Information Based Learning and Knowledge Structure

Chris Tompsett

1. Introduction and background

The systems that we term 'Information Based Learning', apply knowledge-based structuring techniques to the design of hypertext and hypermedia systems. Such systems (described more fully in [Tompsett, 1991]) are intended to provide a learner with an inter-linked set of information resources to explore. Implicit in these systems is the recognition that information, whether presented as photograph, text or other medium, selected for one step in educational activity can be informative for other educational activities.

This places emphasis on the information that is conveyed by a chunk (photograph, extract of text or other medium) at a number of levels : the restricted, context-free information that is conveyed by the chunk itself ; the richer but bounded information that is conveyed by the context defined by the other information that is available within the resource system and the placing of that item within a sequence or set of items explored, and the potential of the information contributing to the subjective experience of an individual at a particular place and point in time. Concomitant with this emphasis is a limited specification of learning goals, allowing a learner to explore resources according to their subjective view of what they wish to learn and, therefore, a separation of the information content of a resource from a description of its original intended use.

Hypertext (and similarly hypermedia) systems depend for their effectiveness on the ease with which users are able to identify and absorb relevant information from the available resources. With the increasing variety of media and volume of resources that are available, providing effective mechanisms to link relevant resource items together becomes critically important. This requires access techniques that are more focused than generic information retrieval approaches (for example those on the WWW) but less constrained than those that are individually authored.

The core of the design process is the use of knowledge representation, to encode the content of each item and also to determine the basis on which links should be created. However knowledge representation on its own, whilst achieving benefits for identifying appropriate connections between 'items', may fail to encode significant information content and we therefore place high value on the difference between the information content within any particular chunk and its representation within the formalism. For this reason, the design process proceeds from authoring of the text, to generation of links (section 3).

This paper describes the design and evaluation of a hypertext system with an underlying link structure created on a representation of the domain similar to a weakened version of KL-one. This system has been evaluated on a national trial and is now to be released following a reassessment of the knowledge content. The effects of these changes are assessed. The potential additional benefits that can be gained from a stronger representation are considered along with the opportunities afforded by further exploitation of the existing formalism.

2. Architectures of hypertext

The generic notion of a hypertext system, as a set of nodes of information connected by links affording a choice of exploration paths to the reader, offers a powerful unification at the computing level but a dangerous simplification in information system design. Traditional publishing is unified at the level of producing and reading from the printed page but authors of technical manuals are unlikely to have much to discuss with writers of children's stories or authors of educational text books. However, technical issues that relate to the creation of a hypertext/hypermedia system, and its subsequent 'reading', by a learner hide the need for radically different solutions for different information system purposes. This is not to suggest that various models of hypertext are incompatible, but to indicate that different models be designed for different styles of reading and reflect different levels of analysis.

The system that we describe here is designed for an educational system, the major implication of which is that the learner's cognitive ability in a defined area of knowledge should have been changed through the process of reading the hypertext material. This does not preclude other uses of the same material, for example in decision support. Exploring the hypertext structure results, of course in a linearised experience for the student — the underlying hypertext links merely allow different students to have a different linearised experience. However, in exploring the hypertext structure, readers will be developing their own model of the hypertext structure and it is our thesis that it is the student's development of this model which represents the process of learning (the relational stage in the SOLO Taxonomy [Biggs, Collis, 1982]).

The readers/learners will be generally in later secondary education or above and working in domains that are typically semi-technical for which we can build a weak problem solving model. More technical or less technical domains can still be analysed using this approach but the link structure of the system will have different properties. We also expect that readers have some familiarity with the domain as we currently do not expect the system to function *in ab initio* mode. The learning task we seek to support is the development of deeper conceptual knowledge, through active exploration of the hypertext material. The readers are assumed to be committed to learning so the design does not take into account the problem of motivating the students.

The hypertext material in the original system was edited into one corpus by a single author, but that is not a requirement of our approach. The consequences of using different authors have an impact on the style of the text material, which we believe will always remain an issue, but will have a significantly less impact on the generation of appropriate links. The design allows for a maximal set of typed links to be created, creating a tightly linked hypertext structure. Typing of the links then allows both open exploration of the material according to the interests of the student, or a more restricted exploration through increasing selectivity of available links from any one node. Implicit in this is a requirement that a weak knowledge representation of the domain can be created that will not impede development of a learner's understanding. This removes the limitations of flexibility inherent in author created links, but does not necessarily preclude their inclusion.

2. 1. Characteristics of knowledge-based analysis

The intention is to step back from the authorship process and to ask why a link should exist between two particular chunks of hypertext. This approach seeks to provide an underlying semantic framework (as with Thoth-II, [Marshall, Irish, 1989]) within which the intention of the author in providing specific links and the intention of the learner in selecting particular links can be identified as rational choices. This separates out the two aspects of linking — deciding why a link should be created from an anchor in any one chunk from deciding the chunks, or sub-component in a chunk to which it is linked. As a corollary, it becomes possible to provide a mechanism whereby any new chunk of material can be automatically integrated into an existing hypertext system.

