

# Distributed Intelligence for Multi-Agent Systems in Search and Rescue

Chinmaya Patnayak

Thesis submitted to the Faculty of the  
Virginia Polytechnic Institute and State University  
in partial fulfillment of the requirements for the degree of

Master of Science  
in  
Computer Engineering

Ryan K. Williams, Chair

Haibo Zeng

Amos L. Abbott

October 9, 2020

Blacksburg, Virginia

Keywords: Distributed Computing, Multi-Agent Systems, Search and Rescue, Unmanned

Aerial Vehicles, Deep Learning, Inference

Copyright 2020, Chinmaya Patnayak

# Distributed Intelligence for Multi-Agent Systems in Search and Rescue

Chinmaya Patnayak

(ABSTRACT)

Unfavorable environmental and (or) human displacement may engender the need for Search and Rescue (SAR). Challenges such as inaccessibility, large search areas, and heavy reliance on available responder count, limited equipment and training makes SAR a challenging problem. Additionally, SAR operations also pose significant risk to involved responders. This opens a remarkable opportunity for robotic systems to assist and augment human understanding of the harsh environments. A large body of work exists on the introduction of ground and aerial robots in visual and temporal inspection of search areas with varying levels of autonomy. Unfortunately, limited autonomy is the norm in such systems, due to the limitations presented by on-board UAV resources and networking capabilities.

In this work we propose a new multi-agent approach to SAR and introduce a wearable compute cluster in the form factor of a backpack. The backpack allows offloading compute intensive tasks such as Lost Person Behavior Modelling, Path Planning and Deep Neural Network based computer vision applications away from the UAVs and offers significantly high performance computers to execute them. The backpack also provides for a strong networking backbone and task orchestrators which allow for enhanced coordination and resource sharing among all the agents in the system. On the basis of our benchmarking experiments, we observe that the backpack can significantly boost capabilities and success in modern SAR responses.

# Distributed Intelligence for Multi-Agent Systems in Search and Rescue

Chinmaya Patnayak

(GENERAL AUDIENCE ABSTRACT)

Unfavorable environmental and (or) human displacement may engender the need for Search and Rescue (SAR). Challenges such as inaccessibility, large search areas, and heavy reliance on available responder count, limited equipment and training makes SAR a challenging problem. Additionally, SAR operations also pose significant risk to involved responders. This opens a remarkable opportunity for robotic systems to assist and augment human understanding of the harsh environments. A large body of work exists on the introduction of ground and aerial robots in visual and temporal inspection of search areas with varying levels of autonomy. Unfortunately, limited autonomy is the norm in such systems, due to the limitations presented by on-board UAV resources and networking capabilities.

In this work we propose a new multi-agent approach to SAR and introduce a wearable compute cluster in the form factor of a backpack. The backpack allows offloading compute intensive tasks such as Lost Person Behavior Modelling, Path Planning and Deep Neural Network based computer vision applications away from the UAVs and offers significantly high performance computers to execute them. The backpack also provides for a strong networking backbone and task orchestrators which allow for enhanced coordination and resource sharing among all the agents in the system. On the basis of our benchmarking experiments, we observe that the backpack can significantly boost capabilities and success in modern SAR responses.

# Dedication

*To Late Bibekanand Patnaik, for being a guiding light, for your service to the nation and science.*



# Acknowledgments

I would like to express my sincere gratitude to my advisor, Dr. Ryan K. Williams for his thorough guidance, support and simply being an inspiration. Multi-agent systems represents a very small part of all what I could learn from him. To my mentor Dr. James McClure, I owe all success in this project, for without his immense knowledge, patience, enthusiasm and mutual love for computers, this document would have described a brick.

To my committee, Dr. Haibo Zeng and Dr. Amos L. Abbott, I am extremely grateful for your constant support and encouragement throughout my program. I am fortunate to have this opportunity to learn from you.

This work represents efforts from Dr. Nathan Lau and Dr. Nicole Abaid's group at Virginia Tech, Black Diamond Search and Rescue, Virginia Department of Emergency Management and Bob Koester, Ph.D., to whom I owe a great debt for their knowledge, resources and support.

I thank my fellow lab mates, Larkin Heintzman, Kevin Smith, Pratik Mukerjee, Akshay Kumar, Amanda Hashimoto, and Tianzi Wang for the stimulating discussions, and for all the fun we had for the duration of the project.

Last but not the least, I would like to thank my family and friends, especially Dhruv and Teeru, for being with me through good times and odd (Data Labelling and Logging, mostly). Thank you Virginia Tech, for this rewarding journey.

# Contents

<b>List of Figures</b>	<b>ix</b>
<b>List of Tables</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Search and Rescue . . . . .	1
1.1.1 Teams and Operations . . . . .	2
1.1.2 Challenges . . . . .	3
1.2 A Significant Use case for Robotics . . . . .	5
1.2.1 Related Work . . . . .	6
1.2.2 Challenges in SAR with UAVs . . . . .	8
1.2.3 A Strategy for Integration of Robots in SAR . . . . .	10
1.3 Summary of Contributions . . . . .	11
1.4 Thesis Outline . . . . .	12
<b>2 Distributed Intelligence for Multi Agent Robotic Systems</b>	<b>14</b>
2.1 Evolution of Computing Systems . . . . .	14
2.1.1 Software Systems . . . . .	15
2.1.2 Computational Hardware . . . . .	18

2.1.3	Networking and the Cloud . . . . .	19
2.2	Emergence of Distributed Systems and Intelligence . . . . .	21
<b>3</b>	<b>A Multi-Agent Approach to SAR</b>	<b>24</b>
3.1	Computational Objectives for Multi-Agent Systems in SAR . . . . .	24
3.2	A New Multi-Agent Approach to SAR . . . . .	26
3.2.1	Agents . . . . .	27
3.2.2	Interactions . . . . .	29
<b>4</b>	<b>WASP: A Surface for Edge Intelligence</b>	<b>33</b>
4.1	Physical Specification and Design Decisions . . . . .	33
4.1.1	Physical Specifications . . . . .	33
4.1.2	Design Objectives . . . . .	37
4.1.3	Design Process, Integration and Iterations . . . . .	38
4.2	Software Infrastructure and Setup . . . . .	41
4.2.1	Containerization . . . . .	41
4.2.2	Orchestration . . . . .	42
4.2.3	Storage . . . . .	44
<b>5</b>	<b>System Benchmarking and Discussions</b>	<b>46</b>
5.1	Computation . . . . .	46
5.2	Networking and Communications . . . . .	50

5.3	Power and Battery . . . . .	52
5.4	System Integration and Field Studies . . . . .	53
<b>6</b>	<b>Future Work</b>	<b>55</b>
6.1	Realization of Multi-Agent SAR . . . . .	55
6.1.1	Hardware Integration and Platform Engineering . . . . .	57
6.1.2	Human Computer Interaction . . . . .	58
6.2	Mobile Compute Cluster for Other Use Cases . . . . .	58
<b>7</b>	<b>Summary</b>	<b>60</b>
	<b>Bibliography</b>	<b>62</b>

# List of Figures

1.1	Robotic Agents in SAR	8
1.2	DJI Matrice 600: Payload vs Flight Time	10
2.1	Microsoft's AirSim Simulator with multi-modal visualizations	17
3.1	Custom Drone Configuration for DJI Matrice 600	28
3.2	WASP (Wearable Mobile Computing Cluster)	29
3.3	Agents and Interactions	30
3.4	Proposed GUI with interactive agents, sensor and agent info	32
4.1	Computing Platforms for the MA-SAR System	34
4.2	Mophie Powerstation AC (24000 mAh)	35
4.3	Netgear Nighthawk XS6 AC4000	37
4.4	Physical Backpack Assembly and Electrical and Networking Interconnects	39
4.5	3D Printed Custom Mounts	40
4.6	Docker Container Architecture for SAR	42
4.7	Kubernetes System Architecture for SAR	43
5.1	Detection and Identification in Computer Vision	47
5.2	Deep Learning Inference Benchmarking	48

5.3	Distributed Deep Learning Inference Benchmarking . . . . .	49
5.4	Visualizing GIS Maps and Sector Assignments . . . . .	51
5.5	Bandwidth Tests over communication channels across nodes . . . . .	52
5.6	Mobility Testing with WASP . . . . .	54

# List of Tables

4.1	Key Specifications for Processor Kits. . . . .	34
4.2	Key Specifications for Battery Packs. . . . .	35
4.3	Key Specifications for Networking Router. . . . .	36
5.1	Power Benchmarking for Electrical Loads. . . . .	53

# List of Abbreviations

AI Artificial Intelligence

CPU Central Processing Unit

DNN Deep Neural Networks

GPU Graphic Processing Unit

HITL Human In The Loop

MAR Multi-Agent Robotics

SAR Search and Rescue

UAV Unmanned Aerial Vehicles

WASP WearAble Supercomputing Platform



# Chapter 1

## Introduction

### 1.1 Search and Rescue

At the time of compiling this thesis, in the backdrop of COVID-19, devastating wildfires engulfed much of the three states of Oregon, Washington, and California in the United States. Over 3 dozen lives were lost across the three states and 22 people were recorded missing just in Oregon, with casualties expected to increase till the fire subsides [1]. A whopping 3.6 million acres were charred across the state of California. Not only wilderness but towns and suburbs were affected, destroying buildings and public infrastructure. Dangerous smoke, the sheer expanse, and access to areas did not just render human SAR missions futile, but nearly impossible. More than 20,000 firefighters continue to engage with the fires.

87,438 active cases of missing people in the United States were reported as of 31st December, 2019 [2]. Among the states, Alaska, as per National Missing and Unidentified Persons System statistics, reported the highest fraction of missing people with respect to the population at 41.8 cases for every 10,000 people. The fraction stands more than double than the rest of the US and could be largely attributed to people occupying harsh natural environments. The 2017 Annual SAR Report by National Park Service (NPS) of United States highlighted 4,194 SAR incidents in the National Parks alone with a mere 41% saves in such missions [3].

Unfavourable changes in environment, either with respect to human(s) subjects or itself, in an

usually short span of time, engenders the requirement for a SAR response. Such unfavourable changes could be associated with either natural or man-made events and disasters demanding search and rescue. Wildfires, avalanches, landslides and floods have a high likelihood of displacing people from the usual in a short period of time, or severely curtailing mobility in dangerous environments. Voluntary human movement for purposes including recreation or non-urban habitation presents another significant fraction of the SAR incidents. State and community-sponsored programs supported by volunteers have helped extensive creation of artificial trails. As a result, more people have found themselves engaging in outdoor recreation in and around state parks, natural reserves and mountains. This, unfortunately has opened up more opportunities for people to venture into unmapped or unmarked areas where people tend to lose their tracks or get disassociated from groups [4]. Immobilization due to medical emergencies also factors in while accounting for numbers. War, acts of terrorism and collapsing infrastructure also substantially add numbers to the count globally.

### 1.1.1 Teams and Operations

The scope of SAR in such responses is not limited to spotting the victim and delivery, but may also extend to mapping, observation, communication relay and public awareness. SAR teams, in its most common form, assumes a structure or team mainly consisting of trained volunteers and/or agency level responders from local, state or federal bodies. Such teams could be often supported by specialised and skilled human resources like mountaineers, climbers and the military. Trained animals such as Dogs (K9 teams) also form an essential part of some teams and have shown a noteworthy impact in success of SAR missions. SAR responses might often include over 200 human responders in the operation with responsibilities that might be medical, exploratory and/or tooling in nature. The expenditure on such missions could rise to as high as tens of thousand US dollars.

