

CL-PPIIMS: A Cognitive Load-Based Privacy-Preserving Intelligent Interruption Management System

Moritz Maleck^[0009-0006-5551-8386], Tom Gross^[0000-0001-8353-7388]

Human-Computer Interaction Group, University of Bamberg, 96045 Bamberg, Germany
hci@uni-bamberg.de

Abstract. Managing interruptions effectively is a key challenge in modern work environments, as poorly timed disruptions during periods of high cognitive load can severely reduce team productivity and individual focus. While Intelligent Interruption Management Systems aim to solve this by assessing user states, they often face a critical trade-off between the precision of physiological sensing and the imperative of user privacy. We introduce CL-PPIIMS, a system built on a privacy-preserving architecture that resolves this conflict by applying Soft Computing principles—specifically, classifying imprecise cognitive states from noisy physiological data—to enhance Human-Computer Interaction. It leverages a stationary eye tracker for high-fidelity cognitive load data but performs all sensitive computations strictly on the user’s local machine. Only a highly abstracted, non-identifiable classification of the user’s state is then transmitted via Bluetooth Low Energy (BLE) to a mobile client, enabling an intelligent, privacy-aware management of team-based communication requests that enhances productivity without compromising user data sovereignty.

Keywords: Effortless Coordination; Team Meetings; Arrangements; Interruptions; Interruption Management Systems; Cognitive Load; Eye Tracking; Privacy by Design; Bluetooth Low Energy (BLE).

1 Introduction

Interruptions—such as private notifications from messengers, or in-persona interruptions during work—impede primary tasks, leading to longer completion times and a reduction in both accuracy and efficiency of work [15; 30; 38]. While the disruptive nature of interruptions is well-documented, it can also yield positive effects, such as fostering creativity or enhancing social connections [15; 37]. The process of managing an interruption involves several distinct stages, from the initial notification to engagement with the secondary task and the subsequent effort to resume the original work [34]. The degree of disruption is influenced by factors such as the complexity of the interruption and its similarity to the ongoing task [11]. This effect is particularly pronounced when an interruption exhibits high complexity and bears little resemblance to the primary activity.

A consistent finding in related research is that interruptions are more harmful when they happen during periods of high cognitive load [3; 30]. Cognitive load, which refers

to the mental effort engaged in working memory to complete a task, has consequently become a key predictor of a user's receptiveness to interruption [3]. In response, a variety of Interruption Management Systems (IMS) have been created to optimise the timing of notifications and minimise their disruptive effects. Some of those systems infer cognitive load from consumer wearables to signal a user's availability [30], while alternative approaches use workload-aligned task models to identify opportune moments for interruption [17]. These methods, however, often present a dilemma: systems relying on simple, indirect measurements may not be precise enough, whereas those using direct physiological sensors introduce substantial privacy risks if they necessitate transmitting sensitive data to a central server for analysis.

Eye tracking, specifically the measurement of pupil diameter, has been established as an effective and reliable technique for the real-time assessment of cognitive load [10; 39]. The pupil's tendency to dilate in response to heightened mental effort, a phenomenon called the Task-Evoked Pupillary Response (TEPR), offers a direct insight into a user's cognitive state [4]. While this technique provides high precision, it also produces a continuous stream of sensitive biometric data, making its secure handling a primary design consideration.

A common software architectural pattern for analysing complex physiological data relies on a server-centric model, where the user has no direct control over their raw data. A key reason for this approach is the high computational demand of many machine-learning algorithms, which often requires outsourcing the processing logic to powerful central servers. For example, the work of Afzal *et al.* [1] describes a system where EEG data is transmitted to a server for classification using a deep neural network.

This pattern is also prevalent in other areas of physiological monitoring. Systems presented by Jeong *et al.* [19] and Chang *et al.* [6] demonstrate a typical workflow where data from physiological sensors are forwarded to a mobile device, which then acts as a gateway to a cloud server for detailed analysis.

With respect to such sensitive biometric data, correct handling requires the implementation of the principle of data minimisation, as outlined in frameworks like 'Privacy by Design' by Cavoukian [5]—ensuring that raw data remains entirely under the user's local control and only highly abstracted, non-identifiable information is ever transmitted, thereby significantly reducing the risk of misuse. An approach that preserves user privacy and data control by performing all calculations on a local device and only transmitting highly abstracted data represents a novel contribution to the field.

This paper introduces CL-PPIIMS (pronounced 'CL-Pims'), a novel concept and system engineered to resolve this precision-privacy trade-off. Our work leverages a privacy-preserving architecture to make the following key contributions:

- A novel, distributed Intelligent Interruption Management System (IIMS)—the CL-PPIIMS—designed to optimise the timing of interruptions in the desk work context. The system features a cognitive load calculation station that processes eye-tracking data strictly on the user's local machine. In line with the principle of data minimisation, it transmits only a highly abstracted classification of the user's state via Bluetooth Low Energy (BLE) to their mobile iOS device. The mobile application, in turn,

forwards this classification to a central server, which manages the timing of interruptions in means of communication requests within a work team.

