Multi-Robot Trajectory Generation using Market-Based Task Allocation for Robotic Search in an Urban Environment

B.A.Sc. Thesis

Yawar Ashraf (1006293672)

Supervisor: Professor Beno Benhabib
University of Toronto Department of Mechanical and Industrial Engineering
Computer Integrated Manufacturing Lab (CIMLab)

Division of Engineering Science
University of Toronto Canada

2024–04-12

Table of Contents

List of Figures

List of Tables

Abstract

Multi-Robot Task Allocation (MRTA) presents a significant challenge in the context of trajectory planning for robotic applications, with its inherently complex and NP-Hard nature. This research focuses on solving the MRTA problem in urban search missions, where the objective is to locate lost individuals using a Multi-Robot System (MRS). To achieve this, the research introduces a foundational lost person model that overlays a probabilistic distribution on a detailed urban map. This model informs the generation of coordinated, computationally efficient trajectories for each robot in the MRS, facilitating optimal task allocation within the search space. The study proposes a Clustered Contract Net Protocol (CNP) Task Allocation Algorithm, where robots are assigned to clusters of tasks rather than individual tasks, promoting area-based searches. The algorithm's approach results in improved search efficiency, enhanced task coordination, and scalability, making it suitable for dynamic urban environments. The research also explores an augmented approach with a frustration index, which helps reduce redundancy and improve task allocation by aligning robots with areas of similar density. The findings from extensive simulations across nine cities demonstrate a high success rate in locating lost persons. The Clustered CNP Task Allocation Algorithm significantly outperforms previous methods, reducing search times and increasing coverage. The research concludes with a discussion on scalability, robustness, and future directions for refining the algorithm, emphasizing the potential of MRS for search and rescue operations in urban settings.

Introduction

Trajectory planning and updating in the context of robotic applications represent fundamental challenges that extend across diverse domains within robotics [1]. One of the major areas of research in the field of robot motion and trajectory planning, due to its relevance to real world problems, is the Multi-Robot Task Allocation (MRTA) problem. This problem is characterized by the challenge of assigning particular tasks, from a larger set of tasks, to all robots in the MRS such that the success of a larger and more complex objective is optimized [2]. Upon closer examination of the problem, if we consider a set of Robots R, set of tasks T and attempt to assign the T to R while attempting to minimize the time it will take to complete all the tasks, this problem presents itself as being NP-Hard [2].

This research aims to coordinate a Multi-Robot System (MRS) for Urban Search missions, specifically aiming to locate lost individuals. Leveraging a foundational lost person model, the primary objective is to generate coordinated, computationally efficient and optimal trajectories for each robot in a Multi-Robot System over an urban search space. Furthermore, as the model incorporates real-time updates, the research focuses on generating updated trajectories for the Multi-Robot System to maximize the likelihood of successfully locating the lost person. The lost person's model will provide a detailed map of the urban space overlayed with the probabilistic distribution of the position of the lost person, constraining the scope of this research to an area coverage problem enveloped in a task allocation problem.

Research Objectives

Key elements of Multi-Robots Systems used in Search and Rescue (SAR) operations encompass collaborative mapping and situational assessment [3], distributed and cooperative area coverage [4] and cooperative search [5]. In a significant portion of the existing SAR literature, a centralized approach is taken toward task allocation [6]. However, this thesis will focus on exploring distributed task allocation algorithms which have the general advantage of enhanced robustness in challenging environments prone to experience agent loss or have unstable communication with the central decision making node. A potential decentralized algorithm that would work well for this thesis is a market-based generic algorithm which relies on utility functions, quantifying each robot’s fitness for specific search tasks modelled as trajectories [7].

Additionally, path planning and area coverage algorithms are primarily concerned with the shape of the survey area; nonetheless, a number of other key variables will be considered such as the environmental conditions and importantly the probability of failure to assess the effectiveness of the solution.

Methodology

This research will involve working with a Multi-Robot System and attempt to locate a missing individual in an urban environment using a decentralized search algorithm. The algorithm used in this research will utilize a market-based approach in an attempt to optimally calculate trajectories and update trajectories as new information becomes available in the lost person's model for each robot in the system [6]. The research will include modelling and simulating a Multi-Robot System and creating a model of the market-based approach algorithm. The primary objective will be to analyze the success rate of finding the lost person. Time permitting, next steps for the thesis will be to perform a complexity analysis against an increase in the number of robots, task allocation schedule depth and search area.

Urban Search for a Heterogeneous Multi-Robot System

This section focuses on attempting to understand the Urban Search and Rescue (USAR) problem in a deeper context. With this context at hand, the problem will be more easily understood as a task allocation problem. The Urban Search and Rescue problem lends itself a lot from the Wilderness Search and Rescue (WiSAR) problem. The USAR problem at a high level can be thought of as a WiSAR problem with more constraints placed on the search space which manifests itself in a more structured urban environment relative to the pseudo-unstructured environment nature of wilderness environments.

The Target

As framed earlier, the current problem is to identify a lost individual, often referred to as the target in this thesis, in an urban environment. In such a situation, there is little but extremely valuable information at hand typically. The key piece of information to initiate the search is the last known position (LKP) of the target. This can help confine the search space depending on how recent this position was. Moreover, information about the target's demographic data is also available and relevant to help categorize the target into expected behaviour patterns. A key piece of insight is that the LKP is the last true known position of the target as any more information about the target's position will promptly conclude the search and label it a success.

As it can be seen in the literature, it is very important to spend time gathering information about the environment as this allows the particular search and rescue (SAR) application to perform better as it is the foundation of the planning process [8]. In a similar fashion, search agents have access to the static environment data before the search begins. Additionally, search agents also have to collect information to develop a better understanding about the environment around them and hence be able to make more informed decisions. Based on all the above collected information, a lost person's model can be created which can generate a probabilistic distribution of the expected behaviour and consequential motion of the target.

Search Resources

In a search and rescue mission, both Unmanned Aerial Vehicles (UAVs) and unmanned Ground Vehicles (UGVs) can be utilized depending on the specific challenges and needs of the situation. Although both UAVs and UGVs present their own strengths in SAR operations, employing a heterogeneous multi-robot system made up of UAVs and UGVs has shown to be very effective in both USAR [9] and WiSAR operations [10].

Search Tasks

In order to inform the movement of a heterogeneous multi-robot system a foundational understanding of the target's movement is provided by the lost person's model. By adjusting the appropriate parameters of such a model, the expected behaviour of the target can be decently approximated. Under the assumption that such a model is an acceptable approximation of the target, its results can be leveraged to advise the search path planning of the heterogeneous MRS. Notably, the lost person's model can be used to generate tasks to be executed by the heterogeneous MRS. It is important to reference this key insight to understand the following discussions regarding the multi-robot task allocation problem.

Defining the Multi Robot Task Allocation Problem

To recap, the problem of assigning a set of tasks to a Multi-Robot System while attempting to minimize the time taken to complete all required tasks is NP-Hard [2].Resultantly, numerous alternative methodologies have been suggested over the years. Before exploring any of these techniques in further detail, it is important to understand how the multi-robot task allocation (MRTA) problem can be defined and consequently be modelled as per the nature of our problem.

For MRTA setups, there are a few key actors in the general case. For the problem mentioned above, the key actors needed can be defined as follows:

1. $R\text{: } Mobile\ Robots\ r_i\text{: } \{i = 1,\ 2,\ \dots n\}$
2. $T\text{: } Tasks\ t_j\text{: } \{j = 1,\ 2,\ \dots n\}$
3. $Cost\text{: } A\ set\ of\ the\ robots'\ costs.\ c_{ij}\ is\ the\ cost\ for\ robot\ r_i\ to\ execute\ task\ t_j$

For the most generalized case, the problem collapses into finding the most optimal assignment of tasks T to a subset of robots R which will attempt to accomplish these tasks [11].

$$A: T \rightarrow R \quad (1)$$

Additionally, we will be exploring market-based approaches and it is important to understand the slightly augmented formulation of such a problem. In market-based approaches, each robot $r \in R$ can perform a task $t \in T$ or a bundle of tasks $d \in D$, where each bundle is a subset of the tasks available $D \subseteq T$. Each robot, depending on the problem constraints, then may bid on either a single task or a bundle of tasks, $b_r(t)$ or $b_r(d)$ respectively.

In the case that a robot can bid on a bundle of tasks, the cost of each bundle can simply be calculated as follows:

$$b_r(d) = \sum_{i=1}^{m} c_{ri} \quad (i \in D) \quad (2)$$

In the equation above, i is the task number and c_{ri} is the cost of the task i performed by robot *r*, and *m* is the total number of tasks present in this bundle of tasks.

Modelling the MRTA Problem

The MRTA problem can be modelled in many different ways and the following subsections will take a deeper dive into some of the leading modelling approaches.

Optimal Assignment Problem (OAP)

The MRTA can be modelled as an Optimal Assignment Problem (OAP) where the goal is to assign the robots R to a set of tasks T in such a manner that the profit made by the assignment combination is maximized [13]. The objective to be maximized therefore is the matrix $P(rt)$ as shown and discussed below in equation (3). As is the case with most real world applications of the MRTA problem, the number of robots and the number of tasks is not the same, however, this problem can be overcome by creating virtual tasks or virtual robots with zero profitability such that the number of robots and tasks equalizes. Due to this variation, it can be assumed that the sets R and T have the same size N.

The OAP problem can mathematically be stated as: given an $n \times n$ matrix P, find the permutation α of 1, 2, 3, … n such that the following is maximized:

$$\sum_{i=1}^{n} P(r_i t_{\alpha(i)}) \qquad (3)$$

Such a task allocation matching between robots R and tasks T is called an Optimal Assignment.

Alliance Efficiency Problem

The alliance algorithm, initially developed to tackle NP problems [14], has undergone generalization for addressing a wide range of mono-objective optimization issues [15]. The Alliance algorithm can be formulated as both a Multi-Objective and a Mono-Objective Alliance algorithm [16]. The Algorithm has been inspired by the idea of multiple tribes working together to achieve a common goal in an environment that will provide them with the resources they need to survive. There are two key features that characterize each tribe, namely, the skills of each tribe and the resources they need to survive. In our case, each tribe can either be represented by a single robot or by a group of robots attempting to complete a singular task or a bundle of tasks D. The key idea of tribes becomes helpful when the tribes form alliances by pooling their skills in an attempt to achieve the common goal. This approach can be especially helpful for a heterogenous MRS where the intrinsic contrast in the robots' abilities becomes a beneficial feature to complete the task rather than a potential drawback. To formulate, a tribe t can be represented as a tuple composed of the following [17]:

- A point of the solution space to represent the location of the tribe x_t
- A set of skills $s_t = \{s_{t,1}, s_{t,2}, s_{t,3}, \dots, s_{t,N}\}$. These skills depend on the value of N_s of the objective function S_i evaluated at x_i

$$s_{t,i} = S_i(x_t) \qquad \forall\ i = 1, 2, \dots N_s \qquad (4)$$

- A set of resource demands $r_t = r_{t,1}, r_{t,2}, \dots, r_{t,N_R}$ that depends on the N_R constraint functions:

$$r_{t,i} = R_i(x_t) \qquad \forall\ i = 1, 2, \dots N_R \qquad (5)$$

- An alliance vector that contains the IDs of all the tribes that are allied to the tribe t. It is important to note that initially, every alliance is made up of a single tribe.

Discrete Fair Division Problem

The MRTA problem can also be formulated as a fair division problem [18]. The idea behind this approach is quite simple, given that there are a finite number of tasks and a finite number of robots, the goal is to divide up the tasks in such a way that the utility or the work being done by each robot is according to its fair share. A fair share for each robot can be characterized by the worth of the tasks being performed by the robot. Given a set of N robots and a set of tasks worth V, the fair share to be received by a homogenous group of robots can be concluded to be $1/N$ of the value V assigned to the complete set of tasks. The fair share for a robot in a heterogeneous MRS will require mappings between different robot types to determine a fair share for each robot.

Furthermore, the Fair division approach can be further subdivided into two classes depending on how a task has been defined for the specific problem:

1. Indivisible Tasks: Such tasks have to be given entirely to a single robot and can be considered to be atomic in nature.

2. Divisible Tasks: Such tasks can be divided into smaller tasks and such a division can result in a homogenous or heterogenous subset of tasks. Accordingly, the subset of these tasks can be given to a homogenous or a heterogeneous MRS.

