

Managing Information Overload in Large-Scale Distributed Mixed-Reality Meetings

KATJA KRUG*, WOLFGANG BÜSCHEL, and MATS OLE ELLENBERG*,

Interactive Media Lab Dresden, TUD Dresden University of Technology, Germany

Large-scale distributed mixed-reality meetings involve many people and their audiovisual representations. These collaborative environments can introduce challenges such as sensory overload, cognitive strain, and social fatigue. In this paper, we discuss how the unique adaptability of Mixed Reality can be leveraged to weaken these stressors by managing information overload.

CCS Concepts: • **Human-centered computing** → **Mixed / augmented reality**; **Computer supported cooperative work**.

Additional Key Words and Phrases: Mixed Reality, Collaboration, Adaptive Systems

ACM Reference Format:

Katja Krug, Wolfgang Büschel, and Mats Ole Ellenberg. 2025. Managing Information Overload in Large-Scale Distributed Mixed-Reality Meetings. In *Proceedings of ACM CHI 2025 Workshop WS 12: Scaling Distributed Collaboration in Mixed Reality (CHI'25)*. ACM, New York, NY, USA, 5 pages.

1 Introduction and Background

In recent years, commercially available head-mounted displays (HMDs) have brought Virtual and Mixed Reality (VR and MR) technology into homes and workplaces. With video see-through MR being supported by most current-gen devices, applications can blend physical and virtual worlds, paving the way for immersive, mixed-reality collaborations, where distributed users come together as virtual avatars. In a future where such devices are more widespread, supporting usage scenarios with larger groups of users will become crucial. One fundamental challenge of scaling up these mixed-reality collaborations is to support human cognition and limit information overload. Real-world, large-scale meetings with dozens or even hundreds of participants are often seen as stressful and unpleasant because of factors such as environmental overstimulation, i.e., information overload and noise, which can lead to disorientation and simplified information processing [6], hindering complex collaboration. The social overload of large-scale meetings can further induce fatigue and reduced engagement [10]. In MR, perceptual issues such as mismatching virtual and real spaces [3], and insufficient HMD ergonomics [2] can be further additional stressors.

In the well-established field of research of adaptive user interfaces, many efforts have been made to support cognition and optimize interaction, e.g., data in MR interfaces can be filtered or clustered to reduce information density [7] and semantic level-of-detail rendering [9] can be used to optimize content to the available screen space. Context-aware adaptations for Mixed Reality have been investigated, among others, by Lindlbauer et al. [5] and Lie et al. [4]. While most of these adaptations at least implicitly support cognition, some systems directly aim at supporting users in challenging contexts, such as when driving [8]. However, we see untapped potential in researching adaptations for large-scale

* Also with Centre for Tactile Internet with Human-in-the-Loop (CeTI)

Authors' Contact Information: [Katja Krug](mailto:katjakrug@acm.org), katjakrug@acm.org; [Wolfgang Büschel](mailto:bueschel@acm.org), bueschel@acm.org; [Mats Ole Ellenberg](mailto:mellenberg@acm.org), mellenberg@acm.org,
Interactive Media Lab Dresden, TUD Dresden University of Technology, Dresden, Germany.

distributed MR systems specifically, where any adaptive change to how and what content is shown can not be seen in isolation. Instead, the impact on collaboration and communication has to be considered as well.

In this paper, we explore and discuss such context- and preference-dependent adaptations and techniques in relation to the presentation and distribution of individuals in large-scale MR meetings. We further describe methods to control such functionality, including user-driven and automatic approaches.

2 Adaptations and Techniques

In the following, we list examples of techniques MR systems could employ to reduce cognitive overload and support collaboration in large-scale MR meetings. This list is not comprehensive, instead it is aimed to be an inspirational starting point for discussions about MR adaptations. Many of these techniques are based on the perceived relevance of people to individual users. This relevance is highly personal and can mean a pre-determined person of interest, an active speaker, a person in one's line of sight, or a potential collaborator, etc. We further discuss these dependencies in [section 3](#).

2.1 Visual Rendering, Filtering, and Grouping of Avatars

Enhancing focus and reducing distractions, adaptive visual filtering techniques can be employed to prioritize information and minimize visual complexity in distributed MR meetings. One approach is **idle avatar reduction**, where inactive avatars could fade out or be replaced by minimal icons until they re-engage. To reduce the visual noise of excessive movement, **key gesture filtering** could limit animations to important communicative gestures like pointing or nodding, and otherwise only display subtle idle animations to avoid the eeriness of completely still avatars. **Relevance-based avatar scaling**, i.e., adjusting avatar size based on importance, could make key individuals more prominent and easy to find while reducing visual clutter from bystanders. Similarly, **selective rendering and transparency** can de-emphasize less relevant avatars by abstracting or fading them out, and enhance important ones by highlighting them with a certain color or additional visualizations, or adding fine details like facial expressions. **Visual avatar grouping** can help structuring the social space by clustering related individuals together, e.g., people from the same work group. This can be done through proximity or adding shared visual traits such as color-coding or uniform clothing. Clusters of avatars could also be combined into cohesive "crowd blobs" that only disperse upon focus, or into a single representative avatar that splits into multiple when triggered.

