

Context-Aware Adaptive Visualizations for Critical Decision Making

Angela Lopez-Cardona^{a,*}, Mireia Masias Bruns^{a,1}, Nuwan T. Attygalle^{b,1}, Sebastian Idesis^a, Matteo Salvatori^a, Konstantinos Raftopoulos^c, Konstantinos Oikonomou^c, Saravanakumar Duraisamy^d, Parvin Emami^d, Nacera Latreche^b, Alaa Eddine Anis Sahraoui^b, Michalis Vakalellis^c, Jean Vanderdonckt^b, Ioannis Arapakis^a and Luis A. Leiva^d

^aTelefónica Scientific Research, Spain

^bUniversité catholique de Louvain, Belgium

^cAEGIS IT Research GmbH, Germany

^dUniversity of Luxembourg, Luxembourg

Abstract. Effective decision-making often relies on timely insights from complex visual data. While Information Visualization (InfoVis) dashboards can support this process, they rarely adapt to users' cognitive state, and less so in real time. We present SYMBIOTIK, an intelligent, context-aware adaptive visualization system that leverages neurophysiological signals to estimate mental workload (MWL) and dynamically adapt visual dashboards using reinforcement learning (RL). Through a user study with 120 participants and three visualization types, we demonstrate that our approach improves task performance and engagement. SYMBIOTIK offers a scalable, real-time adaptation architecture, and a validated methodology for neuroadaptive user interfaces.

1 Introduction

In today's information-saturated environment, users must decide what to process, filter, and prioritize. Although using visualizations to address this challenge has been widely discussed [22], Information visualization (InfoVis) systems have yet to reach their full potential to support critical decision-making. These systems leverage visual elements—charts, graphs, networks, and maps—to facilitate intuitive insights by engaging perceptual mechanisms rather than requiring laborious data comparisons [28]. The effectiveness of InfoVis in responding to critical events hinges on the visual representation of data and the speed at which actionable insights can be derived. We argue that physiological computing holds significant promise for innovation in InfoVis by incorporating human body signals, which can lead to more precise user interface adaptations through closed-loop implicit monitoring [7]. This is supported by evidence that cognitive processes are affected by user interface aesthetics, detectable through physiological signals [12, 21, 19, 11].

We argue that future solutions should consider a user-centric approach, premised on context awareness and emotion-sensing capabilities. In this regard, we propose an innovative adaptive framework inspired by symbiosis principles, where humans and machines cooperate and co-evolve to support decision-making processes [16]. To

this end, we envision a bidirectional channel of communication between AI and human users, which taps directly into human cognitive resources to facilitate an effortless and natural dialogue.

Recent literature highlights the growing need for adaptive visualization systems that consider individual differences in cognitive traits and mental states. According to Yanez et al. [30], each user experiences visualizations through their own lens, shaped by factors such as past experiences, personality, and cognitive abilities—all of which can significantly influence attention, speed, and task accuracy. Furthermore, Chiossi et al. [4] emphasizes the importance of multi-modal adaptive systems capable of capturing the full context of the user.

Prior studies have explored physiological signals to support attention modeling and adaptive interfaces. Salous et al. [24] developed *SmartHelm*, a smart helmet for cargo cyclists that uses multimodal signals and LSTM models to detect attention lapses. Toreini et al. [27] proposed attentive dashboards that rely solely on real-time Eye-tracking (ET) to deliver individualized visual attention feedback, improving attentional resource management during data exploration. Karran et al. [13] used EEG-based neurofeedback in a business context, applying an engagement index and a calibrated classifier to trigger adaptive interface changes, which led to better sustained attention and reduced error. More recently, Gehrke et al. [10] demonstrated a neuroadaptive haptics system that uses Reinforcement Learning (RL) to adapt haptic feedback in XR environments based on either explicit user ratings or EEG-based neural decoders.

Among the methodologies aimed at measuring brain activity and cognitive processes in the field of neuroimaging, Electroencephalography (EEG) stands out for its ability to capture electrical activity from the scalp with high temporal resolution, making it suitable for studying real-time neural dynamics. EEG is a non-invasive and portable method, offering practical advantages over other techniques. By leveraging EEG as a neurophysiological signal for adaptation, our system contributes to this vision—enabling more effective, personalized, and cognitively aligned data exploration experiences.

