Knowledge Representation and Activation

Abstract

Knowledge in any technical field rapidly becomes dense and complex, to the point where manual efforts to represent knowledge are inadequate. This paper discusses the problems in creating a knowledge representation that is self-extensible, that is, the knowledge can be used in the reading of text and its own extension, using the knowledge found in the text. The representation is now being read by itself, requiring its activation, there being no automaton or external program that understands the text, as there can't be if the knowledge is to be extended. The approach is intended to result in an accurate and complete semantic structure, leaving no nuance unrepresented. It requires a much more thorough analysis of the meanings of words than has been previously attempted, and a holistic integration of many aspects – existential, propositional, relational, temporal, locational, grammatical, activational, that have previously been held separate. The resulting structure may be contrasted with an ontological approach as to depth of semantics.

Introduction

First, why attempt to represent knowledge, when people can hold it in their heads? For knowledge that is strongly visual, people can manage very well. As knowledge becomes more logical, people struggle. It is estimated that people can hold no more than about six pieces of information in play at once – more than that and they become confused, or chunk things inappropriately. A knowledge representation that would allow people to observe and manipulate more complex situations would be of benefit, but only if the representation is valid – if it is forced to chunk, or throw away important factors (and if you can't represent all the interactions, you can't know what is important), then the representation has been a failure. This is an argument against manual construction of dense knowledge representations – if the person doing the modelling can't handle more than six pieces of information in play, then they are likely to fail two ways - not to fully understand the original knowledge, and to chunk their model by leaving out factors whose interactions they can't comprehend or that are difficult or tedious to represent.

We can approach knowledge representation using a method which is efficient for certain sorts of knowledge (and is clumsy or impossible for other sorts of knowledge) and requires skilled interpretation, or we can approach the problem using a representation with sufficient flexibility to represent whatever is required, and which can be used equally by a person or by a machine – that is, it is an activatable representation which can be extended by machine. We will attempt to show we have taken the latter approach.

Knowledge in text comes in many forms – two examples will be used here.

Clinical Note

A clinical medical note may connect complex concepts, but does so with simple links. It is not expected that some new concept will emerge from reading the note, only that the huge range of possibilities is greatly reduced for a particular patient by describing facts and surmises. There is some need for activation to eliminate alternative meanings, but otherwise the note uses clear and simple text, and is short.

Engineering Specification

An engineering specification can run to many pages, and describe a new concept that has not existed before – a flying submarine, say. The text is dense, typically has defined terms, either gathered at the front or scattered throughout the text, can make use of noun phrases that run to nine words, and can make references forward and backward across many pages, to the point where the reader is confused.

Knowledge Facets

Knowledge has many facets that need to be represented accurately, and function in concert with other facets. Some of them:

Existential Logic

Examples:

He can't swim.
No evidence of carcinoma was found.
The system shall be capable of self-correction.
The boat shall have the ability to right itself.
Section 9 is void under these circumstances.

Existential logic operates as the lowest layer, supporting all the layers on top of it. It is commonly ignored in representations, or only crudely represented. Existential logic is orthogonal to propositional logic, so attempts to represent it at the propositional level fail, and yet it is easily and simply represented in text. It is also easy to control the existence of entire areas of knowledge, one aspect of topological change.

Propositional Logic

Examples:

He swam the river.
Venous invasion has occurred.
The contractor shall provide adequate facilities.
We looked for evidence but found none.

We are using propositional logic to cover all uses of simple logic except existential and object logic, so it covers logic states of relations, logical connectives between statements and clauses, and logical connection to the discourse. It also covers operators like "but", whose job it is to handle conflicting logical states without error. There are frequent occasions where propositional and existential states are connected:

If John can't swim and the bridge is damaged, he will have to wait.

The appropriate logical state is returned for the next logical operator – in this case a propositional logical state for the implication operator.

The Logic of Groups (not Group Logic)

Examples:

> Jack and Jill went up the hill.
> The cost of welding the hub of the rotor and the shaft was too high.
> It is characterised by increasing and spreading muscular weakness.

