




External Labeling Techniques: A Taxonomy and Survey

Michael A. Bekos¹, Benjamin Niedermann² and Martin Nöllenburg³

¹Institut für Informatik, Universität Tübingen, Tübingen, Germany

²Institute of Geodesy and Geoinformation, University of Bonn, Bonn, Germany

³Algorithms and Complexity Group, TU Wien, Vienna, Austria

Abstract

External labeling is frequently used for annotating features in graphical displays and visualizations, such as technical illustrations, anatomical drawings, or maps, with textual information. Such a labeling connects features within an illustration by thin leader lines with their labels, which are placed in the empty space surrounding the image. Over the last twenty years, a large body of literature in diverse areas of computer science has been published that investigates many different aspects, models, and algorithms for automatically placing external labels for a given set of features. This state-of-the-art report introduces a first unified taxonomy for categorizing the different results in the literature and then presents a comprehensive survey of the state of the art, a sketch of the most relevant algorithmic techniques for external labeling algorithms, as well as a list of open research challenges in this multidisciplinary research field.

CCS Concepts

• **Human-centered computing** → *Visualization techniques; Information visualization*; • **Theory of computation** → *Computational geometry*;

1. Introduction

Displaying textual annotations (or *labels*) for graphical features in maps, technical illustrations like assembly instructions or cutaway illustrations, or anatomy drawings is an important part of information visualization. One can broadly distinguish *internal* and *external* labeling, where internal labels are placed inside or in the direct neighborhood of a feature, whereas external labels are placed in the margins outside the illustration, where they do not occlude the illustration itself; see Fig. 1a for an example of external labeling used in anatomy drawings. In order to link each label to its respective feature, thin leader curves are used that connect label-feature pairs in the form of *callouts*; see Fig. 1.

Over the last twenty years various aspects of the external labeling problem have been considered in computer science – both from a theoretical and from a practical point of view. The range of topics is so wide that the term “external labeling” mostly serves as an umbrella term, which covers several different labeling approaches, such as *contour labeling* [AHS05,NNR17] (see Fig. 1b), *boundary labeling* [BKSW07] (see Fig. 1c), and *excentric* or *focus-region labeling* [FP99,BG05] (see Fig. 1d). Preliminary surveys on external labeling take a rather narrow view on selected labeling models, and focus either on selected theoretical results [Kau09,Wol13,KT18], or on practical works [OJP14].

In this state-of-the-art report, we introduce a unified and extensible taxonomy for the different parameters and models found in the external labeling literature (Section 2). Even though the focus

of this report is on the algorithmic techniques for external labeling, Section 3 first gives a summary of the most important visual aspects to be considered in external labeling. Then, based on the new taxonomy, we present a comprehensive literature survey of more than 50 research papers[†] on external labeling algorithms (Section 4), ranging from theoretical and algorithmic contributions to practical research and results in visual computing, including some empirical studies. From this literature, we collected a set of the most frequently applied algorithmic techniques, from exact, geometric algorithms to flexible, heuristic approaches, that are well suited to address external labeling problems. We provide sketches of the main techniques (Section 5) and discuss their properties and limitations as a guideline for researchers and practitioners interested in external labeling. We conclude with a discussion of ten future challenges in external labeling from various technical perspectives (Section 6).

The purpose of our state-of-the-art report is to provide a broad overview on this diverse and multi-disciplinary research topic, both for researchers entering the area or those extending their scope and for practitioners and domain users searching for suitable labeling techniques for their particular data visualization tasks. Our taxonomy will help to give structure to existing and future research results in external labeling, where previous attempts have been too narrow (e.g., only addressing boundary labeling) or even missing completely.

[†] On www.externallabeling.net we maintain a visual bibliography on external labeling techniques.

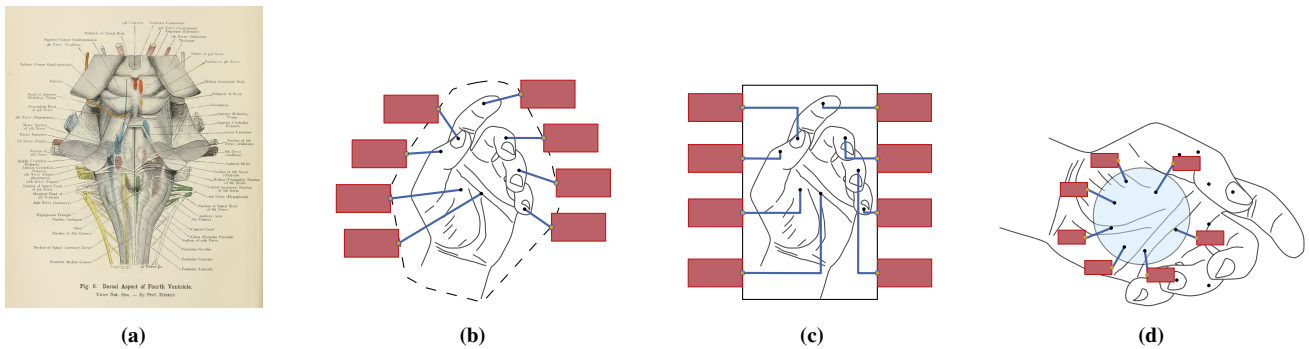


Figure 1: External labelings. (a) Anatomical illustration from: Atlas of applied (topographical) human anatomy for students and practitioners. K. H. v. Bardeleben, E. Haeckel, Rebman Company, 1906. (b) Contour labeling: the positions of the reference points of the leaders are restricted to a predefined contour (dashed). (c) Boundary labeling: the positions of the reference points of the leaders are restricted to a rectangle. (d) Excentric labeling: the labeled features are contained in a circle (blue), while the labels are placed around the circle.

The detailed literature survey gives a systematic overview of the existing knowledge collected in more than 50 publications. The high-level discussion of the most used algorithmic techniques in external labeling serves as an orientation aid when deciding which existing algorithm to select and implement or when designing a new labeling algorithm for a particular labeling task. Typically generic techniques are simpler to adapt to new settings and faster to implement, whereas specialized algorithms show better performance but are more difficult to adapt and implement. Finally, our collection of open research challenges provides ample opportunities for future work on external labeling. The challenges range from algorithmic problems and less explored labeling models, e.g., combining external and internal labeling, to perceptual questions and studies of human comprehensibility of the labeled data.

2. A unified taxonomy

This section introduces a unified and extensible taxonomy of different labeling models proposed in the literature and may serve as an entry point for the interested reader who wants to get familiar with the basic concepts of external labeling. In contrast to previous, narrower taxonomies (e.g., on boundary labeling [BKPS06b]), we provide a common, general structure for classifying existing and future results in the wider field of external labeling. It is designed to be general and extendable enough for addressing new challenges in external labeling and covers all aspects of existing work. First, we formally introduce the most important terminology and concepts on external labeling (Section 2.1). Afterwards, we present distinctive features that we use to characterize and unify the different models for external labeling found in the literature (Section 2.2). Finally, we formally define a generic optimization problem that underlies approaches for creating external labelings automatically (Section 2.3).

2.1. Terminology and concepts

In general a figure with external labeling can be decomposed into three overlaid layers; see Fig. 2. The bottom layer defines the image area containing the illustration and the labeling area. The middle layer contains the features to be labeled. The top layer finally contains the annotations for the features.

Background layer The bottom layer partitions the available *drawing area* D along a simple, closed curve C_R into an *image region* R and a *labeling region* $A = D \setminus R$; we call this layer the *background layer*. The image region contains the unlabeled illustration, while the labeling region defines the empty space surrounding the image, which is reserved for placing the annotations. For instance, the image region could be a bounding box or a convex hull of the illustration.

Feature layer The middle layer defines a set F of pairwise disjoint features of interest to be labeled; we call this layer the *feature layer*. Each feature is a small sub-region of R , which can be a point, a curve or an area. Further, each feature has some attached information, e.g., its name, an icon or a short textual description, which will later be used as the content of its label. Geometrically, we define the *label* as the axis-aligned bounding box of the feature's attached information. We denote the set of all labels by L . Typically, each feature has its own label, but in some special cases multiple features can share the same label or one feature can have several labels.

Labeling layer The top layer finally adds the labels and leaders to the figure; we call this layer the *labeling layer*, refer also to Fig. 3 for a detailed example. More precisely, the label ℓ of a feature f is placed within the labeling region A and connected to f by a simple curve λ ; we call λ the *leader* of ℓ and f . One end point of λ , the so-called *site* (sometimes also called *anchor*) $\sigma(\lambda)$, is a point contained in f . The other end point of λ , the so-called *reference point* $a(\lambda)$, is a point on the boundary of ℓ . We distinguish different ways of placing sites, reference points of leaders and labels. If the site σ can be any point of f , then σ is a *free site*. If σ is restricted to be any point of a pre-defined curve in f , then σ is a *sliding site*. Finally, if σ is restricted to be any point of a finite set of points in f , then σ is a *fixed site*. Note that the site of a point feature is always fixed, in fact the site and the feature coincide. Similarly, if the reference point a can be any point in A , then a is a *free reference point*. If a is restricted to be any point of a curve C_A in A , then a is a *sliding reference point*. Finally, if a is restricted to be any point of a finite set P_A of points in A , then a is a *fixed reference point*. The most common case is the use of sliding or fixed reference points, where the curve C_A or the set P_A either coincide with C_R or are offset by a small distance from




Figure 2: Layer decomposition. The figure (left) is decomposed into the background, feature, and labeling layer.




Figure 3: Terminology used throughout this paper.

C_R . The case of free reference points is less relevant, because labels should generally be placed close to the image region. Recall that the reference point a is a point on the boundary of the label ℓ . From the perspective of ℓ we call this attachment point between the leader λ and ℓ the *port* $\pi(\ell)$ of ℓ , which implies $a(\lambda) = \pi(\ell)$. If the port $\pi(\ell)$ can be any point on the boundary of ℓ , then $\pi(\ell)$ is a *sliding* port. If $\pi(\ell)$ is restricted to be any point of a finite set of points on the boundary of ℓ , then $\pi(\ell)$ is a *fixed* port. The most common fixed-port model restricts the ports to the corners or midpoints of the edges of ℓ . If also the label positions are fixed, the reference point of λ is limited to the ports of ℓ . We call the geometric composition $\gamma = (\lambda, \ell, \pi(\ell))$ of a leader λ and a placed label ℓ attached to each other at the port $\pi(\ell)$ a *callout*. A set \mathcal{L} of callouts is called an *external labeling* of the feature set F ; for brevity we also call it *labeling*. We say that \mathcal{L} is *plane* or *crossing-free* if no two callouts in \mathcal{L} intersect.

2.2. Distinctive features

In this section we discuss three distinctive features of external labeling that are used to characterize the different labeling models found in the external labeling literature.

Admissible positions of labels The admissible positions of the reference points of the leaders and accordingly of the labels significantly influence the overall appearance of the labeled figure. While in general models for external labeling the reference points may be placed anywhere in the labeling area, models for *contour labeling* restrict the positions of the reference points to a predefined contour C_A ; see Fig. 1b. This contour is often chosen in such a way that it roughly matches the shape of the illustration in the background layer. This is a common technique to achieve a label placement that blends in with the illustration.

An even more restrictive model is that of *boundary labeling*, which requires that the image region R is a rectangle and C_A coincides with the boundary of R ; see Fig. 1c. Hence, excluding the corners of R for the placement of reference points, each label intersects exactly one edge of the image boundary in more than one point; we say that the label *touches* that edge. Depending on the number of edges that are touched by labels, we distinguish *1-sided*, *2-sided*, *3-sided* and *4-sided* boundary labeling models. We observe that for 2-sided boundary labeling models the edges that are touched by labels are either *adjacent* or *opposite* edges of R ; Figure 1c shows a 2-sided boundary labeling with opposite edges touched by labels.

In *excentric labeling* only features in a restricted region of the illustration are labeled with their labels placed outside that region; see Fig. 1d. In spite of labels possibly overlapping the actual illustration in the background, we consider excentric labeling to be a special case of external labeling, because the labels are still clearly separated from their features. In excentric labeling the restricted region is typically described by a circle of fixed radius and implements the metaphor of a *lens* that can be moved over the illustration to explore its details. Alternatively, one can think of this restricted region as a *focus region* of the user. In the literature, two main variants are found: labels that are placed freely in the surroundings of the focus region and the special case of contour labeling requiring that all labels are placed along the boundary of the focus region.

Leader type The model that is introduced in Section 2.1 allows general curves for leaders without any restriction to their shape. However, in actual externally labeled figures only leaders with simple, schematic shapes, such as straight lines and polylines, are used

in order to avoid cluttered illustrations and to sustain legibility. In rare cases, *curved* leaders such as Bézier arcs are used. Next, we present a systematic classification of different polyline leader types. Let λ be a leader that consists of k segments, which are ordered from its site to its reference point. We describe the shape of λ by a string $z \in \{s, r, d, o, p\}^k$ of k symbols, where the i -th symbol of z describes the orientation of the i -th segment of λ . In the following we list the meaning of each symbol.

- s** An *s*-segment is a straight-line segment with unrestricted slope.
- r** An *r*-segment lies on a ray that emanates from a given center point M .
- o** An *o*-segment is orthogonal to a reference line or line segment m .
- p** A *p*-segment is parallel to a reference line or line segment m .
- d** A *d*-segment l (for a given angle $0^\circ < \alpha < 90^\circ$) is a diagonal segment defined relative to its preceding segment l_p and its succeeding segment l_s as follows. The turning angle between l_p and l as well as the turning angle between l and l_s is $\pm\alpha$, where the turning angle of two segments sharing a common endpoint is the minimum angle in which one has to rotate one of the two segments around their common endpoint to reach the second.

Due to their shape, leaders consisting only of *o*- and *p*-segments are called *orthogonal* leaders, when the reference line or line segment is either vertical or horizontal. If one additional type of *d*-segments is allowed, then the resulting leaders are called *orthodiagonal*. In the special case, in which the turning angles of each *d*-segment with its preceding and succeeding segments are $\pm 45^\circ$, the obtained leaders are called *octilinear* following the established terminology in network visualization.

In models for boundary labeling the shape of a leader is typically described with respect to the edge of the boundary rectangle R that is touched by the leader's label. Following this convention, Figure 4a illustrates the most frequent leader types in the literature. The labels A, E, F, G and H are connected to *s*-leaders. More precisely, while A is connected to an arbitrary *s*-leader, the labels E, F, G, and H have *r*-leaders, because their straight-line segments lie on rays emanating from the center M . Further, the labels D and M have *o*-leaders.

The remaining labels B, C, I, J, K, and L have leaders with more than one segment. Assuming again that the reference line segment of each leader is the side of the rectangle containing its label, B has a *po*-leader, C has an *opo*-leader, and L has an *po*-leader; recall that the segments of a leader are ordered from its site σ to its reference point a . Further, label I has a *do*-leader, label J has a *pd*-leader, and label K has a *do*-leader. In all three cases the *d*-segment has a turning angle α with respect to its preceding and succeeding segment, respectively; with rare exceptions $\alpha = 45^\circ$ is used for *d*-segments in the literature.

In some applications multiple features can share the same label, i.e., a single label is connected to multiple features via multiple leaders; see Fig. 4b. Using this labeling style the reader can easily identify features with the same meaning. Further, it saves space in the labeling area. We refer to this labeling technique as *many-to-one labeling*. Further, by connecting all leaders to the same reference point of the label, the leaders can be bundled; we call the set Λ of these leaders a *hyperleader* of the label; see Fig. 4c. The literature on labeling with hyperleaders requires that no two leaders in Λ

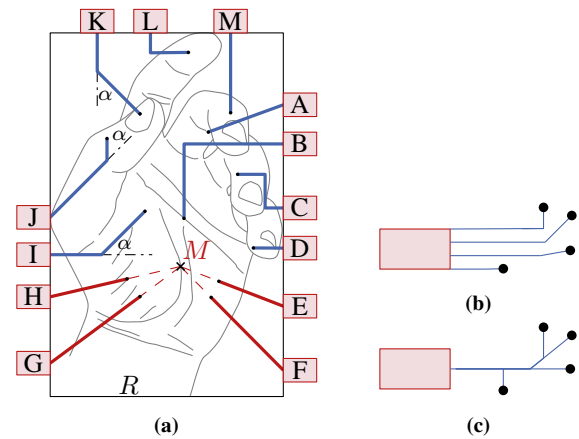


Figure 4: Types of leaders. (a) Leader types mostly used in literature. (b) A label connected to multiple sites. (c) A hyperleader, i.e., a set of leaders having the same reference point, but different sites.

intersect apart from a common suffix. This leads to the desired visual impression that Λ is a single leader forking to multiple sites.

Temporal dimension A widely used application of external labeling are digital visualizations that change over time. For example, consider a medical visualization system that provides the interactive exploration of a 3D body model and that uses external labeling to explain the features of that model. When the user changes the view of the 3D model, the labeling changes correspondingly. From a technical point of view, this process can be seen as an animation consisting of a temporal sequence of 2D images that arise from projecting the 3D model onto the screen space. For each of these 2D images an external labeling is created to name the model's features. Hence, we obtain a temporal sequence of external labelings; we call that sequence a *dynamic labeling*. In contrast, we call a single labeling in that sequence a *static labeling*. Hence, we can distinguish models for *static external labeling* that aim at individual labelings, and models for *dynamic external labeling* that aim at sequences of labelings. Models on dynamic external labeling may define constraints enforcing temporal coherence between consecutive labelings in order to avoid distracting effects such as flickering and jumping labels. In most of the literature on labeling of 3D models the labels are placed in the projection space, which is also our default here, but there are also some papers that place labels in the object space.

2.3. Optimization problem

Next we introduce the formal definition of the external labeling problem. As input we are given a drawing area $D = (R, A)$ partitioned into the image region R and the labeling region A , a feature set F and a set M of model parameters to be specified below. We call the tuple $I = (D, F, M)$ an *instance* of external labeling. A labeling \mathcal{L} of F that satisfies all model parameters of M is called a *feasible* labeling of I . We denote the set of all feasible labelings of I by \mathcal{S} .

The model parameters are a set of layout-specific criteria that must be satisfied. Typical parameters are the selected leader types, site, port, and reference point constraints, and hard constraints that

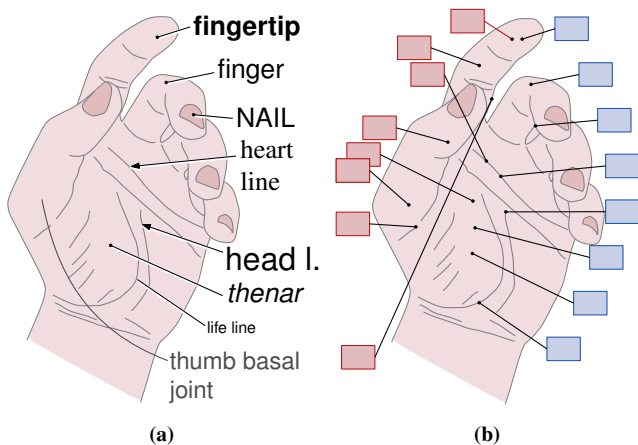


Figure 5: (a) Different examples of rendering and typesetting callouts. (b) The labels on the right-hand side (blue) comply with Criteria C1–C5, the labels on the left-hand side (red) do not.

restrict the mutual interplay of the callouts or their interference with the illustration in the background layer. While the leader types have been specified in Section 2.2, the site, port, and reference point constraints define the admissible placement of the site (fixed, sliding, free), of the port (fixed, sliding), and the reference point (fixed, sliding, free). Examples of hard constraints are crossing-free leaders, non-overlapping labels, and avoidance of obstacles in the illustration. The actual set of model parameters strongly depends on the application and is determined by the designer, see Section 3 for some commonly used guidelines.

