A Location Representation for Generating Descriptive Walking Directions

Gary Look, Buddhika Kottahachchi, Robert Laddaga, and Howard Shrobe
MIT Computer Science and Artificial Intelligence Laboratory
77 Massachusetts Ave, 32-224
Cambridge MA 02139
{garyl,buddhika,rladdaga,hes}@csail.mit.edu

ABSTRACT

An expressive representation for location is an important component in many applications. However, while many location-aware applications can reason about space at the level of coordinates and containment relationships, they have no way to express the semantics that define how a particular space is used. We present LAIR, an ontology that addresses this problem by modeling both the geographical relationships between spaces as well as the functional purpose of a given space. We describe how LAIR was used to create an application that produces walking directions comparable to those given by a person, and a pilot study that evaluated the quality of these directions. We also describe how LAIR can be used to evaluate other intelligent user interfaces.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Help systems, User-centered design*

General Terms

Human Factors, Design

Keywords

Location Ontologies, Navigation Directions

1. INTRODUCTION

We present LAIR (Location Awareness Information Representation), a model of space that can be used to create location-based services. LAIR can be used to represent not only where a person is, but also what a person is near and what he can do at those nearby places. LAIR incorporates concepts that people commonly use when thinking about space. Current representations model either the physical relationships between different spaces or the functional purpose of a given space. LAIR models both of these aspects, and this allows us to build applications that make fuller use of knowledge about a person's current location.

To demonstrate this, we describe how LAIR was used to build the Stata Walking Guide. This is an application that generates walking directions that are similar to those a person would give. LAIR can be used to do more than just describe how to get from one place to another, however. We also describe a tool we created that uses LAIR to provide additional information about the place you are in. This tool allows a person to expand his cognitive map of his current location and help situate himself with respect to other places.

In the remainder of this paper, we describe LAIR and compare it to other location representations. We then describe the Stata Walking Guide and an evaluation of the directions it produces. We follow with a discussion of how LAIR can be used to answer questions about the route the Stata Walking Guide produced. We conclude with a discussion of future work and a summary.

IUI'05, January 10–13, 2005, San Diego, California, USA.

2. RELATED WORK

As the field of ubiquitous computing has matured, the need for a common and expressive representation for location is becoming more apparent. Both instrumented environments and location-based services running on mobile devices can better address user needs by having a way of inferring how a person's presence in a particular location influences what he does at that place [9]. One example of this is a tourist application running on a visitors' kiosk that uses information about your interests to suggest a travel itinerary. This information could be supplied by a software agent running on your cell phone that has been programmed to cooperate with the certified kiosk. Another example is a kiosk network in which each kiosk uses knowledge about the area in which it is installed to determine what information to display [23].

The application-specific nature of many location representations limits knowledge sharing to other applications with highly similar location requirements. A better approach to supporting knowledge sharing between distributed systems is the one taken by the Semantic Web [2]. The Semantic Web is becoming the standard means by which information is shared because it provides a framework in which ontologies – formal descriptions of the concepts and relationships in a particular domain – can be defined.

There are a number of different languages for writing ontologies for the Semantic Web. DAML [10], the DARPA

Agent Markup Language, and OWL [19], the W3C's Web Ontology Language, are two examples. Using ontology languages such as DAML and OWL, a number of ontologies to describe location, such as OpenCyc [16] and OpenGIS [21], have been written. There have also been efforts such as SOUPA [3] that have attempted to combine many of these different ontologies to create a "best-of-breed" ontology for building ubiquitous computing applications.

However, these ontologies only capture geographical and geopolitical properties such as spatial containment, distances, and latitude and longitude. The canonical use of these ontologies is to make inferences based on spatial containment in order to describe the location of a person at different levels of granularity [3]. However, there aren't any applications that use these location ontologies to make other sorts of inferences. What is missing are the means to express what activities are carried out at a given set of geographic coordinates. This functionality is what LAIR provides, and it allows for applications to infer activity from location.

As Hightower notes [9], there are services such as Mapquest that use location to conduct "Nearest X" searches (for example, find the restaurants or hotels within a certain distance of a given address). These services, however, are based on yellow-page type listings, which only associate a label to a place but not any sort of semantic meaning on possible activities that occur at that place.

In addition to Mapquest, there have been other projects that have addressed the problem of generating route instructions. The project most similar to the Stata Walking Guide is CORAL [4]. CORAL uses natural language generation techniques to produce written instructions similar in content and form to those a person would give. CORAL takes data from geographical databases and uses it in a bottom-up approach to produce instructions. The Stata Walking Guide also uses bottom-up techniques but it is also able to use LAIR's ontological structure in performing top-down reasoning on how route instructions should be segmented for presentation.

