



U.S. DEPARTMENT OF
ENERGY

PNNL-21072

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Nuclear Nonproliferation Ontology Assessment Team – Final Report

JD Strasburg
RE Hohimer

January 2012



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.** Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <<http://www.ntis.gov/about/form.aspx>>
Online ordering: <http://www.ntis.gov>



This document was printed on recycled paper.

(8/2010)

Nuclear Nonproliferation Ontology Assessment Team – Final Report

JD Strasburg
RE Hohimer

January 2012

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Summary

The Ontology Assessment Team began in May 2009 and concluded in September 2011. During this two-year time frame, the Ontology Assessment team had two objectives:

- Assessing the utility of knowledge representation and semantic technologies for addressing nuclear nonproliferation challenges;
- Developing ontological support tools that would provide a framework for integrating across the Simulation, Algorithm and Modeling (SAM) program.

The SAM Program was going through a large assessment and strategic planning effort during this time and as a result, the relative importance of these two objectives changed, altering the focus of the Ontology Assessment Team. In the end, the team conducted an assessment of the state of art, created an annotated bibliography, and developed a series of ontological support tools, demonstrations and presentations.

A total of more than 35 individuals from 12 different research institutions participated in the Ontology Assessment Team. These included subject matter experts in several nuclear nonproliferation-related domains as well as experts in semantic technologies. Despite the diverse backgrounds and perspectives, the Ontology Assessment team functioned very well together and aspects could serve as a model for future inter-laboratory collaborations and working groups.

While the team encountered several challenges and learned many lessons along the way, the Ontology Assessment effort was ultimately a success that led to several multi-lab research projects and opened up a new area of scientific exploration within the Office of Nuclear Nonproliferation and Verification.

Acronyms and Abbreviations

AS	Advanced Spectroscopic Analysis
ASV	Advanced Spectroscopic Analysis Validation
DOE	U.S. Department of Energy
FM	Facility Modeling
FMV	Facility Modeling Validation
GS	Geospatial Analysis
GSV	Geospatial Analysis Validation
IM	Integrated Modeling
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
NFC	Nuclear Fuel Cycle
ORNL	Oak Ridge National Laboratory
PNNL	Pacific Northwest National Laboratory
SAM	Simulation, Algorithm and Modeling
SME	Subject Matter Expert
SNM DO	Special Nuclear Materials Detection Ontology
SNM	Special Nuclear Materials
STL	Special Technologies Laboratory
V&V	Verification and Validation

Contents

Summary	iii
Acronyms and Abbreviations	v
1.0 Introduction	1.1
2.0 Project Organization	2.1
3.0 Technology Assessment	3.1
3.1 Description of the State of the Art	3.1
3.1.1 Top Down verses Bottom Up.....	3.1
3.1.2 Available Ontologies.....	3.1
3.2 Identification of Research Needs	3.2
3.2.1 Data Fusion and Integration	3.2
3.2.2 Common Vocabulary	3.2
3.3 Annotated Bibliography	3.2
4.0 Potential Ontology Applications within NNSA	4.1
4.1 Data Fusion and Data Integration.....	4.1
4.2 Model Integration.....	4.1
4.2.1 Data Integration versus Model Integration.....	4.1
4.3 Other Areas of Interest	4.2
4.3.1 Common Language or Translations	4.2
4.3.2 Information Search and Retrieval.....	4.2
4.3.3 Semantic Framework for Understanding the Nuclear Proliferation Problem	4.2
5.0 Ontological Support Tools.....	5.1
5.1 Text Extraction Tools.....	5.1
5.2 Lexical Analysis Tools.....	5.1
5.3 Collaborative Environments.....	5.1
5.3.1 SAM SharePoint.....	5.1
5.3.2 SAM ChipFabrication Wiki	5.1
5.3.3 SAM Portal.....	5.2
5.4 Ontologies/Taxonomies/Demonstrations	5.2
5.4.1 Ontology Development with the Assessment Teams.....	5.2
5.4.2 Ontologies	5.2
5.4.3 Demonstrations.....	5.5
6.0 Project Outcomes.....	6.1
6.1 Follow-on Projects	6.1
6.1.1 Scenario Generation.....	6.1
6.1.2 Small Business and University Research Projects	6.1
6.1.3 Nuclear Fuel Cycle Reasoning.....	6.1

6.2	DOE Presence in the Field Established through STIDS 2011	6.1
6.2.1	Use of Ontology to Facilitate the Creation of Synthetic Imagery of Industrial Facilities	6.2
6.2.2	Ontology-based Software for Generating Scenarios for Characterizing Search for Nuclear Materials	6.2
6.2.3	Nuclear Fuel Cycle Reasoning	6.2
6.2.4	A Framework for Ontology-Supported Intelligent Geospatial Feature Discovery Services	6.2
7.0	Lessons Learned	7.1
7.1	Project Management.....	7.1
7.2	Use-Case Focus.....	7.1
7.3	Collaborative Ontology Development.....	7.1
7.4	Knowledge Elicitation.....	7.2
8.0	Conclusions	8.1
	Appendix A – Annotated Bibliography	A.1

Figures

5.1	Ontology Development Process.....	5.3
5.2	Ontology Development Workflow Within the Ontology Assessment Team.....	5.4
5.3	Visualization of a Portion of One Ontology	5.5

Tables

2.1	Ontology Assessment Team Members	2.1
-----	--	-----

1.0 Introduction

In the spring of 2009, the Office of Nuclear Nonproliferation and Verification's (NA22) Simulation, Algorithm and Modeling (SAM) program embarked on a program-wide assessment process. Seven different assessment teams were formed to explore the state of the art in various research areas: Advanced Spectroscopic Analysis (AS), Geospatial Analysis (GS), Facility Modeling (FM), Advanced Spectroscopic Analysis Validation (ASV), Geospatial Analysis Validation (GSV), Facility Modeling Validation (FMV) and Integrated Modeling (IM). Each of these seven areas was led by a different laboratory and was made up of team members from across the U.S. Department of Energy (DOE) laboratory complex. That summer another team was formed that was focused on ontologies, semantic technologies, and knowledge representation systems. In the framework of the SAM program at the time, ontologies would serve as the connection between the seven different SAM assessment teams and the broader mission of the Office.

During the next two years, the initial seven assessment teams worked to assess the state of the art and develop a gap analysis that could drive future research and development investments. In addition, each team would develop an annotated bibliography of references. Unfortunately, the initial objectives of the Ontology Assessment team were not as clearly defined. The manner in which the initial effort was funded and the accompanying aggressive timeline put the multi-laboratory team in a reactive, as opposed to proactive, state that continued to be a challenge throughout the life of the project. The Ontology Assessment team ended up with two objectives:

- Assess the utility of knowledge representation and semantic technologies for addressing nuclear nonproliferation challenges
- Develop ontological support tools that would provide a framework for integrating across the SAM program.

The relative importance of these two objectives repeatedly changed depending on the status of the other assessment efforts and the direction of the SAM program manager. This dynamic work scope introduced several challenges and reduced the efficiency of progress in any one of the areas.

By the end of the roughly two-year effort, the Ontology Assessment team conducted an assessment of the state of art, identified potential applications within NA22, created an annotated bibliography and developed a series of ontological support tools, demonstrations and presentations. The Ontology Assessment team effort directly led to several follow-on projects and established a DOE/NNSA presence in the area of semantic technology research. Along the way, the team encountered and overcame challenges and learned several lessons.

This document begins by acknowledging the Ontology Assessment team members and the roles they played on this project. This is followed by an assessment of the technical challenges in semantic technologies, a listing of potential applications within NA22 and a description of the various ontological support tools developed under this effort. The report ends with outcomes, lessons learned and conclusions. The annotated bibliography is included as an Appendix.

2.0 Project Organization

During the course of the two years, more than 35 individuals from 12 different research institutions participated in the Ontology Assessment Team. These included subject matter experts in several nuclear-nonproliferation-related domains as well as experts in semantic technologies. Table 2.1 lists team members, their respective laboratories and their primary roles on the Ontology Assessment team. The core contributors were actively involved throughout the life of the project and were responsible for different tasks; other team members participated on an as-needed basis.

Table 2.1. Ontology Assessment Team Members. Names in bold represent the core contributors.

Name	Laboratory	Role
Bruce Bargmeyer	LBNL	Knowledge Engineer
Robert Bean	INL	Subject Matter Expert
Karen Bills	Y-12	Subject Matter Expert
Craig Blackhart	LANL	Knowledge Engineer
Jim Bollinger	SRNL	Subject Matter Expert
Steven Bradley	LLNL	Knowledge Engineer
John Christiansen	ANL	Knowledge Engineer
Helen Cui	LANL	Knowledge Engineer
Carl Czajkowski	BNL	Subject Matter Expert
George Davidson	SNL	Knowledge Engineer
Regina Ferrel	ORNL	Subject Matter Expert
David Gerts	INL	Subject Matter Expert
Vinita Ghosh	BNL	Subject Matter Expert
Max Gorenssek	SRNL	Subject Matter Expert
Larry Hamm	SRNL	Subject Matter Expert
Alex Heifetz	ANL	Subject Matter Expert
Ryan Hohimer	PNNL	Knowledge Engineer
Cliff Joslyn	PNNL	Knowledge Engineer
Kevin Keck	LBNL	Knowledge Engineer
Jim Kornell	NSTEC/STL	Knowledge Engineer
John Krebs	ANL	Subject Matter Expert
Kara Kruse	ORNL	Subject Matter Expert
Glenn May	LLNL	Knowledge Engineer
Liam McGrath	PNNL	Knowledge Engineer
Laura Musgrave	LANL	Subject Matter Expert
James Nutaro	ORNL	Subject Matter Expert
Kelly O'Hara	PNNL	Knowledge Engineer
Paul Pope	LANL	Subject Matter Expert
Oleg Roderick	ANL	Subject Matter Expert
Dave Schoenwald	SNL	Subject Matter Expert
Alex Slepoy	NA22	SAM Program Manager
Alex Sorokine	ORNL	Knowledge Engineer
Jana Strasburg	PNNL	Project Lead
Sandy Thompson	NA22	SAM Technical Advisor
Tim Voss	LLNL	Knowledge Engineer
Richard Ward	ORNL	Subject Matter Expert

The Ontology Assessment team organized itself around the seven other SAM assessment teams. Four sub-groups were formed addressing the focus of one or more of the seven other teams. Members of the four groups were a combination of Subject Matter Experts (SMEs) familiar with the assessment team's topic, and knowledge engineers familiar with ontologies. The four groups were referred to as the Ontology SME teams. Members of the Ontology SME teams interfaced with the other SAM assessment teams and served as ambassadors and educators for and about semantic technologies. There was also an Ontology Development sub-team that was made up of knowledge engineers and focused on developing ontological support tools.

The full team held conference calls on a periodic basis ranging from every few weeks to once a month. There were also opportunities for five face-to-face meetings—two in Washington, D.C., one in Seattle, one at Brookhaven National Lab, and one at Lawrence Berkeley National Lab.

Despite the diverse backgrounds and perspectives, the Ontology Assessment team functioned very well together and could serve as a model for inter-laboratory collaborations and working groups. During the course of the two years, team members gained a new respect for each other's capabilities and a better understanding of the strengths of different labs. The labs often have a reputation for territorial behavior and in-fighting amongst themselves, but this was not the case on the Ontology Assessment team. Several inter-laboratory relationships and collaborations grew out of the Ontology Assessment team and still continue today.

3.0 Technology Assessment

The Ontology Assessment team used literature reviews, discussions, and collaboration to assess the technology and develop a description of the state of the art.

3.1 Description of the State of the Art

3.1.1 Top Down verses Bottom Up

In *A Connecticut Yankee in King Arthur's Court*, Mark Twain penned: “she was wise, subtle, and knew more than one way to skin a cat.” This particular adage seems very fitting in disclosing that very wise ontologists have multiple approaches to solving problems. Over-simplifying, there are two major approaches: 1) top-down design, and 2) bottom-up design.

This debate often played out amongst members of the Ontology Assessment team and reflects the ongoing philosophical debate between these two approaches in the broader community of semantic technologists.

3.1.1.1 Top Down

In this approach, the knowledge engineer makes use of an upper ontology (top-level ontology) which defines the general concepts that are the same in all knowledge domains. The premise here is that the upper ontology will support the interoperability between the various specific domain ontologies used.

To date, there is no one upper ontology considered or accepted as the de facto standard. Different organizations are making an effort to define standards in specific domains.

3.1.1.2 Bottom Up

The bottom-up supporters typically approach problems by asking what specific purpose is the ontology to fulfill, and what specifically is the ontology to do for us to make a difference in our problem space.

This approach generally starts by defining the classes of “primitives” in the data of the domain of discourse. This builds the ontology one concept brick at a time from the bottom foundational concepts up. In other words, the approach is to define the minimal number of concepts in the ontology to solve the defined problem.

3.1.2 Available Ontologies

Whether one takes a top-down or bottom-up approach, there is no shortage of existing ontologies to consider for inclusion into the solution. However, there is a broad range of quality in the selection.

During the span of the Ontology Assessment Team’s research we found ontologies that were not well formed, and some that were. The team even experienced some ontology resources that transitioned from “horribly” formed to “reasonable” over the duration of our investigations – for example, the SWEET ontologies from NASA.

3.2 Identification of Research Needs

The identification of research needs within NA22 turned out to be relatively straightforward. Likely, this is due to the fact that NA22 faces the same problems that challenge so many other enterprises. While considering potential application spaces for semantic technologies within NA22, two topics emerged:

- 1) Data Fusion and Integration;
- 2) Common Vocabulary.

3.2.1 Data Fusion and Integration

The origins of the Ontology Assessment Team were rooted in SAM's Integrated Modeling Team. With multiple models being considered and developed, questions of interoperability arose. This led to the discussion of ontologies for the purpose of data integration between models.

Interoperability between domains is an admirable goal. Ontologies are a means of reaching toward this goal. As an example of interoperability, consider a scenario where a nation state has signed an agreement. As a party to that agreement, what policies have they agreed to, what metrics are available to monitor compliance, and what confirmational information is available to support the metrics?

Ambiguities are also abundant in the area of dual-use technology. Ontological definitions that clarify the relationship between a concept relevant to legitimate activities and that same concept relevant to proliferation activities are valuable. As an example of this, consider the interpretation of a manifest listing high-speed valves used in the semiconductor industry. Knowing that this particular high-speed valve can be used in an enrichment facility is important information.

The output of one model can also serve as the input to another. In order for this to occur, the data models of the two models must be aligned. In simpler terms, if the two models are using the same ontology they can pass data between each other.

3.2.2 Common Vocabulary

Natural language is extremely expressive. However, with that expressiveness, an increase in ambiguity arises. Using formal ontological language, more structure is brought to bear on the definitions of concepts, thereby reducing the ambiguity. In fact, the ambiguity can be decreased to the point where data can be computationally interpreted via the ontological definitions.

This means that by agreeing to use ontological/semantic technologies to represent our data, program managers, researchers, product users, and computers can communicate with less ambiguity.

3.3 Annotated Bibliography

Team members identified and compiled several useful papers during the course of the assessment process. The Annotated Bibliography is included as an Appendix to this document.

4.0 Potential Ontology Applications within NNSA

Members of the Ontology Assessment team concluded that there are most definitely NA22 needs and challenges that can be addressed using ontologies and other semantic technologies.

4.1 Data Fusion and Data Integration

Current state of the art in the semantic technologies is particularly good at addressing data fusion and data integration problems. This is a matter that is not always easy, but straightforward. By communications with SMEs crossing multiple domain spaces, ontologies can be generated that align their data models. The use of “namespaces” to allow each SME the autonomy to define concepts and properties exactly as desired within their own namespace proved to be a powerful tool. Once the ontologies are formed, mappings between separate namespaces can be used to bridge the ambiguities.

The results of the Ontology Assessment team show that for data integration, data fusion, and common language support, ontologies and other semantic technologies can bear fruit. However, it is noted that for quite disparate domains, for example the meaning of treaty and contract clauses in relation to advanced spectroscopy results, the relationships between concepts can be convoluted and complex.

Although working in a surrogate domain space allows for capability development with lower security clearance concerns, it does distance the research products from obvious relevance to the proliferation-detection domain space.

4.2 Model Integration

When talking about “model integration” the level of difficulty rises. One would think that the problem of model integration is the same problem of data integration; however, model integration introduces new challenges.

4.2.1 Data Integration versus Model Integration

Tersely stated, data integration is about defining data types at a relatively primitive level. In contrast, model integration is about defining data types that are more abstract in nature. As an example, consider an object-oriented program that has within its code integers, floats, strings, and Boolean values. This same program can have classes defined that are much more abstract in nature, such as staff, contracts, agreements, bills of lading, spectrograph, signature, and dual use item.

When trying to integrate data from different namespaces, the closer the data is to primitive form, the easier the ontologies are to align. The namespaces may define different units, for instance, but the integration problem can be addressed with conversion factors. However, if the data are more abstract, the ontologies can be very difficult to align. There may or may not be an overlap in the meaning of abstract concepts from two different namespaces even if the lexical overlap is great.

4.3 Other Areas of Interest

4.3.1 Common Language or Translations

Within NA22 there are many disciplines practiced in many domains. Each discipline/domain has its own vocabulary. Even within the same discipline, there exist different dialects. Take for example the concept of a “signature.” “Signature” is a concept that is overloaded, meaning that the term has many meanings, varying with the context within which it is used. NA22 would benefit from the disambiguation of terms.

4.3.2 Information Search and Retrieval

Program managers, researchers, analysts, and policy makers would benefit from being able to search for existing capabilities and previous efforts within the scope of the NA22 programs. The ability to do this would assist in many of the decision-making processes: for example, answering questions such as, “Does a certain capability exist, and if so, what resources have been invested in the capability?”

