International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Comparative Analysis of Edge AI Hardware Platforms for Wearable Assistive Technologies

Chaitanya Mule[1], Dweep Vartak[2], Devang Wangde[3], Prof. Vidya Sagvekar[4]

[1,2,3,4]Department of Artificial Intelligence and Data Science
K. J. Somaiya Institute of Technology, Mumbai, India
[1]chaitanya.mule@somaiya.edu, [2]dweep.vartak@somaiya.edu, [3]devang.w@somaiya.edu, [4]vsagvekar@somaiya.edu

Abstract

This review paper presents a comparative analysis of AI hardware platforms for edge devices, focusing on their suitability for wearable assistive technology applications. As edge computing in AI continues to expand, numerous hardware solutions have emerged, each optimized for specific perfor- mance metrics such as inference speed, energy consumption, and processing capability. By reviewing six recent studies on edge AI hardware, this paper assesses the adaptability of these platforms for energy-constrained environments—crucial for real- time applications such as AI-based wearable devices for visually impaired users. The paper offers key insights into different hardware selections, their performance efficiency and future advancements required to boost the capability of wearable AI solutions.

Index Terms: Edge AI, Wearable Device, Assistive Technology, Hardware Accelerators, Real-Time AI, Energy Efficiency

1. Introduction

With Edge AI technology, the method of operation of assistive devices based, especially on wearable application for visually-impaired individuals is being gradually changed. Tra- ditional assistive technologies like canes or audio guides lasted for limited functions and mostly required continuous manual engagement. However, modern wearable devices benefit from Edge AI and can perform tasks such as real-time object detec- tion, text recognition, and even NLP without prompting from the user, allowing for a much more interactive and seamless experience. These devices identify obstacles, recognize text or face, and respond to voice commands, increasing autonomy and quality of life for visually impaired individuals.

Edge AI, such as the cloud or others, provides low latency processing while conforming to wearable form factors. These devices are relied on by visually impaired users to relay critical information in real time since even the smallest of latencies can inhibit the effectiveness of object detection or navigation assistance. Where real-time processing is neces- sary, edge computing would be useful since this preserves dependency on cloud systems. Although the cloud could offer such advanced capabilities, cloud computing would often be burdened with latency issues and faultiness from relying on networked infrastructures, thus making them very unreliable options for real-time wearable applications.

Wearable devices are to work in spaces with severe con- traints in volume and energy consumption, hence the choice of AI hardware is extremely important. The hardware should maintain a well-determined balance between computation capabilities and consumption of electricity in order for the device to provide high-performance AI-processing without depleting the battery or generating excessive heat. The balance therefore becomes even more paramount in wearables since compact design and prolonged life of the battery are the two keystones in user acceptance and practical use. Many growing varieties of hardware-from general-purpose processors to specialized AI hardware solutions-offer developers multiple alternatives but the delicate task of identifying an optimal solution with respect to performance-versus-portability is indeed critical for the creation of successful wearable assistive devices.

This paper is reviewing advances in edge AI hardware with respect to their usefulness in developing wearable assistive technologies for the visually impaired. This review discusses critical aspects of the process of selecting such hardware- pervasive devices. Given those are considerations in terms of inference speed, energy efficiency, and compatibility with AI frameworks, these aspects are key for anyone working at the intersection of edge AI precisely and developing wearable assistive technologies to better enable equalization and inter- activity in the lives of their visually impaired users.

2. Literature Review

The need for low-power implementations of high- performance modules capable of working in energy-constrained environments propelled the development of edge AI accelerators. Various aspects of these accelerators have been explored by different researchers focusing mainly on their applications in wearable assistive devices.

Pomsar et al. [1] provide a comprehensive overview of edge AI accelerators, underscoring the critical importance of energy efficiency when deploying wearable devices. They give a classification of various accelerators based on various dimen- sions such as power consumption and processing capability while stressing the trade-off between performance and energy efficiency. It is for this reason that power efficiency became the foundation for how hardware selections affect real-time applications, most generally in cases that require the capability to run for a sustained period.

Building on this, Arnautovic and Teskered ´ ziˇ c [2] delve deeper into the performance metrics of embedded systems, specifically focusing on inference speed and energy con- sumption. Their work complements Pomsar et al. by offering a detailed analysis of how AI hardware can be optimized for low-power scenarios while maintaining the computational speed necessary for real-time processing. Both studies under- line the challenge of balancing computational demands with power limitations in wearable devices, making their insights particularly relevant for assistive technologies that rely on continuous operation.

