

Secondary Publication



Förster, Henry; Klesen, Felix; Dwyer, Tim; u. a.

GraphTrials : Visual Proofs of Graph Properties

Date of secondary publication: 19.08.2026

Version of Record (Published Version), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-116874x

Primary publication

Förster, Henry; Klesen, Felix; Dwyer, Tim; u. a. (2025): GraphTrials: Visual Proofs of Graph Properties, in: IEEE transactions on visualization and computer graphics: TVCG, New York, NY: Institute of Electrical and Electronics Engineers (IEEE), vol. 31, no. 10, pp. 8767–8781, doi: 10.1109/tvcg.2025.3577533

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

GRAPHTRIALS: Visual Proofs of Graph Properties

Henry Förster¹, Felix Klesen², Tim Dwyer³, Peter Eades⁴, Seok-Hee Hong⁴, Stephen Kobourov,
Giuseppe Liotta⁵, *Senior Member, IEEE*, Kazuo Misue⁶, *Member, IEEE*, Fabrizio Montecchiani⁷,
Alexander Pastukhov⁸, and Falk Schreiber⁹

Abstract—Graph and network visualization supports exploration, analysis and communication of relational data arising in many domains: from biological and social networks, to transportation and powergrid systems. With the arrival of AI-based question-answering tools, issues of trustworthiness and explainability of generated answers motivate a significant new role for visualization. In the context of graphs, we see the need for visualizations that can convince a critical audience that an assertion (e. g., from an AI) about the graph under analysis is valid. The requirements for such representations that convey precisely one specific graph property are quite different from standard network visualization criteria which optimize general aesthetics and readability. In this paper, we aim to provide a comprehensive introduction to visual proofs of graph properties and a foundation for further research in the area. We present a framework that defines what it means to visually prove a graph property. In the process, we introduce the notion of a visual certificate, that is, a specialized faithful graph visualization that leverages the viewer's perception, in particular, pre-attentive processing (e. g., via pop-out effects), to verify a given assertion about the represented graph. We also discuss the relationships between visual complexity, cognitive load and complexity theory, and propose a classification based on visual proof complexity. Then, we provide further examples of visual certificates for problems in different visual proof complexity classes. Finally, we conclude the paper with a discussion of the limitations of our model and some open problems.

Received 31 January 2025; revised 9 May 2025; accepted 21 May 2025. Date of publication 6 June 2025; date of current version 5 September 2025. Recommended for acceptance by Y. Hu. A preliminary version of this article was presented at The 32nd International Symposium on Graph Drawing and Network Visualization (GD2024) [29]. The program committee chairs of GD2024 invited us to submit a revised version of this conference paper to TVCG. (*Corresponding author: Henry Förster.*)

Henry Förster and Stephen Kobourov are with the School of Computation, Information and Technology, Technical University of Munich, Heilbronn 74076, Germany (e-mail: henry.foerster@tum.de; stephen.kobourov@tum.de).

Felix Klesen is with the Institute of Computer Science, University of Würzburg, 97070 Würzburg, Germany (e-mail: felix.klesen@uni-wuerzburg.de).

Tim Dwyer is with the Faculty of Information Technology, Monash University, Clayton, VIC 3800, Australia (e-mail: tim.dwyer@monash.edu).

Peter Eades and Seok-Hee Hong are with the School of Computer Science, University of Sydney, Camperdown, NSW 2050, Australia (e-mail: peter.eades@sydney.edu.au; seokhee.hong@sydney.edu.au).

Giuseppe Liotta and Fabrizio Montecchiani are with the Department of Engineering, University of Perugia, 06123 Perugia, Italy (e-mail: giuseppe.liotta@unipg.it; fabrizio.montecchiani@unipg.it).

Kazuo Misue is with the Department of Computer Science, University of Tsukuba, Ibaraki 305-8577, Japan (e-mail: misue@cs.tsukuba.ac.jp).

Alexander Pastukhov is with the Department of Psychology, University of Bamberg, 96047 Bamberg, Germany (e-mail: alexander.pastukhov@uni-bamberg.de).

Falk Schreiber is with the Department of Computer Science, University of Konstanz, 78464 Konstanz, Germany (e-mail: falk.schreiber@uni-konstanz.de). Digital Object Identifier 10.1109/TVCG.2025.3577533

Index Terms—Graph visualization, theory of visualization, visual proof.

I. INTRODUCTION

WHILE state-of-the-art graph and network visualization techniques do a reasonable job of untangling graphs to convey meaning and support free-form exploration, there are certain application scenarios where these algorithms fall short. Namely, we focus on applications where it is necessary to *convince* a (possibly non-expert) audience that a particular graph has some structural property. We emphasize that this kind of application scenario differs significantly from the traditional usage of visualization to generate new knowledge. Namely, existing graph and network visualization techniques have sought mainly to represent all aspects of a graph or network structure as faithfully as possible such that a user can explore the visualization, identify structures, and gain insights about the underlying data. These traditional visualization techniques can be sufficient for journalists and other communicators to support a narrative in print or on-line media [12] by showing only selected views of graphs.

However, novel approaches are required in our scenario in which a specific property of the data is to be conveyed in an adversarial setting where the validity of the evidence presented may be questioned (see also the defense lawyer role in Fig. 1 which may, e. g., represent doubts of the audience). For example, the investigative activity of the Italian Revenue Agency (IRA) exploits the visual analysis of social networks whose nodes are the actors of potential fraudulent activities and whose edges represent financial/legal transactions between the actors. The investigators of IRA who suspect a group of persons or a single individual/company of tax evasion submit a case to the Italian financial Police for possible prosecution, which also implies showing some structural properties of the network beyond reasonable doubts. See, e. g., [26], [27], [28] for references about the use of visual analytics in the context of contrasting tax evasion in Italy.¹ Below we describe introductory examples.

Example 1. A network administrator discovers that two critical parts of the infrastructure would not be able to communicate with each other if a particular switch fails. To increase the robustness of the network, new hardware is needed. They have to convince the manager, who has no background in network security, to fund new hardware.

Example 2. In a legal court case, the prosecution discovers that money acquired in black market sales was laundered through a laundromat chain, as evidenced by money provably transferred

¹ One of the authors has also been approached by the Australian Security and Investments Commission (a governmental regulator for stock exchange) inquiring about visualizations to convince a court about illegal trades.

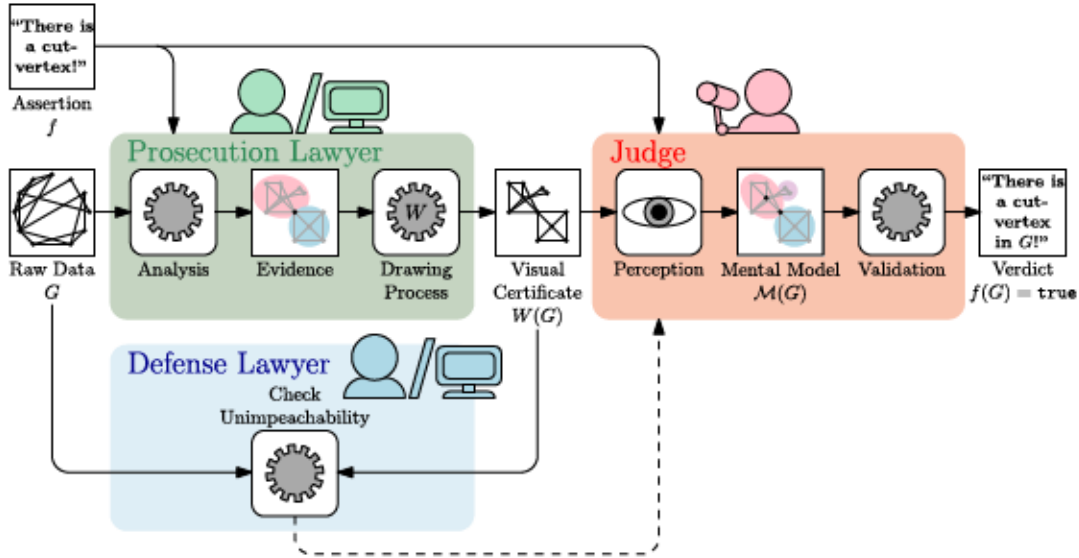


Fig. 1. Our proposed model GRAPHTRIALS identifies key processes for visually proving an assertion about a given graph in an adversarial setting. The *prosecution lawyer*, a software or a human (assisted by software), intends to highlight evidence for a graph being accused of satisfying an assertion by usage of a visual certificate drawing. To convince the *judge*, the human audience of the drawing, the visual certificate must guide the judge's perception to form a mental model which makes the assertion easy to validate. Further, the visual certificate must be unimpeachable as a *defense lawyer*, yet another piece of software or a human adversary, and checks for reasons to doubt the validity of the certificate which may influence the judge's verdict.

via a complicated network from the dealers to the laundromats. The prosecution has to convince the judge that all suspects belong to the criminal syndicate.

Example 3. A new artificial intelligence-based heuristic (AI-based, for short) is able to efficiently decide if a given graph is Hamiltonian, i. e., to test if it contains a cycle traversing all its vertices exactly once². However, false positives must be filtered out. A human operator needs to perform this task as there is no efficient algorithm. To facilitate this, the new version of the algorithm should also create a visualization of the graph making the Hamiltonian cycle obvious to the operator.

Such scenarios have key differences to standard motivations for graph visualization. Typical graph visualization techniques (node-link layout algorithms [23], [69], matrix ordering approaches [7] and mixed approaches which either include features of different paradigms [3], [4], [32], [40] or show different visualizations side-by-side [14], [39]) usually seek a representation showing as many graph properties as possible simultaneously (by trading off aesthetic and readability criteria [11], [8], [9], [22]). However, for the scenarios above it is better to focus on showing optimally and faithfully just one specific property, i. e., we want a *visual proof* for that property.

More precisely, a visual proof is a proof given by the use of a graphical or visual representation called a *visual certificate*. A good visual proof should be clear and concise, conveying the main idea in an easy-to-understand way. It should be able to effectively communicate the desired message without being overly complex or cluttered. Additionally, the visual certificate should be aesthetically pleasing and easy to interpret. Somehow,

it should be able to provide evidence to support the argument being made. Thus, a good visual certificate should be accurate, concise, and free of errors or mistakes.

In fact, visual proofs are already used in mathematics and other areas such as logic, graph theory, computer science, and physics [59], [81]; visual proofs are often easier to understand than algebraic proofs, as they are less abstract and easier to follow. Accessible proofs are often considered more beautiful by mathematicians; e. g., Appel and Haken employed a computer-assisted proof of the long-open four-color theorem in 1976 [5]. This new type of proof sparked philosophical debates [73] and while the theorem is broadly accepted as proved³, researchers still desire a more elegant proof [2]. Thus, we expect that visual proofs are appealing and even more convincing to experts also in fields other than mathematics.

Visual proofs can also convey properties to non-expert users or explain correctness of AI-generated solutions. As powerful chat-based interfaces are capable of generating plausible sounding – but difficult to verify – explanations of complex phenomena, we believe that there is a requirement to understand what makes a graph representation a proper visual certificate.

a) Contribution: We introduce a model identifying important steps and their interactions in a visual proof of a graph property. Based on this model, we formalize the concept of visual certificates and give requirements for a visualization to qualify as such. We then provide further examples of visual proofs for widely used graph properties that showcase our model. Finally, we discuss limitations of our model and, as a result, identify open research questions that should be answered to better understand visual proofs and make them algorithmically usable.

