A Quantitative Similarity Measure for Maps

Richard Frank, Martin Ester

School of Computing Science Simon Fraser University Burnaby B.C., Canada V5A 1S6 rfrank@cs.sfu.ca, ester@cs.sfu.ca

ABSTRACT

In on-demand map generation, a base-map is modified to meet the user requirements on scale, resolution, and other parameters. Since there are many ways of satisfying the requirement, we need a method of measuring the quality of the alternative maps. In this paper, we introduce a uniform framework for measuring the quality of generalized maps. The proposed Map Quality measure takes into account changes in all local objects (Shape Similarity), their neighbourhoods (Location Similarity) and lastly across the entire map (Semantic Content Similarity). These three quality aspects measure the major generalization operators of simplification, relocation and selection, exaggeration and aggregation, collapse and typification. The three different aspects are combined using user-specified weights. Thus, the proposed framework supports the automatic choice of the best alternative map according to preferences of the user or application.

1. INTRODUCTION

With on-demand maps attracting increasingly more attention from large organizations such as Google (maps.google.com), MapQuest (www.mapquest.com) and Microsoft (mappoint.msn.com), the quality of maps generated by these companies becomes a competitive factor. Some of these systems allow for user customization through selection of different object-types while others allow for the visualization of buildings in 3D (Google Earth). Within a single system these maps can be generated in many different ways and from many different resources, each of which

must satisfy a different set of requirements from the client. Previously, paper maps were evaluated during the actual design and creation of the map, by the human geographer, to match the requirements as best as possible. Now, with large quantities of on-demand maps being generated almost instantly without human involvement, the quality of the final output must also be determined automatically by a computer. [16]

In on-demand map generation, a given base-map is modified to the user requirements on scale, resolution, and other parameters such as themes or symbolic representation [11]. Due also to constraints on the output medium, compromises have to be made during this process: shapes must be modified, combined or simply not displayed [12]. Since there are many ways of doing generalization and satisfying the user requirement, we need a method of determining the best alternative map generated so that user preferences can be met. Current map creation processes do not include a way of measuring the quality of the result, nor do they allow customization in the way of preferences [3, 5, 13]. Currently no consolidated measure exists, one that incorporates the shapes of the individual objects, the relationship between them, and their distribution on a map. The methods that have been suggested in the literature [5, 13, 14 and 18] measure some or all of these aspects, but they lack a uniform framework and cannot combine the individual aspects into a single metric. Existing methods display the results of the individual aspects, sometimes as many as 7 [13], to the user for interpretation. These methods are not applicable in the increasingly important scenario where an algorithm automatically needs to select the bestgeneralized map. The goal of this paper is to find the best map from alternatives by establishing a uniform framework which allows the calculation of a single metric quantifying the quality of a map generated regardless of source.

In [13] the map quality aspects have been categorized according to the three scales available on maps: micro (individual), meso (local) and macro (global). While the way map quality measure was split seems to be logical, most of the calculations they employ use very different methods for measuring the changes. [13], [14] and [18] all propose measures, none of which are based on a uniform conceptual framework.

Our approach also categorizes the measure into three aspects which are similar to those in [13]. When analyzing the shapes in isolation, we propose to measure the change by measuring the differences in the outline of the object. Then, instead of measuring the absolute changes in location [13, 14], we propose a method to measure the change of location relative to its immediate neighbours. Neighbourhood relationship is determined through the use of adjacent Voronoi regions. Finally, we've modified the standard entropy metric so that it's more suitable for information displayed on a map when compared to the entropy metric developed in [18]. Taking into account the preferences of the user, these three aspects will then combine into one metric representing the quality of a map.

Different applications will assign different importance to the above three aspects and hence it's not trivial to combine them. Some applications might require that the shape of objects be most accurate, neglecting their exact location. Due to this variance, our approach allows the users to input their preference to which aspect should be given priority during generalization. With the help of preferences, it's now possible to tailor the mapgeneration process, by generating several alternative maps meeting the user requirements of scale, area and resolution. After comparing them to the base map, we select the map with the highest quality out of the alternatives generated; hence the user is presented the best map that was tailored to their desires rather than a generic map that only meets the user specifications of scale and resolution.

