Audio Display of Spatial Information:

Analyzing Movement Strategies in Exploration of Thematic Maps

Megen Brittell

September 20, 2013

Abstract

Auditory displays can improve accessibility of geospatial data and geographic information systems (GIS) for people who are blind. This paper characterizes and quantifies patterns in stylus movements that users who are blind employed to interact with a minimal Geographic Information System (mGIS) and complete a Region Lab educational intervention. Analysis of data collected during behavioral testing that represents participants responses to two tasks (exploration and selection) reveals a tendency for participants to move the stylus along orthogonal axes and identifies a *check-neighbors* gesture that represents participants actions in response to the two tasks. A GOMS-style model describes observed patterns and variation in performance across the two tasks and across repetitions within each task are explored in a quantitative analysis.

1 Introduction

An increasing volume of geospatial data that describes and annotates the world around us is available for public consumption. Maps that present such data are commonly used to teach spatial concepts [13, 15] and to illustrate spatial trends that influence policy making (especially in combination with geographic information systems, e.g. [16, 28]). Traditional printed maps and graphical displays for software systems that process spatial data render the information inaccessible to people who are blind or low vision. Limited access to geospatial data in both raw and interpreted forms limits opportunities to participate in educational, professional, and civic activities.

A three year collaborative project titled "Exploiting the Power of GIS to Enhance Spatial Thinking" (henceforth, "Spatial Thinking Project") explored spatial thinking skills among map readers who are blind.

of the software implementation are given in Section 6. The paper concludes with a summary of the findings and potential future directions (Section 7).

2 Related Literature

The study presented in this paper relies on many disciplines and fields that support the study of accessible displays of spatial data from psychoacoustics and spatial cognition to human computer interaction. This section focuses on literature from two specific areas: accessible displays of spatial data and gesture recognition.

2.1 Approaches to Accessible Map Displays

Recognizing the importance of accessible displays of spatial data, researchers have explored ways to provide non-visual feedback to users who are blind. Instead of relying on vision alone, computer displays can target alternative modalities to improve accessibility. Two common approaches target the human somatosensory system, the senses of touch and proprioception, and auditory displays. Tactile graphics represent an effective way to convey map data to users who are blind. Static production methods include printing on microcapsule paper, embossing, and, more recently, fabrication on 3D printers. These production methods, however, create materials that may be inconvenient to carry in bulk. Dynamic tactile displays are under investigation, but are not yet widely available outside the research community (e.g. BrailleDis 9000 a pin-matrix [30], MudPad actuated by electromagnets [17], and TeslaTouch electrostatic feedback [4]). Haptic devices provide force feedback when moving the point of a stylus over a virtual three dimensional surface (e.g. Geomagic Touch Haptic Device,[1] previously SensAble Phantom) or through vibrations (e.g., Nintendo Wii Remote Plus controller[2]). While haptic feedback can convey spatial location through proprioception and an attribute value through texture, providing sufficiently high resolution spatial information to represent map data is still a challenge.

Communicating two dimensional spatial data[3] through a temporally linear auditory display is non-trival and users have difficulty understanding the overall layout and patterns of data [1, 8]. Research prototypes have instead used haptic devices (e.g., Omero [6] and Haptic Soundscapes [27]) and tablet with stylus input

[1] Geomagic, `http://www.geomagic.com/en/products/phantom-omni/overview`

[2] Nintendo, `http://www.nintendo.com/wii/what-is-wii/#/tech-specs`

[3] The intention in this discussion is to explore displays for users who are blind or low vision. While auditory and haptic feedback have also been explored to augment a visual display (e.g., to augment "visually dominant geographic information systems" [18]), such displays still rely strongly on vision and do not address accessibility needs for users who are blind.

those movements occur. The analysis presented in this paper looks again at trends in direction of movement by evaluating the proportional frequency of observed directions, measuring direction with higher resolution than was reported in previous studies and considering speed of the movements.

While existing accessible map displays successfully communicate some aspects of spatial data, previous research reports that users of audio interfaces have difficulty understanding the overall layout and patterns of data [8, 1]. The analysis described in this paper looks at the actions that users employ to respond to two tasks and how those actions may reflect or influence synthesis of sequential auditory feedback into an internal (mental) representation of spatial data. Methods to describe overall trends in movement direction are presented in Section 4.

2.2 Gesture Recognition

As an input method for touch devices, the term *gesture* often refers to actions that are distinguished by their location, duration, shape, and features (e.g., [2, 32, 20, 29]). Gesture recognition algorithms come in a wide variety from statistical classifiers (e.g., [29]) to template-based point location matching (e.g., the $1 recognizer [32]), and are chosen based on the input device [2] and maximum latency requirements [32]. Some implementations also incorporate gestures that facilitate dynamic feedback during execution of the gesture (e.g. responding to intermediate states of gesture execution [23]) and research has explored ways to reduce the amount of time that there is ambiguity in interpreting gesture input [31].

The analysis in this paper uses the term gesture to refer to any sequence of movements that satisfy an unordered set of content-specific features that are defined by relative location and attribute value. The inferred intention behind the *check-neighbors* gesture is to determine that no neighbors of a target county share a common attribute value. This condition can be determined by visiting all neighbors at least once each,[6] but a target county can also be rejected if a neighbor with a similar attribute value is observed prior to checking all neighbors (the feature set is violated). Existing gesture recognition algorithms detect various components of this gesture (see Table 1), but we are not aware of a single tool that combines all of the necessary components. An algorithm to identify this gesture must be independent of the order in which the features are satisfied (similar to Rubine's algorithm [29]), independent of the exact shape of the gesture and repetition within the gesture (similar to Gesture Coder [23]), and aware of the location of the cursor relative to display contents (e.g., the commercial product Swipe[7]). Further, it must detect the gesture

[6]When performing the *topology-check* gesture, participants in the behavioral testing sessions visited just four neighbors. This observation is discussed further in Section 6.1.1.

[7]Nuance, `http://www.swype.com`

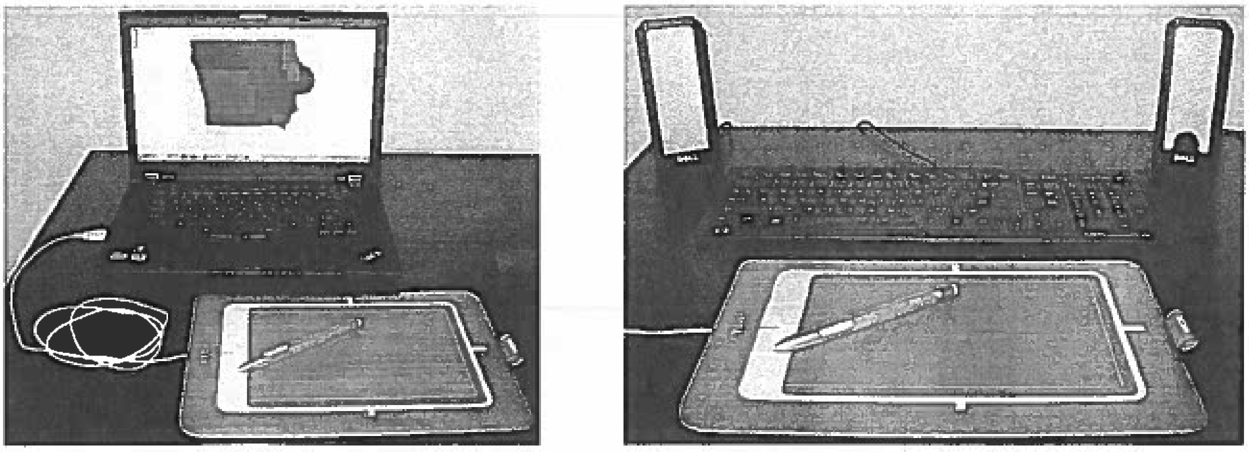

Figure 1: The hardware required for the minimal geographic information system (mGIS) consists of a laptop and a tablet and stylus (left). A full size keyboard and external speakers are optional additional hardware. The tablet, stylus, keyboard, and speakers (right) represent the hardware with which participants interacted during the behavioral testing.

a laptop (with built-in speaks and keyboard) and a tablet with stylus input device (Figure 1, left). The tablet used during behavioral testing sessions is a Bamboo Fun^{TM} Pen & Touch tablet (active area 7.5 x 5.1 inches, 190 x 130 mm) with four hardware buttons on the left-hand side. Users press one of the four buttons to select map elements in the Region Lab (see Section 3.4). Moving the stylus on the tablet provides proprioceptive (kinesthetic) feedback and elicits auditory feedback.