Srija et al. [3] give additional clarification to such ideas by emphasizing the significance of application-specific hardware selection. Instead of resenting cost or technical specifications, it supports a decision-making process driven by benchmarking performance and definite needs of every application. This perspective is consistent with other previous studies but adds an element of practicality; thus urging developers to prioritize real-world usability when choosing their edge AI hardware.

Firmansyah and Paul [4] further heretofore introduced a benchmark framework for evaluation of the performance of deep neural networks across different hardware accelerators. This comparative analysis provides useful information about the strengths and weaknesses of different devices and solidifies the point that the hardware needs to be selected, which needs right balancing of performance and energy efficiency needs. This benchmark approach further enhances the effort laid by Arnautovic and Pomsar by proposing a suitable tool that measures the fitness of hardware to certain use cases.

Khan and Paul [5] shifted the aim toward mobile hardware stating that even smartphones can conduct complicated AI tasks if optimized correctly. The conclusion of their work relates much to wearable assistive devices in that compact low-power hardware will give way to certain applications. Their study underscores the prospects of mobile devices as AI platforms, and thus also adds to the repertoire of hardware options available to developers in this arena.

Lastly, Pomsar and Brecko [6] undertake a detailed bench- mark analysis of the performance of YOLO on edge intelli- gence devices. This research establishes the very possibility of deploying advanced algorithms like YOLO in real-time applications on wearable devices. This work complements the benchmark efforts of Firmansyah and Paul and points out the importance of optimizing any AI model toward a hard- ware platform specifically where real-time object detection is needed in energy-constrained environments.

In the aggregate, this provides a comprehensive understand- ing of the present landscape, Edge AI hardware for wearable assistive devices. Together they point out further the need for concurrence between performance, energy efficiency, and us- ability, accelerating future efforts in research and development in this vibrant field.

3. METHODOLOGY

To evaluate the feasibility of selected AI hardware platforms for wearable assistive devices, we systematically reviewed literature. Information provided in earlier studies pointed to an importance of balancing energy efficiency with performance and usability in edge AI applications. This informed our production of a comprehensive evaluation framework that focuses on the specifics that could determine a successful wearable application for the blind and visually impaired. The methodology involved the following steps:

- **Selection Criteria:** We have taken interest in the studies published in the last three years, directly regarding edge hardware designs of AI in wearable devices. Studies that emphasize performance parameters such as speed of inference, energy consumption, cost, and compatibility with AI models were taken into account with a view to ensure all-naturalness for a real-time, energy-constrained environment normally seen in wearable assistive tech- nologies.
- **Data Extraction:** From each selected study, we extracted key metrics, which included:
  - **Inference Speed:** Measured in frames per second (FPS), inference speed is critical for real-time per- formance in tasks such as object detection and text recognition. Higher FPS values indicate a device’s ability to process more frames per second, which translates to smoother real-time processing for wear- able applications.
  - **Energy Consumption:** Measured in watts (W), en- ergy consumption reflects the importance of power efficiency, particularly for wearable devices where battery life is a major constraint. Devices that offer high performance with lower power consumption are more suitable for continuous operation in assistive technology applications.
  - **Cost:** The overall cost of the hardware platform was considered, as affordability is a key factor for the widespread adoption of assistive technology, espe- cially in resource-limited settings.
  - **Compatibility with AI Models:** We evaluated the ability of each hardware platform to support pop- ular AI frameworks (e.g., TensorFlow, PyTorch, ONNX) and models like convolutional neural net- works (CNNs), which are essential for a wide variety of AI tasks, including object recognition, natural language processing, and speech synthesis.
- **Comparative Analysis:** A detailed comparative analysis was conducted by categorizing the hardware platforms based on specifications such as processing power, mem- ory, storage, and AI framework compatibility. Perfor- mance benchmarks from each study were reviewed to create a comprehensive comparison of the devices. These benchmarks were systematically presented in tables to al- low for a clear side-by-side evaluation of each platform’s capabilities.
- **Usability Evaluation:** To assess the practical applica- bility of each hardware platform for wearable assistive devices, we evaluated usability in real-world scenarios, particularly for visually impaired users. This evaluation included:

– **Real-Time Performance:** They assessed the per- formance of how efficiently each platform worked under real-time conditions focusing on design delay-latency-and its ability to carry out data processing rapidly enough for tasks involving detection of ob- jects and natural language processing.

– **Energy Efficiency:** The study examined the compro- mising of performance against power consumption, thereby ensuring the hardware could adequately op- erate under energy-constrained environments-a key metric for battery-powered wearable devices.

