

A General Cartographic Labelling Algorithm

SHAWN EDMONDSON, *Baker Hill Corporation*

JON CHRISTENSEN, *Painted Word Incorporated*

JOE MARKS, *MERL—A Mitsubishi Electric Research Laboratory*

and STUART M. SHIEBER, *Harvard University*

Abstract Some apparently powerful algorithms for automatic label placement on maps use heuristics that capture considerable cartographic expertise but are hampered by provably inefficient methods of search and optimization. On the other hand, no approach to label placement that is based on an efficient optimization technique has been applied to the production of general cartographic maps—those with labelled point, line, and area features—and shown to generate labellings of acceptable quality. We present an algorithm for label placement that achieves the twin goals of practical efficiency and high labelling quality by combining simple cartographic heuristics with effective stochastic optimization techniques.

Introduction

MANY apparently compelling techniques for automatic label placement use sophisticated heuristics for capturing cartographic knowledge, but, as noted by Zoraster (1991), also use inferior optimization strategies for finding good trade-offs between the variety of competing concerns involved in typical labelling problems. These techniques use procedural methods or “if-then” production rules to represent cartographic knowledge about good label-placement practice, and a variant of depth-first search to explore different labellings (Doerschler and Freeman 1992; Ebinger and Goulette 1990; Jones 1989).

However, depth-first search is now known to be a markedly inferior technique for finding near-optimal labellings in the set of all possible labellings. More powerful optimization strategies have been proposed, including:

- *Physical relaxation* While remaining tethered to the symbols they tag, labels are moved smoothly in response to virtual forces generated by label-label and label-symbol overlaps (Feigenbaum 1994; Hirsch 1982).
- *Zero-one integer programming* Scores associated with each label’s candidate positions (a discrete set) are refined iteratively to better reflect the relative desirability of the different positions (Zoraster 1986; Zoraster 1990).
- *Gradient descent* A randomly generated labelling is improved monotonically by considering all alternative positions for each label (again chosen from a discrete set) and making the single label move that most improves the quality of the whole labelling (Christensen, Marks, and Shieber 1993).
- *Simulated annealing* A generalization of gradient descent in which single label moves that worsen the quality of the labelling are performed occasionally in the hope of avoiding bad labellings that happen to be locally optimal (Christensen, Marks, and Shieber 1993).

These techniques all outperform depth-first search by a wide margin (Christensen, Marks, and Shieber 1995).

So why have these supposedly superior techniques not been adopted widely? In general, implementations of the better optimization strategies have not dealt with the full range of cartographic features (that is, point, line, and area features), nor have they incorporated sufficient cartographic knowledge about text placement. For these reasons, the maps they produce have not been persuasive. However, this shortcoming is not intrinsic, at least for some of the aforementioned techniques. In this paper we show how detailed cartographic knowledge and powerful optimization can be combined effectively. The resulting algorithm is general (it can label point, line, and area features), efficient (it can label dense, page-sized maps in seconds on a PC), and effective (the resulting labellings are visually appealing and consistent with good cartographic practice). Furthermore, compared to existing algorithms, it is concisely stated and easily implemented.

The key to our approach is a careful separation of the cartographic knowledge needed to recognize a good labelling from the optimization procedure required to find one. Our notion of cartographic knowledge is a procedure for computing an absolute numeric score that is properly indicative of a labelling’s quality; no other properties of this scoring function (such as continuity, dif-

Biographical and contact information needed: about 50 words

Revised manuscript received December, 1996.

ferentiability, or definability in terms of production rules) are assumed. We treat the scoring function as an arbitrary objective function to be optimized, and apply a powerful optimization procedure to find its near-optimal values. These values correspond to near-optimal labellings.

To establish the viability of this concept, we describe a particular optimization procedure and a particular scoring function that can be used together in the framework we propose above. Neither are necessarily optimal: it was not our intent to provide the ultimate solution for automatic map lettering, a long-term goal that is probably best left to professional cartographers, not computer scientists. However, we believe that we have identified the framework within which the ultimate solution probably lies. As evidence for this claim, we present a simple system prototype that generates visually compelling labellings for dense page-sized maps in under ten seconds on a personal computer.

We begin by describing our technical approach in sufficient detail to enable our work to be replicated in its entirety. We then present some sample maps labelled using an implementation of our method, and conclude with an analysis of present and future work on label placement.

Technical Approach

Our approach is predicated on a division of the label-placement task into three essentially independent sub-tasks. They are:

- *Candidate-position generation* Given a point, line, or area feature, identify a set of candidate locations for its label. A *labelling* is then a set of label positions, one drawn from each feature's set of candidate positions.
- *Position evaluation* Given a labelling, efficiently compute a score that indicates its quality with respect both to the position of labels relative to the tagged symbol, and to spatial contention among the label and other features and feature labels.
- *Position selection* Given a set of candidate label positions for each map feature, choose one label position from each set so that the overall quality of the labelling, as determined by the evaluation method, is as high as possible.

