

Designing Visual Languages for Description Logics

Brian R. Gaines
Department of Computer Science
University of Victoria
Victoria BC V8W 2Y2, Canada

Abstract

Semantic networks were developed in cognitive science and artificial intelligence studies as graphical knowledge representation and inference tools emulating human thought processes. Formal analysis of the representation and inference capabilities of the networks modeled them as subsets of standard first-order logic (FOL), restricted in the operations allowed in order to ensure the tractability that seemed to characterize human reasoning capabilities. The graphical network representations were modeled as providing a visual language for the logic. Sub-sets of FOL targeted on knowledge representation came to be called *description logics*, and research on these logics has focused on issues of tractability of subsets with differing representation capabilities, and on the implementation of practical inference systems achieving the best possible performance. Semantic network research has kept pace with these developments, providing visual languages for knowledge entry, editing, and presenting the results of inference, that translate unambiguously to the underlying description logics. This paper discusses the design issues for such semantic network formalisms, and illustrates them through detailed examples of significant generic knowledge structures analyzed in the literature, including determinables, contrast sets, genus/differentiae, taxonomies, faceted taxonomies, cluster concepts, family resemblances, graded concepts, frames, definitions, rules, rules with exceptions, essence and state assertions, opposites and contraries, relevance, and so on. Such examples provide important test material for any visual language formalism for logic.

Keywords: Visual languages, description logics, semantic networks

1 Introduction

This paper reports on part of a long-term research program aimed at the development and application of knowledge acquisition tools for knowledge-based systems that has been concerned with formal visual languages representing logical models of knowledge structures. Our early studies in the 1980s targeted *expert system shells* as inference engines (Gaines and Linster, 1990) and found that many extraneous structures had to be added to manage what was not logically well-founded inference (Gaines and Shaw, 1993). With the development of logical foundations for KL-ONE semantic network systems (Brachman, 1977), leading through various implementations to the well-defined CLASSIC specification (Borgida, et al., 1989), it became attractive to target logically sound knowledge representation systems, and we implemented a formal visual language (Gaines, 1991b) for the semantics of CLASSIC, and a knowledge representation server, KRS, that was a modular, extensible, C++ implementation of CLASSIC based on an algebraic model of the logical semantics (Gaines, 1993).

Since CLASSIC did not support negation or disjunction, we developed an integrated inference engine for rule-based reasoning, focusing particularly on rules with exceptions since these

proved to be natural to our experts (Gaines, 1991a). This system was used extensively to develop a number of knowledge based systems solving a range of types of problem (Gaines, 1994), and versions of it are in widespread use to support constructivist educators and our current research on the representation of knowledge structures in the arts, sciences and engineering (Gaines, 2003; Gaines, 2004).

As CLASSIC-like systems evolved through the formalization of description logics (Donini, 2003), the development of sound and complete algorithms for those supporting negation and disjunction (Horrocks, 1998), and the implementation of inference engines for them, such as Racer (Haarslev and Moller, 2001) and Pellet (Parsia and Sirin, 2004), we extended and changed the visual language to provide a knowledge acquisition front-end to any description logic inference engine. This has involved many design decisions between competing visual representation schema, each having some merits, and requiring trade-offs between different ways of representing the same underlying logical knowledge structure.

This paper analyzes the design issues involved in specifying a visual language that supports semantic network representations of knowledge structures which map precisely into description logic primitives, providing them with formal semantics and a two-dimensional logical notation that translates to standard logical expressions. The issues are illustrated using a visual language editor supported by scripts that translate visual specifications of logical structures into a range of standard description logic notations, but the emphasis is not so much on a particular visual language as on the family of possible visual language and the issues involved in the design decisions relating to them. The overall objectives are to present a logico-linguistic analysis of the syntactic, semantic, pragmatic and semiotic issues involved in the design of visual languages for description logics, and to demonstrate that semantic networks, rigorously formalized in logical terms, provide useful tools for ontology engineering activities in a wide range of applications including education, psychology, philosophy, modeling conceptual change, and developing knowledge-based systems.

Section 2 discusses eight major issues in designing visual language formalisms for logics. Section 3 details the design requirements for each logical construct in a semantic network visual language, illustrating their use to represent significant generic knowledge structures and the translations to description logic formalisms that are generated. Section 4 summarizes the way in which the issues of Section 2 are addressed in Section 3, and outlines future research directions.

2 Issues in Designing Visual Languages for Description Logics

In analyzing the design of visual languages for description logics, it is useful to clarify a number of issues.

2.1 Relationship of Semantic Networks to Other Logical Formalisms

Semantic networks of labeled nodes of various types connected by lines of various types to represent knowledge were seen in the early years of natural language processing (Richens, 1956; Quillian, 1967) not only to represent knowledge but also as providing extra-logical inference mechanisms. Such connotations were criticized by Hayes (1977) who argued that “modern formal logic is the most successful language ever developed to express human thought and inference...yet recent writers in the AI field have been almost unanimous in their condemnation of logic as a representational language.” He demonstrated that the frame structures of semantic

networks could be seen as syntactic variants of standard FOL formulae (Hayes, 1979). Woods (1975) had already criticized the early semantic network literature for lack of unambiguous interpretations, and for logical confusions between intension and extension. He notes that, “there is currently no theory of semantic networks,” and “even the question of what the networks have to do with semantics is one which takes some answering.”

Brachman (1977) addressed Woods’ criticism by developing logical foundations for interpreting semantic networks, in particular making a clear distinction between the intensional logic of terminological definitions, and the extensional logic of asserting properties of individuals. This distinction led to the separation of the ‘T-box’ and ‘A-box’ in later implementations of representation and inference systems deriving from semantic networks. It became accepted in the 1980s that semantic networks should be treated as a visual language for some subset of FOL without any extra-logical inference capabilities, and that inference systems based on them should be assessed in terms of soundness and completeness in the same way as an any other proof system for a formal logic.

The studies reported in this paper have taken the position that semantic networks should be treated as a formalism for a subset of FOL, and hence one design issue is:

Design issue 1: There should be unambiguous translations from a semantic network visual language to standard logical formalisms.

2.2 Value of Semantic Network Formalisms

If semantic networks provide a formalism for a subset of FOL, given the existence of well-established and widely used logical formalisms, one has to make a case for introducing yet another one, in particular taking into account that all logical formalisms are difficult to learn and employ effectively. The effort to learn to use a visual language needs justification in terms of added value, which, as Fricke (2003) has noted in the context of Hyperproof, is extra-logical and does not increase logical power. The issue becomes a psychological one, of whether a two-dimensional diagrammatic presentation of logical expressions is useful to people. This in turn raises further questions about: ease of learning, ease of understanding and ease of creation; interaction with preferred modes of cognition, appropriate to everyone, or just to ‘visual thinkers’; different applications to expressing knowledge, understanding arguments, proofs, explanations, and so on; and the pragmatic and semiological aspects of existing practices in the use of diagrams:

Design issue 2: The design of semantic network visual language constructs needs to take into account the human factors of individual users, and the pragmatic and semiological aspects of their use for communication in user communities.

There is a rich literature on the psychology (Bauer and Johnson-Laird, 1993) and semiology (Bertin, 1983) of diagrams, their roles in logical and mathematical proofs (Jamnik, 2001; Stenning, 2002), and knowledge visualization (Keller and Tergan, 2005), the educational merits of visual presentations (Williams, 1983), and so on, but it largely addresses visual images and geometric constructions that are more naturalistic than the artificial constructions of semantic networks. A smaller literature does specifically address semantic networks in applications such as professional communication (Khalifa and Liu, 2006), education (Jonassen, 2005), experimental cognitive modeling (Barsalou, 1992), modeling conceptual change in the history of

science (Andersen, et al., 2006; Hyman, 2007), the use of concept mapping tools (which are essentially semantic networks without formal semantics) to express knowledge (Novak, 1998), and so on. However, so far, operationalization of the results through the computational application of the networks is uncommon.

At the human factors level, we have found that it is important to provide users with a clear and simple model of the meaning of each graphical construction used in a semantic network. There may be several possible explanations of the meaning of a particular construction in a particular situation, only one of which generalizes to all situations in which it is used.

Design issue 3: Each graphical construction used in a semantic network should have a well-defined meaning that can be simply explained and applies in all situations.

The research on description logics that evolved as part of the formalization of semantic networks has become focused on logical, complexity and inference algorithm issues, and is now indifferent to the form of language used to represent logical formulae. Semantic network tools continue to be developed as adjuncts to description logic servers, and used: to provide a visual language for knowledge entry (Eisenstadt, et al., 1990; Clark, et al., 2001; Gennari, et al., 2003; Ernst, et al., 2005; Hayes, et al., 2005; Fox, et al., 2007); to display the subsumption structure computed by inference systems; and to explain the inference processes and display the inferences by annotation and animation of the semantic network (Gaines and Shaw, 1999). This suggests:

Design issue 4: The design of visual languages for semantic networks needs to take into account the various roles that diagrams in the language will play in problem solving activities, such as knowledge entry, knowledge base representation and the explanation of inference.

One needs to assess alternative design choices by comparing their impact on the visual representation of a range of significant knowledge structures. One way of doing this in a principled fashion that is not domain specific is to take the various fundamental generic knowledge structures studied in metaphysics, and compare these in different semantic network and conventional logical formalisms. This suggests:-

Design issue 5: The design of visual languages should be tested by comparing alternative representations, both for realistic examples and for generic knowledge structures such as determinables, contrast sets, genus/differentiae, taxonomies, facets, cluster concepts, family resemblances, graded concepts, frames, definitions, rules, rules with exceptions, essence and state assertions, opposites and contraries, relevance, and so on.

2.3 Description Logics as Subsets of FOL with Known Tractability

Another question that arises relates to why one should target description logics rather than FOL in general, and this relates to another parallel development in the history of description logics. While Hayes' arguments for FOL as a knowledge representation language were highly influential (Charniak and McDermott, 1986), its undecidability made it an unrealistic model for the emulation of human cognitive processes and even for attempting to go beyond them with artificial intelligence. Moreover, studies of computational complexity had raised issues of computational intractability with similar implications to undecidability (Cook, 1971; 1983). Brachman and Levesque (1984) showed that the tractability of subsumption in description logics was highly sensitive to small changes in the representational capabilities of the language, and

“even a seemingly simple frame-based description language can pose intractable computational obstacles.”

The research strategy for the semantic network community became: to define semantic networks with limited representation capabilities such that an inference system could be developed that was sound and complete; for each possible combination of representation capabilities to determine the worst case computational complexity; to implement representation and inference systems with optimal performance given the complexity bounds; and to use them in artificial and real-world test cases to see if meaningful problems could be solved (Donini, 2003). It is this rich, and still developing, formal framework for the tractability of subsets of FOL as knowledge representation languages, coupled with available inference engines, such as the open-source Java implementation of Pellet (Sirin, et al., 2007) and advances in first order extensions of the Davis-Putnam-Logemann-Loveland (DPLL, Davis, et al., 1962) theorem prover algorithms (Baumgartner and Tinelli, 2008), that makes description logic subsets of FOL an attractive target for semantic networks. All the designer of the network formalism need do is ensure that there is unambiguous mapping between the visual formalism and that of the description logic system, and soundness and completeness of inference become issues for that system, not the semantic network.

However, where semantic networks are used for knowledge structure entry it would be useful to encourage users to avoid representation capabilities that lead to increased computational complexity if it is possible to do so.

