Perceptual Scalability for Systemic Risk Visual Analytics

Richard Brath, Holland Vasquez, Michael Crouch, Martin Matusiak,
Nathan Kronenfeld, Derek Gray, Chris Carmichael, David Vergel de Dios, David Jonker

Uncharted Software Inc.
Toronto, ON, Canada
{rbrath, hvasquez, mcrouch, mmatusiak, nkronenfeld, dgray, ccarmichael, dvergeldedios, djonker}@uncharted.software

***Abstract*— Traditional risk assessments often focus on entity-level metrics, overlooking cascading failures and structural interdependencies. While computational models can simulate these dynamics, they frequently lack interpretability. Likewise, large-scale network visualizations often struggle with perceptual scalability, making it difficult for analysts to extract insights from dense, interconnected data. We present a human-in-the-loop (HITL) visual analytics approach that enhances systemic risk assessment through the domain-specific integration of hierarchical grouping, clustering, edge bundling, labelling, and zooming techniques. Our contribution is the deliberate combination of these methods to support perceptual scalability, enabling users to explore large, complex networks more effectively and reason about systemic vulnerabilities. By grounding our design in gestalt principles and analyst workflows, the system supports tracing cascading failures, validating computational models, and interacting with network structures across multiple levels of abstraction. Expert user studies demonstrate the system's effectiveness in improving interpretability, usability, and risk analysis workflows in complex domains.**

***Keywords*—graph visualization, perceptual scalability, systemic risk visual analytics**

I. Introduction

To anticipate and mitigate disruptions in highly interconnected systems, risk professionals must assess vulnerabilities across entire networks. A single point of failure can lead to widespread consequences. Yet, traditional risk assessments often focus on entity-level metrics—such as reliability scores or historical data—which overlook cascading failures and structural dependencies.

Systemic risk refers to disruptions that propagate across a network, driven by its topology. Events like the 2008 financial crisis and the 2020 supply chain crisis demonstrated how localized shocks escalate into near system-wide breakdowns. While computational models can simulate these effects, they often obscure structural relationships, preventing analysts from interpreting outputs or exploring scenarios interactively [1], [2].

Large-scale network visualization introduces additional challenges. As network size and complexity increase, clarity and interpretability often degrade [3]. Though techniques like clustering, edge bundling, and hierarchical grouping exist to support exploration, they are typically applied in isolation. Moreover, navigating large-scale networks requires seamless, high-speed transitions between global and local views, and persistent perceptual anchors—such as meaningful labels at all zoom levels—to prevent disorientation. Without these, analysts are forced into slow, memory-intensive workflows that fragment attention and obscure insight. Without coordination, these methods fail to achieve perceptual scalability—the ability of a visualization to preserve analytical utility at scale [4].

To address these gaps, we present a human-in-the-loop (HITL) visual analytics system for systemic and scenario-based risk analysis in large-scale supply-demand networks. Integrating computational modeling with perceptually optimized visualization, the system incorporates entity-level risk metrics, systemic models to capture cascading failures, and scenario simulations to explore counterfactual interventions **Our contribution is the purposeful, perceptually aware integration of established techniques—hierarchical grouping, clustering, edge bundling, subgraph overlays, labeling at all levels of abstraction, and rapid zooming—into a cohesive, task-driven workflow.** Informed by Gestalt principles and analyst workflows, our system enables users to trace disruptions and fluidly navigate networks at multiple levels. By ensuring perceptual continuity and interaction speed at scale, the system supports sensemaking, model validation, and decision-making in domains where conventional tools fall short. Expert evaluation confirms its effectiveness in enhancing usability, interpretability, and systemic risk reasoning in complex networks.

II. Background

A. Network Risk Modeling

Entity-level risk measures the reliability or vulnerability of individual nodes using pre-defined metrics to answer questions like, "How reliable is company X?" (Fig. 1 left). An entity can have dozens of measures that provide insights into its overall risk profile. A common approach for entity-level risk assessment systematically collects and analyzes metrics for each entity, enabling aggregation, filtering, and drill-down analysis. For example, the Department of Defense (DoD) Supply Chain Risk Management (SCRM) taxonomy defines 96 risk metrics across 12 categories [5], [6]. However, collecting high-quality,

This material is based in part upon work supported by the Defense Advanced Research Projects Agency (DARPA) contract HR001123C0030. Any views, opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. This work has been Approved for Public Release, Distribution Unlimited.

comprehensive risk data at scale is challenging, and many risks (e.g., legal, labor, corruption) are difficult to quantify. Assessing individual entities alone is insufficient; failures in interconnected systems often spread unpredictably, triggering cascading disruptions. To capture these, analysts can leverage systemic risk modeling methods.

Systemic risk modeling examines network-wide dependencies and cascading effects throughout networks. This allows decision-makers to answer questions such as "Which entities are structurally most important to network stability?" and "How will a disruption in one area of the network cascade to others?" (Fig. 1 middle). Network science provides analytic frameworks for assessing systemic fragility by analyzing the structure and dynamics of multilayer networks [7]. Many real-world networks exhibit non-normal degree distributions, where highly connected hubs amplify risk by facilitating rapid contagion [8]. Several computational methods exist for modeling systemic risk. Network centrality measures—degree, betweenness, and eigenvector centrality—help identify structurally critical nodes that trigger cascading failures if disrupted [9]. However, these static measures do not account for dynamic risk propagation. Monte Carlo simulations model probabilistic failure cascades, running thousands of randomized simulations to estimate risk distributions under uncertainty [10]. While they provide quantitative insights into systemic vulnerabilities, they are computationally expensive for large-scale networks and highly sensitive to input assumptions.

Systemic risk approaches like algorithmic techniques for robustness and resilience simulate failure cascades by removing nodes and measuring network recovery under alternative connection pathways [7]. These models assess a network's ability to absorb, recover from, and adapt to adverse events, though defining resilience remains inconsistent across domains [10], [11]. More advanced techniques, such as agent-based models (ABMs), simulate individual behaviors within economic, financial, or supply chain networks to model cascades, ripples, and bullwhip effects [12], [13]. While ABMs provide insights into micro-level interactions, simulating millions or billions of heterogeneous agents requires significant computational resources [14].

Scenario risk modeling builds on systemic risk by introducing counterfactual conditions to explore potential futures and intervention outcomes [15] (Fig. 1 right). Systemic risk modeling techniques naturally extend to scenario analysis by introducing controlled modifications to the network. These can be further extended to model different forms of resilience tactics, such as alternative suppliers or product substitutions. Monte Carlo simulations can generate distributions of possible future states under varying conditions. ABMs further enhance scenario risk analysis by simulating adaptive behaviors of network participants, allowing decision-makers to explore how firms, financial institutions, or suppliers react to policy changes, supply chain disruptions, or economic shocks.