The usage of the evaluation framework allowed us to ensure that our comparative analyses were governed by the most relevant factors pertaining to assistive technologies. Insight was thus made available for which AI hardware platforms have optimal overall performance-energy efficiency-cost combina- tion for wearable applications.

I. COMPARATIVE ANALYSIS OF HARDWARE CAPABILITIES

Upon selection of these edge AI devices for the purposes of a comparative analysis, the prime focus has been put on four major options: Raspberry Pi 5 + Hailo-8L AI Kit, NVIDIA Jetson Nano Dev Kit, Google Coral Dev Board, and Intel NCS

2. These devices are chosen keeping into consideration vari- ous popular, performance-oriented attributes focused on the conveniences available for the developers building wearable AI applications. All four devices fits well under the $300 price tag, reasonable for projects on a budget while provid- ing ample features, numerous processing powers, and energy consciousness. Various facets concerning their specifications, performance metrics, and their overall efficiency for edge AI applications in wearable devices are evaluated among these devices.

A. Hardware Specifications Comparison

TABLE I: Comparative Hardware Specifications

Device	**CPU**	**RAM**	**Storage**	**Supported Frame- works**	**GPU or Acceler- ator**
Raspberry Pi 5 + Hailo-8L AI Acceler- ator	Broadcom BCM2712, Quad-core Cortex-A76	8GB LPDDR4	microSD, USB	TensorFlow Lite, PyTorch, ONNX	Hailo-8L AI Acceler- ator
Nvidia Jetson Nano Dev Board	Quad-core ARM Cortex-A57	4GB LPDDR4	microSD, USB	TensorFlow, PyTorch, TensorRT	NVIDIA Maxwell Archi- tecture GPU
Google Coral Dev Board	NXP i.MX 8M, Quad-core Cortex-A53	1GB or 4GB LPDDR4	eMMC, mi- croSD	TensorFlow Lite, Edge TPU	Edge TPU AI Acceler- ator
Intel NCS 2 AI Ac- celerator (USB stick)	Requires a host device CPU (e.g Rasp- berry Pi)	1GB LPDDR3	N/A (USB stick)	TensorFlow, Caffe	Intel Mo- vidius Myriad X VPU

In this, the Hardware Specifications Comparison highlights the key differences in number of cores, CPU power, memory, storage, and compatibility across various frameworks of these edge AI devices. The Raspberry Pi 5 + Hailo-8L AI Kit takes the lead with a quad-core Cortex-A76 processor with 8 GB of RAM and compatibility with more than one framework; hence it is a high-performance choice for real-time tasks. At the same time, the NVIDIA Jetson Nano and Google Coral Dev Board offer solid alternatives, with comparatively less memory but enough processing power for some vision-based tasks. The Intel NCS 2 requires a host device such as Raspberry Pi 5, limiting its use in the self-contained wearable applications. Definitions of Key Terms:

- **AI Accelerator:** An AI accelerator is the hardware specifically built for the effective execution of tasks in artificial intelligence such as machine learning and deep learning. These chips are meant to provide parallel processing, granting faster training and inference times than common CPUs do.
- **GPU (Graphics Processing Unit):** A GPU is primarily designed for rendering graphics but can also perform AI computations. While it handles parallel tasks well, its ar- chitecture is not exclusively optimized for AI workloads. AI accelerators typically outperform GPUs in speed and efficiency for AI-specific operations.

B. Performance Benchmarks

TABLE II: Performance Metrics Comparison

Metric	**Raspberry Pi 5 + Hailo-8L AI Kit**	**NVIDIA Jetson Nano Dev Kit**	**Google Coral Dev Board**	**Intel NCS 2 with a Host device (Pi 5)**
Inference Speed (FPS)	30-60 FPS	30 FPS	15-30 FPS	15 FPS
Computational Power (TOPS)	13 TOPS	0.472 TOPS	4 TOPS	4 TOPS
Latency (ms)	30-50 ms	30-40 ms	100-150 ms	70-100 ms
Energy Efficiency (FPS/W)	8 FPS/W	4 FPS/W	5 FPS/W	2 FPS/W
Cost	~ $150	~ $258	~ $150	~ $160

Note: All values for latency and power consumption may vary based on use case, model, and testing conditions; thus, they should be considered approximate rather than definitive.

The Performance Metrics Comparison table outlines key benchmarks across four edge AI devices, focusing on in- ference speed, computational power (TOPS), latency, energy efficiency, and cost.

The Raspberry Pi 5 + Hailo-8L AI Kit demonstrates im- pressive capabilities, achieving an inference speed of 30-
60 FPS and offering a robust computational power of 13 TOPS. With a latency between 30-50 ms, it is well-suited for real-time applications and exhibits high energy efficiency at approximately 8 FPS/W, making it a leading choice for performance-intensive tasks, priced at around $150.