At closer inspection of the literature, two different schemes emerge for the Fair Division problem. The first method is called the method of sealed bids [19]. The robots submit a bid for each task and the bids remain private until the end of the bidding period. Once the bidding period has ended, the auctioneer assesses the bids and assigns a winner based on the submitted bids. The winner of this auction will be the highest bid price based on the auction criteria.

The method of markers is the second scheme used for the Fair Division problem [20]. This method is particularly useful when the tasks at hand can be arranged in a linear fashion, this situation typically arises when there is a large number of small tasks that need to be assigned amongst a smaller set of robots. The N robots voice their opinion by placing $N - 1$ markers based on their assessment of what the fair share is and consent to take any segment of tasks that lie between any pair of their own consecutive markers. After this is complete, the robot with the leftmost consecutive markers is assigned the first segment and all its other markers are removed. This process is then repeated until all the robots receive what they perceive to be a fair share for themselves.

Just as a short note, the Fair Division Problem is a well studied problem and is known as the cake cutting problem when the resource R is heterogeneous and is infinitely divisible [21, 22]. A key concept to take away from the above discussion is the concept of *Envy-Freeness*. The concept of Envy-Freeness can be considered the epitome of fairness as put by [23]. Envy-Freeness requires that every agent prefers their own bundle of tasks more than that of any other agent.

MRTA Taxonomy

The following subsections describe MRTA problem approaches and the second builds on top of this by including a class that captures the interdependence of tasks.

Gerkey and Matarić

This section will take a deeper look into the taxonomy of MRTA as categorized by Gerkey and Matarić [24]. The Multi-Robot Task Allocation problems were categorized along three axes. The first Axis distinguishes between robots that can perform Single-Tasks (ST) or Multiple-Tasks (MT) at the same time. The second axis distinguishes between tasks that are characterized as Single Robot (SR) tasks or Multi Robot (MR) tasks. This axis helps identify tasks that only require a single robot to complete them or multiple robots to complete the task, intrinsically capturing the varying complexity of tasks. The third axis deals with Instantaneous Assignment (IA) or time-extended assignment (TA). This axis captures the scheduling complexity of tasks, IA tasks can be allocated without any consideration to future scheduling and TA tasks on the other hand must be assigned to a robot which must also follow a schedule according to which these tasks should be executed. Figure 1 presents a visual portrayal of the above taxonomy presented by Gerkey and Matarić [24, 25].

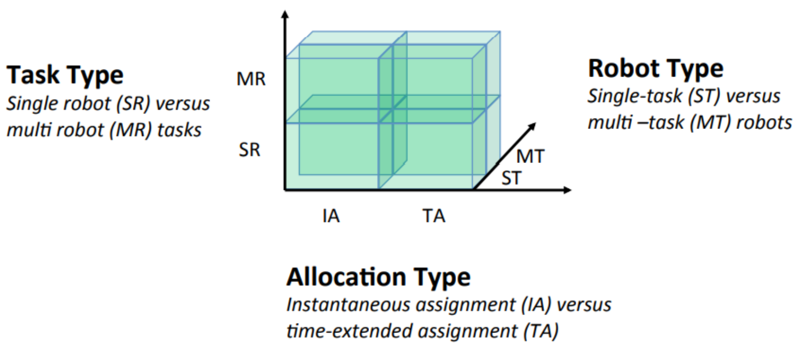

Figure 1: Visual Representation of the Gerkey and Matarić's MRTA taxonomy, from [25]

In presenting their taxonomy, Gerkey and Matarić noted that the ST-SR-IA problem is an optimal assignment problem in combinatorial optimization and can be solved in polynomial time, unlike the other problems which are all strictly NP-Hard [25]. They outline the ST-SR-TA problem, framing it as a variant of a machine scheduling challenge that entails devising task schedules for individual robots [25]. The ST-MR-IA problem, recognized as more challenging, is alternatively referred to as coalition formation. It involves organizing robots into distinct sub-teams without overlap for task execution, sharing mathematical similarities with the set-partitioning problem in combinatorial optimization [25]. And similarly, the MT-SR-IA problem is equivalent with the roles reversed between robots and tasks [25]. The ST-MR-TA problem is once again equivalent to the MT-SR-TA problem with the roles of tasks and robots reversed and involves both coalition forming and scheduling [25]. Moreover, to tackle the MT-MR-IA problem, the goal is to compute a tribe or coalition of robots to perform each task and it allows for a robot to be a part of more than one coalition [25]. Ultimately, the MT-MR-TA problem is an exceptionally challenging problem and can be understood as a scheduling problem with multipurpose machines and multiprocessor tasks [25].

Interdependence

A key limitation in the taxonomy provided by Gerkey and Matarić, as they pointed out, is that the interrelated task utilities and constraints are not captured in their taxonomy [25]. It is important to consider this since interrelatedness of tasks and utilities is especially relevant to real-world problems since it is akin to a robot routing problem. The utility function of each task assigned to a robot is related to the routing cost. If tasks are primarily spatially distributed but otherwise homogeneous, as is often the case with robot routing problems, in such domains, there is a harmony between tasks that are closer together. In such a case the utility of a robot that performs tasks that are spatially nearby is not equal to the sum of the utilities for performing these tasks individually [25]. For this reason Korsah et al have extended the framework presented by Gerkey and Matarić by adding a categorical variable with four unique values to represent the interdependence of tasks. The following subsections will briefly define and outline what these categories are.

- **No Dependencies (ND)**: Such tasks do not affect the utility of other tasks or agents in the system [25].
- **In-Schedule Dependencies (ID)**: Such tasks have agent-task utilities that have intra-schedule dependencies. Therefore, the tasks depend on the task performed by the agent and the other tasks also present within the agent's schedule [25]
- **Cross-Schedule Dependencies (XD)**: Such tasks have agent-task utilities that have inter-schedule dependencies. Therefore, the effective utility of an agent or task depends on both its own schedule and the schedule of other agent-task pairs in the system [25].

- **Complex Dependencies (CD)**: Such tasks have agent-task utilities that have inter-schedule dependencies specifically for complex tasks. Complex tasks can be decomposed into subtasks that may be multiagent-allocatable subtasks. Therefore, the utility of an agent depends on both the schedule of other agents in the system and the task decomposition that is chosen for the complex task.

Below is a visualization of the four possible values of the interdependence class presented by Korsah et al [25].

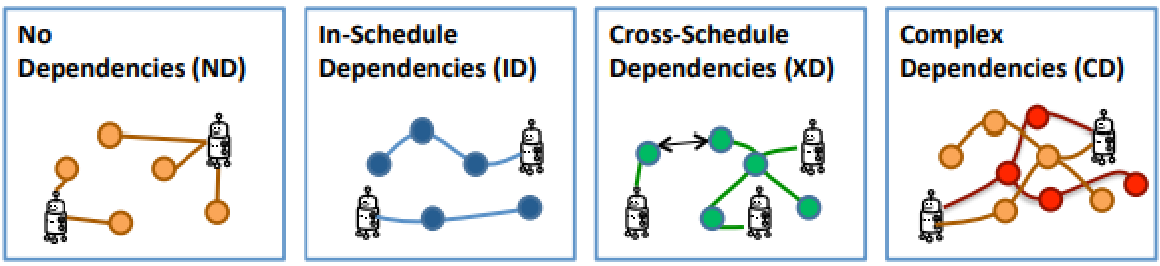

Figure 2: A visualization of the four categories for the interdependence class. Shaded circles represent tasks, solid lines represent agent routes and arrows represent constraints. The multiple colours or superimposed routes in the rightmost figure illustrate multiple task decompositions of complex tasks. From[25]

MRTA Solution Approaches

Two of the most commonly used MRTA approaches will be explored in the following subsections. Namely optimization based approaches including deterministic, stochastic and heuristic approaches and market based approaches after which the pros and cons of market-based approaches will be discussed.

Optimization Based Approaches

Optimization deals with attempting to find the optimal solution to a problem in a solution space given restrictions placed by solution constraints. The optimal solution is then chosen based on a criteria which informs the objective function that is to be maximized, quantifying the goal of the system. A range of optimization approaches can be explored, and the appropriate approach can be selected based on the context and complexity of the problem. Optimization based approaches are distinguished by their ability to explore new areas in the search space due to the inherent randomness of algorithmic variables and can benefit from noisy data as this may allow them to explore novel search spaces and consequently solutions [27, 28].

On the other hand, deterministic approaches follow a rigid path, always producing the same outputs with the same inputs, giving them a repeatable nature. Deterministic methods include numerical methods, classical methods such as gradient and hessian based approaches along with quadratic programming among others [20]. Graph based methods such as uninformed and informed search are also deterministic in nature [20].

Stochastic techniques inherently have randomness associated with their approach and these approaches can largely be classified into trajectory-based and population-based approaches. Simulated annealing is a metaheuristic algorithm and uses a single agent to search the space. Through the simulated annealing process, good and improving moves are accepted while less favorable moves have a much lower chance of being accepted. This mechanism will slowly favour the survival of the fittest and explore favorable solutions spaces.

Conversely, population based algorithms such as ant colony optimization and particle swarm optimization algorithms use multiple agents to help search the space collectively and leverage their large numbers. Population based algorithms have also been used for various MRTA problems. For example, the genetic algorithm used in [29] was capable of tracking a group of targets rather than a single target. Moreover, the genetic algorithm has also been used to provide a solution for the time extended task allocation of multi-robot systems in a disaster scenario [30]. Hybrid optimization approaches were also used to tackle the task allocation problem. In [31], a Simultaneous Task Allocation and Motion Coordination (STAMC) algorithm was proposed and explored the simulated annealing, ant colonization and an auction algorithm for task allocation in the STAMC approach. It was shown that the makespans achieved from all three algorithms while used in combination with STAMC were comparable to each other [31].

In [32], trajectory-based and population-based metaheuristics were introduced and extensively evaluated across various test scenarios. The study shows that such approaches are effective in tackling the complex and often heavily constrained nature of MRS applications as most real-world applications require a team of heterogeneous robots to successfully achieve a goal. Given that real-world MRS applications often involve diverse robots with varying capabilities, the proposed approach in [32] prioritized accounting for robot heterogeneity. Four key features—robot velocity, capabilities, energy level, and aging factor (efficiency)—were incorporated during the implementation phase and integrated into the traveling salesman model (TSP) [32]. Consequently, additional features—task requirements and minimum time needed to complete tasks—were introduced to enhance the representation of cities/tasks in the formulation [32].

Market Based Approaches

The market based approach is an economically inspired idea which attempts to mimic and model the free market by coordinating the activities between the tasks/goods and the robots/agents. This approach is largely based on the idea of auctions to simulate the economically competitive environment. The process of an auction can be described by the mechanism of trading rules, the process of assigning goods and/or services to a set of bidders based on the submitted bids and auction criteria. It is an intuitively simple method of allocating resources to the highest bidder and this concept can be extended to the MRTA problem where the most qualified and the most fit robots to perform a certain task will be able to submit the highest bid for a particular task. Therefore, following through completely with this mechanism will result in all the tasks being auctioned off or allocated between the set of Robots available.

Such an approach requires there to be communication between the robots and the tasks so the robots can submit bids for the tasks based on their fitness and ability to complete them. The negotiation process follows market theory principles, with the team aiming to optimize an objective function based on the utilities of robots in performing specific tasks [33].

Auctions in general allow for competition of resources between different market players and allows any interested party to indicate their interest in any number of resources available in the auction. Auctions are a simple and intuitive method of allocating goods to interested parties and there are multiple types of auctions that can achieve this goal. The types of auctions can largely be divided into two categories, namely Simple-good auctions and Combinatorial auctions. Simple auctions allow bidders to place bids on single items, however, combinatorial auctions allow bidders to place bids on a combination of items, called “packages”. In the coming subsections some of the leading auction designs for the Multi-Robot Task Allocation problem are discussed.

Contract Net Protocol (CNP)

The Contract Net Protocol is a task sharing mechanism for multiagent systems and allows for autonomous competitive negotiation between agents through the use of contracts in the auction space. Information can be distributed among agents using 3 channels[34]:

1. Nodes can request information directly from another node.
2. Nodes can broadcast information about a task.
3. Nodes can specify in their bids for a particular task what further information they need in order to execute the task.

The CNP algorithm can be broken down into 4 stages which provides a high-level structure to the interactions between the nodes, tasks and the coordination necessary to execute complex tasks [34].

