

Correspondence Principles in Workspace Coordination

David Woods
The Ohio State University

Copyright © 1991, 1993, 1994, 1999, 2003 by David D. Woods
Draft, 8-4-03

In complex systems human activity ebbs and flows, with periods of lower activity and more self paced tasks interspersed with busy, high tempo, externally paced operations where task performance is more critical. These domains are characterized by multiple interleaved physical and cognitive activities where a shift from one activity to another may be triggered by knowledge or by new events in the world. In addition, the processes in the work domain are complex and can be modeled in multiple ways.

As a result, there are different kinds of views or perspectives one can take on the processes, there is a need to shift among these different views over task sequences and as situations evolve. The question then becomes how to design representations that support achieving correspondence in terms of shifting focus within a perspective or across perspectives and contrasting perspectives.

The workspace coordination level of analysis is concerned with coordinating multiple views into the processes of the work domain to create a virtual perceptual field or a workspace in which practitioners carry out their domain activities. Developing what kinds of views should be available in parallel or serially, deciding how to coordinate different views in pace with the interleaved and changing tasks of the practitioner, providing means to contrast and shift perspectives are all aspects of Design for Correspondence at this level of analysis.

Before considering the correspondence principles at this level of analysis we need to step back and examine a human skill competency that defines direct correspondence at this level of analysis—how people are able to focus attention on potentially “interesting” parts of a perceptual field despite the fact that what is interesting depends on context.

Direct Correspondence in Shifting Perspective or Knowing Where To Look Next

A fundamental competency of human perceptual systems is the ability to orient focal attention to “interesting” parts of the natural perceptual field (Rabbitt 1984; Wolfe 1992).

“The ability to look, listen, smell, taste, or feel requires an animal capable of orienting its body so that its eyes, ears, nose, mouth, or hands can be directed toward objects and relevant stimulation from objects. Lack of orientation to the ground or to the medium surrounding one, or to the earth below and the sky above, means inability to direct perceptual exploration in an adequate way (Reed, 1988, p. 227 on Gibson and perceptual exploration in Gibson, 1966).”

Both visual search studies and reading comprehension studies show that people are highly skilled at directing attention to aspects of the perceptual field that are of high potential relevance given the properties of the data field and the expectations and interests of the observer. Reviewing visual search studies, Woods (1984) commented, “When observers scan a visual scene or display, they tend to look at ‘informative’ areas . . . informativeness, defined as *some relation between the viewer and scene*, is an important determinant of eye movement patterns” (p. 231, italics in original). Reviewing reading comprehension studies, Bower and Morrow (1990) wrote, “The principle . . . is that readers direct their attention to places where significant events are likely to occur. The significant events . . . are usually those that facilitate or block the goals and plans of the protagonist.”

How do we focus in on areas of high potential relevance when what is relevant depends on context? How do we shift attentional focus in a changing environment where new events may require a shift in attentional focus (or action) at indeterminate times? Basically, the ability to notice potentially interesting events and know where to look next (where to focus attention next) in natural perceptual fields depends on the coordination between orienting perceptual systems (i.e., the auditory system and peripheral vision) and focal perception and attention (e.g., foveal vision). The coordination between these mechanisms allows us to achieve a “balance between the rigidity necessary to ensure that potentially important environmental events do not go unprocessed and the flexibility to adapt to changing behavioral goals and circumstances” (Folk et al. 1992, p. 1043).

The search for solutions to data overload shifts when we re-formulate the problem of navigation as supporting how one knows where to look next in a changing perceptual field. To accomplish this at the workspace coordination level we must think first about sequences of tasks, about what data needs to be seen in parallel in different contexts, and about how attention shifts to interrupt signals and new events.

Costs of low or indirect correspondence in workspace coordination

Human competency at shifting focus defines direct correspondence. But when we shift to situations where computerized systems mediate our access to the world in question, the question becomes how to achieve the kind of directness that characterizes our perception of meaning in natural situations. Indirect correspondence or clumsy workspace coordination place new cognitive burdens on practitioners and increase the risk of cognitive breakdowns.

Studies indicate that when workspace coordination is clumsy:

- Across display transitions can be costly cognitively. These costs can include new knowledge demands, new memory demands, new tasks such as decluttering, and these new burdens can congregate at the wrong times.
- The balance between these costs and potential information gain affect how practitioners interact with a virtual data space available through a narrow keyhole.

- Supporting shifts between different views or types of views is important in workspace coordination. Designing the set of views and how they can be orchestrated is based on what tasks need to be done in parallel versus serially and how one moves from one task to another.
- Effective workspace coordination is more than navigation efficiency. Coordination of views within the virtual workspace is tied to the content of a particular work domain. Cognitive work analysis, workload, timeline or other analyses sensitive to task transitions and incident evolution across different scenarios are needed to provide the information base for the design of a coordinated virtual workspace.
- Effective workspace coordination helps practitioners focus on the underlying process not the interface management tools per se.
- Deciding what to call up, where to look next in the virtual data space is a cognitive task. Part of effective workspace coordination is aiding this process.

