



# DESIGNING THE NEXT GENERATION OF INTELLIGENT PERSONAL ROBOTIC ASSISTANTS FOR THE PHYSICALLY IMPAIRED

Basit Ayantunde

Department of Mechanical Engineering,  
University of Ilorin, Kwara, Nigeria.

Fadlullah Olawumi

Department of Mechanical Engineering  
University of Ilorin, Kwara, Nigeria.

Jane Odum

Department of Computer Science,  
University of Ilorin, Kwara, Nigeria.

Joshua Olalekan

Department of Electrical Engineering  
University of Ilorin, Kwara, Nigeria.

**Abstract**—The physically impaired commonly have difficulties performing simple routine tasks without relying on other individuals who are not always readily available and thus make them strive for independence. While their impaired abilities can in many cases be augmented (to certain degrees) with the use of assistive technologies, there has been little attention to their applications in embodied AI with assistive technologies. This paper presents the *modular framework, architecture, and design* of the *mid-fidelity prototype* of *MARVIN*: an artificial-intelligence-powered robotic assistant designed to help the physically impaired in performing simple day-to-day tasks. The prototype features a trivial locomotion unit and also utilizes various state-of-the-art neural network architectures for specific modular components of the system. These components perform specialized functions, such as automatic speech recognition, object detection, natural language understanding, speech synthesis, etc. We also discuss the constraints, challenges encountered, potential future applications and improvements towards succeeding prototypes.

**Keywords:** embodied AI, automatic speech recognition, object detection, natural language understanding, speech synthesis.

## I. INTRODUCTION

There have been tremendous discoveries and breakthroughs in the field of artificial intelligence and robotics in the past few decades, and technology has accordingly grown from being seen as eliminating the need for humans to augmenting our abilities and in the process, making life easier. Consequently, this poses a challenge in bringing these technologies to their useful applications in our daily lives.

A reported estimate of 1 billion people live with one or more disabilities around the globe, in accordance with a population estimate of 9.6 billion people (World Health Organization and others, 2011). The aged also possess slowly

degrading physical functionalities and functional impairments which varies from one individual to another (Janssen, Samson, and Verhaar, 2002). The global population aged 60 years and above were estimated to be 962 million in 2017 and expected to double by 2050, during which it is projected to reach nearly 2.1 billion (United Nations, 2017).

These impairments make performing day-to-day tasks difficult for these individuals and thus affect the quality of their lives, i.e. hygiene, feeding, alimentations, work, etc (Valenã, Santos, Lima, Santana, and Reis, 2017). The aged phase of life is often made more difficult due to social factors such as subjectivity to discrimination, constructed social exclusion and stigmatization (Rapolienė, 2015).

Physical impairments can be caused by natural aging, diseases (hypertension, glaucoma, cataracts, etc.), accidents, injuries, natural impairments since birth, etc.

While assistive technologies addressing some of these physical impairments have been undergoing rigorous research and development for decades, they mostly take the form of aids and are often incorporated or attached to the impaired individual's body, which in many cases do not suffice well to enable them to perform their daily activities effectively. Some previous assistive technologies include electric wheelchairs, electroencephalography, digital walking sticks, bone-anchored hearing aid (Wikipedia contributors, 2019b), mobility scooters (Wikipedia contributors, 2019f), cochlear implants, prosthesis, text-to-voice wands, intelligent personal digital assistants, etc.

Current IPDAs such as *Alexa*, *Siri*, and the *Google Assistant* are not readily integrable and customizable to run on specialized hardware especially due to the closed-source natures of their architectures and software systems. Other challenges include immobility, non-interactivity with the physical world around them, and the few user interfaces



they readily support leading to inaccessibility for people with the counter-impairments (relative to the user interface).

Our approach tries to help them perform some of these tasks. We propose a simple and modular robotic system utilized toward this end. We also propose an Intelligent Personal Robotic Assistant (IPRA) architecture, which in contrast to the Intelligent Personal Digital Assistant architectures (IPDA) (Sarikaya, 2017), is highly modular, readily integrable with various hardware components for robotics applications, and highly prioritizes the modular, multimodal and fusible nature of the user interfaces.

While the current prototype of the robotic system is basic, does not yet provide support for dynamic environments and specialized use cases, this only serves as a mid-fidelity prototype and at the same time, a proof-of-concept for our framework and architecture, and serves as a basis for future prototypes. The system can be easily built on, improved and used in coordination with pre-existing assistive technologies.

Our current system prototype utilizes a mixture of lightweight state-of-the-art neural network architectures towards a problem decomposition approach, which in turn makes it more reliable in comparison to fully End-to-End (E2E) AI robotic systems that are in many cases questionable due to the non-interpretability of the representations of the world around them, thus making their decisions highly difficult to probe.

Running modern state-of-the-art neural network architectures often requires huge computation power and large memory requirements due to the hundreds of millions and often billions of parameters they possess and in effect need to be deployed on costly hardware devices to meet real-time requirements.

While it is noteworthy that for the affordability and general acceptability of these systems, there is an impending need for them to be deployable on cost-effective hardware platforms often consequently possessing constrained resources as these cost money. This is addressed mostly by building a series of prototypes, and balancing resource consumption against the quality of solutions provided by these systems (Lugaresi et al., 2019) whilst paying attention to cost. Nonetheless, a common practice is to start with a simple system with optimal resources, incrementally add more features and accordingly add more resources as required.

