{x})) = \frac{1}{N} \sum_{i=1}^N (y(\mathbf{x}_i) - \hat{y}(\mathbf{x}_i))^2, \quad (14)$$

where y is the ground truth and $\hat{y}$ the GP estimate, (ii) absolute error at contamination peaks,

$$Error_{\text{peak}}(y(\mathbf{x}), \hat{y}(\mathbf{x})) = |y(\mathbf{x}_{\text{peak}}) - \hat{y}(\mathbf{x}_{\text{peak}})|, \quad (15)$$

(iii) the coefficient of determination R^2 ,

$$R^2 = 1 - \frac{\sum_{i=1}^N (y(\mathbf{x}_i) - \hat{y}(\mathbf{x}_i))^2}{\sum_{i=1}^N (y(\mathbf{x}_i) - \bar{y})^2}, \quad (16)$$

where $\bar{y} = \frac{1}{N} \sum_{i=1}^N y(\mathbf{x}_i)$ is the mean of the ground-truth values over the evaluated grid cells.




Figure 5. Example of a ground truth.

Together, these metrics characterize both global map reconstruction and hotspot fidelity, providing a consistent and quantitative basis for comparing planner performance across methods.

5.4. Statistical Significance Analysis

To complement the comparison of the evaluation parameters, statistical tests are used to determine whether the differences observed in MSE_{map} , $Error_{\text{peak}}$, and R^2 reflect consistent performance differences rather than variability between scenarios. First, a one-way analysis of variance (ANOVA) is performed independently for each metric to test the null hypothesis that all planners have the same mean performance ($H_0 : \mu_1 = \mu_2 = \dots = \mu_k$). In one-way ANOVA, evidence against H_0 is summarized by the F statistic, which compares the variance explained by differences between planners with the residual variance within planners [40]:

$$F = \frac{MSB}{MSW} = \frac{\text{variance between groups}}{\text{variance within groups}} \quad (17)$$

All tests are performed with a significance level of $\alpha = 0.05$, and differences are considered statistically significant when $p < 0.05$ [40].

Since ACO-Path and AquaFeL-PSO [16] are evaluated on the same set of reference scenarios, their direct comparison is further evaluated using a paired t -test [41]. For each

scenario i , the difference is defined as $d_i = x_i^{\text{ACO-Path}} - x_i^{\text{AquaFeL-PSO}}$. Next, the paired t statistic is calculated as follows:

$$t = \frac{\bar{d}}{s_d / \sqrt{gt}}, \quad df = gt - 1 \quad (18)$$

where $\bar{d}$ is the mean of the differences, s_d is the standard deviation of the differences, gt is the number of paired scenarios ($gt = 10$), and df represents the degrees of freedom.

5.5. Parameter Settings

The classical AS is used as the baseline ACO variant [29,31]. Parameter ranges follow the literature (Table 1). Following prior studies, α is kept fixed at $\alpha = 1$, and the remaining ACO hyperparameters are taken from reported settings to maintain comparability. Because the heuristic weight β shows larger variability across related works and directly changes the bias of the transition rule, a focused sensitivity check is carried out by evaluating $\beta \in \{2, 5\}$ while keeping the other parameters unchanged. Pheromone evaporation is set to $\rho = 0.7$, and the number of ants is $K = 20$.

Table 1. ACO parameters.

Parameter	Values	References
α	1	[42–48]
β	2	[42,43]
	5	[44–48]
ρ	0.5	[44,47]
	0.7	[49]
	0.9	[42,43,48]
K	20	[42,44]
	30	[43,46,50]
Iteration	30	[42]
	50	[43,44]
Q	1	[43]

The environment and fleet configuration are summarized in Table 2 and follow the settings in [16,35]. The fleet includes $P = 4$ ASVs with maximum speed $v_{\max} = 2$ m/s. The grid resolution is $\Delta = 100 \text{ m} \times 100 \text{ m}$. The sampling distance is l with $\lambda = 0.3$ [16], and the GP length scale is set to $\ell_0 = 10$ as in [35]. Missions terminate when each ASV reaches its distance budget of 20 km, consistent with Assumption 3.

Table 2. Simulation environment parameters.

Parameter	Description	Assigned Value
Δ	Grid resolution	100 m $\times$ 100 m
P	Number of ASVs in the fleet	4
$\mathbf{x}_0$	Initial ASV positions (grid coordinates)	$\begin{pmatrix} 8 & 56 \\ 37 & 16 \\ 78 & 81 \\ 74 & 124 \end{pmatrix}$
$dist_{\max}$	Maximum distance traveled per ASV	20 km
$v_{\max}$	Maximum ASV speed	2 m/s
λ	Scale factor for sampling distance	0.3
ℓ_0	GP length scale	10

5.6. Evaluation of the Proposed System

This section assesses the proposed IPP under varying exploration distances $dist_{exploration}$ and ACO settings. Performance is analyzed on the contamination map and the uncertainty map, and with quantitative metrics reported later.

5.6.1. Results with $\alpha = 1$ and $\beta = 2$

Three exploration distance were tested—5 km, 10 km, and 15 km—following prior practice in [16]. The ACO-Path uses the GP outputs to generate action zones. At each iteration, it selects the next zone center $\mathcal{C}$ using AS. The decision rule is adaptive: while the cumulative traveled distance is below the exploration threshold, only uncertainty-driven centers ($\mathcal{C}_\sigma$) are admissible. Once the threshold is reached, both mean and uncertainty centers ($\mathcal{C}_\mu \cup \mathcal{C}_\sigma$) are considered. Newly acquired WQP measurements update the GP online, reducing local uncertainty and reshaping subsequent action zones.

Figure 6 summarizes the final models after the monitoring task. For each exploration distance $dist_{exploration}$, the subfigure above (orange scale) shows the GP predictive uncertainty map, and the subfigure below shows the GP predictive mean (the estimated WQP or contamination map). In the uncertainty maps, darker regions indicate areas that were insufficiently explored and thus remain poorly informed. The ASV trajectories are also overlaid: each colored polyline denotes the path of a different ASV. The black dot marks the initial position and the red dot marks the final position at the end of the mission. In the mean maps, higher values indicate the more contaminated regions of the water body.

