

ECE 445
SENIOR DESIGN LABORATORY
PROJECT PROPOSAL

Visual Chatting and Real-time Acting Robot

HAOZHE CHI
(haozhe4@illinois.edu)
MINGHUA YANG
(minghua3@illinois.edu)
JIATONG LI
(jl180@illinois.edu)
ZONGHAI JING
(zonghai2@illinois.edu)

March 10, 2024

1 Introduction

1.1 Problem Statement

Blind individuals often face significant difficulties when navigating unfamiliar environments, such as finding water dispensers in large public spaces. Additionally, there is a risk of injury from interacting with devices that dispense hot water. The emergence of large language models (LLMs) and large visual language models (LVLMs) offers a promising avenue for developing innovative solutions to these challenges.

1.2 Solution Overview

We propose to create an AI-enhanced robotic service system designed to assist blind people by guiding them to water dispensers, providing real-time vocal instructions for navigation and safety, and autonomously refilling water bottles. This system will integrate a camera mounted on the user's head for visual input and use speech-to-text AI technology to interpret verbal commands. The core of our solution is the BLIP-2 visual language AI model, which will process both visual and textual inputs to generate actionable guidance. Additionally, a text-to-speech AI will transform text outputs into auditory instructions, thereby facilitating the user's interaction with their environment. Our system comprises the following key components:

- **Real-Time Visual and Verbal Input Processing:** A combination of a head-mounted camera and speech-to-text AI captures and analyzes the user's surroundings and voice commands.
- **Dynamic Guidance and Interaction:** The BLIP-2 model will provide navigation assistance, warn of potential dangers, and instruct on interacting with a water dispenser.
- **Autonomous Assistance:** A Universal Robot Arm UR3e, controlled by the Robot Operating System and instructed by the Vision Language AI model, will autonomously refill the user's water bottle.
- **User Communication:** Audio feedback and instructions will be delivered through a Bluetooth headset, ensuring clear and effective communication.

Operational Process When a blind individual approaches a water dispenser, the system triggers a specific sequence of actions:

1. The Vision Language AI model guides the user to sit and place their water bottle in a designated location.
2. Subsequently, a robot arm, following instructions from the Vision Language AI model, securely grasps the bottle, fills it with water from the dispenser, and then returns the filled bottle to the user.

Challenges and Innovation Integrating LVLMs within a robotic framework presents unique technical challenges, especially in accurately translating AI-generated instructions into precise robotic movements. This project not only aims to address a tangible issue but also seeks to advance the field of AI and robotics by exploring new applications of visual language models in assistive technology.

1.3 Visual Aid

The visual illustration of our AI-enhanced robotic service system is shown in Figure 1.

