

Exploring Semi-Automatic Map Labeling

Fabian Klute

Algorithms and Complexity Group,
TU Wien, Vienna, Austria
fklute@ac.tuwien.ac.at

Guangping Li

Algorithms and Complexity Group,
TU Wien, Vienna, Austria
guangping@ac.tuwien.ac.at

Raphael Löffler

Algorithms and Complexity Group,
TU Wien, Vienna, Austria
raphael.loeffler@aon.at

Martin Nöllenburg

Algorithms and Complexity Group,
TU Wien, Vienna, Austria
noellenburg@ac.tuwien.ac.at

Manuela Schmidt

Research Group Cartography, TU
Wien, Vienna, Austria
manuela.schmidt@tuwien.ac.at

ABSTRACT

Label placement in maps is a very challenging task that is critical for the overall map quality. Most previous work focused on designing and implementing fully automatic solutions, but the resulting visual and aesthetic quality has not reached the same level of sophistication that skilled human cartographers achieve. We investigate a different strategy that combines the strengths of humans and algorithms. In our proposed labeling method, first an initial labeling is computed that has many well-placed labels but is not claiming to be perfect. Instead it serves as a starting point for an expert user who can then interactively and locally modify the labeling where necessary. In an iterative human-in-the-loop process alternating between user modifications and local algorithmic updates and refinements the labeling can be tuned to the user's needs.

We demonstrate our approach by performing different possible modification steps in a sample workflow with a prototypical interactive labeling editor. Further, we report computational performance results from a simulation experiment in QGIS, which investigates the differences between exact and heuristic algorithms for semi-automatic map labeling. To that end, we compare several alternatives for recomputing the labeling after local modifications and updates, as a major ingredient for an interactive labeling editor.

CCS CONCEPTS

• **Human-centered computing** → **Geographic visualization**; *Human computer interaction (HCI)*; Empirical studies in visualization; • **Theory of computation** → *Design and analysis of algorithms*; *Computational geometry*.

KEYWORDS

human-in-the-loop map labeling, algorithm experiments

ACM Reference Format:

Fabian Klute, Guangping Li, Raphael Löffler, Martin Nöllenburg, and Manuela Schmidt. 2019. Exploring Semi-Automatic Map Labeling. In *27th ACM*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGSPATIAL '19, November 5–8, 2019, Chicago, IL, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6909-1/19/11...\$15.00

<https://doi.org/10.1145/3347146.3359359>

SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '19), November 5–8, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3347146.3359359>

1 INTRODUCTION

Label placement is an important step in map production, both manual and automatic, and it can require up to 50 percent of the total map production time for manually created maps [33]. Imhof's 1975 statement "Good form and placing of type make the good map. Poor, sloppy, amateurish type placement is irresponsible; it spoils even the best image and impedes reading." [14] has not lost its validity until today. Yet, with more and more automation in cartography and fully digital map production, one can argue that the quality of label placement has not improved or even diminished compared to skilled, but tedious manual label placement [23]. While practical label placement algorithms are typically very fast and can compute overlap-free positions of thousands of labels within seconds, the resulting maps usually do not meet the highest quality standards but must be carefully post-processed in tedious work by human cartographers. There is a lack of algorithmic support of human involvement in an automated labeling workflow. Ideally, a responsive labeling algorithm would be able to react on human interaction and respect any added constraints, e.g., choosing an alternative position of a label or changing its font size, by locally updating the solution while keeping the already placed labels as stable as possible. Such a semi-automatic map labeling tool would allow for a much more comfortable and intelligent human-in-the-loop label placement process in digital map production, where neither human alone can deal with the full data complexity, nor machine alone can deal with the fine-tuning and optimization of mathematically somewhat ill-defined map aesthetics. Instead, the tool should combine the computational power and mathematical rigor of geometric labeling algorithms with the expertise and sense of aesthetics of experienced domain experts in cartography.

In this paper we present a prototype for a such a human-in-the-loop label placement approach that supports first applying selected labeling algorithms for placing an initial set of non-overlapping labels for point features that, in our case, maximizes an objective function counting the (weighted) number of labeled features. To proceed from this initial solution our prototype implements several editing and interaction tools for modifying the labeling according to the needs of an expert user, e.g., changing the visibility or position of a label, or its size, shape, and weight. Upon each of those modifications, a second algorithm is re-optimizing and refining

the labeling by taking into account both the user input and the existing solution so as to satisfy the new constraints and maximize the stability with respect to the previous solution. That is, a labeling is computed that contains as many as possible of the previously displayed labels and at the same time resolves any new label conflicts resulting from the user input. Our prototype is designed to be flexible as to which algorithm to actually use for computing initial solutions and iterative updates. Secondly, we take an algorithm engineering perspective on the map labeling problem and perform an experimental simulation study in the GIS software QGIS. The aim is to investigate differences between and suitability of heuristic and exact algorithms (including those provided by QGIS itself) for the envisioned interactive labeling workflow, which requires frequent recomputation of labelings after local modifications.

Related Work. Based on general cartographic guidelines [14, 32], the first algorithmic solutions to the label placement problem have been studied in the cartographic literature in the 1970s and 1980s [13, 33, 34]. In the early 1990s the problem has been introduced as a geometric independent set problem to the computational geometry community [10, 20], where it was recognized as an important application challenge in the computational geometry task force report [5]. It was quickly shown that almost all variants of label placement and label selection problems are NP-hard [10, 20]. Therefore, researchers focused on special cases, approximation algorithms and heuristics for label number maximization or label size maximization problems, predominantly for point features, e.g. [6, 24, 28–30]; for surveys and general introductions see, e.g., [15, 24, 31]. More recent works introduced advanced multi-criteria optimization models [12, 21, 27] that can express more accurately several established cartographic principles, but still with the aim of a full automation of the map labeling process. While progress is made by incorporating more comprehensive cartographic rules for label placement, none of the above approaches includes decisions made by human experts – other than setting preferences, parameters, and priorities in the different scoring functions that control a single optimization run of the respective algorithm. A notable exception is the UserHints framework [7], where human interaction was integrated into solving the label number maximization problem in a fixed-position point labeling setting. In that system, two heuristic methods were implemented as labeling algorithms, and hence the evaluation could not assess the deviation from optimal solutions with respect to the objective function. Moreover, the authors did not consider the stability of the labeling under user interaction. Beyond the label placement problem, interactive optimization [22] and human-guided search [16] are of course techniques that are of general interest and more broadly applicable.

Popular GIS software like *Mapbox*, *ArcGIS Pro*, or *QGIS* also provide labeling algorithms. *Mapbox* allows customized label modifications with data conditions, but no manual selection or drag-and-drop placement. The *ArcGIS Pro* documentation states “Label positions are generated automatically. Labels are not selectable. You cannot edit the display properties of individual labels.” To allow for manual adjustment, labels can be converted to annotations. If the labels are stored in a database, the annotations can be feature-linked, i.e., the annotations update in case features are added or changed. However, after converting to annotations, all positioning

needs to be done manually. Other proprietary software developers like *1Spatial* and *Lorienne* advertise features to modify labels in a more advanced manner. Though, there seems to be no focus on how to better integrate the automatic labeling process into a more interactive approach, especially from an algorithmic perspective. Finally, in *QGIS 3* some advanced labeling tools were introduced. For example, it is possible to manually drag and reposition labels; other labels will be re-placed accordingly. Labels, which were manually edited, can be highlighted and reversed to their default position. While this is a good example that demonstrates the awareness and practical need for semi-automatic labeling solutions, prior to this paper no experimental studies on the performance of different labeling algorithms under interactive editing have been published that evaluate such an approach and guide further development in *QGIS* and other systems.

