BELLARMINE COLLEGE PREPARATORY
Institute for Innovation & Entrepreneurship

Bellarmine College Preparatory

FALL 2026

The Proceedings of the

Bellarmine Academic Research Symposium

STEM in Society, Volume 1

Saturday, October 3, 2026
Bellarmine College Preparatory, San Jose, California

PAPERS BY

Nine Student Research Papers

Ethan Chan
Isabella Chau
Ethan Cheung
Arjun Goli
Samyak Jain
Abhijay Konduru
Aryan Kulkarni
Alexandra Lu & Athena Li
Neethika Vijay, Justin Cho & Wyatt Eros

Foreword

Welcome to the inaugural Bellarmine Academic Research Symposium, and to our first topic, STEM in Society.

Every paper in this collection began the same way: a student noticed a problem, asked a question about it, and then did the slow work of trying to answer it. They built apparatus, wrote code, gathered data, ran experiments, and wrote up what they found, including the parts that did not go to plan.

The range of their curiosity is striking. In these pages you will find a low-cost hearing aid that reads brain signals to focus on the voice its wearer is attending to, a method for mapping floating trash on a lake from aerial images, a phone app that detects and deters phone use behind the wheel, and coatings designed to keep robots working on the Moon and Mars. Other authors model involuntary eye movement, measure what rock climbing does for memory, rebalance the load in a hiking backpack, compare how neural network architectures trade accuracy for speed, and ask when time online makes prejudice better or worse.

The word that matters most in the title STEM in Society is the last one. The authors were asked to consider not only what they discovered but who it affects, and nearly every paper closes with a section on its impact on society. Science and engineering are never done in isolation from the people they serve, and these students have taken that seriously.

My thanks to the authors, the families who supported them, and the teachers and mentors who gave their time. To our guests: read these papers closely, and give this work the attention it has earned.

Dr. Geoff Barrall
Director, Institute for Innovation and Entrepreneurship
Bellarmine College Preparatory

Symposium Committee

STUDENT SYMPOSIUM CHAIR

James Che ’27

ORGANIZING COMMITTEE

Dr. Sonya Arriola
Paul Ballesteros
Dr. Geoff Barrall

REVIEWING PANEL

David Dutton
Brad Lindemann
Casey McCullough
Sachin Salgaonkar
Brooke Slater
Wes Winter

Contents

Cellphone Implementation to Address Distracted Driving Using AI 5
Ethan Chan

Why Digital Exposure Can Backfire: A Contact-Architecture Framework for Social Media and AI 8
Isabella Chau

Automated Aerial Estimation of Floating Trash Density Using Computer Vision Techniques 11
Ethan Cheung

A Low-Cost Multi-lingual Hearing Aid Powered by Neuro-Adaptive System Integrating EEG-Driven Attention Detection with AI Beamforming to Improve Speech Clarity in Noisy Environments 14
Arjun Goli

Pupil Tracking and Spiking Neural Network Modeling of Horizontal Nystagmus 18
Samyak Jain

Multilayer Coatings for Thermal Management for Robots in Outerspace 26
Abhijay Konduru

Comparative Analysis of Modular, Sequential, and Parallel Neural Architectures on Inference Accuracy and Latency 30
Aryan Kulkarni

Effect of Rock Climbing on Neurocognitive Performance and Memory 33
Alexandra Lu and Athena Li

Data-Optimized Rebalancing Apparatus for Load-Balance Monitoring in Hiking Backpacks 36
Neethika Vijay, Justin Cho and Wyatt Eros

Cellphone Implementation to Address Distracted Driving Using AI

Ethan Chan
28echan@gmail.com

1. Introduction

1.1 Background

Automobile accidents are the global leading cause of youth death aged 5 to 29, resulting in over a million deaths every year [1]. In fact, California is the second highest total ranked place for these accidents, due to the increasing number of distracted drivers on the road [2], [3]. In my research, I explored a unique way to detect and deter phone usage while driving by leveraging AI in combination with IOS software implemented in a cellphone.

1.2 Problem

I developed a solution in which the very same device (e.g., cellphone) that causes distracted driving is utilized to detect and deter the distracted driving (Figure 1). Such novel solution requires the cellphone itself to determine (1) whether the cellphone itself is in a moving vehicle, (2) whether the cellphone is being used, and (3) whether the person using the cellphone is a driver. My research aims to evaluate whether a cellphone can be configured, via IOS app development in combination with AI training to accurately determine three metrics, i.e., speed, phone usage, and driver identification, so that they can be combined in a processing architecture implemented in a cellphone to detect phone usage by driver.



Figure 1. Cell phone causing the distracted driving can be utilized to detect and deter distracted driving.

2. Investigation

2.1 Methodology

I developed a distracted-driving detector that is run on a cellphone using IOS software and a trained AI model. The processing architecture includes three detectors (kinematic detector, usage detector, and driver detector) and an evaluator (Figure 2).

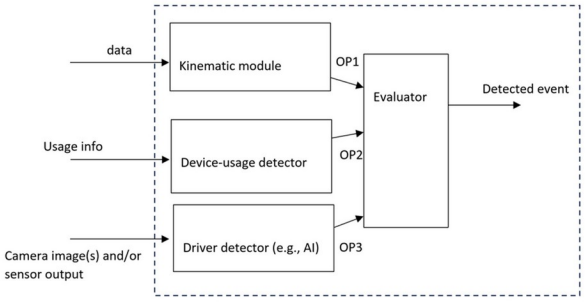


Figure 2. Processing architecture of distracted-driving detector

The AI implements the driver detector and is trained using 329 training images of different drivers and passengers. The training images captured environmental cues (e.g., car seat, steering wheel, window, visor, etc.) so that the

AI model can learn to distinguish a driver from passenger based on these cues.

During use, the kinematic detector obtains speed information from the cellphone's GPS unit. The usage detector obtains touchscreen usage from the touchscreen sensors. If the cellphone's speed is above 5 mph, and if the usage metric indicates cellphone usage, then the cellphone camera captures an image of the cellphone user. The AI model analyzes the image to determine whether the person using the moving cellphone is a driver or not. If so, the distracted-driving detector concludes that there is distracted driving (Figure 3).

```
while app Running:
    gate(1) check speed > 20mph:
        fetch speed data from Apple's location manager
        if speed > 20 = true:
            gate(1) check device usage:
                fetch user's input on screen using
                simultaneousGesture property (identifies
                touches to the screen)
                if user currently swiping on phone:
                    gate(3) use neural network to check
                    driver vs passenger:
                        if neural network returns "driver":
                            distracted event must have
                            occured -> administer
                            consequence screen:
                                pops up screen and force
                                user to wait 3 seconds
                                after 3 seconds, detect
                                touch to deescalate back
                                to gate(1)
                        else:
                            deescalate to gate(1)
                else:
                    deescalate to gate (1)
        else:
            do not escalate to next step and continue
            checking speed

# Figure 3. Pseudocode
```

I tested the accuracy of the driver detection (AI model) using 103 testing images. Since GPS unit and touch screen sensors are well known to be accurate [4], [5], the speed and usage metrics are not tested.

2.2 Results

Figure 4 shows the testing results of the AI model. The testing images include 66 driver images, 21 passenger images, and 16 uncertainty images. The AI model correctly identifies 61 driver images, 18 passenger images, and 10 uncertainty images, with 86.4% accuracy.

	DRIVER	PASSENGER	UNCERTAIN
TRUE DRIVER	61	4	1
TRUE PASSENGER	2	18	1
TRUE UNCERTAIN	1	5	10

$$Accuracy = \frac{61 + 18 + 10}{103}(100\%) = 86.4\%$$

Figure 4. Testing result showing AI model accuracy

3. Conclusion

3.1 Related Work

A previous approach employed by fleet companies relies on physically mounting a dashcam in front of the employee-driver [6]. My approach is different because it does not require a separate dashcam. Instead, my solution utilizes the cellphone (the item causing the driver's distraction) itself to detect distracted driving, which is advantageous because drivers who are distracted by cellphone usage certainly will have the cellphone.

3.2 Limitations

A limitation to my research is the size of training data utilized to train the AI. The training data set includes only 329 training images. The accuracy of the AI model is expected to increase if larger training data set (covering more different situations, such as different persons, different vehicles, different poses, etc.) is utilized.

3.3 Societal Impact

My innovation will help reduce distracted driving and death tolls in major cities. Insurance companies and government will benefit from reduced expenses due to car

accidents. Also, parents can utilize this innovation to prevent their children from distracted driving.

4. References and Additional Details

4.1 References

[1] World Health Organization, “Road traffic injuries,” World Health Organization, Jul. 20, 2026. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries.

[2] UC Berkeley Safe Transportation Research and Education Center (SafeTREC), “2025 SafeTREC Traffic Safety Facts,” Univ. California, Berkeley, 2025. [Online]. Available: https://safetrec.berkeley.edu/2025-safetrec-traffic-safety-facts.

[3] National Safety Council, “Distracted driving,” *Injury Facts*. [Online]. Available: https://injuryfacts.nsc.org/motor-vehicle/motor-vehicle-safety-issues/distracted-driving/.

[4] N. Gershenfeld, “How do touch-sensitive screens work?,” MIT School of Engineering, Jun. 7, 2011. [Online]. Available: https://engineering.mit.edu/ask-an-engineer/how-do-touch-sensitive-screens-work.

[5] V. Gikas and H. Perakis, “Rigorous performance evaluation of smartphone GNSS/IMU sensors for ITS applications,” *Sensors*, vol. 16, no. 8, p. 1240, Aug. 2016. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/27527187/

[6] Samsara, “Samsara AI Dash Cams,” YouTube. [Online]. Available: https://www.youtube.com/watch?v=PazzdlNtB-k.

4.2 Disclosures & Acknowledgements

Claude AI was used in conjunction with my prompting to create parts of an application. It was also utilized to artificially generate synthetic images for training and testing a newly developed AI model

Why Digital Exposure Can Backfire: A Contact-Architecture Framework for Social Media and AI

Isabella Chau
ichau27@priorypanther.com

Keywords / Index Terms
Intergroup contact, digital contact, algorithmic bias, artificial intelligence, prejudice reduction

1. Introduction

1.1 Background

The algorithms and artificial intelligence programs behind social media platforms influence which social groups people interact with digitally, as well as how such groups are represented. Although contact theory suggests that these encounters can mitigate prejudice, digital interaction may instead worsen stereotypes, deepen divisions, and silence marginalized voices.

1.2 Problem

In this paper, the core question addressed is: What design elements differentiate digital contacts that combat prejudice from those that perpetuate it? This research proposes a contact-architecture approach, taking psychological literature on contact theory and applying it to technological algorithms.

2. Investigation

2.1 Methodology

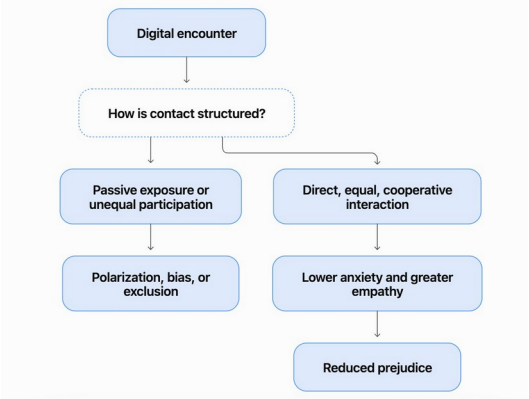
In this systematic review, seven focal sources were analyzed using narrative analysis, selected from a larger pool of forty papers [1]–[7]. The evidence set included two meta-analyses of intergroup contact, a theoretical assessment, a social media field experiment, a language-model bias audit, a meta-analysis of digital contact, and a controlled AI-facilitated contact experiment. The sources provided quantitative evidence on prejudice, its mediators and mechanisms, and technologically mediated interaction. Information extracted from every source included the design of the study, the sample size, the context of contact, the intensity of participant engagement, the psychological variables involved, and the presence of four established contact conditions: equal status, cooperative goals, institutional support, and repeated meaningful contact. Interactions mediated via technological devices were further subdivided into passive and structured direct contact. More emphasis was placed on studies involving controlled experiments than on those based on correlational analysis. Finally, patterns identified across the different sources were condensed into Table 1 by translating each factor into corresponding requirements applicable to online platforms.

