



Distributed Multi-Robot Exploration With Connectivity Maintenance Under QoS Constraints

Hazem Chaabi

► To cite this version:

Hazem Chaabi. Distributed Multi-Robot Exploration With Connectivity Maintenance Under QoS Constraints. Computer Science [cs]. Université de Lille, 2025. English. ⟨NNT : 2025ULILB024⟩. ⟨tel-05379346v2⟩

HAL Id: tel-05379346

<https://hal.science/tel-05379346v2>

Submitted on 3 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



UNIVERSITY OF LILLE

DOCTORAL SCHOOL MADIS 631

RESEARCH INSTITUTE INRIA CENTRE OF THE UNIVERSITY OF LILLE

Thesis prepared by **Hazem Chaabi**
and will be publicly defended on **November 17, 2025**
for obtaining the degree of **PhD in Computer Science**
Discipline **IT and Applications**

Distributed Multi-Robot Exploration With Connectivity Maintenance Under QoS Constraints

In front of the jury composed of:

Isabelle Guerin-Lassous

Full Professor, UCBL

Bertrand Ducourthial

Full Professor, UTC

Andreas Rauh

Full Professor, University of Oldenburg

Marcelo Dias de Amorim

Research Director, CNRS

Nathalie Mitton

Research Director, Inria

Reviewer

Reviewer

Examiner

President

Thesis Supervisor



UNIVERSITÉ DE LILLE

ÉCOLE DOCTORALE MADIS 631

INSTITUT DE RECHERCHE INRIA CENTRE DE L'UNIVERSITÉ DE LILLE

Thèse préparée par **Hazem Chaabi**
et soutenue publiquement le 17 novembre 2025
pour l'obtention du grade de **Docteur en Informatique**
Discipline **Informatique et Applications**

Exploration Multi-Robot Distribuée avec Maintien de la Connectivité sous Contraintes de Qualité de Service

Devant le jury composé de :

Isabelle Guérin-Lassous

Professeure des Universités, UCBL

Bertrand Ducourthial

Professeur des Universités, UTC

Andreas Rauh

Professeur des Universités, Université
d'Oldenburg

Marcelo Dias de Amorim

Directeur de Recherche, CNRS

Nathalie Mitton

Directrice de Recherche, Inria

Rapporteuse

Rapporteur

Examineur

Président

Directrice de Thèse

Résumé

Les Systèmes Multi-Robots (MRS) sont devenus essentiels pour les missions autonomes dans des environnements inconnus ou dangereux, notamment dans des scénarios critiques tels que la recherche et le sauvetage, où une cartographie efficace et une communication robuste sont cruciales. Cependant, équilibrer efficacement une exploration rapide avec une connectivité réseau fiable demeure un défi, particulièrement pour les systèmes décentralisés opérant dans des conditions dynamiques sans contrôle centralisé.

Cette thèse présente un nouvel algorithme d'exploration multi-robots distribué, l'Exploration Dynamique Basée sur les Rôles avec Maintien de la Connectivité (DRBECM), spécifiquement conçu pour répondre à ces défis. L'algorithme proposé utilise une prise de décision décentralisée basée sur le partage d'informations locales, permettant une attribution autonome des rôles entre robots sans s'appuyer sur des informations globales ou une supervision centralisée. Les robots adoptent dynamiquement soit des rôles «d'explorateurs», se concentrant sur la maximisation du gain d'information à travers des stratégies basées sur les frontières, soit des rôles de «supporteurs», employant un positionnement inspiré du comportement d'essaim pour maintenir des liens de communication robustes à travers l'équipe. La sélection des voisins et le maintien de la connectivité sont gérés efficacement en utilisant le Graphe de Voisinage Relatif (RNG).

Pour améliorer davantage l'efficacité d'exploration sous des contraintes de communication réalistes, nous étendons DRBECM en un framework amélioré par apprentissage automatique, DRBECM-ML (Exploration Multi-Robots via Coordination d'Essaim et Évaluation de Connectivité Pilotée par l'Apprentissage Automatique). DRBECM-ML intègre une dynamique d'essaim distribuée avec des modèles d'apprentissage automatique légers, entraînés en utilisant des données de propagation de signal du monde réel provenant de la plateforme d'essai FIT-IoT-Lab, pour estimer avec précision les valeurs d'Indicateur d'Intensité du Signal Reçu (RSSI) en temps réel. Ces prédictions améliorent significativement la prise de décision autonome liée au changement de rôles et à la sélection des frontières, assurant des réseaux de communication stables et résilients tout au long des tâches d'exploration. Les évaluations comparatives indiquent que les algorithmes basés sur les arbres, incluant les Arbres de Décision et l'Amplification de Gradient Extrême (XGBoost), offrent un équilibre optimal entre précision d'estimation et efficacité computationnelle adapté au déploiement sur robots mobiles.

De plus, cette thèse examine des stratégies de communication alternatives en comparant les performances des k plus proches voisins (KNN) par rapport au RNG pour la communication inter-robots, utilisant des données générées via le simulateur réseau ns-3 pour analyser leur efficacité sous diverses configurations. Nos résultats de simulation montrent de manière constante des améliorations significatives du temps d'exploration et une réduction de l'exploration redondante comparé aux approches de référence, tout en maintenant efficacement la connectivité réseau. Enfin, ce travail vise à fournir des solutions de systèmes multi-robots robustes, adaptables et décentralisées, appropriées

pour le déploiement dans des scénarios réels complexes, dynamiques et à infrastructure limitée.

Abstract

Multi-Robot Systems (MRS) have become essential for autonomous missions in unknown or hazardous environments, notably in critical scenarios like search and rescue, where efficient mapping and robust communication are crucial. However, effectively balancing rapid exploration with reliable network connectivity remains challenging, especially for decentralized systems operating under dynamic conditions without centralized control.

This thesis introduces a novel distributed multi-robot exploration algorithm, Dynamic Role-Based Exploration with Connectivity Maintenance (DRBECM), specifically designed to address these challenges. The proposed algorithm utilizes decentralized decision-making based on local information sharing, enabling autonomous role assignment among robots without relying on global information or centralized oversight. Robots dynamically adopt either "explorer" roles, focusing on maximizing information gain through frontier-based strategies, or "supporter" roles, employing flocking-inspired positioning to sustain robust communication links across the team. Neighbor selection and connectivity maintenance are efficiently managed using the Relative Neighborhood Graph (RNG).

To further enhance exploration efficiency under realistic communication constraints, we extend DRBECM into a machine learning-enhanced framework, Multi-Robot Exploration via Flocking Coordination and Machine Learning-Driven Connectivity Assessment (DRBECM-ML). DRBECM-ML integrates distributed flocking dynamics with lightweight machine learning models, trained using real-world signal propagation data from the FIT-IoT-Lab testbed, to accurately predict Received Signal Strength Indicator (RSSI) values in real-time. These predictions significantly improve autonomous decision-making related to role-switching and frontier selection, ensuring stable and resilient communication networks throughout exploration tasks. Comparative evaluations indicate that tree-based algorithms, including Decision Trees and Extreme Gradient Boosting (XGBoost), offer optimal balance between prediction accuracy and computational efficiency suitable for deployment on mobile robots.

Furthermore, this thesis investigates alternative communication strategies by comparing the performance of K-Nearest Neighbors (KNN) against the RNG for inter-robot communication, utilizing data generated via the Network Simulator 3 (NS-3) to analyze their effectiveness under various configurations. Our simulation results consistently show significant improvements in exploration time and reduction in redundant exploration compared to baseline approaches, while effectively maintaining network connectivity. Ultimately, this work aims to provide robust, adaptable, and decentralized multi-robot system solutions suitable for deployment in complex, dynamic, and infrastructure-limited real-world scenarios.

Contents

Contents	vii
1 Introduction	1
1.1 Context and Motivation	1
1.1.1 The Critical Need for Autonomous Exploration Systems	1
1.1.2 The Exploration-Connectivity Trade-off	2
1.1.3 Limitations of Existing Approaches	3
1.1.4 Towards Intelligent, Adaptive Multi-Robot Systems	4
1.2 The NEPHELE Project Context	5
1.3 Problem Statement	7
1.3.1 Core Research Challenges	7
1.3.2 Integration and Implementation Challenges	8
1.4 Research Objectives and Contributions	9
1.4.1 Research Questions and Objectives	9
1.4.2 Summary of Contributions	10
1.5 Thesis Organization	11
1.6 List of publications	12
2 Background and Related Work	13
2.1 Multi-Robot Exploration Strategies	13
2.1.1 Centralized Exploration Approaches	14
2.1.2 Decentralized Exploration Approaches	15
2.1.3 Hybrid Exploration Approaches	18
2.2 Connectivity-Aware Exploration and Maintenance	20
2.2.1 Exploration-Communication Trade-Offs	20
2.2.2 Distributed Control and Coordination Mechanisms	22
2.2.3 Local Information-Based Solutions	24
2.2.4 Fault Tolerance and Robustness	25
2.3 Bio-Inspired Coordination	28
2.3.1 Flocking as a Distributed Coordination Mechanism	28
2.3.2 Swarm Intelligence Applications	29
2.4 Machine Learning in Robotics	31
2.4.1 Learning-Based Navigation and Exploration	31
2.4.2 Predictive Models for Connectivity and Communication	33
2.4.3 Challenges of Distributed Learning	34
2.5 Summary and Research Gaps	36
2.5.1 Key Insights from the Literature	36
2.5.2 Identified Research Gaps	37
2.5.3 Motivation for This Research	38
3 Comparative Analysis of Dynamic Communication Topologies for MRS	39
3.1 Introduction	39
3.2 Background on Graph-Based Topology Strategies	40
3.2.1 Relative Neighborhood Graph (RNG)	40
3.2.2 k-Relative Neighborhood Graph (K-RNG)	40
3.2.3 K-Nearest Neighbors (KNN)	40
3.2.4 Detailed Theoretical Comparison	41
3.3 Simulation and Evaluation Methodology	43
3.3.1 Scenario 1: Peer-to-Peer Communication	43
3.3.2 Scenario 2: Multi-hop Routing	44

3.3.3	Performance Metrics	44
3.4	Results of Comparative Analysis	44
3.4.1	Peer-to-Peer Performance (Scenario 1)	44
3.4.2	Multi-hop Routing Performance (Scenario 2)	45
3.5	Chapter Summary and Implications	47
3.6	Limitations and Future Directions	48
3.6.1	Current Limitations	48
3.6.2	Future Research Directions	49
4	DRBECM: A Distributed Multi-Robot Exploration Approach with Connectivity Maintenance	51
4.1	Introduction	51
4.2	System Model and Background	51
4.3	The DRBECM Algorithm	55
4.3.1	Dynamic Role Assignment	55
4.3.2	Explorer Robot Behavior: Frontier-Based Exploration	56
4.3.3	Supporter Robot Behavior: Flocking-Inspired Positioning	56
4.3.4	Stagnation Detection and Collision Avoidance	57
4.3.5	Information Sharing and Map Updating	59
4.3.6	Main Algorithm	59
4.4	Performance Evaluation	61
4.4.1	Experimental Setup	61
4.4.2	Results and Analysis	62
4.5	Limitations and Future Enhancements	65
4.5.1	Discussion and Current Limitations	65
4.5.2	Algorithmic Enhancements	65
4.5.3	Deployment and Validation Requirements	66
5	DRBECM-ML: Enhancing Exploration with Machine Learning-Driven Connectivity	69
5.1	Introduction	69
5.2	The DRBECM-ML Algorithm	70
5.2.1	Data-Driven Dynamic Role Assignment	70
5.2.2	Enhanced Frontier Selection	70
5.3	Machine Learning Model for RSSI Prediction	71
5.3.1	Data Collection Using FIT IoT-Lab	71
5.3.2	Data Preprocessing	72
5.3.3	Model Training and Hyper-parameter Tuning	73
5.4	Performance Evaluation	73
5.4.1	Machine Learning Model Evaluation	74
5.4.2	Comparative Simulation of Exploration Algorithms	75
5.5	Limitations and Research Directions	78
5.5.1	Discussion and Current Limitations	78
5.5.2	Advanced Machine Learning Integration	79
5.5.3	Real-World Validation and Extension	80
6	Conclusion	81
6.1	Summary of Contributions	81
6.1.1	Comparative Analysis of Communication Topologies	81
6.1.2	Dynamic Role-Based Exploration Framework (DRBECM)	82
6.1.3	Machine Learning-Enhanced Framework (DRBECM-ML)	82
6.2	Discussion of Results	83
6.2.1	Performance Analysis and Trade-offs	83
6.2.2	Scalability and Robustness Characteristics	85
6.2.3	Limitations and Considerations	85

6.3	Perspectives	86
6.3.1	Algorithmic and Theoretical Extensions	86
6.3.2	Technological Enhancements	88
6.3.3	Application-Driven Research Directions	90
6.3.4	Ethical and Deployment Considerations for Emergency Response Systems	91
	List of Terms	93
	Bibliography	95

List of Figures

1.1	NEPHELE Use Case #1 (Emergency/Disaster Recovery) overview	6
1.2	Compute-network continuum for Use Case 1	7
3.1	Lune region for RNG: the shaded intersection of disks of radius $d(p, q)$	40
3.2	K-RNG lune with $k = 1$: the single interior point c does not block the edge. .	40
3.3	Graph comparison: RNG, k-RNG and k-NN on the same 15 points example. . .	41
3.4	Packet Delivery Ratio (PDR) as a function of the number of robots for all evaluated strategies in the peer-to-peer communication scenario.	45
3.5	Average peer-to-peer delay as a function of the number of robots for all evaluated strategies.	46
3.6	Box plot distribution of PDR per strategy in the multi-hop data collection scenario.	46
3.7	RNG multi-hop PDR heatmap across network scales and deployment areas. . .	46
3.8	KNN (K=1) multi-hop PDR heatmap.	47
3.9	KNN (K=2) multi-hop PDR heatmap	47
3.10	KNN (K=3) multi-hop PDR heatmap	47
3.11	1-RNG multi-hop PDR heatmap	48
3.12	2-RNG multi-hop PDR heatmap	48
3.13	3-RNG multi-hop PDR heatmap	48
4.1	RNG network topology in DRBECM algorithm	55
4.2	Illustration of safe and unsafe frontiers	56
4.3	Progression of DRBECM.	58
4.4	Collision avoidance repulsion function $\phi(d)$	59
4.5	Distribution of exploration times across different methods	63
4.6	Distribution of exploration times across different methods	63
4.7	Comparison of exploration efficiency.	64
4.8	Redundant exploration across different methods.	64
5.1	The IoT-LAB M3 board	71
5.2	FIT IoT-LAB Lille	71
5.3	Representation of the Interquartile Range	72
5.4	Box plot of exploration times for DRBECM-ML and other methods	77
5.5	Comparison of redundant exploration for DRBECM-ML and other methods . .	78

List of Tables

2.1	Comparison of <i>decentralized</i> multi-robot exploration approaches	18
3.1	Theoretical Comparison of KNN, RNG, and K-RNG	42
4.1	DRBECM Parameter Specification	54
5.1	Reception log raw data snippet	72
5.2	Sample of preprocessed dataset with normalized coordinates and RSSI measurements	73
5.3	Model Performance Comparison - Grouped by Algorithmic Approach . . .	75
5.4	Simulation Parameters	76
5.5	Percentage of successful exploration runs for each method.	76

The emergence of multi-robot systems as a transformative solution for autonomous exploration in challenging and dangerous environments represents a significant advancement in robotics and artificial intelligence [1, 2]. This introductory chapter establishes the fundamental context, motivation, and scope of this thesis, which focuses on the development of distributed multi-robot exploration frameworks that can effectively balance exploration efficiency with communication connectivity maintenance.

The chapter begins by examining the critical need for autonomous exploration systems in emergency response scenarios [1, 3], where traditional human-based approaches face significant limitations due to safety concerns, time constraints, and accessibility challenges. It then delves into the fundamental exploration-connectivity trade-off that lies at the heart of multi-robot coordination, analyzing how existing approaches fail to adequately address the complex requirements of real-world deployments.

Building upon this foundation, the chapter introduces the specific research problem addressed in this thesis and presents the NEPHELE project context that motivates this work. The discussion progresses to articulate the core research challenges in distributed multi-robot exploration, leading to the formulation of specific research objectives and contributions. The chapter concludes with an overview of the thesis organization, providing a roadmap for the technical developments presented in subsequent chapters.

1.1	Context and Motivation .	1
1.1.1	The Critical Need for Autonomous Exploration Systems	1
1.1.2	The Exploration-Connectivity Trade-off . .	2
1.1.3	Limitations of Existing Approaches	3
1.1.4	Towards Intelligent, Adaptive Multi-Robot Systems	4
1.2	The NEPHELE Project Context	5
1.3	Problem Statement	7
1.3.1	Core Research Challenges	7
1.3.2	Integration and Implementation Challenges	8
1.4	Research Objectives and Contributions	9
1.4.1	Research Questions and Objectives	9
1.4.2	Summary of Contributions	10
1.5	Thesis Organization . . .	11
1.6	List of publications	12

1.1 Context and Motivation

The increasing frequency and severity of natural disasters, industrial accidents, and emergency situations worldwide have underscored the critical importance of rapid, autonomous assessment and exploration capabilities in affected areas. This section establishes the fundamental context for this thesis by examining the limitations of traditional human-based response approaches and the transformative potential of multi-robot systems for addressing these challenges.

The discussion progresses from the specific motivations arising from real-world emergency scenarios to the fundamental technical challenges that emerge when deploying multi-robot teams in unknown, infrastructure-denied environments. Through this analysis, we establish the critical need for systems that can balance aggressive exploration with reliable communication maintenance - a challenge that forms the central focus of this research.

1.1.1 The Critical Need for Autonomous Exploration Systems

In the aftermath of natural disasters, industrial accidents, or emergency situations, rapid assessment and exploration of affected areas is critical for effective response operations. The 2011 Fukushima nuclear disaster, the

[3]: Shen et al. (2021), 'Robots Under COVID-19 Pandemic: A Comprehensive Survey'

[4]: Nagatani et al. (2013), 'Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots'

[5]: Verma et al. (2021), 'Multi-Robot Coordination Analysis, Taxonomy, Challenges and Future Scope'

[6]: Brambilla et al. (2013), 'Swarm robotics: a review from the swarm engineering perspective'

[7]: Amigoni et al. (2017), 'Multirobot Exploration of Communication-Restricted Environments: A Survey'

[8]: Gielis et al. (2022), *A Critical Review of Communications in Multi-Robot Systems*

[9]: Hollinger et al. (2012), 'Multirobot Coordination With Periodic Connectivity: Theory and Experiments'

[10]: Yamauchi (1997), 'A frontier-based approach for autonomous exploration'

2010 Haiti earthquake, and more recent events such as the COVID-19 pandemic have highlighted the urgent need for autonomous systems capable of operating in environments that are too dangerous, inaccessible, or contaminated for human intervention [3, 4]. Traditional approaches relying on human teams face significant limitations due to hazardous conditions, time constraints, radiation exposure risks, structural instability, and the physical impossibility of accessing all affected areas simultaneously within critical time windows [4].

Multi-Robot Systems (MRS) have emerged as a transformative solution to these challenges, offering the potential for coordinated, autonomous exploration of unknown and potentially dangerous environments while minimizing human exposure to risk [5]. The advantages of multi-robot systems over single-robot solutions are particularly pronounced in exploration scenarios: enhanced robustness through redundancy, improved efficiency through parallel operation, increased coverage capability, and the ability to perform multiple complementary tasks simultaneously [5, 6].

However, the deployment of multi-robot systems in real-world emergency scenarios introduces a complex set of technical challenges that extend far beyond those encountered in controlled laboratory environments. These include dynamic environmental conditions, unpredictable obstacles, limited computational resources, energy constraints, and most critically, the need to maintain coordinated behavior despite unreliable communication infrastructure [7, 8].

1.1.2 The Exploration-Connectivity Trade-off

The fundamental challenge in multi-robot exploration lies in balancing two competing and often contradictory objectives: maximizing exploration efficiency while maintaining reliable communication connectivity among team members. This represents a classical optimization problem with multiple conflicting objectives that cannot be simultaneously optimized without trade-offs [7, 9].

Exploration Efficiency demands that robots venture into unknown territories to gather information, build comprehensive environmental maps, and maximize the rate of information acquisition in a given amount of time. Optimal exploration strategies would direct robots toward the most informative locations, typically at the frontiers between known and unknown space, regardless of communication constraints [10]. This objective favors aggressive, far-reaching exploration patterns that can rapidly cover large areas and discover critical information such as survivor locations, structural damage assessments, or environmental hazards within time-constrained mission scenarios.

Communication Connectivity, conversely, requires robots to remain within communication range of their teammates and base stations to enable coordination, real-time data sharing, collaborative decision-making, and mission control. Maintaining connectivity is essential for sharing map information, coordinating exploration assignments to avoid redundant coverage, enabling collaborative behaviors such as collaborative localization and mapping, and ensuring that critical information reaches human operators for decision-making [7, 8].

This inherent tension becomes particularly acute in infrastructure-denied environments where no pre-existing communication infrastructure exists to support robot operations. In such scenarios, robots must create and

maintain their own communication network while simultaneously accomplishing their exploration objectives [2, 11]. The challenge is further complicated by the dynamic nature of the environment, where robot movement constantly changes the network topology, and environmental factors such as obstacles and signal attenuation create unpredictable communication conditions [7, 12]. The complexity of this trade-off also depends on the nature of the data being collected and the number of robots [13, 14].

The exploration-connectivity trade-off manifests differently depending on the operational context. In time-critical scenarios such as search and rescue operations, delays due to connectivity constraints can severely impact mission outcomes; conversely, in scenarios where real-time coordination is essential for safety or mission success, connectivity failures can lead to redundant work or conflicts between robots operating without coordination [9, 15].

1.1.3 Limitations of Existing Approaches

Traditional centralized approaches to multi-robot exploration, while potentially optimal in terms of global coordination and information utilization, suffer from critical vulnerabilities that make them impractical for real-world emergency response deployments. **Single points of failure** represent the most significant weakness: the loss of the central controller due to hardware failure, communication breakdown, or physical damage would immediately incapacitate the entire robotic team [16]. **Communication bottlenecks** arise when all robots must constantly communicate with the central controller, creating network congestion and limiting scalability [8]. These bottlenecks are further exacerbated by the lack of reactivity in case of congestion, increased latency due to centralized processing, and elevated energy consumption from constant communication requirements [7, 13, 17]. The **requirement for global environmental knowledge** assumes perfect information sharing, which is unrealistic in dynamic environments with unreliable communication links [7].

Furthermore, centralized approaches typically require high-bandwidth, low-latency communication links to function effectively, these requirements are often impossible to satisfy in disaster scenarios where communication infrastructure has been damaged or destroyed [8]. The computational burden on the central controller also scales poorly with team size, creating scalability limitations that prevent the deployment of large robot teams when they are most needed [18].

These fundamental limitations have driven researchers toward decentralized solutions that distribute decision-making among individual robots, thereby eliminating single points of failure and reducing communication overhead. However, existing decentralized approaches face their own significant challenges that limit their practical applicability.

Idealized Communication Models: Many current decentralized methods rely on oversimplified assumptions about communication ranges. In reality, wireless communication in unknown environments is affected by path loss, shadowing, multipath fading, interference, obstacles, and non-line-of-sight propagation that cannot be captured by simple geometric models [7, 12]. This disconnect often leads to failures when robots lose connectivity due to unpredicted signal degradation or make suboptimal decisions based on inaccurate connectivity assessments.

[2]: Rooker et al. (2007), 'Multi-robot exploration under the constraints of wireless networking'

[11]: Pei et al. (2013), 'Connectivity and bandwidth-aware real-time exploration in mobile robot networks'

[12]: Bonilla Licea et al. (2024), 'When Robotics Meets Wireless Communications: An Introductory Tutorial'

[13]: Saboia et al. (2022), 'ACHORD: Communication-Aware Multi-Robot Coordination With Intermittent Connectivity'

[14]: Halsted et al. (2021), *A Survey of Distributed Optimization Methods for Multi-Robot Systems*

[15]: Tranzatto et al. (2022), *Team CERBERUS Wins the DARPA Subterranean Challenge: Technical Overview and Lessons Learned*

[16]: Prorok et al. (2021), *Beyond Robustness: A Taxonomy of Approaches towards Resilient Multi-Robot Systems*

[17]: Kabir et al. (2020), *Efficient Multi-Robot Exploration with Energy Constraint based on Optimal Transport Theory*

[18]: Korsah et al. (2013), 'A comprehensive taxonomy for multi-robot task allocation'

[19]: Hoog et al. (2009), ‘Role-Based Autonomous Multi-robot Exploration’
 [20]: Emam et al. (2020), ‘Adaptive Task Allocation for Heterogeneous Multi-Robot Teams with Evolving and Unknown Robot Capabilities’

[21]: Fink et al. (2010), ‘Online methods for radio signal mapping with mobile robots’
 [22]: Quattrini Li et al. (2020), ‘Multi-robot online sensing strategies for the construction of communication maps’
 [23]: Miyagusuku et al. (2018), ‘Precise and accurate wireless signal strength mappings using Gaussian processes and path loss models’

Static Role Assignment: Many approaches assign fixed roles to robots (e.g., leader, follower, relay) that do not adapt to changing mission requirements or environmental conditions, reducing effectiveness in dynamic scenarios. Here, effectiveness refers to the system’s ability to achieve exploration objectives within time and resource constraints while maintaining operational reliability. This limitation becomes particularly pronounced in cases of robot malfunction, where static assignments cannot redistribute tasks among remaining operational robots [19, 20].

Limited Adaptability: Static exploration strategies cannot respond effectively to the changing conditions typical of disaster scenarios, where obstacles, communication dead zones, varying terrain, and emerging priorities require continuous adaptation of individual and team behaviors [5, 6].

Coordination Overhead: While decentralized approaches eliminate single points of failure, they often require extensive inter-robot communication and coordination protocols that can overwhelm limited bandwidth resources and create computational bottlenecks on individual robots, and lead to increased energy consumption due to frequent message exchanges and processing overhead [8, 11].

1.1.4 Towards Intelligent, Adaptive Multi-Robot Systems

The limitations above highlight the need for a new generation of multi-robot exploration systems that combine the robustness and scalability advantages of decentralized architectures with intelligent, adaptive behaviors that can respond quickly to real-world deployment challenges. Such systems must incorporate several key characteristics:

Dynamic Adaptability: The ability to automatically adjust robot roles, exploration strategies, and coordination mechanisms in response to changing environmental conditions, mission requirements, and system state [16].

Realistic Communication Modeling: Integration of realistic wireless communication models (e.g., data-driven RSSI/throughput maps) that account for real-world propagation characteristics, enabling more accurate connectivity assessment and better decision-making [21–23].

Scalable Coordination: Coordination mechanisms that scale with team size while maintaining computational efficiency and avoiding excessive communication overhead [6, 18].

Robust Operation: The ability to maintain effective operation despite individual robot failures, communication link failures, and unpredictable environmental conditions, as emphasized in recent fielded systems [15, 16].

This thesis addresses these requirements through novel algorithmic frameworks that integrate bio-inspired coordination mechanisms, dynamic role adaptation, and machine-learning-driven communication modeling to create more intelligent, robust, and practical multi-robot exploration systems [6, 22, 23].

1.2 The NEPHELE Project Context

This PhD has been funded in the context of the NEPHELE project. NEPHELE is a Horizon Europe project that targets the efficient, reliable, and secure orchestration of hyper-distributed applications across the compute continuum, from IoT devices to edge and cloud. Its vision is operationalized through:

- ▶ A multi-layered *Virtual Object Stack* (VOStack) that virtualizes heterogeneous IoT devices and functions at the edge, enabling openness and interoperability
- ▶ A *synergetic meta-orchestration framework* that deploys and manages application graphs spanning IoT–edge–cloud, leveraging 5G, distributed AI, and cybersecurity mechanisms to bring computation and intelligence close to where data is produced.

Beyond technology, NEPHELE promotes an open development ecosystem (based on Eclipse tooling) and a software repository of reusable components (VOs, composite VOs, and generic IoT enablers), with extensive validation across multiple verticals, including logistics, smart buildings, healthcare, and emergency response.*

Use Case: Emergency/Disaster Recovery

This use case focuses on post-disaster operations (e.g., at a container terminal), aiming to enhance situational awareness for first responders and support the prioritization of rescue actions. The demonstrator includes a simulated partial deployment in the Port of Koper (container fall and suspicious-substance monitoring).[†] Figure 1.1 provides an overview of the operational scene, while Figure 1.2 highlights the data, computation, and network paths across the continuum.

Operational scenarios.

- ▶ **Mapping:** Multi-robot mapping of affected areas with ground robots and/or drones using cameras and Light Detection and Ranging (LiDAR).
- ▶ **Victim detection & injury assessment:** Perception-driven victim localization and assessment with aerial/ground platforms.
- ▶ **Risk prediction:** Anticipation of hazards via robots, drones, and deployed sensor nodes.
- ▶ **Device deployment & sampling:** Ground robots physically deploy sensors and collect liquid samples.
- ▶ **Network & device monitoring:** Continuous monitoring of connectivity for robots, drones, and sensor nodes.

Technical constraints (compute–network–application).

- ▶ **Device management:** Pre-deployment of functionalities on devices/edge, secure bootstrapping/self-configuration, hot add/remove, heterogeneity support, and software self-healing.
- ▶ **Software orchestration:** Dynamic placement and redeployment of components according to service goals and resource availability, with multi-level performance monitoring.

* <https://nephele-project.eu/objectives>

† <https://nephele-project.eu/use-case-1-emergencydisaster-recovery>

- **Low-latency communication:** Support for mobility, intermittent connectivity, and low-delay links to/from the disaster area (5G/Wi-Fi/Ad-hoc).
- **Dynamic multi-robot mapping & fleet management:** Coordination, monitoring, and task optimization across mobile robots via their VOs.
- **Computer vision/AI:** Information extraction and position detection from images and video streams.
- **Smart data handling:** Filtering, down-sampling, and aggregation of heterogeneous, high-volume sensor/robot/camera data before edge/cloud offload.
- **Robots-WSN interaction:** Guidelines for direct communication between robots and wireless sensor nodes.
- **Monitoring dashboard:** A GUI at the operations base for data visualization and mission control.

The distributed multi-robot exploration frameworks developed in this thesis directly address critical requirements of NEPHELE's emergency response use case. The dynamic role-based coordination mechanism aligns with NEPHELE's need for adaptive fleet management and self-configuration capabilities, while the connectivity-aware exploration strategies support the project's requirements for reliable low-latency communication in infrastructure-denied scenarios. The machine learning-enhanced connectivity prediction contributes to NEPHELE's distributed AI objectives, enabling intelligent orchestration decisions at the edge without relying on centralized cloud resources. Furthermore, the decentralized architecture ensures system resilience and supports NEPHELE's vision of robust operation across the IoT-edge-cloud continuum, particularly in disaster scenarios where traditional infrastructure may be compromised.




Figure 1.1: NEPHELE Use Case #1 (Emergency/Disaster Recovery) overview. Source: NEPHELE project website [24].




Figure 1.2: Compute-network continuum for Use Case #1: IoT devices/robots/drones, edge servers, and cloud interconnected via 5G/Wi-Fi/IoT gateways for orchestration, monitoring, and data processing. Source: NEPHELE project website [24].

1.3 Problem Statement

Building upon the context established by both the general challenges in multi-robot exploration and the specific requirements of the NEPHELE project, the central problem addressed in this thesis is the development of a distributed multi-robot exploration system that can effectively balance exploration efficiency with communication connectivity maintenance in unknown, infrastructure-denied environments typical of emergency response scenarios [7, 8, 25, 26].

This overarching problem can be decomposed into several interconnected research challenges that must be addressed simultaneously to create a practical and effective solution:

1.3.1 Core Research Challenges

Challenge 1: Dynamic Role Adaptation Under Uncertainty

How can robots autonomously adapt their roles and behaviors in response to changing exploration and communication conditions without relying on centralized coordination or global environmental knowledge? This challenge is particularly complex because:

- ▶ Role decisions must be made based on limited, local information that may be incomplete or outdated [7, 27].
- ▶ The optimal role assignment changes dynamically as the environment is explored and robot positions evolve.
- ▶ Role transitions must be coordinated among team members to avoid conflicts or gaps in coverage [25, 26, 28].
- ▶ The system must maintain stability and avoid oscillatory behavior where robots continuously change roles [28, 29].

Challenge 2: Connectivity-Aware Exploration with Realistic Constraints

How can exploration strategies account for realistic communication constraints while maintaining system-wide connectivity and avoiding network fragmentation? This requires addressing:

- ▶ The development of exploration strategies that explicitly consider communication requirements as a constraint [7, 8, 30].

[25]: Burgard et al. (2005), ‘Coordinated multi-robot exploration’

[26]: Zlot et al. (2002), ‘Multi-robot exploration controlled by a market economy’

[27]: Gerkey et al. (2004), ‘A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems’

[28]: Choi et al. (2009), ‘Consensus-Based Decentralized Auctions for Robust Task Allocation’

[29]: Hunt et al. (2014), ‘A Consensus-Based Grouping Algorithm for Multi-agent Cooperative Task Allocation with Complex Requirements’

[30]: Banfi et al. (2018), ‘Strategies for coordinated multirobot exploration with recurrent connectivity constraints’

[31]: Stump et al. (2008), ‘Connectivity management in mobile robot teams’

[32]: Michael et al. (2008), ‘Maintaining Connectivity in Mobile Robot Networks’

[33]: Ghaffarkhah et al. (2011), ‘Communication-Aware Motion Planning in Mobile Networks’

[34]: Abu-Aisheh et al. (2022), ‘CARA: Connectivity-Aware Relay Algorithm for Multi-Robot Expeditions’

[35]: Silva et al. (2024), *Communication-Constrained Multi-Robot Exploration with Intermittent Rendezvous*

[36]: Min et al. (2017), *A Directional Antenna based Leader-Follower Relay System for End-to-End Robot Communications*

[37]: Yan et al. (2014), ‘To Go or Not to Go: On Energy-Aware and Communication-Aware Robotic Operation’

[38]: Lindhé et al. (2007), ‘An experimental study of exploiting multipath fading for robot communications’

[39]: Muralidharan et al. (2021), ‘Communication-Aware Robotics: Exploiting Motion for Communication’

[40]: Penumarthi et al. (2017), ‘Multirobot exploration for building communication maps with prior from communication models’

[41]: Otte et al. (2017), ‘Multi-robot task allocation with auctions in harsh communication environments’

[42]: Dias et al. (2006), ‘Market-Based Multirobot Coordination: A Survey and Analysis’

- ▶ Methods for predicting and preventing network partitioning before it occurs [31–33].
- ▶ Techniques for balancing aggressive exploration with conservative connectivity maintenance [30, 34].
- ▶ Mechanisms for recovering from connectivity failures when they inevitably occur [35, 36].