Most commonly, external labeling is seen as an optimization problem to find the best among all feasible labelings (or at least a reasonably good one). Hence, it remains to define a cost function $c: \mathcal{S} \rightarrow \mathbb{R}^+$ that measures the quality of the feasible labelings in \mathcal{S} , e.g., the total length of all leaders. For a detailed discussion on typical optimization criteria see Section 3.2.

Problem 1 (EXTERNAL LABELING)

Input: Instance $I = (D, F, M)$, cost function $c: \mathcal{S} \rightarrow \mathbb{R}^+$

Output: Optimal labeling $\mathcal{L} \in \mathcal{S}$, i.e., $c(\mathcal{L}) \leq c(\mathcal{L}') \forall \mathcal{L}' \in \mathcal{S}$.

3. Visual Aspects of External Labeling

In this section we discuss various visual aspects that affect the aesthetic quality and usability of external labeling. In general, it is far from obvious, which type of external labeling is the best choice and this highly depends on the application. We discuss several important criteria that must be taken into account when deciding on a type of external labeling. We distinguish between the *general style* of the labeling and the actual *placement* of the labels based on it.

3.1. General Style

The visual appearance of a callout in an externally labeled illustration depends on many aspects, which influence the legibility and readability, as well as the overall aesthetics of the illustration to

various degrees. While the focus of this survey is on the algorithmic aspect of optimizing the positioning of labels and leaders, this section gives a high-level overview of other practically important parameters affecting the quality of a labeling. Figure 5a illustrates the visual effects of these design choices.

First of all, most of the leaders used in external labeling algorithms are either straight-line leaders (see Section 4.1) or polylines whose segments have restricted slopes (see Section 4.2). Only few papers also considered curved leaders [KLW14]. The choice of the leader type has a strong influence on the overall appearance. In terms of readability, it is important to consider the overlay of the illustration in the background layer and the leader curves in the labeling layer. One must ensure that lines in the background layer and leader lines can be easily differentiated [Ric17], which means that in illustrations with arbitrary and irregular shapes (e.g., maps or anatomy drawings such as Fig. 1a and 5a), straight and polyline leaders are preferable, whereas in illustrations with many straight lines of regular slopes (e.g., floorplans), curved leaders or straight and polyline leaders with non-aligned slopes should be considered [WTLY12].

Richards [Ric17] presented a list of best-practice guidelines originally for labeling black-and-white technical drawings in aerospace engineering, but most of them should also generalize to other domains. Leader lines should be straight wherever possible and they should consist of a black line and right next to it, on one side, a thin white line to separate the leader from the background. Moreover, leader lines (for area features) should terminate with a terminator symbol, either a small dot if the feature is large enough, or an arrowhead for small features. According to Richards [Ric17], polyline leaders producing staircase effects, possibly for aesthetic reasons, are acceptable as long as there are not too many leaders, but should be avoided for contexts where safety or efficiency of use are relevant.

The graphical rendering of leader lines themselves is also of relevance for the visual quality of a labeled image. Wu and Dalal [WD05] considered different quality features of rendering lines on electronic displays, including the raggedness, blur, waviness, and lumpiness of lines, depending on the method of digitizing Euclidean straight lines and using anti-aliasing effects.

In addition to the leader, the label itself needs to be properly configured and displayed. Foremost, the typesetting of the label text must be done appropriately. For example, Degani [Deg92] gave a detailed list of guidelines and principles for typesetting of flight-deck documentation, whose use of text is comparable to labeling. We mention the most important ones. While one can choose from a large variety of fonts, which should match the context of the illustration, for labels and short annotations generally sans-serif fonts are more legible than serif fonts. The selected font size depends on the estimated viewing distance and can be calculated according to standard guidelines. The same is true for non-textual labels such as icons. Lower-case words are more legible than upper-case words. Regular face is more legible than italic or bold face, which should only be used for highlighting specific words. Finally, a black font over white or lightly colored background has the best legibility.

For the text itself, one consideration is whether abbreviations or line breaks should be used, e.g., Shimabukuro and Collins [SC17] presented a method for on-demand abbreviations of text labels.

3.2. Placement

From a visual point of view the general labeling style surely influences the appearance of the image most. However, once the labeling style has been fixed, the actual placement of the labels comes to the fore. Choosing the most legible and visually pleasing one among all feasible placements is a complex task that requires highly elaborated algorithmic techniques when done automatically, and a good balance of possibly contradicting placement criteria.

Static Labelings For static labelings we have identified the following decisive criteria based on literature that justifies them by interviews with experts and manual analysis of labeled drawings; see Tables 1 and 2 for the according references.

- C1 *The leaders have small length.*
- C2 *The number of leader crossings and label overlaps is small.*
- C3 *The labels mimic the shape of the image.*
- C4 *The labels are distributed evenly.*
- C5 *The directions of the leader segments comply with a set of preferred directions.*
- C6 *The leaders have a small number of bends.*

Criteria C1 and C2 reduce visual clutter and ambiguities and ease association between label and feature; a crossing- and overlap-free solution is often preferred or even required [AHS05, BKS07]. The importance of both criteria is strengthened by the fact that the majority of research takes into account both; see Tables 1 and 2.

Criteria C3–C6 aim at labelings that blend in with the image. For example, evenly distributed labels lead to a homogeneous appearance of the labeling, and labels mimicking the shape of the figure repeal the clear distinction between labeling layer and background layer. They are justified in multiple references (e.g., [AHS05, MTSP12, NNR17]). Concerning the number of occurrences in literature, the following criteria are deemed less important, but still they are verified by interviews with experts [NNR17].

- C7 *There is sufficient space between leaders.*
- C8 *Labels consist of single text lines if possible.*
- C9 *Labels that are semantically related are grouped.*

Criterion C7 reduces ambiguities between leaders, while Criteria C8 and C9 aim at a better legibility and comprehensibility of the labeling, respectively.

We stress that there is not a clear, empirically confirmed ranking among these criteria, which is one of the research challenges discussed in Section 6; their respective relevance and relative priorities also depend on the application domain or even on the image to be annotated. Still they are supported by general labeling and drawing principles such as proximity, planarity, straightness/low detour, and small ink-ratio [Imh75, FW91, Pur97, WPCM02, Tuf01]. We also stress that some of these criteria might be in conflict with each other. As a result, a good labeling is usually a compromise of different criteria. To this end, one usually scales the different criteria by suitable weight factors to obtain a multi-criteria cost function c in Problem 1.

Dynamic Labelings Dynamic labelings are especially found in view management systems, in which the user can interactively explore 3D models. Hence, the placement of the labels becomes an *online* problem for which the labeling algorithm needs to create a

dynamic labeling without exactly knowing the upcoming changes. As the frames of a dynamic labeling form a sequence of static labelings, the criteria for the static case also apply for the dynamic case, but are extended by the general requirement of temporal coherence between two consecutive frames. Following general dynamic map labeling principles [BDY06], *flickering* and *jumping* labels should be prevented, while optimizing and preserving the placement criteria listed above. The proposed approaches found in literature can be categorized into two groups [MTM*16]. The approaches of the first group consider continuously moving labels, which clearly prevents flickering and jumping labels. However, the placement criteria such as overlap-free labelings and sufficient spacing between labels, cannot be completely maintained. In contrast, the approaches of the second group use *hysteresis*, i.e., the displacement of the labels is delayed to stabilize the motion of the labels and to enforce selected placement criteria. However, this easily leads to labels that jump from one position to the next. Preserving the order of the labels and animations visualizing the changes are possible tools to soften this effect. In general, animations have been considered as an important tool in interactive visualizations to visualize the transition from one state to the other [EMJ*11]. We further note that Madsen et al. [MTM*16] showed in a detailed user study on labelings in augmented reality that the participants performed better in the conducted tasks when the labels were directly placed in the 3D space of the object instead of the 2D space. This result is further supported by the user study by Peterson et al. [PAE09] showing that incorporating depth information into the labeling increased the performance of the participants in selection tasks. They particularly showed that this technique softens the negative effects of overlapping labels. Still, only few approaches [PHTP*10, TKGS14] that aim at a clear spatial distinction between objects and placed labels have been developed for external labeling in object space so far.

4. State of the art

In this section we describe a comprehensive collection of the state of the art on external labeling algorithms. We discuss in total 54 references, which we collected in two main steps. Starting from a list of references that we have collected over the years, we enriched this list by a careful search in Google Scholar using different search terms closely related to external labeling. Among others we used the following terms: *external labeling*, *boundary labeling*, *contour labeling*, *excentric labeling*, *callouts*, *labeling with leaders*, *external labels*, *labeling with focus regions*, *hyperleaders*, *automatic text annotations*, *label layout*, *many-to-one labeling*, *labeling of visualizations*. We further enriched the list by parsing the references and the citations of each of the papers that we obtained in the first step as well as by taking into account existing surveys. The result of our search consists of references from journals and conference proceedings related to the visual computing community (including information visualization, geovisualization, augmented reality and human computer interaction) and related to the algorithms community (including computational geometry and graph drawing). We stress that our search was limited to scientific results and did not go beyond this area (e.g., to patents). The goal of this section is to acquaint the reader with the state of the art in both communities and to intensify the exchange of ideas between theory and practice.

In order to structure this section, we group the results by their distinctive features. At the top level we use the leader type as distinctive feature, because it substantially determines the visual appearance of an external labeling. We have identified 29 references using straight-line leaders (Section 4.1) and 25 references using polyline leaders (Section 4.2), which partitions the results into two approximately equally sized groups. At the same time this roughly groups the results into practical and theoretical ones, which can be of more interest for researchers in the visual computing and in the algorithms communities, respectively. We further review the current state on curved leaders comprising only few references (Section 4.3), and empirical evaluation techniques in external labeling (Section 4.4).

We note that our taxonomy as defined in Section 2 does not include all the labeling models that use labels with leaders in general. It requires that the labels in the labeling region A are clearly separated from the underlying illustration in the image region R . Note that there exist labeling models in the literature using *internal* labels with leaders (e.g. [BFH01, AF03, MD06, SMHW15]), where this clear spatial separation between labels and illustration is not required. Since such labeling models raise very different algorithmic problems due to the strong interference between the image and the labels, we exclude their discussion from this survey.

4.1. Straight-line leaders

Straight-line leaders are the simplest way to establish the visual association between features and their external labels. Further, they can be easily traced by the reader [BGNN18]. We first give an overview on the current research (Section 4.1.1) and then discuss the individual papers concerning their particular contributions (Section 4.1.2).

4.1.1. Overview

Using straight-line leaders for external labeling requires a broad variety of design decisions. As a result plenty of different approaches have been developed. We summarize the most important properties of the approaches in Table 1 and discuss them in the following.

Label positions The high-level grouping of the relevant literature as seen in the top row of Table 1 considers three general approaches of placing labels. The first group of approaches considers contour labeling assuming that the labels are placed along a pre-defined contour, which means that labels mimic the silhouette of the illustration. Especially in professional applications such as atlases of human anatomy domain experts require this property [NNR17]. From an algorithmic point of view, pre-defined contours have the advantage that they reduce the search space by one dimension. On the other hand, this also implies that parts of the labeling region are entirely omitted for label placement, while the chosen contour may not host all labels. The second group of approaches relaxes this problem by freely placing the labels using the entire labeling area. We distinguish between free label placement within a continuous labeling region [FP99, HAS04, GHS06b, VVAH07, TKS13, BČM14, HGK14, TKGS14], and free label placement within a discretized labeling region [FLH*06, WTLY11, GHN15]. Finally, the last group of approaches relaxes the requirement that labels may not be placed inside the image region. As these approaches still place the labels in the outer regions of the image as good as possible (in contrast

to internal labelings with leaders), we consider them to be in the corpus of literature on external labeling, although they do not fully satisfy our definition of external labeling; we call this placement technique *non-strict* external labeling accordingly.

S1: labeling contour For research assuming a pre-defined labeling contour, the shape of the contour has an essential impact both on the visual appearance of the labeling as well as on the algorithmic properties. We distinguish between bounding rectangles (S1.1), circles (S1.2) and (buffered) convex hulls (S1.3) as possible shapes, but in rare cases also more complex hulls (S1.4) are taken into account.

S2: site placement In illustrations of human anatomy the locations of the sites are often pre-defined by a domain expert before the labeling is done [NNR17]. However, for example, in interactive view systems, this is often not adequate, and as a result the positions of the sites need to be calculated. The sites should be placed on salient points of the features such that they do not form dense clusters. Further, considering a 3D object, the projection of the object into the image space is not necessarily connected, but may consist of multiple parts. Hartmann et al. [HAS04] suggested to label the biggest part or the part that promises the shortest leader. In most of the research that explicitly considers the placement of sites, this is done in a pre-processing step before the actual labels are placed [HAS04, AHS05, HGAS05, GHS06b, VVAH07, MP09, TKGS14]; only in two cases placing sites and labels is considered simultaneously [ČB10, ČB18]. Hence, most of these techniques can also be installed upstream for the other approaches that require pre-defined sites. Moreover, some of the approaches also adapt the positions of the sites to improve temporal coherence in dynamic scenarios [AHS05, HGAS05, GHS06a, MP09, ČB10, TKS13].

S3: drawing criteria All approaches pay special attention to intersections between callouts, because they easily distract the reader and let the layout appear cluttered; 18 references guarantee that the created labeling is plane, while for the other eleven approaches intersections (S3.1) might occur, e.g., if too many labels need to be placed. Similarly, the length of the leaders (S3.2) is considered to be an important drawing criterion to enforce unambiguity; only two references do not explicitly take leader length into account. Moreover, the spacing between overlapping-free labels (S3.3) is considered by penalizing distances between labels [HAS04, AHS05, HGAS05, GHS06a, VVAH07, TKS13, BČM14, KLW14, NNR17], by stacking them [FP99, GHS06b, BKS07, AHS05, HGAS05, HGK14, GHN15], or by distributing them evenly around the image region [BSF*11]. Further, in 15 references the direction of the leaders (S3.4) is considered as drawing criterion. Either the angles of the leaders are optimized with respect to some pre-defined directions [HAS04, GHS06a, GHS06b, VVAH07, SD08, ČB10, HGK14, GHN15, NNR17, ČB18], the leaders radially emanate from a common center [AHS05, HGAS05, FHS*12, HGK14, TKGS14], or monotonically increasing angles are required when radially ordering the leaders [MTSP12, NNR17]. A vertical and horizontal alignment of the labels or of groups of labels (S3.5) is achieved by stacking the labels horizontally or vertically aligned [FP99, GHN15, GHS06b, VVAH07, HGK14], or by prescribing a rectangular labeling contour [BKS07, ČB10, MTSP12, KLW14, FS16, BGNN18]. Further, the labels are either grouped or clustered (S3.6), and occlusion of the

	[BKS04]	[AHS05]	[BG05]	[HGAS05]	[GHS06a]	[ČB10]	[BSF*11]	[FHS*12]○	[HH14]○	[KLW14]	[FS16]	[NNR17]	[BGN18]	[ČB18]	[FP99]○	[GHS06b]	[FLH*06]	[VVAH07]	[WTLY11]	[HGK14]○	[GHN15]	[HAS04]	[SD08]	[MP09]	[PHTP*10]	[MTSP12]	[TKS13]	[BČM14]○	[TKGS14]	Σ
S1 labeling contour	x	x	x	x	x	x	x	x	x	x	x	x	x	x									x	x					16	
S1.1 rectangle	x	x	x	x	x	x	x	x	x	x	x	x	x	x												x			8	
S1.2 circle																										x			6	
S1.3 convex hull																													6	
S1.4 other hull					x																			x					2	
S2 site placement	x		x	x	x									x		x	x					x	x	x	x	x	x	x	12	
S3 drawing criteria																														
S3.1 crossing-free callouts	x	x	x	x	x	x	x	x	x	x	x	x	x	x			x					x	x	x					18	
S3.2 leader length	x	x	x	x	x	x	x	x	x	x	x	x	x	x		x	x	x	x	x	x	x	x	x	x	x	x	x	27	
S3.3 label spacing	x	x		x	x	x	x			x		x	x		x			x	x	x	x	x				x	x		16	
S3.4 leader direction				x	x	x		x	x			x	x					x	x	x	x	x	x						15	
S3.5 vert./horz. aligned labels	x					x				x	x		x		x	x		x	x	x	x	x			x				11	
S3.6 groups/clusters					x		x					x				x		x					x						7	
S3.7 leader-image occlusion											x											x							3	
S3.8 mixed labeling				x	x														x										3	
S4 temporal coherence	x		x	x	x										x								x	x	x	x	x	x	12	
S5 labeling techniques																														
S5.1 init and improve		x	x	x	x									x				x		x		x	x	x	x	x	x	x	14	
S5.2 greedy			x		x	x										x			x	x	x	x	x			x	x	x	12	
S5.3 force-based		x		x																		x							8	
S5.4 dynamic programming								x	x			x	x									x							5	
S5.5 meta-heuristic		x																x	x										3	
S5.6 matching	x							x																					2	
S5.7 plane sweep								x																					1	
S5.8 mathematical prog.													x									x							2	
S6 contributions																														
S6.1 implementation		x	x	x	x	x	x	x	x	x		x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	27	
S6.2 analysis: exist. drawings		x		x														x				x				x			6	
S6.3 formal proofs	x							x	x		x	x										x							6	
S6.4 expert interviews													x	x										x					4	
S6.5 user study													x	x	x									x	x			x	6	
S7 algorithm type																														
S7.1 heuristic		x	x	x	x	x				x		x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	24	
S7.2 exact	x							x	x	x	x	x	x									x							8	
S8 community																														
S8.1 algorithms	x									x	x																		3	
S8.2 visual computing		x	x	x	x	x	x	x							x	x	x	x	x	x	x	x	x	x	x	x	x	x	26	

Table 1: Properties of the approaches for straight-line leaders. References are partitioned into three top-level groups by contour labeling, free label placement and non-strict external labelings. References considering excentric labeling are marked with “○”.

image background by leaders (S3.7) is taken into account. Finally, three references consider mixed labelings (S3.8) supporting internal and external labels. Other criteria such as the distance between leaders and sites are found in literature; for detailed lists of possible criteria see for example [AHS05, ČB10, MTSP12, NNR17].

S4: temporal coherence In general all of the approaches can be used in dynamic scenarios, as each displayed frame of an animation can be labeled independently. However, this easily creates flickering and jumping sites and labels. Therefore techniques that minimize

the movement of the sites and labels [AHS05, HGAS05, GHS06a, SD08, MP09, ČB10, PHTP*10, TKS13, BČM14] as well as hysteresis techniques that temporally freeze the current labeling [GHS06b, MP09, MTSP12, TKGS14, BČM14] have been developed to soften this effect.

S5: labeling techniques The approaches for obtaining labelings with straight-line leaders are rarely limited to one specific labeling technique, but they use combinations of techniques to obtain an overall strategy. In particular, 14 references use the scheme of

first computing an initial labeling and then improving that labeling successively (S5.1). We observe that greedy algorithms (S5.2) and force-based (S5.3) algorithms are used by twelve and eight references, respectively. They are easy to implement, run fast and provide the possibility of integrating multiple drawing criteria. Other labeling techniques such as dynamic programming (S5.4), meta-heuristic (S5.5), matchings (S5.6) and sweep-line algorithms (S5.7) are rarely used. Fekete and Plaisaint [FP99], Battersby et al. [BSF*11] Kindermann et al. [KLW14], and Götzelmann et al. [GHS06b] used highly specialized algorithms that do not match any of the listed techniques; details are given in Section 4.1.2.