3. LAIR

LAIR is an ontology inspired by Ben Kuipers' TOUR model of a person's cognitive maps of large-scale spaces [13]. The TOUR model is based not just on metric distances but on higher-level concepts such as places and paths between places. This allows us to model the topological and geographical relationships between different places. As in the TOUR model, LAIR represents a location in the real world by a Place and streets and pathways with a Path. LAIR supplements the TOUR model by associating to a Place a description of its functional purpose. For example, a particular building may function as any number of the following: grocery store, bank, or shoe store. A certain area of a building may be a meeting area, lounge, or kitchenette. These descriptions are represented in LAIR by a Functional Place. Instances of locations modelled in LAIR are stored in a semantic network [22]. The semantic network allows us to make inferences about the relationships between different places and the paths between them.

We now describe some of the concepts that LAIR models. Due to space limitations, we limit the discussion to the concepts that are most relevant to the development of the Stata Walking Guide. A complete description of LAIR can be found elsewhere [12].

Places in LAIR have the following properties:

Name. A way to refer to this Place.

On. A list of the Paths this Place is on.

Star. A list of triples (Path, heading, Path direction), that describe the geometry of the intersection formed by the Paths that meet at this Place. The value for heading ranges from 0 to 360. The zero mark for each Place is arbitrary; the Stata Walking Guide sets the zero mark for each Place to be cardinal north. The value for path direction is either +1 or -1, and indicates the direction of travel along the Path if a person were to travel from this Place along the given heading. Path direction will be discussed shortly.

View. A list of triples (Place, heading, distance) describing the other Places that can be observed from this Place. The means by which Places can be observed are not limited to visual detection [14], but for the Stata Walking Guide, we assume a line of sight between a Place and the other Places in its list of Views.

Contained. An unordered list of Places. The geographical extent of a Place is not limited by size, and a Place p may be geographically subsumed in any number of other Places. The Places that subsume p are in p's Contained list. Contained is an unordered list because there may not necessarily be any sort of strict containment ordering between the Places listed in Contained. For example, the state of Connecticut is contained within both New England and the "Tri-State Area." However, since only parts of the Tri-State Area are in New England (namely, Connecticut) there is no strict ordering between these two Places. Of course, for those relationships with a strict containment hierarchy, it is sufficient to only list the immediately subsuming Place.

Function. A list of the Functional Places that describe what can be done at this place and how this place is used.

Paths are one dimensional abstractions for the streets, roads, and hallways that are used to travel from one Place to another. Travel along a Path can be made in either the +1 direction or in the -1 direction. Determining which direction is +1 is arbitrary. Paths in LAIR have the following properties:

Name. A way to refer to this Path.

Row. A list where each element is an ordered list of Places. Each ordered list contains a sequence of Places that would be encountered when travelling along the Path in the +1 direction. To support incomplete knowledge of all Places along the path, Row is a list of these ordered lists.

Functional Places have a name and a list of Actions that can be performed at that Functional Place. The Stata Walking Guide uses the name of a Functional Place to describe waypoints along the route it produces. The list of Actions can also be used to manage access control to the Functional Place [11].

An abbreviated example showing how LAIR can be used to model an area of Cambridge, Massachusetts is shown in Figure 1.

4. THE STATA WALKING GUIDE

To demonstrate how LAIR can be used in an application, we created the Stata Walking Guide. The Stata Walking Guide provides both a route map and written walking directions between different places in MIT's Stata Center. This building is home to MIT's Computer Science and AI Lab

PLACE1:
Name: MIT
On: { Massachusetts Avenue (Mass Ave), Memorial Drive }
Star: { (Mass Ave, 315, +1), (Mass Ave, 165, -1), (Memorial Drive, 65, +1), (Memorial Drive, 245, -1) }
View: { (John Hancock Tower, 135, 1 mile) }
Contained: { Cambridge, Massachusetts, New England }
Functions: { Teaching Institute, Research Institute }

PLACE2:
Name: Central Square Transit Point
On: { Mass Ave, Prospect Street, Western Avenue, River Street }
Star: { (Mass Ave, 135, -1), (Mass Ave, 315, +1), (Prospect Street, 30, +1), (Western Avenue, 270, -1), (River Street, 225, +1) }
Contained: { Central Square, Cambridge, Massachusetts, New England }

PLACE3:
Name: Harvard Yard
On: { Mass Ave, John F. Kennedy Street }
Star: { (Mass Ave, 160, -1), (Mass Ave, 350, +1), (John F. Kennedy Street, 190, +1) }
Contained: { Harvard Square, Cambridge, Massachusetts, New England }

PATH1:
Name: Massachusetts Avenue
Row: { {MIT, Central Square Transit Point, Harvard Yard} }

FUNCTIONAL PLACE1:
Name: Teaching Institute
Actions: { Lecturing }

Figure 1: An abbreviated example of modelling an area of Cambridge, Massachusetts using the LAIR ontology

(CSAIL). This application is especially useful for the Stata Center because many visitors (and current occupants!) find the building's irregular floorplan bewildering (see Figures 2, 3, and 4 for actual floorplans.) The written directions produced from our representation also include landmarks. These directions are in contrast to the type of directions produced by Mapquest, which only list a sequence of "go-to" and "turn" instructions. The written directions supplement information about the route represented on the map and provide further information to help a person develop his own cognitive map of the Stata Center.

