
Bootstrapping the applied ontology practice: Ontology communities, then and now

Peter P. Yim

ONTOLOG and CIM Engineering

San Mateo, California, USA

peter.yim@cim3.com

Abstract. After more than two decades since the interdisciplinary field of *ontology* emerged for information science, it is now poised to make a huge impact. The ontological approach aims at making intended meaning explicit and computable. By enabling better human-machine “understanding” across complex socio-technical chains, ontologies are facilitating the vision of co-evolution of humans and their tools. Applying ontology requires a very broad spectrum of disciplines and skills, and is therefore best tackled not by individuals, but by communities of people working collaboratively. This position paper attempts to trace the brief history of ontology communities from the author’s vantage point, identifying some of the key dates and key players (mostly familiar names, but some unsung heroes as well) that have made things happen for the field of applied ontology. Inspired by Doug Engelbart’s “bootstrap” strategy, the author has developed and operated (until his recent retirement) a virtual collaborative environment ([CIM3.net](http://cim3.net)), on which some ontology communities (such as Ontolog and Ontology Summit) have thrived. Observing that the funding and career opportunities have changed for the better in the past few years, the author reflects on how the ontology communities of practice might leverage what they have built so far, and take applied ontology into an even better future.

Keywords: Ontology, applied ontology, ontology community, applied ontology practice, community of practice, history of ontology community, future of ontology community, bootstrap, collaboration, collaborative work environment

This position paper has been published in 2015 Volume 10, No. 3-4 of the *Applied Ontology* journal (IOS Press). Citation:

//

Yim, Peter P. *Bootstrapping the applied ontology practice: Ontology communities, then and now*. *Applied Ontology*, vol. 10, no. 3-4, pp. 229-241, 2015

//

1. Introduction

The core concepts of ontology and its application have been around for more than 2300 years. For the purpose of this paper, though, we will look only at the development of the field of *ontology* and *applied ontology* since the early 1990s. That was the time when pioneers of this discipline co-opted the term *ontology* from philosophy, for use in information systems, knowledge engineering, and artificial intelligence (Gruninger et al., 2008).

By *applied ontology practice*, we mean the acts of those who enable and promote in the science and engineering disciplines the productization and dissemination of ontology and ontology-based technologies. Therefore, one should expect the members of the community for this practice to include, not only researchers and developers, but also those who enable the practice – from the policy decision makers who fund R&D projects, to the ontology standards developers, to the educators and trainers who ensure an adequate supply of skilled professionals for the field, to the entrepreneurs and innovators, and to the

marketing and sales people who encourage the market to “buy” the ontology-based product or solution, to name a few.

The author comes from a vastly different background than do most of the current members of the ontology community. Before devoting his late career to *applied ontology*, he was in many different fields.¹ This difference in background accounts for perspectives and approaches he has taken that may make his account of interest to the applied ontology community.

The common thread that ties him to everyone else in the field is that he believes “ontology and its applications” is an important pursuit. He believes that applied ontology will improve the world for humanity, and that the goals of applied ontology are achievable. To him, the ontological approach makes intended meaning explicit and computable. By enabling better human–machine “understanding” across complex socio-technical chains, ontologies are facilitating the vision of co-evolution of humans and their tools. On that belief, he dedicated the last 15 years of his working career to facilitating applied ontology. He has hosted, supported, and even co-convened some interesting ontology communities, collaboration platforms, and projects during that period.

2. Ontology and ontology communities

One hears much about the roots of today’s *ontology* concepts (“formal ontology” especially) dating as far back as Aristotle (384–322 BC). We will, however, fast forward to the time when the term *ontology* (in the sense that we are using it in information science today) first emerged. That takes us to the early 1990s (with seminal publications in the 1991 to 1993 timeframe and onwards), notably with work by researchers in the United States, Italy, the United Kingdom, and Germany. One of the early papers from Tom Gruber (1993) contains an extensive bibliography that helps us to identify some of the early pioneers.

The field of *ontology* began as a basic-research pursuit, supported by government research funding (from the likes of the US Defense Advanced Research Project Agency, the National Institutes of Health, the National Science Foundation, the European Union and various research centers in Europe), and at research departments of major corporations. Work was typically carried out at academic and research institutions. Government- or corporate-funded work created teams and communities, and some of those teams were restricted to participants of specific projects or programs. Some early community efforts, such as the Knowledge Sharing Effort (Neches et al., 1993), moved the field forward significantly, especially in research. For this paper, though, we will focus more on the emergent, open, international, community efforts.

2.1. Year 1993: The First Formal Ontology Workshop

In the debut presentation of the International Association for Ontology and Its Applications (IAOA), Nicola Guarino identified a 1993 workshop as the first among key events that marked the development of ontology communities (Guarino et al., 2009). For the first time, researchers from different countries and disciplines convened physically to have a discourse on *ontology*. This was the “International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation” that took place 17–19

¹Some of the very diverse fields that the author has previously worked in, include mathematical modeling (analog and digital), electronic design, manufacturing, software engineering, system integration, professional services, management, future studies, technology policy, information technology (IT) infrastructure, and collaboration technology.