S6: contributions The contribution of research on the straight-line leaders is manifold. Firstly, 27 references describe implementations (S6.1), which are also used for evaluation purposes. All of them discuss examples generated with the presented approaches. Further, six references justify their design decisions on the analysis of existing drawings (S6.2), e.g., atlases of human anatomy and visual dictionaries, and four of them on interviews with domain experts (S6.4) such as surgeons [MP09], radiologists [MTSP12] in the context of medical imaging, designers and editors working on atlases of human anatomy [NNR17] as well as on interviews with experts of related fields (user interface and graphics design) [ČB18]. Further, six of them additionally evaluate the resulting labelings by user studies (S6.5). Moreover, six references give formal guarantees on certain optimization criteria (S6.3), e.g., the total leader length. Apart from analyzing the quality of the labelings, the performance of the presented approaches is discussed by nine references in detail [BG05, GHS06a, VVAH07, ČB10, FHS*12, GHN15, NNR17, ČB18].

S7: algorithm type We consider two different types of algorithms found in the literature. Most of the references (24) consider heuristics, i.e., approaches that find good but not necessarily optimal solutions. Only eight approaches are exact algorithms in the sense that they yield optimal solutions for pre-defined cost-functions.

S8: community We partition the research into references that either stem from the algorithms community or the visual computing community. More precisely, for each reference we take into account in which journal or conference proceedings it has appeared and relate it to one of two communities accordingly. We observe that for straight-line leaders the majority of references (26) stem from the visual computing community, while only three references come from the algorithms community.

4.1.2. Detailed discussion

In this part we discuss the research on straight-line leaders in detail. To that end, we group the results with respect to the allowed label positions, as apart from the chosen leader type the position of the labels strongly determines the visual appearance of the labeling.

Contour labeling Bekos et al. [BKS04, BKS07] considered rectangles as labeling contour and coined the term *boundary labeling*. For straight-line leaders they considered uniform labels on all four sides of the rectangle and presented algorithms that find for n sites a plane labeling in $O(n \log n)$ time and a plane labeling minimizing the total leader length in $O(n^{2+\delta})$ time for any $\delta > 0$; for

details refer to Sections 5.2.4 and 5.2.2, respectively. Kindermann et al. [KLW14] considered boundary labeling for straight-line leaders in the context of annotating texts. They placed the labels either in the left or right margin of the text. For uniform labels they utilized the approach by Bekos et al. [BKS07]. In contrast to \circ -leaders running in between the text lines, straight-line leaders intersect and easily clutter the annotated text. Fink and Suri [FS16] presented dynamic programming approaches that can deal with obstacles, which may not be crossed by leaders. However, their approaches have asymptotically high running times, e.g., for the one-sided case they obtained a running time of $O(n^{11})$ and for the two-sided case they obtained a running time of $O(n^{27})$ when minimizing the total leader length. Barth et al. [BGNN18] used an integer linear programming formulation to create one-sided boundary labelings with straight-line leaders (Section 5.2.5). They used these labelings in a user study on the readability of different leader types. They particularly showed that straight-line leaders perform well when the user is asked to associate labels with their sites and vice versa. However, concerning the aesthetic preferences the users favored \circ -leaders and \diamond -leaders.

Ali et al. [AHS05] suggested circular contours and contours that mimic the silhouette of the illustration (e.g., by a buffered convex hull). In an initial procedure the labels are placed along these contours such that they are stacked and no label-label overlaps occur. Afterwards the positions of the labels are switched until all leader-leader intersections are resolved. In a final compaction step the requirement of being placed on a contour is relaxed and the labels are pushed towards the illustration to reduce the lengths of the leaders. In a companion paper Götzelmann et al. [GAHS05] described a system for labeling 3D models with internal and external labels. In this system the external label placement is done by the approach by Ali et al. [AHS05]. Hartmann et al. [HGAS05] presented a high-level description of that system. The work by Götzelmann et al. [GAHS05] is not listed in Table 1 since it focuses on internal label placement, but the work by Hartmann et al. [HGAS05] gives a high-level description of the same system.

Battersby et al. [BSF*11] proposed *ring maps* considering a circle as contours. Uniform labels are evenly distributed around the circle such that each label is oriented with the ray that emanates from the center of the circle and goes through the center of the label. Hence, the labels are not axis-aligned but radially aligned. The order of the labels around the circle is chosen such that the resulting labeling is plane. Fink et al. [FHS*12] also considered circles as contours in the context of excentric labeling. In contrast to related work [FP99, HGK14], the labels are required to touch the circle. For all presented algorithms they gave formal guarantees on the drawing criteria. In particular they presented dynamic programming approaches that support general weighting functions. Haunert and Hermes [HH14] considered a closely related setting using radial leaders but horizontal aligned labels. They reduced the problem to finding maximum weight independent sets in conflict graphs of the labels, which they solved by means of dynamic programming.

Similar to Ali et al. [AHS05], Čmolek and Bittner [ČB10] considered rectangles, circles, and convex hulls as pre-defined labeling contours. For each area to be labeled they generated a set of candidate callouts that differ in the chosen sites and the chosen position of the labels on the contour. Using a rating of the candidates based on

fuzzy logic, they utilized a greedy algorithm that selects a candidate for each feature. In a successive work Čmolík and Bittner [ČB18] adapted these techniques to label *ghosted views* of 3D models, i.e., views that use transparency to show occluded parts of the models. They particularly focused on the placement of the sites on salient parts of the model. They evaluated the approach in an extensive user study. Similarly to the user study by Barth et al. [BGNN18], the users were asked to associate labels with their highlighted features. The stimuli were created by their approach as well as by experts.

Bruckner and Gröller [BG05] presented an approach for labeling 3D models using a simple iterative algorithm that places the labels along the convex hull of the projected model. Overlaps are resolved by moving the labels along the contour and by exchanging the positions of overlapping labels. If not all overlaps can be resolved, a greedy algorithm selects the labels by priority excluding overlapping labels. Intersections between callouts are resolved in a post-processing step based on the approach presented by Bekos et al. [BKSW07]. Niedermann et al. [NNR17] used convex hulls to prescribe the labeling contours of illustrations in atlases of human anatomy. They presented a general dynamic programming approach that labels a given set of sites with respect to a set of hard and soft constraints. They proved that the calculated solution respects all hard constraints and is optimal with respect to all soft constraints. In contrast to other dynamic programming approaches, the cost-function not only supports the rating of single leaders, but also of consecutive leaders. Thus, they could incorporate drawing criteria such as distances between labels and monotonically increasing angles between leaders.

Götzmann et al. [GHS06a] considered external labeling in interactive 3D visualizations. They used the *orbit* of a figure as contour, i.e., a hull around the figure such that each point on the hull has the same distance to the figure. They computed the initial positions of the labels using a simple greedy algorithm. The movement of the labels and the sites is controlled by a multi-agent system such that for each label there are agents, i.e., local strategies, that control the behavior of the labels. They paid special attention to temporal coherency and combined their approach with internal label placement.

Free label placement Fekete and Plaisant [FP99] introduced *ex-centric labeling*. For straight-line leaders they presented a simple approach that stacks the labels on the left and right hand side of the focus region either according to their vertical or to their horizontal order. To improve temporal coherence, labels are only placed when the focus region is not moved. Heinsohn et al. [HGK14] also considered excentric labeling for straight line leaders, but in contrast to Fekete and Plaisant [FP99] the labels are displayed all the time. Heinsohn et al. suggested four approaches to place labels: (i) an approach that places all labels on a stack on the left hand side of the focus region, (ii) a radial approach, i.e., the leader of each label is part of the ray that emanates from the center of the focus region and goes through the point feature of the labels, (iii) a force-based approach that prefers labels with radial leaders (see Section 5.1.2), and (iv) a *cake-cutting* approach that places the labels equally distributed around the focus region. In case that not all overlaps can be resolved, the focus region is shrunk to reduce the number of point features within the focus region.

Götzmann et al. [GHS06b] extended the system by Hartmann et

al. [HGAS05] by supporting grouping of labels. More precisely, for each group of features they first compute the sites of the features and determine the centroid of those sites. For each centroid they create rectangle candidates in the labeling region that can host the grouped labels. Within a group the labels are vertically aligned. Vollick et al. [VVAH07] formalized the external labeling problem as an energy function that is composed by non-linear terms each representing one drawing criterion. The exact weighting of the terms was learned from existing visualizations (e.g., from hand-made drawings) by applying *nonlinear inverse optimization*. The authors used simulated annealing to find a locally optimal labeling from the energy function.

While those approaches make use of a continuous labeling region, the following discretize the labeling region to find appropriate label positions. Fuchs et al. [FLH*06] partitioned the labeling region into small rectangular regions. Using a simple greedy algorithm the labels were assigned to these regions such that the length of the leaders is minimized. Apart from labeling all features of the illustration, they also considered a *labeling lens*, which is conceptually identical to the focus region considered in excentric labeling. In this case the same approach was applied but restricted to that region. Further, the approach was also used for more general settings, where the image region consisted of multiple parts. Wu et al. [WTLY11] used external labeling to annotate a metro map with station names and photos. To that end, they underlied the metro map with a fine raster describing the free and occupied regions of the map. Only cells that are not occupied can be used for the label placement. They applied a genetic algorithm combined with a greedy algorithm to determine the position of the labels. In particular, they placed textual labels and image labels independently allowing label-leader intersections between the two groups. Gemsa et al. [GHN15] considered the special case of panorama labeling, e.g., for labeling the skyline of a city. They showed that in the proposed model it suffices to consider a finite set of candidates per feature to cover the optimal intersection-free labeling. Further, they presented a dynamic programming approach and mixed integer linear programming formulations for computing optimal labelings with respect to different cost functions and settings.

Non-strict external labeling Hartmann et al. [HAS04] presented an approach for labeling 3D objects. As they considered complex geometric features their approach first determines the location of the sites by shrinking the feature until a single point remains. They argued that this point is an appropriate choice of a non-convex object, because it is “placed at a visual dominant region”. For the actual labeling they presented an approach that consists of two phases. In the first phase an initial layout is computed using a force-based approach (see Section 5.1.2). In the second phase, possible overlaps between labels are minimized using a combination of the same force-based approach and a greedy algorithm. While an external labeling is preferred, labels may overlap the image region. Tatzgern et al. [TKS13] adapted the approach by Hartmann et al. [HAS04] to label 3D explosion diagrams. They first computed an initial layout by forming clusters of labels with similar texts, and placing each label using the first phase of the approach presented by Hartmann et al. [HAS04]. Then, in a second step they selected for each cluster its best label with respect to a set of optimization criteria. Afterwards, they applied the second phase of Hartmann et al. [HAS04]

to minimize overlaps. They applied their algorithm in a dynamic 3D scene. However, they did not enforce a strict separation between labeling and image region such that the labels could overlap the labeled 3D object. To ensure temporal coherence, the order of the labels and sites is fixed during the movement of the camera. Afterwards, when the camera stops possible intersections are resolved. Tatzgern et al. [TKGS14] presented another force-based approach for 3D objects, which they called *hedgehog labeling*. The labeling is done in 3D space using 3D poles instead of straight-line leaders. The forces are used to resolve overlaps between labels. Again the labels are not clearly separated from the labeled object, but the forces are defined such that labels occluding the image are avoided. Pick et al. [PHTP*10] also placed the annotations within the 3D space. They placed the labels directly at their point features and used a force-based approach to adapt the labels dynamically. They did not explicitly exclude intersecting leaders, but they used forces that prevent label occlusion. Further, they introduced a force that reduces oscillating movements of labels.

Stein and Décoret [SD08] formalized the external labeling problem of 3D models in the 2D projection space as an energy function whose minimum yields a labeling with short leaders, crossing-free callouts, and labels that avoid the image region. They used a greedy strategy to find a locally optimal labeling. They improved temporal coherence by penalizing positions of labels that lie too far away from the previous positions. They also avoided flickering effects that occur when sites become invisible during interaction.

In the context of excentric labeling Balata et al. [BČM14] presented an approach for labeling moving unmanned aerial vehicles. They did not explicitly distinguish between image and labeling region, but they used a force-based approach that separates the labels from their sites. This approach ensures that leaders do not intersect, but labels may overlap each other. To improve temporal coherence they considered forces that soften distracting flickering effects, and the possibility of temporally freezing labels. In a conducted user study they showed that freezing labels outperformed the other approach in terms of users' error rates.

We observe that two approaches [MP09, MTSP12] of non-strict external labeling also considered labeling contours. If necessary (e.g., due to missing free space), Mühler and Preim [MP09] relaxed the requirement on placing labels along a contour, but allowed the labels to overlap unimportant structures of the image. To that end, they extended the approach of Ali et al. [AHS05] such that it also supports the annotation of ghost views used in surgical planning. Further, they conducted a user study in which the participants were asked to judge 24 images using different drawing styles. Among others, the user study showed that the participants preferred grouped labels. Mogalle et al. [MTSP12] placed the labels inside of the image boundary such that they touch the boundary. They presented both a greedy algorithm that selects labels candidates and an approach that shifts the labels along the contours. In order to avoid overlaps with other pictorial elements (e.g., legends) they relaxed the requirement that the labels are placed along a rectangle. Instead, they integrated pictorial elements into the shape of the contour.

4.2. Polyline leaders

Most of the polyline leaders in the literature have a schematic and tidy appearance caused by the restriction to segment orientations that are aligned with the Cartesian coordinate axes or possibly with their two bisecting diagonals. Moreover, almost all work on polyline leaders requires crossing-free leaders and considers the boundary labeling setting, where the contour C_A is a rectangle. Again, we first give an overview (Section 4.2.1) of the existing literature on labeling with polyline leaders in terms of several characteristic properties and grouped by different types of input specifications. This is followed by a more detailed discussion of the individual papers and their contributions (Section 4.2.2).

4.2.1. Overview

We first group the polyline-leader algorithms by their input specifications and then classify each paper along eleven properties as listed in Table 2 and briefly discussed in the following.

Input specifications The high-level grouping of the relevant literature as seen in the top row of Table 2 first considers three types of static external, mostly boundary labeling problems. The majority of papers consider point feature labeling, some of which assume the points to be word positions in line-based text documents. While in the former case, labelings are usually unconstrained in the background and feature layers, text labeling often requires leaders that run horizontally between two text lines and may use a track-routing area in a narrow strip between text and labels to move vertically. Another set of papers studies labelings of area features, in which each leader's site within the feature can still be chosen freely by the labeling algorithm. The next group of papers considers mixed labeling problems, which combine the placement of internal labels with external labels. Here internal labels form obstacles that restrict the feasible positions of the leaders. Next, the group of dynamic labeling algorithms take into account point features that move over time as the user may shift a focus region or change the zoom level. While a static labeling algorithm usually does not preserve temporal coherence when applied to dynamic inputs, an algorithm for a dynamic labeling problem, conversely, can easily be applied to a static input. Finally, three papers consider many-to-one (X-to-1) labeling problems, where a single label refers to multiple features.

P1: leader type One of the primary properties to distinguish different labeling algorithms and their results is the used leader type. Most polyline leaders found in the literature use orthogonal segments and are either of type p_0 (P1.1) with a single bend inside the image region or they are of type o_{p_0} (P1.2) with two bends in a narrow track routing area outside the image region but inside the contour C_A . Several papers employ orthodiagonal leaders, usually of type d_0 or p_d with a turning angle of $\pm 45^\circ$ (P1.3). These are shorter than the respective p_0 -leaders and have been empirically confirmed to be aesthetically more pleasing, yet slightly less readable than p_0 -leaders [BGNN18]. Few papers do not use fixed slopes for all polyline leader segments, but rather include segments of type s between the site, the bends, and the label (P1.4).

P2: objective The most common objective function (see also Section 3.2) for polyline leaders is the total leader length (P2.1). A

	Points							Text			Areas			Mixed			Dynamic			X-to-1			Σ				
	[BKS07]	[BKPS06a]	[BHK09]	[BKS08]	[BKNS10]	[HPL14]	[KNR*16]	[BCK*18]	[BGN18]	[LWY09]	[KLW14]	[LPT*11]	[BKPS10]	[FS16]	[YDGM17]	[LN10]	[BKPS11]	[LNS16]	[FP99]○	[AHS05]	[BRL09]○	[NPS10]		[LKY08]	[Lin10]	[BCF*15]	
P1 leader type																											25
P1.1 po-leaders	x	x	x			x	x	x		x		x	x			x					x		x	x	x	x	15
P1.2 opo-leaders	x	x		x		x		x	x	x	x	x	x											x	x		14
P1.3 do- and pd-leaders			x		x								x	x													6
P1.4 other									x												x		x				3
P2 objective																											
P2.1 length minimization	x		x	x	x	x	x	x	x	x	x	x	x								x	x		x	x	x	19
P2.2 bend minimization	x			x		x		x			x													x	x		6
P2.3 multicriteria														x							x	x	x				4
P2.4 other		x	x				x		x						x	x	x							x	x	x	10
P3 contour																											
P3.1 boundary 1-sided	x	x	x	x	x	x		x	x	x	x		x	x	x	x					x		x	x	x	x	19
P3.2 boundary 2-sided (opposite)	x		x	x	x	x			x	x		x	x			x						x		x	x	x	14
P3.3 boundary 2-sided (adjacent)							x	x																			2
P3.4 boundary 3-sided							x	x																			2
P3.5 boundary 4-sided	x				x		x	x				x												x			6
P3.6 non-rectangular																x		x	x	x							4
P4 label type																											
P4.1 uniform labels	x	x	x	x	x	x	x	x			x	x	x	x	x	x	x				x		x	x	x	x	21
P4.2 non-uniform labels	x	x		x	x	x			x	x	x		x								x		x	x			12
P5 label positions																											
P5.1 fixed positions	x		x		x		x	x			x	x	x	x		x								x	x		13
P5.2 sliding positions		x		x		x			x	x			x		x	x	x	x	x	x						x	13
P6 label ports																											
P6.1 fixed ports	x	x		x	x	x	x	x	x	x	x		x	x	x	x	x	x	x	x	x	x	x	x	x	x	23
P6.2 sliding ports	x	x	x	x	x	x	x		x	x	x		x			x											13
P7 leader obstacles							x		x	x			x		x	x	x										7
P8 contributions																											
P8.1 formal proofs	x	x	x	x	x	x	x		x		x	x	x		x	x	x							x	x	x	19
P8.2 implementation		x	x					x	x	x		x		x		x					x	x	x	x	x		14
P8.3 user study								x						x													3
P9 algorithm type																											
P9.1 exact	x	x	x	x	x	x	x	x	x	x	x	x	x		x	x	x							x	x	x	21
P9.2 approximation																x								x			2
P9.3 heuristic									x					x							x	x	x	x			6
P10 techniques																											
P10.1 dynamic programming	x	x	x			x	x	x	x	x	x		x		x		x									x	13
P10.2 plane sweep	x		x		x					x	x					x									x		8
P10.3 weighted matching	x				x								x											x		x	5
P10.4 scheduling				x	x	x					x																4
P10.5 init and improve																		x	x	x							3
P10.6 mathematical programming								x						x		x											3
P10.7 meta heuristics									x																x		2
P10.8 greedy																								x			1
P10.9 other algorithm	x							x								x								x			4
P10.10 NP-hardness	x	x		x	x	x					x		x											x		x	9
P11 community																											
P11.1 algorithms	x	x	x	x	x	x	x		x	x		x	x		x	x	x							x	x	x	17
P11.2 visual computing								x	x			x		x							x	x	x	x			8

Table 2: Properties of the approaches for polyline-line leaders. References are partitioned into six top-level groups by feature type (points, text, areas), use of internal and external labels (mixed), features changing over time (dynamic), and many-to-one labeling (X-to-1). References considering excentric labeling are marked with “○”.

small leader length implies a short distance between site and label and thus optimizes the locality of labels. The criterion also minimizes the ink used to plot the leaders, and thus also the amount of overplotting of the image region. A number of papers further considers bend minimization as an objective (P2.2). Note that a leader without bends is usually also a shortest possible leader, whereas for leaders with bends no such claim about their length can be made. Further, some methods use a multicriteria objective function (P2.3) taking into account leader lengths, but also separation between callouts and displacement of features or labels [YDGM17, FP99, AHS05, BRL09]. Other objective functions (P2.4) are either arbitrary quality measures based on a single leader [BHKN09], finding a feasible labeling at all [LWY09], maximizing the number of labeled features [KNR*16], maximizing the label size [BKPS06a], maximizing the number of internal labels in a mixed model [LN10, BKPS11, LNS16], or minimizing the number of leaders in many-to-one labeling [BCF*15, Lin10, LKY08].

