

Reasoning from Imperfect Knowledge

Stephen K. Reed

Psychology and CRMSE, San Diego State University
Department of Psychology, University of California, San Diego

Adam Pease

Articulate Software

Address Correspondence to:

Stephen K. Reed

Department of Psychology
University of California, San Diego
La Jolla, CA 92093

858 456-6686

sreed@mail.sdsu.edu

Abstract

One reason why there is a lack of cross-references between articles on knowledge representation in the Cognitive and the Information Sciences is that cognitive scientists are interested in descriptive models of how people reason whereas information scientists are interested in prescriptive models to help people reason. Formal ontologies such as the Suggested Upper Merged Ontology aid human reasoning by providing (1) an accurate knowledge base, (2) a formalization of the knowledge base as axioms, and (3) a logic to derive new information through deductive reasoning. However, all systems confront obstacles when reasoning from imperfect knowledge consisting of ambiguous, conditional, contradictory, fragmented, inert, misclassified, or uncertain knowledge. We use work from the Cognitive and the Information Sciences to analyze obstacles for both computers and people when confronted with ambiguous, contradictory, misclassified, and uncertain knowledge. A concluding section considers both practical and theoretical applications of ontologies in the Cognitive Sciences.

Reasoning from Imperfect Knowledge

1. Introduction

Building descriptive models of deductive reasoning has been a major topic of investigation in Cognitive Science, including the construction of natural deduction rules to manipulate propositions in working memory (adapted from Rips, 1983). Other theoretical approaches of deduction advocate mental models that propose people reason by manipulating tokens (Johnson-Laird, 2012). A goal for all descriptive models of reasoning is to account for errors. For instance, participants in a study described by Rips (1983) had to decide which propositional arguments were valid. Acceptance of validity ranged from 17% to 92% across 32 valid arguments.

Reasoning is also a major topic in the Information Sciences, including the construction of formal ontologies that use logic-based programming languages to draw deductive conclusions (Pease, 2011). Perhaps the major cause of the lack of integration of work on deductive reasoning in the Cognitive and the Information Sciences is that Cognitive Science is primarily interested in descriptive models of how people reason whereas the Information Sciences are primarily interested in prescriptive (normative) models to enable machine reasoning.

Formal ontologies help people by providing both an accurate knowledge base and deductive rules to derive conclusions. Knowledge should be based on the latest discoveries in science (Smith & Ceusters, 2010) and deductions based on first- or higher order logics that derive valid inferences (Pease, 2011). When these objectives are achieved, the potential is enormous for supporting unambiguous and machine-readable documentation, consistency verification, data classification, querying, and further ontology development (Hoehndorf, Dumontier, & Gkoutos, 2012).

Although formal ontologies are better in deductive reasoning and people are better in possessing an extensive knowledge base, both ontologies and people face the same challenges when reasoning from an inadequate knowledge base. Knowledge in these situations

does not entail ‘truth’ but requires search for accuracy. Our objective in this article is therefore to explore, at Marr’s (1982) computational level of analysis, the difficulties created for computers and people when reasoning from imperfect knowledge. The computational level “provides an understanding of how a mechanism functions in broader environments that determines the computations it needs to perform (and may fail to perform)” (Bechtel & Shagrir, 2015, p. 31). Marr was concerned with why a function needs to be computed, which requires examining the tasks that need to be performed in the environment.

A primary role of the computational level is to constrain functions that must be computed at the algorithmic and representational level (Cooper & Peebles, 2015). This role enables the computational level to serve as an abstract theory that deliberately avoids lower-level commitments. Cooper and Peebles (2015) proposed that, in addition to Bayesian approaches, a variety of other approaches including formal logic can serve as a foundation for formulating problems at the computational level. Others have shown how logical theories can constrain hypothesis about processing all the way down to the neuroscience level (Baggio, van Lambalgen, & Hagoort, 2015)

Our objective is to analyze the challenges for both people and machines when reasoning from imperfect knowledge. Section 2 describes three major tools in the Information Sciences – SUMO, WordNet, and FrameNet – that can support the construction of computational-level analyses of reasoning. Section 3 summarizes 7 types of imperfect knowledge to illustrate the problems created when reasoning from ambiguous, conditional, contradictory, fragmented, inert, misclassified, and uncertain knowledge. Sections 4 through 7 elaborate on 4 of these types (ambiguous, contradictory, misclassified, uncertain) that are particularly important for constructing ontologies. Section 8 discusses using ontologies in cognitive science as practical tools and as idealized cognitive architectures. The final section contains concluding remarks.

2. Information Science Tools

In a previous article we provided a general introduction to how tools in the Information

Sciences – WordNet, FrameNet, and the Suggested Upper Merged Ontology (SUMO) – can be applied to psychology (Reed & Pease, 2015). WordNet is a large, lexical database that, although initially developed by George Miller to study language acquisition, has become the major source of linguistic definitions for the Information Sciences (Fellbaum, 2010; Miller & Fellbaum, 2007). FrameNet is a database of event schemas based on a theory of frame semantics developed by the linguist Charles Fillmore (Fillmore & Baker, 2010). It is closely related to the theoretical construct of scripts developed by Schank and Abelson (1977). SUMO is a formal ontology of concepts expressed in mathematical logic that enables deductive reasoning (Niles & Pease, 2001; Pease, 2011). We first provide an overview of SUMO before explaining why its connection to WordNet and FrameNet enhances its capabilities for prescriptive and descriptive modeling.

2.1 Suggested Upper Merged Ontology (SUMO)

SUMO (<http://www.ontologyportal.org>) is a large corpus of concepts with definitions in mathematical logic. It covers a very broad range of topics at a high level of generality and also includes significant details of many specific domains as diverse as finance, biological viruses, geography, and automobile parts. Fig. 1 shows the hierarchical organization of some of the major terms in SUMO. There are obvious parallels between the organization of knowledge in Fig. 1 and important theoretical concepts in psychology. For instance, SUMO partitions **Entity** into **Physical** and **Abstract**. A **Physical** entity has a location in space/time and is partitioned into **Object** and **Process**. An **object** is a physical thing like a chair or a glass of water. A **Process** is an action that occurs over time like a university lecture. An **Abstract** entity cannot exist at a particular place in space/time without some physical encoding or embodiment. Subcategories of **Abstract** include **Quantity**, **Attribute**, **Relation**, and **Proposition**.

SUMO's hierarchical organization can be viewed either in a top-down or a bottom-up

manner on its web site Clicking on [a tree view \(http://54.183.42.206:8080/sigma/TreeView.jsp?kb=SUMO&term=Entity\)](http://54.183.42.206:8080/sigma/TreeView.jsp?kb=SUMO&term=Entity) enables a top-down exploration of the hierarchy. A user can see the subclasses of a term by clicking on the + sign in front of a term. Clicking on the + in front of the top-level term **Entity** shows its two subclasses **Physical** and **Abstract**. Clicking on the + in front of the term **Physical** shows its two subclasses **Object** and **Process**. Continuing to click on a + in front of a term reveals the next subclass.

In contrast to most other products called ontologies that are often simple taxonomies with natural language definitions, SUMO is richly axiomatized in higher-order logic. SUMO contains 20,000 named concepts with 80,000 statements that define those concepts in a computable logic. It has been the basis of experiments in automated inference, automated natural language processing, and a range of applications in commercial domains to formalize and organize diverse databases.

SUMO has also been mapped manually to all of the WordNet lexical-semantic database as well as to the FrameNet corpus of lexical-semantic schema. These mappings help to make SUMO useful in the semantic interpretation of language, and also clarify the separation between (1) models of the lexicon in different languages (Multi-Lingual WordNet), (2) models of the semantic restrictions between linguistic elements (FrameNet) and (3) a logical model of reality independent of language (SUMO). This comprehensive model allows us to describe the world, relate the world to how we communicate about it, and draw new conclusions from descriptions using the model with automated deductive inference.

2.2 Connecting WordNet to SUMO

SUMO has been manually mapped to all 117,000 synonym sets in WordNet (<https://wordnet.princeton.edu/>). The mapping was first performed in 2002 (Niles & Pease, 2003) and has been continually updated in the subsequent 14 years as SUMO has expanded. This work acted as a check on the conceptual coverage of SUMO because WordNet is approximately the

size and scope of a standard collegiate English dictionary. Any gaps in coverage were solved by adding new concepts and definitions to SUMO.

The SUMO Search Tool also supports bottom-up viewing by finding WordNet definitions and showing their link to an equivalent or subsuming term in SUMO. For instance, WordNet has six definitions of the word *table* (<http://54.183.42.206:8080/sigma/WordNet.jsp?word=table&POS=1>). One definition is ‘a piece of furniture having a smooth flat top that is usually supported by one or more vertical legs’. This definition has an equivalent mapping to **Table** in Sumo, defined as ‘a piece of **Furniture** with four legs and a flat top. It is used either for eating, paperwork, or meetings’. Clicking on the highlighted term **Furniture** reveals that it is a subclass of **Artifact**, which can be clicked to reveal that it is a subclass of **Object**. Another WordNet definition of *table* is ‘a set of data arranged in rows and columns’. This definition is subsumed by **ContentBearingObject** in SUMO, which is part of the SUMO hierarchy displayed in Fig. 1.

The distinction between an electronic dictionary such as WordNet and a formal ontology such as SUMO can be illustrated by the syllogism on a shirt worn by a Silicon Valley gas-station attendant:

I am nobody.

Nobody is perfect.

I am perfect.