² Note that neural network approaches for NP-hard problems have been described, e. g., in [10]. In addition, the need for visualizations in the context of explainable deep learning has been described, e. g., in [21].

³ According to the *Oxford English Dictionary*, it is yet to be proven as a “mathematical theorem” [63].




Fig. 2. Three layouts generated with yEd [84]. (a) and (b) Circular and organic layouts generated with standard settings, resp. (c) Manually created layout highlighting the cut-vertex.

II. FIRST EXAMPLES OF VISUAL PROOFS

A. Example 1: The Graph Contains a Cut-Vertex

First, we revisit the situation in Example 1. In this communication network there are two distinct parts such that all connections between them traverse a single switch. This corresponds to the graph underlying the network containing a *cut-vertex*, whose removal separates the remainder of the graph into at least two distinct components. Hence, in order to convince the manager, the network administrator has to point out that the graph underlying the network can be separated by the removal of the vertex corresponding to the switch. So, they first layout the graph using a circular layout, which is a wide-spread all-purpose visualization style [67], and point the manager to the fact that the vertex highlighted in red is a cut-vertex; see Fig. 2(a).⁴

Unfortunately, the circular layout does a poor job at highlighting the cut-vertex. While it is evident to the manager that there are a top and a bottom component connected by some edges, they explain that they are not sure if all connections between both components use the suggested cut-vertex or not. Hence, the network administrator prepares a second drawing using a force-directed organic layout where the cut-vertex is clearly visible; see Fig. 2(b). However, the engineer who designed the network becomes defensive and claims that there could be another edge hidden behind the alleged cut-vertex. This argument can be easily disproven by the network administrator as they move the cut-vertex down, obtaining the drawing in Fig. 2(c). Presented with this new line of evidence, the engineer stops arguing and the manager agrees that the network has to be made more robust.

a) Discussion: This example illustrates how standard layout techniques may be unable to highlight even simple properties. In the circular layout, it is not easy to verify even when the cut-vertex is highlighted; see Fig. 2(a). This is due to the Gestalt principle of grouping [75], [76]. Here, the initial perception is guided by continuity and closure of node positions, leading to the perception of a single circular component. As a second step, an observer may see two separate components with edges biasing perception due to connectedness grouping. Thus, the observer has to analyze the entire graph, going node-by-node, to negate the automatic perceptual grouping induced by the layout to verify that there is a cut-vertex. The issue with the second illustration in Fig. 2(b) is of different nature. Namely, the force-directed layout does a much better job at highlighting the cut-vertex. In fact, the observer discovers two dense salient features which are




Fig. 3. (a)–(b) Organic layout generated with standard settings by yEd [84], with a spanning tree highlighted in (b). (c) A manually created layout highlighting the spanning tree.

the two components separated by the cut-vertex and immediately notes that they are connected at a single vertex. Nevertheless, if there is an overlapping edge behind the cut-vertex, the drawing may look the same, challenging the human observer to identify that the vertex is not a cut-vertex. The drawing in Fig. 2(c) avoids this problem by explicitly highlighting the cut-vertex via pre-attentively perceptible patterns (i. e., pop-out effects) [76]. The singular goal of highlighting the cut-vertex is achieved at the cost of traditionally accepted aesthetic metrics [65], as — compared to the circular and force-directed layouts — the general layout is unbalanced, with many crossings and poor resolution; see Table I. Thus, visual certificates may not be useful in traditional exploratory applications, instead they focus on highlighting a specific property.

We remark that a cut-vertex proves non-2-connectivity and a similar approach can be used if we want to highlight that the graph contains a *bridge*, i.e., an edge whose removal separates the graph. Namely, we can use the same visual proof strategy, placing both endpoints of the bridge below the two components separated by its removal.

B. Example 2: The Graph is Connected

In Example 2, to convince the judge, the prosecution lawyer decides to visualize the network of criminals induced by the connections of provable money transfers. The prosecution lawyer draws it with a force-directed approach; see Fig. 3(a). While Fig. 3(a) shows that there are many connections in the graph, it does not emphasize that there is only a single connected component. Hence, the defense lawyer argues that the component containing their client may have been drawn on top of the component with all the convicted criminals. Hence, the prosecution lawyer has to improve their visual proof. To do so, they include a highlighted *spanning tree* that shows that every vertex can be reached from every other vertex; see Fig. 3(b). Although the defense lawyer now has to admit that there is a smaller portion of the drawing to check, i. e., the highlighted edges, their argument stays more or less the same: that there are still crossings between edges of the spanning tree, which may be due to two different highlighted components drawn on top of each other. Thus, the prosecution lawyer creates a third drawing in which the spanning tree is crossing-free; see Fig. 3(c). Here the spanning tree is rooted at the central vertex and vertices are drawn on concentric circles depending on distance from the root. Given this visualization, the defense rests, and the judge decides quickly that indeed all members of the network are affiliated.

a) Discussion: While in Example 1 we have seen that the drawing style of the entire graph can be important to visually prove a property, here we added another dimension. Namely, a

⁴ Unless specified otherwise, the layouts of all visualizations in this paper have been created by the authors.

TABLE I

SELECTED AESTHETIC METRICS OF THE NODE-LINK DRAWINGS IN SECTION II. RED NUMBERS IN PARENTHESES GIVE THE CORRESPONDING VALUES FOR SUBGRAPHS HIGHLIGHTED IN RED IN THE CORRESPONDING FIGURE. FOR STRESS AND CROSSING NUMBER (RED BACKGROUND) LOWER NUMBERS ARE BETTER, OTHERWISE (GREEN BACKGROUND) HIGHER NUMBERS ARE BETTER

	Fig. 2: Cut-Vertex			Fig. 3: Spanning Tree		Fig. 4: Hamilton Cycle	
	(a) Circular	(b) Force-Dir.	(c) Proof	(a) – (b) Force-Dir.	(c) Proof	(a) Force-Dir.	(b) Proof
STRESS [46]	132.1	9.2	36.7	13.6 (31.3)	32.5 (19.6)	9.3 (51.9)	45.5 (12.6)
CROSSING NUMBER [62]	81	19	26	28 (8)	63 (0)	29 (4)	45 (0)
JACCARD INDEX [1]	.274	.332	.283	.330 (.132)	.285 (.160)	.407 (.167)	.339 (.232)
EDGE-LENGTH RATIO [49], [50]	.174	.333	.127	.360 (.404)	.149 (.382)	.401 (.649)	.185 (.542)
NODE RESOLUTION [1]	.174	.104	.100	.181 (.181)	.119 (.119)	.164 (.164)	.184 (.184)
ASPECT RATIO [62]	.985	.774	.966	.884 (.884)	.796 (.796)	.871 (.871)	.987 (.987)
CROSSING RESOLUTION [24], [25]	20.0	27.1	7.2	37.1 (37.1)	20.5 (N/A)	31.4 (48.2)	22.5 (N/A)
ANGULAR RESOLUTION [62]	10.0	.20	.56	1.68 (20.0)	4.44 (21.4)	0.40 (15.93)	12.08 (141.0)




Fig. 4. Four layouts of a graph with a Hamiltonian cycle (red).

subgraph is explicitly color-highlighted for pre-attentive perception. In addition, the drawing of this subgraph was very important in creating a convincing argument. In Fig. 3(b) the drawing of the spanning tree is not very readable. Thus, even with the attention drawn to this portion of the drawing, it remains time consuming to check that a single tree connects all vertices. But when the tree is laid out in a concise and readable fashion as in Fig. 3(c), it is quite evident that it spans all the vertices, as the highlighted edges induce automatic grouping via similarity [75] and act as guidance for attention spread [43]. Similar to Example 1, while the quality of the drawing of the spanning tree is improved, the drawing of the rest of the graph does not measure well on the usual metrics; see Table I.

C. Example 3: The Graph Has a Hamiltonian Cycle

We may train an AI to produce a good node-link drawing for Example 3. As in Section II-B, we observe that the quality of the drawing of the evidence relevant to the property under consideration is more important than the drawing of the full graph (Fig. 4(a)), thus we select a circular layout with the Hamiltonian cycle forming the outer face (Fig. 4(b)). However, the human operator needs to check that all edges of the highlighted outer cycle are indeed present which can become increasingly difficult for larger graphs where resolution may become problematic. We can improve upon these issues by instead using an adjacency matrix representation. While an arbitrary permutation (see Fig. 4(c)) does not provide any insights, an appropriate sorting of rows and columns makes the cycle composed of three components: one red diagonal and two red-cells (top-right and bottom-left); see Fig. 4(d).

a) *Discussion:* We observed that different visualization paradigms may perform better or worse for visually proving a property. While the node-link drawing in Fig. 4(b) already highlights the cycle well, the adjacency matrix representation in Fig. 4(d) composes the Hamiltonian cycle in three components. The perception of the red diagonal is facilitated by figure-ground separation via connectedness and similarity [75],

and the two corner cells stand out due to both color difference and symmetry [82]. The particular advantage of this representation is that the used visual cues scale nicely even for very large matrices (up to the pixel resolution of the screen) [82]. Thus, an important criterion for judging the quality of a visual proof should be the workload required by the observer to evaluate the correctness. As checking for Hamiltonicity is a difficult task with an all-purpose visualization (see also Section III-C), both visualizations should be regarded as valid visual certificates albeit of different quality.

III. RELATED THEORIES, FRAMEWORKS AND MODELS

A. Certifying Algorithms

The concept of visual certificates is related to *certifying algorithms* popularized by McConnell et al. [53], which seek to provide short and easy-to-check *certificates* for the correctness of an algorithm. Let $f : X \rightarrow Y$ be a computable, surjective function for input set X and output set Y and let W be a set of *witnesses*. Intuitively speaking, a witness describes a simple proof certifying that the output y of an algorithm for f on input x satisfies $f(x) = y$. The validity of a witness for a certain combination of inputs and outputs is assessed via the witness predicate $\mathcal{W} : X \times Y \times W \rightarrow \{\text{true}, \text{false}\}$ that fulfills:

- 1) *Witness property:* Given $(x, y, w) \in X \times Y \times W$, it holds $f(x) = y \Leftrightarrow \mathcal{W}(x, y, w) = \text{true}$.
- 2) *Checkability:* Given $(x, y, w) \in X \times Y \times W$, it is trivial to determine $\mathcal{W}(x, y, w)$.
- 3) *Simplicity:* $\mathcal{W}(x, y, w) \Rightarrow f(x) = y$ has a simple proof.

An algorithm for f is now called *certifying algorithm* if for any input $x \in X$ it computes the output $y = f(x) \in Y$ and a witness $w \in W$ such that $\mathcal{W}(x, y, w) = \text{true}$. It is worth noting that Properties 2 and 3 of the witness predicate are vaguely formulated. McConnell et al. [53] suggest that Property 2 can be formalized by requiring that there must be a decision algorithm for $\mathcal{W}$ that runs in a certain time (such an algorithm is called a *checker*). On the other hand, they emphasize that Property 3 is intentionally left subjective as it relies on what is considered common knowledge.

As an example, reconsider Example 2. Here, our function f takes a graph G as input and should output either *true* if the graph is connected or *false* otherwise. A witness for a positive instance would be a spanning tree T as discussed in Section II-B. Given T , we can easily see that G is connected and we can determine this by checking that every vertex is part of T ; i. e., Property 3 holds. In fact, this checking can be done efficiently by running a BFS on T . That is, the time for checking the correctness of the certificate depends only on the number of vertices, not on the number of edges; i. e., Property 2 holds. On

the other hand, if G is not connected, the certifying algorithm for f can provide a partition of the vertices into two disconnected sets V_1 and V_2 as a witness. We may check that we cannot reach any vertex in V_2 if we run a BFS starting from a vertex in V_1 to establish that f was computed correctly; i. e., Property 3 holds. If it is the case that $|V_1| \leq |V_2|$ we end up considering at most half of the vertices in this process, hence we may argue that Property 2 holds. Since we found a good witness for positive and negative instances, we conclude that Property 1 holds.