P3: contour Most of the literature on polyline leaders considers boundary labeling, where the contour is the bounding box of the image region R . The distinction within different algorithms is on the number of sides of R that can be used for attaching labels. The simplest and most-studied case is 1-sided boundary labeling (P3.1). From an algorithmic point of view the 1-sided problem reduces to finding the best ordering of the labels on that side. This is a less complex situation than multi-sided labelings, where one needs to determine a side assignment together with an ordering for the labels on each side. Interestingly, 2-sided settings having labels on two opposite sides (P3.2), say left and right of R , are well studied in 14 papers and admit (so far) more efficient algorithms than the case of labels on two adjacent sides (P3.3) (or three sides (P3.4)) of R , say top and right, with only two references in the literature. In applications with labels being names or other short text strings, using the two opposite sides left and right of the image region is more natural, as the height of a horizontal text label is small and thus many such labels can fit. The top and bottom sides, however, offer space for much fewer labels as the widths of the labels quickly add up to the image width, unless the labels are rotated by 90° . Only if the labels have aspect ratio closer to 1, e.g., for icons or acronyms, or are rotated accordingly, then all four sides (P3.5) are equally well suited for placing labels. In fact, placing labels on all four sides of R has been studied in six papers. Few papers study polyline leaders in the context of other contours (P3.6) like circles in excentric labeling [FP99, BRL09], convex polygons [LNS16, AHS05], or arbitrary silhouettes [AHS05].

P4: label type We distinguish two types of labels. Uniform labels (P4.1) are the most commonly studied (21 references) and all have the same size (or at least the same height as found in single-line text labels) and thus can be represented as a set of empty boxes distributed along the contour. Algorithmically speaking one has to find a matching that assigns each feature to one of these empty boxes and the label text is guaranteed to fit into the box. The more general case of non-uniform labels (P4.2, 12 references) is usually more difficult to solve as the placement of the labels now depends on the ordering of the leaders around the contour, e.g., imagine a large label consuming more space and pushing neighboring labels aside. Many results on non-uniform labels comprise NP-hardness proofs relating

label orderings to job scheduling problems (see Section 5.3). Still, for a number of problem settings, efficient algorithms exist even for non-uniform labels [LWY09, KLW14, LPT*11, NPS10, HPL14].

P5: label positions The label positions are either fixed on the contour (P5.1) or they can slide along it (P5.2). The former case is more restrictive but also simplifies the algorithmic problem in many cases as one degree of freedom is removed. Fixed positions are mostly used for uniform labels (P4.1), since for non-uniform, individually sized labels fixed positions already dictate the assignment of features to labels. In combination with property P6.1, fixed-position labels give rise to leaders whose geometry is fully specified once a feature is assigned to a label. On the other hand, in some situations sliding labels are needed to compute feasible labelings at all, e.g., when labels of non-uniform size can be stacked or when leaders are blocked from reaching certain label positions due to obstacles.

P6: label ports Label ports can either be at fixed positions (P6.1) or they can slide along one side of the label box (P6.2). Again, fixed port positions (in combination with fixed reference points (P5.1)) are algorithmically often easier to handle because the geometry of leaders is fully determined by the position of the site and the position of the fixed port. Moreover fixed ports, e.g., centered at the label, often produce tidier labelings as in consequence the distance between neighboring leaders is uniform. Sliding ports may produce shorter leaders, e.g., when connecting a site that sits below a label to the bottom corner of its label rather than to its midpoint. Many algorithms even work for both types of ports.

P7: leader obstacles Obstacles are areas in the image region that cannot be crossed by a leader. These can be important parts of the image, the individual words in a text document, or internal labels in case of mixed labeling. Obstacles are usually modeled as polygons. They restrict the options of connecting a feature to a label position and thus increase the difficulty of the algorithmic problem, e.g., they can make an instance infeasible. A special situation arises in mixed labeling [LN10, LNS16, BKPS11], where the algorithm can decide whether an internal label, i.e., an obstacle, is placed.

P8: contributions We distinguish three (non-exclusive) types of results that are found in the literature on polyline leaders. Most of the collected work provides exact algorithms (see also property P9.1) accompanied with formal proofs of correctness, optimality of the solution, and asymptotic running times (P8.1). Not all of these algorithms have been implemented and some even have running times that are too high to be of practical relevance. On the positive side, they provide certificates that the corresponding problems are polynomial-time tractable. Eight papers provide formally analyzed algorithms together with proof-of-concept implementations (P8.2). Some other papers do not contain formal results and instead report about implemented systems and algorithms for polyline leaders that are justified by case studies or experiments. Finally, three papers include results of user studies on the computed labelings (P8.3).

P9: algorithm type Here we consider three different types of algorithms found in the literature. Most algorithms (21 references) are exact in the sense that they provide mathematically optimal solutions to their respective labeling problems (P9.1). Approximation

algorithms (P9.2), found in two papers, cannot guarantee to find optimal solutions, but they are accompanied by formal proofs that provide an approximation guarantee that quantifies how far a solution can be from the optimum in the worst case. Finally, heuristics (P9.3) may compute good and sometimes optimal solutions, yet no formal proofs and guarantees are provided in the respective papers.

P10: techniques This property refers to the algorithmic techniques and principles used in the different papers. The main features and some generic descriptions of these techniques are sketched in more detail in Section 5. Most algorithms use dynamic programming (P10.1), which is based on subdividing a problem instance into two (or more) independent subinstances and then creating a solution from compatible subsolutions. Especially 1-sided boundary labeling problems can also be solved by sweep-line algorithms (P10.2) that visit the sites and labels in a suitable sequence, say top to bottom, and compute the leaders at discrete events, e.g., when new points and labels are reached. The weighted matching technique (P10.3) is well suited for problems optimizing over a finite set of possible site–label assignments modeled as a weighted bipartite graph, e.g., for instances with fixed label positions, fixed ports, and uniform labels. A scheduling algorithm (P10.4) is used for the 1-sided labeling of collinear points [BKS08] and the scheduling technique is also the basis for a number of NP-hardness reductions [BKNS10, LPT*11, HPL14]. Further, some papers use global optimization techniques such as SAT solving [BKPS11] and mathematical programming (P10.6), which includes integer linear programming [BGNN18, BKPS11] and quadratic programming [YDGM17]. Other papers use meta heuristics (P10.7) such as genetic algorithms [LWY09] and simulated annealing [Lin10] or heuristics that first construct an initial solution and then improve it by local modifications (P10.5) [FP99, BRL09, AHS05]. A greedy algorithm (P10.8) is used for a crossing-minimization heuristic in many-to-one labeling [LKY08]. Among other techniques (P10.9) are algorithms that resolve leader crossings by swapping some site–label assignments that do not affect the sum of leader lengths [BKSW07], recursive approaches [BKPS11], and a reduction to finding weighted independent sets in outerstring graphs [BCK*18]. Finally, some papers contain NP-hardness reductions (P10.10, see also Section 5.3), mostly for labeling problems with non-uniform labels. These are used to justify heuristics and approximation algorithms or simplifications of the model, e.g., the use of uniform labels.

P11: community This property is identical to S8 for straight-line leaders. Interestingly, polyline leaders are studied primarily in the algorithms community (P11.1, 17 references), but also in several papers in the visual computing community (P11.2, 8 references).

4.2.2. Detailed discussion

We start our discussion with two types of orthogonal leaders, namely p_o -leaders (Section 4.2.2.1) and $o_p o$ -leaders (Section 4.2.2.2). Then we discuss 1-bend orthodiagonal leaders (Section 4.2.2.3). Since the vast majority of the literature considers rectangular contours, i.e., the boundary labeling problem, we assume that the contour C_A is an axis-aligned rectangle, unless stated otherwise.

4.2.2.1. p_o -leaders Orthogonal p_o -leaders are the simplest and most commonly used polyline leader type in external labeling. They

are frequently found in the algorithmic literature, but also appear in professional information graphics and are implemented in some practical labeling methods.

1-sided boundary labeling In the 1-sided setting, one of the simplest boundary labeling problems is to assign external labels of uniform size with fixed ports and thus fixed reference points aligned on a single, say vertical, side of the image region R to a set of n point features in R such that the p_o -leaders are crossing-free. The most commonly used objective function is to minimize the total leader length. For this problem Bekos et al. [BKSW07] presented an $O(n^2)$ -time algorithm to find a length-minimal solution. Their algorithm runs in two phases. First, sites and labels are matched so that they have the same vertical order. Second, leader crossings are iteratively resolved without changing the total leader length. Later Benkert et al. [BHKN09] presented an algorithm based on the sweep-line paradigm that visits all sites and labels in their vertical order. Whenever a new label becomes available, a site is selected such that their leader is guaranteed to be crossing-free and the total length remains minimum. Their algorithm improved the running time to $O(n \log n)$. A very different method to compute length-minimal solutions is a simple integer linear program (ILP) proposed by Barth et al. [BGNN18]. This approach generalizes to arbitrary leader shapes and label positions, as it basically computes a minimum-length crossing-free perfect matching between the given sites and fixed label positions. While in general ILP solving does not scale well, it is still a practical approach for smaller instances. Nöllenburg et al. [NPS10] then relaxed the restriction to fixed reference points with uniform labels and allowed sliding reference points and non-uniform labels. They presented another sweep-line algorithm that computes a length-minimal and crossing-free solution for this problems in $O(n \log n)$ time. Huang et al. [HPL14] additionally relaxed the constraint to have fixed-port labels and allowed sliding ports. In this model, they proposed an $O(n^3)$ -time algorithm for minimizing the total leader length. A heuristic framework for external labeling has been proposed and implemented by Ali et al. [AHS05]. They considered rectangular contours as well as other silhouettes and use rectilinear leaders as one option. The algorithm first determines suitable label positions, e.g., by projecting each site to the closest point of the contour or using force-based methods, and then tries to iteratively resolve leader crossings until a crossing-free solution is found. While this can produce reasonable labelings, no guarantees on optimality or termination are given.

Fink and Suri [FS16] proposed a collection of dynamic programming algorithms for boundary labeling with obstacles. Each obstacle is modeled as a rectilinear polygon and no leader may cross any obstacle. With their algorithm for 1-sided boundary labeling with p_o -leaders and uniform labels, they could minimize the total leader length in $O(n^4)$ time if label positions are fixed and in $O(n^7)$ time otherwise. Their method also extends from point to area features at the cost of another linear factor. If, however, the labels have non-uniform size, they proved NP-hardness for deciding the existence of a crossing-free solution. A more restricted case of boundary labeling with obstacles has been studied by Löffler and Nöllenburg [LN10]. They made the simplifying assumptions that all obstacles are congruent rectangles (which may model a set of internal labels). They provided dynamic programming algorithms for placing non-crossing

p -leaders such that the number of intersected obstacles is minimized. For some cases, e.g., when the obstacles can intersect each other or leaders may cross, they also proved NP-hardness results.

For objective functions other than length minimization Benkert et al. [BHKN09] proposed a general $O(n^3)$ -time algorithm using a dynamic programming approach that optimizes the labeling using uniform, fixed-position labels with sliding ports. Kindermann et al. [KLW14] generalized this algorithm to instances with non-uniform labels using rasterized label positions, where one label may occupy multiple raster slots. This increases the running time to $O(n^4 m^3)$ and the memory space to $O(n^3 m^2)$, with $m > n$ being the number of raster slots. Kindermann et al. used this algorithm for text labeling and made some adaptations to preferably route the p -segments between two lines of text rather than striking them out. Another interesting variation is 1-sided multi-stack labeling [BKPS06a], where labels take positions in two or more label stacks on the same side of C_A . However, using p -leaders the authors only showed that for non-uniform labels it is NP-hard to decide the existence of a crossing-free two-stack labeling. Feasible algorithms in this model, e.g., for uniform labels, are missing.

Finally, two papers considered many-to-one boundary labeling, where sets of point features can share the same label. Lin et al. [LKY08] considered individual p -leaders and labels with multiple ports, one for each feature with that label. They proved that minimizing crossings in 1-sided many-to-one boundary labeling with p -leaders is NP-hard, but they also presented a greedy heuristic iteratively assigning label positions with locally fewest leader crossings. Bekos et al. [BCF*15] investigated a different labeling style using hyperleaders (recall Section 2.2). Here a p -hyperleader for k sites consists of k vertical p -segments connecting the sites to a single horizontal o -segment, which connects to the label. They considered three problem variants. For crossing-free labelings, the problems are to minimize the number of labels or the total leader length, and for labelings with crossings, the problem is to minimize their number. Both problems with crossing-free leaders can be solved in polynomial time using dynamic programming. Crossing minimization can be solved in $O(n \cdot c)$ time, where n is the number of sites and c is the number of labels, if the label order is given; if the order is unconstrained, the problem becomes NP-hard.

2-sided boundary labeling A natural extension of the 1-sided case is to admit 2-sided placement of labels, either on two opposite or two adjacent sides of the bounding rectangle. This makes the algorithmic optimization problem more complex, as the distribution of labels to either side becomes an additional degree of freedom.

For uniform labels with fixed positions on two opposite sides of C_A , Bekos et al. [BKSW07] presented an $O(n^2)$ -time algorithm for computing a length-minimal solution. Using dynamic programming (Section 5.2.1) they obtained a length-minimal assignment of sites and labels, from which crossings were resolved independently on both sides as no pair of leaders to opposite sides can intersect in a length-minimal solution. When replacing point features by a certain type of polygonal area features, Bekos et al. [BKPS10] gave an $O(n^2 \log^3 n)$ -time algorithm for length-minimal two-sided boundary labeling, in which the sites of the leaders are located along the sides of the area features. The difference of this algorithm is that the first step to compute the length-minimal site-label assignment uses the

weighted matching technique (Section 5.2.2). The crossing removal is done as before. For uniform labels with sliding positions, Huang et al. [HPL14] presented an $O(n^5)$ -time algorithm as an extension of their algorithm for the 1-sided case. The heuristic framework of Ali et al. [AHS05], already mentioned for 1-sided labeling, can be applied equally well to 2-sided instances.

For instances with obstacles, the dynamic programming algorithms presented by Fink and Suri [FS16] for minimizing the total length of a labeling with 1-sided p -leaders can be extended to 2-sided instances, at the cost of a significant increase in time complexity from $O(n^4)$ to $O(n^9)$ for fixed label positions and from $O(n^7)$ to $O(n^{15})$ for sliding label positions.

For arbitrary objective functions, Benkert et al. [BHKN09] generalized their dynamic-programming algorithm for 1-sided p -leader labeling to two-sided boundary labeling, yet with a worst-case running time of $O(n^8)$. The reason for the algorithm's much higher complexity compared to $O(n^2)$ for length minimization [BKSW07] is that for non-length-minimal solutions the sites can no longer be easily split into two independent subproblems for the two opposite sides. The extension of Kindermann et al. [KLW14] for non-uniform labels in 1-sided boundary labeling is not properly generalized to optimize 2-sided instances. Instead, for their text labeling application, the authors proposed to simply split the sites at the horizontal median position into two balanced parts and then run the 1-sided algorithm for each side independently.

Many-to-one labeling has also been studied for the 2-sided setting with labels on opposite sides. Since the crossing minimization problem considered by Lin et al. [LKY08] remains NP-hard for the 2-sided case, they again described a heuristic algorithm to obtain a labeling, in which all features sharing the same label connect with their own p -leader to a single label box on one of the two opposite boundaries of C_A . Their algorithm constructs a weighted graph, in which they (heuristically) determine a vertex bisection of minimum weight, which induces the partition of labels to the two boundary sides. The weights are chosen such that the computed bisection has few leaders possibly crossing each other. Then each side is solved with their 1-sided algorithm. The hyperleader labeling by Bekos et al. [BCF*15] is also considered in a 2-sided setting, but using duplicated labels that are vertically aligned on both sides. This actually makes the problem algorithmically simpler because now each hyperleader spans the entire width of the image region and splits the instance into two parts. Therefore the algorithms for crossing-free labelings with minimum number of labels or minimum length are improved to $O(n)$ and $O(n^2)$ running times, respectively. The crossing minimization problem that was shown to be generally NP-hard for the 1-sided case remains open for 2-sided labelings. For a fixed label order, the polynomial-time 1-sided algorithm for crossing minimization can also be applied to the 2-sided case.

A different 2-sided setting is to have labels on two adjacent sides of C_A , e.g., the top and the right side. Here it is no longer the case that a crossing-free solution always exists. Kindermann et al. [KNR*16] presented an $O(n^2)$ -time dynamic programming algorithm to compute a valid crossing-free solution if one exists, both for labels with fixed or sliding ports. Further, they presented an algorithm to maximize the number of labeled sites in a crossing-free labeling in $O(n^3 \log n)$ time. Finally, they modified their algorithm so that it

can minimize the total leader length rather than computing some crossing-free solution, however, at the cost of an increase of the time complexity to $O(n^8 \log n)$ and a space bound of $O(n^6)$. Recently, Bose et al. [BCK*18] improved this to an $O(n^3 \log n)$ -time algorithm, again using the dynamic programming technique. An interesting algorithmic open problem is to investigate the remaining difference in the $O(n^2)$ running time for length minimization with labels on two opposite sides of C_A and the $O(n^3 \log n)$ running time for labels on two adjacent sides.

Multi-sided boundary labeling Three-sided labeling is explicitly considered only in two papers [KNR*16, BCK*18]. Kindermann et al. [KNR*16] showed that this problem can be reduced to splitting an instance into two 2-sided subproblems. This results in an $O(n^4)$ -time algorithm working in linear space for finding a crossing-free solution (if one exists) or for maximizing the number of labeled features. For 4-sided instances, they proposed an algorithm that considers all possibilities to split into two special 3-sided instances to be solved as before. This results in an increase of the running time to $O(n^9)$. Bose et al. [BCK*18] improved upon this by presenting new algorithms for 3- and 4-sided boundary labeling that can even find length-minimal or bend-minimal solutions, for fixed and sliding ports and optionally with the leaders avoiding a set of obstacles. This is achieved by reducing the boundary labeling problem to a weighted independent set problem in outerstring graphs, for which a recently published algorithm is used [KMPV17]. From a heuristic perspective, again the labeling framework of Ali et al. [AHS05] is flexible in terms of the contour C_A and can thus be applied to place labels at three or four sides of a rectangle, as well as around a more complex silhouette contour of the image.

4.2.2.2. o_{po}-leaders Orthogonal o_{po}-leaders have a slightly more complex shape than the corresponding p_o-leaders since they contain an additional segment. As a result, the readability of the labelings obtained by using such leaders decreases with respect to the ones obtained by using s-, p_o- or d_o-leaders, as observed by Barth et al. [BGNN18]. On the positive side, however, the central p-segment of each o_{po}-leader lies outside the image region R (which, for this particular type of leaders, is usually assumed to be rectangular, as otherwise the readability of the obtained labelings further decreases) routed within a so-called *track routing area* that is wide enough to accommodate all p-segments. This implies that o_{po}-leaders are particularly useful for text annotation purposes [LWY09, KLW14], since the first o-segment of each o_{po}-leader, which is the only one that interferes with the underlying text, can be drawn between the lines of the text.