- **Announcement stage**:An agent is assigned to be the auctioneer and is given the task of announcing the set of tasks that are available for bidding.
- **Submission stage**: Individual robots or agents can place bids on each of the tasks after calculating their individual utility values based on their utility functions.
- **Selection stage**: This stage commences once the bids are received from all the interested agents, at this point the auctioneer evaluates the bids based on the auction criteria in order to determine the winner of each task
- **Contract Stage**: Once the auction criteria has been calculated the winning agent is assigned the task at hand and the process continues to loop again until the set of tasks is exhausted and completed.

As the keen eye will identify, the main disadvantage of this approach is that the tasks are assigned to the robots/agents based on their individual self-interests. By extension, it can be concluded that the final solution may be optimal locally for each of the robots/agents, however it may or may not be the optimal for the complete multi-robot system as a whole [34]. An illustration of the CNP stages can be found below in Figure 3:

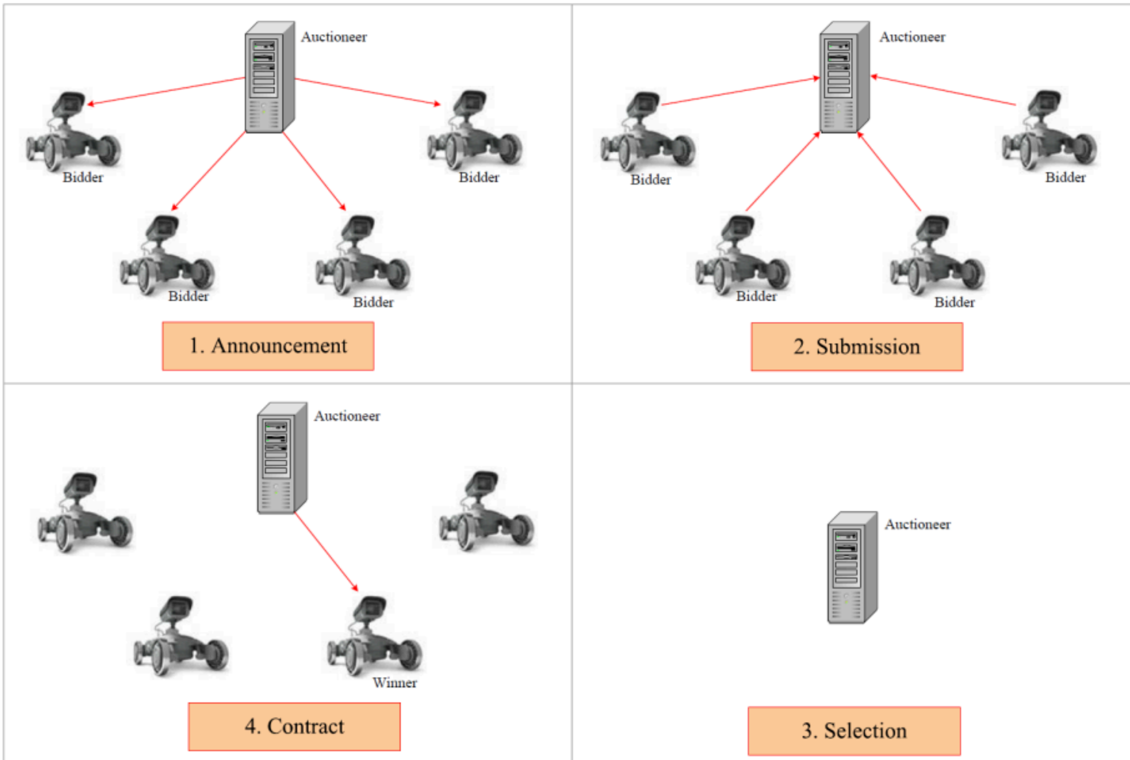

Figure 3: A visualization of the Contract Net Protocol (CNP) algorithm as shown by [14]

Trader-Bots

The Trader-bots approach also applies free market economy concepts for generating efficient and robust multi-robot coordination behaviours via the underlying market-based task allocation mechanism. The top layer of this approach is made up of trader bots which are assigned as a 1:1 ratio for each robot. Additional trading bots are also initialized for representing operators and other resources in the environment as deemed necessary. Each trader in this approach has the ability to negotiate and ensure that it is able to act in its self-interest by getting the best resources or tasks that it needs and is able to make rational decisions while negotiating its contract [35]. The high level goal of this algorithm is to help ensure the completion of the tasks at hand while also maximizing the utility of all the robots in the set. The main advantages of such an algorithm is that it is self-organizing in nature, making it robust and adaptable to unknown environments [35]. To ensure smooth organization and behaviour of the MRS, the Trader-Bots employ two mechanisms, namely the subcontracts and transfers.

- **Subcontract**: In such a contract the bidder agrees to complete a task for a certain price and then agrees to report back to the seller once the task has been completed.
- **Transfer**: In this case, the right to perform the task entirely is sold to the bidder. The bidder provides payment to the seller when it is awarded the contract through the transfer and is consequently not required to report back to the seller once the task has been completed.

Analysis of Market Based Approaches

This section will be dedicated to discussing the Pros and Cons of the Market Based approaches.

- **Efficiency**: Market based approaches are distinguished in their ability to capitalize on the locality of information, and can produce very efficient solutions by leveraging both centralized and decentralized elements of mechanism [36]. Market Based approaches have been shown to provide efficient solutions using varying objective functions [36, 37].
- **Robustness**: Market based approaches have the advantage of all agents being capable of making rational and self-interest oriented decisions as mentioned earlier. Due to this, it allows market-based approaches to not be susceptible to single point failures. Therefore, leveraging a decentralized paradigm can be made robust to several types of failures [36, 38].
- **Scalability**: The computational and communicational requirements of market-based approaches usually do not become prohibitive and therefore allow for efficient solutions to be computed even as the system scales [17]. Notably, market-based approaches scale especially well in cases where large complex tasks can be decomposed into small tasks that can be performed in parallel by smaller subsets of the robots [39].
- **Adaptability**: Market based approaches are usually capable of including new tasks into the task space seamlessly by auctioning them as they are introduced into the system or as they are generated by the agents themselves [39]. Additionally, market based approaches can also be configured to operate in unknown and dynamic environments by allowing agents to modify cost estimates of certain tasks overtime, and hence reallocating them to a more suitable agent if necessary [40].

While market-based approaches offer advantages, their drawback lies in the informal design of cost and revenue functions, as well as complexities in negotiation protocols and penalty schemes

[41]. Additionally, in cases where a centralized approach is able to suffice, market based approaches can often be more complicated to implement and may produce worse results as well [41]. Similarly, when completely decentralized approaches can be taken, market based approaches can be needlessly complicated in implementation and design on top of the additional computational and communication requirements [41].

Problem Framework

As described earlier in this report, the research problem at hand is to find an efficient and effective way to locate a lost person in an urban environment using a heterogeneous Multi-Robot System. The following section will attempt to solidify the framework of this problem in addition to clearly defining the problem scope.

A lost person can be defined as a mobile individual who does not have reliable knowledge of their current urban environment and requires external support to navigate to a safe location. Examples of an individual like this include an elderly person with dementia or a child which is unfamiliar with the environment. It is important to reiterate that the scope of this research is not designed to respond to a disaster situation such as a post earthquake lost person search effort. However, the urgency and swiftness required to rapidly act on the scoped non-disaster situation should not be underestimated. The objective of such a mobile target search effort should be to locate the lost person as quickly as possible while maximizing the chance of locating the lost person. It is important to appreciate the nuance of maximizing the chance of locating the lost person as quickly as possible since it is important to locate the person as soon as possible however not while significantly increasing the risk of failing to find the lost person at all. Hence, the strategy proposed in this paper is cognizant of this nuance.

Both human and robot agents can be used to search for the lost person. This research as mentioned above will be focused on developing a strategy for coordinating the robot agents. Notably, information from human actors can be used to advise the coordination of the robot agents by updating the weights of the lost person model. Adjustment will be explained in a later section. Robot agents will be tasked with advising their movements based on the information from the lost person's model.

Utilizing a multi-robot system is crucial due to the expansive nature of the problem at hand. By employing multiple robots, we not only enhance our ability to cover a larger area efficiently but also significantly elevate the probability of successfully locating the lost individual, thus maximizing the effectiveness of the search and rescue operation.

Lost Person's Model

The lost person model is a model developed by Cameron Haigh from the CIM Lab. The model works to model the movement of the lost person using a point cloud method. This model uses the environmental factors such as time of day, city structure and biographic information about the lost person including their Last Known Position (LKP) to tune the model to closely model the movement of the lost person. The model uses a point cloud method to model the target and a stochastic approach to calculate the expected behaviour of each of these points. A higher density of points in a certain location points towards there being a higher likelihood of the lost person being in the surrounding area according to the Lost Person model and vice versa. This idea is the key piece of information extracted from the Lost Person Model. Notably, the Lost Person model is able to adapt to suburban search areas as well. Suburban search areas present an additional layer of complexity where the Lost Peron may cease to follow paved roads or paths and wander off into large fields. Moreover, the model's output data resolution can be modified, however it is chosen to output data at the granularity of every minute.

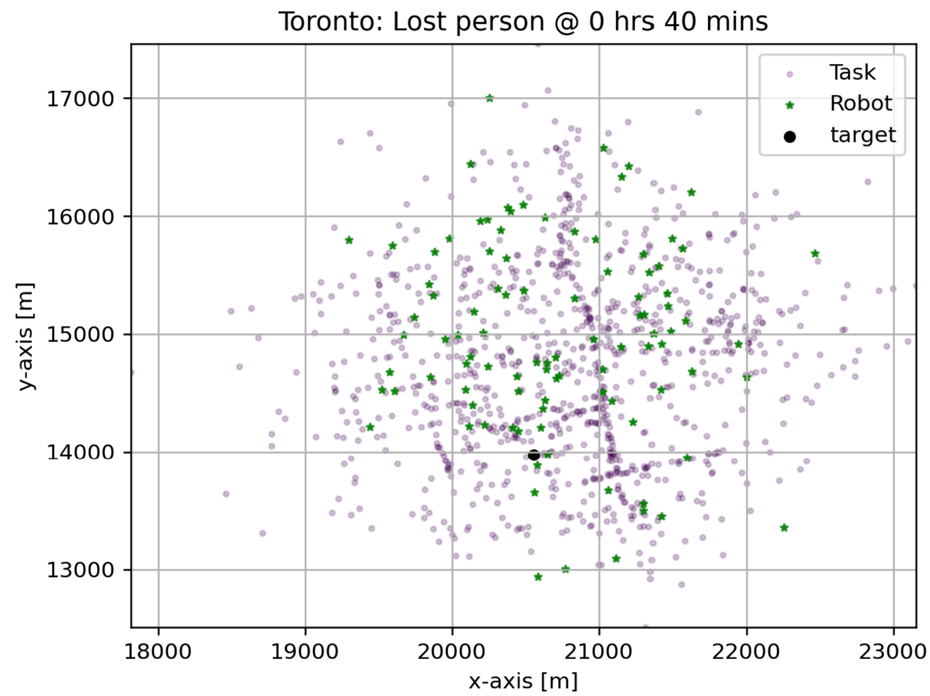

Figure 4: Visualization of the Lost Person model output

Robot Descriptors

While picking the robot actors, it is important to understand the tradeoffs between choosing an Unmanned Aerial Vehicle (UAV) and an Unmanned Ground Vehicle (UGV). UGVs take advantage of the urban road network of the urban/suburban area, however their speed and mobility can be majorly affected by ground factors such as car traffic and human safety considerations. This key drawback can limit the search capacity of UGVs and limit the rate at which the search operation can expand to external factors out of their control. On the other hand, UAVs can circumnavigate this problem since traffic and human safety constraints are much more relaxed. This would allow UAVs to rapidly expand their search area and help significantly increase the chances of finding the Lost Person. Additionally, as mentioned earlier, the lost person search could take place in a suburban environment where the lost person would have the chance to veer off dedicated roads and pathways into fields and large unstructured areas. UAVs would be instrumental in aiding the search in such scenarios which is why they were chosen to be the primary robot agent for this thesis.