While Example 1 also admits a certifying algorithm, for Example 3, two issues emerge. (1) Computing a Hamiltonian cycle is NP-complete, hence, the algorithm for f may not run in polynomial time. However, even if we ignore computational complexity, a problem remains: (2) for a negative instance no easy witness is known, i. e., we do not even know how to guarantee Property 1. This latter appears to be a general issue for NP-complete assertions since a witness for a negative instance would be a short certificate for a co-NP-complete problem.

B. Perception

Visual proofs are concerned with a design of visual evidence for an existence of a specific property, such as the presence of a cut-vertex, a Hamiltonian cycle, and so on. In principle, a proof for such a property can be reduced to a program that returns a binary outcome, affirming or rejecting the claim. This may be sufficient for specialists who are familiar with the property itself, understand and trust the algorithm behind the code, and trust that the code is valid. However, such evidence may not be convincing to a non-specialist (a judge, a stockholder, etc.), particularly because the proof itself will be just one piece of evidence among many. Prior research shows that in such cases, presenting evidence *per se* is not enough, as information can be discounted as confusing, unimportant, or, given the wrong context, even misleading [72], as the accessibility and clarity of evidence could be as important as evidence itself [37].

Due to the diversity of graph properties there can be no general solution. Visual proof design might be guided by the principle of optimizing the *data-ink* ratio [72]. Thus, instead of optimizing overall aesthetic metrics [65], one should minimize the required number of *visual queries*, i. e., attention orientation, driving eye movements, and pattern/object recognition [77].

The human visual perception system consists of three stages: (1) rapid parallel processing involving billions of neurons, e. g., extraction of orientation, texture, color, and motion features; (2) slower processing than Stage 1, e. g., detection of 2D patterns, contours and regions; (3) slow serial processing, involving both working and long-term memory, e. g., object identification [76]. As in Stage 1 the entire visual field is processed quickly in parallel, information that can be captured in this stage can be easily distinguished. Thus, pre-attentive (pop-out) patterns such as color, size, orientation, shapes, etc. should be utilized.

In other words, a good visual proof must ensure that a focal piece of evidence is a visual “pop-out” feature that automatically attracts viewer attention and that the visual layout is parsed and grouped into patterns that express the evidence. In case of the former, studies on visual search provide a comprehensive list of useful pop-out features such as color, size, contrast, or location [82]. Regarding the latter, one can rely on a large body of literature on principles of perceptual organization, commonly known as Gestalt principles [75]. However, yet another constraint is placed by our working memory that limits the number of

nodes, edges, and components that can realistically be assessed at any single time [52].

The examples above illustrate the importance of this approach for visual proofs. For instance, consider the visual evidence for the existence of a cut-vertex in Fig. 2(a). While it uses color to attract the viewer’s attention to the cut-vertex and spatial arrangement to visually separate the two components, it still leads to an excessive number of visual queries, requiring multiple scans of individual vertices to ensure that they are connected only to the cut-vertex and the nodes within the component. In turn, there is a memory bottleneck that is likely to prevent a viewer from being completely certain about the validity of the proof. In contrast, in Fig. 2(c) the graph layout groups the entire evidence in just three components and clearly shows lack of inter-component edges, so that very few visual queries are required to confirm the vertex is indeed a sole connector between the components. In short, although there cannot be a single one-size-fits-all approach for constructing visual proofs, their critical role in aiding the cognition of the viewer means they should be built based on principles of perceptual organization and around the limitations of attention and memory [77].

C. Computational Complexity

To evaluate the amount of the cognitive workload, we will apply concepts from complexity theory [6], [30]. It is also worth mentioning that the examples discussed so far differ in terms of their computational complexity. Namely, all cut-vertices of a graph and a spanning-tree can be found in time $O(n + m)$ based on BFS traversals where n is the number of vertices and m the number of edges while determining a Hamiltonian cycle is NP-complete [30]. Thus, in Example 3, we have visually proven an algorithmically difficult to solve problem.

However, there may be graph properties that cannot be visually proven. We first have to discuss how a human observer interacts with a visual certificate. In Example 1, the human observer identified two connected components and then saw that they can be separated by the removal of their shared vertex. Such a procedure could be seen as an $O(1)$ time algorithm, where the observer determined that there is only a single point where both components touch. Similarly, in Examples 2 and 3, the observer may have checked for every vertex if it was part of the highlighted structure. Even if they were to check this for every vertex one at a time, the resulting algorithm would still run in linear time. Hence, an observer is actually performing a *deterministic* validation algorithm for establishing that a certificate is correct.

Now, consider the complementary question to Example 3, i. e., we want to determine whether a graph does *not* contain a Hamiltonian cycle. This is a *co-NP-complete* problem as it is the complement to an NP-complete problem. For co-NP-complete problems it is likely that there are instances for which there is no certificate that can be checked in polynomial time [6], i. e., if we assume that a human observer deterministically analyzes a visualization (as could be recreated by computer vision), we have to assume that we cannot visually prove a co-NP-complete problem in general.

D. Related Visualization Models

Aside from graph visualizations, the concept of visually enhancing a proof is wide-spread. In mathematics, visual proofs

for theorems have been used since ancient times [81] and there is a plethora of examples [59]. The question if such proofs can be regarded as such also has been discussed philosophically [13]. Also in computer science, visualizations are heavily used to convey knowledge, e. g., while not necessarily proving, an interactive sequential art by Bret Victor [74] beautifully explained an algorithm from a Nature paper [79].

Overall, there is a trend of increasingly sophisticated models considering a holistic integration of visualization into the sense-making process, typically with the goal of informing the design of interactive systems for data exploration. Early models considered a linear pipeline, from data, via various transformations, to a visual display [15]. Visual analytics seeks to apply visualization to support the entire human sense-making loop [64]. More recent models aim to connect sense-making from interactive data visualization, via hypothesis formation and testing, to knowledge generation [66]. An underlying theme across most of this work is the role of computational guidance in the analytics process, and how algorithms can support the various loops in the sense-making process [18]. By contrast, we consider a different model to conceptualize the role of algorithms, and AI, in supporting data (specifically network data) understanding. Our model for visual proofs (Fig. 1) does not seek to replace the traditional sense-making/knowledge-generation loop, but to support humans in situations where the result of a complex algorithm or property needs to be explained and justified.

There are also models related to ours from information visualization research. Song et al. [68] considered a problem that may be seen as a complementary question to the one studied in this paper: They investigated how computer vision can understand network visualizations optimized for human users. Wickham et al. [80] proposed a two-phase procedure to convince a human observer that a data set contains statistically significant difference from randomly generated data. The human observer is first exposed to several randomly generated data sets (similar to a Rorschach test) before being exposed to a line-up consisting of the real data set and a couple randomly generated data sets. The first phase primes the human viewer for statistically insignificant variations so that, in the second phase, statistically significant differences clearly pop out from the noise. Another related model are *Gragnostics*, which are ten features suggested by Gove [34], [35], that are fast to compute and provide a quantification of structural graph properties. In contrast to our model that aims to prove structural properties of graphs, Gragnostics provides the human user with a first impression of the structure of the graph at hand which may be helpful for initiating a thorough investigation. Finally, our model may also be seen as a visual communication of structural graph properties. Visual communication has been investigated in other settings for several decades, see e. g. [70], [71].

IV. THE GRAPHTRIALS MODEL

We are now ready to discuss our formalization of *visual proofs*. For this, we first abstractly outline the process of visually proving properties of graphs in an adversarial setting using a model that we call GRAPHTRIALS; see also Fig. 1. The model includes three distinct roles that have already appeared in our discussion of Example 2 in Section II-B: The *prosecution lawyer* must convince the judge that a certain assertion regarding a graph is true, the *defense lawyer* may raise doubts about the validity of the prosecution lawyer's claims, and the *judge* will determine

the truth of the assertion. The roles *prosecution lawyer*, *judge* and *defense lawyer* are to be seen as abstract descriptions of the different actors in the process; e. g., in Example 1 and 3, the prosecution lawyers were the network administrator and the AI-based algorithm, respectively. The latter example further indicates that not all roles have to be assigned to a human. In fact, we only require that the judge corresponds to the human audience of the visual certificate whereas each lawyer may be either human, software or a human assisted by software. Moreover, as we have seen in Section II-A, it can also occur that a critical audience can act as both the judge and defense lawyer roles simultaneously.

To convince the judge of a valid assertion f for the input graph G , the *prosecution lawyer* draws a visual certificate $W(G)$. To do so, they first analyze the raw data G to reveal evidence that proves the assertion f . The evidence is then embedded in $W(G)$: a visual representation of G that in some way emphasizes the evidence. Note that in the scope of our model we treat the analysis of the raw data and extraction of the evidence as a closed box, i. e., we may assume that the prosecution lawyer already knows that the assertion f is true for the input graph G and may also be given the evidence as input. This allows us to *efficiently* visually prove algorithmically difficult assertions (such as the existence of a Hamiltonian cycle as in Section II-C) and to ignore how the evidence is gathered (either algorithmically or by human interaction) in our model. The latter aspect also provides the possibility to separate the evidence gathering from the visualization process W , i. e., W could be a reusable program that embeds the evidence according to a specification⁵.

The *defense lawyer* checks the *unimpeachability* of $W(G)$ as a visual representation of G certifying $f(G)$. Thus, they may question whether the graph represented in the visualization actually corresponds to the input and they may also raise concerns if $W(G)$ is not distinguishable from a slightly different non-certificate (e. g., in Section II-A we encountered the case where an edge may have been hidden making it invisible to the judge's perception).

The *judge*, the human audience of the visual certificate $W(G)$, will validate the claim $f(G)$ using $W(G)$. In this step, the visual certificate $W(G)$ must guide the judge's perception so that they are able to form a mental model $\mathcal{M}(G)$ that facilitates confirmation of the validity of the assertion $f(G)$. For instance, the guidance can be formed by a suitable choice of topology which leads the judge to identify clusters of the graph as distinct salient features (as in Section II-A) or by adding additional features such as color to draw attention to certain parts of the graph (as in Section II-B). We discuss the judge's mental model in Section IV-C.

It is noteworthy that aside from the input graph and the verdict of the judge, the only information shared by all three roles is the *visual certificate* $W(G)$. In particular, it is the only medium that can be used by the prosecution lawyer to communicate the gathered evidence to the judge, i. e., the evidence is hidden information only accessible by the prosecution lawyer. Similarly, the judge is not communicating its mental model $\mathcal{M}(G)$ to the prosecution or defense lawyer, yet as we discussed above both roles might want to *estimate* what the mental model will look like. Furthermore, the nature of the mental model plays an important role in the validation step performed by the judge.

⁵ To this end, notice that the examples in Sections II-C and V-A both use visual certificates that highlight cycles.

Namely, the cognitive load put on the judge in this step depends hugely on how *complex* $\mathcal{M}(G)$ is. Finally, the defense lawyer's checking for unimpeachability is a process that is independent of the judge and prosecution lawyer and for a *faithful and readable* visual certificate we demand that there is no reason for the defense lawyer to raise doubts to the judge. As a result, there are several properties that we require from a visualization in order to call it a visual certificate and it could occur that an assertion cannot be *visually proven* for every graph for which the assertion is true (for instance we discussed issues related to scalability in Section II-C). To this end, we also state when we want to say that a certain assertion can be visually proven for arbitrary graphs.

