

Hierarchical multi-touch selection techniques for collaborative geospatial analysis

Thomas Butkiewicz, Dong Hyun Jeong, Zachary Wartell, William Ribarsky, and Remco Chang
University of North Carolina at Charlotte

ABSTRACT

In time critical visual analytic environments collaboration between multiple expert users allows for rapid knowledge discovery and facilitates the sharing of insight. New collaborative display technologies, such as multi-touch tables, have shown great promise as the medium for such collaborations to take place. However, under such new technologies, traditional selection techniques, having been developed for mouse and keyboard interfaces, become inconvenient, inefficient, and in some cases, obsolete. We present selection techniques for multi-touch environments that allow for the natural and efficient selection of complex regions-of-interest within a hierarchical geospatial environment, as well as methods for refining and organizing these selections. The intuitive nature of the touch-based interaction permits new users to quickly grasp complex controls, while the consideration for collaboration coordinates the actions of multiple users simultaneously within the same environment. As an example, we apply our simple gestures and actions mimicking real-world tactile behaviors to increase the usefulness and efficacy of an existing urban growth simulation in a traditional GIS-like environment. However, our techniques are general enough to be applied across a wide range of geospatial analytical applications for both domestic security and military use.

Keywords: Multi-touch, hierarchical selection, geospatial analysis, GIS, collaborative analysis

1. INTRODUCTION

As visual analysis applications migrate to the emerging next generation displays, the radical changes in the amount of information and the ways in which it is presented will require that existing interaction techniques be likewise upgraded or replaced with new techniques. These techniques should not merely accommodate the changes, but seek to embrace the new paradigms with full support for non-traditional interaction and promote collaboration. The multi-touch table is currently a very popular technology that has been receiving much media attention, much of which originated from Han's presentation^[4] of an inexpensive and easy to build Frustrated Total Internal Reflection and computer vision based system. These multi-touch tables allow multiple users to interact simultaneously in the most natural and intuitive way, with their bare hands. Standard interaction techniques, developed for mouse-and-keyboard interfaces, translate poorly to this new environment.

In Section 2 we use the multi-touch table environment as an example to discuss ways in which traditional interaction techniques can be improved and extended for visual analytic systems dealing with geospatially tied data. We look at ways in which traditional selection techniques can be improved to take advantage of the areas in which these new display technologies excel. We also consider the new requirements imposed upon interaction techniques and interface design in order to accommodate collaborative environments with multiple simultaneous users.

We demonstrate the applicability and effectiveness of our techniques by taking an existing simulation and a traditional mouse and keyboard interaction technique and redesigning them for use in a multi-touch table environment. Here we implement some of the concepts explored in the theory section to not only accommodate the new input and output characteristics of the multi-touch table, but to take full advantage of the new technology's ability for intuitive interaction and collaborative deployment. Our improvements allow for untrained users to quickly understand and master the interface and spontaneously experiment with different scenarios.

In addition to implementing our techniques within a multi-touch environment, we provide a discussion reflecting on the benefits we gained by using our techniques over the traditional mouse-and-keyboard approach within an actual

geospatial analysis application under real world use with target audiences. We propose that while there is inherent computation overhead in a multi-touch environment, the benefit of having multiple users simultaneously interacting with the same data space for the purpose of collaboration and analysis sharing outweighs the cost. Within our geospatial analysis application, using our proposed interaction technique, a user can quickly identify and select regions-of-interest and pass the analytical results to others for further exploration and investigation. We conclude that the combination of our interaction techniques and new display technologies such as a multi-touch table allows for more intimate collaborations between analysts and decision makers by leveraging the knowledge of multiple users and allowing them to work together more efficiently.

2. THEORY

Certainly the use of a multi-touch table environment has much potential for use with visual analytics. The most exciting prospect is the ability for multiple analysts to simultaneously explore the same data/information space. This can lead to far more comprehensive collaboration, as each is continuously updated on progress with the group's task and can immediately examine not only others' findings but the origins and context of those findings with regards to their own tasks. Shared hypothesis analysis can flourish as a true collaborative process, with the catalysts of natural communication and direct interaction.

For time-critical analysis in geospatial environments one of the most important questions to be answered is Where? Where has something new occurred? Where are the after-effects likely to emerge? Where are the resources we currently need for this situation? These questions can almost always be answered with regions-of-interest. As such, selecting regions-of-interest for analysis is an essential operation that must be streamlined to be efficient but accurate. In a collaborative environment, one must also focus on refinement of these regions, as well as the inherent behavioral characteristics of table-top collaboration that must be accommodated in the interface's design.