- A hybrid architecture for managing interruptions that combines the high precision of stationary, research-grade eye tracking with a robust ‘privacy by design’ processing model. All sensitive data analysis is performed strictly on a user’s local machine, ensuring raw biometric data never leaves their control.
- A flexible and scalable system that can be easily integrated with existing workplace tools to facilitate less disruptive communication within the work domain.

To the best of our knowledge, we are the first to introduce an IIMS that uses cognitive-load-based eye tracking in a privacy-preserving architecture designed for mobile integration. The structure of this article is as follows: first, we provide an overview of related work. Then, we present our CL-PPIIMS concept, followed by a description of its implementation. Finally, we conclude with an outlook for future work.

2 Related Work

The background for our work includes research on interruptions in the workplace, cognitive load theory and its measurement via pupilometer, and architectures for physiological computing.

2.1 Interruptions and Interruption Management

Interruptions disrupt primary tasks, increasing task completion time and affecting accuracy and efficiency [15]. This is especially true in team settings, where collaborative processes can be delayed, affecting the performance of the entire team [13; 15]. The degree of disruption is further influenced by factors such as the complexity of the interruption and its similarity to the ongoing task [11]. Yet, interruptions also yield positive effects, such as fostering creativity or strengthening social connections [15]. The key challenge for IMS is therefore to balance the right moment of delivering interruptions to minimise negative effects and maximise positive effects.

To signal availability in co-located work environments, individuals often rely on social cues, such as maintaining an open or closed office door [30]. These physical cues are not applicable in remote or hybrid settings, which have become more prevalent. The trend towards working from home has introduced new challenges, including an increase in non-work interruptions and a shift in collaboration patterns [23; 35]. In digital tools setting a personal availability status such as ‘Do Not Disturb’ (DND) offers some control but is limited; users may forget to activate them or may configure them too permissively, fearing they will miss critical information [14; 18]. Recent research highlights that users adopt distinct strategies for managing digital notifications, categorisable as ‘proactive’ or ‘reactive’ [27]. Proactive users actively manage settings but can still feel overwhelmed, while reactive users often apply no strategies (and may be more susceptible to constant disruption) [27]. Moreover, users often prefer to suppress an alert (i.e., silent delivery) rather than defer it, e.g., because of fear of missing out [24].

To overcome these limitations, more advanced IMS aim to automate the process of finding opportune moments. Some systems use physiological data from consumer wearables (e.g., heart rate variability, skin conductance) to estimate cognitive load and signal a user's availability via a physical light display [12; 30]. Others analyse task structure to deliver interruptions at natural breakpoints, such as the boundaries between subtasks, where cognitive load is demonstrably lower [3]. Other approaches infer interruptibility from different modalities; e.g., through the usage of speech recognition to detect lexical affirmation cues (like 'got it'), naturally signalling task completion in collaborative settings [31]. Other related interruption management systems mediate communication requests by assessing cognitive load leveraging on the pupil diameter to find moments of low cognitive load for border-crossing between life domains [26], arranging meetings within teams [25], or timing of interruptions in a single user scenario with an email-answering task [21].

2.2 Cognitive Load and Pupilometer

Cognitive Load Theory (CLT) posits that working memory has a limited capacity, and that effective instruction or interface design should minimise extraneous load to free up resources for learning and problem-solving [32]. Eye tracking has proven to be a highly effective, non-invasive method for measuring cognitive load in real time [10; 39]. The pupil diameter is a particularly reliable indicator because it cannot be consciously controlled and is tightly coupled with cognitive processes [4; 7; 10]. [40].

This coupling is rooted in the autonomic nervous system (ANS). Increased mental effort activates the sympathetic nervous system (SNS), a process modulated by the Locus Coeruleus-Norepinephrine (LC-NE) system, which in turn causes the pupil to dilate [33; 36]. This neurophysiological link is fundamental, as the LC-NE system actively optimises performance by regulating 'neural gain'—a process that amplifies task-relevant information while suppressing distractors [16]. The Task-Evoked Pupillary Response (TEPR) is a robust measure of cognitive load [4]. However, the pupil also constricts in response to light—known as the Pupillary Light Reflex (PLR). The PLR is the primary challenge for cognitive pupilometer, as its effect on pupil size is orders of magnitude larger than the subtle changes induced by cognitive load, masking the signal of interest [22; 28]. Pupil diameter vary up to 7.5 millimetres due to light changes, and only up to 0.5 millimetres due to cognitive load changes [10].