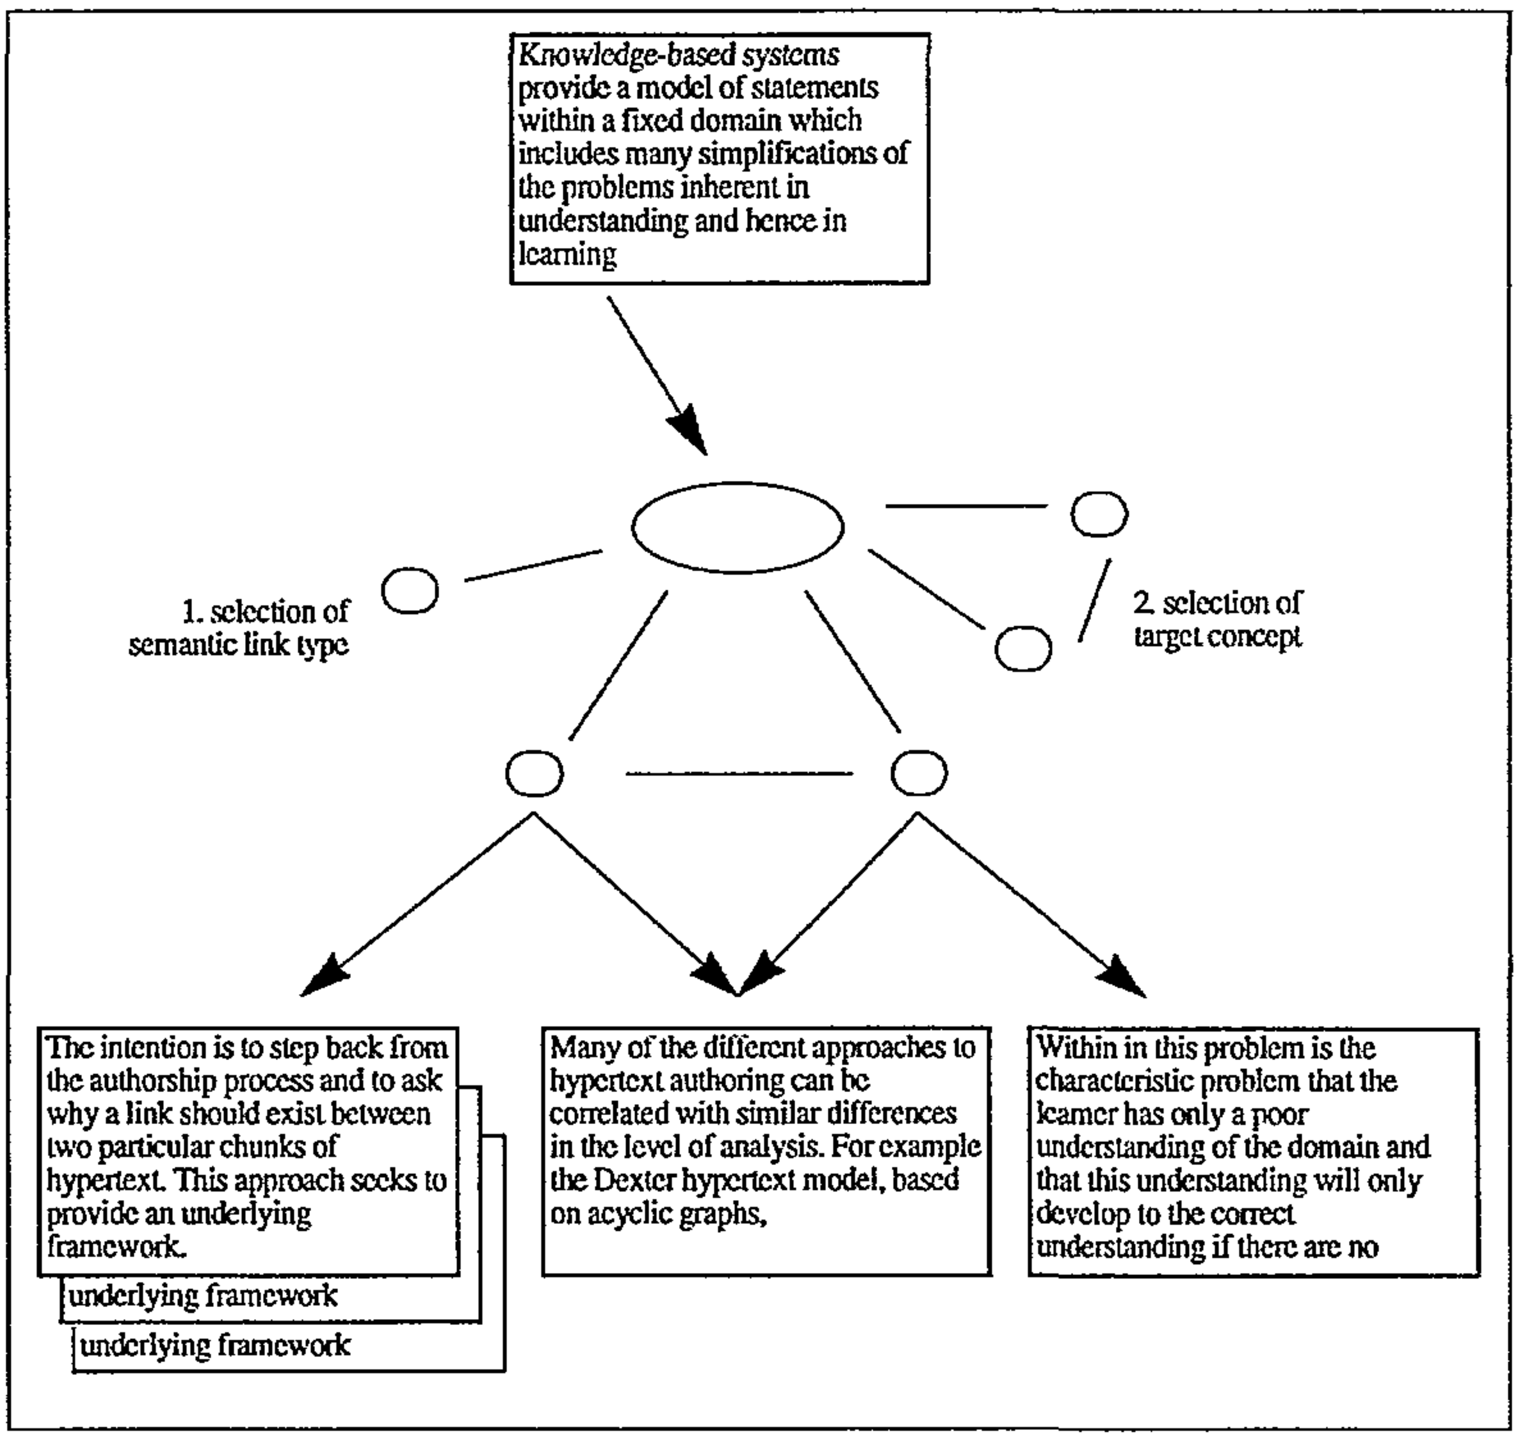

Figure 1

Two stage selection process

Knowledge-based analysis and representation of a particular domain typically functions to provide a computational model of that domain. This introduces three aspects, all of which are central to the design of our current system. These are : the semantic level of representation, the compromises made in developing a model and the potential to capture reasoning within the model.

2. 1. 1. — Semantic level

The notion of semantic level, arose out of the early work in the field of knowledge-based systems (*e. g.* [Bobrow, Winograd, 1977], Hayes [1981] on frame representation, and Brachman [1979] on semantic nets). Brachman's levels, typified by the ontology of the concepts and semantic links used, range from the implementational, the level closest to computation, through the logical level, the epistemological level — representing concepts and their interrelationships, and the cognitive level — concerned with representing human behaviour (and human mistakes).

Although the various authors do not agree on the division between semantic levels, the need to link the relevant levels is accepted. Each level can be considered as providing a model of the domain and the development of a computational model can be seen as the result of a series of transformations from the human behaviour to the resultant computational system with its associated mechanical or mathematical [Popselov, 1979] behaviour. Although the system we describe here has been implemented and evaluated, the intention in the paper is to describe the implications of modelling the creation of a hypertext system at the one particular level.