During behavioral testing, the laptop screen is turned away from the participants to minimize any visual cues perceived by participants who have residual vision. In this arrangement the built in speakers also face away from the participant, effectively reversing the relative position of the left and right speakers. External speakers are connected to the laptop and positioned on the table facing the participants (Figure 1, right) to correctly render the horizontal panning.[9] A full size keyboard is also used during the testing sessions (Figure 1, right). Users give keyboard commands to access menu options and move the stylus across the active area of the tablet to interact with the map. The primary feedback from keyboard input is verbal (delivered through text-to-speech; see 3.2).

3.2 Software

As a research instrument, software design was subject to the constraints of an external experimental design (comparison of the effects on spatial thinking skills of two educational interventions - one using tactile graphics and the other an auditory display). Owing to these constraints, the experimental design did not

[9]The built-in speakers do not emit sound when the external speakers are connected.

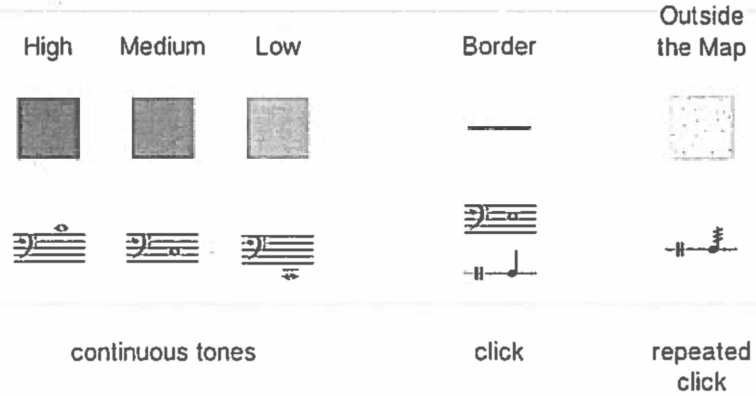

Figure 2: Map symbology specifies the characteristics of elements in both the visual and auditory displays. The attribute value (population density) are represented as a fill color with varying shades of green and as a continuous tone with varying frequency (left). The boundaries are shown with a black line and, in the auditory display, are represented with both a continuous tone (when the stylus is near the boundary line) and a short duration click (when the stylus crosses the boundary line; center). The area outside the map is symbolized with a grey fill color and dot pattern and as a series of short duration clicks (right).

plays only when the stylus is touching the tablet, in motion, and recently moved across a border. Due to the dependence on the stylus being in motion, the click symbol is sensitive to any latency in the feedback, particularly when the stylus is moving quickly. The click symbol is rendered using the "woodblock" instrument from the default MIDI soundbank in the Java Sound API. The MIDI woodblock sound provides short attack and decay durations, while still having a library-implemented smooth transition between the note sounding and not sounding. A continuous tone augments the boundary click symbol, indicating whether or not the stylus is close to a boundary (within 10 map units of the boundary line, approximately 7 mm) when the stylus was stationary.[15]

Continuous tones (symbolizing both attribute value and location of the cursor near a boundary) are synthesized using the Java Soundbank *Syn Square Wave* instrument. The choice to use the square wave to synthesize a continuous note provided a complex waveform (additive combination of multiple harmonics; recommended by Patterson [26]), but at the same time allows us to postpone decisions about how to select a specific timbre. The three levels of the attribute (e.g. low, medium, and high) were separated by octave intervals using notes C2 (approximately 65 Hz), C3 (approximately 130 Hz), and C4 (approximately 261 Hz) along with the associated harmonics that comprise a square wave. The tone symbolizing locations near a boundary line sounded at E3 (approximately 165 Hz), creating an interval of a major third with attribute values that may play simultaneously. These symbols were chosen based on feedback from collaborators who

[15]The continuous tone was chosen in favor of additional clicks at the edge of a buffer around the boundary line, such as that proposed by Evreinova for mathematical line graphs [11], specifically so that participants received feedback about proximity to the line regardless of whether or not the stylus was moving.

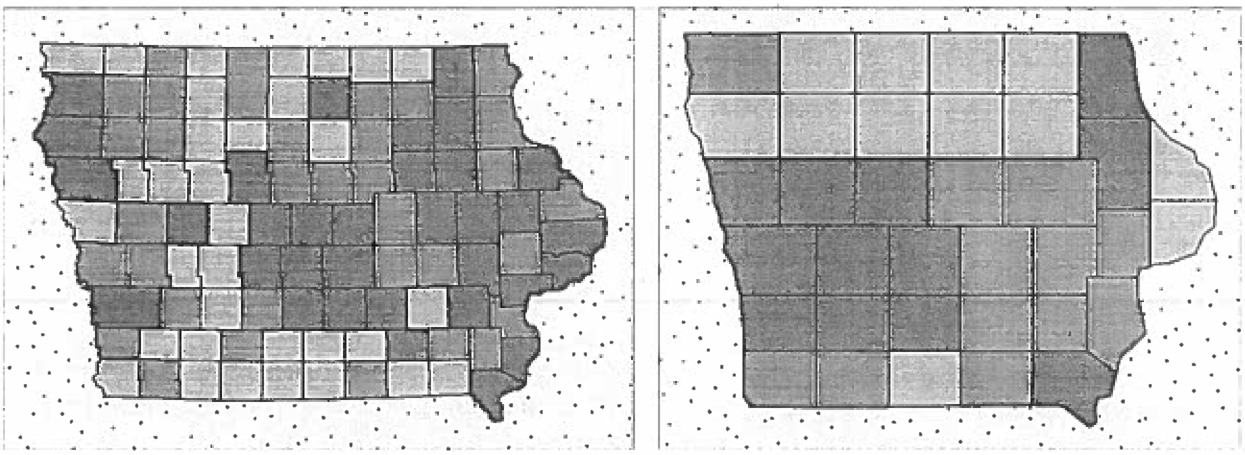

Figure 3: The counties of Iowa[20] (left) inspired the layout of the sample data set used in the Region Lab (right). Population density is shown at intensity of the fill color, classified in three levels: high, medium, and low.

The sample data was inspired by real-world data (for external validity) and modified for experimental control. The counties in the state of Iowa were chosen for their relatively simple shape (primarily straight county boundaries that largely lie on a horizontal/vertical grid).[18] The extent of the map data was marked by a simplified version of the Iowa state boundary. The number of counties was reduced, modifying the original layout of county boundaries from 99 to 36 to further simplify the display and meet experimental balance constraints (e.g., the number of counties was a multiple of three). Population densities were fabricated for the experiment. This accommodated the experimental constraints that the three levels of population density were approximately equivalent. Figure 3 shows the true counties of Iowa (left) and the fictitious data used in the experiment (right). Reflection on this choice of experimental data is presented in the discussion section (Section 6.1).

The mGIS provides a minimal set of spatial analysis tools (supported by functions from the GeoTools library). These tools are available through application menus (and submenus) and implement concrete examples of the four spatial concepts that, in combination, we posit make up the complex concept of *region*: classification, proximity, cluster, and boundary. The *Classification* tool applies (or removes) a filter to the map contents based on values in the attribute table of the spatial data file. In the sample data, this attribute is population density classified as one of three levels (low, medium, or high). The *Proximity* and *Cluster* tools similarly filter display contents, but instead of using values from the original attribute table, these filters are applied to values computed during a preprocessing step that extracts topological relationships between the

[18]The regular roughly square shape of the counties in Iowa are a result of using range and township lines to determine county boundaries: `http://historical-county.newberry.org/website/Iowa/documents/IA_Commentary.htm`

patterns of interaction with the software based Region Lab in an exploratory analysis takes advantage of that log data collected during behavioral testing.