At runtime, the ACO variant AS assigns the next action zone center $\mathcal{C}$ for each ASV using the current ASV position $\mathbf{x}$ and a set of candidate action zone centers derived from either the GP mean or the GP uncertainty (depending on whether the $\mathcal{C}_\sigma$ or $\mathcal{C}_\mu \cup \mathcal{C}_\sigma$ criterion is active). After assigning targets, ASVs move toward their corresponding zones, and ACO-Path accumulates traveled distance per vehicle p . Measurements collected by the on-board sensors are used to update the GP online, which in turn refines both the mean estimate of contamination and the uncertainty field, especially in newly sampled regions.

When the $dist_{exploration}$ is 5 km, Figure 6a, the ASV trajectories cover substantial portions of the map, with emphasis around contamination peaks. ASVs begin by prioritizing high uncertainty regions to improve the GP model. After reaching the 5 km threshold, the criterion switches to $\mathcal{C}_\mu \cup \mathcal{C}_\sigma$, which promotes targeted sampling near suspected hotspots while retaining some exploration. Unvisited dark regions in the uncertainty map indicate areas without samples, resulting in higher residual uncertainty. The GP mean successfully recovers the principal contamination peaks, two prominent peaks, and a third with smaller elevation, showing qualitative similarity to the ground truth (Figure 5).

With a 10 km as $dist_{exploration}$ (Figure 6b), the uncertainty is notably higher than in the 5 km and 15 km cases. The trajectories do not adequately traverse one of the simulated hotspots (approximately in the southwest region), and the model retains large poorly informed areas, which elevates the uncertainty. Consistently, the GP mean captures only two of the three high peaks. The third one is underestimated, yielding lower resemblance to the ground truth.

With a $dist_{exploration}$ equal to 15 km (Figure 6c), ASVs undertake longer trajectories before the switch to $\mathcal{C}_\mu \cup \mathcal{C}_\sigma$, which extends the exploration phase. As a result, the uncertainty is lower than in the 10 km case, and the GP mean recovers all three contamination peaks, closely matching the ground truth. However, because the switch to exploitation is delayed by the longer exploration requirement, there is less early intensification on hotspots compared to the 5 km case.




Figure 6. Contamination map and uncertainty map for exploration distances of 5 km, 10 km, and 15 km with $\beta = 2$.

Table 3 reports the MSE_{map} , the $Error_{\text{peak}}$, and the R^2 for the three exploration distances. The 5 km distance achieves the lowest MSE_{map} and $Error_{\text{peak}}$, indicating the best joint coverage and hotspot exploitation under $\beta = 2$. In contrast, the 10 km distance yields the largest MSE_{map} and $Error_{\text{peak}}$ among the three settings, aligning with the qualitative observation that one hotspot remained insufficiently sampled. The 15 km distance improves over 10 km, particularly in recovering all three peaks and reducing uncertainty, but remains inferior to 5 km in terms of early exploitation due to the delayed switch from C_σ to $C_\mu \cup C_\sigma$.

Table 3. Evaluation metrics for $\alpha = 1$ and $\beta = 2$.

Exploration Distance	MSE_{map}	$Error_{\text{peak}}$	R^2
5 km	0.00144 ± 0.00348	0.19896 ± 0.39400	0.96066 ± 0.09861
10 km	0.01600 ± 0.03525	0.35472 ± 0.29516	0.60847 ± 0.70179
15 km	0.01546 ± 0.03558	0.35376 ± 0.29780	0.62142 ± 0.71184

Regarding R^2 , the highest value is obtained at 5 km with $\beta = 2$ (≈ 0.9767), indicating that, on average, the model explains roughly 97.7% of the variance relative to the ground truth in this setting. For 10 km and 15 km, the 95% confidence intervals are wider and the average R^2 lower, reflecting less consistent performance across simulations. In particular, the drop in R^2 at 10 km is consistent with excessive exploration and insufficient exploitation, which can introduce redundancy and hinder convergence.

In summary, under $\alpha = 1$ and $\beta = 2$, the 5 km distance provides the most favorable balance between exploration and exploitation, yielding the lowest map error and peak error and the highest average R^2 . The 10 km distance over explores relative to exploitation, misses a hotspot, and performs worst across metrics. The 15 km distance extends exploration sufficiently to recover all peaks and reduce the uncertainty compared to 10 km, but its delayed exploitation makes it less effective than 5 km for early hotspot intensification.

5.6.2. Simulation Results with $\alpha = 1$ and $\beta = 5$

This subsection reports results for the same exploration distances $dist_{exploration}$ (5 km, 10 km, and 15 km) using the parameter setting $\alpha = 1$ and $\beta = 5$ (see Table 1). Setting $\beta = 5$ strongly amplifies the influence of the distance heuristic in the AS transition rule, i.e., ants have a much higher preference for moving toward nearby action zone centers.

With $\beta = 5$, AS prioritizes proximity almost exclusively: it tends to choose the nearest action zone centers $\mathcal{C}$. This emphasis reinforces exploitation around already contaminated areas while down weighting exploration of highly uncertain, distant regions. Due to the high value of β , ACO-Path assigns lower probabilities to faraway centers even when those areas are known to be under explored.

Figure 7a shows the uncertainty and mean maps for $dist_{exploration} = 5$ km. The largest uncertainty appears near the boundaries of the simulated lake. Because the required exploration distance is short, ACO-Path switches to the $\mathcal{C}_\mu \cup \mathcal{C}_\sigma$ criterion relatively early. Combined with the high β , this promotes repeated path around nearby sectors. Even so, the predictive mean recovers three contamination peaks, achieving an effective balance between exploration and subsequent exploitation. The mean map exhibits noticeable similarity to the ground truth in Figure 5, aside from localized traces of uncertainty visible in the subfigure Figure 7a.

For $dist_{exploration} = 10$ km, the uncertainty and mean maps are shown in Figure 7b. The overall pattern of uncertainty resembles the 5 km case but with persistently high uncertainty in the central and southeastern regions, indicating insufficient coverage there. Nevertheless, the two principal peaks are detected, and the mean retains qualitative similarity to the ground truth. In terms of areal coverage, however, the 10 km setting with $\beta = 5$ performs worse than 5 km, reflecting the strong bias toward nearby routes and the reduced variety in exploration.

For $dist_{exploration} = 15$ km (Figure 7c), the preference for short moves remains dominant and leads to higher uncertainty compared with the two shorter exploration distances. ASVs continue to select close by paths, frequently reusing previously traveled routes. The contamination map shows two clear peaks and a third weaker peak, yielding a map that is less consistent with the ground truth than in the shorter distance cases. Despite the longer exploration distance, the large β keeps the next target selection tightly focused on nearby action zone centers $\mathcal{C}$. Consequently, the mean estimate does not improve substantially relative to shorter distances, and uncertainty stays elevated over wide areas.

