

Ontology Summit 2018 Communiqué

Contexts in Context

Kenneth Baclawski¹, Mike Bennett², Gary Berg-Cross³,
Cory Casanave⁴, Donna Fritzsche⁵, Joanne Luciano⁶,
Todd Schneider⁷, Ravi Sharma⁸, Janet Singer⁹, John Sowa¹⁰,
Ram D. Sriram¹¹, Andrea Westerinen¹², David Whitten¹³

¹Northeastern University, Boston, MA USA,

²Hypercube Limited, London, UK,

³RDA US Advisory Group, Troy, NY USA,

⁴Model Driven Solutions, Herndon, VA USA,

⁵Hummingbird Design, Chicago, IL USA,

⁶Indiana University, IN USA,

⁷Engineering Semantics, Fairfax, VA USA,

⁸Senior Enterprise Architect, Elk Grove, CA USA,

⁹INCOSE, Scotts Valley, CA USA,

¹⁰Kyndi, Inc. USA,

¹¹National Institute of Standards & Technology, Gaithersburg, MD USA,

¹²Two Six Labs, LLC USA,

¹³WorldVistA USA

Abstract

It is generally agreed that the interpretation of information, in any form, is context-dependent. The goal of the recent Ontology Summit 2018 was to explore the various relationships spanning ontology and context. This article is the end product of the summit and associated symposium. It describes motivation for creating explicit, formal context specifications, and discusses approaches for finding, understanding and formalizing context. We present this work and associated materials with the goal to foster the research and development of approaches to contexts and to drive towards context-aware solutions which can be incorporated into both knowledge engineering processes and ontology design best practices.

Keywords: ontologies; context; situations; microtheories; Open Knowledge Network

1 Introduction

In just the last few years, more data has been produced, at a far greater rate, and with far more complexity, than in all of previous human history. The vast majority of this data is context-dependent. While there has been progress on the development and deployment of ontology-based methods for dealing with data semantics, there has not been as much progress on specifying context. Given that data is nearly always context-dependent, specifying data formally, even with very rich ontologies, will have limited usefulness if the context is only informally known, if it is known at all.

The purpose of this Communiqué is to identify some of the prevailing viewpoints and the major issues and research problems of the formalization of contexts. While any subject matter will have a context, the main focus of the Communiqué is on the contexts of digital information and data. We begin with some background for contexts in Section 2 and the motivations for formalizing a context in Section 3. These sections are followed by a survey of some of the prevailing viewpoints and major approaches to formalizing context in Section 4. The Open Knowledge Network (OKN) is an initiative that makes use of the major approaches to context, and so it furnishes a kind of case study for the tools and techniques that can be used for contexts. The OKN is described in Section 5. The major issues and research questions are summarized in Section 6. This Communiqué is the end product of the Ontology Summit 2018, and we acknowledge the many individuals and organizations who contributed to the summit in Section 7.

2 Background

In general, a context is commonly understood to be the circumstances that form the setting for an event, statement, process, or idea, and in terms of which the event, statement, process, or idea can be understood and assessed. Thus for utterance statements we often talk of the linguistic context of what is being expressed. In addition, there may be a physical context, circumstance or state of affairs in the real world that provides context for uttered statements. Some examples of synonyms or alternate terms that have the flavor of context include circumstances, conditions, factors, perspective, scope, state of affairs, situation, background, scene, setting, and frame(s) of reference.

The context needed to understand any subject matter may include information of any kind, general or specific. For this reason, any reasoning about context is at a metalevel: it's not about the current subject matter, but about the methods for finding some implicit information that should be added to the interpretation of the subject.

The information needed for context can come from several sources. The immediate context includes the sentences that precede or follow the current sentence. The background knowledge includes information about the subject matter that is assumed by the speaker, listener, viewer, author or reader. The situation includes the time, place, and audience or readers. All these sources of information may change at different points in a document or discourse.

We may speak of physical situations as the context for events, and ontologies can model the concept of "situation" using, for example, the Situation Theory Ontology (STO) (Kokar, Matheus, & Baclawski, 2009). However, a situation may not be adequate as a context by itself. Usually, this means that one must at least specify answers to the six basic questions; namely, Who, What,

When, Where, Why and How.¹ Moreover, there are many senses of “context” beyond physical situations (e.g., social context), and thus situations may require much more than just the answers to the six basic questions.

One approach to understanding a context is to consider attitudes and perspectives. Pat Hayes distinguishes two intellectual traditions, each of which brings a different collection of unspoken assumptions: semantic linguistics and cognitive linguistics (Hayes, 1997). Within each of these traditions, Pat Hayes identified two senses for contexts, for a total of four senses:

1. Physical Context,
2. Linguistics/Topic Context
3. Conceptual Context
4. Deductive Context

John Sowa further distinguishes four senses of linguistic context as follows: (Sowa, 2017)

1. the text or discourse;
2. the situation;
3. common background knowledge; and
4. the intentions of the participants.

However, John Sowa admitted that these senses could be subdivided endlessly for any purpose. He also recognized that there are Actual, Modal, and Intentional Contexts (Sowa, 2017).

As we have noted above, a context is anything that impacts the interpretation or truth value of something else. Cory Casanave proposed a pattern for understanding and formalizing context as a mediator (Casanave, 2018), illustrated in Figure 1. In this approach, a context acts as a mediator between a set of propositions (or rules) and the things that are contextualized. The essential relations are that a rule or proposition holds within a context, and that the context provides the framework to understand a set of things (which may also be propositions). If context C is true for an act of interpretation, the propositions that hold within C hold for all things that C contextualizes. This pattern works for many contextual dimensions such as time, location and provenance.

An example that illustrates the potential use of context as a mediator is the distinction between reference and application ontologies in the biomedical domain. An application ontology is one that has been developed by domain experts for use in specific types of applications. Unlike application ontologies, reference ontologies are not designed for any specific application, but are intended to be re-used in multiple application contexts. One can then, ideally, generate application ontologies from one or more reference ontologies using formal methods for specifying the transformation from reference ontologies to application ontologies (Brinkley, Suci, Detwiler, Gennari, & Rosse, 2006). The formal specification of the transformation represents the context of the application ontology as in Figure 1. Unfortunately, in practice, this process can be difficult to achieve as noted in (Malone & Parkinson, 2010).

¹These are sometimes called the “six Ws” in spite of one of them not starting with W.

