

Structuring Coordination to Provide Intelligent Assistance in Groupware

Joshua Introne and Richard Alterman
Volen Center for Complex Systems
Brandeis University
[jintrone, alterman@cs.brandeis.edu]

Abstract

In groupware, users must communicate about their intentions and maintain common knowledge via communication channels that are explicitly designed into the system. Depending upon the task, generic communication tools like chat or a shared whiteboard may not be sufficient to support effective coordination. We have previously reported on a methodology that helps the designer develop task specific communication tools, called coordinating representations, for groupware systems. Coordinating representations lend structure and persistence to coordinating information. We have shown that coordinating representations are readily adopted by a user population, reduce coordination errors, and improve performance in a domain task.

As we show in this article, coordinating representations present a unique opportunity to acquire user information in collaborative, user-adapted systems. Because coordinating representations support the exchange of coordinating information, they offer a window onto task and coordination-specific knowledge that is shared by users. Because they add structure to communication, the information that passes through them can be easily exploited by adaptive technology. This approach provides a simple technique for acquiring user knowledge in collaborative, user-adapted systems.

We document our application of this approach to an existing groupware system. Several empirical results are provided. First, we show how information that is made available by a coordinating representation can be used to infer user intentions. We also show how this information can be used to mine free text chat for intent information, and show that this information further enhances intent inference. Empirical data shows that an automatic plan generation component, which is driven by information from a coordinating representation, reduces coordination errors and cognitive effort for its users. Finally, our methodology is summarized, and we present a framework for comparing our approach to other strategies for user knowledge acquisition in adaptive systems.

Keywords: *Groupware, Knowledge Acquisition, Adaptive User Interfaces, Coordinating Representations, Plan Recognition*

1 Introduction

Acquiring enough run-time information about the user to provide intelligent support is a central issue in building any adaptive system (Fischer, 2001; Jameson, 2002; St. Amant & Young, 2001). To tutor a student, a system needs to know the student's strengths and weaknesses; to offer context sensitive help, it needs to know what the user wants to accomplish; to help find relevant content, it needs to know the user's interests. For a collaborative system to adapt to its users' needs, it must have information that is relevant to collaboration – information like the users' shared goals, common knowledge, or roles.

A major hurdle in the design of a user knowledge acquisition strategy is how to encourage the user to furnish the information that the system needs without impairing the overall usability of the system. There are numerous approaches to the problem. Mixed initiative systems (e.g. Maes, 1994; Horvitz, 1999) often obtain information through agents that engage the user to jointly solve a domain task (e.g. TRAINS, Ferguson, et al., 1996; LUMIERE, Horvitz, et. al., 1998; COLLAGEN, Rich & Sidner, 1998). Intelligent tutorial systems exploit the information gained during the testing process to derive models of students (e.g., Anderson, et. al., 1995). Multi-modal systems (e.g., Bosma & Andre, 2004) and affective interfaces (e.g., Kaiser, E., 2005) integrate other sources of information, like gestures and biometrics, to enhance the information available to a back-end inference process.

In groupware systems, much of the information required to drive intelligent support is already available in the communication that goes on between coordinating users. Unfortunately, most user communication occurs via unstructured channels such as chat or message boards, making it difficult for the system to extract useful information. One way to address this is to introduce structured communication channels (e.g. Malone, et. al, 1996; Soller, 2004). When introducing structure, however, care must be taken so as not to impair the users' ability to communicate with one another.

This article describes a method for introducing structure that people will use, that helps them stay coordinated, and at the same time provides useful information to the adaptive system. The method hinges upon the development of task specific communication tools that support the exchange of structured information relevant to coordination. Because these artifacts add structure to communication, the information that passes through them can be easily exploited by runtime algorithms. Because the information is relevant to coordination in the task domain, it will be useful for domain-specific intent inference. Here, we present a series of case studies demonstrating the approach in a testbed groupware environment.

1.1 Coordinating Representations

There is a rich body of ethnographic research that describes how people design and use structured artifacts to support specific and complex coordination as part of an established work practice (Suchman & Trigg, 1991; Goodwin & Goodwin, 1996; Hutchins, 1995; Schmidt & Simone, 1996). These kinds of artifacts have several common properties. They are external representations that address recurring problems of coordination within a community. They make it easier for individuals to align their private views of the world. There is a protocol that describes their typical use within a particular community of practice. Following Suchman & Trigg (1991), we refer to these kinds of artifacts as *coordinating representations* (CRs) (Alterman, et al., 2001).

There are many examples of CRs in everyday life. The stop sign is a CR that helps people to coordinate their behaviors at an intersection. The arrivals and departures board at a train station supports coordination between many people and the operators of a railway terminal. The grocery list on the wall in the

is used to improve the system from the users' perspective, and from the perspective of the adaptive component designer.

To serve as a roadmap, an outline of the argument is provided here as a series of questions and answers.

1. **Question:** How can adaptive support be added to a groupware system?
Answer: We infer users' intentions.
2. **Question:** Where can the run-time user information required for intent inference be obtained?
Answer: This information is available in unstructured communication channels, but because it is unstructured it is difficult to access.
3. **Question:** How can this information be transformed so as to make it available to the system?
Answer: By introducing a structured communication channel in the form of a coordinating representation.
4. **Question:** If users are provided with this more highly structured communication channel, will they use it?
Answer: Yes, because the structure that is provided by the coordinating representation makes it easier for users to exchange some kinds of information.
5. **Question:** Can a coordinating representation be developed that does not impair user performance?
Answer: Yes; in fact, CRs improve domain performance.
6. **Question:** Is the information that becomes available via the use of the CR useful for intent inference?
Answer: Yes, because it offers a window onto task and coordination-specific knowledge that is shared by users, and this knowledge is highly structured.

The results from case studies will be presented that confirm the answers to the above questions, demonstrating that:

1. Users choose to use the coordinating representation over chat to exchange some kinds of information (*Section 4.1*)
2. Using the coordinating representation helps users perform their task better. (*Section 4.1*)
3. The information provided by the coordinating representation enables intent inference with good accuracy. (*Section 4.2*)
4. An adaptive component that is driven by the information made available is heavily used, improves performance in the domain task, and reduces cognitive effort. (*Section 5*)

The article proceeds as follows. First, the domain, groupware platform, and associated intent inference technique are presented. Some of the representational deficiencies of the groupware system are then described; first from the perspective of the user, and then from the perspective of the developer of an adaptive component. It is then shown how a CR addresses these problems. Empirical data is presented that demonstrates how the CR benefits the users, and how information derived from its use improves intent inference. An adaptive component that uses information from the CR is described, and further empirical data confirming the utility of the adaptive component is presented. After a summary of the overall methodology, we provide a framework for comparing this methodology to other approaches to building adaptive systems. The article concludes with some final thoughts about the generality of the approach and future work.

2 VesselWorld

The experimental platform used in these studies is a groupware system called VesselWorld, in which three participants collaborate to remove toxic waste from a harbor. VesselWorld was demonstrated at CSCW 2000 (Landsman, et al., 2001). The platform was designed with the goal of studying the coordination problems commonly faced by groupware users. The domain task requires varying degrees and types of coordination, collaborators have different roles and responsibilities, and coordinating information is exchanged via a chat window. These features are highly relevant to the study of groupware systems in general. VesselWorld proved to be very challenging for its users; in our case studies, we found that the performance of a given user group usually did not stabilize until after roughly seven hours of use (including a two-hour training session).

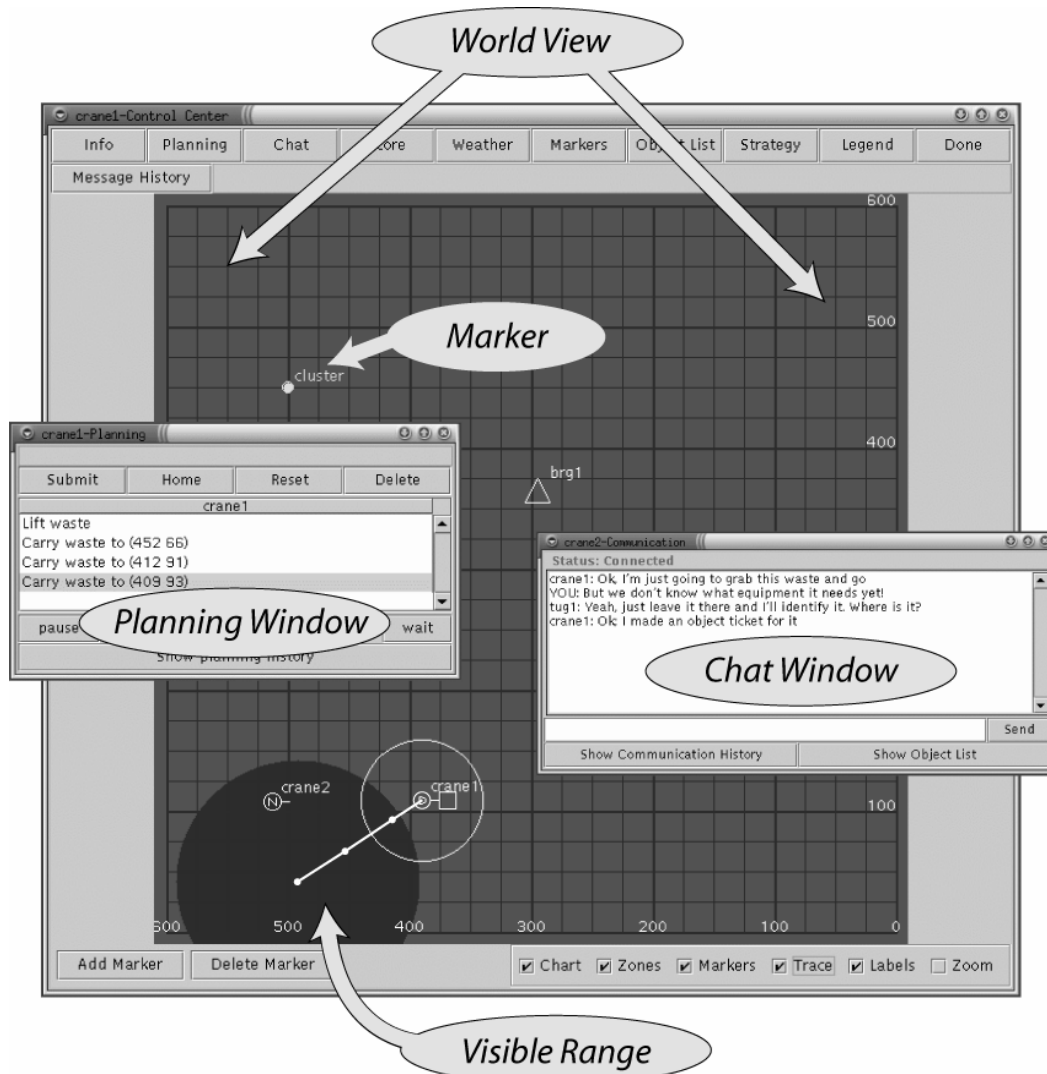


Figure 2: The VesselWorld system

VesselWorld presents a relaxed WYSIWIS environment, shown in Figure 2, in which three participants play the role of ship's captains, and their joint goal is to remove toxic waste barrels from a harbor without spillage. The main interface ("World View" in the figure) is a shared map; the x (increasing to the west) and y-axes (increasing to the north) indicate latitude and longitude respectively. Users are pro-

vided a private “marker” facility so that they may annotate the world view with pertinent information (see the “Marker” in the figure). Each ship can only “see” a small radius around its current location (the darker circle marked “Visible Range” in the figure), so each user has different directly observable domain information; hence users must communicate to maintain shared information about wastes in the environment. Communication may occur at any point, but all communication occurs through a text-based chat window that is part of the system (“Chat Window” in the figure). VesselWorld logs complete interaction data that can be used to replay user activity. This later feature is discussed further in Section 6.

