Legible Cities: Focus-Dependent Multi-Resolution Visualization of Urban Relationships

Remco Chang*
UNC Charlotte
Department of
Computer Science

Ginette Wessel†
UNC Charlotte
College of
Architecture

Robert Kosara‡
UNC Charlotte
Department of
Computer Science

Eric Sauda§
UNC Charlotte
College of
Architecture

William Ribarsky¶
UNC Charlotte
Department of
Computer Science

ABSTRACT

Numerous systems have been developed to display large collections of data for urban contexts; however, most have focused on layering of single dimensions of data and manual calculations to understand relationships within the urban environment. Furthermore, these systems often limit the user's perspectives on the data, thereby diminishing the user's spatial understanding of the viewing region. In this paper, we introduce a highly interactive urban visualization tool that provides intuitive understanding of the urban data. Our system utilizes an aggregation method that combines buildings and city blocks into legible clusters, thus providing continuous levels of abstraction while preserving the user's mental model of the city. In conjunction with a 3D view of the urban model, a separate but integrated information visualization view displays multiple disparate dimensions of the urban data, allowing the user to understand the urban environment both spatially and cognitively in one glance. For our evaluation, expert users from various backgrounds viewed a real city model with census data and confirmed that our system allowed them to gain more intuitive and deeper understanding of the urban model from different perspectives and levels of abstraction than existing commercial urban visualization systems.

Keywords: Urban models, information visualization, multi-resolution.

Index Terms: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Hierarchy and geometric transformations; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

1 INTRODUCTION

Most existing urban model visualization systems focus on layering a few dimensions of data over a 2D map or a 3D model with a limited number of buildings. Often the layering uses colors to depict the data, which quickly limits the number of layers that the user can see at the same time before the combinations of layers become too complex to understand. More importantly, existing systems limit the user's interactions when focusing on specific regions of interest. Specifically, many systems allow the user to drag a bounding box around the area of interest for zooming in. This interaction diminishes the user's understanding of the selected region of interest in relation to the rest of the city both in the sense of spatial relationships and the underlying depicted data.

From interviews with architects and urban planners, we recognize that visualization of an urban model must occur on all lev-

*e-mail: rchang@uncc.edu
†e-mail: gwessel@uncc.edu
‡e-mail: rkosara@uncc.edu
§e-mail: esauda@uncc.edu
¶e-mail: ribarsky@uncc.edu

els of abstraction. For example, when the architects and planners are asked to describe New York City, the descriptions always range from a global level such as "New York is large, compact, and crowded," to the local level such as "the area that I lived in had a strong mix of ethnic groups." Furthermore, there is often a strong sense of relationship in the localized descriptions, "the community that I lived in is more heterogenous than the surrounding neighborhoods." These comments combined indicate that not only do urban visualization tools need to be multi-resolution, the tools also need to show relationships among neighborhoods in a focus-dependent manner.

Our approach is therefore quite different from existing ones. We build on the idea of urban legibility, which is a concept made famous in the 1960s by Lynch [11]. Rather than being just random collections of buildings, a city has certain parts that people intuitively understand and aggregate when describing it from different levels of abstraction. These understandings and aggregations are often based on people's tendencies towards neighborhoods of similar ethnicities and social backgrounds. Together, they form parts of the basis of the elements of legibility as defined by Lynch. Using these legibility elements, we build a tool that provides not only the spatial view but also an information display depicting abstract data such as demographic information, land use, etc. The spatial data is linked with the abstract data so that they maintain and provide the same understanding and aggregation through all levels of abstraction.

Using the tool we developed, *UrbanVis*, the user can find parts of a city that are defined in terms of their spatial layout or boundaries, and then explore their properties. How similar are the people living in a borough, district, or neighborhood? What is the distribution of ethnic groups throughout a city? Through these explorations, the user can begin to understand the properties of the city and envision how changes would impact the urban environment, not just in terms of the physical buildings, but also how such changes affect the social infrastructure. What happens to surrounding neighborhoods if we put a school here? How will changing an area from residential to commercial zoning affect the local economy?

Our approach is unique in that it builds an urban visualization on a clustering algorithm with the goal of providing physical and informational views to the user that are easy to understand from all levels of abstraction. By aggregating the data based on the elements of legibility, UrbanVis opens up many possibilities for exploration and re-examination of existing understandings of a city.

