

An AI-Based System Utilizing IoT-Enabled Ambient Sensors and LLMs for Complex Activity Tracking

Yuan Sun ^a, Jorge Ortiz ^b

WINLAB, Rutgers University Piscataway, NJ, USA

^a ys820@soe.rutgers.edu, ^b jorge.ortiz@rutgers.edu

Abstract: Complex activity recognition plays an important role in elderly care assistance. However, the reasoning ability of edge devices is constrained by the classic machine learning model capacity. In this paper, we present a non-invasive ambient sensing system that can detect multiple activities and apply large language models (LLMs) to reason the activity sequences. This method effectively combines edge devices and LLMs to help elderly people in their daily activities, such as reminding them to take pills or handling emergencies like falls. The LLM-based edge device can also serve as an interface to interact with elderly people, especially with memory issue, assisting them in their daily lives. By deploying such a system, we believe that the smart sensing system can improve the quality of life for older people and provide more efficient protection.

Keywords: Ambient Sensors; Complex Activity Recognition; Elderly Care; Edge Devices.

1. Introduction

Non-intrusive sensors are crucial for modern sensing applications, particularly in fields requiring continuous and unobtrusive monitoring, such as elderly care [11]. These sensors, which do not rely on cameras, offer significant advantages, including enhanced comfort and privacy for users. By seamlessly integrating into the environment without requiring direct interaction or visible placement, they minimize aesthetic and social intrusiveness[7]. Additionally, non-intrusive sensors reduce deployment and maintenance efforts, leveraging existing infrastructure and eliminating the need for frequent battery replacements[17]. This ease of installation and low maintenance makes them practical for long-term use. Furthermore, they provide flexibility and scalability in sensing applications, capturing a broad range of environmental data indirectly, thus enabling comprehensive monitoring across diverse contexts.

In this work, we build a non-intrusive smart sensing system that relies on the reasoning ability of large language models (LLMs) to assist in elderly care. The system can detect complex activities composed of more than two atomic activities. An atomic activity refers to unit-level activities that can be captured by sensors within a short time window and cannot be broken down further. Besides detecting normal atomic activities, we also use LLMs for high-level explanations and reasoning. To adapt to the memory constraints of our edge devices, we employ a local inference model on the IoT devices before sending the data to the LLM. This approach reduces the transmission burden via the wireless network, ensuring efficient and real-time processing.

2. LLM and Complex Reasoning

Our system first collects non-intrusive sensor data from ambient sensors. These sensors detect atomic-level activities. Due to the limited memory capacity of the sensors, a small model runs locally to perform initial inferences. Deploying a small model locally allows for real-time processing, reducing latency and ensuring immediate response to critical activities. Additionally, local processing minimizes the amount of data

transmitted to the cloud, enhancing privacy and reducing bandwidth usage. Once an atomic activity is detected, it is sent to a cloud server where an LLM is deployed. The LLM applies its reasoning and interaction capabilities with the user. For instance, if the sequence of activities indicates eating and drinking without sanitizing, the system can remind the user to wash their hands. Another example is if the user drinks water but forgets to take their pills; the LLM will interact with the user to remind them to take their medication. Additionally, if a user is detected as getting dressed but skipping a critical item like a coat on a cold day, the LLM can prompt them to wear it, ensuring their well-being. Furthermore, in the event of a medical emergency, the system can recognize distress signals or calls for help, and promptly alert emergency services, providing vital assistance when needed.




Fig 1. The system employs both local inferences using ambient sensors and reasoning via a cloud-based LLM. The sensors detect atomic activities, and the cloud server receives these activity sequences as context to further detect higher-level meanings, make decisions, and interact with the user.

3. Sensor Board Setup

Our goal is to construct a device that enhances the generalizability of sensor data explanation while safeguarding user privacy. The device (figure 3), installed in the environment, utilizes non-invasive sensors to detect human activities. These sensors include PIR (motion), IMU (accelerometers), audio, RGB, pressure, humidity, magnetometer, gas, and temperature sensors. They are widely available and well-suited for generalizing smart space sensing capabilities.

