

# Context as a Dynamic Construct

Saul Greenberg

*University of Calgary*

## **RUNNING HEAD: CONTEXT AS A DYNAMIC CONSTRUCT**

***Corresponding Author's Contact Information:*** Saul Greenberg,  
Department of Computer Science, University of Calgary, 2500 University  
Drive N.W., Calgary, Alberta, CANADA T2N 1N4. +1 403 220 6087

***Brief Authors' Biographies:*** Saul Greenberg is a computer scientist with an interest in studying individual and group behavior, articulating interface design principles, prototyping novel systems, developing software infrastructures and physical devices for rapid prototyping, and evaluating system effectiveness through user testing; he is a Professor in the Department of Computer Science at the University of Calgary. See [www.cpsc.ucalgary.ca/grouplab/](http://www.cpsc.ucalgary.ca/grouplab/) for details.

### **Abstract**

Context is a dynamic construct. While some contextual situations are fairly stable, discernable and predictable, there are many others that are not. Similar-looking contextual situations may actually differ dramatically, due perhaps to people's previous episodes of use, the state of their social interactions, their changing internal goals, and the nuances of local influences. The consequence is that, for all but simple cases, the designer of a context-aware application may find it difficult or even impossible to: enumerate the set of contextual states that may exist; to know what information could accurately determine a contextual state within that set; and to state what appropriate action should be taken from a particular state.

### **Contents**

1. Introduction: Revisiting The Definition of Context
2. Theories of Context
  - 2.1. Situated Action
  - 2.2. Activity Theory
  - 2.3. The Locales Framework
3. Implications to Practitioners

4. Conclusions.

5. Notes.

6. References.

# Context as a Dynamic Construct

Saul Greenberg

*University of Calgary*

## Abstract

Context is a dynamic construct. While some contextual situations are fairly stable, discernable and predictable, there are many others that are not. Similar-looking contextual situations may actually differ dramatically, due perhaps to people's previous episodes of use, the state of their social interactions, their changing internal goals, and the nuances of local influences. The consequence is that, for all but simple cases, the designer of a context-aware application may find it difficult or even impossible to: enumerate the set of contextual states that may exist; to know what information could accurately determine a contextual state within that set; and to state what appropriate action should be taken from a particular state.

## 1. INTRODUCTION: REVISITING THE DEFINITION OF CONTEXT

In the anchor article, Dey, Salber and Abowd (2001) ponder the definition of *context*. They first give the Webster's Dictionary definition: "...the whole situation, background or environment relevant to some happening or personality...", and then argue that this definition is too general to be useful in context-aware computing. After reviewing other definitions, they craft a new one that operationalizes the concept in terms of the actors and information sources involved in creating context:

"Any information that can be used to characterize the situation of an entity, where an entity is a person, place, or object that is considered relevant to the interaction between a user and its application, including the user and the application themselves. Context is typically the location, identity and state of people, groups and computational and physical objects."

From this definition, the authors present a context-aware computing framework. The premise behind the framework is that a set of toolkit components can be combined to determine a contextual state: the components capture, transform and aggregate the raw information. If the designer can articulate the set of interesting contextual states ahead of time, then programming for context is simply a matter of determining which of these contextual states we are in, and then having the system take an appropriate action.

Unfortunately, this definition glosses over one critical point. Context is a *dynamic construct* as viewed over a period of time, episodes of use, social interaction, internal goals, and local influences. While some contextual situations are fairly stable, discernable

and predictable, there are many others that are not. The result is that similar-looking contextual situations may actually differ dramatically. That is, in all but the simplest cases it may be difficult or impossible for a designer/programmer to list *a priori*:

- the set of contextual states that may exist.
- what information could accurately determine a contextual state within that set,
- what appropriate action should be taken from a particular state.

This has a severe consequence to context-aware computing, for it implies that the designer/programmer cannot always compose a correct ‘set of rules’ (or program or formal model) for determining context and appropriate actions.

