

Designing for Interpretation

Phoebe Sengers

Information Science and
Science & Technology Studies
Cornell University
301 College Ave.
Ithaca, NY 14850 USA
sengers@cs.cornell.edu

Bill Gaver

Interaction Design
Royal College of Art
Kensington Gore
London SW7 2EU, UK
w.gaver@rca.ac.uk

Abstract

HCI strategies for designing and evaluating systems rest on philosophical stances about how meaning is conveyed through systems. Traditionally, HCI has often conceived of systems as one-way vehicles for conveying a designer's meaning to users. Designing a system means deciding what its meanings should be, and evaluating a system consists, in part, of determining the extent to which those meanings arrive intact to the user. In this paper, we argue that it is time for HCI to draw on a more complex notion of meaning-making, in which not only designers but users and other actors are understood - and encouraged - to play an active role in shaping the meanings of system. Taking user interpretation seriously raises substantial questions about our design and evaluation practices— What does it mean to design a system that can be appropriated? Is a system a success when it is used in ways that contrast markedly with the designer's intentions? Drawing on the traditions of philosophical hermeneutics, we lay out a framework for thinking through the complex ways in which designers, users, and other factors shape the meaning of systems that can begin to guide answers to these questions.

1 Introduction

In this paper, we discuss interpretation as a central concept in newly emerging forms of Human Computer Interaction. We suggest that, rather than conceiving of HCI as a matter of crafting interfaces to convey predetermined possibilities and functions as clearly as possible, it is more realistic to think of technologies creating situations and providing resources that people appropriate into their own systems of meaning. In other words, users can be seen as making the meaning of technological systems as much as designers do. As a consequence, the role of designers must shift from determining the meaning of a system to encouraging, shaping and occasionally disrupting users' processes of meaning-making. Concomitantly, evaluators must move from establishing whether a designer's meaning is, in fact, conveyed, to a more challenging and nuanced evaluation of the ways in which a variety of meanings, not all of which are intended by the designer, are established around systems.

There are three main reasons - historical, cultural and sociological - to consider interpretation central to current HCI. The historical reason is that HCI, and digital artifacts in general, have increasingly decentralized authority over both the narratives and the realities of computation. HCI's original role was as a translator, designing interfaces to help users understand the purposes, capabilities and functions of computer programs. Crucial to this role was the fact that these attributes of computer systems were not determined by HCI experts or users themselves, but instead by some mix of computer engineers and programmers, managers, marketing teams and businessmen. Thus HCI was in the business of conveying an established, authoritative meaning to users, helping them to use computer systems in the 'right' ways. Over time, however, the balance of existential control shifted, as HCI experts started to understand usability not only narrowly, in terms of an individual confronting a computer screen, but more broadly as depending on existing work practices and communities in which computer systems are used (c.f. Grudin 1990). With this shift in emphasis came a shift of inventiveness, as HCI practitioners were increasingly in the position to claim expertise about what sorts of systems people need and what functions those systems should provide. Moreover, it has become increasingly accepted that *users* are the experts about the situated domain knowledge that determines usability, and they have been given (or taken) a growing role as designers themselves,

either through processes such as participatory design (e.g. Bødker et al. 1985), techniques such as end-user programming (e.g. Maclean et al. 1990) or simply through appropriation and misuse (e.g. Dunne & Raby 2001). Overall, the story of human computer interaction is one of decreasing authority about the identity and meaning of computer systems, and a concomitant increase in the role of interpretation of what computer systems are and should be.

The second main reason to consider interpretation a key element in understanding current human computer interaction is cultural. While control over the design of computer systems has increasingly been democratized, the assumption that given systems embody a fixed, authoritative meaning has tended to persist within HCI, even as the culture of use has shifted from the workplace to our everyday lives. This move is a historical development, of course, but we believe it also represents a qualitative shift in the role and purpose of computers: as computers enter everyday life, everything changes. Most fundamentally, their purpose can no longer be expressed simply in terms of utility, efficiency and usability, and their success no longer simply evaluated along those dimensions. Values such as enjoyment, status, aesthetics and expressivity become increasingly important, and success must be assessed along such “subjective” dimensions as well. Moreover, traditional concepts of functionality founder as it becomes clear that successful systems may not be *for* anything in the traditional sense, but instead resources for playful exploration and curiosity-driven engagement (see Gaver et al. 2004). The result is that the experience and meaning of computational systems is increasingly complex, and may be perceived and enacted in very different ways by different people in different situations.

