

OntoPhoto and the Role of Ontology in Organizing Knowledge

Allen C. Benson

School of Information Sciences, University of Pittsburgh,
135 North Bellefield Ave., Pittsburgh, PA 15260, <acb56@pitt.edu>

Allen Benson is currently a Doctoral Student, School of Information Sciences, University of Pittsburgh. He has directed public and academic library systems since 1990. Benson earned his BFA Degree in music from the University of Minnesota in 1974 and his MLIS degree from The University of Alabama in 1993. He is the author of *The Complete Internet Companion for Librarians* (1995), *Securing PCs and Data in Libraries and Schools* (1997), and *Connecting Kids and the Web* (2003). His dissertation research focuses on content analysis of relationship types expressed in photograph descriptions and image searchers' queries.



Benson, Allen C. *OntoPhoto* and the Role of Ontology in Organizing Knowledge. *Knowledge Organization*, 38(2), 79-95. 32 references.

ABSTRACT: This article is concerned with ontology and its applications in Knowledge Organization (KO) activities. Connections are drawn between efforts in artificial intelligence (AI) to capture the meaning of information and make it accessible to machines and the efforts made in libraries to use KO tools in machine-based record building and search and retrieval systems. The practices used in AI that are of interest here include ontology and ontology-based knowledge representation. In this article their applications in KO are directed towards a particularly problematic document type—the photograph. There are two arguments motivating this article. First, ontology-based KO systems that join AI techniques with library cataloging practices make it possible to utilize higher levels of expressivity when describing photographs. Second, KO systems for photographs that are capable of reasoning over concepts and relationships can potentially provide richer, more relevant search results than systems utilizing word-matching alone.

Received 24 February 2010; Revised 7 May 2010; Accepted 16 July 2010

1.0 Introduction

The goal of this article is to draw connections between efforts in artificial intelligence (AI) to capture the meaning of information and make it accessible to machines and the efforts made in libraries and archives to use Knowledge Organization (KO) tools in machine-based record building and search and retrieval systems. The practices used in AI that are of interest here include ontology and ontology-based knowledge representations. In this article, applications of ontology and knowledge representation in libraries and archives are directed towards a particularly problematic document type, the photograph.

There are two arguments motivating this article. The first is that ontology-based KO systems that join AI techniques with library cataloging practices make

it possible to utilize higher levels of expressivity when describing photographs. The second argument is that KO systems for photographs that are capable of reasoning over concepts and relationships can potentially provide richer, more relevant search results than systems designed for word-matching alone.

I begin by describing some of the problems associated with the traditional models used for representing photographs in library information systems. Following this, I describe the key terms and concepts necessary for understanding the topics introduced in this article—terms commonly found in AI literature, but less often in library-and-information science (LIS) literature. I give particular attention to the notions of ontology, knowledge representation, reasoning, and the distinction made between words and the concepts they represent. Describing an ontology of the photo-

graph in the context of computational ontology entails describing a specific KO system in which ontology-based representations can be stored, queries made and inferences drawn. It also means aligning oneself with a particular upper level ontology. I have adopted Scone as a knowledge base, search and inference engine, upper level ontology and home for an ontology of the photograph. Scone and its creator, Scott Fahlman, are introduced and details surrounding the basic workings of the system are explained.

Next, I explore the notions of concepts, properties, and relation types. I proceed by considering how these can be applied in the domain of the photograph and use these devices to limit the scope of the ontology. Finally, I shift the article's focus to the process of modeling an ontology of the photograph and the questions that must be addressed when determining what a thing must be to be considered a photograph. Within this framework, I propose an alternative, improved system of representation that brings a greater degree of logical and ontological rigor to the cataloging of photographs. This is accomplished by modeling a new knowledge base that extends the Scone upper-level ontology to include a general content theory of the photograph called *OntoPhoto*.

The primary purpose of this article, therefore, is to use the photograph as a focal domain for examining how KO systems can be enhanced through applications of ontology, knowledge representation, and reasoning. There is a critical need to evaluate *OntoPhoto*'s goals and the actions it takes to attain those goals, as well as the ontology's soundness, validity, and its behavior in question and answer processes. Instantiating individual photographs and conducting evaluations of *OntoPhoto* will be the central topics of future papers.

2.0 Background

The organization and representation of photographs in twentieth century archives have focused on the objectives of bibliographic systems, and retrieval has focused on matching indexed terms. It appears we are approaching the limits of what this model can accomplish. More recently, attempts have been made to automatically extract and acquire knowledge about digital photographs through content-based image retrieval (CBIR). While it is desirable in many applications, CBIR is limited to digital images and must overcome several challenges including recognizing image features and their semantics that match higher-level concepts humans tend to use. Knowledge-based

data structures, and, in particular, ontology-based knowledge representation, hold promise for image information organization and retrieval in the twenty-first century. But as the research shows, the concepts and relation types associated with photographs are subtle and complex and present serious challenges for computers to recognize and process effectively.

There are three factors that have held back progress in the effective use of knowledge-based information systems for describing and retrieving photographs. First, there is the problem of recognizing useful applications that require high-level semantics, richer than those provided by existing controlled vocabularies such as *Thesaurus for Graphic Materials* (Parker and Zinkham 1995). Second, there is the problem of "the semantic gap," which Laura Hollink defines as, "the discrepancy between the information that can be derived from the low-level image data and the interpretation that users have of an image" (Hollink 2006, 3-4). Finally, the use of ontology-based representation of images is still in its infancy and focuses primarily on the content of specific images in specialized domains, thus limiting reuse.

One promising application area for the use of rich semantic image descriptions is to improve upon precision and recall currently provided by word matching systems of finding, collocating, and retrieving photographs. A second, even more promising application area is question-answering. That is, the ability to ask questions about history, interpretation, aesthetics, a photograph's relation to other photographs, and the relationships among entities within image content. To address the semantic image gap problem, provide a framework for the question-answer system, and address the problem of reusability, I propose developing an alternative system for cataloging photographs. The new system utilizes ontology-based representations in a knowledge-based environment.

3.0 Computational ontology, knowledge representation, and reasoning

Applying concepts of ontology to KO serves numerous purposes. For example, ontology supports the capture and standardization of terms and their meaning and makes this information broadly accessible to multiple organizations engaged in describing, cataloging, and classifying photographs. The advantages of accessing a shared vocabulary multiply when one considers the benefits of reusability. There are many characteristics specific to the domain of the photograph that are independent of individual photo-

graphic images and these can be shared and reused. Rather than encoding these features in one or more thesauri and duplicating this process within each new setting, an ontology of the photograph collects this common, reusable information into a structured, domain-specific library. This section introduces and defines the topics of computational ontology, knowledge representation, and reasoning. These are central to how information in the domain of the photograph is represented in the Scone knowledge base and how meaning is captured in ontologies. (Scone is introduced later in section 4.0 *Scone Knowledge Base System*.)

