Understanding Embodied Communication Affordances in Team Perception and Participation

要　旨

【キーワード】
Telepresence, robotics, collaboration technologies

我々は、会議の遠隔参加者を具現化するDesktop Telepresenceロボット*Jarvis*を提案する。Jarvisと、従来型TV会議や対面会議との比較研究を実施した。この比較で、会議室にいる参加者と遠隔参加者が知覚する参与度合いに差があることがわかった。Jarvisを使用した遠隔会議では、従来型と比較し、会議室にいる参加者の発話のオーバラップ量が有意に多かった。（それは、対面会議と同等であった。）我々の技術が、会話を行うためにどのように利用されたのか、その違いについて記述する。研究結果では、遠隔参加者の動きを実現した技術が、会議室にいる参加者への存在感を増加させたが、遠隔参加者の参与感を増加させることはなかった。この検証を通じて、エージェントやプレゼンスという観点で、Telepresenceソリューションの可能性について議論する。

Abstract

【Keywords】
Telepresence, robotics, collaboration technologies

We present *Jarvis*, a desktop meeting telepresence robot that embodies a remote participant. We conducted a study that compared Jarvis to basic video conferencing and face-to-face meetings. We show that local and remote participants perceived differently their own contributions and others' across conditions. Local participants exhibited significantly more overlapping talk with remote participants who used Jarvis, than with remote participants in basic-video conferencing (and at a rate similar to overlapping speech for co-located groups). We describe differences in how the technologies were used to follow conversation. Our findings indicate that while the kinetic embodied technology increased local participants' perceived presence of remote teammates, it did not enhance remote participants' own sense of telepresence. We discuss our findings in the context of theories of agency and presence, and discuss how these findings extend our understanding of the promise and limitations of embodied video-conferencing solutions.

Author

Jacob T. Biehl
Daniel Avrahami
Anthony Dunnigan

FX Palo Alto Laboratory, Inc

1. Introduction

Advances in processing hardware, encoding techniques, and broadband networks are driving a rise in the adoption of video-based communication technologies, especially in the workplace. These technologies have been shown to increase the frequency of communication between co-workers and, in some cases, improve productivity. A recent industry survey of human-resource executives illustrates the disruptive nature of this technology and boldly postulates that video-based communication tools will be preferred over email in the workplace by 2016 [6]. Although world travel has been steadily growing, recent trends on corporate travel budgets show an overall decline in travel. This suggests an economic driving force for adoption of video-based communication.

Despite the benefits of video-based communication, limitations still exist. Past work has shown that remote users still face a disadvantage when compared to their co-located colleagues. Remote collaborators participate less in conversations, take less dominant roles in groups, and feel less connected to distant coworkers. A new class of technologies has emerged that seeks to mitigate the social disadvantages of video-based communication by providing remote users with a local embodiment. These technologies range from fully mobile robots, to smaller devices that sit on a table.

Studies of these embodied telepresence devices show that they provide distributed teams with an increased sense of their remote colleagues' presence in the local environment and a reciprocal sense of "being there" for that remote worker. Studies have also shown that using these devices can change collaborative behavior. For example, [10] found that their use led to more opportunistic interactions among distant workers.

While the impact of introducing these technologies has been explored at a group level, few studies have investigated how they impact an actual collaborative engagement (e.g. a workplace meeting).

In this paper, we present a lab study that compares the use of an embodied telepresence device named *Jarvis*, a traditional video-conferencing configuration, and fully co-located baseline to complete a specific, collaborative task. In agreement with prior work, we first show the adverse effect of being remote on actual and perceived contributions. We then describe qualitative and quantitative shifts in group interaction and attitudes related to the use of embodied technology. In particular, our results show that local participants interact with a remote person in a fashion more similar to face-to-face when using an embodied technology – potentially as a result of increased perceived agency of the remote person. Our results also highlight the difficulty providing participants a sense of agency in the remote space.

We believe the results and insights from our work strengthen the findings of past studies, while providing a comparative baseline across technologies. Using our comparative result, we also offer design insight for improving the current state-of-the-art embodied telepresence devices.

2. Study

To capture group communication and task engagement, we conducted a between-subject group-collaboration study in which some participants were co-located and others remote. Similar to [4, 8] our study employed a confederate. In our study, groups of three participants evaluated a junior sales trainee (the confederate) on a short sales presentation.

