Tracking Task Progress with a Proactive Video Recording System

Michael Brilka
University of Siegen
Siegen, Germany
michael.brilka@uni-siegen.de

Kristof Van Laerhoven
University of Siegen
Siegen, Germany
kvl@eti.uni-siegen.de

Abstract

Reliving past events is a crucial aspect of visual research, enabling deeper insights into both the sequence and context of actions. Ideally, a continuous recording would ensure that no moment is lost; however, this approach quickly becomes impractical due to the immense effort required to review and analyze large amounts of footage. To address this challenge, prior work has introduced devices equipped with sensors and cameras that automatically trigger recordings when predefined conditions are met. While this selective recording approach reduces unnecessary footage, it is limited to capturing only the present and lacks retrospective context, which is essential for analyzing complex scenarios. We introduce ChronoVault, an open-source framework designed for simple setup, configuration, and recording of retrospective videos. ChronoVault continuously monitors hardware- and software-signals, buffering a video-stream and saving said buffer, when specific conditions are detected. The system is built on the Raspberry Pi Zero 2, which supports camera integration and connectivity with peripheral devices. To evaluate ChronoVault we performed a user study. In it, participants were tasked to build a model out of building blocks and document the ends of certain subassembly tasks. Compared to an external reference camera, we reduce the footage amount, while retaining all relevant events.

CCS Concepts

• **Human-centered computing → Ubiquitous and mobile computing**.

Keywords

wearable, proactive video recording, data reduction

ACM Reference Format:
Michael Brilka and Kristof Van Laerhoven . 2025. Tracking Task Progress with a Proactive Video Recording System. In *Companion of the 2025 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '25), October 12–16, 2025, Espoo, Finland*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3714394.3756257

UbiComp Companion '25, Espoo, Finland

ACM ISBN 979-8-4007-1477-1/2025/10
https://doi.org/10.1145/3714394.3756257

Figure 1: A fully assembled ChronoVault prototype, which has the form factor of a clip-on module with wide-angle front-facing camera and large-surface button (bottom front).

1 Introduction

Documenting critical steps of a task is essential—but how can this be achieved effectively? A straightforward solution is to continuously record all activity. While this ensures no information is lost, it results in large volumes of footage that must be manually reviewed to extract relevant content. Moreover, the presence of a camera may alter participant behavior, potentially affecting the authenticity of the recorded procedure. One possible workaround is to allow participants to privately review and delete unwanted sections of the video. However, we argue this is suboptimal, as it does not address the core issues of the inefficient recording method and instead burdens participants with additional tasks. To overcome these limitations, prior research has proposed custom recording systems that rely on sensor data [18] or live image recognition [14] to automatically detect and record relevant events. These systems are effective when the conditions that warrant recording can be well-defined. Alternatively, some approaches have leveraged subconscious bodily responses [9, 15] to trigger recording events. While these devices reduce the volume of video that must be analyzed and therefore reflect participants' implicit intentions better, they suffer from a key limitation: they do not incorporate active and consensual input from the participant. Furthermore, because recording begins only after certain conditions are met, the systems often miss important contextual information preceding the event.

To address these shortcomings, we propose a novel recording system that empowers participants with the ability to retroactively save video based on their own proactive intent. Our system maintains a rolling buffer of video frames, continuously overwriting the oldest frame with new ones. When the participant decides to save an event, the system writes the buffered frames to persistent memory, allowing past moments to be recovered and relived.

This approach not only preserves relevant context but also results in a collection of smaller, self-contained video clips rather than a single, continuous stream. This segmentation can significantly reduce the cognitive load on researchers during analysis.