3.1 Computational ontology and its role in KO

The photograph and its image content can be viewed and described from an unlimited number of perspectives. From the viewpoint of LIS, KO provides a general framework for describing and organizing individual photographs and collections of photographs. Hjørland (2008) describes KO in terms of three components: 1) activities such as cataloging and indexing performed by information professionals; 2) a field of study concerned with KO processes and systems; and, 3) a social process engaged in producing and disseminating knowledge. Computational ontology, the central theme in this paper, is an evolving, interdisciplinary specialization that supports a knowledge base approach to these KO activities.

I turn to John F. Sowa and Brian Smith for my working definitions of ontology and computational ontology. Sowa provides an interesting starting point: He states that an ontology is “a catalog of the types of things that are assumed to exist in a domain of interest *D* from the perspective of a person who uses a language *L* for the purpose of talking about *D*” (Sowa 2000, 492). Note that in this definition, an ontology is language-based and domain-specific. An information systems ontology is a computational implementation of such an ontology. According to Barry Smith (2003), the term “ontology” in the computer and information science literature first appeared in 1967, in a work by S. H. Mealy (1967) on the foundations of data modeling. Smith (2003, 22) describes information systems ontology as “a software (or formal language) artefact designed with a specific set of uses and computational environments in mind, and often ordered up by a specific client or customer or application program in a specific context.” Computational ontology brings together under a single rubric several components of KO. These include ontological com-

mitments and a conceptual model that makes explicit relationships among concepts; a system of symbolic representation that provides for organizing and giving structure to these concepts and relations; and a computational system that can support instantiating individuals, question-answer systems, and reuse.

Like cards in a card catalog or MARC records in a bibliographic database, both of which are surrogates for documents stored in a library’s collections, there is a distinction made between the words stored in a knowledge-based ontology and what they signify. This is an important distinction between knowledge-based ontologies and traditional KO systems. The words typed on a catalog card or keyed into a MARC record are not intended to stand in as symbolic representations of extralinguistic objects. Ontological commitments in a knowledge base, on the other hand, adhere to principles of knowledge representation. There is a distinction between the word and the concept it represents, or, as Noy and McGuinness (2000, 13) describe it:

Classes represent concepts in the domain and not the words that denote these concepts. The name of a class may change if we choose a different terminology, but the term itself represents the objective reality in the world.

This leads to the subject of knowledge representation and its role in knowledge-base systems like Scone.

3.2 Knowledge representation and reasoning

In their introduction to *Knowledge Representation and Reasoning*, Brachman and Levesque (2004, xvii) state: “Knowledge representation and reasoning is the area of Artificial Intelligence (AI) concerned with how knowledge can be represented symbolically and manipulated in an automated way by reasoning programs.” Symbolic representation distinguishes computational ontology from KO systems that use natural language representation, for example classification systems and thesauri. Knowledge and its representation in computational environments was discussed at length by Allen Newell nearly a quarter of a century earlier when he presented his presidential address to the American Association for Artificial Intelligence held at Stanford in 1979. Newell (1982) defines the knowledge level as the system level that lies above the symbol level in the hierarchy of levels in computer systems. He defined knowledge as the medium of this level. The central feature of the knowledge level is that

an agent behaves according to the content of its knowledge, taking whatever actions attain its goals given this knowledge (Rosenbloom, Newell, and Laird 1991). In other words, the knowledge level draws its knowledge completely from process, representation, and structure, and process at lower system levels.

Sowa (2000) has also suggested there is something more to machine intelligence than knowledge representation alone. The ultimate justification for choosing one system of ontological categories over another, he says, “must be its applicability to language and reasoning about the world” (p. 68). Thus, knowledge representation and reasoning at the knowledge level work hand-in-hand in knowledge-base ontologies. The Scone knowledge base and *OntoPhoto* should tell users of the system what it believes is true about the world of the photograph, not simply recite what has been explicitly represented. Or, as Lakemeyer and Nebel state it (1994, 4 emphasis original), “one wants to reason about these representations to uncover what is *implied* by them.”

Davis, Shrobe, and Szolovits, specialists in knowledge representation, try to define knowledge representation in terms of its roles. They assert that its

most fundamental role is that of a surrogate, a substitute for the thing itself (Davis, Shrobe, and Szolovits 1993). The Pittsburgh Photographic Library (PPL), housed in the Carnegie Library of Pittsburgh, maintains an extensive picture card catalog system. The catalog is a surrogate for the photograph collections, glass-plate negatives, lantern slides, and other photographic materials that cannot be directly stored in the card catalog. Each individual card in the catalog, like the one shown in Figure 1, stands in for a physical object. This particular card is a surrogate for a photographic print with accession number P-1844, arranged and stored in the archives along with other prints. The symbols that are typed and handwritten on the card are natural language representations describing the photograph, the collection to which it belongs, date captured, and subject heading. The card catalog system—the cabinet and its contents—is thus a model of the PPL archives. The physical holdings of the archives can neither be stored in the card catalog nor a computer, but surrogate records representing these holdings can be easily stored in either one.

The second role holds that a knowledge representation “is a set of ontological commitments” (Davis,




Figure 1. “Carnegie Library of Pittsburgh Microfilm Reader.” Index card with attached gelatin silver contact print. From the Pittsburgh Photographic Library card catalog (P-1844). (Reprinted with permission.) Carnegie Library of Pittsburgh. All Rights Reserved. Unauthorized reproduction or usage prohibited.

Shrobe, and Szolovits 1993, 17). In the ontology of the photograph *OntoPhoto*, the categories of entities and the relationships among them represent ontological commitments, that is, the terms that describe how I think about the world.

The third role states that a knowledge representation “is a fragmentary theory of intelligent reasoning” (Davis, Shrobe, and Szolovits 1993, 17). Ontologies, which Chandrasekaran and others describe as quintessentially content theories, define concepts in terms of such hypothetical constructs as the classes, subclasses and instances that make up the domain of interest (Chandrasekaran, Josephson, and Richard 1998, 2). The theory is fragmentary in the sense that it is an incomplete picture of the domain being described. Further, there is a set of inferences that a representation sanctions and a set that it recommends. Perhaps the important point here is that in order for a knowledge base to support reasoning it must do more than simply list things in a hierarchical taxonomy. The ontology must also state explicitly how these things interact with one another to form propositions.