In contrast to previous work—none of which has applied EEG-based mental workload (MWL) estimation to InfoVis—our approach uses this signal to drive the real-time adaptation. We continuously track cognitive load and use RL agents to adapt the dashboard in real time. This enables more granular, workload-sensitive adapta-

* Corresponding Author. Email: angela.lopezcardona@telefonica.com, nuwan.attygalle@uclouvain.be

¹ Equal contribution.

tions, offering a novel contribution to interface personalization. To this end, we first studied which types of adaptations in different kinds of InfoVis dashboards result in enhanced performance and quicker response times in critical situations. We then trained RL agents with this data and implemented a neuro-feedback loop with context-aware RL agents during inference.

Based on an investigation case study involving different types of visualizations (Section 2), the proposed framework provides RL agents (Section 3.4) with environmental awareness through an embedding space of neurophysiological descriptors (Section 3.2), captured using commodity sensors (Section 3.1). These descriptors drive adaptations of the dashboard interface (Section 3.5) within a closed-loop, implicit monitoring framework (Section 3).

2 Use Case

The use case we address involves the investigation of a criminal network by a Law Enforcement Agency (LEA). We created an adaptive Crime Investigation Dashboard (CID), built upon AEGIS's Advanced Visualization Toolkit (AVT)². The formulation of relationships across individuals, the underlying factors determining these connections, and the emergence of cohesive groups (i.e. cliques) with dense links are analysed to uncover hidden patterns and understand behaviours and relations within the network.

To this end, information foraging of data involving interactions/communications within a specific period is investigated. Call Data Records (CDRs) are acquired from relevant telecommunications operators after identifying telephone numbers associated with the group i.e. person(s) of interest of specific phone numbers and devices. This is achieved by analyzing the available historical information extracted from files across multiple modalities, all grounded in the CDR-based interaction patterns previously described.

The CID includes a distribution chart, a timeline, and a network graph (Figure 1). A graphical representation of the volume of calls within the preferred time window is provided by the distribution chart. The timeline offers a breakdown of time-related metrics, focusing on the frequency and duration of calls, and the involved parties. Finally, the network graph visualization, using undirected graphs, presents information about criminal records, including arrest numbers, gender, birthplace, and incarceration states, thereby aiding in the visual identification of latent relationships and outliers.



Figure 1. Crime Investigation Dashboard.

These visualisation layouts (i.e., graph, distribution, and timeline) were selected for their ecological validity they provide (as they often appear in production-ready visualisation platforms such as the AVT), supported by a range of real-world visualization scenarios. Each visualization was developed following user-interface adaptation and information visualisation principles, systematically introducing variations to study their effect on participants' interactions with InfoVis environments and subsequent task performance. Said visualization

principles include well-studied visual primitives such as color, size, area, and line patterns, which map to different layout-specific adaptations to visual elements (e.g., color, shape, size, thickness, position), depending on the context and setting. Subsequently, the RL agent is trained to select the optimal adaptation in real time.

2.1 Data Collection

A total of 120 participants (71 male, 49 female, aged 17–54) were recruited for a 1h study with a 25€ voucher as compensation. The goal is to train a dedicated RL agent for each visualization type. A dataset per type was crafted, each with data from disjoint sets of 40 participants (between-subjects design). The visualizations included the original versions and modifications designed to improve performance, featuring fully and partially adapted content. After each presentation, participants completed a question-answering task based on the displayed content. This allowed the final dataset to be categorized according to combinations of question and visual complexities (i.e., low vs. high), and enabled quantification of task execution efficiency in terms of accuracy and reaction time. Neurophysiological and behavioral signals (i.e. EEG and ET) were recorded jointly for the neurobiological definition of a MWL proxy.