Overall, the contributions of this work are threefold:

- We developed a recording framework to save *retroactively* short, egocentric video clips with accompanying sensor data
- Prototypes were used in an in-lab environment, in which 10 participants were instructed to build a 57-piece lego model
- Results show that this leads to a reduction of produced video footage for a manual construction task, at a low cognitive cost for the participants

2 Related work

Many researchers are confronted with the dilemma of video recording. While the ability to rewatch a scene offers significant advantages, it also raises numerous legal and ethical concerns. Researchers such as Wolf et al. have highlighted the complexities of recording, particularly in public spaces, citing issues related to privacy, consent, and data protection [16, 17, 19]. Although delegating the responsibility for recording to participants may seem like a convenient solution from a researcher's perspective, this approach often proves unreliable in practice. Participants may forget to carry or activate the device [15], leading either to missed recordings or to the unintended capture of private moments.

To address these challenges, several systems have been developed that reduce participant burden and enhance privacy protections. These approaches typically rely on neural networks to detect whether a relevant event is in view [14] or whether specific contextual conditions are met [15]. While effective in predefined scenarios, such methods require a clear and consistent definition of the observable event, something that may not be available in exploratory research contexts. Furthermore, advances in these systems have focused on improving energy efficiency [18] and security features [8, 11, 13], but they often reduce the participant to little more than a passive "camera on legs."

A key limitation of these systems is their dependence on existing datasets. When researchers seek to investigate novel areas without established data, these solutions become impractical.

Commercial products like dashcams offer a partial solution. They continuously record and overwrite older footage unless a manual input—typically the press of a button—marks a clip as protected. This mechanism enables auto-labeling of key events. However, dashcams are designed to record only video and GPS data, with no support for external sensors or additional context [7]. Modifying them for research purposes would require significant effort and would likely not meet GDPR or institutional ethical standards.

A related commercial solution is the bodycam. For instance, a bodycam by Motorola offers the capability to record retroactively: pressing a button saves the two minutes before and after the event. It also includes features such as automatic triggering when a weapon is unholstered or when nearby devices start recording, along with real-time tracking and remote maintenance [3]. However, its high price (approximately €840, excluding tax) [2], and inclusion of features unnecessary or undesirable for research,such as group synchronization and remote control, make it less suitable. Moreover, being a closed and proprietary system, it lacks extensibility.

We propose with ChronoVault a new recording method that saves video footage after the event has passed. It does share some functionality with commercial bodycams, but is specifically designed to support multiple recording triggers and only saves videos after the user explicitly wishes to do so. This, we argue, limits the risk of collecting privacy-intrusive data to an absolute minimum, and avoids having to browse larger video streams afterwards.

3 Framework design of ChronoVault

The recording process is triggered exclusively by a button press from the participant, which serves as an implicit form of consent and ensures that only relevant, intentionally approved clips are retained. This method results in highly focused recordings that preserve only the essential content required for analysis. Unlike continuous recording systems, which burden researchers with the need to view, interpret, and annotate lengthy footage, ChronoVault simplifies the process by producing a collection of short, meaningful clips. This reduces cognitive fatigue and streamlines the data review process. Additionally, by building the system on the Raspberry Pi Zero 2, we enable expansion through its available GPIO pins. These can be used by other researchers to integrate alternative input mechanisms for triggering recordings or to connect auxiliary sensors for enriched data capture.

Basic System Structure. The core of our video capture device is the Raspberry Pi Zero 2 W (RP Zero), which features an ARM processor and a GPU capable of encoding 1080p video at 30 frames per second, which is sufficient for our requirements. However, the RP Zero lacks a real-time clock (RTC) module, resulting in the loss

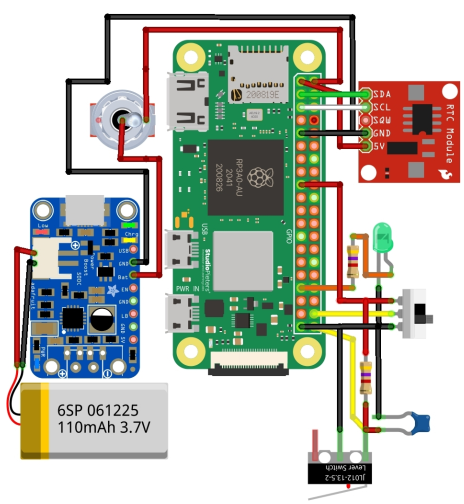

Figure 2: ChronoVault's hardware components are largely commercially available: A Raspberry Pi Zero 2, the Pi HQ Camera with a 185 degree lens, a realtime clock module, switches, LED, and a Li-ion charging board.

of accurate timekeeping between device restarts and making it impossible to timestamp recordings reliably. To resolve this, we integrated an external RTC module into the system.

