

Peacock Bundles: Bundle Coloring for Graphs with Globality-Locality Trade-Off

Jaakko Peltonen^{1,2(✉)} and Ziyuan Lin¹

¹ Helsinki Institute for Information Technology HIIT,
Department of Computer Science, Aalto University, Espoo, Finland
ziyuan.lin@aalto.fi

² School of Information Sciences, University of Tampere, Tampere, Finland
jaakko.peltonen@aalto.fi

Abstract. Bundling of graph edges (node-to-node connections) is a common technique to enhance visibility of overall trends in the edge structure of a large graph layout, and a large variety of bundling algorithms have been proposed. However, with strong bundling, it becomes hard to identify origins and destinations of individual edges. We propose a solution: we optimize edge coloring to differentiate bundled edges. We quantify strength of bundling in a flexible pairwise fashion between edges, and among bundled edges, we quantify how dissimilar their colors should be by dissimilarity of their origins and destinations. We solve the resulting nonlinear optimization, which is also interpretable as a novel dimensionality reduction task. In large graphs the necessary compromise is whether to differentiate colors sharply between locally occurring strongly bundled edges (“local bundles”), or also between the weakly bundled edges occurring globally over the graph (“global bundles”); we allow a user-set global-local tradeoff. We call the technique “peacock bundles”. Experiments show the coloring clearly enhances comprehensibility of graph layouts with edge bundling.

Keywords: Graph visualization · Network data · Machine learning · Dimensionality reduction

1 Introduction

Graphs are a prominent type of data in visual analytics. Prominent graph types include for instance hyperlinks of webpages, social networks, citation networks between publications, interaction networks between genes, variable dependency networks of probabilistic graphical models, message citations and replies in discussion forums, traces of eye fixations, and many others. 2D or 3D visualization of graphs is a common need in data analysis systems. If node coordinates are not available from the data, several node layout methods have been developed, from constrained layouts such as circular layouts ordered by node degree to unconstrained layouts optimized by various criteria; the latter methods can be based on the node and edge set (node adjacency matrix) alone, or can make use of

multivariate node and edge features, typically aiming to reduce edge crossings and place nodes close-by if they are similar by some criterion.

In layouts with numerous edges it may be hard to see trends in node-to-node connections. Edge bundling draws multiple edges as curves that are close-by and parallel for at least part of their length. Bundling simplifies the appearance of the graph, and bundles also summarize connection trends between areas of the layout. However, when edges are drawn close together, the ability to visually follow edges and discover their start and end points is lost. Interactive systems [10] can allow inspection of edges, but inspecting numerous edges is laborious.

Comprehensibility of edges can be enhanced by distinguishing them by visual properties, such as line style, line width, markers along the curve, or color. Following an edge by its color can allow an analyst to see where each edge goes, but poorly assigned colors can make this task hard to do at a glance. **We present a machine learning method that optimizes edge colors in graphs with edge bundling, to keep bundled edges maximally distinguishable.** We focus on edge color as it has several degrees of freedom suitable for optimization (up to three continuous-valued color channels if using RGB color space), but our method is easily applicable to other continuous-valued edge properties. We call our solution **peacock bundles** as it is inspired by the plumage of a peacock; our method results in a fan of colors, reminiscent of a peacock tail, at fan-in locations of edges arriving into a bundle and fan-out locations of edges departing from a bundle. Figure 1 (middle) illustrates the concept and how it can help follow edges. We next review related works and then present the method and experiments.




Fig. 1. Illustration of peacock bundle coloring. **Left:** A graph with node groups A–F, drawn with hierarchical edge bundling. With plain gray coloring finding the connecting vertex pairs is not possible. **Middle:** Peacock bundle coloring reveals that nodes in group A connect to nodes in group E in order, and similarly B to D in order, and C to F in reverse order. The connections are easily seen from the optimized coloring produced by our Peacock Bundles method: bundled edges traveling from and to close-by nodes get close-by colors. **Right:** Pairwise bundling detection as described in Sect. 3.1, for three edges, control points z_{ij} shown as circles, distance threshold T as the radius of light gray circles (small threshold used for illustration). Control points z_{12}, \dots, z_{15} of edge 1 are near control points of edge 2, but only z_{13} is near a control point of edge 3. If, e.g., $K_{ij} = 2$ nearby control points are required between edges, edge 1 is considered bundled with edge 2 but not edge 3; edge 2 is considered bundled with edges 1 and 3, and edge 3 with edge 2 but not edge 1. Since edges 1 and 3 are not considered bundled they could be assigned a similar color. (Color figure online)