We discuss these subtasks in reverse order in the next three sections. Because of the simplicity of the position-selection method, we discuss it briefly; full details have been presented elsewhere (Christensen, Marks, and Shieber 1995). We then consider the problems of label-position evaluation and generation, which form the core of our contribution in this paper.

Position Selection

Given a set of generated candidate positions for each label and an overall evaluation function, selecting positions for all the labels so that the evaluation function is globally minimized¹ is an optimization task. Although many different methods have been proposed for this task, we favour a method based on *simulated annealing* (Kirkpatrick, Ge-

latt Jr., and Vecchi 1983):

- For each feature, place its label randomly in any one of the candidate positions for that feature.
- Initialize a "temperature" T to an initial high value.
- Repeat the following steps until the rate of improvement falls below a given threshold:
 - Decrease T according to an annealing schedule
 - Pick a feature randomly and move its label to a new position randomly chosen from that feature's set of candidate positions.
 - Compute ΔE , the change in the overall labelling's evaluation caused by repositioning the label.
 - If the new labelling is worse, undo the label repositioning with probability $P = 1.0 - \exp(-\Delta E/T)$.

The algorithm is almost completely specified in the outline above. The remaining details concern the initial temperature (chosen so that $P=2/3$ when $\Delta E=1$), the annealing schedule (periodically the temperature decreases by 10 percent), and the termination condition (evaluation of $5n$ consecutive repositionings with no changes made, where n is the number of labelled features). Full details concerning these issues are provided elsewhere (Christensen, Marks, and Shieber 1993; Christensen, Marks, and Shieber 1995). It suffices here to note that the parameters are such that 100,000 computations of ΔE are typically required for convergence on problems involving up to 1,500 labelled features. Thus the overall efficiency of the label-placement algorithm depends crucially on efficient computation of ΔE .

This optimization method has several advantages. First, it is effective at finding near-optimal solutions to label-placement problems. Christensen, Marks, and Shieber (1995) demonstrate through exhaustive empirical testing of point-feature label-placement algorithms that simulated annealing dominates all previously reported practical algorithms for this problem. The results were robust over both actual cartographic data and randomly generated data, and thus provide some foundation for the use of artificial data in the present work. Second, the method is efficient; labellings are generated in an amount of time that is reasonable and competitive with other algorithms.² Finally, it makes no assumptions about the evaluation function beyond the efficient computability of the numeric

1 Our statement of the evaluation function associates lower values with better labellings, so that the optimization problem is one of minimization. This choice is, of course, arbitrary.

2 Both of these findings are remarkable. Simulated annealing is usually thought of as a very inefficient algorithm of "first resort" that is used only until a better approach is devised. However, for point-feature label placement (the purest form of the problem, and the most useful for comparing algorithms) it takes only a few seconds to label dense maps with hundreds of point features, and it produces better labellings than any algorithm developed during the 15 years or so for which automatic label placement has been studied seriously (Christensen, Marks, and Shieber 1995b).

difference ΔE . This last advantage suggests that the method should generalize beyond the point-feature label-placement problem to which it was initially applied. The remainder of this paper can be seen as a verification of this suggestion.

Position Evaluation

We now turn to the task of evaluating label positions. In order to compare various labelling solutions, we require the ability to compute a single numeric score, E , that indicates the quality of a labelling. Furthermore, the change in score, ΔE , that results from repositioning a single label must be efficiently computable. In our algorithm, each label's contribution to E is computed as a weighted sum of simple *metrics*, which are described below.

For each metric, we define an *ideal case* and a *borderline case*. In the ideal case, the candidate position is ideally located as far as the given metric is concerned; in the borderline case, the position is poor, but barely acceptable. Metrics will be designed to yield a value of 0.0 for ideal cases, 1.0 for borderline cases, and higher values for objectionable cases.

First, we consider metrics that quantify spatial crowding and overlap. Next, we examine positioning metrics that depend on the spatial relationship between a label and the feature it tags. We conclude with a description of how the various metrics are combined into a single evaluation function.

OVERLAP METRICS

Feature Overlap

In order to discourage label-symbol overlaps, we need to know the number of feature symbols (the actual circles, polylines, and polygons comprising a map's nontextual symbology) that a given label overlaps. We consider three metrics, *PointOver*, *LineOver*, and *AreaOver*, which measure overlaps with point, line, and area symbols, respectively. When there are no overlaps of a given kind, we have the ideal case. We will take the borderline case to be a single overlap. Thus, one reasonable metric is to simply count the number of overlaps. This is exactly what is done for the metric *PointOver*, which is defined as the number of overlaps between a selected label position and all point-feature symbols.

