Body-Based Augmented Reality Feedback During Conversations

HENNING POHL, Aalborg University, Denmark

Fig. 1. We investigate the design space of displaying information on and around the bodies of other people. In conversational situations, this allows users to keep focus on the conversation partner while still being able to receive incoming information, including information relevant to the conversation itself. Shown here are five augmentations we evaluated in a video-call study: name tag, emotion tracker, weather notification, message notification, and conversation clock. All visualizations shown as direct screen captures from within Snap's Lens Studio with AI-generated persons used for illustrative purposes.

Engaging with our devices as we engage with each other is problematic as it distracts us and diminishes our social interactions. Subtle interactions have been presented as an approach to reconcile personal and computing interactions, through less disrupting technology. Along those lines, we investigate showing information right on and next to the people we are engaging with. Body-based data visualization allows us to maintain our attention with others, but to also receive information at the same time. We explore potential designs of such body-based and especially on-face visualizations and create a set of five prototype visualizations in a Snapchat lens. We use these prototypes in a video call study with 16 participants to evaluate how body-based visualizations affect actual conversations.

CCS Concepts: • **Human-centered computing** → **Mixed / augmented reality**; *Information visualization.*

Additional Key Words and Phrases: Notifications, Body-Based Visualization, Augmented Reality

ACM Reference Format:
Henning Pohl. 2024. Body-Based Augmented Reality Feedback During Conversations. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 246 (September 2024), 22 pages. https://doi.org/10.1145/3676491

1 Introduction

A core aspect of personal computing has been to extend ways of access to digital information. We look up a contact, glance at our next appointment, set out to find a good nearby restaurant, or just check the time. Other information is pushed to us, such as incoming emails and notifications. Devices like phones and smartwatches provide near-ubiquitous access and are with us most of the day. Our engagement with information also is not a solitary activity and commonly occurs together

Author's Contact Information: Henning Pohl, henning@cs.aau.dk, Aalborg University, Aalborg, Denmark.

ACM 2573-0142/2024/9-ART246
https://doi.org/10.1145/3676491

with or in the presence of others. Yet, this introduces competing concerns: how to reconcile our desire to engage with others with retaining access to relevant information?

A range of approaches have emerged to allow users to interact with their devices while also engaging with others. For example, Hansson and Ljungstrand's *Reminder Bracelet* [16] used discrete notification cues so that device use is less intrusive. Moreover, Anderson et al. have coined the term "deceptive devices" and described several prototypes that allow users to hide their interactions with computers from others [2]. Several forms of subtle interaction [41] strive to make interactions less disruptive and unobtrusive, for example, by enabling input through small-scale gestures [10]. Underlying intentions here range from active deception of others to more benevolent consideration of them (e.g., by desiring to pay them attention and not to disturb them). Current approaches to less-disruptive devices commonly offer a lower bandwidth for interaction. Subtle vibrations, a visual display nudged away where others cannot see it, small gestural interactions, and messages hidden in plain sight all are limited compared to standard device use. If we want access to important and rich information, but also want to stay engaged with each other, a different approach is needed.

While subtle feedback is suitable for low-fidelity information it often has its limitations, particularly when it is non-visual. For example, while a small vibration can alert us of an incoming message it is not enough for glancing at it or even reading it in full. Yet, if this visual information is displayed on a screen it inevitably will distract us from others in a noticeable way. As an alternative, previous work has suggested using the space around people we interact with to display information [8, 21]. Mediated through augmented reality (AR) devices, such "paracentral visualizations" allow users to stay focused on a person in front of them, while still receiving rich visual information. Another approach is glanceable AR [30], where virtual content does not take over the center of the visual field but instead resides in the periphery. While the space of such more peripheral visualizations is thus well explored already, we investigate an alternative approach where information is displayed more closely to others. We posit that moving visualizations closer reduces the need for glancing away and allow better focus and engagement with what is in front of us. With respect to the social situations described above, we hence ask the question: *What if information was displayed right on and next to the people we converse with, in order to allow our attention to remain on them?*

In this paper we investigate *body-based visualization* as a means for subtle and non-intrusive access to information. We construct visualizations for a range of pieces of information that are common in mobile interaction or that users might want to have access to during a conversation. This includes temporal information, notifications, information on the person they are engaged with, and task-based support information. We prototype five forms of such body-based visualization in a Snapchat lens. In an evaluation with 16 participants, we then investigate the impact of body-based visualizations in a conversational situation. We investigate how participants perceive the overall conversation as well as the individual visualizations. We find that the visualizations add helpful information and that participants found that they were not distracted by these additions.

In summary, we contribute: (1) a design exploration of body-based visualizations, and (2) an evaluation of how such visualizations impact actual conversations.

2 Related Work

Our work aims to appropriate the bodies of others for visualization during social interaction. As such, we build upon prior work on body-based visualizations as well as technology support for conversational situations. Given that we focus on supporting participants in conversations individually, this support has to be subtle and not disruptive to the situation. Hence, we also draw upon the subtle interaction literature.

2.1 Body-Based Visualization

Given the limited size of mobile devices, the body has attracted interest as a place where to show output. For example, *ScatterWatch* [40] and *LumiWatch* [58] do so for watches with scattered and projected light respectively. Instead of using the skin for output, Schneegass et al. [50] investigated clothing with embedded displays and which locations of these participants would prefer. Similarly, Harrison and Faste [17] studied acceptability of projected on-body interfaces. This included an investigation into how participants would feel about others looking at an interface projected on the participant's body. They identified concerns around steering other's gaze to sensitive areas as well as discomfort with having others use their body as an interface, particularly as this involved touch in their scenario.

Body-based visualizations are not only used for user interfaces and have also seen application in areas such as health and medicine. For example, Hoang et al. [20] investigated projected on-body visualizations in physiotherapy education. Similar display of anatomical information on the body also was used in *BodyVis* [35] and for medical images in *ProjectDR* [55]. Visualizations tied to bodies are also being developed for sports viewing [29], such as to highlight offense trajectories or key players in a defense scheme.

With AR technology, body-based visualizations can also extend to the space around the body, such as with the *Body LayARs* toolkit [39]. As demonstrated by the *RealityTalk* system [27], such visualizations can enhance storytelling and presentations. RealityEffects [28] generalizes this to object-centric visual effects, also describing a corresponding taxonomy based on online videos featuring such effects. Bhatia et al. [5] showed that everyday activities, such as eating and running, can also be enhanced through AR visualizations. Other times, AR on bodies means augmenting people's faces as demonstrated, for example, by the *HoloFace* [23] and *Makeup Lamps* [3] projects. Face-changing AR can be a medium for self-expression or performance, but can also be used to alter how we see the people around us. This was explored by Leong et al. [26], who used AR face filters to alleviate public speaking anxiety in an online presentation scenario. Another use of body-based AR is to show information on the people we meet and ourselves. As Rixen et al. [45] showed, not all such information displays are acceptable to users, in particular non-self-disclosed information of more private nature.

2.2 Supporting Conversational Situations

There have been a range of technologies developed to aid in conversational situations, from wearables to purely software solutions. For example, *EMOTE* [11] is a speculative design for AR visualizations that help users better express their emotions, such as by signifying sadness with virtual rain over ones head. A concrete system is *Cardiolens* [31], which senses and overlays heart rate information on oneself and others and thus, for example, can augment social interactions.

Other systems focus on accessibility support, such as for blind and autistic users. An example of the former is *Expression* [1], which tracks the person the user talks to, infers social signals (e.g., smiling and yawning), and relates these to the user. A similar system was also presented by Tanveer and Hoque [52], who specifically targeted small talk in groups and relating information on the group composition. Williams et al. [57] developed a smart glasses prototype that provides glanceable vocabulary cues to support people with aphasia in conversations. Furthermore, *Wearable Subtitles* by Olwal et al. [38] are a system that helps users who are deaf or hard-of-hearing by transcribing what is spoken around them. Support for people with autism also often builds on facial emotion recognition systems, such as through Washington et al.'s [54] *SuperpowerGlass*. A different case is the *SayWAT* [7] system, which helps users better control the volume and pitch of their voice in conversations.

Within conversational support, there has been continued interest in aiding conversations with strangers. One aspect there is the topic of the conversation and how to help strangers find something to talk about. For example, Nguyen et al. [33] showed participants personalized topic suggestions in their glasses and found that introverted participants in particular benefited from them. Similarly, Ogawa and Maes [37] identified commonalities between people to then show topic suggestions on their watches, with promising initial results. Instead of topic recommendations, another approach is to give people tools to better present what they are interested in to strangers. An example of this is work by Kytö and McGookin [24] around *Digital Selfs*, which facilitated better interactions between strangers, such as by acting as an ice-breaking conversation starter. Changes to the environment can also aid conversations, such as with *SocialStools* [15], which combine movement and projection to foster togetherness between strangers.