Challenge 3: Bridging the Gap Between Theory and Reality

How can multi-robot systems move beyond idealized geometric communication models to incorporate realistic signal propagation characteristics that reflect real-world deployment conditions? This involves:

- ▶ Developing models that capture the complexity of real-world wireless communication including path loss, fading, and interference [37, 38].
- ▶ Integrating these models into real-time decision-making processes without overwhelming computational resources [33, 39, 40].
- ▶ Validating these models using real-world data from representative deployment environments [38, 40].
- ▶ Creating systems that can adapt to varying communication conditions across different environments [8, 30].

Challenge 4: Scalable Coordination in Resource-Constrained Environments

How can coordination strategies scale effectively with team size while maintaining computational efficiency and avoiding communication overhead that could overwhelm limited bandwidth resources? This challenge encompasses:

- ▶ Designing coordination algorithms with computational complexity that scales gracefully with team size [27, 28].
- ▶ Minimizing communication overhead while maintaining effective coordination [8, 28, 41].
- ▶ Ensuring that coordination mechanisms remain effective as communication quality degrades [35, 41].
- ▶ Balancing local decision-making autonomy with global coordination requirements [26, 42].

1.3.2 Integration and Implementation Challenges

Beyond these core technical challenges, the development of a practical multi-robot exploration system must also address several integration and implementation challenges:

Real-Time Operation: Any practical solution must operate in real-time on resource-constrained robotic platforms while maintaining high reliability under uncertain and dynamic environmental conditions. This requires careful attention to computational complexity, memory usage, and algorithmic efficiency [8, 33].

Robustness and Fault Tolerance: The system must continue to operate effectively despite individual robot failures, communication link failures, sensor malfunctions, and unexpected environmental conditions. This requires redundancy, graceful degradation mechanisms, and automatic recovery capabilities [28, 41, 42].

Human-Robot Integration: While the system must operate autonomously, it must also provide appropriate interfaces for human operators to

monitor system status, adjust mission parameters, and intervene when necessary. This requires balancing autonomy with controllability and providing meaningful feedback about system state and performance [43].

Deployment Practicality: The system must be practical to deploy in real-world scenarios, considering factors such as setup time, calibration requirements, environmental adaptability, and maintenance needs. This requires moving beyond laboratory demonstrations to consider the full lifecycle of system deployment and operation [1].

[43]: Crandall et al. (2005), ‘Validating Human–Robot Interaction Schemes in Multitasking Environments’

[1]: Nagatani et al. (2009), ‘Multi-robot exploration for search and rescue missions: A report of map building in RoboCupRescue 2009’

1.4 Research Objectives and Contributions

This thesis seeks to contribute to the field of distributed multi-robot exploration by developing and evaluating two algorithmic frameworks: Dynamic Role-Based Exploration with Connectivity Maintenance (DR-BECM) and its machine learning-enhanced version, Multi-Robot Exploration via Flocking Coordination and Machine Learning-Driven Connectivity Assessment (DRBECM-ML). The research is designed to address some of the challenges identified in the previous sections and aligns with the broader objectives of the NEPHELE project.

1.4.1 Research Questions and Objectives

The main objectives of this research were structured to systematically investigate the identified challenges in distributed multi-robot exploration:

1. **Develop a Dynamic Role-Based Framework:** To design a distributed exploration algorithm where robots can dynamically switch between exploration and communication support roles. The goal was to enable this adaptation based on local information, without requiring a central coordinator.
2. **Investigate Communication Topology Strategies:** To compare dynamic communication topology strategies to identify a suitable approach for maintaining network connectivity during exploration missions. This was intended to find a balance between communication overhead and network reliability.
3. **Incorporate a Realistic Communication Model:** To explore a machine learning-based approach for using more realistic signal propagation models in the team’s decision-making process, aiming to bridge the gap between idealized models and real-world conditions.
4. **Design for Scalability and Efficiency:** To develop algorithmic solutions that remain computationally manageable as the size of the robot team increases, with a focus on algorithmic complexity and resource use.
5. **Conduct a Comprehensive Evaluation:** To perform a thorough experimental evaluation to compare the proposed methods against established baselines, using both simulations and real-world data to assess their performance.

1.4.2 Summary of Contributions

The work presented in this thesis aims to contribute to several aspects of distributed multi-robot exploration. The main contributions are summarized below.

Contribution 1: A Dynamic Role-Based Exploration Framework (DRBECM)

A distributed algorithm, DRBECM, is proposed to help manage the coordination between exploration and communication tasks.

- **Method:** The framework allows robots to autonomously switch between "explorer" and "supporter" roles. The decision is based on a local assessment of network connectivity, using properties from the Relative Neighborhood Graph (RNG) to help preserve the network. The approach integrates frontier-based exploration with flocking-inspired positioning to form adaptive communication chains that can follow the exploring robots. The algorithm also includes mechanisms for stagnation detection, recovery, and collision avoidance.
- **Performance:** In our simulation studies, this framework was compared to several other decentralized approaches. The results showed a reduction in exploration time and redundant exploration. For instance, compared to a Frontier-based method without map sharing, DRBECM showed a median exploration time that was 21.49% faster while operating in a fully decentralized manner. Its performance was comparable to that of a centralized Frontier-based approach but without the need for a central point of control.

Contribution 2: A Comparative Analysis of Communication Topologies

To inform the design of the DRBECM framework, a comparative analysis of dynamic topology strategies was conducted.

- **Method:** We developed a framework to compare three dynamic topology strategies: K-Nearest Neighbors (KNN), Relative Neighborhood Graph (RNG), and k-Relative Neighborhood Graph (K-RNG). Using NS-3 simulations, we evaluated their performance with metrics like Packet Delivery Ratio (PDR) and latency across various network densities and scales.
- **Findings:** The empirical results suggest that RNG is a suitable choice for many distributed multi-robot exploration scenarios due to its parameter-free nature and its performance in our tests. For example, in direct communication tests, RNG's PDR ranged from 95% at low density to 33% at high density, compared to K-RNG strategies which ranged from 70-80% down to 22-28%. The findings offer some guidelines for selecting a topology and provide the rationale for using RNG as the communication backbone in the DRBECM framework.

Contribution 3: A Machine Learning-Enhanced Framework (DRBECM-ML)

This contribution explores moving from simple geometric communication models to a data-driven approach by incorporating a machine learning model for signal propagation.

- **Method:** We integrated an RSSI prediction model trained on data from the FIT IoT-Lab testbed. After evaluating several regression algorithms, a Decision Tree model was selected for its balance of accuracy (R^2 score of 96.6%) and low inference time (0.0025 seconds), making it suitable for on-board robot computation. This model replaced the geometric communication model in DRBECM, allowing for role-switching and frontier selection decisions based on predicted signal strength.
- **Performance:** The use of this model in our test scenarios resulted in an increase in the exploration success rate, from 83.6% for DRBECM to 99.3% for DRBECM-ML. This improvement in mission completion was achieved while maintaining competitive exploration efficiency.

Contribution 4: Experimental Validation and Open Science

To validate the proposed methods, a series of experiments was conducted.

- **Evaluation:** The experiments compared the proposed approaches against five other multi-robot exploration algorithms across several metrics, including exploration time, efficiency, redundancy, and success rate. The statistical significance of the results was assessed over multiple runs (30-100 per configuration).
- **Reproducibility:** The work incorporated real-world communication data from the FIT IoT-Lab testbed. To support reproducibility and encourage further research, the implementation is available as an open-source project at <https://github.com/HazemCHAABI/DRBECM>, along with documentation of the experimental setup.

1.5 Thesis Organization

This thesis is organized into six chapters. It begins by establishing the context and background, then introduces and evaluates the proposed methods, and concludes with a summary and directions for future work.

Chapter 2 - Background and Related Work reviews the theoretical and empirical background for this research. This chapter surveys existing approaches to multi-robot exploration, communication topology management, and connectivity maintenance. By identifying gaps in current research, this chapter provides the motivation for the approaches developed in the thesis.

Chapter 3 - Comparative Analysis of Dynamic Communication Topologies for MRS addresses the question of network topology selection. This chapter presents an empirical comparison of three topology construction strategies: K-Nearest Neighbors (KNN), Relative Neighborhood Graph (RNG), and k-Relative Neighborhood Graph (K-RNG). The chapter's findings inform the design decisions in the rest of the thesis.

Chapter 4 - DRBECM: A Distributed Multi-Robot Exploration Approach with Connectivity Maintenance introduces the DRBECM framework. This chapter presents the algorithmic design, including the mechanisms for dynamic role assignment, exploration, and supporter positioning. Experimental evaluation is presented, comparing the framework's performance against other decentralized approaches.

Chapter 5 - DRBECM-ML: Enhancing Exploration with Machine Learning-Driven Connectivity presents an extension of the DRBECM framework that incorporates a machine learning-based communication model. The chapter details the data handling, model selection, and integration into the framework. Experimental results show an improvement in mission success rates.

Chapter 6 - Conclusion and Future Work summarizes the research, discusses how the proposed approaches address the initial research objectives, and considers the potential implications of the findings. The chapter also reflects on the limitations of the proposed methods and outlines several possible directions for future research.

1.6 List of publications

International conferences

Hazem Chaabi and Nathalie Mitton. 'Multi-Robot Exploration via Flocking Coordination and Machine Learning-Driven Connectivity Assessment'. In: *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*. 2025

Hazem Chaabi and Nathalie Mitton. 'Distributed Multi-Robot Exploration Approach with Connectivity Maintenance'. In: *International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. 2025

Workshops

Hazem Chaabi and Nathalie Mitton, "Comparative Analysis of KNN, RNG and K-RNG for Inter-Robot Communication," The 18th International Workshop on Selected Topics in Wireless and Mobile computing (STWiMob), Marrakech, Morocco, 2025

Background and Related Work

2

The field of multi-robot exploration encompasses a rich body of research spanning distributed systems, robotics, graph theory, bio-inspired coordination, and machine learning [6, 46, 47]. This chapter provides a comprehensive survey of the theoretical foundations and empirical findings that inform the development of the DRBECM and DRBECM-ML frameworks presented in this thesis.

The chapter is structured to progressively build understanding from fundamental exploration strategies through advanced coordination mechanisms. It begins with an examination of multi-robot exploration strategies, comparing centralized, decentralized, and hybrid approaches to understand their respective strengths, limitations, and applicability to different operational scenarios.

The discussion then turns to connectivity-aware exploration and maintenance, examining how multi-robot systems can balance the competing objectives of exploration efficiency and communication reliability. This section establishes the theoretical foundation for understanding exploration-communication trade-offs and reviews existing approaches to distributed control and coordination under communication constraints.

Building upon these foundations, the chapter explores bio-inspired coordination mechanisms, with particular emphasis on flocking behaviors and swarm intelligence principles that enable effective distributed coordination through simple, locally-executed rules. Finally, the chapter examines the role of machine learning in robotics, focusing on learning-based navigation, predictive connectivity models, and the unique challenges associated with distributed learning in multi-agent systems.

Through this comprehensive literature review, the chapter identifies key gaps in current research and establishes the motivation for the novel algorithmic frameworks developed in subsequent chapters.

2.1 Multi-Robot Exploration Strategies

The fundamental challenge of coordinating multiple autonomous robots to efficiently explore unknown environments has given rise to a diverse array of algorithmic approaches, each with distinct advantages and limitations. This section provides a comprehensive analysis of the three primary paradigms that have emerged in multi-robot exploration: centralized, decentralized, and hybrid approaches.

The evolution of multi-robot exploration strategies reflects the ongoing tension between the desire for optimal global coordination and the practical requirements of robustness, scalability, and real-world deployability. Centralized approaches leverage complete environmental knowledge and global optimization to achieve highly efficient exploration patterns, but suffer from inherent vulnerabilities related to single points of failure and communication bottlenecks [42, 46]. Decentralized approaches prioritize system robustness and scalability through local decision-making, but often struggle to achieve the coordination efficiency of their centralized counterparts [6, 46].

2.1	Multi-Robot Exploration Strategies	13
2.1.1	Centralized Exploration Approaches	14
2.1.2	Decentralized Exploration Approaches	15
2.1.3	Hybrid Exploration Approaches	18
2.2	Connectivity-Aware Exploration and Maintenance	20
2.2.1	Exploration-Communication Trade-Offs	20
2.2.2	Distributed Control and Coordination Mechanisms	22
2.2.3	Local Information-Based Solutions	24
2.2.4	Fault Tolerance and Robustness	25
2.3	Bio-Inspired Coordination	28
2.3.1	Flocking as a Distributed Coordination Mechanism	28
2.3.2	Swarm Intelligence Applications	29
2.4	Machine Learning in Robotics	31
2.4.1	Learning-Based Navigation and Exploration	31
2.4.2	Predictive Models for Connectivity and Communication	33
2.4.3	Challenges of Distributed Learning	34
2.5	Summary and Research Gaps	36
2.5.1	Key Insights from the Literature	36
2.5.2	Identified Research Gaps	37
2.5.3	Motivation for This Research	38

[46]: Dudek et al. (1996), 'A taxonomy for multi-agent robotics'

[48]: Parker (1998), ‘ALLIANCE: an architecture for fault tolerant multirobot cooperation’

The emergence of hybrid approaches represents an attempt to capture the benefits of both paradigms while mitigating their respective limitations. These methods combine elements of centralized planning with decentralized execution, or integrate deterministic strategies with adaptive optimization techniques to achieve more balanced performance across diverse operational scenarios [42, 48].

This section examines each paradigm in detail, analyzing their theoretical foundations, practical implementations, and empirical performance characteristics. The analysis provides the foundation for understanding how the DRBECM framework developed in this thesis builds upon the strengths of decentralized coordination while addressing the limitations that have historically constrained purely distributed approaches to multi-robot exploration.

2.1.1 Centralized Exploration Approaches

Centralized exploration approaches in multi-robot systems are characterized by the presence of a central coordinator or a global state space that orchestrates the actions of all robots in the team. In these frameworks, a central entity typically collects information from all agents, processes the global map or state, and then assigns exploration tasks or navigation goals to individual robots. This paradigm enables the exploitation of complete or near-complete knowledge of the environment and robot states, which can lead to highly efficient task allocation and path planning, especially in structured or moderately complex environments [49].

[49]: Bezcioglu et al. (2021), ‘Self-Organised Swarm Flocking with Deep Reinforcement Learning’

[50]: Stachniss et al. (2003), ‘Exploring unknown environments with mobile robots using coverage maps’

[51]: Deng et al. (2020), *Frontier-based Automatic-differentiable Information Gain Measure for Robotic Exploration of Unknown 3D Environments*

[52]: Hu et al. (2020), ‘Voronoi-Based Multi-Robot Autonomous Exploration in Unknown Environments via Deep Reinforcement Learning’

A key advantage of centralized strategies is the ability to optimize exploration based on global utility functions. Optimizing exploration is the process of strategically directing robots to map an unknown environment to maximize coverage and information gain while minimizing costs like time, travel distance, and energy [50, 51]. For instance, utility-based frontier assignment can be formulated where each robot is assigned to explore frontiers that maximize a utility function $\Omega_{ik} = \lambda d_{ik} + (1 - \lambda)\phi_{ik}$, with d_{ik} representing the distance from robot R_i to frontier k , and ϕ_{ik} denoting the distance from the frontier to the robot’s initial position. The scalar λ allows tuning between breadth-first and depth-first exploration, providing flexibility in how the exploration unfolds [52]. Such centralized assignment ensures that robots are systematically distributed across the environment, minimizing redundant coverage and improving overall efficiency. Centralized approaches also facilitate the integration of advanced machine learning techniques, such as Reinforcement Learning (RL), by leveraging a global state representation accessible to all agents. Bezcioglu et al. [49] state that in a centralized multi-agent RL setup, actor and critic networks can access a shared global state space matrix, enabling the system to learn coordinated behaviors such as flocking—a collective motion where individual agents align their velocity and position with their neighbors to move as a coherent group while avoiding collisions[53, 54]—or exploration, regardless of the density or sparsity of the robot population. This centralized learning framework can accelerate convergence and improve the quality of learned policies, as the central controller can directly observe and optimize the collective behavior of the swarm.

However, the reliance on a central coordinator introduces several challenges. The central node becomes a potential single point of failure, and the system’s scalability is inherently limited by the communication and computational bottlenecks at the central entity. As the number of robots

increases or the environment becomes more complex, the volume of data to be processed centrally can grow rapidly, leading to latency and reduced responsiveness [55, 56]. Moreover, maintaining robust communication links between all robots and the central coordinator is often impractical in real-world scenarios, especially in environments with obstacles, interference, or large spatial extents [54, 55].

Despite these limitations, centralized methods remain attractive for scenarios where communication infrastructure is reliable and the environment is not excessively large or cluttered. They are particularly effective when combined with deterministic task allocation algorithms and global optimization techniques. For example, deterministic Coordinated Multi-Robot Exploration (CME) methods can be enhanced by integrating metaheuristic algorithms, such as the Frequency-Modified Whale Optimization Algorithm (FMH-WOA), within a centralized framework. This hybridization leverages the strengths of both deterministic planning and stochastic optimization, resulting in improved exploration outcomes in complex environments [56]. Centralized strategies also provide a natural platform for incorporating predictive models and deep learning-based mapping techniques. For instance, neural network architectures can be employed to jointly perform mapping and planning, allowing the system to predict unseen regions of the environment and prioritize exploration based on expected information gain. This predictive capability, when managed centrally, can significantly enhance exploration efficiency by guiding robots toward areas with the highest potential for new discoveries [57].

In summary, centralized exploration approaches offer strong coordination and optimization capabilities by leveraging global knowledge and centralized decision-making. They are well-suited for environments where communication is reliable and the number of robots is manageable. The integration of utility-based task allocation, reinforcement learning, and hybrid optimization techniques further augments their effectiveness, although care must be taken to address scalability and robustness concerns inherent to centralized architectures [49, 52, 55–57].

While centralized methods demonstrate high efficiency by leveraging global knowledge, their architectural design introduces significant vulnerabilities. The reliance on a central controller creates a single point of failure that is ill-suited for the unpredictable nature of emergency response scenarios [58]. Our work diverges from this paradigm by embracing a fully distributed architecture. We address the challenge of balancing exploration efficiency and network robustness not through global optimization, but through a dynamic, local decision-making framework. This approach avoids the bottlenecks and failure points of centralized systems, offering a more resilient solution for real-world applications where infrastructure is unreliable or non-existent.

2.1.2 Decentralized Exploration Approaches

Decentralized exploration approaches in multi-robot systems are characterized by the absence of a central coordinator, with each robot making decisions based on local information and, in some cases, limited communication with nearby peers. This paradigm is particularly attractive for large-scale, dynamic, or communication-constrained environments, as it enhances system scalability, robustness, and adaptability to failures or environmental changes [56, 59, 60]. A fundamental advantage of decentralized strategies is their ability to operate without global knowledge or persistent connectivity. Robots typically rely on local sensing and, when

[54]: Kim (2023), ‘Leader-Based Flocking of Multiple Swarm Robots in Underwater Environments’

[55]: Shaw (2024), *Autonomous Multi-Robot Exploration Strategies for 3D Environments with Fire Detection Capabilities*

[56]: El Romeh et al. (2023), ‘Theoretical Framework and Practical Considerations for Achieving Superior Multi-Robot Exploration: Hybrid Cheetah Optimization with Intelligent Initial Configurations’

[57]: Shrestha et al. (2019), ‘Learned Map Prediction for Enhanced Mobile Robot Exploration’

[58]: Groenendaal et al. (2013), ‘A Critical Examination of the Assumptions Regarding Centralized Coordination in Large-Scale Emergency Situations’

[59]: Schilling et al. (2019), ‘Learning Vision-Based Flight in Drone Swarms by Imitation’

[60]: Ghedini et al. (2018), ‘Toward efficient adaptive ad-hoc multi-robot network topologies’

available, Ad-hoc communication to coordinate their actions. For instance, in vision-based swarms, agents can learn to coordinate their motion in three-dimensional space by mimicking bio-inspired flocking behaviors, relying solely on visual inputs rather than explicit position sharing. This approach eliminates the need for centralized control and enables fully decentralized, vision-based swarms capable of collision-free and coherent group motion [59]. The resulting trajectories are often smooth, and the system can naturally handle navigation tasks by incorporating migration terms into the predicted velocities of neural controllers.

Swarm intelligence models, such as those inspired by Particle Swarm Optimization (PSO) or virtual pheromone trails [61], have been widely adopted in decentralized exploration. These models allow robots to self-organize, balance exploration and exploitation, and adapt to noisy or uncertain environments. For example, modified PSO algorithms have been used to optimize both the control schemes and their parameters in multi-robot target tracking and collective construction tasks, leveraging local signal intensities or virtual pheromones as fitness measures for the swarm. Such bio-inspired coordination mechanisms facilitate distributed task allocation, cooperative transportation, and adaptive localization of multiple targets, even in the presence of environmental noise or partial observability [62].

Decentralized approaches are also well-suited for scenarios where robots must maintain network connectivity or system efficiency under dynamic conditions. Mechanisms that assess and optimize communication efficiency or energy consumption can be integrated into the control logic, enabling the system to adjust its topology and maintain desired efficiency properties without centralized oversight [60]. This is particularly relevant in applications where communication resources are limited or intermittent, as decentralized strategies can dynamically adapt to changing connectivity conditions and environmental constraints [63]. The integration of Machine Learning (ML), especially reinforcement learning (RL), has further advanced decentralized exploration. Distributed RL techniques enable teams of robots to learn optimal exploration strategies through direct interaction with their environment, enhancing adaptability and robustness over time [56]. For example, neural network-based controllers can map environmental features to control actions, even as the number of moving robots varies, and can directly address collision avoidance by incorporating reward functions that balance collision risk and travel time [64]. However, these methods often require extensive training and may face scalability challenges in high-dimensional or complex environments.

Hybrid decentralized approaches have emerged, combining deterministic exploration strategies with metaheuristic algorithms to leverage the strengths of both paradigms. For instance, frameworks that integrate deterministic Coordinated Multi-Robot Exploration (CME) with metaheuristic algorithms such as Frequency-Modified Whale Optimization Algorithm (FMH-WOA) or the Salp Swarm Algorithm (SSA) have demonstrated improved performance in complex environments. These hybrid methods typically use deterministic strategies for initial environment evaluation, followed by metaheuristic optimization for robot movement, resulting in more robust and effective exploration. The initial configuration of robots has also been identified as a critical factor influencing exploration efficiency, with intelligent initial placement strategies leading to significant performance gains.

Market-based decentralized approaches, inspired by economic models,

[61]: Inahara et al. (2022), 'Research on Search Algorithm Using Particle Swarm Optimization with Virtual Pheromone for Swarm Robots'

[62]: TAN et al. (2013), 'Research Advance in Swarm Robotics'

[63]: Lu et al. (2025), 'Multi-Robot Collaborative Exploration on Communication-Constrained Environments'

[64]: Han et al. (2022), 'Reinforcement Learned Distributed Multi-Robot Navigation With Reciprocal Velocity Obstacle Shaped Rewards'

have also been explored. In these methods, robots act as bidders for exploration tasks, evaluating costs and dynamically distributing tasks based on local assessments. While these techniques can improve adaptability and task distribution, they may introduce additional computational overhead and require complex negotiation protocols [56]. Decentralized exploration is particularly advantageous in real-world applications such as search and rescue, environmental monitoring, and mapping of unknown or hazardous areas. Each robot can operate independently, responding to unique local stimuli and adapting to varying sensing ranges or environmental conditions [65]. The ability to share partial maps or information, even intermittently, can further enhance collective understanding and reduce exploration time, provided robust map-merging and uncertainty-handling techniques are employed [55]. The work of Guangdong Lu et al. [63] demonstrates that decentralized frameworks can maintain high exploration efficiency even as communication conditions degrade, by prioritizing the transmission of spatio-temporal trajectory information. This adaptability is crucial for maintaining system performance in environments where communication is unreliable or costly. In summary, decentralized exploration approaches offer a flexible and robust framework for multi-robot systems operating in complex, dynamic, and communication-constrained environments. By leveraging local sensing, bio-inspired coordination, machine learning, and hybrid optimization techniques, these systems can achieve effective coordination, maintain network robustness, and adapt to a wide range of operational challenges [55, 56, 59, 60, 62–65].

[65]: Bakhshipour et al. (2017), ‘Swarm Robotics Search and Rescue; a Novel Artificial Intelligence-Inspired Optimization Approach’

Table 2.1 compares decentralized exploration approaches along six axes: *Coordinated Exploration* (explicit redundancy reduction via task/goal assignment), *Connectivity Maintenance* (active preservation of multi-hop links or related graph properties), *Robust if Communication Fails* (operation with purely local information), *Bio-Inspired* (flocking, pheromones), and *Learning-Based* (imitation/RL controllers). Two broad patterns emerge. First, methods that are intrinsically robust to outages (random walk; frontier without sharing; vision-based flocking) typically lack explicit coordination and therefore risk redundant coverage and slower convergence [10, 59]. Second, approaches that improve coordination (market/auction mechanisms or distributed RL) rely on communication and generally do not *natively* enforce connectivity, unless augmented with dedicated graph-maintenance logic [26, 42, 64, 66]. Connectivity-aware frameworks deliver stronger network guarantees but may trade off some frontier efficiency if connectivity constraints dominate motion decisions [13, 67, 68]. Hybrids partially mitigate these trade-offs by decoupling mapping and motion optimization stages [56].

Our work builds directly on the principles of decentralized exploration, where autonomy and local decision-making are paramount. While many decentralized methods like Random Walk or Frontier Exploration without Map Sharing offer robustness, they often result in uncoordinated movements and redundant exploration. We address these limitations by introducing a more structured form of coordination within a decentralized framework. By assigning dynamic roles based on local network conditions, our approach moves beyond simple uncoordinated exploration. This role-based specialization allows the system to simultaneously pursue efficient frontier-based mapping and active communication maintenance, achieving a level of coordination that is often lacking in purely decentralized schemes.

Table 2.1: Comparison of *decentralized* multi-robot exploration approaches

Approach	Coordinated Exploration	Connectivity Maintenance	Robust if Comm Fails	Bio-Inspired	Learning-Based
Random Walk (baseline) [69]	✗	✗	✓	✗	✗
Frontier Exploration (no map sharing) [10, 70]	✗	✗	✓	✗	✗
Vision-Based Flocking Swarm [59, 71]	✗	✗	✓	✓	✓
PSO / Pheromone Swarm Models [62, 72]	✗	✗	✓	✓	✗
Market-Based Task Allocation [26, 42]	✓	✗	✗	✗	✗
RL-Based Distributed Exploration [64, 66, 73]	✓	✗	✓	✗	✓
Connectivity-Aware Coordination [13, 67, 68]	✓	✓	✗	✗	✗
Deterministic + Metaheuristic Hybrid [56, 74]	✓	✗	✗	✓	✗

2.1.3 Hybrid Exploration Approaches

[75]: Julia et al. (2010), ‘A hybrid solution to the multi-robot integrated exploration problem’

Hybrid approaches in multi-robot exploration have emerged as a promising direction, aiming to combine the strengths of deterministic and metaheuristic strategies to achieve more robust and efficient exploration outcomes [56, 75]. The integration of these methods is motivated by the limitations observed in purely deterministic or purely stochastic approaches, particularly when dealing with complex, dynamic, or unknown environments. In coordinated multi-robot exploration (CME) [25, 30], deterministic strategies often provide reliable coverage and predictable behavior, but may lack adaptability and can be suboptimal in highly dynamic or cluttered scenarios. Metaheuristic algorithms, inspired by natural processes or collective animal behaviors, offer adaptability and the ability to escape local optima, but may suffer from slower convergence or lack of guarantees on coverage [62].

Recent research has focused on the development of hybrid methods that blend deterministic CME with metaheuristic optimization algorithms. For instance, the Hybrid Cheetah Exploration Technique with Intelligent Initial Configuration (HCETIIC) [56] exemplifies this trend by combining the cheetah optimization algorithm with a strategy for optimizing the initial deployment of robots. This approach has demonstrated significant improvements in exploration efficiency and robustness, highlighting the critical role of initial robot configuration in maximizing coverage and minimizing redundancy. The authors of HCETIIC indicate that the initial configuration can have a profound impact on the performance of the exploration process, especially in environments with varying complexity and communication constraints.

Hybrid methods typically employ a two-step process: an initial deterministic evaluation of the environment, followed by the application of a metaheuristic algorithm to guide robot movement. This structure allows the system to leverage the systematic coverage of deterministic methods while benefiting from the adaptive search capabilities of metaheuristics. For example, a framework integrating the deterministic CME technique [76] with the frequency-modified whale optimization algorithm (FMH-WOA) [77] has been shown to outperform conventional deterministic techniques, particularly in complex environments. The FMH-WOA mimics the predatory behavior of whales, enabling robots to adaptively explore and exploit regions of interest, thus improving overall exploration outcomes.

[76]: Albina et al. (2019), ‘Hybrid Stochastic Exploration Using Grey Wolf Optimizer and Coordinated Multi-Robot Exploration Algorithms’

[77]: Gul et al. (2021), ‘Novel Implementation of Multi-Robot Space Exploration Utilizing Coordinated Multi-Robot Exploration and Frequency Modified Whale Optimization Algorithm’

Another notable hybrid approach is the CME-SSA method, which combines deterministic CME with the salp swarm algorithm (SSA) [56]. The SSA is inspired by the swarming behavior of salps, marine organisms known for their efficient foraging strategies. In this method, the deterministic component manages cost and utility values on the grid map, while the SSA component optimizes the movement of robots to enhance coverage and coordination. This combination has been shown to improve temporal planning and inter-agent coordination, especially in unfamiliar or cluttered environments. The integration of reinforcement learning into hybrid exploration frameworks further enhances adaptability and learning capabilities [49, 64]. RL-based hybrid methods enable robots to learn optimal exploration strategies through interaction with the environment, adjusting their behavior based on feedback and experience. However, these methods often require extensive training and may encounter scalability challenges in large or highly dynamic environments. Distributed reinforcement learning, in particular, allows teams of robots to learn collectively, improving robustness and adaptability over time, but can suffer from the curse of dimensionality as the number of agents and environmental complexity increases [56].

Hybrid approaches are not limited to the combination of deterministic and metaheuristic algorithms, they also encompass the integration of centralized and decentralized strategies. For example, some systems utilize a centralized global state for initial planning, followed by decentralized execution to enhance scalability and robustness. This is particularly relevant in scenarios with restricted communication ranges or dynamic obstacles, where maintaining a central coordinator is impractical. By combining centralized planning with decentralized adaptation, hybrid systems can achieve a balance between global optimality and local responsiveness [49, 55].

The effectiveness of hybrid approaches is influenced by several factors, including robot capabilities, environment complexity, communication mechanisms, and the choice of optimization algorithms. Comparative analyses have demonstrated that hybrid methods can outperform both purely deterministic and purely metaheuristic strategies, particularly in terms of exploration rate, coverage uniformity, and adaptability to changing conditions [56]. Jamshidpey et al. [78] state that hybrid formation approaches achieve more uniform coverage than purely decentralized methods, though predetermined strategies may still offer the highest uniformity in certain scenarios. In summary, hybrid approaches in multi-robot exploration represent a synthesis of deterministic, metaheuristic, and learning-based strategies, often incorporating both centralized and decentralized elements. These methods are characterized by their ability to adapt to complex, dynamic environments, optimize coverage and resource utilization, and maintain robustness in the face of uncertainty. The ongoing development of hybrid techniques, including the refinement of initial configurations, integration of advanced optimization algorithms, and incorporation of machine learning, continues to push the boundaries of what multi-robot systems can achieve in real-world exploration tasks [56].

[78]: Jamshidpey et al. (2024), 'Centralization vs. decentralization in multi-robot sweep coverage with ground robots and UAVs'

Our proposed framework, DRBECM-ML, can be classified as a novel hybrid approach that uniquely integrates three distinct components: established frontier-based exploration, bio-inspired flocking dynamics, and machine learning-driven connectivity prediction. Unlike the hybrid models described, which primarily blend deterministic and metaheuristic algorithms for navigation (e.g., HCETIIC), our system creates a hybrid at a behavioral level. Robots are not locked into a

single strategy but dynamically switch between exploration-focused behavior (frontier-based) and connectivity-reinforcement behavior (flocking-inspired).

2.2 Connectivity-Aware Exploration and Maintenance

The exploration of unknown environments by multi-robot systems presents fundamental challenges that extend beyond mere coverage and mapping. In many practical applications, particularly those involving emergency response, search and rescue, and environmental monitoring, maintaining reliable communication connectivity among team members is equally critical as the exploration objective itself. This section examines the intricate relationship between exploration efficiency and communication maintenance, exploring how multi-robot systems can effectively balance these competing requirements.

The connectivity-aware exploration paradigm recognizes that robot teams operating in unknown environments must simultaneously achieve two potentially conflicting objectives: maximizing information acquisition through aggressive exploration while preserving network integrity to enable coordination and data sharing. This dual requirement introduces complex trade-offs that have motivated extensive research into distributed control mechanisms, local information-based solutions, and robust coordination strategies that can operate effectively under communication constraints and environmental uncertainties.

2.2.1 Exploration-Communication Trade-Offs

Exploration-communication trade-offs are a central concern in the design of multi-robot systems, particularly when operating in environments where both efficient area coverage and reliable inter-robot connectivity are required. The interplay between these objectives often leads to competing demands: maximizing exploration speed and coverage can necessitate robots spreading out, while maintaining robust communication typically requires them to remain within a certain proximity or network topology.

Centralized and decentralized approaches to multi-robot coordination each present distinct trade-offs in this context. Centralized strategies can achieve globally optimized solutions for exploration, as a central planner can allocate tasks and routes to maximize coverage efficiency. However, this comes at the cost of high communication overhead, since the central entity must maintain constant contact with all robots, leading to increased bandwidth requirements and a single point of failure. In contrast, decentralized approaches reduce communication costs by relying on local interactions, which enhances robustness and adaptability to dynamic environments but often results in suboptimal global exploration performance due to the lack of a unified view [79]. Hybrid approaches attempt to balance these extremes, leveraging partial centralization or hierarchical structures to improve scalability and resilience [74].

