2024 IEEE 14th Annual Computing and Communication Workshop and Conference (CCWC) | 979-8-3503-6013-4/24/$31.00 ©2024 IEEE | DOI: 10.1109/CCWC60891.2024.10427713

Towards Distributed Learning to Support Situational Awareness for Robotic Team Augmented Humanitarian Disaster Response

Mark Allison
Department Of Computer Science
University of Michigan - Flint
Flint, MI
markalli@umich.edu

Michael Farmer
Department of Computer Science
Kettering University
Flint, MI
mfarmer@kettering.edu

Zheng Song
Department Of Computer Science
University of Michigan - Dearborn
Dearborn, MI
zhesong@umich.edu

***Abstract*—The use of robots to assist first responders in disaster response has seen increasing adoption as underlying technologies mature. They are inherently adept at real-time knowledge acquisition and are able to perform a myriad of pre-stabilization tasks within hazardous environments in lieu of jeopardizing human lives. Heterogeneous robots operating in teams can collaborate to facilitate first responders' decision making by providing *situational awareness* (what is happening and when). However, the *cognitive load* placed on the human must be carefully managed; too much information and the human becomes overwhelmed, too little, and the human becomes over-reliant on the autonomy and hence complacent. Both lead to poor outcomes, with the potential loss of lives. Our work in progress's approach to this shared autonomy problem is to apply distributed machine learning to identify behaviors of interest from temporal changes to LiDar or video images. The overarching goal is a conceptual framework where a heterogeneous team of robots (aerial, ground robots equipped with different sensing, and computational capabilities) may rapidly learn the pertinent aspects of an unfamiliar dynamic terrain, and accordingly, improve the decision making and projection capabilities of human first responders. Essentially, robots working together to collect and disseminate actionable knowledge in the expedited manner to save human lives.**

***Index Terms*—distributed learning, disaster response, human-robot interaction.**

Introduction

Disasters are discrete naturally occurring or man-made events that surpass local resources purposed for responding to, and containing its consequence [1]. While there are no single accepted model for the management of disasters, activities commonly cluster about four phases: prevention, preparedness, response, and recovery [2]. The application of robotics to this domain historically concerns response efforts [3], although there are emergent applications in recovery efforts to reestablish normal operations. As such, we curtail our discourse to robots in response work. *Response activities* typically progresses in the systematic manner with distinct stages that include damage assessment, and search and rescue as priorities. After forty-eight hours the mortality rate peaks; so response is viewed as a race against time, balancing the urgency to reach as many survivors, while being deliberative as not to incur additional risks to responders and victims alike [1]. These activities are data intensive, time constrained and may be unsafe for humans and animals. The domain's requirement characteristics that informs our approach include:

- Incomplete and uncertain information from disparate sources required for decision making from situational awareness.
- An overloaded or damaged communication infrastructure.
- Limited resources with rapidly changing needs.

These characteristics combined with the nature of the work has attracted the attention of the robotics research community to leverage the potential of the domain to save and protect human lives.

One of the first uses of robots in disaster response was in the aftermath of the world trade center bombings, searching for victims of the attack [4]. They are primarily employed in acquiring situational and environmental awareness (reconnaissance and mapping), interacting with survivors and structures, and in support activities such as creating and sustaining a communication backbone. In order to cover the most area under these timing constraints, disaster robotics researchers are exploring the use of large dissimilarly capable and equipped teams of robots working in concert with responders. These heterogeneous robotic teams collaborating with humans has the potential to gather actionable data and provide services far in excess of the same number of machines operating independently. Their ability to keep human responders out of harms way combined with the proficiency in providing specialized capabilities is becoming more recognized and valued. This value is not without its challenges. To realize competency, these teams should autonomously form coalitions to merge individual capabilities, and coordinate tasks allocation towards a central goal. Furthermore, failures, once detected, must be handled in a methodical and systematic manner. The probability that a failure is detected and correctly handled is called *coverage*. The inclusion of coverage for autonomous multi-robot teams in crisis response is critical to mission success yet remains largely unexplored. This work specifically addresses

979-8-3503-6013-4/24/$31.00 ©2024 IEEE

dissimilarly equipped terrestrial and aerial robots within the response phase, focusing on environmental mapping, and identifying behaviors of operational interest as primary thrusts.

