Capturing Performative Actions for Interaction and Social Awareness

Julie R. Williamson and Stephen Brewster

University of Glasgow, Glasgow UK

Abstract. Capturing and making use of observable actions and behaviours presents compelling opportunities for allowing end-users to interact with such data and eachother. For example, simple visualisations based on on detected behaviour or context allow users to interpret this data based on their existing knowledge and awarness of social cues. This paper presents one such "remote awareness" application where users can interpret a visualization based on simple behaviours to gain a sense of awareness of other users' current context or actions. Using a prop embedded with sensors, users could control the visualisation using gesture and voice-based input. The results of this work describe the kinds of performances users generated during the trial, how they imagined the actions of their fellow participants based on the visualisation, and how the props containing sensors were used to support, or in some cases hinder, successful performance and interaction.

1 Introduction

Capturing and using actions and behaviours for interaction has seen a wide variety of applications, from replacing traditional buttons with gestures (Crossan *et al.*, 2008) (Strachan *et al.*, 2007), to supporting self expression through performance in public places (Perry *et al.*, 2010) (Sheridan *et al.*, 2011), to creating remote awareness of friends and family through ambient interfaces (Dey and de Guzman, 2006). In the area of social signal processing, previous research has focused on creating a foundation of work aimed at sensing and detecting social signals, where effectively using or applying those signals for interaction remains an open challenge (Vinciarelli *et al.*, 2009). This paper presents a possible application area for these signals where simple actions and behaviours are sensed and visualized for intepretation by the users themselves in a remote awareness scenario.

An important aspect of this remote awareness scenario is that the actions and behaviours sensed by the system can be understood as *performances.* Indeed, nearly any action completed in a public place can be considered a performance of some kind. Goffman (1990) describes a wide variety of "performances" that people produce everyday, ranging from unconciously performed actions to specifically designed and directed personas and impression management. Goffman (1966) also describes social contexts as either focused or unfocused, where focused interaction are those with a single point of attention and

R. Murray-Smith (Ed.): MSSP 2010, LNCS 8045, pp. 51–63, 2014.

involve cooperation as opposed to unfocused interaction where people might be in the same place but not actively cooperating or interacting together. This performative perspective helps organize behaviour in public places into relevant categories and highlight behaviours of interest to social signal processing.

This paper presents an application that makes use of basic social signals to allow users to interpret these signals through a simple visualization. In the application, called MuMo, each user is represented by a fish in a virtual fish tank. Users' actions are displayed in the visualisation through their fish, where gestures or movements cause the fish to swim faster and audio or speech cause the fish to blow bubbles. Each user can view the fish tank visualisation as the wallpaper on a mobile phone. Thus, each user can gain some idea of other users' current context by looking at the visualisation. Ambiguity in the visualisation means that users can make a wide variety of interpretations based on what they see. Users can also perform intentionally for the interface knowing that others may be watching, and must balance the concerns of both their physically co-located and remote spectators. The MuMo system was evaluated in an "on-the-street" user study where pairs of users interacted with the system in both public and private spaces for two sessions spaced on week apart. The results of this study show what kind of actions users developed *in situ*, how they considered the influence of spectators, and how they interpreted the visualisation.

2 Background

Using a performative perspective on interaction, human actions can be viewed as a performance of some kind where people are constantly adjusting their own behaviours based on the presence (real or imagined) of spectators. Understanding action in this way has interesting implications for designers of interactive systems, where users can be viewed as actors, interaction spaces as stages, and spectators as the audience. Such a performative perspective can be used to leverage this perspective in design and how such performances can be captured.