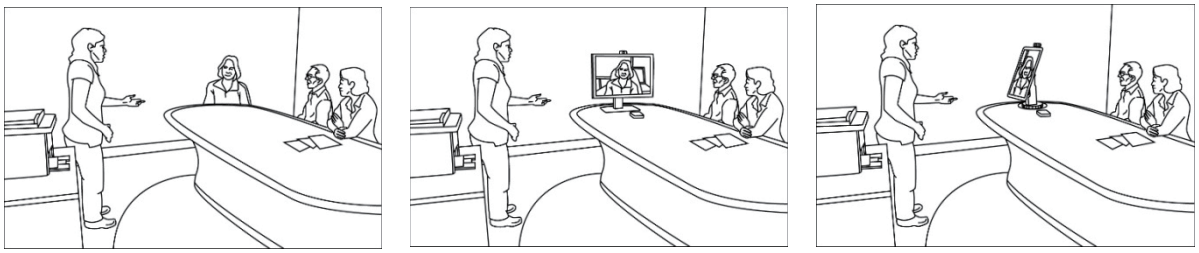

(a) F2F (b) Static (c) Kinetic

Figure 1. Study setup with a presenter and three participants. In the Static (b) and Kinetic (c) conditions, one participant is remote.

2.1 Task

The groups' task – evaluating a junior sales trainee – consisted of three phases: *Presentation*, *Discussion*, and *Feedback*.

In the *Presentation* phase, the sales trainee gave a short sales presentation about a specific multifunction printer for home and small-office use. In order to keep the task consistent across all groups and conditions, we hired a professionally trained actor to play the role of the trainee. The actor was instructed to "play the same character" in all sessions, including consistently performing the same non-verbal behaviors, body posture, and vocal delivery. Scripted content and behaviors gave the appearance of a young, naïve sales trainee who had many faults and struggles in her presentation.

After the presentation was done and the presenter excused from the room, a *Discussion* phase began, in which participants privately discussed the presenter's performance. Finally, the presenter was brought back into the room for the *Feedback* phase, and participants provided her with feedback.

2.2 Conditions

Our study consisted of three group-level conditions used to understand how groups performed an authentic task where group communication was mediated through different technologies: a baseline condition (*F2F*), with all 3 participants co-located, and two conditions (*Static* and *Kinetic*) where one participant was "remote" and the other two in the same room. This is similar to the *hub-and-satellite* meeting paradigm used in [10].

In the *F2F* condition, which served as baseline, all 3 participants (say, participants A, B, and C) were co-located and sat around a medium-sized meeting table (see Figure 1a). In the *Static* and *Kinetic* conditions, however, participant C was in a different room (as a "*Remote*" participant), and communicated via a two-party video call.

In the *Static* condition, which represents current practices for business video conferencing, participant C was in a different room, and communicated with the group via a two-party video call. A large, 30" display was placed on the meeting table in the approximate position of participant C in the F2F condition (Figure 1b). A camera was positioned at the top of the display.

In the *Kinetic* condition (Figure 1c), participant C communicated with the group using a prototype desktop telepresence device called *Jarvis* (see Figure 2). The Jarvis device was chosen as it faithfully represents the emerging class of embodied telepresence devices, which include commercial devices. Jarvis uses a large, portrait-oriented 18" tablet to display the remote participant. This larger display represents the remote participant at "human scale" and proportional to how remote participants were represented in the Static condition. A camera is mounted to the top of the device, creating a direct mapping between display orientation and remote view direction and angle. Jarvis enables the remote participant

Figure 2. Jarvis: A desktop pan & tilt telepresence robot equipped with an 18" screen and HD video camera.

to pan and tilt the display by clicking on the remote video, which re-centered the camera view at the click point. Similar to the Static condition, Jarvis was placed on the meeting table in the approximate position of participant C in the F2F condition.

Video and audio transmission was handled using custom software. All of the clients and infrastructure were installed on our gigabit local network. Video was captured and streamed from both locations using a camera running at 960x540 resolution, 30 frames per second, and at an average bit rate of 1200 Kbps. Remote video that appeared on Jarvis was cropped to fit the aspect ratio of the device's screen. In both conditions, the audio channel was open (i.e. no "push-to-talk" interaction was needed).

