Balancing "Individual Privacy" and "the Public's Right to Know" or How to Properly Display the Location of Confidential Crime Data on Publicly Available Internet Sites

Michael Leitner
E111 Howe-Russell Bldg.
Department of Geography and Anthropology
Louisiana State University
Baton Rouge, LA 70803
Office: (225) 578-2963
Fax: (225) 578-4420
E-mail: mleitne@lsu.edu

Abstract
This research proposes cartographic guidelines for presenting confidential point data on maps. Such guidelines do not currently exist, but are important for governmental agencies that disseminate personal data to the public because these agencies have to balance between the citizens' right to know, and preserving a citizen's right to privacy.

In an experiment, participants compared an original point pattern of confidential crime locations with the same point pattern being geographically masked. Three different local masking methods at three different resolutions were tested. The objective was to identify appropriate geographic masking methods that preserve both the confidentiality of individual locations, and the essential visual characteristics of the original point pattern. The empirical testing reported here is a novel approach for identifying various map design principles that would be useful for representing confidential point data on a map.

1. Introduction
The advent of crime mapping via Internet sites has allowed law enforcement agencies to quickly and easily inform citizens about crime occurrences in their neighborhoods. Although this has satisfied the public's right to know, it has also created new problems, especially when confidential crime data are being mapped and citizens fear their rights to privacy being violated.

The purpose of this paper is to discuss appropriate guidelines to visualize the location of crime records that protect the confidentiality of individuals, and at the same time preserve essential visual characteristics of the original spatial distribution. This is important because the violation of an individual's privacy could have legal implications for the agency distributing this information. Additionally, consequences of not preserving the essential visual characteristics of the actual crime distribution may lead to unfair informal redlining employed by insurance and banking companies.

A common strategy used by many law enforcement agencies to protect the location of confidential crime data is by manipulating the location of such data, a method that is referred to as geographic masking. The goal of geographic masking is to manipulate crime location records just enough to protect the confidentiality of individuals, but not too much to change essential visual characteristics of the true, original spatial distribution of those records.

Armstrong, Rushton & Zimmerman (1999) introduced the term 'geographic masking' into the literature, suggesting similar geographic masking methods that are applied in this research. Whereas their research discusses the influence geographically masked data have on the results of geographically based analyses, the research reported here identifies acceptable design solutions for presenting confidential point data on a map. Acceptable design solutions define any geographic masking method that would preserve as many visual spatial characteristics as possible, while reducing the likelihood of individual identification to an acceptable level.

Leitner & Curtis (2003, 2004) tested ten different geographic masking methods and found only two of them to provide satisfactory design solutions. These include local masking methods that aggregate incident locations at either (1) the midpoint of the street segment or (2) at the closest street intersection. Their research found global geographic masking methods to be inappropriate for displaying the location of confidential point data. In contrast, this research tests three different local masking methods that displace incident locations within grid cells of a regular grid overlaid on top of the study area. Local masking methods include random translation, random rotation and random flipping. Three different regular grids are used with increasing cell sizes, including 100m, 200m, and 350m per side. The objective is to explore the influence different local masking methods, cell sizes and background information ("empty" background, census tracts, street network) has on people's perception when comparing the locally masked with the original, unmasked point pattern. The research reported here applies the same data, study area, and testing instruments as previous research by Leitner & Curtis (2003, 2004) in order to allow direct comparison of the results.

A review of the literature suggests that the methods reported here and in previous work (Leitner & Curtis 2003, 2004) are the first to utilize empirical perceptual studies to assess methods for presenting confidential point data on maps. This work also continues a long-standing tradition of empirical research in map design as a paradigm for eliciting and formalizing cartographic design knowledge (Leitner 1997, Leitner & Buttenfield 2000, Aerts, Clarke & Keuper 2003).

The subject matter of this research relates to an area of geography and related disciplines that has been receiving a fair amount of attention lately, especially as it pertains to the mapping of disease and crime information (Armstrong, Rushton & Zimmerman 1999, Wartell & McEwen 2001, Monmonier 2002, Leitner & Curtis 2003, 2004). The intention of the research reported here is to develop appropriate guidelines for mapping the location of individual-level data that is considered to be confidential. Such guidelines do not currently exist, but their development becomes increasingly important as more and more governmental agencies disseminate their data on maps via the Internet. To put it differently, this research is concerned with masking locations of individual-level data, rather than the attribute information that could be associated with such locations. Masking strategies for attribute data have already been widely discussed in the literature for some time (Duncan and Pearson 1991, Cox 1994, 1996).