B. *Visualization*

Graph visualization is highly relevant to systemic and scenario risk analysis due to the complex structure of network topology, including the arrangement, flow of interactions, and hierarchical dependencies [7]. While computational techniques provide quantitative insights into systemic fragility, effective risk assessment requires visual tools for searching, inspecting, and interacting with networks at multiple levels—from aggregations to individual entities. The following graph visualization techniques can be leveraged to support these abilities. Numerous techniques have been explored to enhance both the scalability and visual comprehension of large-scale graphs. This area has been extensively studied, yet many challenges in graph perception remain [16].

1) Perceptual scalability by overall graph organization: Graphs become difficult to parse when all nodes and edges look alike and there are no landmarks. Numerous graph organization techniques have been applied to address these issues. These include many different approaches for hierarchical grouping of nodes [1], [17], such as color-coding by partition or splitting groups apart (juxtaposition). Containment – i.e. Gestalt common region – is a strong cue used in many kinds of diagrams outside of graph visualization. Furthermore, the magnitude of entities within a group can be expressed by the size of the visual indicating the group. However, as the quantity of groups and their constituents increases, the diagram can become cluttered and difficult to interpret, with unavoidable overlapping containers obscuring element relationships [18]. Containment can be expressed as structure and partitions in one visualization [19], e.g. node glyphs in successive containers, rectangles [20], convex hulls [21], or nodes as treemaps [18].

Clustering [22] is typically done with force-directed layouts [23] (or techniques such as multidimensional reduction via tSNE or UMAP). Fundamentally, this relates to Gestalt proximity—things that are strongly connected are closest—and aligns with Tobler's first law of maps—everything is related, but close things are more related. The unintended consequence is that when all edges are drawn, the longest edges are most visually salient even though they might represent weak connections. Containment and clustering can be combined—for example, by applying clustering within partitions and using containment to

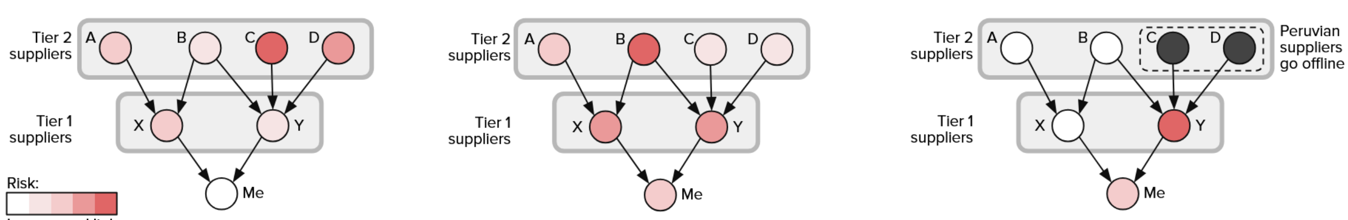

Fig. 1. Left: The network identifies entities; risks are intrinsic to each entity in isolation. Middle: Systemic risk identifies risks that can cascade through networks. Right: Scenario risk stresses the network with specific shocks.

express group boundaries [24]. This allows the visualization to leverage complementary perceptual cues: clustering supports structural comprehension through spatial layout, while containment reinforces semantic grouping, aiding interpretation of network topology and categorical membership. However, there are limitations in scaling to large-scale networks.

Edge bundling simplifies graphs by grouping related connections together. Edge bundling relates to Gestalt continuity, aiding perception of connections aligned with others. However, it also introduces ambiguity, as crossing edges appear aligned but are difficult to visually disentangle without interaction. Moreover, bundling depends on the spatial layout of the graph; it is not inherent to the data and is highly sensitive to positioning, which can impact interpretability [25].

2) Local perception: Using visual attributes to convey data on nodes and edges (size, hue, brightness, shape), is well studied and understood in graph visualization [26].

Labels are important for identification. Many tasks need to identify entities or groups thereof (selections, subgraphs, groups, individual nodes). Similarly, in set visualization, 12 of 26 set analysis tasks require identification [27]. Furthermore, orientation in a large information space depends on landmarks [28], which can be aided with labels [29]. Labels may be embedded within every node [30], which can restrict the size of the graph and hinder text readability. Some network visualizations place labels as a post process overtop the graph [31], occluding nodes or labels, thus requiring extra computationally expensive passes to shift nodes and edges for readability. For example, scale-free graphs have a small number of very important hubs whose occlusion would be detrimental. Labels are important in much of cartography; various heuristics are used to manage label visibility and placement [32], [33].

Subgraphs and selections. To understand subgraphs, a simple approach superimposes the graph subset over a dimmed full graph, leveraging Gestalt figure-ground separation [34].

3) Zoomable graphs for overview and context: Overview plus context via zoom is a common pattern in visualization. Some recommendations suggest including a separate small overview [35], but this may not be required if the information space includes cues to aid navigation [29]. Zooming interfaces are a simple approach for focus plus context in graph analysis (i.e. accessing detail situated with a larger context) particularly when zooming is rapid and easy and working memory is not overloaded. Fast, smooth animated pan and zoom is key to maintaining consistent identification and location of objects and landmark features in context [27] [36]. In practice, fast zoom is integral to online maps, which have migrated from from image-based tile maps to faster vector-based maps [37].

III. Requirements, design and development

The requirements for this work were guided by key research questions central to the Defense Advanced Research Projects Agency (DARPA) Resilient Supply and Demand Networks (RSDN) program and were developed through initial discussions with Subject Matter Experts (SMEs) in supply chain risk, economic analysis, and systemic resilience. Discussions identified the critical analytical needs of risk analysts, which informed the design of our graph-based visualization approach.

A. User analytical needs

Risk analysts require tools for multi-scale exploration, pattern recognition, and interactive risk assessment. They must:

1) **Validate data and network structure**: Analysts must confirm that a derived network aligns with expert expectations, ensuring that key suppliers, industry groupings, and critical dependencies are contained before conducting risk assessment. For example: *Does the data-derived network match expert expectations? Are expected groupings and suppliers present?*
2) **Assess systemic risk**: Analysts must identify weak points and high-risk dependencies by examining failure spread across interconnected supply chains. Understanding the broader impact of disruptions beyond individual entities is essential. For example: *Where are the weak points in the network, and what is their immediate neighborhood?*
3) **Analyze scenario-based risk**: Analysts must be able to model potential disruptions (e.g., geopolitical events, supply shortages, or economic shocks) to anticipate their impact across industries and supply chains. For example: *What is the impact of a sanction, merger, natural disaster, or export restriction for a specific industry?*
4) **Determine systemic risk mitigation strategies**: Analysts need to evaluate interventions and alternatives, such as switching suppliers, modifying trade routes, or reinforcing critical network points to minimize risk exposure. For example: *Given alternative suppliers to choose from, which option reduces overall risk in the network?*

B. Graph-based visualization approach

To effectively support these analytical needs, we hypothesized that a graph visualization approach, integrated within a broader visual analytics system, would provide the best means of exploring and interpreting large-scale networks. Our focus is on supply chain networks, where nodes represent companies or facilities; hierarchical groupings correspond to industries, geopolitical hierarchies, or product hierarchies; and links capture relationships between companies, such as contracts or the flow of goods. Graph-based representations reveal hidden structures, dependencies, and propagation effects, allowing analysts to investigate how disruptions impact supply-demand relationships. However, the sheer number of entities and connections (tens of thousands to hundreds of thousands of nodes) in global company-level supply chain networks can overwhelm traditional visualization methods [38].