Also worth noting is the nature of complexity that arises from designing these systems, due to the non-deterministic timing of events often received by the system and real-time constraints frequently imposed in the process. In a lossless data processing pipeline, backpressure is always avoided and monitored, this is worsened when the designated hardware doesn't have enough resources to counteract the effects of these liable limitations. Our framework provides a few basic components to address these complexities and limitations.

The proposed architecture also highly prioritizes the mod-

ular, multimodal, and fusible nature of the user interfaces to enable it to serve a wide variety of people with/without disabilities, which most current IPDA architectures do not prioritize but instead possess non-modular unimodal user interfaces with limited interactivity and accessibility. The framework's flexibility also allows for the easy deployment of engineered, E2E, and reinforcement learning models within the architecture. The multimodal nature and fusible user interfaces allow the system to fuse and contextualize semanticized sampled data such as visual cues, gestures, and speech into higher-level semantic data. Another benefit of its multimodal nature is that it gives the system a better understanding of context.

## II. CAPABILITIES

Listed below are the assistive capabilities (Skills) implemented into the current prototype of the robot:

- Facial Recognition
- Time Querying
- Schedule Organization
- Phone Call Initiation
- Object Finding
- People Finding
- Intruder Detection
- Visually Impaired Guidance
- Smart Home Devices Control (Bulbs, Fans, etc.)

## III. HARDWARE SPECIFICATIONS

- Acrylic Plastic: Used for the robot's physical enclosure. This was elected due to inaccessibility to 3d Printers.
- Raspberry Pi 3B: 1.2GHz, Quad-core Broadcom BCM2837 64bit CPU, 1GB RAM (Raspberry Pi Foundation, n.d.). This is the main computation powerhouse. This has a minimal computational power that is good enough for the initial prototype.
- Raspberry Pi Camera: 0.5MP, RGB-888. For capturing images used in object detection and facial recognition.
- STM32 Nucleo-F302R8T6: 72MHz, 16KB SRAM, 32 KB Flash (STMicroelectronics, n.d.). For controlling the obstacle detection unit and other physical components.
- HC-SR04 Ultrasonic Sensor: For obstacle detection.
- MG996R Brushed DC Servo Motor: 120° range, 0.14s/60° @ 6V (Electronicos Caldas, n.d.). This tilts the ultrasonic sensor during beamforming.
- L293D Dual H-bridge DC Motor Driver.
- Generic Mono-channelled 16-bit USB Microphone: For keyword spotting and speech interaction.
- Generic USB Speaker: For audio and speech output.
- 74HC-595 8-bit Shift Register: For parallel-series output.
- 5V and 12V Brushless DC Fans: For forced-air convection cooling of the whole system.



- Heat Sinks: For convection cooling of the driver chips, microcontrollers, and microcomputers.

#### IV. SOFTWARE ARCHITECTURE

In this section, we describe the layers in our IPRA architecture. The architecture is shown in Figure (1).

##### A. Input Devices Layer

This is part of the system that accesses/interfaces various hardware components of the system needed to get necessary input data samples from the environment. This is readily available in the Linux kernel (implemented via *udev*) with various drivers that manage the hardware and resources used in the system.

##### B. Input Manager

This collects the various data samples received from the Input Devices Layer and dispatches them to the appropriate User Interfaces. This is beneficial especially when there are multiple User Interfaces in use and a need to share the incoming raw samples or route them to the appropriate User Interfaces arises.

##### C. User Interface

This processes the data samples collected from the environment or generated by the system to add various additional semantics required in its utilization. For example in a voice interactivity setting, it is often required to merge various incoming audio samples from microphone arrays into a frame containing multi-channel samples. This is also useful in a setting where we need to merge/fuse camera and microphone samples. They can take the form of a neural network, perform simple signal processing and/or chained together to obtain different levels of semantics.

##### D. Semantic Interpreter

This models the semanticized data samples and interprets them into commands represented as Skill identifiers and its Entities that are passed to the Skill Manager. This can be implemented as an Automatic Speech Recognizer, Gesture Recognizer, Sign Language Recognizer, etc. depending on the nature of the incoming data and end-user-interface.

##### E. Skill Manager

This part of the system dispatches the recognized Skills, evaluates the Skill's Entities (tokens/parameters) and ensures the Skill's Entities and requirements are completed and fulfilled by interacting with the user via the output User Interface (if any).

##### F. Skill Registry

This part of the system that contains the Skills' definitions and requirements.

##### G. Skill

This unit defines how the robot is meant to execute the semanticized instruction received through the User Interface. This can also be developer or end-user defined via a generic though restricted interface. Skills are similar to Alexa Skills and Google Assistant Actions.

#### V. FRAMEWORK

Our framework increases the speed of developing these systems and also makes them easier to conceptualize since they perform lots of real-time processing and involve complex data-flows often in a multi-threaded fashion. These components are loosely coupled to allow for flexibility, composability, and suitability for unit testing. The system is represented as a directed acyclic graph consisting of Nodes with multi-input/output channels termed Streams. Some of the implemented components are based on some real-time design patterns such as *Observer* and *Watchdog*.