2 Background: Node Layout, Edge Bundling, and Coloring

Node layouts of graphs have been optimized by many approaches, see [9] for a survey. Our approach is not specific to any node layout approach and can be run for any resulting layout. Several methods have been proposed for edge bundling [4, 6–8, 13, 16, 17, 19]. For example, Cui et al. [4] generate a mesh covering the graph on the display based on node positions and edge distribution. The mesh helps cluster edges spatially; edges within a cluster are bundled. Hierarchical Edge Bundling [13] embeds a tree representation for data with hierarchy onto the 2D display. Tree nodes are used as spline control points for edges; bundles come from reusing control points. See Zhou et al. [22] for a recent review and taxonomy.

Unlike node layout and edge bundling, relatively little attention has been paid to practical edge coloring; while graph theory papers exist about the “edge coloring problem” of setting distinct colors to adjacent edges with a minimum number of colors, that combinatorial problem does not reflect real-life graph visual analytics where a continuous edge color space exists and the task is to set colors to be informative about graph properties. Simple coloring approaches exist. A naive coloring sets a random color to each edge: such coloring is unrelated to spatial positions of nodes and edges and is chaotic, making it hard to grasp an overview of edge origins and destinations at a glance. Edge colors are sometimes reserved to show discrete or multivariate annotations such as edge strengths; such coloring relies on external data and may not help gain an overview of the graph layout itself. A simple layout-driven solution is to color each edge by onscreen position of the start or end node. If edges have been clustered by some method, one often sets the same color to the whole cluster [7, 21]; this simplifies coloring, but prevents telling apart origins and destinations of individual edges.

Hu and Shi [15] create edge colorings with a maximal distinguishability motivation related to ours, but their method does not consider actual edge bundling and operates on the original graph; we operate on bundled graphs and quantify edge bundling. Also, instead of only using binary detection of bundled edge pairs (a hard criterion whether two edges are bundled) we differentiate all edges, emphasizing each pair by a weight that is high between strongly bundled local edges and smaller between others, with a user-set global-local tradeoff. Lastly, their method tries to set maximally distinct colors between all bundled edges, needing harder compromises for larger bundles: we quantify which bundled edges need the most distinct colors by comparing their origin and destination coordinates in the layout, and thus can devote color resources efficiently even in large graphs.

Our algorithm is related to nonlinear dimensionality reduction. Although dimensionality reduction has been used in colorization for other domains [3, 5], to our knowledge ours is the first method to optimize local graph coloring with edge bundling as a dimensionality reduction task.

3 The Method: Peacock Bundles

Bundle coloring has several challenges. **1.** Efficient coloring should depend not only on high-dimensional graph properties on the low-dimensional graph layout: if two edges are spatially distinct they do not need different colors. **2.** *Bundles are typically not clearly defined:* the curve corresponding to an individual edge may become *locally* bundled with several other edges at different places along the curve between its start and end node, and edges cannot be cleanly separated into groups that would correspond to some *globally* nonoverlapping bundles. Solutions requiring nonoverlapping bundles would be suboptimal: they would either not be applicable to real-life edge-bundled graphs or would need to artificially approximate the bundle structure of such graphs as nonoverlapping subsets. **3.** The solution should scale up to large graphs with large bundles. In large bundles it is typically not feasible to assign strongly distinct colors between all edges; it is then crucial to quantify how to make the compromise, that is, which edges should have the most distinct colors within the bundle.

