A Task Taxonomy for Edge and Trail Bundling

Markus Wallinger* TU Munich

Stephen Kobourov† TU Munich

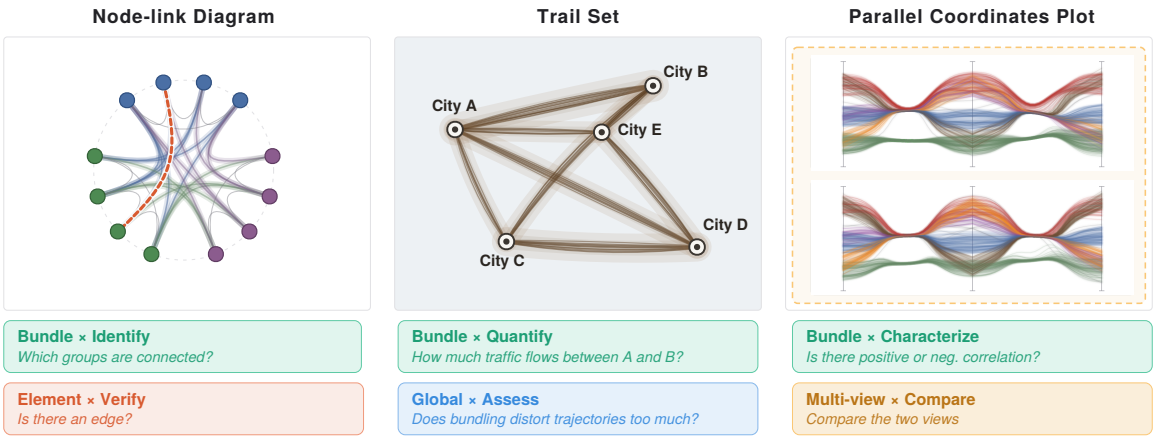

Figure 1: Edge bundling illustrated on three representation types, each supporting different tasks from our Scope × Action taxonomy. (a) A node-link diagram: bundling reveals which groups of modules are strongly connected (Bundle × Identify) but makes it harder to verify whether two specific modules share a direct dependency (Element × Verify). (b) A geographic trail set: bundling supports estimating traffic volume between cities (Bundle × Quantify) but raises the question of whether routes have been shifted away from their true paths (Global × Assess). (c) A bundled parallel coordinate plot: bundling makes it possible to see whether variables increase or decrease together (Bundle × Characterize) and to compare patterns across two datasets (Multi-view × Compare). Different colors in the figure correspond to the four different scope levels: Element, Bundle, Global, and Multi-view.

Abstract

Edge bundling reduces visual clutter by aggregating similar edges, yet practitioners lack a structured vocabulary for reasoning about the tasks that bundled visualizations support. Such a vocabulary is needed both to evaluate the general utility of bundling and to compare different bundling approaches. We address this gap by assembling a corpus of 102 papers, 49 of which contain explicit bundling tasks, spanning node-link diagrams, geographic trail sets, and parallel coordinate plots. From this corpus, we derive a task taxonomy organized as a matrix of *scope* (Element, Bundle, Global, Multi-view) crossed with *action* (Verify, Identify, Characterize, Quantify, Compare, Assess), instantiated across the three representation types. We show that bundling simultaneously *enables* tasks (bundle-level and global reasoning) and *disables* others (element-level precision), a duality not captured by existing task frameworks. Our coded corpus and taxonomy are released as supplemental material on OSF (osf.io/23r67).

Index Terms: Task taxonomy, edge bundling, trail bundling, parallel coordinates, network visualization.

1 Introduction

Edge and trail bundling [9, 12, 17, 24] spatially aggregates similar edges and paths in a drawing, trading visual clutter for overdraw and abstraction. Bundling techniques apply to node-link diagrams, geographic trail sets, and parallel coordinate plots (PCPs) alike. Yet, their evaluation has centered largely on algorithmic quality metrics and visual aesthetics [17, 21, 29] rather than on a systematic understanding of the analytic tasks bundling is meant to support, with perceptual task-based studies remaining rare [26].

