Research Techniques for Augmented Reality Experiences

Stuart Reeves

str@cs.nott.ac.uk

May 11, 2004

Abstract. We contrast the methods and techniques used by Augmented Reality and museum studies, and discuss the history and development of research in both fields. With this in mind, we follow with a more detailed look at the techniques involved in this program of research — ethnography in particular — and the associated theoretical considerations. The practical implications of future research are then considered with reference to previous work, and finally a discussion of the validity and target of analysis.

Social Interaction at the Exhibit Face

Research conducted in museums varies wildly in approach, but is foremostly concerned with the evaluation of existing exhibitions and collections rather than the development of new displays to study. Directly in contrast to this is research into Augmented Reality (AR), which typically consists of the construction of experimental systems, but with little detailed reflection on the process by which people use these systems. In this section we shall briefly examine the history and approaches used by these two areas.

Museum Studies

There is a substantial body of museum studies literature which covers aspects of the material environment, visitor behaviour ('visitor studies'), collaboration and learning in museums. These studies range from very early investigations by psychologists such as Melton and Robinson in the 1920s and 1930s, to other psychologically or sociologically informed research, such as behaviourist [30], cognitive models [1] and environmental psychology [4], and ethnography [24]; however, this set of highly varied approaches do not form a strict time line. For example, the more quantitative methods typically found in earlier visitor studies — using measures such as 'dwell time,' number of objects viewed and routes traced during visits, in order to assess the quality of each exhibit — are still prevalent in some museum journals. Lawrence has provided a detailed history of the tensions and conflicts between many of these methods and background epistemologies [23], and as such an in depth exposition of the history of museum research would not be useful here. She noted, however, the positivism that was still present in much museum research in 1991:

> Useful as typical museum evaluations based on behavioural, cognitive or survey work may be to administrators, we may wish to think extremely carefully about whether or not they really do illuminate the process by which meaning is produced in our museums, or whether they rather serve to perpetuate, largely by ignoring them, society's obfuscations about cultural process. [23]

What Lawrence claims is at issue, then, is the prevalence of reductionistically focussing on only one aspect of the museum ecology, a claim which is reflected by various other researchers' assessments of the state of museum studies. It has been remarked, for instance, that the meaning interpreted by visitors is "neither in the eye of the beholder nor in the structure of the display itself or the situation of the encounter between the two, but in a lively triadic interplay of all three elements." [31] This reflects the emerging notion that museum studies often do not account for the emphasis placed discretely on visitors or objects, and therefore implicitly favour either the environment *or* the visitor exclusively; as a result, the examination of the intrinsic sociality of the museum visit suffered. Blud noted that previous and contemporary studies "have tended to ignore the nature of the visitor, or *visitor group*, and have focused instead on the nature of the exhibit, and how effective different types of exhibit are in stimulating learning." [5] (Emphasis added.) Similarly to Blud, McManus also comments on the importance of social groupings and conduct as having a bearing on learning, speaking here of the separation created between cognitive and behaviourist research, revealing the need for a "model of learning in science museums based on ways that people think *and* behave" [25]. Learning research itself also indentifies the fundamentally social aspects of exhibitions that this separation obfuscates. Crowley and Galco, in their studies of scientific learning at museum exhibits state [7]:

> The complexity, extended time-scale, and socially-embedded nature of *in vivo* learning and development are not just annoyances to be controlled by an experimenter; they are fundamental, irreducible characteristics of how learning actually occurs.

It is the ongoing, in situ and social construction of meaning which naturalistic museum studies, such as those of vom Lehn et al [33, 32], have begun to expose. This work has emphasised the vitally missing aspect of many investigations that have drawn conclusions primarily from survey data, inevitably coupled — as previously mentioned — with behaviourist or cognitive models, for example. Such observations can expose certain features of visitors' interaction with the artefact that is the focus of their attention, and to some extent an approximate measure of attributes such as the 'attractiveness' of a particular exhibit, or the effect the layout of an exhibition space might have. However, the resulting analysis omits central elements, due to its concentrating on individuals or the material environment, but typically disregarding both the ongoing accomplishments of visitors and the effect that co-visitors and other visitors have on configuring exhibit participancy [16, 33]. By drawing attention to the bodily conduct (in addition to conversations) of visitors through the analysis of video data, the emerging organisation of the museum visit may be examined and therefore inform the design and construction of exhibits. The process by which this analysis takes place is informed by ethnomethodology and conversation analysis rather than any particular method (a discussion of this appears later). The use of video data as a primary resource contrasts with much previous research — including many ethnographic studies — which has drawn out conclusions from audio recordings and field observations [32].