Practitioners learn and adapt to cope with the burdens imposed by low or indirect correspondence at the workspace coordination level.

- For experienced practitioners responsible in some field of activity, disorientation or getting lost represents an unacceptable and abrupt breakdown in their ability to meet their goals as responsible agents in that field of activity. Regardless of how cumbersome the virtual workspace and the navigation mechanisms, real practitioners cannot afford to get lost. Instead, they learn to tailor the system and their strategies in ways to try to cope with the clumsiness of the system. A common tactic is to achieve simplicity of operation by ignoring system features--if there are too many possibilities, users drastically prune the ones they invoke.
- If the burdens of switching across displays or reconfiguring the windows are too great, the users try to develop their own fixed spatially dedicated tableau. They will put together a set of views that seem to balance observability of the underlying process or device with the need to avoid distracting interface management burdens that congregate at high criticality or high tempo periods.
- Ineffective designs at the level of workspace coordination force serial access to highly related data, add new interface management tasks that tend to congregate at high criticality and high tempo periods of the task, increase demands on user memory, undermine attentional control skills involved in knowing where to focus when, create the potential for getting lost or disoriented in the large space of options and displays.

Given these kinds of potential problems, what principles help establish direct correspondence at the level of a coordinated virtual workspace?

Visual Momentum

Visual momentum is the impetus or continuity across successive views that supports the comprehension of data following the transition to a new display in a human-computer interface as a practitioner carries out their domain tasks and activities (Woods, 1984).

The concept of Visual Momentum is borrowed from perception and cinematography (Hochberg, 1986) and refers to the impact of a transition from one view to another (a cut in cinematography) on the cognitive processes of the observer, in particular on the observer's ability to extract task-relevant information. As applied to HCI, "the amount of visual momentum supported by a display system is inversely proportional to the mental effort required to place a new display into the context of the total data base and the user's information needs. When visual momentum is high, there is an impetus or continuity across successive views that supports the rapid comprehension of data following the transition to a new display" (Woods, 1984, p. 231).

At one end of the dimension lies poor transitions which consist of (a) total replacement of one view for another and (b) the absence of any visible cues to the virtual field of possible views. When Visual Momentum is low, each "glance" into the artificial data field is independent of previous glances so that the observer must reorient from scratch to each new view as it is called into the limited viewport.

At the other end of the dimension the observer works within a conceptual space in which individual views are grounded. A conceptual space depicts relationships in a frame of reference (Woods, 1995a; 1996). In between lies a variety of techniques for building a sense of a conceptual space analogous to a physical space so that orienting and moving about the virtual perceptual field can employ the same perceptual and cognitive processes that allow us to fluently explore and reorient to new events and changing views in naturally occurring physical spaces. Some of the techniques to increase Visual Momentum are longshots, landmarks, content-laden cues to structure, spatial dedication, coordinating what can be seen in parallel and what in series as a function of task demands (Henderson and Card's "rooms"), side effect views, cues to status. All of these techniques and many others (e.g., trails, bookmarks, 3D spatial metaphors) function as visible cues to the structure of the space of possibilities and cues to the status of those different parts of the domain represented by the artificial data field behind the keyhole. Designers can orchestrate these kinds of techniques to create what the user experiences as a tangible conceptual space to support effective workspace coordination. Overall, techniques for visual momentum coordinate different kinds of detailed or focus views with broader context views. The model for this focus plus context principle is the human perceptual system which shifts focus fluently in pace with the changing tempo of situations based on the center-surround or focal attention plus orienting perceptual systems

Creating Center-Surround Coordination or parallel focus plus context views is the basic means to establish direct correspondence in workspace coordination.

Center-Surround: Parallel focus plus context views

In a natural perceptual field, people are able to switch attention and re-orient to interesting information in their environment (Rabbitt, 1984). The coordination of human focal attention and orienting perceptual functions such as peripheral vision supports the process of knowing where to look when. Orienting perceptual functions provide information about broad patterns in the surrounding environment and pick up changes that might warrant a shift of attention away from the current focus. Woods (1984) suggested that designers of virtual perceptual fields can model their systems on this characteristic of the human perceptual system--a center-surround technique. To support knowing where to look when, designers should surround a highly detailed central view coordinated with lower resolution, i.e., more **distilled**, views of physically or functionally **related** data. Lamping et al. (1995) have called this technique "focus plus context." Furnas (1986) proposed supporting navigation in a similar way by using an optical analogy, the fisheye lens. All of these labels have in common the technique of balancing a high resolution detailed view with summary views of related parts of the virtual perceptual field.

To develop center-surround workspaces requires definitions of what is distilled and what is related given the users expression or definition of a focus of interest.

Surround views are distillations

The point of a surround view is to support the potential to shift attentional focus fluently with change. To accomplish this, surround views need to distill the relevant information that characterizes the situation in a concise, recognizable form. The relevant information must be represented in a way so that observers can size up the state of affairs at a glance.

One indication of a concise distillation of relevant factors is that it is informative even if it is shrunk to a relatively small size. Surround views in which state is solely conveyed through the digital display of elemental data are rarely adequate. Similarly, designers can easily over-summarize and provide too little data to be of value to experienced practitioners. The classic example of the latter vulnerability in representation design is group alarms.