Results of this research are of interest to any agency needing to display confidential data at neighborhood levels. As data become more widely available on the WEB, and as the public increasingly realizes the power of local level rather than global level aggregate maps, the need for accurate and easy geographic masking of confidential data will be of utmost importance to all governmental agencies. According to a website updated by the Mapping and Analysis for Public Safety (MAPS) program (sponsored by NIJ), about 50 local law enforcement agencies in the US now provide online data/maps (see http://www.ojp.usdoj.gov/nij/maps/weblinks.html). Access to this type of data will certainly continue to increase.

2. Local Geographic Masking

The incident locations displayed in the test maps were a subset of all residences of homicide victims in the city of Baton Rouge between 1991 and 1997. This is the same data set and study area used in Leitner & Curtis (2003, 2004). This subset of incident locations was displayed in yellow on top of a blue background. In one-third of the test maps, no additional information was shown. In the second one-third, census tract boundaries were added to the test maps in black. The remaining one-third of the test maps included the complete street network in dark blue.

Incident locations were geographically masked at the local level at three different resolutions. This was accomplished by superimposing regular grids with increasing cell sizes on top of the study area. The finest grid possessed cells measuring 100 meters per side; a second grid had cells with 200 meters per side; and the coarsest grid measured 350 meters per grid cell side. Local geographic masking means that incident locations falling into a cell of the same regular grid were spatially displaced by the same exact amount, but this displacement vector changed randomly between grid cells. Three different local masking methods were tested and analyzed in this research, including random flipping, random rotation, and random translation (Figure 1).

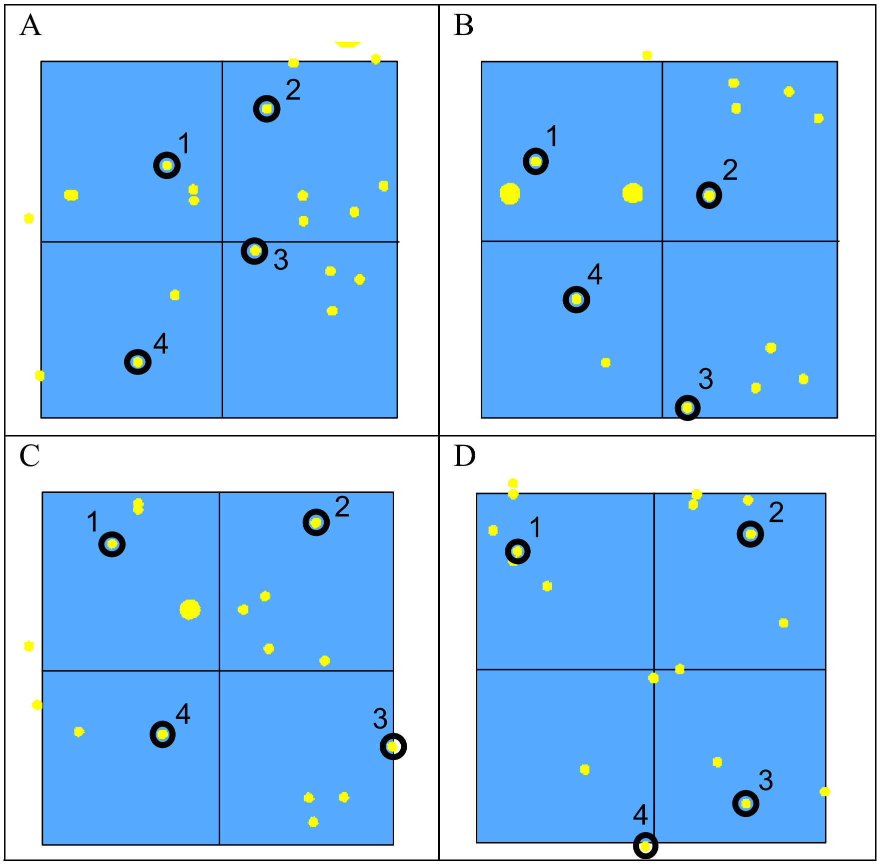

Figure 1. Local geographic masking methods used in the experiment (only a portion of the entire test map is shown) A) Original incident locations, B) flipping randomly either about the vertical, horizontal or both central axes of each grid cell, C) rotating by some random degree around the center of each grid cell, D) translating by some random distance. Adapted from Leitner & Curtis (2004).