Global organizations such as International Search and Rescue Advisory Group (INSARAG) have been responsible for shaping guidelines, methodologies and minimum international standards for SAR response in urban settings and training SAR responders for about 30 years [5]. In the United States, the American Society for Testing and Materials (ASTM) International is responsible for developing the guidelines for SAR. Committee F32 in the ASTM International specifically focuses on developing standards pertaining to equipment, testing, maintenance, SAR management and operations, personnel, training and education [6]. Other notable non-governmental organizations such as the National Association for Search and Rescue (NASAR), Explorer Search and Rescue (ESAR), Mountain Rescue Association (MRA) are engaged in supporting SAR in the US alongside other state and community level organizations. Each SAR incident might present a different set of requirements or structure in their response. The location of SAR responses defines the scope for many of those. Hence, the ASTM International classifies SAR efforts in four self-explanatory categories based on the location of their origin:

- Structural Collapse or Urban SAR
- Waterborne SAR
- Wilderness SAR
- Aeronautical SAR

### 1.1.2 Challenges

Despite efforts to keep the community well informed through resources, guides and a strong SAR organizational body and execution capabilities, modern-day search and rescue operations face some uphill challenges. The most prominent of those could be identified as:

- **Inaccessibility**

SAR missions originating from incidents such as natural disasters or war might render sites dangerous for the rescue personnel or equipment to access. Natural environments like wilderness and mountains innately make hard to access regions limiting mobility. Not just wilderness or natural habitats, even urban search and rescue by causes such as infrastructure collapse and fire, might significantly cripple level of access. Unusually large search spans might also be challenging to overcome with limited responder headcount.

- **Connectivity and Limited Information**

A general heavy reliance on abundant information via the internet is not an affordance in wildness or disaster struck settings as most of such areas are disconnected due to geographic profiles and networking barriers. This affects both, the victims and the rescue teams. Forested lands and mountains cause heavy attenuation of electromagnetic signals and restrict possible channels of communication.

- **Sensing**

Capturing vital cues are integral to the success of any SAR mission. Limited sensory capabilities, either qualitative or quantitative, of human responders in terms of vision, smell or touch might also result in missed opportunities in tracking lost person. The scope and variations in SAR incidents eliminates the possibility of one-modality-to-solve-it-all philosophy and might simultaneously expect various sensory cues such as vision in multiple wavelengths, wind, moisture, smell etc. to drive operational success.

- **Resources and Equipment**

Volunteers are a vital cog in the SAR response as they provide the head count to scan larger areas in a shorter span of time. Availability and organization of volunteers during a SAR response may vary quite significantly and poses a significant challenge.

In 2020, COVID-19 saw a massive surge in trail hikers and a significant decline in number of volunteers [7]. The contagious nature of the virus required special measures and equipment to reduce volunteers' exposure to risks. Equipping the responders with adequate gear and training to ensure their safety from injuries also opens up a challenge for the SAR bodies. [8]

- **Time to rescue**

Time is of significant value in SAR missions and the chances of success of a mission reduce exponentially with the passing hour. For every hour elapsed during a SAR response, the search radius is estimated to increase by another 3 km [9]. Movement in unwieldy terrain, managing and organizing communication and equipment can take away crucial periods of time initiating or during the SAR missions.

## 1.2 A Significant Use case for Robotics

Robots have remarkably extended human capabilities and augmented understanding in a diverse set of tasks ranging from nano-robots for medical diagnosis to rovers for space exploration. By the virtue of their aerial and ground agility, integration of robots might significantly enhance the range and level of access in a situation like response for a SAR Missions. Notably, Unmanned Aerial Vehicles (UAVs), by extending human capabilities through the aerial media, has captured remarkable interest with their impact on fields such as transportation, surveillance, defense, agriculture, disaster management and photography. Given a significantly unhindered media and the ability to sense and scan large surface areas does establishes UAVs as an able ally in SAR missions. With a growing UAV market and manufacturers, acquisition and training costs are expected to decline. A wide range of sensors such as cameras working on a diverse range of wavelengths not only aids the human

perception but also allows the ability to leverage latent information for processing.

### 1.2.1 Related Work

Given the diversity of challenges, tasks and demands associated with SAR, remarkable progress in allied fields have significantly tricked down to a SAR. Developments in integration of robots for SAR has historically spanned four overlapping themes which include (1) evaluation of capabilities of robots in SAR like operations, (2) effectiveness of algorithms such as tracking, path planning and computer vision in SAR, (3) efficiencies in offloading tasks from on-board computers and (4) multi-agent collaboration and interfacing.

Initial body of research in robots for SAR were shaped by development of mobile agents with varying levels of autonomy, primarily for applications in military, and identification of SAR as an application [10, 11, 12]. VGTV-Xtreme, an early ground rover used for SAR applications during Hurricane Katrina in 2005, was developed as a variant to a rover for inspecting air-conditioning ducts [13]. With growing capabilities of robotic systems in early 2000s, integration with human teams in terms of coordination and collaboration was carefully evaluated to create a road-map for future by listing vital challenges for human-robot systems [14]. Early simulation studies in Urban SAR, surprisingly presented negative impact on efficiencies with multiple ground and aerial robotic agents with human operators [15]. With growing prowess and intelligence in robots, researchers increasingly reported developments which swiftly overcame the challenges of integrating robots in SAR [16].

The act of initial or dynamic task planning and designation is often the most formidable challenge in search and rescue [17]. The ability to quickly generate a task order for SAR can have a significant impact on the success of the operation. Robotic and computer systems tend to be of high utility in presenting recommendations by crunching historical and situa-

tional data to generate optimised trajectories and search routes without manual inputs [18]. Localising and tracking of a mobile target with multiple agents turns out to be an NP-hard problem dependent exponentially with respect to number of agents. Various optimization approaches have brought down the complexity in reaching sub-optimal solutions [19] [20]. The tasks defined for aerial agents may also be dependent on trigger events. Work in dynamic work loads for aerial agents and optimization of task delegation strategies with focus on connectivity and coverage has been evaluated in approaches such as Multi-Objective Path Planning (MOPP) algorithm [21].

The capability to carry payloads in the form of visual sensors such as camera to hard to access areas extends their in-field utility. Extensive usage and subsequent advantages of cameras in different spectra in SAR have been demonstrated in past work [22]. Powerful embedded computers opened a massive opportunity in presenting latent yet significant information to the humans. This could be information seeking attention, an alternate visual representation to help the human responder understand and evaluate the environment better than relying on his own affordances. Pioneering developments in utilization of UAVs for semi-autonomous SAR were made by Goodrich et al. where they argue that the utilization of camera-equipped UAVs improved the probability on finding the missing person and accomplishing the goal in a shorter span of time [23] [24]. Agcayazi et. al. demonstrated the utilization of an on-board computer to run real-time anomaly detection algorithms in a semi-autonomous drone setting [25].

The successful deployment of robotic agents gradually transitioned into a more recent trend in utilization of robotic agents in SAR in multi-agent coordination and collaboration settings. Significant work has been done in the area of human-robot collaboration in SAR involving semi-autonomous agents such as rovers, microcopters and UAVs collaborating with off-site command post personnel and on-site human responders [14] [26] [27] [28]. Distribution of

responsibilities and tasks in SAR and Human Robot Interaction and Interfacing [29] has also been extensively evaluated in terms of feasibility and impact. Modern technology interfaces such as Augmented Reality and Virtual Reality are expected to shape the future of human robot interactions in SAR [30] [31].

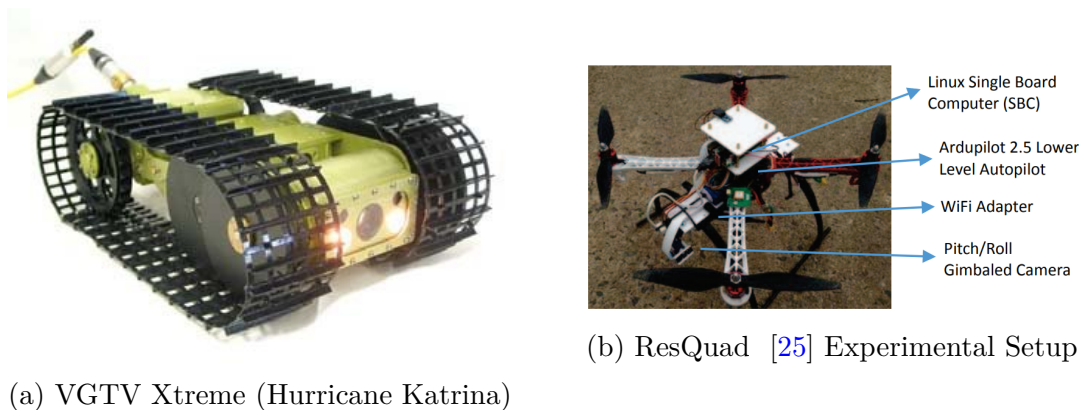


Figure 1.1: Robotic Agents in SAR

DARIUS and ICARUS Project presents one of the more recent architectures and platform of unmanned systems in SAR [32] integrating the ground station and first responders with semi-autonomous agents in a comprehensive manner with actual SAR deployments. Integration and testing of multiple pieces for SAR missions has also been presented with fixed wing aircraft with modifications to off-the shelf action cameras were made to test and benchmark target identification algorithms [22].

### 1.2.2 Challenges in SAR with UAVs

Despite sufficient evidence of UAV's capability to augment Human-In-The-Loop (HITL) SAR missions, swarms continue to battle issues with efficiency, practicality and cost effectiveness in modern implementations and prototypes. To achieve higher levels of autonomy in such a system, we expect the system to host a larger range of sensors for perception and higher computing capabilities to enable a sufficient level of understanding and decision making for



the agents. Considering the current state of research in multi-agent robotic systems in SAR and commercial grade UAV options, we identify certain challenges and trade-offs that limit the viability of such an approach:

- **Flight Times (Time vs Payload)**

Electrical motors on the UAVs work to propel and hover the UAVs in the air against gravity and air resistance. Installing more add-ons to the UAVs result in higher gross weight and surface area that results in higher inertia for movement. The higher inertia results in reduced flight time for the UAVs. As is, the most advanced and powerful commercial drones available in the market present a meagre flight time of 30 - 40 minutes of hover time. A further reduction in hover time would make UAVs impractical for deployments. Fig 1.2 depicts a significant drop in hover time with every added kilogram of payload.

- **On-board Resources**

The capability to process information from a wide range of sensors and present information to the responder is expected to consume tremendous compute resources. Bigger and more powerful processors result in higher consumption of energy from the power banks which is shared with motors responsible for flight and hover time, in addition to increasing payload weight. Computational power, for current systems and approaches, becomes a trade-off with hover times. The commercial drones offer limited on-board computations resources which tend to be insufficient for a use case as complex as search and rescue and developing intelligent systems. Limited on-board data storage also presents a challenge in deploying large systems with high frequency and high resolution sensors.

- **Connectivity / Span / Coverage**




Figure 1.2: DJI Matrice 600: Payload vs Flight Time

With limited computational capabilities of a single drone, sharing of sensing and computational resources becomes key in a robotic system for SAR. This in turn increases reliance on strong networking (high bandwidth, low latency and jitter) backbone to allow intelligent coordination among agents. Scaling the system to span larger areas with weak networking infrastructure presents a huge bottleneck in such deployments.

### 1.2.3 A Strategy for Integration of Robots in SAR

With years of on-field experience, research, understanding of the dynamics of environment and the ability to sift through complex situation, humans continue to be unmatched in terms of efficiency and success of modern SAR operations. A considerable window of im-

provement still remains open in reducing dependence on human resources for hazardous tasks and augmenting the team's understanding of the SAR grounds. This opens up a significant opportunity to equip responders with resources to reduce operational error and risks involved. Given the advances in areas of cyber-physical systems such as computing, motion and intelligence, reinforcements in SAR response with robots develops into a realistic goal. A systematic approach to weaving robots and swarms into human driven missions would be to:

- Carefully evaluate the responders (or team's) strengths and affordances to thoughtfully add unobtrusive technology to aid his natural perception, add critical latent information or help comprehend information in an intuitive way
- Evaluate SAR team's and personnel's limitations, delays and scope of errors and systematically replace the parts a computer has demonstrated ability to perform that task better.

### 1.3 Summary of Contributions

The body of historical work makes the utility of aerial and ground robotic agents evident in SAR. Enabling coordination among semi-autonomous agents, which to a large extent mimics human behavior opens up some exciting avenues to develop practically viable systems. Open opportunities, the limitations presented by current systemic patterns and commercial equipment encourages us to pursue gaps to develop field-ready infrastructure for SAR. This work presents a small step forward in enhancing autonomy in such systems in a hope of increased success in SAR missions. The major contributions of work presented in this thesis are highlighted below:

- This work proposes a new approach and redefines interactions in Search and Rescue through a multi-agent system comprising of human responders, autonomous aerial agents and a wearable mobile compute cluster.
- As a part of this work, a prototype for the wearable compute cluster was developed to augment coordination among the other agents in the system.
- A thorough design process and objectives were framed to justify selection of components and form factor for the current and future iterations of the backpack.
- The capabilities of the wearable compute cluster were benchmarked with respect to compute capabilities on-board offered on commercial UAVs.
- Through utilization of latest and mature open source software frameworks, we demonstrate higher autonomy and collaboration among the nodes and scalability of the system in the future.
- We evaluate the in-field utility of the backpack by rudimentary, yet reasonable in-field testing and simulation.