The theme of this essay pursues this idea of context as a dynamic construct. To develop this theme, I will first review several theories that suggest the dynamic nature of context as a natural part of human-human-artifact interaction. I will then discuss the implications of this to Dey et. al., (2001) context framework, and to context-aware computing in general.

## **2. THEORIES OF CONTEXT**

Three reasonably well-known theories are highly relevant to context, for they describe what comprises context and how people work within these contexts. These are Situated Action (Suchman 1987), Activity Theory (Nardi 1997a) and the Locales Framework (Fitzpatrick 1998; Fitzpatrick, Mansfield, and Kaplan 1996; Fitzpatrick, Kaplan and Mansfield 1996). These theories help us understand the limits of context-aware computing, for all suggest that we must consider context as a continually evolving and highly situation-dependent construct.

### **2.1. Situated Action**

Suchman’s (1987) book *Plans and Situated Action* is a critique of plan-based approaches to cognition, which described activities as a set of well-defined goals and plans that people determine ahead of time. As an alternative, Suchman offers her *theory of situated action*, which takes the perspective that ‘every course of action depends in essential ways upon its material and social circumstances’ [p50, Suchman 1987], i.e., that human cognition and subsequent action is an emergent property of the moment-by-moment interaction of an individual with his or her physical and social environment. This theory does not imply that plans do not exist. Rather, and as nicely summarized by Dourish and Button (1998): “plans are one of a range of resources which guide the moment-by-moment organization of an activity, rather than laying out a sequence of work which is then blindly interpreted.”

Suchman’s situations are analogous to context; although unlike the definitions of context provided by Dey et. al. (2001), she is explicit about its fluid and ever-changing nature. Nardi (1997b) highlights this when she recapitulates that “the inquiry [of analysis of situated action] is meant to take place at a very fine-grained level of minutely observed activities, inextricably embedded in a particular situation”. Nardi continues “...it forces the analyst to pay attention to the flux of the ongoing activity, to focus on the unfolding

of real activity in a real setting”. Lave (1988) also points to the dynamic nature of context as defined by situated action: “[it] emphasizes responsiveness to the environment and the improvisatory nature of human activity.”

Situated action does not imply complete randomness and unpredictability. Suchman (p179, 1987) explains that the goal of studies of situated action is to explicate the relationship between structures of action and the resources and constraints afforded by physical and social circumstances. However, this is far from easy to do, and she cautions against the institutionalized (and often simplistic) social norms that are often used to predict and prescribe all actions (e.g., Suchman 1994). Nardi (p84 1997b) summarizes this as ‘a tension between an emphasis on that which is *emergent, contingent, improvisatory* and that which is *routine and predictable*’.

## 2.2. Activity Theory

*Activity theory* (Nardi 1997a) claims that activity defines context, where an activity comprises a *subject* (the person or group doing the activity), an *object* (the need or desire that motivates the activity), and *operations* (the way an activity is carried out)<sup>1</sup>. Artifacts and environment are seen as entities that mediate activity. Nardi (1997b) then argues that activity-as-context includes not only external resources (people, artifacts, settings) but internal processes (objects and goals) as well.

Activity theory also recognizes the dynamic nature of context. As Nardi emphasizes: ‘activity theory holds that the constituents of activity are not fixed but can dynamically change as conditions change’ (p75, 1997b). Contributing to this change is the transformative relations between people and artifacts: this transformation can change the internal objectives and thus the course of the activity. Thus context cannot be inferred simply by enumerating the external set of people and artifacts (Nardi 1997b), for it must also include people’s internal (and perhaps changing) states. This sentiment is also echoed in Bellotti and Edward’s essay (2001 [this special issue]) who state, “there are human aspects of context that cannot be sensed or even inferred, so context aware systems cannot be designed simply to act on our behalf”.