2.2 Audio Management

To avoid users becoming overwhelmed by the competing voices and background noise of large crowds, adaptive audio management techniques can help prioritizing important sounds and suppress distractions. **Audio prioritization and speech isolation** could emphasize relevant voices while fading out or muting less important ones. To prevent auditory overload, **conversational blur** can merge excessive background voices into a subtle, incomprehensible background noise, ensuring that users remain aware of if and where active conversations are taking place without being overwhelmed. The system can then dynamically clarify speech, e.g., when only few people are talking and most are listening, or when important speakers are talking. Additionally, **adaptive noise suppression** can help regulating background sound and maintain group awareness, e.g., as a large group gradually quiets down, the system can dynamically strengthen the noise suppression to facilitate upcoming announcements.

2.3 Environment Partitioning and Crowd Distribution

Unlike physical environments, Mixed Reality spaces are flexible and can dynamically adapt to user needs, improving coordination and cognitive ease. **Dynamic room partitioning** can help distribute avatars evenly to prevent overcrowding by adjusting the layout of the virtual space. **Virtual borders** are an alternative to the visual grouping mentioned in subsection 2.1, emphasizing these partitions by erecting virtual walls or subtle floor markings that encircle cohesive groups, which could trigger functionality upon crossing, like making certain conversations audible. In-between, **generated pathways** can ensure smooth navigation by establishing routes where movement is allowed and gathering is restricted. To further facilitate discussions, **auto grouping** could detect common topics and dynamically form discussion clusters, gently relocating them to designated areas. Additionally, **elastic room scaling** can leverage MR's lack of physical constraints by expanding the virtual space past the physical one, subtly pushing less relevant groups outwards to maintain focus on key conversations. **Density balancing encouragements**, in the form of subtle visual cues, could encourage users to spread out in crowded areas, promoting a more comfortable and breathable virtual space. **Dedicated privacy zones** could allow users to retreat into a quiet, shielded area for distraction-free conversation. For those feeling overwhelmed, a **virtual palate cleanser environment** can offer a calming space to mentally reset, similar to stepping outside for fresh air. In extreme cases of mental overload, a **quick escape mechanism** should allow users to temporarily leave the virtual environment without exiting the application, keeping their avatar in an idle state, labeled for return.

2.4 AI Assisted Attention Management

In real-life interactions, body language, context, and physical proximity can naturally guide attention. In virtual environments, these cues are limited, requiring additional guidance support to avoid cognitive overload. The system could provide on-demand **conversation summaries**, i.e., brief text or audio recaps of relevant conversations, helping users decide if they want to join a discussion group. Similarly, **content labels** could continuously display dynamically adapting key conversation topics above groups, along with optional icons indicating the emotional tone of the conversation, e.g., friendly, focused, or intense. To aid the discovery of relevant discussions, suggestions for relevant groups or individuals can be made through **engagement predictions**, generated based on user interests, profiles, or prior interactions. Additionally, **keyword-based filtering** can refine the ongoing auditory input, ensuring that users only hear conversations containing relevant terms that are either user-defined or AI generated. **Smart notifications** can dynamically alert and potentially guide users to new participants or relevant discussions, timing notifications strategically to minimize disruptions while enhancing awareness.

3 Adaptation Method

The effectiveness of adaptive techniques for managing information overload in MR depends not only on what is adapted but also on how these adaptations are triggered. These features range from fully user-driven (full control) to fully system-driven (automation without user intervention).

3.1 User-Driven

Here, users manually trigger and customize adaptations. This provides maximum control and transparency but requires a conscious effort and decision-making. User-driven systems can become complex and hard to generalize.

Explicit User Decisions: The user maintains full control throughout and actively makes decisions, such as determining people of interest and their rendering adaptations, or hand-drawing privacy bubbles. While the customizability makes the system highly inclusive and accessible, this approach risks decision fatigue [1] and cognitive overload, since the system quickly becomes complex and hard to operate.

Implicit System Reactions: The adaptations respond automatically to user actions, such as adjusting avatar relevance based on the user’s current Field of View (FoV) or initializing calm-down techniques when detecting a rising heart rate. This reduces manual user effort while maintaining some level of control. However, compared to the explicit decisions, this raises privacy and security concerns, as it requires the processing of sensitive data. It also risks misinterpretation, such as unintentional audio shifts due to FoV movement (Midas touch effect).

3.2 System-Driven

In system-driven adaptations, the system autonomously adjusts the settings in real time based on what it learns from user behavior, social dynamics or environmental factors. This allows for seamless integration and reduced cognitive load, but might cause a feeling of loss of control and disorientation, raises privacy concerns, and might even negatively influence collaboration, if the system’s predictions and conclusions are inaccurate.