3 System Description

The proposed framework employs the Event-Condition-Action paradigm to model and execute context-aware UI adaptations [20] for the AEGIS' AVT dashboard (Section 2). This loop consists of a driver that connects, in real time, to the measurement devices and minimally preprocesses the raw data in 2-second epoch cycles—the Physiological Sensing (PS) module (Section 3.1) and the Behavior Sensing (BS) module that captured events from the dashboard such as the layout used by the user in real time. This is followed by the Compact Data Representation (CDR) module (Section 3.2), which is responsible for engineering the cognitive features from the raw signals. Next, the Neural-Mechanistic Understanding (NMU) module (Section 3.3) infers the MWL state from the derived features, which is sent to the RL module (Section 3.4), which will compute the adaptation and send it to the Adaptation Engine (AE) module (Section 3.5) responsible for executing the changes in the dashboard.




Figure 2. SYMBIOTIK high-level software architecture.

3.1 Physiological Sensing Module

This module interfaces with measurement devices, transmitting raw data in 2-second epochs. Specific Python libraries, like *socket*, enable the connection to different EEG devices and other neurophysiological signal sources, including ET that allow future integration.

² <https://aegisresearch.eu/solutions/advanced-visualization-toolkit/>

3.2 Compact Data Representation

The raw data are then processed in the compact data representation module. Minimal preprocessing is performed to support real-time inference. Band-pass filtering (0.5–40 Hz) is performed, followed by power spectral density which is computed by taking the squared magnitude of the Fast-Fourier transformation (FFT). The resulting power is computed for the frequency bands of interest using electrodes P3 (alpha), C3 (beta), and Fz (theta and delta) (Figure 3).




Figure 3. EEG preprocessing pipeline: (A) Raw signal recording; (B) Bandpass filtering; (C) FFT to convert to frequency domain; (D) Time-frequency analysis to compare every 2 seconds the power of the following frequency bands: (E) Delta, Theta, Alpha, and Beta.

3.3 Neural-Mechanistic Understanding Module

Building on prior work [14, 23, 25], this module defines a preliminary MWL construct by categorizing each band's activity into low, medium, and high levels based on population-level quantiles from 120 samples. The 25th and 75th percentiles from the training data set thresholds for inference. For bands positively correlated with MWL (theta, beta), values below the 25th percentile are considered low, above the 75th are high, and in between are medium or optimal. For negatively correlated bands (alpha, delta), this classification is reversed. Categories per band are linearly combined with fixed weights, pending further optimization.

3.4 Reinforcement Learning Module

A RL framework was formulated as a smart system, where an RL agent interacts with the user and the dashboard, recommending adaptation actions to maximize cumulative rewards. Specifically, a separate agent is trained for each type of layout (graph, timeline, or distribution). To determine which agent to use at inference time, the system detects which representation the user is currently viewing through the BS module. The RL module retrieves a configuration file via an API to determine the current dashboard adaptation and the question being addressed. Additionally, a dedicated classifier is trained to predict the difficulty level of the question at inference time, using popular pre-trained lightweight models for sentence classification fine-tuned on our original questions based on BERT [5] and RoBERTa [17]. For each representation, a behavior policy is learned using the corresponding dataset (Section 2.1). The reward measures whether the selected action results in an optimal MWL. During training, an importance-weighted estimate of the reward is used, accounting for differences between the behaviour policy that generated the dataset and the target policy the agent is learning [8]. The agent learns an optimal policy using a tabular Q-learning method [29] with

epsilon-greedy exploration. The algorithm operates in an environment where episodes consist of a single transition, and learning happens iteratively through a sequence of updates based on observed rewards. During training, the agent updates its Q-table (a matrix of size $n_{states} \times n_{actions}$) based on the experiences in the dataset. The trained agent receives the current state as input and outputs the most appropriate action. The action space (**A**) encompasses all possible adaptations within each representation. The selected action, with its active representation, is then forwarded to the adaptation engine.