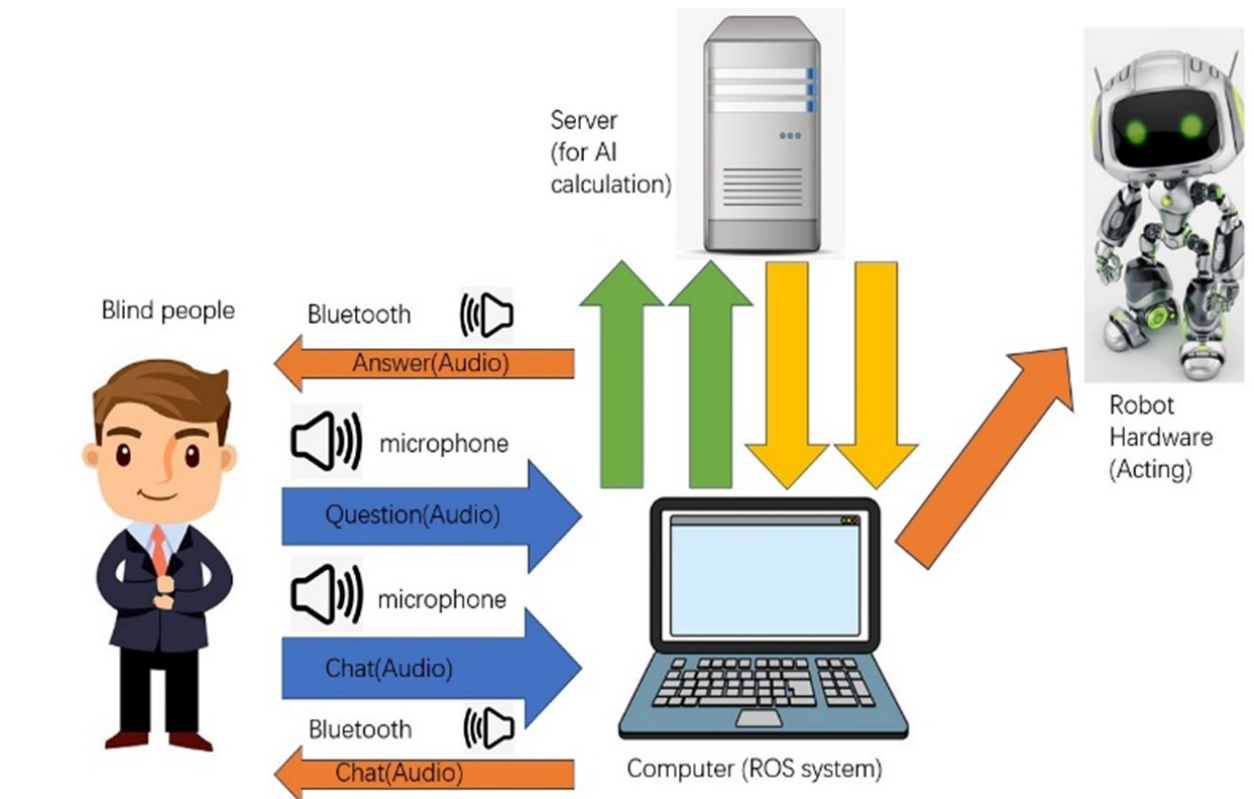


Figure 1: Visual Illustration of the AI-enhanced Robotic Service System

1.4 High-level Requirements List

1. **Server Specifications:** A high-performance server equipped with a GPU boasting a minimum of 24 GB of memory is essential for efficiently running the Visual Language AI model.
2. **Robotic Arm:** The project demands a robotic arm with a minimum of six joints for versatile movement and object manipulation capabilities. An example of such hardware is the Universal Robot UR3e.

3. **Camera System:** We require two cameras capable of capturing real-time images. The primary camera will provide visual feedback from the perspective of blind users, while the secondary camera will be used to detect the presence and position of water bottles. Both cameras must be compatible with the Robot Operating System (ROS), such as the CK camera model.

2 Design

2.1 Block Diagram

The overall block diagram of our AI-enhanced robotic service system is shown in Figure 2.