2.1 Action and Performance

Goffman (1990) describes how everyday life can be understood from a performative perspective, a view that has seen growing popularity in human computer interaction. Goffman describes peoples' everyday behaviours as a performance, where people are constantly adjusting their actions based on feedback from spectators, using places as stages and making use of their appearance and props to support their intended impressions. Goffman describes a wide range of performances, from implicit performances of everyday action and impression management to explicit performances such as giving a formal presentation to an audience. These concepts can be further refined as impressions or performances *given* and *given off*, where impressions *given* relate to those explicit performances and impressions *given off* relate to implicit performances (Goffman, 1990). Implicit performances might be actions that are performed without being explicitly aware of them, but which are unconsciously adjusted constantly throughout the day as feedback is gathered from spectators. More explicit performances carry with them significantly more intention from the performer and more clearly defined performer/spectator roles. Both impressions *given* and *given off* are interesting from a social signal processing point of view, where this performative perspective gives a sociologically motivated approach to organising these behaviours.

2.2 Performative Perspectives in HCI

The concept of interaction as a performance (Jacucci, 2004) provides a way of understanding interaction as the presentation of self and the experience of interacting in front of others. In interactive systems research, this means that performative concepts can be leveraged in design, such as the influence of spectators, users' percieved images of themselves, and narratives within interaction. Reeves *et al.* (2005) describes how the presence of spectators changes how people interact with systems in public places based on the size of their manipulations and the resulting effects. Dalsgaard and Hansen (2008) add to the performative perspective by developing the concept of "performing perception," describing in great detail the experience of performing with respect to the roles users must adopt throughout an interactive experience. Benford *et al.* (2012) describes how traditional narrative structures from theatre can be used to design uncomfortable but rewarding or fulfilling interactions. Each of these examples demonstrate how a performative perspective can be leveraged to inform the design of interactive systems.

2.3 Capturing Performances for Interactive Systems

There are a wide variety of sensors that have been used in activity recognition and social signal processing. Although accelerometers have been widely used in interactive systems, they have seen less action in social signal processing (Vinciarelli *et al.*, 2009). However, there are a several important signals that can be sensed through accelerometers and present interesting opportunities for visualisation and interaction. Crossan *et al.* (2005) demonstrate that accelerometers can be used to sense gait phase during mobile interaction for increased understanding of users' mobile context. This approach has also been used as a means to "instrument" users during evaluations to gain additional data about interaction context in the wild. Microphones have also been used as a mobile form of input in the instrumented usability scenario. Hoggan and Brewster (2010) used a phone's built in microphone to gather data about ambient noise levels to better understand users' current context during an in-the-wild study. However, accelerometers and microphones are not only used for such passive input. Jones *et al.* (2010) showed the possibilities of accelerometers for sophisticated input in a gesture-based text entry system. Scheible *et al.* (2008) created a system where throwing gestures performed on a mobile phone could "toss" content from that phone onto a large public display. Mobile phone sensors like these could be used to capture both actions "given" and "given off" to bring sophisticated social signal processing to a mobile context (Vinciarelli *et al.*, 2010).

3 Exploring Performance for Social Awareness

The study presented in the paper explored how simple behaviours and actions in a mobile context could be used in a remote awareness application. The evaluation explored not only how users interpretated this data but also how they experienced performing and generating this data, particularly when extravagent and exagerated actions were encouraged. During the study, participants were required to generate simple gesture and voice input *in situ* in public and private locations using a mobile remote awareness application with a partner over repeated trials. This application was designed to support divergent multimodal inputs with a high level of flexibility, create the experience of performing in different settings and participate as a distant audience member for a familiar other's performances.

This application, called MuMo, included a visualisation of a virtual fish tank where each user was represented by a fish in the tank that could be controlled using multimodal input. Users generated input by interacting with a small prop embedded with sensors. MuMo was designed to explore the issues of performance and the usage of props when the user was performing for two different audiences: one audience was the fellow participant watching the performance through the fish tank visualisation and the other was the immediately co-located spectators watching the live performance without necessarily being aware of its purpose or the interface itself. This application used highly flexible input methods, where participants were required to create their own performance style in real world locations using gesture and voice. Using this application, users were free to create a variety of performances to suit their current context and could participate as an audience member by watching the visualisation, where divergent imagined interpretations of the visualisation were possible. The possibility of this kind of extravagant performance Jones (2011) creates the opportunity for expression and imagination in real world contexts.