A. Visual Certificates and Visual Provability

We give formal requirements inspired by the concept of certifying algorithms discussed in Section III-A. Let $f: \mathcal{G} \rightarrow \{\text{true}, \text{false}\}$ be an *assertion function* for the set of graphs $\mathcal{G}$, i. e., for some graphs the assertion $f(G)$ is *true* while for others it is not. For instance, if f is the existence of a cut-vertex, some graphs do contain one ($f(G) = \text{true}$) while others do not ($f(G) = \text{false}$). Consider a graph G with $f(G) = \text{true}$ and let $W(G)$ be a visualization of G . We call $W(G)$ *visual certificate* for $f(G)$ if and only if the following hold:

- 1) *Unimpeachability*: We call $W(G)$ *unimpeachable*, if it satisfies the following two properties. First, $W(G)$ should provide *information faithfulness* [60], i. e., it displays the ground truth properties and structures in G . Second, $W(G)$ should provide *task readability* [60], i. e., the judge can *perceive* enough information for validating the assertion.
- 2) *Checkability*: Given $W(G)$, it is trivial to decide that $f(G) = \text{true}$. In particular, this means that the judge's perception leads to the formation of a *mental model* $\mathcal{M}(G)$ that makes it possible for the judge to *efficiently* validate the assertion. The number of distinct observations made by the judge in the process is called the *perceptual complexity*.
- 3) *Simplicity*: Given $\mathcal{M}(G)$, there is a *simple formal proof* for $f(G) = \text{true}$ that relies solely on conclusions that the judge may deduce using $\mathcal{M}(G)$. In particular, this means that $W(G)$ is *perceptually distinguishable* from any possible wrong visual certificate $W'(G)$.

If a visual certificate $W(G)$ exists for each $G \in \mathcal{G}$ with $f(G) = \text{true}$, we call f *visually provable*. Note that the complementary function f^c (which is *true* if and only if $f(G) = \text{false}$) needs not necessarily be visually provable. For instance, we were able to visually prove the assertion that G contains a Hamiltonian cycle in Section II-C but we argued that the absence of such a cycle cannot be visually proven in Section III-C. This and requiring unimpeachability are clear differences to the concept of certifying algorithms whereas checkability and simplicity occur in both models, here considering the perceptual abilities of the judge; see also Section III-A.

We are also interested in *how efficiently* the judge is able to validate $f(G) = \text{true}$ based on $\mathcal{M}(G)$. To this end, we define the *perceptual complexity* as the time that the judge needs to check the assertion given $\mathcal{M}(G)$. The perceptual complexity may depend on the size of the graph, however, in some scenarios (e.g., Example 1) it may be independent of it. Since we assume the judge to make an objective judgment based on the evidence, we can treat the thought process as a deterministic algorithm and

apply methods from complexity theory to evaluate the perceptual complexity.

Reconsider Examples 1 to 3. All drawings shown in Section II have been *unimpeachable* (Property 1), possibly with a single exception. In the drawing in Fig. 2(b), one may argue that it is impossible for the judge to perceive whether there is an edge hidden behind the alleged cut-vertex bringing into question the unimpeachability of the layout. Meanwhile in Fig. 2(c) it is even more difficult to follow all edges. However, by visually disambiguating the precise set of edges connected to the alleged cut-vertex, this version makes it easy to perceive all parts of the graph necessary for the validation.

On the other hand, the drawing in Fig. 2(b) also fails the *simplicity* requirement (Property 3). Namely, one could actually hide an edge behind the cut-vertex that connects the alleged left and right component. Most likely, $\mathcal{M}(G)$ would consist of two connected components connected at a single vertex, as the drawing shows two vaguely compact salient features touching at a single point. Hence, if we showed the judge two drawings of that type where in one of the drawings an edge was hidden behind the cut-vertex, the judge would be unable to *distinguish* the correct certificate from the wrong one. In contrast, the drawing in Fig. 2(c) circumvents this problem, as the cut-vertex is the only vertex at the bottom of the drawing. Similarly, the highlighted parts in Figs. 3(b), (c), and 4 draw the judge's attention towards them, hence, these structures would become part of the mental model. Given the highlighted parts, it is easy to prove connectivity or the existence of the Hamiltonian cycle.

Finally, for the drawings in Figs. 3(b), and 4(a) it is not trivial to *check* whether the assertion is correct (Property 2). The mental model $\mathcal{M}(G)$ will consist of a tangle of highlighted edges embedded in an even larger tangle of edges. While the judge will most likely succeed to establish that the highlighted parts of the graph form a tree or a cycle, they might have to perform a BFS. Hence, especially in the case of Example 2, this would be as efficient as computing the validity of the assertion from scratch and these visualizations are not to be considered visual proofs. On the other hand, the visualizations in Figs. 3(c), and 4(b), 4(d) lead to formation of mental models that highlight the tree or cycle in few components of the mental model. Hence the judge only needs to check if the mental model includes any vertices not belonging to the tree or cycle. Moreover, in Example 3 the judge also has to check that the salient feature representing the cycle in fact contains all edges. This can be easily done in Fig. 4(b) and (d), however, if the size of the graph was larger, the process would be more efficient for the visualization in Fig. 4(d) where it is sufficient to check that the long diagonal and the two single cells are present, i.e., the judge has to check only for a constant number of components yielding constant perceptual complexity. On the other hand, for a visualization in the style of Fig. 4(b), it is required to assert the existence of each presumed edge along the cycle. Thus, more components of the mental model must be checked and the perceptual complexity, that is, the time that the judge needs to check the assertion given the mental model, depends on the number of vertices in the graph.

B. Unimpeachability and the Defense Lawyer

Recall that the task of the defense lawyer is to establish *unimpeachability* (Property 1), i. e., we require that the visualization should display the ground truth properties of the graph, and

that the judge can clearly perceive the parts of the visualization required for validating the assertion. The latter property is required so that the judge can extract the evidence embedded in the visual certificate whereas the former property ensures that a non-certificate which showcases evidence embedded in a visualization of a graph that is not identical to the input can be detected as such. In the literature, these concepts are known as *faithfulness* and *readability*, respectively [60]. *Faithfulness* refers to whether a visualization of a graph faithfully displays its ground truth properties and structures in a logically consistent manner, and the *readability* refers to the perceptual and cognitive interpretation of the visualization by the viewer [60].

More specifically, *task faithfulness* means that visualization contains enough information about a graph G to correctly perform the task. Faithfulness metrics are defined based on the type of ground truth structures of graphs, such as *shape* [42], *cluster* [55], *symmetry* [56] and *change* [54] and can be appropriately selected depending on the application scenario. Similarly, *task readability* means that the user can perceive enough information from the visualization to correctly perform the task, here validating the assertion. In our examples for good visual certificates, task readability was improved at the cost of information readability compared to standard layouts; see Table I.

Our definition of *unimpeachability* is purposely quite objective so that it can be discussed efficiently when establishing new visual proof techniques. A purely logically acting defense lawyer should have no reason to raise doubts about the visual certificate if we obey these requirements and no feedback between defense lawyer and judge is necessary (hence the dashed connection in Fig. 1)

C. the Mental Model and Perceptual Complexity

The mental model $\mathcal{M}(G)$ formed by the judge is an important component in assessing the usefulness of a visual certificate. Since cognition [57] and perception [36] differ from user to user, we must *predict* the *expected* mental model $\mathcal{M}(G)$ instead of assessing the mental model of a specific user. Hence, to influence the mental model, we have to carefully design visual certificates so as to exploit known features of visual perception. In our analysis of Example 1 in Section II-A, we relied on the effect of salient features being automatically grouped and perceived as cohesive components [75]. Thus, we have good reason to assume that a human observer would see two components glued together at a vertex located below both components. In Sections II-B and II-C, we exploited the fact that salient red-colored components naturally draw the attention of the user [44], so that they form a distinct shape in the foreground with the rest of the graph in the background. Thus, the graph layout and rendering style guides the judge's perception and simplifies the analysis.

Another important aspect regarding the judge's mental model is that we expect them to make an objective judgment based on the evidence encoded in the mental model. Thus, in our model the judge is not influenced by any prior knowledge or hypotheses regarding the data but at the same time will only accept the assertion if it has become irrefutable.

V. VISUAL PROOFS FOR GRAPH PROPERTIES

We provide visual proofs for further widely used assertions. We also discuss properties that can be visually proven using

TABLE II
OVERVIEW OF GRAPH PROPERTIES WITH VISUAL PROOFS PRESENTED IN THIS PAPER, AND CORRESPONDING COMPUTATIONAL COMPLEXITY OF UNDERLYING PROPERTIES AS WELL AS PERCEPTUAL COMPLEXITY OF PRESENTED VISUAL PROOFS (n AND m DENOTE THE NUMBERS OF VERTICES AND EDGES, RESP.). SHADED CELLS INDICATE OPEN PROBLEMS, – THE ABSENCE OF VISUAL PROOFS

Assertion	Complexity		Sec.
	Comp.	Percep.	
G is connected	$O(n+m)$	$O(n)$	II-B
G is not (2-)connected		$O(1)$	II-A
G is k -connected	$O(k^3 n^2)$	$O(kn)$	V-C
G is not k -connected		$O(k)$	II-A
G has a Hamilt. cycle (path)	NP-comp.	$O(1)$	II-C
G has a length- k cycle (path)	NP-comp.	$O(k)$	II-C
G is (not) bipartite	$O(n+m)$	$O(1)$	V-A
G is k -colorable	NP-comp.	$O(k)$	V-A
G is (outer-)planar	$O(n)$	$O(n)$	V-B
G is not (outer-)planar	$O(n)$	$O(n)$	V-C
G is (not) complete	$O(n^2)$	$O(1)$	V-A
G contains a size- k clique	NP-comp.	$O(k)$	V-C
G has a size- k independent set	NP-comp.	$O(k)$	V-C
G has a size- k dominating set	NP-comp.	$O(n)$	V-C
Queue number of G is $\leq k$	NP-comp.	$O(n+m)$	V-B
Stack number of G is $\leq k$	NP-comp.		V-B
Diameter of G is $> k$	$O(nm)$	$O(k)$	V-D
$d(u, v) = k$	$O(n+m)$	$O(k)$	V-D
Diameter of G is $\leq k$	$O(nm)$	$O(n+m)$	V-D
Center of a graph	$O(nm)$		V-D
centrality, e.g., betweenness	$O(nm)$		V-D
k -core analysis	$O(n+m)$	$O(n+m)$	V-D
Automorphic equivalence	Isomorp.-c.		V-D
co-NP-complete assertions	co-NP-c.	–	III-C V-C

canonical visualizations and motivate further research by discussing assertions that seem quite challenging to visually prove and providing examples for applying visual proofs in the context of network analysis. For a summary of our discussion, refer to Table II.