The proposed task taxonomy for bundling serves three main purposes. First, it provides a basis for *evaluating* whether bundling is useful in a given analytical setting. Second, it makes it possible to compare different bundling methods. Third, it helps *connect* existing bundling quality metrics (e.g., ink reduction, distortion measures, ambiguity) to the tasks they are meant to serve, clarifying which metrics are relevant for which analytical goals.

A 2017 edge bundling survey by Lhuillier [17] includes a brief discussion of tasks, mapping bundling capabilities to the general task frameworks of Lee et al. [15] and Brehmer and Munzner [3]. The treatment is compact and largely conceptual, and the survey itself concedes that no accepted metrics exist to substantiate claims of task support. We identify three specific shortcomings.

First, existing task discussions fail to account for the fact that bundling introduces a new class of perceptual objects, the *bundles*, which support tasks absent from more general taxonomies. Identifying the major connection groups in a network, comparing the relative thickness of two bundles, or identifying which edges belong to a bundle make sense only after bundling has been applied. No prior framework provides vocabulary for these bundle-level tasks.

Second, task semantics shift substantially across representation types. Tracing in a node-link diagram (following an edge between named endpoints), in a trail set (following a spatial trajectory), and in a PCP (reading a multivariate profile) are fundamentally different operations despite sharing an abstract verb.

Third, prior discussions overstate what bundling supports. The trade-off between precision and aggregation means that some tasks are *enabled* by bundling (bundle-level and global reasoning) while others are *disabled* by it (individual element tracing and attribute reading). Our corpus analysis confirms that tasks such as characterizing or comparing individual elements are missing from the

*e-mail: markus.wallinger@tum.de
†e-mail: stephen.kobourov@tum.de

bundling literature – not because they are meaningless, but because bundling makes them impractical. At the same time, tasks that bundling clearly enables, such as global faithfulness assessment, were absent from prior task discussions.

We address these gaps through a systematic coding study. We assemble a corpus of 102 bundling papers spanning node-link diagrams, trail sets, and PCPs. Unlike prior work, which maps bundling *onto* existing generic frameworks, we derive a task vocabulary *from* bundling practice: from how tasks are actually stated, evaluated, and implied across the literature. From this corpus, we derive a two-dimensional taxonomy of *scope* × *action* (verify, identify, characterize, quantify, compare, assess), instantiated across the three representation types; see Fig. 1 for examples.

2 Background

Bundling across Representation Types. Bundling operates on three data types with different semantics.

A **node-link diagram** maps an abstract graph $G{=}(V,E)$ to the plane with a layout algorithm determining node positions and edge routing. Bundling modifies edge routing, but preserves node positions and the underlying relational data [16, 27].

A **trail set** is a collection of spatially embedded trajectories [13] (flights, eye tracking, vehicle paths). Here, the geometry *is* the data: deforming a trail through bundling distorts real spatial information, making faithfulness assessment qualitatively different from the graph case. However, bundling does not alter the underlying origin-destination data or the recorded positions; it changes only the visual routing between them.

A **parallel coordinate plot** (PCP) represents multivariate records as polylines crossing parallel axes. Bundling aggregates records with similar value profiles [7, 18, 20]. The "edges" here are data records, not relationships. As in the trail case, bundling modifies the visual routing of polylines but does not alter the underlying data values at the axes; the distortion occurs in the inter-axis space.

In all three cases, bundling changes *how* elements are drawn without changing *what* they represent. The severity of the resulting visual distortion, however, differs across representation types. We use the term *element* to refer to the atomic unit in each context: an edge in a graph, a trail in a trajectory set, or a polyline in a PCP.

Task Frameworks. Our taxonomy builds on Lee et al.'s graph task taxonomy [15], Brehmer & Munzner's multi-level typology [3], the clutter-reduction taxonomy of Ellis & Dix [6], and group-level graph visualization taxonomies [22, 25], which introduced tasks for reasoning about *groups* of nodes and edges. Bundle-level reasoning is closer to this group-level analysis than to the node- and edge-level focus of Lee et al.

