Available online at www.sciencedirect.com
ScienceDirect
Int. J. Human-Computer Studies 67 (2009) 269–278
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

A tangible user interface for assessing cognitive mapping ability

Ehud Sharlin[a], Benjamin Watson[b,*], Steve Sutphen[c], Lili Liu[d], Robert Lederer[e], John Frazer[f]

[a]*Department of Computer Science, University of Calgary, AB, Canada*
[b]*Department of Computer Science, North Carolina State University, NC, USA*
[c]*Department of Computing Science, University of Alberta, AB, Canada*
[d]*Department of Occupational Therapy, University of Alberta, AB, Canada*
[e]*Department of Art and Design, University of Alberta, AB, Canada*
[f]*School of Design, Queensland University of Technology, Australia*

Received 1 May 2008; received in revised form 18 September 2008; accepted 21 September 2008
Available online 11 October 2008

Abstract

Wayfinding, the ability to recall the environment and navigate through it, is an essential cognitive skill relied upon almost every day in a person's life. A crucial component of wayfinding is the construction of cognitive maps, mental representations of the environments through which a person travels. Age, disease or injury can severely affect cognitive mapping, making assessment of this basic survival skill particularly important to clinicians and therapists. Cognitive mapping has also been the focus of decades of basic research by cognitive psychologists. Both communities have evolved a number of techniques for assessing cognitive mapping ability. We present the Cognitive Map Probe (CMP), a new computerized tool for assessment of cognitive mapping ability that increases consistency and promises improvements in flexibility, accessibility, sensitivity and control. The CMP uses a tangible user interface that affords spatial manipulation. We describe the design of the CMP, and find that it is sensitive to factors known to affect cognitive mapping performance in extensive experimental testing.

Keywords: Cognitive maps; Wayfinding; Cognitive assessment; Neuropsychological assessment; Tangible user interfaces; Constructional ability; Spatial ability

1. Introduction

Almost every day, people find their way from home to any of a myriad of destinations, and then back again. Most take this skill for granted, but is an amazingly complex ability that has been the subject of decades of research by cognitive psychologists, who call it *wayfinding*. Injury or disease can so impair this ability that many become homebound, and for some unfortunate people, catastrophic failure of their wayfinding ability has lead to death from exposure. Thus medical researchers and clinicians also have a very strong interest in wayfinding.

A crucial component of wayfinding ability is cognitive mapping. A *cognitive map* is a mental representation of a person's environment, relied upon during wayfinding. Researchers have developed many techniques over the years for measuring and assessing this ability. Map drawing or placement is quite common, but is difficult to score consistently, wholly two-dimensional (2D) and necessarily quite abstract in representation. A few researchers have assessed cognitive mapping by asking patients or study participants to arrange three-dimensional (3D) objects representing elements of their environment. The neuropsychological assessment literature (Lezak, 1995) suggests that moving from abstract to concrete, and from 2D to 3D representation will be helpful in increasing assessment sensitivity. Unfortunately, previous manual implementations of this approach were quite unwieldy and difficult to score, generally requiring a very trained and alert assessor (Lezak, 1995).

*Corresponding author. Tel.: +1 919 513 0325; fax: +1 919 515 7896.
E-mail addresses: ehud@cpsc.ucalgary.ca (E. Sharlin), bwatson@ncsu.edu (B. Watson), steve@cs.ualberta.ca (S. Sutphen), lili.liu@ualberta.ca (L. Liu), rlederer@ualberta.ca (R. Lederer), j.frazer@qut.edu.au (J. Frazer).

1071-5819/$ - see front matter
doi:10.1016/j.ijhcs.2008.09.014

To address these problems in assessment, we have designed the Cognitive Map Probe (CMP), an automated tool for the measurement of cognitive mapping ability. The CMP makes use of the Segal model (Frazer 1982, 1995), a tabletop tangible user interface (TUI) originally designed for the input of architectural models. CMP users view a drivethrough of a neighborhood on a large screen perspective display (Fig. 1), and then input their recollection of that neighborhood by arranging 3D building models on the Segal model's tabletop input surface (Fig. 2). The CMP automatically records and scores each change the user makes to the model configuration. The CMP is the first TUI for assessment of cognitive mapping ability, combining the increased sensitivity of 3D input and affordances with the improved consistency, efficiency, flexibility and high-resolution data collection of computerization.

We begin the remainder of this paper with a review of cognitive maps, including their importance in everyday life and their measurement. A detailed description of the CMP follows, including comparisons to related TUIs. We conclude with a rigorous experimental examination of the sensitivity of the CMP to age and task difficulty, two factors that have a well-known relationship to cognitive mapping performance.

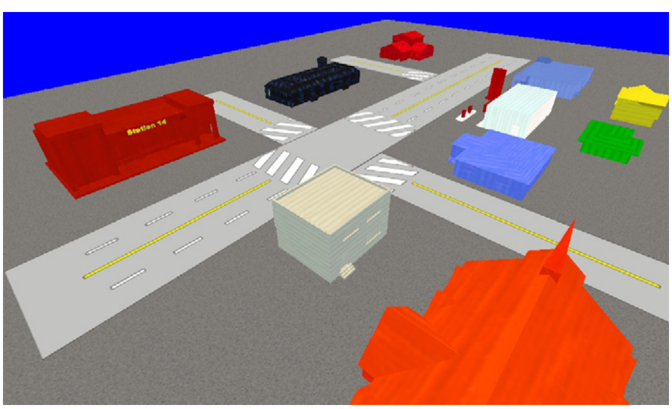

Fig. 1. Virtual neighborhood (exocentric view used only for illustration).

Fig. 2. A participant interacting with the CMP.

2. Measuring cognitive maps

In his pioneering paper, Tolman (1948) argues that rats, like humans, have a mental representation of the world he called a cognitive map. These maps hold detailed spatial information that individuals collect, integrate and use while interacting with the environment. Tolman's work has led to the modern psychological definition of a cognitive map: *an overall mental image or representation of the space and layout of a setting* (Arthur and Passini, 1992).