## 2.3. The Locales Framework

The *Locales Framework* was developed as a principled approach to help people understand the nature of social activity and work, and how a locale (or place) can support these activities (Fitzpatrick 1998; Fitzpatrick, Mansfield, and Kaplan 1996; Fitzpatrick, Kaplan and Mansfield 1996). It is a descriptive theory based on Strauss’ Social Worlds concept (1993), which in turn defines groups of people who share some commitment and collective action. Fitzpatrick observed real world situations: how people created social worlds and how they worked together within them. A main outcome is how she saw *locales* arise as a social world appropriates and uses particular sites and means for

---

<sup>1</sup> It is this inclusion of *object* or intention that separates this theory from situated action (p89 Nardi 1997b). Also, Activity theory’s use of the term *object* is quite different than how the same term is used in Dey et al’s definition presented in Section 1: the former uses it as a synonym for intention, the later as a synonym for artifact.

pursuing work. Locales are not necessarily fixed entities with fixed meanings. For example, two completely separate locales may be realized by a single physical meeting room space and how two different groups use it for the duration of their interaction. A locale does not even have to be associated with a physical space. Rather, it may arise in a virtual domain by the tools that are ready to hand for the group: email, MUDs, instant messaging services, and so on. The important point is that locales arise from social worlds, and that locales are associated with the site and the means for doing work.

Other relevant parts of the Locales Framework describe the subtleties of how locales are used. First are *individual views*. Because individuals can belong to many social worlds, a person may have one's own partial view into multiple locales at any time. As Fitzpatrick, Kaplan and Mansfield (1996) state "her involvement i.e., the current state of her interaction with and in the various worlds' locales, will determine which aspects of locales and indeed which locales are most relevant at any point in time." Second is *mutuality*, a term she uses to describe how people maintain a sense of shared place and that keeps them informed about shared activity. Mutuality includes one person's awareness of others, the artifacts comprising the locale, where things are located, and how things are changing within it. Third are *interaction trajectories*. This recognizes that interactions evolve over time, and that considerable effort is spent managing this evolution. Interaction trajectories include a group's control over past, present and future aspects of routine and non-routine work; how people coordinate and negotiate plans and activities over time; how people leverage past experiences; how breakdowns are noticed and repaired; and how processes are supported.

In many ways, the Locales Framework is about the social construction and use of context, and it too recognizes its dynamic properties. Locales (the site and means) are the external contributors to context; while locales can be fixed, most are fluid. The choice and makeup of these sites and their means change with the moment-by-moment needs of individuals and the social world as a whole. That is, the dynamics affect what facets of the setting might be relevant to context. What is relevant at one time may be irrelevant at other times. Mutuality is the mechanism whereby people recognize both context and subtle contextual changes. Individual views acknowledges that people can be involved in many contexts at once, and may selectively attend to them. Interaction trajectories highlight that context changes and evolves over time, and that part of people's ability is to recognize and manage that change.

### **3. IMPLICATIONS TO PRACTITIONERS**

The theories above describe how context is created from a variety of internal and external factors, how it changes moment by moment, and how people interpret it to perform actions. These theories make it quite clear that that designers must consider several non-trivial aspects when building systems that recognize and take advantage of context.

1. *Determining an appropriate set of canonical contextual states may be difficult or impossible.* It would be very convenient if a designer of a context-aware application could enumerate *a priori* a limited set of likely contexts and what comprises them. Programming the application then becomes a matter of determining which contextual

state best matches the current real context, and then taking appropriate action. While there may be a variety of settings where such canonical contextual states can be articulated and applied, there are also many settings where this will not be possible simply because no canonical set exists. There is also a temporal aspect: a canonical set that may seem an appropriate today may be completely inappropriate tomorrow because of internal and external changes in the social and physical circumstances.