Different models of hypertext can similarly be correlated with similar differences in the level of analysis. For example, the Dexter hypertext model, based on acyclic graphs, [Halasz, Schwartz, 1994] provides detail about the common aspects of the implementation of hypertext systems and deliberately avoids discussion about the content, reflecting the implementational approach. Argumentation systems, such as gIBIS [Remde *et al.*, 1987], with a fixed set of links with a clearly defined behaviour, reflect a logical level, in capturing the particular logic of a design formalism or a framework for construction of a range of formalisms [Marshall *et al.*, 1991]. The underlying link architecture of the system described here is based on epistemological semantic level as epitomised in KL-One [Brachman, 1979].

The effect of selecting this level of modelling of the text as the focus for the link authorship process has two immediate effects. Firstly the knowledge domain underlying the hypertext material has to be modelled conceptually. This requires, within the KL-one formalism, that a semantic net, based on hierarchical conceptual structures needs to be developed, together with a set of role descriptors and structural descriptors (see section **3**). Secondly, the initial focus of attention in generating the underlying conceptual model becomes the individual statements in the domain. This initial focus in generating the conceptual model is later complemented by more general hypertext considerations (see section **3**).

The function of the epistemological level is to guide the answers to the question : *'If I am reading this text, what could be relevant to read next ?'*.

Particular statements can be metalevel statements, describing the conceptual structure, role descriptors and structural descriptors or the statements that represent the consequences of reasoning in the domain. Similarly the occurrence of a term within a statement can function either as a condition within that statement, or more generally, act as a sign for the concept in general, the first interpretation emphasizing the relevance of the statement to reasoning within the domain, the second relating to a broader reflection.

2. 1. 2. — Compromises in knowledge representation

The modelling process, from the real world to a knowledge representation, in order to gain the benefits of automatic knowledge processing, inevitably introduces compromises. Whilst these may be critically significant in restricting the potential of intelligent tutoring systems [Tompsett, 1988], applying knowledge-based techniques to support hypertext, rather than to replace it can ameliorate these problems. Two interrelated aspects of this analysis are particularly relevant : granularity and circumscription.

Granularity arises from the need to limit the level of knowledge-based analysis to one that is productive. Knowledge representation gains no benefit in differentiating between terms where the concepts required for the differentiation fail to link with other terms of the knowledge-base [Clancey, 1983]. Invariably this set of terms will be narrower than the range of base terms used by an expert, or those existing in a series of texts covering knowledge in the domain. For example several terms may be freely used as synonyms within text, reflecting usage of the concept in different communities or distinguishing between formal and non-formal usage of a concept. Since a hypertext system can retain the full range of terminology used, providing the underlying representation allows for the benefits of knowledge-based representation to be realised whilst making the mapping between text and knowledge-base terms explicit — a mapping that is essential for any automatic generation of links.

A similar compromise exists in knowledge representation in determining the depth of representation that is used (for example, the layered heart model [Bratko *et al.*, 1989]). This assumes both that a deeper level of analysis is achievable and that it provides a better explanation, neither of which may be the case [Minsky, 1975]. Retaining the text version allows any explanations that are relevant to be presented in equivalent manner. The depth of knowledge representation is required to provide only the links between text that should be linked, rather than logical explanations and allows the representation to be performed at different levels across the domain.

2. 1. 3. — Reasoning in knowledge-based systems

Within traditional expert system development, the power of the system is dependent on the extent to which the system can make valid deductions based on the information about particular problems presented to it. The power of the reasoning system is therefore based on three elements, the inherent knowledge encoded in the knowledge representation, the faithfulness with which the inference engine mimics reasoning within the domain and the compatibility of the inference engine itself. Generalising from early results in logic, the greater the potential of the first two elements, the less likely it is that the results are computable [Gödel, 1931]. However, in terms of representing the links between chunks in a hypertext system, we do not require that the logic of the representation is fully computable, only that we can compute the links between the chunks. This changes the computation problem from discovering all potential inferences to validating a particular inference. This is then affected by other factors, such as the size of chunk that is linked, since this might include a number of different statements, and the conceptual distance between the current chunk and the target chunk representing the conclusion. Learners are likely to become disoriented if the system creates links between chunks that the learner does not recognise as being linked.

2. 2. Traditional architectures

The low level architecture provides a rationale for why chunks could be linked, but does not necessarily produce a good hypertext system. Aspects such as higher level structuring, information retrieval and navigability are conventionally seen as essential system qualities and are well developed in other systems.

2. 2. 1. — Higher level structuring and continuity

Architectures based on conventional book constructs with additional features (*e. g. Superbook, More, Guide)* provide a higher level structure corresponding to an outline of the complete text. This can be used both as an abstract map of the system, to avoid the 'lost in hyperspace problem', as well as to provide for continuity. These linked properties have different consequences for the reader however. Whereas the outline structure provides support for the reader to select chunks of text in a sequence that is not the one determined by the author, continuity allows the user to avoid choosing the next piece of text. For information retrieval tasks the first aspect is more important, for learning a lack of continuity may prove an impediment. Construction of a continuous piece of text by an author could be seen as merely requiring the use of one or more guided tours [Marshall, Irish, 1989]. Although these typically provide the sequence of chunks of information, the commentary that glues the sequence together is left to the reader [Landow, 1987]. The size of chunk and number of available links become central features in determining the continuity of the hypertext system. Large chunks and a small number of links per chunk in a guided tour allow the original author to establish continuity within the chunks and require little effort on the reader's behalf to make sense of the hypertext links that are followed. Small statement size chunks and a high degree of linking increases the difficulty for the reader.

Lower level continuity within the current system is retained by starting with sources that already exist in linear text form and adopting a hypertext chunk size that is larger than the level used for knowledge representation. The continuity allows an author's original interest in juxtaposition of ideas, lost in the representation process, to be retained within chunks. The knowledge representation system also allows for description of the text chunk at a level higher than the component statements, reflecting higher level aspects of the semantic representation. Additional higher level structuring is provided through hierarchical features of the knowledge representation.

2. 2. 2. — Textual/information retrieval

Although conventional information retrieval processes may be efficient in extremely large corpuses, their dependency on textual similarities and frequency is weak. Similarities, contrasts, contexts or significance are lost — suggesting that the reader must understand the domain to a high degree in order to be able to select the relevant texts from those that are retrieved. More flexible retrieval methods, such as hierarchical/lattice indices [Bruza, 1990] or context dependent techniques, such as the fish-eye approach [Remde *et al.,* 1987], can provide more efficient search techniques but still require a comprehensive understanding of the domain.

The approach we have taken in our current project has been primarily determined by the requirement that our readers are learners. Rapid identification of a small number of hits from the entire system is of less relevance than a focused selection of hits from within any particular point within the system, reinforcing the potential interconnections within the domain, rather than selecting specific ones. Conventional effectiveness of information retrieval might result from combination of the techniques that we have applied, but would require further extensions to the size of the system for any technical analysis to produce significant measures of accuracy.

2. 2. 3. — Navigability

The thesis underlying our design is that existence of a link between two chunks should be rational. In developing the particular knowledge-based rationale, it is clear that there are many justifiable (and hence implicitly typed) reasons. The design is therefore based on providing a restricted three step information retrieval process to generate links.