## 1.4 Thesis Outline

This thesis is organised into 6 chapters including this introduction to the problem statement, existing challenges and our proposed contributions.

The following chapter ' 2, Distributed Intelligence for Multi Agent Systems ' acts as a primer and explores significant computational developments that has shaped the development of modern day distributed systems in robotic applications.

The chapter ' 3, A Multi-Agent Approach to SAR ' assess the major computational demands for a modern SAR response infrastructure. The chapter develops a new approach to SAR by integration of a mobile computing agent to augment coordination among robotic agents and human responders.

Chapter ' 4, WASP: A Surface for Edge Intelligence ' represents the most significant contribution of this research. The chapter describes the key specifications, design objectives, process of development of the mobile computing cluster. The chapter further lays out the software infrastructure setup for the Multi-Agent system proposed in 3.

The following chapter ' 5, Evaluation and Discussions ' highlights various benchmarking experiments undertaken to establish the advantages the backpack has to offer to the system for SAR.

The chapter on ' 6, Future Work ' describes the open areas that require attention to develop our initial prototype into a more field-ready product. The chapter also dwells into the extension of the backpack to aid and enhance systems in other allied applications. The final chapter summarises the work presented in this thesis.

# Chapter 2

## Distributed Intelligence for Multi Agent Robotic Systems

### 2.1 Evolution of Computing Systems

Developments in computing over the past several decades has increasingly supplemented human capabilities and understanding of the world. Diminishing cost of production, advances in fabrication techniques, computer architectures and power management has been transforming computers towards diminishing size and burgeoning computing capabilities. Subsequent reduction in costs of integrating computer systems in various tasks has further fueled the ubiquity of such systems. Areas such as mining, agriculture, education etc. which traditionally have been human-oriented and have been increasingly utilizing computer systems with remarkable success. Computers today have managed to attain an invisible presence in our daily lives and come in form factors ranging from nano-shaped tubules to self-driving cars. Rapid developments in paradigms and capabilities in three major areas viz. software systems, computational hardware and networking could be attributed to the success of modern day computing systems as we see today and also for the enablement of work highlighted in this thesis.

### 2.1.1 Software Systems

Software systems and algorithms have enjoyed the farthest reaching impact among all the aforementioned factors. Not only have developments and optimizations through software enhanced the ability to run applications efficiently, algorithms, per se have paved the way of reshaping the development of hardware systems more recently. Software development workflows and a vast body of open source software have enabled lightning fast progress in software development, re-usability of software, error corrections and maintenance. Two key critical area in software systems viz. Artificial Intelligence, and Software for Simulation and Visualization, arguably have enjoyed the highest traction in research and industry in the last decade of computing.

#### **Artificial Intelligence and Deep Learning**

Building on ideas from over the past 7 decades, recent developments in the field of Artificial Intelligence has revolutionized the traditional computing pipelines. AI can be thought of as a paradigm in software focused on developing devices capable of performing tasks requiring human level intelligence, without being explicitly program to do so. Having proven its merit for overcoming human capabilities in image classification and other vision based tasks [33], research and industry has increasing restored to utilization of machine learning solutions for problems spanning multiple domains. Enormous amounts of structured and unstructured data and increased computing capabilities not only fueled the developments, but have also rendered AI as a viable alternative to most traditional computing methodologies.

AI, innately a broad field of research, is further divided into classes of algorithms and principles. Machine Learning (ML), a branch of Artificial Intelligence, leverages statistical measures to model relationship between data points. The technique relies on availability of large

amounts of data to develop models for specific tasks. Success of AI in vision related tasks was largely attributed to Deep Learning (DL), a branch of Machine Learning that utilises multiple layers of neurons to hierarchically learn patterns and relationship in data [34] [35]. Deep Neural Networks (DNN) tend to replicate and mimic information flow in the human brain with the help of multiple layers of neurons (perceptrons) and hierarchically learning relationships through such layers.

Most algorithms in AI, specifically ML or DL adapt either a supervised or unsupervised learning behaviour depending on the task at hand. Reinforcement learning is another flavor of AI algorithms based on optimizing rewards directed towards a certain goal. The most common pipelines in Machine Learning (Deep Learning) algorithms follow three stages of processing:

- **Data Ingestion and Preprocessing**

Identifying task-specific data source, aggregation, storage, transformation and cleaning of data to allow for formation of meaningful correspondences between the distribution and the task. The stage might expect certain level of domain expertise.

- **Model Training**

The process of discovery and modelling of evident or latent relationships between data and underlying task (classes, labels or predictions)

- **Inference (or Serving)**

Process of productionalizing the model to serve a the expected purpose or task by generating predictions for an input etc.

A large body of work including reusable models and tools have been made available by researchers and industry that has significantly reduced the efforts for adapting (transfer learning), developing, deploying and maintaining such models. Advances in Deep Learning



has exposed some low hanging fruits in extremely stable and robust deployments of object and anomaly detection algorithms [36], significant for a SAR-like use case.

### Simulation and Visualization

Developments in simulation and visualization software have allowed users to model processes, inefficiencies and possibilities over time without actually designing and deploying the system. We have observed increasing emergence of photo-realistic simulation software that are able to mimic real world dynamism and interactions with high fidelity and accuracy. Developments in this area has not only made resource-oriented technology more accessible but have arguably reduced costs associated with development and production too [23]. Simulation software such as ROS Gazebo, Microsoft AirSim, NVIDIA Issac etc., have allowed modelling of robotic agents, sensors and interactions without the need to acquire UAVs or sensors while sufficiently accounting for real world dynamics. Simulation and visualization have also opened doors for more elaborate ways on interacting with multiple streams of data and information through technologies such as Augmented and Virtual reality.




Figure 2.1: Microsoft's AirSim Simulator with multi-modal visualizations

### 2.1.2 Computational Hardware

Increasing number of objects around us, in order to integrate certain level of intelligence, are becoming hosts to processors ranging in size and capacity. In their most common and identifiable format, General Purpose Computers take the shape of the Personal Computers (PC), laptops, and mobile phones which enable users to run a variety of computational loads on them. Manufacturers and hardware designers, depending on certain objectives such as capabilities and cost, strive to find a balance between on-chip resources like CPU, GPU, Memory etc. Very often the general purpose computers do not make good options for task specific requirements, which might require quantity of unique computational resource to perform a specialized operation rather than general purpose cores. This creates a classification based on the architecture, type of resources available on a chip and configurability. Processing units, based on their underlying architecture, can be broadly classified into four major categories:

- **CPU (Central Processing Unit):** Processors intended for general purpose computing and most commonly found in desktop PCs and cellphones. CPU processes information sequentially and supports very limited parallelism. CPU oriented software are very easy to write, develop and maintain and a large number of high level abstractions for writing applications (sequence of instructions) for CPUs do exist.
- **GPU (Graphic Processing Unit):** Processors comprising of a combination of thousand of identical cores, developed natively for rendering graphics. The utility of GPUs have now extended beyond graphic computing to accelerate specialised compute operations and loads. The higher number of cores allow for massive parallelism. GPU oriented software require hardware specific knowledge and usually are sections that run in unison with CPU oriented software.

- **FPGA (Field Programmable Gate Arrays):** Combination of logical elements and IP blocks that can be configured and reconfigured for a specific tasks. FPGA also allow for high levels of parallelism. Programming for FPGA is much more complex, and requires knowledge of Hardware Description Languages (HDL). Recent SDK offerings by FPGA vendors have made writing, compiling, linking and flashing applications easier, but a lot of work needs to be done to make it programmer-friendly.
- **ASIC (Application Specific Integrated Circuits):** Custom integrated circuit which are optimised to run a specific application(s). ASICs are developed with performance and power in mind and offer the most efficient platforms for application-specific tasks. The biggest drawbacks with ASICs are that they are non configurable and have usually very high associated development costs.

Besides developments in underlying processor architecture and fabrication processes, we also have observed phenomenal advances in storage capacities, memory access speeds and peripheral interfaces that have allowed extremely complex applications to generate results in real-time for cyber-physical and robotic systems. In the absence of one-good-solution, each system on a case basis is expected to be developed with an interplay of such hardware offerings.

### 2.1.3 Networking and the Cloud

Evolution of networking technologies has had far reaching consequences on scalability of systems as we see today and would significantly impact the systems of tomorrow. The primary boosters in driving such advances are advances in wired and wireless communication and the evolution of cloud services.

## Wired and Wireless Connectivity

Although the protocols and topologies in computer networking have retained a similar shape over time, advances in networking standards, data compression, error checking and security has enhanced communication between nodes. Modern standards in wireless networking highlighted by low latency and high bandwidths have enabled systems to grow wide and physically less complex. The latest IEEE 802.11 ac standard boasts of about theoretical max speed of approximately 6.933 Gbps. Such standards have trickled down already to cellphones, end points and routers for the end users. At the time of compiling this thesis, commercial roll out of 5G (cellular) devices had begun. The standard enables non line of sight speeds of 200 Mbps and and 600 Mbps to 1.2 Gbps of line of sight bandwidth [37]. Future-ready networking standards such a IPv6 addressing enable addressing 340 trillion-trillion devices to be connected and interfaced together, opening up new lines of interconnected intelligence.

## Cloud Services

Networking and connectivity, over the last decade has made mass under-utilization of computing infrastructure evident. This catapulted the emergence of cloud computing. Paul et al. define cloud computing as a model for sharing a pool of computational resources on an on-demand basis [38]. With the only requirement of sufficient channel for connectivity to such services, users can now take advantage of specialised hardware and preinstalled software without the need to own, manage and maintain the computing infrastructure.

Cloud services have increasingly driven applications out of on-premise systems and extended the capabilities to reuse infrastructure and reduce development costs. Cloud offerings for solutions today exist in three variations, depending on the level of management that is

required on-premise:

- **Infrastructure-as-a-Service (IaaS):** Managed of backend infrastructure.
- **Platform-as-a-Service (PaaS):** Managed infrastructure, OS and runtime.
- **Software-as-a-Service: (SaaS):** Managed infrastructure, OS, runtime, data and applications.

Cloud computing, improved addressing and networking capabilities have enabled a new generation of systems with extremely low cost of adding capabilities to devices and objects around us. Without the need for equipping robots, or agents, with large and powerful computers, such agents can now enjoy theoretically infinite resources as their disposal just by being connected to the internet. These capabilities have opened up avenues for improved efficiencies, form factors and scalability for robots.

## 2.2 Emergence of Distributed Systems and Intelligence

A remarkable impact of developments in accelerated hardware architecture, increasingly modular software algorithms and networking has been the enablement of computing models to go beyond monolithic systems, which were defined as self sufficient computing entities. Systems are developed to meet certain computational requirements which might keep changing over time. Scalability, which can be defined as the ability of a system to expand in response to increasing demands, becomes important in developing long lasting solutions. The two approaches to scaling a system are by either vertical or horizontal scaling. Vertical scaling refers to the approach of adding more resources to a single computing device, or node. Horizontal scaling can be viewed as an approach of adding more instances of the node

that can distribute the computational load. The nodes could be similar (homogeneous) or different (heterogeneous) depending on the use case and requirements. This approach of horizontal scaling paves the way for a much bigger area of work i.e. Distributed Systems. Steen et al. define distributed systems as a collection of computing entities working independently that appears to its users as a single coherent system [39].

Distributed systems present several advantages that can boost the utility of such systems in practical applications. The distributed approach to a large extent eliminates a single point of failure problem, where the system can continue to work in a limited capacity during downtime of node. Peripherals including sensors or devices, in isolated systems, are limited by the bus channel, address space and interfacing ports. Distributed systems, with higher node count, offers the capability to interface higher number of peripherals. A different combination of nodes might also supporting task specific customization with multiple hardware and software combinations.