3.1 The MuMo Application

In the MuMo application, participants were each represented by a fish in a virtual fish tank, as shown in Figure 1. This visualisation could be seen as an animated background on each users' mobile phone and controlled using multimodal input. The application used a server/client architecture where each client updated the server with its current input values and pulled updates from every other user from the server roughly once per second. Thus, users could see the effects of their own actions in the visualisation alongside those of their fellow participants. Participants were told they could use gestures or motions to make their fish swim faster or use audio and voice input to make their fish blow more bubbles. In each case, the fish behaviour was based solely on the magnitude of input, although this was not explained to the participants. For audio input, the louder the sound level the more bubbles the corresponding fish would create. Thus, participants could perform any kind of speech or sound-based action and see the result in the fish tank. Similarly, changes in swimming movements were based

Fig. 1. Screenshots of MuMo application as an active wallpaper. Left shows fish tank visualisation as wallpaper, right shows visualisation with phone widgets.

on the magnitude of acceleration of the gesture performed. This type of sensing was designed specifically to support both extravagant and subtle input, meaningful and abstract input, or simply environmental input that could be reflected in the fish tank visualisation in real time. This flexible style of input afforded unconstrained interaction in order to encourage participants to generate creative methods of controlling the visualisation. This also allowed for imaginative interpretations for those watching the visualisation since the observed output in the visualisation could be generated in a variety of ways.

The interface was controlled using the SHAKE sensor pack[1] to collect accelerometer data with an added microphone as shown in Figure 2. This was then embedded into the various objects or props shown in Figure 3. These props were chosen to provide a variety of objects that could facilitate performance or demonstrate interaction in different ways. These included playful objects, an abstract object, an everyday object, and an object that displayed the bare electronics of the sensors. The playful objects included two plush toys and one solid toy in order to allow for enjoyable and playful interactions. The abstract object was a hollow red mould that would simply act to conceal the sensors. The everyday object was a book with a space hollowed out to conceal the sensors in order to disguise the interactive prop. The final prop was a clear glass jar that exposed the bare electronics of the system as a method for demonstrating the interactive purpose of the prop. These props were selected to provide different visual or cognitive clues for spectators about the performance in order to give performers different methods of exaggerating, disguising, or explaining their performance.

