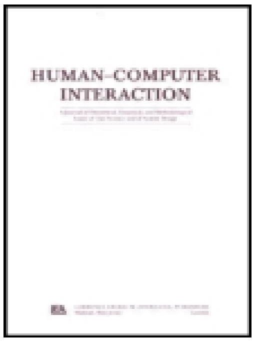

Human–Computer Interaction

ISSN: 0737-0024 (Print) 1532-7051 (Online) Journal homepage: https://www.tandfonline.com/loi/hhci20

Defining, Designing, and Evaluating Peripheral Displays: An Analysis Using Activity Theory

Tara Matthews PhD , Tye Rattenbury PhD & Scott Carter PhD

To cite this article: Tara Matthews PhD , Tye Rattenbury PhD & Scott Carter PhD (2007) Defining, Designing, and Evaluating Peripheral Displays: An Analysis Using Activity Theory, Human–Computer Interaction, 22:1-2, 221-261

To link to this article: https://doi.org/10.1080/07370020701307997

Published online: 05 Dec 2007.

Submit your article to this journal

Article views: 414

View related articles

Citing articles: 3 View citing articles

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=hhci20

HUMAN–COMPUTER INTERACTION, 2007, Volume 22, pp. 221–261
Copyright © 2007, Lawrence Erlbaum Associates, Inc.

Defining, Designing, and Evaluating Peripheral Displays: An Analysis Using Activity Theory

Tara Matthews, Tye Rattenbury, and **Scott Carter**
University of California Berkeley

ABSTRACT

Peripheral displays are an important class of applications that improve our ability to balance multiple activities. However, peripheral display innovation and development has suffered because much of the past work has been technology driven: There exists little theoretical understanding of how they operate in relation to people's everyday lives. In response to this, we present a framework for understanding, designing, and evaluating peripheral displays based on Activity Theory. We argue that peripheral displays are information displays that become unobtrusive to users. As this quality depends on the context of use, we present a framework for describing peripheral displays based on the number and types of activities they support. Furthermore, we argue that different types of displays require different approaches to evaluation. From our own work and a review of related literature we derive a set of general evaluation criteria for peripheral displays (appeal, learnability, awareness, effects of breakdowns, and distraction). We

---

Tara Matthews is a computer scientist with interests in peripheral displays, glanceability, evaluation, multitasking, accessibility, and ubiquitous computing; she is a PhD candidate at the University of California Berkeley. **Tye Rattenbury** is a computer scientist with an interest in models of human behavior, particularly in applications of models in work support systems; he is a PhD candidate at the University of California Berkeley. **Scott Carter** is a computer scientist with interests in ubiquitous computing, peripheral displays, accessibility, and social media; he is a PhD candidate at the University of California Berkeley.

CONTENTS

then describe approaches for evaluating these criteria for different types of peripheral displays and present a case study to illustrate the value of our Activity Theory evaluation framework in practice.

1. INTRODUCTION

Everyday environments include many sources of information that we monitor with little to no conscious effort. Windows reveal the weather, the approximate time of day, and the level of business nearby. A first impression of a restaurant can be gleaned from aromas. Time is passively monitored using clocks, which help us manage when and how to switch between the various activities in our lives. At any given moment, most of the information in our environment is *peripheral* to our main focus of attention. It follows that in a ubiquitous computing environment involving more than one device, most devices and displays will be peripheral. Similarly, when balancing multiple activities on or off the desktop, only one action (or task) will be the focus of a person's attention. *Peripheral displays* are an important class of applications that allow a person to be aware of information from multiple aspects of one or more activities without being overburdened (Weiser & Brown, 1996).

Although the requisite of demanding little to no conscious attention provides a general description of peripheral displays, researchers have nonetheless employed a variety of refinements and generalization in their own definitions. The lacking consensus on a single definition contributes to the difficulty of evaluating peripheral displays, an acknowledged problem in the field (Mankoff et al., 2003). Common terminology and meaning when discussing peripheral displays would enable better design and evaluation. One contribution of this article is to present a way for researchers to understand and describe how peripheral displays are embedded in everyday contexts.

The descriptive theory of peripheral displays we propose and discuss in this article is based on Activity Theory (Leontiev, 1978). Our decision to use Activity Theory is rooted in the observation that any peripheral display may operate in different contexts (both socially and physically defined). Activity Theory provides a framework for describing user context (Nardi, 1996) and consequently provides a framework for describing how people and peripheral displays interact in various situations. Specifically, a peripheral display is any information display that is (1) a *tool* in at least one activity of its user and (2) is used primarily at the *operation*[1] level rather than the action level (i.e., usage requires relatively low cognitive cost[2] due to practiced, low-effort use, a process that can be helped by specific design toward easy interpretation). This definition is user-centered: A display is peripheral when a particular person uses it in an operational or automatic way.

---

1. *Tool* and *operation* are terms used in Activity Theory that are defined in this article.
2. By "low cognitive cost" we mean that monitoring the peripheral display causes minimal distraction from the user's focus of attention, which, by definition, is not the peripheral display.

The contribution of our Activity Theory approach to peripheral display design and evaluation is its explicit focus on operationalization and the set of activities that contextualize the usage of the display. This focus suggests important criteria to measure for evaluations and indicates how those criteria change in importance in different contexts. Some of these criteria–appeal, awareness, effects of breakdowns, and distraction–are present in prior work (Mankoff et al., 2003; McCrickard, Chewar, Somervell, & Ndiwalana, 2003; Shami, Leshed, & Klein, 2005). A standard approach to designing peripheral displays to meet these criteria involves (a) interviewing and/or surveying people to find the location of the display and type of information to display, (b) designing a display for that location and information, and (c) conducting a summative evaluation of the display (Mankoff & Dey, 2003). However, this approach does not recognize that a peripheral display is fundamentally operationalized and that its use depends on the type of activities it supports. Therefore, our work extends prior criteria to include *learnability*, or how readily the display is operationalized. We also highlight evaluation difficulties related to operationalization. For example, some evaluation methods, such as experience sampling, inherently call attention to the tool being evaluated, which is particularly problematic for peripheral display evaluation because it would violate the very nature of the tool. Furthermore, we present design dimensions that characterize how different contexts affect evaluation criteria. Our design dimensions are derived from Activity Theory and include *scope* (will the display support one or many of the target user's activities?), *class(es) of supported activities* (are the supported activities primary, secondary, or pending for the user?), and *criticality* (what is at stake if the user is not aware of the information in the display?). For example, our work clarifies the importance of demonstrating that a high-criticality information display fits into its environment, is empirically easy to learn, provides appropriate levels of notification, is explicit about breakdowns, and is only as distracting as the importance of the information being shown. On the other hand, a display of low-criticality information need not be easily learnable or display breakdowns, but it should be aesthetically pleasing, minimally distracting, and convey information in accordance with the user's interests.

In the next section, we briefly discuss how peripheral awareness relates to other forms of awareness that have been studied and deployed in prior work. Then we introduce Activity Theory and define peripheral displays based on their use. Then we present three design dimensions that describe the use of peripheral displays in context. Next, we discuss important evaluation criteria for peripheral displays based on interviews with peripheral display creators, our own design and evaluation experience, and existing literature. Then we use our design dimensions and criteria to discuss guidelines for focusing evaluations on issues important to peripheral displays. We conclude with a case study of the design and evaluation of a peripheral sound display for the deaf, which highlights the strengths and weaknesses of our Activity Theory framework.

2. RELATED WORK ON AWARENESS

There have been a number of research projects attempting to understand and improve awareness. These projects have utilized different, although related, definitions (Schmidt, 2002), resulting in a variety of approaches to improving awareness. We provide a detailed discussion of closely related work near the pertinent sections in the body of this article. In this section, we focus on more historical and loosely related work to help the reader situate peripheral displays in the wider area of awareness research.

At a high level, awareness support can be divided into two categories depending on the content of what is displayed. On one hand, there are systems and studies of group awareness where the content of information conveyed or analyzed relates to the question, What have my colleagues been working on lately? This question is generally answered either on short time scales like minutes, hours, or days, or at longer time scales like weeks, months, or possibly years. Most of the work on these types of systems has been in the Computer Supported Cooperative Work (CSCW) area of human–computer interaction (HCI), whereas many of the studies on group awareness have been performed outside of HCI in areas like organizational science. On the other hand, there are systems and studies that look at personal task management where the content of information displayed or analyzed relates to the question, What do I need to be aware of to get my work done? Most of the systems that support personal task management have been developed within HCI, whereas studies on personal task management have been a focal point in related areas like psychology and anthropology. We discuss each of these categories in turn.

Group, or more generally social, awareness work can be subdivided on the timescale of phenomena that are studied. Support for shorter timescale (minutes, hours, days) awareness focuses on the coordination of joint-work (Begole, Tang, Smith, & Yankelovich, 2002; Dourish & Bellotti, 1992; Erickson & Kellogg, 2000), for example, who is editing the paper now, when is a good time to ask Bob a question about our project, or who is planning on attending the symposium talk this afternoon? In some cases these systems are essential for collaborative, collocated activities (Heath & Luff, 1992). Other systems provide awareness of distributed activities by providing real-time views on coworkers (Dourish & Bly, 1992; Greenberg, 1996; Gross, Wirsam, & Graether, 2003; Pedersen & Sokoler, 1997) or views of shared resources (Gross et al., 2003). Support for awareness on longer timescales (weeks, months, years) is often concerned with building transactive memory (Mohammed & Dumville, 2001), for example, who knows what in an organization. Some systems do this by presenting work summaries (Huang, Tullio, Costa, & McCarthy, 2002; McCarthy, Costa, & Liongosari, 2001). Other systems focus less on awareness and more on archiving with the goal of knowledge management and transfer (Davenport & Prusak, 1998). Still other systems are concerned with increasing social ties and displaying common interests (Carter, Mankoff, & Goddi, 2004; McCarthy et al., 2001; Sawhney, Wheeler, & Schmandt, 2001). In most of these cases, the awareness supported by the systems is meant to be consciously and focally (in terms of attention) acquired. Although peripheral displays can clearly play a part, they have not been the focus of group awareness systems.

Personal awareness systems tend to focus on task management. They often maintain a representation of the many tasks people work on (e.g., MacIntyre et al., 2001) and support organizing notifications (e.g., van Dantzich, Robbins, Horvitz, & Czerwinski, 2002) and keeping track of potentially interesting events (e.g., Cadiz, Venolia, Jancke, & Gupta, 2002). Some of the systems designed for personal awareness have been peripheral (e.g., McCarthy et al.,, 2001). We discuss personal awareness systems in more detail later in this article (specifically in regards to their design and evaluation).