Many modern systems, both complex and trivial, as we see today leverage the distributed approach to computing as it becomes easier to manage different pieces of application or storage across distributed heterogeneous systems. Key in defining the success of such an approach was also driven by popular software implementations and architectures such as MapReduce [40] and Distributed File Systems. Cloud computing can be considered as the single biggest manifestation of distributed computing.

While the connected ecosystem continues to evolve, it is not uncommon for robots and applications to span disconnected or network-affected areas. A growing body of work has focused their efforts on conditionally shifting towards or away from the cloud ecosystem as required. Edge computing, in contrast to cloud computing, is a model to which brings storage and computational resources closer to the sensors or point of application [41, 42]. Applications driven by low latency, poor reliance on network availability, has boosted interest in edge

computing and intelligence. Despite the contrasting differences in models of computing, Edge and cloud computing tend to complement each other so as to offer scalability and robustness to real time cyber-physical systems. Areas such as Artificial Intelligence and Security have benefited significantly due to the levels of convergence between edge and cloud computing. The six levels of edge intelligence proposed by Zhou et al. represents the convergence of artificial intelligence, cloud and edge computing and successful use cases [43].

With the driving factors described in the above section, efficient tools and frameworks and prolific community contributions in the form of open source software, systems have increasingly been able to offer higher capabilities with reduced energy needs. The renewed ability to scale and control implementations more tightly along the needs of the application is expected to redefine task-specific robotic agents as we see today.

# Chapter 3

## A Multi-Agent Approach to SAR

### 3.1 Computational Objectives for Multi-Agent Systems in SAR

Multi-Agent Robotic Systems present a challenging use case in terms of computing that starkly differs from traditional general purpose computing. Physical and electrical limitations, strict feedback timelines, and dynamics of the environment and agents make SAR response a hard problem from a computational standpoint. Multi-modal nature of sensory information, burgeoning data footprint and complex computational loads demand a hybrid approach tailored for the requirements. Given the interplay of all the aforementioned factors, coordination and collaboration among agents develops into imperative while building a solution for SAR. A few major guiding requirements for a multi-agent system solution for SAR can be identified as follows:

- **Scalable Infrastructure**

Multi-Agent Systems should have the capability to scale according to the demands of the application. Varying level of autonomy, search spans or mission objectives might need to be supported with a unique combination of resources. The ability to scale the system, in terms of addition or removal of nodes while maintaining a common backbone could significantly reduce cost and time associated with development and



deployments.

- **Resource Sharing**

In a collaborative system, it can be expected of modules to reuse information on a different node, to be processed or interpreted in different ways. Multiple or Multi-modal sensor inputs might necessitate sharing of information for a improved situational understanding from a data perspective. Sharing of computational elements can help prevent system or node level bottlenecks and significantly enhance information throughput by clearing congestion.

- **System and Data Redundancy**

Variability in environment and fragility of agents in the system may demand data duplication and high availability in situations arising from failure of agents or components. Systemic Redundancy adds robustness to the system and permits reduced power operation even in case of node loss.

- **On-board Footprint Reduction**

With enhanced frequencies and resolution of modern sensors, data storage and management emerges as a crippling challenge. A 4K video shot at 60 FPS, can amount to about 750 MB in a just one minute. An effort to reduce redundant information and piping them to the information processing pipeline can reap great rewards in faster scene understanding and improved cognition.

- **Power Constraints**

SAR missions utilizing mobile agents do not have access to constant and high capacity power sources. Applications and hardware have to be selected carefully to ensure that the most viable product can be developed out of practical power availability.

## 3.2 A New Multi-Agent Approach to SAR

Based on our SAR strategy highlighted in [1](#). Introduction, utilization of multiple agents, aerial or ground, to work collaboratively with human responders opens up a viable direction of progress. Coupling of humans and robots in Search and Rescue operations not only renders SAR responses safer but also has the potential to significantly boost success in such missions. Developments in distributed computing, artificial intelligence and success of existing systems, as described in [2](#), strongly back our hypothesis.

An evaluation of above computational demands for a Multi-Agent SAR, points us to a observable deficit with respect to available on-board resources on modern UAVs in an isolated mode of operation and coordination through ground control systems. Current approaches also tend to overburden the UAVs by computational inefficiencies in a collaborative system subsequently introducing latency and bottlenecks in processes. Additionally, over-reliance on communication channels impair the capabilities and range of such systems.

We observe that offloading various tasks and demands from resources available on the UAV has significant advantages in terms of longevity of SAR Missions dependent on aerial agents [\[44\]](#). Offloading tasks from the UAVs opens up the opportunity to reduce the utilization of on-board resources including memory, processor and power. It also enable execution of applications on a more suitable external computing surface capable of handling unique compute loads. Disruptive or limited range of communication channels restrict the role of cluster installation in ground control stations. We observe that mobilizing the compute cluster can solve several of coordination and collaboration challenges multi-agent SAR missions face today. After a thorough investigation of real SAR responses and responder interactions, we believe that a wearable cluster in the form of backpack can augment multi-agent collaboration and computational capabilities by a significant magnitude. This chapter describes

efforts in developing new multi-agent approach to SAR by integrating a a wearable back-pack cluster for SAR missions including design decisions, key specifications, and experimental benchmarking.

Our new multi-agent approach is best explained in terms of agents and their interactions as follows:

### 3.2.1 Agents

- **Unmanned Aerial Vehicles (UAVs)**

UAVs play the role of augmenting situational understanding from an aerial POV by acting as 'eyes-in-the-air'. Air being a significantly unobstructed media for motion, allows for easier navigation and routing. UAVs also extend situational access in areas which were otherwise inaccessible or dangerous for humans responders to venture. Besides access, the UAVs are expected to host a range of sensors including Global Positioning System (GPS), Cameras in the visible and thermal range, Ultrasonic Range Finders and LiDARs. The scope of this research included provisioning of a set of four DJI Matrice 600 UAVs with custom sensor carrier mounts [3.1](#). The custom mounts allow mounting of RGB Camera, Thermal Camera, and a LiDAR (Velodyne VLP16). The proprietary nature of the UAV manufacturer infrastructure and the need to extend on-board computational requirements for tasks like data reduction, storage, transmission and perception, an add-on embedded computer was mounted on the UAVs.

- **Mobile Computing Cluster**

The most significant contribution to this system from previous work in Multi-agent SAR systems is the introduction of a Mobile Computing Cluster. The mobile com-



Figure 3.1: Custom Drone Configuration for DJI Matrice 600

puting cluster is tasked to provide the central processing pipeline and computational surface for tasks including behavioral modelling, path planning, perception and visualization. The backpack is powered by an battery array. The latest iteration of the prototype of the computing cluster takes the form of a generic hiking backpack. The design and attributes of the cluster are provided in greater detail in the following chapter 4. For the remainder of this work, the WASP (WearAble Supercomputing Plaform), a nomenclature used for the backpack in its first iteration would be used interchangeably with the 'Mobile Computing Cluster', 'Wearable Mobile Computing Cluster', or simply the 'Cluster'.

- **Human Responders**

The HITL nature of the system integrates a human responder into the system. The Human responder or operator is responsible for managing autonomy of the system on a high level. This involves coordination with ground control station, controlling and

monitoring the operation of the agents by assimilating systemic information from the agents, software visualization and feedback. The responders must also monitor the health and status of the backpack and UAVs.



Figure 3.2: WASP (Wearable Mobile Computing Cluster)

### 3.2.2 Interactions

Interactions define how information, resources and tasks are migrated from one agent to the other. Below is the proposed behaviour of the system:

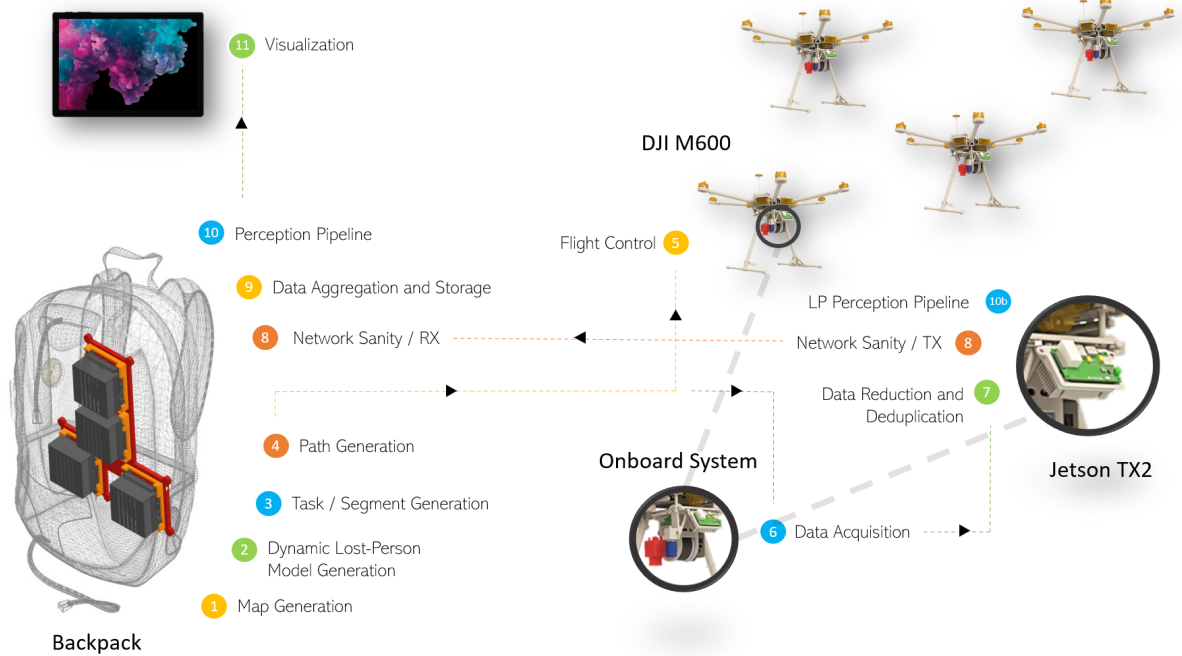


Figure 3.3: Agents and Interactions

- The master node on the backpack initializes other nodes and the networking backbone for the system. Distributed resource manager, containers and container orchestrators are also initiated in this stage.
- The backpack initialises a web server to present info-graphics for the user to access on a handheld device. The web interface enables mission-significant inputs and configurations from the user
- Geographic Information System (GIS) information and layers are fed through a pipeline on the backpack, to generate a high definition map of the region of interest
- A Lost Person Behavioral Modelling application [45] is then executed on the high definition map to develop a heat map representing the possibility of locating the lost person in the region of interest

- The region of interest is broken down into sectors or segments on the basis of heat maps from the lost person behavior modelling algorithm and geometrical preferences. A task order is then generated with prioritization of search sectors
- A planning algorithm generates waypoints for navigation [46] of the UAVs based on the task order and control parameters such as risk reduction, area coverage or internetworking span
- The compute nodes on the UAVs are initiated. The backpack and the UAVs establish communication and continually test for network sanity and communication stream. Containers expose REST APIs to expose underlying algorithms for distributed inference.
- Once the waypoints are generated and issued, the UAVs execute the flight control algorithm to navigate through the waypoints.
- The sensor processing pipeline are initiated. The processing pipeline aggregates data, execute on-board data reduction algorithms, stores a local copy and puts the sensor stream on the network channel. Any critical information in terms of search or detection are conveyed back to the backpack.
- On receiving the data, the backpack stores the data in a persistent storage, and runs detection (object and anomaly) algorithms on the data stream.
- In case of network disruption between the UAV and the backpack, the UAVs have the capability to independently execute a low-power perception pipeline.
- The backpack updates visualization based on lost person heat maps and the predictions of the detection algorithms, 3.4.





Figure 3.4: Proposed GUI with interactive agents, sensor and agent info

The interactions in our proposed approach for SAR exploits advantages offered by coordination among the agents in order to strengthen collaboration. Well defined responsibilities of each agent-ground in our system allows us to plan for contingencies in terms of fail-safe modes of operations.



# Chapter 4

## WASP: A Surface for Edge Intelligence

### 4.1 Physical Specification and Design Decisions

The choice of the form factor, computational, networking and power devices were derived from an in-depth evaluation of requirements posited by a SAR response and a set of design objectives. Technical requirements helped understand how to utilise the surface area on the chassis and available volume. The design objectives set the underlying premise for construction of the current version of the compute cluster.