It is important to distinguish between the psychological concepts of wayfinding and of cognitive maps. Wayfinding refers to the overall process of reaching a destination (Darken and Peterson, 2002; Bowman et al., 2005), while cognitive maps underlie the wayfinding process and enable making and executing decisions about the environment.

The most widely accepted theory of cognitive mapping is the Landmark-Routes-Survey (LRS) model (Darken and Peterson, 2002). The model divides our environmental understanding into three levels—landmark, route and survey—that can be integrated into a single comprehensive cognitive map (Golledge, 1991; Colle and Reid, 1998).

Cognitive maps can often be imprecise. We tend to classify and cluster the massively detailed cognitive spatial information we encounter using simplifications, such as the gathering of objects and landmarks into hierarchies and regions. Cognitive maps also suffer from geometrical scaling and regularization problems (Golledge, 1991).

We acquire cognitive maps through at least two types of environmental interaction: direct physical interaction, for example, by looking around when walking down a street, or by tapping with a cane in case of a visually impaired person; or mediated interaction, for example through maps and virtual environments (VEs).

A number of cognitive- and task-related factors affect cognitive mapping ability, including age, task difficulty and dementia. In fact, dementia such as Alzheimer's disease (AD) affects wayfinding to such an extent that officials typically waive the waiting period for acting on missing person reports for diagnosed dementia patients, who have died from exposure when they become lost and disoriented. Some have proposed assessment of cognitive mapping ability as a form of AD diagnosis (Liu et al., 1991).

2.1. Cognitive maps in virtual environments

Many researchers have explored techniques for improving wayfinding aids in VEs (e.g. Bowman et al., 2005; Sadeghian et al., 2006; Smith and Hart, 2006; Cliburn et al., 2007), giving special attention to VEs as wayfinding training tools (Koh et al., 1999; Darken and Peterson, 2002). The technology that supports these VE applications is diverse and ranges from low-end desktop PCs to CAVEs employing treadmills as travel interfaces (Stanney, 2002).

A concern that overshadows VE-based wayfinding trainers is the problem of training transfer, that is determining whether the cognitive map acquired in the VE is useful in the physical world. Currently there is no clear-cut answer to these questions (Darken and Allard, 1999; Koh et al., 1999). In fact, previous research has shown that wayfinding training with VEs might actually hinder the development of practical survey knowledge (Darken and Goerger, 1999). While the quality of the resulting survey knowledge was questionable, the participant's navigation strategy in VE seems to be effective in predicting her real-world wayfinding ability (Darken and Goerger, 1999). Thus VEs may be more promising for wayfinding assessment than for training.

2.2. Techniques for probing cognitive maps

There are several techniques for assessing cognitive mapping ability. Verbal techniques simply ask a person to describe the environment. These techniques suffer from the subjective nature of the reported information and from natural variability in communicative ability. However, verbal techniques can achieve deep insight into cognitive mapping through use of verbs of motion rather than just dry description of physical locations (Axia et al., 1991).

The bearing and distance technique (Baird, 1979; Colle and Reid, 1998; Goerger et al., 1998; Koh et al., 1999) places a person at a certain location in the environment, and asks him or her to point to other objects in the environment and estimate the distances to them. These inter-object distances and directions are then compared to the distances in the original environment. The bearing and distance technique is easy to implement, but the technique suffers from scale problems and may not be very sensitive to survey knowledge enabling generation of new paths through the environment (Darken and Allard, 1999).

Map drawing (Darken and Allard, 1999) or placement (Piaget and Inhelder, 1956; Baird, 1979; Goerger et al., 1998) techniques ask a person to describe his or her cognitive map through sketching or by spatially placing physical objects. Drawing techniques are sensitive to variation in sketching ability. In work of particular relevance here Piaget's "Model Village" employed 3D objects for map placement, using cardboard models of a church, houses and trees to help children input cognitive maps (1956).

Functional assessment techniques position a person in a previously studied spatial environment and assess the person's ability to perform a novel navigation (Goerger et al., 1998). This technique can provide excellent insight into the user's survey knowledge, but requires considerable assessor time and can raise a psychological Heisenberg-like principle as the ability being measured can be altered by its measurement process (Darken and Allard, 1999).

The use of computers in cognitive mapping ability assessment is very limited. The first use dates to the late 1970s when Baird (1979) designed an automated map placement technique. Study participants located buildings on a 13×13 matrix, displayed on a monitor. Researchers have also collected bearing estimates using computers in an automated variant of the bearing and distance technique (Colle and Reid, 1998).

3. The Cognitive Map Probe

The CMP is an automated system for assessment of cognitive mapping ability. During the first phase of each trial, the participants view a virtual neighborhood displayed with a digital projector (Fig. 1). Viewing can be passive, similar to riding in a bus; or active and more akin to participants driving the bus. Viewing can also be egocentric, with participants seeing a street level view; or exocentric, with participants seeing a bird's eye view. This flexibility enables the CMP to accommodate participants with varying cognitive ability, including possible dementia. On the one hand, a completely passive drivethrough may hinder the development of participants' cognitive maps (Koh et al., 1999). On the other hand, an involved drivethrough interface may pose a challenge to some of the CMP's potential users (for example, elderly users), effectively preventing them from being able to learn the VE, and consequently causing them to fail the experiment's cognitive mapping probing phase.

In the trials reported here, viewing followed an egocentric and largely passive "bus ride" metaphor. The participant played a passenger who is allowed to start and stop the bus, and during stops, can slowly rotate her viewpoint through 360 horizontal degrees. Participants only requested a rotation, but did not control or manipulate it; the rotation itself had constant direction and velocity.