In our context, behaviors that would foster situational awareness are those about human and animal movement within a search space or structural safety concerns such as slight movements in the support system of a building that may signal collapse.

Our overarching goal is autonomy assisted crisis response using a peer-to-peer distributed learning approach under coverage constraints. For a large team of robots, we need to optimize coverage and task assignments within the search space. The distributed learning paradigm uses a large number of nodes coordinating to train a model, orchestrated by a central server. While the training data is decentralized, a solitary central server introduces a single point of failure. Additionally, the communication from the nodes to the central server incurs high costs. Distributed learning has a high potential to congeal knowledge for human situational awareness, and to automate task allocation, however its low tolerance to failure and high communication overhead would render it infeasible for the demands of a typical crisis scenario. Recently there have been approaches to decentralize the global model by offloading to a random subset of nodes [5] or to distribute over a network/graph where node only communicate with one-hop neighbors [6]. In our approach, robots are (re)dispatched from a base that serves to replenish power. We address the central server concern by utilizing these bases as computational resources. Decentralization occurs by using dynamic selection of bases, based on evolving capabilities (communication range, power reserve, location) to support global knowledge distribution. The realization of optimal subset network approach may be impactful to several other domains beyond crisis response.

The remainder of this paper is organized as follows. In the next section we present an indicative scenario to contextualize the approach. Next we introduce the approach itself in section II, then conclude.

I. Indicative Scenario

Consider a situation in the aftermath of an earthquake in a densely populated city. The initial damage that has occurred is subject to a cascading effect as the systems upon which the city depends (utilities, transportation, communication, and residential and commercial buildings) are interconnected [7]. This has led to maps and way-points of the infrastructure to be rendered useless. As is typical, the situation is highly dynamic, hazardous and evolving. The primary focus is damage assessment and search and rescue; success is predicated on minimizing the time taken to accomplish sub-tasks.

Several search and rescue teams comprising of human,animals and robots will deployed to the disaster area, also known as the *hot zone*. Their initial task is to assess the new landscape with the forefront concern to save and preserve human life. Search and rescue activities are performed along with damage assessment. The robots will need to assess structural integrity and debris, identifying areas of risk and safe routes for humans. As such there is a necessity to ascertain ground truths as to risks to humans to begin, and while, conducting the work. To exacerbate the challenge we can assume communication has been disrupted and alternative methods will need to be employed.

As such, a team of dissimilarly equipped robots (heterogeneous) is first deployed to search for survivors by detecting biological motion and collect situational data via SLAM to determine safe paths for first responders incursion. Each robot gathers data, incrementally refining the model and map until their power resources are near depletion. They then return to a base to be recharged, their data downloaded, and they are re-dispatched. To get the most pertinent information as soon as possible, the teams should be first sent to high-value locations that are most likely to be risky. A Deep learning model will be employed, to predict the risk at each location overlayed with probable survivors. The results will inform subsequent redeployments for confirmation using alternate or more precise sensing.

After collecting sensing data at a new location, a team needs to decide whether or not to start the model re-training procedure. Although updating the machine learning model can improve situation awareness for the robotics team, the model training procedure is time and resource intensive. The decision should be made based on several factors, including the new collected sensing data, how the current model's prediction results match the collected ground truth, the available resources and capability of the team, and the the probability of meeting another robot for exchanging updated the local model. The problem is non-trivial as the disaster response teams may not be able to communicate with each other and each team needs to make decisions towards realizing global optimization based on their limited local information.