To enhance perceptual scalability and analytical efficiency, we incorporate hierarchical grouping, clustering, edge bundling, and labelling (coupled with fast pan and zoom). These techniques enable users to navigate, interpret, and analyze complex network structures without losing clarity. By balancing high-level overviews with granular detail [38] our approach supports the full range of analytical tasks needed for systemic risk assessment, including:

1) **Comprehensive graph overview**: The ability to rapidly grasp the overall structure of the network, identify key regions of interest, and detect high-level patterns.
2) **Intuitive navigation**: Seamless graph exploration at multiple levels, from a high-level summary to specific entities and connections, without adding to cognitive load.
3) **Detailed entity inspection**: The ability to drill down on individual nodes and edges to examine their attributes, relationships, and role in the network.
4) **Subgraph assessments**: Tools to isolate, highlight, and analyze portions of the graph to evaluate local risks, dependencies, and vulnerabilities within a broader system.

While graph visualization is key to our approach, it operates in a larger system that supports data extraction, search, filtering, selection, geographic mapping, and analytical management (Fig. 2). These ensure seamless transition between data validation, systemic risk assessment, and mitigation planning.

C. Design and prototyping

1) Hierarchical node-and-link. In system designs, we considered maps, treemaps, and linear flows. We implemented node-link diagrams because of their familiarity to our users, and their individual components—dots, size, color, containment and proximity— were easy to explain, facilitating learning.

Our **node hierarchy** aligns with expert user conceptualization of low-level facilities and companies and high-level industries, goods, or countries. These hierarchical groupings are supported in data—Global Industry Classification Standard (GICS) for grouping companies into industries, internationally standardized Harmonized System (HS) codes for classifying traded products when determining tariffs, or geopolitical administration for grouping regions and countries.

Within node-link diagrams, hierarchies can be represented as trees, treemaps, circle-packing, or other layouts. We prototyped hierarchically-grouped Sankey-style linear left-to-right linkages, but real-world graphs did not conform to the user-expected directionality. Attempting to force these layouts obscured key lateral relationships, demonstrating that layout selection could not be solved independently of graph structure and task intent. We also explored treemaps, but their tight packing introduced misleading spatial relationships and made it difficult to incorporate edge connectivity—highlighting an early tradeoff between space efficiency and structural fidelity.

We applied clustering via force-directed layouts for each hierarchical group of nodes, then placed each group into a containing circle, and repeating this up the hierarchy. This creates a layout which visually resembles **hierarchical circle-packing**. It enabled us to position closely related entities near one another while maintaining separation between less-related groups unlike circle packing [39]. Containing circles encode metadata analogous to the leaf nodes—dot size reflects magnitude (e.g., trade value), ring size corresponds to the number of enclosed entities, and ring color maps to normalized aggregate metrics such as risk score. While the hierarchy provides semantic grouping, clustering adds structural cues by using whitespace to convey relative proximity (as can be seen in

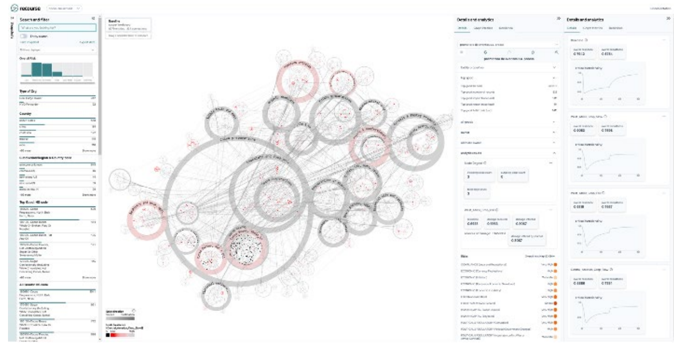

Fig. 2. Supply chain systemic risk system with a graph visualization, supporting components for search, filtering and profiling (left), detail inspection (right), and systemic/scenario risk modeling (far right).

the layout of rings within a larger ring, or dots within a ring). Integrating the two required careful balancing: clustering needed to surface emergent patterns without visually overpowering the meaningful groupings defined by the hierarchy. To make this feasible at scale, we apply force-directed clustering of peers separately within each hierarchical level, significantly reducing computational demands (as opposed to a single full graph force-directed layout). Even so, for very large graphs (e.g., >10,000 nodes), current layout generation occurs at non-interactive speeds (tens of minutes). As this is one-time calculation per dataset, we have not parallelized it at this time.

Fig. 3 shows a graph with individual entities as small dots, successive levels of hierarchy as rings, with the largest visible rings labelled. Dot and peer ring adjacencies use force-directed clustering, revealing closer relationships. This depiction indicates the strongest relative relationships even if the edges are not visible (Fig. 4). Ring sizes and white space result in unique regions, facilitating location and orientation in the graph (e.g. near the big circle or the circle with three equally sized children).

Fig. 4 shows how different risk analyses generate different results, and how these are evident in a copper supply and demand network. In the top left, traditional risk metrics (e.g. legal, economic, labor, political) identify the aluminum sector being the highest risk. In the lower left, simple country affiliation risks show electrical machinery as the highest risk concentration. The right side shows systemic risk and scenario risk—revealing risks to the network not surfaced with the traditional models. In the upper right, a systemic risk analysis identifies a weak spot in iron. Finally, in the lower right, a scenario risk analysis results in high impact in refined copper.

2) Edge analysis: Edges are a primary source of visual overload, especially in large graphs where dense connectivity dominates node and group perception. Simply reducing edge count is insufficient without undermining analytical tasks like tracing propagation paths.