Paper Structure. In Section 2 we introduce our model for semi-automatic map labeling, which combines the classic point-feature label placement with a dynamic update problem. Section 3 introduces our prototype tool and describes a sample map labeling workflow using interactive modifications by a cartographic expert. Finally, Section 4 describes our simulation experiment in *QGIS* to analyze the performance of several labeling algorithms.

2 LABELING MODEL

In this paper we restrict our attention to the *point feature label placement* (PFLP) problem, which is defined as follows. Let P be a set of n feature points in \mathbb{R}^2 . For each point $p \in P$ we are given a finite set L_p of *label candidates*, where each label candidate $\ell \in L_p$ is represented by the bounding box R_ℓ of the feature name placed at a particular position. While in general the label candidates in L_p can be arbitrary label positions, we focus on the standard 4- and 8-position models, where either one of the corners of R_ℓ coincides with p (4-position model) or one of the corners or midpoints of the edges of R_ℓ coincides with p (8-position model). Let $\mathcal{L} = \bigcup_{p \in P} L_p$ be the union of all label candidates.

We say that two label candidates ℓ and ℓ' are in *conflict*, if the two rectangles R_ℓ and $R_{\ell'}$ intersect. Since the names of two conflicting label candidates would overlap, the goal in map labeling is to find a *conflict-free* solution set $S \subseteq \mathcal{L}$ of label candidates. In particular, we require that any two label candidates of the same point p are in conflict so that each point receives at most one label. To optimize the labeling we define a quality function $w: \mathcal{L} \rightarrow \mathbb{R}^+$ that assigns a weight to each label candidate. Then in its basic form, which we implemented for our prototype tool and is also used equivalently in *QGIS* through a cost model for non-labeled features, the PFLP optimization problem is defined as follows.

PROBLEM 1 (PFLP). *Given a set P of n points in \mathbb{R}^2 with a set of label candidates \mathcal{L} and a weight function w , find a conflict-free set $S \subseteq \mathcal{L}$ of label candidates for P such that the weight $W(S) = \sum_{\ell \in S} w(\ell)$ is maximized.*

The simplest weight function is $w \equiv 1$, which just counts the number of selected labels. But more advanced weight functions, defined for single labels, pairs of labels, or even larger subsets, in order to model various cartographic principles are possible [12, 21, 27]. We want to explore user modifications in our semi-automatic

labeling process, which, for example, change the set \mathcal{L} of label candidates to a set \mathcal{L}' or the weight function w to a function w' . Therefore we define the following PFLP update problem.

PROBLEM 2 (PFLP-UPDATE). *Given a set P of n points in \mathbb{R}^2 with a set of label candidates \mathcal{L}' , a weight function w' , and a previous labeling \mathcal{S} , compute a conflict-free solution $\mathcal{S}' \subseteq \mathcal{L}'$ that maximizes the number $|\mathcal{S} \cap \mathcal{S}'|$ of stable labels as well as the weight $W(\mathcal{S}') = \sum_{\ell \in \mathcal{S}'} w'(\ell)$.*

We note that there may be a trade-off between the stability of the new solution \mathcal{S}' and its weight $W(\mathcal{S}')$ that can be adjusted by the user. For solving the two labeling problems algorithmically, we model the conflicts and the label candidates as a weighted *conflict graph* $G = (V, E)$, where the vertex set $V = \mathcal{L}$ consists of all label candidates and the edge set E consists of all pairs of conflicting label candidates. Then, in graph-theoretic terms, an optimal labeling corresponds to a *maximum weight independent set* in G , i.e., a subset $V' \subseteq V$ of vertices such that no two vertices $u, v \in V'$ are adjacent and the weight $\sum_{v \in V'} w(v)$ is maximum. The problem of computing maximum independent sets in graphs is a classic NP-hard problem [11], even in its unweighted form.

2.1 Data

For our computational experiments and the evaluation of the prototype we extracted points-of-interest data from the [OpenStreetMap](#) (OSM) project, filtered it for certain categories or properties, and then stored the name and location of the remaining points as ESRI shapefiles or in a simple JSON file format to be read by our tool. Using data from the OSM project guarantees that the feature distribution is realistic, even if the particular data sets are simplified and not cartographically sound use cases.

We compiled five different datasets, all with unit weights, whose properties are summarized in Tab. 1. The first one consists of all mountain peaks above 2,499 meters in the mountain range “Hohe Tauern” in Austria. It consists of 1,278 homogeneously distributed, natural features and is on average less dense than the other four data sets. We compiled this dataset to use it in the sample workflow with a zoom level set to 12, see Section 3.2. The number of conflicts and conflicts per feature hence are measured in the applied 4-position model of the prototype.

The other four datasets are man-made features taken from Vienna, the province of Lower Austria, and Austria itself. Here we use QGIS 12-position model to measure the conflicts. In case of Austria we filtered all places marked as town or city, resulting in a total of 301 features with 16.58 conflicts per feature. For Lower Austria we took all villages, towns, and cities inside the state boundaries of Lower Austria and Vienna. This resulted in a set with 2,260 features with 22.6 conflicts per feature. These settlement features are irregularly distributed according to the physical geography of an alpine country. Finally we considered all bus, tram, and subway stops inside the state boundaries of Vienna. These features are more dense in the city center and thin out towards the periphery. One set, “Vienna Train”, consists of 1,001 tram and subway stops with 23.27 conflicts per feature. The last data set we call “Vienna Bus/Train”. It adds also all bus stops inside Vienna to the tram and subway stops. These are 3,939 features and 33.66 conflicts per feature, hence it is by far the most densely packed set of features. Note that for all data

Table 1: Test data sets and their properties.

data set	features	conflicts	conflicts/feat
Mountain Peaks	1,278	15,416	3.01
Austria	301	4,991	16.58
Lower Austria	2,260	47,269	22.60
Vienna Train	1,001	23,294	23.27
Vienna Bus/Train	3,939	132,571	33.66

sets the number of conflicts includes the conflicts between label candidates of the same feature, which yields a lower bound of 3 or 11 conflicts per feature in the 4- or 12-position model, respectively.

3 SEMI-AUTOMATIC LABELING PROTOTYPE

We developed a prototype tool that includes four labeling algorithms and provides a proof-of-concept GUI to test user interaction with the system. For the implementation of the backend, especially the algorithms, we used *Java 8* in conjunction with the *Play Framework* (version 2.6) and the *JGraphT* library (version 1.0.1). For displaying the labels we built a web interface using the Javascript libraries *Leaflet* (version 1.0.3) and *D3* (version 4.9.1).

Our application is a one-page design, i.e., the page does not require any reloads while working with it. The user interface, shown in Fig. 1, consists of the large map area in the middle of the window. Here the current labeling is displayed on background map tiles. The labels are drawn as white rectangles with black text and are initially attached to the features according to a 4- or 8-position model. All feature points are displayed as filled circles. A blue circle indicates a labeled feature, while a red circle corresponds to an unlabeled one. On the left-hand side we find a sidebar, containing most of the input controls, e.g., for loading a file or choosing the algorithms. Above the map area, there are three toggle buttons. Two of them manipulate what is shown on the map, while the rightmost button toggles if the labels we just modified are kept as fixed labels or if they can be deleted from the solution by the next update. By clicking on a label, its background color changes to green as seen in Fig. 1 and the label properties area pops up as a sidebar on the right-hand side of the window. Here the current values are displayed, e.g., font size, weight, or margin, and the user can modify them accordingly. Shifting a label candidate position by drag-and-drop is also possible. Concerning user feedback, the application informs the user via small message boxes on the lower right corner. Longer computations will block the user interface, and a progress bar pops up on the top.