Conversation starters and self-presentation are not just relevant for conversations with strangers, but play into conversations in general. For face filter technology, Javornik et al. [22] showed that initiating social interaction is a core use case and that altering ones appearance facilitates fun interactions. As Hetting et al. [18] described, virtual changes to ones appearance—through animated avatars in their case—can also alter behavior through a kind of Proteus effect. Furthermore, as Noh et al. [34] found, hiding behind face filters in remote counseling sessions increased users' level of self-disclosure.

2.3 Subtle Interaction During Conversations

Another facet of interaction with computers during conversations is to design for subtle interactions [41], so that focus on the conversation itself can be more easily maintained. For example, Schiavo et al. [49] built a system to balance participation in a group discussion through subtle nudging, as not to distract from the conversation. Not to disrupt the interaction with one another, even though one engages with technology, is a common goal in many subtle interaction systems. Another example of this are *Deceptive Devices* [2], which allow for interaction that others cannot easily observe. This includes a mug with a display inside, which—when tested in a meeting situation—effectively hid the fact the user was reading from others.

Distraction is a critical aspect of subtle augmentation and one question here is how to show information without overburdening the user. *StARe* by Rivu et al. [44], for example, showed how gaze can be used to control which information overlays are displayed during a conversational situation. Similarly, Ofek at al. [36] studied how much and in which ways information can be taken in during a conversation while not appearing disengaged. Rzayev et al. [47] evaluated different placements of AR notifications during conversations, including how much they distract from the other person. The feasibility of glanceable information in AR was also investigated by Davari et al. [12], who found that reducing the intrusiveness of virtual content when in a conversation made it easier for participants to notice social cues.

Janaka et al. [21] as well as Cai et al. [8] have investigated paracentral (i.e., just outside of central vision, but not quite peripheral vision yet) visualizations for subtle interactions in conversational situations. Both show information around a person the user is speaking to, with the latter project also including a means to control the system. The results from video and in-person studies show that the used visualizations and interactions did not have a negative impact. For example, users could maintain eye contact with each other, attention to each other was not strained, and they rated the interaction as polite and natural. In contrast, use of a phone or issuing voice commands in the same situation was found to be much less polite.

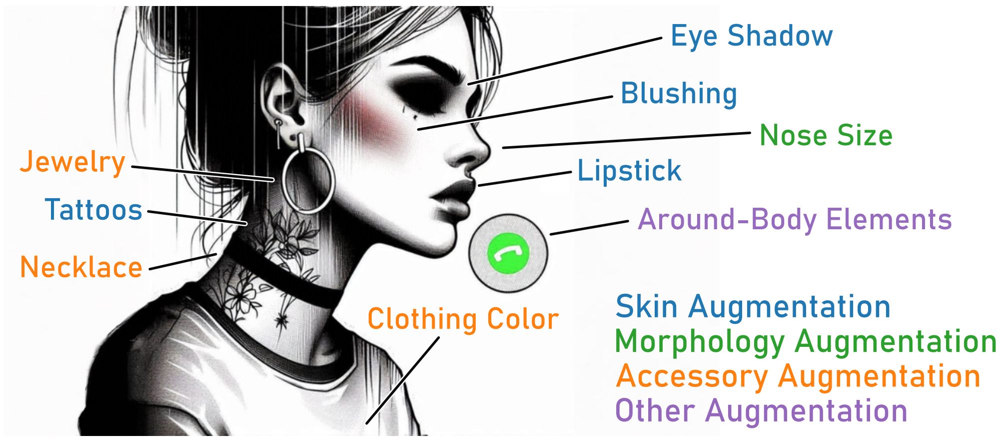

Fig. 2. Body-based visualizations apply augmentations to the skin, morphology, and clothing/accessories, as illustrated here through several examples. Visual elements can also attach as around-body augmentations.

3 Types of Body-Based Visualization

While we have briefly described previous body-based visualization work above, we here take a closer look at such visualizations, particularly the potential mappings of visualization parameters to bodies. To begin with, a body is different from a piece of paper or a screen. Hence, the elements a body-based visualization should consist of also need to differ, are they to fit the target "medium". We take inspiration from Bertin's [4] work on visual variables as we explore how such visual variables can be applied in body-based visualizations. This covers visual elements appearing on and around the body, with a main focus on the face. An alternative design space for situated visualization in AR was described by Lee et al. [25]. While our focus is on visualization directly on the body, the three-dimensional nature of the body lends itself to a broader definition that includes this adjacent space as well. Bertin's variables of size, intensity, granulation, orientation, color, and form all can be applied to marks used in body-based visualizations. For example, a stripe pattern shown on the cheek can vary in these dimensions just as it would on more traditional media. Here, the body acts as a canvas for visualization, with the structure of the body influencing the placement and design of the visuals. However, some forms of higher level mappings in body-based visualization more are directly tied to the body. As illustrated in Figure 2, body-based visualizations add skin, morphology, and clothing/accessory augmentations.

With morphology augmentation we describe visualizations that adapt the structure of the body, treating it as a visual variable on its own. Noteworthy in this respect are *Chernoff faces* [9], where data is represented through variations in facial features of a cartoon character. For example, the overall shape of the face, the way the eyes are drawn, or the width of a smile can all encode data. However, as we want to visualize on real bodies, many of these mappings are not possible or sensible. Yet, some could be incorporated into body-based visualizations through mimicry of the effects of prosthetics or cosmetic surgery. For example, AR face filters already can distort the face to "beautify" it, such as by shrinking a nose. Similarly, visualizations could change the shape of the eyes, add or remove wrinkles, or make the neck fatter. While this could be used to encode data (e.g., mapping a continuous variable to earlobe length), these also are mappings that are not easy to decipher. As Chernoff pointed out, his faces work well for seeing patterns, but useful information might get lost in the noise. Furthermore, visual alterations to the body shape are also rather intrusive and potentially disrespectful to others.

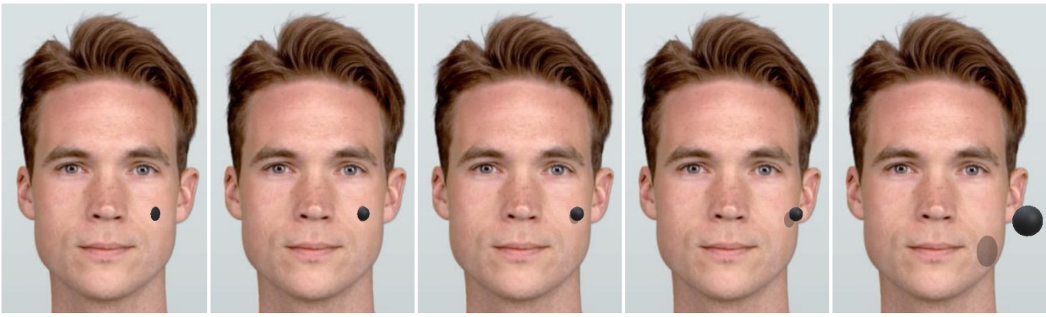

Fig. 3. One dimension of body-based visualization elements is how much they are attached or detached from the body. From left to right, this figure shows a small black mark on the surface of the face, sticking out slightly from the surface, sitting right on the surface, hovering above the surface, and floating off the face.

Fig. 4. Virtual makeup and tattoos allow for a range of visualization. Based on the original view on the left, the variant in the middle darkens the eyeshadows, adds blush, and extends a real tattoo with a virtual one. The right variant also adds a fully virtual name tattoo and colors the eyeshadows.

While we mostly treat on-body and around-body visualization separately, there is not necessarily a binary distinction between an element being fully 2D and fully 3D. Instead, another approach is to differentiate visual elements by how attached or detached they are from the body. For example, there can be a continuous transition from a circle on the skin, via a bulge on the skin, a sphere stuck on the skin, a sphere floating above the skin, to a sphere floating around the body at a distance (see Figure 3). Fundamentally, detached visual elements next to the body can be anything: 3D models, 2D callouts, emoji sprites, or text labels. It is their connection to the body that can make them part of a body-based visualization. This bodily connection can be a graphical element (e.g., a connecting line) or just aligned movement (e.g., tracking a person's hand). For example, we consider a callout showing a person's name and connected to their head (i.e., following them around and annotating them) as a body-based visualization.

Skin augmentations can tap into the makeup design space and apply virtual makeup (illustrated in Figure 4) to show data. Types of makeup include eyeliner, rouge, and highlighter, with each accentuating different parts of the face. The application of makeup can range from subtle to overt. Body-based visualizations can make use of this gamut of expressions as well. For example, information could be shown through virtual makeup making cheeks redder, changing the color of the eye shadow, or thickening the eyelashes.

Related to makeup is permanent pigmentation of the skin through tattoos. In fact, the latter is also used for permanent makeup applications. Tattoos allow for a wide range of visual additions to the skin and are applied to all areas of the body. This does include the face, even though that application area is not very common. Body-based visualization can build upon tattoos (also illustrated in Figure 4) to visualize information, either by adapting their style independently or by adding to existing tattoos.