The rest of the paper is organized as follows: Section 2 surveys related work while Section 3 details the specific approach that we propose along with detailed discussion of each aspect that it measures. Section 4 experimentally evaluates our approach and highlights the benefits. Section 5 concludes the paper with a summary and ideas for future work.

2. RELATED WORK

[13] proposes an approach for the quality assessment based on comparing data characterization, both before the generalization process and after it, at different levels of analysis. It describes 'micro', 'meso' and 'macro' level analysis, which is similar in breakdown to the 'shape', 'location' and 'semantic' similarity measure of our research, respectively. Their 'micro' level contains individual objects of all classes. These objects are described by their properties, such as area, orientation, position, concavity and elongation. Our approach treats location as relative (to neighbouring objects) and hence measures it on the level corresponding to the meso level while [13] treats it as an absolute measure without considering neighbouring objects. This micro level yields 5 individual quality measurements which individually are easy to interpret but fairly difficult when trying to view together. The meso level considers groups of objects, a city block for example, and they propose to measure such properties as density and proximity relations within these groups. Although no specific proximity or density function is given, their definition of a neighbourhood is not as robust as defining neighbours to be the adjacent Voronoi cells of an object. Their grouping of objects limits the usefulness of, for example, a density function, since density is now not calculated across all objects or classes but only in a group, whatever a group happens to be. Hence their quality function partially depends on the quality of, or meaning behind, the grouping. Their macro properties are based on all objects and the distribution of the properties of the individual objects, with the possibility of a restriction to only a single class of object. Unlike our proposed Semantic Content measure, they have no 'global' measure that takes into account, for example, the spatial distribution of the objects (not their properties) across the map. At the end, their algorithm provides them with 5 measures from the micro level and 1 from the meso level. The authors themselves state that the measures need to be aggregated somehow since a human cannot understand the relevance of so many numeric measures at a single time.

[18] proposes a set of new quantitative measures: 'Entropy of Symbol Types' and 'Entropy of Neighbourhood'. Entropy of symbol types does not take into account the spatial distribution and hence the authors state that this measure is not useful for maps. The inherent weakness of the entropy of neighbourhood is that the entropy can be the same for two maps with radically different distribution of object classes, but having the same neighbourhood relationships. The authors introduce 'Region of Influence' by using Voronoi diagrams to create 'Entropy of Voronoi Regions'. Instead of using the probability of an object of class i existing, they use the area of the Voronoi regions. This measure does not make a distinction between a class that is made of one object with Voronoi area of 10, and n objects with a total Voronoi area of 10.

[14] discusses how agents can co-ordinate and co-operate during the generalization process. Their quality measures are simple statistics regarding the objects (like size, minimum width, orientation, position, angle deviation and separation from other objects). This paper presents only the framework; it however does state the need for a good map quality measure.

[8] presents a method comparing and matching objects by proposing a shape description function which is based on the curvature and not the structure of the object. The paper discusses the 'turning function', a function which is a plot of the length vs. slope of the curvature of the original object. The turning function allows the comparison of two shapes by numerically quantifying the difference between two objects by finding the area between the two turning-functions. Typically, the turning-function is normalized to have the same length, that is, the original function must have the same perimeter; but this is not acceptable in most cases because small changes to the object can lead to dramatic changes in the perimeter. To deal with this scenario, the authors propose a matching process whereby pairs of identical sides are matched up and removed from the normaliza-

tion procedure which isolates the sides that have been modified. These sides are then compared by finding the difference between the two turning-functions.