The fourth role states that a knowledge representation “is a medium for pragmatically efficient computation” (Davis, Shrobe, and Szolovits 1993, 17). The knowledge expressed in *OntoPhoto* is encoded in a way that enables Scone to process knowledge efficiently. Efficient computation is not the same as machine understanding. These terms are sometimes confused in the literature. To say that machines understand conveys the impression that machines possess the ability to think and reason like people. Grigoris Antoniou and Frank van Harmelen (2008, 3) state emphatically that “this gives the wrong impression” and prefer using the term machine-processable. They argue that, at least in the current state of understanding intelligent systems, “it is not necessary for intelligent agents to understand information; it is sufficient for them to process information effectively, which sometimes causes people to think the machine really understands.”

The fifth and final role states that a knowledge representation “is a medium of human expression,” which addresses the communications linkage between ontologists and domain experts (Davis, Shrobe, and Szolovits 1993, 17). In modeling *OntoPhoto*—an iterative, collaborative process—the dialogue expands outward to include librarians, photographers, historians, AI researchers, archivists, and others.

To summarize, computational ontologies are language-based and domain-specific. They support

richer, higher-level semantics than is possible in current controlled vocabularies by making concepts, properties, and relationships among concepts explicit and readable by machines. The five roles of knowledge representation as described by Davis, Shrobe, and Szolovits (1993) help guide information organization in KO systems by facilitating reasoning and by providing the means to symbolically represent knowledge. Symbolic representations distinguish between words and what they signify. Describing an ontology of the photograph in the context of an information system requires aligning oneself with an existing upper ontology. I adopt Scone knowledge-base system for this purpose. In the section that follows, Scott Fahlman, the creator of Scone, is introduced along with details surrounding the basic workings of the Scone knowledge-base system.

4.0 Scone Knowledge-Base System

Scone is an open-source knowledge-base system currently under development in the Language Technologies Institute of Carnegie Mellon University (Fahlman n.d.). Scone supports representation, searching and limited forms of common sense reasoning, all features that are useful for expressing and using knowledge relating to photographs. Knowledge is stored in Scone as a set of files written in a computer language called Common Lisp (Fahlman 1982). A core set of these files makes up Scone’s upper-level ontology. The knowledge represented in *OntoPhoto*, which is also expressed using Common Lisp, is saved to a file and uploaded to the server where the Scone engine and knowledge base files reside. Together, all of these components form a single knowledge-base system.

OntoPhoto functions as a mid-level ontology, representing objects in the domain of the photograph at a more granular level than what is represented in the upper ontology of Scone. The boundaries are not clearly drawn between upper-, mid-, and lower-level ontologies. The upper ontology is limited to generic, very abstract concepts such as “physical object” and “action” and the relationships that link them together. As one moves down through the taxonomic relationships, concepts take on narrower, more specific meanings. Figure 2 presents an example of how a photographic concept might be mapped across these three broad ontological layers.

In one sense, *OntoPhoto* can be likened to a data structure like MARC21. Where MARC21 provides a template for describing facts about bibliographic arti-




Figure 2. Example of the vague conceptual boundaries that exist in the domain of the photograph between what is commonly called upper-, mid-, and lower-level ontologies.

facts stored in a library information system, *OntoPhoto* defines the landscape that makes up interesting distinctions between different aspects of the photograph and defines how these relate to each other in a knowledge-base system. Unlike MARC21 and natural language database systems, however, knowledge stored in *OntoPhoto* is represented in a formal language and organized around a hierarchical framework of concepts and relationships among concepts.

Situated at the heart of the Scone Project is Scott Fahlman, its creator, as well as graduate students who are members of the Scone Research Group at Carnegie Mellon University. Proponents of Scone argue that what makes Scone different from other knowledge-base ontologies is how it performs search and inference. Fahlman explains: “Scone uses marker-passing algorithms originally designed for a hypothetical massively parallel machine” (Fahlman 2006a). The massively parallel machine Fahlman is referring to is NETL, a hardware architecture that he began developing during the early 1970s as a doctoral student at the MIT Artificial Intelligence Laboratory (Fahlman

1979). However, to enter knowledge into Scone, one does not have to be familiar with NETL or be concerned with marker-passing algorithms and parallel network theory. The Scone project is devoted to providing support for representing and using knowledge and *OntoPhoto* offers a test case for how structured knowledge in library collections—information relating to photographs in particular—can be represented with precision in a knowledge-base ontology.

4.1. Motivating example

This section presents a brief example of what I call a classical flat model—the system of representation used in current library catalogs for describing photographs. The term “flat” is used to describe a typical bibliographic record where the content is specified in advance of the query, and can only be accessed in an inflexible and rather brittle way. The goal of this example is to search a catalog, matching search terms in the query with indexed terms describing a photograph. Consider the photograph shown in Figure 3. Searching the Library of Congress online catalog using the query “destitute pea pickers” retrieves a description of this photograph. The content of the record is flat in the sense that its content is determined and fixed before the query is submitted. The data are structured using MARC21 formatting standards and *Anglo-American Cataloguing Rules, 2nd edition* (AACR2) content standards. The record states, among other things, the following:

Type of Material: Photograph, Print, Drawing.
 Personal Name: Dorothea Lange, Photographer.
 Main Title: Destitute pea pickers in California.
 Mother of seven children. Age 32.
 Created/Published: 1936 Feb.

To assist in evaluating what happens in this sample search, I turn to Levesque and Lakemeyer (2000, 6) who offer a useful framework for explaining why researchers in AI “want their systems to know a lot ... and want their systems to represent that knowledge symbolically.” The Library of Congress record states that Dorothea Lange is a photographer and that she is the creator of a photograph named “Destitute pea pickers ...” without claiming that there is anything represented in the Library’s catalog system that corresponds to these propositions. The Library has programmed its system to behave in a way that processes queries by matching search terms with index terms and when a match occurs, the system returns a preex-




Figure 3. Dorothea Lange, “Destitute pea pickers in California; a 32 year old mother of seven children, February 1936.” Library of Congress, Prints and Photographs Division, FSA/OWI Collection, [LC-USF34-T01-009058 (black and white film duplicate negative)]

isting record describing Dorothea Lange’s photograph. Stated simply, the success of the search depends on the presence or absence of words.

Fahlman (2009) compares knowledge-assisted searching to “bag-of-words” searching, which he asserts “is fundamentally based on the presence or absence of specific words (or multi-word phrases) in a given document.” Bag-of-words in the strict sense of the term is not a label that fits a MARC21 bibliographic record. Each MARC21 field holds a piece of bibliographic information such as author, physical description, or topical subject heading. The record is not a document in the strictest sense, but is certainly a collection of categorized, unordered sets of words.