A more speculative form of skin augmentation is to mimic material changes of the skin. One example of this is wetness, such as when standing in rain or sweating, which shows as reflectance change of the skin, but also in water beads on the skin. Another example are goose bumps, which show as erect body hairs and bumpy skin at the base of these hairs. Mimicry of natural on-skin phenomena could also include rashes, scaring, or wrinkles. However, such virtual changes of the skin's materiality are currently not well supported by AR technology, For example, tracking and rendering resolutions need to be higher to support small-scale features, such as goose bumps. Hence, while we mention these here for completeness, we do not pursue such changes further in this paper and leave this for future work.

Finally, much of the body is usually covered in clothing which can also be augmented. This includes virtually changing the color of clothes (such as envisioned by Rixen et al. [46]), their texture (e.g., with plaid), or adorning them with marks (e.g., in the form of patches). Similarly, accessories such as jewelry lend themselves to augmentation, such as with virtual extensions of their design. For example, Rantala et al. [42] explored jewelry with built-in markers and digital overlays. Furthermore, clothing and accessories can also be fully virtual and in that form allow for a wide spectrum of visualizations, be it through shape, color, movement, or other attributes. For example, brands like *DRESSX*[1] demonstrate the expressivity of such digital-only fashion.

Together, these body-centric augmentations provide a rich vocabulary for expression. Within each, visual variables can be mapped to data, such as by changing the size and orientation of a tattoo, or the color of virtual make-up. Furthermore, in additions to such mappings, body-based visualizations can also include textual or iconographic elements directly. For example, a tattoo can encode a numeric value in its size, but also show the literal number embedded within it.

4 Prototyping Body-based Visualizations

As we set out to evaluate the impact of body-based visualizations on conversational situations, we needed a set of working prototypes of such visualizations. Hence, we picked five pieces of information relevant to conversational settings and instantiated them with the visual elements described in Section 3. Specifically, we cover (1) name information on, and (2) status of the person one talks to, notifications on (3) weather, and (4) messaging events, as well as (5) timing information for the conversation. While the notification data is not directly pertaining to the conversation, it represents a common interruption thereof, and hence a potential disrupting factor in such situations.

We used Snap's *Lens Studio*[2] (version 4.55.0) to prototype body-based visualizations. While there are some approaches from research for live AR filter effects [23] and body-based AR experiences [39], these do not offer the fidelity and depth of commercial solutions. The *OpenFilter* [43] framework also has bridged the gap between these commercial offerings and research, however it does not allow for live use of filters. Implementing our prototypes within Lens Studio and then using them directly in Snapchat allows us to test their effects in live conversations within video calls. We release[3] the source code of the prototypes and related assets to enable follow-up work.

[1] https://dressx.com/
[2] https://ar.snap.com/lens-studio
[3] https://github.com/henningpohl/Body-Based-Visualization

Lens Studio provides tooling for a wide range of body-based effects. The underlying software stack has components for world, body, hand, face, eye, head, and expression tracking. Furthermore, custom segmentation and classification machine learning models can also be integrated. Tracked bodies can be augmented with 2D and 3D elements in free space as well as anchored to specific locations on the body. Faces can also be warped, stretched, textured, retouched, and adorned with sprites. This allows for implementation of many of the potential body-based visualizations described in Section 3.

We built a custom Snapchat lens to showcase several examples of body-based visualizations within a conversational scenario and allow for evaluation of their overall impact on conversations. For selecting what information to visualize, we draw on previous work on mobile notifications [48] as well as on the conversational setting itself. We cross this with the range of potential body-based visualizations for a set of five visualizations to implement. For this we did not systematically match types of information to visual elements, as that set is too large in scope, but followed an iterative design approach. We considered the larger space of potential mappings and then balanced which ones to use so different visualization types are adequately represented. All visualizations appear on and around the body of the person the user is speaking to, so users do not need to look away from them to see the visualizations.

The visual elements used in our custom lens are shown in Figure 1 and are: (1) a virtual name tag that is affixed to the chest, (2) an emotion tracker that determines how much the conversation partner enjoys the chat and shows this data in a bar chart, (3) a weather widget that appears above the conversation partner's shoulder and provides notifications on incoming clouds and rain, (4) a message widget that provides notifications on incoming messages through a chat bubble emerging from the side of the conversation partner's head, and (5) a conversation timer that shows how many seconds and minutes the user has already talked with the conversation partner. We describe each of these visual elements in more detail below. In terms of types of body-based visualization, these elements cover objects around a person, objects closely attached to a person, and objects directly affixed to a person, as well as virtual tattoos and skin color alterations.

4.1 Conversation Partner Information

In many social situations and casual encounters, the people talking do not know each other beforehand. For better conversations, a fundamental piece of information to know about the other person hence is their name. Consequently, events like conferences or corporate meetings often use name tags to make it easier for strangers to interact. Other information relevant in this context would be affiliations, roles, contact information, or information related to previous encounters. We decided to adopt the basic name tag in a virtual version that then adorns the person one talks to.

For this purpose we mimic the physical name tag in a classic design (i.e., "Hello, my name is…"). We pin the name tag to the upper body to further imitate the wearing of a physical one. It thus follows the person's movement and becomes a part of their outfit.

4.2 Conversation Partner State

Engaging in a conversation also means reading the other person to adjust the interaction. For example, one might try to get them to laugh, to avoid angering them, or to keep a relaxed mood. People are generally good at inferring other's emotional state and reacting to social cues. Furthermore, facial expressions have robust and stable emotional interpretations across cultures [14]. However, this can be a challenge for some, such as people on the autism spectrum. Providing visual cues of others' emotional state could help in those situations and we investigate one particular instance of this: a mood tracker.

Our mood tracker utilized facial expression tracking, which provides weights for a set of facial muscle movements. Specifically, we track the `MouthSmileLeft`, `MouthSmileRight`, `MouthFrownLeft`, and `MouthFrownRight` weights to determine an overall smile score. While not all smiles signify enjoyment [13], positive emotion correlates with smiling [13] and is reliably interpreted by others [32]. Alternatively, one could also leverage AI technology to derive emotion labels from images of a person's face. We compute a running average of the last 50 frames to derive an overall emotional state estimate: the more someone smiled in that time, the higher the score. This is then mapped to a 0–1 range and visualized on a tattoo-like bar graph element on the forehead above the left eye. A tattoo-like appearance is emulated by a multiply blend of dark colors with the skin. The visual element always shows a smiling and frowning face, with more bars filling up towards the former as the determined smile score increases.

4.3 Weather Notifications

There are several scenarios where people might need access to incoming information while engaged in a conversation, even though that information is not pertaining to that conversation directly. For example, based on incoming public transport information, traffic alerts, or breaking news one might want to cut a conversation short and leave early. Within this space, we focus on weather notifications where users might want to react to an incoming rain shower or raised tornado alert.

We utilize the space around the body to manifest weather notifications as iconic representations. As clouds are moving in, a three-dimensional cartoon cloud moves into a position over the conversation partner's shoulder. This is accompanied by a textual warning stating that there are "Clouds incoming". The cloud can also change appearance to a dark and stormy version that also shoots lightning bolts. In this state, the text also changes to "Caution: Rain!" and virtual rain pours down from the cloud.

4.4 Messaging Notifications

Another common form of notifications are those we receive from our social media and messaging apps. Here notifications represent new messages, posts, or corresponding events such as when someone else liked a post or message. There is usually no direct connection of these kinds of information to a conversation one is in, but one might want to keep track of those notifications anyway. This results in a conflict between attention to the conversation and access to incoming information. We focus on notifications for incoming messages and specifically just that there was a new message coming in, not its content.

We picked a common text bubble icon with message count overlay to show these message notifications. To tie this more closely to the body of the person one is interacting with, we make this icon emerge from them. In a playful fashion, a pole pops out from their right ear, with the bubble icon then unfolding from that pole. This was also meant to mimic signal flags on mailboxes as well as popular magic tricks where items are pulled out of and from behind others' ears.

4.5 Conversation Timer

As our conversational situations were scheduled and within time constraints, the time passed during them is information that is worth keeping track of. In everyday life, users might want to limit their time in a given conversation due to a desire to talk with many people at an event, not to appear too engaged, or due to a tight schedule. Similarly, they might desire the opposite, to spend at least a minimum amount of time in a conversation. In any case, currently keeping track of the time requires a glance at a watch, a phone, or a wall clock. As these options all potentially alert the other person that one is checking the time, there can be a social cost to them.