3.1 Workflow and User Interaction

Starting from an unlabeled map, the first step is to import a set of data points to be labeled. Then one of the implemented algorithms (see Section 4.1) is selected to compute an initial solution that can subsequently be modified. The core of the proposed user-centered labeling process consists of a number of implemented modification tools, that were designed according to the needs of a human cartographer and are summarized in Tab. 2. Most of the modifications have a direct effect in the corresponding conflict graph, e.g., deletion or insertion of conflict edges, deletion of vertices, forced selection of

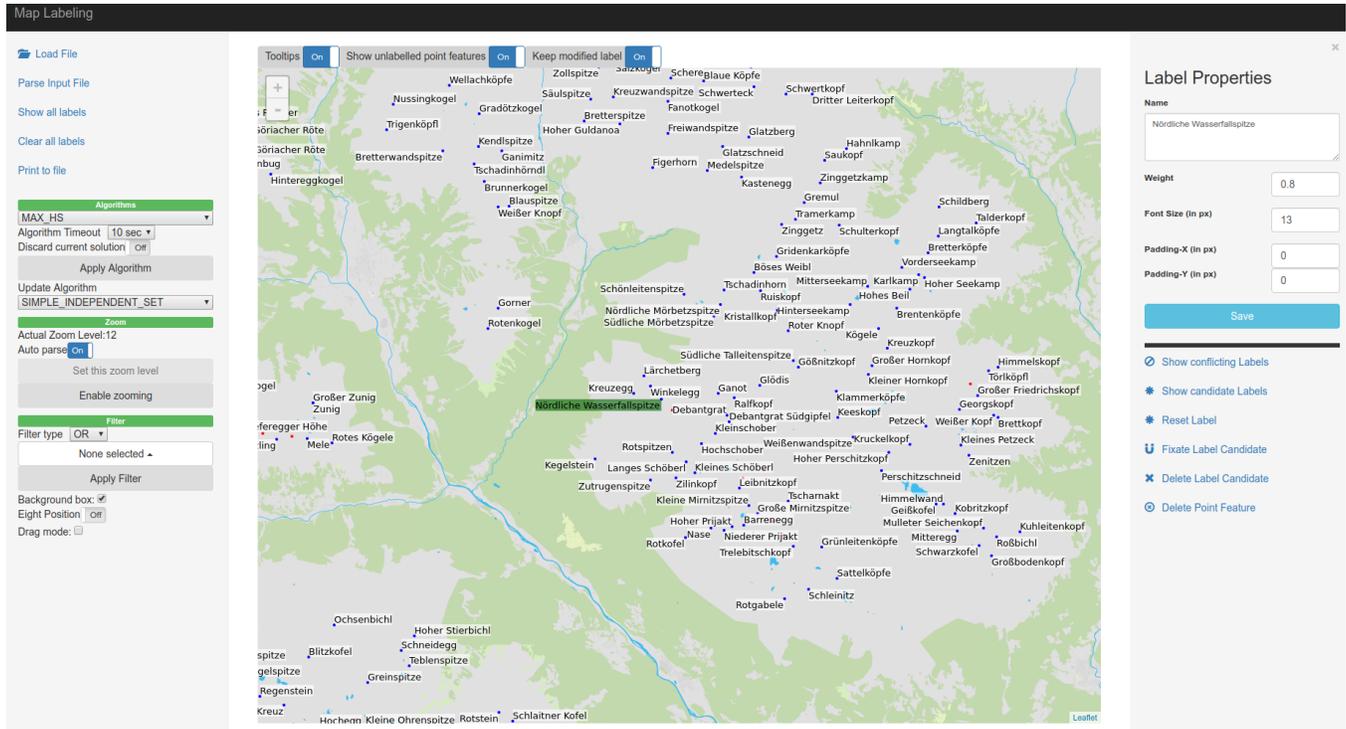


Figure 1: The graphical user interface of the developed prototype.

Table 2: Implemented user modifications.

text	solution	label
change font size	delete point features	drag label in the map
change text	delete label candidates	change padding
add line breaks	fixate label candidates	toggle label box visibility
	change cand. weight	

vertices, or changes of vertex weights. Lastly, the user can select the algorithm for solving the PFLP-Update problem following the modifications; it can be the same algorithm as for computing the initial labeling or a different one, where aspects such as computation speed, stability, and optimality must be taken into account. Our simulation experiments in Section 4.3 provide some empirical guidance for choosing an update algorithm.

3.2 Realistic Sample Workflow

In this subsection, we describe an example of a realistic map labeling workflow that has been performed by a cartographer using our prototype tool; the process was protocolled and video-captured. In the first step, the point features from the Mountain Peaks dataset are added to the map. After loading and parsing the features, the user zooms and pans to the area of interest. Once the desired zoom level is set, an initial labeling can be produced by any of the four algorithms outlined in Section 4.1. Here we used the exact algorithm MHS. After it found a solution, the labeled features are shown with

a blue symbol and the corresponding label candidate; unlabeled features are indicated by a red symbol.

Usually, a cartographer would now try to group and prioritize features according to their attributes first. In this sample scenario, though, we are treating all features with the same importance and are only using visual cues to manually refine the results of the automatic labeling. Reasons for necessary manual adaptations include cases, in which the visual connection between symbols and labels is ambiguous or in which the label does not correspond well with underlying map features (e.g., labels covering lakes). While some cases could be avoided by implementing more sophisticated placement rules (e.g., assigning label and feature weights), other cases seem difficult to automate – either due to their subjective nature or due to the complexity of demands.

After identifying an improvable label, the cartographic expert would click on the label, which activates the view of four (or eight) alternative label positions. The preferred label candidate can be fixated by marking it and activating “Fixate Label Candidate” in the modification sidebar. Alternatively, in “drag mode”, the label can be dragged to the preferred position. Further label edits include the option to change the font size and weight as well as the label name, e.g., using abbreviations or stacking long label names. Based on the manually selected label, the positions of surrounding labels are re-calculated with the selected update algorithm.

In our test scenario, about 15 improvable labels (less than 10% of the initially placed labels) were identified. Not all of them had to be changed manually, since updates of neighboring labels and subsequent re-calculations fixed some of the issues. Of course, the

instant updates sometimes also resulted in new improvable labels. All in all, the option to select and adjust individual labels while maintaining all automatic labeling functionality was considered highly useful by the expert to speed up the label optimization process in comparison to current cartographic workflows without customized algorithmic support for interactive labeling.

4 ALGORITHMS AND EXPERIMENTS IN QGIS

Quantum GIS (QGIS) is one of the most popular open-source geographic information system application in recent years. By using the PAL [8] local search labeling library in its labeling engine, QGIS provides an automated layer labeling. The label placement can be customized by choosing labeling algorithms, adjusting styling, etc. The newly updated labeling toolbar in QGIS 3 provides more tools for manual label placement, which include changing individual label attributes and moving labels in the map. This indicates that the QGIS developers see a clear demand for support of interactive labeling, too. Our goal in this section is to investigate the effectiveness and stability of several algorithms (described in Section 4.1) applicable to solve problems PFLP and PFLP-Update, as well as the algorithms from PAL (see Section 4.2) that are included in QGIS. To this end we integrated our algorithms in the labeling engine of QGIS and performed several simulation experiments. The results are reported in Section 4.3.