Surround views to provide distillations, they must meet the following criteria:

1. For Information to be distilled it must be abstracted.

To broaden the view of the system, information in the surround view should be abstracted to a higher level of information than that of raw data and other details about the system. It is important to note that abstracted information is not simply a lack of detail, but rather an integration of details that informs the observer about the relevant higher level questions (Vicente and Rasmussen, 1992). It involves collecting information from various areas of the system that speak to broader issues about contrasts between routine and exceptional conditions. To do this, the designer may need to transform lower level data, integrate these data and contrast them to related values. Examples of abstracted information include answers to questions like: what mode is the system

working in? is the system functioning normally, or is there a malfunction in one of the subsystems? what activities are taking place at this time?

2. The surround view must *include information about change and sequence.*

If practitioners are assessing the status of a system, they must be able to recognize change against a background. The value of orienting perceptual system is directly based on their sensitivity to pick up change or difference against a background. In addition to information about the current behaviors and states of the system, information about what has happened recently, and what kinds of trends may be developing often contributes to the overall assessment of status. For example, in event-driven worlds faults often result in a cascade of disturbances. A surround view will help the practitioners keep track of the big picture of disturbance evolution while they are pursuing specific portions of the cognitive work in diagnosis and response.

3. *The surround view must show information that is relevant to the viewer's current focus.*

The information in the longshot view should answer questions and provide information that makes sense to a practitioner in his or her task context. Examples of potentially relevant information include the activities that are currently ongoing and helping practitioners see if events are developing in accordance with their expectations. The longshot provides the larger context about the semantics of the field of practice in which one examines different and more focused views.

4. *The surround view should help the viewer shift focus if the current focus is inappropriate or if changes occur that suggest a new focus is relevant.*

The purpose of the surround view as a orienting perceptual function is to support re-focusing attention as situations change and given the potential to be focused in the “wrong” place.

5. *The surround view should support “check reading.”*

An abstracted and distilled overview should help users pick out what conditions or changes are potentially interesting quickly and in a mentally economical way given the current task context and ongoing lines of reasoning.

Surround Views as Longshots

It is commonly accepted that an overview display will support coordination or navigation across the many views available within the virtual data space. In Woods and Watts-Perotti (1997) summarize cases where designers provided what they thought were summary displays, but these displays did not support practitioners and were rejected or little used. What do designers need to do to create effective overviews?

Woods (1984) referred to summary displays that actually serve as effective orientation and navigation aids as longshot displays. In cinema, a longshot is an establishing view that shows relationships between characters and summarizes relevant information. It keeps the viewer involved in the flow of the plot by allowing him/her to step back from

the details, discover why these details are important, how they relate to previous views and to establish a frame of reference to help the observer comprehend upcoming views.

The successful longshot in computer based information systems serves the same purposes. In essence it is a kind of map that captures global relationships in the area synthesized. It helps practitioners step back from the details of the monitored process to assess overall system status. It helps them decide where to look next within the system. It helps them relate the view currently under examination to previous views and to integrate new views into their assessment of process state as they are called into a visible viewport.

Three functions contribute to the effectiveness of a longshot -- the status summary function (cues to status), the orienting function (cues to structure and mapping the structure to domain semantics), and the movement function (direct manipulation).

The Status Summary Function: Status at a Glance

Longshots containing status summary information allow users, in a mentally economical way, to step back and assess their overall situation with respect to the underlying process, device, or activity they are engaged in. By providing task relevant status information, longshots are content-laden. By showing the status of the process or activity behind the computer interface, longshots can help users decide where to look next in the virtual field. The concept of content-laden navigation aids which provide cues to status as well as to structure will echo throughout several of the techniques to enhance Visual Momentum across displays.

In a physically distributed control center (an open workspace), an experienced operator can stand at the back of the room and gain enough information to describe the status of the system, i.e., the current state, the current epoch or phase of operation, the direction things are headed in (e.g., deteriorating or recovering), the stance of actors towards the system (such as routine operations or a tense critical period). If a longshot includes the status summary function, it will externalize pertinent summary information found within the display structure, and thus allow an operator the opportunity to step back and assess the status of the monitored process and to quickly see how the system is behaving. The longshot provides cues to status across the entire system so that practitioners can remain in tune with changing conditions and the “big picture” while they are focused on a detailed task or part of the underlying process. Cues to status are an important technique to enable practitioners to easily “check read” or peripherally pick up what might be interesting changes in other parts of the underlying process that should guide a shift in attentional focus (Woods, 1995b).

Status summary information can also support tasks in domains that are more self-paced. Tasks in these domains often include updating and maintaining information within the display structure. For example, some spreadsheet users update lists of information or financial figures each month. A longshot could help spreadsheet users by providing status summary information about these updating activities (i.e., which

areas of the sheet have been updated, and which ones have not; Watts-Perotti and Woods, 1999).