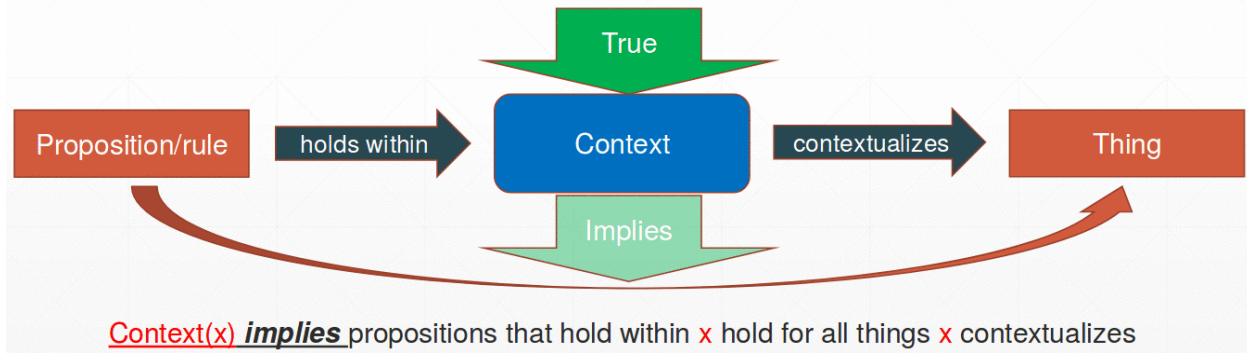


Figure 1: Context as a Mediator from (Casanave, 2018, Slide 15)

There are various sources of information which could be employed for determining the context of some subject matter. These include the scope, competency questions, business requirements, “use cases,” provenance and other documents that were used in the subject matter development process. Still other potential inputs include the intent and perspective of the stakeholders and developers. Such inputs should, in theory, be useful for formally specifying the context, as well as developing the subject matter, since these inputs convey some understanding of what the subject matter is intended to represent. Unfortunately, there is little systematic experience with the use of such inputs for developing a context.

Even for formal artifacts such as ontologies, context is a vaguely defined notion; and a context, when defined at all, is usually specified with informal documentation. As long ago as 1997, Patrick Hayes observed that “there are many ideas about what [a context’s] structure might be” (Hayes, 1997). Later, in a keynote at the CogSIMA 2012 conference, he said a bit more, “Everyone agrees that meaning depends on context, but not everyone agrees what context is.” He then continued with “... any theory of meaning will focus on some of the things that influence it, and whatever is left over gets to be called the ‘context’ ... so the ‘context’ gets to be a trash-can term. It means all the rest, whatever that is.” (Hayes, 2012) It seems that, while context is important, all one can say in general about a context is that it is at a metalevel relative to the subject matter and that the context affects the interpretation of the subject matter. However, there could be guidelines to help make explicit some aspects of a context when using a particular development methodology. Moreover, pragmatic considerations should, ultimately, drive the development of research and development of approaches to contexts. We now consider some of the motivations for why contexts should be formalized.

3 Drivers for Formalizing Contexts

We start with the premise that one role for an ontology is to take contexts that are implicit and make them more explicit. That said, it is still necessary to have a good rationale for devoting resources to develop ontologies for such a role. In this section, we consider several motivations for developing an ontology which will play the role of the context for some subject matter.

3.1 Domain Specific Drivers

Many domains have specific needs for explicit context information. For example, decision support rules for healthcare require context items such as: who (patient identity and demographics), when (age and time related medical events, admission and discharge dates), what (vital signs, behavior, provider information), why (purpose of visit and encounter), etc. In order to be able to effectively answer queries that involve such context items, it is necessary to specify them in some formal manner, ideally using an ontology.

The financial industry is another example where context information is routinely obtained, such as the provenance and details for the individual, institutional and financial information for a business loan or mortgage. Formalizing the context is important for validation of legal requirements. The lack of this capability may have been a contributing factor in the mortgage related financial recession of 2008.

3.2 Integration and Interoperation

Beyond such domain specific needs for formalizing context, there is a need for integration and interoperability, which was studied during the Ontology Summit 2016 - Framing the Conversation: Ontologies within Semantic Interoperability Ecosystems (Fritzsche et al., 2017). Indeed, integration has become a way of life for many organizations, and interoperation of systems across departments and organizations has become essential. Systems that provide or support information have been created based on the prevailing needs of a domain, organization, or application. Time constraints and limited resources prevent integration and reuse from being a priority, in spite of the recognition of the need for them. As stated by Hans Polzer, different systems embed different contexts, purposes, and scope decisions by different institutional sponsors (Polzer, 2018).

For example, a company may have multiple databases for different customer or product relationships, supplier relationships, personnel data and so on, all of which are within various settings that are usually implicit. Interoperability among the company's data resources, or the reusability of the data resources, can only be made possible if these implicit contextual matters are dealt with. In this example, the contexts are roles (such as customer, supplier, employee), relationships (such as `isCustomerOf`, `isEmployeeOf`, `isProductionManagerOf`) or products (such as bicycle, bolt, shoe). Other kinds of contextual matter may run to the whole range of the six basic questions. The use of ontology for context is a unifying conceptual model: a common language across the enterprise. In order to enable interoperability among applications or re-use of data across the enterprise, what were implicit contexts for each set of data must become explicit ontological classes and relations within the ontology.

Each system, organization, community, database or message format is thus defined in its own, too often implicit, context which might, in turn, depend on other contexts. Integration and interoperability require the sharing of information or instructions across these different, independently conceived, system contexts. The context in which these systems are conceived assumes a variety of contextual dimensions, many of which are unstated. Different unstated contexts result in different unstated assumptions in interpreting information and instructions across systems, resulting in error and risk.

While these systems are defined and built independently, systematic integration of their information and processes is essential for collaboration, shared services, information sharing and

analytics. These capabilities are not optional in today’s world; they are essential for the continued existence of commercial enterprises and the effectiveness of government. This implies the practice of integration and interoperability is one of dealing with multiple contexts, understanding their similarities, differences and relationships, and mitigating those differences and potential error and risk. Current practice depends on ontologists to understand and bridge these contextual differences, which can be successful for a particular problem (Allemang, 2018). However, integration and interoperability at scale suggests that the applicable context and their implications be more formally stated such that automated reasoning can support, validate, and in some cases replace, human intervention.

Specific contextual assumptions that may differ across systems and data sets, which may need to be made explicit, include but are not limited to:

- **Time:** Consider a database that represents history integrated with one that assumes the “current” point in time, and the time assumptions of one must be integrated with the other. Alternatively, consider an organization that has a policy for how long a measurement (e.g. of weight) may be considered the weight of an individual before the measurement expires.
- **Spatial Frame:** Consider ontologies that are built assuming a context of the surface of the earth vs. the needs of a space agency. On the earth, a location can be specified by latitude, longitude and altitude, but this does not work well in space. On earth there is a (reasonably) constant acceleration of gravity, obviously not so in space where there is considerable variation.
- **Trust:** While it is common for ontologies to consider assertions as absolutes, information differs in its trustworthiness based on both the source and the interpreter, as well as inherent uncertainties in the information. An example is a medical diagnosis at a particular time. With new patient observations and measurements, the trust in the diagnosis may change dramatically at a later time. Provenance and the relationship between provenance and trust have an impact on what information can be integrated and with what expectations.
- **Terminology:** Different communities and systems will use different terms to refer to the same concept and the same terms may refer to different concepts. Some communities are imprecise as to the meaning of terms. The term “bear,” for example, has very different meanings for finance and wildlife management, both of which are imprecise.

The above are examples of some of the many dimensions of context that may differ across systems being integrated. If ontologies are able to formally specify their contextual assumptions, then logic can be applied to cross the contextual boundaries.

Methods for representing and reasoning about context have improved, as is shown in other sections of this Communiqué. Indeed, there has been over a decade of relevant work, as discussed in (Chen, Finin, & Joshi, 2003) and (Baldauf, Dustdar, & Rosenberg, 2007). More recently, there has been work on context for the Internet of Things, which require high degrees of interoperability and answers to the six basic questions (Perera et al., 2014). Nevertheless, we have yet to see well defined and generally agreed upon best practices with a formal grounding for representing and reasoning about context.