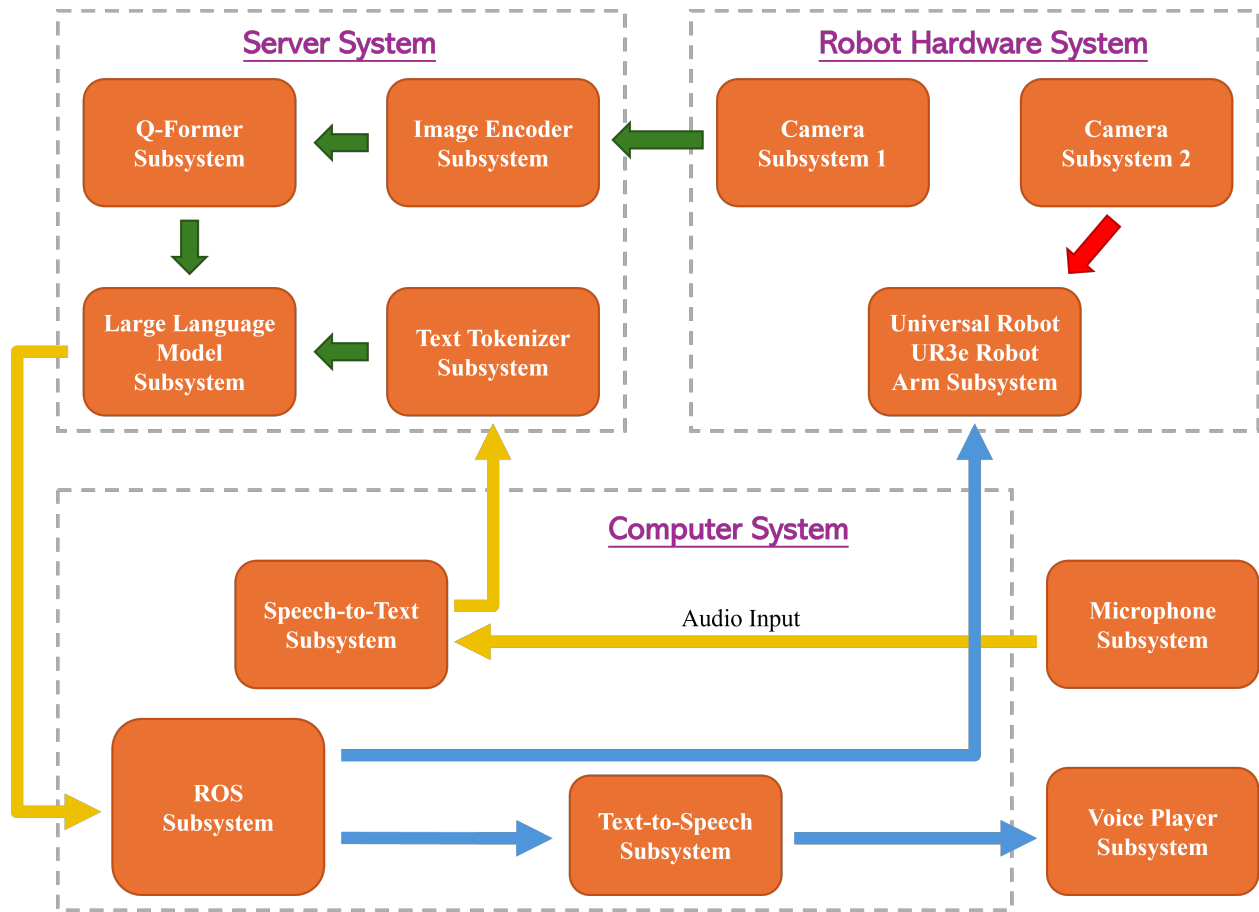


Figure 2: Block diagram of the AI-enhanced robotic service system (Green arrow: visual flow; Yellow arrow: text flow; Red arrow: attach itself to; Blue arrow: instruction flow), in which Camera Subsystem 1 is for blind people and Camera Subsystem 2 is for robot arm.

2.2 Subsystem Overview

Q-Former Subsystem The Q-Former subsystem, based on transformer architecture, transforms visual features into more abstract representations by incorporating attention mechanisms. It receives visual features from the Image Encoder subsystem and relays enhanced features to the Large Language Model subsystem for further processing.

Image Encoder Subsystem Utilizing a Vision Transformer (ViT) structure, this subsystem encodes images into visual features for analysis. It captures images from the Camera subsystem and forwards the processed visual features to the Q-Former subsystem.

Text Tokenizer Subsystem This subsystem tokenizes and converts input text into embeddings, facilitating textual analysis. Although it is mentioned that it receives images from the Camera subsystem, it should correctly receive text inputs, likely from the Speech-to-Text subsystem, and send the resulting embeddings to the Large Language Model subsystem.

Large Language Model Subsystem Processing both visual and textual embeddings, the Large Language Model subsystem generates text outputs that encompass instructions for robotic actions and responses to queries. These outputs are directed to the ROS subsystem for action and response articulation.

Camera Subsystem Operating continuously, the Camera subsystem captures images every second, providing real-time visual data to the Image Encoder subsystem for processing.

Universal Robot UR3e Robot Arm Subsystem Receiving instructions from the ROS subsystem, the UR3e robot arm interacts with its surroundings, leveraging data from the Camera and Laser subsystems to move and manipulate objects as directed.

Speech-to-Text Subsystem This subsystem converts spoken language into textual format, supplying the Text Tokenizer subsystem with text for embedding and further processing.

ROS Subsystem As the control hub, the ROS subsystem manages hardware operations based on instructions and responses generated by the Large Language Model subsystem. It orchestrates the execution of robotic actions and relays responses to the Text-to-Speech subsystem.

Text-to-Speech Subsystem Converting textual responses into audible speech, this subsystem ensures the robot can verbally communicate with users, passing audio outputs to the Voice Player subsystem for playback.

Voice Player Subsystem This subsystem is responsible for audibly playing back responses generated by the Text-to-Speech subsystem, enabling interactive communication between the robot and its human users.

Microphone Subsystem Captures audio from users, sending it to the Speech-to-Text subsystem.

2.3 Subsystem Requirements

Q-Former Subsystem The Q-Former subsystem acts as an integrative bridge, merging the capabilities of a static image encoder with a static Large Language Model (LLM). It is designed to extract and refine output features from the image encoder, irrespective of the input image's resolution. This subsystem comprises two transformer submodules utilizing shared self-attention layers, facilitating efficient processing of visual features and textual data. The image transformer submodule is dedicated to processing visual information received from the image encoder, while the text transformer submodule functions dually as an encoder and a decoder for textual information. This dual functionality enables the Q-Former to seamlessly integrate visual and textual inputs for comprehensive processing.