Anchors exist either from terms within a chunk, or from the chunk itself, based on the higher level aspects of the semantic representation. Depending on the type of anchor, a range of link types are available. Through systematic typing of these link types based on the knowledge representation the reader is presented with the context within which to explore the next chunks. Given the anchor type and the link type, a range of targets is identified and sequenced according to their potential relevance.

3. KL-One representation

The base level KL-One knowledge representation, even without the application of deduction provides a straightforward mechanism for associating any concept with semantically related concepts and provides the basic typing of potential links. In addition to providing a semantic association for each individual concept, and hierarchical structuring of particular concepts through generalisation and specialisation, it also generates a hierarchical clustering of the complete set of concepts which are used to define the context for a set of statements within a chunk. Classification of particular statements within the rule based system allows for more detailed linking and sequencing of cards.

The underlying logic then provides an additional mechanism to increase selectivity, based on a game theoretic model we term 'conversational diagnosis' in the current system.

3. 1. Roles and structure statements

The basic KL-One representation links each concept with its potential component concepts. As an example from the current system for community pharmacists, the concept of a 'case' has named roles of 'patient problem', 'therapy', 'diagnosis' and 'advice'. The epistemological level typically classifies each of these roles according to whether the role must be filled, their cardinality and the potential concepts that can fill the role. For example, the role of patient problem is required and singular, advice is required but plural (*i. e.* non-zero), whilst therapy and diagnosis are optional and plural (*i. e.* may be zero, one or many).

Typically role names and concept names are the same, differing when the same type of concept fulfils two different roles. For example, the concept 'patient problem' has roles of 'patient', 'present medication', 'previous medication', 'presenting symptom' and 'other symptoms'. Both medication roles require a drug as fillers and two symptom roles are filled by 'symptom' concepts. For any particular concept, the role names provide one of the basic semantic links that could be created by the system.

Structure statements provide the intra-conceptual constraints, ensuring that the sub-components of a concept are consistent with each other.

Specialisation of particular concepts is created through matching specialisation of role fillers or through additional intra-conceptual constraints. Specialisations can act as partitions, requiring that a member of the concept must belong to one of the sub-dividing concepts, or simply represent a possible sub-division. As a simple example, a 'cough' symptom is partitioned into 'non-productive cough' and 'sputum production', and one of the partitions of 'cough preparations' is 'expectorants'. If a symptom is classified as a 'non-productive cough', this excludes any medicine from consideration that is in the class of 'expectorants'. The specialisation hierarchies provide two more forms of semantic link, both to a deeper level and generalising to any of the higher levels.

In addition, the partitions provide a further semantic link — between concepts and direct alternatives to that concept within the hierarchy.

3. 2. The context of a chunk

The semantic links described are used to link at the chunk level, by identifying one concept as the *primary context* of the chunk. This is based on the hierarchical clustering of concepts created through the role structure.

The *primary context* for a statement is defined as the enclosing concept required to include all the roles that are included in the statement. These will normally be the roles of a concept, sub-roles of the roles and so on, and the context is the concept that includes those top-level roles. Where the statement explicitly defines aspects of a role, such as 'presenting symptom', where all the other roles are roles or sub-roles of a 'symptom', then the *primary context* will be the parent concept for the 'presenting symptom' role, *i. e.* a 'patient problem'. The context of a set of statements, *e. g.* A, B and C, is defined as the *primary context* for the combined statement (A & B & C).

The definition of a context for a chunk based on analysis of the collection of statements within the chunk is critical for providing higher level continuity within the information retrieval process. It defines a level of linking between chunks that exists at the chunk level, rather than through reference from individual terms within the chunk.

Each of these approaches to creating a semantic link provides a recognisable link type. In browsing using these chunk level links, the reader is therefore provided with a one or two stage menu, forcing the reader to consider the particular change in context that they are intending to follow.

In all cases, defining the *primary context* in this way results in a context that is a specialisation of a more general concept. In some of these cases, the specialisation statements will relate this specialisation hierarchy to another specialisation hierarchy. So if the *primary context* for a chunk is defined to be 'expectorants', then it can be presumed that the context should also include the fact that the symptoms will include 'sputum production'. In this way we define the *extended context* of a chunk to be the *primary context* and any directly implied specializations.

The potential to extend this definition of context further, to attempt to compute all possibly implied specializations would require more specific analysis of computational performance determined by the full definition of KL-One. This would require the system to exclude information through combining the information from several chunks, a level of complexity that is not directly afforded to the reader by the system. Since this reduced scoping is available through an alternative mechanism (section **3. 5.**), this is not currently used.

3. 3. Links from individual terms

The creation of semantic links based on the context does not imply that anchors cannot exist within a chunk. Any term in a chunk, that corresponds to a semantically significant concept or role can be used as an anchor to generate further links. However, links created from these anchors are seen as branching away from browsing at the chunk level and do not generate an intermediate choice menu. This is designed to separate the semantic links at the chunk level, and the links made from references within chunks.

Within the existing project, the interface was required to function without a pointing device. This forced anchors from terms to be subsumed within the chunk level links, and, as a consequence, introduced an additional point of reflection in following these links.

3. 4. Sequencing of target chunks

The approach described above defines the method for considering the links that could be generated from a single chunk and are not dependent on the definition of the knowledge structure and not dependent on the existence or availability of targets for any of these links. There is, of course, an implicit dependence, since the knowledge structure itself will be dependent on the content of the other chunks.

To complete the process of defining the links we also apply knowledge-based techniques to selecting and sequencing the potential target chunks on the basis of the same within chunk description and some additional chunk level classifications. This allows the linking process to be dependent only on the knowledge structure and descriptions of each chunk. Links can automatically be generated both from and to a new chunk without full consideration of the content of other chunks, unless the new chunk alters the underlying knowledge domain.

Selection of target chunks is based on the same criterion as for linking. Each semantic link generates a target concept and we select chunks that refer to the target concept. This is either because the target concept is within the *extended context* for the chunk, or occurs in a knowledge statement within the chunk. This process provides a variable set of target chunks representing a range of closeness in linking to the semantic link.

Sequencing these chunks is based on three primary levels of priority, with further sequencing within each level. The primary sequencing follows the occurrence of the target concept as the *primary context,* then within the *extended context,* and then as a concept referenced within the chunk. Within the *primary context,* further sequencing is effected through the use of simple labelling of the function of the chunk, currently restricted to either 'introduces' and 'describes'. A 'describes' label identifies a chunk that is not replicated, for that concept at that particular level of the concept's specialisation hierarchy is not provided in more detail on other chunks, otherwise the 'introduces' chunk is applied. 'Introduces' chunks are sequenced before 'describes' chunks.