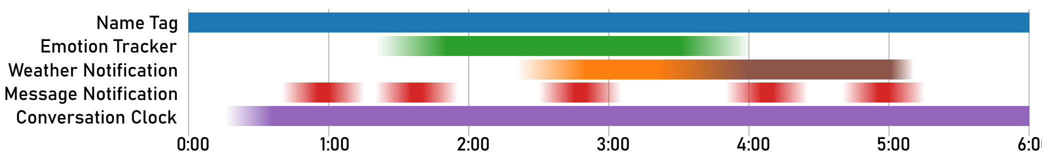

Fig. 5. Five different visualizations are present over the duration of the study. The Name tag was always visible, while other elements appear randomly within predefined intervals. We use gradients to show the presence density. The weather notification at some point switches from a cloud warning to a rain warning, similarly shown as a gradual shift in likelihood.

As an alternative, we implement a body-based timer with visual elements around and above the right eye. We use translucent materials that blend with the body to make these elements appear as light on the skin or painted on colors. Specifically, we have a green annulus element around the eye that fills up in a clockwise fashion with every passing second. Once full, it resets and a blue minute mark is added across the eyebrow line. Thus, users can see how much time has passed by the number of these blue checks and the degree to which the green annulus is filled. In contrast to, say, a textual element displaying the time, we found this kind of visual augmentation blends in more naturally with the face.

5 Evaluation

We evaluated the impact of body-based visualization on conversations with an online study. Our aim was to create a casual conversational setting that closely mimics in-situ conversations, while still allowing us control over the visualizations. As we built our prototype as a Snapchat lens, the study had to run online and we had each participant join a video chat with a confederate on the Snapchat app. Only the confederate was augmented and participants assumed the role of the user of a system that shows information for them on the bodies of others. Their task was just to have a conversation and we subsequently asked them for their opinions via a post-call questionnaire. That questionnaire focused on the value they saw in the body-based visualizations as well as their perception of the visualizations' impact on the conversation. In particular, we wanted to find out whether the felt they were being distracted and whether their attention suffered.

5.1 Apparatus

We used *Snapchat* (version 12.53.0.46) as the platform for this experiment. The confederate used the Snapchat app on a phone (Pixel 4a running Android 13), while participants could join from either their phone, tablet, or a computer (via *Snapchat for Web*). We mounted the phone on a holder and used a ring light for consistent framing and illumination. In the confederate's app we activated the custom Snapchat lens described in Section 4. The consistent framing ensured that five implemented visual elements would be in view during the conversation.

To reduce visual overload, not all visual elements were active the whole time. Instead, we used a mix of constant and intermittent elements that change over the duration of the conversation. Figure 5 shows an overview of which elements were active when. As can be seen, the name tag was always visible and the conversation clock remained visible after it appeared early in the conversation. We randomized appearance and disappearance of some elements in pre-defined intervals to create more variation. Thus the emotion tracker was only active in the middle of the conversation and the weather notifications appear later on, subsequently changing from a cloud warning to a rain warning. Message notifications appeared in five defined intervals during the conversation for a fixed length.

5.2 Procedure

Participants were first redirected to a meeting booking website to schedule a video call with us and provide informed consent alongside. To make the calls somewhat comparable, we asked participants to pick one out of four topics of conversation for that call: food, shows/movies, the weather, and video games. However, if they did not like any of the topics, we also allowed them to pick one themselves. They could also provide a nickname for us to address them with during the call.

We then sent participants a message (see Appendix A) with further instructions as well as information on the visual elements they would see in the call. This included textual descriptions of the elements as well as an image (see Figure 1) showing what these elements look like. We also informed the participants that the notifications they would see are simulated and that they should imagine that this information would be generated by their own device and thus be personal to them. We also told them that the person they would talk to would not be able to see these visual elements and that the elements are personal to them, even though they are shown on and around someone else's body.

The calls themselves were scheduled to be brief chats, not longer conversations. Hence we told participants that this chat would be for 5–10 minutes. At the start of each call we first confirmed they had read the information material we sent them and understood the task. If they failed to read the instructions beforehand, we went over the instructions together with them. We then activated our study lens, which started in a paused state, only showing the name tag. After 20 seconds the lens then commenced along the timeline described earlier and shown in Figure 5. This was timed so there would be a short window to transition the call to the casual conversation, the start of which then coinciding with the start of our lens timeline. The confederate and the participant then had a casual conversation on the participant's chosen topic. During this stage, no comments on or acknowledgment of the visual elements was made from our side, nor did we have participants remarking on them. All conversations went over the five minute minimum and were ended by the confederate at a suitable later moment. We then thanked participants and let them know we would send them a link to the post-chat questionnaire afterwards. After the call had ended, participants could then fill out that questionnaire on their own. The calls themselves were not recorded, as the Snapchat app does not allow for this, so only the questionnaire responses informed our subsequent analysis. Overall, each participant thus spend about 20 minutes on this study.

5.3 Participants

We recruited 16 participants (9 female, 7 male, age 20–52, M=29.2, SD=8.6) via the Prolific crowd-sourcing platform. Participants needed to be fluent in English and be willing to participate in a video call to be eligible for this study. We did not filter participants by whether they already use Snapchat (as that might bias their perception of the used visualizations) and instead just informed them they would need to have or create a Snapchat account. Most of our participants were from South Africa (12), with further individual participants from Poland, Portugal, Spain, and the UK. We paid participants 5 GBP for their participation.

5.4 Measures

In our questionnaire (see Table 1), participants first rated each of the five visual elements with respect to likability, usefulness, and distraction. We then asked them for overall ratings of the conversation, as well as the integration of information with it. Finally, participants answered a set of open ended questions. All ratings were on 7-point Likert scales, repeated for each visual element where a question asks about those. Except for the final question for other comments, participants were required to answer all questions.

Table 1. Items of the post-call questionnaire. We denote 7-point Likert scales with † and open ended questions with ‡. Questions on specific elements are repeated for each.

Question	
I liked the <*element*>	†
I think the <*element*> was useful	†
The <*element*> was distracting	†
I was able to give enough attention to the conversation	†
I was easily distracted by the incoming notifications	†
The person I talked to probably felt ignored	†
It felt wrong to have information displayed on their body	†
The information I got could be helpful in that moment	†
How would you describe the effect of those visualizations on the conversation from your perspective?	‡
What do you think would be information that could be helpful for you to have in a conversation?	‡
What would, in your opinion, be the best way to visualize the information you got during the conversation?	‡
Do you have any other comments for us?	‡

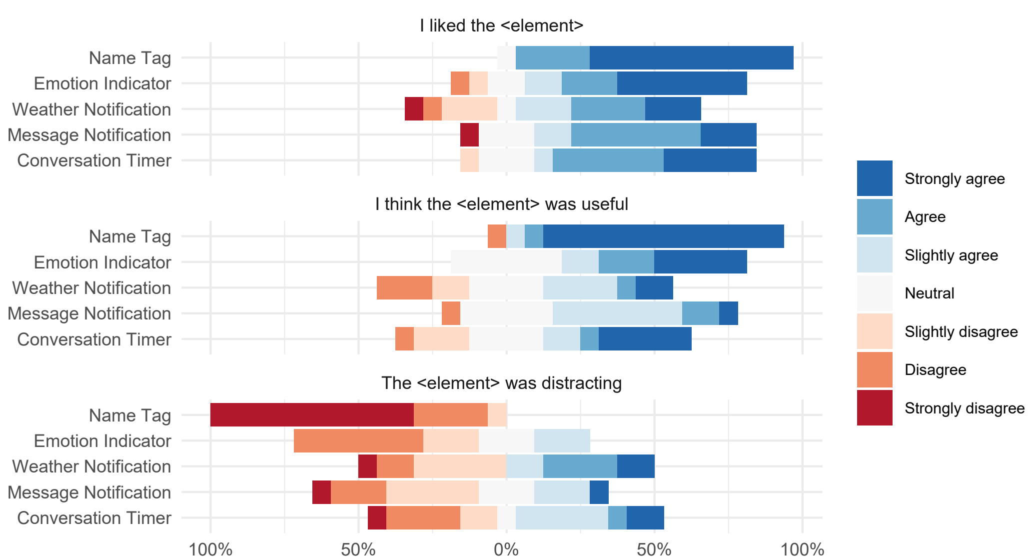

Fig. 6. Participants rated each visual element along likability, usefulness, and distraction dimensions. Overall, most elements were liked and found useful with more mixed responses for distraction.

6 Results

We show the raw response data (see Figures 6 & 7) for all questions and supplement this with statistical analyses. The data from this study is available in the earlier referenced repository as well as as supplemental material to this paper. As all ratings were on 7-point Likert scales, we use non-parametric tests for this analysis. Specifically, we use Wilcoxon signed-rank tests to analyze whether responses to a given question differ from a neutral response, and cumulative link mixed models to analyze differences between ratings across all visual elements. We used an alpha level of 0.05 for all statistical tests. Finally, we analyze the responses to the open-ended questions.

Table 2. Results of Wilcoxon signed-rank tests on all per-element responses. The null hypothesis used in these tests is that the true response is neutral.