Our coloring solution neatly solves these challenges, by posing the coloring as an optimization task defined based on local bundling between *each individual pair of edges*. Our solution is applicable to all graphs and takes into account the full bundle structure in a graph layout without approximations. For any two edges it is easy to define whether their curves are bundled (close enough and parallel) for some part of their length, without requiring a notion of a globally defined bundle; we optimize the coloring to *tell each edge apart from the ones it has been bundled with*. Such optimization makes maximally efficient use of the colors: two edges need distinct colors only if they are bundled together, whereas two edges that are not bundled can share the same color or very similar colors. Moreover, even between two bundled edges, how distinct their colors need to be can be quantified in a natural way based on the node layout: the more dissimilar their origins and destinations are, the more dissimilar their colors should be. Differentiating origins and destinations helps analysts assuming the node layout is meaningful. Computation of peacock bundles requires two steps:

1. *Detection of which pairs of edges are bundled together* at some location along their curve. We solve this by a well-defined closeness threshold of consecutive curve segments. An edge may participate in multiple bundles along its curve.
2. *Definition of the color optimization task.* We formalize the color assignment task as a dimensionality reduction task from two input matrices, a pairwise edge-to-edge bundling matrix and a dissimilarity matrix that quantifies how dissimilar colors of bundled edges should be, to a continuous-valued low-dimensional colorspace, which can be one-dimensional (1D) to achieve a color gradient, or 2D or 3D for greater variety. (Properties like width or continuous line-style attributes could be included in a higher than 3D output space; here we use color only.) We define the color assignment as an edge dissimilarity preservation task: colors are optimized to preserve spatial dissimilarities of start and end nodes among each pair of bundled edges, whereas no constraint is placed between colors of non-bundled edge pairs.

Peacock bundle coloring can be integrated into edge bundling algorithms, but can be also run as standalone postprocessing for graphs with edge bundling, regardless of which algorithms yielded the node layout and edge curves. Peacock bundles optimize colors taking both the graph and its visualization (node and edge layout) into account: color separation needs to be emphasized only for edges that appear spatially bundled. We demonstrate the result on several graphs with different node layouts and a popular edge bundling technique.

3.1 Detection of Bundled Pairs of Edges

Let the graph contain M edges $i = 1, \dots, M$, each represented by a curve. If the curves are spline curves, let each curve be generated by C_i control points; if the curves are piecewise linear, let each curve be divided into C_i segments represented e.g. by the midpoint of a segment. For brevity we use the terminology of control points in the following, but the algorithm can be used just as well for other definitions of a curve, such as midpoints of piecewise linear curves or equidistributed points on the curves if getting such locations is convenient.

Let B_{ij} be a variable in $[0, 1]$ denoting whether edge i is bundled together with edge j . If the edge bundling has been created by an algorithm that explicitly defines bundle memberships for edges, B_{ij} can simply be set to 1 for edges assigned to the same bundle and zero otherwise. However, for several situations this is insufficient: (i) sometimes the bundling algorithm is not available or the bundling has e.g. been created interactively; (ii) some bundling algorithms only e.g. attract edge segments and do not define which edges are bundled; (iii) an edge may be close to several different other edges, so that no single bundle membership is sufficient to describe its relationship to other edges. For these reasons we provide a way to define pairwise edge bundling variables B_{ij} that does not require availability of any previous bundling algorithm.

We set $B_{ij} = 1$ if at least K_{ij} consecutive control points of edge i are each close enough to one or more control points of edge j . Intuitively, if several consecutive control points of edge i are close to edge j , the edges travel close and parallel (as a bundle) at least between those control points. Since our choice of control points does not allow the curves to change drastically between two consecutive control points, the defined B_{ij} is stable when the control point densities between curve i and curve j do not differ too much. In practice, we set K_{ij} to an integer at least 1, separately for each pair of edges, as a fraction of the number of available control points as detailed later in this section.

Formally, for edge i denote the on-screen coordinates of the C_i control points by $\mathbf{z}_{i1}, \dots, \mathbf{z}_{iC_i}$, and similarly for edge j . Let $d(\cdot, \cdot)$ denote the Euclidean distance between two control points, and let T be a distance threshold. Then

$$B_{ij} = \max_{r_0=1, \dots, C_i - K_{ij} + 1} \prod_{r=r_0}^{r_0 + K_{ij} - 1} 1\left(\min_{s=1, \dots, C_j} d(\mathbf{z}_{ir}, \mathbf{z}_{js}) \leq T\right) \quad (1)$$

where $r = r_0, \dots, r_0 + K_{ij} - 1$ are indices of consecutive control points in edge i . The term $1(\cdot)$ is 1 if the statement inside is true and zero otherwise: that is, the

term is 1 if the r th control point of edge i is close to edge j (to some control point s of edge j). The whole product term is 1 if the K_{ij} consecutive control points of i from r_0 onwards are all close to edge j . Finally, the whole term B_{ij} is 1 if edge i has K_{ij} consecutive points (from any r_0 onwards) that are all close to edge j . Figure 1 (right) illustrates the pairwise bundling detection.