The NVIDIA Jetson Nano provides a steady performance with an inference speed of 30 FPS and a computational power of 0.472 TOPS. Although it operates with a similar latency between 30-40 ms, its energy efficiency is lower at approximately 4 FPS/W. Its cost is around $258, reflecting its advanced capabilities for various model flexibility needs.

The Google Coral Dev Board delivers an inference speed of 15-30 FPS with 4 TOPS of computational power, though it has a higher latency ranging from 100 to 150 ms. It offers a balanced energy efficiency of approximately 5 FPS/W and is priced at about $150, making it suitable for low-power vision- based tasks.

In comparison, the Intel NCS 2, when paired with a host device like the Raspberry Pi 5, provides an inference speed of 15 FPS and 4 TOPS of computational power, with a latency between 70-100 ms. However, its energy efficiency is lower, at approximately 2 FPS/W, and it is priced around $160, positioning it as a less favorable option in terms of performance per watt compared to the other devices.

Note: Prices are approximate and may vary based on the retailer and availability.

C. Energy Efficiency Visualization

The following table represents the power consumption (in Watts) of each device during inference:

TABLE III: Computational Power and Efficiency Comparison

Device	**Computational Power (TOPS)**	**Power Consumption (W)**	**TOPS/Watt**
Hailo-8L AI Accelerator with Raspberry Pi 5	13 TOPS	5-10 W	1.3 - 2.6
Nvidia Jetson Nano Dev Kit	0.472 TOPS	5-10 W	0.047 - 0.094
Google Coral Dev Board	4 TOPS	5 W	0.8
Intel NCS 2 with a Host Device (Pi 5)	4 TOPS	5-10 W	0.8 - 0.4

Note: Power consumption values are approximate and may vary based on specific tasks or workloads used during testing.

This analysis highlights the energy requirements and effi- ciency of various AI devices. The Hailo-8L AI Accelerator paired with the Raspberry Pi 5 achieves the highest compu- tational power at 13 TOPS with variable power consumption of 5-10 W, resulting in an efficiency of 1.3 - 2.6 TOPS/Watt, making it ideal for demanding applications. The Google Coral Dev Board operates at a maximum of 5 W, delivering 4 TOPS and 0.8 TOPS/Watt, effectively handling quantized models due to its energy-efficient Edge TPU. The Intel NCS 2 with a host device like the Raspberry Pi 5, also provides 4 TOPS with variable power consumption (5-10 W), yielding an efficiency range of 0.8 - 0.4 TOPS/Watt, which varies based on the host device’s power consumption. In contrast, the Nvidia Jetson Nano Dev Kit offers the lowest performance at 0.472 TOPS and an efficiency of 0.047 - 0.094 TOPS/Watt, making it less suitable for energy-sensitive tasks.

This comparison highlights that the performance and energy consumption figures pertain to the entire edge devices (e.g., Jetson Nano, Raspberry Pi 5 with Hailo-8L, NCS2 with Raspberry Pi, etc.), encompassing both the AI accelerators and their respective host devices, rather than solely reflecting the specifications of the individual accelerators.

Balancing Latency, Power, and Model Complexity: Each device’s consumption of power and efficiency in completing a specified task reflect its suitability for specific wearable applications. For example, a device to be used by a blind user may need great speed and low latency in order to instantly provide feedback about mobility or object recognition, thus the Raspberry Pi 5 + Hailo-8L, despite the greater power draw, would fit such applications. Or if the device needs simple OCR and no real-time performance is required, then the Google Coral Dev Board or the Intel NCS 2 would be preferable because of their lower power consumption, thus extending battery life. It is also quite possible to select simpler and quantized versions of the models in order to further reduce power, thus allowing devices like the Coral Dev Board to operate acceptably in those energy-constrained environments.