Design issue 6: The design of visual languages and their support systems should encourage the use of representational capabilities that minimize the complexity of inference and, if possible, discourage the use of those that increase such complexity.

2.4 Support of Services Beyond Deduction: Abduction, Induction, Analogy

There are debates about the appropriateness of standard logic as a basis for metaphysics (Mason, 2000), psychology (Macnamara, 1986; Hanna, 2006) and artificial intelligence (McDermott, 1987), which are relevant to the decision to target a subset of FOL. If an inference system for what appeared to be a more appropriate deviant logic became available it would be simple to amend the translator from semantic networks to target it. What one would require of any inference system is that it can compute whether a set of logical propositions is coherent or not. Other deductive services can be based on this, and it also provides the foundation for a range of other reasoning activities, such as abduction, induction, analogy, planning, counterfactual creative imagination, and so on. Human reasoning uses inconsistency to manage knowledge processes, never accepting *ex falso quodlibet*, but adjusting the knowledge structure to avoid it. If these paraconsistent procedures are axiomatized as an integral part of the proof system of an underlying logic then that logic will deviate from FOL but still be based on it and revert to it when there are no inconsistencies in the knowledge structure. A good example is the adaptive logic (Meheus, 2000), that has been used to model the management of inconsistency in scientific innovation (Meheus, 2003). These considerations suggest:-

Design issue 7: The design of visual languages for semantic networks needs to take into account not only their roles in supporting processes of deductive inference, but also their roles in managing and explaining the use of such inference in abduction, induction, analogy, and so on.

2.5 Support of Modularity

The adjustments to knowledge structures to achieve paraconsistency may be regarded as the replacement of one module of the knowledge structure by another. Those for defeasability may be viewed as including knowledge modules if they are coherent with the known facts. Many other applications to planning, design, and so on, may be modeled as a search over combinations of knowledge modules to determine one or more that are coherent and satisfy some additional criteria. Similar processes occur in the study of conceptual change in science (Andersen *et al.*, 2006) or developmental learning (Pearsall, et al., 1997), where one module of a knowledge structure may be modified, leaving the remainder unchanged. They also arise in the integration of ontologies from different sources (Grau, et al., 2007). This suggests:-

Design issue 8: The design of visual languages for semantic networks needs to take into account the need for modularity such that knowledge structures may be readily split and recombined. The visual language support system should be able to manage such recombination taking into account feedback from the inference system.

3 Design of Semantic Network Formalisms

This section discusses in detail the design of semantic network formalisms taking into account the issues noted above. It introduces the visual constructs used in networks in the context of their use to support the various fundamental knowledge structures noted in design issue 5 above. This makes it clear what can be achieved with each construct, and motivates the incremental addition of other constructs for what cannot be achieved. Examples are given of the translation to a standard logical formalism so that the visual expression can be compared with a conventional one. For the purposes of this paper the abstract syntax of description logics will be used since it is more succinct than the normal syntax for FOL and readily translates to it (noting that, since the description logic syntax has extensional connotations implication is represented to indicate inclusion, i.e. \sqsubseteq corresponds to \supset which can be confusing).

The ensuing discussion will use philosophical and psychological examples to exemplify issues in the design of visual languages for knowledge representation. To the extent that it comments on the associated literature it is only doing so to illustrate language issues, not attempting to encompass the massive literatures on the philosophical and psychological issues involved, but suggesting that visual representation may be useful in clarifying such issues.

The semantic network tool used in the examples is KNet, which is part of the Rep IV family of knowledge acquisition and inference tools (Gaines and Shaw, 2007). The RepNet graphic tool used to implement KNet allows node and line types to be defined and assigned shapes and colors, which enables one to emulate existing node-link formats such as semantic networks, concept maps, conceptual graphs, bond graphs, and so on (Gaines and Shaw, 1995). An associated script language with a just-in-time compiler that receives events from user interaction with the graphic tool enables semantic networks to be created, managed in relation to one another, exported through translation scripts to inference engines, and the results of inference to be shown in the networks. It also supports a DPLL inference engine written in the script language that enables simple expert system problems to be solved and explained without use of more powerful external services.

3.1 Families of Properties, Constructs, Contrast Sets

A fundamental knowledge structure that is discussed in Aristotle's *Categories* and has been given many different terminologies in later literatures is that of a family of disjoint concepts all of which imply some superordinate concept that names the family. Carnap (1950) in his work on the foundations of probability terms it a *family of properties*. Kelly (1955) in his work on the foundations of constructivist psychology terms such a family a *construct*, with the superordinate concept being the construct's *range of convenience*, and the subordinate concepts being its *poles*. In the cognitive anthropology literature the set of disjoint concepts is called a *contrast set* (Frake, 1969), and Kuhn uses such sets to model Wittgenstein's *family resemblances* in terms of a combination of similarity and dissimilarity (Andersen, 2000).

Figure 1 shows part of a taxonomy of color naming as a hierarchy of two families of properties, with the top two levels being a contrast set of *chromatic* and *achromatic* modes of *diffuse reflection*, and the lower two levels being a contrast set of shades from white to black describing achromatic *diffuse reflection* (Hardin, 1989). The diagram adopts the usual conventions for drawing taxonomies, using undecorated text for nodes, nodes in contrast at the same vertical level, descending lines to lower levels of the taxonomy, and a tacit assumption that all nodes in a contrast set represent disjoint predicates.




Figure 1 Contrast sets in part of a color taxonomy

These conventions make the logical structure of the diagram clear to people and could be translated automatically to logical form. To represent the structure as semantic network two cosmetic changes are desirable: to box the concepts (predicates) in some way since unboxed text is generally used to indicate a binary relation; and to replace the undirected lines with arrows of implication so that the spatial relationships are not part of the formal representation, giving greater freedom in layout. A more significant change is to make the disjoint relations between concepts explicit, not only to avoid encoding it in a spatial relationship, but also because the imposition of such a relationship is a feature of scientific taxonomies which are *designed* to be disjoint, and is only partially present in folk taxonomies (Kay, 1975).

The encoding of disjointness requires an indication that one concept implies the negation of the other. One could indicate a negated concept graphically through a different surround, or lexically through a symbol for negation before the concept name. However, this requires both negative and positive forms of the concept to be included in the diagram which clutters it, and also introduces the general negation of any concept which is a problematic notion that is more powerful than needed to express disjointness or opposition (Horn, 1989). Hence, it seems better to introduce another arrow symbol indicating that one concept implies the negation of another. This is a symmetric relationship and at one time we used an undecorated line for disjointness (Gaines, 1991b). However, the logical symmetry masks an underlying cognitive asymmetry that is reflected in linguistics as the distinction between the *marked* and *unmarked* sides of a contrast

(Lyons, 1968), and in Kelly's (1955) psychology as the distinction between the *implicit* and *emergent* poles of a construct. Hence, we have introduced an arrow with a line across it to indicate that one concept is the antonym of other, that is implies its negation, and explain this to users as an arrow of *opposition*. Figure 2 shows the contrast sets of Figure 1 in a semantic network.




Figure 2 Semantic network for part of a color taxonomy

The description logic propositions output by the description logic translator script for this network are:-

```

chromatic ⊆ diffuse_reflection
achromatic ⊆ diffuse_reflection ⊓ ¬chromatic
white ⊆ achromatic
black ⊆ achromatic ⊓ ¬white
gray ⊆ achromatic ⊓ ¬black ⊓ ¬white
  
```

Or, by the Racer translation script:

```

(in-knowledge-base Determinables2 Determinables2)
(implies chromatic diffuse_reflection)
(implies achromatic (and diffuse_reflection (Not chromatic)))
(implies white achromatic)
(implies black (and achromatic (Not white)))
(implies gray (and achromatic (Not black) (Not white)))
  
```

Note that the contrast set under achromatic can be expanded indefinitely by adding additional concepts that have a normal arrow to achromatic and opposition arrows to the prior concepts in the set. This procedure is logically sound but does not encode some additional relationships expected in a contrast set, such as gradability, and better ways of achieving this will be illustrated in the following sections.

The user may make inferences from this network by setting concepts to be *true* (for some implicit individual), *false*, *open*, or open and of *interest*. The user interface is illustrated in Figure 3. Clicking on a concept node activates a popup menu with options dependent on the state of the concept. The truth value is indicated by a vertical bar for *true*, a horizontal one for *false*, neither for *open*, and a dot (not shown) for the concept being of *interest*. This last state places the concept on the agenda of the inference process which normally will not compute the truth value of a node unless it is relevant to some inference goal. These features are used to demonstrate to

students how taxonomic structures may be used to manage relevancy in such a way as to minimize inferential effort.




Figure 3 User interface controlling inference

Left: popup menus for setting truth values when clicking on a concept node

Right: popup menu for controlling inference when clicking outside all nodes

Clicking outside all nodes activates the popup menu shown on the right which controls the inference process: *Infer* propagates truth values; *Infer All* makes all nodes of interest and propagates truth values; *Explore* activates full inference in which possible worlds are explored by hypothesizing possible truth values for open nodes relevant to inference; *Explore All* does this for all open nodes; *Clear Inferences* and *Clear All* set the states of concepts to *open*; *Explanation On* outputs a logical trace of the inference process; *Debug On* outputs a more detailed trace; *Structure* outputs the logical graph structure in textual form; *Unlock* makes the network editable.

Figure 4 is a screen dump of the inferences when the concept “gray” is set true in Figure 2 and *Infer All* is selected in the inference menu. Nodes whose truth value has been inferred are highlighted with a thicker surround.




Figure 4 Inference with the semantic network for part of a color taxonomy

The system explains its reasoning in terms of the inferences made:

Infer from: gray

gray \sqsubseteq achromatic \sqcap \neg black \sqcap \neg white [gray]
achromatic \sqsubseteq diffuse_reflection \sqcap \neg chromatic [achromatic]

It lists the nodes with truth values on the initial line, and then lists each inference made. The concept indicated in square brackets on the right is that being investigated when the inference is made. In this example it is a node whose truth value is being propagated but in other examples it may be an open node whose truth value is being inferred.

3.2 Taxonomies, Determinables, Genus and Differentiae, Frames, Faceted Taxonomies

Figure 2 show how taxonomies normally represented as in Figure 1 can be encoded in a semantic network, and the construction clearly extends to taxonomies of any width or depth. The inference process is not of particular interest because nodes in contrast sets in scientific taxonomies are mutually exclusive and a node being applicable entails that other members of the contrast set are not applicable, and the node above it in the hierarchy is applicable. However, this is a design feature characterizing scientific taxonomies, and hierarchical structures of implication and opposition may be built from the graphic constructs defined above that are not fully taxonomic and model features of ‘folk taxonomies’ (Conklin, 1969). They can also be used to model non-taxonomic relationships such as comparability and gradability between members of a contrast set as will be illustrated in section 3.4.

In the above example the disjointness in the contrast sets is modeled as explicit opposition. This is necessary for what Johnson (1921) termed a *determinable* with *determinates*, where the determinates are the contrast set and the determinable its superordinate concept. He contrasts this with Aristotle’s construction of taxonomies where the superordinate concept is termed the *genus* and contrast set members are disjoint because they have mutually exclusive properties termed the *differentiae*. Color is the stereotypical example of a determinate with determinables, for example, the colors of the rainbow, each of which may itself be a determinate of more specific colors in a contrast set below it (Funkhouser, 2006).