#### 4.1.1 Physical Specifications

##### Processors

Three embedded computers from the NVIDIA Jetson family of devices, Xavier AGX, were integrated into the backpack using custom 3D printed mounts. High volume of GPU cores and a strong functional developer community makes the execution of common high performance compute loads feasible and reduces deployment turn around times. Optimizations offered by NVIDIA TensorRT supported for Jetson devices allow for high levels of optimization of

deep neural network compute loads resulting in faster inference. The devices combined to 840 grams when attached to the backpack. The devices were powered through USB-C port and were interconnected via on-board Ethernet ports. The UAVs were also mounted with NVIDIA Jetson Xavier NX, which are a stripped down version of the Xavier AGX.

Table 4.1: Key Specifications for Processor Kits.

<b>Make</b>	NVIDIA Jetson	NVIDIA Jetson
<b>Model</b>	AGX Xavier	Xavier NX
<b>Quantity</b>	3	1 per UAV
<b>GPU</b>	512 Cores, Volta@1377MHz	384 Cores, Volta@1100MHz
<b>CPU</b>	8 Core Carmel @ 2.26GHz	6 Core Carmel @ 1.4GHz
<b>Memory</b>	16GB LPDDR4X, 256-bit bus (137 GB/sec)	8GB LPDDR4X, 128-bit bus (51.2 GB/sec)
<b>AI Perf</b>	32 TOPS	21 TOPS
<b>TDP</b>	30 W	15 W
<b>Dimensions</b>	105 x 105 x 65 mm	103 x 90.5 x 34 mm
<b>Weight</b>	0.28 kg	0.172 kg



(a) NVIDIA Jetson AGX Xavier



(b) NVIDIA Jetson Xavier NX

Figure 4.1: Computing Platforms for the MA-SAR System

## Battery

Four high capacities lithium polymer battery packs are used to power the electrical components on the backpack. Lithium Polymer batteries are much safer, lightweight and more

compact alternatives to lithium-ion batteries. Each battery pack allows for powering devices through the standard US AC connector, USB-C and USB-A up to a maximum of 100 W per battery. The ports are provisioned to be used simultaneously as well. The batteries add a combined weight of 3.024 kg. Each battery in the set allows 500 recharge cycles.

Table 4.2: Key Specifications for Battery Packs.

<b>Unit</b>	Mophie Powerstation AC
<b>Capacity</b>	24,000 mAh
<b>Output</b>	100 W AC (110 V)
	30 W DC via USB -C (9V/3A, 5V/3A, 15V/2A, 20V/1.5A)
	12 W DC via USB-A (5V/2.4A)
<b>Dimesions</b>	190 x 114 x 28 mm
<b>Weight</b>	0.756 kg




Figure 4.2: Mophie Powerstation AC (24000 mAh)

## Router

A long-range router forms the networking head attached to the backpack frame. The router established wired connectivity routes to the on-board computing devices through 4 x LAN ports and extends wireless connectivity through 3 available bands (2.4 GHz / 5 GHz / 5 GHz) aerial and external computing devices (UAVs). The variant of router carries a in-build processor and USB 3.0 storage capabilities that can be utilised as network storage for the system. The processor on the router implements load balancing to distribute Wi-Fi utilization among all three WiFi bands, to ensure that all available bands are adequately utilized and avoids overcrowding a single band. The router also supports implicit Beamforming, which allows the router to actively tracks devices and directs Wi-Fi to clients using 6 on-device antennas instead of omni-directional relay.

Table 4.3: Key Specifications for Networking Router.

<b>Unit</b>	Netgear Nighthawk X6S AC4000
<b>Ethernet</b>	1 WAN + 4 LAN
<b>Bandwidth</b>	10/100/1000 Mbps
<b>Wireless Bandwidth</b>	4 Gbps (combined)
<b>Standards</b>	IEEE® 802.11 a/b/g/n/ac
<b>Bands</b>	2 x 5 Ghz at 1.625 Gbps
	1 x 2.4 Ghz at 750 Mbps
<b>Gain</b>	1.76 dBi (2.4 - 2.483 Ghz)
	3.12 dBi (5.15-5.25 GHz)
	3.11 dBi (5.25 - 5.35 GHz)
<b>Processor</b>	64-bit
<b>Memory</b>	128 MB Flash 512 MB RAM
<b>Dimesions</b>	295.5 x 226.8 x 54.5 mm
<b>Weight</b>	1.1 kg
<b># Antennas</b>	6



Figure 4.3: Netgear Nighthawk XS6 AC4000

### 4.1.2 Design Objectives

A huge non-tangible merit in design and success of any device lies in how easily it integrates into the existing scheme of things, making its presence invisible while extending intended benefits. Backpacks tend to be the most ubiquitous gear for SAR personnel, predominantly used to carry SAR equipment, communication radios and first aid or specialized medical kits. The ubiquity and utility of a backpack as a carrying equipment and innovation in terms of ergonomics over the past few decades, inspires the form factor for our proposed computing cluster [47]. A final prototype for the proposed cluster was constructed after multiple iterations of design through CAD based renderings and physical assembly. Choices about the design were made based on our three pillars of design decisions, which were identified as: Ergonomics, Modularity and Customizability and Safety.

#### 1. Ergonomics

Construction of the backpack must ensure minimal obstruction in common movement patterns such as walking, running, climbing and crouching for the SAR responder for the duration of the mission. The backpack must not weight more than maximum backpack weight i.e. 30% of body weight for a reasonably fit adult as per previous work by O'Shea [48]. Design must integrate padding and cushioning as required and seen fit for carrying it comfortably. The backpack must ensure physical and structural integrity during motion to eliminate any jerks and reduce any discomfort.

## 2. Safety

The equipment must be safety to use in SAR like situation. The equipment must prevent and / or restrict any electrical and fire hazard. Best efforts to be made to render the backpack as rugged and weather proof as possible. Fail safe modes to be implemented to eliminate accidents due to heat build up and provisions for air venting.


## 3. Customizability

The equipment shall ensure easy maintenance and repairs. To ensure that the system is adaptable as per situation, the backpack should be modular enough to upgrade components or increase computing, power or networking capabilities as required.

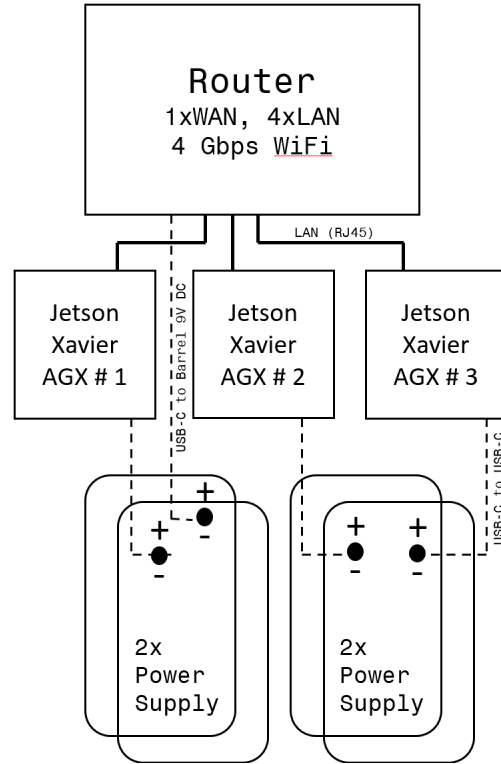
### 4.1.3 Design Process, Integration and Iterations

In order to reduce the number of physical integration iterations and material wastage, we adhered to a evaluation-simulation-evaluation strategy before manufacturing or acquiring components. A (1) requirements document listing out the physical and technical requirements of the item was produced for each part. Based on our design objectives and (2) brainstorming, a (3) preliminary sketch was developed and then translated into a (4) 3D CAD model using Autodesk Inventor. After another round of (5) design review and shredding, parts were either (6) 3D printed or acquired followed by (7) integration in the backpack. In the case of incompatibility or incongruity to requirements, steps (3) through (7) were repeated.

To harmonize with our design objectives, an ergonomic light-weight aluminium hiking backpack frame was borrowed from professional hiking gear. The frame allows for modification of height above the average shoulder level. A harness was attached to the frame with ample support and cushioning for the shoulders, spine and waist. The harness allows for adjust-



(a) Frame with Mounted Components



(b) Backpack Interconnects

Figure 4.4: Physical Backpack Assembly and Electrical and Networking Interconnects

ments in lengths using durable straps and buckles.


A custom chassis to be supported by the backpack frame, was developed using a Polyvinyl Chloride (PVC) sheet. The sheet was perforated with circular holes in order to eliminate extra weight and add flexibility to the sheet so as to follow a slight curvature of the backpack frame. The circular holes also allowed for attachment of devices and equipment to the chassis using suitable nut and bolts. PVC provides great electrical insulation, known to be an intrinsic fire retardant, and for its good dimensional stability. The design process ensured that no component fell out of the contour of the frame reducing the possibility of damage during the mission. The form and ergonomics of a backpack ensure that least degrees of freedom during movement are lost for the carrier. A thoughtful effort was made to drop the centre of gravity for the backpack as low as possible. The batteries and the computing

devices, which tend to be heavier, were attached to the lower half of the backpack. Symmetry to the best possible extent was maintained. This lowering of such equipment allow for translation of load bearings to the waist and away from the shoulders and improved balance for the carrier.

Three major set of devices, battery packs, computers and network router, were mounted on the chassis using custom 3D printed brackets to reduce vibration and improve fit. Special fitments were developed for creation and development of racks for battery to ensure sufficient airflow between batteries and the battery and the chassis. Specialized fasteners were build to press the batteries against the rack to eliminate movement of batteries within the rack.

The chassis was mounted to the backpack frame by 8 specialized mounts with two ball joints sufficient to align and stabilise the chassis on the frame. The design of removable right and flat chassis allows to mount the computing cluster to a ground vehicle, if need be, with sufficient ease.

The gross weight of our final prototype weighed 10.89 kg.



(a) Custom Platform Mount



(b) Custom Battery Mount

Figure 4.5: 3D Printed Custom Mounts



## 4.2 Software Infrastructure and Setup

Our proposed system allows for computational loads to be run at 4 different surfaces which includes:

- On Board UAV Computer [4.1b](#)
- WASP [4.1a](#)
- Handheld Device
- Cloud Infrastructure

### 4.2.1 Containerization

Virtualization is a method in computing to mask the underlying hardware for software application. Two modern methods of software based virtualization are through Virtual Machines (VMs) and containers. Containers are isolated and self sufficient silos for running applications in a environment of their own. Containers are usually built on top of the underlying host kernel supports addition or reuse of libraries on a per container basis. This is in contrast to Virtual Machines, in which each instance carries its own copy of the kernel or operating system. The re-utilization of the kernel and supporting libraries allows each container to use a significantly smaller memory footprint as compared to VM instances. Containers are also quicker to spawn and kill and allow for easy and independent CI/CD pipeline for different applications required in the project.

Given our requirements to execute a set of applications significantly different in nature and requiring a unique set of dependencies, we use containers to isolate our applications. We use container offerings from Docker for this project, which permits open source usage

of Docker containers. Docker daemons were instantiated in all nodes that defaults to the *nvidia runtime*. This allows the compute nodes to access the DLA and GPU cores available on the NVIDIA Jetson Devices.

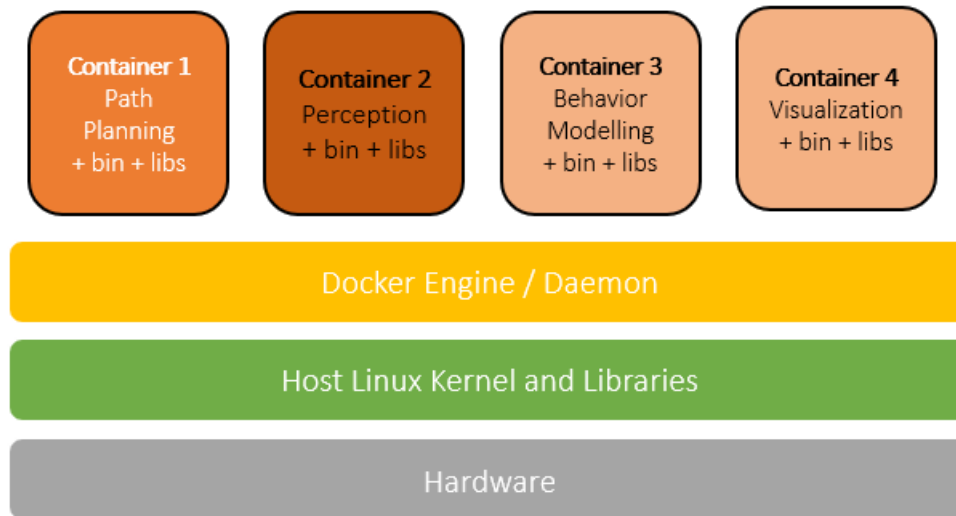


Figure 4.6: Docker Container Architecture for SAR

Open source nature of Docker has allowed the creation of a massive repository for pre-designed docker containers with the ability to be used out of the box for various application.