The conclusion is logically valid but semantically suspect. Why? WordNet defines *nobody* as a person without influence, which is subsumed by **Human** in SUMO. In the first premise the word *nobody* refers to a particular human – the person wearing the shirt in this situation. In the second premise the word *nobody* refers to the empty set of **Humans**. The empty set excludes any individual from being perfect. The conclusion therefore does not follow from the premises when based on the different ontological uses of the word *nobody*.

Fig. 2 shows three SUMO axioms that formalize this analysis. The first two axioms interpret the English word “nobody” as the class of UnimportantPerson, a specific kind of Human, and that the particular station attendant belongs to that class. The third axiom states that there does not exist any Human that has the attribute of being Perfect. The query then asks whether the StationAttendant has the attribute of Perfect and a logical reasoner such as E (Schulz, 2002) or Vampire (Riazanov & Voronkov, 2002) will easily conclude not.

2.3 Connecting FrameNet to SUMO

FrameNet (<https://framenet.icsi.berkeley.edu>) exists between the level of language and formality that exemplify WordNet and SUMO. It addresses the restrictions among words in different contexts of communication. It does not address the meaning of individual words, nor the formal definitions and the implications of utterances. But it does address mutual constraints that are both conventional and semantic in the English language. As such, it can serve to inspire further logical formalization in SUMO.

An example is the Adopt_selection frame:

An Agent has an Attribute that can take different Values, and the Agent selects and then begins to use some Value for that Attribute. The Attribute includes both physical (shape, color) and abstract (belief, practice) types.

Core values, such as Agent, Attribute, and Value, are shown in capital letters. An example is that there are times when teachers should adopt a neutral stance on morally sensitive issues. Frames also have noncore values that include Circumstances, Duration, Manner, Purpose, Role, and Time for the Adopt_selection frame. The noncore frame values are helpful for identifying the context in which a frame is evoked.

One advantage of both FrameNet and SUMO is that there has been work on linking them (Scheffczyk, Pease, & Ellsworth, 2006). Reed and Pease (2015) provide more details of this integration in their article on constructing cognition ontologies using WordNet, FrameNet, and SUMO.

3. Imperfect Knowledge

Constructing and maintaining a prescriptive ontology requires periodic updating by experts to accurately reflect the most recently accepted scientific knowledge (Smith & Ceusters, 2010). In contrast to the expertise encoded into an ontology, Keil (2012) argued that people of all ages have an impoverished understanding of the world. This impoverished understanding, which he labels 'folk science', progresses from identifying causal properties to organizing those properties to identifying functional relations to possessing a full mechanistic understanding. People are nonetheless overconfident in their understanding of mechanism for devices such as a zipper, sewing machine, speedometer, and VCR because they confuse function with mechanism (Rozenblit & Keil, 2002).

Descriptive models of folk science should include reasoning from imperfect knowledge. Table 1 lists seven kinds of imperfect knowledge that pose challenges for both people and machines. The challenge of reasoning from *ambiguous* knowledge is to resolve the ambiguity. Ambiguous knowledge lacks clarity. As stated by the author of a book on the history of housing in a reference to Central America:

Academics call this portion of the New World Mesoamerica, an accomodatingly vague term that could fairly be defined as Central America plus as much or as little of North and South America as needed to support a hypothesis (Bryson, 2010, p. 38).

Resolving an ambiguity often requires selecting an interpretation that best fits the context. We will look at contextual effects in Section 4.

The challenge of reasoning from *conditional* knowledge is to identify the conditions that distinguish between correct and incorrect applications of the knowledge. A special case is transfer. Although much research has investigated situations in which transfer is appropriate, overzealous transfer occurs when there is an inappropriate focus on prior knowledge (Schwartz,

Chase, & Bransford, 2012). The pedagogical challenge is therefore to create instruction that encourages positive transfer without creating overgeneralizations that block new learning.

The challenge of reasoning from *contradictory* knowledge is to discover contradictions when they are not immediately obvious. Formal ontologies search for contradictions among the axioms encoding the knowledge (Pease, 2011). People face the same challenge when examining their own knowledge. For propositions such as 'John smokes' and 'John does not smoke' the contradiction is transparent. In other cases the contradiction is less transparent and may go undetected. Undetected contradictions are particularly likely when deduction is required to reveal the contradiction. We examine reasoning from contradictory knowledge in Section 5.

The challenge of reasoning from *fragmented* knowledge is to integrate related knowledge stored in LTM (long-term memory). An example of fragmentation is diSessa's (2013) knowledge-as-pieces theory, which shows how isolated knowledge limits understanding principles in physics. One consequence of fragmented knowledge is that observations about the world, such as daytime occurs when a region faces the sun and nighttime occurs when a region does not face the sun, are not connected to causal interpretations.

The challenge of reasoning from *inert* knowledge is to activate it when it is needed. Like fragmented knowledge, inert knowledge has difficulty making connections. Unlike fragmented knowledge, inert knowledge is not well connected to the current situation rather than to other knowledge in LTM. Inert knowledge is knowledge stored in LTM that is relevant for interpreting a situation but is unnoticed because it is not perceived as relevant. An example occurred during a training session for management consultants who failed to relate their own experiences to a lesson on contingent contracts (Gentner, Lowenstein, Thompson, & Forbus, 2009). Those consultants who compared the similarities of two contingent contracts were successful in creating a more abstract schematic understanding that helped them retrieve other (episodic) cases based on their experiences.

The challenge of reasoning from *misclassified* knowledge is to identify and reclassify it. Misclassified knowledge is more difficult to correct when it is organized into different intuitive ontologies. Students may know both the attributes of objects and the attributes of processes but mistakenly identify a process as an object (Chi & Slotta, 1993). Experts in their study appropriately used process-based attributes to explain physics problems and matter-based attributes to explain material substance problems. Novices incorrectly used matter-based attributes to explain both types of problems. Section 6 discusses reasoning from misclassified knowledge.

The challenge of reasoning from *uncertain* knowledge is to incorporate probabilities into reasoning. Probabilities can be either merged or integrated with logic. Psychology's new paradigm of reasoning merges probabilities with logic by assessing the degree of confidence in logical assertions (Elqayam & Over, 2013). The computer extraction of meaning from text integrates probability with logic by combining statistics with ontologies (Pujara, Miao, Getoor, & Cohen, 2015). Both the merge and integration of probabilities with logic is the topic of Section 7.

We decided to focus on ambiguous, contradictory, misclassified, and uncertain knowledge because these are central problems that must be solved when constructing ontologies. After discussing these three types of knowledge in more detail, we conclude by commenting on the practical use of ontologies by cognitive scientists.

4. Ambiguous Knowledge

Ambiguous knowledge, such as the geographic region encompassed by MesoAmerica, lacks clarity. The challenge in reasoning from ambiguous knowledge is to achieve clarity by resolving the ambiguity. The two cases discussed in this section – language and categories – are primary sources of ambiguous knowledge. Inferring the intended meaning often depends on contextual constraints as illustrated by the following examples.

4.1 Language Ambiguity

A word often has more than one meaning, as illustrated by the many senses of a word in WordNet. A classic study of how word ambiguity influences comprehension occurred in

research that used the following pair of sentences as a typical example:

Rumor had it that, for years, the government building had been plagued with problems. The man was not surprised when he found several “bugs” in the corner of his room (Swinney & Hakes, 1976, p. 686)

The ambiguous word “bugs” slowed comprehension relative to a control condition in which the unambiguous word “insects” replaced the ambiguous word “bugs”. However, comprehension was not delayed when the phrase “roaches, spiders, and other bugs” replaced the word “bugs”. In this case the context made it clear that bugs referred to insects rather than to an electronic device used by spies.

Subsequent research established that the effect of a clarifying context did not eliminate the activation of an unintended meaning of an ambiguous word. Rather, it enabled the listener to more quickly select the intended meaning following the activation of both meanings (Swinney, 1979). This finding influenced the formulation of the Construction-Integration model of comprehension in which multiple meanings of a word are activated during the initial construction phase (Kintsch & Mangalath, 2011). The meaning that best fits the context of the sentence is then selected during the subsequent integration phase.

WordNet and SUMO can combine to explain the effect of a clarifying context in this example. WordNet has 5 senses of the word “bug”, including 2 in which **Insect** is the subsuming class in SUMO. Another sense, a small hidden microphone, has **Microphone** as a subsuming class. The word “roach” has 5 senses in WordNet, one of which is subsumed by **Insect**. The word “spider” has 3 senses in WordNet, one of which is also subsumed by **Insect** in SUMO. The phrase “roaches, spiders, and other bugs” should therefore clarify that “bugs” refers to “insects” because all three terms share **Insect** as a subsuming class.

Ambiguities also occur for phrases. “The tall girl is the only one that...” is ambiguous with respect to whether the word *one* refers to “the only girl that” or “the only individual that”. Such

statements have been studied for how variables such as context and informativeness influence the interpretation of the ambiguity (Foppolo, Marelli, Meroni, & Gualmini, 2015). Linguistic context has been explored in the ontology literature but rarely with a computable representation because providing such representations without external recourse to human intuition requires expressive logics. Expressive logics are first-order logic and above (Enderton, 2001), including various modal logics (Chellas, 1980) but especially higher-order logics (Enderton, 2009).