For area and line features, however, we need a more subtle metric. A label position that intersects and is parallel to the polyline of a line feature or an area border is completely unacceptable. However, a position that intersects the polyline at right angles might qualify as barely acceptable. We define *LineOver* and *AreaOver* accordingly: Let p_1 and p_2 be the points where a polyline enters and exits the label's bounding rectangle (Figure 1). Let the vector \hat{b} point in the direction of the label's baseline (east in Figure 1) and the vector $\hat{v} = (p_2 - p_1) / |p_2 - p_1|$ approximate the direction of the polyline where it intersects the label. The value of the metric (*AreaOver* or *LineOver*) is then $1 + 9 |\hat{v} \cdot \hat{b}|$.




Figure 1 Calculation of ϕ for the overlap metric *LineOver*

This metric can have value 0.0 (no intersection between label and line: ideal), or can range from 1.0 (label and line intersect and are perpendicular: barely acceptable) to 10.0 (label and line intersect and are parallel: unacceptable). In the example of Figure 1, *LineOver* is 2.56 for (a), and 9.86 for (b). These values are appropriate, since (a) is a much better placement than (b). Of course, neither is as good as a placement that does not intersect the line at all, which would yield zero for *LineOver*.

AreaOver and *LineOver* discard much detail—everything but the direction of the polyline and the label. However, they can be computed very efficiently. In experiments, a more elegant metric based on the area of intersection between the label and the polyline proved to do a slightly better job of evaluating label-line intersections, but at a high cost in efficiency.

Label Overlap

Mutually overlapping labels are clearly undesirable. As with the *PointOver* metric, we simply count the number of overlaps: we define *LabelOver* to be the number of overlaps between one selected label position and all others.

It is possible to precompute values for almost all of our metrics prior to position selection. For example, *LineOver* can be precomputed for each candidate position without any knowledge of actual selected label positions. The exception to this rule is *LabelOver*, which must be recomputed continually during position selection. Fortunately, we are able to accomplish this re-evaluation quickly by first computing and storing all pairs of intersections between candidate label positions.

PointOver, *LineOver*, *AreaOver*, and *LabelOver* together comprise all the metrics concerned with spatial contention among labels and symbols. For simplicity, we have not included metrics for other types of label-symbol overlap, e.g., mountain ranges, vegetation, etc. However, the introduction of additional overlap metrics is straightforward. We now turn to metrics that govern the positioning of a label with respect to the feature that it tags.

POINT-POSITIONING METRICS

The sole metric for point-feature labels, *PointPos*, is determined by a straightforward ranking of a discrete set of 17 candidate positions surrounding a point. The actual values used are illustrated in Figure 2, and described more fully in “candidate label positions for point features” on page 6. These values are based loosely on Imhof's well-




Figure 2 An example of 17 candidate label positions for a point feature and their PointPos values

known guidelines (Imhof 1962):

- Label positions to the right of a point are preferred to those on the left.
- Labels above a point are preferred to those below.
- The more a label's baseline is offset from a horizontal line through the centre of its associated point, the less favoured it is.
- Finally, the three positions above the point (with penalties 0.8, 0.825, and 0.875) may be shifted up slightly if a descender would otherwise overlap the point feature. In that case, an additional penalty of 0.25 is added.

LINE-POSITIONING METRICS

Before describing the various line metrics that figure in the evaluation function, we must first define some useful terms and measures, illustrated in Figure 3. A line feature




Figure 3 Terminology for line-feature labels. The swath is bounded on the left and right by the vertical dashed lines.

is represented as a long, thin polygon. However, we refer to such a polygon as a “polyline” throughout, since any computation is done with respect to the actual polyline that forms the nearest side of the line feature's polygon (whether above or below the label). The *baseline* of a label position is the line upon which the characters are drawn. However, we also take into account the shape of the text, defining the *skyline* of a label to be the union of the bounding boxes of the label's characters. The skyline extends both above and below the baseline. The *ideal distance* δ from the baseline of a label to a perfectly straight line feature below it is included in several of the line-metric formulas. The ideal distance should vary with both the

thickness of the line feature and the type size, so we set $\delta = \text{ascent} / 4 + \text{thickness} / 2$, where the ascent is the distance from the baseline to the top of capital letters of the label. The *swath* is an infinitely long strip which is perpendicular to the baseline and centred about the label. The width of the swath is 20% greater than the width of the label. The part of the polyline near the label that intersects the swath is termed the *swath line*.