4.1 Labeling Algorithms

We selected four labeling algorithms as representatives of existing labeling approaches, three increasingly sophisticated heuristics and one exact algorithm. The simplest algorithm (called GREEDY) is an easy-to-implement and fast greedy approach with little to no optimization. The second algorithm (MIS) aims to obtain a good approximation of the maximum independent set in the conflict graph. Given that a lot of graph libraries provide independent set algorithms it is also very easy to apply by taking an existing implementation. The algorithm KAMIS [19] computes large independent sets by a combination of several advanced algorithmic techniques such as graph partitioning and kernelization. Finally we designed a new exact algorithm (MHS) based on a MAXSAT formulation. To solve our MAXSAT formulation we use the solver *MaxHS*, which also gives the name to our algorithm. *MaxHS* is a freely available solver which ranks highly in the competition held at the annual SAT conferences.

All algorithms except KAMIS support weighted labels, but this does not directly affect our experiments as we initially use unweighted labels. Maintaining a previous solution can be done with GREEDY, MIS and MHS. In the case of GREEDY we just keep the conflict-free subset of the old solution completely and try to extend it. For MIS and MHS we adjust the weights of the old labels so that they have higher priority to be included in the solution. Since KAMIS currently does not support weighted instances, prioritizing labels of a previous solution via weights is not possible with KAMIS. From a theoretical perspective, specialized algorithms that respect a previous solution have been considered recently for the *dynamic independent set* problem [4]. So far, this was not investigated in the light of the PFLP problem.

GREEDY. Algorithm GREEDY computes a maximal independent set $D \subseteq V$ in the conflict graph $G = (V, E)$ using a greedy approach. It starts by picking a random vertex $u \in V$ of maximum weight and adds it to D . All neighbors of u and u itself are then marked. Next we pick an unmarked vertex $v \in V$ of maximum weight, add it to D and mark v and its neighbors. We repeat this until no unmarked vertex remains. The constructed set D is a maximal independent set of G , i.e., it cannot be extended. But there is no guarantee that it is a maximum weight independent set. The selected set of label candidates corresponds to the set D . Our implementation is also able to take as input an independent set $D' \subseteq V$ and guarantee that for the new solution D we have $D' \subseteq D$.

MIS. Like GREEDY, the algorithm MIS builds a maximal independent set, but in contrast to GREEDY it tries to find a good approximation to a maximum weight independent set in G . Such approaches are well known in the map labeling literature [1, 25]. One way to find a large maximal independent set is to find a small minimal *vertex cover* $D \subset V$ of the conflict graph G . A vertex cover D has the property that for every edge $(u, v) \in E$ at least one of the two vertices u, v is contained in D . Consequently, the set $D' = V \setminus D$ is an independent set, as by definition of D no two vertices in D' can be neighbors in G . In our implementation we use the greedy vertex cover heuristic as implemented in the graph library JGraphT. For a vertex u let $\text{deg}(u)$ be its degree in the conflict graph, then in each step the algorithm picks the vertex with the smallest ratio of $w(u)/\text{deg}(u)$, adds it to D , and removes u together with all edges incident to u .

KAMIS. The third algorithm KAMIS is based on the maximum independent set solver framework *KaMIS* [19]. By combining kernelization, local search, an evolutionary algorithm, graph partitioning and other techniques, this advanced maximum independent solver can very successfully find large independent sets in huge sparse graphs. There are three algorithm components in this framework. Firstly, to create initial solutions, it uses a swap-based local search called *ARW* [19]. After a greedy insertion of vertices with small residual degrees in the independent set, the local search applies (1, 2)-swaps. In particular, if two vertices have only one common neighbor in the current solution, the optimization search then inserts these two vertices and removes its neighbor to increase the size of the independent set locally. The second algorithm component is the evolutionary algorithm *EvoMIS* [18]. It employs a multi-way partitioning on the graph to make the exchange of sub-solutions in graph components possible. After the recombination, newly generated offsprings are locally optimized using swaps. The third component is a kernelization technique [19] with both exact and inexact kernels. Exact kernelization applies reduction rules to decrease the problem size without affecting solution quality. For example, isolated vertices can always be added to the independent set. Besides exact kernelization, inexact kernelization rules are used to reduce the search space in the optimization phase. Intuitively, we choose vertices with very small degree which are in the current candidate independent set and fixate them. Consequently their neighborhood is then deleted, which reduces the size of the remaining instance.

The whole algorithm works as follows. The first phase is to reduce the problem size by exact kernelization. After an initialization phase using *ARW*, the evolutionary procedure *EvoMIS* combines

sub-solutions of components and optimizes the newly generated solutions. Inexact kernelization of the fit solutions make further exact kernelization possible, and the algorithm repeats the whole process from there.

MHS. Finally we introduce MHS, an approach based on satisfiability testing of Boolean formulas. A Boolean formula ϕ in conjunctive normal form consists of a logical conjunction of clauses $\phi = c_1 \wedge \dots \wedge c_m$. Each clause c_i is a logical disjunction of one or more Boolean variables x_1, \dots, x_n and their negations $\neg x_1, \dots, \neg x_n$; these are called *literals*. Each *variable* can take the value *true* or *false*. For a particular *truth assignment* φ of all variables, one can evaluate the truth values of all clauses. The formula ϕ is *true* if every clause is *true*. A truth assignment for which ϕ evaluates to *true* is called a *satisfying assignment*.

The satisfiability problem (SAT) asks for the existence of a satisfying assignment, given a Boolean formula ϕ , and is one of the fundamental NP-complete problems [11]. If every clause consists of at most two literals, the restricted problem is known as 2-SAT and can be solved in polynomial time [17]. Formann and Wagner [10] modeled the labeling problem for a 2-position model as a 2-SAT formula, which allowed them to test in polynomial time whether a conflict-free labeling exists that assigns a label candidate to every point.

Let P be a set of point features, \mathcal{L} the label candidates and $G = (V, E)$ the conflict graph of P . We first construct a 2-SAT formula ϕ that guarantees the solution set to be conflict free. Let Λ be the set of variables of ϕ and Γ the set of clauses. To build our formula, we introduce for each vertex $u \in V$ a Boolean variable $\lambda_u \in \Lambda$, and for each conflict edge $(u, v) \in E$ the clause $\gamma(u, v) = (\neg \lambda_u \vee \neg \lambda_v) \in \Gamma$. We derive a solution set $\mathcal{S} \subseteq \mathcal{L}$ from a truth assignment φ of ϕ , by choosing a label candidate $\ell \in \mathcal{L}$ to be added to \mathcal{S} if and only if for the corresponding vertex $u \in V$ the variable λ_u is *true* in φ .

Now \mathcal{S} is conflict-free if and only if φ is a satisfying assignment of ϕ . This can be seen by remembering that for every conflicting pair of labels we find an edge $(u, v) \in E$ between the vertices $u, v \in V$ corresponding to the conflicting labels. For such an edge we introduced the clause $\gamma(u, v)$ which states that it is not possible to set $\lambda_u = \lambda_v = \textit{true}$ in any satisfying assignment. While such an assignment leads to conflict-free label sets it does not maximize the number of labels. In particular the assignment φ mapping all variables to *false* is a satisfying one, but the solution set \mathcal{S} resulting from φ is empty.