To address this, advanced frequency-domain metrics have been developed. The Index of Pupillary Activity (IPA) uses a wavelet transformation to count high-frequency pupil oscillations, which are assumed to reflect cognitive activity, while filtering out low-frequency changes associated with ambient light [9]. A more recent and robust evolution is the Low/High Index of Pupillary Activity (LHIPA) [8]. Instead of just counting high-frequency events, LHIPA computes the ratio of energy in the low-frequency bands to the high-frequency bands [8; 29]. This ratio-based approach is more closely tied to the underlying push-pull dynamic of the ANS. LHIPA has been shown to be more effective than IPA at detecting cognitive load in several contexts, particularly in less restrictive tasks [8].

2.3 Privacy-Preserving Physiological Computing

The use of sensitive physiological data necessitates architectures that prioritise user privacy. The common approach in mobile health and other fields is to send sensor data to a powerful cloud server for analysis [1; 6; 19]. While this offers high computational power, it creates a central point of failure and a high-value target for data breaches, forcing the user to trust a third-party provider with their most sensitive data.

An alternative paradigm is Edge Computing, which advocates for moving computation closer to the source of data generation [2]. This reduces latency and bandwidth usage while significantly improving privacy and security by keeping data local. Such an approach is a direct implementation of the ‘Privacy by Design’ framework, which mandates that privacy be built into a system’s architecture from the outset, rather than being added as an afterthought [5]. Conceptually, this is related to Federated Learning, where machine learning models are trained on local user data without the raw data ever being shared with a central server [20].

In the next section we introduce our concept for providing real-time cognitive-load data in a privacy-preserving way by sending abstracted classifications via Bluetooth Low Energy to mobile phone devices, which are then using these in the context of an intelligent interruption management system—our CL-PPIIMS.

3 A Concept for a Privacy-Preserving Mobile Cognitive Load-Based Intelligent Interruption Management System

We introduce CL-PPIIMS, an approach for supporting less disruptive interruptions by means of meeting requests within a work team, determining opportune moments based on the user’s cognitive load. The system is designed upon a multi-tiered architecture separating concerns to maximise both measurement precision and user privacy. This architecture acts as a ‘privacy firewall’, abstracting sensitive information at each stage to ensure that only the necessary, non-identifiable data is transmitted across systems.

The CL-PPIIMS concept is realised through a four-tiered architecture managing the flow of information from raw physiological signals to actionable interruption management decisions. It contains local, high-fidelity sensing and processing to ensure data quality and privacy; efficient, abstracted state transmission to mobile devices; and a flexible application logic handled by clients and a backend server.

Tier 1: High-Fidelity Local Sensing With Eye Tracking. The first tier contains a stationary eye tracker, equipped with at a local user’s workstation. A stationary tracker is inherently less invasive in a professional context, as it does not need to be worn, causes physical discomfort, or needs to be managed by the user throughout the day. Instead, it integrates seamlessly into the existing work environment, operating unobtrusively in the background for many hours without restriction. Still, this user-centric setup concurrently provides ideal conditions for high-fidelity data acquisition, allowing high precision for the data input.

Tier 2: Local Cognitive Load Assessment. All computationally intensive and privacy-sensitive analysis is performed in this tier, entirely on a separate local PC in the pure control of the user, serving as *Cognitive Load Calculation Station*. This can be a low-budget, affordable Windows PC (e.g., the Gazepoint GP3 eye tracker only requires a Windows 10/11 PC with at least Intel Core i5 2022 or newer, and 8 GB RAM at the time of writing this article). A service, the *Cognitive Load Calculation Service*, ingests the raw pupil diameter stream from any connected eye tracker (support for various manufacturers—Tobii and Gazepoint—is implemented). It ingests the raw pupil diameter stream and executes a pipeline including a feature extraction method inspired by the Low/High Index of Pupillary Activity (LHIPA). To meet the demands of a real-time system by reducing computationally expensive calculations, our approach uses a modified version that focuses on the energy ratio between the low and high-frequency components of the pupillary signal. This local processing model is the cornerstone of the system’s privacy guarantee. By confining the raw biometric data to a local, offline operating machine in control of the user, the risk of unauthorised disclosure or misuse by third parties is minimised. The user retains absolute data sovereignty over their most sensitive information.

Tier 3: Abstracted State Transmission via BLE. Once the cognitive load has been processed, the system performs a critical data abstraction and minimisation step by classifying the detailed cognitive load and responding raw data into a discrete state (i.e., ‘Low’, ‘Medium’, ‘High’). Only this final, non-biometric classification is provided to the *Cognitive Load Classification BLE Sender* to act as Bluetooth Low Energy (BLE) peripheral. BLE is a low-power, secure, and ubiquitous protocol ideal for broadcasting small amounts of data over short distances. This design elegantly circumvents the problem of transmitting high-bandwidth eye-tracking data wirelessly.