To transfer recorded video files to a computer, we avoided cumbersome methods such as disassembling the device or connecting it to a monitor. Instead, we opted for a wireless solution via SSH. A Wi-Fi hotspot is created on the RP Zero, which is assigned the static IP address 25.25.25.25. Another PC can then connect to this network and initiate an SSH session. To conserve energy, the hotspot is not active by default. A hardware switch determines the startup mode: either the capture service is launched, or the device starts the hotspot for data transfer. The camera is connected to the RP Zero via the CSI interface. We selected the official Raspberry Pi High Quality Camera, which supports interchangeable lenses—including a 185° wide-angle lens, though other compatible camera modules and lenses can also be used. We added an LED indicator to display the system's state; The LED stays lit while the system is actively recording and turns off automatically in case of an error. The design leaves 23 available GPIO pins (excluding power and ground), which offers flexibility for additional features or custom input mechanisms. For power, we chose a Li-Po battery. To maximize energy efficiency, the battery's voltage is supplied directly to the RP Zero without boosting it to 5V, which the RP Zero can tolerate. Charging is handled by an off-the-shelf charging board, and a physical switch is included for toggling power to the system. The schematic for the system can be found in Figure 2.

Software. Our software is built on top of the existing Picamera2 Python package, which provides core functionality for previewing, controlling the camera, saving footage, and performing basic image manipulations. Leveraging this foundation, we developed a lightweight, modular package designed for easy integration, configuration, and execution, requiring only a few lines of code to deploy. We structured our implementation around three main components:

Settings. Configuration is handled via a YAML file, which is loaded automatically during package initialization. The YAML file defines key parameters such as output folder paths, camera settings, and a user-specific identifier. This approach facilitates centralized configuration and simplifies the management of multiple devices in parallel.

Record. Video capture is managed through the Picamera2 library, which is pre-installed in the current Raspberry Pi OS. To enable retroactive recording, we utilize the CircularOutput feature, which maintains a buffer in RAM where incoming frames continuously overwrite the oldest ones. A blocking queue is used to listen for save requests from the Save-Event module. Upon receiving a trigger, the system begins encoding the buffered frames. Optionally, it can continue recording for a user-defined duration before stopping and writing the final H.264-encoded video to disk. The resulting file is stored in the configured capture folder.

Save-Event. This module is responsible for triggering the recording process by signaling the Record module to save the frames inside the CircularOutput. It pushes arbitrary data into the communication queue and can optionally execute a user-predefined callback function. Save events can be triggered via software or hardware mechanisms. For software triggers, we use a blocking queue with no timeout. For hardware triggers, we use a GPIO pin monitored by the rpi-lgpio library (developed by Dave Jones)[10]. Specifically, the system waits for rising or falling edges using the GPIO.wait_for_edge function. The architecture supports multiple instances of the Save-Event module running concurrently, allowing a combination of hardware and software triggers as needed.

In summary, this software package enables researchers to reliably deploy and customize the video recording system with minimal effort, while still supporting extensive flexibility and customization.

Possible extensions. Due to the availability of unused GPIO pins and the unoccupied USB connector, other researchers can easily extend the system to capture additional data from external sources. When using the USB port, it is important to ensure that a regulated 5V supply is provided on the main power bus, otherwise certain USB peripherals may not function reliably.

The GPIO header offers versatile expansion options. Several pins support standard communication protocols such as SPI or UART, which enable the transmission of more complex or structured data between the RP Zero and peripheral devices. The remaining GPIO pins can be used to read or write simple digital signals, allowing for basic binary sensing or actuation.