2.1 Selection

A number of researchers have devised ways of making multi-touch selection more precise, thus overcoming the inherent lack of positional precision common in many vision based systems with lower camera/display resolution ratios. Benko^[1] presents a number of Dual Finger Selection techniques which allow, for example, one finger to control the magnification of a target for the other finger to touch, or the speed of a cursor following the other finger. The selection techniques this paper, however, are concerned not with providing means of enhancing precision of manual selection, but on facilitating fast and efficient selection of regions in a hierarchical geospatial environment. We do provide a free form selection method that might technically benefit from these precision enhancing techniques, but it is *intended* to be an imprecise method of selection, good for circling roughly "around" objects of interest (a very broad and general action, but one that is also very commonly desired during exploratory analysis.)

When a more precise and accurate selection is required, we rely on the abundance of geospatial records delineating everything from individual tax parcels, school districts, to entire counties. These pre-defined regions, imported from existing GIS databases, are easily formed into a hierarchy with a range of granularities that make excellent building blocks for selections to be accurately created and refined. During selection and refinement, we chose not to implement complex gestures that require memorization and training, but instead to treat the users fingers as small tools that, when used together, form a more powerful and flexible instrument, much like individual tines forming a rake. A single finger can perform a simple action, in this case bringing pre-defined areas (e.g. tax parcels) into a region-of-interest. When this behavior is multiplied across an entire hand, the user can rake or sweep whole neighborhoods of tax parcels in and out of, and back and forth between multiple regions-of-interest in a natural and intuitive manner. While the individual finger behaviors themselves could each be replicated individually with a mouse doing a "drag-and-drop" action, the cumulative effect of being able to complete hundreds of these simple actions in just a few seconds by sweeping your hands around the map is like nothing one can do with a mouse. In effect, it provides a sandbox like feel to the interface, with the added bonus of being able to dynamically vary the granularity of the medium to iteratively refine your selections.

2.2 Collaboration

The main challenge that must first be overcome when attempting to handle multiple simultaneous users in a collaborative application on a multi-touch table such as ours is determining the ownership of touch events. That is, when a finger touches the surface, the system needs to be able to distinguish which user has initiated the event.

Early systems, such as DiamondTouch^[3] achieve this by having finger touches complete capacitive electrical circuits between the table's surface, through the user, and into the users chair. This places restrictions on the surface's material construction and translucency, thus limiting the projection methods that can be used for the display. (Generally only projection from above is suitable here). By making the chairs an integral component of the tracking system, it limits the number, arrangement, movement, and positions of users. We argue that placing these limitations on users is unacceptable for most use real-life usage where people will unpredictably change stances and move about. Furthermore, the above projection systems suffer from the serious issue of occlusion, which worsens with each additional user.

The more promising option is to use computer vision methods to track users interacting with the table. By placing a camera above the table, one can segment the arms and hands of individual users and monitor their locations. When a touch event occurs, its location can be compared to the known positions of the users' hands, and thus the user who initiated the touch event can be determined. While vision based solutions generally have significant computational cost, however the tracking of arm and hand locations can be done at relatively low resolution due to the high ratio of size between the objects being tracked and the tracking area, and there is a much lower requirement for frame rates since one is not concerned with tracking each fine movement, only general locations. Some have found successful methods which allow collocating the user tracking camera(s) behind the screen with the touch tracking camera(s) through the use of DNP HoloScreen, as in the TouchLight^[7] system, which allows projection onto transparent surfaces, or with a active surface material, such as used by Kunz^[5] for CAVE walls, that can alternate its opacity fast enough to allow both projection of the display image and visual tracking of the users on the other side.

Another set of issues that arises when designing collaborative multi-user systems such as ours is workspace territoriality. Decisions to be made here include what distinctions, if any are to be made between "personal space" and "shared/group space." Scott^[6] observed tabletop collaboration in casual and formal settings and found participants would intuitively divide their workspaces into three general types of interaction areas: personal, group, and storage. These areas would expand and contract dynamically as the activities continued, but there were several constant themes that must be considered when designing a collaborative workspace. Areas directly in front of participants were always used as personal workspaces in which participants would individually execute their own tasks. Areas within in the center of the tabletop were used as a "shared" or "group" workspace. It was here that participants collaborated on the main tasks that contributed toward the overall group's goals. It was also here that participants would help each other with their individual tasks. The peripheral edges of the table tended to be used as a sort of "storage area" in which items were loosely organized and moved further away from the active workspaces as they became less often used. We further discuss Scott's conclusions about design implications in respect to our particular application in Section 3.4.

