

# A COMMON SENSE VIEW OF KNOWLEDGE GRAPHS

Michael K. Bergman<sup>1</sup>, Coralville, Iowa USA

July 1, 2019

*AI3:::Adaptive Information* blog

Of late, it seems everywhere I turn some discussion or debate is taking place as to what a ‘knowledge graph’ is and how to define one. The semantic Web community not infrequently engages in such bursts of definition and vocabulary harangues. For a community ostensibly focused on ‘semantics’, definitions and terminology disputes sometimes seem all too common. Still, the interest in knowledge graphs (KGs) is quite real and represents an important visibility point for the semantic Web.

For example, in the last year there have been major sessions or workshops on KGs at the major semantic technology conferences of [WWW](#), [ESWC](#), [ISWC](#), and [Semantics 2019](#), and special conferences in [Dagstuhl](#), New York City ([Columbia](#)), [China](#) and [Cuba](#), among others I have surely missed. Ongoing for the past month has been a lengthy discussion on what is a knowledge graph on the W3C’s [semantic Web mailing list](#). Google Scholar, as another example, now lists over [16,000 academic papers](#) on knowledge graphs. These listings do not include the many individual blog and company posts on the subject. This blossoming of attention is great news.

In tracing the lineage of the ‘knowledge graph’ term, many erroneously date it to 2012 when Google (some claim) “coined” it for its structured entity-attribute information, now a common view on its search results page (with [Marie Curie](#) being one of the first-cited examples [\[1\]](#)). Though Google’s embrace has been important from a marketing and visibility viewpoint, the term ‘knowledge graph’ goes back to the 1970s, and the ideas behind a KG goes back even further. According to Pat Hayes [\[2\]](#), “The idea of representing, or at least displaying, knowledge as a graphical diagram (rather than as, say, a set of sentences) has a very old history. In its modern sense it goes back at least to 1885 (C S Peirce ‘existential graph’) and can probably be traced into medieval writings and earlier. (The Torah version of Genesis refers to a ‘tree of knowledge’.) It has been re-invented or rediscovered many times since, and seems to blossom in public (or at least academic) discussions with a periodicity of roughly 40 years.”

As someone who has written extensively on ‘[ontologies](#)’, the kissing cousins of knowledge graphs, and who has adopted the KG term as my preferred label, I thought it would be useful to survey the discussion and frame it in terms of some of these historical perspectives. Curious of how I was using the term myself, I searched my recent book and found that I use the term ‘ontology’ some 568 times, while using ‘knowledge graph’ 226 times [\[3\]](#). Though I can acknowledge some technical bases for distinguishing these two terms, in practice, I tend to use them interchangeably. I think many practitioners, including many ‘experts’ and academics with direct ties to the semantic Web, tend to do the same, for reasons I outline below. Though I am sure the world is not clamoring for still another definition of the term, I think we can benefit from a broad understanding of this increasingly important term from common-sense and pragmatic grounds. That is what I try to provide in this article.

---

<sup>1</sup>Email: [mike@mkbergman.com](mailto:mike@mkbergman.com)

## What is a Knowledge Graph?

If one searches on the phrase ‘what is a knowledge graph’, one finds there are some [99 references](#) on Google as a whole (my article will make it a cool 100!) and some [22 academic papers](#) on Google Scholar. Those are impressive numbers for such a specific phrase, and reflect the general questioning around the concept. Still, for those very interested, these numbers are also quite tractable, and I encourage you to repeat the searches yourself if you want to see the diversity of opinions, though most of the perspectives are captured herein.

I present and compare various definitions of ‘knowledge graph’ three sections below. However, to start, from a common-sense standpoint, we can easily decompose an understanding of the phrase from its composite terms. ‘Knowledge’ seems straightforward, but an inspection of the literature and philosophers over time indicates the concept is neither uniformly understood nor agreed. My own preference, which I discuss in Chapter 2 in my book [\[3\]](#), is from [Charles S. Peirce](#). His view of knowledge had a similar basis to what is known as justified, true belief ([JTB](#)). Peirce extended this basis with the notion that knowledge is information sufficiently believed to be acted upon. Nonetheless, we should assume that not all who label things as knowledge graphs embrace this precise understanding.

The concept of ‘graph’, the second composite term, has a precise and mathematical understanding as nodes (or vertices) connected by edges. Further, since knowledge is represented by language in the form of statements or assertions, the graph by definition is a directed and labeled one. In precise terminology, a knowledge graph has the form of a directed (mostly acyclic) graph (or [DAG](#)). (There can be cyclic or transitive relationships in a KG, but most are subsumptive or hierarchical in nature.)