For this article we treat peripheral displays as output devices that provide *peripheral awareness*, which we define as the amount of information shown by the display that people are able to register and use without focal attention. Note that by this definition, peripheral awareness can be considered a subset of situational awareness, which involves all information perceived, comprehended, and acted on in a certain context (Endsley, 1995). Although all of the work mentioned in this section could support situational awareness to some degree, only those that are operationalizable, such as (MacIntyre et al., 2001), could provide peripheral awareness without requiring focal attention.

3. ACTIVITY THEORY AND PERIPHERAL DISPLAYS

Peripheral displays are often thought of in perceptual terms. In these definitions, peripheral displays are those that literally sit in the periphery of a person's field of vision or that issue subtle auditory cues. Although these perceptual interpretations of peripheral displays are initially useful, they quickly break down. What happens if a person changes the angle of their gaze or starts attending to the auditory cue? Does the device cease to be a peripheral display? To fully analyze peripheral displays, one should also consider the "messy" scenarios when a peripheral display becomes the user's focus of attention. The theory we chose for this purpose is Activity Theory. Activity Theory is an expansive framework for describing human behavior, but for our purposes there are four particularly important points:

1. Activities, objects, and motives: *Activities* correspond to long-term (weeks, months, years) projects of a person or group of people. These projects are directed toward some object in the world like a new product design or a social relationship. An object may be physical or conceptual, but is always coupled with a motive – a driving force that seeks to satisfy some need(s) of the people in the activity (Kuutti, 1995; Leontiev, 1978, 1981).
2. Actions and operations: For execution, activities are composed of goal-directed *actions* that are themselves composed of environmentally contingent *operations.* Actions are equivalent to the classic HCI notion of tasks, operations are equivalent to the operations from GOMS (John & Kieras, 1996).
3. Tool artifacts: In terms of elements in the world, activities are composed of an object, the people involved, and the *tools* that the people use to carry out and support the operations and actions in the activity. Tools are socially constructed artifacts that both encode the operations in which they are used and guide the user in formulating goals and actions for using them.
4. Multiple, ongoing activities: People have multiple, ongoing activities at any given time. However, they are generally working on only a subset of these activities through their current actions and operations.

These points are made by Leontiev (1978) in his initial formulation of Activity Theory. There have been a number of extensions to this original formulation, most notably by Yrjö Engeström and his colleagues (Engeström, Miettinen, & Punamäki, 1999). This extension has two primary additions to Leontiev's original formulation. First, the notions of community, rules, and division of labor are included with the notions of tools, objects, and motives. Although these additional notions clearly influence how an activity is carried out, they are most useful in analyzing how an activity changes and evolves to handle various breakdowns and contradictions that arise (Engeström, 1987). Second, groups of activities are often modeled as systems of codependent activities (whose objects and motives constrain one another) where the relationships between the activities are delineated and analyzed. Both of these extensions have provided insights into the nature of activities, particularly in how they change and evolve (e.g., Nardi & Redmiles, 2002). However, as we demonstrate in this article, the framework proposed by Leontiev, captured in the aforementioned four points, provides enough paradigmatic scaffolding for designers to implement a more context sensitive approach to the design and evaluation of peripheral displays. We discuss these four points in more detail in the following subsections.

3.1. Activities, Objects, and Motives

An *activity* is a long-term project of a person or group of people. People engage in activities to satisfy their needs, for example, the need to eat, to be social, or to accomplish (Leontiev, 1978). Although a person's needs have biological as well as social origins, their expression is largely determined by that person's sociocultural setting. Need expressions are captured in the object and motive of the activity. For example, in attempting to satisfy the needs previously listed, a person might orient their actions toward the yield of a harvest, a familial relationship, or a doctorate degree. Roughly, an object is the topic of an activity, the entity in the world that gives the activity temporal and spatial coherence as well as semantic grounding (Collins, Shukla, & Redmiles, 2002). An object can be material (e.g., a handcrafted chair), semitangible (e.g., a plan for building a chair), or abstract (e.g., a vague but developing design idea for a chair; Kuutti, 1995; Leontiev, 1978).

The object of an activity is always associated with a motive. The motive of an activity summarizes the key aspects or transformations of the object that will satisfy the person's underlying needs. Hence, the motives of activities are generally referred to as *life-forces.* Lasting multiple days, months, or years, these motives constitute a major part of every person's personality (Leontiev, 1978). Common motives include seeking social companionship by developing a relationship and achieving desired transformations of self by learning a new skill set or by gathering knowledge.

As we discuss in this article, peripheral displays are artifacts that augment and support one or more of a user's activities. Practically, this means that the peripheral display promotes and empowers the user in their efforts to satisfy the motives of their activities. Hence, peripheral display designers need to be aware of their target user's activities and how a peripheral display could appropriately influence them.

3.2. Actions and Operations

Activities are performed through multiple shorter term (seconds, minutes, hours) processes called *actions.* Actions are similar to what HCI literature calls goal-directed tasks. For example, programming a module might be an action within the activity of completing a software project. In HCI literature regarding peripheral display use, the actions and goals a user is currently focusing her attention on is often called the *primary task* of the user. We call this the *primary action.* Actions often service multiple activities. This statement has two distinct implications: (a) A generic action can service different activities at different times (e.g., writing a passage of text could be a useful action in many different activities), and (b) a specific action can simultaneously services multiple activities (e.g., writing a passage of a research paper might service a research project activity, a developing understanding of the language in which the paper is written, and providing a paycheck with which to support a family).

The activities that are serviced by the user's primary actions provide a context for understanding it (Nardi, 1996). For example, if a person is writing a passage of text, the context surrounding this action is dependent on whether she is writing an e-mail to a relative under the activity of maintaining a familial relationship, she is a researcher writing a section of a research publication, she is writing part of a fictional short story for fun, and so on.

Actions involve multiple *operations*, which are well-defined, habitual routines used while a person is performing an action. The concept of operations is similar to *automatized* behavior defined in cognitive science literature. *Automaticity* describes skilled behavior that can be performed easily with little effort and attention, generally as a result of practice or learning. Operations are directly influenced by the *conditions* of the environment in which they take place. An important distinction between the operation level and the action level for peripheral displays is that operations require low cognitive load to execute, whereas actions require high cognitive load. This distinction is a consequence of the creative nature of setting and accomplishing action goals versus executing a habitual routine at the operation level. This distinction practically means that people are generally working on at most one action at a time (their primary action) due to limitations of cognitive resources like working memory. Conversely, a person can carry out many operations at once in appropriate environments (*e.g.*, drinking and eating while having an intense discussion about a research project).

Thus, if a display is to be classified as peripheral it should be used primarily at the level of operations (and not at the action level). This does not mean, however, that a peripheral display will never reach the action level. Even well-designed peripheral displays will occasionally become the focus of their user's attention, but reaching operational use is a central objective in peripheral display design. Also, a new peripheral display must be learned and appropriated by its user, which is generally done at the action level or even as a separate activity (Leontiev, 1978). (Learning how to use a new peripheral display is one of the evaluation criteria that we discuss later in this article.) For example, many peripheral displays abstract the information they present in some way. The user generally has to learn this abstraction to be able to easily (i.e., at the operation level) interpret the information presented.

Operations are sequenced to complete an action, but this sequencing is not always optimally efficient or constructive (e.g., trial-and-error methods of problem solving). An important type of operation sequencing in peripheral displays design involves chaining together operations that do not build on one another. For example, most people will subconsciously glance at a clock or out a window while they are working on various actions. These glancing operations may not contribute to the completion of the user's primary action, but they occur nonetheless. Later in this article we define a class of peripheral displays that rely on operations that do not service the user's primary action.

3.3. Tool Artifacts

A fundamental pillar of Activity Theory is that people's interaction with the world is mediated by physical and psychological tools (Leontiev, 1978). Tools are artifacts that enable people to act on the objects of their activities. In other words, these tools allow users to accomplish, understand, motivate, or see the future transformations of their activities.

Tools are socially constructed. Hence they are subject to the trends and fashions of cultures, which are constantly evolving and transforming. Individually, tools are appropriated and adapted for various actions and operations. Activity Theory is vague when defining tools, because it tries not to bias or limit the full variety of forms and functions that mediating artifacts take.

Peripheral displays are tools. Moreover, the importance of a peripheral display is determined by its importance in the activity or activities that it supports. This could range from low (e.g., a single clock in a room with multiple clocks) to high (e.g., the altimeter in a cockpit).

3.4. Classification of a Person's Multiple Ongoing Activities

As previously noted, a person's actions can service multiple activities simultaneously. The sequence of operations performed by a person may also include some operations that do not service the goal of their primary action. Relying on these two points, we posit a set of four classes of activities. Note that this classification is not standard in Activity Theory: We derive them here to categorize the types of activities peripheral displays are likely to support. Also note that these classes of activities are not universal or static. They not only differ from person to person, they change as the person's primary action changes (e.g., a pending activity may become a primary activity).

1. *Dormant activities.* This class consists of activities that are not serviced by any operation performed in the user's current sequence of operations. Generally, there will be activities that a user is likely not to work on in a particular setting (e.g., developing a relationship with a distant relative while in their office setting). However, many of these restrictions are based on choice, and the user may choose to incorporate these activities in the future. In designing peripheral displays, we recommend that

designers only treat activities that cannot be worked on in their current setting as dormant. For example, the activity of learning to ski is hard to work on in an office cubicle and can reasonably be ignored in the design of an office peripheral display. We do not discuss dormant activities further because their impact on the design of displays is limited.
2. *Primary activities.* This class includes activities that are serviced by the user's primary action. In other words, it is the class of activities that are serviced by operations needed to complete the user's primary action. We refer to tools that support this class of activities as primary tools.
3. *Secondary activities.* This class includes the activities that are serviced by operations that are in the user's primary action but do not promote the attainment of the primary action's goal. In other words, these activities are not the focus of the user's current action. In contrast to the next class of activities, secondary activities are not likely to become primary in the near future.
4. *Pending activities.* This final class of activities is similar to secondary activities with one important distinction. Whereas both secondary and pending activities are monitored by the user, pending activities are monitored with the intent that they will become the primary activity in the near future. Generally, a pending activity was once a primary activity that has been temporarily "set aside." The user is monitoring some aspect of the activity to decide when to start working on it again.

An important caveat to the activity classification just described concerns the notion of serendipity. Many peripheral displays have been designed to support serendipitous information gathering (Carter et al., 2004; Kaye, 2001; Mynatt, Back, Want, Baer, & Ellis, 1998; Pedersen & Sokoler, 1997), often as their main purpose. As we argue later in this article, serendipity can be understood within our Activity Theory framework. Consequently, our primary focus is not on defining or understanding serendipity but on how to design and evaluate peripheral displays that support it. Accordingly, we discuss serendipity in the upcoming peripheral display design and evaluation sections.