These behaviors indicate that a high β causes the ACO to select the shortest path. While this can intensify exploitation near known hotspots, it limits the ability of the ACO-Path to explore new regions and to build a more complete picture of the environment. The effect is that some peaks are captured well, but valuable information in other lake regions is missed. The metrics in Table 4 confirm this limitation, with moderate R^2 values and higher errors in both the GP map (MSE_{map}) and the detected contamination peaks ($Error_{peak}$).

Examining R^2 with $\beta = 5$ reveals strong performance at $dist_{exploration} = 5$ km. At 15 km, the average R^2 drops to 0.81766 with a wide 95% confidence interval (± 0.44586), indicating substantial variability across simulated runs. This variability is consistent with

the $\beta = 5$ sensitivity to proximity between action zone centers, which reduces exploration diversity and can degrade model fit in parts of the map.




Figure 7. Contamination map and uncertainty map for exploration distances of 5 km, 10 km, and 15 km with $\beta = 5$.

Table 4. Evaluation metrics for $\alpha = 1$ and $\beta = 5$.

Exploration Distance	MSE_{map}	$Error_{peak}$	R^2
5 km	0.00234 ± 0.00800	0.21827 ± 0.36820	0.94094 ± 0.20054
10 km	0.00285 ± 0.00431	0.25406 ± 0.35760	0.91963 ± 0.11475
15 km	0.00803 ± 0.02301	0.28244 ± 0.37209	0.81766 ± 0.44586

5.6.3. Discussion of ACO-Path Evaluation Results

This section reviews the behavior of the proposed IPP with $\beta = 2$ and $\beta = 5$ across the three exploration distances (5 km, 10 km, and 15 km).

With $\beta = 2$, the distance heuristic has a moderate effect. ASVs tend to choose nearby action zone centers $\mathcal{C}$, but they can still move to farther ones when the pheromone level τ_{ij} supports that choice. As a result, paths are relatively uniform. In principle, $\beta = 2$ where an exploration distance $dist_{exploration}$ is equal to 5 km allows broader coverage. However, the actual coverage depends on how diverse the generated action zone centers are. When the exploration distance is 10 km or 15 km, ASVs often follow similar paths and collect repeated measurements. In the mean maps, this leads to only partial reconstruction of the contamination peaks because the GP receives less diverse measurements.

With $\beta = 5$, the distance heuristic dominates. In the selection rule, the nearest action zone center $\mathcal{C}$ is chosen most of the time. The pheromone update reinforces this behavior by

adding pheromone to short and frequently used paths. Paths become local and are visited repeatedly. Distant areas remain unexplored even if their uncertainty is high. In some cases (10 km and 15 km), the planner still identifies nearby peaks, but the reconstruction is local rather than global.

By contrast, with $\beta = 2$ and exploration distance ≥ 5 km, the planner reaches a better balance between exploration and exploitation. Figure 6a shows a consistent reduction in uncertainty over the map and correct estimation of the three main ground truth peaks (Figure 5). This setting also yields the lowest errors for the full map and the peak estimates, together with an R^2 close to 1, supporting its overall performance (see Tables 3 and 4).

5.7. Comparison with Other Planners

This section compares the best-performing configuration of ACO-Path (identified in the previous section) against five path planners from the literature using four ASVs. The evaluation focuses on three metrics: the MSE_{map} , the $Error_{\text{peak}}$, and the R^2 . All values are reported as mean $\pm$ 95% confidence interval over repeated simulations, which also indicates the stability of each planner.

The planners are as follows: Lawnmower, which executes a uniform sweep with parallel tracks and does not adapt to information gathered during the mission; Classical PSO, a PSO scheme inspired by collective motion [51] that explores the space but does not tightly couple waypoint selection to the evolving predictive model; Random Path, which assigns a random heading and directs the vehicles to continue on that course following each measurement; Random Grid, which likewise randomizes headings but restricts them to right angles ($0^\circ, 90^\circ, 180^\circ, 270^\circ$); AquaFeL-PSO [16], which integrates GP feedback to place waypoints in informative zones; and ACO-Path, which combines an inverse-distance heuristic η_{ij} with the probabilistic ACO rule and pheromone deposition. In ACO-Path, the candidate action zone centers are generated from the GP and updated as new measurements arrive.

All planners were evaluated under identical conditions: the same discretized environment, fleet size ($P = 4$), initial ASV positions, sampling model, and maximum distance of the mission (20 km per vehicle). For each planner, the GP was updated with the measurements collected along its trajectories to produce the estimated map $\hat{y}(\mathbf{x})$.

Table 5 reports the quantitative results. AquaFeL-PSO achieves the lowest MSE_{map} and the highest R^2 , indicating the most accurate global reconstruction. ACO-Path achieves the lowest $Error_{\text{peak}}$, showing greater consistency in detecting contamination hotspots, while ranking second on MSE_{map} and R^2 . On other hand, Classical PSO exhibits the largest MSE_{map} and the lowest R^2 , ranking last overall. Lawnmower attains a moderate MSE_{map} and a high R^2 , but its $Error_{\text{peak}}$ is larger than that of ACO-Path and Random Path. Similar to Lawnmower, Random Path has a low MSE_{map} with the second-lowest $Error_{\text{peak}}$, though it remains non-adaptive to the model. Random Grid (right-angle random walk) underperforms in MSE_{map} and R^2 and shows higher variability across runs.

To assess whether these differences are statistically significant, Table 6 reports a one-way ANOVA conducted independently for each metric across all planners ($\alpha = 0.05$, 10 scenarios). Statistically significant differences are observed for MSE_{map} and R^2 ($p < 0.05$ and $F > F_{\text{critical}}$). For $Error_{\text{peak}}$, the ANOVA result is not significant at $\alpha = 0.05$ ($p = 0.05051$ and $F < F_{\text{critical}}$), and therefore, the null hypothesis of equal means cannot be rejected for this metric.

Table 5. Comparison of performance metrics across path planners with four ASVs.