4.3.3 Semantic Framework for Understanding the Nuclear Proliferation Problem

The nuclear nonproliferation communities have been constructing an understanding of the nuclear fuel cycle and weaponization processes for a significant amount of time. However, the constructs have not been represented in formal ontological languages. Primarily, this is due to the fact that the formal languages did not exist at the time of concept formation.

5.0 Ontological Support Tools

In addition to the technology assessment, the Ontology Assessment team developed and utilized several ontological support tools over the course of the two-year effort.

It should be noted that in order to make unhindered progress toward assessing ontologies and semantic technologies for use in the nuclear proliferation domain, the team used the surrogate domain of semiconductor manufacturing. The semiconductor domain was suggested as an analogy to the nuclear proliferation domain and freed the team from potential classification concerns.

The use of a surrogate domain did cause some issues when trying to encode the knowledge elicited from the SAM Assessment Teams. Knowledge engineers interpreted the analogous concepts from the nuclear proliferation domain and attempted to represent them in the semiconductor domain.

5.1 Text Extraction Tools

Sandia National Laboratory hosted and managed the Chip Wiki. This wiki was used to capture the subject matter expertise of the semiconductor manufacturing domain. Periodically, the web content was crawled and text extraction of key words and concepts was applied. The resulting extractions were used for lexical analysis.

5.2 Lexical Analysis Tools

Lexical analysis, led by Pacific Northwest National Laboratory (PNNL), was done to compare the SME's key words and concepts (extracted from wiki content and collections of documents) with the concepts represented in the current ontologies under development. The analysis allowed the team to assess the alignment and overlap statistics between the SME's concepts and the Ontology concepts.

5.3 Collaborative Environments

Several collaborative environments were utilized or created during the course of this effort.

5.3.1 SAM SharePoint

Oak Ridge National Laboratory (ORNL) hosted a SharePoint site which allowed members of the Ontology Assessment team to collaborate and share ideas. The site also provided the functionality to collect relevant research documents and their references. Additionally, the site provided discussion forums for the geographically dispersed team to discuss the state of the Ontology Assessment team's efforts.

5.3.2 SAM ChipFabrication Wiki

Los Alamos National Laboratory (LANL) hosted a wiki for the Assessment Teams to collect information about the Semiconductor industry. The wiki was a knowledge elicitation tool to capture

subject-matter expertise regarding semiconductor manufacturing processes and the facilities central to those processes.

5.3.3 SAM Portal

Lawrence Berkeley National Laboratory (LBNL) hosted a portal for ontology team members to assess ontology and knowledgebase management tools. The portal provided the repository for ontologies under development by the assessment team members. The SAM Portal also provided a means to access other ontologies developed by the broader community of ontologists.

5.4 Ontologies/Taxonomies/Demonstrations

During the course of the project several ontologies and taxonomies were created. No matter which ontology language is used by the ontologist, the basic steps in ontology creation are the same.

The domain of the ontology needs to be selected. The intended use of the domain ontology should be characterized. The terms identified as being in the domain need to be enumerated. The key concepts need to be defined in classes. The key properties need to be defined. Restrictions on the properties and classes need to be set. The relationships between classes need to be defined using the properties. This process is illustrated in Figure 5.1.

5.4.1 Ontology Development with the Assessment Teams

Within the Ontology Assessment team, the general approach to ontology development was to have the SAM assessment teams compile a corpus of seminal papers and augment the corpus with a summary of each paper and a key word list. These materials from the assessment team would be examined with lexical analysis tools to compare (check the alignment with) the developing ontologies. Based on the results of the lexical analysis, the knowledge engineers in the appropriate Ontology SME team incorporated the changes into the ontology. SMEs in the Ontology SME team would internally vet the changes to the current ontology. In addition to the Ontology SME team, the SAM Assessment team SME would then vet the ontology. This workflow is shown in Figure 5.2.

5.4.2 Ontologies

During the course of this effort, several small ontologies were developed by various team members. Figure 5.3 shows a portion of one of these ontologies. It describes the concepts and relationships between various types of facilities, their components, and different applicable detection strategies.

5.4.2.1 Semiconductor Ontologies

The semiconductor industry turned out to be a very useful surrogate domain for the proliferation domain. This surrogate provided the complexity necessary to realistically model the real domain. It included policy agreements, import and export controls, material flows, processes, and even dual-use examples.

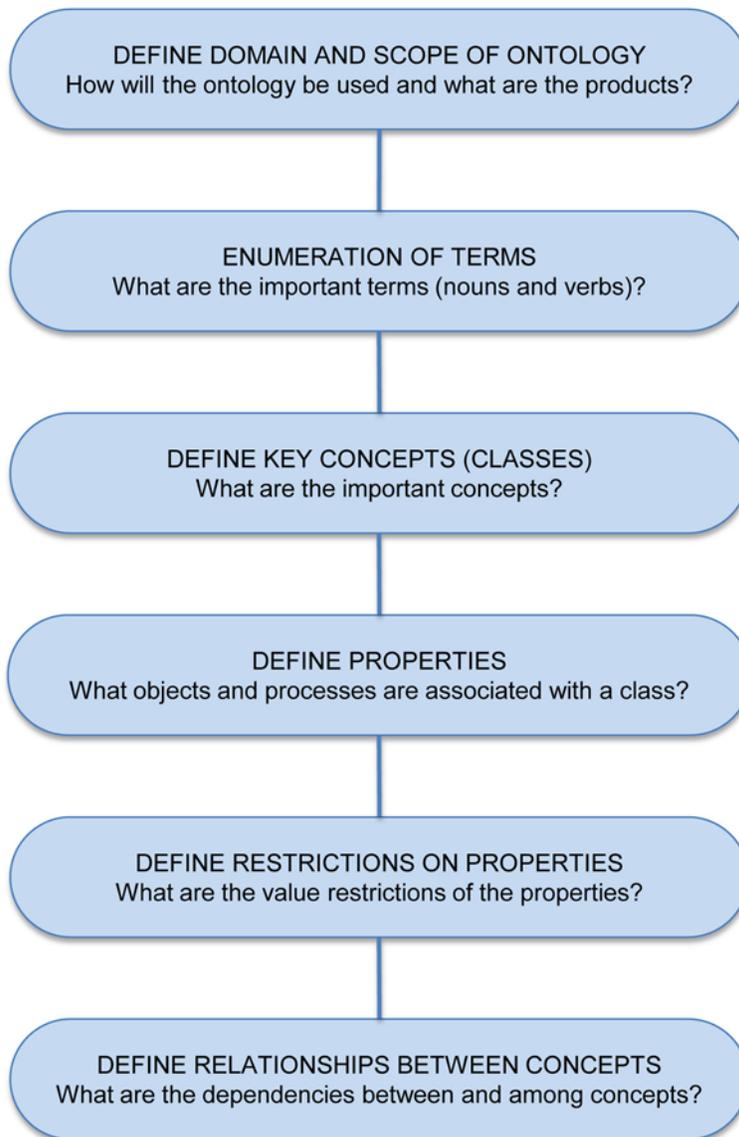


Figure 5.1. Ontology Development Process

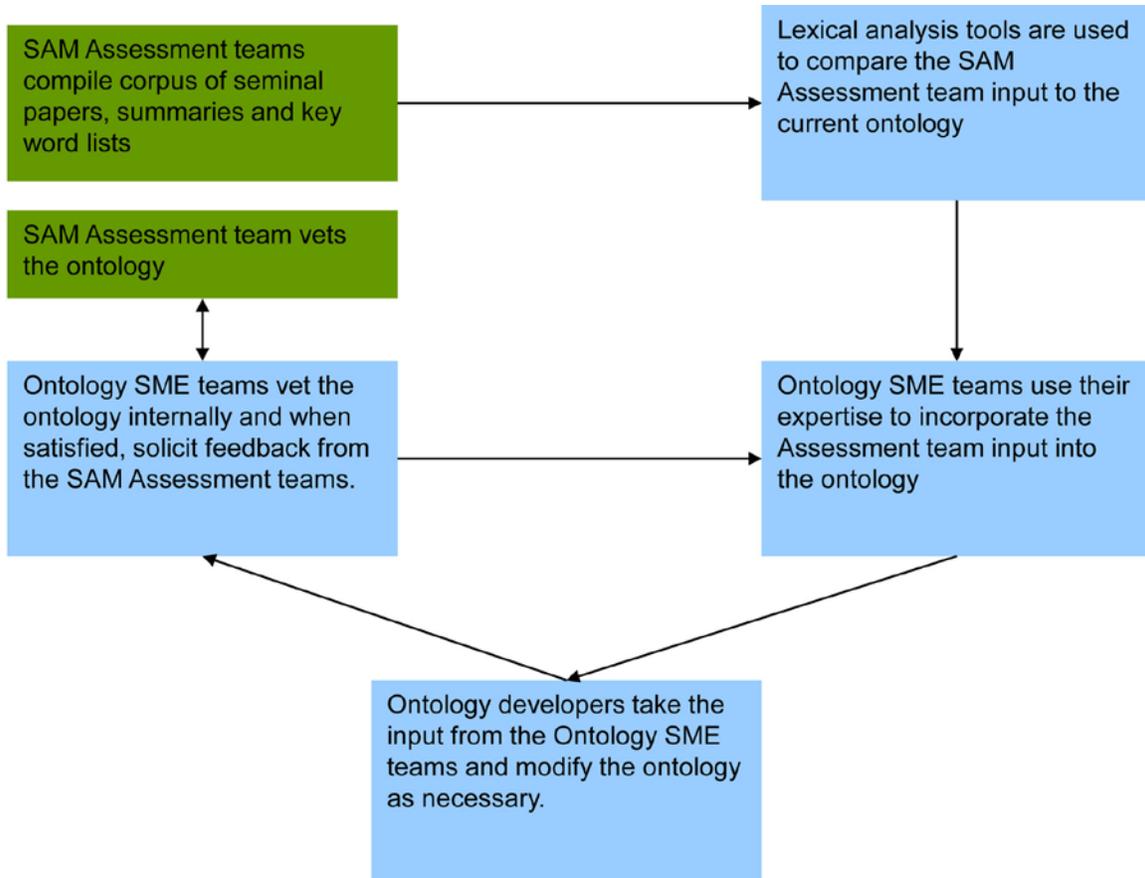


Figure 5.2. Ontology Development Workflow Within the Ontology Assessment Team

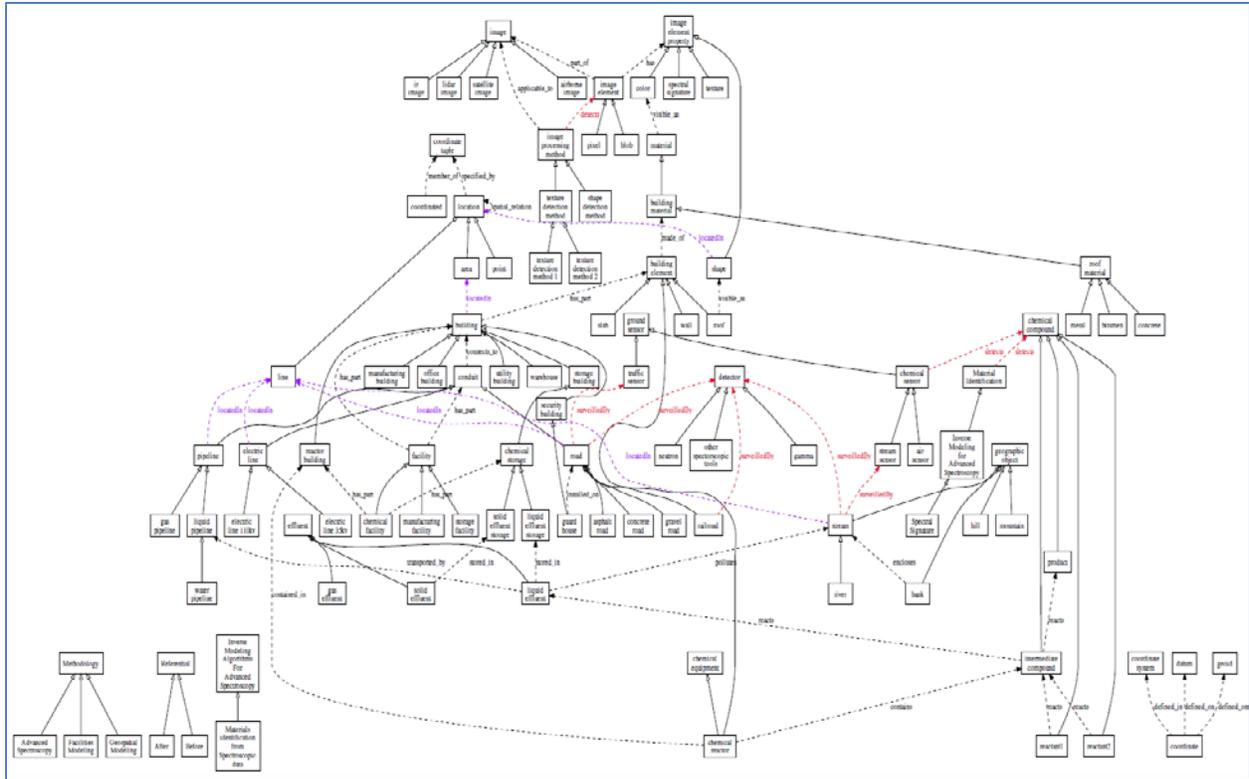


Figure 5.3. Visualization of a Portion of One Ontology

5.4.2.2 Spatial Ontologies

The representation of spatial relationships was investigated. These relationships are important to consider in several areas of interest in the proliferation domain. Image analysis and facilities modeling are two of the domains the ontology team found the most obvious.

5.4.2.3 Process Ontologies

Form does follow function. The form of a facility is a function of the processes occurring around and within it. The processes of the nuclear fuel cycle require a significant infrastructure to support them.

5.4.3 Demonstrations

As a demonstration of how ontologies and knowledgebases could be used to reason about a proliferation scenario a use-case was contrived to provide a specific problem definition. The use-case described a fictitious nation state suspected of planning a secret space-based weapons program using illicitly manufactured semiconductors. This demonstration was shown as part of a project update in June 2010.

Facets of the use-case scenario included intellectual property, export control regulations, special material flows, licensing agreements, facilities modeling, chemical processing, and personnel.

Within the bounds of the use-case, known facts about the scenario were asserted into a knowledgebase. Inferred facts (entailed by the known facts) were examined for feasibility. As an example, if the assertion were made that a chemical process was occurring at a specific location, can a path be inferred along which an airborne plume may carry effluents?

6.0 Project Outcomes

6.1 Follow-on Projects

6.1.1 Scenario Generation

As a direct result of the June 2010 briefing, a specific call for proposals was issued to develop an ontologically-driven scenario generator. Members of the Ontology Assessment team submitted several proposals in response to this call.

A team led by ORNL was awarded the project and used both domain-specific ontologies and knowledgebases containing latent-background information to enable a software environment to generate scenarios from a general set of user input variables. The resulting scenarios could be used for testing algorithms used in the detection of special nuclear materials (SNM). An ontology was developed called the Special Nuclear Materials Detection Ontology (SNM DO). The software used the ontologies to configure the architecture and to drive inferencing. This software greatly expanded the number of scenarios generated from a single set of user input.

6.1.2 Small Business and University Research Projects

The Ontology Assessment team provided input for the SBIR and University BAA calls for proposals and several team members served as collaborators or reviewers on several subsequent proposals. Dr. Liping Di and his research team from George Mason University were funded by the SAM program to develop a framework for intelligent geospatial feature discovery services.

6.1.3 Nuclear Fuel Cycle Reasoning

This FY12–FY13 follow-on project is pursuing the goal of representing domain knowledge of the nuclear fuel cycle (NFC) and weapons development processes to support useful reasoning. To implement such a reasoning framework for the entirety of the NFC is an extremely large effort, so the scope of this project will be focused on proof of concept and usefulness.

PNNL and Special Technology Laboratory (STL) staff will take a two-level approach. A “broad and shallow” ontological representation of the major steps in the NFC will be implemented. Additionally, a “deep and narrow” ontological representation of the knowledge necessary to support one or two of the major steps in the NFC will be created.

6.2 DOE Presence in the Field Established through STIDS 2011

DOE had a special session at the Semantic Technologies for Intelligence, Defense, and Security 2011 (STIDS 2011) conference at George Mason University November 15–17. Four SAM-related efforts presented papers during the session.

6.2.1 Use of Ontology to Facilitate the Creation of Synthetic Imagery of Industrial Facilities

Paul Pope and Randy Roberts presented their paper “Use of Ontology to Facilitate the Creation of Synthetic Imagery of Industrial Facilities.” The paper presents a system to create benchmark imagery for conducting verification and validation (V&V) of auto-annotation algorithms. Algorithms which perform auto-annotation of remotely sensed imagery need to undergo V&V such that end users can make a fitness-for-use judgment regarding their particular application and can be assured of a high level of confidence in achieving success.

6.2.2 Ontology-based Software for Generating Scenarios for Characterizing Search for Nuclear Materials

Richard Ward talked about *Ontology-based Software for Generating Scenarios for Characterizing Search for Nuclear Materials*. Utilizing both domain-specific ontologies and those containing latent-background terminology, Richard and his team created a software environment that generates an expanded number of scenarios from a general set of user input variables. The resulting scenarios can be used for purposes of testing algorithms for detection of SNM. The specific ontology developed, the SNM DO, was built using knowledge of the detection process for searchers in an urban setting. By using ontologies both to configure the software architecture and to drive inferencing based on ontological reasoning, they greatly expanded the number and variety of scenarios generated from a single set of user input.

