# Ontology Summit 2017 Communiqué – AI, Learning, Reasoning and Ontologies

Kenneth Baclawski* Mike Bennett† Gary Berg-Cross‡
Donna Fritzsche§ Todd Schneider¶ Ravi Sharma‖
Ram D. Sriram** Andrea Westerinen††

**Abstract**

There are many connections among artificial intelligence, learning, reasoning and ontologies. The Ontology Summit 2017 explored, identified and articulated the relationships among these areas. As part of the general advocacy of the Ontolog Forum to bring ontology science and engineering into the mainstream, we endeavored to abstract a conversational toolkit from the Ontology Summit sessions that may facilitate discussion and knowledge sharing amongst stakeholders concerned with the topic. Our findings are supported with examples from the various domains of interest. The results were captured in the form of this Communiqué, with expanded supporting material provided on the web.

Keywords: artificial intelligence; machine learning; reasoning; ontologies

## 1 Introduction

We are currently witnessing increasingly widespread applications of Artificial Intelligence (AI), which deals with intelligent behavior, learning and adaptation in computational systems. Three of the most significant drivers and enablers of AI technology are the availability of increasingly massive amounts of data (Big Data); the rapidly dropping cost of storing and processing data; and advances in machine learning (ML) techniques (Wactlar, 2017). This situation has made it possible to exploit sophisticated ML techniques that require large amounts of data to be effective. The applications of ML have "boosted Android's speech recognition, and given Skype Star Trek-like

*Northeastern University, Boston, MA USA
†Hypercube Limited, London, UK
‡RDA US Advisory Group, Troy, NY USA
§Hummingbird Design, Chicago, IL USA
¶Engineering Semantics, Fairfax, VA USA
‖Senior Enterprise Architect, Elk Grove, CA USA
**National Institute of Standards & Technology, Gaithersburg, MD USA
††Nine Points Solutions, Potomac, MD USA

instant translation capabilities. Google is building self-driving cars, and computer systems that can teach themselves to identify cat videos. Robot dogs can now walk very much like their living counterparts" (McMillan, 2015, paragraph 4). It is worth noting that while these may be very useful, they are types of what are called "narrow AI" technologies. These are AI applications that allow computers to solve specific problems, like image recognition, or perform reasoning tasks that do not emulate the full range of human, self-directed, cognitive abilities, which is referred to as "general AI".

The Ontology Summit 2017 was an attempt to survey the ways in which the AI techniques of ML, reasoning and ontology engineering are being used for their mutual benefit. These uses were classified into three tracks, but it was soon clear that the different tracks had significant overlap with each other, and there was considerable variety within each track. It was also noted that the terms "reasoning" and "learning" have many interpretations. In the Ontology Summit, we chose to restrict "learning" to machine learning because the context of the Summit's theme was AI. However, human learning was also explored in the Summit to some extent, especially as it relates to machine learning. The term "reasoning" was not intended to be restricted to formal, logical reasoning. In general, the Summit found that classifying the many techniques, determining the best practices, and identifying synergies among technologies for ML, reasoning and ontologies have emerged as three key challenges for the exploitation of the relationships among ML, reasoning and ontology engineering.

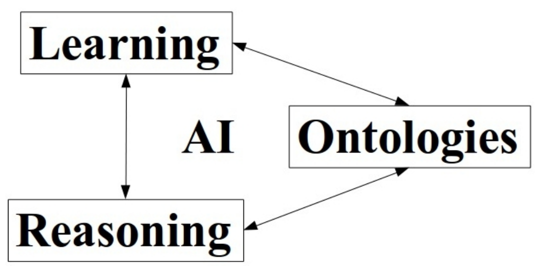


Figure 1: Ontology Summit 2017 Logo

The Ontology Summit 2017 surveyed the current state of the art among the major AI topics of learning, reasoning and ontologies with three tracks, summarized in Section 2 below. Each track focused on one of the relationships between two of the three AI topics, as illustrated in the Ontology Summit 2017 logo in Figure 1. The relationship between learning and reasoning was addressed indirectly via ontologies. Some of the background for the theme of the Ontology Summit is presented in Section 3. This is followed, in Section 4, by a survey of some of the opportunities and challenges of the relationships among learning, reasoning and ontologies.

The final publication is available at IOS Press through http://dx.doi.org/10.3233/AO-170191

## 2 Track Summaries

### 2.1 Overview Session

The Overview Session began the Ontology Summit with a presentation introducing many of the themes that were expanded in the track presentations and discussions that followed. In particular, the overview presented the Ontology Learning Layer Cake in Figure 2, which was used as a common touchpoint across all of the tracks. The highest layer represents logic and axioms, with lower layers depicting schemata, relations, concept hierarchies, and synonyms. Terms are at the lowest layer. The layer cake presents a framework for describing the process of knowledge extraction. An example is the process of acquiring a concept hierarchy which can be depicted by graphs representing relationships among elements (Buitelaar, 2017; Getoor, 2017).

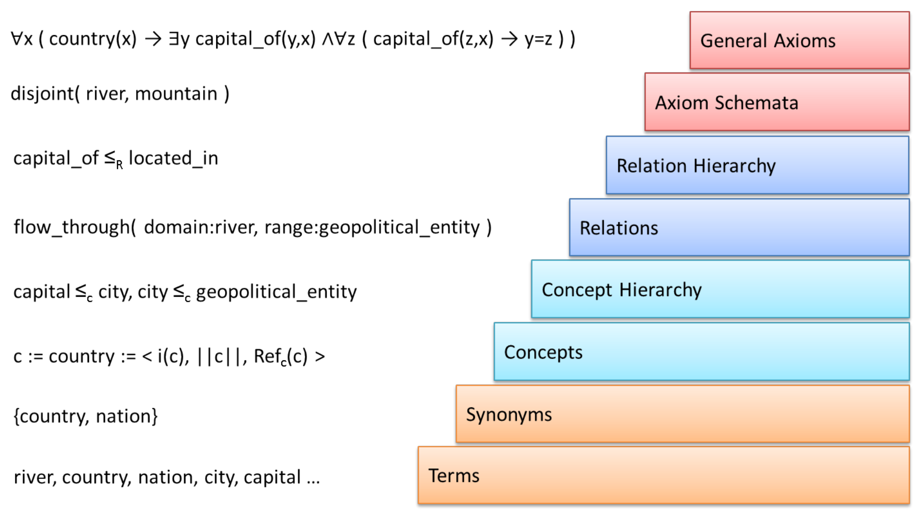


Figure 2: Ontology Learning Layer Cake (Buitelaar, 2017)

### 2.2 Track A: Using ML to extract knowledge and improve ontologies

This track addressed one of the big bottlenecks for AI; namely, how to create sufficient knowledge about the world for a truly intelligent agent. For the most part, knowledge bases and ontologies for such purposes have been largely handcrafted, which is both time and resource consuming. This track explored the use of automation and various ML approaches to extract knowledge and improve ontologies including populating them (Berg-Cross, 2017). The following are some of the highlights of the presentations in this track.