Hierarchical semantic bundles. Traditional edge-bundling algorithms group edges by spatial proximity. Instead, we leverage the data hierarchy itself, first creating a stable hierarchy of nested graphs, then aggregating quantifiable flows at each level [18]. This forms semantically meaningful, layout-independent bundling that reflects the domain structure. Bundles are computed through each hierarchical level and visualized as

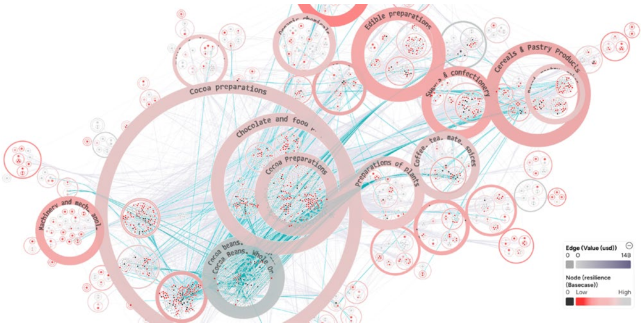

Fig. 3. Overview and systemic risk of 9,000 companies in the global cocoa supply chain. Individual companies (dots) are grouped in product rings. Supplies flow (grey edges) from bottom left to top right. Red saturation show greater weakness in the network. Selections (green) reveal two downstream levels of connections to exporters of Cocoa Beans from one country.

low-curvature paths with directional flow encoded as a color gradient. This hierarchy-aware approach ensures stability and interpretability—bundles retain meaning as users zoom, filter, or navigate the graph. This required deliberate coordination with the layout: we used hierarchical circle-packing and clustering to position nodes and drew bundles accordingly. This tight coupling between bundling and hierarchy—rather than layout alone—preserves the analytical utility of flow tracing at scale. Fig. 3 shows the cocoa supply chain with bundled flows highlighting dominant left-to-right propagation.

Hierarchical edge levels. Because our bundling approach is driven by the underlying data hierarchy rather than layout geometry, it supports dynamic edge selection and filtering by any level of hierarchy without re-bundling. Analysts can select and view high-level relationships (e.g., industry-to-industry), intermediate groupings (e.g., sub-industries), or detailed leaf-node connections (e.g., company-to-company), all while preserving visual coherence. This flexibility is made possible by the stability of the hierarchical bundling: edges are precomputed through each level of the nested graph, allowing fast switching and consistent interpretation. As shown in Fig. 5, this multilevel edge rendering helps analysts explore connectivity across scales without sacrificing clarity or context.

3) Subgraphs: are essential for addressing specific analytic tasks such as identifying particular entities in the graph and related connections upstream or downstream. In designs and prototypes, we explored juxtaposing side-by side subgraphs but ultimately chose superimposed subgraphs. This decision reflected a tradeoff: juxtaposition supports comparison but fragments context, whereas superposition risks clutter. We mitigated this through saturation shifts and boundary highlighting to reinforce foreground-background separation.

Selection: search, facets, expansion. To analyze what-if scenarios—such as mergers, sanctions, or regional disruptions—users must easily define and manipulate entity selections. Our system supports search, faceted filtering, and topological expansion, enabling analysts to construct targeted subgraphs. Starting from a seed (e.g., a company(s) or group(s)), users can expand selections hierarchically (e.g., all children within a group) or relationally (e.g., next-level upstream or downstream entities). These expansions surface dependency

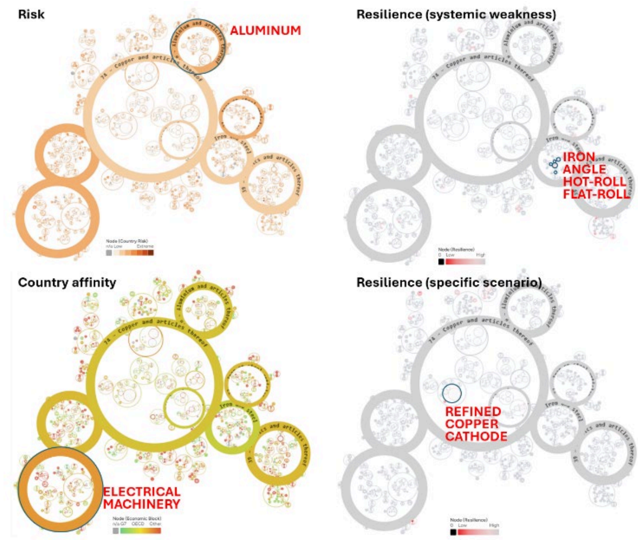

Fig. 4. Copper industry network analysis comparing traditional risk assessments (left) with network-based risk models (right). The left side illustrates risk concentrations based on predefined metrics: darker red indicates higher risk, while lighter colors represent lower risk. The right side presents systemic and scenario-based resilience assessments, identifying weak points in the network and potential cascading failures. Labels highlight the most affected commodities or industries under different risk perspectives.

chains and exposure pathways critical for scenario modeling. As shown in Fig. 6, users can progressively explore levels of connection— starting with a few antimony processing companies and progressing through successive levels of downstream connections. Highlighted rings of selected entities aid quick location of entities of interest and identify the industries in which they are contained.

This approach helps analysts understand the up and downstream ecosystem surrounding an entity or product. For instance, it enables tracing successive customer dependencies from a mine or mapping all required inputs back to raw materials for a specific product, such as an air conditioner. These topological expansions often reveal hidden or non-obvious interdependencies within the network. In Fig. 7, we highlight upstream paths for three seemingly unrelated copper-based products. While ammunition (middle) shows distinct sourcing, overlapping upstream structures in the supply chains for air conditioners (left) and electrical transformers (right) indicate shared vulnerabilities that might otherwise go unnoticed.

4) Labels: are critically important to orientation and identification. They function as persistent perceptual anchors, enabling users to maintain mental continuity as they explore between abstraction levels [36]. Rather than relying on slow, directed, motor-based interactions (e.g. tooltips or drill-down), labels can be accessed with low attention and read automatically [3], thereby providing low-effort landmarks. This required balancing perceptual accessibility with performance constraints, particularly when dynamically determining which labels to show at each zoom level in real time.

Label visibility. With hundreds of thousands of potential labels, we need to determine the most relevant ones. We explicitly wanted the largest elements displayed at any zoom

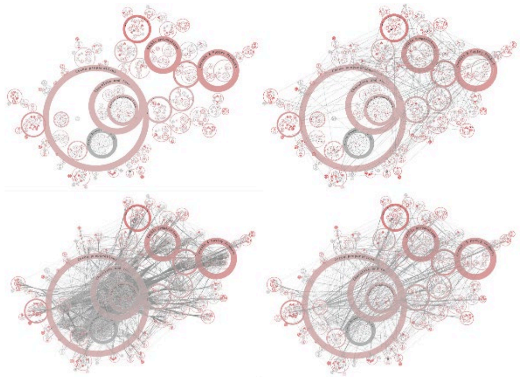

Fig. 5. Cocoa industry graph with successive levels of edges. Clockwise from top-left: none; connecting top-level industries; connecting mid-level industries; connections between individual companies).

level to be labelled (much like a political map or roadmap). We associate labels with nodes and rings and display labels for the largest rings within the current view and zoom level (Fig. 8). This zoom-sensitive visibility provides a scalable alternative to static labeling, but required careful tuning to avoid visual instability and ensure semantic relevance across transitions.