[5] presents a method to generalize by attempting to calculate all possible maps that can be generalized. It tries to move each object into k possible places, for n objects, there are k^n possible maps. The evaluation function they adopt is based on minimizing the total number of conflicts (pairwise overlaps) within a particular region. As stated by the authors, they "require the development of more advanced evaluation functions that take account of a wider range of constraints, including those of form and structure".

3. OUR APPROACH

During a typical generalization process multiple maps can be considered acceptable even when generated from the same source, but a single map has to be selected to be returned to the user. In this section, we propose a framework for comparing two maps generated from the same source, resulting in one metric. The final results of each comparison can be evaluated with respect to the user preferences and the optimal map then given to the user in response to their original request. This process is illustrated in Fig. 1. Since we are given multiple maps generated from the same source, we assume that correspondence between objects from different maps is established by using the underlying Object-ID's retrieved from the database when the map is created. If the method is applied to maps generated from different sources than extra metadata is required to establish correspondence between objects.

Since all computer-based output medium is made of pixels, be it a monitor or printer, the map that is presented to the user must be in raster format, regardless of whether the source data was bitmap-data or vector-data. The generalization procedures vary depending on whether the data is sourced from bitmaps or vector data, but the result must be bitmap based. This paper presents a method of measuring not the quality of the individual generalization operators but the quality of the map that is presented to the user. When a map is put through a generalization process, there are several operations that can be used to solve a conflict between objects on the final map. This conflict occurs because it is not possible to display enough detail on a small-scale map in order to visualize all of the original objects. In a not-so-hypothetical example, two objects of length 15 pixels each are

6 Richard Frank, Martin Ester



Fig. 1. Comparison of two maps

separated by a single pixel, but after a scale-change of a factor of 10, those 2 objects now occupy a total area of 3 pixels (1.5 pixels each). Since the space between them cannot be displayed because, by definition, the output medium cannot display 0.1 pixels, the 2 objects can overlap. The generalization method has several choices [1, 5, 12, 14, 16, 17], it can:

- not display one of the objects (selection/elimination)
- enlarge the objects so they can still be displayed properly (exaggeration)
- combine multiple objects (aggregation)
- move the objects a bit further apart so that they do not overlap (relocation/displacement)
- remove some features from the original object (simplification/reduction)
- reduction in the dimensionality of the object, from spatial to a line, for example (collapse)
- represent using typical distribution pattern (typification)

Which operation deals with the conflict depends on the algorithm. There are multiple algorithms that have been proposed, such as 'least squares with constraints' or 'simulated annealing' [1], but regardless of the algorithm, the above types of change can be classified into 3 different aspects.

The first aspect is shape similarity and is derived from a change where the shape or size of the object is changed. Shape Similarity would measure the changes from exaggeration, collapse and the simplification operations defined above. This measure takes place on the object level because only the single object is changed.

If the object is moved or the relative distances change due to a change of shape of one of the objects, then the location of it, and hence relatively its neighbours', is changed, but this change is restricted to the local neighbourhood. A change on this level impacts the location similarity and measures location changes due to collapse, displacement, exaggeration and simplification.

The last type of measure, Semantic Content Similarity, measures changes due to aggregation, omission and typification. These changes all impact the entire map and hence have a global effect. Also, each object belongs to a class and the removal of an object will influence the entire class distribution across the map hence impacting the importance of the objects.

In order to combine the different aspects of map quality, it is desirable that the underlying conceptual framework for all aspects is uniform. When working with maps and other spatial objects, Voronoi diagrams and Delaunay triangulations are prevalent since they clearly represent the spatial relationships between the objects [4, 7, 9]. These Voronoi diagrams split an area of space into regions, called Voronoi cells, which contain all the points that are closest to the object contained in the Voronoi cell. The size of these cells gives an indication of how dense an area a certain object is in or how large an object is and hence is a good aid when it is required to calculate the amount of information contained in a certain area. The Voronoi cell structure also yields the Delaunay triangulation which easily allows calculating an object's immediate neighbourhood, a prerequisite to our calculation on the local level.

