Generating Trails Automatically, to Aid Navigation when you Revisit an Environment

[The final version of this article was published in Presence, Vol. 17, Issue 6, pages 562-574, published by The MIT Press.]

Roy A Ruddle

School of Computing

University of Leeds

Leeds

LS2 9JT

England

email: royr@comp.leeds.ac.uk

tel : +44 (0)113 343 1711

fax : +44 (0)113 343 5457

Abstract

A new method for generating trails from a person's movement through a virtual environment (VE) is described. The method is entirely automatic (no user input is needed), and uses string-matching to identify similar sequences of movement and derive the person's primary trail. The method was evaluated in a virtual building, and generated trails that substantially reduced the distance participants traveled when they searched for target objects in the building 5-8 weeks after a set of familiarization sessions. Only a modest amount of data (typically five traversals of the building) was required to generate trails that were both effective and stable, and the method was not affected by the order in which objects were visited. The trail generation method models an environment as a graph and, therefore, may be applied to aiding navigation in the real world and information spaces, as well as VEs.

1 Introduction

Trails are an integral part of how humans navigate. In the real world we routinely use trails left by other people to minimize the physical effort of traveling through snow or forest, and to reduce the cognitive effort involved when traveling from one place to another in the countryside. Modern technology such as GPS systems now allows us to record our own movements through the real world, our travels through information spaces are captured by mechanisms such as web browser history lists, and equivalent functionality would be straightforward to implement for virtual environments (VEs). This means that it is now possible for an individual to capture all of their life’s travels, but how could this vast quantity of raw movement information be presented in a form that the individual can utilize?

One thing that is certain is any trail system must consolidate people’s raw movements so they are not overwhelmed by masses of detail (Lamming et al., 1994; O’Hara et al., 2006). In this paper we describe a new method for automatically generating trails from a person’s movements, which helped them remember where they had previously traveled and made it significantly easier to return to places after a lengthy absence. At the core of our method is the use of string-matching to identify sections of paths that are traveled frequently.

This paper is divided into two main sections. The first describes our trail generation method, starting with related research into trails in VEs, the real world and information spaces, all of which provide foundations for our method. The second describes the evaluation of our trails using a study in which participants first familiarized themselves with the layout of a virtual building and then, 5-8 weeks later, returned to it.

2 Automatically generated trails

Trails are "beaten paths" that arise from people's incidental navigation, as opposed to *tours* which are equivalent to perfectly laid out walkways designed by one person for the use of others (Reich, Carr, De Roure, & Hall, 1999). VE trails have presented user's raw movements in the form of virtual footprints, vapor trails, balls of string or breadcrumbs (Darken & Sibert, 1993; Grammenos, Filou, Papadakos, & Stephanidis, 2002; Pettifer, Cook, & Marsh, 2004; Ruddle, 2005). In most cases the trail data has been recorded automatically, rather than expecting a user to specify each waypoint, and this echoes findings from research into information spaces which showed that expecting users to manually identify each item that formed part of a trail was a major limitation of an otherwise successful system (De Roure et al., 2001). Examples of tours in VEs include presenting the path to be followed by a "control" group of participants as a line on the floor of a VE, so that navigation was trivial (e.g., Hartley, Maguire, Spiers, & Burgess, 2003).

For trails to be implemented, it is commonplace for the environment to be modeled as a graph. For environments like buildings and street networks every junction and dead end is a *node* (e.g., Ruddle, 2005), whereas for open plan environments nodes correspond to areas of the environment that are derived using knowledge of its content and structure (e.g., Peponis, Conroy Dalton, Wineman, & Dalton, 2004). For all these types of environment the path segment that connects a pair of nodes is a *link*, and it follows that the route a person travels between two places is a sequence of links.

Most existing trail systems record and store a person's raw movements, but where the systems differ is in how that movement information is subsequently made available. Some systems seek to match a person's requirements to one specific trail, and then make available the raw movements of that trail. Examples are systems that allow people to share GPS trails for hiking and mountain biking routes (e.g., www.memory-map.co.uk), and browser plug-ins that allow a person to follow a particular search path that someone took on the WWW (e.g., www.trexy.com). In VEs, specific trails have been used to illustrate an individual's path when reporting studies of navigation (e.g., Gamberini, Cottone, Spagnolli, Varotto, & Mantovani, 2003; Lessels & Ruddle, 2005).

Other trail systems merge different paths together, for which there are three main approaches. The most basic of these just provides a Boolean flag for each node/link, which indicates whether it has ever been visited/traversed. In studies of Web navigation, portraying this Boolean information on a graphical overview helped participants to return quickly to webpages (nodes) that had recently been visited (Hightower, Ring, Helfman, Bederson, & Hollan, 1998; Milic-Frayling & Sommerer, 2003; Utting & Yankelovich, 1989), but over an extended period of time the number of pages visited becomes large and, therefore, finding the one you wish to return to becomes very difficult. Web browser history lists attempt to overcome this by allowing page visits older than a certain number of days to be forgotten, but this just means you can only return to pages visited in the recent past.

A second approach is to portray a user's actual movements (e.g., see Figure 2a), which allows the number of times each link has been traversed to be determined and also potentially allows particular links to be recognized from the shape of the user's movements. Evaluations showed that vapor trails helped users find each other in a collaborative VE (Pettifer et al., 2004) and the build up of the trails allowed people to gain an impression of which places have been visited frequently vs. rarely, but over an extended period of time a spider's web of trails is created (see Figures 1 & 2a) and this "trail pollution" impedes rather than aids navigation

(Darken & Peterson, 2002; Grammenos et al., 2002; Ruddle, 2005). All three of these groups of researchers have attempted to overcome pollution by supplementing trails with detailed information about the time a path was traversed, in which direction and/or by whom, but in no case did the supplementary information have a significant effect on participants navigation. Portraying every path is also the approach adopted for information spaces, with systems such as Google’s Web History also allowing navigation within individual websites to be consolidated into a single item that provides a high level view.