Question	Element	V statistic	p-value	
Liked	Name Tag	120	<0.001	***
Liked	Emotion Indicator	98	0.004	**
Liked	Weather Notification	87	0.127	
Liked	Message Notification	80	0.017	*
Liked	Conversation Timer	90	0.002	**
Useful	Name Tag	134	<0.001	***
Useful	Emotion Indicator	55	0.005	**
Useful	Weather Notification	46	0.632	
Useful	Message Notification	57	0.031	*
Useful	Conversation Timer	62	0.067	
Distracting	Name Tag	0	<0.001	***
Distracting	Emotion Indicator	10	0.013	*
Distracting	Weather Notification	80	0.546	
Distracting	Message Notification	26	0.172	
Distracting	Conversation Timer	58	0.931	

Note: *p<0.05; **p<0.01; ***p<0.001

We first analyze the per-element ratings. As shown in Figure 6, our participants overall liked all elements and generally also found them useful. Wilcoxon signed-rank tests show that several ratings differ significantly from a neutral response (see Table 2). The name tag, emotion indicator, and message notifications were rated as significantly higher in usefulness. While elements were overall liked, one participant strongly disliked the weather and message notifications. Specifically, the name tag, emotion indicator, message notification, and conversation timer were significantly liked, but not the weather notification. In terms of distraction, the results were more mixed and responses also varied more strongly between the elements. Only the name tag and emotion indicator were rated as significantly less distracting with opinions on all other elements split. The name tag is noteworthy here as nobody found it distracting, while the weather notification and conversation timer were found the most distracting (albeit non-significantly). Separate cumulative link mixed models fitted per question with participant as random term showed significant differences between the elements with respect to being liked ($\chi^2(4) = 16.9, p = 0.002$), useful ($\chi^2(4) = 26.2, p < 0.0001$), and distracting ($\chi^2(4) = 53.6, p < 0.0001$).

While some participants found individual visual elements distracting, their overall impression (see Figure 7) was one of low distraction. This manifests in their perceived high ability to maintain attention to the conversation, and disagreement with being easily distracted by the incoming notifications. Furthermore, the participants strongly disagreed that the person they talked to probably felt ignored, indicating that they did not feel they themselves behaved in a way that would lead to such a feeling on the other side. We ran one-sample Wilcoxon signed-rank tests for each question to check whether the responses significantly differed from neutral. We found significant differences for all questions, with attention ($V = 136, p = 0.0003$) and helpfulness ($V = 102, p = 0.02$) rated better than neutral and distraction ($V = 17, p = 0.008$), wrongness ($V = 28, p = 0.040$) and ignoring ($V = 0, p = 0.0005$) rated lower than neutral.

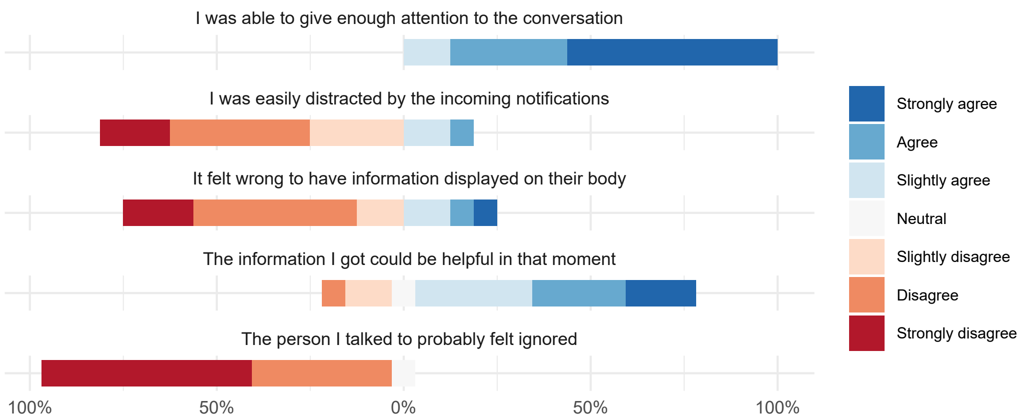

Fig. 7. Participants overall did not see the body-based visualizations as distracting, morally wrong, or negatively influencing the conversation they had. Instead, they found they could keep their attention on the conversation while receiving helpful information.

The information conveyed by the visual elements itself was overall deemed potentially helpful and participants thus did see value in the proposed approach. We also checked whether they would have ethical or moral issues with overlaying information on the body of another (i.e., in a sense appropriating another's body for their own benefit). Yet, this notion was rejected by the majority of participants who did not see such an issue with the proposed body-based visualizations.

Asked to describe the effects of the visual elements on the conversations, participants mostly pointed to their helpful properties. For example, the name tag was mentioned positively several times, including because with it *"you don't have to make the person feel bad if you forgot their name"* (P13) and as it *"could decrease cognitive load"* (P15). Similarly, participants saw the utility of the emotion tracker, noting that *"it's good to know how the other person is feeling"* (P12) and how it makes one *"aware if the person is enjoying the conversation or getting bored"* (P13). One participant also made an overall statement along those lines, noting that the visualizations *"gave me a clear indication of how the person I was engaging with felt"* (P14) and another stated that the visualizations *"were quite useful especially the name tag and the emotion tracker"* (P16). The conversation timer was also mentioned, but while one participant felt it *"helped with keeping our time"* (P12), another thought that *"the more distracting thing was the timer on the eye"* (P1).

Distraction, in fact, was mentioned in many comments with no clear consensus on whether the visualizations were or were not distracting. For example, one participant stated that *"the weather notification was a bit annoying as it started raining on the person's face it was rather distracting"* (P13), which is echoed by another participant who pointed out *"they were a bit distractive but I was able to keep up with the conversation"* (P6). On the other hand, we got a comment that *"[the visualizations] were not that distracting"* (P3) and another that they *"were rather helpful/informative and I never felt any distractions"* (P11). Some participants remarked on the interaction between distraction and the quality of the conversations. For example, one noted that *"since the conversation was good, they were not a distraction, but more of an added benefit to the conversation just to know the surroundings of the person and the mood they are in"* (P7), which is also echoed in another comment pointing out the visualizations *"were unnoticeable for the most part and maybe that could have been due to the conversation we were having"* (P4). Finally, understanding of the visualizations also matters as one participant noted: *"To an extent it was very disturbing more so as I didn't really understand* *what each meant but with time I understand and was able to ignore them and continue to have a conversation"* (P5). Not understanding a visualization can lead to more distraction, as visible with one participant who commented: *" I'm not sure what the worm crawling on the person's face was for but that was also a bit distracting"* (P13). In this case, while the study instructions described that this was a conversation timer and how it works, the participant did not recognize this element and thus misread the green circle around the eye.

Asked what information they think would be helpful to have during a conversation, many participants referred back to the visualization they just experienced. In particular, the name of the person one talks to, their emotional state, message notifications, and the time passed were mentioned. For example, one participant noted the timer was helpful as *"if you have another appointment you won't be late or maybe you have a limited amount of data so you let the person know instead of getting cut off"* (P13), while another said *"the conversation clock is helpful to keep up with what I have to say, and finish"* (P6). On the other hand, the emotion tracker also was helpful *"so you can change the topic if the conversation is going down hill"* (P13). The latter point was also commented on by others who noted *"[the emotion tracker] is a nice effect because it let's you know when the conversation is going good or not"* (P7), going as far as thinking this technology would enable one to *"know how they truly feel"* (P16). With respect to the weather notifications, one participant pointed out they *"like the weather effect, if only it showed they actually [sic] temperature of the place they are currently in"* (P7). But participant also mentioned information to visualize that were not part of the evaluation. For example, one would like to *"[know] what [the other person] thinks of every subject I talk about"* (P1) and another desired *"very brief meeting agenda set as a transparent background text"* (P15), which would then serve as *"a tick-able checklist"* (P15). Thinking of remote rather than in-person conversations, other participants were interested in information on the location of the conversation partner as well as the time difference and distance to them. However, one participant also had a more negative take here, noting that *"Genuinely I do not think I would want to have a conversation like that again because with the changes it makes it can get distracting and the conversation could not reach it's full potential based on my partial distraction"* (P5).

Asked what they thought would be the best way to visualize the information they got during the conversation, many commented they would rather like the information on the side instead of on the face. However, a few had a more positive take on this, such as noting: *"How it was, was perfect. I just would change the position of timer"* (P1). Similarly another comment was that *"The idea of displaying it the way it was displayed is pretty much good and might just need to be tweaked a bit for best visual image"* (P11). The short duration of the call and thus the condensed presentation of the visualization was seen as problematic by one participant, who stated that *"I would have liked if it was a pop up that happened 1 by 1 and not everything at once"* (P7). While we are ultimately interested in real-life conversations, the current format was a video call and this did bias participants in their responses somewhat. For example, the desire to have things on the side commonly reference the screen and its edges. We also had one participant state the best way to visualize would be *"as a part of reality and not as distracting filters"* (P3).