Unmanned Ground Vehicles (UGV)	Unmanned Aerial Vehicles (UAV)
Easier to detect Lost Person due to closer proximity on the ground	Challenging to detect Lost person from a high altitude
Search speed limited by ground traffic and human safety factors	Search speed in bounded by the UAVs motor controls
Strictly requires paved surfaces to navigate	Can perform search over unstructured areas such as fields and farms
Largely unaffected to weather and environmental conditions	Susceptible to weather and environmental conditions, namely precipitation and high wind

Table 1: Discussing the Advantages (Green) and Disadvantages (Red) between UGVs and UAVs

Pseudo Iso-Probability curves

Iso-probability curves encompass the target's Last Known Position (LKP) and indicate the maximum distance achievable by the slowest Pth percentile target in any direction after a specific duration [43]. Iso-probability curves are not used as the defining feature of this algorithm, rather they are used as an augmenting feature to advise the search trajectories of the search agents after information from the Lost Person model has been exhausted after one cycle. The Pseudo-Iso-Probability curve is defined by the circular motion the search agent adopts in order to search for the target. The reason why it is labelled the Pseudo-Iso-Probability curve is because it assumes the pseudo-correct assumption that the robot is at the LKP position of the target and performs an iso-probability curve search motion of the 99th percentile of the target location. The Iso-Probability curve in this curve grows in radius as time passes to account for the fact that it is a mobile target and the 99th percentile contour of where the target could be present would grow with time as well.

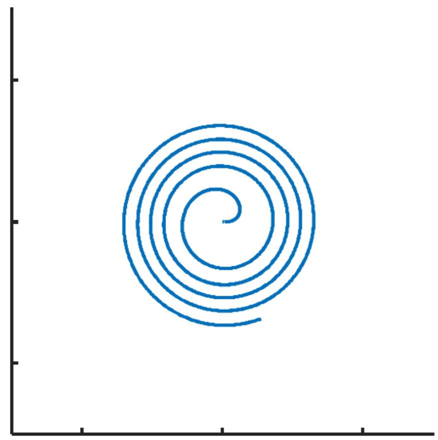

Figure 5: Visualization of the Pseudo-Iso-Probability trajectory as it grows with time by making concentric circles with a growing radius. The y and x axis are the x and y coordinates of the trajectory [43]

Targets and Actors

While extracting data from the Lost Person model, the data points along with the search agents must be classified into functional groups in order to set up a robust and effective simulation environment.

Task

The data received from the Lost Person model is temporal in nature and for the purpose of this paper, simulations using 1000 points were used. Each of those 1000 points were available as x and y coordinates over a period of 12 hours but target pose information only being available at the granularity of every minute. Each point was defined as a *Task* which must be completed by the search agents. Furthermore, a task would be classified as completed when a search agent robot is within a certain detection threshold of the *Task*.

Cluster

A Cluster is defined as a group of N tasks. The grouping is done by grouping the nearest neighbours together. For each cluster, the centroid of the cluster is also defined and the area coverage of each cluster is also defined. Clustering provides the benefit of grouping together certain *Tasks* and allowing a more efficient and effective distribution of search agents over the search area. Specifically, clustered *Tasks* allow a single search agent to effectively focus on *N* tasks rather than assigning *N* robots to search over *N* tasks. The reason this would not be optimal is because these tasks are closest to each other in proximity and there will be a larger time the search agents spend moving between tasks than focusing on performing the search at the requisite areas.

Robot

The search agent is defined as a robot. Each robot can have unique characteristics such as its maximum velocity and binary detection threshold. This allows for a heterogeneous system of search robots to be involved in the search. Robots are mobile UAVs that have direct communication with a centralized node to help coordinate the Multi-Robot System in the most effective manner to ensure the successful search of the Lost Person.

Target

The Target defined for each experiment is chosen at random from amongst the total tasks or lost person's simulated in the Lost Person model. The motion of the target, fittingly, is defined by the Lost Person model. Once a random target is picked, the pose of the target is masked from the search agent's, effectively reducing the number of available Tasks to search to become *N-1*. Where *N* is the total number of lost persons simulated by the Lost Person model.

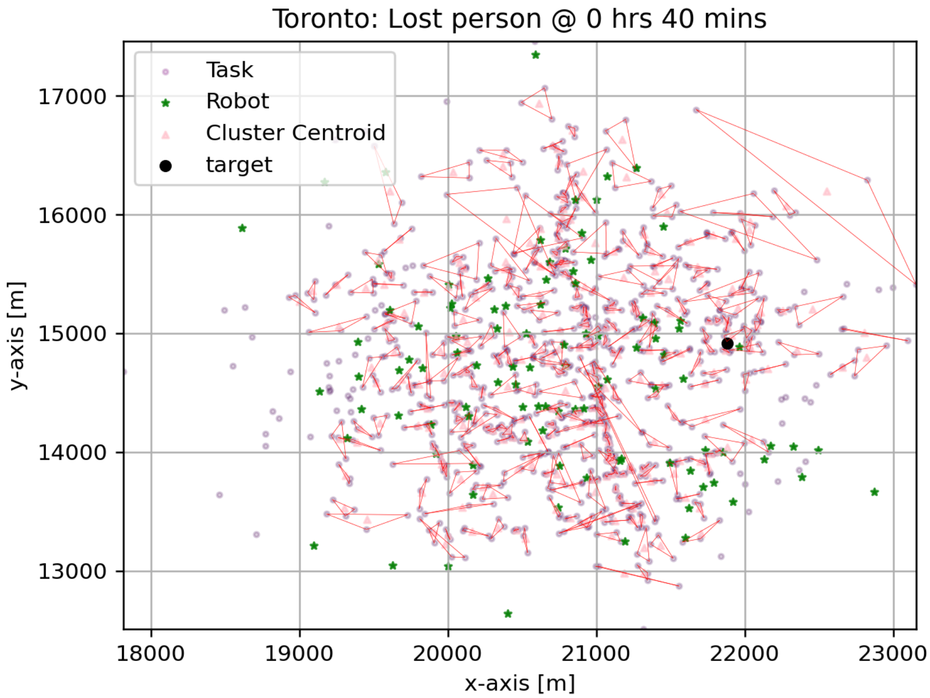

Figure 6: Visualization of the MRS Task Allocation Simulation

Identifying the Robot Taxonomy Region

We can identify the taxonomy that this problem belongs to, using the Gerkey and Mataic matrix. This identification will help us with understanding the core applications of this new method. As discussed earlier, it only requires a single robot to perform a single task and due to the nature of the tasks, the robots can only perform a single task at any given moment. It is impossible for a robot to be present at more than one location at a given moment, therefore we require a single robot to be assigned to a single task at any given moment. Furthermore, the proposed method in this paper classifies the tasks as being part of the instantaneous class of task assignment. The output of the Lost Person model is not locally calculated by each robot independently, rather the output is communicated to all the robots from a central node of communication. The reason for this centralized mode of communication is to ensure that the proposed method has the flexibility to account for any new information received from human search agents (i.e. any information collected from agents other than the UAVs).

Notably, the tasks do not have an interdependent nature, however the tasks are not mutually independent given the very nature of how these tasks are generated and processed. To complete any arbitrary task, no other task needs to be completed however, due to the stochastic nature of the way the task points are created, it is expected that tasks will be found in much higher densities closer to the global mean of the task distribution and the density of tasks will be lower as we deviate from the global mean. The proposed method aims to exploit this relationship between tasks, however, for the purpose of this paper the tasks can be classified as not having any dependencies between each other.

Multi Robot Task Allocation Problem Taxonomy	
Task Type	Single Robot (SR)
Robot Type	Single-Task (ST)
Allocation Type	Instantaneous Allocation (IA)
Interdependence	No Dependance (ND)

Table 2: Multi Robot Task Allocation Problem Taxonomical Classification

Defining as an Alliance Efficiency Problem

The problem can be defined as an Alliance Efficiency problem with a mono-objective optimization issue. Its inspiration stems from the concept of diverse tribes collaborating to achieve a shared objective within an environment that sustains them with necessary resources. Each tribe, represented by a single robot, is characterized by its unique set of skills and resource requirements. We can define tasks as resources that robots require to survive, hence each robot is required to complete tasks in order to "survive". The Alliance Efficiency formulation becomes beneficial when we consider the case where the amount of tasks available are less than the number of robots to help complete these tasks. In this case, we can create an Alliance A_f which is the final alliance before a new cycle of the control loop is executed. This final alliance will grow in size as the number of robots not assigned to a Task generated by the Lost Person model increases. In order to help provide resources to all tribes in this alliance, each robot which is a part of this alliance will follow the pseudo-iso-probability trajectories described earlier to collect their resources.

Contract Net Protocol (CNP)

The proposed method in this paper will be using the Contract Net Protocol (CNP) method as the task sharing mechanism for multiagent systems to allow for autonomous competitive negotiation between agents through the use of contracts in the auction space. To recap, information can be distributed among agents using 3 channels [34]:

1. Nodes can request information directly from another node.
2. Nodes can broadcast information about a task.
3. Nodes can specify in their bids for a particular task what further information they need in order to execute the task.

The CNP algorithm can be broken down into 4 stages which provides a high-level structure to the interactions between the nodes and tasks and a detailed discussion of this can be found earlier in this paper.

Experimental Setup

Some of the key metrics to be extracted from this study is to understand how successful each method is and to also understand how swift each method is in delivering a successful search of the Lost Person. The importance of both of these key metrics cannot be understated. In addition to both these key metrics, it is important to evaluate the robustness of these methods. Therefore, in order to evaluate the robustness of these methods, 9 cities across Ontario were chosen to represent 3 classes of urban centers:

1. Large Metropolitan Cities
2. Mid-size suburban Town/City
3. Small Rural Town

The cities were classified into each of the 3 different bins by looking at the population of each of these cities. Table 3 below shows a breakdown of how these cities were selected and which bins they have been classified into. By testing the efficacy of the proposed method on different urban areas we can gain a much more nuanced understanding of the performance and robustness of the proposed method.

Class	City	Population
Large Metropolitan City	Toronto	2,930,000
	Ottawa	995,000
	Mississauga	830,000
Mid-size suburban City	Oakville	210,000
	London	400,000
	Kitchener	240,000
Small Rural Town	Belleville	50,000
	Cobourg	19,500
	Orillia	31,000

Table 3: City classification and populations

Methodologies

This section will introduce the three methodologies analyzed in this paper and present and analyze experimental results from each of them.

Tracking Agents

Before discussing how each method works, it is important to understand how each Task and Cluster is classified throughout the lifecycle of the algorithm. An appreciation of the internal bookkeeping of the algorithm will permit a deeper understanding of the methods proposed.

All *Tasks* and *Clusters* are classified as at least one of the following throughout the lifecycle of each method mentioned below:

- **Unassigned:** An item which has not been assigned to be searched by a search agent.
- **Assigned:** An item which has not been assigned to be searched by a search agent.
- **Complete:** An item which has been searched by a search agent.
- **Incomplete:** An item which is currently being searched by a search agent

Furthermore, each robot can be classified as assigned or unassigned:

- **Unassigned:** A robot which has not been assigned an item (*Task*, *Cluster*)which it is tasked with searching.
- **Assigned:** A robot which has been assigned an item (*Task*, *Cluster*)which it is tasked with searching.

Closest Task Completion (Baseline)

The Baseline method was chosen to provide a foundational lower bound to understand how well we can expect the problem to be solved without using any deep reasoning. The basic mechanism of this method works by using the simple concept of assigning available search agents to the closest available *Task*. Similar to the latter methods, the baseline method functions in a cyclical method. A single cycle is characterized by the exhaustion of all available tasks. All available tasks are deemed exhausted when there are no more tasks left to be assigned (i.e. unassigned tasks). Once this happens, a new cycle of the algorithm is started where all tasks are classified as incomplete, hence, there are no complete tasks either. Moreover, the incomplete tasks at the end of the previous cycle are classified as assigned tasks and the remaining tasks are flagged as unassigned. The algorithm continues to function until either the Lost Person (i.e. Target Task) is located successfully or the search termination threshold of 12 hours has been reached. The scope of this research was to successfully and rapidly locate the Lost Person, hence the scope of the search was limited to 12 hours since the LKP of the Lost Person.

As seen below in Figure 7, the success rate of using the baseline method was quite unsatisfactory and would result in most of the search attempts ending in a failure. The highest rate of success was seen when using 100 robots for the search task and the success rate can be seen to steadily decrease as the number of robots decreases. The figure below is the culmination of a total of 900 simulations and is able to provide a high level overview of the performance of the Closest Distance Method. This performance was expected from the above method since it does not take advantage of the inherent structure built into the problem as we can leverage information from the Lost Person model to characterize areas of high likelihood and areas of lower likelihood and assign search resources appropriately.