The Orienting Function

The orienting function of a longshot helps operators orient to where they are (the currently visible views) relative to the set of views that they could examine in this context. The orienting function helps them comprehend cuts from one view to another. Longshots contain map-like characteristics which show users where they are located in relation to the important parts or landmarks within the virtual perceptual field. The map can serve as a representational framework for capturing what options are relevant to the current situation, which options have been recently selected/inspected (a trail), and support user browsing through potentially relevant views. The latter raises the question of how a longshot can be developed to invite exploration (Norman, 1988).

The structure of the overview display should reflect the structure of the views within the workspace as they represent the semantics of the field of practice. This type of conceptual map makes movement in the large network of possible displays the equivalent of moving in the semantics of the domain from the point of view of a practitioner. When a summary supports both this function and provides information about significant status and change in the underlying process or device, (the system status function), the longshot helps the operators formulate relevant questions and helps them decide where to direct their attention next. Supporting the orienting function helps produce interface transparency where the practitioners can concentrate on their activities and goals in their field of practice instead of being focused on the interface control mechanisms themselves.

The following principles should be followed to support the orienting function of a longshot:

1. The longshot must be coordinated with other views.

The longshot needs to be coordinated with the other more local kinds of views available to users and the reverse as well. A longshot serves as a visual representation of the virtual data space otherwise hidden behind the narrow keyhole. This means the longshot must capture the structure of the virtual data space (cues to structure). When it fulfills this purpose, the longshot can serve as a reminder of what types of local views are available and which of these views are relevant to the user's task context. The goal of coordinating the different views is to enhance the user's ability to extract and integrate information across cuts or transitions from one view to another (the fundamental definition of high visual momentum).

2. The overview display should include relevant frames of reference.

Coordination between longshot and other views depends on choosing, coordinating or creating frames of reference related to the practitioner's tasks in the field of practice. A frame of reference specifies relationships between parts and part-whole relationships. Depicting these relationships for a field of practice creates the longshot. The longshot defines and makes apparent the frames of reference within which other views exist and can be integrated. The other views have a meaning as a portion, neighborhood or

perspective of the larger frame of reference defined by the longshot. The portion currently visible becomes a place (where the user is “located”) within the larger space defined by the longshot. However, system designers sometimes create “overview” displays that do not show the operators where they are in relation to where they could be. Problems range from showing a list of menu options without indicating where the operator is located within the display structure, to showing a broad view of the structure of the monitored process without indication of how the process is broken up into displayable chunks (Woods, 1991).

3. The overview display must always be available in parallel with other views.

Experience from several projects (Woods, 1984; Bolt, 1984; Case 4) all point out that a longshot is effective only if it is available physically in parallel with other types of views. The point of a longshot is to help the practitioner know where to look next in the virtual field. If the longshot is not constantly available, then the practitioner has to decide when it is appropriate to consult it with the associated cognitive burdens. If other constraints make it impossible to constantly show the longshot, relax those constraints. An effective longshot is one fundamental kind of view which should be available in parallel in a well coordinated workspace.

4. A longshot combines information about the state and behavior of the monitored process with indication of the different types of available views relevant to the context.

When the representation of the types of available views is combined with the information about the monitored process, the overview display shows where operators could shift their attention within the display structure by showing what is happening in other areas. Therefore, a longshot integrates abstracted information about the status of the system while indicating the set of options or other views that are of potential relevance (Mitchell & Saisi, 1987). In effect, a longshot serves as a map of possibilities annotated with distilled data about the state of the process.

The Movement Function

In addition to showing overall system status and giving clues for where to look next within the system, the overview display should indicate how an operator can move to an area of interest within the display structure. The longshot functions as a map of other display possibilities in part. The map should make it apparent how one calls up or moves to or navigates to other views. Norman (1988) refers to this problem as the gulf of execution.

As a map of possibilities, a longshot provides a view of multiple options in parallel and shows the relationships between these options at multiple levels. This has led Billingsley (1982) and Woods (1984) to emphasize the need for maps or map-like structures to aid movement through a virtual space over a series of menu choices. Too often, systems of menus are designed in such a way as to force users through a long series of menu options which provide few options in parallel and which produce multiple across-display transitions without providing any content to the observer (see many Web sites). The longshot can support the movement function by allowing users

to select from it the views for display in other viewports. Kahn (1995, 1996) shows a variety of ways that maps can be used to support movement within a Web site.

But a longshot can be more than a structured menu with multiple options in parallel. Since the longshot is a map that organizes parts and relationships between parts for the domain of interest, it is one of the prerequisites for providing a direct manipulation interaction (Hutchins, Hollan and Norman, 1986). Designers can use this map to allow users to directly specify where they want to go or how to instruct the system without thinking in detail about the interface mechanisms. Many of the interface and navigation problems in Example 2, the infusion device for home care of pre-term labor, discussed at the opening of this chapter disappear once a meaningful frame of reference for user activities is found (Obradovich and Woods, 1996). In this particular case, user activities are all about different rates of infusion over different time intervals and different bolus sizes at different points in time. In other words, what is informative in this field of practice are dose-time relationships. Since none of this is visible in the infusion device, the designers are forced to create awkward, arbitrary, and cognitively demanding sequences of displays and interaction to program the device. Not surprisingly, the result is classic HCI deficiencies in the device producing typical user problems and errors (Norman, 1988).