Objects and relations may be gathered into groups, with different types of operations on those groups – effectively object logic - A and B, C or D, all, any. There can be considerable ambiguity - A of B and C - only mitigated by using the semantics of the objects.

Relational Logic

Examples:

> The tumour may have metastasised to the liver.
> John went to the pictures.
> John wanted Fred to go to the pictures.
> Olga thought that John wanted Fred to go to the pictures.
> It appears the trial drug regime has failed to prevent the clustering of macrophages at sites of inflammation.

Relations allow the interconnection of objects, including relations (this is unusual in knowledge representation, trivially easy in text). Relations are controlled by existential and logical states. They also allow layering of relations without limit. Clausal or TransClausal relations, like “think” or “prevent” or "told", provide logical states to relations they control, taking the place of the discourse (providing a “virtual discourse”, as "He told John the bridge was damaged").

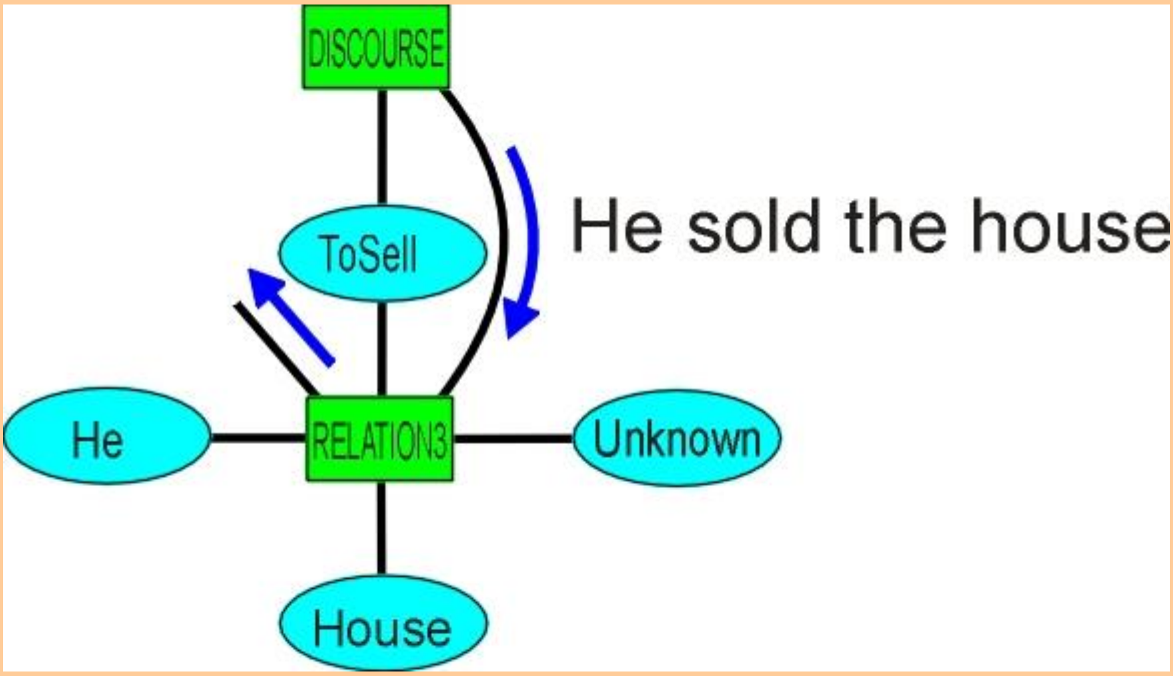

In the figure, the relation ToSell receives a propositional logical True from the Discourse – the statement is True. The RELATION3 operator manages the correlation between the existential and propositional connections – if the propositional connection is true, the existential connection must be true, whereas if the existential connection had an incoming True, the propositional connection could be anywhere from False to True. The ToSell relation is a reflection of ToBuy – a mapping operator is used to link them.

States

Examples:

> The door is shut because the patient died.
> He has a slightly elevated temperature.
> He was determined to succeed.