Local random flipping means that incident locations were flipped either about the vertical, horizontal, or both central axes of each grid cell and then moved on top of its closest street segment. The type of flipping changed randomly between the cells of the regular grid. For example, in Figure 1-B the incident location labeled as #1 was flipped vertically; the other three incident locations (#2, #3 and #4) were each flipped horizontally. Local random rotation involved rotating incident locations by some random degrees around the center of each grid cell and then placing them on top of the closest street segment. In Figure 1-C the incident location #1 was rotated by 120º to the left; #2 by 60º to the right; #3 by 240º to the right; and #4 by 120º to the right. Local random translation means that incident locations were first translated by some random distances and then moved on top of their closest street segment. In Figure 1-D, the incident location labeled as #1 was translated 150 meters in x- and 350 meters in y-direction from the lower left corner of its grid cell; #2 was translated 300 meters in x and 400 meters in y; #3, 450 meters in x and 50 meters in y; and #4, 300 meters in x and 200 meters in y.

The angles of rotation and the translation distances were chosen arbitrarily when masking point patterns. Clearly, different choices might have yielded different results. For example, smaller rotation angles and shorter translation distances in Figures 1-C and 1-D would have resulted in shorter point movements. In addition, smaller cell sizes lead to shorter point movements and vice versa. Investigating the impact differently masked point patterns and different sizes of grid cells have on people's visual perception should be understood as an exercise in exploratory analysis. It is hoped that results of this and similar previous research (Leitner & Curtis 2003, 2004) will create a base line for the development of appropriately displaying the location of confidential point data in maps that are available on publicly available Internet sites.

3. Structure of the experiment

The experiment included 42 map pairs. One map of each pair always showed all incidents in their correct location. In the second map, all incident locations were geographically masked at the local level using different cell sizes. All possible combinations of three different backgrounds (no background, census tract boundaries, and street network), three local masking methods (random flipping, rotation, and translation) and three grid cell sizes (100m, 200m, and 350m per side) were tested in the experiment. All combinations accounted for 27 of the 42 map pairs included in the experiment. Each of the remaining fifteen map pairs included an unmasked version of the incident locations together with the same incident locations being masked globally. Global geographic masking means that every single incident location in the study area was spatially displaced by the same exact amount. Five different global masking methods were used, including global flipping around the horizontal central axis, the vertical central axis, both central axes, and rotating around the map center either by 60º to the right or by 120º to the left. These were the same global masking methods analyzed in Leitner & Curtis (2003, 2004). They were included in this experiment in order to provide a similar variability in masking methods allowing results of this research to be compared to results from previous research (Leitner & Curtis 2003, 2004). Each of the five global masking methods was tested for each of the three different backgrounds. For the 27 map pairs that included the locally masked incident locations, participants were asked to identify hot spot areas in the masked point pattern. However, for the remaining fifteen map pairs participants were asked to identify hot spots in the original or unmasked point pattern, rather than in the globally masked point pattern. For all map pairs in the experiment, participants did not know which point pattern was masked and which one was unmasked. The order of presentation of the map pairs in the experiment was random.

Each experiment was put together as a power point presentation. In addition, all participants received a color printout (hard copy) of the entire presentation. The first slide of the presentation included the instructions (see Figure 2); all other slides included the map pairs; one map pair per slide (Figure 3). Instructions were briefly repeated at the bottom of each slide showing a map pair. Overall, 62 participants completed the experiment. All participants were students from four different Geographic Information (GI) Science classes being taught in the Department of Geography and Anthropology at Louisiana State University in Baton Rouge, LA during the Fall 2004 semester.

Experiments were conducted during class time, lasting between 10 and 35 minutes. For each map pair, participants were asked to complete two tasks. The first task was to compare the two point pattern maps and decide if the two patterns were similar or different. Participants were asked to rank the two point patterns as "very similar" (rank value of 1) to "very different" (rank value of 7). This number was recorded on the hardcopy maps. The second task involved identifying areas within one of the two point pattern maps that showed a high concentration of incidents. This was either the map for which the incident locations were geographically masked at the local level or the original unmasked map for each of the fifteen map pairs that included the globally masked incident locations. Test participants were asked to mark those areas, if they thought they existed, with a pen or pencil directly onto the hardcopy maps. Also, test participants did not know, which map from each pair was geographically masked and which one was unmasked.

Each slide in this presentation shows two point pattern maps that are displayed next to each other.

Your first task is to compare the two point patterns and decide how similar or different they are.

Please make this decision by choosing a whole number between **1** and **7.** The number **1** means that the two point patterns are **VERY SIMILAR** and the number 7 means that they are **VERY DIFFERENT.**

Your second task is to identify areas within one of the two point pattern maps that show **a high concentration of points**. Mark those areas with a pen or pencil directly on the hard copy maps provided.

There is no time limit to complete this experiment.

Thank you for your participation!

Figure 2. First slide of the experiment, including the instructions. Courtesy of Leitner & Curtis (2004).