### 4.2.2 Orchestration

A multi-agent system like ours, demands the ability to not only launch containers across multiple systems or swarms, but also to orchestrate them on the basis of computational demands and available resources. A family of software tools called Container Orchestrators helps achieve our goals. The orchestrators help in provisioning of nodes using configurations to make resources available to the swarm for deployments, and manage their life cycle. They also help in managing networking among nodes and managing container allocations and deployments across them depending on load balancing criteria set by the user.

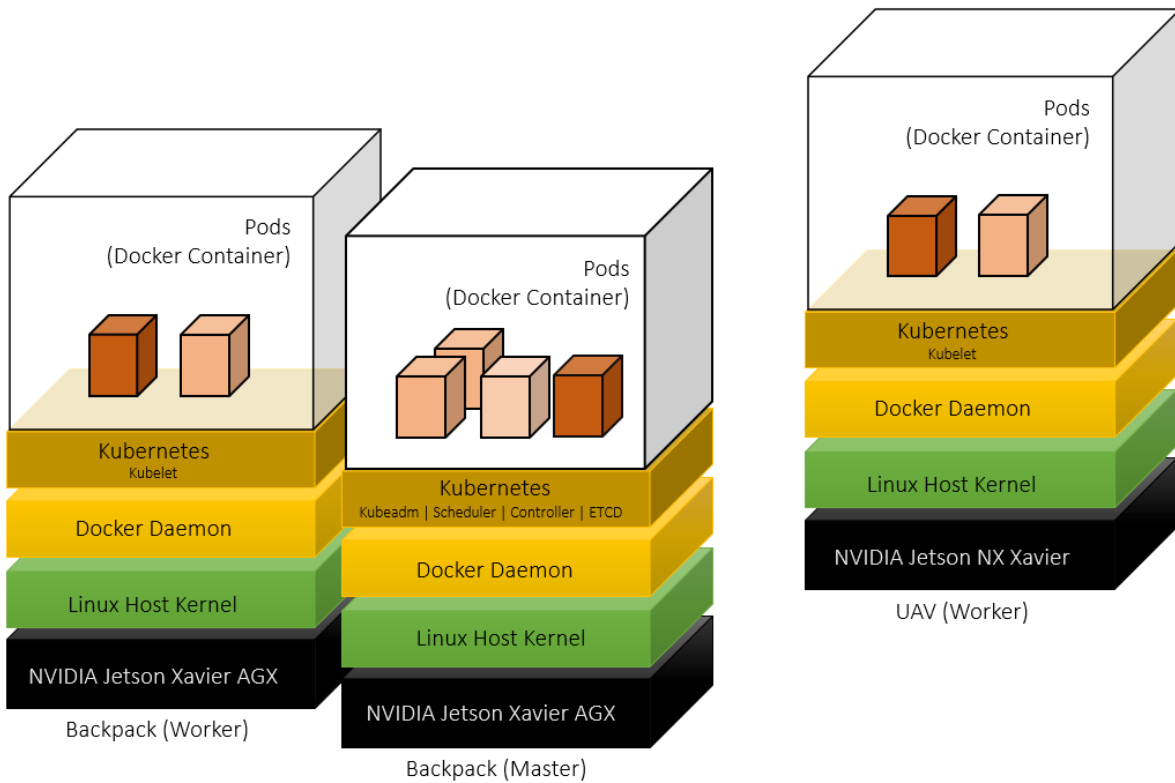


Figure 4.7: Kubernetes System Architecture for SAR

After careful evaluation of two popular offerings viz Docker Swarm and Kubernetes, we chose the latter for our deployments. Kubernetes is a mature open-source software framework developed and maintained by Google for container management and deployments and is successfully integrated in pipelines for few of the biggest technology organizations in the world. The ability to scale applications uniformly to cloud based services, available infrastructure, marginally better resources and documentations as compared to docker-swarm were the primary reason driving our choice of the orchestrator.

Kubernetes is set up across all nodes in the system in a Master-Worker configuration. One of the nodes on the backpack is configured as the master node, which houses the Kubernetes Control plane which includes the following major components:

- kube-apiserver
- kube-scheduler
- kube-controller-manager
- etcd
- Container Runtime

The Kubernetes infrastructure is set up in a high-availability setup, which means in the case of a failure of the master node, the master responsibilities are conveniently transitioned to another selected node. The worker nodes in the system run the kubelet agent, kube-proxy and the container runtime.

Each node in the system runs Pods, which are the smallest unit of computation that could be created and managed in the Kubernetes environment. Each pod could run one or more (Docker) containers. Pods could also be replicated in number on or across nodes. Multiple pods can be instantiated on a single node.

Configurations and definitions for Pods and replica sets are usually written through YAML or JSON file format which make it easy to write and understand. An noteworthy configuration offered by Kubernetes is to assign affinity for certain tasks to specific nodes. This comes of huge value in orchestrating tasks with preference to available nodes on the backpack instead of the UAV nodes.

### 4.2.3 Storage

The storage needs for the system arise from the requirement to aggregate multi-modal sensor data arriving from on-board cameras and sensors in the UAVs. Multiple sources of data streams and distributed computing warrants the need for distributed storage. For our setup,

a distributed data store is managed across all nodes in the cluster using the Hadoop Distributed File System (HDFS) infrastructure. HDFS, yet another tested and mature open source software framework, enables easy access to data in a network, data replication, data resiliency through failures and maintains data integrity through transactions. HDFS allows the convenience of any generic file system in terms of addition, deletions and moving of files. Each files is stored in the HDFS as a sequence of blocks which can be replicated across one or more nodes through a configuration on a per file basis. The HDFS is set up with a replication parameter of 2, which implies the number of copies of the data fragment maintained across the cluster.

# Chapter 5

## System Benchmarking and Discussions

### 5.1 Computation

We anticipate the system to be running three major flavours of SAR application software:

- **Deep Learning Inference**

Inference for Deep Learning (or Machine Learning) models posits a unique set of compute requirements and differ from training infrastructure. Inference process while having relaxed computational needs, tend to have lower tolerance for latency. After a network is trained, the model architecture and weights are passed through one or more variants or levels of network optimization which may include techniques such as optimization and pruning of the network to reduce the execution footprint in terms of memory and complexity of operations [49]. The output of the training and optimization stage is the model architecture and a set of weights associated with each node in the model. The Deep learning inference process vectorizes the input data and passes it through the network and set of weights to generate a prediction.

We deployed four Deep Neural Network models for the purpose of benchmarking inference capabilities on the backpack as well as on-board computers. We identify segmen-

tation, object detection and key-point detection as key objectives for visual cognition in a SAR context. While, classification models in computer vision tend to enlist the contents of an image, Object detection models go a step beyond to generate bounding boxes around classified objects. Image segmentation is a much finer approach to object detection, where classification happens on a per pixel basis, and hence could also be as per-pixel classification models [36]. Of many proposed approaches in DNN-based object detection, a remarkable variant is object detection and pose estimation by detecting reference key points, which we believe would be significant in a SAR use-case.

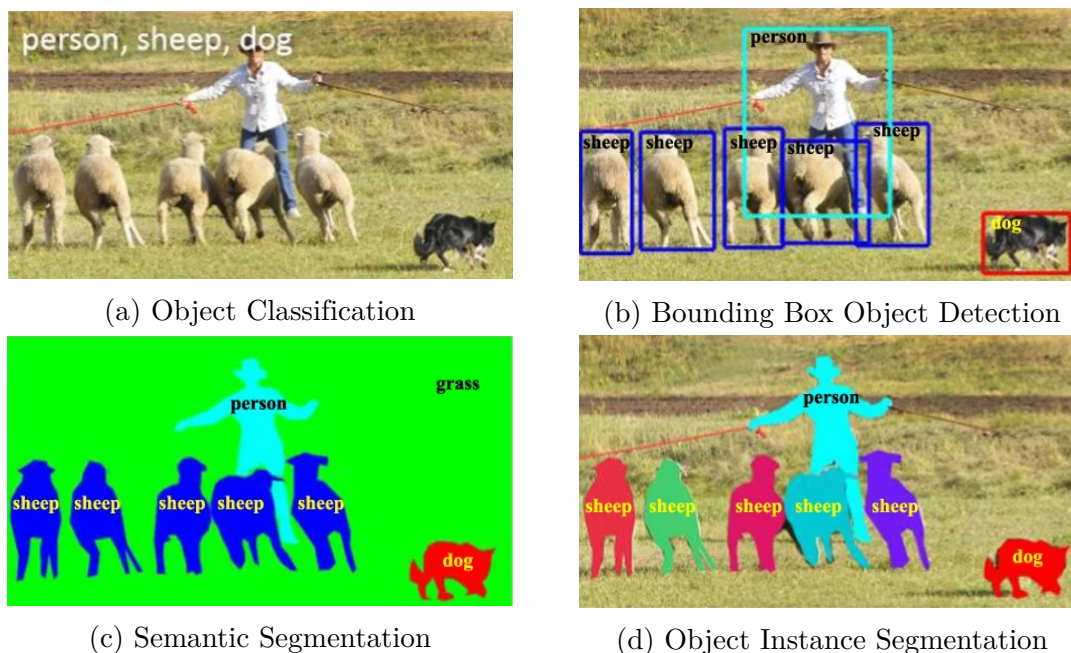


Figure 5.1: Detection and Identification in Computer Vision

To benchmark segmentation, point-based pose estimation, fast and high-accuracy object detection models on our hardware, we reuse the open source models and weights available for U-Net [50], OpenPose [51], SSD MobileNet V1 [52] and Inception V2 [53]. The models were optimised using a NVIDIA TensorRT pass-through. Frames per Second (FPS) were recorded and averaged for two runs for all four networks on a single host, and presented in 5.2. The tests were conducted with running background pro-




Figure 5.2: Deep Learning Inference Benchmarking

cesses and supporting applications including containers and orchestrators to emulate actual response time scenarios.

With the fastest on-board sensor generating 60 FPS of data as reference, we observe significantly high detection throughput from all the 4 evaluation networks. Even with the slowest network (Image segmentation with U-Net) we were able to extract 2.4x and 3.96x throughput with respect to the reference sensor frequency (60 FPS) on the Jetson NX and Xavier AGX respectively. We also observed throughput as high at 26.41x the reference frequency with SSD MobileNet based Object Detection on the Jetson Xavier AGX. This evaluation represents high computing capabilities allowing for ample leeway to run DNN-based computer vision algorithms on high resolution



images to achieve high fidelity detection.




Figure 5.3: Distributed Deep Learning Inference Benchmarking

To evaluate distributed deep learning inference, we ran 8 container instances (pods) of ResNet-Inception-V2 object detection network across three nodes in our system. Two of the nodes were Jetson Xavier AGX (AGXa and AGXb) mounted on the backpack and interconnected through the networking router. The third node was a Jetson Xavier NX (NXa) to emulate the UAV and connected to the backpack over Wi-Fi. We utilized the Deep Stream Containers provided by NVIDIA which runs the object detection network on 1080p sample video streams. The instances were scaled across the system through Kubernetes (kubeadm). We recorded the FPS throughput from each of the device in the system to generate 5.3. With just three heterogeneous devices (AGXa, AGXb and NXa), we observed an Object Detection throughput consistently exceeding 300 FPS on 1080p frames.

- **Path Planning and Decision Making**

Generating a task list and recommendation of search sectors are integral precursor to SAR missions. Recommendation and prioritization of task sectors includes generating Voronoi like search sectors and determining the priority of traversal and search. The task generation is followed by path planning for the multiple UAVs with several considerations. The path planning considerations might involve various indicators such as maximization of search space, maximization of connected coverage or other patterns. Multi agent path planning, historically seen and executed as CPU bound processes, opens up a massive opportunity to integrate parallelism.