3. APPLICATION

Our specific implementation comes in the form of an interactive viewer for the prediction results from an urban growth model/simulation developed by The Center for Applied Geographic Information Science at UNC Charlotte. The underlying data to be presented consisted of a number of large (8545x6200) raster maps covering a 24-county region surrounding Charlotte, North Carolina. Maps showing land use, specified as either natural/undeveloped, developed/impervious surfaces, water, or protected space (e.g. parks) were given for a number of historical time steps starting from 1976, the present, and for predictions of future land use every 5 years until 2030.

The existing method of presenting the results of the urban growth model was to form them into a simple slide show. We sought to develop a new interface that could be used to both interactively explore and present these results to politicians, community groups, policy makers, and the general public. For small audiences we wanted to not only be able to present the results, but also engage them with the ability to play with the data themselves, leading to the discovery of findings of personal interest. As such, we recognized the promise of a multi-touch table to help us accomplish these objectives.

Although this application was developed specifically for deployment on our touch-table, we attempted to maintain a certain level of backwards compatibility with the traditional mouse/keyboard interface. In this pursuit we developed a number of touch-emulation key-mouse combos that would simulate the various gestures needed to navigate the map, complete selections, etc. We ensured that all selection and refinement techniques presented here can be accomplished, albeit slowly and only one user at a time, through the use of a traditional mouse alone. This backwards compatibility allows the system to be demonstrated on a laptop when and where the table itself cannot be physically present, as well as for remote testing and development.

3.1 Probe-based interaction

It was important to both allow multiple users to collaboratively operate the system simultaneously while also giving personal “work-spaces” to each user. As such, we chose to implement a probe-based method of interacting with the underlying data. Probes^[2], which consist of user-defined regions-of-interest tied to individual, floating coordinated visualizations, allow any number of users to each focus in on areas of their own choosing and receive customized visualizations and interfaces giving them both viewports and control panels into the underlying data/models. The portion of the probe-interaction we will concentrate on in this paper is the set of tools provided to allow users to select regions-of-interest and the those provided to allow refinement of already selected regions-of-interest.

3.2 Selection

We provide users with a number of different methods of selecting regions-of-interest. Since our target users are primarily people who are unfamiliar with the system and will not have time to be taught how to use it, we decided to allow users to choose a method from an onscreen “toolbox,” as opposed to implementing a system of gestures (which would need to be learned) that would initiate the different selection modes. We placed multiple instances of the toolbox along the lower-right (from each user’s perspective) corners of the screen. This allows quick access (as most users rest their dominate hand nearby), eliminates the need to reach across the shared workspace, and as our tables touch-resolution is lower around the edges (due to extreme-wide angle camera lens), this space is better suited for blunt button presses than precise manipulation of data.