Tier 4: Mobile Application Logic. The *Intelligent Interruption Management System Client Application* on the user’s mobile iOS device acts as a BLE central, continuously listening for the cognitive load state broadcast by the cognitive load calculation station. A key function of the client application is to enhance team awareness by translating the received data into a visible status for colleagues, which acts as a digital social cue. This status can reflect several states (cf. Figure 1): ‘Working / Not at desk’ if the client is not connected via BLE, ‘Working (Eye Tracker Data Unavailable)’ if the connection exists but no valid eye-tracking data is received, or a specific cognitive load level (‘Low’, ‘Medium’, ‘High’) when data is available. This awareness information allows colleagues to make more considerate decisions about interrupting other colleagues even when not initiating meeting requests through the provided system, discouraging them from interrupting a team member who is currently experiencing a high or medium cognitive load. In parallel to providing this visual cue, the iOS client relays the cognitive load classification to a central backend server—the *Intelligent Interruption Management System Server*—via a standard HTTPS request. The backend server is responsible for the final application logic. It manages a priority queue of pending meeting requests from colleagues, using the received cognitive load status to decide when to release them based on their priority, while ignoring users with invalid eye tracking data or not connected to the BLE sender (e.g., because not being at the work desk). For example, it delivers high-priority requests during a ‘Medium’ load state,

while holding all non-urgent requests until the user’s status changes to ‘Low’. This final tier provides flexibility and allows leveraging the cognitive load information without accessing the underlying physiological data.

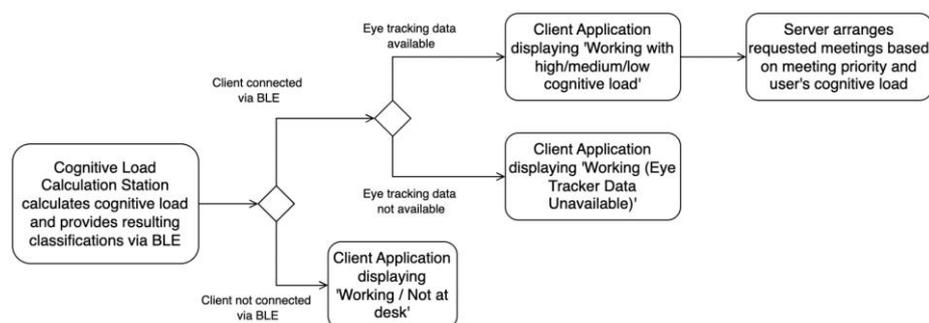


Fig. 1. Schematic representation of the client application’s state logic. The displayed user status depends on the BLE connection and the availability of eye-tracking data, serving as social awareness information (a social cue) for team members. Meeting requests are only processed for users with an available cognitive load classification.

This tiered approach ensures a clear separation of concerns, leveraging the best technology for each task: precision hardware for sensing, local compute for privacy-preserving analysis, BLE for efficient wireless communication, and standard mobile/backend technologies for application logic.

In the next section, we present the implementation of our CL-PPIIMS system, containing a cognitive load calculation station, a mobile iOS client application and a backend server application to manage interruptions between registered users.

4 CL-PPIIMS Implementation

In this section, we describe the implementation and functionality of the CL-PPIIMS system. The implementation consists of three main components: a Windows-based cognitive load calculation station, a native mobile client application for iOS, and a backend server for managing interruption requests. Each component was developed to adhere to the conceptual design outlined earlier, focusing on privacy-by-design, real-time cognitive load assessment, and seamless integration into a typical desk-based work environment.

4.1 System Overview

We implemented the CL-PPIIMS system using a distributed client-server architecture (cf. Figure 2). The core of the system is the Cognitive Load Calculation Station, a Windows 10 PC (Intel Core i5-8500T, 8 GB RAM) connected to a Gazepoint GP3 eye tracker (GP3V2). This station runs several services developed in Python (versions 3.10.11 and 3.13.5) for data acquisition and calculation, and a .NET (SDK version

9.0.302) application for BLE broadcasting. The Intelligent Interruption Management System Client Application is a native iOS 18.5 application running on an iPhone 16e. The final component, the Intelligent Interruption Management System Server, operates on a Node.js (version 23.5.0) server. This architecture facilitates a one-way flow of cognitive load data from the local station to the client, and bi-directional communication for managing meeting requests between the client and the server.

