







Quantitative Comparison of Time-Dependent Treemaps

E. Vernier^{2†}, M. Sondag^{1†}, J. Comba², B. Speckmann¹, A. Telea³, and K. Verbeek¹

†The first two authors contributed equally to this paper.

¹TU Eindhoven, Netherlands

²UFRGS, Brazil

³Utrecht University, Netherlands

Abstract

Rectangular treemaps are often the method of choice to visualize large hierarchical datasets. Nowadays such datasets are available over time, hence there is a need for (a) treemaps that can handle time-dependent data, and (b) corresponding quality criteria that cover both a treemap's visual quality and its stability over time. In recent years a wide variety of (stable) treemapping algorithms has been proposed, with various advantages and limitations. We aim to provide insights to researchers and practitioners to allow them to make an informed choice when selecting a treemapping algorithm for specific applications and data. To this end, we perform an extensive quantitative evaluation of rectangular treemaps for time-dependent data. As part of this evaluation we propose a novel classification scheme for time-dependent datasets. Specifically, we observe that the performance of treemapping algorithms depends on the characteristics of the datasets used. We identify four potential representative features that characterize time-dependent hierarchical datasets and classify all datasets used in our experiments accordingly. We experimentally test the validity of this classification on more than 2000 datasets, and analyze the relative performance of 14 state-of-the-art rectangular treemapping algorithms across varying features. Finally, we visually summarize our results with respect to both visual quality and stability to aid users in making an informed choice among treemapping algorithms. All datasets, metrics, and algorithms are openly available to facilitate reuse and further comparative studies.

CCS Concepts

• Human-centered computing → Treemaps; • Information systems → Temporal data;

1. Introduction

Treemaps are one of the best-known methods for visualizing large hierarchical datasets. Given an input tree whose leaves have several attributes, treemaps recursively partition a 2D spatial region into cells whose visual attributes (area, color, shading, or annotation) encode the tree's data attributes. Compared to other methods such as node-link techniques, treemaps effectively use all available screen pixels to show data, and thus can display trees of tens of thousands of nodes on a single screen. Most treemaps use rectangles, although there are alternative models such as Voronoi treemaps [BDL05], orthoconvex and L-shaped treemaps [dBSvdW14], and Jigsaw treemaps [Wat05]. In this paper, we focus exclusively on rectangular treemaps.

The input for a rectangular treemap is a rectangle R and a set of non-negative values a_1, \dots, a_n together with a hierarchy on these values (represented by a tree). The output is a treemap T , which is a recursive partition of R into a set $\mathcal{R} = \{R_1, \dots, R_n\}$ of interior-disjoint rectangles, where (a) each rectangle R_i has area a_i , and (b) the regions of the children of an interior node of the hierarchy form a rectangle (associated with their parent). Such a partition of a rectangle into a set of disjoint rectangles is also called a *rectangular*

layout, or *layout* for short. Typically the input values are *normalized*, that is, the sum $A = \sum_i a_i$ corresponds to the area of R .

Nowadays, large hierarchical datasets are also available over time. Hence, there is a need for *time-dependent* treemaps which display changing trees and data values. Ideally, such time-dependent treemaps enable the user to easily follow structural changes in the tree and in the data. In a time-dependent setting, the input values become functions $a_i: [0, X] \rightarrow \mathbb{R}_{\geq 0}$ for each i , where the discrete domain $[0, X]$ represents the different time steps in the data. We assume that the hierarchy on the values and R are not time-dependent, and that the values a_i are properly normalized for each time step separately. We use the special value $a_i(t) = 0$ to represent that data element i is not present at time t ; and we speak of insertions or deletions if $a_i(t)$ starts or stops to be nonzero, respectively.

The *visual quality* of rectangular treemaps is usually measured via the aspect ratio of its rectangles. This indicator can become arbitrarily bad: Consider a treemap that consists of only two rectangles. If the area of one of these rectangles tends towards zero, then its aspect ratio tends towards infinity. Nagamochi and Abe [NA07] describe an algorithm (APP) which computes, for a given set of values and a hierarchy, a treemap which provably approximates the optimal

aspect ratio. De Berg *et al.* [dBSvdW14] prove that minimizing the aspect ratio for rectangular treemaps is strongly NP-complete. Kong *et al.* [KHA10] propose perceptual guidelines to improve treemap design and Zhou *et al.* [ZCYT17] perform user studies to test the effectiveness of different rectangular treemapping algorithms. Recently Lu *et al.* [LFH*17] argue that the optimal aspect ratio for treemaps should, in fact, be the golden ratio. In Section 2, we describe the state-of-the-art of rectangular treemaps in detail along with the various characteristics of rectangular treemaps.

For time-dependent treemaps, a second quality criterion is *stability*. Ideally, small changes in the data should result only in small changes in the treemap. Such stable behavior ensures that the only changes the user sees are due to the data, and not due to the decisions the algorithm makes. In recent years a few non-rectangular treemaps were specifically developed for time-dependent data. Hahn *et al.* [HTMD14] and Hees and Hage [vHH17] describe stable versions of Voronoi treemaps. Chen *et al.* [CDY17] propose a small-multiple metaphor to visualize time-dependent hierarchies. Their algorithm computes a global layout for all time steps simultaneously, but does not handle insertions or deletions. Scheibel *et al.* [SWD18] give an algorithm that maps changes in the data onto an initial layout. However, “treemaps” of subsequent time steps are not proper rectangular layouts as white space is introduced when resolving overlaps between rectangles.

A different approach to visualizing time-varying hierarchical data is taken by Lukaczyk *et al.* [LWM*17], Köpp and Weinkauff [KW19], and Li *et al.* [LZD*19], who show how to compute static overviews of the entire evolution of the tree. Alternatively, Guerra-Gómez *et al.* [GGPPS13] and Card *et al.* [CSP*06] use interactivity to explore time-varying data. For a broader perspective on tree visualizations, Graham and Kennedy [GK10] present a survey of visualizations that compare multiple trees, while Schulz *et al.* [SHS11] and Scheibel *et al.* [STL*20] present, respectively, a survey and a taxonomy for the visualization of a single tree.

Contribution. Despite their enduring popularity, a comprehensive evaluation of treemaps is currently lacking, even more so for the time-dependent case. Individual papers tend to report on only a few algorithms and evaluate only a few datasets, often without a principled discussion of quality metrics. To provide insights to both researchers and practitioners and to allow them to make an informed choice when selecting a treemap for their specific application and data, we perform an extensive quantitative evaluation of rectangular treemaps for time-dependent data. Our three main contributions are:

(1) We introduce a new method to measure the stability of time-dependent treemaps which explicitly considers the input data (Section 3.2). An algorithm is stable if small changes in the input data result in small changes in the layout, that is, data change and layout change correlate positively. Previously proposed stability metrics measure only the layout change and conclude that small layout changes are a sign of a stable algorithm. However, to properly measure stability, we also need to capture the data change and then correlate data and layout change. Here, we have to overcome the difficulty that the data and the layout space are a priori incomparable. We solve this problem by introducing the concept of a *baseline treemap* T^* which represents the minimum amount of change that

any time-dependent treemap must incur (given the input data) when moving from treemap T to the next treemap T' .

(2) We propose a novel classification scheme for time-dependent datasets. Specifically, based on our discussion of the state-of-the-art of treemaps in Section 2, we observe that the performance of treemaps depends on the characteristics of the datasets used. We identify four potential representative features that characterize time-dependent hierarchical datasets and classify all datasets used in our experiments accordingly. We experimentally test the validity of this classification on 2405 datasets, and analyze the relative performance of 14 state-of-the-art rectangular treemapping algorithms across varying features. Generally we conclude that our proposed features do indeed have predictive value, both with respect to visual quality and stability. We also observe that algorithms that are designed to be stable tend to in fact be more stable across features.

(3) We perform a quantitative evaluation of 14 rectangular treemapping algorithms on more than 2000 datasets. We visually summarize our results with respect to both visual quality and stability to aid users in making an informed choice among treemaps. All datasets, metrics, and algorithms are openly available [URLb]. Section 5 reports on our experimental results, we conclude in Section 6.