The distance threshold T should be set to a value below which line segments appear very similar; a rule of thumb is to set T to a fraction of the total diameter (or larger dimension) of the screen area of the graph. Similarly, a convenient way to set the required number of close-by control points K_{ij} is to set it to a fraction of the maximum number of control points in the two edges, requiring at least 1 control point, so that for each pair of edges i and j we set $K_{ij} = \max(1, \lfloor \max(C_i, C_j) K_{min} \rfloor)$ where $K_{min} \in (0, 1]$ is the desired fraction.

Detected pairwise bundles match ground truth in all simple examples we tried (e.g. Fig. 1 left); in experiments of Sect. 4 where no ground truth is available the bundling is visually good; edges bundled with any edge of interest can be interactively checked at <http://ziyuang.github.io/peacock-examples/>.

3.2 Optimization of Edge Colors by Dimensionality Reduction

Our coloring is based on dimensionality reduction of bundled edges from an original dissimilarity (distance) matrix to a color space; we thus need to define how dissimilar two bundled edges are. We aim to help analysts differentiate where in the graph layout each edge goes; we thus use the node locations of edges to define the similarity. Denote the two on-screen node layout coordinates of edge i by \mathbf{v}_i^1 and \mathbf{v}_i^2 . We first define

$$d_{ij}^{\text{original}} = \min(\|\mathbf{v}_i^1 - \mathbf{v}_j^1\| + \|\mathbf{v}_i^2 - \mathbf{v}_j^2\|, \|\mathbf{v}_i^1 - \mathbf{v}_j^2\| + \|\mathbf{v}_i^2 - \mathbf{v}_j^1\|). \quad (2)$$

Denote the set of p features for edge i as a vector $\mathbf{x}_i = [x_{i1}, \dots, x_{ip}]$, and denote the low-dimensional output features for edge i as a vector $\mathbf{y}_i = [y_{i1}, \dots, y_{iq}]$ where $q \in \{1, 2, 3\}$ is the output dimensionality. We define the dimensionality reduction task as minimizing the difference between the endpoint dissimilarity of bundled edges and dissimilarity of their optimized colors. This yields the cost function

$$\min_{\{\mathbf{y}_1, \dots, \mathbf{y}_M\}} \sum_i \sum_j B_{ij} (d_{ij}^{\text{original}} - d^{\text{out}}(\mathbf{y}_i, \mathbf{y}_j))^2 \quad (3)$$

where $d^{\text{out}}(\mathbf{y}_i, \mathbf{y}_j)$ is the Euclidean distance between the output features. The terms B_{ij} are large for only those pairs of edges that are bundled, thus minimizing the cost assigns colors to preserve dissimilarity within bundled edges, but allows freedom of color assignment between non-bundled edges. The cost encapsulates that greater difference of edge destinations should yield greater color difference, and that color differentiation is most needed for strongly bundled edges. While alternative formulations are possible, (3) is simple and works well.

From Local to Global Color Differentiation. The weights B_{ij} detect edges according to thresholds T and K_{ij} . Some edge pairs that fail the detection might still visually appear nearly bundled: instead of differentiating only within