[1] More information: `http://code.google.com/p/shake-drivers/`

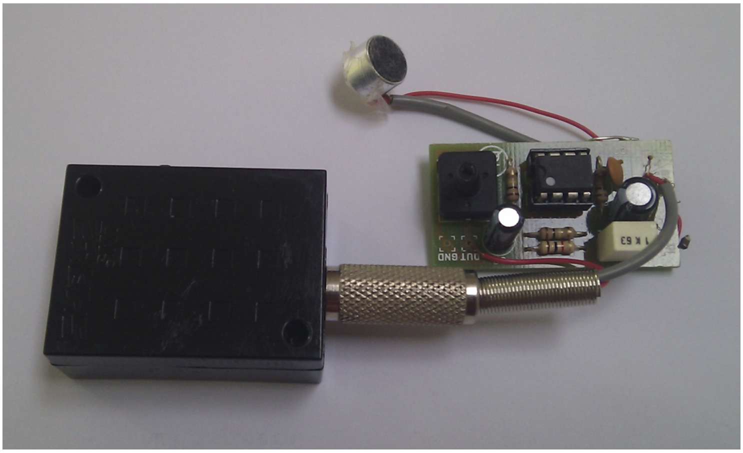

Fig. 2. The SHAKE sensing device with an added external microphone

3.2 The Study

Participants were recruited in pairs, where each pair completed two usage sessions spaced about one week apart. These sessions were repeated to give participants multiple chances to interact with the system and develop performance preferences based on multiple experiences. Before each session, participants were told only that they could control their fish's swimming behaviour using gestures and the bubbling behaviour using sound and were given a chance to briefly experiment with the system. Then, the session began with the first participant being taken to a public location, a busy pavement, while the second remained in a private indoor location. Once both were ready to begin, the first participant was asked to complete three performance tasks, such as creating more bubbles, while the second was asked to interpret the first participant's actions by watching the visualisation on the phone. After these performance tasks were complete, the first participant was then asked to interpret the other's actions while the second participant completed three performance tasks. The participants would then switch locations and the tasks were repeated. Each task lasted two minutes. This study design allowed participants to perform actions in both the public and the private setting as well imagine how their partner would perform actions in both settings. Once both participants had completed their tasks in each location, they were interviewed together about their experiences.

4 Results

The study involved eight participants recruited in pairs. The pairs included two couples and two pairs of friends, with four females and four males. The participants ranged in age from 20 to 28. The results focus on an in depth analysis of a relatively small user group in order to gain a highly personal qualitative

Fig. 3. Participants could select one of six objects containing an embedded sensor pack to control their fish in the tank

insight into the use and experience of this application. These results are based on the observation of the participants, recorded observations provided by the participants and transcripts of the interviews.

4.1 Creating Performances

Given that participants were allowed to create open ended performances using gesture and speech, it is not surprising there were a wide variety of styles and actions that resulted in the different locations where this study was completed. In each case, these actions can be analysed from a performative perspective to better understand how people generate actions and behaviours in this interactive context.

Performative Actions – Even though the sensors were contained solely within the prop, performances were not limited to interactions with that prop and often involved additional interactions purely as an enhancement to the experience and appearance of performing. For example, one participant chose to sing to the prop for voice input in the private indoor setting. Although this was an unnecessarily extravagant interaction, this participant found that performance enjoyable and amusing, especially when his partner imagined this performance. Another participant performed swimming motions with both hands while outdoors. Even though the prop would only sense the movement of one arm in this case, the participant still enjoyed performing with both arms. Perhaps this action better demonstrated the purpose of the participant's actions, where spectators might more easily understand the action of mimicking swimming with both arms. In these cases, the experience of performance was augmented with either playful or meaningful actions to add to the functional aspects of interaction to make interaction more fun, more enjoyable, or more socially acceptable.

Hidden/Subtle Actions – Participants found ways of performing input that were subtle or hidden from passersby while still giving their fellow participant highly visible actions on the visualisation. Because the system was flexible enough to support both extravagant and subtle actions, participants could exploit this to balance their desire to perform for their partner while also considering the immediately co-located spectators around them. The hidden/subtle actions included input such as tapping the prop to make noise, fidgeting with the prop in hand, and using environmental noise to create input. For example, one participant chose to use the music of an outdoor performer as the input for their performance when audio was needed. These types of actions allowed participants to create meaningful input to the system without performing highly visual actions.

Functional Actions – In some cases, participants chose only to perform actions that completed the task without adding any additional performance or play. For example, participants would simply shake or wave the sensor to create gestures or say things like "I'm creating test speech for a system" or "I'm talking into the sensor now to see if something happens." In this case, participants did not try to actively hide or disguise their performance, but instead tried to demonstrate the purpose of the performance clearly by using "test speech" or rigid, purposeful actions. In this approach to impression management, participants aimed to make it clear they were interacting with a system by keeping the phone or prop visible and performing noticeably rigid actions in order to call attention to the action as purposeful input.

These different styles of performance were influenced as much by location as personality. For example, one participant performed purely functional actions while outside and highly performative actions inside. Another participant completed highly performative actions both inside and outside. Yet another participant completed hidden or subtle actions both inside and outside. Because the interface supported a variety of actions, participants were able to change their performance style as needed in order to continue participating and feel comfortable about interaction. These decisions varied between participants, depending on personal preferences and personality. These factors represent an interesting influence on social acceptability that needs further exploration.