March 1993, in Padova, Italy (Guarino and Poli, 1993). Collected among the historic artifacts from that workshop was the preface which outlined a new direction that some among us in the ontology community have been pursuing since. The contributors of that workshop also constituted some of the field's pioneers – a good number of whom are still with us today and who are still helping to guide our way forward.

2.2. Year 1995: The Protégé community

One of the oldest and largest communities in the ontology-related domain is the Protégé user community. Protégé is an ontology-development tool and a platform for the construction of knowledge-based systems and ontology-based solutions. It has its roots in Mark Musen's PhD research work in the 1980s (Musen, 1989). Musen worked on a domain-specific knowledge-acquisition tool, known as OPAL, to build the complex knowledge bases required in a cancer chemotherapy advisor called ONCOCIN. That work was completed in the early 1980s. OPAL and ONCOCIN soon transitioned into Protégé-I. The notion of ontology was introduced to the Protégé tool (in Protégé-II, early 1990s) to provide a more generic, component-based approach that would allow for automatic generation of knowledge-entry forms and run-time selection of alternative problem-solving methods (Musen, 2015).

As the Protégé tool began to be adopted, a user community emerged. In this case, the community was created around the tool/technology, and served to foster collaborative research, development, education, and user support. When asked by the author, when and how did the community get started, Musen recalled, "Four of us met at a bar at a conference in Pavia in 1995. We called it 'The first Protégé User Group Meeting.' Our e-mail lists started at around the same time too". As of this writing, one can see on the Protégé website² that the tool has more than 290,000 registered users (with unknown numbers having downloaded the software without registering); the team has 3500 subscribers on the [protege-user] mailing list, as well as more than 18,000 participants on their announcement list.³

2.3. Year 1998: The GO and OBO communities

Communities also formed around domain-specific ontologies. The Gene Ontology⁴ (GO) community is a case in point. In an exchange with the author, Chris Mungall (one of the community administrators) recounts: "The GO community started back in 1998, marked by a commitment of GO to the Concurrent Version System (CVS) repository. That year also marked the coming together of model organism databases around the initial version of GO that Michael Ashburner and Suzanna Lewis produced". As Mungall points out, the various GO mailing lists are primarily for GO Consortium members and people applying GO terms to genes as part of their bio-curation work. While the size of the GO discussion list is 100–200 subscribers, the size of the wider community of users is much larger, as there are hundreds of software applications doing GO-based analyses, and each of these has its own community of users (although those communities are no longer tightly knit together).

Successful communities also tend to extend themselves in positive ways. As an example, in 2001, Ashburner and Lewis started the OBO (Open Biomedical Ontologies) initiative.⁵ They aimed to coordinate

²See: <http://protege.stanford.edu/>.

³See: <http://protege.stanford.edu/community.php>.

⁴See: <http://www.geneontology.org/>.

⁵See: <http://www.obofoundry.org/>.

with other ontology developers for the life sciences so that they would apply the key principles underlying the success of the GO – namely, that ontologies be open, orthogonal, instantiated in a well-specified syntax, and designed to share a common space of identifiers (Smith et al., 2007). The OBO community thus emerged. As of this writing, 135 ontologies are included in the OBO Library (the repository for OBO ontologies), and their mailing list has more than 300 subscribers.

2.4. Year 1998: The First Formal Ontology in Information Systems (FOIS) Conference

Five years after their first international workshop, the Formal Ontology community gained enough momentum and support to initiate its international conference series. FOIS'98 – the First International Conference on Formal Ontology in Information Systems was held in Trento, Italy, from June 6 to 8, 1998 (Guarino, 1998). Like other scientific conferences, the FOIS conferences provide the ontology research and development community with a familiar platform for collaboration. FOIS addresses diverse domains, such as conceptual modeling, database design, software engineering, organizational modeling, artificial intelligence, computational linguistics, the life sciences, bioinformatics, geographic information science, knowledge engineering, information retrieval, and the Semantic Web. FOIS'98 marked the beginning of an important collaborative effort of the formal ontology community. Eight FOIS Conferences⁶ have been held since 1998, and the FOIS conference is currently held every other year.

2.5. Year 1998: The Heidelberg workshop

Another landmark event for the ontology community took place during 1998. While unrelated, this event was held in Germany immediately after the FOIS Conference in Italy that year. It was the Heidelberg workshop sponsored by the Klaus Tschira Foundation, and convened at Villa Bosch, Heidelberg, Germany, June 10–16, 1998. It was a first attempt towards achieving some sort of convergence on basic ontological categories and relations among representatives of a broad interdisciplinary community. The author only first heard about this event from Nicola Guarino when they met at the US National Institute of Standards and Technology for the Upper Ontology Summit almost eight years later. To piece things together for this paper, the author received much help from Pat Hayes and John Sowa who provided details and artifacts relating to the workshop (Spillers et al., 2015). They all remembered that it was Robert Spillers (a retired IBM manager) who did most of the organizational work to make this event happen. Unfortunately, those in attendance did not arrive at a consensus that was publishable. That said, the mail thread among participants of this workshop initiated an informal virtual collaboration of formal ontology researchers, and inspired the emergence of the Standard Upper Ontology (SUO) effort in the following years.