Another potential ambiguity is the adjective “tall” in the noun phrase “The tall girl”. Pease (2011) presents an interpretation of Parsons' (1990) context parameter for adjectives within SUMO and higher-order logic. This approach formalizes the context of the adjective as being over a constructed class. So, a tall girl may be tall for a 12-year-old girl but short compared to the set of female humans. Fig. 3 displays a class-forming operator **KappaFn** to relate the class of all 12-year-old girls to the proposition that one particular girl is tall. The axiom states that the girl is tall in comparison to the class of entities that are 12-year-old human females. So “tall”, which is a nearly meaningless adjective without context, now is given a very precise context over which it applies.

4.2 Category Ambiguity

Another source of ambiguity that has been extensively studied by cognitive scientists is category membership. In a formal ontology class membership can be broad or general but not ambiguous. For example **Process** is a precise and general term for an action that takes place at a particular place and time (which can be a point or a range). It is very general, but not ambiguous. A class in SUMO can have particular instance members, but they do not comprise the definition of a given class. This is in contrast to **Sets**, which are defined by their members.

Hampton (2007) defines a vague concept as one that picks out a category with no clearly defined boundaries. He links vagueness to the theoretical construct of typicality, which reflects how representative an exemplar is of its category. Hampton proposed that vagueness is

the inevitable result of a knowledge system that stores the centers, rather than the boundaries, of conceptual categories. SUMO provides explicit extensions definitions of its terms, using rules that state the boundaries of each class. It does not rely on typicality or confuse the concept of ambiguity with generality. The notion of Mammal, for example, is general but not ambiguous, probabilistic, or vague.

Ambiguous knowledge can have practical consequences. A famous example in the legal literature is a hypothetical rule prohibiting vehicles in public parks (Spellman & Schauer, 2012). Although cars are clearly prohibited, what about bicycles and baby carriages? What about trucks that need to make deliveries? A judge might try to determine the purpose behind the rule but she would still have to struggle with interpreting the meaning of *vehicle*. The lack of research on the professional judgments of lawyers and judges (Kahan, 2015) contributes to the mystery.

Decisions based on ambiguous knowledge might benefit from a computational-level analysis of the ambiguity. The judge could begin with a definition of *vehicle*. The noun *vehicle* in WordNet has four senses, the most relevant for this context is 'a conveyance that transports people or objects'. It maps onto the equivalent term **Vehicle** in SUMO, which provides a more specific definition: 'Vehicle is the subclass of **TransportationDevices** that transport passengers or goods from one place to another by moving from one place to the other with them, e.g., cars, trucks, ferries, and airplanes. Contrast with devices such as pipelines, escalators or supermarket checkout belts, which carry items from one place to another by means of a moving part, without the device removing from the origin to the destination.'

The SUMO definition provides a helpful beginning but does not address the category typicality issue discussed by Hampton (2007). Rosch and Mervis (1975) investigated typicality by asking participants to rank order 20 vehicles, all of which are consistent with SUMO's definition of **Vehicle**. The top ten in rank order were car, truck, bus, motorcycle, train, trolley car, bicycle, airplane, boat, and tractor. Most typical vehicles are good candidates for exclusion

except perhaps for tractor because, like wheelbarrow ranked 19th, could be used to maintain the park. Interpretation of 'excluded vehicles' therefore would need to examine how the category exemplars are used in this situation.

A law could better be written with respect to a precise definition, rather than a notion of typicality that could easily be dependent upon different contexts and experiences of different individuals. A law should apply equally to people regardless of their intuitions. Laws need not have exhaustive enumerations to avoid recourse to intuition. In this example, the prohibition against vehicles could specifically exclude any vehicle that is used to maintain the park. A rich formal theory such as SUMO can address what objects are *for* – their teleology – not just what they are.

Ross (1997) raised the general issue of whether actions should depend on classification alone or on the combination of classification and use. He argued that including use as an attribute in classification does not necessarily lead to a more complicated theory of categorization because it may help people focus more quickly on relevant features. Focusing on relevant features is also related to Barsalou's (1991) theory of ad hoc categories – categories assembled to fulfill a particular goal. Although a goal (going on a camping trip) may require collecting exemplars across superordinate categories such as food and shelter, the goal in this case (using vehicles to maintain a park) requires selecting exemplars within a superordinate category.

WordNet has 10 senses of the word *maintain*. The one most appropriate for this context is to 'maintain for use and service', which maps onto **obligation** as a subsuming term in SUMO. In addition, case roles are an essential tool for describing processes in SUMO (Pease, 2011). Some case roles are **agent, destination, experiencer, origin, patient, instrument, resource, result, path, and direction**. Consider the sentence "A park worker uses a wheelbarrow to maintain the rose garden". The statement is a 'maintenance'

event that has ‘park worker’ in the **agent** slot, ‘wheelbarrow’ in the **instrument** slot, and ‘rose garden’ in the **destination** slot. Notice that case roles can be combined as in FrameNet to specify co-occurrences among words.

We can use those **CaseRole(s)** in definitions that employ a higher-order logical relation such as **hasPurpose** and **hasPurposeForAgent** that relate objects to formulas stating their purpose, or more narrowly we can simply say that only vehicles actually engaged in maintenance are allowed (excluding perhaps someone just riding their tractor through a park for the fun of it). For example, we might define a park rule that has an exception to the “no vehicles” rule, when the vehicle in question is meant for maintaining the park. The axiom in Fig. 4 defines the relations among Policy, Park, Vehicle, and Maintenance.

A common denominator in many attempts to clarify ambiguities is to use the context to infer the intended meaning. A knowledge-based system called Scone has been designed for the specific purpose of emulating human-like reasoning in multiple contexts (Fahlman, 2011). Each context represents a different world-model but a context inherits most of its knowledge from another context, explicitly representing only the differences. Every concept and every link in Scone has a connection to a context node. This multi-context mechanism supports reasoning about hypothetical or counterfactual situations, changes over time, and the knowledge and beliefs of other agents.

5. Contradictory Knowledge

Another type of imperfect knowledge occurs when statements are unambiguous but contradictory. The challenge in reasoning from contradictory knowledge is to recognize and correct harmful contradictions but utilize helpful contradictions. Section 5.1 discusses discovering harmful contradictions in a faulty knowledge base. Section 5.2 shows how helpful contradictions can aid theorem proving.

5.1 Discovering Harmful Contradictions

Failure to recognize contradictions occurred for a computer program called STUDENT that was studied by Paige and Simon (1966). STUDENT directly translates text into equations but does not have semantic knowledge. It therefore could not detect a contradiction in the following problem:

The number of quarters a man has is seven times the number of dimes that he has. The value of the dimes exceeds the value of the quarters by two dollars and fifty cents. How many has he of each coin?

STUDENT can nonetheless serve as a model of those students who also miss a contradiction because they use a direct translation strategy rather than a mental model of the problem (Mayer & Hegarty, 1996).

Fig. 5 shows 9 premises that capture algebraic knowledge and information in the problem. A student who used a direct translation strategy could (in principle) construct the following equation by applying Premises 3 to 6:

$$\text{Number of dimes} \times \$0.10 = \$2.50 + (7 \times \text{Number of dimes}) \times \$0.25$$

Solving the equation for the Number of dimes yields -1.52 dimes, which contradicts Premise 1 that the number of dimes cannot be a negative number.

An alert student might notice a contradiction in the problem without constructing an equation by reasoning from Premises 7 to 9. Application of Premises 7 and 8 yields Total value (quarters) > Total value (dimes), which contradicts Premise 9 from the problem that Total value (dimes) > Total value (quarters). The contradiction in this case depends on qualitative reasoning that if there are more quarters than dimes, the value of the dimes cannot exceed the value of the quarters.

A strength of ontologies such as SUMO is that they can employ automated logical reasoning to find faulty knowledge. Logical reasoning uses basic logical symbols of “-” for “not”, “^” for “and”, and “v” for “or”. It also uses “=>” for “implies”. This symbol can also be read as “if ___ then ___”, so “A => B” when read as “if A then B”. These notations are used to construct

logical rules.

A simple rule is double negation in which $\neg\neg A$ is the same as A . There are also transformation rules. In standard first-order logic, a rule that states “ A implies B ” (or “if A then B ”) is logically equivalent to “ $\neg A \vee B$ ”. Another transformation is deMorgan's rule that $\neg(A \wedge B)$ is the same as $(\neg A) \vee (\neg B)$. The [statement says](#) that [either A or B \(or both\)](#) must be false. Other inference rules combine two statements. If we know A or B is true, and we also know that $\neg A$ is true, we can conclude B is true. We can also think of this from the perspective of an implication (rule). If we know $A \Rightarrow B$ and we know A , then we can conclude B . This rule is known as Modus Ponens.

The example in Fig. 6 (adapted from Nelson, 1999) shows how a machine can be useful in finding faulty knowledge that might not be immediately obvious to a human. The notation near the top of the figure makes the statements more succinct and illustrates how linguistic statements can be transformed into mathematical expressions that a machine can follow without recourse to human intelligence and knowledge about words, language and the world. The proof finds a contradiction in the premises by concluding that (1) Heather is going to be fired and (2) Heather is not going to be fired.

The proof in Fig. 6 reveals that premises 3 (The boss likes Heather) and 6 (the boss doesn't fire anyone he likes) ultimately lead to the conclusion that Heather will not be fired ($\neg F$) whereas premises 1, 2, 4, and 5 lead to the conclusion that Heather will be fired (F). So, we might consider correcting the knowledge by removing premises 3 and 6, since perhaps the boss believes more in merit than just having likable employees. Removing premises 3 and 6 removes lines 14 and 15 in the proof resulting in the conclusion in line 13 that Heather will be fired.