So, while there is some dispute around what constitutes ‘knowledge’, there should be a pretty common-sense basis for what a knowledge graph is, at least at a general level.

Some attempt to further categorize KGs into general and domain ones (sometimes using the terms ‘vertical’ or ‘enterprise’). In the real world we also see KGs that are heavily biased toward instances and their attributes (the Google Knowledge Graph, for one) or reflect more complete knowledge structures with instances and a conceptual framework for organizing and reasoning over the knowledge space (such as our own [KBpedia](#) or the KGs used by [virtual intelligent agents](#)). I should note we see this same diversity in how one characterizes ‘ontologies’.

## Knowledge Graphs: Citations Over Time

As noted, the use of the phrase ‘knowledge graph’ did not begin with Google. In fact, we can trace back the earliest references to the phrase to 1973. Here is a graph of paper citations for the phrase in the period 1970 – 2019 [\[4\]](#):




Figure 1: 'Knowledge Graph' Citations, 1970 – 2019

Coincident with Google’s adoption, we see a pretty significant spike in citations beginning in 2012. Today, we are seeing about 5,000 citations yearly in the academic literature.

Because later citations swamp out the numbers of earlier ones, we can plot these citations on a log scale to see earlier references:




Figure 2: 'Knowledge Graph' Citations (log basis), 1970 – 2019

The shape of this log curve is consistent with a standard power-law distribution. The earliest citations we see are from 1973, when there were two. One was in the context of education and learning [5]; the other was based on linguistic analysis [6]. The first definition of ‘knowledge graph’ was presented by Marchi and Miquel in 1974 [7]. There were few further references for the next decade. However, that began to change in the mid-1980s, which I elaborate upon below.

## Knowledge Graphs: A Marketing Perspective

As we will see in a few moments, not all knowledge graphs are created the same. Yet, as the charts above show, one KG was the first gorilla to enter the room: Google. Though there were many mentions for 40 years prior, and quite a few in the prior decade, the announcement of Google’s [Knowledge Graph](#) (title case) in 2012 totally changed perceptions and visibility. It is an interesting case study to dissect why this large impact may have occurred.

The first obvious reason is Google’s sheer size and market clout. Further, Google’s announcement was accompanied by some nifty graphics, which can be [viewed on YouTube](#), that showed connections between entities that conveyed a sense of dynamism and navigation. Still, not all announcements or new initiatives by Google gain the traction that knowledge graphs did. Clearly, more was at work.

The first additional reason, I think, was the groundwork for the ideas of graphs and connections that had come from the [linked data](#) and [semantic Web](#) communities in the 5-10 years prior. People had been prepped for the ideas of connections and that, when visualized, they would take on the form of a graph.

Another additional reason came from within the community itself. The first term for a knowledge graph was ‘ontology.’ (Though as we will see in the next section, not everyone agrees that KGs and ontologies can be used synonymously.) Yet all of us who were advocates and practitioners well knew that the idea of an ontology was difficult to explain to others. The term seemed to have both technical and abstract connotations and, when first introduced, always required some discussion and definition. The notion of a knowledge graph, on the other hand, is both descriptive and intuitive.

Lastly, I think there was another factor at work in the embrace of knowledge graphs. The semantic Web community had been split for some time between two camps. One camp saw the need for formal ontologies and were often likely to use the OWL language, at least for formal conceptual schemas. Another camp leaned more to diversity and decentralized approaches to vocabularies, and tended to prefer RDF (and, later, linked data). At the same time, this latter camp recognized the need for connectivity and data integration. The idea of a ‘knowledge graph’ was able to sidestep these differences and provide a common terminology for both camps. I think each camp embraced the term, in part, for these different reasons. Both camps, however, seemed to share the view that the term ‘knowledge graph’ offered some strong marketing advantages, not least of which was the implicit endorsement of Google.