##### A. Node

A node is the unit of the graph that performs computations. Not to be confused with ROS Nodes (ROS contributors, 2016), this can contain a whole program and does not rely on inter-process communication. This is highly similar to Mediapipe's Calculators (Lugaresi et al., 2019). A Node runs on at most one thread at an execution timepoint. Nodes operating on different threads can communicate via Streams.

##### B. Packets

This is a basic data unit of a generic type that flows from one Node to another. It is implemented as a cheaply and efficiently copied data structure.

##### C. Streams

This component of the framework serves as a connection between two or more components of the framework, this either serves as a multi or mono output channel, therefore either serving one or more Packets per execution timepoint. A Stream of Packets is implemented as a single-ended thread-safe queue. Streams can be lossy or lossless. Lossy Streams can miss a specified number of successive Packets in case it is unable to keep up with the data flow rate and thus avoiding backpressure, while Lossless Streams are not allowed to miss any Packets and must be processed within a specified time frame.

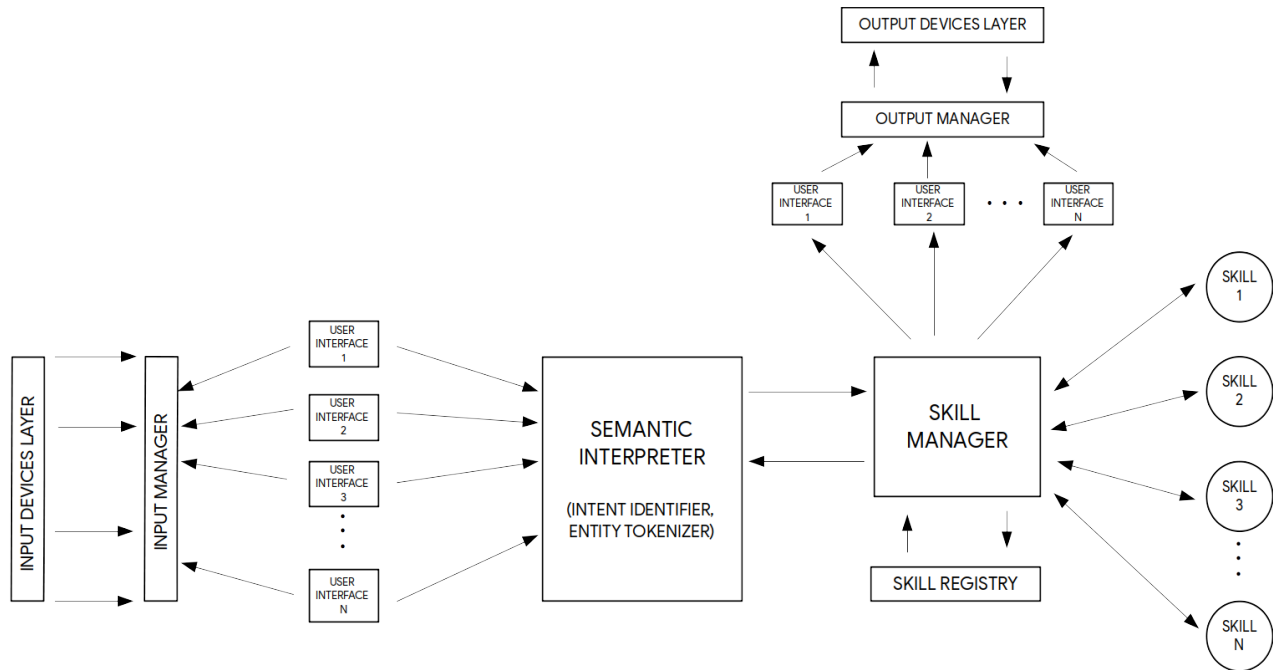


Fig. 1. IPRA Architecture

#### D. Watchdog

The Watchdog monitors the rate of data flow from one component to another along a Stream, ensures real-time constraints are met and monitors backpressure along the Stream. This is configured via a data-only structure specifying the requirements i.e latency, and throughput.

#### E. Latch

This is a mechanism that helps control the flow of Packets along the Stream, this is useful when serving multiple Streams and there is a need to control the flow of Packets to some or all of the Streams, this is controlled via a binary-state input Stream. This stops or resumes an already stopped flow of Packets along a Stream.

#### F. Aggregator

This component of the framework continuously slices accumulated sampled data across a specified interval (time-window), this is especially useful in aggregating audio or video samples for performing deep learning inference.

#### G. Attention Node

This is a special type of Node, this serves as a state Stream to the system, this is useful when it is not necessary to perform actions on every received input. An Attention Node is implemented as a Node whose inputs are generic input Streams and outputs are Streams of bits. This can be integrated with Deep Neural Networks and other algorithms.

#### H. Skill Registry

This implements the Skill Registry component of the IPRA architecture IV-F. This serves as a router from one data type to another, this includes integers, strings, pointers, function pointers, functors, etc. For example, a Skill Registry is implemented as a hash table of string keys to functor values which are executed on the same thread or another thread depending on the specified execution policy.

#### I. Skill

This implements the Skill component of the architecture IV-G. This can be as simple as a query-only Skill which only needs to respond to a user via the selected interface (e.g. speech synthesizer) upon request. Skills are of two types, namely:

- 1) *High-Level Skills*: Both developers and end-users can create this type of Skill. They only have access to the abstracted utilities and device platform's features exposed to them by the developer.
- 2) *Low-Level Skills*: This type of Skill is intended for developer utilization. They typically have access to the lower level system utilities and hardware resources.