4.2 Imagining Others

Because the MuMo application required participants to create their own input, fellow participants watching the interface could not be sure what kinds of actions their partner was performing given the current output. Participants had to imagine how they thought their partner might be performing based on what they could see in the visualisation and their knowledge about their partner's current social context. This was both a positive and a negative aspect of this application, where some participants found it difficult to attach meaning to the interface while others enjoyed the process of imagining their fellow participant performing highly energetic, silly, or emotional behaviours. These imaginings not only contributed to the spectator experience of this application through the visualisation but also provided motivation for participants to generate performative input to the system.

For those participants that enjoyed imagining their partner performing through the interface, participants allowed and encouraged their partner to imagine highly divergent performances, even when this was not realistic or likely. For example, some participants imagined their partners singing or dancing as input for the visualisation even though their partner was in the outdoor setting and it was unlikely they would be singing or dancing there. Even though participants knew such energetic and performative actions where unlikely, participants were able to suspend their disbelief and enjoyed imagining these kinds of actions anyway. These creative imaginings occurred both when pairs of participants used highly visible, performative interactions and when pairs of participants used the most subtle and discreet methods of interaction. For example, one participant imagined her partner "singing a relaxing song" and "jumping with it [turtle] on one leg." These interpretations were recorded even though both participants used extremely subtle actions for input, such as microphone tapping. Participants enjoyed imagining these playful actions, even if they did not perform these kinds of actions themselves.

4.3 Props and Performance

During each of the two sessions, participants could select an object of their choice as their prop. The prop was an important part of the types of actions and behaviours participants would perform because the prop would be highly visible during interaction and could both support and hinder performative actions. For example, a playful prop like a toy might encourage fun interactions because toys are made to be played with while an everyday prop like a book may be more acceptable to carry around in public places but not typically be viewed as an "interactive" object. Of the props including in this study, the turtle object was chosen eight times, the dolphin was chosen five times, the book, jar, and owl were chosen once and the red mould was never chosen. When discussing their choices of these objects, participants described how the objects worked and failed as props.

Props as Toys – The most commonly picked objects were the turtle and dolphin plush toys. Participants favoured these props for their playful nature and their ability to relate to the lively and lighthearted application. These props were often used in a playful manner, even though participants knew that these kinds of actions would not provide any additional input to the application. For example, participants would move the fins of the turtle or cover it's eyes as part of their performance even though this did not generate additional effects.

Props as Pairs – Participants often chose their props based on their partner even though they knew the props would not be used together. Choosing props together allowed participants to better understand what kinds of performances their partner might complete and also provided a better connection between partners. For example, one participant stated that “first I wanted to pick the glass jar, but when I saw he picked a toy I wanted to pick a toy as well.” Another participant stated that “I picked the dolphin because you picked a toy, so it’s two soft toys. Otherwise, I would’ve picked the book.”

Props as Everyday Objects – Although some objects, such as the book, represented common objects one might normally carry around, participants felt less comfortable using these objects when interacting with the application. While using the book as a prop, one participant stated that “when I was inside I sang a song, I just made it up. But when I was outside I tried to talk very quietly. It wasn’t as normal as I thought it would be.” When discussing other everyday objects that might be used as props, one participant stated that “you might put it [sensors] into an object that you walk around with, like a coffee cup, but you wouldn’t talk into a coffee cup.” Although these props might disguise or hide sensors effectively, they make poor interactive objects when it comes to performance.

Participants also discussed the benefits of different props with respect to physical attributes like size or texture. For example, when describing why the dolphin was a useful prop, one participant stated that “it’s easier to hold than one of the hard objects, nicer to hold.” When describing objects that would make the most desirable props, participants stressed the importance of using soft or flexible objects. The ability to manipulate the props and the comfort of holding a soft object made them easier to use. Participants also described the benefits of using different props to conceal the sensors. When describing why a prop would be better than simply holding the sensor pack, one participant stated that “it’s bigger, so there’s more you can do with it.” Participants also described how the prop makes performance more comfortable. For example, one participant stated that ”it was much easier to just wave around the turtle than it would’ve been to wave a bunch of sensors”. Other participants would have preferred a more anonymous object. When discussing negative aspects of using props, one participant stated that “it made me more conscious of it, holding the object. If I just had the sensor in my hand people might not have noticed what I was doing.” Because this application clearly had a playful nature, participants often chose props that encouraged this playfulness. However, props that are more abstract or anonymous were still desirable and in a different application area might have been more popular.