6.2.3 Nuclear Fuel Cycle Reasoning

Ryan Hohimer and Jana Strasburg presented a paper, “Intelligent Hierarchical Reasoning Agents for Enhanced Decision Support.” This paper set forth a reasoning framework that analyzes data according to ontological configuration. They presented how this framework functions, and indicated the intent to use the framework to reason over nuclear fuel cycle data.

6.2.4 A Framework for Ontology-Supported Intelligent Geospatial Feature Discovery Services

Dr. Liping Di presented a paper, “A Framework for Ontology-Supported Intelligent Geospatial Feature Discovery Services.”

7.0 Lessons Learned

7.1 Project Management

The particular manner in which this multi-laboratory team was created and funded introduced challenges and inefficiencies that made project management difficult. Funding was often sent to the labs without a proposal, life cycle plan, or statement of work in place. As a result, a significant portion of the resources were used to try to organize the effort and develop a scope of work that was commensurate with the funding and skills available at each of the labs. At the same time, the objectives of the Ontology Assessment team were also changing as the program manager emphasized and de-emphasized certain aspects of the effort. The loosely defined and dynamic expectations as well as the distributed and hard-to-track resources across multiple labs made it very difficult to manage the scope, schedule and budget of any individual laboratory, much less the multi-lab effort as a whole.

Any future efforts would benefit greatly from a well-defined statement of work in place for both the individual laboratories and the overall multi-lab effort prior to the initiation of funding. Developing this after the fact for such a large effort, while trying to make progress on objectives, is exceedingly difficult.

7.2 Use-Case Focus

Progress is best achieved by focusing on a single, well defined use-case. During the course of the assessment, several different potential use-cases and approaches were suggested. For example, an ontology could be used for

- Creating a framework for communication between the various SAM assessment teams
- Cataloging the capabilities within the SAM program to identify what models, algorithms, etc. have been developed in the past and could be leveraged in the future
- Helping an all-source analyst connect different pieces of disparate data
- Understanding where a particular signature or observable was relevant within the nuclear fuel cycle
- Generating scenarios to support simulations

Instead of focusing on an application, the initial emphasis was on building an ontology with as many concepts and connections as possible. However, without a specific use -case driving the development, the utility of the resulting ontology is very limited.

7.3 Collaborative Ontology Development

Having a collaborative and distributed team working to build a moderate to large ontology requires a significant investment in resources. One individual can develop a small ontology in a relatively short time; getting consensus and buy-in from a larger group takes considerably more effort. Initial time and resources must go toward development of a specific use-case; adoption of standards to assure interoperability; and creating a collaborative environment in which multiple developers can share and critique each other's efforts. There also needs to be frequent communication among the various teams to assure that progress is being made in a unified direction. Due to a variety of factors, the Ontology

Assessment team was not able to get the needed infrastructure in place in advance and resorted to several parallel efforts, but this created much inefficiency and reduced the overall effectiveness of the project.

7.4 Knowledge Elicitation

Ontologies are best developed by close collaboration between subject matter experts and knowledge engineers. The expectations and experience of the two different groups initially caused them to talk past one another. The subject matter experts were very excited about the potential of an ontology and would provide a corpus of relevant documents to the knowledge engineers, expecting an ontology in return. On the other hand, the knowledge engineers needed the subject matter experts to provide their conceptual understanding of the domain to identify terms, definitions, and related concepts that they could then encode into an ontological framework.

There was also an impression that the ontology building could be automated. While there are a variety of programs available to mine documents for concepts, almost all of those tools use a statistical or word-frequency approach to identify terms of interest. Taking a collection of terms and constructing a consistent ontology or knowledge representation still requires a great deal of human intervention. Once an initial ontology has been developed, it can be used to analyze a collection of documents to identify items of interest and bootstrap the ontology-building process.

8.0 Conclusions

There were several very positive intangible outcomes from the Ontology Assessment team. Prior to the SAM assessment process, very few NA22 researchers were familiar with ontologies or other semantic technologies, and had little concept of how they might relate to their field of study. The Ontology Assessment effort provided funding for researchers from diverse domain backgrounds an opportunity to learn about a new field, creating capability where none previously existed. By the end of the two years, each of the institutions had at least one or two champions for semantic technologies who understood the potential value and impact ontologies could have on nonproliferation research.

The interactions between subject matter experts and knowledge engineers was also of great benefit. While many labs had pockets of expertise in semantic technologies, nearly all of these knowledge engineers were new to the domain of nuclear nonproliferation. Understanding the needs of a new group of potential users can also lead to enhancements in existing tools and the development of new approaches in knowledge representation and reasoning.

In addition, several strong inter-laboratory relationships and collaborations grew out of Assessment Team participation. Several of these collaborations continue today, well after the original Ontology Assessment team activities have come to a close.

Historically NA22 has been focused on developing sensors and hardware to make measurements. The challenge facing many government organizations is not identifying new technology to fill gaps in their information, but instead trying to understand the meaning embedded in the information they already have. As we look toward the future, the challenges of data management and meaning extraction are growing in priority. Ontologies and semantic technologies are one approach for addressing these challenges. Ontologies and semantic technologies are also becoming more mainstream as evidenced by the recent development of IBM's Jeopardy-playing Watson computer or even the new Siri application on the iPhone.

The full impact of the SAM Program's investment in the Ontology Assessment team may not be known for quite some time. Through collaboration and discussions, the assessment process planted seeds throughout the DOE complex that will take time to germinate and grow. Some of those seedlings may bear fruit and address the future needs of the nuclear nonproliferation and national security community.

Appendix A
Annotated Bibliography

Appendix A

Annotated Bibliography

Albright D and P Brannan. April 23, 2009. *Update on Khushab Plutonium Production Reactor Construction Projects in Pakistan*. Institute for Science and International Security.

Abstract: ISIS has obtained commercial satellite imagery from DigitalGlobe taken January 30, 2009 of the Khushab plutonium production reactors in Pakistan. The imagery shows that major construction of the buildings associated with the second Khushab reactor is likely finished and that the roof beams are being placed on top of the third Khushab reactor hall (figure 1). The operational status of the second reactor is unknown, but it could start in the near future.

This paper discusses nuclear production facilities. Imagery is semantically tagged. Additionally, the concept of Fissile Material Cutoff Treaty is discussed, making a connection between facilities development and treaty status.

Albright D and P Brannan. May 12, 2008. *The Al Kibar Reactor Extraordinary Camouflage Troubling Implications*. Institute for Science and International Security.

Abstract: The U.S. intelligence community's release of information about the Al Kibar reactor in Syria settles at least part of a public debate that has swirled since Israel bombed the site in the early hours of September 6, 2007. However, the new information again raises difficult questions about the ability of governments and international institutions to detect secret nuclear facilities in their early stages of construction. Early detection is critical to developing effective diplomatic means of preventing the operation of unsafeguarded nuclear facilities. The information released on April 24, eight months after the strike, does not settle all the major questions about Al Kibar and Syria's nuclear program. But it does address the fundamental question of the purpose of this facility. Faced with intense secrecy from Israel, the United States, and Syria, the public was largely left to speculate about this site through the fall and winter. After ISIS successfully located the attacked site in late October 2007 and tentatively confirmed it as a reactor, the public debate centered on whether it indeed was a nuclear reactor. Many disagreed with ISIS's assessment, citing the lack of visible attributes expected for a reactor project. Foremost among the skeptics was the International Atomic Energy Agency, whose head Mohammed El Baradei said, "Our experts who have carefully analyzed the satellite imagery say it is unlikely that this building was a nuclear facility." Popular magazines such as *The New Yorker*, also opined that the facility had apparently little to do with nuclear reactors, relying partially on the IAEA for confirmation of that mistaken view. Many experts from the nonproliferation community were drawn into the debate, which was widely recognized as one of the most mysterious nuclear stories in recent memory. One drawback of this debate is that it diverted critical public discussion of the wider implications of the strike. Few support the use of military force to settle international nuclear disputes. Director General El Baradei was particularly critical of a country taking military action over pursuing established channels for adjudicating compliance or lack thereof with the Nuclear Nonproliferation Treaty. He rightly criticized the United States and Israel for not presenting any evidence about the site to the IAEA and allowing a diplomatic process to resolve the situation peacefully. In addition, after the fiasco of alleged Iraqi WMD, many were skeptical of media reports that sourced anonymous Bush administration officials identifying the site as a reactor. Many critiques were based on a belief that a reactor construction project would have a range of signatures, and that many of these indicators would be visible in commercial satellite imagery.

These include physical protection, distinctive shapes or heights of a reactor building, power lines, water cooling system, a tall stack, specialized or additional transportation arrangements, and nearby housing. This report focuses on the astonishing lengths to which Syrian engineers and architects went to hide these commonly expected attributes and conceal the building's true purpose. This case serves as a sobering reminder of the difficulty of identifying secret nuclear activities and how too often debate about the veracity of technical assessments is subordinated to political or ideological goals on both the right and left. It should also serve as a call to bolster national and international methods to better detect these facilities. The current domestic and international capabilities to detect nuclear facilities and activities are inadequate to prevent more surprises in the future.

This paper discusses concepts related to "salient features of a nuclear facility". Additionally, the idea of obfuscation and camouflaging salient features is considered. Geospatial analytics relevance.

Albright D and P Brannan. November 13, 2008. *Arak Heavy Water Reactor Construction Progressing* Institute for Science and International Security.

Abstract: Much of the focus on Iran's nuclear facilities remains fixed on the Natanz enrichment plants and Iran's progress toward mastering the gas centrifuge uranium enrichment process. While seldom noticed, Iran continues with construction of a heavy water reactor at Arak, called the IR-40, which has a declared power of 40 megawatts-thermal.¹ ISIS has obtained commercial satellite imagery from DigitalGlobe of the Arak heavy water reactor taken on October 7, 2008 (see figure 1). A comparison of this image to one taken by GeoEye OrbView-3 (see figure 2) on February 25, 2007 shows significant progress. Construction of the reactor dome is nearing completion and other buildings adjacent to the reactor appear much further along. Construction of the cooling towers also appears complete in the October 2008 image, whereas its foundation is absent in the February 2007 image.

This paper discusses salient features of a nuclear facilities. Change detection over temporal differences used to estimate completion times.

Albright D, J Shire and P Brannan. February 11, 2009. *Is Iran Running Out of Yellowcake*. Institute for Science and International Security.

Abstract: Iran could be close to exhausting its supply of uranium oxide while lacking the adequate resources to sustain indigenous commercial-scale uranium processing and enrichment. Our conclusion, echoed in a recent report by Mark Hibbs in *Nuclear Fuel*,¹ is based on an examination of Iran's uranium reserves, its stocks of yellowcake (uranium oxide) acquired from overseas sources and, the requirements to sustain a commercial nuclear power program. The absence of activity at one of Iran's two uranium mines casts further doubt on its claims that it can establish independence in the fuel cycle required for a civil nuclear energy program.

This paper discusses image analysis and annotations. The analysis is in the context of quantifying the production from the known uranium mines in country. The paper discusses the mines' production and the relevancy to safeguards agreements. Domains of interest in this paper include; the nuclear fuel cycle, image annotation, safeguard agreements, treaties, and treaty violation.

Amitani S and E Edmonds. 2007. "Algorithms for Scenario Generation Systems." In *The 21st Annual Conference of the Japanese Society for Artificial Intelligence*.

Abstract: This paper presents design rationales and prototype of a generative system for visualising possible sequences of information to stimulate human creative thinking. Possible sequences of information segments, usually called contexts, scenarios, narratives, or storytelling, have been used as a tool for exploring and stimulating thinking about possible events, assumptions relating to these events, and courses of actions for a broad range of information designers from public audiences searching on the internet to analysts and policy makers. Thanks to the search technology, it has become easier to obtain information that a person is looking for, however, it is a laborious task to grab an overview of information space so that an information designer can find contextually relevant information pieces and sequence them into contextually meaningful ways. In this research, we design and develop a generative system that visualises possible contexts.

This paper is one of the more relevant paper with regard to reasoning. It discusses the nature of composing basic concepts into sequences of patterns.

Assistant Secretary of Defense. November 12, 1968. *Memorandum of Conversation Negotiations with Israel F-4 and Advanced Weapons*. I-35993/68.

Abstract not available.

The reporting of conversational negotiations is an example where more formal semantic technology language may be of use for the purposes of disambiguation. For example, parties in this conversation spent a great deal of energy and time in the definition of terms such as "introduce". This was done in the context of 'introducing' an area to nuclear weapons.

Baglioni M, MV Masserotti, C Renso and L Spinsanti. 2007. "Building Geospatial Ontologies from Geographical Databases." In *GeoS 2007*, pp. 195-209.

Abstract: The last few years have seen a growing interest in approaches that define methodologies to automatically extract semantics from databases by using ontologies. Geographic data are very rarely collected in a well organized way, quite often they lack both metadata and conceptual schema. Extracting semantic information from data stored in a geodatabase is complex and an extension of the existing methodologies is needed. We describe an approach to extracting a geospatial ontology from geographical data stored in spatial databases. To provide geospatial semantics we introduce new relations which define geospatial ontology that can serve as a basis for an advanced user querying system. Some examples of use of the methodology in the urban domain are presented.

This paper investigates the extraction of semantics from the structure in database schema.

Baker AB, RJ Eagan, PK Falcone, JM Harris, GV Herrera, WC Hines, RL Hutchinson, AK Moonka, ML Swinson, EK Webb, TD Woodall and GD Wyss. 2002. *A Scalable Systems Approach for Critical Infrastructure Security*. SAND2002-0877, Sandia National Laboratories, Albuquerque, NM.

Abstract: Critical infrastructures underpin the domestic security, health, safety and economic well being of the United States. They are large, widely dispersed, mostly privately owned systems operated under a mixture of federal, state and local government departments, laws and regulations. While there currently are enormous pressures to secure all aspects of all critical

infrastructures immediately, budget realities limit available options. The purpose of this study is to provide a clear framework for systematically analyzing and prioritizing resources to most effectively secure US critical infrastructures from terrorist threats. It is a scalable framework (based on the interplay of consequences, threats and vulnerabilities) that can be applied at the highest national level, the component level of an individual infrastructure, or anywhere in between. This study also provides a set of key findings and a recommended approach for framework application. In addition, this study develops three laptop computer- based tools to assist with framework implementation—a Risk Assessment Credibility Tool, a Notional Risk Prioritization Tool, and a County Prioritization tool. This study's tools and insights are based on Sandia National Laboratories' many years of experience in risk, consequence, threat and vulnerability assessments, both in defense- and critical infrastructure-related areas.

The paper is a resource for understanding a particular viewpoint about risk, consequences, threats and vulnerabilities. These concepts, along with their relationships to one another are defined and used in the context of infrastructure security.

Beck H, K Morgan, Y Jung, J Wu, S Grunwald and H-y Kwon. 2009. "Ontology-Based Simulation Applied to Soil, Water, and Nutrient Management." Springer, Berlin.

Abstract: Ontology-based simulation is an approach to modeling in which an ontology is used to represent all elements of a model. In this approach, modeling is viewed as a knowledge representation problem rather than a software engineering problem. Ontology-based techniques can be applied to describe system structure, represent equations and symbols, establish connections to external databases, manage model bases, and integrate models with additional information resources. Ontology reasoners have the potential to automatically compare, organize, search for, and discover models and model elements. We present an environment for building simulations based on the Lyra ontology management system, which includes Web-based visual design tools used for constructing models. An example application based on a model of soil, water, and nutrient management in citrus that uses the approach is also presented.

This paper addresses the differences between traditional modeling and simulation with programming languages and modeling and simulation with ontology languages. The authors describe their implementation of an Ontology Management System they call Lyra. This is an excellent background paper for integrated modeling, simulations, processing modeling, etc. There is a good literature review section on state of the art in ontology tools.

Bench-Capon TJM. 1998. "The role of ontologies in the verification and validation of knowledge based systems." In *Ninth International Workshop on Database and Expert Systems Applications*, pp. 64-69. 25-28 Aug 1998.

Abstract: The author gives some preliminary examination of the ways in which an ontology-an explicit specification of the conceptualisation of the domain-can support the verification and validation of a knowledge based system. The discussion is focussed on a simple, well known example, relating to the identification of animals. Key elements of the support provided by the ontology relate to: attempting to give coherence to the domain conceptualisation; making the role of experts in verification and validation more structured and less at the mercy of interpretation; constraining the number of test cases required to give good coverage of the possible cases; and structuring the testing to give better assurance of its efficacy, and a possible basis for greater automation of the testing process. Finally the author makes some brief remarks on the relation between the ontology of a knowledge based system, and a database and its schema.

This paper discusses the viability of using ontologies as a means or enhancement to the V&V process.

Bescond P. 1988. "Public Verification: The SPOT Satellite Technology." In *AIP Conference Proceedings, Volume 178*, pp. 149-164.

Abstract: The successful launch and operation of the SPOT satellite system has raised significant interest from independent arms-control and third-party treaty-verification organizations. SPOT is the first system to make detailed satellite images of any location on earth available to anyone on an unrestricted commercial basis. The operation of the satellite and the commercial distribution of the imagery follow the United Nations' resolution referred to as "open skies, open access," which recognizes space as an open resource, the peaceful use of which is openly available to everyone. For the first time ever, SPOT combines an unrestricted image acquisition and distribution policy with advanced capabilities including: 10-meter resolution panchromatic images; the ability to acquire repeat images of any area several times per month; and the production of stereoscopic imagery which can be viewed and analyzed in three dimensions. SPOT's characteristics have significantly decreased the minimum object size which can be detected, recognized, identified and described by interpreters of civilian satellite imagery. SPOT imagery has already been used several times by the news media to investigate and portray various strategic facilities around the world which are significant to arms control and treaty verification activities, including the Chernobyl nuclear plant, the Krasnoyarsk radar facility, the Semipalatiusk nuclear testing area, and other locations. The real significance of this usage is that, for the first time ever, detailed images of these types of facilities are openly available to anyone, and are not restricted solely to military intelligence organizations. In addition to direct observation of facilities, mapping and monitoring of environmental conditions relevant to nuclear-test monitoring can also be performed. Examples include accurate mapping and location studies for instrumentation siting, and analysis of environmental conditions which could affect monitoring activities and instrumentation results. Because of its commercial distribution policy and advanced capabilities, SPOT represents an excellent opportunity for governments, corporations and individuals to develop the potential for monitoring arms control, treaty verification and geopolitical situations.