2.3 Procedure

Upon arriving at our lab, participants were shown to a meeting room. For the Kinetic and Static conditions, the first participant to arrive at the lab was whisked to a separate room and assigned the role of *Remote* participant. They all completed informed consent, demographic, and pre-study questionnaires.

Participants were told that their task was to evaluate a junior sales trainee who, as part of her training, had to make a short sales pitch in front of a live audience. The three phases of the task were described and the presenter was brought to the room.

After a short introduction, the confederate began her presentation, following a script that covered the features, functionality, and characteristics of the printer. Participants were allowed to interrupt and ask questions, but the actress was instructed to keep to and return to her script as much as possible. The presenter was then excused from the room and the group began their private discussion. Participants were told that they are allowed to move around the table or reposition the display if they so wished and that they could take as much time as they needed. When participants informed the researcher that they were done, the presenter was brought back and the group provided feedback.

Participants were then taken to separate rooms to complete post-study questionnaires where they rated the performance of the presenter, as well as their own and fellow evaluators' performance. Finally, we asked remote participants to comment on the challenges of working as a group through mediated communication.

2.4 Participants

We recruited participants via fliers and postings to community bulletin boards (e.g. Craigslist). The call indicated that participants would be given an Amazon gift card in return for an hour of their time to evaluate a sales presentation. 27 participants were recruited (11 women), with an average age of 44 (SD=13.8).

A pre-study questionnaire was administered to determine all participants' familiarity and experience using video communication tools (e.g. video chat). Responses varied from no experience to daily use. Most users reported using video chat at a frequency of 1-3 times per month across a variety of work and personal

uses. No meaningful differences in frequency or type of use were seen across conditions.

Participants were randomly assigned to one of the three Group conditions, with three groups per condition. Participants did not know each other.

2.5 Measures

We collected a variety of subjective and objective measures, which included:

- *Performance questionnaire.* Participants rated the quality of the trainee's presentation across several dimensions. They also rated their own performance as evaluators, as well as the performance of fellow evaluators. The questions were modeled after existing shared identity, reciprocity, and group efficacy questionnaires.
- *Video analysis.* We recorded each session using multiple cameras (local and remote, when applicable) and coded for gestures and manipulation of artifacts and devices (including display and robot).
- *Turn-taking and Overlapping Talk.* We hand-coded video and audio recordings for start and end of speech segments ("turns"). Additionally, we coded turns that began when another participant was already talking as "Overlapping Talk" and noted the role (*remote* or *local*) of the already-speaking participant.
- *Open-ended post-study questions.* We examined comments from participants in the Static and Kinetic conditions on the challenges of working as a group through mediated communication.

3. Results

A total of 4.5 hours of video were recorded (30 minutes per group, on average). The average length, in minutes, for Discussion and Feedback were 9:47 (SD=5.2) and 10:26 (SD=7.6), respectively. Length varied across groups because we allowed participants to take as long as they wished for their discussion and the feedback.

We now report observations from the performance questionnaires, analysis of the videos, and the post-study interviews. We compare the behavior and attitudes of remote and co-located participants then report some surprising behavioral differences between our conditions. Finally, we examine whether (and how) participants took advantage of the ability to adjust the point-of-view (POV).

3.1 Participation and Perceived Performance

We examined if being a remote participant affected participation and perception, and looked for differences between the Kinetic and Static conditions. Remote participants are often at risk of becoming isolates. Indeed, we found that participation levels of remote participants were significantly lower than their co-located counterparts. Remote participants took the floor 22% of the time, on average, compared to 37% for local participants ($t(25)=2.184$; $p=0.039$). We saw no significant difference in floor time between the Kinetic and Static conditions (although differences in turn length and overlapping talk are reported later).

The reduced participation (or perceived ability to contribute) was further reflected in participants' responses in the performance questionnaire; Remote participants rated their own contribution (out of 7) significantly lower than local participants did (4.3 vs. 6.1; $F(1,5)=11.41$, $p=0.02$, with Group as a random factor, see Figure 3). Additionally, looking at how local participants in the Kinetic and Static conditions rated the contribution of teammates showed a marginally significant effect with remote teammates receiving lower ratings than co-located (5.1 vs 6.25; $t(5)=1.99$, $p=.052$, see Figure 4). However, no significant differences were found between the Kinetic and Static conditions. Finally, when asked to rank order