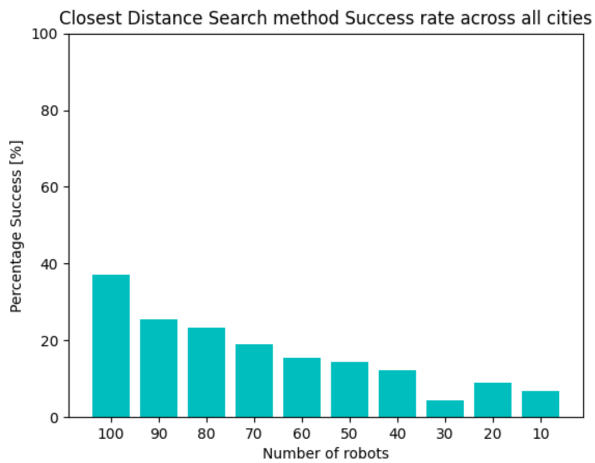

Figure 7: Success rate of using the Closest Distance method across cities of all sizes using a varying number of robots.

Looking deeper at the results, unique trends can be uncovered with regards to the performance of the method on cities of different scales. Figures 8 till Figue 10 illustrate these results.

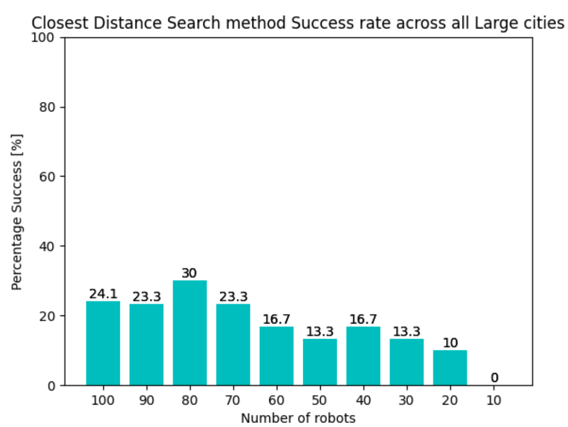

Figure 8: Success rate of using the Closest Distance method across Large cities using a varying number of robots.

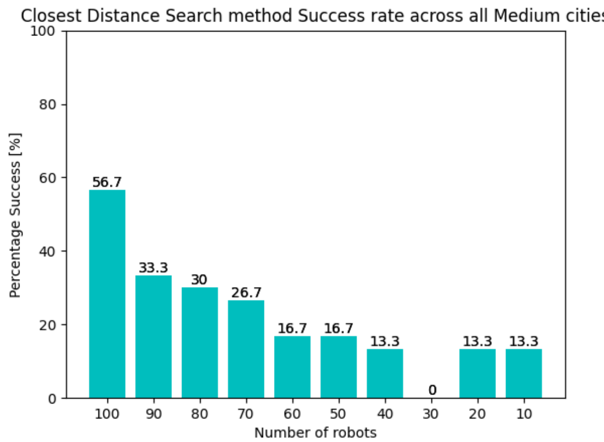

Figure 9: Success rate of using the Closest Distance method across Medium size cities using a varying number of robots.

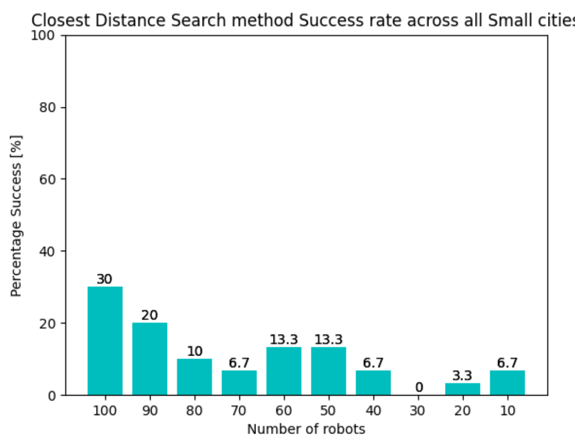

Figure 10: Success rate of using the Closest Distance method across Small size cities using a varying number of robots.

Interestingly, the performance of the method was largely invariant to the number of robots deployed in the search effort. This phenomenon can be explained by the fact that due to the large size of the city it is unlikely that the Lost Person will be found after a shorter time threshold than a smaller city. The reason can be explained to be that for a city with a large area, each robot has a larger area that it is effectively in charge of. Therefore, the most likely manner in which the Lost Person is located is if the robot is fortunate enough to be in relatively close proximity to the target. If not, then the robot is unable to locate the Lost Person.

Furthermore, results from searching small cities are slightly paradoxical in nature. There is a clear trend in between the success rate of the search and the number of robots used. The success rate gradually goes down as the number of robots used in the search is reduced. However, the low success rate of the method is a result of the fact that due to the fact that all the tasks are clumped together in much closer proximity relative to larger cities, the robots are biased towards searching areas that are present in the highest density of tasks. This bias creates a window for the target task to potentially escape farther away from areas of high density and can allow the target task to wander much farther away. Hence, the surprisingly low efficacy of this method on smaller size cities. Moreover, the Closest Distance method has had the highest success in medium sized cities as can be seen in the figures above.

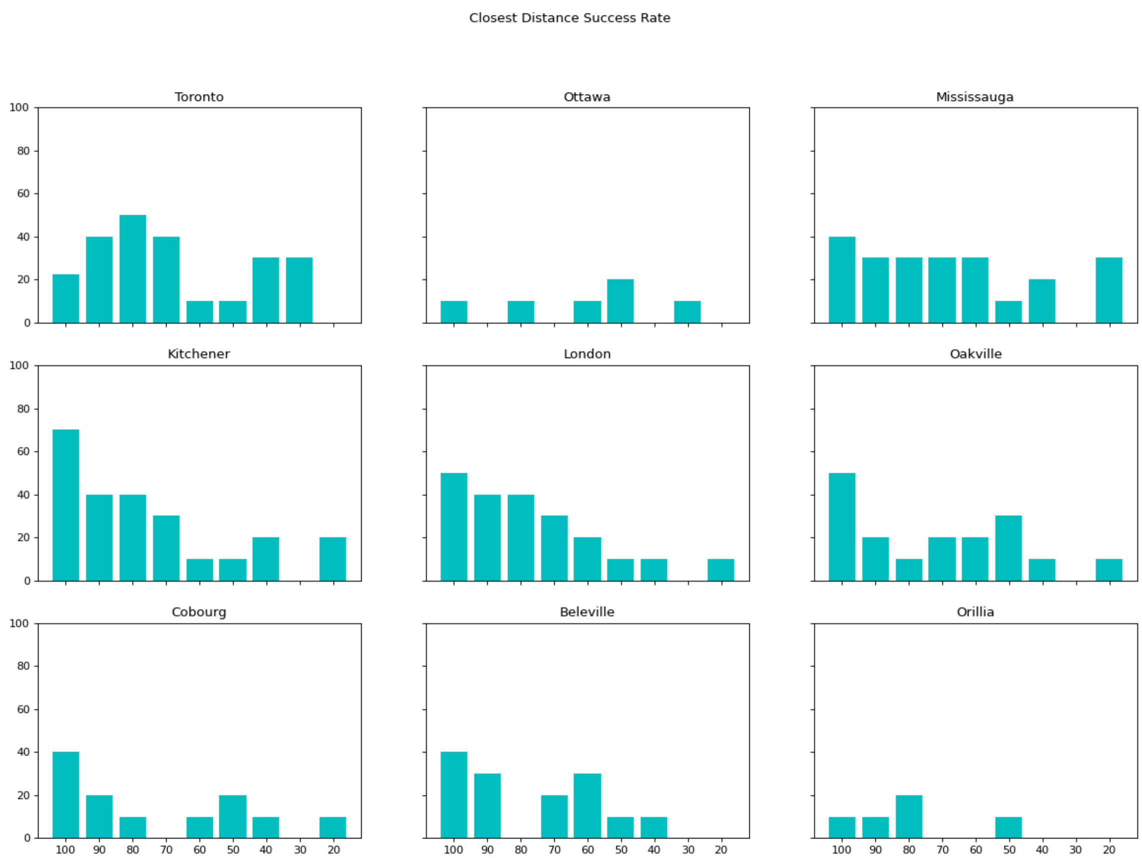

Figure 11: Closest Distance Success Rate across all cities with varying number of robots

Clustered CNP Task Allocation

This section will discuss the results and findings from the Clustered CNP (Contract Net Protocol) algorithm which presents part of the main algorithm of this research. Figure 19 below shows a visualization of the Clustered CNP Task Allocation Algorithm at a snapshot in time.

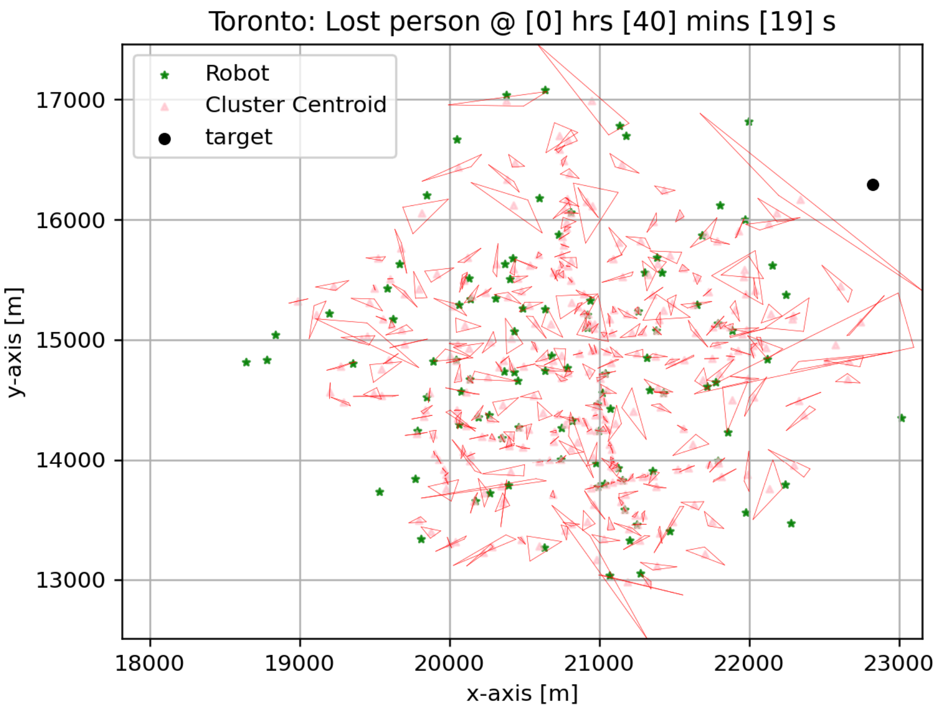

Figure 12: Visualization of the Clustered CNP Task Allocation Method Simulation

Algorithm Description

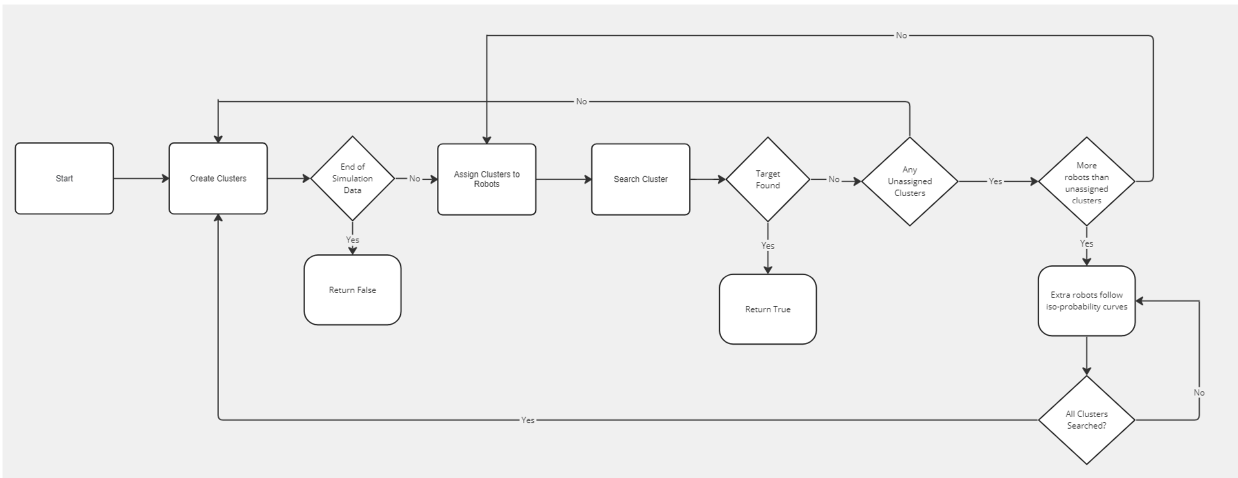

Figure 13: Clustered CNP Task Allocation state diagram

This algorithm uses the CNP Task Allocation algorithm described earlier to coordinate tasks between search agents and the framing has not been replicated here to prevent duplicate descriptions.