Many actions change the state of an object – it may be how much a door is open, whether one is logged on to a computer, the state of someone's health. These are complex states, that cannot be represented without reference to the object –a large bacterium is smaller than a small child, a rapid crawl is slower than a slow run. There is a network of active interactions among states – their state of health is irrelevant if the person has died.

Temporal Logic

Examples:

> In the last six weeks, there has been a sharp increase in infections.
> Influenza causes a rapid onset high fever which lasts three to four days after infection.
> The goods must be delivered today.
> Testing shall occur in parallel with assembly.
> He went out at 6 pm, shutting the door behind him.

Relations carry time information – start, duration or termination – and these times can be poked in anywhere in text. Times that are not explicit can be inferred from relations around them.

Locational Logic

Examples:

> The equipment was site tested.
> Non-neoplastic lung tissue shows centrilobular emphysema.
> The bronchial resection margin shows squamous metaplasia.

All physical objects (including relations) have locations. Location information can include explicit location, being near something else (on the margin of) or proximal, or a heading (northeast or anteroposterior).

Numbers and Units

Examples:

The antenna shall operate in the X-Band.
There shall be less than two defects per thousand.
WBC – 3.5-9.0x10^9/L.

Numbers form a separate network structure, controlled where appropriate by propositional and existential operators, and the existence of the objects of which they may be attributes. The interaction of numbers, in turn, can provide logical values. In most text, there are surprisingly few numbers, with almost all meaning being found in relations among objects. However, getting the units for dimensions right is rather important. The system is adept at ranges.

Grammatical Knowledge

There is a network for parts of speech, with meanings where a word can be a noun or a verb or an adjective. Sentences are turned into chains of objects, and grammatical patterns are used to build grammatical structures on the parse chains. The grammatical patterns make frequent reference to semantic structures – both those pre-existing, and those being built by the text. We have not yet made the system capable of extending its grammatical patterns by reading text, although the system has been fitted with a word unraveller, as terms newly constructed from known prefixes, stems and suffixes are often found in medical text.

General Stuff

We haven't mentioned inheritance, or hardwired relations such as **INVOCATIONS** and **COMPONENTS**, where efficiency and speed override flexibility – we are not expecting a kidney or a vehicle not to be a noun. A large document is itself a complex structure, with the ability to point at any part of itself, and switch parts of itself on and off.

The dictionary is a dynamic construct, so it can extend when new definitions are encountered in the text. The dictionary has absolute collocations, called wordgroups - "bus driver" - and non-absolute collocations - "bus stop" - where the concept can be described, but further checking is required to ensure the concept is being referenced.

Putting It All Together

All of the knowledge facets emerge in even very small pieces of text, making it inconceivable that anything more than trivial or extremely surprising non-visual knowledge can be represented by hand.

Conclusion

A system which attempts to handle all the complexities of text, and the objects and relations described therein, requires a large closely integrated set of abilities for the representation of knowledge. Without all these abilities, inadequacies rapidly surface and representation of knowledge is poor.