II. Approach

A. Distributed Learning of Temporal Changes to Images

In the context of machine learning to support situational awareness for disaster response, we seek to learn the safe paths for human responders and the existence of biological movement (survivors) based on direct sensing and deriving knowledge from temporal changes to SLAM liDar or video images. For example, if a sequence of image shows rapid changes in a small area for a relatively smooth object, this may indicate the presence of a human. Similarly slower changes observed in larger objects may indicate imminent collapse. Since the movement of biological and structural elements are non-exact, it is well suited to be framed as an semi-supervised classification problem. Discerning between biological (human and animal) and object movement is based on the biological systems typically having more degrees of freedom therefore exhibiting more complex patterns of motion [9].

The distributed learning of what image changes mean depends on refining models across multiple robots and their bases as nodes. In a large multi-dimensional data-set such as that in the scenario, we necessarily require this parallelized approach to overcome the computational limitations of the size of the dateset.

Framing the Problem: If we divide the area A in a grid pattern (MxN) the approach is to: (1) dispatch robots minimizing redundancy to tile m, n where $m \in M$ and $n \in N$; (2) optimize the correct classification of temporal anomalies as α, biological , β structurally relevant shifts, or other; and (3), facilitate situational awareness by complete a coverage map of A in a manner that tiles are notated by probabilities of the existence of α and β . In a naive manner we may structure the problem at (2) as one of logistical regression to binary classify α the existence of biological movement or not.

Data Partitioning : Each tile is assigned to a different robot as a computing node. Let D_{ij} represent the dataset on tile i, j.

Local Model Updates: Each robot will independently computes local updates to the model parameters based on its tile anomaly presence. In some occasions the robot may call for assistance from another (see [10] for a more granular treatment of assistive behavior) In this case the sub-team will collaborate to ensure coverage. To compute gradients for the log-likelihood:

$$-1/r_i \sum_{j=1}^{r_i} (y_{ij} - \rho_{ij}) X_{ij} \tag{1}$$

Where: r_i is the number of instances on robot i; y_{ij} is the true label of instance j on robot i; ρ_{ij} is the predicted probability of instance j on node i; and, X_{ij} is the feature vector.

Communication and Aggregation: Robots communicate their local updates to their bases that act as a central coordinator or parameter server. This involves aggregating the local updates to update the global model parameters.

B. Assuring Coverage

Within a particular tile, especially in a disaster response scenario, a particular terrestrial robot may not have the locomotion capability (wheeled, tracked, legged). When this occurs blindspots may occur in the LiDar image due to an obstruction as in Figure 1 In this event a robot will summon assistance to complete coverage as in Figure 2. In our simulation we demonstrate the feasibility of Lidar coverage through coordinated relocation as in Figure 3.

Once coverage is achieved the subsequent task is to effectively merge sensor readings.

Figure 4 provides the processing schema for our multi-robot LiDAR based situational awareness strategy. The workflow begins with locally optimized `RANSAC` processing [11] to detect lines for constructing landmarks. The detected landmarks are compared with existing previously detected landmarks through the Association Processing. The Landmark Tracking establishes tracks and feeds these into the SLAM processing which integrates the Robot IMU data with the track positions of the landmarks.

The LiDAR point clouds are also passed to the `Blob Detection` after the RANSAC removes the larger landmark feature points to detect moving targets. The core of the work associated with identification of biological and structural behaviors occurs in the `moving object association` and `moving object tracking`. This is implemented by modifying the python based *Motion Tracker Beta* [12]. Although this platform is purposed for video, we have had some promising initial results with Lidar temporal differentiation.