Augmented Reality in Museums

Whilst museum studies typically centre on evaluation, Augmented Reality research is usually concerned with the development of novel technologies. It is often noted that such technologies could be applied to museum settings, however there is frequently little regard for the impact such devices have on the social ecology of exhibitions, and often only the most cursory evaluation of such technology if any observations are present at all. Studies in Computer-Supported Cooperative Work (CSCW) are targeted towards understanding collaboration in many domains (often informed by ethnomethodological techniques of analysis), and although there is an increasing corpus of research devoted to studying innovative technologies in the social context of museums [12, 20, 36, 11, 14, 34], other public experiences [10, 9] and usability studies [29, 3, 2], little literature exists within CSCW that examines Augmented Reality systems in particular.

The insertion of novel technologies into exhibits is fraught with problems both of misunderstanding how technology can be used and interactional issues arising from the introduction of pieces of equipment in an exhibition space. Often we see the placing of virtually 'naked' (rather than in-context) computer systems in museums which have been observed to configure the visitors for the same individual-to-artefact interaction that would be expected in computer users in an office. The natural collaborative ecology engendered by the majority of exhibits is clearly not assisted by this effect, even though visitors have been found to attempt to organise collaboration in spite of these hindrances [19]. Technology integrated into exhibitions, therefore, must take into account both the context and the aesthetic.

The focus and study that such systems require is dependent on this role the system is stationed in for the exhibition. For AR systems which form a component part of the exhibit, it is vitally important that the context of the environment be compromised as little as possible by the insertion of such technologies. In designing such devices that form an integral part of an experience, this respect for the local aesthetic must become central to the design process, and may become a powerful influence over technical decisions, such as registration issues (discussed in [26]). Conversely, for AR systems which occupy a more conventional, auxiliary function in the exhibition space, such as providing a guidance role (e.g., Sotto Voce [2]), design considerations might be more appropriately focussed on usability issues, for instance, rather than attempts to shroud the technology because of its place in the exhibit context. Here we can see that, when the aesthetic and contextual issues are accounted for, a device may be subsumed into the environment to become part of an exhibit, or, for the opposite extreme, a device provides an 'annotation' to the public display, distinctly separate from the content of the display. This place occupied by Augmented Reality systems is often challenged by the aesthetic hurdles imposed by technology itself. AR systems conventionally feature Head Mounted Displays (HMDs) or PDAs, and less frequently, specialist content overlay systems such as half-silvered screens. It is usual in AR research for the overlaid, digital content to take precedence over the aesthetic considerations both of the raw device and the target of the augmentation. It has been demonstrated in our research that the analysis of social interaction at the exhibit face has revealed, at least for our device, this broader nature of AR that encompasses not only the content delivered, but the augmented target and the augmenting device [26].

Evaluating Augmented Reality Experiences

Given this body of evaluation present in the museum and CSCW literature, what, then, is an appropriate way to evaluate AR experiences? In this section we will discuss the varied approaches, their shortcomings, and how they are to be collated into a coherent picture under the general theme of ethnography. Ethnography is perhaps the most appropriate technique for museum-based research that plunges technology 'in the wild.' We shall discuss this matter, however, in the conclusion.

Approaches, Techniques and Methods

Taking a cue, then, from the developments of the sociological literature on interaction in museums, and the existing research on technology in public exhibitions, we are faced with a cluster of ethnographic techniques and methods which may be employed in the study of public exhibitions. We shall assess several of these in turn, with reference to their use in evaluating technologies.