There has been research into calculating the similarity of two shapes, one of these techniques is outlined in [2, 10 and 15] which uses the skeleton of objects, also called shock graphs. The skeleton is very similar to a Voronoi diagram (but within an object) and [2, 10] discusses a procedure of calculating the difference between two such skeletons. For certain objects, such as those with indentations, the shock graph is very different although the original shapes can be considered the acceptable result of shape simplification, as in Fig. 2.



3.1. Shape Similarity

In order to compare the individual objects, some kind of descriptor is needed which describes the amount of change between two shapes. In the methods used previously, shapes are described by certain properties which yield a number, and either the absolute or relative change between the two values results in the Shape Similarity measure. [13] describes properties such as *area, perimeter* or *length/width* and applies them to calculate Shape Similarity, but none of these are unique to a single shape and hence are not the best descriptors. The *Turning Function* is a step function which describes a shape by its perimeter vs. slope; the *x*-coordinate denotes the distance along the perimeter and the *y*-coordinate denotes the value of the slope [8]. The difference in *y*-coordinates of adjacent line segments represents the turn-angle of the corresponding line-segments in the original object.

On maps, which are bitmaps when they are presented to the user, all the objects on the map are pixel based and hence have outlines composed of 45° or 90° angles (at the pixel level) when using the 8 neighbourhood definition, hence the turning-function will be a step-wise function. This function uniquely identifies the shape and is rotation independent since rotation of the original shape is equivalent to a translation in the turning-function [8]. However, in order to be able to compare two shapes, their turning-function must have the same length, i.e. we must normalize with respect to the perimeter. A small disturbance in the perimeter of the object could cause the normalized turning-function to not align properly during matching because sides that should be the same will not be due to one object now having a larger perimeter than the other even if the two shapes are very similar otherwise.

To isolate disturbances, a matching procedure must be performed [8]. The goal is for the shapes to be compared and the largest identical segment (including the turns in between them) removed from the normalization and distance calculations. Detecting the identical segments is simple if only the



absolute lengths are considered but can be problematic if the scale of the objects are varied since then the right scale-change has to be found before matching can take place but finding the right scale-change relies on finding an optimal matching first. Our matching algorithm compares the lengths and turns of all the edges in both shapes and removes the largest identical subset.

After matching, the non-overlapping area, $Area(O_i)$, between the two turning-functions is calculated, where O_i represents object *i* on both maps. For an example see solid areas in the graph of Fig. 3. This yields the absolute difference between the two shapes with a numerical value that can theoretically be fairly large. In order to restrict the result to be between 0 and 1, the relative *Shape Similarity* (SS) for object *i* on Map *A* and Map *B* is defined as:

$$SS_i(O_i) = 1 - \frac{Area(O_i)}{Max[Area(TF_{i,A}), Area(TF_{i,B})]}$$

• where *Area*(*TF*_{*i*,*A*}) is the area under the turning-function for object *i* on map *A*.

If the two shapes being compared are identical, then their turningfunctions will be identical, resulting in an SS measure of 1. Globally, each shape on Map A is compared to its corresponding shape on Map B and each shape will yield one Shape Similarity number. The SS across an entire Map A, when compared to Map B, is defined as the average of all the shape difference numbers:

$$SS(Map_A, Map_B) = \sum_i SS_i(O_i) / n$$

3.2. Location Similarity

Objects can be displaced during generalization and this displacement becomes greater as the scale decreases [6]. This change can be reflected and measured in the change in distance relative to other objects. A local neighbourhood for an object is composed of the object itself and its immediate neighbours, where the neighbours are defined based on adjacency of Voronoi cells. Calculating the pair-wise distances for all objects is computationally expensive and is not necessary since changes in the Voronoi structure do not propagate past the immediate neighbouring cells, hence, when calculating the Location Similarity, all calculations can be restricted to the immediate neighbourhood, as illustrated in Fig. 4.