detected bundles, it is meaningful to differentiate other edges too. The simplest way is to encode a *tradeoff* between local (within-bundle) and global differentiation in the B_{ij} : we set $B_{ij} = 1$ if edges i and j are bundled, otherwise $B_{ij} = \epsilon$ where $\epsilon \in [0, 1]$ is a user-set parameter for the preferred global-local tradeoff. When $0 < \epsilon < 1$, the cost emphasizes achieving desired color differences between bundled edges (where $B_{ij} = 1$) according to their dissimilarity of origins and destinations, but also aims to achieve color differences between other edges ($B_{ij} = \epsilon$) according to the same dissimilarity. As the optimization is based on desired dissimilarities between edges, it intelligently optimizes colors even when all edge pairs can have nonzero weight: $\epsilon = 0$ means a pure local coloring where only bundled edge pairs matter, and $\epsilon = 1$ means a pure global coloring that aims to show dissimilarity of origin and destination for all edges regardless of bundling. In our tests coloring changes gradually with respect to ϵ . In experiments, when emphasizing local color differences, we set $\epsilon = 0.001$ which achieved local differentiation and formed color gradients for bundles in most cases.

A way to set a more nuanced tradeoff is to run edge detection with multiple settings and set weaker B_{ij} for edges detected with weaker thresholds; in practice the above simple tradeoff already worked well.

Relationship to Nonlinear Multidimensional Scaling. Interestingly, minimizing (3) can be seen as a specialized weighted form of *nonlinear multidimensional scaling*, with several differences: unlike traditional multidimensional scaling we treat edges (not data items or nodes) as input items whose dissimilarities are preserved; our output is not a spatial layout but a color scheme; and most importantly, the cost function does not aim to preserve all “distances” but weights each pairwise distance according to how strongly that pair of edges is bundled. The theoretical connection lets us make use of optimization approaches previously developed for multidimensional scaling, here we choose to use the popular stress majorization algorithm (SMACOF) [1] to minimize the cost function.

Color Range Normalization. After optimization, output features \mathbf{y}_i of each edge must be normalized to the range of the color channels (or positions along a color gradient). Simple ideas like applying an affine transform to the output matrix $Y = (\mathbf{y}_1, \dots, \mathbf{y}_M)$ would give different amounts of color space to different bundles, thus colors within bundles would not be well differentiated. We propose a normalization to maximally distinguish edges within each bundle. Let Col denote the color matrix to be obtained from normalization. For each \mathbf{y}_i , let $\{\mathbf{y}_{il}\}_{l=1}^{M_i}$ be the set of output features where each edge l is bundled with edge i . We assemble \mathbf{y}_i and $\{\mathbf{y}_{il}\}_{l=1}^{M_i}$ into a matrix $Y^i = (\mathbf{y}_i, \mathbf{y}_{i_1}, \dots, \mathbf{y}_{i_{M_i}})$, affinely transform Y^i to $\tilde{Y}^i = (\tilde{\mathbf{y}}_i, \tilde{\mathbf{y}}_{i_1}, \dots, \tilde{\mathbf{y}}_{i_{M_i}})$ so that each entry in \tilde{Y}^i is within the allowed range (say, $[0, 1]$), then set Col_i , the i -th column of Col (color vector for edge i) as $\tilde{\mathbf{y}}_i$. This normalization expands the color range within bundles.

Where to Show Colors. The optimized colors can be shown along the whole edge, or at “fan-in” segments where the edge enters a bundle and “fan-out” segments where it departs a bundle. Edge i is bundled with j if several consecutive curve segments of i are close to j ; the last segment before the close-by ones is the fan-in segment; the first segment after the close-by ones is the fan-out

segment. In experiments we show color along the whole edge for simplicity. Note that, as with any edge coloring, colors of close-by edges may perceptually blend, but our optimized colors then remain visible at fan-in and fan-out locations.

4 Experiments

We demonstrate the Peacock bundles method on five graphs (Figs. 2, 3 and 4): two graphs with hierarchical edge bundling [13], and three with force-directed edge bundling [14]. The two graphs with hierarchical edge bundling are created