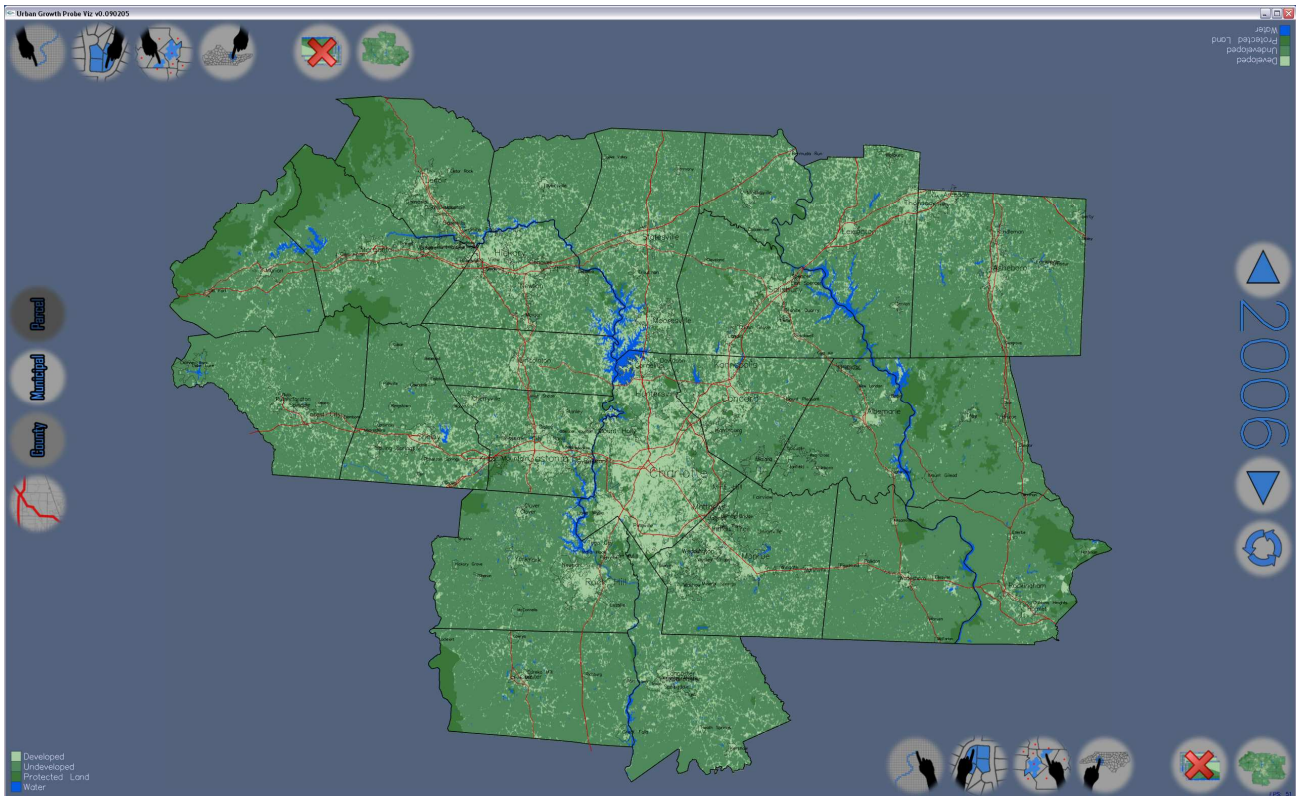


Fig. 1. The default view of the system’s interface. Notice that the commonly used selection toolbox buttons are conveniently presented both in the lower-right for the Southern users, and the upper-left for the Northern users. Map legends are also replicated for each side in the opposite corners. Less commonly used buttons to control map drawing parameters are presented midway on the rarely (due to table shape) occupied East and West sides of the table, putting them easily within reach of both Southern and Northern users.

The first selection tool is a free-form, or lasso selection. This tool is used commonly to select irregular regions, such as proposed development areas, along features (roads, waterways, etc), and any other areas not already delineated by pre-loaded geospatial records. When the free-form selection mode is entered, the user can drag their finger anywhere on the map, acting like a pen (with visual feedback in the form of a bold red line). The points that are recorded in between the finger-down and finger-up events are transformed into a list of lines, and we check for intersections that create valid polygons. We also check to see if the addition of a short connecting line could “close up” a near-polygon. This ensures that the user does not need to fully complete/connect a circling gesture around a region in order for it to register. The polygonal regions generated are then used to generate the desired region-of-interest.

The next set of selection tools take advantage of the pre-loaded hierarchical geospatial data the system has associated with the survey extent. These tools allow the user to select these pre-defined sub-regions at multiple granularities. The basic levels we provide in our system are: parcel, municipal, and county. The municipal selection tool, for example, allows the user to quickly and easily select the area within a municipal (city, town, etc) boundary, including any unconnected outparcels. When entering the selection state with one of these tools, the geospatial boundaries associated with the desired granularity are highlighted on the map. You can then either make a single “click” with your finger, which will select the encompassing region at the desired granularity, or you can drag your fingers across multiple regions, which will select all regions (at the selected granularity) that encompass any portion of your fingers’ trail. We accomplish this by taking the list of touch-points and checking to see if they fall within the bounding circles of valid boundaries, then checking any matching points against the actual boundaries with a simple point-in-polygon test. All of the polygons from any matching regions (not just those from the matching boundary, which is a subset) are then compiled to form a single region-of-interest.

A second, slightly modified version of this type of selection mode is possible for these particular tools, in which the completion of a touch-event (finger-down to finger-up) does not automatically exit the selection state, instead each touch-event spawns a separate region-of-interest until the selection state is manually exited. In this manner, one is able to, for example, quickly tap a number of counties and have regions-of-interest for each spawn immediately without the iterative process of entering the selection state each time.