- Bottlenecks due to the complexity of ontology engineering remain a challenge but a variety of ML approaches and tools, including statistical, linguistic and bio-inspired, have been built and can be utilized for a variety of purposes. One can extract information and structured knowledge from a variety of sources to facilitate the development and maintenance of ontologies (Buitelaar, 2017; Corcoglioniti, 2017; Hruschka, 2017). One can filter noisy data to further improve the quality of developed ontologies (Aasman, 2017; Pafilis, 2017). One can translate linguistic realizations of ontology entities from one language to another (Buitelaar, 2017).

- The ontology learning layer cake provides a conceptual framework for discussion of what types of knowledge are being built. Automated knowledge development (such as taxonomy development) should be gauged against manually constructed knowledge (such as hierarchies) (Buitelaar, Cimiano, & Magnini, 2005).

- There has been significant recent progress with supervised learning. Some approaches to machine learning rely on the accessibility of large amounts of labeled training data, the creation or curation of which can be resource-intensive, and time-consuming (*Machine learning*, 2017). Partly in response to these resource challenges, there are also complementary, less resource intensive prospects from semi- and unsupervised learning approaches that offer the possibility of progress without the bottleneck of needing a large number of training examples. Proper seeding of these, however, may be needed, as was the case with the Never-Ending Language Learning (NELL), along with cumulative use of knowledge to support learning to learn (Hruschka, 2017).

- In some domains, such as the biomedical, there are sufficient ingredients for discovering new knowledge by leveraging ML, quality ontologies, and a wealth of domain data about genes and cell functions (Yu, 2017).

## 2.3 Track B: Using background knowledge to improve machine learning results

The mission of this track was to scope out challenges and opportunities when using background knowledge to improve machine learning results, the role of ontologies and comparable resources in achieving this, and the requirements for ontologies that may be used in these ways. Five speakers provided examples and insights, and there was lively discussion of the issues and ideas they raised. The sessions provided insights on the semantic, syntactic and contextual aspects of machine learning using ontologies. Relations with the themes of other tracks were discussed, along with common ideas and opportunities (Bennett & Westerinen, 2017). The following are some of the highlights of the presentations and discussions in this track.

- Background knowledge is useful for making ML results understandable, as well as providing the human qualities of sentiment and intuition (Davidson, 2017; Presutti, 2017).

- There is a bewildering array of model choices and combinations (Bennett & Westerinen, 2017).

- Background knowledge could improve the quality of ML results by using reasoning techniques to select learning models and prepare the training and test data (reducing large, noisy data sets to manageable, focused ones) (Davidson, 2017; Erekhinskaya, 2017; Falk, 2017).
- The ontologies used to enhance NLP results need not be the same ontologies that are synthesized by NLP tools. Indeed, they may not even be the same kinds of ontology (Bennett & Westerinen, 2017). The challenge is to ensure that the ontologies are compatible, so that one can iteratively improve the same ontology using these two activities.
- Context is important for disambiguating terms such as "bear" and "cheese". For example, "bear" can be a noun (an animal), a verb (such as to bear arms) or an adjective (a bear market). Similarly, "cheese" can refer to a food or is slang for a drug (Bell & Kendall, 2017).
- Combining ontology engineering with ML can improve decision support, including improving the quality of decisions, making the reasons for a decision more understandable, and adapting the decision making process to changing conditions (Baclawski, 2017).
- The Financial Industry Business Ontology$^{TM}$ (FIBO) and corporate taxonomies can help extract and integrate information from data warehouses, operational stores and natural language communications (Bell & Kendall, 2017).
- It is important not only to be able to extract knowledge graphs from multilingual text but also for the ontology itself to be multilingual (Buitelaar, 2017).

## 2.4 Track C: Using ontologies for logical reasoning and vice versa

The goal of track C was to discuss the techniques developed for reasoning using ontological foundations. Any intelligent agent has four basic components:

1. sensors which receive external signals in various forms;
2. knowledge, which manifests in various forms (e.g., qualitative, quantitative and combinations of both);
3. inference mechanisms which reason about the world, given the sensor input, using knowledge; and
4. actuators, which execute various forms of action (e.g., physical and mental) (Fritzsche & Sriram, 2017).

However, human intelligence also requires:

1. feedback mechanisms, including the ability to learn new behaviors (Baclawski, 2017; Gil, 2017);
2. emotion and sentiment analysis (Davidson, 2017; Presutti, 2017);

3. social behaviors (Davidson, 2017).

The following are some of the highlights of the presentations and discussions in this track.

- Ontologies form the core for knowledge representation, which is used by the appropriate inference strategy (Kuksa, 2017).
- Many forms of inference strategies exist (e.g., backward and forward reasoning, inexact reasoning, constraint satisfaction, theorem proving) (Kuksa, 2017).
- By progressively adding axioms (premise selection) one can improve the speed of inference (Kuksa, 2017).
- Other forms of reasoning are also important inference mechanisms. Probabilistic is one such form of reasoning (Aasman, 2017; Getoor, 2017). Analogical reasoning is another (Rugaber, 2017).
- Ontology interoperability ranges from lowest expressivity (e.g., taxonomy) at the syntactic level to intermediate (e.g., thesaurus) at the structural level with various levels of expressivity (e.g., conceptual models and logical theory), and to the highest semantic level with First Order Logic. See Figure 2.
- Tools and technologies also follow this progression (Kuksa, 2017).
- Reasoning can help disambiguate terms in overlapping domains (Hitzler, 2017).
- Ontologies can aid in the discovery of scientific knowledge and can help automate the discovery workflow. Workflow analysis can help understand the results of the scientific discovery process (Gil, 2017).
- Design engineering examples applicable for ontology inference relate structural strategies, requirements and use cases, including ecosystem requirements (Rugaber, 2017).
- Ontologies can be used for finding analogies in the biological domain to problems in the engineering domain (Rugaber, 2017).