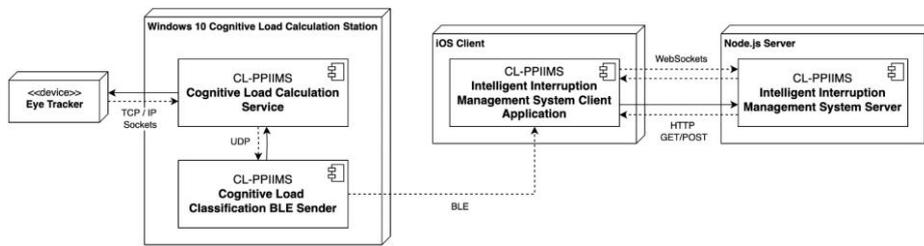


Fig. 2. System overview of CL-PPIIMS, showing the local calculation station, the iOS client, and the backend server.

4.2 Cognitive Load Calculation Station

The Windows 10 station is the privacy-centric core of the system. Its sole purpose is to calculate a user's cognitive load locally and broadcast an abstracted classification. Upon startup, a batch script initiated by the Windows Task Scheduler first launches the Gazepoint Control software (v7.0.0), followed by our eye tracker connection service, the calculation service, and finally the .NET BLE sender. For the first five minutes of operation, the system collects pupilometry data to establish a robust rolling baseline for its calculations. After this initialisation phase, it begins to calculate and broadcast cognitive load classifications in real time. Our experience shows that a single eye tracker is sufficient for a dual-monitor setup, as the calculation relies only on pupil diameter, which can be reliably measured even when the user is looking at a mobile device in front of their monitors.

4.3 iOS Client Application

The iOS client application is the primary user interface for the system. When first launching the app, the user is asked to grant permissions for notifications and for finding nearby Bluetooth devices. They then connect to a working space by providing their name, a planned quitting time, the name of their local BLE sending device, and the URL for the backend server. If the user starts the app after their specified quitting time has passed, a message informs them that they have been automatically logged out. The quitting time is used by the server to automatically remove users from the active pool, preventing meetings from being scheduled after the end of their workday.

Once connected, the app's main screen (cf. Figure 3 A) displays the user's own status, including their current cognitive load classification received via BLE. A key feature of the main screen is the 'Colleagues' list, which provides team awareness by showing the

real-time status of other active users. This functions as a digital social cue, allowing users to respect their colleague's current workload; for example, they can see if a colleague is experiencing medium or high cognitive load and decide against a physical interruption.

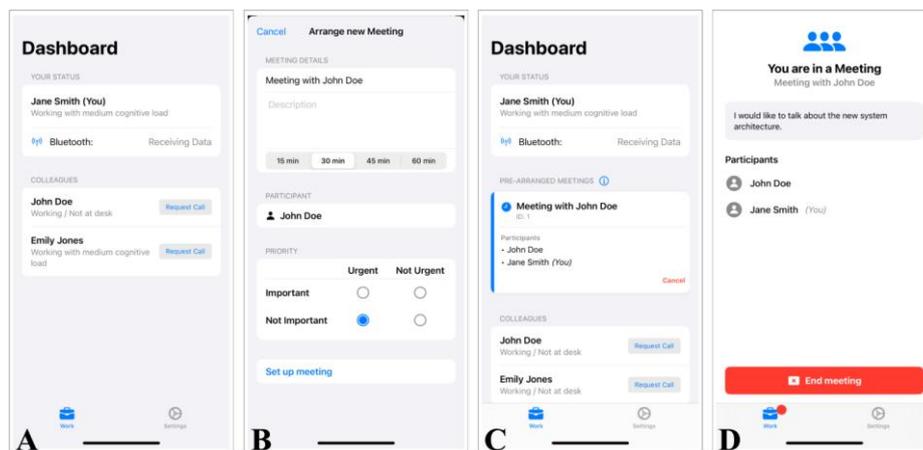


Fig. 3. Screenshots of the CL-PPIIMS Intelligent Interruption Management System Client Application implemented as an application for iOS. (A) Main screen of the application, showing the own cognitive load classification and other connected colleagues and their cognitive load classifications. (B) Screen to set up a meeting request, providing a title, an optional description and a priority. (C) Main screen with pre-arranged (i.e., pending) meeting requests listed. (D) Screen informing about initiated meeting by the backend server after opening the app via the sent out push notification.

From the main screen, users can initiate 1:1 meeting requests with their colleagues. The 'Arrange new Meeting' screen (cf. Figure 3 B) allows the user to specify a title, description, estimated duration, and a priority for the meeting. The priority is defined on a two-axis scale: 'Important/Not Important' and 'Urgent/Not Urgent'. Once a meeting request is set up, it appears in the 'Pre-arranged Meetings' section on the main screen, where the user has the possibility to cancel it again (cf. Figure 3 C).

When the server determines an opportune moment and initiates a meeting, the participants receive a push notification via the Apple Push Notification service. The app then transitions to an active meeting screen (cf. Figure 3 D), which shows the meeting details and participants. To ensure the system knows when users are available again, one of the participants has to manually end the meeting within the app.