In the trial's second phase, participants move to a 2D input surface and tangibly construct their cognitive map of the neighborhood they have just visited (Fig. 2). Participants accomplish this by arranging physical, 3D models of buildings on the Segal model's 2D board. When participants place or remove buildings from the board, the system records the building ID, its 2D location and the time of the event. During placement, the system also records the building's orientation. When participants are satisfied that the constructed configuration accurately represents their cognitive map, they signal the assessor who advances the system to the next trial.

3.1. System and interface

We printed the CMP's user interface by creating 10 virtual building models in a software package, and then outputting them in 3D using rapid prototyping technology. The resulting polyester objects are quite sturdy and mounted on flat bases, under which is a single connector for the Segal model's board. Aligning the base with the board's slots aligns the connector to its matching slot and eases insertion of the model. The CMP limits the orientation input of the models to the four orthogonal angles only (0°, 90°, 180° and 270°). All the models are of similar scale and users can arrange them easily with two hands. We spray-painted the models with primary colors for easy viewing by the elderly, but hand-painted important details such as store signs in contrasting colors. The models have finely detailed shapes, and include doors, windows, and even the patterns of wood siding. We also attached a simple street pattern to the board (one four-way and one "T" intersection, see Figs. 1 and 2); the pattern remained there throughout assessment. All 10 models and the street pattern can fit onto the board at the same time.

The virtual versions of these physical models also populate the virtual neighborhoods shown to participants in the first phase of each assessment trial. Thus buildings in the displayed virtual neighborhoods match the physical models used for tangible interaction exactly in shape and nearly exactly in color.

The CMP uses the Segal model, a pioneering TUI named in memory of architect and advocate of home self-design Walter Segal. John Frazer and his colleagues built the Segal model (Frazer, 1982, 1995) in collaboration with Segal to support his work. They designed the device to enable direct, tangible interaction with architectural floor plans and their components, such as walls, doors, windows, plumbing fixtures and furniture. It is a 102 cm × 71 cm board covered with an array of 768 edge connector slots arranged in 24 columns of 16 vertical slots and 16 rows of 24 horizontal slots. Each slot has contacts enabling recognition of 127 different connector types, after accounting for symmetries in orientation. They represented architectural components with physical 3D models, each coupled to a unique connector type. Since our application required tangible, tabletop interaction very similar to that supported by the Segal model, we converted it for our use.

Early TUIs demonstrated the concept of coupling physical objects with digital information and function (Fitzmaurice et al., 1995; Ullmer and Ishii, 2001). Other efforts focused on TUIs as spatial interfaces (Sharlin et al., 2004), or explored their social design implications (Hornecker and Buur, 2006). A recent TUI effort had applied design goals not that different from the original Segal model, but was implemented using current technology (Hosokawa et al., 2008).

We designed the CMP carefully, matching its interface to our cognitive mapping assessment task. The design followed three heuristics: intuitive spatial mapping between interface and task, I/O unification and support for trial-and-error actions (outlined in Sharlin et al., 2004). A major consideration for our choice of the Segal model for the CMP implementation was its convenience and accessibility to us. However, we believe that the Segal model's limitations (e.g. its coarse sampling of location and orientation) likely acted as valuable input constraints, and may have actually benefited the CMP. Indeed, Sharlin et al. (2002) reported that TUI physical constraints that arguably hinder design flexibility can benefit cognitive assessment, simplifying the task and allowing the elderly and participants with mild dementia to participate. By embodying the assessment task in its simple physical representation, the CMP allows us to approach participants with limited skills and capabilities and automatically assess their cognitive mapping ability via intuitive tangible interaction.

3.2. Assessment measures

After assessment, the CMP analyzes the data it has collected to score the participant's performance. As we discussed above, one may score cognitive maps in a number of ways. Our measures (listed below) all originate from existing, manual methods for probing cognitive mapping ability, with the simpler measures (*number* and *difference*) used widely in both verbal (e.g., Axia et al., 1991) and survey-based protocols (e.g., Piaget and Inhelder, 1956). Our more elaborate measures (*distance*, *orient* and *interbuilding*) are directly based on the frequently employed bearing and distance techniques (e.g., Baird, 1979).

All of our measures involve comparisons of the actual map M to the participant's cognitive map C, and are clamped to the range [0,1], with 1 being the ideal result. Measures that disregard position and treat M and C only as sets of buildings are

$$Number = 1 - abs(|M| - |C|)/|M|$$
$$Difference = 1 - (|M - C| + |C - M|)/(|M| + |C|)$$

Here *number* compares the number of buildings $|M|$ and $|C|$ in the actual and cognitive maps, without ensuring that the maps contain the same buildings. In contrast, *difference* checks that the buildings in the actual and cognitive maps match by counting $|M{-}C|$, the number of buildings in actual but not in the cognitive map, and $|C{-}M|$, the number in the cognitive but not in the actual map. Measures that compare actual and cognitive maps' building positions within the set of matching buildings $M \cap C$ include:

$$Distance = 1 - \sum_i (dist(M_i, C_i)/d_{\max})/m_{\max}$$
$$Orient = 1 - \sum_i (odiff(M_i, C_i)/180)/m_{\max}$$
$$Interbuilding = 1 - \sum_i \sum_j (abs(DM_{ij} - DC_{ij})/d_{\max})/m_{\max}^2$$

where all sums range over the set $M \cap C$, $d_{\max}$ is the length of the CMP board diagonal and $m_{\max}$ is the maximum of $|M|$ in the entire assessment. *Distance* returns the average distance between matching buildings in the actual and cognitive maps, with *dist* referring to Euclidian distance. *Orient* measures the average difference in orientation between matching buildings in the actual and cognitive maps, with *odiff* measuring the angular difference between the orientations of two buildings in degrees. Note that *orient* is not sensitive at all to location, and that because of the Segal model's input constraints, the orientation difference between any one pair of matching buildings can only be 0°, 90° or 180°. *Interbuilding* compares the arrangements of the buildings in the actual and cognitive maps, first storing building-to-building distances within each map in square matrices D_M and D_C, then finding the average difference between corresponding entries in these matrices. D_M and D_C have entries $dist(M_i,M_j)$ and $dist(C_i,C_j)$, respectively (i and j again range over $M \cap C$).