4 Methodology

Stylus movements constitute an observable measure of participants' actions in response to the tasks in the Region Lab (Section 3.4). The analysis presented in this paper focuses on exploration and selection associated with the concept of proximity. Data was collected in log files during behavioral testing with the mGIS. First, preprocessing isolates the stylus movements associated with the tasks of interest and resamples the data to minimize artifacts of a mismatch between the speed with which participants moved the stylus and the sample rate and sampling resolution. The transition between the two target tasks was not included in the log file for two of the four occurrences. A keystroke level model (KLM) is used to infer the position in the log at which the omitted transitions were likely to have taken place.

Analysis of patterns of stylus movement captured in the logs addresses the first research question: How can observed trends in stylus movement be characterized? The two tasks are broken down with goals, operators, methods, and selections rules (GOMS) models to formalize anticipated and observed behaviors. The GOMS models of a *check-neighbors* gesture is translated into a state-based representation of the actions, which in turn is applied to automate detection of the gesture in log data. Trends in movement are also evaluated by representing speed and direction[22] in polar coordinates, performing a point pattern analysis, and identifying peaks in the kernel density estimate (KDE). The GOMS model, the identified gesture, and the number of peaks in the KDE serve as three ways to characterize stylus movements.

Quantitative measures are used to explore differences between the two tasks and over time in response to the second research question: How do movement strategies differ between tasks and over time? The quantitative analysis includes metrics used in a previous study with iSonic [8] as well as the output from the point pattern analysis. Computed values are used in linear models to determine the significance.

4.1 Goals, Operators, Methods, and Selection Rules (GOMS) Model

This section proposes a theoretical model of the processes involved in completing two tasks in the Region Lab: exploration and selection (see Section 3.4). It first describes the two tasks, providing details of anticipated and observed responses. Next, it introduces a GOMS-style model (goals, operators, methods, and selection

[22]Direction of movement is recorded relative to the active area of the tablet. Participants were free to adjust the position of the tablet (and keyboard) so the orientation of the tablet relative to the participant varied across participants.

feedback indicating that it was incorrect and asks the participant to try again. The session does not advance until a correct selection is made. All participants ultimately made a correct selection.

The training interface did not teach any specific movement strategy. Participants were instructed to touch the stylus to the tablet and move the stylus while touching the tablet (`drag`) to elicit feedback. Each auditory symbol was named, but participants were not given explicit instructions regarding how to interpret the stream of auditory feedback. The instructions given to participants for this task were intentionally vague to avoid biasing their self-selected strategies for extracting information from the auditory display. The Region Lab requires users to query the display at least once during the exploration task. This means that although participants may have been satisfied with their mental map from previous exploration (or deferring exploration until they knew the next objective, see also Section 6.1.2), participants had to at least touch the stylus to the tablet before the Region Lab advances to the next task.

While participants independently developed movement strategies, their choices in adopting strategies that governed movements depended on the task and there were similarities in the movement patterns across participants. A formalized model of the procedural knowledge involved in completing the two tasks contributes to our understanding of the interface's usability. Different movement patterns provide varying levels of effectiveness in addressing the objectives of the two different tasks. To articulate the observed movements and reason about their relationship with the respective tasks, each task is decomposed into a series of goals, operators, methods, and selection rules.

4.1.2 Models

The goals, operators, methods, and selection rules (GOMS) model, introduced by Card, Moran, and Newell in 1983 [5], represents a way to formally describe user behavior. GOMS models have been applied in the field of human-computer interaction to describe interaction with a software interface with varying levels of granularity. As part of the analysis described in this paper, a GOMS model is used to think through and reason about the possible tasks and subtasks involved in completing the Region Lab. Although participants' performance in the map-based tasks does not qualify as error free (see Section 5.1.5) or even "routine cognitive skilled behavior," the GOMS models help organize our understanding of the task.

One complication in the model of the exploration task is that there was no explicit objective or proficiency that participants were required to attain to complete the task. Movements may be executed and repeated any number of times. Such non-deterministic behavior adds complexity to the model, but with sufficient granularity (and practice!), map exploration may still be a skilled activity that conforms to many of the direction that did not follow an obvious pattern (in contrast to the regular reversals of a zig-zag movement). These patterns reflect participants strategies to structure their responses to the exploration task and guide their movements in the selection task.

A GOMS-model is used to formalize the observed actions. It describes stylus movement strategies as the combination of a gross movement pattern and a fine movement pattern. The gross movement patterns describe generalized trends in movement across the active area of the tablet. The fine movement patterns represent higher resolution variation in stylus movement. In addition to this model that describes exploration behaviors, a key stroke level model – a specific application of the broader GOMS model – is used to predict the amount of time required to make a menu selection (see Section 4.4.1).

4.1.3 Translate to state model

From the GOMS model of fine resolution movements, a state based model of a *check-neighbors* gesture is defined and applied to identify occurrences of the gesture in the streams of sampled cursor locations. Without trying to model memory capacity, the state based model assumes that during a fine movement pattern the user can remember only two features beyond those that belong to a candidate gesture (i.e., knowledge from one candidate gesture does not translate to another; see Section 4.2, Figure 5).[25] This assumption applies to recalling spatial attributes during the selection and neither applies to the exploration task nor precludes users' ability to match temporal auditory feedback with a pre-planned linear sequence.

4.2 Gesture Identification

Several participants developed an ordered sequence of movements in response to the instruction to "select a county with no similar neighbors," i.e., identify a county based on the relationship between its attribute value and those of adjacent counties. Participants who employed this approach located a county with the desired attribute level, then moved the stylus laterally (right and left) and radially (up and down), receiving feedback related to the attributes of the adjacent counties. Such a cross-shaped gesture represents an approach to interrogating the attribute-based topology of the counties. A state-based model was developed to automate the process of objectively identifying this *check-neighbors* gesture within the stream of sampled stylus locations in each log file.

[25]Attention influences memory [3], and the selection task focuses topology in the immediate vicinity (not the broader spatial layout). During a gross movement pattern, users may be more likely to synthesize various parts of the temporal auditory stream into a mental representation that incorporates two dimensional spatial components, as evidenced by the participants' ability to return to previously explored areas of the display when prompted by a description of the display elements that occupied that space in the verification task (Section 3.4).

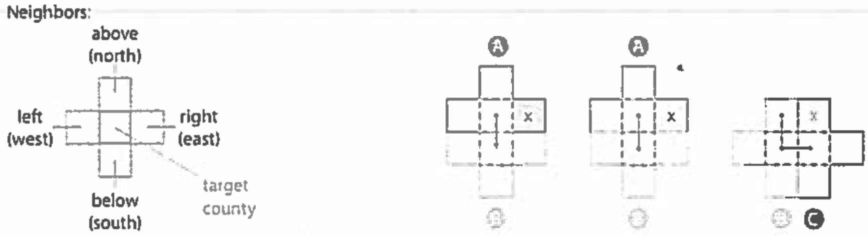

Figure 5: Relative to a target county in a candidate *check-neighbors* gesture, there are four neighbors (left). The attribute of each neighbor determines the values of the four features in the automated gesture recognition. The state model assumes that participants' memory of features does not transfer from one candidate gesture to another (right). For example, the county marked with an x may have been visited during candidate gesture A. When the stylus moved down and to the right, however, the value of county x is unknown. The stylus would need to revisit county x to satisfy the *above* feature for gesture C.

4.3 Behavioral Testing

Researchers from both Geography and Computer Science traveled to national conferences of organizations for the blind (American Council for the Blind and National Federation of the Blind) to conduct testing. The data collected during these test sessions provide an objective quantitative record of participants' interaction with the mGIS. An exploratory analysis identifies and characterizes patterns in this data.

Nine participants (age range 29-61, mean = 44) completed the software based intervention (i.e. the Region Lab implemented in the mGIS). These participants were recruited among attendees of the National Federation of the Blind annual conference. Seven participants were blind from birth, one became blind at age six, and one became blind at age 15. Use of the JAWS screen reader was included in the recruiting criteria to ensure that interaction with the screen reader would not confound interaction with the research instrument software interface. All participants were frequent computer users (daily use) and had between 7 and 21 years of experience using JAWS (mean = 14 years). Including JAWS proficiency among the recruiting criteria also had the effect of targeting participants who did not rely on screen magnifiers or residual vision for interaction with computers. Two participants reported having used a tablet input device prior to the test session. Seven participants had previous music experience and two reported having perfect pitch.