The examples above are of determinable/determinate taxonomies, but the graphic primitives defined can also be used to represent genus/differentiae taxonomies. For example, Figure 5 shows *the tree of Poryphry*, a representation of Aristotle’s taxonomy of substance from the third century used by Sowa (2000) as an early example of a semantic network.

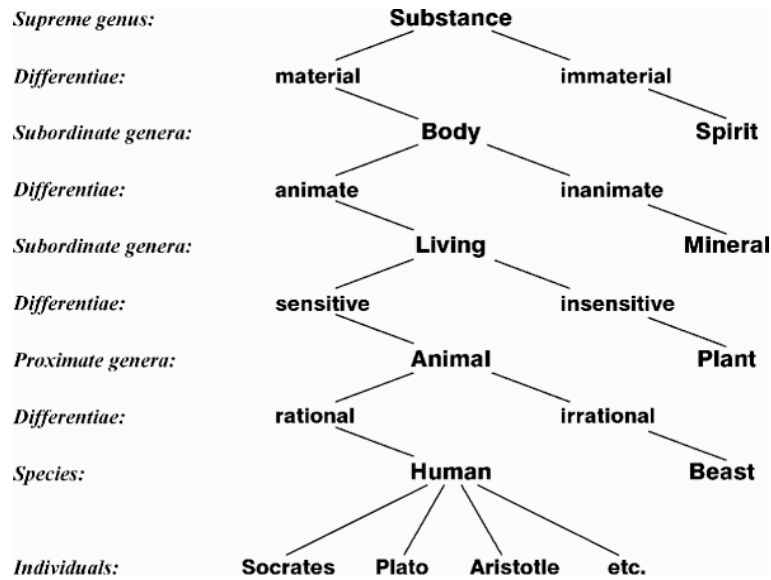


Figure 5 Tree of Porphyry (Sowa, 2000)

Figure 6 shows this taxonomy of substance represented as a semantic network. The disjointness of concepts at each level of the taxonomy, body and spirit for example, is not represented directly but through the disjointness of their *differentiae*, material and immaterial.

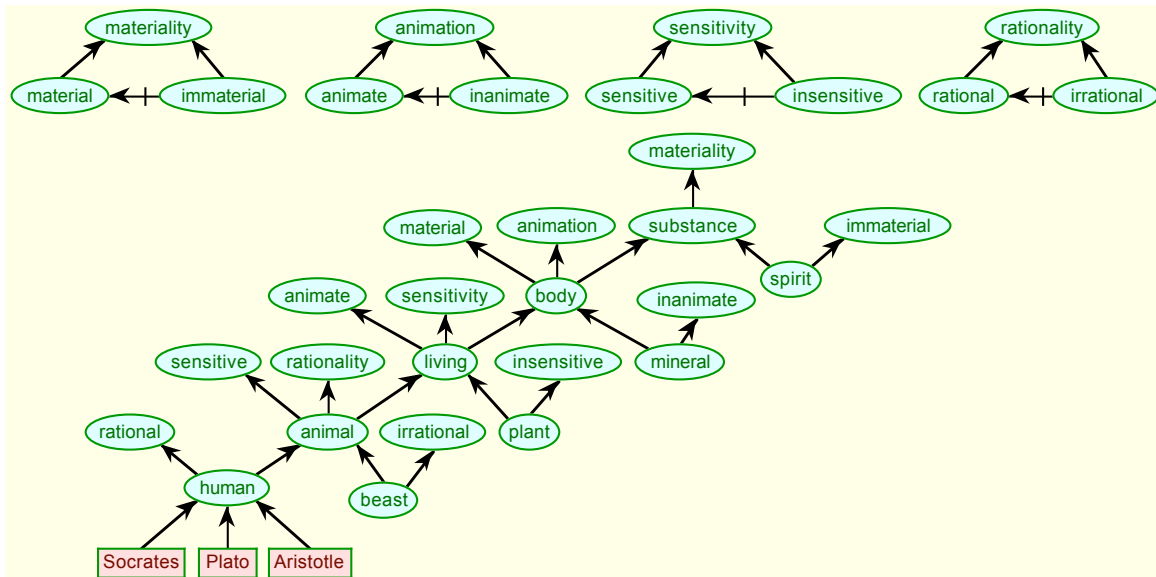


Figure 6 Tree of Porphyry as a semantic network

The description logic translator script outputs:-

material ⊆ materiality	immaterial ⊆ materiality ⊓ ¬material
animate ⊆ animation	inanimate ⊆ animation ⊓ ¬animate
sensitive ⊆ sensitivity	insensitive ⊆ sensitivity ⊓ ¬sensitive
rational ⊆ rationality	irrational ⊆ rationality ⊓ ¬rational
substance ⊆ materiality	
spirit ⊆ substance ⊓ immaterial	body ⊆ substance ⊓ material ⊓ animation

living \sqsubseteq body \sqcap animate \sqcap sensitivity	mineral \sqsubseteq body \sqcap inanimate
animal \sqsubseteq living \sqcap sensitive \sqcap rationality	plant \sqsubseteq living \sqcap insensitive
beast \sqsubseteq animal \sqcap irrational	human \sqsubseteq animal \sqcap rational
human(Socrates) human(Plato)	human(Aristotle)

I have shown the *differentiae* at the top not just as oppositions but as determinables under a determinate. This may seem anachronistic but, as Salmieri (2007) has noted, Aristotle had already discussed the notions underlying Johnson’s determinable/determinate construction in his *Categories*, and it seems intrinsic to the semantics of contraries. The disjunction of opposite concepts is itself a concept that may be explicit in the lexicon or covertly implied (Berlin, et al., 1968). This has the consequence of adding the determinables as superordinate concepts in the taxonomy where they may be seen as *dispositions* (Mumford, 1998), for example substance has the potential to be material. This knowledge structure may be historically and philosophically debatable, and the determinable dispositions can be removed without affecting the logical structure of the taxonomy, but it illustrates the kind of issues that the formal logical expression of a knowledge structure can raise. Even more interesting is the possibility of the four dispositional concepts being at other levels in the hierarchy, for example Figure 7 shows them as all being dispositions of substance.




Figure 7 Possible dispositions of substance

Again this does not change the logical structure *qua* taxonomy, but it highlights many metaphysical issues that are debated in the literature, and similar issues arise in modern developments of taxonomies for knowledge-based systems.

Figure 7 also illustrates another important knowledge structure, that of a *frame* (Minsky, 1974), having *slots* for the four determinables to be filled by *values* from their corresponding determinates. Since each of the determinables could also be expanded to have underlying taxonomic structures, not just a simple opposition, it also illustrates another significant knowledge structure, that of *faceted taxonomies* which are important in library science (Broughton, 2006). They are common in all cognitive systems as structured frames (Barsalou, 1992). Kelly (1955) uses the term *constructs* for psychological constructions having the form of the four at the top of Figure 6, and models those at the bottom as a hierarchical *construct system* based on these constructs.

3.3 Cluster concepts

Cluster concepts were introduced as a way of modeling Wittgenstein’s *family resemblances* which cannot be defined in terms of necessary and sufficient conditions (Cooper, 1972). The notion is that the presence of one of many subsets of a range of possible properties is sufficient to define a concept, although no particular property is necessary. They have been used in the philosophy of art to provide a knowledge structure for the concept of an *art object* (Gaut, 2000), which is a stereotypical example of what Weitz (1977) terms an *open concept*, subject to change and supposedly lacking formal definition.

A model of a cluster concept is a frame of determinables some or all of which have preferred determinates that are normally, or usually, expected in exemplars of the concept. This could be turned into a formal definition in terms of the determinables themselves being the defining properties, as will be discussed in Section 3.5. However, there is a more psychologically realistic model in terms of abductive processes where the assertion that an individual has certain properties leads to the deductive inference that it comes under the determinable concepts for those properties which in turn leads to the abductive inference that it comes under a frame concept characterized by those determinables. This inference can be validated by checking whether it comes under the remaining determinables.

For example, Figure 8 presents Gaut's (2000) account of art as a cluster concept in semantic network form. The concept *art object* is represented as a frame based on eleven determinables providing dimensions along which to assess an art work. The concept *stereotypical art object* is represented as a concept implying all the positive determinates of the determinables. Note that the determinables can be seen as dispositions, that if something is construed as an art object then it is also construed as having the potential to be aesthetically positive or negative, to evoke an emotional response, and so on.

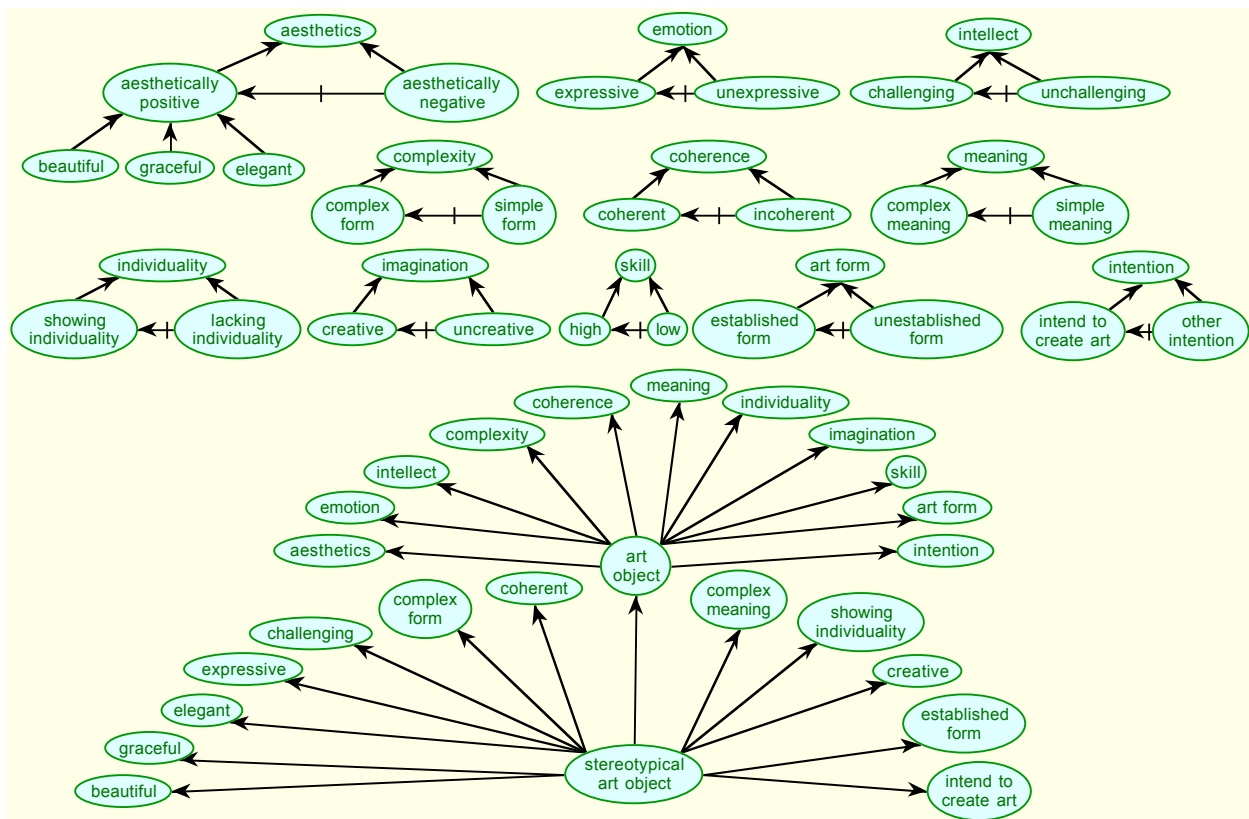


Figure 8 Art as a cluster concept

In terms of an entity being construed as an art object, the abductive reasoning path might be that someone described it as 'creative, beautiful and complex in meaning'. These would be seen as determinates of the determinables imagination, aesthetics and meaning, which form part of the frame for an art object, suggesting abductively that it is being construed as an art object, and raising questions about its status on the other dimensions. If those questions are regarded as

reasonable then the abductive inference is confirmed. but, even if not, the mistake will be seen to be a reasonable one, not a category error. For some philosophers or historians of art, the art object frame may require some specific links at a lower level, for example, that it is essential that the intention be to create art. The network provides a framework in which to clarify and discuss such issues.