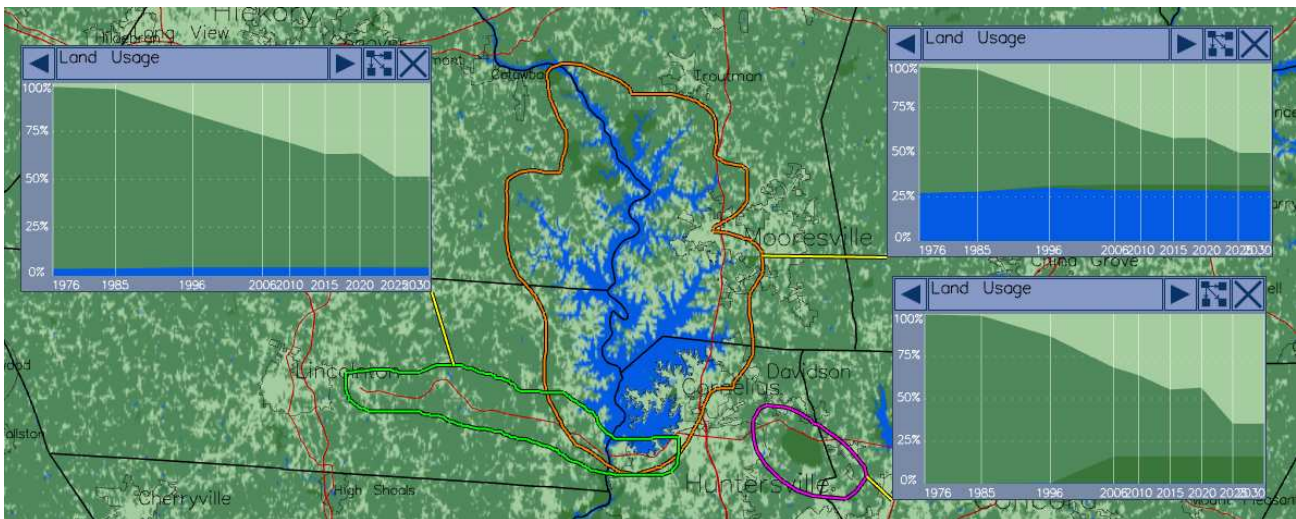


Fig. 2. Example regions-of-interest generated through the use of the free-form, or lasso selection tool. This method of selection allows for the selection of irregular regions-of-interest and those not easily defined by pre-loaded geospatial records. In orange, the user has traced a region roughly around a lake’s shoreline. In magenta, the user has quickly circled a protected area (a park). In green, the user has selected a region along either side of an interstate highway.

