

Quantifying the Space-Efficiency of 2D Graphical Representations of Trees

Michael J. McGuffin and Jean-Marc Robert

Abstract— A mathematical evaluation and comparison of the space-efficiency of various 2D graphical representations of tree structures is presented. As part of the evaluation, a novel metric called the *mean area exponent* is introduced that quantifies the distribution of area across nodes in a tree representation, and that can be applied to a broad range of different representations of trees. Several representations are analyzed and compared by calculating their mean area exponent as well as the area they allocate to nodes and labels. Our analysis inspires a set of design guidelines as well as a few novel tree representations that are also presented.

Index Terms—Tree visualization, graph drawing, efficiency metrics.

1 INTRODUCTION

A variety of graphical representations are available for depicting tree structures (Figure 1), from “classical” node-link diagrams [23, 7], to treemaps [14, 26, 6, 30], concentric circles [2, 27, 31], and many others (see [13] for a survey). A major consideration when designing, evaluating, or comparing such representations is how efficiently they use screen space to show information about the tree. To date, however, it is unclear how to go about evaluating space-efficiency in a way that can be applied to the large variety of tree representations and that enables a fair comparison of them. Space-efficiency might be described in terms of area, aspect ratio, label size, or other measures. However, there is no accepted standard set of metrics for evaluating the space-efficiency of tree representations, and it is unclear what approach would be general enough to be applied to all the forms in Figure 1.

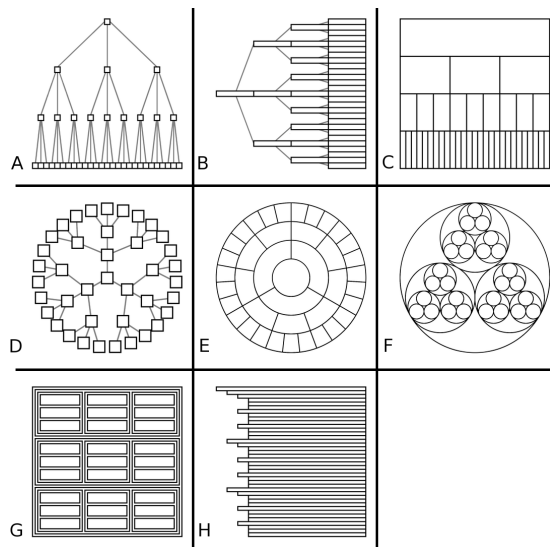


Fig. 1. Several basic kinds of tree representations, here each showing a complete 3-ary tree of depth 3 as an example. All representations are drawn to just fit within a 1×1 unit square. A: classical (layered) node-link [23, 7]. B: a variation on A, where the shape of nodes better accommodates long labels. C: icicle. D: radial [10, 9]. E: concentric circles [2, 27, 31]. F: nested circles, similar to [5, 28]. G: treemap [14, 26]. H: indented outline, sometimes called a “tree list”, and common in file browsers such as Microsoft Explorer.

One basic metric of space-efficiency is the total area of a representation. Assuming the representation is bound within a 1×1 square, both icicle diagrams and treemaps (Figures 1C and 1G) have a total area of 1, and are equally efficient (and both optimal) according to this metric. Likewise, concentric circles and nested circles (Figures 1E and 1F) both have a total area of $\pi/4 \approx 0.785$ (the area of a circle of diameter 1), and are also equally efficient according to the metric of total area. However, experience suggests that the representations within each of these pairs do not scale equally well with larger, deeper trees. This article shows that there are finer ways of distinguishing efficiency, i.e. that there is more to space-efficiency than total area.

Treemaps are often described as optimally space-efficient, not just because they have a total area of 1, but also because they allow for what we call a *weighted partitioning* of the area. Nodes can be allocated more or less area, depending on some attribute such as file size, population, or number of species, and furthermore this weighted partitioning can be done without reducing the total area used. These are indeed desirable properties, however they are not unique to treemaps. Figure 2 shows that icicle diagrams also allow for a weighted partitioning of area, and incidentally have no need for margins between the borders of nodes as treemaps often do.

Furthermore, although a weighted partitioning is useful for showing the relative sizes of nodes in Figures 2A and 2C, an unfortunate side effect is that labels on small nodes are very difficult to read. If users are more interested in seeing the identity of all nodes rather than their relative sizes, an alternative approach would be to give equal weight to each leaf node (Figures 2B and 2D), improving the overall legibility of nodes. (Although not shown in the figure, the labels could also be augmented to numerically show the “size” attribute of each node.) In terms of label size or legibility, Figures 2B and 2D are clearly preferable, but even they still result in much whitespace around certain labels, suggesting that a more space-efficient (in terms of label size) representation might be possible.

Clearly, it would be useful to have some way to quantitatively distinguish the four possibilities in Figure 2, e.g. in terms of their respective scalability and the sizes of their labels. If total area is the only metric of space-efficiency used, and “optimal” space-efficiency is defined as a total area of 1 (possibly partitioned by weight), then we have no way of distinguishing these four cases. If alternative metrics of space-efficiency are used, such as those investigated in this article, it is not clear initially if treemaps, or any other representation, will still turn out to be optimal with respect to such alternative metrics.

This article identifies several metrics related to space-efficiency, and performs the first rigorous analysis and comparison of the space-efficiency of most of the basic tree representation styles in the information visualization literature, including all those in Figure 1. Some of the key ideas involved are (1) the use of a metric of the size of the smallest nodes (i.e. the leaf nodes) in the representation, in addition to a metric of total area; (2) analyzing the area of *labels* on the nodes, which implicitly takes into account both the size and aspect ratio of the nodes, measuring how much “useful” area they contain; and (3) analyzing how these metrics behave asymptotically, as the tree grows

- Michael J. McGuffin is with École de technologie supérieure, Montréal, Canada, E-mail: michael.mcguiffin@etsmtl.ca.
- Jean-Marc Robert is with École de technologie supérieure, Montréal, Canada, E-mail: jean-marc.robert@etsmtl.ca.

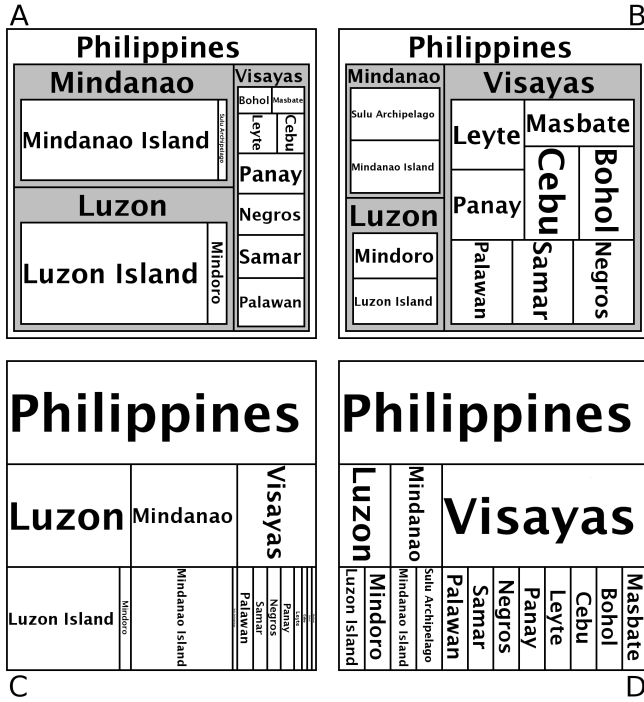


Fig. 2. A tree of regions and major islands of the Philippines, drawn using squarified treemaps (top) and using icicle diagrams (bottom). The two diagrams on the left weight leaf nodes by geographic area, whereas the two diagrams on the right give equal area to leaf nodes. Labels are rotated when necessary to maximize their size.

arbitrarily deep. Our analysis allows us to rank tree representations by their efficiency, which is useful for helping designers choose the most efficient representation allowable within other given constraints.

Our work also quantifies an interesting difference between representations in how they distribute area across nodes. For example, the icicle diagrams in Figures 1C, 2C, and 2D allocate equal area to each level of the tree: the root node has the same area as all the leaf nodes together. Treemaps, in contrast, typically allocate more area to deeper nodes. There is a tradeoff here, since we would like users to be able to see as many deep nodes as possible (which tend to also be the most numerous nodes), while at the same time providing some information about shallow nodes (for example, to give an overview of subtrees, and/or to guide the user in zooming operations). This article develops a new metric, the *mean area exponent*, that describes the distribution of area across levels of a tree representation, to quantify this tradeoff.

Finally, we also present a set of design guidelines for using tree representations, as well as a few novel tree representations, including a variation on squarified treemaps that allows for larger labels within the nodes.

2 RELATED WORK

Different tree representations, including classical node-link, icicle, nested enclosure, and indented outline, were identified decades ago in [4, 16], and an interactive version of the indented outline representation (now popular in file browsers such as Microsoft Explorer) was presented in [11]. Subsequent years have seen variations on these representations proposed. Treemaps are a relatively recent innovation, and are a kind of nested enclosure representation. Treemaps are often described as *space-filling*, a highly desirable property for space-efficiency.

The term “space-filling” can sometimes be problematic, however. For example, a view sometimes expressed [21] holds that tree representations can be divided into two classes: (1) node-link diagrams, that illustrate parent-child relationships with line segments or curves, and

(2) space-filling representations, which include treemaps and concentric circles such as Sunburst (Sunburst was described as space-filling in [27], and [21] similarly describe [2] as space-filling.) However, these two classes seem to not be disjoint, because some node-link diagrams also “fill space” [19, 20]. The 2nd class also ignores an interesting difference between treemaps and concentric circles, namely that parent nodes in treemaps *enclose* their children, whereas parents in concentric circle diagrams are *adjacent* to their children. Finally, the term “space-filling” suggests increased space-efficiency, however it is easy to design a treemap layout algorithm that *occupies* all available space without making good *use* of it, for example, by using excessively thick margins, or by concentrating child nodes in only one corner of their parent, leaving the rest of the parent empty and unused. Would such a treemap cease to be considered space-filling, even though its root node covers all the available space? Without a precise definition of “space-filling”, we recommend being cautious about using this term to refer to a category of tree representations, since the name seems to imply that members of the category are more space-efficient than non-members. As an alternative, categories of representations could instead be based on how the nodes are drawn (e.g. representations where the nodes are mapped to points, and those where the nodes are mapped to areas) or on how parent-child relationships are shown (e.g. through line segments, enclosure, adjacency, or relative positioning). The *space-efficiency* of a given representation can be treated as a separate matter, and evaluated by several metrics, as demonstrated in this article.

Within the graph drawing community, a common approach for evaluating space-efficiency is to compare the total area required by different drawings (i.e., representations) of the same graph or tree. Since any drawing can be scaled arbitrarily in x and y , to ensure a meaningful comparison, the “resolution” of the representations is fixed, often by requiring that nodes be positioned on a grid (i.e. with integer coordinates) [9]. There are problems with this general approach, however, especially when comparing representations of trees rather than graphs. For example, allowing only grid positions may be misleading, because non-integer coordinates can significantly reduce total area without compromising the clarity of the representation or the space available for labels (Figure 3). As a potential remedy, instead of positioning nodes on a grid, we might instead impose a minimum distance between nodes, or a minimum size for non-overlapping labels centered over the nodes. Unfortunately, matters are complicated by the fact that some tree representations (such as Figures 1C, 1E, 1F, 1G) involve nodes that have an area and shape, and there may be nodes and labels of different sizes within a single representation (e.g. deeper nodes may be smaller and have smaller labels). This makes it less clear how to impose a fixed resolution in a way that is fair across tree representations. Note that this issue does not arise in traditional graph drawing, where nodes are typically mapped to points.

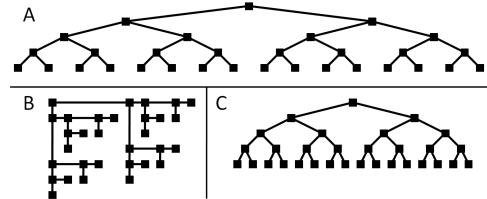


Fig. 3. A and B are adapted from a comparison in Figure 5 of [1], and show two different graphical representations of the same tree where nodes are constrained to positions with integer coordinates. B is clearly more compact than A. In C, however, we have redrawn the representation from A with the integer coordinate constraint relaxed, and the resulting graphical representation has a convex hull whose area is only about 5% greater than that in B. Notice also that the minimum horizontal spacing between nodes in B and C is the same, allowing nodes to be overlaid with horizontally oriented labels of the same size in both cases.

In our work, rather than comparing total area with a fixed resolution, we fix the total area available, and fix its aspect ratio. Representations

are drawn within a fixed, bounding 1×1 unit square, and we analyze the resulting area allocated to individual nodes as well as the resulting size of labels. In a mathematical sense, this can be made equivalent to the fixed-resolution approach, but we find our approach simpler. The aspect ratio of typical computer screens, of typical windows, and of the human visual field is approximately 1. By imposing the same square-shaped total area on each representation, it is clear that comparisons will be fair, while also allowing the representation to position nodes and subdivide area between nodes in many different ways.

The graph drawing community has developed algorithms for drawing trees, with nodes positioned on grid points, with optimal total area [12, 25]. In the output of these algorithms, however, the relative positioning of sibling nodes varies greatly, making it difficult to perceive the overall tree structure. We suspect these algorithms are unlikely to gain widespread use, and focus on analyzing more conventional, visually uniform, symmetrical depictions of trees.

Several algorithms exist for generating treemap layouts, each of which has advantages and disadvantages in terms of criteria such as stability, preservation of ordering of nodes, aspect ratio of nodes, and shape of nodes. One recent example with a mathematical treatment is [30]. All of these layout algorithms allow for a weighted partitioning of the total area (and, interestingly, [30] identifies “space-filling” with the ability to perform a weighted partitioning). In this article, we do not consider stability or preservation of ordering, although these two criteria are trivially met for many of the tree representations we consider. Unlike the previous work, however, we do examine label area, which implicitly takes into account both the size and aspect ratio of nodes.

The efficiency of “inclusion” tree layouts (another nested enclosure representation, somewhat similar to treemaps) has also been investigated [22], but not in a way that is easily compared to other kinds of tree representations.

Previous work on labeling drawings of graphs [15] has assumed the labels to be drawn with constant-size text, outside of the elements to label. In trees, however, the number of elements to label grows exponentially with depth, and labels placed outside nodes can introduce significant clutter and area requirements. Our work assumes labels will be placed inside nodes, and allows for labels of varying size, corresponding to how nodes in trees are typically labeled in practice.

3 METRICS RELATED TO SPACE-EFFICIENCY

We use $n_{i,j}$ to denote the j th node at depth i in the tree. For example, if $n_{i,1}$ has P children, they will be $n_{i+1,1}, \dots, n_{i+1,P}$; and if the next node $n_{i,2}$ has Q children, they will be $n_{i+1,P+1}, \dots, n_{i+1,P+Q}$. We also simply use n_i to denote some node at depth i . For example, n_0 and $n_{0,1}$ both denote the root node. Finally, in the 2D representation of the tree, the region corresponding to a node has area $a(n_{i,j})$ or $a_{i,j}$, or simply a_i if all nodes at depth i have the same area.

3.1 Total Area A

Total area is a basic metric of space-efficiency, which we define here more precisely. To ensure a fair comparison between representations, we require that the 2D representation be bounded by a 1×1 square. The total area A occupied is simply

$$A = \sum_{i,j} a_{i,j} = \sum_{n \text{ node}} a(n)$$

To ensure $0 \leq A \leq 1$, we require that no area be counted twice, i.e. that we subtract away overlap. For example, in representations involving nested enclosure such as treemaps, we require that $a(n)$ be the area corresponding to n after subtracting away any overlapping area assigned to descendant nodes. In a treemap, if node n has thin margins surrounding its children, then $a(n)$ will be the area of the margins, and will not include the children. Not doing this would give an unfair advantage to representations with overlapping areas.

