

Energy-Aware Informative Path Planning for Heterogeneous Multi-Robot Systems

Aiman Munir

Ayan Dutta

Ramviyas Parasuraman

Abstract—Effective energy management is essential for maximizing information gathering tasks with networked mobile robots, particularly for large-scale, energy-intensive tasks such as agricultural monitoring and wildfire mapping. This paper presents a novel framework that integrates robots’ energy profiles with confidence bounds of their assigned regions to optimize sampling targets. Designed for persistent, long-term deployments, the framework employs Gaussian Process Regression (GPR) to maximize data acquisition and accurately reconstruct unknown spatial distributions (e.g., algae outbreaks or humidity maps). The method enables seamless transitions among exploration (mapping uncertain regions at high energy), exploitation (refining maps at moderate energy levels), and recharging (navigating to charging stations at low energy), thereby achieving energy-balanced informative path planning. Experiments demonstrate the effectiveness of the approach against state-of-the-art methods in generating energy-efficient and distinct paths for heterogeneous robots, delivering up to 32% energy savings while maintaining high reconstruction accuracy. Hardware experiments closely matched the performance in simulation.

I. INTRODUCTION

Robotic data collection systems play a critical role in environmental monitoring, precision agriculture, search and rescue, and surveillance [1], [2]. A key component of these systems is Informative Path Planning (IPP), which directs robots to maximize data utility by reconstructing unknown spatio-temporal phenomena [3]. Most IPP methods assume unlimited energy and focus solely on information gain [4], an assumption unrealistic for real-world deployments such as wildfire monitoring or large-scale agriculture, where energy is both limited and mission-critical [5], [6]. In such settings, poor energy management can cause premature mission failure and degraded data quality. Thus, effective IPP must balance information gain with energy efficiency in multi-robot systems [7]. Multi-robot Informative Path Planning (MIPP) extends IPP by enabling teams to collaboratively explore environments for greater information gain, with heterogeneous systems offering benefits from diverse energy, mobility, and sensing capabilities [8]. Yet, most methods assume robot uniformity, causing inefficiencies in heterogeneous settings [3], [6]. Realizing the full potential of such teams requires energy-aware coordination and joint planning, particularly in large-scale, energy-constrained missions.

In heterogeneous systems, efficient energy management is crucial, as each robot’s energy levels and capabilities must be considered to maximize operational time and overall




Fig. 1. Overview of the proposed energy-aware IPP, assigning role to robots based on their current energy budget.

system performance [9], [10]. For example, tasks that require significant energy, such as long-range data collection, should be allocated to robots with greater energy reserves and vice versa, to avoid premature depletion and ensure the long-term survival of the robots in the mission. Furthermore, in a persistent IPP application, such energy conservation is paramount, while autonomous periodic recharging as well as redeployment are equally important. Some prior works have addressed energy constraints in IPP by considering available energy or static depletion rates [10], [11]. However, many real-world scenarios involve dynamic energy factors like terrain, sensor use, and speed [9], [5] that are crucial to consider for deployment.

To address these challenges, we propose a novel IPP methodology for heterogeneous multi-robot systems that operate under energy constraints. Our approach adapts each robot’s future sampling location based on its current energy availability and confidence (or uncertainty) in its information predictions (see Fig. 1 for an illustration). Extensive experimental data demonstrate that the proposed technique maximizes both information gain and energy efficiency, allowing more robots to remain operational by the end of the mission. This study advances the state-of-the-art by presenting a novel energy-aware multi-robot informative path-planning (EA-MIPP) algorithm with the following contributions.

- Incorporate the unique energy characteristics of each robot while adapting the path planning for information sampling. Specifically, our approach leverages exploration-exploitation tradeoffs in IPP and smooth state transitions decided by individual robots driven by their energy levels and sampling uncertainty.
- Ensure that all robots contribute fairly to estimating the underlying spatial process, especially in large-scale scenarios where the involvement of every available robot is critical. At a system level, it ensures that energetic robots contribute more to the exploration process in the IPP process, which requires large movements. Energy-deprived robots congregate at core sampling points, exploiting the predicted spatial distribution model.

- We enable a persistent MIPP formulation for scenarios with accessible recharging stations, optimizing the IPP process to utilize the available charging infrastructure.

These contributions are pivotal in maximizing overall information gain while efficiently managing energy consumption and increasing long-term operational performance in real-world applications such as ocean monitoring, multi-target tracking, surveillance, surveys in large-scale precision agricultural systems, and algae outbreaks.

II. RELATED WORK

IPP techniques include utility maximization within partitioned regions [12], conflict-minimizing algorithms [13], RRT*-based path refinement [14], asynchronous coordination for cooperative exploration [15], and Sparse Gaussian Processes for planning with limited data [16]. Area partitioning (e.g., Voronoi diagrams [17]) is widely used for dividing exploration regions [18], [3], [19], with observations later fused via Gaussian mixture models [2].