Ethnography and Ethnomethodology. Ethnographic approaches have, especially in the field of CSCW, brought about certain amounts of methodological freedom from the trappings of purely cognitive theories of human-computer interaction. Such cognitive theories may be particularly informative when studying certain highly constrained forms of single user human-computer interaction,[1] however people predominantly (and maybe instinctively) collaborate using technology, both in co-located (e.g., use of a shared display) and distributed (e.g., groupware over networks) situations, synchronously and asynchronously. Typically, the nature of these collaborative setups precludes controlled experimentation. It is clear, then, that the analysis of this interaction and collaboration hinges upon ability to examine this conduct in real world environments, be they public exhibitions or computer labs.

Although ethnography does not subscribe to any particular methodological predisposition, certain philosophical underpinnings influence the collection and investigation of data. Ethnomethodology — which is a 'sociological commitment' rather than a specific method — informs the process of analysing ethnographic data, such that interactions amongst people and computers are seen as a process described by several ethnomethodological concepts. Ethnomethodology (and therefore conversation analysis and video-based interaction analysis; see later) assume that all social interaction is, at a fundamental level, structurally organised, and that the study of this interaction should therefore reveal the emergent organisation that is presupposed to be inherent. The program of ethnomethodology reveals this organisation by, for instance, confounding the "background expectancies" present in social interaction. These expectancies are laid bare by performing a "breach" of this orderliness, exemplified by Garfinkel's account of his students pretending to be lodgers in their own home [13]. This perspective provides some unique insights into what is described as a 'moment-by-moment,' 'indexical' and 'reflexive' organisation of the structure of social interaction, that, of course, includes collaborative activities involving technology. The result of such analysis seeks to present a set of generalised, abstract forms of organisation that have emerged from the data in multiple locations. In practical terms for the concerns of Human-Computer Interaction (HCI), this often translates to the issues surrounding reconciling differing perspectives on digital content, and the building of a 'common stock of knowledge'

[1] It is not the intention to portray the study of cognition as unworkable, but rather as inappropriate in particular situations.

related to the technology, such as those demonstrated in [6].

The collection of information that any detailed ethnography draws from is often as broad as possible. Data such as recorded interviews (typically open-ended or semi-structured) and more casual conversations, video footage, copies of internal and external paperwork (e.g., publicity material) and photographs — amongst other data elements — form a rich set that the ethnographer may draw from. It is the job of the ethnographer to collate these often highly disparate collections and assemble an interpretation of how that data may be ordered into documentation of the work practices that emerge from the observations. Any ethnographer is, of course, a filter to this data, selectively retaining or discarding information, and in this respect all ethnography should be treated with care.

Video-Based Interaction Analysis. Ethnographic video-based interaction analysis — which is heavily influenced by ethnomethodology and conversation analysis — provides an ideal framework for studying both captured talk and bodily conduct [16]. Although such recordings only present a version of the events and are subject to what the camera operator deems significant, video data is repeatable and helps reduce the human 'filtering' that inevitably takes place during the gathering of field notes. Video analysis "looks for orderliness and patterns in people's routine interactions, but operates at a finer level of detail than conventional ethnographic observation." [27] In this way, the work practices are allowed to emerge from the interaction repeated by the recording. In terms of interaction with technology, video analysis can reveal how screens, audio channels and perceptual awareness of the use of technology, for example, may feature in interaction [18, 17].

Being concerned with the methods by which people conduct their interaction means that the documenting of video data cannot have partiality to any particular method of encoding. The analysis aims to be set in the context of the recording itself, describing (in contextual language) the themes of conduct and interaction emerging across multiple instances. Nevertheless, the same care required by the recording and interpretation of field observations should also be exercised with video data, as with any ethnographic resources.

Video analysis brings with it other more practical considerations which, although relatively banal, may have great impact on the data set obtained. Permission, the privacy of those observed, camera positionings and background work must all be considered. Jordan and Henderson describe excellently the many facets to the application of video-based interaction analysis [22], and a detailed discussion of this would merely reiterate points previously made.