The distance between two spatially extended objects is typically based on the distance between representative points on these objects. Defining the representative points to be the *center-of-gravities* can lead to the points to be inside other objects. Alternatively, the representative points can be defined to be the points *furthest* from the other objects, i.e., the point on object A is the point which is furthest from object B and vice versa. In instances of very large objects, even when the objects are in contact with each other, the distance can be very large and so this type of measurement does not yield appropriate results. The representative points of A can also



Before Generalization

After Generalization

Fig. 4. Relocation

be ones *closest* to *B*; this measure seems intuitive because this is the way humans also define distance, from the closest edge between two locations. Hence this type of measurement is used.

The pair-wise distances between two objects can be stored in a matrix $dist_A(i,m)$, where *i* and *m* are two objects on Map *A*. The *Location Similar-ity* metric for object *i* is then defined to be:

$$LS_i(A, B) = 1 - \frac{\left[\sum_{m=1}^{p} \left| \frac{dist_A(i,m) - dist_B(i,m)}{max(dist(i))} \right| \right]}{p}$$

• where *max(dist(i))* is the maximum distance between object *i* and all of its Voronoi neighbours and *p* is the number of neighbours.

 $LS_i(A, B)$ can be interpreted as the average change in distances between a central object *i* and the objects around the central object.

Since LS_i is a local neighbourhood measure, it generates a measure for every single object and those measures must be combined into a single global measure describing the Location Similarity for the map. Thus the global *Location Similarity* for the map is defined by

$$LS(Map_A, Map_B) = \left(\sum_{i}^{n} LS_i\right)/n$$

• where *LS_i* is the Location Similarity metric for object *i* and *n* is the number of objects on the map.

This definition is assuming that all the objects are of equal importance, and if this assumption does not hold then weights can be introduced for each class or object.

Since the location-similarity metric is normalized with respect to the largest distance between the object and its neighbouring objects, the largest possible value that can occur is 1. If none of the objects move then both dist(i,m) will have the same value and $LS_i(A,B)$ will equal 1. Since $LS_i(A,B)$ is restricted to a value between 0 and 1, the average of all $LS_i(A,B)$ -s (the Location Similarity index for the entire map) will also be bounded to be between 0 and 1.

Due to operations done during generalization, such as a deletion or move, it is possible that the set of neighbours before and after generalization will not be the same, as illustrated by the central object in Fig. 5, in which case when we calculate the $LS_i(A, B)$ for the shaded object,

 $dist_A(i, m) - dist_B(i, m)$ will not be defined between some objects because some objects will not exist after generalization and hence their pair-wise



distance cannot be determined. Therefore the neighbourhood definition must be restricted to objects that existed before and after generalization. Hence, the Location Similarity of objects (such as the shaded object) is only defined on the common set of neighbors between the two objects.

Another possible outcome of generalization where the neighbourhood set is not the same before and after generalization is aggregation. In this case the aggregate object is not the same as the individual objects (before generalization) and hence would not be placed into the intersection of the two sets. A simple example is shown with the shaded objects in Fig. 6 where the central four objects are aggregated. These differences are not measured by the Location Similarity metric, but by the Semantic Content Similarity.



Fig. 6. Aggregation (different shading represent different classes)

3.3. Semantic Content Similarity

During the generalization process several objects can be aggregated into one, or be deleted. These operations imply a loss of information and there have been statistics developed in [18] for measuring the amount of information in a set of objects. The existing entropy calculation measures the information within a set of non-spatial objects and is based solely on the number of objects within each class on the map. On a map the objects are spatially distributed, two maps with identical number of objects from the same classes could have drastically different spatial distributions as discussed in [18]. The method we propose takes into account spatial information by weighing each entry in the Entropy sum by their respective Voronoi areas, as a percent of the total map area.