D. Cost-Efficiency Analysis for Wearable Assistive Devices

TABLE IV: Key Ratios for Cost-Effectiveness and Power Efficiency

Device	Cost per TOPS (USD/TOPS)	FPS/Watt Ratio	Cost per FPS/Watt (USD/(FPS/W))
Raspberry Pi 5 + Hailo-8L	$11.54	approx. 8	$18.75
NVIDIA Jetson Nano	$604.66	approx. 4	$71.25
Google Coral Board Dev	$37.50	approx. 5	$30
Intel NCS 2	$40.00	approx. 2	$80

Table IV concludes with a general review of the factors which affect the suitability of various devices for wearable assistive applications, targeting the cost per TOPS, FPS/Watt ratio, and cost per FPS/Watt.
It is the most powerful configuration, Raspberry Pi5+Hailo8L, being ongoingly saleable in performing torturous tasks with the cost per Tera Operations Per Second less than $11.54 and a cost of $18.75 per FPS/W, thereby both suitable for real-time object detection on budget-sensitive applications.
The NVIDIA Jetson Nano, on the other hand, ends up with a substantially-export cost per TOPS of $604.66 and $71.25 as a cost per FPS/Watt, designed more for tasks that are highly demanding in terms of the machine's power under non-battery- restricted conditions.
Reasonably priced, the Google Coral Dev Board costs
$37.50 per TOPS and $30.00 per FPS/Watt; therefore it should perform well in low-power, vision-based tasks.

The Intel NCS 2 is too costly for everything other than the algorithm itself, having a cost per TOPS of $40.00, but with $80.00 of a cost per FPS/Watt makes it rather ineffective for budget-oriented low-power tasks.
In general, the Raspberry Pi5 + Hailo8L and the Google Coral Dev Board give better cost-efficiency for wearable assistive technology, while the NVIDIA Jetson Nano provides a specific approach based on computation versatility.

II. Usability in Edge AI Applications

This section majorly evaluates the suitability of various hardware devices for wearable applications, simultaneously focusing on real-time AI performance, ease of integration, and practical deployment challenges, particularly for assistive technologies.

TABLE V: Strengths and Weaknesses of Hardware Devices for Wearable Applications

Device	Strengths	Weaknesses
Hailo-8L with Raspberry Pi 5	High performance, low latency, great support, versatile framework compatibility, compact size, flexible storage options, low power consumption	Moderate power consumption may affect battery life, potential bottlenecks with storage options
Intel Neural Compute Stick2 (NCS 2)	Highly portable, optimized for specific AI tasks, cost-effective	Lower efficiency and power-to-performance ratio, Requires host device, Limited for complex tasks

Nvidia Jetson Nano Dev Kit	Excellent for complex AI tasks, strong ecosystem, supports multiple frameworks	Higher power consumption, less suitable for energy-efficient applications, lower performance rating compared to others
Google Coral Dev Board	Optimized for efficient AI tasks, low power consumption, ideal for vision applications	Limited support for NLP, moderate processing capability, less versatile for multi-modal applications

Real-World Deployment Challenges: For wearable assis- tive devices, the most demanding operational considerations include size, weight, and battery efficiency. Devices must be lightweight and compact while allowing for extended and comfortable wear for users with visual impairment. Some larger setups, such as the Raspberry Pi 5 + Hailo-8L, offer high performance, but sure attention must be paid to their integration into wearables to avoid user discomfort. Con- versely, smaller options such as the Intel NCS 2 with a small host device provide more heavy-duty wearability for discreet applications but offer much less computation power and flexibility for complex tasks.

Battery Life Considerations: In particular, battery capacity is extremely important for wearable technologies to perform continuous AI processing, such as real-time object detection or text recognition functions. The Raspberry Pi 5 + Hailo- 8L has power efficiency of roughly 8 frames per second per watt, meaning it is most suited for applications requiring long battery life and high performance. While the Google Coral Dev Board is efficient at around only 5 FPS/W, it is better suited for more moderate power tasks where a balance between energy use and processing power is important. The power- hungry NVIDIA Jetson Nano, for example, has a typical efficiency of 4 FPS/W. High energy consumption means that these devices require larger batteries or must be charged more frequently; this may be particularly inconvenient for individuals with some degree of visual impairment dependent on long-term uninterrupted work. Therefore, tuning AI models and employing power management techniques are key to achieving performance at battery life in such setups.

Enhancing Usability with Advanced Features: Putting current hardware options to effective use for vision-based tasks, development within the sphere of Natural Language Processing (NLP) offers the potential for higher usable inter- facing, voice command, contextual understanding, and real- time language interaction. These advancements will offer devices an option to be more interactive and more accessible, particularly for users with visual impairment relying on voice- based interfaces. For example:

- **Voice Commands Complementing Object Detec- tion:**NLP could support follow-up queries after text de- tection (e.g., ”What does the sign say?” or ”Translate this text”), improving the device’s functionality in tasks requiring additional context or language support.
- **Enhancing Text Recognition:** When the device scans and reads text, NLP could allow the user to ask follow- up questions for clarification or additional details. For instance, after detecting text on a sign, the user could ask, ”What does the sign say?” or ”Translate this text into another language.”