[79]: Poudel et al. (2022), ‘Decentralized and Centralized Planning for Multi-Robot Additive Manufacturing’

[74]: Mohamed et al. (2018), ‘A Hybrid Decentralized Coordinated Approach for Multi-Robot Exploration Task’

Connectivity-aware exploration methods explicitly address the need to maintain network integrity while exploring. These strategies often employ constraints that ensure each robot remains within communication range of a minimum number of peers, or that the overall network graph remains connected. For example, potential field-based control laws can be

used, where robots are repelled by nearby agents and obstacles, leading to an optimized spatial distribution that respects both coverage and connectivity requirements [80]. The trade-off here is evident: increasing the minimum required connectivity can limit the area that can be explored in a given time, as robots are forced to stay closer together, while relaxing connectivity constraints can risk network fragmentation and loss of coordination.

The structure of the communication network itself plays a significant role in these trade-offs. Sparse networks, with a minimal number of links, are energy-efficient and reduce communication overhead, but are more vulnerable to failures, loss of a single link can disconnect the network. Conversely, denser networks enhance robustness and security but at the cost of increased energy consumption and communication complexity. The optimal network topology thus depends on the number of robots, the intended application, and the required speed of consensus or information propagation. Griparic et al. [81] indicate that models inspired by real-world social networks, such as the Albert-Barabási preferential attachment model, can be used to dynamically adjust connectivity in a way that balances robustness and efficiency. Decentralized estimation and maintenance of network connectivity have been advanced through distributed algorithms that allow each robot to estimate global properties, such as the algebraic connectivity of the network, using only local information. These methods enable robots to make informed decisions about movement and communication, trading off between the precision of connectivity estimation and the computational or communication resources expended [82]. Such decentralized strategies are particularly well-suited for dynamic environments, where the network topology may change frequently due to robot movement or environmental obstacles. Energy consumption is another critical factor in the exploration-communication trade-off. Algorithms that minimize unnecessary movement and avoid collisions not only preserve energy but also reduce the risk of communication loss due to robots straying out of range.

However, in complex environments, strategies that prioritize connectivity may require additional steps or detours, increasing the total energy expenditure and exploration time. The relationship between team size and exploration efficiency is also non-linear, while adding more robots can decrease exploration time, it can also increase the complexity of maintaining connectivity and coordinating actions [74].

Machine learning techniques, such as reinforcement learning and predictive models, are increasingly being integrated into multi-robot exploration frameworks to dynamically balance exploration and communication objectives. These approaches can learn to assign credit to individual robots for their contributions to exploration, adapt to varying field-of-view constraints, and optimize long-term outcomes rather than short-term gains [83]. By leveraging attention mechanisms and hierarchical control architectures, these systems can intelligently fuse information about frontiers, robot orientations, and connectivity status to make context-aware decisions [52, 83].

Bio-inspired coordination mechanisms, including swarm intelligence and flocking behaviors, offer additional strategies for managing exploration-communication trade-offs. Swarm algorithms often rely on simple local rules that lead to emergent global behaviors, such as aggregation or dispersion, which can be tuned to maintain connectivity while maximizing coverage [72]. For instance, robots may adjust the randomness of their

[80]: Bayındır (2015), 'A Review of Swarm Robotics Tasks'

[81]: Griparic et al. (2022), 'Consensus-Based Distributed Connectivity Control in Multi-Agent Systems'

[82]: Gasparri et al. (2017), 'Bounded Control Law for Global Connectivity Maintenance in Cooperative Multirobot Systems'

[83]: Chiun et al. (2025), *MARVEL: Multi-Agent Reinforcement Learning for constrained field-of-View multi-robot Exploration in Large-scale environments*

movement based on local density, settling in high-density regions to maintain communication links, or use evolved neural controllers to aggregate when within sensing range of peers [80]. These methods are inherently robust to individual failures and can adapt to changing environmental conditions. The integration of connectivity-aware exploration strategies with robust communication maintenance mechanisms is essential for effective multi-robot operation in complex, dynamic environments. The challenge lies in designing algorithms and network structures that can dynamically adapt to changing conditions, balancing the competing demands of exploration efficiency, communication reliability, energy consumption, and robustness [72, 84]. The ongoing development of distributed estimation, bio-inspired coordination, and machine learning-enhanced control architectures continues to push the boundaries of what is achievable in multi-robot exploration.

[72]: Duan et al. (2023), 'From animal collective behaviors to swarm robotic co-operation'

[84]: Guo et al. (2021), 'Imitation Learning with Graph Neural Networks for Improving Swarm Robustness under Restricted Communications'

Our research confronts the exploration-communication trade-off not as a static constraint but as a dynamic problem to be actively managed. Rather than simply constraining exploration to maintain connectivity, our framework introduces a dynamic role-assignment mechanism that explicitly addresses this balance.

2.2.2 Distributed Control and Coordination Mechanisms

Distributed control and coordination mechanisms are fundamental to enabling robust and scalable multi-robot exploration, particularly in scenarios where maintaining network connectivity is essential for effective operation. In distributed systems, each robot typically relies on local information and interactions with neighboring agents, which reduces the need for centralized control and lowers communication overhead. This approach is inspired by natural swarms, such as flocks of birds or schools of fish, where collective behavior emerges from simple local rules and interactions, leading to self-organized and adaptive group dynamics. A key advantage of distributed mechanisms is their ability to maintain network connectivity even as the swarm size increases or the environment becomes more complex. In self-organizing network architectures, robots communicate primarily with their immediate neighbors, allowing the system to scale efficiently and adapt to dynamic changes in topology. The selection of leader agents within the swarm can further enhance coordination, enabling the group to accomplish specific tasks without requiring all members to maintain direct links to a central controller [53].

[53]: Liu et al. (2020), 'Leader-following flocking for unmanned aerial vehicle swarm with distributed topology control'

[85]: Ghedini et al. (2017), 'Toward fault-tolerant multi-robot networks'

Ghedini et al. [85] state that algebraic connectivity maintenance laws can be employed to ensure that the communication graph remains connected, even as robots move and the network topology evolves. By representing the communication network as an undirected graph, it becomes possible to monitor and control the system's connectivity properties in real time.

Distributed strategies also benefit from advanced graph-based methods, such as importance-correlated directed graph convolutional networks (IDGCNs) [84], which facilitate multidimensional feature extraction and aggregation under restricted communication conditions. These methods enable the system to maintain a backbone network for robust information flow, even as the swarm's topology and scale change over time. The use of weighted link contraction degree matrices allows for real-time assessment of graph importance, supporting the maintenance of network robustness and scalability. Furthermore, imitation learning can be leveraged to generate distributed strategies that closely approximate the performance of centralized expert templates, ensuring balanced

distribution of communication and control responsibilities across the swarm [84].

In the context of exploration, distributed coordination mechanisms must balance the competing objectives of maximizing area coverage and preserving network connectivity. Robots often need to make local decisions about movement and task allocation based on partial observations, which can lead to suboptimal global outcomes if not properly coordinated. To address this, distributed algorithms may incorporate communication protocols that allow agents to share critical information about their local environment, enabling more informed decision-making and improved collective performance. The use of mean embedding distributions, where each agent is treated as a sample, has been shown to enhance the representation of global and local states, supporting more effective coordination in pursuit-evasion and rendezvous tasks [49].

Bio-inspired approaches, such as flocking and swarm intelligence, provide additional insights into distributed coordination. These methods draw on principles observed in natural systems, where simple behavioral rules, such as alignment, cohesion, and separation, enable groups of agents to move cohesively while avoiding collisions and maintaining connectivity [53]. By emulating these behaviors, multi-robot systems can achieve robust and adaptive exploration in complex environments without relying on centralized control.

Machine learning techniques, particularly reinforcement learning (RL), have been integrated into distributed control frameworks to further enhance adaptability and performance. RL-based approaches allow robots to learn optimal policies for navigation and exploration through trial and error, even in the presence of imperfect sensing and dynamic obstacles. In distributed settings, connectivity constraints can be incorporated directly into the learning process, ensuring that robots maintain communication links while exploring unknown environments. The authors of [86] indicate that providing navigation demonstrations and imposing connectivity constraints during RL training can yield flexible policies capable of handling diverse sensor configurations and environmental challenges. Distributed control mechanisms are also designed to be resilient to failures and communication disruptions. By relying on local interactions and decentralized decision-making, the system can continue to function effectively even if individual robots lose connectivity or become inoperative. This robustness is critical for real-world applications, where unpredictable events and environmental hazards are common [53, 85].

In summary, distributed control and coordination mechanisms offer a scalable, robust, and adaptive framework for multi-robot exploration, particularly in connectivity-aware scenarios. By leveraging local interactions, advanced graph-based methods, bio-inspired coordination, and machine learning, these systems can maintain network robustness and achieve effective exploration in complex, dynamic environments [49, 53, 84–86].

Our proposed DRBECM and DRBECM-ML frameworks build directly upon these distributed control principles. Rather than relying on centralized coordination or global state information, our approach employs a fully decentralized architecture where each robot makes autonomous decisions based solely on local information from its immediate neighbors. The dynamic role-assignment mechanism at the core of our framework exemplifies distributed coordination, allowing robots to adaptively switch between exploration-focused and connectivity-maintenance behaviors based on local network conditions. This design philosophy

[86]: Li et al. (2022), ‘Decentralized Global Connectivity Maintenance for Multi-Robot Navigation: A Reinforcement Learning Approach’

aligns with the scalability and robustness advantages of distributed systems while addressing the specific challenge of maintaining network connectivity during exploration missions in unpredictable environments.

2.2.3 Local Information-Based Solutions

Local information-based solutions in multi-robot systems are characterized by agents making decisions based solely on data available from their immediate environment or direct neighbors, without access to global knowledge of the system. This approach is particularly relevant in connectivity-aware exploration, where maintaining network connectivity and robustness must be achieved under constraints of limited communication and computational resources. In decentralized multi-agent systems, each agent typically possesses only partial knowledge of the overall network topology, being aware only of its direct connections. This limitation complicates the task of maintaining or optimizing global properties such as algebraic connectivity, as agents cannot directly assess the state of the entire communication graph. Instead, strategies are developed to ensure that local actions collectively preserve desired connectivity properties, often by maintaining a certain level of local connectivity or by using consensus algorithms that rely on information redundancy within local neighborhoods. Griparic et al. [81] indicate that consensus-based methods are widely adopted for real-time, distributed connectivity maintenance, leveraging local information exchange and redundancy to enhance resilience against network disruptions. The challenge of maximizing algebraic connectivity in weighted undirected graphs, where edge weights represent inter-agent relationships, is further exacerbated by the limited computational and power capabilities of individual agents. As a result, the focus shifts from global maximization to maintaining connectivity above a critical threshold, balancing the trade-off between network robustness and communication overhead. Excessive efforts to maximize connectivity can introduce delays and redundant messaging, which may degrade overall system performance [81].

Local information-based solutions are also prominent in swarm robotics, where coordination emerges from simple rules and local interactions rather than centralized control. For example, in collaborative manipulation tasks inspired by ant behavior, robots coordinate object transport using only local cues and interactions, achieving robust group behavior without global oversight. This decentralized approach enhances fault tolerance and adaptability, as the system does not rely on a single point of failure and can function effectively with simple agents. Baygındı et al. [80] state that local reinforcement paradigms, where rewards are assigned only to agents directly accomplishing objectives, align well with swarm intelligence principles and reduce the need for global information sharing.

In connectivity-aware exploration, local information is often used to guide motion planning and maintain network structure. For instance, when planning trajectories, robots may discretize their reachable sets and select feasible paths that ensure the induced communication graph remains connected after movement. This process relies on local observations and motion primitives, allowing agents to adapt their actions based on immediate connectivity constraints [87]. The use of local heuristics, such as scoring potential intermediate nodes based on environmental features, further supports navigation and exploration in unknown or dynamic

[87]: Shi et al. (2021), 'Communication-Aware Multi-robot Coordination with Submodular Maximization'

environments, especially when global information is unavailable or unreliable [88].

Distributed control strategies based on local information, while scalable and robust, can be susceptible to suboptimal performance due to limited observation and communication ranges. This can result in the system becoming trapped in local optima, highlighting a performance gap compared to centralized approaches. To address this, hybrid methods such as imitation learning have been explored, where distributed agents learn to mimic centralized expert strategies using only local observations, thereby improving the global effectiveness of local policies [84].

In practical applications, local information-based solutions are essential for maintaining connectivity and robustness in environments with communication constraints, such as underground tunnels or large-scale manufacturing floors. Robots may retrieve missing information from neighbors within communication range, assemble partial maps, or adapt their behavior to compensate for dynamic changes and uncertainties in the environment [63, 79]. The absence of global knowledge and reliance on local interactions not only supports scalability but also enhances the system's ability to cope with failures and dynamic conditions.

Bio-inspired coordination mechanisms, such as those observed in ant colonies or flocking behaviors, further exemplify the power of local information-based strategies. These systems achieve complex group objectives through simple, locally informed rules, demonstrating that effective global coordination can emerge from decentralized, local interactions. However, certain tasks, such as clustering, may experience diminishing returns with increased agent numbers due to inter-robot interference, underscoring the importance of carefully designing local rules to balance cooperation and competition [80]. Machine learning techniques, particularly reinforcement learning, have been integrated with local information-based approaches to enhance adaptability and performance. For example, parameter sharing schemes and multi-agent policy training enable robots to learn collision avoidance and socially-aware navigation policies using only observable states of nearby agents and environmental cues [86]. These methods allow for efficient policy learning and execution in complex, dynamic settings, further bridging the gap between local decision-making and global system objectives. Overall, local information-based solutions form the backbone of scalable, robust, and adaptive multi-robot systems, enabling effective connectivity-aware exploration and maintenance in scenarios where global information is inaccessible or impractical to obtain [79–81, 87].

Our DRBECM and DRBECM-ML frameworks are fundamentally local information-based solutions. Each robot's decisions are made autonomously based on its own sensor readings and information received from its direct neighbors.

[88]: Ou et al. (2024), 'Autonomous Navigation by Mobile Robot with Sensor Fusion Based on Deep Reinforcement Learning'

2.2.4 Fault Tolerance and Robustness

Fault tolerance and robustness are fundamental requirements for multi-robot systems operating in dynamic and uncertain environments, particularly when connectivity-aware exploration and maintenance are critical. The ability of a multi-robot network to withstand failures, whether due to individual robot malfunctions, communication breakdowns, or environmental disturbances, directly impacts the system's effectiveness and reliability. A key aspect of robustness in multi-robot exploration is the maintenance of network connectivity under adverse conditions.

The work of Cinara Ghedini et al. [60] highlights that effective inter-robot communication is essential for coordinated exploration, and that network topologies must be designed to address not only connectivity maintenance but also collision avoidance, fault tolerance, and area coverage improvement. Their model demonstrates that simultaneous control strategies can yield resilient networks capable of enhancing sensing coverage while avoiding collisions and maintaining connectivity, even in the presence of faults. This approach is experimentally validated in both fault-free and fault-prone scenarios, showing that the system can adapt to failures without significant degradation in performance.

The probabilistic control law proposed by Ghedini et al. [60] further supports robustness by dynamically driving vulnerable robots toward the barycenter of their local group, thereby reducing the risk of network fragmentation. This decentralized methodology allows each robot to evaluate its own vulnerability and act accordingly, increasing the likelihood of maintaining a connected network even as individual robots experience failures or communication losses [85]. Such decentralized strategies are particularly advantageous in large-scale systems, where centralized control may become a bottleneck or a single point of failure.

From a graph-theoretic perspective, maintaining the algebraic connectivity of the communication graph is a widely adopted strategy for ensuring robustness. Sabattini et al. [89] discuss decentralized connectivity maintenance control strategies that exploit local estimations of the second smallest eigenvalue ($\tilde{\lambda}_2$) of the Laplacian matrix, which serves as a measure of network connectivity. Even though the exact value of $\tilde{\lambda}_2$ may not be available to each robot, partial derivatives can be computed locally, enabling decentralized implementation of connectivity-preserving control. This approach ensures that the network remains robust to node or link failures by continuously adapting the robots' positions to maximize connectivity.

In addition to connectivity maintenance, fault tolerance is enhanced by strategies that allow the system to continue functioning despite the loss or malfunction of individual robots. Shaw emphasizes the importance of developing exploration strategies that enable the remaining robots to adapt and continue their tasks in the event of failures. This adaptability is crucial for increasing the overall resilience of the system, especially in environments that are dynamic or time-varying, where obstacles and conditions may change unpredictably [55]. The integration of such strategies ensures that the exploration process is not halted by isolated failures, but rather that the system can reconfigure and redistribute tasks as needed.

Robustness is also closely linked to the ability to detect and avoid collisions in decentralized settings. Nagavarapu et al. [90] propose a technique for decentralized inter-robot collision avoidance that relies on minimal sensing and communication, using only bearing information to detect potential collisions before they occur. This approach reduces the dependency on high-bandwidth communication and complex sensing, making the system more robust to communication constraints and sensor failure. The resilience of multi-robot systems is further supported by the use of self-organizing principles inspired by biological swarms. Duarte et al. [91] note that swarm robotic systems inherently possess high levels of scalability, versatility, and robustness to individual failures, as collective behavior emerges from simple local interactions. However, they also point out that real-world deployment of such systems remains

[89]: Sabattini et al. (2011), 'On decentralized connectivity maintenance for mobile robotic systems'

[90]: Nagavarapu et al. (2016), 'Multi-Robot Graph Exploration and Map Building with Collision Avoidance: A Decentralized Approach'

[91]: Duarte et al. (2016), 'Evolution of Collective Behaviors for a Real Swarm of Aquatic Surface Robots'

challenging, with most advances demonstrated in controlled laboratory environments.

This observation underscores the need for continued research into robust coordination mechanisms that can operate reliably outside of idealized settings. Connectivity-aware exploration strategies must also account for communication constraints, as highlighted by Yang et al. Their experiments show that maintaining strong connectivity often imposes motion constraints on the robots, which can limit their ability to perform exploration tasks efficiently. Balancing the trade-off between connectivity and task performance is therefore essential for robust operation, especially when communication ranges are limited or variable [68].

The integration of machine learning techniques, such as reinforcement learning and predictive models, offers additional avenues for enhancing robustness. Lu et al. [63] discuss the potential for curriculum learning and the introduction of malfunctioning robots during training to improve the system's ability to handle failures and adapt to communication bottlenecks. These approaches can lead to more resilient exploration strategies that maintain performance even as the environment or the robot team changes.

Hybrid approaches that combine centralized and decentralized elements can also contribute to fault tolerance. For example, Mohamed et al. describe a system in which robots construct and broadcast global maps, allowing for distributed decision-making and redundancy in information sharing. This ensures that even if some robots fail, the remaining agents can access the necessary information to continue exploration [74]. Finally, the importance of robust initial configurations is emphasized by El Romeh and Mirjalili [56], who propose an optimization-based approach to selecting starting positions for robots. By leveraging advanced optimization algorithms, such as the cheetah optimizer, their method aims to maximize the effectiveness of exploration from the outset, thereby reducing the impact of early failures or suboptimal deployments. Taken together, these strategies illustrate that fault tolerance and robustness in connectivity-aware multi-robot exploration are achieved through a combination of decentralized control, adaptive communication, collision avoidance, self-organization, and intelligent task redistribution. The interplay of these mechanisms ensures that multi-robot systems can maintain effective coordination and performance, even in the face of failures and dynamic environmental challenges [92, 93].

Robustness is a core design principle of our framework, emerging directly from its decentralized and adaptive nature. Unlike centralized systems, the failure of a single robot does not cause the entire system to collapse. If an explorer fails, other explorers will eventually cover its assigned frontier. If a supporter fails, the flocking-inspired behavior of other supporters and the dynamic role-switching mechanism will allow the network to reconfigure and bridge the resulting communication gap. Furthermore, our use of RNG for neighbor selection contributes to robustness by creating a sparse network; the system is not overly reliant on any single link. This inherent redundancy and adaptability provide a significant degree of fault tolerance, which is critical for the intended application in emergency response scenarios.

[68]: Yang et al. (2024), 'Integrating On-line Learning and Connectivity Maintenance for Communication-Aware Multi-Robot Coordination'

[92]: Bai et al. (2025), *Realm: Real-Time Line-of-Sight Maintenance in Multi-Robot Navigation with Unknown Obstacles*

[93]: Tran et al. (2022), 'Frontier-led swarming: Robust multi-robot coverage of unknown environments'

2.3 Bio-Inspired Coordination

Nature provides compelling examples of collective intelligence through the coordinated behavior of biological systems such as bird flocks, fish schools, and insect swarms. These natural phenomena demonstrate how simple, locally-executed rules can give rise to sophisticated group behaviors that exhibit remarkable adaptability, robustness, and efficiency. The study of bio-inspired coordination in multi-robot systems seeks to harness these principles to create artificial systems capable of achieving complex collective objectives through decentralized decision-making and local interactions.

This section explores the application of bio-inspired coordination mechanisms in multi-robot exploration, with particular emphasis on flocking behaviors and swarm intelligence. These approaches offer significant advantages in terms of scalability, fault tolerance, and computational simplicity, making them particularly well-suited for resource-constrained environments where centralized control is impractical or impossible. The investigation of these natural coordination strategies provides the foundation for understanding how distributed multi-robot systems can achieve coherent group behavior while maintaining individual autonomy and adaptability.

2.3.1 Flocking as a Distributed Coordination Mechanism

Flocking represents a distributed coordination mechanism inspired by the collective behaviors observed in natural systems, such as bird flocks and fish schools. In multi-robot systems, flocking enables a group of relatively simple agents to achieve complex, coordinated movement through local interactions and decentralized control. Each robot, or agent, typically follows a set of simple rules based on the positions and velocities of its neighbors, resulting in emergent group-level behaviors that are robust and adaptive to environmental changes [62].

The fundamental principles of flocking in robotics are rooted in the concept of self-organization, where global order arises from local interactions without centralized oversight. Individual robots are often equipped with only basic sensing and communication capabilities, allowing them to perceive the relative positions of nearby agents and adjust their own motion accordingly. This approach not only reduces the computational and communication burden on each robot but also enhances the scalability and fault tolerance of the system [94]. The emergent intelligence of the swarm is not a property of any single agent but rather a consequence of their collective interactions, enabling the group to adapt to dynamic environments and unforeseen obstacles [62, 94]. Distributed flocking control strategies are typically designed to achieve objectives such as velocity alignment, collision avoidance, and cohesion. For instance, robots may adjust their velocities to match those of their neighbors, maintain a preferred distance to avoid collisions, and steer towards the average position of the group to preserve cohesion. These behaviors can be encoded using artificial potential fields or consensus algorithms, which are computed locally by each agent based on information from its immediate neighbors [84]. The decentralized nature of these computations ensures that the system remains operational even in the presence of communication dropouts or individual robot failures, as each agent relies only on local information [94].

[94]: Aggravi et al. (2021), 'Connectivity-Maintenance Teleoperation of a UAV Fleet With Wearable Haptic Feedback'

Imitative learning from centralized expert templates has been explored to enhance the robustness of distributed flocking strategies. In such approaches, a centralized controller is used to generate optimal behaviors in a fully connected, small-scale swarm, and these behaviors are then distilled into distributed policies that can be executed by individual robots with only local information [84]. This method leverages the strengths of centralized optimization while retaining the scalability and resilience of distributed control. The adaptability of flocking mechanisms is particularly advantageous in complex and cluttered environments. For example, when robots encounter obstacles or hazardous configurations, local interaction rules can prompt them to adjust their trajectories, such as by changing altitude or lateral position, to maintain group cohesion and safety. Consensus on leader velocity is another feature that can be integrated into flocking frameworks, allowing the entire group to track a designated leader's motion and avoid excessive stretching or fragmentation, especially in challenging terrains. Flocking as a distributed coordination mechanism is not limited to theoretical constructs; it has been validated in practical scenarios such as surveillance, search-and-rescue, and disaster response, where robustness, scalability, and adaptability are critical [94]. The ability of simple agents to collectively achieve sophisticated behaviors through local rules and interactions underscores the potential of bio-inspired approaches in advancing the capabilities of multi-robot systems [62, 84, 94].

We draw inspiration directly from flocking behaviors, but apply them in a novel context: not for general group cohesion, but specifically for the task of dynamic connectivity maintenance. In our framework, only the robots in the supporter role employ a flocking-inspired algorithm. Their objective is not simply to move as a cohesive group, but to position themselves as an adaptive relay chain that connects the forward-positioned explorers back to the base station. The movement of a supporter is influenced by the positions of both nearby supporters and the explorers they are supporting. This specialized application of flocking principles as a targeted mechanism for network support, rather than a general mobility model for the entire swarm, is a key distinction of our work.

2.3.2 Swarm Intelligence Applications

Swarm intelligence has emerged as a foundational paradigm for coordinating large groups of autonomous robots, drawing inspiration from the collective behaviors observed in natural systems such as flocks of birds, schools of fish, and colonies of social insects. The core principle underlying swarm intelligence is the reliance on simple local rules and decentralized control, which collectively give rise to complex, adaptive, and robust group behaviors without the need for centralized oversight [80, 95, 96]. One of the most prominent applications of swarm intelligence in multi-robot systems is the implementation of consensus and agreement protocols. In these approaches, individual robots iteratively update their internal states, such as heading, position, or behavioral mode, based on information exchanged with their immediate neighbors. This process enables the entire swarm to converge towards a common objective or shared understanding of the environment, even in the presence of uncertainty or incomplete information [95, 97]. Consensus-based strategies are particularly advantageous in scenarios where the environment is dynamic or partially observable, as they allow the swarm to adaptively coordinate actions and share critical information in real time.

[95]: Antonio Luca et al. (2018), 'Design and simulation of the emergent behavior of small drones swarming for distributed target localization'

[96]: Zhang et al. (2008), 'Collective behavior coordination with predictive mechanisms'

[97]: Giordano et al. (2013), 'A passivity-based decentralized strategy for generalized connectivity maintenance'

Swarm intelligence also facilitates advanced collective behaviors such as self-assembly and morphogenesis. In these scenarios, robots physically connect to form larger structures or adapt their configuration to accomplish tasks that would be infeasible for individual agents. The s-bots project exemplifies this capability, where robots autonomously decide when and how to attach to one another using only local cues, such as colored LEDs, to signal attachment readiness. This self-assembly has been leveraged for collaborative object transportation, navigation in hazardous environments, and overcoming obstacles like holes in the terrain [80]. The ability to autonomously form and reconfigure physical structures enhances the versatility and resilience of robotic swarms in complex environments.

Bio-inspired coordination techniques extend beyond physical assembly to encompass collective search and exploration. Particle Swarm Optimization (PSO), originally developed as a computational algorithm inspired by social foraging behaviors, has been adapted for direct use in swarm robotics. In this context, robots use fitness-based strategies to search for targets, dynamically adjusting their trajectories based on both their own experiences and those of their peers [62]. Such approaches enable efficient coverage and target localization in large, unknown environments, leveraging the distributed sensing and decision-making capabilities of the swarm.

The integration of machine learning, particularly reinforcement learning and predictive modeling, has further expanded the potential of swarm intelligence in multi-robot systems. Reinforcement learning architectures have been employed to optimize agent trajectories in real time, allowing swarms to adapt to model mismatches and significant environmental disturbances. These methods have been validated both in simulation and in physical experiments with quadrotor flocks navigating three-dimensional spaces [98]. Predictive mechanisms, such as decentralized model predictive control (MPC), enable each robot to estimate the future states of its neighbors using only locally available information. This enhances the swarm's ability to coordinate motion and maintain cohesion without requiring global knowledge of the system state [96]. The use of predictive intelligence, inspired by observations of natural flocks, has been shown to improve the robustness and adaptability of collective behaviors, especially in leaderless groups.

Swarm intelligence also plays a crucial role in maintaining network connectivity and robustness during exploration tasks. Distributed agreement protocols and connectivity-aware strategies ensure that information can propagate throughout the swarm, even as individual robots move or encounter communication constraints [13, 95]. By encoding both physical interaction qualities and higher-level behavioral constraints into the inter-agent communication model, swarms can dynamically balance the competing demands of exploration, task execution, and network maintenance [13, 97]. This adaptability is essential for effective operation in environments characterized by obstacles, variable radio propagation, and fluctuating network loads. The application of swarm intelligence in multi-robot systems thus encompasses a wide spectrum of capabilities, from consensus formation and self-assembly to adaptive exploration and robust communication. These bio-inspired strategies, augmented by advances in machine learning and predictive modeling, provide a powerful framework for achieving scalable, resilient, and efficient coordination in complex, dynamic environments [80, 95, 96, 98].

Our DRBECM and DRBECM-ML frameworks are practical applications of

[98]: Beaver et al. (2021), 'An overview on optimal flocking'

swarm intelligence principles, designed for the specific task of exploration with connectivity maintenance. The emergent behavior of the swarm is not centrally planned but arises from a set of simple, locally-executed rules.

2.4 Machine Learning in Robotics

The integration of machine learning techniques into robotic systems has emerged as a transformative approach for addressing the complexity and uncertainty inherent in real-world environments. In the context of multi-robot exploration, machine learning offers powerful tools for enhancing decision-making capabilities, improving environmental understanding, and adapting to dynamic conditions that traditional algorithmic approaches may struggle to handle effectively.

This section examines the role of machine learning in robotic exploration, focusing on learning-based navigation strategies, predictive models for connectivity assessment, and the unique challenges associated with distributed learning in multi-agent systems. The discussion encompasses both the opportunities presented by data-driven approaches and the practical considerations that must be addressed when deploying machine learning models in resource-constrained, real-time robotic applications.

2.4.1 Learning-Based Navigation and Exploration

Learning-based navigation and exploration have become central to advancing multi-robot systems, particularly as environments grow in complexity and unpredictability. The integration of machine learning, especially Deep Reinforcement Learning (DRL), has enabled robots to autonomously acquire navigation and exploration policies that adapt to dynamic and partially observable settings. The versatility of neural networks, combined with the ability to optimize behavior based on reward-driven feedback, has attracted significant research attention for applications in control, localization, and planning. DRL approaches allow robots to learn navigation strategies directly from sensory data, such as camera images, without requiring explicit environmental models. For instance, policy models trained with Double Deep Q-Network (DDQN) can guide robots through unknown spaces by mapping visual input to discrete action choices, demonstrating the potential of end-to-end learning for robust navigation [99].

The use of learned generative models, such as variational autoencoders (VAEs), further enhances exploration by enabling robots to predict and infer the structure of unseen regions in an environment. By generating interpretable map outputs, these models support more efficient path planning and facilitate human understanding of the robot's decision-making process. The separation of the generative mapping component from the planning module allows for modular system design, where the learned model provides probabilistic predictions that inform exploration strategies. This approach not only improves coverage efficiency but also offers a transparent framework for integrating information-theoretic objectives into navigation tasks [57].

Reward design is a critical aspect of learning-based exploration. The reward function must balance safety, efficiency, and task completion. For example, a terminal action can be introduced to halt exploration once a predefined coverage threshold is achieved, with positive rewards

[99]: Miranda et al. (2024), 'Generalization in Deep Reinforcement Learning for Robotic Navigation by Reward Shaping'

for sufficient exploration and penalties for premature termination. The action space is often defined over occupancy grid maps, where safe and dangerous actions are distinguished based on the robot's knowledge of free, unknown, and occupied spaces. This structure ensures that the learning process prioritizes both effective coverage and collision avoidance.

[73]: Li et al. (2020), 'Deep Reinforcement Learning-Based Automatic Exploration for Navigation in Unknown Environment'

The authors of [73] indicate that such reward-driven frameworks are essential for enabling robots to autonomously build maps of unknown environments while maintaining operational safety. Graph-based attention networks and other neural architectures have also been explored for their ability to learn spatial relationships and optimize exploration trajectories. These methods can outperform traditional nearest-neighbor or sampling-based approaches in terms of coverage speed, particularly in large-scale environments. However, challenges remain in handling fragmented frontiers and suboptimal viewpoint selection, which can lead to redundant exploration and increased trajectory lengths. According to [83], while learned greedy methods excel in rapid coverage, they may struggle with efficiency when the environment contains many small, scattered frontiers.

The integration of learning-based methods with classical robotics techniques is an active area of research. For example, combining DRL with information-theoretic planning or frontier-based exploration can yield hybrid systems that leverage the strengths of both paradigms. Such systems can adaptively balance exploration and exploitation, dynamically adjusting their strategies based on environmental feedback and learned experience [99]. Additionally, predictive models can be used to anticipate the outcomes of actions, further improving the efficiency and robustness of multi-robot exploration [57].

[100]: Jang et al. (2020), 'Multi-Robot Active Sensing and Environmental Model Learning With Distributed Gaussian Process'

Communication and coordination among multiple robots introduce additional complexity. Learning-based approaches can be extended to multi-agent settings, where policies are trained to account for the actions and observations of other agents. This enables distributed exploration, where robots share information and coordinate their movements to maximize collective coverage while minimizing redundancy. In distributed Gaussian process (GP) frameworks, for example, robots actively sense and share environmental data to collaboratively build accurate maps, all while considering communication constraints and collision avoidance [100]. The ongoing development of learning-based navigation and exploration methods is closely tied to advances in neural network architectures, reward engineering, and multi-agent coordination. As these techniques mature, they promise to deliver multi-robot systems that are not only more autonomous and adaptable but also capable of operating effectively in complex, dynamic, and uncertain environments [57, 73, 83, 99].

[101]: Xu et al. (2021), 'Improving exploration efficiency of deep reinforcement learning through samples produced by generative model'

Our research diverges from the common application of machine learning of exploration policy generation [101]. Instead of learning how to move, our DRBECM-ML framework uses a machine learning model for a distinct, targeted purpose: connectivity prediction. We integrate a pre-trained regression model that, given the relative coordinates between two robots, predicts the expected Received Signal Strength Indicator (RSSI). This predictive capability is then used to inform the robot's existing decision-making logic. For example, an explorer robot will evaluate candidate frontiers not only based on their location but also on the predicted signal quality at that location. This allows the system to make more robust, data-driven decisions that account for realistic signal propagation, moving beyond purely geometric or heuristic-based connectivity criteria.

2.4.2 Predictive Models for Connectivity and Communication

Predictive models for connectivity and communication in multi-robot systems have become increasingly sophisticated, leveraging advances in machine learning to anticipate and adapt to dynamic network conditions. These models are essential for ensuring that distributed robotic teams maintain robust communication links, especially in environments where network topology is subject to frequent changes due to robot mobility, obstacles, or varying task requirements. A key challenge in multi-robot coordination is the maintenance of network connectivity while enabling efficient exploration and task execution. Traditional approaches often rely on explicit control laws or heuristics to preserve communication links, but these can be limited in their adaptability to unforeseen changes in the environment or team composition. Recent research has introduced predictive models that utilize data-driven techniques, such as Artificial Neural Network (ANN) and reinforcement learning, to forecast connectivity states and optimize communication strategies accordingly [49, 102].

