

# Knowledge Graphs in Practice: Characterizing their Users, Challenges, and Visualization Opportunities

Harry Li , Gabriel Appleby , Camelia Daniela Brumar , Remco Chang , Ashley Suh 

**Abstract**— This study presents insights from interviews with nineteen Knowledge Graph (KG) practitioners who work in both enterprise and academic settings on a wide variety of use cases. Through this study, we identify critical challenges experienced by KG practitioners when creating, exploring, and analyzing KGs that could be alleviated through visualization design. Our findings reveal three major personas among KG practitioners – KG Builders, Analysts, and Consumers – each of whom have their own distinct expertise and needs. We discover that KG Builders would benefit from schema enforcers, while KG Analysts need customizable query builders that provide interim query results. For KG Consumers, we identify a lack of efficacy for node-link diagrams, and the need for tailored domain-specific visualizations to promote KG adoption and comprehension. Lastly, we find that implementing KGs effectively in practice requires both technical and social solutions that are not addressed with current tools, technologies, and collaborative workflows. From the analysis of our interviews, we distill several visualization research directions to improve KG usability, including knowledge cards that balance digestibility and discoverability, timeline views to track temporal changes, interfaces that support organic discovery, and semantic explanations for AI and machine learning predictions.

**Index Terms**—Knowledge graphs, visualization techniques and methodologies, human factors, visual communication

## 1 INTRODUCTION

Knowledge graphs have emerged as a popular approach to represent and manage complex data from a variety of domains [1]. Due to their ability to encode semantically rich information in the form of entities and the relationships between them, knowledge graphs are now an industry standard for data unification [80], question-answering [41], recommendation systems [32], explainable AI [47], and many other practical applications [38]. However, despite the growing popularity of knowledge graphs (KGs), there is a limited understanding of the *types* of KG users, the challenges they face, and the limitations of current tools and visualization designs for KGs used in practice.

To address this gap, we conducted an interview study with 19 KG practitioners across eight different organizations. The participants of our study come from a broad background of both industry and academic experiences, representing a diverse set of domains – including biology, finance, health, drug development, software, cybersecurity, information science, and materials science. From the analysis of our interviews, we provide a characterization of the common personas of KG users, their expertise, tool usage, the obstacles they face when using KGs, and their unmet visualization needs. We then propose new directions for visualization research that leverages the semantic richness of KGs to both improve upon existing designs and address common challenges.

We identify three personas for KG users: *Builders* who construct and maintain KGs, *Analysts* who explore and extract insights from KGs, and *Consumers* who use insights from KGs for downstream tasks and decision-making. Across these three personas, we find that their expertise, tasks, and needs can vary drastically. Common challenges experienced by these personas include: difficulty querying KGs, poor data quality, evolving and mismanaged data provenance, schema inconsistencies, and a lack of organizational KG standardizations.

Several participants stressed sociotechnical challenges in understanding the desired outcomes of a knowledge graph. When this happens, Builders can overcomplicate the construction of the KG (e.g., with too many features), rendering it unsustainable. Analysts can struggle to acquire and deliver relevant insights to stakeholders, while Consumers

may ultimately perceive no utility in the KG's use for their downstream tasks. Consequently, organizations fail to adopt the KG in practice.

In addition to sociotechnical challenges, issues were raised related to current visualization methods. These challenges include: scalability, limited support for organic discovery, unaddressed domain-specific needs, and overly complex visualizations for end users. We find that node-link diagrams, although frequently used as a visual medium, are often ineffective for both generating and delivering insights from large KGs. For instance, one practitioner told us their software team's end users ultimately preferred simple table-based KG representations over custom-built interactive graph interfaces. Overall, we identify a need for visualization solutions that are targeted towards *each* KG persona, particularly those employing higher levels of abstraction to facilitate communication with diverse audiences.

When asked about what types of visual interfaces could be beneficial for exploring KGs, we found that participants widely praised Wikipedia's ability to support insight generation through on-the-fly data and entity hopping. In particular, “Wikipedia-style” interfaces are desirable for users when engaging in *open-ended KG exploration*, where there is no specific analysis target. Conversely, search engines like Google can help users when they want to pose direct and precise questions to the KG, representing *goal-oriented KG exploration*. As a whole, participants emphasized that there is currently no universally accepted solution for either type of KG-based exploratory analysis.

The remainder of our paper is structured as follows: Section 2 provides a background on knowledge graphs, as well as previous work on visualization practices for KGs. Section 3 outlines our protocol and analysis methodology for our interviews with KG practitioners. In Sections 4, 5, and 6, we respectively characterize the users, challenges, and visualization opportunities of KGs based on our findings. Finally, in Section 7 we discuss the limitations and future work for this study.

To summarize, the major contributions of this paper are:

- A thematic analysis of interviews conducted with 19 KG practitioners across eight different organizations who regularly create, explore, analyze, and deliver insights from KGs.
- A characterization of KG users, their common organizational roles, areas of expertise, tool usage, and visualization needs.
- Directions for future KG visualization research, along with design sketches iterated on with domain experts, aimed at alleviating KG-related challenges identified from our interview study.

## 2 BACKGROUND & RELATED WORK

We begin with a brief background on knowledge graphs and introduce terminology used throughout this paper. We then discuss related work

• Harry Li and Ashley Suh are with MIT Lincoln Laboratory and Tufts University. E-mail: {harry.li, ashley.suh}@ll.mit.edu.  
• Gabriel Appleby, Camelia Daniela Brumar, and Remco Chang are with Tufts University. E-mail: {gabriel.appleby, camelia\_daniela.brumar, remco.chang}@tufts.edu.

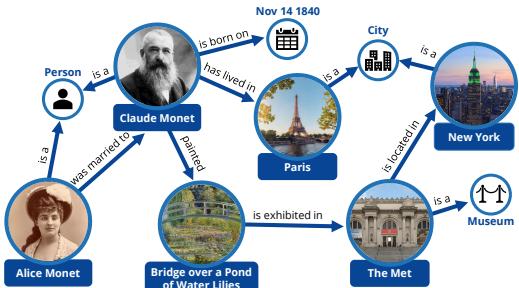


Fig. 1: An illustrative example of a knowledge graph (KG). In a KG, different types of entities (nodes) can have different types of relationships (edges) defined between them. We further discuss KGs in Section 2.1.

in KG system design, as well as KG visualization tools and practices.

## 2.1 Knowledge Graphs

Fundamentally, a knowledge graph is a data model that represents knowledge in the form of nodes (i.e. *entities*), edges (i.e. the *relations* between entities) and properties (i.e. *attributes*) that can be defined for both nodes and edges. A visual illustration of a simple KG can be seen in Figure 1. The concept of a knowledge graph dates back to the rise of the semantic web [10] in 2001; however, KGs gained notable popularity after the launch of Google's Knowledge Graph [21] in 2012, along with the demand for more sophisticated representations of data.

KGs store their information as triplets (e.g.,  $\{head, relation, tail\}$ , or  $\{subject, predicate, object\}$ ), providing a robust (and often hierarchical) structure to reason about data. The entire ‘blueprint’ that defines the KG’s structure (e.g., its nodes, edges, properties) is its *schema*. A schema specifies the different node types, edge types, properties, and constraints in the KG. By doing so, the schema ensures consistency, thereby facilitating data integration and providing a shared understanding of the data model for effective querying and analysis.

There are several well-known knowledge graph databases commonly used today (e.g., Stardog [69], Neo4j [53]) with respective graph query languages (e.g., SPARQL [22], CYpher [24]). By leveraging the underlying structure of the graph, queries on knowledge graphs can retrieve specific relationships, navigate through different nodes, and uncover meaningful patterns in the data. Graph analytic algorithms can also be applied to generate new features, as well as to learn and make predictions about data [11]. This enables knowledge graphs to support classification and regression analysis on both its nodes and edges.

The semantic nature of KGs makes them well-suited for enhancing language tasks, particularly when combined with large language models (LLMs) [3, 15, 61]. KGs are also used to manage diverse data sources, including data lakes, data warehouses, and knowledge bases (e.g., WikiData). For a complete background on KGs, including their common applications, we point to Hogan et al.’s survey paper [38].

## 2.2 Studies in Knowledge Graphs

Understanding the benefits, shortcomings, and future research directions for knowledge graphs was outlined in a 2019 Dagstuhl seminar report [11]. In the same year, another Dagstuhl report provided new directions in visual analytics research for multilayer networks [44], which have similar properties to KGs. A 2022 Dagstuhl report on the intersection of graph databases and network visualization [45] is likely the most similar in motivation to our work here.

However, to the best of our knowledge, our interview study is the first attempt to characterize the practitioners of KGs, their challenges, their tool usage, and their visualization needs. Similar qualitative work has been done within the visualization research community to understand the obstacles faced by ML experts [39, 60] (e.g., *data cascades* [64]), client-facing data scientists [55], exploratory data analysts [4, 43, 78], and ML stakeholders [72]. In addition to improving collaborations, these studies are similarly conducted to discover new opportunities for visualization to alleviate challenges faced by practitioners.

A recent position paper by Lissandrini et al. [50] called for better KG exploration tools. The authors highlight important tasks and use cases for KG systems, particularly for KG creators and maintainers. From our interviews, we identify similar challenges and needs for users (e.g., maintaining KG schemas, scalability, and the demand for interim query results); however, we also identify unmet visualization needs for KG builders, analysts, and consumers. While the authors provide an excellent foundation for database system designs, our work differs in that we interview KG practitioners directly to understand their needs across various technologies and visualization tools.