The stereotypical art object is of interest because, while many art objects will conform to the stereotype, there will also be many exceptions that are of particular interest just because they exemplify the boundaries of the norm. People will use the stereotype to model that something is an art object but not, say, very challenging or aesthetically pleasing. “But not” can be modeled as the dynamic generation of a concept that has a common superordinate with *stereotypical art object* and is subordinate to *art object*. One will usually not wish to populate a knowledge base with all such possible concepts, although some of them may become so commonly used as to be given specific names. This potential explosion of “but not” concepts is one reason why the original cluster concept notion has been criticized (Boër, 1974), but the dynamic creation of transient concepts may be an appropriate model for cognitive systems.

3.4 Graded Concepts

One important aspect of determinates under a determinable is that they are not only disjoint but also that “determinates under the same determinable admit of comparison in a way unavailable to pairs of properties with no determinable in common” (Funkhouser, 2006). For example in Figures 1 and 2 gray may be seen as more similar to black and white than they are to one another. It is ‘between’ them and this is shown in the diagram. It can also be modeled in the logic. If we take each determinate to be characterized by the set of salient predicates implied by it, and black, achromatic and white to be those predicates in this case then the sets generated are:

$$\begin{aligned} \text{white} &\sqsubseteq \text{achromatic} \sqcap \text{white} \sqcap \neg \text{black} \\ \text{gray} &\sqsubseteq \text{achromatic} \sqcap \neg \text{white} \sqcap \neg \text{black} \\ \text{black} &\sqsubseteq \text{achromatic} \sqcap \neg \text{white} \sqcap \text{black} \end{aligned}$$

The cardinality of the symmetric difference between these sets is a distance measure, δ , such that: $\delta(\text{white}, \text{gray}) = 2$, $\delta(\text{gray}, \text{black}) = 2$, $\delta(\text{white}, \text{black}) = 4$, which places them along a graded scale from white through gray to black. This construction models graded or fuzzy concepts (Lu, et al., 2008) through the existing logical structure without introducing additional concept types.

Note that the three concepts white, gray and black are identical in their relative structures, and that it is the choice of salient concepts that introduces the differentiation in the metric which places gray between white and black and generates a three-point scale. There is a psychological or ontological choice involved in the selection of the salient concepts that determine the nature of the scale.

This three-point model can be extended to more gradations by increasing the levels in the taxonomy. For example, Figure 9 shows an additional level introducing five determinates at the lowest level, white, light gray, mid gray, dark gray and black, in terms of four salient concepts, white, very light gray, achromatic, very dark gray, and black. A computation similar to that above places the five lowest determinates along a five point linear scale. It also places very light gray and very dark gray on the same scale, between white and light gray, and dark gray and black respectively, indicating the potential use of a seven point linear scale.

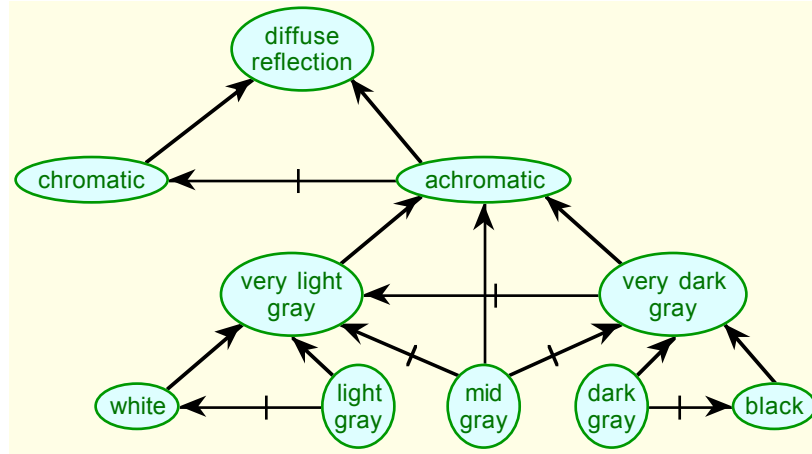


Figure 9 A graded concept with five or seven points on a linear scale

One can also add two additional concepts between light gray and mid gray, and mid gray and dark gray, respectively as shown in Figure 10, which shows how the scales are derived. Such constructions seem to underlie the gradability of the majority of concepts (Rosch and Lloyd, 1978).




Figure 10 A graded concept with up to nine points on a linear scale

The topology of a linear scale is just one example of a graded concept. Other topologies can be generated by alternative structures, for example a circular scale appropriate to hues.

3.5 Definitions, rules, exceptions

The issue of definitions has been left to a late stage in this paper because, while they are fundamental to mathematical logic, and description logics were developed to provide a mechanism for the logical definition of terminology, their role in human knowledge structures is much debated (Fodor, et al., 1980; Pitt, 1999). None of the examples given so far has been definitional, but, nevertheless, they show that many significant knowledge structures can be modeled, and that abductive processes, supported and validated by deductive processes, may account for many apparently definitional constructions.

However, any visual language for description logics must also be able to represent definitions, at the very least, conjunctive definitions, for example of the form: $Woman \equiv Female \sqcap Person$, that it is not only necessary for a woman to be a female person, but that is also sufficient to define a woman as a female person. In earlier versions of the visual language we made a distinction between *base*, or *primitive*, concepts and defined concepts by using different surrounds (Gaines, 1991b). This seemed appropriate because CLASSIC made this distinction. However, as description logics came to encompass disjunction and negation, and it became possible to represent rules directly rather than as an additional specialist inference process, it also began to make sense to distinguish the necessary and sufficient links from the ones that were necessary only. That is to use two forms of arrow of implication, one of which was definitional and the other only implicational, and to distinguish a primitive concept as being one with no definitional links.

This is illustrated in Figure 11 where the structure on the left is implicational generating $Woman \sqsubseteq Person \sqcap Female$, that in the center definitional generating $Woman \equiv Female \sqcap Person$, and on the right rule-like generating $Pensionable_woman_rule \equiv Woman \sqcap Over_pensionable_age_for_women$ and $Pensionable_woman_rule \sqsubseteq entitled_to_pension$. The heavier arrow with the head to some extent pointing back is explained to users as providing part of a concept definition together with any other such arrows from the concept.




Figure 11 Implicational (left), definitional (center), and rule-like knowledge structures

Negative definitional arrows enable complements and negated concepts to be defined. Figure 12 shows man defined as the complement of woman as defined in the center of Figure 11, translating to $Man \equiv Person \sqcap \neg Woman$ and $Not_Man \equiv \neg Man$. From these definitions the system can infer that a man is not female and that a person must be either a woman or a man. Negations concepts also be defined, such as “not man” shown, but these, as noted by Aristotle (Horn, 1989), convey little information as concepts *per se*, and their use in an ontology is usually an error in not defining a more restricted complement through relative negation.




Figure 12 Complement and negation

Definitional arrows could be used to replace the implicational arrows for “art object” and “stereotypical art object” in Figure 8 making the concepts involved jointly sufficient as well as necessary. There would be no change in the psychological and logical processes described in the use of these concepts, that is, given the number of defining concepts involved, an abductive inference would be made that the defined concept might reasonably be the explanation why some of the dimensions were being used, and then the applicability of the others would be checked. This helps to clarify the sense in which cluster concepts can be regarded as definitional (Boër, 1974), and provide further interpretation of psychological studies of concepts and their necessary and sufficient features (Pothos and Hahn, 2000).

People often use rules with *exceptions*, that a general rule applies except in particular circumstances, and representing knowledge in this way can make the structure more comprehensible (Gaines, 1996), and also supports *default reasoning* where a general rule is applied when it is not known if the exception applies provided this does not lead to a contradiction (Besnard, 1989).

Figure 13 shows a simple expert system based on rules with exceptions derived from data in the UCI machine learning repository (Asuncion and Newman, 2007).




Figure 13 Contact lens prescription problem as semantic network of rules with exceptions

The problem is one of prescribing hard or soft contact lenses for a client having certain attributes (Cendrowska, 1987), and the four determinables under “person” at the upper left are those whose determinates are necessary to characterize the client and allow the appropriate determinant of the determinable, “prescription,” at the upper right to be inferred. The “prescription context” frame at the middle right activates the determinables, making their determinates relevant to the problem. The definitions/rules at the bottom solve the problem, and this particular rule set is activated by the “prescription rules” concept at the bottom. Setting this concept true of a client brings these rules into play and also sets the relevant determinables true raising the question as to which of their determinates applies. Providing truth values for the determinates appropriate to the client and running inference enables the prescription to be inferred.

The description logic translator script outputs (indentations added for clarity):-

```

astigmatism ⊆ person
  astigmatic ⊆ astigmatism          not_astigmatic ⊆ astigmatism ⊓ ¬astigmatic
tear_production ⊆ person
  normal ⊆ tear_production          reduced ⊆ tear_production ⊓ ¬normal
presbyopia ⊆ person
  young ⊆ presbyopia                old ⊆ presbyopia ⊓ ¬young
  presbyopic ⊆ old                  pre-presbyopic ⊆ old ⊓ ¬presbyopic
myopia ⊆ person
  myope ⊆ myopia                    hypermetrope ⊆ myopia ⊓ ¬myope
prescription ⊆ person
  lens ⊆ prescription               none ⊆ prescription ⊓ ¬lens
  hard ⊆ lens                       soft ⊆ lens ⊓ ¬hard
prescription_context ⊆ astigmatism ⊓ tear_production ⊓ prescription ⊓ myopia ⊓ presbyopia
prescription_rules ⊆ prescription_context
prescribe_soft ⊆ soft
  prescribe_soft ≡ not_astigmatic ⊓ normal ⊓ prescription_rules ⊓ ¬exception_soft
  exception_soft ≡ myope ⊓ presbyopic
prescribe_hard ⊆ hard
  prescribe_hard ≡ astigmatic ⊓ normal ⊓ prescription_rules ⊓ ¬exception_hard
  exception_hard ≡ old ⊓ hypermetrope
prescribe_none ⊆ none
  prescribe_none ≡ prescription_rules ⊓ ¬prescribe_hard ⊓ ¬prescribe_soft

```

The default rules are simple, that a patient whose tear production is normal will be prescribed a hard lens if astigmatic and a soft lens if not. However, there is an exception to the soft prescription if the patient is presbyopic and myopic, and to the hard if hypermetropic and old. This example is a nice one for tutorial purposes since the dataset of 24 cases is small enough to be understood, and yet was problematic for early machine learning algorithms using greedy uncertainty reduction heuristics (Cendrowska, 1987). Students can identify the default rules fairly readily, test them on the dataset, and then add the appropriate exceptions. They can examine default reasoning by setting the exception concepts false and finding that they solve 21 cases correctly. They can also see possible worlds search in action by setting “prescription rules” and “lens” true and running inference to determine that the tear production must be normal in all possible cases where a lens is prescribed as shown in Figure 14.

