ABSTRACT

Title of thesis:	ARTIFICIALLY INTELLIGENT MEDICAL ASSISTANT ROBOT: AUTOMATING DATA COLLECTION AND DIAGNOSTICS FOR MEDICAL PRACTITIONERS
	Paulos Daniel, Nina Horne, Kevin Kuo, Michelle Marsandi, Natalie Offenberg, Dana Ronin, Ryan Utz, and Johan Vandegriff
Thesis Directed By:	Dr. Anil Deane Institute for Physical Science and Technology

Healthcare providers face financial, regulatory, and logistical obstacles in supplying quality care. A robotic system can improve patient outcomes and reduce demands on providers by automating data collection and supplementing medical diagnoses. Team AIMAR (Artificially Intelligent Medical Assistant Robot) constructed such a system focusing on three core features: natural language interaction, computer vision, and mobility. Thus, in addition to developing a robotic base with navigational and conversational abilities, Team AIMAR implemented two prototype modules: a skin lesion image classifier and a medical chatbot. Additionally, Team AIMAR created a framework to test and assess the functionality of the fully integrated system in a simulated environment. Several directions exist for future work, including expanding the user interface, improving navigation and sensing capabilities, communicating with electronic health record systems, and the integration of a physical arm.

ARTIFICIALLY INTELLIGENT MEDICAL ASSISTANT ROBOT: AUTOMATING DATA COLLECTION AND DIAGNOSTICS FOR MEDICAL PRACTITIONERS

by

Paulos Daniel, Nina Horne, Kevin Kuo, Michelle Marsandi, Natalie Offenberg, Dana Ronin, Ryan Utz, and Johan Vandegriff

Thesis submitted to the Gemstone Honors College of the
University of Maryland, College Park in partial fulfilment
of the requirements for the honors citation of
the Gemstone Honors College
2021

Advisory Committee:
Dr. Anil Deane, Institute for Physical Science and Technology
Dr. James Borelli, Department of Physical Therapy and Rehabilitation Science
Dr. Mark Fuge, Department of Mechanical Engineering
Dr. Marc Levitt, Children's National Hospital
Dr. Ian White, Fischell Department of Bioengineering

Table of Contents

List of Figures

Chapter 1: Introduction

Access to healthcare is a global issue. Healthcare systems in the U.S. are inadequately prepared to handle the high influx of patients. In a single year, individuals spent a total of 1.1 billion hours seeking healthcare [1]. The average patient spends 121 minutes on a routine doctor's office visit, a large portion of which is spent waiting to see a doctor and filling out paperwork. This equates to an average loss of $43 per doctor's visit for the patient's employer [1]. Time constraints imposed by many other aspects of daily life prevent patients from receiving proper medical attention. Moreover, the COVID-19 pandemic has shown the value of decreasing unnecessary face-to-face interactions between healthcare professionals with infectious patients. Technology is becoming more integrated with everyday life than ever before and has the potential to address these concerns by assisting with and even automating various healthcare tasks. A system with such capabilities would greatly reduce administrative burden and streamline the delivery of healthcare services, thereby improving patient outcomes.

Artificially intelligent (AI) systems and algorithms have a wide array of applications. AI is now as accurate as medical experts at diagnoses in several areas. Noting the potential to automate diagnosis and widely increase access to medical consultation, several companies are developing AI systems for decision-making in healthcare. One example is Brightlamp, a startup that developed a cloud-based mobile application to diagnose concussions through the measurement of pupil dilations [2]

[3]. These measurements are then pushed to a cloud-hosted neural network to determine if the person is exhibiting characteristic markers of a concussion. This example also highlights the ability of medical AI systems to both build and utilize digitized medical records and data sets.

Team AIMAR aims to build a modular, robotic, artificially intelligent, healthcare system to stream diagnostics and data collection in a doctor's office setting. The potential outcomes of this system include reducing the workload on doctors and nurses, improving patient outcomes while lowering healthcare costs, decreasing the amount of time spent in a doctors office, enabling easy access to fast medical care regarding skin diseases and common medical diseases, and providing medical assistance for areas with limited healthcare access. Direct objectives related to improving patient outcomes are 1) reduced errors in diagnosis, 2) reduced number of diagnostic tests needed, and 3) reduced demand on individuals making the diagnosis. There is a great disparity of quick, easy and cost effective access to healthcare; the potential outcomes listed above could greatly improve equity in healthcare.

Prior research has been conducted by former Gemstone team SCOPE (Sensory Computing and Object Processing Entity) in creating an assistive robot for healthcare delivery that is mobile, responds to spoken commands, and demonstrates artificial intelligence by extracting meaning about patient's health from conversations and visual interactions, and summarizing these observations into reports that could be merged with patients' electronic health records [4]. SCOPE has indicated that many further features, functionality, automation, and integration are necessary, such as securely connecting to electronic health records, expanding conversation capabilities, and improving obstacle avoidance. SCOPE integrated features such as visual object recognition, mobility, and natural language processing into the main program, using such facets to streamline triage and routine interactions between provider and patient.

"The AIMAR robot" refers to the entire system of hardware and sensors, interface and control software, and medically assistive capabilities. Navigation and mobility are core features of the robotic base, and extending capabilities to arm movement is a highly desired goal. Interface and control softwares, which handle basic voice interaction and navigation are also considered part of the robotic base. Modules for specific physical or medical tasks can be added to or removed from this robotic base. In general, a module's purpose is to carry out a data collection task, medical analysis task, or both. The action of taking and diagnosing a picture of a patient's skin condition is one such example. This modular organization of software allows distribution of computational tasks to various platforms, whether it be parsing speech onboard the robotic base, organizing data on a local network machine, or analyzing medical data in the cloud. Each module can provide specialized diagnostic information as a supplement to the standard information a healthcare worker receives about a patient. This data can then be used by the practitioner to arrive at the best course of action.

Team AIMAR chose to focus on two modules as examples of relatively routine screening. The first module is a medical chatbot, which converses with a patient, asks what symptoms they are experiencing, and determines a medical condition. The module heavily utilizes natural language interaction software and clearly reflects the capability to autonomously collect medical data. The second module is a skin lesion image classifier, which takes pictures of a patient's skin and classifies them using a deep learning model. Though skin lesion diagnosis is a highly specialized task, current computational methods exceed the accuracy of general practitioners due to recent advances in deep learning and availability of quality datasets. Furthermore, implementing a skin lesion classifier will lay a foundation for future modules that require a camera and execute computer vision tasks.

Team AIMAR envisions the ability to add or interchange modular components for the base robotic system depending on the need within the hospital setting. Other modular components could include an EEG monitor, heart rate tracker and pulse oximeter. Each of these modules adds a unique element to the AIMAR robot that could potentially aid medical professionals depending on the right context. But again to reiterate, due to the time and resource constraints Team AIMAR has decided to focus on the medical chat and a skin image detection algorithm as the potential modular components for the AIMAR robot. In the context of image recognition, the integration of machine learning programs could assist with the processing and analysis of images of skin pathologies. In conjunction with both local processing and cloud computing, the modularity of Team AIMAR's proposed product will increase the adaptability of the base model to meet a wider host of challenges in healthcare, including diagnostic support. Therefore, the guiding research questions for Team AIMAR are as follows:

1. How can modular sensors and a user interface be integrated with a robotic system to streamline data collection?
2. How can machine learning algorithms be leveraged to provide accurate diagnoses of specific diseases?
3. How can natural language technologies be paired with a robotic system to assist medical professionals?

Chapter 2: Literature Review

Team AIMAR reviewed the literature concerning research on robotic components, generalized medical knowledge, and image processing to analyze skin conditions. Many existing robots have robotic bases that allow the robot to move and a user interface. The incorporation of several data collection modules was not seen in an existing robot. Team AIMAR assessed several existing robots in order to determine which components would be best integrated into the AIMAR robot. This provides a framework upon which Team AIMAR will build to narrow the project direction through additional research in mobility, robotic arms, computer vision for awareness of the environment, and sensor calibration for different environments and individuals. AIMAR will host a variable number of medical modules; AIMAR's skin image analysis is a proof-of-concept for one such module.

2.1 Examples of Medical Assistant Robots

Many medically assistive robots of varying specifications already exist, and the research supporting these robots contains a wealth of information on their design and implementation which has been taken into consideration in regards to the construction of AIMAR. These robots also showcase the various applications that a medically assistive robot can be used for, as well as the areas in the field that could use improvement.