Finally, the CMP forms a composite measure that responds to participant building, location and orientation matching abilities:

$$Similarity = difference \times distance \times orient$$

We use a product rather than a sum because we view these abilities as independent and describing a 3D ability space—and because the range of the corresponding measures is [0,1], a unit cube. *Similarity* therefore represents the proportion of this cube spanned by a participant's abilities. A 10% reduction in a component measure also reduces the composite measure by 10%, not 3.33%, as in an unweighted sum. We do not weigh the component measures non-uniformly, because we have no prior knowledge of the measures' independence.

Recall that the CMP also records the time of each action on the board. This allows us to add *totalTime*, the time it takes to complete one assessment trial, to our suite of measures. We can also probe the progress participants make during the assessment by comparing our measures to the current time. Fig. 3 graphs *similarity* vs. time for all participants in one assessment trial. We construct the additional measure *dSim* by finding the differences between consecutive measurements of *similarity* divided by the time elapsed between those measurements, and averaging the resulting "local slopes" over all such pairs in an assessment trial.

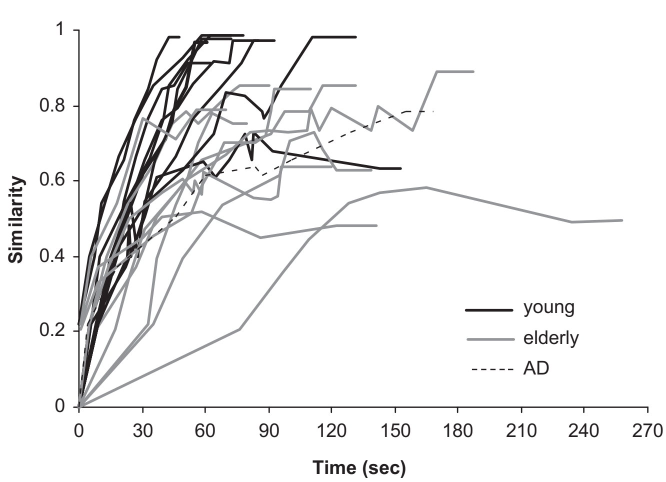

Fig. 3. *Similarity* vs. time (the most complex world (eight buildings), all participants).

3.3. System strengths

The CMP offers the following advantages over existing methods for assessing cognitive mapping skill.

Sensitivity: The CMP monitors participant progress (or lack thereof) throughout map construction. In contrast, existing methods assess cognitive mapping only when the map is complete. In addition, the CMP's 3D tangible interface allows a much more direct translation of cognitive maps into physical representations, with fully detailed buildings viewable in perspective from all sides, much as they are during travel through the represented neighborhoods themselves. Commonly used 2D assessment methods offer only highly abstracted 2D projections of the represented environment and its buildings. Ultimately, it should be possible to add adaptivity to the CMP, focusing more quickly and completely on the limits of participant ability, and improving sensitivity further.

Accessibility: Many of the populations commonly given cognitive mapping assessments face cognitive, visual or motor challenges. Unlike traditional 2D assessment techniques, the CMP uses an interface that is intuitive, easy to see, and simple to manipulate. This proved invaluable during our work with the elderly.

Consistency: If an assessment is to have meaning outside of its original context, all assessors must perform its procedures consistently and reliably. Existing manual 2D assessments are consistent, but achieving this consistency requires that the assessments be simple to perform, reducing assessment sensitivity. Because it is automated, the CMP achieves the highest level of consistency while at the same time improving sensitivity with complex tasks and very frequent measurement of the participant.

Control: The CMP's virtual neighborhood display will always be simpler than real-world stimuli. On the other hand, virtual display offers an amazing degree of control in assessment. Assessors can change climates, rotate or remove landmarks, display buildings located incorrectly by the participant translucently on top of correctly located buildings, and place neighborhood viewpoints in midair—effects extremely difficult if not impossible to achieve in the real world.

4. Assessing the Cognitive Map Probe

How sensitive is the CMP to well-known cognitive factors in practice? In this section, we describe the experiment we performed to find answers to this question. We also describe what we learnt about the accessibility and consistency of the CMP as we put it through its paces.

4.1. Methodology

We designed the CMP to support a wide range of cognitive mapping tasks. In this experiment, we sampled this range by varying the *number of buildings* in the virtual neighborhood participants viewed and attempted to recreate.

We expected that cognitive mapping performance would worsen by all measures as the *number of buildings* in the mapped environment increased. We also anticipated that performance among our elderly participants would be worse than the performance of our young participants, reflecting the natural effects of *age* on cognitive ability.

4.1.1. Participants

Our experiment had 20 participants, ranging in *age* from 25 to 81. Ten of the participants, ranging in age from 22 to 50 years old were treated as our under 55 group, or the "young" participants group. The average age of this group was 30.5 years, with a standard deviation of 8.31 years. The other ten participants, ranging in age from 55 to 81 years old were treated as our above 55 group, or the "elderly" group. The average age of this group was 68.9 years, with a standard deviation of 10.86 years.

We balanced both groups in gender. As a preliminary study, we also worked with one additional participant diagnosed with mild AD. Experimental results do not include this single participant unless otherwise noted.