Machine learning models, particularly ANNs, have demonstrated the ability to predict system-level properties such as communication quality and network robustness based on sensory and positional data from the robots. For instance, by training on historical data of robot positions and communication outcomes, an ANN can learn to estimate the likelihood of link failures or the emergence of network partitions under various configurations. This predictive capability enables proactive adjustments to robot trajectories or communication parameters, reducing the risk of disconnection and improving overall system resilience. The authors of [102] indicate that such models can achieve high prediction accuracy with low computational overhead, making them suitable for real-time deployment in robotic swarms.

[102]: Tsai et al. (2023), ‘Applying a Neural Network to Predict Surface Roughness and Machining Accuracy in the Milling of SUS304’

Reinforcement Learning has also been applied to the problem of connectivity-aware communication. In these approaches, robots learn policies that balance the trade-off between exploration and the preservation of communication links. For example, RL agents can be trained to select actions that maximize coverage or information gain while minimizing the probability of network fragmentation. This is achieved by incorporating connectivity metrics into the reward function, guiding the learning process toward behaviors that maintain robust communication [49, 57]. Bezcioglu et al. [49] state that combining bio-inspired collective motion with RL allows for the automatic tuning of parameters that influence both flocking behavior and network connectivity, reducing the need for manual intervention and improving adaptability to new scenarios. Connectivity maintenance is further enhanced by predictive models that estimate the impact of robot movements on the underlying communication graph. By modeling the network as a dynamic graph and using machine learning to predict changes in connectivity, robots can make informed decisions about when to move, when to relay messages, or when to adjust their communication range. This approach is particularly effective in scenarios where the network topology is highly dynamic, such as during large-scale exploration or when robots must frequently reconfigure to avoid obstacles [82]. According to [84], strategies based on graph neural networks can dynamically adjust swarm morphology to achieve a balanced distribution of graph importance, thereby improving robustness during topology mutations.

Hybrid strategies that integrate predictive models with traditional control laws have also been explored. For example, a system might use a predictive model to forecast potential connectivity losses and then invoke a control law to steer robots back into communication range if a risk is detected. This layered approach leverages the strengths of both data-driven prediction and formal control theory, resulting in more robust and flexible multi-robot systems [82, 84]. In addition to maintaining connectivity, predictive models can optimize communication efficiency by anticipating congestion or interference in the network. By analyzing patterns in message traffic and environmental factors, machine learning algorithms can suggest adjustments to communication schedules, routing protocols, or transmission power, thereby reducing latency and packet loss. These optimizations are particularly valuable in dense or cluttered environments where communication resources are limited [87]. For instance, decentralized frameworks that rely on local information exchange can benefit from predictive models that estimate the availability and reliability of communication links, enabling agents to coordinate more effectively without centralized oversight.

The integration of predictive models into multi-robot communication frameworks not only enhances robustness but also supports scalability. As the number of robots increases, the complexity of maintaining global connectivity grows, making it impractical to rely solely on centralized control or exhaustive communication. Predictive models enable robots to make local decisions that collectively preserve network integrity, facilitating the deployment of large-scale swarms in complex, dynamic environments [79]. Hyondong Oh et al. [103] outline that behavior-based methods, which combine multiple controllers for tasks such as collision avoidance and formation keeping, can be augmented with predictive models to anticipate and mitigate connectivity disruptions. Overall, the use of predictive models for connectivity and communication represents a significant advancement in the field of multi-robot systems. By leveraging machine learning techniques, these models provide the foresight and adaptability required to maintain robust, efficient, and scalable communication networks in the face of environmental uncertainty and operational complexity [49, 79, 84, 102].

[103]: Oh et al. (2017), ‘Bio-inspired self-organising multi-robot pattern formation: A review’

Our work directly contributes to this area by developing and integrating a predictive model for communication quality. We leverage the large-scale FIT IoT-LAB testbed to collect extensive real-world signal propagation data, which is then used to train and evaluate a suite of regression models. The goal is to predict RSSI values based on positions.

2.4.3 Challenges of Distributed Learning

Distributed learning in multi-robot systems introduces a set of unique challenges that stem from the inherent characteristics of decentralized architectures, the dynamic nature of real-world environments, and the limitations of onboard computation and communication. One of the primary difficulties is the need for each robot to learn and adapt its behavior based on partial, local observations, often without access to global state information. This partial observability complicates the learning process, as agents must infer the state of the environment and the intentions of other robots from limited sensory data and intermittent communication. In highly stochastic and partially observable Markov decision processes (POMDPs), such as those encountered in autonomous exploration, model-free deep reinforcement learning (DRL) methods

are often employed due to their capacity to extract implicit patterns from past experiences. However, these methods require substantial data and computational resources, which may not be readily available on resource-constrained robotic platforms [66, 104].

Another significant challenge is the coordination of learning across multiple agents. In distributed settings, robots must not only learn effective individual policies but also ensure that their actions contribute to the collective objective, such as maximizing coverage or minimizing exploration time. Achieving this coordination is complicated by the non-stationarity introduced when each agent's policy evolves independently, potentially destabilizing the learning process. The dynamic partitioning of the environment, as seen in Voronoi-based approaches, can help distribute tasks among robots, but the underlying learning algorithms must be robust to frequent changes in task allocation and environmental disturbances [52]. Furthermore, the need to update shared representations, such as information graphs or exploration trees, in real time places additional demands on communication bandwidth and synchronization [52, 75].

Communication constraints are a persistent obstacle in distributed learning. Limited bandwidth, delays, and the possibility of information loss can hinder the timely exchange of learned models, policy updates, or sensory data among robots [54, 94]. In cluttered or large-scale environments, multi-hop communication may be required, introducing further latency and increasing the risk of outdated or inconsistent information propagating through the network [54, 74]. These issues can degrade the performance of distributed learning algorithms, especially those that rely on frequent parameter sharing or consensus mechanisms. Stability-maintenance control techniques, such as energy-bounding algorithms, have been proposed to mitigate the effects of communication delays and information loss, but integrating these with learning-based approaches remains an open research problem [94].

The heterogeneity of robot capabilities and the diversity of tasks further complicate distributed learning. Robots may differ in sensing, actuation, or computational power, necessitating adaptive learning strategies that can accommodate such variations. For instance, attention-based neural networks have been explored to model multi-scale dependencies and prioritize relevant information, but their effectiveness depends on the quality and consistency of the data exchanged among agents [66]. Additionally, the integration of learned motion policies with global navigation strategies requires careful design to ensure that local learning does not conflict with global objectives, particularly in goal-driven exploration scenarios [104].

Robustness and adaptability are critical in dynamic, unpredictable environments. Distributed learning algorithms must be resilient to robot failures, sensor noise, and unforeseen obstacles. Hybrid architectures that combine reactive and deliberative components have shown promise in enabling robots to escape local minima and adapt to new situations, but balancing the trade-off between exploration efficiency and safety remains challenging. The uncertainty in robot pose and the need to analyze exploration decision trees in real time add further complexity to the learning process [75]. Finally, the evaluation and validation of distributed learning strategies in realistic settings present methodological challenges. Simulations can provide valuable insights, but transferring learned policies to physical robots often exposes discrepancies due to unmodeled dynamics, sensor inaccuracies, or environmental variability [94, 105]. Ensuring that distributed learning algorithms generalize well

[66]: Cao et al. (2024), 'Deep Reinforcement Learning-Based Large-Scale Robot Exploration'

[104]: Cimurs et al. (2022), 'Goal-Driven Autonomous Exploration Through Deep Reinforcement Learning'

[105]: Shetty et al. (2023), 'Decentralized Connectivity Maintenance for Multi-robot Systems Under Motion and Sensing Uncertainties'

across different platforms and scenarios is essential for their practical deployment in multi-robot systems.

We consciously sidestep many of the challenges associated with online, distributed learning by adopting a different approach. Instead of having robots learn policies during deployment, we use a machine learning model that is trained offline on a comprehensive, real-world dataset. This pre-trained model is then deployed as a static, lightweight prediction module on each robot. This strategy offers several advantages. It avoids the immense computational and communication overhead required for online training and parameter sharing [106, 107].

[106]: Chen et al. (2021), ‘Communication-efficient federated learning’

[107]: Zong et al. (2024), ‘Fedcs: Efficient communication scheduling in decentralized federated learning’

2.5 Summary and Research Gaps

This comprehensive literature review reveals the rich theoretical landscape of multi-robot exploration and highlights the complex interplay between exploration efficiency, communication reliability, and system robustness. Through the examination of centralized, decentralized, and hybrid approaches, along with connectivity-aware strategies, bio-inspired coordination mechanisms, and machine learning techniques, several key insights and research gaps emerge that directly motivate the contributions of this thesis.

2.5.1 Key Insights from the Literature

The survey of multi-robot exploration strategies demonstrates a fundamental tension between coordination efficiency and system resilience. Centralized approaches achieve optimal global coordination through complete environmental knowledge and sophisticated optimization techniques [49, 52], but their reliance on central coordinators introduces critical vulnerabilities and scalability limitations that render them unsuitable for real-world deployment in unpredictable environments [55, 58]. Conversely, decentralized approaches prioritize robustness and scalability through local decision-making [59, 60], but often struggle to achieve effective coordination, leading to redundant exploration and suboptimal coverage [10, 59].

Hybrid approaches have emerged as a promising compromise, attempting to capture the benefits of both paradigms [56, 74, 75]. However, existing hybrid methods primarily focus on combining deterministic and metaheuristic algorithms for navigation optimization rather than addressing the fundamental challenge of dynamic role specialization in multi-robot teams. The integration of reinforcement learning into hybrid frameworks shows promise [49, 64, 66] but introduces significant computational overhead and training complexity that limits practical deployment.

The connectivity-aware exploration literature reveals the critical importance of maintaining network integrity during exploration missions [74, 79]. Existing approaches typically treat connectivity as a constraint to be satisfied rather than as a dynamic property to be actively managed. This static perspective limits the ability of multi-robot systems to adapt to changing environmental conditions and communication challenges. Furthermore, most connectivity-aware methods rely on geometric connectivity models that fail to capture the complex propagation characteristics of real-world wireless communication [81].

Bio-inspired coordination mechanisms demonstrate the power of simple, local rules in generating sophisticated collective behaviors [62, 84, 94]. However, the application of these principles to connectivity-aware exploration remains largely unexplored. Existing bio-inspired approaches focus primarily on general flocking or swarm behaviors without addressing the specific requirements of maintaining network connectivity during exploration tasks. This represents a significant opportunity for innovation in applying natural coordination strategies to engineered systems.

The machine learning literature in robotics shows tremendous potential for enhancing navigation and exploration capabilities [57, 73, 99]. However, most learning-based approaches focus on policy generation rather than leveraging machine learning for specific prediction tasks that support existing algorithms. The challenges of distributed learning [66, 104] have led many researchers to avoid real-time learning in favor of pre-trained models, but the application of such models to connectivity prediction remains underexplored.

2.5.2 Identified Research Gaps

Through this comprehensive analysis, several critical gaps emerge in the current state of multi-robot exploration research that limit the practical deployment of these systems in real-world scenarios:

Dynamic Role Specialization: Existing multi-robot exploration frameworks assign static roles to robots or rely on centralized task allocation mechanisms. There is a lack of approaches that enable robots to dynamically adapt their roles based on local network conditions and exploration requirements. This limitation prevents systems from achieving the flexibility needed to respond to changing environmental conditions and communication challenges.

Integration of Exploration and Connectivity Maintenance: Current approaches typically treat exploration and connectivity maintenance as separate, competing objectives. The literature lacks frameworks that view these objectives as complementary aspects of a unified system design. This gap prevents the development of truly integrated solutions that can simultaneously optimize both exploration efficiency and network robustness.

Bio-Inspired Connectivity Maintenance: While bio-inspired coordination has been extensively studied for general swarm behaviors, its application to the specific problem of dynamic connectivity maintenance during exploration remains largely unexplored. The potential for flocking-inspired algorithms to adaptively position relay robots and maintain network bridges represents an untapped research opportunity.

Practical Machine Learning Integration: The machine learning literature in robotics focuses primarily on complex policy learning approaches that require significant computational resources and training infrastructure. There is a gap in research that explores how simple, pre-trained predictive models can be integrated into existing algorithms to enhance specific capabilities without introducing the complexity of online learning.

Robustness in Unpredictable Environments: Many existing approaches are designed and evaluated in controlled environments with predictable conditions. The literature lacks comprehensive treatment of system

robustness in truly unpredictable scenarios where infrastructure is unreliable, communication is intermittent, and environmental conditions change rapidly.

2.5.3 Motivation for This Research

The identified gaps collectively point toward the need for a fundamentally different approach to multi-robot exploration that addresses the core challenge of maintaining network connectivity while exploring unknown environments. The motivation for this research stems from the recognition that existing solutions fail to meet the requirements of real-world deployment scenarios, particularly in emergency response and disaster management contexts where reliability and adaptability are paramount.

The DRBECM and DRBECM-ML frameworks developed in this thesis address these gaps through a novel combination of dynamic role assignment, bio-inspired connectivity maintenance, and practical machine learning integration. By enabling robots to dynamically switch between exploration-focused and connectivity-maintenance roles based on local conditions, our approach transcends the traditional static role assignments that limit existing systems. The integration of flocking-inspired algorithms specifically for connectivity maintenance provides a natural and robust mechanism for adaptive relay positioning. Finally, the incorporation of realistic communication prediction models, trained on real-world data, enables more accurate connectivity assessment and decision-making.

This research contributes to the advancement of multi-robot exploration by demonstrating how carefully designed combinations of established techniques can address fundamental limitations in existing approaches. The emphasis on practical deployment considerations and real-world validation ensures that the proposed solutions are not merely theoretical contributions but represent genuine advances toward the deployment of robust multi-robot exploration systems in challenging environments.

Through this comprehensive foundation, the thesis establishes both the theoretical context and practical motivation for the novel algorithmic frameworks presented in subsequent chapters. The integration of insights from distributed systems, bio-inspired coordination, and machine learning provides the necessary background for understanding how the proposed approaches advance the state of the art in connectivity-aware multi-robot exploration.

Comparative Analysis of Dynamic Communication Topologies for MRS

3

3.1 Introduction

Effective coordination in Multi-Robot Systems is fundamentally dependent on the underlying communication strategy, which significantly impacts the efficiency of area coverage, mapping accuracy, and collaborative decision-making [5]. In dynamic and unknown environments, such as those encountered in search and rescue or post-disaster assessment, robot mobility causes constant and unpredictable changes in network topology, requiring adaptive and localized communication protocols that do not rely on static infrastructure.

A central challenge in designing these protocols is managing the fundamental trade-off between comprehensive information sharing and the associated communication overhead in terms of bandwidth, energy, and computation [108]. Furthermore, real-world deployments must contend with communication-restricted environments, where obstacles, signal attenuation, and non-line-of-sight (NLOS) paths lead to intermittent connectivity and potential data loss. These factors demand robust strategies that can maintain network integrity amidst uncertainty [7, 109]. Graph-based approaches have emerged as a primary method for managing these dynamic inter-robot communications, offering a structured way to build and maintain a communication graph that is sparse enough to minimize overhead yet robust enough to ensure reliable data exchange. [110].

This chapter presents a rigorous, quantitative comparison of three prominent graph-based strategies for dynamic topology construction: K-Nearest Neighbors (KNN), Relative Neighborhood Graph (RNG), and the k-Relative Neighborhood Graph (K-RNG). The objective is to determine which approach provides the most suitable foundation for decentralized multi-robot exploration and which one offers the best delay/PDR compromise. To this end, the strategies are evaluated through extensive simulations against critical Quality-of-Service (QoS) metrics, namely Packet Delivery Ratio for reliability and end-to-end latency for timeliness, to assess their reliability and timeliness.

The selection of these three methods is motivated by their complementary approaches to the fundamental challenge of balancing network connectivity with communication overhead. KNN represents the distance-based paradigm, offering computational simplicity through pure proximity-based neighbor selection. RNG embodies the geometric approach, providing parameter-free operation with formal connectivity guarantees through lune-based constraints. K-RNG bridges these approaches by combining geometric rigor with tunable density control. Together, these methods span the complete spectrum from sparse to dense topologies, covering the core design space for graph-based communication in autonomous multi-robot systems. Their widespread adoption in the literature and established theoretical foundations make them the natural benchmark for evaluating communication strategies in distributed robotics applications.

3.1	Introduction	39
3.2	Background on Graph-Based Topology Strategies	40
3.2.1	Relative Neighborhood Graph (RNG)	40
3.2.2	k-Relative Neighborhood Graph (K-RNG)	40
3.2.3	K-Nearest Neighbors (KNN)	40
3.2.4	Detailed Theoretical Comparison	41
3.3	Simulation and Evaluation Methodology	43
3.3.1	Scenario 1: Peer-to-Peer Communication	43
3.3.2	Scenario 2: Multi-hop Routing	44
3.3.3	Performance Metrics	44
3.4	Results of Comparative Analysis	44
3.4.1	Peer-to-Peer Performance (Scenario 1)	44
3.4.2	Multi-hop Routing Performance (Scenario 2)	45
3.5	Chapter Summary and Implications	47
3.6	Limitations and Future Directions	48
3.6.1	Current Limitations	48
3.6.2	Future Research Directions	49

3.2 Background on Graph-Based Topology Strategies

This section introduces the three graph-based topology strategies evaluated in this comparative analysis: K-Nearest Neighbors, Relative Neighborhood Graph, and k-Relative Neighborhood Graph. Each strategy manages network links based on distinct geometric or proximity-based criteria, offering different approaches to balancing network connectivity and overhead. These strategies determine the set of neighbors to which a robot sends its data, directly shaping the flow of information through the system.

3.2.1 Relative Neighborhood Graph (RNG)

Relative Neighborhood Graph, first proposed by Toussaint [111], is a parameter-free graph, a characteristic that makes it highly attractive for autonomous systems as it eliminates the need for manual tuning. Within the MRS context, an edge is established between two nodes, p and q , if and only if no third robot is closer to both p and q than they are to each other. Formally, this condition is met when the "lune," defined as the intersection of two open disks centered at p and q with a radius equal to the Euclidean distance $d(p, q)$, contains no other robot, as shown in Figure 3.1. This creates a sparse network that efficiently manages network connectivity while reducing redundant data exchanges. As a subgraph of the Delaunay triangulation and a supergraph of the Euclidean Minimum Spanning Tree, RNG is guaranteed to preserve network connectivity if the initial graph is connected [111, 112].

[111]: Toussaint (1980), 'The relative neighbourhood graph of a finite planar set'
 [112]: Jaromczyk et al. (1992), 'Relative neighborhood graphs and their relatives'




Figure 3.1: Lune region for RNG: the shaded intersection of disks of radius $d(p, q)$.

[113]: Foster et al. (2022), *Generalized Relative Neighborhood Graph (GRNG) for Similarity Search*




Figure 3.2: K-RNG lune with $k = 1$: the single interior point c does not block the edge.

[114]: Zardini et al. (2024), 'A quantum k-nearest neighbors algorithm based on the Euclidean distance estimation'

3.2.2 k-Relative Neighborhood Graph (K-RNG)

The K-RNG is a generalization of the RNG that affords tunable control over network density [113]. An edge is formed between nodes p and q if their lune contains at most k other nodes from the set, as shown in the example Figure 3.2. The parameter k , a non-negative integer, allows for adjustable neighborhood definitions, when $k=0$, the K-RNG is identical to the standard RNG. As k is increased, the resultant graph becomes progressively denser, incorporating more edges to potentially enhance redundancy and capture broader neighborhood relationships. As k is increased, the resultant graph becomes progressively denser, potentially enhancing redundancy. However, this flexibility comes at the cost of increased computational complexity compared to RNG and requires careful tuning of the k parameter for optimal performance.

3.2.3 K-Nearest Neighbors (KNN)

The K-Nearest Neighbors (KNN) strategy dictates that each robot establishes communication links with its k closest neighbors, determined by Euclidean distance [114]. In contrast to RNG and K-RNG, which employ a geometric lune-based criterion that accounts for the spatial configuration between nodes, KNN is predicated exclusively on a simple distance metric. Although this facilitates rapid adaptation to local density fluctuations, it offers no formal guarantee of graph connectivity. This is a significant limitation, as it may result in network partitioning, especially

in sparse or irregularly distributed robot formations, severely impacting communication reliability[115].




Figure 3.3: Graph comparison: RNG, k-RNG (top) and k-NN (bottom) on the same 15 points example.

Figure 3.3 illustrates an example of the key structural differences between these graph families on the same point set. The RNG graph (top row) shows how relaxing the empty region constraint from circles (RNG) to larger regions (k-RNG) systematically increases connectivity while preserving important geometric properties. Notably, the 2-RNG maintains planarity despite the increase in edge count. However, as the number of edges increases, k-RNG loses its planarity as seen in the $k=3$ case.

In contrast, the k-NN graphs (bottom row) exhibit fundamentally different connectivity patterns. For small values of k , k-NN graphs may produce disconnected components, as seen in the $k=2$ and $k=3$ cases where several isolated clusters exist. This connectivity issue is a key limitation of k-NN graphs for applications requiring connected proximity structures. However, as k increases, k-NN graphs typically achieve connectivity at the cost of losing planarity unlike the RNG.

3.2.4 Detailed Theoretical Comparison

Having detailed the individual geometric and proximity-based criteria for each strategy, it is useful to consolidate their theoretical properties for a direct comparison. The following table 3.1 summarizes these key attributes, contrasting the three approaches based on characteristics such as graph planarity, connectivity guarantees, and algorithmic complexity. These parameters are critical for multi-robot exploration applications: *connectivity* ensures reliable information flow for coordinated exploration, *planarity* affects routing protocol performance and computational efficiency (non-planarity increases routing algorithm complexity because

[116]: Ansari et al. (2005), 'A Generalization of the Face Routing'
 [117]: Clouser et al. (2013), 'Void traversal for efficient non-planar geometric routing'

it requires additional memory, computation, and algorithmic adaptations to handle edge crossings and guarantee reliable message delivery, unlike planar networks where simpler algorithms suffice [116, 117]), *energy efficiency* determines operational lifetime in resource-constrained environments, *robustness* impacts system resilience to robot failures, and *algorithmic complexity* affects real-time implementation feasibility on embedded platforms.

Table 3.1: Theoretical Comparison of KNN, RNG, and K-RNG

Attribute	KNN	RNG	K-RNG
Node Degree	Fixed at k (if reachable)	Variable, generally between 2-3	Variable
Planarity	Not planar	Guaranteed	Often planar for small k
Global Connectivity	Not guaranteed	Guaranteed	Guaranteed
Energy Efficiency	Moderate to high (depends on k)	High	Moderate to high (depends on k)
Robustness	Moderate	Low	Tunable via k
Algorithmic Complexity	Low	Moderate	Moderate

Connectivity and Robustness

[118]: Balister et al. (1995), 'Connectivity of random k -nearest-neighbour graphs'

KNN: Ensures connectivity with a high probability, provided that the value of k is sufficiently large [118]. However, if k is too small, the network's robustness to node failures is reduced. A key drawback is that KNN graphs are generally not planar, which can complicate certain routing protocols. This non-planarity occurs because KNN connections are based purely on distance ranking without geometric constraints, allowing communication links to cross arbitrarily in the plane. The resulting edge crossings significantly complicate routing protocols, which must handle overlapping communication paths and maintain more complex routing tables. Non-planar topologies require routing algorithms with higher memory consumption for storing multiple path alternatives and increased computational overhead for path selection decisions, contrasting with planar networks where efficient face-routing or geographic forwarding algorithms can guarantee delivery with minimal resource requirements.

RNG: Always a planar graph [111]. It maintains an average node degree of 2-3 neighbors, resulting in highly efficient memory utilization as each robot stores minimal neighbor relationships and routing information. RNG is a subgraph of the Delaunay triangulation and is guaranteed to be connected if the underlying point set is connected, as it always contains a Minimum Spanning Tree (MST) [111]. Due to its sparse nature, it has few redundant paths, making it less robust to node or link failures.

K-RNG: A connected graph as it is a supergraph of RNG, it provides a tunable trade-off between network density and connectivity. Increasing the parameter k adds more edges, which improves robustness and connectivity at the cost of creating a denser graph [119]. This increased

[119]: Maw-Shang et al. (1991), '20-relative neighborhood graphs are hamiltonian'

density progressively compromises the planarity property as k values grow, with edge crossings becoming evident for higher k values as demonstrated in Figure 3.3. The loss of planarity incurs specific computational penalties: memory requirements scale approximately $O(k)$ neighbors per node compared to RNG's constant 2-3 degree, communication overhead increases due to packet collisions and channel contention from denser connectivity patterns, and routing complexity escalates as algorithms must evaluate multiple overlapping paths rather than employing the simpler geometric routing strategies available in planar topologies.

$\text{RNG (1-RNG)} \subseteq \text{2-RNG} \subseteq \dots \subseteq \text{K-RNG}$

Energy Efficiency and Interference

KNN: May result in higher energy consumption and interference due to redundant links, especially in dense networks.

RNG: Minimizes energy consumption and interference by reducing the number of edges, but at the cost of potential disconnections.

K-RNG: Balances energy efficiency and connectivity; as k increases, energy consumption rises but so does network robustness.

Algorithmic Complexity and Locality

KNN: Simple, fully distributed, and localized; each node only needs local neighbor information.

RNG: Can be constructed locally with one-hop neighbor information, but requires geometric computations.

K-RNG: More complex than RNG due to the additional parameter k , but still feasible for distributed implementation.

3.3 Simulation and Evaluation Methodology

The comparative analysis was executed within the NS-3 discrete-event network simulator [120]. All robot nodes were equipped with an IEEE 802.11n wireless stack operating in ad-hoc mode, utilizing a YansWiFiChannel with a LogDistance propagation loss model (path loss exponent 3.0). Each robot ran a custom application responsible for neighbor discovery and topology management. Neighbor discovery was facilitated by periodic "Hello" messages broadcast every second, containing the sender's current position. Based on the selected strategy (RNG, K-RNG, or KNN), each robot would then periodically send 512-byte data packets via User Datagram Protocol (UDP) to its calculated neighbor.

[120]: Riley et al. (2010), 'The ns-3 Network Simulator'

3.3.1 Scenario 1: Peer-to-Peer Communication

The primary objective of this scenario was to assess the efficiency of direct, single-hop communication among neighboring robots. The experimental parameters were as follows:

- **Network Scale:** 5 to 60 robot nodes.
- **Area:** 100×100 meters.
- **Mobility:** Random Walk2D mobility model (1-3 m/s).

[121]: Oluwafemi et al. (2024), 'Random walk theory and application'

- **Simulation Duration:** 10 seconds per run.

The Random Walk2D mobility model [121] simulates unpredictable movement patterns where each robot independently selects a random direction (0-360°) and speed within the specified range, moves for a random time duration, then pauses briefly before selecting new movement parameters. This model effectively captures the erratic mobility patterns typical in exploration scenarios.

3.3.2 Scenario 2: Multi-hop Routing

This scenario was formulated to evaluate the scalability and routing efficacy of each strategy in a data collection application requiring data forwarding to a central sink. The experimental parameters included:

- **Network Scale:** 10 to 60 robot nodes.
- **Area:** Ranging from 200×200 to 400×400 meters.
- **Base Station:** A stationary sink node at the center of the operational area.
- **Mobility and Duration:** RandomWalk2D mobility over 300 seconds.
- **Routing Protocol:** Packet routing was governed by a greedy geographic forwarding protocol, wherein packets are relayed to the neighbor that minimizes the Euclidean distance to the base station.

3.3.3 Performance Metrics

To facilitate a direct comparison, the following QoS metrics were employed across both scenarios:

- **Packet Delivery Ratio (PDR):** The fraction of data packets successfully received at their final destination relative to the total number transmitted.
- **End-to-End Delay:** The total time elapsed from packet generation at the source to its successful reception at the final destination, encompassing multiple hops where applicable.

3.4 Results of Comparative Analysis

The simulations revealed distinct performance characteristics for each strategy, with outcomes varying significantly between the peer-to-peer and multi-hop routing scenarios.

3.4.1 Peer-to-Peer Performance (Scenario 1)

Packet Delivery Ratio

In direct communication, RNG demonstrates superior PDR performance across all network densities, as illustrated in Figure 3.4. It begins with a PDR of approximately 95% for 5 robots, declining to 33% with 60 robots, consistently outperforming the other strategies. This highlights its effectiveness in creating sparse, efficient topologies that minimize packet collisions and interference. The K-RNG strategies consistently

underperform, with initial PDR values between 70-80% declining to 22-28%. This can be attributed to the increased network density created by keeping neighbors within the lune, which leads to a higher rate of packet collisions and channel contention. While KNN strategies are competitive (85-90% PDR) for a low number of robots, and they ultimately converge to similar PDR values as network density increases.




Figure 3.4: Packet Delivery Ratio (PDR) as a function of the number of robots for all evaluated strategies in the peer-to-peer communication scenario.

Delay

Analysis of peer-to-peer delay, as shown in Figure 3.5, reveals that while RNG experiences the steepest relative increase in delay as the network grows, it consistently maintains the **lowest delay values** throughout the evaluation. In contrast, K-RNG strategies exhibits higher baseline delays that remained relatively consistent, suggesting that their increased connectivity does not compensate for the additional network overhead.

3.4.2 Multi-hop Routing Performance (Scenario 2)

Packet Delivery Ratio

In the multi-hop data collection scenario, the performance trends shifted, as demonstrated in Figure 3.6. The K-RNG strategies achieve the highest median PDR, with K-RNG (k=3) showing the best performance at approximately 32-33%. This suggests that the geometric constraints of K-RNG, with a moderate tolerance for interior points, create more robust multi-hop paths for the geographic forwarding protocol. RNG remains highly competitive with a median PDR of 31-32%, demonstrating stable routing behavior. In contrast, KNN strategies consistently show lower median performance, highlighting how the lack of connectivity guarantees becomes more pronounced in multi-hop scenarios.




Figure 3.5: Average peer-to-peer delay as a function of the number of robots for all evaluated strategies.




Figure 3.6: Box plot distribution of PDR per strategy in the multi-hop data collection scenario.




Figure 3.7: RNG multi-hop PDR heatmap across network scales and deployment areas. Shows consistent performance with graceful degradation, reaching 55.1% peak PDR in multi-hop routing scenarios.

Detailed Multi-hop Performance Analysis via Heatmaps: To provide a more comprehensive understanding of the multi-hop PDR performance characteristics across different operational conditions, we present detailed heatmap analyses that visualize performance variations across different network scales (10-60 robots) and deployment areas (100×100 to 400×400 units).

The RNG heatmap (Figure 3.7) reveals consistent multi-hop routing performance with: (1) Strong performance in smaller deployment areas (100×100m) with PDR values reaching 55.1% through effective sparse topology construction that reduces routing overhead; (2) Predictable degradation patterns as network density increases, maintaining reasonable multi-hop performance (26.0-35.9%) across larger deployment areas; (3) Stable routing behavior that scales well with increasing network parameters.

In contrast, KNN ($K=1$) exhibits highly variable multi-hop behavior (Figure 3.8) with performance fluctuating dramatically based on network conditions. The purely distance-based neighbor selection creates unpredictable routing topologies that suffer from poor multi-hop path reliability, with performance dropping as low as 14.0% in challenging multi-hop scenarios.

The comparison between different KNN configurations for multi-hop routing (Figures 3.9 and 3.10) illustrates how increased connectivity can improve routing path diversity and reliability. KNN ($K=2$) shows enhanced multi-hop stability compared to $K=1$, providing more routing options through additional neighbor connections. KNN ($K=3$) achieves the best performance among KNN variants in multi-hop scenarios (reaching 65.4% PDR), but still exhibits the fundamental weaknesses of distance-based selection when routing paths become complex.

The K-RNG family demonstrates controlled evolution of multi-hop routing characteristics (Figures 3.11, 3.12, and 3.13). K-RNG ($K=1$), equivalent to standard RNG, provides the baseline geometric multi-hop performance. K-RNG ($K=2$) shows modest improvements in multi-hop scenarios, particularly beneficial for medium-density networks where controlled redundancy enhances routing path reliability. K-RNG ($K=3$) achieves the highest peak multi-hop performance (71.1% PDR) among all evaluated strategies, validating the benefits of controlled geometric constraint relaxation for multi-hop routing applications.

Multi-hop Performance Insights: The heatmap analysis reveals several critical insights for multi-hop routing in multi-robot systems:

- **Routing Topology Impact:** Geometric constraints (RNG, K-RNG) create more stable multi-hop paths compared to purely distance-based approaches (KNN), leading to more predictable routing behavior across different network conditions.
- **Density-Dependent Routing:** Multi-hop performance shows strong dependency on network density, with geometric strategies maintaining better routing reliability as density increases, while distance-based strategies suffer from increased routing instability.
- **Controlled Redundancy Benefits:** K-RNG strategies demonstrate that controlled geometric redundancy significantly improves multi-hop routing performance, with K-RNG ($K=3$) providing optimal balance between connectivity and routing efficiency.
- **Scalability Characteristics:** RNG and K-RNG strategies exhibit superior scalability in multi-hop scenarios, maintaining consistent routing behavior across different deployment scales, while KNN strategies show more variable performance patterns.

This detailed multi-hop performance characterization validates the selection of geometric topology strategies for multi-robot systems requiring reliable multi-hop communication, particularly in scenarios where routing reliability is critical for mission success.

3.5 Chapter Summary and Implications

This chapter presented a direct comparative analysis of the RNG, K-RNG, and KNN dynamic topology strategies for multi-robot communication in terms of delay and PDR. The results from extensive NS-3 simulations




Figure 3.8: KNN ($K=1$) multi-hop PDR heatmap showing erratic performance patterns in multi-hop routing. Performance ranges from 58.7% in optimal conditions to 14.0% in challenging scenarios.




Figure 3.9: KNN ($K=2$) multi-hop PDR heatmap showing improved multi-hop stability over $K=1$ with peak performance of 65.1%. Enhanced connectivity provides better routing path diversity.




Figure 3.10: KNN ($K=3$) multi-hop PDR heatmap achieving highest KNN multi-hop performance (65.4% peak) through increased connectivity redundancy, but still suffers from routing instability in dense scenarios.