2.2 Results

The connection between intergroup contact and decreased levels of prejudice has been demonstrated across multiple analyses, including 515 separate studies and 713 individual sample groups. In addition, anxiety

reduction and increased empathy and perspective-taking were shown to mediate this effect. In relation to technology use, digital evidence showed distinctions based on mere exposure versus active interaction. For instance, digital contact decreased prejudice among a sample group of 9,385 participants across 88 separate studies (effect sizes d = 0–0.25). Exposure to conflicting viewpoints led to high levels of polarization in one study of 1,652 individuals, while another study found that a language-model AI reproduced biases against minority dialects. On the other hand, a mediated AI contact experiment lowered interracial prejudice and anxiety among 128 participants compared with non-mediated control groups. Ultimately, technology's effects depended less on exposure itself than on how the encounter was structured.

Figure 1. How contact architecture shapes digital outcomes



3. Conclusion

3.1 Related Work

Past research has demonstrated both general and digital contact effects [1], [4]. This study extends that work by converting psychological conditions into platform design requirements and measurable outcomes.

3.2 Limitations

Because this was not a preregistered systematic review, source selection may be incomplete, and the proposed framework has not yet been directly tested across platforms or demographic groups.

3.3 Societal Impact

This framework could potentially assist developers, schools, and Bay Area institutions in differentiating meaningful contact from exposure that may worsen prejudice. Balanced participation and consistent moderation of platforms could benefit marginalized users, while poorly designed systems may amplify harassment, stereotypes, or exclusion. The framework should therefore be tested for effectiveness, accessibility, privacy, and unequal outcomes before implementation.

Table 1. Proposed solution: contact conditions translated into testable platform requirements.

Evidence-based condition	Technical or policy requirement	Measurable outcome
Equal status	Balance visibility and speaking power	Participation gaps, exposure diversity
Shared goals and cooperation	Build collaborative tasks, not solo debate	Reciprocity, task completion, empathy
Institutional support	Publish rules and enforce moderation consistently	Harassment, safety, appeal outcomes
Repeated meaningful contact	Support persistent groups and follow-up	Trust and attitude durability

4. References and Additional Details

4.1 References

[1] T. F. Pettigrew and L. R. Tropp, “A meta-analytic test of intergroup contact theory,” Journal of Personality and Social Psychology, vol. 90, no. 5, pp. 751–783, 2006, doi: 10.1037/0022-3514.90.5.751.
[2] T. F. Pettigrew and L. R. Tropp, “How does intergroup contact reduce prejudice? Meta-analytic tests of three mediators,” European Journal of Social Psychology, vol. 38, no. 6, pp. 922–934, 2008, doi: 10.1002/
ejsp.504.
[3] E. L. Paluck, S. A. Green, and D. P. Green, “The contact hypothesis re-evaluated,” Behavioural Public Policy, vol. 3, no. 2, pp. 129–158, 2019, doi: 10.1017/bpp.2018.25.
[4] L. P. da Costa, K. Bierwiaczonek, and M. Bianchi, “Does digital intergroup contact reduce prejudice? A meta-analysis,” Cyberpsychology, Behavior, and Social Networking, vol. 27, no. 7, pp. 440–451, 2024, doi: 10.1089/cyber.2023.0591.
[5] C. A. Bail et al., “Exposure to opposing views on social media can increase political polarization,” Proceedings of the National Academy of Sciences, vol. 115, no. 37, pp. 9216–9221, 2018, doi: 10.1073/
pnas.1804840115. [6] V. Hofmann, P. R. Kalluri, D. Jurafsky, and S. King, “AI generates covertly racist decisions about people based on their dialect,” Nature, vol. 633, pp. 147–154, 2024, doi: 10.1038/
s41586-024-07856-5. [7] S. Sahab et al., “E-contact facilitated by conversational agents reduces interethnic prejudice and anxiety in Afghanistan,” Communications Psychology, vol. 2, art. no. 22, 2024, doi: 10.1038/ s44271-024-00070-z.

Automated Aerial Estimation of Floating Trash Density Using Computer Vision Techniques

Ethan Cheung
ethan.cheung27@bcp.org

1. Introduction

1.1 Background

Waterborne trash is an environmental concern not only because of its impact on ecosystems, but also because it is often concentrated in specific areas rather than evenly distributed across the water surface. Without reliable information about where trash is most concentrated, cleanup efforts within lakes become inefficient and difficult to prioritize.

At present, estimating the density and location of waterborne trash usually relies on manual visual inspection at multiple sites. Cleanup prioritization is therefore based on limited, localized observations, and because floating trash can shift over time, hotspot locations may change, making it challenging to consistently monitor high-density areas.

1.2 Problem

This research aims to support volunteers and trash cleanup organizers by helping identify high-impact cleanup areas through automated analysis of aerial-style images. By providing a clearer and more consistent view of trash distribution over time, this approach has the potential to improve cleanup planning and support more effective monitoring of water resources.

2. Investigation

2.1 Methodology

Because of constraints relating to the flight of real drones, this project will focus on creating a mock-up system to prove feasibility. The model of the lake is constructed from a sheet of blue paper measuring 28 by 22 inches, with random assorted Legos scattered across its surface to represent floating trash. On the 4 corners and the middle of the areas between them, I printed out 8 ArUco markers to allow the camera to triangulate its position relative to the blue paper. The corner markers were 2 inches wide and 3 inches tall. The other markers were 3 inches wide and 4 inches tall. They were laid out as shown in Figure 1.

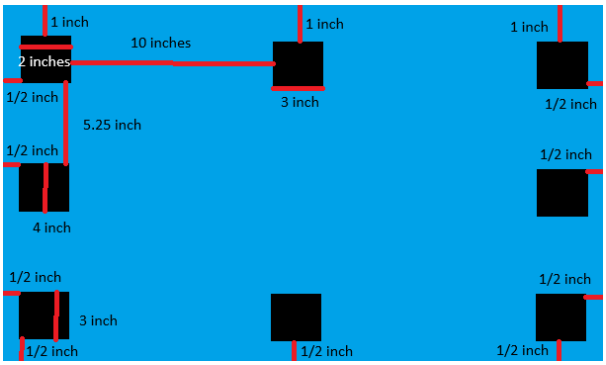


Figure 1. The layout of the ArUco markers. It is important to note that the corner markers are slightly smaller than the other ones and that all the markers are not equidistant from the edge of the paper because they were printed out on small rectangles.

If I were to use a real drone, the drone's position would be instead discovered using an onboard GPS system. To simulate a drone flying across the water to scan the surface of the lake, I held my phone roughly 3 feet above the sheet of paper and slowly moved it back and forth over the paper, thoroughly scanning it. I made sure to record the entire area of the blue paper and the process took from 30

seconds to a minute to complete. The resulting .MOV video was then sent to the Raspberry Pi for processing.

The video frames were analyzed one by one. First, the image needed to be aligned using the ArUco markers and homography transformation. Homography is essentially a series of matrix transformations used to map points in two image planes with one another. In this case, it was used to ensure the camera's video was always taken from a consistent top-down view, no matter if the camera was held at a slight angle.

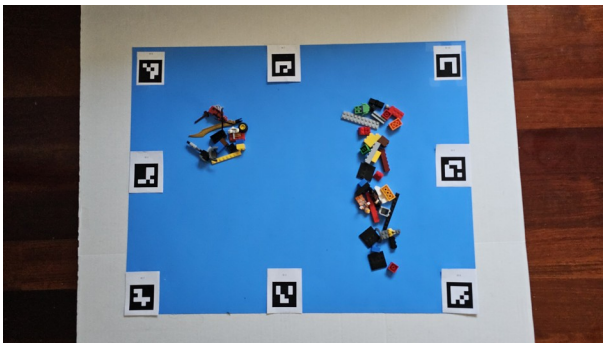


Figure 2. An example setup. The ArUco markers can clearly be seen along the edges of the blue paper and the Lego bricks representing trash sit in the middle.

It was also used to align each frame relative to each other in order to ensure the moving camera wouldn't cause a single cluster of trash to be counted twice because it appeared to have moved when the camera moved. computer vision was used to determine the color of the background by examining the most common color in the image. Any pixels not of that color were designated to be trash. Next, a grid was created over the image. Computer vision was then used to determine where the centers of the trash clusters were, and the grid cells the centers were in were marked as having trash. The more centers were in a cell, the more trash-dense that cell was designated as. After this process is repeated with every image frame, the resulting heatmap is exported as a PNG file, with an overlay option also available to see the heatmap over the trash.

2.2 Results

As a whole, my project was a success. The heatmap was generally aligned to the trash's location, with hotspots appearing over the trash's location. The reproducibility was 100%, meaning when the algorithm was fed the same video, it outputted the same result every time. I met my FPS and latency targets very easily, with the average FPS being 52.78 and the average latency, or the time taken to process one frame, being 18.97 milliseconds.

For hotspot agreement between the heatmap and the trash's actual location, I had to visually compare the two images and rate their similarity. I would look at both the heatmap and the picture of trash position, along with the heatmap overlaid atop the trash picture. I would then estimate how closely each heatmap was. The agreement ranged from 70% to 100%, with the average being 85% agreement, a respectable number.

3. Conclusion

3.1 Related Work

Methods for scanning bodies of water for trash already exist, but they have drawbacks that this project solves. Satellites such as the Copernicus Sentinel-2 can scan large swathes of water at a time, but they can take days to get into the right location to perform a scan. Drone-based systems like MARLIT can use deep learning to detect floating trash, but it is not a real-time system. Users also need to use R/RStudio, install code, and figure out how to use nadir-position images, presenting a significant barrier to use. In contrast, this project is real-time and cheap and simple to run. It only requires a camera-mounted drone and a Raspberry Pi, and can return an output in minutes.

3.2 Limitations

While conducting this project using a mock-up provided valuable data, it would have been ideal to use a real drone on a real body of water. Furthermore, I would like to use AI in the future. For this project, we found the Hailo AI model wasn't consistent, but I feel it could help reduce the problems shadows, reflections, and non-trash objects might cause.

3.3 Societal Impact

This project would most benefit smaller volunteer groups such as school clubs. Larger organizations could use more effective methods, such as the use of satellites or drone scanning that isn't real-time. However, because my project is real-time and low-cost, smaller groups who still want to clean up waterborne trash can go out with a cheap drone and raspberry pi, find where the trash is, then clean it up, hopefully all in a single day.

4. References and Additional Details

4.1 References

[1]
O. Garcia-Garin et al., "Automatic detection and quantification of floating marine macro-litter in aerial images: Introducing a novel deep learning approach connected to a web application in R," Environmental Pollution, vol. 273, p. 116490, Mar. 2021, doi: https://doi.org/10.1016/j.envpol.2021.116490.
[2]
"Copernicus Sentinel-2 helps advance marine debris detection - Sentinel Success Stories - Sentinel Online," Sentinel Success Stories, Mar. 26, 2020. https://sentinels.copernicus.eu/web/success-stories/-/copernicus-sentinel-2-helps-advance-marine-debris-detection

4.2 Disclosures & Acknowledgements

The Hailo AI module was originally used to try and detect trash clusters. Chat GPT was also used to research homography and Aruco markers.

A Low-Cost Multi-lingual Hearing Aid Powered by Neuro-Adaptive System Integrating EEG-Driven Attention Detection with AI Beamforming to Improve Speech Clarity in Noisy Environments

Arjun Goli
goli.arjun@gmail.com

Keywords / Index Terms
Neuro-adaptive systems, EEG, beamforming, hearing aids, auditory attention decoding, assistive AI

1. Introduction

1.1 Background

The human brain has the ability to selectively focus on a single speaker in noisy environments, a phenomenon known as the "cocktail party effect." However, conventional hearing aids primarily amplify surrounding sounds without understanding which speaker the listener intends to attend to, creating major communication challenges for seniors and individuals with hearing impairments in crowded environments such as restaurants, hospitals, and senior living communities across the globe. Recent advances in artificial intelligence, EEG-based neural decoding, and adaptive beamforming create the opportunity for neuro-adaptive hearing systems that can align audio amplification with a user's cognitive focus in real time at low cost. This research explores how low-cost AI-driven hearing assistance technologies can improve accessibility, speech clarity, and quality of life for individuals struggling with hearing loss in noisy social settings.

1.2 Problem

Conventional hearing aids and acoustic beamforming systems cannot determine which speaker a listener is cognitively focused on in multi-speaker environments, often resulting in amplification of unwanted background noise. This research investigates whether EEG-based auditory attention decoding combined with real-time SRP-PHAT and MVDR beamforming can accurately identify a listener's directional attention and increase signal-to-noise ratio (SNR) compared to traditional amplification methods. System performance was evaluated using directional classification accuracy, SNR improvement, and end-to-end latency measurements collected during controlled auditory attention experiments.