Smart space sensors produce two types of outputs: binary and continuous. Binary sensors, like PIR motion sensors, provide "on" and "off" states. Although useful for indicating environmental status, they may be inadequate for analyzing complex activities [8, 2]. Our device uses PIR sensors for motion detection, transmitting "ON" and "OFF" signals. In contrast, accelerometers offer continuous outputs, tracking movement along the X, Y, and Z axes. We utilize them to detect environmental vibrations, such as different vibration frequencies for pouring water versus running [1, 16, 14].

The RGB sensor detects colors within a specific range, identifying the presence of colors based on the activated channels. Pressure sensors measure barometric pressure, providing insights into weather patterns and indoor air quality, especially in HVAC operations. Magnetometers detect magnetic fields, but readings can be influenced by nearby electronic devices and human activity due to RF electromagnetic fields [12, 5]. Gas sensors detect substances like ethanol, alcohol, and carbon monoxide, useful for monitoring air quality changes due to activities such as heavy running.




Fig 2. The non-intrusive sensor board we design for our system

The Raspberry Pi Model B+ (Figure 3) serves as the central processing unit, interfacing with a custom sensor board via GPIO pins. The sensor board, equipped with various ambient sensors, collects data such as temperature, humidity, and motion. The Raspberry Pi, powered by a Broadcom BCM2837B0 quad-core ARM Cortex-A53 64-bit SoC at 1.4GHz with 1GB LPDDR2 SDRAM, runs a Python script to read sensor values at regular intervals, process the data, and store it for further analysis. This setup, featuring dual-band 802.11ac wireless LAN, Bluetooth 4.2, Gigabit Ethernet, and multiple USB ports, enables real-time monitoring and data collection essential for elderly care assistance and activity tracking. Although our current sensor setup is insufficient to run a model for data embedding locally, we plan to implement local atomic activity detection in the next stage of this work.

4. Initial Experiment

In our initial experiment, we will use the ambient sensor to collect atomic activities and test the LLM model on the cloud. The target is to determine if it possesses the reasoning capabilities that can assist in the daily lives of elderly individuals.



Fig 3. Raspberry Pi Model B+ used in our ambient sensor setup, facilitating seamless integration for elderly care assistance and activity tracking

4.1. Data Collection

The atomic level activities we collect include 20 activities: eat, paperdis, write, chop, hand wash, pour water, clean floor, knock, run, curtain, light switch, type, door pass, wipe desk, chat, basketball, saw, shave, wash dish, and brushing teeth. The sampling frequency is 90Hz for each sensor channel. There are discussions about whether the high-fidelity audio sensor channel [10] is the main channel for detecting patterns. Audio data can also be intrusive [3, 15]. To avoid privacy issues associated with the audio channel, its frequency is decreased from 16kHz to 90kHz. According to Shannon's sampling theorem, information contained in frequencies greater than half the sampling rate cannot be recovered [13, 9, 4]. Figure 4 visualizes the sensor data collected during eating activities. Figure 5 visualizes the data collected during chopping activities.

4.2. Initial Results

We developed a data encoder to process the sensor data before sending it to the LLM. Our model primarily uses a channel-wise MLP to extract features from each channel, followed by a sensor fusion module also based on MLP. Additionally, we incorporated a fast Fourier convolution module [6] to improve the accuracy rate.

Table 1 presents the initial results of F1 scores, precision, and recall for detecting various activities using the model.

Most activities, such as 'eat', 'write', 'chop', 'hand wash', 'clean floor', 'knock', 'run', 'curtain', 'light switch', 'type', 'door pass', 'wipe desk', 'basketball', 'saw', 'shave', 'wash dish', and 'teeth', achieve perfect scores of 1.00 across all metrics. However, 'paperdis', 'pour water', and 'chat' show lower performance with F1 scores of 0.43, 0.48, and 0.32, respectively. These variations indicate that while the model performs exceptionally well for most activities, there is room for improvement in detecting certain activities.

We also demonstrate the robustness of our model when the noise level 7 and frequency6 of the data vary.

4.3. Applying LLMs to Sequence Detection and Complex Activity Recognition

After detecting atomic activities, we send the sequence of detected activities to the LLM. For example, to remind elderly people to take medication before eating, we send a sequence of activities such as "brushing teeth → hand wash → pour water → eat". The corrected sequence from the LLM is "brushing teeth → hand wash → take medication → pour water → eat", recognizing the complex activity as "forgetting medication".

