Hybrid Transformation for Indexing and Searching Web Documents in the Cartographic Paradigm

Fiona Lee, Stéphane Bressan, Beng Chin Ooi

School of Computing. National University of Singapore. 3 Science Drive 2, Singapore 117543. email:{flee,steph,ooi}@comp.nus.edu.sg

Abstract

As information becomes available on the World Wide Web in larger quantities and in more disparate formats and media, adequate search engines and portal services providing search and filtering modalities tailored to the needs of the various communities of users become essential components of the global information infrastructure. Global-Atlas is a geographical search engine leveraging the cartographic paradigm for the indexing and searching of World Wide Web documents. Documents are indexed according to their geographical footprint, i.e. the bounding box of the geographical region to which they are related. Documents are searched by interactively drawing bounding boxes on maps. The effectiveness of Global-Atlas depends on the capability to quickly and to accurately index documents and maps. Maps, however, come in a variety of unspecified coordinate systems and projections. In this paper, we evaluate various surface-fitting techniques for the calibration of maps and devise an original hybrid calibration method from the empirical results obtained.

Key words: Web-based Geographical Information System, Map Calibration, Map Registration

1 Introduction

The quantity and diversity of information readily available from the World Wide Web compels more adequate search engines and portal services providing search and filtering modalities tailored to the information needs of the various communities of users. In several application domains, one very natural paradigm for indexing, searching, or sharing information is the cartographic paradigm. Indeed, in domains such as transportation, environment protection, tourism, information is best presented and searched by means of maps. In such applications, the geographical communities (tourists visiting the same region and service providers for these regions, local community service providers and inhabitants, professionals engaged in collaborative engineering projects) are best assisted in their needs by systems whose interfaces, query languages, and data models support graphical map-based interactions.

Starting from the observation that a large corpus of World Wide Web document could be usefully indexed and searched according to a cartographic paradigm, and recognizing the weaknesses of existing services and products, we have design and developed Global-Atlas, a geographical search engine. Information available on the World Wide Web is indexed according to its geographical "footprint" in place or in addition to the traditional keywords or categories used by most search-engines such as Yahoo, AltaVista, Hotbot, or NorthernLight. In practice, Global-Atlas indexes the URL of the documents (not only maps in GIF or JPEG format but also other Web documents such as HTML documents, movies, or sound bytes) by the bounding box of the region to which the document is topically related. Queries to the search engine are expressed as regions drawn on a map. A region is expressed as a four-sided polygon, which are bounded in terms of longitude and latitude values. The motivation of using graphical query is that it is more intuitive and it also provides the users with more freedom and ease in their search. In the process of searching for the most adequate map, a desirable feature for the Global-Atlas is to support multiple levels of region map search with the wide range of maps on the WWW. Hence, allowing users to reuse the maps returned as a user interface for refinement of queries. This can be done if enough information has been collected about a map image or a satellite picture, thus potentially enabling Global-Atlas to reuse each geographical image available on the Internet as query interface.

Finally and in practice, other constraints can be added to the query such as keywords, the required precision of the map, size of the maps, etc. The indexing of the data and the efficient processing of the queries in Global-Atlas rely on spatial indexing data structures [1,2]. Global-Atlas is extending the capabilities of its predecessor the World Travel Editor [3]. One of the main problem encountered in the World Travel Editor was the difficulty for the administrator of the system, in charge of indexing new documents, to handle sufficiently large quantities of documents to make the system effectively useful. Therefore, much effort has been spent in the development of more convenient tools and techniques for the semi-automatic indexing of documents.

One of the most difficult issues in the process of registering documents in the Global-Atlas index is the calibration of maps, satellite pictures and aerial pictures. In order to reuse maps obtained from different sources over the Internet, the geographical coordinates (longitude, latitude) given in a map have to be recognized. As the number of mapping projections is extremely large, it is not easy to detect each projection. Hence, there is a need to find a good generic mapping transformation to calibrate different maps to a single reference coordinate system.

In this paper we present our ideas on sharing of global map documents, images and references on the web with the Global-Atlas Search Engine and the Administrator Toolkit. We will also present an evaluation on the results of a comparative empirical analysis of various surface fitting techniques for the calibration and reuse of maps and pictures for Global Atlas. The surface fitting techniques we compare have been used in remote sensing, however, as we will see, the indexing of Internet documents compels the fine tuning of the combination of these techniques. Based on the empirical results obtained using representative samples of maps (different regions and different projections), we have devised an original hybrid calibration method. Once maps are calibrated and registered, they may be used as query interfaces. Furthermore, in order to provide users with an accurate interactive query drawing and feedback, we devised an original double transformation method and perform a second study on the performance of the various techniques for a good reverse transformation technique.

2 Maps, Information and Digital Communities

Let us first look at the different types of information systems using the cartographic paradigm. A Geographical Information System (GIS) is a system that supports the management (storage and retrieval) of information according to its geographical properties and facilitates the decision and design processes based on this geographical information. For instance GIS are used to manage environmental projects, monitoring of the wild life in the state of Bavaria, or plan city developments, public housing construction in Singapore. In most cases, the geographical information is defined and represented with respect to a coordinate system and, most likely to a reference map of the terrain (rivers, coastlines, and elevation). As a consequence, GIS technology's focus on the management of geographical objects, i.e. 2D or 3D featured graphical objects. An example of a GIS system available for the personal computers users is ArcView or ArcInfo (URL: http://www.esri.com/).

Geographical Digital Libraries (geolibraries) are different from GIS as their main role is to serve as a catalog and possibly a reference database for a collection of documents. The documents can be digitized or electronic books and texts, photographs, or maps. Digital Libraries are "Geographical" when the content or the organization of the documents is geographical. The Alexandria [4] or the Georep [5] digital libraries are examples of geolibraries. It is not uncommon however that geolibraries extend the notion of document to many kinds of geographical objects bringing them closer to large multipurpose GIS.

This is the case for the two above-mentioned projects.

When the primary goal of geographical information systems is to deliver maps, they could be qualified as Digital Atlas (DA). It is nowadays very common to find such systems for car navigation (digitized road-Atlas). Microsoft Encarta is probably today's most popular general purpose DA. Encarta is a comprehensive world atlas and geographic reference guide. It delivers high quality detailed maps, up-to-date statistical data, and the rich cultural information. Users can customize maps by adding or removing place names, borders and topographic detail. Still from Microsoft but in the professional side, Microsoft TerraServer [6] is a DA of satellite images (URL: http://terraServer.microsoft.com/).

Our approach is closest to the one of the Alexandria Digital Library (ADL) project at the University of California at Santa Barbara. ADL is a geolibrary where a primary attribute of collection objects is their location on Earth, represented by geographic footprints. A footprint is the latitude and longitude values that represent a point, a bounding box, a linear feature, or a complete polygonal boundary. As in Global-Atlas, footprints can be attached to any type of information about a geographic location, including maps, aerial photographs, and remote sensing images as well as text items, museum objects, specimens, data set, music – any geo-referenced piece of information. In Global-Atlas, however, the documents themselves are not gathered in local repository but managed independently by their creator or distributor on the World Wide Web.

The ADL project has developed a content standard for gazetteer objects and a hierarchical type scheme for geographic features. Both of these developments are based on ADL experience with an earlier gazetteer component for the Library, based on two gazetteers maintained by the U.S. federal government. A gazetteer, also called a geographical name server, is a list of geographic names, together with their geographic locations and other descriptive information. A geographic name is a proper name for a geographic place and feature, such as Santa Barbara County, Mount Washington, St. Francis Hospital, and Southern California. This project covers a wide aspect of the geographical information field. The process of loading new gazetteer database to its system is very labor intensive, and there are also several issues, concerning fuzzy areas, control of gazetteer data, and obtaining bounding boxes and polygonal extents, which the ADL are still researching.

Other related and recent projects attempting to organize the information on the web in a geographic paradigm includes Hitachi's Cyberwalk [7–10]. Cyberwalk is a virtual reality environment in which an avatar wanders the virtual town of Kyoto. As the avatar takes a walk on the maps, geographically related web pages are pushed to the users. The web pages are being indexed by hand.