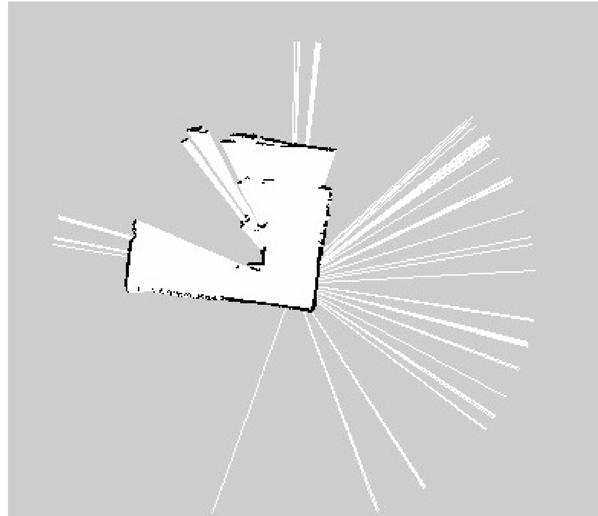

Fig. 1. Lidar image showing blind spot

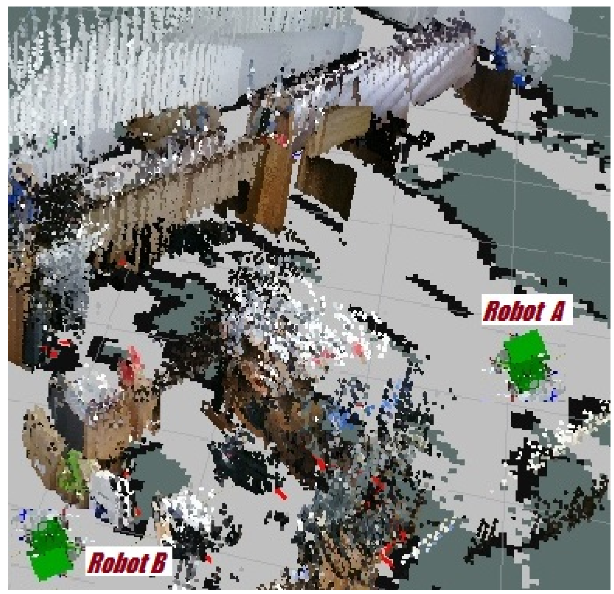

Fig. 2. SLAM

Another class of moving objects that is of key importance are fellow robots. They identify each other through their tracks and in the Multi-Robot fusion the maps of the various robots are integrated into a single map that is in a centralized coordinate system rather than the individual robot-centric systems. As discussed in subsection II-B, robots will coordinate relocation to assure coverage at this juncture. Both the landmarks tracking and the object tracking are performed via Kalman filters. As the RANSAC processing converts the

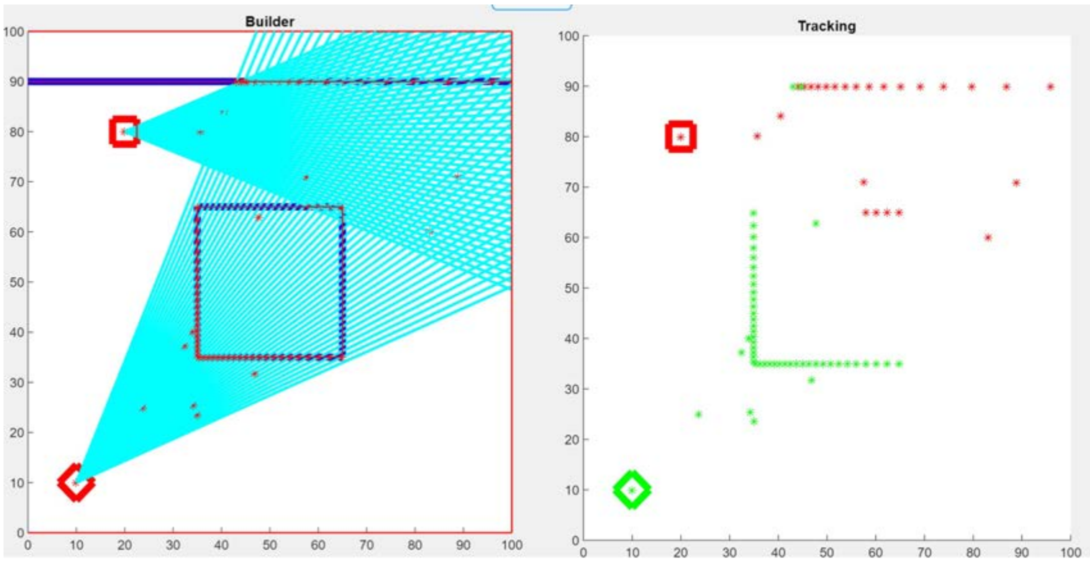

Fig. 3. Simulation Platform showing two robots coordinating coverage by relocation