## VI. MECHANICAL DESIGN

### A. Locomotion Unit

We utilize a trivial 4-wheel locomotion unit, this not only simplifies our initial prototype but also enables us to reduce its cost. This is controlled by the *STM32 Nucleo F302R8T6* microcontroller running at 72kHz.

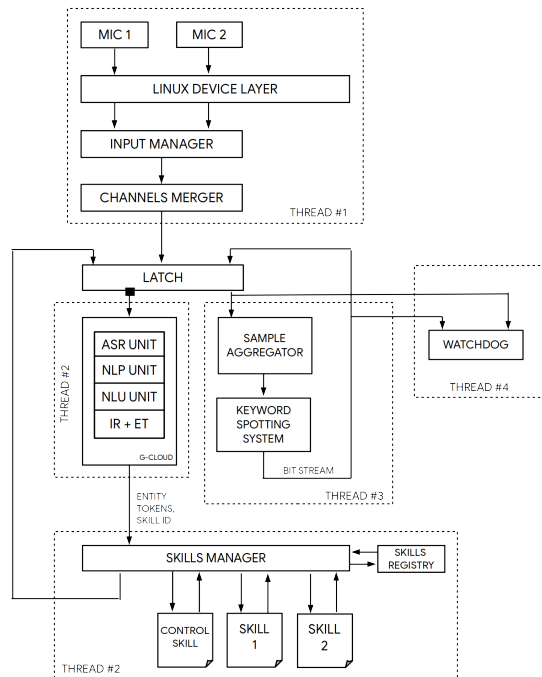


Fig. 2. Speech Only Interface Architecture using our IPRA framework. This is a streamlined version of our original deployed architecture. This highlights the key components of our framework used in implementing the architecture.

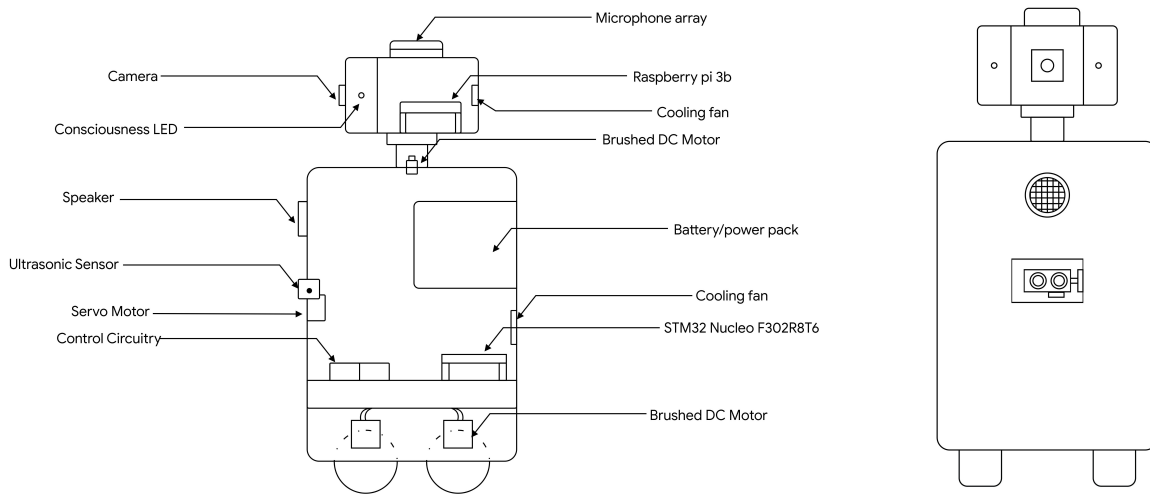


Fig. 3. L-R. Left Sectional and Front View of Marvin

The microcontroller receives a 16-bit data packet over UART at a baud rate of 115200. The received data contains 10 bits of information consisting of the direction bits and the speed bits, this is received via interrupts and executed immediately by a self preempting high priority daemon.




Fig. 4. Locomotion Unit Control Packets

$D$  : 2 bits, represents the following directions:

- Left Forward
- Left Backward
- Right Forward
- Right Backward

$V$  : 8 bits, represents speed (normalized to  $[0, 255]$ ).

The locomotion unit is stopped by setting the speed bits to zero in any direction. The microcontroller then transmits the received data to the L293D motor controller via PWM pulses which in turn varies the voltages passed to the brushed DC motors.

### B. Obstacle Detection Unit

We utilize a time-of-flight (ToF) based sensor for obstacle detection, namely the cost-effective HC-SR04 ultrasonic sensor which utilizes ultrasonic sound waves. We employ a simple beamforming technique, by rotating the sensor vertically with an MG996R brushed DC servo motor. The maximum detectable obstacle distance is restricted to a threshold of  $2.5m$  to allow for faster sampling which is a common trade-off in ToF technologies that utilize slow flight mediums. The ideal obstacle distance can simply be computed from the time-of-flight as:

$$D_{ideal} = \frac{C_{air} \times T}{2} \quad (1)$$

where  $D_{ideal}, C_{air}, T \in \mathbb{R}$ ,  $D_{ideal}$  denotes flight distance (ignoring acoustic attenuation (Wikipedia contributors, 2019a)) of the sound wave,  $C_{air}$  denotes the speed of sound in air ( $C_{air} \approx 346ms^{-1}$  at room temperature) and  $T$  denotes time taken for the sound wave to reach an obstacle and return to the sensor. Thus, by limiting  $D_{ideal}$  to a maximum of  $2.5m$ , we can ideally obtain a sampling frequency of  $69.2Hz$ .