User Parameter Dependent The system automatically personalizes adaptations based on the user profile, such as highlighting co-workers, emphasizing discussion contributions by higher ranking roles, alerting users about conversations adjacent to their expertise. This facilitates networking while reducing its awkwardness, but risks echo chambers and neglects potentially shifting goals and interests during conversations.

Context-Triggered Adaptations respond to the user’s current context, such as visually emphasizing a speaker with relevant contributions to the topic, or muting background noise when the user is part of a discussion group. This subtly improves focus and enables smooth transitions between contexts, but may feel unexpected if the user is not aware of the context changing. The system might also misinterpreted the context.

Socially-Adaptive The adaptations are based on interpersonal relationships and group behavior, such as visually emphasizing a speaker that is being looked at by many people, or clustering habitual conversation partners and guiding them towards each other in the crowd. This helps strengthen interpersonal relationships and leadership emergence, but can reinforce bias, suppress less prominent voices and limit new connections.

4 Conclusion

Large-scale MR meetings can induce diverse cognitive challenges from information overload to social fatigue. To address these, we explored adaptations and techniques that dynamically manage sensory input and support communication. These techniques exist along the user agency spectrum, ranging from complete user control to fully automated adaptations. While the user-driven methods offer transparency and customization, they can lead to decision fatigue and interaction complexity. On the other hand, fully automated solutions reduce cognitive effort but risk misinterpretation, loss of control, and raise privacy concerns. Hybrid user-system solutions could act as a promising middle-ground between these ends of the spectrum, balancing automation with user agency. AI could suggest optimizations that users can refine or override. Users can make active selections of people and spaces and AI adaptations honor and emphasize these selections. Adaptations should be unobtrusive, assistive, and customizable. They should enhance collaboration and keep users empowered without dictating their experience. Transparency is crucial and users should always understand how and why their environment is adapting. The thoughtful design of these MR adaptations can reduce cognitive strain and create more comfortable and intuitive distributed large-scale MR meetings.

Acknowledgments

This work has been funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden. The authors further acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of “Souverän. Digital. Vernetzt.”. Joint project 6G-life, project identification number: 16KISK001K.

References

- [1] Roy F Baumeister, Ellen Bratslavsky, Mark Muraven, and Dianne M Tice. 2018. Ego depletion: Is the active self a limited resource? In *Self-regulation and self-control*. Routledge, 16–44. doi:10.1037/0022-3514.74.5.1252
- [2] Yumiao Chen, Xin Wang, and Huijia Xu. 2021. Human factors/ergonomics evaluation for virtual reality headsets: a review. *CCF Transactions on Pervasive Computing and Interaction* 3, 2 (2021), 99–111. doi:10.1007/s42486-021-00062-6
- [3] David Drascic and Paul Milgram. 1996. Perceptual issues in augmented reality. In *Stereoscopic Displays and Virtual Reality Systems III*, Mark T. Bolas, Scott S. Fisher, Mark T. Bolas, Scott S. Fisher, and John O. Merritt (Eds.), Vol. 2653. International Society for Optics and Photonics, SPIE, 123 – 134. doi:10.1117/12.237425
- [4] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Streli, and Christian Holz. 2024. SituationAdapt: Contextual UI Optimization in Mixed Reality with Situation Awareness via LLM Reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. ACM, New York, NY, USA, Article 43, 13 pages. doi:10.1145/3654777.3676470
- [5] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-Aware Online Adaptation of Mixed Reality Interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (UIST '19)*. ACM, New York, NY, USA, 147–160. doi:10.1145/3332165.3347945
- [6] Peter Suedfeld. 1985. Stressful Levels of Environmental Stimulation. *Issues in Mental Health Nursing* 7, 1-4 (1985), 83–104. doi:10.3109/01612848509009451
- [7] Markus Tatzgern, Valeria Orso, Denis Kalkofen, Giulio Jacucci, Luciano Gamberini, and Dieter Schmalstieg. 2016. Adaptive information density for augmented reality displays. In *2016 IEEE Virtual Reality (VR)*. IEEE, 83–92. doi:10.1109/VR.2016.7504691
- [8] Shaoyue Wen, Songming Ping, Jialin Wang, Hai-Ning Liang, Xuhai Xu, and Yukang Yan. 2024. AdaptiveVoice: Cognitively Adaptive Voice Interface for Driving Assistance. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, New York, NY, USA, Article 253, 18 pages. doi:10.1145/3613904.3642876
- [9] Abby Wysopal, Vivian Ross, Joyce Passananti, Kangyou Yu, Brandon Huynh, and Tobias Höllerer. 2023. Level-of-Detail AR: Dynamically Adjusting Augmented Reality Level of Detail Based on Visual Angle. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. ACM, 63–71. doi:10.1109/VR55154.2023.00022
- [10] Xin Zhang, Xiao yan Ding, Gao shan Wang, and Liang Ma. 2020. Investigating the influences of social overload and task complexity on user engagement decrease. *Total Quality Management & Business Excellence* 31, 15-16 (2020), 1774–1787. doi:10.1080/14783363.2018.1509698

Received 14 March 2025; accepted 24 March 2025