Figure 3.11: K-RNG (K=1) multi-hop PDR heatmap, equivalent to standard RNG, providing baseline geometric multi-hop performance with 65.7% peak PDR and consistent routing behavior.




Figure 3.12: K-RNG (K=2) multi-hop PDR heatmap showing modest improvements over K=1 in multi-hop scenarios, particularly in medium-density networks where controlled redundancy enhances routing reliability.




Figure 3.13: K-RNG (K=3) multi-hop PDR heatmap achieving highest overall multi-hop performance (71.1% PDR) among all evaluated strategies, demonstrating superior multi-hop routing through optimal geometric constraint relaxation.

show that no single strategy is optimal for all scenarios, but clear patterns emerge.

For direct peer-to-peer communication, RNG provides the best balance of performance and simplicity, achieving superior PDR and low delay. In multi-hop routing scenarios, K-RNG demonstrates superior performance, though at the cost of increased computational complexity and the need for parameter tuning. KNN consistently underperforms in both scenarios, as its purely distance-based selection and lack of connectivity guarantees make it less robust for reliable communication tasks.

The detailed heatmap analysis of multi-hop performance reveals that geometric topology strategies (RNG, K-RNG) provide significantly more stable and predictable routing behavior compared to distance-based approaches. K-RNG (K=3) achieves the highest multi-hop performance with 71.1% peak PDR, demonstrating the benefits of controlled geometric redundancy for complex routing scenarios.

For practitioners designing distributed multi-robot systems, on which connectivity should be maintained, RNG stands out as a good choice for most general-purpose applications. Its parameter-free nature eliminates the need for complex tuning, and it delivers strong, balanced performance across different network densities and communication patterns. For applications specifically requiring robust multi-hop routing, K-RNG (K=3) provides superior performance at the cost of increased computational complexity. This empirical evidence provides a solid foundation for the selection of RNG as the core communication backbone for the DR-BECM exploration framework, which will be detailed in the subsequent chapter.

3.6 Limitations and Future Directions

3.6.1 Current Limitations

The presented comparative analysis operates under several methodological constraints that define the boundaries of its applicability. The NS-3 simulation environment, while providing controlled experimental conditions, cannot capture all complexities of real-world signal propagation including multipath fading, time-varying channel conditions, and environmental interference patterns that significantly affect communication quality in deployed systems.

The evaluation focused primarily on PDR and latency metrics, which provide essential but incomplete characterization of network performance. Additional QoS parameters such as energy consumption per successful transmission, network resilience to node failures, and adaptive capacity under varying traffic loads would provide more comprehensive performance assessment for practical deployment scenarios.

The mobility model employed (Random Walk2D) represents simplified movement patterns that may not accurately reflect the structured exploration behaviors exhibited by coordinated robot teams. More realistic mobility models incorporating frontier-based movement, obstacle avoidance, and formation maintenance behaviors could reveal different performance characteristics for each topology strategy.

3.6.2 Future Research Directions

Several promising research directions emerge from this comparative analysis that could significantly enhance the practical applicability of dynamic communication topologies for multi-robot systems.

Adaptive Topology Strategies

Future research should investigate dynamic parameter adjustment mechanisms for K-RNG that can automatically optimize the interior point tolerance based on real-time network conditions, robot density, and communication quality metrics. Such adaptive systems could leverage reinforcement learning or control-theoretic approaches to continuously tune the k parameter, potentially achieving the benefits of geometric redundancy without requiring manual parameter optimization.

Machine learning approaches could predict optimal topology configurations based on mission phase, environmental conditions, and team composition. By learning from historical performance data, adaptive systems could proactively adjust communication strategies to maintain optimal performance across diverse operational scenarios.

Hybrid Communication Approaches

Integration of multiple topology strategies within the same network represents a significant research opportunity. Individual robots could dynamically select among RNG, K-RNG, and KNN based on their local network conditions and mission requirements. For example, robots in sparse regions might employ K-RNG for enhanced connectivity, while those in dense areas could use RNG for efficiency.

Hierarchical communication architectures combining different topology strategies at different network levels could optimize performance across multiple scales. Local clusters might use dense KNN connections for rapid information sharing, while inter-cluster communication could employ sparse RNG links for efficient long-range connectivity.

Extended Performance Evaluation

Future evaluations should incorporate additional QoS metrics including energy consumption patterns, network lifetime under battery constraints, and graceful degradation characteristics under progressive node failures. Such comprehensive assessment would better inform deployment decisions for resource-constrained robotic applications.

Investigation of topology performance under realistic robot exploration behaviors, including frontier-based navigation, formation maintenance, and obstacle avoidance, would provide more accurate performance characterization for actual deployment scenarios. Integration with realistic robot dynamics and sensor limitations would further enhance evaluation validity.

Real-World Validation

Extensive field experiments across diverse environments are essential to validate simulation findings and identify environmental factors that significantly impact topology performance. Urban environments with complex signal propagation, natural terrain with vegetation effects, and indoor scenarios with multipath interference each present unique challenges that could favor different topology strategies.

Long-term deployment studies could evaluate topology stability and adaptation under changing environmental conditions, robot failures, and evolving mission requirements. Such studies would inform the development of robust topology management strategies for practical applications.

DRBECM: A Distributed Multi-Robot Exploration Approach with Connectivity Maintenance

4

4.1 Introduction

This chapter introduces the **Dynamic Role-Based Exploration with Connectivity Maintenance (DRBECM)** framework. DRBECM is a novel, fully distributed multi-robot exploration algorithm designed to operate effectively in unknown environments where maintaining uninterrupted communication is as critical as the exploration task itself. This is particularly relevant in scenarios such as search and rescue, where reliable data flow and team coordination are paramount.

Traditional exploration algorithms often face a trade-off between maximizing the speed of discovery and preserving network connectivity. Many approaches rely on centralized control or global environmental knowledge, which are impractical in infrastructure-denial scenarios [122]. Conversely, methods that prioritize connectivity often suffer from reduced exploration efficiency, leading to slower mission completion [123].

DRBECM addresses these challenges through a decentralized framework that allows robots to dynamically adapt their behavior. The core contribution is a role-switching mechanism that assigns each robot one of two roles: **explorer** or **supporter**. Explorers actively seek to expand the known map using a frontier-based strategy, while supporters position themselves to act as communication relays, ensuring the entire fleet remains connected to the base station. This distribution of tasks, managed autonomously by each robot based on local information, creates a system that is robust, scalable, and efficient.

This chapter details the theoretical underpinnings and mechanics of the DRBECM algorithm, from its system model to the specific behaviors governing robot actions and interactions. We will describe the rules for dynamic role assignment, the frontier-based exploration logic, the flocking-inspired supporter positioning, and the mechanisms for collision avoidance and stagnation recovery. The chapter concludes by presenting a performance evaluation, comparing DRBECM against established centralized and decentralized methods to empirically validate its effectiveness.

4.2 System Model and Background

The development of effective distributed multi-robot exploration algorithms requires a clear specification of the system architecture, robot capabilities, and environmental assumptions that govern system operation. This section establishes the fundamental system model underlying the Dynamic Role-Based Exploration with Connectivity Maintenance (DRBECM) framework, providing the mathematical foundations and operational constraints that guide the algorithm design.

The system model defines a multi-robot team operating in an unknown two-dimensional environment, where each robot possesses limited sensing and communication capabilities and must make decisions based

4.1	Introduction	51
4.2	System Model and Background	51
4.3	The DRBECM Algorithm	55
4.3.1	Dynamic Role Assignment	55
4.3.2	Explorer Robot Behavior: Frontier-Based Exploration	56
4.3.3	Supporter Robot Behavior: Flocking-Inspired Positioning	56
4.3.4	Stagnation Detection and Collision Avoidance	57
4.3.5	Information Sharing and Map Updating	59
4.3.6	Main Algorithm	59
4.4	Performance Evaluation	61
4.4.1	Experimental Setup	61
4.4.2	Results and Analysis	62
4.5	Limitations and Future Enhancements	65
4.5.1	Discussion and Current Limitations	65
4.5.2	Algorithmic Enhancements	65
4.5.3	Deployment and Validation Requirements	66

solely on local information. By explicitly modeling robot state, communication constraints, and environmental interaction, this framework provides the necessary foundation for understanding how distributed coordination can emerge from local decision-making processes while maintaining global system objectives of exploration efficiency and network connectivity.

We consider a multi-robot system composed of a fleet of N homogeneous mobile robots, $\mathcal{R} = \{r_1, r_2, \dots, r_N\}$, operating in an unknown two-dimensional environment. A stationary base station, B , is positioned at a known location $\mathbf{p}_B$ and serves as the primary communication hub and the origin point for the exploration mission. The system is designed to be fully decentralized, with each robot r_i making decisions based on its local perceptions and information exchanged with its direct neighbors.

Each robot r_i is defined by the following state and capabilities:

- **Position and Velocity:** $p_i(t), v_i(t) \in \mathbb{R}^2$ at a given time t . We assume that each robot has access to its own position through localization systems such as GPS (in outdoor environments), visual-inertial odometry, or Simultaneous Localization and Mapping (SLAM).
- **Role:** $\rho_i(t)$, which can be either ‘explorer’ or ‘supporter’.
- **Sensing Range (R_s):** A radius within which the robot can map its immediate environment.
- **Communication Range (R_c):** A radius within which the robot can exchange data with other robots or the base station.
- **Local Map and Frontiers:** $M_i(t)$ represents the robot’s local map of the environment, and $F_i(t)$ is the set of detected frontiers (boundaries between explored and unexplored space).

We denote the Euclidean distance between two robots r_i and r_j at time t by $d_{ij}(t) = \|\mathbf{p}_i(t) - \mathbf{p}_j(t)\|$.

As established in Chapter 3, network connectivity is managed using the RNG. The selection of RNG over alternative dynamic topology strategies is based on our comparative analysis presented in Chapter 3, which demonstrated RNG’s superior balance between communication reliability and computational efficiency in multi-robot exploration scenarios. Specifically, RNG showed the best Packet Delivery Ratio performance while being parameter-free, making it particularly suitable for dynamic environments where optimal parameter tuning is challenging. Each robot r_i determines its set of neighbors, $N_i(t)$, based on the RNG criterion, ensuring a connected and sparse communication graph without the need for parameter tuning. This local neighbor selection minimizes communication overhead and computational complexity, which is essential for a scalable distributed system.

The construction of the RNG neighbor set is performed locally by each robot. Algorithm 1 describes the process where robot r_i initially considers all robots within communication range as potential neighbors, then removes those that violate the RNG lune criterion.

Algorithm 1: RNG Neighbor Set Construction for Robot r_i **Input** : Robot position $\mathbf{p}_i$, Communication range R_c **Output**: RNG neighbor set $N_i(t)$

```

1  $N_i(t) \leftarrow \{r_j \mid d_{ij}(t) \leq R_c, j \neq i\};$  // Start with all robots in
   range
2 foreach  $r_j \in N_i(t)$  do
3   foreach  $r_k \in N_i(t), k \neq j$  do
4      $d_{ij} \leftarrow \|\mathbf{p}_i - \mathbf{p}_j\|;$ 
5      $d_{ik} \leftarrow \|\mathbf{p}_i - \mathbf{p}_k\|;$ 
6      $d_{jk} \leftarrow \|\mathbf{p}_j - \mathbf{p}_k\|;$ 
7     if  $d_{ik} < d_{ij}$  and  $d_{jk} < d_{ij}$  then
8        $N_i(t) \leftarrow N_i(t) \setminus \{r_j\};$  // Remove  $r_j$ , as  $r_k$  is in the
         lune
9       break;
10    end
11  end
12 end
13 return  $N_i(t);$ 

```

The algorithm evaluates the lune criterion: robot r_k is inside the lune of r_i and r_j if it is closer to both than they are to each other ($d_{ik} < d_{ij}$ and $d_{jk} < d_{ij}$). If such a robot exists, r_j is removed from the neighbor set. The complexity is $O(|N_i(t)|^2)$ where $|N_i(t)|$ remains bounded at 2-3 neighbors on average in two-dimensional RNG graphs.

The DRBECM algorithm relies on multiple interacting mechanisms, each defined by specific parameters that govern system behavior. Table 4.1 provides a comprehensive specification of all algorithm parameters, organized by functional category including system architecture, communication constraints, role assignment criteria, supporter dynamics, collision avoidance, stagnation recovery, and frontier exploration. The specified value ranges represent typical operational configurations used in the experimental evaluation presented in Section 4.4.

Table 4.1: DRBECM Parameter Specification

Parameter	Definition	Value/Range
SYSTEM ARCHITECTURE		
N	Total number of robots in the exploration team	$3 \leq N \leq 15$
$\mathcal{R}$	Set of robots $\{r_1, r_2, \dots, r_N\}$	–
B	Base station with fixed position $\mathbf{p}_B$	–
$\mathbf{p}_i(t)$	Position vector of robot r_i at time t	$\mathbb{R}^2$
$\mathbf{v}_i(t)$	Velocity vector of robot r_i at time t	$\mathbb{R}^2$
COMMUNICATION AND SENSING		
R_c	Communication range radius	20 units
R_s	Sensing range radius	4 units
$d_{ij}(t)$	Euclidean distance between robots r_i and r_j	$\ \mathbf{p}_i(t) - \mathbf{p}_j(t)\ $
$N_i(t)$	RNG neighbor set of robot r_i	$\subseteq \mathcal{R}$
ROLE ASSIGNMENT MECHANISM		
$\rho_i(t)$	Role assignment: explorer or supporter	{explorer, supporter}
$E_i(t)$	Set of neighboring explorer robots	$\subseteq N_i(t)$
$S_i(t)$	Set of neighboring supporter robots	$\subseteq N_i(t)$
$L_i(t)$	Complete set of robots within communication range	$\subseteq \mathcal{R}$
α	Safety margin coefficient for communication range	$\alpha \in [0.8, 0.9]$
$C_n(i)$	Boolean connectivity indicator to base station	{True, False}
FLOCKING-INSPIRED SUPPORTER DYNAMICS		
γ	Weighting factor balancing supporter influences	$\gamma \in [0.3, 0.7]$
β_1	Scaling coefficient for supporter-supporter cohesion	$\beta_1 \in [0.3, 1.0]$
β_2	Scaling coefficient for supporter-explorer alignment	$\beta_2 \in [0.3, 1.0]$
$\mathbf{p}_{i,s}(t)$	Position influence from neighboring supporters	$\mathbb{R}^2$
$\mathbf{p}_{i,e}(t)$	Position influence from neighboring explorers	$\mathbb{R}^2$
COLLISION AVOIDANCE SYSTEM		
R_{avoid}	Collision avoidance activation radius	$[2, 4]$ units
k_{avoid}	Repulsive force scaling parameter	$k_{\text{avoid}} \in (0, 1)$
$\phi(d)$	Distance-based repulsion function	See Equation 4.9
STAGNATION DETECTION AND RECOVERY		
T_{stagnant}	Time window for stagnation assessment	$[10, 20]$ steps
D_{stagnant}	Displacement threshold for stagnation detection	$[1, 2]$ units
γ_{rec}	Recovery movement scaling factor	$\gamma_{\text{rec}} \in [0.1, 0.3]$
FRONTIER-BASED EXPLORATION		
$M_i(t)$	Local environmental map maintained by robot r_i	Grid cells
$F_i(t)$	Set of detected frontier points	$\subseteq \mathbb{R}^2$
$F_{\text{safe}}(t)$	Subset of frontiers within communication range	$\subseteq F_i(t)$

4.3 The DRBECM Algorithm

The originality of DRBECM lies in its ability to let robots intelligently switch between exploring and supporting tasks while operating in unknown, potentially hostile environments with obstacles. The primary objective is to explore these areas as rapidly as possible using limited resources, with the capability to strategically position robots as relay stations for continuous monitoring and maintaining communication links to the base station. This dynamic adaptation is governed by a set of distributed rules that respond to local network conditions and exploration progress.

4.3.1 Dynamic Role Assignment

The cornerstone of DRBECM is the mechanism that allows robots to autonomously switch between ‘explorer’ and ‘supporter’ roles. This decision is made locally by each robot based on its current situation, primarily its connectivity to the network. The goal is to ensure that explorer robots do not lose their connection to the base station as they venture into unknown territory.

Initially, all robots begin as explorers. A robot r_i currently in the ‘explorer’ role will switch to the ‘supporter’ role if it detects that moving further might jeopardize its connectivity to the base station. This is determined by checking if all its neighbors (including the base station) are beyond a certain threshold of its communication range, R_c . An explorer becomes a supporter to bridge a potential communication gap.

Conversely, a ‘supporter’ robot can switch back to being an ‘explorer’ if its role as a communication relay becomes redundant. This occurs when it has no explorers to support, is close to other supporters, and is not the critical link for another supporter further from the base. This ensures that robots do not remain passive supporters unnecessarily and can rejoin the exploration effort when possible.

The formal rule for a robot r_i to update its role ρ_i at time $t + 1$ is defined as follows:

$$\rho_i(t + 1) = \begin{cases} \text{supporter,} & \text{if } \rho_i(t) = \text{explorer, } E_i(t) \neq \emptyset, \\ & \forall r_j \in S_i(t) \cup B, d_{ij}(t) > \alpha R_c \\ \text{explorer,} & \text{if } \rho_i(t) = \text{supporter, } C_n(i) = \text{True,} \\ & E_i(t) = \emptyset, |S_i(t)| = 1, \forall r_j \in L_i(t), \\ & C_n(j) = \text{True, } \exists r_k \in L_i(t), \\ & \rho_k(t) = \text{supporter, } d_{kB}(t) < d_{iB}(t) \\ \text{otherwise,} & \rho_i(t) \end{cases} \quad (4.1)$$

where $E_i(t)$ and $S_i(t)$ are the sets of connected explorers and supporters, $L_i(t)$ is the set of all robots in communication range, $C_n(i)$ is a boolean indicating connection to the base station, and $\alpha \in [0, 1]$ is a safety margin parameter.

The neighbor sets $E_i(t)$ and $S_i(t)$ are fundamental to the role assignment mechanism, as they represent the RNG-based connections between robot r_i and other robots with specific roles. Figure 4.1 illustrates the RNG network topology where each robot maintains connections only to neighbors that satisfy the geometric RNG criterion. For instance,




Figure 4.1: RNG network topology in the DRBECM algorithm showing the distributed multi-robot system. Explorers (triangles) and supporters (diamonds) form RNG connections based on geometric proximity criteria. Purple edges represent explorer-supporter connections, orange dotted edges show supporter-supporter links, and black edges connect robots to the base station.

supporter S2 in the figure has neighbor sets $E_i(t) = \{E2\}$ (connected explorers) and $S_i(t) = \{S1, S4\}$ (connected supporters), forming its complete neighbor set $N_i(t) = \{E2, S1, S4\}$ used for role assignment decisions. The RNG ensures that $N_i(t) = E_i(t) \cup S_i(t) \cup \{B\}$ (when the base station is within range) forms a sparse, connected subgraph that facilitates efficient decision-making while minimizing communication overhead. Each robot uses only its local RNG neighborhood to determine role transitions, enabling fully distributed operation without requiring global network knowledge.

The role assignment mechanism in Equation 4.1 represents a formal finite state machine with two states $\{explorer, supporter\}$ and well-defined transition conditions based on local network topology. This formulation guarantees deterministic behavior where each robot's role transition depends solely on locally observable variables: neighbor connectivity ($E_i(t), S_i(t)$), base station connectivity ($C_n(i)$), and distance relationships ($d_{ij}(t), d_{iB}(t)$).

The state machine exhibits several critical properties:

- **Local Decidability:** All transition conditions can be evaluated using information available within the robot's communication range R_c .
- **Bounded State Space:** The binary role assignment prevents state explosion common in multi-agent systems [124].

[124]: Mosteo et al. (2017), 'Optimal role and position assignment in multi-robot freely reachable formations'

4.3.2 Explorer Robot Behavior: Frontier-Based Exploration

Explorer robots are tasked with maximizing information gain. They employ a frontier-based exploration strategy, where frontiers are the boundaries between known and unknown areas. However, to maintain connectivity, explorers do not simply move to the nearest frontier. Instead, they restrict their choices to a set of **safe frontiers**.

A frontier point f is considered safe for an explorer r_i if it lies within the communication range R_c of at least one existing supporter robot or the base station. This constraint ensures that by moving to the new location, the explorer can establish a reliable communication link, thereby extending the network's reach without becoming isolated. The set of safe frontiers, $F_{\text{safe}}(t)$, is formally defined as:

$$F_{\text{safe}}(t) = \left\{ f \in F_i(t) \mid \begin{array}{l} \exists r_j \in N_i(t), \rho_j(t) = \text{supporter}, \\ \|f - p_j(t)\| \leq R_c \end{array} \right\} \quad (4.2)$$

By selecting targets only from this subset, the system expands its explored area in a controlled manner, with supporters creating a communication backbone that follows the explorers.

4.3.3 Supporter Robot Behavior: Flocking-Inspired Positioning

The primary function of supporter robots is to maintain network connectivity. They act as mobile relay nodes, positioning themselves dynamically to bridge the distance between explorer robots and the base station. DRBECM uses a flocking-inspired model to govern supporter movement, which is computationally simple and highly adaptive.




Figure 4.2: Illustration of the safe and unsafe exploration areas based on communication range. The green cells represent safe frontiers within the robot's communication range of a supporter, while red cells are unsafe.

A supporter robot r_i adjusts its position based on two main influences: the positions of neighboring supporter robots and the positions of neighboring explorer robots. The final position is a weighted combination of these two factors, allowing it to balance between maintaining a strong chain with other supporters and moving to better support the explorers.

The position update is calculated as follows :

$$p_i(t+1) = \gamma p_{i,s}(t) + (1-\gamma)p_{i,e}(t) \quad (4.3)$$

where:

- $\gamma \in [0, 1]$ is a weighting factor that balances the influence of supporters and explorers.
- $p_{i,s}(t)$ is the position influence from other supporters and the base station, encouraging cohesion.

$$\mathbf{p}_{i,s}(t) = \mathbf{p}_i(t) + \beta_1 \frac{1}{|N_{\text{supporters}}(t)| + 1} \left(\sum_{r_j \in N_{\text{supporters}}(t)} (\mathbf{p}_j(t) - \mathbf{p}_i(t)) + (\mathbf{p}_B - \mathbf{p}_i(t)) \right) \quad (4.4)$$

- $p_{i,e}(t)$ is the position influence from explorers, drawing the supporter towards the active exploration front.

$$\mathbf{p}_{i,e}(t) = \mathbf{p}_i(t) + \beta_2 \frac{1}{|N_{\text{explorers}}(t)|} \left(\sum_{r_k \in N_{\text{explorers}}(t)} (\mathbf{p}_k(t) + \mathbf{v}_k(t) - \mathbf{p}_i(t)) \right) \quad (4.5)$$

- β_1 and β_2 are scaling constants.
- $N_{\text{supporters}}(t)$ and $N_{\text{explorers}}(t)$ are the sets of neighboring supporters and explorers, respectively.
- $\mathbf{v}_k(t)$ is the velocity of explorer r_k at time t .

This decentralized, nature-inspired approach enables the communication network to fluidly adapt its shape to follow the explorers as they expand into unknown territory.

4.3.4 Stagnation Detection and Collision Avoidance

To enhance the robustness of the exploration process, DRBECM incorporates two additional mechanisms:

Stagnation Detection: An explorer robot may get stuck due to obstacles or challenging terrain or being at the edges of the communication range, hindering exploration progress. DRBECM includes a stagnation detection mechanism where a robot monitors its displacement over a time window. A robot is considered *stagnant* if the maximum displacement over the last T_{stagnant} time steps is less than a threshold D_{stagnant} as detailed in Eq. 4.6.

$$\max_{t-T_{\text{stagnant}} \leq \tau \leq t} \|\mathbf{p}_i(t) - \mathbf{p}_i(\tau)\| \leq D_{\text{stagnant}} \quad (4.6)$$




Figure 4.3: Progression of DRBECM. The eight subfigures illustrate an example exploration scenario with $N = 4$ robots: (1) Explorer robots are initially deployed near the base station; (2)–(4) Early exploration, where the robots disperse and expand the explored area until they reach the base station’s communication limit; (5) Dynamic role assignment, as two robots switch to “supporter” roles; (6) Supporters move to maintain strong communication links; (7) “Support supporting,” where one supporter aids another to reach farther areas; (8) Demonstration of the flocking-inspired support mechanism, where supporter robots strategically position themselves to maintain connectivity.

where:

- T_{stagnant} is the chosen stagnation detection time window.
- D_{stagnant} is the chosen stagnation distance threshold.

Upon detecting stagnation, the explorer initiates a recovery behavior by moving towards the nearest supporter or the base station to reset its position and resume exploration:

$$\mathbf{p}_i(t+1) = \mathbf{p}_i(t) + \gamma_{\text{rec}}(\mathbf{p}_{\text{target}} - \mathbf{p}_i(t)) \quad (4.7)$$

where:

- $\mathbf{p}_{\text{target}}$ is the position of the closest supporter or the base station.
- γ_{rec} is a movement scaling factor for recovery.

Collision Avoidance: To ensure safe operation, each robot implements a collision avoidance function. This mechanism applies a repulsive force based on the proximity of its direct RNG neighbors. By focusing only on this small set of neighbors, the system efficiently prevents collisions without the computational overhead of considering all robots in the vicinity.

To ensure safe operation, each robot implements a collision avoidance mechanism that maintains a minimum distance from its RNG neighbors (repulsive force). Figure 4.4 illustrates the repulsion function $\phi(d)$ that

generates repulsive force based on the proximity of neighboring robots. By focusing on RNG neighbors, the robots can efficiently manage collision avoidance without excessive computational overhead.

The acceleration $\mathbf{a}_i(t)$ of robot r_i is adjusted as:

$$\mathbf{a}_i(t+1) = \mathbf{a}_i(t) + \sum_{r_j \in N_i(t)} \phi(d_{ij}(t)) \frac{\mathbf{p}_i(t) - \mathbf{p}_j(t)}{d_{ij}(t)} \quad (4.8)$$

where:

- $\phi(d)$ is a repulsion function defined as:

$$\phi(d) = \begin{cases} k_{\text{avoid}} \left(\frac{R_{\text{avoid}} - d}{R_{\text{avoid}}} \right), & \text{if } d < R_{\text{avoid}}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.9)$$

- R_{avoid} is the collision avoidance range,
- $k_{\text{avoid}} \in (0,1)$ is a scaling factor for the repulsion force,

4.3.5 Information Sharing and Map Updating

A critical component of DRBECM's decentralized strategy is the local sharing of information. Instead of relying on a central server, robots exchange data directly with neighbors within their communication range ($L_i(t)$) to collaboratively build an understanding of the environment. Each robot shares its local map and its list of known frontiers.

The map and frontier update rules for each robot r_i are as follows:

$$M_i(t+1) = M_i(t) \cup \text{SensorData}_i(t) \cup \bigcup_{r_j \in L_i(t)} M_j(t) \quad (4.10)$$

$$F_i(t+1) = \left(F_i(t) \cup \bigcup_{r_j \in L_i(t)} F_j(t) \right) \setminus M_i(t+1) \quad (4.11)$$

Equation 4.10 states that a robot's map, M_i , is updated by merging its current map with new data from its own sensors ($\text{SensorData}_i(t)$) and the maps received from its neighbors. Equation 4.11 updates the frontier set, F_i , by first taking the union of its own frontiers and those received from neighbors, and then removing any points that are now part of the newly updated map. This ensures frontiers are accurately managed as new areas become explored. This localized sharing strategy keeps robots synchronized without overwhelming the network, enhancing both scalability and overall exploration performance.

4.3.6 Main Algorithm

The distinct mechanisms described above (neighbor selection, role assignment, movement logic, and information sharing) are integrated into a single coherent operational loop executed by each robot. Algorithm 2 provides a consolidated view of this process, formalizing the steps a robot takes in each iteration of the exploration task.




Figure 4.4: Collision avoidance repulsion function $\phi(d)$ showing linear decrease in repulsive force with distance. No repulsion occurs beyond R_{avoid} .

Algorithm 2: DRBECM Main Loop (run at each robot r_i)**Input** : Constants: R_c, R_s **Output**: $p_i, v_i, \rho_i, M_i, N_i, F_i$

```

1  $p_i \leftarrow p_{\text{initial}};$ 
2  $v_i \leftarrow 0;$ 
3  $\rho_i \leftarrow \text{explorer};$ 
4  $M_i \leftarrow \{c \in [0, X_{\max}] \times [0, Y_{\max}] \mid \|p_i - c\| \leq R_s\};$ 
5  $F_i \leftarrow \partial M_i;$ 
6 while  $F_i \neq \emptyset$  do
7    $N_i \leftarrow \text{GetRNGNeighbors}(p_i);$ 
8    $\rho_i \leftarrow \text{NewRole}(N_i, \rho_i);$ 
9    $(p_i, v_i) \leftarrow \text{MoveRobot}(N_i, \rho_i, F_i, R_c);$ 
10   $(M_i, F_i) \leftarrow \text{MapUpdateShareInfo}(N_i, M_i, F_i, R_s);$ 
11 end

```

Algorithm 2 presents the proposed approach. Each robot r_i maintains its state (position p_i , velocity v_i , role ρ_i), a local map M_i , and a frontier set F_i . The algorithm initializes these variables (lines 1–5) and then enters its main exploration loop (lines 6–10).

The exploration process is driven by the frontier set F_i , which represents unexplored areas. In each iteration, the robot determines its RNG neighbors N_i (line 7). This step is crucial, as it defines the local network topology that informs subsequent decisions.

Based on this network information, the robot updates its role ρ_i (line 8, Equation (4.1)). This dynamic role assignment allows the system to adapt to the current exploration state, balancing between active exploration and network maintenance. Specifically, if an explorer robot detects that its movement might break network connectivity, it can switch to a supporter role, acting as a relay to maintain the communication link and further support the other explorers, as shown in the example scenario depicted in Fig. 4.3, specifically sub-figure (5), where two robots switch to the supporter role. Conversely, if a supporter robot determines that it's no longer needed as a relay (e.g., when explorers have moved closer to the base or other supporters), it can switch back to an explorer role to continue active exploration.

The robot's movement is then computed based on its current role and the positions of its neighbors (line 9, Equations (4.2)-(4.7)). Explorers move towards frontier points, while supporters adjust their positions to maintain network connectivity, ensuring a cohesive exploration effort. As illustrated in Fig. 4.3, sub-figures (6) and (7), the two explorer robots advance toward the frontiers, aided by supporters who maintain connectivity back to the base station.

Finally, the robot updates its local map and frontier set based on new sensor data and information shared with neighbors (line 10, Equations (4.10) and (4.11)). This step is critical as it integrates new information, potentially revealing new frontiers or closing existing ones, which directly influences the next iteration of the algorithm.

This cycle continues until the frontier set is empty, indicating complete exploration or the robots cannot move further due to the communication constraints.

Computational Complexity Analysis

The computational complexity of DRBECM per robot per time step is analyzed as follows:

- **Neighbor Selection (RNG):** $O(|N_i|^2)$ where $|N_i|$ is the number of robots within communication range
- **Role Assignment:** $O(|N_i|)$ for local connectivity assessment
- **Movement Planning:** $O(|F_i|)$ where $|F_i|$ is the number of local frontiers
- **Information Sharing:** $O(|M_i| \cdot |N_i|)$ where $|M_i|$ is the map size

The overall complexity is $O(|N_i|^2 + |F_i| + |M_i| \cdot |N_i|)$, which scales well with team size due to the locality of operations.

This computational complexity demonstrates significant efficiency advantages for distributed deployment. The quadratic term $O(|N_i|^2)$ for RNG neighbor selection is bounded by the geometric properties of RNG graphs, where nodes in two-dimensional space maintain an average degree of 2-3 neighbors regardless of network density [125]. This bounded-degree property ensures that $|N_i|$ remains small and approximately constant as the team size scales, effectively reducing the RNG computation to $O(1)$ for practical deployments.

[125]: Melchert (2013), 'Percolation thresholds on planar Euclidean relative-neighborhood graphs'

The locality-based complexity structure provides three key scalability advantages:

- **Distributed Load:** Computational burden is evenly distributed across all robots rather than concentrated at central controllers
- **Parallel Execution:** All robots execute their decision loops simultaneously without coordination overhead
- **Bounded Memory:** Local map storage $|M_i|$ grows with explored area per robot, not total team size

4.4 Performance Evaluation

The effectiveness of the DRBECM framework must be demonstrated through rigorous empirical evaluation that compares its performance against established multi-robot exploration approaches. This section presents a comprehensive experimental study that evaluates DRBECM's performance across multiple dimensions including exploration efficiency, network connectivity maintenance, and system robustness.

The evaluation methodology employs simulation-based experiments that systematically vary team size, environmental complexity, and operational parameters to assess algorithm performance under diverse conditions. Through comparison with both centralized and decentralized baseline approaches, this evaluation provides quantitative evidence of DRBECM's ability to achieve the dual objectives of efficient exploration and reliable connectivity maintenance while operating in a fully distributed manner.

4.4.1 Experimental Setup

All simulations were conducted with Python scripts using a grid-based simulation. The source code and configuration files for replicating these experiments are available at <https://github.com/HazemCHAABI/DRBECM>.

We conducted a series of simulations to evaluate the performance of our proposed method against existing multirobot exploration algorithms. The experiments were carried out on a 120×120 grid map, with a sensing range of 4 units and a communication range of 20 units. Each algorithm was tested with varying numbers of robots, ranging from 11 to 15, on 100 runs. We stop the simulation after 3000 steps if it is not complete. The following methods were compared:

[69]: Xia et al. (2020), 'Random Walks: A Review of Algorithms and Applications'

[126]: Jain et al. (2017), 'Comparative study of frontier based exploration methods'

- **DRBECM (Proposed Method)**: A decentralized method inspired by the flocking behavior, focusing on efficient exploration and connectivity maintenance.
- **Random Walk [69]**: A baseline decentralized approach where robots move randomly without coordination. In this method, each robot independently selects a random direction and moves in that direction for a predetermined number of steps before choosing a new random direction.
- **HCETIIC (Hybrid Cheetah Exploration Technique with Intelligent Initial Configuration) [56]**: Utilizes a central planner to coordinate robot movements and exploration tasks, considering the critical impact of initial robot positions. The algorithm aims to maximize exploration efficiency across different start configurations, including uniform, centralized, random, perimeter, clustered, and strategic positions.
- **Frontier Exploration (No Map Sharing) [126]**: A decentralized method in which robots individually explore frontiers without sharing information. Each robot maintains its own map of the environment and identifies frontier cells (unexplored areas at the boundary of known space) independently. Robots select frontiers to explore based on criteria such as distance and potential information gain, without coordinating their choices with other robots.
- **Frontier Exploration (Map Sharing) [126]**: A centralized solution that enables robots to share maps and coordinate their exploration. In this approach, robots periodically communicate their local maps to a central server. The shared information is used to construct a global map of the environment, allowing for more informed decision-making when selecting frontiers to explore.