The MG996R servo motor has a minimum operating latency of  $0.14s$  for every  $60^\circ$  at a running voltage of  $6V$  (Electronicos Caldas, n.d.), this can turn at a maximum angle of  $120^\circ$ , therefore  $\theta \in [0.0, 120.0]$ . Though it does not have a feedback control system, future versions of the prototype will address this. There are more accurate servo motors with a larger  $\theta$ , but due to cost reasons, we utilized the MG996R servo motor

Figure 5 Illustrates how we perform beamforming by varying the inclination angle of the ultrasonic sensor  $\theta$  through the servo motor. The unit is placed such that  $D_y^1$  is small enough to be safely climbed by the robot and the  $120^\circ$  coverage of the ultrasonic servo covers our area of interest (the dotted arc).

Thus, the component distances relative to the ultrasonic can be computed using the simple relation:

$$D_y = D_{ideal} \times \tan \theta \quad (2)$$

$$D_x = D_{ideal} \times \cos \theta \quad (3)$$

While it is clear that the mechanism is highly dependent on the robot's body configuration, our evaluation also revealed that the mechanism is highly subject to environmental distortions, which can be compensated for using filtering algorithms. Most of our design decisions are greatly influenced by the limited hardware at our disposal, future versions hope to address this.




Fig. 5. Obstacle Detection Technique using A Servo and an Ultrasonic Sensor

### C. Trunk

The trunk has a cuboidal structure (See Figure 3). This houses the batter/power pack, control circuitry, locomotion unit, and the STM32 Nucleo F302R8T6 microcontroller. Attached to its front is a mono-channeled USB speaker and the obstacle detection unit described in Subsection VI-B.

### D. Head

The Head houses the Raspberry Pi 3B microcomputer. Attached to its top is a microphone array used for sampling audio waveforms. To its front are attached two LEDs that illustrate the consciousness state of the robot i.e. Red for awaiting a command, Green for receiving a command, Yellow for executing a command. An RGB-888 0.5MP camera is also attached to its front for performing object detection, facial recognition, and other computer-vision-related tasks. The Head is connected mechanically to the Trunk via the high-torque MG996R servo which tilts the Head and is controlled by the STM32 Nucleo F302R8T6 microcontroller.



### E. Cooling

We perform forced air convection cooling by using heat sinks and high-velocity brushless 12V and 5V DC fans as the temperature of the system tend to rise to 80°C during prolonged usage.

## VII. KEYWORD SPOTTING

This part of the system runs continuously on-device and detects for a Hotword. The major implication arising from this is that the Keyword Spotting (KWS) system needs to have minimal latency, enough to meet real-time requirements whilst maintaining decent false-accept and recall rates (He et al., 2018; Sun et al., 2017).

The raspberry pi 3 typically runs at 1.2GHz (Quad-core) with a Broadcom BCM2837 64bit CPU having 1GB of RAM (Raspberry PI Foundation, n.d.), this limits the intensity of computation and memory size of the model that can be run on this device, and thus the architectures need to be carefully evaluated and benchmarked.

Though Traditional systems employ Hidden Markov Models in KWS tasks to model both the keyword and background noise, they were superseded by Recurrent Neural Networks and Convolutional Neural Networks (He et al., 2018; Wikipedia contributors, 2019e). We utilize the *cn-trad-fpool-3* architecture which is one of the popular small-footprint keyword spotting architectures that have been increasingly gaining traction due to the limited number of multiplies and parameters they possess (He et al., 2017; Sainath and Parada, 2015; Sun et al., 2017).

### A. Dataset

For our experiment, the word "Marvin" is chosen as the Hotword.

We utilize the *Environmental Noise Classification (ESC)* dataset (Piczak, 2015) to augment the training data and make the model rely on more robust features and be less susceptible to background noises. We also used the *Speech Commands Dataset* (Warden, 2018). Next, we manually crowd-sourced 569 audio samples from 125 volunteers within our university. These datasets primarily consisted of single-channel waveforms at 48kHz sampled under quiet conditions. To reduce the computational load, we resampled the audio samples from 48kHz to 16kHz, at this sampling frequency, the intelligibility of human speech is retained and the discretized-time sampled signal also retains spatio-temporal information useful enough for inference (Wikipedia contributors, 2019g). Next, the resulting corpus is converted into a signed 16-bit PCM encoded format, clipped to 1s, then shuffled and split into training, test and evaluation sets in the ratio of 10:1:2 respectively.

### B. Architecture

The maximum duration  $D_{max}$ , allowed for an utterance of the keyword is 1000ms. Our implementation of *cn-trad-fpool-3* has an approximate size of 8.5 MB with an average running latency of 11.8 ms and a 336μs standard deviation per unit-batch inference, which is sufficient for our use case.

type	m	r	n	p	q	Par.
conv	24	10	64	1	3	15.4K
conv	12	5	64	1	1	164.8K
lin	-	-	32	-	-	65.5K
dnn	-	-	128	-	-	4.1K
softmax	-	-	4	-	-	0.5K
Total	-	-	-	-	-	250.3K

where:

$m$  = kernel size along the temporal dimension.