2. Rectangular Treemaps

We next discuss the most well-known rectangular treemapping algorithms. For a fair comparison during our experiments, we require that treemap rectangles have exactly the correct areas and partition the input rectangle. Algorithms that do not satisfy these requirements are not included in our evaluation. Recall that the input for a rectangular treemap is a rectangle R and a set of non-negative values a_1, \dots, a_n together with a hierarchy on these values (represented by a tree). The children of a node in this hierarchy are given in a particular order in the input. We distinguish two classes of treemaps, which either do or do not use this order. For time-dependent data we also distinguish between state-aware and stateless treemaps. Contrary to stateless treemaps, state-aware treemaps do not compute the treemap separately at a time step, but (can) use the layout of the previous time step to compute a new layout. Most treemaps are stateless; we discuss the state-aware algorithms separately.

Unordered treemaps do not (need to) adhere to the input nodes’ order when computing the layout. Typically, input weights are sorted to help the algorithm achieve good visual quality. Unordered treemaps in our evaluation include Squarified treemaps (SQR) [BHvW00] and Approximation treemaps (APP) [NA07]. APP comes with a guaranteed upper bound on the worst-case aspect ratio, while SQR often achieves near-optimal aspect ratios in practice. The visual quality of unordered treemaps is relatively unaffected by high *weight variance*, as reordering weights allows the layout to group similar-size rectangles in the treemap, typically leading to better aspect ratios. Yet, the sorted order of the weights may change rapidly over time, especially if the *weights change* much over time or if the *weight variance* is low. This can negatively affect the stability of the treemaps.

Ordered treemaps are required to adhere to the order of nodes as given in the input, which roughly ensures that rectangles close to each other in the input are close to each other in the resulting treemap.

This typically improves the stability of treemaps, but may worsen visual quality. We include nine ordered treemaps in our evaluation. The first ordered treemaps [SW01] include the Pivot-by-Middle (PBM), Pivot-by-Size (PBZ), and Pivot-By-Split (PBS) algorithms. Similar algorithms are the Strip algorithm (STR) [BSW02] and the Split algorithm (SPL) [Eng05]. Other algorithms, like the Spiral algorithm (SPI) [TS07], and the Hilbert (HIL) and Moore (MOO) algorithms [TC13], lay out rectangles following a space-filling curve. Finally, the very first treemap algorithm (Slice-and-Dice (SND)) by Shneiderman [Shn92] can also be considered an ordered treemap. While not ordered by design, the resulting (combinatorial) layout depends only on the hierarchy and not on weights. In fact, SND uses the depth in the hierarchy to compute the layout (slicing vertically on even depth and horizontally on odd depth), rather than simply applying the same algorithm recursively. Hence, SND's visual quality strongly depends on the *number of levels* in the input hierarchy. Typically, laying out large rectangles near small rectangles leads to poor aspect ratios. Hence, the visual quality of ordered treemaps is negatively affected by high *weight variance*. However, ordered treemaps are relatively stable over time compared to unordered treemaps, as order is maintained. Finally, *insertions and deletions* may affect the visual quality and stability of ordered treemaps to varying degrees, depending on how they are handled exactly.

State-aware treemaps can use the layout of the previous time step to compute a new layout and so can largely control their stability. The treemap for the first time step is typically an existing unordered treemap. The first state-aware treemap was introduced by Sondag *et al.* [SSV18]. Their Local Moves algorithm (LM) is initialized with APP, and allows only a small number of local modifications to the (combinatorial) layout between time steps. They also show how to update areas between time steps without significantly changing the layout (layouts remain “order-equivalent”). In our evaluation we include the Local Moves algorithm with 4 local moves between time steps (LM4), and without local moves (LM0). A similar algorithm is the Git algorithm (GIT) by Vernier *et al.* [VCT18b], which is initialized with SQR, and does not allow any changes to the (combinatorial) layout between time steps. Both state-aware treemaps also support insertions and deletions, updating the layouts locally where necessary (for insertions, the position in the layout can be chosen to maximize visual quality). By design, the stability of state-aware treemaps is relatively unaffected by frequent *weight changes* over time. Also, the visual quality of the initial treemaps should be relatively high. However, since the layouts cannot change much over time, the visual quality of state-aware treemaps will decrease over time if weights change significantly. Frequent *insertions and deletions* may also cause treemaps with poor visual quality, as treemaps are not recomputed as a whole. However, many insertions can help to correct rectangles with bad aspect ratio caused by weight changes over time. This is especially helpful for state-aware algorithms that do not allow any changes to the layout, like LM0 and GIT. Note that SND has a fixed layout if the input hierarchy does not change and is hence very stable, but it does not explicitly use the previous state.

Finally, note that the *number of levels* in the input hierarchy can have a strong effect on all classes of algorithms. In general, more levels imply less freedom in the layout strategy. As a result, unordered treemaps become more similar to ordered treemaps. Overall, the visual quality tends to decrease and the stability tends to increase.

3. Metrics

Wattenberg [Wat05] identifies several desirable properties of treemaps: (1) nicely shaped regions (visual quality), (2) stability with regard to changing leaf values, (3) stability with regard to changing tree structure, and (4) preservation of order information. Regarding Property (3), the tree structure can change in various ways; for example, nodes can merge or split, nodes can change parents, or there are general insertions and deletions. In our experiments we do not make any assumptions on the type of changes to the tree structure, and hence treat them as general insertions and deletions. Furthermore, we do not assume that the order of the values in the data is meaningful in general. Thus, we consider the following two important criteria to evaluate treemaps: *visual quality* and *stability*. We discuss well-established metrics for both below. We also introduce a new method to measure the stability of time-dependent treemaps which captures inherent data changes. We compute metrics for each leaf rectangle separately and then aggregate these values for each algorithm and dataset (see Section 5 and [URLb] for details). Note that we do not compute metrics for non-leaf nodes.

3.1. Visual quality

The weight information in a treemap is conveyed by the areas of its rectangles. Since areas of rectangles closer to squares are visually easier to estimate than areas of elongated rectangles, visual treemap quality is commonly measured by the aspect ratio of its rectangles. Although it has been proposed that the ratio should be close to the golden ratio [LFH*17] instead of the minimum aspect ratio of 1, it is commonly accepted that strongly elongated rectangles hinder readability of treemaps. We thus aim for the overall goal of making rectangles as square as possible, or similarly, minimizing the number of elongated rectangles. For a rectangle R_i of width $w(R_i)$ and height $h(R_i)$, we define the aspect ratio $\rho(R_i)$ as

$$\rho(R_i) = \min(w(R_i), h(R_i)) / \max(w(R_i), h(R_i)). \quad (1)$$

Observe that this definition is the inverse of the usual definition for aspect ratio. Its values range from 0 to 1, where values of ρ close to 0 are considered “bad” and values close to 1 are considered “good”. The bounded range allows for easy aggregation. Note that, compared to the usual definition of $1/\rho$, rectangles with larger aspect ratios have a smaller influence on the aggregated score.

3.2. Stability

Evaluating the stability of a treemap is more involved than evaluating visual quality. Consider treemaps at two consecutive time steps $T(t)$ and $T(t+1)$. Since stability does not *explicitly* depend on the value of t , we denote the former and the new treemap by T and T' respectively, to simplify notation. We also denote the rectangle areas in T and T' by $\{a_1, \dots, a_n\}$ and $\{a'_1, \dots, a'_n\}$, respectively. For a stable treemapping algorithm, the (visual) difference between T and T' should roughly correspond to the difference between $\{a_1, \dots, a_n\}$ and $\{a'_1, \dots, a'_n\}$. Note that the combination of large changes in data values and small changes in the layouts is unlikely since rectangle areas in treemaps must exactly match the data values. Hence, we actually want to measure *instability*, that is, large layout changes that are not caused by large data changes.

Most existing treemap stability metrics consider only the visual change in the treemap's layout $d(T, T')$, usually computed by evaluating the change $\delta(R_i, R'_i)$ for each rectangle separately and aggregating it over all rectangles. Shneiderman and Wattenberg [SW01] define δ as the Euclidean distance between the vectors $(x(R_i), y(R_i), w(R_i), h(R_i))$ and $(x(R'_i), y(R'_i), w(R'_i), h(R'_i))$, where x , y , w , and h are the coordinates of the top-left corner, width, and height of a rectangle, respectively. They then define d as the average over all rectangles. Hahn *et al.* [HTMD14, Hah15] simplify this metric by defining δ as the distance moved by the centroid of a rectangle, again defining d as the average. Tak and Cockburn [TC13] use the same δ as [SW01], but define d as the variance over all values computed by δ . They also propose a drift metric, which measures how much a rectangle moves away from its average position over a long period. Recently, Scheibel *et al.* [SWD18] introduced two new layout change metrics: The *average aspect ratio change* defines δ as the relative change between the aspect ratios of R_i and R'_i , and defines d as the average. The *relative parent change* defines δ as the relative change of the distance between the center of a rectangle and the center of its parent, again defining d as the average. Chen *et al.* [CDY17] propose a metric to quantify the ability of users to track time-dependent data in treemaps, which is closely related to the drift metric [TC13]. A different approach measures layout change using pairs of rectangles. Hahn *et al.* [HBD17] introduce the *relative direction change*, which, for every pair of rectangles R_i and R_j , measures how much the angle from the center of R_i to the center of R_j changes. Finally, Sondag *et al.* [SSV18] proposed the *relative position change*, which, for every rectangle pair (R_i, R_j) , measures how much the relative position of R_i with respect to R_j changes. The distance d is then defined as the average over all rectangle pairs.