4.1.2. Design

All participants performed 10 trials, beginning with three practice trials. During the seven recorded trials, participants viewed the same virtual neighborhoods in the same order, with *number of buildings* in virtual neighborhood increasing from two to eight across trials. Fig. 4 shows the 10 virtual neighborhoods used for the practice and actual trials. We were concerned about the transfer and practice effects that might result from this consistent increase in difficulty, but gave more importance to avoiding participant confusion and frustration, especially among the elderly. The steady increase in difficulty enabled us to identify the limit of each participant's wayfinding ability without requiring them to repeat more than a few disorienting trials beyond those limits. Indeed, three participants not analyzed in our discussion or results were unable to complete all 10 trials. Moreover, a steady increase in stimulus intensity is not uncommon

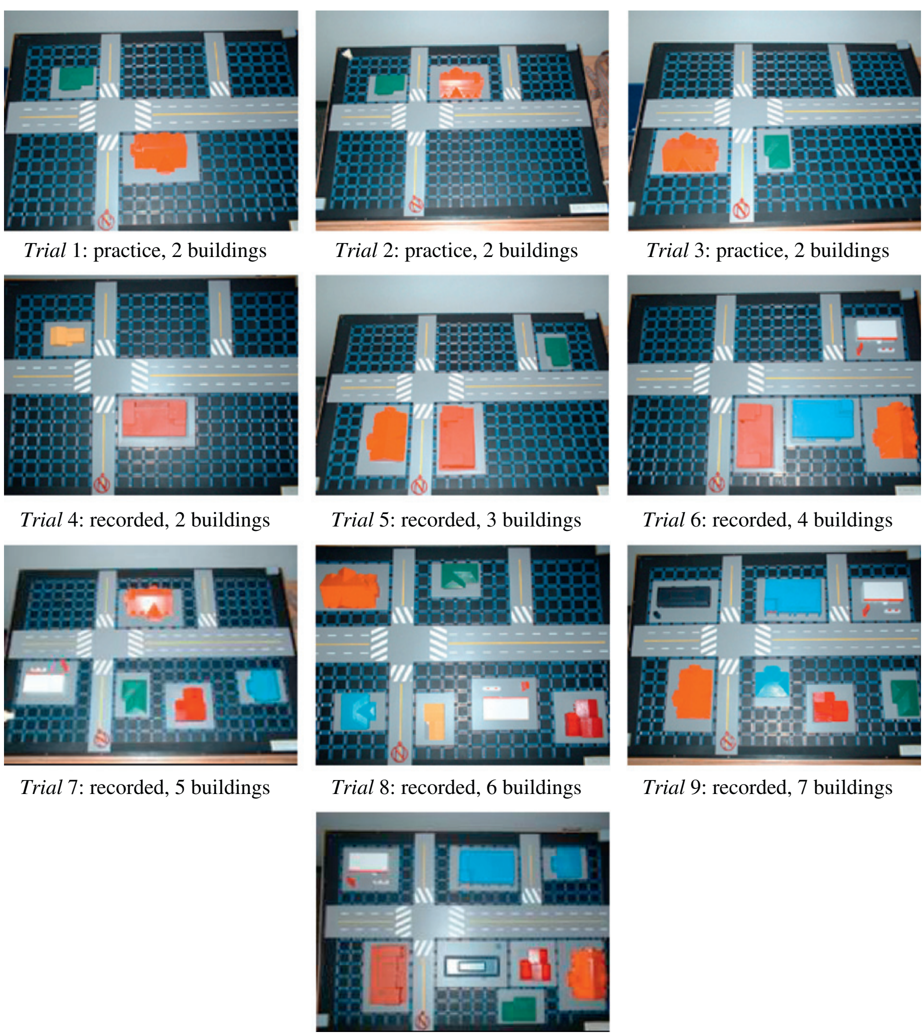

Trial 1: practice, 2 buildings
Trial 2: practice, 2 buildings
Trial 3: practice, 2 buildings
Trial 4: recorded, 2 buildings
Trial 5: recorded, 3 buildings
Trial 6: recorded, 4 buildings
Trial 7: recorded, 5 buildings
Trial 8: recorded, 6 buildings
Trial 9: recorded, 7 buildings
Trial 10: recorded, 8 buildings

Fig. 4. The 10 experimental neighborhoods correctly laid out on the CMP.

in assessment and psychophysics, being a central component of the experimental method of limits.

4.1.3. Apparatus and procedure

We conducted all experiments according to a strict written protocol, and with a script read aloud to each participant. The script introduced the participants to the CMP, the experiment, and its purpose, told them that they might stop the experiment at any time, and asked them to sign a consent form. We then interviewed participants quickly, obtaining answers to questions concerning age, education and occupation. Participant anonymity was always preserved.

We emphasized accuracy over speed in instruction, asking participants to be as precise as possible, but reminding them that the CMP was recording the speed of their actions. We told participants that there was no time limit and that they may decide when they had finished each task, but that they should do the best they could in reconstructing each neighborhood.

The assessor guided participants through the three initial practice trials to train them in the use of the CMP. All practice trials used simple, two building neighborhoods (Fig. 4, trials 1–3. In the first trial, the assessor introduced the CMP board and its models, as well as the ''bus ride'' metaphor for the largely passive, egocentric viewings participants would have of virtual neighborhoods. The assessor then took participants through a viewing of the virtual neighborhood that corresponded to the map already on the board. The assessor made certain that the participant understood this virtual–physical correspondence, and demonstrated that the passive viewing might be paused at will for a panoramic viewing (see below). In the second trial, the assessor introduced board interaction to the participant by asking the participant to identify a slight change to the virtual neighborhood during a new virtual tour. The assessor then turned off the virtual neighborhood display and asked the participant to adjust the CMP board to match this changed virtual neighborhood. In the third trial, the assessor confirmed that participants completely understood typical interaction by having participants view a completely new virtual neighborhood, and asking them to recreate it on the CMP board, again after the assessor turned off the virtual neighborhood display.