We include five positioning metrics for line-feature labels in our evaluation function. The first three measure the label's relationship to the swath line: *AveDist* and *MinDist* measure the average and minimum distance from the label to the swath line; *Flatness* measures the degree of curvature of the swath line. The final two metrics measure aspects of the relationship between the label and the line feature as a whole. *Centredness* looks at the proximity of the label to the centre, measured end-to-end, of the line feature. *Aboveness* simply indicates whether the label is above or below the line feature. We discuss these metrics in more detail below.

In the following discussion we will assume we are measuring a label that is above the line feature. The rules and methods work the same if the label happens to be below the line; we simply use the other side of the skyline.

Average distance

Positioning a label at an appropriate distance from its line feature is crucial. To compute a local measure of the average distance d from a label to its associated line feature, we take the area between the swath line and the lower side of the skyline, then divide by the width of the swath. The relevant area is shaded darkly in Figure 4.




Figure 4 Computing average distance

Next we incorporate the average distance d into a useful metric with the following characteristics. First the metric should yield the optimal measure of 0.0 for an average distance δ , the ideal distance defined above. Next, let the borderline case be an average distance of either 0.0 or 2δ . We choose to let *AveDist* grow as the square of the deviation from δ , since experiment has shown this to work better than a simple linear model. Thus, we arrive

at $AveDist = (d - \delta/2)/\delta/2$. By way of example, Figure 5 shows 10 label positions with low *AveDist* values for a given line feature. The position drawn in bold has the best value of the 10.

Minimum Distance

AveDist measures the average distance from a label to its line feature in the label's neighbourhood. However, due to curvature of the line or descenders in the text, some feature segments might lie very close to the label, or actually on it, without making *AveDist* large. We need a metric to quantify this potential anomaly.




Figure 5 Label positions with good *AveDist* values

A reasonable measure is the minimum distance d' between any two points p_1 on the swath line and p_2 on the skyline. Ideally, d' is the ideal distance δ ; barely acceptable values are $d'=0$ and $d'=2\delta$. A function that obtains these values is $MinDist = (d' - \delta)/\delta/2$. However, as we will see, we do not need to consider the *MinDist* metric explicitly when evaluating a labelling, because our generation algorithm only creates candidate positions with perfect *MinDist* values.

Flatness

It is better to locate labels where the associated line feature is relatively flat. To quantify flatness, we compute the deviation of the swath line from a straight line L that is parallel to the baseline and offset the ideal distance from it. We sum up the area between the swath line and L , then divide by the width of the swath. The area sum is shaded darkly in Figure 6.




Figure 6 Computing Flatness

Let the quotient—representing the swath line's deviation from a straight line—be d'' . In the ideal case, d'' is zero; we will pick the borderline case to be when $d''=\delta$. As with *AveDist* and *AreaPos*, we let *Flatness* grow as the square




Figure 7 Label positions with good Flatness values

of d'' . Hence, $Flatness = d''^2/\delta^2$. Several label positions with low *Flatness* values for a sample line feature are shown in Figure 7. The label appears in the lowest-valued position.

Aboveness

Consider a horizontal line feature and a label running left to right: the label may be placed above or below the line. Following Imhof, we prefer labels above the line to those below. Therefore the metric *Aboveness* has a value of 0.0 if the position is above, 1.0 if below.

Centredness

Currently, our heuristics assume that each line feature will have only one label. Therefore it is important that labels lie near the centres of their associated lines. To evaluate this quality, we need a global metric *Centredness*, which we compute as follows. We find the point p on the polyline closest to the midpoint of the label baseline. Let l_1 be the distance along the polyline from one end to p ; let l_2 be the polyline's total length. Then $l = l_1/l_2$ ranges from 0.0, if p lies at an end, to $1/2$ if p lies in the middle, to 1.0, if p lies at the other end. The ideal case is when $l=1/2$; the borderline case is when l is 0.0 or 1.0. A suitable metric is thus $Centredness = |2l - 1|$.

AREA-POSITIONING METRICS

Area-specific metrics quantify the relationship between a label position and its area feature. We will assume that area features are usually large enough to accommodate their labels, which consist of closely spaced horizontal text. Given this assumption, a single, simple metric is generally all that is required for satisfactory area-feature labelling. The metric, *AreaPos*, measures the proximity of a label to its area's centroid.¹ Let c be the distance from the

1 This metric alone is completely adequate for the sample maps we consider. However, it will clearly not suffice for all area-feature shapes and all position-generation algorithms. If necessary, additional metrics that are more sensitive to the shape of the area feature, e.g., metrics that encourage alignment with an area's medial axis (Ahn and Freeman 1984), can be incorporated into the evaluation function. We omit such metrics from this discussion.