In addition, the on-board Bluetooth adapter remains accessible. Using Python libraries such as Bleak [5], researchers can interface with a wide variety of Bluetooth-enabled sensors and devices, further expanding the potential for multimodal data collection.

Energy efficiency. To evaluate the energy impact of our software on the RP Zero, we conducted a series of experiments focused on power consumption. A constant voltage power supply was used, delivering 4V at a maximum of 1.5A. We chose 4V as our supply voltage to better reflect the li-po battery characteristics, which are commonly used in portable systems. While commercial power banks typically provide 5V, they do so using DC-DC converters that introduce additional power loss, something we aimed to avoid.

For current measurement, a 0.5Ω resistor was placed in series with the RP Zero. We measured the voltage drop across this resistor to calculate the current, using Ohm's Law.

To analyze the impact of our software, we ran it in a minimal configuration, continuously while directing the camera toward a dynamic video feed consisting of rapidly changing scenes. We also included a test scenario in which a software-based Save-Event was triggered every 20 minutes to simulate periodic clip saving under real-world usage.

For comparison, we evaluated three system configurations: OS Lite (terminal-only interface, no camera or additional software running), Full Raspberry Pi OS Desktop, and our full system (including camera, software, and active Save-Event module)

The resulting measurements are summarized in Table 1.

When comparing our power consumption results to those reported by Jean-Luc Aufranc [4], we observe similar idle power usage under baseline conditions. During full operation—while running our software—the system exhibits a power consumption of

	Os Lite	Os Desktop	Software running	Software recording
Total Wh	14.57	15.80	34.38	34.20
W	0.60	0.65	1.43	1.42

Table 1: Energy comparison of different software configurations, using a power supply at 4V.

approximately 1.43 watts. This level of efficiency supports the feasibility of portable operation using relatively inexpensive battery solutions, making the system practical for mobile or long-duration deployments.

To validate our design choices and assumptions regarding voltage levels, we conducted an additional experiment in which we gradually reduced the supply voltage from 5.0V to 3.3V in 0.05V steps. For each step, we measured the voltage drop across the 0.5 Ω series resistor over a 10-minute period, allowing us to monitor system behavior and power draw under various conditions.

The experiment showed that operating at lower voltages reduces power consumption. Instead of drawing 1.63W at around 5V, the system only uses 1.50W at 3.8V, a savings of 0.13W. This does not yet account for the additional power loss typically introduced by a 5V step-up converter. Notably, when supply voltage drops below a threshold, the RP Zero's under-voltage detection circuit reduces CPU clock speed, further lowering power usage. This threshold isn't clearly defined and lies form 3.8V down to 3.4V. In our battery-powered recording test, this down clocking had no noticeable effect on video quality. However, it brought a significant benefit: reduced power consumption of just 0.78W. With these benefits we like to suggest that limiting the RP Zeros clock speed, when longer run times are needed.

In our battery tests, we used a 1500 mAh LiPo battery. Under normal operation, the system was able to record for approximately 3.5 hours. With the reduce clock speed we are able to get an additional 0.5 to 1 hour. Overall we get around 4 hours of recording time. If we set the RP Zero in the reduced power mode from the beginning, we estimate that the recording time could be doubled, offering up to 7–8 hours of continuous use on a single charge.

Enclosure Design. To make the system portable and suitable for field use, we designed and assembled a custom lanyard-mountable case. This form factor provides us with a stable and elevated viewpoint, ensuring a clear overview of the surrounding environment.

To maintain ease of use, the mode-selection and power switches were placed on the side of the case for accessibility. For triggering the video save event, we designated the lower front section as the button area. A micro switch is positioned underneath this area, allowing participants to save the buffered video to flash memory with a single press. To address mechanical issues caused by the weight of the front button (e.g., accidental triggering or sluggish rebound), we integrated an additional micro switch on the opposite side to provide increased counterforce and improve tactile feedback.