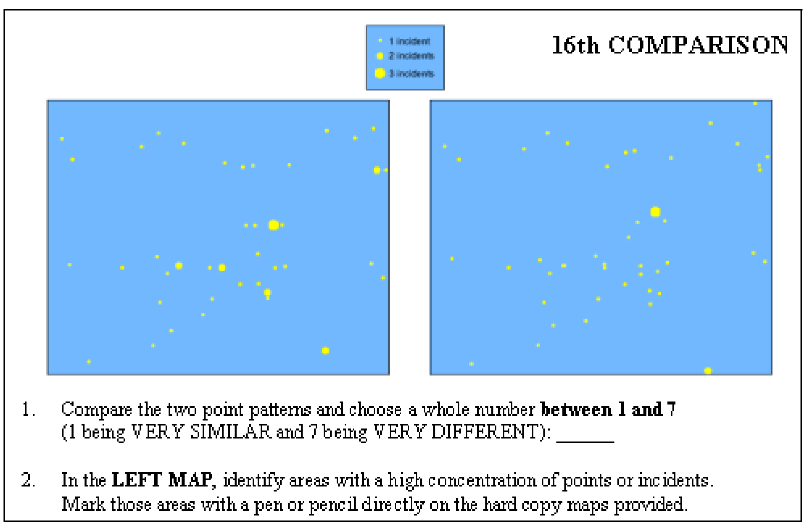

Figure 3. Example of a pair of test maps included in the experiment. All incident locations in the left map are locally masked by spatially aggregating them to their closest street intersection. All incidents in the right map are shown with their true, unmasked location. Courtesy of Leitner & Curtis (2004).

4. Analyzing the influence of different local geographic masking methods

One objective of this research is to investigate the influence different local geographic masking techniques, map backgrounds and cell sizes has on the perceived similarity between a map of original (i.e., geographically unmasked) incident locations that are confidential and the same incident locations being geographically masked. This objective is addressed in two ways: first, by comparing the participant's perceived similarity between the masked and the unmasked incident locations for each map pair; secondly, by comparing the participant's perceived number, sizes, shape and location of hot spots between the same two incidents maps for each map pair. Test participant's responses to how the original incident locations compared to their masked counterparts are summarized in Table 1.

Cell Size	Translating by some random distance within each grid cell	Rotating by some random degree around the center of each grid cell	Flipping randomly either about the vertical, horizontal or both axes of each grid cell
500x500m[1]	4.99	3.64	4.34
350x350m	5.08	4.08	4.20
200x200m	3.60	3.08	3.25
100x100m	3.01	2.77	2.94
	Spatial aggregation at midpoint of street segment[1]	**Spatial aggregation at street intersection[1]**	
	1.80	2.43	

Table 1. Test participant's perceived similarity between the original and the masked point pattern for different local geographic masking methods and cell sizes, irrespective of background information. All entries are mean ranks based on a scale of perceived similarity ranging from very similar (mean rank of 1) to very different (mean rank of 7). [1] Results were taken from Leitner & Curtis (2004).

Two main observations can be made from the results in Table 1. First and as expected, the mean rank measuring the perceived similarity between the original and the locally masked incident locations decreases with smaller cell sizes. The closer the mean rank is to 1, the more similar the original and masked incident locations were perceived and the closer the mean rank is to seven, the more different the two point patterns were perceived. A mean rank lower than four indicates that the two point patterns were perceived to be more similar than different from each other. Accordingly, mean ranks above four indicate that the two point patterns were perceived to be more different than similar. The results in Table 1 clearly show that when geographic masking is performed within cells that are 200 meters per side or shorter than the masked incident locations are always perceived to be more similar (than different) compared to the original (unmasked) incident locations. The same is true for incident locations that are geographically masked by aggregating them at the midpoint of their street segment or by aggregating them to their closest street intersection. The results from the latter two local masking methods were taken from Leitner and Curtis (2004). To sum up, when incident locations are geographically masked by overlaying a regular grid on top of the study area, then the grid cell should not exceed 200 meters per side. Grid cell sizes of 350 and 500 meters per side are too coarse and would yield inappropriate masking results. This is an important finding that should be considered when masking the location of confidential point data.

A second observation is that the selection of the masking method is important, insofar as the mean rank varies greatly between the three masking methods. However, these differences become smaller with smaller cell sizes. A Kruskal-Wallis H test reveals that significant differences exist between the three masking methods for each cell size, with the exception of the smallest cell size. The Kruskal-Wallis H test yielded a chi-square statistic (p-value) of 88.835 (0.000), 49.332 (0.000), 14.748 (0.001), and 2.746 (0.253) for cell sizes of 500x500, 350x350, 200x200 and 100x100 meters, respectively. Irrespective of cell size, local rotation consistently yielded the lowest mean rank, whereas local translation always yielded the highest mean rank. This could be explained by the fact that among the three local masking methods, rotation preserved the relative positioning among the incident locations within each grid cell the most, translation preserved it the least.