Operationally, interoperability has perspectives or ‘dimensions’, beyond technical ones (such as social and cultural dimensions). There are several models that attempt to describe interoperability for sharing across agencies. Notably, these include the Levels of Information Systems Interoperability (LISI) model from MITRE (LISI, 1998); the Systems, Capabilities, Operations, Programs, and Enterprises (SCOPE) model from the Network-Centric Operations Industry Consortium (NCOIC) (Creps et al., 2008); and the National Information Exchange Model (NIEM) from the Office of the Director of National Intelligence (*NIEM Website*, 2017). All of these models attempt to capture aspects of the entity or entities that are needed to interoperate and their context. However, though specific technical or syntactic problems can be overcome, there are continuing issues concerning expression of, and reasoning about, context. As contexts tend to be higher order, some first-order logics may not be sufficient, further complicating reasoning tasks.

The formal representation of context, along with logical representations that enable automated reduction of time cost and risks associated with integration and interoperability, remains an open topic of research. The fundamental question is: How can the symbols and terms used in the communication among entities and different contexts of usage be made sufficiently explicit and usable for both machines and humans to ensure consistency of interpretation? An additional problem occurs when ontologies or systems are independently defined in different contexts using different terminologies. How can the contextual assumptions of the various systems be made explicit? How can processes and information defined in separate contexts be joined into a common “system of systems” that retain semantic integrity while reducing time, cost and risk? The evolution of systems and their contexts also have a profound impact on integration. Notable efforts to handle these challenges include more formal approaches to the alignment of ontologies. See for example, (Kachroudi, Diallo, & Ben, 2017; Buttigieg et al., 2016). For other approaches, see Section 4. However, there remain significant challenges.

3.3 Natural Language

The research literature on language is probably the earliest place where contexts were studied, and natural language utterances are heavily contextual in practice. Natural Language Processing (NLP) depends on context for disambiguation, so formalizing the context would be beneficial for NLP as well as many other tasks. Common-sense, for example, has been noted as a key requirement for properly deconstructing natural language text according to different contexts. An example given by (Cambria & White, 2014) concerns the difference in appraising the concept “small.” This is negative when describing rooms in a hotel review, but it is positive when describing the queues in a post office. As another example, “go read the book” is positive for a book review but negative for a movie review.

Large amounts of important data are represented using natural language, and extracting knowledge to knowledge graphs (KG) using NLP and Machine Learning (ML) is an active field. In Section 5, we will discuss an initiative for making such KGs available for public purposes. Formal context is central for the viability of such initiatives. Indeed, one would expect that linguistics in general, beyond only NLP, should provide a useful area for understanding contexts and indeed discourse analysis does help, but this is still an open issue (Bärenfänger et al., 2008).

3.4 Big Knowledge

Big Knowledge, with its heterogeneity, depth, and complexity, may be as difficult as Big Data, especially if we are leveraging heterogeneous, noisy and conflicting data to create knowledge. Single ontologies fail to scale as work expands and more contexts are encountered, thus creating the need for additional ontologies and the capability to bridge between them.

One ontology approach to Big Knowledge is to select some part of the real world suited to a particular interest or purpose, and model that as a module. Modular approaches to building ontologies fit for purpose and designed to be expandable as well as alignable are one best practice to consider (“Ontology Summit”, 2014). Differences of selection and interpretation are impossible to avoid in order to meet the intended purpose, and different external factors will generate a different context for each intelligent agent doing the interpretation. Multiple, formal contexts are likely to be the only means of dealing with the large, highly complex problems of Big Knowledge. However, as with Big Data, one must ensure that the contexts and their processing are scalable as the amounts and types of data and knowledge grow.

3.5 Knowledge Graph Building

Context could have a significant benefit for KG development. Large scale KG development must contend with noisy, malicious, missing and incomplete information. Generally, converging and redundant evidence is needed before a fact is believed and incorporated into a KG. There are many issues here, including in the extreme, how to protect a KG from being targeted by a pernicious source. Because people may differ about beliefs, there may be multiple perspectives about certain entities and their relations. One can imagine large KGs growing around such differences as we learn how to safely improve the quality of what is learned and organized. Automated cleaning of data remains tentative and manual curation of such data is often needed. Explicitly specifying the context, especially the differing beliefs, could be used as the basis for improving automated cleaning of KGs.

4 Approaches for Formalizing Context

Having made a case for the formalization of context, one can now ask how one might find and specify a context. Ontologies are built within particular perspectives which may not be shared and may not be made explicit. Thus we need a way to qualify and specify what is meant by context and its associated perspective of some subject, no matter how formal the subject might be, for example, using first order logic or mathematical category theory.

In a simple formulation, we may think of adding something about some object of attention as metadata about the object. We connect to ontologies as part of this view because a particular ontology can tell us something important about some domain which is the scope of an ontology’s coverage. So an ontology might be used to express some background knowledge about some topic, data, object etc. But, in turn, there may be a context for some subject by which we mean that we can say something about it, outside of what it says itself.

We note that different speakers during the summit expressed somewhat different views about what candidates for a relevant context might consist of. However, since context is at the metalevel,

i.e. it is something about something, then we can ask what metadata are needed to discuss the object and what level of expressiveness would be required for that discussion.

Examples include the idea that a concept, for example, a medical illness, could and should be modeled differently depending on the contextual view in which it is considered. Thus, a contextual view of illness might depend on a spatio-temporal coordinate (first world vs. third world), the thematic focus (research or treatment), a subjective perspective of agent (patient or therapist), some adopted level of granularity for the representation (cellular or organismic), and the intended application of the subject matter to be contextualized. Thus, there are many approaches one might take. This section surveys the major approaches that were presented during the summit. These approaches range from lightweight, metadata oriented methods using RDF to specify the context of RDF statements, to much richer formal mechanisms for which contexts are formal objects over which one can quantify and express first-order properties. The various approaches are not exclusive and in fact are complementary. It can be advantageous to use more than one or even all of them, where appropriate.

We begin in Section 4.1 with top-level ontologies as a means of specifying context. Provenance is a promising approach for formalizing context and is described in Section 4.2. We then, in Section 4.3, discuss relatively small, highly reusable ontologies, called microtheories. As with top-level ontologies, relating a term to a microtheory is a mechanism for specifying an aspect of the intended context for that term. In fact, since top-level ontologies are relatively small, one may regard a top-level ontology as being a kind of microtheory; however, this is not usually done because a top-level ontology has a privileged position at the most general (or “top”) of any hierarchy. After the microtheories subsection, we mention some extensions of logic that have the potential for being standard mechanisms for specifying logical contexts. Finally, there are still other proposals for representing contexts, and we briefly describe some of them in the last subsection.

4.1 Top-Level Ontologies

A top-level ontology (TLO) is one that consists of very general terms (such as “object” or “property”) that are intended to be common across all domains. The terms of a top-level ontology are very general. By choosing to specialize from one of these general terms, one is explicitly specifying one aspect of the context for that term; namely, its relative place in the knowledge graph and also that it inherits all of the properties of the parent term. For an introduction to TLOs and a list of the major TLOs, see (Bennett, 2018). In this section we consider the possible role that TLOs can have for representing context. In particular, we consider whether or not TLOs may be used to provide a set of common, organizing theories that serve to partition the various kinds of context for a subject.