2.6. Year 2000: Start of the IEEE: SUO working group

Inspired by the 1998 Heidelberg workshop and supported by Robert Spillers and others, James Schoening initiated the IEEE P1600.1: Standard Upper Ontology Working Group (SUO WG) in June 2000. Schoening led that effort for six years (Schoening et al., 2015). The community worked and thrived on a set of mailing lists hosted by IEEE and administered by Schoening. The SUO Working Group grew to 88 voting members and logged 25,000 postings. They went as far as shortlisting six candidate upper ontologies, but never came to building a consensus on the one “standard upper ontology” they had set out

⁶See: <http://iaoa.org/fois/>.

to create. The SUO mailing list, however, set the stage for virtual collaboration among geographically distributed members of the community, and paved the way for some of the same players to collaborate in the Ontolog Forum, after the Ontolog community emerged in 2002.

2.7. Year 2002: The ONTOLOG community of practice comes into being

In 2001, the author was on an OASIS Technical Committee (TC) led by Jon Bosak (widely recognized as “the father of XML”), developing an XML standard known as the Universal Business Language (UBL).⁷ The author suggested then to develop UBL as an ontology. Along with Leo Obrst and other former colleagues,⁸ a virtual presentation was made in April 2002, a “UBL-Ontolog” mailing list was started in mid-May, and a face-to-face presentation was made in June that year, to make the case to the UBL-TC members. However, two factors made the move toward an “ontology” approach problematic for UBL: (i) early feedback from the participants indicated more interest in learning about ontologies than applying experience with ontological development, and (ii) tight timetables, deadlines, and priorities within the UBL effort made that inappropriate. To get the sense of urgency across to the team, Jon Bosak had, on one occasion, said, “if we don’t come up with a UBL invoice standard next month, another 5 people will have created their own invoice syntax and vocabulary”. Ontology clearly wasn’t mature enough to deliver what was needed then. Everyone reached consensus that the effort would best be spun off and made into an independent mailing list, which would garner even more support from a broader community. Bosak (who was actually a philosophy major at college) was very supportive. He kept reminding everyone that they should remain pragmatic, and strive to address application-oriented issues – something that the author took to heart. The Ontolog community (also known as the “Ontolog Forum”) emerged, co-convened by Kurt Conrad, Leo Obrst and Peter Yim, in September 2002. An independent mailing list was created, hosted on CIM3.NET, and an invitation was disseminated (Conrad et al., 2002) to invite participation from people in the broader ontology community, the standards development community, and the knowledge management community. Ontolog was envisioned as an “open dialog in ontology”. It was designed to be a Community of Practice (CoP) in the sense that John Seely Brown (of Xerox PARC) would have it: “a small group of people who have worked together over a period of time. Not a team, not a task force, not necessarily an authorized or identified group. They are peers in the execution of “real work”. What holds them together is a common sense of purpose and a real need to know what each other knows” (Brown and Gray, 1995). There seems to have been a lot of pent-up energy for such a community during those early years. Coupled with the support of a virtual Collaborative Work Environment (CWE) by way of the CIM3.NET infrastructure, Ontolog attracted some very talented and dedicated individuals, who pushed it forward with exciting online discussions, virtual activities, and projects. The Ontolog community grew into, arguably, the strongest community of its kind – an open, international, virtual, community of practice devoted to ontology, ontological engineering, and semantic technologies. Ontolog developed processes that were regarded as community best practices that others were following. With a richer platform for collaboration, some of the discussions that formerly took place on the aforementioned SUO WG mailing list gradually migrated over

⁷See: https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ubl.

⁸These were his former colleagues from VerticalNet (a company that was among the first to employ an “ontology” approach in business-to-business e-commerce applications). They included Adam Cheyer (who used to run Software Engineering), Jack Park (a senior software engineer), Leo Obrst (who used to run Ontology Engineering) and Peter Yim (who used to run Program Management) at that company. After they all left VerticalNet (which went bust during the bursting of the 2000 dot-com bubble) they started a side-project called “SOUL” – an acronym for “Semantics, Ontology and Universal Language” – which was behind the push to get UBL to adopt an ontology approach.

to the Ontolog-forum list (the discussion list for Ontolog community members). Ontolog also went on to provide the virtual workspace for, and played host to, other ontology- or ontology-standards-related communities and their efforts, on its virtual CWE. The CWE also served as the community's Dynamic Knowledge Repository (DKR), capturing and archiving a huge body of knowledge, as Ontolog members interacted to share their expertise and experience. The community has grown from dozens of people back in 2002, to include about a 1000 members, and to about 3000 subscribers on its event announcement mailing list. After Peter Yim, who used to run the day-to-day operations of Ontolog, announced in 2014 that he would be phasing into his retirement, a Board of Trustees was formed to provide leadership, and to administer the transition and the continuing operations of Ontolog.