This paper describes the features and capabilities of the SPOT satellite system. Domain knowledge about satellite imagery and that imagery's attributes is well laid out. In connection to ontology research, the paper discusses facilities and treaty verification as uses for the imagery.

Breen C, L Khan and A Ponnusamy. 2002. "Image Classification Using Neural Networks and Ontologies." In *13th International Workshop on Database and Expert Systems Applications, 2002*.

Abstract: The advent of extremely powerful home PC and the growth of the Internet have made the appearance of multimedia documents a common sight in the computer world. In the world of unstructured data composed of images and other media types, classification often comes at the price of countless hours of manual labor. This research aims to present a scalable system capable of examining images and accurately classifying the image based on its visual content. When retrieving images based on a user's query, the system yields a minimal amount of irrelevant information (high precision) and ensures a maximum amount of relevant information (high recall).

Authors describe their work in using trained neural networks to classify image content into ontologically defined class categories.

Burr T, WD Stanbro and W Charlton. 2001. "An Evaluation of Safeguards Approaches for Neptunium." *Journal of Nuclear Science and Technology* 38(3):209-216.

Abstract: The International Atomic Energy Agency has recently drawn attention to the fact that neptunium (Np), a byproduct of the nuclear power industry, can be used to make nuclear weapons. Current monitoring approaches for Np do not rely on material balance accounting as is used for uranium and plutonium. In the future this may change. Although full material balance accounting is not anticipated for Np, it is informative to evaluate the impact and benefit of full material balance accounting when considering other options. Therefore, this paper will apply systems analysis to evaluate ways to convert the current system to full materials balance accounting that will minimize the intrusiveness of the verification system and minimize costs to both the facility operator and the inspection agency. We then compare full material balance accounting to partial material balance accounting and to a ratio-monitoring technique referred to as flow sheet verification. We conclude that sampling approximately 25% of the batches is likely to be adequate and that Pu (or perhaps ¹³⁷Cs) will be the most effective surrogate for estimating the Np in the input accountability tank.

This paper is relevant to follow on efforts to model the Nuclear Fuel Cycle. The IAEA is noting the increased interest in Np and Am, and the significance to the production of nuclear weapons.

Camara G, M Egenhofer, FT Fonseca and AM Vieira Montero. 2001. "What's In An Image?" In *Spatial Information Theory, Foundations of Geographic Information Science International Conference, COSIT 2001*, pp. 474-488. September 19-23, 2001, Morro Bay, CA.

Abstract: This paper discusses the ontological status of remote sensing images, from a GIScience perspective. We argue that images have a dual nature—they are fields at the measurement level and fiat objects at the classification level—and that images have an ontological description of their own, distinct and independent from the domain ontology a domain scientist uses. This paper proposes a multi-level ontology for images, combining both field and object approaches and distinguishing between image and user ontologies. The framework developed contributes to the design of a new generation of integrated GISs, since two key benefits are achieved: (1) the support for multiple perspectives for the same image and (2) an emphasis on using images for the detection of spatial-temporal configurations of geographic phenomena.

This paper worth the read. Authors note that images can have multiple interpretations depending on the context applied.

Carley KM, M Dombroski, M Tsvetovat, J Reminga and N Kamneva. 2003. "Destabilizing Dynamic Covert Networks." In *Proceedings of the 8th International Command and Control Research and Technology Symposium*, Washington, DC.

Abstract: Most commanders, politicians and intelligence agents have at least an intuitive understanding of hierarchies and how to affect their behavior. However, covert organizations, such as terrorist organizations, have network structures that are distinct from those in typical hierarchical organizations. In particular, they tend to be more cellular and distributed. This makes it difficult to apply the lessons of experience in determining how best to destabilize these groups. This problem is further compounded by the vast quantities of, yet incomplete, information. What is needed is a set of tools and an approach to assessing destabilization strategies that takes these difficulties in to account and provides analysts with guidance in assessing destabilization tactics. Such an approach is forwarded in this paper. In addition, initial lessons learned are discussed. The particular approach is extensible and scales well to groups composed of 1000's of members.

This paper expands on subject matter expertise in the differences between traditional command hierarchies and non-traditional command hierarchies (hierarchies employed by covert organizations)

Casely-Hayford L and S Sufi. 2006. "The ISIS Facilities Ontology and OntoMaintainer." In *Proceedings of the UK e-Science All Hands Meeting 2006*.

Abstract: This paper presents an ISIS facilities ontology based on keywords in the ISIS Metadata catalog (ICAT) and OntoMaintainer, a web application to facilitate the collaborative development of ontologies. The ISIS facilities ontology aims to organize and make keywords in the ISIS Metadata Catalogue more explicit which will improve the search and navigation of data by category. The Ontology Maintainer allows users to view current versions of the ontology, and send feedback on concepts modeled back to the maintainers of the ontology. To ensure ease of use, this service has been made available to the user community via the World Wide Web.

Ontology maintenance is addressed in this paper. As with the original knowledge elicitation, knowledge maintenance is not an easy interaction between knowledge owners and knowledge engineers. A tool is discussed to aid in the maintenance effort (OntoMaintainer).

Center for Strategic Intelligence Research. 2004. *Global War on Terrorism: Analyzing Strategic Threat (Discussion Paper 13)*. Joint Military Intelligence College, Washington, DC.

Abstract: Islam generally divides the world into two "houses," the "house of Islam" and the "house of war." Although not of its own choice, the U.S. as the sole remaining superpower, and a non-Muslim state, is the modern leader of the "house of war." The context of Islam's duality and the unavoidable threat it poses to U.S. national security can be understood through the lens of lifelong scholars. The Intelligence Community must orient itself to this menace. Generally, Islam is peaceful within the "house of Islam"; however, it is prepared to extend itself into the "house of war" by force or "Jihad" (extensively discussed later). "This bipartite division of the world into an abode of peace and an abode of war finds a parallel in the communistic theory of Soviet Russia." Historically, there have been periods when Islam has abandoned forceful expansion, or has made accommodations that, at least temporarily, provide a middle ground between the two "houses." Generally, when Islam is on the advance, Jihad accelerates, and when the spread of Islam is on the decline, it stops. Given the undesirability of a conflict between the U.S. and Islam, it is imperative that U.S. policymakers understand the threat that comes with leading the "house of war," while maintaining awareness of the historical middle-ground.

From an ontological point of view this document is relevant in that it contains subject matter expertise in a number of domains. The centric theme of these domains is Islam and Jihad.

Central Intelligence Agency. 2005. *Glossary of Geospatial Terms (UNCLASSIFIED/FOUO)*.

Abstract: This glossary contains a collection of geospatial-related words and explanations.

This glossary's content is a resource for representing geospatial concepts in an ontology.

Cochran TB and MG McKinzie. January 5, 2005. "Satellite Views of the Hermit Kingdom New Perspectives on North Korea." In *Carnegie Endowment for International Peace*, Washington, DC.

Abstract not available.

This 113 page resource has materials useful for creating a knowledgebase (ontology + fact assertions) about North Korea.

Comber A, P Fisher and R Wadsworth. 2005. "You know what land cover is but does anyone else?...an investigation into semantic and ontological confusion." *International Journal of Remote Sensing* 26(1):223-228.

Abstract: Information derived from remotely sensed data is increasingly being used to describe land cover and landscape structure for a range of applications. Users in different disciplines may use this information, although they may have different perceptions of land cover from those who created the information. They may be unaware of the origin of the information or its meaning and they may treat it as if it were data. Current paradigms for reporting meta-data and data quality do not adequately communicate the producer's knowledge, and should be extended to describe the conceptual, semantic and ontological meanings of land cover.

A topic of discussion in this paper is "data pedigree", or provenance. Algorithms applied to remotely sensed data to classify land use generate an interpretation of the data that may not be fully represented in derivative products.

Cragin K and SA Daly. 2004. *The Dynamic Terrorist Threat: An Assessment of Group Motivations and Capabilities in a Changing World*. RAND, Santa Monica, CA.

Abstract: Following the terrorist attacks of September 11, 2001, the U.S. government became engaged in a war on terrorism. Such a war has already required substantial military and diplomatic resources, and it is likely to require even more. Moreover, the war on terrorism will continue in the face of other competing U.S. strategic pursuits. It is essential, therefore, that the U.S. government prioritize its counterterrorism activities and conduct the war on terrorism as efficiently as possible. The purpose of this report is twofold: first, it attempts to develop a matrix that helps policymakers identify the threat that terrorist groups pose to the United States; second, it assesses how terrorists adapt and change, to identify such groups' vulnerabilities. By combining these two approaches, the authors are able to suggest ways that the U.S. government can refine its counterterrorism policies. Thus, the report has direct relevance not only to the ongoing war on terrorism and those involved, but also to other audiences interested in the dynamic threat of terrorism.

Daly S, J Parachini and W Rosenau. 2005. *Aum Shinrikyo, Al Qaeda, and the Kinshasa Reactor: Implications of Three Case Studies for Combating Nuclear Terrorism*, RAND, Santa Monica, CA.

Abstract: Revelations about A.Q. Khan's global nuclear marketing efforts and Osama bin Laden's contact with Pakistani nuclear scientists have raised concerns about the prospects of terrorists acquiring a nuclear or radiological weapons capability. During the Cold War, there were few instances where subnational groups seemed motivated and competent to acquire such capabilities. Over the past 15 years, terrorist capabilities, motivations, and opportunities have changed. Potential terrorist acquisition of a nuclear or radiological weapons capability poses a grave danger to U.S. national security and to the security of the international system of nation-states. The historical record of terrorists pursuing nuclear and radiological capabilities is small in size, complicated by significant information gaps, and not well understood. The size of the dataset and the considerable unknowns about the cases where groups have sought these capabilities make it difficult to assess the nature of the danger and to anticipate new developments in the nature of the threat. However, given the potential consequences of terrorist theft of a nuclear weapon or indigenous development of a nuclear device—even one employing a crude design that produces

only a small nuclear yield—poses a serious danger that the United States and other allied nations must take extraordinary measures to thwart. Developing an effective and comprehensive strategy to prevent terrorist acquisition of nuclear and radiological weapons capabilities must begin with a thorough understanding of the historical record of terrorist efforts and opportunities to acquire these capabilities. The three case studies described herein were undertaken to contribute to an understanding of this phenomenon. This documented briefing is part of a larger project for the U.S. Air Force entitled “Denying Armageddon: Preventing U.S. Adversaries from Acquiring Nuclear Weapons.”

Di L, P Zhao, W Yang and P Yue. 2006. "Ontology-driven Automatic Geospatial-Processing Modeling based on Web-service Chaining." In *Earth Science Technology Conference*.

Abstract: Earth System Science (ESS) research and applications often involve in collecting, analyzing and modeling with distributed heterogeneous geospatial data. Those data are processed step-by-step in geospatial analysis systems to extract information and knowledge products for applications and decision makings. Conceptually, such a step-by-step process forms a geospatial processing model that represents the knowledge of geospatial domain experts. This paper presents a study on ontology-driven automatic creation and execution of geospatial processing models in GeoBrain, a Web-service based geospatial knowledge system, to produce user-specific products. Web Services and Service-Oriented Architecture (SOA) provide a framework to support interoperable machine-to-machine interaction over a network. Web service chaining aimed to solve complex application tasks is changing the way of developing and deploying applications. By wrapping data and processes with Web services, it is easy to transform a geospatial processing model into a service chain. From design (knowledge), instantiation (information) to execution (data), this paper illustrates the whole life cycle of the geospatial-processing modeling and relevant implementation in GeoBrain. Ontology is usually used to capture domain knowledge. In this paper, a number of ontologies, including geospatial scientific ontology, geospatial data ontology and geospatial processing ontology, are introduced as the knowledge base to present geospatial domain terms and concepts, linkage between concepts and datasets, relationships among heterogeneous data, and associations between processes and data. By these ontologies, data and processes can be used for more effective discovery, automation, integration, and reuse across multiple diverse applications.

Durak U, S Guler, H Oguztuzun and SK Ider. 2007. "An Exercise in Ontology Driven Trajectory Simulation with MATLAB SIMULINK." In *21st European Conference on Modelling and Simulation*, pp. 435-440.

Abstract: We demonstrate an application of the ontology driven methodology to develop trajectory simulations in a function-oriented style. We adopt a model based approach to software development, guided by the domain engineering process, to promote knowledge and software reuse. MATLAB Simulink block definitions have been generated from the function specifications in the Trajectory Simulation Ontology, called TSONT. MATLAB implementations of the blocks have been generated from the DAVE-ML definitions of the functions, which are incorporated in TSONT. Finally, the simulation has been put together by manually connecting the blocks.

Dzemydiene D and E Kazemikaitiene. 2005. "Ontology-Based Decision Support System for Crime Investigation Processes Information Systems Development." In *Information Systems Development: Advances in Theory, Practice, and Education*, pp. 427-438. eds: O Vasilecas, W Wojtkowski, J Zupančič, A Caplinskas, W Wojtkowski and S Wrycza. Springer US, New York.

Abstract: The characteristics of complex and dynamic crime investigation domain require new ways of information extraction and knowledge representation. The crime analysis information system, based on ontology, ensures the proper application of the structural model of crime information, determines the main rules, how to acquire the important forensic and crime investigation information about crime from the primary sources. The goal of developing a decision support system (DSS) is to act "helpfully" in terms of reliably reproducing the cognitive behaviour of crime investigators demands thoroughly considered concepts, principles, and crime scenarios. We can classify the problems arising in the investigation of a crime into two main groups: the problems of a purely criminal nature and those associated with the criminal law, penal process, organization, and criminology. One can solve the first group of problems by employing the crime analysis information, while for the problems of the other group we need information on the criminal law, penal process and the like, that would be significant for the crime investigation.

This paper is very relevant to the follow on work of the NFC Reasoning project. The concepts of supporting a criminal investigation are parallel with those of investigating proliferation. Authors use UML to illustrate their points.

Echevarria III AJ. 2005. *Fourth Generation War and Other Myths*. U.S. Army War College, Carlisle, PA.

Abstract: Fourth Generation War (4GW) emerged in the late 1980s, but has become popular due to recent twists in the war in Iraq and terrorist attacks worldwide. Despite reinventing itself several times, the theory has several fundamental flaws that need to be exposed before they can cause harm to U.S. operational and strategic thinking. A critique of 4GW is both fortuitous and important because it also provides us an opportunity to attack other unfounded assumptions that could influence U.S. strategy and military doctrine. In brief, the theory holds that warfare has evolved through four generations: 1) the use of massed manpower, 2) firepower, 3) maneuver, and now 4) an evolved form of insurgency that employs all available networks -political, economic, social, military- to convince an opponent's decisionmakers that their strategic goals are either unachievable or too costly. The notion of 4GW first appeared in the late 1980s as a vague sort of "out of the box" thinking, and it entertained every popular conjecture about future warfare. However, instead of examining the way terrorists belonging to Hamas or Hezbollah (or now Al Qaeda) actually behave, it misleadingly pushed the storm-trooper ideal as the terrorist of tomorrow. Instead of looking at the probability that such terrorists would improvise with respect to the weapons they used-box cutters, aircraft, and improvised explosive devices-it posited high-tech "wonder" weapons. The theory went through a second incarnation when the notion of nontrinitarian war came into vogue; but it failed to examine that notion critically. The theory also is founded on myths about the so-called Westphalian system and the theory of blitzkrieg. The theory of 4GW reinvented itself once again after September 11, 2001 (9/11), when its proponents claimed that Al Qaeda was waging a 4GW against the United States. Rather than thinking critically about future warfare, the theory's proponents became more concerned with demonstrating that they had predicted the future. While their recommendations are often rooted in common sense, they are undermined by being tethered to an empty theory. What we are really seeing in the war on terror, and the campaign in Iraq and elsewhere, is that the increased dispersion and democratization of technology, information, and finance" brought about by globalization has given terrorist groups greater mobility and access worldwide. At this point,

globalization seems to aid the nonstate actor more than the state, but states still play a central role in the support or defeat of terrorist groups or insurgencies. We would do well to abandon the theory of 4GW altogether, since it sheds very little, if any, light on this phenomenon.

Fisk MD. 2002. "Accurate Locations of Nuclear Explosions at the Lop Nor Test Site Using Alignment of Seismograms and IKONOS Satellitel Imagery." *Bulletin of the Seismological Society of America* 92(8):2911-2925.

Abstract: Locations are estimated of 11 underground nuclear explosions conducted between 26 May 1990 and 29 July 1996 at the Lop Nor test site in western China. To perform the analysis, seismic recordings were obtained from the Center for Monitoring Research (CMR), U.S. National Data Center, and Incorporated Research Institutions for Seismology (IRIS) Data Management Center. Using recordings with adequate signal-to-noise ratios, P-wave arrivals are picked manually by aligning waveforms of the events at common stations. This provides very precise relative phase timings (i.e., with relative errors typically much less than 0.1 sec). Based on analysis of commercial IKONOS satellite imagery, the locations of three explosions are fixed and used as master events. The other events are located relative to the nearest master event after correcting their travel times by the travel-time residuals (relative to IASP91) of the master events. Location estimates of the 25 September 1992 and 29 July 1996 nuclear tests are within about 100 m of each other. The 90%-confidence error ellipse for the 29 July 1996 event has a semimajor axis of less than 0.3 km. The results indicate that these two tests were likely conducted in the same tunnel. Location estimates of nine explosions conducted at the eastern zone of the Lop Nor test site are within 140 to 490 meters of vertical boreholes and infrastructure to support nuclear testing observed in IKONOS imagery. The location estimates of the eleven Chinese explosions are all considered to be well within one km of the true detonation points and, hence, are categorized as GT1 events (as defined by Yang et al., 2000a). These solutions allow seismic travel times to be calibrated accurately for the Lop Nor region at seismographic stations of the International Monitoring System (IMS).