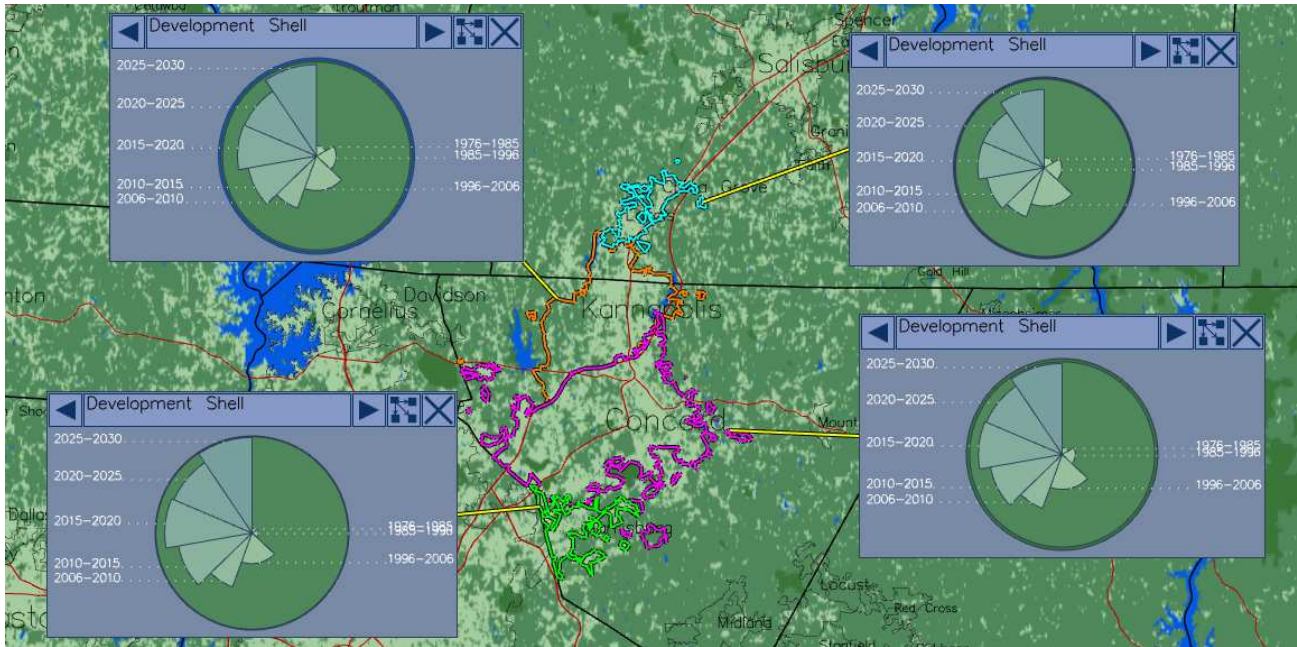


Fig. 3. Example regions-of-interest generated through the use of the municipal boundary selection tool. This method allows for quick and easy selection of complicated political boundaries including unconnected outparcels. Here the user has separately selected four adjacent cities, each of which's boundaries are translated into their own regions-of-interest. Had the user instead selected all four cities in a single finger-stroke, a single region-of-interest encompassing all four's boundaries would have been generated.

3.3 Refinement of Selections

The refinement mode provides the user with a highly interactive method of iteratively adding and removing spatial regions at multiple levels of granularity to and from existing regions-of-interest. The gestures for adding, removing, and transferring sub-regions are highly intuitive “pull” and “push” motions. The user can “pull” unselected sub-regions into existing regions-of-interest with the drag of a finger, and likewise can “push” selected sub-regions outside of or into other regions-of-interest. These gestures can be singularly replicated with a mouse interface, but their real power comes from the multi-touch abilities of the system: The user can utilize all their fingers to make refinements in a natural, instinctive manner. Simultaneously, a sweeping whole hand motion can pull broad swaths of territory into a region-of-interest, while a lone index finger accurately flicks individual parcels out to refine the border region.

When the user initiates the refinement mode, touches are tracked and processed as follows until the refinement mode is exited: Upon each finger-down event, a new touch object is created. We first check to see if the down-point is inside or outside of any current regions-of-interest that can be refined.

If a down-point is inside of an existing region-of-interest, we mark this touch as a removal in progress and wait for updates. As finger-movement events associated with this touch-event are received, we store them to examine, if needed later, to determine which sub-regions the touch-event crossed in between the finger-down and finger-up events. Finally, the location of the finger-up event then dictates the action to be taken. If the finger-up point is within the same region-of-interest as the finger-down point, the event is simply discarded. If the finger-up point is outside the region-of-interest, a list of sub-regions is constructed based on the initial finger-down location and the intermediate finger-movement locations along the way to the finger-up location. If the finger-up event is located inside of a separate region-of-interest, the sub-regions in the list are removed from the original region-of-interest and added to the second region-of-interest. If the finger-up event is located outside any existing region-of-interest, the sub-regions in the list are simply removed from the original region-of-interest and discarded.

If a down-point is outside of any existing region-of-interest, we mark this touch as a addition in progress and wait for updates. Once again we record the subsequent locations of finger-movement events as they are received, so that we can examine them later to determine which sub-regions outside of the region-of-interest the touch has crossed in between the finger-down and finger-up events. When the terminal finger-up event is received, we check to see if it is within an

existing region-of-interest. If it is not, we simply discard the entire touch event. If it is, we again build a list of sub-regions based on the initial finger-down location and all intermediate points up to the finger-up location. Any of these sub-regions that are not already in the region-of-interest (containing the finger-up location) are then added to the region-of-interest.