The example in Fig. 6, however, does not show the most current modern approach to theorem proving that involves proof by refutation. The next section explains this approach.

5.2 Formulating Useful Contradictions

A general obstacle to finding contradictions is the large number of required comparisons

as the knowledge base increases in size. Although there have been significant advances in the speed of theorem proving on large knowledge bases for both first-order and higher-order logics, significant challenges remain. The scope of the search space is such that it is not guaranteed that a solution will be found in bounded time, even if such a solution exists. The SInE algorithm for axiom relevance detection (Hoder & Voronkov, 2011) has resulted in [one of the most significant recent advances in theorem proving over large knowledge bases](#) and the annual Automated Systems Competition on Automated Deduction (Sutcliffe & Suttner, 2006) continues to drive the research community to higher performance.

Modern theorem proving is proof by refutation that begins by assuming the opposite of a question and then finding a resulting contradiction (Harrison, 2009). Proofs resulting from these methods are in the form of counterfactuals. For example, assuming the scenario of an Agatha Christie mystery, if we ask "Did the butler do it?" we wind up with a chain of reasoning that says things like "If we assume the butler didn't do it, then Ms. Marple must have been in the dining room." Discovering that Ms. Marple was not in the dining room would implicate the butler.

Proof by refutation has been known for many years. What has changed in recent years is that implementations of automated theorem proving on modern computers have become fast enough to be practical, having response times often of just a few seconds even for theories like SUMO having tens of thousands of logical statements (Pease, Sutcliffe, Siegel, & Trac, 2010). And having SUMO, WordNet and FrameNet together provides resources for an automated or semi-automated approach for expressing linguistic information in a computable logical form, with a significant amount of background knowledge that doesn't have to be re-encoded for every new problem that researchers want to solve.

There are nonetheless challenges in presenting such proofs for understanding by nonexperts. Proofs by refutation are unnatural for human reasoning. This has led to the field of "natural deduction" (Prawitz, 2006), which unfortunately has not had a degree of computational success comparable to that of modern proofs by contradiction.

The greater success of computers than people in executing proofs by refutation illustrates that computers are often more skillful than people at complex problem solving. The success of computer programs was demonstrated when IBM's Deep Blue defeated Gary Kasparov in chess, IBM's Watson defeated experts on Jeopardy, and DeepMind's AlphaGo defeated the European Go champion (Kitano, 2016). However, achieving human-level general intelligence is still a distant goal for AI (Adams et al., 2012).

6. Misclassified Knowledge

The previous section discussed one important aspect of ontologies – deductive reasoning from a knowledge base. This section discusses another important aspect of ontologies – the classification of knowledge. The challenge in reasoning from misclassified knowledge is to discover and reclassify it. Misclassifications are evident in intuitive ontologies as documented by Chi (2008, 2013). Her distinction between hierarchical, lateral, and ontological errors determines the organization of this section.

6.1 Hierarchical Misclassifications

One of the six types of imperfect knowledge identified in Section 3 was fragmented knowledge as represented by diSessa's (2013) knowledge-as-pieces theory. His theory contrasts with theories that propose knowledge involves a more coordinated use of ideas. Examples of the latter approach include Carey's (2009) work on conceptual development, Vosniadou's (2013) work on conceptual change, and Chi's (2013) work on ontological categories. diSessa, Sherin, and Levin (2014) refer to these coordination classes as integrated systems consisting of strong internal constraints that are unlike the fragmentation and independence of phenomenological primitives.

An intuitive ontology proposed by Chi (2013) to describe novice knowledge distinguished between three distinct ontological trees depicting *Entities*, *Processes*, and *Mental States* (Fig. 7). *Entities* in her ontology consist of *Objects* and *Substances*, rather than serve as the top node

of a single hierarchy as they do in SUMO. *Processes* occur over time and can be *Sequential* or *Emergent*. *Mental States* are abstract, in one's mind. Chi's knowledge-as-coherence perspective shows the importance of ontologies as theoretical constructs. Her goal is "not to lay out the exact ordering and structure of hierarchical and lateral categories and trees, nor to decide which categories deserve the name ontology, or how many ontologies there are" (Chi, 2013, p. 65). Rather, it is to provide an illustration of intuitive categorical structures.

One type of error that can result from intuitive ontologies is a *hierarchical* error (Chi, 2008). Hierarchical errors occur within the same branch of a hierarchy. Categories within a branch – such as chairs, furniture, artifacts, and objects – enable new categories to inherit properties of the more generic categories higher in the hierarchy. However, the inheritance of properties is imperfect and can result either in failing to inherit, or incorrectly inheriting, a property of a higher category. Failing to inherit a property is demonstrated by how beliefs about matter can differ from beliefs about a particular instance of matter (Shtulman & Valcarcel, 2012). This can result in contradictions such as "all matter has heat" but "ice does not have heat".

Another example, illustrated in Fig. 7, involves identifying situations in which both *Artifacts* and *Living Beings* should inherit the properties of *Objects*. While many studies of early knowledge development investigated learning the differences between animate and inanimate objects (Carey, 2009; Mandler, 2004) some studies investigated situations in which animate and inanimate objects should be classified as equivalent. For example, 7-year-old children were more proficient than 5-year-old children in appropriately transferring physics principles between animate and inanimate objects when instructed on either a child or a rock (Heyman, Phillips, & Gelman, 2003). Property inheritance from the **Object** category would occur in SUMO because **SentientAgent**, defined as an **Agent** that is capable of perception and experiences some level of consciousness, is subsumed by **Object** in Fig. 1.

Misclassifications are also likely to occur when a category does *not* inherit all the proper-

ties of its subsuming category. In 2006, for the first time in its history, the International Astronomical Union defined the characteristics of a 'planet' in our Solar System and consequently demoted Pluto from 'planet' to 'dwarf planet' (Messeri, 2010). The society defined a 'planet' as a celestial body that: (a) is in orbit around the Sun, (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, and (c) has cleared the neighborhood around its orbit. A 'dwarf planet' has the first two attributes but not the third. The classification of Pluto as a dwarf planet creates ontological problems because 'dwarf planet' would typically be a subclass of 'planet' and therefore inherit its attributes, including a cleared neighborhood around its orbit.

There is a precedent in Cognitive Science, however, for a subclass not inheriting all of the attributes of its subsuming class. An attribute of a category exemplar can override the attributes of its category in the hierarchical network model proposed by Collins and Quillian (1969). For instance, the attribute 'cannot fly' would be stored at the exemplar level for penguins and ostriches, overriding the attribute 'fly' stored at the 'bird' level. Classification of other birds that do not fly would therefore benefit from comparing their similarity to an instance such as penguin rather than to the attributes of a subsuming category (Johansen, Savage, Fouquet, & Shanks, 2015).

Instead of overriding what is essentially a linguistic ambiguity, SUMO resolves it with deeper and more precise modeling of the real world, rather than conforming to how informal natural language has evolved to refer to entities in the world. **Bird** is simply a **WarmBloodedVertebrate** and not defined as having the **capability** of **Flying**. SUMO then employs a further division of subclasses to create **FlyingBird** and **FlightlessBird** where only the former has that capability. Multiple inheritance allows us to have both the conventional biological taxonomy in which the biological order of Sphenisciformes (penguins) and groups such as Ratites are flightless and are from different branches of the taxonomy. This illustrates well how taxonomies can only classify with respect to one characteristic at each subdivision.

With modern advances in DNA analysis, the biological taxonomy is undergoing significant transformation, but without a richer representational formalism, biologists will be forced to classify with respect either to evolutionary progress (as expressed in DNA) or by physical features or capabilities, as has been the case prior to DNA based analysis.

6.2 Lateral and Ontological Misclassifications

Hierarchical misclassifications that occur within a branch are either too specific or too general. In contrast, lateral misclassifications occur across branches within an ontology such as classifying an *Artifact* as a *Living Being* in Fig. 7. An additional type of misclassification, ontological errors, occurs across ontologies such as classifying a *Process* as an *Object* (Fig. 7). Chi (2008) argued that the consequences of hierarchical misclassifications are mild when compared with lateral or ontological mistakes:

The central question to pose about lateral and ontological categories is the cost of category mistakes. We define a category mistake as the case when a concept has been assigned inappropriately to a lateral or alternative ontological category. In contrast to incorrect hierarchical categorization, category mistakes are damaging in that categorical inferences and attributions will be erroneous, creating a barrier to correct learning with deep understanding (Chi, 2008, p. 65).

There are both similarities and differences between the three ontological trees in Fig. 7 and the SUMO ontology. One major difference is that SUMO and other formal ontologies consist of a single ontological tree. The advantage is that it is therefore possible to see where Chi's top nodes – *Entities*, *Processes*, and *Mental States* – fit within a single ontology. As indicated in Fig. 1, both **Object** and **Process** are a **Physical Entity** in SUMO. A *Mental State* is an **Abstract Entity** that corresponds to **StateofMind**, which is subsumed by **PsychologicalAttribute** in SUMO. Although embedding **Object** and **Process** in distinct ontological trees highlights their differences, embedding them in a single ontological tree does not weaken Chi's

argument regarding category mistakes. Both terms occupy lateral categories in SUMO. In addition, SUMO clarifies their similarity. Both are a **Physical Entity**.