The solution to these problems is to provide a dose-time structure as the basic frame of reference. Users will be unable to see all of the relevant dose-time picture in detail at one time because of multiple therapy plans and because the relevant information may stretch out into the future or back into the past. But a longshot of the dose-time frame of reference can support navigation across more detailed views, and it can support direct manipulation mechanisms for instructing the system about desired therapy plans. This radical restructuring of the basic concept behind the interface, based on the semantics of the field of practice, would completely restructure the nature of the interaction. Users would no longer have to learn or remember the syntax of the interface; instead, the interface would match the semantics of their activity--setting up, modifying and monitoring dose-time relationships.

Longshots contain two kinds of information: information about the underlying process or field of practice (e.g., state information) and information about the virtual data space (e.g., what views are available). It is important to perceptually distinguish these two different categories of information. Case 1, earlier in this chapter, is one example where the visual indications of where one can click to open new views/windows are incomplete and ambiguous. Similar confusions about which visual marks are "clickable" and which are not have been found in usability tests of Web sites (Nielsen, 1996).

Status summary information, combined with the multiple parallel options and structured relationships between parts, then can act like a preview function to help the observer to quickly find and focus in on what is interesting given their context and goals.

Longshots allow practitioners to step back from the details of a part of the underlying process or one subtask to see important trends and activities about the process as a

whole. This role means that longshots can be an important element in creating open workspaces that support cooperative activity if the longshot is available in common to the entire work team. Just because a display is available to all through a large screen display or other shared display mechanism is not sufficient for it to function as a longshot. The criteria in this section need to be met for the overview to contribute effectively to the shared workspace.

When an overview display provides the three functions discussed above--(1) status summary information, (2) visible cues to the semantic structure of the virtual field, and (3) allows practitioners direct mechanisms to shift their "gaze" within this space--then the overview functions as a longshot. The internet currently contains many sites which lack effective longshots and demonstrate the cognitive and performance consequences as users are forced through multiple steps blindly trying to get to information relevant for their purposes. Because effective navigation is central for users to be able to take advantage of the Web's power, we have seen a wave of innovation in a number of techniques, most notably techniques for mapping web sites. Kahn (1995, 1996) provides many examples, principles and a method to help designers create effective longshots and increase visual momentum across displays for Web sites and other human-computer interfaces.

Summarizing the above in a single principle of design for correspondence principle:

The navigation structure should be a model of the topic being navigated and summarize the status of the parts of the topic at a glance.

Surround Views as Neighborhood

In addition, to depicting global relationships, a surround can be designed as a view of the neighborhood around the current view of focus. This requires thinking about relationships across views to be able to define neighborhoods—what is "nearby" given different kinds of relationships in the processes or activities in question.

One question becomes what criteria should be used to define nearby related material. In many designs, developers use the physical topology of the underlying process to define nearby material. But representing functionally inter-related material in the surround is more powerful from the point of view of aiding practitioner extraction of meaning at this level of analysis (knowing where to look next). In this technique, the surround represents information that may be distant by one criterion (e.g., physical connectivity) but is close related to another criterion (e.g. diagnostic value). This technique is model-based as it presupposes a model of relationships between different types of views and information about the domain in question.

The selection of "nearby" material may be based on the focus view selected by the practitioner given a model of the domain or process which specifies what must be potentially relevant if one is focused on the material in the central view. For example, Woods and Hollnagel (1987) used one cognitive task analysis technique to map meaningful contexts for process control domains. As a result, they were able to

develop an information system concept where the computer system automatically displayed distilled views of functionally related topics or areas (e.g., other related goals, mechanisms and requirements). When a practitioner selected a topic or area of interest for display in the high resolution viewport, he or she had also defined a focus of attention. A functional model of the domain plus information on current status allowed the system to know what other topics are relevant to the primary view of interest. Based on this model, summaries about the status of these contextually relevant topics appeared in the surround automatically. These functionally related summaries in the surround could vary with context, i.e., the type and level of information displayed could depend on the state of the underlying system or on the state of the problem solving process.

Side Effect Views

A specific kind of neighborhood surround view is side effect views (Woods and Hollnagel, 1987). Side effect views provide users with information about distant areas of the information space that might be affected by actions they are taking or activities going on within their high-resolution central view. The goal is to give users a global picture of the state changes, both main effects as well as side effects, that occur as a result of actions taken.

Designers may want to consider side effect views whenever a domain has multiple interconnections between systems and functions so that actions can have multiple effects--the intended or main effects as well as other unintended "side" effects. In this kind of situation, it is relatively easy for practitioners to err by missing side effects of their plans and activities. Examples of domains with multiple interconnections are common in process control. For example, in nuclear power plants changing net water inflow to the reactor has multiple effects--coolant inventory changes, but reactivity (the nuclear reaction) can change as well because boron is dissolved into the coolant water to act as a moderator of the nuclear reaction. Increasing or decreasing net water inflow affects coolant level and boron concentration.