Summarizing the above, we distinguish two types of layout change metrics: (1) *absolute* metrics measure how much individual rectangles move/change, and (2) *relative* metrics measure how much positions of pairs of rectangles change relative to each other. For our experiments, we use both an absolute and a relative metric. In particular, as an absolute metric, we use the *corner-travel distance*, which is a well-known metric used in computer vision to quantify change between two shapes using feature points [TM07, Sze10]. In the vision community, it was established already many years ago [S*94, Bie87] that corners are a perceptually useful feature to identify and track. Besides this perceptual validation, the corner-travel metric lies also within a small bounded factor of the original metric introduced by Shneiderman and Wattenberg [SW01]. Specifically, let $w(R)$ and $h(R)$ be the width and height of an input rectangle R , respectively. Let p_i , q_i , r_i , and s_i (p'_i , q'_i , r'_i , and s'_i) be the positions of the corners of a rectangle R_i (R'_i). We define the normalized corner-travel (CT) distance for a rectangle as

$$\delta_{CT}(R_i, R'_i) = \frac{\|p_i - p'_i\|_1 + \|q_i - q'_i\|_1 + \|r_i - r'_i\|_1 + \|s_i - s'_i\|_1}{4 \sqrt{w(R)^2 + h(R)^2}}. \quad (2)$$

where $\|x\|_1$ denotes the ℓ_1 norm. Simply put, δ_{CT} is the corner-to-corner correspondence distance between R_i and R'_i . Note that $0 \leq \delta_{CT}(R_i, R'_i) \leq 1$, since a rectangle corner can travel by at most the length of the diagonal of R .

As a relative metric, we use the relative position change [SSV18]. We established experimentally that the corner-travel and the relative position change metric correlate clearly on more than 2000 data

sets. Hence in Section 5 we report only on experiments using the corner-travel distance. All other data can be found here [URLb].

Data change. The stability metrics discussed above do not take data change into account. If data changes by a large amount, then the layouts should be allowed to change significantly without considering this to be instability. To add data change to a stability metric, one can consider the difference or ratio between the layout change and the data change [VCT18a, VCT18b]. However, there are two problems: (1) we need a way to measure data change, and (2) the metric spaces for data and layouts need to be comparable. For example, data change can be measured in terms of changes of rectangle *areas* (since these correspond to the data). However, layout changes such as the corner-travel distance measure *lengths*, not areas. Areas and lengths are not directly comparable, and thus their ratios or differences may not be meaningful. Although such metrics could be made comparable by suitable normalization, such adaptations are necessarily metric-specific and ultimately result in numbers whose meaning is not clear.

Baseline treemap. We overcome the above issues with a new method that captures data change *in the layout space*. For this, we define a *baseline* treemap T^* with respect to T and T' . The layout of T^* (that is, the combinatorial structure of the rectangular subdivision which constitutes T^*) is based on the layout of T . However, the areas of the rectangles in T^* are the areas $\{a'_1, \dots, a'_n\}$ of T' . The idea is that T^* aims to minimize the layout distance to T among all treemaps with the areas of T' . Put differently: T^* approximates the minimum amount of change that any time-dependent treemap must incur when moving from T and its associated area values $\{a_i\}$ to the next treemap T' and its area values $\{a'_i\}$. As a result, $d(T, T^*)$ is a good metric for data change in the layout space.

We construct T^* for each tested algorithm and each time step using a hill-climbing algorithm, which was proven to converge in [EMSV12]. For a rectangular layout (treemap) T , a *maximal segment* is a maximal contiguous horizontal or vertical line segment contained in the union of the borders of all rectangles in T (for example, the green segments in Figure 1). Put simply, a horizontal maximal segment (which is not part of the input rectangle R) always has endpoints on the interior of two vertical segments and vice versa. For two horizontal maximal segments s_1 and s_2 , we say that $s_1 < s_2$ if there is a rectangle in T whose bottom side coincides with s_1 and whose top side coincides with s_2 . This defines a partial order on horizontal maximal segments. We define a partial order on vertical maximal segments analogously (Figure 1). We say that T is *order-equivalent* to T^* if the corresponding partial orders on maximal segments are isomorphic. For every possible set of areas, there exists an order-equivalent treemap to T that correctly represents those areas [EMSV12]. In particular, we can initially define T^* as the treemap order-equivalent to T (computed with any of the tested algorithms) with the areas $\{a'_1, \dots, a'_n\}$ of T' .

If rectangles are inserted or deleted, the baseline treemap cannot be order-equivalent to T , so we handle insertions and deletions separately. Dealing with deletions is easy: we simply let the areas go to zero. For insertions, we must be more careful. Indeed, while we consider only rectangles present in both T and T' when measuring stability (R_i and R'_i in Equation 2), inserted rectangles can strongly

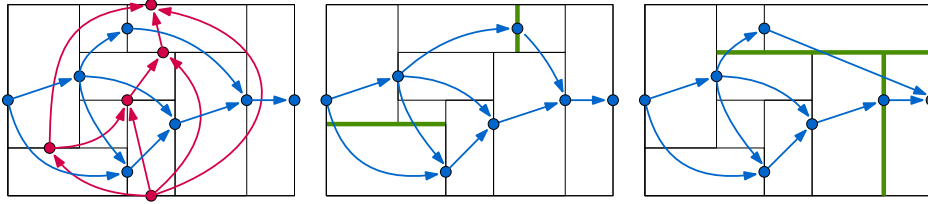


Figure 1: Left: Partial orders of the maximal segments. Middle: a layout order-equivalent to the left figure, changed maximal segments highlighted in green. Right: a layout not order equivalent to the other two figures. Red/blue arrows: relations between maximal segments.

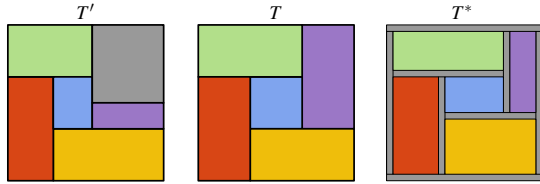


Figure 2: Treemaps T' (with gray rectangle inserted), T , and T^* (with gray area spread over maximal segments).

impact the positions of rectangles in T^* . We observe that the baseline treemap does not need to be a proper treemap: it only needs to capture how much rectangles must minimally move to update to the new data. To minimize the movement of the rectangles due to insertions (and hence be as stable as possible), we distribute the cumulative area of the inserted rectangles over the “walls” (borders) of treemap T as evenly as possible. To do so, we replace every maximal segment e in T by a rectangle, and assign each such rectangle a portion of the inserted area corresponding to the length of e (Figure 2). Hence all walls become equally thick and the original rectangles of T need to move as little as possible to yield T^* .

The baseline treemap T^* as proposed here is not a perfect baseline, as it does not always minimize the movement of every rectangle. However, the layout change between T and T^* is still a very good estimate for the minimum necessary layout change between T and T' , and thus a good measure for data change (see Figure 3: nearly all points lie on or below the diagonal).

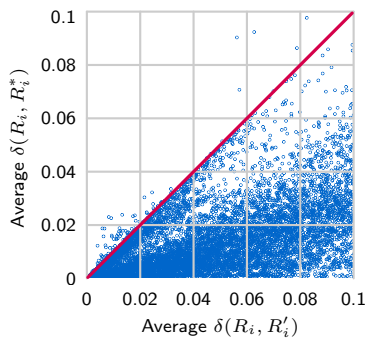


Figure 3: Scatter plot of the average layout change between T and T' or T^* for a random 25% sample of all algorithms and datasets.