2.8. Year 2006: The start of the Ontology Summits

An important event was conceived and co-convened in 2006 by a few of the Ontolog members – Patrick Cassidy, Leo Obrst, Steve Ray, and Peter Yim. This was another attempt to bring together upper-ontology gurus. Pat Cassidy suggested that the event be called an “Ontology Summit”. With the support of Steve Ray, then a division chief at the US National Institute for Standards and Technology (NIST), Ontolog and NIST collaborated in 2006 to organize the “Upper Ontology Summit”. In the spirit of the Heidelberg workshop and the IEEE SUO Effort, this Summit attempted to get the stewards of the major open upper ontologies to come together and converge on some consensus. A ten-point joint communiqué (Obrst et al., 2006) was released by the representatives of eight upper ontologies well known at that time, namely: the Generalized Upper Model (GUM) by Bateman et al. (1995); the ontology of Descriptions and Situations (D&S) by Gangemi and Mika (2003); the Process Specification Language (PSL/ISO 18629) by Gruninger and Menzel (2003); the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) by Masolo et al. (2003); OpenCyc by Lenat and Guha (1989); the Suggested Upper Merged Ontology (SUMO) by Niles and Pease (2001); the Basic Formal Ontology (BFO) by Arp et al. (2015); the upper ontology based on ISO 15926 by Batres et al. (2007).

Since then, Ontology Summits have been held each year to explore a theme of choice. Not only does this structure provide a chance for the community to have its top minds spend time together to address an issue that is of importance to the advancement of the field of *ontology*, it also creates a platform for the ontology community to work collaboratively with other related communities to jointly pursue in depth a chosen theme. The “Ontology Summit” process matured as time progressed, and this annual program now comprises almost four months of virtual discourse (over archived mailing lists) and virtual panel presentations and discussion sessions (over augmented conference calls), and culminates in a two-day face-to-face workshop, during which the community, among other things, shares its findings and present its distilled thoughts in a collaboratively developed communiqué. Over the past ten years, the Ontology Summit program offered themes that included:

- Ontology Summit 2006 – “The Upper Ontology Summit”⁹ (Obrst et al., 2006).
- Ontology Summit 2007 – “Ontology, Taxonomy, Folksonomy: Understanding the Distinctions”¹⁰ (Gruninger et al., 2008).
- Ontology Summit 2008 – “Toward An Open Ontology Repository”¹¹ (Obrst et al., 2008).
- Ontology Summit 2009 – “Toward Ontology-based Standards”¹² (Ray et al., 2009).

⁹See: <http://ontolog.cim3.net/wiki/UpperOntologySummit.html>.

¹⁰See: <http://ontolog.cim3.net/wiki/OntologySummit2007.html>.

¹¹See: <http://ontolog.cim3.net/wiki/OntologySummit2008.html>.

¹²See: <http://ontolog.cim3.net/wiki/OntologySummit2009.html>.

- Ontology Summit 2010 – “Creating the Ontologists of the Future”¹³ (Neuhaus et al., 2011).
- Ontology Summit 2011 – “Making the Case for Ontology”¹⁴ (Uschold, 2011).
- Ontology Summit 2012 – “Ontology for Big Systems”¹⁵ (Schneider et al., 2012).
- Ontology Summit 2013 – “Ontology Evaluation Across the Ontology Lifecycle”¹⁶ (Neuhaus et al., 2013).
- Ontology Summit 2014 – “Big Data and Semantic Web Meet Applied Ontology”¹⁷ (Obrst et al., 2014).
- Ontology Summit 2015 – “Internet of Things: Toward Smart Networked Systems and Societies”¹⁸ (Underwood et al., 2015).

Co-organizing institutions for the Ontology Summits have also grown to include six core organizers: the Ontolog Forum (Ontolog), the National Institute of Standards and Technology (NIST), the National Center for Ontological Research (NCOR), the National Center for Biomedical Ontology (NCBO), the International Association for Ontology and its Applications (IAOA), and the US National Coordination Office for Networking and Information Technology Research and Development (NCO-NITRD). At times, other institutions and communities relevant to a particular year’s theme have also participated (Yim et al., 2015).

2.9. Year 2009: The International Association for Ontology and Its Applications (IAOA) was founded

Several participants at the 2006 FOIS Conference in Baltimore, USA, suggested that a formal professional membership organization be incorporated to advance the field of ontology. The suggestion was more adequately discussed at the next FOIS conference in 2008 in Saarbrücken, Germany, and a consensus was reached that such an endeavor be pursued. The International Association for Ontology and Its Applications (IAOA)¹⁹ was proposed, and Nicola Guarino was requested by the community to lead the planning and organization work. IAOA (Guarino et al., 2009) was formally established in Trento, Italy, on April 29th 2009 as a non-profit corporation under the laws of Italy, with the following mission statement:

The International Association for Ontology and its Applications is a non-profit organization the purpose of which is to promote interdisciplinary research and international collaboration at the intersection of philosophical ontology, linguistics, logic, cognitive science, and computer science, as well as in the applications of ontological analysis to conceptual modeling, knowledge engineering, knowledge management, information-systems development, library and information science, scientific research, and semantic technologies in general.

The leadership of IAOA is provided by an Executive Council, with the support of an Advisory Board. IAOA has adopted the FOIS Conference series (started 1998, as described above) as its flagship activity, and it has designated the journal *Applied Ontology*²⁰ (which was started in 2005) an affiliated publication. The IAOA President co-chairs the annual Ontology Summit. The organization runs a Summer

¹³See: <http://ontolog.cim3.net/wiki/OntologySummit2010.html>.

¹⁴See: <http://ontolog.cim3.net/wiki/OntologySummit2011.html>.