# 3 Background

Early success with machine translation, as well as machine "learning" using statistical methods, suggested that some progress could be made sub-symbolically, i.e., without specific representations of knowledge (Koehn, 2010). Recent AI techniques are dominated by sub-symbolic ML. However, sub-symbolic ML generally works by solving classification or regression problems on uninterpreted raw data. Systems devised to solve these problems can be said to "learn" in the sense of optimizing a set of model parameters to increase performance over time. Calling it "learning" makes it sound cognitive and mind-like, but computationally it generally has no resemblance to how humans think, learn and understand or how ontologies represent knowledge.

The most common recent AI techniques use biologically inspired neural net architectures along with optimizing and statistical approaches. In addition to being subsymbolic, these techniques proceed "bottom up" from data such as text and images. Interacting with text and images like this is very different from the much broader, biologically inspired, experience of interacting with the world. While there are some smart systems, such as self-driving cars, that have a limited range of such interactions, learning on the job here seems risky (Baclawski, Gross, et al., 2017).

Current AI progress has yet to master the broader forms of learning and understanding that comes from real-world, embodied experience. Some think that such embodied learning requires starting with a cognitive core and then successively developing more sophisticated cognitive models (Hruschka, 2017). The social aspect of real-world, embodied experience includes learning common knowledge from other intelligent agents, along with their information bearing products, such as text, data and physical actions. While acquisition of domain knowledge and domain reasoning methods continue to improve, it has proven very hard to "code" into machines or to learn bottom up without some seed knowledge. Some automated help is needed to handle the major bottleneck issues of domain and general knowledge, which can help with common sense and reasoning, such as the many daily inferences that humans make (Wactlar, 2017).

ML techniques are primarily statistical, so incorporating uncertainty into ontologies would be useful when ontologies are used with ML. The following are two of the techniques for integrating probability with semantics:

- Statistical Relational Learning. General AI needs to deal with both relational structure and uncertainty. A particular line of work focusing on the combination of probabilistic models with description logic is known as Probabilistic Semantics (Pileggi, 2016).
- Probabilistic Soft Logic (PSL) is a machine learning framework for developing probabilistic models. PSL uses first order logic rules as a template language for graphical models over random variables with soft truth values from the interval [0, 1] (Getoor, 2017).

In addition to uncertainty reasoning, there are many fields that are necessary for achieving AI that is general and robust. Figure 3 illustrates these components and their interactions with one another.

# 4 Opportunities and Challenges

A great many techniques have been developed that make use of learning, reasoning and ontologies. This section is an attempt to survey some of the problems and opportunities of the relationships among learning, reasoning and ontologies.

## 4.1 Predictions using Biological Organization

The cell, along with body systems, is usually modeled using levels of organization, each of which can be represented using an ontology – from the molecular and biochemical level, to the cellular and tissue level, to the organ and organ system level, and

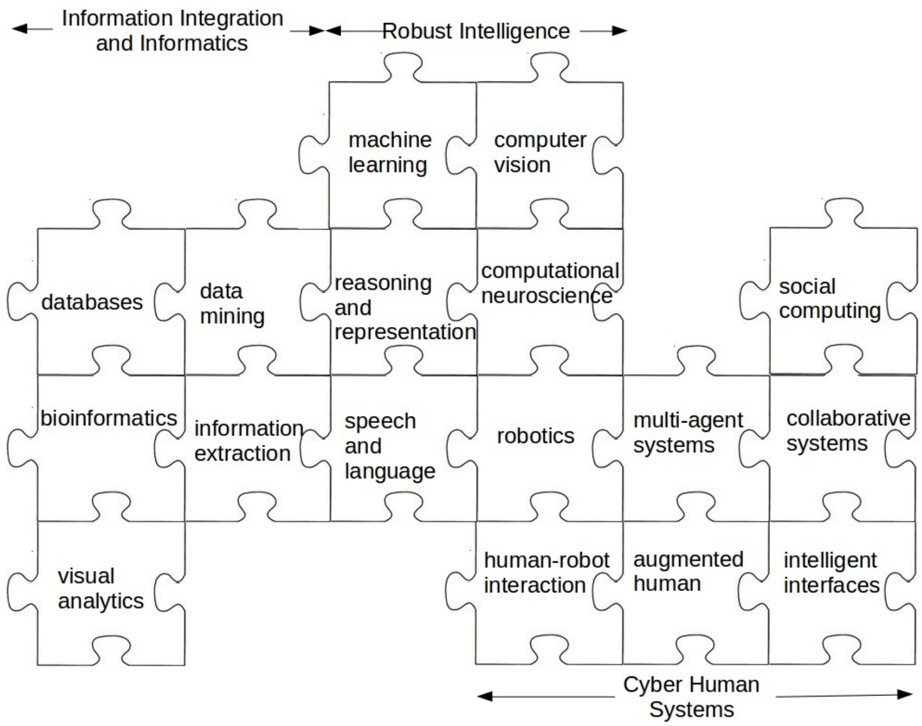


Figure 3: Robust Intelligence Components and Interactions (Wactlar, 2017)

to the level of biospheres. Biological organization is often called a "hierarchy", but it is not a hierarchy in the ontological sense. Functionalizing the levels of organization, i.e., mapping from one level to another, is a form of reasoning, called "biological inference". An example is the genotype-phenotype map (*Genotype-phenotype*, 2017). ML guided by ontologies, such as the manually curated Gene Ontology (GO), is the underlying technique for biological inference from genes to their protein products (*Gene Ontology*, 2008). This use of ML employs extensive measurement data on the network relations between genes and their protein products. GO is structured to specify these network relations by means of the three GO domains: Biological Process, Cellular Component, and Molecular Function (Yu, 2017).