Discourse. Recorded conversation from video provides a rich complimentary resource to bodily interaction. Techniques for encoding and analysing discourse vary reflexively with exactly what extracted information is desired. It may be appropriate, for instance, to encode talk by characterising phrases into categories assigned with different meanings, and then performing statistical analysis. Such methods handle large quantities of data well, however, talk captured from video will often require analysis that fits with the "finer level of detail" of repeated, short audiovisual segments. A conversation analytic approach compliments the observation of the methods used in the organisation of physical interaction that is the target of video analysis techniques.

Conversation analysis provides a framework for studying the emergent organisation of talk and the methods used by speakers to engage in social interaction, and may be used to inform design as well as providing a particular approach to evaluation. De-

signing, for instance, a system that respects the organisational storytelling elements of Sacks's *preface*, *telling* and *response sequence*[2] [28] that occur in everyday talk, a phenomena particularly pertinent in museums and other public experiences, has been observed to assist collaboration [35]. Woodruff et al note that:

> ...it should be remembered that CA [Conversation Analysis] is about human-human interaction. We do not try to design computer systems that purport to interact with humans according to the "rules" of CA ... Rather, we design systems to operate in a manner that reflects human practices that are likely to arise in particular situations. [35]

Shortcomings

There are some distinctly lacking features when only observing interaction embodied in conversation and movement. The structure of overarching practices, and the reasons behind particularly obscure actions undertaken by those being observed may not be apparent without consultation of more orthodox ethnographic data sources. Extra clarification from interviews, open-ended questionnaires and workplace documents will assist the development of a broader picture and context for interaction. Particularly important to video may be elucidating certain actions and exchanges by presenting the participants with the video data of themselves and requesting some form of commentary.

It is important to realise that these extra pieces of data are not merely enhancements to the video data, but central parts to the assembly of a broad view of the interaction that has taken place. Furthermore, for studies which centre around the use of technology, the data set is widened by the fact that computers are information processing machines. Detailed log files of the digital side of interaction often prove to be invaluable when constructing a clear picture of the situation. The evaluation of performance art such as *Can You See Me Now?* [10] and *Uncle Roy All Around You* [9] illustrates well this pluralist approach to reflection. Informed by ethnographic reports [6], sensor logs, interviews, radio logs and player text messages, the resulting analysis drew conclusions that might not have been possible without a rich data set. Since the attributes of performance art are often reflected in museum exhibits, it seems sensible to formulate similar methods of evaluation. (Historically, artists have progressively blurred the boundaries between exhibitions, interactive and performance art [21].)

Practical Implications

We consider here some of the practical implications that currently undertaken research has on the probable path future work will follow, with reference to some of the lessons learned from the One Rock project [26].

[2]The methods used by conversationalists are characterised by Sacks as an ongoing accomplishment during the course of talk. *Turn taking machinery* describes the methods used to negotiate order in conversation such that some basic implicit agreements of talk are satisfied, for example, that only one person should speak at a time, or the correct demarcation of the end of a 'turn' of talk. The preface of a telling consists of the negotiation of the commencement of a telling of a particular teller's story with co-conversationalists. A telling is the section that contains the actual story (potentially with interruptions and subsequent embedded talk). Finally, a response sequence represents any ensuing talk (or lack thereof) between teller and co-conversationalists.

Deriving Design from Frameworks and Guidelines

The aesthetic design of a system's physical interface, and by inference, elements of its internal construction, may be informed by frameworks that provide guidelines for the design space, such as what is "sensible, sensable or desirable" [3]. More specific to Augmented Reality, our analysis of One Rock demonstrated that AR systems typically background the target of the augmentation and the augmenting device in favour of the content that will be overlaid. Assessing carefully the relative weights the design of an AR system (and the context in which it resides) gives to these three considerations may help one uncover what is *appropriate*. For example, it was not appropriate that the AR device in One Rock use a fiducial marker form of registration[3], since the aesthetic context of the exhibition precluded this. Furthermore, the strict single user access provided by the AR device was appropriately employed, whereas for systems such as *Sotto Voce*, sharing content was key.