Sequencing within the lowest level is based on two further criteria. The first criterion sub-classifies the target chunks on the basis of whether the target concept occurs in the conclusion or the condition part of the inference. In order to correspond to logical inference, chunks in which the target concept occurs as conditions are sequenced before those that only occur as conclusions. These further sub-groups can then sequenced according to the frequency of occurrence.

3. 5. Local and global selectivity

The decision to define the target of a link to be a set of chunks can either be seen as an information retrieval process or as an extension of the hypertext concept. Whilst it is interesting to note that this semantic approach places only minimal significance on frequency of occurrence, of more concern in this paper is the impact of semantic selection on the learning process and its extension from the local selectivity, described so far, to a notion of global selectivity for focusing the exploration.

When browsing through the system, each chunk defines a set of possible links, and hence a set of sets of chunks that are potentially linked. Each powerset, therefore, partially constrains the reader's exploration of the information. This constraint is weak, since there may be many different links from any one chunk and many target chunks for each link. Further, if the reader continues to browse from one of these target chunks, there will be a new powerset of destination chunks. Since there is no imposed continuity, other than that links exist, browsing only imposes a local selectivity although the potential to pursue a consistent [Chelnokov, Zephyrova, 1993] line of reading is enhanced through appropriate selection of named links by the reader. In order to provide better support for more directed exploration by the reader, a mechanism has been provided to allow the reader to define a semantically coherent subset of the domain for exploration This applies constraints to all browsing — introducing a degree of global selectivity that persists between choices of links.

The intention of this global selectivity is to apply inferencing within the domain to eliminate further chunks from being potential targets. In order to do so in a manner that is comprehensible to the reader (and also pertinent to the domain in which we originally worked) the concept of 'conversational diagnosis' was introduced.

3. 6. Conversational diagnosis

The use of case studies is a common feature in the design of learning environments and this is taken as the underlying model for conversational diagnosis. However, the presentation of a case study as a chunk fails to provide the interactivity that would promote learning through exploration of the domain knowledge represented by the chunks. In particular it would fail to allow for exploration of the range of potential solutions to problems and the factors that an expert would consider and reject in reaching the solution unless the process of solving was directly described — reverting to text rather than exploiting hypertext.

The basis for conversational diagnosis is modelled on a game-theoretic approach to diagnosis. Two or more experts hold a conversation describing cases that they have known. The aim of each expert is to hold the conversation and to demonstrate knowledge within the domain. To control the conversation they each attempt to steer the discussion towards their own cases. Underlying each discussion is a 'most general case', which starts as a case with no details and becomes more and more detailed as the game progresses. At any stage statements within the domain can be classified into three modal types, based on two distinctions. The first distinction is between those that are impossible, and those that remain consistent with the case. A further distinction can be made for the consistent statements. The interesting statements can either be those that must be true for any case fitting the current description ('everybody knows that ...') and statements representing inferences that could still apply to some of the cases fitting the description.

The rules of the game require that statements can only be made if they are consistent with this case description (*i. e.* it is possible that the statement could be true about a case that fits the current description of the 'most general case'). Similarly, if an expert chooses to add details to the case in order to steer the conversation, it can only be done if the new details are consistent with the current description of the case. In this way, as more detail is added to the case, the relevant statements become more restricted and the 'most general case' becomes more and more detailed. The game ends when adding more detail to the case would fail to change the set of statements that can be made.

In the system a mechanism is provided to describe the case through the progressive addition of details. After adding detail to the case the reader can return to browsing the hypertext system with the added constraint that the only chunks that are accessible are those that could be made by an expert in the game. This eliminates all of the information that is irrelevant to the given case. As the game proceeds, the number of statements that can be made reduces, providing the selectivity that is required. The distinction between 'everybody knows that' statements is provided through an additional set of links named as 'details about the case'.

3. 7. Learning in the model

The SOLO model provides four levels of learning within the target range, pre-structured, uni-structured, multi-structured and relational and it is at the latter stage of learning that the system is aimed. Earlier stages could also be covered but were not appropriate to the experiment that was conducted. The assumption made of the readers was that the content was known (pre-structured) and that the relevance of specific aspects were understood (multi-structured) but that the interaction between different elements were not developed. The design of the hypertext system is, therefore, to allow the reader to explore the interaction between different features of cases and the underlying knowledge. Making evident the interactions between different aspects of the knowledge requires that the range of potential associations is made possible as well as the cumulative effect of combined details (local and global selectivity). This general description of the learning process is based on further models, both of conceptual development and of expertise.

There are several different levels at which exploration of the knowledge could be said to lead to conceptual development, depending on the extent to which the structure of the model reflects the conceptual model that is developed by the learner [Tompsett, 1988]. The model proposed here could function at different levels. At a weak level it functions as any hypertext system, providing links to explore a domain of knowledge. A stronger claim would be that the structure of the knowledge facilitates development of a conceptual model in the reader. Even stronger still would be the claim that the structure is a strong model of the mental model that would be developed by a learner. The last of those positions would reflect a Vygotskian viewpoint if the semantic model proposed acts as a sound model of expert understanding. Even though such an underlying assumption is untestable, since mental models are impossible to inspect, the Vygotskian model has been taken as an underlying model for the work.

In this way, basing global selectivity on game theory designed around the behaviour of experts, can be seen as a coherent extension of the hypertext system, though the extent to which the reader can 'participate' in discussions with the expert are clearly restricted. However, the model of diagnosis behaves in many ways like real experts, for example, there is no guarantee that there is a single solution — the diagnosis process proceeds until no further information can be eliminated.

Evaluation of the model would evidently require comparative analysis of different models of hypertext [Jonassen, 1993] and some measure of the ability to make the correct semantic associations [Schvaneveldt *et al.*, 1985]. The context of the experiment, however, required that the current system was evaluated with real learners and independently of the system developers.

4. The learning experiment

The current system was developed to provide professional training to community pharmacists [Briggs *et al.*, 1993]. Professional updating normally occurs through seminars (usually arranged in large towns), professional journals, distance learning packages and commercial leaflets. Seminars attract a well defined group (approximately 40% within the catchment areas of the trial group [Trill, Dyke, 1991, p. 44]) but their location, coupled with a legal requirement that a qualified pharmacist is in the pharmacy at all times restricts access. The UK Department of Health, through the College of Pharmacy Practice who are responsible for post-graduate pharmacy education, funded a project to determine if knowledge-based systems would provide a mechanism for complementing the existing provision of professional development and exploiting the computers that are available within the pharmacies. Although the brief appears amenable to a straightforward expert system approach, we successfully argued that such approaches fail to provide an effective learning environment.

4. 1. The 'over the counter drugs' trial

The requirement was to produce a demonstration system for professional updating within the field of 'over the counter drugs'. As a demonstration domain this is less dangerous than those that are available on prescription and, following an unpublished government study, had been identified as an area in which accuracy could be improved. The project was to lead to an initial prototype, to be tested in a limited first trial, leading then to an extended demonstrator trialled with 100 users. As system designers we were excluded from management and evaluation of the trials and were only allowed to provide user support should problems arise.