3.2 Leaf Node Area

Another simple metric is based on the size of the smallest nodes in the representation, which typically means the leaf nodes. This metric

could be defined as the average (or minimum) area of the leaf nodes in the representation. In the case of a tree where all the leaf nodes are at the same depth D , and they all have the same area, then the metric would simply be $a_D = a(n_D)$.

In typical representations, as a tree grows deeper, eventually the leaf nodes become so small that they are barely visible, and any labels they contain are illegible. The leaf node area $a(n_D)$ can be thought of as a metric of roughly how much zooming the user must perform to make the leaf node labels legible.

3.3 Mean Area Exponent ρ

Next, we define a novel metric ρ that quantifies the distribution of area across levels of a tree representation, independent of the total area A . In this subsection, we limit our consideration to “homogeneous” tree representations (with constant branching factor, leaf nodes all at the same depth, and nodes within the same level having the same area), and in Section 3.6 we generalize to other representations.

Consider a complete B -ary tree of depth $D \geq 1$, with $B \geq 2$. In other words, a tree where every non-leaf node has B children, and the path between every leaf node and the root has D edges. Assume also that nodes within the same level have the same area, i.e. for each $i = 0, 1, \dots, D$, all nodes at depth i have the same area a_i .

We first define a local metric $\rho(p;c)$ associated with a parent node p and one of its children c , or in other words associated with the edge between p and c . We seek a definition of $\rho(p;c)$ that satisfies the following conditions:

$$\rho(p;c) = \begin{cases} -1 & \text{if } a(c) = a(p) \cdot B \\ 0 & \text{if } a(c) = a(p) \\ +1 & \text{if } a(c) = a(p)/B \end{cases}$$

$$\rho(p;c) < 0 \quad \text{if } a(c) > a(p)$$

$$\rho(p;c) > 0 \quad \text{if } a(c) < a(p)$$

A value of -1 corresponds to a child occupying B times the area of its parent; a value of 0 corresponds to the child and parent occupying equal area; and a value of $+1$ corresponds to a child occupying $1/B$ the area of its parent. (In this 3rd case, the B children of a parent collectively occupy the same total area as their parent, which is seen, for example, in Figure 1C.) The definition of the local metric that we will use is

$$\rho(p;c) = \log_B \frac{a(p)}{a(c)}$$

(Intuitively, the use of a logarithm here is related to the fact that trees grow exponentially with depth.)

In the tree we are considering, because all nodes within the same level have the same area, we can associate the same ρ_i with all parent nodes at depth i and their children at depth $i+1$, where $\rho_i = \rho(n_i; n_{i+1}) = \log_B(a_i/a_{i+1})$. We then define the global metric ρ for the entire tree simply as an arithmetic mean:

$$\rho = \frac{1}{D} \sum_{i=0}^{D-1} \rho_i = \frac{1}{D} \sum_{i=0}^{D-1} \log_B \frac{a_i}{a_{i+1}}$$

We will refer to this global metric as the *mean area exponent* ρ (this name is explained in Section 3.5), and we can rewrite it as

$$\rho = \frac{1}{D} \sum_{i=0}^{D-1} \log_B \frac{a_i}{a_{i+1}}$$

$$= \log_B \sqrt[D]{\prod_{i=0}^{D-1} \frac{a_i}{a_{i+1}}} \quad (1)$$

$$= \log_B \sqrt[D]{\frac{a_0}{a_D}} \quad (2)$$

Interestingly, the expression (1) has a simple interpretation: it is the logarithm of the geometric mean of the ratios of areas of nodes on

consecutive levels. Expanding the product in expression (1) allows many terms to cancel, yielding the rather simple expression (2) that can be used to compute the metric. As will be seen, such canceling of terms will not be possible in the more general case of heterogeneous tree representations.

3.4 Example Analysis of Figure 4

To illustrate the calculation of the metrics A , $a(n_D)$, and ρ , we evaluate them in the cases shown in Figure 4.

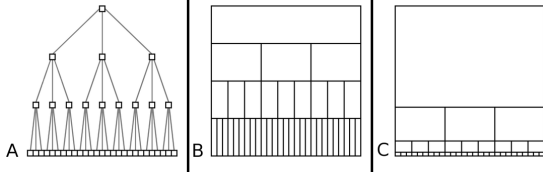


Fig. 4. Example representations of a tree with $B = 3$ and $D = 3$. A: classical node-link. B: icicle. C: a variation on icicle, where each child node is scaled-down with respect to its parent in both the x and y directions by the same factor, allowing for a larger label on the root.

Because the tree being considered is complete and B -ary, there are B^i nodes at depth i for $0 \leq i \leq D$, and $(B^{D+1} - 1)/(B - 1)$ nodes in total. Also, for each representation in Figure 4, all nodes at a given depth have congruent rectangular regions. Let w_i , h_i , and $a_i = w_i h_i$ be the width, height, and area of the rectangular region corresponding to a node at depth i .

In Figure 4A, because there are B^D leaf nodes, and they are squares, they each have width and height $w_D = h_D = 1/B^D$. Since all nodes are the same size, $a_i = w_i h_i = a_D = w_D h_D = 1/B^{2D}$ and the total area occupied is

$$A = \sum_{i=0}^D B^i a_i = \sum_{i=0}^D B^i \frac{1}{B^{2D}} = \frac{(B^{D+1} - 1)}{(B - 1)} \frac{1}{B^{2D}} = \Theta(1/B^D)$$

where $\Theta()$ denotes an asymptotically tight bound¹. The mean area exponent, which quantifies the distribution of area across levels of the tree, is

$$\rho = \log_B \sqrt[D]{\frac{a_0}{a_D}} = \log_B \sqrt[D]{\frac{1/B^{2D}}{1/B^{2D}}} = 0$$

Next, consider Figure 4B. Each node has $1/B$ the area of its parent. Clearly, $w_i = 1/B^i$, $h_i = 1/(D+1)$, $a_i = 1/((D+1)B^i)$, and

$$A = \sum_{i=0}^D B^i a_i = \sum_{i=0}^D B^i \frac{1}{(D+1)B^i} = 1$$

The leaf node area $a(n_D) = 1/((D+1)B^D) = \Theta(1/(DB^D))$, and the mean area exponent is

$$\rho = \log_B \sqrt[D]{\frac{a_0}{a_D}} = \log_B \sqrt[D]{\frac{1/((D+1)B^0)}{1/((D+1)B^D)}} = \log_B \sqrt[D]{B^D} = 1$$

Finally, consider Figure 4C. Although this representation might not be used in practice, it illustrates a situation where the distribution of area is more heavily towards the root. Each node now has $1/B$ the width of its parent, $1/B$ the height of its parent, and $1/B^2$ the area of its parent. Specifically, $w_i = 1/B^i$, $h_i = (B-1)B^{D-i}/(B^{D+1}-1)$, and $a_i = (B-1)B^{D-2i}/(B^{D+1}-1)$. It is easy to check that $A = 1$, $a(n_D) = (B-1)/((B^D)(B^{D+1}-1)) = \Theta(1/B^{2D})$, and $\rho = 2$.

To summarize the results:

¹Asymptotic analysis in this article is always with respect to D , i.e., expressions that approach zero as $D \rightarrow \infty$, with B held constant. If $f(D)$ and $g(D)$ both approach zero as D approaches infinity, then the statement $f(D) = \Theta(g(D))$ means there exist constants $k_1, k_2 > 0$ such that $k_1 g(D) < f(D) < k_2 g(D)$ for all sufficiently large D .

Figure	4A	4B	4C
total area A	$\Theta(1/B^D)$	1	1
leaf node area $a(n_D)$	$1/B^{2D}$	$\Theta(1/(DB^D))$	$\Theta(1/B^{2D})$
mean area exponent ρ	0	1	2

In terms purely of total area A , Figures 4B and 4C are preferable to 4A. However, total area does not help us distinguish 4B and 4C. This distinction is provided by the other two metrics. The representation with the most slowly shrinking leaf area (as a function of D) is 4B. Notice also that 4A and 4C have the same leaf area (in terms of Θ -complexity, ignoring hidden constants), but different ρ values, because the value of ρ is based on the areas of leaf nodes *and* of the root node. Larger values of ρ mean that more area is allocated to the root. In other words, ρ is a measure of how heavily the representation's total area is weighted toward the root, or of how quickly node area decreases with depth.

3.5 Other Metrics Related to ρ

A deeper understanding of ρ can be achieved by briefly considering some closely related metrics of area distribution. We can define an alternative local metric $R(p; c) = a(p)/a(c)$ that is simply a ratio of areas, and a global metric $R = (\prod_{i=0}^{D-1} a_i/a_{i+1})^{1/D}$ that is the geometric mean of these ratios. It is easy to check that $\rho = \log_B R$, i.e. $B^\rho = R$, which motivated calling ρ the mean area *exponent*. Furthermore, whereas R and ρ are defined in terms of ratios of areas, we can define analogous metrics that are ratios of lengths. The four resulting global metrics R , ρ , S , and σ are summarized hence:

	ratios	logarithms of ratios
2D area	mean area ratio $R = B^\rho$	mean area exponent $\rho = \log_B R$
1D length	mean scale factor $S = B^\sigma = \sqrt{R}$	mean scale exponent $\sigma = \log_B S = \rho/2$

These four metrics all measure the distribution of area in the tree representation. Of these, $S = \sqrt{R}$ has perhaps the most familiar geometric interpretation, since it is simply an average scale factor. However, we prefer to use ρ in our analyses, firstly because it is more often independent of the particular branching factor B than R or S are, and secondly because ρ is more often an integer than σ (due to the fact that we are analyzing 2D tree representations).

Incidentally, we also experimented with using an *arithmetic* mean of ratios, i.e. $(1/D) \sum_{i=0}^{D-1} a_i/a_{i+1}$, and found that such a formulation resulted in less elegant properties than will be presented in Section 3.6.

3.6 Generalizing ρ to Heterogeneous Tree Representations

We now extend the definition of ρ to complete B -ary trees where the areas of sibling nodes are not equal. (Later we extend to incomplete trees with variable branching factor.) In such a tree, let $n_{i,j}$ denote the j th node at depth i , where $0 \leq i \leq D$ and $1 \leq j \leq B^i$. The children of $n_{i,j}$ are $n_{i+1,(j-1)B+1}, \dots, n_{i+1,(j-1)B+B}$. We define the value of ρ associated with any non-leaf node as

$$\rho(n_{i,j}) = \frac{1}{B} \sum_{k=1}^B \rho(n_{i,j}; n_{i+1,(j-1)B+k})$$

This is simply the arithmetic mean of the local $\rho(p; c)$ defined earlier, over all the node's children.

Next, the value ρ_i characterizing the distribution of area across the levels i and $i+1$ of a tree representation is itself an arithmetic mean of the previous quantity:

$$\rho_i = \frac{1}{B^i} \sum_{j=1}^{B^i} \rho(n_{i,j})$$

Finally, the global metric ρ is defined as

$$\rho = \frac{1}{D} \sum_{i=0}^{D-1} \rho_i \quad (3)$$

Combining these three equations yields an expression that can be rewritten in a few different ways. (In reading the following, it helps to keep in mind that i is an index over levels, j an index over parents in a given level, and k an index over children.)

$$\rho = \frac{1}{D} \sum_{i=0}^{D-1} \frac{1}{B^i} \sum_{j=1}^{B^i} \frac{1}{B} \sum_{k=1}^B \log_B \frac{a_{i,j}}{a_{i+1,(j-1)B+k}} \quad (4)$$

$$= \frac{1}{D} \sum_{i=0}^{D-1} \frac{1}{B^i} \sum_{j=1}^{B^i} \log_B \sqrt[B]{\prod_{k=1}^B \frac{a_{i,j}}{a_{i+1,(j-1)B+k}}} \quad (5)$$

$$= \frac{1}{D} \sum_{i=0}^{D-1} \frac{1}{B^i} \sum_{j=1}^{B^i} \log_B \frac{a_{i,j}}{\sqrt[B]{\prod_{k=1}^B a_{i+1,(j-1)B+k}}} \quad (6)$$

$$= \frac{1}{D} \sum_{i=0}^{D-1} \log_B \frac{\sqrt[B^i]{\prod_{j=1}^{B^i} a_{i,j}}}{\sqrt[B^{i+1}]{\prod_{j=1}^{B^{i+1}} a_{i+1,j}}} \quad (7)$$

$$= \log_B \sqrt[D]{\frac{a_{0,1}}{\sqrt[B^D]{\prod_{j=1}^{B^D} a_{D,j}}}} \quad (8)$$

Each of the above expressions has a fairly simple interpretation. The radical in expression (5) is the geometric mean of the ratio of the (i, j) th node's area and each of its children's areas. The inner ratio in expression (6) is between the (i, j) th node's area and the geometric mean of its children's areas. The inner ratio in expression (7) is between the geometric mean of areas of all parents on level i and the geometric mean of areas of all children on level $i+1$. Expanding expression (7) results in many terms canceling, yielding expression (8), which involves a ratio between the root's area and the geometric mean of the leaf areas. Note that an alternatively defined metric, based on *arithmetic* means of areas or *arithmetic* means of ratios, would not lead to the kind of multiple, equivalent expressions shown above.

Notice that in equation (3) (and expression (7)), each level's ρ_i contributes with equal weight, even though deeper levels have many more nodes. The root's area has a much larger impact on ρ than any of the B^D leaf nodes. To make the unequal weighting of nodes more apparent, we can rewrite equation (4) as

$$\rho = \frac{1}{DB^D} \sum_{i=0}^{D-1} \sum_{j=1}^{B^i} \sum_{k=1}^B B^{D-(i+1)} \log_B \frac{a_{i,j}}{a_{i+1,(j-1)B+k}} \quad (9)$$

which is simply the *weighted* arithmetic mean of the local metric $\rho(n_{i,j}; n_{i+1,(j-1)B+k})$, where the local metric for an edge between levels i and $i+1$ has weight $B^{D-(i+1)}$, and the sum of the weights is DB^D .

There is another way we can define ρ , that is equivalent to the previous definitions, but that is more easily extended to other kinds of trees such as incomplete trees. In equation (9), the weight $B^{D-(i+1)}$ assigned to an edge between levels i and $i+1$ is equal to the number of paths through that edge that connect the root node to a leaf node. If we consider all paths between the root and leaf nodes, and for each path we find the arithmetic mean of $\rho(p, c)$ along that path, and then find the (unweighted) arithmetic mean over all the paths, we end up with the same quantity ρ defined earlier. The weighting of edges seen in equation (9) arises “for free” because many paths pass through shallow edges, causing them to contribute many times.

To make this precise, let $\text{path}(a, b)$ be the set of nodes along the path from node a to node b , and let $p(n)$ be the parent of node n . If n is one of the B^D leaf nodes, then $\text{path}(n_0, n)$ is the path from the root to leaf n . We define

$$\rho = \frac{1}{B^D} \sum_{n \text{ leaf}} \frac{1}{D} \sum_{m \in \text{path}(n_0, n) \setminus \{n_0\}} \log_B \frac{a(p(m))}{a(m)} \quad (10)$$

Not only is equation (10) equivalent to equation (9), it is also straightforward to generalize to incomplete trees with varying branching factor. Given a node n , let $B(n)$ be the branching factor at n , i.e. the

number of children of n , and let $D(n)$ be the depth of n , i.e. the number of edges in the path from n to the root n_0 . Also, let \mathcal{L} be the set of leaf nodes in the tree. Then we define a more general mean area exponent as

$$\rho = \frac{1}{|\mathcal{L}|} \sum_{n \in \mathcal{L}} \frac{1}{D(n)} \sum_{m \in \text{path}(n_0, n) \setminus \{n_0\}} \log_{B(p(m))} \frac{a(p(m))}{a(m)} \quad (11)$$

which subsumes the previous definitions as special cases, and which can be applied to a fully heterogeneous tree representation². We note that the canceling of terms that resulted in simpler expressions (2) and (8) cannot be performed here, partly because the logarithm in equation (11) does not have a constant base.