3.5. An Activity Theory Definition of Peripheral Displays

Based on the four characteristics of Activity Theory that we just introduced, we arrive at the following definition:

> A peripheral display is any information display that (a) is a tool in at least one activity of its user and (b) is used primarily at the operation level rather than the action level.

This definition is framed by an understanding that the peripheral display user has multiple, ongoing activities. Depending on how the peripheral display is used, it may be a primary tool in the user's primary activities or a nonprimary tool supporting the user's secondary and/or pending activities.

Notice in our definition of peripheral displays that they are primarily used at the operation level. This means that usage requires relatively low cognitive cost due to operationalization (i.e., automaticity). The amount of learning required by a user to operationalize their use of a peripheral display will depend in part on how well the display adheres to ease of use design principles like glanceability (i.e., enabling quick and easy visual information uptake).

For the scope of this article, we use this definition to distinguish peripheral displays from those displays that are designed to attract focused attention and work at the action level, which we label *interruption displays.* Specifically, whereas peripheral displays are designed not to interrupt the flow of a user's primary action, interruption displays are disruptive and generally cause a user to concentrate on a different action and/or activity.

To illustrate the distinction between peripheral and interruption displays, consider Figure 1a. This figure schematically illustrates two consecutive actions taken by a user (Jane). In this case, Jane has three activities. Two of these activities are primary (A and C), and one is a secondary activity (B). At the start of this diagram, Jane consciously selects an action to work on that is related to activity A (e.g., writing a paragraph in a research paper). While she performs the operations that compose action A, she also performs operations related to B (e.g., glancing at a repetitive stress monitor). Note that some of Jane's operations combine into a conscious action (A), whereas other operations do not (B). When Jane completes the action associated with activity A she consciously selects a new action related to activity C (e.g., reading an email from her child's daycare center). Again, while she performs this action she performs operations related to activity B. The tool that allows Jane to complete operations on activity B while completing actions related to different primary activities is *a peripheral display.*

Figure 1. **Distinguishing peripheral and interruption displays.**

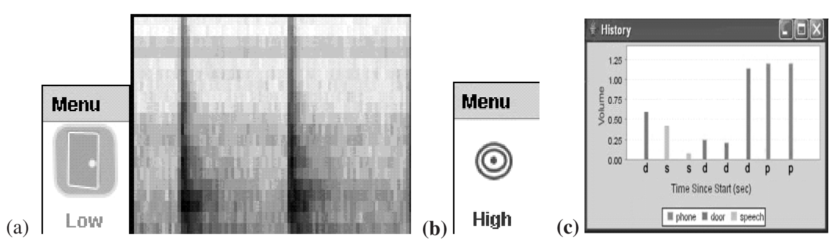

Now consider Figure 1b. In this case, Jane receives an interruption (e.g., a phone ring) that pertains to activity D (e.g., her child's health) while she is completing an action related to activity A. This interruption forces Jane to consciously switch her action. In fact, most interruption displays present information that is intended to instigate action/task switches. Although the information presented in a peripheral display might result in the user switching actions, the peripheral display should be designed such that a user is not interrupted and can choose to finish her current action first. We illustrate this type of interaction in Figure 1c, where the peripheral display shows the new state of activity D, but Jane finishes her action servicing activity A before switching.

Note that people will tend to operationalize routine processes over time, potentially mitigating the interpretability of a display. Tolmie, Pycock, Diggins, Maclean, and Karsenty (2001, p. 401) observed a mother adapt her use of an alarm indicating that it was time for her children to prepare for school. As the alarm sounded, the mother would continue her primary action, "translating [a text] from English into French," uninterrupted. In that case, a display that was initially designed to be interruptive became, through routine use, a peripheral display.

In practice, most displays will be interruptive initially, but a goal for peripheral display designers is to create a display that lends itself to operationalization. Conversely, operationalization presents a challenge for designers of interruptive displays.

3.6. Comparison With Existing Terminology

Existing literature has used various terms to describe displays operating on the periphery, including *ambient* and *notification* display. Definitions for both of these vary widely, though some common themes can be discussed within our framework.

Ambient displays are typically defined as aesthetic displays (Mankoff et al., 2003; Stasko, Miller, Pousman, Plaue, & Ullah, 2004), often integrated with the environment (Ishii & Ullmer, 1997; Stasko et al., 2004), and designed to convey information subtly (Ishii & Ullmer, 1997; Mankoff et al., 2003; Mynatt et al., 1998; Stasko et al., 2004). Because of their subtlety and focus on aesthetics, ambient displays tend to convey less critical information, for example, the Dangling String (Weiser & Brown, 1996) shows network traffic; AROMA (Pedersen & Sokoler, 1997) conveys remote person awareness; and InfoCanvas (Miller & Stasko, 2002) depicts traffic, weather, airfare, and stock information. Ambient displays can have any scope and can support primary, secondary, and/or pending activities. In general, ambient displays are a subset of peripheral displays since designs are typically intended for operational use (i.e., low-effort interpretation).

Pousman and Stasko (2006) defined *ambient information systems* as displays that (a) display information that is important but not critical, (b) can move from the periphery to the focus of attention and back again, (c) focus on tangible representations in the environment, (d) provide subtle changes to reflect updates in information (should not be distracting), and (e) are aesthetically pleasing and environmentally appropriate. Using this definition, they propose a taxonomy with four design dimensions that categorizes existing ambient displays and points to open areas for new design research. Unlike the framework we propose here, this ambient display taxonomy is not intended to address evaluation issues. Furthermore, this taxonomy is based solely on an examination of existing displays, whereas our definition and framework are based on theory regarding the way people use peripheral display.

The term *notification* display or system is often used as an umbrella term including peripheral displays. For example, McCrickard, Czerwinski, and Bartram (2003) define notification systems as applications that

> attempt to deliver current, important information to users in an efficient and effective manner without causing unwanted distraction to ongoing tasks. … These types of displays share the common design goal of providing the user with access to additional information without requiring excessive levels or prolonged periods of attention. (p. 510)

This definition implies that peripheral displays of more critical information and interrupting displays as we define them are notification systems. However, because low- to high-criticality displays are included in our definition of peripheral display and this definition specifies that displays deliver "important information," it is unclear if peripheral displays are a proper subset.

4. PERIPHERAL DISPLAYS IN CONTEXT: DESIGNDIMENSIONS

To help characterize how different contexts affect evaluation criteria, we derive three design dimensions for peripheral displays based on Activity Theory: scope of use, class(es) of supported activities, and criticality. These dimensions are primarily concerned with the activities that are influenced and serviced by the peripheral display. As previously discussed, these activities are the usage context for the display and consequently influence the success of a peripheral display design. After describing each design dimension, we list example peripheral displays that highlight the range of the dimensions.

4.1. Scope of Use

Peripheral displays are operation level tools in their user's activities. The scope of use of a peripheral display refers to the number of activities the display is designed to support. Note that users will likely appropriate peripheral displays to support a variety of activities not originally intended by the designer. But this dimension refers specifically to the number of activities that the display is designed to support, rather than future activities that the display may support.

Designers should consider whether they are creating a display that is either applicable to one specific activity or applicable to more than one activity. The former refers to displays that enable the user to perform operations that service only one activity. The latter includes peripheral displays whose associated operations service more than one activity. For example, a peripheral display that allows its user to monitor e-mail—for example, the e-mail Orb (Hsieh & Mankoff, 2003)—can support both work and personal activities. Note that we are concerned not with the ease of appropriating the display for a different activity but only with the number of activities the display is intended to support.

We exclude displays that do not service any activity (i.e., those displays that do not promote or empower the user in their efforts to satisfy the motives of any of their activities). For example, very low-fidelity monitoring devices, such as simple fire alarms that detect fires but do not provide the user with any additional information on how to deal with the fire, would categorically fall into this excluded group of displays. Ultimately, however, whether a display can service an activity depends on the user. For example, the altimeter in a cockpit is not a peripheral display for someone who does not pilot planes.

4.2. Class(es) of Supported Activities

People are generally involved in multiple ongoing activities but are focused on a small set of these through their primary action. The class(es) of supported activities design dimension captures whether the peripheral display supports primary activities, secondary activities, and/or pending activities (see the prior description of these classes of activities). Although it is natural to consider only two cases—peripheral displays that support primary activities and peripheral displays that do not—we further distinguish between displays that support secondary activities and displays that support pending activities.

Practically, this distinction allows the peripheral display designer to assess whether they need to design for transitions that enable a pending activity to become a primary activity or whether they need to design only for the monitoring of some aspect of a secondary activity. Scalable Fabric (Robertson et al., 2004) is an example of a peripheral display designed to support pending activities. Scalable Fabric users interact with windows in a central focus area of the screen in a normal manner, but when a user moves a window into the periphery, the window shrinks. Placing a set of windows near each other in the periphery dynamically creates a task composed of them. Window management is supported on a task level: Clicking on a task name in the periphery restores all the task windows to the focus area. Scalable Fabric enables users to monitor the status of pending actions/tasks that have been temporarily set aside but could be resumed at any time via the shrunken windows. The display also supports transitioning between different activities (by clicking on peripheral tasks), appropriately modifying its view when pending activities become primary activities (by restoring task windows to the focus area).

Even when a peripheral display does not support the user's primary activities, a designer should have some understanding of the user's primary activities and how the user's primary action supports them. For example, a peripheral display designer needs this understanding to choose an appropriate interaction modality. In some situations, interaction modalities carry different restrictions (e.g., many researchers have shown that when driving a vehicle, conveying various types of information audibly leads to better driving and in-vehicle task performance than conveying it visually; see Wickens & Seppelt, 2002, for a survey). Without assessing the user's primary action—a major component of the current situation and context—a designer may be unable to make appropriate design decisions regarding the display.

4.3. Criticality

The final design dimension is criticality, which refers to how critical or important the activities the display is designed to support are for the user. Criticality represents a continuous range from *low criticality* to *high criticality.* For simplicity, we compare and contrast these two endpoints but note that activity importance could fall anywhere in the range. Although some criticality distinctions are largely universal (e.g., most people would consider activities involving life or death situations as high criticality), Activity Theory does provide a handle on a more subjective sense of criticality. Because each activity for a user has an associated motive, one can gauge criticality by assessing the importance of the appropriate motives for the user—generally speaking, the more important the motive the more critical the activity.[3]

3. Leontiev (1978) addressed the gradation of criticality of all the different activities (and motives) that a person has by sorting them into a hierarchy. As we are only interested in the activities influenced by the peripheral display, this framework is beyond the scope of this article.

4.4. List of Example Peripheral Displays

Our design dimensions have meaning only in relation to specific people and their sets of activities. Next we describe two simple personas. Then we present example displays for each of our three design dimensions and how they might be used. These examples serve to illustrate the connection between a peripheral display, its use, and the activities of the user.

Personas

Bob is a social services coordinator who exercises regularly, owns a home, and flies single-engine planes in his spare time. **Jane** is a computer science professor who is partially deaf, a single mother of a toddler, and a frequent user of public transportation.