Path Planner	MSE_{map}	$Error_{\text{peak}}$	R^2
Lawnmower	0.00250 ± 0.00659	0.29046 ± 0.40836	0.93806 ± 0.14418
Classical PSO	0.01154 ± 0.02739	0.28896 ± 0.37876	0.65691 ± 0.83469
Random Path	0.00294 ± 0.01159	0.25148 ± 0.38566	0.93115 ± 0.22776
Random Grid	0.00698 ± 0.01188	0.29169 ± 0.45653	0.78334 ± 0.45105
AquaFeL-PSO	0.00077 ± 0.00089	0.26066 ± 0.25235	0.97909 ± 0.02565
ACO-Path	0.00144 ± 0.00348	0.19896 ± 0.39400	0.96066 ± 0.09861

Table 6. Results of the one-way ANOVA test across path planners.

Metric	p -Value	F	F_{critical}
MSE_{map}	0.01190	3.26921	2.38607
$Error_{\text{peak}}$	0.05051	2.37988	2.38607
R^2	0.00919	3.42900	2.38607

In addition, since ACO-Path and AquaFeL-PSO are the two best-performing planners according to Table 5 and are evaluated on the same scenarios, Table 7 reports a paired t -test to directly assess whether the performance gap between them is statistically significant. No statistically significant differences are found for any metric (two-sided $p > 0.05$).

Table 7. Paired t -test between ACO-Path and AquaFeL-PSO.

Metric	t	df	p -Value
MSE_{map}	1.2458	9	0.244
$Error_{\text{peak}}$	−1.0971	9	0.301
R^2	−1.1898	9	0.261

In Figure 8, each figure shares the same layout. The top panel shows the GP uncertainty over the water resource (darker orange means higher uncertainty) with ASV trajectories overlaid. The black dots mark initial positions and red markers indicate final positions of the ASVs. The bottom panel shows the GP mean or the estimated contamination map.

In the lawnmower path planner (Figure 8a), the ASVs travels in parallel to obtain uniform coverage. In the uncertainty graph, the long, straight paths reduce the uncertainty in the interior but leave higher uncertainty along corners. As a consequences, in the mean graph, the field is smooth and coherent, but the hotspots are diluted: peaks are present yet underestimated and slightly displaced with respect to the ground truth. This behavior explains the larger peak error.

The classical PSO (Figure 8b) updates candidate waypoints using social and cognitive terms, which tends to concentrate sampling near locally attractive regions without an explicit rule to avoid overlap. In the uncertainty graph, paths cluster around a few local optima, leaving other areas with high uncertainty. The mean map recovers parts of the main peaks but misses structure where the fleet did not sample.

In Random Path (Figure 8d), allowing arbitrary course changes produces more dif-fused, winding coverage. The ASVs visit several sectors and reduce σ across much of the interior, though some high-uncertainty regions remain without a targeting rule. In the mean map, the main hotspots are reconstructed with reasonable location and contrast. This is reflected in a competitive R^2 (around 0.93 in our runs) and a peak error lower than Lawnmower.




Figure 8. Results of all path planners: uncertainty map (**top**) and estimated WQPs map (**bottom**).

In Random Grid (Figure 8c), constraining the directions to $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ yields straight trajectories, similar to blocks. This pattern often repeats segments and leaves gaps, especially near the boundaries. In the uncertainty map, these gaps persist as areas of high σ . On the mean map, the hotspots appear diluted or slightly shifted, consistent with a higher MSE_{map} , a lower R^2 than in Random Path, and a higher mean maximum error due to missing or misdected peaks.

In the AquaFeL-PSO (Figure 8e), the PSO objective is shaped by the GP so that sampling increases information while still tracking high mean values. In the uncertainty graph, paths spread more effectively and reduce uncertainty more uniformly, including near boundaries. In the contamination map, the reconstruction closely matches the ground truth (Figure 5): the two dominant hotspots are well located and the weaker eastern peak appears clearly. This balance explains the low map-level error and stable performance across runs.

ACO-Path (Figure 8f) selects the next action zone center using an AS rule that combines pheromone (memory of successful choices) with the distance-based heuristic. The policy prioritizes uncertainty to explore first. Once an ASV reaches the exploration distance, it switches to $C_\mu \cup C_\sigma$ criterion to generate the action zones. In the uncertainty map, vehicles first visit high-uncertainty regions and later concentrate sampling in informative areas with high mean values, leaving some peripheral areas with moderate uncertainty values. In the mean map, the main hotspots are estimated with good location and contrast: the southeast and central west peaks stand out and the weaker eastern feature is visible. This targeted behavior yields low peak error and high R^2 .

In addition to reconstruction accuracy, the smoothness of the resulting trajectories was also analyzed, since frequent sharp heading changes can increase the burden on the low-level controller and make practical execution more demanding in terms of control effort and energy consumption. In practice, abrupt turns translate into large and fast-varying turning commands, which may drive the commanded inputs closer to actuator limits (magnitude, rate constraints), thereby reducing tracking authority and increasing the required control action [52,53].

This effect was quantified by computing the absolute heading change between consecutive trajectory segments, $|\Delta\alpha|$, as well as by counting the number of turns above a threshold (here, $|\Delta\alpha| \geq 45^\circ$) along the mission distance. Figure 9 reports the evolution of the turning demand versus distance as mean $\pm$ 95% confidence interval over the simulations, and Table 8 summarizes the cumulative number of sharp turns at the end of the analyzed segment. The values are reported as mean $\pm$ 95% confidence interval over 10 mission tasks.

Table 8. Cumulative number of sharp turns ($|\Delta\alpha| \geq 45^\circ$).

Path Planner	Turns ($ \Delta\alpha \geq 45^\circ$)
AquaFeL-PSO	145.10 $\pm$ 13.82
ACO-Path	100.22 $\pm$ 31.67

From an execution-oriented perspective, reducing sharp turns is relevant not only for control effort but also for energy usage. In mobile robots, power or energy demand is strongly dependent on the motion regime: straight segments typically yield the lowest consumption, whereas trajectories with pronounced rotational components require higher actuation due to increased lateral slip and frictional losses during maneuvering [54,55].




Figure 9. Evolution of turning effort with mission distance.

Although AquaFeL-PSO attains the best global reconstruction (lowest MSE_{map} and highest R^2), it consistently produces trajectories with a higher incidence of sharp turns than ACO-Path. At the end of the segment, AquaFeL-PSO accumulates 145.10 ± 13.82 turns with $|\Delta\alpha| \geq 45^\circ$, whereas ACO-Path accumulates 100.22 ± 31.67 , which corresponds to approximately 44.9% fewer sharp turns for ACO-Path. This difference is also observed earlier in the mission. Overall, these results indicate that ACO-Path yields smoother trajectories (fewer abrupt heading changes) while remaining competitive in map-level metrics, suggesting a more favorable trade-off when practical execution constraints are considered.