To reduce errors involving missed side effects, one nuclear control room computer-based display system included a specific side effects view (Woods and Hollnagel, 1987). If an operator is interested in one function (e.g., calls up displays on coolant inventory as the primary focus), the system also provides a summary of the status of other functions which could be influenced by changes in the primary area of interest in a side effects window (e.g., the reactivity/boron concentration function).

These kinds of complex interconnections happen in many other domains as well. For example, extended spreadsheets couple together multiple data tables. Values from one table may be used in calculations in other data tables. If a user of an extended spreadsheet changes the value of a cell on the screen, it may impact calculations and values in other data tables which are not visible because the extended spreadsheet is much larger than the available keyhole.

A side effect view can indicate whether a change in the table on a display causes numbers in distant areas of the extended spreadsheet to change in important ways

(Watts-Perotti and Woods, 1999). In this study, side effect views were proposed to explicitly represent the interconnections within and between multiple data tables in extended spreadsheets, and to provide immediate feedback about how changes in one spreadsheet cell affects other cells and data tables. Side effect views can help users notice errors quickly, and can help users find the cause of errors by showing which cells contribute to an erroneous cell. Note that for a side effect view to provide immediate feedback about user actions, the view must be available in parallel with the active area of the spreadsheet. Also, the feedback must provide information that is relevant to user tasks and goals. For example, if users have the goal of keeping specific cells in their spreadsheet positive (i.e. if users do not want the cell representing a project budget total to be less than zero), a side effect view can aid this goal by showing when this event is in danger of occurring (a formula result goes negative).

Overall, surround views should cue the observer to related views within the display space given the context and should provide status information as well. Just like the orienting function of peripheral vision, status information can help users decide when to change their focus of attention. The neighborhood surround also has a status summary function which requires (as for longshots) the data displayed should be distilled and abstracted, reveal activities and change, support shifting attention, and support check reading.

Designers have to decide what kind of information at what level of summary and abstraction is appropriate for a surround. What is nearby needs to be thought of as a semantic property of the conceptual space if one is to use the center-surround technique to help observers know where to look when (linking domain semantics to the structure of the virtual perceptual field). In many domains this is complicated by the fact that there are multiple semantic relationships between topics, areas and data.

There have been a large number of innovations which are based, in part, on a center-surround concept (Mackinlay, Robertson and Card, 1991; Mitta and Grunning, 1993; Robertson and Mackinlay, 1993; Sarkar and Brown, 1994; Rao and Card, 1994; Bartram et al., 1995; Mackinlay, Robertson and Deline, 1996; Greenberg, 1996). These innovations go by varied names depending on their individual history of development but they all depend on the basic concept of providing a higher resolution focused view surrounded by contextually relevant information and views.

Coordinating Perspectives

Knowing where to look next is also concerned with coordinating different perspectives on the processes and topics of interest. Different views can be a focus on different regions of a single perspective on the processes of interest. For example, the focus could be a portion or single region taken from a larger functional or physical topology. Different view can also refer to taking a different perspective on the topic or process of interest. For example, effective diagnosis in some areas requires contrasting or shifting between the physical topology representing sensors relative locations and the functional topology which captures how disturbances can propagate (e.g., Rasmussen, 1983).

Achieving direct correspondence involves providing cues that support shifts within or across perspectives. One technique for coordinating perspectives is landmarks.

Landmarks

Hochberg and Gellman (1977) define landmarks as "features that are visible at a distance that provide information about location and orientation." In the computer medium, landmarks are features in the interface that are visible at a glance and provide information about location and orientation from the current focal view to other possible views within or across perspectives.

Landmarks play an important role in navigation in physical space (e.g., Passini, 1992). For example, the Space Needle, which can be seen from a distance from most areas of downtown Seattle, can be used to keep track of where people are in relation to where they are going. A landmark provides information about both the relative distance they must travel to reach their destination, as well as the direction they must travel to get there. Once people become familiar with the area around the Space Needle, the structure can also serve as a reminder of the buildings and areas that surround it. In this case, the Space Needle serves as a summary of the surrounding area, as well as a frame of reference to help people establish their location relative to destinations. Landmarks support repair processes as the observer risks becoming disoriented. If people lose their sense of where they are located in the virtual data space, they can look for familiar landmarks to help them re-orient to their surroundings of possible views and perspectives.

Content free landmarks are structuring cues; content-laden landmarks provide perceptual cues to structure and signify something about the content of that area (e.g., a summary of ongoing status). Both types play a role in building a virtual perceptual space.

Watts-Perotti and Woods (1999) followed up field work on the navigation strategies adopted by extended spreadsheet users with a simulation study that examined the role of landmarks and longshots as well as other cues to structure in supporting navigation. In this study, experienced users of large extended spreadsheets were given data related in a large set of interconnected data tables and asked to carry out various tasks modeled on the tasks user perform in actual organizations. The investigators varied the cues to support navigation such as landmarks, longshots and others (e.g., spatial dedication). The results showed that landmarks are important cues for aiding navigation. Participants more directly targeted relevant data with landmarks than without. In fact, they considered landmarks so important that five of the seven participants when working with an extended spreadsheet without landmarks stopped performing the task and actually added landmarks to the extended spreadsheets before continuing with their tasks. Both content-laden and content-free were important to the participants in the study.