Peripheral Display Examples

Scope of 1, Primary Activity, Low Criticality. It is nearly time to head home, and Jane is wrapping up loose ends and getting organized to leave work. Her office has an information ticker that indicates bus arrival and departure times. Because Jane is planning to ride the bus, she peripherally monitors the ticker to make sure she finishes her wrap-up in time to catch the next bus. Here, the ticker supports one specific activity for Jane: managing her work–life balance. This is Jane's primary activity, and the information displayed in the ticker is low criticality (Jane could walk home if she misses the bus).

Scope of 1, Secondary Activity, Low Criticality. Jane is concentrating on writing a research paper, and a repetitive-stress monitoring program presents the length of time she has gone without taking a break. Here, there is one activity that the display supports (maintaining her health), the supported activity is secondary to her primary activity (writing a research paper), and the supported activity is relatively low criticality.

Scope of 1, Primary Activity, High Criticality. Bob is monitoring an altimeter while flying a plane. Here, there is one activity that the display supports (flying a plane), it is Bob's primary activity, and it is highly critical (negligence of the activity would risk Bob's life). Note that the altimeter is only peripheral for trained pilots who have operationalized its interpretation. The altimeter is an example of a more complex visual display that may require more learning to reach operational or peripheral use.

Scope of 1, Secondary/Pending Activity, High–Criticality. While Jane is cooking, she monitors a high-fidelity visual display showing her baby playing in another room. In this case, the display supports one activity (maintaining the health of her baby),which is highly critical. This activity could be secondary if Jane's baby is relatively independent and Jane is not expecting to directly attend to her baby in the near future. Or, if Jane expects her baby to need her direct assistance in the near future, this could be a pending activity.

Scope of 2, Primary and Secondary Activities, Low–Criticality. Jane is using a version of IC2Hear (Matthews, Fong, Ho-Ching, & Mankoff, 2006), an application that visually displays sound information, to monitor audience noise and to gain feedback of her own voice level while she teaches a class. Although IC2Hear can support low- and high-criticality activities, here the display supports two low-criticality activities: managing a class and practicing public speaking. The former activity is primary, and the latter is secondary.

Scope of 2, Pending Activities, Low Criticality. Jane is using Scalable Fabric (Robertson et al., 2004) to manage research projects that all have approaching deadlines. Her primary action is writing part of a research paper in Activity A, but she is also waiting for an e-mail to finish a grant proposal for Activity B and for some data-processing algorithms to finish in Activity C. Jane is using Scalable Fabric to monitor the arrival of the e-mail she needs and to determine when the data processing has finished. Here, there are two supported activities, both of which are low-criticality, pending activities for Jane. The display also supports transitioning between different activities. For example, when the e-mail arrives for Activity B (Jane's grant proposal) Scalable Fabric shows that there is an e-mail in Jane's inbox. But when Jane clicks on the inbox to read the e-mail, Activity B becomes the primary activity, and Scalable Fabric appropriately modifies its view of Activity B.

Scope of 2, Primary and Secondary Activities, High and Low Criticality. Bob uses his heart monitor while at work to monitor his stress and fitness level. He also uses his heart monitor while flying to gauge his susceptibility to g-force blackouts. Here, the display supports two activities: health monitoring and flying. The former is nonprimary and low criticality, the latter is primary and high criticality.

Scope of Many, Secondary Activities, High and Low Criticality. When Bob is busy talking with people over the phone or during meetings, he monitors a display that indicates the number of high-priority e-mails remaining in his inbox. Here, the display supports more than one specific activity (Bob's high-priority e-mails relate to different work activities as well as some of his personal activities), which are secondary and of varying levels of criticality.

5. EVALUATION CRITERIA

Traditional user interfaces that support primary activities are typically evaluated on criteria related to efficiency and effectiveness, such as time to complete supporting actions, success rate of completing actions, number of errors, and quality of the resulting object. However, because peripheral displays are not designed for direct interaction, it is often harder to assess the influence of these displays on a person's efficiency or overall work effectiveness. Hence, different evaluation criteria are required for peripheral displays.

This section presents five evaluation criteria for peripheral displays that we derived from interviews with peripheral display creators, from our own design and evaluation experience, and by reviewing existing literature. The criteria are *appeal*, *learnability*, *awareness*, *effects of breakdowns*, and *distraction*. Our review of past literature focused on peripheral display evaluation research, namely, the ambient heuristics (Mankoff et al., 2003); the Context of Use Evaluation for Peripheral Displays (CUEPD) method (Shami et al., 2005); and the Interruption, Reaction, and Comprehension (IRC) classification model for notification systems (McCrickard, Chewar, et al., 2003). Criteria suggested by these works fall into a subset of our criteria: appeal, awareness, effects of breakdowns, and distraction. We add learnability because it is essential to operationalization.

In the following subsections, we discuss related peripheral display evaluation literature, define our evaluation criteria, and discuss how to evaluate the criteria based on Activity Theory for various peripheral display designs.

5.1. Related Work in Peripheral Display Evaluation

Past peripheral display research has provided several specialized evaluation methods (Mankoff et al., 2003; Shami et al., 2005) and a framework for evaluating notification systems (McCrickard, Chewar, et al., 2003). The specialized methods provide heuristics (Mankoff et al., 2003) or metrics (Shami et al., 2005) either to measure the success of a display or to detect problems. Our goal is to generalize and extend these approaches by formalizing a more general set of evaluation criteria for peripheral displays.

Examining the heuristics and metrics embedded in existing methods reveals underlying criteria. For example, the CUEPD method (Shami et al., 2005) includes survey questions for display evaluators that ask about various concrete issues related to the general criteria of noticeability, comprehension, relevance, division of attention, and engagement. The ambient heuristics

(Mankoff et al., 2003) similarly focus on concrete usability issues that relate to more general criteria influencing the adoption and use of peripheral displays. By specifying evaluation criteria instead of specific metrics (that target various concrete issues related to these criteria), we can highlight common evaluation issues across varying contexts and situations. The evaluation criteria we present in the rest of this section are based partially on an analysis of existing methods, with revisions and extensions that are grounded in our Activity Theory framework.

Our Activity Theory approach defines a set of design dimensions and evaluation criteria that is explicitly dependent on the user's multiple, ongoing activities. This distinguishes our work from prior work, which has not provided clear guidelines for peripheral display evaluation that considers the user's activities. One model, IRC, classifies notification systems and is inclusive of peripheral displays (McCrickard, Chewar, et al., 2003). McCrickard, Chewar, et al. applied the IRC model to evaluation by using levels of interruption, reaction, and comprehension as general criteria for evaluation metrics (e.g., they modified a user study questionnaire to ask about these criteria). Like our framework, the IRC model describes displays varying in criticality. However, the IRC model does not explicitly contain any notion of the user's context (i.e., their activities). Although they recognized the importance of context, they assumed that the designer can aptly map his or her understanding of the user's activities to a position in the IRC space unassisted. Instead of supporting this mapping process, the IRC framework focuses on and supports the evaluation of low-level interaction with a display—with the primary objective of determining how effective the display is from a human information processor perspective.

Our approach to peripheral display evaluation relies on five criteria: appeal, learnability, awareness, effects of breakdowns, and distraction. The importance of these evaluation criteria depends on dimensions that capture relevant activity-based context: scope, activity class(es), and criticality of the display. We describe our evaluation criteria in the following sections.

5.2. Appeal (Usefulness, Aesthetics)

Appeal refers to users' qualitative enjoyment of a display. In other words, this criterion represents their overall feelings about the display. This criterion can be broken down into *usefulness* and *aesthetics*, which represent different aspects of users' qualitative feelings about a display.

Designers we interviewed mentioned usefulness and aesthetics as important when evaluating peripheral displays. They noted that adoption of a peripheral display depends on its appeal to users. In addition, this criterion is informed by the ambient display heuristics presented in (Mankoff et al., 2003), several of which affect the appeal of a display: "aesthetic and pleasing design," "useful and relevant information," and "match between design of ambient display and environment" (p. 172). Appeal is also related to the "engagement" survey category that is part of the CUEPD method (Shami et al., 2005, p. 582), which suggests that user feedback be gathered on the display's attractiveness and usage enjoyment.

Activities supported by peripheral displays are typically already ripe with other tools, social contexts, and semantic themes. By understanding these aspects of the user's activities the designer can conceptualize what displays might be useful and aesthetically pleasing. This type of user profiling is common practice in many design methods, but we believe designers will be better equipped in assessing the appeal of their displays by explicitly considering a user's activities.

5.3. Learnability

Learnability is the amount of time and effort required for users to operationalize their use of a peripheral display. Although the notion of reaching operational use of a system is not novel to interaction design (John & Kieras, 1996), the systems designed using these methods are often focused on conscious, action-level interaction. Hence they are concerned with the user's ability to form and execute goals (John & Kieras, 1996). Alternatively, our learnability criteria specifically exclude action-level interaction–if the user is making conscious goals related to the interaction with or content of a display, it is not being used peripherally. Of course, even well-designed peripheral displays will occasionally become the focus of their user's attention. The design goal, as emphasized by our Activity Theory framework and captured by the learnability evaluation criterion, is for peripheral display users to be able to use the display predominately at the operation level. In general, displays that are quick and easy to interpret are more likely to be used peripherally, as they will require less practice to become operational. The learnability of a display may influence its adoption, because users may be less likely to use a display they find difficult to learn, unless interpreting it is meant to present a challenge and users expect this (Hallnäs & Redström, 2001).

Past literature on the design of peripheral displays has not explicitly focused on learnability as an evaluation criterion. Instead the focus was on the design qualities that enable quick or easy operationalization, rather than on the user's operationalization process. For example, the CUEPD survey (Shami et al., 2005) asks if the user was able to understand information just by glancing at it. This implies that glanceability, or quick and easy interpretation, is important. Likewise, the ambient heuristics (Mankoff et al., 2003) call for a "consistent and intuitive mapping," so that users spend less effort learning the mappings (p. 172). Although easy operationalization is appropriate for many peripheral displays, we acknowledge that effective peripheral displays may require effort to learn to use.

Activity Theory suggests that it would be most beneficial to evaluate learnability based on the activities in which the display is used and on feedback from users. Every person will require a different amount of practice to use a new display peripherally (i.e., at the operation level). In Activity Theory, new tools are often traced from the activity level (when the tool is first introduced and the user is trying to learn about it and how it relates to other tools) to the action level (when the user still has to direct his or her attention to the tool to use it appropriately) to the operation level (when the user can use the tool without high cognitive load; Leontiev, 1978). The rate of this learning depends on the user as well as the situation.