In the related problem MAXSAT we do not ask for a satisfying assignment of a given formula, but instead for an assignment that maximizes the number of clauses that evaluate to *true*. MAXSAT, as well as MAX-2-SAT, are well known to be NP-complete [11]. To model the PFPL problem as a MAX-2-SAT formula we add a literal for every $\lambda_u \in \Lambda$ as a separate clause $\gamma(u) = \lambda_u \in \Gamma$. However, if we simply maximize the number of satisfied clauses, we have no guarantee whether some of the clauses $\gamma(u, v)$ evaluate to *false* – which would imply that two conflicting label candidates can be selected. Hence we use a version of the MAXSAT problem, the *partial maximum satisfiability problem* (PMAX-SAT) [3]. In PMAX-SAT the set of clauses is partitioned into a set of *hard clauses* Γ_H and a set of *soft clauses* Γ_S . We must find a truth assignment φ such that any clause $\gamma \in \Gamma_H$ evaluates to *true*, while the number of clauses

$\gamma' \in \Gamma_S$ that evaluate to *true* is maximized. In our case we define the clauses $\gamma(u, v) \in \Gamma$ as hard clauses and the clauses $\gamma(u) \in \Gamma$ as soft clauses. Now for any solution the clauses expressing a conflict have to be satisfied. We note that a similar formulation has been used to model the maximum clique problem [9], which is equivalent to the maximum independent set problem in the complement graph.

As a next step we extend our model to also solve the PFLP-Update problem. In *weighted* PMAX-SAT a soft clause $\gamma \in \Gamma_S$ can be assigned a weight $w(\gamma) \in \mathbb{R}^+$, which can be seen as a penalty for falsifying γ . The aim is to find an assignment φ of ϕ that satisfies all the clauses in Γ_H and minimizes the sum of penalties of the unsatisfied clauses in Γ_S . Let $\ell \in \mathcal{L}$ be a label candidate, w its weight, and $\mathcal{S}' \subset \mathcal{L}$ the previous solution. We have to specify the weights for the soft clauses. For every $\gamma \in \Gamma_S$ that corresponds to a label candidate from \mathcal{S}' we set $w(\gamma) = w + \varepsilon$ for some parameter $\varepsilon \geq 0$, which gives higher priority to selecting a previously displayed label candidate over a previously unused one. In our implementation we used $w = 1$ and $\varepsilon = 1$. If we want to strictly prioritize maximum solutions over solutions with fewer labels (but possibly more from \mathcal{S}') we need to compute a suitable ε . Let $k = |\mathcal{S}'|$ and assume all label weights are uniform $w \equiv 1$. Then we can choose $0 < \varepsilon < 1/k$.

4.2 QGIS labeling algorithms

For our experiments we further selected three representative algorithms from the labeling engine of QGIS. The FALP algorithm is the simplest greedy algorithm, which is also used to build the initial solution in other optimization algorithms. The algorithm CHAIN is a local search algorithm with chained neighborhood moves. The most advanced algorithm in QGIS is POPMUSIC (called `pop_tabu_chain`). It combines the general paradigm of POPMUSIC [26] with CHAIN and tabu search.

FALP. The fast procedure FALP builds an unweighted solution in two steps: First, all label candidates are ordered by increasing number of conflicts with other label candidates. By using the ordered position mode in QGIS, the ordering respects the common preference of candidate positions [14] for tie-breaking. Then the label candidates are visited in this order. Once a label position is chosen, other candidates of its feature and other label candidates in conflict with this label are removed from the ordering, and all conflict numbers of remaining label positions are updated and re-sorted. The implementation in QGIS is incremental. It maintains the conflict number of each candidate and updates the values accordingly.

CHAIN. The local search approach CHAIN is chaining multiple modifications in order to escape from local minima. After building an initial solution with FALP, improvements of the current labeling are searched by applying a sequence of chained modifications. A chain of modification is formed as follows. First, a (seed) feature is chosen randomly, and it will be labeled (unlabeled) or modified in the current solution. This move may create new overlaps and therefore may lead to a chain of modifications of the current solution. The chain search temporarily applies these new changes and the local search process continues. The chain search will stop after a specified number of modifications is reached. Once the chain is stopped, the best solution reached along the chain will replace the current one.

Table 3: Average running times (in ms) for computing an initial solution (init) or an update (upd), number of initially labeled features (init f), and number of labeled features on average over all runs (feat). Best values are printed in bold.

Algorithm	Austria				Lower Austria				Vienna Train				Vienna Train/Bus			
	init	upd	init f	feat	init	upd	init f	feat	init	upd	init f	feat	init	upd	init f	feat
CHAIN	13	13	179	179	140	146	497	534	38	39	282	288	263	266	351	370
FALP	5	4	176	177	80	79	494	529	19	19	280	285	192	194	349	367
GREEDY	0.9	0.9	152	159	13	14	388	414	3	3	242	248	18	19	277	283
MIS	14	13	171	175	261	250	457	509	56	57	272	279	701	668	325	347
KAMIS	123	121	182	182	581	863	523	560	248	249	289	295	57746	6272	373	391
MHS	181	164	182	179	11821	4257	523	537	734	561	289	288	20937	7212	373	373
POPMUSIC	1599	1448	181	181	14684	15723	504	541	3915	3961	282	289	17820	18020	355	373

POPMUSIC. The POPMUSIC algorithm is an implementation of the POPMUSIC framework presented by Taillard and Voss [26]. The basic idea applied to the PFPL problem can be summarized as follows: We begin with an initial solution generated by FALP. Every feature is now considered as a sub-part of the instance. We also maintain a list L of features which is initially empty. The algorithm iteratively considers features not in L and executes the CHAIN algorithm starting from this feature with a bound on the maximum number of labels CHAIN is allowed to consider in its search. If this procedure leads to an improvement we remove all labels considered by CHAIN from L . In case the run did not further improve the solution we add the feature from which we started CHAIN to L . POPMUSIC terminates once L contains all features. Additionally, PAL implements a tabu search paradigm for the CHAIN procedure to avoid cyclically visiting the same solutions when optimizing from some feature. For further details see [2] and [8].

4.3 Simulation Experiments

Our experiment focuses on two aspects. The first aspect asks how viable our advanced heuristic KAMIS and the exact approach MHS are compared to the heuristics implemented in QGIS. The second aspect asks which combination of algorithms performs best in an interactive scenario as defined by Problem 2. All runs of the algorithms produce valid, overlap-free labelings. Our main interest in this experiment lies in the computational performance, the objective value of the solution, and the stability of the updated solutions. We are not claiming that the resulting labelings are competitive with manually labeled maps, as the applied modifications in this simulation study are generated by a random process and not by a cartographic expert. Yet the algorithmic performance is expected to be comparable for modifications made purposefully to improve a labeling.

In this section, we focus on the data sets for Lower Austria and the denser transport network of Vienna, since they are the largest and most dense label sets, respectively. For the two other data sets our findings are similar.

All experiments were run on a standard desktop computer with an eight-core Intel i7-860 CPU clocked at 2.8 GHz and 8 GB RAM, and running Archlinux kernel version 5.1.4. We compiled KAMIS, as well as our code together with QGIS 3.7, using gcc version 8.3 and cmake 3.14.4. The compile flags for QGIS and KAMIS were set

to Release. Every test was performed 50 times for each combination of initial and update algorithm.

Initial Solutions. To compare our algorithms with the existing labeling algorithms as found in QGIS, we considered four of the datasets presented in Section 2.1. Our expectation would be that the greedy heuristics CHAIN, FALP, GREEDY, and MIS have clearly faster running times compared to POPMUSIC, MHS, and KAMIS, while the number of labeled features is expected to be higher for the latter algorithms. We present our findings in Table 3. Our intuitive expectations get largely confirmed for the greedy heuristics. GREEDY is clearly the fastest algorithm, but also leaves a large gap in the solution quality, even compared to the other greedy heuristics. In terms of running times we have to consider the overhead of building the conflict graph for MIS. For the Austria data set this took around 11ms, for Lower Austria 100ms, and for the two Vienna data sets it took 49ms for the sparse and 279ms for the dense one. If this overhead could be removed, e.g., in an update run we would not need to recompute the full graph each time, MIS would run about as fast as FALP. In terms of solution quality we also see that MIS stays behind CHAIN and FALP. Comparing FALP and CHAIN we see that FALP runs roughly twice as fast as CHAIN, but the solutions are slightly worse. In the end FALP and CHAIN provide similar results in terms of quality and speed.