The multiple equivalent formulations and interpretations we have presented show that the mean area exponent ρ exhibits a promising degree of conceptual versatility as a metric.

As a minor additional extension to our definitions, if the root node's area is zero and the leaf nodes all have positive area, then we define $\rho = -\infty$. In the opposite case, if the root node has positive area and the leaf nodes all have zero area, then we define $\rho = +\infty$.

We conclude this section with a simple example. Consider the icicle diagram in Figure 5A, where the width of each parent node is evenly divided amongst its children. As can be confirmed by the reader, applying Equation 11 reveals that $\rho = 1$ (exactly). In Figure 5B, however, each node is given a width proportional to the number of leaf nodes under it, causing most leaf nodes to be larger than in the former case, hence the metric $\rho \approx 0.884$ is lower.

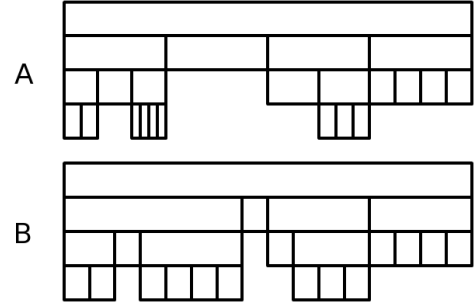


Fig. 5. Two icicle diagrams of the same tree. The top diagram (A) allocates more of its total area to shallow nodes, as quantified by its higher value of ρ .

3.7 Additional Metrics, and Example Analysis of Figure 2

So far, our metrics have considered the area of nodes in the representation. In many situations, however, the labels within the nodes (whether they consist of strings, images, or some combination) are at least as important to the user as the structure of the tree. The legibility of these labels depends on their size, which is not generally equal to the size of the nodes they appear on. As can be seen in Figure 2, a mismatch in the aspect ratios of a label and its node can leave much whitespace unused within the node. Hence, for each of the three metrics seen so far (total area, leaf node area, and mean area exponent), we define a variant based on the areas of labels, resulting in six metrics: the three original metrics based on node areas, and three label-oriented metrics.

Let $a_l(n) \leq a(n)$ be the area of (the rectangular bounding box of) the label embedded in the area of node n . The first label-oriented metric we define is the total label area

$$A_l = \sum_{n \text{ node}} a_l(n) \leq A \leq 1$$

²A minor caveat: due to the base of the logarithm, ρ is not defined if any node n has only one child, i.e. $B(n) = 1$. This problem does not arise in our analyses, but if it had to be circumvented, the metrics R or S (Section 3.5) could be defined in the fully heterogeneous case without any logarithms, and used instead of ρ .

which is of course analogous to total area A .

The second label-oriented metric quantifies the size of the smallest labels in the representation, which typically means the leaf labels (i.e. labels on the leaf nodes). This second metric could be defined as the average (or minimum) leaf label area, or in the case of a homogeneous representation, it would simply be $a_l(n_D)$.

The third label-oriented metric is a variant of ρ that quantifies how quickly the labels become illegibly small with depth. This variant ρ_l is defined simply by replacing the areas of nodes with the areas of labels in either equation (11) for the general case or equation (2) for the homogeneous case. For example, in the homogeneous case,

$$\rho_l = \log_B \sqrt[p]{\frac{a_l(n_0)}{a_l(n_D)}}$$

The following table lists the values of all six metrics for the 4 heterogeneous cases in Figure 2. In the table, rows 1-3 are for metrics involving node area, and rows 4-6 are for the analogous label-oriented metrics. Row 5 contains 3 different metrics of the size of the smallest labels which will be explained shortly.

	Figure	2A	2B	2C	2D
1	Total area A of nodes	1.000	1.000	1.000	1.000
2	Average area of leaf nodes	0.0465	0.0505	0.0278	0.0278
3	Mean area exponent ρ	0.842	0.462	1.262	0.895
4	Total area A_l of labels	0.366	0.419	0.368	0.591
5	Average area of leaf labels	0.0171	0.0219	0.00524	0.0177
	Average leaf label height	0.0586	0.0791	0.0294	0.0699
	Minimum leaf label height	0.0217	0.0360	0.00821	0.0418
6	Mean area exponent ρ_l for labels	0.755	0.498	1.797	0.871

The metrics based purely on node area (rows 1-3) are of limited value and can even be misleading. For example, the total area $A = 1$ in all four cases, failing to identify any difference in space-efficiency between them. Also, the average area of leaf nodes is exactly equal between cases 2C to 2D, even though it is plain from Figure 2 that there is a difference in how easily the labels can be read.

We recommend instead using metrics based on labels (rows 4-6). Examining these metrics, we notice that, as expected, from Figure 2A to 2B, and again from 2C to 2D, there is an increase in A_l as well as in the average area of leaf labels, and also a decrease in ρ_l . These effects are desirable if we want to improve legibility, rather than show the relative sizes of nodes (as is usually done in treemaps, and done in Figures 2A and 2C).

Comparing all four cases (2A, 2B, 2C and 2D) in terms of label size, we see that 2D has the best A_l value, but 2B has the best value for average area of leaf labels. The higher A_l value for 2D seems due largely to the larger labels it has on shallow nodes, which do not contribute to the legibility of the deep nodes. On one hand, shallow nodes might be considered more important than deep nodes, because they can give the user a high-level overview of the tree structure when there is not enough space to give every individual label a legible size. On the other hand, most nodes in a typical tree are deep, and the shallow nodes only constitute a minority, so making them bigger will impair the legibility of a majority of nodes. Clearly, there is a tradeoff (quantified by ρ_l) in allocating area to shallow or deep labels.

Nevertheless, if our goal is to make *all* the labels as legible as possible (which should be the case at least for static, non-zoomable tree representations), then we should prefer representations where the smallest labels are as large as possible (as quantified in the first sub-row within row 5 of the table). We can further refine this idea for *text* labels (as opposed to labels consisting of images): the legibility of a label is not

best quantified by its area, but rather by its area per character, or alternatively by its font size (hence the other two sub-rows within row 5, concerning “label height”). With this in mind, the most space-efficient representation in the table would either be 2B (whose leaf labels have the best *average* height) or 2D (which has the best *minimum* leaf label height).

Stepping back from the cases in Figure 2, we recapitulate the 6 kinds of metrics, numbered the same way they are in the above table:

1. **Total node efficiency:** the amount of available space that is occupied, as quantified by the total area A . Treemaps and icicle diagrams, for example, are optimal in this sense, since $A = 1$ for them.
2. **Smallest node efficiency:** the degree to which the smallest node(s) are made as large as possible. Since leaf nodes are typically the smallest, this efficiency can normally be quantified by the average (or minimum) area of leaf nodes. In the case of a tree where all leaf nodes n_D have the same area, it can be quantified by their area $a_D = a(n_D)$.
3. **Node area distribution:** a measure of the area tradeoff between shallow and deep nodes. We quantify this with ρ , though the other global metrics in Section 3.5 could play the same role.
4. **Total label efficiency:** the amount of available space devoted to labels, as quantified by the total label area A_l .
5. **Smallest label efficiency:** the degree to which the smallest label(s) are made as legible as possible. This can be quantified by the average (or minimum) area of leaf labels, or, in the case of a tree where all leaf labels have the same area, by $a_l(n_D)$. In the case of labels that are strings of text of variable length, we can instead quantify this with the average (or minimum) area of leaf labels per character, or the average (or minimum) font height of leaf labels.
6. **Label area distribution:** a measure of the area tradeoff between shallow and deep labels, which we quantify with the mean area exponent ρ_l for labels. A value close to zero indicates that all labels are the same size, which is often desirable in static, non-zoomable representations. However, in zoomable representations of deep trees, a value between 1 and 2 is better (Section 5.4 discusses why).

For completeness, we discuss and report all 6 kinds of metrics in this section and the later sections of the article. However, as already mentioned, we recommend using label-oriented metrics (numbers 4-6 in the above list), because they implicitly take into account the aspect ratio of the node area and how well it matches the label’s aspect ratio. An important lesson from our work is that only considering total area A , without regard to labels, is a crude way of quantifying space-efficiency, that fails to distinguish between any of the representations in Figure 2.

Within the label-oriented metrics, even A_l alone is not the best way to quantify the space-efficiency of labels, since a large A_l may be due to a few large labels on shallow nodes while most deep labels are very tiny. Metrics of **smallest label efficiency** should arguably take precedence over A_l . Finally, the mean area exponent ρ_l for labels is not a measure of space-efficiency *per se*, but rather a measure of a tradeoff, and will be discussed more in the later sections of the article.

The metrics listed above can each be evaluated for any given instance of a tree’s representation (i.e. for any given *drawing*). However, we would like to gain insight into the behavior of these metrics for large trees (i.e. as $D \rightarrow \infty$), for a variety of different representation styles (or drawing *conventions*, to use graph drawing terminology), as a function of the aspect ratio of the labels involved. The next section performs an analysis of such behavior, which requires some simplifying assumptions, but which yields more general results than could be obtained by analyzing individual instances of representations such as those in Figure 2.

Since we will be interested in trends in the performance of a representation for trees whose depth D is large, we finally define

$$A_\infty = \lim_{D \rightarrow \infty} A \quad ; \quad \rho_\infty = \lim_{D \rightarrow \infty} \rho$$

$$A_{l,\infty} = \lim_{D \rightarrow \infty} A_l \quad ; \quad \rho_{l,\infty} = \lim_{D \rightarrow \infty} \rho_l$$

Of course, in practical settings, D is always limited by the size of data sets and the resolution of output media. However, the resolution and physical size of common display devices have increased steadily over recent years, and many tree data sets are deep enough to require some amount of zooming before viewing leaf nodes. It seems reasonable to us to evaluate the limit of various metrics as $D \rightarrow \infty$, because representations with better asymptotic efficiency should perform better at a given display resolution and require less zooming.

4 ASYMPTOTIC ANALYSIS OF TREE REPRESENTATIONS

This section analyses the representation styles in Figure 1 and a few other figures. The analysis considers complete B -ary trees with depth $D \geq 1$ and $B \geq 2$, where all nodes at the same depth i have the same area $a_i = a(n_i)$. In addition, we assume that each node contains a rectangular label with aspect ratio $L \geq 1$, where the label must fit within the confines of the node without overlapping descendants, and the label is made as large as possible, rotating the label by an arbitrary angle if necessary to better fit the node. For example, in the leaf nodes of 1C, labels would be oriented vertically, and in the leaf nodes of 1E they would be oriented along radial lines. Note also that L is constant for the entire tree, hence it might be thought of as the maximum length of strings, or the maximum aspect ratio of images.

Of course, in practical settings, real trees seldom have the properties of being complete, with a fixed branching factor B , and labels of fixed aspect ratio. However, these assumptions will allow our analysis to shed significant light on basic differences between the various styles of tree representations.

The results are presented in Table 1. For each representation, the table reports the mean area exponent, area associated with leaf nodes, and total area, for both nodes and labels, including limits as $D \rightarrow \infty$. Tree representations are also ranked by the various metrics.

The last row in Table 1 lists some additional simplifying assumptions that were used in certain cases. These simplify the expressions appearing in the table, focusing attention on behavior in the most important cases. In almost all cases, it would be an easy matter to remove these assumptions and work out a more complete analysis, however this would clutter the table with more complicated results and would probably not yield any important insights.

In the following subsections, we discuss our choice of the representations that were analyzed, and then present the mathematical details of the analyses, case-by-case. Readers uninterested in these details may prefer to skip ahead to Section 5, and the later sections, where the results and their implications are discussed.

4.1 Choice of Representations to Analyze

Initially, we chose to analyze the representations in Figures 1 and 4, because they include key representations from previous visualization literature, as well as some interesting variants. In particular, 1B is a variant of 1A that better accommodates long labels, 4C serves as an example where ρ has a high value (and turns out to have an excellent $A_{l,\infty}$ value, as will be seen), and 1G is an “extreme” treemap where the margins have zero thickness and labels only appear on leaf nodes. (Although such a treemap is ambiguous for homogeneous trees, treemaps without margins and without labels on non-leaf nodes are used in practice, for example in software such as SequoiaView.)

The initial analysis of this set of representations had two outcomes that led to analyzing additional representations. First, it became apparent that leaf nodes are larger if they are laid out along a longer curve (e.g. the leaf nodes in 1E, and their labels, are larger than those in 1C/4B, because in the former case they are laid out along a curve of length π , and in the latter case they are laid out along a curve of length 1). This led us to invent and analyze the representation in Figure 8,

Figure	1(a)/4(a)	1(b)	1(c)/4(b)	4(c)	1(d)	1(e)	1(f)	1(g)	1(h)	8	9	10
ρ	0	0	1	2	0	$1 - \Theta(\frac{1}{B} \log_B D)$	$\rho_\infty - \Theta(\frac{1}{B})$	$-\infty$	$\Theta(\frac{1}{B^D})$	$1 - \Theta(\frac{1}{B} \log_B D)$	$\rho_\infty - \Theta(\frac{1}{B})$	$\log_8 9 - \Theta(\frac{1}{B})$
ρ_∞	0	0	1	2	0	1	$2 \log_B(1/\phi) \leq 2$	$-\infty$	0	1	$\log_{B+1/L} > 1$	$\log_8 9 \approx 1.057$
$a(n_D)$	$\frac{1}{B^D}$	$\Theta(\frac{1}{D B^D})$	$\Theta(\frac{1}{D B^D})$	$\Theta(\frac{B-1}{B B^D})$	$\Theta(\frac{\pi^2}{2 B^D})$	$\Theta(\frac{\pi}{2 B^D})$	$\frac{\pi}{4(1/\phi)^D}$	$\frac{1}{B^D}$	$\Theta(\frac{B-1}{B B^D})$	$\Theta(\frac{2}{D B^D})$	$\left(\frac{1-m}{B+1/L}\right)^D$	$\frac{1}{(B+1)^D}$
Rank	9	5	5	10	8	4	7	1	2	3	6	6
A	$\Theta(\frac{B}{(B-1)B^D})$	$\Theta(\frac{B}{(B-1)B^D})$	1	1	$\Theta(\frac{\pi^2 B}{2(B-1)B^D})$	$\pi/4 \approx 0.785$	$\pi/4 \approx 0.785$	1	$1 - \Theta(\frac{K(B-1)D}{B B^D})$	$1 - \Theta(\frac{1}{B^D})$	1	1
A_∞	0	0	0	0	0	$\pi/4 \approx 0.785$	$\pi/4 \approx 0.785$	1	1	1	1	1
Rank	7	5	1	1	6	4	4	1	2	3	1	1
ρ_l	0	0	$2 - \Theta(\frac{2}{B} \log_B D)$	2	0	$2 - \Theta(\frac{2}{B} \log_B D)$	$\rho_{l,\infty} - \Theta(\frac{1}{B})$	$-\infty$	0	$2 - \Theta(\frac{2}{B} \log_B D)$	$\rho_{l,\infty} - \Theta(\frac{1}{B})$	$\log_8 9 - \Theta(\frac{1}{B})$
$\rho_{l,\infty}$	0	0	2	2	0	2	$2 \log_B(1/\phi) \leq 2$	$-\infty$	0	2	$\log_{B+1/L} > 1$	$\log_8 9 \approx 1.057$
$a_l(n_D)$	$\frac{1}{L B^D}$	$\frac{L}{B^D}$	$\frac{L}{B^D}$	$\frac{1}{L B^D}$	$\Theta(\frac{\pi^2}{2 L B^D})$	$\Theta(\frac{\pi^2 L}{2 B^D})$	$\frac{L}{(L^2+1)(1/\phi)^D}$	$\frac{1}{L B^D}$	$\Theta(\frac{L(B-1)^2}{B B^D})$	$\Theta(\frac{16L}{B^D})$	$\frac{1}{L(B+1)^D}$	$\frac{1}{L(B+1)^D}$
Rank	9	6	6	9	8	5	3	1	7	4	2	2
A_l	$\Theta(\frac{B}{L(B-1)B^D})$	$\Theta(\frac{B L}{(B-1)B^D})$	$\Theta((2 - \frac{1}{L})(\frac{B+1}{B-1})\frac{1}{B})$	$A_{l,\infty} - \Theta(\frac{1}{B^D})$	$\Theta(\frac{\pi^2 B}{2 L(B-1)B^D})$	$\Theta(\frac{\pi}{L^D})$	$A_{l,\infty} + \Theta((B\phi^2)^D)$	$1/L$	$\Theta(\frac{L(B-1)}{B B^D})$	$\Theta(\frac{4}{D^2})$	$A_{l,\infty} + \Theta(k^D)$	$1/L$
$A_{l,\infty}$	0	0	0	$\frac{B}{(B-1)L} \leq 1$	0	0	$\frac{L}{(L^2+1)(1-B\phi^2)} < \frac{1}{2}$	$1/L \leq 1$	0	0	$\frac{(1-m)^2(B+1)L}{(BL+1)^2 - (BL(1-m))^2} < \frac{3}{4}$	$1/L \leq 1$
Rank	11	8	5	1	10	7	4	2	9	6	3	2
Notes				assumes $L \geq 2$			assumes $B \geq 4$	assumes D even		assumes B multiple of 4	assumes D even	$B = 8$

Table 1. The space-efficiency of several tree representations. In each row of rankings, 1 means best. Rankings are based first on upper bounds and limits, and then on Θ -complexity, comparing “hidden” constants (shown explicitly in many cases to aid comparison) to break ties.

where leaf nodes are laid out along a curve of length almost 4 (making its leaf nodes and leaf labels even larger), and whose total area A is nearly optimal. Second, it became apparent that treemaps have promising space-efficiency properties. (Although it has long been acknowledged that treemaps are optimal in terms of total area A , it was initially unclear to us how efficient treemaps would be in terms of label size, and our evaluation of multiple metrics clarified this.) It became clear that we should analyze a treemap where there are margins between nodes and labels on every node, hence we added the case in Figure 9. Finally, we added the case in Figure 10, which is inspired by a kind of fractal, and which we suspected would have excellent space-efficiency and be interesting to compare to treemaps and other representations.