Consider the following three possible approaches to the issue of the representation of context, which may or may not prove to be compatible with each other:

1. Context as Class: One of the TLO categories is the overarching category of ‘Context’ for some subject matter.
2. Kinds of Context: Several high level categories partition the kinds of context, such as one category for each of the six basic questions.

3. Everything as Context: For any element in some subject being contextualized, the contextually defined ‘meaning’ or semantics of that element is the sum of all the other elements to which it is related.

In terms of (1) ‘Context as Class’ one approach is to represent the Peircean notions of Firstness, Secondness and Thirdness (Categories, nd; Firstness, nd; Peirce Triad, nd) in the TLO. In this approach, any ‘Secondness’ category is one that brings two or more things together in some context, that context being the ‘Thirdness’ category. In the example at the beginning of Section 3.2, this approach would treat the categories of customer, supplier etc. (which are “firstness”) as existing in the contexts (“thirdness”) of customer relationship, supplier relationship and so on (“secondness”).

In terms of (2), the possible kinds of context, each TLO has its own approach for representing context. In addition to roles, relationships, products and the like, there is a range of contexts or perspectives in which the semantics of some model element, or the overall application and use of the model element or some set of terms, could be contextualized. To illustrate these, consider how some of the basic questions could be addressed in a TLO.

- Where: examples could include: geographic region, location, named place, niche, environment, biological cell, organ, etc.
- When: time, date, era, epoch, etc. Considering the commonly used TLO partitioning of ‘Continuant’ versus ‘Occurrent’ (things which exist in all their parts or in their identity across time versus things that are a feature of time), one can consider some continuant as being the context in which some occurrent is understood, and conversely the occurrent as the context in which some continuant is viewed. That is, continuant and occurrent may each be considered as viewpoints for something of the other kind. An example would be that a person’s life and times would be the context in which to understand that person.
- What: systems may be considered as the context for something, with environments being considered as kinds of system, so that any organism needs to be understood within its environment. For example, the environment ontology developed by the ENVO Consortium defines the concept of a material system as “A material entity consisting of multiple components that are causally integrated” (Buttigieg et al., 2013).
- Who: roles are considered as kinds of context, within a general TLO partitioning of Player - Role - Context. Role itself, as an TLO construct, may be divided into Relational Role, Processual Role and Social Role. In the social case this reflects Searle’s Ontology of Social Constructs, where the formulation of ‘X counts as Y in C’ is a specialization of this general pattern.

A further type of contextualization, not included in the basic questions, is the notion of granularity, either in time or space. Granularity is the way that some part of the world is conceptualized by some agent and, therefore, the way that each concept would be represented in the agent’s ontology. In other words, granularity is the way the world is ‘carved up at the edges’. This will depend on the appropriate scale or granularity at which that carving up takes place, such as atomic, molecular, cellular, animal, regional, galactic and so on. This kind of contextualization would potentially be a further feature of the Where, the When or the What.

For each of the above basic questions and their granularities, a TLO will usually have concepts that form the broad categories of which different kinds of these contextual notions may be categorized - the holes, roles, occurs and so on. Absent so far in these explorations has been the ‘Why?’, i.e., the purpose, and it is not clear if this represents a special challenge to this treatment, something that did not belong in an ontology, or something that could only be applied to the ontology as a whole. For example, why are we using this ontology, this set of representations, to do something? An example of a purpose or function as part of an explicitly contextualized activity is given in the General Ontology Evaluation Framework (GOEF) (Luciano, 2013, Slides 7 and 8).

It is possible that approaches (1) and (2) could be combined, or could be considered as the same kind of idea; that the various basic questions and their granularities, being partitions in many TLOs, might either be or (as in the case of roles, be defined within) some context. The contexts of client relationships versus specific product customer relationships point to a possible hierarchy of roles and a corresponding hierarchy of contexts. One research problem that deserves further exploration is whether this pattern forms a common organizing principle for all or most of the identified context types, perhaps with granularity or scale as a further, separate distinguishing feature.

The third approach suggested above is that everything in a given subject is the context for everything else. Given the wide-ranging nature of the basic questions, it is possible that most or all of the concepts to which any given concept relates would in any case fit under one or another of basic questions and their granularities. That would make (3) compatible with approaches (1) and (2). Again, this deserves further exploration, with realistic examples.

Given the premise that an enterprise or reference ontology can make previously implicit concepts explicit, finding the entire context can never realistically be completed. To include every possible ‘contextual’ matter as something within the ontology would be to complete an ontology of everything - the mythical and entirely useless 1:1 scale map of the world. Therefore, having hauled a bunch of contextual things into the ontology and made them part of it, that ontology itself still exists within some context and is interpreted in that context. It seems that all that has been accomplished is to make the ontology broader or narrower. However, the TLO and hence the context that it provides, is still distinguishable as an identifiable part of the overall ontology. Accordingly, what were seen as two separate notions of context - classes of things that are kinds of Context arranged under some TLO, versus the context in which the ontology itself is used - are in fact the same basic notion of context. Different uses of the ontology would take place in different contexts, and many (but not all) things that might be considered as context, might be included within some ontology in order to contextualize other elements of that ontology.

These considerations lead naturally to the “hub and spoke” approach to distributed ontology development popularized by Barry Smith (Smith, 2018, Slide 40). This approach builds a tree of ontologies where each ontology extends an existing one. In some domains, most notably biomedicine, this approach has been very successful. While the root of the tree is regarded as being ‘the’ TLO; in fact, each ontology is the TLO for its child ontologies in the tree. For example, a chemistry ontology could be used for structural chemistry or for chemical safety. However, as mentioned above, the extent to which this is an effective organizing principle for specifying context in general is a research issue.

Summarizing, there is the ontology of contexts and there is also the context of an ontology. The ontology of contexts is simply an ontology; the broadest categories of the kinds of thing that may be considered as the context for other things, are what makes up the TLO. These are the notions of Who, What, When, Where, and hoW, which, taken along with other aspects of the ontology

like the scales or granularities, hierarchies of roles, relations, descriptions of systems, processes, functions and so on, make up the context for any given concept. What you do with the ontology - the Why - remains a matter of broader perspective.

4.2 Provenance

Provenance is information about entities, activities, instruments, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness. PROV is a W3C recommendation for representing provenance. PROV also provides definitions for accessing provenance information, validating it, representing processing steps such as derivation, and the provenance of provenance. The goal of PROV is to enable the wide publication and interchange of provenance on the Web and other information systems (PROV-Overview, 2013).

PROV consists of a core ontology, called the Starting Point, a set of what are called the Expanded classes and properties, and a set of Qualified classes and properties. The core consists of three classes: `prov:Entity`, `prov:Activity` and `prov:Agent`; and nine properties such as `prov:wasGeneratedBy`, `prov:wasDerivedFrom` and so on. Given the very general nature of the core classes, PROV may be regarded as being a TLO. Indeed, given that essentially any data can be provenance data, PROV must be a TLO. Many of the Expanded classes and properties are subclasses and subproperties of the core classes and properties. The expectation is that specific uses of PROV will add additional subclasses and subproperties. The Qualified classes are reifications of the core properties to support specifying provenance of provenance (PROV-O, 2013).