¹⁵See: <http://ontolog.cim3.net/OntologySummit/2012/>.

¹⁶See: <http://ontolog.cim3.net/OntologySummit/2013/>.

¹⁷See: <http://ontolog.cim3.net/OntologySummit/2014/>.

¹⁸See: <http://ontolog.cim3.net/OntologySummit/2015/>.

¹⁹See: <http://iaoa.org>.

²⁰See: <http://www.iospress.nl/journal/applied-ontology/>.

Institute every year, hosts various workshops, operates a range of technical committees and special interest groups, provides scholarships and other incentives to upcoming scholars in the field, and supports other related professional events.

2.10. Ontology standards communities

Over the years, many project teams and small communities have formed around the task of developing ontology-related standards, and they have contributed significantly to the advancement of the field. These groups generally work under the auspices of a standards development organization such as the International Organization for Standardization (ISO), International Electrotechnical Commission (IEC), the United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT), the Institute of Electrical and Electronics Engineers (IEEE), the Organization for the Advancement of Structured Information Standards (OASIS), the Object Management Group (OMG) or the World Wide Web Consortium (W3C).

The discourse for the Ontology Summit of 2009 revolved around ontology and standards. Some of the motivations and efforts of ontology standards communities are described in the Ontology Summit Communiqué from that year.²¹

In their editorial “Applied Ontology: The next decade begins”, Guarino and Musen (2015) describe some standards and de facto standards. They reflect some of the results of dedicated efforts by members of these ontology standards communities over the years.

3. Behind CIM3.NET

The notion of a CIM3 system was first introduced by the author in the late 1980s and early 1990s while he was still a practitioner in manufacturing (Yim, 1991). The CIM3 acronym originally stood for “Computer Integrated Man-Machine Manufacturing”, but it evolved by the late 1990s to stand for “Collaboration in huMan-Machine-Methodology”. CIM3 is an approach to holistic organizational improvement through extending the traditional “automation system” architecture to incorporate computer-supported cooperative work (CSCW) techniques, such as the speech-act-based “language/action perspective” in the design of a system (Winograd et al., 1991).

The CIM3.NET is a cloud computing infrastructure, designed and implemented by the author and his colleagues on the [cwe-dev] team, to provide Collaborative Work Environment (CWE) support and services to virtual communities (Yim, 2009). A prototype system of this kind was first implemented, presented, and employed at the Doug Engelbart Colloquium held at Stanford University (January–March 2000), which the author co-organized, when he was then a staff advisor at Engelbart’s Bootstrap Institute (Engelbart et al., 2000). Features of such a system (with enhancements from Adam Cheyer) were also implemented for company internal collaboration purposes, during the author’s next stint at VerticalNet. Unfortunately, VerticalNet business failed as a result of the dot-com stockmarket crash and the 2001 “9/11” terrorist attack in the US.²² After VerticalNet, the author made the CIM3.NET system and ser-

²¹See: http://ontolog.cim3.net/wiki/OntologySummit2009_Communique.html.

²²It was unfortunate because VerticalNet was a promising commercial enterprise, among the first of its kind to employ an applied ontology approach, with a reasonably sized pool of talent (at its peak, VerticalNet had an ontology engineering staff of 47, made up of ontologists, computer scientists and programmers, and subject matter experts in many commercial domains). The company failed mainly due to uncontrollable external economic and social conditions, and as a result, “applied ontology” was never given a proper chance to prove itself back in that setting then.

vices available for business and government clients. He also started providing it as a pro bono service to the ontology community.

The CIM.NET architecture was inspired by Doug Engelbart's vision on how humanity may be able to tackle its urgent and complex problems through the "Bootstrap" strategy (Engelbart et al., 1999). Engelbart argues for an infrastructure that can foster the co-evolution of the capabilities of humans and those of tools. In a networked improvement community (NIC), one might boost the performance and Collective IQ of the organizations involved (Engelbart, 1992). By improving on the ability to improve, the bootstrap approach allows the human-machine system to improve exponentially, and hence allow humanity to cope more effectively with the unavoidable exponential growth in the changes and their complexities that are happening around them. Engelbart's quest and approach has remained a key influence on the author's work. Besides those on the CIM3 [cwe-dev] team, colleagues that are inspired by Engelbart's vision include folks such as Alan Kay, Terry Winograd, Tom Gruber, Stefan Decker, and Adam Cheyer. In particular, Engelbart encouraged everyone to identify an "improvement vector" he or she deems most important and to work on it passionately with the bootstrap approach. Sometime in the year 2000, the author decided to choose "ontology" as his improvement vector. With ontology making intended meaning explicit and computable, it will facilitate mutual (machine-machine, machine-human, human-machine-machine-human, and even human-human) "understanding" in a human-machine system. The author concludes that this makes *ontology* a most promising piece in the "capability infrastructure" that might facilitate the co-evolution of human capabilities and tools capabilities (i.e., the co-evolution of humans and technology). That choice motivated what followed in his pursuit to help advance ontology (in whatever way he can contribute) in the ensuing 15 years. During that period, the author leveraged the CIM3.NET infrastructure, as well as his own technical, operations and management skills to support ontology-related communities²³ and helped them to advance their causes.