Microevaluations

Although the development of Augmented Reality systems can be greatly informed by ethnographies and related studies, it is not always practical or permitted by time to create AR systems and only evaluate after the event. It may be instructive to conduct some 'pre-test' *microevaluations*, at least to provide a basic trajectory of development. Ideally, such a microevaluation could be conducted in situ (e.g., in the exhibition), but to a limited set of participants (and therefore not technically 'in the wild'). This may, however, not always be possible. Lab-based microevaluation may at the very least provide some indications of a design direction to pursue, or highlight easily overlooked issues. For example, during the One Rock project, it was discovered that an electronic compass that was being used was affected by magnetic variations in the exhibition space.

The methods to assess criteria like usability may require techniques such as video analysis or observation of interaction with a prototype system. The way users respond to and are configured by particular interface paradigms is a vast factor for this interactivity, and depends greatly on the centricity with which an interface is thought of [15]. Designing AR interfaces that are not just restricted to what Computer Science conventionally considers to be 'the interface' — namely input/output devices — means that museum curators, labels and those maintaining the system become subsumed into the interface. It becomes more apparent, then, that even though miniature evaluations before actual implementation are desirable, it is difficult to predict how a particular system will fare if the context is not present.

Record and Replay

The study of any computing technology has, in particular, one major boon. Storing log data can potentially provide incredible resources with which to construct a richer understanding of the computer's 'perspective' on a given interaction. From this power, however, emerges a subtle matter worthy of consideration: what log data will be needed for analysis? Instinctively, it would appear sensible to unconditionally log everything. Even if this is a feasible strategy in terms of bandwidth, there are often states which may not be possible to record without extra measures. Equally, analysis of reconstructed

[3]Fiducial markers allow, through the use of vision techniques, the three dimensional registration of the system observing those markers.

system logs may require that the simulation be run at variable speed (i.e., faster or slower against the 'real-time' of the original system's execution), or even backwards. Is is beneficial — but not always practically possible[4] — to create a system architecture that supports this nonlinear traversal of the time line.

Drawing again on the analysis of One Rock, log data capture was performed, however some vital system state elements were not recorded, such that the reconstruction created from a session of use could not convey whether the users saw content which was affected by lighting in the exhibition. Compounding this was the fact that the AR system featured modified streams of video captured from the front of device, and therefore recording this data was unfeasible. In spite of this, however, the synchronising of recorded video footage (used for interaction analysis) with this reconstructed, simulated view of the device was vital in creating a larger, richly detailed picture of visitor interaction.

Conclusion

As we have detailed, the history of techniques for analysis used in museums and other public experiences must inform the selection of the methods and approaches that shall be used for future study. Public experience scenarios are, by nature, 'in the wild,' and as such will be favoured by the philosophical principles of ethnography, namely that of immersing the study in terms of its environment. Augmented Reality, situated in the context of an exhibit, would be subject to the same 'attack on data' experienced by traditional museum studies if AR research were to rely on those orthodox measures of exhibit success. As the methods described are fiercely qualitative, the most obvious alternative to this is quantitative measurement and a more 'experimental' test design. This method would presuppose the focus of study, and in doing so would preclude the more free detection of emerging phenomena that are seen to be associated with complex public settings. Quantitative study should be used in according measure to the level of appropriate and legitimate control achievable over the environment, and would struggle to provide a similar level of granularity for analysis of bodily conduct, simply because, although conduct *could* be encoded, the resulting conclusions would not address so thoroughly the ongoing context of such actions due to its reductionist approach.

Ethnographic studies of ethnomethodological basis typically seek to detect, characterise and draw general themes from the emergent phenomena that arise from interaction. If we believe that the assumptions of ethnomethodology and related subdisciplines are reasonable, then it follows that general themes drawn out from analysis will be useful in understanding and supporting, justifying or discouraging certain facets of the design of Augmented Reality experiences.

References

[1] M. B. Alt. Evaluating didactic exhibits: A critical look at Shettel's work. *Curator*, 20(3):241–258, 1977.