$r$  = kernel size along the frequency dimension.

$n$  = number of kernels or hidden units.

$p$  = downsampling stride size along the temporal dimension.

$q$  = downsampling stride size along the spatial dimension.

This model architecture only pools in frequency and not along the temporal dimension.

### C. Augmentation

To make the model more useful in real-world scenarios and more robust to noises, distortions, and compensate for informalities not reflecting in the original datasets, we applied a few augmentation techniques on the dataset, namely:

- Pitch Augmentation
- Speed Augmentation
- Pitch-Speed Augmentation
- White-Noise Augmentation
- Background-Noise Augmentation
- Value Augmentation
- HPSS Augmentation
- Random Shift Augmentation

The noise augmentations were applied at SNRs randomly sampled between  $[-5dB, 10dB]$

### D. Preprocessing

To eliminate temporal redundancy in the sampled waveforms, we apply a 40-dimensional log-Mel energy filter bank feature extraction, this also makes training easier as many repetitive temporal features are removed in the process.

### E. Training

A multi-category classification approach is employed to allow the model rely on more robust features of the input distribution, rather than in a binary classification scheme where the model is more liable to rely on non-robust features of the input distribution. The vanilla categorical cross-entropy loss function is used in this effect.

The learning rate and batch sizes are fine-tuned on the test, training, and evaluation sets. Figure 7 exhibits the training and test inferences' accuracy and loss graphs.




Fig. 6. Sample Keyword Log-Mel Energy Filter Bank Feature

#### F. Deployment

To further reduce the computation load, we quantized the model from 32-bit floating-point precision to 8-bit signed-integer precision using the quantization regime specified in Subsection 4, this comes with minimal accuracy tradeoffs. We continuously shift the discretized-time sound sampling aggregation window by 250ms and aggregate 1s of samples, this enables us to have decent coverage of the incoming data samples as the position of the Hotword is non-deterministic. The aggregated sound samples are then passed to the interpreter to perform inference.

#### VIII. SPEECH SYNTHESIS

We utilized the Google Cloud Platform to synthesize natural-sounding high-fidelity speech via a pretrained generative autoregressive neural network, namely Wavenet (van den Oord et al., 2016; Wikipedia contributors, 2019d). Wavenet especially takes advantage of dilated and casual convolutions which makes it efficient and model spatio-temporal features of speech at a remarkable accuracy. This model generates natural-sounding audio whilst modeling linguistic features of a given text to a remarkable extent. A downside to the model is its large size and inference latency which inherently made us deploy it to the Google Cloud Platform. With sufficient hardware upgrade we can overcome this limitation.

#### IX. INTENT RECOGNITION

This part of the system recognizes *Skill Intents* and *Tokens* from a given speech input. This is performed using the Google Cloud platform's Dialogflow Speech-Intent engine (Wikipedia contributors, 2019c). This is primarily because models yielding decent performance on intent recognition and/or natural language understanding tasks are often computationally intensive and have large memory requirements as the probabilistic distribution of the modelled data is highly multi-dimensional. We utilized pre-existing DialogFlow *Skill* definitions and implemented custom ones such as Phone Call Initiation, Object Finding, and People Finding.

#### X. FACIAL RECOGNITION

This part of the system is used for facial recognition, emotion detection, and intruder detection. As such the system also needs to execute at a decent speed and accuracy, to this end, we utilize the populous Facenet neural network architecture (Schroff, Kalenichenko, and Philbin, 2015). Facenet performs at a very high accuracy on real faces (99.65% on the LFW dataset (Huang, Ramesh, Berg, and Learned-Miller, 2007), 95.12% on YouTube Faces Database (Lior Wolf and Maoz, 2011)).

Facenet's approach yields state-of-the-art performance on face-recognition tasks which is due to the great representational efficiency of the architecture whilst using only 128-bytes embeddings per face (Schroff et al., 2015), with this approach we were able to eliminate the need to retrain our models every time a new face is added into the system as this would be very computationally intensive and impractical for real-world use cases. However, the saved 128-bytes embeddings are not used to train any additional classification model instead the face recognition module employs the use of the  $L_2$  metric to calculate the distances between two face embeddings and identifies the subject. The faces are recognized if the  $L_2$  distance satisfies a specified unary predicate. The relation can be expressed as:

$$S = \sum_{n=1}^{128} (\hat{E}_i - E_i)^2$$

$$IdentifierPredicate(S) = \begin{cases} 1, & \text{if } S \geq threshold \\ 0, & \text{otherwise} \end{cases}$$

where  $\hat{E}_i, E_i \in \mathbb{R}^{128}, S \in \mathbb{R}^1$  and

$\hat{E}_i$  = Face Embedding

$E_i$  = Reference Face Embedding

$S$  = Squared Euclidean Distance Between  $\hat{E}_i$  and  $E_i$

We made use of the serially layered-pipeline approach and at the first layer of the pipeline the image taken from the camera is passed through a Haar Feature-Based Cascade Classifier which then gives a bounding box to each of the faces detected in the image, the closest face to the camera (or the face with the biggest bounding box) is then sent down through the pipeline for further processing. At the second layer, the output of the first layer is resized to  $\mathbb{R}^{1 \times 160 \times 160 \times 3}$ , which is the tensor input shape of the Facenet model. At the third layer, the output of the second layer is normalized to a mean of 128 and standard deviation of 128 and then converted from 32-bit floating precision to its unsigned 8-bit equivalent which is the expected input type of the quantized Tflite model, the output of this layer is then passed to the Tflite Facenet model which then computes the feature encoding ( $\in \mathbb{R}^{128}$ ). We employed the use of quantization-aware training in order to minimize accuracy degradation when the model is converted to a quantized