Biological inference leverages knowledge of gene co-expressions and protein-protein interactions to create a data-derived gene similarity network. An alignment process identifies which data terms are new and which recapitulate existing knowledge in GO. Taken together, this knowledge can be reduced to an ontological hierarchy and aligned with the GO, suggesting names for new, data-driven terms. A majority, about 60%, of relations found by data derivation for cellular components are already in the GO-Cellular Component, but only about 25% of the derived terms for Biological Process and Molecular Function were already found in GO, indicating that much potentially useful knowledge can be uncovered using such techniques (Hahmann, Stephen,

& Broderic, 2016; Yu, 2017).

Not every domain has the degree of agreement on base ontologies that has been achieved in the BioMedical domain. Likewise, not every area has the degree of agreement on levels of organization. Therefore, it is an open research issue whether something like this be achieved in other domains.

## 4.2 Context Identification

Words like "stock" and "bear" have many different interpretations that depend strongly on the context (Bell & Kendall, 2017). Thus in the context of a financial discussion that touches on stocks, the term "bear" is likely referring to investor or stock market attitude as opposed to a type of mammal. Context is fundamental to interpretation, yet it is difficult to formalize the notion of context. Perhaps the best stance is to agree with Pat Hayes (Hayes, 1997) that there is no single notion of context. At its core, a context determines whether a proposition can be said to be true or false. For example, saying "There are lots of bears in the stock market" is reasonable, while a proposition stating that "There are lots of polar bears in the stock market" is likely false, due to the context. What we understand as the context can depend on context itself. Important work in this area was done by John McCarthy and Pat Hayes in their situation calculus (Hayes, 1997). Another important technique is situation theory (Devlin, 1991). Situation theory has been formalized as an ontology expressed in the Web Ontology Language (OWL) (Baclawski, Malczewski, Kokar, Letkowski, & Matheus, 2002; Kokar, Matheus, & Baclawski, 2009). Situation theory is very popular in many domains, especially military and business domains.

While situation theory is an effective formalization for context in many cases, it is not a complete solution to the notion of context. The challenge is to develop an effective formal notion of context that can be used to disambiguate the interpretation of words in human discourse and potentially lead to what is called "understanding". Both situation calculus and situation theory support reasoning processes. Situation theory is especially versatile in this respect, allowing many forms of reasoning (Baclawski, Chan, et al., 2017).

## 4.3 Cognitive Scaffolding

Machine learning has come a long way since Arthur Samuel's 1959 definition of ML as a sub-field of computer science that gives computers the ability to learn without being explicitly programmed (Samuel, 1959). ML, in this sense, has become much more feasible now that more and richer data has become available. However, as noted in Section 3 above, we are still far from achieving ML without any explicit programming. Modern machine learning workflows often include routine tasks for: problem evaluation, data exploration, data pre-processing, and model training, followed by testing and deployment, all of which are supervised by humans. Nevertheless, some applications of ML have become more cognitive, contextual and holisitic rather than being purely bottom up. Achieving these features requires more intelligent processing and more knowledge.

For example, some knowledge is needed to handle ambiguity for words such as "pen", which have many senses. One sense of "pen" is of a writing instrument, but another is of a small enclosure for holding animals or children, depending on context. But the context here isn't really a statistical one. Understanding a sentence with "pen" in it often requires real world knowledge about the relative sizes of boxes and pens. One way to provide such knowledge is to begin with a knowledge "starter" or "seed" which can be expanded by applying AI processes, such as ML. An example of seeding knowledge for intelligent process was illustrated in the NELL system for one important cognitive task: reading (Hruschka, 2017). The inputs to NELL include:

1. an initial ontology defining hundreds of categories (e.g., person, sports team, fruit, emotion) and relations (e.g., plays on team (athlete, sports team), plays instrument (musician, instrument)) that NELL was expected to read about, and

2. 10 to 15 seed instance examples of each category and relation.

Seeding knowledge in this case was helped by the topical focus of what NELL's reading task was, for example, reading about sports or music.

A more general question is the ontological basis of sufficient knowledge needed by an autonomous, intelligent agent which observes and acts in and on an environment in a directed way to achieve goals. This remains an abiding question. No single architecture, technique or tool is currently available for developing an intelligent agent even for relatively simple information agents such as envisioned in the original DAML effort ("DARPA Agent Markup Language", 2006). Candidate approaches, however, abound from disciplines as diverse as Cognitive Development, Cognitive Science, Developmental Robotics and AI. One example is the Soar cognitive architecture for general intelligent agents (Laird, Newell, & Rosenbloom, 1987; Newell, 1990; Wray & Laird, 2003). Another example is the ACT-R cognitive architecture for simulating and understanding human cognition (Anderson & Lebiere, 1998). A third example is the belief-desire-intention software model for programming intelligent agents (Georgeff & Lansky, 1987; Huber, 1999; Rao & Georgeff, 1991).

But such agent systems have difficulty accommodating common sense things like diverse spatio-temporal information, including quantitative and qualitative assessments within a single analytic context in a suitable period of time. Yet, as part of the analytic process for understanding a situation, humans easily integrate both quantitative and qualitative information assessments to arrive at conclusions, and this happens before humans, for example, learn to read. That is, the seeding for something like NELL appears to be part of human development. How is this? It seems reasonable to assume that some degree of innate structure is needed to develop a cognitive system and relevant knowledge for some common things as part of an agent's experience. Such cognitive development, particularly in the context of general intelligence, is sometimes discussed in terms of an early scaffolding with a core set of cognitive abilities providing a temporary structure to afford organizing more general knowledge and learning during the progressive development into a richer cognitive skill system.

In Cognitive Science, knowledge is conceived as the main outcome of the process of understanding: by interacting with the environment, intelligent agents are able to interpret and represent world facts, suitably acting to preserve themselves and pursue specific goals accordingly (Albertazzi, 2000; Neisser, 1987). Representing knowledge is a necessary step for communication, but knowledge can be properly represented only insofar as world phenomena are previously presented to humans, namely experienced through cognitive structures. Such a cognitive scaffolding could be understood as a starter set: a type of dynamic building block.