The third major reason for considering interpretation key to a correct understanding of human-computer interaction is sociological. The sociology of technology has generated a wealth of empirical data which underscores the essential role of users and other mediators in defining the meaning of technological products (e.g. Oudshoorn & Pinch 2003). While it may be natural for us, as the generators of technology, to presume that the meanings we attach to them are final, in fact substantial historical and sociological studies of technology have demonstrated that while designers may have the first word, they hardly have the last. Bijker, for example, has demonstrated that emerging technologies like the bicycle go through periods of interpretive flexibility, when their meanings are fluid and under debate, with social groups redefining those meanings based on their own interests until one meaning may come to be stabilized (Bijker 1995). Kline and Pinch have established that the meaning of automobiles shifted through a series of social battles among manufacturers, users, and distributors for decades before a car came to be what we know today: a black-boxed technology for transportation (Kline & Pinch 1996). Lindsay, in a study of the TRS-80, has demonstrated that such meanings continue to be developed by users and hackers long after designers, manufacturers, and distributors have abandoned a technology (Lindsay 2003). In our own time, we have seen cell phones shift from the designer’s intention of a device for business work to an essential part of teenage identity formation (e.g. Ito & Akabe).

Such cases make it increasingly difficult to maintain the position that a given system does or should embody a fixed, authoritative meaning to be communicated through its interface. The producers of computational systems are not necessarily authorities on their creations on the one hand, and users are not passive consumers on the other. In practice, of course, producers may argue persuasively for a preferred interpretation, and socially constructed interpretations tend to spread and reify over time. But appreciating the role and power of interpretation in determining the meaning of computational systems is important both in reconceptualizing the *process* of design, from initial user studies to final evaluations, and the *products* we create as well, allowing us to understand the value of systems that thwart easy interpretation or invite many competing ones.

In the rest of this paper, we build on these initial observations, seeking to establish new ways to interpret interpretation. We start by describing a number of systems that emphasize the role of interpretation in order to contrast the view of meaning as conveyed by a system with one of meaning as created in interaction with a system. We then review a number of theoretical standpoints from non-scientific fields that are usually overlooked by HCI, showing how they may allow us to better understand design for interpretation, and suggest that new accounts of interaction might best be based on these sorts of accounts rather than more traditional ones.

2 Design Studies

In this section, we ground our interest in and approach to designing for interpretation by demonstrating several cases from our own practice in which the model of systems as conveying the designer's meaning broke down in interesting ways. In each case, we expected the systems to be interpreted in a particular, albeit open-ended way, then found that users' actual interpretation of these systems was not only substantially different but arguably more interesting than what we had expected to find. These cases raise for us serious questions about how design practice can take account of such appropriation.

2.1 Double-Deck Desk



Figure 1: The Double-Deck Desk

The Double Deck Desk (Dddesk) was a two-storied work station with accompanying software that its designers installed in the foyer of a large office building over a two-week period (Gaver, Boucher, & Martin 2003). It was intended to address the need for contemplation within the workplace by literally lifting people from everyday commotion and by providing software that encouraged people to reflect on their activities and aspirations. The Dddesk was a standard-sized office desk at ground level with a raised platform and integrated table about 3.5 meters above the ground (see Figure 1). Prototype 'Reflection Engine' software projected onto the lower and upper desks allowed people to create 'mindmaps' – words arranged to show patterns and relations among concepts – across the two desks using keywords automatically culled from their existing files. The completed Dddesk was installed in the atrium of a large industrial research facility, and interested passersby were given a tour of the Dddesk so the designers could promote the work and assess reactions to it.

Dddesk was intended to raise issues about the nature of modern knowledge work, in particular to suggest that achieving contemplation is a value to be supported. Once the Dddesk had been explained to people, they engaged with this story by interacting with the software and the Dddesk itself. Indeed, sitting on the upper desk did seem to evoke feelings of peace and contemplation, removed from activities below. But it was not spontaneously perceived this way by uninformed observers. For while the Dddesk was successful in prompting viewers' narrative interpretations, the stories they told were not what had been expected.