The Scope–Action–Target (SAT) framework of Oddo et al. [19] provides the closest structural template, organizing tasks into triplets of scope (the scale of data considered), action (a verb defining the goal), and target (the data entity of interest). Our framework adopts SAT's scope and action dimensions but differs in the role of the target. In the context of unlabelled graphs, the target dimension addresses interesting graph-theoretic structures. In bundling, the target is simpler: at the element scope, it is always an individual edge, trail, or polyline; at the bundle scope, it is the perceptual aggregate; at the global scope, it is the entire visualization. Different bundling techniques (directional, density-based, hierarchical) do not constitute distinct targets, since humans reason about bundles as aggregate visual objects regardless of how they were produced. The variety of what we can *do* with bundles is captured by the action dimension. We therefore organize our taxonomy as a **Scope** × **Action** matrix, instantiated by representation type.

Bundling Methods and Quality Metrics. Representative algorithmic families include hierarchical edge bundling [9], force-directed bundling [10], image-based methods [12, 17, 24], geometry-based approaches [4, 14], confluent drawings [2], and Edge-Path bundling [1,27,28,30]. For a comprehensive algorithmic treatment, we refer to [17].

Several quality metrics [29] have been proposed, including ink reduction [5], edge displacement [21, 27], and ambiguity [27]. However, these metrics are typically discussed in isolation from underlying tasks, with few exceptions [26]. Our taxonomy provides a principled way to connect them. For example, ambiguity metrics benefit element- or bundle-level Identify and Verify tasks, and ink reduction metrics serve bundle-level and global-level Identify tasks.

3 Methodology

We assembled a corpus of 102 papers by collecting papers cited in a prior bundling survey [17], supplemented by a forward-citation search on bundling techniques. We included papers that either (a) propose or evaluate a bundling method, (b) use bundling in an application, or (c) conduct user studies involving bundled visualizations. While this approach does not guarantee exhaustive coverage, it captures the core bundling literature and a broad cross-section of application papers spanning all three representation types.

Coding procedure. Two authors developed the coding scheme on a random subset of 30 papers using open coding, identifying every passage where a paper states, evaluates, or implies a task, and agreeing on a reference set of codes for that subset. For each task mentioned, we recorded: the scope level, the action type, whether the mention was *explicit* (stated as a task, evaluated in a study, or listed as a design requirement) or *implicit* (inferred from system design, usage scenario, or figure captions), the representation type, and the paper's original terminology. After the initial round, we organize tasks along two dimensions:

- **Scope:** Element, Bundle, Global, or Multi-view (the perceptual granularity at which attention operates).
- **Action:** verify, identify, characterize, quantify, compare, or assess (the analytical goal).

Because coding proceeded by consensus on this shared reference set rather than by independent double-coding, we do not report a formal inter-rater reliability coefficient; the agreed codes served as the reference standard, and a human coder examined every retained mention to confirm that the raw text and its evidence supported the assigned tag.

Bundling attribution filter. Many papers use bundling as one component within a larger visualization system (e.g., NodeTrix [8] combines adjacency matrices with bundled edges). We applied a bundling attribution filter during coding: a task was attributed to bundling only if bundled edges or bundling-created aggregates were the visual component that supported or hindered the task. Tasks performed on non-bundled components (e.g., reading an adjacency matrix cell) were excluded. This ensures that our corpus analysis reflects what people can do specifically *with bundles*, not what they do with visualization systems that happen to include bundling.

LLM-assisted screening. Following recent work showing that LLMs can support qualitative coding when given a clear codebook [23], we used an LLM (Claude Sonnet 4.6) as a first-pass screener. Each paper was submitted with our coding scheme, synonym lists mapping literature phrasings to our canonical labels, and the bundling attribution filter with worked examples; the model returned candidate task mentions with proposed classifications and attribution evidence. Then, a human coder reviewed every annotation, correcting approximately 15% of the LLM's codes. These were mostly over-attributions from related- or future-work passages, duplicates, and scope confusions between Global and Multi-view (e.g., a cross-view comparison miscoded as Global × Compare), rather than genuine misclassifications. The 15% reflects the first-pass screening filter, not residual error in the human-verified corpus. The prompt is available in the supplemental material.