Unfortunately, there are currently no accepted starter sets within most domains nor a general theory as to what a starter set is or what the first recognizable knowledge and reasoning components are. The Cognitive Linguistics Hypothesis, for example, suggests that likely, common human experiences with the world are simple, limited and constrained (Johnson, 1990). Given this, a core part of understanding is grounded in perception and action. This core semantics is represented in what some call "image schemata", which act as metaphorical frames and cognitive building blocks. Candidates for what image schemata could be include such familiar ontological foundation notions as: Objects, Process and Part-Whole relations, Motion, Full-Empty, Container, Blockage, Surface, Path, Link, Collection, Merging, Scale and Emerge (Oltramari, 2011). Some ontology design pattern work and reference ontologies have leveraged these notions, such as work on containment, motion and path.

This leads to many challenges:

- What are candidates for a set of knowledge that could provide adequate cognitive scaffolding? Various possibilities were discussed above.
- As part of scaffolding, an intelligent agent must represent relevant knowledge so that it is accessible and usable for achieving the agent's purpose. How is such meta-knowledge about representation learned?
- As part of scaffolding, intelligent agents need control mechanisms to find relevant pieces of knowledge in particular contexts. How is this learned and what knowledge is involved?
- How can an agent recognize existing patterns and entities, even with partial and/or noisy input?
- How can an agent determine what existing categories a pattern belongs to and how well it fits the categories?
- Predicting the near-term future is an important requirement for agent situation awareness. Even if a pattern has been recognized, how can it be projected into the future?
- How can new patterns and entities be learned and categorized?
- How can an agent select pertinent information at the input level as well as during learning and cognition?
- How can an agent learn new skills, both mental and physical?

## 4.4 Ontology Alignment

Currently, there is no general agreement on ontologies to handle the range of heterogeneous information in the Big Data age. While we have seen numerous efforts to create domain ontologies, the vocabularies and ontologies behind various data sources are not generally interoperable. General methods to merge and align ontologies include such things as PROMPT, an algorithm for semi-automatic merging and alignment of ontologies (Noy & Musen, 2000). But as noted more recently in Ontology Summits and related sessions, there are issues in reconciling and aligning ontologies with different assumptions and concepts (Hashemi et al., 2012; Stephen and Hahmann, 2016). There is work on ontology integration which has produced algorithms and heuristics with some success in making such computations tractable (Euzenat, 2004). However, the effective use of ontology formalisms (i.e., rules and axioms) as part of an integration process remains an open question (Udrea, Getoor, & Miller, 2007).

One of the things that makes the ontology integration process difficult is that as part of the process we need to understand the relationship between knowledge structures (classes and properties) and instance data in target ontologies. Existing ontology matching and alignment techniques are very restricted. They find similarities, equivalences and subsumption relations between two (or more) ontologies which must, at least, be syntactically and schematically integrated, have similar scope and context, and be no more expressive than OWL. In reality, semantic integration between ontologies of even a single domain (such as in hydrology) is much more problematic (Stephen and Hahmann, 2016). It requires translation of the ontologies' languages and a more rigorous specification of the semantics in each ontology. This can currently be done only by manual integration of the ontologies, but use of a suitable reference ontology may help automate this as in (Stephen and Hahmann, 2016). In addition, integration must have the capacity to use the semantics of the ontology to model the relationships between the ontologies being integrated, and to create a coherent and consistent integrated or aligned ontology.

Another abiding source of difficulty for matching parts of ontologies is that an ontology is designed with certain background knowledge (axiomized or not), for a purpose, and within a specific context (explicit or implicit). The context for an ontology can include the experience of the ontologists who developed the ontology, their preference for particular upper level ontologies, domain vocabularies, ontology design patterns or source data used in the development of the ontology. These may not be part of an ontology specification, and, thus, are not available to aligning tools or entity/relation matchers. This lack of background knowledge and context can lead to ambiguities (Shvaiko & Euzenat, 2008).

One example of an attempt to deal with the challenges of ontology alignment is the Ontology Alignment Evaluation Initiative (OAEI), which runs contests on ontology alignment ("Ontology Alignment", 2016). In the OAEI, ontology matchers are challenged with a robust set of ontology and data sources to be matched. For example, match the Adult Mouse Anatomy (2744 classes) with the National Cancer Institute Thesaurus (3304 classes) which describes the human anatomy. As in past campaigns, they use a systematic benchmark series to be matched. The work of this benchmark series has been to identify the areas in which each alignment algorithm is strong and weak.

## 4.5 Knowledge Graph Identification and Extraction

The reality of Big Data allows querying from massive repositories of potentially interrelated facts. Unfortunately, as noted in prior Ontology Summits, representing this information in rich formation to make it useful knowledge is a formidable challenge (Ray, Grüninger, Mason, & West, 2009; Hashemi et al., 2012). One interesting thrust is to transform source material (typically natural language text) into a knowledge graph form. A knowledge graph is a structure where entities are graph nodes, categories are word labels associated with each node, and relations are directed edges between the nodes. A knowledge graph is thus one simplified version of an ontology and something less formal than Sowa's conceptual graphs (Sowa, 1976). Such efforts to build even this simple structure require resolving entity identification and entity relationships. There is a degree of uncertainty and noise in and about such relationships targeted in these extractions as well as the need to infer missing information, and determining which candidate facts should be included into a knowledge graph as part of the identification process. One approach is to:

1. associate extraction confidences along with candidate facts,
2. identify co-referent entities, and
3. incorporate ontological constraints.

This approach relies on probabilistic soft logic (PSL), a recently introduced probabilistic modeling framework which easily scales to millions of facts such as demonstrated with extractions from the NELL project containing over 1M extractions and 70K ontological relations (Pujara, Miao, Getoor, & Cohen, 2013). The underlying mathematical framework of PSL supports extremely efficient inference continuous optimization task, which can be solved efficiently. PSL includes the ability to reason holistically about both entity attributes and relationships among the entities, along with ontological constraints. In practice, PSL has produced state-of-the-art results in many areas spanning NLP, social-network analysis, and computer vision. With PSL, large-scale knowledge graph extraction problems with millions of random variables can be orders of magnitude faster than existing approaches (Getoor, 2017).