Most MIPP studies assume unlimited energy, with variants exploring communication and coordination constraints [1], [6], [2]. Di Garo et al. [1] propose a leader-follower framework with time budgets and collision avoidance. Dutta et al. [6] focus on maintaining network connectivity during joint path planning, while Luo and Sycara [20] address online sampling and coverage. Benevento et al. [21] assign one robot to sampling and others to coverage, later refined by Santos et al. [22] to balance roles among all robots. These approaches overlook energy constraints, prioritizing only information gain.

Research on heterogeneous teams remains sparse. Manjanna et al. [23] deploy robotic boats with fixed roles, while Praneel et al. [24] use a hierarchical system with varying computational capacities. Neither considers energy heterogeneity in planning.

Recent work has begun addressing energy disparity. Newaz et al. [25] decouple sampling and navigation, which can lead to suboptimal routes. Shi et al. [11] incorporate heterogeneous robot costs (distance, battery, terrain) but rely on fixed parameters, limiting adaptability. Munir et al. [26] adapt partition sizes to robot energy levels but focus on coverage, not IPP. Energy-aware planning with recharging stations has only recently been explored [27].

Our research closes this gap by introducing a distributed controller that adapts sampling to real-time energy and exploration progress. Unlike prior work, it integrates energy-aware exploration, recharging strategies, and persistent deployment, enabling long-term multi-robot operations. This improves energy efficiency and robustness in modeling spatial distributions, particularly for robots operating under size, weight, and power constraints in large-scale, dynamic missions.

III. PRELIMINARIES

A. Multi-robot Informative Path Planning (MIPP)

This problem focuses on n robots cooperatively exploring an environment $\mathcal{Q} \in \mathbb{R}^2$ and collecting valuable information. Each location $x \in \mathcal{Q}$ has an importance value $\phi(x)$. The $\phi(\cdot)$

denotes the underlying spatial distribution that is unknown to robots but can be observed with a sensor at the current location. Let $\mathcal{T}$ denote the set of feasible candidate locations. The i th robot at a position $x_i \in \mathcal{T}$ can observe samples along its path using $OBS(x_i, \phi(x_i))$. The objective is to maximize the information gain $\mathcal{I}$ while satisfying trajectory constraints C with bounds B :

$$\mathbf{x}^* = \operatorname{argmax}_{\{x_i \in \mathcal{T}\}} \mathcal{I} \left(\bigcup_{i=1}^n OBS(x_i, \phi(x_i)) \right), \quad \text{s.t. } \mathcal{C}(x_{1:n}) \leq B.$$

Here, $\mathbf{x}^*$ is the resulting set of robot states to visit next to maximize information gain before fully exploring the environment. This global objective is implemented in a distributed sense by using Voronoi partitioning to assign areas to robots optimally, where each robot's region $\mathcal{V}_i$ is defined as:

$$\mathcal{V}_i = \{q \in \mathcal{Q} \mid \|q - x_i\| \leq \|q - x_j\|, \forall j \neq i\}. \quad (1)$$

Each robot then individually selects the most informative point within its Voronoi region:

$$x_i^* = \operatorname{argmax}_{x \in \mathcal{V}_i} \mathcal{I}(x), \quad (2)$$

The robots are then controlled by navigating to the chosen informative point using any optimal path planning algorithms like A*, RRT*, for example, [28].

$$\dot{x}_i(t) = -k_p (x_i(t) - x_i^*), \quad (3)$$

where k_p is a control gain.

B. Gaussian Process Regression for IPP

Gaussian Processes (GPs) [29] are non-parametric models used to estimate unknown functions by defining a prior over the function space. Using observed data, GPs provide predictions along with uncertainty estimates. Given a set of observations $Y = \{y_1, y_2, \dots, y_n\}$ at input locations $X = \{x_1, x_2, \dots, x_n\}$, the process value at any point x is modeled as $y = \phi(x) + \epsilon$, where ϵ represents Gaussian noise. GPs are fully defined by a mean $\mu(x)$ and covariance $k(x, x')$, capturing the function's expected value and correlations between inputs: $\phi(x) \sim \mathcal{GP}(\mu(x), k(x, x'))$. Gaussian Process Regression (GPR) leverages this to model the unknown function $\phi(x)$ based on noisy data, often using the squared exponential kernel, with hyperparameters optimized via maximum log-likelihood [30]. Our aim is to learn $\phi(x)$ for all the unobserved data points $X^* = \{x_1^*, x_2^*, \dots, x_m^*\} \in \mathcal{Q} \setminus X$. For a given test location x^* , the mean and variance (uncertainty measure) of the posterior distribution of $\phi^*(x^*)$ is derived as follows:

$$\mu[x^*] = \mu(x) + k_*^T (K + \sigma_n^2 I)^{-1} (y - \mu(x)) \quad (4)$$

$$\sigma^2[x^*] = k_{**} - k_*^T (K + \sigma_n^2 I)^{-1} k_* \quad (5)$$

Here, $K = K(X, X)$ represents the covariance matrix obtained from the training data, k_* is the covariance between the training points and test points, and k_{**} is the covariance between only the test points. This GPR model is trained