4.2 Case-By-Case Analysis

4.2.1 Case 1A/4A

Section 3.4 already found the values of ρ , a_i , $a(n_D)$, and A for Figure 1A/4A, from which the values of ρ_∞ and A_∞ follow immediately, allowing us to fill several cells in the first column of Table 1. (Note that the asymptotic expressions in the table sometimes explicitly show the (normally hidden) coefficient of the leading term, e.g. showing $A = \Theta(B/(B-1)B^D)$ rather than $A = \Theta(1/B^D)$, to aid in comparing the metrics of various representations.)

We must next determine the area $a_l(n_i)$ of a label on a node at depth i . Because each node is a square with area $a(n_i)$, a label with aspect ratio $L \geq 1$ within the square will have area $a(n_i)/L$ (see Figure 6, top right and bottom right, for examples). Hence $a_l(n_i) = a(n_i)/L = 1/(LB^{2D})$, and the values of ρ_l , $\rho_{l,\infty}$, $a_l(n_D)$, A_l , and $A_{l,\infty}$ all follow by straightforward application of their definitions, completing the entries in the first column of Table 1.

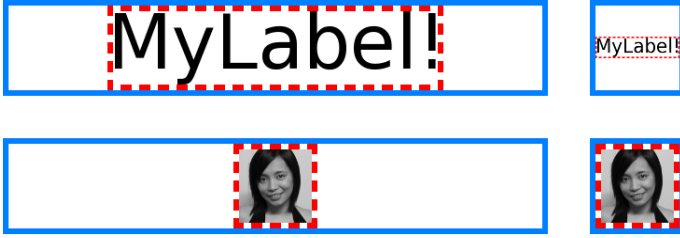


Fig. 6. Labels are shown in dashed red, enclosed by nodes in solid blue. Labels have a fixed aspect ratio L (equal to 1 in the examples of the bottom row, and greater than 1 in the top row). If the aspect ratio of nodes grows arbitrarily large with depth D (the examples at left), then label area $a_l(n_i)$ will be proportional to L . On the other hand, if the aspect ratio of nodes stays fixed as D increases (the examples at right), then label area $a_l(n_i)$ will be inversely proportional to L .

4.2.2 Case 1B

In Figure 1B, as with the previous case, we have nodes all of equal area, and leaf nodes tiling one edge of the unit square (hence $h(i) = 1/B^D$), but now the nodes are rectangles with width $w(i) = 1/(D+1)$, allowing more area to be filled. So $a_i = \frac{1}{(D+1)B^D}$ and

$$A = \sum_{i=0}^D B^i \frac{1}{(D+1)B^D} = \frac{(B^{D+1} - 1)}{(B-1)} \frac{1}{(D+1)B^D} = \Theta\left(\frac{B}{(B-1)D}\right)$$

and $A_\infty = 0$. Also, because all nodes have the same area, we again have $\rho = \rho_\infty = 0$.

To find the area of labels, we note that, unlike the previous case, the aspect ratio of nodes grows arbitrarily large as D increases. This is because the height $h(i)$ of each node decreases much faster than the width $w(i)$. Hence, labels will look like those in Figure 6, top left and bottom left, and will have an area proportional to L (assuming a sufficiently large D , to ensure that the node's aspect ratio is greater than L). In

particular, $a_l(n_i) = L(h(i))^2 = L/B^{2D}$. Because all labels have the same size, it follows that $\rho_l = \rho_{l,\infty} = 0$. Finally

$$A_l = \sum_{i=0}^D B^i \frac{L}{B^{2D}} = \frac{(B^{D+1} - 1)}{(B-1)} \frac{L}{B^{2D}} = \Theta\left(\frac{BL}{(B-1)B^D}\right)$$

and $A_{l,\infty} = 0$.

4.2.3 Case 1C/4B

Section 3.4 found the values of ρ , a_i , $a(n_D)$, and A for this case, from which the values of ρ_∞ and A_∞ follow immediately. Similar to the previous case, nodes become arbitrarily “skinny” as D increases. For sufficiently large D , the labels on leaf nodes will be rotated to be oriented vertically and have area $a_l(n_D) = L(w(D))^2 = L/B^{2D}$, whereas the label on the root node will be oriented horizontally and have area $a_l(n_0) = L(h(0))^2 = L/(D+1)^2$. Hence,

$$\begin{aligned} \rho_l &= \log_B \left(\frac{a_l(n_0)}{a_l(n_D)} \right)^{1/D} = \log_B \left(\frac{B^{2D}}{(D+1)^2} \right)^{1/D} \\ &= \log_B B^2 - \log_B (D+1)^{2/D} = 2 - \Theta\left(\frac{2}{D} \log_B D\right) \end{aligned}$$

and $\rho_{l,\infty} = 2$.

Next, we must calculate A_l , which requires carefully considering the labels that can fit within the nodes. Let $i_1 = \log_B((D+1)/L)$, $i_2 = \log_B(D+1)$, and $i_3 = \log_B(L(D+1))$. If i_2 is an integer, nodes at depth $i = i_2$ are squares. At depth $i < i_2$, nodes are wider than high (i.e. $w(i) > h(i)$), and labels are oriented horizontally. At depth $i > i_2$, we have the opposite situation, and labels are rotated to be oriented vertically. And at depth $i < i_1$ or $i > i_3$, the aspect ratio of the nodes is greater than the aspect ratio L of the label. This leads to the following piece-wise expression for label area:

$$a_l(n_i) = \begin{cases} L(h(i))^2 = L/(D+1)^2 & \text{if } i \leq i_1 \\ (w(i))^2/L = 1/(LB^{2i}) & \text{if } i_1 \leq i \leq i_2 \\ (h(i))^2/L = 1/(L(D+1)^2) & \text{if } i_2 \leq i \leq i_3 \\ L(w(i))^2 = L/B^{2i} & \text{if } i_3 \leq i \end{cases}$$

Hence, the summation for calculating A_l can be split into four summations, one for each of the pieces defining $a_l(n_i)$, which we summarize here without the full details:

$$\begin{aligned} A_l &= \sum_{i=0}^D B^i a_l(n_i) \\ &= \sum_{0 \leq i \leq i_1} B^i \frac{L}{(D+1)^2} + \sum_{i_1 < i \leq i_2} B^i \frac{1}{LB^{2i}} \\ &\quad + \sum_{i_2 < i \leq i_3} B^i \frac{1}{L(D+1)^2} + \sum_{i_3 < i \leq D} B^i \frac{L}{B^{2i}} \\ &\approx \frac{(B+1)(2-1/L)}{(B-1)(D+1)} - \frac{L}{B-1} \left(\frac{1}{(D+1)^2} + \frac{1}{B^D} \right) \\ &= \Theta\left(\left(2 - \frac{1}{L}\right) \left(\frac{B+1}{B-1}\right) \frac{1}{D}\right) \end{aligned}$$

The step in the above calculation that involves an approximation is exact if i_1 , i_2 , and i_3 are integers and D is sufficiently large.

Finally, we clearly have $A_{l,\infty} = 0$.

4.2.4 Case 4C

Section 3.4 found the values of ρ , a_i , $a(n_D)$, and A for this case, from which the values of ρ_∞ and A_∞ follow immediately. If we assume that $L \geq 2$ (the more typical case), we find $a_l(n_i) = (w(i))^2/L = 1/(LB^{2i})$. It is easy to confirm, then, that

$$A_l = \sum_{i=0}^D B^i a_l(n_i) = \frac{B}{(B-1)L} - \frac{1}{(B-1)LB^D} = \frac{B}{(B-1)L} - \Theta\left(\frac{1}{B^D}\right)$$

hence $A_{l,\infty} = B/((B-1)L)$. The values for ρ_l , $\rho_{l,\infty}$, and $a_l(n_D)$ listed in the table are also easy to verify.

Actually, the above results (that are reported in Table 1) do not only hold for $L \geq 2$, but also hold more generally if the label's aspect ratio L is greater than or equal to the node's aspect ratio (i.e. $L \geq (B^{D+1} - 1)/((B-1)B^D) < 2$). In the opposite case, where $1 \leq L \leq (B^{D+1} - 1)/((B-1)B^D)$, we instead find that $a_l(n_D) = \Theta(L(B-1)^2/(B^2 B^{2D}))$ and $A_l = L(B-1)/B + \Theta(1/B^D)$. The results for this opposite case are omitted from Table 1 for brevity, but are reported in Table 2 where we need them to rank representations as a function of L .

4.2.5 Case 1D

All nodes in this case are squares of equal size. The leaf nodes have their centers on a circle with radius almost $1/2$ and circumference almost π (with this approximation becoming arbitrarily good as D increases). The side length of each node is well approximated by $\pi/((\sqrt{2})B^D)$, where the factor of $\sqrt{2}$ ensures that corners of adjacent leaf nodes need never overlap. It follows that $a(n_i) \approx \pi^2/(2B^{2D})$, from which it is easy to confirm the values of the metrics listed in Table 1 for this case.

4.2.6 Case 1E

In a circle of radius r , a sector (i.e. a "pie slice"-shaped region) with angle θ radians has area $\frac{\theta}{2\pi}\pi r^2 = \frac{1}{2}\theta r^2$. The area of each node in Figure 1E can be expressed as the difference in area between two overlapping sectors with the same angle but different radii. The angle corresponding to a node at depth i is $2\pi/B^i$, and the radii of the corresponding sectors are $\frac{i}{2(D+1)}$ and $\frac{i+1}{2(D+1)}$. So the area of the node is

$$a(n_i) = \frac{1}{2} \frac{2\pi}{B^i} \left(\left(\frac{i+1}{2(D+1)} \right)^2 - \left(\frac{i}{2(D+1)} \right)^2 \right) = \frac{\pi(2i+1)}{4(D+1)^2 B^i}$$

It is then straightforward to verify that $A_\infty = A = \sum_{i=0}^D B^i a_i = \pi/4$, as we should expect from the figure, and also that $a(n_D) = \Theta(\pi/(2DB^D))$. Next, we find

$$\begin{aligned} \rho &= \log_B \left(\frac{a(n_0)}{a(n_D)} \right)^{1/D} = \log_B \left(\frac{B^D}{2D+1} \right)^{1/D} \\ &= 1 - \log_B(2D+1)^{1/D} = 1 - \Theta\left(\frac{1}{D} \log_B D\right) \end{aligned}$$

Unlike the previous cases where ρ is a constant, here ρ is a function of D , with $\rho < 1$. Evaluating the limit, we find $\rho_\infty = 1$.

Next we analyze the size of labels. For the label on the root node, we make use of the fact that the area of a rectangle (such as a rectangular label) of aspect ratio $L \geq 1$ inscribed within a circle of radius r is $4Lr^2/(L^2+1)$. Since the radius of the root node is $1/(2(D+1))$, we deduce that the area of the root node's label is $a_l(n_0) = L/((L^2+1)(D+1)^2)$. As for leaf labels, we know that if D is large, the leaf nodes can be well approximated as skinny rectangles that are much longer than wide, with width $(D/(D+1))(\pi/B^D)$. Hence $a_l(n_D) \approx L\pi^2/B^{2D}$ for large D . Calculating ρ_l and $\rho_{l,\infty}$ is then straightforward.

To find A_l , we first note that A_l is bounded below by the area of the root label $a_l(n_0) = L/((L^2+1)(D+1)^2) = \Theta(1/D^2)$. Next, an upper bound on A_l can be found by approximating all nodes (except for the root) with rectangles containing appropriately oriented labels. We can then find an approximating upper bound on $a_l(n_i)$ defined in four pieces, closely analogous to that used in Section 4.2.3. Once again, we can split A_l into four summations, to find an upper bound on A_l that is $\Theta(1/D^2)$. Because the lower and upper bounds have the same Θ -complexity, this establishes the asymptotically tight bound of $A_l = \Theta(1/D^2)$. Furthermore, the detailed calculations involved in the approximating upper bound suggest that the hidden constant is π , hence A_l is given as $\Theta(\pi/D^2)$ in Table 1.

Note that, although the value of the hidden constant for A_l may only be approximate, its value only makes a difference when comparatively ranking representations with the same complexity. In Table 1, the only other representation whose A_l has the same complexity is that in Figure 8, whose A_l value was found with an approximation so similar to 1E's that any error in the hidden constants would not change their relative ranking.

4.2.7 Case 1F

In this representation, we assume that the circles corresponding to sibling nodes are laid out with their centers in a circle, so they are inside, and tangent to, the parent node's circle.