2. Investigation

The neuro-adaptive hearing system integrated EEG-based auditory attention decoding with SRP-PHAT direction-of-arrival estimation and MVDR adaptive beamforming. The hardware setup included a BioAmp EXG Pill EEG module, Ag/AgCl electrodes placed at O1 and O2 with mastoid reference, a ReSpeaker 4-Mic Pi HAT v3.0 microphone array, and a Knowles balanced armature receiver. EEG signals were sampled at 250 Hz with electrode impedance maintained below 10 kΩ.

50 participants completed controlled auditory attention experiments using two simultaneous speech streams delivered through stereo headphones. Participants were instructed to focus on either the left or right speaker while minimizing movement. EEG preprocessing included notch filtering, bandpass filtering, alpha-band isolation, and artifact rejection. Hemispheric alpha asymmetry relative to baseline activity was used to

classify attention direction over short decision windows, with approximately 12% of noisy EEG windows excluded during processing. The classified attention direction dynamically steered beamforming toward the cognitively attended speaker in near real time. System performance was evaluated using directional classification accuracy, signal-to-noise ratio (SNR) improvement, artifact rejection percentage, and end-to-end latency measurements across repeated trials. Baseline comparisons were performed between normal amplification, beamforming-only enhancement, and beamforming combined with EEG-guided steering.

EEG preprocessing included notch filtering, bandpass filtering, alpha-band isolation, and artifact rejection. The following equations describe the core computations of the neuro-adaptive pipeline.

Equation 1 - Hemispheric Alpha Asymmetry (EEG Attention Score)

$$A(t) = \frac{P_R(t) - P_L(t)}{P_R(t) + P_L(t)}$$

Where:

A(t)A(t) = normalized auditory attention score at time tt

PR(t)PR(t) = alpha-band EEG power from the right occipital hemisphere

PL(t)PL(t) = alpha-band EEG power from the left occipital hemisphere

This equation represents the hemispheric asymmetry computation described in the Methods section for classifying listener attention direction.

Equation 2 — Signal-to-Noise Ratio (SNR)

$$\text{SNR}_{dB} = 10 \log_{10} \left(\frac{P_{signal}}{P_{noise}} \right)$$

Where:

Psignal = power of the attended speech signal

Pnoise = power of background noise

This equation corresponds to the SNR improvement graph shown in the Results section.

Equation 3 — MVDR Beamforming Weight Vector

$$\mathbf{w}_{MVDR} = \frac{\mathbf{R}^{-1}\mathbf{a}(\theta)}{\mathbf{a}^H(\theta)\mathbf{R}^{-1}\mathbf{a}(\theta)}$$

Where:

wMVDR = adaptive beamforming weight vector

R = spatial covariance matrix of microphone signals

a(θ) = steering vector for direction θ

aH = Hermitian transpose of steering vector

This equation represents the MVDR adaptive beamforming method referenced in the poster bibliography and methodology.

Equation 4 — SRP-PHAT Localization Objective

$$\hat{\theta} = \arg\max_{\theta} \sum_{i<j} R_{ij}^{PHAT}(\tau_{ij}(\theta))$$

Where:

θ̂ = estimated sound source direction

RijPHAT = PHAT-weighted cross-correlation between microphones i and j

τij(θ) = predicted time delay for direction θ

This equation corresponds to the SRP-PHAT direction-of-arrival estimation pipeline described in the project.

Code

```
for each trial in experiment:

    acquire EEG(O1, O2) and microphone array signals

    preprocess EEG:
        apply notch filter
        apply bandpass filter
        isolate alpha band (8–12 Hz)

    compute hemispheric alpha asymmetry

    classify attention direction:
        LEFT or RIGHT speaker

    estimate sound source direction:
        run SRP-PHAT localization

    apply MVDR beamforming toward attended direction

    enhance attended speech signal

    evaluate:
        classification accuracy
        SNR improvement
        end-to-end latency

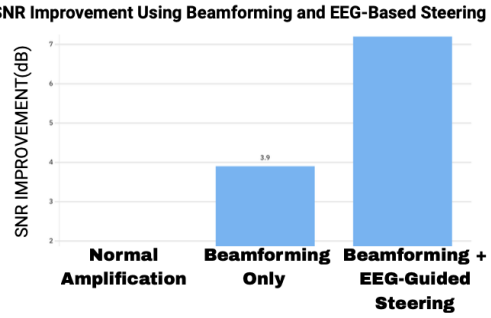
report:
    mean accuracy
    mean SNR improvement
    average latency
    artifact rejection percentage
```

2.2 Results

The neuro-adaptive hearing system achieved above-chance auditory attention decoding using only two occipital EEG channels. Across five participants, the system reached a mean directional classification accuracy of 68%, with best-case performance of 75%, compared to a 50% chance baseline. Average decision latency was approximately 1.4 seconds, with end-to-end latency of 1.6 seconds. Approximately 12% of EEG windows were excluded during artifact rejection.

SNR analysis showed that beamforming alone improved directional noise suppression, while beamforming combined with EEG-guided steering achieved the highest enhancement by aligning amplification with the listener's attentional focus. Results demonstrated that EEG-guided steering provided stronger directional enhancement than acoustic-only beamforming approaches.

Figure 1. Figure X. SNR Improvement Using Beamforming and EEG-Based Steering.



This graph compares signal-to-noise ratio (SNR) improvement across three hearing enhancement approaches: normal amplification, beamforming-only enhancement, and beamforming combined with EEG-guided steering. Conventional amplification provided minimal directional noise suppression, while beamforming alone achieved an SNR improvement of approximately 3.9 dB through spatial filtering. The highest performance was achieved using beamforming combined with EEG-guided steering, which increased SNR improvement to approximately 7.2 dB by dynamically aligning audio amplification with the listener's cognitively attended sound source. These results demonstrate that integrating EEG-based auditory attention decoding with adaptive beamforming can substantially improve speech clarity in noisy multi-speaker environments compared to traditional acoustic-only approaches.

Table 1. Performance comparison of conventional amplification, beamforming-only enhancement, and EEG-guided neuro-adaptive beamforming evaluated across five participants in controlled auditory attention experiments. Higher SNR improvement and directional classification accuracy indicate better performance.

Method	Directional Classification Accuracy (%, ↑)	SNR Improvement (dB, ↑)	End-to-End Latency (s, ↓)
Normal Amplification	50.0 (chance baseline)	0.0	0.0
Beamforming Only	N/A	3.9	0.4
Beamforming + EEG-Guided Steering	68.0 (Best: 75.0)	7.2	1.6
Δ (Improvement vs. Beamforming Only)	+18.0	+3.3	+1.2

These findings support the feasibility of low-cost neuro-adaptive hearing systems capable of selectively enhancing attended speech in noisy multi-speaker environments using near real-time EEG-guided audio steering.

3. Conclusion

3.1 Related Work

Previous studies by Nima Mesgarani, Edward F. Chang, and James A. O'Sullivan demonstrated that auditory attention can be decoded from neural activity in multi-speaker environments. This project extends prior work by combining EEG-based attention decoding with SRP-PHAT localization and MVDR beamforming in a low-cost neuro-adaptive hearing system for real-time cognitively guided audio enhancement.

3.2 Limitations

This study used a small sample size of five participants and only two occipital EEG channels, limiting generalizability across broader populations and real-world environments. EEG recordings were also sensitive to motion artifacts, electrode placement variability, and environmental noise during testing.

3.3 Societal Impact

This research has the potential to benefit seniors, individuals with hearing impairments, and people who struggle to communicate in noisy environments by enabling hearing systems that selectively amplify the speaker a listener is cognitively focused on rather than amplifying all surrounding sounds equally. The low-cost and

multi-lingual hardware architecture used in this project may improve accessibility for underserved communities that cannot afford advanced commercial hearing aids, while the integration of AI and EEG-based attention decoding demonstrates how neuro-adaptive technologies can improve quality of life and social participation.

However, reliance on EEG acquisition may exclude users who are uncomfortable wearing neural sensing devices or individuals for whom reliable EEG signals are difficult to obtain due to physiological or medical factors. Additionally, future deployment of brain-guided assistive systems must carefully address ethical concerns related to neural data privacy, informed consent, and responsible handling of cognitive information collected during real-time operation.

4. References and Additional Details

4.1 References

[1] N. Mesgarani and E. F. Chang, "Selective cortical representation of attended speaker in multi-talker speech perception," Nature, vol. 485, no. 7397, pp. 233–236, 2012.

[2] J. A. O'Sullivan, A. J. Power, N. Mesgarani, S. Rajaram, J. J. Foxe, B. G. Shinn-Cunningham, M. Slaney, S. A. Shamma, and E. C. Lalor, "Attentional selection in a cocktail party environment can be decoded from single-trial EEG," Cerebral Cortex, vol. 25, no. 7, pp. 1697–1706, 2015.

[3] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," Proceedings of the IEEE, vol. 57, no. 8, pp. 1408–1418, 1969.

[4] J. DiBiase, H. Silverman, and M. Brandstein, "Robust localization in reverberant rooms," in Microphone Arrays: Signal Processing Techniques and Applications, Berlin, Germany: Springer, 2001.

[5] J.-M. Valin, "A hybrid DSP/deep learning approach to real-time full-band speech enhancement," in IEEE Multimedia Signal Processing Workshop (MMSP), 2018.

[6] J. Benesty, J. Chen, and Y. Huang, Microphone Array Signal Processing. Berlin, Germany: Springer, 2008.

[7] Upside Down Labs, "BioAmp EXG Pill." [Online]. Available: https://upsidedownlabs.tech

[8] Seeed Studio, "ReSpeaker 4-Mic Pi HAT v3.0." [Online]. Available: https://www.seeedstudio.com

[9] Arduino, "Arduino UNO R4." [Online]. Available: https://www.arduino.cc

[10] Knowles Corporation, "Balanced Armature Receivers for Hearing Applications." [Online]. Available: https://www.knowles.com

4.2 Disclosures & Acknowledgements

AI-assisted development tools, including Cursor and Google Gemini, were used during portions of the software development and technical writing process. These tools assisted with brainstorming, code prototyping, debugging, grammar refinement, and organization of research materials related to the EEG processing and beamforming pipeline. All experimental methodology, system implementation, analysis, and interpretation of results were independently designed, validated, and reviewed by the authors.

This work was partially supported through a $10,000 grant awarded by Lead4Change, which helped advance community outreach, prototyping, and research activities related to neuro-adaptive hearing assistance technologies. The team also appreciates the valuable feedback provided by seniors and participants from local community and senior living centers throughout Santa Clara County during system demonstrations and testing sessions.

The EEG acquisition system used lightweight wearable electrodes for non-invasive neural signal collection.[1]

[1]EEG signals were collected only during controlled experimental sessions and did not store personally identifiable participant information.

Pupil Tracking and Spiking Neural Network Modeling of Horizontal Nystagmus

Samyak Jain
sam.j.272009@gmail.com

1. Introduction

1.1 Background

Nystagmus is a congenital or acquired ocular-motor condition of involuntary, often rhythmic eye movements [1], [8]. The oscillations may be horizontal, vertical, or torsional and vary in waveform, frequency, and amplitude, sometimes reducing visual acuity or causing oscillopsia. Jerk nystagmus has a slow phase with a rapid corrective movement, while pendular nystagmus is more nearly sinusoidal [2]. It can arise from instability in the systems that stabilize gaze, including the vestibulo-ocular reflex, the gaze-holding integrator, brainstem nuclei, and the cerebellum [1], [8].

1.2 Problem

Nystagmus intensity changes with gaze direction, convergence, attention, and fatigue [2], so quantitative measurement is important for characterizing the condition and evaluating interventions. Video-oculography gives objective measures of waveform, amplitude, and frequency that supplement clinical examination [5], [7], and computational models can then test how alternating activity in ocular-motor pathways may generate the oscillations [3]. This study used video-based pupil tracking across several viewing and physiological conditions, together with a two-population neural model, to examine a possible mechanism for the observed rhythms.

2. Methods

2.1 Experimental Recordings

Videos were recorded with a stationary camera at 1920 × 1080 pixels and about 30 frames/s. One participant with nystagmus was recorded under the conditions in Table 1, and one without nystagmus served as a negative control (Experiment 0). Experiment 1 was the straight-ahead reference and the comparison for Experiments 6–8, since separate pre-screen, rested, and without-glasses recordings were unavailable.