4. Where are we heading?

Whether or not people are consciously aware that the "bootstrap" strategy is at play, open communities collaborating in cyberspace have ended up performing (at least, anecdotally) more effectively than those operating in more traditional (non-open, non-networked, non-co-evolutionary) settings. The phenomena brought about by integrated circuits, or by the Internet, are cases in point. The author expects the improvements that will be brought about by applied ontology in the future to be similarly revolutionary. The "Ontology Summit 2011 Communiqué: Making the Case for Ontology" (Uschold, 2011) details some of the value propositions of applied ontology, and describes how a member of the ontology community can make the case for ontology to a wider audience.

In the last few years, *ontology* has had its share of success stories – in projects such as Apple's acquisition of Siri²⁴ (a natural language understanding, ontology-based mobile app) for hundreds of millions of dollars (April 2010); the IBM Watson²⁵ computer prevailing over its human counterparts on National

²³Ontology-related communities and projects that were hosted/supported on CIM3.NET over the years included UBL, ONTOLOG, OOR, OOR-IPR, Protégé-wiki and old-release repository, CODS, SUMO, OKMDS, DAO, IAOA, CLv2, OntoIOP, SOCoP, OntologySummit, OntologyBasedStandards, UoM, SemanticWiki, RulesReasoningLP, EarthScienceOntolog, DAO; as well as more knowledge/collaboration driven ones, such as eGov-CWE, NIST-Interop, KMatKent, Bootstrap, CHM, MP, SOFI, etc. More details on some of these work are available under: <http://ontolog.cim3.net/wiki/WikiHomePage.html#nidF>.

²⁴See: <http://www.apple.com/ios/siri/> & http://ontolog.cim3.net/wiki/ConferenceCall_2010_02_25.html.

²⁵See: <http://www.ibm.com/smarterplanet/us/en/ibmwatson/>; http://ontolog.cim3.net/wiki/ConferenceCall_2014_01_30.html#nid455F & http://ontolog.cim3.net/wiki/ConferenceCall_2006_05_11.html.

TV in the game of “Jeopardy!” (February 2011); the announcement that Google, Microsoft, Yahoo, and Yandex are teaming up to develop [Schema.org](https://schema.org/)²⁶ (June 2011), and the launch of Google Knowledge Graph²⁷ (May 2012) (Singhal, 2012).

Over the last couple of years, we have also been seeing headlines such as:

- “IBM announced plans to invest \$1 billion to begin a separate business unit around its super-computer Watson, with \$100 million earmarked for venture investment”.
- “Artificial Intelligence Startups See 302% Funding Jump in 2014!”
- “AI Startup Sentient Technologies emerged with \$103.5 million in a new series C round of funding”.
- “Google Acquires London-based AI Startup, DeepMind, for more than \$400 Million”.
- “Car companies are scrambling for AI talent as Google and Apple’s driverless cars loom”.
- “Toyota plans to spend at least \$1 billion on a Silicon Valley research center to study autonomous driving and robotics”.

It is not just the individual headline (though every single one of them is significant) but the trend established from these news items collectively that is worth noting. Artificial intelligence, after 60 years of development, is being given the attention it deserves, finally. While applied ontology is only one sub-discipline within the broad scope of artificial intelligence, we know, though, at the end of the day, that all the AI technologies – whether they be machine learning, natural language processing, knowledge representation, reasoning, or planning – will end up finding their way into the hybrid solutions needed for our typical “complex” real live problems. Applied ontology will definitely be there and it will play key roles in some of the solutions.

The trend is unmistakable: Collectively, we have been doing something right. Applied ontology has turned the corner and has a very exciting time ahead. The good applied ontology practitioners are now presented with an unprecedented availability of funding and investments, as well as career opportunities. These are the kinds of opportunities that we weren’t able to see just, say, six years ago.

Can the ontology community leverage what it has built in the last 20–25 years and capitalize on these available opportunities to deliver a meaningful and lasting impact? The author believes the answer is a resounding “Yes!” and that eventuality and impact will be inevitable, although the timing may still depend on how effectively we can advance things.

We might be able to consciously speed things up, allowing applied ontology to make a stronger impact, by addressing some of the challenges that face us today:

- *Diversify the membership of our community of practice* to cover the full range of stakeholders. We need more than just researchers and theoreticians. Let’s bring in everyone else who can fill the capability gaps. We need to consider “the practice” as a whole. We need to strengthen our outreach, and to start collaborating with other segments of the AI space. We want more exposure to, and collaboration with, solution partners (the system analysts and architects, data modelers, and software engineers). We want to invite others into the CoP, and to create exposure of applied ontology to wider (but related) professional services people, corporate decision makers, and policy makers, too.
- *Expand our penetration into education* for the typical university graduate today (and also, of course, for the younger students). These trainees are probably still unaware of what “an ontology” or “applied ontology” is all about. Our influence on education still has much room to improve. The corol-

²⁶See: <https://schema.org/>.

²⁷See: https://www.google.com/intl/en_en/insidesearch/features/search/knowledge.html.

lary to this issue is that if we don't resolve this problem, we will not have enough ontologists to meet our needs, even if the opportunities are right in front of us.