3.5 Adaptation Engine

The Adaptation Engine (AE) receives an adaptation strategy suggested by the RL module along with a visualization type (e.g., network, timeline, distribution). Adaptation strategies are categorized as "No Adaptation", "Partial Adaptation" (i.e., with attributes such as color, shape, size, proximity, and thickness), or "Full Adaptation" [6, 26]. It then queries the Adaptation Catalogue, which maps a strategy to actionable adaptation operations based on ergonomic criteria [1] and linguistic levels [2]. The AE encodes the most appropriate operation in a JSON configuration file to be sent to the AVT dashboard via the InfoVis Gateway to adapt the final user interface accordingly, which is in turn subject to evaluation by EEG [9].

3.6 Infovis Gateway

This component serves as a middleware between the AP and the Dashboard. It serves as a centralized orchestrator, combining Web-Socket protocols for real-time dashboard updates with Kafka³ message broker to manage high-volume data streams between the components. Kafka Broker decouples services, enabling asynchronous event handling. Each component publishes its information to Kafka topics that are organized by conceptual categories (e.g., biosignals, strategies, adaptations). Any component interested in a particular type of data can subscribe to the relevant topic, consume updates as they are published, and react autonomously. This event-driven architecture significantly improves system scalability, flexibility, observability, real-time adaptability, and modularity. Components can operate independently without requiring knowledge of each other's existence, facilitating easier integration of new modules and allowing the system to dynamically adapt to evolving requirements.

4 Conclusions and Future Work

SYMBIOTIK implements a scalable, real-time adaptation architecture, and a validated methodology for neuroadaptive user interfaces, by training RL agents with real-world data across different layout types. An important avenue for future work involves developing a robust and personalized MWL metric. Preliminary analysis of our dataset revealed that behavioural indicators (e.g., accuracy and response time) show great promise in predicting question and visual complexities, as well as the effects of adaptation. These findings suggest that a multimodal approach leveraging multiple signals may help define a better MWL proxy with real-time applicability, e.g., through strategies such as latent variable modelling. Also, using ET to indicate visual attention is a key research direction [18]. It reflects the perceived importance of design elements [3] and supports the development of task-specific cognitive state models generalizable across participants [15] and layouts. The final aim is to create a common and layout-agnostic agent, able to generalize in diverse settings.

³ <https://kafka.apache.org/>

Acknowledgements

This research is supported by Horizon Europe's European Innovation Council through the Pathfinder program (SYMBIOTIK project, grant 101071147) and by the Industrial Doctorate Plan of the Department of Research and Universities of the Generalitat de Catalunya, under Grant AGAUR 2023 DI060.