The third approach merges trails by counting the number of times each node/link was visited/traversed. These data may be used either as a filter, so only nodes/links visited/traversed more than a certain number of times are displayed (Cockburn & Jones, 1996), or to affect how a node/link is displayed (e.g., using link thickness to indicate “traversal frequency” (how often a link has been traversed); Wexelblat & Maes, 1999). When evaluated in a 20 minute web browsing task, participants who had access to their predecessors’ trails found information significantly faster than participants who had to search by themselves (Wexelblat & Maes, 1999). However, usage of this approach in VEs has so far been limited to the analysis of navigational paths (e.g., Chittaro & Ieronutti, 2004; Elvins, Nadeau, Schuk, & Kirsh, 2001).

2.1 A new method for automatic trail generation

To be successful, a trails system needs to carefully select which information to provide, and what to “forget” (Lamming et al., 1994; O’Hara et al., 2006). The raw trails and traversal frequency approaches described above overload a person with so much information that, in practice, they are unlikely to make use of it. To counteract this, our approach generates trails that present a small amount of carefully chosen information which on the one hand reinforces a person's memory for where they have traveled, and on the other complements it. The simplicity of the information that is presented outweighs the loss of detail.

The rationale for our approach is as follows. It is generally accepted that route memory is central to people's ability to navigate successfully in everyday settings, and it has been suggested that a network of paths is one of the key elements of people's mental models of the places in which they live ("for most people interviewed, paths were the predominant city element"; p49, Lynch, 1960). Therefore, a person's spatial knowledge could be reinforced if the paths they frequented most often were presented as a *primary trail* that also then provides a framework for the person's navigation. Other paths the person had traversed could be presented as a *secondary trail*, complementing memory by ensuring they didn't forget anywhere they had been (e.g., as in Skopik & Gutwin, 2005).

Generating a primary and secondary trail automatically involves three steps: (1) a scoring method that identifies how traveled each part of the environment is, (2) a filtering method to choose which links are included in a trail, and (3) how this trail should then be presented. These are described in the following subsections. Scoring uses a novel application of string-matching to "reward" sequences of links that are traversed together, works for sets of destinations rather than only for a path between a pair of specified end points, and is entirely automatic. The filtering method is novel because it permits branches in an otherwise linear trail, and the presentation method is also new.

2.1.1 Scoring methods

The traversal frequency approach (Cockburn & Jones, 1996; Wexelblat & Maes, 1999) treats each link separately and, therefore, immediately loses the concept of a person's paths (i.e., sequences of links that were followed) through a space. More sophisticated solutions identify identical or similar sequences of movement. This may be achieved using a probabilistic approach (Borges & Levene, 2000) but string-matching has been used much more widely.

String-matching is a technique that can be used to find the parts of two strings that are identical (e.g., identifying that the longest common subsequence of “color” and “colour” is “colo”), or quantify the overall similarity between two strings (e.g., the number of characters that need to be inserted, deleted or edited to change one string into another). The technique is one of the core methods that underpins spellchecking in word processors, and is routinely used to analyze gene data (Krane & Raymer, 2003). String-matching has also been successfully used to analyze navigation patterns in the real world (Conroy Dalton, 2003; Peponis et al., 2004) and in websites (Barra, Malandrino, & Scarano, 2003; Ruddle, 2006; White & Drucker, 2007), and find similarities in paths through a computer file system (Gams & Reich, 2004).

We have investigated a number of methods of using string-matching to generate trails. One method used a pre-processing stage to find identical sub-sequences and then scored these using the inexact gene-alignment approach. However this led to trail fragmentation because the pre-processing cleaned the data too much (a certain amount of “noise” seems to be important to glue trails together). A superior method applied string-matching directly to the raw trail data, and was implemented according to the following pseudocode:

```
#
# Initialization
#
INITIALIZE each link’s score to zero
```

```
    STORE path person followed in each navigation session as
sequence of links
    #
    # Calculate score for each link using
    # string-matching
    #
    FOR each pair of paths

      REPEAT

        FIND longest, unmatched subsequence
        of links that appears in both paths

        SET subseqLength TO number of links
        in subsequence

        SET pathLength TO number of links in
        most recent of the two paths

        CALCULATE score for subsequence,
        S = subseqLength / pathLength

        FOR each link in subsequence
          INCREASE link’s score by S
```

```
        ENDFOR

        INDICATE subsequence has been
        matched

      UNTIL no unmatched link appears in
      both paths

    ENDFOR
    #
    # Adjust scores to take account of each
    # link’s physical length
    #
    MULTIPLY each link’s score by length of
    the link (in meters)
```

The method generates scores in the range 0.0 (two paths that have no links in common) and 1.0 (the paths are identical), and is best illustrated using an example. Consider a person who first travels the path ABCWXYDE and subsequently travels the path ABCDE (each letter denotes a segment of the path). The longest subsequence, ABC, would have a score of 3 / 5 = 0.6 (subsequence length / length of the second path). The only other common subsequence, DE, would have a score of 0.4.

Alternatively, scoring may be performed by adopting the type of inexact string-matching that is used in gene alignment and allows different values for the match bonus, and delete, substitute, insert origin and insert continuation penalties (for details, see Krane & Raymer, 2003). If the match bonus equaled 1.0 and all the penalties equaled -1.0 then the example above would generate two subsequences with the same scores as exact string-matching. However, if the insert continuation penalty was -0.25 then there would be one subsequence with a score of (5 x match bonus - 1.0 x insert origin - 2.0 x insert continuation) / (path length x match bonus) = 0.7.

2.1.2 Filtering methods

From raw movement data (e.g., see Figures 1 & 2a) it is difficult to identify a person's primary pathway. Therefore, the next stage in our trail generation process is a filtering method that determines which links should make up a person's primary trail, and which should be relegated to the secondary trail. The method used to determine a primary trail has to make a balance between making the trail easy to follow vs. providing wide coverage so the trail reliably includes all of a person's frequent destinations.

Wide coverage occurs if the primary trail is allowed to contain branches and loops. However, a pilot study showed that this made the trail too complex to be beneficial during navigation. There was minimal chance of a participant choosing the correct trail at each junction to find an efficient path to the target destinations, so participants tended to traverse the trail in fragments rather than in its entirety.