## Knowledge Graphs: Definitions

Once Google popularized the ‘knowledge graph’ term in 2012, a number of authors began to assemble various definitions of the concept. One could argue that this scramble to define the term was in part an attempt to bound the term consistent with prior prejudices. Another possible motivation was the desire to give precision to a heretofore imprecise term. In any case, the proliferation of definitions led some researchers to compile recent and historical examples. Notable among these have been McCusker *et al.* [8], Ehrlinger and Wöß [9], Gutierrez [10], and Hogan *et al.* [11]. I have built upon their compilations, plus added new entries based on the research noted above, to produce the following table of ‘knowledge graph’ *definitions*, reaching back to the earliest ones:

Source	Year	Definition
Marchi and Miquel <a href="#">[7]</a>	1974	A mathematical structure with vertices as knowledge units connected by edges that represent the prerequisite relation
Hoede <a href="#">[12]</a>	1974	See <a href="#">[12]</a> and text and Bakker and de Vries below
Bakker <a href="#">[13]</a>	1987	A knowledge graph is a labeled directed graph $D(P,A)$ with (a) $P$ is a set of points that represent concepts, relations or frameworks; (b) $A \subseteq P \times P$ is a set of arcs that form the connections between the entities. An arc exists for $p_1 \in P$ to $p_2 \in P$ if: (b1) $p_1$ represents a concept that is the tail of a relation represented by $p_2$ ; (b2) $p_2$ represents a concept that is the head of a relation represented by $p_1$ . (b3) $p_1$ represents an element in the contents of a framework and $p_2$ represents the FPAR-relation. (b4) $p_2$ is a framework and $p_1$ represents the accompanying FPAR-relation
de Vries <a href="#">[14]</a>	1989	A directed graph that distinguishes three types of asserted relationships: (1) an object has a certain property, (2) an object is an instance of another object, (3) a change in a property of an object leads to a change in another property of that object
James <a href="#">[15]</a>	1992	A knowledge graph is a kind of semantic network. . . . One of the essential differences between knowledge graphs and semantic networks is the explicit choice of only a few types of relations
Zhang <a href="#">[16]</a>	2002	A new method of knowledge representation, [which] belongs to the category of semantic networks. In principle, the composition of a knowledge graph is including concept (tokens and types) and relationship (binary and multivariate relation)
Popping <a href="#">[17]</a>	2003	A particular kind of semantic network
Singhal (Google) <a href="#">[1]</a>	2012	A graph that understands real-world entities and their relationships to one another: things, not strings
Ehrlinger and Wöß <a href="#">[9]</a>	2015	A knowledge graph acquires and integrates information into an ontology and applies a reasoner to derive new knowledge
Krötzsch and Weikum <a href="#">[18]</a>	2016	Knowledge graphs are large networks of entities, their semantic types, properties, and relationships between entities
Krötzsch <a href="#">[19]</a>	2017	Are characterized by several properties that together distinguish them from more traditional knowledge management paradigms: (1) Normalization: Information is decomposed into small units of information, interpreted as edges of some form of graph. (2) Connectivity: Knowledge is represented by the relationships between these units. (3) Context: Data is enriched with contextual information to record aspects such as temporal validity, provenance, trustworthiness, or other side conditions and details
Paulheim <a href="#">[20]</a>	2017	A knowledge graph (i) mainly describes real world entities and their interrelations, organized in a graph, (ii) defines possible classes and relations of entities in a schema, (iii) allows for potentially interrelating arbitrary entities with each other and (iv) covers various topical domains
Shao <i>et al.</i> <a href="#">[21]</a>	2017	A knowledge graph is a graph constructed by representing each item, entity and user as nodes, and linking those nodes that interact with each other via edges.
Villazon-Terrazas <i>et al.</i> <a href="#">[22]</a>	2017	A knowledge graph is a set of typed entities (with attributes) which relate to one another by typed relationships. The types of entities and relationships are defined in schemas that are called ontologies. Such defined types are called vocabulary. A knowledge graph is a

structured dataset that is compatible with the RDF data model and has an (OWL) ontology as its schema.