A paired statistical test was conducted to verify whether the difference in sharp turns between the two planners is significant under matched scenarios. Using the common runs that reached the 20 km of traveled distance ($n = 9$, $df = 8$), a paired two-sided t -test indicates a statistically significant reduction in the number of sharp turns for ACO-Path compared to AquaFeL-PSO ($t = -4.6317$, $p = 0.00168$).

In summary, across all planners, a clear trade-off appears between broad spatial coverage and hotspot fidelity. The random baselines provide a useful reference: allowing continuous headings (Random Path) generally reconstructs the field better than restricting motion to right angles (Random Grid), although both remain behind planners that explicitly exploit GP feedback. Lawnmower enforces systematic coverage, which reduces uncertainty over large areas but also tends to smooth and attenuate peaks. Classical PSO can concentrate sampling around locally attractive regions, but without an explicit mechanism to avoid overlap, it may leave parts of the domain insufficiently explored. Focusing on the two best methods, the paired comparison in Table 7 shows that the differences between ACO-Path and AquaFeL-PSO in MSE_{map} , $Error_{\text{peak}}$, and R^2 are not statistically significant (two-sided $p > 0.05$), even though their mean values differ. Overall, AquaFeL-PSO delivers the most accurate global reconstruction, whereas ACO-Path achieves the lowest hotspot error and, as the turning analysis shows, does so with smoother trajectories.

The evaluation relies on ground truth fields and an idealized execution model (e.g., noise-free measurements and perfect trajectory tracking). This setting was chosen to focus on the planning component and to separate it from sensing and control effects. Assessing robustness under measurement noise, localization errors, and communication constraints will be addressed in future work.

6. Discussion of the Results

Across the tested configurations, the best performance is obtained with $\alpha = 1$, $\beta = 2$, and an exploration distance of 5 km (Table 3). This setting reduces uncertainty early and then concentrates samples on high-mean regions, which aligns with the qualitative maps in Figure 6.

When the exploration distance is set to 10 km with $\beta = 2$, the proposed IPP remains too long in exploration. One hotspot is undersampled and large regions keep high uncertainty. This aligns with the increase in MSE_{map} and $Error_{\text{peak}}$, as well as the drop in R^2 . At 15 km, coverage improves compared to 10 km and all peaks appear, but the late switch to exploitation leaves less distance to exploit them, so results remain weaker than at 5 km.

With $\beta = 5$, the distance term dominates the choice of targets. Ants almost always prefer nearby centers. This sharpens local sampling around known areas but limits coverage, especially at 10–15 km (Figure 7). Uncertainty remains high in parts of the map, and the confidence intervals in Table 4 widen. The 5 km setup is again the most solid in this group, but it does not surpass the best $\beta = 2$ setting. In summary, too much weight on distance leads to repeated short moves and missed regions.

The multi-planner comparison puts these findings in context (Table 5, Figure 8). AquaFeL-PSO, which also uses the GP to guide sampling, achieves the lowest MSE_{map} and the highest R^2 . Therefore, the most accurate global map. ACO-Path obtains the lowest $Error_{\text{peak}}$ and ranks second on MSE_{map} and R^2 , which indicates reliable hotspot detection with competitive global fit. Lawnmower ensures coverage but tends to smooth peaks. Random Path improves peak error relative to Lawnmower with similar MSE_{map} . Random Grid, constrained to right angles, repeats segments and leaves gaps, which is reflected by its higher MSE_{map} and R^2 . Classical PSO explores the space, but without a tight coupling to the GP during waypoint selection, it ends up with the highest MSE_{map} and the lowest R^2 on average.

The interpretation mentioned above is supported by inferential analysis. The one-way ANOVA among all planners (Table 6) detects statistically significant differences for the map-level metrics MSE_{map} and R^2 ($p < 0.05$), reinforcing the conclusion that the ranking observed in Table 5 reflects systematic differences in performance rather than fluctuations between scenarios. However, for $Error_{\text{peak}}$, the ANOVA result does not meet the criterion $\alpha = 0.05$ ($p = 0.05051$), so it does not provide sufficient evidence to claim that there is a clear separation between all planners in terms of hotspot error in this test.

Since ACO-Path and AquaFeL-PSO are evaluated on identical reference scenarios, a paired analysis provides a more appropriate direct comparison. As indicated in Table 7, the paired t tests do not indicate statistically significant differences between these two planners for MSE_{map} , $Error_{\text{peak}}$, or R^2 (two-tailed $p > 0.05$). Therefore, although AquaFeL-PSO achieves the most robust global reconstruction and ACO-Path achieves the lowest average error in hotspots, the current set of scenarios does not support a conclusive claim of superiority of one method over the other in these accuracy metrics.

Beyond the accuracy of the reconstruction, the analysis of turns adds an execution-focused dimension to the comparison between planners. Using the absolute change in heading between consecutive segments together with the count of sharp turns, it has been observed that AquaFeL-PSO generates more aggressive maneuvers than ACO-Path for comparable mission distances. At 20 km, the average number of turns exceeding 45° is higher for AquaFeL-PSO, while ACO-Path produces smoother heading profiles. This difference is relevant from an operational point of view: frequent sharp turns imply larger and more rapidly varying turn commands, which reduces the tracking margin as the demanded inputs approach the magnitude and/or speed constraints of the actuator [52,53]. Furthermore, trajectories with pronounced rotational components tend to require more energy

than nearly rectilinear motions, as turns increase actuation requirements and associated losses [54,55]. A paired two-sided t -test on the runs that reach 20 km of traveled distance confirms that this reduction is statistically significant ($t = -4.6317$, $df = 8$, $p = 0.00168$).

These results were obtained under controlled assumptions (noise-free measurements, synchronized motion, centralized coordination, and synthetic ground truths adapted to the lake grid). Under these controlled assumptions, the results consistently show that coupling Ant System with GP-driven action zones and an explore then exploit policy improves hotspot reconstruction while maintaining a competitive global fit against the selected baselines.