Image Encoder Subsystem The Image Encoder subsystem employs a transformer-based architecture, specifically a ViT-L/14 model, to encode images into a compact, feature-rich representation. It utilizes 32 distinct queries, each with a dimension of 768, matching the Q-Former's hidden dimension. The output, denoted as Z , adopts a 32×768 dimension, significantly reducing the representation size compared to the original image features extracted by the ViT-L/14 (257×1024), thus enhancing processing efficiency.

Text Tokenizer Subsystem The Text Tokenizer subsystem is responsible for converting raw text input into structured embeddings, using a tokenizer consistent with the BERT model. This approach ensures compatibility with widely utilized Large Language Models, such as Llama, facilitating seamless integration within the broader architecture and promoting effective textual analysis.

Large Language Model Subsystem This subsystem leverages an open-source Large Language Model, such as Llama, serving as a downstream decoder to process the combined input features from both visual and textual sources. The choice of Llama aligns with the architectural principles of BLIP-2 [1] based systems, enabling sophisticated text output generation that includes both instructions for robotic actions and responses to user inquiries.

Camera Subsystem The camera subsystem incorporates a CK camera, selected for its compatibility with the ROS ecosystem. It is tasked with capturing real-time images, ensuring the robot can react to its environment with updated visual information.

Universal Robot UR3e Robot Arm Subsystem Utilizing the Universal Robot UR3e, this subsystem is central to the robot's physical interactions with its environment. The UR3e is chosen for its flexibility and precision, featuring six joints that facilitate a wide range of movements and tasks.

Speech-to-Text Subsystem This subsystem employs an open-source model, such as those available from Google, to convert human speech into text. The processed textual data is then relayed to the Text Tokenizer subsystem, ensuring that voice commands are accurately interpreted and acted upon.

ROS Subsystem The Robot Operating System (ROS) subsystem serves as the robot's meta-operating system, offering essential services including hardware abstraction, device control, and inter-process communication. ROS's comprehensive ecosystem supports the development and execution of robotic applications across diverse hardware setups.

Text-to-Speech Subsystem Employing *pyttsx3*¹, an open-source Python library for offline text-to-speech conversion, this subsystem transforms text outputs from the ROS subsystem into audible speech. This allows the robot to communicate responses to user queries verbally, enhancing the interactive experience.

Voice Player Subsystem The Voice player subsystem is responsible for the audible output of the robot's responses. It receives audio files from the Text-to-speech subsystem and plays them out loud, enabling the robot to communicate effectively with its human users. Needs to incorporate a Bluetooth headset for clear audio playback.

Microphone Subsystem Should use a Bluetooth microphone for flexible and reliable audio capture from users.

2.4 Tolerance Analysis

A key design consideration is the latency involved in data transfer between the robot and the server, which is critical for real-time interaction and control. Through simulations, we have identified that data transfer delays can be minimized to approximately 10 milliseconds. Additionally, employing Python libraries such as flash-attention can significantly enhance the processing speed of our AI models. Experimental results suggest that the overall processing time, though variable, remains within a few seconds based on the complexity of input data, which is within acceptable limits for real-time operation.

¹<https://pyttsx3.readthedocs.io/en/latest/>

3 Ethics and Safety

As we work on our robot project, we're keeping a close eye on doing things in the right way according to the IEEE Code of Ethics [2].

3.1 Ethics

Privacy (ACM 1.7: Respect the privacy of others) We're making sure that any personal information the robot picks up is kept safe and private. Only people who really need to see this info can access it, and we're careful about how we store and handle it.

Fairness (IEEE - Avoiding Real or Perceived Conflicts of Interest) Sometimes, AI can be biased, making unfair decisions. We're working hard to train our robot with a wide variety of data so it treats everyone fairly.

Being Open (ACM 1.2: Avoid harm) We want everyone to understand how our robot makes decisions. We're keeping records and will be open about how the robot works and what it's been taught.

3.2 Safety

Avoiding Accidents (IEEE - Priority to Public Welfare) Our robot is built to be safe around people and objects, with emergency stops and sensors to keep it from bumping into things.

Staying Secure (ACM 3.7: Recognize the need to protect personal data) Since our robot is smart and connected, we're putting in strong security to protect it from hackers.

Dealing with Mistakes (ACM 2.5 & IEEE - Acknowledge and Correct Mistakes) If something goes wrong, the robot is designed to handle it safely and let us know so we can fix the problem.

References

- [1] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [2] Institute of Electrical and Electronics Engineers. (2016) IEEE Code of Ethics. Accessed on: 2024-03-07. [Online]. Available: <https://www.ieee.org/about/corporate/governance/p7-8.html>