Fig. 7. L-R Training Loss and Accuracy Graphs for Our **cnn-trad-fpool3** Implementation

format. To improve the model's speed and also reduce its size, we quantized it using the symmetric 8-bit signed integer quantization with the zero-point set to 0. The 8-bit signed quantization scheme approximates floating-point values using the relation (Google Developers, n.d.):

$$real\_value = (int8\_value - zero\_point) \times scale \quad (4)$$

#### XI. OBJECT DETECTION

This part of the system is intended for Detecting Objects and People. This is capable of detecting common items such as Kettles, Cups, Chairs, Beds, Tables, Laptops, Books, etc. It is especially useful for the visually impaired. We utilize a pretrained Mobilenet model (Howard et al., 2017).

The Mobilenet architecture is a streamlined and small-footprint architecture that performs remarkably well in mobile and embedded vision applications. The model is lightweight and highly efficient due to its extensive utilization of depthwise-separable convolutions which are a form of factorized vanilla convolutions (Howard et al., 2017) whilst still modeling the data distribution remarkably though with a minimal accuracy tradeoff. This thus greatly reduces the number of multiplies and the model's size relative to a vanilla convolutional neural network. The model was quantized using the quantization regime described in 4.

#### XII. CONCLUSION AND FUTURE WORK

In the literature, we proposed a novel architecture called IPRA. We also illustrated how we can implement a variety of user interfaces using the multi-modal user interface and hardware agnostic architecture. We investigated some of the important design decisions leading to a cost-effective and interactive modular robot. As a next step to help the adoption, exploration, and advancement of this Architecture, we plan on releasing the framework and robotic assistant system implementation. In future high-fidelity prototypes, we plan on implementing and integrating more useful features, integrating high-quality sensors and actuators, and upgrading our hardware to enable it to run the whole

software fastly and locally without dependence on cloud services. Further, we plan to extend our experiments to different environmental conditions in the wild and integrate filtering and sensor fusion algorithms.

#### XIII. ACKNOWLEDGEMENTS

We would like to thank Prof. R. G. Jimoh for his feedbacks on each draft of the paper. We would also like to thank Yusuf Oladejo from the University of Ilorin for his valuable inputs and feedback on this project.

#### REFERENCES

- Electronicos Caldas. (n.d.). *Mg996r high torque metal gear dual ball bearing servo*. Retrieved from <https://www.electronicoscaldas.com/datasheet/MG996R.Tower-Pro.pdf>
- Google Developers. (n.d.). Tensorflow lite 8-bit quantization specification. Retrieved from [https://www.tensorflow.org/lite/performance/quantization\\_spec](https://www.tensorflow.org/lite/performance/quantization_spec)
- He, Y., Prabhavalkar, R., Rao, K., Li, W., Bakhtin, A., & McGraw, I. (2017). Streaming small-footprint keyword spotting using sequence-to-sequence models. *CoRR, abs/1710.09617*. arXiv: 1710.09617. Retrieved from <http://arxiv.org/abs/1710.09617>
- He, Y., Sainath, T. N., Prabhavalkar, R., McGraw, I., Alvarez, R., Zhao, D., ... Gruenstein, A. (2018). Streaming end-to-end speech recognition for mobile devices. *CoRR, abs/1811.06621*. arXiv: 1811.06621. Retrieved from <http://arxiv.org/abs/1811.06621>
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... Adam, H. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Huang, G. B., Ramesh, M., Berg, T., & Learned-Miller, E. (2007). *Labeled faces in the wild: A database for studying face recognition in unconstrained environments* (tech. rep. No. 07-49). University of Massachusetts, Amherst.