1-sided boundary labeling In the 1-sided setting, computing a feasible labeling, in which no two o_{po}-leaders intersect, is an easy task. Assuming without loss of generality that the external labels are attached on a vertical side of the image region R , we observe that the vertical order of the sites must be identical to the vertical order of their corresponding labels. This observation directly gives rise to an $O(n \log n)$ -time constructive algorithm [BKSW07]. For the problem of finding labelings, in which the total number of bends is minimized, Bekos et al. [BKSW07] proposed an $O(n^2)$ -time dynamic programming based algorithm to find bend-minimal solutions, assuming that the ports of the labels are sliding. For the case that

the ports of the labels are fixed, Huang et al. [HPL14] proposed a slightly less efficient algorithm (which is also based on dynamic programming) to find a bend-minimal solution in $O(n^3)$ time. In the context of collinear sites, Bekos et al. [BKS08] proposed an $O(n \log n)$ -time algorithm to compute length-minimal solutions by employing a linear-time reduction to a single machine scheduling problem. Notably, their algorithm is not difficult to be adjusted to the more general 1-sided boundary labeling setting assuming that the ports of the labels are sliding; for details refer to Section 5.2.3.

For the case that rectangular-shaped obstacles (e.g., internal labels) are allowed within the image region R , which must not be crossed by the leaders, Fink and Suri [FS16] presented an $O(n^{11})$ -time algorithm to compute a length-minimal solution, under the assumption that the labels are of uniform size. Bekos et al. [BKPS11] and Löffler et al. [LNS16] studied more general mixed labelings, in which each site can be labeled either with an internal label that is located directly to its top-right or with an external label through an o_{po}-leader. When labels are of uniform height and the external labels are to the right side of the image region R , Bekos et al. [BKPS11] provided an $O(n \log n)$ -time algorithm for maximizing the total number of internal labels. Note that the case where the external labels are to the left side of R is not symmetric and it turns out to be more difficult. For this case, Bekos et al. [BKPS11] provided a quasi-polynomial algorithm of $O(n^{\log n + 3})$ time complexity, a more efficient 2-approximation and an integer linear programming formulation for maximizing the number of internal labels. Löffler et al. [LNS16] presented an improved algorithm to solve the case, in which the external labels are on the left side of R and they are connected to their corresponding sites with type-o leaders, in time $O(n^3(\log n + \delta))$, where δ denotes the minimum of n and the inverse of the distance of the closest pair of points.

1-sided boundary labelings with o_{po}-leaders have also been studied in the many-to-one setting, in which a point can be attached to more than one label [LKY08, Lin10]. Lin et al. [LKY08] observed that in this setting crossings will inevitably occur, and proved that minimizing their number is NP-hard. Finally, they presented a 3-approximation algorithm, which places the labels in “median order” as introduced by Eades and Wormald [EW94]. In a subsequent work, Lin [Lin10] allowed more than one occurrence of the same label, in order to avoid crossings that negatively affect the readability of the produced labelings. Naturally, the focus of his work was on minimizing the total number of used labels; a task which in the 1-sided case can be solved in $O(n \log n)$ time by a simple top to bottom traversal of the sites, assuming that the labels are along a vertical side of the image region R .

We conclude this subsection by mentioning that 1-sided boundary labeling with o_{po}-leaders has also been studied when more than one stack of labels are allowed along the side s of the image region R containing the labels, under the additional assumption that the labels are all of uniform size and their ports are sliding [BKPS06a]. In this setting, Bekos et al. [BKPS06a] studied the problem of finding boundary labelings, in which the height of each label is maximized, assuming that the side s is vertical. Polynomial-time algorithms are given for the case of two stacks of labels, and for the case of three stacks of labels under the additional assumption that each o_{po}-leader that connects a point to a label of the second (third)

stack must have its p -segment in the gap between the first and the second (second and third, respectively) stack of labels.

1.5-sided boundary labeling In this setting, the labels are attached along a single side of the image region R , say without loss of generality the right, as in the ordinary 1-sided boundary labeling. However, an op -leader can be routed to the left side of R temporarily and then finally to the right side, that is, its p -segment is allowed to be on the left side of R . Lin et al. [LWY09] presented an $O(n^5)$ -time algorithm, which – based on dynamic programming – computes a length-minimal solution, under the assumption that the position of the labels are fixed and their sizes are uniform; the ports of the labels can be either fixed or sliding. Lin et al. [LPT*11] extended this algorithm to the bend-minimization problem keeping the time complexity unchanged, and they presented an improved algorithm to compute length-minimal solutions in $O(n \log n)$ time.

2-sided boundary labeling In the 2-sided setting, the algorithmic optimization becomes more complex due to the additional degree of freedom raised by the distribution of the labels to two sides. For uniform labels with either fixed or sliding ports on two opposite sides of the image region R , Bekos et al. [BKSW07] presented an $O(n^2)$ -time algorithm for computing a length-minimal solution. Their approach is based on dynamic programming and exploits the fact that in order to avoid crossings, the vertical order of the sites must be identical to the vertical order of their corresponding labels at each of the two sides of R . Under the same set of assumptions, Huang et al. [HPL14], and Fink and Suri [FS16] proposed less efficient algorithms (which are also based on dynamic programming) to find bend-minimal solutions in $O(n^5)$ time, and length-minimal solutions in the presence of rectangular-shaped obstacles within the image region R in $O(n^{27})$ time, respectively. The time complexity of the algorithm by Fink and Suri [FS16] was improved by Bose et al. [BCK*18] to $O(n^9)$, under the additional assumption that the position of the labels are fixed. Note that if the labels are of non-uniform heights and their positions are not fixed, then both optimization problems mentioned above (i.e., the minimization of the total leader length and the minimization of the total number of leader-bends) become NP-hard, even in the absence of obstacles [BKSW07, HPL14]. They also remain NP-hard in the special case in which the input sites are collinear [BKS08].

In the many-to-one setting, computing a 2-sided boundary labeling with op -leaders and minimum number of leader crossings is NP-hard, even in the case in which the same number of labels must be attached on two opposite sides of the image region R and the labels are of uniform heights with either fixed or sliding ports [LKY08]. For this particular setting, Lin et al. [LKY08] suggested an algorithm with an approximation factor of at least three that also depends on the structure of the input instance. If more than one copy of the same label is allowed, Lin [Lin10] provided an $O(n^2)$ -time algorithm which yields a crossing-free routing of the leaders while simultaneously minimizing the number of used labels.

Finally, in the 2-sided setting, in which the labels are given along two adjacent sides of the image region R , Bose et al. [BCK*18] gave an $O(n^{12})$ -time algorithm to compute a length-minimal solution (if any), under the assumption that the position of the labels are fixed. Notably, their algorithm works also in the presence of obstacles.

4-sided boundary labeling In the 4-sided setting, there are significantly fewer results, and they concern boundary labelings of minimal length. To the best of our knowledge, there are no results concerning bend-minimal boundary labelings or boundary labelings that are optimal in terms of some different objective function.

Bekos et al. [BKSW07] presented a polynomial-time algorithm to compute length-minimal labelings, assuming that the position of the labels around the image region R are fixed (i.e., specified as part of the input). Their algorithm consists of two steps. In the first step, a (not necessarily crossing-free) length-minimal site-label assignment is computed by finding a minimum-cost perfect matching on an appropriately defined bipartite graph (for details refer to Section 5.2.2). In the second step, the crossings are eliminated such that the total leader length of the assignment computed in the previous step is not affected (i.e., by keeping the solution optimal). The time complexity of the algorithm is dominated by the time needed to compute the minimum-cost perfect matching in the first step and depends on whether the ports of the labels are fixed or sliding. More concretely, in the former case the bipartite graph is geometric, which allows for an efficient minimum-cost perfect matching computation in $O(n^2 \log^3 n)$ time [Vai89]. In the latter case, however, the bipartite graph has no special property, and therefore a minimum-cost perfect matching must be computed using the Hungarian method, which needs $O(n^3)$ time; see, e.g., [Kuh55].

We conclude this section by mentioning that the aforementioned algorithm was extended to the case of area features by Bekos et al. [BKPS10]. In the case of area features with constant number of corners, the computation of the minimum-cost perfect matching, which determines the complexity of the algorithm, is done on a bipartite graph that is not necessarily geometric, which implies that the Hungarian method must be used (as in the case of sliding label ports above), yielding again a time complexity of $O(n^3)$.

4.2.2.3. do - and pd -leaders Orthodiagonal leaders with one bend have a simple shape, which seems not to introduce clutter in the obtained labelings [BGNN18]. In the framework of boundary labeling, in which the image contour C_A is an axis-aligned rectangle, these types of leaders were introduced and first studied by Benkert et al. [BHKN09] back in 2009. On the other hand, if C_A is a circle, then early works that adopt the paradigm of do -leaders date back to excentric labeling for interactive labeling of focus regions [FP99].

Let us first assume that C_A is an axis-aligned rectangle. As already mentioned, Benkert et al. [BHKN09] introduced boundary labelings with do -leaders and preliminarily observed that a crossing-free solution might not always be feasible. For a general objective function, they proposed a dynamic programming approach, which in $O(n^5)$ time (in $O(n^{14})$ time) determines whether there exist an optimal crossing-free boundary labeling with do -leaders, assuming that the ports of the labels are sliding while the positions of the labels are fixed along a single side (along two opposite sides, respectively) of the image region R ; in the positive case, the labeling can be reported without increasing the time complexity. A significantly improved algorithm was given, when the objective function is the minimization of the total leader length and the labels are along a single side of R . For this case, the authors suggested an algorithm, which adopting the sweep-line paradigm determines whether there exists an optimal crossing-free boundary labeling with do -leaders in $O(n^2)$ time.

In order to overcome infeasibility issues, in a subsequent work, Bekos et al. [BKNS10] extended the study of boundary labelings with do -leaders by also incorporating od - and pd -leaders, and proved that a combination of these types of leaders guarantees the existence of a crossing-free solution. More precisely, they proved that a crossing-free solution (i.e., that is not optimal under some objective function) with od - and pd -leader can always be derived by a plane-sweep technique in $O(n \log n)$ time, assuming that the position of the labels are fixed along the image region R . In the same setting, a length-minimal solution can be found in $O(n^3)$ time by appropriately adjusting both steps of the matching-based algorithm given in [BKSW07] (see also Section 4.2.2.2).

Yang et al. [YDGM17] observed that while the algorithm by Bekos et al. [BKSW07] can ensure the absence of crossings in a length-minimal solution, it cannot guarantee adequate separation between leaders and connection sites or other leaders, which may lead to ambiguity. Hence, they proposed a quadratic program to further improve the output labelings of the algorithm of Bekos [BKSW07] and they used them in a framework, called `MapTriX`, for visualising many-to-many flows between different geographic locations by connecting an OD matrix [Voo13] with origin and destination maps [WD08]. Löffler et al. [LNS16] extended the study of boundary labelings with do -leaders in the mixed setting, where also internal labels are allowed, and provided a dynamic programming based algorithm to maximize the number of internal labels.

4.2.2.4. Other polyline leaders Few works studied external labelings with polyline leaders that are neither orthogonal nor orthodiagonal. In text annotation, Kindermann et al. [KLW14] proposed boundary labelings with os -leaders, where the first segment of each leader is used to reach the margin of the text, while the second one has a slope such that all labels can be arranged on the margin of the text. Obviously, such labelings can be derived from opo -labelings by removing the last bend of each leader. Fekete and Plaisant [FP99], and Bertini et al. [BRL09], who studied excentric labeling, also proposed site-label connections with os -leaders. Here, the first segment of each leader is orthogonal to the boundary of the lens, while the slope of the second segment is chosen such that all labels can be arranged on the boundary of the circular focus region.

4.3. Curved leaders

In this section we discuss research that uses smooth curves as leaders. In the context of automatically creating metro maps, Wu et al. [WTLY12] suggested to use 2-sided boundary labeling and curved leaders based on B-splines to annotate stations with images. They required that the resulting labeling is plane, while also minimizing the number of intersections between leaders and the metro map. At the core of their approach, they used a general flow-network model to construct such leaders. In a user study based on eye-tracking, they showed that the curved leaders outperform straight-line leaders and orthogonal leaders. Wu et al. [WPT*15] further extended this approach to 4-sided boundary labeling. To that end, they first partitioned the image region into four regions using the method by Bekos et al. [BKSW07] and labeled each region independently with the approach by Wu et al. [WTLY12]. Kindermann et al. [KLW14] used Bézier curves to connect text annotations with

their sites. Their method first computes a plane straight-line labeling. Afterwards, each leader is interpreted as a cubic Bézier curve whose control points are iteratively moved using a force-based method.

4.4. Empirical evaluation techniques

In this section we briefly discuss the state of the art on different evaluation techniques for empirically examining design decisions and evaluating labeling techniques.

The analysis of professional (handmade) drawings, e.g., technical drawings, visual dictionaries, medical illustrations, is a task that is neither complicated nor expensive and helps in extracting drawing criteria for a planned labeling technique. Niedermann et al. [NNR17] extracted drawing criteria in a semi-automatic way from atlases of human anatomy to generate similarly looking labelings. Vollick et al. [VVAH07] proposed a system that automatically learns the drawing style from existing systems. While both approaches require that drawings in the specific style already exist, the analysis of existing drawings can also be used to justify default drawing criteria satisfied in almost all labeling styles [HAS04, AHS05, HGAS05, MTSP12]. Moreover, interviews with domain experts help to justify the extracted criteria and to make further design decisions [MP09, MTSP12, NNR17, ČB18].

Certainly the most extensive evaluation technique are user studies. In order to obtain expressive results, an elaborated design, equipment and sufficiently many participants with different backgrounds are required. In research on external labeling, user studies are conducted to assess labeling approaches. Fekete et al. [FP99] investigated the question whether excentric labeling is a reasonable alternative for zooming. Čmolfík and Bittner [ČB18] compared their algorithm against label layouts created by humans. Madsen et al. [MTM*16] presented a study comparing different labeling approaches in dynamic settings; to the best of our knowledge this is the only investigation that takes multiple existing labeling algorithms into account. In contrast, Bertini et al. [BRL09], Mühler and Preim [MP09], Pick et al. [PHTP*10], Wu et al. [WTLY12] and Balata et al. [BČM14] did not compare their results with other approaches, but investigated how participants perform when using their approaches and asked them concerning the experiences made when using their systems. Instead of evaluating a concrete labeling approach, Barth et al. [BGNN18] conducted a user study to investigate general properties of external labeling, namely in this particular case the readability of different leader types. Depending on the study, the participants were asked to conduct different types of tasks. As the main purpose of a labeling is to unambiguously relate features and labels to each other, the most often proposed task is to associate labels with their features and vice versa. This is either done by reading tasks [BRL09] or by selecting particular labels [BČM14, BGNN18, ČB18]. Alternatively, the participants are asked to check the existence of particular labels [FP99] or to solve more complex tasks on the illustrations that require the use of the labeled features [YDGM17]. Only the user study by Wu et al. [WTLY12] used eye-tracking. Typical measures are the response times and error rates of the participants. Finally, except for the user study by Čmolfík and Bittner [ČB18], all presented studies take the participants' preferences and personal experiences with the labeling layouts into account. In particular Pick et al. [PHTP*10] assessed

their approach by conducting an expert walkthrough in which the participants answered few questions after using their system.

5. Labeling techniques

In this section we discuss algorithmic techniques that are frequently applied in external labeling. We split them into non-exact (Section 5.1) and exact algorithms (Section 5.2). While exact algorithms solve Problem 1 and provide an optimality guarantee for the solution, non-exact approaches refrain from such guarantees in favor of the running time. In Section 5.3 we discuss complexity results.

5.1. Non-exact algorithms

First we describe algorithmic techniques to solve Problem 1 without necessarily yielding optimal labelings with respect to a cost function.

5.1.1. Greedy algorithms

Greedy algorithms are simple local optimization strategies that iteratively construct solutions by *greedily* extending them by elements that promise the best improvement of the solution in each step. For an introduction to greedy algorithms see, e.g. [CLRS09, Chapter 17].

Characteristics

- Support general multi-criteria cost functions, but typically find only local optima.
- Hard constraints can be easily enforced, but may lead to unlabeled features.
- Easy to implement and typically fast in practice.

Sketch Greedy algorithms in external labeling can be summarized by the following scheme. For each feature a candidate set of callouts is parameterized; let \mathcal{C} denote the union of all those sets. Starting with an empty labeling \mathcal{L} , the callouts are *greedily* added to \mathcal{L} with respect to some given cost function and hard constraints. More precisely, the callout $c \in \mathcal{C}$ with lowest cost among all callouts in \mathcal{C} is added to \mathcal{L} . Afterwards, the callout c and all callouts that cannot be added to the current labeling \mathcal{L} without violating a given hard constraint are removed from \mathcal{C} . The procedure is repeated until the candidate set \mathcal{C} is empty. As result of the procedure the labeling \mathcal{L} is returned. Since in each step it is guaranteed that \mathcal{L} is a labeling of the given instance satisfying all hard constraints, the resulting set \mathcal{L} is also a labeling. However, it is not guaranteed that it is the cost-optimal labeling among all possible labelings. Further, depending on the hard constraints, in some cases it may happen that not all features are labeled, even if a complete labeling exists.

Publications: [HAS04], [BG05], [FLH*06], [GHS06a], [LKY08], [SD08], [ČB10], [WTLY11], [MTSP12], [TKS13], [HGK14], [GHN15], [ČB18].

5.1.2. Force-based approach

Force-based approaches are local optimization strategies that have been originally introduced in the field of graph drawing to construct layouts of graphs [Ead84, FR91]. According to this approach, a graph is modeled as a physical system with forces acting on it, and a good layout is obtained by an equilibrium state of the system. In the context of external labeling, the sites, labels, leaders, and features

become the *objects* in the physical system that interact with each other by forces. The approaches mainly differ in the objects that are fixed or can freely move, the definition of the forces, and how an equilibrium is found.

Characteristics

- Supports general multi-criteria cost functions, but typically find only local optima.
- Needs an initial labeling as starting point.
- Intuitive formulations as it is based on a mechanic analogy.

Sketch A common technique to search for an equilibrium of the system is to iteratively apply two steps. Firstly, compute for each object o in the system the sum f_o of the forces that act on it. Secondly, move each object o along the direction of f_o by an amount that is proportional in the magnitude of f_o and a predefined constant Δt that may change over time. If after a number of iterations the objects do not move any more or a maximum number of iterations is reached, the system is assumed to be in an equilibrium state. We note that the initial configuration of the system, i.e., the one in the first iteration, strongly influences the quality of the overall result of the procedure. For example Ali et al. [AHS05] use this iterative approach as a post-processing step to improve a layout created beforehand. They introduce forces that optimize the angles between leaders and the distance between labels and their sites. Similarly as done in simulated annealing, they start with a large Δt and decrease it with each iteration. Hence, in each iteration the possible displacement of the labels is reduced.

Another approach defines potential functions for the objects based on force fields. The positions of the objects are then chosen such that they minimize the total potential of the objects. Hartmann et al. [HAS04] introduce for each label a static force field such that the attracting force between a label and its feature increases with increasing distance. Further, a label is repelled by other objects and the border of the drawing area. The local minima of the potential function are found by using particles that move within this force field. Each local minimum is interpreted as a label candidate of the same feature. Multiple labels are placed iteratively. In each iteration a label is placed and fixed at its preferred position in its force field as described beforehand. For the remaining labels their force fields are adapted incorporating the newly placed label as repulsive force.