## 2.3 Visualization Solutions for Knowledge Graphs

There are an increasing number of systems that aim to visualize and explore the data in knowledge graphs. Latif et al. contributed *VisKonect* [46], a multi-coordinate visualization system for EventKG [30] that analyzes the connections of historical figures based on the events they participated in. VisKonect includes an NLP-based panel that lets users ask templated questions to the KG with the GPT-3 language model [15]. Ahmad et al. [2] contributed a visualization that maps data from a KG for patients with inflammatory bowel disease to compare their history, progressions, and administered treatments. Husain et al. contributed a multi-scale visual analytics approach for exploring biomedical knowledge graphs [42]. Their approach includes three types of views, a ‘global’ (high-level) view, a local ‘drilled down’ view, and a text-evidence document view. Partl et al. [59] contributed a scalable path finding approach with multiple views that queries and ranks candidate paths by topological features, as well as node and edge properties. While more visualization tools are being designed for and around KGs, they are not without their own set of limitations [50]. KGs are typically very large in size, may have repeat entities or attributes with similar names (i.e. entity ambiguation [27]), and may contain obsolete or out-of-date data [38]. For this paper, we target the identification of similar issues preventing wider visualization adoption for KGs, and potential opportunities to alleviate those challenges.

## 2.4 Knowledge Graph Solutions for Visual Analysis

In addition to visualization solutions targeted for knowledge graphs, the research community has looked into how we can leverage KGs to build better visualizations and interfaces. Dating back to 2008, Chan et al. presented Vispedia [17], an interactive visual exploratory tool that allows users to integrate and visualize data tables from DBpedia [7]. KG4VIS [49] recommends visualizations using a knowledge graph constructed from a large corpus of data-visualization pairs. KG4VIS can also generate rules from a KG embedding to “explain” why the model recommends certain visualizations given the user’s data. Cashman et al.’s CAVA system [16] utilizes KGs to help users interactively perform data augmentation on their existing datasets. Specifically, users are able to automatically join new attributes from a KG to improve performance on analysis tasks, e.g., improving a model’s predictive power.

While many of the tools discussed in this section are robust to their particular data domain and use case, it is currently unclear whether current solutions and visualization techniques can cover broader KG practitioner needs. In the following section, we describe an interview study to investigate these questions. In Section 5.4, we distill the potential benefits and tradeoffs of current KG visualization designs.

## 3 METHODOLOGY

To better understand the users of KGs, their use cases, challenges, and visualization needs, we conducted an interview study with KG practitioners from both research and enterprise settings. All interview and supplemental materials can be accessed at <https://github.com/TuftsVALT/KGsInPractice>.

### 3.1 Participant Recruitment

We recruited interview participants via emails to the authors’ professional contacts. In the recruitment email, we specified that interviews would be conducted with practitioners who had experience working with knowledge graphs in some capacity (e.g., creating, maintaining, querying). The demographics for our final 19 interview participants

PID	Education	Job Title	Company Domain	KG Persona(s)	Years of Experience	Familiarity with KGs	Familiarity Creating KGs	Familiarity Analyzing KGs	Familiarity Querying KGs	Familiarity Visualizing KGs
01	MS	Research Scientist	FFRDC	Builder, Analyst	4	4 (Moderate)	4 (Moderate)	4 (Moderate)	3 (Some)	3 (Some)
02	MS	Research Scientist	FFRDC	Builder, Analyst	20	3 (Some)	3 (Some)	3 (Some)	3 (Some)	4 (Moderate)
03	PhD	Research Scientist	FFRDC	Analyst	2	3 (Some)	3 (Some)	3 (Some)	3 (Some)	3 (Some)
04	MS	Research Scientist	FFRDC	Analyst	5	4 (Moderate)	3 (Some)	4 (Moderate)	3 (Some)	3 (Some)
05	MS	Research Scientist	FFRDC	Builder, Analyst	3	4 (Moderate)	4 (Moderate)	4 (Moderate)	3 (Some)	3 (Some)
06	MS	Research Scientist	FFRDC	Analyst	5	3 (Some)	3 (Some)	3 (Some)	3 (Some)	2 (Slight)
07	MS	Research Scientist	FFRDC	Analyst	5	4 (Moderate)	4 (Moderate)	4 (Moderate)	3 (Some)	2 (Slight)
08	MS	Software Developer	FFRDC	Builder, Analyst	3	2 (Slight)	2 (Slight)	2 (Slight)	3 (Some)	2 (Slight)
09	MS	Research Scientist	FFRDC	Builder, Analyst	2	4 (Moderate)	5 (Extreme)	3 (Some)	4 (Moderate)	2 (Slight)
10	PhD	Research Scientist	FFRDC	Builder, Analyst	5	3 (Some)	2 (Slight)	3 (Some)	4 (Moderate)	3 (Some)
11	MS	Data Analyst	Enterprise (Finance)	Builder, Analyst	2	4 (Moderate)	4 (Moderate)	4 (Moderate)	3 (Some)	1 (None)
12	PhD	Director	Enterprise (Health)	Builder, Analyst	8	4 (Moderate)	4 (Moderate)	4 (Moderate)	4 (Moderate)	4 (Moderate)
13	BS	Industry Analyst	Enterprise (Consulting)	Analyst, Consumer	2	5 (Extreme)	3 (Some)	3 (Some)	3 (Some)	4 (Moderate)
14	MS	PhD Student	Academia	Builder, Analyst	5	4 (Moderate)	4 (Moderate)	4 (Moderate)	4 (Moderate)	4 (Moderate)
15	PhD	Data Scientist	Enterprise (Health)	Analyst	4	4 (Moderate)	2 (Slight)	5 (Extreme)	5 (Extreme)	4 (Moderate)
16	PhD	Comp. Biologist	Enterprise (Health)	Builder, Analyst	5	4 (Moderate)	5 (Extreme)	4 (Moderate)	4 (Moderate)	4 (Moderate)
17	PhD	Principal Scientist	Enterprise (Tech)	Builder, Analyst	15	5 (Extreme)	5 (Extreme)	5 (Extreme)	5 (Extreme)	4 (Moderate)
18	MBA	Digital Lead	Enterprise (Health)	Consumer	1	3 (Some)	2 (Slight)	2 (Slight)	2 (Slight)	3 (Some)
19	MS	PhD Student	Academia	Builder	2	4 (Moderate)	4 (Moderate)	3 (Some)	3 (Some)	2 (Slight)

Table 1: Participant demographics for our interview study, described in Section 3. From left to right: the participant's ID; job title; the organization they work in (FFRDC stands for Federally Funded Research and Development Center); their primary persona(s) as KG users (further explained in Section 4.1); years of experience with KGs; overall familiarity working with KGs, creating or maintaining KGs, exploring or analyzing KGs, querying KGs, and visualizing KGs. Familiarity was self-reported on a Likert Scale from (1) not at all familiar to (5) extremely familiar.

is shown in Table 1. P1-10 come from varying divisions within the same FFRDC. P12, P15, and P16 come from the same company. The remaining participants are from 6 different organizations.

Participants self-reported their demographics, including their highest education level, job title, company domain, and subjective familiarity with knowledge graphs. Familiarity was selected on a Likert *Level of Familiarity* Scale from 1-5, where 1=Not at all familiar, 2=Slightly familiar, 3=Slightly familiar, 4=Moderately familiar, and 5=Extremely familiar. The following questions were asked regarding familiarity:

1. How would you rate your familiarity with KGs **in general**?
2. How would you rate your familiarity **creating or modifying** KGs?
3. How would you rate your familiarity **exploring, analyzing, and gaining insights** from KGs?
4. How would you rate your familiarity **querying** KGs?
5. How would you rate your familiarity **visualizing** KGs?

### 3.2 Protocol

Two authors conducted all interviews during a six month period. Most (18/19) interviews were conducted virtually through video conferencing software, with one conducted in person. We recorded and transcribed 15/19 interviews, and took detailed notes for the remaining where recording was not possible due to the interviewee's company policies.

Each interview lasted roughly one hour. Three of our interviews were conducted as focus groups (Group 1: P2, P3, P4; Group 2: P5, P6; Group 3: P9, P10) while the remaining were individual interviews. Participants were walked through the same set of PowerPoint slides prepared by the authors. This slide deck contained questions to help elicit the participants' KG experience, roles, projects, challenges, and visualization needs. The slides are provided as supplemental material.

After each interview, the authors met to discuss emergent themes, following a thematic analysis process [13]. We determined our sample size by reaching thematic saturation, where consistent themes emerged, and no new findings or potential codes were identified. In the end, we concluded our study after interviewing a total of 19 KG practitioners.

### 3.3 Interviews

At the start of each interview, participants were given a high-level definition of a knowledge graph and an illustrative example image (Figure 1) of one containing historical figures as nodes, relationships between figures as edges, and attributes belonging to both. In Section 7.2 we discuss our participants' own definitions of a knowledge graph.

Following this briefing, slides were shown asking participants to describe their experience with KGs: how they were created, how they

were queried, maintained, and their characteristics (e.g., size, schema, domain). We asked which types of questions participants were trying to answer with the KG, which data domain they worked in, which tools were commonly used, and which (if any) visualization solutions or tools had been helpful in the past.