As discussed earlier, at the start of the search operation all tasks are characterized into four bins. Namely, assigned tasks, unassigned tasks, complete tasks and incomplete tasks. Similarly, Clusters are created to help aid this method. The clusters are created by clustering using the K-Nearest Neighbours approach and in the following experiments each cluster is made up of 3 Tasks. Moving forward, each robot is assigned to a Cluster instead of a Task. The key difference in the search task of the robot is now the robot has to search over a triangular area as opposed to visiting a single coordinate. The task of searching over an area forces the robot to generate complex trajectories as opposed to driving on a line to arrive from one Task to the next. Similar to tasks, once a cluster is completely searched, the robot is assigned to search the next closest Cluster. As one can predict, since searching large cluster areas requires a non-trivial amount of search time, there will consequently be a non-trivial amount of time where the number of clusters available to search is less than the number of search agents available (UAV robots). In this case, the robots not assigned to search a cluster are tasked with following iso-probability curves at their local coordinate. Since their location was advised by the Lost Person model, the surplus robots follow a trajectory of growing concentric circles similar to Figure 5 at their current location until all clusters have been searched. Once all clusters have been searched, the cycle is repeated until either the Lost Person is found or the cutoff time of 12 hours is reached.

Results

Due to the larger area cover of the Clustered CNP Task Allocation approach, the algorithm is expected to outperform the Closest Distance approach mentioned earlier. The data below has been collected by conducting 900 simulations over 9 cities and 10 simulations per each bin of robot numbers for each city (i.e. 100, 90, 80, etc). As illustrated in Figure 14 below, the success rate has drastically increased as the success rate of locating the Lost Person can be seen to be nearly 80% successful even with only 30 robots over the search space of 9 cities of varying size. The success rate of this method can be largely attributed to the fact that the robots now perform searches over an area as opposed to focusing their search efforts on a narrow strip of land while travelling from one task coordinate to the next. This mechanism allows the robots to take advantage of the fact that there is a high probability of the Target Task being located in regions with high Task density (i.e. smaller cluster sizes) and tasks the robot to search the space in between tasks in this manner rather than only focusing on search from Task coordinate to the next coordinate.

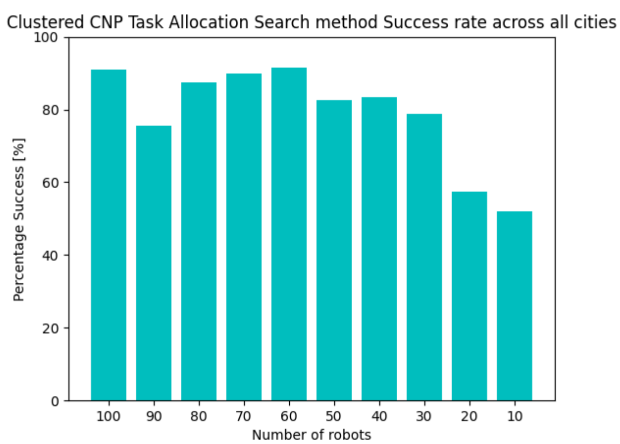

Figure 14: Success rate of using the Closest Distance method across cities of all sizes using a varying number of robots.

Looking deeper at the results, unique trends can be uncovered with regards to the performance of the Clustered CNP Task Allocation method on cities of different scales. Figures x-x+3 illustrate these results.

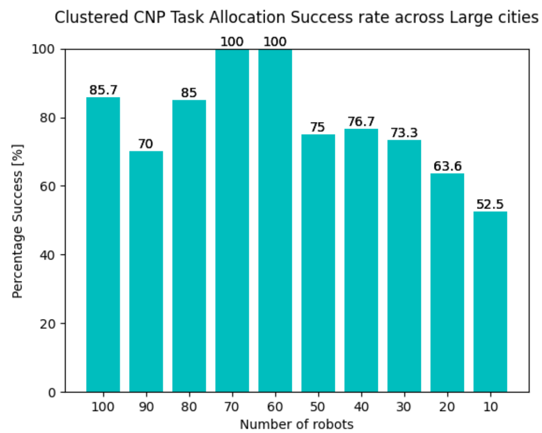

Figure 15: Success rate of using the Clustered CNP Task Allocation method across Large cities using a varying number of robots.

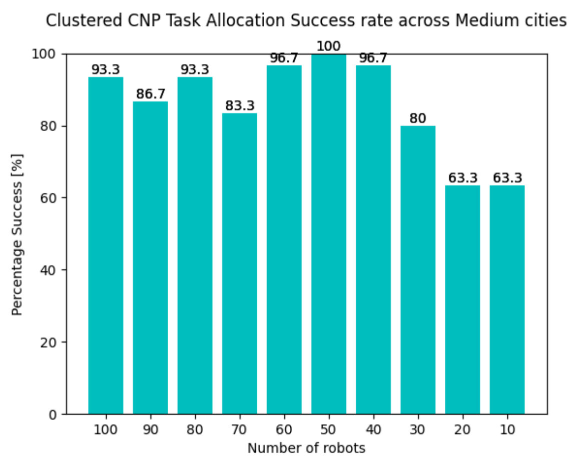

Figure 16: Success rate of using the Clustered CNP Task Allocation method across Large cities using a varying number of robots.

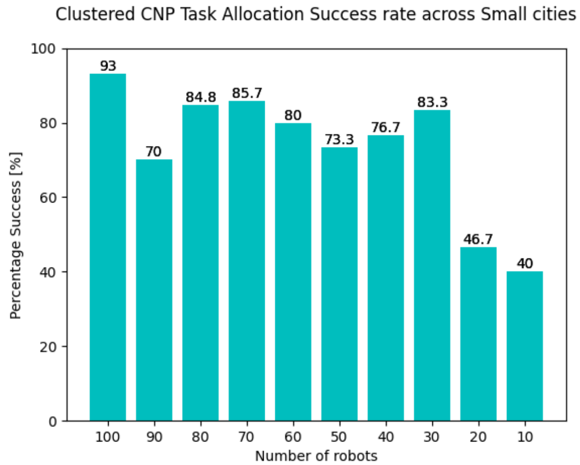

Figure 17: Success rate of using the Clustered CNP Task Allocation method across Large cities using a varying number of robots.

As we can see from the graphs above, there is a significant increase in the performance of the Clustered CNP Task Allocation algorithm relative to the Closest Distance method. All three categories exhibit the behavior of locating the Lost Person with lower success as the number of robots decreases. Additionally, as expected, smaller cities have a higher success rate of locating the Lost Person relative to the medium and large urban centers. This shows that the method is able to mitigate the bias present in the Closest Distance method described earlier.

Concludingly, this method is able to provide satisfactory results of successfully locating the Lost Person up to as low as using 30 robots across cities of varying sizes. Figure 18 below shows the performance results of the clustered CNP Task Allocation algorithm across all 9 cities.

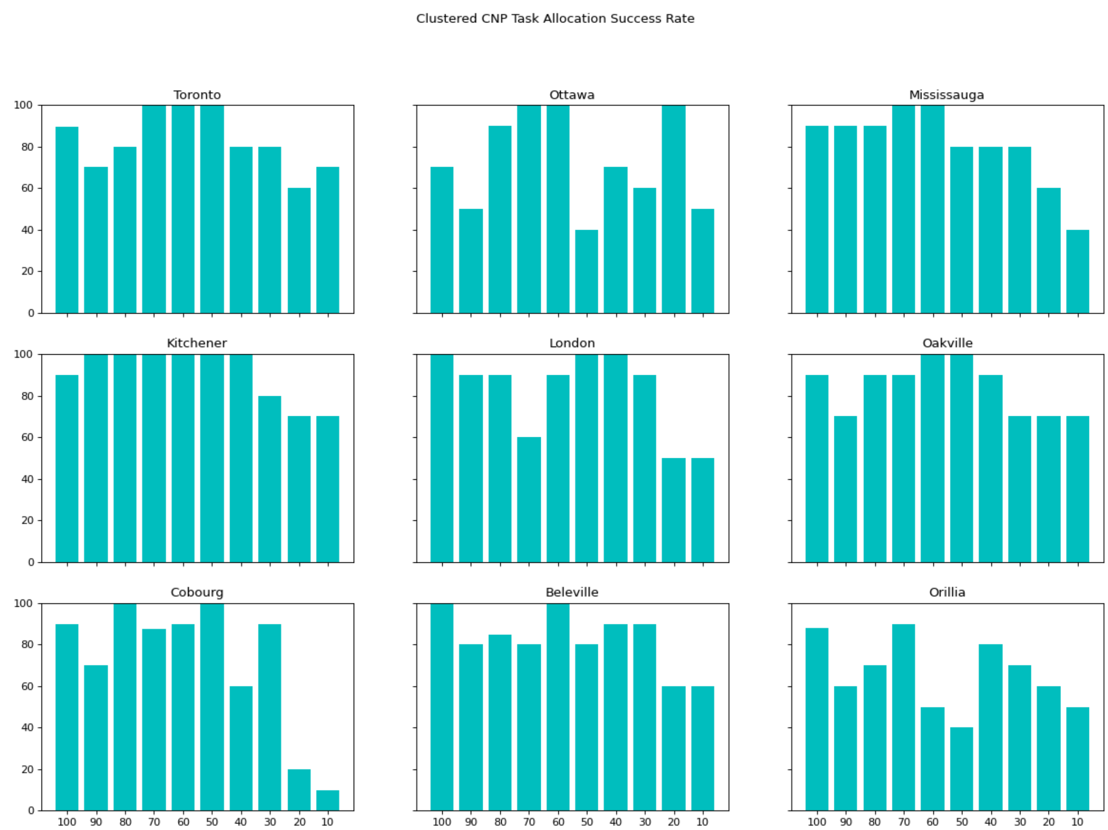

Figure 18: Clustered CNP Task Allocation Success Rate across all cities with varying number of robots

Clustered CNP Task Allocation with Frustration Index

Similar to the Clustered Task Allocation method described above, the method presented in this section is augmented with a frustration index associated with each robot to help assign the robot to the best fit. The task assignment protocol is still dictated by CNP however, robots are now able to be paired up with clusters that match their frustration index. A new variable is introduced in this method, namely a frustration index which is associated with each robot and each cluster. The idea behind the introduction of this new variable is that each robot will have a relaxed constraint to be assigned to a specific area of the search space characterized by the density of tasks in that area. Namely, some robots will be constrained to search high cluster density areas and some robots will be tasked with searching lower cluster density areas. This will minimize the time robots spend commuting in between these areas and will consequently spend a higher amount of time performing useful search and consequently not creating large windows for the target task to escape from high cluster density areas to the search frontier where the presence of search robots is highly sparse.

Each cluster has an associated dynamic frustration index calculated using the following formula:

$$Frustration\ Index\ =\ \frac{Cluster's\ centroid\ distance\ from\ map\ mean}{Maximum\ distance\ from\ map\ mean\ to\ map\ boundary}$$

The map mean is calculated using the average of all task coordinates. The map boundary is defined while initializing the simulation and it is decided to be the farthest a task will travel from the initial map mean. Furthermore, each robot is initialized with its frustration index at 0. Its index grows as it searches more clusters and the robot frustration index is updated to be the average frustration index of all clusters it has searched. Consequently, while assigning a robot to its appropriate cluster, the 5 closest clusters measured using euclidean distance are chosen and the robot is assigned to the cluster with the closest frustration index to itself. Ties are broken randomly.