It can be shown that the ratio $\phi < 1$ of a node's circle's radius to its parent's circle's radius is

$$\phi = \frac{\sin(\pi/B)}{1 + \sin(\pi/B)}$$

(This is shown in [17]). Note that since $B \geq 2$, we have $0 < \phi \leq 0.5$. Since the radius of the root node's circle is 0.5 , the radius $r(i)$ of a node at depth i is $r(i) = 0.5\phi^i$, and its area is

$$\pi(r(i))^2 = \frac{1}{4}\pi\phi^{2i}$$

Recall that $a(n_i)$ is defined as the area of the node after subtracting away the overlapping area of descendants, which in this case means subtracting away the area of the B children whenever $i < D$:

$$a(n_i) = \begin{cases} \frac{1}{4}\pi\phi^{2i} & \text{if } i = D \\ \frac{1}{4}\pi\phi^{2i}(1 - B\phi^2) & \text{otherwise} \end{cases}$$

It is easy to check that $A = A_\infty = \pi/4$. The mean area exponent is

$$\begin{aligned} \rho &= \log_B \left(\frac{a(n_0)}{a(n_D)} \right)^{1/D} = \log_B \left(\frac{1 - B\phi^2}{\phi^{2D}} \right)^{1/D} \\ &= 2\log_B \frac{1}{\phi} - \frac{1}{D} \log_B \frac{1}{1 - B\phi^2} = 2\log_B \frac{1}{\phi} - \Theta\left(\frac{1}{D}\right) \end{aligned}$$

As with the previous case, ρ is a function of D , now with $\rho < 2$. We note that $\phi \approx \pi/B$ for large B , and

$$\rho_\infty = \lim_{D \rightarrow \infty} \rho = \log_B \frac{1}{\phi^2} = 2\log_B \left(1 + \frac{1}{\sin(\pi/B)} \right) \leq 2$$

The formula used in the previous subsection (1E), for the area of a rectangle inscribed within a circle, can again be applied to show that the area of leaf labels in this case is $a_l(n_D) = L/((L^2+1)(1/\phi)^{2D})$. For labels on non-leaf nodes, we simplify calculations by assuming that labels are inscribed within a circle that is itself inscribed within the child nodes (Figure 7). The circle inscribed within the children, i.e. the dashed red circle in Figure 7, has radius $r(i) - 2r(i+1) = \phi^i(0.5 - \phi)$, hence the label inscribed within it has area $a_l(n_i) = 4L\phi^{2i}(0.5 - \phi)^2/(L^2+1)$. This is not the optimal label size, but is an excellent approximation of the optimal size for large B . It is then a straightforward exercise to calculate ρ_l , $\rho_{l,\infty}$, A_l and $A_{l,\infty}$.

(The results for Case 1F in Table 1 assume $B \geq 4$, as indicated in the last row of the table. For $B = 2$ and $B = 3$, a different choice of label position makes the non-leaf labels significantly larger, but this would not change $a_l(n_D)$ nor the upper bound of $A_{l,\infty} \leq 1/2$ that we give in Table 1, and therefore would not change the rankings given in Table 1.)

4.2.8 Case 1G

The representation shown in Figure 1G is a slice-and-dice treemap [14]. The figure shows thin margins around each set of child nodes for illustration purposes. However, here we consider the extreme case where the margins have thickness zero and can be ignored. Because a_i is defined as the area of a node after subtracting away the area of

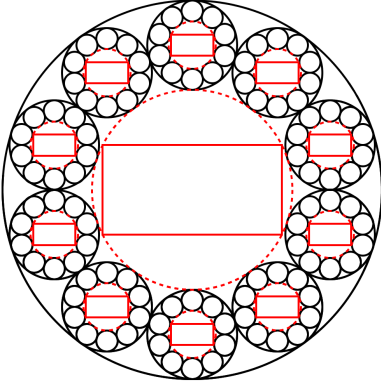


Fig. 7. The representation from Figure 1F with $B = 10$ and $D = 2$. Nodes are shown in black. Red rectangles, inscribed within the dashed red circles, show where we assume the labels are for the root node and for the root's children. The labels' aspect ratio $L = 2$ in this example. (For simplicity, labels on leaf nodes are not shown.)

descendant nodes, the leaf nodes are the only nodes with non-zero area, and the leaf nodes furthermore cover all the available area:

$$a_i = \begin{cases} 1/B^D & \text{if } i = D \\ 0 & \text{otherwise} \end{cases}$$

It follows that $A = A_\infty = 1$. Also, because $a_0 = 0$ and $a_D > 0$, we have $\rho = \rho_\infty = -\infty$.

As indicated by $\rho = -\infty$, the treemap allocates all area to its leaf nodes. We could of course modify the treemap to have margins, which would increase the value of ρ . We consider such a case later.

The analysis of label area is simple if we assume D is even (ensuring leaf nodes are square), yielding the results in Table 1. For D odd, the results are slightly different, and omitted for brevity.

4.2.9 Case 1H

This representation corresponds to the “indented outline” tree views used in many file browsers and other software. Nodes are rectangles of constant height, stacked in a single column. The top-to-bottom ordering of rectangles corresponds to a depth-first traversal of the tree. Each rectangle is indented to the right by a distance proportional to its depth.

The height of each node is $h(i) = (B - 1)/(B^{D+1} - 1)$. Let the indentation per level be $Kh(i)$, a constant multiple of the height, where $K > 0$. Then the width of the nodes is $w(i) = 1 - iKh(i)$, and

$$a_i = \frac{(B - 1)(1 - iKh(i))}{B^{D+1} - 1} = \frac{(B - 1)(B^{D+1} - 1 - K(B - 1)i)}{(B^{D+1} - 1)^2}$$

We omit the details of calculating A , and simply give the result:

$$A = 1 - \frac{KB(DB^D(B - 1) - B^D + 1)}{(B^{D+1} - 1)^2} = 1 - \Theta\left(\frac{K(B - 1)D}{B^D}\right)$$

hence $A_\infty = 1$. Also, it is straightforward to show that

$$\rho = -\frac{1}{D} \log_B \left(1 - \frac{KD(B - 1)}{B^{D+1} - 1}\right) \geq 0$$

Next, consider that for any quantity x that approaches zero, we have $\log_B(1 + x) = \Theta(x/\ln B) = \Theta(x)$. This fact can be used to simplify the above expression for ρ , yielding $\rho = \Theta(1/B^D)$. We also have $\rho_\infty = 0$.

Finally, it is straightforward to also calculate the metrics related to label area, reported in Table 1.

4.2.10 Case 8

One result of our analysis which we noticed early on is that the leaf nodes (and labels) in 1E are larger than those in 1C, because the leaf nodes are laid out along a curve of length π rather than a curve of length 1. This led us to invent the representation in Figure 8, where leaf nodes are laid out along (most of) the 4 edges of the bounding unit square, having a length of almost 4. This *concentric squares* representation leaves small corners of the unit square unused, but nevertheless outperforms the 1E representation in terms of all the area metrics.

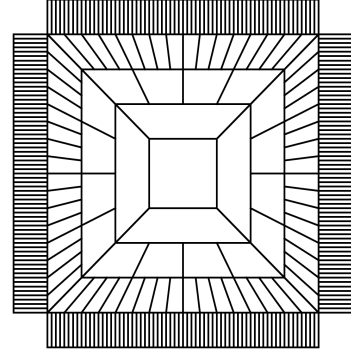


Fig. 8. A concentric squares representation, with $B = 4$ and $D = 4$. Notice that $A < 1$ because of the small square regions in the corners that are not used.

To simplify the analysis of this representation, we assume that B is a multiple of 4 (note that any B value could be used with concentric squares, albeit perhaps less efficiently). We also assume the concentric squares increase in size with a constant step, i.e. the side length of the squares is $(i + 1)/(D + 1)$ for $i = 0, \dots, D$, analogous to the concentric circles in 1E. It follows that the area of the root node is $a(n_0) = 1/(D + 1)^2$, the area of the four unused corners is equal to $a(n_0)$, and the area of each leaf node is $a(n_D) = 2D/((D + 1)^2 B^D)$. We omit the further details of the calculations, which are simple for most of the metrics. The calculation of A_l is non-trivial, but can be done in a very similar manner to how A_l was calculated in Sections 4.2.3 and 4.2.6, i.e. with a summation split into four parts. The non-leaf nodes are trapezoids in this case, but can again be approximated with rectangles, leading to an upper bound on A_l , as was done in Section 4.2.6, and a lower bound of the same complexity is found by simply considering the root node's label.

4.2.11 Case 9

Because treemaps (e.g. representation 1G) seem promising, we analyze a treemap with margins and with labels on all its nodes. We consider the kind of treemap shown in Figure 9, where labels alternate in orientation from one level to the next, to save space and simplify analysis. Nodes at even depths $0, 2, 4, \dots$ are squares (hence $w(i) = h(i)$ for even i) and have horizontally oriented labels. Nodes at odd depths are rectangles and have vertically oriented labels. All labels have aspect ratio L . Given i even, the labels at depths i and $i + 1$ are the same size, and the size of labels is chosen so that each label at depth $i + 1$ matches the height of the nodes at depth $i + 2$. The sizes and relative positions of labels and margins within nodes at depths i and $i + 1$ are simply scaled down for each of the nodes at depths $i + 2$ and $i + 3$, respectively. The only exception to this is at $i = D$, where labels are at maximum size within leaf nodes.

Let $m \geq 0$ be the small fraction of the height (or width) of each node at even depth that is devoted to margins, with m independent of i or L . Next, consider the labels at depth 0 and 1, and how their dimensions along the vertical axis must add up to the total length $1 - m$ not used for margins. There is 1 label at depth 0, with vertical size $h_l(0)$, and B labels at depth 1, each with vertical size $w_l(0)$ (because they have the same dimensions as the root label, only rotated). Hence, $h_l(0) + Bw_l(0) = 1 - m$. Since $w_l(0) = Lh_l(0)$, it follows that $h_l(0) =$

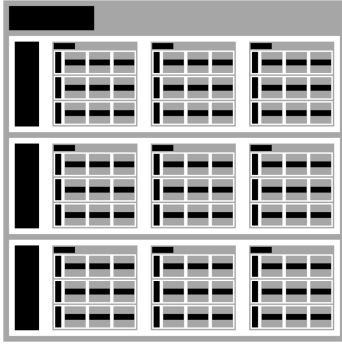


Fig. 9. A treemap with margins and labels, with branching factor $B = 3$ and depth $D = 4$. Rather than show actual labels, the bounding rectangles for the labels are shown in black, and have an aspect ratio of $L = 3.5$ in this example. The root node's label is at upper left; there are also $3^4 = 81$ leaf nodes, each with their own label.

$(1 - m)/(1 + BL)$. More generally, for i even, it can be shown that

$$\frac{h_i(i)}{h(i)} = \frac{1 - m}{BL + 1} \quad ; \quad \frac{h(i+2)}{h(i)} = \frac{w(i+2)}{w(i)} = \frac{1 - m}{B + 1/L}$$

The area of each square corresponding to a node at even depth i is

$$w(i)h(i) = (w(i))^2 = \left(\frac{1 - m}{B + 1/L} \right)^i$$

(Note that the above expression is not equal to a_i , because a_i is defined as the area after subtracting away the area of descendants.) Assuming D even, it is straightforward to confirm the values given in Table 1 for ρ , ρ_∞ , and $a(n_D)$, and clearly $A = A_\infty = 1$.

The area of a label is

$$a_l(n_i) = \begin{cases} a_l(n_{i-1}) & \text{if } i \text{ odd or } i = D \\ \frac{1}{L} \left(\frac{1 - m}{B + 1/L} \right)^{i+2} & \text{otherwise} \end{cases}$$

Assuming D even, the values of ρ_l , $\rho_{l,\infty}$, and $a_l(n_D)$ given in Table 1 follow. Though we omit the demonstration, for D even it can also be shown that

$$A_l = \frac{(1 - m)^2(B + 1)L}{(BL + 1)^2 - (BL(1 - m))^2} + k^D \left(\frac{1}{L} - \frac{(1 - m)^2(B + 1)L}{(BL + 1)^2 - (BL(1 - m))^2} \right)$$

where $k = (1 - m)/(1 + 1/(BL)) < 1$. (Note that A_l is close to 1 for m small, B large, D small, and $L = 1$.) Evaluating the limit, we find

$$A_{l,\infty} = \lim_{D \rightarrow \infty} A_l = \frac{(1 - m)^2(B + 1)L}{(BL + 1)^2 - (BL(1 - m))^2}$$

Keeping in mind that $B \geq 2$, $L \geq 1$, and $m \geq 0$, it can be shown that the above is bounded by $A_{l,\infty} < 3/4$. (Note that $A_{l,\infty}$ is close to $3/4$ for m small, B small, and L large). Thus, the treemap in Figure 9 fills all available area ($A = 1$), and has excellent leaf label area with $a_l(n_D) = \frac{1}{L} \left(\frac{1 - m}{B + 1/L} \right)^D$; but in the limit, as $D \rightarrow \infty$, its labels use less than 75% of the available area ($A_{l,\infty} < 3/4$). In contrast, the treemap in Figure 1G also has $A = 1$, has better leaf label area $a_l(n_D) = \Theta(\frac{1}{LB^D})$, and its $A_{l,\infty}$ can be as high as 1 (depending on L), but it has the significant disadvantage of having no labels on non-leaf nodes.

As a side note, there is an interesting variation on 1G that allows all nodes to be labeled. The idea is to temporarily add an extra child to each non-leaf node, yielding a $(B + 1)$ -ary tree, and then generate a layout using the representation in Figure 1G, and then use the space allocated for the temporary children to display labels of their parent nodes. This allows all of the original nodes to be labeled, and yields a

leaf label area $a_l(n_D) = \frac{1}{L} a(n_D) = \frac{1}{L(B+1)^D}$ (assuming D even). Further analysis also shows that $A_{l,\infty} \leq 11/15$ (with $A_{l,\infty} = 11/15$ when $B = 2$ and $L = 3$). The bound on $A_{l,\infty}$ here is not as high as in Figure 9, so we did not include this case in Table 1, but it is interesting that $11/15$ is so close to $3/4$.

4.2.12 Case 10

Is it possible for a representation to allow labels on all nodes ($a_l(n_i) > 0$ for all i), and not only occupy all available area ($A = 1$), but also use all available area by filling it with labels ($A_l = A = 1$)? Figure 10 shows an example of a tree representation with these properties for the special case $B = 8$ and $L = 1$. Consider first the case where $B = 8$ and $L \geq 1$ is variable. All nodes in the representation are squares, the total area $A = A_\infty = 1$, total label area $A_l = A_{l,\infty} = 1/L$, and the leaf label area $a_l(n_D) = \frac{1}{L} a(n_D) = \frac{1}{L(B+1)^D}$ is comparable to that in Figure 9. It can also be shown that $\rho = \rho_l = \frac{D-1}{D} \log_8 9$ and $\rho_\infty = \rho_{l,\infty} = \log_8 9 \approx 1.057$. If $L = 1$, then $A_l = A = 1$, as we sought.

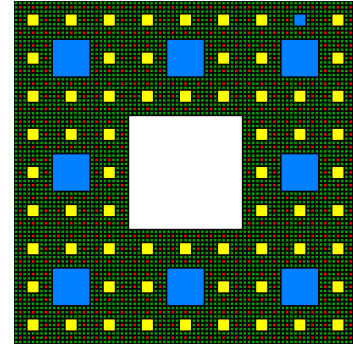


Fig. 10. A tree representation with $B = 8$ and $D = 4$ based on the Sierpiński carpet fractal. Each non-leaf node is surrounded by its 8 children. The root node is white, surrounded by its 8 children in blue, and the leaf nodes are green, surrounding their parents in red. Notice that the bottom two levels consist of nodes of the same size.

Note that the representation in Figure 10 can be extended to cases where $B = (a^2 - 1)b^2$, for integers $a \geq 2$ and $b \geq 1$. To see how, imagine a “generating pattern” for the representation in Figure 10 based on a grid of 3×3 cells, where the central cell is the root node and each of the other 8 cells are recursively subdivided. Now, imagine instead a grid of $a \times a$ cells, where one cell is the root node and each of the other $(a^2 - 1)$ cells are further divided into a $b \times b$ smaller cells, each of which is then recursively subdivided. Similar tricks can be used to extend the representation to other values of B if rectangular nodes are used (appropriate for labels with $L > 1$) and/or a small amount of area is left unused after tiling most of the grid with nodes.

5 DISCUSSION

5.1 General observations

Perhaps the most important thing that Table 1 allows us to do is compare different representations. When Stasko and Zhang [27] described Sunburst (essentially 1E) as “space-filling”, with the qualification that it does “not completely occupy the display space as does the Treemap”, they were presumably referring to the fact that $A = \pi/4 < 1$ for 1E. Our analysis allows us to compare 1E with many other representations (including treemaps), not just in terms of total area A , but according to all the metrics we have considered. Doing so reveals some subtle issues that go beyond the fact that $A < 1$ for 1E. For example, we could compare 1E with 1C which, like treemaps, have optimal total area $A = 1$. (1E and 1C are also interesting to compare because the former is a polar coordinate version of the latter.) We find that 1E has leaf nodes that are (asymptotically) $\approx \pi/2$ times larger, and leaf labels $\approx \pi^2$ times larger, than those in 1C, despite the fact that both representations must fit within the same 1×1 square, and 1E doesn’t

even take up all the available space! So, there are clearly some important advantages in 1E that might make it preferable over 1C, despite the fact that 1E “does not completely occupy the display space” like 1C does.