5 Discussion

This study provides some interesting insights into the ways in which these participants created performances in the wild, used props to enhance their interactions and demonstrated their intentions to co-located spectators. By performing through the MuMo interface, participants were performing for the immediately co-located spectators as well as their fellow participant watching their actions through the visualisation. Thus, participants in this study were constantly performing for two audiences and had to balance the needs and expectations of these spectators simultaneously. For example, participants had to balance their desire to generate energetic or amusing input for their fellow participant with their desire to perform socially acceptable interactions in public places. In some cases, this meant that participants chose to limit their performance and the resulting output of the system, limiting the spectator experience for their fellow participant. In other cases, participants found ways of performing that were both comfortable for themselves and created ample output in the application for their fellow participant to enjoy.

Because this application required only basic actions but also supported extravagant ones, participants took full advantage of this flexibility and generated a wide variety of behaviours and actions through the system. The types of performances created were highly dependant on the location of the performance, with participants actively making decisions about their adoption of different performances based on their current location. In general, participants were more likely to perform highly visible or noticeable actions in the indoor location as compared to the outdoor location, which is in line with the results discussed in the previous chapter. Additionally, participants often adjusted their performances to match their fellow participant. Because the first session ended with an interview, participants learned what kind of actions their fellow participants had imagined them doing and what actions their fellow participant had actually performed during the first session. This was reflected in the second session where pairs of participants performed actions that were discussed during the first session. This included actions that might be amusing to their fellow participant or actions they thought the other participant might be performing as well. This demonstrates how social influence can affect adoption, even though this example is on a very small scale. For example, usage over time might allow constantly evolving practices and behaviours as the users of the application respond to each other and learn how to interpret the visualisation based on their knowledge of each other. Interpretations that come out of familiarity and extended use of an ambient display have been seen before Brewer *et al.* (2007), and certainly this emerging behaviour is an important aspect of these types of applications and how people might make use of sksocial signals in their everyday lives over time.

Participants' awareness of their partner watching the visualisation provided motivation for participants to perform amusing actions but also led participants to perform extremely subtle actions and simply allow or encourage their fellow participant to imagine more entertaining actions. Pairs of participants had varying degrees of enjoyment imagining the performance of their fellow participant, with the two couple pairs being the most imaginative. Even when both participants performed subtle actions in the outdoor settings, both participants enjoyed imagining amusing performances. Although these imaginings provided some motivation to perform amusing actions, these participants were still highly aware of the co-located spectators, or passersby. In some cases, participants modified their performance when outside. For example, one participant used singing input while inside and conversational speech while outside. Both of these actions generated similar output in the visualisation, but participants used these different kinds of actions in order to maintain their comfort, experience, and enjoyment of the application. These adjustments show how considerations for both audiences must be balanced while using this application in public contexts.

6 Conclusion

The user study presented in this paper explored how participants generated and interpreted basic social actions and behaviour in real world settings. This involved using a remote awareness visualisation on a mobile device that could be controlled with gesture or voice based input. During the study, participants demonstrated three methods for generating multimodal output for the visualisation. Participants used highly performative actions, hidden or subtle actions, and simply functional actions when generating input for the system. Because the system supported both extravagant and subtle input equally, participants could perform a wide variety of actions as input and adjust these actions fluidly based on their current context. The variety of possible actions and the purposeful ambiguity in the visualisation also meant that participants could interpret the visualisation in many different ways, incorporating their knowledge of their fellow participant's personality, current context, and previous actions and inputs. These results demonstrate an interesting scenario for making use of basic social signals as part of a remote awareness application.