7. Conclusions

This work presented an IPP based on ACO for monitoring WQPs with a fleet of ASVs. The proposed planner, ACO-Path, couples a GP surrogate with Ant System so that decisions are guided not only by geometric proximity but also by the GP mean (contamination estimate) and variance (model uncertainty). From these maps, action zone centers are generated online, and the policy follows an explore then exploit policy: it prioritizes uncertainty ($\mathcal{C}_\sigma$) until an ASV reaches an exploration distance $dist_{exploration}$, then switches to a combined $\mathcal{C}_\mu \cup \mathcal{C}_\sigma$ criterion to concentrate sampling on informative regions.

In simulations inspired by Lake Ypacaraí, ACO-Path was assessed under three exploration distances (5, 10, and 15 km) and two heuristic weights ($\beta \in \{2, 5\}$). The configuration $\alpha = 1$, $\beta = 2$, with a 5 km exploration distance yielded the most consistent performance, achieving the best trade-off across map-level reconstruction and hotspot fidelity. Increasing the exploration threshold to 10–15 km under $\beta = 2$ delayed the exploitation stage, which reduced refinement near hotspots within the available mission distance. For $\beta = 5$, the inverse-distance heuristic dominated target selection, strengthening local sampling but limiting coverage and leaving higher uncertainty in less visited regions.

A comparative study against Lawnmower, Classical PSO, Random Path, Random Grid, and AquaFeL-PSO further highlighted these trends. AquaFeL-PSO achieved the strongest global reconstruction on average, while ACO-Path achieved the lowest mean error at contamination peaks and remained competitive in map-level metrics. Statistical analysis supports these conclusions: differences between planners are significant for global reconstruction metrics, while hotspot error shows weaker separation under the same experiments. When directly comparing the two best planners in equal scenarios, ACO-Path and AquaFeL-PSO, the paired analysis does not provide sufficient evidence to claim a consistent advantage of one method over the other in reconstruction metrics, suggesting that both methods are competitive under the evaluated conditions.

Beyond reconstruction accuracy, the analysis of turns adds an execution-oriented perspective. At comparable mission distances, ACO-Path consistently produces smoother trajectories than AquaFeL-PSO, with fewer abrupt course changes. This feature is relevant for practical deployments, where frequent abrupt turns can increase low-level control effort and energy consumption.

If the method is taken to deployment, the IPP should be integrated with a local path planning layer in charge of obstacle avoidance and inter-ASV collision avoidance during execution, so that safety is handled at control level while the IPP provides the nominal path.

As future work, the authors propose to (i) enable online adaptation of key ACO parameters; (ii) extend ACO-Path to multi-objective planning that jointly considers field accuracy, energy consumption, and travel time; (iii) adopt spatiotemporal GP so that the planner adapts to the changing dynamics of a water resource rather than a static field; (iv) conduct a more extensive sensitivity analysis of the ACO hyperparameters and evaluate the ACO-Path with additional IPP planners, including more recent ACO variants

and learning-based approaches, under a unified benchmarking protocol; (v) consider more realistic sensing and navigation conditions (e.g., measurement noise and localization uncertainty) and additional deployment constraints to further assess robustness in real-world monitoring missions.

Author Contributions: Conceptualization, M.J.T.K. and N.B.; methodology, M.J.T.K. and D.G.R.; software, M.J.T.K. and N.B.; validation, M.J.T.K., M.A. and D.G.R.; resources, M.A. and D.G.R.; writing—original draft preparation, M.J.T.K. and N.B.; writing—review and editing, M.J.T.K., M.A. and D.G.R. All authors have read and agreed to the published version of the manuscript.

Funding: M.A. acknowledges financial support from the National Science and Technology Council (CONACYT), Paraguay, through project PINV01-819, “Técnicas de Planificación de Rutas en enjambre de Vehículos Autónomos de Superficie a partir de parámetros fisicoquímicos de un cuerpo de agua”. This work was funded by the the Spanish Ministry of Science and the European Union with “Next Generation EU / PRTR” funds under Grant PID2024-158365OB-C21.

Data Availability Statement: The data that supports the results can be obtained using the main script found in the Github repository https://github.com/Natitesis/ACO_Nati.git, accessed on 1 February 2026.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Ogidi, O.I.; Akpan, U.M. Aquatic biodiversity loss: Impacts of pollution and anthropogenic activities and strategies for conservation. In *Biodiversity in Africa: Potentials, Threats and Conservation*; Springer Nature: Singapore, 2022; pp. 421–448.
- López Moreira M., G.A.; Hinegk, L.; Salvatore, A.; Zolezzi, G.; Hölker, F.; Monte Domecq S., R.A.; Bocci, M.; Carrer, S.; De Nat, L.; Escribá, J.; et al. Eutrophication, research and management history of the shallow Ypacaraí Lake (Paraguay). *Sustainability* **2018**, *10*, 2426. [CrossRef]
- Wang, M.; Bodirsky, B.L.; Rijneveld, R.; Beier, F.; Bak, M.P.; Batool, M.; Droppers, B.; Popp, A.; van Vliet, M.T.; Stokal, M. A triple increase in global river basins with water scarcity due to future pollution. *Nat. Commun.* **2024**, *15*, 880. [CrossRef]
- Zhao, S.; Kvale, K.F.; Zhu, L.; Zettler, E.R.; Egger, M.; Mincer, T.J.; Amaral-Zettler, L.A.; Lebreton, L.; Niemann, H.; Nakajima, R.; et al. The distribution of subsurface microplastics in the ocean. *Nature* **2025**, *641*, 51–61. [CrossRef]
- Jiang, L.Q.; Dunne, J.; Carter, B.R.; Tjiputra, J.F.; Terhaar, J.; Sharp, J.D.; Olsen, A.; Alin, S.; Bakker, D.C.; Feely, R.A.; et al. Global surface ocean acidification indicators from 1750 to 2100. *J. Adv. Model. Earth Syst.* **2023**, *15*, e2022MS003563. [CrossRef]
- Kroeker, K.J.; Kordas, R.L.; Crim, R.; Hendriks, I.E.; Ramajo, L.; Singh, G.S.; Duarte, C.M.; Gattuso, J.P. Impacts of ocean acidification on marine organisms: Quantifying sensitivities and interaction with warming. *Glob. Change Biol.* **2013**, *19*, 1884–1896. [CrossRef]
- Bednaršek, N.; Feely, R.A.; Reum, J.C.; Peterson, B.; Menkel, J.; Alin, S.R.; Hales, B. Limacina helicina shell dissolution as an indicator of declining habitat suitability owing to ocean acidification in the California Current Ecosystem. *Proc. R. Soc. B Biol. Sci.* **2014**, *281*, 20140123. [CrossRef] [PubMed]
- Bindoff, N.; Cheung, W.; Arístegui, J.; Guinder, V.; Hallberg, R.; Hilmi, N.; Williamson, P. Chapter 5: Changing ocean, marine ecosystems, and dependent communities. In *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*; Cambridge University Press: Cambridge, UK; New York, NY, USA, 2019.
- Feng, L.; Wang, Y.; Hou, X.; Qin, B.; Kutser, T.; Qu, F.; Chen, N.; Paerl, H.W.; Zheng, C. Harmful algal blooms in inland waters. *Nat. Rev. Earth Environ.* **2024**, *5*, 631–644. [CrossRef]
- Song, T.; Zhang, H.; Xu, Y.; Dai, X.; Fan, F.; Wang, Y.; Liu, G. Cyanobacterial blooms in Lake Taihu: Temporal trends and potential drivers. *Sci. Total Environ.* **2024**, *942*, 173684. [CrossRef] [PubMed]
- Jiang, M.; Cannizzaro, J.; McFarland, M.N.; Wistort, Z.; Beckler, J.S.; Hu, C.; Moore, T. Modeling water quality and cyanobacteria blooms in Lake Okeechobee: II. Dynamics of diurnal cycles and impacts of cyanobacteria diel vertical migration. *Ecol. Model.* **2025**, *505*, 111107. [CrossRef]
- Chaffin, J.D.; Mishra, S.; Kane, D.D.; Bade, D.L.; Stanislawczyk, K.; Slodysko, K.N.; Jones, K.W.; Parker, E.M.; Fox, E.L. Cyanobacterial blooms in the central basin of Lake Erie: Potentials for cyanotoxins and environmental drivers. *J. Great Lakes Res.* **2019**, *45*, 277–289. [CrossRef]
- United Nations. *Transforming Our World: The 2030 Agenda for Sustainable Development*; United Nations: New York, NY, USA, 2016.
- Luis, S.Y.; Reina, D.G.; Marín, S.L.T. A Multiagent Deep Reinforcement Learning Approach for Path Planning in Autonomous Surface Vehicles: The Ypacara-C-Lake Patrolling Case. *IEEE Access* **2021**, *9*, 17084–17099. [CrossRef]