- **Visualization**

A critical challenge in HITL SAR missions is to convey multi-modal, information rich data streams in a meaningful way to the user or operator. We identify live sensor streaming and rendering high definition GIS maps as two key operations driving the responder UI.

To evaluate graphical capabilities, we leverage a web-based GIS mapping and task order generation application developed by Wang and Lau at Virginia Tech. The application was ported and executed on a backpack node and visualizations were captured on a tablet connected to the backpack [5.4](#)

## 5.2 Networking and Communications

All the nodes in the system are interconnected through the network router mounted on the backpack. The nodes mounted on the drones and the hand held device connect to the networking hub (router) through Wi-Fi. The bands are managed by the load-balancer on the network router and can either be 5 GHz or 2.4 GHz depending on the proximity of the node to the backpack. The computing nodes mounted on the backpack are connected to the

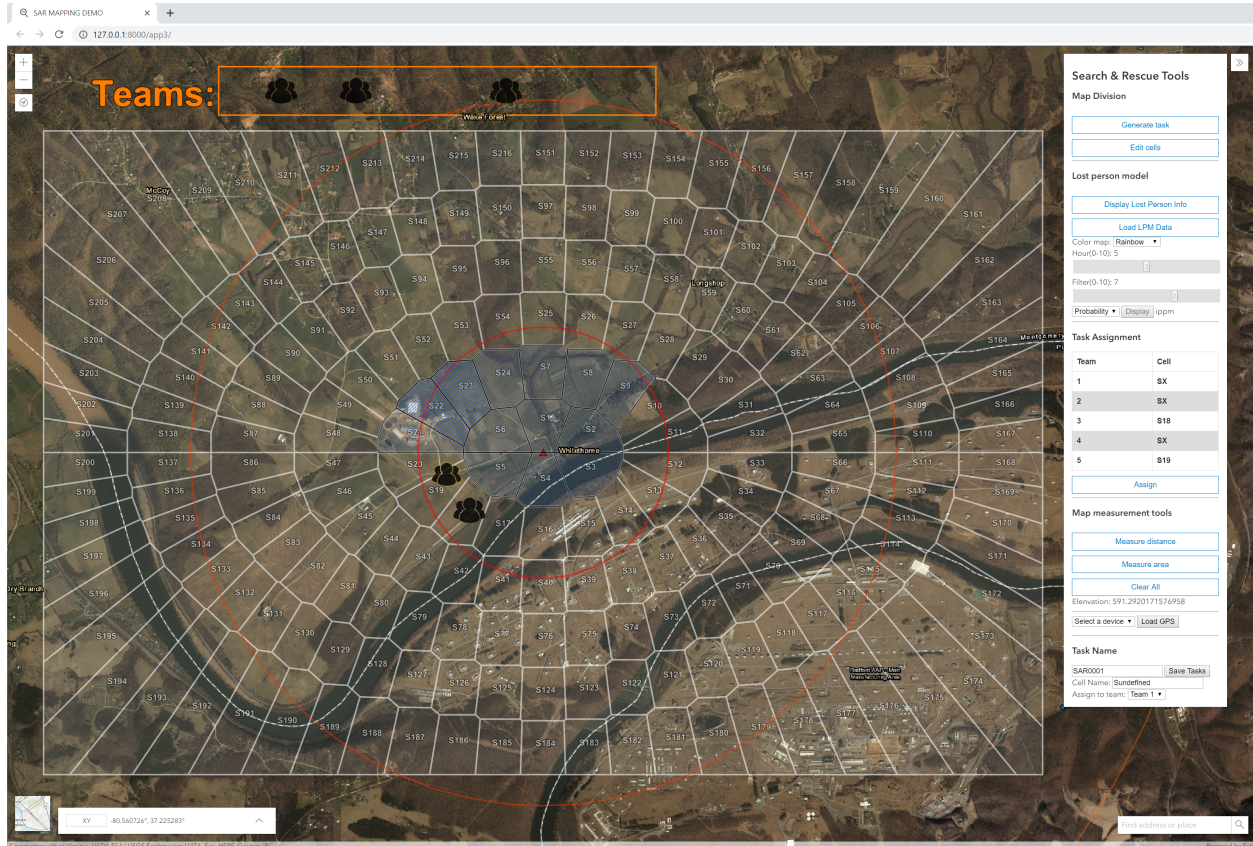


Figure 5.4: Visualizing GIS Maps and Sector Assignments

router through Ethernet interface. Static IPs are assigned to each node in the system.

To test networking and communications across the nodes, we used the final revision of our prototype with all three nodes active and interconnected via Ethernet. To emulate the UAV in a field setting, we used two standalone NVIDIA Jetson NX Xaviers. The Jetson Xavier NX's were connected to the backpack's router via Wi-Fi.

Exhaustive bandwidth benchmarking was performed across all the nodes. The evaluation tool used here was **iperf3** utility for Linux. Bidirectional tests were conducted across all communication channels WASP-WASP (Ethernet), WASP-UAV (Wi-Fi), UAV-UAV (Wi-Fi) for a period of 30 seconds each.

The figure makes it evident that UAV-UAV connection using an UAV as an Access Point




Figure 5.5: Bandwidth Tests over communication channels across nodes

(AP) has the lowest bandwidth. By adding a high-performant router as an access point, we significantly increase the bandwidth. Bandwidth significantly rises as we move data across the UAV nodes to WASP nodes. This strengthens our research hypothesis of augmenting collaboration among nodes by introducing a mobile compute cluster.

Once information has been passed over to the backpack, data can be routed through high throughput and bandwidth Ethernet channels for distributed intelligent applications via data and container orchestrators.

### 5.3 Power and Battery

Portability and mobility of the backpack levies severe power limitations and in a way determines the longevity of each session. Our current setup allows the distribution of power

Table 5.1: Power Benchmarking for Electrical Loads.

	<b>Session 1</b>	<b>Session 2</b>	<b>Session 3</b>
<b>Compute Load</b>	Idle	Deep Learning Inference	Netgear Nighthawk XS6 AC4000 via AC
<b>Type</b>	N/A	DetectNet Object Detection (SSD-Mobilenet-v2 /ssd_mobilenet_v2_coco)	1 Hz ping across devices
<b>Peripherals</b>	None	USB Webcam	3 x Wi-Fi; 3 x LAN
<b>Duration on Single Charge</b>	10h: 27m : 18s	06h: 52m : 54s	05h: 15m : 23s

through all four batteries, or powering each component through a unique battery. To evaluate the capacity of the batteries and transmission losses, we benchmark them with respect to power loads.

We conducted the benchmarks in three different experimental loads. In our first setup, a battery powered a backpack node which ran idle with Linux 18.04 and default processes. In the second setup, a single node was powered by a battery and DNN-based object detection network was run on a 1080p camera feed from a USB based camera. In our third setup, we connected our backpack router to a single battery. The router was connected to all three backpack nodes, and three Wi-Fi agents. Two pinging services were launched across devices connected through Wi-Fi and Ethernet to keep the network traffic active. Due to experimental limitations, we used the AC adapter for our final experimental session, which incurred an additional electrical loss.

## 5.4 System Integration and Field Studies

The structural integrity and usability of the backpack was evaluated through a on-field exercise. Common movements associated with a SAR response viz. brisk walking, crouching, climbing, leaping (-1 m), climbing (1 m) were exercised while wearing the backpack. We

observed and reported no significant vibrations or disassociation of components during the drill.




Figure 5.6: Mobility Testing with WASP

From our benchmarking efforts, we were able to theoretically establish the practical limitations in terms of computational surfaces, mission longevity, interfacing, networking and physical mobility. The experimental setups were constructed in order to mimic in-field testing as closely as possible and what was in the scope of this work. With the backpack, we were able to demonstrate significant computational capabilities to run complex computer vision algorithms to serve at a minimum of 60 FPS for at least 3 Full-HD sensor streams, and render intuitive and interactive user interfaces. With networking benchmarks we were able to make advantages of integrating a mobile compute cluster more evident in comparison to UAV based SAR response. We were able to establish the maximum mission time based on the lowest power benchmarking run at 5 hours and 15 minutes.



# Chapter 6

## Future Work

The new Multi-agent approach to SAR discussed in this work and the development of prototype of wearable compute cluster is a modest first step in realization of such a system and a lot of work remains to be done. This section on future work is divided into two sections viz. realization of the proposed multi-agent approach and the extension of the backpack into a mobile compute cluster for allied use cases.

### 6.1 Realization of Multi-Agent SAR

A successful realization of the proposed system is heavily dependent on eliminating structural design hurdles, robust software infrastructure and deployments and human centric interfaces.

#### Structural Design

- Weatherproofing

For the backpack to be practically viable, weather-proofing is a must. A study and evaluation of protecting material for the backpack was deemed out of scope for this work. Fabric membranes like Gore-Tex tend to be ideal candidates for developing a custom sheath or fabric housing for the backpack. Gore-Tex, a patented fabric membrane repels liquid water while allowing water vapor and air to permeate. The membrane has

successfully been utilized commercially for making breathable and waterproof clothing apparel and footwear.

- Integrity

The components on the backpack are currently attached through stainless steel nuts and bolts. Though structurally stable, they do have the tendency to loosen over time. Physical testing such as vibration tests hold huge value in evaluation of the structural integrity of the prototype assuming relatively rough usage during SAR missions. Such tests are commonly deployed for structurally complex machines like automobiles and UAVs.

- Human Factors and Ergonomics

From the human factors and ergonomics perspective, surveys in terms of success of such a product and usability should be carried, studied and evaluated. Thoughtful surveys, when conducted for the correct user group, play a pivotal role in bringing such equipment out of the prototyping stage into production.

## Software Infrastructure

- Developing Optimized Workloads for SAR

Though investigation and benchmarking through open source reviewed research algorithms help establish a basis for comparison, they are far from presenting an accurate measure of actual in-field applications. Specialized algorithms for Lost person Behavior Modelling [45], Risk Averse Path Planning [46] and Recommendation and Neural Network based multi-modal Aerial Object and Anomaly Detection algorithms are currently being worked on as a part of the umbrella project corresponding to this work. A significant effort lies ahead in optimising workloads for the hardware to leverage the



accelerated compute capabilities and parallelism offered by the system.

- Cloud as an Agent

With large-scale roll out of advanced cellular standards such as 5G, a possible line of work opens up in integrating cloud as an agent in the computational infrastructure. The choice of software frameworks theoretically make the current system conveniently extensible to utilize major cloud services such as Google Cloud Platform, Microsoft Azure or Amazon Web Services. Significant benchmarking and investigation would be required in terms of network bandwidth, latency and backend infrastructure to develop a minimum viable product.

- Resiliency and Managing Systemic Failures

Due to experimental limitations, the current software infrastructure lacks adequate in-field testing for resiliency in terms of node failure or data loss through communication channels. In addition to that, fail safe modes for all nodes need to be developed in order to push the system to a safe operating state in the situation of equipment malfunction or damage.

### 6.1.1 Hardware Integration and Platform Engineering

- Accelerated Hardware for SAR

Configurable and accelerated Hardware architectures like FPGAs have increasingly become more convenient to program and translate software code to hardware. Such devices and custom ASICs for Deep Learning Inference could be further investigated to enhance computing capabilities on the backpack.

- Contingency and Fail-Safe

Efficiency of processing and management in our proposed system is highly dependent

on the communication bandwidth and latency. Integration of active antennas could massively aid the otherwise restricted range of communication offered by our on-board networking router. A move away from proprietary frameworks offered by commercial drones makers could also help in eliminating redundant system on the UAV and unnecessary channel interference.

### 6.1.2 Human Computer Interaction

Working with simultaneous multi-modal sensors streams and prediction might suggest an cognitive overhead for the human operator. Intuitive and minimal displays that represent the most critical latent and evident information, while eliminating redundant information, is key to human understanding in SAR and follow-up response.

## 6.2 Mobile Compute Cluster for Other Use Cases

The mobile compute cluster was designed with an objective of enhancing compute capabilities of existing systems owing to increased application complexity and eliminating dependence on cloud and internet based services. With impact of modern software systems, a lot of areas open interesting avenues to advance this work.