TABLE 1
EXPERIMENTAL CONDITIONS AND THEIR PURPOSES

EXPERIMENT	CONDITION	PURPOSE
0	Participant without nystagmus	Negative control
1	Straight-ahead gaze	Baseline nystagmus
2L, 2R, 2U, 2D	Left, right, up, and down gaze	Gaze dependence
3	Reading text	Functional impact
4N, 4F	Near and far fixation	Accommodation effect
5M, 5N	Morning and evening recordings	Fatigue and time-of-day effects
6	After 3 hours of screen use	Visual strain effect
7	After 2 hours of physical activity	Physiological state effect
8	With glasses	Optical correction effect
9	Viewing moving black-and-white stripes	Optokinetic response

2.2 Pupil Tracking

Pupils were extracted with Python and OpenCV: eye regions were converted to grayscale, smoothed, and thresholded, and candidate contours were scored by size, shape, darkness, and proximity to the previous position (Fig. 1). Small facial movements were corrected with sparse Lucas–Kanade optical flow, and pupil coordinates were mapped into a stabilized facial frame. Frames without a confident detection were treated as missing; tracking error is the percentage of such frames (Fig. 2).

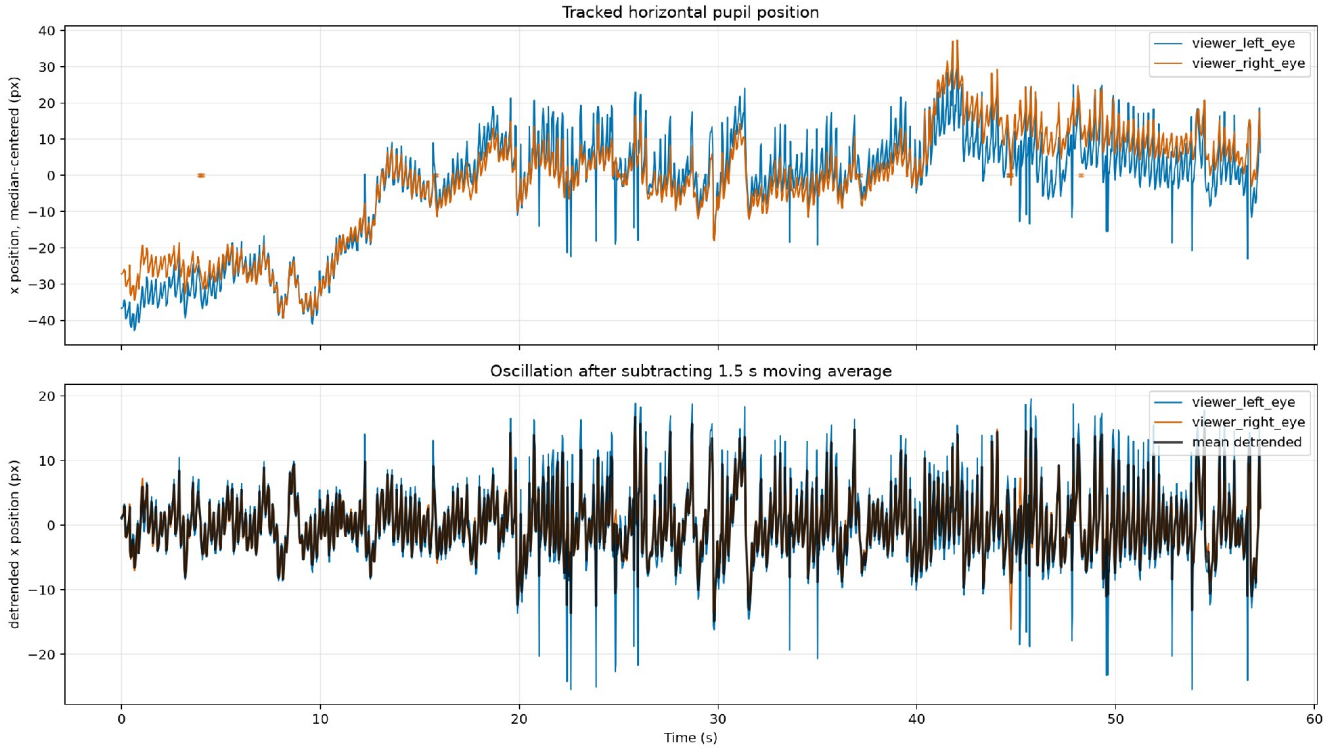


Fig. 1. Representative OpenCV tracking in Experiment 1. Top: median-centered left- and right-pupil positions. Bottom: signals after removing a 1.5-s moving average, with the mean detrended binocular trace in black.

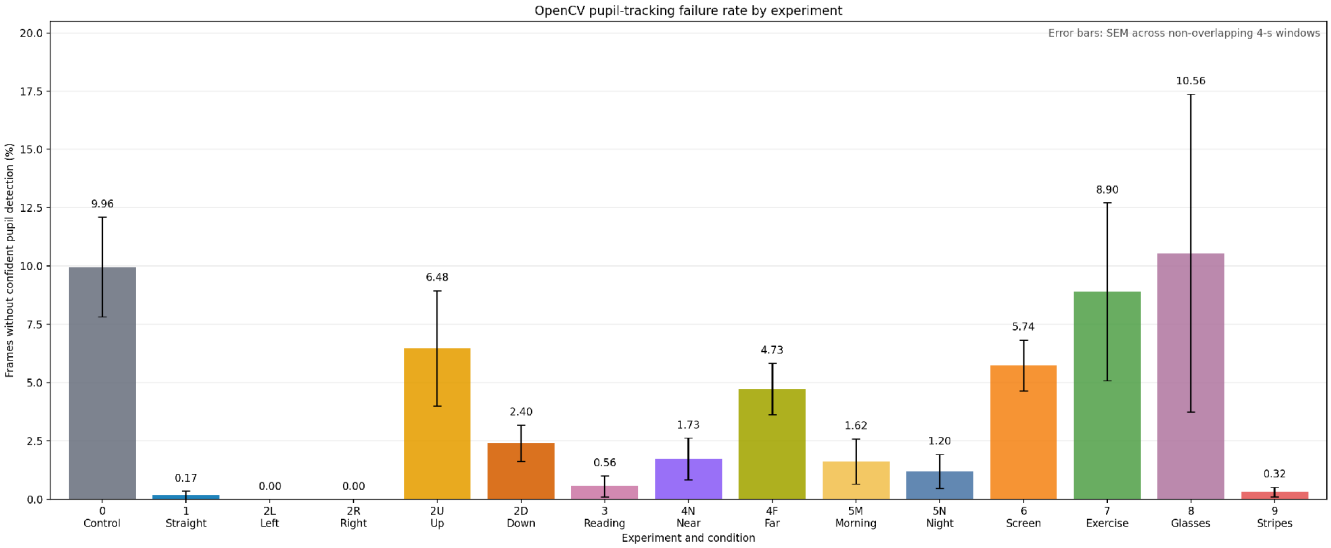


Fig. 2. OpenCV pupil-tracking failure rate by condition (percentage of frames without a confident detection). Error bars show the standard error across 4-s windows.

2.3 Oscillation Analysis

Stabilized horizontal traces were band-pass filtered from 0.5 to 8 Hz with a fourth-order zero-phase Butterworth filter, and a binocular trace was formed by averaging the eyes. Dominant frequency came from Welch spectra and left–right coherence; a rhythm was accepted when coherence was at least 0.35 and its spectral score at least 6 dB above background, otherwise reported as not detected (ND). Amplitude was the 95th-minus-5th-percentile displacement, measured in non-overlapping 4-s windows.

2.4 Neural Oscillator Model

A two-population model tested whether alternating activity in rightward- and leftward-preferring populations, with reciprocal inhibition and adaptation, could reproduce the eye movements. Their dynamics were

$$\begin{aligned} \tau_r \frac{du_R}{dt} &= -u_R - g\,r_L - b\,a_R + I \\ \tau_r \frac{du_L}{dt} &= -u_L - g\,r_R - b\,a_L + I \\ \tau_a \frac{da_R}{dt} &= -a_R + r_R \\ \tau_a \frac{da_L}{dt} &= -a_L + r_L \end{aligned} \tag{1}$$

with rectified population rates

$$r_R = \max(0,\, u_R), \quad r_L = \max(0,\, u_L \tag{2}$$

where u_R, u_L are population activities, a_R, a_L adaptation variables, g reciprocal inhibition, b adaptation strength, and I tonic input. An adaptive leaky integrate-and-fire version used 120 neurons per population, and the difference between the smoothed rates drove a first-order ocular plant,

$$\tau_e \frac{dx}{dt} = -x + K\,(r_R - r_L) \tag{3}$$

Experiment 1 calibrated baseline frequency and pixel gain. To represent conditions, the baseline output x_0 was rescaled by a fitted gate, timescale, and gain,

$$x_c(t) = q_c\,G_c\,x_0\!\left(\frac{t}{s_c}\right) \tag{4}$$

with the gate zero for ND recordings, the timescale setting frequency, and the gain setting amplitude. Because these were fit to the same recordings, the results are in-sample calibration rather than independent prediction.

3. Results

3.1 Measured Pupil Oscillations

Coherent oscillation was detected in 11 of 15 recordings; none met the criteria for the control (Experiment 0), left gaze (2L), up gaze (2U), or far fixation (4F). The straight-ahead baseline (Experiment 1) measured 4.10 Hz and 3.14% of inter-pupil distance. Similar frequencies occurred for down gaze (4.30 Hz), near fixation (3.96 Hz), morning and evening (3.85 and 3.99 Hz), screen exposure (4.23 Hz), and moving stripes (3.77 Hz); slower rhythms occurred for right gaze (0.73 Hz), reading (0.71 Hz), exercise (1.26 Hz), and glasses (2.60 Hz). The largest amplitudes were in the morning (4.46%), reading (4.05%), and evening (4.03%), and the smallest was right gaze (0.43%) (Fig. 3).