Wilcke <i>et al.</i>	<a href="#">[23]</a>	2017	A data model used in the Semantic Web . . . based on three basic principles: 1. Encode knowledge using statements. 2. Express background knowledge in ontologies. 3. Reuse knowledge between datasets
Bianchi <i>et al.</i>	<a href="#">[24]</a>	2018	<ul style="list-style-type: none"> <li>• Large knowledge bases</li> <li>• Entities classified using types</li> <li>• Types organized in sub-types graphs</li> <li>• Binary relationships between entities</li> <li>• Semantics and inference via rules/axioms</li> <li>• Semantic similarity with lexical, topological and other feature-based approaches</li> </ul>
McCusker <i>et al.</i>	<a href="#">[8]</a>	2018	An Unambiguous Graph [a graph where the relations and entities are unambiguously identified] with a limited set of relations used to label the edges that encodes the provenance, especially justification and attribution, of the assertions
Xiong	<a href="#">[25]</a>	2018	A data resource that contains entries (‘entities’) that have their own meanings and also information (‘knowledge graph semantics’) about those entries
Bellomarini <i>et al.</i>	<a href="#">[26]</a>	2019	A semi-structured datamodel characterized by three components: (i) a ground extensional component, that is, a set of relational constructs for schema and data (which can be effectively modeled as graphs or generalizations thereof); (ii) an intensional component, that is, a set of inference rules over the constructs of the ground extensional component; (iii) a derived extensional component that can be produced as the result of the application of the inference rules over the ground extensional component (with the so-called “reasoning” process)
d’Amato <i>et al.</i>	<a href="#">[27]</a>	2019	A graph-based structured data organization, endowed with formal semantics ( <i>i.e.</i> , a graph-based data organization where a schema and multiple labelled relations with formal meaning are available)
Columbia University	<a href="#">[28]</a>	2019	An organized and curated set of facts that provide support for models to understand the world
Hogan <i>et al.</i>	<a href="#">[11]</a>	2019	A graph of data with the intent to compose knowledge
Kejriwal	<a href="#">[29]</a>	2019	A graph-theoretic representation of human knowledge such that it can be ingested with semantics by a machine; a set of triples, with each triple intuitively representing an ‘assertion’
Morrison	<a href="#">[30]</a>	2019	A knowledge base that’s made machine readable with the help of logically consistent, linked graphs that together constitute an interrelated group of facts
PoolParty	<a href="#">[31]</a>	2019	A graph-theoretic representation of human knowledge such that it can be ingested with semantics by a machine; a set of triples, with each triple intuitively representing an ‘assertion’
Tran and Takasu	<a href="#">[32]</a>	2019	A collection of triples, with each triple $(h,t,r)$ denoting the fact that relation $r$ exists between head entity $h$ and tail entity $t$
Vidal <i>et al.</i>	<a href="#">[33]</a>	2019	A knowledge graph is presented as the intersection of the formal models able to represent facts of various types and levels of abstraction using a graph-based formalism

*Table 1: Selected Definitions of ‘Knowledge Graph’ Over Time*

As the usage charts showed, more definitions have been forthcoming since the Google KG announcement than came before. There has also been a concerted effort to provide more “precise” understandings of what the KG concept means since its broad embrace by the semantic Web community. That effort really began in earnest about 2016. A few conferences have been devoted to the topic.

## Knowledge Graphs: Concepts, Connections, Contradictions

The first definition of ‘knowledge graph’ by Marchi and Miquel in 1974 [7] found so far in the literature is really not too far from our common-sense understanding of knowledge assertions in the form of a directed graph. What constituted knowledge was left undefined in this paper.

A more focused treatment of knowledge graphs began in 1982 with efforts by the Universities of Twente and Groninger [12]. Cornelis Hoede was the lead researcher on an effort to capture science or medical knowledge in the form of graphs. A key component of this research, which stretched over decades, was the parsing and analysis of human language. Some of the key students and their theses that derived from these efforts were Bakker [13], de Vries [14], van de Berg [34], and Zhang [16], among others. (The preface to Zhang’s thesis provides a nice overview of the history of these students.) Popping was another of the co-researchers [17]. Many of these provided definitions of KGs, as captured in Table 1 above.

These researchers readily acknowledged that their efforts were a derivation of [semantic networks](#), a graph-based AI approach dating from the mid-1950s, and had close relationships to the [conceptual graphs](#) developed by John Sowa [35], though they were developed independently. Interestingly, both Hoede and students and Sowa traced the lineage of their respective efforts back to the [existential graphs](#) of Charles Peirce from 1885 [36 for Hoede and his students] to 1909 [37 and MS 514 for Sowa]. Eventually Hoede and Sowa published together and Hoede came to describe knowledge graphs as a subset of conceptual graphs.

Though the knowledge graphs of Hoede and his students initially employed a relatively small number of relationships (such as part of, a kind of, a cause of), they came to understand that the graphs should be embedded in “frameworks” (what today we call ‘ontologies’) and that they could be constructed of multiple sub-graphs [13]. According to one of Hoede’s earliest students, René Ronald Bakker, “The choice of a graph for the representation of empirical scientific knowledge is appealing because, besides having the advantage of being an ‘obvious representation’, the well developed mathematical theory of graphs can be used for the analysis and structuring of the represented knowledge.” [13, p. 22]. The group also acknowledged that the graphs could be used as the knowledge base of an expert system or for knowledge extraction.