Fonseca FT, MJ Egenhofer and P Agouris. 2002. "Using Ontologies for Integrated Geographic Information Systems." *Transactions in GIS* 6(3).

Abstract: Today, there is a huge amount of data gathered about the Earth, not only from new spatial information systems, but also from new and more sophisticated data collection technologies. This scenario leads to a number of interesting research challenges, such as how to integrate geographic information of different kinds. The basic motivation of this paper is to introduce a GIS architecture that can enable geographic information integration in a seamless and flexible way based on its semantic value and regardless of its representation. The proposed solution is an ontology-driven geographic information system that acts as a system integrator. In this system, an ontology is a component, such as the database, cooperation to fulfill the system's objectives.

This paper, written in 2002, is a bit dated. It does reflect the increasing thought at the time that ontology for data fusion of diverse representation types is the way to go...

Frishkoff G, P LePendu, R Frank, H Liu and D Dou. 2009. "Development of Neural Electromagnetic Ontologies (NEMO): Ontology-based Tools for Representation and Integration of Event-related Brain Potentials." In *Proceedings of the International Conference on Biomedical Ontologies*.

Abstract: We describe a first-generation ontology for representation and integration of event-related brain potentials (ERPs). The ontology is designed following OBO "best practices" and is

augmented with tools to perform ontology-based labeling and annotation of ERP data, and a database that enables semantically based reasoning over these data. Because certain high-level concepts in the ERP domain are ill-defined, we have developed methods to support coordinated updates to each of these three components. This approach consists of “top-down” (knowledge-driven) design and implementation, followed by “bottom-up” (data-driven) validation and refinement. Our goal is to build an ERP ontology that is logically valid, empirically sound, robust in application, and transparent to users. This ontology will be used to support sharing and meta-analysis of EEG and MEG data collected within our Neural Electromagnetic Ontologies (NEMO) project.

General Accounting Office. 2002. *Export Controls Department of Commerce Controls Over Transfers of Technology to Foreign Nationals Need Improvement*. GAO-02-972, U.S. General Accounting Office, Washington, DC.

Abstract: To protect its national security and foreign policy interests, the United States controls exports of civilian technologies that have military uses. U.S. firms may be required to obtain a license from the Department of Commerce before exporting these "dual-use" technologies from the United States to many other countries, including countries of concern. Since Commerce regulations also deem domestic transfers of controlled dual-use technologies to citizens of these countries to be exports, Commerce may require firms that employ foreign nationals working with these technologies in this country to obtain "deemed" export licenses. The firms should, in many cases, hold a deemed export license, and the foreign nationals should have an appropriate visa classification, such as an H-1B specialized employment classification. Commerce issues deemed export licenses to firms that employ or sponsor foreign nationals after consulting the Departments of Defense, State, and Energy. Deemed export licenses are generally valid for 2 years and comprise almost 10 percent of all export licenses approved by Commerce. In fiscal year 2001, Commerce approved 822 deemed export license applications and rejected 3. Most of the approved licenses allowed foreign nationals from countries of concern to work with advanced computer, electronic, or telecommunication and information security technologies in the United States. To better direct its efforts to detect possible unlicensed deemed exports, in fiscal year 2001 Commerce screened thousands of applications for H-1B and other types of visas submitted by foreign nationals overseas. From these applications, it developed 160 potential cases for follow-up by enforcement staff in the field. However, Commerce did not screen thousands of H-1B change-of-status applications submitted domestically to the Immigration and Naturalization Service for foreign nationals already in the United States. In addition, Commerce could not readily track the disposition of the 160 cases referred to field offices for follow-up because it lacks a system for doing so. Commerce attaches security conditions to almost all licenses to mitigate the risk of providing foreign nationals with controlled dual-use technologies. However, according to senior Commerce officials, Commerce staff do not regularly visit firms to determine whether these conditions are being implemented because of competing priorities, resource constraints, and inherent difficulties in enforcing several conditions.

Gomez Jr. J. 2005. "Terrorist Motivations for the Use of Extreme Violence." *Strategic Insights* 4(5).

Abstract: Islamist violence must be measured and well thought out. Misconstruing the underlying causes of Islamist rage or overreacting to Islamist violence may only intensify militancy, not temperate it.” For policy-makers, understanding the source or cause of discontent serves as the best hope to remedy the ills that lie beneath what some refer to as “sacred” terrorism or religious terrorism. Yet a basic question is what would extremists hope to achieve by resorting to such violent acts? Is terrorism ever rational? Can terrorism be deterred? All too often, our analysis of extremist motives begins with our reaction to the terrorist act itself. However, for certain

extremist organizations, channeling efforts to identify and isolate the root cause for such events requires a deeper understanding of the intricacies that foment such profound actions, specifically suicide terrorism and the use of a weapon of mass destruction. In a recent review by Jeff Goodwin of Jessica Stern's *Terror in the Name of God*, he stated that few studies probe deeply into the cause of terrorism and, as a result, "it remains a mystery. A contributing factor is that social movement scholars with very few exceptions have said little about terrorism. Nor have they paid sustained attention to the more general question of how movement organizations make strategic choices, of which terrorism is one." [2] To resolve group level problems we need to view it from a group level or movement level perspective. It is at that point that we can start crafting more tailored solutions to counter the extremist threat. The central theme of this essay is to examine the rationale behind terrorist use of weapons of mass destruction. Here, I compare Aum Shinrikyo and al Qaeda to analyze to what extent their actions were the result of strategic choice or group behavior. In brief, Aum Shinrikyo's decision to attack civilians on a Japanese subway reflects an organization in a desperate fight for survival. According to Martha Crenshaw's organizational perspective, "terrorist actions often appear inconsistent, erratic, and unpredictable" and terrorist acts occur as a result of internal group dynamics. [3] The group's ultimate decision to strike a Tokyo subway system was as much an attack on Japan's political culture as it was an act by a desperate group. In contrast, al Qaeda's methodical planning and extensive preparation reflects an instrumental approach where the act of terrorism is that of strategic choice founded on the basis of collective values. According to Crenshaw, such an organization ultimately fails when the group is unable to reach its political objectives or when the cost of conducting such terrorist acts exceeds any foreseeable benefits.

Gruber T. 1995. "Toward Principles for the Design of Ontologies Used for Knowledge Sharing." *International Journal of Human-Computer Studies* 43(5-6):907-928.

Abstract: Recent work in Artificial Intelligence (AI) is exploring the use of formal ontologies as a way of specifying content-specific agreements for the sharing and reuse of knowledge among software entities. We take an engineering perspective on the development of such ontologies. Formal ontologies are viewed as designed artifacts, formulated for specific purposes and evaluated against objective design criteria. We describe the role of ontologies in supporting knowledge sharing activities, and then present a set of criteria to guide the development of ontologies for these purposes. We show how these criteria are applied in case studies from the design of ontologies for engineering mathematics and bibliographic data. Selected design decisions are discussed, and alternative representation choices are evaluated against the design criteria.

Gruber T. 2009. "Ontology." In *Encyclopedia of Database Systems*, pp. 1963-1965. eds: L Liu and MT Ozsu. Springer-Verlag.

Abstract: In the context of computer and information sciences, an ontology defines a set of representational primitives with which to model a domain of knowledge or discourse. The representational primitives are typically classes (or sets), attributes (or properties), and relationships (or relations among class members). The definitions of the representational primitives include information about their meaning and constraints on their logically consistent application. In the context of database systems, ontology can be viewed as a level of abstraction of data models, analogous to hierarchical and relational models, but intended for modeling knowledge about individuals, their attributes, and their relationships to other individuals. Ontologies are typically specified in languages that allow abstraction away from data structures and implementation strategies; in practice, the languages of ontologies are closer in expressive power to first-order logic than languages used to model databases. For this reason, ontologies are

said to be at the "semantic" level, whereas database schema are models of data at the "logical" or "physical" level. Due to their independence from lower level data models, ontologies are used for integrating heterogeneous databases, enabling interoperability among disparate systems, and specifying interfaces to independent, knowledge-based services. In the technology stack of the Semantic Web standards [1], ontologies are called out as an explicit layer. There are now standard languages and a variety of commercial and open source tools for creating and working with ontologies.

This is a good paper for the ontological newbie. It is a good reference to use if you are laying the background of using ontologies.

Gupta V and F Pabian. 1997. "Investigating the allegations of Indian nuclear test preparations in the Rajasthan desert." *Science & Global Security* 6(2):101-188.

Abstract: This paper describes how commercial satellite imagery was used along with news reports and published scientific articles to investigate the December 1995 allegations of Indian nuclear test preparations in the Rajasthan Desert. Taking the allegations against India as an example of a future CTB compliance dispute, the investigation was conducted to test the utility of commercial satellite imagery for CTB verification. The technical inquiry produced a series of findings on India's nuclear testing history and on the recent nuclear test allegations. These findings included the exact location of the subsidence crater created by the May 18, 1974 nuclear test, the discovery of an adjacent military range near Khetolai village, and the observation of recent large-scale, unusual activity at this military range in the immediate vicinity of the 1974 test site. The image-derived information was used to sift fact from fiction in the conflicting media reports. It was then integrated into the collection of credible evidence and analyzed to determine whether the observed activity at the Khetolai military range was conventional, missile testing, nuclear, or innocuous.

Hamazaki T and T Sakata. Winter 2009. "Japanese Earth Observation Satellite Program and Its Potential Contribution to IAEA Safeguards Activities." *JNMM* 27(2):25-28.

Abstract: This paper intends to show the outline of Japanese Earth-observation satellite program and its potential contribution to the International Atomic Energy Agency's safeguards activities. It introduces a summary of some Japanese Earth-observation satellite program and the overview of the Advanced Land Observing Satellite. Data availability and restrictions are discussed, and some recommendations to the IAEA are described.

Hecht G. 2006. "Nuclear Ontologies." *Constellations* 13(3):320-331.

Abstract: In October 2002, George W. Bush claimed that Saddam Hussein had "recently sought significant quantities of uranium from Africa." Pressed for details, his administration cited CIA intelligence that Iraq had tried to purchase 500 tons of uranium from Niger. Skepticism mounted, with the CIA itself expressing doubts about the intelligence. By the 2003 State of the Union address, senior officials instead credited the information to British intelligence. But the source mattered less, according to Bush, than the inescapable conclusion: Iraq planned to build nuclear weapons. Since then, both the uranium and the weapons claims have been decisively refuted. In February 2002, the CIA sent former diplomat Joseph Wilson to investigate whether Niger had indeed concluded a deal with Iraq. Wilson found no trace of the alleged sale. When he heard Bush's statement a few months later, he initially assumed that the president meant some other uranium-producing African nation. Upon realizing that Bush really did mean Niger, an appalled Wilson went public in the New York Times. In an attempt to discredit him, an "anonymous

source” from the Bush administration outed Wilson’s wife as a CIA operative. Meanwhile, it turned out that the proof that Saddam sought uranium “from Africa” consisted of forged documents peddled by an Italian businessman. Most analysts now agree that there simply was no Iraq-Niger uranium deal.¹ The re-emergence of this story in the media produced a cacophony of opinion, mostly pitting Bush stalwarts against defenders of Wilson. To the administration’s obvious delight, bickering over personal credibility drowned out the pivotal issue of whether Bush misled the nation into war. It also bypassed the real significance of the Niger episode for contemporary global nuclear relations.

Horowitz BM and YY Haimes. 2003. "Risk-Based Methodology for Scenario Tracking, Intelligence Gathering, and Analysis for Countering Terrorism." *Systems Engineering* 6(3):152-169.

Abstract: Disruption of a terrorist attack depends on having information facilitating the identification and location of those involved in supporting, planning, and carrying out the attack. Such information arises from myriad sources, such as human or instrument surveillance by intelligence or law enforcement agencies, a variety of documents concerning transactions, and tips from a wide range of occasional observers. Given the enormous amount of information available, a method is needed to cull and analyze only that which is relevant to the task, confirm its validity, and eliminate the rest. The risk-based methodology for scenario tracking, intelligence gathering, and analysis for countering terrorism builds on the premise that in planning, supporting, and carrying out a terrorist plot, those involved will conduct a series of related activities for which there may be some observables and other acquirable evidence. Those activities taken together constitute a threat scenario. Information consistent with a realistic threat scenario may be useful in thwarting an impending attack. Information not consistent with any such scenario is irrelevant. Thus, the methodology requires a comprehensive set of realistic threat scenarios that would form a systemic process for collecting and analyzing information. It also requires a process for judging the validity and usefulness of such information. The key questions for intelligence gathering and analysis are: how to produce a comprehensive set of threat scenarios, how to winnow that set to a subset of most likely scenarios, what supplementary intelligence is worth pursuing, how to judge the relevance of available information, and how to validate and analyze the information. The methodology presented in this paper can serve as a vehicle with which to enable the intelligence community to better: (a) assess the intent and capabilities of terrorist groups, (b) develop and compare terrorist scenarios from different sources and aggregate the set that should guide decisions on intelligence collection, (c) assess the possible distributions of responsibility for intelligence gathering and analysis across various homeland security agencies at the federal, state, and local levels, and (d) establish effective collection priorities to meet the demands of counterterrorism. Some of the critical issues addressed in this paper include: (1) how to create a reasonably complete set of scenarios and filter it down to a more manageable set to establish intelligence collection priorities, (2) how to integrate the wide variety of intelligence sources associated with monitoring for terrorism and analytically account for the corresponding disparities in information reliability, and (3) how to incorporate these new methodologies into existing information management efforts related to protecting our nation’s critical infrastructures.

Hsin-Nan L, T Shian-Shyong, W Jui-Feng, L Huan-Yu and S Jun-Ming. 2007. "An Iterative, Collaborative Ontology Construction Scheme." In *Second International Conference on Innovative Computing, Information and Control, ICICIC '07*, pp. 150-150. 5-7 Sept. 2007.

Abstract: Nowadays, ontology of subject knowledge is applied for intelligent e-learning to provide learners with adaptive learning guidance and efficient learning content management. However, it is difficult and costly to construct an ontology for a given domain even though using

the GUI tools. Therefore, how to propose an easy way to acquire the knowledge of experts, how to integrate these diverse opinions of experts into the consensus, and how to incrementally refine and converge the ontology to maintain the most suitable one have been paid attention recently. Accordingly, in this paper, we propose an Iterative, Collaborative Ontology Construction (ICOC) scheme. In each iteration, our proposed wild-like online ontology editor can help users collaboratively contribute their knowledge. Next, it automatically generates an appropriate questionnaire to converge the opinions of users to a new version of ontology using a Delphi-like method. The convergence process stops when all the relations are converged or eliminated by the questionnaire analysis. According to this ICOC scheme, the created ontology can be more acceptable.

IAEA. 2008. *Fast Reactor Knowledge Preservation System: Taxonomy and Basic Requirements*. IAEA Nuclear Energy Series No. NG-T-6.3.

Abstract: The IAEA has taken the initiative to coordinate efforts of Member States in the preservation of knowledge in the area of fast reactors. In the framework of this initiative, the IAEA intends to create an international database compiling information from different Member States on fast reactors through a web portal. Other activities related to this initiative are being designed to accumulate and exchange information on the fast reactor area, to facilitate access to this information by users in different countries and to assist Member States in preserving the experience gained in their countries. The purpose of this publication is to develop a taxonomy of the Fast Reactor Knowledge Preservation System (FRKPS) that will facilitate the preservation of the world's fast reactor knowledge base, to identify basic requirements of this taxonomy on the basis of the experience gained in the fast reactor area, as well as results of previous IAEA activities on fast reactor knowledge preservation. The need for such taxonomy arises from the fact that during the past 15 years there has been stagnation in the development of fast reactors in the industrialized countries that were involved, earlier, in intensive development of this area. All studies on fast reactors have been stopped in countries such as Germany, Italy, the United Kingdom and the United States of America and the only work being carried out is related to the decommissioning of fast reactors. Many specialists who were involved in the studies and development work in this area in these countries have already retired or are close to retirement. In countries such as France, Japan and the Russian Federation that are still actively pursuing the evolution of fast reactor technology, the situation is aggravated by the lack of young scientists and engineers moving into this branch of nuclear power. Appreciation is expressed to all the participants who contributed to this publication.

IAEA. May 16, 2003. *Communications Received from Certain Member States Regarding Guidelines for Transfers of Nuclear-Related Dual-Use Equipment, Materials, Software and Related Technology*. INFCIRC/254/Rev.5/Part 2.

Abstract: With the objective of averting the proliferation of nuclear weapons and preventing acts of nuclear terrorism, suppliers have had under consideration procedures in relation to the transfer of certain equipment, materials, software, and related technology that could make a major contribution to a "nuclear explosive activity," an "unsafeguarded nuclear fuel-cycle activity" or acts of nuclear terrorism. In this connection, suppliers have agreed on the following principles, common definitions, and an export control list of equipment, materials, software, and related technology. The Guidelines are not designed to impede international co-operation as long as such co-operation will not contribute to a nuclear explosive activity, an unsafeguarded nuclear fuel-cycle activity or acts of nuclear terrorism. Suppliers intend to implement the Guidelines in accordance with national legislation and relevant international commitments.