To summarize, the overarching questions for our interviews were:

1. What is your experience with KGs (what do you use them for, how are they created, what are common challenges faced)?
2. What kinds of questions do you try to answer with your KGs?
3. What tools or techniques do you use, and what is the typical outcome of using them?
4. What visualization tools or methods do you use for KGs? What do you look for in a visualization tool, and what is missing?
5. At what granularity do you want to see or view the KG?

We also walked participants through several examples of knowledge graph visual analysis tools (specifically, [30, 46, 57]) to understand what could be potentially helpful or not helpful for different KG use cases. Participants' feedback, in addition to their responses to the above questions, helped to inform the opportunities for future KG visualization research we distill in Section 6. A write-up of their full responses to the presented tools is provided in our supplemental.

### 3.4 Analysis

The goal of our analysis was to qualify the users of KGs, their applications, their frequently experienced challenges, and visualization needs that are not satisfied by current technologies. We carried out a qualitative coding process to analyze our interview data. Our codebook was developed iteratively between all authors, and followed the protocols in [20, 52] for good codebook development. After each interview, the authors met to recap the discussion and note the most common uses cases, tools, issues, and needs experienced by the participant(s). Through this process, the strongest themes that emerged from our analysis related to: (1) why KGs are used over other solutions (*use cases*); (2) tools used for and with KGs (*tools*); (3) problems working with KGs (*challenges*); (4) lack of visualization support (*visualization needs*).

The interview data was coded by the original interviewers. We determined an utterance in the interviews to be a participant's turn in conversation, where a turn was a response to a single question from the slide deck. An individual code could only be assigned once per utterance, but many codes could be assigned in an utterance. Both authors coded the set of interviews, swapped to agree or disagree with each other's coding, then swapped again to make resolutions. Our entire coding process was thoroughly collaborative and conducted to

holistically categorize the interview findings, rather than rigorously compare the frequency of code usage. Therefore, disagreements were discussed and ultimately decided between both coders.

To remain consistent with our analysis procedure, we present the qualitative findings of our interviews in the following section in terms of our themes, and only use code counts to report the total number of participants experiencing a particular challenge, use case, and so on.

## 4 KNOWLEDGE GRAPH USERS & CURRENT PRACTICES

First we distill personas for the users of knowledge graphs, their common use cases, data sources, reasons for using a KG over other data models, as well as their frequently used tools and technologies.

### 4.1 KG Personas

From our interviews, we identify three major KG practitioner personas, highlighted in Figure 2. While each persona comes with distinct expertise, responsibilities, and needs, we often found that one person could step into multiple personas depending on their use case or organizational role. We assigned personas to each participant in Table 1 based on the variety of KG tasks they regularly perform.

**KG Builder:** The KG Builder is usually an expert in database systems, data management, or data modeling. Builders are responsible for creating a KG from its source data, deciding which database or representation method to use for storing the KG (discussed more in Section 4.4), and developing the KG's schema. Builders are typically an expert with one or more graph querying languages, particular those associated with their chosen KG representation. Many (12/19) of our interview participants fit into the Builder persona.

**KG Analyst:** The KG Analyst is typically an expert in data science or ML. Analysts are responsible for generating insights from the KG either as artifacts (e.g., reports) for end users, or as input for downstream analysis and AI/ML tasks. KG Analysts have familiarity with extracting information from the KG via its querying language, but may not necessarily be experts. Most of our participants who fit into the Builder persona also fit into the Analyst persona (11/12), as naturally Builders either are themselves the Analysts, or work closely with Analysts to construct the KG to meet the analysis use case. The majority (17/19) of our participants fit into the Analyst persona.

**KG Consumer:** The KG Consumer is generally an expert in the data domain, business, overarching use case, or the KG's sociocultural context (i.e. milieu [72]). While Consumers typically do not interact directly with a KG database or its querying language, they are still a stakeholder or end user of the KG, and know what "types" of insights would be valuable to extract. Consumers tend to rely on KG Analysts, query building GUIs, or automated reporting systems to generate those insights. Few (2/19) of our participants fit appropriately into the Consumer persona; we discuss this limitation in Section 7.3.

### 4.2 KG Use Cases

Our participants are using KGs for a wide variety of use cases (see Table 2 for examples, and the supplemental for a full list):

- 5/19 use KGs as enterprise data catalogs or data warehouses;
- 9/19 use KGs for path discovery;
- 15/19 use KGs with AI/ML to improve data quality, to analyze data, and to improve their predictive models. 6/19 specifically use KGs for node classification or regression, i.e. predicting a class label or numeric score for a set of nodes in the graph.

In terms of scale, 8/19 of our participants work with KGs containing thousands of nodes, 7/19 work with millions of nodes, and 4/19 work with billions of nodes. The number of edges ranged from millions to billions, with 3 to 10,000 properties on nodes and edges.

### 4.3 Benefits & Affordances of Using KGs

We also asked participants why they preferred to use KGs over other data structures, like traditional graphs or relational databases.

**Schema Flexibility:** KGs were praised as versatile due to their flexible schema structure, particularly when compared to other data structures:

Use Case	Description
Data cataloging	Creating a query-able knowledge base for data scientists and developers to quickly find data they need Standardizing and de-duplicating terminology and data usage Modeling the organization's business logic, e.g. "the organization has X branch, which has Y types of employees" Modeling facilities, their security systems, lights, fire suppression, etc. to connect to physical floor layouts Managing global web content, e.g., showing information about a movie based on the website visitor's country
Path Discovery	Finding new cyber threat pathways in a cyber KG [36] connecting computer systems and known exploits Finding new treatment pathways in a KG connecting diseases and possible treatments (e.g., for drug discovery) Finding new materials synthesis pathways in a KG connecting different chemical compositions Identifying user workflows in a KG connecting user actions in an enterprise network
AI/ML	Explaining why an anomaly was predicted from a model using data that is also connected to the KG Predicting stock prices for publicly traded companies in a KG connecting companies, industries, and supply chains Using NLP to process text-based data sources (social networks) to detect author profiles and authoring changes

Table 2: A select set of our participants' KG use cases. KGs are used for data cataloging, pathway discovery, as well as training, understanding, and improving AI/ML models. More details in Section 4.2.

If you think about a relational database, the column number must be the same for all the records, right? But for our KG...if you want to add a new type of data to a certain record, you can simply add in an edge. -P19

**Integration Across Multiple Data Sources & Domains:** Participants reported integrating both public (13/19) and non-public (10/19) data to generate and contextualize their KGs for end users. This allows Analysts and Consumers to discover previously unknown relationships:

The individual associations are insignificant on their own, it's only when we look at the multiple associations in the context of the graph do we... really find clusters. -P18

**Semantic Encodings:** A KG's semantic nature goes hand in hand with its robust ability to manage data. One participant told us a KG's ability to encode semantic-based relationships makes its usage worth the additional complexity of graph modeling and graph query languages:

We like the semantic explicitness of them [knowledge graphs], even if again, they're still hard to work with. -P9

In addition the above affordances, participants told us they distinctly use KGs to perform data augmentation, generate concept maps or ontologies, and to organize libraries of data assets via data catalogs.

### 4.4 Usage of Existing KG Tools

Our participants find success in several databases, tools, and methods for representing KGs:

- 7/19 use Neo4j with CYPHER;
- 6/19 use the Resource Description Framework with SPARQL;
- 4/19 use the NetworkX [33] Python package;
- 2/19 use SQLite3 with SQL;
- 1/19 use spreadsheets and adjacency matrices;

Our participants also use a combination of tools, libraries, and interfaces for visualizing their KGs:

- 7/19 use Neo4J Bloom [57];
- 5/19 use Gephi [9];
- 4/19 use NetworkX;
- 4/19 use Cytoscape [66];
- 2/19 use D3.js [12];
- 6/19 did not use any visualization tools.

One participant explained the need to create custom visualization tools for Consumers:



Fig. 2: Three personas we identified for the users of knowledge graphs from our interviews, described in Section 3. From the left, a user can be a KG Builder (e.g., database administrator), an Analyst (e.g., data scientist), or a Consumer (e.g., stakeholder). All three types of KG users have distinct roles, tasks, needs, and expertise – however, it is possible a user can belong to more than one persona. For example, a user that creates their own KG of companies (“KG Builder”) to predict which to invest into (“KG Analyst”). We further describe these personas in Section 4.1.

Typically I find interfaces like *Gephi* good for your initial exploration, but once you start wanting to put together a dashboard for an end user to use... you're going to find you want features in there that these tools don't have. I've always found that they're good to start with, and then I have to make custom tools after that. -P2

We discuss the benefits and tradeoffs of these visualization tools for KGs in Section 5.4, as well as in our supplemental material.

## 5 KNOWLEDGE GRAPH CHALLENGES

Our participants reported several broad challenges with using KGs in practice, many of which are fundamentally rooted in data sourcing and data quality issues. We outline the most common challenges experienced by our participants, and use these challenges to motivate directions for visualization research in Section 6.

### 5.1 Data Quality

The most common challenge faced by our participants (15/19) are problems surrounding data quality. These challenges include:

- Sparse or missing data [74], i.e. nodes or links that participants know should appear, but for one or more reasons do not.
- Incorrect or unverifiable data [19], i.e. nodes or links that participants know should not be in the dataset.
- Obsolete data [50], i.e. nodes or links in the dataset are no longer valid or relevant.
- Duplicate entities [25], i.e. multiple nodes or links in the dataset that actually should be combined into a single node or link.