PROV shows great potential as an approach to formalizing context or at least to be a good start and even a foundation for formalizing context.

4.3 Microtheories

The notion of a microtheory (MT) was introduced by Cycorp who developed a large library of MTs (Cycorp, nd). An MT is a relatively small, highly reusable ontology. MTs can distinguish the origins of the facts, and provide meta-statements about the facts, exactly as required for specifying contexts. In the Cyc Knowledge Base, the interpretation of every fact and every inference is localized to a specific region of “context space,” and all conclusions that can be inferred involve only facts that are visible from that region of context space.

An exploration of the notion of MTs as exemplified by the Cyc Project is helpful for providing a language with which to consider the question of dealing with the seemingly impossible task of specifying the entire context of a subject. Microtheories are not, as the name might suggest, some siloed representation of subject matter that is incompatible with other representations of the same or adjacent subject matter. Rather an MT takes the form of a module of the overall ontology, such that in order to reason over some specific topic (say human liver cells or credit default swaps), one needs only the broader concepts related to that subject matter and to those concepts referred to in its properties; and one does not need to make reference to ontologies of adjacent subject matter (skin cells, interest rate swaps, etc.).

An analogy that has been proposed here is that if the overall reference ontology is like the walls of a darkened room, then shining a light on parts of that ontology equates to considering the context of that part of the ontology. The light may cover a broader or a narrower area but will

not cover the whole ontology at one time. The features of the world that have been represented as explicit ontological classes may not all need to be referred to at the same time. And the room itself of course ends somewhere.

Unfortunately, to date, outside of Cycorp and its Cyc Knowledge Base which is structured as a hierarchy of MTs, there still seems to be little consistent use of MTs.

4.4 Extending First-Order Logic to Specify Contexts

The Cycorp MTs are usually specified as logical theories. First-order logic (FOL) is the fundamental technique for specifying logical theories.² Since an ontology is a logical theory, FOL is the language for rich ontologies. The standard syntax for FOL is ISO Common Logic (CL) (Delugach, 2005). While one can specify any ontology (or, more precisely, any first-order ontology), in principle, using CL, it does not have a mechanism for specifying context, other than informal comments and annotations. IKL is an extension of CL, extended with the ability to talk about the propositions that its own sentences express, and to describe its own referring names as character strings (Hayes & Menzel, 2006).

During several sessions of the summit, John Sowa presented an introduction to IKL and showed how it could be used to specify contexts (Sowa, 2018). While IKL could potentially be an effective language for expressing context information, it is not yet commonly used so its potential remains unrealized.

Another extension of logic for context is called Description Logics of Context (DLC) which extend Description Logic (DL) for context-based reasoning (Klarman & Gutiérrez-Basulto, 2016). This approach descends from J. McCarthy’s tradition of treating contexts as formal objects over which one can quantify and express first-order properties, which is similar to the approach used by the Situation Theory Ontology (STO), except that STO is not first-order. The DLC is founded on a kind of two-dimensional possible world semantics, where one dimension represents the usual object domain and the other dimension is a domain of contexts. In this approach, there are two interacting DL languages, the object and the context language, interpreted over their respective domains. The DLC differs from IKL primarily in being founded on DL rather than CL. This has the advantage of being more compatible with the Semantic Web. The DLC is relatively recent, and there have been some notable examples such as (Stephen & Hahmann, 2017), but like IKL, DLC is not yet commonly used, so its potential also remains unrealized.

4.5 Other Languages

There are other approaches than the ones described above. It is beyond the scope of this Communiqué to list every proposal; however, we mention a few of them here.

The Hierarchy of Templates (HTemp) ontology consists of about 150 templates that can be instantiated to incorporate context information.(Zarri, 2017).

RDF++ extends RDF to allow one to add metadata to each RDF fact. The metadata can include the provenance, time, information location and so on of a target fact triple in the form of metadata RDF triples (Nguyen, 2017). RDF++ is a lightweight extension of RDF, so it has excellent potential

²Note that by “fundamental” we mean that FOL has central importance, and not that FOL is all encompassing. Higher order logic may be needed for some contexts.

for being generally supported by Semantic Web and Linked Data tools. RDF++ is part of the OKN initiative and is discussed in Section 5.1.3.

Other examples include RDF quadruples, named graphs, annotated RDF, and contextualized knowledge repositories. These are still relatively new paradigms which introduce a new factor into knowledge engineering practice.

5 Open Knowledge Network

The Open Knowledge Network (OKN) initiative grew out of the observation that entity-oriented knowledge bases in the form of “graphs” now seem ubiquitous. Structured “knowledge” is used in personal assistants and consumer/search applications, but this knowledge is private and thus can be hard to extend for other uses. In response, the OKN vision is one of constructing a public knowledge space in knowledge graph (KG) form consisting of a sustainable data ecosystem. Based on this store of public knowledge, applications may be built on top of this for such domain areas as biomedicine, manufacturing, and geoscience. Broadly, the OKN process can be visualized as the building of an application using knowledge formalized in a KG in three broad steps each of which involves some context as follows:

1. Data acquisition. Use web crawlers to find relevant web pages and extract the required content from these sources (Web Crawler, nd).
2. Structure the data. Map extracted data to one or more ontologies.
3. Build the knowledge graph. As related material is found, identify the similarity and build a link to extend knowledge graph.

Thus, acquisition can start with crawling information sources like the Web to identify and extract relevant information. To aid in the reuse and understanding of the data, some information about the sources needs to be contextualized. In addition, extracted material needs to be “structured” for applications, which means that information about additional structure becomes part of the context of the refined information. The subsequent step of entity identification leverages background ontologies and a KG construction which involves selecting a minimal knowledge “tree” that connects all semantic types and, in turn, has its own construction context. Thus OKN and its resulting KG involves many contexts that need to be documented (NITRD, 2017).

In Section 5.1, we elaborate on how the OKN initiative approaches the acquisition, structuring, and linking steps described briefly above. In Section 5.2, we describe the issues and research problems that emerged in the application of these steps.

5.1 OKN Acquisition Techniques

5.1.1 Application Focus, Knowledge Breath and Depth

The summit sessions on the OKN described the basic approach used to extract information from web and public resources and leverage lightweight methods and tools, such as schema.org to create a network of open knowledge. This knowledge then becomes a resource for other applications built on top of the network. Unfortunately, this often means that not a great many contextual

issues involving deep ontologies are taken into account. Because an application is being served, the ontology used to provide a schema for the KG may be relatively simple and not based very extensively on prior ontology work. This can be expected to change as deeper problem areas such as healthcare are targeted.

5.1.2 Semantic Web and Other Sources

As part of a linked data approach, OKN knowledge may be published in lightweight form as RDF and RDF graphs. However, a richer approach is likely to be needed, such as using Cycorp microtheories. For OKN it then becomes particularly important to understand how to do this and such efforts may affect resulting ontology development, use and maintenance. An important goal of OKN going forward is to identify some of the major research problems, such as the scope, nature and precision with which context should be specified when information is extracted. While lightweight efforts may be adequate for starting the experience, in some areas a more formal semantics will be necessary to avoid semantic issues and provide better reuse.