4.2 Scone’s features

This same photograph represented in Scone and described within the framework of an ontology allows for specifications that are not available in the classical flat model described earlier. As with a traditional library catalog, the searcher may begin with one or more concepts such as documentary photography,

photojournalism, Great Depression, Migrant Mother, or Dorothea Lange. An ontology also makes it possible for the searcher to reason with a variety of objects such as people in various roles, camera equipment, events, institutions, and so on. A query might also be built around different relation types that exist among objects, for example, *photographers* who are affiliated with the *Farm Security Administration* under the direction of *Roy Emerson Stryker*. These features are made possible through a variety of services offered by Scone including multiple inheritance, default reasoning with exceptions, and multiple contexts.

4.2.1 Multiple inheritance and default reasoning with exceptions

Scone supports multiple inheritance through an is-a hierarchy. If “knife” is represented as a kind of eating utensil, “knife” has its own unique properties and inherits more general properties from its superclass “eating utensil.” Multiple inheritance means that “knife” can also be a subclass of “weapon” and “tool.” Inheritance also applies to properties, roles and relations.

A common example of default reasoning with exceptions given in AI literature relates to birds that do not fly. While the default bird is a “flying thing,” Scone can make exceptions for penguins and ostriches, which are birds that cannot fly. In Scone, one can add properties to individuals or categories and these can be default categories with exceptions. Florence Own Thompson, the iconic figure in Dorothea Lange’s FSA photograph pictured in Figure 3, can inherit all the default properties belonging to all women, and also be given the unique property of “pea picker” and “destitute,” both of which are terms included in Lange’s original caption assigned to this photograph.

Scone supports systems in which searchers design queries in terms of a particular set of properties and class memberships. In the near future, a query could combine searching with browsing and provide the searcher with multiple paths for pursuing new knowledge (Fahlman 2009). For example, the searcher might enter a query, “Exhibits of Dorothea Lange Japanese internment photographs.”

4.2.2 Making recommended inferences

A powerful feature of Scone is its search and inference engine. For example, we can formalize in Scone the proposition that “Dorothea Lange is a photographer”

and “Photographers use cameras,” and, after some manipulation, Scone reasons that “Dorothea Lange uses a camera.” This kind of reasoning is called logical inference because the proposition represented by the last proposition logically entails the propositions represented by the initial ones. Fahlman offers a good example of this feature when he describes a scenario in which a photograph is given the caption “Dog catching a Frisbee” and the photograph is retrieved using the query “Pets playing with toys” (Fahlman 2008a).

This demonstrates how the category “dog” is expanded to the category “pet” (“dog” is-a kind of “pet”). The inference capabilities these examples demonstrate can be accomplished by other knowledge-base systems and are not unique to Scone. What sets Scone apart is its scalability, expressive power, and the speed with which it can perform these searches—design decisions Fahlman made realizing that he would have to give up something in exchange for these goals. Brachman and Levesque (2004) explain the consequences of choosing one property over another and the tensions that exist between the degrees of expressiveness offered by a representational language and the ability to reason with that language. They suggest that, as a representational language’s expressive power increases, one’s ability to handle that system and effectively reason with it decreases. Fahlman (2008b) is well aware of these tradeoffs, consciously making the choice “to give up on logically complete proof procedures,” in favor of scalability and expressiveness. Fahlman concludes that, while most knowledge representation systems reason by a process of logical theorem proving, Scone’s less formal reasoning method is good enough for the kind of everyday reasoning needed in its applications. Put another way, Scone facilitates making inferences to a level Fahlman believes is sufficient for human-like reasoning and then stops.

4.2.1 Multiple contexts

One final characteristic worth noting for its relevance to the photograph is a feature Fahlman calls multiple contexts. A context in Scone is a labeled node that represents some state (real or hypothetical) of the world, or some viewpoint. Every referent in a knowledge base is connected to a context node, and every context is interlinked through a hierarchy of is-a links just like other nodes in the knowledge base (Fahlman 2006a, 123). Fahlman describes driving from home to the airport, an action that creates two contexts, “one representing the world before the event and the other

representing the world after the event” (Fahlman 2006a, 124). Applied to photographs, multiple contexts can be used effectively for representing a collection of photographs that are first located in an individual’s safety deposit box and later moved to an archive, where the collection is accessioned, processed, and cataloged. Both settings share some features in common, for example, restrictions on access and image content. There are specific aspects that change after the move, such as the geographic location and relation to other photograph collections that are part of an archive.

Another interesting use of context applies to a practice known as repeat photography—that is, the replication of photographs to show change over time. The images shown in Figure 4 illustrate a plot of land being monitored over time in the Mojave Desert (Webb et al. 2001). These two images describe two different contexts, but share certain data in common.




Figure 4. A plot of land in the Mojave Desert provides an example of replication photography. Photograph A was taken in 1964 and photograph B was taken 36 years later. Courtesy of the U. S. Geological Survey.

Scone representation, which supports multiple contexts, enables reuse of the shared data including location of the plot and certain geographical features, and it supports the ability to make distinctions about new data. New data includes the date photograph B was taken, relevant information associated with changes in shrubs, rainfall, soil moisture, and so on.

The added knowledge resulting from both a principled approach to ontology development and the background information that matches the meaning of words in a query to the meaning of a photograph and entities related to that photograph expands Scone's knowledge in multiple directions. Scone's knowledge base makes it possible for a more precise, but casual and non-linear, means of exploring information.

5.0 Knowledge bases, concepts, roles, and relation types

Photographs are objects in the world. As far as we know, any photograph can be modeled or represented in a library catalog system. Whether a library catalog embodies or materializes a photograph is a different matter. The word "photograph" is not the object. It is a representation of the object, and, as Timothy Binkley (1997) portrays this notion, using words to represent image content can be problematic. When Binkley claimed (p. 107), "pictures are superbly demonstrative and eminently computable, transcending the parochial limitations of natural languages," he was asserting that visual information supplies its own content. The purpose of this section is to begin working through these problems of representing the meaning of photographs in the context of constructing an ontology of the photograph. Binkley's claim may be true when viewing the photograph itself, but in library catalogs and knowledge-base systems, knowledge about the photograph is most often re-presented in linguistic forms. This section introduces the nomenclature and conceptual framework used for making this transition from being an object in the world to becoming a concept and then a representation in a knowledge base.