4.1 Content of Written Directions

There have been many studies in cognitive psychology [1, 5, 8, 18, 20], that have looked into what characteristics good route directions have. The results are consistent with expectations: good directions should be presented in the correct temporal-spatial order, places where turns are required should be clearly identified, and the instructions should include some indication that a person is travelling in the right direction.

However, since most of the aforementioned studies took place in outdoor or underground environments [7], as a first step in designing the Stata Walking Guide, we wanted an understanding of the vocabulary people use when giving written directions in an indoor environment. To this end, we collected a corpus of written directions from current Stata Center occupants. This approach is different from other studies of route knowledge [5, 18] which generate corpora by collecting verbal directions that are later transcribed. Since we were interested in generating written directions, we wanted the corpus we analyzed to reflect this.

We collected our corpus by asking 10 people (7 men and 3 women) to provide 3 sets of directions each. These 10 people work on different floors of the Stata Center, and as such, we could not assume they all had the same knowledge about the layout of the Stata Center. So, instead of asking our volunteers to provide directions for the same three routes, we asked each volunteer to provide 3 route descriptions where the first route remained entirely on one floor, the second started on one floor and ended on an adjacent floor, and the third started on one floor and ended on a floor two or more floors away from the starting floor. We thus collected 30 route descriptions in all, with each route being distinct from all the others. We regarded the diverse knowledge from which these directions were drawn as a benefit because we were not interested in studying how different people described the same route. Rather, we wanted to see how routes through different parts of the Stata Center are described.

Our examination of the directions revealed a useful list of characteristics. As a result of this analysis, we decided the Stata Walking Guide should produce walking directions with the following properties:

- **Directions do not use metric distances or cardinal directions.** Metric distances and cardinal directions were rarely used in the directions in our corpus. Only one person mentioned a metric distance at all, and in this case, it was just to say, "Turn right after a few yards." The lack of mention of specific distances is replaced by descriptions of where the next turn should be made, for example "go to the end of the hall" or "go to this landmark". The lack of cardinal directions is understandable since it is easier to refer to and orient yourself with respect to places within the local environment, rather than the cardinal directions.

- **Directions are more complex than a sequence of "go-to" and "turn" directives.** "Go-to"-type directions are not always needed. Half of our volunteers described at least one route in which one turn direction was followed by another. For parts of routes in which turns are required in quick succession, it does not make sense to insert a "go-to" instruction because a person may walk past the place they need to turn.

- **Use landmarks to identify places to turn.** Consistent with results from cognitive psychology, we found that all of our volunteers used landmarks to identify points in a route where a change of direction had to be made. While landmarks were not used in every route, they were used to identify places to turn in 23 of the 30 routes.

- **Use landmarks to verify travel in the right direction.** In addition to identifying places to turn, landmarks are also used to indicate that a person is traveling in the right direction. Eight of our ten subjects used landmarks in this manner, and half of all route directions mentioned a landmark as a waypoint to confirm that a person was traveling in the right direction. This is in line with Lovelace's results [18] for routes that had a number of turns and path segments.

- **Describe the physical spaces a route passes by, or passes through.** Nine of our ten volunteers used descriptions of areas as landmarks when describing a route. There were 20 route descriptions which included some description of a space as a landmark ("You are now in a small open area with some large printers on your right", "[go] straight across the common area"), and in 16 of these route descriptions, the functional purpose of the place was used ("lounge", "tea kitchen").

- **Doors are useful landmarks.** From our corpus of 30 routes, 23 of them required passing through at least one doorway. When describing these routes, our volunteers mentioned a large majority of the doorways (45 out of 50) a person had to pass through.

- **Describe hallway intersections.** The description of at least one hallway was included in eleven of the route directions we collected. The descriptions included phrases such as "head down the narrow hallway" and "follow the hallway until it comes to a T."

4.2 System Architecture

The Stata Walking Guide consists of two parts, a route-finding component that produces the graphical route map between two places and a translation component that converts the route information into written directions. The route-finding component uses A^* search to produce a series of waypoints that represents a route from one place to another. This route-finding component was an existing piece of software, written before LAIR, so the A^* search is done over a simple coordinate system. The resulting path is represented as a sequence of vertices in this coordinate system. We included the graph used by the route-finding component into our LAIR representation of the Stata Center, but instead of rewriting the route-finding component to search over the LAIR representation, we decided to use the LAIR representation to infer the written directions from just the sequence of vertices produced by the route-finding component. Being able to infer descriptive written directions from this lower-level abstraction illustrates the richness of our LAIR representation.