The majority of these problems stem from incomplete or in-progress enterprise KGs. Data quality issues negatively impact AI/ML collaborations (e.g., data cascades [64]), making it difficult to account for a model's true robustness to missing data, noise, duplications, and so on.

While open-source KGs may be “complete” and useful for testing AI/ML models, they can also be unrealistic when compared to a real-world KG: “*The drawback of using WikiData is that it's too good... there are no holes you'd otherwise find in practice*” (P1).

**Manual Data Updates:** As with many kinds of data sources, KGs incur problems related to manually entering, validating, and invalidating data. While some KGs can be automatically generated, many still require manual human data entries to curate [77], which can be extremely burdensome on Builders who create enterprise data catalogs:

When you have 10 thousand attributes, it's not humanly possible to sit down and define all of them... Then if a system changes, the tags become obsolete. -P18

P19 described the challenge of validating KGs with millions of nodes, particularly when Consumers must manually perform this validation:

Domain scientists need to manually validate whether this extraction makes sense or not, whether it's completely nonsense or it is correct.

So we have to use human experts to validate randomly selected sample data. -P19

P16 shared they must periodically rebuild their KG from updated source data. During this rebuilding process, P16 and their team know that certain nodes and connections from the source data are invalid, however, it is difficult to manually annotate and ‘integrate’ these invalidations:

We often see you know, an association that doesn't make sense, or that we've already invalidated internally... it would be nice to have an easy way to basically flag that for all future versions of the KG that are built. -P16

**KG Entity and Path Challenges:** Challenges strongly associated with KGs include entity disambiguation and chokepoints in path discovery. Several (9/19) participants said they face challenges with entity ambiguity, in which a node or edge has multiple meanings in the KG:

One problem is entity disambiguation. We take the data as is, we're not sure that two nodes are actually different. -P8

From the participants using KGs for path discovery, 4/9 have problems with chokepoint nodes through which many paths converge then diverge. When running path-finding algorithms to discover new connections, densely connected nodes in between the target nodes can grossly inflate the number of discovered paths, leading to irrelevant outputs:

A big problem is some of these intermediate layers are a lot smaller than the layers they're connected to, so they're chokepoints. If you do that two step linkage between the two nodes, you end up with probably a lot of things that are irrelevant and more nodes than you actually want... it would be super cool if we could use machine learning to infer the appropriate linkages. -P9

Some participants told us their team is experimenting with hyperedge representations, i.e. edges that connect more than just two nodes, in an attempt to avoid path chokepoints and problematic path convergence.

### 5.2 Querying

Querying is the most challenging problem faced by both KG Analysts and Consumers (11/19), particularly because each KG representation method typically has its own unique querying language [38]. Learning a graph query language is also difficult for end users:

It's a hard sell to get them [end users] to spend the time to invest in learning those query languages without knowing that they're going to get something out of it. -P16

Even though KG Builders and Analysts remarked that end users should not have to (nor want to) learn a KG query language, one KG Consumer told us that they still need the ability to “ask the KG” questions for their own downstream tasks:

As an end user who doesn't write SPARQL, I would like to ask [the KG] who my best customers are. I'd like to visually explore

the graph to determine that, since I know there are many possible answers to that question. -P13

**Lack of Interim Results:** Often as a user is developing a query, they need an interim subset of results to determine if the query is pulling relevant information. However, many querying systems instead wait until all the information is ready before returning the full set of results:

It was very frustrating because you'd construct a query, it would take like 15 minutes to load, and then you'd get no results after it finished. So having some type of interim result... like 'kept alive' with examples of the stuff it's bringing back, that kind of interaction would be helpful. At least just to make sure that you're on the right track. -P4

Long wait times caused by computation is frustrating for users [67] and interrupts their analysis workflows [51]. Work in progressive visualization [5] could alleviate similar issues for KGs.

### 5.3 Socio-Technical Problems

Two of our participants mentioned that many of their challenges have both social and technical aspects that stem from difficulties in interpersonal communication and collaboration.

**Incomplete Understanding of End Users' Needs:** P17 has observed that many people are drawn towards creating KGs before properly understanding the overarching use case and needs of the end users:

People always want to go build a knowledge graph of everything... They've defined success from a technical point of view. But what are you going to do with it? Why are you doing it? Who's going to go use it? What's the value it's going to produce? -P17

P18 (a KG Consumer) told us a similar story about a developer (Builder) who unnecessarily over-complicated the construction of a KG:

I have a firm belief that a developer added all those features [to the KG] because they thought, 'the more features, the better,' instead of considering what we actually needed for analysis. -P18

In the end, an enormous amount of time and effort is spent to create a KG that might not necessarily have utility for its users. P18 was adamant that a simpler version of their company's KG would have better met the end users' needs – leading to its wider adoption.

**Non-standardized Nomenclature:** Another challenge is a lack of standardized nomenclature, in which different groups of people may use one word to describe multiple meanings, or alternatively use various terms to describe the same concept:

There's this concept of profit. What does profit mean? Well, it's actually this complicated math that's not in the source. So we go talk to different people, and they're gonna have different answers. One word can mean multiple things to multiple people. -P17

**Organizational Politics and Unsustainability:** Across enterprise settings, KGs can fail to be adopted for political reasons, or fail to be maintained due to its long-term unsustainability:

Was our knowledge graph successful? No. The reason for failure was more political. You need funding and resources from leadership, but the interest died out. -P18

Organizational issues are often cited as a major reason for AI and ML models failing to be adopted in industry settings [60].

**Too Many Hammers, Not Enough Talking:** While research continues to focus on optimizing knowledge graph databases and query languages, P17 told us that many technological problems are already solved. Instead, P17 believes computer scientists need to be open to addressing the social problems related to KGs:

We can build faster graph database systems. Is it intellectually challenging to go do that? Definitely. But is that going to push the barrier for the world to take unknowns and turn them into knowns? Honestly, no. We don't need more hammers. We need to go figure out how to use the hammers... This is where computer scientists

can get uncomfortable, doing qualitative methods. -P17

### 5.4 Current KG Visualization Designs

A major point of discussion in our interviews was how current KG visualization designs either meet or fail to meet users' needs. By far, node-link diagrams [37] (NLDs) were the most commonly used KG visualization across all three KG user personas (18/19). However, we find that NLDs have shortcomings for the (albeit) many challenges they tackle. We discuss those challenges below.

**Lack of Scalability for Visual Sanity Checking:** Node-link diagrams (NLDs) are commonly used by KG Builders (9/12) as "sanity checks" to *"make sure nothing weird is going on... that the graph is connected as expected"* (P14). For creating NLDs, most of our participants used Gephi [9]. When asked about the limitations of node-link diagrams for sanity-checking, participants told us that scalability was the biggest issue: *"I don't want to have to see the full graph... I'd just like to see something like, a quick sanity check"* (P6).

Scalability is an issue at two levels. First, when the knowledge graph is dense, it can be difficult to make sense of the resulting NLD. To alleviate this problem, an interactive NLD (e.g., force-directed layout [28, 71]) is used; however, when the KG is very large, it can be computationally difficult to render the entire graph – an optimization problem that is still an ongoing area of graph visualization research [29, 75].

**Lack of Efficacy for KG Consumers:** Many (12/19) participants explicitly mentioned that NLDs are impossible to interpret by end users at the scale of thousands, millions, or billions of nodes. Several (9/19) participants criticized NLDs for quickly turning into "hairballs," making it difficult for end users to digest meaningful information:

Because graphs are very visual, I hypothesize that there's something socially in our brain that automatically says, I want to see it. I hear this all the time, 'I just want to see it.' Okay, so you see this small graph, so what? 'Show me something bigger.' Then it turns into a hairball. This is the story of graph visualization. -P17

P12 told us about an interactive graph interface that his development team built, only for end users to reject it in favor of a table diagram:

We put a lot of software developers' effort into this GUI that showed our analysis as a graph to the user. But the results that came back always looked like a ball of yarn, unstructured... Users couldn't make sense of it. In the end, they preferred a table. We played around with different ways to clean up and summarize the graph, but we never found a good way to visualize the "graph-ness" of the data in a way that the users could navigate. -P12

From our interviews, we believe KG Consumers tend to prefer tables over other representations due to: (1) their simplicity and familiarity across multiple domains, and (2) the Consumer's task at hand is often straightforward (e.g., data retrieval). This is in contrast to Builders and Analysts, who tend to prefer more comprehensive visualizations (e.g., NLDs) when exploring the data or completing more complex tasks.

Regardless of the shortcomings of NLDs, two participants specifically mentioned that they make for good eye candy: *"I will argue that they make very pretty pictures... they're great slide decoration"* (P13).

## 6 VISUALIZATION OPPORTUNITIES FOR KNOWLEDGE GRAPHS

Finally, we present directions for visualization research that can begin to alleviate many of the challenges identified in Section 5. Where appropriate, we include participant quotes or references to related literature that motivates each recommendation.

### 6.1 Graph-Abstracted Visualizations

A recent Dagstuhl report [45] discusses future directions of research at the intersection of graph databases and network visualization. With respect to the current limitations identified in this work, it is critical that a relevant design space for knowledge graph visualizations (and their end users) is explored and contributed. Building on previous graph visualization work can provide a starting point [48, 58, 59].