- *Keep our IPR open.* Ontology is so fundamental, that allowing private enterprises to monopolize the wrong elements of the Intellectual Properties Rights (IPR) will eventually hamper the future for everyone. Good lessons can be learned from how people have been trying to fight for “net neutrality” and appropriate software patents. We need to do even better.
- Lastly, we should *continue “bootstrapping”* by leveraging the ontology and ontology-engineering approaches to improve on how ontology communities and their members can better collaborate and advance the field.

In short, the opportunities we have been waiting for are finally here. Let us keep up the good work, and make applied ontology deliver.

Acknowledgements

We are indebted to Doug Engelbart and John McCarthy for showing us the way. The author would like to acknowledge collaborators in the various ontology communities and projects he has had the privilege to work with over the years. Last but not least, thanks are due to other [cwe-dev] team members. In particular, major contributors to the CIM3.NET infrastructure design and implementation were due to Jonathan Cheyer (who architected and designed the IT infrastructure), Shinya Yamada (who maintained the cloud-based system, and transitioned it to a virtual machine infrastructure), Eugene Kim (who introduced us to the wiki, and implemented purple-numbers to provide the fine-grained addressability in the CWE, called for by Engelbart's design), as well as Tejas Parikh, Ken Baclawski, Jonathan Cheyer and others (who implemented the purple-mediawiki-extension, and transitioned the Ontolog wiki-content to the Purple Semantic MediaWiki (PSMW) platform).

References

- Arp, R., Smith, B. & Spear, A.D. (2015). *Building Ontologies with Basic Formal Ontology*. MIT Press.
- Bateman, J.A., Magnini, B. & Fabris, G. (1995). The generalized upper model knowledge base: Organization and use. *Towards Very Large Knowledge Bases* (pp. 60–72). Amsterdam: IOS Press.
- Batres, R., West, M., Leal, D., Price, D., Masaki, K., Shimada, Y., Fuchino, T. & Naka, Y. (2007). An upper ontology based on ISO 15926. *Computers & Chemical Engineering*, 31(5), 519–534.
- Brown, J. & Gray, E. (1995). The people are the company. Available at: <http://www.fastcompany.com/26238/people-are-company>. Accessed: 25 November 2015.
- Conrad, K., Obrst, L. & Yim, P. (2002). Ontolog invitation letter. Available at: <http://ontolog.cim3.net/forum/ontolog-forum/2002-09/msg00000.html>.
- Engelbart, D. (1992). Toward high-performance organizations: A strategic role for groupware (AUGMENT,132811). Available at: <http://www.dougenelbart.org/pubs/augment-132811.html>. Accessed: 25 November 2015.
- Engelbart, D., et al. (1999). The bootstrap vision and mission statements. Available at: http://bootstrap.cim3.net/vision_mission.html. Accessed: 25 November 2015.
- Engelbart, D., et al. (2000). The Doug Engelbart colloquium at Stanford: An in-depth look at “The Unfinished Revolution”. (UnRev-II). Available at: <http://dougenelbart.org/colloquium/sessions.html>. Accessed: 25 November 2015.
- Gangemi, A. & Mika, P. (2003). Understanding the Semantic Web through descriptions and situations. In *On the Move to Meaningful Internet Systems 2003: CoopIS, DOA, and ODBASE* (pp. 689–706). Berlin: Springer.
- Gruber, T. (1993). Toward principles for the design of ontologies used for knowledge sharing. Available at: <http://tomgruber.org/writing/onto-design.pdf>. Accessed: 25 November 2015.
- Gruninger, M., Bodenreider, O., Olken, F., Obrst, L. & Yim, P. (2008). Ontology Summit 2007 – Ontology, taxonomy, folksonomy: Understanding the distinctions. *Applied Ontology*, 3(3), 191–200.