15. Arzamendia, M.; Gutierrez, D.; Toral, S.; Gregor, D.; Asimakopoulou, E.; Bessis, N. Intelligent online learning strategy for an autonomous surface vehicle in lake environments using evolutionary computation. *IEEE Intell. Transp. Syst. Mag.* **2019**, *11*, 110–125. [\[CrossRef\]](#)
16. Ten Kathen, M.J.; Peralta, F.; Johnson, P.; Flores, I.J.; Reina, D.G. AquaFeL-PSO: An informative path planning for water resources monitoring using autonomous surface vehicles based on multi-modal PSO and federated learning. *Ocean Eng.* **2024**, *311*, 118787. [\[CrossRef\]](#)
17. Popović, M.; Vidal-Calleja, T.; Hitz, G.; Chung, J.J.; Sa, I.; Siegart, R.; Nieto, J. An informative path planning framework for UAV-based terrain monitoring. *Auton. Robot.* **2020**, *44*, 889–911. [\[CrossRef\]](#)
18. Rückin, J.; Jin, L.; Popović, M. Adaptive informative path planning using deep reinforcement learning for uav-based active sensing. In Proceedings of the 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23–27 May 2022; pp. 4473–4479.
19. Rückin, J.; Magistri, F.; Stachniss, C.; Popović, M. An informative path planning framework for active learning in uav-based semantic mapping. *IEEE Trans. Robot.* **2023**, *39*, 4279–4296. [\[CrossRef\]](#)
20. Vashisth, A.; Rückin, J.; Magistri, F.; Stachniss, C.; Popović, M. Deep reinforcement learning with dynamic graphs for adaptive informative path planning. *IEEE Robot. Autom. Lett.* **2024**, *9*, 7747–7754. [\[CrossRef\]](#)
21. Popović, M.; Ott, J.; Rückin, J.; Kochenderfer, M.J. Learning-based methods for adaptive informative path planning. *Robot. Auton. Syst.* **2024**, *179*, 104727. [\[CrossRef\]](#)
22. Yanes Luis, S.; Shutin, D.; Marchal Gómez, J.; Gutiérrez Reina, D.; Toral Marín, S. Deep reinforcement multiagent learning framework for information gathering with local Gaussian processes for water monitoring. *Adv. Intell. Syst.* **2024**, *6*, 2300850. [\[CrossRef\]](#)
23. Diop, D.S.; Luis, S.Y.; Esteve, M.P.; Marín, S.L.T.; Reina, D.G. Decoupling patrolling tasks for water quality monitoring: A multi-agent deep reinforcement learning approach. *IEEE Access* **2024**, *12*, 75559–75576. [\[CrossRef\]](#)
24. Alabdabari, A.A.; Abed, I.A. New robot path planning optimization using hybrid GWO-PSO algorithm. *Bull. Electr. Eng. Inform.* **2022**, *11*, 1289–1296. [\[CrossRef\]](#)
25. Heng, H.; Rahiman, W. ACO-GA-Based Optimization to Enhance Global Path Planning for Autonomous Navigation in Grid Environments. *IEEE Trans. Evol. Comput.* **2025**, *30*, 226–240. [\[CrossRef\]](#)
26. Liu, Y.; Guo, S.; Tang, S.; Song, J.; Zhang, J. Path Planning for Robots Based on Adaptive Dual-Layer Ant Colony Optimization Algorithm and Adaptive Dynamic Window Approach. *IEEE Sens. J.* **2025**, *25*, 19694–19708. [\[CrossRef\]](#)
27. Peng, T.J.; Chen, P.J.; Su, T.L.; Lin, C.C. Automated spray path planning based on Bayesian optimization and ant colony optimization. *Int. J. Adv. Manuf. Technol.* **2025**, *139*, 5491–5509. [\[CrossRef\]](#)
28. Xiong, C.; Chen, D.; Lu, D.; Zeng, Z.; Lian, L. Path planning of multiple autonomous marine vehicles for adaptive sampling using Voronoi-based ant colony optimization. *Robot. Auton. Syst.* **2019**, *115*, 90–103. [\[CrossRef\]](#)
29. Dorigo, M.; Maniezzo, V.; Colnari, A. Ant system: Optimization by a colony of cooperating agents. *IEEE Trans. Syst. Man, Cybern. Part B (Cybern.)* **1996**, *26*, 29–41. [\[CrossRef\]](#) [\[PubMed\]](#)
30. Dorigo, M.; Birattari, M.; Stutzle, T. Ant colony optimization. *IEEE Comput. Intell. Mag.* **2007**, *1*, 28–39. [\[CrossRef\]](#)
31. Alonso, S.; Cordón, O.; Fernández de Viana, I.; Herrera, F. La Metaheurística de Optimización Basada en Colonias de Hormigas: Modelos y Nuevos Enfoques. In *Optimización Inteligente: Técnicas de Inteligencia Computacional Para Optimización*; Servicio de Publicaciones y Divulgación Científica de la Universidad de Málaga: Málaga, Spain, 2004; pp. 261–314.
32. Rasmussen, C.E. Gaussian processes in machine learning. In *Proceedings of the Summer School on Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2003; pp. 63–71.
33. Williams, C.K.; Rasmussen, C.E. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
34. Schulz, E.; Speekenbrink, M.; Krause, A. A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *J. Math. Psychol.* **2018**, *85*, 1–16. [\[CrossRef\]](#)
35. Peralta, F.; Reina, D.G.; Toral, S.; Arzamendia, M.; Gregor, D. A bayesian optimization approach for multi-function estimation for environmental monitoring using an autonomous surface vehicle: Ypacarai lake case study. *Electronics* **2021**, *10*, 963. [\[CrossRef\]](#)
36. Peralta, F.; Reina, D.G.; Toral, S. Water quality online modeling using multi-objective and multi-agent Bayesian Optimization with region partitioning. *Mechatronics* **2023**, *91*, 102953. [\[CrossRef\]](#)
37. Dorigo, M.; Stutzle, T. *Ant Colony Optimization*; MIT Press: Cambridge, MA, USA, 2004.
38. Dorigo, M.; Blum, C. Ant colony optimization theory: A survey. *Theor. Comput. Sci.* **2005**, *344*, 243–278. [\[CrossRef\]](#)
39. Ministerio del Ambiente y Desarrollo Sostenible (MADES). *Análisis Trimestral: Informe Laboratorial de la Calidad de Agua del Ypacarai*; Ministerio del Ambiente y Desarrollo Sostenible (MADES): Asunción, Paraguay, 2022.
40. Larson, M.G. Analysis of variance. *Circulation* **2008**, *117*, 115–121. [\[CrossRef\]](#)
41. Mishra, P.; Pandey, C.M.; Singh, U.; Keshri, A.; Sabaretnam, M. Selection of appropriate statistical methods for data analysis. *Ann. Card. Anaesth.* **2019**, *22*, 297–301. [\[CrossRef\]](#) [\[PubMed\]](#)