Before attempting to define a modified entropy measure, it is important to define what the entropy measure is based on, i.e.: what an object is. Objects that intersect or behave as a network that span the entire map (ex: road network) should be treated as segments of individual objects such as individual roads or river-segments. For example: an object with a general shape of '+' can be split into 4 non-intersecting segments. Similarly, when objects 'branch' (eg: a river splits), the object can be broken into the component line segments: a river segment the shape 'Y' will have 3 components. This allows any shape-similarity measure to work since there are no objects that are not closed or contain holes. This also is required in order to measure the disappearance or displacement of road segments of a roadnetwork since a disappearance of a segment would be detected by the Semantic Information measure while a displacement would be captured in the Shape Similarity measure.

We need to determine the change of semantic content between two maps, the original and the generalized. As discussed in [18], the size of the Voronoi cell is a good indicator of how large the object is and also indicates to a certain degree the distribution on the map. It would be beneficial to modify the existing Entropy measure by incorporating the relevance of objects according to their Voronoi regions. True, information is lost when something disappears, but the objects are more influential. We define the *Voronoi Entropy* of Map *A* using:

$$VE(Map_A) = \sum \left[P_i \times \ln(P_i) \times \% V \right]$$

where

- V_i is the total area of the Voronoi regions of objects of class i,
- *K_i* is number of objects of class *i*,

- *N* is total objects on the map,
- and $P_i = K_i/N$.

A class is a collection of objects that have the same semantic attributes but can vary in position, size and other attributes, for example 'banks', 'roads'. As an example, using Fig. 6, the information in Table 1 can be collected.

		Map 1		Map 2	
Object	Class	# objects	% Voronoi area	# objects	% Voronoi area
Clear	Stores	4	50%	4	50%
Dark	Residence	4	12.5%	1	12.5%
Pattern	Parking-lot	3	37.5%	3	37.5%
Table 1. Semantic Content Similarity					

Hence, calculating the VE for both maps, $VE(Map_A)$ and $VE(Map_B)$:

$VE(Map_A)$	$= [4 \times \log(4) \times .5] + [4 \times \log(4) \times .125] + [3 \times \log(3) \times .375]$ = 2.0419
$VE(Map_B)$	$= [4 \times \log(4) \times .5] + [1 \times \log(1) \times .125] + [3 \times \log(3) \times .375]$ = 1.7408

The largest value that the entropy measure of a single map can take can be arbitrarily large, but the change between the two maps, when expressed relative to the map with the larger entropy, is between 0 and 1. The *Semantic Content Similarity* (SCS) similarity can hence be defined as the amount the two entropy measures have changed:

 $SCS(Map_{A}, Map_{B}) = \frac{MIN[VE(Map_{A}), VE(Map_{B})]}{MAX[VE(Map_{A}), VE(Map_{B})]}$

The relative similarity can be I when Maps A & B are identical in composition, although they don't have to be identical, just the number and class of objects along with their Voronoi Areas have to be the same. This type of change, however unlikely, will give a result of I, but these changes would be measured by the other map quality measures.

3.4. Combining the three quality aspects

Although the results from the three components are meaningful individually, in order to calculate one metric for each map, the three results have to be consolidated into one. Since all three numbers purposefully have the same range of values (0 to 1) and all behave in a similar fashion (i.e.: small change is indicated by a small value) it is possible to combine them. The simple average can be a good indicator of all three values; it however gives equal weight to all components when that assumption might not be appropriate for the application or user.

Since the importance of the different aspects of map quality depends on the user or application, the weights for each can be user defined.

> $Quality(Map_A, Map_B) = w_1 * SS(Map_A, Map_B) +$ $+ w_2 * LC(Map_A, Map_B) + w_3 * SCS(Map_A, Map_B)$

Where w_i is the weight for the metric and $\sum w_i = 1$

The resulting quality value would be associated with the map that it was calculated from. A lower quality value would indicate that the map had to undergo a larger amount of change and hence the quality is worse than a map with a higher measure. By comparing the result of the calculation and choosing the generalized map with the highest similarity value, it is possible to select the best quality maps.