It is worth noting how misleading metrics that ignore label size can be. For example, cases 1C, 1H, and 8 all have optimal asymptotic total area $A_\infty = 1$, however their asymptotic total *label* area $A_{l,\infty} = 0$ is as bad as can be. Case 1H is a particularly dramatic example: in Table 1, it is ranked 2nd in terms of both leaf area $a(n_D)$ and total area A , and the complexity of its leaf area $a(n_D) = \Theta(1/B^D)$ is as good as can be (ignoring leading coefficients), yet this representation is ranked 9th in terms of total *label* area A_l . More generally, any representations that have leaf nodes that become arbitrarily “skinny” with depth D (i.e. 1B, 1C, 1E, 1H, 8) will have leaf area $a(n_D)$ values much larger than their leaf *label* area $a_l(n_D)$, because the mismatch between the aspect ratios of the leaf nodes and labels grows arbitrarily bad with depth D , hence most of the area within each leaf node will not be used by the label. This clearly demonstrates the importance of taking into account label size to evaluate the space-efficiency of tree representations.

One pattern that can be seen in the leaf label area values is that representations with arbitrarily skinny leaf nodes (1B, 1C, 1E, 1H, 8) have $a_l(n_D)$ proportional to L , because they correspond to the situation in Figure 6, top left and bottom left. On the other hand, representations whose leaf nodes have a roughly fixed aspect ratio correspond to the situation in Figure 6, top right and bottom right, and have $a_l(n_D)$ *inversely* proportional to L — this is true, or approximately true, of all the other representations analyzed.

It is also interesting that, within the set of representations with arbitrarily skinny leaf nodes (1B, 1C, 1E, 1H, 8), the leaf nodes are laid out, or tiled, along a 1-dimensional curve. In cases 1B, 1C, and 1H, this curve is an edge of the 1×1 bounding square, and has length 1. In case 1E, the curve is the circumference of an inscribed circle, and has length π . In 8, the curve is (most of) the 4 edges of the 1×1 bounding square, with length almost 4. In all these cases, the labels on the leaf nodes are oriented perpendicular to the direction of the tiling, and as $D \rightarrow \infty$ there is plenty of room for the labels in this perpendicular direction. The size of the labels is only constrained along the direction of tiling, and hence the *length* of the curve is the limiting factor for the size of leaf labels. The complexity of $a_l(n_D) = \Theta(L/B^{2D})$ is the same (ignoring leading coefficients) for all these cases, but the leading coefficients reveal that the representations with longer curves have larger $a_l(n_D)$ values. This explains why 1E has a better $a_l(n_D)$ value than 1B or 1C, and also why the representation in Figure 8 is the best of the set. In fact, our invention and analysis of representation 8 was motivated by a desire to tile the leaf nodes along a longer curve than in 1E, as well as a desire to fill more of the corners of the available area than does 1E. Finally, it is also clear why 1H has the worst $a_l(n_D)$ value of the set of representations: it tiles *all* nodes along the same curve of length 1 rather than just tiling leaf nodes.

Certain other representations (i.e. 1A, 4C, 1D) also tile their leaf nodes along a curve, but are worse because their leaf nodes have a fixed aspect ratio. Hence, labels are constrained both along the direction of tiling and perpendicular to it, and $a_l(n_D) = \Theta(1/(LB^{2D}))$ is inversely proportional to L for these representations.

In addition to tiling nodes along a longer curve (Figure 8), another approach we experimented with to create more space for leaf nodes is shown in Figure 11, right. This is a variation on 1E where the last level of nodes is “rotated” to lie along radial lines instead of along a circumference; we call this an *asymmetrical* representation. (Such a rotation of nodes is comparable to the “folding” in section 4.1 of [8].) To avoid confusion, in an incomplete tree, such rotation should only be applied to the deepest level of leaf nodes, and not to leaf nodes on more shallow levels.

5.2 Comments on 1H

The indented outline representation (1H) is interesting for a few reasons. Prior to our analysis, we expected this representation to perform rather well, because it allows nodes to be tiled vertically while allowing long labels to extend horizontally without introducing any extra

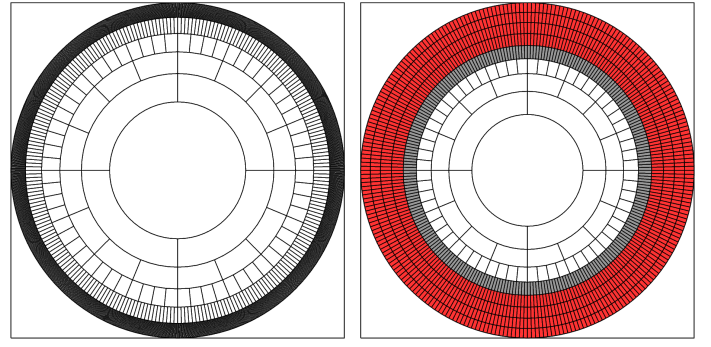


Fig. 11. Two different representations with $B = 4$ and $D = 5$. At left is a normal case, with radii of circles chosen so $\rho = 1$. As can be seen, the leaf nodes, appearing in the outermost ring, are crowded beyond legibility, making the ring appear almost black. At right, nodes on levels 0 through 4 are shown as usual with the central circle and the inner 4 rings, however the leaf nodes (in red) are then distributed over the subsequent 4 outer rings. Each set of B leaf nodes with a common parent (in gray) is effectively “rotated” and laid along a radial line from its parent. This allows leaf nodes to retain a useful size, unlike the case on the left.

space between nodes. Now, in a practical setting, the nodes in such a representation are typically not nearly as wide as the representation is tall; for example, the tree might contain 200 hundred nodes, each of which has a label that is a mere 25 characters long, so the aspect ratio of a rectangle bounding the representation might be $\approx 200/25$. For the sake of our analysis, however, to enable fair comparison with other representations, we imposed the same 1×1 bounding square as with all the representations, and therefore we made the nodes in 1H extend horizontally as far as possible to fill the available space. The results partially confirmed our expectations: $a_l(n_D)$ is proportional to L (indicating that labels have room to “grow” without any penalty) and, as $D \rightarrow \infty$, we have $A_\infty = 1$, i.e. the space left unused by indentation becomes negligible and the nodes fill all the available space. (Interestingly, 1H is the only representation we have found for which $A_\infty = 1$ while all the nodes have asymptotically the same size, i.e. $\rho_\infty = 0$). However, somewhat to our surprise, the total space used by labels goes to zero (i.e. $A_{l,\infty} = 0$) and 1H ranks relatively poorly in terms of leaf label area and total label area. As explained in the previous section, this is because the mismatch in aspect ratios of labels and nodes grows without bound as $D \rightarrow \infty$.

Now, what if we had instead modeled this representation as it is typically used in practice, i.e. with nodes that are just wide enough to contain their labels instead of extending horizontally as far as possible? The bounding rectangle around the representation would then have been tall and narrow instead of a square, and in the limit (as $D \rightarrow \infty$), all of the bounding rectangle’s area would be used by the labels. But, in the limit, the aspect ratio of the bounding rectangle would also go to infinity, so if a user zoomed out to fit the entire representation within their (square) screen, most of the available screen space would be wasted. So, the end result of the analysis would be the same, but the analysis would be complicated by having to consider the aspect ratio of the bounding rectangle of the representation. This illustrates why it is simpler and easier to bound all representations within a 1×1 bounding square from the start of analysis: if a tree representation performs poorly within such a square (i.e. $A_{l,\infty}$ is small), there’s no point in relaxing the constraint of a bounding *square* and re-analyzing, since the poor performance may simply be re-expressed as a poor aspect ratio of the bounding rectangle, resulting in yet another metric to consider.

Despite the relatively poor space-efficiency we found for 1H, this indented outline representation is popular in user interfaces such as file browsers. This popularity may partly be due to the fact that when users are using an indented outline, they are typically zoomed in on some subregion of the tree rather than looking at the entire tree, and

during scrolling they mostly only need to scroll vertically. It may also have something to do with how users can expand or collapse subtrees, and work their way down a path, by clicking on nodes and performing small, similar, repeated mouse motions, e.g. motions down and right. Such details are not captured in our analysis (though the incorporation of such interactive aspects would be an interesting direction for future research). However, in examining the results of our analysis, we found a way to improve the space-efficiency of 1H (in terms of $a_l(n_D)$), while remaining true to the basic idea of an indented layout. The key was to combine 1B with 1H in a way that is exactly as space-efficient as 1B (Figure 12). Such a hybrid might prove to be useful in interactive software, combining advantages of 1H while being more space-efficient.

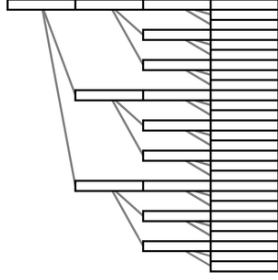


Fig. 12. A hybrid of the representations in Figures 1B and 1H. Each non-leaf node in 1B is shifted upwards to be aligned with its top-most child. The result imitates the indentation in 1H, but requires no change in node (or label) sizes, hence the labels have the same, larger size they have in 1B.

5.3 Optimal Space-Efficiency?

Examining the rankings by label-oriented metrics in Table 1, we see that the most space-efficient representations are 1G, 9, and 10. The only representations that allow $A_{l,\infty}$ to achieve a value of 1 are 1G, 10, and 4C (the last of which has one of the worst $a_l(n_D)$ values of all). 1G is unique in having the best possible $a(n_D)$, but has the important disadvantage of not allowing labels on non-leaf nodes. 10 has itself the severe disadvantage of only being efficient for certain values of B , such as $B = 8$.

The remaining representation, 9, is perhaps the most space-efficient “usable” representation. It also has a disadvantage, though arguably less severe: our analysis showed that $A_{l,\infty} < 3/4$ for this representation. Thus, in the limit, it is impossible to fully use all the available space for labels. The value of the bound $3/4$ is perhaps an artifact of the size and placement we chose for nodes and labels in our treemap with margins. However, the fact that there is a bound less than 1 probably reflects a deeper limitation of treemaps with margins, and achieving $A_{l,\infty} = 1$ may be impossible without “filling up” the space to the right of labels with nodes. It is an intriguing challenge to try to find an alternative way to layout a treemap that would improve on our bound of $3/4$, as this could have both theoretical and practical implications.

A more ambitious goal would be to find a tree representation for which $A_{l,\infty} = 1$ and $a_l(n_D) = \Theta(1/B^D)$, that (unlike 1G) allows labels on all nodes. Ideally, such a representation would also be equally efficient for all values of B and L .

One objection that might be raised against seeking optimal space-efficiency is the following: even if some representation can be found that has optimal $A_{l,\infty}$ and $a_l(n_D)$, there will likely be a perceptual benefit in introducing some whitespace or margins between elements of the representation, to help users better perceive appropriate groupings of elements. This is a valid point, however, it would be preferable for the amount of whitespace to be under the control of the user (or the designer) and to be something that can be added to the representation (e.g. in the form of margins), rather than there being a minimal amount of whitespace imposed by the representation. Hence we still see value in seeking representations that are optimal, i.e. that use all available space.

5.4 Observations concerning ρ and ρ_l

Regarding the mean area exponent ρ , we point out that representations can often be adjusted to yield different ρ values. For example, the horizontal lines in 1C/4B are shifted downward to produce 4C, changing the heights of nodes and the ρ value from 1 to 2. In an analogous manner, Figure 13 shows how 1E can be adjusted to yield different ρ values. Thus, one could imagine exposing ρ as a controllable parameter to a designer or user, to “tune” a tree representation.

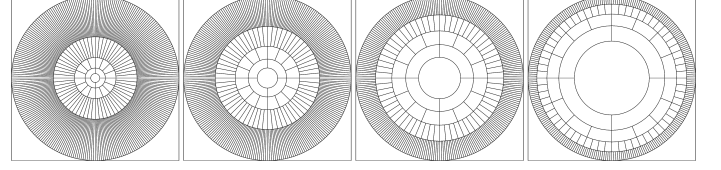


Fig. 13. Changing the radii of circles in Figure 1E allows us to create a representation with any ρ value. Here, $B = 4$, $D = 4$, and from left to right, ρ is 0, 0.33, 0.66, and 1, respectively, gradually allocating more area to shallow nodes.

In terms of label-oriented metrics, the most space-efficient representations in Table 1 that allow labels on all nodes (namely, 9, 10, and 1F, but not 1G since it only allows labels on leaf nodes) have $\rho_{l,\infty}$ values between 1 and 2. Representations with $\rho_{l,\infty} < 1$ or $\rho_{l,\infty} \geq 2$ are ranked as less efficient.

Why would $1 \leq \rho_{l,\infty} < 2$ correspond to more efficient representations? Well, as $D \rightarrow \infty$, the number of leaf nodes increases exponentially, so their size (and label legibility) decreases exponentially. If $\rho_{l,\infty}$ is low ($\rho_{l,\infty} = 0$, for example), then this exponential decrease in size will also happen to labels on shallow nodes, making the representation less efficient in terms of $A_{l,\infty}$. In contrast, if $\rho_{l,\infty} \approx 1$, shallow labels are more likely to remain legible as D increases, which would make a zoomed-out overview of the tree more useful, with shallow labels guiding the user during zoom-in operations. Finally, if $\rho_{l,\infty}$ is high ($\rho_{l,\infty} \geq 2$), then the exponential decrease in size of deep labels will be faster, making the representation less efficient in terms of $a_l(n_D)$. Representations 9 and 10, arguably the 2 best representations overall, have $\rho_{l,\infty}$ only slightly larger than 1. This is not surprising, because the best leaf label size $a_l(n_D)$ we could hope for would be something slightly less than $1/B^D$, for which we expect $\rho_{l,\infty}$ to be close to 1 (assuming the root label’s area $a_l(n_0)$ does not go to zero exponentially fast as $D \rightarrow \infty$). So, $\rho_{l,\infty}$ values close to 1 seem to characterize the most space-efficient representations in terms of label size.

We also note that neither $a_l(n_D)$ nor A_l alone is always a good indicator of space-efficiency. For example, 1G has optimal $a_l(n_D)$, but does not allow labels on any non-leaf nodes, and 4C has the best $A_{l,\infty}$ value among the representations analyzed, but has a poor $a_l(n_D)$ value. In contrast, the condition $1 \leq \rho_{l,\infty} < 2$ consistently corresponds only to the most space-efficient forms we have analyzed. So, if the space-efficiency of a representation is to be assessed with just a *single* metric, it may be feasible to use $\rho_{l,\infty}$ (or ρ_l , whose value as a function implies the value of $\rho_{l,\infty}$) as the metric of choice, rather than $a_l(n_D)$ or A_l . However, if *two* metrics can be used, we recommend using $a_l(n_D)$ and A_l , with the former taking precedence over the latter, because together they provide more information than ρ_l alone. This is the approach taken in Table 2 which is described next.

6 DESIGN IMPLICATIONS AND GUIDELINES

6.1 Further comparisons of representations

In Table 2, we have attempted to distill the most important information in Table 1. We eliminated representations 10 (which we judge to be impractical) and 1G (which, while used in practice in software such as SequoiaView, has very limited support for labels). Representation 4C, which we initially analyzed to illustrate a high- ρ case, turns out to have an excellent total label area A_l , and could be useful in a zoomable user interface (ZUI) [3], and so is not eliminated in Table 2. We also report only the most important label-oriented metrics in Table 2, and give 2

rankings for different values of L . In both rankings, representations are ranked by $a_l(n_D)$, using A_l only to break ties. In the first ranking, L is assumed to be significantly larger than 1, for example, $L > 5$. In the second ranking, the labels are squares (such as image thumbnails) with $L = 1$. There's nothing particularly special about the value 5 in the $L > 5$ case; 5 is simply chosen to make the ranking for that case more clear, since the relative performance of certain representations can be more subtle for values of L between 1 and 5.

