Final Report for the DFG-Project
Algorithms for Interactive Variable-Scale Maps

Jan-Henrik Haunert, Alexander Wolff, Thomas C. van Dijk

October 2, 2015

Project Identifier (DFG-Geschäftszeichen)

Ha5451/3-1

Applicants

Prof. Dr.-Ing. Jan-Henrik Haunert
Institut für Geoinformatik und Fernerkundung
Universität Osnabrück

Prof. Dr. Alexander Wolff
Institut für Informatik
Lehrstuhl für effiziente Algorithmen und wissensbasierte Systeme
Universität Würzburg

Project Title

Algorithms for Interactive Variable-Scale Maps
(Algorithmen für interaktive Landkarten mit gleitendem Maßstab)

Funding Period

January 15, 2012 – January 15, 2015

Project Publications

Journal Articles

[A] J.-H. Haunert and L. Sering.
Drawing road networks with focus regions. *IEEE Transactions on Visualization and Computer Graphics (Proc. Information Visualization 2011)*, 17(12):2555–2562, 2011.

[B] T. C. van Dijk and J.-H. Haunert.
Interactive focus maps using least-squares optimization. *International Journal of Geographical Information Science*, 28(10):2052–2075, 2014.

Papers in Peer-Reviewed Conference and Workshop Proceedings

[C] T. C. van Dijk and J.-H. Haunert.
A probabilistic model for road selection in mobile maps. In *Proc. 12th Int. Sympos. Web and Wireless Geograph. Inform. Systems (W2GIS'13)*, volume 7820 of *Lecture Notes in Computer Sci.*, pages 214–222. Springer-Verlag, 2013.

[D] T. C. van Dijk, K. Fleszar, J.-H. Haunert, and J. Spoerhase.
Road segment selection with strokes and stability. In *Proc. 1st ACM SIGSPATIAL Workshop MapInteraction*, pages 72–77, 2013.

[E] M. Chimani, T. C. van Dijk, and J.-H. Haunert.
How to eat a graph: Computing selection sequences for the continuous generalization of road networks. In *Proc. 22nd ACM SIGSPATIAL Int. Conf. Advances in Geograph. Inform. Systems (ACM-GIS'14)*, pages 243–252, 2014.

[F] T. C. van Dijk, A. van Goethem, J.-H. Haunert, W. Meulemans, and B. Speckmann.
Map schematization with circular arcs. In *Proc. 8th Int. Conf. Geograph. Inform. Sci. (GIScience'14)*, volume 8728 of *Lecture Notes in Computer Sci.*, pages 1–17. Springer-Verlag, 2014.

[G] N. Schwartges, J.-H. Haunert, A. Wolff, and D. Zwiebler.
Point labeling with sliding labels in interactive maps. In *Proc. 17th AGILE Int. Conf. Geograph. Inform. Sci. (AGILE'14)*, Lecture Notes in Geoinform. and Cartography, pages 295–310. Springer-Verlag, 2014.

[H] N. Schwartges, A. Wolff, and J.-H. Haunert.
Labeling streets in interactive maps using embedded labels. In *Proc. 22nd ACM SIGSPATIAL Int. Conf. Advances in Geograph. Inform. Systems (ACM-GIS'14)*, pages 517–520, 2014.

[J] N. Schwartges, B. Morgan, J.-H. Haunert, and A. Wolff.
Labeling streets along a route in interactive 3d maps using billboards. In *Proc. 18th AGILE Int. Conf. Geograph. Inform. Sci. (AGILE'15)*, Lecture Notes in Geoinform. and Cartography, pages 269–287. Springer-Verlag, 2015.

1 Project Goals

Nowadays, interactive maps are widely used in navigation systems or on the Internet. Most online map systems (such as products by Google or Open Street Map) allow for zooming and panning, that is, changing the scale and the position of the current map view, respectively. Typically, the user gets to see a map from a small set of preprocessed maps. When a user calls a zoom function, the map viewer switches to another map, which may cause abrupt changes to what the user sees. As a result, the user may lose context.