For the remaining three algorithms we see that POPMUSIC can provide pretty good solutions in terms of number of labeled features, even compared with the optimal approach MHS and clearly better ones compared to the greedy heuristics. For KAMIS it turned out that its initial solution was in fact always an optimal solution equivalent to MHS. Looking at the running times, we see that surprisingly MHS and KAMIS outperform POPMUSIC by one to two orders of magnitude for all but one data set. The data set on which POPMUSIC surpasses KAMIS and MHS is also the most dense one, namely the full transport network of Vienna. Likely this behavior is due to the fact that especially KAMIS uses small cuts in the graphs very well, where a cut $E' \subseteq E$ in a graph $G = (V, E)$ is a set of edges such that G decomposes into two or more independent parts after removal of E' . Intuitively the sparser instances also have smaller cuts, while the denser instances lead to more highly connected conflict graphs with large cuts.

Modifications. For the experiments with modifications we consider five runs of labeling algorithms. In the first run we produce an initial labeling. As seen in the previous paragraph, KAMIS, MHS,

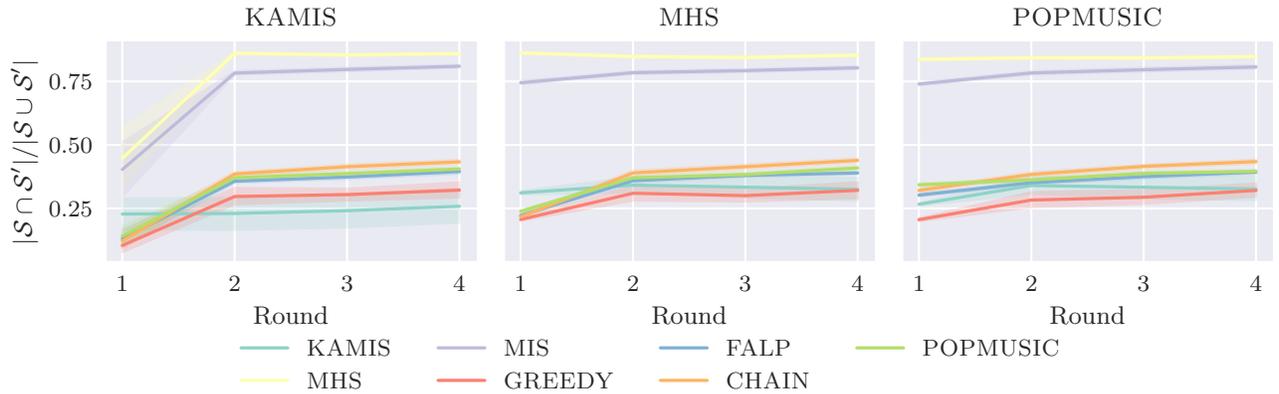


Figure 2: Stability over four rounds of modifications in the Lower Austria data set.

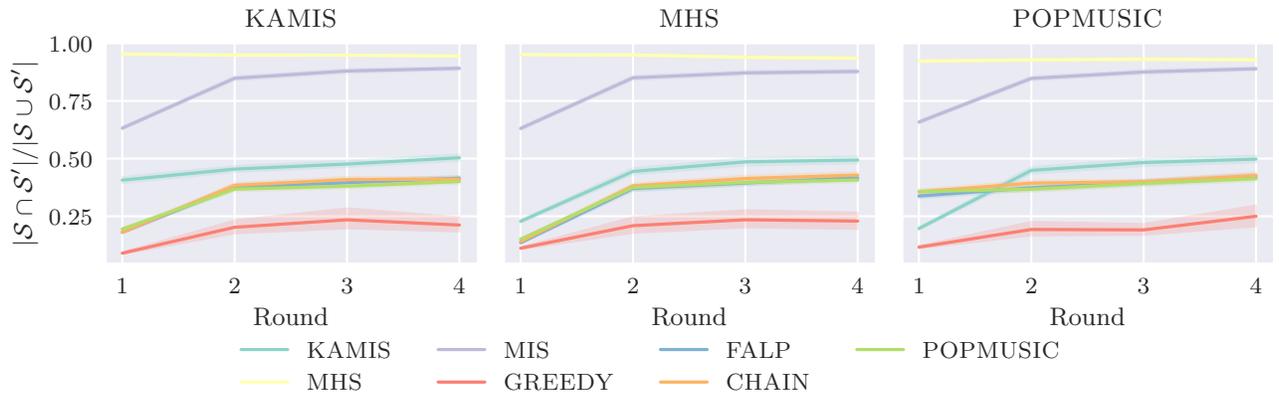


Figure 3: Stability over four rounds of modifications in the Vienna data set of bus, tram, and subway stops.

and POPMUSIC are the natural candidates for this, since they provide the best solution while still running sufficiently fast. To simulate manual modifications of labels in the current solution, we choose after each run labels from the complete set of labels uniformly at random. In each round we will change the font size to 20 for one percent of the labels, and for another three percent set it to 5. Note that our initial font size is set to 10. Furthermore, we pick one percent of the labels and delete them from further consideration. From the perspective of the conflict graph, these modifications cover all relevant changes: insertion and deletion of conflict edges and deletion of vertices. It should further be noted that shrinking takes precedence over enlarging, i.e., if we decide to enlarge a label, then shrink it and again enlarge it, it will still remain at a font size of 5. Considering Figures 4 and 5 this explains immediately why on average the number of labeled features increases after each round. For computing the updates we are especially interested in MHS, GREEDY, and MIS, as these algorithms consider weights and thus can optimize stability of the previous solution. To measure stability of an updated labeling we compute the following ratio. Let \mathcal{S} be a labeling and \mathcal{S}' a labeling of the modified instance, then a stable solution keeps $|\mathcal{S} \cap \mathcal{S}'| / |\mathcal{S} \cup \mathcal{S}'|$ as close as possible to 1.0.

Figures 2 and 3 show our findings in regard to this measure. We used POPMUSIC, KAMIS, and MHS as initial algorithms and all algorithms as update algorithms, to explore if by random chance the

solutions stay stable as well. Clearly we see this is not the case when comparing with MHS and MIS, as for nearly all possible starting algorithms they manage to keep the ratio at a value of 0.8 for the transport network of Vienna and 0.97 for the lower Austria data set. In general it seems that MIS is worse at maintaining stability, while MHS is computing more stable labeling. Interestingly, for KAMIS as a starting algorithm and the dense data set of Vienna MHS heavily changes the solution after the first modification. This might point to a difference in the solution between KAMIS and MHS. Without further investigation it is hard to determine the exact cause of this change. If it turns out that MHS and KAMIS find very different high quality solutions this could be interesting to cartographers independently of our interactive approach. The very naive approach of GREEDY seems in fact to be too simple to keep the solution stable.

Turning to the running times and number of labeled features, consider Table 3 again. Comparing MHS and MIS we see that MHS is at a disadvantage in terms of running time, but keeps up more labels on average and as just seen produces more stable results. On the other hand MIS is very fast and for not too dense input sets seems to be fine in terms of stability.