Fig. 4. Diagrams of the three valid refinement gesture cases (removal, addition, and transfer) and their results. In (a) the user puts their finger down (green dot) within the existing selected region-of-interest (light gray), drags their finger (blue dotted line) across a number of sub-regions, and removes their finger (red dot) outside of the region-of-interest. The result of this gesture is that any already selected sub-regions that were passed over are removed from the region-of-interest. (i.e. they are “pushed” out.) In (b) the user’s touch instead begins outside of any existing region-of-interest, drags into one, then releases. The result is that any sub-regions passed over on the path inward are added to the region-of-interest. (i.e. they are “pulled” in.) Finally, in (c) the user has started within one region-of-interest and ended in another. The result is that any sub-regions passed over along the path are transferred to the second region-of-interest. Each finger can simultaneously be generating one of these gestures, allowing hands to sweep large swaths of territory into and out of regions-of-interest.

By handling individual touches in this manner, the refinement mode allows users to reshape regions-of-interest in a natural and intuitive way. Because each finger stroke is handled individually we can update the contribution of that action to provide immediate visual feedback of the resulting changes. Not only can multiple users immediately evaluate their refinements as they are making them, but this type of instant feedback encourages exploratory use and supports complex analytical usage. For example, the user might want to equalize a particular demographic statistic (e.g. income) between multiple adjacent regions-of-interest. Using our techniques, the user can watch the visualizations in the probe interfaces tied to those regions while iteratively manipulating the borders between them at increasingly smaller granularities to find optimal divisions between the regions..

3.4 Collaboration

For our particular application, our target users are expected to come and go randomly and there is limited opportunity for intricate manipulation of the data that would require *protective* personal space to complete. As such we chose to not enforce any strict limits on who can do what where. Returning to Scott’s^[6] observations about intuitive division of collaborative space, we can recognize that the map in our application is obviously the center point and the “group workspace” area. It is here that users share regions-of-interest and where finished glyph-like probe interfaces are often overlaid atop their associated regions. The “personal spaces” for each user are the initially empty spaces around the map directly within their arms reach. It is here that users privately analyze the contents of their regions-of-interest within their associated probe interfaces, before moving them into the shared workspace for distribution. Also, as Scott predicts, the periphery of the table, areas along the edges and out of users’ easy reach becomes a “storage space”. This space tends to fill up over time with “monitor” probe interfaces that have been set to keep an eye on particular regions for long periods of time. Those “monitor” probe interfaces that concern a single user tend move from their personal workspace into the periphery to make way for new personal analysis tasks, and those that are of interest to the entire group tend to be moved to the periphery from the center as the data they monitor becomes stale or less relevant to newer analytical results being shared in the central “group workspace.”

Relevant design implications identified in Scott’s research correspond with design choices made in this application. Providing appropriate table space is a concern. While the physical size of our table top is fixed beyond our control, we

include simple methods of resizing all interface elements to allow more content in less space. There are of course limits on how much one can shrink and stack interface elements to fit more onscreen, but fortunately there is also an inherent limitation on how many users can physically fit comfortably around a table.

It is noted that one should provide functionality appropriate to each type of workspace. We accomplish this by: tying all probe interfaces, in which analyses are conducted, back to their associated regions-of-interest, providing methods of reducing complexity and size for interfaces that one wishes to “store” on the periphery, and the ability to reduce probe interfaces to a glyph like state for overlay on the central map. Finally, the freedom of our probe-based interface allows for the casual grouping of items in the workspace to both allow users to customize their personal workspaces and freely share their findings with the group.

For our specific application, formally tracking *specific* users is not necessary; we are far more concerned with the location of users as they casually interact with the table. The general location of users around the table is needed to determine how interface elements, probe interfaces, and visualizations are oriented with respect to the table. Given this straightforward requirement, and taking into account our multi-touch table is based on the direct IR illumination technique, placing a camera with an IR-pass filter above the work surface is the most attractive solution. With an IR-pass filter installed, the camera, mountable on a telescoping arm or hanging from the ceiling will only see a brightly illuminated rectangle corresponding to the table’s surface. It is then simply a case of automatically calibrating a homography from this rectangle to screen coordinates and then watching with background subtraction for shadows entering the rectangle. When a touch event spawns or modifies an interface element, the shadow nearest to it can be followed back to its intersection with the edge of the rectangle/screen and this point on the perimeter can then be used to determine the necessary orientation for that interface element. Thus, when a user interacts with an interface element, it can automatically rotate to face them properly. This process could be further expanded to allow more complicated actions that would be needed in a collaborative environment, such as the sharing of information with other users: One user could drag a probe interface “into the hands” of another user, at which point it would then orient itself to face its new observer.