after every new sample is collected. We use the mean and uncertainty outputs of the GPR model ($\mu[\cdot], \sigma^2[\cdot]$) in our information function (2) of the MIPP acting as a surrogate $\phi_*(\cdot)$ for the unknown spatial distribution $\phi(\cdot)$. Specifically, building on the established information function [31], we set

$$\mathcal{I}(x) = \alpha\mu(x) + \beta\sigma(x), \quad s.t., \alpha + \beta = 1, \quad (6)$$

where $\mu(x)$ and $\sigma^2(x)$ are the predicted mean and variance at point x of the surrogate function $\phi_*(\cdot)$, estimated via GPR. The parameter $\alpha = 1 - \beta$ is typically set based on the temporal constraints [26], [22] or regret bounds [32].

C. Energy Consumption Model

To represent the energy depletion of a robot, we adopt a widely recognized energy consumption model for mobile robots [33], [34], [26]. This model defines two key parameters:

- *Temporal Energy Cost*: E_{static} (energy consumed while the robot is idle, e.g., for sensing and computing, which can vary based on robot type and payload)
- *Spatial Energy Cost*: E_{motion} (energy used during movement depending on speed, altitude, and energy efficiency.)

The energy level of the robot i at time t is expressed as:

$$E_i(t) = E_i(t-1) - \dot{E}_i(t) \quad (7)$$

$$\dot{E}_i(t) = E_{static} \cdot dt - E_{motion} \cdot dx, \quad (8)$$

We assume that the robot does not know these energy parameters (E_{static} and E_{motion}) but can only measure its current energy level $E_i(t)$ (%), which is commonly available in most battery-operated devices.

IV. EA-MIPP CONTROLLER FORMULATION

The proposed Energy-aware MIPP (EA-MIPP) approach dynamically adapts each robot's role in exploration and recharging based on its energy level and confidence (variance) in prediction estimates, contributing optimally to the IPP process with heterogeneous robots. The controller assigns robots to one of three phases—exploration, energy conservation, or exploitation/recharging—by considering both the robot's energy and the variance in its assigned region.

At each time step t , each robot i at position x_i calculates its sampling goal location and applies velocity control $\dot{x}_i(t) \in \mathbb{R}^2$ according to the following control law that adapts between the current robot goal x_i^{goal} between the learning (IPP) objective (going to sampling goal x_i^{learn}) and the recharging objective (going to recharge station $x_i^{station}$):

$$\dot{x}_i = -k(x_i - x_i^{goal}), \quad (9)$$

$$\text{where, } x_i^{goal} = \tau_i(t) \cdot x_i^{learn} + (1 - \tau_i(t)) \cdot x_i^{station}. \quad (10)$$

Here, $\tau_i(t) \in [0, 1]$ is a dynamic weighting function that balances the robot's priorities between IPP and the need to recharge (if charging stations are available). Specifically, the value of $\tau_i(t) = S(E_i^{\%}, E_{low})$, where $E_i^{\%} = \frac{E_i(t)}{E_{max}}$ represents the robot's current energy as a percentage of its maximum

capacity, and E_{low} is the energy threshold for recharging necessity. The sigmoid function $S(x, \epsilon)$ is defined as:

$$S(x, \epsilon) = \frac{1}{1 + e^{-s(x-\epsilon)}}, \quad (11)$$

where $s > 0$ controls the steepness of the curve, and ϵ is a threshold parameter. When this measure $S(E_i^{\%})$ falls below a defined lower energy threshold E_{low} , the sigmoid function will smoothly transition to $\tau_i(t)$ from 1 to 0 when the threshold is hit, making the robot prioritize its recharging destination ($x_i^{station}$). When the energy is well above the threshold, it would instead guide the robot toward learning or information collection (x_i^{learn}). i.e., when $\tau_i \rightarrow 1$, the robot i aims to learn the underlying distribution by optimizing information collection within its designated region while managing energy efficiently, x_i^{learn} . If a recharging station is available, the robot moves toward the station by transitioning its $\tau_i \rightarrow 0$. Otherwise, the robot either selects another available station or waits for its turn (see Sec. IV-B).

A. Balancing exploration and Exploitation

Per the IPP process in (6), a robot's target learning position, x_i^{learn} , can be determined by:

$$x_i^{learn} = \arg \max_{x \in \mathcal{V}_i} \mathcal{I}^{learn}(x) \quad (12)$$

$$\mathcal{I}^{learn}(x) = \alpha_i^{learn}(t)\mu(x) + (1 - \alpha_i^{learn}(t))\sigma^2(x) \quad (13)$$

The parameter $\alpha_i^{learn}(t)$ has significant ramifications to the energy budget, which we leverage by further transitioning the IPP learning process between exploration (i.e. if $\alpha = 0$; robot goes to the location that has maximum uncertainty) and exploitation (i.e., if $\alpha = 1$; robot goes to the location that has maximum information value) states.