2. *Determining what information is necessary to infer a contextual state may be difficult.* The above theories suggest that many things contribute to context, and that the relevance of any bit of information is highly dependant on the particular situation. External things—the artifacts, the physical environment, the people—are relatively simple to capture (although mapping that information onto a context may be hard). In contrast, internal things—an individual’s interest in that contextual setting, their history of interaction, their current objectives, and the state of the activity they are pursuing—are extremely difficult to capture. The problem is that there is likely no guaranteed way for a system to infer a correct contextual state by just collecting external information. At best, the system can only provide an approximation or educated guess of the real current context.
3. *Determining an appropriate action from a given context may be difficult.* The actions that people do, or the desired responses people expect from a context-aware application are also highly situation-dependent. The various theories suggest that even if two contextual situations appear almost identical, the desired action in one may differ substantially from the other simply because a different series of events may have led to those situations. People’s internal states—their objectives, their social constructs—may differ as well. Consequently, there is a chance that the action performed by the system may be the wrong one.

With these design caveats in mind, let us return to Dey, Salber and Abowd’s (2001) framework of generic toolkit components that developers can use to build context-aware applications. Their basic idea is that:

- *context widgets* capture the raw information used to characterize and environment,
- *interpreters* can transform the raw information captured by one or more context widgets into another more useful abstraction,
- *aggregators* collect related information into a single context and offer it to the application,
- *services* execute actions on behalf of the application, and
- *discoverers* maintain are registry of the above components

Their many examples then suggest how component architectures for context-aware applications are constructed: the context widgets and interpreters are used to collect information, the interpreters determine the contextual state, and the services take actions based on that state. The simplicity of this component architecture is very appealing to designers and programmers: they can use the conceptual framework to design the architecture, and the context-aware toolkit to implement the context-aware application.

While these ideas are good, the design trap should now be obvious. The framework and toolkit is an elegant way to design and implement context-aware applications for simple and highly routine contextual situations. However, it does not include anything to inform the designer about what contextual situations are appropriate to it, i.e., whether useful canonical contextual states exist, whether information can be captured to infer that state, and whether resulting actions are meaningful.

One can, of course, argue that the purpose of the framework and toolkit is not to inform designers about context, but to speed actual development. Yet the very fact that building context-aware applications becomes easy also means that we will see many uninformed designers build inappropriate ones. We have already seen this happened in other computer application domains. Workflow systems, for example, began with a fairly naïve view of how procedures in an office setting could be captured and automated. A variety of systems were built, where the job of the workflow analyst became a routine matter of recognizing (what appeared to be) canonical workflow states and the activities that should happen from those states. While this worked for certain types of settings and situations, workflow systems were applied to entirely inappropriate settings as well. People began to fight the system, for the system view of context (in this case the workflow context) did not fit with what was actually happening. It took some time for the community to recognize the problem, and even longer for commercial systems designers to accept the limitations of procedural workflow (Grinter 2000). Fortunately, newer workflow systems do recognize this problem, and are far more flexible in dealing with particular situations that do not fit the canonical workflow norm. Of course, workflow systems are not unique: we have seen similar problems with email-like systems based on speech-act theory (Suchman 1993).

The danger of building similar misguided context aware applications is not just an academic cautionary note, for we can already find many instances of inaccurate context-aware systems in even very simple settings. For example, one can now buy off-the-shelf lights that turn on when it detects that someone is within a room (or outside security lights that turn on when any motion is detected), or toilets that flush when someone stands up, or automatic sink taps that turn on when it detects something in the sink. Other examples are listed by Bellotti and Edwards' essay (2001). It would be humorous to recount the variety of times these systems get it wrong: but I am sure that the reader has experienced these and similar annoying situations themselves.

Our own experiences (by myself and my group) in building more sophisticated context-aware applications also echo how easy it is to get it wrong. For example, we built an always-on media space that tries to balance privacy and distraction concerns between distance-separated users of these spaces (Greenberg and Kuzuoka 2000). The idea is fairly simple: what people see through the video channel is a reciprocal function of how far away people are from the displays. If both are close together, they see and hear each other in full fidelity. If one moves away, sound is turned off and image fidelity is reduced through blurring (Boyle, Edwards and Greenberg 2000). If both move away, the image and frame-rate is further reduced. That is, we wrote a state table of possible contexts, instrumented the environment to determine which context the actors were in, and had a programmed action that the system would take on the actors' behalf. While the system



worked very well in certain settings (e.g., to connect two distant offices whose occupants wished to work together), we found it inappropriate and inaccurate for other settings. For example, when one camera was situated in a home office, it was possible for the system to inadvertently capture and present in full fidelity the image of another home occupant (perhaps in a state of undress) as they walked by the open door. Also, the simple rules did not work as well: there were times when the home worker did not want others to see them in full fidelity, perhaps because they had just woken up and were working behind the computer in their pajamas. People's internal states also affected this desire: sometimes a person just wanted some solitude.