In order to overcome this problem, our project aimed at algorithms for the generation of *variable-scale maps*. Under a variable-scale map we understand either (i) a sequence or continuum of maps that covers an interval of scales or (ii) a single map that displays different regions at different scales. A compilation of good single-scale maps is not the same as a good variable-scale map since variable-scale maps need to satisfy additional constraints. In order to produce good variable-scale maps, algorithmic solutions for continuous map generalization, variable-scale label placement, and map-layout generation are needed.

An important project goal was to formalize these tasks as optimization problems and, thereby, to make the quality of variable-scale maps measurable. Furthermore, the project aimed at efficient exact algorithms and heuristics for application in real-time systems.

The focus of the project was on maps of road networks, which are widely used in navigation systems. Those systems usually have a small screen, which allows either for a large-scale representation of a user's direct vicinity or for an overview map with very limited detail. Neither type of map, however, is well suited for navigation. An important goal of the project was, therefore, to develop graphical representations that provide both focus and context information. We term such representations *focus-and-context maps*. An existing approach to computing focus-and-context maps is to enlarge the regions in the focus of the user with a fish-eye projection. This usually leads to large distortions and thus to rather illegible maps. One goal of this project was to overcome this issue, by computing focus-and-context maps with minimum distortion. Furthermore, our aim was to decide automatically which line segments of a road network to display in a single map with multiple different scales and in a map that allows a user to zoom in and out. When the user zooms out, the segments should disappear in a consistent manner, thus we were faced with the problem of computing an appropriate selection sequence. We investigated possibilities of combining our techniques for focus-and-context maps with map schematization.

An important part in making traditional paper maps is label placement. While some first theoretical results [1] for labeling interactive maps had been established by the start of the project, we wanted to explore the practical difficulties that have to be solved in real time in dynamic scenarios such as in navigation systems (which additionally have very restricted resources in terms of display size and computational power).

2 Project Development and Results

We summarize our results concerning minimum-distortion focus-and-context maps of road networks (Sect. 2.1), the selection of road segments for such maps (Sect. 2.2), the computation of selection sequences for maps supporting continuous zooming (Sect. 2.3), map schematization (Sect. 2.4), and label placement (Sect. 2.5). The structuring of the discussion in Sections 2.1 to 2.4 has been chosen to clarify the motivations for the different problems as much as possible but also reflects the progress of the development over time. The methods for map labeling (Sect. 2.5) have been developed independently of the other methods and in parallel to them.

2.1 Focus-and-context maps with minimum distortion

Figure 1(a) shows a map of Boston in which a user has selected a circular focus region F. We suppose that the user wants to enlarge this focus region by a factor of 3, for example, to better understand the complex highway interchange that is very cluttered. After the enlargement of the focus region, the map should still show the same part of the network as before and use no more space than the original map. Importantly, the enlargement process shall not introduce new edge crossings. Distortions of the map are tolerable, since we assume that the map is not used for measuring distances, but the distortion should be minimized. Figure 1(b) shows the result of applying an existing fish-eye projection [6] that achieves these requirements; especially northeast of the focus region large distortions occur, in the form of bended lines that were straight in the input map. In order to minimize such distortions, we developed an optimization method [A]. More specifically, we formalized the problem as a convex quadratic program that consists of a convex quadratic objective function and a set of constraints in the form of linear inequalities. This type of problem can be solved efficiently; commercial solvers for convex quadratic programming are available. The objective function models the total distortion of the map as the sum of a local measure of distortion over all nodes of the network. For this local measure we take the star-shaped graph into account that contains a node, its incident edges, and its adjacent nodes; we consider and quantify every geometric change occurring to this graph as a distortion, except for uniform scaling and translation. This in essence means that different parts of a map can be moved and have different scales without being considered as distorted; distortions unavoidably occur at the connections between those parts. Figure 1(c) shows the result that we obtained with this method for the map of Boston. We observe that most of the originally straight lines remain straight in the output map. This achievement can also be expressed quantitatively: the total distortion is 75% lower in the map obtained by optimization (Fig. 1(c)) than in the map obtained with the fish-eye projection (Fig. 1(b)). Our method based on convex quadratic programming needed 6.7 seconds to compute the map in Fig. 1(c) on a state-of-the-art desktop PC.

We consider the running time of the method based on convex quadratic programming acceptable for the production of static maps, but it is not for dynamic maps in which the focus region can change. For example, if the focus region is

(a) original map and focus region F

(b) fish-eye projection [6], F scaled by 3

(c) our result [A], F scaled by 3

Figure 1: Three maps of Boston.