5.1 Knowledge-base systems

The Artificial Intelligence community offers various definitions of knowledge base in the literature. A knowledge base, according to John F. Sowa (2000, 495), is "an informal term for a collection of information that includes an ontology as one component." Brachman and Levesque (2004, 7) tell us that a knowl-

edge base is "a collection of symbolic structures representing what it believes and reasons with during the operation of the system." Fahlman (2003-2008, 8, emphasis original) describes a Scone knowledge base as being "made up of *elements* that are connected together to form a *semantic network*." To bring these ideas into focus, it is helpful to draw comparisons with KO systems and point out what, if any, similarities exist. Ontology of the photograph is the focus of this article, so that is the formal knowledge-base structure we will use for comparison.

An ontology consists of formal statements that represent meanings of concepts, properties, and certain relationships in a domain of interest. To begin with, formal statements in a knowledge base are quite different from statements made in a typical library catalog record; they are more formal than the language used for entering bibliographic information into a MARC21-formatted record. While it is true that librarians follow certain data content standards when building MARC21 records, for example, AACR2, the language, itself, is natural language. In MARC21, there is a field for author, title information, and so on. Three-digit numbers called "tags" represents these fields. For example, the 100 tag marks or identifies the kind of data that follows the 100 tag as being personal name main entry (author). While the OPAC label for this field—what a library patron sees on their computer monitor—varies from library system to library system, the underlying records all follow the same data structure standards. Authors, titles, and publishers can be compared to classes of objects in an ontology. There is an implicit relationship in library catalogs between author and book: Authors write books. This same relationship in a knowledge base, however, is formally defined, made explicit, and can be reasoned over.

The idea that knowledge statements in a knowledge base can have "meaning" is the most important characteristic that distinguishes ontology-based KO systems from, for example, library catalogs built around MARC21-formatted records, archives using EAD (Encoded Archival Description) finding aids, or search and indexing systems like Google. Fahlman, in his paper "Natural Language: It's All About Meaning," aggressively argues this point when he states (2006b, 2), "most linguists, computational linguists, and AI researchers have accepted for a long time that meaning—both the knowledge conveyed by an utterance and a large amount of background knowledge—is central to the enterprise of producing and understanding natural language." Fahlman acknowledges the astounding, but partial, success of these systems

that rely on “statistical means—without really understanding anything” or what he has called the “meaning-free bag-of-words model” (2006b, 2). Where these systems fall short is in the percentage of documents returned that are not relevant to the searcher’s query. In the final analysis, KO systems like the Library of Congress online catalog depend on the searchers, themselves, to read the documents in search result sets to determine, as Fahlman puts it (2006b, 2), “whether a document really contains the answer to their question, and what that answer is.”

5.2 Concepts

Concepts are described in Cruse’s *Meaning in Language* (2000, 127) as “organized bundles of stored knowledge representing an articulation of events, entities, situations, and so on in our experience.” The IDEF5 Method, a highly formalized method for capturing a set of ontological commitments, refers to concepts as kinds and defines kinds as “an objective category of objects that are bound together by a common nature, a set of properties shared by all and only the members of the kind” (Benjamin et al. 1994, 13). In an ontology of the photograph, concepts represent classes of objects such as “the typical *Cart de visite*” or “the typical inkjet print.” In Scone, concepts are enclosed in curly brackets. For example, {photograph} refers to the class of objects that represents the typical photograph; it “represents a specific *concept* or *meaning* ... not a word or word-definition” (Fahlman 2003-2008, 8).

Once again, understanding the word-concept distinction is important for understanding how library catalogs differ from knowledge-base systems. The meaning of the concept {photograph} is not derived from the word “photograph.” Rather, it derives its meaning from its position in a hierarchical taxonomy, from relationships it forms with other concepts in the ontology and from the properties it inherits from its superclass (parent class). A relationship can be drawn between natural language and concepts in Scone, but, as Fahlman explains it (2003-2008, 8-9), making this connection “is the job of external software (an English-language front-end to Scone) to resolve any ambiguity.”

5.3 Relations

Relations are the connections or associations among concepts. For example, in *OntoPhoto* a photographic print is a type-of photograph. The type-of relation es-

tablishes a class hierarchy or taxonomy of hierarchical relationships. For instance, “Cyanotype print” is a type-of “photographic print,” and “photographic print” is a type-of “photograph.” Terms commonly associated with taxonomical hierarchies are class, subclass, superclass, and subsumption. Noy and McGuinness describe classhood like this (2000, 12): “A subclass of a class represents a concept that is a ‘kind of’ the concept that the superclass represents.” In the earlier example, “photograph” can be described as a superclass of “photographic print” and the class “photograph” subsumes “photographic print,” which is a subclass of “photograph.” All of the relations described thus far are binary, that is, the relation holds between two entities. A is a kind-of B. Theoretically relations could have any number of arguments, but binary and ternary relations are the most common in ontologies. The knowledge base sample presented below includes both binary and ternary relations. “Dorothea Lange is-a person” constitutes a binary relation. “Dorothea Lange is-the photographer of LC-USF34-T01-009058” is a ternary relation.

The example of a ternary relation, “x-is-the-y-of-z,” presents an interesting study of relations in how it treats the preposition of as a kind of role relation. Roles are described in more detail later in this section. The theme that has been pervasive throughout this article is that knowledge bases and ontologies are formal, explicit, and precise accounts of things in the world—objects, their properties, and relations. What makes the ontologist’s work particularly difficult is expressing with exactness terms such as “of.” Of is a preposition and has multiple meanings depending on the context. In this example, the photographer Dorothea Lange is an animate being acting intentionally, so the role relation between photographer and the specific photograph the photographer captures is an AGENT relationship. In the statement, “Migrant Mother’ is part of the FSA/OWI Collection,” the preposition of takes on a different meaning. Here of can be interpreted as a PART-OF relation within the context of library collections.

To say something interesting about the specific photograph shown in Figure 3 requires making more complex knowledge statements made up of a variety of concepts and relations. For example, we could begin by creating a subtype of photograph called a “photographic print” and state that there is an individual node (an instance of a specific photographic print) “LC-USF34-T01-009058.” We could add that the photographer who took this photograph is named “Dorothea Lange” and the name of the photograph is

“Migrant Mother.” This small snippet of knowledge can be represented as shown in Figure 5 using Scone’s knowledge representation system.

```
(new-type {photographic print} {photograph})
(new-indv {LC-USF34-T01-009058} {photographic print})
(new-indv {Dorothea Lange} {person})
(new-indv-role {image name} {photograph} {string})
(new-indv-role {photographer} {photograph} {person})
(x-is-the-y-of-z {"Migrant Mother"} {image name} {LC-
USF34-T01-009058})
(x-is-the-y-of-z {Dorothea Lange} {photographer} {LC-
USF34-T01-009058})
```

Figure 5. Knowledge expressed in Scone representation.