Divergent interpretations of the Dddesk emerged from the very beginning of its development. As it was constructed within the designers' studio space, for instance, some of their colleagues complained that it suggested surveillance over the closely packed working environment. Within the office building, the Dddesk's appearance was initially greeted with bemused speculations about vertical office-sharing, prompted by its coincidence with a programme to reduce office space. Finally, requests for stories about the desk prompted accounts of its being used in a child's room by parent below and child above, or as a setting in which defence lawyers could practice dominating the opposition. As far as the designers could tell, nobody spontaneously interpreted the Dddesk according to the original intentions of providing an opportunity for contemplation and reflection.

Despite the Dddesk's failure to communicate unequivocally, it remained a powerfully suggestive object, provoking people to tell stories about it that revealed their desires and fears about the workplace. While the Dddesk did not tell the story the designers had envisioned for it, it did evoke powerful stories from the workers who puzzled over it – stories that were more meaningful for them than the one it had been designed to tell.

2.2 Miro

Miro was a graphical, interactive system installed by Boehner, Chen, and Liu in the Cornell Information Science building to provide building occupants with a sense of the overall emotional climate in the office. The designers surveyed the office for a week prior to installation to get a sense of the overall emotional rhythms during the day. They installed emotion entry stations in several locations that allowed users to input their emotions during the day. The emotional data collected through these two techniques was aggregated and displayed through the movements and colors of the objects in the display. The goal was for users in the office space to be able to develop a sense of the lab's emotional climate by interpreting the display, learning the language of the display over time.



Figure 2: Miro (left) and users puzzling over its meaning (right)

In practice, users did develop a sense of the lab's emotional climate by interpreting the display. This interpretation, however, did not consist in developing an understanding of what the display was intended by its designers to communicate. As one of its users said, "Uh, I have no idea what it means." Users would stand in front of the display, developing a sense of its meaning; "it's clearly displaying the stress levels related to that NSF deadline next week," users would say - even when the display was actually displaying happiness. As an object to be decoded, Miro was a clear failure. Nevertheless, users *did* develop a sense of the office's emotional climate from the discussions that Miro's presence incited. Users created interpretations of the system that were often more correct than the system itself, based on their background knowledge of what was happening in the office. Miro acted as a trigger for interpretation but did not actually directly transmit information. Oddly, Miro did fulfill its designers' intentions of encouraging reflection on emotional climate, but not in the way the designers intended.

2.3 Key Table

For about a month in 2003, Gaver, Boucher, Pennington, and Walker gave the Key Table (Figure 1) to a London family to live with in their own home. The Key Table was designed to infer people's emotions from the way they set down their accessories (e.g. keys, change, mobile phone) when they come home. Load sensors supporting the tabletop measured the force with which things were placed on it, and, by analogy to the heuristic that slamming doors spell trouble, guessed that forceful placements equaled emotional upsets. A wirelessly linked picture frame signaled the Table's mood estimations, swinging out of kilter to warn other household inhabitants to tread carefully.



Figure 3: The Key Table as envisioned (left) and evaluated (right)

The designers hoped the Key Table would emphasise the importance of emotional communication in the home and encourage people to think about how they use objects and spaces to express their moods (perhaps to the extent of *purposefully* using the Key Table to express their moods). The family who used the table drastically reinterpreted the designers' original suggestion that the table would reveal their moods, however. Instead, their imaginations were captured by the portrait of a dog that had been placed, with little reflection, in the picture frame. Based on this picture, the users anthropomorphised (or perhaps 'caninomorphised') the table as an animated presence in their home. Instead of seeing the table as reflecting their own moods, they saw it as having moods of its own. This ended with them rechristening the table as 'Terrence the Table', playing games with it "just as we do our cats," and dressing the table in unusual materials.

So was the Key Table a success? Again, it depends on the perspective from which success is to be judged. Seen as an object embodying a narrative about emotional interfaces, the user trial was a failure. Seen as an object lending itself to interpretation, appropriation, and attachment, on the other hand, it appears to have been very successful.

3 Rethinking Interpretation in HCI

By traditional HCI standards, each of the preceding cases documents a design failure, since in each case we failed to anticipate the ways in which users would make sense of the systems we had created. But in another sense, each of the systems presents a very real design success, since the systems engaged users, remained compelling over time, and the meanings and practices that developed around them were at least as and often more interesting than our original intentions. Still, as designers we are left in a difficult position. If these systems worked precisely because they did not work as intended, how could we make them better? If we focus on user appropriation of systems as an evaluation metric, is every system a success, since every system, perhaps especially unusable ones, can trigger new interpretations and can be used in ways we do not intend?