2.1.1 Pearl the Nursebot

Some prior robotic approaches include Pearl the NurseBot. Pearl assists the elderly by providing daily reminders and navigation guidance [5]. Pearl relies heavily on probabilistic AI techniques at virtually all levels of perception and decision-making. The robot successfully demonstrated that it could escort people in an assisted living facility, a time-consuming task carried out by nurses, interacting through speech and visual displays. Because many elderly people have difficulty understanding even simple sentences and articulating a response in a computer-understandable way, the robot uses efficient particle filter techniques to track and adapt to individuals. A partially observable Markov decision process algorithm performs high-level control, arbitrating information gathering and performance-related actions, and safety considerations are incorporated into simple perceptual modules through a risk-sensitive robot localization algorithm. Pearl features an off-the-shelf autonomous mobile robot navigation system, speech recognition software, speech synthesis software, fast image capture, and compression software for online video streaming, face detection tracking software, and other various software modules. Pearl was tested in various experiments by Montemerlo et al. [5]. First, it interacted verbally and spatially with a large number of elderly users with the task of delivering sweets. Second, the robot autonomously led 12 full guidances involving 6 different elderly people in which it contacted the person, explained the reason for the visit, walked them through the facility, and provided information on the weather. The data were collected through the elderly person's initial reactions to the robot. After interaction with Pearl, post experimental debriefings were performed with the elderly. The main strengths of Pearl were its ability to communicate and guide the elderly, and the most limiting flaw was its inability to adapt its velocity to people's walking pace. These results show promise in the use of AI to adapt to individual people, specifically through conversations and navigation through facilities. Pearl is an example of an existing robot that uses natural language interaction to aid in a healthcare environment. However, one of the main issues with Pearl is that it was unable to change its velocity to adapt to the specific elderly person it was assisting. This is something that Team AIMAR will take into consideration when determining whether the robot will be navigating alongside the person it is assisting or if it will be autonomous.

2.1.2 Care-O-Bot II

Care-O-Bot assists the elderly and handicapped with household tasks by fetching objects, serving as a mobility aid, and integrating communication and social tasks. It aids patients in watching TV, acts as a day-time manager that provides medication reminders, supervises vital signs and emergency calls, and calls physicians or authorities using a speech interface or touch-screen [6]. Care-O-Bot is equipped with a hybrid system architecture, containing deliberative and reactive components. The robotic assistant can be instructed interactively through the user interface. Each user command is transferred to a symbolic planner which then generates a list of actions necessary to fulfill the given task. The robot control system executes these tasks consecutively using a previously learned world model provided in a database. The assistant's knowledge about its environment is updated continuously based on sensor readings. One of Care-O-Bot's most important skills is the ability to manipulate objects in its environment. A manipulator arm, developed specifically for mobile service robots, provides the possibility of handling typical objects in the home environment. The flexible gripper attached to the manipulator is suitable for grasping objects found typically in a household e.g. mugs, plates, and bottles.

The most relevant aspects of the Care-O-Bot include its manipulator arm, communication and social integration features, and sensor-based knowledge of its environment.

2.1.3 Social Robot

Another example of a robotic medical assistant for home use is Social Robot. A multidisciplinary team of researchers designed a robot that aimed to provide patients comforting social stimulation as well as practical health care assistance at home. Adapting to the ever-changing medical needs of elderly patients is an important aspect of such a technology. Specifically, Social Robot's medical functionalities exist in two main formats: in the cloud as applications that can easily be updated or modified, and through a communication interface that connects a patient with a virtual care team of medical specialists [7]. On board, Social Robot is equipped with a variety of hardware components that allow it to navigate and effectively address patient needs. Navigation is supported through the use of an infrared sensor, an inertial measurement unit, and ultrasonic sensors. Communication with the virtual care team is facilitated by an High Definition (HD) camera and microphone. Interaction with the patient is provided through a touch screen tablet, a speaker, and a set of colored Light Emitting Diodes (LED) that form Social Robot's face (Figure 1). All of these tools, paired with hardware, allow Social Robot to provide useful assistance to patients. For example, Social Robot can identify a patient through facial recognition software, ask the patient questions and determine the patient's emotional state, and then based on the response, can direct the patient to an appropriate activity (e.g. social activity: call friend, or needs support: call caregiver).

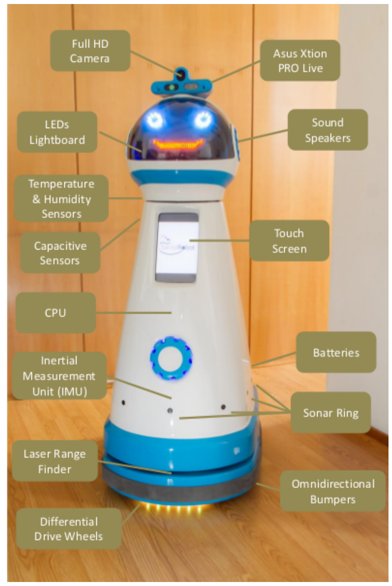

Figure 1. The Social Robot and its capabilities are shown.

2.2 Healthcare Robot Design Concepts

Various existing medical functionalities in robots were analyzed; the constraints considered were feasibility, time constraints, and individual interest. Team AIMAR studied designs with a few basic sensors and a connection to electronic health records, and add-on modules with specific use cases, rather than a monolithic design with all the sensors built-in. The applications of such concepts would reduce the overall cost of the base unit and would allow different sub-teams to work on individual modules. Mobile robots, compared to a stationary kiosk, enable flexibility wherever the robotic system is deployed. The use of the LIDAR sensor to map out objects calls for a controlled environment, allowing for easier mapping. Doctor's offices are much smaller and easier to map with the LIDAR sensor compared to a hospital. One of the largest decisions was whether to focus on diagnosis of one or more diseases, or assistance with tasks. The AIMAR robot would collect patient information to help inform a doctor's diagnosis rather than assist with tasks since people are much better at assisting patients, but not as good at processing large amounts of data. This would help reduce the amount of time a doctor would spend per patient, as they would already have the diagnostic data to analyze. For future research, possibly by a Gemstone different team, AIMAR as a home robot would assist a client with a specific disease to complete tasks and would monitor their specific condition and alert a doctor if any issues arose.

2.2.1 Drive System

University of Maryland Gemstone Team SCOPE constructed a robotic platform to help analyze data about a patient's condition. Their system used four holonomic wheels from the VEX EDR Robotics platform to drive forward, backward, and turn. Holonomic wheels are wheels with two degrees of freedom, allowing a robot to move in two dimensions without having to rotate. In a healthcare facility setting, to avoid obstructions, it is ideal for a robot to have similar movement capabilities as a human to move in tight spaces and avoid turning. Ideally, it should be able to move sideways and diagonally to maneuver difficult spots. Holonomic wheels can be used for this purpose; however, most holonomic wheels are thin and use independent rollers to satisfy two degrees of freedom. As a result, they have difficulty traversing non-flat or slick terrain, and the thin wheels can get stuck on holes in the ground (e.g. the gap between the elevator and the floor in a healthcare facility setting).

Another possible solution to navigating healthcare facility floors is a ball drive system. Ball drive systems use balls for wheels, while rollers transfer force to the ball to rotate the ball in all directions. One type of ball drive system uses only one ball to move but is too complex and expensive to implement within a 3 years [8]. A system similar to this one using 3-4 smaller ball drive units has the potential to reduce design complexity and fit a crowded healthcare facility setting better, as the wheels are spheres instead of thin cylinder shapes, allowing it to travel over obstacles easier.

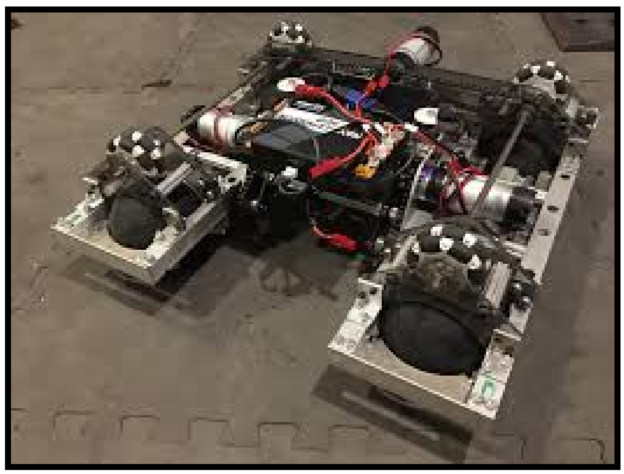

Figure 2. A sample ball drivetrain was developed for use in a robotic competition. Four balls keep the robot balanced on the ground while being able to rotate and move in all 8 cardinal directions.

2.2.2 Robotic Arm

To increase the functionality of the AIMAR robot as well as add the ability to manipulate objects, Team AIMAR considered implementing some form of robotic arm. Initial tests have been performed using the LCV Lab's UArm swift pro that has 4 degrees of freedom. However for the final design, AIMAR would ideally require more mobility to grab delicate and oddly placed objects such as medical equipment like thermometers. A robotic arm with six degrees of freedom would allow the arm to grab an object at any point in its range, allowing for the imitation of movement of a human arm [9]. The cost of constructing or even purchasing a 6-axis robotic arm is prohibitive, thus requiring Team AIMAR to utilize and run tests with the UArm provided by the LCV Lab instead.