Overlap metrics	
<i>PointOver</i>	10
<i>LineOver</i>	15
<i>AreaOver</i>	10
<i>LabelOver</i>	40
Positioning Metrics	
Point positioning	
<i>PointPos</i>	1
Line positioning	
<i>AveDist</i>	1
<i>Flatness</i>	1
<i>MinDist</i>	NA
<i>Centredness</i>	3
<i>Aboveness</i>	0.25
Area positioning	
<i>AreaPos</i>	10

Table 1 Metric weights

Notice that the weight given to the *MinDist* metric is irrelevant, since the position-generation procedure for line-feature labels is guaranteed to generate only positions for which the *MinDist* value is perfect (0.0).

Position Generation

While we express most of the system’s cartographic knowledge in the evaluation function, it is clearly helpful for the generation routines to be created with some knowledge of the evaluation metrics. We have identified three qualities of a good position-generation method:

- 1 For efficiency during position selection, the number of generated positions should be relatively small. For example, in our system, the generation algorithm never provides more than 32 candidate positions for each feature.
- 2 The generation method should strike an appropriate balance between efficiency and quality: identifying only high-quality positions is too expensive and difficult, whereas naive methods tend to produce too many low-quality candidates. An appropriate balance can be achieved by keeping the position-evaluation function in mind when formulating the position-generation procedure.
- 3 The candidate positions for a given label should occupy a variety of different locations near its feature in order to give the selection algorithm the opportunity to exploit trade-offs among label positions for different features in its quest for a globally optimal labelling.

Note that it is not important that all the generated positions be of the highest quality, nor even that the best imaginable position be included. Instead, it is important that the set of positions includes a variety of locations, even if that requires inclusion of some mediocre positions. This

centre of a label position to the centroid of its area feature. In the ideal case, $c=0$. Let s be the distance from the area’s centroid to its furthest vertex. We will take the borderline case to be when $c=s$. A suitable linear function is therefore $AreaPos=c/s$.

THE OVERALL

EVALUATION FUNCTION

The overall evaluation function is a weighted sum of the metrics described above for each label on the map. Suitable values for the weights were set intuitively and refined empirically; they are summarized in Table 1.

is in contrast to other systems that first identify a single ideal position for a feature, and then choose minor perturbations of the position when attempting to resolve conflicts. Such strategies, by constraining the search space excessively, make the search for a globally optimal labelling much harder.

We describe position-generation procedures for the three fundamental types of map feature: points, lines, and areas. These types are not entirely disjoint; for example, line features can sometimes be best described as area features, as with a river on a large-scale map. However, this is merely a problem of choosing the most appropriate fundamental type as a representation, not a new problem in position generation.

CANDIDATE LABEL POSITIONS FOR POINT FEATURES

Point-feature labels are generated according to the fixed pattern illustrated in Figure 2. Conceptually, this pattern is defined by matching anchor points on the label to anchor points on the point feature.

Anchor points of a point feature are at multiples of 45 degrees on a circle centred on the feature. The radius of this circle is given by $\max(r * 1.3, r + 0.1 * f)$, where r is the radius of the point feature (or its effective radius if noncircular), and f is the width of ‘x’ in the label’s font.

Anchor points of a label are on a box formed by the top and sides of the label’s bounding box, and the baseline. Horizontally, a point is located at an end, at the midpoint, or one-third of the way from an end. Vertically, a point is aligned with one of four lines: the top of the bounding box, a line through the tops of lower-case letters (the “x-height line”), a line centred between the x-height line and the baseline, or the baseline. Not all combinations of anchor points are generated as candidate positions; the 17 positions used are shown in Figure 2.

In addition, label positions are altered to take advantage of the kerning possibilities of particular letters (the information is stored within the program). For instance, if a label is in the most favourable position, it will be slightly more to the left if the first letter is ‘V’ rather than ‘D’. As noted in the section on point-positioning metrics, candidate positions may also be altered to take descenders into account, thereby affecting the point-positioning metric values shown in the figure.

CANDIDATE LABEL POSITIONS FOR LINE FEATURES

The problem of generating positions for lines is thornier: unlike point features, line features appear in a rich variety of curves and orientations. We would prefer all candidate positions to avoid overlapping the features they tag. For point features, we can easily generate candidate positions that are guaranteed to be valid and of reasonable quality. With line features, however, we must be slightly more industrious in order to meet this proviso without paying an unreasonable price in performance.

Our approach is to apply a limited version of our overall generate/select paradigm before the general label-

selection procedure is begun. Here is a summary:

- 1 Generate many candidate positions without worrying about their validity.
- 2 Adjust them in simple ways that optimize some of the more easily computed line metrics.
- 3 Evaluate them according to all precomputable line metrics.
- 4 Cull all but the k best positions for some k .