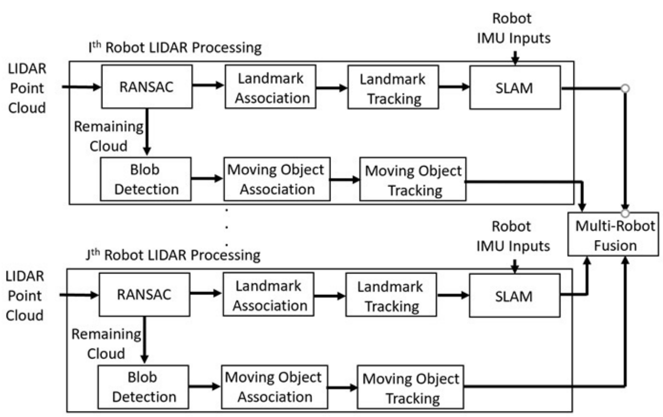

Fig. 4. Processing for Multi-robot LiDAR

r-theta information into x-y linear objects a traditional Kalman filter can employed for the landmarks. The blob detections are tracked using the Extended Kalman filter since their detections remain in r-theta space.

C. Survivor Social Aspects - Framing

Within search and rescue, survivors more likely will be cognitively distressed due to injury and lack of contact with the outside world. The social interaction with a rescuing robot should be carefully managed as these dependents may experience scare and isolation. As such social role framing as a structure of expectation is required to ensure the appearance of support, companionship and calming [13].

D. Risk

We introduce a very simple model for predicting the risk of a location, which can be specified as $(x, y, z, \vec{G}, \vec{L}) \rightarrow R$, where (x, y, z) denotes the longitude, latitude, and altitude of the location, $\vec{G}$ denotes gas transmission pipelines, and $\vec{L}$ denotes hazardous liquid transmission pipelines. When a team finishes assessing the risk at its current location, it runs the risk prediction model, calculates the risks of nearby locations, and select one with the highest risk as it next target. As the model is developed for general purpose and cannot fit exactly into the city after the earthquake, a team t_1 will use its risk assessment results to re-train the model by updating its parameters w_1. When t_1 meets with another team t_2 whose updated parameters are denoted as w_2, the two teams exchange their parameters and use the federated average algorithm [8] to calculate a new optimal parameter set $w' = FedAvg(w_1, w_2)$. The procedure is shown by Figure 5.

III. Conclusion

In this paper we introduced our approach to foster situational awareness in first responders using a distributed learning methodology with sensor fusion. We have presented some preliminary simulation results performed on real lidar equipped robots towards proof of principle. There are several critical aspects that we intend to address or are being investigated in parallel.

IV. Acknowledgments

This work is being supported by the National Science Foundation under award #2220513.

References

[1] R. R. Murphy, S. Tadokoro, and A. Kleiner, "Disaster robotics," in *Springer handbook of robotics*. Springer, 2016, pp. 1577–1604.

[2] D. A. McEntire, *Disaster response and recovery: strategies and tactics for resilience*. John Wiley & Sons, 2021.

[3] S. Park, Y. Oh, and D. Hong, "Disaster response and recovery from the perspective of robotics," *International Journal of Precision Engineering and Manufacturing*, vol. 18, no. 10, pp. 1475–1482, 2017.

[4] R. G. Snyder, "Robots assist in search and rescue efforts at wtc," *IEEE Robotics and Automation Magazine*, vol. 8, no. 4, pp. 26–28, 2001.