5.1.3 Contextual Knowledge Engineering

As argued in Section 3, there is a need for a formal mechanism to specify a context. In efforts to satisfy this need, we hope to arrive at an understanding of contextualization that can be incorporated into engineering practices. To some extent, lightweight efforts to associate contextual information is already underway. RDF++, for example, is one such effort as discussed in Section 4.5 above.

All of the approaches in Section 4 extend the ability to represent individuals, concepts, properties and their relations, with the ability to document some relevant selection of contexts. In addition to these capabilities, we need to separate ontological knowledge of entities, concepts and the like between and among these contexts (Homola, Serafini, & Tamilin, 2010). In light of all of this future work, the OKN will need to refine the tools and technology to make it easier and faster to build and validate knowledge graphs and will also need new applications based on this knowledge. One important challenging goal is to contextualize knowledge bases, possibly by adding MTs quickly and with little or no human interaction (Taylor et al., 2007). However, it will still be necessary to contextualize the added MTs at the metalevel.

5.1.4 Possible Use of Microtheories

As noted in Section 4.3, there is little consistent use of MTs outside Cycorp. There is an opportunity for OKN to employ microtheories to frame a small number of very broad reasoning contexts starting with high-level, abstract knowledge, that fan out into progressively more specific contexts for application use. A small research project to test whether this idea would be valuable for providing focus to early OKN work, both to determine the amount of effort required, and to determine the benefits it would provide.

5.2 Research Issues Raised by the OKN

As noted above, the OKN initiative recognizes the importance of context, just as the Cyc Project did. We now describe some of the major issues and research problems for contexts that were raised by the OKN initiative and that are worthy of further investigation.

5.2.1 Selecting Sources with Context and Annotation

One research issue considered by the OKN and by other efforts is the problem of finding the context. There are many data sources and types to consider as sources for an OKN KG even within the Semantic Web. As a research tactic, low hanging fruit might be gained by harvesting triple stores and linked data since their structuring helps to extract facts relatively easily. To some degree these come with some contextual information such as the URL location of the “fact.”

However, producing a refined high quality KG may require investigating the origin of the different data targets and determine which extractions should be added to the knowledge graph and which ones should be discarded to avoid conflicts. In some cases, classification of the extracted information is important and some sources may be better documented with metadata to make this an easier task. However if a hierarchy of MTs is used, it may require a fair amount of context to determine where to fit the subject information into the overall hierarchy. To address such confliction problems, it will be necessary to record provenance about every node and edge in an OKN KG. Provenance at this level of detail may be difficult to manage and sophisticated approaches, such as the PROV Ontology (PROV-O, 2013), have not been routinely used in KG construction. Faithful use of PROV would capture information about the data/information entity, the agent that has modified the data entity and the activities of the modification.

Data annotated with Schema.org may also be considered a useful and opportunistic place to start investigations and tests. But limitations of the semantics of Schema.org have been noted before.

5.2.2 Depth and Formality of Representation

Another issue of interest to the semantic community concerns the question of the best representation language, as well as the depth and detail, for context (Paulheim, 2017). The evolutionary path to developing context ideally should result in a rich and detailed ontology. What is possible using current technologies is an open question since data may be defined in varying contexts or refined using various ontologies and processes with differing assumptions.

Since a major goal of OKN is to make rich knowledge openly available to a wide audience, an issue is how to organize and store the data-based knowledge for efficient access. A lightweight path might use RDF triple stores, but a richer representation may be achievable and useful.

5.2.3 Enhancing Engineering Practices

As noted in Section 5.1.3, it is important to incorporate contextualizations into engineering practices. For efforts like OKN, this should include guidance and best practices for the extraction and building of KGs as well as how to clean, refine and organize them with suitable robust and rich KBs and ontologies. OKN is building a suite of tools and technologies to make the lifecycle of KGs and their associated contexts easier and faster to build.

6 Summary of Issues and Research Questions

The Communiqué concludes by summarizing the most important issues and research questions for contexts that were raised during the summit.

6.1 Engineering Practice

In spite of the clear need for the formal, explicit specification of context, such information is largely informal and ad hoc in modern engineering practice. Aside from developing languages, paradigms and tools, it is essential to provide well tested, well founded best practice guidelines based on both theoretical considerations and industrial experience. Generally speaking, it is much easier to capture context during the development process than to attempt to discover context after the fact. An example of this is to formally specify provenance information using PROV during development. More generally, engineering practices should be extended to incorporate contextualizations as part of the development process.

6.2 Finding the Context

Another common issue mentioned in the summit is finding the context of some subject matter. More precisely, how can one determine a sufficient level of context description for operational purposes? There are many, more specific issues for this general problem. Could the subject matter development materials, such as the scope, competency questions, business requirements and “use cases” be used? If so, then how? What kind of metalevel reasoning about context could be used to determine what is relevant? Where and how do we look for what is relevant? The physical situation? Some explicit or implicit agenda? The general goals or purpose that led us to our current activity? Future developments of contexts will need to be capable of handling the dynamic nature of context as situations, events and scopes of ontologies and embedded or related data change.

6.3 Limiting the Context

The “flip” side of finding the context is ensuring that it is not too large and avoids “noise.” As noted repeatedly, there is no limit, in principle, for what and how detailed information could conceivably be as part of a context for some subject matter. How does one limit and clean the context so that it can be reasoned about effectively, yet still have sufficient coverage for operational purposes? What is the appropriate level of granularity to use? As with finding the context, it is more effective to select the level of granularity and to limit and clean the context during the development process rather than after the fact. See Sections 3.5, 4, 4.1, 4.3 and 5.2.1.

6.4 Scalability of Context

How can one ensure scalability as larger and more complex contexts are needed? The Ontology Summit 2014: Big Data and Semantic Web Meet Applied Ontology considered the problems presented by Big Data, and some of the problems and proposed solutions of that summit may be applicable for big contexts (Grüninger et al., 2014). However, the scalability issue of large contexts will have its own features, since even very small subjects could, in principle, have very large contexts. See Sections 3.2, 3.4 and 5.2.2.

6.5 Context Language

What language should be used for representing contexts? One could use existing ontology languages. This is the approach of TLOs, PROV and MTs. One could also extend existing ontology languages. This is the approach taken by IKL, DLC, RDF++ and many others. Whatever language is used, it is important to distinguish the context from its subject and to explicitly specify the relationship between them. More experimentation and experience is needed to determine the advantages and disadvantages of the various possibilities. See Sections 2, 4.1, 4.2, 4.4, 4.5, 5.1.3 and 5.2.2.

6.6 Crossing contextual boundaries

Where ontologies or systems are independently defined with differing context and terminologies, how can one harmonize diverse conceptualizations in such fields as systems engineering? The problem of interoperability and the idea of bridging ontologies was examined in the Ontology Summit 2016 (Fritzsche et al., 2017). See Sections 3.2 and 4.2.

6.7 Relevance of Linguistics

Work in linguistics in general, and NLP in particular, is well known to be relevant to the development of ontologies (Baclawski et al., 2018). The issue is how much of the work done in linguistics concerning context is relevant to formal contexts for information systems. The challenge is drawing the big lessons from these fields and integrating them in a general (rough) framework. One example of work on this problem is discussed in Section 3.3 where ML techniques are used for discovering context.