Here, we propose to develop a new model for meaning-making around technological systems that will address this complexity of interpretation and allow us to develop answers to these questions. The goal for this model is to support and motivate a new range of design and evaluation in HCI. Some of these practices already exist in HCI, but through our model their justification and coherency becomes clearer. In developing this model, we are drawing inspiration from the humanities, in which interpretation is a central trope. In particular, we here argue that the rich traditions of hermeneutics, or the philosophy of interpretation, can provide a new framework for understanding how designers, users, and other mediators create meaning and interpretations around technological systems. We propose this as the first step towards a comprehensive theory of human meaning-making around machines that can guide design and evaluation strategies.

3.1 Hermeneutics as a Frame for Co-Interpretation

In moving from a simple model in which designers are responsible for the meaning of systems to a more complex understanding of how meaning is created by many participants around machines, we are guided by a few central questions. How can designers shape the meanings of the systems they build, whether to favor particular ones or encourage a broad range of interpretations? How do users create meanings from technology, in ways that may receive, exceed or circumvent the meanings intended by the designers? What role is played in this process by designers, users, the technology itself, and other factors which are normally seen as external to the system design and evaluation process?

In order to answer these questions, we draw from hermeneutics, or the philosophy of interpretation. Classically, hermeneutics faces as its central conundrum the question of how a document or text (or, in our case, a technology) which is created in one context can become meaningful in a very different context from that in which it is created. We are by no means the first to draw on hermeneutics as inspiration for HCI specifically or computer science more generally; some of the more influential examples include Winograd & Flores's critique of AI systems as being unable to manage the infinite horizon of background assumptions that inform human interpretation (Winograd & Flores 1987) and Dourish's development of design principles based on an understanding of how people use social and tangible resources to make the world meaningful to them (Dourish 2001). Many approaches that draw on hermeneutics contrast formal models which aim to completely represent all aspects of context which are pertinent to the system's functioning with a (generally preferred) hermeneutic or interpretational approach which sees meaning as fundamentally undecidable in a formal, a priori manner and therefore requiring new, nonformalized approaches to systems design, e.g. (West 1997), (Chalmers & Galani 2004).

While we are sympathetic to this project, our aim here is different. Rather than drawing on a specific hermeneutic theory in order to provide one alternative to the formal project, we are stepping back to look at the possibilities for hermeneutic theories in general to frame the projects of HCI, including both formal and more situated or interpretational approaches. In doing so, we are not privileging one approach to hermeneutics, whether Heideggerian, Wittgensteinian, or Shannon-Weaverian, but rather suggest the potential for a wide range of interpretational strategies to rethink HCI, some of which, such as ethnography, already play a central role in HCI, and some of which, such as psychoanalysis, are as yet relatively untapped. The framework we are using to analyze these hermeneutic theories is drawn from the work of philosopher Paul Ricoeur (Ricoeur 1970), who identified two major alternative classes of hermeneutics: the *hermeneutics of faith*, in which the object is to uncover the meaning intended by the author in the text's/technology's new context, and the *hermeneutics of suspicion*, which strips away the author's intended meaning in search of *unintended* meanings which may be more true than those intended by the author. Here, we will unpack in turn the implications of each of these styles of hermeneutics for HCI. Again, we are not privileging one particular approach to

hermeneutics as preferable for HCI, but trying to understand how HCI can address the space of possible hermeneutic approaches.

3.2 The Hermeneutics of Faith

The hermeneutics of faith, as defined by Ricoeur, is a class of interpretational strategies whose goal is to reproduce as accurately as possible the original meaning of a text. The classic example of a hermeneutics of faith is biblical exegesis, in which interpreters are taxed to understand the meaning of a text written thousands of years ago for the modern context. Because the context in which the text is created is very different from the context in which the text is read, readers must do substantial work to unpack the meaning of the text for them. “Faith” in this case refers to the faith that there is an original meaning intended by the author, which it is the responsibility and duty of the reader to discover.

This model of interpretation, although derived from the humanities, is in some ways very similar to the model of interpretation used in HCI. In the hermeneutics of faith, a text is seen as springing from the author’s context and conveying meaning to the reader’s context. Similarly, in HCI we think of a technology as springing from the context in which it is designed, with the hope that it will convey a similar meaning in its context of use. So in evaluating a system, for example, we judge a system according to the designer’s original intention, defining a system as a success if the meanings intended by the designer in one context are successfully conveyed to the user in another.