The progression of a VesselWorld session is turn-based, such that every user must submit a single step to the server before it evaluates them and updates the world on each client screen. Users may plan any number of steps in advance, although each step can only involve objects that are currently visible. Plans can be managed (steps may be deleted or plans reset) via a separate planning window (“Planning Window” in the figure). Users’ plans are not visible to each other, again requiring explicit communication to manage coordinated plans.

A VesselWorld session is complete when all toxic waste barrels have been moved to a large barge, which has a fixed position and unlimited capacity. Each ship has a different role in this process. Two of the ships have cranes that can be used to lift toxic waste barrels from the harbor and load them onto a barge. The third user is a tugboat that can be used to drag small barges (which have limited capacity) from one place to another. The crane operators can load multiple wastes onto the small barge, and at least one of them must also be present to unload the barrels and place them on the large barge. For notational convenience, we will occasionally refer to the crane operators as “cranes,” the tugboat operator as the “tug,” and toxic waste barrels as toxic wastes or simply wastes.

Wastes are of different types and require different coordination strategies to be removed from the harbor. A single crane may lift a small or medium waste, but two cranes must join together to lift and carry a large waste, and an extra large waste may be jointly lifted but can only be carried on a small barge by the tug. Wastes may require specialized equipment to be moved, and the cranes carry different types of equipment. The tug is the only actor who can determine the type of equipment a waste requires.

The users are scored by a function that takes into account the number of barrels cleared, the number of steps this took, the number of errors (dropped waste barrels) made, and the difficulty of the problem. In all user studies, the users were instructed to try to maximize their score.

2.1 Intent Inference in VesselWorld

Planning in VesselWorld is a laborious and error prone operation (Alterman, et al., 2001). User errors often occur because of forgotten plan steps or joint plans that have become unsynchronized. We sought to develop an automatic plan generation tool to address these problems. A hurdle in making such a tool useful is that there are an overwhelming number of potential goals for each user at any given time. Thus, an intent inference procedure was developed to reduce the number of possible goals to a manageable list from which users could then make a selection.

Bayesian Networks (BNs; Pearl, 1988) were used to infer user intentions. BNs have been used in numerous systems and in a variety of ways to perform plan recognition (e.g. Charniak & Goldman, 1993; Horvitz, et al., 1998; Albrecht, et al. 1997). The approach taken here is straightforward. The BNs used in VesselWorld are static and model the relationship between information about the state of the domain and users’ likely intentions. At runtime, information about the domain is posted to these models for each agent-waste pair separately. The likelihood an agent has an intention with respect to a given waste is

read off one of the nodes in the network, and the intention with the highest value for a given agent-waste pair is taken to be the most likely one for that agent. New information is posted to the BN whenever a relevant change in the world is detected.

Two BNs were developed, one which infers crane intentions and one which infers tug intentions. Together, the two BNs can infer seven goal types; JOINT LIFT (<crane><waste>), LIFT (<waste>), JOINT LOAD (<crane><waste>), and LOAD (<waste>) for the cranes, and BRING (<small barge><waste>), ID (<waste>), and TRANSFER (<small barge ><large barge>) for the tug. In this paper, we restrict our analysis to the portion of the crane network that predicts the Cranes’ intentions to lift wastes. This BN is shown in Figure 3; it models the likelihood that a crane operator has the intention to lift (or jointly lift with the other crane operator) a specific toxic waste.

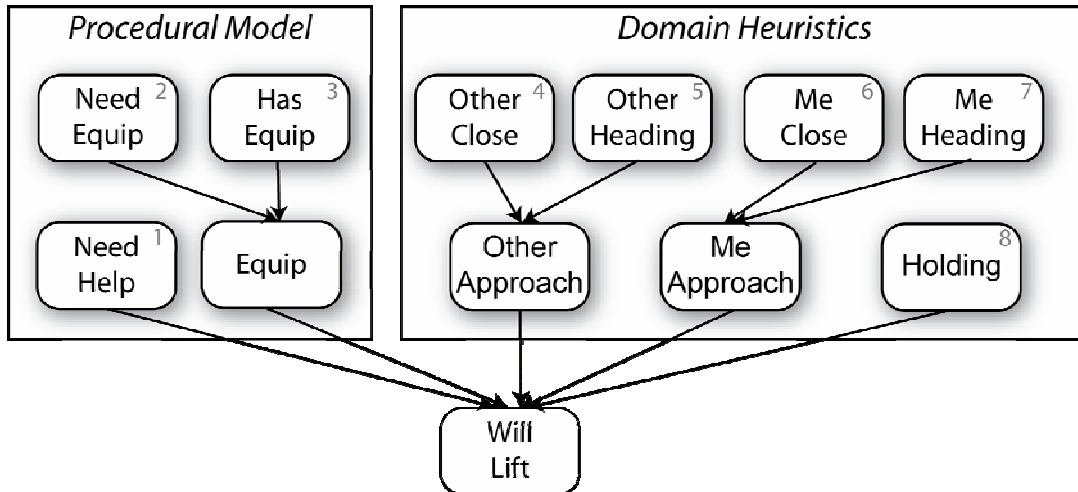


Figure 3: Schematic of BN used to infer Crane Lift Intentions (information is posted to numbered nodes)

The information used to determine if a crane has an intention to lift a waste is:

- The size of the waste (which determines whether a single crane can lift the waste, or the support of another crane is required). (Node 1)
- The type of equipment required by the waste and the type of equipment the crane has. (Nodes 2 & 3).
- Whether the cranes are close to or heading towards the waste. (Nodes 4-7).
- If the crane actor is currently holding a waste. (Node 8).

As portrayed in Figure 3, the BN combines a procedural model of the task-domain with domain-specific heuristics that indicate which of those actions are most likely. For example, a procedural constraint is that an individual crane cannot lift a waste unless it either has the right type of equipment, or the waste does not require equipment; this information is captured in three nodes, “Need Equip,” “Has Equip,” and “Equip.” These three nodes do not represent a simple rule, because the equipment requirements for a waste might not be known, and hence it is necessary to explicitly handle this kind of uncertainty.

Heuristic factors interact with the procedural model to differentially weight those wastes that are possible lift candidates. For the cranes, these factors are how close the ship is to a waste, how close the current heading of the ship is to a path that will intersect with a waste, and whether or not the operator is already holding an object. This heuristic information in part reflects the physical constraints of the domain; a ship must be close to a waste before it can be lifted.

As a whole, the BN models users that have a certain degree of proficiency with the domain. It is expected that such users understand the rules of the game, and will only try to lift wastes that they can. It is also expected that these users will proceed in a generally predictable fashion. That is, they will not zig-zag wildly upon an approach to a domain object, and they will not randomly pick up wastes and put them down again. It would be possible to incorporate additional models for less proficient users, or model users' domain concepts more precisely (e.g., Horvitz, 1998). The focus of this work, however, is upon solving the difficulties in acquiring the user information necessary to drive the user model at runtime.

3 Representational problems in VesselWorld

Any piece of software may be considered to be part of a representational system. A representational system has three essential ingredients (Norman, 1991):

1. The represented world (that which is to be represented);
2. The representing world (a set of symbols);
3. An interpreter (which includes procedures for operating upon the representation).

The representing world mediates the interaction between the interpreter and the represented world, and in so doing constrains the types of interaction that are possible between the two. As applied to software, the software system is a representational medium that constrains both how information in the represented world is viewed, and what procedures are available to the user for interacting with this world.

For adaptive systems, software can be seen to function as a mediating layer in two distinct but intersecting representational systems. First, and most clearly, it mediates the interaction between the user and domain (which may include other users), transforming the represented world so as to (hopefully) better support domain activity. The software system also mediates the interaction between the designer of the adaptive component and the user. In this later case, software very much constrains what information the designer can acquire about the user at runtime.

Here, representational deficiencies in VesselWorld are considered from both perspectives. The problem from the user's perspective, which is most traditionally thought of as an HCI problem, is considered first. The problem from the perspective of the adaptive component designer, which is a user knowledge acquisition problem, is then described. Both these kinds of problems are exactly those which Coordinating Representations can solve. Following the analysis of these problems, it is shown how the introduction of a CR in VesselWorld does just that.

3.1 Problems from the user's perspective

During a VesselWorld problem solving session, users must search for and discover information about the waste barrels that are floating in the harbor. Because each user has different locally available information, and recovering wastes requires the coordinated efforts of multiple users, it is necessary that participants communicate to establish mutual knowledge about the wastes. Managing and maintaining this shared information comprises a significant portion of the work that users must do during a VesselWorld session.

VesselWorld initially provided two tools to support this work. The chat window allows users to communicate information about objects in the domain. It provides a scrollable history that records the entire dialogue during a given chat session, so that users can more easily recover information that may have been missed or is forgotten. The other tool is a "marker" facility, which allows individual users to anno-

tate their own private maps (each user can only see their own markers). A marker consists of a highlighted point on the map, and a free-form text message for recording information (see Figure 2).