The backpack has the ability to be used as a small-scale on-premise data centre for applications ranging from research, field testing, exploration, mining and construction, medical and educational camps.

The available networking channels can also help establish a temporary communication channel for military response, secondary disaster relief or event management purposes.

The ability to harness large computational resources make this a viable product to exe-

cute local simulations or computation for time or privacy sensitive applications such field operations in construction, mining, testing, scientific research and experimentation.

# Chapter 7

## Summary

Search and Rescue has historically been a significant challenge owing to high variations in the circumstances they are executed in and dangers associated with it. SAR teams, which largely comprise of trained volunteers often supported by skilled resources have increasingly seen the usage of technology equipment to counter unwieldy challenges, large region of interest, augment perception in environments and mitigate risks for human responders. Introduction of UAVs have had a remarkable impact on the success of SAR and has received a lot of attention in the research community, for advancing the utilization of aerial agents in SAR responses. The current patterns in integration, level of autonomy and capabilities of the modern drone limit the utility of UAVs in the field.

Significant advances in computing technologies and newer paradigms now has made enhanced autonomy possible for a SAR like use case. Available open source frameworks and tools which have matured significantly for a production setting can be easily translated to a SAR response use case. Newer fields in computing such as Artificial Intelligence, especially Deep Learning has enabled machines to learn to do a certain tasks without explicitly programming them to do so. Paradigms such as cloud and distributed computing have resulted in highly performant resources with remarkably less complexity.

This work proposes a new architecture for multi-agent search and rescue. The most integral contribution of this work is the development of a wearable mobile compute cluster. The compute cluster is responsible for providing a larger compute surface to a SAR Mission,

managing communication and resource sharing between the agents in a multi-agent SAR system. Evaluation of various aspects of the the backpack in enabling enhanced understanding by offering more powerful compute surface was undertaken as a part of this research. By theoretical evidence and benchmarking, we were able to conclude that the backpack was capable of producing real time and high accuracy results on common SAR compute loads and is a modest step in boosting success rates in SAR missions. This work also develops future course of action to transform the backpack from a prototype to production by carefully analysing structural and software limitations, and brings to light opportunities and low hanging fruits that could enable a swift absorption of such a platform in future SAR responses.

# Bibliography

- [1] US Fires, The Guardian. <https://www.theguardian.com/us-news/2020/sep/14/us-west-fires-nearly-all-missing-people-accounted-for-deaths-smoke-air-quality>.
- [2] National Crime Information Center (NCIC). Missing person and unidentified person statistics 2019. *Pursuant to the Requirements of the Crime Control Act of 1990*.
- [3] 2017 Annual SAR Report, National Park Service (NPS). <https://nps.maps.arcgis.com/apps/opstdashboard/index.html#/b526c87ae21f4a669eb6c9238c2c4bcf>.
- [4] Travis Heggie and Michael Amundson. Dead men walking: Search and rescue in us national parks. *Wilderness environmental medicine*, 20:244–9, 02 2009.
- [5] International Search And Rescue Advisory Group (INSARAG). <https://www.insarag.org>.
- [6] ASTM Committee F32 on Search and Rescue. <https://www.astm.org/COMMITTEE/F32.htm>.
- [7] The PEW Charitable Trust. Search and Rescue Teams Already Stretched Thin See Surge in Calls. <https://www.pewtrusts.org/en/research-and-analysis/blogs/stateline/2020/07/06/search-and-rescue-teams-already-stretched-thin-see-surge-in-calls>.
- [8] Kenneth Iserson. Injuries to search and rescue volunteers. a 30-year experience. *The Western journal of medicine*, 151:352–3, 10 1989.
- [9] Mario Silvagni, Andrea Tonoli, Enrico Zenerino, and Marcello Chiaberge. Multipurpose

- uav for search and rescue operations in mountain avalanche events. *Geomatics, Natural Hazards and Risk*, pages 1–16, 10 2016.
- [10] Karl Murphy, Maris Juberts, Steven A. Legowik, Marilyn Nashman, Henry Schneiderman, Harry A. Scott, and Sandor Szabo. Ground vehicle control at nist: from teleoperation to autonomy. 8 1993.
- [11] Magali Barbier, Jean-Francois Gabard, and Herve Ayreault. A generic architecture for autonomous uninhabited vehicles. page 17, 11 2007.
- [12] Sonia Waharte and Niki Trigoni. Supporting search and rescue operations with uavs. 09 2010.
- [13] Jill Drury, Holly Yanco, Whitney Howell, Brian Minten, and Jennifer Casper. Changing shape: improving situation awareness for a polymorphic robot. pages 72–79, 01 2006.
- [14] D. D. Woods, J. Tittle, M. Feil, and A. Roesler. Envisioning human-robot coordination in future operations. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2):210–218, May 2004.
- [15] Roger A. Chadwick. The impacts of multiple robots and display views: An urban search and rescue simulation. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 49(3):387–391, 2005.
- [16] Michail Chalaris, George Pallis, G. Statheropoulos, I. Balta, K. Mikedi, Pierina Julia Vamvakari, A. Pappa, Gavriil Xanthopoulos, Milt Statheropoulos, Sofia Karma, and Eleni Zorba. Use of unmanned vehicles in search and rescue operations in forest fires: Advantages and limitations observed in a field trial. *International Journal of Disaster Risk Reduction*, 13:307–312, 09 2015.

- [17] M Lewis, H Wang, Shih-Yi Chien, P Scerri, P Velagapudi, K Sycara, and B Kane. Teams organization and performance in multi-human/multi-robot teams. In *2010 IEEE International Conference on Systems, Man and Cybernetics*, pages 1617–1623. [ieeexplore.ieee.org](http://ieeexplore.ieee.org), October 2010.
- [18] B. Joseph, Sabine Langevin, O. Amram, Patrick Le Blaye, and N’Kaoua Bernard. Cognitive task analysis: A crucial step to improve the usability of a ground control station for unmanned air system mission. 06 2009.
- [19] Geoffrey Hollinger, Sanjiv Singh, Joseph Djugash, and Athanasios Kehagias. Efficient multi-robot search for a moving target. *I. J. Robotic Res.*, 28:201–219, 02 2009.
- [20] T. Furukawa, F. Bourgault, B. Lavis, and H. F. Durrant-Whyte. Recursive bayesian search-and-tracking using coordinated uavs for lost targets. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 2521–2526, May 2006.
- [21] S. Hayat, E. Yanmaz, T. X. Brown, and C. Bettstetter. Multi-objective uav path planning for search and rescue. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5569–5574, May 2017.
- [22] Jingxuan Sun, Boyang Li, Yifan Jiang, and Chih-yung Wen. A camera-based target detection and positioning uav system for search and rescue (sar) purposes. *Sensors*, 16:1778, 10 2016.
- [23] Michael Goodrich, Bryan Morse, Damon Gerhardt, Joseph Cooper, Morgan Quigley, Julie Adams, and Curtis Humphrey. Supporting wilderness search and rescue using a camera-equipped mini uav: Research articles. *Journal of Field Robotics*, 25:89 – 110, 02 2008.



- [24] I R Nourbakhsh, K Sycara, M Koes, M Yong, M Lewis, and S Burion. Human-robot teaming for search and rescue. *IEEE Pervasive Comput.*, 4(1):72–79, January 2005.
- [25] M. T. Agcayazi, E. Cawi, A. Jurgenson, P. Ghassemi, and G. Cook. Resquad: Toward a semi-autonomous wilderness search and rescue unmanned aerial system. In *2016 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 898–904, 2016.
- [26] G J M Kruijff, M Janíček, S Keshavdas, B Larochele, H Zender, N J J, T Mioch, M A Neerinx, J V Diggelen, F Colas, M Liu, F Pomerleau, R Siegwart, V Hlaváč, T Svoboda, T Petříček, M Reinstein, K Zimmermann, F Pirri, M Gianni, P Papadakis, A Sinha, P Balmer, N Tomatis, R Worst, T Linder, H Surmann, V Tretyakov, S Corrao, S Pratzler-Wanczura, and M Sulk. Experience in system design for Human-Robot teaming in urban search and rescue. In *Field and Service Robotics*, Springer Tracts in Advanced Robotics, pages 111–125. Springer, Berlin, Heidelberg, 2014.
- [27] A Freedy, O Sert, E Freedy, J McDonough, G Weltman, M Tambe, T Gupta, W Grayson, and P Cabrera. Multiagent adjustable autonomy framework (MAAF) for multi-robot, multi-human teams. In *2008 International Symposium on Collaborative Technologies and Systems*, pages 498–505. [ieeexplore.ieee.org](http://ieeexplore.ieee.org), May 2008.
- [28] Lanny Lin, Michael Roscheck, Michael Goodrich, and Bryan Morse. Supporting wilderness search and rescue with integrated intelligence: Autonomy and information at the right time and the right place. volume 3, 01 2010.
- [29] J Scholtz. Theory and evaluation of human robot interactions. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*, pages 10 pp.–. [ieeexplore.ieee.org](http://ieeexplore.ieee.org), January 2003.
- [30] Kam Tso, Gregory Tharp, Wayne Zhang, and Ann Tai. A multi-agent operator interface for unmanned aerial vehicles. volume 2, pages 6.A.4–1, 02 1999.

- [31] Aaron Hitchcock and Kelvin Sung. Multi-view augmented reality with a drone. pages 1–2, 11 2018.
- [32] P. Chrobocinski, E. Makri, N. Zotos, C. Stergiopoulos, and G. Bogdos. Darius project: Deployable sar integrated chain with unmanned systems. In *2012 International Conference on Telecommunications and Multimedia (TEMU)*, pages 220–226, July 2012.
- [33] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge, 2015.
- [34] Y. Bengio. Learning deep architectures for ai. *Foundations*, 2:1–55, 01 2009.
- [35] J. Schmidhuber. Draft : Deep learning in neural networks : An overview. 2014.
- [36] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. 09 2018.
- [37] 5g performance, coverage and power efficiency, qualcomm. <https://www.qualcomm.com/news/onq/2019/10/23/benchmark-global-study-proves-5g-performance-coverage-and-power-efficiency>.
- [38] Prantosh Kumar Paul and Mrinal K. Ghose. Cloud computing: Possibilities, challenges and opportunities with special reference to its emerging need in the academic and working area of information science. *Procedia Engineering*, 38:2222 – 2227, 2012. INTERNATIONAL CONFERENCE ON MODELLING OPTIMIZATION AND COMPUTING.
- [39] Maarten van Steen and Andrew Tanenbaum. A brief introduction to distributed systems. *Computing*, 08 2016.

- [40] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. In *OSDI'04: Sixth Symposium on Operating System Design and Implementation*, pages 137–150, San Francisco, CA, 2004.
- [41] M. Satyanarayanan. The emergence of edge computing. *Computer*, 50(1):30–39, Jan 2017.
- [42] P. Tsung, T. Chen, C. Lin, C. Chang, and J. Hsu. Heterogeneous computing for edge ai. In *2019 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, pages 1–2, 2019.
- [43] Zhi Zhou, Xu Chen, En Li, Liekang Zeng, Ke Luo, and Junshan Zhang. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. 05 2019.
- [44] Qinglin Yang, Xiaofei Luo, Peng Li, Toshiaki Miyazaki, and Xiaoyan Wang. Computation offloading for fast cnn inference in edge computing. pages 101–106, 09 2019.
- [45] Amanda Hashimoto and Nicole Abaid. An agent-based model of lost person dynamics for enabling wilderness search and rescue. 10 2019.
- [46] Barnabas Gavin Cangan, Larkin Heintzman, Amanda Hashimoto, Nicole Abaid, and Ryan K. Williams. Anticipatory human-robot path planning for search and rescue, 2020.
- [47] M. Gandy, D. Ross, and T. E. Starner. Universal design: lessons for wearable computing. *IEEE Pervasive Computing*, 2(3):19–23, 2003.
- [48] Michael O’Shea. Backpack weight and the scaling of the human frame. *The Physics Teacher*, 52:479–481, 11 2014.

- [49] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016.
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. volume 9351, pages 234–241, 10 2015.
- [51] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019.
- [52] A. Howard, Menglong Zhu, Bo Chen, D. Kalenichenko, W. Wang, Tobias Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *ArXiv*, abs/1704.04861, 2017.
- [53] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. *AAAI Conference on Artificial Intelligence*, 02 2016.