A key difference, however, between the hermeneutics of faith approach and that of HCI is that the former highlights the labour required of the reader (user) in identifying the author’s (designer’s) original intention. In the hermeneutics of faith, this takes on a sense of moral duty: the reader is duty-bound to labour to understand the author’s intention. In HCI, the correct communication of the designer’s meaning is seen instead as the moral duty of the designer, with the goal of making correct interpretation of the system as easy as possible on the user. Nevertheless, HCI implicitly recognizes this labour in designing artefacts such as user manuals which can aid users in their quest to make sense of the system. Ethnographic approaches to HCI explicitly recognize this labour; Suchman’s classic *Plans and Situated Actions*, for example, can be read as an extended unpacking of all the labour that users must do to make sense of a technology in their own concrete situation. In studies, users frequently seem to operate with a moral economy derived from the hermeneutics of faith, working hard to try to understand what the system is *really* supposed to mean and feeling bad if they don’t succeed – even when designers intend many readings to be possible (e.g. Höök et al. 2003).

A second key difference between the hermeneutics of faith and HCI models of meaning-transmission is in the notion of what is communicated. HCI models generally see meaning as fairly clear-cut, informational units that are transferred from the designer to the user. This notion of information transmission is coupled with a model of reading as decoding of the information corresponding to each sign or symbol, as in, for example, the Informative Art project, which systematically manipulates attributes of apparent artworks to allow informed users to read them as bus schedules or weather maps (Redström et al. 2000). In the hermeneutics of faith, in contrast, symbols communicate more than they hold – they have a revelatory power that allows the reader to access a greater meaning than is contained in them as mere chunks of information. As Ricoeur describes, “[t]he revealing power of symbols opposes symbols to technical signs, which merely signify what is posited in them and which, therefore, can be emptied, formalized, and reduced to mere objects of a calculus” (Ricoeur 1970, p. 31). This hermeneutic conception of meaning as more than mere informational units is what informs the critiques of formal approaches discussed in the previous section, since if meaning is more than information and interpretation more than decoding, then purely formal models will never be able to fully address the meanings technologies in use have for people.

3.3 The Hermeneutics of Suspicion

Despite these differences, the hermeneutics of faith largely follows the implicit hermeneutics operative in HCI. While it, like HCI, aims to develop a maximal fidelity of meaning transmission, another alternative is provided by the *hermeneutics of suspicion*, which Ricoeur draws out of the works of Freud, Nietzsche, and Marx. Each of these thinkers seeks to interpret texts, not by searching for the author's original meaning, but by seeking meanings and subtexts that are *not* intended by the author to be the point or goal of the text. Psychoanalysis, for example, seeks to understand phenomena such as dreams, free association, and slips of the tongue as indicative of cloaked, subconscious conflicts, while Marxism sees all cultural phenomena as indicative of a masked struggle for power over material resources. Symbols, which carry revelatory power in the hermeneutics of faith, instead distort and hide meaning in the hermeneutics of suspicion. Under the hermeneutics of suspicion, therefore, “[t]o interpret is to understand a double meaning” (Ricoeur 1970, p. 8): both the superficial, intended meaning, and the hidden, unintentional, but more true meaning that lurks underneath.

If we apply the hermeneutics of suspicion to HCI, a substantially new picture emerges, in which authorial intent is no longer the gold standard for design. For example, if we apply a psychoanalytic model to systems evaluation, we may seek to identify how systems reflect their designers’ subconscious concerns and how these unintentionally shape user experience; or, alternatively, we may look at how users’ interpretations of systems is grounded in their own subconscious conflicts and what factors in the system design lead users to project these meanings onto the system. Similarly, if we were to design from a psychoanalytic perspective, the goal of our design may be to create a kind of “Rorschach” system that maximally supports users in projecting their own personal meanings onto it (cf. Gaver et al. 2003). Evaluation of such a system may examine how wide the range of possible meanings of the system is, rather than whether the author’s one intended meaning won out.