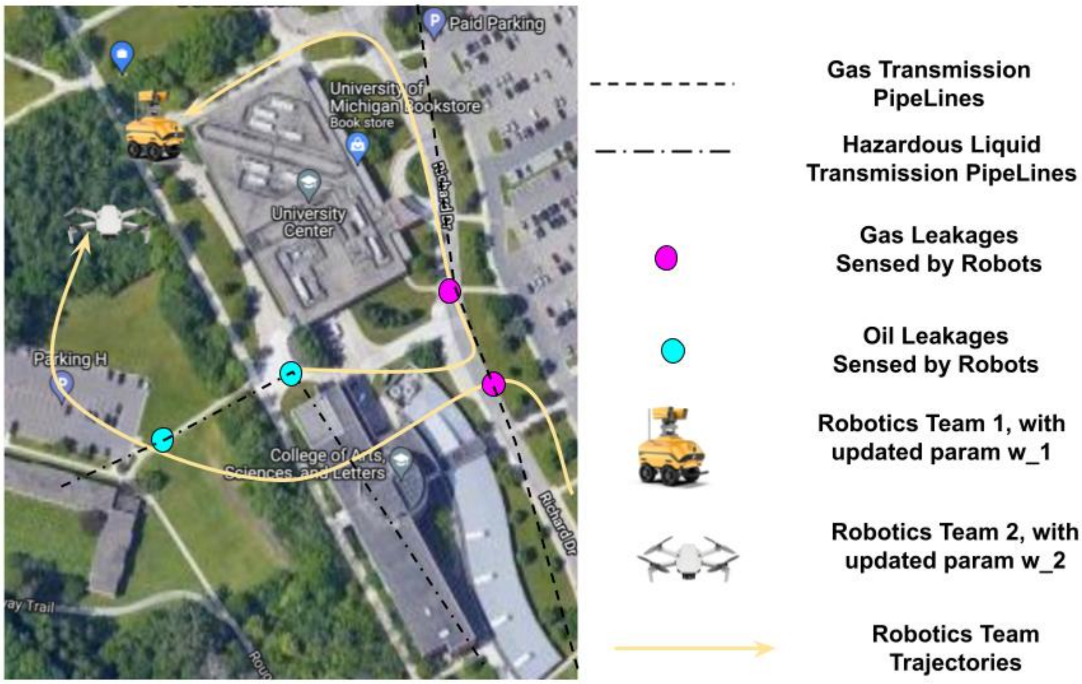

Fig. 5. Risk Prediction and P2P Federated Learning-based Model Update

[5] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.

[6] A. Lalitha, O. C. Kilinc, T. Javidi, and F. Koushanfar, "Peer-to-peer federated learning on graphs," *arXiv preprint arXiv:1901.11173*, 2019.

[7] P. Price, "Depends on what you mean by "disaster"," *Mississippi Libraries*, vol. 70, no. 3, pp. 56–57, 2006.

[8] J. Mills, J. Hu, and G. Min, "Communication-efficient federated learning for wireless edge intelligence in iot," *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 5986–5994, 2019.

[9] A. Chouchourelou, A. Golden, M. Shiffrar, and A. Chouchourelou, "What does "biological motion"really mean? differentiating visual percepts of human, animal, and nonbiological motions," *People watching: Social, perceptual, and neurophysiological studies of body perception*, pp. 63–81, 2013.

[10] M. Allison, M. Spradling, and N. Knock, "Uav collaborative search using probabilistic finite state machines," in *International Command and Control Research and Technology Symposium-Knowledge Systems for Coalition Operations (ICCRTS-KSCO 2017)*, 2017.

[11] O. Chum, J. Matas, and J. Kittler, "Locally optimized ransac," in *Pattern Recognition: 25th DAGM Symposium, Magdeburg, Germany, September 10-12, 2003. Proceedings 25*. Springer, 2003, pp. 236–243.

[12] K. Floch and A. Kossa, "Motion tracker beta: A gui based open-source motion tracking application," *SoftwareX*, vol. 23, p. 101424, 2023.

[13] V. Groom, V. Srinivasan, C. L. Bethel, R. Murphy, L. Dole, and C. Nass, "Responses to robot social roles and social role framing," in *2011 International Conference on Collaboration Technologies and Systems (CTS)*. IEEE, 2011, pp. 194–203.