These tools reflect the designers' initial vision of users' representational needs for managing domain information and establishing mutual knowledge about waste barrels. It was expected that users would publish local information via the chat window, and use the marker facility to record all waste information. During actual system use, it was found that these tools did not provide users with sufficient support. These representational deficiencies were most clearly manifest in specific features of the users' runtime dialogue.

One such feature was explicit talk about coordination, as shown in Figure 4. In the example Crane2 suggests that the participants work out a shorthand for referring to waste information (line 3). Such explicit talk reveals a perceived need by the users for some form of coordinating structure.

1. Tug1: There are two waste sites near me, one at 120/415, one just SE of it.
2. Crane1: what size/ equi[p?
3. Crane2: hmm shall we come up with a shorthand for a waste info?
4. Tug1: The Se one is small, needs a Dredge, the first is small, and need NO equipment. Does that mean I can pick it up?

Figure 4: Explicit talk about coordination

Another feature of user chat was the emergence of explicit conversational patterns for managing domain information. For example, one group would perform a "marker check" when there were discrepancies among individual's local waste information, in which all of the known wastes would be validated. First, an actor would announce a marker check, and then proceed to transcribe all of her local waste information into the chat window. The other actors would confirm or contradict this information according to their own private views of the world. The transcription process that this group of actors undertook involved a significant amount of mundane work. However, users made the decision that this was less work than trying to recover from errors that were caused by inconsistent information about the world.

Recurring coordination problems were also indicative of a representational deficiency in the system. Frequently, one actor's reported perceptions raised questions from another actor, possibly about the size or location of a waste. This was usually caused by a lack of precision in communicating about locations, leading to difficulties in resolving references to objects. Sometimes, these errors would lead to one actor moving to the area of the waste in dispute to resolve the problem.

The coordination problems and dialogue features described above can be generally described as grounding problems (Clark, 1995). The creation and maintenance of common ground forms a substantial portion of the work that goes into conversation (Fergusson, et al., 1996; Traum 1994). Research has also demonstrated that technological media have impacts on grounding (Clark & Brennan, 1990; Brennan, 1998; see also Clark & Wilkes-Gibbs, 1990).

To a degree, our analysis methodology (see Alterman, et al., 2001 and Feinman & Alterman, 2003 for more detail) builds upon some of this earlier work on the grounding process. Our analysis elucidates those coordination problems that lead to the design of domain specific CRs, and grounding problems are one type of coordination problem that CRs can address. Critically, CRs solve two problems in the design of adaptive groupware. They address coordination problems for collaborating users, and they can also

address a knowledge acquisition problem. In the following, the knowledge acquisition problems in VesselWorld are examined more carefully.

3.2 Problems from the adaptive component designer’s perspective.

The BN introduced in Section 2 is a model of the behavior the designer expects of users in different situations at runtime. For instance, the designer expects that if the two cranes are approaching a large waste and are not carrying anything, there is a good chance they will attempt to jointly lift the waste. In order for this model to function, information about the user and the user’s runtime context is required. Some of this information can be derived from the plan steps users submit (where the users are, what direction they are heading in, what they are holding). However, a significant portion of the required information about wastes (where they are, what size they are, what equipment they require) is only available in the users’ discussion about the domain. This information is very hard to extract automatically.

Figure 5 is a portion of the chat logs taken from a typical planning session in VesselWorld. Note that users sometimes refer to toxic wastes by coordinates on the shared map (e.g. “105, 420”). In the first line of the example, Crane2 announces a waste at (120, 420). In lines 2-4, Crane1 asks for clarification about the specifics of the waste. In lines 5-6, the tug replies (having apparently already investigated that toxic waste barrel) with corrected coordinates (105, 420) and specific information about the waste. In line 8, Crane2 thanks the Tug operator for the clarification, and the Tug closes the conversational turn in line 9.

```
1. Crane2: I found a waste at 120
           420
2. Crane1: ok
3. Crane1: what type of waste?
4. Crane1: large, small?
5. Tug1:   105 420 needs a dredge,
           i think that is where
           you are
6. Tug1:   small
7. Crane1: ok
8. Crane2: Thanks for checking
9. Tug1:   no problem
```

Figure 5: Excerpt from chat during VesselWorld session

This dialogue illustrates some of the problems in automatically extracting the domain information required as input to our intent inference procedure. In order to associate information appearing in separate utterances with a single concrete waste, it is necessary to correctly resolve references. However, because the dialogue occurs between three active participants, the conversational turns that might be used to narrow the reference resolution scope are hard to identify. Furthermore, referring expressions can change from utterance to utterance even within the same conversational turn. For example, line 1 refers to the waste as “120 420” and line 5 refers to the same waste as “105 420.” People can sometimes handle such ambiguities, but this is problematic for automatic reference resolution algorithms.

Within user groups, some structural conventions for referring to domain objects do emerge, but the type and use of this structure varies widely between and within groups. Some groups are explicit in defining a shorthand; other groups converge to common conventions over time. Within groups, referential styles vary between participants and over time. Oftentimes, multiple conventions will co-exist. In Figure 6, a section of a transcript is shown for a group that exhibits this behavior. In the segment of dialogue shown,

the players are announcing waste barrels they have found. In this group Crane2 always includes an “@” between the type and the location portions of the referring expression, and a question mark is used to indicate that equipment requirements for the waste are unknown. Crane1 usually places the type description before the location, and the tug places a type description after the location. The participants never converge to a common convention

- | |
|--------------------------|
| 1. Crane1: l 543 204 |
| 2. Crane1: s 562 150 |
| 3. Crane2: X?@150,559 |
| 4. Tug1: 1 395 lg dredge |
| 5. Crane2: s?@190,434 |
| 6. Crane2: s?@202,336 |
| 7. Crane1: sm 394 71 |
| 8. Crane1: large 395 22 |

Figure 6: Different ways of referring to wastes, within one group

For purposes of knowledge acquisition, the problem with such inconsistencies is that it is very difficult to produce a single rule or set of rules to identify referring expressions within chat. Rather than developing specialized algorithms to deal with the nuances of three-way, live chat in the VesselWorld domain, it would vastly simplify our task if users were to enter all the information the system needs in a common, structured form. Although this might seem like it would unnecessarily burden the user, below we will see that this is not the case.

4 Using CRs to fix representational problems

VesselWorld has representational deficiencies from the both the user’s and the adaptive component designer’s perspectives. In mediating the interaction between the users and domain, it does not offer the right kind of support for managing information about wastes. In mediating the interaction between the designer and the users, it does not provide structured access to the runtime information required for successful intent inference. The analyses performed above led to the design of a Coordinating Representation, which is a solution to both of these problems.

There are other approaches to designing of collaborative artifacts that address various issues in collaboration. Cognitive work analysis (Vincente, 1999) advocates a host of ethnographic methods for developing interfaces that address collaborating users’ needs. Activity theorists focus on the need for mediating artifacts that address tensions in the larger activity system (e.g. Engeström, 2000). Suthers (2003) approaches the design problem from the perspective of the teacher, as a way to improve the process of collaborative inquiry. However, no existing framework draws together design problems from the perspectives of both the user and the adaptive component designer.

Several CRs were developed for VesselWorld (Alterman, et al., 2001). The CR that addresses the representational problems described above is called the Object List (Figure 7). The Object List is a WYSIWIS (What You See Is What I See) component that helps users to manage domain information and coordinate references. It displays the same data in tabular format for each user. Users enter and maintain all of the data in the Object List. Each row of data contains several fields of information, including a user assigned name, the status, and the location of the associated object. The location field may be filled in by clicking first on the field and then on the object on the map. The size, equipment, action, and leak fields are filled in using drop-down menus. A free text field (“Notes”) is also provided for each entry so that any other relevant information may be communicated. Entries in the Object List

can be displayed on the World View (Figure 2) as icons that are annotated with the name that is in the “Name” field.

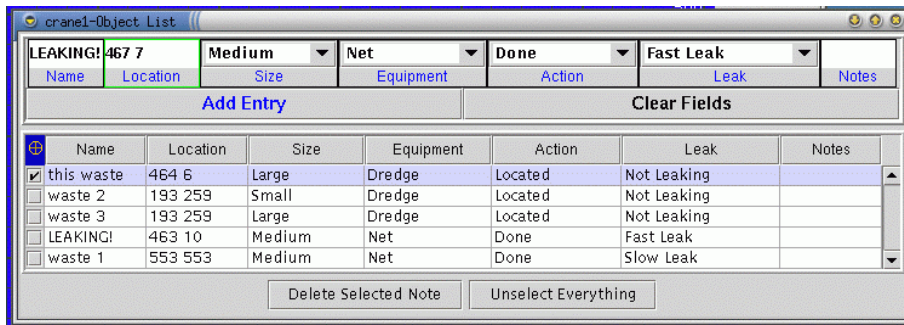


Figure 7: The Object List

There is nothing exotic in the design of the Object List. It is remarkable as the product of a repeatable design methodology that is explicit guidance for the creation of shared artifacts which simplify coordination for users and make intent inference easier. In the following, empirical data is provided that demonstrates the Object List indeed achieves both of these goals.

4.1 Improvements from the user’s perspective

In all user experiments, the Object List was heavily used, and the coordination problems described above were no longer observed. The Object List had significant qualitative impacts upon the way users coordinated information about toxic waste barrels. Figure 8 compares sample dialogues from users who did not have access to the Object List (the left column) to those that did (the right column). It is immediately apparent that the group that had the Object List spent far less time talking about the details of each waste. None of their discussion appears to be designed to exchange information about the wastes per se, but rather is focused on strategizing.

Without the Object List	With the Object List
1. Crane2: what was small at 275, 400?	1. Crane1: I got an XL!
2. Tug1: sX	2. Tug1: I got nothing, you luck basrstartd.
3. Crane2:ahve sX at 450 above that	3. Crane2: I got an Xl and an L, mommy! ;)
4. Tug1: mD at 400 350	4. Tug1: Merry christmas kids...
5. Tug1: yes, there is an sX at 275 250 as well	5. Crane1: I'll map North third?
6. Crane2:I have two amall at 275, 400 and 275, 450 are these the same?	6. Tug1: I'll take middle 3rd.
7. Tug1: no, there are two sX there	7. Crane2: I'm at south-central. Tug, where are you?
8. Tug1: well, there are actually three in almost a line from n-s	8. Tug1: I'm jus nw of the barge, let me put that on the map...
9. Crane2:at 400 and 450? what about the 275, 250?	9. Tug1: actually more w than n.
10.Crane2:ok, so the southern exists?	10.Crane2: With the LB in the corner, maybe our best bet is moving the SB NW and loading it with all the NW corner's goodies, which CRANE1: can map
11.Crane2:I'm almost there, nothing at 275, 250	11.Crane1: not a bad plan...
12.Tug1: 300 350, 250 400, and 275 250 are my markers	12.Tug1: Ok, I'll make a bit of a sweep around here while CRANE1: looks around.
13.Tug1: argh	13.Crane1: Tug, can you pick up the SB at your earlier opp?
14.Tug1: I mean 275 450	14.Tug1: CRANE2: can map up on the way?
15.Crane2:ok, those sound good	
16.Tug1: i don't know why I kee doing that.	