The issue of a single versus multiple ontologies is also an issue in Vosniadou and Skopelli's (2014) framework theory. At times they refer to multiple ontologies, as in "requires the creation of new ontologies" (p. 1436), but at other times refer to changing classifications within a single ontology as in "undergo fundamental ontological re-categorization" (p. 1429) and "assigned to the wrong ontological category" (p. 1439). Relating concepts across multiple intuitive ontologies is a problem for the learner who has to learn how concepts in one ontology are related to concepts in another ontology. It is also a problem for the knowledge theorist.

The mapping and merging of ontologies is a major computational challenge for the Information Sciences. As suggested by its name, the Suggested Upper Merged Ontology began as a project by IEEE to merge publicly available ontological content into a single, comprehensive, and cohesive structure (Niles & Pease, 2001). Research on ontological alignment has produced a wide range of methods that include estimating lexical similarity between concepts and identifying structural similarities between ontologies using tree-based representations (Raad & Evermann, 2015).

Raad and Evermann (2015) argue that, although these current approaches at alignment are based on well-justified logical, linguistic, and statistical procedures, the methods lack an underlying theoretical foundation that can be provided by work on human analogical reasoning. In particular, they propose that LISA (Learning and Inference using Schemas and Analogies) offers such a foundation (Hummel & Holyoak, 1997). Raad and Evermann link ontology alignment with analogical reasoning by mapping terms in LISA such as Proposition, Predicate, Role, and Object onto logical statements that are suitable for the Web Ontology Language OWL. Although tools in the Information Sciences should inform work in the Cognitive Sciences (Reed & Pease,

2015), the Cognitive Sciences can also inform work in the Information Sciences as illustrated by the use of LISA as a theoretical foundation for ontology alignment.

7. Uncertain Knowledge

A type of imperfect knowledge that is currently receiving extensive attention within both psychology (Elqayam & Over, 2013) and computer science (Jordan & Mitchell, 2015) is uncertain knowledge. Uncertain knowledge requires incorporating probabilities into reasoning and we focus on combining probabilities with logic. Section 7.1 discusses merging probabilities with logic by considering the probability that a logical statement is true. Section 7.2 discusses integrating probabilities with logic by combining the bottom-up processing of uncertain data with the top-down constraints imposed by ontologies.

7.1 Merging Probabilities with Logic

In 2013 a special issue of *Thinking & Reasoning* introduced a new paradigm of reasoning to merge logical deduction with probability judgments (Elqayam & Over, 2013). The deduction paradigm in psychology had typically required drawing conclusions about logical validity. Logical validity is based on the binary distinction between true and false. In contrast, probabilistic validity is more general and captures the confidence of beliefs as represented by various degrees of subjective probability.

There is an approach to reasoning with uncertainty that is compatible with deductive reasoning. It is a more common solution to reason about uncertainty numerically, but that is not the only solution. SUMO supports saying that one proposition is more likely than another, as opposed to saying that some proposition has a likelihood value of 0.765 and another 0.423, for example. Creating what is called a "partial order" of likelihood allows reasoning within a normal two-valued logic, and making assertions about the relative a priori and a posteriori likelihoods of particular concluded expressions (Cohen, 1985). This does however require a logic that is beyond first-order logic. To handle heuristic uncertainties as well as modal and epistemic logical constructs, SUMO employs a higher-order logic.

Evans and Over (2013) argued that the distinction between induction and deduction should nonetheless be maintained in the new paradigm of reasoning. Induction adds new beliefs and deduction draws out implications from those beliefs. Valid deductive inferences preserve truth by arriving at conclusions that are only implicit in the premises. Strong inductive inferences, in contrast, go beyond the premises to conclusions that only likely follow. The conclusions in this case are necessarily more uncertain than their premises.

An example of the merge of probability with logic is that the conditional statement *if p then q* becomes a probability statement $P(\text{if } p \text{ then } q)$. For example, what is the probability that if global warming continues London will be flooded? An empirical prediction, attributed to Rips and Marcus (1977), is that people will evaluate $P(\text{if } p \text{ then } q)$ as equivalent to $P(q|p)$. A number of studies have supported the Rips and Marcus hypothesis (Elqayam & Over, 2013). $P(q|p)$ acquires an important status in Bayes theorem for revising an hypothesis (h) based on new evidence (e). The revised probability of h , $P(h|e)$, depends in part on the probability of obtaining the evidence if the hypothesis is true, $P(e|h)$.

The contributors to the special issue identified several key questions for the new paradigm of reasoning. One is the role of normative theories. Logic provides a normative theory for deductive reasoning and Bayes theorem provides a normative theory for revising probabilities. Should normative theories be incorporated into the merge and, if so, how? Another question asks how to combine the separate literatures on deductive and inductive reasoning.

One promising approach to answering these questions is to embed induction and deduction within the dual process framework of reasoning (Evans & Over, 2013). The authors propose that use of the deductive reasoning paradigm in contemporary research encourages slow, deliberate Type 2 reasoning. This type of reasoning, however, would be difficult to carry out for ordinary reasoning from beliefs. Everyday heuristic reasoning in their proposal primarily occurs at the fast and intuitive Type 1 level.

7.2. Integrating Probabilities with Logic

The separate literatures on deductive and inductive reasoning in psychology have a parallel in computer science. Inductive reasoning to organize data has been the domain of machine learning (Jordan & Mitchell, 2015). Deductive reasoning to draw inferences from knowledge has been the domain of ontology. How can these two literatures be integrated to join the bottom-up analysis of data with the top-down constraints imposed by ontologies?

We intentionally use the term “integrate” to distinguish it from the term “merge” in the previous section. The new paradigm of reasoning merges probability and logic into a single mathematical entity such as $P(\text{if } p \text{ then } q)$. In contrast, machine learning and ontology can work together while preserving their separate identities.

A recent effort on deriving knowledge from text provides an excellent case study (Pujara et al., 2015). Information-extraction techniques face an overwhelming problem of interpreting erroneous and incomplete information. The use of semantic information provided by ontologies can aid the interpretation.

Fig. 8 illustrates how this occurs through a knowledge graph construction process. Box 1 shows conflicting candidate facts produced by an information-extraction system. Continuous truth-values represent the degree of confidence for the various candidate facts such as ‘Kyrgystan’ and ‘Kyrgystan Republic’ are countries. Box 2 uses these data to construct an uncertain extraction graph. Entities are shown in rectangles, labels are shown in circles, and uncertain information is shown as dashed lines.

Semantic dependencies are added in Box 3. Ontological information, shown by dotted lines, reveals that ‘bird’ and ‘country’ are mutually exclusive (Mut) categories so ‘Kyrgystan’ cannot be both a ‘bird’ and a ‘country’. Domain (Dom) links countries to continents. The relation ‘Same’ establishes the degree of confidence that two different names refer to the same entity. When two entities are very similar any label assigned to the first entity will also be assigned to the second entity. Box 4 shows the completed knowledge graph in which both ‘Kyrgyzstan’ and the ‘Kyrgyz Republic’ are labeled as a country located in Asia.

8. Using Ontologies in Cognitive Science

There are two primary approaches to using ontologies in cognitive science. The more traditional approach is to use an ontology as an applied tool for organizing and retrieving knowledge about cognition. A second approach is to use an ontology as a theoretical model. This use may be somewhat unique to the field of cognitive science because cognitive science is also concerned with the organization of knowledge. This section discusses both approaches.

8.1. *Ontologies as Information Science Tools*

Electronic devices provide overwhelming amounts of data about human behavior but the data require organization, analysis, and interpretation. To better prepare psychology students for this new revolution, Griffiths (2015) proposed that learning programming skills will be as important as learning experimental design. Others have argued that the emphasis on calculus as the culminating mathematical experience leaves U. S. students ill-equipped at posing questions that lead to innovations in computation (Li & Bishop, 2016).

Ontologies provide examples of innovations in computation. An advantage of constructing a formal ontology is its application as an information-retrieval and decision-support tool for a wide variety of knowledge domains (Hoehndorf et al., 2012). In order to be compliant with the use of SUMO as an upper ontology, a mid-level ontology on cognition would have to satisfy several criteria that are listed in 'Conformance testing for SUMO' <http://www.ontologyportal.org/conform.html>. A conforming implementation (1) uses terms as defined by SUMO, (2) or uses terms that are defined entirely by other terms in SUMO, and (3) is consistent; a contradiction cannot be derived by means of first-order logic from the set of statements in its implementation of SUMO.

Although the axioms of an ontology enable deductive reasoning they will only be as accurate as the knowledge base. As stated at the beginning of this article, ontologies strive to serve as a prescriptive model of how knowledge should be organized based on the most recent advances in science. Achieving accuracy is particularly demanding in areas of 'unsettled

science' in which there are multiple camps of experts who have different perspectives (Smith & Ceusters, 2010). Psychology would likely be a strong candidate for the 'unsettled science' classification.

Psychologists should nonetheless attempt to build knowledge bases that document what they know about psychology. In our introductory article (Reed & Pease, 2015) we proposed approximately eight premises for each of seven categories (attention, external attention, internal attention, capacity, cognitive load, selection, conscious awareness) to illustrate a sample knowledge base for the concept 'attention'. A taxonomy of attention proposed by Chun, Golomb, and Turk-Browne (2011) provided our starting point. The premises included definitions, empirical findings, and theoretical arguments. An example of each includes:

- Vigilance is the process of paying close and continuous attention (WordNet).
- Selecting stimuli at an early stage based on sensory information requires less mental effort than selecting stimuli at a late stage based on meaning (Johnston & Heinz, 1978).
- The central executive in Baddeley's working memory model controls attention (Baddeley, 2000).