- Exploration (low $\alpha_i^{learn}(t)$) encourages robot i to seek new information at time t .
- Exploitation (high $\alpha_i^{learn}(t)$) focuses on refining the existing knowledge in known areas.

Adjusting $\alpha_i^{learn}(t)$ can significantly impact the distance traveled by the robot, which is one of the largest contributors to energy depletion. Specifically, increasing $\alpha_i^{learn}(t)$ from 0 to 0.25 can halve the distance traveled while maintaining acceptable mapping accuracy, as demonstrated in [7]. However, excessively increasing $\alpha_i^{learn}(t)$ risks reducing the accuracy of reconstructing spatial distributions.

To achieve this energy-aware strategy, each robot dynamically adjusts its exploitation parameter $\alpha_i^{learn}(t)$ based on the robot's current energy level, $E_i(t)$, and the variance $\sigma_{\mathcal{V}_i}^2$ within its assigned Voronoi region:

$$\alpha_i^{learn}(t) = \alpha_{max} - (\alpha_{max} - \alpha_{min}) \cdot S(E_i^{\%}, E_{thresh}) \cdot S(\sigma_{\mathcal{V}_i}^2, \sigma_{thresh}^2), \quad (14)$$

where α_{min} and α_{max} define the bounds for the exploration-exploitation tradeoff control gain, and the sigmoid functions $S(E_i^{\%}, E_{thresh})$ and $S(\sigma_{\mathcal{V}_i}^2, \sigma_{thresh}^2)$ are similarly defined as in (11) to adjust $\alpha_i^{learn}(t)$ dynamically based on real-time energy and variance levels. Based on the findings in [7], we set the bounds for the tradeoff gain as $\alpha_{min} = 0.25$ and $\alpha_{max} = 0.75$ for achieving an energy-balanced IPP.

We set thresholds for energy and variance, E_{thresh} and σ_{thresh}^2 , to further guide the robot's path. When energy or variance is low, $\alpha_i(t)$ increases, shifting the focus to conserve energy. Either of the thresholds (low energy E_{thresh}) and (low variance σ_{thresh}^2) needs to be hit by a robot to start conserving its energy. For example, reaching an uncertainty below the variance threshold (regardless of the energy level) means the robot's region is well explored, and the information gain would not be high enough if it keeps choosing the location with maximum uncertainty. In this case, the robot can choose the location of the peak distribution value it has predicted so far and reach saturation of its IPP objective, since it has learned the GPR models well. Another example is when the energy level is below the threshold (regardless of the uncertainty threshold), the robot benefits from transitioning into an energy-conservation mode by sampling in the neighborhoods of the peak of the distribution values while leaving the responsibility of intense exploration to other robots that have higher energy levels.

Prior studies have used variance thresholds to determine convergence and to guide the transition from exploration to exploitation [31], [35]. The proposed adaptive parameter $\alpha_i^{\text{learn}}(t)$ enables each robot to optimize its energy by sampling selection for information gain within its assigned region. This allows robots to transition between exploration and exploitation as needed, thereby sustaining persistent operations.

B. Recharging Strategy

When $\tau_i = 0$, the robot prioritizes recharging and moves towards the nearest charging station x_i^{station} , which is the location of the recharging station j assigned to robot i . To ensure robots recharge efficiently and promptly resume environmental modeling, we select the charging station that minimizes total downtime, accounting for both travel and waiting times. Given a set of charging stations $S = \{1, \dots, s\}$, the optimal station j is selected by minimizing: $\min_{j \in S} \left(\frac{d_j}{v_i} + w_j \right)$, where d_j is the distance to station j , v_i is the robot i 's velocity, and w_j represents the waiting time at station j . We assume that each robot must wait for a fixed number of iterations at the recharging station to achieve a full charge. If a robot is already charging, another robot assigned to the same station will wait nearby until the station is available.

C. Robot States and Transitions

The robot smoothly transitions through three operational phases based on energy and information conditions:

- **Exploration:** The robot actively explores when energy is sufficient and environmental variance is high, aiming to maximize information gain from uncertain regions.
- **Exploitation:** When energy or variance thresholds are met, the robot prioritizes regions with higher mean information, optimizing energy use by reducing movements.

- **Recharging/Waiting:** When energy is low, the robot seeks a charging station or enters a waiting state until it can recharge, ensuring continued operation.

V. EXPERIMENTS AND RESULTS

We conducted Python-based simulations implementing the proposed controller and rigorously evaluated it against relevant controllers from the literature. We consider two primary tasks.

- 1) **Task 1** focuses on energy management in environments without charging stations, so we can analyze the sole effect of the proposed controller on the energy conservation characteristics.
- 2) In contrast, the **Task 2** examines system performance in environments where recharging stations are available and enables us to analyze the applicability in persistent deployment scenarios.