- Janssen, H. C., Samson, M. M., & Verhaar, H. J. (2002). Vitamin D deficiency, muscle function, and falls in elderly people. *The American Journal of Clinical Nutrition*, 75(4), 611–615. doi:10.1093/ajcn/75.4.611. eprint: <http://oup.prod.sis.lan/ajcn/article-pdf/75/4/611/23657721/611.pdf>
- Lior Wolf, T. H., & Maoz, I. (2011). *Face recognition in unconstrained videos with matched background similarity*.
- Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., ... Grundmann, M. (2019). Mediapipe: A framework for building perception pipelines. *CoRR, abs/1906.08172*. arXiv: 1906.08172. Retrieved from <http://arxiv.org/abs/1906.08172>
- Piczak, K. J. (2015). Version V2. doi:10.7910/DVN/YDEPUT
- Rapolienė, G. (2015). Old age stigmatization. *European Quarterly of Political Attitudes and Mentalities*, 4, 63–81.
- Raspberry PI Foundation. (n.d.). Raspberry pi 3 model b. Retrieved from <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>
- ROS contributors. (2016). Nodes. Retrieved from <http://wiki.ros.org/Nodes>
- Sainath, T., & Parada, C. (2015). Convolutional neural networks for small-footprint keyword spotting. In *Interspeech*.
- Sarikaya, R. (2017). The technology behind personal digital assistants: An overview of the system architecture and key components. *IEEE Signal Processing Magazine*, 34, 67–81. doi:10.1109/MSP.2016.2617341
- Schroff, F., Kalenichenko, D., & Philbin, J. (2015). Facenet: A unified embedding for face recognition and clustering. *CoRR, abs/1503.03832*. arXiv: 1503.03832. Retrieved from <http://arxiv.org/abs/1503.03832>
- STMicroelectronics. (n.d.). Stm32f302r8 - mainstream mixed signals mcus arm cortex-m4 core with dsp and fpu, 64 kbytes flash, 72 mhz cpu, 12-bit adc 5 msp, comparator, op-amp. Retrieved from <https://www.st.com/en/microcontrollers-microprocessors/stm32f302r8.html>
- Sun, M., Raju, A., Tucker, G., Panchapagesan, S., Fu, G., Mandal, A., ... Vitaladevuni, S. (2017). Max-pooling loss training of long short-term memory networks for small-footprint keyword spotting. *CoRR, abs/1705.02411*. arXiv: 1705.02411. Retrieved from <http://arxiv.org/abs/1705.02411>
- United Nations. (2017). World population ageing. Retrieved from [https://www.un.org/en/development/desa/population/publications/pdf/ageing/WPA2017\\_Highlights.pdf](https://www.un.org/en/development/desa/population/publications/pdf/ageing/WPA2017_Highlights.pdf)
- Valenã, T. D. C., Santos, W. d. S., Lima, P. V., Santana, E. d. S., & Reis, L. A. (2017). Deficiãfãsica na velhice: um estudo estrutural das representaã§ãsociais. *Escola Anna Nery*, 21. Retrieved from [http://www.scielo.br/scielo.php?script=sci\\_arttext&pid=S1414-81452017000100208&nrm=iso](http://www.scielo.br/scielo.php?script=sci_arttext&pid=S1414-81452017000100208&nrm=iso)
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., ... Kavukcuoglu, K. (2016). Wavenet: A generative model for raw audio. *CoRR, abs/1609.03499*. arXiv: 1609.03499. Retrieved from <http://arxiv.org/abs/1609.03499>
- Warden, P. (2018). Speech commands: A dataset for limited-vocabulary speech recognition. *CoRR, abs/1804.03209*. arXiv: 1804.03209. Retrieved from <http://arxiv.org/abs/1804.03209>
- Wikipedia contributors. (2019a). Acoustic attenuation. [Online; accessed 15-November-2019]. Retrieved from [https://en.wikipedia.org/wiki/Acoustic\\_attenuation](https://en.wikipedia.org/wiki/Acoustic_attenuation)
- Wikipedia contributors. (2019b). Bone-anchored hearing aid — Wikipedia, the free encyclopedia. [Online; accessed 15-October-2019]. Retrieved from [https://en.wikipedia.org/w/index.php?title=Bone-anchored\\_hearing\\_aid&oldid=915698226](https://en.wikipedia.org/w/index.php?title=Bone-anchored_hearing_aid&oldid=915698226)
- Wikipedia contributors. (2019c). Dialogflow — Wikipedia, the free encyclopedia. [Online; accessed 15-November-2019]. Retrieved from <https://en.wikipedia.org/w/index.php?title=Dialogflow&oldid=902286097>
- Wikipedia contributors. (2019d). Google cloud platform — Wikipedia, the free encyclopedia. [Online; accessed 15-November-2019]. Retrieved from [https://en.wikipedia.org/w/index.php?title=Google\\_Cloud\\_Platform&oldid=925790056](https://en.wikipedia.org/w/index.php?title=Google_Cloud_Platform&oldid=925790056)
- Wikipedia contributors. (2019e). Hidden markov model — Wikipedia, the free encyclopedia. [Online; accessed 15-October-2019]. Retrieved from [https://en.wikipedia.org/w/index.php?title=Hidden\\_Markov\\_model&oldid=912778235](https://en.wikipedia.org/w/index.php?title=Hidden_Markov_model&oldid=912778235)
- Wikipedia contributors. (2019f). Mobility scooter — Wikipedia, the free encyclopedia. [Online; accessed 15-October-2019]. Retrieved from [https://en.wikipedia.org/w/index.php?title=Mobility\\_scooter&oldid=911850748](https://en.wikipedia.org/w/index.php?title=Mobility_scooter&oldid=911850748)
- Wikipedia contributors. (2019g). Nyquist–shannon sampling theorem — Wikipedia, the free encyclopedia. [Online; accessed 8-October-2019]. Retrieved from [https://en.wikipedia.org/w/index.php?title=Nyquist%E2%80%9393Shannon\\_sampling\\_theorem&oldid=919845373](https://en.wikipedia.org/w/index.php?title=Nyquist%E2%80%9393Shannon_sampling_theorem&oldid=919845373)
- World Health Organization and others. (2011). *World report on disability 2011*. World Health Organization.