For the user evaluation, we surveyed fourteen experts with different occupational backgrounds ranging from real estate developers and urban planners to geographic information system (GIS) users. From this user evaluation we formally identify features of the system that were most useful to these professional urban experts as well as a range of possible future directions. We concluded that a majority of the participants believed our visualization tool enabled them to better perform their daily tasks as it provided new features that were not available in current commercial software systems.
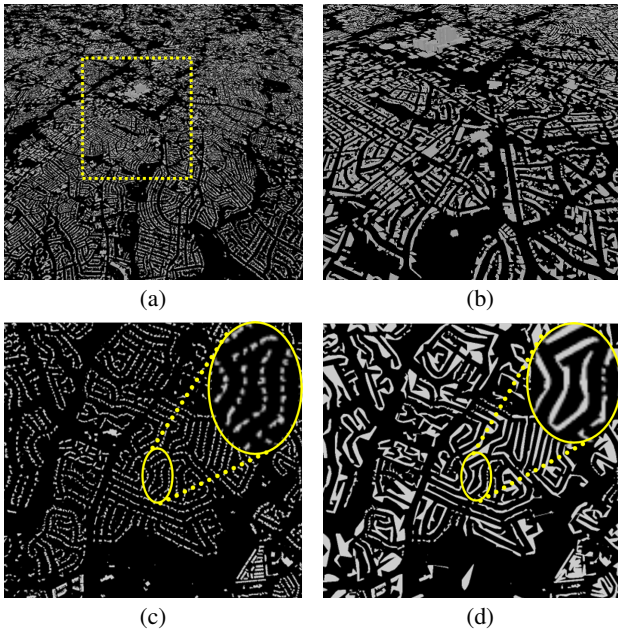
(a)   (b)

(c)   (d)

Figure 1: The building model aggregation used by UrbanVis. a) View of downtown Charlotte and its surrounding regions from afar; b) what is actually being rendered when the selected yellow box region is enlarged; c) using a pixel tolerance ε of 50; d) setting ε to 500: notice that as ε increases, so does the amount of aggregation. This results in larger, but fewer, clusters.

2   RELATED WORK

We build on work in urban planning and urban legibility, model simplification, and the connection between information visualization and geographical views.

2.1   Urban Legibility

Urban planning has focused largely on the use of social, economic, and political factors in evaluating urban growth and development [7]. The methodologies are adapted from the social sciences, and involve accumulation and analysis of complex data. There is relatively little emphasis given to the detailed form or geometry of the city. The work that does seek to connect social and political factors to urban form does so on a local rather than a city or regional scale [9].

Alternatively, urban design has focused on the form and geometry of the city. Traditionally, simple geometric models of the city have been the basis of discussion and design, either in planimetric view [13], in sequential perspective view [4] or using cognitive mapping [11]. There has been relatively little emphasis on policy factors in urban design, and little way to relate these issues to urban form.

There has been recent work seeking to explicitly link urban morphology and underlying economic, social and politics at all scales of urbanism. The work done at Harvard directed by Rem Koolhaas has sought to weave the economic, political and social factors explicitly into the development of urban form [3]. Mitchell, in his book *City of Bits* [12], wrote about the emergence of urban forms that will change from fixed ideas of space to shifting realms of intersecting digital and spatial networks. Neither of these efforts has been explicit about tools that will enable these new insights.

2.2   Model Clustering and Aggregation

Typical work in simplifying models does not take the specifics of buildings into account, changing their shapes in arbitrary ways (this

is the case for most traditional decimation algorithms [6]). Work that is more specific to models of urban environments is usually geared towards a specific use, like walk-throughs at ground level, but does not perform very well for overviews [17].

We utilize the urban model clustering and simplification method proposed by Chang et al. [1]. This method consists of two parts. During preprocessing, it computes a hierarchical spatial clustering of building models based on Lynch's ideas. This method is specific to city models, and produces much more recognizable results for this type of data than general mesh decimation or simplification algorithms (a more in-depth discussion is available elsewhere [2]).

During runtime, a view-dependent level-of-detail algorithm chooses the appropriate clusters to render based on a pixel error metric (Figure 1a/b). In a rendering application, this is measured relative to the eye point of the camera, which the user controls, as well as a global quality parameter ε. Low values of ε mean lower pixel errors and therefore higher accuracy and similarity to the original models; whereas high values of ε create larger and fewer clusters and lower number of polygons at the cost of visual fidelity (Figure 1c/d).

2.3   Information and Geographical Visualization

A survey of GIS tools is outside the scope of this section, but Geo-VISTA Studio [5] needs to be mentioned for its integration of classical GIS visualization with information visualization views like parallel coordinates [8].

A classic (and early) example of the combination of a geographical view with interactive querying is HomeFinder [16], which lets users find houses that fulfill certain criteria. The geographical view is only a map though, and there is no aggregation of data.

Shanbhag et al. [14] use visualization of demographic data over time to validate partitionings. This is very close in spirit to our work, but lacks any data on the physical layout of buildings, separators, etc.