While there have been advances in knowledge graph identification, it remains an open problem to extend iterative knowledge extraction and learning techniques of systems like NELL to systems that are capable of learning, retaining and using knowledge over a lifetime.

## 4.6 Processes

It should be obvious that learning, reasoning and ontologies occur within larger processes where they, and the relationships between them, form steps. John Sowa proposed that "For intelligent systems, the cognitive cycle is more fundamental than any particular notation or algorithm". Then he concluded that "By integrating perception, learning, reasoning, and action, the cycle can reinvigorate AI research and development" (Sowa, 2015, p. 56). Several of the summit presentations emphasized the importance of processes, which were mostly in the form of some kind of "feedback loop" (Aasman, 2017; Baclawski, 2017; Gil, 2017; Oltramari, 2017; Yu, 2017). We will refer to such loops as "cognitive cycles". There are many examples of cognitive cycles where learning, reasoning and ontologies all occur. The scientific discovery process is an example with a long history (Gil, 2017). A great many activities can be regarded as decision making cycles in which each iteration improves understanding and awareness by finding new knowledge as well as by rejecting some previous knowledge (Baclawski, 2017; Hitzler, 2017).

In the past, a single iteration of a cognitive cycle, such as the scientific discovery process, could take decades. Today, cognitive cycles occur more quickly, much more data must be processed and the data is more complex. Learning, reasoning and ontologies, and the relationships between them that are the subject of this Ontology Summit, can play important roles in cognitive cycles. Several previous summits are also relevant to the cognitive cycle. There is a need to deal with massive amounts of data (Grüninger et al., 2014). The data come from large collections of sensors (Underwood et al., 2015). Finally, the processing of the data requires many steps that must interoperate (Fritzsche et al., 2017).

While combining learning, reasoning and ontologies within cognitive cycles has potential advantages, it is not commonly practiced. To the extent that such processes are automated at all, they are generally ad hoc and informal. To automate the scientific discovery process, it is necessary to use NLP to extract the workflow of experimental activities that are performed, and the scientific hypotheses that are generated (Baclawski, Futrelle, Fridman, & Pescitelli, 1993; Gil, 2017). Another requirement is to record provenance, i.e., the origin of facts and knowledge. The challenge is to develop the required ontologies, to standardize them, to formulate best practices, and to convince communities to use them. In some cases, such as PROV-O for provenance and the Open Provenance Model for Workflows for the scientific discovery workflow, ontologies have been standardized (Garijo & Gil, 2014; PROV-O, 2013). However, other requirements of the cognitive cycle are less advanced. None of the existing ontologies are frequently used, and best practices are only starting to emerge (Baclawski, 2017; Gil, 2017).

There are many ways that learning, reasoning and ontologies can be combined with one another in a cognitive cycle. Learning and reasoning are fundamental for each of the steps in a cognitive cycle as well as the transition from one step to another step. Ontologies can be used to organize these computational processes. Such techniques can be used for processing relevant sensor data, which both make use of ontologies for organizing the data and also help develop these ontologies. The ontologies can have the further benefit of helping to make the data and the processing of the data more understandable. Ontologies can be the basis for explaining results thus helping to build trust in a system. Learning and reasoning could also be used at a meta-level to optimize a cognitive cycle, to detect problems with the cycle, and to help correct problems. Together, ontologies, learning and reasoning can be used to ensure that the components of a system interoperate with one another in the intended manner.

## 4.7 Dealing with Criticisms and Fluctuations

It is well known that the history of AI research is one of a series of boom and bust fluctuations ("AI History", 2017). The subfields of AI, as well as the field of AI as a whole, have exhibited extreme fluctuations, and there is risk that this will happen again. The field of ontology engineering is not immune to this same risk. While some domains, such as the biomedical domain, use ontologies heavily and very successfully, other domains, such as manufacturing, have had no significant applications of ontologies in spite of the many ontologies that have been developed (Smith, 2017). AI is expanding so rapidly in so many sectors of the economy that the mainstream media are beginning to question whether there are sufficient benefits derived from these investments (Hardy, 2016). One possible way to mitigate the risk is to take the criticisms seriously and to address them, rather than to ignore or ridicule the criticisms as has happened in the past ("AI Critiques", 2017).

Big Data in general has been criticized for several years, and ML is subject to some of the same criticisms (Marcus & Davis, 2014). While improving the relationships between ontologies, learning and reasoning cannot completely address these issues and concerns, they may be able to help.

1. Understanding the results is still essential. Indeed strong AI, thinking like a person, ultimately faces the challenge of representing and using the knowledge available to people. This issue was discussed extensively in Sections 4.2, 4.3 and 4.6, and many of tracks and presentations in the summit emphasized the importance of this issue and proposed methods for dealing with it by strengthening the connections between ML and ontologies (Baclawski, 2017; Davidson, 2017; Gil, 2017; Presutti, 2017).

2. ML results most commonly are essentially a collection of correlations. One criticism is the common mistake of presuming that a correlation automatically implies causation. Understanding the results of ML in human terms can help eliminate at least the more implausible examples of inferring causation from correlation.

3. There is often a lack of consistency and interoperability of the data being used in Big Data applications. Interoperability was the topic of the Ontology Summit 2016. We discussed this issue in Sections 2.4, 4.4 and 4.6 above.

4. A common criticism of Big Data applications is that the questions being asked are often too imprecise. In other words, there is a disconnect between the queries being performed on the data and the interpretation of the results. Precision is one of the goals of ontologies. Improving the relationship between ontologies and ML would help ensure that interpretations are consistent and understandable.

The problem of disenchantment with ontology engineering in some domains, such as manufacturing, could potentially be addressed by establishing a "foundry" initiative in each domain, similar to the Open Biological Ontology (OBO) Foundary that has been very successful in the biomedical community ("Open Biological Ontologies", 2003). However, it is an open problem whether one can replicate the success of the OBO Foundary in other domains (Smith, 2017).