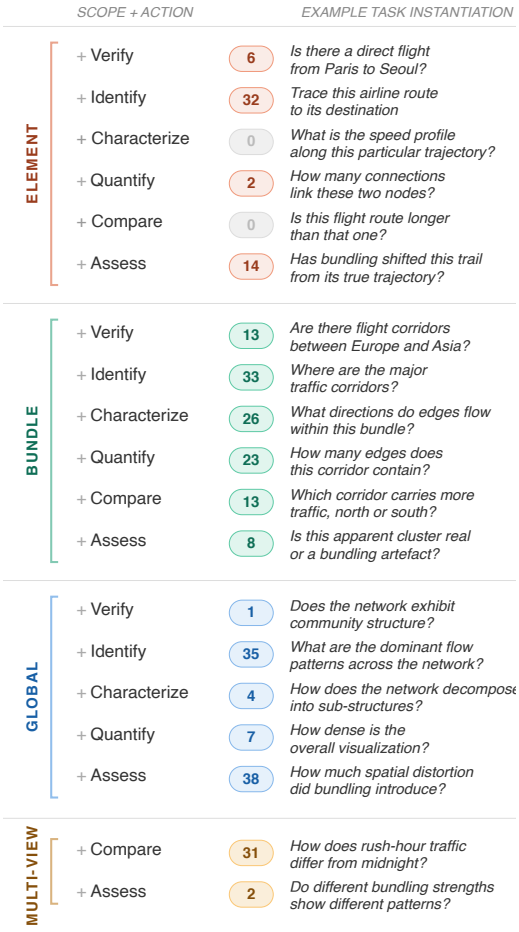

Figure 2: The Scope × Action taxonomy. Each row pairs an action type with an example task instantiation. Colored ovals indicate the number of unique papers that mention each combination. Gray ovals mark tasks that bundling hinders: coherent analytical questions that received zero mentions because bundling visually merges the individual elements needed to perform them. Structurally empty cells (e.g., Multi-view × Verify) are omitted entirely. The full coded corpus is available in the supplemental material (Table 1).

4 Task Taxonomy

Our taxonomy is a matrix of four *scope levels* crossed with six *action types*; see Fig. 2. Each combination represents an abstract task and the scope level determines the target. Concrete instantiations vary by representation type. Not every cell is populated: some combinations are *structurally empty*, meaning no meaningful task exists (e.g., Multi-view × Verify collapses into within-view verification), while others are *hindered by bundling*, meaning the task is coherent but bundling makes it difficult by merging the elements needed to perform it (e.g., Element × Characterize).

4.1 Scope: Four Levels of Attention

Bundling produces a hierarchy of perceptual objects that determines where human attention is directed.

At the **element level**, the focus is on single connections, trails, or records. Because bundling visually merges elements, tasks at this level often require mental disentanglement or interaction support such as lensing [31] or selective unbundling [11]. Two element-level tasks, Characterize and Compare, are hindered by bundling: characterizing the profile of one edge or comparing two individual edges requires isolating them from the bundle first, effectively undoing what bundling provides. Verify and Quantify are also difficult here, as both require disentangling overlapping paths.

At the **bundle level**, one engages with the perceptual aggregates that bundling creates. Bundles are emergent visual constructs not present in the raw data as they arise from the algorithmic transformation. Tasks at this level include identifying the major connection groups or comparing bundle thickness. These are related to the group-level tasks discussed in Sec. 2 but are specific to the bundling context, where the groups are defined by spatial proximity rather than by data attributes.

At the **global level**, one reasons about the entire visualization to perceive macro-patterns (communities, flows, coverage, anomalies) that would be obscured by clutter in an unbundled view. Global × Compare is structurally empty: there is only one global view per visualization, so any comparison requires a second view, which falls under the multi-view scope.

At the **multi-view level**, we compare two or more bundled views showing different datasets, time steps, parameter settings, or spatial extents. This scope is inherently one of comparison. One might argue that "finding the same cluster in both views" is an Identify task rather than a Compare task, but locating a structure in view B that was first seen in view A is operationally a comparison: we need to match across views, which requires assessing similarity. Verify, Identify, Characterize, and Quantify, when applied across views, all reduce to either a within-view operation at a lower scope or a comparison of what is seen in each view. We therefore populate only Compare and Assess at this level. Multi-view × Assess captures the specific case where a user evaluates whether the choice of bundling parameters or representation type affects which patterns are visible.

4.2 Action: Six Question Types

The action dimension captures the analytical goals. We identify six actions, each defined by a distinctive verb.

Verify (*Does X exist?*) Check for the presence of a specific element, bundle, or structural pattern. At the element level, Verify often depends on Identify: confirming that an edge exists typically requires tracing it first. This dependency helps explain why Verify appears less frequently than Identify in the corpus.