The evaluation focused on three key metrics: exploration time (time to 100% coverage), exploration efficiency (coverage per unit distance traveled), and redundant exploration (overlap in explored areas).

4.4.2 Results and Analysis

The simulation results highlight the effectiveness of DRBECM in balancing the competing demands of exploration and connectivity.

Exploration Time

The exploration time, defined as the time required to achieve 100% coverage of the map, was analyzed across the exploration methods. Figures 4.5 and 4.6 depict the distribution of exploration times for the proposed method and baseline algorithms. The results demonstrate that the proposed method achieves exploration times comparable to the centralized Frontier with Map Sharing method. This is particularly notable given the decentralized nature of the proposed approach. The exploration time is significantly reduced for the proposed method compared to

decentralized baselines such as Random Walk and Frontier without Map Sharing. Our proposed method achieves a median exploration time that is 21.49% faster than Frontier without Map Sharing and 80.26% faster than Random Walk. Whilst, the centralized approach Frontier with Map Sharing, achieves the lowest median exploration time but requires global map sharing.




Figure 4.5: Distribution of exploration times across different methods




Figure 4.6: Distribution of exploration times across different methods

Exploration Efficiency

Measured as the ratio of total area covered to the total distance traveled by all robots, exploration efficiency reflects how effectively the system coordinates its movements. DRBECM achieved significantly higher efficiency than the decentralized baselines. Specifically, it **outperformed Frontier without Map Sharing by 126.8% and Random Walk by 291.2%.**

Although the centralized Frontier with Map Sharing method achieved the highest efficiency, it comes at the cost of requiring global information and a central controller, making it less suitable for infrastructure-less environments.




Figure 4.7: Comparison of exploration efficiency.

Redundant Exploration

Redundant exploration, measured as the total area explored by more than one robot, is a key indicator of coordination inefficiency. DRBECM achieved a substantial reduction in redundancy compared to other methods. The levels of redundant exploration were **55.93% lower than Frontier without Map Sharing** and **74.10% lower than Random Walk**. This demonstrates that the combination of frontier-based exploration and flocking-inspired support effectively minimizes overlapping work, allowing the system to cover the environment much more efficiently.




Figure 4.8: Redundant exploration across different methods.

4.5 Limitations and Future Enhancements

4.5.1 Discussion and Current Limitations

The collective results affirm the efficacy of the DRBECM method. By dynamically adapting robot roles based on local conditions, the framework achieves an effective balance between scalable, decentralized operation and the high performance typically associated with centralized in terms of exploration time reduction, redundant exploration minimization, and computational efficiency while maintaining network connectivity. Unlike methods such as HCETIIC or Frontier with Map Sharing, DRBECM does not depend on a central controller, making it inherently more robust and suitable for real-world applications where communication infrastructure is unreliable or non-existent.

The dynamic role-switching mechanism ensures that robots maintain connectivity while exploring unknown regions, outperforming decentralized baselines in both efficiency and redundancy. However, the framework operates under several constraints that limit its current applicability to idealized scenarios.

The role-switching logic can occasionally be overly conservative, leading to more supporters than necessary and slightly slower coverage than theoretically optimal. This conservative bias, while ensuring network connectivity, represents a trade-off between reliability and exploration speed that may not be optimal for all mission scenarios. This trade-off could be limited through adaptive parameter tuning based on real-time system performance, predictive modeling to anticipate connectivity requirements, and context-aware optimization that adjusts the balance between exploration aggressiveness and network maintenance based on mission criticality and environmental conditions.

The current model assumes stable communication links with predictable geometric properties, whereas real-world conditions involve intermittent connectivity, signal fading, and environmental interference patterns that cannot be captured by simple geometric models. Additionally, the framework assumes homogeneous robot capabilities and does not address teams with varying sensing ranges, communication equipment, or mobility constraints. Nevertheless, these assumptions provide a solid theoretical foundation that enables the framework to achieve its primary objective of balancing exploration efficiency with connectivity maintenance in a fully distributed manner. The demonstrated performance improvements over existing decentralized approaches validate the effectiveness of the dynamic role-based coordination mechanism, while the identified limitations serve as clear directions for algorithmic enhancements and real-world adaptations.

4.5.2 Algorithmic Enhancements

Several algorithmic improvements could significantly enhance the framework's capabilities and applicability across diverse deployment scenarios.

Heterogeneous Robot Integration

Extending the algorithms to accommodate robots with different sensing ranges, communication capabilities, processing power, and mobility constraints would require developing specialized role assignment mechanisms that leverage the strengths of different robot types. Robots with superior communication capabilities might be preferentially assigned supporter roles, while those with advanced sensing equipment could focus on exploration tasks.

Mixed-capability teams could employ hierarchical coordination where high-capability robots provide strategic coordination while resource-constrained robots execute local tasks. Such heterogeneous integration would significantly increase the practical applicability of the framework to real-world deployment scenarios where uniform robot teams are impractical.

Multi-Objective Optimization

Integration of explicit multi-objective optimization frameworks could better balance exploration efficiency, communication quality, energy consumption, and mission completion time. Current algorithms make these trade-offs implicitly through their design, but explicit formulations could enable better performance tuning for specific mission requirements and constraints [127, 128].

Pareto optimization approaches [129] could identify optimal trade-off frontiers for different operational scenarios, allowing mission planners to select configurations that best match their specific requirements. Dynamic weight adjustment could enable real-time optimization as mission conditions evolve.

[127]: Li et al. (2024), 'Large Language Model Based Multi-Objective Optimization for Integrated Sensing and Communications in UAV Networks'

[128]: Concha et al. (2024), 'Bayesian Optimization Framework for Efficient Fleet Design in Autonomous Multi-Robot Exploration'

[129]: Zhao et al. (2021), 'Pareto Optimal Multirobot Motion Planning'

Advanced Coordination Strategies

Development of hierarchical coordination mechanisms for mixed-autonomy scenarios, where some robots have higher-level planning capabilities or human operators provide strategic guidance, could bridge the gap between fully autonomous and supervised operations. Such systems could maintain the robustness of distributed operation while leveraging centralized intelligence when available.

Predictive models for robot trajectory and network evolution could enable more proactive decision-making, particularly for supporter positioning and frontier selection. By forecasting network topology changes based on current movement patterns, the system could make more informed decisions about role assignments and positioning strategies.

4.5.3 Deployment and Validation Requirements

Comprehensive field experiments are essential to validate the algorithms' performance under realistic conditions and identify necessary adaptations for practical deployment. The transition from simulation to real-world operation typically reveals challenges that cannot be anticipated in controlled environments.

Systematic field testing should encompass diverse environmental conditions including urban environments with signal interference, natural

terrain with communication range variations, and indoor scenarios with limited localization accuracy. Such studies would establish the framework's robustness and identify environmental factors that significantly impact performance.

Long-term autonomy experiments could evaluate performance in extended missions where battery life, hardware reliability, and environmental changes become significant factors. Understanding system behavior as robots are lost due to failures would guide the design of resilience mechanisms for practical deployments.

While these deployment challenges highlight important areas for future development, several of the fundamental limitations identified in this chapter—particularly the reliance on geometric communication models and the assumption of predictable signal propagation—are directly addressed through the machine learning-enhanced framework presented in the following chapter. By integrating real-world signal propagation data and RSSI prediction models, the enhanced approach moves beyond the idealized geometric assumptions of the current framework, demonstrating how data-driven connectivity prediction can significantly improve both decision-making accuracy and mission success rates in realistic deployment scenarios.

DRBECM-ML: Enhancing Exploration with Machine Learning-Driven Connectivity

5

5.1 Introduction

Building upon the distributed framework of DRBECM introduced in the previous chapter, this chapter presents **DRBECM-ML**, an enhanced multi-robot exploration system. This evolution of the algorithm addresses a key limitation of the original model: its reliance on purely geometric criteria for connectivity maintenance. While effective, using a fixed communication range (R_c) does not always capture the complexities of real-world wireless signal propagation, which can be affected by environmental factors not visible to the robots' sensors [130].

DRBECM-ML overcomes this by integrating a lightweight, data-driven machine learning model to predict network quality in real time. The core innovation is the replacement of abstract distance-based rules with realistic, predicted Received Signal Strength Indicator (RSSI) values. This allows the system to make more informed and robust decisions, ensuring network stability while efficiently exploring unknown areas.

The specific contributions of DRBECM-ML, as detailed in this chapter, are fourfold:

1. **A refined frontier selection process** that uses an RSSI prediction model trained on real-world signal propagation data to evaluate candidate exploration targets based on predicted connectivity metrics, with the goal of maintaining the quality of the connection while efficiently exploring unknown areas.
2. **An enhanced role-switching mechanism** that builds upon DRBECM's dynamic role assignment by incorporating RSSI-based evaluation to more effectively preserve communication in practical environments. This allows robots to transition between roles based on realistic signal strength assessments rather than purely geometric criteria.
3. **A lightweight machine learning integration** designed to operate within the strict computational constraints typical of embedded robotic platforms, leveraging offline-trained models that enable real-time connectivity predictions without the overhead of distributed learning.
4. **Integration of bio-inspired flocking dynamics** with data-driven connectivity prediction, addressing the challenge of maintaining network stability while exploring unknown areas efficiently through adaptive relay placement.

By embedding machine learning directly into the robotic control loop, DRBECM-ML represents a significant step towards creating more resilient and realistic multi-robot systems, particularly for emergency response scenarios where network reliability is critical. This chapter will detail the algorithmic enhancements, the development of the predictive model, and a comparative performance evaluation against the original DRBECM and other established methods.

5.1	Introduction	69
5.2	The DRBECM-ML Algorithm	70
5.2.1	Data-Driven Dynamic Role Assignment	70
5.2.2	Enhanced Frontier Selection	70
5.3	Machine Learning Model for RSSI Prediction	71
5.3.1	Data Collection Using FIT IoT-Lab	71
5.3.2	Data Preprocessing	72
5.3.3	Model Training and Hyper-parameter Tuning	73
5.4	Performance Evaluation	73
5.4.1	Machine Learning Model Evaluation	74
5.4.2	Comparative Simulation of Exploration Algorithms	75
5.5	Limitations and Research Directions	78
5.5.1	Discussion and Current Limitations	78
5.5.2	Advanced Machine Learning Integration	79
5.5.3	Real-World Validation and Extension	80

5.2 The DRBECM-ML Algorithm

DRBECM-ML retains the decentralized philosophy and the dynamic role-assignment framework of its predecessor. However, it significantly modifies the decision-making logic by incorporating a machine learning module for connectivity prediction. This allows for a more nuanced and adaptive approach to both role assignment and frontier selection.

5.2.1 Data-Driven Dynamic Role Assignment

The fundamental enhancement in DRBECM-ML is the shift from geometric-based connectivity checks to a data-driven model based on RSSI predictions. Instead of assuming a perfect circular communication range, each robot now assesses the quality of its communication links based on predicted signal strength. This allows the role-switching mechanism to react to potential signal degradation before a link is lost entirely.

The rule for updating a robot's role is formalized as follows: An *explorer* robot r_i continuously evaluates the predicted RSSI to its nearest *supporter* (or the base station). If the signal strength is predicted to fall below a predefined **strong** threshold (T_{strong}), it transitions to a *supporter* role to reinforce the network. Conversely, a *supporter* can switch back to an *explorer* if its presence as a relay is no longer critical and its link quality is above a **weak** threshold (T_{weak}), ensuring it can rejoin the exploration effort when safe to do so.

The formal rule for robot r_i to update its role $\rho_i(t+1)$ is:

$$\rho_i(t+1) = \begin{cases} \text{supporter,} & \text{if } \rho_i(t) = \text{explorer}, E_i(t) \neq \emptyset, \forall r_j \in S_i(t) \cup B \\ & \text{such that } \mathbf{RSSI}_{ij}(t) \leq \mathbf{T}_{strong} \\ \text{explorer,} & \text{if } \rho_i(t) = \text{supporter}, Cn(i) = \text{True}, E_i(t) = \emptyset, \\ & |S_i(t)| = 1, \forall r_j \in L_i(t), Cn(j) = \text{True}, \exists r_k \in L_i(t) \\ & \text{with } \rho_k(t) = \text{supporter,} \\ & d_{kB}(t) < d_{iB}(t) \text{ and } \mathbf{RSSI}_{ij}(t) \geq \mathbf{T}_{weak} \\ \rho_i(t), & \text{otherwise} \end{cases} \quad (5.1)$$

where:

- ▶ $E_i(t)$ and $S_i(t)$ are the sets of neighboring explorers and supporters.
- ▶ $\mathbf{RSSI}_{ij}(t)$ is the measured RSSI between robots i and j .
- ▶ T_{strong} and T_{weak} are the RSSI thresholds for role switching.
- ▶ $Cn(i)$ indicates if robot i is connected to the base station B .
- ▶ $L_i(t)$ is the set of all neighboring robots.

5.2.2 Enhanced Frontier Selection

In DRBECM-ML, the process of selecting an exploration target is enhanced with connectivity awareness. After identifying frontier cells, an explorer robot performs a post-processing step to evaluate each candidate frontier based on its connectivity potential.

The enhanced frontier selection mechanism evaluates each candidate frontier point f by predicting the RSSI value between f and the robot's connected supporters or the base station. The set of safe frontiers is defined as:

$$F_{safe}(t) = \left\{ f \in F_i(t) \mid \max_{r_j \in S_i(t)} \widehat{RSSI}_{fj}(t) \geq T_{rssi} \right\} \quad (5.2)$$

where:

- ▶ $\widehat{RSSI}_{fj}(t)$ is the predicted RSSI value between frontier f and supporter robot r_j
- ▶ T_{rssi} is the minimum acceptable predicted RSSI threshold

The frontier with the maximum predicted RSSI among all candidates is selected as the exploration target, ensuring that exploration proceeds along paths that maintain strong communication links.

5.3 Machine Learning Model for RSSI Prediction

The development of accurate connectivity prediction models requires the integration of real-world communication data into the algorithmic framework. This section details the comprehensive process of developing, training, and validating a machine learning model for Received Signal Strength Indicator (RSSI) prediction based on extensive experimental data collected from a large-scale wireless sensor network testbed.

The methodology encompasses the entire machine learning pipeline, from data collection and preprocessing through model selection and performance evaluation. By leveraging real-world signal propagation data rather than idealized theoretical models, this approach enables DRBECM-ML to make connectivity decisions based on realistic assessments of communication quality. The resulting model provides the foundation for enhancing exploration decisions with data-driven insights while maintaining the computational efficiency required for real-time robotic applications.

5.3.1 Data Collection Using FIT IoT-Lab

The machine learning model for RSSI prediction is trained using real-world data collected from the FIT IoT-Lab testbed [131] in Lille. The experimental setup involved 126 M3 nodes. Each M3 node (visible in figure 5.2 as small devices attached to the ceiling with LED indicators) features a 32-bit ARM Cortex-M3 (STM32F103REY) microcontroller operating at up to 72 MHz, with 64 KB of RAM and 256 KB of ROM. The radio interface uses an AT86RF231 chip supporting 2.4 GHz IEEE 802.15.4 communication at a maximum bandwidth of 256 kbit/s.

Experimental Parameters

The data collection process employed the following protocol:

- ▶ One node designated as broadcaster transmitting at 0dBm power
- ▶ 125 remaining nodes acting as listeners recording RSSI measurements
- ▶ Periodic message transmission with timestamp, node identifier, and channel information




Figure 5.1: The IoT-LAB M3 board [131]

[131]: Adjih et al. (2015), 'FIT IoT-LAB: A large scale open experimental IoT testbed'



Figure 5.2: FIT IoT-LAB Lille ©Inria / Photo C. Morel

- Systematic recording of received signal strength across various node positions

Each measurement record contains timestamp, node identifier, message ID, channel number, and measured RSSI in dBm, providing a comprehensive dataset of real-world signal propagation characteristics. Table 5.1 shows a representative sample of the collected data structure.

Table 5.1: Reception log raw data snippet

Timestamp	Node ID	MsgId	Len ^a	Rssi ^b	X	Y	Z	Distance ^c
1712131915.240118	m3-182	65	121	-54.0	29.31	5.3	9.6	11.3842
1712131917.181086	m3-237	66	121	-58.0	31.71	12.5	9.6	13.6821
1712131919.186740	m3-228	67	121	-55.0	25.71	11.3	9.6	7.5895
1712131921.177211	m3-138	68	121	-52.0	28.11	0.3	8.5	12.9356
1712131923.185310	m3-142	69	121	-65.0	32.91	0.3	8.5	16.8086
1712131925.171761	m3-135	70	121	-55.0	24.51	0.3	7.6	10.6752
1712131927.182682	m3-149	71	121	-55.0	33.79	0.5	9.4	17.4378
1712131929.182901	m3-141	72	121	-67.0	31.71	0.3	7.6	15.8808
1712131931.184665	m3-256	73	121	-64.0	32.91	16.1	9.6	16.0997

^a Message length in bytes

^b RSSI in dBm

^c Distance from the broadcaster node

5.3.2 Data Preprocessing

To ensure the quality and consistency of the data used for training the machine learning model, the raw logs from the FIT IoT-Lab testbed were subjected to a rigorous preprocessing pipeline. This process involved three key steps:

1. **Cleaning and Parsing:** The initial step involved filtering the raw data to discard any incomplete or invalid entries. For each valid log entry, the RSSI value, which was saved as a string (e.g., "-60dBm"), was parsed and converted into a numerical format (e.g., -60) suitable for model training.
2. **Outlier Removal:** Wireless signal propagation is inherently noisy, leading to sporadic and extreme RSSI readings caused by hardware anomalies or transient environmental interference [133]. To create a robust model based on typical signal patterns, we employed an Interquartile Range (IQR) filter. The first (Q1, 25th percentile) and third (Q3, 75th percentile) quartiles of the RSSI distribution were calculated. Any data point falling outside the range of $[Q_1 - 1.5 \times \text{IQR}, Q_3 + 1.5 \times \text{IQR}]$ was classified as an outlier and excluded from the training dataset.
3. **Coordinate Normalization:** The model is intended to predict RSSI based on the relative positions of robots, not their absolute coordinates in the environment. To facilitate this, the raw (x, y) coordinates of each listening node were normalized by shifting and scaling them relative to the position of the broadcasting node. This ensures that the model learns distance-based signal attenuation patterns effectively and independently of the system's global position.




Figure 5.3: Representation of the Interquartile Range - Wikipedia [132]

[133]: Tuset-Peiro et al. (2020), 'A Dataset to Evaluate IEEE 802.15.4g SUN for Dependable Low-Power Wireless Communications in Industrial Scenarios'

The final preprocessed dataset comprises approximately 15,000 measurement records, each containing normalized X,Y coordinates relative to the broadcaster position and the corresponding measured RSSI values

in dBm. Table 5.2 shows a representative sample of the processed data structure.

Following these steps, the refined dataset was split into training (80% of the collected data) and testing (20%) subsets, with the test data being withheld from the training process to ensure an unbiased evaluation of the model's performance.

X	Y	RSSI (dBm)
8.4	-8.60	-66.0
4.8	0.00	-56.0
3.6	0.46	-55.0
7.2	0.96	-59.0
13.2	-1.20	-64.0

Table 5.2: Sample of preprocessed dataset with normalized coordinates and RSSI measurements

5.3.3 Model Training and Hyper-parameter Tuning

We evaluated a suite of regression algorithms (described in subsection 5.4.1) to identify the most suitable model for predicting RSSI values based on the relative coordinates of the transmitter and receiver. The evaluated models included Decision Tree, Random Forest, Gradient Boosting, LightGBM, CatBoost, XGBoost, Support Vector Regressor (SVR), and K-Nearest Neighbors Regressor. Each model's hyper-parameters were optimized using a 3-fold cross-validated grid search, with the objective of minimizing the Root Mean Square Error (RMSE). In addition to accuracy metrics like Mean Absolute Error (MAE) and the R^2 score, we also measured two critical resource-related metrics:

- **Inference Time:** The average time required for the model to generate a prediction. This is a critical factor for real-time decision-making on robotic platforms.
- **Memory Usage:** The memory footprint of the deployed model. This is crucial for ensuring the model can run on resource-constrained embedded systems.

Inference time and memory usage were respectively measured with Python's `time` and `sys.getsizeof` libraries.

The final selected model was chosen based on its ability to provide a strong balance between high prediction accuracy and low computational overhead, making it suitable for integration into the DRBECM-ML framework.

5.4 Performance Evaluation

The performance of the DRBECM-ML framework was evaluated in two stages using Python's *scikit-learn* library for machine learning model assessment and a custom built multi-robot exploration simulator implemented in Python. First, we assessed the performance of the various machine learning models for connectivity prediction using standard regression evaluation metrics. Second, we conducted a comparative simulation study to measure the exploration performance of DRBECM-ML against the original DRBECM and other established algorithms in a grid-based simulation environment.

5.4.1 Machine Learning Model Evaluation

To rigorously assess the performance of the connectivity prediction models, we used several standard regression metrics. Good performance is characterized by low error values (RMSE and MAE), an R^2 score close to 1, and minimal inference time and memory usage.

- **Root Mean Square Error (RMSE):** Measures the square root of the average of the squared differences between predicted and actual RSSI values. A low RMSE indicates a small average error.

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2} \quad (5.3)$$

- **Mean Absolute Error (MAE):** Measures the average absolute difference between predicted and actual values.

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i| \quad (5.4)$$

- **Coefficient of Determination (R^2 Score):** Represents the proportion of the variance in the dependent variable that is predictable from the independent variable(s). An R^2 score close to 1 indicates that the model explains most of the RSSI variability.

$$R^2 = 1 - \frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{\sum_{i=1}^n (\bar{y} - y_i)^2} \quad (5.5)$$

where $\hat{y}_i$ is the predicted RSSI, y_i is the actual RSSI, $\bar{y}$ is the mean of the actual RSSI values, and n is the number of samples.

Inference time and memory usage were measured with Python's `time` and `tracemalloc` libraries, respectively, to ensure the selected model could run efficiently on resource-constrained robotic platforms.

Results of Model Comparison

Table 5.3 summarizes the performance of the evaluated regression models. The results clearly indicate that tree-based methods significantly outperform other models like Linear Regression and SVR in terms of accuracy. For instance, the Decision Tree model achieves an RMSE of 1.194825 and an R^2 score of 0.966, indicating a very strong fit to the data. This represents a 76.1% reduction in RMSE compared to AdaBoost and a 79.8% reduction compared to Linear Regression. Critically, the tree-based models also offer excellent computational performance. The Decision Tree model has an exceptionally low inference time of 0.0025 seconds, making it ideal for real-time applications. While XGBoost offers the lowest memory footprint (16,548 bytes), the Decision Tree model's memory usage (24,953 bytes) is also modest and well within the capabilities of modern robotic platforms. Given its excellent balance of high accuracy, extremely low inference time, and manageable memory usage, the Decision Tree model was selected for integration into the DRBECM-ML framework.

The evaluated regression models encompass diverse algorithmic approaches to capture the relationship between spatial coordinates and

RSSI values. **Tree-based methods** including Decision Tree, Random Forest, Extra Trees, Gradient Boosting, LightGBM, CatBoost, and XGBoost construct hierarchical decision structures that can capture non-linear spatial patterns in signal propagation [134–138]. **Linear methods** such as Linear Regression and Elastic Net assume linear relationships between position and signal strength, providing computational efficiency but potentially limited modeling capacity for complex propagation environments [139]. **Ensemble methods** like AdaBoost combine multiple weak learners to improve prediction accuracy through sequential error correction [140]. **Instance-based methods** such as K-Nearest Neighbors predict RSSI values based on spatial proximity to training samples, making no explicit model assumptions [141]. **Kernel methods** including Support Vector Regression employ high-dimensional transformations to capture complex non-linear relationships while maintaining computational tractability through kernel functions [142, 143].

[134]: Chen et al. (2016), 'XGBoost: A Scalable Tree Boosting System'
 [135]: Ke et al. (2017), 'LightGBM: A Highly Efficient Gradient Boosting Decision Tree'
 [136]: Prokhorenkova et al. (2018), 'CatBoost: Unbiased boosting with categorical features'
 [137]: Breiman (2001), 'Random Forests'
 [138]: Geurts et al. (2006), 'Extremely randomized trees'
 [139]: Zou et al. (2005), 'Regularization and variable selection via the elastic net'
 [140]: Freund et al. (1997), 'A decision-theoretic generalization of on-line learning and an application to boosting'

Table 5.3: Model Performance Comparison - Grouped by Algorithmic Approach

Model	RMSE	MAE	R ² Score	Inference Time (s)	Memory Usage (bytes)
Tree-based Methods					
Decision Tree	1.194825	0.549787	0.966369	0.002500	24953
Random Forest	1.194740	0.549707	0.966374	0.251301	28806
Extra Trees	1.194825	0.549787	0.966369	0.062986	28694
Gradient Boosting	1.194830	0.549920	0.966369	0.112972	25863
LightGBM	1.194848	0.550685	0.966368	0.005000	25016
CatBoost	1.194825	0.549787	0.966369	0.007545	26706
XGBoost	1.194824	0.550040	0.966370	0.007999	16548
Linear Methods					
Linear Regression	5.922185	4.839405	0.173792	0.002001	24894
Elastic Net	5.922161	4.839469	0.173799	0.001003	25063
Ensemble Methods					
AdaBoost	5.013599	4.249386	0.407860	0.004002	25566
Instance-based Methods					
K-Nearest Neighbors	1.262389	0.561584	0.962459	0.013499	29979
Kernel Methods					
Support Vector Regression	5.105113	3.741817	0.386046	1.938107	25022

5.4.2 Comparative Simulation of Exploration Algorithms

To evaluate the impact of the ML-based connectivity model on exploration performance, we conducted a series of simulations comparing DRBECM-ML with the original DRBECM and the same set of baseline algorithms from Chapter 4.

Simulation Settings

The experiments were conducted in a 120x120 grid environment. The number of robots varied from 11 to 15, and each configuration was run 30 times to ensure statistical significance. A run was terminated if it reached the maximum of 3000 simulation steps. Based on [144], the RSSI thresholds for role-switching were set to $T_{strong} = -65dBm$ and $T_{weak} = -75dBm$. Key parameters are summarized in Table 5.4.

The methods compared were:

Table 5.4: Simulation Parameters

Parameter	Value
Grid Size	120 × 120
Sensing Range	4 units
Number of Robots	11–15
Maximum Simulation Steps	3000
T_{strong}	-65 dBm
T_{weak}	-75 dBm
Number of Runs per Configuration	30

- ▶ **DRBECM-ML:** The proposed machine learning-enhanced framework that uses RSSI prediction models for connectivity-aware role assignment and frontier selection decisions.
- ▶ **DRBECM:** The original dynamic role-based exploration framework that relies on geometric distance criteria for connectivity assessment and role switching between explorer and supporter modes.
- ▶ **Random Walk** [69]: A baseline decentralized approach where robots move in randomly selected directions for predetermined durations, providing a lower-bound performance benchmark without any coordination or intelligent exploration strategy.
- ▶ **HCETIIC** [56]: A centralized method that combines deterministic exploration planning with metaheuristic optimization algorithms and strategic initial robot positioning to maximize coverage efficiency.
- ▶ **Frontier Exploration (No Map Sharing)** [126]: A decentralized approach where robots independently identify and explore frontier cells (boundaries between known and unknown areas) using only local environmental knowledge without coordinating exploration targets with teammates and without any communication constraints.
- ▶ **Frontier Exploration (Map Sharing)** [126]: A centralized variant that enables global map sharing and coordinated frontier assignment through a central server, allowing robots to avoid redundant exploration and optimize coverage through informed decision-making.

Successful Exploration Rate

We define a "successful" exploration as one that covers at least 99% of the environment within the 3000-step limit. As shown in Table 5.5, DRBECM-ML achieves a successful exploration rate of 99.3%. This is a dramatic improvement over the original DRBECM's 83.6% success rate and highlights the primary benefit of the ML-driven approach. By making more realistic connectivity assessments, robots avoid paths that would lead to network fragmentation, allowing them to complete the exploration task far more reliably. The performance of DRBECM-ML is comparable to the Frontier-based methods, which also achieve very high success rates.

Table 5.5: Percentage of successful exploration runs for each method.

Method	Successful Exploration (%)
DRBECM	83.6
DRBECM-ML	99.3
Frontier (Map Sharing)	100.0
Frontier (No Map Sharing)	99.6
HCETIIC	45.6
Random Walk	61.4

Exploration Time

Figure 5.4 shows the distribution of exploration times for the successful runs. Interestingly, DRBECM exhibits a slightly lower median exploration time than DRBECM-ML. This suggests that the purely geometric criteria of the original algorithm can lead to more aggressive, and therefore faster, exploration. However, this speed comes at the cost of robustness, as reflected in its lower success rate. DRBECM-ML accepts a modest trade-off in exploration speed for a significant gain in reliability. It is more cautious, as it is guided by a more realistic model of signal propagation, but this caution ensures mission completion. As expected, the centralized Frontier (Map Sharing) approach remains the fastest, leveraging its global map to optimize exploration, but at the cost of high communication overhead and a single point of failure.

Redundant Exploration

Redundant exploration measures the degree of overlap in the areas covered by different robots. Figure 5.5 compares the redundancy levels across methods. The results show that DRBECM-ML has a slightly higher median redundancy than the original DRBECM. This is a direct consequence of its more conservative nature; by incorporating real-world RSSI predictions, the algorithm guides robots along paths that prioritize connectivity, which may lead to slightly more overlap compared to the more direct, geometrically-driven paths of DRBECM. However, despite this minor increase, DRBECM-ML still demonstrates significantly lower redundancy than other decentralized methods like Random Walk and Frontier (No Map Sharing), indicating its overall efficiency in coordinating exploration.




Figure 5.4: Box plot of exploration times for the different exploration methods. The y-axis represents the steps needed to achieve the target coverage, while the x-axis lists the methods under comparison.




Figure 5.5: Comparison of redundant exploration across the different methods.

5.5 Limitations and Research Directions

5.5.1 Discussion and Current Limitations

The development and evaluation of DRBECM-ML demonstrate the significant benefits of integrating data-driven machine learning models into distributed robotic exploration frameworks. By replacing the idealized geometric connectivity model of DRBECM with a realistic RSSI prediction model trained on real-world data, DRBECM-ML achieves a far more robust and resilient system.

The primary finding is the substantial increase in successful exploration rate from 83.6% for DRBECM to 99.3% for DRBECM-ML. This confirms that more accurate assessment of communication links prevents robots from becoming isolated, which was a key failure mode in the original algorithm. This enhancement comes at the cost of modest increases in exploration time and redundant coverage, reflecting the inherently more cautious approach that prioritizes network integrity over raw speed.

However, several limitations constrain the current applicability of the machine learning integration. The RSSI prediction models, though trained on real-world data from the FIT IoT-Lab testbed, were developed using indoor measurements with controlled conditions. Signal propagation characteristics in outdoor environments, underground scenarios, or cluttered industrial settings may differ significantly from the training data, potentially affecting prediction accuracy and decision-making quality.

The current framework relies on a single connectivity metric (RSSI) and does not incorporate other important communication quality indicators such as packet loss rates, latency variations, or bandwidth availability. Additionally, the static nature of the pre-trained model limits adaptability to novel environments or changing conditions that differ significantly from the training scenarios.

5.5.2 Advanced Machine Learning Integration

The machine learning component offers numerous opportunities for enhancement that could significantly improve the framework's adaptability and performance across diverse deployment scenarios.

Extended Quality-of-Service Prediction

Expanding the connectivity prediction models to incorporate additional QoS metrics such as latency, bandwidth availability, and packet loss rates could enable more sophisticated decision-making that considers the full spectrum of communication requirements. Multi-output regression models could simultaneously predict multiple communication quality indicators, providing richer information for role assignment and frontier selection decisions.

Time-series prediction models could forecast communication quality trends, enabling proactive decision-making that anticipates degrading conditions rather than merely reacting to current measurements. Such predictive capabilities would be particularly valuable for supporter positioning and exploration path planning.

Online Learning and Adaptation

Development of online learning mechanisms would allow robots to adapt their connectivity models based on local experience, improving performance in novel environments or scenarios that differ significantly from training conditions. Incremental learning algorithms could continuously update prediction models using locally observed signal propagation patterns.

Transfer learning approaches could enable rapid adaptation of connectivity models to new environments by leveraging knowledge gained from previous deployments. This could significantly reduce the data collection requirements for new deployment scenarios while maintaining prediction accuracy.

Collaborative Learning Approaches

Investigation of federated learning approaches could enable collaborative model improvement across robot teams while preserving data locality and reducing communication overhead. In large-scale deployments where centralized model training is impractical, federated approaches could enable continuous improvement of connectivity predictions through distributed learning from local experiences.

Consensus-based learning could ensure that robots operating in the same environment converge to consistent connectivity models, improving coordination effectiveness while maintaining the benefits of distributed operation.

Advanced Decision Integration

Deep reinforcement learning could be applied to optimize the role-switching and frontier selection strategies based on mission-specific objectives and environmental characteristics. Rather than using fixed decision rules, RL agents could learn optimal policies that adapt to specific deployment scenarios and mission requirements.

Multi-modal sensor fusion could enhance connectivity prediction accuracy by incorporating additional environmental sensors beyond simple position information. Factors such as local obstacle density, electromagnetic interference levels, or atmospheric conditions could be integrated into prediction models to improve accuracy in diverse environments.

5.5.3 Real-World Validation and Extension

The machine learning enhancements require extensive validation under realistic deployment conditions to establish their effectiveness beyond controlled simulation environments. Field experiments across diverse environments would test the robustness of the prediction models and identify necessary adaptations for practical applications.

Integration with physical robot platforms would reveal implementation challenges related to computational constraints, sensor noise, communication delays, and model deployment complexity. Such practical considerations often require algorithmic modifications not apparent in simulation studies but crucial for successful deployment.

Extension to dynamic environments with obstacles, time-varying conditions, and complex terrain would demonstrate the framework's applicability to realistic emergency response scenarios where environmental condition change rapidly and unpredictably.