5. GENERATING WRITTEN DIRECTIONS

5.1 Describing Turns and Landmarks

To generate written directions from a sequence of waypoints, we first group the waypoints into sets, with each set representing part of a LAIR Path along the route. The points that are in the intersection of two sets represent LAIR Places where two or more Paths intersect in reality. Since we model the geometry of intersections and we know the previous and the next waypoint on the route, it is straight-forward to determine where to turn and in which direction.

To describe where to turn, we use a number of different rules. The first rule uses the fact that sometimes, a particular path will come to an end, forcing a person to turn one way or another. Since we model where paths begin and end, it is straight-forward to generate statements such as "turn right at the end of the hallway," when appropriate. The second rule includes landmarks in the description of where to turn. At these decision points we insert a functional description of the area that a person has entered, for example, "when you enter the lobby, turn left."

When describing landmarks that identify places in the route to turn, the Stata Walking Guide makes use of a landmark's visibility to condense the directions produced. It does this by starting at the landmark, l and going backwards along the route to find the longest sequence of waypoints (LAIR Places) that have l in their View. Any turns along this sequence of waypoints are not mentioned, and instead are replaced by the phrase "You will see l on your right (or left, as the case may be); walk towards it."

Landmarks are also used to confirm that a person is headed in the right direction. For example, the Stata Walking Guide could generate a directive like "walk down the hall, pass the copy room on your right..." The translation component inserts a landmark if the length of a particular path on a route, as measured by the number of waypoints provided by the route-finding component, is above a particular threshold. For these long paths, the translation component searches for landmarks along the path and inserts the landmark that is nearest the midpoint of the path segment.

The generated directions also indicate where on the route a person walks through doors. Places that are on either side of a doorway are labeled as such, and when a route takes a person from one side of a doorway to the other, a "Walk through the doors" statement is included in the directions. If there is a description of the Functional Place that the person has entered, this is also included in the directions.

5.2 Grouping the Directions

The last part of the generation process groups directions into segments to make it easier for a person to understand the directions as a whole. The grouping is a recursive procedure that counts how many instructions are presented in a given segment. If the number is above a certain threshold, the instructions are split into two groups, with the starting and ending points in separate groups. Points along the route (and their corresponding directions) are grouped together based not on metric distances, but rather on geographic similarities, as modeled in LAIR's ontological structure. For example, depending on the length of a route, two waypoints may be grouped together if they are on the same floor of a building. A third grouping may be required for intermediate points that differ greatly from both the start and end points.

For example, consider a route from my office to the CSAIL front office. Figure 2 shows the maps and walking directions the Stata Walking Guide produces for this route. In LAIR, office 221 is contained in the AIRE research neighborhood, which is contained in the second floor, which is contained in the Stata Center (we denote this by writing 221 $<$ AIRE neighborhood $<$ Second floor $<$ Stata Center). CSAIL's front office is located in CSAIL Headquarters on the fourth floor of the Stata Center's Gates Tower (Front office $<$ CSAIL Headquarters $<$ Fourth floor $<$ Gates Tower $<$ Stata Center).

In our representation, the most significant geographic difference between 221 and CSAIL's front office is that 221 is located on the second floor whereas the front office is located in the Gates Tower. This difference forms the basis of our first split. This procedure is then applied recursively to the resulting two sets of directions and terminates when either the number of directions is below threshold or when the start and end points of a segment differ only at the level of spatial containment immediately one level above them. While the recursion is not very deep for routes within the Stata Center (usually, only one segmentation is needed for routes between floors, and none for intra-floor routes), this algorithm can segment directions for longer routes, such as those used for coast-to-coast road trips.