In our experiments, we set the E_{thresh} parameter to the 60% of the maximum initial energy capacity in the system, and the $\sigma_{\text{thresh}}^2 = 3dB^2$. To validate our approach, we compared it against three multi-robot informative sampling strategies:

MIPP: A multi-robot Informative Path Planning (MIPP) method that selects goal points based on locations with the highest variance within each robot's Voronoi region [8].

H-MIPP: The heterogeneous multi-robot IPP approach described in [11] employs multiplicative Voronoi partitioning, where the region partitioning among robots is weighted based on user-defined parameters (weights).

EAC-MIPP: An adaptation of the energy-aware additive Voronoi partitioning, as outlined in [26], where the robot partitions (coverage areas) are influenced by real-time energy characteristics. To incorporate the IPP objective, goal points are selected based on the maximum variance within each robot's assigned region.

To quantify the disparity in robots' energy levels, we introduce a heterogeneity score defined as the sum of squared differences among the robots' initial energy levels E_i .

$$\text{Score HS} = \sqrt{\frac{1}{n} \sum_{i=1}^n (E_i - \bar{E})^2}, \quad (15)$$

where $\bar{E} = \frac{1}{n} \sum_{i=1}^n (E_i)$ is the mean of robot energy levels.

A. Performance Metrics

We evaluate performance using several metrics. Accuracy is measured by the Root Mean Square Error (RMSE) between the predicted and ground truth spatial distributions, while uncertainty is quantified by the variance of Gaussian Process Regression (GPR) predictions. Energy efficiency is captured through normalized cumulative energy availability, expressed as the ratio of available to maximum system energy, and by the percentage of energy conserved relative to baseline consumption. System activity is assessed through the number of active robots during operation, the number of robots still alive at the end of a trial, and the cumulative distance traveled. Convergence iteration denotes when RMSE

no longer improves, and recharging cycles indicate how often robots require recharging. Finally, robot states track the proportion of time spent in exploration, exploitation, waiting, and recharging phases.

B. Results and Discussion: Task 1

The results for the two scenarios evaluated in this Task 1 in a $40 \times 40\text{m}^2$ environment are shown in Figs. 2 and 3. Scenario 1 involves a heterogeneous energy configuration, where robots begin with either high (100) or moderate (60) energy levels. Specifically, the initial energy levels are given by $\{60, 60, 60, 100, 100, 100\}$.

H-MIPP algorithm accounts for the heterogeneity in initial energy levels by partitioning accordingly, while the baseline MIPP approach ignores energy considerations and focuses solely on path optimization. The results indicate that H-MIPP outperformed MIPP and EAC-MIPP by leveraging energy heterogeneity, whereas our proposed EA-MIPP clearly outperformed all three by incorporating real-time route adjustments based on available energy.

EA-MIPP approach balances exploration and exploitation effectively, leading to faster RMSE convergence. With an RMSE of 1.10, it outperforms HMR-MIPP (2.05), EAC-MIPP (2.49), and MIPP (2.49). The proposed EA-MIPP excels in energy efficiency and accurate environment modeling, conserving $\sim 7\text{x}$ more energy than the other approaches.

Scenario 2 explores a more diverse distribution of initial energy levels among the robots, ranging from 40 to 90, with the specific configuration given by $\{40, 40, 80, 80, 90, 90, 50, 50\}$.

As in the previous scenario, the H-MIPP algorithm effectively captures the underlying initial energy heterogeneity and demonstrates improved performance over both MIPP and EAC-MIPP. However, despite its advantage in balancing information gain with energy constraints, H-MIPP does not achieve the highest performance in terms of robot survival or overall energy conservation. EAC-MIPP continues to emphasize energy preservation, at the expense of mapping accuracy. In contrast, the baseline MIPP approach performs exhaustive exploration without considering energy constraints, resulting in premature energy depletion across multiple robots. In this scenario, the proposed EA-MIPP approach adaptively assigns roles based on the energy levels of the robots: robots with lower energy are directed toward exploitation, focusing on nearby informative regions, while those with higher energy are allocated to exploration tasks involving longer travel paths. The corresponding α plots offer further insight into the temporal evolution of these energy-aware role assignments.

EA-MIPP approach achieved an RMSE of 1.29, similar to that of H-MIPP, while outperforming EAC-MIPP and MIPP, which produced RMSE values of 2.31 and 1.96, respectively. Furthermore, EA-MIPP conserved 9x more energy than H-MIPP and MIPP and 1.5x more energy than EAC-MIPP. This underscores the effectiveness of the EA-MIPP approach in large-scale multi-robot applications, demonstrating its ability to optimize both informative sampling and energy efficiency.