4. 1. 1. — The subject matter

Two experts were appointed to provide knowledge and guidance within the domain, Dr. Alison Blenkinsopp and Dr. John Purves, at Bradford University. The sub-domain of 'upper respiratory tract symptoms' was identified as representative of the complexity of decision taking in the field. Any further extension of the domain would lead to an increase in volume within the system without providing a further test of the model. Since the intention had been to retain features of normal script an existing corpus of material was identified (extracted from [Blenkinsopp, Paxton, 1989]). This was already conveniently 'chunked' into a paragraph structure appropriate to the system constraints.

4. 1. 2. — Characteristics of the user group

The selection of community pharmacists as the target group provided some interesting challenges and constraints. Age range varies from 25 to 65 which implies a range of initial professional development from apprenticeship to graduate entry followed by supervised practice. Pharmacists may be individual owners working by themselves, one of a group working in a larger pharmacy, or working in a major chain of pharmacies. In addition to the pharmacists themselves, the intention was that less qualified assistants should also be able to learn from the system, although their background knowledge was considerably less.

4. 1. 3. — Characteristics of the original system

The 'PCs' available ranged from monochrome CGA CP/M wordprocessors, used for labelling of drugs, to PC-486 machines with super-VGA and large hard disks, used for patient records and drug databases. This wide range required the minimum specification for a working system to be set as an AT equivalent, with a minimum of 1M ram, monochrome CGA screen without a pointing device.

This restriction placed significant constraints on the design of the interface. Anchors within the chunks were not possible and thus were compressed into a single, two stage menu. Definition of the patient to enable global constraints on browsing was provided by a further menu, although this introduced an unnatural 'hitch' when switching back to browsing. Further menu options allowed for more conventional hypertext operations, such as forwarding through the selection of chunks, backtracking and book-marking. Without colour, it was not possible to provide a direct indication of which links had been followed and which had not. To circumvent the problem of circular exploration of a subdomain without effecting a complete exploration an alternative strategy, based on short term memory, was adopted, removing recently visited chunks from becoming targets. The number was kept to $7^{+}/-2$, those which could be recalled and quickly be revisited by backtracking.

In addition to the semantic browsing, an index was provided to allow direct entry at specific points within the hypertext system. This was generated from the specialisation hierarchies but restricted to a small number of entry points to encourage browsing as much as possible.

The limitation of screen size and definition forced us to adopt a chunk size of approximately 14 lines of 50 characters. This matched well with the paragraph size of the primary source text, but was smaller than the paragraph size of some of the more academic texts that became available during the work. A prototype, comprising only 96 chunks, based on the symptom of 'cough' was conducted during January/February 1991 with 10 practising pharmacists each of whom was familiar with computers. This validated the approach to link construction and the sufficiency, with minor changes, of the interface design. Perhaps unsurprisingly, the limitations of the interface design from our own point of view, was less noticeable to the pharmacists who, working on their own machines, did not have such high expectations as we did ourselves.

4. 2. Example of knowledge structuring

An example of the knowledge structuring for one particular card is shown below (**Figure 2** and **Figure 3**), together with the knowledge as it is represented. The primary concept of the card is defined from the roles in the semantic representation of the individual statements in the chunk (see **Figure 3**, keywords are emphasized) :

Condition Croup
This usually occurs in infants, with a peak incidence of 18 months to 3 years. The cough has a harsh barking quality. It develops a day or so after the onset of cold-like symptoms. It is often associated with difficulty in breathing, and an inspiratory stridor (noise in throat on breathing in). GP referral is necessary.

Figure 2
Typical chunk content

typically if condition **is** croup **then**
'patient restrictions' **includes** [age **of** patient **is** 'young child'].

if condition **is** croup **then**
nature **of** cough **is** 'harsh and barking'.

typically if condition **is** croup **then**
'associated symptoms' **includes** 'difficulty in breathing'.

typically if condition **is** croup **then**
'associated symptoms' **includes** 'inspiratory stridor'.

if condition **is** croup **then**
advice **includes** 'GP referral'.

Figure 3
Semantic representation of chunk content

In this case, 'patient restrictions' and 'associated symptoms' are roles of a condition, as are 'site of action' (an anatomical structure), 'cause' (a stimulus), incidence, prognosis and 'potential management' ; 'advice' is a sub-role of 'potential management'. This combination of roles indicates that 'conditions' is the context of the chunk, specialised as 'croup'. The lattice of 'conditions' that includes croup, generated from other chunks, is shown in **Figure 4.** Links will automatically be generated to higher level generalisations (in this case : 'acute conditions' 'upper respiratory tract conditions' and 'respiratory tract conditions'. In addition, links will be created to other 'acute conditions' and other 'upper respiratory tract conditions'.