This type of segregation into distinct areas can also be

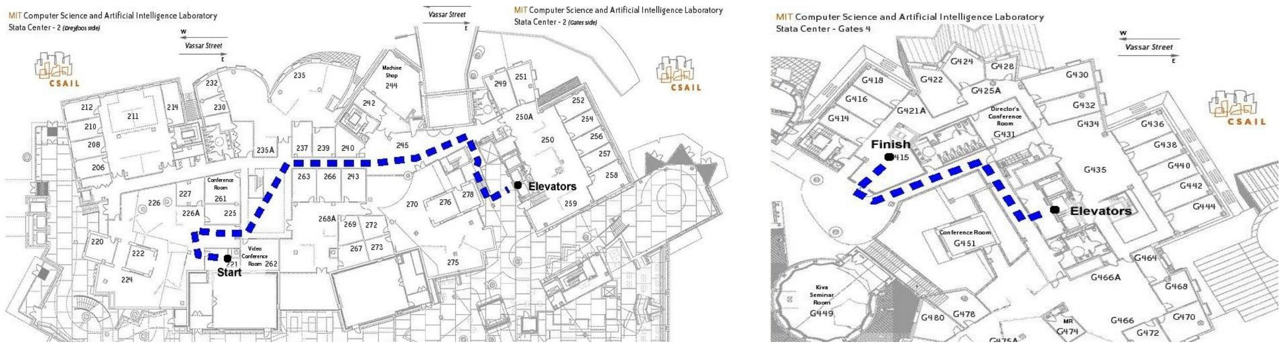

Head out of 221 and
Turn right.
You will see a set of double doors on your right; walk towards it.
Walk through the doorway into the Lounge.
Turn left at Lounge.
Walk into the main hallway.
Turn right at main hallway.
Walk forward.
Turn right at near (sic) the end of the hall.
Walk into the Elevator Lobby.
You will see elevators on your left; walk towards it.

Take the elevator to the fourth floor.

Head out of Elevator Lobby and
Turn right.
Turn left.
Walk forward.
Turn right.
Walk through the doorway into a new area (CSAIL_Headquarters).

Walk out of room 221
Turn right and right again, facing the AIRE doors
Walk through the AIRE doors into the lounge
Take a left through the lounge
Take a right into the main corridor
Walk down the main corridor past the MERS (robot) area
Take a right and go through the glass door and down the steps to the 1st floor elevator lobby
Take an elevator to the fourth floor
Exit the 4th floor elevator lobby and turn right
Turn left at the end of the corridor and enter the R&D Lounge
Turn right and enter CSAIL HQ.

Figure 2: An example of a route (top) and written directions (lower left) produced by the Stata Walking Guide. A person's description of the same route is shown in the lower right. The route shown was the first route used in our pilot study.

used to produce higher-level directions that abstract away the individual turn-by-turn instructions. These higher-level directions could be used by current occupants with some knowledge of the Stata Center.

6. EVALUATION

We ran a pilot study to evaluate the subjective quality of the directions the Stata Walking Guide produces. In this study, seven volunteers, three women and four men, each of whom were not familiar with the Stata Center, were given a set of three route descriptions to follow. The first route was chosen because it spanned multiple floors. The second and third routes were chosen to represent within-floor routes. Each route description included a statement of the starting and ending points, a map of the building with the route highlighted, and written directions generated by the Stata Walking Guide. Participants were told that the purpose of the study was to identify characteristics of good directions. They did not know the written directions they were following were generated by a computer program. Figures 2, 3, and 4 show the route maps, the directions generated by the Stata Walking Guide, and the same route described by a person.

The amount of overlap between routes was limited to only two instances where the routes briefly overlapped portions of hallways. The routes did not share any common turns. Since the goal of the pilot study was to assess the subjective quality of the directions, and not the effect learning had on being able to get to the destination, we did not consider the small amount of overlap between routes to be an issue.

Each of the routes began and ended in a different place. Participants were taken to each of the starting points via routes that did not intersect the route they were to follow. An observer followed each participant as he walked the route. If a participant made a wrong turn and did not realize this error before wandering 15 feet from the correct choice point, they were told of the error and reoriented at the correct choice point. After walking all three routes, participants filled out a questionnaire where they were first asked to subjectively rate the quality of each set of directions on a scale of 1 to 5, with 5 being the highest. They were then asked to compare the directions they received to human generated descriptions of the same routes. For each of the three routes, participants were asked whether the set of directions they followed were worse than, comparable to, or better than the second set of directions they were shown. Figures 2, 3, and 4 show the directions provided by people along with the ones the Stata Walking Guide produced.

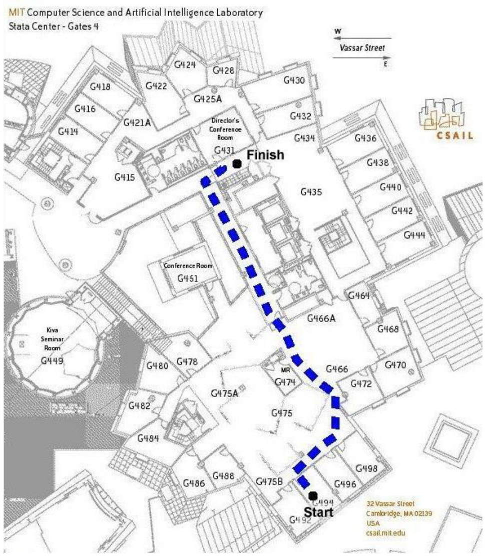

Head out of G494 and
Turn right.
You will see a set of double doors on your left; walk towards it.
Walk through the doorway into a new area (4th_Floor_Gates).
Walk forward.
You will pass the Elevator Lobby on your right along the way.
Turn right.
Walk through the doorway into a new area (4G_Kitchenette).