Stability metric. We can now define a stability metric that takes data change into account. Consider a rectangle R_i and the corresponding rectangles R'_i and R_i^* in T' and T^* , respectively, and let δ be the layout change function for single rectangles. Two natural choices for spatial stability are the difference or ratio between $\delta(R_i, R'_i)$ and $\delta(R_i, R_i^*)$. Our experiments showed that the difference is typically more informative, that is, it exhibits clearer, more pronounced patterns, than the ratio. Hence, we define the stability of a single rectangle as

$$\sigma(R_i) = \max(0, \delta(R_i, R'_i) - \delta(R_i, R_i^*)) \quad (3)$$

Note that $\sigma(R_i) = 0$ if $\delta(R_i, R'_i) \leq \delta(R_i, R_i^*)$, which is possible. Indeed, a value of 0 for $\sigma(R_i)$ represents “very stable”, and R_i^* is considered to be (roughly) as stable as possible.

Limitations. The stability metrics we use focus only on consecutive time steps. The stability of time-varying treemaps could conceivably be influenced by effects that span multiple time steps, which our metrics do not capture directly. However, we believe that the most salient events influencing stability occur between consecutive time steps and hence we focus on this scenario.

4. Data

The visual quality and/or stability of treemaps clearly depends on the datasets used. Simply measuring the average performance over a (large) collection of datasets does not reveal such information. We aim to provide sufficient insight so that both practitioners and researchers can make informed choices about which algorithm to use for their data. For this, we study the performance of treemaps as a function of the characteristics of the input data. We classify the datasets into data classes along with explicit features and evaluate the metrics of different treemapping algorithms for each class.

4.1. Data features

Our methodology is inspired by the framework proposed by Smith-Miles *et al.* [SMBWL14] to objectively measure the performance of algorithms across datasets. For each dataset, we compute a number of features that (hopefully) capture the characteristics influencing the relative performance of treemapping algorithms. As a result, every dataset is represented by a point in a low-dimensional *feature space* \mathcal{F} . Similar feature-based approaches are also used to measure the relative performance of dimensionality-reduction methods [EMK*19] or in machine learning [Bis06]. Based on the discussion of treemapping algorithms in Section 2, we identify the

following four features: **1.** Levels of hierarchy, **2.** Variance of node weights, **3.** Weight change, and **4.** insertions and deletions.

Obviously, other features could be used to characterize (time-dependent) trees, such as the minimum, maximum, and average node degrees, the (im)balance of the tree structure [BO73, KY15], and the number of nodes. Two seemingly obvious candidates for features that we do not currently consider are the *number of nodes* and the *branching factor* (i.e., the average internal node degree). Arguably the number of levels in the hierarchy, the branching factor, and the number of nodes correlate to some degree. For example, if the hierarchy has only one level, then the branching factor and the number of leaves are the same. Hence, we should include at most two of these features in our analysis. Among these three features, the number of levels is with certainty a discriminating factor between algorithms, see our discussion in Section 2. Furthermore, all algorithms we consider, with the exception of SND, are recursive and treat each level independent from the preceding ones. Hence one can argue that the branching factor, which determines the number of nodes that have to be handled during a single step of this recursion, is a more relevant data feature than the total number of nodes. Nevertheless, we decided not to include the branching factor in our experiments, for the following two reasons. First of all, from the description of the algorithms, it seems that the branching factor is likely less relevant for their *relative* performance than the other four chosen features. That is, the descriptions of the algorithms do not give an indication that the branching factor is able to predict if an algorithm *A* will perform better than an algorithm *B* on a given dataset. Second, it is very difficult to define meaningful value-ranges for the branching factor and then to find datasets that cover these ranges in combination with all other data features. Given that the number of data classes and, correspondingly, the number of datasets needed for a meaningful evaluation, grows exponentially with the number of features chosen (see Section 4.2), we decided to restrict ourselves to four features. While we cannot exclude that the branching factor may influence relative performance, we do believe that the four features chosen have higher predictive value.

4.2. Data classes

Using the feature space \mathcal{F} , we partition all datasets into classes. For each feature we define a small number of subclasses based on only that feature. The data class of a dataset is then defined as the combination of the subclasses for each feature. We determined the value-ranges defining the subclasses by analyzing the distribution of feature values over our 2405 real-world tree datasets.

Levels of hierarchy (3 subclasses). We use three ranges for classification: 1 level (1L), 2 or 3 levels (2/3L), and more than 3 levels (4+L). Most hierarchical datasets we have analyzed have 2 or 3 levels. This number of levels is quite common for datasets that are visualized via treemaps, since they frequently concern geo-spatial subdivisions such as countries, continents, and their subregions, grouped by a classification scheme, such as the World Bank regional classification. Furthermore, visually understanding the node nesting in deeper treemaps becomes difficult [VvWvdL06, BHvW00]. A special case are datasets with only 1 level, that is, sets of weight values. Such datasets are also often visualized by treemaps, as these are more space-filling than alternatives such as bar charts [VvWvdL06].

These datasets are challenging for treemaps that implicitly use the depth of the hierarchy. Finally, we consider datasets with more than 3 levels, which correspond to deep hierarchies such as, for example, file systems or software architectures [HTMD14, HBD17, VCT18a].

Variance of node weights (2 subclasses). We distinguish between low variance (LWV) and high variance (HWV). To ensure that the total number of tree nodes does not strongly influence our classification, we use the coefficient of variation σ/μ to determine the subclass. The standard deviation σ and the mean μ are computed over all leaf weights over all time steps. We say that there is low variance if $\sigma/\mu \leq 1$ and high variance if $\sigma/\mu > 1$, respectively.

Weight change (3 subclasses). We distinguish between low weight change (LWC), regular weight change (RWC), and spiky weight change (SWC). The weight change of a single rectangle is measured by the absolute difference in the relative area (with respect to the input rectangle R) between time steps. The weight change of a treemap between two time steps is defined as the sum of weight changes of all rectangles. To determine the subclass of a dataset, we use the distribution of weight changes between time steps over all time steps in the dataset, specifically the mean μ and the standard deviation σ . Datasets with low weight change have $\mu < 5\%$ and $\sigma < 5\%$. Datasets with a larger mean ($5\% \leq \mu < 20\%$) and a relatively small coefficient of variation ($\sigma/\mu \leq 1$) are classified as having regular weight change. The weights of these datasets steadily change over time, without any extreme changes. Remaining datasets are classified as having spiky weight change. In those datasets weights change drastically ($\mu > 20\%$), or there is large variation ($\sigma/\mu > 1$) along with substantial changes ($\mu > 5\%$ or $\sigma > 5\%$).

Insertions and deletions (3 subclasses). We distinguish between low insertions and deletions (LID), regular insertions and deletions (RID), and spiky insertions and deletions (SID). We measure the impact of insertions and deletions between two time steps t and $t+1$ as the cardinality of the symmetric difference between the two sets of rectangles with non-zero weights at t and $t+1$, divided by the number of rectangles with non-zero weights at t . We again classify the datasets based on the distribution (μ and σ) of impact values over all time steps. Same as for the weight change, LID is defined by $\mu < 5\%$ and $\sigma < 5\%$, RID is defined by $\mu < 20\%$ and $\sigma/\mu \leq 1$, and the remaining datasets are in SID.

The full classification results in $3 \times 2 \times 3 \times 3 = 54$ data classes. In Section 5, we evaluate how the performance of treemapping algorithms depends on the classes, that is, if the classification is sensible.

4.3. Datasets

We collected a total of 2405 time-dependent hierarchical datasets from a variety of sources, detailed below. We found at least one dataset for 46 (out of 54) instances of our proposed data classes. See Figure 4 for the distribution of datasets over classes: clearly not all classes arise with equal frequency in our data sources.

World Bank [URL18c]: (2142 datasets) World development indicators such as agriculture, rural and urban development, education, trade and health. Hierarchy either according to the World Bank regional classification, grouping countries into subregions and continents, or no hierarchy present.

	LWV			HWV			
	LWC	RWC	SWC	LWC	RWC	SWC	
1L	SID	153	58	4	66	114	69
	RID	1	14	9	2	15	53
	LID	58	56	10	84	158	132
2/3L	SID	152	58	4	68	118	81
	RID	1	16	10	3	20	55
	LID	58	65	9	97	180	185
4+L	SID	1			4		
	RID		2				
	LID	1	6	3	56	19	76

Figure 4: Distribution of datasets over classes.