[2] P. M. Aoki, R. E. Grinter, A. Hurst, M. H. Szymanski, J. D. Thornton, and A. Woodruff. Sotto Voce: exploring the interplay of conversation and mobile

[4] True computational reversibility would require computers to be structured in a fundamentally different way, either using energy lost as an informational resource, or restructured computer organisation [8].

audio spaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 431–438. ACM Press, 2002.

[3] S. Benford, H. Schnädelbach, B. Koleva, B. Gaver, A. Schmidt, A. Boucher, A. Steed, R. Anastasi, C. Greenhalgh, T. Rodden, and H. Gellerson. Sensible, sensable and desirable: a framework for designing physical interfaces. Technical Report EQUATOR-03-003, School of Computer Science & IT, Nottingham University, 2003.

[4] S. Bitgood. Environmental psychology in museums, zoos, and other exhibition centers. In R. B. Bechtel and A. Churchman, editors, *Handbook of Environmental Psychology*, pages 461–480. John Wiley & Sons, 2002.

[5] L. M. Blud. Social interaction and learning among family groups visiting a museum. *Museum Mangement and Curatorship*, 9(1):43–51, 1990.

[6] A. Crabtree. Informing the evaluation of *Can You See Me Now?* in Rotterdam: Runners' and control room work. Technical Report Equator-03-004, School of Computer Science & IT, Nottingham University, May 2003.

[7] K. Crowley and J. Galco. Everyday activity and the development of scientific thinking. In K. Crowley, C. D. Schunn, and T. Okada, editors, *Designing for Science: Implications from Everyday, Classroom, and Professional Settings*, chapter 13, pages 393–413. Lawrence Erlbaum Associates, 2001.

[8] R. P. Feynman. *Feynman Lectures on Computation*, chapter 5: Reversible Computation and the Thermodynamics of Computing, pages 137–184. Perseus Books, 1996.

[9] M. Flintham, R. Anastasi, S. Benford, A. Drozd, J. Mathrick, D. Rowland, N. Tandavanitj, M. Adams, J. Row-Farr, A. Oldroyd, and J. Sutton. *Uncle Roy All Around You*: Mixing games and theatre on the city streets. In *Proceedings of Level Up: The First International Conference of the Digital Games Research Association (DIGRA)*, November 2003.

[10] M. Flintham, S. Benford, R. Anastasi, T. Hemmings, A. Crabtree, C. Greenhalgh, N. Tandavanitj, M. Adams, and J. Row-Farr. Where on-line meets on the streets: experiences with mobile mixed reality games. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 569–576. ACM Press, 2003.

[11] M. Fraser, J. Bowers, P. Brundell, C. O'Malley, S. Reeves, S. Benford, L. Ciolfi, K. Ferris, P. Gallagher, T. Hall, L. Bannon, G. Taxén, and S. O. Hellström. Retracing the past: Mixing realities in museum settings. To appear in to ACE 2004, 2004.

[12] M. Fraser, D. Stanton, K. H. Ng, S. Benford, C. O'Malley, J. Bowers, G. Taxén, K. Ferris, and J. Hindmarsh. Assembling history: Achieving coherent experiences with diverse technologies. In *Proceedings of European Conference on Computer Supported Cooperative Work (ECSCW)*, pages 179–198. Oulu University Press, 2003.

[13] H. Garfinkel. *Studies in Ethnomethodology*, chapter 2: Studies of the Routine Grounds of Everyday Activities, pages 35–75. Prentice-Hall Press, 1967.

[14] R. E. Grinter, P. M. Aoki, M. H. Szymanski, J. D. Thornton, A. Woodruff, and A. Hurst. Revisiting the visit: understanding how technology can shape the museum visit. In *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW)*, pages 146–155. ACM Press, 2002.

[15] J. Grudin. Interface. In *Proceedings of the 1990 ACM Conference on Computer Supported Cooperative work (CSCW)*, pages 269–278. ACM Press, 1990.

[16] C. Heath and J. Hindmarsh. Analysing interaction: Video, ethnography and situated conduct. In T. May, editor, *Qualitative Research in Practice*, pages 99–121. Sage, 2002.

[17] C. Heath, J. Hindmarsh, and P. Luff. Interaction in isolation: The dislocated world of the London Underground train driver. *Sociology*, 33(3):555–575, 1999.