4.4 Pre-processing

In addition to general formatting, pre-processing includes two major steps. First, the the logs are partitioned based on the task to which participants were responding when the sampled stylus positions were recorded. Second, the log files are resampled to mitigate the disproportionate number of observed directions that fall required could represent the interval in which the transition actually occurred. A potential confound to this approach is that long intervals also occurred if the participant lifted the stylus from the tablet when they were asking a question or just pausing the feedback to think for a moment. Although there is a risk of incorrectly interpreting a long interval as a transition, the number of occurrences of long intervals within the data set is small (see Section 5.2.1).

Relative navigation using the arrow keys and keyboard short-cuts are two standard strategies for menu interaction with JAWS. Two KLM models capture these strategies and produce lower-bound estimates of the execution time for this menu selection. These estimates are compared against the observed intervals between stylus input actions (i.e. during which the stylus was not in contact with the tablet) to determine when the transition occurred and to partition the log files into subsets that each contain sampled stylus positions in response to exactly one task – either exploration or selection.

4.4.2 Resample Log Data

The speed at which participants moved the stylus was relatively slow compared to the spatial resolution of the tablet with stylus and the high sample rate. As a result, sequential samples tend to lie across pixel boundaries. When considering direction of movement, the high sample rate creates a bias in the observed directions of movements toward the directions orthogonal to pixel boundaries. To mitigate this artifact of the sample rate, the data is resampled at a rate with a fewer samples per second. Using an integer multiple of the original rate means that all points included in the output are also members of the input set (a strict subset). This avoids the need to interpolate spatially and temporally between input points (as was done, for example, to mitigate differences in gesture execution speed in the $1 recognizer [32]).

4.5 Quantitative Metrics from Related Work

The proportion of enumeration units visited, number of steps (where *step* is defined as a transition from one enumeration unit to its neighbor [8]), and direction consistency (i.e., *keep-a-direction* steps that are defined as "keeping a clear direction for more than two steps" [8]) provide metrics with which to compare the two tasks (exploration and selection). Differences in context and task instructions, however, prevent a direct comparison of the results across the studies. For example, the states displayed in the iSonic interface used in the Delogu study were contiguous [8], and the subsets of counties using in the Region Lab were not: after applying filters as part of the spatial anlayisis tools, some counties were removed from the display leaving a blank area. As a result, the way that steps and keep-a-direction steps are counted differ.

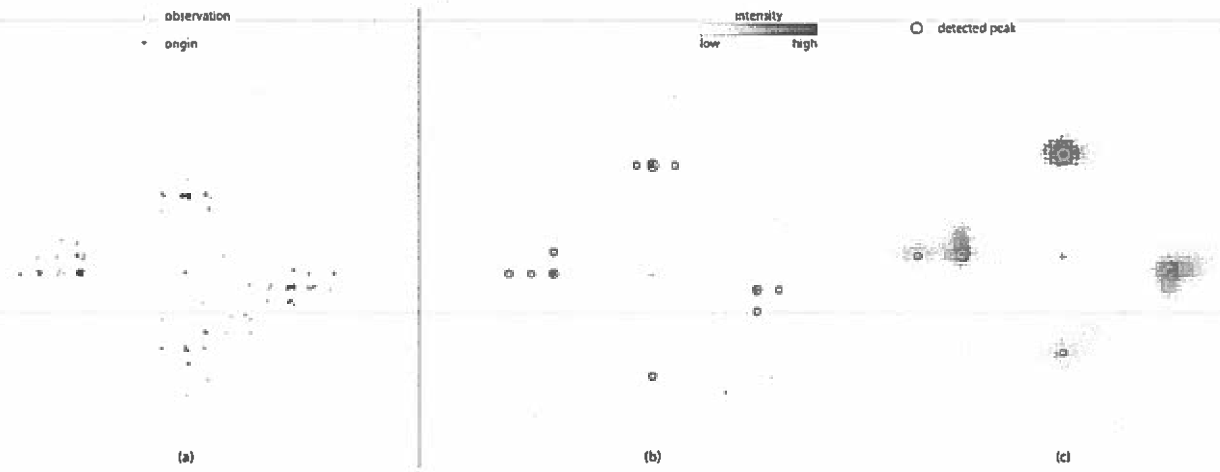

Figure 6: The approach to bandwidth estimation selected a value that minimizes the mean squared error. As shown in this example (P276, step 40), the overall minimum captures fine scale variation in intensities (bandwidth=0.0098, left). To capture the more general pattern, the bandwidth was increased by a factor of 50 (bandwidth=0.4925,right).

depending on the number of samples in the log file. The heights of the peaks are not directly comparable across log files (even within participant); the number of peaks is used to compare trends in direction of movement across tasks and over time. The threshold for classifying local maxima as peaks was set at 20% of the maximum intensity value in the kernel. Both the scale factor for the bandwidth (above) and the peak threshold were selected based on visual inspection of the results and qualitative judgement that the output was representative of the underlying data. While the values varied from one observation to the next, the process for determining the values was systematic.

To address the second research question (How do movement strategies differ between tasks and over time?), the number of peaks found in the the output of the kernel density estimate and three metrics defined in the literature (number of steps, proportion of *keep-a-direction* steps, and visit rate) are used to explore differences between the two tasks. Linear regression approximates any change over time (i.e., the three repetitions of the explore + selection pair of tasks). A paired t-test is used to measure the strength of the difference.

5 Results

Following the methodology outlined in the previous section, the pre-processing produces an estimate of the time required to execute a menu selection and conducting the analysis provides the number of occurrences of the *check-neighbors* gesture, quantifies trends in direction of movements, and computes measures of statistical

Task_item: Menu
    mode is `keyboard`
    keyboard_strategy is `direct_key_bindings | traverse_menu`

Task_item: Map
    mode is `stylus`
    stylus_strategy is `tap | drag`

Task_item: Select
    mode is `hardware_button`

Figure 7: Three GOMS-style task items define the hardware interaction necessary to complete the various tasks in the Region Lab.

Selection Rules for Goal: **Set Input_Device**
    If `<mode>` of `<current_activity>` is *keyboard*, Then
        Step: Move hand(s) to keyboard
    Else, If `<mode>` of `<current_activity>` is *stylus*, Then
        Step: Pick up stylus
        Step: Move to `<initial_location>` of `<current_movement_strategy>`
    Else, If `<mode>` of `<current_activity>` is *hardware_button*, Then
        Step: Hold stylus position constant
        Step: Move hand to corner of the tablet
    Return_with_goal_accomplished

Figure 8: Tasks in the Region Lab require users to switch between different hardware components including the keyboard, the stylus, and the tablet (the active area and the hardware button).

```
Method for Goal: Perceive Auditory_feedback
    Step: Determine <number_of_pitches> in feedback
    Step: Determine <pitch> of feedback
    Step: Determine <number_of_clicks> in feedback
    Decide: If <number_of_pitches> is two, Then
            Step: Store true under <is_near_boundary>
        Else, Step: Store false under <is_near_boundary>
    Decide: If <number_of_pitches> is greater_than_zero, Then
            Step: Store <pitch> under <pitch_class>
    Decide: If <number_of_clicks> is greater_than_zero, Then
            If <action_strategy> of <user> is drag, Then
                Step: Store true under <crossed_boundary>
    Accomplish Goal: Update Mental_representation
    Return_with_goal_accomplished

Method for Goal: Update Mental_representation
    Accomplish Goal: Perceive Auditory_feedback
    Step: Perceive kinesthetic feedback
    Step: Integrate auditory and kinesthetic feedback to interpret and encode stimulus
    Step: Append perceived auditory event to auditory working memory
    Step: Synthesize new input with contents of the visuo-spatial sketchpad
    Return_with_goal_accomplished
```

Figure 10: Although not an exhaustive treatment of the cognitive processes involved in using the auditory display, perceiving the pitch, number of pitches, and number of clicks are critical to interpreting the auditory symbology in the mGIS. Similarly, the integration and of multi-modal feedback and synthesis that feedback with information in memory is essential.

The details of cognitive processes related to auditory perception are beyond the scope of this investigation, but the elements of that process are thought to be involved and that influenced design decisions for the mGIS are listed among the operators in Table 10. Similarly, the processes involved in storing and modifying a mental representation of spatial data is beyond the scope of this project, but the model includes a brief listing of basic steps that may be involved to help guide our reasoning about user behaviors (Table 10).