Half the participants made use of the final open question to provide further comments. These were all positive, noting, for example, that *"It was a great experience. I would like to see some of these tools available in the future"* (P13) and that *"The timer clock is very much useful. I find it innovative"* (P6). Two also commented again on distraction, noting *"This was an interesting study and I learnt that I'm not easily distracted"* (P10) and *"I liked the effects on the conversation, the effects were not distracting at all"* (P7).

7 Discussion

Our participants felt they were able to keep their attention on the conversation and did not get distracted by the visual elements. They did not feel the person they talked to would have a reason to feel ignored, further supporting the notion that, at least in their mind, they devoted considerable attention to that person. Furthermore, the participants saw the potential of the used visual elements and augmentation of conversational situations in general. This overall supports our initial assumption that body-based visualizations would be a suitable approach for providing information in conversational settings. Gaze plays an important role in conversations (see, e.g., [6, 19]) and information access can be disruptive [53] and have people miss social cues [12]. Moving information closer to the people we converse with counters these issues, as also indicated by our participants' attention and distraction responses.

The most relevant work to compare our results to is Janaka et al.'s [21] paper on paracentral and near-peripheral visualizations. In particular they also evaluated a body-anchored circular visualization element, but in contrast to our eye clock, theirs was enveloping the whole head instead of appearing on the face. The associated degree of distraction in their first study was low, in line with our results. In a second study, they also evaluated this element in an in-person conversation, where some wandering gazes were noticed, but overall no lack of eye-contact was felt by the participant around whom the visualization was shown. In general, our visualizations were more invasive as several were shown directly on the face and many were animated. Yet, our participants did not report levels of distraction or effects on the conversation out of line with the ones reported by Janaka et al.

Building on Janaka et al.'s work, *ParaGlassMenu* [8] by Cai et al. evaluated interactive elements placed in the same paracentral space. The menus here are all circular arrangements around the head of a person the user is conversing with. Their second study evaluated this again in an in-person conversational setting, finding no negative effect of interacting with such a menu on the quality of the conversation. As Cai et al. note, some participants actually remarked that this way of interacting helped increase their engagement in the conversation compared to past experiences. In our study we only used the visual elements for feedback so it is not clear whether interaction with them would similarly not have impacted the conversations. The effects of their presence itself though, are in line with Cai et al.'s findings as well.

Finally, our participants assumed a role equal to the receiver in Rzayev et al.'s [47] study on notification placements during social interaction. They found that notifications that were aligned with the face of the other person were less intrusive than notifications locked to their view. Overall, their receivers reported medium intrusiveness scores, which is partly comparable to our participants' distraction scores of similar magnitude.

Overall, our results point to the suitability of body-based visualizations in conversational situations. While no element was rated as significantly more distracting than neutral, we still observe that distraction ratings vary between them. Most importantly, the name tag was rated much less distracting than the other elements, which is likely because it also is the only completely static element. On the other hand, the state transitions of the two notification elements contained noticeable movement, which could explain why they were rated higher in distraction. While these animations make change more apparent and the visualizations more playful, when optimizing for subtleness it likely is a good idea to reduce or even eliminate such movement. Participants still liked those visualizations, though, so some settings might also benefit from leaning more into the playfulness and expressivity of body-based visualizations. For example, we tested their effect in a one-on-one conversation online, but during a party chat with several people more wandering gaze might well be more acceptable.

7.1 Ethical Considerations

Is it problematic to display information on and around others? While the user can see this information, the person it is shown on does not, so if they cannot tell, is there even an issue here? We believe there can be, as there is a loss of control over self-presentation, and instead another person is manipulating how we appear. For considerations how to view this situation we can turn to other technologies where this kind of control loss occurs. Photo editing is one such area where photographs and editors commonly change how people appear. What kinds and how much editing is ethical varies, but commentary and guidelines on this have developed [56]. A more stark case is the recent emergence of "deepfakes", where a person's face is superimposed on another video or image, or content is just fully generated using artificial intelligence. In either case, deepfakes are often non-consensual and used in inappropriate and unethical ways, such as described by Story and Jenkins [51]. In both the above cases, other people can see and broadly distribute the altered or generated imagery. On the other hand, the visualizations we describe here are private and ephemeral and as such also pose a less severe threat. Our participants overall did not see such information display as inherently "wrong", but they also took on the role of recipients. How people would feel after learning that others use them as part of information displays remains to be more closely investigated. We would expect this to depend on a multitude of factors, such as relational distance, format and style of visualization, whether such visualizations can be recorded, and the quality of their interaction with the user.

7.2 Limitations

We see the potential for body-based visualization primarily in face-to-face interactions, yet due to technological limitations our study used a video call format. Thus these results might not directly translate as screen size, resolution, or network delay likely have an effect and in-person interaction is richer in comparison. In some of their comments our participants referred to the used medium (e.g., mentioning screen edges), so they might also have evaluated the quality of a remote conversation experience, not the potential of an in-person one.

Another limitation of the study format is that we had no way to directly measure participant distraction or engagement. Instead, we rely on self-reported distraction measures that depend on participants' abilities of introspection. With lack of an objective distraction measure, there also might be differences in participants' understanding of what constitutes distraction and how to relate their experience to the given scale. We could also only evaluate one side of the conversation, as the other conversation partner had to be a confederate. For the participant to see the conversation partner augmented, that partner needed to run the study lens and thus also always saw themselves with the added augmentations.

8 Conclusions and Future Work

We have investigated how body-based visualizations can display information and specifically how they impact conversations with others. Showing visual elements right on and around people we interact with allows to maintain focus on them while still receiving notifications and other information, including conversation-specific ones like mood indicators. We have prototyped five concrete visualizations as a Snapchat lens and tested them in short conversations via video calls with participants. Our participants overall found these visualizations helpful, useful, and not coming at a cost to their ability to focus on the conversation partner and the conversation in general. They also did not see such augmentations as inherently problematic or rude to the person these visualizations are shown on and next to.

These results point to promising use cases for visual feedback to occur not just in free space or around people, but also right on and close to people. In line with previous work, our results provide further evidence that augmentation of other people does not result in a deterioration of interactions with them. We focused on general notifications and information relevant to a casual chat conversation, but there are many other interpersonal scenarios were body-based visualizations might also be suitable and helpful. For example, future work might want to explore commercial scenarios (e.g., bartering), security applications (e.g., entry controls), or uses in entertainment (e.g., comedy shows). The devices for using such visual elements in real-life situations are still bulky and intrusive, but with advances in wearable AR, we hope body-based visualizations to become a suitable output format for interactive systems.

References

[1] ASM Iftekhar Anam, Shahinur Alam, and Mohammed Yeasin. 2014. Expression: A dyadic conversation aid using Google Glass for people who are blind or visually impaired. In *6th International Conference on Mobile Computing, Applications and Services*. IEEE, 57–64. https://doi.org/10.4108/icst.mobicase.2014.257780

[2] Fraser Anderson, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2015. Supporting Subtlety with Deceptive Devices and Illusory Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1489–1498. https://doi.org/10.1145/2702123.2702336

[3] Amit H. Bermano, Markus Billeter, Daisuke Iwai, and Anselm Grundhöfer. 2017. Makeup Lamps: Live Augmentation of Human Faces via Projection. *Computer Graphics Forum* 36, 2 (2017), 311–323. https://doi.org/10.1111/cgf.13128

[4] Jacques Bertin. 1983. *Semiology of graphics*. University of Wisconsin press.