Michael Boyle, one of our lab members, suggests that context-aware systems such as ours should be designed with the premise that there is a strong likelihood that the system will get things wrong. Consequently, context-aware systems: should be fairly conservative in the actions it takes; should make these actions highly visible; and should leave 'risky' actions to user control (Bellotti and Edwards (2001) take a similar point of view in their essay). For example, consider how we are redesigning our reactive media space (Boyle 2001). First, to minimize privacy risks, the system is now far more cautious about revealing video than masking it. When a full privacy condition is inferred, the system physical rotates a motorized camera away from the person to point to a blank wall. If this inference is incorrect, the parties are inconvenienced but not put at risk. Because the opposite action of inferring that a person wants to regain their full video connection is far riskier; we make this a manual action; the system does not rotate the camera back to face the person automatically. Second, system presents noticeable feedback to the actor. We use camera rotation instead of a lens shutter to make the camera position highly visible, and we attach a small display to the camera so one can always see exactly what is passed on to others (i.e., whether full or blurred video is being transmitted), Third, it presents simple mechanisms for the people to adjust and/or override its context-awareness properties i.e., a person can momentarily place their thumb over the camera lens can to block the video completely (Boyle, Edwards and Greenberg 2000).

The main point behind the above examples is that building inappropriate context-aware software and hardware is already happening. The challenge is to educate designers about the subtleties of context, and to build software that can cater to the changing nature of context.

#### **4. CONCLUSION**

Context-aware computing is an important evolutionary step in computer use. Understanding context is vital if we are to build effective ubiquitous computing systems, information appliances, reactive environments, and computer devices that fit within architectural settings. However, we as designers must realize that it is all too easy to trivialize context, and as a result we will end up building inappropriate applications.

We are still coming to grips with what we mean by context, and it will take some time before we have good design principles for context-aware computing (see Bellotti and

Edwards' essay (2001) for a starting point). For now, we present three high-level ideas about what is needed to get context right.

Getting context right means that we must study the expected context-of-use carefully. One approach is to use ethnographic methods to observe and analyze many contextual episodes, and to determine (using Nardi's words in 1997b) what contexts are mostly emergent, contingent, improvisatory and what are routine and predictable. Similarly, we must carefully understand the effects of getting context wrong: in some situations, a single inappropriate system action may be enough to preclude people from using it further. This implies that contextual systems should be fairly conservative in the acts it takes: 'risky' actions should be taken only if there is compelling evidence of correctness.

Getting context right also means that our toolkits must incorporate flexibility. Systems that infer context will hopefully get contextual guesses correct most of the time, but failure is inevitable. Because the information used to infer context might change, people should be able to adjust what information is collected and how it is interpreted so that context is inferred more accurately. Because expected responses to a context may change, people should be able to adjust the actions the system takes so that they are appropriate. Because some contexts cannot be inferred (perhaps because they appear only once), people should be able to over-ride the system altogether.

Getting context right is also about good interface design. People should be able to see what context the system thinks it has inferred. Through feedback, actions taken by the system should be clearly linked to that context. Through light-weight interface mechanisms, people should be willing to adjust contextual information, over-ride the system, or adjust the way actions are taken—if a person considers any of the above too hard to do and if consequences of wrong contextual guesses are annoying, they will likely just turn off the system altogether. Finally, observations of how the context-aware application performs in real world setting should inform its further design and redesign.