Inokuchi M and M Mekawa. 2005. "Knowledge Creation Semiconductor Factories by User-Oriented Yield Management System." In *Engineering Management Conference*, pp. 720-724.

Abstract: This paper presents a method of creating and managing advanced engineering organizational knowledge. For a semiconductor manufacturer to remain competitive in the marketplace, it must have a highly effective semiconductor yield management system (YMS). The construction and maintenance of such a YMS is a typical example of the creation and management of organizational engineering knowledge, and also an example of how to support it by information technology. This paper proposes a new concept—user-oriented engineering. User-oriented engineering is an approach based on the user centric concept and end user programming, applying ontology and problem solving method, to optimize the end user's knowledge and to share the knowledge. YMSs based on user-oriented engineering (UO-YMS) have been developed and have been in operation in several semiconductor factories for several years. The tacit knowledge of expert engineers has been transformed into explicit knowledge, then shared by the organization and transformed into systemic knowledge. The result has been a reduction in yield - management tool development time, shorter yield analysis time required studying loss, and a greatly improved yield.

Jasani B. Winter 1999. "Civil Radar Observation Satellites for IAEA Safeguards." *JNMM* 27(2):29-34.

Abstract: The usefulness of commercial remote sensing satellite imagery to monitoring nuclear arms related treaties and conventions is briefly considered in this paper. It is argued that observations from satellites can now contribute to the International Atomic Energy Agency's (IAEA) safeguards procedures. Now that the IAEA may soon embark upon the use of important measures are the Comprehensive Nuclear Tests Ban Treaty (CTBT) and a possible Fissile Material Cut-off Treaty (FMCT). Verification of these treaties using commercial satellites is explored briefly in this paper. Not only the arms control treaties but the question of the management of the environment after a nuclear accident is also briefly discussed here. Such activities are being facilitated by the increasing availability of relatively high-resolution (0.8-1m panchromatic and 4m multi-spectral) satellite Imagery.

Jones SA. 2004. *Current and Future Challenges for Asian Nonproliferation Export Controls: A Regional Response*. U.S. Army War College, Carlisle, PA.

Abstract: The proliferation of weapons of mass destruction (WMD) and advanced conventional arms remains one of the gravest threats to the security of the international community. Countries of concern continue to pursue WMD by purchasing related technologies and components from foreign suppliers. Export controls represent one of the key elements of a comprehensive nonproliferation strategy. They include procedures adopted by countries to regulate and monitor trade in weaponry and weapons-related (dual use) technologies. However, the effectiveness of export control as a tool for limiting the spread of sensitive technologies and weaponry has been compromised by globalization and a complex array of international developments. The distinction between military and commercial products, for example, has become less clear. Therefore, it is likely that export control policies and institutions need to be continually adjusted if they are to serve international security objectives. How countries in the Asia region respond to the relentlessly changing nature of the proliferation challenge will affect profoundly the shape of global security. In many instances, the countries of the region are major transshipment and assembly points for critical strategic dual-use goods and technologies. Some of these countries are major producers of strategic items, while others are or have potential to become suppliers. Yet, national export control systems in the region, with a few exceptions, remain rudimentary and resource-poor. This monograph examines the current state of export control system development

in the greater Asia region, with particular emphasis on the economic and security environment in which these systems operate. Identification then is made of the gains and remaining deficiencies in export control development. The monograph concludes by examining the applicability of the European Union's effort to coordinate export controls to the regional forces shaping the regional trade and security dynamics in Asia and what the United States can do to facilitate greater export control development and cooperation.

Kessler C, K Janowicz and M Bishr. 2009. "An Agenda for the Next Generation Gazetteer Geographic Information Contribution and Retrieval." In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '09)*.

Abstract: Gazetteers are key components of georeferenced information systems, including applications such as Web-based mapping services. Existing gazetteers lack the capabilities to fully integrate user-contributed and vernacular geographic information, as well as to support complex queries. To address these issues, a next generation gazetteer should leverage formal semantics, harvesting of implicit geographic information -- such as geotagged photos -- as well as models of trust for contributors. In this paper, we discuss these requirements in detail. We elucidate how existing standards can be integrated to realize a gazetteer infrastructure allowing for bottom-up contribution as well as information exchange between different gazetteers. We show how to ensure the quality of user-contributed information and demonstrate how to improve querying and navigation using semantics-based information retrieval.

Kessler C, M Raubal and C Wosniok. 2009. "Semantic Rules for Context-Aware Geographical Information Retrieval." In *EuroSSC'09 Proceedings of the 4th European conference on Smart sensing and context*.

Abstract: Geographical information retrieval (GIR) can benefit from context information to adapt the results to a user's current situation and personal preferences. In this respect, semantics-based GIR is especially challenging because context information { such as collected from sensors { is often provided through numeric values, which need to be mapped to ontological representations based on nominal symbols. The Web Ontology Language (OWL) lacks mathematical processing capabilities that require free variables, so that even basic comparisons and distance calculations are not possible. Therefore, the context information cannot be interpreted with respect to the task and the current user's preferences. In this paper, we introduce an approach based on semantic rules that adds these processing capabilities to OWL ontologies. The task of recommending personalized surf spots based on user location and preferences serves as a case study to evaluate the capabilities of semantic rules for context-aware geographical information retrieval. We demonstrate how the Semantic Web Rule Language (SWRL) can be utilized to model user preferences and how execution of the rules successfully retrieves surf spots that match these preferences. While SWRL itself enables free variables, mathematical functions are added via built-ins { external libraries that are dynamically loaded during rule execution. Utilizing the same mechanism, we demonstrate how SWRL built-ins can query the Semantic Sensor Web to enable the consideration of real-time measurements and thus make geographical information retrieval truly context-aware.

Kitamura Y, M Ikeda and R Mizoguchi. 1996. "A Qualitative Reasoning Based on an Ontology of Fluid Systems and Its Evaluation." In *PRICAI'96: Topics in Artificial Intelligence*, LNCS 1114, pp. 288-299.

Abstract: This research is concerned with causal understanding and qualitative reasoning of behavior of physical systems, which are crucial issues of model based problem solving. In this paper, a new method of qualitative reasoning and causal ordering is proposed and its application

to a power plant is presented. The method is based on our kernel ontologies of causality and time-resolution and a domain ontology of fluid systems. These ontologies help make the design rationales of our method explicit and facilitate reusability of our models. The whole of the target system is represented by combining a set of local component models and global constraints. The component models include local and causal characteristics of each component which are independent of context for their reuse on the basis of the ontology of causality. Global constraints with time-scales are derived according to the general properties of the physical entity which are prepared beforehand as a part of the domain ontology. They contribute to providing intuitive causal ordering of complex behavior originated in various configurations of components, including inter-component negative feedback. Furthermore, the method has been successfully applied to a power plant. All the reasoning results matched those obtained by a domain expert including their ambiguities.

Kokar MM, JA Letkowski, CJ Matheus, M Hinman, K Baclawski and J Salerno. 2004. "Use Cases for Ontologies in Information Fusion." In *Proceedings of the Seventh International Conference on Information Fusion*, pp. 415-421.

Abstract: Ontologies are becoming increasingly popular due to recent efforts to extend the capabilities of the World Wide Web through the addition of formal semantics. While ontologies have traditionally been used as precise languages to facilitate efficient exchange of information among people, the "Semantic Web" is extending this role to software agents. For this to be possible, ontologies must be formalized in languages processable by computers, such as OWL, the W3C's Web Ontology Language. The purpose of OWL ontologies is to permit software agents to understand web content and to interact intelligently with Web services (which may themselves be software agents). The use of such ontologies, however, need not be constrained to the Web. Recently, ontologies have found their way into higher-level information fusion where they are providing a means for describing and reasoning about sensor data, objects, relations and general domain theories. To the best of our knowledge, there is as of yet no documented effort to capture the main uses of ontologies in information fusion. In this paper, we start filling this void by presenting a number of "use cases," i.e., scenarios of the use of ontologies in the context of higher-level information fusion. In this paper, we develop use cases in which ontologies are used both for the fusion process itself and for the development of fusion systems. The use cases cover scenarios in which the agent roles are played by people, software or both.

This paper discusses data fusion and information fusion primarily. It is one of the few papers mentioning consistency checking amongst the advantages to using ontologies. Concept in this paper illustrated with UML diagrams.

Kwon H-Y, S Grunwald, HW Beck, Y Jung, SH Daroub, TA Lang and KT Morgan. 2010. "Ontology-based simulation of water flow in organic soils applied to Florida sugarcane." *Agricultural Water Management* 97(1):112-122.

Abstract: An ontology-based simulation (OntoSim) is a unique data modeling environment where soil-plant nutrient processes are represented as database objects and the user-defined relationships among objects are used to generate computer code (Java) for running the simulation. The aim of this study was to model hydrologic processes of sugarcane-grown organic soils utilizing OntoSim in the Everglades agricultural Area (EAA) of South Florida. This OntoSim-Sugarcane model describes the complex hydrology of subirrigation and open ditch drainage commonly used on Florida farms. Model calibration was conducted by (i) selecting rectangular farm water management units (<12 ha), which are encompassed with farm ditches, from two farms in the EAA, (ii) assembling all relevant input data including water tables (WT) recorded at

the monitoring farm well of each unit, and (iii) optimizing the fits between the simulated and observed daily WT during two consecutive water years (WY). By calibrating two site-specific parameters – lateral saturated hydraulic conductivities of soil profiles and vertical saturated hydraulic conductivity of the underlying limestone bedrock – good agreement between simulated and observed daily WT was obtained (Nash–Sutcliffe efficiency coefficient >0.65 ; coefficient of residual mass $<1\%$) within the units during WY96–97 (May 1995–April 1997). The validation of the model during subsequent WY98–99 at both units also showed Nash–Sutcliffe efficiency >0.55 and coefficient of residual mass $<3\%$. It indicated that OntoSim-Sugarcane is able to simulate daily fluctuations of WT within the farm units and estimate lateral drainage/sub-irrigation and deep seepage that significantly contribute to the water balance at farms in the EAA. Thus, it can be a promising management tool to provide farmers with accurate assessment of water movement in this agricultural area.

Lagos N and RM Stetchi. 2007. "A Manufacturing Ontology for eLearning." In *IPROMS 2007*.

Abstract: The paper presents the common ontology being developed in the ISAR project that integrates different aspects of the e-learning ISAR system. First existing ontologies developed in the manufacturing domain are reviewed and the need for an integrated approach towards developing e-learning lessons based on a common ontology is established. The presented ontology defines the training material, context, and manufacturing domains, as well as their interactions. An illustrative example demonstrates the use of the common ontology as part of the ISAR e-learning system.

Lemaignan S, A Siadat, JY Dantan and A Semenenko. 2006. "MASON: A Proposal For An Ontology Of Manufacturing Domain." In *Proceedings of the IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS'06)*.

Abstract: This paper presents a proposal for a manufacturing upper ontology, aimed to draft a common semantic net in manufacturing domain. Usefulness of ontologies for data formalization and sharing, especially in a manufacturing environment, are first discussed. Details are given about the Web Ontology Language (OWL) and its adequation for ontologies in the manufacturing systems is shown. A concrete proposal named MASON (MANufacturing's Semantics ONtology) is presented and two applications of this ontology are exposed: automatic cost estimation and semantic-aware multiagent system for manufacturing.

Lugar RG. June 2005. *The Lugar Survey Proliferation Threats and Responses*. Washington, DC.

Abstract: Since the fall of the Soviet Union, vulnerability to the use of weapons of mass destruction has been the number one national security dilemma confronting the United States and much of the world. After many years, the events of September 11, 2001, and the subsequent public discovery of al-Qaeda's methods, capabilities, and intentions brought our vulnerability to the forefront. The War on Terrorism proceeds in a world awash with nuclear, chemical and biological weapons and materials. Most of these weapons and materials are stored in the United States and Russia, but they also exist in India, Pakistan, Iran, Libya, North Korea, Syria, Sudan, Israel, Great Britain, France, China, and perhaps other nations. We must anticipate that terrorists will use weapons of mass destruction if allowed the opportunity. The minimum standard for victory in this war is the prevention of any terrorist cell from obtaining weapons or materials of mass destruction. We must make certain that all sources of WMD are identified and systematically guarded or destroyed. With this in mind, I am hopeful that this study will contribute to the discussion inside and outside of governments about how we can strengthen non-proliferation efforts, improve safeguards around existing weapons and materials, bolster

intelligence gathering and interdiction capabilities, and expand international cooperation in dealing with a threat that should deeply concern all governments and peoples.

Lyons MM, MP Hall, J Wilson, S Saxena and E Sherman. Winter 2005. *Export Compliance Prevention and Enforcement*.

Abstract: Companies that wish to import or export goods and technology to or from the United States face a complex regulatory landscape. The import/export laws have garnered even greater attention from federal authorities in recent years, after their importance was heightened by the events of September 11, 2001. This article focuses on the export side, where businesses must comply with the dictates of three separate executive departments: Commerce, Treasury and Justice. It discusses the different compliance and enforcement roles of these agencies, highlights applicable laws and regulations and outlines the potential civil and criminal penalties imposed for their violation.

Masolo C, S Borgo, A Gangemi, N Guarino, A Oltramari and L Schneider. 2003. *The WonderWeb Library of Foundational Ontologies Preliminary Report*. ISIB-CNR (Institute for Biomedical Engineering), Padova, Italy.

Abstract: Ontologies are the basic infrastructure for the Semantic Web. Everybody agrees on this, as the very idea of the Semantic Web hinges on the possibility to use shared vocabularies for describing resource content and capabilities, whose semantics is described in a (reasonably) unambiguous and machine-processable form. Describing this semantics, i.e. what is sometimes called the intended meaning of vocabulary terms, is exactly the job ontologies do for the Semantic Web. But what kinds of ontologies do we need? This is still an open issue. Some people believe that upper level ontologies are important, others think they are a waste of time, and prefer to concentrate on lightweight ontologies, focusing on the minimal terminological structure (often just a taxonomy) which fits the needs of a specific community. The point is that ontologies can be used in different ways within the Semantic Web. On one hand, for instance, they can be used for semantic access to a specific resource; in this case the intended meaning of a single term is more or less known in advance, and the ontology can be limited to those structural relationships among terms which are relevant for the query (in many cases, taxonomic relationships are enough). On the other hand, ontologies can be used to negotiate meaning, either for enabling effective cooperation between multiple artificial agents, or for establishing consensus in a mixed society where artificial agents cooperate with human beings. This is a completely different task for ontologies, which requires the explicit representation of ontological commitment in terms of a rich axiomatization. The axiomatization's purpose is to exclude terminological and conceptual ambiguities, due to unintended interpretations. In general, meaning negotiation is of course an extremely hard task (both conceptually and computationally), but it only needs to be undertaken once, before a cooperation process starts. The quality of a meaning negotiation process may drastically affect the trust in a service offered by the Semantic Web, but not the computational performance of the service itself. For example, a product procurement process involving multiple agents with distributed lightweight ontologies may be carried out in an efficient way by using simple terminological services, but the risk of semantic mismatch can be minimized only if the agents share a (more or less minimal) common ontology.

McMaster RB and EL Usery, Eds. 2005. *A Research Agenda for Geographic Information Science*. CRC Press, Boca Raton, FL.

Abstract: A close relationship exists between GIS and numerous applications, including cartography, photogrammetry, geodesy, surveying, computer and information science, and

statistics, among others. Scientists coined the term "geographic information science (GIScience)" to describe the theory behind these fields. A Research Agenda for Geographic Information Science extensively details the issues and fundamental scientific problems that must be solved if the use of GIS in these and other fields is to advance. Immediately following the founding of the University Consortium for Geographic Information Science (UCGIS), the group identified in a Research Agenda the topics that represented major challenges to the GIS research community. The first chapter of this book delivers an introduction to the agenda and to the collective guidance that the agenda provides to researchers. Chapters 2-10 discuss nine original research challenges. Chapters 11-14 provide the basis of the agenda's four "Emerging Themes." Each chapter, written by researchers involved in the conception of the topics, discusses basic research elements, the UCGIS approach, the need for the National research agenda, contributions to knowledge and society, and offers a complete set of references. The final section draws general conclusions about the UCGIS approach and the defined research challenges.

Mendes O and A Abran. 2004. "Software Engineering Ontology: A Development Methodology." *Metrics News* 9.

Abstract: According to Gruber's definition an ontology [1] is "a formal specification of a conceptualization". A conceptualization being a simplified, abstract way of perceiving a segment of the world (a piece of reality), for which we agree to recognize the existence of a set of objects and their interrelations, as well as the terms we use to refer to them and their agreed meanings and properties. Thus, ontologies represent a consensual, shared description of the pertinent objects considered as existing in a certain domain of knowledge (the domain of discourse). They constitute a special kind of software artifact conveying a certain view of the world (conceptualization), specifically designed with the purpose of explicitly expressing the intended meaning of a set of agreed existing objects. Ontologies could play an important role in Software Engineering, as they do in other disciplines, where they: 1) provide a source of precisely defined terms that can be communicated across people, organizations and applications (information systems or intelligent agents); 2) offer a consensual shared understanding concerning the domain of discourse; 3) to render explicit all hidden assumptions concerning the objects pertaining to a certain domain of knowledge [2]. Despite some initial efforts to develop partial (sub domain) ontologies [3] [4] [5] [6], as a field of knowledge, Software Engineering still does not have a comprehensive detailed ontology which describes the concepts that domain experts agree upon, as well as their terms, definitions and meanings. Such ontology would also need to look at the more pertinent interrelations where concepts participate in the creation of the semantic network in which they are inserted. The development of a "software engineering domain ontology" will allow us to: 1) share and reuse all knowledge accumulated until now in the Software Engineering field; 2) open news avenues to automatic interpretation of this knowledge, using information systems or intelligent software agents.