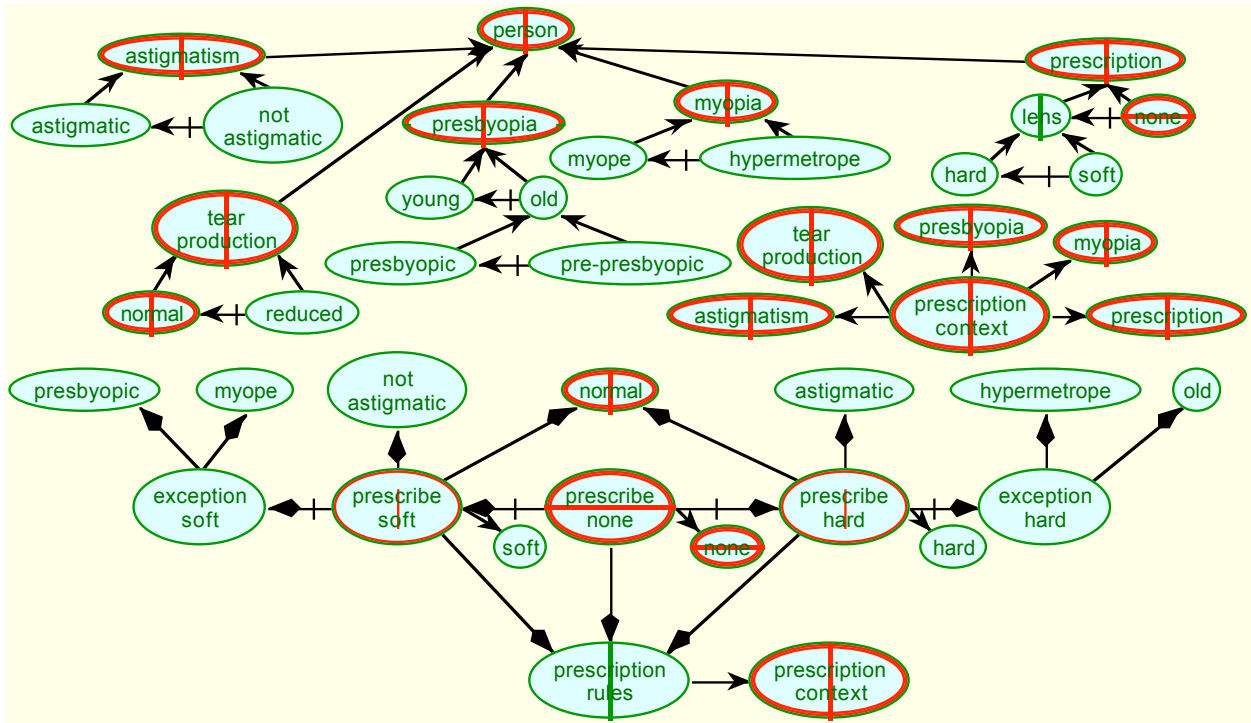


Figure 14 Results of inference for lens problem shown on network

The explanation generated for these inferences is:-

Explore from: lens \sqcap prescription_rules

lens \sqsubseteq prescription [lens]

prescription_rules \sqsubseteq prescription_context [prescription_rules]

lens \sqsubseteq \neg none [?none]

prescription \sqsubseteq person [prescription]

prescription_context \sqsubseteq astigmatism \sqcap myopia \sqcap presbyopia \sqcap tear_production [prescription_context]

\neg none \sqsubseteq \neg prescribe_none [?prescribe_none]

Explore Level 1.1: \vdash prescribe_hard \sqsubseteq \neg prescribe_none

prescribe_hard \sqsubseteq astigmatic \sqcap hard \sqcap normal \sqcap \neg exception_hard [prescribe_hard]

astigmatic \sqsubseteq \neg not_astigmatic [?not_astigmatic]

hard \sqsubseteq \neg soft [?soft]

\neg not_astigmatic \sqsubseteq \neg prescribe_soft [?prescribe_soft]

Explore Level 2.1: \vdash \neg old \sqsubseteq \neg exception_hard

\neg old \sqsubseteq \neg presbyopic [?presbyopic]

\neg presbyopic \sqsubseteq \neg exception_soft [?exception_soft]

Possible World 1: astigmatic \sqcap hard \sqcap normal \sqcap prescribe_hard \sqcap \neg exception_hard \sqcap

\neg exception_soft \sqcap \neg not_astigmatic \sqcap \neg old \sqcap \neg presbyopic \sqcap \neg prescribe_soft \sqcap \neg soft

Explore Level 2.2: \vdash hypermetrope \sqsubseteq \neg exception_hard

Possible World 2: astigmatic \sqcap hard \sqcap normal \sqcap prescribe_hard \sqcap \neg exception_hard \sqcap

\neg hypermetrope \sqcap \neg not_astigmatic \sqcap \neg prescribe_soft \sqcap \neg soft

Explore Level 1.2: \vdash prescribe_soft \sqsubseteq \neg prescribe_none

prescribe_soft \sqsubseteq normal \sqcap not_astigmatic \sqcap soft \sqcap \neg exception_soft [prescribe_soft]

not_astigmatic \sqsubseteq \neg astigmatic [not_astigmatic]

soft \sqsubseteq \neg hard [soft]

\neg astigmatic \sqsubseteq \neg prescribe_hard [?prescribe_hard]

Explore Level 2.1: $\vdash \neg \text{myope} \sqsubseteq \neg \text{exception_soft}$
 Possible World 3: $\text{normal} \sqcap \text{not_astigmatic} \sqcap \text{prescribe_soft} \sqcap \text{soft} \sqcap \neg \text{astigmatic} \sqcap \neg \text{exception_soft}$
 $\sqcap \neg \text{hard} \sqcap \neg \text{myope} \sqcap \neg \text{prescribe_hard}$

Explore Level 2.2: $\vdash \neg \text{presbyopic} \sqsubseteq \neg \text{exception_soft}$
 Possible World 4: $\text{normal} \sqcap \text{not_astigmatic} \sqcap \text{prescribe_soft} \sqcap \text{soft} \sqcap \neg \text{astigmatic} \sqcap \neg \text{exception_soft}$
 $\sqcap \neg \text{hard} \sqcap \neg \text{presbyopic} \sqcap \neg \text{prescribe_hard}$

Exploration Inferences: normal
 $\neg \text{prescribe_none} \sqsubseteq \text{prescribe_hard} \vee \text{prescribe_soft} [\neg \text{prescribe_none}]$

The system generates the first five inferences listed by truth value propagation and, when it can make no further inferences in this way, starts making hypotheses to resolve unresolved disjunctions and hence explore possible worlds. Firstly, asserting “prescribe hard” is true in order to infer “prescribe none” is false, and then propagating truth values until it needs a second hypothesis, that “old” is false in order to infer that “exception hard” is false. This is sufficient to infer a complete set of truth values for “Possible World 1.”

It then backtracks to its second hypothesis and hypothesizes that “hypermetrope” is false in order to infer that “exception hard” is false, which is sufficient to generate “Possible World 2.” It backtracks again to its first hypothesis, hypothesizing instead that “prescribe soft” is true in order to infer “prescribe none” is false, and through these processes generates a set of four possible worlds that cover all relevant situations. It then outputs as “Exploration Inferences” any concept whose truth value has not been inferred by truth value propagation but has been inferred to be the same in all possible worlds, in this case that “normal” is true. Finally it outputs any unresolved disjunctions, in this case that “prescribe none” being false implies either “prescribe hard” or “prescribe soft” is true.

The following sections introduce the remaining graphic primitives needed to support description logics.

3.6 Individuals

Individuals are represented by a rectangular surround as already shown at the bottom left of Figure 6. An arrow from an individual to a concept indicates that the concept applies to the individual, that is the concept as a predicate is true of the individual. In the lens example above one could enter a case graphically as an individual with its properties as shown in Figure 15, for the example of a client known to have been prescribed contact lenses as in Figure 14 on the left and a client for whom a prescription is required on the right. These translate to $(\text{lens} \sqcap \text{prescription_rules})(\text{client_known_to_have_been_prescribed_lenses})$ and $(\text{prescription_rules} \sqcap \text{astigmatic} \sqcap \text{normal} \sqcap \text{presbyopic} \sqcap \text{myope})(\text{client_needing_a_contact_lens_prescription})$, respectively.




Figure 15 Properties being asserted of individuals

The semantics of individuals is that their labels are unique identifiers of an existing individual. As Zalta (1988) nicely expresses it: individuals *exemplify* properties whereas concepts *encode*

them. The properties exemplified can be factored into those that are essential, being rigid designators across all possible world, and those which are state descriptions and can vary across possible worlds. We use the definitional arrow for the essential properties and the implicational one for the state properties as shown in Figure 16.




Figure 16 Distinguishing essence and state

The semantics of essential properties is that they, together with its unique identifier, define the individual *qua* concept, that is the *concept* of the individual includes its essence as well as its identification and is satisfied by the individual, and only that individual, in any state. This model serves to distinguish ‘kinds’ or ‘sorts’ as rigid designators in an ontology, helping avoid some of the issues that Guarino and Welty (2004) have documented in the original WordNet and CYC upper-level ontologies. It is also useful for student exercises on the metaphysical nature of individuals, through inference that the knowledge structure is incoherent when the state violates the essence, indicating that the existence of the individual should no longer be asserted, the essence is inappropriate or the state erroneous.

3.7 Binary relations, roles

One restriction that description logics place on subsets of FOL is that they support only binary predicates, termed *roles*, originally treating these in impoverished fashion as having no inter-role inferential constraints, and now usually supporting role inclusion specification but not role chain inclusions or equalities since these are known to lead to undecidability (Donini, 2003). Roles have been classically shown in semantic networks as unboxed text and we have adopted this convention as illustrated in Figure 17. On the left “Mother without daughter” is defined as a mother none of whose children are female, translating to $\text{Mother_without_daughter} \equiv \text{Mother} \sqcap \forall \text{has_Child.}\neg \text{Female}$. On the right Mary is shown as having two children, Susan and Arthur, translating to $\text{has_Child}(\text{Mary}, \text{Susan}), \text{has_Child}(\text{Mary}, \text{Arthur})$, or, for export to Racer, $(\text{related Mary Susan has_Child}), (\text{related Mary Arthur has_Child})$.




Figure 17 Binary relations as roles

One design decision apparent in Figure 17 is that a role is treated as a node in its own right, not as a label on a link. This avoids the clutter of having multiple links to represent multiple role constraints or fillers, and makes the grouping of aspects of the domain of a role visually apparent.

Figure 18 illustrates how relationships between roles are represented.