Photographs—the physical manifestations— can exist in various relations with other objects. A photograph can be mounted in a frame, hanging on a wall, three feet below the ceiling, next to the exit in an art gallery. The preceding sentence describes spatial relation types, all of which can be defined and represented in the knowledge base. Picture elements can also exist in various relations with other picture elements. In Dorothea Lange’s photograph shown in Figure 3, the main subject, Florence Owen Thompson, holds relations with other objects in the image including each one of her children. Locative expressions like “child on the left” can assist historians in identifying individuals pictured in photographs. Furthermore, the relationship between the unseen camera and objects in the scene can be represented in terms of point of view—a feature of aesthetics important to professional photography teachers teaching composition.

5.4 Roles

Concepts by themselves tell us very little about things. Roles, sometimes called attributes or slots, provide us with additional information by describing the characteristics of things. In the domain of the photograph, these include, for example, a photograph’s material base, chemistry, color, dimensions, and capture date. Terms that designate roles can be considered functions and are often combined with “of.” The *IDEF5 Method Report* describes function within this context as “a mapping that takes each member of a given set of individuals to a single specific value” (Benjamin et al. 1994, 13). Gregory McCulloch (1989, 8) describes a function “as a sort of abstract machine or processor which needs to be fed

with certain things, like numbers, and which duly extrudes others.” For example, in the Farm Security Administration photograph shown in Figure 3, the role *emulsion-chemistry-of* maps the photograph to the light-sensitive chemistry applied to its surface. To say “The photographic print known as Migrant Mother has an *emulsion-chemistry-of*” requires an answer. The light-sensitive chemistry used on the surface of this print, for example, has *emulsion-chemistry-of* gelatin and silver halide crystals. A measurement of 8 x 10 inches is a standard photographic paper size. This feature of a photograph comprises two roles: *height-length-of* and *width-length-of*, both of which yield values. The point here is that ontology-based representation enables one to formally express roles in KO systems. This can enhance traditional KO systems because roles derive meaning from their placement within a semantic data structure, whereas, in MARC records, roles are treated as meaning-free index terms.

6.0 Ontology of the photograph: *OntoPhoto*

OntoPhoto—a portmanteau word blending ontology and photograph—stands for the research project I am undertaking to examine important ontological analysis and conceptual modeling questions. *OntoPhoto* offers a test case of how a photograph’s structured knowledge can be represented in an ontology-based approach to KO. There are no specialized corpora of data pertaining to the photograph from which to draw information for modeling an ontology. However, the concept of a “photograph” is included in many already existing information system taxonomies. Examples include Princeton WordNet®, Library of Congress Thesaurus for Graphic Materials, and the Getty Art and Architecture Thesaurus. The data used here are extracted from literature on the photograph pertaining to history, general theory, and practice, as it offers the richest source of photography-related information. There are broad methodological challenges that must be overcome to create a corpus of photograph-related data for use in creating conceptual model. These challenges are beyond the scope of this article. The discussion here focuses on a limited set of concepts for illustrative purposes only. The main goal of this final section is to illustrate applications of the principles introduced in earlier sections. This is done through diagrammatic views of some of the ontological commitments made in *OntoPhoto* accompanied by explanations of the analysis behind their conceptualizations. Only a small sampling of concepts are discussed and presented graphically. Throughout this section, entity

types are enclosed in curly brackets { }, and relation types are represented in italics.

6.1 *The medium-image-context distinction*

A photograph's essence in *OntoPhoto* originates in the notion that it is a medium for carrying or displaying; it is an image, that is, a picture of something; and its meaning derives from the contexts in which it is made and viewed. This section is concerned with defining the landscape that makes up interesting distinctions between these concepts and defining how they relate to each other in a knowledge base. The labels used in *OntoPhoto* for these three concepts are {photograph}, representing the physical medium; {photographic image} representing the content of the photograph; and {photograph context}, representing the information that exists outside the frame of the photograph. The differences between these concepts and their relationship to each other are nontrivial. An ontology makes explicit what connects one to the other, but what is it precisely that is being connected and how are they related? There are potentially an infinite number of objects, relationships among objects, subject themes, and so on, that could be represented. The {photographic image} of Florence Owen Thompson represented in Figure 3 (here I am referring to the physical artifact stored in the archives, not the image presented in this journal) is a physical thing lying upon the surface of a piece of photographic paper housed in the Library of Congress. That is, the {photographic image} is the sum total of all the silver halide crystals suspended in gelatin on the surface of the photographic paper. While the image is made of physical stuff, its information content is abstract. The following sections explore more closely the concepts of medium and image content.

6.1.1 *The medium*

The concept {photograph} links to Scone's upper ontology at the node {image object}. {Photograph} represents the medium and is a type-of {image object}. The concept {image object} represents things created for communicating meaning, primarily visually and nonverbally. The concept {image object} refers to the medium, the physical manifestation, base, or support, not the information content. It is the physical stuff in which the intellectual content is stored. The medium may be made of paper, glass, wood, metal, cloth, cellulose acetate (film base), or other materials suitable for displaying a (photographic im-

age}. In film photography, the medium is the material on which light-sensitive emulsion is applied. The surface is exposed to light, and the subsequent development of an image takes place.

6.1.1.1 *The split- and complete-split-subtypes*

Two critical design concepts adopted by *OntoPhoto* and made possible by Scone's representation system are the notions of new-split-subtypes and new-complete-split-subtypes. The new-split-subtype creates a set of any number of disjoint subtypes under a parent. For example, the parent {photographic paper} may have split-subtypes {art paper}, {resin coated paper}, {fiber paper}, and so on. The new-complete-split-subtypes makes the additional claim that the split subtypes under the parent taken all together represent a complete set of possible subclasses. In other words, every member of the parent class must belong to one of the subclasses listed in the complete split. For example, the parent class {person} may have the new-complete-split-subtypes {female} and {male}. Any member of the class {person} must belong to either {female} or {male} (Fahlman 2003-2008, 30-32).

OntoPhoto applies the split-subtype function to the classification hierarchy at the level of {photograph}. At the node {photograph} branching occurs in several dimensions. The concept {photograph}, for example, cross-classifies to {reflective} and {transparent}. These concepts and the split-subtype relations are illustrated in Figure 6.




Figure 6. Snippet view of *OntoPhoto* showing the concept {photograph} and its complete-split-subtypes.

These dimensions—the two pairs of split-subtypes—are essentially parallel in the taxonomy and various combinations of values along these dimensions specify subclasses. For example, Dorothea Lange's photograph archived at the Library of Congress and represented in Figure 3 is an {analog photograph} and a {reflective photograph}. Another instance of Lange's

“Migrant Mother,” for example, the original black and white negative, would be {analog photograph} and {transparent photograph}.