(a) original map with 300,000 edges

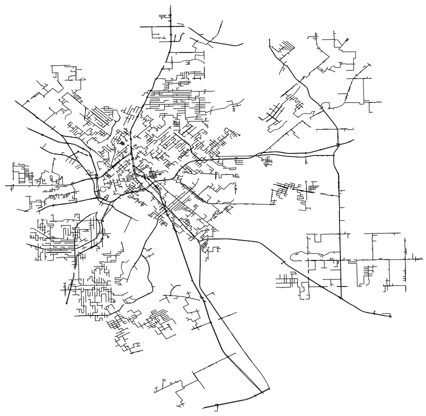

(b) after 90,000 edge deletions

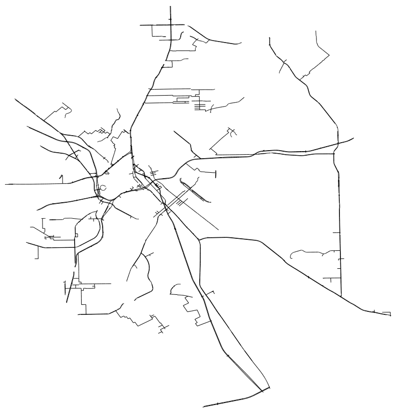

(c) after 101,000 edge deletions

Figure 2: Three maps of Dallas.

supposed to follow the path of a moving user, the focus-and-context map needs to be updated in real time. For this purpose, we developed a second optimization method based on least-squares adjustment [B]. In order to avoid abrupt changes between two subsequent maps, we continuously morph between them. To speed up the computation, we relax the hard constraints of our original model and integrate them into the objective function. Furthermore, we first solve the model without the constraints forbidding edge crossings. The solution that we obtain is not a good end result, but it can be used already to start the morphing process and to identify a small set of constraints that are needed to avoid edge crossings in the output map. By repeatedly solving the model with additional constraints, we can get rid of all edge crossings and, finally, obtain a crossing-free map of low distortion. This defines the end of the morphing process—and the beginning of a new morphing process if the user continues moving. With this approach we are able to respond to user interactions (i.e., selections of focus regions) in real time and, in particular, to let the focus region follow a user's trajectory continuously; see our video[1].

2.2 Selecting edges for focus-and-context maps of road networks

Enlarging dense regions in a map is not the only possibility of increasing the map's legibility. Instead, we can reduce the visual clutter by map generalization, which subsumes processes that adapt a map to a particular scale. This includes geometric simplification, selection, and aggregation, for example. In this project, we developed algorithms for selecting a subset of all edges of a given road network for a visualization. We subdivided the selection problem into two tasks. The first task is to compute a measure of importance (or weight) for each edge of the road network. The second task is to select a maximum-weight set of edges that satisfies certain constraints.

We developed a probabilistic model to solve the first task (defining the edge weights) for variable-scale maps [C]. In this model, the weight of each edge reflects the probability that, within a certain time, that edge will be visited by the user. We have shown that our measure is closely related to PageRank, which Google famously used to measure the importance of websites, and can be computed similarly. Our model allows us to address a set of different scenarios, ranging from the scenario that a user explores an unknown environment to the scenario that a user is following a precomputed route. When applying a naive algorithm for the selection task (that is, when selecting a user-specified number of edges of largest weight) we obtain a selection around the user's current position in the first scenario and a selection of the precomputed route with some context information in the second scenario.

For the second task (computing a selection of edges) we also developed efficient algorithms that either take the connectivity of the selection or the stability of the selection over time into account [D]. More specifically, we group the edges into sequences (so-called strokes) of low turning angles and require that every

[1] https://www.youtube.com/watch?v=fXxN8pT6i0c

stroke has to remain connected in the reduced graph, meaning that one is not allowed to remove edges from the interior of a stroke. With this requirement, we can compute a solution of maximum weight with an efficient algorithm based on dynamic programming. Furthermore, assuming that we can predict the changes of the edge weights over time, we considered the problem of computing selections of edges for a sequence of time steps. In this problem, two aims are modeled and integrated in an overall cost function: The total weight of the selected edges (i.e., the sum over all time steps) has to be large and the number of switches between the two possible states of an edge (being selecting or not) has to be small. For this problem we developed an efficient flow-based algorithm, with which we were able to reduce the number of switches drastically (by more than 80%) while still getting 90% of the weight that is achievable without considering the stability of the selection. An open question is whether the problem that respects both connectivity of strokes and stability over time can be solved efficiently. We were, however, able to prove NP-hardness for the case that the grouping of the edges into strokes changes over time, or the switch cost can depend on the edges involved.