Figure 18 Binary relations as roles

The inverse relationship between roles is symmetric and represented through an undecorated line between them, as shown for “has child” and “has parent”, translating to $\text{has_parent} \equiv \text{has_child}^-$ or (define-role `has_parent (inverse has_child)`). The inclusion relation is represented as for concepts, as shown for “acquaintance” as a binary relation including both “friend” and “co-worker” and translating to $\text{friend} \sqsubseteq \text{acquaintance}$, $\text{co-worker} \sqsubseteq \text{acquaintance}$. The definition of “uncle” as someone who is a male sibling of a parent shown on the right, which translates to $\text{uncle} \equiv \text{has_parent} \circ \text{has_brother}$ or (define-role `uncle (compose has_parent has_brother)`), is not currently supported by most decision logics because the role chain equality construction can lead to undecidability (Donini, 2003). This is a problem for the semantic web applications of description logics and has led to the *ad hoc* addition of the Semantic Web Rule Language (SWRL) to OWL, the Web Ontology Language (Horrocks, et al., 2005). SWRL enables one side of the definition of uncle to be expressed in DATALOG form as $\text{parent}(?a, ?b) \wedge \text{brother}(?b, ?c) \rightarrow \text{uncle}(?a, ?c)$. However, as already noted, rules modeled by additional necessary properties of a definition can already be modeled by description logics and it seems unfortunate that modeling those involving role chains requires the addition of a completely different framework.

The problems with role constructs in semantic networks and description logics may be seen as resulting from the separation of unary predicates as *concepts* from binary predicates as *roles*, and we have experimented with visual language representations that treat n-ary predicates in a uniform fashion. However, this results in more complex graphical representations which seem less natural than those of Figure 18. At this stage, it is reasonable to conclude that the representation of some aspects of relations is problematic both for practical inference systems and visual languages, and that this is an area where further research is required.

3.8 Cardinality constraints, concrete domains

Many description logics support *cardinality constraints* on the number of fillers of a role, and *concrete domains* of sets of strings, integers, ranges of integers, sets of ranges of integers, and equivalent constructs for numbers in general, dates, and so on, ordered by set inclusion. They can be modeled for purposes of representation, inference and implementation as operations upon arbitrary, ‘plugged-in’ lattices of subsets of differing types (Gaines, 1993), and general complexity results have been derived for inference with different classes of operation in such domains (Baader and Sattler, 2003).

As illustrated in Figure 19, a rectangle with doubled vertical sides is used for cardinality constraints and concrete domain specifications, and explained to users as representing set-theoretic constraints. The symbol, “ \exists ” is supported as an alias for “ ≥ 1 ”, and values in a concrete domain are represented by the same rectangular surround that is used for individuals.




Figure 19 Cardinality constraints and concrete domains

The description logic translator script outputs:-

Parent $\equiv \exists \text{has_Child}$ Person $\sqsubseteq [\text{Integer } 0 \leq 150] \text{ age}$
 Person(Fred) age(Fred, 45)

Further examples of cardinality constraints in use in ontologies are given in Figures 20 and 21.

Figure 20 is a semantic network representing an example a SNePS case study (Shapiro, 1991), a puzzle where 4 people each hold two of eight jobs, no job is held by more than one person, and various constraints are given on who can hold what. It shows the use of cardinality constraints and an inverse role relationship to capture the problem, together with other constraints such as no person is both chef and police officer, no male is a chef, Pete is not a nurse or police officer, and so on.

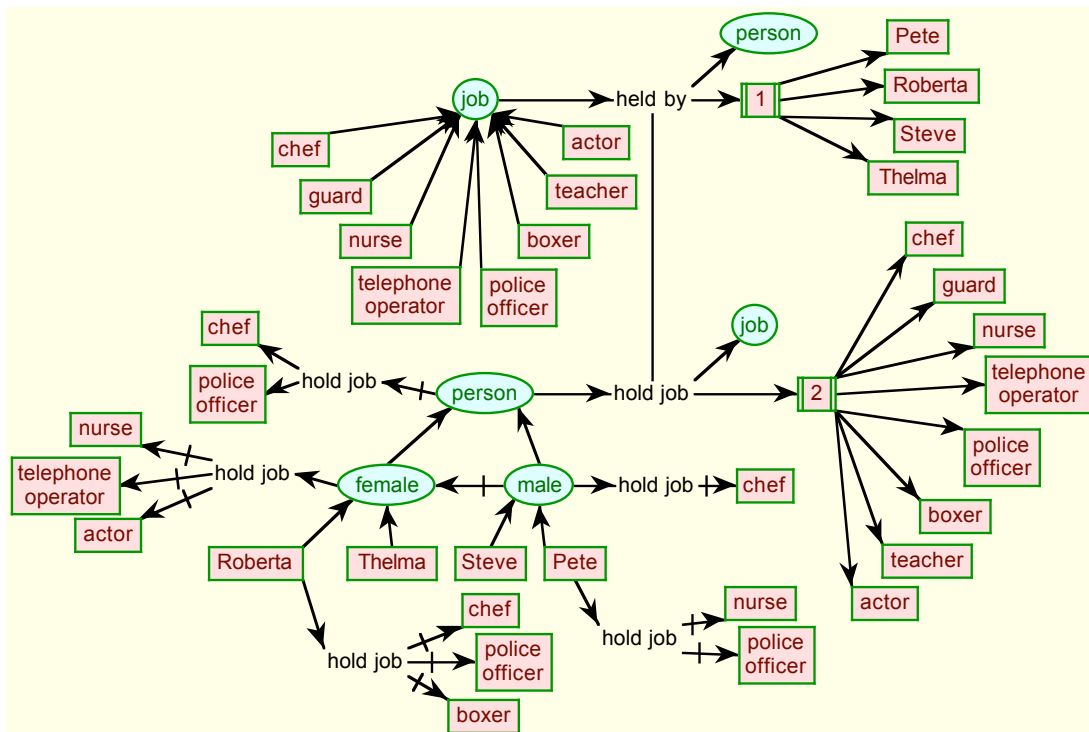


Figure 20 Semantic network representing a puzzle

The description logic translator script outputs:-

```

job ⊆ ∀held_by.person ⊆=1 held_by.(Roberta ⊔ Steve ⊔ Thelma ⊔ Pete)
person ⊆ ∀hold_job.job ⊆=2 hold_job.(guard ⊔ nurse ⊔ telephone_operator ⊔ police_officer ⊔ boxer ⊔
  teacher ⊔ actor ⊔ chef) ⊆ ¬(∃hold_job.chef ⊆ ∃hold_job.police_officer)
male ⊆ ⊆ person ⊆ ¬female ⊆ ¬∃hold_job.chef
female ⊆ person ⊆ ¬∃hold_job.actor ⊆ ¬∃hold_job.telephone_operator ⊆ ¬∃hold_job.nurse
hold_job ≡ held_by
female(Thelma)    male(Steve)    (male ⊆ ¬∃hold_job.police_officer ⊆ ¬∃hold_job.nurse)(Pete)
(female ⊆ ¬∃hold_job.boxer ⊆ ¬∃hold_job.police_officer ⊆ ¬∃hold_job.chef)(Roberta)
job(chef)         job(guard)      job(nurse)       job(telephone_operator)
job(police_officer) job(boxer)     job(teacher)     job(actor)
  
```

KNet’s DPLL inference system readily solves the problem to infer that Roberta is a guard and teacher, Thelma is a chef and boxer, Pete is a telephone operator and actor, and Steve is a nurse and police officer, and the trace of its solution can be compared by students with their own.

3.9 Disjunction

Disjunctively defined concepts are represented by a concept oval with a doubled boundary, and inference with disjunctive definitions is the dual of that with conjunctive definitions. Disjunctively defined concepts may be seen as propagating falsity in the same way as conjunctive ones propagate truth. However, they present interesting issues because they are very problematic for human reasoning (Johnson-Laird and Byrne, 1991), and, while we provide them, we discourage their use in ontologies because it is often possible, and clearer, to define them non-disjunctively. For example Figure 21 shows a family ontology used as an example in the *Description Logic Handbook* (Donini, 2003). “Parent” is defined disjunctively as “Mother” or “Father”, and, as shown below, the translator outputs this as $\text{Parent} \equiv \text{Father} \sqcup \text{Mother}$, or (equivalent $\text{Parent}(\text{or } \text{Father } \text{Mother})$). However, the logically equivalent non-disjunctive definition shown at the top left of Figure 19, which translates to $\text{Parent} \equiv \exists \text{has_Child}$, is better since it captures the essence of the notion of parent without introducing the complication of inferring this from the disjunction of mother and father.




Figure 21 Family ontology

The description logic translator script outputs:-

```

Woman ≡ Female ⊓ Person
Man ≡ Person ⊓ ¬Woman
Wife ≡ Woman ⊓ ∃has_Husband.Man
Mother ≡ Woman ⊓ ∃has_Child.Person
Father ≡ Man ⊓ ∃has_Child.Person
Grandmother ≡ Mother ⊓ ∃has_Child.Parent
Mother_Without_Daughter ≡ Mother ⊓ ∀has_Child.¬Woman
Mother_With_Many_Children ≡ Mother ⊓ ≥3 has_Child
Parent ≡ Father ⊔ Mother

```


3.10 Modularity, management of ontology integration

It will already be apparent that the visual language constructs shown support *modularity* in much the same way as a linear notation, in that a concept can be defined in one place and used in another, as exemplified, for example, in Figures 6 and 8. The visual language supports the definition of necessary and/or sufficient conditions in one place and their use in others, and also the partial definition of such conditions in several places and the management of the consequences of merging the partial definitions.

The KNet visual language support system allows knowledge structures defined in one file to be combined with those defined in others, and the combinations to be changed to experiment with alternative definitions or to model conceptual change. A visual symbol is provided for links to external knowledge structures. Figure 22 shows a variant of Figure 13 where the determinable structures at the top of the figure are kept in an external knowledge structure, and a link is provided to the knowledge base containing it. This is implemented as a URL to the knowledge base kept in the hidden *annotation field* of the node designed to hold URL's and associated data in hypermedia applications (Gaines *et al.*, 1995; 2007).




Figure 22 Link to definitions in another knowledge base

Thus, any concept used in one knowledge structure may be defined or further refined in another. This enables ontologies to be entered with concepts that are initially treated as primitives but later have necessary and/or sufficient conditions defined for them. For example, the concept “has body” on the left of Figure 23 may later be given more explicit meaning through a role constraint as shown in the center, and the concept “body” may itself be attributed properties elsewhere as shown on the right.




Figure 23 Use of a concept in one knowledge structure and refinement in others

Abbreviation of concept labels may well make the original label misleading, for example “body” instead of “has body”, and lead to the kind of problems in the ontology that Guarino and Welty (2004) have documented in their development of *OntoClean* or some variant of what Bennett and Hacker (2003) have termed the “mereological fallacy” in the neurosciences of attributing to a part of a system that which can only meaningfully be attributed to the whole.

A good example of modular knowledge structures are the *differentiae* associated with the different species in a biological taxonomy. It is customary to show the bare taxonomy diagrammatically as in Figure 1, and to document the *differentiae* as cluster concepts in

voluminous associated material. This documentation describes the complex knowledge structures supporting the recognition of members of each species throughout their life cycles, and paying particular attention to cases that are difficult to distinguish.

Modularity does not apply only to the conceptual relationships. The assertion of the properties of individuals is a modular aspect of the knowledge structure, subject to change when states change or inconsistencies are found. The choice of what determinate to assign to an individual is a human decision subject to constraint by its consequences, but otherwise open as reflected in Kelly's (1970) *choice corollary* that "A person chooses for himself that alternative in a dichotomized construct through which he anticipates the greater possibility for the elaboration of his system." When people are faced with an inconsistency they often prefer to change the previously asserted 'facts' rather than the ontology being used (Revlis and Hayes, 1972).