The work of Hoede and students was not casual and extended over at least 25 years [12]. Though terminology was different then, and today’s standards were lacking, it is striking the degree to which those efforts both reflect Peirce and many current efforts. Bakker, in the first thesis for the group in 1987, which mentions KGs more than 150 times, presented the following diagram:



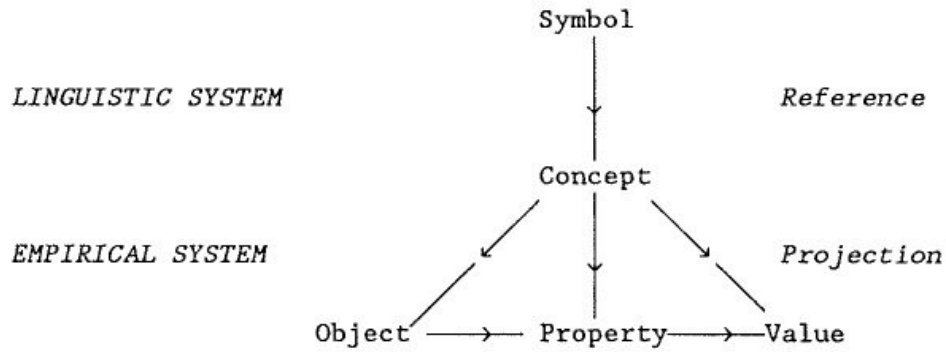


Figure 3: Knowledge Graph Schema per Hoede, approx. 1982 [13]

We can see the reflection of a ‘triple’ at the bottom of the diagram, including how objects may roll up into concepts. The figure also acknowledges the relationship to language and symbol representations, striking in its relation to current information extraction and consistent with Peirce’s [semiosis](#). The group further recognized the importance of objects being either subjects or objects of assertions, as well as how to deal with language ambiguity [38].

Much of this perspective appears to have been overlooked by the time Google made its KG announcement in 2012, and thus by many subsequent researchers.

Subsequent to 2012 definitions of KGs have added other considerations. Knowledge bases, large scales, use of particular languages ([RDF](#) or [OWL](#), for example), provenance, an entity focus, type classification, context, machine readability, dataset linkages, data consistency, formal models, formal semantics, use of reasoners, use of triples, use of rules or axioms, etc., are all aspects mentioned by one or more of the definers listed in Table 1. Granted, any of these may be useful additions to make a knowledge graph more usable or defensible, but it is unclear if any of them are a strict requirement to be a member of the genus.

Partially, in reaction to this tightening of requirements, Hogan *et al.* [11] have taken the approach to broaden and make the definition of KGs as flexible as possible as “a graph of data with the intent to compose knowledge.” We seem to be resurrecting the old tensions between formality and flexibility that characterized the early evolution of the semantic Web.

Since Peirce has clearly influenced many of the ideas leading to knowledge graphs, their use, and their representation, we could perhaps look to his ethics of terminology to provide guidance for how to adjudicate these different definitions [39]. Peirce grants primacy to the first formulator of a concept and defers to her naming conventions. By this basis, we would perhaps look to Hoede and his group for our guidance. However, given much of this work was neglected, and the Google announcement set a new basis for a large-scale entity repository characterized by types and attributes, our guidance may not be so clear-cut.

## Knowledge Graphs: What They Are Not

One observation arising from this survey is that knowledge graphs are not precise. They are not precise in scope, intent, use, or construction. That means, to me, that certain proscriptions offered in KG-wide definitions are not appropriate:

- Provenance is not required, though what makes information believable so as to engender action is a possible condition for ‘knowledge’



- Context is useful, and can help in disambiguation and understanding knowledge, but is not strictly required
- A specific language or data model, such as RDF, concept graphs, or OWL, is not required
- Types or type classification are not required
- Neither instances, nor attributes, nor concepts, nor specific relations are required, but one or two is
- A knowledge graph need not be both machine- and human readable
- No specific schema or formal logic is required
- Curation is not required, though again believability of the information is important
- A specific scope, broad or narrow, is not required, and
- Statements in the knowledge graph need not be ‘triples’, but they do need to be some form of knowledge assertion.

## Knowledge Graphs: Best Practices and a Spectrum of Uses

This survey has pointed out a few things regarding knowledge graphs. We see that the term has been around for many years, and has been used differently and with imprecision by many authors. The diversity of knowledge graph use extends from mostly attribute characterizations of instances (such as the Google Knowledge Graph) to mostly conceptual frameworks akin to upper-level ontologies. Attempts to provide a “precise” definition of KGs appear likely doomed to failure, and appear to be more a reflection of prior prejudices than a derivation of some discovered inherent meaning.