5. Analyzing the influence of different base map information

The previous section identified local masking methods and cell sizes that were appropriate for the display of confidential point locations, if the (perceived) similarity between original and masked incident locations is a priority. This section analyzes the influence different base map information might have on the participants' perception of the similarity between the original and locally masked point patterns. Results in Table 2 reveal that the mean rank changes somewhat across the three base maps for each of the three masking methods and four cell sizes. There does not appear to be a clear trend of one particular base map consistently showing the lowest or highest mean rank. Base maps that include census tracts possess a (relatively) high number of low mean ranks, whereas empty background maps show a (relatively) high number of high mean ranks. The results in Table 2 further show that for each masking method, the mean rank variability across the three base maps increases with larger cell sizes. The question that needs to be answered now is, are these base map variations statistically significant for any of the three masking methods and cell sizes?

	Translating by some random distance within each grid cell			**Rotating by some random degree around the center of each grid cell**		
Cell size	**No base map**	**Census tract boundaries**	**Street network**	**No base map**	**Census tract boundaries**	**Street network**
500x500m[1]	4.61	4.66	5.28	3.10	3.93	3.90
350x350m	5.15	4.98	5.10	4.08	4.08	4.06
200x200m	3.79	3.48	3.52	2.79	3.24	3.20
100x100m	2.97	3.10	2.95	2.84	2.65	2.82
	Flipping randomly either about the vertical, horizontal or both axes of each grid cell			**Alternative local masking methods that are independent of grid cell sizes[1]**		
Cell size	**No base map**	**Census tract boundaries**	**Street network**	**No base map**	**Census tract boundaries**	**Street network**
500x500m	5.02	3.72	4.70	**Spatial aggregation at midpoint of street segment**		
350x350m	4.18	4.05	4.39	2.04	1.67	1.70
200x200m	3.32	3.19	3.24	**Spatial aggregation at street intersection**		
100x100m	2.90	2.90	3.00	2.40	2.32	2.59

Table 2. Test participant's perceived similarity between the original and the masked point pattern for different local geographic masking methods, cell sizes and background maps. All entries are mean ranks based on a scale of perceived similarity ranging from very similar (mean rank of 1) to very different (mean rank of 7). [1] Results were taken from Leitner & Curtis (2004).

To this end, a Kruskal-Wallis H test was applied to explore if test participants' responses varied significantly due to different base map information. This test was carried out for any combination of cell sizes and masking methods included in Table 2. Results revealed that only two of all fourteen combinations (three local masking methods multiplied by four cell sizes, plus two alternative local masking methods) showed statistically significant differences (p-value < 0.01) across the three different base maps. The two combinations included local rotation and local flipping for grid cells with 500 meters per side. All other twelve combinations were not statistically significant. These results mean that for most local masking methods analyzed, test participants do not perceive any differences whether a base map includes political-administrative boundaries, or the street network, or no base map at all.