The easiest type of trail to follow is a linear sequence of links, and the benefit of such a structure has been shown by research into the design of tours ("Walden's paths") through hypertext (Dave et al., 2003; Furuta, Shipman, Marshall, Brenner, & Hsieh, 1997). We took guidance from these findings and developed a method of generating a primary trail that was mainly linear, but allowed side branches along sequences of links that a person had often backtracked. This reflected the fact that people often deliberately employ a backtracking strategy to avoid becoming disoriented, and in most environments there are places that lie in a dead end and can only be reached by backtracking.

The method worked by finding the highest scoring “linear” trail through the environment, with any branches only contributing half of their links’ scores because a person following the trail would have to backtrack once they reached the end of a branch. The link scores were calculated using the method outlined in the previous section (see pseudocode). The resulting trail captured a person’s primary path through the environment and was straightforward to follow (e.g., see Figures 4a & 4b in the Results section).

2.1.3 Presentation methods

The most basic presentation method is to show a person’s raw movements (see Figure 2a). Compared with having no trail at all, a raw trail halved the distance that participants traveled the first time they searched an environment for targets, but on subsequent occasions the benefit of having a trail was substantially reduced (Ruddle, 2005). Participants paid little attention to the trail when there was a lot of trail pollution and, instead, relied on their memory for the targets’ locations, which of course was error-prone.

To overcome the effects of pollution our primary trail identified a person’s main pathway through an environment, and all the other links they had traversed were included on the secondary trail. Through pilot studies, we identified important issues in the way these trails are presented.

Representing the primary and secondary trails as solid and dashed cyan “pavement”, respectively, produced two major problems (see Figure 2b). First, when seen from a distance (e.g., looking down a corridor to plan one’s route) the primary and secondary trails were difficult to distinguish from each other because the dashes of the secondary trail appeared to merge together. Second, we had reasoned that the primary trail would contain only short branches formed by backtracking, so when presented with a branching trail a person would notice which section was short and search that first, before continuing on the main part of the trail. In practice this was not the case, so participants sometimes missed some of the targets.

These problems were solved by using different colors for the primary and secondary trails, boosting the saliency of the primary trail by increasing its width, and explicitly marking the sections of the primary trail that were side trails (see Figure 2c). This meant it was trivial for a person to follow the whole primary trail, including branches, from beginning until end.

The primary/secondary trail was recomputed at the end of each navigation session, so a participant’s movements during the current session were overlaid as a raw trail. Although it was commonsense to distinguish between previous/current movement, the benefit it provided was probably small because a pilot study had shown no significant difference between raw trails that were always one color vs. a different color for each navigation session.

3 Experiment

The latest version of our trails was formally evaluated in an experiment. A mixed factorial design was used with each participant navigating a 24-room virtual building in a total of eight sessions, using either a *string-matched trail* that automatically highlighted a participant’s primary path through the building (see Figure 2c) or a *raw trail* which was equivalent to a ball of string that unrolled behind a participant while they traveled (see Figure 2a).

The eight sessions were divided into six “familiarization” sessions that took place on one day, and two “revisit” sessions that took place 5-8 weeks later on another day. In each session, the task was to find six target objects in any order and participants were informed that the targets were always in the same place.

3.1 Method

3.1.1 Participants

A total of 20 people (14 men and 6 women) took part in the experiment. Their mean age was 23.4 years (SD = 4.3). Participants were randomly assigned to the two groups, subject to the constraint of an equal number of men/women in each group. One man from each group did not return for the revisit sessions. All the participants volunteered for the experiment, gave their informed consent, and were paid an honorarium for their participation.

3.1.2 Materials

The VE software was an OpenGL application that ran on a Linux PC. The display was a 17-inch monitor that had a resolution of 1280 x 1024 pixels. The horizontal graphical FOV was 90 degrees, which allowed participants to look simultaneously down two corridors that intersected at 90 degrees, albeit with the disadvantage of introducing some distortion from the physical FOV. The application update rate averaged 50 Hz. and data concerning a participant’s position and orientation was recorded to disk in real-time.

Three virtual buildings were designed, each with a set of rooms bounded on two sides by a wide entrance hall. The rooms were all the same size and were laid out using a regular grid pattern. Each room contained one entrance, which was always “open” (no door) and allowed participants to look inside. In each building some of the rooms contained a target object (a picture of an object, texture mapped onto a 1.5m side cube that was positioned in the center of the room), and the remainder of the rooms were empty. Different objects were used as targets in the three buildings. To provide visual landmarks, pictures of other objects were texture mapped onto the wall at every corridor junction and each corner of the building.

The building that was used to allow participants to learn the user interface contained four rooms and four target objects. The practice building contained eight rooms arranged in a 4 x 2 grid, and three targets. The test building contained 24 rooms (6 x 4 grid) and six targets (see Figure 4). In the test building, two of the targets (dog and strawberry) were accessed from dead-end corridors, but the others were on through sections of corridor (car, clock, piano and saucepan). The buildings' graph structures had nodes at every corridor intersection, corner and dead end, and in every room (each room had one entrance and so was also a dead end).

The user interface utilized the mouse and keyboard. To move forward, backward, left or right a participant held down the appropriate cursor key, which moved them at 2 m/s (a fast walk). If a participant held down pairs of keys such as forward/left then they moved diagonally. Moving the mouse up/down allowed the participant to look up to 90 degrees up/down. Moving the cursor to the left or right rotated the view direction in that direction at a rate that increased with the horizontal offset of the cursor from the center of the screen, to a maximum rate of 60 degrees/s. Movement was always in the participant's direction of view.

At all times in a given session, the targets that remained to be found were indicated by a text message on the screen. To indicate that a target had been found the participant pressed the <enter> key when they were in the target's room, which caused the name of that target to be removed from the display.

3.1.3 Trail design

The raw trail (control) group was provided with a trail that showed exactly where a participant had traveled in all sessions, including the current one (see Figure 2a).

The trail for the string-matched group was generated solely from a participant's movement data, and presented using the method shown in Figure 2c. The primary/secondary trail was calculated from the participant's movements in all previous sessions, with movements in the current session overlaid as a raw trail. The scoring method used to generate the primary trail was as previously described (see pseudocode, above), with the string-matching implemented using an efficient, established algorithm (Smith & Waterman, 1981). The filtering method used the linear trail approach (see Filtering Methods, above) and, to be allowable as a branch, more than 50% of a link's traversals had to have been either the outward or return leg of when the participant was backtracking.