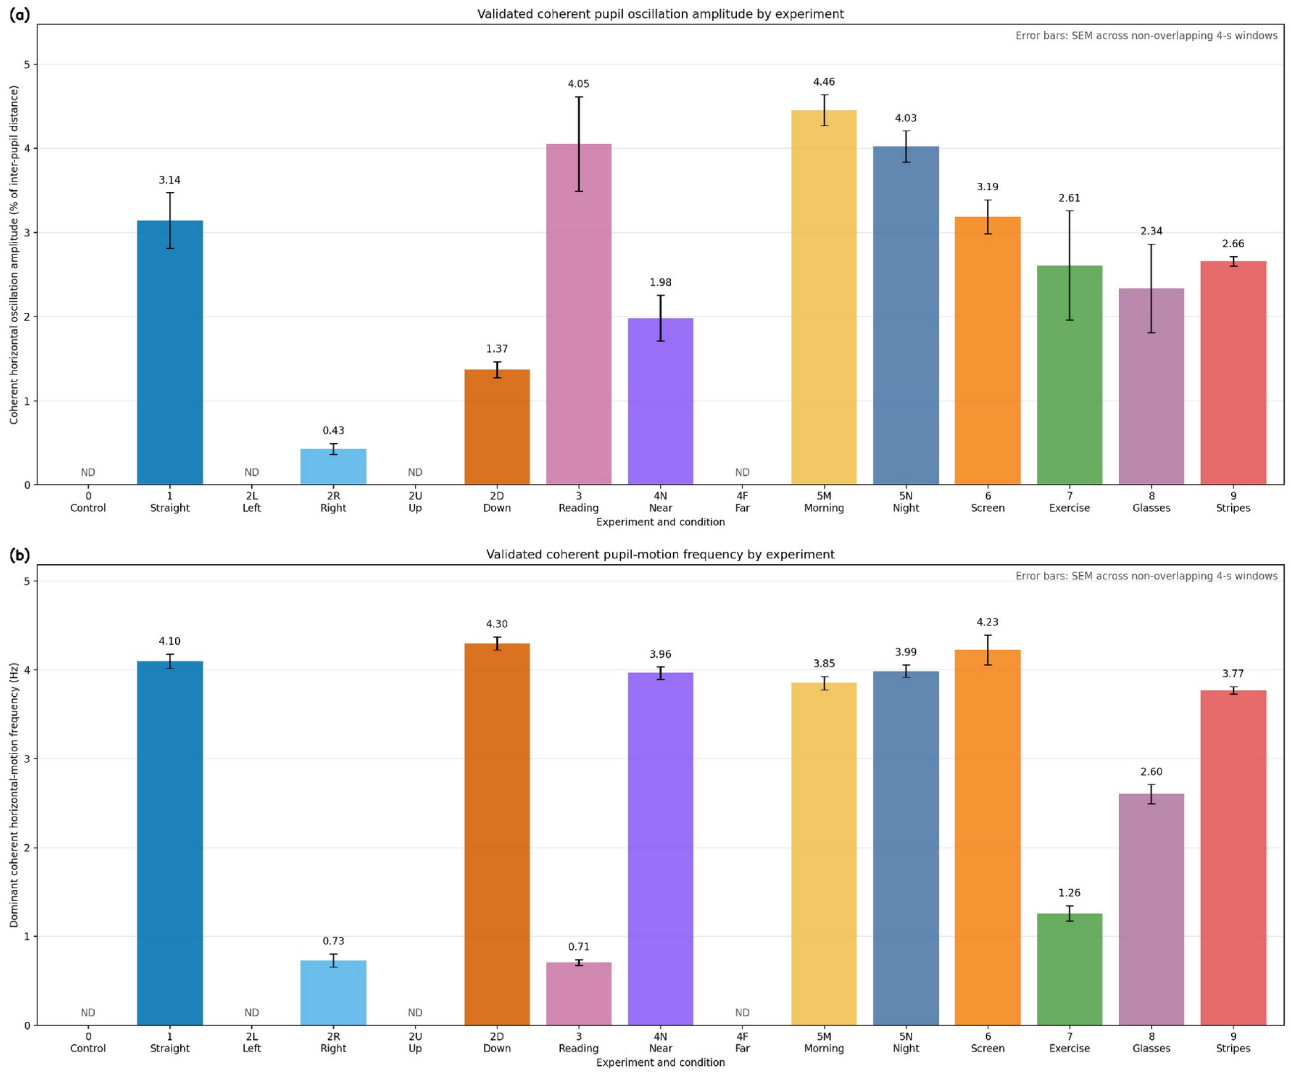


Fig. 3. Measured pupil-oscillation amplitude (a) and dominant frequency (b) by condition. Amplitude is a percentage of inter-pupil distance; error bars show the standard error across 4-s windows. ND indicates no detected oscillation.

3.2 Condition-Aware Model Fit

The model reproduced the detected or ND state of every recording, with amplitude set by the condition gain. Fitted frequencies stayed within two standard errors of measurement (mean absolute error 0.081 Hz across the ten non-baseline detected recordings), and mean phase-aligned waveform correlation rose from 0.357 for the fixed baseline model to 0.663 (Fig. 4). The best matches were Experiment 8 ($r = 0.87$), 9 (0.83), 6 (0.82), and 1 (0.81), while reading was weakest (0.11), consistent with its nonstationary movements.

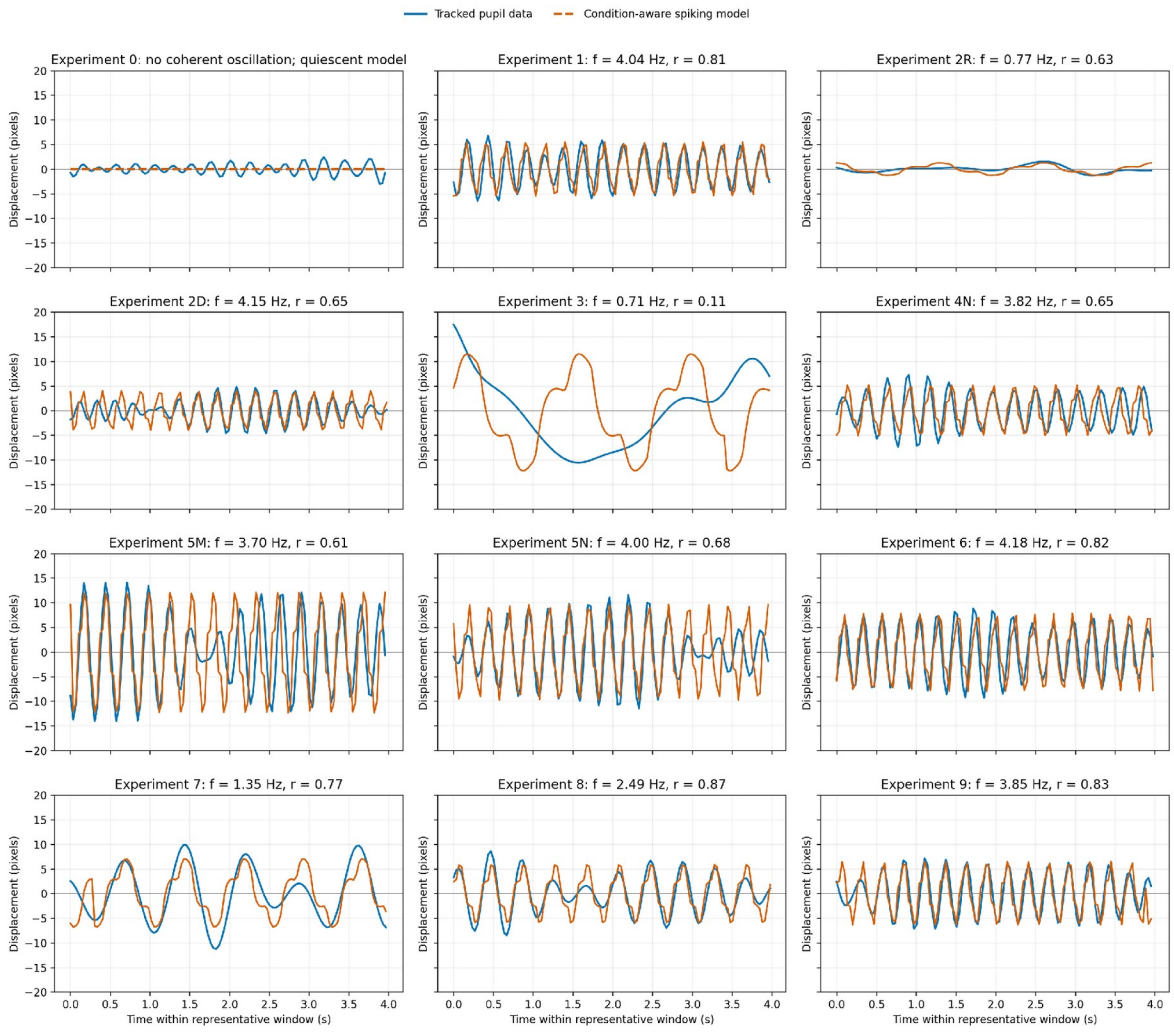


Fig. 4. Tracked pupil displacement (blue) and condition-aware model fits (orange), all at the same vertical scale. Experiment 0 shows the control and quiescent model; f and r are given for detected oscillations.

3.3 Spiking-Network Activity

The spiking network produced alternating population activity that generated oscillatory eye output (Fig. 5). Experiment 1 gave a 4.04-Hz rhythm (r = 0.81), and applying the slower Experiment 7 timescale gave a 1.35-Hz rhythm (r = 0.77), showing that one mechanism can span timescales.

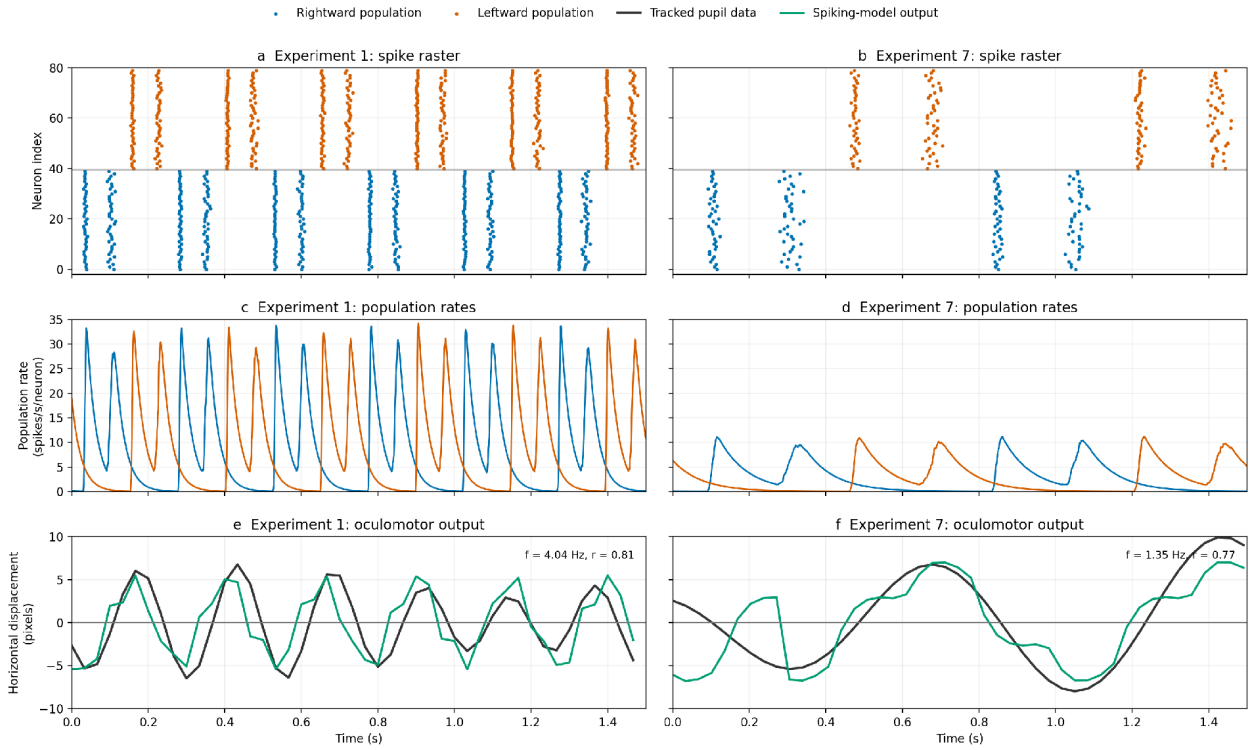


Fig. 5. Spiking-network simulations for Experiments 1 and 7. Alternating population activity drives the modeled eye displacement, giving a 4.04-Hz rhythm (Exp. 1) and a 1.35-Hz rhythm (Exp. 7), correlating with tracked motion at r = 0.81 and r = 0.77.

Increasing reciprocal inhibition drove a transition from balanced to alternating activity (Fig. 6). At g = 0.4 the rates were balanced and output stayed below the control threshold; oscillation appeared at g = 0.8, and at g = 2.0 the model matched the Experiment 1 amplitude and produced a 4.0-Hz rhythm, near the measured 4.10 Hz.