Figure 8: Dialogue from users before an after the introduction of the object list

Figure 9 depicts the information users entered into the Object List during the dialog in the right hand column of Figure 8. This demonstrates that the information previously exchanged via the chat tool has now been offloaded into the Object List.

⊕	Name	Location	Size	Equip	Action	Leak	Notes
⊗	G1	556 465	XLarge	Unknown	Located	Not Leaking	
⊗	A1	186 107	XLarge	Unknown	Located	Not Leaking	
⊗	m	550 447	Small	None	Located	Not Leaking	
⊗	A2	249 21	Large	Unknown	Located	Not Leaking	
⊗	m	250 149	Small	None	Located	Not Leaking	
⊗	m	449 349	Small	None	Located	Not Leaking	
⊗	SB	305 310	XLarge	None	Located	Not Leaking	

Figure 9: Object List created during dialogue in Figure 8

To assess the impact of coordinating representations quantitatively, a formal study was performed comparing the performance of teams with and without CRs. Each test group consisted of three teams of three subjects. Subjects were a mix of area professionals, mostly in computer-related industries, and undergraduate students; all were paid a flat fee for the experiment.

The results of this study, shown in Table 1, compare the final five hours of play for each group (by which time performance had stabilized), and these results were normalized over a general measure of complexity for each of the problems solved. The performance of the test group that had the CRs was better across several measures: amount of chat, number of errors committed, number of system events generated (an indicator of interface work), and clock time.

Table 1: Improvement of CR groups over non-CR groups; final 5 hours of play

Indicator	Improvement (reduction)
Communication	57% ($p < 0.01$)
Domain Errors	61% ($p < 0.2$)
System Events	38% ($p < 0.06$)
Clock time	49% ($p < 0.01$)

The most significant effect is the 57% reduction in communication generated. This confirms the observation that a significant amount of communication was offloaded into the CRs. Also highly significant is the 49% reduction in clock time. Only slightly less significant is the reduction in system events (mouse clicks, etc.), down 38%. Additionally, overall domain errors (errors in performing domain actions which led to a toxic spill) were reduced by 61%. The variance of this measure was quite high due to the overall low frequency of this kind of error; this reduced its confidence below statistical significance ($p < 0.2$).

The above results demonstrate that the introduction of the Object List modifies the representational system from the user’s perspective in ways that improve the ability of the collaborators to coordinate information. It is shown in the following how the Object List also addresses the representational needs of the designer.

4.2 Improvements from the adaptive component designer’s perspective

The Object List captures information about the world that would otherwise be exchanged in chat, and transforms it into a form that can be easily used in the intent inference process described above. In addition to information about where the wastes are, what equipment they need, and what size they are, the Object List also provides the system with a set of user-assigned labels. These labels can be used to mine chat for additional information about users’ intentions. Detailed information about how chat information was used can be found in Appendix A.

The relative utility of information from the Object List for performing intent inference is evaluated in the following. This is done by comparing the ability of the BN to predict user lift actions in historical log files across several information conditions. Before presenting results, the evaluation methodology is described in detail.

4.2.1 Evaluation Methodology

To evaluate the utility of information from the Object List, the BN was trained and tested under four information conditions using a single dataset (described in Table 2), and its performance was compared across these conditions. The dataset consists of time-stamped information derived from the log files collected from each use session in the dataset. It contains the state information required by the BN. State information includes information about each waste (position, size, and equipment needs), the users’ locations and heading, and wastes as they are lifted. The dataset also contains information about references made by users during chat to objects that are listed in the Object List.

Table 2: Dataset for the evaluations

Team	Sessions	Avg. # of wastes per problem	Total Hours
Team 1	10	11.7	9.9
Team 2	6	11	8.4
Team 3	9	14.3	9.1
Team 4	16	14.5	8.7
All	41	13.5	34.3

Each condition changes the amount and quality of information about wastes in the dataset. The four information conditions, and the information available under each condition, were:

1. **Perfect Information** – Under this condition, the dataset contains complete and accurate information about the position, size, and equipment needs for every waste at each point in time. Perfect information is not typically available at runtime, but it can be derived from the problem files which define the initial problem configuration for a usage session. This is an ideal case, and provides a baseline against which to compare the other information conditions.
2. **Object List** – Under this condition, the dataset contains information about the position, size, and equipment needs for wastes that are entered in the Object List, as it becomes available. This information is subject to errors, duplication, and omissions. Compared to the Perfect Information condition, performance of the BN under this condition is indicative of the quality of the information users provide via the Object List.
3. **Perfect info + Chat** – All of the information in the Perfect Information condition, plus the occurrence of names of objects (as entered by users into the Object List) in chat. More precisely, there is an entry in the dataset for each time a word appears in chat matches one of the labels for a waste in the Object List (see Appendix A). This condition is used to investigate how much predictive information references to wastes in chat add.
4. **Object List + Chat** – Information from the Object List condition, plus the occurrence of names for objects in chat. This condition represents the best possible runtime performance of our inference procedure without information from problem files.

Because the dataset contains different information under each condition, the probability distribution stored in the BN may be optimal for one information condition but not another. If the evaluation were to be performed using a single probability distribution for all data conditions, results might reflect the “preference” of that particular probability distribution, rather than the utility of information contained in the dataset.

To overcome this problem, the probability distribution of the BN was fit to the dataset in each information condition before testing. To train the network, each dataset is compiled into a series of cases. A case consists of values for each non-hidden node in the BN (see Section 2.1), and each case represents a change in the relevant information about the world. During the training phase, values for the “Will Lift” node were extended back in the dataset from each actual lift for a ten minute time window. The size of the time window was based on initial experimentation, and chosen to maximize the performance of the BN.

The $EM(\eta)$ algorithm (Bauer, et al., 1997) was used to train the network, and the same starting parameters for the BN were used in each case. The $EM(\eta)$ algorithm is a generalized version of the standard

EM algorithm with a learning parameter. When $1.0 < \eta < 2.0$, the $EM(\eta)$ algorithm is significantly faster than the standard *EM* algorithm, but still has good convergence properties. For all of our training sessions, $\eta=1.8$. Algorithm performance was further optimized by posting only unique cases to the network, and weighting each unique case in the training set according to the number of times it occurred.

For the evaluation, two performance metrics were calculated: the correct goal rate (CGR), which is the proportion of correctly guessed goals; and the false positive rate (FPR), which is the proportion of guesses that were false. A guess is made whenever a relevant state variable changes. Any uninterrupted sequence of correct guesses leading up to the step immediately preceding the execution of the predicted goal is counted as a single correct goal. The total number of goals is the number of wastes lifted. Thus,

$$CGR = \text{correct goals} / \text{total goals}$$

$$FPR = \text{incorrect guesses} / \text{total guesses}$$

4.2.2 Evaluation

The evaluation was performed solely to compare the utility of different information sources for intent inference, and in particular, the utility of the information derived from the object list. These performance metrics are not indicative of how well the intent inference procedure performs within the context of a specific adaptation. Studies of our adaptive component are documented in the next section.

The results of the evaluation are shown in Table 3, and reflect average performance across all teams for each data condition. A single factor ANOVA demonstrated that differences between information conditions were highly significant for both CGR ($F(131,3)=10.84, p<.0001$) and FPR ($F(131,3)=3.98, p<.01$).

Table 3: Intent inference results for different information sources

Condition	Correct Goal Rate (StdDev)	False Positive Rate (StdDev)
Perfect Info	.83 (.14)	.53 (.13)
Object List	.70 (.17)	.60 (.16)
Perfect Info + Chat	.87 (.12)	.51 (.11)
Object List + Chat	.77 (.15)	.58 (.15)

The “Perfect Info” case, in the top row of the table, provides a baseline against which results for the other conditions may be compared. It is a rough indicator of the best the intent inference procedure can do, given only complete information about the state of the world. Across the four teams in the dataset, the Correct Goal Rate for the “Perfect Info” case ranged from .77 to .91 for this condition. There was a weak correlation between problem size (number of toxic wastes) and individual team performance ($r=.23$), reflecting the fact that it is more difficult to make good guesses when there are more options to choose from. In general, these metrics indicate that the inference procedure is effective.

As expected, the intent inference procedure does not perform as well with information from the Object List alone. However, results from the “Object List” condition were still good, and demonstrate that use of the Object List was reliable enough to be useful for intent inference.

The “Perfect Info + Chat” condition demonstrates that references add significant information that cannot be derived from knowledge about the state of the domain. Thus, regardless of access to state information (for instance, if there were intelligent sensors placed in the world) the Object List adds information that still improves intent inference.

The combination of reference information from chat and domain information from the Object List (the “Object List + Chat” condition) improves the performance of the procedure to a point where it is nearly

as good as the “Perfect Info” condition. These results confirm that, by modifying the representational properties of the media between the user and adaptive component designer, high quality runtime information can be made available for intent inference.

To demonstrate that this information is sufficient to drive a useful adaptive support mechanism at runtime, an adaptive component was implemented and its use evaluated. These results are described in the following section.

5 An Adaptive Component

As described in Section 2.1, planning in VesselWorld is a laborious and error prone operation, and user errors are frequently the result of forgotten plan steps or joint plans that have become unsynchronized. The adaptive component was thus designed to help users formulate basic individual and joint plans. Using the intent inference procedure, the system can provide a small set of likely plans from the hundred or so that are possible at any point in time.