Adding an additional knowledge structure may be *conservative* in not changing the consequences of the knowledge structures to which it is added, but only defining new ones. Or it may result in changes to what was already inferable, including those leading to inconsistency. Such changes may be precisely what is of interest when one is examining conceptual change or the impact on innovation of the inconsistencies resulting from combining ontologies. A good example is the instability inferred for the orbiting electron model of the atom when Newton's ontology for motion and Maxwell's for electrodynamics were combined, which Bohr resolved by introducing quantized energy levels (Heilbron and Kuhn, 1969). In other applications it may be a design requirement that the merging of ontologies is conservative and constraints must be applied to achieve this (Grau *et al.*, 2007).

Ontologies are intrinsically incremental, distributed, subject to variation, and to the availability of many alternative structures dependent on factors such as the context of use, for example, social and legal notions of "family," and individual and cultural knowledge, preferences and milieu. The visual language support system must be able to support the management, and tracking, of such variations. Some consequences of integration may be made immediately apparent by an analysis of whether syntactically new definitions of already defined concepts are introduced. However, these need not necessarily involve semantic change if, for example, they repeat existing definitions, and there may also be semantic consequences beyond those the syntactic analysis indicates. Managing this requires close integration of the visual language system and the inference engine services and is a topic requiring substantial further research.

In the context of this paper, the grouping of related knowledge structures in the visual language seems to support human comprehension of the possibilities, and the consequences, of bringing together various knowledge modules.

4 Conclusions

This is a practice and experience paper attempting to highlight issues that will need to be addressed by any designer of a semantic network-like visual language for some subset of FOL. The general design issues have been discussed in Section 2 in the context of semantic networks but stated as generally as possible to make them relevant to other visual languages for logic.

4.1 Addressing the Design Issues

Section 3 presents a detailed example of the design of a visual language for description logics illustrating how the design issues may be addressed. Issue 1, of unambiguous translation from the visual language to standard logical formalism, is addressed by algorithms that translate any diagram in the visual language to statements in one of the standard description logic formalisms. Issue 2, of taking into account current practice in user communities, is addressed by the use of existing conventions for semantic networks wherever possible. Issue 3, of each construction in the visual language having a well-defined and simply explained meaning, has been addressed through examples of how the constructions described are explained to users. Issue 4, of taking into account the different roles the visual language will play, has been addressed by designing a single language that can be used for knowledge entry, knowledge base representation and, through visual annotation generated in the inference process, for explanation of inference.

Issue 5, testing the design through a wide range of test cases, including both real problems and generic knowledge structures, has been addressed through the diversity of examples given to illustrate various visual language constructs in Section 3. Issue 6, encouragement of the use of representational capabilities that minimize the complexity of inference has been addressed in terms of the provision of relative negation, depreciation of disjunction, and so on. Issue 7, the need to take into account the use of the visual language not only in deductive reasoning but also in various ampliative reasoning patterns has been addressed by examples where, for example, abductive reasoning based on partial definitions provides a better model of human practice than would deductive reasoning based on full definitions. Issue 8, the need for modularity, has been addressed by a design which allows diagrams to be split into components and recombined as needed.

These issues are relevant to the design of other forms of visual languages for logics, such as those based on Venn and Euler diagrams, with issues 1 through 4 being normally addressed but consideration of issues 5 through 8 perhaps providing some interesting research issues. The examples of visual language constructs and their applications given in Section 3 also provide useful test cases for other visual languages. For example, the representation of disjointness and the differentiation of sufficient and necessary properties in definition through different arrow types provide interesting approaches to the representation of important constructs in knowledge representation, and also raises the question of how these would be represented in other visual languages.

4.2 Future Directions

The exact expression of human knowledge, its communication, understanding and use in inference are difficult tasks at the apex of human achievement. Linguistic, visual and logical formulations of knowledge are complementary rather than competitive. All three require a high degree of skill to be fluent in their application. Natural language is not particularly ‘natural’ when it is used technically and precisely, and translations between equivalent expressions in linguistic, visual and logical forms can provide insights in relation to knowledge structures and argument forms beyond those of any single mode of expression.

I believe that the most significant issue for visual languages for logic is what contribution they make to the human understanding of knowledge structures. In this they are not so much to be compared with other logical formalisms as with natural language representations which are the

common way in which people think and communicate about knowledge. The capability of the visual language to translate precisely to a logical formalism which itself can be used for automatic inference needs to be evaluated for its possible contribution to conventional linguistic discussion about knowledge.

This paper has drawn upon a range of philosophical literatures for its examples of fundamental knowledge structures and shown that the visual representation and its precise logical consequences may be helpful in discussing the issues raised in those literatures. This is reminiscent of Leibniz's notion that the arithmetization of syllogistic logic would enable philosophical disputes to be resolved in a friendly fashion by computation (Sotirov, 1999). Perhaps it is an idea whose time, with the advances in modern theorem provers, has come, and visual languages may provide the bridge between the verbal expression of ideas and their logical explication. Perhaps not, but it is still an interesting and worthy objective.

In many disciplines such as electronic, mechanical and chemical engineering, diagrams have long played a central role in design, not just as aids but as 'truth bearers' (Perini, 2005). One can see analogies between electronic circuit or molecular structure diagrams and semantic networks. The generic knowledge structures discussed in this paper are similar to electronic modules built of basic components, or molecules built of atoms. In these other disciplines, once basic structures have been designed, fabricated and tested, they are used as modules to build larger sub-systems and eventually particular systems. Computer-aided design systems for engineering application operate with sub-systems as units and can show the structure of the underlying modules, and so on, down to the components. Alternative modules and sub-systems can be developed to achieve the same functionality or improve existing functionality.

This is an attractive model for the design of knowledge structures, whether they represent philosophical arguments, models of human reasoning, or improved modes of accessing knowledge at large such as the semantic web. Research on knowledge entry through the graphical assembly of components (Clark *et al.*, 2001) and design patterns for semantic web content (Gangemi, 2005) are already moving in this direction. We need further research on visual languages that can match the capabilities of the knowledge representation and inference systems we now have, and keep pace with and support their ongoing evolution.

We also need to facilitate the adoption and use of formal visual languages and their associated inference capabilities in scholarly communities that may have little knowledge, or sympathy with, formal logic, but might find the insights it can provide of value to their own endeavors. This is certainly a worthy and realistic objective, and one to which this paper is intended to contribute.

References

- Andersen, H. (2000). Kuhn's account of family resemblance: a solution to the problem of wide-open texture. *Erkenntnis* 52, 313-337.
- Andersen, H., Barker, P. and Chen, X. (2006). **The Cognitive Structure of Scientific Revolutions**. Cambridge University Press, Cambridge.
- Asuncion, A. and Newman, D. J. (2007). **UCI Machine Learning Repository**. University of California, School of Information and Computer Science <http://www.ics.uci.edu/~mlern/MLRepository.html>, Irvine, CA.

- Baader, F. and Sattler, U. (2003). Description logics with aggregates and concrete domains. **Information Systems** 28, 979-1004.
- Barsalou, L. W. (1992). Frames, concepts and conceptual fields. In **Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization** Lehrer, A. and Kittay, E. F., eds., pp. 21-74. L. Erlbaum, Hillsdale, N.J.
- Bauer, M. I. and Johnson-Laird, P. N. (1993). How diagrams can improve reasoning. **Psychological Science** 4, 372-378.
- Baumgartner, P. and Tinelli, C. (2008). The model evolution calculus as a first-order DPLL method. **Artificial Intelligence** 172, 591-632.
- Bennett, M. R. and Hacker, P. M. S. (2003). **Philosophical Foundations of Neuroscience**. Blackwell, Malden, MA.
- Berlin, B., Breedlove, D. E. and Raven, P. H. (1968). Covert categories and folk taxonomies. **American Anthropologist** 70, 290-299.
- Bertin, J. (1983). **Semiology of Graphics**. University of Wisconsin Press, Madison, Wis.
- Besnard, P. (1989). **An introduction to Default Logic**. Springer, Berlin.
- Boër, S. E. (1974). Cluster-concepts and sufficiency definitions. **Philosophical Studies** 26, 119-125.
- Borgida, A., Brachman, R. J., McGuinness, D. L. and Resnick, L. A. (1989). CLASSIC: a structural data model for objects. **ACM SIGMOD Record** 18, 58-67.
- Brachman, R. J. (1977). What's in a concept: structural foundations for semantic networks. **International Journal of Man-Machine Studies** 9, 127-152.
- Brachman, R. J. and Levesque, H. J. (1984). The tractability of subsumption in frame-based description languages. **Proc. of the 4th National Conference on Artificial Intelligence (AAAI-84)** 34-37.
- Broughton, V. (2006). The need for a faceted classification as the basis of all methods of information retrieval. **Aslib Proceedings: New Information Perspectives** 58, 49-72.
- Carnap, R. (1950). **Logical Foundations of Probability**. University of Chicago Press, Chicago.
- Cendrowska, J. (1987). An algorithm for inducing modular rules. **International Journal of Man-Machine Studies** 27, 349-370.
- Charniak, E. and McDermott, D. V. (1986). **Introduction to artificial intelligence**. Addison-Wesley, Reading.
- Clark, P., Hayes, P., Reichherzer, T., Thompson, J., Barker, K., Porter, B., Chaudhri, V., Rodriguez, A., Thomere, J. and Mishra, S. (2001). Knowledge entry as the graphical assembly of components. **Proceedings of the International Conference on Knowledge Capture** 22-29.
- Conklin, C. O. (1969). Lexicographical treatment of folk taxonomies. In **Cognitive Anthropology** Tyler, S. A., ed, pp. 28-41. Holt, New York,.
- Cook, S. A. (1971). The complexity of theorem-proving procedures. In **Proceedings of the Third Annual ACM Symposium on Theory of Computing**, pp. 151-158.
- Cook, S. A. (1983). An overview of computational complexity. **Communications ACM** 26, 401-408.
- Cooper, D. E. (1972). Definitions and clusters. **Mind** 81, 495-503.
- Davis, M., Logemann, G. and Loveland, D. (1962). A machine program for theorem proving. **Communications of the ACM** 5, 394-397.