Minimizing node occlusion. Labels that occlude nodes can be source of critical errors in a visual analysis task. To prevent this we overlaid curved labels within rings so that they do not interfere with nodes (Fig. 3). This also alleviates potential confusion as to which entity the label is associated with [40].

Ring thickness and labels. We smoothly increase ring width as labels reach the threshold of visibility, embedding labels within the rings instead of above them. Compared to consistently thin rings, thick, visually dominant rings draw attention to labeled regions.

Zoom level encoding via label specificity. Label content signals zoom depth; more specific labels appear at deeper zoom levels. Users, expert or novice, can intrinsically apply real-world knowledge to aid orientation (Fig. 8). For example, *Waffles* is more specific than *Cereals & Pastry* (HS codes); *Software* is more specific than *Information Technology* (GICS); and *Rotterdam* and *The Hague* are more specific than *European Union*. This layered specificity ensures users retain a meaningful frame of reference—reducing cognitive load and avoiding disorientation common in large network navigation.

Short labels. Quick orientation is better aided by short text. HS codes can be highly descriptive: *Electrical machinery and equipment and parts thereof* or *sound recorders and reproducers, television image and sound recorders and reproducers, and parts and accessories of such articles*. In these cases, we truncated or used an LLM to reduce labels to the most relevant words (e.g. *Electrical machinery…*).

Labels have other limitations. Language compatibility was not an issue, as our target system was designed by and for English speakers. The readability of curved text labels was identified as a possible perceptual issue, but no users reported such difficulties.

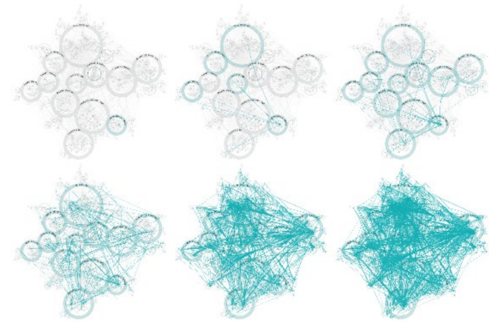

Fig. 6. Successive levels of connection through the antimony supply demand network.

5) Fast zoom: The entire design is predicated on fast zoom to maintain orientation within the larger graph [29] [32]. This choice constrained nearly every aspect of our implementation, from rendering technology to encoding strategies, as the system needed to maintain fluid responsiveness without sacrificing clarity. A tile-based rendering approach, wherein new tiles at successive levels of zoom are loaded and cross-faded, was considered too slow. Instead, we implemented a WebGL vector-based approach, preprocessing as much as possible and putting this into the GPU scenegraph. For example, the appearance/disappearance of labels on zoom is implemented within GL to maintain high-speed performance on interaction. While this requires a longer initial load time, it allows smooth, uninterrupted interaction without dynamically loading or managing children, edges, or labels during pan and zoom.

We evaluated the system on a large-scale real-world dataset with more than 100,000 nodes and 500,000 edges. It consistently maintained frame rates above 50 frames per second (fps). Additional scalability tests with synthetic data modeled on real-world graph properties confirmed the system remains performant (above 50 fps during both panning and zooming) with up to 3 million nodes and 3 million edges. This performance is essential not just for usability, but for preserving context across interactions—allowing users to visually track targets and changes without interruption or re-orientation.

By contrast, many open-source JavaScript graphing libraries struggle with interactivity beyond 1 million nodes. Our approach does incur a one-time loading cost due to data transfer to the GPU and the preprocessing of label metadata. However, once loaded, the system remains efficient and responsive throughout analysis (i.e. >50fps). Even at million-node scale, the initial load time remains under one minute. This delay can be mitigated through user interface cues that indicate loading status and allow navigation to begin as soon as interaction becomes available.

IV. Case Studies

The system has been used to respond to quick risk assessment analysis (within one business day). While it supports a wide variety of analytic questions, most analyses to date have examined what-if shocks. These assessments have been estimated to require a week or more using previous techniques (a data science workflow using tools such as Jupyter notebooks).

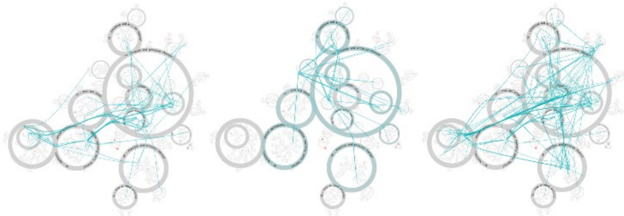

Fig. 7. Copper network with two level upstream subgraphs for air conditioners; ammunition; and electrical transformers highlighted.

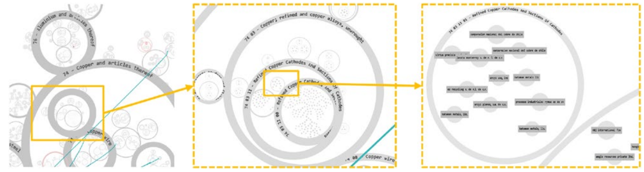

Fig. 8. Labels fade in on smaller items with increasing specificity on zoom, from left: 74 Copper and articles thereof, to 7403 Copper refined and copper alloys, to 74031101 Refined copper cathodes and sections of cathodes, and to individually labelled nodes.

The analyst receives a request to model a hypothetical scenario to model such as *A) The consequences of a production halt at a major copper mine; B) The effects of sanctions against specific copper product manufacturers; C) The breadth of industries affected by a specialized copper product requiring a critical material in another country; and D) The downstream effects of a cocoa crop failure in a region. Given these modeling tasks, the analyst uses our system as described below.*

1. Create an industry network. Scenarios A–C, require a global copper network, while D requires a global cocoa network. The system provides rudimentary tools for extracting a desired industry network. For example, the analyst could search for target industry codes or line-item terms in a repository of more than 2 billion documents and then expand +/- a few levels of connectivity from all the matches. Various parameters and thresholds could be set on this retrieval (e.g., only relationships with a value of more than $10m USD). The result would typically be a network from 10,000–100,000 or more entities.