3   SYSTEM OVERVIEW

The system uses two views (Figure 2): a 3D model view and a multi-dimensional data view. The views each have their own window, making window management easier on setups with two screens or projectors. The two views are fully linked and each accepts user interaction. The 3D model view shows clusters of buildings based on legibility elements and provides spatial awareness within the urban environment. The data view displays abstract information of the clusters shown in the 3D model view and adds an extra perspective for understanding the city. Together, the views allow the user to explore the urban model from both the geographical and the informational angles.

3.1   3D Model View

The 3D model view (Figure 2, right) shows the geometries of the buildings in the city, and thus acts as a navigation tool and the display for building clusters at the same time. The user can interactively navigate the city using either mouse or keyboard and view the city from any view distance or angle.

The focus which guides the aggregation of buildings is normally the eye point of the camera which the user controls. To decouple the clustering from the viewing, the eye point is represented by a yellow sphere which is connected to the ground with a thin line. The user can move the sphere around the map and also up and down to change the clustering: when the sphere is high above the ground, the cluster sizes are larger, allowing the user to see overviews of the entire area. When the sphere is lower to the ground, cluster sizes under the sphere are finer, allowing the user to inspect a specific local region (Figure 3). The focus region is thus not a fixed area, but varies with distance from the focus point directly under the sphere. The degree of focus is shown on the buildings themselves as a color
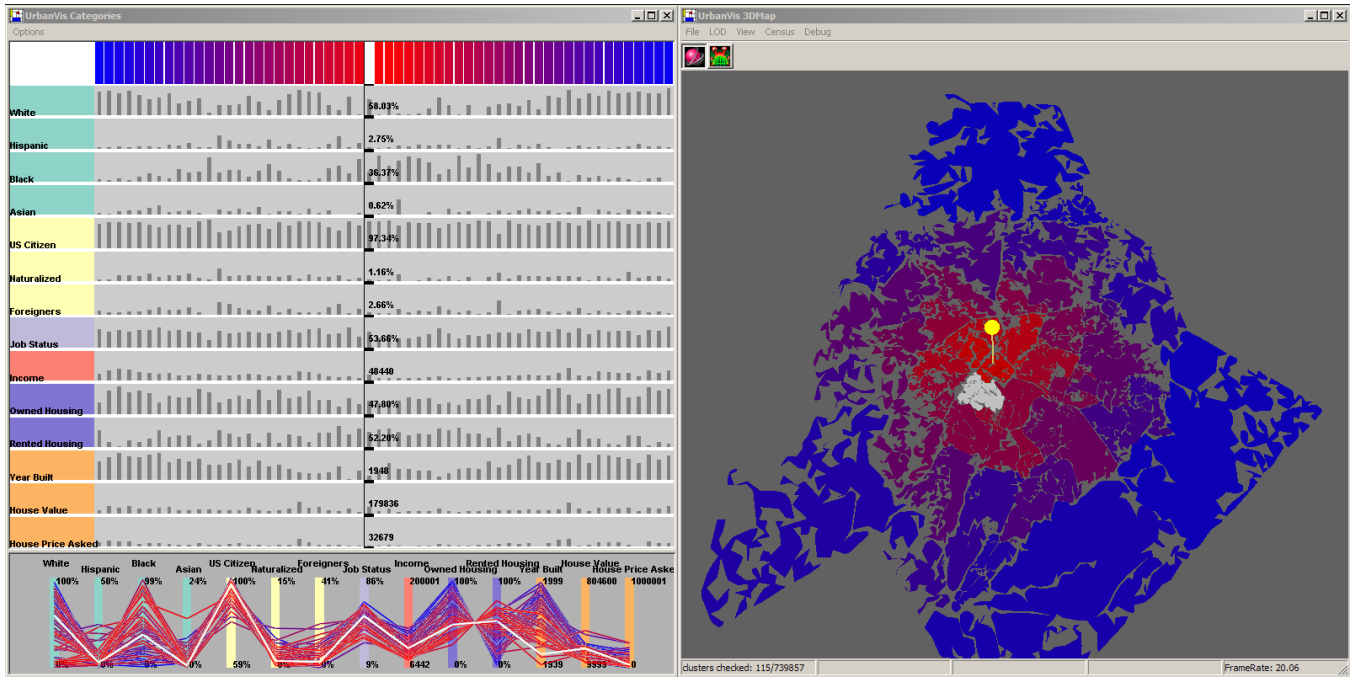
Figure 2: UrbanVis overview. The data view on the left shows demographic data of the areas around the focus point (focus in the middle). The model view on the right shows the clustered building models. The color gradient indicates the distance from the focus point, and provides a visual link between the two different data views (matrix view and parallel coordinates) and the model view. The data shown is census data for the city of Charlotte in Mecklenburg county, North Carolina. The straight lines in the lower half of the model view are where the city and county border South Carolina.