2.2.3 Comparison of User Trust Levels in Different Robotic Interfaces

When developing a patient-facing medical interface that is intended to simulate interaction with a real human, three main options exist-- an avatar, a robot, or video footage. Though the choice between these three options is specific to each project, the user's level of trust should be a universal consideration. In a study conducted to assess user's level of trust in expertise between these three mediums, researchers gave participants a series of extremely difficult questions, provided them two advisors (an avatar, a robot, or video footage of a human) and rewarded correct answers with chocolate [10]. The researchers found that participants sought advice more often from the robot and video experts, and less often from the avatar. These results suggest, at least in this specific scenario, that a lack of personification does not harm user trust in the robot. Furthermore, the use of an avatar should be discouraged in most cases to maximize the user's trust in their robotic caretaker. In regards to the limitations of the study, the authors acknowledged that their results may have suffered due to participants' prior knowledge of question material as well as the inability to gauge a specific participant's desire for chocolate. Influenced by this study, as well as the technical constraints for implementing avatar features such as lip synchronization, Team AIMAR elected to not use SCOPE's animated avatar interface, shown in Figure 3 below.

Figure 3: SCOPE's Avatar projected onto a computer screen. They selected an animated avatar interface.

The AIMAR robot does not aim to replace or imitate a human, but rather embrace that it is a functional robotic assistant. As later detailed in the methodology, Team AIMAR's use of text and images in Mycroft GUI has no consideration for lip synchronization or avatar animation; such a design choice will not detract from the AIMAR robot's functionality and user reception.

Comparison of robot size is also important when making design decisions based on user trust levels. In a study on the effect of robot size on the anxiety felt by a human subject the robot approaches, researchers found that subjects felt most anxious with the larger robots at 1.8m tall and smaller robots at 0.6m tall compared to a robot with 1.2m height [11]. Guided by this study, team AIMAR decided on having a medium-sized robot at a height of 4 feet. This size will enable the AIMAR robot to perform its necessary modular functionalities while helping surrounding patients feel comfortable.

2.2.4 Healthcare Robot Capabilities

In addition to the capabilities discussed above, some other important characteristics for the AIMAR robot are affordability, safety, sufficient battery life, and additional sensors. More affordability means increased accessibility to healthcare facilities with limited funding. Safety features such as smooth edges and a stable center of mass to prevent falling are important for the AIMAR robot to interact daily with patients and healthcare workers [12]. Battery life is key for the robot's use in a healthcare facility. Healthcare workers often work twelve-hour shifts, experience shift overruns due to staffing shortages and the need to complete tasks before leaving, may have to work overtime, and have to ensure 24/7 care [13]. To help alleviate stress, demand, and workload, the AIMAR robot should ideally have a battery life of at least 24 hours, double the time of a nurse's shift, and capable of being charged within a couple of hours. Ideally, the facility might have a couple of AIMAR robots that take different "shifts," as well as hold different "roles" within the facility as needed.

2.3 Medical Diagnostic Modules

2.3.1 Natural Language Interaction

AI also directly assists medical professionals, through tools such as medical reference chatbots and text analysis [14], [15].

Team AIMAR's primary concerns in the development of a medical chatbot are automated dialog design, structured data collection, and quality of patient interaction.

Two types of underlying chatbot frameworks were considered: dialog trees and generative text models. Dialog trees lay out a clear mapping between input text and chatbot behavior (i.e. data collection state and final diagnosis), which is desirable from the standpoint of model explainability. However, literature on such frameworks has shortcomings in end-to-end implementation. Some present dialog which is less "conversational" than hoped for, while others only accept Yes/No responses (except for an initial input query) [16], [17]. Additionally, dialog trees require review of medical literature and are commonly manually created on a cloud platform such as Alexa Developer or Google Dialogflow [18]. These concerns are partially addressed by novel machine learning and natural language processing-based systems which use a corpus of conversational data to generate human-like patient-doctor conversations [19], [20]. With these approaches, however, medical and conversational insights in the dataset are implicitly encoded in the final model, which presents challenges to structuring dialog and systematically collecting patient data. Finally, obtaining or leveraging existing medical data to suit a specific use case is a difficult task, and is a prime caveat of the automated modeling offered by machine learning systems. Following these considerations, the AIMAR robot's diagnostic chatbot uses a dialog-based on online "symptom checker" tools such as those provided by WebMD and Mayo Clinic, which enables the AIMAR robot to produce more human-like conversation while maintaining a standard list of questions to answer for each patient.

2.3.2 Medical Image Processing

A constant stream of research at the intersection of machine learning and computer vision has shown promise in real-time image classification and feature extraction tasks across a variety of domains. Team SCOPE's final product incorporated visual object recognition to healthcare-oriented items, providers, and patients. SCOPE used the IBM Watson cloud platform for VOR, applying image segmentation preprocessing to increase IBM's interpretive accuracy and allow the robot to differentiate between different visual features in the same frame.

More specifically, medical images have provided opportunities for healthcare providers to extract certain features regarding a patient's pathology and inform diagnostic and prognostic decisions. While technical machinery like an MRI machine or ultrasound monitor can provide relevant data for certain diagnoses, doctors can rely on features visible to the naked eye to inform other diagnoses. Automated diagnostics is a valuable tool in medicine, such as in the case of melanoma, the cancer of pigment-producing skin cells, where the prognosis is significantly impacted by the time of diagnosis. Unfortunately, false diagnoses account for nearly $700 million US dollars in medical costs, as general practitioners perform such diagnoses at an accuracy around 62-63% [21].

Computer vision has demonstrated promise in accurately diagnosing whether or not a given patient has a certain condition based on observable symptoms. In classifying skin lesions, many different algorithms and datasets have been combined to produce a variety of classification models. For instance, the largest online dataset of melanocytic skin lesion images is DermNet Skin Atlas, which previous research has used small subsets of for simple clustering and linear classifier techniques [21]. More advanced classification techniques may draw from medical intuitions; one feature which is commonly used to determine malignancy is the abruptness of a skin lesion [22]. Automated image processing can be used to determine exact abruptness levels for the cutoffs of skin lesions. This data can be used to train non-linear probabilistic classification models known as neural networks; if aspects of the neural network, such as the learning rate and input layer representation, are carefully designed, the resulting model yields significant improvement over the aforementioned linear models. Besides using the abruptness feature to diagnose skin conditions, other image processing techniques have yielded higher sensitivities and accuracy in correctly identifying the skin condition. The analysis of shape, color and texture features has led to specificity of 92.34% and sensitivity of 93.33%. These percentages of high specificity and high sensitivity essentially lead to higher precision diagnoses for the patient. An artificial neural network has even been created that analyzes both the geometries and textures of the lesions and yielded an accuracy rate of about 94% [22]. These image processing techniques all have promising results, as they yield an accuracy that rivals experts (90% detection rate) in the matter and greatly exceed that of general practitioners (62-63%) [21]. An expert system has been implemented to diagnose psoriasis, eczema, ichthyosis, acne, meningitis, measles, scarlet fever, warts, insect bites and stings, using a C Language Interpreted Production System (CLIPS). The expert system categorized these skin diseases into skin rashes without fever, skin rashes with fever, and skin infections [23]. A user could associate a number with the option displayed by the program; after diagnosis, the system would provide the user with an overview of the disease, causes, treatment, and general advice. This system does not require a significant amount of training, and can be a valuable aid in making diagnoses efficiently and effectively.

The current state of the art in image processing tasks, including skin lesion classification, is a deep learning architecture called a convolutional neural network (CNN). These types of models are enormously complex in many ways: the required dataset size, the number of parameters learned, and the time it takes to train. One of the earliest models to exceed dermatologist levels of accuracy was based on "Inception v3", a deep learning architecture developed by Google and trained using ImageNet, a dataset of approximately 1.28 million images from 1,000 (not necessarily medical) categories [24] [25]. To adapt Inception for skin lesion classification, researchers further tuned it on 127,463 skin lesion images curated from 18 different online medical repositories [26].

Team AIMAR intends to build on past research by implementing a similar pre-trained and fine-tuned skin lesion classifier as a module within a larger robotic system. In doing so, an end-to-end example of automated disease recognition was demonstrated, which could help medical professionals observe changes in the severity of a patient's condition and facilitate the personalization of care.

2.4 Health Data Format

A primary concern of handling patient data is security and privacy, considering the importance of the information being stored and transmitted. As of March 2003, the federal standard format for primary clinical messaging is Health Level 7 (HL7) V2.x, utilized in 90% of large hospitals [27]. The most common standard for transmitting health-related images uses Digital Imaging and Communications in Medicine (DICOM) files.

2.5 EEG Signals

Electroencephalography (EEG) refers to the real-time processing and visualization of detectable changes in neuronal electrical activity. EEGs monitor the fluctuations of electrical potentials from the cerebral cortex. Particular neurological disorders, including epilepsy, depression, and autism have been associated with specific frequencies or signatures. Although certain disease signatures are currently used for diagnostic healthcare purposes, the use of machine learning algorithms to recognize such signatures has been shown to increase diagnostic accuracy and identify new relationships between patterns in EEG data [28].