- Gruninger, M. & Menzel, C. (2003). The Process Specification Language (PSL) theory and applications. *AI Magazine*, 24(3), 63–74.
- Guarino, N. (1998). *Formal Ontology in Information Systems: Proceedings of the First International Conference (FOIS'98)*, Trento, Italy, June 6–8. Amsterdam: IOS Press.
- Guarino, N. & Musen, M.A. (2015). Applied ontology: The next decade begins. *Applied Ontology*, 10(1), 1–4.
- Guarino, N. & Poli, R. (1993). International workshop on formal ontology in conceptual analysis and knowledge representation. Preface and table of contents from the proceedings. Available at: http://ontolog.cim3.net/file/resource/historic-archives/FOIS-community/International-FormalOntology-Workshop_Padova-It_1993-03.pdf. Accessed: 25 November 2015.
- Guarino, N., Yim, P., Galton, A., et al. (2009). The international association for ontology and its applications (IAOA) – Ontolog virtual panel presentation. Available at: http://ontolog.cim3.net/wiki/ConferenceCall_2009_06_18.html and, in particular, slide #4 on http://ontolog.cim3.net/file/resource/presentation/IAOA_20090618/IAOA-presentation_20090618.pdf. Accessed: 25 November 2015.
- Lenat, D.B. & Guha, R.V. (1989). *Building Large Knowledge-Based Systems; Representation and Inference in the Cyc Project*. Addison-Wesley Longman Publishing Co., Inc.
- Masolo, C., Borgo, S., Gangemi, A., Guarino, N. & Oltramari, A. (2003). WonderWeb deliverable D18. Ontology library (final). ICT project, 33052.
- Musen, M.A. (1989). *Automated Generation of Model-Based Knowledge Acquisition Tools*. London: Pitman.
- Musen, M.A. (2015). The Protégé project: A look back and a look forward. *AI Matters*, 1(4), 4–12.
- Neches, R., Gruber, T., et al. (1993). The knowledge sharing effort. Available at: <http://www-ksl.stanford.edu/knowledge-sharing/papers/kse-overview.html>. Accessed: 25 November 2015.
- Neuhaus, F., Florescu, E., Galton, A., Gruninger, M., Guarino, N., Obrst, L., Sanchez, A., Vizedom, A., Yim, P. & Smith, B. (2011). Creating the ontologists of the future. *Applied Ontology*, 6(1), 91–98.
- Neuhaus, F., Vizedom, A., Baclawski, K., Bennett, M., Dean, M., Denny, M. & Yim, P. (2013). Towards ontology evaluation across the life cycle. *Applied Ontology*, 8(3), 179–194.
- Niles, I. & Pease, A. (2001). Towards a standard upper ontology. In *Proceedings of the International Conference on Formal Ontology in Information Systems* (Vol. 2001, pp. 2–9). New York: ACM.
- Obrst, L., Cassidy, P., Ray, S., Smith, B., Soergel, D., West, M. & Yim, P. (2006). The 2006 Upper Ontology Summit joint Communiqué. *Applied Ontology*, 1(2), 203–211.
- Obrst, L., Musen, M., et al. (2008). The Ontology Summit 2008 Communiqué: Towards an Open Ontology Repository. Available at: http://ontolog.cim3.net/wiki/OntologySummit2008_Communique.html. Accessed: 25 November 2015.
- Obrst, L., Gruninger, M., Baclawski, K., Bennett, M., Brickley, D., Berg-Cross, G. & Yim, P. (2014). Semantic Web and Big Data meets Applied Ontology: The Ontology Summit 2014. *Applied Ontology*, 9(2), 155–170.
- Ray, S., Gruninger, M., Mason, H., West, M., et al. (2009) The Ontology Summit 2008 Communiqué: Toward Ontology-based Standards. Available at: http://ontolog.cim3.net/wiki/OntologySummit2009_Communique.html. Accessed: 25 November 2015.
- Schneider, T., Hashemi, A., Bennett, M., Brady, M., Casanave, C., Graves, H. & Yim, P. (2012). Ontology for Big Systems: The Ontology Summit 2012 Communiqué. *Applied Ontology*, 7(3), 357–371.
- Schoening, J., Yim, P., et al. (2015). Collection of artifacts relating to the “IEEE P1600.1: Standard upper ontology working group (SUO WG) effort” which included the author’s notes, a brief first-person account by James Schoening, and the 2001 IJCAI Paper “Origins of the IEEE standard upper ontology” by Ian Niles and Adam Pease. Available at: <http://ontolog.cim3.net/file/resource/historic-archives/IEEE-SUO-WG/>. Accessed: 25 November 2015.
- Singhal, A. (2012). Introducing the knowledge graph: Things, not strings. *Official blog* (of Google). Available at: <https://googleblog.blogspot.co.uk/2012/05/introducing-knowledge-graph-things-not.html>. Accessed: 25 November 2015.
- Smith, B., Ashbner, M., Lewis, S., et al. (2007). The OBO foundry: Coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*, 25, 1251–1255.
- Spillers, R., Hayes, P., Sowa, J., et al. (2015). Collection of artifacts relating to the “1998 Heidelberg Workshop” which included the author’s notes, Pat Hayes’ collection of the related email messages exchanged and a brief first-person account about the event by John Sowa. Available at: http://ontolog.cim3.net/file/resource/historic-archives/Heidelberg-workshop_1998/. Accessed: 25 November 2015.
- Underwood, M., Gruninger, M., Obrst, L., et al. (2015). Internet of things: Toward smart networked systems and societies. The Ontology Summit 2015 Communiqué. *Applied Ontology*, 10(3,4), 355–365.
- Uschold, M. (2011). Making the case for ontology. *Applied Ontology*, 6(4), 377–385.
- Winograd, T., Newman, N. & Yim, P. (1991). Including people in computer integrated manufacturing design. In *Proceedings of the International Conference on Computer Integrated Manufacturing*, Singapore, October 2–4, 1991. Available at: <http://www.cim3.com/whitepaper/lapcim1a.htm>. Accessed: 25 November 2015.
- Yim, P. (1991). CIM3: Computer integrated manufacturing systems: An introduction. In *Proceedings of the International Conference on Computer Integrated Manufacturing*, Singapore, October 2–4, 1991. Available at: http://www.cim3.com/whitepaper/cim3_1a1.htm. Accessed: 25 November 2015.

Yim, P. (2009). The CIM3.NET collaborative work environment. Available at: <http://www.cim3.com/CIM3-CWE.html>. Accessed: 25 November 2015.

Yim, P., Ray, S., Obrst, L., Cassidy, P., et al. (2015). The Ontology Summit program. Available at: <http://ontolog.cim3.net/wiki/OntologySummit.html>. Accessed: 25 November 2015.