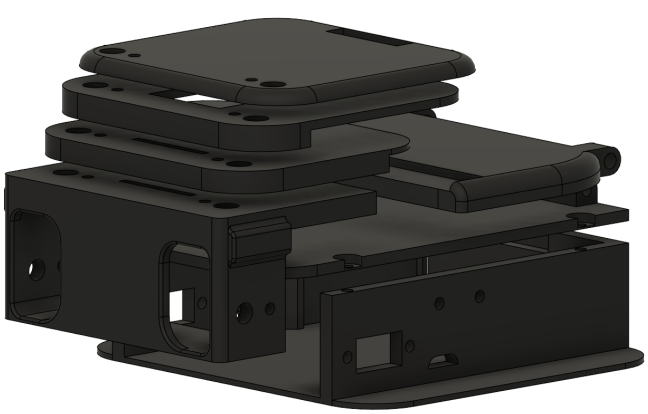

Figure 3: Our 3D printable designed case

In a later iteration, the micro switches were repositioned to the exterior of the case to maximize internal space for components. The final enclosure design can be seen in Figure 3. For a fully assembled device, reference Figure 1.

4 User study

To evaluate the functionality of our platform and demonstrate its advantages during both data capture and post-analysis, we designed a user study. In this experiment, we compare the footage captured by ChronoVault with that of an external reference camera, while also collecting qualitative feedback from participants regarding their experience using the system.

4.1 Design

In the user study, participants are instructed to assemble a small flower model (upper left in figure 4) using building blocks. The model, designed for children aged five and above, consists of 57 individual pieces and is constructed following a 16-step illustrated manual [1].

Each participant receives a bag containing the building blocks along with a printed assembly guide. Simultaneously, they wear the ChronoVault device throughout the task. To obtain a secondary perspective, an external camera is positioned to record the entire scene from a distance. Participants are seated at a table in a small corner of a room, as illustrated in Figure 4.

Figure 4: Participant working on the task while wearing ChronoVault, referencing the completed model in the upper left corner.

The building process is divided into two subtasks: steps 1–6 and steps 7–16, with a coffee break separating the two phases. Participants are instructed to press the button on the ChronoVault device whenever they feel they have completed the subtask. This allows us to evaluate the effectiveness of intent-driven capture and the participants' ability to identify meaningful recording moments.

4.2 Participant recruitment

We successfully recruited a total of 10 participants from computer science department, each with diverse academic backgrounds and research interests. Notably, 50% of the participants originated from a single country, while the remaining individuals came from various other nations and continents, offering an international perspective.

It is important to acknowledge that the participant pool consisted exclusively of young male individuals. While this limits the demographic diversity of our study, we believe that the simplicity and accessibility of both the task and the ChronoVault system reduce the likelihood that factors such as gender or age would significantly influence the outcome.

In future ongoing studies we aim to get a more diverse participant groups to validate generalizability.

4.3 Evaluation

By analyzing the externally recorded footage, we were able to visualize the participants' workflow. As shown in Figure 5, the total task completion time ranged from 10 minutes to 24.5 minutes, with an average duration of 16.65 minutes. The length of the building pause varied between 3.36 and 10.66 minutes.

During our analysis, we made several key observations:

- Participants demonstrated varying levels of familiarity with building blocks. Some completed the first six steps quickly, while others required more time, suggesting little to no prior experience.
- All participants followed the provided instructions to the best of their ability, although some occasionally overlooked smaller details.
- Nearly all participants successfully completed the model. In one case, four pieces remained unconnected, but the overall structure was assembled correctly.
- All participants successfully pressed the ChronoVault button.

When analyzing the data recorded by ChronoVault, we found it significantly easier to identify the endpoints of the building task. This efficiency is due to our design decision to only record the most recent 20 seconds of activity, triggered by user input. However, we occasionally observed missing footage at the beginning of some clips. This issue was largely mitigated by adding additional filtering capacitors to the power supply and the 3.3V bus, which helped stabilize the system during recording.