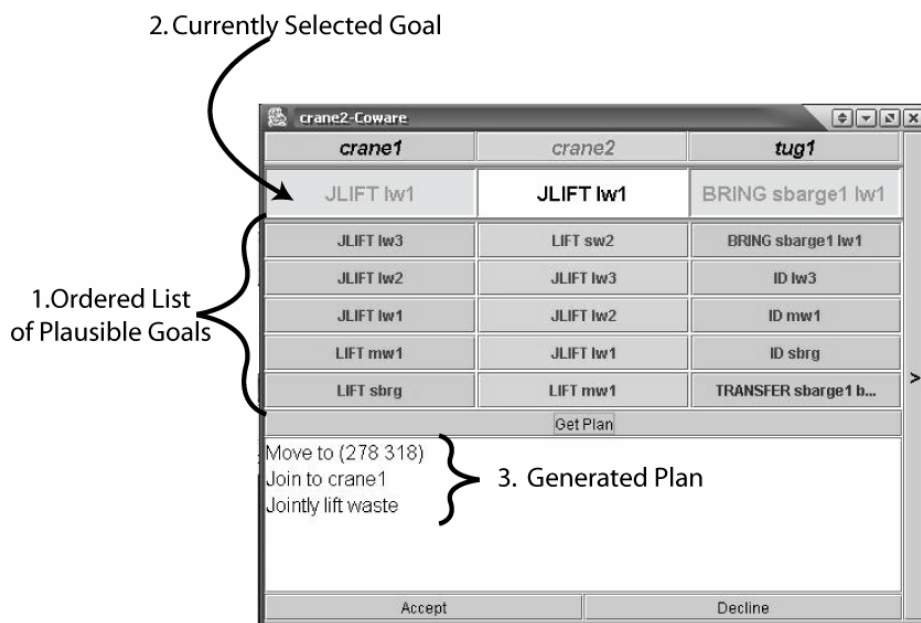


Figure 10: The adaptive component

The interface to the component is shown in Figure 10. It displays inferred plausible domain goals for each of the users. Like the Object List, it is a WYSIWIS component, so each user can see the other users’ goals as well. Users can select plausible goals from their own column to have the system automatically generate a plan. The component can generate individual as well as shared plans. If the information in the Object List is correct, the generated plans are guaranteed to be correct.

The detailed function of the adaptive component is as follows:

1. Each time state information is updated (e.g., when plans are executed, information is added to the Object List, an object reference appears in chat, etc.) the system offers each user up to five plausible goals, displayed in order of priority and color-coded according to the system’s “confidence” in that prediction. Each user can only select goals from their column.

2. When the user selects a goal, it is copied into the top row, which displays each user's currently confirmed goal. The user then has the option to request an automatically generated plan that will accomplish the selected goal (the "Get Plan" button shown in Figure 10).
3. The system generates a plan that the user can inspect. In cases where the goal involves multiple actors, the other actors are invited to join the plan. If all invited actors accept the invitation, a plan is generated; if invited actors do not accept the invitation, the requesting user is so informed.
4. If the user accepts the plan (by clicking the "Accept" button in Figure 10), it is automatically copied into the user's planning window for execution.

If the plan is generated from correct state information, no user modifies the state in such a way that conflicts with the generated plan, and the plan is executed to completion, it will succeed in achieving the desired goal.

5.1 User Studies

We performed a formal study with four teams to evaluate the use of the adaptive component and its effects on performance. Two of the teams were given access to the adaptive component, and the others were not. The participants were a mix of students and local-area professionals, with varying degrees of computer proficiency. Each team was trained for two hours in use of the system, and then solved randomly chosen VesselWorld problems for approximately ten hours. To alleviate fatigue concerns, the experiment was split into four three-hour sessions.

For the teams with the component, the inference procedure used information from the Object List and chat (the Object List + Chat condition) to infer user goals and construct plans. The following results report on the last 5 hours of play time for each group, by which time performance of the users had stabilized.

5.1.1 The component was heavily used

All groups used the component to generate plans within the system. On average, users confirmed a goal every 1.5 minutes (SD=46 seconds), requested a plan for each confirmed goal, accepted 71% of plans requested (SD=19%), and completed the execution of 83% (SD=6.75%) of these plans. Overall, this indicates that roughly 59% of confirmed goals resulted in a plan that was executed to completion. For each problem solving session, one quarter of all plan steps submitted to the server were generated by the component (SD=8%).

The component generated plans for 43% (SD=15%) of the domain goals it could have predicted for the Cranes (some goals, like "search the harbor" were not inferred). It was not possible to obtain a similar statistic for the Tug operator because it is difficult to recognize the tug's goals in the collected log files (goals for the tug are not bracketed by easy to detect plan steps like "LIFT" and "LOAD").

5.1.2 The component improved users' performance in the domain task

The groups that had the component had 45% ($p < .10$) fewer joint errors (failures during joint actions) per minute than the groups that did not. This difference is not significant at the .05 level because of the small sample size and overall low proportion of joint errors. A reduction in joint errors corroborates prior analysis of use of the VesselWorld system, which indicated that joint errors were usually the result of plan submissions becoming unsynchronized (Alterman, et. al, 2001). Because the component generates coordinated plans in advance, users could simply submit each step and be assured that actions would be coordinated.

5.1.3 The component reduced cognitive effort

To measure the change in cognitive effort between the two populations, the amount of time it took users to execute plans and the amount of interface work were evaluated. It was found that the amount of clock time taken by users between submitting steps of automatically generated plans was 57% less ($p < .01$) than in groups without the adaptive component, but that there were no significant differences in the number of mouse clicks per waste. Because the reduction in clock time for groups with the component cannot be explained by a reduction in the amount of interface work, we conclude that the component reduced the cognitive effort of the collaborators.

As stated in Section 1.2, these studies show that the adaptive component was heavily used, improved user performance, and reduced cognitive effort during plan execution. This is verification that collaborating users can generate enough structured information when using a coordinating representation to drive useful intelligent support. We conclude that the approach to adding intelligent support demonstrated here was successful for this domain.

6 Summary of the Methodology

This work has demonstrated how the design of a groupware system can be modified to address the needs of two intersecting representational systems. This is done by adding structure to the system in the form of a Coordinating Representation. The CR helps users coordinate and improves performance in the domain task. It also renders information previously exchanged via an unstructured communication channel accessible for use in an adaptive support system.

Clearly, much in this approach rests on the design of the CR. To facilitate the application of the approach to other domains, a repeatable design methodology for CRs with the necessary properties has been developed (Alterman et al., 2001, Landsman & Alterman, 2005; Feinman & Alterman, 2003). This methodology can be summarized in four steps:

1. Transcripts are collected of runtime user behavior.
2. These transcripts are analyzed in order to identify weak spots in the user's representational system.
3. This analysis is used to develop coordinating representations that people will use, and that improve coordination.
4. Information collected by the CRs is leveraged to drive adaptive components.

Transcript collection has roots in ethnography and interaction analysis. Traditionally, ethnography has relied upon *in situ* observation. Suchman & Trigg (1991) and Goodwin & Goodwin (1996) explored the use of video recordings in capturing and analyzing naturally occurring workplace activity. While video is a very powerful technique for studying a work practice, it is a resource intensive process. For the ethnographer, groupware makes life much easier, because a large portion of the relevant activity is mediated by the system itself, and can thus be passively recorded and analyzed offline.

To perform the kind of ethnographic analysis required to identify representational deficiencies in groupware, it is helpful to be able to replay transcripts, rewind when necessary, easily locate high-level task events (e.g., the submission of plans), and add annotations. To provide this functionality, a component-based software framework called THYME was developed (Landsman and Alterman, 2005). Groupware systems built using THYME automatically log all usage data, and replay tools can be rapidly

generated. The replay tools provide a console similar to that on a VCR, but with all of the above features.

Two techniques have been developed to analyze transcripts. One of them, which was described briefly in Section 3.1, examines users' recurrent activities to identify weak spots in an existing representational system (Alterman et al., 2001). The other technique examines co-reference chains (Feinman & Alterman, 2003). To perform this later analysis, the analyst identifies and tags all references to objects, plans, and other entities of interest in a logged dialogue. Metrics are then calculated for the distribution and lifespan of each type of reference and these metrics are used to identify the need for supporting structure. Empirical evidence has been collected demonstrating that these analysis techniques can be used to make predictions about the type of structure that will be useful (Feinman, 2006).

These analyses guide the development of CRs that address representational needs of both the user and the designer. Because CRs fix problems with an existing representational system for the users, they will improve the users' ability to stay coordinated, and users will use them. This article has demonstrated how the structure introduced by CRs can then be leveraged to introduce adaptive support.

7 Balancing User Effort and Adaptive Component Designer Effort

As discussed above, adaptive software can be viewed as the mediating layer in two intersecting representational systems. This situation is depicted in Figure 11. In the diagram, the software system is presented as a composition of three components. This is intended as a descriptive conceptual breakdown only, and is based loosely upon early adaptive software models that grew out of user-interface management systems (e.g. Hefley & Murray, 1993).




Figure 11: Two intersecting representational systems

The three components of this model are;

The interface - Contains the signals and controls that the user interacts with directly. It is analogous to Norman's (1991) surface representation.

The domain model - The system's internal representation of the task-domain. It defines the functionality of the system. The domain model is very often implicit in the code that makes up a system, rather than represented concretely.

The adaptive component - May contain a user model, domain knowledge, an inference engine, and other technology related to run-time adaptation. It serves as a proxy that supports the time delayed interaction between the designer and user. It encapsulates the designer's assumptions about the user's potential needs at runtime, what information is needed to recognize those needs, and how they should be met.

There are two representational systems in the diagram. In one, software mediates the interaction between the user and the domain; this is the user's perspective. In the other, the software mediates the interaction between the user and designer of the adaptive component; this is the designer's perspective. This conceptualization highlights what are often considered competing design requirements for adaptive systems (Schneiderman, Maes; 1997). The choices that are made with respect to the balance between these requirements are reflected in the design of the interface. These choices have consequences for the amount of work the user must do, and the resources expended in developing the system.

Several other systems are examined here with respect to the balance struck between these competing design requirements. This balance is characterized by the degree to which specific interface features are incorporated to meet the design requirements from either perspective. The set of features meeting design criteria from the user's perspective are referred to as a *task-language*. A task-language is devoted to supporting user's manipulation and interpretation of the domain-task. In VesselWorld, a user plans a step to move his ship closer to a waste using the task language.

The set of features meeting the designer's criteria (from the perspective of the intent inference engine) are referred to as a *meta-language*. A meta-language is not the task-language itself, but is rather a set of interface features that allow the user to tell the system *about* the task being performed. VesselWorld does not have a meta-language.

7.1 Task Languages

Task-languages are designed to mediate and support user action in a domain. Examples of task-languages include direct manipulation and command-line interfaces. In the case of groupware, a chat window is part of the task-language, because it mediates the user's coordination with other users (part of the domain task). Some user-adaptive systems infer user's needs by monitoring their task interaction alone. These systems contain few or no meta-language operators (Figure 12).