Hybrid AI Solutions: A promising approach to improve usability in wearable AI devices is to leverage hybrid AI solutions. For instance, combining multiple compact AI accel- erators with a single low-power host device can achieve both energy efficiency and high processing capability. A lightweight host, such as the Raspberry Pi Zero 2 W, paired with compact accelerators like the Google Coral USB Accelerator for vision tasks and the Intel Neural Compute Stick 2 for language processing, provides an efficient balance between power us- age and computational strength. This configuration allows wearable devices to handle both lightweight and intensive tasks, enhancing the experience for users who depend on both visual and conversational interaction in assistive devices. Additionally, developing standardized

APIs and frameworks for seamless integration across hardware platforms will be essential to streamline AI deployment and promote flexibility across device ecosystems.

Also, creating **standard APIs and frameworks** to connect different hardware platforms will play a key role to make it easier to roll out AI apps. This would help developers work together and make sure AI models can adapt to various devices, which makes the whole system more flexible and able to grow.

III. Analysis

This comparison features the importance of selecting right hardware for wearable assistive technologies, focusing on its alignment with specific application requirements.

The Hailo-8L-Enhanced Raspberry Pi 5 has an influence on inference speed reaching 30 to 60 FPS and a strong computational operation ability of 13 TOPS. This allows it to run complex tasks with ease, like convolutional neural networks (CNNs) to detect objects and recognize text in real-time, with little delay between 30-50 ms. This kind of performance is key for wearable assistive devices that need to process data and give instant feedback for tools designed to help impaired users. The Raspberry Pi 5 gives users different storage choices, including microSD and USB. While microSD cards are less costly, they can be slower than built-in storage. USB storage lets you add more space and update , but picking the right storage type is essential to access large datasets .

Considering user experience, the Pi 5 and Hailo-8l AI Kit's compact design makes it easy to integrate with different wear- able formats; either belts or handheld devices. With relatively less power usage, about 5-10 W , it ensures longer battery life which is important for Edge AI devices. Also, the Pi 5 works well with popular machine learning tools like TensorFlow Lite, PyTorch, and ONNX. This lets developers use all kinds of pre- trained AI models for object detection, language translation, or sentiment analysis.

To efficiently manage AI tasks, adequate RAM is essential, impacting the device's ability to handle large datasets and mul- tiple processes simultaneously. The Raspberry Pi 5, with its higher RAM capacity (up to 8GB), can support more complex models and larger datasets, ensuring smooth operation during real-time processing tasks.

The Intel NCS 2 is a very small, low-power device (4 TOPS) that provides AI performance of the latest generation. It is a USB stick, and hence requires a host (for example a laptop or other edge platform such as Raspberry Pi) to communicate with. This host dependency has the effect of making it infrequent for wearables to be totally self-contained, whereas its 1 W power consumption is mostly utilized for battery life extension of the paired system.

Even if the NCS 2 prides itself on its energy efficiency, its capabilities are determined by the performance of the host device's RAM. Perfection would be guaranteed by 4GB of RAM if requirements are fulfilled for AI tasks, which can be a constraint for devices that are supposed to process bigger data sets and/or more complex models. The best fit for NCS 2 is dealing with less complex scenarios such as object detection or simple AI inference, but its limitations surface when it is used on more advanced applications like natural language processing (NLP).

The Nvidia Jetson Nano Dev Kit leaves other AI products in the dust when comes down to dealing with AI applications that involve sequential data processing, such as voice recognition and NLP. Its adaptable GPU has the capacity of both doing paralleling tasks and general-purpose processing which makes it a very versatile option for several types of AI workloads. Nonetheless, although it serves as the bulwark of the AI task, its design does not allow it to be the superior at efficiency in terms of the specialized AI hardware.

Although the Jetson Nano's power consumption (5–10 W) is the only factor that limits its use as a battery-powered wearable device where low energy consumption is of major concern. Likewise, the Jetson Nano, like the Intel NCS 2 too, needs at least 4GB of RAM to run successfully, especially, for the data-intensive applications. Even though one of its strong features is the diversity of AI frameworks the KIT supports, such as TensorFlow, PyTorch, and TensorRT, it may become less preferable because of high energy consumption for long- term operation wearables.

The Google Coral Dev Board comes with a special Edge TPU, the very best part being the fact it is optimized for the conduction of quantized models making the board very efficient for real-time tasks like object detection and image classification. It is a difficult test upon which no challenger will survive if this is to be compared with its durability. Google Coral, despite its high efficiency in consuming only 5W of power for the smooth operation of portable, battery- powered devices, finds it hard to apply the software to the more demanding domains such as NLP, thus minimizing its applicability in devices that are supposed to encompass different modes of operation like voice interaction and text analysis.