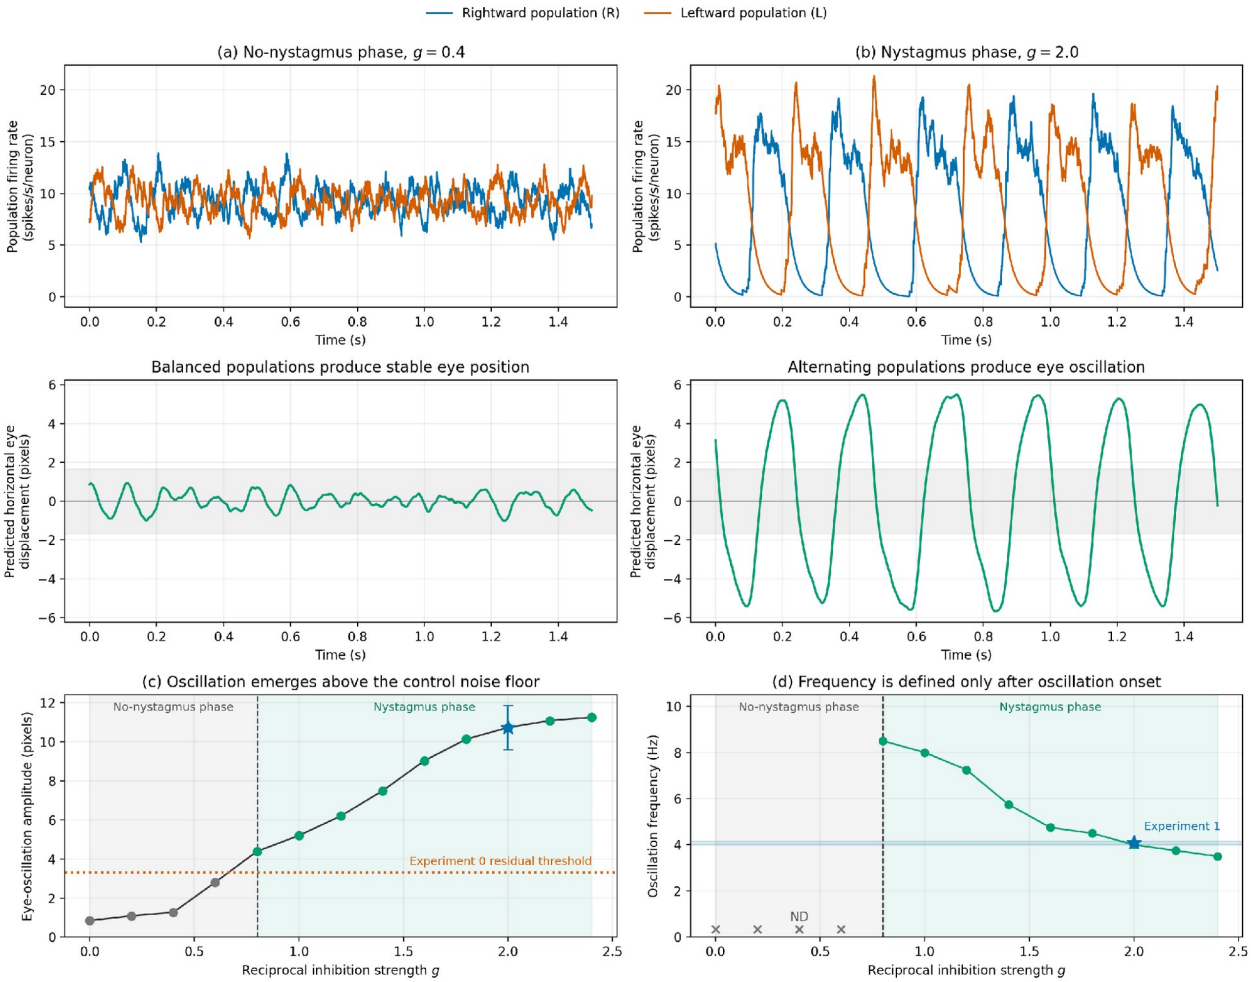


Fig. 6. Transition from no detectable nystagmus to nystagmus-like output. Below g = 0.8 the populations stay balanced and output stays within the control range; at g = 2.0 the simulated amplitude and frequency match Experiment 1.

4. Conclusion

Video-based pupil tracking quantified condition-dependent changes in horizontal nystagmus. The participant with nystagmus showed a 4.10-Hz baseline rhythm, absent in the control, and amplitude and frequency varied with gaze direction, fixation distance, reading, time of day, screen exposure, exercise, optical correction, and moving stimulation, so nystagmus is not a fixed oscillation and should be evaluated under more than one condition.

The two-population model gave a plausible mechanism: alternating population activity produced oscillation while balanced activity produced stable position, and increasing reciprocal inhibition moved the spiking network from a no-detection state to a nystagmus-like state that matched Experiment 1. Limitations include the small number of participants, pixel-based measurement, and condition parameters fit to the same data; future work should add participants, calibrated eye-position measurement, repeated sessions, and independent validation.

References

[1] R. V. Abadi, “Mechanisms underlying nystagmus,” J. R. Soc. Med., vol. 95, no. 5, pp. 231–234, 2002, doi: 10.1177/014107680209500504.

[2] R. V. Abadi and A. Bjerre, “Motor and sensory characteristics of infantile nystagmus,” Br. J. Ophthalmol., vol. 86, no. 10, pp. 1152–1160, 2002, doi: 10.1136/bjo.86.10.1152.

[3] L. F. Dell’Osso, “Biologically relevant models of infantile nystagmus syndrome: The requirement for behavioral ocular motor system models,” Semin. Ophthalmol., vol. 21, no. 2, pp. 71–77, 2006, doi: 10.1080/08820530600613795.

[4] V. L. N. Fu et al., “Visual acuity development of children with infantile nystagmus syndrome,”

Invest. Ophthalmol. Vis. Sci., vol. 52, no. 3, pp. 1404–1411, 2011, doi: 10.1167/iovs.09-4686.

[5] M. Tekavcic Pompe, A. Liasis, and R. Hertle, “Visual electrodiagnostics and eye movement recording: World Society of Pediatric Ophthalmology and Strabismus consensus statement,” Indian J. Ophthalmol., vol. 67, no. 1, pp. 23–30, 2019, doi: 10.4103/ijo.IJO_1103_18.

[6] M. D. Richards and A. Wong, “Infantile nystagmus syndrome: Clinical characteristics, current theories of pathogenesis, diagnosis, and management,” Can. J. Ophthalmol., vol. 50, no. 6, pp. 400–408, 2015, doi: 10.1016/j.jcjo.2015.07.010.

[7] J.-R. Rizzo et al., “Eye movement recordings: Practical applications in neurology,” Semin. Neurol., vol. 39, no. 6, pp. 775–784, 2019, doi: 10.1055/s-0039-1698742.

[8] M. Strupp et al., “Central oculomotor disturbances and nystagmus: A window into the brainstem and cerebellum,” Dtsch. Arztebl. Int., vol. 108, no. 12, pp. 197–204, 2011, doi: 10.3238/arztebl.2011.0197.

Multilayer Coatings for Thermal Management for Robots in Outerspace

Abhijay Konduru
abhijay.konduru28@bcp.org

Keywords / Index Terms
Thermal Management, Space Robotics, Multilayer coatings

1. Introduction

1.1 Background

Human-operated robotic systems are important for exploring the Moon and Mars. These controlled robots can perform tasks such as digging and collecting samples, providing insights into geology and surface environments. During such expeditions, one factor that can affect the durability of a robot on the Moon and Mars is the lack of a protective atmosphere to shield it from extreme temperatures. The Moon lacks an atmosphere, which causes temperatures to swing from a high of 400 K (260°F, 127°C) during the day to a low of 100 K (−280°F, −173.15°C) at night [1]. Similarly, the thin atmosphere on Mars, composed primarily of carbon dioxide, nitrogen, and argon, causes heat to be lost, leading to temperature fluctuations from 70°F (20°C) to as low as −225°F (−153°C), which can cause robots to freeze if not properly protected [2].

To protect against such extreme temperatures, various types of coatings are used on space robots. White coatings, including those made with zinc oxide, titanium dioxide, or silicate, help to reject heat. These materials have low absorptance and high emittance, allowing them to reflect sunlight and release heat. A few white coatings also offer resistance to atomic oxygen. These coatings have an absorptivity ranging from 0.15 to 0.25 and an emissivity greater than 0.85 [3]. A high emissivity (close to 1) means a material absorbs and emits thermal radiation very well, like that of blackbody. A low emissive value (near zero) means that the material reflects or scatters more radiation and absorbs less.

In contrast to white coatings, black coatings are designed to absorb and retain heat, which helps robots stay warm and regulate their temperature. Materials such as carbon-based coatings, black anodized aluminum, and MLS-85SB silicates perform well because they have high absorptance and emissivity. These materials have absorptivity and emissivity values higher than 0.9 [3]. It should, however, be noted that not all materials are reliable for use as coatings. Polyurethane-based blacks do not hold up well against atomic oxygen and are not suitable for all missions. Choosing the right coating depends on the mission environment and the specific temperature required to maintain for each robotic system.

Advanced materials for extreme environments are often described in terms of a "life ring" of three key capabilities: multi-functionality, adaptivity, and self-restoration. In the present work, the proposed multilayer coatings primarily demonstrate multi-functionality, simultaneously providing low solar absorptivity and high mid-infrared emissivity to regulate thermal loads on robotic systems. This enables passive thermal management without additional active cooling systems. Lessons from spacecraft thermal control technologies provide useful guidance

for the design of coatings intended for robotic systems operating on the Moon and Mars. Modern spacecraft commonly employ optical solar reflectors and multilayer dielectric coatings that combine low solar absorptivity with high infrared emissivity to regulate thermal loads [4]. Materials such as SiO_2, Al_2O_3, and TiO_2 are widely used in these coatings due to their optical stability and durability in space environments [4].

1.2 Problem

For this research paper, WPTherml was used [5]. This Python package simulates the optical and thermal properties of multilayer surface coatings to determine how they can aid in thermal management in space. This package uses the transfer matrix method to calculate quantities such as absorptivity and emissivity. It enables the design of materials to control how much heat they emit, which is important for allowing robots to withstand extreme temperatures in space. This research investigates how various factors can be adjusted to design surface coatings that regulate heat and radiation exposure on the Moon and Mars.

2. Investigation

2.1 Methodology

To address the extreme temperature conditions faced by robots in space, simulations of emissivity and absorptivity were carried out to identify candidate materials for thermal management coatings for space robots. WPTherml (v1.1.1) [5], a Python-based package, was used to simulate how different combinations of materials behave under temperatures across the infrared spectrum and visible light range.

This research includes running simulations on both individual and multilayer combinations of the materials (SiO_2, HfO_2, Au, Al_2O_3, TiO_2, Ag) and evaluating their performance across different thicknesses ranging from 10 nm to 1000 nm. This approach allowed the exploration of how specific configurations could enhance thermal emission control and facilitate the identification of suitable materials for protecting robots on the Moon or Mars.

The key performance metric analyzed was net cooling power (P_{cool}), which indicates whether a coating causes cooling (positive P_{cool}) or heating (negative P_{cool}). This metric was used to assess each material's effectiveness in managing heat under space conditions [5]. The transfer matrix method implemented within WPTherml was used to calculate quantities such as absorptivity and emissivity for the multilayer coating stacks [5].

2.2 Results

The objective of this study is to identify a novel surface coating that can help space robots manage heat in extreme environments, such as the Moon or Mars. Graphs generated from the simulations were used to assess whether each coating met the thermal management criteria. Greater absorptance was indicated by a larger area under the absorptance curve, while a larger area under the emissivity curve corresponds to greater emissivity [5].

Most single materials did not perform well for cooling. For example, gold (Au) at 100 nm produced a negative net cooling power, indicating significant heating. In contrast, SiO_2 at 200 nm showed a moderately effective cooling result, meaning it cooled by radiating heat away. However, if only SiO_2 were used, achieving sufficient thermal control would require relatively thick layers, increasing the overall mass of the coating, which is not ideal for space missions where minimizing weight is essential.

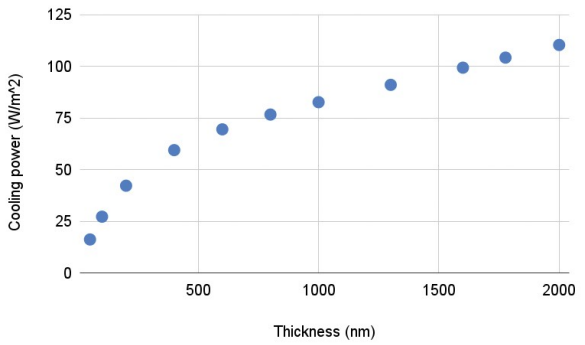


Figure 1. The relationship between thickness and cooling power (P_{cool}) for SiO_2 at various thicknesses. Increasing SiO_2 thickness improves cooling performance, although the rate of improvement decreases at larger thicknesses.

Some dielectric-only multilayer combinations showed improved cooling performance. For example, HfO_2 + TiO_2 at 100 nm each produced positive cooling values, indicating effective cooling behavior. During the exploration of three-layer stacks, Al_2O_3 + TiO_2 + HfO_2, each with a thickness of 100 nm, demonstrated improved cooling performance with low absorptance and high emissivity. These results suggest that multilayer dielectric coatings can improve passive thermal regulation for robots operating on the Moon and Mars.