Turn right, then left, and go straight down the hall through the doors and past the elevators. Go through the first door on your right.

Figure 3: The second route used in our pilot study. Below the route map are the directions the Stata Walking Guide produced. A person's description of the same route follows below that.

The results of the pilot study are both promising and give insight on the limitations of the Stata Walking Guide. Of the three routes, the directions the Stata Walking Guide produced for Route 1 received the lowest mean score, 3.5; this set of directions was unanimously rated as worse than the human provided directions. Routes 2 and 3 received mean scores of 4.0 and 4.4, respectively. For Route 2, 6 of the 7 participants rated the directions produced by the Stata Walking Guide as better than those produced by a person, while the other person rated them as worse. For Route 3, 5 of the participants rated the two sets of directions as comparable. One participant considered the Stata Walking Guide's directions to be better than the directions produced by a person while the last participant considered them worse.

The high mean scores and favorable comparisons to the human directions in Routes 2 and 3 are evidence that the Stata Walking Guide can use LAIR to produce a serviceable set of directions. While Routes 2 and 3 were shorter than Route 1, the longer route distance was not the main reason why Route 1 received lower ratings. One-on-one interviews with the study participants revealed a consistent explanation.

In Route 1, the route takes a person from the second floor of the Stata Center to the fourth. In our LAIR model of that part of the fourth floor, many of the Places on the route did not have any landmarks listed in their View property. As a result, the Stata Walking Guide was limited to describing that part of the route as a sequence of turns, along with a mention of the doorway entrance leading into CSAIL headquarters. This limitation alone would not have been so bad, but the difficulty this caused was compounded by the fact that the floorplan opens up at the same place where the last "Turn right" instruction occurs (see Figure 2). When participants arrived at this waypoint, they had the tendency to veer slightly to the right, due to both the floorplan and the "Turn right" instruction. As a result, they would walk right past the entrance to the CSAIL main office, which was, technically, on their right.

What Route 1 illustrates is the importance of having landmarks and describing sharp turns. The LAIR model didn't have enough information at that point, and so while it did the best it could with the limited information available, the result was far from ideal. This lack of detail in the directions was the reason why the Stata Walking Guide's output was unanimously considered worse than what a person could produce. The ability to describe sharp turns is a feature that will be incorporated into the next version of the Stata Walking Guide.

One of the main reasons why the LAIR model of the Stata Center was lacking in that area of the fourth floor is that we currently create these LAIR models by hand. We discuss the problem of automatically creating rich LAIR representations in Section 8.

7. ISLE

To further explore what could be done with LAIR representations of space, we created ISLE, the Interactive Simulator for LAIR Exploration. Originally conceived as a debugging tool for the Stata Walking Guide, ISLE can be used to verify the underlying relationships in the semantic network used to represent LAIR. However, the sorts of debugging assistance that ISLE provides for checking the correctness of the underlying representation can also be used in a number of other ways. A person running ISLE on a mobile device could use ISLE to get more information about his current location. Recording what queries ISLE is given makes it a tool that can help answer questions about what a person is interested in and what may be confusing him.

The input to ISLE is the same sequence of waypoints that the Stata Walking Guide uses. Using LAIR, ISLE looks at the different relationships between the waypoints and other nearby places and allows a person to ask different variations of the same question: "Where am I?" The following is a list of questions that ISLE can answer:

1. "Describe this place." *Answer:* "You are in a lounge, which is in the AIRE research neighborhood."

2. "What can I do here?" *Answer:* "You are in a kitchenette where you can get coffee and microwave a meal."

3. "How will I know I'm going in the right direction?" *Answer:* "You will see a copy machine on your left, then a kitchenette on your left, and then the second floor lounge on your right."