C. Results: Task 2 - Persistent Deployment with Recharging

This task assesses the effectiveness of the proposed approach in sustaining MIPP operations in environments with charging stations. To emulate real-world battery swapping or replacement, a recharging delay of 10 iterations is introduced at the charging station. We incorporate the standard MIPP as a baseline, using the same E_{low} threshold as the proposed EA-MIPP approach, which directs robots to the charging station when their energy falls below this threshold. In this experiment, ten robots are deployed in a $60 \times 60\text{m}^2$ environment, with half having high energy levels (100) and the other half having moderate energy levels (60). As shown in Fig. 4, the EA-MIPP method not only reaches a lower RMSE more quickly but also demonstrates greater energy conservation, requiring 8 recharging cycles compared to the 11 cycles needed by the MIPP approach. These results further strengthen the findings from the previous setting experiments, reinforcing the EA-MIPP's superior performance and energy efficiency in environments with charging stations.

Persistent Deployment Analysis: We conducted additional experiments to evaluate the controllers under various persistent deployment settings, considering variations in robot energy heterogeneity and recharging station setting, as summarized in Table I. Specifically, we tested configurations with 2 and 4 charging stations using 10 robots, and 4 robots with 2 stations. The results demonstrate that the proposed approach reduced convergence iterations by 2 to 40, while maintaining the same RMSE as the MIPP. Furthermore, it decreased recharging cycles by 1 to 3, enhancing its effectiveness in persistent deployment scenarios. This improvement ensured higher robot availability, ranging from 2% to 6%, for the underlying path planning tasks.

D. Hardware Experiments

Here, we present the experimental results with an in-house heterogeneous swarm robotics testbed. Four robots, initialized with heterogeneous energy levels $\{100, 80, 60, 40\}$, are deployed in a $2.5 \times 1.5\text{m}^2$ environment, each operating at a constant velocity of 0.1 m/s. To emphasize the effects of energy-aware decision-making, the environment is scaled by a factor of 10. Two charging stations are positioned along the lower boundary of the workspace to facilitate energy replenishment. As shown in Fig. 5, the proposed persistent EA-IPP approach required only two charging cycles, compared to four cycles used by the baseline persistent IPP method. Additionally, the EA-IPP approach demonstrated faster RMSE convergence, indicating improved energy efficiency while maintaining accurate reconstruction of the environmental process, which is consistent with the trends observed in simulation studies.

E. Scalability Analysis

We evaluate the effectiveness of the proposed approach in larger environments by varying the number of robots and adjusting the dimensions of the workspace accordingly. The dimensions of the workspace ranged from $20 \times 20\text{m}^2$ to $100 \times 100\text{m}^2$, with the number of robots varying from 4 to




Fig. 2. Performance results of Task 1, Scenario 1 with initial energy levels $\{60, 60, 60, 100, 100, 100\}$. The top plots depict the evolution of RMSE, variance, cumulative energy, alive robots across different approaches, and the α parameter evolution in the proposed EA-MIPP. The bottom plots illustrate the ground truth and final predicted mean for all the approaches and robot trajectories. The "x" and "o" symbols mark the initial and the final robot locations, respectively.




Fig. 3. Performance results of Task 1, Scenario 2 with initial energy levels $\{40, 40, 80, 80, 90, 90, 50, 50\}$. The top plots depict the evolution of RMSE, variance, cumulative energy, alive robots across different approaches, and the α parameter in the proposed EA-MIPP. The bottom plots illustrate the ground truth and final predicted mean for all the approaches and robot trajectories. The "x" and "o" symbols mark the initial and the final robot locations, respectively.




Fig. 4. Demonstration of the MIPP and recharging objective for persistent IPP in environments with recharging stations (Task 2, Setting 2). The top plots illustrate the evolution of RMSE, variance, cumulative energy, and distance traveled by the robots for the proposed EA-MIPP approach compared with the MIPP baseline with our recharging strategy. The bottom plots display the final predicted mean, robot trajectory, and the state transitions in the proposed approach. The magenta squares at the left and right corners denote the charging stations.

TABLE I

COMPARISON OF MIPP AND EA-MIPP PERFORMANCE FOR DIFFERENT PERSISTENT DEPLOYMENT CONFIGURATIONS WITH RECHARGING.

Robots	Recharging Stations	Heterogeneity Score	RMSE		Available Robots (%)		Recharging Cycles		Exploration Time (%)		Convergence Iteration	
			MIPP	EAMIPP	MIPP	EAMIPP	MIPP	EAMIPP	MIPP	EAMIPP	MIPP	EAMIPP
10	4	0.67	0.85	0.7	91.9	95.7	10	7	91.9	54.2	232	192
10	4	0	0.57	0.53	86.9	88.7	13	10	86.84	37.73	296	271
10	2	0.67	0.49	0.76	88.66	89.93	10	8	88.67	55.27	276	227
10	2	0.77	0.62	0.54	91.67	92.8	10	9	91.67	53.37	299	281
4	2	0.89	0.67	0.60	83.75	87.5	4	3	83.75	44.38	77	75
4	2	1.0	0.81	0.66	86.25	92.2	4	2	85.94	36.88	72	65





Fig. 5. Demonstration of the MIPP and recharging objective for persistent IPP on swarm testbed. The top plots illustrate the evolution of RMSE, variance, cumulative energy, and distance traveled by the robots for the proposed EA-MIPP approach compared with the MIPP baseline with our recharging strategy. The bottom plots display the final predicted mean, robot trajectory along with final robots' positions, and the state transitions in the proposed approach. The magenta squares at the left and right corner denote the charging stations.