[18] C. Heath and P. K. Luff. Collaboration and control: Crisis management and multimedia technology in london underground line control rooms. *Journal of Computer Supported Cooperative Work*, 1(1–2):69–94, 1992.

[19] C. Heath and D. vom Lehn. Misconstruing interactivity. In M. Hinton, editor, *Interactive Learning in Museums of Art and Design*. Victoria and Albert Museum, May 2003.

[20] J. Hindmarsh, C. Heath, D. vom Lehn, and J. Cleverly. Creating assemblies: Aboard the ghost ship. In *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW)*, pages 156–165. New York: ACM Press, 2002.

[21] E. Huhtamo. On the origins of the virtual museum. In *Nobel Symposium (NS120) on Virtual Museums and Public Understanding of Science and Culture*, May 2002.

[22] B. Jordan and A. Henderson. Interaction analysis: Foundations and practice. *Journal of the Learning Sciences*, 4(1):39–103, 1995.

[23] G. Lawrence. Rats, street gangs and culture: evaluation in museums. In G. Kavanagh, editor, *Museum Languages: Objects and Texts*, pages 9–32. Leicester University Press, 1991.

[24] S. MacDonald. *Behind the Scenes at the Science Museum*. Berg, 2002.

[25] P. M. McManus. It's the company you keep ...the social determination of learning-related behaviour in a science museum. *The International Journal of Museum Management and Curatorship*, 6:263–270, 1987.

[26] S. Reeves, M. Fraser, H. Schnädelbach, C. O'Malley, and S. Benford. Content, target and device: Exhibiting augmented reality. Submitted to CSCW, 2004.

[27] K. Ruhleder and B. Jordan. Capturing complex, distributed activities: Video-based interaction analysis as a component of workplace ethnography. In A. S. Lee, J. Liebenau, and J. I. DeGross, editors, *Information Systems and Qualitative Research*, pages 246–275. Chapman and Hall, 1997.

[28] H. Sacks. An analysis of the course of a joke's telling in conversation. In R. Bauman and J. Sherzer, editors, *Explorations in the Ethnography of Speaking*, pages 337–353. Cambridge University Press, 1974.

[29] H. Schnädelbach, B. Koleva, M. Flintham, M. Fraser, S. Izadi, P. Chandler, M. Foster, S. Benford, C. Greenhalgh, and T. Rodden. The Augurscope: A mixed reality interface for outdoors. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 9–16. ACM Press, 2002.

[30] H. H. Shettel. Exhibits: Art form or educational medium? *Museum News*, 52(1):32–41, 1973.

[31] J. Umiker-Sebeok. Behavior in a museum: A semio-cognitive approach to museum consumption experiences. *Signifying Behavior: A Journal of Research in Semiotics, Communication Theory and Cognitive Science*, 1(1):52–100, 1994.

[32] D. vom Lehn, C. Heath, and J. Hindmarsh. Exhibiting interaction: Conduct and collaboration in museums and galleries. *Symbolic Interaction*, 24(2):189–216, 2001.

[33] D. vom Lehn, C. Heath, and H. Knoblauch. Configuring exhibits. In H. Knoblauch and H. Kotthoff, editors, *Verbal Art Across Cultures: The aesthetics and proto-aesthetics of communication*, pages 281–297. Gunter Narr Verlag Tubingen, 2001.

[34] R. Wakkary, M. Hatala, K. Newby, D. Evernden, D. Leo, and G. Mah. Echoing the conversational space of museums through audio augmented reality and adaptive information retrieval. In *Submitted to CHI04*, 2004.

[35] A. Woodruff and P. M. Aoki. Conversation analysis and the user experience. In *Proceedings of Workshop on Exploring Experience Methods Across Disciplines, ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2004)*, April 2004.

[36] A. Woodruff, P. M. Aoki, A. Hurst, and M. H. Szymanski. Electronic guidebooks and visitor attention. In *Proceedings of International Cultural Heritage Informatics Meeting (ICHIM)*, volume 1, pages 437–454, september 2001.