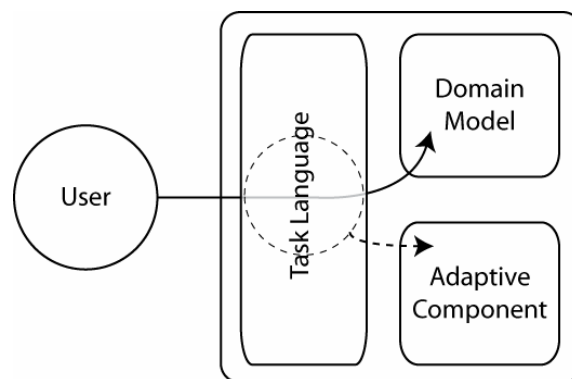


Figure 12: Adaptive system with task-language

In such systems, the feasibility of and effort involved in developing the adaptive component depend upon how easily interaction data can be interpreted against the backdrop of available task and domain knowledge. For example, the SmartEDIT system (Lau et al., 2000) is a programming by demonstration

system that can infer user plans for making syntactical edits to a document. The system does employ a minimal meta-language that requires the user to segment the execution trace into training examples, but its inference mechanism interprets task-level information directly.

SmartEDIT's inference mechanism (a version space algorithm) is successful for two reasons. All relevant user information is readily available by monitoring changes to the text buffer, the cursor location, and the clipboard contents. Furthermore, a set of domain operators that cover a large portion of the syntax editing task-domain are readily specified. For such well-defined domains, the costs associated with knowledge engineering can be relatively small.

Recommender systems are a different class of user-adaptive software that require minimal or no meta-languages. Among the most widely used of these (e.g. Amazon.com; Google), knowledge acquisition depends only upon a user's task level behavior – browsing, purchasing, or searching. The two most commonly used inference techniques in widely deployed recommender systems are collaborative filtering (Goldberg, et. al, 1992) and text-based information retrieval. For collaborative filtering, domain knowledge is based upon statistical patterns of many users performing similar (browsing, purchasing, or searching) behaviors. Text-based information retrieval depends upon statistical patterns of words in a large corpus of documents that others have created. In both cases, the bulk of the knowledge engineering effort is offloaded to external sources.

Generally then, for adaptive systems that employ a task-language:

User effort to engage the adaptive component is low. The user is willing to use the interface; this is the sole interface through which the user interacts with the domain. This does not imply that the system is easy to use; it only means that knowledge acquisition does not introduce much extra work for the user.

Adaptive component design effort is domain dependent. Often, importing a sufficient amount of domain knowledge into the system requires a significant investment of engineering effort, though sometimes this domain knowledge is available from external sources.

7.2 Meta Languages

When the system cannot extract enough information from the user's interaction with a task-language to support adaptation, a meta-language can be included. Examples of meta-languages include preference dialogs used to configure a system's behavior, or a natural language interface to an agent. The defining feature of a meta-language is that its primary role is to address the knowledge acquisition needs of the designer.




Figure 13: Adaptive system with meta-language

Meta-languages cover the spectrum from formal to natural, and may support a more or less objectified interaction with the adaptive component. The Collagen system (Rich & Sidner, 1998; Lesh, et al., 1999) is an example of a system that used a formal meta-language to support an objectified interaction with an agent. Collagen is a middleware framework for the development of collaborative agents that can provide guidance and offload some user work in the domain task. In its initial formulation, Collagen included a formal meta-language called ADL (Artificial Discourse Language) (Sidner, 1994). The user provided the system with information about her goals via a menu-driven interface to ADL (which provided a list of valid dialogue moves), and the intelligent component then used this information to provide support in the form of suggestions about next steps and contextual reminders.

In Collagen, ADL reduces development costs in two ways. It provides the system with structured information about the user's intentions. It also enforces a particular interaction between the user and adaptive component, so that domain knowledge and the algorithms associated with its use can be represented in a portable and domain independent way. In this manner, Collagen forms an adaptive systems framework that can greatly reduce development costs in the long-term.

A drawback with the formal / objectified approach is that the user is required to manage two tasks instead of one. One is her desired domain task, and the other is explaining her actions to the system. This second task is made even more cumbersome when it requires the use of a formal language. As was noted by Collagen's developers, "it is often more efficient and natural to convey intentions by performing actions" (Lesh et al., 1999; page 23).

One way to avoid requiring that the user manage two tasks is to combine the meta-language with the task-language, so that the interaction no longer objectifies the adaptive component and the user is free to focus on the domain task. For example, the Epsilon collaborative learning environment requires interlocutors to use sentence openers that represent speech acts during chat, and this structure is in turn used to identify problems with the collaborative process (Soller, 2004). The task and meta-language are coincident in Epsilon, but still distinct, as the structure that is added to the task-language is done so solely to address designer's knowledge acquisition needs. This kind of approach can lead to interference with the user's domain-task.

Natural meta-languages (which rely upon natural language understanding) can be employed to avoid problems with formal languages; they provide a rich source of information about the user, and potentially require little or no training to use. Interaction via a natural meta-language usually requires an objectified adaptive component (a conversational "agent"), which potentially raises the same "two-task" problem as above. However, sufficiently advanced natural language would allow the user to choose the best modality for the task at hand without ever having to focus on the meta-task of making herself understood.

The major drawback with natural language understanding is that technology is not yet advanced enough to scale beyond special purpose applications (Zukerman & Litman, 2001). This results in additional work from both the user and designer's perspective. To the user, the interface may appear to support a wider range of conversational moves and vocabulary than it actually does, and the user is left to navigate a dialog through trial and error (Schneiderman, 1998). For the designer, a good deal of effort must go into building the language recognition facility, and input to the intelligent component cannot be easily constrained, and so may be more difficult to process.

Numerous approaches have been developed to manage the tradeoffs when employing meta-languages. For example:

- Mixed initiative systems, such as Lookout (Horvitz, 1999) try to calculate the trade-off between the benefit provided by the adaptive system and the additional work the meta-language requires from the user. An assumption in this approach is that at certain points, the user may *want* to engage an intelligent agent, and that this can be recognized with a sophisticated enough domain task model and processing of the task-language.
- Some systems adapt the meta-language itself over time by allowing the user to introduce or re-define terms. For instance Pan, et al., (2005) describe a system in which the system and the user adapt to one another to achieve a common vocabulary. This form of two-way adaptation is an interesting take on end-user-programming (Nardi, 1993), and may lower the bar for casual users that would not otherwise be interested in customizing the system's behavior.
- Multi-modal systems incorporate other input modalities along with limited natural language processing to improve system comprehension. The CASIS system (Leong, et al., 2005) uses contextual cues like room temperature, ambient noise, and room activity to reduce error rate in speech recognition. Such systems are still in early phases of research, and require specialized hardware in controlled settings, but offer significant performance gains over mono-modal meta-languages.

Generally then, for adaptive systems that employ a meta-language:

User effort may be too high. Whether or not people are willing to use a meta-language depends on the potential benefit and the design of the language. Using a meta-language can disrupt the domain task, either by requiring the user to manage an additional task, or by interfering with the task-language. Meta-languages can trade reduced design costs for additional user effort.

Adaptive component design effort can be reduced by implementing a meta-language, although this depends on the type of language used. Often, a designer's efforts to make a meta-language more appealing to the users, as with the case of natural languages, ultimately results in more overall work (for both the designer and the user) than the meta-language saves in terms of making inference easier.

7.3 Coordinating Representations: Combining Task and Meta languages

The approach described in this article employs a coordinating representation to serve the needs of both the designer and the user. The task and meta-languages are not merely coincident (as with the Epsilon system above). Instead, the CR combines the task-language and meta-language into a single interface mechanism that serves both roles (Figure 14).

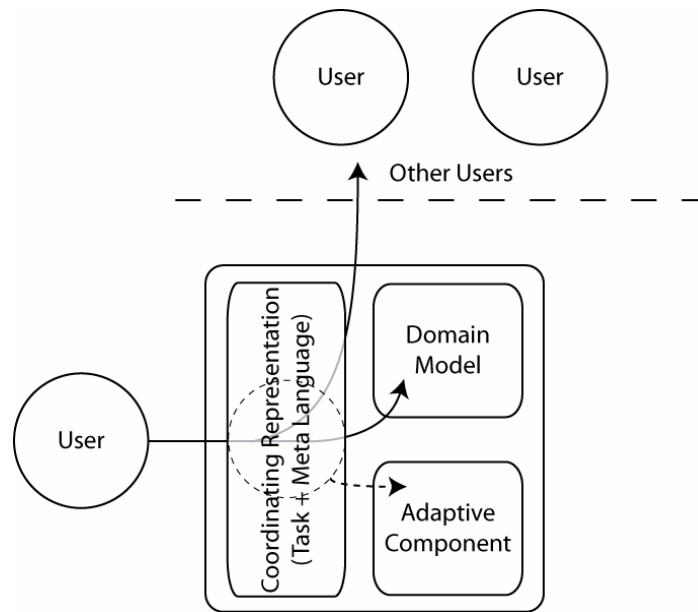


Figure 14: Adding structure to a task language to enable intelligent support in groupware.

The general idea of combining user needs and designer needs into a common interface tool is not restricted to collaborative systems (see St. Amant, et al., 2003, for a discussion). However, as we have shown, groupware is particularly amenable to the approach because of the difficulty that users have in coordinating their activities. For VesselWorld, it was precisely the difficulty in maintaining common ground that makes the structure of the Object List so useful for users. This structure is equally useful from the designer's perspective, because it renders the users' task-relevant common ground accessible to runtime inference mechanisms.

There are other systems that use coordinating representations in a similar manner, although their role is rarely given such a central focus. Jameson, et. al. (2003) describes a collaborative multi-agent system for scheduling a business trip, in which agents serve as proxies for users in a scheduling scenario. The role of the agent, and one of the main foci of that work, is to provide a social face to help users stay aware of others' preferences and motivations. However, most of the information that is derived from the users comes from a coordinating representation that allows them to express their preferences and explore possible solutions to the constraint problem set up by the collaborative task.

Concept maps are a class of CRs that have proved useful in adaptive systems. Concept maps are visual representations of interlinked concepts that can be useful for people trying to communicate about large conceptual structures or coordinate knowledge engineering efforts (Novak & Gowin, 1984). Cañas & Carvalho (2004) describes several adaptive systems that use concept maps in collaborative settings. One system uses the structure in a concept map, along with WordNet (Miller, 1990) to disambiguate word senses. Another system draws conclusions from the concept maps shared by students in a collaborative learning environment. These conclusions can help the students to identify weaknesses or inconsistencies in their own conceptual structures.