## 5. NOTES

*Acknowledgments.* Dey, Salber and Abowd wrote a fine anchor article that helped provoke thoughts about context in computing: this essay would not be possible without them. As well, the systems produced in my laboratory by Michael Boyle, Michael Rounding and Hideaki Kuzuoka, and our group's discussions about the successes, compromises and failures when using these systems contributed significantly to the thoughts expressed in this article. Revisions between the draft and the final version of this essay were influenced by discussions with Michael Boyle, as well as by comments made by Tom Moran and Paul Dourish.

*Support.* Partial funding of provided by the National Sciences and Engineering Research Council of Canada (NSERC), the Alberta Software Engineering Research Consortium (ASERC), and Microsoft Research.

*Authors' Present Addresses.* Saul Greenberg, Department of Computer Science, University of Calgary, 2500 University Drive N.W., Calgary, Alberta, CANADA T2N 1N4. +1 403 220 6087

*HCI Editorial Record.* (supplied by Editor)

## 6. References

- Bellotti, V. & Edwards, K. (2001). Intelligibility and Accountability: Human Considerations in Context Aware Systems. *Human-Computer Interaction, 16*, xxx-xxx. [this special issue]
- Boyle, M. (2001) *Ubiquitous Awareness Spaces*. (Yellow Series Report 2001-682-05). Department of Computer Science, University of Calgary, Alberta, Canada
- Boyle, M., Edwards, C. & Greenberg, S. (2000). The Effects of Filtered Video on Awareness and Privacy. *Proceedings of the CSCW'00 Conference on Computer Supported Cooperative Work*, ACM Press.
- Dey, A. K., Salber, D. & Abowd, G. D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction, 16*, xxx-xxx.
- Dourish, P. & Button, G. (1998). On "Technomethodology": Foundational Relationships between Ethnomethodology and Interactive System Design. *Human-Computer Interaction, 13*(4), 395-432.
- Fitzpatrick, G. (1998) *The Locales Framework: Understanding and Designing for Cooperative Work*. Unpublished doctoral dissertation. Department of Computer Science and Electrical Engineering, The University of Queensland, Brisbane, Australia.
- Fitzpatrick, G., Mansfield, T. & Kaplan, S. (1996) Locales framework: Exploring foundations for collaboration support. *Proceedings of the OzCHI '96 Sixth Australian Conference on Computer-Human Interaction*, 34-41, Hamilton, New Zealand.
- Fitzpatrick, G., Kaplan, S. & Mansfield, T. (1996) Physical Spaces, Virtual Places and Social Worlds: A Study of Work in the Virtual Places for Collaboration. *Proceedings of ACM CSCW'96 Conference on Computer-Supported Cooperative Work* p.334-343, ACM Press
- Greenberg, S. & Kuzuoka, H. (2000). Using Digital but Physical Surrogates to Mediate Awareness, Communication and Privacy in Media Spaces. *Personal Technologies, 4*(1), Elsevier.

- Grinter, R. (2000) Workflow Systems: Occasions for Success and Failure. *Computer Supported Cooperative Work* 9(2), 189-214, Kluwer Academic Press.
- Lave, J. (1988) *Cognition in practice*. Cambridge University Press.
- Nardi, B. (1997a) *Context and Consciousness: Activity Theory and Human Computer Interaction*, MIT Press.
- Nardi, B. (1997b) Studying context: A comparison of activity theory, situated action models, and distributed cognition. In Nardi (1997a).
- Strauss, A. (1993) *Continual Permutations of Action*. Aldine De Gruyter, New York.
- Suchman, L. (1987) *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.
- Suchman, L. (1994) Speech Acts and Voices: Response to Winograd et al. *Computer Supported Cooperative Work*, 3(1), 85-95, Kluwer.
- Suchman, L. (1993) Do Categories Have Politics? The Language/Action Perspective Reconsidered. *Proceedings of the Third European Conference on Computer-Supported Cooperative Work*, 1-14, Kluwer Academic Press.