Control of Point of View or Contrasting Perspectives

Coordinating different perspectives on the processes and topics of interest can be based on different conceptual points of view each defined by different frames of reference for defining relationships. But with the explosion of technology for creating three-dimensional (3-D) virtual environments, coordinating multiple perspectives also applies to spatial environments and data.

Technology for creating 3-D virtual environments is not about a contrast between 2-D and 3-D visualizations (Roesler et al., 2001; Tittle et al., 2001). Rather the new technological capabilities represent the ability of practitioners to control point of view by which they observe a scene (and the activities that go on in that scene).

The history of technological capabilities have biased our models of how to represent scenes: we tend to call line of sight egocentric views 3-D and top down views on a scene as 2-D views. However, these are just two of many points of view for an observer to examine the **same** scene. No one point of view is ideal for picking up task relevant information about the scene (e.g., studies trying to decide if 3-D or 2-D views are better). Rather, there can be many points of view or transitions across points of view that provide the observer the critical information for that task and goal (Roesler et al., 2001).

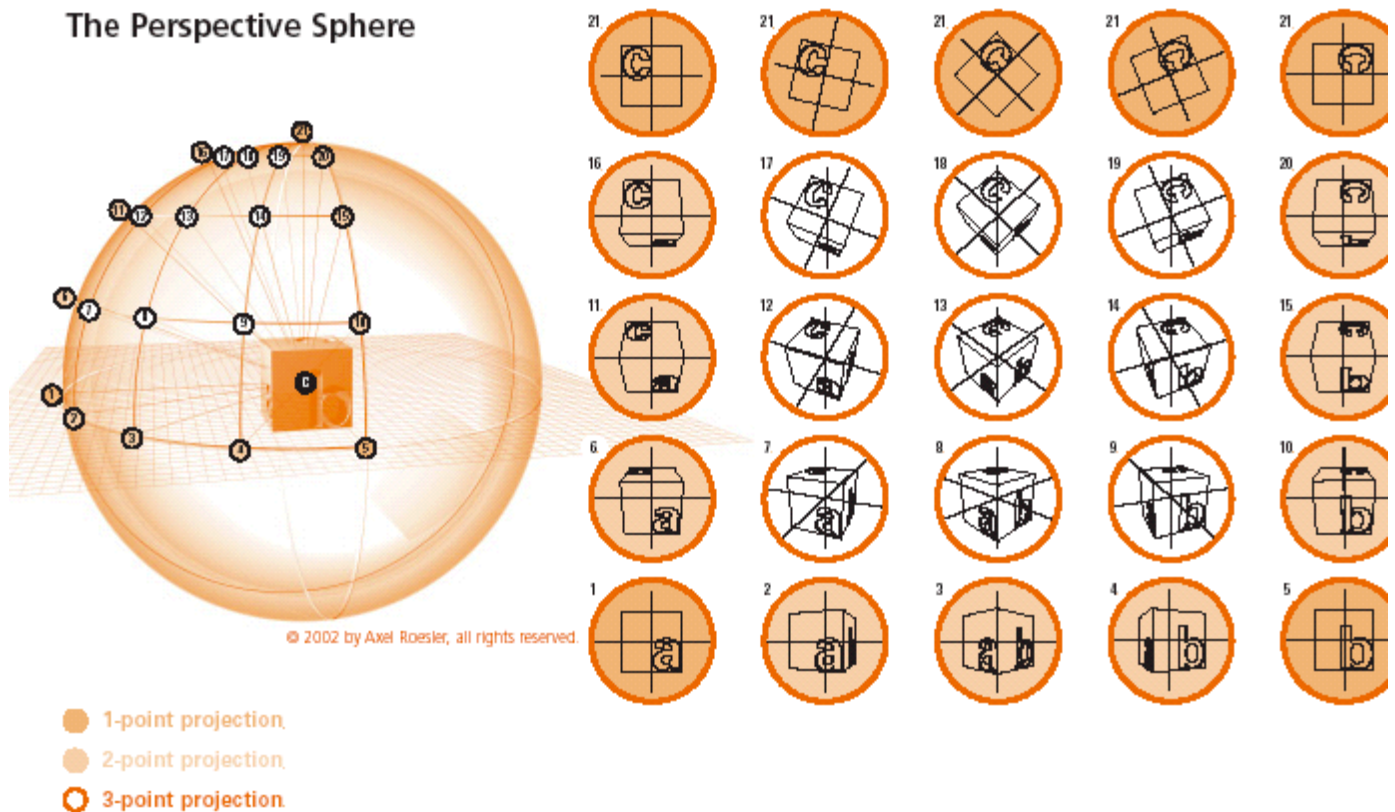
When designing views based on a 2-D perspective, there is a tendency to think of each view as a unique display or page. When working with 3-D virtual environments each unique camera position or viewpoint on a 3-D scene defines a unique perspective or view on that scene.