- GitHub [URL18e]:** (150 datasets) Hierarchies of folders, files, and classes, weighted by the number of code lines, extracted from all revisions of several GitHub repositories using Scitools [URLa].
- Movies:** (107 datasets) Movies from MovieLens [HK16] and TMDb [URL18d]. We constructed a time-dependent hierarchy using the group-rows-by-attribute-value partitioning method [Tel06, VvWvdL06]. The hierarchy groups movies based on their genres, actors, release date, and keywords. Each leaf is a movie, whose weight corresponds to its rating over a given period of time.
- Custom:** (6 datasets) Several individual datasets were added: *Dutch Names* [Mee16] contains the frequency of popular baby names in the Netherlands per year; *UN Comtrade Coffee* [URL17] contains the amount of coffee each country imported per year; *ATP* contains personal information, historical rankings, and match results from 1968 to 2018 for ATP tennis players [URL18a]; and *Earthquakes* contains the time, location, depth and intensity of seismic phenomena provided by the USGS Earthquake Hazards Program [URL18b].

Importantly, note that the above selection of dataset sources is *orthogonal* to the description of the feature space \mathcal{F} . The former covers the *origin* of data (which may cover application-specific aspects not captured by our feature space); the latter covers application-independent data aspects as captured by the data classes of \mathcal{F} . The distributions of data classes covered by our different data sources can be found in the supplementary material. The large and varied collection of World Bank datasets is able to cover all data classes with at most 3 levels of hierarchy (to which it is inherently limited). The GitHub and Movies datasets further cover a number of data classes with 4+ levels of hierarchy.

5. Experimental Results

We ran all 14 algorithms on all time steps of all 2405 datasets, generated the baselines for all these instances (Section 3), and recorded all layouts, that is, the positions of all rectangles $R_i(t)$ at all time steps t . Per dataset we aggregate our results for all metrics and algorithms by first by taking the mean over all rectangles in a single time step, and then by taking the mean again over all time steps. This is necessary since the number of rectangles may differ per time step.

We focus on two specific questions: We first explore the *validity* of our data classification (Section 5.1) and then we study the *performance* of all algorithms with respect to visual quality and stability across varying data features (Section 5.2). In the supplementary material we additionally compare the performance of all algorithms on each data *class* separately. We believe that the resulting visual summary will help researchers and practitioners choose a suitable treemapping algorithm for their data.

5.1. Data classification analysis

We evaluate if the relative performance of treemapping algorithms is more consistent within a data class than for an arbitrary collection of datasets. To perform this analysis we need to establish how we can capture the consistency of relative performance for a collection of datasets, and how we can compare this consistency between multiple collections. We restrict our analysis to data classes that contain at least 50 datasets, for otherwise the observed consistency is not sufficiently reliable. For each such data class, we randomly sample 50 datasets to use in this analysis. We also randomly sample 50 datasets among all 2405 datasets (all classes) as a baseline for comparison. Note that all collections must have the same number of datasets in the analysis to ensure that the comparisons are fair.

Now consider a single collection of datasets. To measure the consistency of relative performance among different datasets in this collection, we cannot directly use the computed metrics for visual quality and stability, as these values may differ greatly between datasets. Alternatively, we could rank the algorithms per dataset, but then algorithms with very similar performance may imply a greater variance in relative performance than is the case. Instead, we define the *relative performance* (separately for visual quality and stability) per dataset as follows. We compute both the best value (maximum for visual quality, minimum for stability) and the median value over all algorithms over this dataset. The *relative performance score* for each algorithm on this dataset is then computed by linearly

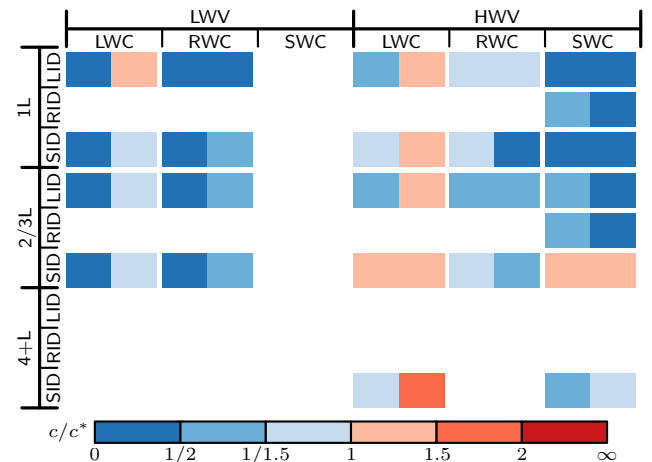


Figure 5: For each data class with at least 50 datasets, the ratio of the consistency score (visual quality on the left, stability on the right) between the data class and the baseline.

interpolating between these two values, where the best algorithm receives score 0, and the median algorithm receives score 0.5. The relative performance score is capped at 1, to avoid outliers. The resulting scores are comparable between different datasets.

We next analyze the consistency of relative performance within collections of datasets in two different ways. First, we use a quantitative approach: for each algorithm we compute the variance of the relative performance scores over all datasets in a collection, and we sum up these variances over all algorithms. This results in a

consistency score c for a collection of datasets. Figure 5 displays the consistency scores (for visual quality and stability) of all data classes (with at least 50 datasets) compared to the consistency scores c^* of the baseline collection (created by random sampling). A cell is colored blue (more consistent) if c is smaller than c^* ; a cell is colored red (less consistent) if c is larger than c^* .

Nearly all data classes for visual quality and most data classes for stability are more consistent than the baseline. This indicates that our features are splitting the datasets into valid data classes where the

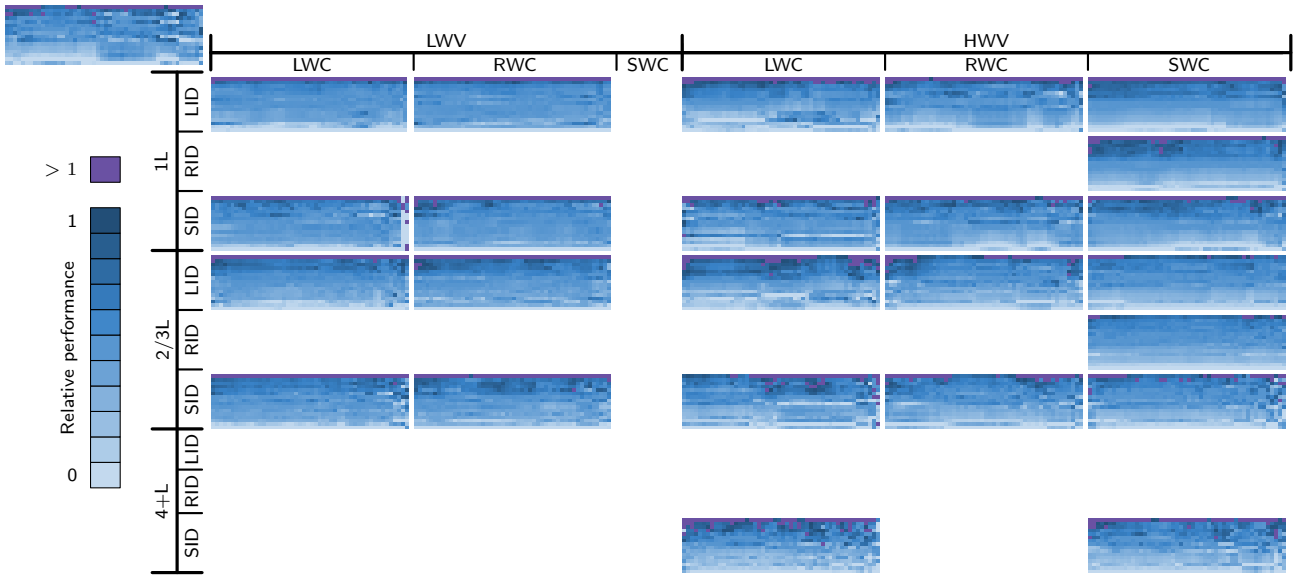


Figure 6: Visual quality: matrix plots for each data class with at least 50 datasets plus baseline (left top). In each matrix plot, rows correspond to algorithms, columns to datasets. The lighter the color, the better the relative performance, capped at 1 (purple).