3.4 Implications for HCI

When applied to HCI, the distinctions raised by the hermeneutics of faith vs. the hermeneutics of suspicion suggests that we need to address both types of meaning in our practices: the *manifest* meanings – those intended by the designer – and the *latent*, or found meanings – those uncovered by the users or the analysts of the system. Identifying latent meanings, practitioners of the hermeneutics of suspicion argue, requires special apparatuses of interpretation that can dig beneath the obvious surface meanings to reveal other meanings lurking below. Here, we would like to suggest candidate theories HCI may draw on for developing a better understanding of how users and designers co-interpret systems. We use them to unpack the design case studies with which we began.

One example of such a theory is *ethnomethodology*, a set of methods for understanding how people make everyday sense of the world. Like other methods based in the hermeneutics of suspicion, it begins from an understanding that the apparent orderliness and meaning of the everyday world is illusory. It seeks ways to reveal how that orderliness and meaning is actually constructed by people while trying to make sense of the world. In many ways, Miro was an object study of ethnomethods, as users strove to make what was fundamentally incomprehensible meaningful for them.

Ethnomethodology has already been applied with substantial success in HCI (e.g. Button & Dourish 1996). There is room, we believe, for many more theories to help us rethink how users and designers co-interpret systems. As discussed in the prior section, one of these untapped theories is *psychoanalysis*. The Ddesk can be seen as a ‘projective object,’ similar to the projective tests used by psychoanalysts. Just as projective tests employ ambiguous imagery to elicit people’s interpretations

and thus reveal their preoccupations, so the Dddesk impelled but did not constrain interpretation. It served to raise issues regarding work without resolving them. With people reading it in terms of surveillance, space efficiency, play, or power, as well as our suggested interpretation of contemplation, the Dddesk demanded explanation but did not provide it. It was the Dddesk's very unclarity as a narrative object that made it effective in eliciting a wide range of interpretations. In this it resembles art pieces that invite explanation only to thwart it. Its power to impel numerous narratives seems to come from the juxtaposition of a normal desk with the use of height, a psychologically powerful dimension with many implications.

The *sociology of technology*, whose recent work on users was highlighted in the introduction, provides us with an apparatus for uncovering how social interaction between and among users and designers establishes particular meanings for systems. Sociology of technology questions the common assumption that technology's meaning is set by its designer, and instead interprets technology's meaning as the consequence of a process of negotiation between users, designers, and other institutions. The Key Table illustrates some of the major mechanisms by which this renegotiation works; the system provided unanticipated affordances for interpretation, which the household members used over time to negotiate a complex, unexpected, and somewhat uncomfortable interpretation. This new, socially negotiated meaning came, through the medium of the documentary filmmaker, to clash with the interpretation intended by the designers. Yet even as the users revelled in their freedom to reinterpret the technology, they were aware that their interpretation was idiosyncratic and probably, from the point of view of the designer, wrong. More work is required to develop design strategies that allow users to appropriate systems without this sense of transgression.

4 Conclusion

The hermeneutics of suspicion suggests a set of new tactics we may use to try to unpack the meaning of systems. Designers may have hidden assumptions we need to unpack. Designers may hold subconscious assumptions that are built into the system without his or her awareness. The medium chosen may shape or alter the meaning in ways the designer did not predict; i.e., users may interpret the system based on factors in the design that were thought by the designer to be irrelevant. The context in which the user interprets the system may substantially influence the meaning s/he creates. The message embodied in the system may not be limited to the system itself; it may also contain messages about cultural values we need to unpack.

As Ricoeur describes it, the hermeneutics of suspicion is an approach to interpretation in which one meaning - that intended by the author - is rejected in favor of a hidden, latent meaning that the reader may uncover. An alternative to this extreme position is offered by more recent developments in media theory in which the goal is not to uncover one true meaning but to identify multiple possible meanings, none of which is considered a priori to be "better." Derrida's (1978) notion of *deconstruction*, for example, sees texts as necessarily embodying contradictions and containing many conflicting viewpoints. Deconstructive readings reveal some of the viewpoints that are normally suppressed, but these conflicting views are no more 'true' than those intended by the authors. By grounding in these more relativist interpretational strategies, we may be able to better identify, design for, and evaluate the multitude of possible meanings that authors, designers, and other actors construct around systems.

5 Acknowledgements

This work is sponsored by NSF Award # IIS-0238132, an Intel grant for Supporting Well-Being and Reflection in the Digital Home, and the Equator IRC (EPSRC GR/N15986/01); the Double Deck Desk was sponsored by Hewlett Packard. Our thoughts on this work have been influenced by Kirsten

Boehner, Tom Jenkins, Michael Golembewski, and our other Affective Presence partners: Geri Gay, Kristina Höök, Joseph "Jofish" Kaye, Michael Mateas, and the Intel People & Practices Group. Special thanks to Brooke Foucault for her support of our project.