Some examples of types of {photograph} include {photographic print}, {lantern slide}, and {photographic film}. Some subtypes of {photographic print} include {calotype}, {platinum}, and {autochrome}. Like {photographic print}, the entity type {photographic film} has several subtypes such as {Kodacolor}, {Ektachrome}, {cellulose nitrate}, and {cellulose diacetate sheet film}.

There are two relations that define the entity {photographic print}: 1) A {photographic print} is-a {analog photograph}, and 2) A {photographic print} is-a {reflective photograph}. The inheritance principle states that anything true about the entities {analog photograph} and {reflective photograph} must also be true of {photographic print}. Likewise, typical members of the class {photographic film} share

two attributes in common: 1) A {photographic film} is-a {analog photograph}, and 2) A {photographic film} is-a {transparent photograph}.

Based on what has been discussed thus far, in addition to some roles and instances, a small ontology of the photograph can be mapped out as shown in Figure 7. Snippet views are abbreviated views of an ontology’s concepts and relations. Snippet views of *OntoPhoto* are presented here because it is impossible to show on a single journal page the full-scale anatomy of the ontology.

The distinctions *OntoPhoto* makes between physical manifestation {photograph} and content of {photographic image} extends into the realm of {digital photograph}. The {digital photograph} is highly problematic, however, because the nature of a digital image—how it is created, presented, stored, viewed, and so on—is quite different from an analog photograph. Still, there are some salient properties that are




Figure 7. Snippet view of *OntoPhoto* showing some of the concepts and relation types found in the domain of physical manifestation.

shared in common, the most important being that both are representational technologies. Further, classifying a digital photograph as a kind of photograph is in accordance with current views in archival studies and information science. There is literary warrant justifying doing this. Other properties inherent in a digital photograph, alluded to earlier, are not those of its apparent kind. Until there is a useful theoretical structure available to assist in classifying digital photographs as something other than photographs, they remain as concepts in the current model.

6.1.2 *The image*

Extending ontology-based representation to the visual modality—image content or the semantic content of images—and mapping this to the physical manifestation of the photograph requires a radically different approach to representation than is currently applied in descriptive cataloging methods. One must make explicit in the knowledge base the relationship that exists between the physical manifestation of the photograph and the image content. The manner in which this relationship is represented determines, among other things, the knowledge that can be extracted from the knowledge base as it gets populated. This section first provides examples of interesting concepts and relation types pertaining to the image and then explains how *OntoPhoto* makes the conceptual link between image and medium.

6.1.2.1 *Describing the image*

In these early stages of development, *OntoPhoto* regards the subject of the photograph as the content of the image and groups all of the attributes an image may have under two broad categories: physical and abstract. This may be viewed as corresponding to the classic dichotomy in subject analysis that asks the questions: what is the picture about and what is it a picture of. These two entity types are expressed in Figure 8 as {subject matter physical} and {subject matter abstract}. *OntoPhoto* does not yet make explicit the subject facets of who (person), when (time), where (place), what (activities, events and objects), and so on. Facets would describe homogeneous classes of concepts, the members of which share characteristics that distinguish them from members of other classes. Distinctions could also be made between generic and specific. For example, a given photograph could be described generically as a “holiday event” or specifically as a “Fourth of July parade.”

Figure 8 illustrates a simple ontology describing {photographic image} and its hierarchical connection to {physical object}. Once again, the concept of split-subtypes is applied. I apply this function to the classification hierarchy at the level of {photographic image}. At the node {photographic image}, branching occurs in several dimensions. First, the concept {photographic image}, like its physical carrier {photograph}, cross-classifies as either {analog} or {digital}. Second, the concept {photographic print} cross-classifies along another dimension as either {negative} or {positive}. Finally, {photographic print} cross-classifies along a third dimension {chromatic} or {monochromatic}. These various entity types and relations are presented diagrammatically in Figure 8.

6.1.2.2 *Linking the medium to the image*

The last point covered in this article concerns how to represent, in a knowledge base, the relationship between the medium and its image content. Current library information systems implicitly encode linguistic relations about the physical manifestation of photographs and their images. For example, the specialized *Thesaurus for Graphic Materials I: Subject Terms (TGM I)* makes distinctions between the photograph’s subject content and artifactual value (Parker and Zinkham 1995). Visual researchers, the Library of Congress asserts, “know with great certainty whether they wish to see examples of formats and physical types or images in which formats or physical types constitute the subject” (Parker and Zinkham 1995). If a particular photograph is an example of a physical type such as cartes de visite—say a photograph showing Chief Archivist Scott Yoss describing the cartes de visite collection at Photo Antiquities in Pittsburgh—a cataloger could index the photograph with the term cartes de visite by entering the term in MARC field 655 (for the physical characteristic of the item).

In Scone representation, the relationship between the image and its carrier is made explicit with the following form:

```
(new-indv-role {photographic image} {photograph} parent: {graphic image})
```

This form can be read as follows: A {photographic image} is a type-of {graphic image} and every {photograph} has-a {photographic image} (in the absence of an explicit cancellation). A diagrammatic view of this relationship is shown in Figure 9.




Figure 8. Snippet view of *OntoPhoto* showing a sample of the concepts




Figure 9. This figure shows the mapping of physical manifestation to image content.

7.0 Conclusion

I have here offered a detailed view of an alternative approach to representation of photographs first introduced in 2009 as the *Semantic Archives* model (Benson 2009). The question posed at that time was whether archivists engaged in describing photographs needed a more formalized system of representation than those offered by traditional description standards. The argument was put forth that higher-level semantics would support more meaningful represen-

tations, effectively reaching the level of semantic Web utility. Also, I have introduced *OntoPhoto*, an ontology of the photograph, and described knowledge-base systems, their component parts, and how they can be utilized to organize information about the photograph's physical manifestation and image content. Together these tools offer high-level semantics in a formalized system of notation that humans can read and machines can process effectively. Snippet views of *OntoPhoto*'s backbone taxonomy helped illustrate how an ontology-based system of representation provides an effective and efficient environment for organizing information, one that supports the features of intelligent reasoning and can potentially improve the breadth and accuracy of search results.

While ontologies strive to present viable alternatives to traditional data structure standards, they do not intend to replace them. Rather, it is another critical role of ontologies such as *OntoPhoto* domain ontology to collect together, utilize, and make interoperable the terms and semantics found in existing thesauri such as *Art and Architecture Thesaurus (ATT)*, *Thesaurus for Graphic Materials (TGM)*, and wordnets such as Princeton University's WordNet®. This capability begins with a rich ontology of photograph concepts and relation types. The meanings of terms and relationships in a given thesaurus are then mapped to various points in the ontology. There will be instances where one thesaurus's definition of a "photograph" includes a "film negative" and another thesaurus does not.