4.4 Intelligent Interruption Management Server

The backend server is implemented in Node.js and acts as the central coordinator for all meeting requests. It maintains a real-time list of all active users and their current cognitive load classifications, as relayed by their respective iOS clients. The server's

primary logic is to process the queue of pending meeting requests based on both the request's priority and the cognitive load of all participants.

The decision-making logic is as follows:

- A meeting classified as 'Important and Urgent' can be initiated if all participants have a 'High', 'Medium', or 'Low' cognitive load.
- A meeting classified as 'Important and Not Urgent' or 'Not Important and Urgent' will only be initiated if all participants have a 'Medium' or 'Low' cognitive load.
- A meeting classified as 'Not Important and Not Urgent' requires all participants to be in a 'Low' cognitive load state.

Furthermore, a meeting can only be initiated if all participants have a status that indicates they are at their desk and available (i.e., their status is not 'Working / Not at desk' or 'Eye Tracker Data Unavailable'). Communication between the iOS clients and the server is handled via a combination of bi-directional WebSockets for real-time status updates and standard HTTP GET/POST requests for actions like setting up or cancelling meetings.

5 Summary and Outlook

We introduced CL-PPIIMS as a novel system designed to optimise interruptions within a work team by scheduling meeting requests based on real-time cognitive load assessments. By aligning the timing of interruptions with moments of lower cognitive load, this approach aims to minimise disruption during demanding tasks while enhancing team awareness and productivity. The system's design, including its privacy-preserving architecture with local data processing, a dedicated cognitive load calculation station, and a mobile client for providing social cues, demonstrates its potential to balance productivity with user privacy and data sovereignty. The implementation shows the feasibility of combining high-precision, stationary eye tracking with a mobile, flexible interruption management framework.

However, limitations such as the dependency on a stationary eye tracker and the potential for environmental factors like changing light conditions to influence the data suggest room for improvement. The system's reliance on a stationary eye tracker confines its use to a single workstation; future work may explore extending its functionality to mobile contexts, for instance by enriching the system with the integration of data from consumer wearables when the user is away from their desk. Future work could focus on addressing these limitations through a comprehensive, long-term user study to evaluate the system's impact on productivity and user satisfaction in real-world scenarios. Another key development could be the enhancement of the server-side decision logic to learn user preferences over time, and the potential integration of group meeting functionalities beyond the current 1:1 implementation. Moreover, the current logic, which requires all participants to be simultaneously below a cognitive load threshold, could lead to impractical waiting times and thus reduce user acceptance, necessitating more flexible options such as defining optional attendees. Furthermore, one could ex-

plore additional factors influencing interruptibility, such as integrating with user calendars to account for scheduled commitments. These enhancements could aim to increase the system's robustness, intelligence, and versatility, potentially improving its broader applicability across diverse professional contexts.

An important next step apart from new integrations is a comprehensive evaluation of the tool, preferably a real-world *in-situ* evaluation (i.e., not in a lab); such a study would provide the necessary results for a detailed discussion, critically analysing the system's real-world effectiveness, usability, and user experience to derive both practical implications for workplace settings and theoretical contributions to the field.

Ultimately, CL-PPIIMS represents more than just a tool for interruption management; it serves as a blueprint for a future generation of human-centred adaptive systems. The architectural principles of local processing, data minimisation, and user sovereignty are critical for building trust in systems that interact with our most sensitive physiological data. As technology becomes more integrated into our daily lives, the development of systems that can intelligently and respectfully adapt to a user's cognitive state will be paramount. By demonstrating a viable path that reconciles high-fidelity sensing with robust privacy protections, this work contributes to the broader goal of creating technology that not only enhances productivity but also supports user well-being in an increasingly demanding digital world.

Acknowledgments. We thank the members of the Cooperative Media Lab at the University of Bamberg. We also thank the anonymous reviewers for insightful comments.

References

1. Afzal, Muhammad Abrar, Gu, Zhenyu, Bukhari, Syed Umer and Afzal, Bilal. (2024). Brainwaves in the Cloud: Cognitive Workload Monitoring Using Deep Gated Neural Network and Industrial Internet of Things. *Applied Sciences* 14, 13.
2. Bablu, Tarek Aziz and Rashid, Mohammad Tanvir. (2025). Edge Computing and Its Impact on Real-Time Data Processing for IoT-Driven Applications. *Journal of Advanced Computing Systems* 5, 1. pp. 26-43.
3. Bailey, Brian P and Iqbal, Shamsi T. (2008). Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. *ACM Transactions on Computer-Human Interaction (TOCHI)* 14, 4. pp. 1-28.
4. Beatty, Jackson. (1982). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin* 91, 2. pp. 276-292.
5. Cavoukian, Ann. (2009). Privacy by Design: The 7 Foundational Principles. *Information and privacy commissioner of Ontario, Canada* 5, 2009. pp. 12.
6. Chang, Chao-Shu, Wu, Tin-Hao, Wu, Yu-Chi and Han, Chin-Chuan. (2023). Bluetooth-Based Healthcare Information and Medical Resource Management System. *Sensors* 23, 12. pp. 5389.
7. Dionisio, Daphne P., Granholm, Eric, Hillix, William A. and Perrine, William F. (2001). Differentiation of deception using pupillary responses as an index of cognitive processing. *Psychophysiology* 38, 2. pp. 205-211.
8. Duchowski, Andrew T, Krejtz, Krzysztof, Gehrer, Nina A, Bafna, Tanya and Bækgaard, Per. (2020). The Low/High Index of Pupillary Activity. In *Proceedings of the 2020 CHI*