A major question is who should create the premises and then possibly the axioms. One promising approach in philosophy is the Indiana Philosophy Ontology project (Buckner, Niepert, & Allen, 2011) that uses the *Stanford Encyclopedia of Philosophy* as an online, open access knowledge base that is continuously updated (<http://plato.stanford.edu>). *Encyclopedia* entries have been submitted by over 1,300 volunteer authors and then reviewed by approximately 115 volunteer subject editors. The Indiana Philosophy Ontology project uses a combination of automated methods and expert feedback to create an ontology for this knowledge base.

Cognitive Psychology currently lacks such a major effort. Based on the assumption that group knowledge is less subjective than individual knowledge (Aroyo & Welty, 2015) organizations such as the Cognitive Science Society and the Association for Psychological

Science should take the initiative by forming working groups to formulate and organize knowledge about Psychology. Progress in building ontologies in other domains, particularly Bioinformatics, is rapidly leaving Psychology behind (Yarkoni, 2012).

8.2. Ontologies as Idealized Computational Architectures

Although ontologies have been designed to aid human reasoning by organizing knowledge, they can also support theory development in Cognitive Science by serving as idealized computational architectures. Griffiths, Lieder, and Goodman (2015) distinguished between three levels of cognitive architectures based on an idealized process model at the computational level, a more realistic process model at an intermediate level, and an actual process model at the algorithmic level. They proposed beginning with an ideal solution that has no resource limits and then explore the consequences of this assumption by adding resource limits to produce a more psychologically plausible model.

An example of a resource limitation is the limited capacity of working memory. However, the major resource limitation is a lack of knowledge. An unlimited capacity working memory would be wasted if it lacked accurate knowledge. The role of an ontology as an ideal computational architecture would be to provide a knowledge hierarchy sufficient for executing a task, axioms that formally specify relations among the terms of the ontology, and a logic that enables deductions from the axioms. Adding resource limitations would place limitations on either the knowledge hierarchy, the axioms expressing relations among terms, or the logical rules for deriving new knowledge.

We have seen how limitations of the knowledge hierarchy in Chi's (2013) intuitive ontologies can result in misclassifications at an inappropriate level in the same vertical branch. A more severe limitation can result in misclassifications in different vertical branches. An even more severe limitation can result in misclassifications in different ontologies, as illustrated in Fig. 7.

Knowledge limitations in using logical rules also exist at various levels of expertise.

Some people could construct the proof in Fig. 6 to find a contradiction in the premises. Others would be unable to construct the proof but could understand the steps in the displayed proof. Still others would have difficulty following the steps in Fig. 6. For both the knowledge base and the deductive rules that draw inferences from the knowledge base, formal ontologies can serve as idealized computational architectures. Adding resource limitations to the ontology can then produce descriptive models of human behavior (Griffiths et al., 2015).

8. Concluding Remarks

Our goal in this article has been to articulate some of the challenges faced by both people and machines when confronted with imperfect knowledge. Making imperfect knowledge less imperfect can benefit from input by both sources. Machines can express knowledge in a clear and logical form but may require human intervention to aid interpretation of the input.

For instance, the axiom in Fig. 3 makes it clear that a particular girl is tall compared to other 12-year-old girls but that may not have been the intention of the speaker. The axiom in Fig. 4 states that maintenance vehicles are exceptions to the no-vehicle policy but the policy's intention may be to reduce pollution by prohibiting all gas-powered vehicles. The contradiction in Fig. 6 can be resolved by eliminating axioms 3 and 6 (as we suggested) but eliminating other axioms could also resolve the contradiction. Reasoning from axioms creates a need to make knowledge less ambiguous, contradictory, and uncertain through the joint efforts of people and machines.

Davenport and Kirby (2016) discuss this joint effort in their recent book *Only humans need apply: winners and losers in the age of smart machines*. Three categories of winners that are most relevant to reasoning from imperfect knowledge are described in chapters on stepping up, stepping in, and stepping forward. People who step up manage problem solving at a higher cognitive level than machines because they understand the bigger picture. People who step in

rely on machines to make routine decisions but correct their mistakes and tweak them for better performance. People who step forward build the next generation of smart machines.

Although both the Information and the Cognitive Sciences contribute to reasoning from knowledge, including imperfect knowledge, there are few cross-references between these two domains. Given the central importance of knowledge in both disciplines, we find this lack of communication surprising. This article reflects on reasoning from imperfect knowledge from the perspectives of both the Information and the Cognitive Sciences to establish a link between the two domains.

Acknowledgements

The authors would like to thank reviewers who made many helpful suggestions. The first author was a visiting scholar in the Department of Psychology at UCSD while writing the manuscript.

Table 1

Challenges for Reasoning from Imperfect Knowledge

Ambiguous knowledge. The challenge is to recognize and resolve the ambiguity. Resolving the ambiguity can be difficult but context can provide a hint.

Conditional knowledge. The challenge is to identify conditions that distinguish between correct and incorrect applications of knowledge. These conditions support the appropriate transfer of knowledge.

Contradictory knowledge. The challenge is to discover a contradiction when it is not immediately obvious. It can be difficult to find contradictions in a large knowledge base.

Fragmented knowledge. The challenge is to integrate related knowledge stored in LTM. One consequence of fragmented knowledge is that observations about the world are not connected to their causes.

Inert knowledge. The challenge is to activate relevant knowledge when it is needed. Inert knowledge is knowledge stored in LTM that is relevant for interpreting a situation but is unnoticed because it is perceived as irrelevant.

Misclassified knowledge. The challenge is to identify and reclassify it. Misclassified knowledge is difficult to correct when it is organized into intuitive ontologies that are nonetheless partially effective in producing plausible explanations.

Uncertain knowledge. The challenge is to incorporate probabilities into reasoning. Probabilities can be either merged or integrated with logic.

Figure Captions

Figure 1. Hierarchical organization of (a) selected Physical Entities in SUMO (b) selected Abstract Entities in SUMO Indentation indicates subclasses.

Figure 2. SUMO axioms that conclude that the station attendant is not perfect.

Figure 3. A SUMO axiom that states a particular girl is tall compared to other 12-year-old girls.

Figure 4. A SUMO axiom that states maintenance vehicles are an exception to the no-vehicle rule.

Figure 5. A SUMO axiom that states maintenance vehicles are an exception to the no-vehicle rule.

Figure 6. Logical proof establishing a contradiction that Heather is, and is not, going to be fired.

Figure 7. Intuitive ontologies. From Chi (2013).

Figure 8. An illustration of the knowledge graph construction process. Based on Pujara, Miao, Getoor, & Cohen (2015)

Physical Entities**Object****SelfConectedObject****Substance****CorpuscularObject****ContentBearingObject****SymbolicString****Icon****LinguisticExpression****Agent****SentientAgent****Region****Process****InternalChange****QuantityChange****StateChange****IntentionalProcess****OrganizationalProcess****DualObjectProcess****Combining****Motion****Path**

Figure 1a.

Abstract Entities**Quantity****Number****PhysicalQuantity****ConstantQuantity****LengthMeasure****MassMeasure****FunctionQuantity****Attribute****RelationalAttribute****NormativeAttribute****TruthValue****SetOrClass****Relation****SpatialRelation****Traverses****BinaryRelation****Causes****Proposition****Procedure****Plan**

Figure 1b.

1. The station attendant is an unimportant person.

```
(instance StationAttendant UnimportantPerson)
```

2. An unimportant person is human.

```
(subclass UnimportantPerson Human)
```

3. There does not exist any human who is perfect.

```
(not  
  (exists (?X)  
    (and  
      (instance ?X Human)  
      (attribute ?X Perfect))))
```

```
query: (attribute StationAttendant Perfect)
```

Figure 2.

**(mannerInContext Tall Sarah
(KappaFn ?X
(and
(instance ?X Human)
(age ?X (MeasureFn 12 Years))
(attribute ?X Female))))**

Figure 3.

**(exists (?P)
(and
(instance ?P Policy)
(containsInformation ?P
(modalAttribute
(exists (?V ?P)
(and
(instance ?V Vehicle)
(instance ?P Park)
(located ?V ?P)
(not
(exists (?M)
(and
(instance ?M Maintaining)
(patient ?M ?P)
(instance ?M ?V))))))
Prohibition))))**

Figure 4.

Coin Problem

The number of quarters a man has is seven times the number of dimes that he has. The value of the dimes exceeds the value of the quarters by two dollars and fifty cents. How many has he of each coin?

Premises

1. Number of dimes ≥ 0
2. Number of quarters ≥ 0
3. Total value (dimes) = Number of dimes x Value of a dime (\$.10)
4. Total value (quarters) = Number of quarters x Value of a quarter (\$.25)
5. Number of quarters = 7 x Number of dimes [from problem]
6. Total value (dimes) = \$2.50 + Total value (quarters) [from problem]
7. If Number of quarters > Number of dimes then Total value (quarters) > Total value (dimes)
8. Number of quarters > Number of dimes [from problem]
9. Total value (dimes) > Total value (quarters) [from problem]

Figure 5.

Premises

1. Either Heather attended the meeting or Heather was not invited.
2. If the boss wanted Heather at the meeting, then she was invited.
3. The boss likes Heather.
4. Heather did not attend the meeting.
5. If the boss did not want Heather there, and the boss did not invite her there, then she is going to be fired.
6. The boss doesn't fire anyone he likes.