This procedure is clearly efficient, involves no time-consuming search, yet produces a variety of relatively good positions. Here is a more detailed version of the procedure.

- 1 Generate multiple positions along the length of the line feature:
 - a) Let *start* be the point at one end of the polyline.
 - b) Let *inc* be one eighth of *width*, the width of the label.
 - c) Repeat until *start* is less than *width* from the end of the polyline, measured as distance along the line:
 - i) Find a point *end* on the polyline that is a distance *width* from *start*.
 - ii) Generate two coincident positions with baselines that run from *start* to *end*; mark one “above”, the other “below”.
 - iii) Increment *start* by *inc*.

(At this point, the algorithm will have produced a large number of candidate positions, none of which will be positioned at a good distance from the line. In fact, most will intersect the line. However, the next step corrects this problem by applying a translation.)

- 2 Adjust all generated candidate positions to achieve the ideal value for the *MinDist* metric by applying appropriate translations perpendicular to their baselines. One position of each pair is moved to a position above the line; the other, below.
- 3 If a line feature is shorter than its own label, no positions will have been generated in steps 1 and 2. If this is the case, pick a point at the centre of the line feature and generate positions as if the given point were a point feature.
- 4 Score the generated positions according to all precomputable metrics—everything but *LabelOver*.
- 5 Delete all but the k best positions.

Typically, we use $k=32$, which is a good compromise allowing a sufficient variety of positions while not overly expanding the search space. Figure 8 shows a typical line feature labelled at the best




Figure 8 Generated label positions for a line feature

generated position, as well as the skylines of the 15 next-best positions.

CANDIDATE LABEL POSITIONS FOR AREA FEATURES

For many simple applications (car-navigation systems and the like), area features can be identified adequately by simple horizontal labels. We utilize the same strategy for areas that we used for lines: generate a large number of candidate positions, then cull by precomputable metrics.

Given a polygon P representing the area's boundary, we find the inset polygon P' such that if a label is centred at any point in P' , the label will lie entirely inside P . (If no such P' exists, then it is impossible to place a label inside P that does not overlap the boundary of P . In this case, the area is labelled as if it were a point feature.) We then generate n (typically 200) quasi-random points evenly distributed throughout P' using a Sobol' quasi-random sequence (Press et al. 1992).¹ Candidate positions are generated centred at these 200 points. These positions are scored and all but the best-scoring k positions (again, we use $k=32$) are eliminated.

Figure 9 shows the top 20 generated positions for the label of a given area feature; the label itself is shown at the most favourable position of the 20.



Figure 9 Candidate label positions for an area feature

1 The n points are most useful if no two points are close together, because labels for nearby points are likely to overlap the same features and labels. The Sobol' quasi-random sequence is well suited for achieving a dispersed distribution of points.

As evidence that the labelling algorithm that we have presented has the potential to generate cartographically plausible labellings fully automatically and in reasonable time, we present a selection of sample maps with randomly generated point, line, and area features as labelled by our prototype software system.¹ The sample maps exhibit the ability of the algorithm to trade off the evaluation criteria for point, line, and area features in a reasonable manner.

PC. Only three of the eight seconds were spent selecting positions.²

Figures 11 through 13 show a typical progression of the position-selection process. Initially, labels are placed in random locations (Figure 11); in under half a second, most of the labels are positioned satisfactorily (Figure 12); and in under one second, the final labelling has been found (Figure 13).

A different type of map is shown in Figure 14. In contrast to the map in Figure 10, this map contains mostly line and area features, and shows the ability of the algorithm to find a plausible labelling under the tight constraints presented by the interaction of these features.