3.1.4 Procedure

Each participant came to our laboratory on two occasions, first to become familiar with the test building and then, 5-8 weeks later, to revisit it. The procedure was as follows.

The first occasion lasted up to two hours, during which a participant learned the user interface controls, searched the practice building four times for its three targets, and then searched the test building six times for its six targets. Participants were informed that, in each building, the searches always started in the same place, the building remained unchanged from session to session, the targets were always in the same place, and could be found in any order. In the practice building participants were provided with the same type of trail (raw vs. string-matched) that they would have in the test building, so they gained experience of how the trail developed and could be used to assist navigation.

The second occasion lasted up to 30 minutes, during which a participant just searched the test building twice for its targets. Participants were informed that the building was identical to when they performed they familiarization sessions, the trails were their own from before and would be added to during the two revisit sessions. The only thing that had changed in the 5-8 weeks since familiarization was participants' own memory for where things were in the building.

3.2 Results and discussion

The goal of our research is to develop trails that help people navigate when they return to a place that they used to be familiar with. The effect of trails on navigation was determined by analyzing participants' navigational performance and behavior. Assessment of the characteristics of the trails themselves was performed by analyzing factors such as the quantity of data required to generate a useful trail, stability (i.e., how a trail changed over time), and a comparison of our method with the more basic approach of traversal frequency. The data presented in this section are for the 18 participants who returned for the revisit sessions.

3.2.1 Effect of trails on navigation

Navigational performance was measured by expressing the distance, *d*, that participants traveled in each session as a percentage in excess of the shortest path, *s*, from the start point to all the targets. Thus, 0% meant a participant traveled to the targets by the shortest possible route, which was 281 meters and took approximately two minutes.

$$\text{Percentage extra distance traveled} = 100 * (d - s) / s$$

The performance data were analyzed using analyses of variance (ANOVAs) that treated the type of trail (raw vs. string-matched) and gender as between participants factors, and the sessions as a repeated measure. The time that elapsed before revisiting the environment was partly dictated by participants' availability, and preliminary analyses confirmed that this elapsed time had no effect on participants' performance.

Two ANOVAs were performed. The first just analyzed the familiarization sessions (Sessions 1-6 in Figure 3), to check whether the type of trail affected the rate at which participants learned the locations of the targets. Participants did travel a shorter distance as familiarization progressed ($F(5, 70) = 9.77$, $p < .001$), but the degree to which the target locations were learned was not affected by the type of trail being used ($F(1, 14) = 0.24$, $p = .63$) or gender ($F(1, 14) = 3.70$, $p = .08$) and there were no significant interactions (specifically, for the trail type vs. session interaction, $F(5, 70) = 0.70$, $p = .63$).

The second ANOVA analyzed participants' ability to navigate when they revisited the environment, which was the main purpose of our trails. Participants' mean performance in the last two familiarization sessions was compared with their mean performance in the two revisit sessions (Sessions 5 & 6 vs. 7 & 8 in Figure 3). Overall participants' performance deteriorated ($F(1, 14) = 18.10$, $p = .001$), but there was no main effect of type of trail ($F(1, 14) = 0.89$, $p = .36$) or gender ($F(1, 14) = 2.43$, $p = .14$). However, there were significant interactions between the type of trail and sessions ($F(1, 14) = 5.81$, $p = .03$) and gender and sessions ($F(1, 14) = 5.46$, $p = .04$). With a string-matched trail, participants only traveled 4% further during revisitation than they had during the last two familiarization sessions (men traveled 4% less far, but women traveled 20% further). However, with a raw trail participants traveled 23% further during revisitation (men and women traveled 21% and 39% further, respectively).

So how did the string-matched trail aid participants' navigation when they revisited the environment? This may be answered by looking in detail at the paths participants followed during the first revisit (Session 7).

Three participants in the string-matched group simply followed their respective trails from the start point to all six targets in turn (see Figure 4a), and another three participants followed their trails but took a shortcut along the way (see Figure 4b). The trails of the other three participants began in the main part of the building, rather than where the navigation session started (so that trail generation was completely automatic, our method was provided with no information about key locations such as the start position or location of the targets). These participants joined/left their trail on up to seven occasions, piecing together segments until all the targets were found. For two of these participants all six targets were on the trail, but the other participant had to make a detour off the trail to one target (the Saucepan; see Figure 4c).

By contrast, six of the nine participants in the raw trail group had notable difficulty when they revisited the building in Session 7. The root causes of these difficulties were that participants: (a) repeatedly traveled through some parts of the environment instead of searching in places that were untouched in that session, and (b) failed to notice a target when they were in its vicinity. The former could have been prevented simply if the participant had been able to focus their search on the sections of the building that were covered by a string-matched trail, as was done by the participant who followed the path shown in Figure 4c. The latter was much less likely to occur with a string-matched trail because of its visual salience.

3.2.2 Trail characteristics

The previous section provided quantitative evidence about the benefits of trails when participants revisited the environment, and explained why the string-matched trails were useful. Movement data from the raw trails group were generated into “string-matched” trails at the end of the experiment, and found to have similar characteristics to the trails of the string-matched group. The present section analyses these characteristics.

To determine how much data were required to generate a useful trail we counted the number of targets that lay on a trail at the end of each session. After three sessions 33% of the trails contained all the targets, rising to 78% after Session 5 and 83% from Session 6 onward (see Figure 5). From Sessions 5 to 8 an average of 97% of each trail remained the same from one session to the next, meaning that the trails provided stable cues that could be depended on for navigation.

Two other points should be emphasized. First, trail quality did not depend on how quickly a participant had learned the targets' locations. This is shown by the fact that by the end of Session 5 the trails of the six best-performing participants and three of the four worst-performing participants were all trivial to follow (each trail began at the start point and visited all the targets; e.g., as in Figures 4a & 4b).