Monch L and M Stehli. 2003. "An Ontology for Production Control of Semiconductors Manufacturing Processes." In *MATES 2003*, LNAI 2831, pp. 156-167.

Abstract: In this paper, we describe an ontology for a hierarchically organized production control system in semiconductor manufacturing. The semiconductor manufacturing domain is characterized by reentrant product flows, sequence dependent setup-times, prescribed due-dates, a diverse product mix, a mix of different process types including batch processes and preventive maintenance issues because of complex technological processes. Starting from a hierarchical decomposition of the manufacturing system, we use an agent-based architecture for implementing the resulting production control system. In order to coordinate the autonomous entities of the

hierarchy, we suggest an ontology that is appropriate to the hierarchical control approach. We illustrate the use of the suggested ontology.

Mönch L and J Zimmermann. 2008. "An Ontology to Support Adaptive Agents for Complex Manufacturing Systems." In *32nd Annual IEEE International Computer Software and Applications, COMPSAC '08.* , pp. 531-536.

Abstract: In this paper, we extend the FABMAS ontology towards its usage in adaptive production control systems for complex manufacturing systems. FABMAS is a hierarchically organized multi-agent-system for production control in the semiconductor manufacturing domain. It consists of decision-making and staff agents. Staff agents support the decision-making agents. We introduce adaptive staff agents. Adaptive staff agents make sure that the parameters of production control algorithms are properly adjusted in a situation-dependent manner. We derive necessary extensions of the FABMAS ontology based on the architecture of adaptive staff agents and their interaction with other agents of the manufacturing control system. We present an example for using the ontology within FABMAS.

Mowatt-Larssen R. April 2, 2008. *Statement of Rolf Mowatt-Larssen Director of The Office of Intelligence and Counterintelligence United States Department of Energy Before the Homeland Security and Governmental Affairs Committee United States Senate.*

Abstract not available.

No author listed. June 2005. *Suggested Reading List for New Analysts.*

Abstract not available.

A list of references for new geospatial intelligence analysts compiled in November 2005. An interesting collection of potential resources including books, journal articles, and online databases and journals

No author listed. n.d. "Bibliography - Imagery Analysis."

Abstract not available.

Bibliography and Appendix for an unknown document. Perhaps from Frank Pabian's Commercial Satellite Imagery, Another Tool in the Nonproliferation Verification and Monitoring Toolkit chapter of Nuclear Safeguards, Security, and Nonproliferation: Achieving Security with Technology and Policy edited by James Doyle.

No author listed. n.d. "Presentation, Iran Peaceful Use of Nuclear Energy?"

Abstract not available.

A high-level presentation discussing Iran's Nuclear Program. Similar in content to the U.S. Department of Energy and National Nuclear Security Administration (2006) Presentation on Iran's Nuclear Program

Noy NF and DL McGuiness. 2000. *Ontology 101: A Guide to Creating Your First Ontology*. KSL-01-05, Stanford Knowledge Systems Laboratory.

Abstract: Ontologies have become core components of many large applications yet the training material has not kept pace with the growing interest. This paper addresses the issues of why one would build an ontology and presents a methodology for creating ontologies based on declarative knowledge representation systems. It leverages the two authors experiences building and maintaining ontologies in a number of ontology environments including Protege-2000, Ontolingua, and Chimaera. It presents the methodology by example utilizing a tutorial wines knowledge base example. While it is aimed at users of frame-based systems, it can be useful for building ontologies in any object-centered system.

Obrst L. 2008. "Challenges of Spatial Semantics: Nearness, Parts, Connectivity, Vagueness." In *Spatial Ontology CoP Workshop Presentation*. October 17, 2008.

Abstract not available.

Office of the Under Secretary of Defense For Acquisition Technology and Logistics. 2004. *Report of the Defense Science Board Task Force on Preventing and Defending Against Clandestine Nuclear Attack*. Washington, DC.

Abstract: The DSB addressed this threat in previous studies conducted in 1997 (also chaired by Richard Wagner) and 1999/2000 (chaired by Roger Hagenruber). Much has changed since then. The 11 Sept. 2001 attacks demonstrated the intent of terrorists to inflict massive damage. Nuclear proliferation has proceeded apace, with North Korea and Iran achieving nuclear weapon capability or coming closer to it, and it could spread further. The United States is engaged in a war against terrorism, and DoD is beginning to devote significant effort to combating WMD. The Department of Homeland Security (DHS) has been established. Thinking about the threat of clandestine nuclear attack has changed, and some efforts to explore defenses have begun. However, one thing has not changed: little has actually been done against the threat of clandestine nuclear attack. The DSB Summer Study on Transnational Threats (1997) first developed the ambitious idea of a very large, multi-element, global, layered civil/military system of systems of scope sufficient to have some prospect of effectively thwarting this threat. There was little resonance with this vision (outside of the Task Forces in 1997 and 2000), but since then, and especially since the attacks of 11 Sept. 2001, it has begun to be discussed more widely. This report will revisit such a national/global system, largely as context for the main focus of the Task Force: DoD's roles and capabilities. Following briefings from many government agencies and subject matter experts, the Task Force arrived at its basic findings and recommendations in early 2003. Since then, those results have been discussed in over 40 meetings within DoD and elsewhere, leading to certain refinements. This report reflects the outcomes of that process and weaves together viewgraphs used in the discussions with elaborating text.

Pabian F. 2008. "Chapter 12 - Commercial Satellite Imagery: Another Tool in the Nonproliferation Verification and Monitoring Toolkit." In *Nuclear Safeguards, Security, and Nonproliferation*, pp. 221-248 ed: JE Doyle. Elsevier Inc.

Abstract: This chapter highlights the evolutionary and revolutionary role that commercial satellite imagery (as another independent "open source") is playing on the international stage in providing a heretofore unimaginable basis for greater global transparency and the way it has helped, and will continue to help, detect and monitor undeclared nuclear-related facilities and activities. In that role, such imagery serves as an enabler of transparency for nation states

previously lacking their own “ national technical means ” (NTM, a euphemism for “ spy satellites ”), sometimes even being referred to as a “ Poor Man’s NTM, ” 1 but also, perhaps more significantly, for multinational organizations, nongovernmental organizations, the news media, academics, or even just plain individual interested parties or hobbyists. The opportunity now exists for anyone with relatively modest computer access (and a credit card for supplemental imagery purchases when necessary) to view, identify, and monitor nuclear facilities and associated activities that have the potential to help threaten international peace and security and that might otherwise have remained clandestine.

Parisi SB, Jochen; Berssenbrugge, Jan; Radkowski, Rafael. 2007. "Using Ontology to create 3D Animations for Training Purposes." *International Journal of Software Engineering and Its Applications* 1(1):67-78.

Abstract: 3D Computer-based animations are nowadays used for training purposes in a wide range of industrial applications like assembly, maintenance and operations. Their creation, also known as authoring, is usually a time consuming task, demanded to professional 3D designers, who need at first a good understanding of the involved entities and actions in order to realize the customized animations. The proposed methodology deals with the use of an ontology in order to filter and understand generic natural language training requests; once identified the proper actors and actions, they are associated to the corresponding models and movements to be performed in the virtual environment and translated in a 3D graphic format template. The result is a customized animation which can be created by a non-expert designer and then visualized by the worker through many modalities, on a desktop computer or on a PDA for example. Role of the ontology is to reduce the overall complexity of the animation authoring process by assuring the necessary comprehension of customized training requests as well as reusability and extensibility of the structure of the modeled object and of animations’ components in different domains.

Plant R and G Antoniou. 1999. "Editorial Special issue on verification and validation." *Knowledge-Based Systems* 12:1-2.

Abstract: The ability of software developers to be certain of the correctness of the code and systems that they create is central to their task. It is the universal focus that spans all techniques, styles and development methodologies employed by software and knowledge engineers. Assuring the reliability and correctness of knowledge-based systems has become an important focus within software engineering, as these systems become ever more sophisticated and increasingly embedded within other systems and environments. Through Validation and Verification (commonly referred to as V&V), knowledge engineers attempt to create techniques to ensure that systems meet the users requirements (validated) and that the systems execute the requirements of the specification correctly (verified). A central exchange point for research within the V&V community has been the annual workshops at AAAI held since 1988. The purpose of this special issue is to provide an overview of the state of the art in KBS V&V, through a collection of six papers from the 1997 AAAI workshop held in July at Providence, RI, USA.

Poucet A. 2006. "Arms Control and Non-Proliferation Treaties: An Ontology of Concepts and Characteristics." In *Verifying Treaty Compliance: Limiting Weapons of Mass Destruction and Monitoring Kyoto Protocol Provisions*, pp. 41-60 eds: Avenhaus, R., N Kyriakopoulos, M Richard and G Stein. Springer-Verlag, Berlin.

Abstract: To reduce the risk and proliferation of arms, particularly weapons of mass destruction, a number of international treaties and other agreements and export control regimes have been concluded. Their purpose is to reduce or eliminate certain weapons or weapon systems, to curb

the proliferation of weapons and of sensitive, dual-use technologies, or to increase security and build confidence in other ways. Some of these treaties are multilateral and intended to cover all States, such as the Treaty on the Non-Proliferation of Nuclear Weapons (NPT), the Conventions on Biological (BWC) and Chemical Weapons (CWC), the Comprehensive Nuclear Test Ban Treaty (CTBT). Others are applicable to a particular region, such as the Conventional Forces in Europe Treaty (CFE), the Open Skies Treaty, the Antarctic Treaty, the Treaty of Tlatelolco, covering Latin America, and other nuclear weapon-free zone treaties. Some others are bilateral at origin, such as the Intermediate Nuclear Forces Treaty (INF) and the Strategic Arms Reduction Treaty (START). In [1] an extensive overview of arms control treaties is presented and their characteristics analyzed. Much of the analysis in that paper served as a basis for this chapter. In the following sections, the most important nonproliferation treaties, NPT, CWC, BWC and CTBT are presented in detail. Also the CFE, an arms control treaty, and the Kyoto protocol (an example of verification to protect the environment) are presented in detail.

Probst F. 2008. "Observations, Measurements and Semantic Reference Spaces." *Applied Ontology* 3(1-2):63-89.

Abstract: What is needed to enable communication about observation and measurement results in information systems? Information system ontologies make a certain conceptualization explicit and partially account for the meanings of symbols associated with this conceptualization. Yet, the meaning of signs denoting measurement results such as "10 m", "red" or "high" cannot be specified with currently available ontologies. They fail to separate the ontological nature of some observable quality from the specification of how to observe and name the measurement result. We employ the foundational ontology DOLCE for characterizing the ontological nature of observable magnitudes. This involves dealing with ontological questions like "What kinds of observable qualities exist, in which entity does the observed quality inhere and how are the magnitudes of the observed quality structured?". Then, in order to capture the semantic aspects of an observation result, we introduce semantic reference spaces, which help deal with semantic questions like "Do the signs "10 m", "33 feet" or "shallow" have the same meaning? Do these signs refer to the same entity, e.g. the depth magnitude of a lake? How to establish a unit of measure?". We posit that the semantic questions can be approached efficiently only if agreement is reached on the ontological questions, and show that the specification of the meaning of signs denoting measurement results is enabled via the extension of the foundational ontology DOLCE with semantic reference spaces.

Pumphrey CW. 2000. *Transnational Threats: Blending Law Enforcement and Military Strategies*. U.S. Army War College, Carlisle, PA.

Abstract: On February 2-3, 2000, the U.S. Army War College, the Triangle Institute for Security Studies, and the Duke University Center for Law, Ethics, and National Security co-sponsored a conference in Chapel Hill, North Carolina. The conference examined transnational threats, including terrorism involving weapons of mass destruction, cyber threats to the national infrastructure, and international organized crime. The goal was to evaluate the seriousness of such threats and discuss strategies for dealing with them. In particular, the conference sought to address the question of how military and law enforcement could blend their strategies to better counter transnational threats. A secondary purpose was to clarify the role of the military in meeting challenges that transcend national borders and threaten our national interests. This book highlights some of the main issues and themes that ran through the conference. After looking at the various threats and undertaking a risk assessment, the book considers the unique aspects of transnational threats, and then identifies the key challenges facing the United States, paying particular attention to the role of the military. To conclude, the book discusses some of the steps

that should be taken to secure ourselves against transnational threats. The Strategic Studies Institute is pleased to publish this volume as a contribution to the debate on important national security issues.

Ressler J and M Dean. 2007. "Geospatial Ontology Trade Study." In *Ontology for the Intelligence Community (OIC-2007)*. November 28-29, 2007, Columbia, Maryland.

Abstract: This short paper summarizes a survey of ontologies relevant to geospatial intelligence. 45 geospatial and temporal ontologies, in 11 categories, were assessed against 3 use cases: annotation, qualitative reasoning, and information integration. Specific recommendations and more general conclusions are provided. The paper presents an illustration of a feature with several different ontology representations.

Schlenoff CI, RW Ivester, D Libes, PO Denno and S Szykman. 1999. *An Analysis of Existing Ontological Systems for Applications in Manufacturing and Healthcare*. NIST Interagency/Internal Report (NISTIR) - 6301, National Institute of Standards and Technology.

Abstract: The objective of the work described in this paper is to move closer to the ultimate goal of seamless system integration using the principle behind ontological engineering to unambiguously define domain-specific concepts. Current integration efforts are usually based solely on how information is represented (the syntax or terminology) without a description of what the information means (the semantics). This paper documents the results of the first phase of this project, that of analyzing existing ontological systems to determine which is most appropriate for the manufacturing and healthcare domains. In particular, this phase involved the exploration of efforts that are studying both the uses of ontologies in the general sense and those that are using ontologies for domain specific purposes.

Shankar M, A Sorokine, B Bhaduri, D Ressiguie, S Shekhar and JS Yoo. "Spatio-Temporal Conceptual Schema Development for Wide-Area Sensor Networks." In *GeoS'07 - 2nd international Conference on GeoSpatial Semantics*.

Abstract: A Wide-Area Sensor Network (WASN) is a collection of heterogeneous sensor networks and data repositories spread over a wide geographic area. The diversity of sensor types and the regional differences over which WASNs operate result in semantic interoperability mismatches among sensor data, and a difficulty in agreeing on methods for sensor data access and exchange. We assume that sensors and their associated data have an explicit spatio-temporal basis (or tagging) in their representation. In this paper, we describe a spatio-temporal loosely coupled federated database model for the WASN data storage problem - that of unifying query and data representation given a heterogeneous WASN - and propose a conceptual schema to ease the problem of integration of sensor data representations. This is a continuing and critical challenge as sensor networks become more ubiquitous and data interoperation becomes increasingly vital for a variety of applications (such as homeland security, transportation, environmental monitoring, etc.). We employ a top-down ontology-driven software development methodology. We use the SNAP/SPAN ontology as a sample framework for the conceptual schema. We compare our methodology of conceptual schema development with a bottom-up entity-oriented schema construction and discuss the differences in the two approaches. A unique contribution is the discussion of deployment experiences to evaluate proposed approaches in the context of a concrete WASN testbed.

Smith B and DM Mark. 2003. "Do Mountains Exist? Towards an Ontology of Landforms." *Environment and Planning B: Planning and Design* 30(3):411-427.

Abstract: The authors begin the paper with the question 'Do mountains exist?' They show that providing an answer to this question is surprisingly difficult and that the answer that one gives depends on the context in which the question is posed. Mountains clearly exist as real correlates of everyday human thought and action, and they form the archetype for geographic objects. Yet individual mountains lack many of the properties that characterize bona fide objects, and 'mountains' as a category also lacks many of the properties that characterize natural kinds. In the context of scientific modeling of the environment, especially of such phenomena as surface hydrology and fluvial erosion and deposition, mountains are not picked out as constituents of reality in their own right at all; rather, they are just parts of the field of elevations whose gradients shape the direction of runoff and influence the intensity of erosion. Thus, although an object-based ontology of mountains and other landforms is required to do justice to our everyday conceptions of the environment and to support spatial reasoning and natural language processing, topographic databases designed to support environmental modeling can be field-based at geographic scales.

Sokolski H. 2000. *Prevailing in a Well-Armed World: Devising Competitive Strategies Against Weapons Proliferation*. U.S. Army War College, Carlisle, PA.

Abstract: The Strategic Studies Institute is pleased to publish *Prevailing In A Well-Armed World: Devising Competitive Strategies Against Weapons Proliferation*. This work provides insights into the competitive strategies methodology. Andrew Marshall notes that policymakers and analysts can benefit by using an analytical tool that stimulates their thinking—more directly—about strategy in terms of long-term competition between nations with conflicting values, policies, and objectives. Part I of this work suggests that the competitive strategies approach has value for both the practitioner and the scholar. The book also demonstrates the strengths of the competitive strategies approach as an instrument for examining U.S. policy. The method in this book focuses on policies regarding the proliferation of weapons of mass destruction. In “shaping” the international environment in the next millennium, no other national security issue seems as complex or important. The imperative here is to look to competitive strategies to assist in asking critical questions and thinking broadly and precisely about alternatives for pitting U.S. strengths against opponents’ weaknesses. Part II uses the framework to examine and evaluate U.S. nonproliferation and counterproliferation policies formed in the final years of the 20th century. In Part III, the competitive strategies method is used to analyze a regional case, that of Iran. The insights contained in this book provide an opportunity to pause and consider alternative and innovative approaches to strategic thinking and proliferation policy. The Strategic Studies Institute is pleased to offer this book to assist practitioners and scholars in thinking strategically about U.S. defense policies and priorities.

Sokolski HD and T Riisager. 2002. *Beyond Nunn-Lugar: Curbing the Next Wave of Weapons Proliferation Threats from Russia*. Nonproliferation Policy Education Center.