6.8 Nested Contexts

How can one distinguish what belongs in the context from what belongs in the subject matter being contextualized? A complex system can have multiple levels of abstraction and hence multiple levels of context. They may nest in a linear fashion, but more complex inter-relationships could, in principle, also occur. For example, independently developed ontologies can easily contradict each other, but one may still consider them to be contexts for some subject matter. For example, STO allows seemingly contradictory statements so long as the statements are in different situations. In one situation, a cat could be alive and in another situation the same cat is dead.³ See Section 4.1.

6.9 Purpose

One of the basic questions is ‘Why?’, and answering it provides the purpose or intention for the subject, or at least for why one is focusing on the subject. Representing a purpose in a formal manner may be useful, although there is some controversy about how useful it might be. Nevertheless, it is an interesting research problem to formally specify purpose. Whether it is truly useful will require more study and experience. See Section 4.1.

³This assumes that the context is non-probabilistic classical mechanics. In quantum mechanics, it seems that a cat could be both alive and dead in the same situation.

6.10 Higher Order Logic

Some ontologies for specifying context, such as the STO, use higher order logic. However, in practice, it is usually possible to avoid higher logic when building systems that make use of context, such as situation awareness systems. Whether it is sometimes necessary to use higher order logic for specifying context remains unsolved. See Sections 2 and 3.2.

7 Acknowledgments

Certain commercial software systems are identified in this paper. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology (NIST) or by the organizations of the authors or the endorsers of this Communiqué; nor does it imply that the products identified are necessarily the best available for the purpose. Further, any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NIST or any other supporting U.S. government or corporate organizations.

We wish to acknowledge the support of the ontology community, and the support of The Networking and Information Technology Research and Development (NITRD) Program who generously contributed the use of their facilities for the Ontology Summit Symposium.

We also wish to thank the session co-champions, invited speakers and participants who contributed to the Ontology Summit. In a departure from recent summits, the Ontology Summit 2018 encompassed sessions for a much longer period of time, starting in September 2017 and continuing until April 2018. There were 23 invited speakers, some of whom gave presentations at more than one session. The session chairs and web links to the session pages and any slide presentations are listed in Table 1 after the References Section below. For simplicity, the web links have been abbreviated: one should prepend “<http://bit.ly/>” before each of the seven-character codes in the table. The sessions were recorded, and the recording are available on the session web pages. The invited speakers and web links to their slide presentations or video recordings are listed in Table 2. The session chairs and invited speakers covered far more material than it was possible even to summarize in this Communiqué, so the reader is encouraged to review the slides and listen to the presentations.

References

- Allemang, D. (2018). *FIBO and Cesium*. Retrieved on May 15, 2018 from <http://bit.ly/2Kkfm4m>
- Baclawski, K., Bennett, M., Berg-Cross, G., Fritzsche, D., Schneider, T., Sharma, R., ... Westertinen, A. (2018). Ontology Summit 2017 Communiqué: AI, Learning, Reasoning and Ontologies. *Applied Ontology*, 13(1), 3–18.
- Baldauf, M., Dustdar, S., & Rosenberg, F. (2007). A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(4), 263-277.
- Bärenfänger, M., et al. (2008). OWL ontologies as a resource for discourse parsing. *LDV Forum*, 23(1).

- Bennett, M. (2018). *Upper ontologies for specifying context*. Retrieved on May 15, 2018 from <http://bit.ly/2FDL5te>
- Brinkley, J., Suciu, D., Detwiler, L., Gennari, J., & Rosse, C. (2006). A framework for using reference ontologies as a foundation for the semantic web. *AMIA Annu Symp Proc*, 96-100. Retrieved on June 3, 2018 from <http://bit.ly/2JnXRwV>
- Buttigieg, P., et al. (2013). The environment ontology: contextualising biological and biomedical entities. In *Journal of Biomedical Semantics* (Vol. 4). <https://doi.org/10.1186/2041-1480-4-43>
- Buttigieg, P., et al. (2016). The environment ontology in 2016: bridging domains with increased scope, semantic density, and interoperability. *Journal of Biomedical Semantics*, 7(1), 57.
- Cambria, E., & White, B. (2014). Jumping NLP curves: A review of natural language processing research. *IEEE Computational Intelligence Magazine*, 9(2), 48–57.
- Casanave, C. (2018). *Context aware ontologies*. Retrieved on May 15, 2018 from <http://bit.ly/2Epjyrl>
- Categories* (Peirce). (nd). Wikipedia. Retrieved on May 23, 2018 from <http://bit.ly/1x81Q5A>
- Chen, H., Finin, T., & Joshi, A. (2003). An ontology for context-aware pervasive computing environments. *The Knowledge Engineering Review*, 18(3), 197–207.
- Creps, R., Polzer, H., Yanosy, J., Sutton, F., Jones, A., Ford, R., ... Schneider, T. (2008). *Systems, capabilities, operations, programs, and enterprises (scope) model for interoperability assessment version 1.0*. Retrieved on April 27, 2018 from <http://bit.ly/2JznDK0>
- Cycorp Website*. (nd). Retrieved on May 23, 2018 from <http://www.cyc.com>
- Delugach, H. (Ed.). (2005). *Common Logic - A framework for a family of logic-based languages*. ISO/IEC JTC 1/SC 32N1377, International Standards Organization Final Committee Draft
- Firstness*. (nd). Merriam-Webster. Retrieved on May 23, 2018 from <http://bit.ly/2KPgEOI>
- Fritzsche, D., Grüninger, M., Baclawski, K., Bennett, M., Berg-Cross, G., Obrst, L., ... Westertinen, A. (2017). Ontology Summit 2016 Communiqué: Framing the Conversation: Ontologies within Semantic Interoperability Ecosystems. *Applied Ontology*, 12(2), 91–111.
- Groth, P., & Moreau, L. (Eds.). (2013). *An overview of the PROV family of documents*. Retrieved on May 17, 2018 from <http://bit.ly/2L46MSd>
- Grüninger, M., Obrst, L., Baclawski, K., Bennett, M., Brickley, D., Berg-Cross, G., ... Yim, P. (2014). Ontology Summit 2014 Communiqué: The Semantic Web and Big Data meet Applied Ontology. *Applied Ontology*, 9(2), 155-170.
- Hayes, P. (1997). Contexts in context. In *AAAI Fall Symposium*. Retrieved on May 15, 2018 from <http://bit.ly/2jf0BBa>
- Hayes, P. (2012). Situations, contexts, states of affairs, and the limits of formalization. In *IEEE Conference on Cognitive and Computational Aspects of Situation Management*. Retrieved on May 15, 2018 from <http://bit.ly/2EKdVTq>
- Hayes, P., & Menzel, C. (2006). *IKL specification document* (Tech. Rep.). IKRIS Interoperability Group. Retrieved on May 15, 2018 from <http://bit.ly/2rEif2D>
- Homola, M., Serafini, L., & Taminin, A. (2010). Modeling contextualized knowledge. In *Proc. 2nd Workshop on Context, Information and Ontologies (CIAO 2010)* (Vol. 626).
- Kachroudi, M., Diallo, G., & Ben, S. Y. (2017). On the composition of large biomedical ontologies alignment. In *Proceedings of the 7th International Conference on Web Intelligence, Mining and Semantics*. ACM.