The Center of Spatial Information Science of the University of Tokyo has also recently designed a prototype system. The system is capable of indexing Japanese Web pages (such as restaurant and other stores homepages), by identifying addresses in the text and matching them with a gazetteer if Japan Street addresses to obtain coordinates [11]. The availability of portable and wireless device with the ability of geographical positioning tracking (GPS), such as mobile phones and personal digital assistant (PDAs) will compel a new generation of location dependent applications. Global Atlas is a first step in this direction.

3 Global Atlas

Global-Atlas is a Geographical Search-Engine. In Global Atlas, as in its predecessor the World Travel Editor [3], references to documents published on the Web (URLs) by third parties are stored in a database and associated to a geographical region. Examples of these documents are image maps, satellite pictures or even text documents about the geographical regions. Bounding boxes of the region are used to index documents in place or in addition to the traditional keywords or linguistic concepts used by most existing search engines (Altavista, HotBot, Northern Light, or Yahoo). Web users can search the database using a client interface (currently a Java applet) and "draw" queries.

As illustrated in Figure 1, besides the usual keyword search offered by a conventional search engine, the Global Atlas search engine provides users with a graphical map interface whereby a world map is loaded on it by default. The user can search different map regions graphically by drawing a query (rectangular bounding box) over the world map. The Global Atlas search-engine supports three different types of query search, namely inclusion, covered and intersection searches. An inclusion search is to look for regions that includes or contains the query's bounding box. While, the covered (by) search returns regions that are covered by or inside the query's bounding box. The intersection search returns regions, which intersects with the query's bounding box. The reader notices the geoprint of one of the returned document, a map of Europe on Figure 1.

The user can choose to browse the retrieved results as in a conventional search engine and view them in her favourite Web-browser. Alternatively, with the retrieved results, which are maps, the user has the option of loading the image map onto the map interface of the G-Atlas search engine for further refinement of search. For, in the process of registering and indexing the image of a map, it is often possible to gather sufficiently enough information to allow the map to be reused as a search interface. This feature is very powerful as it allows the

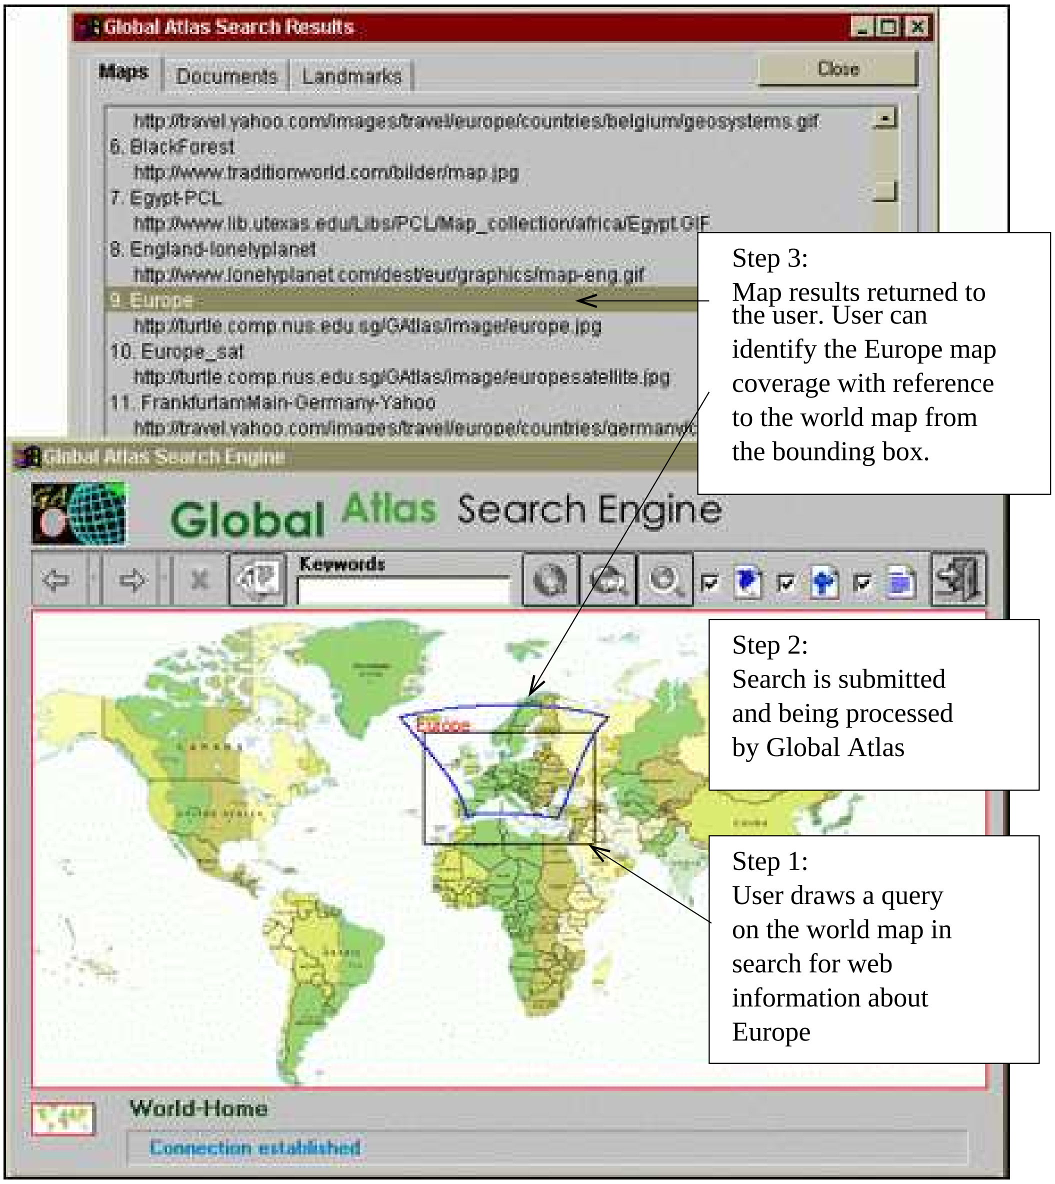

Fig. 1. Performing a query search over the Europe region

user of Global Atlas to choose the most adequate map return by the search engine to formulate and refine her queries. Aerial photos and satellite picture can also be reused once it is registered with enough information. Figure 2 illustrates the reuse of a retrieved map to further pose queries.

The most intuitive examples can probably be drawn from tourism industry. Let us imagine, for instance, a tourist or a business traveller planning a trip to Singapore. A geographical search engine could best help her to explore the World Wide Web resources available to her (not only maps and street plans but also pictures, guides, landmarks or public transportation on-line routes and schedules) by providing a very natural and intuitive cartographic interface. The user can express her queries by graphically interacting with maps