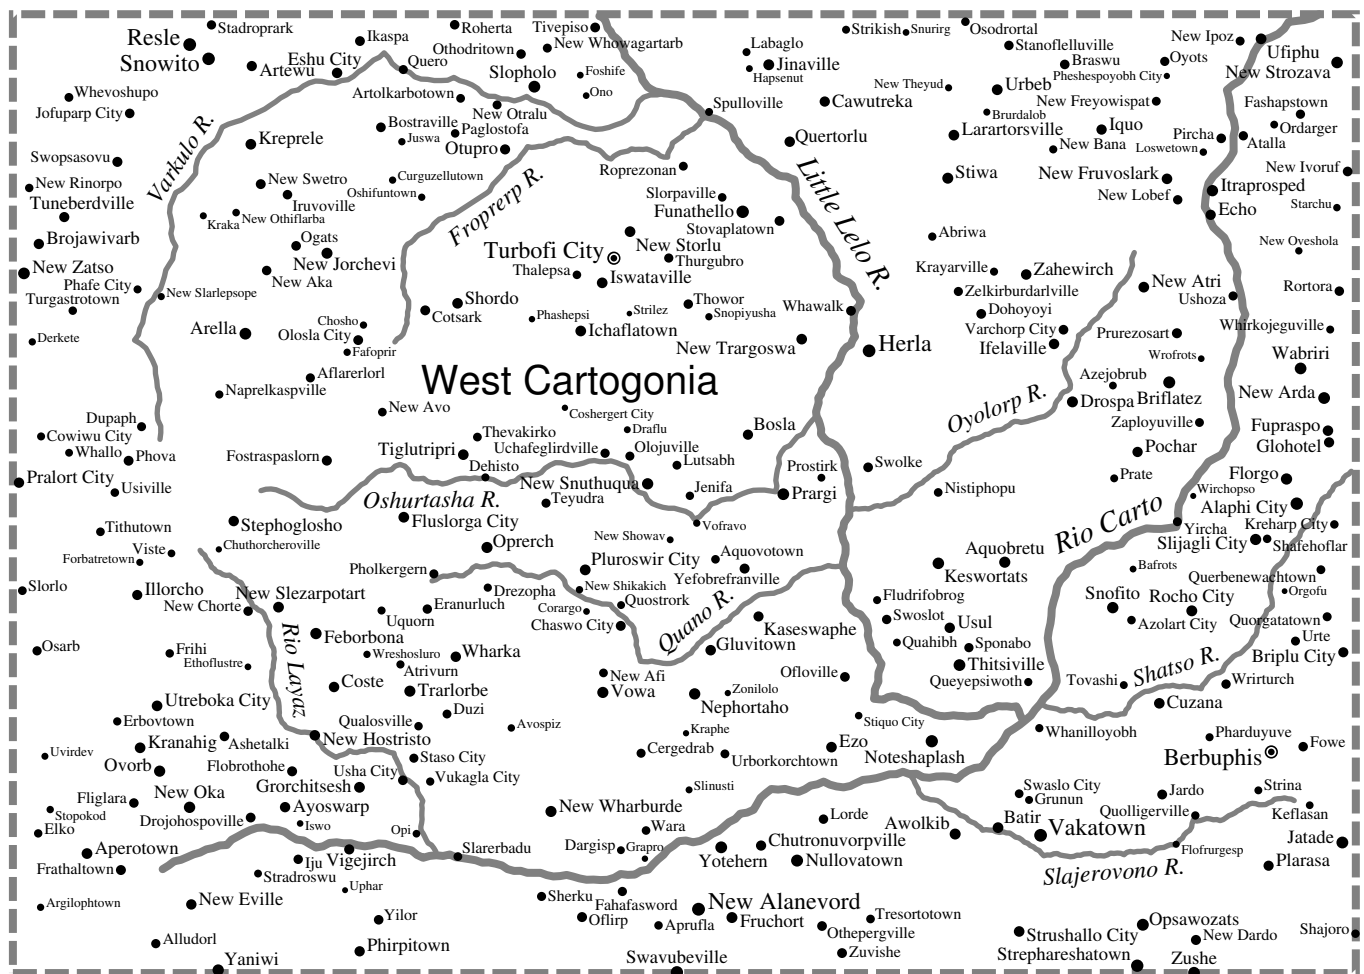


Figure 10 A randomly generated map labelled by the presented algorithm

1 The point-feature coordinates were picked randomly from a uniform distribution. Points less than 0.1 inches from the centre of any existing point feature were rejected so as to eliminate feature-feature overlaps. Additionally, a small number of points in configurations that precluded any valid labelling were removed or adjusted by hand. Line features were generated automatically by a fractal algorithm. Having a parametrized map-generation procedure that could produce maps with a variety of feature densities and characteristics was extremely useful in developing and testing our algorithm.

2 Most of the time is spent on precomputation that facilitates efficient position evaluation. Several time-consuming aspects of this precomputation, primarily pairwise intersection testing, use inefficient algorithms that were selected primarily for their ease of implementation. A commercial-grade implementation on a state-of-the-art workstation should be able to label maps like the one in Figure 10 virtually instantaneously.




Figure 11 The initial random labelling of a randomly generated map




Figure 12 The map from Figure 11, halfway through position selection




Figure 13 The final labelling of the map from Figure 11

Conclusions and Future Work

We have presented a cartographic labelling algorithm that we believe is:

- *general*—point, line, and area features are labelled using a unified methodology;
- *effective*—it produces compact, visually appealing maps;
- *efficient*—text placement is accomplished in less than

10 seconds on a standard personal computer for page- or screen-sized maps with up to 300 labelled features;

- *accessible*—it is simply explained, easily implemented, and readily extended.

Our implementation of the algorithm illustrates the advantages to be gained from a framework in which cartographic knowledge is summarized numerically in a scoring function that is then optimized by powerful heuristic search techniques.