2.5.1 Future Direction of EEG Signal Processing Module

While digital processing of EEG signals has a variety of applications, including seizure detection and sleep state classification, team AIMAR decided not to include the development of the EEG signal processing module in the robotic prototype to focus more on the skin image analysis module as well as the medical chat bot module. Given that EEG signals are nonlinear and non-stationary, empirical mode decomposition and wavelet decomposition have been used to conduct data analysis on such signals [29]. Given the span of time dedicated to the development of each module, team AIMAR did not have sufficient time to learn the highly technical and complex algorithm classification methods for EEGs. EEGs are an important tool for studying the human brain and epileptogenic behavior and the abundance of available data makes them a good candidate for machine learning analysis. By further extracting and refining information from the wavelet coefficients, a reliable computer-aided diagnosis system can be created and added as a modular AIMAR component in a future iteration of the robot.

2.6 Conclusion of Literature Review

Team AIMAR reviewed the literature concerning research on robotic components, the use of EEGs for data collection and epileptic seizure detection, and image processing to analyze skin conditions. Based on this research, Team AIMAR utilized a robotic base that houses data collection modules and a user interface. A skin image processing module and a medical assistant chatbot module serve as proof-of-concept implementations. The skin image processing module detects skin conditions using an on-board camera combined with a deep learning model. The natural language interaction module allows the patient to have a conversation with the robot regarding any questions he/she may have regarding symptoms experienced. The AIMAR robot records conversations with the patient for future reference, as well as performs a basic medical diagnosis using conversation data and a medical condition database.

Many medically assistive robots exist, providing a basis for team AIMAR to build upon their physical capabilities.. They all specialize in different areas of expertise, but share a fundamental structure. Team AIMAR analyzed these robots for their implementation of the fundamental basics of their build structure, navigation techniques, and communication techniques with their patients to determine which method would best suit AIMAR. Through this research it was determined that the AIMAR robot should be a medium-sized mobile robot designed to be used in a medical setting and should include an arm capable of manipulating objects. Additionally, the robot should be capable of verbally interacting with patients, but should not use a visual avatar to try to imitate human interaction. The robot will also be designed to be a modular system, allowing for additional diagnostic tools to be integrated into the robot as they are developed and as needed by the user.

Team AIMAR also analyzed existing research on machine learning for skin image analysis. A variety of methods have been used in past research, ranging from feature detection tests to color detection tests. The feasibility and practicality of the creation of these methods have been used to inform Team AIMAR's decision to use a neural network to create an algorithm that allows for easier and quicker detection of skin lesions that could be potentially harmful. On top of this Team AIMAR did not want to compromise the accuracy of detection and hoped to even exceed the detection of skin ailments compared to general practitioners. Team SCOPE incorporated visual object recognition (VOR) to healthcare-oriented items, providers, and patients. The purpose of this research is to integrate a robotic element, natural language interaction, and image processing techniques through machine learning to suggest diagnoses, and aid medical professionals.

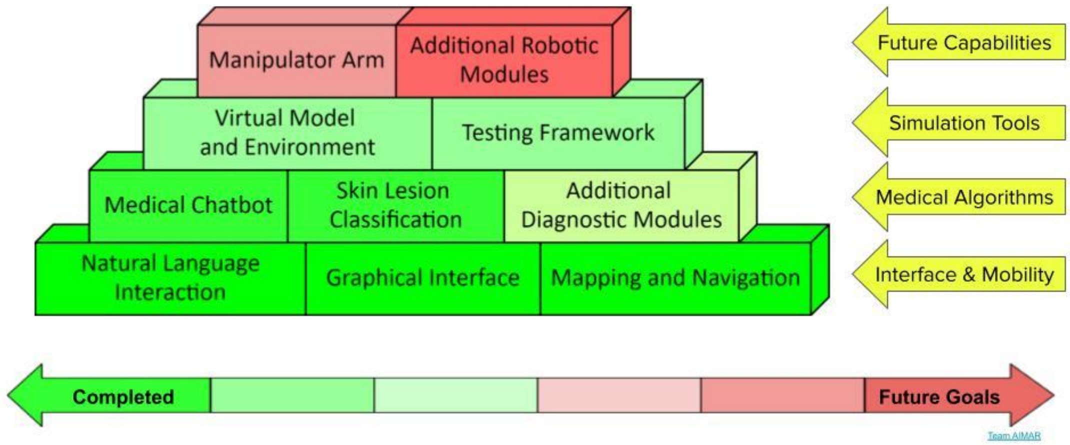

Figure 4. The module hierarchy of a fully functional AIMAR prototype. This hierarchy is presented as a heat map to demonstrate the level of completeness of different modular functional capabilities. Bright green shows completed modules and bright red shows potential future developments.

Chapter 3: Methodology

3.0 Overview

Team AIMAR's methodology aims to integrate a variety of technologies in robotics, machine learning, and natural language interaction. In doing so, automating certain tasks specific to the administration of a protocol to assist medical professionals. In the context of diagnosis, tasks could include the transport of a hardware appliance or specialized diagnosis (i.e searching for a specific body region and taking a measurement). The AIMAR robot can also be integrated into a clinical workflow to streamline the monitoring of an individual's status. Machine learning algorithms deployed on the AIMAR robot will identify clinically relevant patterns from feedback to facilitate diagnostic or prognostic judgment. It is important to note that the AIMAR robot itself does not make medical decisions, such as diagnosing conditions or prescribing medicine, but rather assists a trained professional in doing so by collecting data and providing analyses. Additionally, the AIMAR robot communicates with a central computer in its local network for storage and processing purposes. The robotic prototype incorporates two main independent modular components: a medical chatbot and a skin lesion classification model.

3.1 Robotic Design

The AIMAR robot intakes data from its independent modular components and provides information to assist patients and healthcare providers in the detection and diagnosis of disease. The design of a base model includes a user interface screen and modular components such as a robotic arm with 6 degrees of freedom and different end effectors for different uses, such as one with a smart camera at the end of it for image processing as addressed in 3.4, one for cleaning floors or disinfecting spaces, one for transporting meals, one for assisting the doctor with tools, etc. The arm includes sensors (e.g. shaft encoders, accelerometers, others as necessary) to maintain precise control over the robot while interacting with a user, as well as a force-torque sensor on the end effector to ensure safe interactions with the patient. Additional modules include the ball drive, LIDAR for navigation, voice interaction features such as a speaker and microphone, and an interactive screen, as addressed in sections 3.2 and 3.3. All these modules could be switched out easily as shown in the Computer Aided Design (CAD) below by simply sliding the top cylinder into the one below.