Generally then, for adaptive systems that employ coordinating representations:

User effort is reduced. Communicating coordinating information, such as required in maintaining common ground, is part of the domain task in a groupware system. As set forth in Section 4.1, our

empirical data proves that CRs can make this easier, and because CRs actually reduce user workload, users are willing to use them.

Adaptive component design effort is low. Information can be easily leveraged for adaptation because it is highly structured, and germane to the users' joint task. Our approach in VesselWorld used widely available AI techniques, and the data discussed in Section 5.1 validates the effectiveness of the resultant adaptive system.

8 Generality & Future Work

Coordinating representations offer a means for merging two traditionally competing design requirements for user-adapted systems into a single representational medium. On the one hand, CRs can reduce user workload and improve coordination. On the other, by structuring coordinating information, CRs provide the designer with a rich source of user and context information that can be leveraged by an adaptive support system at runtime.

In our methodology, coordination problems become opportunities for knowledge acquisition. An analytical procedure has been provided to identify these problems, and it can be applied to any media that a community of actors uses to coordinate a recurring activity. Relatively complete usage data is required to perform the prescribed analysis procedure. In the case of groupware that mediates the majority of the users' interactions, this may be achieved through the addition of logging utilities to the groupware system itself. For collaborations which extend outside of the media under investigation, more traditional forms of data collection (e.g. videotape) will be required.

We have validated our approach via case studies that demonstrate its effectiveness in VesselWorld. The coordination problems experienced by VesselWorld users are highly representative of those problems in other collaborative environments. Users have different information about the world and need to share this information to accomplish their joint task. Individuals change objects in the world and need to inform others of these changes. Users have different roles and responsibilities. The task requires close coordination at times, but long term commitments are also a critical part of the activity.

In this article, we've focused primarily upon problems related to grounding, but our analysis has revealed other types of coordination problems as well. Alterman et al. (2001) document the problems users faced in sequencing closely coordinated activities and tracking long-term commitments, and describe CRs that were developed to address these problems.

Similar kinds of problems have been described in the context of modern military applications. As networked technology has become ubiquitous in the military, electronic chat has found its way into mission-critical applications. Usage patterns reminiscent of the coordination problems in VesselWorld have been documented. Chat history is heavily relied upon to recall missed or forgotten information, and there are problems keeping track of mission orders (Heacox, et al. 2004). The need for adaptive chat mechanisms that are more aware of the user's task so as to minimize interruptions (Cummings, 2004) has also been noted.

Our approach is agnostic with regards to the particular inference technique that is used. Bayesian Networks were employed because they are easy to use, were sufficiently powerful for our needs, and many open-source and off-the-shelf implementations are available. Other techniques might have been used instead of or in addition to BNs. Sophisticated NLP algorithms could likely derive more mileage from references to wastes in chat than we were able to with our procedure. The training algorithms used to analyze the BN's performance could be employed at runtime to adapt to user behavior over time. More

sophisticated planning algorithms and optimization routines could have been employed to produce efficient plans given available waste information. The point is that a range of possibilities for building adaptive support cascade from the increased availability of structured information about the user. The approach is then limited by the relevance of the information to the adaptation goal, the power of available inference algorithms and expertise of the designer.

The end result of applying our methodology to VesselWorld was a successful adaptive plan-generation component. This component is entirely domain specific, and may not be necessary or feasible in other domains. A natural extension to the existing methodology would offer guidance to the designer regarding the type of adaptive support that may be desirable or feasible using coordinating information from CRs. The development of such an extension is our current and future focus.

An avenue we are actively exploring is the generation of awareness information that is generally not available in collaborative applications. Much work on awareness has focused upon replicating information that is lost in distributed settings, such as awareness of what is going on (Dourish & Belotti, 1992), who is around (Dourish & Bly, 1992), and what has changed in a shared environment (Gutwin & Greenberg, 1998). Recently, Carroll, et al. (2003) has investigated the potential role of “activity awareness” in collaborative applications. The term activity awareness describes awareness of how work is embedded within the context of the overall activity. Coordinating representations are useful for generating the information required to provide this kind of awareness because they capture and structure information that constitutes the users’ shared context. The computer, as a mediating artifact with significant abilities to summarize, sort, and synthesize structured data is in a good position to automatically combine and provide this information as feedback to users.

Acknowledgments

This research was supported by the Office of Naval Research under grants No. N00014-96-1-0440 and N66001-00-1-8965.

References

1. Albrecht, D., Zukerman, I. and Nicholson, A.: 1998, “Bayesian models for Keyhole Plan Recognition in an adventure game.” *User Modeling and User-Adapted Interaction* 8(1-2), 5-47.
2. Alterman, R., Feinman, A., Landsman, S. and Introne, J.: 2001, “Coordinating Representations in Computer-Mediated Joint Activities.” *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, Boston, MA, 43-48.
3. Anderson, J. R., Corbett, A. T., Koedinger, K. R. and Pelletier, R.: 1995, “Cognitive Tutors: Lessons Learned.” *The Journal of the Learning Sciences*, 4(2), 167-207.
4. Bosma, W.; Andre, E : 2004, “Exploiting Emotions to Disambiguate Dialog Acts.” In *Proceedings of IUI’04*. Funchal, Madeira, Portugal., 85-92.
5. Bauer, E., Koller, D., & Singer, Y.: 1997. “Update rules for parameter estimation in Bayesian networks.” In *Proceedings 13th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*., 3–13
6. Brennan, S., “The Grounding Problem in Conversations With and Through Computers.” In Fussell, S. & Kreuz, R. (Eds.), *Social and cognitive psychological approaches to interpersonal communication*. Hillsdale, NJ: Lawrence Erlbaum. 201-225.

7. Cañas A. J. Carvalho, M.:2004, “Concept Maps and AI: an Unlikely Marriage?” *Proceedings of SBIE 2004: Simpósio Brasileiro de Informática na Educação*, Manaus, Brasil.
8. Carroll, J., Neale, D., Isenhour, P., Rosson, M.B., McCrickard, D.S.:2003, “Notification and Awareness: Synchronizing Task-Oriented Collaborative Activity.” *Int J. Human-Computer Studies*, 58, 605-632.
9. Charniak, E., and R. Goldman: 1993, “A Bayesian Model of Plan Recognition.” *Artificial Intelligence* 64:53-79.
10. Clark, H. H., and Brennan, S. A. (1991). Grounding in communication. In L.B. Resnick, J.M. Levine, & S.D. Teasley (Eds.). *Perspectives on socially shared cognition* . Washington: APA Books.
11. Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition* , 22, 1-39.
12. Clark, H. :1996. “Using Language.” Cambridge, Great Britain. Cambridge University Press.
13. Cummings, M.:2004, “The Need for Command and Control Instant Message Adaptive Interfaces: Lessons Learned from Tactical Tomahawk Human-in-the-Loop Simulations.” *CyberPsychology & Behavior*. 7(6):656-661
14. Dourish, P. and Bellotti, V.:1992, “Awareness and Coordination in Shared Workspaces.” *Proceedings of CSCW*, Toronto, 107–114.
15. Dourish, P., Bly, S., 1992. Portholes: supporting awareness in a distributed work group. *Proceedings of the ACM CHI '92 Conference on Human Factors in Computing Systems.*, pp. 541-547.
16. Engeström, Y.: 2000, “Activity theory as a framework for analyzing and redesigning work.” *Ergonomics*, 43:7. 960-974.
17. Feinman, A., and Alterman, R.: 2003, “Discourse Analysis Techniques for Modeling Group Interaction.” In Brusilovsky, P., Corbett, A., and de Rosis, F, editors, *Proceedings of the Ninth International Conference on User Modeling*, pages 228-237.
18. Feinman, A.:2006. "From Discourse Analysis to Groupware Design". PhD Thesis; Computer Science Department, Brandeis University.
19. Ferguson, G., Allen, J., Miller, B.:1996 “Trains-95: Towards a mixed initiative planning assistant.” In *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems*, 70-77.
20. Ferguson, G. and Allen, J.:1998. “TRIPS: An Intelligent Integrated Problem-Solving Assistant”, in *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, WI. 567-573.
21. Fischer, G.: 2001, “User Modeling in Human-Computer Interaction.” *User Modeling and User-Adapted Interaction*. 11:1-2, 65-86.
22. Goldberg, D., Nichols, D., Oki, B., and Terry, D. (1992) “Using collaborative filtering to weave an information tapestry. *Communications of the ACM*. 35(12):61-70.
23. Goodwin, C. and Goodwin, M: 1996, “Formulating Planes: Seeing as a situated activity.” In D. Middleton & Y. Engeström (Eds.) *Cognition and communication at work*, Cambridge University Press, 61-95.

24. Gutwin, C., Greenberg, S.: 2002, "A Descriptive Framework of Workspace Awareness for Real-Time Groupware." *Journal of Computer Supported Cooperative Work*. 11:411-446.
25. Heacox, N., Moore, R., Morrison, J., Yturalde, R.: 2004. "Real-time Online Communications: 'Chat' User in Navy Operations." *Proceedings of Command and Control Research and Technology Symposium*. San Diego, CA. (Available online at: http://www.dodccrp.org/events/2004/CCRTS_San_Diego/CD/papers/086.pdf).
26. Hefley, W., Murray, D.: 1993, "Intelligent User Interfaces." *International Workshop on Intelligent User Interfaces*. Orlando, FL. 3-10.
27. Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, D.: 1998, "The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users," In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. Madison WI. Morgan Kaufmann Publishers, pages 256-265.
28. Horvitz, E.: 1999, "Principles of Mixed-Initiative User Interfaces." In *Proceedings of CHI' 99*. Pittsburgh, PA. ACM Press, NY, 159-166
29. Hutchins, E.: 1995, "Cognition in the Wild." Cambridge, MA. MIT Press.
30. Jameson, A.: 2002. "Adaptive interfaces and agents." In J. A. Jacko & A. Sears (Eds.), *Human computer interaction handbook*. Mahwah, NJ: Erlbaum.
31. Jameson, A., Baldes, S., and Kleinbauer, T.: 2003, "Generative Models of Group Members as Support for Group Collaboration." *Workshop on User and Group Models for Web-Based Adaptive Collaborative Environments: Proceedings of the 9th International Conference on User Modeling*, 1-14.
32. Kaiser, E.: 2005. "Multimodal New Vocabulary Recognition through Speech and Handwriting in a Whiteboard Scheduling Application." In *Proceedings of IUI'05*. San Diego, CA., 51-58.
33. Landsman, S., Alterman, R., Feinman, A., Introne, J.: 2001, "Vesslworld and ADAPTIVE", *Brandeis University Tech Report CS-01-213*; Presented as a demonstration at Computer Supported Cooperative Work, 2000.
34. Landsman, S., and Alterman, R.: 2005, "Using Transcription and Replay in Analysis of Groupware Applications". *Brandeis University Technical Report CS-05-259*.
35. Lau, T., Domingos, P., and Weld, D.: 2000 "Version space algebra and its application to programming by demonstration." In *Proceedings of the Seventeenth Int'l. Conf. on Machine Learning*, 527-534.
36. Leong L.H., Kobayashi, S., Kshizuka, N., Sakamura, K.: 2005 "CASIS: A Context Aware Speech Interface System." In *Proceedings of the 10th International Conference on Intelligent User Interfaces*. 231-238.
37. Lesh, N, Rich, C., and Sidner, C.L.: 1999, "Using plan recognition in human-computer collaboration." In *Proc. 7th Int. Conf. on User Modeling*, pp. 23-32.
38. Maes, P.: 1994, "Agents that Reduce Work and Information Overload." *Communications of the ACM*, 37(7). 30-40.