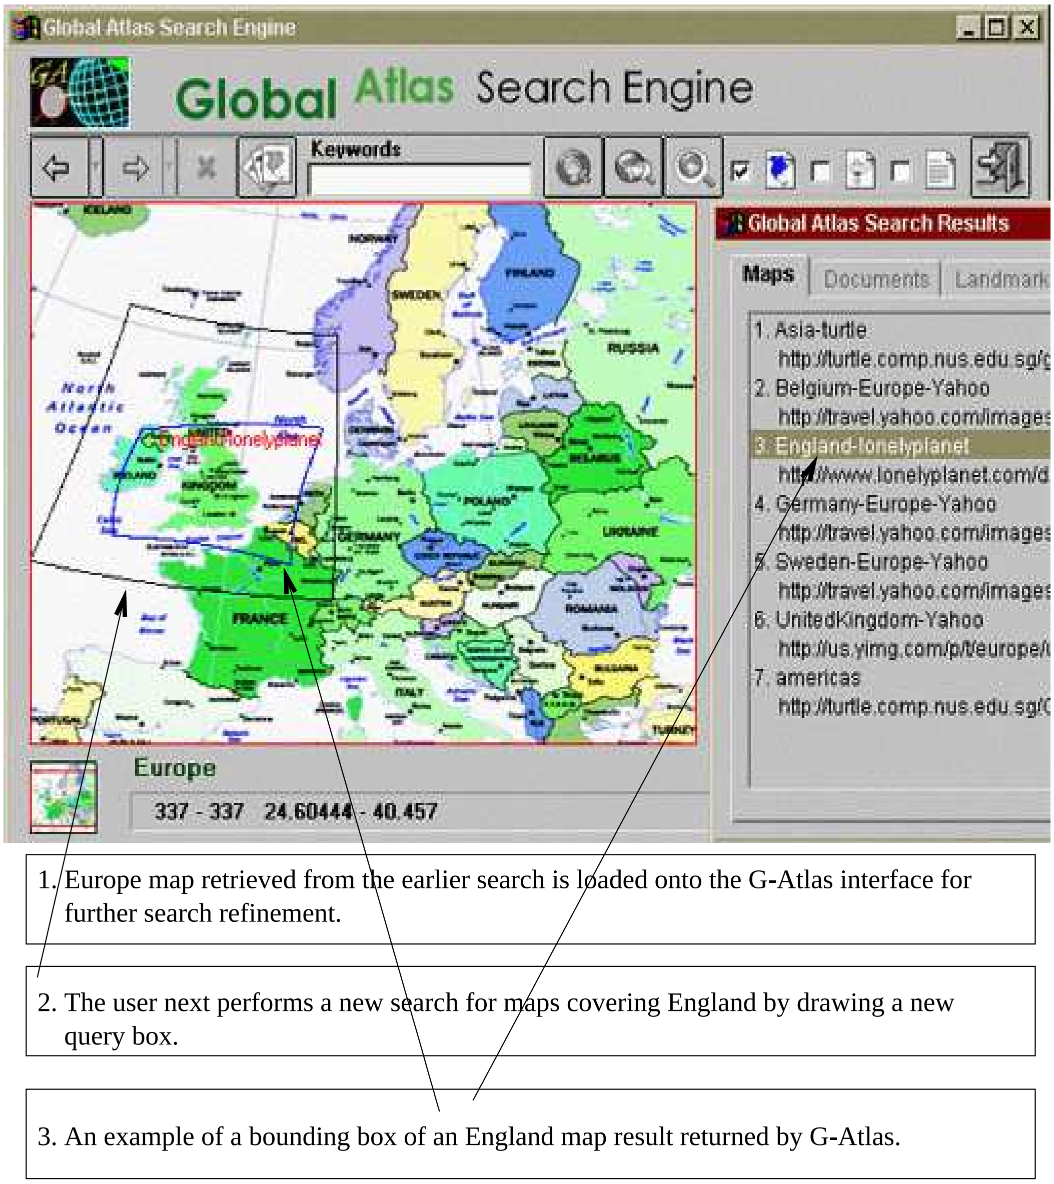

Fig. 2. Retrieved Europe map can be used as an interface for new searches

and selecting areas of interest. A search engine would then retrieve reference information indexed as being related to a location in the neighbourhood, overlapping or included in the selected areas. See figure 3 for an illustration of places of interest landmark provided by G-Atlas.

3.1 Global-Atlas Architecture

Global-Atlas consists of three main components: the Client Interface, the Administrator Tool, and the Server. The Global-Atlas Client interface is a Java applet. Global-Atlas supports three different types of query search, namely in-

Fig. 3. Query of landmarks (places of interest) from a Singapore Map

clusion, covered and intersection searches. Queries are expressed as polygons drawn on the map and the polygon shape is base on the map's projection type. Polygon drawing of query boxes will be explained in section 5. Results are classified under three categories: maps, documents and landmarks. The user can choose to browse the retrieved documents and maps in any web browser. Map results can also be reused as new search interfaces for refine queries. Landmark information consists of the name and the longitude and latitude coordinate point. It is being visualize as a point on the map interface of the Global Atlas search engine.

The Global-Atlas Administrator Tools, illustrated on figure 4 is a suite of tools that facilitate the population and maintenance of the database. It supports the automatic indexing of text (HTML) documents, landmarks as well as the semi-automatic calibration of map images (by co-registration [12]). Each HTML document and map image stored in the database has a GeoPrint attribute. A GeoPrint attribute contains the name of the document/map, related keywords, latitude and longitude values of the bounding box and also the set of ground control points (for maps). Ground control points are points where the (longitude, latitude) values are known together with the pixel coordinate of the map.

The Global-Atlas search engine Server works as the backbone of the whole Global-Atlas system and supports storage of the document references. It fol-

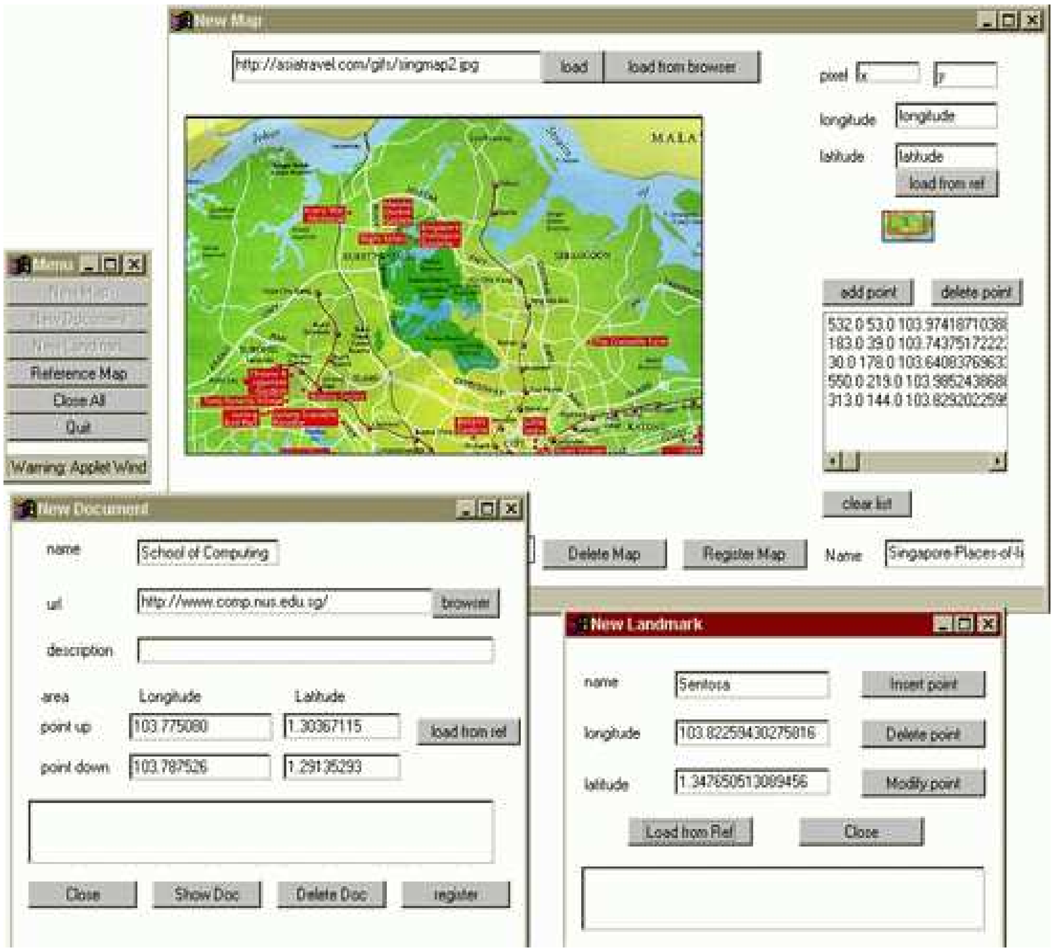

Fig. 4. Global-Atlas Administrator Tools

lows the now traditional three-tiered architecture of Web-based database applications: Web client (browser +Applet)/web server/database server. The Web server is Microsoft Internet Site Server. The server processing is done by Java Servlets using Jrun. The Java DataBase Connectivity package is used to establish the database server connection. The database server is Oracle release 8.0.5. Spatial data storage and spatial query processing is handled by Oracle Spatial Cartridge extension. Keyword indexing is based on Oracle Context Cartridge.

4 Registration of Geographical Images

The study of the development of a geographical search engine leads to a major issue of calibrating different projected map images to a single global projection coordinate system. Map calibration or registration is a fundamental task in map processing used to match two or more map images of different map projections. World maps in atlas provide a spatial referencing system, but the same pieces of the earth can look different according to the map projection used for the map. [13]. A map projection is the manner in which the spherical surface of the Earth is represented on a two dimensional surface. The three most common map projections are the azimuthal, cylindrical and conical. Azimuthal projection may be imagined as the transformation to a projection plane, which is tangential to the generating globe. Cylindrical projection may be imagined as the transformation to the plane if this is wrapped round the globe in the form of a tangent cylinder. Conical projection may be imagined as the transformation from the sphere to the plane through the medium of a cone wrapped round the globe. Virtually, all large systems, which evaluate images, require the registration of images, or a closely related operation, as an intermediate step. In this section, we review the broad range of registration techniques surveyed and determine the most appropriate technique for the registration of maps obtained from the Internet.