2.3 Computing selection sequences to support smooth zooming

In a map supporting continuous zooming, the road network should become successively more generalized when zooming out. Roads that have vanished at some point in time during this process must not reappear when zooming out further. To avoid sudden changes, the generalization should proceed in very small steps. Therefore, we consider removing the edges of the road network one by one. We have developed algorithms to compute appropriate selection sequences [E], that is, to decide about the order in which the edges of the road network disappear. After computing the selection sequence once and storing it in an appropriate data structure, we can quickly select the edges for a map of any user-specified scale. As in our work on edge selection in focus-and-context maps (Sect. 2.2), we assume that the edges of the road network are weighted according to their importance. The edge weights can be computed using any method, e.g., PageRank. Our algorithms only require the weights to be positive real numbers. To ensure that important edges remain in the network relatively long, we maximize the sum of the total weight of all visible edges over all steps. As an important constraint we require that, after every edge removal, the network remains connected. We showed that, under the connectivity constraint, the problem is NP-hard, and we developed an integer linear program (ILP) that solves it optimally. Since solving this ILP is only feasible for small instances, we also developed constant-factor approximation algorithms and heuristics. We experimentally demonstrated that these heuristics perform well on real-world instances. Figure 2 on page 5 shows an input map of Dallas with two maps from the selection sequence that we computed.

2.4 Map schematization

Map schematization is the process of simplifying a map by enforcing unnatural regularities. In schematic metro maps, for example, it is common to constrain the

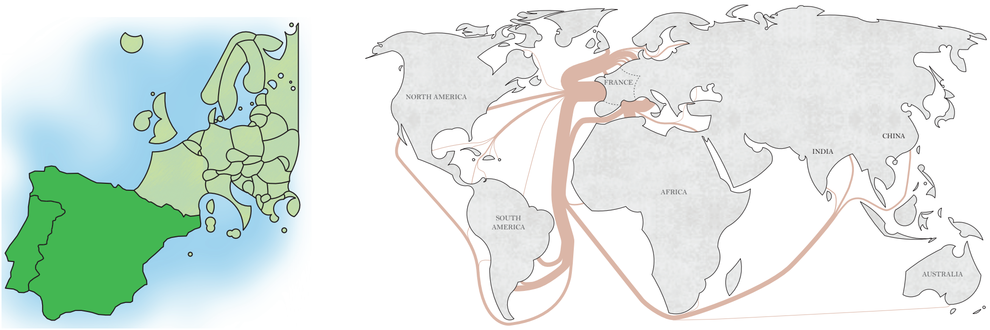

(a) a map of Europe (b) a map visualizing the export of wine from France

Figure 3: Examples of combining map distortion and map schematization.

orientation angles of the line segments to multiples of 45 degrees. Another form of schematization, which is less well explored in the cartographic literature, is to represent linear features using curves. We developed a method that schematizes a map automatically using circular arcs and explored possibilities of combining this method with our method for focus-and-context maps [F]. One motivation for this work was to allow map users to better recognize those parts of a focus-and-context map that have been enlarged. Unless one is very familiar with maps of Boston, it is hardly recognizable that the map in Figure 1(c) has been distorted. In the map of Europe in Fig. 3(a), on the other hand, the user is given very clear hints that the map is not to scale. In particular, we suggest schematizing and thereby de-emphasizing regions that are represented at a small scale and thus emphasizing the focus region. Our schematization method allows us to define different degrees of schematization for different lines in the same map. The degree of schematization can be used to convey information about the scale that is locally valid in a focus-and-context map. Figure 3(b) shows another map that we produced with our method. It depicts volumes of wine exported from France and resembles a famous 19th-century map[2] by the French engineer Charles Joseph Minard. In this map, we did not enlarge a certain country but introduced distortion to widen narrow passages such as the English Channel. This allowed us to represent the export quantities using line width.