- Conference on Human Factors in Computing Systems (CHI '20)* (Honolulu HI, USA). Association for Computing Machinery. pp. 1-12.
9. Duchowski, Andrew T, Krejtz, Krzysztof, Krejtz, Izabela, Biele, Cezary, Niedzielska, Anna, Kiefer, Peter, Raubal, Martin and Giannopoulos, Ioannis. (2018). The Index of Pupillary Activity: Measuring Cognitive Load vis-à-vis Task Difficulty with Pupil Oscillation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)* (Montreal QC, Canada). Association for Computing Machinery. pp. 1-13.
 10. Eckstein, Maria K., Guerra-Carrillo, Belén, Miller Singley, Alison T. and Bunge, Silvia A. (2017). Beyond eye gaze: What else can eyetracking reveal about cognition and cognitive development? *Developmental Cognitive Neuroscience* 25. pp. 69-91.
 11. Gillie, Tony and Broadbent, Donald. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological research* 50, 4. pp. 243-250.
 12. Goyal, Nitesh and Fussell, Susan R. (2017). Intelligent interruption management using electro dermal activity based physiological sensor for collaborative sensemaking. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 1, 3. pp. 1-21.
 13. Gross, Tom. (2013). Supporting Effortless Coordination: 25 Years of Awareness Research. *Computer Supported Cooperative Work (CSCW)* 22. pp. 425-474.
 14. Gross, Tom and Mueller, Anna-Lena. (2021). NotificationManager: Personal boundary management on mobile devices. In *Human-Computer Interaction–INTERACT 2021: 18th IFIP TC 13 International Conference, August 30–September 3, 2021, Proceedings, Part IV 18* (Bari, Italy). Springer. pp. 243-261.
 15. Gross, Tom and von Kalben, Michael. (2023). A Literature Review on Positive and Negative Effects of Interruptions and Implications for Design. In *Human-Computer Interaction – INTERACT 2023: 19th IFIP TC13 International Conference* (York, United Kingdom). Springer. pp. 373-379.
 16. Huang, Ringo and Clewett, David. (2024). *The Locus Coeruleus: Where Cognitive and Emotional Processing Meet the Eye*. In Papesh, Megan H. and Goldinger, Stephen D. , eds. *Modern Pupillometry*. Springer, Cham, Switzerland. pp. 3-75.
 17. Iqbal, Shamsi T and Bailey, Brian P. (2005). Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)* (Portland OR USA). pp. 1489-1492.
 18. Jarupreechachan, Weerachaya, Kitchat, Kotcharat and Surasak, Thattapon. (2023). ‘Do Not Disturb’: An Implication Design To Be Alerted But Less Stress. In *2023 International Seminar on Application for Technology of Information and Communication (iSemantic)* (Semarang, Indonesia). IEEE. pp. 176-181.
 19. Jeong, Jin-Woo, Lee, Woochan and Kim, Young-Joon. (2021). A Real-Time Wearable Physiological Monitoring System for Home-Based Healthcare Applications. *Sensors* 22, 1. pp. 104.
 20. Kairouz, Peter, McMahan, H Brendan, Avent, Brendan, Bellet, Aurélien, Bennis, Mehdi, Bhagoji, Arjun Nitin, Bonawitz, Kallista, Charles, Zachary, Cormode, Graham and Cummings, Rachel. (2021). Advances and Open Problems in Federated Learning. *Foundations and Trends® in Machine Learning* 14, 1–2. pp. 1-210.
 21. Katidioti, Ioanna, Borst, Jelmer P, Bierens de Haan, Douwe J, Pepping, Tamara, van Vugt, Marieke K and Taatgen, Niels A. (2016). Interrupted by your pupil: An interruption management system based on pupil dilation. *International Journal of Human–Computer Interaction* 32, 10. pp. 791-801.
 22. Klingner, Jeff, Kumar, Rakshit and Hanrahan, Pat. (2008). Measuring the task-evoked pupillary response with a remote eye tracker. In *Proceedings of the 2008 symposium on Eye*