Appendix

Layered Relations

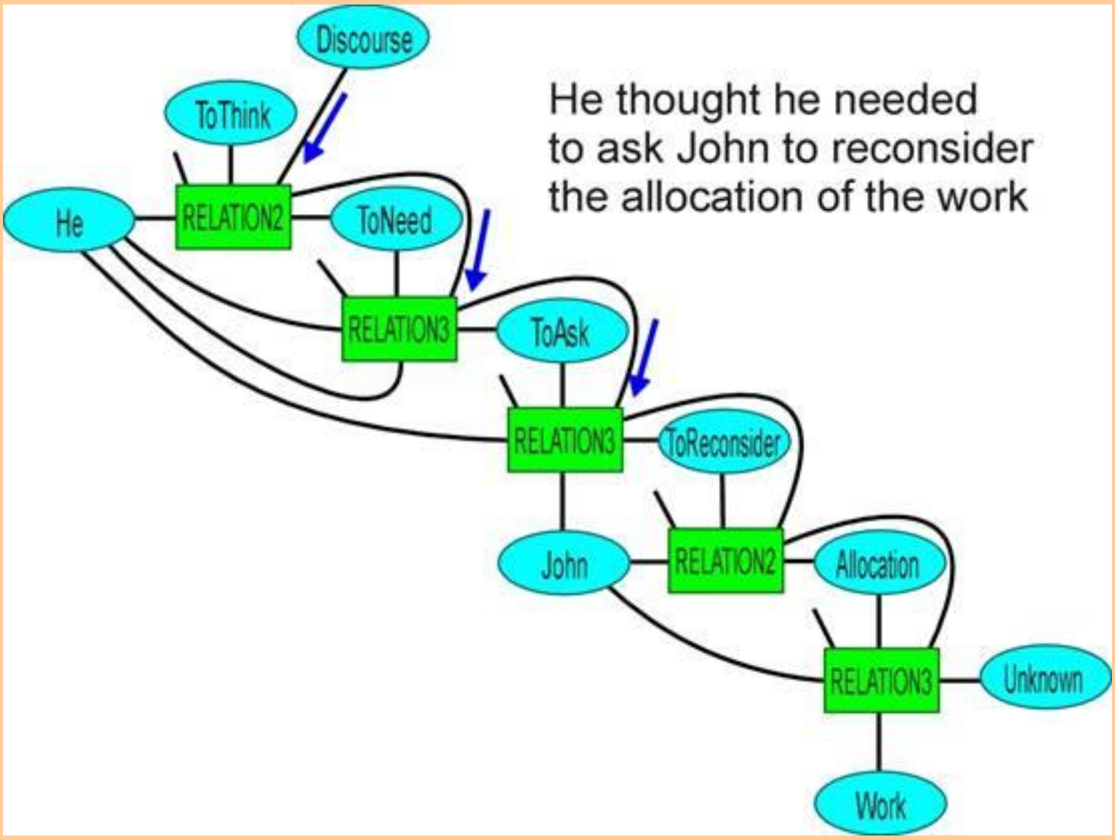

Figure 2: Layering

Complex text can have many layers, where a relation like “ToThink” controls the context of the relations below it. With layering, the relations below may return a False, but it does not override or cause an inconsistency with the logical state coming from above –

John thinks Fred is guilty.

John can still think Fred is guilty, even if he is found innocent.

Object Groups

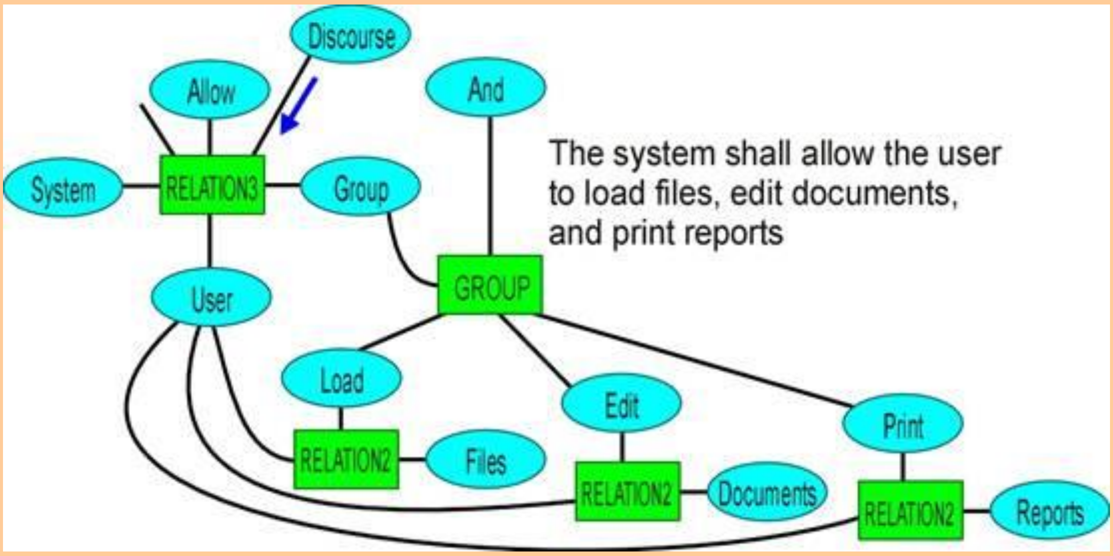

Figure 3: An Object Group

Object groups can be all objects – a car, a plane or a train -, or all relations – they ran and jumped and frolicked – or a mixture – It was either a new car or a trip to France. The object group radiates the properties of its elements.