As a result, we no longer needed to review the full 166 minutes of continuous footage to locate the relevant 5-minute overlapping segments, representing a 96.99% reduction in the total video material requiring analysis. This substantial decrease in review time highlights ChronoVault's potential to improve researcher efficiency and reduce the cognitive burden associated with video analysis, particularly in observational or behavioral studies. While this approach offers a clear advantage for researchers, it also shifts the cognitive burden onto the participants. They must now actively assess whether their current situation is included in the situations to record. This may work well if the tasks are simple and straightforward. However, in more complex scenarios, where participants are engaged in demanding tasks or deep thought processing, this reliance could lead to missed or incomplete recordings, as participants may forget or miss to trigger the recording.

Thanks to the ChronoVault capture device, we receive a clear, intentional signal from the participant indicating when they consider the task to be complete. This allows us to accurately timestamp the participant's perceived endpoint of the activity. In contrast, relying solely on footage from a secondary camera would require manual interpretation and guessing, which can vary significantly between observers and introduce subjective bias and uncertainty. Moreover, ChronoVault eliminates the need to manually search or scrub through extended video recordings to locate moments of interest, such as task completion. This targeted recording approach streamlines the analysis workflow, enabling researchers to focus on contextual interpretation rather than time-consuming footage navigation.

4.4 Visual inspection of captured video

By reviewing the video footage captured by ChronoVault, we are able to reliably assess the current progress of the model assembly, as illustrated in subfigure A of Figure 6. However, depending on how participants move during the task, the ChronoVault camera may not always remain perfectly centered, as shown in subfigure B. Despite this, the use of a 185-degree wide-angle lens ensures that the model remains visible even in cases of suboptimal camera alignment. This setup proves effective even with our current recording settings of 720p at 30 fps, as demonstrated in the 2× zoomed-in example in subfigure C. The footage is generally sufficient for identifying key actions and understanding task progression. That said, the quality and framing of the captured footage can vary depending on participant height. For some participants, the ChronoVault device may hang closer to table level, reducing the top-down visibility of the task area (see subfigure D). This limitation could be addressed by adjusting the lanyard length to better match participant height, or by integrating the recording hardware into alternative wearable designs, especially when deploying ChronoVault for long term field studies.

4.5 Survey

After completing the experiment, participants were asked to fill out a survey consisting of 33 questions. These covered a range of topics, including personal background, technology experience, the System Usability Scale (SUS) [6], Trust in Automation Scale (TIAS) [12], our own custom questions related to the use and experience of ChronoVault in manual tasks based on a Likert-scale system, and whether they would be willing to use the device again.

For the SUS, participants achieved an quiet good average score of 77.5, with a Cronbach's Alpha of 0.791, indicating good internal consistency. In the TIAS section, the average score was 4.9, with a Cronbach's Alpha of 0.782. This shows to us that ChronoVault is in general a usable, satisfactory and trustable system.

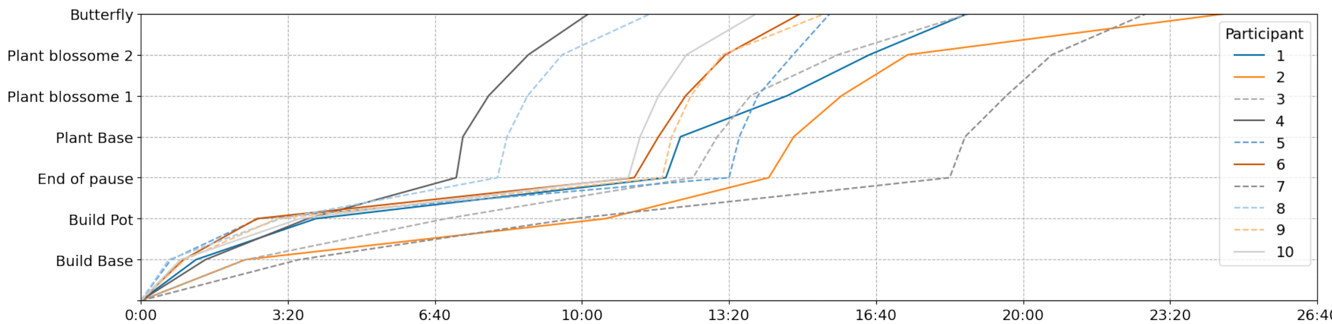

Figure 5: Participants' subtask duration and step progression shows a large variability in the time taken to assemble the model.