4.1 Image Registration Methods

The fundamental characteristic of any image registration technique is the type of spatial transformation or mapping used to register two images. Transformation is a mapping of locations of points in one image, to new locations of points in another. In global transformation, the map is composed of a single equation that maps each point in the first image to new location in the second image. In local transformation, the mapping of points in the image depends on their location - the map is composed of several smaller maps (several equations) for each piece of the image that is considered. Most registration techniques involve searching over the space of transformations of a certain type to find the optimal transformation for a particular problem. Polynomial transformation is one of the most general global transformations. It can account for many types of distortions so long as the distortions do not vary too much over the image.

In the paper by Brown L. G. [14], several registration methods are categorized according to the type of variations between images. As maps obtained from the Internet are generally of unknown projection type and misalignment, modelling the transformation is difficult. At most we could obtain varying number of control points either manually or semi-automatically from expert recognition of the image maps. Hence, the Point- or Landmark-mapping technique, is a more appropriate registration methods for calibration of Internet maps. The point- or landmark-mapping technique is the primary approach currently taken to register two images whose type of misalignment is unknown [15]. This occurs if the class of transformations cannot be easily categorized, but we can use the landmarks that can be found in both images and match them using a general transformation. The general method for point mapping consists of three stages: (1) Compute Features in the image. (2) Feature points in the reference image (control points) are corresponded with feature points in the data image. (3) A spatial mapping, usually two 2D polynomial functions of a specified order (one for each coordinate in the registered image) is determined using these matched feature points.

However, for images with transformation which are truly difficult to model and match, point-mapping methods could be used together with a user feedback between the stages of finding the correspondence between control points and finding the optimal transformation. These techniques rely on more sophisticated search strategies including relaxation [16], cooperation, clustering [17], hierarchical search, and graph matching. Image maps obtained from the Internet have unknown projections, and hence have unknown misalignment. The most appropriate technique to calibrate or register these image-maps, are the point- or landmark mapping technique which is the general technique for registering images of unknown misalignment. In this paper, we present our research findings on the study of various transformation techniques and our own original hybrid calibration method and double transformation technique for the Global Atlas search engine.

4.2 Transformation Methods

The total number of map projections, which can be described, is infinitely large [18]. From this population about 400 have been described, though less then one quarter of them have been named and used. With this vast amount of different projections of maps, it is truly difficult to determine the projection type of a map image independently published on the World Wide Web. Hence, there is a need to find a suitable transformation method to refer to any map image into a single global coordinate system for the Global Atlas search engine. The transformation from one map projection to another is essentially the coordinate conversion of a point on one plane to the corresponding point on another plane. The Cartesian coordinates (x,y) of a point on a map are functionally related to position on the earth's surface expressed in geographical coordinates (Φ, λ), where

$$x = f_1(\Phi, \lambda); \qquad y = f_2(\Phi, \lambda)$$

There are three basic methods of transforming (x,y) coordinates to (Φ, λ) coordinates [19,20]. They are the Analytical transformation, General Numerical transformation and Direct or grid-on-grid transformation. The analytical method is the most obvious and straightforward solution to the problem when it applies. It simply uses the analytical formulae for f_1 and f_2. Of course it assumes that the map coordinate system has been identified and that the formulae are known, which is impractical in our case. Both the Grid-on-grid and the general numerical methods rely on the knowledge of ground control points. Ground control points are points where the longitude and latitude values (Φ, λ), together with the pixel values (x,y) of the image, are known. If the ground control points are known to be exact, interpolation methods are used. However, if the ground control points are approximate values, Least Square Regression (LSR) methods are more suitable. By making assumptions about the general form of the functions f_1 and f_2 for which we know the exact or approximate values at the ground control points we are left with solving a system of equations. The functions are generally chosen to be polynomials of a fixed degree in x and y but this is not necessary.

The standard methods discussed below work well as long as the resulting system of equations in the parameters of the function is linear. It is also possible, of course, to consider functions that lead to non-linear systems of equations provided that a method for their resolution is known and practical. The reader may consult [20] for a detailed description of these methods, and their variants (e.g. Weighted and Moving Least Square) or extensions (B-Spline method). The required number of exact (interpolation) or minimum (LSR) number of ground control points needed to determine the coefficients vary according to the order or degree of the polynomial. For a polynomial of degree, n, the number of ground control points required is $\frac{1}{2}(n+1)(n+2)$.

The Grid-on-Grid transformations are special cases of generic interpolation methods. However they are specific for they make specific assumptions about the types of geometric transformations or invariants the functions f and g must embed. The most commonly used Grid-on-Grid methods are the linear conformal, similarity or Helmert transformation; and the Affine transformation. Both use linear polynomials for f_1 and f_2, and attempt to capture three basic operations: scaling, rotation, and translation. We also designed and used an original method, which can be classified as a Grid-on-Grid method and is called the Proportional method [21]. The Proportional method uses linear polynomials.

5 Map Images Registration and Reuse in Global-Atlas

The motivation of the study of the calibration of a map image serves a dual purpose: to determine a precise bounding box of the region represented for indexing, and to determine the transformations necessary to reuse the map image as a query interface. Our study focuses on the Point Mapping -Global Polynomial technique. In order to reuse the map image as the query interface we need to compute both a direct and a reverse transformation between the map image and the reference coordinate system. This is to provide user with a better visual feedback of the query rectangle.

Figure 5 illustrates the principle of the two transformations. On the left hand side is the indexed map image. The map is in an unknown projection, although we can see the curvature of the parallels and meridians drawn on the map. On the right hand side is a representation of the reference coordinate system in an orthogonal longitude/latitude projection. Two points set by the user, when pressing, dragging, and releasing the mouse defines the query polygon on the left image. It corresponds to a query rectangle in the reference system. The direct transformation is used to map the query points (in pixel values) selected on the map image to the longitude/latitude coordinate system and define query rectangle. The accuracy of this transformation impacts directly the recall and precision of the query. The reverse transformation is used to provide to the user a visual feedback of the actual region being defined by the query. Indeed, as illustrated on figure 5 because of the possible mismatch of projection, the query rectangle may not be a rectangle in the original map coordinate system.

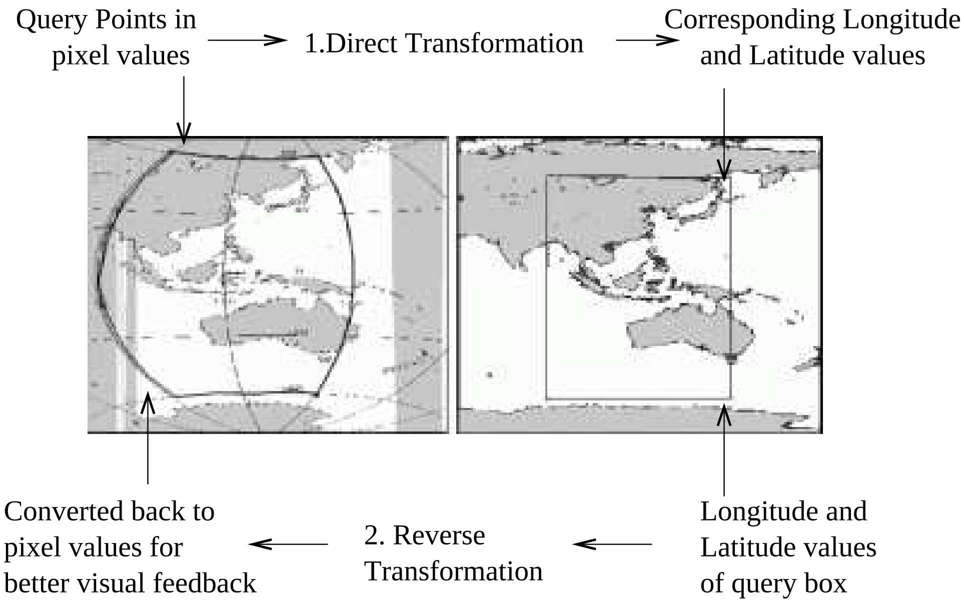

Fig. 5. Direct and Reverse Transformations for a Query

Depending on the type of function obtained by the direct transformation, it may be too difficult or impossible to obtain analytically an exact inverse (e.g. for polynomials of degree two and above). Solving the corresponding system of non-linear equations for each point may also be too expensive [19]. The simplest idea is to use a similar approach as for the direct transformation and estimate the reverse transformation. However, we wish to take one additional criterion into account in order to provide the user with more visual comfort. Namely, as illustrated by figure 6 it is possible that an approximate reverse transformation does not exactly match the query corner points. This results in a query box being drawn on the query map image at an offset position from the start and end points selected by the user (mouse down and mouse released). This phenomenon, as we have experienced, can be rather disturbing for the users. Therefore we wish to reduce its amplitude or eliminate it without losing too much of the accuracy of the query feedback. In other words, we are looking for a method to compute an accurate reverse transformation that respects as much as possible the identity of the query corner points.