During the first phase of a recorded trial, participants viewed a virtual neighborhood from a passive, egocentric perspective, moving through the neighborhood at street level. The drivethrough of the virtual neighborhood followed the same predefined path in all the trials. A virtual compass appeared at the beginning of each trial on the ground plane, indicating which direction was north, and disappeared a short while after the drivethrough had began. Participants could optionally halt their motion at any time, and could ask the ''bus driver'' (the experiment administrator) to rotate slowly through 360° for a panoramic viewing before continuing along the viewing path. The assessor then turned the virtual neighborhood display off and participants moved into the trial's second phase, during which they interacted with the board and attempted to reconstruct the neighborhood they had just viewed from memory. A physical pointer similar to the compass seen in the first phase indicated which direction was north. Participants never received any feedback or comments about their performance from the CMP or the assessor. Participants required $1\frac{1}{2}$ h on average to complete the full set of 3 practice and 7 recorded trials, as well as a short post-assessment interview.

4.2. Results

Since the Segal model is a historic interface, we fully expected some noise in data collection. However, the CMP performed relatively well. Most importantly, noise did not force any participant to repeat a trial. The CMP also made no errors when reporting location. Nevertheless, there were errors when reporting the identity of the buildings attached or detached from the board. Unidentified buildings made up 18% of all actions on the board; the assessor corrected these interactively during the trial. Misidentified buildings made up less than 2% of all actions (21 actions total), but had to be corrected after assessment by manually matching CMP data to video recordings of the assessment. Though annoying, both types of errors occurred at rates quite manageable for our purposes and we are confident that a more polished implementation, possibly using different base technology, could eliminate most if not all of these problems.

Fig. 5 presents our experimental results by all dependent measures. We analyzed these results with one mixed analysis of variance (ANOVA) for each dependent measure. Each such analysis was two-way (2 *age* $\times$ 7 *num buildings*), with *age* a between subjects factor, and *num buildings* a within subjects factor. We present the results of these analyses in Table 1.

The CMP responded very much in line with our expectations to the cognitive factor *age* and the task factor *num buildings*. In the seven measures that responded significantly to *age*, the elderly were uniformly worse in cognitive mapping performance. In the seven measures that responded significantly to *num buildings*, response was more complex, with measures worsening initially as the number of buildings increases, then reaching a plateau or even improving slightly as the number of buildings reached maximum. It may be that when the number of buildings was high, the additional location constraints imposed by the physical street pattern on the board limited the number of possible configurations and made cognitive mapping easier. Alternatively or additionally, since participants always encountered trials with larger neighborhoods later in the assessment, participants may simply have been more practiced by the time these larger neighborhoods were encountered (one possible transfer effect).

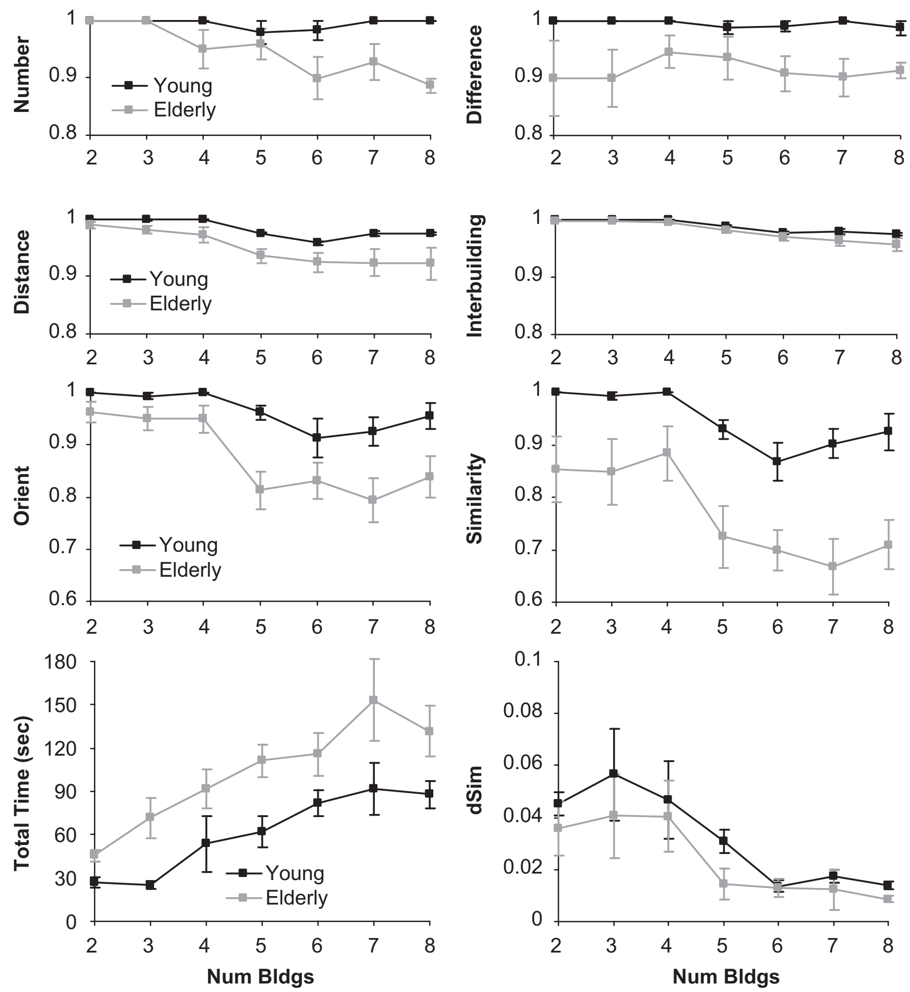

Fig. 5. Means and standard errors of the CMP's assessment measures, as a function of *age* and *number of buildings*. AD participant excluded.