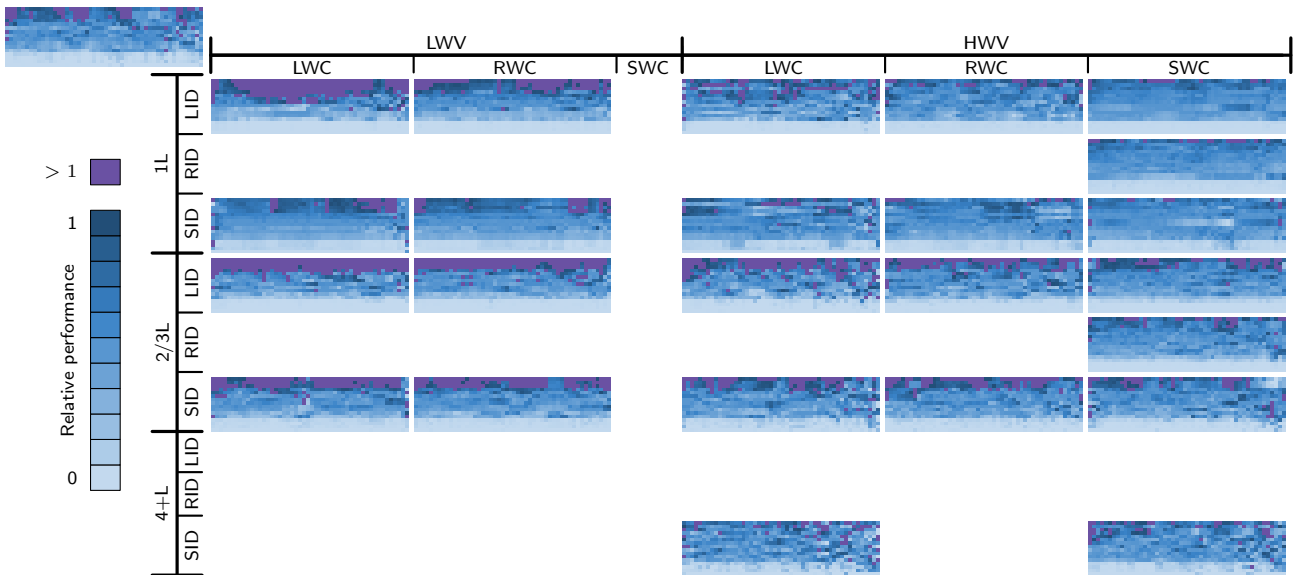


Figure 7: Stability: matrix plots each data class with at least 50 datasets plus baseline (left top). In each matrix plot, rows correspond to algorithms, columns to datasets. The lighter the color, the better the relative performance, capped at 1 (purple).

relative performance of an algorithm is easier to predict than in the baseline. However, the stability column for high weight variance and low weight change is less consistent than the baseline. As discussed in Section 2, the stability of unordered treemaps becomes worse compared to ordered treemaps when the weight variance is low or the weight change is high, due to reordering of the input weights. As a result, the difference with respect to stability between ordered and unordered treemaps is less pronounced for these data classes; the relative performance is hence influenced more by accidental details of individual datasets and less by structural differences between the algorithms. Additionally there are two data classes where the visual quality is less consistent than the baseline. It is not clear to us at this point what the cause of these inconsistencies is; one possibility are hidden correlations in the data classes.

Second, we use a more qualitative approach to assess the consistency of relative performance. For each data class we create a matrix plot, that shows the relative performance scores of all algorithms for all datasets in the collection (see Figures 6 and 7). Each column in the matrix plot represents a dataset, and each row represents an algorithm. The color of every “cell” in the matrix plot indicates the relative performance score of an algorithm on a dataset, where lighter colors indicate better (lower) relative performance scores. Relative performance scores that were capped at 1 are indicated with purple. To better enable the visual assessment of consistency among the different datasets in a collection, we order the datasets (columns) so that those with similar scores are next to each other as much as possible. Also, we order the algorithms (rows) so that the algorithms with better average score are lower in the matrix plot. In particular, the order of algorithms in the matrix plots for different data classes can be different. Figure 6 shows the matrix plots for visual quality, with the corresponding matrix plot for the baseline collection at the left-top. Figure 7 shows the matrix plots for stability.

First consider the matrix plot for visual quality (Figure 6). For the low weight variance subclass we indeed see that the matrix plots are much smoother than the baseline, which confirms the results in Figure 5. We also observe an increasing number of irregularities when going from 1 level treemaps to 2/3 levels or 4+ levels, since more levels impose more restrictions on the layout and hence all algorithms perform more similarly.

Consider now Figure 7. First of all, we notice that there is a set of four algorithms at the bottom of every matrix plot. These are the state-aware algorithms and SND. For nearly all datasets, regardless of the specific data class, these four algorithms are much more stable than any of the others. There is a large difference between the low weight variance and high weight variance subclasses. For low weight variance there is a set of algorithms that perform consistently much worse than the median (purple cells). These include the unordered treemaps which are particularly sensitive to changes in such data.

5.2. Performance analysis across features

The analysis in Section 5.1 shows that our data classification is valid. We now study how visual quality and stability depend on the *features* of the datasets. We aim to understand how sensitive a given algorithm is to variations in one or several features of its input data. For each data class we calculate the average visual quality and

stability. For each subclass of a feature we then take the average over all data classes that belong to it. This ensures that even though we have different numbers of datasets in each data class, they are all weighted equally. We show this data in Figures 8,9,10,11. Each point in each figure represents the score for one algorithm on one subclass of the feature, for example, low weight variance. We draw a polyline that connects the points of one algorithm and use glyphs to indicate the different subclasses. The different algorithms are indicated with categorical colors (see figure legends).

Recall that a low value for the stability metric indicates a *stable* algorithm and that the visual quality metric (aspect ratio) is bounded between 0 and 1. In particular, note that visual quality (ρ) of 0.5 for a single rectangle indicates a 2-by-1 rectangle. A ρ of 0.25 however is perceptually much worse than a ρ of 0.5 in terms of area perception as can be inferred from Kong *et al.* [KHA10], coming close to their “extreme aspect ratios” of 4.5.

Levels of hierarchy. Figure 8 considers the levels of hierarchy feature, which has three values: 1L, 2/3L, and 4+L. From Figure 8, we see that all algorithms, in particular the stateless ones, are more stable as the number of levels increase. In contrast to most other algorithms, the visual quality of state-aware algorithms (LM0, LM4, GIT) as well as SND increases with the number of levels. We also see that SQR and PBS have the longest polylines, that is, they are the most sensitive to the number of levels.

Variance of node weights. Figure 9 considers the weight variance feature, which has two values: LWV and HWV. Increasing the weight variance decreases the visual quality for all algorithms, except for APP (and SND). Additionally we see that the unordered treemaps are indeed more sensitive to this feature in terms of stability compared to the other algorithms. These algorithms reorder the data based on the weight to determine their layout, and if the weight are close to each other this happen more often.

Weight change. Figure 10 considers the weight change feature, which has 3 values: LWC, RWC, and SWC. The near-vertical polylines for the stateless algorithms show that visual quality seems to be largely unaffected by this feature. The stability however decreases quickly. Conversely, for the state-aware algorithms the polylines are mostly near-horizontal: the stability is largely unaffected, but the visual quality decreases. As the only state-aware algorithm that allows changes to the layout, LM4 makes an explicit tradeoff between stability and visual quality (see the slightly sloping line).

Insertions and deletions. Finally, Fig. 11 considers the insertions and deletions feature, which has three values: LID, RID, and SID. The plot shows a similar variation of visual quality and stability as seen for the weight change feature (Fig. 10). Yet, the polylines for the stateless algorithms now show a ‘kink’ at the midpoint (RID, regular insertions/deletions). Hence these algorithms are most unstable for regular insertions/deletions, and stabler for linear and spiky insertions/deletions. Interestingly, the state-aware methods (LM0, LM4, GIT) show a similar kink but oriented differently. These methods thus achieve poorest visual quality for regular insertions/deletions and highest quality on the other two values of this feature.

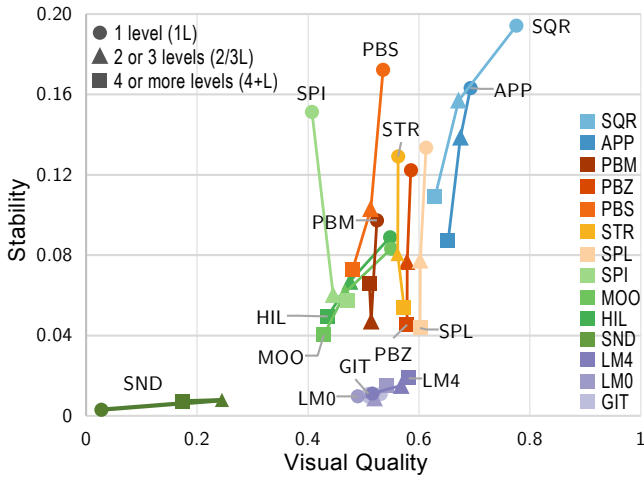


Figure 8: Visual quality vs stability as function of the levels of hierarchy feature.