Beyond standard network visualizations, some (5/19) participants advocated for KG applications and visualizations that are *not* graphs:

Just because you have a knowledge graph doesn't mean the visualization has to be a graph... Under the hood, the end user doesn't even have to know it's there. But they're getting the benefits of having it there in terms of improved results... I've come to see that the most effective way for non-technical experts to engage with KGs is through specific applications that are powered by a KG, rather than directly tangling with a ball of yarn visualization. -P13

P16 found that end users are not only confused by NLDs, but trust the analysis results less when they are presented as a graph visualization:

In our experience, the more we can shelter the end user from the underlying graph structure, the better their willingness to interact with the data and accept the results that come out of it. As soon as the level of complexity of the graph reaches a certain level on the screen, users really tend to shut down and not trust any of it. -P16

P15 agreed, telling us that, “*there's a big difference between the format that machines want to read data, versus the format that humans want to read data.*” While KGs store complex data (knowledge) that is interpretable to both humans and machines, this feedback suggests that users do not necessarily prefer to *see* this knowledge in graph form.

**What if I have to (or want to) use a graph visualization?** As discussed in Section 5.4, many of our participants still have to use NLDs (or general graph visualizations) for a variety of use cases and applications, particularly Builders and Analysts. Our participants identified several capabilities, often lacking in current tools, that should be supported when interacting with or consuming KG visualizations:

- The ability to immediately begin analysis at the user's desired point or region of interest (i.e. drilled drill down into the KG).
- The ability to filter, bundle, condense, collapse, or expand areas (regions) of the KG during open-ended exploration.
- The ability to switch views, while maintaining context, depending on the KG data type (e.g., from graph view to table view).

## 6.2 Balancing Digestibility and Discoverability

One of the advantages to node-link diagrams is the ability for users to traverse the KG from node to node to discover new information. However, users find NLDs ineffective in practice because the scales of KGs they work with (Section 4.2). Consequently, there is a need for visual interfaces that balance both digestibility and discoverability, to allow users to properly process information and explore the KG:

I think the key with [KG interfaces] is not to lose the pivotability and discoverability that is fairly unique to graphs... If you can maintain this sort of continuous reference to the graph that allows for that organic discovery process as opposed to like static results... that could be really good. -P13

**Wikipedia as an EDA Tool:** Four of our participants expressed that Wikipedia functions as an effective graph exploration tool, since Wikipedia simultaneously presents detailed information about a specific article (i.e. node) and also embeds hyperlinks to other related articles (i.e. edges connecting to another node). In this interface setup, users can both explore the graph structure of Wikipedia while also gaining valuable insights of their choosing by interacting with articles. We discuss the design of potential KG interfaces to support seamless knowledge discovery in Section 6.3.

**Contextual Knowledge Cards:** Similar to a *baseball card*, a knowledge card can give a high-level summary of the most essential data for that particular node or entity in the KG. Five of our participants said they use similar visualizations to deliver KG context to Consumers:

A knowledge card can be a really powerful as a visualization in use cases that I want to understand the context of something. This is where knowledge graphs are really powerful, where they have the advantage over like traditional relational databases... being able to understand context and relationships. -P13

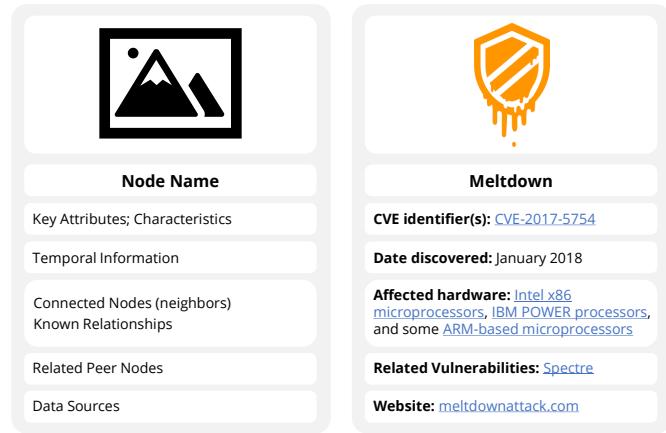


Fig. 3: Left: an example knowledge card template with node and edge information that may be relevant to a KG end user. Right: an example knowledge card of a cybersecurity vulnerability that we iterated on with one of our participants to understand what might be useful to a cyber analyst. We describe knowledge cards in Section 6.2.

A template for the basic information expected in a knowledge card is provided in Figure 3. To create this template, we had follow-up calls with P5. During these calls, we iterated on knowledge cards that would be useful for their particular use case. In general, the information needed was highly specific to the domain and use case. General information included: an image of the entity being queried from the KG, key attributes or identifying characteristics of the entity (including temporal data e.g., updates to the entity or when the entity was created), known paths, relevant entities, and data sources or credentials.

Connecting this research opportunity to Section 6.1's, another interesting direction could include the investigation of using knowledge cards as the “nodes” in a node-link diagram. Further, when clustering groups of nodes in the KG for a ‘global view,’ a knowledge card could represent a high level abstraction for collections of nodes (e.g., a card representing a cluster of universities or sports teams).

## 6.3 KG-Based Interfaces that Support Organic Discovery

There are a variety of opportunities for visual interfaces to align with the capabilities of KGs, thereby facilitating data discovery and exploration. In particular, interfaces that allow end users to “ask the KG questions” that they would not have thought to ask in the first place:

When I'm looking at an entity in the knowledge graph, I want all the interesting questions and answers about that entity available... So, for example, if there's a drug undergoing a particular clinical trial, I want to be able to quickly have all the interesting properties about that drug on a page... [to support] very fast discovery of insights about individual instances of data that I probably wouldn't have found without the tool showing me. -P15

We identified two types of exploration use cases for KG visual interfaces: those for unguided, unstructured, and organic exploration (i.e. “**open-ended KG exploration**”), and those for directed and targeted exploration (i.e. “**goal-oriented KG exploration**”):

One is a very exploratory type of search, just navigating, clicking around and trying to learn things with no clear goal. Have you ever clicked on something on Wikipedia, then you end up going down this rabbit hole? Did you have an objective? No, but you did a bunch of work, and you probably learned something. The other [type] is something very specific... Think of ‘Googling’ something. This is a search problem, but with intention. -P15

We believe there are two directions for KG-based visualization interface design to support both types of exploration.

**Visual Interfaces Built on Top of KGs:** One direction is for interfaces built “on top of the KG,” that is, users interact with an interface to explore and ask the KG questions directly. In these interface designs, interactive visualization can help users make sense of the KG and query it, possibly without any knowledge of the KG’s inherent query language (e.g., similar to using Wikipedia or Google’s search engine). There has been prior work done in the KG community to support KG exploration through query-building GUIs [23, 31]. Similar efforts have been made to lower the barrier for querying relational databases, such as NoSQL [34], as well as natural language queries (e.g., NL-to-SQL) [26], and natural language interfaces (NLIs) [8, 56]. While NLIs are becoming more popular in the visualization research community (e.g., [40, 54]), we see an untapped opportunity to integrate these techniques with KG exploratory visualization tools.

**Using KGs to Augment Exploratory Visualization Tools:** Another direction is utilizing KGs altogether to enhance current visualization interfaces. With the growing popularity of public KGs [76], visualization system designers can consider adding the ability for users to “connect” to a KG relevant to their own domain or task. For example, given a generalizable exploratory visualization tool like Tableau or Voyager [79], an additional widget to “ask a KG” may be implemented to enrich exploration with data that the user is not currently connected to, increasing in-situ discovery and data integration (similar to [16]). Given a KG’s semantic richness and well-defined (often hierarchical) structure, KGs could seamlessly provide additional context, annotations, supplementary visualizations, etc. to an analysis session.

#### 6.4 KGs for Explainability

The semantic nature of KGs (Section 4.3) can be leveraged for explainable AI (XAI) to help model creators debug their model during training, and also to help end users trust a model’s predictions [73]. This is particularly valuable for deep neural networks, which are often regarded as black boxes with limited interpretability [63]. Lecue [47] describes opportunities for using KGs to help encode the semantics of inputs, outputs, and their properties in a neural network.

**XAI for Model Debugging:** Three of our participants wished there were tools to help them debug and improve ML and KG model training:

Since machine learning is mostly my focus, visualization would have been helpful... When the data is in this intermediate knowledge graph form, it's really hard to debug and visualize performance... it's hard to tell like, is my model doing worse because the data is different, or because the data is exactly the same except for some small detail? -P9

**XAI for KG Analysts and Consumers:** KGs can serve as an effective tool for providing predictive explainability for Analysts and Consumers. For instance, in image analysis, saliency maps can highlight the specific areas in the input image that a neural network focused on, providing insights into why the model arrived at a particular prediction [68]. Similarly, ML models trained on KGs could offer meaningful semantic explanations for their predictions by highlighting the relevant nodes and edges that played a role in shaping their predictions. By leveraging KGs, we can empower users to understand the reasoning behind a model’s predictions and build greater trust in its outputs.

As an illustrative example, P9 is developing an anomaly detection algorithm that evaluates the semantic proximity of objects using a KG connecting household objects and their locations. Their model should ideally be able to detect that a hammer is an anomaly in the context of a kitchen. In their case, the KG helps “explain” the detected anomaly:

Hammers are normally in a shed. And our hope was basically that by using a source of context, we could do anomaly prediction for things like, should a hammer be in the kitchen? The answer is no here. So the way we chose to get that context was through a knowledge graph. -P9

P9 described two types of explanations that the KG can provide for why the model classifies a hammer in a kitchen as an anomaly:

1. Nodes in the KG like pots and plates appear in this context, but are not closely related to hammers in the KG.
2. Nodes in the KG like wrenches and shovels are closely related to hammers, but do not appear in this context.