Table 1
Results of two-way ANOVAs in CMP evaluation

Independent measures	Dependent measures	ANOVA
Age	*Totaltime*	$F(1,18) = 9.242$, $p = .007$
Age	*Number*	$F(1,18) = 14.797$, $p = .001$
Age	*Difference*	$F(1,18) = 14.928$, $p = .001$
Age	*Orientation*	$F(1,18) = 15.73$, $p = .001$
Age	*Distance*	$F(1,18) = 7.2$, $p = .015$
Age	*Interbuilding*	$F(1,18) = 10.29$, $p = .005$
Age	*Similarity*	$F(1,18) = 18.68$, $p < .0005$
# bldgs	*Totaltime*	$F(6,108) = 15.432$, $p < .0005$
# bldgs	*Number*	$F(6,108) = 3.400$, $p = .004$
# bldgs	*Orientation*	$F(6,108) = 3.537$, $p = .003$
# bldgs	*Distance*	$F(6,108) = 6.64$, $p < .0005$
# bldgs	*Interbuilding*	$F(6,108) = 15.789$, $p < .0005$
# bldgs	*Similarity*	$F(6,108) = 5.33$, $p < .0005$
# bldgs	*dSim*	$F(6,108) = 3.374$, $p = .004$
Age × # bldgs	*Number*	$F(6,108) = 2.884$, $p = .012$

Only *dSim* failed to respond significantly to *age*. Trends in the data indicated that rates of mapping progress for the young might become larger than rates for the elderly, where experimental sample size increased. Similarly, only *difference* did not vary significantly as *num buildings* changed. Here the null hypothesis—that the normalized set difference is simply not sensitive to the size of the map participants are attempting to reproduce—likely provides the best explanation of this result. However, an interesting reflective symmetry in the young and elderly curves (Fig. 5) may indicate opposite and canceling responses to the number of buildings.

The effects of *age* and *num buildings* interacted only in the *number* measure. While *num buildings* had little effect on the young, the mapping performance of the elderly dropped significantly by this measure as the number of buildings increased. This is likely due to an age-based difference in recall.

4.3. Discussion

In this section, we review the broader implications of our results for the CMP. We begin, however, by noting again that because of our need to find the cognitive thresholds of our participants quickly, we ordered experimental trials so that the *num buildings* factor increased steadily. Because of this pointed lack of counterbalancing or randomization in *num buildings*, practice effects are confounded with the observed effects of *num buildings*.

4.3.1. Confirmations

Our experimentation confirms that the CMP is statistically sensitive to age and environmental complexity, factors known to affect cognitive mapping performance. Elderly participants built less accurate maps than young participants, while larger maps were more difficult for participants to reproduce than smaller.

We were also pleased with the match of the CMP interface to the mapping task, and its accessibility to the elderly population. Almost all of our participants were able to complete all trials—and most reported they had fun doing so. This was true whether participants were university students or World War II veterans.

Our results are very preliminary, but we were also gratified to see that our single AD participant was among the worst performers, tentatively indicating possible use of the CMP for palliative care of persons with AD. Much more research is required before we can realize this application.

4.3.2. Surprises

We expected that assessment performance would worsen as *num buildings* increased. Instead, *num buildings* had a much more complex impact. While confounding practice effects certainly had an influence on this result, the initial *decrease* in mapping performance as the number of buildings increases (the opposite of a practice effect) leads us to believe that the increasing constraints on placement as the CMP was filled played a larger role. This suggests that researchers might control mapping difficulty in future experiments by varying proportion of the map used for street cues.

We did not expect the *age* $\times$ *num buildings* interaction we saw in our results. It would be interesting to see if performance in the *number* measure also declines for the young as the number of buildings increases further.

4.3.3. Implications

While our results indicate great promise for applications like the CMP, there is much work that remains if its assessment paradigm is to become common in clinical and research settings. First, the measurement sensitivity and reliability of CMP-like tools must be probed further, with comparisons made to existing assessment techniques, and typical score distributions found so that unusual assessment results might quickly be recognized. Second, tangible and tabletop interaction must become cheaper and more reliable, so that newer versions of the CMP will be more cost effective.

The applied promise of the CMP is the result of a design approach from which we hope others might benefit. Perhaps most important is the careful match of interface limitations to application constraints. TUIs are extremely innovative and bring with them new interactive freedoms, but like any other interface, they rarely have the expressiveness of corresponding manual tools. This leaves many TUI applications with a sense of unfulfilled promise. Fortunately, probing of cognitive maps does not require high expressiveness, enabling the CMP's tabletop TUI to capture most if not all of the applied input domain. In fact, for the elderly high expressiveness can become a barrier to successful interaction. This makes the CMP sufficient for real-world application, unlike many TUIs. Moreover, through careful use of our design heuristics (Sharlin et al., 2004), the CMP is also well suited to its application. It offers an intuitive spatial mapping for the task, providing a physical, 3D embodiment of the objects users arrange (buildings), and the environment in which they must be located (streets). It unifies the input and output spaces, permitting users to examine the results of their work (a map) in the same sensory space in which they create that work. Finally, it supports trial-and-error, permitting users to rearrange the elements of interaction (buildings) in any order, whenever they like. The CMP is the second outcome of our design approach: we found similar promise for a different sort of cognitive assessment in Cognitive Cubes (Sharlin et al., 2002).

5. Future work and conclusion

In this paper, we have presented the CMP, a TUI for the assessment of cognitive mapping ability. In experimentation, the CMP proved to be sensitive to factors known to affect cognitive mapping ability.

Our work on the CMP will continue. There are many interesting opportunities for improving its sensitivity. For example, the CMP could be used iteratively, with visual feedback given to the participant about the accuracy of the currently reproduced map, enabling the participant to attempt to correct their map. We might explore active or exocentric viewing modes. Wayfinding researchers might analyze the detailed histories of map building compiled by the CMP to find the decision trees formed by participants. Ultimately, the CMP might also prove useful in therapeutic applications.

We are interested in exploring how well the CMP matches measures obtained with established cognitive mapping assessment methods, such as distance and bearing. A comparative study of the CMP and a range of such methods would reveal the advantages and disadvantages of the CMP.