5.4. Awareness

Definitions of awareness widely vary in past literature (Schmidt, 2002). For this article we focus on peripheral awareness, which we define as the amount of information shown by the display that people are able to register and use without focal attention. For our purposes, peripheral awareness includes both consciously and unconsciously registered information. Because the purpose of a peripheral display is to convey some information, it follows that the user's awareness of that information can be used to judge the effectiveness of the display.

Interview participants said user awareness of information is an important criterion for evaluating peripheral displays. Our past studies of Hebb (Carter et al., 2004), sound displays for the deaf (Matthews, Fong, et al., 2006), and Scalable Fabric (Matthews, Czerwinski, Robertson, & Tan, 2006) have explored measuring awareness. Many studies in peripheral display literature have also evaluated awareness as an important criterion (Hsieh & Mankoff, 2003; Mankoff et al., 2003; McCrickard & Zhao, 2000; Mynatt et al., 1998; Mynatt, Rowan, Craighill, & Jacobs, 2001; Pedersen & Sokoler, 1997; Plaue, Miller, & Stasko, 2004).

Evaluation literature also considers awareness as an important criterion. In the IRC model from McCrickard, Chewar, et al. (2003), both reaction and comprehension are similar to awareness. Reaction is the speed and accuracy of the user's response to the information provided by the display. Comprehension refers to the user's ability to make sense of the information displayed and remember it at a later time. Because these criteria focus on immediate response and conscious registration, they are more appropriate for notification and interruption displays than for peripheral displays. In particular, peripheral displays often do not result in an observable reaction and reaction and comprehension are not applicable to unconsciously registered information.

The ambient heuristics (Mankoff et al., 2003) also indicate that awareness is important and prescribe that "useful and relevant information" is visible (p. 172). Similarly, the CUEPD survey (Shami et al., 2005) asks if users were "able to understand the information in the display" (p. 582).

Like appeal and learnability, awareness will depend on the user's activities. It is important to focus on how the information in the peripheral display relates to the other information in the user's activities. Understanding this relationship will help the designer assess how aware the user needs to be about information in the display, how often it should be updated, and how often the user should be monitoring the display.

5.5. Effects of Breakdowns

The effects of breakdowns refer to how apparent breakdowns are to users and how easily users can recover from them. It is apparent when most tools suffer breakdowns. For example, when a mouse suddenly stops responding, it suddenly becomes for the user a tool that requires repairing, whereas before it had been an unremarkable means of production–to use Heidegger's (1927/1962) phrases, the mouse becomes *present-at-hand* when it had been *ready-to-hand.* However, by their nature peripheral displays are unlikely to become present-at-hand even with they suffer breakdowns. Thus, people may use the display even when the information it presents is misleading, which eventually can lead to disruptions and ultimately rejection of the display. When breakdowns are made obvious and recovery is straightforward, users are more likely to adopt displays.

The ambient heuristics point to the effects of breakdowns, saying that "error prevention" is an important design consideration, because "users should be able to distinguish between an inactive display and a broken display" (Mankoff et al., 2003, p. 172). In an evaluation of the Bus Mobile, which shows bus schedule information, the visibility of breakdowns was shown to be a problem. The state signified by the bus tokens still underneath the white skirt had two possible meanings: no buses are scheduled or a motor is broken. Users could not tell the difference. One interview participant described a situation in which a different peripheral display designed to show levels of activity on a main server broke down and caused mild panic in a lab group. At one point, an error in the display caused it to freeze. Users interpreted this to mean that the main server had frozen, and they frantically searched for problems with it. Also, in a past field study of the e-mail Orb (Hsieh & Mankoff, 2003), the Orb was not displaying anything for half a day before users noticed. These examples show the importance of making breakdowns apparent and recovery easy.

As with the earlier criteria, understanding the effects of breakdowns depends on the user's activities. If the designer understands what aspects of the user's activity are influenced by the peripheral display and how these aspects relate to the activity as a whole, the designer can make more appropriate decisions on how to expose breakdowns and how to support recovery. Understanding the criticality of the activity and the scope of the display are also important. Breakdowns in the gauges in a cockpit could represent a serious concern that needs to be communicated to the pilot. Likewise, if the display supports multiple activities (scope greater than one) it is important to consider interactions between activities that could influence how to make users aware of the breakdown and to support efficient recovery. For example, in the scenario where Bob uses a heart monitor to check on his health as well as factors that could influence his ability to fly, a breakdown in the heart monitor needs to be conveyed in a manner that alerts Bob to the failure of the device but does not distract him from flying.

5.6. Distraction

Distraction is the amount of attention the display attracts away from a user's primary action. Distraction is important because it affects the user's ability to carry out his or her primary action and will likely influence their qualitative reactions to using the display.

Interview participants named distraction as an important criterion, saying that a crucial measure of success for *peripheral* displays is that they *be peripheral* and not unnecessarily distract the user. Past studies of the Bus Mobile (Mankoff et al., 2003), Hebb (Carter et al., 2004), sound displays for the deaf (Matthews, Fong, et al., 2006), and e-mail displays (Hsieh & Mankoff, 2003) have all evaluated distraction, often (but not always) with a goal of minimizing it.

Past research also implies that distraction is an important criterion. In the IRC model (McCrickard, Chewar, et al., 2003), interruption describes the event that causes the user to switch their focal attention to the notification. Interruptions cause distraction from a primary task. The ambient heuristics (Mankoff et al., 2003) prescribe that a display "should be unobtrusive and remain so unless it requires the user's attention" (p. 172). This implies that measuring distraction is useful. Finally, the CUEPD survey (Shami et al., 2005) asks several questions about distraction (e.g., did the user notice the display and was the user able to adequately focus on his or her primary task).

Distraction is a natural criterion given our Activity Theory analysis. Because the display is monitored at the same time as the user is performing an action, it will require some user attention. This does not necessarily result in the user being less efficient. For example, the display could provide information that is useful and/or important to the user's action, enabling him or her to perform better with the display than without.

In the next section we describe how Activity Theory and our design dimensions constrain the application and evaluation of these criteria.

6. GUIDELINES FOR DESIGNING AND EVALUATING DISPLAYS

In this section we discuss the evaluation of peripheral displays relative to our design dimensions: scope, classes of activities supported, and criticality. In particular, we discuss how criteria will vary in importance depending on a display's position along each design dimension (e.g., for displays with *high criticality*, *awareness* is more important, whereas certain aspects of *appeal*, like aesthetics, are less important). We also discuss the primary issues involved in the evaluation of peripheral displays that support serendipity. There are two reasons for specifically addressing serendipity: First, some peripheral displays specifically target serendipitous interaction (Carter et al., 2004; Kaye, 2001; Mynatt et al., 1998; Pedersen & Sokoler, 1997); second, because serendipity implies novelty for the user at some level (action and/or activity), knowledge of their existing activities may not provide enough structure for designing and evaluating serendipitous peripheral displays.

For scope, we distinguish between displays that are associated with one activity and displays that are associated with more than one activity. For classes of supported activities, we discuss peripheral displays that support the user's primary activity (e.g., the altimeter in a cockpit), displays that support secondary activities (e.g., a bus schedule display), and displays supporting pending activities–for example, Scalable Fabric (Robertson et al., 2004). For criticality, we discuss peripheral displays that are associated with highly critical activities (e.g., the altimeter) and displays associated with low-criticality activities (e.g., a bus schedule display).

In our experience, scope, classes of supported activities, and criticality can be treated as independent dimensions. Accordingly, we present how each dimension influences peripheral display design separately. We proposed that evaluators combine the evaluation recommendations presented next depending on where their display falls within these dimensions.

6.1. Scope

In general, changes in scope do not change the relative importance of our evaluation criteria; for example, knowing that a peripheral display supports one activity instead of two does not make appeal any more or less important. However, as scope increases, evaluating each criterion can become significantly more complicated. In this section, we describe additional complications.

- *Appeal*: As scope increases, it becomes increasingly important to consider the impact of the peripheral display in more situations and contexts. Consequently, the appeal of the display will have more constraints. For example, if a peripheral display supports both work and personal activities, it is important to evaluate the extent to which it harmonizes with the potentially different aesthetics of both domains.
- *Learnability*: As scope increases, it is important to consider how the user will learn to use the peripheral display in each of the activities that are supported. In some cases, a single method of learning is sufficient. However, there are situations where learning to use the display would differ depending on the user's situation. For example, InfoCanvas displays different streams of information, each supporting a different activity and each with different graphical representations that users need to learn (Stasko et al., 2004).
- *Awareness*: As scope increases, the number of situations and contexts that the peripheral display is used in also increases. Thus, measuring a user's awareness could become more complicated. For example, if an e-mail monitor is designed to support both work and personal activities, the user's ability to maintain awareness of displayed information may change as the relative importance of the supported activities changes. In this case, it is important to understand how different situations (e.g., upcoming work deadlines, important personal matters, lifestyle changes) alter the user's awareness of the display.
- *Effects of breakdowns*: As scope increases, the number of potential usage situations also increases. To assess the effects of breakdowns, evaluators need to consider how a breakdown influences each supported activity individually. For example, in the earlier scenario of Jane using IC2Hear, if the display malfunctions and indicates that she is talking far too quietly, Jane's reaction will influence not only her primary activity of managing her class but also her public speaking activity.
- *Distraction*: Generally, as scope increases, assessing distraction becomes more complicated. If a peripheral display conveys information about many activities, there are more opportunities for distractions. However, with any scope, it is important to assess not only activity-specific distractions (e.g., IC2Hear distracting Jane from managing her class) but cross-activity distractions like an important personal e-mail causing Bob to step out of a meeting.

Lab experiments can be particularly useful in examining support for a single activity in depth. But as scope increases, controlling lab experiments to handle all the possible interactions of the supporting activities becomes difficult.

6.2. High-Criticality Displays

Peripheral displays associated with a high-criticality activity are not tools for opportunistic information, curiosity, or aesthetic appeal. Thus, in terms of appeal, usefulness is more important than aesthetics. Below we describe how appeal, and the other evaluation criteria, should be adjusted for displays supporting high-criticality activities.