From our custom questions, most participants reported that recording a video was easy (3 fully support this giving a 1 on a Likert-scale of 7); however, 1 participant didn't agree (6). This highlights a possible need to further improve the clarity of the instructions we provide. Regarding the building task, nearly everybody (83%) said the device did not interfere with their ability to assemble the model.

Participants expressed mixed feelings, 2 slightly against it (2,3), 2 (4) unsure and two fully in favor of it (7), about who should control the recording process and whether automatic recording via body cues would be desirable. Five participants also reported uncertainty (scoring 5 to 7) about whether the system had successfully registered their input. This suggests a hardware revision may be necessary—potentially adding feedback mechanisms like a vibration motor or similar indicator. Lastly most participants strongly indicated (1) that they feel in control of the device. The rest of the participants had mixed feeling about it.

Overall, participants indicated to us that they found the system to be easy to use, trustworthy and reliable. They also indicated that they would use ChronoVault again.

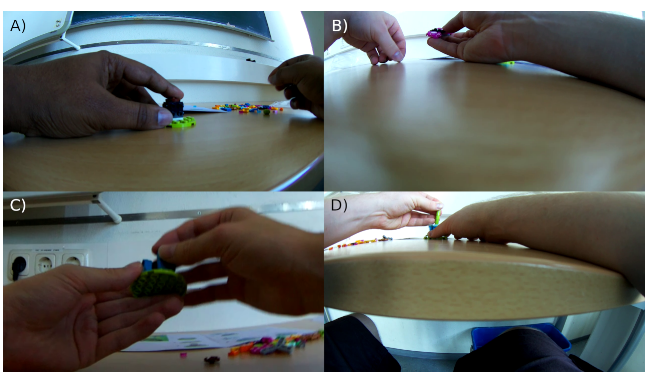

Figure 6: Some exemplary frames from the captured ChronoVault videos from our assembly task.

5 Conclusion

We have demonstrated that our system can reliably record video using the Raspberry Pi camera module. By capturing only short, event-triggered segments rather than recording continuously for extended periods, we significantly reduce the volume of data requiring analysis. This approach not only enhances efficiency but also aligns more closely with GDPR principles by limiting the collection of excessive or potentially sensitive footage. The system is designed to be both adaptable and extensible. Researchers can easily replicate and customize the setup to suit their specific use cases. Leveraging Python and a widely supported GPIO library simplifies the integration of additional sensors or input mechanisms. For more complex configurations, the Raspberry Pi Zero 2 W offers multiple connectivity options, including Wi-Fi, Bluetooth, and USB, enabling seamless communication with external devices.

We have also shown that the RP Zero can be powered directly from a battery, removing the need for a 5V step-up converter. This not only avoids the energy losses typically associated with voltage conversion but also results in a slight overall reduction in power consumption.

The video capture mechanism can be triggered through various means, including autonomous sensor-based detection or manual activation. This flexibility allows participants to act as either passive contributors or active decision-makers in the recording process. However, delegating control to participants may increase their cognitive load, particularly in complex or demanding situations.

To evaluate the capabilities of ChronoVault, we conducted a small user study. Compared to footage from a secondary camera, ChronoVault enabled a 96% reduction in video content requiring review. Participant feedback was generally positive, indicating that the device rarely interfered with task performance. Nevertheless, the feedback also revealed areas for improvement in both the hardware and usability.

In order to support reproducing our ChronoVault prototype, the Python code, 3D enclosure designs, schematics, and the bill of materials can be publicly downloaded from:
https://michael-brilka.github.io/ChronoVault

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 262513311 – SFB 1187.

References

[1] 2021. Lego Friends - 30417. https://www.lego.com/cdn/product-assets/product.bi.core.pdf/6401056.pdf.