Moreover, the 1GB or 4GB of RAM may not be enough for the more complex AI models or multitasking and this could limit its efficiency in the most resource-consuming applications. Despite the bestowed gravity of his argument, the Coral Dev Board's high level of energy efficiency, in addition, to a special focus on vision-based applications is a line that will carry the wearables with the ability to detect and categorizing objects with a sequence of priorities.

If we think about the future, combining vision-based tasks with natural language processing in those hybrid technologies will give a more broad solution for assistive devices. For example, the blind person could be familiar with the object through the device that comes up with real-time object detec- tion and voice interaction. In other words, the device could not only identify objects but also deal with a conversation about described objects.

The choice of the hardware for wearable assistive de- vices mostly depends on processing capacity, power which is available, budget constraints, preferred functionalities, and RAM sufficient for running properly. The Hailo-8L AI Kit- Enhanced Raspberry Pi 5 is a flexible and reliable tool for both the performance and power efficiency side as well as expandability, thus, it can be a very strong contender for future edge AI implementations.

4. FUTURE WORK

Edge AI development, as one aspect of this whole move- ment, has lots of potential for the ways in which assistance devices' capabilities would be improved and several aspects could be investigated further in that regard. This part lists the main areas for future research and development, as well as the challenges in realizing and the strategies for integration.

Enhanced Natural Language Processing (NLP) Capabili- ties:

Even though many edge AI devices are the best in the vision- based tasks, they often find it difficult to perform the complex NLP operations which are very much needed for reliable assistive applications. Here we must emphasize research on the optimization of algorithms which will facilitate low power devices to process languages in real time for applications such as speech-to-text conversion, language translation, and sentiment analysis.

Implementation Challenges: Building lightweight models such as DistilBERT will be based on the trade-off between model size and accuracy. Scientists will have to play with on-device processing to make sure that low latency and responsiveness are achieved.

- **Roadmap:** In the first stages, the process may involve benchmarking existing models both against edge hard- ware, and across processing units, after which it would be done through a number of iterations and feedback loops that would prove the performance metrics of the case in hand, this is finally leading to the deployment of language models that have been tailored to wearable applications.
- **Hybrid Model Deployment for Multi-Modal Tasks:** Current solutions often only target one part of multi- modality, namely vision, or speech recognition. The near future would probably require hybrid architectures that can carry out multi-modal tasks through the integration of vision, NLP and audio information technology into a single system.

Implementation Challenges: Creating a seamless connec- tion between different AI modalities requires a single infor- mation format and a good way to distribute the resources for processing to be sure that all the processors work an equal amount.

- **Roadmap:** Co-researching with teams of experts from different fields might be a way of trying out and verifying multi-modal solutions that then can be introduced in a stepwise manner to real-life situations, which will prove their usefulness in assistive technologies.
- **Energy Optimization Techniques:** In AI wearable tech- nology, the transition between functionality and power consumption is, on the one hand, the principal obstacle. Future studies should focus on power management, for instance, with dynamic power scaling, more optimized inference scheduling, and better cooling solutions.

Implementation Challenges: These methods that are really good must be experimentally shown to be very efficient without any drop in performance.

- **Roadmap:** First, analyze the power consumption patterns for different types of applications to spot the areas of energy optimization. After the initial step of introducing and testing the options in laboratory conditions, the possible solutions are to be deployed in actual wearable devices.
- **Development of Edge-Optimized AI Models:** Because the efficient models are in high demand, the work on edge-optimized AI models is very urgent. The research should be oriented towards the introduction of custom architectures as well as the use of the model compression techniques, for instance, knowledge distillation which allows significant reduction of the model size while maintaining the performance.

Implementation Challenges: Ensuring that compressed models maintain accuracy across diverse tasks can be chal- lenging.

- **Roadmap:** Bring the Raspberry Pi 5 community and model developers together to develop special models for the Hailo-8L AI Kit, tested and improved many times through the user feedback loop.
- **Integration with Cloud-Based Solutions:** Research into hybrid edge-cloud approaches that utilize the cloud for in- tensive computing and the local node for critical real-time processing process should help overcome the limitations of edge devices.

Implementation Challenges: Developing a robust com- munication protocol is critical to minimizing latency while ensuring security and privacy.

- **Roadmap:** Prototype a hybrid model by shifting specific tasks to the cloud and then do elaborate real-world applications testing to gauge effectiveness and choose the information to users for feedback to make improvements.
- **Improved Interoperability and Standardization:** The elevating range of edge AI hardware requires technology- based integration and standardization to simplify devel- opment and deliver compatibility.