Grammatical Patterns

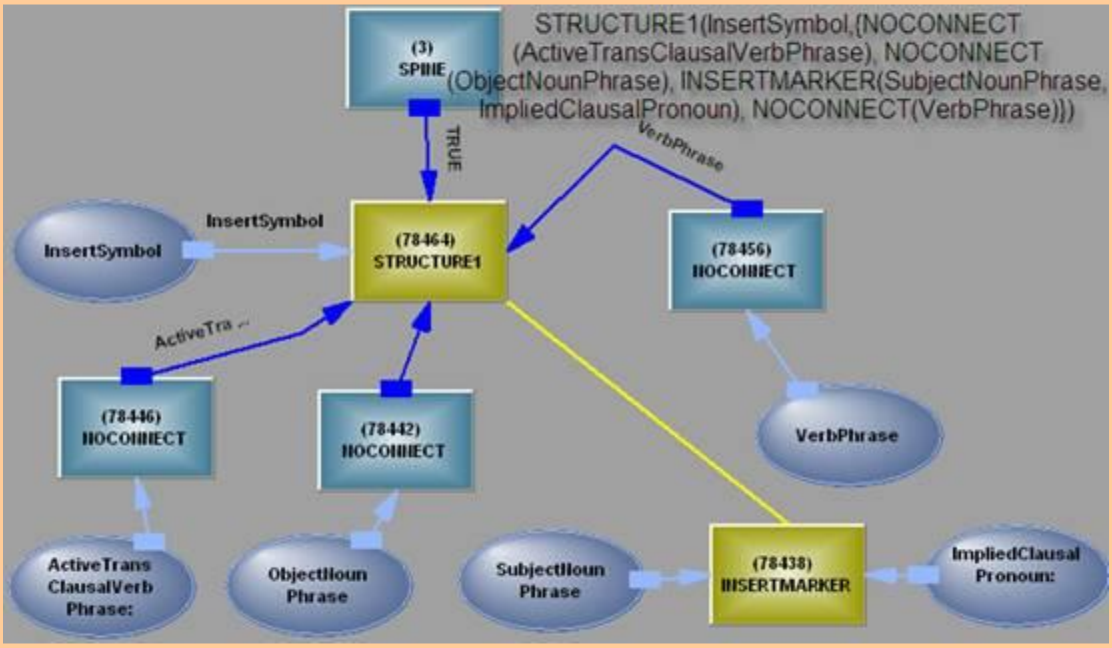

Figure 4: An active Grammar Pattern

STRUCTURE1 patterns are used to build the grammatical structure – they mostly clone themselves onto the grammatical structure, and new patterns access the newly combined structure. Some patterns are used to modify the current structure – the one in Figure 4 adds a Clausal Pronoun to the parse chain.