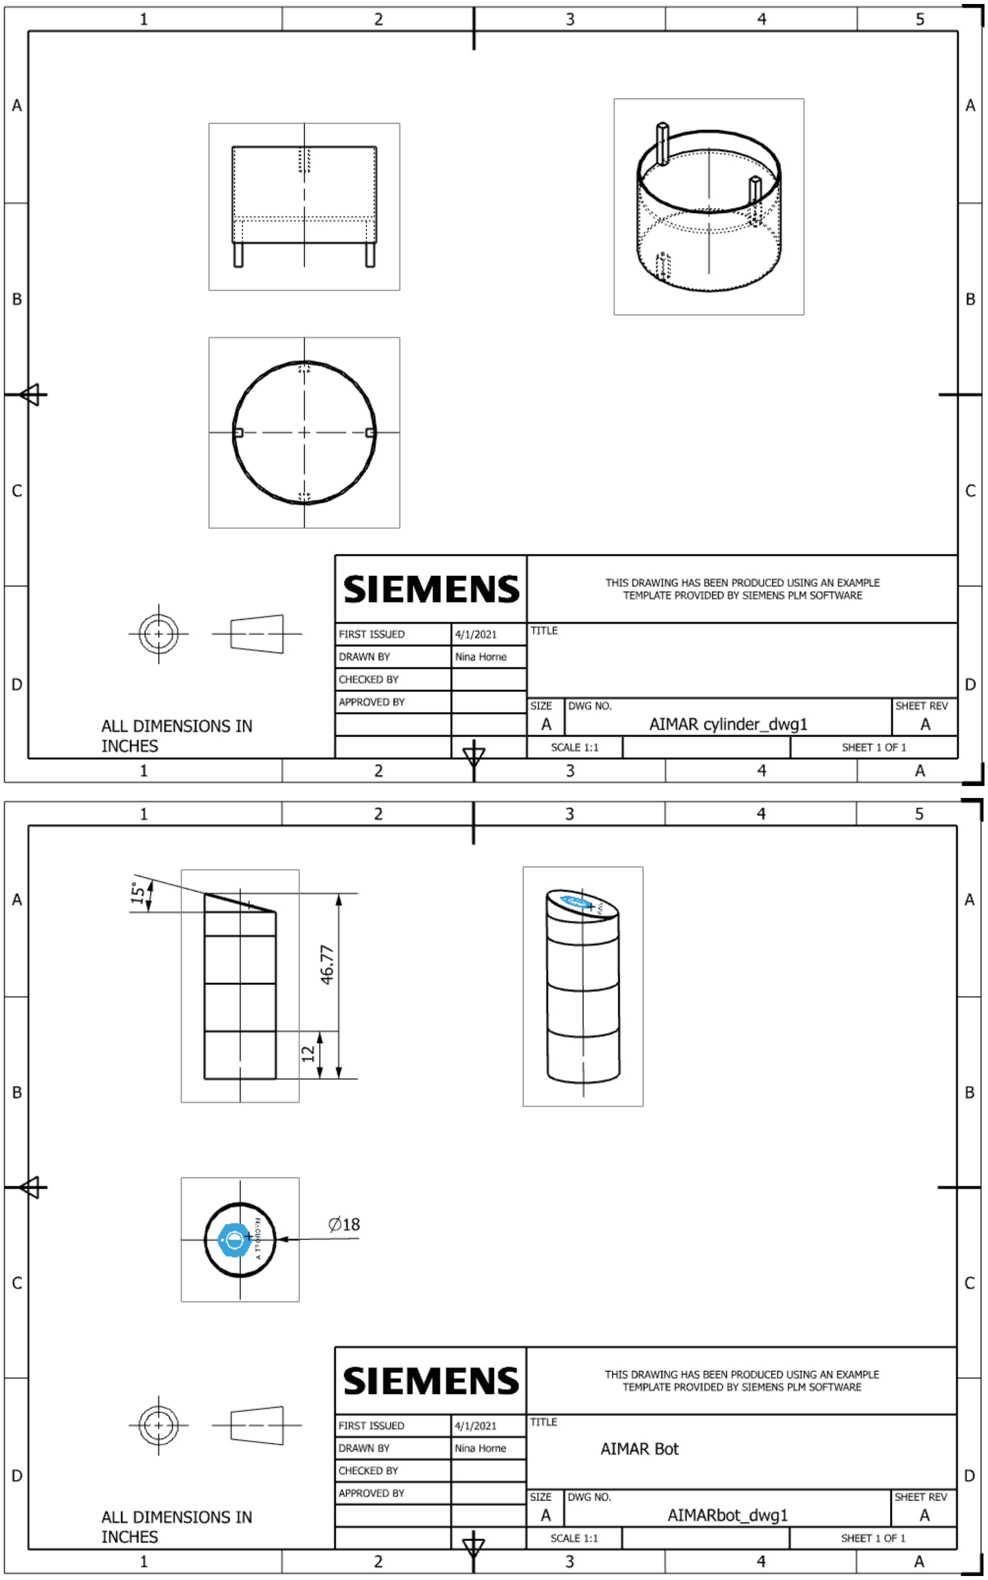

Figure 5: The CAD model for the AIMAR modular design.

3.2 Voice Assistant

AIMAR interacts with users through a voice assistant program which handles text-speech conversion and parsing. Early prototypes of AIMAR used Alexa Developer, but commercial solutions such as Alexa, Google Dialogflow and IBM Watson have dependencies on hardware or cloud backends. Therefore, Team AIMAR opted to use Mycroft, an open-source voice assistant for Linux systems. A Mycroft “Skill” is a collection of Python scripts that are triggered by pre-defined phrases, also called “intents”. A Skill that houses all of the AIMAR robot’s capabilities was implemented, allowing functionality from other modules (i.e. navigation, skin image classification, and patient database communication) to be controlled through voice commands and monitored through Mycroft’s responses. Mycroft also includes a programmable GUI which is used to display information from AIMAR’s diagnostic modules.

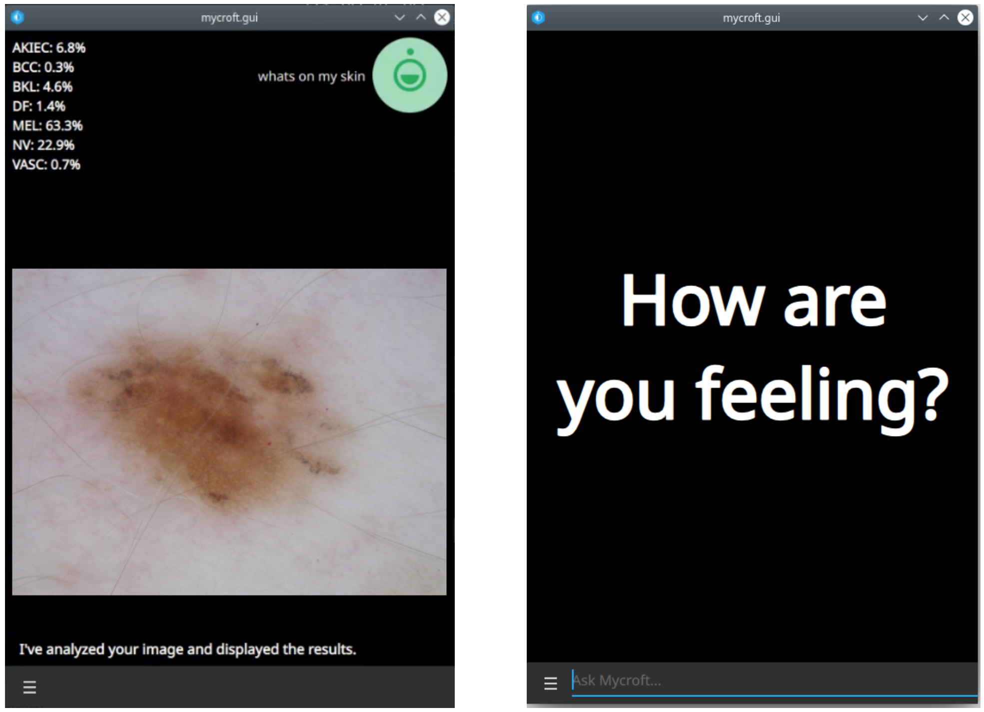

Figure 6. AIMAR's interface when the skin lesion module (left) and chatbot module (right) are in use.

3.3 Medical Chatbot

The AIMAR chatbot module uses data scraped from the Mayo Clinic's Symptom Checker website to generate symptom collection dialogs and perform basic medical text analysis. The Symptom Checker allows users to select a symptom (e.g. cough), then presents a series of multi-answer questions related to that symptom (e.g. Is there phlegm? If so, what color?). This question structure was adopted into a two-leveled dialog tree for the chatbot. First, the chatbot asks the patient for their primary symptom. Thereafter, as shown in Figure 7, the chatbot will ask questions specific to that symptom and identify keywords in the patient's response. However, the patient's response need not generate an exact match; the entire text can be used in a text-matching algorithm. Such an algorithm takes a pair of text corpora and computes a similarity score. In this case, the two corpora are the patient's responses and a summary of a condition's causes and symptoms (e.g. the WebMD info page for "common cold"). Running the algorithm on each condition produces a ranking of conditions by similarity. Three algorithms were tested using the inputs outlined above:

1. Both corpora are tokenized and represented as bag-of-words vectors. A similarity / dissimilarity score is calculated using an inner product of the two vectors.
2. Run a substring-similarity algorithm on each patient response vs. the medical condition corpus for a particular condition. The mean response score is the condition score.
3. Using BERT, a language model known as a "transformer", each sentence is encoded as a vector. The similarity scores between vectors are calculated identically to 1.

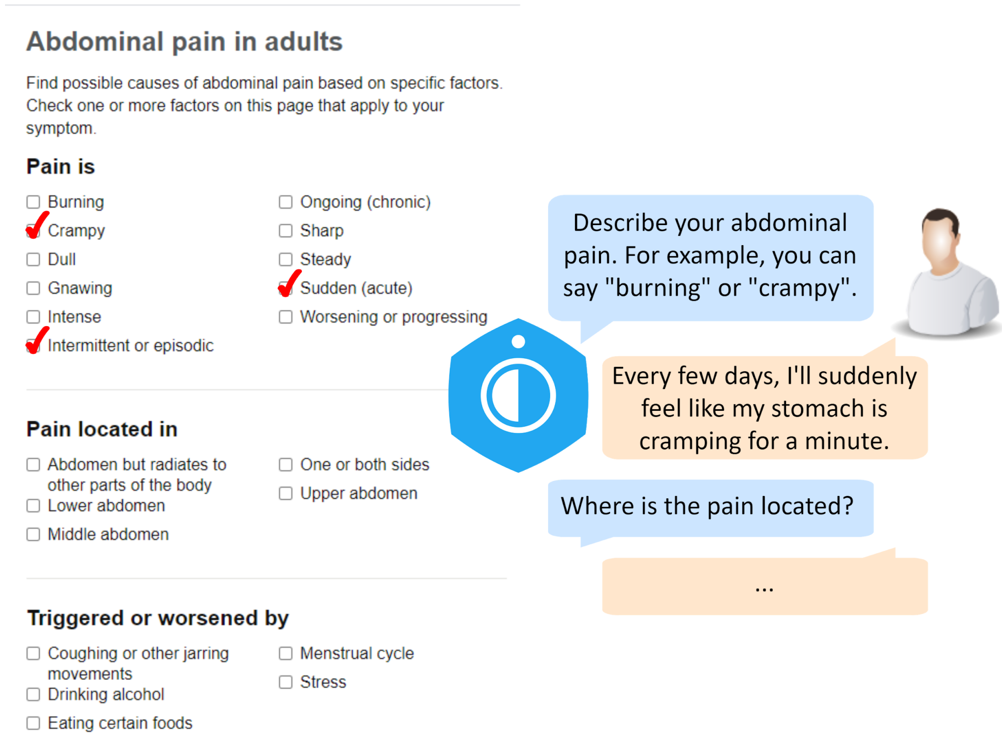

Figure 7. AIMAR chatbot's medical module. Diagnosis questions for abdominal pain in the Mayo Clinic Symptom Checker (left) are converted into full sentences in AIMAR chatbot's dialogue (right).