We posit that these kinds of contextual explanations can be important for KG-based XAI, and potentially powerful for visual analytics.

#### 6.5 KG Timelines: Tracking Evolutions Over Time

We identify two distinct temporal-based directions for knowledge graph visualization research.

Another peculiarity of my data is the entities are all time stamped roughly. So it would be interesting to see the evolution of products or entities over time. - P9

**Visualizing Multi-Attributed Time-Series KG Data:** First, there is a need to consume and understand time series, temporal, or “timestamped” data from knowledge graphs, e.g., the data in EventKG [30]. We observe the need for new visual designs and interfaces that allow users to precisely (and organically) navigate and consume temporal data, events, and relations in a large-scale KG. Work similar to Brehmer et al.’s [14] could be done to contribute design spaces for multi-attributed temporal (time-series) data, since the nodes and edges of knowledge graphs typically contain multiple attributes.

**Tracking KG Data Evolution and Authoring Changes:** KG users need visualization solutions that track how the knowledge graph has changed over time, similar to previous visualization work in tracking software changes [70, 81]. Specifically, users need help tracking and validating in what capacity the KG has been assigned additional information, e.g., through new edges (relations) added. Users also need to know whether new information drastically changes their mental model of the KG, or makes their analyses out of date. For example, if an Analyst is curating a report using news articles contained within a KG, it is important for them know whether their current state of information is obsolete. As discussed in Section 5.1, some users also do not have an effective method of annotating data invalidations to omit certain nodes or links when they rebuild KGs from source data.

#### 6.6 Mapping Dynamic Data onto Static Views

One of the challenges associated with using NLDs for large KGs is that the algorithms used to generate “nice” and computationally fast layouts (e.g., FDLs [28]) often calculate the node and link positions dynamically (or stochastically). Consequently, the graph layout can change drastically each time the KG visualization loads, which can confuse users who must reorient themselves after each change.

One possible solution was posed by P12, whose team of biologists found success in always visualizing KG data on top of the Roche Biochemical Pathways<sup>1</sup> diagram, a standardized graph visualization detailing various biochemical processes:

There are networks in biology that people [biologists] are already familiar with that are best visualized as a graph. The graph becomes static, then you can load data and project it onto that map. This lets me highlight parts in the network that were active in an assay, or played a role in diabetes. -P12

For example, the same method is used when visualizing navigational directions. The network of roadways remains static, while routes and icons can be overlaid on top. While this KG visualization opportunity may be use case specific, it could support Consumers who require the context of their own domain to extract insights from the KG.

#### 6.7 KG Schema Creator and Enforcer

A good, consistent knowledge graph hinges on a good, consistent knowledge graph schema [1]. Interview participants (typically KG Builders) told us that creating a reliable schema can take several months at a time, which halts the actual development process of the KG database:

<sup>1</sup><http://biochemical-pathways.com/#/map/1>

I think another thing that's very much missing from the [KG] landscape is how to even build the graph to begin with, how to put together your schema... It can make querying impossible if you don't build it correctly. Like if I build it in *this* way, I can get *this* information out of it. But if I build it in this other way, I won't ever be able to do *this* query. -P2

**KG Schemas as Visual Maps:** There are many graph and tree visualizations that could act as a starting point for schema visuals [37]. KG Builders need concise but detailed views of the schemas they are creating, maintaining, and iterating on for their knowledge graph. Moreover, visual designs should consider that schemas may change over time: attributes can be added to nodes or edges, new relations can be created, while others may be removed entirely. A good way to highlight how a schema has changed over time should be integrated.

**Interactive Schema Builder & Enforcer:** Interview participants that spend months building a schema frequently use tools like *Visio* [35] to create the schema framework or template (i.e. a schema visual map). The common usage of *Visio* should make clear what is needed for an interactive schema builder: flexibility, customizability, and a multitude of design tools. However, what is lacking in a generic tool like *Visio* is the ability to: (1) preview the schema, (2) integrate directly into a KG workflow, (3) enforce types and constraints in the schema (e.g., [6]). In an ideal tool, the KG Builder could also preview how different queries could be accomplished given the current schema design.

## 7 DISCUSSION

### 7.1 Domain-Specific KG Visualization Designs

We began these interviews with the assumption that there is a common challenge faced by most KG practitioners which could be addressed by a generalizable KG visualization system. Instead, we quickly found that – even though our participants did share common challenges – their domain-specific needs could not be met by a single visualization solution. Many (14/19) participants explicitly stated they believe that **visualization challenges for KGs are domain specific for end users**.

It's probably going to be really hard to find visual metaphors that work for everybody... and for different tasks, right? I think the domain will prioritize how you visually present the data, and then that's what a useful visual metaphor for the data will be. -P10

As P2 told us: "*It's going to be hard to make a generalization unless you know the exact use cases.*" Creating effective KG visualizations will require additional formal or informal user studies [65] to understand the end users' data domain, applications, questions, and needs.

### 7.2 What Defines a Knowledge Graph?

There is often confusion for what makes a data model a KG beyond storing nodes and edges, a commonality across any graph data representation. We asked our interview participants to provide their own definition and criteria for a knowledge graph, the full list is included as supplemental material. Based on our own participants' collective definitions of a KG, we offer the following description:

*A data model representing entities as nodes, the multi-relationships between those nodes as edges, and properties defining them, such that humans and machines can easily understand the nuances of that data due to its semantics.*

The most important criteria of a KG identified by our participants were: (1) its ability to store different types of nodes (or entities), different types of semantic relationships (or edges) between those nodes, and the attributes (or properties) on them; (2) its ability to help both humans and machines understand what the data "*actually means*" – e.g., in the greater context of the data domain or use case.

### 7.3 Limitations & Future Work

Additional work beyond our interview study is needed to further understand the role of visualization for knowledge graphs used in practice. While KGs have been an active area of research in other communities (e.g., database systems [50] and NLP [18]), they have only recently

become a target of study in visualization research (e.g., [16, 49]). There are a multitude of opportunities for visualization research to leverage the semantic-richness of KGs, as well as application-driven research to augment current KG tools. We presented our participants with three tools [30, 46, 57] discussed in Section 2, with mixed feedback on their perceived helpfulness (see supplemental). Future work should investigate to what extent existing graph visualization research can be applied to KGs, as well as how accessible these systems are to practitioners.

For our study, we interviewed 19 practitioners across eight organizations, with roughly half the participants coming from the same FFRDC. We envision a wide array of future studies conducted to better understand the challenges and needs of KG practitioners – similar to the ongoing user-centered research being done for AI/ML collaborations [64]. As many of our participants discussed with us, robust technical solutions have already been posed related to building and completing KGs [62], however, social challenges related the usability of KGs remains a large-scale issue in many collaborative settings.

Another limitation of our study is the lack of KG Consumers interviewed: 17/19 participants were either KG Builders or Analysts. This is in part because Consumers tend to work with applications that are served from a KG "under the hood," and might have been self-selected out of our interview solicitation process. Consequently, many of the identified challenges and posed solutions in this study relate to practitioners' feedback who work more directly with a knowledge graph. Future work should therefore be done to target the ultimate end users of KGs – even those that might be unaware that they in fact use KGs. Suresh et al. conducted similar research for ML stakeholders [72].

In general, we were not able to identify any single solution for "the best" visual encodings for KGs – instead, we identified shortcomings of current designs. However, we believe these findings underpin the need for further research in visualization for KGs, and KGs for visualization. While we posed a variety of visualization research directions and possible designs in Section 6, future work will need to address the validity of those suggestions. This also opens up visualization research in curating KG task taxonomies (similarly, to compare and contrast to [48]), design guidelines, and potential KG design spaces.

## 8 CONCLUSION

We presented an interview study with 19 practitioners in industry and academic settings across eight organizations who regularly use knowledge graphs. From our interviews, we distilled common knowledge graph practices, uses cases, and tools frequently used by practitioners. We identified three personas for the users of KGs: (1) *Builders* who create and maintain KGs, (2) *Analysts* who explore and analyze the data in KGs, and (3) *Consumers* who use the insights from KGs for downstream tasks. From those personas, we discussed how each KG user has distinct expertise, tasks, and visualization needs. Overall, we found a gap in current tools and visualization methods (e.g., the usage of node-link diagrams for representing and visually communicating large KGs) for the challenges experienced by interview participants. Based on these collective findings, we outlined several directions for visualization research to enable better KG maintenance, "open-ended" and "goal-oriented" data discovery, analyses, and collaborations.

## ACKNOWLEDGMENTS

We sincerely thank each of the practitioners who took the time to participate in our interview study, and the reviewers for their insightful feedback on improving the quality of our paper. This work was supported by National Science Foundation grants IIS1452977, OAC-1940175, OAC-1939945, OAC-2118201, NRT-2021874.

**DISTRIBUTION STATEMENT A.** Approved for public release. Distribution is unlimited. This material is based upon work supported by the Department of the Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Air Force.