A. Important Graph Properties

a) *Non- k -Connectivity*: Recall the discussion of Example 1 in Section II-A. Effectively, we showed here that the graph is *not 2-connected* as the removal of the cut-vertex separates the graph into two pieces. Similarly, we can visually prove the assertion that a graph is *not k -connected*: In this case, there is a set of $k - 1$ vertices whose removal separates the graph and we can layout the graph in such a way that all connections between two clearly separated parts run via this vertex set which we can again place below the remainder of the drawing. If k becomes large, the judge has to actively start counting, however, arguably the task can still be efficiently solved.

b) (Non-)Bipartiteness and k -colorability:

We can use a matrix representation to visually prove bipartiteness; see Fig. 5(a). When sorting the rows and columns according to the two independent subsets, bipartiteness can be simply checked by verifying if the two empty squares are indeed empty [82]. This approach also generalizes to k -colorability as shown in Fig. 5(b) for 4 colors, however, for sparse graphs like in Fig. 5(c) additional highlighting of the (supposedly)

TABLE III

SELECTED AESTHETIC METRICS OF THE STRAIGHT-LINE NODE-LINK DRAWINGS IN SECTION V. RED NUMBERS IN PARANTHESES GIVE THE CORRESPONDING VALUES FOR SUBGRAPHS HIGHLIGHTED IN RED IN THE CORRESPONDING FIGURE. FOR STRESS AND CROSSING NUMBER (RED BACKGROUND) LOWER NUMBERS ARE BETTER, OTHERWISE (GREEN BACKGROUND) HIGHER NUMBERS ARE BETTER

	Fig. 5: k -Col.	Fig. 6: Non-Bipartite			Fig. 7: Non-Complete		Fig. 8: Planarity			
	(d)	(a)	(b)	(c)	(a)	(b)	(a)	(b)	(c)	(d)
STRESS [46]	6635	567.6 (9.1)	651.0 (1.9)	1022.7 (3.3)	58.8	94.4	179.3	258.4	182.8 (.651)	114.2 (1.34)
CROSSING NUMBER [62]	7452	6211 (2)	9493 (0)	9992 (0)	12649	12650	0	0	25 (5)	24 (0)
JACCARD INDEX [1]	.037	.177 (.243)	.180 (.361)	.182 (.361)	.918	.920	.127	.098	.116 (.640)	.170 (.391)
EDGE-LEN. RAT. [49], [50]	.090	.020 (.161)	.023 (.586)	.013 (.472)	.126	.047	.024	.026	.069 (.486)	.068 (.552)
NODE RESOLUTION [1]	.020	.016 (.123)	.016 (.314)	.011 (.307)	.126	.047	.024	.026	.035 (.486)	.052 (.324)
ASPECT RATIO [62]	1	.941 (.606)	.941 (.977)	.994 (.994)	.998	.898	.472	.167	.420 (.996)	.739 (.724)
CROSSING RESOLUTION [24]	2.6	0.79 (73.2)	0.79 (N/A)	0.74 (N/A)	14.4	4.68	N/A	N/A	34.9 (62.3)	20.0 (N/A)
ANGULAR RESOLUTION [62]	0.21	.017 (6.87)	.017 (112.0)	.007 (97.6)	7.20	1.23	0.0	0.0	2.25 (28.5)	0.0 (3.6)




Fig. 5. Visualizing k -colorability. (a) A bipartite graph. (b) A dense 4-colorable graph. (c) and (d) an adjacency matrix and a node-link visualization of a sparse 4-colorable graph.




Fig. 6. Visualizing non-bipartiteness of a graph. In (a) the odd cycle is self-intersecting, making it difficult to certify that it is in fact a cycle. Both in (b) and (c) the cycle is clearly visible where in (c) the cycle forms the outer boundary of the drawing letting it stand out even more compared to (b). Finally, in (d) the odd cycle is represented by a distinguishable pattern in the adjacency matrix.

empty squares might be necessary. For small graphs, a node-link diagram might be easier to read and hence preferable, however the approach does not scale well due to resolution since the judge needs to verify there are no edges within the subsets; see Fig. 5(d).

An odd-length cycle certifies that a graph is not bipartite, so non-bipartiteness can be visually proven by highlighting a shortest odd cycle in a drawing. In an arbitrary drawing, the cycle may be hard to spot, see Fig. 6(a). Redrawing the cycle in convex position makes it easier to read (see Fig. 6(b)), especially if it is the convex outer cycle; see Fig. 6(c) (this makes the rest of the graph harder to read; see Table III). The cycle is now clearly visible and the judge just needs to assert oddness. While depending on the odd cycle length counting may be inevitable, the judge can use the symmetry of the drawing of the cycle to see that the cycle is odd (e.g., in Fig. 6(c), there is a single top-most but no single bottom-most vertex). For larger cycle lengths, an adjacency matrix representation may be beneficial: Sort the rows and columns along the odd cycle and mark it, then append the remaining vertices arbitrarily. Then, alter the spacing of the matrix so that even rows and columns are thicker than odd




Fig. 7. Visualizing non-completeness of a graph. In (a) the missing edge is very difficult to spot. In (b) and (c) on the other hand it is obvious that an edge is missing.

ones; see Fig. 6(d). The cell closing the cycle is a square if and only if the length is odd.

c) Completeness and Non-Completeness:

Non-completeness is evidenced by a single missing edge and can be visually proven with a circular layout with the missing edge on the outer cycle. This approach does not scale well for a larger graphs; see Fig. 7(a). Readability and scalability can be improved by drawing focus to the missing edge, see Fig. 7(b). However, one can also use a matrix representation (see Fig. 7(c)) since spotting a missing square scales well from a perception perspective [82]. This technique can also prove completeness.

B. Visual Proofs Using Canonical Representations

We discuss some graph properties that have historically been associated with certain representations that can serve as visual certificates.

a) *Planarity and Outer-Planarity:* Planar and outer-planar graphs are important as some types of real-world networks are planar or near-planar. For instance, road-networks and train-networks can be interpreted as graphs where junctions/interchanges are nodes and connections between junctions are edges. This is relevant because, certain problems admit efficient solutions for planar graphs but not in general. The defining characteristic of these graph classes is that they admit a crossing-free drawing (for outer-planarity, additionally, every vertex is located on the outer face). There is a plethora of drawing algorithms for these graphs that can be used to create a visual certificate [23].

In order to produce a visual certificate, one may proceed as follows. First, compute a planar embedding; a selection of practical planar embedding algorithms can for instance be found in the work on certifying algorithms [53] where planarity is also discussed as an example problem where a certificate should be




Fig. 8. Visualizing (a) planarity, (b) outer-planarity, (c) non-planarity and (d) non-outerplanarity.

presented as part of the solution to an algorithmic problem. In a next step, we can compute a planar straight-line drawing in worst-case asymptotically optimal area using for instance the algorithm by Kant [47]; see also Fig. 8(a) for an example output of this approach.

In fact, the same approach can be applied if one wants to prove *outerplanarity*, i.e., in order to show that a graph admits a planar drawing where every face is adjacent to the outer face. In this scenario, we can insert an additional vertex v^* into the outer face and connect it to all other vertices. We then choose a triangle that includes v^* as the outer face and apply the algorithm by Kant [47]. Once v^* is deleted again, every vertex is then visible from above, which nicely indicates the adjacency to the outer face; see Fig. 8(b).

While the algorithm by Kant [47] is worst-case optimal regarding the consumed area, the produced drawings suffer from other undesirable properties, e.g., angular resolution and edge-length ratio are very small even for smaller input graphs; see also Table III. As a result, scalability can be problematic and computing drawings *highlighting* planarity for larger graphs is an intriguing open question.

b) Non-Planarity and Non-Outerplanarity: It is well known that a graph is non-planar if and only if it contains either a K_5 or $K_{3,3}$ minor [48] while it is non-outer-planar if and only if it contains either a K_4 or $K_{2,3}$ minor [19]. Namely, a minor is a graph that can be obtained from the initial graph by a series of *edge contractions*, i. e., one identifies both ends of an edge. A possible way to highlight non-planarity is to draw the graph in such a way, that a K_5 or $K_{3,3}$ minor M becomes obvious. To this end, one can group the vertices that will eventually be contracted to a vertex of M and provide a highlighted representation of M surrounded by the remainder of the graph; see, e.g., Fig. 8(c) where the K_5 minor of the displayed graph is visualized in convex position. The same kind of approach can be applied for showing a K_4 or $K_{2,3}$ minor indicating non-outerplanarity. See, e.g., Fig. 8(d) where a $K_{2,3}$ is highlighted in a non-planar drawing of a planar but not outerplanar graph. Once more, these examples exemplify that overall readability can be sacrificed to make a visual proof obvious; see Table III. Moreover, also note that we can actually highlight the two parts of a $K_{2,3}$ and $K_{3,3}$ minor in two different colors as in Fig. 8(d).

For larger graphs, however, each vertex in the minor can correspond to a large subgraphs of the initial drawing. Thus, it




Fig. 9. Visual certificate for stack number 3.

may be infeasible to visualize them in an easy-to-grasp fashion in a static layout. Thus, one may want to animate the edge contraction sequence. We leave it as an interesting open question whether such a visualization would be a convincing visual proof for non-planarity of larger graphs.

c) Stack and Queue Number: Stack and queue number are graph parameters arising for instance in the context of scheduling and VLSI layout (see [38], [83]). If a graph has a bounded stack (queue) number it admits a stack layout (queue layout, resp.) which uses few *pages*, i. e., layers in which the edges are drawn. On each page in a stack layout, no edges cross while in a queue layout, no edges nest. These features can be easily detected, i. e., a stack (queue) layout with k pages visually proves that the stack (queue) number is *at most* k . If the stack (queue) number is small, one can use a projective visualization, displaying the different pages as if they were drawn in different half-planes of the three-dimensional space; see Fig. 9. In contrast, if the stack (queue) number increases, it may be necessary to show each page separately and user interaction may be required.

C. Assertions That are Challenging to Visually Prove

a) k -Connectivity: Visually proving that a graph is k -connected is much more difficult compared to the complementary question discussed in Section V. Specifically, the prosecution lawyer needs to highlight a sparsification G' (i. e., a k -connected spanning subgraph) of G which is k -connected, similar to the visual proof in Section II-B. We can utilize the Nagamochi-Ibaraki algorithm [58], which computes a k -connected spanning subgraph G' with $O(kn)$ edges in linear time. Therefore, the judge may verify k -connectivity in $O(kn)$ time using G' , which is much faster for dense graphs with $O(n^2)$ edges.

However, for highly connected and large graphs (i. e., large values of k and n), the perceptual complexity will be quickly increasing, making the design of effective visual proofs a significantly challenging problem. For small values of k , say $k = 2$, one may be able to design more effective visual proofs by highlighting the structure of the sparsification better. Note that the problems of determining planar [33] or minimum sized [20] 2-connected spanning subgraphs are NP-complete while there is a PTAS for the latter problem [31].

b) Parameterized and co-NP-hard Assertions: Other interesting variants of some problems presented in Sections II, V, and V-B are *parameterized assertions*. That is, instead of asking for a Hamiltonian cycle, we may ask for the existence of a cycle of length k . Similarly, we could be interested to prove that there is an independent set, a clique or a dominating set of size k . Such assertions can be easily visually proven, however, the perceptual complexity depends on the value k as the judge has to count the size of the certificate.

The complementary questions for some of the presented problems are co-NP-complete; e. g., showing the *absence* of a Hamiltonian cycle, proving that the queue (stack) number of a graph is *at least* k , and showing that a graph is *not* k -colorable. This is in line with our conjecture that co-NP-complete problems admit no visual proofs in general. On the other hand, for specific instances, a visual proof might exist. For instance, if a graph contains a $(k + 1)$ -clique, this clique can be highlighted (e.g., similar to our proof of non-planarity in Fig. 8(c)) to indicate non- k -colorability. However, there are also graphs without $(k + 1)$ -cliques that cannot be colored with k colors. (e.g., cycles of odd length).