Since these are average values over all rounds of modification we finally consider boxplots of the percentage of labeled features for KAMIS, MHS, and POPMUSIC as starting algorithms and KAMIS,

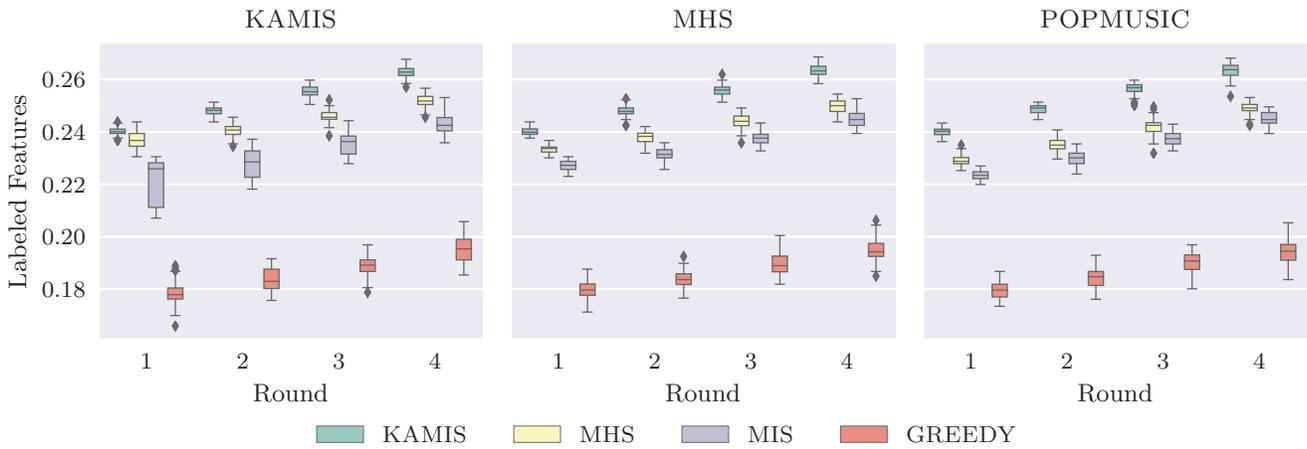


Figure 4: Labeled features in the Lower Austria data set.

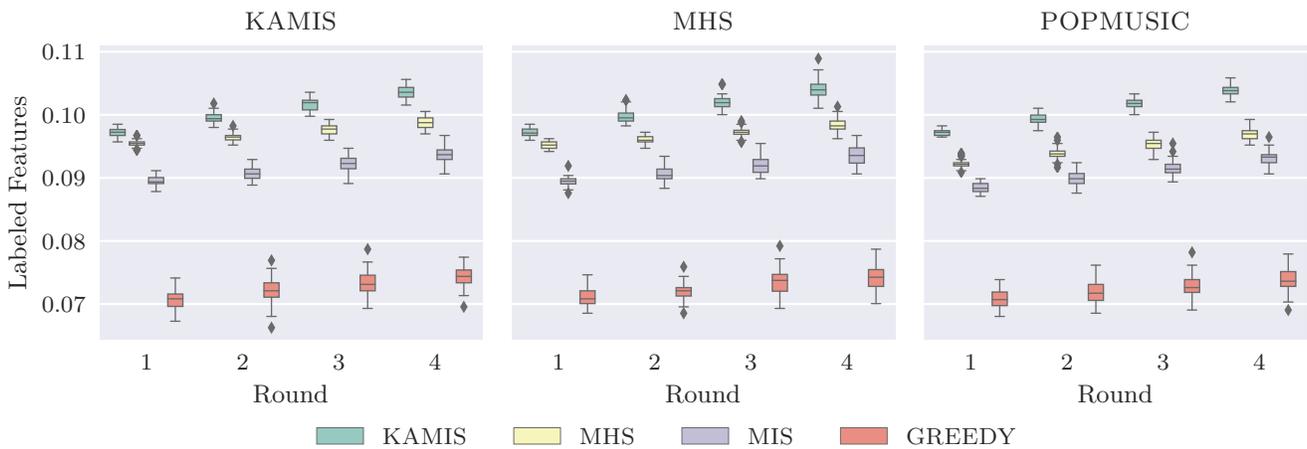


Figure 5: Labeled features in the Lower Austria data set.

MHS, GREEDY, and MIS as update algorithms. Our findings are shown in Figures 4 and 5. Remember that KAMIS does not try to keep the solution stable, hence in these plots it should rather be seen as a potential optimum for the number of labeled features, if the algorithms were allowed to disregard the old solution. As expected we find that MHS comes closest to this theoretical optimum. MIS is not far behind, while GREEDY is several percent behind. Correlating with the drop in stability we also find that for KAMIS as a starting algorithm and MHS and MIS as update algorithms we get a large variation in the solution quality on the lower Austria data set.

Discussion. In conclusion we can say that it is more than viable to use exact or near-exact approaches for map labeling when compared with sophisticated heuristics. Especially the cut based maximum independent set approach of KAMIS is promising in general, as it not only improves the computation times by one to two orders of magnitude over the most sophisticated heuristic POPMUSIC in QGIS, but also performs consistently better in the optimization goal. A combination of KAMIS and a heuristic like POPMUSIC seems like a very useful future approach. In such a combined idea one could use cuts to find dense areas which are solved by POPMUSIC, while

the sparser areas are handled by KAMIS. Also the SAT approach of MHS has potential as SAT-solvers continue to get better and the used solver MaxHS surely will be surpassed in the coming years.

In terms of greedy heuristics we saw that GREEDY and MIS both have their disadvantages. GREEDY is in the end a too trivial approach to realize good solutions, while MIS suffers from the overhead of building the conflict graph. In a labeling framework though, where the conflict graph would be kept readily available as modifications are made, clever independent set heuristics may become a viable approach.

In terms of modifications and interactive labeling we saw that it is possible to use just simple weights to keep the solution stable. MHS and MIS are both suitable algorithms to handle updates, with MIS clearly being the faster approach. A viable combination of the two could be to run MHS only every fifth or tenth iteration of modification. Also it is likely that future maximum independent set frameworks will support weights, bringing the approach of KAMIS also to the table. In case these weighted frameworks exhibit a similar running time to KAMIS they likely would provide the

best of both worlds, fast running times and stable, high quality solutions.

5 CONCLUSION

We have presented a prototype tool for supporting semi-automatic map labeling workflows together with first experimental results on four possible labeling algorithms and on how to combine them for computing initial labelings and updates after user modifications. This is underlined by the QGIS project recently implementing similar ideas into their label placement engine. An immediate consequence from our experiments is that targeted and fast update algorithms that aim for label stability are needed to support interactive modifications. In our future investigations we aim to develop algorithms specifically tailored to optimize the proposed stability criteria. As a first step, we plan to investigate fast dynamic weighted independent set heuristics, where the weights better model the stability of the labeling. Moreover, the labeling algorithms must take into account more advanced and accurate cartographic quality constraints, also in combination with the stability criteria. Ultimately, this may be fully integrated, e.g., into the QGIS project in order to find its way into practical map production. At that point, meaningful formal user studies with GIS experts on usability as well as final labeling quality and required interaction efforts are needed to validate whether human-in-the-loop optimization for label placement is meeting its expectations.

ACKNOWLEDGMENTS

This work is supported by the Austrian Science Fund (FWF) under Grant P 31119.