[5] Arpit Bhatia, Henning Pohl, Teresa Hirzle, Hasti Seifi, and Kasper Hornbæk. 2024. Using the Visual Language of Comics to Alter Sensations in Augmented Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 603, 17 pages. https://doi.org/10.1145/3613904.3642351

[6] Leanne S. Bohannon, Andrew M. Herbert, Jeff B. Pelz, and Esa M. Rantanen. 2013. Eye contact and video-mediated communication: A review. *Displays* 34, 2 (2013), 177–185. https://doi.org/10.1016/j.displa.2012.10.009

[7] LouAnne E. Boyd, Alejandro Rangel, Helen Tomimbang, Andrea Conejo-Toledo, Kanika Patel, Monica Tentori, and Gillian R. Hayes. 2016. SayWAT: Augmenting Face-to-Face Conversations for Adults with Autism. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 4872–4883. https://doi.org/10.1145/2858036.2858215

[8] Runze Cai, Nuwan Nanayakkarawasam Peru Kandage Janaka, Shengdong Zhao, and Minghui Sun. 2023. ParaGlassMenu: Towards Social-Friendly Subtle Interactions in Conversations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 721, 21 pages. https://doi.org/10.1145/3544548.3581065

[9] Herman Chernoff. 1973. The Use of Faces to Represent Points in k-Dimensional Space Graphically. *J. Amer. Statist. Assoc.* 68, 342 (1973), 361–368. https://doi.org/10.1080/01621459.1973.10482434

[10] Enrico Costanza, Samuel A. Inverso, and Rebecca Allen. 2005. Toward Subtle Intimate Interfaces for Mobile Devices Using an EMG Controller. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) *(CHI '05)*. Association for Computing Machinery, New York, NY, USA, 481–489. https://doi.org/10.1145/1054972.1055039

[11] Sinem Şemsioğlu, Pelin Karaturhan, and Asım Evren Yantaç. 2022. EMOTE: An Interactive Online Tool for Designing Real-Time Emotional AR Visualizations. In *13th Augmented Human International Conference* (Winnipeg, MB, Canada) *(AH2022)*. Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. https://doi.org/10.1145/3532525.3532527

[12] Shakiba Davari, Feiyu Lu, and Doug A. Bowman. 2022. Validating the Benefits of Glanceable and Context-Aware Augmented Reality for Everyday Information Access Tasks. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 436–444. https://doi.org/10.1109/VR51125.2022.00063

[13] Paul Ekman, Richard J. Davidson, and Wallace V. Friesen. 1990. The Duchenne smile: Emotional expression and brain physiology: II. *Journal of Personality and Social Psychology* 58, 2 (1990), 342–353. https://doi.org/10.1037/0022-3514.58.2.342

[14] Paul Ekman, Wallace V. Friesen, Maureen O'Sullivan, Anthony Chan, Irene Diacoyanni-Tarlatzis, Karl Heider, Rainer Krause, William Ayhan LeCompte, Tom Pitcairn, Pio E. Ricci-Bitti, Klaus Scherer, Masatoshi Tomita, and Athanase Tzavaras. 1987. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of*

Personality and Social Psychology 53, 4 (1987), 712–717. https://doi.org/10.1037/0022-3514.53.4.712

[15] Ge Guo, Gilly Leshed, and Keith Evan Green. 2023. "I Normally Wouldn't Talk with Strangers": Introducing a Socio-Spatial Interface for Fostering Togetherness Between Strangers. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 272, 20 pages. https://doi.org/10.1145/3544548.3581325

[16] Rebecca Hansson and Peter Ljungstrand. 2000. The Reminder Bracelet: Subtle Notification Cues for Mobile Devices. In *CHI '00 Extended Abstracts on Human Factors in Computing Systems* (The Hague, The Netherlands) *(CHI EA '00)*. Association for Computing Machinery, New York, NY, USA, 323–324. https://doi.org/10.1145/633292.633488

[17] Chris Harrison and Haakon Faste. 2014. Implications of Location and Touch for On-Body Projected Interfaces. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) *(DIS '14)*. Association for Computing Machinery, New York, NY, USA, 543–552. https://doi.org/10.1145/2598510.2598587

[18] Susan C. Herring, Ashley R. Dainas, Holly Lopez Long, and Ying Tang. 2020. Animoji performances: "Cuz I can be a sexy poop". *Language@Internet* 18, 1 (2020). http://nbn-resolving.de/urn:nbn:de:0009-7-50465

[19] Roy S. Hessels. 2020. How does gaze to faces support face-to-face interaction? A review and perspective. *Psychonomic Bulletin & Review* 27, 5 (01 Oct 2020), 856–881. https://doi.org/10.3758/s13423-020-01715-w

[20] Thuong Hoang, Martin Reinoso, Zaher Joukhadar, Frank Vetere, and David Kelly. 2017. Augmented Studio: Projection Mapping on Moving Body for Physiotherapy Education. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 1419–1430. https://doi.org/10.1145/3025453.3025860

[21] Nuwan Janaka, Chloe Haigh, Hyeongcheol Kim, Shan Zhang, and Shengdong Zhao. 2022. Paracentral and Near-Peripheral Visualizations: Towards Attention-Maintaining Secondary Information Presentation on OHMDs during in-Person Social Interactions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 551, 14 pages. https://doi.org/10.1145/3491102.3502127

[22] Ana Javornik, Ben Marder, Jennifer Brannon Barhorst, Graeme McLean, Yvonne Rogers, Paul Marshall, and Luk Warlop. 2022. 'What lies behind the filter?' Uncovering the motivations for using augmented reality (AR) face filters on social media and their effect on well-being. *Computers in Human Behavior* 128 (2022), 107126. https://doi.org/10.1016/j.chb.2021.107126

[23] Marek Kowalski, Zbigniew Nasarzewski, Grzegorz Galinski, and Piotr Garbat. 2018. HoloFace: Augmenting Human-to-Human Interactions on HoloLens. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 141–149. https://doi.org/10.1109/WACV.2018.00022

[24] Mikko Kytö and David McGookin. 2017. Augmenting Multi-Party Face-to-Face Interactions Amongst Strangers with User Generated Content. *Computer Supported Cooperative Work (CSCW)* 26, 4 (01 Dec 2017), 527–562. https://doi.org/10.1007/s10606-017-9281-1

[25] Benjamin Lee, Michael Sedlmair, and Dieter Schmalstieg. 2024. Design Patterns for Situated Visualization in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 1324–1335. https://doi.org/10.1109/TVCG.2023.3327398

[26] Joanne Leong, Florian Perteneder, Muhender Raj Rajvee, and Pattie Maes. 2023. "Picture the Audience…": Exploring Private AR Face Filters for Online Public Speaking. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 512, 13 pages. https://doi.org/10.1145/3544548.3581039

[27] Jian Liao, Adnan Karim, Shivesh Singh Jadon, Rubaiat Habib Kazi, and Ryo Suzuki. 2022. RealityTalk: Real-Time Speech-Driven Augmented Presentation for AR Live Storytelling. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 17, 12 pages. https://doi.org/10.1145/3526113.3545702

[28] Jian Liao, Kevin Van, Zhijie Xia, and Ryo Suzuki. 2024. RealityEffects: Augmenting 3D Volumetric Videos with Object-Centric Annotation and Dynamic Visual Effects. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (IT University of Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 1248–1261. https://doi.org/10.1145/3643834.3661631

[29] Tica Lin, Chen Zhu-Tian, Yalong Yang, Daniele Chiappalupi, Johanna Beyer, and Hanspeter Pfister. 2023. The Quest for Omnioculars: Embedded Visualization for Augmenting Basketball Game Viewing Experiences. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 962–972. https://doi.org/10.1109/TVCG.2022.3209353

[30] Feiyu Lu, Shakiba Davari, Lee Lisle, Yuan Li, and Doug A. Bowman. 2020. Glanceable AR: Evaluating Information Access Methods for Head-Worn Augmented Reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 930–939. https://doi.org/10.1109/VR46266.2020.00113

[31] Daniel McDuff, Christophe Hurter, and Mar Gonzalez-Franco. 2017. Pulse and Vital Sign Measurement in Mixed Reality Using a HoloLens. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*

(Gothenburg, Sweden) *(VRST '17)*. Association for Computing Machinery, New York, NY, USA, Article 34, 9 pages. https://doi.org/10.1145/3139131.3139134

[32] Lynden Miles and Lucy Johnston. 2007. Detecting Happiness: Perceiver Sensitivity to Enjoyment and Non-Enjoyment Smiles. *Journal of Nonverbal Behavior* 31, 4 (01 Dec 2007), 259–275. https://doi.org/10.1007/s10919-007-0036-4

[33] Tien T. Nguyen, Duyen T. Nguyen, Shamsi T. Iqbal, and Eyal Ofek. 2015. The Known Stranger: Supporting Conversations between Strangers with Personalized Topic Suggestions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 555–564. https://doi.org/10.1145/2702123.2702411

[34] Hayoun Noh, Juhee Go, Sophia Song, Songi Kim, and Younah Kang. 2024. Investigating the Possibility of Using an AR Mask to Support Online Psychological Counseling. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 78 (apr 2024), 33 pages. https://doi.org/10.1145/3637355

[35] Leyla Norooz, Matthew Louis Mauriello, Anita Jorgensen, Brenna McNally, and Jon E. Froehlich. 2015. BodyVis: A New Approach to Body Learning Through Wearable Sensing and Visualization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1025–1034. https://doi.org/10.1145/2702123.2702299

[36] Eyal Ofek, Shamsi T. Iqbal, and Karin Strauss. 2013. Reducing Disruption from Subtle Information Delivery during a Conversation: Mode and Bandwidth Investigation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 3111–3120. https://doi.org/10.1145/2470654.2466425

[37] Hiromu Ogawa and Pattie Maes. 2020. Smartwatch-Based Topic Suggestions to Enrich Casual Conversations in Awkward Encounters. In *Proceedings of the 2020 ACM International Symposium on Wearable Computers* (Virtual Event, Mexico) *(ISWC '20)*. Association for Computing Machinery, New York, NY, USA, 68–72. https://doi.org/10.1145/3410531.3414310