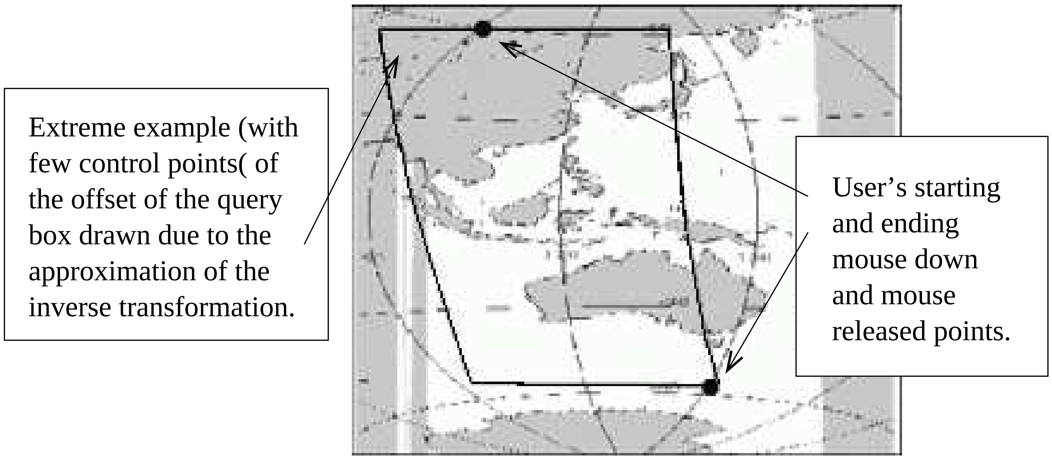

Fig. 6. Possible offset after double transformation

5.1 Options on the selection of transformation methods

For the selection of transformation methods for the Global-Atlas, no constraint is placed on to selecting the type and number of transformation or the degree of the polynomial. In actual fact, an optimum combination of these variables can be selected depending on the ground control points available, so as to obtain the best set of transformation results. The initial idea was to use the Least Square Regression (LSR) technique for the first transformation, as it allows varying ground control points usage. The plan was to use the polynomial of the corresponding degree, once its minimum number of ground control points is available. However, as will see in the next section, the experimental results suggested other criteria. We were also interested in evaluating the effectiveness of the Proportional method [21] we had designed, as we were interested to use it for its simplicity.

For the reverse transformation, we looked into three different methods: the Weighted (or alternatively the Moving) Least Square, the Grid Least Square, and the Constrained Least Square. The weighted least square regression is a variant of the least square method where the role of the different ground control points is controlled by means of weights. We therefore include the query corner points the second transformation giving them a heavy weight.

Notice that the proper weights have to be determined empirically. Weights too light do not achieve the objective of fitting the query corners, while weights too heavy compromise the accuracy of the overall result. The moving least square is a variant where the influence of the weight is greater in the neighborhood of the corresponding ground control points. First tests did not show enough improvement to convince us to try this second method. The Grid least square regression method is based on the belief that the more the merrier. Since we have already computed the direct transformation, we can create a large number of ground control points (for instance a grid to achieve a uniform coverage of the region) complying with the direct transformation. This method has the advantage of confining the computation of the reverse transformation to the registration phase whereas the two other methods require the reverse transformations to be computed in the client interface for each query. Finally the constrained least square regression method consists in forcing the surface computed by the LSR to contain the two query points. As a consequence we are guaranteed that there will be no change in values for these two points. The method can be seen as a mixture between interpolation and LSR. The equations to be manipulated are slightly more complicated but the complexity is essentially the same as the one of the standard LSR. The Grid LSR and the constrained LSR are original methods designed for our purpose; they are described in detail in [21].

5.2 Methodology

We now wish to compare the performance of the various transformation methods for both the forward and reverse transformation. Because of the nature of the methods, we need to study their effectiveness for varying numbers of control points. In order to obtain results which we can quantify we need a data set of which we have full knowledge. We first randomly select some regions for our test. A region is defined by a maximum and a minimum longitude and latitude. For each region we create a grid [12], i.e. a set of equally spaced points in the longitude/latitude coordinate system. For each point in the region, we measure or compute the corresponding pixel or centimeter coordinates in both a slave picture, and a master picture, namely in the master and slave grids. The slave picture and grid represent the region in one of the possible map projection. The master picture and grid represent the same region in the reference map projection, i.e., in our case, the longitude/latitude coordinate system. Some of the points in the region are then randomly chosen to be the ground control points on the slave picture. All the points (including the ground control points) are then used to evaluate the effectiveness of the transformation essentially by comparing their actual image with their image by the transformation.

Obtaining such a data set from real maps or satellite pictures would be tedious and error prone. We therefore use the Generic Mapping Tool (GMT) [22] and its database for the generation of multiple maps of the same region with different projections. In fact we do not even generate actual pictures but just files containing the different coordinates needed for the set of points defining the region. In this approach, we of course make the assumption (which may be erroneous in some extreme cases) that the GMT system is capable itself of sufficiently accurate analytical transformations for the various map projections considered.

In the experiments, we use six different and representative map projections (details of the regions and projections used as well as non-aggregated results are presented and discussed in [21]). We generate a master grid for eighteen different regions. For each region we generate the corresponding slave grids. Of course the decisions taken on the basis of the experiment's results may vary according to the actual frequency of one or the other projection. We believe that our test set is representative of maps found on the Internet.

To determine the precision and accuracy of the transformations, the average Euclidean distance is calculated for each pair of grids. We average the Euclidean distance ED between each point (x2, y2) on the grid and the transformed point (x1, y1). To take into consideration of map regions of different size, the Average Euclidean Distance for each map is normalised so as to obtain a fairer comparison. The normalised average Euclidean distance represents the amount of distance that is away from its actual point in terms of percentage size of the map being tested.

Normalized Average Euclidean Distance,

$$NAvgED = \frac{AvgED}{maxED}, \quad where \quad maxED = \sqrt{width^2 + height^2}$$

5.3 *Experiment on the Forward Transformation to determine a good transformation method*

In our first experiment we evaluate the accuracy of the various forward transformation methods for varying numbers of ground control points. A general distribution rule of the ground control points is that there should be a good distribution around the edges of the image, to be corrected with a scattering of points over the body of the image [12]. This is necessary to ensure that the mapping polynomials are well behaved over the whole image. We aim at this objective by randomly selecting the ground control points in six quadrants visited in a round-Robin manner.

Given the discussion on the different transformations, it must be clear that we are essentially interested in figuring out the best choice of the degree of the polynomial considered in a Least Square Regression technique [20] given a number of available ground control points. However, we suspected that grid-on-grid methods may outperform LSR with polynomial degree 1 when few control points are available and included the Proportional method (for the case of 4 ground control points) in our tests. The results presented average 1000 random selection of 3 to 30 ground control points for each of the 100 pairs of grids.

The results of this experiment show that with increasing number of ground control points, the accuracy of the transformation for each degree polynomial shows an improving trend. However, the improvement is of a decreasing rate, and there comes a point for each polynomial where the improvement is practically of no significance in pixel values.