In conclusion, DRBECM-ML represents a significant advancement toward more realistic and reliable multi-robot exploration systems. The integration of data-driven connectivity prediction addresses a fundamental limitation of geometric models while maintaining the computational efficiency required for real-time operation. The framework provides a solid foundation for incorporating advanced machine learning techniques while preserving the fundamental advantages of distributed coordination and connectivity awareness.

This thesis has addressed the fundamental challenge of distributed multi-robot exploration in unknown environments while maintaining robust communication connectivity. Through a systematic research approach that progressed from communication infrastructure analysis to algorithmic development and real-world enhancement, this work has contributed novel solutions that balance exploration efficiency with network reliability in infrastructure-denied scenarios. The research has demonstrated that distributed multi-robot systems can achieve performance comparable to centralized approaches while providing superior robustness and scalability through innovative role-based coordination and machine learning-driven connectivity prediction.

6.1 Summary of Contributions

6.1.1 Comparative Analysis of Communication Topologies

The first contribution of this thesis is a comprehensive empirical analysis of dynamic communication topology strategies for multi-robot systems. Through extensive NS-3 simulations, we systematically compared three prominent graph-based approaches: K-Nearest Neighbors (KNN), Relative Neighborhood Graph (RNG), and k-Relative Neighborhood Graph (K-RNG). The evaluation was conducted across multiple scenarios to assess their effectiveness under varying network conditions and operational requirements.

The comparative analysis revealed that RNG provides the optimal balance between communication efficiency and network connectivity, achieving superior Packet Delivery Ratio (PDR) performance in peer-to-peer scenarios with initial PDR values of approximately 95% for 5 robots, declining to 33% with 60 robots while consistently outperforming other strategies. In multi-hop routing scenarios, RNG maintained competitive performance with median PDR values of 31-32%, demonstrating stable routing behavior across different network conditions.

K-RNG strategies showed marginal improvements in multi-hop scenarios, with K-RNG ($k=3$) achieving the highest median PDR of approximately 32-33%. However, this performance came at the cost of increased computational complexity due to the geometric lune calculations required for interior point evaluation. The parameter k in K-RNG provides tunable control over network density, but requires careful optimization for specific deployment scenarios.

KNN strategies consistently underperformed across both evaluation scenarios, with the best variant KNN ($k=3$) achieving only 30-31% median PDR in multi-hop routing. The purely distance-based neighbor selection approach, combined with the lack of connectivity guarantees, proved particularly unsuitable for scenarios requiring reliable communication paths.

This empirical evidence established RNG as the communication backbone for subsequent algorithmic developments, providing a parameter-free

6.1	Summary of Contributions	81
6.1.1	Comparative Analysis of Communication Topologies	81
6.1.2	Dynamic Role-Based Exploration Framework (DRBECM)	82
6.1.3	Machine Learning-Enhanced Framework (DRBECM-ML)	82
6.2	Discussion of Results . .	83
6.2.1	Performance Analysis and Trade-offs	83
6.2.2	Scalability and Robustness Characteristics . . .	85
6.2.3	Limitations and Considerations	85
6.3	Perspectives	86
6.3.1	Algorithmic and Theoretical Extensions	86
6.3.2	Technological Enhancements	88
6.3.3	Application-Driven Research Directions . . .	90
6.3.4	Ethical and Deployment Considerations for Emergency Response Systems	91

solution that eliminates the complexity of connectivity parameter tuning while ensuring guaranteed network connectivity properties. The sparse yet connected nature of RNG graphs makes them particularly suitable for resource-constrained robotic applications where minimizing communication overhead is crucial.

6.1.2 Dynamic Role-Based Exploration Framework (DRBECM)

The second major contribution is the development of the Dynamic Role-Based Exploration with Connectivity Maintenance (DRBECM) algorithm, a novel distributed multi-robot exploration framework that addresses the fundamental trade-off between exploration efficiency and communication reliability. The algorithm introduces a dynamic role-switching mechanism where robots autonomously alternate between explorer and supporter roles based on local connectivity assessments and exploration needs.

The DRBECM framework operates through several integrated components. Explorers employ frontier-based exploration strategies enhanced with safe frontier selection mechanisms to maximize information gain while maintaining communication links. The safe frontier selection ensures that exploration targets remain within communication range of supporter robots or the base station, preventing network fragmentation during exploration. Supporters utilize flocking-inspired positioning algorithms to form adaptive relay chains that preserve network connectivity as the team expands into unknown territory.

The framework incorporates robust operational mechanisms including collision avoidance based on RNG neighbor relationships, stagnation detection using displacement monitoring over time windows, and recovery behaviors that guide robots back to communication range when isolated. The information sharing mechanism enables local map updates and frontier discovery through neighbor communication, maintaining system coherence without requiring centralized coordination.

Experimental validation demonstrated that DRBECM achieves significant performance improvements over baseline approaches. The algorithm completed exploration tasks 21.49% faster than decentralized frontier-based methods without map sharing and 80.26% faster than random walk approaches. Redundant exploration was reduced by 55.93% compared to frontier-based approaches without map sharing and by 74.10% compared to random walk methods, demonstrating effective coordination despite the distributed architecture.

The algorithm's computational complexity of $O(|N_i|^2 + |F_i| + |M_i| \cdot |N_i|)$ per robot per time step scales well with team size due to the locality of operations, where $|N_i|$ represents the number of RNG neighbors, $|F_i|$ is the number of local frontiers, and $|M_i|$ is the local map size. This computational efficiency makes the framework suitable for real-time deployment on resource-constrained robotic platforms.

6.1.3 Machine Learning-Enhanced Framework (DRBECM-ML)

The third contribution extends DRBECM through the integration of machine learning-driven connectivity prediction, addressing the limita-

tions of purely geometric communication models in real-world scenarios. DRBECM-ML incorporates a lightweight Received Signal Strength Indicator (RSSI) prediction model trained on comprehensive real-world signal propagation data collected from the FIT IoT-Lab testbed.

The data collection process involved 126 M3 nodes equipped with 32-bit ARM Cortex-M3 microcontrollers and AT86RF231 radio interfaces supporting IEEE 802.15.4 communication at 2.4 GHz. One node served as a broadcaster transmitting at 0dBm power, while 125 listener nodes recorded RSSI measurements along with temporal and spatial metadata. This extensive dataset provided realistic signal propagation characteristics that capture environmental effects not modeled by geometric approaches.

Data preprocessing included cleaning and parsing of raw logs, outlier removal using Interquartile Range (IQR) filtering to eliminate sporadic signal fluctuations, and coordinate normalization relative to broadcaster positions. The refined dataset was systematically evaluated using multiple regression algorithms including Decision Tree, Random Forest, Gradient Boosting, LightGBM, CatBoost, XGBoost, Support Vector Regressor, and K-Nearest Neighbors Regressor.

Comparative evaluation revealed that tree-based methods significantly outperformed other approaches, with Decision Tree models achieving Root Mean Square Error (RMSE) values of 1.194825 and R^2 scores of 0.966, representing a 76.1% reduction in RMSE compared to AdaBoost and 79.8% reduction compared to Linear Regression. The Decision Tree model demonstrated exceptional computational efficiency with inference times of 0.0025 seconds and memory usage of 24,953 bytes, making it ideal for real-time robotic applications.

The framework replaces geometric connectivity criteria with data-driven assessments using empirically determined RSSI thresholds: T_{strong} for role-switching decisions and T_{weak} for recovery transitions. These thresholds were established based on IEEE 802.15.4 communication characteristics.

The enhanced frontier selection process evaluates candidate exploration targets based on predicted connectivity metrics, with frontiers selected based on maximum predicted RSSI values among all candidates. This ensures that exploration proceeds along paths that maintain strong communication links while maximizing information gain.

The integration of machine learning resulted in a dramatic improvement in mission success rates, increasing from 83.6% for the original DRBECM to 99.3% for DRBECM-ML when defining successful exploration as achieving at least 99% environmental coverage within the simulation limit of 3000 time steps. This enhancement demonstrates the significant value of incorporating realistic communication characteristics into robotic decision-making processes, particularly in scenarios where communication reliability is critical for mission success.

6.2 Discussion of Results

6.2.1 Performance Analysis and Trade-offs

The experimental results across all contributions demonstrate consistent improvements over baseline approaches while revealing important trade-

offs between different design choices. The systematic evaluation provides insights into the effectiveness of various algorithmic components and their interactions under different operational conditions.

In the communication topology analysis, RNG's superior performance in peer-to-peer scenarios stems from its creation of sparse, geometrically sound networks that minimize packet collisions and interference while maintaining essential connectivity properties. The parameter-free nature of RNG eliminates the need for optimization in dynamic environments, making it particularly attractive for autonomous systems where adaptive behavior must occur without human intervention.

However, K-RNG variants showed marginally better performance in multi-hop routing scenarios, suggesting that controlled redundancy can enhance path reliability in complex routing situations. The geometric lune criterion with controlled interior point tolerance appears to create more robust multi-hop paths compared to purely distance-based approaches. K-RNG ($k=3$) achieved the optimal balance between connectivity redundancy and computational overhead, though this required careful parameter tuning.

The DRBECM algorithm's effectiveness derives from its dynamic adaptation capability, allowing the system to respond to changing exploration and connectivity requirements without centralized coordination. The bio-inspired flocking mechanism for supporter positioning proved particularly effective, creating adaptive relay chains that maintain network integrity while following exploration progress. This natural coordination emerges from local interactions without requiring global state information or centralized planning.

Performance analysis revealed that DRBECM achieves exploration times comparable to centralized frontier-based methods while operating in a fully distributed manner. This validates the effectiveness of local decision-making based on comprehensive environmental and network state information. The algorithm's ability to balance competing objectives through role-switching demonstrates the potential of adaptive coordination mechanisms in complex multi-robot systems.

The machine learning enhancement in DRBECM-ML addresses a critical limitation of geometric connectivity models by incorporating realistic signal propagation characteristics that account for environmental factors not captured by idealized geometric relationships. The substantial improvement in success rates demonstrates that more accurate connectivity prediction enables better exploration decisions, preventing robots from becoming isolated and ensuring mission completion reliability.

The trade-offs associated with the machine learning integration involve modest increases in exploration time and computational overhead, which proved acceptable given the significant reliability improvements. DRBECM-ML exhibits slightly higher median exploration times compared to the original DRBECM, reflecting the more conservative decision-making enabled by realistic connectivity assessment. However, this conservatism translates directly into dramatically improved mission completion rates, making it highly desirable for critical applications where mission failure has significant consequences.

6.2.2 Scalability and Robustness Characteristics

The proposed frameworks demonstrate favorable scalability characteristics due to their decentralized architecture and local decision-making processes. The RNG-based communication topology naturally limits the number of neighbors each robot must consider, with research showing that nodes in two-dimensional RNG graphs have an average of 2-3 neighbors. This property maintains computational efficiency as team size increases, preventing the quadratic complexity growth that would occur with all-to-all communication approaches.

The role-based coordination mechanism distributes the exploration workload dynamically, preventing bottlenecks that could arise from fixed role assignments or centralized task allocation. As exploration progresses and network topology evolves, robots can seamlessly transition between roles based on local conditions, ensuring that the system adapts to changing requirements without requiring global reconfiguration.

Robustness analysis reveals that the algorithms maintain effectiveness across varying network densities and environmental conditions. The dynamic role-switching mechanism provides resilience to individual robot failures, as supporters can become explorers when needed and vice versa. This adaptability ensures that the system continues operating even when robots are lost due to hardware failures, battery depletion, or environmental hazards.

The integration of multiple recovery mechanisms enhances operational reliability in challenging scenarios. Stagnation detection prevents robots from becoming trapped in local minima, while collision avoidance based on RNG relationships ensures safe operation in dense formations. The combination of these mechanisms creates a robust system that can handle the uncertainties and failures typical of real-world deployment scenarios.

Stress testing with varying network densities (5-60 robots in simulation scenarios) demonstrated that the algorithms maintain their fundamental properties across different operational scales. While performance metrics such as PDR naturally degrade with increased density due to communication contention, the relative advantages of the proposed approaches remain consistent, suggesting good scalability to larger team sizes.

6.2.3 Limitations and Considerations

Several limitations should be acknowledged in interpreting these results, as they provide important context for understanding the scope and applicability of the proposed approaches. The simulation-based evaluation, while comprehensive and systematic, cannot capture all complexities of real-world deployment scenarios. Environmental factors such as physical obstacles beyond simple geometric relationships, varying terrain that affects robot mobility, and complex signal propagation patterns influenced by building structures, vegetation, or weather conditions may impact performance differently than observed in controlled simulations.

The machine learning models, though trained on real-world data from the FIT IoT-Lab testbed, were developed using indoor measurements with controlled conditions. The signal propagation characteristics in outdoor environments, underground scenarios, or cluttered industrial settings may differ significantly from the training data, potentially affecting the accuracy of RSSI predictions and subsequent decision-making quality.

The current framework assumes homogeneous robot capabilities and does not address heterogeneous teams with varying sensing ranges, communication capabilities, processing power, or mobility constraints. Real-world deployments often involve robots with different specifications due to cost considerations, mission requirements, or hardware availability. Extending the framework to handle such heterogeneity would require significant algorithmic modifications and additional validation.

The algorithms' performance in highly dynamic environments with rapid topology changes, significant communication interference, or frequent robot failures requires further investigation. While the simulation studies included mobility and some failure scenarios, the controlled nature of these tests may not fully represent the challenges of disaster response scenarios where environmental conditions change rapidly and unpredictably.

Battery life and energy consumption were not explicitly modeled in the current evaluation framework. Real-world deployments must consider energy constraints that affect both robot movement and communication capabilities. The trade-offs between exploration efficiency and energy consumption may influence optimal role assignment strategies and exploration patterns.

Finally, the framework's dependence on accurate localization capabilities may limit its applicability in GPS-denied environments or scenarios where localization systems are degraded. While the algorithms could potentially operate with relative positioning information, this would require modifications to the geometric calculations and may affect performance characteristics.

6.3 Perspectives

The research contributions presented in this thesis open multiple avenues for future investigation and development, spanning theoretical advances, technological enhancements, and practical applications. This section outlines the most promising perspectives emerging from this work, organized around key research directions that could significantly advance the state-of-the-art in distributed multi-robot systems.

6.3.1 Algorithmic and Theoretical Extensions

Heterogeneous Multi-Robot Teams. A natural extension of the DRBECM framework involves adapting the role-based coordination mechanism to heterogeneous robot teams with varying capabilities, sensing ranges, and communication specifications. Future work could explore dynamic role assignment strategies that account for robot-specific capabilities, such as specialized sensors for different environmental conditions or varying battery capacities that influence exploration duration. The challenge lies in developing fair and efficient task allocation mechanisms that leverage individual robot strengths while maintaining overall system coherence and connectivity.

Our approach to heterogeneous integration leverages the inherent flexibility of the role-based coordination framework. The current role assignment equations (Eq. 4.1 and 5.1) can be extended to incorporate robot capability vectors $\mathbf{c}_i = [c_{sensing}, c_{comm}, c_{battery}, c_{mobility}]$, where role decisions

weight both network conditions and individual capabilities. This adaptation requires no fundamental redesign of the coordination mechanisms, the RNG topology management and flocking-inspired positioning remain unchanged while role assignment becomes capability-aware through enhanced utility functions. The distributed nature of our framework naturally accommodates varying robot specifications without requiring centralized capability management.

Multi-Objective Optimization Integration. The current framework implicitly balances exploration efficiency and connectivity maintenance through role switching. Future research could formalize this as a multi-objective optimization problem, incorporating additional objectives such as energy consumption, mission duration, risk assessment, and information quality. Evolutionary algorithms or Pareto-optimal solution sets could provide more sophisticated trade-off mechanisms, allowing mission planners to adapt system behavior to specific operational requirements.

The DRBECM framework provides an ideal foundation for multi-objective optimization since it already implicitly balances competing objectives through role switching. Formalizing these trade-offs requires extending the current decision mechanisms with explicit objective functions for exploration efficiency, connectivity reliability, energy consumption, and risk assessment. The role assignment process becomes a multi-objective decision problem where robots solve $\min \mathbf{F}(\rho_i, \mathbf{p}_i) = [f_1, f_2, \dots, f_n]$ with dynamic weight adjustment based on mission phase and environmental conditions. Each f function evaluates how well the system performs according to a specific criterion. The DRBECM-ML's RSSI prediction model seamlessly integrates into this framework, providing realistic connectivity assessments for the optimization process.

Advanced Learning Paradigms. While DRBECM-ML demonstrates the value of supervised learning for connectivity prediction, several advanced learning approaches could further enhance system capabilities. Reinforcement learning could enable robots to optimize their coordination strategies through experience, potentially discovering novel coordination patterns not anticipated by human designers. Federated learning approaches could allow robots to collaboratively improve their models while maintaining data privacy and reducing communication overhead. Transfer learning could enable rapid adaptation to new environments by leveraging knowledge gained in previous deployments.

The integration of advanced learning paradigms builds naturally upon our existing framework architecture. For reinforcement learning, the current rule-based role assignment can be transformed into policy networks $\pi_i(s_i) \rightarrow \text{explorer, supporter}$ where state representations include local network topology, frontier information, and connectivity predictions. The DRBECM-ML's lightweight Decision Tree model provides a good foundation for federated learning extensions through model averaging across robots without sharing raw sensor data. While supervised and federated learning require only algorithmic enhancements, reinforcement learning integration necessitates redesigning reward functions and action spaces while preserving the core distributed coordination principles.

Theoretical Guarantees and Formal Verification. Establishing formal theoretical guarantees for exploration completeness, connectivity maintenance, and convergence properties would significantly strengthen the algorithmic foundations. Future work could develop mathematical proofs for system behavior under various conditions, enabling predictable performance and safety certification for critical applications. Model checking

and formal verification techniques could ensure that the distributed coordination mechanisms satisfy specified safety and liveness properties.

Establishing formal guarantees for our framework requires reconstructing the mathematical formulation using formal verification techniques while preserving the algorithmic behavior. The system can be modeled as a distributed finite state machine with states $S = (role_i, pos_i, neighbors_i)$ and transitions governed by our role update equations. Connectivity guarantees can be proven using RNG graph properties, while exploration completeness follows from frontier-based progression with bounded stagnation recovery. This represents a parallel mathematical development that requires complete reformulation of all equations as temporal logic formulas, though the underlying algorithms remain unchanged.

6.3.2 Technological Enhancements

Advanced Sensor Integration and Environmental Modeling. The integration of diverse sensor modalities beyond basic ranging and communication could dramatically enhance environmental understanding and decision-making capabilities. LiDAR-based 3D mapping could improve obstacle avoidance and path planning in complex three-dimensional environments. Chemical sensors could enable detection of hazardous substances, while thermal imaging could enhance victim detection in search and rescue scenarios. The challenge involves fusing heterogeneous sensor data in real-time while maintaining the lightweight computational requirements of distributed systems.

Our framework's sensor integration approach extends the current local map representation $M_i(t)$ to multi-layered environmental models incorporating thermal, chemical, and structural data. With enhanced sensor suites and expanded robot capabilities, our model can be augmented with new parameters for task-specific operations, thermal detection thresholds for search and rescue, chemical concentration gradients for environmental monitoring, or structural integrity assessments for industrial inspection. These enhancements require adding sensor-specific parameters to the role assignment equations and extending the frontier selection criteria to incorporate multi-modal information gain metrics, where frontiers are evaluated based on their potential for different types of sensor data collection. The frontier evaluation function adapts to include sensor-specific information gain metrics while maintaining the core connectivity constraints. This enhancement requires only extending the frontier evaluation function and local map data structures, the fundamental role-based coordination, RNG topology management, and DRBECM-ML's RSSI prediction capabilities continue operating unchanged with the expanded sensor suite.

Deployable Communication Infrastructure. Beyond enhancing individual robot capabilities, future research could explore extending the DRBECM framework's connectivity maintenance strategy through deployable relay infrastructure. This emerging perspective envisions supporter robots strategically positioning autonomous relay nodes at critical network junctions, creating persistent communication bridges that remain operational as teams advance into unexplored territories. These deployable nodes could serve dual purposes: maintaining network connectivity when physical robot positioning becomes insufficient, and enabling continuous monitoring of previously explored areas while the main robot team progresses to new frontiers. This approach represents a natural evolution from purely mobile relay positioning to hybrid architectures

that combine dynamic robot coordination with strategically deployed static infrastructure. The challenge involves developing coordination frameworks that seamlessly integrate mobile robots with deployed infrastructure while preserving the distributed decision-making advantages that make the DRBECM approach inherently robust and scalable.

Edge Computing and Distributed AI. The deployment of edge computing capabilities could enable more sophisticated on-board processing while maintaining the distributed nature of the system. Distributed AI architectures could allow robots to share computational burdens for complex tasks such as simultaneous localization and mapping, object recognition, or path planning. This could significantly enhance system capabilities without requiring centralized infrastructure, making it particularly valuable for infrastructure-denied environments.

The DRBECM framework's distributed architecture aligns naturally with edge computing paradigms. Our current lightweight Decision Tree model (0.0025s inference time, 24,953 bytes memory usage) demonstrates edge-computing readiness, while the role-based structure provides natural load balancing for computational tasks. Supporter robots can extend their relay function to include computational relay with processing data from nearby explorers and sharing results without fundamental algorithmic changes. This requires only enhanced communication protocols for computational task distribution while maintaining the core coordination mechanisms and connectivity management strategies.

Adaptive Multi-Protocol Communication Systems. Integration of multiple communication technologies (WiFi, Bluetooth, LoRa, Zigbee) could enable dynamic protocol selection based on data transmission requirements, environmental conditions, and network topology. High-bandwidth WiFi connections could support intensive data sharing during map updates and sensor fusion, while low-power LoRa protocols could maintain long-range connectivity for critical status updates and coordination messages. The DRBECM-ML framework's RSSI prediction model could be extended to multi-protocol environments, enabling intelligent communication protocol switching based on predicted signal quality, data urgency, and energy constraints. This adaptive communication approach would allow robots to optimize their connectivity strategy in real-time, using high-bandwidth protocols when available and falling back to long-range, low-power alternatives when operating at the edge of the communication network.

Quantum-Enhanced Sensing and Communication. Long-term technological perspectives include the integration of quantum sensing technologies for enhanced precision in navigation and environmental monitoring. Quantum communication protocols could provide unprecedented security for sensitive mission data, while quantum computing algorithms could enable more sophisticated optimization and learning capabilities on future robotic platforms.

Quantum technology integration represents the nearly only perspective requiring fundamental redesign of our framework. While the high-level role-based coordination principles remain valid, quantum sensing and communication necessitate completely different mathematical models for localization accuracy, communication protocols, and error correction mechanisms. The conceptual framework of dynamic role assignment and distributed coordination transfers to quantum-enhanced systems, but all underlying equations, assumptions, and implementation details require quantum-compatible reformulation. This represents a complete technological paradigm shift rather than an algorithmic enhancement.

6.3.3 Application-Driven Research Directions

Domain-Specific Optimization. Different application domains present unique challenges and opportunities for system optimization. Search and rescue operations require rapid deployment and human detection capabilities, suggesting integration with computer vision and thermal sensing. Environmental monitoring applications could benefit from long-term autonomy and adaptive sampling strategies. Industrial inspection scenarios might emphasize precision navigation and defect detection. Each domain presents opportunities for specialized algorithm variants that leverage domain knowledge while maintaining the core coordination principles.

The parametric flexibility of our framework makes domain-specific optimization the most straightforward extension path. Different application domains require only adjusting threshold values ($\alpha, \beta_1, \beta_2, T_{strong}, T_{weak}$), sensor integration modules, and objective function weights while maintaining the core algorithms unchanged. Search and rescue scenarios emphasize rapid exploration parameters and human detection sensors; environmental monitoring prioritizes energy efficiency and long-term deployment configurations; industrial inspection focuses on precision navigation and comprehensive coverage parameters. The role-based coordination remains domain-agnostic, requiring no fundamental modifications to accommodate diverse operational requirements.

Human-Robot Collaboration Frameworks. Future developments should address seamless integration with human operators and decision-makers. This includes developing intuitive interfaces for mission specification and monitoring, adaptive automation that adjusts to human preferences and expertise levels, and collaborative decision-making frameworks that combine human insight with algorithmic efficiency. The challenge involves maintaining system autonomy while providing appropriate human oversight and intervention capabilities.

Our distributed coordination framework could naturally support human oversight through its decentralized communication structure. However, determining the appropriate level of human intervention presents significant challenges in balancing autonomy with controllability. For example, allowing humans to directly override individual robot role assignments could disrupt the distributed coordination logic and create network fragmentation, while restricting input to high-level mission parameters may not provide sufficient control in critical situations. Specific challenges include: deciding whether humans can force an explorer to become a supporter (potentially compromising exploration efficiency), allowing manual frontier selection that might violate connectivity constraints, or permitting real-time parameter adjustments ($\alpha, T_{strong}, T_{weak}$) that could destabilize the coordination mechanism. This human-robot integration leverages the framework's inherent flexibility while maintaining system stability through protected autonomous decision-making layers.

Multi-Scale Coordination. Scaling the coordination mechanisms to very large robot teams (hundreds to thousands of robots) presents both theoretical and practical challenges. Hierarchical coordination structures could enable scalable operation while maintaining the advantages of distributed decision-making. This might involve formation of dynamic clusters or sub-teams that coordinate locally while maintaining global awareness through representatives or specialized coordinator roles.

Scaling our framework to thousands of robots requires hierarchical extensions rather than fundamental redesign. We envision fractal coordination

structures where current DRBECM algorithms operate within clusters (10-50 robots) while cluster representatives coordinate at higher levels using similar role-based principles. Each cluster maintains its own base station or leader, creating a multi-level hierarchy that preserves the distributed coordination advantages at each scale. This approach requires adding inter-cluster coordination protocols and cluster formation mechanisms while maintaining the fundamental role-based coordination and connectivity maintenance principles that naturally scale to hierarchical structures.

6.3.4 Ethical and Deployment Considerations for Emergency Response Systems

The proposed DRBECM frameworks address critical challenges in emergency response robotics, where ethical considerations are intrinsically linked to the technical design choices that affect human safety and mission effectiveness. The decentralized coordination mechanisms developed in this thesis raise important considerations that are directly relevant to responsible deployment in disaster scenarios.

Safety-Critical Decision Making in Distributed Systems. The distributed nature of DRBECM requires individual robots to make autonomous decisions about exploration paths, role assignments, and connectivity maintenance without centralized oversight. In emergency response scenarios, these decisions directly impact human safety outcomes. For instance, when robots must choose between maintaining communication connectivity and pursuing potentially life-saving exploration targets, the embedded decision-making logic must prioritize human welfare. Future implementations should incorporate explicit safety constraints that prevent robots from compromising communication networks that are critical for coordinating human rescue operations.

Reliability and Trust in Life-Critical Applications. The machine learning enhancements in DRBECM-ML, while improving mission success rates from 83.6% to 99.3%, introduce questions about algorithmic transparency and reliability in life-critical scenarios. The Decision Tree model, chosen for its balance of accuracy and interpretability, provides some degree of explainability compared to black-box alternatives. However, deployment in emergency response requires additional validation mechanisms to ensure that connectivity predictions remain accurate across diverse real-world environments and degraded conditions typical of disaster zones.

Human-Robot Coordination and Operational Integration. The NEPHELE project context highlights the need for seamless integration with human first responders and existing emergency management systems. The distributed exploration algorithms must be designed to complement rather than replace human decision-making, providing reliable situational awareness while respecting the critical role of human expertise in emergency response. This includes developing interfaces that allow human operators to understand and trust the autonomous exploration process, intervene when necessary, and integrate robot-collected information into broader response coordination efforts.

Privacy and Data Security in Disaster Response. Multi-robot exploration systems in emergency scenarios collect sensitive information about affected individuals, infrastructure damage, and operational capabilities. The decentralized architecture provides inherent resilience against single

points of failure but requires careful consideration of data handling and sharing protocols. Future implementations must incorporate appropriate encryption, access controls, and data retention policies that protect victim privacy while enabling effective response coordination.

Special Terms

A

Ad-hoc Ad-hoc Network Mode. 6, 16

ANN Artificial Neural Network. 33

ARM Advanced RISC Machine. 71

C

CME Coordinated Multi-Robot Exploration. 15, 16

D

DDQN Double Deep Q-Network. 31

DRBECM Dynamic Role-Based Exploration with Connectivity Maintenance. vi, 9–11, 13, 14, 48, 51, 57, 73

DRBECM-ML Multi-Robot Exploration via Flocking Coordination and Machine Learning-Driven Connectivity Assessment. vi, 9, 11, 13, 69, 73

DRL Deep Reinforcement Learning. 31

F

FMH-WOA Frequency-Modified Whale Optimization Algorithm. 15, 16

H

HCETIIC Hybrid Cheetah Exploration Technique with Intelligent Initial Configuration. 18, 65

I

IoT Internet of Things. 5

IQR Interquartile Range. 72

K

KNN K-Nearest Neighbors. vi, 39

L

LiDAR Light Detection and Ranging. 5

M

MAE Mean Absolute Error. 73

ML Machine Learning. 16

MRS Multi-Robot Systems. vi, 2, 39

N

NS-3 Network Simulator 3. vi

P

PDR Packet Delivery Ratio. 10, 39, 45

PSO Particle Swarm Optimization. 16, 30

Q

QoS Quality-of-Service. 39, 44

R

RL Reinforcement Learning. 14, 16, 33

RMSE Root Mean Square Error. 73

RNG Relative Neighborhood Graph. vi, 10, 39, 40, 55

RSSI Received Signal Strength Indicator. vi, 4, 11, 32, 69, 71, 75

S

SSA Salp Swarm Algorithm. 16, 19

U

UDP User Datagram Protocol. 43

W

Wi-Fi Wireless Fidelity. 6

X

XGBoost Extreme Gradient Boosting. vi

Bibliography

Here are the references in citation order.

- [1] Keiji Nagatani et al. 'Multi-robot exploration for search and rescue missions: A report of map building in RoboCupRescue 2009'. In: *IEEE International Workshop on Safety, Security & Rescue Robotics (SSRR)*. 2009. doi: [10.1109/SSRR.2009.5424148](https://doi.org/10.1109/SSRR.2009.5424148) (cited on pages 1, 9).
- [2] Martijn N. Rooker and Andreas Birk. 'Multi-robot exploration under the constraints of wireless networking'. In: *Control Engineering Practice* (2007). doi: <https://doi.org/10.1016/j.conengprac.2006.08.007> (cited on pages 1, 3).
- [3] Yang Shen et al. 'Robots Under COVID-19 Pandemic: A Comprehensive Survey'. In: *IEEE Access* (2021). doi: [10.1109/ACCESS.2020.3045792](https://doi.org/10.1109/ACCESS.2020.3045792) (cited on pages 1, 2).
- [4] Keiji Nagatani et al. 'Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots'. In: *Journal of Field Robotics* (2013). doi: <https://doi.org/10.1002/rob.21439> (cited on page 2).
- [5] Janardan Verma and Virender Ranga. 'Multi-Robot Coordination Analysis, Taxonomy, Challenges and Future Scope'. In: *Journal of Intelligent & Robotic Systems* (2021). doi: [10.1007/s10846-021-01378-2](https://doi.org/10.1007/s10846-021-01378-2) (cited on pages 2, 4, 39).
- [6] Manuele Brambilla et al. 'Swarm robotics: a review from the swarm engineering perspective'. In: *Swarm Intelligence* (2013) (cited on pages 2, 4, 13).
- [7] Francesco Amigoni, Jacopo Banfi, and Nicola Basilico. 'Multirobot Exploration of Communication-Restricted Environments: A Survey'. In: *IEEE Intelligent Systems* (2017). doi: [10.1109/MIS.2017.4531226](https://doi.org/10.1109/MIS.2017.4531226) (cited on pages 2, 3, 7, 39).
- [8] Jennifer Gielis, Ajay Shankar, and Amanda Prorok. *A Critical Review of Communications in Multi-Robot Systems*. 2022. URL: <https://arxiv.org/abs/2206.09484> (cited on pages 2–4, 7, 8).
- [9] Geoffrey A. Hollinger and Sanjiv Singh. 'Multirobot Coordination With Periodic Connectivity: Theory and Experiments'. In: *IEEE Transactions on Robotics* (2012). doi: [10.1109/TRO.2012.2190178](https://doi.org/10.1109/TRO.2012.2190178) (cited on pages 2, 3).
- [10] Brian Yamauchi. 'A frontier-based approach for autonomous exploration'. In: *IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97*. 1997. doi: [10.1109/CIRA.1997.613851](https://doi.org/10.1109/CIRA.1997.613851) (cited on pages 2, 17, 18, 36).
- [11] Yuanteng Pei, Matt W. Mutka, and Ning Xi. 'Connectivity and bandwidth-aware real-time exploration in mobile robot networks'. In: *Wireless Communications and Mobile Computing* (2013). doi: [10.1002/WCM.1145](https://doi.org/10.1002/WCM.1145) (cited on pages 3, 4).
- [12] Daniel Bonilla Licea, Mounir Ghogho, and Martin Saska. 'When Robotics Meets Wireless Communications: An Introductory Tutorial'. In: *Proceedings of the IEEE* (2024). doi: [10.1109/jproc.2024.3380373](https://doi.org/10.1109/jproc.2024.3380373) (cited on page 3).
- [13] Maira Saboia et al. 'ACHORD: Communication-Aware Multi-Robot Coordination With Intermittent Connectivity'. In: *IEEE Robotics and Automation Letters* (2022). doi: [10.1109/LRA.2022.3193240](https://doi.org/10.1109/LRA.2022.3193240) (cited on pages 3, 17, 18, 30).
- [14] Trevor Halsted et al. *A Survey of Distributed Optimization Methods for Multi-Robot Systems*. 2021. URL: <https://arxiv.org/abs/2103.12840> (cited on page 3).
- [15] Marco Tranzatto et al. *Team CERBERUS Wins the DARPA Subterranean Challenge: Technical Overview and Lessons Learned*. 2022. URL: <https://arxiv.org/abs/2207.04914> (cited on pages 3, 4).
- [16] Amanda Prorok et al. *Beyond Robustness: A Taxonomy of Approaches towards Resilient Multi-Robot Systems*. 2021. URL: <https://arxiv.org/abs/2109.12343> (cited on pages 3, 4).
- [17] Rabiul Hasan Kabir and Kooktae Lee. *Efficient Multi-Robot Exploration with Energy Constraint based on Optimal Transport Theory*. 2020. URL: <https://arxiv.org/abs/2009.00862> (cited on page 3).