5.1.2 Gross Movement Patterns

As a generalized model, each iteration of the movement strategy consisted of a change in orientation (a turn) and a change in location (a movement). The amount of change is determined by a combination of a specific gross movement pattern and a fine movement pattern (see Section 5.1.3). The patterns included in the model were informally observed during the behavioral testing sessions. Due to the cognitive processing involved in making sense of the spatial data presented in the auditory display, this model is insufficient to estimate the time required to complete the tasks or determine whether or not a pattern is optimal. Instead

Gross_movement_item: Undirected
Name is `Unsystematic_order`
Initial_direction is any
Delta_direction is

$$\begin{cases} \texttt{along\_edge\_or\_toward\_center} & : \texttt{<current\_location>} = \textit{on\_edge} \\ \texttt{toward\_tablet} & : \texttt{<current\_location>} = \textit{off\_the\_tablet} \\ \texttt{any\_direction} & : \textit{otherwise} \end{cases}$$

Initial_location is `random`
Step_distance is `random`

Figure 12: Participants tended to keep the stylus on the tablet (reflected by the edge checks in this item), but movement patterns within the extent of the active area of the tablet did not always display obvious organization.

Gross_movement_item: Direct
Name is `Direct_pointing`
Initial_direction is `[direction of target relative to current location]`
Initial_location is `<current_stylus_location>`
Delta_direction is *zero*
Step_distance is `[distance to target]`

Figure 13: A straight path from the current location of the stylus to a known target location is the most efficient gross movement pattern.

refine stylus position. In practice, however, participants used gross movement patterns to locate the target counties. Fine movement patterns were then used to augment the search or investigate the characteristics of the elements in the display that were close to the current stylus position.

5.1.3 Fine Movement Patterns

Fine movement patterns can be used to augment a gross movement pattern or to query the display for information that could inform a decision about the display contents. During a gross movement (i.e., movement in a general direction), some participants used repeated short movements (Figure 14). For example, movements across short distances and in multiple directions increase the amount of feedback for a small area of the display and the *check-neighbors* gesture describes structured movements within a small area. Similar to the gross movement patterns, these fine movement patterns could be characterized by their geometric shape. They were conducted to systematically explore without a-priori knowledge of the display contents and seemed to continue independent of the sequence of auditory events in the feedback.

Two fine movement patterns could be characterized by their relationship with the display contents (Fig-

```
Model: "Learn spatial layout"
      Starting_goal is Explore Map

Method for Goal: Explore Map
      Step: Store <Map> under <current_activity>

      Step: Store <gross_movement_strategy> from <user_preference>
            under <current_gross_movement_strategy>

      Step: Store <fine_movement_strategy> from <user_preference>
            under <current_fine_movement_strategy>

      Accomplish Goal: Set Input_Device
      Step Explore_loop:
            Accomplish Goal: Query Display
            If <state> of <mental_model> is satisfied, Then
                  Step: Delete <current_task>
                  Return_with_goal_accomplished
            Goto: Explore_loop
      Return_with_goal_accomplished
```

Figure 16: Exploration is one of the four tasks in the Region Lab and consists of a sequence of queries to the auditory display. Participants were asked to explore the map and familiarize themselves with the layout. They decided independently when they were satisfied with their level of familiarity with the display contents.

A GOMS-style model that uses the goals and items defined in Figures 7 through 15, describes the organization of the steps involved in the exploration and selection tasks from the Region Lab.

5.1.4 Exploration Task Strategies

The exploration task is characterized by gross movement patterns. Participants tended to employ broad strokes that covered a large proportion of the active area of the tablet (although the observed difference was not found to be significant, see Section 5.4). The model depicted in Table 16 conditions continuation on whether or not the user is satisfied with their exploration (see Section 5.1) without attempting to define how individual users determined satisfaction.

5.1.5 Selection Task Strategies

The selection task was designed to test participant's mental representation of the map contents. As shown by the number of attempts taken, however, response accuracy was low and and completion of the selection task involved a substantial exploration effort. On average, participants took four attempts to correctly select the target county. Two participants made correct selections on the first attempt for all repetitions of the task

```
Model: "Selection Task"
    Starting_goal is Select Target_county

Method for Goal: Select Target_county
    Step: Store <Single_Low_County> under <current_task_name>
    Step: Store <Value of current_task> under <current_target_value>
    Step: Store <target_pitch of Current_Task> under <current_target_pitch>
    Step: Store <empty> under <pattern>
    Step: Store <same_plus_(different_same)_x_4> under <target_pattern>
    Accomplish Goal: Set Input
    Step Search_loop:
        Accomplish Goal: Locate a County
        Decide:
            If <strategy> is Guess_and_check, Then
                Accomplish Goal: Perform selection action
            Else, If <strategy> is Check_neighbors, Then
                Accomplish Goal: Check Neighbors
                If <state> of <mental_model> is satisfied, Then
                    Step: Delete <current_task>
                    Return_with_goal_accomplished
                Else,  Goto: Search_loop
            // Other strategies (such as circumnavigation could be added here)
    Return_with_goal_accomplished

Selection Rules for Goal: Perform selection action
    // Set the mode to hardware button,
    Accomplish Goal: Set Input

    // and perform a button press
    Step: Press button
    Return_with_goal_accomplished
```

Figure 19: In response to the selection task, participants developed strategies to verify that the conditions of the task were met. Two such strategies were "guess and check" and a *check-neighbors* gesture

Direct menu navigation

Time (s)	Operator	Action	Description
1.35	M	Mental preparation	Plan the selection
0.40	H	Homing	Move hand(s) to the keyboard
0.20	K	Keypress 'Alt'	Send focus to the menus
0.20	K	Keypress 'u'	Move focus to the *Tutorial* menu
0.20	K	Keypress 'i'	Select the *Instruction* menu item
0.40	H	Homing	Move hand(s) to the tablet and stylus

2.75 Total time

Relative menu navigation

Time (s)	Operator	Action	Description
1.35	M	Mental preparation	Plan the selection
0.40	H	Homing	Move hand(s) to the keyboard
0.20	K	Keypress 'Alt'	Send focus to the menus
0.20	K	Keypress 'right arrow'	Move focus to the *Tools* menu
0.20	K	Keypress 'right arrow'	Move focus to the *Tutorial* menu
0.20	K	Keypress 'down arrow'	Move focus to the *Instruction* menu item
0.20	K	Keypress 'enter'	Select the *Instruction* menu item
0.40	H	Homing	Move hand(s) to the tablet and stylus

3.15 Total time

Figure 20: Two keystroke level models (KLM) of the actions required to select the *Instruction* menu item, which initiates the transition from the exploration task to the selection task provide slightly different time estimates. Participants could either use the mnemonic keyboard shortcuts (direct navigation; top) or use the arrow keys (relative navigation; bottom) to navigate the menu.

	Interval durations that exceed predicted value (seconds)								Mean duration of known transition intervals (seconds)
Participant	Step 3			Step 6					
109	**24.5**			**29.7**					13.3
125	**6.8**			**51.5**					23.0
200	**26.7**			**26.9**	3.3				23.0
276	**14.1**			**11.8**					8.7
518	**63.2**	7.9	5.5	**81.5**	3.6				22.7
628	**17.1**			**23.2**					13.0
712	**25.9**			**41.5**					28.1
757	**19.4**			**31.0**	3.7	20.3	4.8		20.7
975	**47.5**			4.4	**25.0**	20.8	3.8	4.3	23.9

Table 2: Candidate intervals to represent the execution of a menu selection are those intervals whose duration exceed the theoretical values predicted with a KLM model. The duration of intervals that both exceeded the minimum threshold predicted by the model and passed the two exclusion criteria are highlighted in bold.

shared pixel boundaries (movement orthogonal to the pixel boundary; 50.84%) and across a shared point at the corner of the pixel extent (movement at a 45° angle to pixel boundaries; 17.7%). The disproportionate representation of directions that correspond to movement across pixel boundaries (Figure 22) is a result of the discretization of samples in a digital display and an imbalance between resolution in time and in space. While movement in the four or eight evenly spaced directions (i.e., cardinal and sub-cardinal directions that could be mapped to the arrow keys or numbered keypad) has been emphasized in the literature (e.g., [34, 8]), raw data collected during behavioral testing sessions shows an artificially strong trend manifesting itself as a stair-step pattern for paths along diagonals (Figure 23, center column). Resampling the raw data (24 Hz) gives a more representative measure of movement direction (Figure 23, right column).