20, increasing proportionally to the size of the environment. In these experiments, we consider a heterogeneous scenario in which half of the robots started with 100% energy, while the other half began with 60% energy.

The percentage of energy difference, calculated for the conserved energy, allow us to determine how much the proposed EA-MIPP improves over the state-of-the-art methods. Specifically, we compare against the H-MIPP as it directly addresses the energy heterogeneity among robots. Simulations were run until either the RMSE saturated or the robots' energy was depleted. A noticeable trend of increasing energy gain with the number of robots in different environment sizes emerged, showing 22 to 32% energy savings. The results demonstrate the applicability of the proposed approach in large-scale environments (Fig. 6(a)).

F. Influence of the Energy Threshold

Here, we analyze the influence of the energy threshold parameter (E_{thresh}) in Eq. 14 evaluated on the setting similar to Scenario 1 in Task 1 and Setting 1. The energy threshold parameter is varied from 1 to 0, where $E_{\text{thresh}} = 1$ prioritizes quick exploitation based solely on energy levels. In contrast, $E_{\text{thresh}} = 0$ enables full exploration until the variance threshold is met, after which it switches to exploitation. Although similar to MIPP, this approach still transitions to exploitation. We aim to balance RMSE and energy consumption, optimizing the accuracy of IPP per unit of energy spent.



(a) Scalability (b) E_{thresh} effect on RMSE and available Energy

Fig. 6. Performance evaluation of energy-aware planning: (a) shows energy savings with increasing environment size and robot count; (b) demonstrates the effect of varying E_{thresh} on RMSE and energy levels.

From Fig. 6(b), we observe that prioritizing exploitation early (e.g., $E_{\text{thresh}} = 1$ or 0.8) conserves energy but results in higher RMSE. In contrast, continuous exploration ($E_{\text{thresh}} = 0$) leads to full energy depletion for all robots. The threshold values between 0.6 and 0.2 offer a good trade-off between RMSE and energy efficiency. Among these, $E_{\text{thresh}} = 0.6$ achieves the best balance, minimizing RMSE while conserving energy, making it the optimal choice for our experiments.

VI. CONCLUSION

This paper presents a novel energy-aware informative path-planning approach for multi-robot systems that optimizes data collection while considering each robot's distinct energy budget. By integrating energy characteristics with the confidence bounds of the assigned regions, the pro-

posed method effectively balances energy consumption and information gain. This ensures that the multi-robot system can efficiently and reliably model the underlying spatial distribution, even under energy constraints. The approach demonstrates its potential for improving the performance of energy-limited data collection and increasing survival capacity in long-term and large-scale missions. Furthermore, the scenario that extends the proposed approach to include recharging station availability highlights its applicability for persistent IPP, ensuring that the robots remain operational without significantly impacting their task.