References

Benford, S., Greenhalgh, C., Giannachi, G., Walker, B., Marshall, J., Rodden, T.: Uncomfortable interactions. In: Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems, CHI 2012, pp. 2005–2014. ACM, New York (2012)

Brewer, J., Williams, A., Dourish, P.: A handle on what's going on: combining tangible interfaces and ambient displays for collaborative groups. In: Proceedings of the 1st International Conference on Tangible and Embedded Interaction, TEI 2007, pp. 3–10. ACM, New York (2007)

Crossan, A., Murray-Smith, R., Brewster, S., Kelly, J., Musizza, B.: Gait phase effects in mobile interaction. In: CHI 2005 Extended Abstracts on Human Factors in Computing Systems, CHI EA 2005, pp. 1312–1315. ACM, New York (2005)

Crossan, A., Williamson, J., Brewster, S., Murray-Smith, R.: Wrist rotation for interaction in mobile contexts. In: Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI 2008, pp. 435–438. ACM, New York (2008)

Dalsgaard, P., Hansen, L.K.: Performing perception—staging aesthetics of interaction. ACM Trans. Comput.-Hum. Interact. 15, 13:1–13:33 (2008)

Dey, A.K., de Guzman, E.: From awareness to connectedness: the design and deployment of presence displays. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2006, pp. 899–908. ACM, New York (2006)

Goffman, E.: Behavior in public places: notes on the social organization of gatherings. Free press paperback. Free Press (1966)

Goffman, E.: The presentation of self in everyday life. Penguin Psychology. Penguin (1990)

Hoggan, E., Brewster, S.A.: Crosstrainer: testing the use of multimodal interfaces in situ. In: Proceedings of the 28th International Conference on Human Factors in Computing Systems, CHI 2010, pp. 333–342. ACM, New York (2010)

Jacucci, G.: Interaction as Performance. Cases of configuring physical interfaces in mixed media. PhD thesis, University of Oulu (2004)

Jones, E., Alexander, J., Andreou, A., Irani, P., Subramanian, S.: Gestext: accelerometer-based gestural text-entry systems. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2010, pp. 2173–2182. ACM, New York (2010)

Jones, M.: Journeying toward extravagant, expressive, place-based computing. Interactions 18, 26–31 (2011)

Perry, M., Beckett, S., O'Hara, K., Subramanian, S.: Wavewindow: public, performative gestural interaction. In: ACM International Conference on Interactive Tabletops and Surfaces, ITS 2010, pp. 109–112. ACM, New York (2010)

Reeves, S., Benford, S., O'Malley, C., Fraser, M.: Designing the spectator experience. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2005, pp. 741–750. ACM, New York (2005)

Scheible, J., Ojala, T., Coulton, P.: Mobitoss: a novel gesture based interface for creating and sharing mobile multimedia art on large public displays. In: Proceedings of the 16th ACM International Conference on Multimedia, MM 2008, pp. 957–960. ACM, New York (2008)

Sheridan, J., Bryan-Kinns, N., Reeves, S., Marshall, J., Lane, G.: Graffito: crowd-based performative interaction at festivals. In: Proceedings of the 2011 Annual Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA 2011, pp. 1129–1134. ACM, New York (2011)

Strachan, S., Murray-Smith, R., O'Modhrain, S.: Bodyspace: inferring body pose for natural control of a music player. In: CHI 2007 Extended Abstracts on Human Factors in Computing Systems, CHI EA 2007, pp. 2001–2006. ACM, New York (2007)

Vinciarelli, A., Murray-Smith, R., Bourlard, H.: Mobile social signal processing: vision and research issues. In: Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI 2010, pp. 513–516. ACM, New York (2010)

Vinciarelli, A., Pantic, M., Bourlard, H.: Social signal processing: Survey of an emerging domain. Image Vision Comput. 27(12), 1743–1759 (2009)