# 5 Conclusion

The Ontology Summit 2017 has examined a wide range of issues, opportunities, challenges and future prospects for the interconnections among learning, reasoning and ontologies in the context of AI. We learned that one can uncover useful knowledge about biological organization using ML and GO, but it is open whether this achievement is possible in other domains. We discussed the problem of identifying the notion of context, which remains an unsolved problem, although some progress has been made in some domains. We noted that AI in general, and ML in particular, is not possible without some kind of starter knowledge; and we examined the many challenges in developing the "cognitive scaffolding" for such knowledge. We revisited the issue of interoperability that was covered in last year's summit, but for the specific problem of ontology alignment, and determined that it is a significant open problem, although progress is being made. We examined the problem of extracting knowledge graphs from source material for use in ML and reasoning, and noted the opportunities for advances in AI using iterative knowledge extraction and learning techniques for general AI applications. We explored the processes that organize applications of learning, reasoning and ontologies as well as the relationships between them; and we proposed opportunities for addressing many important problems relevant to general AI applications. We showed how one can address, to some extent, some of the recent criticisms of ML and Big Data by strengthening the connections between learning, reasoning and ontologies. We also considered how one may be able to address the problem of the acceptance of ontology engineering techniques in domains that have not been receptive to them.

# 6 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 NIST who generously contributed the use of their facilities for the Ontology Summit Symposium.

# References

Aasman, J. (2017). Cognitive probability graphs need an ontology. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xOVeOM`

Albertazzi, L. (2000). Which semantics? In L. Albertazzi (Ed.), *Meaning and cognition. a multidisciplinary approach* (p. 1-24). Amsterdam: Benjamins Publishing Company.

Anderson, J., & Lebiere, C. (1998). *The Atomic Components of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Baclawski, K. (2017). Toward combining ontologies and machine learning for improving decision making. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2DVUS9D`

Baclawski, K., Chan, E., Gawlick, D., Ghoneimy, A., Gross, K., Liu, Z., & Zhang, X. (2017). Framework for ontology-driven decision making. *Applied Ontology*, *12*(3-4), 245–273. http://dx.doi.org/10.3233/AO-170189

Baclawski, K., Futrelle, R., Fridman, N., & Pescitelli, M. (1993). Database techniques for biological materials & methods. In *First Int. Conf. Intell. Sys. Molecular Biology* (pp. 21–28).

Baclawski, K., Gross, K., Chan, E., Gawlick, D., Ghoneimy, A., & Liu, Z. (2017). Self-adaptive dynamic decision making processes. In *IEEE Conference on Cognitive and Computational Aspects of Situation Management.* Retrieved December 9, 2017 from `http://bit.ly/2fOG9G2`

Baclawski, K., Malczewski, M., Kokar, M., Letkowski, J., & Matheus, C. (2002). Formalization of situation awareness. In *Eleventh OOPSLA Workshop on Behavioral Semantics* (pp. 1–15). Seattle, WA.

Bell, B., & Kendall, E. (2017). Leveraging FIBO with semantic analysis to perform onboarding, KYC and CDD. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xYf9dI`

Bennett, M., & Westerinen, A. (2017). Track B Summary Report: Using background knowledge to improve machine learning results. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2fPCKa6`

Berg-Cross, G. (2017). Track A Summary Report: Using automation and machine learning to extract knowledge and improve ontologies. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2fN7rws`

Buitelaar, P. (2017). Ontology learning – some advances. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xOKH6p`

Buitelaar, P., Cimiano, P., & Magnini, B. (Eds.). (2005). *Ontology Learning from Text: Methods, Evaluation and Applications*. IOS Press.

Consortium, T. G. O. (2008). The Gene Ontology project in 2008. *Nucleic Acids Research*, *36 (Database issue)*, D440-4. doi:10.1093/nar/gkm883. PMC 2238979Freely accessible. PMID 17984083

Corcoglioniti, F. (2017). Frame-based ontology population from text with PIKES. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2ymshKG`

*Critiques of Artificial Intelligence.* (2017). Wikipedia. Retrieved December 9, 2017 from `http://bit.ly/2ymeeVF`

*DARPA Agent Markup Language homepage.* (2006). Retrieved December 9, 2017 from http://www.daml.org

Davidson, S. (2017). The Investigator's Toolkit: Deriving immediate, actionable insights from unstructured data. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xOGGyH

Devlin, K. (1991). *Logic and Information*. Cambridge, U.K.: Cambridge University Press.

Erekhinskaya, T. (2017). Converting text into FIBO-aligned semantic triples. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xOUBoo

Euzenat, J. (2004). An API for ontology alignment. In *International Semantic Web Conference* (pp. 698–712). Springer Berlin Heidelberg.

Falk, C. (2017). The Meaning-Based Machine Learning Project. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xYfaOO

Fritzsche, D., Grüninger, M., Baclawski, K., Bennett, M., Berg-Cross, G., Obrst, L., … Westerinen, A. (2017). Ontology Summit 2016 Communiqué: Framing the Conversation: Ontologies within Semantic Interoperability Ecosystems. *Applied Ontology*, *12*(2), 91–111.

Fritzsche, D., & Sriram, R. (2017). Track C Summary Report: Ontologies and Reasoning. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2ylLf4i

Garijo, D., & Gil, Y. (2014). *Open Provenance Model for Workflows (OPMW).* Retrieved on May 15, 2018 from http://bit.ly/1L8XFrg

*The genotype-phenotype distinction.* (2017). Wikipedia. Retrieved December 9, 2017 from http://bit.ly/sA5sj

Georgeff, M., & Lansky, A. (1987). *Procedural knowledge* (Tech. Rep. No. Technical Report 411). Artificial Intelligence Center, SRI International.

Getoor, L. (2017). Combining statistics and semantics to turn data into knowledge. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xdX07H

Gil, Y. (2017). *Reasoning about scientific knowledge with workflow constraints: Towards automated discovery from data repositories.* Retrieved December 9, 2017 from http://bit.ly/2yllNMa

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.

Hahmann, T., Stephen, S., & Broderic, B. (2016). *Domain reference ontologies vs. domain ontologies: What's the difference? Lessons from the water domain.* Retrieved on May 15, 2018 from http://bit.ly/2fPfEjR

Hardy, Q. (2016). *Artificial intelligence software is booming. But why now?* NY Times. Retrieved December 9, 2017 from http://bit.ly/2d0jz71

Hashemi, A., Schneider, T., Bennett, M., Brady, M., Casanave, C., Graves, H., … Whetzel, T. (2012). Ontology for Big Systems: The Ontology Summit 2012 Communiqué. *Applied Ontology*, *7*(3), 357–371.