A first look at figure 7 reveals that each polynomial of the least square regression requires a certain minimum number of ground control points before it is able to produce a relatively accurate transformation. This minimum number is different from the minimum requirement of ground control points for each polynomial. The minimum number of ground control points for polynomial degree 2 is 6, but at ground control point 6, the transformation results are better if polynomial of degree 1 is used instead. Only at ground control point 7 is the polynomial of degree 2 better than that of degree 1. The same goes of polynomials of degree 2 and 3. Studying the results, polynomial 3 is better than polynomial 2 only for ground control points of 14 and above. This hence invalidates the initial idea of selecting the polynomial degree based on its minimum requirement of ground control points. Transformation results using its higher order polynomial are not better than if the lower order polynomial is used.

For polynomial of degrees 4 and 5, the transformation accuracy shows continuous improvement with increasing ground control points. However, the results obtained for up to 40 control points, are still not as good as for polynomial of degree 3. Hence, the conclusion is that, higher polynomial degree requires even greater number of ground control points before the transformation is able to show accurate results. The belief is that, if more ground control points are added for testing, polynomials of degree 4 and 5 will eventually reach a better accuracy than polynomial 3. However, it is not cost effective, as more computation power is required for polynomials of higher degrees and the improvement may be neither noticeable by the user nor relevant for the search.

Comparing the Proportional transformation method against LSR polynomial

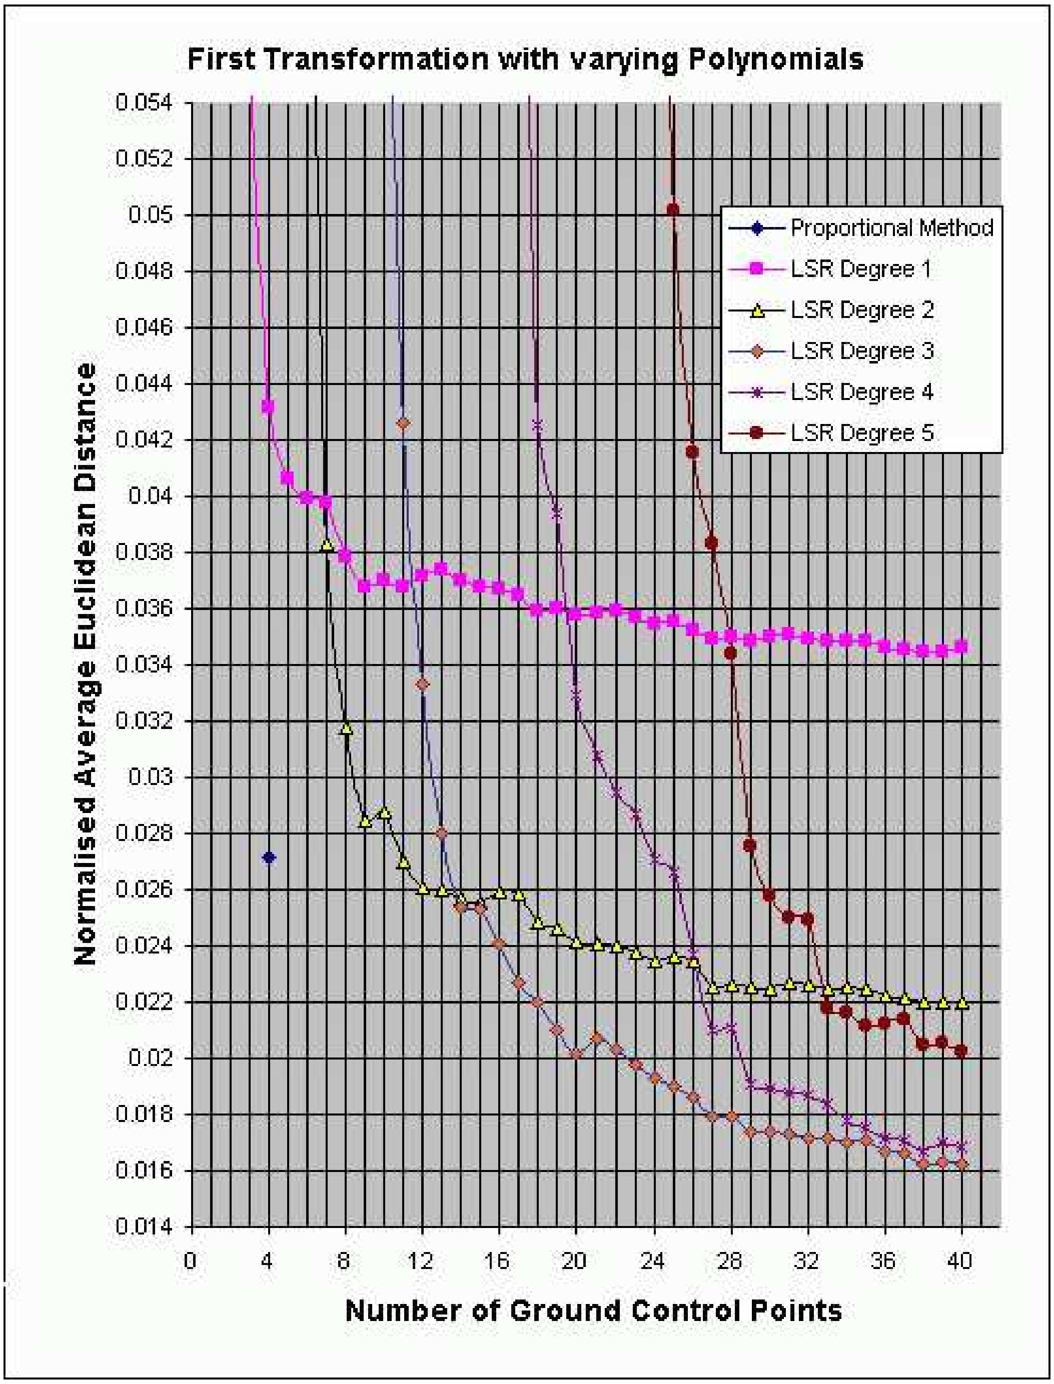

Fig. 7. First transformation results with varying Ground Control Points

of degree 1 and 2, we realize that the Proportional method is a better transformation method for ground control points of up to 10 (we need however to select four out of ten).

Hence, to improve the mapping transformation accuracy, we have modified our initial idea of polynomial degree usage with the amount of ground control points as illustrated by the following table:

Table 1
Number of ground control points vs LSR Polynomial degree usage

	Number of Ground Control Points	Transformation Method
Inital Plan		
	3 to 5	LSR Polynomial Degree One
	6 to 9	LSR Polynomial Degree Two
	10 and above	LSR Polynomial Degree Three
New Modified Plan		
	3	LSR Polynomial Degree One
	4 (to 10)	Proportional Transformation
	10 to 13	LSR Polynomial Degree Two
	14 and above	LSR Polynomial Degree Three

For the For the range of 5 to 10 ground control points, 4 of the best distributed ground control points is to be selected so as to perform the Proportional Transformation. As illustrated by figure 8 and 9, it is concluded that the new modified method of determination of polynomial degree usage would have a better mapping transformation performance.

Fig. 8. Initial transformation method using LSR with a greedy increment of the polynomials' degree

As mentioned above, our conclusion is based on the assumption that the dis-

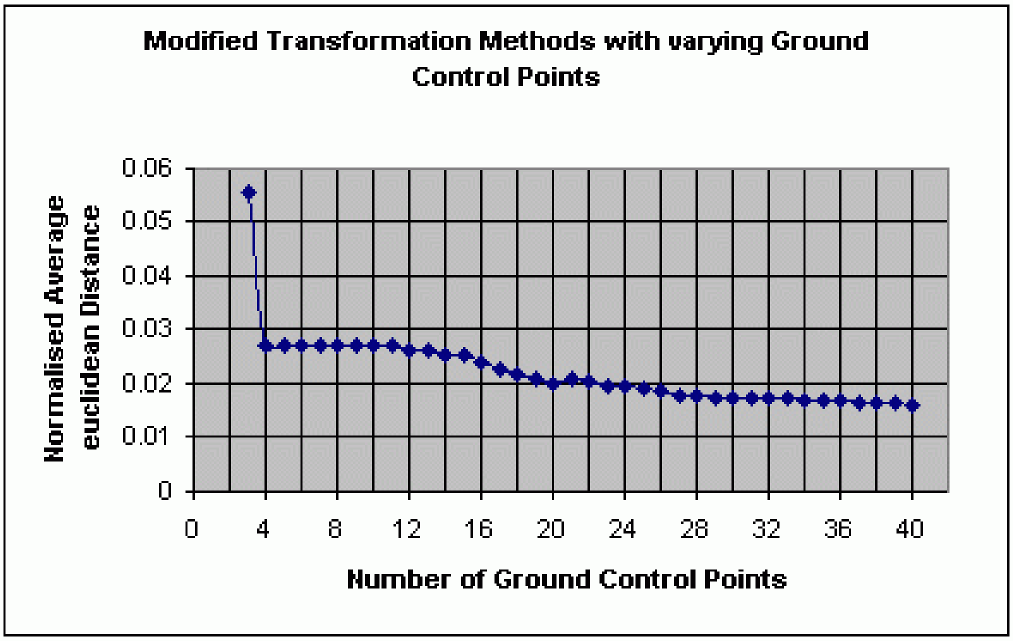

Fig. 9. New hybrid transformation method tuned according to empirical results

tribution of our sample is representative. Our conclusion may be refined, using the same methodology, as we gain more insight from the experience of operating Global Atlas about the most commonly found projections. However, so far our experience with Global Atlas did not contradict our choice of a balanced sample.