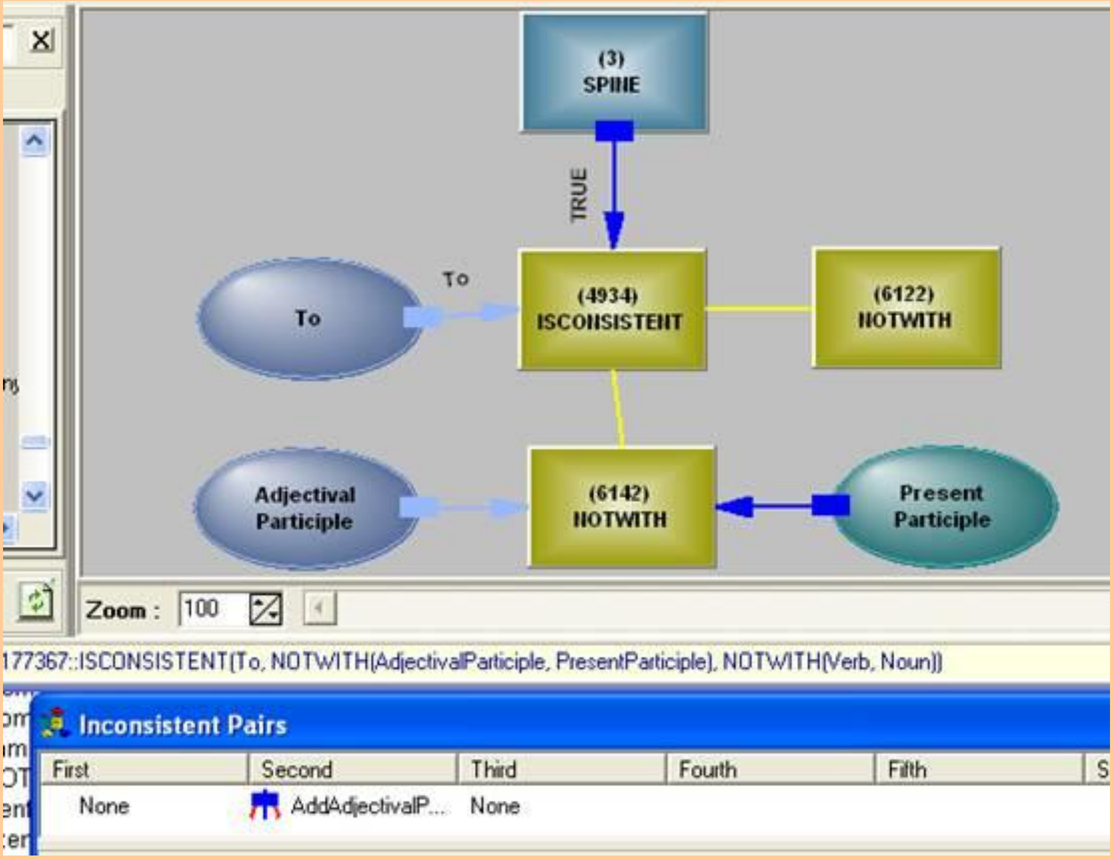

Figure 7: An ISCONSISTENT operator

The dictionary has absolute collocations, called wordgroups – "by virtue of", "descending colon". The COLLOCATION operator allows more processing around groups of words that can't be guaranteed to be absolute collocations – "ASTM 2501A:2007" or "from 6 to 9 metres" for example, but only require a little more analysis. Some possible collocations need analysis of verbs and noun phrases before it can be determined that there is a noun phrase – "The red bus stops at the green bus stops", so can't be handled with COLLOCATION operators.

Document Structure

Documents have structure too, and it becomes important to represent that structure, so the descriptions in the document are used exactly as the writer intended. The writer can then reference some part of the document, as

> 5.7.3.4. The On-Board Audio Announcer shall provide …as described in clause 7.10.3.19.3.c.