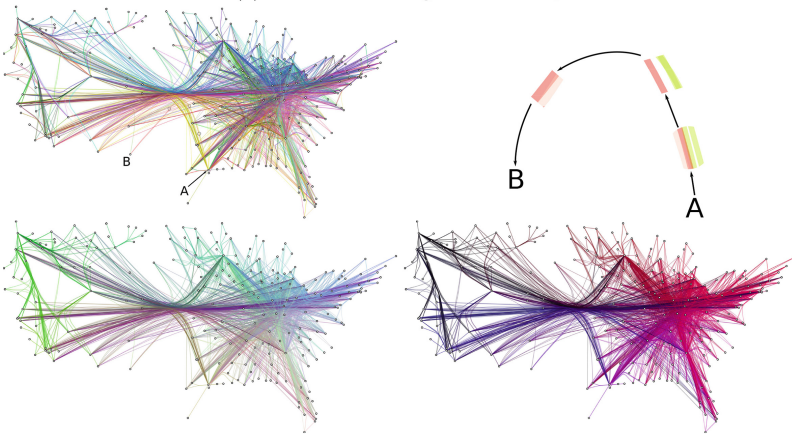


Fig. 2. Colorings for the graph “radial”. **Top left:** the coloring from Peacock with $\epsilon = 0.001$. **Top right:** zoomed-in versions of the parts within dashed-line circles in the top-left figure as examples of local coloring. In the four zoomed-in parts, colors show a linear gradient and vary in yellow-red-blue, thus (1) local colors are differentiated, and (2) they span roughly the same full color range. The local colors also help follow edges at bottom right of the graph, where colors are homogeneous in the baseline coloring. **Bottom left:** coloring from Peacock, $\epsilon = 1$. The bundles are colored into 3 parts: the blue-ish upper half, green-ish lower-left part, and yellow-ish lower-right part. There are also red bundles joining the blue and green parts, differentiating itself from other bundles. **Bottom right:** the baseline coloring. The bundles are colored into the red-ish upper half and blue-ish lower half. Compared with the coloring with $\epsilon = 1$ from Peacock, the bundle from left to right, and the bundle at the top-left corner are less distinguishable. (Color figure online)



(a) Different coloring for “tree map”



(b) Different coloring for “airline”

Fig. 3. Colorings for the graphs “tree map” and “airline”. In both subfigures: **Top left:** the colorings from Peacock with $\epsilon = 0.001$. In Fig. 3a, the local linear gradient is clearer at the ends of the bundles. In Fig. 3b, the large bundle in the middle shows the local coloring, by separating the bundle into the upper blue dominating part, the middle red-ish part, and the lower lighter part. **Top right:** examples of how the colorings enhance readability by investigating the parts within dashed line circles in both top-left subfigures. In Fig. 3a, the colors in bundle A help the user to recognize, for example, 1) the blue-ish part in bundle A leads to the blue-ish part of the top-right “claw” or the right “claw”; 2) the red-ish part in bundle A leads to the red-ish part of the “claw” at the right of bundle B, or the top right “claw”; 3) the yellow-ish half that joins in the middle leads to bundle B or to the “claw” at the right of bundle B. Figure 3b shows how the coloring help a pink edge from A to B stand out against other edges in the same bundle. **Bottom left:** the coloring from Peacock with $\epsilon = 1$. In Fig. 3a, bundles are globally differentiated. In Fig. 3b, for nodes of large degrees, the edges connecting to them have distinct colors for different directions. **Bottom right:** the baseline coloring. In Fig. 3a, the user may mis-recognize that there are edges from bundle A to bundle B. In Fig. 3b, unlike the bottom-left subfigure, edges connecting to the same node tend to have similar colors. (Color figure online)




Fig. 4. Colorings for graph “Jane Austen”. **Top left:** the coloring from Peacock with $\epsilon = 0.001$. The “crossing” at the right half of the figure shows a clear differentiation: red edges go from upper-right to lower-left; green edges go from upper to lower; blue edges go from upper-left to lower-right. Without the local coloring, it is difficult to tell whether the bundles or the edges are crossing or just tangential to one another. **Top right:** another example from the zoomed-in version of the part within the dashed line circle in the top-left figure. Edge colors change from green to purple-ish. The colors help the user follow the edges after the heavily bundled part in the middle: red edges mostly go leftwards, green edges are scattered, and blue edges mostly go rightwards. **Bottom left:** the coloring from Peacock with $\epsilon = 1$. Less locality but more globality. The earlier blue edges in the right “crossing” become purple, but still distinguishable from the other two bundles. **Bottom right:** the baseline coloring, which loses the distinguishability shown in the Peacock coloring. (Color figure online)

from the class hierarchy of the visualization toolkit Flare [11], with the built-in radial layout (graph named “radial”; Fig. 2) and tree map layout (“tree map”; Fig. 3a) in d3.js [2] respectively. The four graphs with force-directed edge bundling are: a spatial graph of US flight connections (“airline”; Fig. 3b); a graph of consecutive word-to-word appearances in novels of Jane Austen (“Jane Austen”; Fig. 4); and a graph of matches between US college football teams (“football”; Fig. 5). The last three graphs are laid out as an unconstrained 2D graph by a recent node-neighborhood preserving layout method [18]. For all