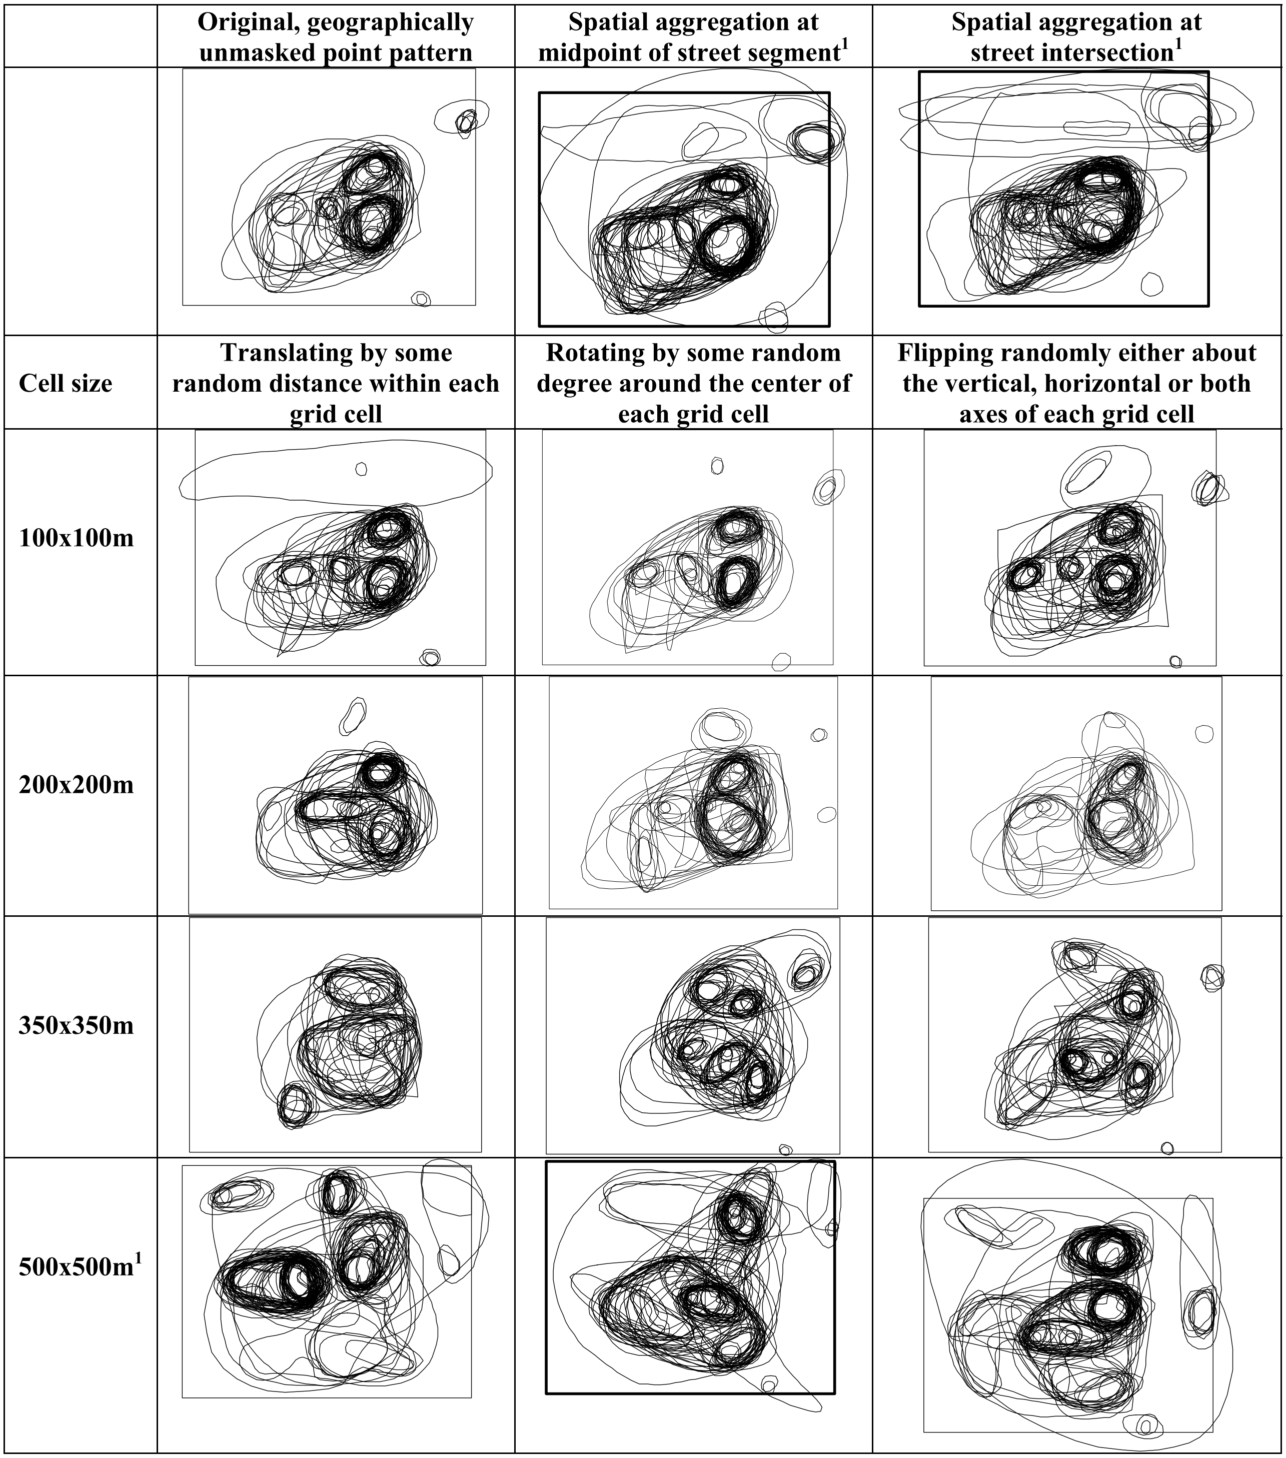

Figure 4. Comparing hot spots between the original (upper left) and different geographically masked point patterns for different cell sizes. All displays were derived from incident locations placed on maps with no background information. [1]Results were taken from Leitner & Curtis (2004).