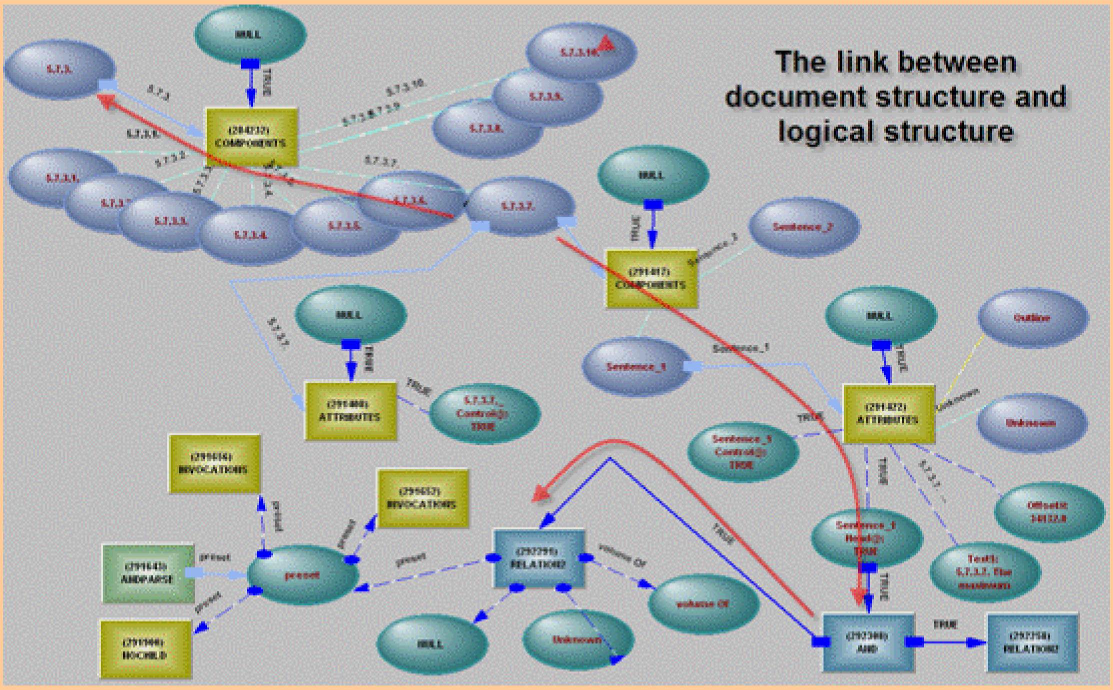

Figure 9: Document reference

Definitions

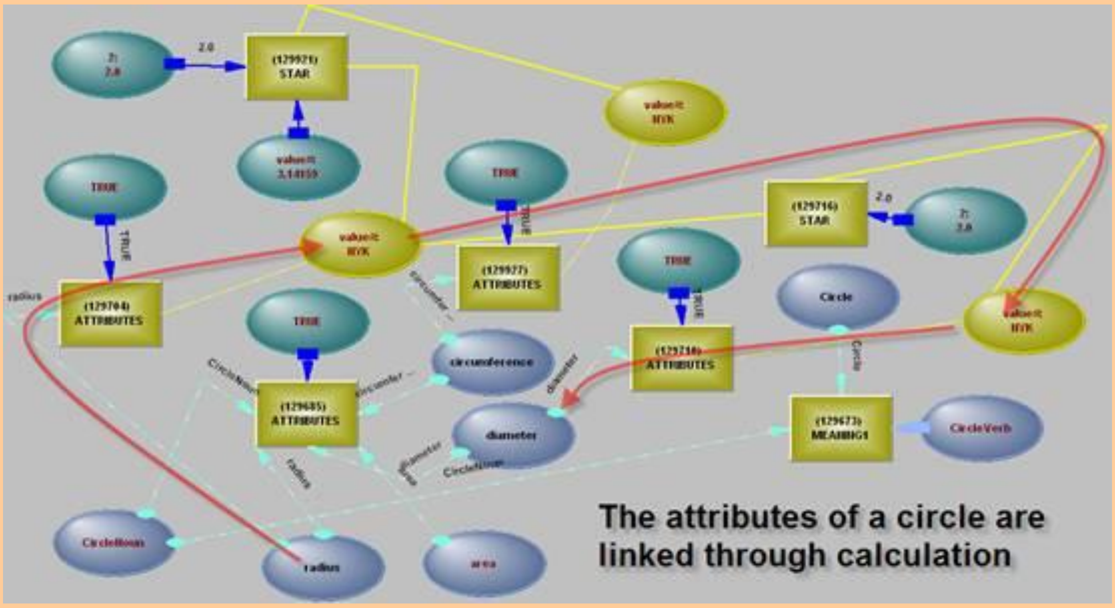

Figure 10: Circle attributes

A circle has attributes which have numerical values – radius, diameter, circumference, area. Knowing the numerical value for any one means that all are known – there is an undirected numerical network that provides this capability.