Second, trail quality was not dependent on participants visiting the targets in a consistent order. The trivial-to-follow trails were generated from paths that visited an average of 2.4 targets in the same order from one session to the next, whereas for all the other trails an average of 2.5 targets had been visited in the same order. The trivial-to-follow trail shown in Figure 4b was for a typical participant, who ranked 10th out of 18 in terms of total distance traveled during the familiarization sessions, and had visited an average of 2.0 targets in the same order from one session to the next.

3.2.3 Alternative scoring methods

Our method used exact string-matching to calculate link scores (see pseudocode in *Scoring Methods*), but what effect would a basic link frequency traversal approach, or inexact string-matching, have had on the trails?

To analyze this, the trails that would have been created at the end of Sessions 5-8 were calculated using: (i) traversal frequency, and (ii) inexact string-matching that used a wide range of parameters for the match bonus, and delete, substitute, insert origin and insert continuation penalties.

Although the number of targets on a trail was not affected by the scoring method, traversal frequency led to primary trails that were substantially longer (M = 527 m) than those generated by exact string matching (M = 422 m). Long, meandering trails are a disadvantage because they are unlikely to be followed in their entirety (see Figure 4c).

Use of inexact, rather than exact, string-matching also had no effect on the number of targets on a primary trail, and caused a modest increase in trail length. However, even with an extreme set of parameters (match bonus = 1.0; delete, substitute and insert origin penalty = -0.1; insert continuation penalty = -0.01) the mean length only increased to 449 m.

4 Conclusion

This paper describes a new method for generating trails. People dislike having to perform a large amount of manual work to record a trail (De Roure et al., 2001), so we developed a completely automatic method that generated a trail solely from a person's movements through an environment. Our method was evaluated in a virtual building, and showed that there was a significant interaction between the distance traveled during familiarization with the building vs. revisiting it 5-8 weeks later, and the type of trail being used (string-matched vs. raw). With our string-matched trail, participants' navigation only deteriorated slightly when the building was revisited (4% increase in distance traveled). By contrast, participants in a control condition (a raw trail) traveled 23% further during the revisits. A comparison was not made with trail-less navigation, because previous research indicates that there is little difference a raw trail and no trail after an environment has been navigated several times (Ruddle, 2005).

Two additional points should be noted. First, most of the difference in distance traveled occurred during the first revisit (Session 7; see Figure 3). If a person has been absent from an environment for years then our trails are likely to be useful for an extended period, while the person reforms/repairs their cognitive map. However, if the absence was shorter (e.g., the few weeks of the present study) then the benefit will be shorter lived. Second, caution should be exercised in interpreting the significant interaction of gender with sessions that occurred for the revisits, because the sample size was small (the study was designed to evaluate trails, not investigate gender) and previous research has shown that interface proficiency and spatial ability account for much of gender’s effect (Waller, 2000).

Our method used string-matching to emphasize long segments of paths that were traversed frequently. From this, the method identified a person’s primary trail through the environment, echoing the importance of route knowledge in everyday navigation and the suggestion that a network of paths is a key element in our mental models of the places in which we live (Lynch, 1960). The method only required a modest quantity of data to generate a useful trail (after five navigation sessions, 78% of the trails contained all six target destinations) which then remained stable and, therefore, could be relied on during navigation.

By using a scoring method that emphasized identical sequences of movement, the method generated trails that were substantially shorter than those that would have been generated by a basic link traversal frequency method. Short trails lead a person directly to a destination whereas, as our pilot studies have shown, a long trail tends to be followed in fragments rather than as a whole.

In addition, the replacement of exact string-matching by inexact matching had little effect on the trails. This shows that our method is robust in terms of the particular string-matching parameters that are chosen.

In terms of application, large-scale VEs such as *Second Life* are part of everyday life for many people, with organizations from universities to news agencies, musicians and retailers making use of such worlds to take advantage of new methods of social communication and for doing business. Given how disoriented many people become when navigating in VEs (Ruddle, 2005), our method of trail generation could be very beneficial. If a person used one of these VEs extensively, visiting some parts for professional purposes and different parts during leisure time, then string-matching would only find similarities in paths to destinations visited for each purpose so each could have its own primary trail. On the other hand, if the sets of destinations overlapped in space then they would all be integrated into one primary trail, which may be a better match for our memory of where we had traveled than if the sets were artificially separated. Alternatively, some contextual information could be provided to allow a small number of mutually distinct trails to be generated. It would also be possible to generate trails from the movements of different people so server-side trails could be generated to assist first-time visitors to a VE that an organization maintains.

Finally, the trail generation method described in this paper is applicable to any environment that may be modeled as a graph, and this includes the real world and information spaces. The availability of low-cost, lightweight positioning devices (e.g., GPS systems used by hikers, and GPS-enabled mobile phones) means that it is already feasible for people to record their travels, graph structures suitable for trails can be generated for both environments that have intrinsic nodes/links (e.g., buildings) and those that are open spaces (Peponis et al., 2004), so our trail generation method could be used a broad range of settings to turn people's raw travel data into an effective aide memoir for future use. Due to their scale and complexity, information spaces provide the greatest navigational challenge that many people face in their everyday lives. Our method has clear potential for addressing this challenge by generating personal trails from data captured automatically by one's web browser. For this, websites are already structured as a graph (each page is a node) and recent advances in computer hardware means it is opportune to consider how browsers could present pages in a more spatial fashion, so that people could see a set of pages laid out in front of them, rather than only viewing each page discretely, and a website becomes an information landscape.

References

Barra, M., Malandrino, D., & Scarano, V. (2003). "Common" web paths in a group adaptive system. *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT'03*, pp. 218-219). New York: ACM.

Borges, J., & Levene, M. (2000). A heuristic to capture longer user web navigation patterns. In K. Bauknecht, S. Kumar Madria, & G. Pernul (Eds.), *Lecture Notes in Computer Science, 1875*, 155-164. Berlin: Springer-Verlag.

Chittaro, L., & Ieronutti, L. (2004). A visual tool for tracing users' behavior in virtual environments. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'04*, pp. 40-47). New York: ACM.

Cockburn, A., & Jones, S. (1996). Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human-Computer Studies, 45*, 105-129.

Conroy Dalton, R. (2003). The secret is to follow your nose. *Environment and Behavior, 35*, 107-131.