References

- [1] J. M. C. Bastien and D. L. Scapin. Preliminary findings on the effectiveness of ergonomic criteria for the evaluation of human-computer interfaces. In *INTERACT '93 and CHI '93 Conference Companion on Human Factors in Computing Systems*, CHI '93, page 187–188, New York, NY, USA, 1993. Association for Computing Machinery. ISBN 0897915747. doi: 10.1145/259964.260198. URL <https://doi.org/10.1145/259964.260198>.
- [2] F. Bodart and J. Vanderdonckt. Using ergonomic rules for evaluation by linguistic ergonomic criteria. In Y. Anzai, K. Ogawa, and H. Mori, editors, *Symbiosis of Human and Artifact*, volume 20 of *Advances in Human Factors/Ergonomics*, pages 367–372. Elsevier, 1995. doi: [https://doi.org/10.1016/S0921-2647\(06\)80244-3](https://doi.org/10.1016/S0921-2647(06)80244-3). URL <https://www.sciencedirect.com/science/article/pii/S0921264706802443>.
- [3] G. Cartella, M. Cornia, V. Cuculo, A. D'Amelio, D. Zanca, G. Boccignone, and R. Cucchiara. Trends, Applications, and Challenges in Human Attention Modelling. volume 9, pages 7971–7979, 2024. doi: 10.24963/ijcai.2024/882. URL <https://www.ijcai.org/proceedings/2024/882>. ISSN: 1045-0823.
- [4] F. Chiossi, J. Zagermann, J. Karolus, N. Rodrigues, P. Balestrucci, D. Weiskopf, B. Ehinger, T. Feuchtnner, H. Reiterer, L. L. Chuang, M. Ernst, A. Bulling, S. Mayer, and A. Schmidt. Adapting visualizations and interfaces to the user. *it - Information Technology*, pages 133–143, 2022. doi: 10.1515/itit-2022-0035.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In J. Burstein, C. Doran, and T. Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- [6] D. Eloi, A. Sahraroui, J. Vanderdonckt, and L. A. Leiva. User-controlled form adaptation by unsupervised learning. In *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction, NordiCHI 2024, Uppsala, Sweden, October 13-16, 2024*, pages 16:1–16:8. ACM, 2024. doi: 10.1145/3677045.3685431. URL <https://doi.org/10.1145/3677045.3685431>.
- [7] S. H. Fairclough. *Current research in neuroadaptive technology*, chapter Designing human-computer interaction with neuroadaptive technology. Academic Press, 2022. doi: 10.1016/B978-0-12-821413-8.00006-3.
- [8] S. Fujimoto, D. Meger, and D. Precup. Off-Policy Deep Reinforcement Learning without Exploration, 2019. URL <http://arxiv.org/abs/1812.02900>. arXiv:1812.02900 [cs].
- [9] D. Gaspar-Figueiredo, S. Abrahão, E. Insfrán, and J. Vanderdonckt. Measuring user experience of adaptive user interfaces using EEG: A replication study. In *Proceedings of the 27th International Conference on Evaluation and Assessment in Software Engineering, EASE 2023, Oulu, Finland, June 14-16, 2023*, pages 52–61. ACM, 2023. doi: 10.1145/3593434.3593452. URL <https://doi.org/10.1145/3593434.3593452>.
- [10] L. Gehrke, A. Koselevs, M. Klug, and K. Gramann. Neuroadaptive Haptics: Comparing Reinforcement Learning from Explicit Ratings and Neural Signals for Adaptive XR Systems, 2025. URL <http://arxiv.org/abs/2504.15984>. arXiv:2504.15984 [cs].
- [11] S. Haddad, K. Latifzadeh, S. Duraisamy, J. Vanderdonckt, O. Daassi, S. Belghith, and L. A. Leiva. Good GUIs, Bad GUIs: Affective evaluation of graphical user interfaces. In *Proc. UMAP*, 2024. doi: 10.1145/3627043.3659549.
- [12] L. M. Hirshfield, R. Gulotta, S. Hirshfield, S. Hincks, M. Russell, R. Ward, T. Williams, and R. Jacob. This is your brain on interfaces: enhancing usability testing with functional near-infrared spectroscopy. In *Proc. CHI*, 2011. doi: 10.1145/1978942.1978996.
- [13] A. J. Karan, T. Demazure, P.-M. Leger, E. Labonte-LeMoine, S. Senecal, M. Fredette, and G. Babin. Toward a Hybrid Passive BCI for the Modulation of Sustained Attention Using EEG and fNIRS. *Frontiers in Human Neuroscience*, 13, 2019. ISSN 1662-5161. doi: 10.3389/fnhum.2019.00393. URL <https://www.frontiersin.org/journals/human-neuroscience/articles/10.3389/fnhum.2019.00393/full>. Publisher: Frontiers.
- [14] E. Kutafina, A. Heiligers, R. Popovic, A. Brenner, B. Hankammer, S. M. Jonas, and J. Zweerings. Tracking of mental workload with a mobile eeg sensor. *Sensors*, 21(15):5205, 2021. doi: 10.3390/s21155205.