Fig. 5. Colorings for graph “football”. **Top left:** the coloring from Peacock with $\epsilon = 0.001$. The uncertainty in this graph is mostly from the small clusters of nodes at the end of the edges. **Top right:** an example showing how the coloring help distinguish heavily bundled edges from A to B. The shown segments are the zoomed-in version of the parts within the dashed line circle or ellipse in the top-left figure. We can see, for example, that the blue edge leads to the top node in B, while the yellow edge leads to leftmost node in B. This is also noticeable in the top-left figure, particularly for the part between A and B is heavily bundled. **Bottom left:** the coloring from Peacock with $\epsilon = 1$. Colors of edges from the same cluster are differentiated (e.g., at the top-left cluster, edge colors vary from red to blue). **Bottom right:** the baseline coloring, which only reflects the locations of the bundles. (Color figure online)

graphs, edge bundles were created by a d3.js plugin implementing the algorithm [20] adapted to splines. All coloring are compared with a baseline coloring from end point positions.

The Baseline. We compare our method with a baseline coloring that directly encodes end point positions into color channels. We choose channel red and blue

for the encoding in the experiments. Let $\mathbf{v}_i^1 = (x_i^1, y_i^1)$ and $\mathbf{v}_i^2 = (x_i^2, y_i^2)$ be the onscreen coordinates of edge i 's two end points as in (2). We first create a 3-dimensional vector $\widetilde{Col}_i^{\text{baseline}}$ as the “unnormalized” color for edge i as

$$\widetilde{Col}_i^{\text{baseline}} = (\min(x_{i,1}, x_{i,2}), 0, \min(y_{i,1}, y_{i,2}))^T \quad (4)$$

then we affinely normalize the matrix $\widetilde{Col}^{\text{baseline}}$ into $[0, 1]$ to obtain the final baseline colors Col^{baseline} .

Choices of Peacock Parameters. The parameters T and K_{\min} in (1) must be chosen to determine B_{ij} . We set T to 2–4% of $\max(\text{graph width}, \text{graph height})$, and fix K_{\min} as 0.4. Experiments show the choices give good results empirically.

Figures 2, 3 and 4 show the results from the proposed method and the baseline. The top-left subfigures are with the tradeoff parameter set to prefer locality in the coloring. The top-right subfigures provide zoomed-in views detailing the local color variation (“peacock fans”) and demonstrating how the coloring improves readability and helps follow edges. The bottom-left figures are optimized to differentiate origins and destinations globally (tradeoff parameter $\epsilon = 1$), hence colors indicate overall trends of connections between areas of the graph layout, at the expense of less color variability within bundles. The bottom-right figures are from the baseline, also aiming to show variability of endpoint positions the coloring but not optimized by machine learning; the simple baseline coloring leaves bundles and within-bundle variation less distinguishable.

5 Conclusions

We introduced “peacock bundles”, a novel edge coloring algorithm for graphs with edge bundling. Colors are optimized both to preserve differences between bundle locations, and differentiate edges within bundles. The algorithm is based on dimensionality reduction without need to explicitly define bundles. Experiments show the method outperforms the baseline coloring with several graphs and bundling algorithms, greatly improving the comprehensibility of graphs with edge bundling. Potential future work includes incorporating color perception models [12], and more nuanced weighting schemes for global-local tradeoffs.

We acknowledge computational resources from the Aalto Science-IT project. Authors belong to the COIN centre of excellence. The work was supported by Academy of Finland grants 252845 and 256233.