REFERENCES

- [1] Pankaj K. Agarwal, Marc van Kreveld, and Subhash Suri. 1998. Label placement by maximum independent set in rectangles. *Computational Geometry* 11, 3 (1998), 209–218. [https://doi.org/10.1016/S0925-7721\(98\)00028-5](https://doi.org/10.1016/S0925-7721(98)00028-5)
- [2] Adriana C.F. Alvim and Éric D. Taillard. 2009. POPMUSIC for the point feature label placement problem. *European Journal of Operational Research* 192, 2 (2009), 396–413. <https://doi.org/10.1016/j.ejor.2007.10.002>
- [3] Carlos Ansótegui, Maria Luisa Bonet, and Jordi Levy. 2010. A New Algorithm for Weighted Partial MaxSAT. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI'10)*. AAAI Press, 3–8. <http://dl.acm.org/citation.cfm?id=2898607.2898608>
- [4] Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. 2018. Fully dynamic maximal independent set with sublinear update time. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25–29, 2018*. 815–826. <https://doi.org/10.1145/3188745.3188922>
- [5] Bernard Chazelle. 1996. *Application Challenges to Computational Geometry: CG Impact Task Force Report*. Technical Report TR-521-96. Princeton University.
- [6] Jon Christensen, Joe Marks, and Stuart Shieber. 1995. An Empirical Study of Algorithms for Point-feature Label Placement. *ACM Trans. Graph.* 14, 3 (July 1995), 203–232. <https://doi.org/10.1145/212332.212334>
- [7] Hugo A.D. do Nascimento and Peter Eades. 2008. User hints for map labeling. *Journal of Visual Languages & Computing* 19, 1 (2008), 39–74. <https://doi.org/10.1016/j.jvlc.2006.03.004>
- [8] Olivier Ertz, Maxence Laurent, Daniel Rappo, Abson Sae-Tang, and Eric Taillard. 2009. PAL-A cartographic labelling library. *Position IT July* (2009), 56–61.
- [9] Zhiwen Fang, Chu-Min Li, Kan Qiao, Xu Feng, and Ke Xu. 2014. Solving Maximum Weight Clique Using Maximum Satisfiability Reasoning. In *European Conference on Artificial Intelligence (ECAI'14)*. IOS Press, 303–308. <https://doi.org/10.3233/978-1-61499-419-0-303>
- [10] Michael Formann and Frank Wagner. 1991. A packing problem with applications to lettering of maps. In *Proc. 7th (SoCG'91)*. ACM, 281–288. <https://doi.org/10.1145/109648.109680>
- [11] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York.
- [12] Jan-Henrik Haunert and Alexander Wolff. 2017. Beyond Maximum Independent Set: An Extended Integer Programming Formulation for Point Labeling. *Int. J. Geo-Information* 6, 11 (2017), 342:1–342:20. <https://doi.org/10.3390/ijgi6110342>
- [13] Stephen A. Hirsch. 1982. An Algorithm for Automatic Name Placement Around Point Data. *The American Cartographer* 9, 1 (1982), 5–17. <https://doi.org/10.1559/152304082783948367>
- [14] Eduard Imhof. 1975. Positioning Names on Maps. *The American Cartographer* 2, 2 (1975), 128–144. <https://doi.org/10.1559/152304075784313304>
- [15] Konstantinos G. Kakoulis and Ioannis G. Tollis. 2013. Labeling Algorithms. In *Handbook of Graph Drawing and Visualization*, Roberto Tamassia (Ed.). CRC Press, Chapter 15, 489–515.
- [16] Gunnar W. Klau, Neal Lesh, Joe Marks, and Michael Mitzenmacher. 2010. Human-guided search. *J. Heuristics* 16, 3 (2010), 289–310. <https://doi.org/10.1007/s10732-009-9107-5>
- [17] Melven R Krom. 1967. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Mathematical Logic Quarterly* 13, 1-2 (1967), 15–20.
- [18] Sebastian Lamm, Peter Sanders, and Christian Schulz. 2015. Graph Partitioning for Independent Sets. In *Experimental Algorithms - 14th International Symposium, SEA 2015, Paris, France, June 29 - July 1, 2015, Proceedings*. 68–81. https://doi.org/10.1007/978-3-319-20086-6_6
- [19] Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F. Werneck. 2017. Finding near-optimal independent sets at scale. *Journal of Heuristics* 23, 4 (2017), 207–229. <https://doi.org/10.1007/s10732-017-9337-x>
- [20] Joe Marks and Stuart Shieber. 1991. *The Computational Complexity of Cartographic Label Placement*. Technical Report. Harvard University. <http://www.eecs.harvard.edu/~shieber/Biblio/Papers/label.pdf>
- [21] Maxim A. Rylov and Andreas W. Reimer. 2014. A Comprehensive Multi-Criteria Model for High Cartographic Quality Point-Feature Label Placement. *Cartographica* 49, 1 (2014), 52–68. <https://doi.org/10.3138/carto.49.1.2137>
- [22] Stacey D. Scott, Neal Lesh, and Gunnar W. Klau. 2002. Investigating human-computer optimization. In *Human Factors in Computing Systems (CHI'02)*. ACM, 155–162. <https://doi.org/10.1145/503376.503405>
- [23] Seth Stevenson. 2012. The Greatest Paper Map of the United States You'll Ever See. *Slate* (Jan. 2012). http://www.slate.com/articles/arts/culturebox/2012/01/the_best_american_wall_map_david_imus_the_essential_geography_of_the_united_states_of_america.html
- [24] Tycho Strijk. 2001. *Geometric Algorithms for Cartographic Label Placement*. Ph.D. Dissertation. Universiteit Utrecht.
- [25] Tycho Strijk, Bram Verweij, and Karen Aardal. 2000. *Algorithms for Maximum Independent Set Applied to Map Labelling*. Technical Report. Utrecht University.
- [26] Éric D Taillard and Stefan Voss. 2002. POPMUSIC—Partial optimization meta-heuristic under special intensification conditions. In *Essays and surveys in meta-heuristics*. Springer, 613–629. https://doi.org/10.1007/978-1-4615-1507-4_27
- [27] Steven van Dijk, Marc van Kreveld, Tycho Strijk, and Alexander Wolff. 2002. Towards an evaluation of quality for names placement methods. *Int. J. Geographical Information Science* 16, 7 (2002), 641–661. <https://doi.org/10.1080/13658810210138742>
- [28] Marc van Kreveld, Tycho Strijk, and Alexander Wolff. 1998. Point Set Labeling with Sliding Labels. In *Proceedings of the Fourteenth Annual Symposium on Computational Geometry (SCG '98)*. ACM, 337–346. <https://doi.org/10.1145/276884.276922>
- [29] Frank Wagner and Alexander Wolff. 1997. A practical map labeling algorithm. *Comput. Geom. Theory Appl.* 7 (1997), 387–404. [https://doi.org/10.1016/S0925-7721\(96\)00007-7](https://doi.org/10.1016/S0925-7721(96)00007-7)
- [30] Frank Wagner, Alexander Wolff, V. Kapoor, and Tycho Strijk. 2001. Three Rules Suffice for Good Label Placement. *Algorithmica* 30 (2001), 334–349. <https://doi.org/10.1007/s00453-001-0009-7>
- [31] Alexander Wolff and Tycho Strijk. [n.d.]. The Map Labeling Bibliography. <http://i11www.iti.kit.edu/~awolff/map-labeling/bibliography/>
- [32] Clifford H. Wood. 2000. A Descriptive and Illustrated Guide for Type Placement on Small Scale Maps. *Cartographic J.* 37, 1 (2000), 5–18. <https://doi.org/10.1179/caj.2000.37.1.5>
- [33] Pinhas Yoeli. 1972. The Logic of Automated Map Lettering. *The Cartographic Journal* 9, 2 (1972), 99–108. <https://doi.org/10.1179/000870472787352505>
- [34] Steven Zoraster. 1986. Integer Programming Applied To The Map Label Placement Problem. *Cartographica* 23, 3 (1986), 16–27. <https://doi.org/10.3138/9258-63QL-3988-110F>