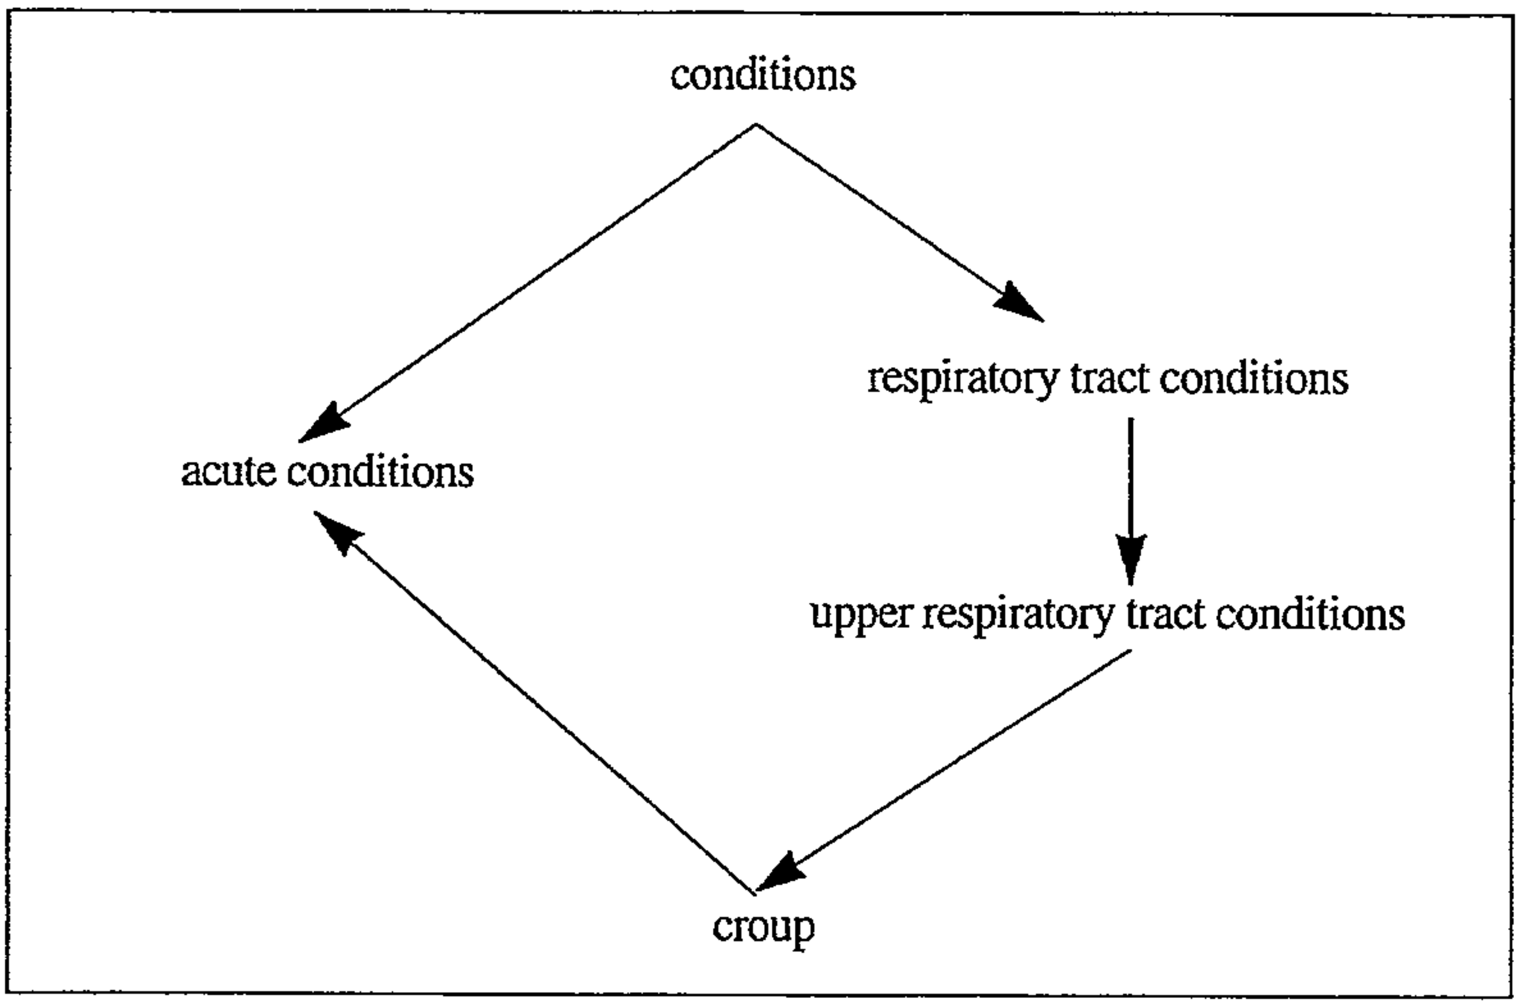

Figure 4
Part of the 'conditions' hierarchy

The two different link names are retained in the first stage of target selection, even though the target sets are not mutually exclusive sets. Use of the names of the typed links, and the names of the possible targets provides both knowledge about the domain and reinforcement of the potential context within which the reader will expect consistency.

The use of keywords such as 'typically', was introduced in recognition of the frequent use of such modal words (also generally, often) within the chunks. There is no theoretical problem with the introduction of modal constructs within the representation, and inference engine, although choice of a suitable modal system is clearly required. The use of 'every-one knows' provided an approximation to the required modality.

In addition to representation of the statements, this chunk was also described as "introducing 'croup'" as more detailed information was available on other chunks.

4. 3. Changes from first prototype

Only three significant issues that arose from use of the prototype system related to the use of language within the chunks, the alacrity with which obvious conclusions were reached and the scope of the knowledge.

The first point identified the importance of uniformity of professional language across the system. It was not difficult to make such changes as they did not have any impact on the link structure but rather highlighted the difficulty of using the same text chunks with readers of different professional standing (for example, excluding some shop assistants).

The second point identified a few simple conclusions that can rapidly be deduced, in particular the need to seek further professional advice from general practitioners. Since these conclusions were well known, the pharmacists were irritated by them being repeated. Although they accepted that such points could not be omitted, it was the potential prognoses underlying the conclusions that was considered of higher significance. As a consequence, the presentation of the 'everybody knows that' conclusions were re-sequenced, to provide the 'advice' role of a case last. Although this would be simple to effect through sequencing of the role definitions, it is not that such use of a sequence, imposed by the linear nature of the definition, should be used to effect this [Clancey, 1984].

The third issue raised by the trial group concerned the additional knowledge that could be introduced. As most of the treatments are not particularly effective, some of the clients have already exhausted conventional treatments and are interested in non-standard effects, such as acupuncture or homeopathic treatments. Extension to cover these treatments, although possible within the system, would not have kept within the remit of the funding agency.

Apart from the changes detailed above, the other change was to extend the domain to cover the full range of upper respiratory tract conditions, including, colds, asthma and allergies etc.

System documentation provided an introduction to the concept of hypertext and described the way in which the interface worked. Learning was supported through the use of a set of case studies, modelled on a format already available in print. These short case studies, typically one page long, introduce a particular case and discuss the issues that a professional might consider in reaching a decision as to the advice to give and possible therapy to recommend. The cases were presented in two ways in relation to the hypertext system. The first set interleaved presentation of case information, with discussion about the case, with guidance as to how the hypertext system could be used to identify the relevant issues. The second set provided a description of a case with a separate analysis of the appropriate advice and therapy. The selection of case studies was made to provide broad coverage of the domain.

4. 4. Results of full trial

The full system was trialled, independently of the system designers, during the period August to November 1991. The context of the project's funding required that the system was evaluated as an integrated form of professional training, with the evaluation being conducted independently [Trill, Dyke, 1991]. Seventy-nine practising pharmacists were recruited by the College of Pharmacy Practice to take part, with 30 completing the two linked questionnaires (A and B) that formed the first stage in the evaluation. Questionnaire A was used as a pre-trial questionnaire to classify the user group and questionnaire B as a post-test to elicit their experiences, perceptions of suitability and accuracy and overall impressions of the learning package.

Of those that did not complete the trial, 24 reported a lack of time, 10 had inappropriate hardware and 8 were excluded as having been in the first trial. This gave a response rate of 40% overall [*ibid.*, p. 7] and 83% excluding those who were unable to participate. Of the 30 who completed the questionnaires, 10 were selected for further in-depth interviews.

Although more general in approach than we would have conducted ourselves, the trial and evaluation allowed us to work with the target group of professionals who would not normally take part in such evaluations. The results from the trial of interest here are those that relate to the professional acceptability of the system and those that provide feedback about the validity of the semantic links and selection mechanisms that were generated.