2. Validate the network. The source data is incomplete, and there is no single definitive source of ground truth for trade activities. Even at the level of country-to-country trade collected by national statistical agencies there are mismatches from each country's' perspective [41]. At a company-to-company level, data is derived from sources such as shipping documents (e.g. bills-of-lading are available from some countries), corporate reports (e.g. the U.S. requires companies to report key suppliers and customers over a threshold), government platforms (e.g. USAspending.gov), and various other sources.

As such, the analyst needs to validate the extracted network. For example, at a high level, they need to verify that the expected industries exist. A copper network (Fig. 7) should contain air conditioners (made with much copper tubing) and electrical transformers (made with much copper wire). Similarly, visible industries should match expectations. A quick scan of Fig. 3 reveals a large cocoa ring (the core of the query) and rings such as Sugar & Confectionery; Coffee, Tea, etc; Cereals & Pastries, etc., all of which are expected based on real-world knowledge. Anomalies are also evident: in an initial version of the copper network, unexpected industries such as concrete and food are present due to logistics agencies—which ship nearly every kind of good—reported as end-customers. These can be subsequently removed on successive iterations. The analyst can further investigate by comparing entities to Internet searches. Minor anomalies can be individually removed. At this point, the analyst has a high-level verification of the network.

3. Network sensitivity can be explored. A quick proxy described by a banking risk manager is to skim the network for entities either too big to fail or too connected to fail. Corresponding data attributes (e.g. trade value, node degree) can be mapped to entity size and saturation and easily perceived within the context of larger network.

Systemic risk can assess resilience across the network structure [7]. The analyst can run a Monte Carlo of 100 simulations, typically in a minute. This calculates normalized damage per node (percent production retained) and aggregated per ring, thereby drawing attention to larger, redder rings (Fig. 3). *Sugar & Confectionery* and *Cereals & Pastries* have more saturated red than *Coffee, Tea, etc.*, meaning that the first two are more exposed to systemic risk in this network than the latter.

4. Set up scenario. After reviewing the overall network, the analyst can begin to investigate the specifics of the scenario. They can search for identified companies (scenario A–B); filter products (scenario C); or select and filter criteria (scenario D).

These selections can be expanded up or downstream for a fast qualitative analysis (e.g. *how connected are these entities to the larger network, within a few steps*). As the initial problem framing might have ambiguity, variations on the selection can be considered, such as adjacent products or related countries.

Selections form the basis for configuring the scenario. Analysts can set initially impacted subgraph (e.g. mine to close, companies to sanction, regions where crops fail) and an optional subset upon which to investigate the impact. Each model has additional parameters that the analyst can configure via a model-driven user-interface, such as how long the impact lasts, options for entity reaction and recovery to failure, and methods for estimating inventory. The analyst then runs the scenario, which might take a minute to hours depending on the model.

5. Scenario impact analysis. Models can generate large amounts of data. The model panel (Fig. 2, far right) shows overall summary model metrics while the large graph visualization shows aggregations through the hierarchy down to the nodes. Users choose model metrics of interest to display, such as robustness (the lowest level of output during the scenario) or resilience (the percent of production retained during the scenario)—which then appear as visual attributes such as saturation (Fig. 2). As there might be a range of values produced, detailed results of the simulation are available in the side panels.

From these, the analyst can assess the network impact. In Fig.7 (bottom-right), the mine closure has the most impact on refined copper. The analyst can also compare relative impact across different industries at the level of constituents (Fig. 9).

The analyst then assembled the snapshots from the above analysis to produce a report that indicates the degree of impact

Fig. 9. An impact result reveals a different distribution of green vs light green nodes in the two larger rings.

that the given scenario has created across the industry, which industries are most impacted, and the affected constituents.

Comments from the stakeholders requesting these reports include “This provides a level of technical due diligence that can aid our strategic decision-making to improve resilience,” and “These analyses provide great short-term impacts. But these results will be useful for institutions that can adaptively manage risk over long time horizons: short term risks can be absorbed, and long-term resilience can be the outcome from these capabilities more broadly applied.”

V. Evaluation Study

A. Method and procedure

A user evaluation was conducted to assess the usability and effectiveness of our system, focusing on its ability to support entity-level, systemic, and scenario risk analysis in supply-demand networks. Six subject matter experts with backgrounds in supply chain management, economic analysis, and finance participated. The virtual evaluation was conducted via Microsoft Teams, where participants were guided through structured tasks designed to reflect real-world analytical workflows.

Each session lasted approximately 90 minutes and followed a structured format. Participants first received an introduction and training using a simplified dataset, where they learned how to navigate the interface, locate entities of interest, and interpret network visualizations. Following the training, participants completed six tasks that evaluated the system’s ability to support different types of risk analysis. For entity-level risk, participants were asked to locate a specific company, such as “Acme Corporation,” and identify its market size, industry group, and geographic location. This task assessed how effectively the tool enabled users to search, filter, and drill down into individual entity details. For systemic risk, participants were given a scenario in which a geopolitical conflict impacted 5% of the supply chain and were asked to assess how this disruption cascaded through the network, ultimately identifying the most vulnerable industry. This task tested the system’s ability to facilitate network-wide analysis of failure propagation. Lastly, for scenario risk, participants simulated an economic recession that further damaged an additional 10% of the network and were asked to determine which industries were most affected and how different interventions could mitigate the impact. Upon completing the tasks, participants were asked to rate their experience using a 5-point Likert-scale questionnaire, evaluating the system’s clarity, efficiency, and usability. They

TABLE I. Ratings from Usefulness Questionnaire

Task	Question	M	SD	Min	Max
1	How well does interface support finding entities of interest?	4.7	0.5	4.0	5.0
2	How well does the interface allow them to assess resilience?	4.6	0.8	3.0	5.0
3	How well does the interface support finding the weakest companies in the SDN?	4.4	0.8	3.0	5.0
4	How usable do you find the edge bundling feature for improving SDN clarity?	4.9	0.4	4.0	5.0
5	How helpful is industry code grouping?	4.6	0.5	4.0	5.0
5	How helpful is location code grouping?	4.2	0.4	4.0	5.0
5	How helpful is no-grouping option?	3.8	1.1	3.0	5.0
5	How well did grouping feature support exploring SDN?	4.8	0.4	4.0	5.0
6	How usable did you find the various data segmentation features in simplifying the ability to search and filter to the	4.4	0.5	4.0	5.0
6	How useful did you find the search bar for finding entites matching your terms?	4.9	0.4	4.0	5.0
6	How useful did you find the data profiling facets for examining data breakdowns and highlighting them in the	4.6	0.5	4.0	5.0
6	How useful did you find the viewing options (view only/hide)?	4.7	0.5	4.0	5.0

also provided qualitative feedback through a post-evaluation discussion, where they reflected on strengths, challenges, and potential areas for improvement.