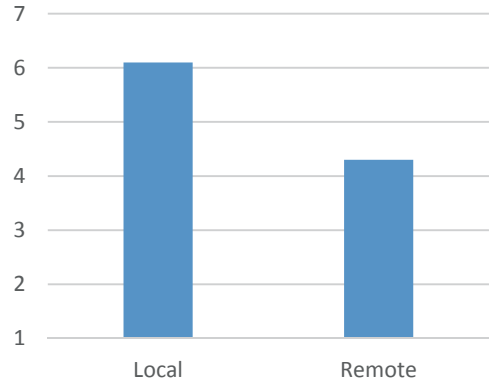

Figure 3. Self-Rating of contribution to the discussion (scale 1-7)

the group from most contributing member (1) to least contributing member (3), not a single remote participant ranked themselves as most helpful, while 62% of local participants did. (In fact, 83% of remote participants ranked themselves as least helpful.)

Interestingly, remote participants in both Kinetic and Static conditions rated the sales presentation significantly less favorably than local participants did (2.3 vs. 3.4; $F(1,5)=12.3$, $p=0.017$, with Group as a random factor). Asked to rate whether they thought the trainee will be a successful sales professional, remote participants gave significantly lower rating than co-located participants (M=3.0 vs. 5.1; $F(1,5)=12.00$, $p=0.018$, with Group as a random factor). In their post-study interview, a remote participants in the *Static* condition stated "*I was able to observe more objectively maybe, listen to what they're saying*." We saw no significant difference in ratings of the presentation and

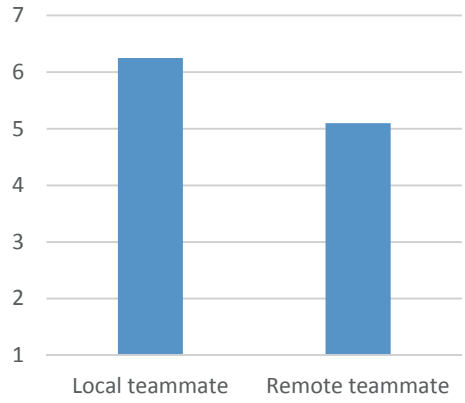

Figure 4. Local participants' rating of the contribution of teammates (scale 1-7)

presenter between the remote conditions and there was no significant difference between remote participants in the Kinetic and Static conditions.

3.2 Talk and Overlapping Talk

Next, we wanted to explore whether the use of the embodied device would produce communication style more similar to face-to-face communication. We used recorded video and audio to code the start and end times of participants' speech. We also coded for overlapping talk (i.e., whether a participant started speaking over a participant that was already speaking). To investigate group turn-taking and overlapping talk absent the presenter's (confederate) influence, we focus our analysis on the Discussion phase (during which, the presenter was not in the room).

3.2.1 Turn-Taking

Analysis of turns taken per minute showed no significant difference across conditions (Kinetic M=1.65, SD=0.93, Static M=1.51, SD=0.76, and F2F M=2.11, SD=1.04). Turn *length* showed a marginal difference across conditions; turn length was greater on average in the Static condition (15.5s, SD=19.41) compared to F2F (11.8s, SD=11.78) and Kinetic conditions (11.7s, SD=14.10; $F(2,410)=2.693$, p=0.07, see Figure 5). Comparing length of turns taken by remote participants showed only a marginally significant difference between conditions (Bonferroni correction, with $\alpha=0.017$), with Kinetic participants taking shorter turns than Static participants (8.5s vs. 20.2s; $t(50)=2.078$, $p=0.043$).

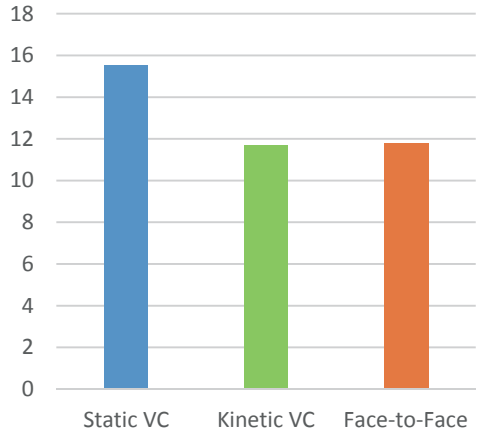

Figure 5. Turn taking length (in seconds).