39. Malone, T., Lai, K., Grant, K.:2001, "Two Design Principles for Collaboration Technology: Examples of Semiformal Systems and Radical Tailorability." In Olson, G., Malone, T., and Smith, J. *Coordination Theory and Collaboration Technology*. Lawrence Erlbaum Associates, NJ. 125-160.
40. Miller, G.:1990. "Five Papers on WordNet." *International Journal of Lexicography*, 3(4).
41. Nardi, B.:1993. "A Small Matter of Programming." Cambridge, MA. The MIT Press.
42. Norman, D. A.:1991. "Cognitive artifacts." In J. M. Carroll, editor, *Designing interaction: psychology at the human-computer interface*, Cambridge University Press. 17-38.
43. Novak, J. D. and Gowin, D. B.:1984. "Learning How to Learn." New York, NY: Cambridge University Press.
44. Pan, S., Shen, S., Zhou, M., Houck, K.:2005 "Two-Way Adaptation for Robust Input Interpretation in Practical Multimodal Conversation Systems." *Proceedings of the 10th International Conference on Intelligent User Interfaces*. 35-42.
45. Pearl, J.:1988, "Probabilistic Reasoning in Intelligent Systems." San Francisco, Calif.: Morgan Kaufmann.
46. Rich, C. and Sidner, C. L.:1998. "COLLAGEN: A Collaboration Manager for Software Interface Agents." *User Modeling and User-Adapted Interaction*. 8(3-4):315-350.
47. St. Amant, R., Dinardo, M., and Buckner, N. Balancing Efficiency and Interpretability in an Interactive Statistical Assistant. In *Proceedings of IUI'03*. (Miami, FL, 2003). ACM Press, New York, 2003. 181-188.
48. St. Amant, R., Young, R.:2001. "Interface Agents in Model World Environments." *AI Magazine* 22(4): 95-108.
49. Schmidt, K., and Simone, C.:1996 "Coordination mechanisms: Towards a conceptual foundation of CSCW Systems Design." *Computer Supported Cooperative Work. The Journal of Collaborative Computing*, 5(2-3), 155-200.
50. Sidner, C.L.:1994. "An Artificial Discourse Language for Collaborative Negotiation." *Proceedings of the Twelfth National Conference on Artificial Intelligence*. AAAI Press. Menlo Park, CA., 814-819.
51. Schneiderman B.:1998, "Designing the User Interface", Addison Wesley.
52. Shneiderman, B., Maes, P.:1997, "Direct Manipulation vs. Interface Agents: A Debate," *Interactions*, 4(6). 42-61.
53. Soller, A.:2004, "Computational modeling and analysis of knowledge sharing in collaborative distance learning" *User Modeling and User-Adapted Interaction*. 14: 351-381
54. Suchman, L. and Trigg, R.:1991, "Understanding Practice: Video as a Medium for Reflection and Design." In Joan Greenbaum & Morten Kyng (eds.), *Design at Work: Cooperative Design of Computer Systems*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.
55. Suthers, D. (2003) "Representational Guidance for Collaborative Inquiry." In *Andriessen, J., Baker, M., and Suthers, D. (eds.), Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments. Netherlands, Kluwer. 27-46.*

56. Traum, D.: 1994. "A Computational Theory of Grounding in Natural Language Conversation." Doctoral Thesis. Technical Report TR545. University of Rochester.
57. Vicente, K.: 1999. "Cognitive work analysis: toward safe, productive, and healthy computer-based work." New Jersey: Lawrence Erlbaum Associates, Inc.
58. Zuckerman, I. and Litman, D.: 2001, "Natural Language Processing and User Modeling: Synergies and Limitations." *User Modeling and User-Adapted Interaction*. 11: 129-158.

Appendix A: Using chat

Users in VesselWorld often refer to wastes in chat using the labels they've assigned to objects in the Object List. To incorporate this information into the intent inference procedure, we modified the BN based on the following analysis.

Table 4 depicts the likelihood that a reference for an object will appear during the three consecutive five minute windows prior to a lift of that object at time t . In the table, "Joint" and "Single" refer to whether a waste requires both or just one crane operator to lift. In the $\sim Lift$ conditions, values reflect the likelihood some waste is referred to prior to a lift of a different waste.

Table 4: Probability of reference preceding a lift at time t

	$t-5$ to t		$t-10$ to $t-5$		$t-15$ to $t-10$	
	Joint	Single	Joint	Single	Joint	Single
Lift	.62	.42	.27	.15	.25	.08
$\sim Lift$.15	.11	.10	.07	.08	.04

There is about a sixty percent chance that waste will be referred to in chat in the five minutes preceding the lift if that waste requires assistance, and about a forty percent chance if that waste can be lifted singly. Chat references are predictive of lift actions for roughly a fifteen-minute window of time preceding a lift. On the basis of this analysis, we expanded our BN to include three five minute windows of chat history, with one node for each five minute window. The expanded network is shown in Figure 15.

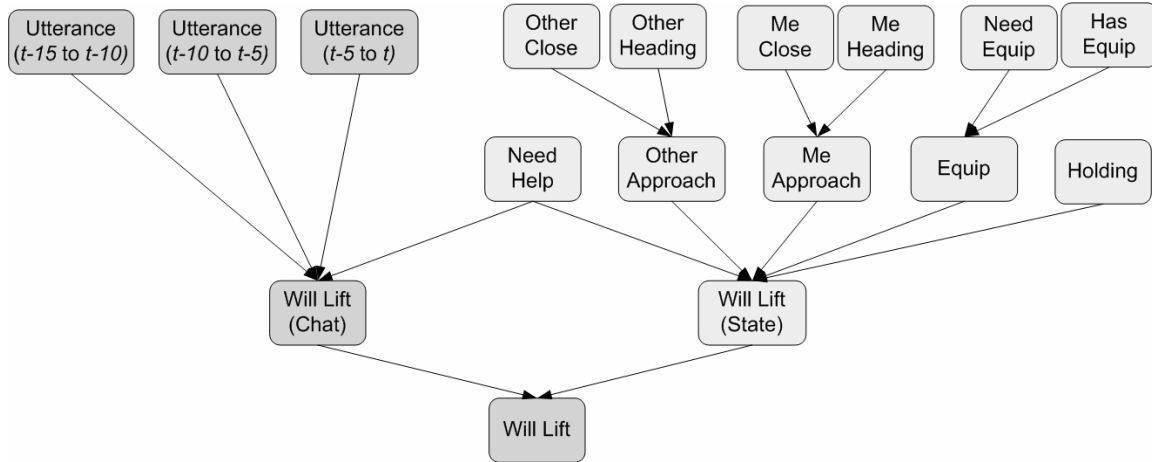


Figure 15: BN with nodes for chat added; added nodes are colored differently

Table of Contents

1	Introduction.....	2
1.1	Coordinating Representations.....	2
1.2	Organization of the argument	3
2	VesselWorld.....	5
2.1	Intent Inference in VesselWorld.....	6
3	Representational problems in VesselWorld.....	8
3.1	Problems from the user’s perspective.....	8
3.2	Problems from the adaptive component designer’s perspective.....	10
4	Using CRs to fix representational problems.....	11
4.1	Improvements from the user’s perspective.....	12
4.2	Improvements from the adaptive component designer’s perspective.....	14
4.2.1	Evaluation Methodology.....	14
4.2.2	Evaluation	16
5	An Adaptive Component	17
5.1	User Studies	18
5.1.1	The component was heavily used.....	18
5.1.2	The component improved users’ performance in the domain task.....	18
5.1.3	The component reduced cognitive effort.....	19
6	Summary of the Methodology	19
7	Balancing User Effort and Adaptive Component Designer Effort.....	20
7.1	Task Languages	21
7.2	Meta Languages.....	22
7.3	Coordinating Representations: Combining Task and Meta languages.....	24
8	Generality & Future Work.....	26
	Appendix A: Using chat.....	32

List of Figures

Figure 1: The complex sheet (Suchman & Trigg, 1991)	3
Figure 2: The VesselWorld system.....	5
Figure 3: Schematic of BN used to infer Crane Lift Intentions (information is posted to numbered nodes)	7
Figure 4: Explicit talk about coordination	9
Figure 5: Excerpt from chat during VesselWorld session	10
Figure 6: Different ways of referring to wastes, within one group.....	11
Figure 7: The Object List.....	12
Figure 8: Dialogue from users before an after the introduction of the object list.....	13
Figure 9: Object List created during dialogue in Figure 8.....	13
Figure 10: The adaptive component	17
Figure 12: Two intersecting representational systems.....	20
Figure 13: Adaptive system with task-language.....	21
Figure 14: Adaptive system with meta-language.....	22
Figure 15: Adding structure to a task language to enable intelligent support in groupware.	25
Figure 16: BN with nodes for chat added; added nodes are colored differently	32

List of Tables

Table 1: Improvement of CR groups over non-CR groups; final 5 hours of play	14
Table 2: Dataset for the evaluations.....	15
Table 3: Intent inference results for different information sources	16
Table 4: Probability of reference preceding a lift at time t	32