B. Results

Every participant successfully completed each task, often requiring minimal guidance. Survey responses were statistically analyzed to determine mean ratings and standard deviations (see Table 1). These ratings were highly positive, though potentially skewed by the small group context.

Written comments and discussion notes were thematically analyzed to identify recurring feedback and improvement suggestions. The following sections detail how participant insights highlighted the tool’s strengths in facilitating efficient entity-level exploration, systemic risk analysis, and interactive network-based decision-making.

Entity-Level Analysis: Participants found locating and drilling down into specific entities to be seamless and efficient. Using the network map and search capabilities, they were able to quickly identify companies, countries, and industry groups of interest. Users described the interface as "very easy," "very clear," and "very intuitive", with minimal effort required to explore entity details and relationships. The ability to group and filter entities based on industry code and location was seen as a valuable feature, enabling targeted investigation of competitive landscapes and supply chain structures. Some participants suggested enhancements, such as Google Maps-style zooming to improve entity searches by dynamically highlighting relevant results, which has since been added to the system.

Systemic Risk Analysis: The tool’s resilience and robustness analytics provided valuable insights into network-wide risk propagation, allowing participants to assess how disruptions impact interconnected industries. Users were able to correctly interpret robustness scores and identify the weakest points in the network, even though some were initially unfamiliar with systemic risk modeling. Training significantly improved participant understanding, demonstrating the system’s accessibility for experts and non-experts. While users successfully found industries with the lowest resilience, they suggested refinements to resilience scoring explanations and visual representations; improving color contrast, clarifying risk scales, and incorporating motion-based visual cues could enhance interpretability of cascading failures and failure paths.

Visualization and User Interaction: Edge bundling received strong positive feedback, with users describing it as "very helpful" and "highly illustrative" for simplifying dense network graphs. The ability to toggle bundling on and off was particularly well received, allowing analysts to switch between a high-level overview and a detailed network examination. We hypothesize that un-bundled edges may be valuable for direct connections, which are otherwise lost when bundled together, but we have not tested this. Grouping by industry code and/or geographic location was considered valuable for enhancing the interpretability of complex supply-demand networks. While most participants preferred grouped views, some valued the no-grouping option for a raw, unstructured network perspective. Data filtering capabilities were seen as useful, though some users found advanced segmentation features complex and suggested more guided interactions or predefined filtering templates.

VI. Conclusion

Designing effective visualizations for large networks is challenging due to perceptual scalability issues. Techniques like hierarchical grouping, edge bundling, and level-of-detail labeling combined with high-performance interaction are crucial to reduce clutter and improve interpretability. We showed that we could achieve an effective balance in very large graphs supporting tasks requiring high-level and low-level perception of metrics, subgraphs, and identification throughout. There is much to consider for future research for perceptual tuning and scaling of graph visualizations.

Fewer edges. We experimented on multiple projects with reducing or removing visible links on overviews and only displaying them on interaction. Anecdotally, we are concerned that when no links are displayed, users do not immediately understand they are looking at a graph, with potential confusion that the display may be a scatterplot or different projection. What is the minimal number of edges that should be displayed and how might those edges be the most relevant edges?

More versatile rings. Why do they need to be rings? Can they support other groupings, such as Euler intersections? Or other arbitrary layouts, such as a grid (e.g. country x industry). Will they work with other clustering algorithms instead of known hierarchies—a challenge will be finding suitable labels to aid orientation when using machine-driven multi-attribute clustering. Users questioned whether rings could be used as a heatmap or paired with contour heatmaps to help draw attention to areas of problems in the network.

Increased label density. We show few labels compared to roadmaps and label-dense cartography from which we take inspiration. How many labels could be displayed and be effective? There are other graph visualizations that use many more labels, but when all labels are displayed with the same font attributes, none take precedence, making it harder to know which to read first. All our labels are the same font size and style. whereas maps make heavy use of font attributes to encode data.

Multiple levels of highlighting: We overload color: a) we use a glow to indicate mouse overs; b) a specific hue to indicate selection of lines and nodes (while decreasing opacity on the background); c) fill color/line saturation to indicate data attributes; and d) two different unique hues indicate different selections (currently incomplete). Inherently, color is a strong cue that can be challenging when overloaded. The tiny size of marks in the overview precludes attributes such as shape or texture, and a requirement for static reports precludes motion.

Identifying communities and representations: We believe that there may be additional use for community detection to find groupings within the network that do not follow the organized hierarchies. We note that some companies have many facilities, which, in turn may be situated in different industries. The relation of those companies could be depicted by grouping, using edges, or other techniques from set visualization [27].

Greater scalability. The core zoom/pan implementation scales well in terms of performance based on our tests. Ideally, the approach should be tested on graphs of much larger scale. Incorporation of vector tiles could lead to higher performance and more flexibility with depiction of local details, either via enhanced labelling or other kinds of landmarks. In addition to zoom, we hypothesize that expand/contract could effectively support greater scalability as some portions of the graph may not be relevant at some times.

ACKNOWLEDGMENT

Many people have contributed to this work including Ahmed Mostafa, Rob King, and David Giesbrecht. The authors thank all RSDN collaborators for their support and encouragement.

References

[1] D. Jonker, S. Langevin, D. Giesbrecht, M. Crouch, and N. Kronenfeld, "Graph mapping: Multi-scale community visualization of massive graph data," *Information Visualization*, vol. 16, no. 3, pp. 190–204, 2017.

[2] A. Perer and B. Shneiderman, “Balancing systematic and flexible exploration of social networks,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 693–700, 2006.

[3] R. Brath, *Visualizing with text*. Boca Raton, FL: CRC Press, first edition ed., 2021.

[4] B. Shneiderman and A. Aris, “Network visualization by semantic substrates,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 733–740, 2006.

[5] S. J. Morani, “Supply chain risk management taxonomy version 2.0,” Jan. 2025.

[6] Office of the Assistant Secretary of Defense for Sustainment (OASD(S)), “DoD SCRM taxonomy version 2.0,” Jan. 2025.

[7] A. A. Ganin, E. Massaro, A. Gutfraind, N. Steen, J. M. Keisler, A. Kott, R. Mangoubi, and I. Linkov, “Operational resilience: concepts, design and analysis,” *Scientific Reports*, vol. 6, p. 19540, Jan. 2016.

[8] A.-L. Barabasi and M. P´ osfai,´ *Network science*. Cambridge: Cambridge university press, 2016.

[9] L. C. Freeman, “A set of measures of centrality based on betweenness,” *Sociometry*, vol. 40, no. 1, pp. 35–41, 1977.

[10] The National Academies, Policy and Global Affairs, Committee on Science, Engineering, and Public Policy, and Committee on Increasing National Resilience to Hazards and Disasters, *Disaster resilience: national imperative*. Washington: National Academies Press, 2012.