- tracking research & applications* (Savannah, Georgia). Association for Computing Machinery. pp. 69-72.
23. Leroy, Sophie, Schmidt, Aaron M and Madjar, Nora. (2021). Working from home during COVID-19: A study of the interruption landscape. *Journal of Applied Psychology* 106, 10. pp. 1448.
 24. Li, Tianshi, Haines, Julia Katherine, De Eguino, Miguel Flores Ruiz, Hong, Jason I and Nichols, Jeffrey. (2022). Alert Now or Never: Understanding and Predicting Notification Preferences of Smartphone Users. *ACM Transactions on Computer-Human Interaction* 29, 5. pp. 1-33.
 25. Maleck, Moritz and Gross, Tom. (2024). TeamMeetingArranger: A Less Disruptive Way of “Do You Have a Minute?”. In *Proceedings of 16th International Conference on Intelligent Human Computer Interaction (IHCI 2024)* (Twente, The Netherlands). Springer. pp. 54-65.
 26. Maleck, Moritz and Gross, Tom. (2025). CLBoundaryManager: A System Facilitating Boundary Management and Border Crossing Between Life Domains. In *Proceedings of the 17th International Conference on Social Computing and Social Media (SCSM 2025), held as part of the 27th HCI International Conference (HCII 2025)* (Gothenburg, Sweden). Springer. pp. 335-350.
 27. Özdemir, Mustafa Can, Mottus, Mati and Lamas, David. (2024). Echoes of the Day: Exploring the Interplay Between Daily Contexts and Smartphone Push Notification Experiences. *Applied Sciences* 15, 1. pp. 14.
 28. Pflöging, Bastian, Fekety, Drea K, Schmidt, Albrecht and Kun, Andrew L. (2016). A model relating pupil diameter to mental workload and lighting conditions. In *CHI '16: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose California, USA). pp. 5776-5788.
 29. Puentes, Paola Ruiz, Soberanis-Mukul, Roger D, Acar, Ayberk, Gupta, Iris, Bhowmick, Joyraj, Li, Yizhou, Ghazi, Ahmed, Kazanzides, Peter, Wu, Jie Ying and Unberath, Mathias. (2023). Pupillometry in telerobotic surgery: A comparative evaluation of algorithms for cognitive effort estimation. *MedRobot* 1.
 30. Schaule, Florian, Johanssen, Jan Ole, Bruegge, Bernd and Loftness, Vivian. (2018). Employing consumer wearables to detect office workers' cognitive load for interruption management. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 2, 1. pp. 1-20.
 31. Shivakumar, Ashutosh, Bositty, Aishwarya, Peters, Nia S and Pei, Yong. (2020). Real-Time Interruption Management System for Efficient Distributed Collaboration in Multi-tasking Environments. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1. pp. 39, 1-23.
 32. Sweller, John. (1994). Cognitive load theory, learning difficulty, and instructional design. *Learning and Instruction* 4, 4. pp. 295-312.
 33. Szulewski, Adam, Kelton, Danielle and Howes, Daniel. (2017). Pupillometry as a Tool to Study Expertise in Medicine. *Frontline Learning Research* 5, 3. pp. 53-63.
 34. Trafton, J Gregory and Monk, Christopher A. (2007). Task interruptions. *Reviews of human factors and ergonomics* 3, 1. pp. 111-126.
 35. Van der Lippe, Tanja and Lippényi, Zoltán. (2020). Co-workers working from home and individual and team performance. *New technology, work and employment* 35, 1. pp. 60-79.
 36. Van der Wel, Pauline and Van Steenbergen, Henk. (2018). Pupil dilation as an index of effort in cognitive control tasks: A review. *Psychonomic Bulletin & Review* 25, 6. pp. 2005–2015.
 37. Wei, Ran and Lo, Ven-Hwei. (2006). Staying connected while on the move: Cell phone use and social connectedness. *New media & society* 8, 1. pp. 53-72.

38. Yuan, Xiaofang and Zhong, Lijuan. (2024). Effects of multitasking and task interruptions on task performance and cognitive load: considering the moderating role of individual resilience. *Current Psychology*. pp. 1-11.
39. Zagermann, Johannes, Pfeil, Ulrike and Reiterer, Harald. (2016). Measuring Cognitive Load using Eye Tracking Technology in Visual Computing. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization (BELIV '16)* (2016, Baltimore MD USA). Association for Computing Machinery. pp. 78 - 85.
40. Zagermann, Johannes, Pfeil, Ulrike and Reiterer, Harald. (2018). Studying eye movements as a basis for measuring cognitive load. In *Extended Abstracts of the 2018 CHI conference on human factors in computing systems* (Montreal QC Canada). pp. 1-6.