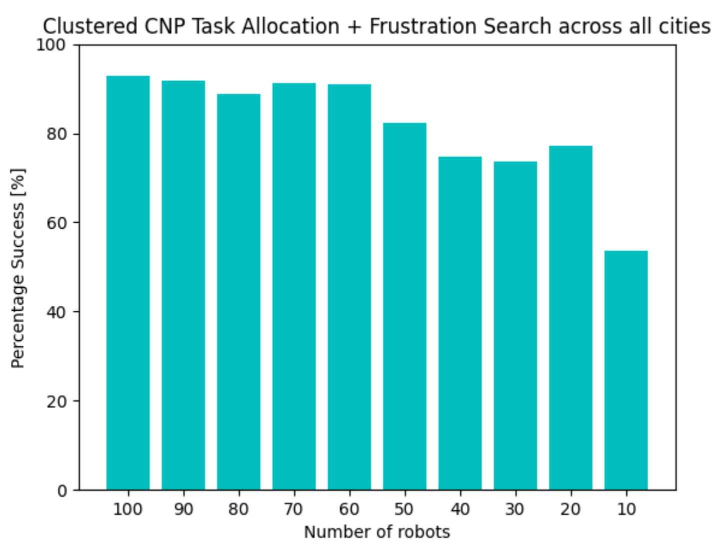

Figure 19: Success rate of using the Clustered CNP Task Allocation with Frustration Index method across cities of all sizes using a varying number of robots.

Looking deeper at the results, unique trends can be uncovered with regards to the performance of the Clustered CNP Task Allocation method on cities of different scales. Figures x-x+3 illustrate these results.

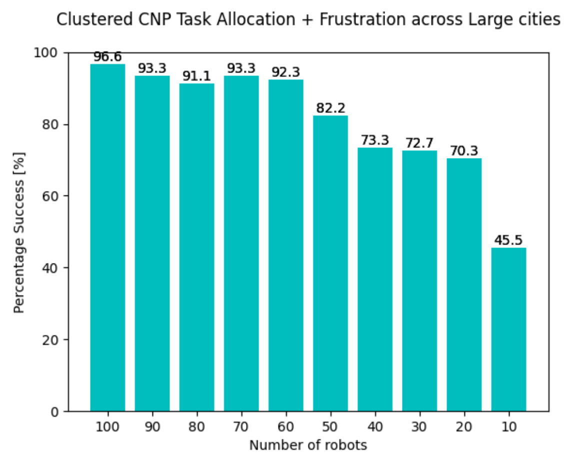

Figure 20: Success rate of using the Clustered CNP Task Allocation with Frustration Index method across Large cities using a varying number of robots.

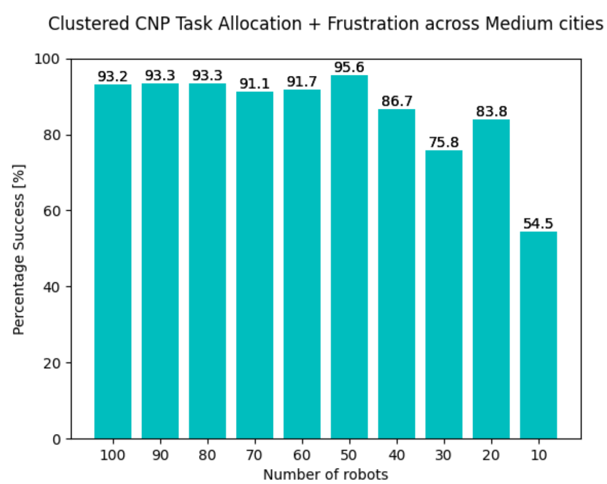

Figure 21: Success rate of using the Clustered CNP Task Allocation with Frustration Index method across Medium cities using a varying number of robots.

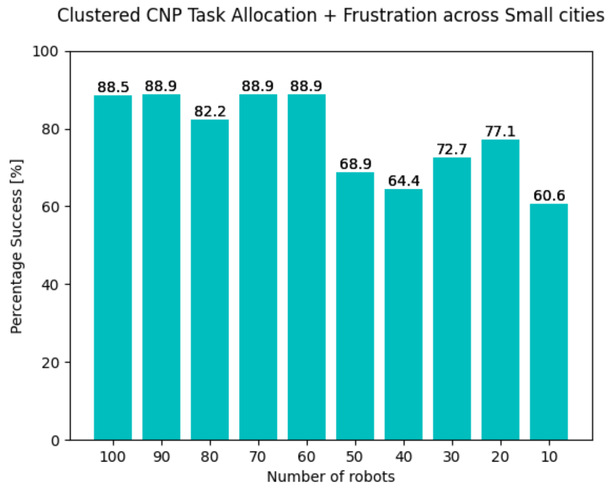

Figure 22: Success rate of using the Clustered CNP Task Allocation with Frustration Index method across Small cities using a varying number of robots.

By inspecting the graphs above, it can be clearly seen that the maximal successful performance peaks of the method have been decreased. However, it can be seen that the frustration index was able to successfully create an expected impact in the results as it has now smoothed out the variation in success rate between varying numbers of robots used for search. While examining graph x, we can notice that there is a decreasing linear relationship between the number of robots used and the success rate of the search.

Additionally, further inspection reveals a similar trend for all three urban center categories. The frustration index has been shown to successfully smooth out the performance of the Clustered CNP Task Allocation described earlier to be more reliable. This increased reliability creates a defined window for the search to be augmented with human search agents and consequently increase the already competitive success rate of the proposed search method. Figure 23 below shows the performance results of the clustered CNP Task Allocation algorithm with Frustration Index across all 9 cities.

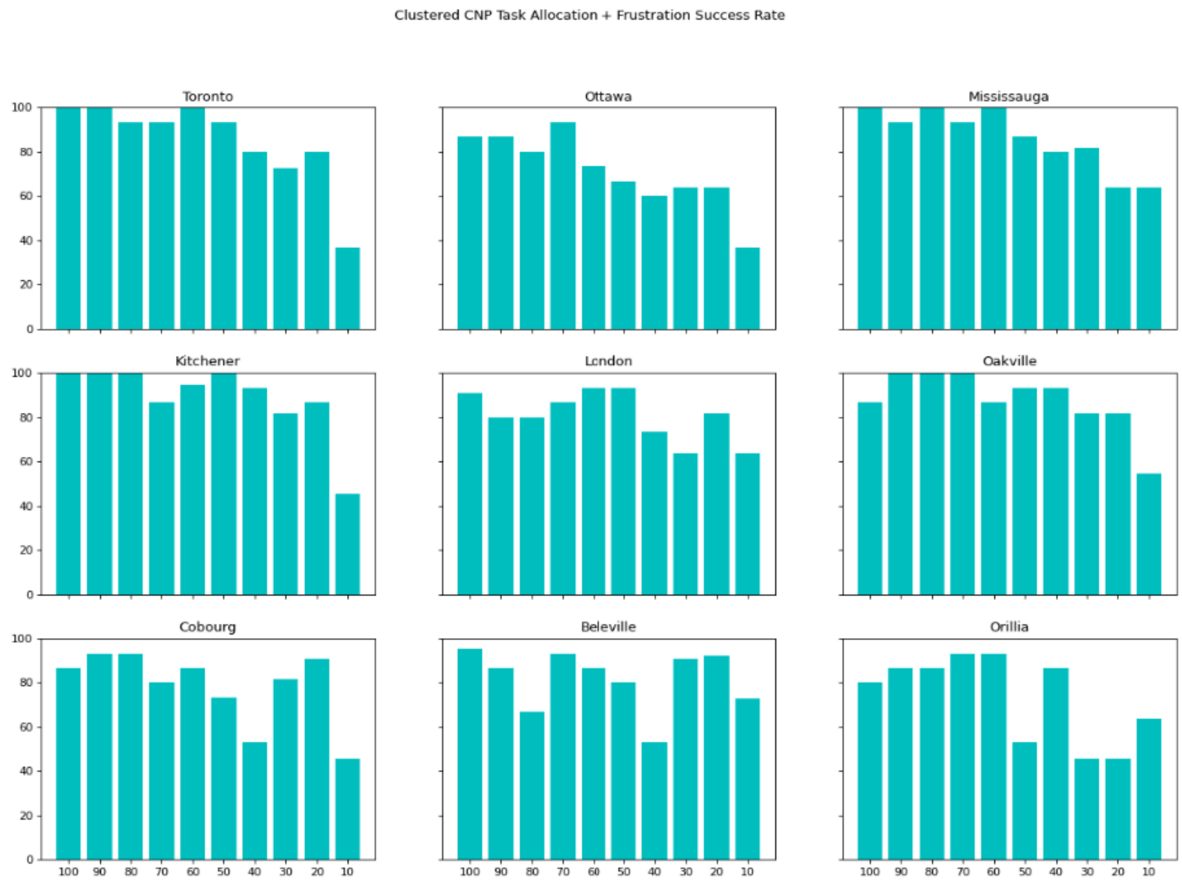

Figure 23: Clustered CNP Task Allocation + Frustration Index Success Rate across all cities with varying number of robots

Limitations

While the proposed solution offers a comprehensive approach to locating a lost person in an urban environment, it comes with several limitations. These limitations may impact the efficiency and reliability of the multi-robot system (MRS) and should be considered when implementing this strategy.

1. Centralized Coordination

The coordination of robot agents relies on a central node for communication and decision-making. This centralized approach creates a single point of failure. If the central node experiences disruptions due to technical issues, network connectivity problems, or cyberattacks, the entire coordination system could be compromised. This risk is heightened in complex urban environments where communication infrastructure might not always be reliable.

2. Weather Conditions

Unpredictable weather can significantly affect the operation of Unmanned Aerial Vehicles (UAVs) within the MRS. Inclement weather such as heavy rain, strong winds, or thunderstorms can disrupt UAV flights, grounding them and reducing the overall efficiency of the search effort. This limitation poses a challenge, especially in regions with volatile weather patterns. Ground-based robots might also face limitations in extreme weather, such as flooding or snow.

3. Resource Constraints

Multi-robot systems require substantial resources, including power, maintenance, and data processing capabilities. Limitations in these resources can restrict the operation duration and scope, affecting the ability to conduct extended searches or cover larger areas. Additionally, the cost of maintaining and operating a MRS might be prohibitive for some organizations or communities.

4. Ethical and Privacy Considerations

Conducting search and rescue operations in urban environments raises ethical and privacy concerns, particularly regarding the collection and use of personal data. Ensuring compliance with privacy regulations, obtaining consent from individuals affected by the search, and minimizing the intrusion into private spaces are essential considerations that may influence the implementation of the proposed solution.

By addressing these limitations, the proposed approach can become more robust, resilient, and adaptable, ultimately increasing the chances of successfully locating lost persons in urban environments. This holistic approach paves the way for future research and applications in non-disaster search operations, contributing to safer and more efficient search and rescue efforts.

Benefits

The proposed solution for locating lost persons in urban environments, utilizing a Multi-Robot System (MRS) and the Lost Person Model, offers multiple benefits, enhancing search efficiency and increasing the likelihood of success.

1. Efficiency

The coordination of multiple robots through a centralized system allows for efficient task distribution and search coverage. This approach reduces overlap and ensures that the robots are searching in areas with the highest probability of finding the lost person. Moreover, the Lost Person Model's probabilistic predictions are key to improving efficiency by guiding robots to focus on key locations, saving time and resources.

2. Scalability

The proposed solution is inherently scalable, allowing for increased search areas or additional robots without a significant decrease in efficiency. This scalability is particularly useful in large urban environments where the search may need to cover extensive areas. The Lost Person Model's flexible data resolution supports this scalability by providing the right level of detail for different scenarios.

3. Adaptability

The adaptability of the solution is a key advantage. The Lost Person Model uses environmental factors and human feedback to adjust search patterns, allowing for rapid response to changing conditions. This adaptability extends to the Multi-Robot System, where human-robot interaction plays a crucial role in refining the search strategy. The system's ability to incorporate new information and adjust the search trajectory ensures that it remains effective even as conditions evolve.

4. Reduced Risk and Safety

The proposed solution minimizes risk by using robots for search operations, reducing the need for humans to enter potentially dangerous or hard-to-reach areas. This approach enhances safety for human rescuers while allowing for thorough exploration of the search area. The centralized communication node further enhances safety by ensuring a coordinated response in case of emergencies.

Future Work

The Clustered CNP (Contract Net Protocol) with Frustration Index Task Allocation Algorithm has shown significant potential in improving search and rescue operations in urban environments. To build upon these findings, the following next steps are proposed:

Further Refinement of the Frustration Index

The introduction of the frustration index has provided a new dimension for task allocation. Future research should explore how different configurations of the frustration index impact search efficiency and success rates. This refinement could involve varying the calculation method, exploring alternative measures of task distribution, or experimenting with different boundaries to create a more balanced assignment of tasks.