However, to fully interpret this image, the observer needs to be able to understand the relationship of the camera or viewpoint to the spatial environment. This is another example of making a process observable - in this case the imaging process - in order to make the view created from the imaging process more informative to tasks/goals in context.

To understand the implications of this observation, consider the Perspective Sphere as a model of the relationship between viewpoint and objects in a scene (Roesler et al., 2001). The Perspective Sphere illustrates perspective projections from camera position if the latter is moved in a 25-point grid on an eighth of a sphere's surface, looking at the center point of a cube (a generic stand in for the scene of interest) located in the center of this sphere. The diagram demonstrates that any single image from any single static virtual camera position presents a limited view of any real scene. While the image of the objects change with viewpoint, human perception generates a coherent percept of the object that is stable over changing viewpoint. To use the power of graphics technology, designers need to support the observers achieve the same stability as they change viewpoint or perspective within a virtual environment.

The Perspective Sphere diagram illustrates regularities in how changing virtual camera position changes an object's representation to an observer located at that camera position. The front, 90 degree side, and 90 degree top views show only one face of the

cube. On the other hand, every position on a direct latitude between top view and front view, or top view and side view, show 2 faces of the cube. The same holds for a camera moving on the equator between front view and 90 degree side view. Every other camera position in the 1/8 sphere's surface area between 90 degree top, front, and side view points reveals 3 faces of the cube. These conditions repeat eight times as the we consider camera position over the entire surface of the sphere.



View Tracks: 3D Virtual Displays are Viewpoint Dependent
 © 2002 by Roesler/Tittle/Woods, all rights reserved. C/S/E/L

The multiple views of the scene (the cube) shown in this diagram provide a simplified, yet comprehensive model of the relationship between observer's point of view on a spatial object in a scene. The relationship is defined by placing a camera everywhere on the surface of a circumscribing sphere (while the sphere is centered in the center point of the cube, and the camera is always looking at this center point). Six camera positions, or six exclusive points on the surface, show examples of one-point perspective where only one face of the cube is visible.

As the camera moves over the surface of the sphere, other patterns occur: At all points on the equator of the sphere and on all points of the four 90 degree latitudes

connecting the one-point camera positions, the view position will result in 2-point projections of the cube, showing two faces of the cube. These view positions are located as orthogonal lines on the sphere's surface.

At all of the remaining points of the surface, the view position will result in three-point projections of the cube, showing 3 faces of the cube simultaneously. These view positions cover the entire surface of the sphere, except the lines along which the 2-point projection views are observed, and at the intersections of these lines (where 1-point views are obtained).

From these observations it can be concluded that all representation schemes for displaying 3-D environments on a 2-dimensional picture plane are constrained and defined by the position of a view point relative to the scene. Working in the other direction, the representation of the scene tells us where the camera is located (view point, view track or region) that generates such views of the object.


The Perspective Sphere model warns designers about possible misinterpretations of virtual 3D scenes (scale ambiguities, optical illusions, distortions of relationships between objects in the scene, mis-estimation of distances and spatial relationships) that come along with the new powers for generating cues to depth in computer visualizations (see Tittle et al., 2001).

There are many lessons one may draw from the example of the Perspective Sphere for developing more direct correspondence:

1. Any single image from any single static virtual camera position presents a limited view of any 3-D scene.
2. Relationship of observer to the scene is a critical variable in defining what is informative in the scene.
3. Provide control of point of view to observers (e.g., provide view tracks as the basic unit of interaction).

The Perspective Sphere model also points to some design concepts that avoid these difficulties and enhance the usefulness of 3-D representations for supporting human performance. For example, one such technique is the center surround central to establishing visual momentum—show the view of the scene and the virtual camera position that provides that view in parallel to support better comprehension of the scene. The figure below illustrates the technique as the left hand perspective represents the view point or camera position used to generate the scene image on the right. Making these two views available in parallel and providing camera movement along this track makes it easy comprehend the scene and to find camera positions that enhance certain judgments (e.g., relative length).

The value of using the center-surround technique in this way is that it gives users guidance about the imaging model that generated the specific view while providing a sense of the possible views that one could take on the scene. Notice that the image on the left does not just show camera location, but it also shows the area of the scene falling within the camera point of view. Such cues to the imaging process can greatly enhance the ability of observers to interpret any specific sequence of views of a 3-D environment.



This technique needs to be supplemented by observer control of the virtual camera position. Having the left hand view of the relationship of the camera to the environment visible, makes the imaging process more observable, but it is only directable when the user can control camera movement and orientation relative to the environment.

Another technique that such a viewing framework can provide is the creation of landmark camera positions along the track. In this case, the two extreme positions along the track (top down and 'line of sight' camera positions) define landmark camera positions. Users should be able to quickly and directly move the camera to these landmark positions. The interface should also allow users to set or modify landmark positions as they discover special points of view that support the demands of the tasks they are carrying out with the aid of that virtual environment.

The above discussion illustrates many of the principles for direct correspondence at the level of workspace coordination. Using 3-D virtual environments to represent activities in a physical space of interest raises issues about coordinating multiple perspectives on the scene. To meet this criterion principles such as center-surround and landmarks are important means to achieve support requirements of observability and directability.