D. Visual Proofs for Network Analysis

So far, we have mainly discussed how to visually prove assertions stemming from graph theory. In network analysis, other analyses are equally important. We now discuss how we can visually prove assertions in this domain.

a) *Diameter and Center*: The *diameter* of a graph is the length of the longest shortest path between any two vertices in a graph, and the *center* of a graph is a vertex whose distance (i. e., shortest path) to all other vertices in the graph is minimized. Both are fundamental properties of graphs based on the distance between the vertices. A visual proof of the diameter can be obtained by the small multiples of the level drawings of the BFS trees rooted at each vertex v , where the BFS tree with the maximum depth is highlighted as the diameter of the graph (in fact this single BFS tree is sufficient for a lower bound for the diameter). Such a level drawing of the BFS tree rooted at a vertex v can be also used as a visual proof for vertices whose distance from v is exactly k , by highlighting the vertices on level k . Similarly, it can be used for a visual proof for the graph center problem, however this is more involved due to the comparison of the sum of all distances from v to other vertices. The perceptual complexity of these visual proofs increases with size and density.

b) *Other Network Analysis Measures*: There is a plethora of important measures in network analysis, e. g. *centrality*, *k-core* and *automorphic equivalence* [78]. The former two can be computed in polynomial time, while the latter is Isomorphism-complete. Possible visual proofs for these measures can be a level drawing or a radial drawing, where the level or concentric circle is defined based on the corresponding measures, such as centrality values or k -core index. Moreover, the k -core analysis can be visualized by a topographic map-style inclusion drawing, which may serve as a visual certificate. Note that these visual proofs require a careful analysis by the judge resulting in linear perceptual complexity, that is, novel visual certificates that the judge can validate more efficiently are of interest.

VI. LIMITATIONS OF THE GRAPHTRIALS MODEL

a) *Scalability*: In the GRAPHTRIALS model, we must not only visualize the evidence represented in the visual certificate, but also display the remainder of the graph faithfully. This may result in higher computational complexity compared to other visualization techniques, e. g., force-directed graph layouts, whose purpose is to create an overall readable representation. Even more, one may argue that it would also be possible to forgo visualization completely and use an assertion software to

validate the evidence computationally. While this could drastically reduce the computation time and require fewer software components, there are in fact real-world application scenarios, e. g., in court, where it may be better to show a visual certificate accompanied by a short explanation why the certificate is indeed establishing the assertion instead of simply telling the audience that a piece of software analyzed the network and found the evidence for the assertion; see Section I.

Another benefit of visual proofs over a non-visual assertion software is that bugs in the visual proof pipeline can be spotted in the visual certificate, i. e., either the represented graph is not the input graph or the evidence is not a true evidence for the claim. Another scalability issue is to display the entire graph faithfully. In Section IV, we assumed that the visual certificate may be represented by few components in the judge's mental model and that the formation of that mental model can be mainly guided by usage of bottom-up and pattern recognition processes. For large input graphs, the screen resolution might not permit an information-faithful representation of the input graph so that one must resort to techniques for displaying larger data, e. g., zooming. The introduction of such modes of user interaction may be problematic for our model as it may lead the judge to increasingly use top-down processes of perception which may influence the formation of the mental model. In this regard, we also note that user interaction might be a useful tool for the judge to check for unimpeachability once they extracted a candidate certificate from the visualization. For instance, in Example 1, an interactive display of Fig. 2(b) might allow the judge to move the potential cut-vertex in order to verify that no edge is hidden behind the vertex. Such features might be particularly useful for visualizations of larger graphs where collinearities of vertices are harder to avoid. We leave it as an intriguing open problem to investigate the application of interactive features in visual certificates further.

b) *Background Knowledge Required for Checkability*: In practical settings, checkability, i. e., the judge's ability to check the validity of an assertion efficiently based on their mental model, also depends on the background knowledge of the judge. For example, in Section II, we required knowledge of the terms cut-vertex, connectivity and Hamiltonian cycle. Thus, we may be tempted to require the judge to know basic notions of graph theory, specifically knowledge of the properties for which the visual proof is exhibited. In addition to graph-theoretic background knowledge, familiarity with the chosen visualization style may be important. For instance, we have seen in Section II-C that less frequently used paradigms such as adjacency matrix visualizations may allow for very efficient visual certificates. These aspects are limitations, when the audience is in fact non-expert and some required background knowledge is lacking. If so, additional explanations may also be presented to the audience examining the visual certificate, similar to court proceedings; see also Section I.

c) *Subjective Aspects of Unimpeachability*: In Section IV-A, we introduced unimpeachability as a combination of information faithfulness and task readability. As a consequence, the defense lawyer has to check properties that can either be objectively fulfilled or violated. Thus, in its most simple form, the defense lawyer can be a computer program that checks whether the visualized graph is actually the input graph (such as in Example 3) and the judge can rely on the fact that the defense lawyer would have rejected the drawing if it was not unimpeachable. Moreover, such a defense lawyer software may be implemented once for a

certain drawing style and reused for other visual proof processes utilizing the same layout technique. We believe that our objective definition of unimpeachability works well in scenarios where we simply want to check that the drawing method worked correctly (as in Example 3) or where the judge has no reason to distrust the prosecution lawyer (as in Example 1). However, the requirements for unimpeachability are context dependent and a visual proof may need to be defended against subjective counterarguments, brought forth by an adversary (as in Example 2) or by the skepticism of the judge (i. e., a single person fulfills both roles judge and defense lawyer). In such scenarios, there may be further subjective aspects of unimpeachability, e. g., the judge may expect the visualization to be similar to an already known layout of the graph, as otherwise the defense lawyer may question that the visual certificate shows the graph in question, i. e., there is feedback from the defense lawyer to the judge; see the dashed connection in Fig. 1.

d) *Uncertainty related to the Mental Model*: The judge's mental model is an abstraction that is difficult to describe as it can vary from judge to judge due to differences in both cognition [57] and perception [36]. Further, it may even be uncertain to the judge how they abstract the visualization. Thus, there are important open questions related to the judge's mental model: It is important to understand how visualization techniques can accurately influence the abstraction of the judge so that the evidence gathered by the prosecution lawyer can be translated into the mental model as unmodified as possible. Gathering empirical evidence that our concept of perceptual complexity captures the reality, i. e., that judges indeed use the mental model to establish the verdict, is an intriguing open problem. Moreover, we may ask which perceptual complexity is still accepted by human judges (linear workload may already be overwhelming). Finally, it may be worth investigating alternative measures for perceptual complexity that are agnostic to whether or not a human judge actually follows a deterministic algorithm. Such measures may be similar to the ones used in predictive models in HCI such as KLM [16] or GOMS [17].

e) *Non-Human Judges*: In this paper, we described the judge as a necessarily human actor in our model. As a result, we developed our requirements for the validity of a visual certificate based on generally accepted properties of the human visual system. Hence, our examples of visual certificates primarily use visual channels and representations suitable for human perception. The advent of multi-modal large language models (MLLMs, for short), however, might give rise to application scenarios where the entity to be convinced of a certain graph property might actually be a generative AI instead of a human. To this end, also note that MLLMs actually already possess some visualization literacy [11], [41], [51].

As a result, it is of interest to also investigate whether our model for visual proofs extends to the requirements of non-human judges, which can for instance be MLLMs. As a preliminary study, we asked the MLLM ChatGPT [62] to assess the visual certificates we presented in our initial examples in Section II. Overall, we observed that our visual proof techniques could also help ChatGPT making a correct assessment. While in general the MLLM expressed high confidence regarding its answers, we observed that for the more challenging tasks of finding a spanning tree and a Hamiltonian cycle it resorted to proxies instead of evaluating the concrete question asked. Moreover, we observed that a suitable visualization could increase the reported confidence. Further, we observed that the MLLM did

not always express doubts about potentially misleading visual certificates (e.g., Figs. 2(b) or 3(b)), which might cause problems in scenarios where MLLMs are actually employed as decision makers. See supplemental material for details. We believe that further evaluation is an important open problem.

VII. OPEN PROBLEMS

The most important open problems motivated by our model are related to empirically evaluating its effectiveness in user studies. In particular, we ask:

- 1) Are visual proofs in fact scalable? How do they extend to geospatial and dynamic graphs where the data are expected to obey spatial and/or temporal constraints?
- 2) Which features contribute to perceptual complexity?
- 3) Do response times depend mostly on perceptual complexity?
- 4) When do human users regard a visual certificate as unimpeachable?
- 5) What are human limits for the perception of graph properties? For instance, the minimum perceivable slope difference is ≈ 2 degrees [45].
- 6) What is the trade-off between perceptual complexity and cognitive load?

Moreover, we believe that a software fulfilling the role of the defense lawyer would be a helpful tool, both for further studying visual certificates but also for evaluating the trustworthiness of graph visualizations.

ACKNOWLEDGMENT

This work was initiated at Dagstuhl seminar 23051 "Perception in Network Visualization". We thank the organizers for making this fruitful interdisciplinary exchange possible and all participants for interesting discussions and insights during the seminar week.

REFERENCES

- [1] A. R. Ahmed, F. D. Luca, S. Devkota, S. G. Kobourov, and M. Li, "Multicriteria scalable graph drawing via stochastic gradient descent, (SGD)²," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 6, pp. 2388–2399, Jun. 2022, doi: [10.1109/TVCG.2022.3155564](https://doi.org/10.1109/TVCG.2022.3155564).
- [2] M. Aigner and G. M. Ziegler, "Five-coloring plane graphs," in *Proofs From the Book* (3rd ed.). Berlin, Germany: Springer, 2004, doi: [10.1007/9783662054123_30](https://doi.org/10.1007/9783662054123_30).
- [3] L. Angori, W. Didimo, F. Montecchiani, D. Pagliuca, and A. Tappini, "ChordLink: A new hybrid visualization model," in *Proc. Graph Drawing Netw. Visualization: 27th Int. Symp.*, 2019, pp. 276–290, doi: [10.1007/9783030358020_22](https://doi.org/10.1007/9783030358020_22).
- [4] L. Angori, W. Didimo, F. Montecchiani, D. Pagliuca, and A. Tappini, "Hybrid graph visualizations with chordlink: Algorithms, experiments, and applications," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 2, pp. 1288–1300, Feb. 2022, doi: [10.1109/TVCG.2020.3016055](https://doi.org/10.1109/TVCG.2020.3016055).
- [5] K. Appel and W. Haken, "Special announcement," *Discrete Math.*, vol. 16, no. 2, pp. 179–180, 1976, doi: [10.1016/0012365X\(76\)901473](https://doi.org/10.1016/0012365X(76)901473).
- [6] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*. New York, NY, USA: Cambridge Univ. Press, 2009, doi: [10.1017/CBO9780511804090](https://doi.org/10.1017/CBO9780511804090).
- [7] M. Behrisch, B. Bach, N. H. Riche, T. Schreck, and J. Fekete, "Matrix reordering methods for table and network visualization," *Comput. Graph. Forum*, vol. 35, no. 3, pp. 693–716, 2016, doi: [10.1111/cgf.12935](https://doi.org/10.1111/cgf.12935).
- [8] M. A. Bekos et al., "A heuristic approach towards drawings of graphs with high crossing resolution," in *Proc. Graph Drawing Netw. Visualization: 26th Int. Symp.*, 2018, pp. 271–285, doi: [10.1007/9783030044145_19](https://doi.org/10.1007/9783030044145_19).