6. Analyzing hot spots for different masking methods and different base maps

Areas within each point pattern that show a high concentration of incident locations are usually of great interest and worthwhile to be studied further. Any good and appropriate masking method needs to preserve the number, locations, sizes and shapes of these 'hot spot' areas. For this reason, participants were asked to identify high concentrations of incident locations in all geographically masked maps and in the original, unmasked maps. Results are displayed in Figure 4, which only includes hot spots that were drawn into point patterns with no background maps. Results for hot spots for the other two background maps (census tract boundaries and street network) are not shown here, because they did not differ much from the displays in Figure 4. In other words, test participants chose similar numbers, locations, sizes and shapes of these 'hot spot' areas irrespective of the background information.

The results in Figure 4 clearly show that the size, shape, location and number of hot spots change with increasing cell sizes but are not influenced by different masking methods. Specifically, there seems to be much agreement in the size, shape, location and number of hot spots between the original, unmasked point pattern and all local masking methods with small cell sizes of 100m and 200m per side. Leitner & Curtis (2004) found the same agreement between the original point pattern and point patterns that were spatially aggregated at the midpoint of their street segment or at their street intersection (Figure 4). This is an important result and it makes these local masking methods very useful and appropriate tools to visualize the location of confidential data. A further observation is that of a hierarchical arrangement of differently sized hot spots. There are two small (local) hot spots to the right of the center of each map. These two are combined into one medium-sized hot spot. Then there is a larger (regional) hot spot apparent across the lower portion of each map that encloses the two small and the one medium-sized hot spot.

However, differences in the size, shape, location and number of hot spots are visible when comparing the original point pattern with locally masked point patterns for the two largest cell sizes (350m and 500m per side). In all these locally masked point patterns, hot spots suddenly pop-up in areas, where they do not exist in the original map. This points to a particular danger of using an inappropriate masking method (cell size), namely the likelihood of an incident cluster to appear in a neighborhood of a city, where such a cluster does not exist in reality. If this information is distributed to the public, consequences may include false perceptions regarding the nature of the distribution and/or unfair informal redlining methods employed by some insurance and banking companies.

7. Preserving the privacy of confidential incidence locations

This section investigates to what extent one might be able to identify an individual incident location after this location has been geographically masked. Recall that each incident location identifies the street address where a homicide victim resided. Consequently, this section discusses to what extent the privacy of the confidential location of an individual residence would be preserved if that location has been geographically masked.

To answer this question, Leitner & Curtis (2004) randomly selected a sample of fourteen different street addresses from a total of 48 addresses used in their research (these are the same addresses that are also used in this study). All addresses were located in a residential neighborhood in Baton Rouge with mostly single-family houses. Churches, schools, office buildings, etc. were interspersed between the residences. Each of the fourteen residences associated with the street address from the sample was visited, and the number of residences on either side of the street segment that included the residence of interest was counted. Accordingly, all residences that were located on either side of all street segments of the street intersection to which the residence of interest was closest to were also counted. The results showed that the number of residences along either side of a street segment ranges from a minimum of two to a maximum count of 29 residences with a mean of 16.3. If the number of residences along the street segments with a common intersection is counted then the minimum number is 15, the maximum is 58 and the mean is 32.6. Both minimum and maximum numbers included the street address where a homicide victim had resided.

The question that needs to be addressed now is: What is the minimum number of all residences, including the residence of interest, so that the privacy of an individual residing at this address is not compromised? For this purpose, Leitner & Curtis (2004) applied "suppression", which is one of the confidentiality rules used by the U.S. Census Bureau of how to protect the privacy of individuals (see http://factfinder.census.gov). According to this suppression rule, the U.S. Census Bureau feels the privacy of an individual sufficiently protected, if the individual belongs to a group that includes four to six individuals.

If the suppression rule would be applied to preserve the privacy of confidential incident locations, then one confidential residence must be among a group of at least seven residences (the base population). According to Leitner & Curtis (2004) four of the fourteen residences for the local masking method 'aggregating incident locations at the midpoint of its street segment' did not have a sufficient number (at least seven) of other residences along the street segment and therefore would have to be suppressed. In other words, such geographically masked locations should not be displayed in a map, because the privacy of individuals living at these residences would not be guaranteed. In cases like this, one solution would be to move each of these four residence locations to their closest street intersection or randomly displace them within the cells of a regular grid. The total number of residences falling within each grid cell (the base population) can be found through direct field observation or by counting the number of residences in high resolution remotely sensed images.