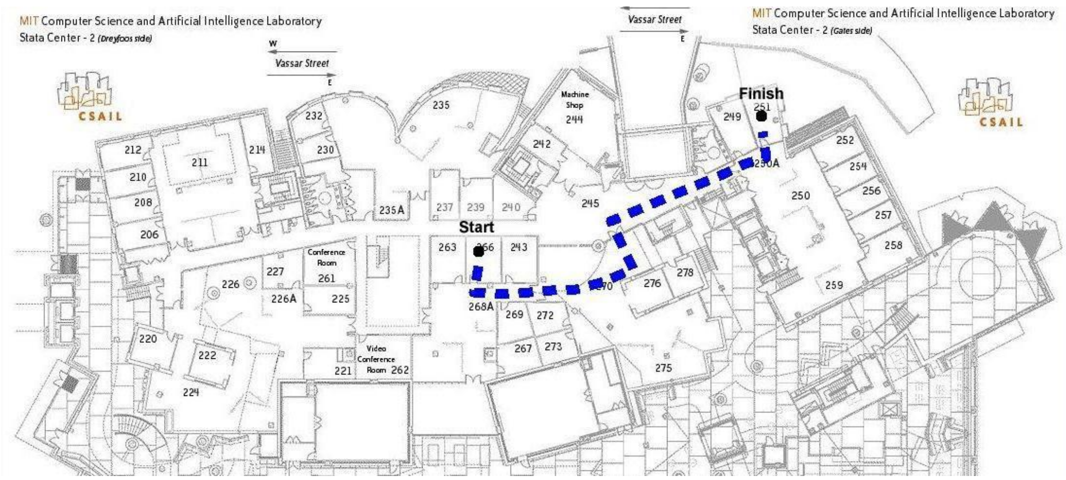

Head out of 266 and
Turn left.
Walk through the doorway into a new area (MERS).
You will see a set of double doors on your left;
walk towards it.
Walk into the main hallway.
Turn right at main hallway.
Walk forward.
Walk through the doorway into a new area
(Genesis_Group).
251 is on your left.

Exit my office (32-266) and take a left. Go through the first set of double doors and then the second set as well. Once in the hallway turn right. Walk straight towards, and go through the door just pass the kitchen. Patrick's office it the second door on the left.

Figure 4: The third route used in our pilot study. The directions produced by the Stata Walking Guide appear on the left. A person's description of the same route appears on the right.

4. "Is place X near place Y?"

5. "Is place P along my route from X to Y?"

It is worth noting that many, but not all, of these questions can be answered in other types of spatial representations. For example, question 1 is the canonical use of the ontologies described in Section 2. As in those applications, ISLE also follows the upward pointers to find and list the containing Places. The answers for questions 2 and 3 use LAIR's representation of Functional Place and the View property of a Place. For the sake of conciseness and to reduce confusion, the Stata Walking Guide does not provide all of this information in the directions it produces. ISLE could also use the View property to help a lost person get his bearings. By presenting a list of likely landmarks that a person could see, ISLE could do a coarse grain form of triangulation based on what landmarks a person actually does see.

The notion of "nearness" (such as that raised in question 4 and its implications for the related question 5) is a difficult one to answer. In a coordinate-based representation, nearness can be determined by comparing a metric distance against some threshold. However, there are other things in addition to metric distance that impact our perception of what is near [6, 17]. Among these factors are whether or not a destination is in the same subsuming area of a given starting point, if there is a line of sight between the two places, and if the route between the two places requires passing through doorways.

Location models that model topological relationships can account for some of these factors but not all of them. These factors are represented in LAIR, and so ISLE can account for these factors when considering how near a place X is to Y. We are trying different ways of using ISLE to address the nearness problem, and one method that we are experimenting with is to take a weighted score consisting of factors that influence the perception of nearness and compare that against some threshold. The following factors go into computing the weighted score: the metric route distance between X and Y, the similarity of their locations (as measured by how many levels up in the spatial containment hierarchy we need to go to find a common Place that both X and Y are in), whether there is a line of sight between X and Y, and how many doorways a person needs to cross along the route from X to Y.

8. FUTURE WORK

A formal evaluation of the quality of directions produced by the Stata Walking Guide is underway. In this study, we expand on the work started in the pilot study by comparing how a group presented with directions produced by the Stata Walking Guide rate those directions against how a group presented with human produced directions rank their directions.

There are two other areas of future work that we are pursuing in the near term. The first is being able to automatically define LAIRPlaces and Paths, from architectural CAD drawings [15]. Currently, most of the LAIR model is created by hand. While we do have scripts to generate much of the data related to Places and Paths, these entities, and Place properties such as View and Contained relationships, still have to be manually defined. The second area of future work is to deploy the Stata Walking Guide to OK-Net, the Stata Center's network of public information kiosks [23].

A longer term area of future work is to use LAIR to develop tools that would provide a person with information that is relevant to his current location, current task, and to other things that are important to him. There are many scenarios in which having this sort of information would greatly support a person's dynamic decision making process. These situations range from everyday ones in which a person decides to stop by the post office because it is on his way to the grocery store to more critical situations such as an ambulance driver deciding to take a patient to a hospital that is not the closest one to his current location because that is the one best equipped to handle this particular trauma.

9. SUMMARY

This paper makes three contributions. First, it argues why location representations that only model geographical relationships restrict the degree of intelligent interaction applications can have with users. Second, it describes LAIR, an ontology that addresses this shortcoming by representing both the geographic relationships between spaces, as well as the functional purpose of a particular space. Third, it presents a pair of applications, the Stata Walking Guide and ISLE. The Stata Walking Guide demonstrates how LAIR can be used to create an application that address a problem domain requiring intelligible user output (walking directions). A pilot study provides preliminary support to the quality of the directions the Stata Walking Guide produces. ISLE shows that LAIR can be used, not only to create applications that require intelligent user interfaces, but that LAIR can also be used to create tools to assess intelligent user interfaces.