If the original order detected is "eat → basketball → brush teeth", the LLM will suggest "brush teeth → wash hands → eat → basketball" to ensure proper hygiene before eating, recognizing the complex activity as "unhygienic behavior". We use GPT-4 to implement sequence verification and reminders. For instance, if the input sequence is "door pass → using paper dispenser", the LLM suggests "door pass → turn

the switch → paper dispenser" to remind the user to turn on the light, identifying the complex activity as "preventing slipping". The order rules can be configured in the prompt to meet detailed requirements. Our current results demonstrate an implementable framework that can help elderly people, especially those with memory issues, to live a high-quality life.




Fig 4. Eating activity collected by the ambient sensor

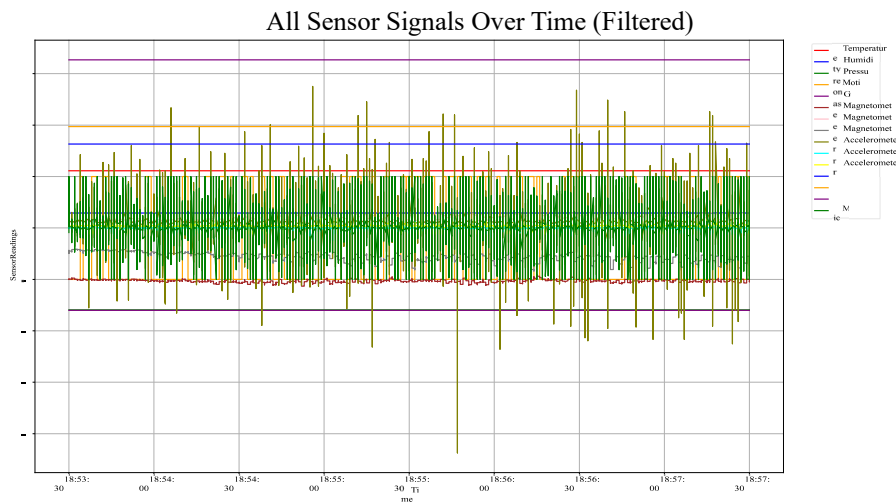


Fig 5. Data collected from chopping activities




Fig 6. Model robustness in detecting data at different frequencies




Fig 7. Model robustness with varying levels of noise in the sensor data

5. Conclusion and Future Work

In this work, we designed a framework leveraging IoT and LLMs to assist elderly people and enhance their quality of life. We employed a board with non-intrusive sensors to avoid discomfort for the elderly while monitoring their behavior to provide assistance. We implemented a low-frequency sensor sampling strategy, particularly for the microphone channel, using an unrecoverable sampling rate that cannot reproduce the original sound information. A smart IoT box was built to collect initial data for experimentation. Our initial results demonstrate that our model can successfully identify daily activities from the sensor data. We further utilized the GPT-4 interface to test the reasoning and complex activity detection capabilities, showcasing a promising framework that can assist elderly individuals.

Table 1. Initial results of F1 scores for each activity

| Activity Name | F1 | Precision | Recall |
|---------------|------|-----------|--------|
| eat | 1.00 | 1.00 | 1.00 |
| paperdis | 0.43 | 0.50 | 0.38 |
| write | 1.00 | 1.00 | 1.00 |
| chop | 1.00 | 1.00 | 1.00 |
| hand wash | 1.00 | 1.00 | 1.00 |
| pour water | 0.48 | 0.50 | 0.47 |
| clean floor | 1.00 | 1.00 | 1.00 |
| knock | 1.00 | 1.00 | 1.00 |
| run | 1.00 | 1.00 | 1.00 |
| curtain | 1.00 | 1.00 | 1.00 |
| light Switch | 1.00 | 1.00 | 1.00 |
| type | 1.00 | 1.00 | 1.00 |
| door pass | 1.00 | 1.00 | 1.00 |
| wipe desk | 1.00 | 1.00 | 1.00 |
| chat | 0.32 | 0.33 | 0.31 |
| basketball | 1.00 | 1.00 | 1.00 |
| saw | 1.00 | 1.00 | 1.00 |
| shave | 1.00 | 1.00 | 1.00 |
| wash dish | 1.00 | 1.00 | 1.00 |
| teeth | 1.00 | 1.00 | 1.00 |

However, the current smart sensor setup uses an older version of the Raspberry Pi. We plan to use a higher version of the Raspberry Pi to enable real-time inference on the edge device. Additionally, we will conduct user studies with more participants to interact with our system, helping us evaluate the design and implementation of the system.