Identify (*What or where is X?*) Locate a target or read a basic property. This subsumes tracing (following an element) and attribute lookup (reading a single value such as weight).

Characterize (*What defines X internally?*) Decompose an entity into its parts or properties. While Identify answers "which one?" or "where is it?" (locating or surfacing a target), Characterize answers "what kind?" by describing internal structure (e.g., edge directions within a bundle or the correlation pattern within a cluster).

Quantify (*How much of X?*) Estimate magnitude; for example, element counts, bundle thicknesses, global density, etc.

Compare (*How does X differ from Y?*) Assess similarities or differences between entities at the same scope level. At the bundle level, this includes comparing thickness or detecting splits and merges; at the multi-view level, comparing bundled views across datasets, time steps, or parameter settings.

Assess (*Is X faithful?*) Evaluate whether the bundled representation truthfully reflects the underlying data. This covers distortion (how far bundling displaced elements), artifact detection (whether perceived patterns are genuine), and parameter sensitivity. This action is specific to bundling.

Compound tasks. Some analytical operations span multiple cells. For example, determining that a particular node bridges two clusters requires first locating the clusters (Bundle × Identify), then tracing individual edges from the node to verify it connects to both (Element × Identify). Following Oddo et al. [19], we treat these as chained sequences of atomic tasks rather than adding combinations.

Table 1: Distribution of task mentions across the Scope × Action matrix (49 papers, 444 mentions). Gray cells mark tasks hindered by bundling; empty cells are structurally impossible and omitted.

	Verify	Identify	Characterize	Quantify	Compare	Assess	
Element	6	52	0	1	0	17	76
Bundle	15	**80**	36	26	20	9	**186**
Global	1	51	5	7		**72**	136
Multi-view					43	3	46
	22	183	41	34	63	101	444

4.3 Representation Type as Instantiation

Each Scope × Action cell defines an abstract task. The representation type determines concrete semantics, because the same abstract task means different things for node-link diagrams, trail sets, and PCPs. Consider three examples.

Element × Verify asks whether a specific connection exists. In a node-link diagram, this means checking whether two nodes share an edge. In a trail set, it means checking whether any trajectory links two regions. In a PCP, it means checking whether a record with a particular value combination exists. All three are existence checks, but the data semantics differ fundamentally.

Bundle × Characterize asks what defines a bundle internally. In a node-link diagram, this means examining the directional mix of edges within a bundle. In a trail set, it means describing the mix of short and long trips that defines an area. In a PCP, it means determining whether the variables within a cluster correlate positively or negatively. This last instantiation is specific to PCPs and has no meaningful counterpart in the other two types.

Global × Assess asks whether bundling is faithful. In a node-link diagram, this concerns edge displacement from the original layout. In a trail set, the question is whether bundling shifted routes from their true spatial paths, a stronger requirement because the geometry carries real-world meaning. In a PCP, distortion occurs only in the inter-axis space, leaving values at the axes intact.

These examples illustrate why a single abstract taxonomy requires representation-specific interpretation. We flag representation-specific instantiations in our coded corpus rather than splitting them into separate tasks.

5 Corpus Analysis

Of the 102 papers, 49 contained at least one bundling-attributed task and form the basis of the taxonomy; the remainder either address bundling algorithmically without discussing tasks, or use bundling so peripherally that no task passed the attribution filter. By representation type, 322 mentions (73%) concern node-link diagrams, 78 (16%) trail sets, and 44 (11%) PCPs. The vast majority of mentions are explicit (92%); only 36 are implicit, a ratio that likely reflects both the prevalence of explicit task claims in the bundling literature and a limitation of LLM-based screening, which may undercount implicit mentions. Tab. 1 shows the per-cell distribution and we highlight the key patterns below.

Identify is the dominant action with 183 mentions (41%). At the bundle level, 31 of 49 papers claim their method "shows major flows," "makes clusters visible," or "surfaces connection groups," all instances of Bundle × Identify (locating which aggregates are present, not describing their internal composition). At the global level, 32 papers frame bundling as "providing an overview" or "reducing clutter to reveal structure" (Global × Identify).