Hayes, P. (1997). Contexts in context. In *AAAI Fall Symposium.* Retrieved on May 15, 2018 from http://bit.ly/2jf0BBa

*History of Artificial Intelligence.* (2017). Wikipedia. Retrieved December 9, 2017 from http://bit.ly/oT29RS

Hitzler, P. (2017). On the roles of logical axiomatizations for ontologies. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xYCV9N

Hruschka, E. (2017). Never-Ending Learning approach for populating and extending an ontology. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xYeEjN

Huber, M. (1999). A BDI-theoretic mobile agent architecture. In *Agents '99. Proceedings of the Third Annual Conference on Autonomous Agents* (pp. 236–243). Seattle, WA, USA: ACM.

Johnson, M. (1990). *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. The University of Chicago Press.

Koehn, P. (2010). *Statistical Machine Translation*. Cambridge, UK: Cambridge University Press.

Kokar, M., Matheus, C., & Baclawski, K. (2009). Ontology-based situation awareness. *Information Fusion*, *10*(1), 83–98.

Kuksa, E. (2017). Reasoning with ontologies in Ontohub. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xevLtD

Laird, J., Newell, A., & Rosenbloom, P. (1987). Soar: an architecture for general intelligence. *Artificial Intelligence*, *33*(1), 1–64.

Lebo, T., Sahoo, S., & McGuinness, D. (Eds.). (2013). *PROV Ontology (PROV-O).* Retrieved on May 15, 2018 from http://bit.ly/2xPcx2k

*Machine learning: the power and promise of computers that learn by example.* (2017). The Royal Society. Retrieved December 9, 2017 from http://bit.ly/1Ng9nHf

Marcus, G., & Davis, E. (2014). Eight (no, nine!) problems with Big Data. *NY Times*. Retrieved December 9, 2017 from http://bit.ly/1kgErs3

McMillan, R. (2015). *AI has arrived and that really worries the world's brightest minds.* Retrieved December 9, 2017 from http://bit.ly/1KPX68u

Neisser, U. (1987). From direct perception to conceptual structure. In *Concepts and Conceptual Development* (p. 11-24). Cambridge: Cambridge University Press.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA, USA: Harvard University Press.

Noy, N., & Musen, M. (2000). PROMPT: algorithm and tool for automated intelligence. In *AAAI-2000.* Austin, TX.

Oltramari, A. (2011). An introduction to hybrid semantics: the role of cognition in semantic resources. In *Modeling, Learning, and Processing of Text Technological Data Structures* (pp. 97–109). Springer Berlin Heidelberg.

Oltramari, A. (2017). From machines that learn to machines that know: the role of ontologies in machine intelligence. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from http://bit.ly/2xOVdKI

*Ontology Alignment Evaluation Initiative.* (2016). Retrieved December 9, 2017 from `http://bit.ly/2fPmwxL`

*Open Biological Ontologies website.* (2003). Retrieved December 9, 2017 from `http://bit.ly/HDFc7`

Pafilis, E. (2017). EXTRACT 2.0: interactive extraction of environmental and biomedical contextual information. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xOpZDA`

Pileggi, S. (2016). Probabilistic semantics. *Procedia Computer Science*, *80*, 1834–1845.

Presutti, V. (2017). Semantic Web machine reading with FRED. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xe6Feo`

Pujara, J., Miao, H., Getoor, L., & Cohen, W. (2013). Knowledge graph identification. *Lecture Notes in Computer Science*, *8218*, 542–557.

Rao, A., & Georgeff, M. (1991). Modeling rational agents within a BDI-architecture. In R. Fikes & E. Sandewall (Eds.), *Proceedings of Knowledge Representation and Reasoning (KR&R-91)* (pp. 473–484). Morgan Kaufmann Publishers, San Mateo, CA.

Ray, S., Grüninger, M., Mason, H., & West, M. (2009). *Ontology Summit 2009 Communiqué: Toward Ontology-based Standards.* Retrieved December 9, 2017 from `http://bit.ly/2ylSDga`

Rugaber, S. (2017). Applications of ontologies to biologically inspired design. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xdOm9u`

Samuel, A. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 535–554.

Shvaiko, P., & Euzenat, J. (2008). Ten challenges for ontology matching. In *Proc. 7th International Conference on Ontologies, Databases, and Applications of Semantics (ODBASE)* (pp. 1163–1181). Monterey, Mexico: Springer-Verlag. Lecture Notes in Computer Science; On the Move to Meaningful Internet Systems: OTM 2008

Smith, B. (2017). The Industry Foundry Ontology. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2xPfoZc`

Sowa, J. (1976). Conceptual graphs for a data base interface. *IBM Journal of Research and Development*, *20*(4), 336-357. doi:10.1147/rd.204.0336

Sowa, J. (2015). *Why has AI failed? And how can it succeed?* Retrieved December 9, 2017 from `http://bit.ly/1jSAQSz`

Stephen, S., & Hahmann, T. (2016). *Semantic alignment of the Groundwater Markup Language with the Reference Hydro Ontology HyFO.* Retrieved December 9, 2017 from `http://bit.ly/2ymd7oX`

Udrea, O., Getoor, L., & Miller, R. (2007). Leveraging data and structure in ontology integration. In *Proc. 2007 ACM SIGMOD International Conference on Management of Data.* ACM.

Underwood, M., Grüninger, M., Baclawski, K., Bennett, M., Berg-Cross, G., Hahmann, T., … Sriram, R. (2015). Ontology Summit 2015 Communiqué: Internet of Things: Toward Smart Networked Systems and Societies. *Applied Ontology*, *10*(3,4), 355–365.

Wactlar, H. (2017). Robust intelligence research at NSF. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2ymUa5p`

Wray, R., & Laird, J. (2003). An architectural approach to ensuring consistency in hierarchical execution. *J. Artificial Intelligence Research*, *19*, 355–398.

Yu, M. (2017). Inferring the hierarchical structure and function of a cell from millions of biological measurements. In *Ontology Summit 2017: AI, Learning, Reasoning, and Ontologies.* Retrieved December 9, 2017 from `http://bit.ly/2ymr0m0`