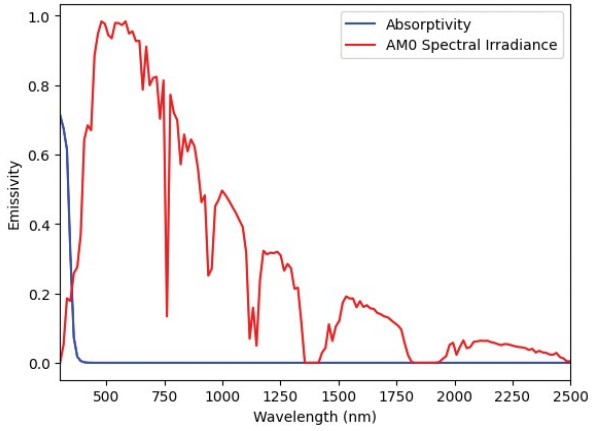


Figure 2: The absorptivity for Al_2O_3 + TiO_2 + HfO_2, each with a thickness of 100 nm

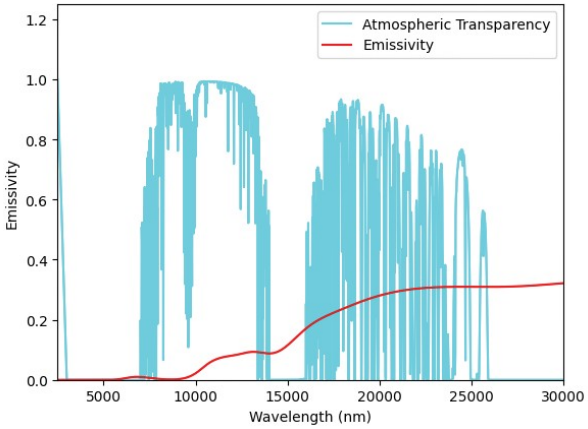


Figure 3: The emissivity for Al_2O_3 + TiO_2 + HfO_2, each with a thickness of 100 nm, overlaid with terrestrial atmospheric transmittance for reference.

3. Conclusion

3.1 Related Work

Previous studies have investigated spacecraft thermal control technologies and multilayer dielectric coatings for regulating heat in extreme space environments. Gilmore described how white coatings with low solar absorptivity and high infrared emissivity are commonly used for spacecraft thermal regulation [3]. Similarly, Baturkin discussed the use of multilayer dielectric coatings and optical solar reflectors for thermal control in small spacecraft systems [4]. Varner et al. developed WPTherml, a Python package that simulates the optical and thermal properties of multilayer materials using the transfer matrix method [5]. Building upon these studies, the present work specifically investigates how multilayer dielectric coatings composed of SiO_2, HfO_2, TiO_2, and Al_2O_3 can improve passive thermal management for robots operating on the Moon and Mars.

3.2 Limitations

This study is limited to simulations performed using WPTherml and does not include experimental validation of the proposed coatings under real lunar or martian

environmental conditions. Additionally, only a limited set of materials and thickness combinations were evaluated, meaning other multilayer configurations could potentially achieve greater cooling performance. Factors such as long-term radiation exposure, atomic oxygen degradation, and manufacturing constraints were also not experimentally investigated.

3.3 Societal Impact

Improved passive thermal management coatings could benefit future robotic exploration missions on the Moon and Mars by increasing the durability and operational lifespan of robotic systems while reducing dependence on active cooling systems. Multilayer dielectric coatings that provide high emissivity and low absorptivity may also reduce overall energy consumption and improve mission efficiency. These technologies could support scientific exploration, autonomous robotic operations, and future space infrastructure development.

However, the development of advanced thermal coatings may increase manufacturing complexity and material costs. Some materials used in multilayer coatings may also require resource-intensive production processes. Additionally, the results presented in this work are simulation-based and further experimental testing is required before these coatings can be implemented in practical robotic systems for extraterrestrial missions.

4. References and Additional Details

4.1 References

[1] A. R. Vasavada, D. A. Paige, and S. E. Wood, "Near-surface temperatures on Mercury and the Moon and the stability of polar ice deposits," Icarus, vol. 141, no. 2, pp. 179–193, 1999. [Online]. Available: https://doi.org/10.1006/icar.1999.6175

[2] G. M. Martínez, N. O. Renno, and E. Fischer, "The modern near-surface martian climate: A review of in-situ meteorological data from Viking to Curiosity," Space Science Reviews, vol. 212, pp. 295–338, 2017. [Online]. Available: https://doi.org/10.1007/s11214-017-0360-x

[3] D. G. Gilmore, Spacecraft Thermal Control Handbook, Volume I: Fundamental Technologies. El Segundo, CA, USA: Aerospace Press, 2002. [Online]. Available: https://doi.org/10.2514/4.989117

[4] V. Baturkin, "Micro-satellites thermal control: Concepts and components," Acta Astronautica, vol. 56, no. 1–2, pp. 161–170, 2005. [Online]. Available: https://doi.org/10.1016/j.actaastro.2004.09.003

[5] J. F. Varner, N. Eldabagh, D. Volta, R. Eldabagh, and J. Foley, "WPTherml: A python package for the design of materials for harnessing heat," Journal of Open Research Software, vol. 7, p. 28, 2019.

Comparative Analysis of Modular, Sequential, and Parallel Neural Architectures on Inference Accuracy and Latency

Aryan Kulkarni
aryanamol10@gmail.com

Bellarmine Academic Research Symposium • STEM in Society, Vol. 1 (2026) • Extended Abstract

1. Introduction

1.1 Background

Large language models (LLMs) and neural networks more broadly have become foundational tools across scientific and commercial domains. As datasets grow in volume and complexity, the choice of architectural framework increasingly determines a model's ability to extract meaningful patterns within acceptable computational budgets. While the literature covers individual architectures, direct comparisons under controlled conditions—holding data, hardware, and optimizer fixed—remain relatively sparse.

A motivating example is the binary classification of speaker loudness. Given plotted data points for large and small speakers, an LLM must scan each data point, recognize patterns in increase and decrease, and output a confident prediction. Figure 1 illustrates this simple linear scenario, where the model outputs a high-confidence prediction (Large: 0.95, Small: 0.05) from a linear decision boundary.

Figure 1: Illustrative example of binary classification using a linear decision boundary. The LLM identifies speaker size (large vs. small) from loudness-vs-instances scatter data.

As data complexity increases for additional categories, overlapping distributions, or nonlinear boundaries, the model must produce multiple ranked predictions. Figure 2 shows this scenario, where three simultaneous predictions (Prediction 1: 0.56, Prediction 2: 0.39, Prediction N: 0.05) are required, and the line of best fit curves to accommodate the more complex structure.

Figure 2: Multi-class prediction scenario with a curved decision boundary. Increased data complexity requires the LLM to maintain multiple ranked probability estimates simultaneously.

The central challenge is the latency-accuracy trade-off, where layer-by-layer processing introduces substantial latency and can cause increased likelihood of error. Architectural decisions about how computation is organized fundamentally affect both the speed and accuracy of inference.

1.2 Problem

This paper investigates three neural network paradigms: (1) sequential (hierarchical layer stacking), (2) parallel (dual-branch processing), and (3) modular (decomposition via neural chains). How do these architectural paradigms compare in final regression accuracy (MSE) and training time across tasks of increasing nonlinear complexity, when all other training variables are held constant?

2. Investigation

2.1 Methodology

To isolate the effect of architecture, all experimental variables except network structure were held constant. Three models were evaluated across three synthetic regression environments, each defined by a distinct generating function of escalating complexity. Common parameters: training samples (N) = 5,000; neurons per layer = 64; optimizer = Adam (default learning rate); training epochs = 100; hardware = A100/L4 GPUs (identical across all runs).

Three synthetic data environments were generated over the domain $X \in [0, 10]$ with N = 5,000 uniformly spaced samples. Environment 1—Linear: $y = 2X + \varepsilon$, $\varepsilon \sim N(0, 1)$, a standard baseline. Environment 2—Curved: $y = 0.5X^2 + \sin(X) + \varepsilon$, $\varepsilon \sim N(0, 2)$, combining polynomial and sinusoidal components. Environment 3—Chained: $y = X \cos(0.5X) + \sin X \cos(0.5X) + 0.5 \operatorname{sgn}(\sin(X)) + \varepsilon$, $\varepsilon \sim N(0, 0.5)$, a multi-stage nonlinear function designed to test subtask decomposition. Figure 3

shows representative scatter plots for each environment.

Figure 3: Scatter plots of the three synthetic regression datasets. Left: Linear ($y = 2X + \varepsilon$). Center: Curved ($y = 0.5X^2 + \sin(X) + \varepsilon$). Right: Chained (multi-stage nonlinear composition with noise).

Sequential model: Implemented in Keras/TensorFlow. Fully connected feed-forward network with decreasing layer widths (64 → 32 → 16 → 1).

Parallel model: Implemented in TensorFlow. Dual-branch architecture, each branch receiving the full input; outputs concatenated before a final linear projection.

Modular model: Implemented in PyTorch. Chain of three sub-networks, each receiving the output of the previous stage, targeting a different challenge of the regression problem.

Two metrics were recorded for each model-environment pair: (1) final training MSE at epoch 100, and (2) total wall-clock training time in seconds. Both metrics were collected across two identical independent runs to verify consistency.

2.2 Results

Figure 4 presents the MSE training curves across all three environments and models over 100 epochs. All models exhibit rapid initial loss reduction followed by stabilization, though the rate of descent and final plateau differ meaningfully across architectures. In the linear environment, the modular model (green) descends to the lowest plateau, while sequential and parallel converge to similar but higher loss values. In the curved environment, a similar pattern holds. In the chain environment, the ranking inverts: the sequential model (blue) converges to the lowest final loss ($\sim 10.4 \times 10^1$), while the modular model converges more slowly to a marginally higher peak ($\sim 10.6 \times 10^1$).

Figure 4: MSE training loss curves over 100 epochs for Sequential (blue), Parallel (orange), and Modular (green) architectures across Linear (left), Curved (center), and Chain (right) environments.

Table 1 reports the final MSE and training time for each model-environment combination.

Table 1. Final MSE and Training Time (s) by Model and Environment

Model	Linear		Curved		Chain	
	MSE	Time (s)	MSE	Time (s)	MSE	Time (s)
Sequential	29.38	37.80	29.46	239.64	30.08	10.39
Parallel	30.31	37.83	30.93	240.53	30.38	10.44
Modular	18.67	38.16	17.72	245.68	18.60	10.55

Note: Lower MSE indicates higher accuracy.

The modular model achieved the best accuracy in two of three environments. In the linear case, it reduced MSE to 18.67—approximately 37% lower than sequential (29.38) and parallel (30.31). In the curved case, it achieved an MSE of 17.72 versus 29.46 (sequential) and 30.93 (parallel)—again roughly 40% lower. The exception is the chain environment, where the modular model's MSE (18.60) was the highest of the three despite training for the longest time (10.55 s). The sequential model achieved its best relative outcome in the chain environment: MSE of 30.08 at 10.39 seconds—the fastest training time across that scenario. The parallel model demonstrated consistent but middling performance across all environments, with the smallest variance in both MSE and training time.

3. Conclusion

3.1 Related Work

Jacobs et al. [1] introduced the Mixture-of-Experts paradigm, which motivates modular decomposition by showing that dedicated sub-networks can specialize effectively on distinct subtasks. LeCun et al. [3] demonstrated that hierarchical feed-forward architectures suffer from error propagation across layers—a limitation the sequential model exhibits in our curved environment. Szegedy et al. [5] showed that parallel multi-branch architectures improve representational efficiency in deep networks; our parallel model's consistency across environments supports this finding, though it does not achieve peak accuracy. Our work differs by directly comparing all three paradigms under identical training conditions on synthetic regression tasks of controlled complexity, isolating the architectural variable.

3.2 Limitations

All experiments used synthetic data with known generating functions; generalization to real-world datasets with unknown structure is not guaranteed. Modular subtask definitions were hand-specified, and an automated decomposition strategy might produce different results. All models were evaluated at a fixed scale (64 neurons, 5,000 samples); architectural rankings may shift at different parameter counts or dataset sizes.

3.3 Societal Impact

This research benefits practitioners who must select neural architectures under compute constraints—particularly in resource-limited settings such as edge devices or under-resourced research labs. By providing concrete accuracy-latency trade-off data, it enables more informed decisions that reduce unnecessary computation and energy use. The primary risk of harm is over-generalization: practitioners may apply these findings to domains (e.g., image classification, NLP) where the synthetic regression results do not transfer, leading to poorly performing or inefficient deployed models. The findings should be treated as directional guidance rather than definitive architectural prescriptions.

4. References and Additional Details

4.1 References

[1] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, “Adaptive mixtures of local experts,” Neural Computation, vol. 3, no. 1, pp. 79–87, 1991.

[2] A. Kulkarni, “Neural architecture comparison experiments [Source code],” GitHub, 2025. [Online]. Available: https://github.com/aryanamol10/Research-Paper-Comparative-Analysis-of-Neural-Models

[3] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, “Backpropagation applied to handwritten zip code recognition,” Neural Computation, vol. 1, no. 4, pp. 541–551, 1989.