REFERENCES

- [1] G. A. Di Caro and A. W. Z. Yousaf, "Multi-robot informative path planning using a leader-follower architecture," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 10 045–10 051.
- [2] A. Dutta, O. Patrick Kreidl, and J. M. O’Kane, "Opportunistic multi-robot environmental sampling via decentralized markov decision processes," in *Distributed Autonomous Robotic Systems: 15th International Symposium*. Springer, 2022, pp. 163–175.
- [3] W. Luo, C. Nam, G. Kantor, and K. Sycara, "Distributed environmental modeling and adaptive sampling for multi-robot sensor coverage," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019, pp. 1488–1496.
- [4] A. Viseras, D. Shutin, and L. Merino, "Online information gathering using sampling-based planners and gps: An information theoretic approach," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 123–130.
- [5] X. Cai, B. Schlotfeldt, K. Khosoussi, N. Atanasov, G. J. Pappas, and J. P. How, "Energy-aware, collision-free information gathering for heterogeneous robot teams," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2585–2602, 2023.
- [6] A. Dutta, A. Ghosh, and O. P. Kreidl, "Multi-robot informative path planning with continuous connectivity constraints," in *2019 International Conference on Robotics and Automation*. IEEE, 2019, pp. 3245–3251.
- [7] A. Munir and R. Parasuraman, "Exploration–exploitation tradeoff in the adaptive information sampling of unknown spatial fields with mobile robots," *Sensors*, vol. 23, no. 23, p. 9600, 2023.
- [8] K.-C. Ma, Z. Ma, L. Liu, and G. S. Sukhatme, "Multi-robot informative and adaptive planning for persistent environmental monitoring," in *Distributed Autonomous Robotic Systems: The 13th International Symposium*. Springer, 2018, pp. 285–298.
- [9] X. Cai, B. Schlotfeldt, K. Khosoussi, N. Atanasov, G. J. Pappas, and J. P. How, "Non-monotone energy-aware information gathering for heterogeneous robot teams," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8859–8865.
- [10] T. Setter and M. Egerstedt, "Energy-constrained coordination of multi-robot teams," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 4, pp. 1257–1263, 2016.
- [11] Y. Shi, N. Wang, J. Zheng, Y. Zhang, S. Yi, W. Luo, and K. Sycara, "Adaptive informative sampling with environment partitioning for heterogeneous multi-robot systems," in *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 11 718–11 723.
- [12] N. Fung, J. Rogers, C. Nieto, H. I. Christensen, S. Kemna, and G. Sukhatme, "Coordinating multi-robot systems through environment partitioning for adaptive informative sampling," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3231–3237.
- [13] B. Woosley, P. Dasgupta, J. G. Rogers III, and J. Twigg, "Multi-robot information driven path planning under communication constraints," *Autonomous Robots*, vol. 44, no. 5, pp. 721–737, 2020.
- [14] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.
- [15] D. Jang, J. Yoo, C. Y. Son, and H. J. Kim, "Fully distributed informative planning for environmental learning with multi-robot systems," *arXiv preprint arXiv:2112.14433*, 2021.
- [16] K. Jakkala and S. Akella, "Multi-robot informative path planning from regression with sparse gaussian processes," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 12 382–12 388.
- [17] C. He, Z. Feng, and Z. Ren, "Distributed algorithm for voronoi partition of wireless sensor networks with a limited sensing range," *Sensors*, vol. 18, no. 2, p. 446, 2018.
- [18] A. Dutta, A. Bhattacharya, O. P. Kreidl, A. Ghosh, and P. Dasgupta, "Multi-robot informative path planning in unknown environments through continuous region partitioning," *International Journal of Advanced Robotic Systems*, vol. 17, no. 6, p. 1729881420970461, 2020.
- [19] A. Munir, E. Latif, and R. Parasuraman, "Anchor-oriented localized voronoi partitioning for gps-denied multi-robot coverage," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 3395–3402.
- [20] W. Luo and K. Sycara, "Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6359–6364.
- [21] A. Benevento, M. Santos, G. Notarstefano, K. Paynabar, M. Bloch, and M. Egerstedt, "Multi-robot coordination for estimation and coverage of unknown spatial fields," in *2020 IEEE international conference on robotics and automation (icra)*. IEEE, 2020, pp. 7740–7746.
- [22] M. Santos, U. Madhushani, A. Benevento, and N. E. Leonard, "Multi-robot learning and coverage of unknown spatial fields," in *2021 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2021, pp. 137–145.
- [23] S. Manjanna, A. Q. Li, R. N. Smith, I. Rekleitis, and G. Dudek, "Heterogeneous multi-robot system for exploration and strategic water sampling," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4873–4880.
- [24] P. Chand and D. A. Carnegie, "Mapping and exploration in a hierarchical heterogeneous multi-robot system using limited capability robots," *Robotics and autonomous Systems*, vol. 61, no. 6, pp. 565–579, 2013.
- [25] A. A. R. Newaz, T. Alam, J. Mondello, J. Johnson, and L. Bobadilla, "Multi-robot information gathering subject to resource constraints," in *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 2021, pp. 1–6.
- [26] A. Munir, A. Dutta, and R. Parasuraman, "Energy-aware coverage planning for heterogeneous multi-robot system," in *Distributed Autonomous Robotic Systems: 16th International Symposium*. Springer, 2024.
- [27] G. Sharma, A. Dutta, and J.-H. Kim, "Optimal online coverage path planning with energy constraints," in *Proceedings of the 18th international conference on autonomous agents and multiagent systems*, 2019, pp. 1189–1197.
- [28] D. Kularatne, S. Bhattacharya, and M. A. Hsieh, "Time and energy optimal path planning in general flows," in *Robotics: science and systems*. Ann Arbor, MI, 2016, pp. 1–10.
- [29] C. E. Rasmussen and H. Nickisch, "Gaussian processes for machine learning (gpml) toolbox," *The Journal of Machine Learning Research*, vol. 11, pp. 3011–3015, 2010.
- [30] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [31] M. Mantovani, F. Pratissoli, and L. Sabattini, "Distributed coverage control for spatial processes estimation with noisy observations," *IEEE Robotics and Automation Letters*, 2024.
- [32] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE transactions on information theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [33] R. Parasuraman, K. Kershaw, P. Pagala, and M. Ferre, "Model based on-line energy prediction system for semi-autonomous mobile robots," in *2014 5th International Conference on Intelligent Systems, Modelling and Simulation*. IEEE, 2014, pp. 411–416.
- [34] Y. Mei and Y.-H. Lu, "Energy-efficient motion planning for mobile robots," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA’04. 2004*, vol. 5. IEEE, 2004, pp. 4344–4349.
- [35] D. Jang, J. Yoo, C. Y. Son, D. Kim, and H. J. Kim, "Multi-robot active sensing and environmental model learning with distributed gaussian process," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5905–5912, 2020.