Interestingly, in the $L > 5$ ranking, we find that the best representations are all those involving nested enclosure of nodes, followed by all those with "skinny" leaf nodes (i.e. leaf nodes with unbounded aspect ratio), followed by all those with "suarish" leaf nodes (i.e. leaf nodes with a fixed aspect ratio). These three categories are indicated in the column headings in Table 2, and representations are arranged left-to-right, from most to least space-efficient, according to the $L > 5$ ranking. In the $L = 1$ ranking, again the best representations are those involving nested enclosure, but after that point the ranking alternates somewhat between the other two categories of representations.

The last two rows in Table 2 contain some additional information of practical value. The second last row concerns the orientation of labels, which can impact readability. The last row indicates which representations lend themselves to a weighted partitioning. As already explained, treemaps are commonly used to create a weighted partitioning of the total area available, however other representations also allow this, such as 1C (as shown in Figure 2), or 1E and 8 (by changing the angles of wedges according to weight). In other representations, such as 1D or 1A, it would be possible to vary the size of nodes according to their weight, but this might reduce the total area A of the representation, so they do not lend themselves as well to weighted partitioning.

As a side note, the second last row of Table 2 indicates that representation 9 allows all labels on leaf nodes to be oriented the same way. This is not necessarily true if the tree is incomplete (i.e. if leaf nodes are not all at the same depth). However, it is easy to modify the treemap in Figure 9 so that *all* labels are oriented horizontally, even if this comes at the cost of a reduction in space-efficiency. Because of this, the information in the second last row is a useful approximation of the truth for treemaps with margins and labels.

6.2 Design Guidelines

We now enumerate some guidelines that may be useful to designers. For brevity, we use the terms *skinny label* to refer to a label whose aspect ratio L is significantly greater than 1, *suarish label* to refer to a label where $L \approx 1$, *skinny leaf* to refer to a leaf node whose aspect ratio grows arbitrarily large as $D \rightarrow \infty$, and *suarish leaf* to refer to a leaf node whose aspect ratio remains fixed or converges to a constant value as $D \rightarrow \infty$.

1. Choose a representation whose leaf node aspect ratio matches the aspect ratio of labels to be used. If the labels to be placed on nodes are skinny, the designer could look in Tables 1 or 2 for representations whose leaf nodes are skinny (indicated in the tables by $a_l(n_D)$ proportional to L), and consider using those first. On the other hand, if the labels are suarish, the designer could consider first using representations with suarish leaf nodes (indicated by $a_l(n_D)$ inversely proportional to L).

2. Modify labels to match the aspect ratio of their nodes. If a designer has chosen a representation whose leaf nodes are suarish, then labels consisting of long strings of text should be wrapped into suarish multi-line blocks of text. On the other hand, if the representation has skinny leaf nodes, then labels consisting of long strings of text should of course *not* be wrapped into multiple lines.

3. Modify nodes to match the aspect ratio of their labels. In the case of representations with skinny leaf nodes, even with skinny labels, eventually most of the available space will not be used by the labels, because the leaf nodes' aspect ratio grows without bound. One possible way to create a better match between aspect ratios is shown in Figure 14. The horizontal lines separating levels of nodes are shifted downward as far as possible without making deep labels any smaller. This allows shallow labels to be made larger, increasing A_l . Analogously, the circles and squares separating levels within Figures 1E

Figure	Nested enclosure of nodes			Leaf nodes have an unbounded aspect ratio				Leaf nodes have a fixed aspect ratio			
	9 treemap with margins and labels	1(f) nested circles	8 concentric squares	1(e) concentric circles	1(c)/4(b) icicle	1(b) classical variant	1(h) indented outline	1(d) radial	4(c) icicle variant	1(a)/4(a) classical (layered)	
leaf label area $a_l(n_D)$	$\frac{1}{L} \left(\frac{1-m}{B+1/L} \right)^D$	$\frac{L}{(L+1)(1/\phi)^{2D}}$	$\Theta\left(\frac{16L}{B^2}\right)$	$\Theta\left(\frac{\pi^2 L}{B^2}\right)$	$\frac{L}{B^2}$	$\frac{L}{B^2}$	$\Theta\left(\frac{L(B-1)^2}{B^2 B^2}\right)$	$\Theta\left(\frac{\pi^2}{2L(B-1)B^2}\right)$	$\begin{cases} 1/(LB^2), L \geq 2 \\ \Theta\left(\frac{(B-1)^2}{B^2 B^2}\right), L = 1 \end{cases}$	$\frac{1}{LB^2}$	
total label area A_l	$A_{l,\infty} + \Theta(L^D)$	$A_{l,\infty} + \Theta((B\phi^2)^D)$	$\Theta\left(\frac{4}{B^2}\right)$	$\Theta\left(\frac{\pi^2}{B^2}\right)$	$\Theta\left(2 - \frac{1}{L}\right)\left(\frac{B+1}{B-1}\right)\frac{1}{B}$	$\Theta\left(\frac{BL}{(B-1)B^2}\right)$	$\Theta\left(\frac{L(B-1)}{B^2 B^2}\right)$	$\Theta\left(\frac{\pi^2 B}{2L(B-1)B^2}\right)$	$A_{l,\infty} \pm \Theta\left(\frac{1}{B^2}\right)$	$\Theta\left(\frac{B}{L(B-1)B^2}\right)$	
total label area $A_{l,\infty}$ as $D \rightarrow \infty$	$\frac{(1-m)^2(\theta+1)L}{(BL+1)^2 - (BL(1-m))^2} < \frac{3}{4}$	$\frac{L}{(L+1)(1-\phi^2)^2} < \frac{1}{2}$	0	0	0	0	0	0	$\begin{cases} 1/(LB^2), L \geq 2 \\ \Theta\left(\frac{(B-1)^2}{B^2 B^2}\right), L = 1 \end{cases}$	0	
Rank if $L > 5$ ("skinny" labels)	1	2	3	4	5	6	7	8	9	10	
Rank if $L = 1$ ("suarish" labels)	1	2	3	4	6	7	9	5	8	7	
Allows all leaf labels to be oriented the same way, for easier reading	YES (see text for qualifications)	YES	no	no	YES	YES	YES	YES	YES	YES	
Allows weighted partitioning without reducing total area A	YES	YES	YES	YES	YES	no	no	no	YES	no	

Table 2. A comparison of representations relevant for practical applications. Rankings are by $a_l(n_D)$ (1 being best), using A_l to break ties.

and 8 can be shifted away from the root node to increase A_l without reducing $a_l(n_D)$.

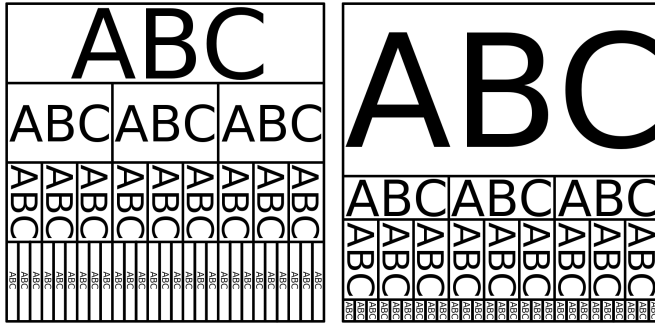


Fig. 14. At left, an icicle diagram like Figure 1C, with labels added. At right, the icicle diagram has been modified to allow shallow labels to be larger, without reducing the size of deep labels.

Another example of modifying nodes to match the aspect ratio of labels is presented in Section 7, where nodes in a treemap are “rectified” rather than squarified, to create rectangles whose aspect ratio is closer to the aspect ratio of their labels.

4. In classical representations, orient labels to be perpendicular to the levels. It is interesting that 1A, the most “classical” of representations, is ranked last in space-efficiency in Table 2 when $L > 5$. Representation 1B is significantly superior, because it uses up more of the space available between the levels of nodes. Rotating skinny labels so their length is perpendicular to the levels of the tree imitates 1B and is more space-efficient. Although there is a strict 1×1 bounding rectangle in Figures 1A and 1B, this guideline also holds when the representation is allowed to have a bounding rectangle that is not 1×1 : in such a case, applying the guideline simply improves the aspect ratio of the bounding rectangle of the representation. Figure 15 illustrates.

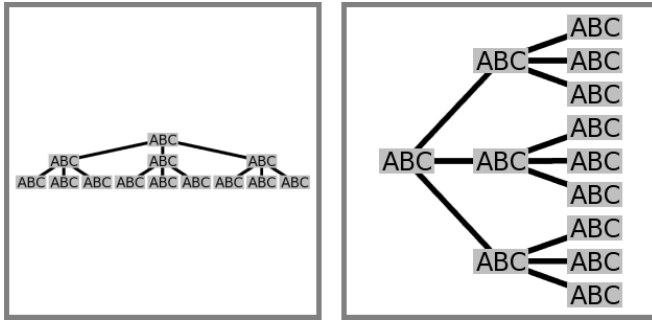


Fig. 15. At right, labels have been rotated (as has the entire representation), allowing labels to be shown at a larger size within the same square viewport.

5. In radial representations, orient labels to be along radial lines. For reasons analogous to those in the previous guideline, if designers wish to use a radial representation, they should consider orienting labels so their length is along radial lines, in imitation of 1E rather than 1D. This will allow labels to be larger within the allowed space. (Note, however, that this can also negatively impact readability, as the user may need to rotate the view to read certain labels.)

6. In representations that do not use nested enclosure, modify the representation to lay out skinny leaf nodes along a curve that is as long as possible. As already discussed, laying out skinny leaf nodes along a longer curve will allow the leaf labels to be larger. This motivated our concentric squares representation (Figure 8). Designers may find other ways to tile skinny nodes along longer curves, perhaps by taking inspiration from space-filling curves.

7. Within the given constraints, choose a representation that maximizes space-efficiency. Designers typically work within several con-

straints beyond just space-efficiency, which might make them choose any of the representations considered. However, we can still identify simple criteria and rules of thumb for choosing a representation that increases space-efficiency.

First, the most basic criterion is that there be labels on all nodes that decrease in size with depth, i.e. $a_l(n_i) > a_l(n_{i+1}) > 0$ or $\rho_l > 0$. Second, the representation should have a leaf label size $a_l(n_D)$ that shrinks as slowly as possible (ideally, $a_l(n_D) = \Theta(1/B^D)$ or $\rho_l \approx 1$). Third, the representation should have a total label area A_l as high as possible (ideally, $A_l = 1$). The representations in Table 2 all satisfy the first criterion, and are ranked according to the degree to which they satisfy the 2nd and 3rd.

Other criteria may also be considered, producing a decision procedure for choosing a representation. For example, a representation that uses nested enclosure may be deemed too confusing in certain contexts, or the designer may want to improve legibility of labels on leaf nodes by having them all oriented the same way — such criteria eliminate subsets of representations from further consideration. The following is an example decision procedure based on the information in Table 2:

```

if okay to use nested enclosure then
  use treemap (Figure 9)
else if leaf labels needn't be all oriented the same way then
  use concentric squares (Figure 8)
else if  $L$  significantly larger than 1 then
  use icicle (Figure 1C)
else if need support for weighted partition then
  use icicle (Figure 1C)
else
  use radial (Figure 1D)

```

The first step in the above procedure is to ask if the use of nested enclosure is acceptable. If it is, no further questions need be asked: a treemap should be used, because it is the most space-efficient, and because it also allows labels to all be oriented the same way, and allows for weighted partitioning of area. If not, then further questions are required to narrow down the choices.

Designers may combine the above criteria with their own criteria that may have little to do with label size, and create their own rules of thumb or decision procedure. Such a decision procedure could also be extended to include novel tree representations not considered here, by performing the same kind of analysis of space-efficiency that we have done in this article.

6.3 Using the metrics to evaluate other representations

Given a tree representation not considered in our analysis, such as one being considered by a designer, or a novel representation being proposed by a researcher, which metrics should be used to evaluate its space-efficiency, and how?

Section 3.7 listed six kinds of metrics, and Sections 3.7 and 5.1 showed how misleading metrics can be that do not take into account label size. We recommend evaluating, most importantly, a metric of **smallest label efficiency** (for example, $a_l(n_D)$), or the average font height in leaf labels), and preferably also a metric of **total label efficiency** (i.e. the total area A_l devoted to labels).

These two kinds of metrics can be applied to any given instance of a tree representation, such as those in Figures 2, and so could be used to evaluate and compare different static diagrams of the same tree.

However, if one wishes to evaluate a given representation style (i.e. drawing convention), then these metrics should ideally be evaluated analytically, as is done in Section 4. If an analytical approach is too difficult or requires introducing unrealistic assumptions, then a numerical evaluation may be performed over many randomly generated trees to find average values of the metrics. To evaluate representation styles, we also recommend finding the asymptotic behavior of these metrics for large values of depth D , ideally evaluating $A_{l,\infty}$, as well as the mean area exponent ρ_l and its limit $\rho_{l,\infty}$. As discussed in Section 5.4, having $1 \leq \rho_{l,\infty} < 2$, with $\rho_{l,\infty}$ close to 1, is a condition associated with highly efficient representations.

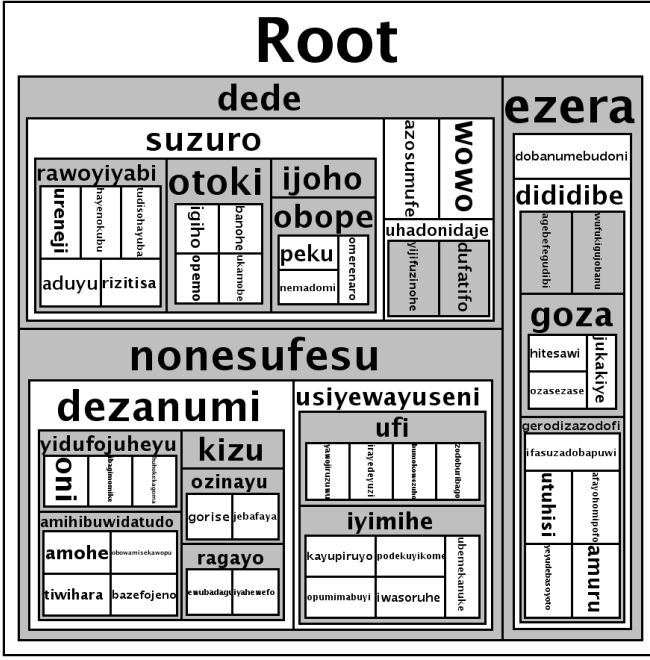


Fig. 16. A squarified treemap of a randomly generated tree.

Finally, we note that there is no need to consider any explicit metrics of size or aspect ratio of nodes, because these are both implicitly taken into account when evaluating the size of labels.

7 RECTIFIED TREEMAPS

Thinking about the 3rd design guideline in Section 6.2 led us to create a variant of squarified treemaps that increases the size of labels. The idea is that, rather than “squarifying” nodes to give them an aspect ratio closer to 1, we “rectify” them (i.e. make them into rectangles) to give them an aspect ratio closer to the aspect ratio of their labels. We implemented this idea with a simple change to the layout algorithm described in Section 3.2 of [6]. Let R be a list of rectangles that are tiled together to form a single row of total length w and total area s , and let $r(i)$ be the area of the i th rectangle in R . The row has thickness s/w , hence the i th rectangle has dimensions $wr(i)/s$ and s/w , and has aspect ratio $\lambda(w^2r(i)/s^2) \geq 1$ where we define the function $\lambda(x) = \max(x, 1/x) \geq 1$ to always have aspect ratios greater than or equal to 1. As defined in [6], the worst (i.e. largest) aspect ratio of the rectangles in the row is

$$\text{worst}(R) = \max_i \lambda(w^2r(i)/s^2)$$

The function worst is used within the algorithm in [6] to generate a squarified layout. Our one modification to the algorithm is to change the function to

$$\text{worst}(R) = \max_i \begin{cases} \lambda(\lambda(w^2r(i)/s^2)/L(i)), & \text{if } i\text{th rectangle is a leaf node} \\ \lambda(w^2r(i)/s^2), & \text{otherwise} \end{cases}$$

where $L(i) \geq 1$ is the aspect ratio of the label on the i th rectangle. In other words, the effective aspect ratio of leaf nodes is relative to the aspect ratio $L(i)$ of their labels. Bringing this *relative* aspect ratio closer to 1, as the squarified layout algorithm does, results in nodes that better match their label’s aspect ratio.