While we believe this framework to be far superior to previous approaches that are based on expert-system production rules and depth-first search, we do not claim that the specific algorithm described here cannot be improved. For example, the following problems warrant further attention:

- *Disambiguation* Our current scoring function classifies the interaction between each label and its feature as either overlapping or non-overlapping. However, more subtlety is necessary to identify candidate positions that do not overlap features, but that come close enough to cause some degree of ambiguous association.
- *Selectivity* A useful option in labelling dense maps is to allow for (occasional) deletion of certain labels or even the features themselves in areas of congestion. As noted above, we had to resort to some manual feature deletion in our sample maps to ensure that an acceptable labelling was feasible. Automation of this selection process has already been demonstrated in the context of point-feature label placement (Langran and Poiker 1986; Christensen, Marks, and Schieber 1995b).
- *Flexibility* When label or feature deletion is not acceptable, the text can be modified to make the labelling problem simpler, or at least to provide additional placement options. For example, labels can be replaced by numbers (with a suitable key shown elsewhere), and text can be broken and stacked. Another option is to use a line to connect displaced labels and the features they tag.
- *Expressivity* Suggestion of candidate positions should allow for greater expressivity. For instance, horizontal labels are not ideal for all area features, curved labels may be better for some line and area features, and long linear features should allow for multiple occurrences of the label and for distribution of multi-word labels along the length of the feature.

These problems present no conceptual difficulty for the algorithm outlined here. We anticipate that they can be addressed by improved generation procedures and evaluation metrics without a significant increase in execution time. In the hope that others might want to build on our work, we are happy to make our software available on request to other researchers. Our testbed is written in C, runs on Unix platforms, and produces PostScript output.

5: 36–47, September.

- KIRKPATRICK, S., C. D. GELATT JR., and M. P. VECCHI. 1983. Optimization by simulated annealing. *Science* 220: 671–680.
- LANGRAN, GAIL E. and THOMAS K. POIKER. 1986. Integration of name selection and name placement. In *Proceedings of the Second International Symposium on Spatial Data Handling*, pages 50–64, Seattle, Washington, July. International Geographical Union and International Cartographic Association.
- PRESS, WILLIAM, SAUL TEUKOLSKY, WILLIAM VETTERLING, and BRIAN FLANNERY. 1992. *Numerical Recipes in C*. Cambridge University Press, Cambridge, England, 2nd edition.
- ZORASTER, STEVEN. 1986. Integer programming applied to the map label placement problem. *Cartographica* 23/3: 16–27.
- . 1990. The solution of large 0–1 integer programming problems encountered in automated cartography. *Operations Research* 38/5: 752–759, September–October.
- . 1991. Expert systems and the map label placement problem. *Cartographica* 28/1: 1–9, Spring

Résumé Certains algorithmes, apparemment puissants et visant le positionnement automatique d'étiquettes sur des cartes, utilisent des méthodes heuristiques qui représentent une expertise cartographique considérable. Par contre, leur efficacité est diminuée par des méthodes de recherche et d'optimisation qui s'avèrent inefficaces. D'autre part, aucune approche de posi-

tionnement d'étiquettes qui utilise une technique efficace d'optimisation n'a vraiment été utilisée lors de la production de cartes générales où l'on retrouve des points étiquetés, des lignes et des surfaces. Aucune méthode n'a démontré sa capacité à générer un étiquetage de qualité acceptable. Nous présentons ici un algorithme de positionnement d'étiquettes qui réussit à rencontrer les deux objectifs d'efficacité et de grande qualité en combinant des méthodes cartographiques simples et des techniques d'optimisation efficaces basées sur des variables aléatoires.

Zusammenfassung Ein allgemeiner Algorithmus für die Platzierung der Kartenschrift Einige scheinbar leistungsfähige Algorithmen für automatische Kartenschriftplatzierung verwenden eine Heuristik, in die beträchtliche kartographische Erfahrung eingegangen ist, die aber von nachweisbar unzulänglichen Methoden der Suche und Optimierung behindert wird. Andererseits hat man zwar Verfahren der Schriftplatzierung, die auf einer leistungsfähigen Optimierungstechnik basieren, in der Produktion allgemeiner Karten—d.h., solcher mit benannten Punkt-, Linien- und Flächenobjekten—eingesetzt, aber nicht zeigen können, daß sie Beschriftungen von annehmbarer Güte erzeugen. Wir führen einen Algorithmus für die Schriftplatzierung vor, der beide Ziele—praktische Leistungsfähigkeit und hohe Qualität der Beschriftung—erreicht, und zwar durch die Kombination einer einfachen kartographischen Heuristik mit wirkungsvollen stochastischen Optimierungstechniken.