[38] Alex Olwal, Kevin Balke, Dmitrii Votintcev, Thad Starner, Paula Conn, Bonnie Chinh, and Benoit Corda. 2020. Wearable Subtitles: Augmenting Spoken Communication with Lightweight Eyewear for All-Day Captioning. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1108–1120. https://doi.org/10.1145/3379337.3415817

[39] Henning Pohl, Tor-Salve Dalsgaard, Vesa Krasniqi, and Kasper Hornbæk. 2020. Body LayARs: A Toolkit for Body-Based Augmented Reality. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) *(VRST '20)*. Association for Computing Machinery, New York, NY, USA, Article 14, 11 pages. https://doi.org/10.1145/3385956.3418946

[40] Henning Pohl, Justyna Medrek, and Michael Rohs. 2016. ScatterWatch: Subtle Notifications via Indirect Illumination Scattered in the Skin. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Florence, Italy) *(MobileHCI '16)*. Association for Computing Machinery, New York, NY, USA, 7–16. https://doi.org/10.1145/2935334.2935351

[41] Henning Pohl, Andreea Muresan, and Kasper Hornbæk. 2019. Charting Subtle Interaction in the HCI Literature. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3290605.3300648

[42] Inka Rantala, Ashley Colley, and Jonna Häkkilä. 2018. Smart Jewelry: Augmenting Traditional Wearable Self-Expression Displays. In *Proceedings of the 7th ACM International Symposium on Pervasive Displays* (Munich, Germany) *(PerDis '18)*. Association for Computing Machinery, New York, NY, USA, Article 22, 8 pages. https://doi.org/10.1145/3205873.3205891

[43] Piera Riccio, Bill Psomas, Francesco Galati, Francisco Escolano, Thomas Hofmann, and Nuria Oliver. 2022. OpenFilter: A Framework to Democratize Research Access to Social Media AR Filters. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 12491–12503. https://proceedings.neurips.cc/paper_files/paper/2022/file/50fd4a244de17f856709036edda9854e-Paper-Datasets_and_Benchmarks.pdf

[44] Radiah Rivu, Yasmeen Abdrabou, Ken Pfeuffer, Augusto Esteves, Stefanie Meitner, and Florian Alt. 2020. StARe: Gaze-Assisted Face-to-Face Communication in Augmented Reality. In *ACM Symposium on Eye Tracking Research and Applications* (Stuttgart, Germany) *(ETRA '20 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 14, 5 pages. https://doi.org/10.1145/3379157.3388930

[45] Jan Ole Rixen, Mark Colley, Ali Askari, Jan Gugenheimer, and Enrico Rukzio. 2022. Consent in the Age of AR: Investigating The Comfort With Displaying Personal Information in Augmented Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 295, 14 pages. https://doi.org/10.1145/3491102.3502140

[46] Jan Ole Rixen, Teresa Hirzle, Mark Colley, Yannick Etzel, Enrico Rukzio, and Jan Gugenheimer. 2021. Exploring Augmented Visual Alterations in Interpersonal Communication. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA,

Article 730, 11 pages. https://doi.org/10.1145/3411764.3445597

[47] Rufat Rzayev, Susanne Korbely, Milena Maul, Alina Schark, Valentin Schwind, and Niels Henze. 2020. Effects of Position and Alignment of Notifications on AR Glasses during Social Interaction. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) *(NordiCHI '20)*. Association for Computing Machinery, New York, NY, USA, Article 30, 11 pages. https://doi.org/10.1145/3419249.3420095

[48] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-Scale Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 3055–3064. https://doi.org/10.1145/2556288.2557189

[49] Gianluca Schiavo, Alessandro Cappelletti, Eleonora Mencarini, Oliviero Stock, and Massimo Zancanaro. 2014. Overt or Subtle? Supporting Group Conversations with Automatically Targeted Directives. In *Proceedings of the 19th International Conference on Intelligent User Interfaces* (Haifa, Israel) *(IUI '14)*. Association for Computing Machinery, New York, NY, USA, 225–234. https://doi.org/10.1145/2557500.2557507

[50] Stefan Schneegass, Sophie Ogando, and Florian Alt. 2016. Using On-Body Displays for Extending the Output of Wearable Devices. In *Proceedings of the 5th ACM International Symposium on Pervasive Displays* (Oulu, Finland) *(PerDis '16)*. Association for Computing Machinery, New York, NY, USA, 67–74. https://doi.org/10.1145/2914920.2915021

[51] Daniel Story and Ryan Jenkins. 2023. Deepfake Pornography and the Ethics of Non-Veridical Representations. *Philosophy & Technology* 36, 3 (26 Aug 2023), 56. https://doi.org/10.1007/s13347-023-00657-0

[52] Mohammad Iftekhar Tanveer and Mohammed Ehsan Hoque. 2014. A Google Glass App to Help the Blind in Small Talk. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility* (Rochester, New York, USA) *(ASSETS '14)*. Association for Computing Machinery, New York, NY, USA, 297–298. https://doi.org/10.1145/2661334.2661338

[53] Mariek M.P. Vanden Abeele, Andrew T. Hendrickson, Monique M.H. Pollmann, and Rich Ling. 2019. Phubbing behavior in conversations and its relation to perceived conversation intimacy and distraction: An exploratory observation study. *Computers in Human Behavior* 100 (2019), 35–47. https://doi.org/10.1016/j.chb.2019.06.004

[54] Peter Washington, Catalin Voss, Aaron Kline, Nick Haber, Jena Daniels, Azar Fazel, Titas De, Carl Feinstein, Terry Winograd, and Dennis Wall. 2017. SuperpowerGlass: A Wearable Aid for the At-Home Therapy of Children with Autism. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 112 (sep 2017), 22 pages. https://doi.org/10.1145/3130977

[55] Ian Watts, Pierre Boulanger, and Greg Kawchuk. 2017. ProjectDR: Augmented Reality System for Displaying Medical Images Directly onto a Patient. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology* (Gothenburg, Sweden) *(VRST '17)*. Association for Computing Machinery, New York, NY, USA, Article 70, 2 pages. https://doi.org/10.1145/3139131.3141198

[56] Thomas H. Wheeler. 2002. *Phototruth Or Photofiction? Ethics and Media Imagery in the Digital Age.* Routledge, New York, NY, USA. https://doi.org/10.4324/9781410613080

[57] Kristin Williams, Karyn Moffatt, Denise McCall, and Leah Findlater. 2015. Designing Conversation Cues on a Head-Worn Display to Support Persons with Aphasia. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 231–240. https://doi.org/10.1145/2702123.2702484

[58] Robert Xiao, Teng Cao, Ning Guo, Jun Zhuo, Yang Zhang, and Chris Harrison. 2018. LumiWatch: On-Arm Projected Graphics and Touch Input. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3173574.3173669

A Participant Instructions

Participants received the following message (personal data omitted) after scheduling a video call with us:

Thanks for signing up for our study! For the video call you will need to connect to the user handle '<handle>' (given name: <name>) on Snapchat. Please add this user to your friend list before the call. After this study has concluded we will delete our account and thus also not appear in your friend list anymore. We will not retain any information on your Snapchat account and no data is recorded during our call. The only data we keep is from the questionnaire you have to fill out afterwards. We will send you a link to that questionnaire after our call has concluded.

At the scheduled time, please video call us via the Snapchat app or Snapchat on the web (https://web.snapchat.com/). We do not know which Snapchat account belongs to your Prolific ID, so we cannot initiate the call ourselves. Should you not be able to make it at that time after all, you can reschedule or cancel via the booking site.

During the video call we would like to just have a casual conversation. What we are trying to evaluate in this study is your opinion of a set of visualization elements that will become active during the call. These simulate augmented reality effects we believe could enhance conversations in real-life in the future. We have implemented five different visualization elements:

- *A virtual name tag that shows you the name of who you are speaking to.*
- *An emotion tracker that analyzes the person you are talking to and gives you an estimate of their current mood. This is shown as a bar graph above their left eye, filling up more the happier they seem.*
- *A weather notification widget that let's you know about upcoming changes to the weather around you. That widget floats next to the person you are talking to.*
- *A message notification that let's you know about incoming text messages. These show up as text bubbles popping out from the side of the person you are talking to and show the number of unread messages.*
- *A conversation clock that informs you how much time has already passed talking to that person. A green ring around the right eye shows seconds and blue marks above the eye show how many minutes have passed.*

For an image of what these effects look like, please see: <link to picture also shown in Figure 1>

Note that the notifications are simulated and do not represent actual incoming messages or changing weather. For the purpose of this study, please imagine that the shown information is generated by your own device and thus personal to you. The person you are talking to is not able to see these effects and these visualization elements are meant to be personal to you, even though they are shown on and around the other person's body.

Received February 2024; revised May 2024; accepted June 2024