Darken R. P., & Peterson, B. (2002). Spatial orientation, wayfinding, and representation. In K. M. Stanney (Ed.) *Handbook of virtual environments: Design, implementation, and applications* (pp. 493-518). Mahwah, NJ: Lawrence Erlbaum.

Darken, R. P., & Sibert, J. L. (1993). A toolset for navigation in virtual environments. *Proceedings of ACM User Interface Software & Technology (UIST '93*, pp. 157-165). New York: ACM.

Dave, P., Karadkar, U. P., Furuta, R., Francisco-Revilla, L., Shipman, F., Dash, S., & Dalal, Z. (2003). Browsing intricately connected paths. *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT'03*, pp. 95-103). New York: ACM.

De Roure, D., Hall, W., Reich, S., Hill, G., Pikrakis, A., & Stairmand, M. (2001). MEMOIR - an open framework for enhanced navigation of distributed information. *Information Processing and Management, 37*, 53-74.

Elvins, T. T., Nadeau, D. R., Schuk, R., & Kirsh, D. (2001). Worldlets: 3-D Thumbnails for Wayfinding in Large Virtual Worlds. *Presence: Teleoperators and Virtual Environments, 10,* 565-582.

Furuta, R., Shipman III, F. M., Marshall, C. C., Brenner, D., & Hsieh, H. (1997). Hypertext paths and the world-wide web: Experiences with Walden's paths. *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT'97*, pp. 167-176). New York: ACM.

Gamberini, L., Cottone, P., Spagnolli, A., Varotto, D., & Mantovani, G. (2003). Responding to a fire emergency in a virtual environment: different patterns of action for different situations. *Ergonomics, 46*, 842-858.

Gams, E., & Reich, S. (2004). Following your colleagues' footprints: Navigation support with trails in shared directories. *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT'04*, pp. 89-90*)*. New York: ACM.

Grammenos, G., Filou, M., Papadakos, P., & Stephanidis, C. (2002). Virtual prints: Leaving trails in virtual environments. *Proceedings of the Eighth Eurographics Workshop on Virtual Reality (EGVE'02*, pp. 131-138). Aire-la-Ville, Switzerland: Eurographics Association.

Hartley, T., Maguire, E. A., Spiers, H. J., & Burgess, N. (2003). The well-worn route and the path less traveled: Distinct neural bases of route following and wayfinding in humans. *Neuron, 37*, 877-888.

Hightower, R. R., Ring, L. T., Helfman, J. I., Bederson, B. B., & Hollan, J. D. (1998). Graphical multiscale web histories: A study of PadPrints. *Proceedings of the ACM Conference on Hypertext and Hypermedia (HT'97*, pp. 58-65). New York: ACM.

Krane, D. E., & Raymer, M. L. (2003). *Fundamental concepts of bioinformatics*. New York: Pearson Education.

Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., Robinson, P., & Sellen, A. (1994). The design of a human memory prosthesis. *The Computer Journal, 37*, 153-163.

Lessels, S., & Ruddle, R. A. (2005). Movement around real and virtual cluttered environments. *Presence: Teleoperators and Virtual Environments, 14*, 580-596.

Lynch, K. (1960). *The image of the city*. Cambridge: MIT Press.

Milic-Frayling, N., & Sommerer, R. (2003). WebScout: Support for revisitation of web pages within a navigation session. *Proceedings of IEEE/WIC Web Intelligence 2003* (pp.689-693). Los Alamitos, CA: IEEE.

O'Hara, K., Morris, R., Shadbolt, N., Hitch, G. J., Hall, W., & Beagrie, N. (2006). Memories for life: A review of the science and technology. *Journal of the Royal Society Interface, 3*, 351-365.

Peponis, P., Conroy Dalton, R., Wineman, J., & Dalton, N. (2004). Measuring the effects of layout upon visitors' spatial behaviors in open plan exhibition settings. *Environment and Planning B: Planning and Design 31*, 453-473.

Pettifer, S., Cook, J., & Marsh, J. (2004). PlaceWorld, and the evolution of electronic landscapes. In D. N. Snowdon, E. F. Churchill, & E. Frécon (Eds.) *Inhabited Information Spaces* (pp. 25-50). London: Springer-Verlag.

Reich, S., Carr, L., De Roure, D., & Hall, W. (1999). Where have you been from here? Trails in hypertext systems. *ACM Computing Surveys, 31*(4es).

Ruddle, R. A. (2005). The effect of trails on first-time and subsequent navigation in a virtual environment. *Proceedings of the IEEE Virtual Reality Conference (VR'05*, pp. 115-122). Los Alamitos, CA: IEEE.

Ruddle, R. A. (2006). Using String-matching to Analyze Hypertext Navigation. *Proceedings of the 17th ACM Conference on Hypertext and Hypermedia (HT'06*, pp. 49-52). New York: ACM.

Skopik, A., & Gutwin, C. (2005). Improving revisitation in fisheye views with visit wear. *Proceedings of the Computer Human Interfaces Conference (CHI'05*, pp. 771-780). New York: ACM.

Smith, F. F., & Waterman, M. S. (1981). Identification of common molecular sequences. *Journal of Molecular Biology, 147*, 195-197.

Utting, K., & Yankelovich, N. (1989). Context and orientation in hypermedia networks. *ACM Transactions on Information Systems, 7*, 58-84.

Waller, D. (2000). Individual differences in spatial learning from computer-simulated environments. *Journal of Experimental Psychology: Applied, 6*, 307-321.

Wexelblat, A., & Maes, P. (1999). Footprints: History-rich tools for information foraging. *Proceedings of Computer Human Interfaces Conference (CHI'99*, pp. 270-277). New York: ACM.

White, R. W., & Drucker, S. M. (2007). Investigating behavioral variability in web search. *Proceedings of the 13th International World Wide Web Conference (WWW'04*, pp. 21-30). New York: ACM.