About the most we can say about knowledge graphs in general, in keeping with their constituent terms, is that they are a representation of knowledge (however defined) in the structural form of a directed (mostly acyclic) graph. Recent attempts to add more rigor and precision to the definition of a KG beyond this do not appear warranted. Because our understanding of KGs remains somewhat vague, we should ask and expect authors of KGs to define the scope and basis of the knowledge used in their graphs.

In this diversity, KGs are not much different than ontologies, with a similarly broad and contextual use. Indeed, in the words of Ehrlinger and Wöß [\[9\]](#), “Graph-based knowledge representation has been researched for decades and the term knowledge graph does not constitute a new technology. Rather, it is a buzzword reinvented by Google and adopted by other companies and academia to describe different knowledge representation applications.”

Still, there is a common-sense intuitiveness to the term ‘knowledge graph’ and it appears far superior to the label ‘ontology’ as a means of describing these knowledge structures to the general public. For these reasons, I think the term has ‘legs’ and is likely to be the term of use for many years to come. Yet, given the diversity of use, I also agree with the advice of Hogan *et al.* [\[11\]](#) who “strongly encourage researchers who wish to refer to ‘Knowledge Graphs’ as their object of study to rigorously define how they instantiate the term as appropriate to their investigation.” While we may not find a precise definition across all uses of knowledge graphs, we should strive to define how they are used in specific contexts.

In our own contexts based on Cognonto’s work and our open-source KBpedia, we tend to define knowledge graphs as complete knowledge bases of concepts and instances and their attributes, coherently organized into types and a logical and computable graph, and written in RDF, [SKOS](#), and OWL 2. (Amongst other researchers, our usage perhaps comes closest to that for Pan *et al.* [\[40\]](#).) The actionability of the ‘knowledge’ in these graphs comes not so much from the explicit tagging of items with provenance metadata as from the use of trustworthy sources and manual vetting and logical testing during graph construction. This definition is appropriate given our use of knowledge graphs, more or less in a centralized form, for data interoperability, knowledge management,

and knowledge-based artificial intelligence. Knowledge graphs designed for other purposes — such as for less discriminant Web-wide navigation or retrieval or simple vocabularies — may employ different constructions.

Thus, knowledge graphs, like ontologies, have a broad range of applications and constructions. When one hears the term, it is less important to reflect on some precise understanding as to realize that human language and knowledge is being presented in a connected, graph form. That, I believe, is the better practical and common-sense understanding of KGs.

## Acknowledgements

This article was originally posted on the *AI3::Adaptive Information* Web site at <http://www.mkbergman.com/2244/a-common-sense-view-of-knowledge-graphs/>. This version has been edited and reformatted slightly for PDF distribution. We thank Cognonto Corporation for making this content freely available.