10. ACKNOWLEDGEMENTS

This research was funded by MIT's Project Oxygen. Special thanks go to Albert Huang for developing the A* path finding component

11. REFERENCES

[1] G. L. Allen. Principles and practices for communicating route knowledge. *Applied Cognitive Psychology*, 14(4):333–359, July/August 2000.

[2] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, May 2001.

[3] H. Chen, F. Perich, T. Finin, and A. Joshi. SOUPA: Standard Ontology for Ubiquitous and Pervasive Applications. In *International Conference on Mobile and Ubiquitous Systems: Networking and Services*, Boston, MA, August 2004.

[4] R. Dale, S. Geldof, and J.-P. Prost. Coral: using natural language generation for navigational assistance. In *Proceedings of the twenty-sixth Australasian computer science conference on Conference in research and practice in information technology*, pages 35–44, 2003.

[5] M. Denis, F. Pazzaglia, C. Cornoldi, and L. Bertolo. Spatial discourse and navigation: An analysis of route directions in the city of Venice. *Applied Cognitive Psychology*, 13(2):145–174, 1999.

[6] M. Duckham and M. Worboys. Computational structure in three-valued nearness relations. In D. R. Montello, editor, *Proc. COSIT 2001*, volume 2205 of *Lecture Notes in Computer Science*, pages 76–91. Springer, September 2001.

[7] S. Fontaine and M. Denis. The production of route instructions in underground and urban environments. In C. Freksa and D. M. Mark, editors, *COSIT 1999*, volume 1661 of *Lecture Notes in Computer Science*, pages 83–94, Stade, Germany, August 1999. Springer.

[8] J. M. Golding, A. C. Graessner, and J. Hauslet. The process of answering direction-giving questions when someone is lost on a university campus: the role of pragmatics. *Applied Cognitive Psychology*, 10:23–39, 1996.

[9] J. Hightower. From position to place. In *Proceedings of The 2003 Workshop on Location-Aware Computing*, pages 10–12, October 2003. part of the 2003 Ubiquitous Computing Conference.

[10] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. Reviewing the design of DAML+OIL: An ontology language for the semantic web. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI 2002)*, Edmonton, Alberta, 2002.

[11] B. Kottahachchi and R. Laddaga. Access controls for intelligent environments. In *Proceedings of the Conference on Intelligent Systems Design and Applications, (ISDA 2004)*, Budapest, Hungary, August 2004.

[12] B. Kottahachchi and G. Look. LAIR - Location Awareness Information Representation. Submitted to Percom 2004.

[13] B. J. Kuipers. *Representing Knowledge of Large-Scale Space*. PhD thesis, Massachusetts Institute of Technology, July 1977. MIT AI/TR 418.

[14] B. J. Kuipers and T. S. Levitt. The spatial semantic hierarchy. *Artificial Intelligence*, 119:191–233, 2000.

[15] V. Y. Kulikov. Building model generation project: Generating a model of the MIT campus terrain. Master's thesis, Massachusetts Institute of Technology, May 2004.

[16] D. Lenat. Cyc: A large-scale investment in knowledge infrastructure. In *Communications of the ACM*, volume 38. ACM, November 1995.

[17] J. Lin, R. Laddaga, and H. Naito. Personal location agent for communicating entities (PLACE). *Interacting with Computers*, 15(4):559–576, August 2003.

[18] K. L. Lovelace, M. Hegarty, and D. R. Montello. Elements of good route directions in familiar and unfamiliar environments. In C. Freksa and D. M. Mark, editors, *COSIT 1999*, volume 1661 of *Lecture Notes in Computer Science*, pages 65–82, Stade, Germany, August 1999. Springer.

[19] D. L. McGuinness and F. van Harmelen. OWL web ontology language overview. Technical report, World Wide Web Consortium, February 2004. http://www.w3.org/TR/2004/REC-owl-features-20040210/.

[20] P.-E. Michon and M. Denis. When and why are visual landmarks used in giving directions? In D. R. Montello, editor, *COSIT 2001*, volume 2205 of *Lecture Notes in Computer Science*, pages 292–305, Morro Bay, California, September 2001. Springer.

[21] Open GIS Consortium. http://www.opengis.org/.

[22] S. Peters and H. Shrobe. Using semantic networks for knowledge representation in an intelligent environment. In *PerCom '03: 1st Annual IEEE International Conference on Pervasive Computing and Communications*, Ft. Worth, TX, USA, March 2003. IEEE.

[23] M. Van Kleek. Intelligent environments for informal public spaces: the Ki/o kiosk platform. Masters of Engineering Thesis, Massachusetts Institute of Technology, February 2003.