6 References

1. Bijker, W. E. (1995). *Of bicycles, bakelites and bulbs. Toward a theory of sociotechnical change*. Cambridge, MA: MIT Press.
2. Bødker, S., Ehn, P., Romberger, S., & Sjögren, D. (Eds.). (1985). *The UTOPIA project: An alternative in text and images (Graffiti 7)*. Swedish Center for Working Life, the Royal Institute of Technology, Stockholm, Sweden and the University of Aarhus, Denmark.
3. Button, G. & Dourish, P. (1996) *Technomethodology: Paradoxes and possibilities*. Proceedings of CHI 1996, pp. 19-26.
4. Chalmers, M. & Galani, A. (2004). *Seamful interweaving: Heterogeneity in the theory and design of interactive systems*. In Proceedings of DIS 04 (pp 243-252).
5. Derrida, J. (1978). A. Bass, Trans., *Writing and difference*. Chicago: U. Chicago Press.
6. Dourish, P. (2001). *Where the action is: The foundations of embodied interaction*. Cambridge, MA: MIT.
7. Dunne, A. & Raby, F. (2001). *Design noir: The secret life of electronic objects*. Basel, Switzerland: Birkhäuser.
8. Gaver, W., Beaver, J., & Benford, S. (2003). *Ambiguity as a resource for design*. Proceedings of CHI'03, Ft. Lauderdale, Florida. April 2003, 233-240.
9. Gaver, W., Boucher, A., & Martin, H. (2003). *Evaluating the Double-Deck Desk*. Proceedings of HCII '03, Crete, Greece. June 2003.
10. Gaver, W., Bowers, J., Boucher, A., Gellerson, H., Pennington, S., Schmidt, A., Steed, A., Villars, N., and Walker, B. (2004). *The Drift Table: Designing for ludic engagement*. Proc. CHI'04 Design Expo. New York: ACM Press.
11. Grudin, J. (1990). *The computer reaches out: The historical continuity of interface design*. Proc. CHI '90. New York: ACM Press, pp 261-268.
12. Holmquist, L.E. & Skog, T. (2003). *Informative art: Information visualization in everyday environments*. Proceedings of 1st international conference on computer graphics and interactive techniques in Australasia and South East Asia, pp. 229-235.
13. Höök, K., Sengers, P. & Andersson, G. (2003). *Sense and sensibility: Evaluation and interactive art*. In 2003 Conference on computer-human interaction (CHI).
14. Ito, M. & Okabe, D. (Forthcoming). *Intimate connections: Contextualizing Japanese youth and mobile messaging*. In R. Harper, L. Palen, & A. Taylor, eds., *Inside the text: Social perspectives on SMS in the mobile age*.
15. Kline, R. & Pinch, T. (1996). *Users as agents of technological change: The social construction of the automobile in the rural United States*. *Technology and culture*, Vol. 37, No. 4., 763-795.
16. Lindsay, C. (2003). *From the shadows: Users as designers, producers, marketers, distributors, and technical support*. In N. Oudshoorn & T. Pinch, eds., *How users matter*. Cambridge, MA: MIT Press, 29-50.
17. Maclean, A., Carter, K., Lövstrand, L., & Moran, T. *User-tailorable systems: Pressing the issues with buttons*. Proc. CHI '90, pp 175-182.
18. Oudshoorn, N. & Pinch, T., eds. (2003). *How users matter*. Cambridge, MA: MIT Press.
19. Redström, J, Skog, T., & Hallnäs, L. (2000). *Informative art: Using amplified artworks as information displays*. Proc. Designing augmented reality environments (DARE) 2000, Elsinore, Denmark, pp. 103 – 114.
20. Ricoeur, P. (1970). D. Savage, Trans., *Freud and philosophy: An essay on interpretation*. New Haven: Yale.
21. Suchman, L. (1987). *Plans and situated actions*. Cambridge: Cambridge.
22. West, D. (1997). *Hermeneutic computer science*. CACM (Vol 40, Issue 4), 115-116.
23. Winograd, T. & Flores, F. (1987). *Understanding computers and cognition: A new foundation for design*. Addison-Wesley.