3.4 Image Processing

In deep neural image classification, a convenient and effective method is to start with a pre-trained model rather than one with randomly initialized parameters. For the base convolutional neural network, Team AIMAR used an architecture named ResNet18 was employed [31]. Next, the model is loaded in PyTorch, a Python deep learning library [30]. We then feed the ISIC 2020 Challenge Dataset into PyTorch to fine-tune and test the model. The publicly available ISIC 2020 dataset contains 33,126 images across 9 classification labels, including melanoma, melanocytic nevi, and keratosis [32]. The highest performing network trained on a 4-fold cross-validation split achieves an accuracy of 93%. The final model runs locally on the AIMAR robot, and may also be interchanged with other PyTorch models. Future modules in AIMAR can similarly use image classification for other conditions.

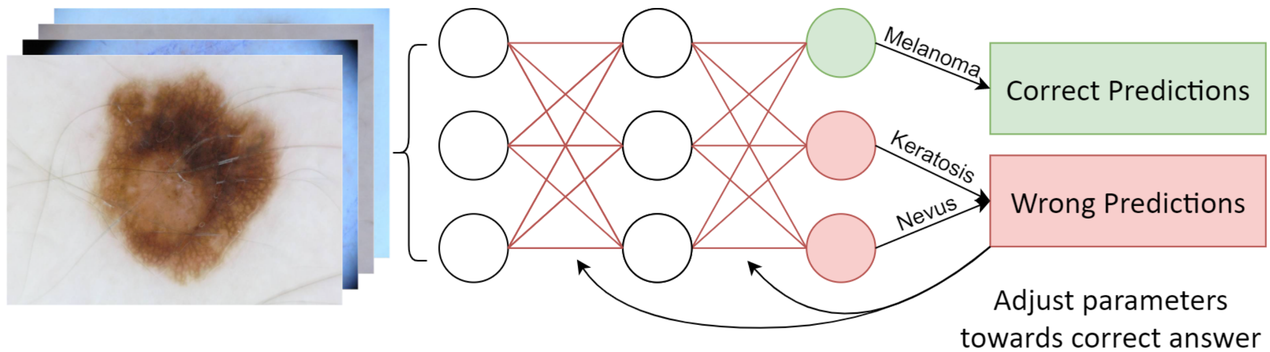

Figure 8. The procedure of training a neural network to classify a dataset of skin lesion images. The final network is deployed on AIMAR chatbot as a diagnostic module.

3.4 Integration

3.4.1 Integration of Subsystems

To support the software outlined above, the Raspberry Pi 3B+ (included in the TurtleBot3 package) was replaced with a Raspberry Pi 4 installed with Ubuntu Server 20.04. Instructions for installing the individual software components are in the technical documentation. By integrating these components, the AIMAR robot will be able to accomplish many tasks. As patients arrive at the healthcare facility, the AIMAR robot can serve as a receptionist and check on patient information. The AIMAR robot can also serve as a nurse by directing patients and collecting medical data through a conversation, camera, or another collection module. Once data collection is complete, the AIMAR robot can directly assist doctors by providing a preliminary analysis of the patient's health condition.

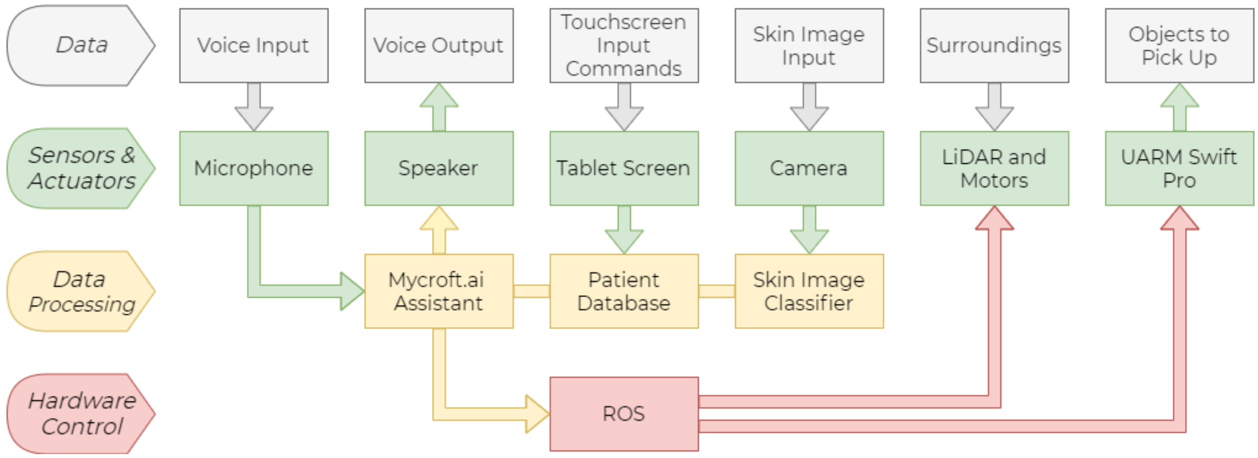

Figure 9. A flowchart showing how the AIMAR system handles interaction with its environment.

3.4.2 Robotic Testing Framework

A virtual testing framework was developed to support the proposed modularity of the AIMAR robot system, as well as verify the integration of components on a desktop computer before deployment to the robot. The framework is a suite of Python scripts that tests the concurrent behavior of the AIMAR robot's diagnostic modules, interface, and physical movement. To accomplish this, the Gazebo software was used to simulate a virtual robot and environment. Simultaneously, all backend software components were run by Team AIMAR. Finally, inputs were simulated that the AIMAR robot would normally detect through its hardware or manual signal, such as an image to feed into the skin lesion classifier or a voice command to check up on a patient. Each script verifies conditions associated with a certain task; for example, the "patient checkup" procedure test script simulates two sets of voice commands: the first is a doctor telling the robot to perform the checkup, and the second is the set of patient responses to the robot's questions. In conjunction, the script verifies if the AIMAR robot arrived at the correct room and if the final diagnostic and recorded patient data matched their expected values. As a whole, the framework allows developers to maintain consistent behavior in the AIMAR robot as modules are added. The framework will still be compatible even if many of the underlying implementations are changed, as long as the core components of the software system (namely ROS and Mycroft) and their input/output structures are kept intact. Furthermore, it outlines a set of tasks that each version of AIMAR is capable of.

3.4.3 Subteams

Team AIMAR had three sub-teams: Team Brain, Team Linguini, and Team Robo. Team Brain focused on skin image analysis and classification, algorithms, and applications of machine learning. Team Linguini focused on natural language processing and the user interface. Team Robo focused on robot hardware, navigation, and the operation of the robotic arm. Team members had the opportunity to work on objectives outside of their subteam to learn more about different topics, and collaborated with each other. The success of each subgroup led to the integration of the three components to ensure the proper interaction among the modules of the robot.

3.5 Conclusion

The goal of Team AIMAR was to create a robotic unit that can streamline data collection and interact with humans in a medical setting. These capabilities could improve the quality and efficiency of healthcare by reducing the time and cost of seeking medical attention, and also by reducing the likelihood of a misdiagnosis that could lead to ineffective treatment. Team AIMAR conceptualized a medium-sized mobile robot with a modular design where each added module increases the functionality and diagnostic capabilities of the robot. Team AIMAR implemented a mobile robotic base, a robotic arm, and two diagnostic modules: a medical chatbot module, and a skin image classification module.