Even after resampling the data, classifying all movements as one of four directions (rightward, leftward, upward, and downward) revealed a higher proportion of lateral movements (right and left across the body) than radial movements (toward and way from the body; Figure 24).

5.3 Gesture Identification

anecdotal observations during behavioral testing sessions suggested that participants used a *check-neighbors* gesture to determine whether or not a county has neighbors that share the same attribute value. A state-based model extracts the gesture from the streams of recorded stylus positions and identifies 13 occurrences of the gesture in logs from seven unique participants, each contributing one to three occurrences. Log files for the exploration task contain three occurrences of the gesture and those for the selection task contain ten occurrences. Each occurrence of the gesture identified by the automated process indicates a sequence of steps (transitions from one county to another) that satisfies the model. The current method does not compute a measure of confidence in the identification. Interpretation of the identified candidate gestures is discussed further in Section 6.3.

The shape of the gesture varied from crossing linear strokes (Figure 25) to strongly circular (Figure 26). Some identified instances of the gesture could be coincidental rather than intentional (i.e. false positive). Without addition information (e.g., explicit input from the participant), it is not possible to discriminate between a false positive that matches the state-based model and an intentional action to check topology.

5.4 Quantitative Metrics

Metrics defined in the literature (visit rate, number of steps, and number of *keep-a-direction* steps [8]) and the number of peaks in the output of the KDE proposed in the Methodology (Section 4.5.1) measure possible

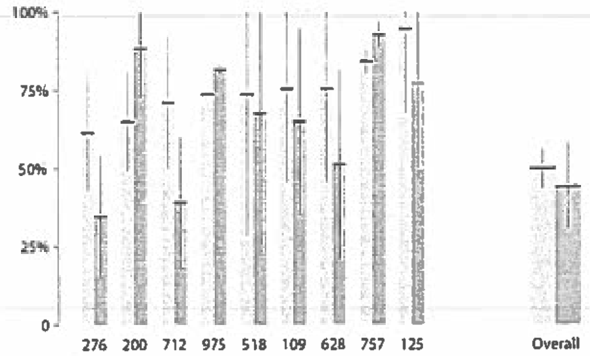

Figure 27: The observed percentages of counties visited (visit rate) were not significantly different between the two tasks. Arranged here from lowest to highest visit rate for the exploration task, there is no obvious trend in the respective visit rates for the selection task.

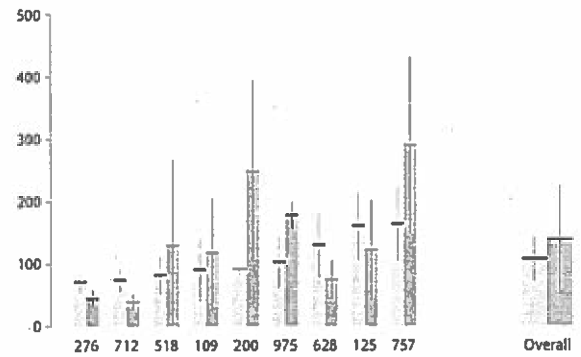

Figure 28: The observed number of times that the stylus transitioned from one county to another county (transition frequency) was not significantly different between the two tasks. Select is Red, Explore is blue

differences between performance in the exploration task and that of the selection task (addressing the second research question).

As part of the Region Lab, participants visited 74.49% of counties during the exploration task and 67.07% of counties during the selection task, on average. The observed difference in the percent of counties visited between the two tasks was not significant (paired T-test: $t = 1.373$, $df = 8$, $p = 0.207$, *n.s.*; Figure 27).

The average number of steps is 85 and 117 for the exploration and selection tasks, respectively. The difference in the counts is not statistically significant (paired t-test: $t = 1.2477$, $df = 8$, $p = 0.247$, *n.s.*; Figure 28).

The average proportion of keep-a-direction steps is 9.00% for the exploration task and 5.11% for the selection task. The difference in the mean proportion of keep-a-direction steps between the two tasks is substantial (paired t-test: $t = 3.5815$, $df = 8$, $p = 0.007$; Figure 29).[31]

[31] Due to the exploratory nature of this analysis, the low probability is not sufficient to conclude a significant difference.

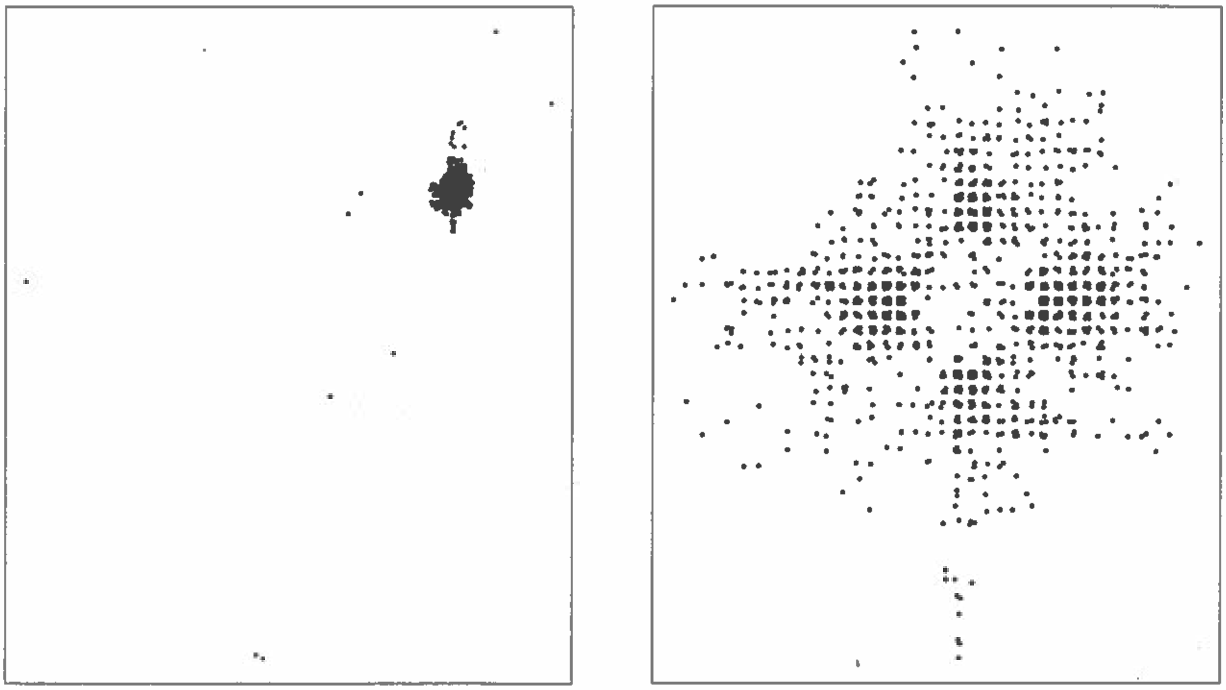

Figure 30: The 99th percentile delineates the speed based outliers from the core cluster of observed speeds. Points whose distance from the origin (speed) is above the 99th percentile are represented in red and are excluded from further analysis.

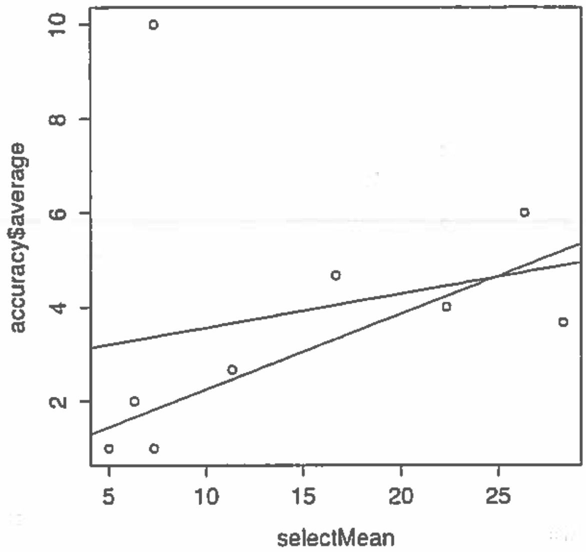

Figure 32: Peak count was observed to be a strong predictor of selection accuracy when an instance of "guess-and-check" selection was excluded.

one participant who used a "guess-and-check" method reveals a much stronger relationship (F-14.51, df=6, p=0.0089; Figure 32).