[4] A. Stöffelbauer, “How large language models work,” Data Science at Microsoft, Medium, Oct. 24, 2023. [Online]. Available: https://medium.com/data-science-at-microsoft/how-large-language-models-work-91c362f5b78f

[5] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2015, pp. 1–9.

4.2 Disclosures & Acknowledgements

AI Use: [Disclose any AI use for brainstorming, editing, or code assistance per paper guidelines if applicable.]

Mentor/Funding Acknowledgements: [Add acknowledgements during camera-ready.]

Effect of Rock Climbing on Neurocognitive Performance and Memory

Alexandra Lu and Athena Li
alexandralu08@gmail.com and athenalinosli@gmail.com

Keywords / Index Terms
Memory, Cognition, Exercise, Climbing

1. Introduction

1.1 Background

Exercise is shown to improve memory and executive functioning. Prior studies show that intermediate-intensity exercise significantly increases performance on working memory tasks [1] and that certain sports led to a more significant increase in memory [2].

Since rock climbing combines exercise with cognitive problem-solving, it may significantly enhance cognitive functioning.

1.2 Problem

Our goal was to test whether exercise improved memory test scores, and whether climbing had a greater impact on scores compared to cardiovascular exercise.

We hypothesized that both rock climbing and cardiovascular exercise would increase number-recall scores, with climbing producing the greater increase.

2. Investigation

2.1 Methodology

Participants were recruited from a Bay Area high school climbing club and completed an informed consent form.

Data collection was performed over three climbing meetups at a climbing gym. For those who attended multiple meetups, their scores were counted separately for each meetup.

Participants were assigned to one of two exercise conditions by counting off in a within-subjects, crossover design. Before exercise, participants completed a baseline number memorization test: reciting a growing sequence of random digits to a partner, with the sequence's contents changing each time. The longest sequence they could recall was their score.

The first exercise group did 15 minutes minute-on, minute-off jumping jacks. The second group alternated between 1 minute of climbing and resting. After the exercise blocks were over, participants took the number recall test again, switched exercise, and took the test a final time.

2.2 Results

We gathered 25 data points over 3 sessions. The average score at baseline was 7.92 (SD 1.52). After climbing the number was 8.60 (SD 1.29), and after jumping jacks the number was 8.44 (SD 1.29) (Table 1).

A Repeated-Measures ANOVA found a statistically significant difference between all three conditions (). Paired Samples T-Tests found a statistically significant difference between the baseline score and post-cardio score (, as well as baseline score and post-climbing score (). There was no statistically

significant difference between the post-cardio and post-climbing scores (4).

2.3 Discussion

In this experiment we found that exercise increased number recall ability. However, there was no statistically significant difference between climbing and cardio, suggesting that the type of exercise does not make a big difference.

3. Conclusion

3.1 Related Work

Prior research found that experienced climbers had higher accuracy on climbing-related memory tests compared to inexperienced climbers, but not on a general memory test [3]. Our data adds to this by testing the differences between climbing and cardio using number recall tasks.

3.2 Limitations

Our small sample size led to less statistical power. Participants were recruited from one school, limiting generalizability. The memory test could be subject to human error or miscounting. Furthermore, students took the baseline test first and could have developed familiarity with the test and thus scored higher, though there was no statistically significant difference between scores of individuals who attended multiple sessions despite prior exposure to the test.

3.3 Societal Impact

Our findings provide evidence that exercise improves memory. These findings could benefit adolescents' education by encouraging schools to incorporate exercise to increase cognitive functioning. Further research could determine the additional benefit of rock climbing compared to other forms of exercise.

4. References and Additional Details

4.1 References

[1] McMorris, T., Sproule, J., Turner, A., & Hale, B. J. (2011). Acute, intermediate intensity exercise, and speed and accuracy in working memory tasks: A meta-analytical comparison of effects. Physiology & Behavior, 102(3–4), 421–428. https://doi.org/10.1016/j.physbeh.2010.12.007

[2] Zach, S., & Shalom, E. (2016). The influence of acute physical activity on working memory. Perceptual and Motor Skills, 122(2), 365–374. https://doi.org/10.1177/0031512516631066

[3] Garrido-Palomino, I., Giles, D., Fryer, S., González-Montesinos, J. L., & España-Romero, V. (2024). Cognitive function of climbers: An exploratory study of working memory and climbing performance. The Spanish Journal of Psychology, 27, e24. https://doi.org/10.1017/SJP.2024.25

The jamovi project. (2026). *jamovi* (Version 2.7.38.0) [Computer software]. https://www.jamovi.org

Table 1. Average score from all participants over all 3 study days (N=26)

Method	Baseline	Post-cardio	Post-climb
Result	7.92	8.44	8.60
Standard Deviation	1.53	1.29	1.29

Table 2. Repeated-measures, within-subjects ANOVA analysis of scores for all participants across all three conditions

	Sum of squares	df	Mean Square	F	p
RM Factor 1	6.32	2	3.16	4.42	.017
Residual	34.35	48	0.72		

Table 3. Paired Samples T-Tests between each pair of conditions for all participants

		statistic	df	p
Baseline vs. post-cardio	Student's t	-2.59	24	.016
Baseline vs. post-climb	Student's t	-2.28	24	.032
Post-cardio vs. post-climb	Student's t	-0.78	24	.444

Data-Optimized Rebalancing Apparatus for Load-Balance Monitoring in Hiking Backpacks

Neethika Vijay, Justin Cho and Wyatt Eros
nvijay27@ndsj.org, yongjustinchoi@gmail.com and wyatt.eros@gmail.com

Keywords / Index Terms

Hiking, backpacks, injury prevention, smart textiles, wearable technology

1. Introduction

1.1 Background

Hiking has become the most popular outdoor activity in the United States, with participation doubling over the past decade to reach 63 million hikers annually. [1] This surge in outdoor recreation has brought a parallel rise in preventable injuries, particularly muscle strains, shoulder pain, and lower back compression caused by improper backpack loading and weight imbalance. [2] About 15% of hikers report back pain during or after their hikes, and roughly 12% experience shoulder or neck injuries directly tied to unevenly distributed loads or poorly adjusted straps. [3]

1.2 Problem

Current hiking backpacks lack real-time load-distribution feedback, leaving hikers reliant on subjective awareness. This product-design-geared project asks: **Can a backpack prototype integrating dual textile-based force sensors, an Arduino Circuit Playground, and NeoPixel/haptic actuators reliably detect shoulder load asymmetry and deliver real-time visual and vibrational feedback to the user?** We evaluate success by whether the sensor produces a measurable resistance drop over 200Ω under 20lbs force, and whether the system triggers alerts when analog-read difference exceeds a calibrated threshold of 25–30 units.

2. Investigation

2.1 Methodology

We modified a commercial backpack with handmade custom straps embedding sensors and feedback. (Figure 1)



Figure 1. The strap side of the V0 prototype.

Straps were made of layered ripstop nylon, TPE, memory foam, and 3D spacer mesh. Textile-based force sensors using three 3D spacer mesh layers with knitted conductive thread sandwiched by silver-plated nylon were inserted into the top of the straps. (Figure 2) Iterative testing of six sensor models (placing a 20lb weight on each sensor; using voltmeter to measure resistance drop) identified this design as optimal in comparison to others tested.

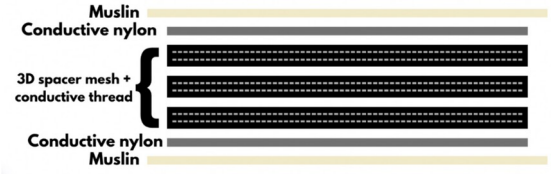


Figure 2. A digital diagram of the material composition of the force sensor.



Figure 2. The six iterations of tested force sensors. (For Table 1: Left to Right A, B, C, D, E, F)

An Adafruit Circuit Playground powered by double-A battery was mounted inside the pack. Sensors connected via 470Ω voltage dividers to analog pins A3 (left) and A2 (right). Two 3-LED NeoPixel strips attached to pins A7/A6. (Figure 3) DRV2605 haptic motor routed along interior back panel. (Figure 4) Firmware read sensors, applied right-side calibration (0.75×), computed difference, and updated outputs. At >25-unit difference, LEDs displayed load-tier colors and haptic motor buzzed until balanced.

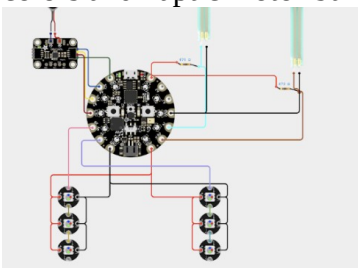


Figure 3. Digital diagram of circuit inbuilt into the back panel/straps of backpack.

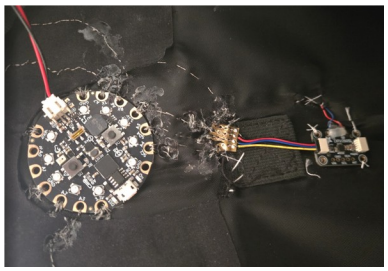


Figure 4. The interior back panel of the backpack. Adafruit Circuit Playground and haptic vibration motor are clearly shown.

2.2 Results

Final sensor achieved 401.2→1.2Ω range, suitable for analog reading. (Table 1) Bench tests confirmed ≥25-unit differences triggered responses within ~50 ms. Adafruit haptic motor provided stronger vibration than Lilypad alternative. No formal human trials conducted; validation limited to controlled compression and serial logging.

Model	Materials	Construction + Design	Range from Voltmeter
A	• Conductive Thread • 3D spacer mesh • Silver-plated woven nylon	• Conductive thread sewn through the face of the 3D spacer mesh • Conductive Fabric Sandwiching the 3D mesh	21-4 Ω Dropped when pressed
B	• Conductive Thread • 3D spacer mesh • Silver-plated woven nylon	• Conductive thread sewn horizontally through the sides • The tail and head of the string is left out to conduct with the fabric • Conductive Fabric Sandwiching the 3D mesh	Didn't work
C	• Velostat • Conductive tape	• Velostat center • Conductive tape on either side	230-265 Ω Increased when pressed
D	3D spacer mesh	• Conductive Fabric Sandwiches the mesh	No significant change in resistance
E	• Conductive Thread • 3D spacer mesh	• Conductive thread sews partially through	Inconsistent range
F	• Velostat • Silver-plated woven Nylon	8 ply sandwich switching between nylon and velostat	432-40.2 Ω Dropped when pressed

Table 1. Composition and voltmeter testing results of six sensor iterations.

3. Conclusion

3.1 Related Work

Prior textile pressure sensors used silver nanowires and Velostat for medical applications. [4], [5] Our work adapts these to hiking, prioritizing low-cost hand-sewn construction and dual-mode feedback.

3.2 Limitations

This project did not include rigorous on-body trials with human participants, so sensor performance under long-term dynamic hiking conditions remains unverified.

3.3 Societal Impact

Beginner hikers and general backpack wearers may benefit from accessible load-balance feedback that could reduce strain injuries. However, reliance on LEDs and vibration excludes users with sensory or visual impairments unless alternative modalities are added.

4. References and

Additional Details

4.1 References

[1] Outdoor Industry Association, "2026 Hiking Report: Executive Summary," 2026. [Online]. Available:https://oia.outdoorindustry.org/l/51282/2026-01-20/f49h7b

[2] Light Hiking Gear, "The Science of Balance: How Load Distribution Changes Hiking Endurance," 2026. [Online]. Available: https://lighthikinggear.com/blogs/hiking/the-science-of-balance-how-load-distribution-changes-hiking-endurance

[3] KURU Footwear, "59 Hiking Injury Statistics and Trends 2026," 2025. [Online]. Available: https://www.kurufootwear.com/blogs/articles/hiking-injury-statistics

[4] M. Kciuk et al., "Intelligent Medical Velostat Pressure Sensor Mat Based on Artificial Neural Network and Arduino Embedded System," Appl. Syst. Innov., vol. 6, no. 5, p. 84, 2023, doi: 10.3390/asi6050084.

A. Chrusch and M. Kavin, "Survey of Musculoskeletal Injuries, Prehike Conditioning, and On-Trail Injury Prevention Strategies Self-Reported by Long-Distance Hikers on the Appalachian Trail," Wilderness Environ. Med., vol. 32, no. 3, pp. 322–331, Sep. 2021, doi: 10.1016/j.wem.2021.04.004.

Y. Lian et al., "Ultrasensitive Wearable Pressure Sensors Based on Silver Nanowire-Coated Fabrics," Nanoscale Res. Lett., vol. 15, no. 1, p. 70, 2020.