[11] I. Linkov, B. D. Trump, M. Golan, and J. M. Keisler, “Enhancing resilience in post-covid societies: by design or by intervention?,” *Environmental Science & Technology*, vol. 55, no. 8, pp. 4202–4204, 2021. PMID: 33739817.

[12] J. Geanakoplos, R. Axtell, J. D. Farmer, P. Howitt, B. Conlee, J. Goldstein, M. Hendrey, N. M. Palmer, and C.-Y. Yang, “Getting at systemic risk via an agent-based model of the housing market,” *American Economic Review*, vol. 102, p. 53–58, May 2012.

[13] G. Van Voorn, G. Hengeveld, and J. Verhagen, “An agent based model representation to assess resilience and efficiency of food supply chains,” *PLOS ONE*, vol. 15, pp. 1–27, 11 2020.

[14] C. M. Macal, “Everything you need to know about agent-based modelling and simulation,” *Journal of Simulation*, vol. 10, no. 2, pp. 144– 156, 2016.

[15] D. Bisias, M. Flood, A. W. Lo, and S. Valavanis, “A survey of systemic risk analytics,” *Annual Review of Financial Economics*, vol. 4, no. Volume 4, 2012, pp. 255–296, 2012.

[16] K. Klein, S. Kobourov, B. E. Rogowitz, D. Szafir, and J. Miller, “Perception in network visualization (Dagstuhl Seminar 23051),” *Dagstuhl Reports*, vol. 13, no. 1, pp. 216–244, 2023.

[17] C. Vehlow, F. Beck, and D. Weiskopf, “The state of the art in visualizing group structures in graphs,” in *Eurographics Conference on Visualization (EuroVis) - STARs* (R. Borgo, F. Ganovelli, and I. Viola, eds.), The Eurographics Association, 2015.

[18] D. Holten, “Hierarchical edge bundles: visualization of adjacency relations in hierarchical data,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 741–748, 2006.

[19] S. Hadlak, H. Schumann, and H.-J. Schulz, “A survey of multi-faceted graph visualization,” in *Eurographics Conference on Visualization (EuroVis) - STARs* (R. Borgo, F. Ganovelli, and I. Viola, eds.), The Eurographics Association, 2015.

[20] U. Dogrusoz, E. Giral, A. Cetintas, A. Civril, and E. Demir, “A compound graph layout algorithm for biological pathways,” in *Graph Drawing* (J. Pach, ed.), (Berlin, Heidelberg), pp. 442–447, Springer Berlin Heidelberg, 2005.

[21] M. Balzer and O. Deussen, “Level-of-detail visualization of clustered graph layouts,” in *2007 6th International Asia-Pacific Symposium on Visualization*, pp. 133–140, 2007.

[22] M. Girvan and M. E. J. Newman, “Community structure in social and biological networks,” *Proceedings of the National Academy of Sciences*, vol. 99, no. 12, pp. 7821–7826, 2002.

[23] E. R. Gansner and S. C. North, “Improved force-directed layouts,” in *Graph Drawing* (S. H. Whitesides, ed.), (Berlin, Heidelberg), pp. 364–373, Springer Berlin Heidelberg, 1998.

[24] S. Chaturvedi, C. Dunne, Z. Ashktorab, R. Zachariah, and B. Shneiderman, “Group-in-a-box meta-layouts for topological clusters and attribute-based groups: space-efficient visualizations of network communities and their ties,” *Computer Graphics Forum*, vol. 33, no. 8, pp. 52– 68, 2014.

[25] A. Lhuillier, C. Hurter, and A. Telea, “State of the art in edge and trail bundling techniques,” *Computer Graphics Forum*, vol. 36, no. 3, pp. 619–645, 2017.

[26] C. G. Healey, “Perception in visualization.”

[27] B. Alsallakh, L. Micallef, W. Aigner, H. Hauser, S. Miksch, and P. Rodgers, “Visualizing sets and set-typed data: state-of-the-art and future challenges,” in *EuroVis - STARs* (R. Borgo, R. Maciejewski, and I. Viola, eds.), The Eurographics Association, 2014.

[28] J. Raskin, The humane interface: new directions for designing interactive systems. Reading, Mass: Addison-Wesley Professional, 2000.

[29] C. Ware, *Information visualization: perception for design*. Waltham, MA: Morgan Kaufmann, 3d edition ed., 2013.

[30] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker, “Cytoscape: a software environment for integrated models of biomolecular interaction networks,” *Genome Research*, vol. 13, no. 11, pp. 2498–2504, 2003.

[31] M. Bastian, S. Heymann, and M. Jacomy, “Gephi: an open source software for exploring and manipulating networks,” *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 3, pp. 361–362, Mar. 2009.

[32] H. Reiterer and T. Buring, “Zooming techniques,” in¨ *Encyclopedia of Database Systems* (L. Liu and M. T. Ozsu, eds.), pp. 3684–3689, Berlin:¨ Springer, 2009.

[33] J. A. Tyner, *Principles of map design*. New York: Guilford Press, 2010.

[34] D. Todorovic, “Gestalt principles,” *Scholarpedia*, vol. 3, no. 12, p. 5345, 2008. revision #91314.

[35] K. Hornbæk, B. B. Bederson, and C. Plaisant, “Navigation patterns and usability of zoomable user interfaces with and without an overview,” *ACM Trans. Comput.-Hum. Interact.*, vol. 9, p. 362–389, Dec. 2002.

[36] A. Cockburn, A. Karlson, and B. B. Bederson, “A review of overview+detail, zooming, and focus+context interfaces,” *ACM Comput. Surv.*, vol. 41, Jan. 2009.

[37] A. Miller, “Under the hood of Google Maps 5.0 for Android,” Dec. 2010.

[38] B. Shneiderman, “The eyes have it: a task by data type taxonomy for information visualizations,” in *The Craft of Information Visualization* (B. B. Bederson and B. Shneiderman, eds.), Interactive Technologies, pp. 364–371, San Francisco: Morgan Kaufmann, 2003.

[39] W. Wang, H. Wang, G. Dai, and H. Wang, “Visualization of large hierarchical data by circle packing,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’06, (New York, NY, USA), p. 517–520, Association for Computing Machinery, 2006.

[40] M. Reckziegel, L. Pfeiffer, C. Heine, and S. Janicke, “Modeling how¨ humans judge dot-label relations in point cloud visualizations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 6, pp. 2144–2155, 2020.

[41] E. Ortiz-Ospina, B. Rohenkohl, V. Samborska, S. V. Teutem, D. Beltekian, and M. Roser, “Trade and globalization,” *Our World in Data*, 2018. https://ourworldindata.org/trade-and-globalization