While the results include statistical probabilities, as part of the exploratory analysis they are not conclusive evidence. Further study controlling for additional factors and targeting additional details related to the observations made in this analysis will be required. An interpretation of the results - both observations and statistical outcomes - are discussed in the next section.

6 Discussion

To complete the Region Lab, participants translate verbal (text-to-speech) instructions into sequences of actions that are executed using the stylus. Stylus movements are observable artifacts of participants' stylus use. The translation is influenced by the design of the software, the display contents, and the participants' strategies. The discussion in this section includes brief reflection on the design of the Region Lab and interprets the results presented in the previous section.

6.1.2 Tasks in the Region Lab

Understanding the patterns in the distribution of population density using an auditory interface is non-trivial and coupled with the the complex geospatial thinking concept of region, the Region Lab presented a challenge to several participants. Wording in the explanations and instructions avoided the term region prior to the introduction of the *Region Tool* (a custom spatial analysis tool). Some participants did not realized that the *Region Tool* is a composite tool that aggregates the functionality of each of the other tools (i.e., classification, proximity, cluster, and boundary) into a single step. Further it was unclear whether or not participants understood that applying the Region Tool (or each of the component tools in sequence) resulted in a representation of the newly identified regions. Although the intention was to teach participants the target concept as they completed the Region Lab, participants could provide correct responses without possessing a correct understanding of the concept. All of the participants successfully completed all of the tasks, however, their comments revealed that not all of them understood the objective of the Region Lab. Some participants did not make a connection between the set of individual concepts (embodied in the individual spatial analysis tools) and the more complex concept of region.

The instructions for the exploration task were intentionally vague and were intended to encourage participant to form an internal (mental) representation of the distribution of counties during the open ended exploration task. During the behavioral testing, however, some participants spent less time exploring than had been anticipated. They engaged in limited interaction with the display before proceeding to the selection task. In the data collected, there was no direct evidence to evaluate whether or not participants were creating, updating, or verifying an internal representation of the distribution.[32]

In contrast to exploration, the selection task does provide an explicit, concrete objective: select a point in the display that satisfies a given condition (e.g. "no similar neighbors"). If a participant has formed a sufficient internal representation during exploration, this task may reduce to a query of that internal representation, a pointing action, and a button press. During the exploration, however, the participants had been given limited information about what they would be asked to do with the information that they gathered from the map or what kind of internal representation of the data would be sufficient. The distinction between the two tasks was not clear. Actions during the selection task involved more extensive search behaviors than

[32]Stylus movement collected during the verification task provide evidence of using an internal representation of the spatial layout. Participants return with high fidelity to the general location on the display that they had previously selected, even without auditory feedback (i.e., after the intervening spatial analysis tool has removed the selected county from the display and that area of the display is represented with silence). During that step, however, some participants realized that they could "appease the wizard" by just tapping the stylus to the display. Quantifying the level of fidelity in the verification task or it's effective role in the educational intervention was beyond the scope of the investigation reported in this document.

6.2 Models

Both a keystroke level model (KLM) and a more general goals, operators, methods, and selection rules (GOMS) model help formalize out understanding of the actions that recorded stylus positions represent and reason about the way that participants may have translated the instructions into those actions. The keystroke level model KLM gives a conservative, lower bound estimate of the time required to execute the selection. Participants did not always perform the optimal sequence of actions (for their chosen approach: absolute keystroke or arrow keys for relative navigation) when making a menu selection. For example, several participants used the key combination 'Alt-f' to enter the menus - specifically to go to the *File* menu - even though they could have used the letter key for the *Tutorial* menu ('u') to go directly to the target menu. The two versions of the KLM model represent two ways to navigate the menus and submenus, but both assume that user actions are optimal. To the contrary, participants exhibited suboptimal behavior that satisfied the task but took more time to complete. The predicted times were much smaller than the observed time intervals in which the stylus was not touching the tablet (approximately a factor of ten). For the data in this analysis, either model is sufficient to identify unique candidate intervals during which the menu selection was likely to have taken place. Additional delays would need to be incorporated into the model to account for the alternative, suboptimal strategies.

The GOMS-style model provides a way to make explicit beliefs and assumptions about how users completed the two target tasks. An objective in their creation is to produce simple models that account for theoretically predicted and empirically observed actions. They cannot reveal motivation and have not been extensively validated, but provide organization and vocabulary to classify and discuss observations from behavioral testing. Complexity or contradiction in the models helps identify gaps in our understanding. Looking at abstract representations of the tasks also helps mitigate the researchers' visual bias when determining interface requirements and evaluating user behavior.

Initially, we expected participants to spend a substantial amount of time exploring the map to build a mental representation. When they were then given a selection task, the response time could be captured in a KLM as the sum of the time spent querying the mental representation for the target, the time to execute a pointing task (i.e. position the stylus over the target), and the time required to press the hardware button on the tablet. In practice, however, we observed participants explore briefly and then proceed directly to the selection task where they were given the constraints of the selection. In the behavioral testing, some participants anticipated the target for a subsequent selection task (as some participants who seem to do) and could attend to and efficiently represent in memory aspects of the display that facilitate decision making

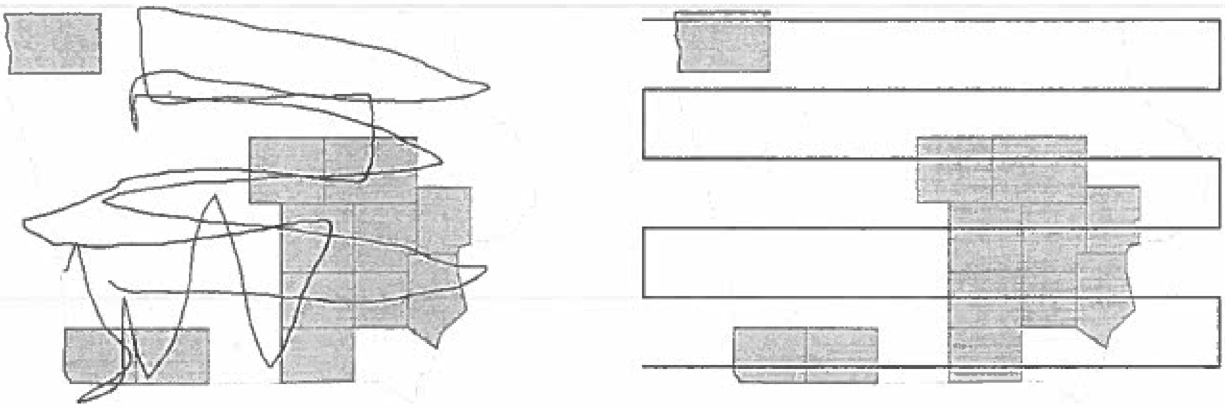

Figure 34: In a comparison of a strongly geometric (right) and empirically observed (left) instances of the serpentine gross movement pattern both show broad coverage of the active area of the tablet. Unlike the theoretical model, the empirically observed instance includes movements along the diagonal and the stylus did not move all the way to the edge of the active area before changing direction.

contribute interface usability. Although it may be difficult to adapt the GOMS model to reflect components of cognitive processing, identifying specific software design choices or assumptions about user behavior that challenge or contradict the framework may also clarify our understanding of observed behaviors, help improve usability, and lead to new questions.

6.3 Gesture Identification

Traditional approaches to gesture recognition (see Section 2.2) rely on an explicit start of the gesture. When extracting instances of the gesture from a stream of stylus locations, the start of the gesture is inferred from the changes in the location of the cursor relative to display elements. Although this approach cannot determine intentionality and is susceptible to false positive identifications (deciding that a gesture was performed when it was not an intentional act by the user), it provides an objective metric with which to investigate the prominence of the gesture.