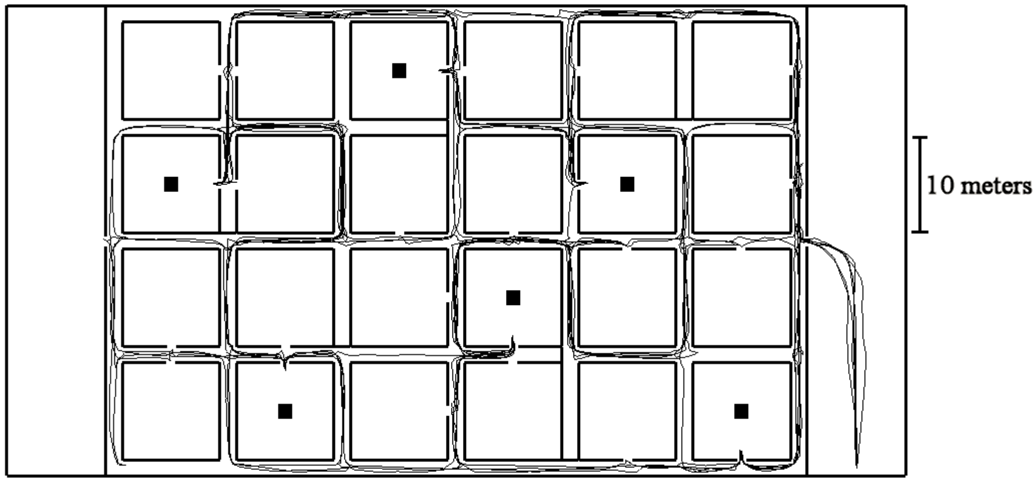

Figure 1. *Trail pollution caused by displaying every path a person has ever followed. The data are from the experiment described in this paper, and are for one participant's familiarization sessions. The primary trail that our method generates from these paths is shown in Figure 4b.*

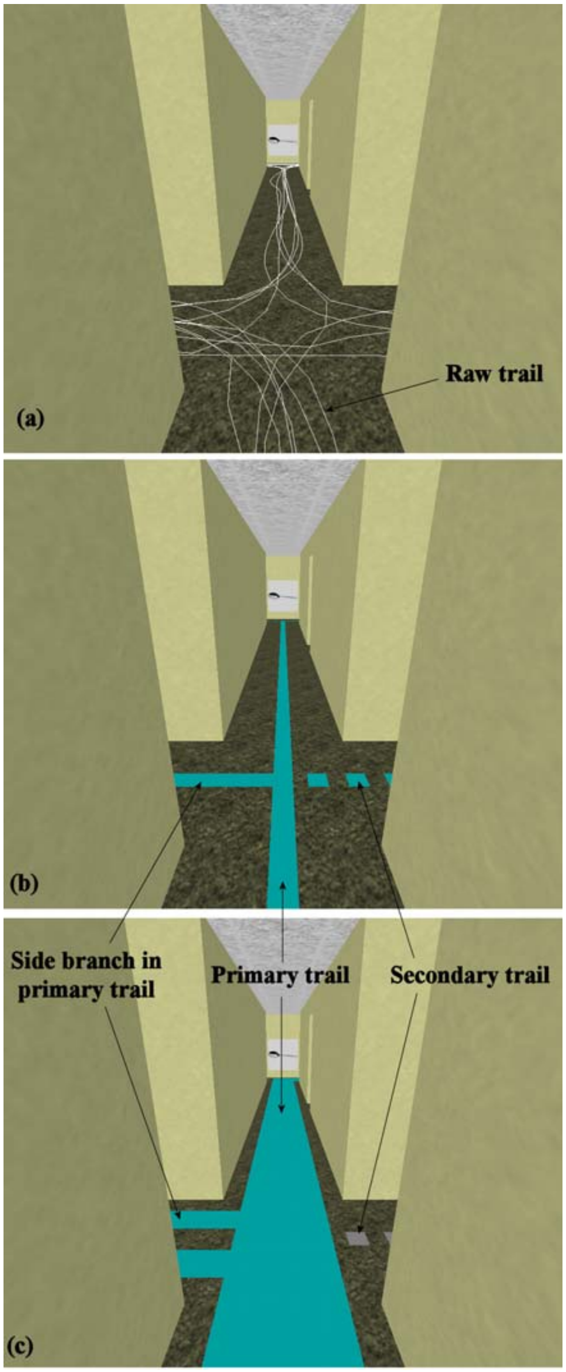

Figure 2. *Three methods of presenting the same movement data: (a) raw trails showing a person's actual movements (this was used as a control condition in the experiment), (b) a method that made it difficult to identify branches in the primary trail, and (c) a solution that also increased the distinction between the primary and secondary trail.*

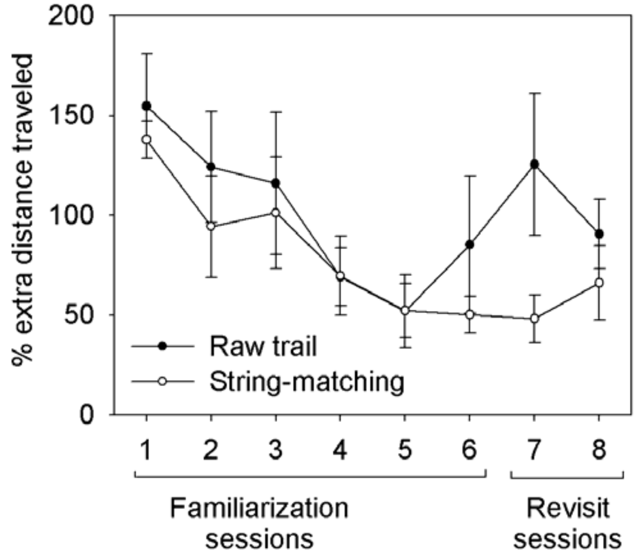

Figure 3. *Percentage extra distance traveled in excess of shortest possible route connecting the start point to all six targets. Error bar show standard error of the mean.*