2.5 Map labeling

With respect to map labeling, we have considered the problem of annotating objects in interactive maps with rectangular labels [G]. The users can zoom continuously, thus the label sizes as well as the size of the visible part of the map change (when measured in world coordinates); since the users are allowed to pan, the visible part of the map is subject to translations. To keep the labeling free of occlusions, we need to react to those changes, which we do by turning labels on or off and by sliding the labels horizontally. More precisely, we allow a label to be placed at any

[2] `https://commons.wikimedia.org/wiki/File:Minard%E2%80%99s_map_of_French_wine_exports_for_1864.jpg`

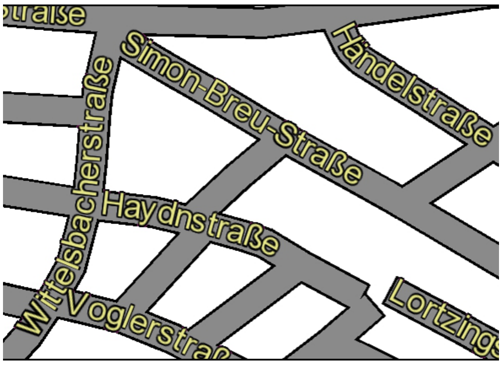

Figure 4: A map with embedded labels.

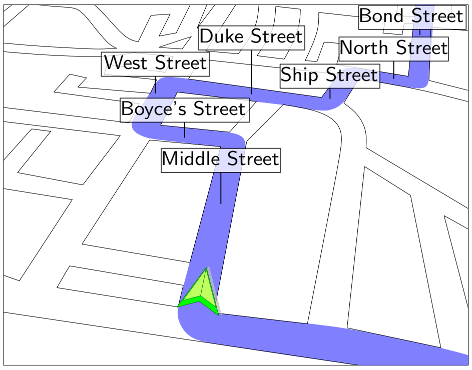

Figure 5: A map with billboard labels.

position where its lower side touches the corresponding point. We assume that each point comes with a priority; the higher the priority the more important it is to label the point. Given a dynamic scenario with user interactions, our objective is to maintain an occlusion-free labeling such that, on average over time, the sum of the priorities of the labeled points is maximized. Since even the static version of the problem is known to be NP-hard [5], we focused on heuristic approaches. We developed an efficient and effective heuristic that labels points with sliding labels in real time. Our heuristic proceeds incrementally; it tries to insert one label at a time, possibly pushing away labels that have already been placed. To quickly predict which labels have to be pushed away, we use a geometric data structure that partitions screen space. With this data structure we were able to double the frame rate when rendering maps with many labels.

In addition to labeling points, we have considered labeling linear objects (such as streets) in interactive maps where the user can pan, zoom (within a certain scale range), and rotate continuously [H]. Our labels contain text (such as street names). They are embedded into the objects they label; see Fig. 4. This means that the labels follow the curvature of the objects, they do not move with respect to the map background, but they scale in order to maintain constant size on the screen. Our objective is to label as many streets as possible and to select label positions of high quality while forbidding labels to overlap at street crossings. We have developed a simple but effective algorithm that takes curvature and crossings into account and produces aesthetical labelings. On average over all interaction types, our implementation reaches interactive frame rates of more than 85 frames per second.

Finally, we have considered labeling linear objects in interactive 3D maps that provide a perspective view of a scene [J]. This type of map is most common in navigation systems that guide a user along a pre-computed route. We assume that the future route of the user is known or well predicted within the currently visible part of the map and suggest that the streets belonging to that route shall be brought to the user's attention. To do so, we annotate those streets with a special type of labels, so-called billboards; see Fig. 5. A billboard is a rectangle holding the annotation text, is oriented towards the user, placed at some distance above the midpoint of the street to be labeled, and connected to that point by a vertical line segment, the *leader*. Our goal is to maintain an overlap-free labeling that reacts to changes of the view in real time. To this end, we dynamically vary the lengths of the leaders. In order to achieve that labels move smoothly, we do not strictly forbid label–label overlaps. We have developed a force-directed algorithm that applies forces to labels to cause overlapping labels to repel each other, while keeping leaders as close to their desired length as possible. On real-world data, with a realistic number of labels, we obtain frame rates of more than 400 frames per second, while drastically reducing the total overlapped area per frame, compared to an algorithm with fixed leader lengths.