The proof-of-concept of this robotic design was done using a TurtleBot3 and a UArm Swift Pro. The Raspberry Pi included on the robot was upgraded to a Raspberry Pi 4 to support the robot's diagnostic modules. The TurtleBot3 includes a LIDAR sensor, which allows the robot to create a map of any new environment that it is placed in and then use this map to navigate its environment. This feature is important to the functionality of the robot because it allows for the robot to navigate a medical office on its own without the need for a human guide. While this TurtleBot serves as the proof-of-concept, Team AIMAR has also designed a CAD model that more closely resembles the final robot that they envision being used. This robot is approximately 4 feet tall and includes a ball drive system that allows it to move in any direction so that it can navigate more easily around obstacles and in tight spaces. The robot is not designed to imitate a human but is rather designed to support patient trust with its functional, non-intimidating robotic design.

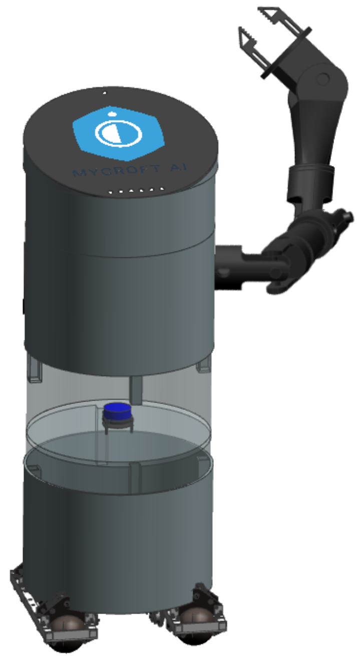

Figure 10: A virtual AIMAR bot is shown. This virtual model is used in the testing framework.

Team AIMAR integrated this TurtleBot with the natural language interaction tool Mycroft. This tool was used to add functionality to the base robot by allowing for voice command navigation of the robot. Through this feature, pre-defined phrases can be used to verbally instruct the robot where to go and the robot will correctly execute these verbal commands. Mycroft was also used to implement the robot's medical chatbot module. This medical chatbot uses symptom data from Mayo Clinic's Symptom Checker to provide a list of possible diagnoses. A patient can then verbally respond to the AIMAR robot's prompts about their symptoms, and the robot will record this data to match the patient with possible conditions. While this is adapted from an existing online tool, this implementation is beneficial as it allows patients to give free form responses, the entirety of which are used to inform the output, rather than forcing the patient to select from a predefined list of symptoms as is done in Mayo Clinic's Symptom Checker.

Team AIMAR also deployed a skin image classification module to run locally on the TurtleBot3's Raspberry Pi 4B+ processing unit. This module was created using the Python deep learning library PyTorch, which allows it to be interchanged with other PyTorch models that either achieve greater accuracy or include a wider variety of conditions. The model uses a ResNet18 base, fine-tuned using the ISIC 2020 Challenge Dataset of 33,126 images across 9 skin condition labels. The highest performing network achieved an accuracy of 93%.

While the AIMAR robot is not intended to be used as a sole form of diagnosis, the features that are included on the bot can be a valuable tool to guide doctors to a correct diagnosis. While doctors and nurses are busy carrying out other vital tasks, the AIMAR robot can be sent to begin collecting necessary information from the patient. The information collected by the medical chatbot can be used to provide doctors with a preliminary list of concerns that can be used to further guide the doctor's questions or the testing that they wish to perform. Additionally, when the patient has a skin concern, the robot can take an image of the affected area and provide a likely diagnosis. The doctor would then use their expert opinion to either verify the diagnosis or provide a diagnosis of their own. Through this process, patients may potentially receive a faster and more accurate diagnosis, and medical professionals can spend more time performing tasks that only they can do instead of collecting routine information. As such this robot has the potential to streamline diagnostics, lower costs of medical care, reduce unnecessary time spent in a doctor's office, and improve patient outcomes.

3.5.1 Innovations

The innovative contribution of Team AIMAR was in the integration of the various technologies on the robot and in the modular design of the robot. While Team AIMAR made use of existing technologies such as the TurtleBot3, UArm, Mycroft, and PyTorch, these technologies were brought together in a unique way to create a robotic device with greater functionality than any of these tools have on their own. Each technology that is used is vital to the functionality of the robot at some point in time between when the robot is sent to see a patient and when the robot provides its diagnosis to the doctor. Additionally, it is the integration of these components that serve as the base of the robot's modular system. This modular system is innovative because it allows for easy additions to the AIMAR robot that reduce possible limitations of a medical robot by allowing new software or hardware components to be incorporated to serve a wide variety of functions. This modular system also gives the choice for medical offices to decide which features they need, and only connect the modules that are of use to them.

3.5.2 Future Directions

One of the key features of the AIMAR robot is its modular design. This modular design allows for the functionality of the robot to be increased by future researchers by adding new modules with diagnostic tools for conditions that have not been included in the initial design. One module that Team AIMAR envisions being implemented on the robot in the future is an EEG module. Such a module would include an EEG headset and could use artificial intelligence to detect seizures, diagnose neurological disorders, or classify sleep stages. Team AIMAR also envisions assistive modules being developed by future researchers that make use of the robot arm to manipulate objects to carry out tasks for either doctors or patients.

Another feature that could be added in the future is the ability to securely access and edit a patient's existing health records. While the AIMAR robot keeps its own records of the information that it collects from patients it would be ideal if relevant information was stored in a standard format and automatically updated in the patient's records as this would avoid the need for a doctor or nurse to update the records themselves. Being able to access existing health records would also allow for an algorithm to be designed that takes into account the patient's known medical history when attempting to diagnose them. Additionally, the privacy of a patient's health data is a major concern, so it would be best if the information collected from the patient is stored in a secure health record system.

Appendices

A.1 Timeline

Fall 2018:

- Complete background research.
- Study Team SCOPE's project and decide which elements will be incorporated into the prototype of AIMAR.
- Determine what medical conditions to focus on and obtain the necessary equipment.
- Identify and possibly obtain viable medical datasets.
- Determine which programming languages and libraries/packages are best suited to the project.

Spring 2019:

- Complete research proposal.
- Decide if doctors will be surveyed about their views of what AIMAR should be. If a survey is planned, obtain IRB approval and perform steps necessary to collect data.
- Start physically building the robot and implement any pieces of Team SCOPE's project that are being kept.
- Begin prototyping the basic robot design and incorporating moving components.

Fall 2019:

- Using the baseline robot design, begin incorporating the image processing sensors.
- Investigate the diagnostic capabilities of different artificial intelligence algorithms and how they can be integrated with the data collection methods.
- Collecting results and presenting findings at Junior Colloquia.
- Develop a detailed thesis outline.
- Research various conferences as a possible forum for scholarly research.
- Produce a first draft of the thesis.

Spring 2020:

- Finalize a prototype that includes the sensors that measures skin image data and image processing.
- Begin data analysis and classification by testing various artificial intelligence algorithms and comparing their diagnostic efficacy.
- Obtain expert feedback regarding the thesis to evolve the future direction of the project.

Fall 2020:

- Analyze the findings of the project.
- Make another draft of the thesis which includes results and discussion of the project.
- Obtain another round of expert feedback regarding the thesis to evolve the future direction of the project.

Spring 2021:

- Update the thesis to include the suggestions made by the experts who attended the thesis conference.
- Defend the final thesis at the Team Thesis Conference.

A.2 Budget

Team AIMAR was provided with an $1800 budget from the Gemstone Program. Additionally, in Spring of 2019, Team AIMAR received the 2nd prize Gemstone Library Award amount of $1000. This totaled to a budget of $2800 for the project, of which $151.99 was allocated for the expenses as listed in Table A2.

Resources provided by IPST LCV Lab

Equipment	Use
TurtleBot3 Burger, Waffle	A mobile base which houses all other physical components of the robot.
uArm Swift Pro	Robotic arm with 3-4 degrees of freedom. Will be used for manipulations.
Amazon Alexa	Used for natural language interaction prototyping.

Table A1: Equipment provided by the IPST LCV lab and its usage.

Expenses

Item Description	**Item Price**	**Qty.**	**Total Purchase Cost**
Battery	$ 49.90	2	$ 99.80
DVI to HDMI Cable	$ 6.99	1	$ 6.99
Pi Cam Frame	$ 12.90	2	$ 25.80
16 gb Micro SD Card	$ 5.46	1	$ 5.46
Raspberry Pi Speaker	$ 7.99	1	$ 7.99
Raspberry Pi Microphone	$ 5.95	1	$ 5.95
		Total Expenses:	$ 151. 99

Table A2: Outlined expenses for Team AIMAR.