Assess is the second-largest action (101 mentions, 23%), but its distribution is uneven. Global × Assess accounts for 72 mentions across 36 papers, covering distortion, parameter sensitivity, and the risk of false patterns. In contrast, Bundle × Assess (9 mentions, 8 papers) is notably sparse: few papers ask whether a perceived bundle is genuine or an artifact, despite this being a known risk [17].

The Bundle row is the richest scope with 186 mentions spread across all six actions, confirming that bundles function as first-class perceptual objects with their own task vocabulary. The Multi-view row is dominated by Compare (43 of 46 mentions, 31 papers), predominantly in trail-set contexts.

Two cells received zero mentions: Element × Characterize and Element × Compare. As discussed in Sec. 4.1, these are coherent tasks that bundling hinders rather than empty cells. Their absence from the corpus is consistent with the enabling/disabling duality and is reinforced by interaction techniques explicitly designed to recover element-level access [11, 31].

6 Discussion

Bundling enables and disables tasks simultaneously. The corpus confirms a duality not captured by prior frameworks: bundling *enables* tasks at the bundle and Global levels by introducing perceptual aggregates, while it *disables* tasks at the element level by merging individual elements. Prior task discussions [17] list capabilities that strong bundling effectively disables (e.g., precise attribute reading) while omitting tasks that bundling clearly enables (e.g., global faithfulness assessment, bundle-thickness comparison). A taxonomy that does not represent both directions will overstate what bundling supports. This dual structure makes the taxonomy usable in three ways: descriptively (classifying the tasks a paper targets), generatively (surfacing gaps), and evaluatively (matching quality metrics to intended tasks).

Evaluation gaps. The empty and sparse cells translate into concrete research questions for future work. *RQ1:* Can viewers reliably distinguish a genuine bundle from a bundling artifact? Bundle × Assess appears in only 8 papers despite this known risk. *RQ2:* How do bundling parameters affect which patterns are visible across views? Systematic Multi-view comparison of parameter settings is almost absent. *RQ3:* Do tasks transfer to trail and PCP bundling as readily as to node-link diagrams? The 73/16/11 split reflects a literature bias. *RQ4:* Which quality metrics predict performance on which tasks? Most papers report metrics without specifying the tasks.

Limitations. Our coding inherits the literature's biases: node-link papers dominate (73% of mentions), and PCP bundling is underrepresented. The LLM-assisted screening may also undercount implicit task mentions, as suggested by the 92%/8% explicit/implicit ratio. We also note the absence of a formal inter-rater agreement coefficient as a limitation of our coding procedure. The taxonomy is grounded in three established bundling settings; some variants (e.g., 3D bundling, bundling on non-Euclidean layouts) may require additional instantiations.

7 Conclusion

We present a task taxonomy for edge and trail bundling, derived from a systematic analysis of 102 papers and organized as a Scope × Action matrix. By extending the analysis beyond node-link diagrams to include trail sets and parallel coordinates, and by recognizing bundles as first-class perceptual objects, the taxonomy provides a structured vocabulary for reasoning about what can be done with bundled visualizations. The two-dimensional structure enables descriptive use (classifying existing tasks), generative use (identifying empty cells as evaluation gaps), and evaluative use (connecting bundling quality metrics to the tasks they serve). We invite the community to use the taxonomy when designing bundling evaluations, selecting quality metrics, and reporting task coverage in future bundling papers.

Supplemental Materials

The supplemental material provides the full coded corpus underlying our taxonomy. It contains the complete set of 444 bundling-attributed task mentions across the 49 papers, each coded by scope, action, and representation type, and annotated with the source paper's original terminology and the evidence attributing the task to bundling. Both a machine-readable table (one row per mention, with paper metadata) and a typeset version of the corpus are included, along with a cross-tabulation of mentions by scope, action, and representation type. We further provide the screening materials used during coding: the coding scheme with synonym lists mapping literature phrasings to our canonical labels, the bundling attribution filter with worked examples, and the scripts used to run the first-pass LLM screening over the paper collection. All supplemental material is available on OSF (osf.io/23r67).