5.4 Quality of the Reverse Mapping Transformations

In order to provide the user with a more relevant query bounding box on the Global-Atlas system, a double transformation procedure is performed. However, as the perfect reverse of the first transformation using the Least Square Regression method could not be obtained, the feedback of the query box for the user becomes visually incorrect. Hence, this drives the motivation for this second experiment, which is to select the best reverse transformation technique to improve the visual query bounding box returned to the user. A number of different map projections are selected for testing to represent the image maps. The best transformation method, based on the results, is to be incorporated into the Global-Atlas search engine.

The transformation techniques to be to be experimented are the Proportional Method, the Least Square Regression, the Weighted Least Square Regression, the Grid Least Square Regression and the Constrained Least Square Regression techniques. Base on these techniques, the following experiment to determine the Recall and Precision of the query bounding box returned to the user is performed. Recall refers to how much the relevant information is returned to the user. Precision refers to the how relevant the results return are as compared to the number of result returned. In our case, Recall refers to how much of the actual region (data points) queried by the user is returned in the query bounding box drawn. Precision refers to how much of the returned (data points) query bounding box is within the relevant region.

The Grid least square regression method is based on the belief that the more the merrier. Since we have already computed the direct transformation, we can create a large number of ground control points (for instance a grid to achieve a uniform coverage of the region) complying with the direct transformation. This method has the advantage of confining the computation of the reverse transformation to the registration phase whereas the two other methods require the reverse transformations to be computed in the client interface for each query.

Finally the constrained least square regression method consists in forcing the surface computed by the LSR to contain the two query points. As a consequence we are guaranteed that there will be no change in values for these two points. The method can be seen as a mixture between interpolation and LSR. The equations to be manipulated are slightly more complicated but the complexity is essentially the same as the one of the standard LSR.

The Grid LSR and the constrained LSR are original methods designed for our purpose; they are described in detail in [21]. Base on these techniques, experiments (see figures 10 and 11) are perform to determine the Recall and Precision of the query bounding box returned to the user is performed. In our case, Recall refers to how much of the actual region (data points) queried by the user is returned in the query bounding box drawn. Precision refers to how much of the returned (data points) query bounding box is within the relevant region.

$$Recall = \frac{AmountofrelevantregionsCoveredbyreturnedquerybox}{Totalamountofrelevantregions}$$

$$Precision = \frac{AmountofrelevantregionsCoveredbyreturnedquerybox}{TotalnumberofregionsCoveredbyreturnedquerybox}$$

Notice that recall and precision as we define them above correspond to the visual feedback and are only indirectly related to the results of subsequent query evaluation, which we do not discuss in this paper.

Base on the data set of 100 maps generated using GMT, a simulation of the

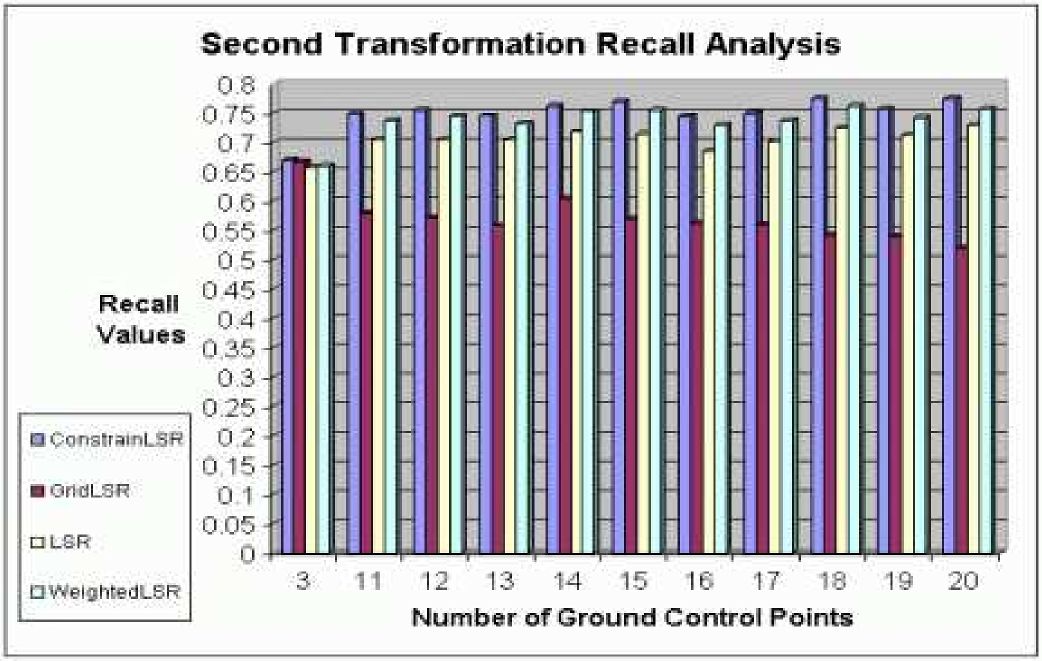

Fig. 10. Reverse transformation recall

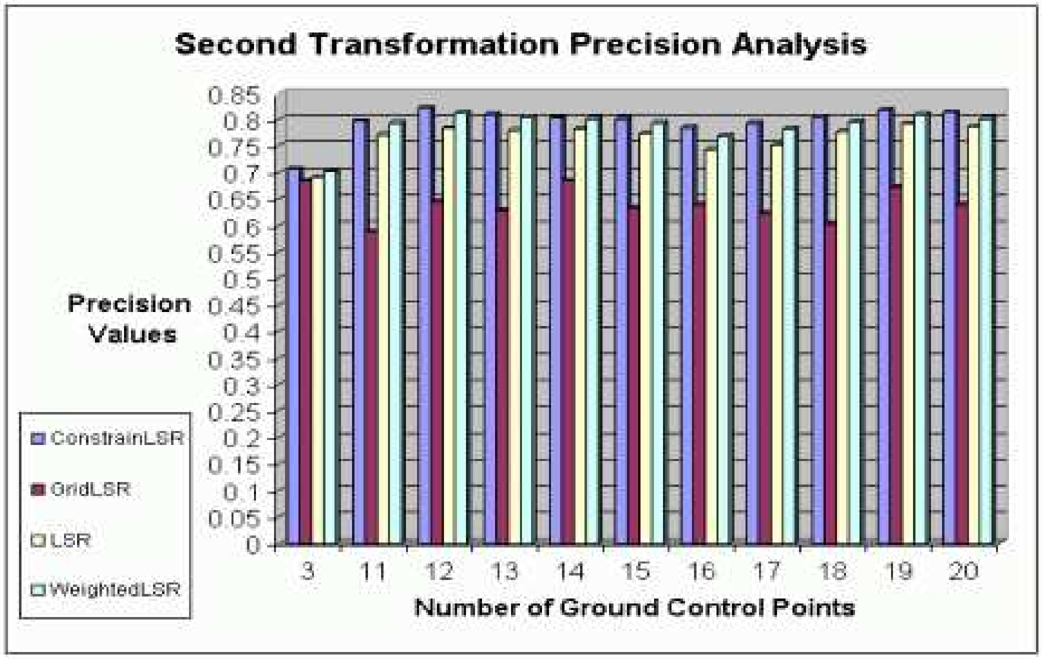

Fig. 11. Reverse transformation precision

user query onto the maps is performed, so as to test the second transformation. This simulation consists of selecting two random query points to represent a user query bounding box. With these two points, the first transformation is performed to obtain the longitude and latitude values of the query. This is to generate a rectangular bounding box based on the longitude and latitude values. The second transformation is then performed for all points on the bounding box to obtain the corresponding user query box on the data map. The Proportional method is tested for the second transformation from 4 to 10 ground control points. Looking at the results for the proportional transformation the recall and precision values is around 0.57 and 0.64 respectively. These values are not as good as the values obtained from the Constrained least square regression transformation, which is discussed below.