- *Appeal*: Highly critical information displays that are perceived as not useful may be less effective. Thus qualitative usefulness is more important to maximize than aesthetics. It is also more important to evaluate how well the display fits with its environment than to evaluate its aesthetic appeal. For example, it is less important that a cockpit gauge is independently aesthetic–it is more important that the gauge not unduly stand out in its environment, making it likely to garner more attention than it deserves.
- *Learnability*: If the display is initially introduced to the user in the context of the high-criticality activity, it is important to demonstrate that the display is readily learnable. This is crucial because high-criticality displays are related to an important motive and should not distract the user from focusing on the associated activity, even during the learning period. For example, the altimeter gauge in a cockpit is meant to support quick access to important information, and changing its design could have grave consequences. Alternatively, if the display is introduced into a simulation of the high-criticality activity and context, then it might be acceptable to have a longer learning curve, provided the user can actually reach operational use before the display is deployed in the real situation.
- *Awareness*: It is important to evaluate the extent to which users register information that the display communicates. For the altimeter example, a well-designed gauge might allow pilots to be aware of normal information as they deem appropriate but alert the pilot when the plane's altitude is abnormal.
- *Effects of breakdowns*: It is crucial to reveal breakdowns to the user, because they may depend on displayed information to support their high-criticality activity. It is important to convey notification of the breakdown quickly and clearly, and recovery should be easy for the user. For example, a pilot would need to know immediately if the altimeter broke and would need a way to remedy the problem.
- *Distraction*: For high-criticality displays, it is particularly important that distractions caused by a high-criticality activity display correspond to the importance of new information being shown (i.e., more important information warrants more distracting updates). It is important that pe-

ripheral displays of highly critical information be more noticeable, but they do not typically force users to switch to a different action. For example, a pilot might be preparing the landing gear when the altimeter indicates that the plane's altitude is too low. This information should be shown in a way that is emphasized so that the pilot is more likely to notice it. In rare cases, the display may be distracting enough to cause the pilot to abandon her current action and start a new one, such as steering the plane upward. However, a display that interrupts the user regularly is not peripheral and therefore requires different approaches to design and evaluation.

6.3. Low-Criticality Displays

A low-criticality activity display is a tool for maintaining awareness related to activities of low importance. For example the Bus Mobile (Mankoff et al., 2003) provides bus schedule information for an activity of managing a work–life balance. Low-criticality activity displays are often characterized as being artifacts for opportunistic information, curiosity, or aesthetic appeal rather than for productivity. This means that is it typically important to maximize appeal and minimize distraction. It also implies different user needs and expectations for learnability, awareness, and knowledge of breakdowns. Next we describe how each criterion should be adjusted for displays that support activities that are less critical.

- *Appeal*: This is an important criterion because the display is low criticality and likely used by choice rather than necessity. Users will not adopt a display unless it is aesthetically pleasing or useful. The level of harmony between the display and its environment is a matter of aesthetics, for example, Informative Art (Redström, Skog, & Hallnäs, 2000).
- *Learnability*: A low-criticality display need not necessarily be readily learnable, but it is important to match the learnability of the display with user expectations. For example, Slow Technology (Hallnäs & Redström, 2001) is designed to be slow to learn and understand to give people time to reflect.
- *Awareness*: Awareness levels will not always be high for low-criticality information. For example, users may be unaware of bus schedule information shown on the Bus Mobile until right before they decide to go home.
- *Effects of breakdowns*: Regardless of information criticality, it is important that users are alerted of breakdowns and there exists a simple, obvious recovery, method. However, it is less crucial that users know about breakdowns immediately than it is for high-criticality activity displays.

- *Distraction*: Similar to high-criticality displays, it is important that distractions caused by a display correspond to the importance of new information being shown. Because low-criticality displays often only show information of low importance, they should rarely distract the user from her primary action.

6.4. Displays Supporting Primary Activities

A primary-activity display supports the user's current activity and may influence the user's primary action (e.g., an altimeter supports the activity of safely flying the plane, which is also the pilot's primary activity). This means that the display is related to a specific context and a relatively stable set of actions that support the user's primary activity, making lab experiments more tenable. Criteria can be evaluated as follows:

- *Appeal*: This criterion is difficult to evaluate fully in controlled settings. One approach is to evaluate appeal qualitatively following use.
- *Learnability* and *awareness*: A user's awareness of information can be tested with knowledge questions regarding the displayed information. Gathering awareness data at various stages of the study gives an indication of how quickly and effectively the user learned to get information from the display.
- *Effects of breakdowns*: This criterion can be evaluated by simulating breakdowns and observing a user's response.
- *Distraction*: This criterion can be evaluated by measuring overall speed and success. For example, in the activity of flying safely, the user's current goal might be to land the plane. The altimeter gauge might indicate that the plane is approaching safely but distract the user so that she makes a mistake.

6.5. Displays Supporting Secondary Activities

Because secondary activities do not typically become primary activities, progress is made by monitoring their status. For example, the Bus Mobile (Mankoff et al., 2003) supports the activity of managing a work–life balance by providing users with help in deciding when to leave work and head home. While waiting for a bus to be close enough, a user could be working on another activity, such as writing a research paper.

Because the user's primary action and activities generally define their context (Nardi, 1996), a peripheral display that supports only secondary activities could be used in a variety of contexts that may vary and change. It follows that the context in which the peripheral display is used can be difficult to simulate realistically and evaluation may require an extended deployment (i.e., weeks). A field study uses primarily qualitative measures to evaluate criteria. For this reason, a display should be deployed for an extended period of time (i.e., weeks) so users can develop accurate qualitative feedback. In such a field study, each criterion can be evaluated via qualitative feedback from users during and after deployment.

Secondary activity displays do not impose any priority on evaluation criteria. Rather, it is important to determine the display's other design dimensions (criticality and scope) before deciding how to evaluate each criterion. For example, ensuring that a display is easy to learn is more important for high-criticality displays than for low-criticality displays; the fact that the display supports secondary activities does not affect this criterion.

6.6. Displays Supporting Pending Activities

Pending activity displays are designed to allow users to monitor nonprimary activities that may become primary in the near future. This type of monitoring does not involve making progress on the activity; rather, the user switches to the activity to make progress on it. For example, Scalable Fabric (Robertson et al., 2004) enables users to monitor the status of actions that have been temporarily set aside but could be resumed at any time. For example, users might set aside activities while waiting for a file to download, for a coworker to provide some information, or for a paper to finish printing.

Switches between a primary activity and the pending activity can be simulated during laboratory tests in a semirealistic way. For example, a study could involve users monitoring a peripheral display of e-mail updates, so that they can switch to an e-mail activity when certain messages arrive, while editing a document as a primary activity. When evaluating displays that support transitioning between multiple primary and pending activities, it is important to evaluate support for the transitions themselves (i.e., a representation change from pending to primary activity). For example, evaluating Scalable Fabric (Robertson et al., 2004) might involve measuring difficulty or distraction caused by moving focal windows into the periphery and bringing peripheral windows into the focal area.

Like secondary activity displays, pending activity displays do not impose any importance level on evaluation criteria. Instead, it is essential to examine the other design dimensions (criticality and scope) to determine the priority of each criterion. However, it is important to evaluate the learnability of the display's support for switching between activities. Ideally, it should be easy to learn how to switch between the primary and pending activities.

6.7. Designing for Serendipity

Designers often want displays they develop to support the discovery of a useful and/or interesting thing by accident (i.e., without a preconceived plan to either search for or find it). The consequence of such a serendipitous encounter is a (hopefully positive) change in the user's behavior (e.g., saying new or different things, making different decisions, enjoying different situations). We argue that serendipitous behavior changes are easier to predict (leading to more effective displays) when they occur at the level of actions rather than activities. We specifically exclude serendipitous behavior changes at the operation level because they are generally initiated at the action or activity levels first (Leontiev, 1978). An example of a serendipitous behavior change at the action level is a knowledge worker who accidentally finds an online database that he begins to check regularly for information, whereas at the activity level a person might start a new lifelong relationship with someone she met at a party.

In general, because activities define a stable context for actions (Leontiev, 1978; Nardi, 1996), they can be used as a structure within which action level behavior changes can be measured and predicted. It is precisely this stable structure that makes serendipity at the action level easier to design for than at the activity level.[4] In this case, the activities whose associated actions have changed could be primary, secondary, or pending activities for the user. For example, suppose a person is focused on his primary work activity and notices on a group awareness display–such as in Carter et al. (2004)—that another group in the office is working on a similar project. This person makes it a goal to talk with someone in that group as soon as possible. In this example, the peripheral display designer could use knowledge of the user's current activities to design for this serendipitous information encounter.

Factors leading to a new action depend on complex interactions between the user, her existing activities (and their associated motives), and the other potentially influential people and artifacts in her environment. It is therefore difficult to simulate serendipity in a way that is realistic enough for users to be-

4. The assumption of a stable structure enveloping a design space underlies many design methods (although the stable structures they posit may differ conceptually). If one assumes that the structure surrounding the designed changes is not stable, either because the designer is attempting to change it or because the user is changing, the complexity of the design process increases. These additional complexities are not necessarily negative. Often, they can push designs in new directions or reveal significant insights (Engeström, 1987; Spinuzzi, 2003). Of course, activities will change. But, they are relatively stable over longer time frames (weeks, months, or years; Leontiev, 1978). We address the inevitable change in a user's activities in the postdeployment assessment step of our design process.

gin new actions. For this reason, most existing peripheral displays intended to support serendipity were evaluated in the field (Carter et al., 2004; Mynatt et al., 1998; Pedersen & Sokoler, 1997).

Alternatively, the formation of new activities, from serendipitous encounters or otherwise, is not well understood. In fact, activity formation is still an open research question (Leontiev, 1978, 1981), although some frameworks have been proposed (Engeström, 1987). Consider the following example: A person notices in the group awareness display that three other groups are working on highly related projects, but in isolation from one another. This fact prompts the person to start looking for a new job in a different company. This new job-finding activity is hard to design for, as it extends beyond the user's current set of activities. Essentially, until the new activity has stabilized, it is hard to predict the motive associated with it.

7. CASE STUDY: IC2HEAR

Here we discuss an example peripheral display design and evaluation process in the context of our Activity Theory–based design dimensions and evaluation criteria. IC2Hear is a visual, peripheral display to help people who are deaf maintain awareness of sounds in their environment (Matthews, Fong, et al., 2006). In our deployment of IC2Hear, it was a high-scope, low- to high-criticality, primary, secondary, and pending activity display.

We analyze the methods we used in Matthews, Fong, et al. (2006) to gather knowledge of user needs and activities, design the IC2Hear application, and evaluate the prototypes. Results reported in Matthews, Fong, et al. focus on presenting knowledge about users' needs for nonspeech sound awareness and *design knowledge* for peripheral displays of sound for the deaf. In this article, we analyze the IC2Hear *design process* with our Activity Theory framework and present *evaluation suggestions* using our proposed criteria and guidelines. The process actually followed in Matthews, Fong, et al. did not use this Activity Theory approach. Thus we highlight mistakes made during the design process that could have been avoided if we had used the approach presented in this article.

7.1. Design Process

The goal of our design process was to understand sound awareness needs and challenges by gathering information about users and their activities. The end result of our design process was a set of three visual peripheral displays that visualized nearby sounds. At the beginning of the IC2Hear project, we interviewed 8 participants who are deaf to gather information about their activities and sound awareness needs. In the same session as the interview, we presented participants with sketches of potential sound awareness displays and asked their opinions. We built three of the most popular sound awareness displays and evaluated them in a formative lab study. In this section, we discuss the design process that occurred before any prototypes were built.