Figures 16 and 17 show the output before and after this change, for a randomly generated tree with random labels. In both cases, each leaf node was given equal weight for the treemap algorithm. As can be seen, the nodes in Figure 16 are often more squarish and contain much unused space, especially when they contain a long label.

Figure 18 shows the output after a 2nd minor change, whereby leaf nodes are given weight proportional to the number of characters in

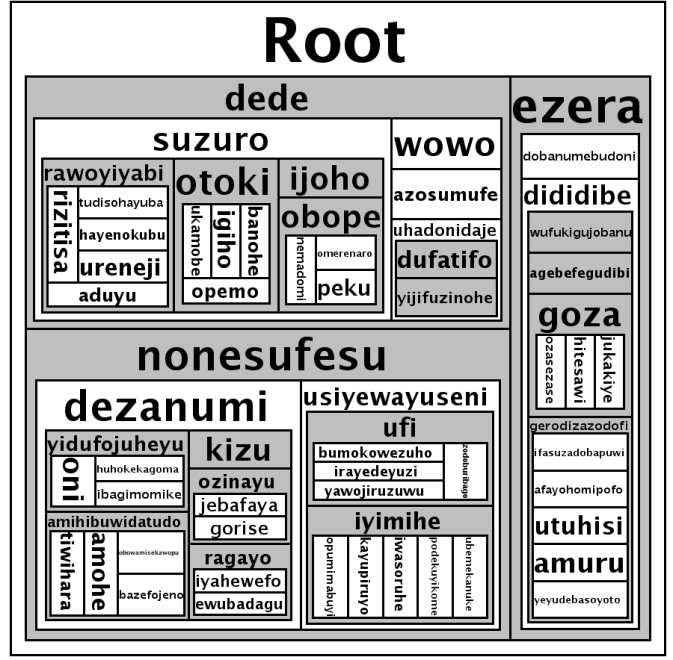


Fig. 17. A “rectified” treemap of the same tree in Figure 16, where leaf nodes are made to more closely match the aspect ratio of their labels.

their label, to give the same weight to each character. The goal, in this case, is to make each label as legible as possible. However, it would also be possible to give leaf nodes a different weight that is a function of some “size” attribute of the nodes (such as file size), and in general, the rectified treemap will have larger, more legible labels than those in a normal squarified treemap. (However, such rectified treemaps with a weighted partitioning will also have a disadvantage: the skinnier the labels are, the skinnier the nodes will be, and hence the more difficult it will be for a user to accurately compare the areas of nodes. So there is a tradeoff between label size and comparability of node areas.)

The following statistics show that our modifications yield an improvement over Figure 16 in all the label-oriented metrics. Comparing Figures 17 and 18, we see that although Figure 18 has smaller average leaf label height, it also has larger minimum leaf label height, and a smaller standard deviation of leaf label height, i.e. the heights are more consistent.

Figure	16	17	18
Total area A of nodes	1.000	1.000	1.000
Mean area exponent p	0.811	0.811	0.799
Average area of leaf nodes	0.00688	0.00688	0.00671
Total area A_l of labels	0.331	0.396	0.402
Mean area exponent p_l for labels	0.657	0.540	0.496
Average area of leaf labels	0.00244	0.00375	0.00342
Average leaf label height	0.0238	0.0297	0.0265
Minimum leaf label height	0.0105	0.0112	0.0172
Standard deviation of leaf label height	0.0115	0.0115	0.00598

8 CONCLUSIONS

We have presented a rigorous, quantitative analysis and comparison of the space-efficiency of several 2D graphical representations of trees, including most of the basic kinds found in the information visualization literature (or, at least, simplified versions of them). We are unaware of any previous similar analysis performed of *any* single tree representation, let alone all the ones considered here collectively. Our analysis makes clear the relative space-efficiency of each representation with respect to various metrics. This allows us to rank the tree representations, which is important for enabling designers to choose the most space-efficient possibility among those not eliminated by various design constraints. One theoretically interesting result of our work is

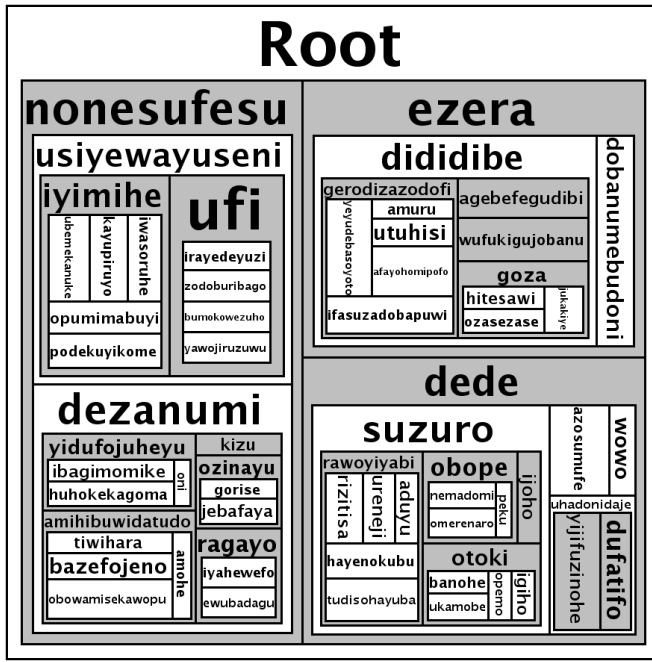


Fig. 18. The same as Figure 17, but now with each leaf node given weight proportional to the number of characters in its label, to give equal priority to each character.

that even treemaps (as modeled in Figure 9) are suboptimal with respect to certain metrics, having $A_{l,\infty} < 3/4$ for example.

Our analysis is performed in terms of both node size and label size, and investigates multiple facets of space-efficiency — a list of six kinds of metrics is given in Section 3.7. We have demonstrated that metrics which do not take into account label size can be misleading; for example, concentric circle representations were shown to be more space-efficient (in terms of leaf label size) than icicle diagrams, contrary to what would be expected from simply comparing their total area A . Fortunately, simply evaluating the area of labels can be done without separately evaluating the size or aspect ratio of nodes, since the latter two are both implicitly taken into account in analyzing the size of labels. Hence, we recommend using label-oriented metrics for evaluating representations. We propose that future work in this area go beyond claiming that a new representation is “space-filling” merely because of the total area A that is occupied, and that researchers apply the recommendations in Section 6.3 to evaluate the performance of new tree representations.

We have also presented a novel metric, the mean area exponent ρ , that quantifies the distribution of area in a representation, i.e. that quantifies the tradeoff between shallow and deep nodes. This metric can be calculated for both homogeneous and heterogeneous tree representations. We have also applied variants of the metric (ρ_l and its limit $\rho_{l,\infty}$) in our analysis of label areas, and discussed how certain values are associated with highly space-efficient representations ($\rho_{l,\infty}$ close to, or slightly greater than, 1). These metrics could be used to classify representations (for example, $\rho_{l,\infty}$ could be a dimension within a taxonomy of representations), or could be used as tools to summarize and better understand the performance of a representation, or as input parameters for tuning a given representation (e.g. Figure 13). Finally, the definition of ρ is not arbitrary, but affords multiple interpretations (Sections 3.3 and 3.6) that we find fairly compelling.

For designers, we have presented the most immediately applicable practical results of our work in Section 6, which includes a set of design guidelines and an example decision procedure for choosing a representation.

Finally, we have shown that our work has generative power, in the novel tree representations developed through this work: concentric

squares (Figure 8), asymmetrical concentric circles (Figure 11), a hybrid of indented outline and classical layered (Figure 12), and the rectified treemap (Figures 17 and 18). The concentric squares representation is more space-efficient than many commonly used representations, such as classical layered and concentric circles, and the rectified treemap improves the legibility of labels on status quo treemaps.

9 FUTURE DIRECTIONS

Future work could perform further analytical or numerical evaluation of the metrics, possibly on a corpus of real-world trees, and/or on other tree representations we have not considered such as the “space-optimized” representation of Nguyen and Huang [19] or hybrid mixtures of representations [32] or 3-dimensional representations such as cone trees [24].

Representations with greater space-efficiency might also be useful to develop, such as treemaps with margins that do not have the limitation of $A_{l,\infty} < 3/4$ of representation 9 (probably by placing some nodes to the right of their parent’s label), or a variation on Figure 10 that generalizes to $B \neq 8$. It is also possible that the efficiency of concentric squares (Figure 8) could be improved by laying out leaf nodes along a curve whose length is greater than 4, perhaps even along a space-filling curve. (Interestingly, space-filling curves have recently been used for graph layout [18].)

A significant issue for future metrics of tree representations is to take into account human perceptual abilities and/or how well various user tasks, such as those described in [29], are supported by each representation. For example, with a classical layout (1A) it is easier to estimate the distance between two nodes and easier to see if two nodes are on the same level than in a treemap — this is ignored by our metrics. As more experimental data emerges from researchers about user performance at various tasks with various tree representations, it would probably be reasonable to extend Table 2 with additional rows indicating the relative support for different tasks, and to refine the decision procedure in Section 6.2.

Our metrics also assume that there is no value in whitespace, whereas in practice whitespace can be very valuable in making groupings of elements easier to perceive. It may be possible to develop new metrics that weigh the value of label size along with the value of some whitespace to aid in the perception of structure and groupings.

Another issue not yet addressed by the metrics presented is how to model interactive displays that change over time, where nodes and labels may appear, move, change in size, or be elided, depending on the user’s input.

ACKNOWLEDGEMENTS

Early beginnings of this work benefited from conversations in 2004 with Shengdong Zhao at University of Toronto about taxonomies of tree representations. Expression (2) was first worked out by McGuffin in February 2008 on a napkin provided by Dunkin’ Donuts. Funding for this research came from a start-up grant from ETS. We also thank the reviewers for their insights and help, in particular for suggesting an analysis of the type of treemap described in the last paragraph of Section 4.2.11.

REFERENCES

- [1] V. Anand, K. Hansen, R. Jianu, and A. Rusu. Area-efficient visualization of web data. In *Proceedings of International Conference on Internet Computing (IC)*, pages 83–89, 2004.
- [2] K. Andrews and H. Heidegger. Information slices: Visualising and exploring large hierarchies using cascading, semi-circular discs. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis), Late Breaking Hot Topics*, pages 9–12, 1998.
- [3] B. B. Bederson and J. D. Hollan. Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of ACM Symposium on User Interface Software and Technology (UIST)*, pages 17–26, 1994.
- [4] J. Bertin. *Sémiologie graphique: Les diagrammes, Les réseaux, Les cartes*. Éditions Gauthier-Villars, Paris, 1967. (2nd edition 1973, English translation 1983).

- [5] R. Boardman. Bubble trees: The visualization of hierarchical information structures. In *Extended Abstracts of ACM CHI 2000 Conference on Human Factors in Computing Systems*, pages 315–316, 2000.
- [6] M. Bruls, K. Huizing, and J. J. van Wijk. Squarified treemaps. In *Proceedings of the joint Eurographics and IEEE TCVG Symposium on Visualization*, pages 33–42, 2000.
- [7] C. Buchheim, M. Jünger, and S. Leipert. Improving Walker’s algorithm to run in linear time. In *Proceedings of Symposium on Graph Drawing (GD)*, pages 344–353, 2002.
- [8] S. K. Card and D. Nation. Degree-of-interest trees: A component of an attention-reactive user interface. In *Proceedings of ACM Advanced Visual Interfaces (AVI)*, 2002.
- [9] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.
- [10] P. D. Eades. Drawing free trees. *Bulletin of the Institute for Combinatorics and its Applications*, 5:10–36, 1992.
- [11] D. C. Engelbart and W. K. English. A research center for augmenting human intellect, 1968. (Video of public demonstration of NLS) <http://www.1968demo.org/>.
- [12] A. Garg, M. T. Goodrich, and R. Tamassia. Planar upward tree drawings with optimal area. *International Journal of Computational Geometry and Applications*, 6:333–356, 1996.
- [13] I. Herman, G. Melançon, and M. S. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 6(1):24–43, January 2000.
- [14] B. Johnson and B. Shneiderman. Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *Proceedings of IEEE Visualization (VIS)*, pages 284–291, 1991.
- [15] K. G. Kakoulis and I. G. Tollis. A unified approach to labeling graphical features. In *Proceedings of Symposium on Computational Geometry*, pages 347–356, 1998.
- [16] D. E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms*, pages 309–310. Addison-Wesley, 1968.
- [17] H. Koike and H. Yoshihara. Fractal approaches for visualizing huge hierarchies. In *Proceedings of IEEE Symposium on Visual Languages (VL)*, pages 55–60, 1993.
- [18] C. Muelder and K.-L. Ma. Rapid graph layout using space filling curves. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, 2008.
- [19] Q. V. Nguyen and M. L. Huang. A space-optimized tree visualization. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, pages 85–92, 2002.
- [20] Q. V. Nguyen and M. L. Huang. EncCon: An approach to constructing interactive visualization of large hierarchical data. *Information Visualization*, 4:1–21, 2005.
- [21] C. Plaisant, J. Grosjean, and B. B. Bederson. SpaceTree: Supporting exploration in large node link tree, design evolution and empirical evaluation. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, pages 57–64, 2002.
- [22] K. Pulo and M. Takatsuka. Inclusion tree layout convention: An empirical investigation. In *Proceedings of Asia-Pacific Symposium on Information Visualisation*, pages 27–35, 2003.
- [23] E. M. Reingold and J. S. Tilford. Tidier drawings of trees. *IEEE Transactions on Software Engineering*, SE-7(2):223–228, March 1981.
- [24] G. G. Robertson, J. D. Mackinlay, and S. K. Card. Cone trees: animated 3D visualizations of hierarchical information. In *Proceedings of ACM 1991 Conference on Human Factors in Computing Systems (CHI)*, pages 189–194, 1991.
- [25] C.-S. Shin, S. K. Kim, and K.-Y. Chwa. Area-efficient algorithms for upward straight-line tree drawings. In *Proceedings of International Computing and Combinatorics (COCOON)*, 1996.
- [26] B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics (TOG)*, 11(1):92–99, January 1992.
- [27] J. Stasko and E. Zhang. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, pages 57–65, 2000.
- [28] S. T. Teoh and K.-L. Ma. RINGS: A technique for visualizing large hierarchies. In *Proceedings of Graph Drawing (GD)*, pages 268–275, 2002.
- [29] Y. Wang, S. T. Teoh, and K.-L. Ma. Evaluating the effectiveness of tree visualization systems for knowledge discovery. In *Proceedings of Eurographics/IEEE-VGTC Symposium on Visualization*, 2006.
- [30] M. Wattenberg. A note on space-filling visualizations and space-filling curves. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, pages 181–186, 2005.
- [31] J. Yang, M. O. Ward, and E. A. Rundensteiner. InterRing: An interactive tool for visually navigating and manipulating hierarchical structures. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)*, pages 77–84, 2002.
- [32] S. Zhao, M. J. McGuffin, and M. H. Chignell. Elastic hierarchies: Combining treemaps and node-link diagrams. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis) 2005*, pages 57–64, 2005.