| (a) | (b) |

Figure 3: Changing the zoom level of the focal point (shown as a yellow sphere and a line connecting it to the ground). The color gradient from red to blue shows the proximity of the clusters to the focus. a) When the sphere is far away from the ground, the region of interest is larger, and the user can see an overview of the area at a glance; b) when the sphere is closer to the ground, the region of interest and clusters are smaller, thus allowing a more detailed inspection.


| (a) | (b) |

Figure 4: (a) User selects and highlights a cluster in the model (shown as white); (b) at any time, the user can change the view to looking at individual buildings instead of clusters.

gradient from red to blue. These colors provide a link between the two views (see below), and give the user an indication of how narrow or wide the focus currently is.

The user can select and highlight any cluster by double-clicking on it (Figure 4a), and also view the urban model as individual buildings rather than clusters for a closer inspection of a neighborhood (Figure 4b).

3.2 Data View

The data view (Figure 2, left) consists of two parts that display the same information in different ways: a matrix panel and a parallel coordinates panel. Both show the data associated with buildings or building clusters relative to the position of the focus point. In the

examples in this paper, the data is demographic data from the 2000 Census, but any geographically linked data (e.g., traffic statistics, crime rates, etc.) could be shown.

The top part of the data view can be switched between bar charts, line charts, or gradient charts (Figure 5). In any case, the view is organized in columns with each linked to a cluster. The columns are labeled with colors that correspond to cluster colors in the model view. The number of columns therefore changes dynamically with the number of clusters that are displayed as the user changes the level of detail or moves the focus around the city.

There are two orderings of columns. Under normal use, the clusters closest to the focal point are drawn in the middle of the view (Figure 6a), which corresponds to the usual way the model view is used, i.e., the user will keep the focus close to the center of the view, and recenter if needed. The user can also sort the columns depending on the values of a selected category of data for quick iden-
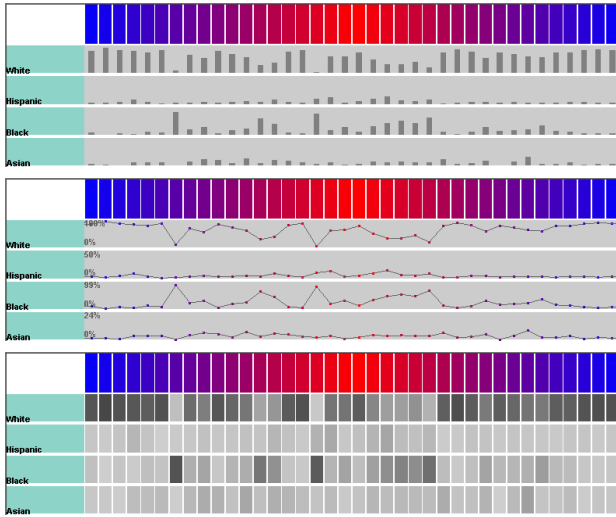
Figure 5: Different displays of the same data; (top) Bar charts; (middle) line charts; (bottom) gradient grid charts.


(a)   (b)

Figure 6: Sorting columns. a) normally, columns are sorted by the distances of their corresponding clusters to the focal point. The closer the clusters are to the focal point, the closer the column is to the middle of the screen (and more red in color); b) the user can also sort the columns based on a specific data dimension.

tification of the clusters with the desired value ranges (Figure 6b).

Each row of the bar/line/gradient charts shows a specific dimension of the represented data. The graphs are color-coded to show groupings of related categories, making quick identification and orientation easier. In Figure 6, there are 14 categories of data, separated into 6 different groups.

The bottom part of the data view shows the same data, but using parallel coordinates [8] to better show relationships between dimensions in the data. Like the matrix panel, the lines in the parallel coordinates view are color-coded to match the cluster colors, and the colors of the axes correspond to the colors of the rows in the matrix view.

Although the two views depict the same data, we find that the different presentations of the data give the user different types of understanding. The matrix view shows the relationship between clusters of buildings that are close to each other. The user can thus quickly see the homogeneity of the neighborhoods around the focal point. Sorting the matrix by a data dimension can also reveal correlations between data properties.

The parallel coordinates view cannot show spatial relationships, but can easily reveal relationships between data dimensions, allowing the user to easily identify positive or negative correlations between categories.

Figure 7: Using a slider to find buildings in the city that fit a specific criterion. In this example, only buildings that are built after 1985 are shown.