As with many user-centered approaches to design, Activity Theory emphasizes the need for a deep understanding of user motives and activities, the tools used, and other people involved. Our framework focuses this process for peripheral display designers on gathering information that will determine the three design dimensions of the display: scope, classes of activities supported, and criticality of information shown. Focusing the formative design work helps streamline the process and draws attention to important issues for peripheral displays. For example, highly critical information is so important to people that they may emphasize it in interviews. Using our framework, we know that low and highly critical information can be displayed in different ways, so it is important to gather detailed information about both to inform the different display mechanisms.

From interviews, we determined that IC2Hear was high scope, supported all classes of activities, and conveyed all criticalities of information. For example, people wanted sound awareness to support their work productivity–a higher criticality, primary activity (e.g., phone calls from clients). They also wanted sound awareness to support their ability to maintain social ties with coworkers–a lower criticality, secondary activity (e.g., coworkers chatting in the hall).

Unfortunately, we did not separate low- and high-criticality display mechanisms in our sketches, something our Activity Theory framework would have encouraged us to do. This limited the feedback we received from users and hence resulted in less successful prototypes. In particular, almost all of our design sketches used the same display mechanism for all types of information (both low and high criticality). For example, all information was displayed abstractly with colors and shapes regardless of importance. A few design sketches only showed important sounds, which were depicted with icons (phones, doorbells, voices, alarms). Participants gravitated toward displays that clearly depicted important sounds, although they actually wanted a combination of display mechanisms that could depict all sounds of interest (as we learned in later evaluations). One display combined both: a spectrograph placed next to an icon, where the spectrograph continuously conveyed literal sound information and the icon showed certain important sounds. This was a favorite among participants. In building the prototype, however, we made another mistake regarding learnability (discussed next).

The implication for our designs was that, because IC2Hear was a many-scope display, using multiple approaches to convey information of different criticality was important to supporting all the target activities. In particular, our designs should have focused on using subtle visuals for continuous sound awareness of low to moderately critical sounds and should have included some more interruptive visuals for highly critical sounds. The most critical sounds that participants listed were interruptions (e.g., the phone, alarms, doorbells, people trying to capture attention, cars honking, etc.). Whereas our visualization of highly critical sounds was designed for a decontextualized aesthetic (based only on the sound itself), we should have focused more on designs that acknowledged and harmonized with the environment in which these sounds occurred. This change in focus could have improved the user's awareness and the appropriateness of interruptions. Our sketches and prototypes might have been improved with careful consideration of the display's goal to make users aware of information ranging from low to high criticality.

Our Activity Theory framework's emphasis on separating information of differing criticality could also have improved feedback gathered in interviews. Participants emphasized highly critical sounds in interviews and in their comments about design sketches. We should have clearly separated more and less critical sounds in our discussion, gathering more data about how to support less critical sound awareness.

In summary, our display would have better supported user needs if we had gathered more information about participants' use of less critical sounds in their activities and had created a larger set of sketches that combined multiple mechanisms for conveying information of differing criticality. In the next sections, we present the IC2Hear prototypes and discuss their evaluation.

7.2. Prototypes

Based on feedback from interviews, we implemented three graphical displays: Spectrograph with Icon (Figure 2a), Single Icon (Figure 2b), and the History Display (Figure 2c). Spectrograph with Icon was a design favored by interviewees that combines continuous, detailed sound information and an icon for more critical sounds (phone rings, nearby voices, opening/closing doors). The icon portion displays a set of rings that indicated gross pitch and volume (pitch: low = blue, red = high; volume: low = few rings, high = many rings). The other portion is a standard spectrograph, a visualization of amplitude (darkness), frequency (y-axis patterns), over time (x-axis). In Figure 2a the icon shows a person knocking on a door. Single Icon is simply the icon portion of Spectrograph with Icon. Figure 2c shows the History Display, a bar graph of past sounds where color identifies the sound, height of the bar indicates the sound's volume (e.g., loud sounds make taller bars), and position along the x-axis indicates time (recent sounds are on the right). It shows higher criticality sounds only.

Figure 2. **Three implemented prototypes of IC2Hear, peripheral sound displays for the deaf. (a) Spectrograph with Icon showing knocking on a door. This initial prototype places the spectrograph and icon windows side-by-side. (b) Single Icon showing a quiet, high frequency, unrecognized sound. (c) The History Display showing phone ringing (labeled *p*), door sounds (*d*), and speech (*s*). The height of bars indicates volume, and the *x*-axis represents time (more recent sounds are on the right).**

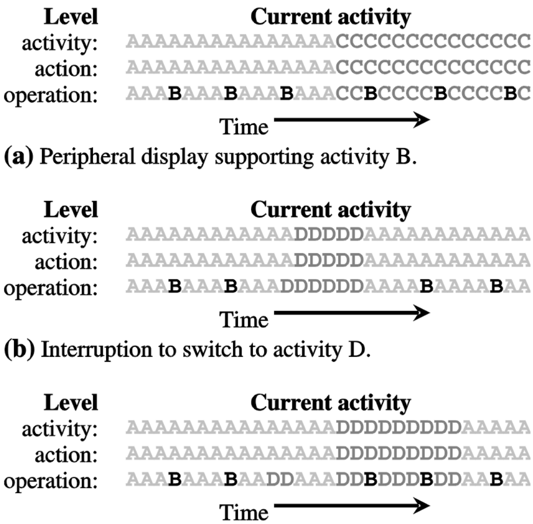

(a) Peripheral display supporting activity B.

(b) Interruption to switch to activity D.

(c) Peripheral display supporting activities B and D.

7.3. Evaluation

Here we discuss a formative, uncontrolled lab evaluation of IC2Hear, highlighting issues raised by our evaluation framework. In particular, the framework encourages designers to consider certain criteria in displays—learnability, awareness, effects of breakdowns and distraction, and appeal—and our evaluation showed that our prototypes did not adequately address all the criteria. A summative evaluation has not yet been conducted.

Four users who are deaf participated in our lab evaluation, intended to gather initial qualitative feedback on IC2Hear. We simulated an office setting, training our system to recognize important office sounds for the evaluation (phone, voices, door opening/closing, and door knocking). We did not explicitly design icons to show only important sounds; rather, we designed them to show any sounds *recognized* by our recognition system. However, in practice this ended up being the more critical sounds, because they tend to be more distinctive. During the study, users monitored the peripheral display while checking e-mail in a relaxed setting without any particular tasks. We simulated sounds in the office and asked users for feedback about the display. Data gathered in this session were entirely qualitative.

Our formative evaluation made it evident that none of our prototypes adequately address all criteria important to peripheral displays, a problem our Activity Theory framework may have directed us to consider. In particular, we may not have prototyped a design as difficult to learn as a spectrograph, we may have included different display mechanisms for information of differing criticality, and we may have better enabled users to detect breakdowns in the History Display.

Learning is necessary for a display's use to become operationalized or peripheral. From interviews, we knew that users expected the display to be easy to learn. In the formative evaluation, users were skeptical about using the spectrograph because they thought it was too difficult to learn. One participant said, "It's just a blurry picture. I have no idea what the sound is." Users thought Single Icon and the History Display were easy to learn, preferring these two displays over Spectrograph with Icon.

Realizing the importance of easy learnability early on could have greatly improved our success with IC2Hear. From the beginning, we should have inspected design sketches and eliminated or modified ones that would be hard to learn. It is well-known in speech therapy literature that "reading" spectrographs requires extensive practice for months (Zue & Cole, 1979). Several of our design sketches were similar to a spectrograph, using movement of colors and abstract shapes to represent sounds. Before building any display, we could have tested user learning and interpretation speeds for various sounds to ensure our designs would enable easy learning and information awareness.

Our interview participants favored the Spectrograph with Icon display on paper, which was likely due to a lack of alternatives that conveyed low- and high-criticality sounds in different ways. Our formative evaluation indicated that users preferred Single Icon and the History Display for learnability. However, the History Display only showed more critical events, and Single Icon showed only limited information about unrecognized events. To address this, we could have tested more designs for conveying a range of less critical sounds in an easy to learn way.

Users in the formative evaluation also felt the displays did a poor job of varying their display mechanisms based on information criticality. Users made it clear that they wanted overt interruptions for highly critical sounds and nondistracting updates for less critical sounds. They suggested that the icon flash when the sound was "really important" and be more "visual[ly] quiet" for less important sounds. The Single Icon and Spectrograph with Icon, in particular, did not adequately distinguish between sounds of varying criticality. For this reason, users favored the History Display, which inadvertently emphasized loud sounds (which users told us were more important) by using larger bars. However, it is likely that varying the sizes of bars would not be distracting enough for very important events like the phone or alarms. The design of interruptions for high-criticality sounds needs more exploration, as we did not adequately address this in our designs.

Finally, the evaluation showed that the History Display did not enable users to easily discover *breakdowns*, an important feature of displays showing high-criticality information. Single Icon and Spectrograph with Icon showed all sounds, so users could see the constantly updating display was working. The History Display, however, only showed recognized sounds (phone, voices, door opening/closing, door knocking), which happen less often. This could impede users from noticing a system breakdown. One possible solution would be to display unrecognized sounds on the History Display as well, an idea that we may have included in the display if we had considered the criteria emphasized by our framework.

To summarize, the three displays were preferred differently for different criteria: The History Display performed better on learnability, distraction, and awareness; Single Icon and Spectrograph with Icon showed more information and better enabled recovery from breakdowns; and Single Icon also enabled better learnability. None of these designs were strong on all criteria. Overall, this analysis using our framework, with its emphasis on design dimensions and criteria that are important to peripheral displays, helped reveal central issues that could guide future redesigns and evaluations.

8. CONCLUSION

In this article, we introduced a descriptive theory of peripheral displays based on Activity Theory. Specifically, a peripheral display is any information display that is a tool in at least one activity of its user and is used primarily at the *operation* level. Our Activity Theory framework extends past work on peripheral displays by highlighting the need for designers to focus on operationalization and on the activities that contextualize the use of the display.

One of the primary advantages of our framework is that it provides practical guidelines for designing and evaluating peripheral displays. Specifically, we used the framework to derive a set of important design dimensions for peripheral displays–scope, class(es) of supported activities, and criticality; to extend peripheral display evaluation criteria–appeal, learnability, awareness, effects of breakdowns, and distraction; and to understand how to evaluate different criteria in different contexts (i.e., in relation to different users and their activities).

The major limitation of our Activity Theory framework is that it does not alleviate the difficulties of applying design and evaluation methods. It guides the design and evaluation processes, but design and evaluation methods remain challenging and time-consuming to employ. Understanding all of the subtleties of user behavior would require an intensive triangulation of methods that concentrates on finding issues less likely to be captured by other methods, such as the display's affects on activities it was not expected to support. For example, it is likely that users will not be consciously aware enough of some behavior changes to provide useful self-report data, log analysis tools inevitably will fall short of describing all aspects of user activities, and participant observation is very time-consuming and subject to observer bias.