Abstract: Although the present book never intended to be quite so timely, *Beyond Nunn-Lugar: Curbing the Next Wave of Weapons Proliferation Threats from Russia* is one book that, coming so soon after the events of September 11, 2001, and shortly before President Bush's Russian summit, should find a ready audience. The essays in the book were originally commissioned by the Nonproliferation Policy Education Center (NPEC) as part of a year-long study on the future of U.S.-Russian nonproliferation cooperation. What makes the book different from other studies of U.S.-Russian cooperation is its reliance on competitive strategies. Originally devised as a

business management tool and subsequently used by the Pentagon to guide its military planning against the Soviet Union, the analytical approach based on competitive strategies requires analysts to evaluate long-term trends, 10 to 20 years out. Rather than beginning with one's own aims and strategies, however, competitive strategies demands that analysts first understand the objectives, strengths, and weaknesses of competitors and other key parties. A key objective of competitive strategies is to detail how best to pit one's strengths against a competitor's enduring weaknesses in a series of moves and countermoves. The goal is to devise strategies that force one's competitor to spend more time and resources shoring up his weaknesses than in taking offensive action.

Sorokine A. n.d. "Ontology-Based Feature Extraction from Remotely Sensed Imagery" (white paper for NGA).

Abstract: This white paper is a result of a series of informal conversations with Dan Adams of NGA during a SensorNet kick-off meeting in Washington, DC on July 25th, 2005 and GeoComputation 2005 conference in Ann Arbor, MI on August 1–3rd, 2005. Most of currently available remote sensing methods are geared toward analysis of multi- and hyper-spectral data, texturing, . . . and largely ignore composition of features as a method for image interpretation and feature detection. Typically methods mis-detect or fail to detect features or arrangements of features that are easily recognized by a human image interpreter. Modern remote sensing methods are good at quantitative tasks of imagery interpretation such as, for example, determining the zone of oil spill, estimating crop damage and location of fires. However, qualitative image interpretation tasks such as, determining the presence of a larger object constituted of smaller ones, or guessing the presence of an object judging from its surrounding are beyond their capabilities. The goal of this project is to develop an automated software system that would be capable of extracting such complex geographic features from multiple sources of data with remote sensing imagery being the primary source.

Sorokine A, T Bittner and C Renschler. 2006. "Ontological Investigation of a Multiscale Ecosystem Classification Using the "National Hierarchical Framework of Ecological Units" as an Example." *GeoInformatica* 10(3):313-335.

Abstract: This paper presents a formalized ontological framework for the analysis of multiscale classifications of geographic objects. We propose a set of logical principles that guide such geographic classifications. These principles are represented as definitions and axioms of the first-order predicate logic. Then we demonstrate an application of these principles with a practical example of the "National Hierarchical Framework of Ecological Units." The study presents a minimal set of entities and relation need for such formalizations and emphasizes the need to distinguish clearly between classes of ecosystems and ecosystems-individuals physically located on the Earth surface. The ontological framework proposed in the paper has the potential to be used to facilitate interoperability between such geographic classifications.

Spatial Ontology Community of Practice. n.d. "SOCop Introduction to Ontologies."

Abstract not available.

A presentation that provides an introduction to the Spatial Ontology Community of Practice. Includes discussions on the value of ontologies, an overview of the CoP, an introduction to ontologies and an example.

Squassoni S. Winter 1999. "Safeguards and Satellite Imagery Potential Applications." *JNMM* 27(2):35-39.

Abstract: Commercial satellite imagery provides the International Atomic Energy Agency with yet another vantage point from which to monitor a state's nuclear activities. The most practical uses of remote sensing for safeguards include serving as an inspection aid, detecting changes at certain types of facilities and sites, and corroborating state declarations on the nuclear fuel cycle and nuclear-related activities as well as open-source information. Although the IAEA's experience with member state information may lead some observers to conclude the high-resolution satellite imagery could yield a significantly better ability to detect undeclared activities, a realistic assessment must conclude that a credible capability for wide-area monitoring with satellite images is not cost-effective. The agency and member states will need to balance carefully the value added by CSI with costs it would entail.

Stasch C, K Janowicz, A Bröring, I Reis and W Kuhn. 2009. "A Stimulus-Centric Algebraic Approach to Sensors and Observations GeoSensor Networks." pp. 169-179 eds: N Trigoni, A Markham and S Nawaz. Springer Berlin / Heidelberg. 5659.

Abstract: The understanding of complex environmental phenomena, such as deforestation and epidemics, requires observations at multiple scales. This scale dependency is not handled well by today's rather technical sensor definitions. Geosensor networks are normally defined as distributed ad-hoc wireless networks of computing platforms serving to monitor phenomena in geographic space. Such definitions also do not admit animals as sensors. Consequently, they exclude human sensors, which are the key to volunteered geographic information, and they fail to support connections between phenomena observed at multiple scales. We propose definitions of sensors as information sources at multiple aggregation levels, relating physical stimuli to observations. An algebraic formalization shows their behavior as well as their aggregations and generalizations. It is intended as a basis for defining consistent application programming interfaces to sense the environment at multiple scales of observations and with different types of sensors.

Supreme Court of the State of New York. 2008. "The People of the State of New York Against Li Fang Wei."

Taylor JM. 2009. "Computational Detection of Humor: A Dream or a Nightmare? The Ontological Semantics Approach." In *2009 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, pp. 429-432. 2009. DOI 10.1109/wi-iat.2009.318.

Abstract: This paper deals with computational detection of humor. It assumes that computational humor is a useful task for any number of reasons and in many applications. It discusses the computational linguistic/ semantic preconditions for computational humor and an ontological semantic approach to the task of humor detection, based on direct and comprehensive access to meaning rather than on trying to guess it through statistical-cum-syntactical keyword methods. The paper is informed by the experience of designing and implementing a humor detection model, whose decent success rate confirmed some of the assumptions while its misses made other ideas prominent, including the necessity of full text comprehension. The bulk of the paper explains how the comprehensive representation of meaning and, most importantly how unstructured natural language text is automatically translated into the ontologically defined text meaning representations that can be used then to detect humor in them, if any, automatically.

Thomas TS and SD Kiser. 2002. *Lords of the Silk Route: Violent Non-State Actors in Central Asia*. INSS Occasional Paper 43, USAF Institute for National Security Studies, USAF Academy, CO.

Abstract: We are pleased to publish this forty-third volume in the Occasional Paper series of the United States Air Force Institute for National Security Studies (INSS). This paper, while it reports the results of research undertaken across the year prior to the events of September 11 and their aftermath, presents an analysis that is both timely and relevant given those events. This important paper represents the kind of original thinking that this Institute was designed in the hope of fostering. The two authors—each of whom is individually the winner of a previous INSS outstanding research award—develop and test a systematic, targeted, and useful methodology for examining the non-state political violence and its practitioner that the United States now faces. Their analysis also is grounded in Central Asia, a new but increasingly important region to United States military interest and presence. The paper stands well on either of those legs—a systematic methodology for violent non-state actors or a detailed and security oriented examination of an emerging critical region. Taken together, the two legs mark it as a singularly significant work, one well worthy of serious study.

Truong QS, R Keefe, P Barnes and JP Paquette. Winter 1999. "Potential Application of Commercial Satellite Imagery in International Safeguards." *JNMM* 27(2):13-18.

Abstract: This paper describes an investigation of the potential use of commercial satellite imagery in the International Atomic Energy Agency's Strengthened Safeguards System. The work is carried out to provide input to IAEA internal discussions regarding potential applications of satellite imagery, including identification of the type of expertise requires to support such activity. The main aims of the study are to identify and to demonstrate, with concrete examples, potential used of satellite imagery in international safeguards. The paper reports on the effect of image resolution, in the range of 120 m (infrared) to simulated 1 m (panchromatic); the use of various spectra (microwave, visible and thermal infrared); some aspects of change detection; and detecting undeclared facilities. We use examples that include a multiunit nuclear complex, a nuclear research site, single-unit nuclear generating stations and mine sites. A draft report documenting the results to date has been completed and submitted to the IAEA. This report will be revised for confirmation of the findings when high-resolution satellite imagery in the range of 1 m becomes available, hopefully in early 1999.

Tudorache T, NF Noy, S Tu and MA Musen. 2008. "Supporting Collaborative Ontology Development in Protege." In *7th International Semantic Web Conference*, pp. 17-32 Karlsruhe, Germany Springer.

Abstract: Ontologies are becoming so large in their coverage that no single person or a small group of people can develop them effectively and ontology development becomes a community-based enterprise. In this paper, we discuss requirements for supporting collaborative ontology development and present Collaborative Protégé—a tool that supports many of these requirements, such as discussions integrated with ontology-editing process, chats, and annotations of changes and ontology components. We have evaluated Collaborative Protégé in the context of ontology development in an ongoing large-scale biomedical project that actively uses ontologies at the VA Palo Alto Healthcare System. Users have found the new tool effective as an environment for carrying out discussions and for recording references for the information sources and design rationale.

U.S. Department of Defense. 1995. *Mapping, Charting & Geodesy Glossary of Feature and Attribute Definitions*. MIL-STD-2408(DMA).

Abstract: This Military Standard, MIL-STD-2408, is a specialized dictionary wherein physical and conceptual earth objects entities and their descriptive characteristics/properties are defined as Mapping, Charting, and Geodesy (MC&G) features and their associated attributes. These features and attributes are codified by the use of the Feature/Attribute Coding Standard (FACS), which is the internal coding schema used by the Defense Mapping Agency's (DMA's) Digital Production System (DPS) to produce numerous hardcopy and digital products.

U.S. Department of Energy and National Nuclear Security Administration. 2006. "Iran's Nuclear Program."

Abstract not available.

A brief, high-level presentation about Iran's Nuclear Program.

United States Department of State, Bureau of Nonproliferation and Office of Export Control Cooperation. June 2-4, 2003. "Weapons of Mass Destruction Proliferation Threat." In *Central Asia and Caucasus Regional Forum*, Almaty, Kazakhstan.

Abstract not available.

This brief, high-level presentation discusses chemical, biological and nuclear WMD proliferation threats and dual use materials and how proliferators might attempt to obtain these materials. It also includes two case studies.

Uschold M and M Gruninger. 1996. "Ontologies: Principles, Methods and Applications." *Knowledge Engineering Review* 11(2):93-136.

Abstract: This paper is intended to serve as a comprehensive introduction to the emerging field concerned with the design and use of ontologies. We observe that disparate backgrounds, languages, tools and techniques are a major barrier to effective communication among people, organizations and/or software understanding (i.e. an "ontology") in a given subject area, can improve such communication, which in turn, can give rise to greater reuse and sharing, interoperability, and more reliable software. After motivating their need, we clarify just what ontologies are and what purpose they serve. We outline a methodology for developing and evaluating ontologies, first discussing informal techniques, concerning such issues as scoping, handling ambiguity, reaching agreement and producing definitions. We then consider the benefits and describe a more formal approach. We re-visit the scoping phase, and discuss the role of formal languages and techniques in the specification, implementation and evaluation of ontologies. Finally, we review the state of the art and practice in this emerging field, considering various case studies, software tools for ontology development, key research issues and future prospects.

Varzi A. 2007. "Spatial Reasoning and Ontology: Parts, Wholes, and Locations." In *Handbook of Spatial Logics*, pp. 945-1038. eds: M Aiello, I Pratt-Hartmann and J van Benthem. Springer-Verlag, Berlin.

Abstract: Spatial reasoning is no abstract business. It is, to a great extent, reasoning about entities located in space, and such entities have spatial structure. If the table is in the kitchen, then it

follows that the table top is in the kitchen, and it follows because the top is part of the table. If the concert took place at the stadium, then it didn't take place in the theater, for concerts are spatially continuous. Even when we reason about empty places, we typically do so with an eye to the anatomy of their potential tenants: space as such is perceptually remote and we can hardly understand its structure without imagining what could fill the void. This tension (not to say this ambiguity) between concrete object-oriented thinking and abstract space-oriented thinking is responsible for many of the philosophical issues that lie behind any formal theory of spatial reasoning. The purpose of this chapter is to take a closer look at these delicate matters. Rather than doing this in general, however, we shall look at how the subtle interplay between purely spatial intuitions and intuitions about concrete spatial entities shows up in the construction of a formal theory. More specifically, we shall consider three sorts of theory, each of which occupies a prominent position in recent literature: (1) mereology, or the theory of parthood relations; (2) topology, broadly understood as a theory of qualitative spatial relations such as continuity and contiguity; and (3) the theory of location proper, which deals explicitly with the relationship between an entity and the spatial region it occupies.

Wei-Tek T, R Vishnuvajjala and Z Du. 1999. "Verification and validation of knowledge-based systems." *Knowledge and Data Engineering, IEEE Transactions on* 11(1):202-212.

Abstract: Knowledge-based systems (KBSs) are being used in many applications areas where their failures can be costly because of losses in services, property or even life. To ensure their reliability and dependability, it is therefore important that these systems are verified and validated before they are deployed. This paper provides perspectives on issues and problems that impact the verification and validation (V&V) of KBSs. Some of the reasons why V&V of KBSs is difficult are presented. The paper also provides an overview of different techniques and tools that have been developed for performing V&V activities. Finally, some of the research issues that are relevant for future work in this field are discussed

Willrich M and TB Taylor. 1974. *Nuclear Theft Risks and Safeguards*. The Ford Foundation, Cambridge, MA.

Abstract: The Energy Policy Project was initiated by the Ford Foundation in 1971 to explore alternative national energy policies. This book, one of the series of studies commissioned by the Project, is presented here as a timely and carefully prepared contribution to today's public discussion about nuclear energy. We analyze the possibility of nuclear violence using fissionable material that might be stolen from the U.S. nuclear power industry, and we discuss what can and should be done to prevent that from happening.

Zaitseva L. Fall/Winter 2002. "Illicit Trafficking in the Southern Tier and Turkey since 1999: A Shift from Europe?" *The Nonproliferation Review*:168-182.

Abstract: After the collapse of the former Soviet Union in 1991, illicit trafficking of nuclear and other radioactive material emerged as a serious international concern. The economic and social conditions that followed the collapse left nuclear and radioactive material often poorly guarded and vulnerable to theft. In the early 1990s, Europe observed a sharp increase in nuclear smuggling incidents, as stolen nuclear and other radioactive material was brought from the former Soviet republics to Western Europe in the hope of finding a market. Since 1994, however, reported illicit trafficking incidents in Europe have declined. By contrast, since 1999 there has been an increase in such incidents in the Caucasus, Central Asia, and Turkey. Analysts have long speculated that nuclear smugglers would exploit this region--known as the southern tier--as a transit corridor. This report looks at the illicit trafficking situation in the southern tier and Turkey

in an attempt to establish whether these regions have become new routes for smuggling nuclear and other radioactive material. It discusses reported incidents of illicit trafficking in these countries, assesses their responses to the threat of trafficking, and evaluates foreign assistance provided to the region to combat this smuggling. The report concludes with recommendations for improving international anti-nuclear smuggling efforts.

Zhang H and FN von Hippel. 1999. *The Application of Commercial Observation Satellite Imagery for the Verification of Declared and Undeclared Plutonium Production Reactors*. Report No. 319, Princeton University, Princeton Environmental Institute.

Abstract: Negotiations on a Fissile Material Cutoff Treaty (FMCT) will soon be underway at the Committee on Disarmament in Geneva, and will include detailed attention to how such a treaty could be verified. This paper explains how commercial observation satellites can be effective in verifying that reactors used to produce weapons plutonium in the past are kept in a shutdown status under a cutoff treaty or moratorium. The satellites considered are the new-generation satellites with fine spatial resolution images in the visible and near infrared band complemented by thermal infrared images with lower-spatial resolution but good temperature resolution. These satellites can also contribute to the detection of undeclared nuclear-reactor sites and suspicious construction activities.

Zhao P, L Di, G Yu, P Yue, Y Wei and W Yang. 2009. "Semantic Web-based Geospatial Knowledge Transformation." *Computers & Geosciences* 35(4):798-808.

Abstract: Earth and space science research and applications typically involve collecting and analyzing large volumes of geospatial data much of which is derived from other existing data by applying a scientific workflow. Such a step-by-step process can be viewed as a process of geospatial knowledge transformation, which often involves hypotheses, inferences and integrations to derive user-specific data products from the knowledge of domain experts. Our research is focused on reducing the transformation effort by providing component inference and integration tools. The Semantic Web envisions a new standardized information infrastructure to enable interoperable machine-to-machine interactions and automatic or semi-automatic service chaining for deriving knowledge over networks. This paper describes a generic framework and implementation of how the Semantic Web proceeds through the life cycle of geospatial knowledge transformation, from geospatial modeling (knowledge formalization), through model instantiation (service chain) to model execution (data product). Our approach relies on semantic integrations. A number of ontologies used to capture domain knowledge are introduced in this paper as the basis of knowledge bases for describing and reasoning geospatial data and services. Also, a semantically enabled geospatial catalog service is described to enable more effective discovery, automation and integration of geospatial data and services.

Zheng W, Y Ouyang, J Ford and FS Makedon. 2003. "Ontology-Based Image Retrieval." In *Proceedings of WSEAS*.

Abstract: The current most desirable image retrieval feature is retrieving images based on their semantic content. Currently there are two major image retrieval paradigms that attempt to provide this: text-based metadata image retrieval and content-based image retrieval. In practical applications, both have limitations. In this paper, we discuss an ontology-based image retrieval approach that aims to standardize image description and the understanding of semantic content. Ontology-based image retrieval has the potential to fully describe the semantic content of an image, allowing the similarity between images and retrieval query to be computed accurately.



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnl.gov



U.S. DEPARTMENT OF
ENERGY