References

1. Borg, I., Groenen, P.J.F.: Modern Multidimensional Scaling: Theory and Applications. Springer Series in Statistics, 2nd edn. Springer, New York (2005)
2. Bostock, M., Ogievetsky, V., Heer, J.: D3 data-driven documents. IEEE Trans. Vis. Comput. Graph. **17**(12), 2301–2309 (2011)

3. Casaca, W., et al.: Colorization by multidimensional projection. In: Proceedings of SIBGRAPI 2012, pp. 32–38. IEEE (2012)
4. Cui, W., Zhou, H., Qu, H., Wong, P.C., Li, X.: Geometry-based edge clustering for graph visualization. *IEEE Trans. Vis. Comput. Graph.* **14**(6), 1277–1284 (2008)
5. Daniels, J., Anderson, E.W., Nonato, L.G., Silva, C.T., et al.: Interactive vector field feature identification. *IEEE Trans. Vis. Comput. Graph.* **16**(6), 1560–1568 (2010)
6. Dwyer, T., Marriott, K., Wybrow, M.: Integrating edge routing into force-directed layout. In: Kaufmann, M., Wagner, D. (eds.) GD 2006. LNCS, vol. 4372, pp. 8–19. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-70904-6_3](https://doi.org/10.1007/978-3-540-70904-6_3)
7. Ersoy, O., Hurter, C., Paulovich, F.V., Cantaneira, G., Telea, A.: Skeleton-based edge bundling for graph visualization. *IEEE Trans. Vis. Comput. Graph.* **17**(12), 2364–2373 (2011)
8. Gansner, E.R., Hu, Y., North, S., Scheidegger, C.: Multilevel agglomerative edge bundling for visualizing large graphs. In: Proceedings of PacificVis 2011, pp. 187–194. IEEE (2011)
9. Gibson, H., Faith, J., Vickers, P.: A survey of two-dimensional graph layout techniques for information visualisation. *Info. Vis.* **12**(3–4), 324–357 (2013)
10. Grossman, T., Balakrishnan, R.: The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor’s activation area. In: Proceedings of CHI 2005, pp. 281–290. ACM (2005)
11. Heer, J.: Flare (2009). <https://git.io/v6buH>
12. Heer, J., Stone, M.: Color naming models for color selection, image editing and palette design. In: Proceedings of CHI 2012 (2012)
13. Holten, D.: Hierarchical edge bundles: visualization of adjacency relations in hierarchical data. *IEEE Trans. Vis. Comput. Graph.* **12**(5), 741–748 (2006)
14. Holten, D., Van Wijk, J.J.: Force-directed edge bundling for graph visualization. *Comput. Graph. Forum* **28**(3), 983–990 (2009)
15. Hu, Y., Shi, L.: A coloring algorithm for disambiguating graph and map drawings. In: Duncan, C., Symvonis, A. (eds.) GD 2014. LNCS, vol. 8871, pp. 89–100. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-45803-7_8](https://doi.org/10.1007/978-3-662-45803-7_8)
16. Hurter, C., Ersoy, O., Telea, A.: Graph bundling by kernel density estimation. *Comput. Graph. Forum* **31**(3pt1), 865–874 (2012)
17. Luo, S.J., Liu, C.L., Chen, B.Y., Ma, K.L.: Ambiguity-free edge-bundling for interactive graph visualization. *IEEE Trans. Vis. Comput. Graph.* **18**(5), 810–821 (2012)
18. Parkkinen, J., Nybo, K., Peltonen, J., Kaski, S.: Graph visualization with latent variable models. In: Proceedings of MLG 2010, pp. 94–101. ACM (2010)
19. Pupyrev, S., Nachmanson, L., Kaufmann, M.: Improving layered graph layouts with edge bundling. In: Brandes, U., Cornelsen, S. (eds.) GD 2010. LNCS, vol. 6502, pp. 329–340. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-18469-7_30](https://doi.org/10.1007/978-3-642-18469-7_30)
20. Sugar, C.: d3.forcebundle (2016). <https://git.io/v6GgL>
21. Telea, A., Ersoy, O.: Image-based edge bundles: simplified visualization of large graphs. *Comput. Graph. Forum* **29**(3), 843–852 (2010)
22. Zhou, H., Xu, P., Yuan, X., Qu, H.: Edge bundling in information visualization. *Tsinghua Sci. Technol.* **18**(2), 145–156 (2013)