Publications: [HAS04], [AHS05], [HGAS05], [PHTP*10], [HGK14], [TKS13], [BČM14], [TKGS14].

5.2. Exact algorithms

In this section we describe algorithmic techniques that are used to solve Problem 1 exactly, i.e., with respect to a given cost function these techniques yield optimal labelings.

5.2.1. Dynamic programming

Dynamic programming is by far the most used technique in the context of external labeling as evidenced by Tables 1 and 2. The general idea is to find an optimal labeling using recursion originating from the input instance. In each recursive step, the optimal labeling of the currently considered instance I is composed of optimal labelings




Figure 6: Decomposition tree of an instance for external labeling. The instance (blue) enclosed by the labels ℓ_1 and ℓ_5 (top) is decomposed into sub-instances enclosed by the labels ℓ_1 and ℓ_3 , and ℓ_3 and ℓ_5 , respectively (middle row). Each of them is again decomposed into two instances, which do not contain any point features anymore.

obtained from smaller disjoint sub-instances of I . The base case of the recursion is reached when I does not contain any feature to be labeled and the optimal labeling of I is trivial. Storing the intermediate solutions in a table for subsequent look-ups ensures that each sub-solution is computed only once, which keeps the running time polynomial. For an introduction to dynamic programming see, e.g., example [CLRS09, Chapter 16].

Characteristics

- Allows general multi-criteria cost functions.
- Requires that feasible labelings are crossing-free.
- Hard constraints can be easily incorporated.
- Mostly applied in boundary labeling (see Tables 1 and 2).
- Depending on the actual setting, it tends to yield algorithms that have high asymptotic running time and storage consumption.

Sketch In the following we sketch the technique in greater detail providing a general blue print. Let $I = (D, F, M)$ be an instance of external labeling. In the remainder, we call an instance I' a *sub-instance* of I if the drawing area of I' is a sub-region of D and the features of I' are a subset of the features in F . Further, two instances are *disjoint* if their drawing areas do not intersect.

The dynamic programming approach recursively constructs an optimal labeling \mathcal{L} of I based on optimal labelings of disjoint sub-instances of I ; see Fig. 6 for an illustration. To that end, a set \mathcal{I} of sub-instances of I is defined (including I) such that I can be recursively decomposed into instances from \mathcal{I} . More precisely, two instances I' and I'' form a *decomposition* of I if

- I' and I'' are disjoint sub-instances of I ,
- each feature of I is contained either in I' or in I'' , and

- both I' and I'' contain fewer features than I .

The recursive decomposition of I is then described as a rooted, binary tree T over \mathcal{I} such that

- each node of T is an instance of \mathcal{I} ,
- each node of T has either no or two children,
- the children of a node form a decomposition of that node,
- I is the root of T ,
- the leaves of T are empty instances containing no features.

We call such a tree a *decomposition tree* of I . The actual choice of \mathcal{I} strongly depends on the concrete problem setting and typically requires fundamental insights into the geometric structure of the problem. In particular, the set \mathcal{I} is chosen such that there is a decomposition tree T with

$$c_{\text{OPT}}(I) = \begin{cases} c_I, & I \text{ is a leaf,} \\ c_{\text{OPT}}(I') + c_{\text{OPT}}(I'') + c_{I', I''} & I \text{ has children } I' \text{ and } I'', \end{cases}$$

where $c_{\text{OPT}}(I)$ denotes the cost of an optimal labeling of I , and c_I and $c_{I', I''}$ are constants. The existence of such a decomposition tree directly leads to the following classic dynamic programming approach. We introduce a table X that contains for each instance $I \in \mathcal{I}$ an entry $X[I]$ representing the cost of an optimal labeling of I , i.e., $X[I] = c_{\text{OPT}}(I)$. Since we consider a minimization problem, we can use the following recurrence to actually compute $X[I]$.

$$X[I] = \begin{cases} c_I, & I \text{ is empty} \\ \min_{(I', I'') \in \mathcal{D}_I} X[I'] + X[I''] + c_{I', I''} & \text{otherwise,} \end{cases}$$

where $\mathcal{D}_I \subseteq \mathcal{I} \times \mathcal{I}$ is the set of all decompositions of I in $\mathcal{I} \times \mathcal{I}$. We observe that it suffices to consider decompositions from $\mathcal{I} \times \mathcal{I}$, because the decomposition tree only consists of instances from \mathcal{I} . Further, we compute each table entry only once – for example in increasing order with respect to the number of contained point features. Hence, each table entry $X[I]$ is computed in $O(|\mathcal{I}|^2)$ time. Since there are $|\mathcal{I}|$ entries in X , we obtain $O(|\mathcal{I}|^3)$ running time and $O(|\mathcal{I}|)$ storage consumption for computing the costs of an optimal labeling of an input instance. Using a standard backtracking approach that reconstructs the solution based on the decomposition tree, we also obtain an optimal labeling for that instance. It is crucial to keep \mathcal{I} as small as possible, because the running time of the approach strongly depends on the size of \mathcal{I} . A common approach is to show that the drawing area of the instances can be specified by only few parameters over small domains and to use these parameters to appropriately define the subinstances.

The majority of approaches that are found in literature specify the drawing area of an instance based on leaders (or on callouts if the positions of the labels are not fixed). For example, consider 1-sided boundary labeling for point features and do -leaders such that the label positions are fixed, the labels touch the right of the boundary rectangle R , and the ports of the labels are also fixed on the midpoint of their left edges; see Fig. 7. Let A be the set of the induced reference points. We assume $|A| \geq |F|$, and that the cost function rates single labels. Similar to Benkert et al. [BHKNO9] we define the drawing area of a sub-instance of that problem setting by a simple polygon specified by two point features p_1, p_2 and two reference points a_1, a_2 ; see Fig. 7a. More precisely, the do -leader λ_1 connecting p_1 with a_1 bounds the area from above, and




Figure 7: One-sided instance of boundary labeling with d_o -leaders. (a) The instance I (blue drawing area) is specified by two point features p_1 and p_2 , and the two reference points a_1 and a_2 . (b) The leader λ with site p and reference point a decomposes the instance I into the two disjoint sub-instances I' (red drawing area) and I'' (orange drawing area).

the d_o -leader λ_2 connecting p_2 with a_2 bounds the area from below. Further, the area is bounded from the right by the bounding rectangle and from the left by the vertical line that goes through the rightmost site of λ_1 and λ_2 . We assume that λ_1 and λ_2 neither belong to the drawing area nor to the instance and observe that all point features of such an instance lie to the right of the sites of both leaders. Such an instance is uniquely defined by the choice of the point features and reference points of λ_1 and λ_2 . Thus, all such instances can be described by the tuples $(p_1, a_1, p_2, a_2) \in \mathcal{I} = F \times A \times F \times A$. Hence, \mathcal{I} has size $O(|F|^2|A|^2) = O(|A|^4)$.

Benkert et al. [BHKN09] prove that for any input instance there is a decomposition tree over \mathcal{I} as defined above. To see that, consider an optimal labeling \mathcal{L} of an instance I in \mathcal{I} ; see Fig. 7b. Let λ_1 and λ_2 be the leaders that bound I . We denote the point features and reference points of λ_1 and λ_2 by p_1, p_2 and a_1, a_2 , respectively. If I does not contain any point features, the instance I is a leaf in the decomposition tree. Further, we have $c_{\text{OPT}}(I) = 0$. Otherwise, if I contains some point features, let λ be the leader with leftmost site in \mathcal{L} . We denote its site by p . The leader splits the instance into the two disjoint sub-instances $I' = (p_1, a_1, p, a)$ and $I'' = (p, a, p_2, a_2)$. The instances I' and I'' are the children of I in the decomposition tree. Further, we obtain $c_{\text{OPT}}(I) = c_{\text{OPT}}(I') + c_{\text{OPT}}(I'') + c(\lambda)$, where $c(\lambda)$ denotes the costs for the leader λ . Hence, we can apply the dynamic programming approach as described above. The set \mathcal{I} contains $n = O(|A|^4)$ instances and each instance has $m = O(|A|)$ decompositions; recall that we assume $|A| \geq |F|$. Hence, we obtain $O(|A|^5)$ running time and $O(|A|^4)$ storage consumption. For p_o -leaders Benkert et al. [BHKN09] further show that an instance can be uniquely specified by two reference points improving the size of \mathcal{I} to $n = O(|A|^2)$, which yields $O(|A|^3)$ running time and $O(|A|^2)$ storage consumption. Similar dynamic programming approaches have been considered for other leader-types; see Tables 1 and 2. Further, while most approaches specify instances merely based on leaders, few approaches also use other separators [KNR*16, BCK*18].

Publications: [BKPS06a], [BKSW07], [BHKN09], [LWY09], [LN10], [LPT*11], [FHS*12], [BCK*18], [HPL14], [KLW14], [GHN15], [BCF*15], [KNR*16], [LNS16], [IFS16], [NNR17].

5.2.2. Weighted matching

Weighted matching has been applied successfully to instances of external labeling, in which the objective is the minimization of the total leader length, under few reasonable assumptions; see, e.g., [BKSW07, BKPS10, BKNS10]. The idea here is, given an instance of the external labeling problem, to construct an edge-weighted bipartite graph, such that a perfect matching of minimum cost for this graph translates back into a labeling of the original input instance, whose total leader length is minimum. For an introduction to matchings see, e.g., [AMO93, Chapter 12].

Characteristics

- Mostly applies to boundary labeling problems where the objective is to minimize the total leader length [BKSW07, BKPS10, BKNS10].
- A post-processing step is usually required to make the computed labelings crossing-free.
- The running time of the algorithm is usually dominated by the time needed to solve the matching problem.
- Requires uniform labels placed at fixed positions.

Sketch We sketch this technique for an instance $I = (D, F, M)$ of the 4-sided boundary labeling model, in which the feature set F consists of n point features, the leaders are of type- s , and both the positions of the labels within the labeling region A and the ports of the labels are fixed [BKSW07].

The main step in this technique is the definition of a weighted bipartite graph $G = (V_F \cup V_L, E, w)$, where $w: E \rightarrow \mathbb{R}^+$. The node-set of G is defined as follows. For each point feature p of F , graph G contains a node $v(p)$ in V_F , and for each label candidate ℓ , graph G contains a node $v(\ell)$ in V_L . The edge-set of G is then defined as the Cartesian product of V_F and V_L , that is, $E = V_F \times V_L$. In other words, graph G is complete bipartite. To complete the description of the definition of graph G , it remains to describe the weights of its edges. To this end, consider an edge e of G . By definition, edge e connects a node $v(p)$ in V_F to a node $v(\ell)$ in V_L . The weight $w(e)$ of edge e is defined as the length of the leader λ connecting point feature p with the port $\pi(\ell)$ of the label candidate ℓ . Equivalently, since leader λ is of type- s , the weight $w(e)$ of edge e is set to the Euclidean distance between point feature p and port $\pi(\ell)$.

A labeling \mathcal{L} for instance I that minimizes the total length of the leaders corresponds to finding a minimum-cost perfect matching for G , which exists since G is complete bipartite; also, the triangle inequality ensures that no two leaders of \mathcal{L} intersect, i.e., labeling \mathcal{L} is plane. In general, the problem of finding a minimum-cost matching of a weighted bipartite graph can be solved in $O(n^3)$ time using the Hungarian method; see, e.g., [Kuh55]. However, graph G has a special property that allows this problem to be solved more efficiently, namely, graph G is geometric (that is, drawn on the plane), and the weight of each edge corresponds to the Euclidean distance between its endpoints. For such graphs, the problem of finding a minimum-cost matching can be solved in time $O(n^{2+\delta})$, where δ can be chosen arbitrarily small [AES99].

The technique can also be applied to labels with sliding ports with a minor modification, namely, that the weight of each edge $e =$

$(v(p), v(\ell))$ of G must be the length of the *shortest* leader connecting point feature p with label ℓ . On the negative side, however, graph G is not geometric anymore, which implies that finding a minimum-cost matching for graph G now takes $O(n^3)$ time.

We conclude this section by noting that this technique has also been applied to boundary labeling with `opo`-leaders [BKSW07, BKPS10], and with a combination of `od`- and `pd`-leaders [BKNS10]. However, in these cases a post-processing step is required to eliminate potential crossings between leaders that may appear in the labeling \mathcal{L} computed by the minimum-cost matching for graph G . It is also worth noting that for the case of `opo`-leaders and labels with fixed ports, the minimum-cost bipartite matching problem can be solved more efficiently as observed in [BKSW07], since the bipartite graph G is geometric with respect to the Manhattan distance, which allows for a solution in $O(n^2 \log^3 n)$ time [Vai89].

Publications: [BKSW07], [LKY08], [BKPS10], [BKNS10], [BCF*15].

5.2.3. Scheduling

Boundary labeling has also been related to scheduling problems, in particular to *single-machine scheduling*. Here, a set of n jobs J_1, J_2, \dots, J_n is given, each of which is to be executed on a single machine. Each job J_i is associated with a processing time ρ_i and a *time window* (b_i, d_i) . If a job J_i is processed within its time window, then it incurs no penalty; otherwise, an *earliness* penalty or a *tardiness* penalty incurs, which is equal to the corresponding deviation. The objective is to determine a schedule, so that either the total earliness-tardiness penalty or the number of penalized jobs is minimized. For an introduction to scheduling see, e.g., [Bru04, Chapter 1].

Characteristics

- Mostly applies to instances of 1-sided of boundary labeling, in which the objective is to minimize the total leader length or to minimize the total number of leader bends [BKS08, GHN15].
- The technique has also been used to prove NP-hardness results [BKNS10, LPT*11, HPL14].
- The main idea is that the length of each label corresponds to the processing time of a corresponding job, while the objective to be optimized is expressed as a sum of the earliness and of the tardiness penalties of the jobs.

Sketch We sketch this technique on an instance $I = (D, F, M)$ of boundary labeling with type-`opo` leaders and sliding labels of non-uniform sizes that must be placed on one side on the boundary of the image region R , say the right side. For each $i = 1, 2, \dots, n$, point feature p_i of I is associated with a job J_i , whose processing time is h_i and whose due window is $(y_i - h_i, y_i + h_i)$, where y_i and h_i denote the y -coordinate of p_i and the height of the label of p_i , respectively.

Assume first that S^* is a schedule that minimizes the number of penalized jobs and consider the labeling \mathcal{L} obtained by setting the x -coordinate of the lower left corner of the label of the point feature p_i to the starting time of job J_i in S^* . Then, it is not difficult to see that \mathcal{L} has the minimum number of leader bends, because for each $i = 1, 2, \dots, n$ the leader connecting point feature p_i to its label is of type `o` (that is, without any bend) if and only if job J_i is completely processed in its time window in S^* . On the other

hand, if S^* minimizes the total earliness-tardiness penalties, then a symmetric argument implies that labeling \mathcal{L} is optimal in terms of the total length of the leaders. These simple linear-time reductions imply that finding 1-sided boundary labelings with sliding labels of non-uniform sizes that are either optimal in terms of the total number of leader bends or in terms of the total leader length can be found in $O(n^2)$ and $O(n \log n)$ time, respectively [BKS08, Kou96].

We finally note that the above reduction has also been used to prove NP-hardness mostly for non-uniform labels. The main idea here is that the order of the labels (and as a result the order of the jobs to be executed on the machine) may be variable when the leaders are not of type `opo`; see, e.g., [BKNS10, LPT*11, HPL14]. Hence, finding a schedule that minimizes the total earliness-tardiness penalty or the number of penalized jobs becomes NP-hard, which gives rise to corresponding NP-hardness reductions.

Publications: [BKS08], [BKNS10], [LPT*11], [HPL14], [GHN15].

5.2.4. Plane sweep

Plane sweep is a well-known technique with several applications, mainly in computational geometry. This technique has been proven to be useful also in boundary labeling since several instances of this problem have purely geometric flavour (in particular, for uniform labels with position and fixed ports); see S5.7 in Table 2. We stress, however, that the technique mostly applies when the goal is to find a crossing-free solution efficiently (that is, without taking into account any optimization criterion). A notable exception is the $O(n \log n)$ -time algorithm by Benkert et al. [BHKNO9] to compute 1-sided length-minimal boundary labelings with `po`-leaders. For an introduction to plane-sweep approaches see, e.g., [dBCvKO08].

Characteristics

- Mostly applies to variants of boundary labeling with uniform labels, fixed reference points and ports.
- Usually results in crossing-free boundary labelings.
- The complexity of the obtained algorithms is usually $O(n \log n)$, since a sorting of the input point features (e.g., from top to bottom or from left to right) is required.

Sketch We describe this technique on a simple variant of 1-sided boundary labeling with `s`-leaders, in which the external labels are of uniform size with fixed reference points and ports aligned on a single, say right, side of the image region R . The algorithm is incremental, as it processes the labels from bottom to top in order to compute a labeling in which no two leaders cross [BKSW07].

More precisely, let $\pi_1, \pi_2, \dots, \pi_n$ be the ports of the labels as they appear from bottom to top. Then, for $i = 1, 2, \dots, n$, the i -th port π_i is connected via an `s`-leader to the first unlabeled point feature that is hit by a ray r_i that emanates from π_i , initially points vertically downwards, and is rotated around π_i in clockwise order. Suppose that a crossing occurs between any two leaders. But then the rotating ray would have found the corresponding point features in the reverse order, which contradicts the fact that the two leaders cross. A straightforward implementation yields a time complexity of $O(n^2)$, which can be improved to $O(n \log n)$ using standard computational geometry techniques.

We conclude the description of this technique by mentioning that the algorithm presented above can be extended to compute corresponding 4-sided boundary labelings, under the same set of restrictions [BKSW07]. The idea is to partition the image region R into a particular number of convex non-overlapping sub-regions (using again appropriate rotating lines), such that the number of point features in the interior of each sub-region equals the number of labels on its boundary, which implies that the 1-sided algorithm can be applied to each of the computed sub-regions independently.

Publications: [BKSW07], [BHKNO9], [BKNS10], [BKNS10], [BKPS11], [Lin10], [NPS10], [LPT*11], [FHS*12], [KLW14].

5.2.5. Mathematical programming

Mathematical programming is a general global optimization approach. A mathematical program consists of a set of variables, a set of inequalities and an optimization function. The goal is to find an assignment of the variables such that all inequalities are satisfied and the value of the optimization function is (depending on the application) minimized or maximized. Often, these approaches make use of integer or binary variables. While this particularly provides the possibility of easily formulating combinatorial problems such as Problem 1, integer variables, generally speaking, turn the optimization problem into an NP-hard one. Still, powerful solvers for mathematical programs have been developed, which can be used to solve a broad variety of instances in reasonable time. For an introduction to mathematical programming see, e.g., [Van14].

Characteristics

- General technique that can be used to formulate many (NP-hard) optimization problems.
- Solvers are available, but their running times are hard to predict.
- Common technique to obtain optimal solutions for the purpose of evaluating optimality gaps of heuristics for the same problem.

Sketch We briefly sketch a basic variant of a mathematical program for Problem 1. More precisely, we present an integer linear programming formulation (ILP), i.e., a mathematical program that consists of integer variables, linear inequalities and a linear objective function. Alternatively, one may also use quadratic programming formulations. Instead of linear cost-functions they also allow the formulations of quadratic cost-functions. While this supports the formalization of more complex problems, this usually also comes with a significantly higher running time.