## REFERENCES

[1] B. Abu-Salih. Domain-specific knowledge graphs: A survey. *J. Netw. Comput. Appl.*, 185:103076, 2021. doi: 10.1016/j.jnca.2021.103076 1, 8

[2] S. Ahmad, D. Sessler, and J. Kohlhammer. Towards a comprehensive cohort visualization of patients with inflammatory bowel disease. In *Proc. VAHC*, pp. 25–29. IEEE Computer Society, Los Alamitos, 2021. doi: 10.1109/VAHC53616.2021.00009 2

[3] B. AlKhamissi, M. Li, A. Celikyilmaz, M. Diab, and M. Ghazvininejad. A review on language models as knowledge bases. *arXiv preprint arXiv:2204.06031*, 2022. doi: 10.48550/arXiv.2204.06031 2

[4] S. Alsbaugh, N. Zokaei, A. Liu, C. Jin, and M. A. Hearst. Futzling and moseying: Interviews with professional data analysts on exploration practices. *IEEE Trans. Vis. Comput. Graph.*, 25(1):22–31, 2019. doi: 10.1109/TVCG.2018.2865040 2

[5] M. Angelini, G. Santucci, H. Schumann, and H.-J. Schulz. A review and characterization of progressive visual analytics. *Informatics*, 5(3), 2018. doi: 10.3390/informatics5030031 6

[6] R. Angles, A. Bonifati, S. Dumbrava, G. Fletcher, A. Green, J. Hidders, B. Li, L. Libkin, V. Marsault, W. Martens, F. Murlak, S. Plantikow, O. Savkovic, M. Schmidt, J. Sequeda, S. Staworko, D. Tomaszuk, H. Voigt, D. Vrgoc, M. Wu, and D. Zivkovic. Pg-schema: Schemas for property graphs. *Proc. ACM Manag. Data*, 1(2), 2023. doi: 10.1145/3589778 9

[7] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives. Dbpedia: A nucleus for a web of open data. In K. Aberer, K.-S. Choi, N. Noy, D. Allemang, K.-I. Lee, L. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber, and P. Cudré-Mauroux, eds., *Semant. Web*, pp. 722–735. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-76298-0\_52 2

[8] J. Aurisano, A. Kumar, A. Gonzales, J. Leigh, B. DiEugenio, and A. Johnson. Articulate 2 : Toward a conversational interface for visual data exploration. In *Proc. VIS*, 2016. 8

[9] M. Bastian, S. Heymann, and M. Jacomy. Gephi: An open source software for exploring and manipulating networks. *Proc. AAAI*, 3(1):361–362, 2009. doi: 10.1609/icwsm.v3i1.13937 4, 6

[10] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Sci. Am.*, 284(5):34–43, 2001. 2

[11] P. A. Bonatti, S. Decker, A. Polleres, and V. Presutti. Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (Dagstuhl Seminar 18371). *Dagstuhl Reports*, 8(9):29–111, 2019. doi: 10.4230/DagRep.8.9.29 2

[12] M. Bostock, V. Ogievetsky, and J. Heer. D3 data-driven documents. *IEEE Trans. Vis. Comput. Graph.*, 17(12):2301–2309, 2011. doi: 10.1109/TVCG.2011.185 4

[13] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qual. Res. Psychol.*, 3(2):77–101, 2006. doi: 10.1191/1478088706qp063oa 3

[14] M. Brehmer, B. Lee, B. Bach, N. H. Riche, and T. Munzner. Timelines revisited: A design space and considerations for expressive storytelling. *IEEE Trans. Vis. Comput. Graph.*, 23(9):2151–2164, 2017. doi: 10.1109/TVCG.2016.2614803 8

[15] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In *Adv. Neural Inf.*, vol. 33, pp. 1877–1901. Curran Associates, Inc., 2020. 2

[16] D. Cashman, S. Xu, S. Das, F. Heimerl, C. Liu, S. R. Humayoun, M. Gleicher, A. Endert, and R. Chang. Cava: A visual analytics system for exploratory columnar data augmentation using knowledge graphs. *IEEE Trans. Vis. Comput. Graph.*, 27(2):1731–1741, 2021. doi: 10.1109/TVCG.2020.3030443 2, 8, 9

[17] B. Chan, L. Wu, P. Hanrahan, J. Talbot, and M. Cammarano. Vispedia: Interactive visual exploration of wikipedia data via search-based integration. *IEEE Trans. Vis. Comput. Graph.*, 14(06):1213–1220, 2008. doi: 10.1109/TVCG.2008.178 2

[18] X. Chen, S. Jia, and Y. Xiang. A review: Knowledge reasoning over knowledge graph. *Expert Syst. Appl.*, 141:112948, 2020. doi: 10.1016/j.eswa.2019.112948 9

[19] P. Cimiano and H. Paulheim. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semant. Web*, 8(3):489–508, 2017. doi: 10.3233/SW-160218 5

[20] J. T. DeCuir-Gunby, P. L. Marshall, and A. W. McCulloch. Developing and using a codebook for the analysis of interview data: An example from a professional development research project. *Field Methods*, 23(2):136–155, 2011. doi: 10.1177/1525822X10388468 3

[21] L. Ehrlinger and W. Wöb. Towards a definition of knowledge graphs. *Proc. ESWC Posters and Demos Track*, 48(1-4):2, 2016. 2

[22] B. Ell, A. Harth, and E. Simperl. Sparql query verbalization for explaining semantic search engine queries. In V. Presutti, C. d'Amato, F. Gandon, M. d'Aquin, S. Staab, and A. Tordai, eds., *Proc. ESWC*, pp. 426–441. Springer International Publishing, Cham, 2014. doi: 10.1007/978-3-319-07443-6\_29 2

[23] S. Ferré. Sparklis: An expressive query builder for sparql endpoints with guidance in natural language. *Semant. Web*, 8(3):405–418, 2017. doi: 10.3233/SW-150208 8

[24] N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer, and A. Taylor. Cypher: An evolving query language for property graphs. In *Proc. SIGMOD*, p. 1433–1445. ACM, New York, 2018. doi: 10.1145/3183713.3190657 2

[25] A. Gal. Uncertain entity resolution: Re-evaluating entity resolution in the big data era: Tutorial. *Proc. VLDB Endow.*, 7(13):1711–1712, 2014. doi: 10.14778/2733004.2733068 5

[26] Y. Gan, X. Chen, J. Xie, M. Purver, J. R. Woodward, J. Drake, and Q. Zhang. Natural SQL: Making SQL easier to infer from natural language specifications. In *Proc. EMNLP*, pp. 2030–2042. ACL, Punta Cana, Dominican Republic, 2021. doi: 10.18653/v1/2021.findings-emnlp.174 8

[27] T. Gao, M. Dontcheva, E. Adar, Z. Liu, and K. G. Karahios. Datatone: Managing ambiguity in natural language interfaces for data visualization. In *Proc. UIST*, p. 489–500. ACM, New York, 2015. doi: 10.1145/2807442.2807478 2

[28] H. Gibson, J. Faith, and P. Vickers. A survey of two-dimensional graph layout techniques for information visualisation. *Proc. InfoVis*, 12(3-4):324–357, 2013. doi: 10.1177/1473871612455749 6, 8

[29] J. Gómez-Romero, M. Molina-Solana, A. Oehmichen, and Y. Guo. Visualizing large knowledge graphs: A performance analysis. *Future Gener. Comput. Syst.*, 89:224–238, 2018. doi: 10.1016/j.future.2018.06.015 6

[30] S. Gottschalk and E. Demidova. Eventkg: A multilingual event-centric temporal knowledge graph. In *Proc. ESWC*, pp. 272–287. Springer, 2018. doi: 10.1007/978-3-319-93417-4\_18 2, 3, 8, 9

[31] P. Graffin, M. Mironov, M. Fellmann, B. Lantow, K. Sandkuhl, and A. V. Smirnov. Sparql query builders: Overview and comparison. In *BIR Workshops*, pp. 255–274, 2016. 8

[32] Q. Guo, F. Zhuang, C. Qin, H. Zhu, X. Xie, H. Xiong, and Q. He. A survey on knowledge graph-based recommender systems. *IEEE Trans. Knowl. Data Eng.*, 34(08):3549–3568, 2022. doi: 10.1109/TKDE.2020.3028705 1

[33] A. Hagberg, P. Swart, and D. S. Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab. (LANL), Los Alamos, NM (United States), 2008. 4

[34] J. Han, H. E. G. Le, and J. Du. Survey on NoSQL database. In *Proc. ICPCA*, pp. 363–366, 2011. doi: 10.1109/ICPCA.2011.6106531 8

[35] S. A. Helmers. *Microsoft Visio 2016 Step By Step: MS Visio 2016 Step by Step*. Microsoft Press, 2015. 9

[36] E. Hemberg, J. Kelly, M. Shlapentokh-Rothman, B. Reinstadler, K. Xu, N. Rutar, and U.-M. O'Reilly. Linking threat tactics, techniques, and patterns with defensive weaknesses, vulnerabilities and affected platform configurations for cyber hunting. *arXiv preprint arXiv:2010.00533*, 2020. doi: 10.48550/arXiv.2010.00533 4

[37] I. Herman, G. Melancon, and M. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Trans. Vis. Comput. Graph.*, 6(1):24–43, 2000. doi: 10.1109/2945.841119 6, 9

[38] A. Hogan, E. Blomqvist, M. Cochez, C. D'Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, and A. Zimmermann. Knowledge graphs. *ACM Comput. Surv.*, 54(4), 2021. doi: 10.1145/3447772 1, 2, 5

[39] S. R. Hong, J. Hullman, and E. Bertini. Human factors in model interpretability: Industry practices, challenges, and needs. *Proc. CHI*, 4(CSCW1), 2020. doi: 10.1145/3392878 2

[40] J. Huang, Y. Xi, J. Hu, and J. Tao. Flownl: Asking the flow data in natural languages. *IEEE Trans. Vis. Comput. Graph.*, 29(1):1200–1210, 2023. doi: 10.1109/TVCG.2022.3209453 8