References

[1] D. Archambault, G. Liotta, M. Nöllenburg, T. Piselli, A. Tappini, and M. Wallinger. Bundling-aware graph drawing. In *32nd International Symposium on Graph Drawing and Network Visualization, GD 2024, September 18-20, 2024, Vienna, Austria*, vol. 320 of *LIPIcs*, pp. 15:1–15:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi: 10.4230/LIPICS.GD.2024.15 2

[2] B. Bach, N. H. Riche, C. Hurter, K. Marriott, and T. Dwyer. Towards unambiguous edge bundling: Investigating confluent drawings for network visualization. *IEEE Trans. Vis. Comput. Graph.*, 23(1):541–550, 2017. doi: 10.1109/TVCG.2016.2598958 2

[3] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2376–2385, 2013. doi: 10.1109/TVCG.2013.124 1, 2

[4] W. Cui, H. Zhou, H. Qu, P. C. Wong, and X. Li. Geometry-based edge clustering for graph visualization. *IEEE Trans. Vis. Comput. Graph.*, 14(6):1277–1284, 2008. doi: 10.1109/TVCG.2008.135 2

[5] G. P. Ellis and A. J. Dix. The plot, the clutter, the sampling and its lens: occlusion measures for automatic clutter reduction. In *Proceedings of the working conference on Advanced visual interfaces, AVI 2006, Venezia, Italy, May 23-26, 2006*, pp. 266–269. ACM Press, 2006. doi: 10.1145/1133265.1133318 2

[6] G. P. Ellis and A. J. Dix. A taxonomy of clutter reduction for information visualisation. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1216–1223, 2007. doi: 10.1109/TVCG.2007.70535 2

[7] J. Heinrich and D. Weiskopf. Parallel coordinates for multidimensional data visualization: Basic concepts. *Comput. Sci. Eng.*, 17(3):70–76, 2015. doi: 10.1109/MCSE.2015.55 2

[8] N. Henry, J. Fekete, and M. J. McGuffin. Nodetrix: a hybrid visualization of social networks. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1302–1309, 2007. doi: 10.1109/TVCG.2007.70582 2

[9] D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Trans. Vis. Comput. Graph.*, 12(5):741–748, 2006. doi: 10.1109/TVCG.2006.147 1, 2

[10] D. Holten and J. J. van Wijk. Force-directed edge bundling for graph visualization. *Comput. Graph. Forum*, 28(3):983–990, 2009. doi: 10.1111/J.1467-8659.2009.01450.X 2

[11] C. Hurter, O. Ersoy, and A. Telea. MoleView: An attribute and structure-based semantic lens for large element-based plots. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2600–2609, 2011. doi: 10.1109/TVCG.2011.223 3, 4

[12] C. Hurter, O. Ersoy, and A. C. Telea. Graph bundling by kernel density estimation. *Comput. Graph. Forum*, 31(3pt1):865–874, 2012. doi: 10.1111/J.1467-8659.2012.03079.X 1, 2

[13] C. Hurter, B. Tissoires, and S. Conversy. Fromdady: Spreading aircraft trajectories across views to support iterative queries. *IEEE Trans. Vis. Comput. Graph.*, 15(6):1017–1024, 2009. doi: 10.1109/TVCG.2009.145 2

[14] A. Lambert, R. Bourqui, and D. Auber. Winding roads: Routing edges into bundles. *Comput. Graph. Forum*, 29(3):853–862, 2010. doi: 10.1111/J.1467-8659.2009.01700.X 2

[15] B. Lee, C. Plaisant, C. S. Parr, J. Fekete, and N. Henry. Task taxonomy for graph visualization. In E. Bertini, C. Plaisant, and G. Santucci, eds., *Proceedings of the 2006 AVI Workshop on BEyond time and errors: novel evaluation methods for information visualization, BELIV 2006, Venice, Italy, May 23, 2006*, pp. 1–5. ACM Press, 2006. doi: 10.1145/1168149.1168168 1, 2

[16] A. Lhuillier, C. Hurter, and A. C. Telea. FFTEB: edge bundling of huge graphs by the fast fourier transform. In *2017 IEEE Pacific Visualization Symposium, PacificVis 2017, Seoul, South Korea, April 18-21, 2017*, pp. 190–199. IEEE Computer Society, 2017. doi: 10.1109/PACIFICVIS.2017.8031594 2