- [15] M. Langner, P. Toreini, and A. Maedche. Cognitive state detection with eye tracking in the field: an experience sampling study and its lessons learned. pages 109–129, 2024. ISSN 2196-6826. URL <https://dl.gi.de/handle/20.500.12116/43992>. Publisher: De Gruyter.
- [16] L. A. Leiva, I. Arapakis, M. Vakalellis, and J. Vanderdonckt. Context-aware Adaptive Visualizations for Critical Decision Making. CHI Workshop on Future of Computational Approaches for Understanding and Adapting User Interfaces, 2023.
- [17] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach, 2019. URL <http://arxiv.org/abs/1907.11692>. arXiv:1907.11692 [cs].
- [18] A. Lopez-Cardona, P. Emami, S. Idesis, S. Duraisamy, L. A. Leiva, and I. Arapakis. A comparative study of scanpath models in graph-based visualization. In *Proceedings of the ACM Symposium on Eye Tracking Research and Applications (ETRA)*. ACM, 2025.
- [19] K. Lukanov, H. A. Maior, and M. L. Wilson. Using fNIRS in usability testing: Understanding the effect of web form layout on mental workload. In *Proc. CHI*, 2016. doi: 10.1145/2858036.2858236.
- [20] V. G. Motti and J. Vanderdonckt. A computational framework for context-aware adaptation of user interfaces. In R. J. Wieringa, S. Nurcan, C. Rolland, and J. Cavarero, editors, *IEEE 7th International Conference on Research Challenges in Information Science, RCIS 2013, Paris, France, May 29-31, 2013*, pages 1–12. IEEE, 2013. doi: 10.1109/RCIS.2013.6577709. URL <https://doi.org/10.1109/RCIS.2013.6577709>.
- [21] E. M. M. Peck, B. F. Yuksel, A. Ottley, R. J. Jacob, and R. Chang. Using fNIRS brain sensing to evaluate information visualization interfaces. In *Proc. CHI*, 2013. doi: 10.1145/2470654.2470723.
- [22] L. Perkhofor, O. Lehner, and H. Treiblmaier. InfoVis: The impact of information overload on decision making outcome in high complexity settings. In *Proc. AIS SIGHCI Pre-ICIS Workshop*, 2016. doi: <https://aisel.aisnet.org/sighci2016/3>.
- [23] S. Puma, N. Matton, P. V. Paubel, É. Raufaste, and R. El-Yagoubi. Using theta and alpha band power to assess cognitive workload in multitasking environments. *International Journal of Psychophysiology*, 123:111–120, 2018. doi: 10.1016/j.ijpsycho.2017.11.002.
- [24] M. Salous, D. Küster, K. Scheck, A. Dikfidan, T. Neumann, F. Putze, and T. Schultz. SmartHelm: User Studies from Lab to Field for Attention Modeling. In *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1012–1019, 2022. doi: 10.1109/SMC53654.2022.9945155. URL <https://ieeexplore.ieee.org/document/9945155>. ISSN: 2577-1655.
- [25] W. So, S. Wong, J. Mak, and R. Chan. An evaluation of mental workload with frontal eeg. *PloS one*, 12(4):e0174949, 2017. doi: 10.1371/journal.pone.0174949.
- [26] K. Todi, G. Bailly, L. A. Leiva, and A. Oulasvirta. Adapting user interfaces with model-based reinforcement learning. In Y. Kitamura, A. Quigley, K. Isbister, T. Igarashi, P. Bjørn, and S. M. Drucker, editors, *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*, pages 573:1–573:13. ACM, 2021. doi: 10.1145/3411764.3445497. URL <https://doi.org/10.1145/3411764.3445497>.
- [27] P. Toreini, M. Langner, A. Maedche, S. Morana, and T. Vogel. Designing Attentive Information Dashboards. *Journal of the Association for Information Systems*, 23(2):521–552, 2022. ISSN 1536-9323. doi: 10.17705/1/jais.00732. URL <https://aisel.aisnet.org/jais/vol23/iss2/4>.
- [28] Y. Wang, W. Wang, A. Abdelhafez, M. Elfars, Z. Hu, M. Bâce, and A. Bulling. SalChartQA: Question-driven Saliency on Information Visualisations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pages 1–14. Association for Computing Machinery, 2024. ISBN 9798400703300. doi: 10.1145/3613904.3642942. URL <https://doi.org/10.1145/3613904.3642942>.
- [29] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3):279–292, 1992. ISSN 1573-0565. doi: 10.1007/BF00992698. URL <https://doi.org/10.1007/BF00992698>.
- [30] F. Yanez, C. Conati, A. Ottley, and C. Nobre. The State of the Art in User-Adaptive Visualizations. *Computer Graphics Forum*, 44(1):e15271, 2025. ISSN 1467-8659. doi: 10.1111/cgf.15271.