Despite this limitation, our Activity Theory framework is an important tool for designers that enables them to tailor their design process and evaluation methods to focus on issues that are important for peripheral displays. By focusing the design and evaluation processes on design dimensions (scope, classes of activities supported, and information criticality) and criteria (learnability, awareness, breakdowns, distraction, and appeal) that are important to peripheral displays, our framework helps designers address issues that are key to the success of their displays.

NOTES

Acknowledgments. We thank Bonnie Nardi, Jennifer Mankoff, and Amanda Williams for valuable feedback.

Support. This work was funded by National Science Foundation Graduate Fellowships.

Authors' Present Addresses. Tara Matthews, Berkeley Institute of Design, CS Department, University of California, Berkeley, CA 94720. E-mail: tmatthew@cs.berkeley.edu. Tye Rattenbury, Berkeley Institute of Design, CS Department, University of California, Berkeley, CA 94720. E-mail: rattenbt@cs.berkeley.edu. Scott Carter, Berkeley Institute of Design, CS Department, University of California, Berkeley, CA 94720. E-mail: sacarter@cs.berkeley.edu.

HCI Editorial Record. First manuscript received November 28, 2005. Revisions received April 12, 2006, August 25, 2006, September 28, 2006, and October 5, 2006. Accepted by Panos Markopoulos, Boris de Ruyter, and Wendy Mackay. Final manuscript received October 12, 2006. —*Editor*

REFERENCES

Begole, J., Tang, J., Smith, R., & Yankelovich, N. (2002). Work rhythms: Analyzing visualizations of awareness histories of distributed groups. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*. New York: ACM.

Cadiz, J., Venolia, G., Jancke, G., & Gupta, A. (2002). Designing and deploying an information awareness interface. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*. New York: ACM.

Carter, S., Mankoff, J., & Goddi, P. (2004). Building connections among loosely coupled groups: Hebb's rule at work. *Journal of Computer Supported Cooperative Work, 13,* 305–327.Collins, P., Shukla, S., & Redmiles, D. (2002). Activity Theory and system design: A view from the trenches. *Journal of Computer Supported Cooperative Work, 11,* 55–80.

Davenport, T., & Prusak, L. (1998). *Working knowledge: How organizations manage what they know.* Cambridge, MA: Harvard Business School Press.

Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW).* New York: ACM.

Dourish, P., & Bly, S. (1992). Portholes: Supporting awareness in a distributed work group. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Endsley, M. R. (1995). Measurement of situation awareness in dynamic systems. *Human Factors, 37,* 65–84.

Engeström, Y. (1987). *Learning by expanding: An Activity-Theoretical approach to developmental research.* Helsinki, Finland: OrientaKonsultit.

Engeström, Y., Miettinen, R., & Punamäki, R. (Eds.). (1999). *Perspectives on Activity Theory.* Cambridge, UK: Cambridge University Press.

Erickson, T., & Kellogg, W. (2000). Social translucence: An approach to designing systems that support social processes. *ACM Transactions on Computer-Human Interaction, 7,* 59–83.

Greenberg, S. (1996). Peepholes: Low cost awareness of one's community. *CHI Conference Companion.* New York: ACM.

Gross, T., Wirsam, W., & Graether, W. (2003). AwarenessMaps: Visualizing awareness in shared workspaces. *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems.* New York: ACM.

Hallnäs, L., & Redström, J. (2001). Slow technology–Designing for reflection. *Journal of Personal Ubiquitous Computing, 5,* 201–212.

Heath, C., & Luff, P. (1992). Collaboration and control: Crisis management and multimedia technology in London Underground line control rooms. *Journal of Computer Supported Cooperative Work, 1,* 69–94.

Heidegger, M. (1962). *Being and time* (Rev. ed., J. Macquarrie & E. Robinson, Trans.). New York: Harper and Row. (Original work published 1927)

Hsieh, G., & Mankoff, J. (2003). *A comparison of two peripheral displays for monitoring email: Measuring usability, awareness, and distraction* (Tech. Rep. No. CSD-03-1286). EECS Department, University of California, Berkeley, Berkeley, California.

Huang, E., Tullio, J., Costa, T., & McCarthy, J. (2002). Promoting awareness of work activities through peripheral displays. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

John, B. E., & Kieras, D. E. (1996). The GOMS family of user interface analysis techniques: Comparison and contrast. *ACM Transactions on Computer-Human Interaction, 3,* 320–351. New York: ACM.

Kaye, J. (2001). *Symbolic olfactory display.* Unpublished master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Kuutti, K. (1995). Activity theory as a potential framework for human-computer interaction research. In B. Nardi (Ed.), *Context and consciousness: Activity theory and human-computer interaction* (pp. 17–44). Cambridge, MA: MIT Press.

Leontiev, A. N. (1978). *Activity, consciousness, and personality.* Englewood Cliffs, NJ: Prentice Hall.

Leontiev, A. N. (1981). *Problems of the development of mind.* Moscow: Progress Press.

MacIntyre, B., Mynatt, E., Voida, S., Hansen, K., Tullio, J., & Corso, G. (2001). Support for multitasking and background awareness using interactive peripheral displays. *Proceedings of the ACM Symposium on User Interface Software and Technology.* New York: ACM.

Mankoff, J., & Dey, A. K. (2003). From conception to design: A practical guide to designing and evaluating ambient displays. In K. O'Hara & E. Churchill (Eds.), *Public and situated displays: Social and interactional aspects of shared display technologies* (pp. 210–230). Dordrecht, The Netherlands: Kluwer.

Mankoff, J., Dey, A. K., Hsieh, G., Kientz, J., Lederer, S., & Ames, M. (2003). Heuristic evaluation of ambient displays. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Matthews, T., Czerwinski, M., Robertson, G., & Tan, D. (2006). Clipping lists and change borders: Improving multitasking efficiency with peripheral information design. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Matthews, T., Fong, J., Ho-Ching, F. W., & Mankoff, J. (2006). Evaluating non-speech sound visualizations for the deaf. *Behaviour & Information Technology*, *25,* 333–351.

McCarthy, J., Costa, T., & Liongosari E. (2001). UniCast, OutCast & GroupCast: Three steps toward ubiquitous, peripheral displays. *Proceedings of the International Conference on Ubiquitous Computing (UbiComp).* New York: Springer.

McCrickard, D. S., Chewar, C. M., Somervell, J. P., & Ndiwalana, A. (2003). A model for notification systems evaluation—Assessing user goals for multitasking activity. *ACM Transactions on Computer-Human Interaction*, *10*, 312–338.

McCrickard, D. S., Czerwinski, M., & Bartram, L. (2003). Introduction: Design and evaluation of notification user interfaces. *International Journal of Human-Computer Studies*, *58*, 509–514.

McCrickard, D. S., & Zhao, Q. A. (2000). Supporting information awareness using animated widgets. *Proceedings of the USENIX Conference on Tcl/Tk (Tcl2K).*,USENIX.

Miller, T., & Stasko, J. (2002). Artistically conveying information with the InfoCanvas. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI).* New York: ACM.

Mohammed, S., & Dumville, B. (2001). Team mental models: Expanding theory and measurement through cross-disciplinary boundaries. *Journal of Organizational Behavior*, *22,* 89–106.

Mynatt, E. D., Back, M., Want, R., Baer, M., & Ellis, J. B. (1998). Designing Audio Aura. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Mynatt, E. D., Rowan, J., Craighill, S., & Jacobs, A. (2001). Digital family portraits: Supporting peace of mind for extended family members. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Nardi, B. (Ed.). (1996). *Context and consciousness: Activity Theory and human-computer interaction.* Cambridge, MA: MIT Press.

39.Nardi, B., & Redmiles, D. (Eds.). (2002). Activity theory and the practice of design [Special issue]. *Journal of Computer Supported Cooperative Work, 11*(1–2).

Pedersen, E. R., & Sokoler, T. (1997). AROMA: Abstract representation of presence supporting mutual awareness. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

Plaue, C., Miller, T., & Stasko, J. (2004). Is a picture worth a thousand words? An evaluation of information awareness displays. *Proceedings of the Conference on Graphics Interfaces (GI).* Wellesley, MA: AK Peters Ltd.

Pousman, Z., & Stasko, J. T. (2006). A taxonomy of ambient information systems: Four patterns of design. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI).* New York: ACM.

Redström, J., Skog, T., & Hallnäs, L. (2000). Informative art: Using amplified artworks as information displays. *Proceedings of the Conference on Designing Augmented Reality Environments (DARE).* New York: ACM.

Robertson, G., Horvitz, E., Czerwinski, M., Baudisch, P., Hutchings, D. R., Meyers, B., Robbins, D., & Smith, G. (2004). Scalable Fabric: Flexible task management. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI).* New York: ACM.

Sawhney, N., Wheeler, S., & Schmandt, C. (2001). Aware community portals: Shared information appliances for transitional spaces. *Personal and Ubiquitous Computing, 5,* 66–70.

Schmidt, K. (2002). The problem with "awareness": Introductory remarks on "Awareness in CSCW". *Journal of Computer Supported Cooperative Work, 11,* 285–298.

Shami, N. S., Leshed, G., & Klein, D. (2005). Context of use evaluation of peripheral displays. *Proceedings of the IFIP TC13 International Conference on Human Computer Interaction (INTERACT).* New York: Springer.

Spinuzzi, C. (2003). *Tracing genres through organizations: A sociocultural approach to information design.* Cambridge, MA: MIT Press.

Stasko, J., Miller, T., Pousman, Z., Plaue, C., & Ullah, O. (2004). Personalized peripheral information awareness through information art. *Proceedings of the International Conference on Ubiquitous Computing (UbiComp).* New York: Springer.

Tolmie, P., Pycock, J., Diggins, T., Maclean, A., & Karsenty, A. (2001). Unremarkable computing. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI).* New York: ACM.

van Dantzich, M., Robbins, D., Horvitz, E., & Czerwinski, M. (2002). Scope: Providing awareness of multiple notifications at a glance. *Proceedings of the ACM Conference on Advanced Visual Interfaces (AVI).* New York: ACM.

Weiser, M., & Brown, J. S. (1996). Designing calm technology. *PowerGrid Journal, 1*(1).

Wickens, C. D., & Seppelt, B. (2002). *Interference with driving or in-vehicle task information: The effects of auditory versus visual delivery* (Tech. Rep. No. AFHD-02-18/GM-02-3). Urbana, IL: University of Illinois Human Factors Division.

Zue, V., & Cole, R. (1979). Experiments on spectrogram reading. *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).* Los Alamitos, CA: IEEE Press.