3.3 Dimension Thresholding

The tools presented so far are mostly tailored towards the exploration of an urban model from the model view. For tasks that have specific search criteria, such as looking for areas with high percentages of certain ethnicities, we employ a simple slider to highlight the clusters that match the given criterion (Figure 7).

As the user moves the slider, the model and data views update interactively, highlighting the clusters that fulfill the criterion. To maintain legibility, the other clusters are shown as well, but in a darker color.

4 APPLICATION SCENARIOS

In order for the participants of our study to understand how the application might apply in real world settings, we provided them with scenarios of how a user might interact with the urban visualization tool. By using actual demographic data taken from the United States Census 2000 [15] for the county of Mecklenburg in Charlotte, North Carolina, we were able to apply the data to the 3D building layout of the area. The demographics utilized in this specific demonstration cover various categories such as ethnicity, citizenship, job status, income, and housing statistics. However, the system is not limited to these categories and is configurable to each user's needs.

A simple scenario provided to the user allowed for an immediate understanding of the possible everyday uses of the visualization tool. For instance, according to Jeff Michaels, director of the Charlotte Urban Institute, the city of Charlotte's annual "Charlotte Neighborhood Qualify of Life Study" looks for areas of high ethnic population with low levels of income to identify possible improvements to these regions through urban planning. Using our system, we can quickly identify the regions in Charlotte that fit the two criteria (Figure 8). Furthermore, upon further inspection, we identify that there are some characteristics of these neighborhoods that are of interest to the Urban Institute. Specifically, by examining the parallel-coordinates, we find that the level of Hispanic populations in these areas have a positive correlation with the percentage of foreigners and the percentage of people who rent housing (Figure 8). The relationships between these categories of data is not easily obtainable using current commercial software. As Victoria Bott, Director of Land Use and Environmental Planning Division at UNC Charlotte said, "using current software requires going back and forth between ten different windows to find these relationships, whereas your system shows all those relationships in one simple view."

Another real life example was given by real estate developer Ed Harris of Harris Associates. In his occupation, identifying areas with homogeneity in demographics is often very important when negotiating radical or new concepts in urban planning. In his experience, local governments of areas with high homogeneity in demographics are more likely to accept new concepts because of their

shared demographic background. However, areas with high heterogeneity often result in disagreements between the different demographic groups due to their differences in perspectives. Using our system, we allow the user to quickly identify the level of homogeneity of downtown Charlotte versus the town of Davidson (20 miles north of Charlotte) where the company is located (Figure 9). The figure indicates that downtown Charlotte is heterogeneous in demographics and citizenship status, whereas Davidson is much more homogeneous.

5 USER EVALUATION

We asked 14 expert users to evaluate our system from their own perspectives and identify the strengths and weaknesses of our system. These 14 experts have disparate backgrounds, ranging from the Center for Real Estate at UNC Charlotte (Dr. Steve Ott), the UNC Charlotte Urban Institute (Dr. Jeff Michael, Dr. Victoria Bott, and Charlynn Burd), Charlotte Mecklenburg County Geographic Information Systems Office (Kurt Olmsted, Todd Wilson, Tobin Bradley, Andy Goretti, and Paul Martin), Planning Department (Lori Quinn), and School System (Christine Hamlett and Jennifer Dean), and independent real estate development (Ed Harris and David Stewart). In the study, we first asked the experts to fill out a pre-test questionnaire that identified their backgrounds in urban studies and their proficiency levels with geographic information systems. Then we demonstrated features of our system, followed by a few simple scenarios in finding interesting characteristics of the census data in Charlotte. After the demonstrations, we asked the experts if our tool could be used in their areas of expertise. Finally, we concluded by asking them to give feedback on the usefulness of the system as well as any suggestions for future improvements. With their consent, the users were tape-recorded during these sessions.

Focus-Dependent and Dynamic Clusters

All but one expert agreed that the focus-dependent view with dynamic clusters helps in understanding not just the region of interest, but also its surrounding areas. Christine Hamlett of Charlotte Mecklenburg Schools commented that using this technique would allow her to focus on the potential sites of a new school, and still show the "projections of future student populations based on surrounding new housing developments." Jennifer Dean continued to add, "new housing developments often impact the existing school systems in terms of student population and demographics" and implied that the dynamic clustering helped in visualizing the changes and seeing the potential new effects.

The one user who didn't find this technique useful commented that most projects he worked on had strict boundary requirements. With these restrictions, it had never been necessary for him to examine surrounding areas.