Once the structures of these knowledge organization systems have been analyzed and clarified, information can be correctly aligned and represented in a uniform manner at the level of concepts and relation types. A common conceptual vocabulary begins to emerge with ontology forming the heart of the KO system. This offers several advantages: First, ontology-based data integration provides more specific and hierarchically structured KO by considering a rich set of relations between concepts. This enables visual researchers to expand their searches and retrieve information from cross-related and intersecting data. Second, an ontology-based conceptual vocabulary system offers a rich cataloging environment through knowledge sharing and reuse. A shared conceptual vocabulary is likely to include a rich set of domain-specific terms such as printing-out paper and emulsion, terms that describe attributes such as dimension and capture date, and general terms that describe spatial relations such as in, on, and between. Sharing the representation of this knowledge among catalogers

increases the potential for knowledge reuse and eliminates the need for replicating the ontological analysis process. Finally, an ontology-based representation technology is a computational environment in which reasoning takes place. Multiple inheritance and inference expands the universe of knowledge beyond what is stated explicitly in the system.

References

- Antoniou, Grigoris and van Harmelen, Frank. 2008. *A semantic web primer*. Cambridge, Massachusetts: MIT Press.
- Benjamin, Perakath C., et al. 1994. *IDEF5 method report*. College Station, Texas: Knowledge Based Systems, Inc.
- Benson, Allen C. 2009. The archival photograph and its meaning: formalisms for modeling images. *Journal of archival organization* 7: 148-87.
- Binkley, Timothy. 1997. The vitality of digital creation. *The journal of aesthetics and art criticism* 55: 107-16.
- Brachman, Ronald J. and Levesque, Hector J. 2004. *Knowledge representation and reasoning* New York: Elsevier.
- Chandrasekaran, B., Josephson, John R. and Benjamins, V. Richard. 1998. Ontology of tasks and methods. In *ECAI98's workshop on Application of Ontologies and Problem Solving Methods*. Brighton: Wiley.
- Cruse, D. Alan. 2000. *Meaning in language: An introduction to semantics and pragmatics* New York: Oxford University Press.
- Cruse, D. Alan. 1999. What are ontologies and why do we need them? *IEEE intelligent systems* January/February 1999: 20.
- Davis, Randall; Shrobe, Howard; and Szolovits, Peter. 1993. What is a knowledge representation? *AI magazine* 14: 17-33.
- Fahlman, Scott E. 1979. *NETL: A system for representing and using real-world knowledge*. Cambridge, Massachusetts: MIT Press.
- Fahlman, Scott E. (n.d.). *The Scone Knowledge-Base Project*. Available <http://www.cs.cmu.edu/~sef/scone/>.
- Fahlman, Scott E. 1982. Common Lisp. *Annual review of computer science* 2: 1-19.
- Fahlman, Scott E. 2003-2008. *Scone user's guide*. Pittsburgh, PA: Carnegie Mellon University, School of Computer Science. Available <http://www.cs.cmu.edu/~sef/scone/Scone-User.htm>.

- Fahlman, Scott E. 2006a. Marker-passing inference in the Scone knowledge-base system. In Lang, J., Lin, F., and Wang, J. eds., *KSEM 2006, LNAI 4092*. Berlin: Springer, pp. 114-26.
- Fahlman, Scott E. 2006b. Natural language: it's all about meaning. In *Position Paper for the 2006 AAAI Fellows Symposium*. Available <http://www.cs.cmu.edu/~sef/scone/publications/AAAI Fellows2006.pdf>.
- Fahlman, Scott E. 2008a. Personal communication (November 27, 2008).
- Fahlman, Scott E. 2008b. In defense of incomplete inference. In *Knowledge nuggets*. Available <http://sef-linux.radar.cs.cmu.edu/nuggets/?m=200807>.
- Fahlman, Scott E. 2009. Personal communication (November 12, 2009).
- Hjørland, Birger. 2008. What is knowledge organization (KO)? *Knowledge organization* 35: 86-101.
- Hollink, L. 2006. *Semantic annotation for retrieval of visual resources* (Doctoral dissertation, Dutch Research School for Information and Knowledge Systems, 2006). Available <http://www.cs.vu.nl/~laurah/1/papers/ProefschriftLauraHollink.pdf>.
- Lakemeyer, Gerhard and Nebel, Bernhard. 1994. *Foundations of knowledge representation and reasoning*. New York: Springer Verlag.
- Levesque, Hector J. and Lakemeyer, Gerhard. 2000. *The logic of knowledge bases*. Cambridge, Massachusetts: MIT Press.
- McCulloch, Gregory. 1989. *The game of the name*. New York: Oxford University Press.
- Mealy, G. H. 1967. Another look at data. In *Proceedings of the Fall Joint Computer Conference, November 14-16, Anaheim, California*. AFIPS conference proceedings, vol. 31. Washington, DC: Thompson Books, London: Academic Press, pp. 525-34.
- Newell, A. 1982. The knowledge level. *Artificial Intelligence* 18n1 : 87-127.
- Nirenburg, Sergei and Raskin, Victor. 2004. *Ontological semantics*. Cambridge, Massachusetts: MIT Press.
- Noy, Natalya F. and McGuinness, Deborah L. 2000. Ontology development 101: A guide to creating your first ontology. Stanford KSL Technical Report KSL-01-05. Available <http://ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness.pdf>.
- Parker, Elizabeth Betz and Zinkham, Helena. 1995. *Thesaurus for graphic materials I*. Washington, D.C.: Library of Congress. Available <http://www.loc.gov/rr/print/tgm1/>.
- Rosenbloom, P. S., Newell, A. and Laird, J. E. 1991. Toward the knowledge level in Soar: The role of the architecture in the use of knowledge. In VanLehn, K., ed., *Architecture for intelligence: The 22nd Carnegie Mellon Symposium on cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 75-112.
- Smith, Barry. 2003. Ontology and Information Systems. Preprint version of longer draft of chapter "Ontology." In *Blackwell guide to the philosophy of computing and information*, ed. L. Floridi. Oxford: Blackwell. Available http://ontology.buffalo.edu/ontology_long.pdf.
- Sowa, John F. 2000. *Knowledge representation: Logical, philosophical, and computational foundations*. Pacific Grove: Brooks/Cole.
- Webb, R. H. et al. 2001. Monitoring of ecosystem dynamics in the Mojave Desert: the Beatley permanent plots. USGS Fact Sheet FS-040-01, U.S. Geological Survey. Available <http://pubs.usgs.gov/fs/FS-040-01/pdf/fs-040-01.pdf>.