[2] 2023. Motorola V500 Bodycam - Comfective. https://comfective.nl/product/motorola-v500-bodycam/.

[3] 2023. V500 Body Camera - Motorola Solutions EMEA. https://www.motorolasolutions.com/en_xu/video-security-access-control/body-worn-cameras/v500.html.

[4] Jean-Luc Aufranc. 2021. A deep dive into Raspberry Pi Zero 2 W's power consumption. https://www.cnx-software.com/2021/12/09/raspberry-pi-zero-2-w-power-consumption/.

[5] Henrik Blidh. 2017. Bleak. https://github.com/hbldh/bleak.

[6] John Brooke. 1995. SUS: A quick and dirty usability scale. *Usability Eval. Ind.* 189 (11 1995).

[7] José Miguel Matos Lopes da Costa. 2019. Dashcam Video Streaming and Storage. (2019).

[8] Erin Griffiths, Salah Assana, and Kamin Whitehouse. 2018. Privacy-preserving Image Processing with Binocular Thermal Cameras. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 4, Article 133 (Jan. 2018), 25 pages. https://doi.org/10.1145/3161198

[9] J. Healey and R.W. Picard. 1998. StartleCam: a cybernetic wearable camera. In *Digest of Papers. Second International Symposium on Wearable Computers (Cat. No.98EX215)*. 42–49. https://doi.org/10.1109/ISWC.1998.729528

[10] Dave Jones. 2019. rpi-lgpio. (2019).

[11] Mina Khan, Glenn Fernandes, Akash Vaish, Mayank Manuja, and Pattie Maes. 2021. Wearable System for Personalized and Privacy-preserving Egocentric Visual Context Detection using On-device Deep Learning. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization* (Utrecht, Netherlands) *(UMAP '21)*. Association for Computing Machinery, New York, NY, USA, 35–40. https://doi.org/10.1145/3450614.3461684

[12] Melanie J. McGrath, Oliver Lack, James Tisch, and Andreas Duenser. 2025. Measuring trust in artificial intelligence: validation of an established scale and its short form. *Frontiers in Artificial Intelligence* Volume 8 - 2025 (2025). https://doi.org/10.3389/frai.2025.1582880

[13] Octav Opaschi and Radu-Daniel Vatavu. 2020. Uncovering Practical Security and Privacy Threats for Connected Glasses with Embedded Video Cameras. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 4, Article 167 (Dec. 2020), 26 pages. https://doi.org/10.1145/3432700

[14] Giovanni Schiboni, Fabio Wasner, and Oliver Amft. 2018. A Privacy-Preserving Wearable Camera Setup for Dietary Event Spotting in Free-Living. In *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 872–877. https://doi.org/10.1109/PERCOMW.2018.8480222

[15] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2018. Privaceye: Privacy-preserving first-person vision using image features and eye movement analysis. *arXiv preprint arXiv:1801.04457* (2018).

[16] Gemma Wilson, Derek Jones, Patricia Schofield, and Denis J Martin. 2016. Experiences of using a wearable camera to record activity, participation and health-related behaviours: Qualitative reflections of using the Sensecam. *Digital Health* 2 (2016), 2055207616682628.

[17] Katrin Wolf, Albrecht Schmidt, Agon Bexheti, and Marc Langheinrich. 2014. Lifelogging: You're Wearing a Camera? *IEEE Pervasive Computing* 13, 3 (2014), 8–12. https://doi.org/10.1109/MPRV.2014.53

[18] Muchen Wu, Parth H. Pathak, and Prasant Mohapatra. 2015. Enabling privacy-preserving first-person cameras using low-power sensors. In *2015 12th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. 444–452. https://doi.org/10.1109/SAHCN.2015.7338345

[19] Ayoung Yoon, Kyle ML Jones, and Lydia Spotts. 2017. Ethical use of lifelogging data for research: Perceived value and privacy concerns of wearable camera users. *Nordic Journal of Information Science and Culture Communication* 6, 1 (2017).