A.3 Additional Resources

Team Website: https://umd-aimar.com/

Github Repository (Code): https://github.com/UMD-AIMAR

Thesis Defense recording:

https://drive.google.com/file/d/1JbjeBbXCfezs5G3sh4fMwnjqszvbVzVu/view?usp=s
haring

Thesis Defense presentation:

https://drive.google.com/file/d/1IS3cIeE2MEE-mhJky5vUi2eOJd9rgS-F/view?usp=s
haring

Uarm How To:

https://docs.google.com/document/d/1nLnjW-WwqBMz6RlxzHTPAPMKNusWprZ6l
AuOPG39ZOM/edit

Demonstration Videos (on Youtube):

https://www.youtube.com/channel/UCieSv8MWX5RHF4hU0_nHUug

A.4 Hardware

Uarm Swift Pro: https://www.ufactory.cc/pages/uarm

TurtleBot3, includes LIDAR: https://www.turtlebot.com/

Raspberry Pi: https://www.raspberrypi.org/

PiCam: https://projects.raspberrypi.org/en/projects/getting-started-with-picamera

References

[1] J. Miller, “Paying for health care with time,” *Harvard Gazette*, Oct. 05, 2015. https://news.harvard.edu/gazette/story/2015/10/paying-for-health-care-with-time/ (accessed Nov. 16, 2020).

[2] C. Slyder, “Student startup plans to offer app that can detect concussions,” 2016. https://www.purdue.edu/newsroom/releases/2016/Q4/student-startup-plans-to-offer-app-that-can-detect-concussions.html (accessed Oct. 19, 2020).

[3] TechCrunch, “Brightlamp wants to use AI to spot concussions,” *TechCrunch*, 2017. https://social.techcrunch.com/2017/05/16/brightlamp-wants-to-use-ai-to-spot-concussions/ (accessed Jan. 23, 2021).

[4] John Bachkosky VI, Alexandra Boukhvalova, Kevin Chou, William Gunnarsson, James Ledwell, Brendan McTaggart, Xiaoqing Qian, Nicholas Rodgers, John Shi, Jason Yon, Anil Deane, “SENSORY COMPUTING AND OBJECT PROCESSING ENTITY (SCOPE): ASSISTIVE ROBOTICS FOR HEALTHCARE,” p. 101, 2017.

[5] M. Montemerlo, J. Pineau, N. Roy, S. Thrun, and V. Verma, “Experiences with a mobile robotic guide for the elderly,” in *Eighteenth national conference on Artificial intelligence*, USA, Jul. 2002, pp. 587–592, Accessed: Oct. 19, 2020. [Online].

[6] M. Hans, B. Graf, and R. D. Schraft, “Robotic home assistant Care-O-bot: past-present-future,” in *11th IEEE International Workshop on Robot and Human Interactive Communication Proceedings*, Sep. 2002, pp. 380–385, doi: 10.1109/ROMAN.2002.1045652.

[7] D. Portugal, L. Santos, P. Alvito, J. Dias, G. Samaras, and E. Christodoulou, *SocialRobot: An interactive mobile robot for elderly home care*. 2015, p. 816.

[8] M. Kumagai, “Development of a ball drive unit using partially sliding rollers — An alternative mechanism for semi-omnidirectional motion —,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 3352–3357, doi: 10.1109/IROS.2010.5651007.

[9] M. Crawford, “How Many Axes Does Your Robot Need?,” 2017. /topics-resources/content/Many-Axes-Does-Robot-Need (accessed Oct. 19, 2020).

[10] Y. Pan and A. Steed, “A Comparison of Avatar-, Video-, and Robot-Mediated Interaction on Users’ Trust in Expertise,” *Front. Robot. AI*, vol. 3, 2016, doi: 10.3389/frobt.2016.00012.

[11] Y. Hiroi and A. Ito, “Are bigger robots scary? —The relationship between robot size and psychological threat—,” in *2008 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, Jul. 2008, pp. 546–551, doi: 10.1109/AIM.2008.4601719.

[12] A. K. Pandey and R. Gelin, “A Mass-Produced Sociable Humanoid Robot: Pepper: The First Machine of Its Kind,” *IEEE Robot. Autom. Mag.*, vol. 25, no. 3, pp. 40–48, Sep. 2018, doi: 10.1109/MRA.2018.2833157.

[13] S. Korol, "Hospital nurse shift patterns and collaboration — Strategies for improvement," 2019. https://www.beckershospitalreview.com/quality/hospital-nurse-shift-patterns-and-collaboration-strategies-for-improvement.html (accessed Apr. 01, 2021).
[14] C. Dreisbach, T. A. Koleck, P. E. Bourne, and S. Bakken, "A systematic review of natural language processing and text mining of symptoms from electronic patient-authored text data," *Int. J. Med. Inf.*, vol. 125, pp. 37–46, May 2019, doi: 10.1016/j.ijmedinf.2019.02.008.
[15] H. Kazi, B. S. Chowdhry, and Z. Memon, "MedChatBot: An UMLS based Chatbot for Medical Students," *Int. J. Comput. Appl.*, vol. 55, no. 17, pp. 1–5, Oct. 2012, doi: 10.5120/8844-2886.
[16] J. Beaudry, A. Consigli, C. Clark, and K. J. Robinson, "Getting Ready for Adult Healthcare: Designing a Chatbot to Coach Adolescents with Special Health Needs Through the Transitions of Care," *J. Pediatr. Nurs.*, vol. 49, pp. 85–91, Nov. 2019, doi: 10.1016/j.pedn.2019.09.004.
[17] K. Rarhi, A. Bhattacharya, A. Mishra, and K. Mandal, "Automated Medical Chatbot," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 3090881, Dec. 2017. doi: 10.2139/ssrn.3090881.
[18] N. Rosruen and T. Samanchuen, "Chatbot Utilization for Medical Consultant System," in *2018 3rd Technology Innovation Management and Engineering Science International Conference (TIMES-iCON)*, Dec. 2018, pp. 1–5, doi: 10.1109/TIMES-iCON.2018.8621678.
[19] K. Oh, D. Lee, B. Ko, and H. Choi, "A Chatbot for Psychiatric Counseling in Mental Healthcare Service Based on Emotional Dialogue Analysis and Sentence Generation," in *2017 18th IEEE International Conference on Mobile Data Management (MDM)*, May 2017, pp. 371–375, doi: 10.1109/MDM.2017.64.
[20] G. Zeng *et al.*, "MedDialog: Large-scale Medical Dialogue Datasets," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, Nov. 2020, pp. 9241–9250, doi: 10.18653/v1/2020.emnlp-main.743.
[21] A. Bhattacharya, A. Young, A. Wong, S. Stalling, M. Wei, and D. Hadley, "Precision Diagnosis Of Melanoma And Other Skin Lesions From Digital Images," *AMIA Jt. Summits Transl. Sci. Proc. AMIA Jt. Summits Transl. Sci.*, vol. 2017, pp. 220–226, 2017.
[22] R. Erol, M. Bayraktar, S. Kockara, S. Kaya, and T. Halic, "Texture based skin lesion abruptness quantification to detect malignancy," *BMC Bioinformatics*, vol. 18, no. Suppl 14, Dec. 2017, doi: 10.1186/s12859-017-1892-5.
[23] S. Abu-Naser and A. Akkila, "A Proposed Expert System for Skin Diseases Diagnosis," Oct. 2008.
[24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," *ArXiv151200567 Cs*, Dec. 2015, Accessed: Apr. 29, 2021. [Online]. Available: http://arxiv.org/abs/1512.00567.
[25] O. Russakovsky *et al.*, "ImageNet Large Scale Visual Recognition Challenge," *ArXiv14090575 Cs*, Jan. 2015, Accessed: Apr. 29, 2021. [Online]. Available: http://arxiv.org/abs/1409.0575.

[26] A. Esteva *et al.*, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, Art. no. 7639, Feb. 2017, doi: 10.1038/nature21056.
[27] P. Aspden, J. M. Corrigan, J. Wolcott, and S. M. Erickson, Eds., *Patient Safety: Achieving a New Standard for Care*. Washington (DC): National Academies Press (US), 2004.
[28] O. Faust, U. R. Acharya, H. Adeli, and A. Adeli, “Wavelet-based EEG processing for computer-aided seizure detection and epilepsy diagnosis,” *Seizure*, vol. 26, pp. 56–64, Mar. 2015, doi: 10.1016/j.seizure.2015.01.012.
[29] M. Kaleem, A. Guergachi, and S. Krishnan, “EEG seizure detection and epilepsy diagnosis using a novel variation of Empirical Mode Decomposition,” in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Jul. 2013, pp. 4314–4317, doi: 10.1109/EMBC.2013.6610500.
[30] A. Paszke *et al.*, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” *ArXiv191201703 Cs Stat*, Dec. 2019, Accessed: Jan. 24, 2021. [Online]. Available: http://arxiv.org/abs/1912.01703.
[31] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” *ArXiv151203385 Cs*, Dec. 2015, Accessed: Jan. 24, 2021. [Online]. Available: http://arxiv.org/abs/1512.03385.
[32] V. Rotemberg *et al.*, “A Patient-Centric Dataset of Images and Metadata for Identifying Melanomas Using Clinical Context,” *ArXiv200807360 Phys.*, Aug. 2020, Accessed: Jan. 24, 2021. [Online]. Available: http://arxiv.org/abs/2008.07360.