Integrated Displays

On the use of the two integrated views between the 3D Model View and the Data View, 13 out of 14 of the participants found the combination of the two to be useful. Tobin Bradley attested that "[the integrated displays] are an asset in handling large amounts of data and faster user production rates because it provides a link between the 3D urban model and the data display." As Jeff Michael succinctly put it, the dual views provided the "here's what I'm looking for, and there's where it is" capability to understanding urban data relationships.

The only expert who did not find the integrated displays to be useful mentioned that GIS experts had been successful for years in using single displays and that having multiple screens sometimes caused confusion in which screen to focus.

Multi-Dimension Visualization

Displaying multiple dimensions of data using the matrix view and the parallel coordinates allow the user to quickly see relationships within the urban data. Lori Quinn asserted that "an asset to this visualization tool is that the selection of data makes demographic relationships instantly apparent in the area of specificity. In current systems, you have to design [the necessary queries], analyze them, and then modify the queries to find the correlations that you are looking for." Kurt Olmsted further added, "Sometimes users have to go through a lot of different sources of data or running [statistical analysis] to find relationships. Your tool is providing an on-the-fly, interactive way of instantly noticing nearby statistical data and their relationships." All of our users shared the same sentiments and found the Data View to be useful, although one user commented that the Data View required some explanation before the relationships in the data became apparent.

Victoria Bott summarized the strengths of our system eloquently. "Essentially what you are providing with this tool is a spatially sensitive graphic display. The strength of this tool is the dynamic table that displays areas in a spatially understandable way. In other software systems, the user is required to scan the tabularly listed rows of a GIS database, which gives no indication of the rows' geospatial locations or their relationships between one another. Another strong aspect is the fact that your focus area and peripheral areas are cohesively orientated. When that aspect is combined with the ability to change the level of detail through clustering, the user gains a new dimension [of understanding]. Changing the level of detail in other software programs becomes cumbersome from running [multiple repetitive] queries."

Collectively, expert users saw the potentials of our system both academically and commercially. Academically, they recognized that our system offered an entirely new perspective in studying urban landscapes and felt the tool provided them with vivid mental maps of their own spatial awareness in an urban environment. Commercially, users believed that this system can help increase productivity and provide better execution of their daily tasks by substantially improving the way they interact with GIS databases.

6 DISCUSSION AND FUTURE WORK

A few very interesting and important topics emerged from the discussions with our expert users. The most controversial is the results of dynamic clustering, on which we received informative but disparate feedbacks. There are a few experts who found the clustering to be confusing. Specifically, experts who work with the census data expected the clusters to follow the boundaries of the census tracts. Similarly, experts who work with zip codes or school districts wanted to see the clusters form boundaries and shapes that they are familiar with. In contrast, other experts praised the results of the dynamic clusters as the clusters provided "possibilities to new district boundaries" that one might not have been aware of.

While all experts agree that they retain spatial understanding of the city of Charlotte through all levels of simplification using our system, the question of what makes an urban model legible remains open. The one key point that everyone can can agree on is that the sense of legibility is very subjective and changes depending on each individual's level of understanding and perspective of a city. As mentioned before, urban experts with specific domain knowledge form their sense of legibility based on their domain of expertise; residents of a city orient themselves based on familiar establishments (such as a local restaurant) that might not be visually or mentally important to others. For tourists or visitors to a city, visually distinctive landmarks such as skyscrapers or major roads are important features for understanding the surroundings. Conversely, soldiers in an urban battlefield require a different set of training and understanding of a city to effectively communicate spatial relationships in a dynamic environment [10].