- [9] M. A. Bekos et al., "A heuristic approach towards drawings of graphs with high crossing resolution," *Comput. J.*, vol. 64, no. 1, pp. 7–26, 2021, doi: [10.1093/comjnl/bxzz133](https://doi.org/10.1093/comjnl/bxzz133).
- [10] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [11] A. Bendeck and J. T. Stasko, "An empirical evaluation of the GPT-4 multimodal language model on visualization literacy tasks," *IEEE Trans. Vis. Comput. Graph.*, vol. 31, no. 1, pp. 1105–1115, Jan. 2025, doi: [10.1109/TVCG.2024.3456155](https://doi.org/10.1109/TVCG.2024.3456155).
- [12] L. Bounegru, T. Venturini, J. Gray, and M. Jacomy, "Narrating networks," *Digit. Journalism*, vol. 5, no. 6, pp. 699–730, 2017, doi: [10.1080/21670811.2016.1186497](https://doi.org/10.1080/21670811.2016.1186497).
- [13] J. R. Brown, *Philosophy of Mathematics: A Contemporary Introduction to the World of Proofs and Pictures*, 2nd ed., Evanston, IL, USA: Routledge, 2008, doi: [10.4324/9780203932964](https://doi.org/10.4324/9780203932964).
- [14] M. Burch et al., "Guiding graph exploration by combining layouts and reorderings," in *Proc. 13th Int. Symp. Vis. Inf. Commun. Interact.*, 2020, pp. 25:1–25:5, 2020, doi: [10.1145/3430036.3430064](https://doi.org/10.1145/3430036.3430064).
- [15] S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., 1999, doi: [10.1111/cgf.13730](https://doi.org/10.1111/cgf.13730).
- [16] S. K. Card, T. P. Moran, and A. Newell, "The keystroke-level model for user performance time with interactive systems," *Commun. ACM*, vol. 23, no. 7, pp. 396–410, 1980, doi: [10.1145/358886.358895](https://doi.org/10.1145/358886.358895).
- [17] S. K. Card, T. P. Moran, and A. Newell, *The Psychology of Human-Computer Interaction*. Mahwah, NJ, USA: Erlbaum, 1983.
- [18] D. Ceneda, T. Gschwandtner, and S. Miksch, "A review of guidance approaches in visual data analysis: A multifocal perspective," *Comput. Graph. Forum*, vol. 38, no. 3, pp. 861–879, 2019.
- [19] G. Chartrand and F. Harary, "Planar permutation graphs," *Annales de l'institut Henri Poincaré. Sect. B. Calcul des probabilités et statistiques*, vol. 3, no. 4, pp. 433–438, 1967.
- [20] J. Cheriyan, A. Sebö, and Z. Szigeti, *An Improved Approximation Algorithm for Minimum Size 2-Edge Connected Spanning Subgraphs*. In R.E. Bixby, E.A. Boyd, and R. Z. Ríos-Mercado, eds., London U.K.: IPCO, 1998, pp. 126–136, 1998, doi: [10.1007/3540693467_10](https://doi.org/10.1007/3540693467_10).
- [21] J. Choo and S. Liu, "Visual analytics for explainable deep learning," *IEEE Comput. Graph. Appl.*, vol. 38, no. 4, pp. 84–92, Jul./Aug. 2018, doi: [10.1109/MCG.2018.042731661](https://doi.org/10.1109/MCG.2018.042731661).
- [22] S. Devkota, A. R. Ahmed, F. D. Luca, K. E. Isaacs, and S. G. Kobourov, "Stress-plus-x (SPX) graph layout," in *Proc. Int. Symp. Graph Drawing Netw. Vis.*, 2019, pp. 291–304, doi: [10.1007/9783030358020_23](https://doi.org/10.1007/9783030358020_23).
- [23] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1999.
- [24] E. Di Giacomo, W. Didimo, G. Liotta, and H. Meijer, "Area, curve complexity, and crossing resolution of non-planar graph drawings," in *Proc. Int. Symp. Graph Drawing Netw. Vis.*, 2009, pp. 15–20, 2009, doi: [10.1007/9783642118050_4](https://doi.org/10.1007/9783642118050_4).
- [25] E. Di Giacomo, W. Didimo, G. Liotta, and H. Meijer, "Area, curve complexity, and crossing resolution of non-planar graph drawings," *Theory Comput. Syst.*, vol. 49, no. 3, pp. 565–575, 2011, doi: [10.1007/s0022401092756](https://doi.org/10.1007/s0022401092756).
- [26] W. Didimo, L. Giamminonni, G. Liotta, F. Montecchiani, and D. Pagliuca, "A visual analytics system to support tax evasion discovery," *Decis. Support Syst.*, vol. 110, pp. 71–83, 2018, doi: [10.1016/j.dss.2018.03.008](https://doi.org/10.1016/j.dss.2018.03.008).
- [27] W. Didimo, L. Grilli, G. Liotta, L. Menconi, F. Montecchiani, and D. Pagliuca, "Combining network visualization and data mining for tax risk assessment," *IEEE Access*, vol. 8, pp. 16073–16086, 2020, doi: [10.1109/ACCESS.2020.2967974](https://doi.org/10.1109/ACCESS.2020.2967974).
- [28] W. Didimo, L. Grilli, G. Liotta, F. Montecchiani, and D. Pagliuca, "Visual querying and analysis of temporal fiscal networks," *Inf. Sci.*, vol. 505, pp. 406–421, 2019, doi: [10.1016/j.ins.2019.07.097](https://doi.org/10.1016/j.ins.2019.07.097).
- [29] H. Förster et al., "GraphTrials: Visual proofs of graph properties," in *Proc. Graph Drawing Netw. Vis.*, pp. 16:1–16:18, 2024, doi: [10.4230/LIPIcs.GD.2024.16](https://doi.org/10.4230/LIPIcs.GD.2024.16).
- [30] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York, NY, USA: W. H. Freeman, 1979.
- [31] L. Georgiadis, G. F. Italiano, and A. Karanasiou, "Approximating the smallest 2-vertex connected spanning subgraph of a directed graph," *Theor. Comput. Sci.*, vol. 807, pp. 185–200, 2020, doi: [10.1016/j.tcs.2019.09.040](https://doi.org/10.1016/j.tcs.2019.09.040).
- [32] E. D. Giacomo, W. Didimo, G. Liotta, F. Montecchiani, and A. Tappini, "Comparative study and evaluation of hybrid visualizations of graphs," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 7, pp. 3503–3515, Jul. 2024, doi: [10.1109/TVCG.2022.3233389](https://doi.org/10.1109/TVCG.2022.3233389).
- [33] O. Goldschmidt, A. Takvorian, and G. Yu, "On finding a biconnected spanning planar subgraph with applications to the facilities layout problem," *Eur. J. Oper. Res.*, vol. 94, no. 1, pp. 97–105, 1996, doi: [10.1016/03772217\(95\)002022](https://doi.org/10.1016/03772217(95)002022).
- [34] R. Gove et al., "Gragnostics: Fast, interpretable features for comparing graphs," in *Proc. 23rd Int. Conf. Inf. Vis.*, 2019, pp. 201–209, doi: [10.1109/TV.2019.00042](https://doi.org/10.1109/TV.2019.00042).
- [35] R. Gove, "Gragnostics: Evaluating fast, interpretable structural graph features for classification and visual analytics," in *Proc. Integrating Artif. Intell. Visual. Vis. Knowl. Discov.*, 2022, pp. 311–336, doi: [10.1007/9783030931193_12](https://doi.org/10.1007/9783030931193_12).
- [36] L. Grzeczowski, A. M. Clarke, G. Francis, F. W. Mast, and M. H. Herzog, "About individual differences in vision," *Vis. Res.*, vol. 141, pp. 282–292, 2017, doi: [10.1016/j.visres.2016.10.006](https://doi.org/10.1016/j.visres.2016.10.006).
- [37] K. Healy, *Data Visualization: A Practical Introduction*. Princeton, NJ, USA: Princeton Univ. Press, 2018.
- [38] L. S. Heath and A. L. Rosenberg, "Laying out graphs using queues," *SIAM J. Comput.*, vol. 21, no. 5, pp. 927–958, 1992, doi: [10.1137/0221055](https://doi.org/10.1137/0221055).
- [39] N. Henry and J. Fekete, "MatrixExplorer: A dual-representation system to explore social networks," *IEEE Trans. Vis. Comput. Graph.*, vol. 12, no. 5, pp. 677–684, Sep./Oct. 2006, doi: [10.1109/TVCG.2006.160](https://doi.org/10.1109/TVCG.2006.160).
- [40] N. Henry, J. Fekete, and M. J. McGuffin, "NodeTriX: A hybrid visualization of social networks," *IEEE Trans. Vis. Comput. Graph.*, vol. 13, no. 6, pp. 1302–1309, Nov./Dec. 2007, doi: [10.1109/TVCG.2007.70582](https://doi.org/10.1109/TVCG.2007.70582).
- [41] J. Hong, C. Seto, A. Fan, and R. Maciejewski, "Do LLMs have visualization literacy? An evaluation on modified visualizations to test generalization in data interpretation," *IEEE Trans. Vis. Comput. Graph.*, early access, pp. 1–13, 2025, doi: [10.1109/TVCG.2025.3536358](https://doi.org/10.1109/TVCG.2025.3536358).
- [42] S. Hong et al., "DSC: New degree-based shape-based faithfulness metrics for large and complex graph visualization," in *Proc. Pacific Vis. Symp.*, 2022, pp. 51–60, doi: [10.1109/PacificVis53943.2022.00014](https://doi.org/10.1109/PacificVis53943.2022.00014).
- [43] R. Houtkamp, H. Spekrijns, and P. R. Roelfsema, "A gradual spread of attention," *Percept. Psychophys.*, vol. 65, no. 7, pp. 1136–1144, 2003, doi: [10.3758/BF03194840](https://doi.org/10.3758/BF03194840).
- [44] L. Itti, "Models of bottom-up attention and saliency," in *Neurobiology of Attention*. Amsterdam, The Netherlands: Elsevier, 2005, pp. 576–582, doi: [10.1016/B9780123757319/500987](https://doi.org/10.1016/B9780123757319/500987).
- [45] M. Janssen, M. Lauvenberg, W. van der Ven, T. Bloebaum, and H. Kingma, "Perception threshold for tilt," *Otol. Neurotol.*, vol. 32, no. 5, pp. 818–825, 2011, doi: [10.1097/MAO.0b013e31821c6c7b](https://doi.org/10.1097/MAO.0b013e31821c6c7b).
- [46] T. Kamada and S. Kawai, "An algorithm for drawing general undirected graphs," *Inf. Process. Lett.*, vol. 31, no. 1, pp. 7–15, 1989, doi: [10.1016/00200190\(89\)901026](https://doi.org/10.1016/00200190(89)901026).
- [47] G. Kant, "Drawing planar graphs using the canonical ordering," *Algorithmica*, vol. 16, no. 1, pp. 4–32, 1996, doi: [10.1007/BF02086606](https://doi.org/10.1007/BF02086606).
- [48] C. Kuratowski, "Sur le problème des courbes gauches en topologie," *Fundam. Mathematicae*, vol. 15, pp. 271–283, 1930.
- [49] S. Lazard, W. J. Lenhart, and G. Liotta, "On the edge-length ratio of outerplanar graphs," in *Proc. Graph Drawing Netw. Vis.: 26th Int. Symp.*, 2017, pp. 17–23, 2017, doi: [10.1007/9783319739151_2](https://doi.org/10.1007/9783319739151_2).
- [50] S. Lazard, W. J. Lenhart, and G. Liotta, "On the edge-length ratio of outerplanar graphs," *Theor. Comput. Sci.*, vol. 770, pp. 88–94, 2019, doi: [10.1016/j.tcs.2018.10.002](https://doi.org/10.1016/j.tcs.2018.10.002).
- [51] Z. Li, H. Miao, V. Pascucci, and S. Liu, "Visualization literacy of multimodal large language models: A comparative study," 2024, *arXiv:2407.10996*.
- [52] S. J. Luck and A. R. Hollingworth, *Visual Memory*. London, U.K.: Oxford Univ. Press, 2008, doi: [10.1093/acprof:oso/9780195305487.001.0001](https://doi.org/10.1093/acprof:oso/9780195305487.001.0001).
- [53] R. M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer, "Certifying algorithms," *Comput. Sci. Rev.*, vol. 5, no. 2, pp. 119–161, 2011, doi: [10.1016/j.cosrev.2010.09.009](https://doi.org/10.1016/j.cosrev.2010.09.009).
- [54] A. Meidiana, S. Hong, and P. Eades, "New quality metrics for dynamic graph drawing," in *Proc. Graph Drawing Netw. Vis.: 28th Int. Symp.*, 2020, pp. 450–465, doi: [10.1007/9783030687663_35](https://doi.org/10.1007/9783030687663_35).
- [55] A. Meidiana, S. Hong, P. Eades, and D. A. Keim, "A quality metric for visualization of clusters in graphs," in *Proc. Graph Drawing Netw. Vis.: 28th Int. Symp.*, 2019, pp. 125–138, doi: [10.1007/9783030358020_10](https://doi.org/10.1007/9783030358020_10).
- [56] A. Meidiana, S. Hong, P. Eades, and D. A. Keim, "Automorphism faithfulness metrics for symmetric graph drawings," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, pp. 3241–3255, Jul. 2024, doi: [10.1109/TVCG.2022.3229354](https://doi.org/10.1109/TVCG.2022.3229354).