For ground control points of 3, 11 and above, the experiment is performed on the least square regression method, weighted LSR method, and the constrained LSR method. The results of the experiment shows that the Constrained Least Square Regression is the best overall technique for both recall and precision, with the Weighted technique coming in a close second. As illustrated by figure 10, it is quite obvious that the Constrained Technique has a much higher recall values than the conventional LSR technique. This implies that by using the Constrained Technique, the visual query box returned to the user will cover fewer irrelevant regions. Figure 11 shows the precision values of the various LSR transformation techniques. It appears that the precision values for the three techniques are quite close, although, the Constrained LSR technique does have a slight edge over the other two techniques and provide the best possible fit for the query corners.

Hence, our conclusion on this experiment is that, although all three techniques have about the same accuracy/precision, the Constrained LSR technique is better at providing the user with a query bounding box that does not cover as many irrelevant regions and provide interaction comfort. The Grid LSR, despite its performance, remains an optional method that can be used for lightweight clients.

6 Conclusion and Future Work

In order to assist the administrator of the Global-Atlas search engine in task of registering maps to be retrieved and reused as graphical query interfaces, we have designed an original double transformation method, which combines several existing surface-fitting methods. The parameters of the combination are the results of an extensive performance evaluation of the methods based on a representative data set.

Its performance has been tuned using maps generated in different projections with the Generic Mapping Tool [22]. As the initial set of control points a user can provide by looking at the map is usually small (3-4 points at most), we use a database or features and remarkable point to propose additional reference points based on the initial approximation. The map registration process in

Global Atlas can be summarized as follow: 1. First the image is uploaded into the administrator tool. The user indicates the initial control points and their coordinates. 2. Iteratively, the system uses its current interpolation to look in its database for additional control points to propose. 3. When the administrator is satisfied, the map is stored with the interpolated transformations and indexed by the inferred bounding box.

As a search engine, Global-Atlas is not bound to providing the exact and complete set of results to a search. It is indeed unlikely that the user will browse the most probably very large amount of candidate documents returned to her. At the same time, by nature of the content of many documents, in particular text documents, the geoprint (or geoprints) is likely to be necessarily vague. Finally, not only the computer screen resolution limit the precision of the query expressed, but also, the precision intended by the user herself is often relatively low. The quality of the registration process methods and its effectiveness will be measured in recall and precision of the search mechanism it enables. We have so far been able to develop and use techniques leading to very satisfactory results.

We are planning to use a geographical name server [23], by matching extracted geographical names or landmarks from map or documents, to automatically obtained the ground control points. It is foreseen that the demand for location dependent application will motivate the information providers to include geographical information in the form of Meta data, eg using emerging standards such as geoXML (url: http://www.geoxml.de/english/index.html). We will also be looking at ways to support geoXML standard in Global Atlas so as to read any information available on the web.

To automate the image registration process for usage in the Global Atlas search engine, more extensive research has to be made on the various techniques of point mapping with feedback. This would require architecture of a user interactive system for registering image maps. In order to semi-automatically increase the number of control points, we plan to connect the administrator tool to a gazetteer (or name server) which will interactively propose known land marks as control points to be selected or re-positioned by the administrator. We also foresee the possibility to use a database of visual features such as coastlines and rivers to match the automatically extracted features on the maps and propose more anchors for the accurate determination of the transformations. Global-Atlas prototype had been on demo at the EDBT2000 [24] and has been presented in the WISE2000 conference [25].

Global-Atlas (URL: http://turtle.comp.nus.edu.sg/gatlas) offers an additional answer to the problem of precision of the search with the possibility to reuse retrieved maps to refine the search.

Acknowledgements

The authors would like to thank Dave Abel for his comments and suggestions on an early draft of this paper.

References

[1] B. Ooi, Efficient query processing in geographic information systems, Lecture Noes in Computer Science 471 (1990) 208.

[2] B. Ooi, E. Bertino, R. Sacks-Davis, K. Tan, J. Zobel, B. Shidlovsky, B. Cantania, Indexing techniques for advanced database applications, in: Monograph Series, Kluwer Academics, 1997.

[3] S. Bressan, P. Bonnet, B. Granveaud, The world travel editor: A geographical search engine for the internet, in: Proceedings of the International Conference on Extending Database Technology, 1996.

[4] T. Smith, A brief update on the alexandria digital library project: Constructing a digital library for geographically referenced materials, D-Lib Magazine .

[5] Y. Be'dard, F. Letourneau, C. Martel, M. Proulx, Georep, a www customizable georeferenced digital library for spatial data, D-Lib Magazine .

[6] T. Barclay, J. Gray, D. Slutz, Microsoft terraserver: A spatial data warehouse, in: 25th VLDB Conference, 1999.

[7] S. Kitazawa, H. Mizuno, T. Tokuda, T. Yamada, Map-based information mediation on the web with float coordinate system, in: IEEE/SMC, 1999, pp. 93–98.

[8] S. Kitazawa, H. Mizuno, T. Tokuda, T. Yamada, On extraction of www pages with spatial information for mediating their distribution, in: IEEE/SMC, 1999, pp. 74–79.

[9] S. Kitazawa, T. Sobue, M. Tamano, H. Tsuji, T. Yamada, Environment for spatial information sharing, in: Digital Cities: Experience, Technologies and Future Perspectives, Springer, 1999, pp. 315–325.

[10] H. Tsuji, T. Yamada, Browsing local community "what's new" in spatial context, in: I'MEDIATE, 1998, http://context.mit.edu/imediat98/paper5.

[11] U. of Tokyo, Center for spatial information science, http://www.csis.u-tokyo.ac.jp.

[12] J. Ricards, Remote Sensing Digital Image Analysis: An Introduction, 3rd Edition, Springer-Verlag, 1993.

[13] A. Robinson, Elements of Cartography, New York, USA: John Wiley, 1984.

[14] L. Brown, A survey of image registration techniques, in: ACM Computing Surveys, Vol. 24, 1992.

[15] A. Kain, J. Ton, Registering landsat images by point matching, in: IEEE Trans. Geosci. Remote Sensing, 1989, pp. 642–651.

[16] S. Ranade, A. Rosenfled, Point pattern matching by relaxation, in: Patt. Recog., Vol. 12, 1980, pp. 269–275.

[17] S. Benett, S. Kopstein, G. Stockman, Matching images to models for registration and object detection via clustering, in: IEEE Trans. Patt. Anal. Machine Intell., 1982, pp. 131–147.

[18] D. Maling, Coordinate Systems and Map Projections, 2nd Edition, Pergamon Press, Oxford, 1992.

[19] D. Ghilani, P. Wolf, Adjustment Computations Statistics and Least Squares in Surveying and GIS, John Wiley and Sons Inc, 1997.

[20] P. Lancaster, K. Salkauskas, Curve and Surface Fitting an Introduction, London: Orlando: Academic Press, 1986.

[21] F. Lee, G-atlas: A geographical search engine for the internet, Honours thesis, School of Computing, National University of Singapore (1999).

[22] P. Wessel, H. Smith, The Generic Mapping Tools GMT. Technical Reference and Cookbook, Geodynamics Branch, Geosciences Laboratory, 1995.

[23] Getty, Geographic name server, http://shiva.pub.getty.edu.

[24] S. Bressan, B. Ooi, F. Lee, Demonstration of the global atlas prototype, Proceedings of the EDBT conference demo track .

[25] S. Bressan, B. Ooi, F. Lee, Global atlas: Calibrating and indexing documents from the internet in the cartographic paradigm, Proceedings of the 1st International Conference on Web Information System Engineering (WISE 2000) .