- [18] G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. 'A comprehensive taxonomy for multi-robot task allocation'. In: *The International Journal of Robotics Research* (2013). doi: [10.1177/0278364913496484](https://doi.org/10.1177/0278364913496484) (cited on pages 3, 4).
- [19] Julian de Hoog, Stephen Cameron, and Arnoud Visser. 'Role-Based Autonomous Multi-robot Exploration'. In: *Computation World: Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns*. 2009. doi: [10.1109/ComputationWorld.2009.14](https://doi.org/10.1109/ComputationWorld.2009.14) (cited on page 4).
- [20] Yousef Emam et al. 'Adaptive Task Allocation for Heterogeneous Multi-Robot Teams with Evolving and Unknown Robot Capabilities'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2020. doi: [10.1109/ICRA40945.2020.9197283](https://doi.org/10.1109/ICRA40945.2020.9197283) (cited on page 4).
- [21] Jonathan Fink and Vijay Kumar. 'Online methods for radio signal mapping with mobile robots'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2010. doi: [10.1109/ROBOT.2010.5509574](https://doi.org/10.1109/ROBOT.2010.5509574) (cited on page 4).
- [22] Alberto Quattrini Li et al. 'Multi-robot online sensing strategies for the construction of communication maps'. In: *Autonomous Robots* (2020). doi: [10.1007/s10514-019-09862-3](https://doi.org/10.1007/s10514-019-09862-3) (cited on page 4).
- [23] Renato Miyagusuku, Atsushi Yamashita, and Hajime Asama. 'Precise and accurate wireless signal strength mappings using Gaussian processes and path loss models'. In: *Robotics and Autonomous Systems* (2018). doi: <https://doi.org/10.1016/j.robot.2018.02.011> (cited on page 4).
- [24] *Nephele Use case1: Emergency/Disaster Recovery*. <https://nephele-project.eu/use-case-1-emergencydisaster-recovery> (cited on pages 6, 7).
- [25] W. Burgard et al. 'Coordinated multi-robot exploration'. In: *IEEE Transactions on Robotics* (2005). doi: [10.1109/TR0.2004.839232](https://doi.org/10.1109/TR0.2004.839232) (cited on pages 7, 18).
- [26] Robert Zlot et al. 'Multi-robot exploration controlled by a market economy'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2002. doi: [10.1109/ROBOT.2002.1013690](https://doi.org/10.1109/ROBOT.2002.1013690) (cited on pages 7, 8, 17, 18).
- [27] Brian Gerkey and Maja Mataric. 'A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems'. In: *The International Journal of Robotics Research* (2004). doi: [10.1177/0278364904045564](https://doi.org/10.1177/0278364904045564) (cited on pages 7, 8).
- [28] Han-Lim Choi, Luc Brunet, and Jonathan P. How. 'Consensus-Based Decentralized Auctions for Robust Task Allocation'. In: *IEEE Transactions on Robotics* (2009). doi: [10.1109/TR0.2009.2022423](https://doi.org/10.1109/TR0.2009.2022423) (cited on pages 7, 8).
- [29] Simon Hunt et al. 'A Consensus-Based Grouping Algorithm for Multi-agent Cooperative Task Allocation with Complex Requirements'. In: *Cognitive computation* (2014). doi: [10.1007/s12559-014-9265-0](https://doi.org/10.1007/s12559-014-9265-0) (cited on page 7).
- [30] Jacopo Banfi et al. 'Strategies for coordinated multirobot exploration with recurrent connectivity constraints'. In: *Autonomous Robots* (2018). doi: [10.1007/s10514-017-9652-y](https://doi.org/10.1007/s10514-017-9652-y) (cited on pages 7, 8, 18).
- [31] Ethan Stump, Ali Jadbabaie, and Vijay Kumar. 'Connectivity management in mobile robot teams'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2008. doi: [10.1109/ROBOT.2008.4543418](https://doi.org/10.1109/ROBOT.2008.4543418) (cited on pages 7, 8).
- [32] Nathan Michael et al. 'Maintaining Connectivity in Mobile Robot Networks'. In: *Experimental Robotics*. 2008. doi: [10.1007/978-3-642-00196-3_14](https://doi.org/10.1007/978-3-642-00196-3_14) (cited on pages 7, 8).
- [33] Alireza Ghaffarkhah and Yasamin Mostofi. 'Communication-Aware Motion Planning in Mobile Networks'. In: *IEEE Transactions on Automatic Control* (2011). doi: [10.1109/TAC.2011.2164033](https://doi.org/10.1109/TAC.2011.2164033) (cited on pages 7, 8).
- [34] Razanne Abu-Aisheh et al. 'CARA: Connectivity-Aware Relay Algorithm for Multi-Robot Expeditions'. In: *Sensors* (2022). doi: [10.3390/s22239042](https://doi.org/10.3390/s22239042) (cited on page 8).
- [35] Alysson Ribeiro da Silva et al. *Communication-Constrained Multi-Robot Exploration with Intermittent Rendezvous*. 2024. URL: <https://arxiv.org/abs/2309.13494> (cited on page 8).
- [36] Byung-Cheol Min et al. *A Directional Antenna based Leader-Follower Relay System for End-to-End Robot Communications*. 2017. URL: <https://arxiv.org/abs/1711.08007> (cited on page 8).
- [37] Yuan Yan and Yasamin Mostofi. 'To Go or Not to Go: On Energy-Aware and Communication-Aware Robotic Operation'. In: *IEEE Transactions on Control of Network Systems* (2014). doi: [10.1109/TCNS.2014.2337971](https://doi.org/10.1109/TCNS.2014.2337971) (cited on page 8).

- [38] Magnus Lindhé, Karl Johansson, and Antonio Bicchi. 'An experimental study of exploiting multipath fading for robot communications'. In: *Proceedings of Robotics: Science and Systems*. 2007. doi: [10.15607/RSS.2007.III.037](#) (cited on page 8).
- [39] Arjun Muralidharan and Yasamin Mostofi. 'Communication-Aware Robotics: Exploiting Motion for Communication'. In: *Annual Review of Control, Robotics, and Autonomous Systems* (2021). doi: [10.1146/annurev-control-071420-080708](#) (cited on page 8).
- [40] Phani Krishna Penumarthi et al. 'Multirobot exploration for building communication maps with prior from communication models'. In: *International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. 2017. doi: [10.1109/MRS.2017.8250936](#) (cited on page 8).
- [41] Michael Otte, Michael Kuhlman, and Donald Sofge. 'Multi-robot task allocation with auctions in harsh communication environments'. In: *International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. 2017. doi: [10.1109/MRS.2017.8250928](#) (cited on page 8).
- [42] M. Bernardine Dias et al. 'Market-Based Multirobot Coordination: A Survey and Analysis'. In: *Proceedings of the IEEE* (2006) (cited on pages 8, 13, 14, 17, 18).
- [43] Jacob Crandall et al. 'Validating Human-Robot Interaction Schemes in Multitasking Environments'. In: *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* (2005). doi: [10.1109/TSMCA.2005.850587](#) (cited on page 9).
- [44] Hazem Chaabi and Nathalie Mitton. 'Multi-Robot Exploration via Flocking Coordination and Machine Learning-Driven Connectivity Assessment'. In: *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*. 2025 (cited on page 12).
- [45] Hazem Chaabi and Nathalie Mitton. 'Distributed Multi-Robot Exploration Approach with Connectivity Maintenance'. In: *International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. 2025. doi: [10.1109/DCOSS-IoT65416.2025.00020](#) (cited on page 12).
- [46] Gregory Dudek et al. 'A taxonomy for multi-agent robotics'. In: *Autonomous Robots* (1996). doi: [10.1007/BF00240651](#) (cited on page 13).
- [47] Yann LeCun, Y. Bengio, and Geoffrey Hinton. 'Deep Learning'. In: *Nature* (2015). doi: [10.1038/nature14539](#) (cited on page 13).
- [48] Lynne E. Parker. 'ALLIANCE: an architecture for fault tolerant multirobot cooperation'. In: *IEEE Transactions on Robotics and Automation* (1998). doi: [10.1109/70.681242](#) (cited on page 14).
- [49] Mehmet B. Bezcioglu, Barry Lennox, and Farshad Arvin. 'Self-Organised Swarm Flocking with Deep Reinforcement Learning'. In: *7th International Conference on Automation, Robotics and Applications (ICARA)*. 2021. doi: [10.1109/ICARA51699.2021.9376509](#) (cited on pages 14, 15, 19, 23, 33, 34, 36).
- [50] Cyrill Stachniss and Wolfram Burgard. 'Exploring unknown environments with mobile robots using coverage maps'. In: *International Joint Conference on Artificial Intelligence (IJCAI)*. Morgan Kaufmann Publishers Inc., 2003 (cited on page 14).
- [51] Di Deng et al. *Frontier-based Automatic-differentiable Information Gain Measure for Robotic Exploration of Unknown 3D Environments*. 2020. URL: <https://arxiv.org/abs/2011.05288> (cited on page 14).
- [52] Junyan Hu et al. 'Voronoi-Based Multi-Robot Autonomous Exploration in Unknown Environments via Deep Reinforcement Learning'. In: *IEEE Transactions on Vehicular Technology* (2020). doi: [10.1109/TVT.2020.3034800](#) (cited on pages 14, 15, 21, 35, 36).
- [53] Chao Liu et al. 'Leader-following flocking for unmanned aerial vehicle swarm with distributed topology control'. In: *Science China Information Sciences* (2020). doi: [10.1007/s11432-019-2763-5](#) (cited on pages 14, 22, 23).
- [54] Jonghoek Kim. 'Leader-Based Flocking of Multiple Swarm Robots in Underwater Environments'. In: *Sensors* (2023). doi: [10.3390/s23115305](#) (cited on pages 14, 15, 35).
- [55] Ankit Shaw. *Autonomous Multi-Robot Exploration Strategies for 3D Environments with Fire Detection Capabilities*. 2024. URL: <https://arxiv.org/abs/2411.15953> (cited on pages 15, 17, 19, 26, 36).
- [56] Ali El Romeh and Seyedali Mirjalili. 'Theoretical Framework and Practical Considerations for Achieving Superior Multi-Robot Exploration: Hybrid Cheetah Optimization with Intelligent Initial Configurations'. In: *Mathematics* (2023). doi: [10.3390/math11204239](#) (cited on pages 15–19, 27, 36, 62, 76).

- [57] Rakesh Shrestha et al. 'Learned Map Prediction for Enhanced Mobile Robot Exploration'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2019. doi: [10.1109/ICRA.2019.8793769](#) (cited on pages 15, 31–33, 37).
- [58] Jelle Groenendaal, Ira Helsloot, and Astrid Scholtens. 'A Critical Examination of the Assumptions Regarding Centralized Coordination in Large-Scale Emergency Situations'. In: *Journal of Homeland Security and Emergency Management* (2013). doi: [doi:10.1515/jhsem-2012-0053](#) (cited on pages 15, 36).
- [59] Fabian Schilling et al. 'Learning Vision-Based Flight in Drone Swarms by Imitation'. In: *IEEE Robotics and Automation Letters* (2019). doi: [10.1109/LRA.2019.2935377](#) (cited on pages 15–18, 36).
- [60] Cinara Ghedini, Carlos H.C. Ribeiro, and Lorenzo Sabattini. 'Toward efficient adaptive ad-hoc multi-robot network topologies'. In: *Ad Hoc Networks* (2018). doi: [10.1016/j.adhoc.2018.03.012](#) (cited on pages 15–17, 26, 36).
- [61] Hiroto Inahara and Naoki Motoi. 'Research on Search Algorithm Using Particle Swarm Optimization with Virtual Pheromone for Swarm Robots'. In: *IEEJ Journal of Industry Applications* (2022). doi: [10.1541/ieejjia.21000994](#) (cited on page 16).
- [62] Ying TAN and Zhong-yang ZHENG. 'Research Advance in Swarm Robotics'. In: *Defence Technology* (2013). doi: [10.1016/j.dt.2013.03.001](#) (cited on pages 16–18, 28–30, 37).
- [63] Guangdong Lu et al. 'Multi-Robot Collaborative Exploration on Communication-Constrained Environments'. In: *IEEE Access* (2025). doi: [10.1109/ACCESS.2025.3539518](#) (cited on pages 16, 17, 25, 27).
- [64] Ruihua Han et al. 'Reinforcement Learned Distributed Multi-Robot Navigation With Reciprocal Velocity Obstacle Shaped Rewards'. In: *IEEE Robotics and Automation Letters* (2022). doi: [10.1109/LRA.2022.3161699](#) (cited on pages 16–19, 36).
- [65] Mohammad Bakhshipour, M. Jabbari Ghadi, and F. Namdari. 'Swarm Robotics Search and Rescue; a Novel Artificial Intelligence-Inspired Optimization Approach'. In: *Applied Soft Computing* (2017). doi: [10.1016/j.asoc.2017.02.028](#) (cited on page 17).
- [66] Yuhong Cao et al. 'Deep Reinforcement Learning-Based Large-Scale Robot Exploration'. In: *IEEE Robotics and Automation Letters* (2024). doi: [10.1109/LRA.2024.3379804](#) (cited on pages 17, 18, 35–37).
- [67] Chendi Lin, Wenhao Luo, and Katia Sycara. 'Online Connectivity-aware Dynamic Deployment for Heterogeneous Multi-Robot Systems'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2021. doi: [10.1109/ICRA48506.2021.9561748](#) (cited on pages 17, 18).
- [68] Yupeng Yang et al. 'Integrating Online Learning and Connectivity Maintenance for Communication-Aware Multi-Robot Coordination'. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2024. doi: [10.1109/IROS58592.2024.10802189](#) (cited on pages 17, 18, 27).
- [69] Feng Xia et al. 'Random Walks: A Review of Algorithms and Applications'. In: *IEEE Transactions on Emerging Topics in Computational Intelligence* (2020). doi: [10.1109/tetci.2019.2952908](#) (cited on pages 18, 62, 76).
- [70] Brian Yamauchi. 'Frontier-based exploration using multiple robots'. In: *International Conference on Autonomous Agents*. 1998. doi: [10.1145/280765.280773](#) (cited on page 18).
- [71] Reza Olfati-Saber. 'Flocking for multi-agent dynamic systems: algorithms and theory'. In: *IEEE Transactions on Automatic Control* (2006). doi: [10.1109/TAC.2005.864190](#) (cited on page 18).
- [72] Haibin Duan, Mengzhen Huo, and Yanming Fan. 'From animal collective behaviors to swarm robotic cooperation'. In: *National Science Review* (2023). doi: [10.1093/nsr/nwad040](#) (cited on pages 18, 21, 22).
- [73] Haoran Li, Qichao Zhang, and Dongbin Zhao. 'Deep Reinforcement Learning-Based Automatic Exploration for Navigation in Unknown Environment'. In: *IEEE Transactions on Neural Networks and Learning Systems* (2020). doi: [10.1109/TNNLS.2019.2927869](#) (cited on pages 18, 32, 37).
- [74] Khalil Mohamed, Ayman El Shenawy, and Hany Harb. 'A Hybrid Decentralized Coordinated Approach for Multi-Robot Exploration Task'. In: *The Computer Journal* (2018) (cited on pages 18, 20, 21, 27, 35, 36).
- [75] Miguel Julia et al. 'A hybrid solution to the multi-robot integrated exploration problem'. In: *Engineering Applications of Artificial Intelligence* (2010). doi: [10.1016/j.engappai.2009.12.005](#) (cited on pages 18, 35, 36).

- [76] Kamalova Albina and Suk Gyu Lee. 'Hybrid Stochastic Exploration Using Grey Wolf Optimizer and Coordinated Multi-Robot Exploration Algorithms'. In: *IEEE Access* (2019). doi: [10.1109/ACCESS.2019.2894524](https://doi.org/10.1109/ACCESS.2019.2894524) (cited on page 18).
- [77] Faiza Gul et al. 'Novel Implementation of Multi-Robot Space Exploration Utilizing Coordinated Multi-Robot Exploration and Frequency Modified Whale Optimization Algorithm'. In: *IEEE Access* (2021). doi: [10.1109/ACCESS.2021.3055852](https://doi.org/10.1109/ACCESS.2021.3055852) (cited on page 18).
- [78] Aryo Jamshidpey et al. 'Centralization vs. decentralization in multi-robot sweep coverage with ground robots and UAVs'. In: *IEEE Robotics and Automation Letters* (2024). doi: [10.1109/LRA.2022.3224667](https://doi.org/10.1109/LRA.2022.3224667) (cited on page 19).
- [79] Laxmi Poudel et al. 'Decentralized and Centralized Planning for Multi-Robot Additive Manufacturing'. In: *Journal of Mechanical Design* (2022). doi: [10.1115/1.4055735](https://doi.org/10.1115/1.4055735) (cited on pages 20, 25, 34, 36).
- [80] Levent Bayındır. 'A Review of Swarm Robotics Tasks'. In: *Neurocomputing* (2015). doi: [10.1016/j.neucom.2015.05.116](https://doi.org/10.1016/j.neucom.2015.05.116) (cited on pages 21, 22, 24, 25, 29, 30).
- [81] Karlo Griparic et al. 'Consensus-Based Distributed Connectivity Control in Multi-Agent Systems'. In: *IEEE Transactions on Network Science and Engineering* (2022). doi: [10.1109/TNSE.2021.3139045](https://doi.org/10.1109/TNSE.2021.3139045) (cited on pages 21, 24, 25, 36).
- [82] Andrea Gasparri, Lorenzo Sabattini, and Giovanni Ulivi. 'Bounded Control Law for Global Connectivity Maintenance in Cooperative Multirobot Systems'. In: *IEEE Transactions on Robotics* (2017). doi: [10.1109/TR0.2017.2664883](https://doi.org/10.1109/TR0.2017.2664883) (cited on pages 21, 33, 34).
- [83] Jimmy Chiun et al. *MARVEL: Multi-Agent Reinforcement Learning for constrained field-of-View multi-robot Exploration in Large-scale environments*. 2025. URL: <https://arxiv.org/abs/2502.20217> (cited on pages 21, 32).
- [84] Ce Guo et al. 'Imitation Learning with Graph Neural Networks for Improving Swarm Robustness under Restricted Communications'. In: *Applied Sciences* (2021). doi: [10.3390/app11199055](https://doi.org/10.3390/app11199055) (cited on pages 22, 23, 25, 28, 29, 33, 34, 37).
- [85] Cinara Ghedini, Carlos Ribeiro, and Lorenzo Sabattini. 'Toward fault-tolerant multi-robot networks'. In: *Networks* (2017). doi: <https://doi.org/10.1002/net.21784> (cited on pages 22, 23, 26).
- [86] Minghao Li et al. 'Decentralized Global Connectivity Maintenance for Multi-Robot Navigation: A Reinforcement Learning Approach'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2022. doi: [10.1109/ICRA46639.2022.9812163](https://doi.org/10.1109/ICRA46639.2022.9812163) (cited on pages 23, 25).
- [87] Guangyao Shi et al. 'Communication-Aware Multi-robot Coordination with Submodular Maximization'. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2021. doi: [10.1109/ICRA48506.2021.9561041](https://doi.org/10.1109/ICRA48506.2021.9561041) (cited on pages 24, 25, 34).
- [88] Yang Ou et al. 'Autonomous Navigation by Mobile Robot with Sensor Fusion Based on Deep Reinforcement Learning'. In: *Sensors* (2024). doi: [10.3390/s24123895](https://doi.org/10.3390/s24123895) (cited on page 25).
- [89] Lorenzo Sabattini, Nikhil Chopra, and Cristian Secchi. 'On decentralized connectivity maintenance for mobile robotic systems'. In: *50th IEEE Conference on Decision and Control and European Control Conference*. 2011. doi: [10.1109/CDC.2011.6161067](https://doi.org/10.1109/CDC.2011.6161067) (cited on page 26).
- [90] Sarat Chandra Nagavarapu, Leena Vachhani, and Arpita Sinha. 'Multi-Robot Graph Exploration and Map Building with Collision Avoidance: A Decentralized Approach'. In: *Journal of Intelligent & Robotic Systems* (2016). doi: [10.1007/s10846-015-0309-9](https://doi.org/10.1007/s10846-015-0309-9) (cited on page 26).
- [91] Miguel Duarte et al. 'Evolution of Collective Behaviors for a Real Swarm of Aquatic Surface Robots'. In: *PLOS ONE* (2016). doi: [10.1371/journal.pone.0151834](https://doi.org/10.1371/journal.pone.0151834) (cited on page 26).
- [92] Ruofei Bai et al. *Realm: Real-Time Line-of-Sight Maintenance in Multi-Robot Navigation with Unknown Obstacles*. 2025. URL: <https://arxiv.org/abs/2502.15162> (cited on page 27).
- [93] Vu Phi Tran et al. 'Frontier-led swarming: Robust multi-robot coverage of unknown environments'. In: *Swarm and Evolutionary Computation* (2022). doi: <https://doi.org/10.1016/j.swevo.2022.101171> (cited on page 27).
- [94] Marco Aggravi, Claudio Pacchierotti, and Paolo Robuffo Giordano. 'Connectivity-Maintenance Teleoperation of a UAV Fleet With Wearable Haptic Feedback'. In: *IEEE Transactions on Automation Science and Engineering* (2021). doi: [10.1109/TASE.2020.3000060](https://doi.org/10.1109/TASE.2020.3000060) (cited on pages 28, 29, 35, 37).

- [95] Alfeo Antonio Luca et al. 'Design and simulation of the emergent behavior of small drones swarming for distributed target localization'. In: *Journal of Computational Science* (2018). doi: [10.1016/j.jocs.2018.09.014](https://doi.org/10.1016/j.jocs.2018.09.014) (cited on pages 29, 30).
- [96] Hai-Tao Zhang et al. 'Collective behavior coordination with predictive mechanisms'. In: *IEEE Circuits and Systems Magazine* (2008). doi: [10.1109/MCAS.2008.928446](https://doi.org/10.1109/MCAS.2008.928446) (cited on pages 29, 30).
- [97] Paolo Robuffo Giordano et al. 'A passivity-based decentralized strategy for generalized connectivity maintenance'. In: *The International Journal of Robotics Research* (2013). doi: [10.1177/0278364912469671](https://doi.org/10.1177/0278364912469671) (cited on pages 29, 30).
- [98] Logan E. Beaver and Andreas A. Malikopoulos. 'An overview on optimal flocking'. In: *Annual Reviews in Control* (2021). doi: <https://doi.org/10.1016/j.arcontrol.2021.03.004> (cited on page 30).
- [99] Victor R. F. Miranda et al. 'Generalization in Deep Reinforcement Learning for Robotic Navigation by Reward Shaping'. In: *IEEE Transactions on Industrial Electronics* (2024). doi: [10.1109/TIE.2023.3290244](https://doi.org/10.1109/TIE.2023.3290244) (cited on pages 31, 32, 37).
- [100] Dohyun Jang et al. 'Multi-Robot Active Sensing and Environmental Model Learning With Distributed Gaussian Process'. In: *IEEE Robotics and Automation Letters* (2020). doi: [10.1109/LRA.2020.3010456](https://doi.org/10.1109/LRA.2020.3010456) (cited on page 32).
- [101] Dayong Xu et al. 'Improving exploration efficiency of deep reinforcement learning through samples produced by generative model'. In: *Expert Systems with Applications* (2021). doi: <https://doi.org/10.1016/j.eswa.2021.115680> (cited on page 32).
- [102] Ming-Hsu Tsai et al. 'Applying a Neural Network to Predict Surface Roughness and Machining Accuracy in the Milling of SUS304'. In: *Electronics* (2023). doi: [10.3390/electronics12040981](https://doi.org/10.3390/electronics12040981) (cited on pages 33, 34).
- [103] Hyondong Oh et al. 'Bio-inspired self-organising multi-robot pattern formation: A review'. In: *Robotics and Autonomous Systems* (2017). doi: <https://doi.org/10.1016/j.robot.2016.12.006> (cited on page 34).
- [104] Reinis Cimurs, Il Hong Suh, and Jin Han Lee. 'Goal-Driven Autonomous Exploration Through Deep Reinforcement Learning'. In: *IEEE Robotics and Automation Letters* (2022). doi: [10.1109/LRA.2021.3133591](https://doi.org/10.1109/LRA.2021.3133591) (cited on pages 35, 37).
- [105] Akshay Shetty, Timmy Hussain, and Grace Gao. 'Decentralized Connectivity Maintenance for Multi-robot Systems Under Motion and Sensing Uncertainties'. In: *NAVIGATION: Journal of the Institute of Navigation* (2023) (cited on page 35).
- [106] Mingzhe Chen et al. 'Communication-efficient federated learning'. In: *Proceedings of the National Academy of Sciences* (2021). doi: [10.1073/pnas.2024789118](https://doi.org/10.1073/pnas.2024789118) (cited on page 36).
- [107] Ruixing Zong et al. 'Fedcs: Efficient communication scheduling in decentralized federated learning'. In: *Information Fusion* (2024). doi: <https://doi.org/10.1016/j.inffus.2023.102028> (cited on page 36).
- [108] Zhi Yan, Nicolas Jouandeau, and Arab Ali Cherif. 'A Survey and Analysis of Multi-Robot Coordination'. In: *International Journal of Advanced Robotic Systems* (2013). doi: [10.5772/57313](https://doi.org/10.5772/57313) (cited on page 39).
- [109] Mohammed Abdulhakim Al-Absi et al. 'Obstacles effects on signal attenuation in line of sight for different environments in V2V communication'. In: *20th International Conference on Advanced Communication Technology (ICACT)*. 2018. doi: [10.23919/ICACT.2018.8323630](https://doi.org/10.23919/ICACT.2018.8323630) (cited on page 39).
- [110] Michael M. Zavlanos and George J. Pappas. 'Distributed Connectivity Control of Mobile Networks'. In: *IEEE Transactions on Robotics* (2008). doi: [10.1109/TR0.2008.2006233](https://doi.org/10.1109/TR0.2008.2006233) (cited on page 39).
- [111] Godfried T. Toussaint. 'The relative neighbourhood graph of a finite planar set'. In: *Pattern Recognition* (1980). doi: [https://doi.org/10.1016/0031-3203\(80\)90066-7](https://doi.org/10.1016/0031-3203(80)90066-7) (cited on pages 40, 42).
- [112] J.W. Jaromczyk and G.T. Toussaint. 'Relative neighborhood graphs and their relatives'. In: *Proceedings of the IEEE* (1992). doi: [10.1109/5.163414](https://doi.org/10.1109/5.163414) (cited on page 40).
- [113] Cole Foster, Berk Sevilimis, and Benjamin Kimia. *Generalized Relative Neighborhood Graph (GRNG) for Similarity Search*. 2022. URL: <https://arxiv.org/abs/2208.10022> (cited on page 40).
- [114] Enrico Zardini, Enrico Blanzieri, and Davide Pastorello. 'A quantum k-nearest neighbors algorithm based on the Euclidean distance estimation'. In: *Quantum Machine Intelligence* (2024). doi: [10.1007/s42484-024-00155-2](https://doi.org/10.1007/s42484-024-00155-2) (cited on page 40).

- [115] Mathew D. Penrose and Xiaochuan Yang. *Largest nearest-neighbour link and connectivity threshold in a polytopal random sample*. 2023. URL: <https://arxiv.org/abs/2301.02506> (cited on page 41).
- [116] Sabeel Ansari, Lata Narayanan, and Jaroslav Opatrny. 'A Generalization of the Face Routing'. In: *International Conference on Mobile and Ubiquitous Systems: Networking and Services*. 2005. doi: [10.1109/MOBILQUITOUS.2005.3](https://doi.org/10.1109/MOBILQUITOUS.2005.3) (cited on page 42).
- [117] Thomas Clouser et al. 'Void traversal for efficient non-planar geometric routing'. In: *Ad Hoc Networks* (2013). doi: [10.1016/j.adhoc.2013.05.013](https://doi.org/10.1016/j.adhoc.2013.05.013) (cited on page 42).
- [118] Paul Balister et al. 'Connectivity of random k-nearest-neighbour graphs'. In: *Advances in Applied Probability* (1995). doi: [10.1239/aap/1113402397](https://doi.org/10.1239/aap/1113402397) (cited on page 42).
- [119] Chang Maw-Shang, Tang. Y. Chuan Yi, and R. C. T. Lee. '20-relative neighborhood graphs are hamiltonian'. In: *Journal of Graph Theory* (1991). doi: <https://doi.org/10.1002/jgt.3190150507> (cited on page 42).
- [120] George F. Riley and Thomas R. Henderson. 'The ns-3 Network Simulator'. In: *Modeling and Tools for Network Simulation*. Ed. by Klaus Wehrle, Mesut Güneş, and James Gross. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 15–34. doi: [10.1007/978-3-642-12331-3_2](https://doi.org/10.1007/978-3-642-12331-3_2) (cited on page 43).
- [121] Oluwarotimi Oluwafemi, Emmanuel Famakinwa, and Ometere Balogun. 'Random walk theory and application'. In: *World Journal of Advanced Engineering Technology and Sciences* (2024). doi: [10.30574/wjaets.2024.11.2.0116](https://doi.org/10.30574/wjaets.2024.11.2.0116) (cited on page 44).
- [122] Facundo Benavides et al. 'An Auto-Adaptive Multi-Objective Strategy for Multi-Robot Exploration of Constrained-Communication Environments'. In: *Applied Sciences* 9 (2019). doi: [10.3390/app9030573](https://doi.org/10.3390/app9030573) (cited on page 51).
- [123] Derek Ming Siang Tan et al. 'IR2: Implicit Rendezvous for Robotic Exploration Teams under Sparse Intermittent Connectivity'. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2024. doi: [10.1109/IR0558592.2024.10801761](https://doi.org/10.1109/IR0558592.2024.10801761) (cited on page 51).
- [124] Alejandro Mosteo, Eduardo Montijano, and Danilo Tardioli. 'Optimal role and position assignment in multi-robot freely reachable formations'. In: *Automatica* (2017). doi: [10.1016/j.automatica.2017.03.040](https://doi.org/10.1016/j.automatica.2017.03.040) (cited on page 56).
- [125] O. Melchert. 'Percolation thresholds on planar Euclidean relative-neighborhood graphs'. In: *Physical review. E, Statistical, nonlinear, and soft matter physics* (2013). doi: [10.1103/PhysRevE.87.042106](https://doi.org/10.1103/PhysRevE.87.042106) (cited on page 61).
- [126] Upma Jain, Ritu Tiwari, and W. Wilfred Godfrey. 'Comparative study of frontier based exploration methods'. In: *Conference on Information and Communication Technology (CICT)*. 2017. doi: [10.1109/INFOCOMTECH.2017.8340589](https://doi.org/10.1109/INFOCOMTECH.2017.8340589) (cited on pages 62, 76).
- [127] Haoyun Li et al. 'Large Language Model Based Multi-Objective Optimization for Integrated Sensing and Communications in UAV Networks'. In: *IEEE Wireless Communications Letters* (2024). doi: [10.1109/LWC.2025.3529082](https://doi.org/10.1109/LWC.2025.3529082) (cited on page 66).
- [128] David Molina Concha et al. 'Bayesian Optimization Framework for Efficient Fleet Design in Autonomous Multi-Robot Exploration'. In: *ArXiv* (2024). doi: [10.48550/arXiv.2408.11751](https://doi.org/10.48550/arXiv.2408.11751) (cited on page 66).
- [129] Guoxiang Zhao and Minghui Zhu. 'Pareto Optimal Multirobot Motion Planning'. In: *IEEE Transactions on Automatic Control* (2021). doi: [10.1109/tac.2020.3025870](https://doi.org/10.1109/tac.2020.3025870) (cited on page 66).
- [130] Alexis Barrios-Ulloa et al. 'Modeling Radio Wave Propagation for Wireless Sensor Networks in Vegetated Environments: A Systematic Literature Review'. In: *Sensors* (2022). doi: [10.3390/s22145285](https://doi.org/10.3390/s22145285) (cited on page 69).
- [131] Cedric Adjih et al. 'FIT IoT-LAB: A large scale open experimental IoT testbed'. In: *IEEE 2nd World Forum on Internet of Things (WF-IoT)*. 2015. doi: [10.1109/WF-IoT.2015.7389098](https://doi.org/10.1109/WF-IoT.2015.7389098) (cited on page 71).
- [132] Wikipedia. *Interquartile range* — Wikipedia, The Free Encyclopedia. <http://en.wikipedia.org/w/index.php?title=Interquartile%20range&oldid=1301081025>. 2025 (cited on page 72).
- [133] Pere Tuset-Peiro et al. 'A Dataset to Evaluate IEEE 802.15.4g SUN for Dependable Low-Power Wireless Communications in Industrial Scenarios'. In: *Data* (2020). doi: [10.3390/data5030064](https://doi.org/10.3390/data5030064) (cited on page 72).

- [134] Tianqi Chen and Carlos Guestrin. 'XGBoost: A Scalable Tree Boosting System'. In: *ACM SIGKDD Conference*. 2016 (cited on page 75).
- [135] Guolin Ke et al. 'LightGBM: A Highly Efficient Gradient Boosting Decision Tree'. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2017 (cited on page 75).
- [136] Liudmila Prokhorenkova et al. 'CatBoost: Unbiased boosting with categorical features'. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2018 (cited on page 75).
- [137] Leo Breiman. 'Random Forests'. In: *Machine Learning* (2001) (cited on page 75).
- [138] Pierre Geurts, Damien Ernst, and Louis Wehenkel. 'Extremely randomized trees'. In: *Machine Learning* (2006) (cited on page 75).
- [139] Hui Zou and Trevor Hastie. 'Regularization and variable selection via the elastic net'. In: *Journal of the Royal Statistical Society: Series B* (2005) (cited on page 75).
- [140] Yoav Freund and Robert E Schapire. 'A decision-theoretic generalization of on-line learning and an application to boosting'. In: *Journal of Computer and System Sciences* (1997) (cited on page 75).
- [141] Thomas M Cover and Peter E Hart. 'Nearest neighbor pattern classification'. In: *IEEE Transactions on Information Theory* (1967) (cited on page 75).
- [142] Harris Drucker et al. 'Support Vector Regression Machines'. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 1997 (cited on page 75).
- [143] Alex J Smola and Bernhard Schölkopf. 'A tutorial on support vector regression'. In: *Statistics and Computing* (2004) (cited on page 75).
- [144] Sven Zacharias et al. '2.4 GHz IEEE 802.15.4 channel interference classification algorithm running live on a sensor node'. In: *SENSORS, 2012 IEEE*. 2012. doi: [10.1109/ICSENS.2012.6411279](https://doi.org/10.1109/ICSENS.2012.6411279) (cited on page 75).