References

- [1] Hamad Ahmed and Muhammad Tahir. 2017. Improving the accuracy of human body orientation estimation with wearable imu sensors. *IEEE Transactions on instrumentation and measurement*, 66, 3, 535–542.
- [2] Luca Arrotta, Gabriele Civitarese, and Claudio Bettini. 2022. Dexar: deep explainable sensor-based activity recognition in smart-home environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6, 1, 1–30.
- [3] Amay J Bandodkar and Joseph Wang. 2014. Non-invasive wearable electrochemical sensors: a review. *Trends in biotechnology*, 32, 7, 363–371.
- [4] Oresti Banos, Juan-Manuel Galvez, Miguel Damas, Hector Pomares, and Ignacio Rojas. 2014. Window size impact in human activity recognition. *Sensors*, 14, 4, 6474–6499.
- [5] Paolo Bernardi, Marta Cavagnaro, Stefano Pisa, and Emanuele Piuze. 2000. Specific absorption rate and temperature increases in the head of a cellular-phone user. *IEEE transactions on microwave theory and techniques*, 48, 7, 1118–1126.
- [6] Lu Chi, Borui Jiang, and Yadong Mu. 2020. Fast fourier convolution. *Advances in Neural Information Processing Systems*, 33, 4479–4488.
- [7] Veena Chiduralla, Xiaodong Wang, Xinrong Li, and Jesse Hamner. 2022. Iot based sensor system design for real-time non-intrusive occupancy monitoring. In *2022 IEEE International Conference on Internet of Things and Intelligence Systems (IoT&IS)*. IEEE, 252–257.
- [8] Munkhjargal Gochoo, Tan-Hsu Tan, Shih-Chia Huang, Shing-Hong Liu, and Fady S Alnajjar. 2017. Dcnn-based elderly activity recognition using binary sensors. In *2017 international conference on electrical and computing technologies and applications (ICECTA)*. IEEE, 1–5.
- [9] Aftab Khan, Nils Hammerla, Sebastian Mellor, and Thomas Plötz. 2016. Optimising sampling rates for accelerometer-based human activity recognition. *Pattern Recognition Letters*, 73, 33–40.
- [10] Gierad Laput, Yang Zhang, and Chris Harrison. 2017. Synthetic sensors: towards general-purpose sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 3986–3999.
- [11] Athanasios Lentzas and Dimitris Vrakas. 2020. Non-intrusive human activity recognition and abnormal behavior detection on elderly people: a review. *Artificial Intelligence Review*, 53, 3, 1975–2021.

- [12] M Martinez-Burdalo, A Martin, M Anguiano, and R Villar. 2004. Comparison of fdt-d-calculated specific absorption rate in adults and children when using a mobile phone at 900 and 1800 mhz. *Physics in Medicine & Biology*, 49, 2, 345.
- [13] Anzah H Niazi, Delaram Yazdaneh, Jennifer L Gay, Frederick W Maier, Lakshmi Ramaswamy, Khaled Rasheed, and Matthew P Buman. 2017. Statistical analysis of window sizes and sampling rates in human activity recognition. In *HEALTHINF*, 319–325.
- [14] Samir A Rawashdeh, Derek A Rafeldt, and Timothy L Uhl. 2016. Wearable imu for shoulder injury prevention in overhead sports. *Sensors*, 16, 11, 1847.
- [15] Ioan Susnea, Luminita Dumitriu, Mihai Talmaciu, Emilia Pecheanu, and Dan Munteanu. 2019. Unobtrusive monitoring the daily activity routine of elderly people living alone, with low-cost binary sensors. *Sensors*, 19, 10, 2264.
- [16] Xuzhong Yan, Heng Li, Angus R Li, and Hong Zhang. 2017. Wearable imu-based real-time motion warning system for construction workers' musculoskeletal disorders prevention. *Automation in construction*, 74, 2–11.
- [17] Xin Yu and Panos D Prevedouros. 2013. Performance and challenges in utilizing non-intrusive sensors for traffic data collection.