We assume that we are given a set F of features and for each feature $f \in F$ a set L_f of callout candidates; we denote the union of all those candidate sets by L . Further, let $c: L \rightarrow \mathbb{R}$ be a cost function that rates each candidate in L . We aim at a labeling $\mathcal{L} \subseteq L$ such that for each feature $f \in F$ exactly one callout $\gamma \in L_f$ is contained in \mathcal{L} and $\sum_{\gamma \in \mathcal{L}} c(\gamma)$ is minimized among all those labelings.

Based on the work by Barth et al. [BGNN18], we use the following ILP formulation to express this optimization problem. We note that this is a standard formulation to express assignment problems. For each candidate $\gamma \in L$ we introduce a binary variable $x_\gamma \in \{0, 1\}$. We interpret the variable x_γ such that $x_\gamma = 1$ if the callout is selected for the labeling and $x_\gamma = 0$ otherwise. We further introduce

the following linear constraints.

$$x_\gamma + x_{\gamma'} \leq 1 \text{ for each } \gamma, \gamma' \in L \text{ that exclude each other} \quad (1)$$

$$\sum_{\gamma \in L_f} x_\gamma = 1 \text{ for each feature } f \in F \quad (2)$$

Constraint (1) ensures that if two callouts exclude each other due to some given hard constraint (e.g., because they intersect), at most one of them is added to the labeling. Further, Constraint (2) enforces that for each feature exactly one candidate is selected. Subject to Constraints (1) and (2) we minimize the total costs $\sum_{\gamma \in L} c(\gamma) \cdot x_\gamma$. The resulting labeling is then defined as $\mathcal{L} = \{\gamma \in L \mid x_\gamma = 1\}$. We note that Barth et al. [BGNN18] use a similar formulation to minimize the total leader length in boundary labeling.

Publications: [BKPS11], [GHN15], [YDGM17], [BGNN18].

5.3. Complexity results

In contrast to internal labeling, where the majority of the problems turn out to be NP-hard, several external labeling problems admit polynomial-time solutions, as we have seen so far. In external labeling the problems become usually NP-hard, when the labels are non-uniform or their positions are not fixed. Next, we present some of these negative results grouped by the problem used to prove hardness. For an introduction to NP-hardness see, e.g., [GJ79].

The well-known, weakly NP-complete PARTITION problem [GJ79] has been used in several reductions in boundary labeling. Bekos et al. [BKSW07] observed that the task of assigning non-uniform labels to two opposite sides of the image region R , say to two vertical sides, is in one-to-one correspondence with the PARTITION problem, if the label heights sum up to twice the height of R . In the context of panorama labelings, Gemsa et al. [GHN15] employ a reduction from the PARTITION problem (that uses the same principle ideas as an earlier one of Garrido et al. [GIM*01]) to prove that determining whether a panorama labeling with vertical leaders exists is NP-hard. A conceptually similar reduction is given by Bekos et al. [BKS08] in the context of labeling collinear sites with non-uniform labels, but directly from the labeling problem by Garrido et al. [GIM*01]. In the context of boundary labeling with obstacles, Fink and Suri [FS16] suggest another reduction from the PARTITION problem to prove that determining whether a 1-sided labeling with po -leaders and non-uniform labels exists is NP-hard, even in the case of a single obstacle. Finally, a slightly more involved reduction from a variant of the PARTITION problem has been suggested by Bekos et al. [BKPS06a] to prove that the multi-stack boundary labeling problem with non-uniform labels is NP-hard.

Bekos et al. [BKNS10], Lin et al. [LPT*11] and Huang et al. [HPL14] proved that different variants of the boundary labeling problem, in which the labels are non-uniform and the objective is the minimization of the total leader length, are NP-hard by employing conceptually similar reductions from an NP-complete scheduling problem commonly known as TOTAL DISCREPANCY [GTW88]. According to this problem, the jobs have different processing times but they all share a common *midtime*, in which the first half of each job must be completed; the goal is to compute a schedule that minimizes the total earliness-tardiness penalty (see also Section 5.2.3).

In the many-to-one boundary labeling setting, Lin et al. [LKY08]

prove that the problem of minimizing the number of crossings between leaders can be expressed as a specific instance of the NP-hard TWO-LAYER CROSSING MINIMIZATION problem [EW94], according to which the goal is to compute a two-layer drawing of a given bipartite graph with minimum number of crossings. Bekos et al. [BCF*15] proposed a reduction from the FIXED-LINEAR CROSSING NUMBER problem [MNKF90] to prove that the problem of minimizing the number of crossings between leaders remains NP-hard, even in the case where several leaders to the same label can share a common (horizontal) segment. Recall that the fixed-linear crossing number problem asks for a linear embedding of a given graph, in which the order of the vertices is fixed along a line L and the edges must be drawn either completely above or completely below L such that the number of edge-crossings is minimum.

Publications: [BKPS06a], [BKSW07], [LKY08], [BKS08], [BKNS10], [LPT*11], [HPL14], [BCF*15], [FS16].

6. Discussion and conclusion

The previous two sections provided an in-depth discussion of the state of the art on external labeling algorithms. Our literature review identified that two mostly independent research communities have studied external labeling problems extensively: the visual computing community as well as the algorithms community. In visual computing the vast majority of research considers straight-line leaders and presents heuristics taking a large set of drawing criteria into account. In these practical works the concrete application and deployment of the approaches play a major role. In contrast, the algorithms community focuses on proving formal quality guarantees in labeling models limited to fewer drawing criteria. Algorithmic and geometric properties and results are primary, while abstraction makes the concrete application become secondary. In these algorithmic works the results are formal models and exact algorithms, mainly for poly-line leaders and boundary labeling. This raises the question what the two communities can learn from each other and how research efforts can be joined to create synergies. In the following we discuss this question in greater detail and state ten open challenges in external labeling. In these challenges we have not included immediate follow-up problems of individual papers (e.g., improve an algorithm's complexity or extend a method to a broader setting). Rather, these challenges emerge from the discussion that we have developed so far. They concern questions that are missing in the existing literature and, according to our expertise, need attention.

As evident from Section 4, a plethora of different labeling approaches have been developed and evaluated in the visual computing community. Still, direct comparisons between different results are missing from the literature, as opposed, e.g., to traditional map labeling, where such comparisons are available both in terms of algorithmic efficiency, but also in terms of the quality of the produced labelings [CMS95]. As a result, the approaches are mostly assessed by discussing labelings of selected example illustrations and 3D models. Hence, even if the approaches consider similar settings (possibly in different applications), one can hardly assess which of the approaches is better in the sense that it yields the more readable or more aesthetically pleasing labeling. Since the evaluations that have been conducted so far are mainly preliminary (see Section 4.4), we deem the systematic and quantitative evaluation of the existing

approaches to be one of the most important challenges. In order to obtain sustainable results, this particularly requires to make the results comparable. Especially practitioners will benefit from empirically founded readability results and guidelines when choosing a labeling style and an algorithm for their labeling tasks.

Challenge 1 *Systematically evaluate and compare different approaches for external labeling.*

This is where the theoretical works on boundary labeling have their particular strength. They use a common formal structure that is divided into two parts: the labeling model (including all parameters, hard and soft constraints) and the algorithms computing labelings within this particular model. Hence, the formalism admits simple comparisons of the models' features. Further, due to the clear distinction between model and algorithm, an evaluation can be split into two independent parts, i.e., the assessment of the labeling model and the assessment of the labeling algorithm. The unified taxonomy presented in Section 2 is the first step towards generalizing this structure from boundary labeling to more general external labeling models, enabling a better comparison of different approaches.

Multiple labeling models have been proposed in both communities. As observed in Sections 4.1 and 4.2, in the algorithms community these models are mostly specified by defining a particular leader type and a precise but limited objective function (e.g., the total leader length) that takes into account general drawing principles (e.g., proximity, planarity, and straightness/low detour [Pur97, WPCM02]). On the other hand, in the visual computing community user studies, experts interviews and the manual analysis of handmade drawings are common tools to extract and assess various additional drawing criteria, which are combined and balanced using multi-criteria optimization heuristics. As evident from Section 4.4, in both communities an extensive and systematic comparison of the labeling models is missing. Currently, criteria are often selected by common sense, experience, and following general cognitive and visualization principles and guidelines [Ric17]. While this is reasonable, we believe that establishing a set of quantitative, well-justified quality measures specific to external labeling would be of great help when comparing and choosing between different labeling models. It lends itself to include the already extracted drawing criteria (e.g., short, crossing-free leaders) as a basis for such measures, but extend this also towards more global measures such as balance, harmony, and interplay between image and callouts. Introducing new measures requires a systematic justification and ranking.

Challenge 2 *Establish and rank quantitative quality and aesthetics measures for evaluating and comparing external labelings.*

Assuming suitable and computable quality measures are established, we can compare different labeling models by computing and evaluating labelings for large sets of input instances. For example, one might aim at labelings that exclusively optimize a single quality measure such that two different labeling models can be compared with respect to this measure. We note that this procedure is independent from the actual labeling algorithm, but only takes the labeling model into account. However, this raises several problems. First of all, we need algorithms that compute mathematically optimal solutions for labeling problems. While for boundary labeling a large variety of such algorithms exist (as evident from Section 4.2, and

Tables 1 and 2), for contour labeling only the algorithm by Niedermann et al. [NNR17] is known. For the general external labeling problem there exist a large number of heuristics that perform well in specific application scenarios, but there is a lack of exact algorithms as evident from Section 4.1. On the other hand, exact algorithms based on established techniques such as dynamic programming or mathematical programming are expected to have high running times and thus limited practical applicability. Hence, the development of alternative algorithmic techniques that perform well in practice is needed. From a theoretical point of view it is also interesting to find lower bounds on the time complexity for both the underlying computational problems and the developed algorithms.

Challenge 3 *Develop efficient algorithms for optimally solving external labeling problems that are applicable in practice.*

We deem such algorithms to be helpful both for evaluating labeling models and heuristics, and for producing high-quality labelings used in professional books such as atlases of human anatomy, where somewhat longer computation times can be tolerated. On the other hand, for dynamic settings such as those that appear in the visualization of 3D models the computation of labelings in real-time is required. Further, complex constraints that improve temporal coherence must be satisfied. For such applications fast heuristics seem to be the first choice. We stress here that the vast majority of the existing heuristics (see Section 4.1) depend highly on the application domain and as a result they are tailored to the specifications of the problem, which they are designed to solve (e.g., [FLH*06, MTSP12, BČM14]). Therefore, it is of importance to develop more generic heuristics of broader scope that can be deployed in several applications.

Challenge 4 *Develop more generic, less problem-specific heuristics and algorithms for external labeling.*

Dynamic and interactive settings have been extensively considered by the visual computing community, while the algorithms community has mainly considered static labelings. A promising research direction is to build up on that knowledge to integrate temporal coherence and user interaction into formal models and to develop algorithms with formal quality guarantees for dynamic settings. The absence of such algorithms is greatly evident from Tables 1 and 2; in Table 1 observe that the intersection of S4 (temporal coherence) and S7.2 (algorithm type: exact) is empty, while in Table 2 there is only one entry (i.e., [NPS10]) in row P9.1 (algorithm type: exact) for dynamic external labelings, which focuses on a very restricted case as discussed in Section 4.2.

Challenge 5 *Develop efficient exact algorithms for computing dynamic, temporally coherent, and interactive external labelings.*

To obtain comparable and reproducible experimental results in labeling research, it is necessary to establish benchmark sets (which are common in several other fields, e.g., traditional map labeling [WW97], graph clustering [BMS*14], route planning [BS18] and graph drawing [DD13]) and make them available for experiments. Depending on the problem setting, these benchmark sets may consist of geographic data (as, e.g., those in [WW97]), illustrations, maps, or 3D models each annotated with feature sets of varying type, density, and complexity. The need of these benchmarks is even more evident from the fact that several papers in the literature

contain sample labelings obtained by implementations of different approaches (refer to S6.1 in Table 1 and to P8.2 in Table 1) but no systematic/quantitative evaluations.

Challenge 6 *Establish general and representative benchmark sets for external labeling.*

Establishing benchmarks will further facilitate the development and the systematic evaluation of (possibly much faster) heuristics, as their performance will be easy to be evaluated by simply comparing the optimal labelings and the results of these heuristics with respect to different quality measures. To simplify the process of such evaluations we suggest to develop a common external labeling framework and interface that allows to easily plug in different models and algorithms. This particularly may motivate researchers to make their source code available to the scientific community (e.g., only two of the 41 implementations reported in Tables 1 and 2 are publicly available in source code [FP99, BG05] and just few others provide JavaScript implementations or legacy Java applets). We note that the open source library *OpenLL* is currently being developed and provides an API specification for label rendering in 2D and 3D graphics [LGB*18]. As the library has its focus on rendering, only simple placement algorithms are provided so far. We deem this library to be one possible starting point for tackling Challenge 7.

Challenge 7 *Develop a generic framework for models and algorithms on external labeling.*

We pointed out that there is a division between strong practical work using straight-line leaders and a significant amount of theoretical work on polyline leaders. This can be partially explained by the fact that straight-line leaders are the shortest and simplest visual association between sites and labels, while polyline leaders possess complaisant geometric properties that can be exploited for exact algorithms. Nevertheless, it is far from clear which leader type is to be preferred. In a first formal user study on leader readability, Barth et al. [BGNN18] showed that *po*-leaders outperform *s*-leaders and that *opo*-leaders lag far behind. However, those experiments were conducted for boundary labeling using blank background images. Hence, in other applications, such as text annotations, *opo*-leaders seem to be a better choice because they can run between the text lines, while *po*- or *s*-leaders pass through the text. Furthermore, *s*-, *po*-, *do*- and *pd*-leaders, as well as curved leaders are all frequently found in professional information graphics. They arguably differ in their aesthetic appeal, which is not only influenced by their geometry, but also by the type and style of the background image. The few existing guidelines as described in Section 3 are valuable resources when assessing different options of external labelings, but they do not cover all types of labelings and are primarily based on practical experience [Woo00, Ric17]. There is a lack of formal empirical studies and readability results to confirm or adapt these guidelines.

Challenge 8 *Establish evidence-based guidelines and rules to determine which leader type suits which application and task best.*

Finally, we consider the combination of internal and external labels to be a fruitful research field, as comparatively little work has been done on this topic. From the algorithmic perspective only few special cases of mixed labeling have been considered so far that mainly combine very simple internal and external labeling models [BKPS11, LN10, LNS16]. From a practical, visual perspective the

question that arises is when to use or not to use external rather than internal labels. In the same direction, it is also worth studying more sophisticated mixed labeling models, which allow combinations of internal and external labels, but also support *internal* labels connected to their features by short leaders; see, e.g., [BFH01, LSC08].

Challenge 9 Develop rules to address the dichotomy between *internal* and *external* labels, i.e., when to use which type of labels.

Challenge 10 Broaden the research on mixed labeling by combining more sophisticated labeling models for *internal* and *external* labels.

In summary, both in the visual computing community and in the algorithms community a multitude of specializations and variants of external labeling have been considered from two different perspectives. This survey gave a comprehensive overview over both fields discussing the research and algorithmic techniques that have been presented over the last two decades. At the same time it shall pave the way to join both research directions creating new synergies and pushing forward the progress on external labeling.

Acknowledgments We thank the anonymous reviewers for their helpful comments and suggestions to improve this survey.

References

- [AES99] AGARWAL P. K., EFRAT A., SHARIR M.: Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM Journal on Computing* 29 (1999), 912–953. 21
- [AF03] AZUMA R., FURMANSKI C.: Evaluating label placement for augmented reality view management. In *Mixed and Augmented Reality (ISMAR'03)* (2003), IEEE, pp. 66–75. 7
- [AHS05] ALI K., HARTMANN K., STROTHOTTE T.: Label layout for interactive 3D illustrations. *Journal of the WSCG* 13, 1 (2005), 1–8. 1, 6, 7, 8, 9, 11, 12, 13, 14, 15, 16, 18, 19
- [AMO93] AHUJA R. K., MAGNANTI T. L., ORLIN J. B.: *Network flows - theory, algorithms and applications*. Prentice Hall, 1993. 21
- [BCF*15] BEKOS M., CORNELSEN S., FINK M., HONG S.-H., KAUFMANN M., NÖLLENBURG M., RUTTER I., SYMVONIS A.: Many-to-one boundary labeling with backbones. *Journal of Graph Algorithms and Applications* 19, 3 (2015), 779–816. 12, 13, 15, 21, 22, 24
- [BCK*18] BOSE P., CARMÍ P., KEIL J. M., MEHRABI S., MONDAL D.: Boundary labeling for rectangular diagrams. In *Algorithm Theory (SWAT'18)* (2018), vol. 101 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 12:1–12:14. 12, 14, 16, 17, 21
- [BČM14] BALATA J., ČMOLÍK L., MIKOVEC Z.: On the selection of 2d objects using external labeling. In *Human Factors in Computing Systems (CHI'14)* (2014), ACM, pp. 2255–2258. 7, 8, 11, 18, 19, 25
- [BDY06] BEEN K., DAICHES E., YAP C.: Dynamic map labeling. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (Sep. 2006), 773–780. 6
- [BFH01] BELL B., FEINER S., HÖLLERER T.: View management for virtual and augmented reality. In *User Interface Software and Technology (UIST'01)* (2001), ACM, pp. 101–110. 7, 26
- [BG05] BRUCKNER S., GRÖLLER E.: Volumeshop: An interactive system for direct volume illustration. In *Visualization (Vis'2005)* (2005), IEEE, pp. 671–678. 1, 8, 9, 10, 19, 25
- [BGNN18] BARTH L., GEMSA A., NIEDERMANN B., NÖLLENBURG M.: On the readability of leaders in boundary labeling. *Information Visualization* (2018). 7, 8, 9, 10, 11, 12, 14, 16, 17, 18, 23, 25
- [BHKN09] BENKERT M., HAVERKORT H. J., KRÖLL M., NÖLLENBURG M.: Algorithms for multi-criteria boundary labeling. *Journal of Graph Algorithms and Applications* 13, 3 (2009), 289–317. 12, 13, 14, 15, 17, 20, 21, 22, 23
- [BKNS10] BEKOS M. A., KAUFMANN M., NÖLLENBURG M., SYMVONIS A.: Boundary labeling with octilinear leaders. *Algorithmica* 57, 3 (2010), 436–461. 12, 14, 18, 21, 22, 23, 24
- [BKPS06a] BEKOS M. A., KAUFMANN M., POTIKA K., SYMVONIS A.: Multi-stack boundary labeling problems. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)* (2006), vol. 4337 of *LNCS*, Springer, pp. 81–92. 12, 13, 15, 16, 21, 23, 24
- [BKPS06b] BEKOS M. A., KAUFMANN M., POTIKA K., SYMVONIS A.: Polygon labelling of minimum leader length. In *Asia-Pacific Symposium on Information Visualisation (APVis'06)* (2006), Misue K., Sugiyama K., Tanaka J., (Eds.), vol. 60 of *CRPIT*, Australian Computer Society, pp. 15–21. 2
- [BKPS10] BEKOS M. A., KAUFMANN M., POTIKA K., SYMVONIS A.: Area-feature boundary labeling. *Computer Journal* 53, 6 (2010), 827–841. 12, 15, 17, 21, 22
- [BKPS11] BEKOS M. A., KAUFMANN M., PAPADOPOULOS D., SYMVONIS A.: Combining traditional map labeling with boundary labeling. In *Current Trends in Theory and Practice of Computer Science (SOFSSEM'11)* (2011), vol. 6543 of *LNCS*, Springer, pp. 111–122. 12, 13, 14, 16, 23, 25
- [BKS08] BEKOS M. A., KAUFMANN M., SYMVONIS A.: Efficient labeling of collinear sites. *Journal of Graph Algorithms and Applications* 12, 3 (2008), 357–380. 12, 14, 16, 17, 22, 23, 24
- [BKSW04] BEKOS M. A., KAUFMANN M., SYMVONIS A., WOLFF A.: Boundary labeling: Models and efficient algorithms for rectangular maps. In *Graph Drawing (GD'04)* (2004), vol. 3383 of *LNCS*, Springer, pp. 49–59. 8, 9
- [BKSW07] BEKOS M. A., KAUFMANN M., SYMVONIS A., WOLFF A.: Boundary labeling: Models and efficient algorithms for rectangular maps. *Computational Geometry: Theory and Applications* 36, 3 (2007), 215–236. 1, 6, 7, 9, 10, 12, 14, 15, 16, 17, 18, 21, 22, 23, 24
- [BMS*14] BADER D. A., MEYERHENKE H., SANDERS P., SCHULZ C., KAPPES A., WAGNER D.: Benchmarking for graph clustering and partitioning. In *Encyclopedia of Social Network Analysis and Mining*. Springer, 2014, pp. 73–82. 25
- [BRL09] BERTINI E., RIGAMONTI M., LALANNE D.: Extended eccentric labeling. *Computer Graphics Forum* 28, 3 (2009), 927–934. 12, 13, 14, 18
- [Bru04] BRUCKER P.: *Scheduling algorithms (4. ed.)*. Springer, 2004. 22
- [BS18] BLUM J., STORANDT S.: Scalability of route planning techniques. In *Automated Planning and Scheduling (ICAPS'18)*. (2018), de Weerd M., Koenig S., Röger G., Spaan M. T. J., (Eds.), AAAI Press, pp. 20–28. 25
- [BSF*11] BATTERSBY S. E., STEWART J. E., FEDE A. L.-D., REMINGTON K. C., MAYFIELD-SMITH K.: Ring maps for spatial visualization of multivariate epidemiological data. *Journal of Maps* 7, 1 (2011), 564–572. 7, 8, 9
- [ČB10] ČMOLÍK L., BITTNER J.: Layout-aware optimization for interactive labeling of 3D models. *Computers & Graphics* 34, 4 (2010), 378–387. 7, 8, 9, 19
- [ČB18] ČMOLÍK L., BITTNER J.: Real-time external labeling of ghosted views. *IEEE Transactions on Visualization and Computer Graphics* (2018). 7, 8, 9, 10, 18, 19
- [CLRS09] CORMEN T. H., LEISERSON C. E., RIVEST R. L., STEIN C.: *Introduction to Algorithms (3. Edition)*. MIT Press, 2009. 19, 20
- [CMS95] CHRISTENSEN J., MARKS J., SHIEBER S. M.: An empirical study of algorithms for point-feature label placement. *ACM Trans. Graph.* 14, 3 (1995), 203–232. 24
- [dBCvKO08] DE BERG M., CHEONG O., VAN KREVELD M., OVERMARS M. H.: *Computational geometry: algorithms and applications*. Springer, 2008. 22
- [DD13] DI BATTISTA G., DIDIMO W.: Gdtoolkit. In *Handbook on Graph Drawing and Visualization*, Tamassia R., (Ed.). Chapman and Hall/CRC, 2013, pp. 571–597. 25