- Klarman, S., & Gutiérrez-Basulto, V. (2016). Description logics of context. *Journal of Logic and Computation*, 26(3), 817–854.
- Kokar, M., Matheus, C., & Baclawski, K. (2009). Ontology-based situation awareness. *Information Fusion*, 10(1), 83–98.
- Lebo, T., Sahoo, S., & McGuinness, D. (Eds.). (2013). *PROV Ontology (PROV-O)*. Retrieved on May 15, 2018 from <http://bit.ly/2xPcx2k>
- LISI. (1998). *Levels of Information Systems Interoperability*. Retrieved on April 27, 2018 from <http://bit.ly/2HUtM7f>
- Luciano, J. (2013). *The General Ontology Evaluation Framework (GOEF): A proposed infrastructure for the ontology development lifecycle*. Retrieved on June 3, 2018 from <http://bit.ly/2JrUXar>
- Malone, J., & Parkinson, H. (2010). *Reference and application ontologies*. Cambridge, CB10 1SD, UK.
- Nguyen, V. (2017). Semantic web foundations for representing, reasoning, and traversing contextualized knowledge graphs.. Ph.D. Thesis
- NIEM website. (2017). National Information Exchange Model of the Office of the Director of National Intelligence. Retrieved on June 2, 2018 from <http://bit.ly/2Lcm1re>
- NITRD. (2017). *Big Data Interagency Working Group: 3rd Workshop on an Open Knowledge Network*. Retrieved on May 15, 2018 from <http://bit.ly/2rDrTm8>
- Ontology Summit 2014: Big Data and Semantic Web Meet Applied Ontology*. (2014). Retrieved December 9, 2017 from <http://bit.ly/2xPFI1D>
- Paulheim, H. (2017). Knowledge graph refinement: A survey of approaches and evaluation methods. In *Semantic web* (Vol. 8.3, pp. 489–508).
- Peirce's Triad*. (nd). Stanford Encyclopedia of Philosophy. Retrieved on May 23, 2018 from <https://stanford.io/2KPhoU0>
- Perera, C., et al. (2014). Context aware computing for the internet of things: A survey. *IEEE Communications Surveys and Tutorials*, 16(1), 414–454.
- Polzer, H. (2018). *Data/entity representation in different contexts and frames of reference*. Retrieved on May 15, 2018 from <http://bit.ly/2KgxDJC>
- Smith, B. (2018). *Response to john sowa at the ontology summit 2018*. Retrieved May 23, 2018 from <http://bit.ly/2row7xT>
- Sowa, J. (2017). *Contexts in language and logic*. Retrieved on April 27, 2018 from <http://bit.ly/2JAYapt>
- Sowa, J. (2018). *Representing and reasoning about contexts using IKL*. Retrieved on May 15, 2018 from <http://bit.ly/2xrTkIa> and <http://bit.ly/2y4h41v>
- Stephen, S., & Hahmann, T. (2017). An ontological framework for characterizing hydrological flow processes. In *LIPICs-Leibniz International Proceedings in Informatics* (Vol. 86). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Taylor, M., et al. (2007). Autonomous classification of knowledge into an ontology. In *FLAIRS Conference*.
- Web Crawler*. (nd). Retrieved on May 23, 2018 from <http://bit.ly/R7Ys74>
- Zarri, G.-P. (2017). Functional and semantic roles in a high-level knowledge representation language. In *Artificial intelligence review*. doi: 10.1007/s10462-017-9571-5

Table 1: Session Chairs of the Ontology Summit 2018

Date	Session	Chair(s)	Slides/Videos
6 September	2vqdcFy	Kenneth Baclawski	2rHDNeU
13 September	2wVi9aM	Kenneth Baclawski	2hcGbnA
20 September	2xrTkIa	John Sowa	2wEBtKz
27 September	2y4h41v	John Sowa	2JAYapt
4 October	2f0Nsxx	David Whitten	2wDGTW1
11 October	2fU17Ge	Jack Ring	2yaEscJ
18 October	2fUBkuT	Mike Bennett	2hNCxRv
25 October	2fUgWkm	John Sowa	2xnHLdy
1 November	2fUKUOH	Kenneth Baclawski	2j56HEb
08 November	2fU1e4C	Kenneth Baclawski	2jbkusZ
15 November	2fU1gcK	Matthew West	2jzmAf7
17 January	2CuIyfU	John Sowa	2DhqYv1
24 January	2CxSe9C	David Whitten and Ravi Sharma	2I19hBJ
31 January	2CwTidV	David Whitten and Ravi Sharma	2DPEeYC
7 February	2CxSghK	Mike Bennett and David Whitten	2FDL5te
14 February	2CwS961	Ram D. Sriram and Gary Berg-Cross	2DG72TZ
21 February	2CumqCf	Cory Casanave and Ravi Sharma	2K1QX81
28 February	2CuAmfA	Kenneth Baclawski	2F94UIk
7 March	2CwS4iJ	Janet Singer and Jack Ring	2DiopIN
14 March	2CwS9mx	Mike Bennett and David Whitten	2FC3qqr
21 March	2CvAuvq	Cory Casanave and Ravi Sharma	2G0EZPo
28 March	2CumqSL	Ram D. Sriram and Gary Berg-Cross	2rFcUbh
4 April	2Cwc440	Kenneth Baclawski	2J1XkbJ
11 April	2CvABqF	Kenneth Baclawski	2JGsPh9
18 April	2CvmmCu	Kenneth Baclawski	2HdYZPf
25 April	2Hf7EAT	Kenneth Baclawski	2KhGo6G
30 April/1 May	2Fy60JR	Kenneth Baclawski	

Table 2: Invited Speakers at Ontology Summit 2018 Sessions

Date	Invited Speaker(s)	Slides/Videos
17 January	David Whitten Ravi Sharma	2DtIpJu
17 January	Mike Bennett David Whitten	2DfNTYb
17 January	Ram D. Sriram Gary Berg-Cross	2DG72TZ
17 January	Cory Casanave Ravi Sharma	2DF30v7
17 January	Janet Singer and Jack Ring	2DaFpRY
24 January	Spencer Breiner	2E4uTNi
31 January	Dov Dori	2DQsR2H
31 January	Cory Casanave	2Epjyr1
7 February	Barry Smith	2FBh501
7 February	Frank Loebe	2FDkGf2
14 February	Vicki Tardif Holland	2GaQYJY
14 February	Ramanathan Guha	2HfD8Yb
14 February	Mayank Kejriwal	2G7YsgT
21 February	Eswaran Subrahmanian and Ira Monarch	2K1QX81
7 March	Janet Singer	2rH4Maq
7 March	Jack Ring	2rGPouJ
7 March	Hillary Sillitto	2rGtR1S
14 March	Vipul Kashyap	2FDL88o
21 March	Hans Polzer	2KgxDJC
21 March	Dean Allemang	2KkfM4m
28 March	Vinh Nguyen	2rGbihx
28 March	Amit Sheth	2rGihHt
28 March	Charles Klein	2Grcfj2
30 April	Bryan A. Biegel	2rHtXti
30 April	Gary Berg-Cross	2rF30Xo
30 April	Vinh Nguyen	2r1CkL2
30 April	Ram D. Sriram	2rmI9ba
1 May	Richard Conroy	2rmI0Va
1 May	Barry Smith	2row7xT
1 May	John Sowa	2rmxvkv
1 May	Sowa-Smith Debate	2JSA19e