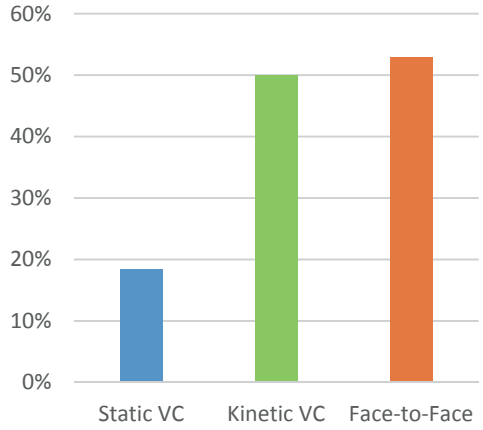

Figure 6. Percentage of overlapping talk when the remote participant was interrupted (reported as percentage of overall turns).

3.2.2 Overlapping Talk

Overlapping talk—when *"more persons than one talk at once"*—is a natural component of turn-taking in conversation. Prior literature found overlapping talk in natural face-to-face conversations in over 40% of turns (*c.f.* [2, 3]) Computer-mediated meetings have been shown to contain less overlapping talk than co-located meetings [7].

In our data, the rate of overlapping talk was high (47.8%, SD=13.8) and there was no main effect of condition. However, looking specifically at the involvement of remote participants in overlapping talk uncovered interesting results: local participants were *three times more likely* to interrupt a remote participant in the Kinetic condition (at 50%, similar to the rate of overlapping talk in the F2F condition) than to interrupt a remote participant in the Static condition (where only 18% of turns were interrupted by a local participant), see Figure 6. This difference, with Bonferroni connection, is significant ($t(4)=7.47$, $p=0.017$).

On the other hand, there was no significant difference when the remote participant was the one interrupting (7.5% and 13.5% for the Kinetic and Static conditions, respectively; $t(4)=1.573$, $p=0.625$).

3.3 Manipulating the Remote View

The ability of a remote participant to adjust their point-of-view (POV) was a key difference between the Kinetic and Static conditions. Unlike in [8], where a confederate operated the embodied proxy, our study allowed participants to use this capability as their own discretion.

Thus, we felt it was interesting to look at when this ability was used and for what purpose. Further we wanted to understand if there was contrasting behavior the Static condition when the capability was unavailable.

Through examination of recorded video, we saw that remote participants exhibited varying degrees of use. While one remote participant changed their POV only twice: at the beginning of the Discussion phase, and at the beginning of the Feedback phase, another remote participant adjusted their POV 22 times in a 40-minute session. (The remaining remote participant changed their POV 6 times in that group's 30-minute session.)

As expected, a large portion (43%) of all changes in POV were made to be able to see the person(s) speaking (since co-located participants rarely moved, orienting the POV to the person speaking was relatively easy). Another trigger to change the POV was when the focus of attention of the co-located participants has shifted. This included non-verbal shift – e.g., when co-located participants look towards the presenter who's returned to the room for the Feedback session – or when the subject of the conversation was

out of view – e.g., when features of the printer were discussed. Finally, in two cases, changing POV was also used to visually track a local participant moving around the room (e.g., to inspect the printer).

At the end of the Presentation phase, participants were told that they should feel free to move around or turn the screen around. However, unlike the Kinetic condition, remote participants in the Static condition could not adjust their POV. As stated by one of the participants "*they can hear me, they can see me, but I could only see [one of the local participants]* ***or*** *[the presenter]*." Thus, these participants had to rely on co-located group members to change their POV. Through examination of recorded video we observed that in all groups in the Static condition, co-located users manually adjusted the display and camera of the remote participant to afford them a better view of the participants (Discussion phase) or presenter (Feedback phase). In fact, one of the Static groups deliberately asked the presenter to sit at a seat at the opposite end from the remote participant for the Feedback phase, such that participants and presenter could be in view. These observations are important as they highlight the groups' need for adjusting the remote person's POV as meeting circumstances change. This supports past findings that local participants moved the embodied remote participant to foster inclusion in the group. Interestingly, manually adjusting a remote person's POV was sometimes accompanied with a statement such as "*We're going to move you now*" (local participant to remote participant).