[41] X. Huang, J. Zhang, D. Li, and P. Li. Knowledge graph embedding based question answering. In *Proc. WSDM*, p. 105–113. ACM, New York, 2019. doi: 10.1145/3289600.3290956 1

[42] F. Husain, R. Romero-Gomez, E. Kuang, D. Segura, A. Carolli, L. Liu, M. Cheung, and Y. Paris. A multi-scale visual analytics approach for exploring biomedical knowledge. In *Proc. VAHC*, pp. 30–35. IEEE Computer Society, Los Alamitos, 2021. doi: 10.1109/VAHC53616.2021.00010 2

[43] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer. Enterprise data analysis and visualization: An interview study. *IEEE Trans. Vis. Comput. Graph.*, 18(12):2917–2926, 2012. doi: 10.1109/TVCG.2012.219 2

[44] M. Kivelä, F. McGee, G. Melançon, N. H. Riche, and T. von Landesberger. Visual Analytics of Multilayer Networks Across Disciplines (Dagstuhl Seminar 19061). *Dagstuhl Reports*, 9(2):1–26, 2019. doi: 10.4230/DagRep.9.2.1 2

[45] K. Klein, J. F. Sequeda, H.-Y. Wu, and D. Yan. Bringing Graph Databases and Network Visualization Together (Dagstuhl Seminar 22031). *Dagstuhl Reports*, 12(1):67–82, 2022. doi: 10.4230/DagRep.12.1.67 2, 6

[46] S. Latif, S. Agarwal, S. Gottschalk, C. Chrosch, F. Feit, J. Jahn, T. Braun, Y. Tchenko, E. Demidova, and F. Beck. Visually connecting historical figures through event knowledge graphs. In *Proc. VIS*, pp. 156–160. IEEE Computer Society, Los Alamitos, 2021. doi: 10.1109/VIS49827.2021.9623313 2, 3, 9

[47] F. Lecue. On the role of knowledge graphs in explainable AI. *Semant. Web*, 11(1):41–51, 2020. doi: 10.3233/SW-190374 1, 8

[48] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry. Task taxonomy for graph visualization. In *Proc. BELIV*, p. 1–5. ACM, New York, 2006. doi: 10.1145/1168149.1168168 6, 9

[49] H. Li, Y. Wang, S. Zhang, Y. Song, and H. Qu. Kg4vis: A knowledge graph-based approach for visualization recommendation. *IEEE Trans. Vis. Comput. Graph.*, 28(01):195–205, 2022. doi: 10.1109/TVCG.2021.3114863 2, 9

[50] M. Lissandrini, D. Mottin, K. Hose, and T. B. Pedersen. Knowledge graph exploration systems: are we lost? In *CIDR*, vol. 22, pp. 10–13, 2022. 2, 5, 9

[51] Z. Liu and J. Heer. The effects of interactive latency on exploratory visual analysis. *IEEE Trans. Vis. Comput. Graph.*, 20(12):2122–2131, 2014. doi: 10.1109/TVCG.2014.2346452 6

[52] K. M. MacQueen, E. McLellan, K. Kay, and B. Milstein. Codebook development for team-based qualitative analysis. *CAM j.*, 10(2):31–36, 1998. doi: 10.1177/1525822X980100020301 3

[53] J. J. Miller. Graph database applications and concepts with Neo4J. *Proc. SAIIS*, 2324(36), 2013. 2

[54] R. Mitra, A. Narechania, A. Endert, and J. Stasko. Facilitating conversational interaction in natural language interfaces for visualization. In *Proc. VIS*, pp. 6–10, 2022. doi: 10.1109/VIS54862.2022.00010 8

[55] A. Mosca, S. Robinson, M. Clarke, R. Redelmeier, S. Coates, D. Cashman, and R. Chang. Defining an Analysis: A Study of Client-Facing Data Scientists. In J. Johansson, F. Sadlo, and G. E. Marai, eds., *EuroVis 2019 - Short Papers*. The Eurographics Association, 2019. doi: 10.2312/evs.20191173 2

[56] A. Narechania, A. Srinivasan, and J. Stasko. NL4DV: A Toolkit for generating Analytic Specifications for Data Visualization from Natural Language queries. *IEEE Trans. Vis. Comput. Graph.*, 2020. doi: 10.1109/TVCG.2020.3030378 8

[57] Neo4J. Neo4j Bloom. <https://neo4j.com/product/bloom/>. Accessed: 2023-03-24. 3, 4, 9

[58] C. Nobre, M. Streit, and A. Lex. Juniper: A tree+table approach to multivariate graph visualization. *IEEE Trans. Vis. Comput. Graph.*, 25(1):544–554, 2018. doi: 10.1109/TVCG.2018.2865149 6

[59] C. Partl, S. Gratzl, M. Streit, A. M. Wassermann, H. Pfister, D. Schmalstieg, and A. Lex. Pathfinder: Visual analysis of paths in graphs. *Comput. Graph. Forum*, 35(3):71–80, 2016. doi: 10.1111/cgf.12883 2, 6

[60] S. Passi and S. J. Jackson. Trust in data science: Collaboration, translation, and accountability in corporate data science projects. *Proc. CHI*, 2(CSCW), 2018. doi: 10.1145/3274405 2, 6

[61] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, and A. Miller. Language models as knowledge bases? In *Proc. EMNLP/IJCNLP*, pp. 2463–2473. ACL, Hong Kong, 2019. doi: 10.18653/v1/D19-1250 2

[62] A. Rossi, D. Barbosa, D. Firmani, A. Matinata, and P. Merialdo. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Trans. Knowl. Discov. Data*, 15(2), 2021. doi: 10.1145/3424672 9

[63] C. Rudin and J. Radin. Why Are We Using Black Box Models in AI When We Don't Need To? A Lesson From an Explainable AI Competition. *Harv. Bus. Rev.*, 2019. doi: 10.1162/99608f92.5a8a3a3d 8

[64] N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. Paritosh, and L. M. Aroyo. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proc. CHI*. ACM, New York, 2021. doi: 10.1145/3411764.3445518 2, 5, 9

[65] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. Vis. Comput. Graph.*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213 9

[66] 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*, 13(11):2498–2504, 2003. doi: 10.1101/gr.1239303 4

[67] B. Schneiderman. Response time and display rate in human performance with computers. *ACM Comput. Surv.*, 16(3):265–285, 1984. doi: 10.1145/2514.2517 6

[68] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2014. doi: 10.48550/arXiv.1312.6034 8

[69] Stardog. Stardog the enterprise knowledge graph platform. <https://www.stardog.com/>. Accessed: 2023-03-24. 2

[70] M.-A. D. Storey, D. Čubranović, and D. M. German. On the use of visualization to support awareness of human activities in software development: a survey and a framework. In *Proc. SoftVis*, pp. 193–202, 2005. doi: 10.1145/1056018.1056045 8

[71] A. Suh, M. Hajij, B. Wang, C. Scheidegger, and P. Rosen. Persistent homology guided force-directed graph layouts. *IEEE Trans. Vis. Comput. Graph.*, 26(1):697–707, 2020. doi: 10.1109/TVCG.2019.2934802 6

[72] H. Suresh, S. R. Gomez, K. K. Nam, and A. Satyanarayan. Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. *Proc. CHI*, 2021. doi: 10.1145/3411764.3445088 2, 4, 9

[73] I. Tiddi, F. Lévé, and P. Hitzler, eds. *Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges*, vol. 47 of *Studies on the Semantic Web*. IOS Press, 2020. 8

[74] E. Toussaint, P. Guagliardo, L. Libkin, and J. Sequeda. Troubles with nulls, views from the users. *Proc. VLDB Endow.*, 15(11):2613–2625, 2022. doi: 10.14778/3551793.3551818 5

[75] T. Von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. J. van Wijk, J.-D. Fekete, and D. W. Fellner. Visual analysis of large graphs: state-of-the-art and future research challenges. *Comput. Graph. Forum*, 30(6):1719–1749, 2011. doi: 10.1111/j.1467-8659.2011.01898.x 6

[76] A. Waagmeester, G. Stupp, S. Burgstaller-Muehlbacher, B. M. Good, M. Griffith, O. L. Griffith, K. Hanspers, H. Hermjakob, T. S. Hudson, K. Hybiske, et al. Wikidata as a knowledge graph for the life sciences. *Elife*, 9:e52614, 2020. doi: 10.7554/elife.52614 8

[77] Wikipedia. Wikidata Statistics. <https://www.wikidata.org/wiki/Wikidata:Statistics>. Accessed: 2023-03-24. 5

[78] K. Wongsuphasawat, Y. Liu, and J. Heer. Goals, process, and challenges of exploratory data analysis: an interview study. *arXiv preprint arXiv:1911.00568*, 2019. 2

[79] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proc. CHI*, p. 2648–2659. ACM, New York, 2017. doi: 10.1145/3025453.3025768 8

[80] G. Xiao, L. Ding, B. Cogrel, and D. Calvanese. Virtual knowledge graphs: An overview of systems and use cases. *Data Intell.*, 1(3):201–223, 2019. doi: 10.1162/dint\_a\_00011 1

[81] Y. Yoon, B. A. Myers, and S. Koo. Visualization of fine-grained code change history. In *IEEE VL/HCC*, pp. 119–126, 2013. doi: 10.1109/ VLHCC.2013.6645254 8