- Donini, F. M. (2003). Complexity of reasoning. In **The Description Logic Handbook** Baader, F., Calvanese, D., McGuinness, D., Nardi, D. and Patel-Schneider, P., eds., pp. 96-136. Cambridge University Press, Cambridge.
- Eisenstadt, M., Domingue, J., Rajan, T. and Motta, E. (1990). Visual knowledge engineering. **Software Engineering, IEEE Transactions on** 16, 1164-1177.
- Ernst, N. A., Storey, M. A. and Allen, P. (2005). Cognitive support for ontology modeling. **International Journal of Human-Computer Studies** 62, 553-577.
- Fodor, J. A., Garrett, M. F., Walker, E. C. and Parkes, C. H. (1980). Against definitions. **Cognition** 8, 263-7.
- Fox, P., McGuinness, D., Raskin, R. and Sinha, K. (2007). A volcano erupts: semantically mediated integration of heterogeneous volcanic and atmospheric data. **Proceedings of the ACM first workshop on CyberInfrastructure: information management in eScience** 1-6.
- Frake, C. O. (1969). The ethnographic study of cognitive systems. In **Cognitive Anthropology** Tyler, S. A., ed, pp. 28-41. Holt, New York,.
- Fricke, M. (2003). **What are the advantages of Hyperproof-like reasoning systems?** British Society Philosophy Science.
- Funkhouser, E. (2006). The determinable-determinate relation. **Noûs** 40, 548-569.
- Gaines, B. R. and Linster, M. (1990). Integrating a knowledge acquisition tool, an expert system shell and a hypermedia system. **International Journal of Expert Systems Research and Applications** 3, 105-129.
- Gaines, B. R. (1991a). Integrating rules in term subsumption knowledge representation servers. In **AAAI'91: Proceedings of the Ninth National Conference on Artificial Intelligence**, pp. 458-463. AAAI Press/MIT Press, Menlo Park, California.
- Gaines, B. R. (1991b). An interactive visual language for term subsumption visual languages. In **IJCAI'91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence**, pp. 817-823. Morgan Kaufmann, San Mateo, California.
- Gaines, B. R. (1993). A class library implementation of a principled open architecture knowledge representation server with plug-in data types. In **IJCAI'93: Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence**, pp. 504-509. Morgan Kaufmann, San Mateo, California.
- Gaines, B. R. and Shaw, M. L. G. (1993). Eliciting knowledge and transferring it effectively to a knowledge-based systems. **IEEE Transactions on Knowledge and Data Engineering** 5, 4-14.
- Gaines, B. R. (1994). A situated classification solution of a resource allocation task represented in a visual language. **International Journal Human-Computer Studies** 40, 243-271.
- Gaines, B. R. and Shaw, M. L. G. (1995). Concept maps as hypermedia components. **International Journal Human-Computer Studies** 43, 323-361.
- Gaines, B. R. (1996). Transforming rules and trees into comprehensible knowledge structures. In **Knowledge Discovery in Databases II** Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. and Uthurusamy, R., eds., pp. 205-226. AAAI/MIT Press, Cambridge, Massachusetts.
- Gaines, B. R. and Shaw, M. L. G. (1999). Embedding formal knowledge models in active documents. **Communications of the ACM** 42, 57-63.
- Gaines, B. R. (2003). Organizational knowledge acquisition. In **Handbook on Knowledge Management: 1** Holsapple, C. W., ed, pp. 317-347. Springer, Berlin.

- Gaines, B. R. (2004). Understanding ontologies in scholarly disciplines. In **Proceedings 2004 International Workshop on Description Logics, DL2004** Haarslev, V. and Möller, R., eds. CEUR-Workshop Proceedings, <http://CEUR-WS.org/>, Whistler, B.C.
- Gaines, B. R. and Shaw, M. L. G. (2007). **Rep IV Research Manual**. Centre for Person-Computer Studies, <http://repgrid.com/RepIV/RepIVManual/>.
- Gangemi, A. (2005). Ontology design patterns for semantic web content. In **Proceedings of the Fourth International Semantic Web Conference: LNCS 3729** Gil, Y., ed, pp. 262 – 276. Springer, Berlin.
- Gaut, B. (2000). "Art" as a cluster concept. In **Theories of Art Today** Carroll, N., ed, pp. 25-44. University of Wisconsin Press, Madison, WI.
- Gennari, J. H., Musen, M. A., Fergerson, R. W., Grosso, W. E., Crubézy, M., Eriksson, H., Noy, N. F. and Tu, S. W. (2003). The evolution of Protégé: an environment for knowledge-based systems development. **International Journal of Human-Computer Studies** 58, 89-123.
- Grau, B. C., Horrocks, I., Kazakov, Y. and Sattler, U. (2007). A logical framework for modularity of ontologies. **Proceedings International Joint Conference Artificial Intelligence** 298–304.
- Guarino, N. and Welty, C. A. (2004). An overview of OntoClean. In **Handbook on Ontologies** Staab, S. and Studer, R., eds., pp. 151-171. Springer, Berlin.
- Haarslev, V. and Moller, R. (2001). RACER system description. **Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2001)** 2083, 701–705.
- Hanna, R. (2006). **Rationality and Logic**. MIT Press, Cambridge, MA.
- Hardin, C. L. (1989). Could white be green? **Mind** 98, 185-288.
- Hayes, P., Eskridge, T. C., Saavedra, R., Reichherzer, T., Mehrotra, M. and Bobrovnikoff, D. (2005). Collaborative knowledge capture in ontologies. **Proceedings of the 3rd international conference on Knowledge capture** 99-106.
- Hayes, P. J. (1977). In defence of logic. In **Proceedings International Joint Conference Artificial Intelligence**, pp. 559-565.
- Hayes, P. J. (1979). The logic of frames. In **Frame Conceptions and Text Understanding** Metzger, D., ed, pp. 46-61. de Gruyter, Berlin.
- Heilbron, J. L. and Kuhn, T. S. (1969). The genesis of the Bohr atom. **Historical Studies in the Physical Sciences** 1, 211-290.
- Horn, L. R. (1989). **A Natural History of Negation**. University of Chicago Press, Chicago.
- Horrocks, I. (1998). Using an expressive description logic: FaCT or fiction. **Proc. of the 6th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'98)** 636–647.
- Horrocks, I., Patel-Schneider, P. F., Bechhofer, S. and Tsarkov, D. (2005). OWL rules: A proposal and prototype implementation. **Web Semantics: Science, Services and Agents on the World Wide Web** 3, 23–40.
- Hyman, M. D. (2007). Semantic networks: a tool for investigating conceptual change and knowledge transfer in the history of science. In **Übersetzung und Transformation** Böhme, H., Rapp, C. and Rösler, W., eds., pp. 355-367. de Gruyter, Berlin.
- Jamnik, M. (2001). **Mathematical Reasoning with Diagrams: From Intuition to Automation**. CSLI, Stanford, CA.
- Johnson, W. E. (1921). **Logic**. Cambridge University Press, Cambridge.
- Johnson-Laird, P. N. and Byrne, R. M. J. (1991). **Deduction**. Lawrence Erlbaum, Hillsdale, NJ.

- Jonassen, D. H. (2005). Tools for representing problems and the knowledge required to solve them. In **Knowledge And Information Visualization: Searching for Synergies, LNCS 3426** Tergan, S.-O. and Keller, T., eds., pp. 82-94. Berlin.
- Kay, P. (1975). A model-theoretic approach to folk taxonomy. **Social Science Information** 14, 151.
- Keller, T. and Tergan, S.-O. (2005). Visualizing knowledge and information: an introduction. In **Knowledge And Information Visualization: Searching for Synergies, LNCS 3426** Tergan, S.-O. and Keller, T., eds., pp. 1-23. Berlin.
- Kelly, G. A. (1955). **The Psychology of Personal Constructs**.
- Kelly, G. A. (1970). A brief introduction to personal construct theory. In **Perspectives in Personal Construct Theory** Bannister, D., ed, pp. 1-29. Academic Press, London.
- Khalifa, M. and Liu, V. (2006). Semantic Network Discussion Representation: Applicability and Some Potential Benefits. **Professional Communication, IEEE Transactions on** 49, 69-81.
- Lu, J., Li, Y., Zhou, B. and Kang, D. (2008). Reasoning within extended fuzzy description logic. **Knowledge Based Systems** to appear.
- Lyons, J. (1968). **Introduction to Theoretical Linguistics**. Cambridge University Press, London,.
- Macnamara, J. (1986). **A Border Dispute: The Place of Logic in Psychology**. MIT Press, Cambridge, Mass.
- Mason, R. (2000). **Before Logic**. State University of New York Press, Albany.
- McDermott, D. (1987). A critique of pure reason. **Computational intelligence** 3, 151-160.
- Meheus, J. (2000). An extremely rich paraconsistent logic and the adaptive logic based on it. In **Frontiers of Paraconsistent Logic** Batens, D., Mortensen, C., Priest, G. and Van Bendegem, J.-P., eds., pp. 189-201.
- Meheus, J. (2003). Inconsistencies and the dynamics of science. **Logic and Logical Philosophy** 11, 129-148.
- Minsky, M. (1974). **A Framework for Representing Knowledge**. MIT-AI Laboratory Memo 306, Boston.
- Mumford, S. (1998). **Dispositions**. Oxford University Press, Oxford.
- Novak, J. D. (1998). **Learning, creating, and using knowledge : concept maps as facilitative tools in schools and corporations**. L. Erlbaum Associates, Mahwah.
- Parsia, B. and Sirin, E. (2004). Pellet: An OWL DL Reasoner. **Proceedings of the International Workshop on Description Logics** 104,
- Pearsall, N. R., Skipper, J. E. J. and Mintzes, J. J. (1997). Knowledge restructuring in the life sciences: A longitudinal study of conceptual change in biology. **Science Education** 81, 193-215.
- Perini, L. (2005). The truth in pictures. **Philosophy of Science** 72, 262-285.
- Pitt, D. (1999). In defense of definitions. **Philosophical Psychology** 12, 139-156.
- Pothos, E. M. and Hahn, U. (2000). So concepts aren't definitions, but do they have necessary or sufficient features. **British Journal of Psychology** 91, 439-450.
- Quillian, M. R. (1967). Word concepts: a theory and simulation of some basic semantic capabilities. **Behavioral Science** 12, 410-430.
- Revlis, R. and Hayes, J. R. (1972). The primacy of generalities in hypothetical reasoning. **Cognitive Psychology** 3, 268-290.

- Richens, R. H. (1956). Preprogramming for mechanical translation. **Mechanical Translation** 3, 20-25.
- Rosch, E. and Lloyd, B. B. (1978). **Cognition and Categorization**. Lawrence Erlbaum, Hillsdale, NY.
- Salmieri, G. (2007). Aristotle on the ontological basis of zoological classification. In **Proceedings of Conference Nature and its Classification**. Bristol University, Bristol.
- Shapiro, S. C. (1991). Case Studies of SNePS. **Special Issue on Implemented Knowledge Representation and Reasoning Systems, SIGART Bulletin** 128-134.
- Sirin, E., Parsia, B., Grau, B. C., Kalyanpur, A. and Katz, Y. (2007). Pellet: A practical OWL-DL reasoner. **Web Semantics: Science, Services and Agents on the World Wide Web** 5, 51-53.
- Sotirov, V. (1999). Arithmetizations of syllogistic a la Leibniz. **Journal Applied Non-Classical Logics** 9, 387-405.
- Sowa, J. F. (2000). **Knowledge Representation: Logical, Philosophical, and Computational Foundations**. Brooks/Cole, Pacific Grove.
- Stenning, K. (2002). **Seeing Reason: Image and Language in Learning to Think**. Oxford University Press, Oxford.
- Weitz, M. (1977). **The Opening Mind: A Philosophical Study of Humanistic Concepts**. University of Chicago Press, Chicago.
- Williams, L. V. (1983). **Teaching for the Two-sided Mind: A Guide to Right Brain/Left Brain Education**. Prentice-Hall, Englewood Cliffs, N.J.
- Woods, W. A. (1975). What's in a link. In **Representation and Understanding** Bobrow, D. G. and Collins, A., eds., pp. 35–82. Academic Press, New York.
- Zalta, E. N. (1988). **Intensional Logic and the Metaphysics of Intentionality**. MIT Press, Cambridge, MA.