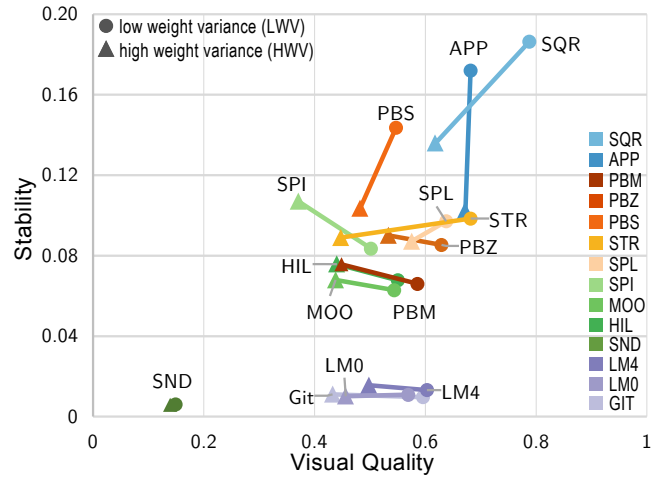


Figure 9: Visual quality vs stability as function of the variance of node weights feature.

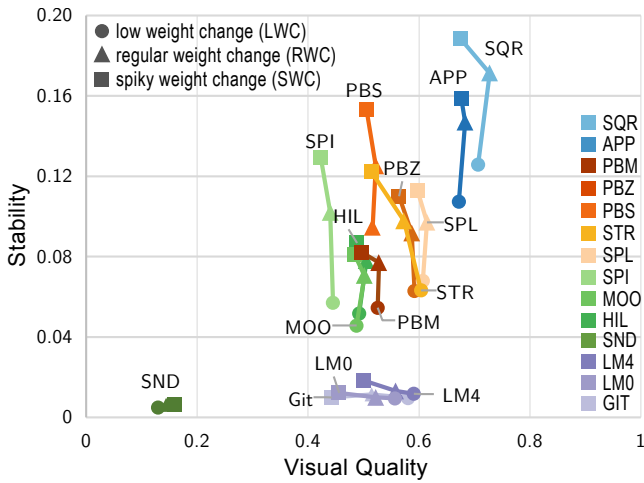


Figure 10: Visual quality vs stability as function of the weight change feature.

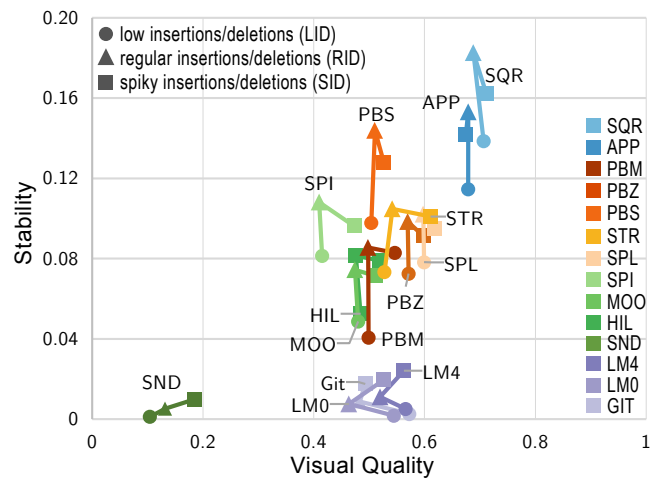


Figure 11: Visual quality vs stability as function of the insertions and deletions feature.

6. Discussion and Conclusion

We performed an extensive quantitative evaluation of rectangular treemaps for time-dependent data. To do so, we introduced a new methodology based on baseline treemaps to measure the stability of time-dependent treemaps. Baseline treemaps enable us to measure the change in the input data in a manner that is mathematically comparable to the measures for the layout change of the corresponding treemaps. Furthermore, we proposed a novel classification scheme for time-dependent data sets via a four-dimensional feature space (weight variance, weight change, tree depth, and the pattern of insertions and deletions). These four features naturally arose from a discussion on various types of state-of-the-art treemapping algorithms. Our experimental analysis shows that our proposed clas-

sification is valid in general and that most data classes are well suited to predict the performance of treemapping algorithms. For most data classes, our visual summary comparing all algorithms across all data classes and both metrics can hence serve as a reliable resource for researchers and practitioners. Last but not least, all datasets, metrics, and algorithms used in our evaluation are openly available [URLb].

Limitations and future work. Our experiments show that our identified features and the resulting feature space generally work well and result in a meaningful classification of datasets. However, there are whole sets of data classes for which we could not find sufficiently many (or even any) datasets. This is partially inherent in the classification and somewhat natural: data sets with low weight variance hardly ever exhibit spiky weight change behavior, so that particular

column in our table is essentially empty. But among the 18 classes of treemaps with 4 or more levels we found a significant number of datasets only for two classes, which both are essentially populated by datasets stemming from software repositories. The question remains if there are other significant types of time-dependent hierarchical datasets which have four or more levels and which escaped our searches. As it is, the results for these two particular classes are representative for only a restricted type of data.

Our classification works well for visual quality, with the exception of two cases (2/3 level, spiky insertions and deletions, high weight variance, and low or spiky weight change). We have a large number and variety of datasets at our disposal for these two classes, but nevertheless, it is unclear to us what causes these inconsistencies in the performance of the tested algorithms. There might be a hidden correlation in these datasets and one or more additional features might be needed to separate these classes further.

While we do have a significant number of datasets at our disposal and hence can validate our claims with some certainty, we still might be observing some bias in our collection. As stated above, essentially all datasets with 4 or more levels stem from software repositories. Furthermore, all World Bank datasets have at most 3 levels. It would be interesting to analyze if and how this bias in the data influences our results. To overcome possible data bias, we would also like to construct, and evaluate on, synthetic datasets. Doing so is not trivial; creating datasets that avoid sampling biases and are representative of real-world datasets (for a suitable definition of “real-world”) is a challenging (but important) question in its own right in information visualization in particular and in data science in general.

To complement our quantitative evaluation it would naturally be of interest to evaluate the performance of treemapping algorithms in various usage scenarios through user studies. The two metrics we use for visual quality and stability are both perceptually salient according to studies performed in previous work. However, a study that evaluates the combination of and the trade-offs between visual quality and stability could deliver important insights as to where on the Pareto-front an optimal treemapping algorithm should lie.

Finally, our evaluation currently does not measure the run-time and correspondingly the scalability of the algorithms used in our experiments. Our implementations are not (equally) optimized and hence a fair comparison is currently impossible. Scalability is clearly an important factor in online usage scenarios, and we hope to be able to complement our current set of implementations with optimized versions in the near future.

Acknowledgments. The Netherlands Organisation for Scientific Research (NWO) is supporting M. Sondag and B. Speckmann under project no. 639.023.208, and K. Verbeek under project no. 639.021.541. This study was also financed in part by CAPES (Finance Code 001) and CNPq (Process 308851/2015-3).