In some instances the generalization operators can be considered to be dependent, for example the change of shape of one object could cause the distance between two objects to change. In these instances, two (or all three) measures would simultaneously change to reflect the change in the object itself. This not only is intuitive but desired.

4. EXPERIMENTAL RESULTS

To illustrate the meaningfulness of the proposed map similarity measure, we present experimental results on some maps generated using generalization methods developed in the GEMURE project [16]. One of the goals of the GEMURE project was to improve on-demand cartographic information delivery through generalization and multiple-representation, and they required a method to evaluate the results, which our method was designed to do.

The first map being a large-scale map (Fig. 7) that act as the 'base-map' while two others (Fig. 8 and Fig. 9) are alternative generalizations of the



same area. These alternative maps are then compared to the 'base-map' in order to find out their accuracy. The maps depict a small section of Quebec City with three different classes of objects: roads, residential buildings and commercial buildings.

The first generalized map, *Map1* (Fig. 8) includes generalization effects of shape simplification along with merges of nearby smaller objects. The second map, *Map2* (Fig. 9), is at a smaller-scale than *Map1* and also uses selection with a large number of objects disappearing or merging with neighbouring objects. The results of the comparison are shown in Table 2. In the following, we discuss the different results for the three aspects of Map Quality.

The SS values are also somewhat intuitive: in *Map1* all the objects that are both in the *Base-Map* and *Map1* are fairly similar (note that this does not mean that all the objects are the same, some small objects were aggregated into bigger objects, but these are not counted in the Shape calculation) whereas the outline of a lot of the objects were drastically changed going from the *Base-Map* to *Map2*. Since the objects on *Map1* are relatively close to their original counterparts, while *Map2* underwent a lot of change, hence the SS quality of *Map1* is measured to be higher than *Map2*.

The location-similarity is relatively high for all the maps. *Map1* is able to preserve the location well since it is not generalized very much, *Map2*

	Comparing		
*weight indicated in parenthesis	Base-Map to Map1	Base-Map to Map2	
Shape Similarity (15%)	0.996	0.285	
Location Similarity (70%)	0.821	0.828	
Semantic Similarity (15%)	0.437	0.270	
Final Similarity	0.790	0.663	

Table 2 - Results using a different set of weights

uses selection in abundance and hence is also able to preserve location relatively well.

Semantic Content Similarity was low for both maps because they underwent significant aggregation during generalization. *Map1* however had the best SCS since it contained the most objects while *Map2* contained fewer. This is also reflected in the SCS measures.

By defining a set of weights, the user is able to place more emphasis on different aspects of map quality. It is possible that a different map is selected depending on what the weights are. In this case, *Map1* is determined to be of better quality.

5. CONCLUSION

In this paper we presented an approach that can calculate the quality of alternative generalized maps using a uniform framework and present a single number that quantifies the quality. The approach takes into account changes in individual objects in the form of the Shape Similarity, groups of objects using the Location Similarity and changes across the entire map using Semantic Content Similarity. The framework also allows for the user to specify a set of preferences which then influences the final metric and hence the choice of the "best" map. Our experimental evaluation on real maps demonstrates that the proposed Map Quality measure produces intuitive results and, thus, supports the automatic map selection according to the preferences of the user.

Our method could also lead to a novel way of performing spatiotemporal data mining by allowing the calculation of the changes that areas of a map have gone through over time. By comparing two maps of the same area, but different time-periods, and applying a data-mining algorithm to sub-areas of the two maps, it would be possible to use the map as the search-space to search out the sub-area with the most change. Each sub-area from time-period T_1 can be compared to the same sub-area in time-period T_2 using the Map Quality calculations discussed above, and a quality metric calculated for each sub-area. The sub-area with the smallest Map Quality metric would indicate the region that has undergone the largest amount of change between T_1 and T_2 . Using this method, it would be possible to answer such queries as "Where did most of the development occur between T_1 and T_2 ?" with nothing more than the temporal database containing the map available.