- 
- [1] Amit Singhal, 2012. “Introduction the Knowledge Graph: Things, Not Strings,” originally published May 16, 2012, on the Google blog. Now moved to <https://www.blog.google/products/search/introducing-knowledge-graph-things-not/>. Retrieved June 28, 2019.
- [2] See <https://lists.w3.org/Archives/Public/semantic-web/2019Jun/0052.html>.
- [3] Michael Bergman, 2018. *A Knowledge Representation Practionary: Guidelines Based on Charles Sanders Peirce*, Springer, 462 pp.
- [4] The data is from citations on Google Scholar using a date-range query for ‘knowledge graph’ for citations up to the current year for each year (ex., [https://scholar.google.com/scholar?q=%22knowledge+graph%22&hl=en&as\\_sdt=0%2C5&as\\_ylo=&as\\_yhi=1973](https://scholar.google.com/scholar?q=%22knowledge+graph%22&hl=en&as_sdt=0%2C5&as_ylo=&as_yhi=1973)), and then subtracting prior years’ citations to get the current year count. Note that Google provides about 15 citations prior to 1973, which are all erroneous. These were subtracted from subsequent counts.
- [5] E.W. Schneider, 1973. “Course Modularization Applied: The Interface System and Its Implications For Sequence Control and Data Analysis,” 21 pp. paper presented at the Association for the Development of Instructional Systems, Chicago, Illinois, April 1972.
- [6] Peter Kümmel, 1973. “An Algorithm of Limited Syntax Based on Language Universals,” in *Proceedings of the 5th Conference on Computational Linguistics,-Volume 2*, pp. 225-247. Association for Computational Linguistics, 1973.
- [7] E. Marchi and O. Miguel, 1974. “On the Structure of the Teaching-learning Interactive Process,” *International Journal of Game Theory*, 3(2):83–99.
- [8] J.P. McCusker, J. Erickson, K. Chastain, S. Rashid, R. Weerawarana, and D. McGuinness, 2018. “What is a Knowledge Graph?,” *Semantic Web J.* (submitted, but not accepted); see <http://www.semantic-web-journal.net/content/what-knowledge-graph>.
- [9] Lisa Ehrlinger and Wolfram Wöß, 2016. “Towards a Definition of Knowledge Graphs,” *SEMANTiCS (Posters, Demos, SuCCESS)* 48.
- [10] Claudio Gutierrez, 2019, “Concise Account of the Notion of Knowledge Graph,” in Piero Andrea Bonatti, Stefan Decker, Axel Polleres, and Valentina Presutti, eds., *Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web*, pp. 43-45, Dagstuhl Seminar 18371, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] Aidan Hogan, Dan Brickley, Claudio Gutierrez, Axel Polleres, and Antoine Zimmerman, 2019, “(Re)Defining Knowledge Graphs,” in Piero Andrea Bonatti, Stefan Decker, Axel Polleres, and Valentina Presutti eds., *Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web*, pp. 74-79, Dagstuhl Seminar 18371, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [12] Sri Nurdianti and Cornelis Hoede, 2008. “25 Years Development of Knowledge Graph Theory: The Results and the Challenge,” *Memorandum 1876*, September 2008, Department of Applied Mathematics, The University of Twente, Enschede, The Netherlands, 10 pp; available from <https://core.ac.uk/download/pdf/11468596.pdf>.
- [13] R. R. Bakker, 1987. *Knowledge Graphs: Representation and Structuring of Scientific Knowledge*, Ph.D. thesis, University of Twente, Enschede, ISBN 9001963-4.

- [14] P.H. de Vries, 1989. *Representation of Science Texts in Knowledge Graphs*, Ph.D. thesis, University of Groningen, Groningen, The Netherlands ISBN 90-367-0179-1.
- [15] P. James, 1992. "Knowledge Graphs," in Reind P. Van de Riet and Robert A. Meersman, eds., *Linguistic Instruments in Knowledge Engineering*, Proceedings of the 1991 Workshop on Linguistic Instruments in Knowledge Engineering, Tilburg, The Netherlands, 17-18 January 1991, Elsevier Science Inc., 1992. [P. James is likely a pseudonym for other authors]
- [16] Liecal Zhang, 2002. *Knowledge Graph Theory and Structural Parsing*, Ph.D. thesis, Twente University Press, 232 pp.
- [17] Roel Popping, 2003. "Knowledge Graphs and Network Text Analysis," *Social Science Information* 42 (1): 91-106.
- [18] Markus Krötzsch and Gerhard Weikum, 2016. "Web Semantics: Science, Services and Agents on the World Wide Web," *Journal of Web Semantics: Special Issue on Knowledge Graphs*, Vol 37-38, pp 53-54.
- [19] Markus Krötzsch, 2017. "Ontologies for Knowledge Graphs?," in Alessandro Artale, Birte Glimm and Roman Kontchakov, eds., *Proceedings of the 30th International Workshop on Description Logics*, Montpellier, France, July 18-21, 2017, Vol. CEUR-1879.
- [20] Heiko Paulheim, 2017. "Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods," *Semantic Web* 8 (3): 489-508.
- [21] Lixu Shao, Yucong Duan, Xiaobing Sun, Honghao Gao, Donghai Zhu, and Weikai Miao, 2017. "Answering Who/When, What, How, Why through Constructing Data Graph, Information Graph, Knowledge Graph and Wisdom Graph," in SEKE, pp. 1-6.
- [22] Boris Villazon-Terrazas, Nuria Garcia-Santa, Yuan Ren, Alessandro Faraotti, Honghan Wu, Yuting Zhao, Guido Vetere, and Jeff Z. Pan, 2017. "Knowledge Graph Foundations," in Jeff Z. Pan, Guido Vetere, Jose Manuel Gomez-Perez, and Honghan Wu, eds., *Exploiting Linked Data and Knowledge Graphs in Large Organisations*, Springer, Heidelberg.
- [23] Xander Wilcke, Peter Bloem, and Victor De Boer, 2017. "The Knowledge Graph as the Default Data Model for Learning on Heterogeneous Knowledge," *Data Science* 1 (1-2): 39-57.
- [24] Frederico Bianchi, Matteo Palmonari, and Debora Nozza, 2018. "Towards Encoding Time in Text-Based Entity Embeddings," in Denny Vrandečić, Kalina Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee and Elena Simperl, eds., *Proceedings of the 17th International Semantic Web Conference*, Monterey, CA, USA, October 8-12, 2018, pp. 56-71, Springer.
- [25] Chenyan Xiong, 2018. "Text Representation, Retrieval, and Understanding with Knowledge Graphs," Ph.D. thesis, University of Massachusetts, Amherst, MA.
- [26] Luigi Bellomarini, Daniele Fakhoury, Georg Gottlob, and Emanuel Sallinger, 2019. "Knowledge Graphs and Enterprise AI: the Promise of an Enabling Technology," in *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pp. 26-37.
- [27] Claudia d'Amato, Sabrina Kirrane, Piero Bonatti, Sebastian Rudolph, Markus Krötzsch, Marieke van Erp, and Antoine Zimmermann, 2019. "Foundations," in Piero Andrea Bonatti, Stefan Decker, Axel Polleres, and Valentina Presutti, eds., *Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (Dagstuhl Seminar 18371)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [28] Columbia University, 2019. Knowledge Graph Conference, May 7-8, 2019, New York City. See <http://sps.columbia.edu/executive-education/knowledge-graph-conference>.
- [29] Mayank Kejriwal, 2019. "What Is a Knowledge Graph?," in *Domain-Specific Knowledge Graph Construction*, pp. 1-7. Springer, Cham.
- [30] Alan Morrison, 2019. From <https://www.quora.com/What-is-the-difference-between-a-knowledge-graph-and-a-graph-database/answer/Alan-Morrison>, June 21, 2019. Retrieved June 28, 2019.
- [31] PoolParty, Ltd., 2019. *Knowledge Graphs: Transforming Data into Knowledge*, white paper from Semantic Web Company, Jan 2019, 22 pp. Downloaded June 28, 2019.
- [32] Hung Nghiep Tran and Atsuhiko Takasu, 2019. "Analyzing Knowledge Graph Embedding Methods from a Multi-Embedding Interaction Perspective." *arXiv preprint arXiv:1903.11406*.