Implementation Challenges: Establishing widely accepted standards can be slow due to the diversity of existing plat- forms.

- **Roadmap:** Collaborate with the industry stakeholders to introduce standard Application Programming Interfaces (API) and cross-compatible libraries that help develop a larger ecosystem for wearable AI applications. This will ease the integration problems between devices and models.
- **Security and Privacy Enhancements:** Since wearables collect confidential information about users, placing a high priority on privacy and security is of the utmost importance. Future research should look into the estab- lishment of secure solutions through data encryption at the electric circuit and the software levels.

Implementation Challenges: Balancing security measures with performance and user experience can be difficult.

- **Roadmap:** Utilize federated learning models that make training possible as a collaborative exercise without any data breaching. Conduct thorough tests to show that the new security features will give way to better functionality and user experience if they are a bit obstructive.

Thus, the ability of wearable devices to work with AI on the edge will demand a community of developers to mitigate the current limitations and come up with visionary solutions. Through the guidelines set by the field of research, researchers and tech developers will be in a position to boost technology awareness and usability in order to increase safety and life quality for users.

5. Conclusion

The Raspberry Pi 5 in variant Hailo-8L uniquely serves as a perfect fit for the wearable artificial intelligence applica- tion because it is high-performance, energy-saving and cost- effective. Its ability to act quickly due to a low delay and compatibility with AI frameworks used by most technology companies makes it the right choice for the development of aids for the blind. Besides, the Pi 5's flexible capacity for storage and its robust community support package the device more, placing it as a multi-functional solution for edge AI applications in wearable devices.

Second is the Intel NCS 2, which is although suitable, suffers from proper energy efficiency when it uses a host computer, thus it appears to be better performing at more discrete, less battery consuming assignments. The Google Coral Dev Board is a suitable product for simple AI devices because it is energy-efficient and task-optimized. NVIDIA is universal but in high-demand wearable applications, the device's efficiency isn't even comparable to that of the other players in the high-demand field. The Pi 5, according to tests, is the best of its kind in terms of analog strength and a micro-community, which proves that the lead was indeed an intelligent one.

Forward-looking technologies, including the combination of hybrid AI that combines vision processing and natural language processing (NLP), have a good future for improving customer relations. This way, devices could give contextual information and make more natural and intuitive interactions possible, which would be a seamless experience for people with visual impairments. In addition to this, the advent of the next-gen AI hardware solutions, including the neuromor- phic chips, possesses the power to trigger a revolution in performance and energy consumption, which in its turn, can allow for the development of more powerful and longer-lasting wearables.

The coordination between hardware manufacturers, AI model developers, and assistive technology experts is crucial for boosting further innovation in this field. By working as a team, these stakeholders can improve the limits of wearable AI, thereby making hardware and software more efficient to the increasing demands of the users. As the industry develops, one of the main focuses should be on energy optimization and hybrid AI functionality in order to increase the penetration of wearable devices, particularly in battery-operated environ- ments.

In the end, the coexistence of advanced AI hardware and user-centered design will bring about a radical change in the field of wearable assistive technologies, which will have the aim of improving accessibility and enhancing the quality of life for visually impaired people.

References

1. L. Pomsar, A. Brecko, and I. Zolotova, "Brief Overview of Edge AI Ac- celerators for Energy-Constrained Edge," in Proc. 20th Int. Symposium INFOTEH-JAHORINA, Kosˇice, Slovakia, 2021.
2. A. Arnautovic´, E. Teskeredzˇic´, "Evaluation of Artificial Neural Network Inference Speed and Energy Consumption on Embedded Systems," in Proc. 20th Int. Symposium INFOTEH-JAHORINA, 2021.
3. S. Srija, P. Kawya, T. A. Reddy, "Choosing Appropriate AI-enabled Edge Devices, Not the Costly Ones," in Proc. 27th Int. Conf. Parallel and Distributed Systems (ICPADS), 2021.
4. M. H. Firmansyah, A. Paul, "DeepEdgeBench: Benchmarking Deep Neural Networks on Edge

Devices," in Proc. Cross Strait Radio Science and Wireless Technology Conf., 2021.

5. M. A. Khan, P. Paul, "AI Benchmark: All About Deep Learning on Smartphones in 2019," IEEE Trans. Human-Mach. Syst., vol. 50, no. 6, 2020.
6. L. Pomsar, A. Brecko, "Benchmark Analysis of YOLO Performance on Edge Intelligence Devices," in Proc. Cross Strait Radio Science and Wireless Technology Conf., 2021.