8. Summary and discussion

The results of this research suggest that confidential incident locations should be geographically masked within cells of a regular grid that are 200 meters per cell side or shorter. This study found larger cell sizes to be inappropriate for geographical masking because larger cell sizes result in masked point patterns that are perceived to be different from the original, unmasked point pattern. Larger cell sizes are also inappropriate, because they "generate" high concentrations of incident locations in neighborhoods of the study area, were they do not exist in the actual point distribution. Whereas the cell size is crucial for appropriately masking confidential incident locations, the type of masking method and the type of background map is not. This study has shown that for cell sizes of 200m per side or shorter, locally masked point patterns do not differ significantly from each other whether they are masked by random translation, random rotation, or random flipping or whether the background map includes census tract boundaries, a street network, or is empty.

This study together with earlier studies by Leitner & Curtis (2003, 2004) have identified a set of local masking methods that could be used by law enforcement and other agencies interested to disseminate confidential point data to the public in the form of maps via the Internet. Which local masking method to preserve an individual's privacy is ultimately selected depends on the level of confidentiality that needs to be preserved and on privacy rules set forth by the agency itself. Agencies such as the U.S. Department of Health and Human Services, the U.S. Census and the U.S. Department of Justice have developed privacy rules for statistical or attribute information, individual information that is usually recorded as text, in a table, or in a spreadsheet format. Unfortunately, no such standards exist for the protection of spatial or locational confidentiality. This type of confidentiality is mostly associated with the visualization of individual's statistical information in maps and hence the importance and appropriateness of this research.

The least amount of confidentiality is preserved, when individual incident locations are aggregated at the midpoint of their street segment followed by aggregating incident locations at their closest street intersection (Leitner & Curtis 2004). It is expected that masking methods based on larger cell sizes will preserve one's individual privacy more than methods based on smaller cell sizes. Current research is underway to identify appropriate data sources in the form of high-resolution imagery that will allow counting the total number of residences (the base population) within individual grid cells. This will provide information on how much confidentiality is preserved with cell-based geographic masking methods. Among the future applications for this research is the development of a mapping and/or GIS module that will enable confidential data to be appropriately masked before it is released to the public.

Acknowledgements

The author would like to thank the Baton Rouge Police Department for providing the crime data used in this research. The author would also like to thank all students from one computer cartography, two introductory geographic information system (GIS) classes and one GIS and spatial analysis seminar, who participated in this experiment. These classes were held in the Department of Geography and Anthropology at Louisiana State University in Baton Rouge, LA during the Fall Semester 2004.

References

Aerts, J. C. J. H. Clarke, K. C. and A. D. Keuper. 2003. Testing popular visualization techniques for representing model uncertainty. *Cartography and Geographic Information Science* 30(3): 249-261.

Armstrong, M.P. Rushton, G. and D. L. Zimmerman. 1999. Geographically masking health data to preserve confidentiality. *Statistics in Medicine* 18(5): 497-525.

Cox, L.H. 1996. Protecting confidentiality in small population health and environmental statistics. *Statistics in Medicine* 15: 1895-1905.

Cox, L.H. 1994. Matrix masking methods for disclosure limitation in microdata. *Survey Methodology* 20(2): 165-169.

Duncan, G.T. and R. W. Pearson. 1991. Enhancing access to microdata while protecting confidentiality: Prospects for the future. *Statistical Science* 6(3): 219-239.

Leitner, M. 1997. *The Impact of Attribute Certainty Displays on Spatial Decision Support*, unpublished dissertation, Department of Geography, State University of New York at Buffalo.

Leitner M. and B. P. Buttenfield. 2000. Guidelines for the display of attribute certainty. *Cartography and Geographic Information Science* 27(1): 3-14.

Leitner, M. and A. Curtis. 2003. Cartographic guidelines for the visualization of confidential data using geographic masking. *21st International Cartographic Conference*, Durban, South Africa, 101-110.

Leitner, M. and A. Curtis. 2004. Cartographic guidelines for geographically masking the locations of confidential point data. *Cartographic Perspectives* 49: 22-39.

Monmonier, M. 2002. *Spying with Maps: Surveillance Technologies and the Future of Privacy*. Chicago, IL: University of Chicago Press.

Wartell, J. and J. T. McEwen. 2001. *Privacy in the Information Age: A Guide for Sharing Crime Maps and Spatial Data*, U.S. Department of Justice, Office of Justice Programs.