The CMP can certainly be improved. For example, replacing the Segal model with a newer tabletop interface would dramatically increase spatial resolution, flexibility and degrees of freedom. Yet, as we pointed out earlier, we are not certain that this input complexity would help us assess the elderly. For example, do we really want an elderly participant to have full, unconstrained control over the model orientation?

Most pressingly however, we plan to compare the sensitivity and utility of the CMP to a cognitive map probing application built using a standard 2D, Windows, Icons, Menus, and Pointing Devices (WIMP) interface. This would tell us what sort of advantage (if any) the CMP's 3D interface offers in practice.

Acknowledgments

Yoshifumi Kitamura and Yuichi Itoh discussed this work with us at length. This work was supported by an NSERC operating grant, and NSF Grant 0112937.

References

Arthur, P., Passini, R., 1992. Wayfinding: People, Signs, and Architecture. McGraw-Hill Ryerson, Toronto.

Axia, G., Peron, E.M., Baroni, M.R., 1991. Environmental assessment across the life span. In: Garling, T., Evans, G.W. (Eds.), Environment, Cognition and Action: An Integrated Approach. Oxford UP, New York.

Baird, J.C., 1979. Studies of the cognitive representation of spatial relations. Journal of Experimental Psychology: General 108 (1), 90–91.

Bowman, D., Kruijff, E., LaViola Jr., J., Poupyrev, I., 2005. 3D User Interfaces: Theory and Practice. Addison-Wesley, USA.

Cliburn, D., Winlock, T., Rilea, S., Van Donsel, M., 2007. Dynamic landmark placement as a navigation aid in virtual worlds. In: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, pp. 211–214.

Colle, H.A., Reid, G.B., 1998. The room effect: metric spatial knowledge of local and separated regions. Presence: Teleoperators and Virtual Environments 7 (2), 116–129.

Darken, R.P., Allard, T., 1999. Spatial orientation and wayfinding in large-scale virtual spaces II. Presence: Teleoperators and Virtual Environments 8 (6), iii–vii.

Darken, R.P., Goerger, S.R., 1999. The transfer of strategies from virtual to real environments: an explanation for performance differences? Virtual Worlds and Simulation (VWsim'99), 17–20 January, San Francisco, CA 31 (2), 159–164.

Darken, R.P., Peterson, B., 2002. Spatial orientation, wayfinding, and representation. In: Stanney, K. (Ed.), The Handbook of Virtual Environments: Design, Implementation, and Applications. Lawrence Erlbaum Associates (Chapter 28).

Fitzmaurice, W.G., Ishii, H., Buxton, W., 1995. Bricks: laying the foundations for graspable user interfaces. In: Proceedings of the CHI '95, pp. 442–449.

Frazer, J.H., 1982. Use of simplified three-dimensional computer input devices to encourage public participation in design. In: Proceedings of the Computer Aided Design 82. Butterworth Scientific, pp. 143–151.

Frazer, J.H., 1995. An Evolutionary Architecture. Architectural Assoc.

Goerger, S., Darken, R., Boyd, M., Gagnon, T., Liles, S., Sullivan, J., Lawson, J., 1998. Spatial knowledge acquisition from maps and virtual environments in complex architectural spaces. In: Proceedings of the 16th Applied Behavioral Sciences Symposium. US Air Force Academy, CO.

Golledge, R.G., 1991. Cognition of physical and built environments. In: Garling, T., Evans, G.W. (Eds.), Environment, Cognition and Action: An Integrated Approach. Oxford UP, New York.

Hornecker, E., Buur, J., 2006. Getting a grip on tangible interaction: a framework on physical space and social interaction. In: Proceedings of the ACM CHI, pp. 437–446.

Hosokawa, T., Takeda, Y., Shioiri, N., Hirano, M., Tanaka, K., 2008. Tangible design support system using RFID technology. In: Proceedings of the ACM Conference on Tangible and Embedded Interaction, pp. 75–78.

Koh, G., Von, V., Thomas, E., 1999. Use of virtual environments for acquiring configurationally knowledge about specific real-world spaces: preliminary experiment. Presence: Teleoperators and Virtual Environments 8 (6), 632–657.

Lezak, M.D., 1995. Neuropsychological Assessment. Oxford University Press, Oxford.

Liu, L., Gauthier, L., Gauthier, S., 1991. Spatial disorientation in persons with early senile dementia of the Alzheimer type. American Journal of Occupational Therapy 45 (1), 67–74.

Piaget, J., Inhelder, B., 1956. The Child's Conception of Space. The Humanities Press Inc., New York, pp. 419–446.

Sadeghian, P., Kantardzic, M., Lozitskiy, O., Sheta, W., 2006. The frequent wayfinding-sequence (FWS) methodology: finding preferred routes in complex virtual environments. International Journal of Human–Computer Studies 64, 356–374.

Sharlin, E., Itoh, Y., Watson, B., Kitamura, Y., Sutphen, S., Liu, L., 2002. Cognitive cubes: a tangible user interface for cognitive assessment. In: Proceedings of the ACM CHI.

Sharlin, E., Watson, B., Kitamura, Y., Kishino, F., Itoh, Y., 2004. On tangible user interfaces, humans and spatiality. Personal and Ubiquitous Computing 8 (5), 338–346.

Smith, S.P., Hart, J., 2006. Evaluating distributed cognitive resources for wayfinding in a desktop virtual environment. In: Proceedings of the IEEE Symposium on 3D User Interfaces.

Stanney, K., 2002. Handbook of Virtual Environments: Design, Implementation, and Applications. Lawrence Erlbaum Associates.

Tolman, E.C., 1948. Cognitive maps in rats and men. Psychological Review 55 (4), 189–208.

Ullmer, B., Ishii, H., 2001. Emerging frameworks for tangible user interfaces. In: John, M. Caroll (Ed.), Human–Computer Interaction in the New Millennium. Addison-Wesley, USA.