Exploration of Advanced Clustering Techniques

The current approach uses the K-Nearest Neighbors method to create clusters. Further research could investigate alternative clustering techniques to enhance the accuracy and flexibility of the clustering process. Techniques such as hierarchical clustering or density-based clustering could offer improved adaptability to different environments and search scenarios.

Development of Adaptive Search Strategies

The adaptability of the Clustered CNP Task Allocation Algorithm is a key strength. To further enhance this adaptability, research should focus on developing dynamic search strategies that can adjust to evolving conditions in real-time. This could involve incorporating machine learning techniques to predict lost person's movements based on previous searches or using real-time data to update search patterns.

Investigation of Scalability and Resource Optimization

Given the scalability of the Clustered CNP approach, future research should explore how to optimize resource allocation as the system scales. This investigation could involve analyzing the computational and communication overhead associated with larger robot fleets, identifying optimal numbers of robots for various search scenarios, and determining the most effective strategies for parallel task allocation.

Evaluation of Ethical and Privacy Considerations

As search and rescue operations become more advanced, ethical and privacy considerations must be addressed. Future research should examine the ethical implications of using robotic agents in urban environments, particularly in terms of data collection, surveillance, and privacy. This evaluation should result in guidelines to ensure compliance with ethical standards and privacy regulations.

Conclusion

This work has explored the Clustered CNP (Contract Net Protocol) Task Allocation Algorithm as a means to optimize search and rescue operations in urban environments. Through extensive simulations and data analysis, several major conclusions were drawn from this research:

Improved Search Efficiency

The Clustered CNP approach, with its area-based task allocation, significantly improves the efficiency of search operations. By allowing robots to search over broader areas, the success rate in locating a lost person has been substantially increased compared to previous methods like the Closest Distance approach.

Enhanced Task Coordination

The introduction of clusters allows for better coordination among search agents. By clustering tasks and assigning robots to search entire areas instead of specific coordinates, the algorithm reduces redundancy and facilitates more complex search trajectories. This coordination ensures that robots cover more ground in a shorter time, increasing the likelihood of successful outcomes.

Versatility and Scalability

The Clustered CNP Task Allocation Algorithm demonstrated versatility in adapting to different urban environments. The ability to scale with varying numbers of robots and search over large or small areas adds to the algorithm's flexibility. This scalability is crucial for managing complex search operations in diverse urban settings.

Introduction of the Frustration Index

The addition of a frustration index in the Clustered CNP Task Allocation with Frustration Index approach offers a novel mechanism for task allocation. This variable, based on the distance from the map's mean to the boundary, provides a dynamic way to assign tasks to robots, helping to minimize commuting time and optimize search efforts. This method also showed a smoother success rate, reducing fluctuations and providing more consistent results across various city scales.

The significance of this research lies in its potential to transform search and rescue operations in urban environments. By introducing a more efficient, coordinated, and scalable task allocation algorithm, this work contributes to the field's broader goal of enhancing search outcomes while minimizing resource use. The approach's adaptability and robustness to failures further underscore its value, suggesting a promising pathway for future search and rescue methodologies.

Moreover, the insights gained from this work offer valuable guidance for real-world applications, particularly in non-disaster scenarios where locating lost persons is critical. The conclusions drawn from this research can inform future studies, promoting a deeper understanding of how Multi-Robot Systems (MRS) can be optimized to support search operations effectively. With the outlined next steps, further exploration and refinement are possible, contributing to a more reliable, ethical, and efficient approach to search and rescue in urban contexts.

Appendix A

The code for this research can be found at the following link:
https://github.com/haighcam/masc/tree/market_tasks

The GitHub Repository is part of the Intellectual Property of the Computer Integrated Manufacturing Lab (CIMLab) at the University of Toronto under the Supervision of Professor Benhabib. All requests to access this repository can be forwarded to the CIMLab.

References

1. A. Gasparetto, P. Boscariol, A. Lanzutti, and R. Vidoni, “Trajectory planning in robotics - researchgate,” Trajectory Planning in Robotics, https://www.researchgate.net/profile/Paolo-Boscariol/publication/257759823_Trajectory_Planning_in_Robotics/links/00b4952fde030402d3000000/Trajectory-Planning-in-Robotics.pdf
2. E. Nunes, M. Manner, H. Mitiche, and M. Gini. A taxonomy for task allocation problems with temporal and ordering constraints. Robotics and Autonomous Systems, 90:55–70, 2017
3. Y. Tian, “Reliable and resource-aware collaborative slam for multirobot search and rescue,” Ph.D. dissertation, Massachusetts Institute of Technology, 2019.
4. I. Maza and A. Ollero, “Multiple uav cooperative searching operation using polygon area decomposition and efficient coverage algorithms,” in Distributed Autonomous Robotic Systems 6. Springer, 2007, pp. 221–230.
5. J. Suarez and R. Murphy, “A survey of animal foraging for directed, persistent search by rescue robotics,” in 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics. IEEE, 2011, pp. 314–320.
6. G.-J. M. Kruijff, I. Kruijff-Korbayova, S. Keshavdas, B. Larochelle, ´ M. Jan´ıcek, F. Colas, M. Liu, F. Pomerleau, R. Siegwart, M. A. ˇ Neerincx et al., “Designing, developing, and deploying systems to support human–robot teams in disaster response,” Advanced Robotics, vol. 28, no. 23, pp. 1547–1570, 2014.
7. T. Jian et al., “Using auction-based Task Allocation Scheme for simulation optimization of Search and rescue in disaster relief,” Simulation Modelling Practice and Theory, https://www.sciencedirect.com/science/article/pii/S1569190X17301934.
8. R. Zlot and A. Stentz, “Market-based multirobot coordination for complex tasks,” The International Journal of Robotics Research, pp. 73–101, 2006.
9. C. Nam and D. Shell, “Assignment algorithms for modeling resource contention and interference in multi-robot task-allocation,” Presentation at IEEE International Conference on Robotics and Automation (ICRA), 2014.

10. C. Nam and D. Shell, “Assignment algorithms for modeling resource contention and interference in multi-robot task-allocation,” Presentation at IEEE International Conference on Robotics and Automation (ICRA), 2014.
11. V. Lattarulo and G. T. Parks, “A preliminary study of a new multi-objective optimization algorithm,” IEEE Congress on Evolutionary Computation (CEC), pp. 1–8, 2012.
12. K. Lerman, C. Jones, A. Galstyan, and M. Mataric, “Analysis of dynamic task allocation in multi-robot systems,” The International Journal of Robotics Research, pp. 225–241, 2006
13. V. Lattarulo, T. Kipouros and G. T. Parks, "Application of the multi-objective Alliance Algorithm to a benchmark aerodynamic optimization problem," 2013 IEEE Congress on Evolutionary Computation, Cancun, Mexico, 2013, pp. 3182-3189, doi: 10.1109/CEC.2013.6557959.
14. Khamis, A., Hussein, A., Elmogy, A. (2015). Multi-robot Task Allocation: A Review of the State-of-the-Art. In: Koubâa, A., Martínez-de Dios, J. (eds) Cooperative Robots and Sensor Networks 2015. Studies in Computational Intelligence, vol 604. Springer, Cham. https://doi.org/10.1007/978-3-319-18299-5_2
15. M. de Longueville, A course in topological combinatorics. Springer, 2012
16. P. Vatsolaki and A. Tsalpatouros, “Ewos: A sealed-bid auction system design and implementation for electricity interconnector capacity allocation,” Information, Intelligence, Systems and Applications (IISA), 2013
17. D. Kim, Y. So, and S. Kim, “Study of marker array list method for augmented reality service based smart home,” International Journal of Smart Home, pp. 51– 64, 2011.
18. Liu, S., Lu, X., Suzuki, M., & Walsh, T. (2023). Mixed Fair Division: A Survey. ArXiv, abs/2306.09564.
19. H. Guo, W. Li, and B. Deng, “A survey on fair allocation of chores,” MDPI, https://www.mdpi.com/2227-7390/11/16/3616 (accessed Jan. 19, 2024).
20. Ariel D. Procaccia. Cake cutting algorithms. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, Handbook of Computational Social Choice, chapter 13, pages 311–329. Cambridge University Press, 2016.

21. Gerkey, B.P., Mataric, M.J.: A formal analysis and taxonomy of task allocation ´ in multi-robot systems. ‘e International journal of robotics research 23(9), 939–954 (2004)

22. Korsah GA, Stentz A, Dias MB. A comprehensive taxonomy for multi-robot task allocation. The International Journal of Robotics Research. 2013;32(12):1495-1512. doi:10.1177/0278364913496484

23. Aziz, H., Pal, A., Pourmiri, A. et al. Task Allocation Using a Team of Robots. Curr Robot Rep 3, 227–238 (2022). https://doi.org/10.1007/s43154-022-00087-4

24. J. C. Spall, “Stochastic optimization,” Handbook of computational statistics. Springer., pp. 173–201, 2012. 48.

25. U. Diwekar, “Optimization under uncertainty,” Introduction to Applied Optimization. Springer., pp. 1–54, 2008.

26. P. J. Shea, K. d Alexander, and J. Peterson, “Group tracking using genetic algorithms,” Proc. of the International Society Information Fusion, 2003.

27. E. G. Jones, M. Dias, and A. Stentz, “Time-extended multi-robot coordination for domains with intra-path constraints,” Autonomous Robots, pp. 41–56, 2011.

28. D. K. Liu and A. K. Kulatunga, “Simultaneous planning and scheduling for multiautonomous vehicles,” Evolutionary Scheduling, Springer, pp. 437–464, 2007.

29. M. Badreldin, A. Hussein, and A. Khamis, “A comparative study between optimization and market-based approaches to multi-robot task allocation,” Advances in Artificial Intelligence, 2013.

30. R. Zlot and A. Stentz, “Market-based multirobot coordination for complex tasks,” The International Journal of Robotics Research, pp. 73–101, 2006.

31. Z. Alibhai, “What is contract net interaction protocol?” IRMS Laboratory, SFU, 2003

32. M. B. Dias, “Trader-bots: a new paradigm for robust and efficient multirobot coordination in dynamic environments,” PhD Thesis, Robotics Institute, Carnegie Mellon University, 2004.

33. B. Coltin and M. Veloso, “Mobile robot task allocation in hybrid wireless sensor networks,” Intelligent Robots and Systems (IROS), pp. 2932–2937, 2010.

34. M. Dias and A. Stentz, “Opportunistic optimization for market-based multirobot control,” International Conference on Intelligent Robots and Systems, pp. 2714– 2720, 2002.
35. M. B. Dias and A. Stentz, “A free market architecture for distributed control of a multirobot system,” International Conference on Intelligent Autonomous Systems, pp. 115–122, 2000.
36. R. M. Zlot, “An auction-based approach to complex task allocation for multirobot teams,” PhD Thesis, Carnegie Mellon University, 2006.
37. R. Zlot, A. Stentz, M. B. Dias, and S. Thayer, “Multi-robot exploration controlled by a market economy,” International Conference on Robotics and Automation (ICRA), 2002.
38. M. Dias, R. Zlot, N. Kalra, and A. Stentz, “Market-based multirobot coordination: A survey and analysis,” Proceedings of the IEEE, pp. 1257–1270, 2006.
39. Koester, R.J. Lost Person Behavior: A Search and Rescue Guide on Where to Look for Land, Air, and Water; dbS Productions: Charlottesville; dbS Productions LLC: Charlottesville, VA, USA, 2008; ISBN 1879471396.
40. Delmerico, J., Mueggler, E., Nitsch, J., Scaramuzza, D.: Active Autonomous Aerial Exploration for Ground Robot Path Planning. IEEE Robot. Autom. Lett. 2, 664–671 (2017). https://doi.org/10. 1109/LRA.2017.2651163
41. Reardon, C., Fink, J.: Air-Ground Robot Team Surveillance of Complex 3D Environments. In: Proc. IEEE Int. Symp. Safety, Security, Rescue Robot. pp. 320–327, Lausanne (2016)
42. OpenStreetMap. OpenStreetMap.org. Retrieved January 19, 2024, from https://www.openstreetmap.org/#map=4/38.01/-95.84
43. Zendai Kashino, Goldie Nejat, and Beno Benhabib. 2019. Multi-UAV-based autonomous wilderness search and rescue using target iso-probability curves. In Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS). IEEE, 636–643.