References

- [BDL05] BALZER M., DEUSSEN O., LEWERENTZ C.: Voronoi treemaps for the visualization of software metrics. In *Proc. ACM Symposium on Software Visualization* (2005), pp. 165–172. 1
- [BHvW00] BRULS M., HUIZING K., VAN WIJK J. J.: Squarified treemaps. In *Proc. Data Visualization* (2000), pp. 33–42. 2, 6
- [Bie87] BIEDERMAN I.: Recognition-by-components: A theory of human image understanding. *Psychological Review* 94, 2 (1987), 115–147. 4
- [Bis06] BISHOP C.: *Pattern Recognition and Machine Learning*. Springer, 2006. 5
- [BO73] BOORMAN S. A., OLIVIERA D. C.: Metrics on spaces of finite trees. In *Journal of Mathematical Psychology* (1973), vol. 10, pp. 26–59. 6
- [BSW02] BEDERSON B. B., SHNEIDERMAN B., WATTENBERG M.: Ordered and quantum treemaps: Making effective use of 2D space to display hierarchies. *ACM Transactions on Graphics* 21, 4 (2002), 833–854. 3
- [CDY17] CHEN Y., DU X., YUAN X.: Ordered small multiple treemaps for visualizing time-varying hierarchical pesticide residue data. *Visual Computer* 33, 6 (2017), 1073–1084. 2, 4
- [CSP*06] CARD S., SUH B., PENDLETON B. A., HEER B., BODNAR J. W.: Time tree: Exploring time changing hierarchies. In *Proc. Symposium on Visual Analytics Science and Technology* (2006), pp. 3–10. 2
- [dBSvdW14] DE BERG M., SPECKMANN B., VAN DER WEELE V.: Treemaps with bounded aspect ratio. *Computational Geometry* 47, 6 (2014), 683–693. 1, 2
- [EMK*19] ESPADOTO M., MARTINS R., KERREN A., HIRATA N., TELEA A.: Towards a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics* (2019). 5
- [EMSv12] EPPSTEIN D., MUMFORD E., SPECKMANN B., VERBEEK K.: Area-universal and constrained rectangular layouts. *SIAM Journal on Computing* 41, 3 (2012), 537–564. 4
- [Eng05] ENGDAHL B.: Ordered and unordered treemap algorithms and their applications on handheld devices. *Master’s degree project, Department of Numerical Analysis and Computer Science, Stockholm Royal Institute of Technology, SE-100 44* (2005). 3
- [GGPPS13] GUERRA-GÓMEZ J., PACK M., PLAISANT C., SHNEIDERMAN B.: Visualizing change over time using dynamic hierarchies: TreeVersity2 and the StemView. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2566–2575. 2
- [GK10] GRAHAM M., KENNEDY J.: A survey of multiple tree visualisation. *Information Visualization* 9, 4 (2010), 235–252. 2
- [Hah15] HAHN S.: Comparing the layout stability of treemap algorithms. *Proc. HPI research school on service-oriented systems engineering 95* (2015), 71–79. 4
- [HBD17] HAHN S., BETHGE J., DÖLLNER J.: Relative direction change – a topology-based metric for layout stability in treemaps. In *Proc. International Conference on Information Visualization Theory and Applications* (2017), pp. 88–95. 4, 6
- [HK16] HARPER F. M., KONSTAN J. A.: The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems* 5, 4 (2016), 19. 7
- [HTMD14] HAHN S., TRÜMPER J., MORITZ D., DÖLLNER J.: Visualization of varying hierarchies by stable layout of voronoi treemaps. In *Proc. International Conference on Information Visualization Theory and Applications* (2014), pp. 50–58. 2, 4, 6
- [KHA10] KONG N., HEER J., AGRAWALA M.: Perceptual guidelines for creating rectangular treemaps. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 990–998. 2, 9
- [KW19] KÖPP W., WEINKAUF T.: Temporal treemaps: Static visualization of evolving trees. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 534–543. 2
- [KY15] KUHNER M. K., YAMATO J.: Practical performance of tree comparison metrics. *Systematic Biology* 64, 2 (2015), 205–214. 6
- [LFH*17] LU L., FAN S., HUANG M., HUANG W., YANG R.: Golden rectangle treemap. *Journal of Physics: Conference Series* 787, 1 (2017). 2, 3
- [LWM*17] LUKASCZYK J., WEBER G., MACIEJEWSKI R., GARTH C., LEITTE H.: Nested tracking graphs. *Computer Graphics Forum* 36, 3 (2017), 12–22. 2

- [LZD*19] LI G., ZHANG Y., DONG Y., LIANG J., ZHANG J., WANG J., MCGUFFIN M. J., YUAN X.: BarcodeTree: Scalable comparison of multiple hierarchies. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 1022–1032. 2
- [Mee16] MEERTENS INSTITUUT, KNAW: Nederlandse voornamenbank. <https://www.meertens.knaw.nl/nvb>, accessed 30-05-2016. 7
- [NA07] NAGAMUCHI H., ABE Y.: An approximation algorithm for dissecting a rectangle into rectangles with specified areas. *Discrete Applied Mathematics* 155, 4 (2007), 523–537. 1, 2
- [S*94] SHI J., ET AL.: Good features to track. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (1994), pp. 593–600. 4
- [Shn92] SHNEIDERMAN B.: Tree visualization with tree-maps: a 2D space-filling approach. *ACM Transactions on Graphics* 11, 1 (1992), 92–99. 3
- [SHS11] SCHULZ H.-J., HADLAK S., SCHUMANN H.: The design space of implicit hierarchy visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics* 17, 4 (2011), 393–411. 2
- [SMBWL14] SMITH-MILES K., BAATAR D., WREFORD B., LEWIS R.: Towards objective measures of algorithm performance across instance space. *Computers & Operations Research* 45 (2014), 12–24. 5
- [SSV18] SONDAG M., SPECKMANN B., VERBEEK K.: Stable treemaps via local moves. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 729–738. 3, 4
- [STL*20] SCHEIBEL W., TRAPP M., LIMBERGER D., DÖLLNER J.: A taxonomy of treemap visualization techniques. In *Proc. International Conference on Information Visualization Theory and Applications* (2020). 2
- [SW01] SHNEIDERMAN B., WATTENBERG M.: Ordered treemap layouts. In *Proc. IEEE Symposium on Information Visualization* (2001), pp. 73–78. 3, 4
- [SWD18] SCHEIBEL W., WEYAND C., DÖLLNER J.: EvoCells – a treemap layout algorithm for evolving tree data. In *Proc. International Conference on Information Visualization Theory and Applications* (2018), pp. 273–280. 2, 4
- [Sze10] SZELISKI R.: *Computer Vision: Algorithms and Applications*. Springer, 2010. 4
- [TC13] TAK S., COCKBURN A.: Enhanced spatial stability with Hilbert and Moore treemaps. *IEEE Transactions on Visualization and Computer Graphics* 19, 1 (2013), 141–148. 3, 4
- [Tel06] TELEA A.: Combining extended table lens and treemap techniques for visualizing tabular data. In *Proc. VGTC Conference on Visualization* (2006), pp. 120–127. 7
- [TM07] TUYTELAARS T., MIKOLAJCZYK K.: Local invariant feature detectors: A survey. *Foundations and Trends in Computer Graphics and Vision* 3, 3 (2007), 177–280. 4
- [TS07] TU Y., SHEN H.: Visualizing changes of hierarchical data using treemaps. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1286–1293. 3
- [URLa] Scitools. <https://scitools.com>. 7
- [URLb] Treemap resources. <https://eduardovernier.github.io/dynamic-treemap-resources-eurovis>. 2, 3, 4, 10
- [URL17] UN Comtrade Database. <https://comtrade.un.org>, accessed 15-02-2017. 7
- [URL18a] ATP Tennis Rankings. https://github.com/JeffSackmann/tennis_atp, accessed 03-07-2018. 7
- [URL18b] USGS Earthquakes. <https://earthquake.usgs.gov/earthquakes/browse/stats.php>, accessed 03-07-2018. 7
- [URL18c] Worldbank indicators. <https://data.worldbank.org/indicator/>, accessed 04-07-2018. 6
- [URL18d] The Movie Database. www.themoviedb.org, accessed 10-02-2018. 7
- [URL18e] Github. <https://github.com>, accessed 16-07-2018. 7
- [VCT18a] VERNIER E., COMBA J., TELEA A.: Quantitative comparison of dynamic treemaps for software evolution visualization. In *IEEE Conference on Software Visualization* (2018), pp. 96–106. 4, 6
- [VCT18b] VERNIER E., COMBA J., TELEA A.: A stable greedy insertion treemap algorithm for software evolution visualization. In *IEEE Conference on Graphics, Patterns and Images* (2018), pp. 158–165. 3, 4
- [vHH17] VAN HEES R., HAGE J.: Stable and predictable Voronoi treemaps for software quality monitoring. *Information and Software Technology* 87 (2017), 242 – 258. 2
- [VvWvdL06] VLIENEN R., VAN WIJK J. J., VAN DER LINDEN E. J.: Visualizing business data with generalized treemaps. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 789–796. 6, 7
- [Wat05] WATTENBERG M.: A note on space-filling visualizations and space-filling curves. In *Proc. IEEE Symposium on Information Visualization* (2005), pp. 181–186. 1, 3
- [ZCYT17] ZHOU M., CHENG Y., YE N., TIAN J.: Effectiveness and efficiency of using different types of rectangular treemap as diagrams in cartography. In *International Cartographic Conference* (2017), pp. 187–206. 2