4. Discussion

Our study compared an embodied technology against both a non-embodied technology and a baseline face-to-face, all co-located group configuration. Our results highlight a likely understood, but not often discussed assessment of the state-of-the-art; despite many advantages, embodied technologies still exhibit many of the known negative effects introduced by a mediated communication channel. Remote participants (in both conditions) contributed less, rated the experience poorer, and generally felt less connected than their co-located peers. Further, these negative findings were not significantly different between the two remote conditions. This suggests that the affordances of the kinetic embodiment are unable to fully overcome the negative impacts of the mediated channel.

In [8] a confederate was used as the operator of a proxy when studying the impact of different forms of kinetic motion, while in [4] a confederate was used as the local participant. In contrast, our study used both local and remote participants. Our observations suggest that, when the use of the kinetic capability is left up to a participant, one will observe a wide range of usages of the medium.

Prior work (e.g. [10]) showed that in addition to improved group collaborations, embodied technologies facilitated social integration and team building over time. Our study, by contrast, used participants who did not know one another, and yielded several different findings. While work team members in the real world are likely to be familiar with one another, the difference in findings suggests that embodied technologies may not directly impact group collaboration, but the stronger social connections fostered by the technology are what direct impacts collaboration quality. This consideration certainly needs further exploration. However, if true, it has strong implications for the use of these technologies. Specifically, loosely affiliated or one time groups (e.g. a product evaluation panel demonstrated in our study) may not be the ideal user population of embodied technologies.

While we did not find many differences

between the Kinetic and Static conditions with respect to perceptions of group performance, we did find several interesting nuanced differences. It was particularly interesting to examine how and why the remote participant adjusted their point of view (POV), contrasted with reliance on local participants in the Static condition. Kinetic remote participants used and appreciated the capability – it assisted in focusing conversation and conveyance of visual attention. Local participants appreciated the impact on group dynamic as well. One local participant noted, *"to use this [referring to the embodied device] in a group setting, for meetings like this, that's great that he can, you know, move around and see everyone."* In comparison, the Static condition placed a burden on local participants, and as local participants explained, the *"[she] was so far away"* and *"[he] only had one point of view."* In contrast to prior findings [5] the differences in behavior and perception we observed suggests that for larger, unbalanced teams relying on local teammates to adjust the POV might not support equitable group interaction.

A particularly interesting difference was found in turn-taking behavior between Kinetic and Static remote participants. Kinetic remote participants' turn-taking behavior more closely matched that of their co-located counterparts. A difference was also found in the frequency of overlapping talk. Particularly, the rate of interruptions of remote participants by local participants in the Kinetic condition was similar to rates of interruption between participants in the co-located condition. This is in contrast to the Static condition where interruptions were far fewer.

These results suggest that the increased agency in the Kinetic condition led participants to interact with the remote person in a fashion more similar to their interaction with co-located participants - that is, interrupt them in the same way they would interrupt a person at the same table. Thus, Kinetic embodiments may have meaningful value in business and other professional meetings where participation equality is important to preserve. It is interesting to consider that the amount of overlapping talk generated by the remote participants themselves did not differ significantly between the Static and Kinetic conditions.

Finally, qualitative feedback in the post-study interviews also indicated a subtle shift in how the disadvantages of the mediated channel were perceived by local participants. In the Static condition, local participants were sympathetic to the remote participants. As one local participant put it, *"We could adjust to the situation. We can just jump in. He [remote participant] can't."* Another stated *"it's harder for the person on that end [remote participant] to get a good idea of what's really going on."*

This was in contrast to the Kinetic condition, where participants' feedback did not directly discount the affordances of the remote participant. For instance, one participant stated he *"…recognized there were three people here, he [remote participant] wasn't more or less included."* Another participant stated *"every time he [remote participant] had a question we would address him without any problem."* Some participants even attributed the reduced remote participation to the personality of the remote participant – *"he [remote participant] was more non talkative than he should have been. Which, of course, this would affect you interacting with anyone. So I don't think that this was the technology, just the person who is using it."* These observations may be early evidence of the embodied device's ability to support stronger social integration and team bonding.

Taken as a whole, our findings may provide broader insights into how embodied telepresence devices impact collaborators' *sense of presence* and ability to achieve *agency*.