Abbreviated notation

A = Heather Attended the meeting

I = Heather was Invited to the meeting

W = the boss Wanted Heather at the meeting

F = Heather is going to be fired

L = Boss Likes Heather

Proof Establishing a Contradiction

- | | |
|---|---------------------------|
| 1. $A \vee \neg I$ | Premise 1 |
| 2. $W \Rightarrow I$ | Premise 2 |
| 3. L | Premise 3 |
| 4. $\neg A$ | Premise 4 |
| 5. $\neg W \wedge \neg I \Rightarrow F$ | Premise 5 |
| 6. $L \Rightarrow \neg F$ | Premise 6 |
| 7. $\neg W \vee I$ | 2: def'n of \Rightarrow |
| 8. $\neg(\neg W \wedge \neg I) \vee F$ | 5: def'n of \Rightarrow |
| 9. $\neg\neg W \vee \neg\neg I \vee F$ | 8: deMorgan's rule |
| 10. $W \vee I \vee F$ | 9: double negation |
| 11. $\neg I$ | 1,4: Modus Ponens |
| 12. $\neg W$ | 7,11: Modus Ponens |
| 13. F | 11,12,10: Modus Ponens |
| 14. $\neg F$ | 3,6: Modus Ponens |
| 15. Contradiction | 13,14 |

Figure 6.

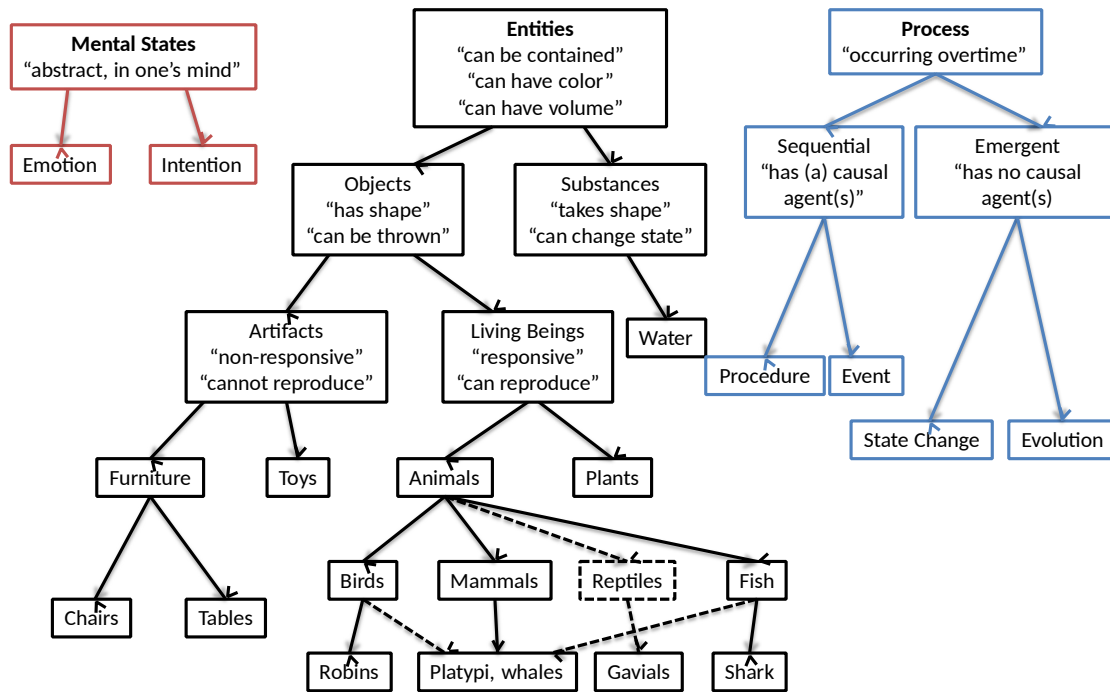


Figure 7.

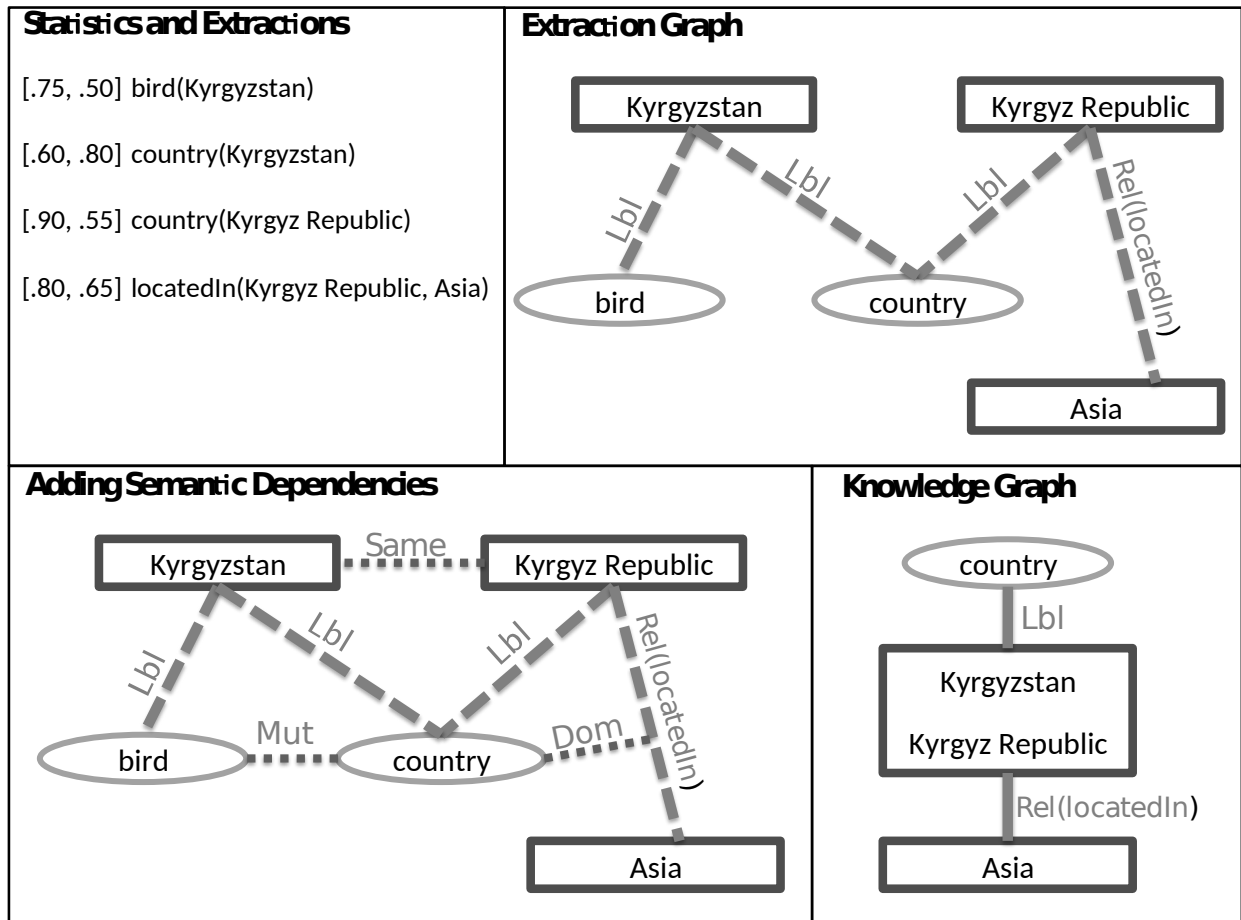


Figure 8.

References

- Adams, S. S., Arel, I., Bach, J., Coop, R., Furlan, R., Goertzel, B., . . . Sowa, J. (2012). Mapping the landscape of human-level artificial general intelligence. *AI Magazine*, 33, 25-41.
- Aroyo, L., & Welty, C. (2015). Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36, 15-24.
- Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4, 417-423.
- Baggio, G., van Lambalgen, M., & Hagoort, P. (2015). Logic as Marr's computational level: four case studies. *Topics in Cognitive Science*, 7, 287-298.
- Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 27, pp. 1-64). San Diego: Academic Press.
- Bechtel, W., & Shagrir, O. (2015). The non-redundant contributions of Marr's three levels of analysis for explaining information-processing mechanisms. *Topics in Cognitive Science*, 7, 312-322.
- Bryson, B. (2010). *At Home: A Short History of Private Life*. New York: Doubleday.
- Buckner, C., Niepert, M., & Allen, C. (2011). From encyclopedia to ontology: toward dynamic representation of the discipline of philosophy. *Synthese*, 182, 205-233.
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
- Chellas, B. (1980). *Modal logic: An introduction*. Cambridge, UK: Cambridge University Press.
- Chi, M. T. H. (2008). Three types of conceptual change: Belief revision, mental model transformation, and categorical shift. In S. Vosniadou (Ed.), *Handbook of research on conceptual change* (pp. 61-82). Hillsdale, NJ: Erlbaum.
- Chi, M. T. H. (2013). Two kinds and four sub-types of misconceived knowledge, ways to change it and the learning outcomes. In S. Vosniadou (Ed.), *International handbook of research on conceptual change* (2nd ed., pp. 49-70). New York: Routledge.
- Chi, M. T. H., & Slotta, J. D. (1993). The ontological coherence of intuitive physics. *Cognition*