To get an impression of how our labeling algorithms perform in a dynamic setting, we suggest watching the videos that we generated with our algorithms for point-feature labeling[3], embedded street labeling[4], and billboard labeling[5].

3 Prospective Applications of Project Results and Future Research

Since we have successfully published our results and have received attention in the public media (in the daily newspaper of Würzburg, Mainpost[6] and in the online platform Scinexx[7]), we expect a significant impact of our project on future developments. Specifically, we have published two articles in highly cited journals, namely the article [A] in the *IEEE Transactions in Visualization and Computer Graphics* (TVCG; 2011 impact factor: 2.215) and the article [B] in the *International Journal of Geographical Information Science* (IJGIS; 2014 impact factor: 1.655). Furthermore, we have published seven papers in peer-reviewed proceedings of conferences or workshops, all of which have become part of the *ACM Digital Library* [D, E, H], the *Springer Lecture Notes in Computer Science* [C, F], or the *Springer Lecture Notes in Geoinformation and Cartography* [G, J].

Our scientific findings have contributed significantly to the rapidly growing market for web cartography, location-based services, and navigation systems. New algorithmic solutions for small-screen cartography and visualization systems allowing for interaction with spatial information are urgently needed. In practice, heuristic algorithms are often developed whose objectives are not formalized well. For a company maximizing its short-term or mid-term profit, this may be reasonable. For us, in contrast, an important task was to define models of cartographic requirements and quality and, thereby, to allow different solutions to be compared and measured against each other. After defining these models, we were able—in some cases—to solve them with efficient and exact algorithms. On the other hand, we found some problems to be NP-hard. For those problems, we compared results of heuristic methods with exact solutions obtained by mathematical programming.

[3] `http://www1.pub.informatik.uni-wuerzburg.de/pub/schwartges/dynapointlab`

[4] `http://www1.pub.informatik.uni-wuerzburg.de/pub/schwartges/dynalinelab`

[5] `http://www1.pub.informatik.uni-wuerzburg.de/pub/schwartges/dynaroutelab`

[6] `http://www.mainpost.de/art735,6697657`

[7] `http://www.scinexx.de/wissen-aktuell-14447-2012-02-15.html`

Not only do we see clear applications of our models and algorithms for the specific cartographic problems that we have studied, we also see the need to promote and establish general approaches of rigorous problem modeling and solving in cartography and GIScience. With this aim, we have contributed a chapter [3] on the use of combinatorial optimization for spatial analysis to a new Handbook on Geodesy, which is about to appear. We have also approached a more application-oriented readership with the idea to use optimization for map generalization [2].

Beyond cartography, our results have applications in information visualization and, in particular, in graph and network visualization. An important open question emerging from our project is how dynamic but reasonably stable (i.e., not abruptly changing) visualizations of networks can be produced.

4 Expected Economic Usage

A direct economic usage of the project results was not expected. Together with our project partner Bosch Car Multimedia, however, we filed a patent [4].

5 Project Team and Research Partners

The project team at the University of Würzburg consisted of Jan-Henrik Haunert and Alexander Wolff, the two applicants; Thomas C. van Dijk, who filled a post-doctoral position that was established with project funds; and Nadine Schwartges, whose PhD studies at the University of Würzburg were funded for two years (50%) by Bosch Car Multimedia, Hildesheim, and for one year with a scholarship of a Bavarian qualification initiative for *Women in Research and Teaching* (Programm der Staatsregierung "Chancengleichheit für Frauen in Forschung und Lehre"). In September 2013, Jan-Henrik Haunert accepted a professorship for geoinformatics at the University of Osnabrück, from where he continued to lead the project. The project team was supported by several student assistants (most notably by Ben Morgan, Leon Sering and Dennis Zwiebler) who contributed to articles [J, A, G] that were published in the course of the project. One publication has resulted from a collaboration with the group of Bettina Speckmann [F], TU Eindhoven, and one from a collaboration with Markus Chimani [E], University of Osnabrück.