42. Fu, C.; Wang, Y.; Gu, Y.; Ma, M.; Xue, T. Routing optimization of high-level orderpickers in a rectangular warehouse. In Proceedings of the 2011 International Conference on Consumer Electronics, Communications and Networks (CECNet), Xianning, China, 16–18 April 2011; pp. 4388–4391.
43. Wang, Y.; Fu, C.; Ma, M.; Wang, L. Routing optimization of high-level order picking truck based on swarm intelligent algorithm. *Appl. Mech. Mater.* **2012**, *101–102*, 414–417. [[CrossRef](#)]
44. Chen, T.L.; Cheng, C.Y.; Chen, Y.Y.; Chan, L.K. An efficient hybrid algorithm for integrated order batching, sequencing and routing problem. *Int. J. Prod. Econ.* **2015**, *159*, 158–167. [[CrossRef](#)]
45. Chen, F.; Xu, G.; Wei, Y. An integrated metaheuristic routing method for multiple-block warehouses with ultranarrow aisles and access restriction. *Complexity* **2019**, *2019*, 1280285. [[CrossRef](#)]
46. Jin, H.; Wang, W.; Cai, M.; Wang, G.; Yun, C. Ant colony optimization model with characterization-based speed and multi-driver for the refilling system in hospital. *Adv. Mech. Eng.* **2017**, *9*, 1–18. [[CrossRef](#)]
47. Li, J.; Huang, R.; Dai, J.B. Joint optimisation of order batching and picker routing in the online retailer’s warehouse in China. *Int. J. Prod. Res.* **2017**, *55*, 447–461. [[CrossRef](#)]
48. De Santis, R.; Montanari, R.; Vignali, G.; Bottani, E. An adapted ant colony optimization algorithm for the minimization of the travel distance of pickers in manual warehouses. *Eur. J. Oper. Res.* **2018**, *267*, 120–137. [[CrossRef](#)]
49. Cano, J.A.; Campo, E.A.; Correa-Espinal, A.A.; Gómez-Montoya, R.A. Optimización por colonia de hormigas para el ruteo de la preparación de pedidos en almacenes de múltiples bloques. *Inf. Tecnol.* **2021**, *32*, 121–130. [[CrossRef](#)]
50. Chen, F.; Wang, H.; Xie, Y.; Qi, C. An ACO-based online routing method for multiple order pickers with congestion consideration in warehouse. *J. Intell. Manuf.* **2016**, *27*, 389–408. [[CrossRef](#)]
51. Eberhart, R.; Kennedy, J. Particle swarm optimization. In Proceedings of the IEEE International Conference on Neural Networks, Perth, Australia, 27 November–1 December 1995; Volume 4, pp. 1942–1948.
52. Falcone, P.; Borrelli, F.; Asgari, J.; Tseng, H.E.; Hrovat, D. Predictive active steering control for autonomous vehicle systems. *IEEE Trans. Control Syst. Technol.* **2007**, *15*, 566–580. [[CrossRef](#)]
53. Liu, C.; Sun, T.; Wang, X. Constrained LOS guidance for path following of underactuated marine vehicle with input saturation. *Ocean Eng.* **2024**, *307*, 118217. [[CrossRef](#)]
54. Effati, M.; Skonieczny, K. Systematic solution for optimally energy-efficient turning radius for wheeled skid-steer rovers. *Robot. Auton. Syst.* **2023**, *159*, 104306. [[CrossRef](#)]
55. Graba, M.; Amamou, A.; Kelouwani, S.; Allani, B.; Zeghmi, L.; Agbossou, K.; Mohammadpour, M. Toward safer and energy efficient global trajectory planning of self-guided vehicles for material handling system in dynamic environment. *IEEE Access* **2023**, *11*, 30753–30767. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.