For instance, the shift in communication behavior that resembled the F2F teams could suggest that local participants perceived an increased sense of presence with remote participant. In contrast, our results also show that, despite an ability to adjust their point-of-view, remote participants in the Kinetic condition did not exhibit detectable behavior shifts, nor provide any qualitative feedback that would point to a change in presence or agency. Findings by Tang et al. [9] may explain that remote users' inability to view their own physical remote representation inhibited their sense of agency in the remote space. Indeed, the environment through which remote participants in both conditions experienced the remote space were nearly identical –as a video on the screen. We feel these findings push forward our understanding of the value and use of embodied telepresence devices, while also outlining the need for much deeper investigations. Based on the current results, we are particularly compelled to explore ways to provide a remote team member with a proper representation of their remote physical manifestation.

5. Conclusions and future work

In this work, we compared basic video conferencing, emergent kinetic video-conferencing techniques, and face-to-face meetings. Our findings build upon previous research, showing the adverse impact of being remote on participation levels and perception of contribution. We also showed an interesting difference in overlapping speech between the conditions. We then showed the potential of kinetic embodied proxies to impact behaviors exhibited by local participants in video conferencing meetings.

Still, a better understanding is needed of the work tasks that benefit (or suffer) from use of kinetic telepresence, and the types of work teams for whom such proxies are most appropriate. As such, we plan to investigate the use of a proxy similar to Jarvis in across a variety of tasks and settings using methodology similar to ESP [10].

Lastly, a key limitation of telepresence proxies, such as the one used in this study and others (e.g., [1,4,8]) is that, while they create an embodiment of the remote person in the local space, the remote person still experiences the space the same way – as a video on the screen (albeit with an ability to adjust their POV). This is supported in part by our results that show difference in turn-taking behavior for local participants with an embodied device present, without a parallel effect on the remote participants.

In future work, we plan to explore mechanisms (such as a 3rd person view of their proxy) to provide the remote person a sense of their representation in the local space, and examine its effect on their sense of presence.

6. References

[1] Adalgeirsson, S. O., & Breazeal, C. MeBot: A Robotic Platform for Socially Embodied Presence. *In Proc*. *HRI 2010*. IEEE Press (2010), 15-22.

[2] Heldner, M., & Edlund, J. Pauses, Gaps and Overlaps in Conversations. *Phonetics*, 38(4), (2010), 555-568.

[3] Kurtić, E., Brown, G. J., & Wells, B. Resources for Turn Competition in Overlapping Talk. *Speech Communication*, 55(5), (2013), 721-743.

[4] Rae, I., Mutlu, B., and Takayama, L. Bodies in Motion: Mobility, Presence, and Task Awareness in Telepresence. *In Proc*. *CHI 2014*. ACM Press (2014), 2153-2162.

[5] Rae, I., Takayama, L. and Mutlu, B. In-Body Experiences: Embodiment, Control, and Trust in Robot-Mediated Communication. In Proc. CHI 2013. ACM

Press (2013), 1921-1930.

[6] Redshift Research. Global View: Business Video Conferencing Usage and Trends. (2014)

[7] Sellen, A. J. Speech Patterns in Video-Mediated Conversations. *In Proc*. *CHI 1992*. ACM Press (1992), 49-59.

[8] Sirkin, D., Venolia, G. Tang, J.C., Robertson, G.G., Kim, T., Inkpen, K., Sedlins, M., Lee, B., and Sinclair, M. Motion and Attention in a Kinetic Videoconferencing Proxy. *In Proc*. *INTERACT 2011*. IFIP (2011), 162-180.

[9] Tang, A., Biocca, F., & Lim, L. Comparing Differences in Presence During Social Interaction in Augmented Reality Versus Virtual Reality Environments: An Exploratory Study. *In Proc*. *Presence 2004*, (2004), 204-208.

[10] Venolia, G., Tang, J., Cervantes, R., Bly, S., Robertson, G., Lee, B., and Inkpen, K. Embodied Social Proxy: Mediated Interpersonal Connection in Hub-and-Satellite Teams. *In Proc*. *CHI 2010*, ACM Press (2010), 1049-1058.

Author's Introductions

Jacob T. Biehl
FX Palo Alto Laboratory
Area of specialty: Computer Science (Ph.D.), Human-Computer Interaction, Collaborative Systems

Daniel Avrahami
FX Palo Alto Laboratory
Area of specialty: Human-Computer Interaction (Ph.D.), Ubiquitous Computing, Collaboration and Communication

Anthony Dunnigan
FX Palo Alto Laboratory
Area of specialty: Visual Design and Communication