- and Instruction*, 10, 249-260.
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, 62, 73-101.
- Cohen, P. R. (1985). *Heuristic reasoning about uncertainty: an artificial intelligence approach*. Marshfield, MA: Pitman Publishing, Inc.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8, 240-248.
- Cooper, R. P., & Peebles, D. (2015). Beyond single-level accounts: The role of cognitive architectures in cognitive scientific explanations. *Topics in Cognitive Science*, 7, 243-258.
- Davenport, T. H., & Kirby, J. (2016). *Only humans need apply: winners and losers in the age of smart machines*. New York: HarperCollins.
- diSessa, A. A. (2013). A bird's-eye view of the "pieces" vs. "coherence" controversy (from the "pieces" side of the fence). In S. Vosniadou (Ed.), *International Handbook of Research on Conceptual Change* (2 ed., pp. 31-48). New York: Routledge.
- diSessa, A. A., Sherin, B., & Levin, M. (2014). Knowledge analysis: An introduction.
- Elqayam, S., & Over, D. E. (2013). New paradigm psychology of reasoning: An introduction to the special issue edited by Elqayam, Bonnefon, and Over. *Thinking & Reasoning*, 19, 249-265.
- Enderton, H. (2001). *A mathematical introduction to logic* (2 ed.). New York: Academic Press.
- Enderton, H. (2009). Second-order and higher-order logic. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. Stanford, CA.: The Metaphysics Research Lab.
- Fahlman, S. E. (2011). Using Scone's multiple-context mechanism to emulate human-like reasoning *Proceedings of the AAAI Fall Symposium on Advances in Complex Systems*. Arlington, VA: Association for the Advancement of Artificial Intelligence.
- Evans, J. S. B. T., & Over, D. E. (2013). Reasoning to and from belief: Deduction and induction

- are still distinct. *Thinking and Reasoning*, 3, 267-283.
- Fellbaum, C. (2010). WordNet. In R. Poli, M. Healy, & K. Achilles (Eds.), *Theory and applications of ontology: computer applications* (pp. 231-243). Heidelberg, Germany: Springer.
- Fillmore, C. J., & Baker, C. F. (2010). A frames approach to semantic analysis. In B. Heine & H. Narrog (Eds.), *The oxford handbook of linguistic analysis* (pp. 313-340). Oxford: Oxford University Press.
- Foppolo, F., Marelli, M., Meroni, L., & Gualmini, A. (2015). Hey little sister, who's the only one? Modulating informativeness in the resolution of privative ambiguity. *Cognitive Science*, 39, 1646-1674.
- Gentner, D., Lowenstein, J., Thompson, L., & Forbus, K. D. (2009). Reviving inert knowledge: Analogical encoding supports relational retrieval of past events. *Cognitive Science*, 33, 1343-1382.
- Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135, 21-23.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7, 217-229.
- Hampton, J. A. (2007). Typicality, graded membership, and vagueness. *Cognitive Science*, 31, 355-384.
- Harrison, J. (Ed.) (2009). *Handbook of practical logic and automated reasoning*. New York: Cambridge University Press.
- Heyman, G. D., Phillips, A. T., & Gelman, S. A. (2003). Children's reasoning about physics within and across ontological kinds. *Cognition*, 89, 43-61.
- Hoder, K., & Voronkov, A. (2011). Sine Qua non for large theory reasoning CADE'11 *Proceedings of the 23rd international conference on automated deduction* (pp. 299-314).

Berlin: Springer-Verlag.

Hoehndorf, R., Dumontier, M., & Gkoutos, G. V. (2012). Evaluation of research in biomedical ontologies. *Briefings in Bioinformatics*, *14*, 696-712.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427-466.

Johansen, M. K., Savage, J., Fouquet, N., & Shanks, D. R. (2015). Salience not status: How category labels influence feature inference. *Cognitive Science*, *49*, 1594-1621.

Johnson-Laird, P. N. (2012). Inference with mental models. In K. J. Holyoak & R. G. Morrison (Eds.), *The oxford handbook of thinking and reasoning* (pp. 134-154). New York: Oxford University Press.

Johnston, W. A., & Heinz, S. P. (1978). Flexibility and capacity demands of attention. *Journal of Experimental Psychology: General*, *107*, 420-435.

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, *349*, 255-260.

Kahan, D. M. (2015). Laws of cognition and the cognition of law. *Cognition*, *135*, 56-60.

Kintsch, W., & Mangalath, P. (2011). The construction of meaning. *Topics in Cognitive Science*, *3*, 346-370.

Kitano, H. (2016). Artificial intelligence to win the Nobel prize and beyond: Creating the engine for scientific discovery. *AI Magazine*, *37*, 39-49.

Li, T. M., & Bishop, A. (2016, March 5-6). Calculus is so last century. *The Wall Street Journal*, p. A11.

Mandler, J. M. (2004). *The foundations of mind: Origins of conceptual thought*. Oxford, England: Oxford University Press.

Marr, D. C. (1982). *Vision: A computational investigation into the human representational system and processing of visual information*. San Francisco: Freeman.

Mayer, R. E., & Hegarty, M. (1996). The process of understanding mathematical problems. In R.

- J. Sternberg & T. Ben-Zeev (Eds.), *The nature of mathematical thinking* (pp. 29-54). Mahwah, NJ: Erlbaum.
- Messeri, L. R. (2010). The problem with Pluto: Conflicting cosmologies and the classification of planets. *Social Studies of Science*, 40, 187-214.
- Miller, G. A., & Fellbaum, C. (2007). WordNet then and now. *Language Resources & Evaluation*, 41, 209-214.
- Nelson, R. (1999). Course Slides for "CSC 173: Computation and Formal Systems". Retrieved March 15, 2016
https://www.cs.rochester.edu/~nelson/courses/csc_173/proplogic/resolution.html
- Niles, I., & Pease, A. (2001). Toward a standard upper ontology. In C. Welty & B. Smith (Eds.), *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)* (pp. 2-9). Ogunquit, Maine.
- Niles, I., & Pease, A. (2003). Linking lexicons and ontologies: Mapping WordNet to the Suggested Upper Merged Ontology. *Proceedings of the IEEE International Conference on Information and Knowledge Engineering*, 412-416.
- Paige, J. M., & Simon, H. A. (1966). Cognitive processes involved in solving algebra word problems. In B. Kleinmuntz (Ed.), *Problem solving: Research, method, and theory*. New York: John Wiley & Sons.
- Parsons, T. (1990). *Events in the semantics of English: A study of subatomic semantics*. Cambridge, MA: MIT Press.
- Pease, A. (2011). *Ontology: A Practical Guide*. Angwin, CA: Articulate Software Press.
- Pease, A., Sutcliffe, G., Siegel, N., & Trac, S. (2010). Large reasoning with SUMO at CASC. *AI Communications*, 23, 137-144.
- Prawitz, D. (2006). *Natural deduction: A proof-theoretic method*. Mineola, NY: Dover Publications.
- Pujara, J., Miao, H., Getoor, L., & Cohen, W. W. (2015). Using semantics and statistics to turn

- data into knowledge. *AI Magazine*(Spring), 65-74.
- Raad, E., & Evermann, J. (2015). The role of analogy in ontology alignment: A study of LISA. *Cognitive Systems Research, 33*, 1-16.
- Reed, S. K., & Pease, A. (2015). A framework for constructing cognition ontologies using WordNet, FrameNet, and SUMO. *Cognitive Systems Research, 33*, 122-144.
- Riazanov, A., & Voronkov, A. (2002). The design and implementation of Vampire. *AI Communications, 15*, 91-110.
- Rips, L. J. (1983). Cognitive processes in propositional reasoning. *Psychological Review, 90*, 38-71.
- Rips, L. J., & Marcus, S. L. (1977). Suppositions and the analysis of conditional sentences. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension* (pp. 185-219). New York: Wiley.
- Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573-605.
- Ross, B. H. (1997). The use of categories affects classification. *Journal of Memory and Language, 37*, 240-268.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science, 26*, 421-562.
- Schank, R., & Abelson, R. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.
- Scheffczyk, J., Pease, A., & Ellsworth, M. (2006). *Linking FrameNet to the Suggested Upper Merged Ontology*. Paper presented at the Proceedings of the 2006 conference on Formal Ontology in Information Systems.
- Schulz, S. (2002). E - A Brainiac theorem prover. *Journal of AI Communications, 15*, 111-126.
- Schwartz, D. L., Chase, C. C., & Bransford, J. D. (2012). Resisting overzealous transfer. *Educational Psychologist, 7*, 204-214.

- Shtulman, A., & Valcarcel, J. (2012). Scientific knowledge suppresses but does not supplant earlier intuitions. *Cognition*, *124*, 209-215.
- Smith, B. F., & Ceusters, W. (2010). Ontological realism: A methodology for coordinated evolution of scientific ontologies. *Applied Ontology*, *5*, 139-188.
- Spellman, B. A., & Schauer, F. (2012). Legal reasoning. In K. J. Holyoak & R. J. Morrison (Eds.), *The Oxford Handbook of thinking and reasoning* (pp. 719-735). New York: Oxford University Press.
- Sutcliffe, G., & Suttner, C. (2006). The State of CASC. *AI Communications*, *19*, 35-48.
- Swinney, D. A. (1979). Lexical access during sentence comprehension: Reconsideration of
- Swinney, D. A., & Hakes, D. T. (1976). Effects of prior context upon lexical access during sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, *15*, 681-689.
- Vosniadou, S. (2013). Conceptual change in learning and instruction: The framework theory approach. In S. Vosniadou (Ed.), *International handbook of research on conceptual change* (2 ed., pp. 11-30). New York: Routledge.
- Yarkoni, T. (2012). Psychoinformatics: New horizons at the interface of the psychological and computing sciences. *Current Directions in Psychological Science*, *21*, 391-397.