Figure 8: Case study 1: Finding neighborhoods with high Hispanic population near the downtown area: (right) the user starts by putting the focal point over the downtown region; (left) using brushing in the parallel coordinates window, the user highlights the regions that have high Hispanic population. Notice the positive correlation between the Hispanic population, the percentage of foreigners, and the percentage of residents who rent their housing in these selected areas.
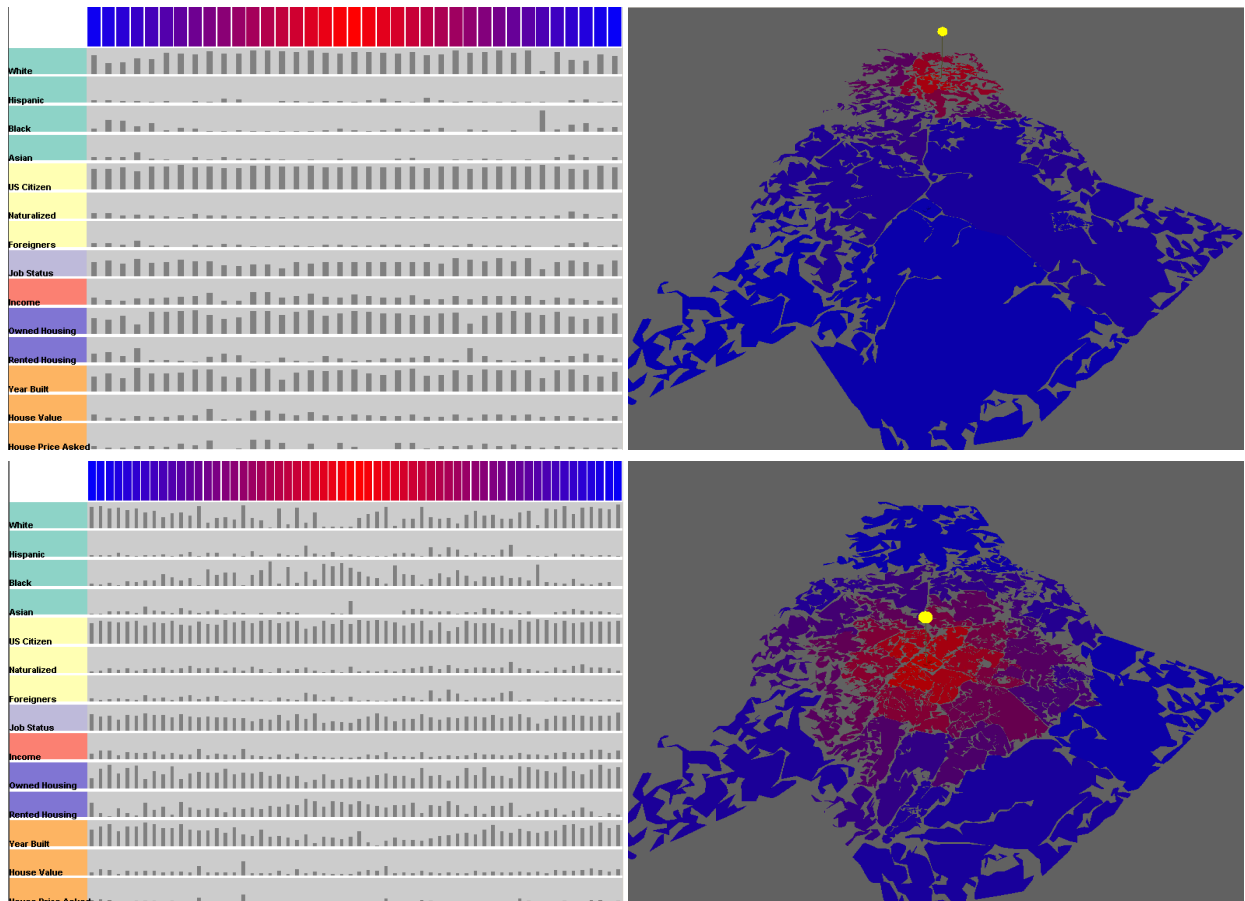

Figure 9: Case study 2: Showing the differences in the amount of homogeneity between downtown Charlotte (bottom) and the Davidson area. (top) Davidson area: notice that the bars in the Data View are all approximately the same height around the red clusters, indicating that the neighborhoods in the Davidson area tend to be more homogenized, whereas (b) in the downtown area of Charlotte, the differences between surrounding neighborhoods are more apparent.

Although our approach uses geometry to create clusters that are understandable to all users, we note that creating *legible* cities to users of all backgrounds is not a trivial task and would require knowledge of the user's perspective of the city prior to creating the clusters. Finding interactive methods so that these clusters could be tailored towards the need of each individual user remains an important future direction of research for us.

Another topic of discussion that was commented on by almost all experts is the need to integrate data from different sources. For real estate developers, seeing how commercial buildings and public establishments such as drug stores and schools intertwine with residential neighborhoods is important in identifying the needs of the neighborhoods. Members of the Charlotte Mecklenburg GIS Office noted that seeing tax records of individual buildings on top of the Census data would give better understanding of the economic development of the city. Similarly, experts from the Charlotte Mecklenburg School system would like to see Census information blended with crime statistics to find better routes for students and school buses. Since our system performs clustering on a per-building level, assigning specific properties from different data types to each building is trivial. However, finding the necessary data and identifying the best way to represent sometimes conflicting sources of data requires more investigation. For example, commercial districts without any residents would not have any census information, and purely residential districts would not contain information on economic growth of the area. Integrating these two orthogonal sources of data into a cohesive view is important in enhancing the user's ability in seeing the patterns and relationships between the data.

Along the same lines, many users noted that seeing temporal changes in a city would be very interesting. While we agree that time is a very relevant factor in urban visualization, the challenge lies in the collection of data such that the 3D models of buildings and the additional sources of data match both spatially and temporally.