By making assumptions about the maximum number of simultaneous users (realistically under 6 for our size table), this solution can also be used to restrict selection and refinement tools/modes to the touches of individual users. For example, in a time-critical situation, User A could be making rough selections of areas they know need particular resources as they are request them, while User B examines and refines these selections to be realistic in terms of resources available and continuously updated return-on-investment figures, simultaneously User C is following along, making final approval or denial decisions as to the actions to take with each of the refined regions. In such a manner, a collaborative decision making process can be continuously undertaken from start to finish in the same analytical environment, in which all Users can see each others work in context with their own.

4. DISCUSSION

We presented our system to members of our target audience (urban planners, politicians, etc) at the 2009 Institute for Emerging Issues Forum, which was focused on growth and infrastructure needs. Individuals across a wide range of familiarity with computers were encouraged to use the system to investigate their “home” districts. Most found our interface straightforward and had no problems understanding the multiple tools available to select regions-of-interest at various granularities. Users were primarily interested in selecting regions-of-interest corresponding to either a municipal boundary or the more general “around” a particular feature (a request the free-form selection tool is invaluable for). Selecting large numbers of bounded regions, which previously (with a mouse) required either repeated clicking or a tiresome back and forth “scribbling” motion, was now possible for people to do with literally a wave of the hand.

The novelty of the touch-based interaction was effective at initially drawing people’s interest into the application, but on a more meaningful level, it allowed people to immediately step up and start using the application, without the social barrier of asking to have a turn using the mouse and/or keyboard. The collaborative nature made accommodating new users easy, as more experienced users could take a moment to guide them along or demonstrate a feature without having to abandon their own analyses already in progress. Reset buttons were included to quickly and conveniently clear out the no longer wanted regions-of-interest and probe interfaces that were left over after a user had walked away from the table, but if the system was to be left out in public unattended, say as part of a museum exhibit, some sort of automatic clean up routine would be most beneficial. This routine would seek to clear deserted personal workspaces while maintaining interesting analyses distributed in the shared workspace.



Fig. 5. Our system being demonstrated and interacted with by members of our target audience at the 2009 Institute for Emerging Issues Forum.

5. CONCLUSIONS

Collaborative visual analytic environments require consideration of a number of factors. One must consider the specific requirements of the problem domain, and what the overriding and imperative questions to be answered are. In our case of geospatial analysis, we are primarily concerned with Where? These questions are predominantly answered spatially in the form of regions-of-interest. Selecting these regions-of-interest is thus the main focal point when designing an interface for a geospatial analysis application. We aimed to design our selection methods to be fast, efficient, and intuitive. By creating simple basic actions, and allowing the user to initiate any number of these simultaneously, we empowered their hands with the ability to create and reshape regions interactively in an instinctive manner.

By understanding the ways in which users divide territory in table-top collaborative spaces, we can predict how the spaces we provide will be used. This allows us to provide features that will promote more efficient collaborative behaviors, such as sharing and storing of results.

The combined result of these considerations is a system that is powerful, but non-intimidating; a system that allows and encourages complex combinations of actions, but can be operated without any training.

REFERENCES

- [1] Benko, H., Wilson, A., and Baudisch, P., "Precise Selection Techniques for Multi-touch Screens," in ACM CHI Conference on Human Factors in Computer Systems, pp. 1263-1272, 2006.
- [2] Butkiewicz, T., Dou, W., Wartell, Z., Ribarsky, W., and Chang, R., "Multi-Focused Geospatial Analysis Using Probes," IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 6, Nov. 2008, pp. 1165-1172.
- [3] Dietz, P., and Leigh, D., "DiamondTouch: a Multi-user Touch Technology," in Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology, pp. 219-226, 2001.
- [4] Han, J. Y., "Low-cost Multi-touch Sensing through Frustrated Total Internal Reflection," in Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology, pp. 115-118, 2005.
- [5] Kunz, A. M., and Spagno, C. P., "Technical System for Collaborative Work," in Proceedings of the Workshop on Virtual Environments 2002, pp. 73-80, 2002.
- [6] Scott, S. D., Sleelagh, M., Carpendale, T., and Inkpen, K. M., "Territoriality in collaborative tabletop workspaces," in Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work, pp. 294-303.
- [7] Wilson, A. D., "TouchLight: an imaging touch screen and display for gesture-based interaction," in Proceedings of the 6th International Conference on Multimodal Interfaces, pp. 69-76, 2004.