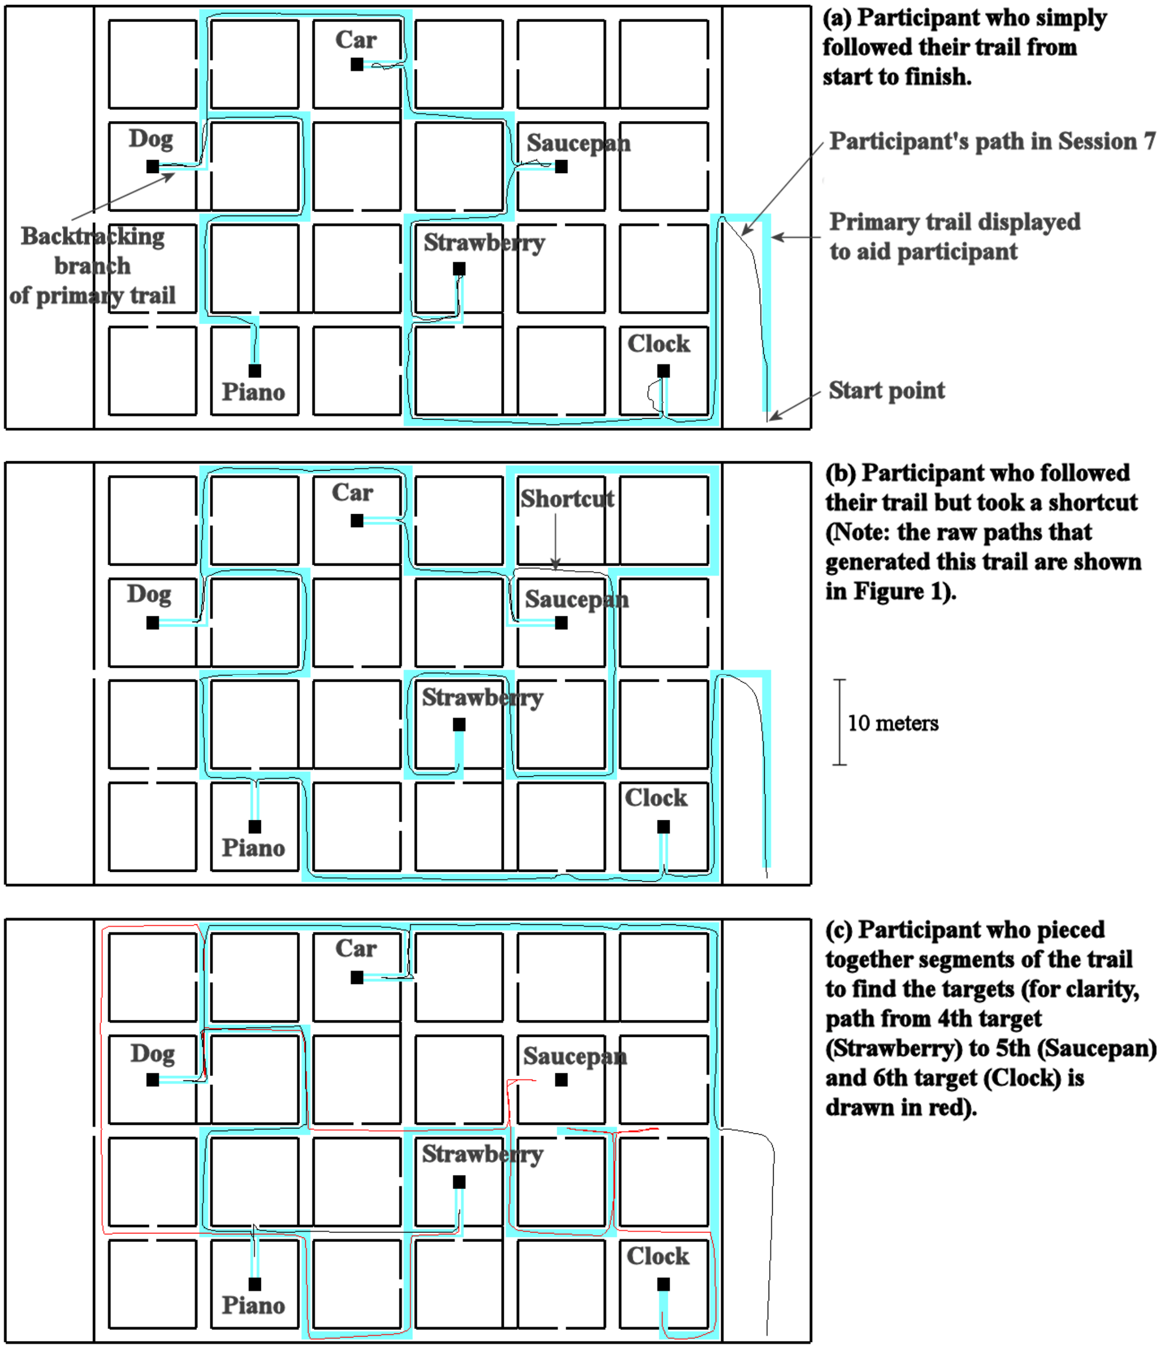

Figure 4. *Examples of paths followed in Session 7 by participants in the string-matched group. Only a participant's primary trail is shown (the secondary trail, which showed all the other places the participant had visited, is omitted for clarity).*

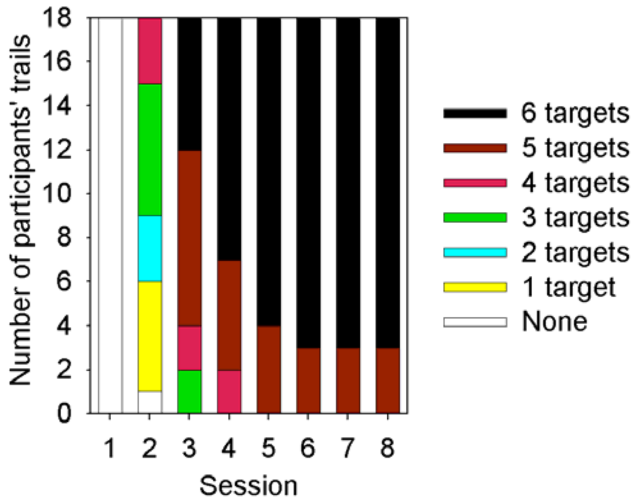

Figure 5. *Number of trails that connected a given number of targets (trails calculated from movement data for all 18 participants in string-matched and raw groups who returned for the revisit sessions (Sessions 7 & 8)).*