- [33] Maria-Esther Vidal, Kemele M. Endris, Samaneh Jazashoori, Ahmad Sakor, and Ariam Rivas, 2019. “Transforming Heterogeneous Data into Knowledge for Personalized Treatments—A Use Case,” *Datenbank-Spektrum*: 1-12.
- [34] H. van den Berg, 1993. *Knowledge Graphs and Logic: One of Two Kinds*, Ph.D. thesis, University of Twente, Enschede, The Netherlands, ISBN 90-9006360-9.
- [35] John F. Sowa, 1976. “Conceptual Graphs for a Data Base Interface,” *IBM Journal of Research and Development* 20, (4): 336-357.
- [36] C. S. Peirce, 1885, “On the Algebra of Logic,” *American Journal of Mathematics*, 7: 180-202, 1885.
- [37] C. S. Peirce (with J. Sowa), 2017. “Existential Graphs (MS 514 of 1909), with commentary by John Sowa,” see <http://www.jfsowa.com/peirce/ms514.htm>, last modified October 29, 2017; retrieved June 30, 2019.
- [38] From Bakker [13, p. 16]: “According to this conception objects, properties and values are base elements of empirically oriented scientific theories and, by consequence, of systems that represent scientific knowledge. Within the empirical system, a property is the result of the identification of an aspect of an object, and a value is produced by the measurement of a property. Symbols and concepts (including subsets of natural language, logic, and mathematics) are the linguistic elements to denote these base elements. Objects, properties and values are related to concepts by a realization or projection of the linguistic concept or symbol into an empirical, ‘real world’ system. Objects, properties and values must have been defined on a linguistic level to be able to speak about empirical elements. Note that a single concept can play different roles. For example, a ‘sphere’ may be an object in relation to its property ‘radius’, or it may be a value in relation to a property ‘shape’.”
- [39] C. S. Peirce, 1903. “The Ethics of Terminology,” in *Syllabus of Certain Topics of Logic*, pp. 1931-1959, Alfred Mudge & Son, Boston; see MS 434 and parts of MS 433 and CP 2.219-226.
- [40] Jeff Z. Pan, Guido Vetere, Jose Manuel Gomez-Perez, and Honghan Wu, eds. *Exploiting Linked Data and Knowledge Graphs in Large Organisations*, Heidelberg: Springer, 2017.