6. REFERENCES

- Christopher B Jones, Alia I Abdelmoty, Michael E Lonergan, Peter van der Poorten and Sheng Zhou. *Multiscale Spatial Database Design For Online Generalisation*. 9th International Symposium on Spatial Data Handling. (10)
- [2] Christopher Gold, *Crust and Anti-Crust: a one-step boundary and skeleton extraction algorithm*. Annual Symposium on Computational Geometry, pp 189-196, 1999. (11)
- [3] Dijk, S.F. van, Kreveld, M.J. van, Strijk, T., & Wolff, A. Towards an evaluation of quality for names placement methods. International Journal of Geographical Information Science, 16(7) 2002, 641-661. (18)
- [4] Franz Aurenhammer. Voronoi diagrams--a survey of a fundamental geometric data structure. ACM Computing Surveys, 23(3):345-405, September 1991. (5)
- J Mark Ware, Christopher B Jones Conflict Reduction in map generalization using iterative improvement. GeoInformatica, Dec 1998, vol. 2, no.4, pp 383-407. (16)
- [6] João, Elsa M., Causes and consequences of map generalisation, Taylor & Francis, London, 1998. (2)
- [7] Kenneth E Hoff III, Tim Culver, John Keyser, Ming Lin, Dinesh Manocha. Fast computation of generalized Voronoi Diagrams using graphics hardware. SIGGRAPH 1999. (1)
- [8] Longin jan Latecki, Rolf Lakamper, Shape similarity measure based on correspondence of visual parts. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 22 No 10. October 2000. (14)
- [9] Niranjan Mayya & V.T. Rajan, Voronoi Diagrams of Polygons: A framework for shape representation. Journal of Mathematical Imaging and Vision, 1996, Volume 6, Issue 4, pp 355 – 378. (12)
- [10] Philip N. Klein, Thomas B Sebastian, Benjamin B. Kimia. *Shape matching using edit-distance: an implementation*. Symposium on Discrete Algorithms 2001. Twelfth annual ACM-SIAM symposium on Discrete algorithms, pp 781 790. (6)
- [11] Sonia Rivest, Yvan Bedard and Pierre Marchand. Toward better support for spatial decision making: defining the characteristics of spatial on-line analytical processing (SOLAP). GEOMATICA vol 55, 2001, pp 539 to 555. (8)
- [12] Stefano Spaccapietra, Christine Parent, Christelle Vangenot. GIS Databases: From multi-scale to multi-representation. In Proceedings of the International Workshop on Emerging Technologies for Geo-Based Applications, LNAI 1864, Springer, 2000. (3)
- [13] Sylvain Bard, Quality Assessment of Generalized Geographical Data. Accuracy 2002 Symposium, 2002. (15)

- [14] Sylvie Lamy, Anne Ruas, Yves Demazeu, Mike Jackson, William Mackaness, Robert Weibel. *The Application of Agents in Automated Map Generalisation*. 19th International Cartographic Conference, 1999. (13)
- [15] Thomas B. Sebastian, Philip N. Klein, Benjamin B. Kimia, *Recognition of Shapes by Editing Shock Graphs*. ICCV 2001, pp 755-762. (4)
- [16] Yvan Bedard and Eveline Bernier. Supporting multiple representations with spatial databases views management and the concept of VUEL. Joint Workshop on Multi-Scale Representations of Spatial Data, ISPRS WG IV/3, ICA Com. on Map Generalization, 2002. (7)
- [17] Z. Li, H. Yan, T. Ai, J. Chen, Automated building generalization based on urban morphology and Gestalt theory. Int. J. Geographical Information Science, Vol 18, no. 5, July-August 2004. pg 513-534 (17)
- [18] Zhilin Li, Peizhi Huang. Quantitative measures for spatial information of maps. International Journal of Geographical Information Science 2002. vol 16, no 7, pp 699-709. (9)