Lastly, some experts suggested a potential use of this system outside of our original design goals. Jeff Michael and Victoria Bott of the Charlotte Urban Institute mentioned using our system as a tool to compare between different cities. Specifically, urbanists have widely accepted that Charlotte as an emerging southern city has mimicked the growth of Atlanta due to their similarities in locale and culture. It would be interesting to juxtapose the two cities in our system and see if such patterns of similarities are apparent. On a local scale, independent real estate developers Ed Harris and David Stewart mentioned a similar use of our system. Namely, it is important for developers to foresee pockets of potential growths in a city. For any given developing region, if a developer can identify another similar but already established region in the city, the developer might be able to project the potential growths of the developing region based on the history of the established one. Although the idea of using our system as a predictive tool is still being investigated, we are very excited about the potential benefits that it could bring.

7 CONCLUSION

We introduce an interactive tool to visualize an urban model in a focus dependent and multi-resolution fashion both geometrically and concerning the underlying data while retaining the legibility of the city. Throughout the user's exploration of the urban model, the system allows the user to maintain spatial awareness of the focus area as well as the peripheral areas.

Cohesively integrating the 3D Model View and the Data View allows the user to see the relationships between the geospatial information of the urban model with the related urban data such as the census information. The Data View further shows multiple categories of data in one glance, which is an improvement over existing commercial software when exploring urban models as it helps the user to easily identify correlations between the categories.

As indicated by the experts in our user study, our system contains features that fundamentally change the way users would interact with urban data, which in turn enhances their ability to better understand the urban model. With the addition of more task specific data, we look forward to the expert users using our system in their daily tasks.

REFERENCES

[1] R. Chang, T. Butkiewicz, C. Ziemkiewicz, Z. Wartell, N. Pollard, and W. Ribarsky. Hierarchical simplification of city models to maintain urban legibility. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches*, page 130. ACM Press, 2006.

[2] R. Chang, T. Butkiewicz, C. Ziemkiewicz, Z. Wartell, N. Pollard, and W. Ribarsky. Hierarchical simplification of city models to maintain urban legibility. Technical Report CVC-UNCC-06-01, Visualization Center, University of North Carolina at Charlotte, 2006.

[3] C. J. Chung, editor. *Great Leap Forward*. Taschen, 2001.

[4] G. Cullen. *The Concise Townscape*. Butterworth Architecture, 1961.

[5] M. Gahegan, M. Takatsuka, M. Wheeler, and F. Hardisty. Introducing GeoVISTA Studio: an integrated suite of visualization and computational methods for exploration and knowledge construction in geography. *Computers, Environment and Urban Systems*, 26(4):267–292, 2002.

[6] M. Garland and Y. Zhou. Quadric-based simplification in any dimension. In *ACM Transaction on Graphics, Vol. 24, No. 2*, pages 209–239. ACM, 2005.

[7] A. Garvin. *The American City: What Works and What Doesn't*. McGraw Hill, 2002.

[8] A. Inselberg and B. Dimsdale. Parallel coordinates: A tool for visualizing multi-dimensional geometry. In *IEEE Visualization*, pages 361–378. IEEE CS Press, 1990.

[9] H. Lefebvre. *The Production of Space*. Blackwell, 1991.

[10] M. Livingston, L. Rosenblum, S. Julier, D. Brown, Y. Baillot, J. Swan, J. Gabbard, and D. Hix. An augmented reality system for military operations in urban terrain. In *Proceedings of the Interservice / Industry Training, Simulation, and Education Conference*, page 89, 2002.

[11] K. Lynch. *The Image of the City*. The MIT Press, 1960.

[12] W. J. Mitchell. *City of Bits: Space, Place, and the Infobahn*. MIT Press, 1996.

[13] C. Rowe and F. Koetter. *Collage City*. MIT Press, 1978.

[14] P. Shanbhag, P. Rheingans, and M. desJardins. Temporal visualization of planning polygons for efficient partitioning of geo-spatial data. In *Proceedings Information Visualization*, pages 211–218. IEEE CS Press, 2005.

[15] U.S. Census Bureau. United States 2000 Census. http://www.census.gov, 2000.

[16] C. Williamson and B. Shneiderman. The dynamic homefinder: Evaluating dynamic queries in a real-estate information exploration system. In *Proceedings Information Retrieval*, pages 338–346. ACM Press, 1992.

[17] P. Wonka, M. Wimmer, K. Zhou, S. Maierhofer, G. Hesina, and A. Reshetov. Guided visibility sampling. In *Proceedings SIGGRAPH*, pages 494–502, 2006.