6 Qualification of Young Scientists

In June 2013, Jan-Henrik Haunert obtained his habilitation at the University of Würzburg, which qualified him for his professorship. Nadine Schwartges received her PhD in April 2015. Thomas C. van Dijk started to pursue his habilitation. He continues his work on algorithms for geographic information systems and gives lectures on that topic.

References (For refs. [A]–[J], see separate list of "project publications".)

[1] K. Been, M. Nöllenburg, S.-H. Poon, and A. Wolff. Optimizing active ranges for consistent dynamic map labeling. In *Proc. 24th Annu. ACM Sympos. Comput. Geom. (SoCG'08)*, pages 10–19, 2008.

[2] J.-H. Haunert. Map generalisation – quality benchmarks through optimisation. *GIM International*, 25(1):31–33, 2011.

[3] J.-H. Haunert and A. Wolff. Räumliche Analyse durch kombinatorische Optimierung. In W. Freeden and R. Rummel, editors, *Handbuch der Geodäsie*. Springer-Verlag, 2015. To appear.

[4] C. Neumann and N. Schwartges. Verfahren und Vorrichtung zum Bestimmen einer Darstellung von Beschriftungselementen einer digitalen Karte sowie Verfahren und Vorrichtung zum Anzeigen einer digitalen Karte, 2014. DE Patent App. DE201,210,213,307.

[5] M. van Kreveld, T. Strijk, and A. Wolff. Point labeling with sliding labels. *Comput. Geom. Theory Appl.*, 13(1):21–47, 1999.

[6] D. Yamamoto, S. Ozeki, and N. Takahashi. Focus+glue+context: An improved fisheye approach for web map services. In *Proc. 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM-GIS'09)*, pages 101–110, 2009.

Summary

We have developed algorithms for continuous map generalization, variable-scale label placement, and map-layout generation. Before we developed an algorithm, we always formalized the respective task as an optimization problem. Therefore, the project has not only resulted in efficient exact algorithms and heuristics for application in real-time visualization systems, but also in models of quality for a *variable-scale map*, with which we refer to a sequence or continuum of maps that covers an interval of scales or to a single map that displays different regions at different scales.

Our algorithms for map-layout generation enlarge user-selected focus regions in a road network map without changing the map's size or the displayed part of the network. They reduce the scale of non-focus regions and introduce some distortion to compensate for the scale differences in the map. The distortion is much lower, however, than with existing fish-eye projections for similar tasks. Our first algorithm relies on convex quadratic programming and produces realistically-sized maps within a few seconds. A better performance, which allows for applications in real-time systems, is obtained by a second algorithm based on least-squares optimization. For the second algorithm we had to relax some of the originally hard constraints of our model. Still, it yields solutions of high cartographic quality.

Regions that are demagnified by our method obviously tend to appear cluttered in the output map. For this reason, but also to generate maps supporting continuous zooming, we have developed methods that select a subset of all road segments for a map. We have developed a probabilistic model of how humans navigate a city, which allows us to quantify the importance of every road segment. The aim is to select segments of preferably high importance while satisfying certain constraints about the output map, for example, to ensure that the network appears not too cluttered and is connected. We have studied different models of connectivity and, in some cases, have been able to prove that the selection problem is NP-hard. For several selection problems of practical relevance, however, we have found efficient exact algorithms. For the NP-hard cases we have developed approaches by integer programming, approximation algorithms, and heuristics.

With respect to map labeling, we have considered the problem of placing text for points and linear objects in interactive maps that allow users to pan, zoom, and rotate continuously. Our general aim is to label as many objects as possible (on average over time in a dynamic setting) such that no two labels overlap. For this task, we have concentrated on the development of real-time-capable heuristics. Finally, we have developed a method for labeling linear objects in interactive 3D maps that provide a perspective view of a scene. Labels are displayed as billboards. Based on a model of forces, the method avoids overlaps between labels. At the same time, the method ensures that labels are close to their corresponding objects.

To conclude, we have contributed many new algorithms to the rapidly growing market for web cartography, location-based services, and navigation systems. An important open question emerging from our project is how dynamic but reasonably stable (i.e., not abruptly changing) visualizations of networks can be produced.