- [57] A. Miyake and N. P. Friedman, "The nature and organization of individual differences in executive functions: Four general conclusions," *Curr. Directions Psychol. Sci.*, vol. 21, no. 1, pp. 8–14, 2012, doi: [10.1177/0963721411429458](https://doi.org/10.1177/0963721411429458).
- [58] H. Nagamochi and T. Ibaraki, "A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph," *Algorithmica*, vol. 7, no. 5/6, pp. 583–596, 1992, doi: [10.1007/BF01758778](https://doi.org/10.1007/BF01758778).
- [59] R. Nelsen, "Proofs without words: Exercises in visual thinking," in *Classroom Resource Materials*. Washington, DC, USA: Mathematical Association of America, 1993.
- [60] Q. H. Nguyen, P. Eades, and S. Hong, "On the faithfulness of graph visualizations," in *Proc. IEEE Pacific Vis. Symp.*, pp. 209–216, 2013, doi: [10.1109/PACIFICVIS.2013.6596147](https://doi.org/10.1109/PACIFICVIS.2013.6596147).
- [61] T. Nishizeki and M. S. Rahman, *Planar Graph Drawing*, Lect. Notes Ser. Computing, Computing, vol. 12. Singapore: World Scientific, 2004, doi: [10.1142/5648](https://doi.org/10.1142/5648).
- [62] OpenAI. Chatgpt. Accessed: May 08, 2025. [Online]. Available: <https://chatgpt.com/>
- [63] Oxford English Dictionary, "Four-colour | four-color, adj," in *Oxford English Dictionary*. London, U.K.: Oxford Univ. Press, 2023.
- [64] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *Proc. Int. Conf. Intell. Anal.*, McLean, VA, USA, 2005, pp. 2–4.
- [65] H. Purchase, "Metrics for graph drawing aesthetics," *J. Vis. Lang. Comput.*, vol. 13, no. 5, pp. 501–516, 2002, doi: [10.1016/S1045926X\(02\)902326](https://doi.org/10.1016/S1045926X(02)902326).
- [66] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim, "Knowledge generation model for visual analytics," *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 1604–1613, Dec. 2014, doi: [10.1109/TVCG.2014.2346481](https://doi.org/10.1109/TVCG.2014.2346481).
- [67] J. M. Six and I. G. Tollis, "A framework for circular drawings of networks," in *Proc. Graph Drawing Netw. Vis.: 28th Int. Symp.*, 1999, pp. 107–116, doi: [10.1007/3540466487_11](https://doi.org/10.1007/3540466487_11).
- [68] S. Song, C. Li, D. Li, J. Chen, and C. Wang, "GraphDecoder: Recovering diverse network graphs from visualization images via attention-aware learning," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 7, pp. 3074–3088, Jul. 2024, doi: [10.1109/TVCG.2022.3225554](https://doi.org/10.1109/TVCG.2022.3225554).
- [69] R. Tamassia, *Handbook on Graph Drawing and Visualization*. London, U.K.: Chapman and Hall, 2013.
- [70] E. R. Tufte, *The Visual Display of Quantitative Information*. Graphics Press, 1992.
- [71] E. R. Tufte, *Visual Explanations - Images and Quantities, Evidence and Narrative*. Cheshire, CT, USA: Graphics Press, 1997.
- [72] E. R. Tufte, *The Visual Display of Quantitative Information*. Cheshire, CT, USA: Graphics Press, 2001.
- [73] T. Tymoczko, "The four-color problem and its philosophical significance," *J. Philosophy*, vol. 76, no. 2, pp. 57–83, 1979.
- [74] B. Victor, "Scientific communication as sequential art," 2011. [Online]. Available: <http://worrydream.com/ScientificCommunicationAsSequentialArt/>
- [75] J. Wagemans et al., "A century of gestalt psychology in visual perception: I. perceptual grouping and figure-ground organization. Psychol," *Bulletin*, vol. 138, no. 6, pp. 1172–1217, 2012, doi: [10.1037/a0029333](https://doi.org/10.1037/a0029333).
- [76] C. Ware, *Information Visualization—Perception for Design*. San Mateo, CA, USA: Morgan Kaufmann, 2004.
- [77] C. Ware, *Visual Thinking for Information Design*. San Mateo, CA, USA: Morgan Kaufmann, 2021, doi: [10.1016/C20160013955](https://doi.org/10.1016/C20160013955).
- [78] S. Wasserman and K. Faust, *Social Network Analysis: Methods and Applications*. New York, NY, USA: Cambridge Univ. Press, 1994, doi: [10.1017/CBO9780511815478](https://doi.org/10.1017/CBO9780511815478).
- [79] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, 1998, doi: [10.1038/30918](https://doi.org/10.1038/30918).
- [80] H. Wickham, D. Cook, H. Hofmann, and A. Buja, "Graphical inference for Infovis," *IEEE Trans. Vis. Comput. Graph.*, vol. 16, no. 6, pp. 973–979, Nov/Dec. 2010, doi: [10.1109/TVCG.2010.161](https://doi.org/10.1109/TVCG.2010.161).
- [81] Wikipedia, "Mathematical proof," 2023. [Online]. Available: https://en.wikipedia.org/wiki/Mathematical_proof#Visual_proof
- [82] J. Wolfe, "Visual search," *Curr. Biol.*, vol. 20, no. 8, pp. R346–R349, 2010, doi: [10.1016/j.cub.2010.02.016](https://doi.org/10.1016/j.cub.2010.02.016).
- [83] M. Yannakakis, "Embedding planar graphs in four pages," *J. Comput. Syst. Sci.*, vol. 38, no. 1, pp. 36–67, 1989, doi: [10.1016/00220000\(89\)900329](https://doi.org/10.1016/00220000(89)900329).
- [84] yWorks. yEd - graph editor. 2023. [Online]. Available: <https://www.yworks.com/products/yed>



Henry Förster received the BSc degree in engineering and computing from the Technical University of Freiberg, and the MSc and PhD degrees in computer science from the University of Tübingen. Currently, he is working as a postdoctoral researcher with Technical University of Munich. His research interests include graph drawing, geometric algorithms and human perception of information visualizations.



Felix Klesen received the BSc degree in aerospace computer science, as well as MSc and PhD degrees in computer science from the University of Würzburg. His research interests include graph drawing, geometric algorithms and algorithms for geoinformation systems. He is currently working as an optimization specialist in industry.



Tim Dwyer received the BCompSci(Hons.) degree from the University of Melbourne, in 2001 and the PhD degree from the University of Sydney, in 2006. He has worked in industry (Microsoft USA) as well as academia. His current appointment is professor with Monash University. His research interests include network and other data visualisation techniques as well as HCI and Immersive Analytics in AR/VR. He is a long time member of the IEEE VIS, ACM CHI conference communities, among others.



Peter Eades received the PhD degree from the Australian National University, in 1978, and has worked in Universities and Research Institutes across the world. He is an emeritus professor in the School of Computer Science with the University of Sydney, known for his research in Graph Drawing. He has been investigating methods for visualization of networks since the 1980 s. His algorithms are currently commonly used in diverse software systems, from social networks, biological networks, CASE tools, to security.



Seok-Hee Hong is a professor and ARC Future Fellow (2013-2016) with the School of Computer Science, University of Sydney. She was a Humboldt Fellow (2013-2014), ARC Research Fellow (2008-2012), and a project leader of VALACON (Visualisation and Analysis of Large and Complex Networks) project with NICTA (National ICT Australia), in 2004-2007. Her research interests include Graph Drawing, Algorithms, Network Visualization and Visual Analytics.



Stephen Kobourov received the BS degrees in mathematics and computer science from Dartmouth College, and the MS and PhD degrees from Johns Hopkins University. He is a professor of computer science with the Technical University of Munich. His research interests include information visualization, graph theory, and geometric algorithms.



Fabrizio Montecchiani received the PhD degree in computer engineering from the University of Perugia, in 2014. He is currently an associate professor with the University of Perugia, Department of Engineering. His research interests include graph drawing and graph algorithms, computational geometry, information visualization and visual analytics, algorithm engineering for Big Data.



Giuseppe Liotta (Senior Member, IEEE) is a professor of computer science with the Department of Engineering of the University of Perugia, Italy. His research interests are mainly directed with the analysis and design of algorithms and systems that have applications in the fields of network science, information visualization, and graph drawing. On these topics he published more than 300 research papers and received several grants. His research for this paper was partly supported by MUR of Italy (PRIN Project no. 2022TS4Y3N – EXPAND).



Alexander Pastukhov is a cognitive neuroscientist with the University of Bamberg. He worked with Volgograd State Technical University, University of Plymouth, and Magdeburg University. His research interests are in behavioral research and dynamic models of visual perception, consciousness, attention, and decision making.



Kazuo Misue (Member, IEEE) received the BSc and MSc degrees from the Tokyo University of Science, in 1984 and 1986, and the PhD degree in engineering from the University of Tokyo, in 1997. He is a professor with the Institute of Systems and Information Engineering, University of Tsukuba, Japan, where he performs research and development on information technology. His research interests include human-computer interaction, information visualization, and automatic graph drawing. He is a member of the ACM, IEICE, IPSJ, JSAI, and JSSST.



Falk Schreiber is full professor and head of the Life Science Informatics group at the University of Konstanz. He worked with universities and research institutes in Germany and Australia and is adjunct professor with Monash University Melbourne. His main interests are network analysis and visualization, immersive analytics as well as multi-scale modeling of processes, all with a focus on life science applications from metabolic processes to collective behavior.