[17] A. Lhuillier, C. Hurter, and A. C. Telea. State of the art in edge and trail bundling techniques. *Comput. Graph. Forum*, 36(3):619–645, 2017. doi: 10.1111/CGF.13213 1, 2, 4

[18] K. T. McDonnell and K. Mueller. Illustrative parallel coordinates. *Comput. Graph. Forum*, 27(3):1031–1038, 2008. doi: 10.1111/J.1467-8659.2008.01239.X 2

[19] M. I. Oddo, R. Smith, S. Kobourov, and T. Munzner. Visualization tasks for unlabeled graphs. *IEEE Trans. Vis. Comput. Graph.*, 2026. 2, 3

[20] G. Palmas and T. Weinkauf. Space bundling for continuous parallel coordinates. In E. Bertini, N. Elmqvist, and T. Wischgoll, eds., *18th Eurographics Conference on Visualization, EuroVis 2016 - Short Papers, Groningen, The Netherlands, June 6-10, 2016*, pp. 61–65. Eurographics Association, 2016. doi: 10.2312/EUROVISSHORT.20161162 2

[21] R. Saga. Quantitative evaluation for edge bundling based on structural aesthetics. In *18th Eurographics Conference on Visualization, EuroVis 2016 - Posters, Groningen, The Netherlands, June 6-10, 2016*, pp. 17–19. Eurographics Association, 2016. doi: 10.2312/EURP.20161131 1, 2

[22] B. Saket, P. Simonetto, S. G. Kobourov, and K. Börner. Node, node-link, and node-link-group diagrams: An evaluation. *IEEE Trans. Vis. Comput. Graph.*, 20(12):2231–2240, 2014. doi: 10.1109/TVCG.2014.2346422 2

[23] R. H. Tai, L. R. Bentley, X. Xia, J. M. Sitt, S. C. Fankhauser, A. M. Chicas-Mosier, and B. G. Monteith. An examination of the use of large language models to aid analysis of textual data. *International Journal of Qualitative Methods*, 23, 2024. doi: 10.1177/16094069241231168 2

[24] M. van der Zwan, V. Codreanu, and A. C. Telea. Cubu: Universal real-time bundling for large graphs. *IEEE Trans. Vis. Comput. Graph.*, 22(12):2550–2563, 2016. doi: 10.1109/TVCG.2016.2515611 1, 2

[25] C. Vehlow, F. Beck, and D. Weiskopf. Visualizing group structures in graphs: A survey. *Comput. Graph. Forum*, 36(6):201–225, 2017. doi: 10.1111/CGF.12872 2

[26] M. Wallinger, O. Akbulut, K. A. Rufai, H. C. Purchase, and D. Archambault. How do people perceive bundling? an experiment. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, CHI 2025, Yokohama, Japan, 26 April 2025- 1 May 2025*, pp. 1113:1–1113:14. ACM, 2025. doi: 10.1145/3706598.3713444 1, 2

[27] M. Wallinger, D. Archambault, D. Auber, M. Nöllenburg, and J. Peltonen. Edge-path bundling: A less ambiguous edge bundling approach. *IEEE Trans. Vis. Comput. Graph.*, 28(1):313–323, 2022. doi: 10.1109/TVCG.2021.3114795 2

[28] M. Wallinger, D. Archambault, D. Auber, M. Nöllenburg, and J. Peltonen. Faster edge-path bundling through graph spanners. *Comput. Graph. Forum*, 42(6):e14789, 2023. doi: 10.1111/CGF.14789 2

[29] M. Wallinger, J. Miller, A. Maftei, and S. G. Kobourov. Quantitative metrics for edge bundling of network visualizations. *Computer Graphics Forum*, n/a(n/a):e70443. doi: 10.1111/cgf.70443 1, 2

[30] M. Wallinger, T. Piselli, A. Tappini, D. Archambault, G. Liotta, and M. Nöllenburg. Bundling-aware graph drawing revisited. *IEEE Trans. Vis. Comput. Graph.*, 31(12):10828–10839, 2025. doi: 10.1109/TVCG.2025.3616583 2

[31] N. Wong, S. Carpendale, and S. Greenberg. Edgelens: An interactive method for managing edge congestion in graphs. In *IEEE Symposium on Information Visualization (InfoVis)*, pp. 51–58, 2003. doi: 10.1109/INFVIS.2003.1249008 3, 4