The *check-neighbors* gesture was observed more often in the selection task (ten times) than it was in the exploration task (three times). However, as discussed above there was not a clear distinction between the behaviors during the exploration task and those in the selection task: participants tended to wait for instructions about the specific selection task before conducting a thorough exploration (Section 6.1.2). Participants who were anticipating the selection task may have performed a gesture to check the attribute values of neighboring counties during the exploration task or the identified gestures may simply be a coincidental match. The fact that multiple participants independently developed a cross-shaped gesture that could pro-

Steps that continued in the same direction as their direct predecessor could indicate straight line movement along a pre-planned trajectory. We observed a relatively higher proportion of "keep-a-direction" steps for the exploration task than for the selection task (9.00% and 5.11%, respectively). Again the difference was not found to be significant. The data does, however, show a trend toward greater differences between the tasks for participants who performed higher proportions of "keep-a-direction" steps during the exploration task.

A non-parametric kernel provides a way to characterize the observed movements at a higher resolution than has been reported in previous studies of stylus movements related to displays of spatial data. Number of peaks is a measure of regularity of movement direction that is independent of the display contents (the users choices in continuing, modifying, or ceasing movement may still be influenced by the feedback, but the analysis does not incorporate information from the relevant area of the display). The average number of peaks varied across participants, but was consistent across the two tasks. While the small number of participants limits the feasibility of a detailed statistical analysis, the difference between the number of peaks tended to increase over time.

The number of peaks in the selection task accounts for a substantial portion of the variability in accuracy responding to the selection task (PRE=0.4485). There is no evidence that this relationship is causal, but may suggest that evidence of systematic exploration strategies could serve as a proxy to identify users who will perform well in the domain of geospatial thinking.

The quantitative analysis did not reveal any statistically significant differences between the two tasks or over time. Differences in performance metrics may have been masked by individual differences, but there was certainly insufficient power for the statistical tests. With only nine participants, the probability of finding a significant results is low ($\approx 41\%$; i.e., a high probability of a Type II error) even if there were a large true difference between the two conditions ($\eta = 0.3$) [19]. In this data set, the number of participants was too small to reach conclusive results for a model of interaction effects; the computed probabilities illustrate the design of the analysis, but additional data collection is required to conclude meaningful significance.

7 Conclusions

This paper presents an analysis to address two research questions. The first research question is methodological: How can observed trends in stylus movement be characterized? A GOMS-style model of user strategies, a state-based model of a gesture for checking topology, and a point pattern analysis of points representing serve as a proxy for users' attention during exploration and selection and movement is a direct representation of the users actions as they completed the given tasks. No significant differences were observed between the two tasks based on the quantitative metrics. Results reveal trends in the direction of movement, with users preferring lateral (right and left) and radial (toward and away from the user's body) movements over movements along diagonals. With a high probability of a Type II error Results from the available data suggest changes in the number of directions observed in participants' free-form movements over time. Additional data is needed to determine if those changes over time are significant.

Through this work, we are striving to create an accessible interface to a tool for processing and understanding geospatial data. Toward this end, we created the minimal Geographic Information System (mGIS) and observed user performance in the Region Lab. The results from this analysis (and the process of its execution) have informed our understanding of how users who are blind interact with an auditory display of spatial data. They serve as a starting point for further design and experimental evaluation of the mGIS interface. The prototype display implemented in the mGIS is based on our experience with printed maps and graphical computer displays, but involving end users as we iterate through the design process will produce a more accessible and usable interface. Our findings will guide future inquiry into auditory interface design. Potential directions for future work include diversifying the spatial thinking tasks that provide context for the analysis and refining the code that supports audio rendering in the software prototype.

Enhance Training Materials The independent development of the *topology check* gesture by several participants suggests that it may be a useful strategy to teach to novice users of the interface or to provide as a tool in future versions of the mGIS display.

Extend Analysis to Additional User Tasks

The analysis described in this paper focuses on only two of the tasks in the Region Lab: exploration that focuses on the layout of the display contents and selection of a county that had no similar neighbors. The same data set used for this analysis also contains data for selection tasks that target different spatial thinking concepts and for a verification task highlighting the effect of spatial analysis tools (e.g., explore the map to "notice that the county by itself in the bottom of the map has been removed"). Considering these additional tasks could help clarify how the task influences user behaviors.

8 Acknowledgements

I would like to thank my advisor Michal Young and committee members Sarah Douglas and Amy Lobben for the guidance, insight, and advice that I have received over the course of this project.

This research is supported in part by the National Institutes of Health (NIH NEI 1RC1EY020316). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NIH.

The Bamboo tablet used in the experiment was generously donated by Wacom.

References

[1] Alty, J. L., and Rigas, D. Exploring the use of structured musical stimuli to communicate simple diagrams: the role of context. *International Journal of Human-Computer Studies 62*, 1 (2005), 21–40.

[2] Appert, C., and Zhai, S. Using strokes as command shortcuts: cognitive benefits and toolkit support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2009), CHI '09, ACM, pp. 2289–2298.

[3] Baddeley, A. Working memory. *Current Biology 20*, 4 (2010), R136 – R140.

[4] Bau, O., Poupyrev, I., Israr, A., and Harrison, C. TeslaTouch: electrovibration for touch surfaces. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology* (New York, NY, USA, 2010), UIST '10, ACM, pp. 283–292.

[5] Card, S. K., Moran, T. P., and Newell, A. *The Psychology of Human-Computer Interaction.* Lawrence Earlbaum Associates, Hillsdale, NJ, USA, 1983.

[6] De Felice, F., Renna, F., Attolico, G., and Distante, A. A haptic/acoustic application to allow blind the access to spatial information. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)* (Tsukuba, Japan, March 2007), IEEE.

[7] Delogu, F., Belardinelli, M. O., Palmiero, M., Pasqualotto, E., Zhao, H., Plaisant, C., and Federici, S. Interactive sonification for blind people exploration of geo-referenced data: comparison between a keyboard-exploration and a haptic-exploration interfaces. *Cognitive Processing 7*, 0 (2006), 178–179.

[19] JUDD, C. M., MCCLELLAND, G. H., AND RYAN, C. S. *Data Analysis: A Model Comparison Approach*, 2nd ed. Routledge, New York, NY, 2009.

[20] KANE, S. K., WOBBROCK, J. O., AND LADNER, R. E. Usable gestures for blind people: understanding preference and performance. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (New York, NY, USA, 2011), CHI '11, ACM, pp. 413–422.

[21] KEELE, S. W. Movement control in skilled motor performance. *Psychological Bulletin 70*, 6 (December 1968), 387–403.

[22] KIERAS, D. GOMS models for task analysis. In *The Handbook of task analysis for human-computer interaction*, D. Diaper and N. Stanton, Eds. Lawrence Earlbaum, Mahwah, New Jersey, 2004, ch. 4.

[23] LÜ, H., AND LI, Y. Gesture Coder: a tool for programming multi-touch gestures by demonstration. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI '12, ACM, pp. 2875–2884.

[24] MAHAN, R. P. Circular statistical methods: Applications in spatial and temporal performance analysis. Tech. rep., University of Georgia, 1991.

[25] NORMAN, D. A. *Psychology of Everyday Things*. Basic Books, Inc., New York, NY, 1988.

[26] PATTERSON, R. D. *Guidelines for Auditory Warning Systems on Civil Aircraft.* (CAA Paper No. 82017). Civil Aviation Authority, Cheltenham, England, 1982.

[27] RICE, M., JACOBSON, R. D., GOLLEDGE, R. G., AND JONES, D. Design considerations for haptic and auditory map interfaces. *Cartography and Geographic Information Science 32*, 4 (2005), 381–391.

[28] RICKETTS, T. C. Geographic information systems and public health. *Annual review of public health 24*, 1 (2003), 1–6.

[29] RUBINE, D. Specifying gestures by example. *SIGGRAPH Comput. Graph. 25*, 4 (July 1991), 329–337.

[30] VÖLKEL, T., WEBER, G., AND BAUMANN, U. Tactile graphics revised: The novel BrailleDis 9000 pin-matrix device with multitouch input. In *Computers Helping People with Special Needs*, K. Miesenberger, J. Klaus, W. Zagler, and A. Karshmer, Eds., vol. 5105 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2008, pp. 835–842.