- [Deg92] DEGANI A.: *On the Typography of Flight-Deck Documentation*. Tech. Rep. 177605, Ames Research Center, 1992. 5
- [Ead84] EADES P.: A heuristic for graph drawing. *Congressus numerantium 42* (1984), 146–160. 19
- [EMJ*11] ELMQVIST N., MOERE A. V., JETTER H.-C., CERNEA D., REITERER H., JANKUN-KELLY T.: Fluid interaction for information visualization. *Information Visualization 10*, 4 (2011), 327–340. 6
- [EW94] EADES P., WORMALD N. C.: Edge crossings in drawings of bipartite graphs. *Algorithmica 11*, 4 (1994), 379–403. 16, 24
- [FHS*12] FINK M., HAUNERT J.-H., SCHULZ A., SPOERHASE J., WOLFF A.: Algorithms for labeling focus regions. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2583–2592. 7, 8, 9, 21, 23
- [FLH*06] FUCHS G., LUBOSCHIK M., HARTMANN K., ALI K., SCHUMANN H., STROTHOTTE T.: Adaptive labeling for interactive mobile information systems. In *Information Visualization (InfoVis'06)* (2006), IEEE, pp. 453–459. 7, 8, 10, 19, 25
- [FP99] FEKETE J.-D., PLAISANT C.: Excentric labeling: Dynamic neighborhood labeling for data visualization. In *Human Factors in Computing Systems (CHI'99)* (1999), ACM Press, pp. 512–519. 1, 7, 8, 9, 10, 12, 13, 14, 17, 18, 25
- [FR91] FRUCHTERMAN T. M. J., REINGOLD E. M.: Graph drawing by force-directed placement. *Software: Practice and experience 21*, 11 (1991), 1129–1164. 19
- [FS16] FINK M., SURI S.: Boundary labeling with obstacles. In *Canadian Conference on Computational Geometry (CCCG'16)* (2016), pp. 86–92. 7, 8, 9, 12, 14, 15, 16, 17, 21, 23, 24
- [FW91] FORMANN M., WAGNER F.: A packing problem with applications to lettering of maps. In *Computational Geometry (SoCG'91)* (1991), ACM Press, pp. 281–288. 6
- [GAHS05] GÖTZELMANN T., ALI K., HARTMANN K., STROTHOTTE T.: Form follows function: Aesthetic interactive labels. In *Computational Aesthetics in Graphics, Visualization and Imaging (CAe'05)* (2005), Eurographics Association, pp. 193–200. 9
- [GHN15] GEMSA A., HAUNERT J.-H., NÖLLENBURG M.: Multirow boundary-labeling algorithms for panorama images. *ACM Transactions on Spatial Algorithms and Systems 1*, 1 (2015), 1–30. 7, 8, 9, 10, 19, 21, 22, 23
- [GHS06a] GÖTZELMANN T., HARTMANN K., STROTHOTTE T.: Agent-based annotation of interactive 3D visualizations. In *Smart Graphics (SG'06)* (2006), vol. 4073 of LNCS, Springer, pp. 24–35. 7, 8, 9, 10, 19
- [GHS06b] GÖTZELMANN T., HARTMANN K., STROTHOTTE T.: Contextual grouping of labels. In *SimVis* (2006), pp. 245–258. 7, 8, 9, 10
- [GIM*01] GARRIDO M. A., ITURRIAGA C., MÁRQUEZ A., PORTILLO J. R., REYES P., WOLFF A.: Labeling subway lines. In *Algorithms and Computation (ISAAC'01)* (2001), vol. 2223 of LNCS, Springer, pp. 649–659. 23
- [GJ79] GAREY M. R., JOHNSON D. S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. 23
- [GTW88] GAREY M. R., TARJAN R. E., WILFONG G. T.: One-processor scheduling with symmetric earliness and tardiness penalties. *Mathematics of Operations Research 13*, 2 (1988), 330–348. 23
- [HAS04] HARTMANN K., ALI K., STROTHOTTE T.: Floating labels: Applying dynamic potential fields for label layout. In *Smart Graphics (SG'04)* (2004), vol. 3031 of LNCS, Springer, pp. 101–113. 7, 8, 10, 18, 19
- [HGAS05] HARTMANN K., GÖTZELMANN T., ALI K., STROTHOTTE T.: Metrics for functional and aesthetic label layouts. In *Smart Graphics (SG'05)* (2005), vol. 3638 of LNCS, Springer, pp. 115–126. 7, 8, 9, 10, 18, 19
- [HGK14] HEINSOHN N., GERASCH A., KAUFMANN M.: Boundary labeling methods for dynamic focus regions. In *Pacific Visualization Symposium (PacificVis'14)* (2014), IEEE, pp. 243–247. 7, 8, 9, 10, 19
- [HH14] HAUNERT J.-H., HERMES T.: Labeling circular focus regions based on a tractable case of maximum weight independent set of rectangles. In *ACM SIGSPATIAL Workshop on MapInteraction* (2014). 8, 9
- [HPL14] HUANG Z., POON S., LIN C.: Boundary labeling with flexible label positions. In *WALCOM: Algorithms and Computation (WALCOM'14)* (2014), vol. 8344 of LNCS, Springer, pp. 44–55. 12, 13, 14, 15, 16, 17, 21, 22, 23, 24
- [Imh75] IMHOF E.: Positioning names on maps. *The American Cartographer 2*, 2 (1975), 128–144. 6
- [Kau09] KAUFMANN M.: On map labeling with leaders. In *Efficient Algorithms*, Albers S., Alt H., Näher S., (Eds.), vol. 5760 of LNCS. Springer, 2009, pp. 290–304. 1
- [KLW14] KINDERMANN P., LIPP F., WOLFF A.: Luatodonotes: Boundary labeling for annotations in texts. In *Graph Drawing (GD'14)* (2014), vol. 8871 of LNCS, Springer, pp. 76–88. 5, 7, 8, 9, 12, 13, 15, 16, 18, 21, 23
- [KMPV17] KEIL M., MITCHELL J., PRADHAN D., VATSHELLE M.: An algorithm for the maximum weight independent set problem on outer-string graphs. *Computational Geometry: Theory and Applications 60* (2017), 19–25. 16
- [KNR*16] KINDERMANN P., NIEDERMANN B., RUTTER I., SCHAEFER M., SCHULZ A., WOLFF A.: Multi-sided boundary labeling. *Algorithmica 76*, 1 (2016), 225–258. 12, 13, 15, 16, 21
- [Kou96] KOULAMAS C.: Single-machine scheduling with time windows and earliness/tardiness penalties. *European Journal of Operational Research 91*, 1 (1996), 190–202. 22
- [KT18] KAKOULIS K. G., TOLLIS I. G.: Automatic placement of labels in maps and drawings. In *Handbook of Approximation Algorithms and Metaheuristics*, Gonzalez T. F., (Ed.), 2nd ed. Taylor & Francis, 2018, ch. 7. 1
- [Kuh55] KUHN H. W.: The Hungarian method for the assignment problem. *Naval Research Logistic Quarterly 2* (1955), 83–97. 17, 21
- [LGB*18] LIMBERGER D., GROPLER A., BUSCHMANN S., WASTY B., DÖLLNER J.: Openll: An api for dynamic 2d and 3d labeling. In *Information Visualisation (IV'18)* (July 2018), pp. 175–181. 25
- [Lin10] LIN C.: Crossing-free many-to-one boundary labeling with hyperleaders. In *Pacific Visualization Symposium (PacificVis'10)* (2010), IEEE, pp. 185–192. 12, 13, 14, 16, 17, 23
- [LKY08] LIN C., KAO H., YEN H.: Many-to-one boundary labeling. *Journal of Graph Algorithms and Applications 12*, 3 (2008), 319–356. 12, 13, 14, 15, 16, 17, 19, 22, 23, 24
- [LN10] LÖFFLER M., NÖLLENBURG M.: Shooting bricks with orthogonal laser beams: A first step towards internal/external map labeling. In *Canadian Conference on Computational Geometry (CCCG'10)* (2010), pp. 203–206. 12, 13, 14, 21, 25
- [LNS16] LÖFFLER M., NÖLLENBURG M., STAALS F.: Mixed map labeling. *J. Spatial Information Science 13* (2016), 3–32. 12, 13, 16, 18, 21, 25
- [LPT*11] LIN C., POON S., TAKAHASHI S., WU H., YEN H.: One-and-a-half-side boundary labeling. In *Combinatorial Optimization and Applications (COCA'11)* (2011), vol. 6831 of LNCS, Springer, pp. 387–398. 12, 13, 14, 17, 21, 22, 23, 24
- [LSC08] LUBOSCHIK M., SCHUMANN H., CORDS H.: Particle-based labeling: Fast point-feature labeling without obscuring other visual features. *IEEE Trans. Vis. Comput. Graph. 14*, 6 (2008), 1237–1244. 26
- [LWY09] LIN C., WU H., YEN H.: Boundary labeling in text annotation. In *Information Visualization (InfoVis'09)* (2009), IEEE, pp. 110–115. 12, 13, 14, 16, 17, 21
- [MD06] MAASS S., DÖLLNER J.: Efficient view management for dynamic annotation placement in virtual landscapes. In *Smart Graphics (SG'06)* (2006), vol. 4073 of LNCS, Springer, pp. 1–12. 7

- [MNKF90] MASUDA S., NAKAJIMA K., KASHIWABARA T., FUJISAWA T.: Crossing minimization in linear embeddings of graphs. *IEEE Transactions on Computers* 39, 1 (1990), 124–127. [24](#)
- [MP09] MÜHLER K., PREIM B.: Automatic textual annotation for surgical planning. In *Vision, Modeling, and Visualization (VMV'09)* (2009), Eurographics Association, pp. 277–284. [7](#), [8](#), [9](#), [11](#), [18](#)
- [MTM*16] MADSEN J. B., TATZQERN M., MADSEN C. B., SCHMALSTIEG D., KALKOFEN D.: Temporal coherence strategies for augmented reality labeling. *IEEE Transactions on Visualization and Computer Graphics* 22, 4 (2016), 1415–1423. [6](#), [18](#)
- [MTSP12] MOGALLE K., TIETJEN C., SOZA G., PREIM B.: Constrained labeling of 2d slice data for reading images in radiology. In *Visual Computing for Biomedicine (VCBM'12)* (2012), Eurographics Association, pp. 131–138. [6](#), [7](#), [8](#), [9](#), [11](#), [18](#), [19](#), [25](#)
- [NNR17] NIEDERMANN B., NÖLLENBURG M., RUTTER I.: Radial contour labeling with straight leaders. In *Pacific Visualization Symposium (PacificVis'17)* (2017), IEEE. [1](#), [6](#), [7](#), [8](#), [9](#), [10](#), [18](#), [21](#), [25](#)
- [NPS10] NÖLLENBURG M., POLISHCHUK V., SYSIKASKI M.: Dynamic one-sided boundary labeling. In *Advances in Geographic Information Systems (ACM-GIS'10)* (2010), ACM Press, pp. 310–319. [12](#), [13](#), [14](#), [23](#), [25](#)
- [OJP14] OELTZE-JAFRA S., PREIM B.: Survey of labeling techniques in medical visualizations. In *Visual Computing for Biomedicine (VCBM'14)* (2014), Eurographics Association, pp. 199–208. [1](#)
- [PAE09] PETERSON S. D., AXHOLT M., ELLIS S. R.: Objective and subjective assessment of stereoscopically separated labels in augmented reality. *Computers & Graphics* 33, 1 (2009), 23–33. [6](#)
- [PHTP*10] PICK S., HENTSCHEL B., TEDJO-PALCZYNSKI I., WOLTER M., KUHLER T.: Automated Positioning of Annotations in Immersive Virtual Environments. In *Joint Virtual Reality Conference of EGVE (JVRC'09)* (2010), Kuhlen T., Coquillart S., Interrante V., (Eds.), The Eurographics Association. [6](#), [8](#), [11](#), [18](#), [19](#)
- [Pur97] PURCHASE H.: Which aesthetic has the greatest effect on human understanding? In *Graph Drawing (GD')* (1997), vol. 1353 of LNCS, Springer, pp. 248–261. [6](#), [24](#)
- [Ric17] RICHARDS C.: Technical and scientific illustration. In *Information Design: Research and Practice*, Black A., Luna P., Lund O., Walker S., (Eds.). Taylor & Francis, 2017, ch. 5, pp. 85–106. [5](#), [24](#), [25](#)
- [SC17] SHIMABUKURO M., COLLINS C.: Abbreviating text labels on demand. In *Information Visualization (INFOVIS'17)* (2017), IEEE. Poster abstract. [5](#)
- [SD08] STEIN T., DÉCORET X.: Dynamic label placement for improved interactive exploration. In *Non-Photorealistic Animation and Rendering (NPAR'08)* (2008), ACM Press, pp. 15–21. [7](#), [8](#), [11](#), [19](#)
- [SMHW15] SCHWARTGES N., MORGAN B., HAUNERT J.-H., WOLFF A.: Labeling streets along a route in interactive 3D maps using billboards. In *AGILE 2015* (2015), LNCG, Springer, pp. 269–287. [7](#)
- [TKGS14] TATZGERN M., KALKOFEN D., GRASSET R., SCHMALSTIEG D.: Hedgehog labeling: View management techniques for external labels in 3D space. In *Virtual Reality (VR'14)* (2014), IEEE, pp. 27–32. [6](#), [7](#), [8](#), [11](#), [19](#)
- [TKS13] TATZGERN M., KALKOFEN D., SCHMALSTIEG D.: Dynamic compact visualizations for augmented reality. In *Virtual Reality (VR'13)* (2013), IEEE, pp. 3–6. [7](#), [8](#), [10](#), [19](#)
- [Tuf01] TUFTE E. R.: *The Visual Display of Quantitative Information*. Graphics Press, 2001. [6](#)
- [Vai89] VAIDYA P. M.: Geometry helps in matching. *SIAM Journal on Computing* 18 (1989), 1201–1225. [17](#), [22](#)
- [Van14] VANDERBEI R. J.: *Linear Programming: Foundations and Extensions*. Springer US, 2014. [23](#)
- [Voo13] VOORHEES A. M.: A general theory of traffic movement. *Transportation* 40, 6 (2013), 1105–1116. [18](#)
- [VVAH07] VOLLIK I., VOGEL D., AGRAWALA M., HERTZMANN A.: Specifying label layout style by example. In *User Interface Software and Technology (UIST'07)* (2007), ACM Press, pp. 221–230. [7](#), [8](#), [9](#), [10](#), [18](#)
- [WD05] WU W., DALAL E. N.: Perception-based line quality measurement. In *Image Quality and System Performance II* (2005), vol. 5668, SPIE. [5](#)
- [WD08] WOOD J., DYKES J.: Spatially ordered treemaps. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1348–1355. [18](#)
- [Wol13] WOLFF A.: Graph drawing and cartography. In *Handbook of Graph Drawing and Visualization*, Tamassia R., (Ed.). CRC Press, 2013, ch. 23, pp. 697–736. [1](#)
- [Woo00] WOOD C. H.: A descriptive and illustrated guide for type placement on small scale maps. *Cartographic J.* 37, 1 (2000), 5–18. [25](#)
- [WPCM02] WARE C., PURCHASE H., COLPOYS L., MCGILL M.: Cognitive measurements of graph aesthetics. *Information Visualization* 1 (2002), 103–110. [6](#), [24](#)
- [WPT*15] WU H.-Y., POON S.-H., TAKAHASHI S., ARIKAWA M., LIN C.-C., YEN H.-C.: Designing and annotating metro maps with loop lines. In *Information Visualization* (2015), IEEE, pp. 9–14. [18](#)
- [WTLY11] WU H.-Y., TAKAHASHI S., LIN C.-C., YEN H.-C.: A zone-based approach for placing annotation labels on metro maps. In *Smart Graphics (SG'11)* (2011), vol. 6815 of LNCS, Springer, pp. 91–102. [7](#), [8](#), [10](#), [19](#)
- [WTLY12] WU H.-Y., TAKAHASHI S., LIN C.-C., YEN H.-C.: Travel-route-centered metro map layout and annotation. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 925–934. [5](#), [18](#)
- [WW97] WAGNER F., WOLFF A.: A practical map labeling algorithm. *Comput. Geom.* 7 (1997), 387–404. [25](#)
- [YDGM17] YANG Y., DWYER T., GOODWIN S., MARRIOTT K.: Many-to-many geographically-embedded flow visualisation: An evaluation. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 411–420. [12](#), [13](#), [14](#), [18](#), [23](#)