In terms of overall impression, 22 reported being enthusiastic with the system, with 4 being unsure and 4 unenthusiastic [*ibid.*, p. 8]. This support was further corroborated despite the limited scope of the system, with 13 finding the material "intellectually stimulating" [*ibid.*]. In general discussion of the benefits, the evaluators highlighted the following :

" • Provided a valuable tool for community pharmacists
- Provided a useful tool for updating essential knowledge
- Assisted community pharmacists to adopt a structured way of dealing with customer symptoms
- Will encourage community pharmacists to actively continue their own personal development" [*ibid.*].

and, in response to the question "Which of the following do you think are assisted by the use of the package" :

" • updating knowledge about the underlying disease states
- understanding the significance of symptoms
- recognising dangerous symptoms
- taking the customer's history
- knowing more about the causes of symptoms
- knowing whether to refer the customer to the GP" [*ibid.*].

In summarising these aspects the reviewers note :

> "Once again, the development of a structured and more comprehensive questioning technique was found to be highly rated thus indicating a change in behaviour, even if this is not described as a change in approach" [*ibid.*].

Although the external evaluation does not provide a direct comparison with alternative approaches, several aspects of the evaluation provide vindication for the underlying design. The comments made by the pharmacists, and supported by the in-depth interviews, suggest that the browsing mechanism generated links that were relevant and comprehensible and no reports were made of being 'lost in hypertext'. Interestingly, the various aspects of system design, the browsing mechanism, index and selection based on 'conversational browsing' were used to a different extent by different users, as were the case studies. The flexibility of access and fluent transfer between them might be seen as more critical than any single design feature.

Overall, the positive comments made by the pharmacists provide exactly the reflection of the knowledge and its interconnections that we would wish to develop through the use of semantic links, developing, in particular, the relational level in the SOLO taxonomy — an understanding of the inter-relationships between associated areas of knowledge.

4. 5. Extension and re-release

This trial system is shortly to be released in an updated version for an open trial within the UK and is being made available at cost to pharmacists who wish to use it. The requirements for release were that the knowledge should be brought up to date and that some more specific details of particular products (as opposed to the possibly active ingredients) should be included.

Bringing the text up to date requires both revision of the text and consequently revision of the links within the system. In this particular instance a more detailed classification of one sub-class of drugs was introduced, which entailed a revision of the underlying knowledge structure and hence revision of some links that already existed in addition to those directly implied by changes in the chunks. The design of the system makes such changes relatively simple to effect.

The inclusion of a new area of knowledge introduced an aspect which had so far been excluded from consideration. The introduction of named products was effected through the introduction of a new role for a particular drug. The system design allowed for critical information about the drug to be missed whilst locating relevant products. Within a decision support system, such system behaviour causes no problems, but, from the point of view of learning, this potentially allows the learner to miss the critical information they need to acquire. This issue for learning was discussed during the original system design and was the reason for exclusion of such information during the first trial. For the re-release, the information has now been introduced and constraints have been introduced, for this group of learners, to restrict certain classes of links.

5. Learning in hypertext

The nature of learning from a hypertext system will vary, depending on the amount of knowledge that the learner has when they start to use the system. The context within which this trial took place — professional updating — was intended to emphasize the knowledge-based links within a well defined corpus of knowledge and not to impart the details of the individual chunks. In the original design, the learner was unconstrained in browsing around the network, with modification of the network restricted only through the use of a global context to reduce the size of the available network. In other contexts, such freedom might be detrimental to acquisition of the underlying knowledge (SOLO levels 1 to 3), and different user models might be more appropriate.

5. 1. Locus of control

This suggests that the control of exploration should be separated out from the underlying semantic architecture, with the currently defined set of semantic links considered as a maximal set. At one extreme level of control we could be totally un-restrictive, as in the full trial system ; at the other extreme we could restrict links to a single anchor per chunk — and impose an implicitly linear structure on the system. Our intention in extending this work is to exploit the conceptual structure and labelling of links to allow for alternative approaches to control. The limited use that has been made of this in the re-release version and the potential of additional features below, suggest that there is scope for considerable further exploration.

The control mechanisms proposed can be treated either as static — *i. e.* additional constraints on the network that exclude certain links, or, more generally dynamic — constraints on exploration based on the previous experience and behaviour of the learner. For example, the constraint introduced in the re-release, to disallow the learner from reaching a product without having encountered the underlying drug information reflected a view of learning that led to omission of product information from the original trial version. The more likely situation, in which dynamic control is exerted, is in monitoring the way in which a learner initially explores a given hypertext corpus.

5. 2. Evaluating sequences of links and nodes

As a learner browses through this hypertext system, a trail is left behind of both the chunks that are 'visited' and the {concept, labelled link} pairs that define the association between these chunks. The most effective patterns of exploration will depend not only on the stage of learning but also on individual learner styles. We contend that monitoring of this latter sequence will provide a model of the learner's preferred style and, together with limited typing of the function of the chunk, the opportunity to facilitate systematic exploration of the corpus of knowledge.

The intention is to reduce the complexity of the domain. A simple example of control, which would still allow for some choice in learning style to be available in the early stages of exploration, would be to limit the link types available to specialisation or generalisation links within a single concept hierarchy (uni-structured exploration) and similarly restricting the targets that can be accessed. Another simple measure of complexity for a specific KL-One representation is based on the longest sequence of role links required to link a currently available set of concepts within different hierarchies. At any given time this will be less than the longest sequence from the maximal concept (the root of the intra-conceptual hierarchy defined by the role links) to any other concept within the representation. Exploration can then be managed through restriction of the role links available, managing the context and hence complexity of the problem that the learner encounters. Indeed a combination of the two is inherent in the global selection process that underlies the 'conversational diagnosis' approach.

5. 3. Expanding focus

The original design of 'conversational diagnosis' was to provide selectivity through focusing in on particular sub-sets of the domain, discarding irrelevant chunks. The reduction in the size of the available corpus promotes focused browsing of more detailed information and assists complete coverage of relevant information. Since managing complexity has already been identified as a potential aid for systematic exploration, an alternative approach would be to systematically extend the information available as the detail provided about a 'case' brings more directly relevant detailed information. This new approach would potentially allow a relatively fixed set of cards to be available at all times within a total set of chunks perhaps two orders of magnitude larger.

5. 4. Making implicit decisions explicit

Within the current system, a relatively weak model of reasoning has sufficed for representing information in the domain, in part because we are only required to eliminate chunks rather than individual statements. The particular logic that is modelled is clearly dependent on the underlying domain, and other logics may be appropriate in other domains. One plausible extension to increase the sequencing of targets, would be to measure the relevance of a target chunk by the degree to which the set of statements on the chunk is possibly valid (given the current 'case definition'). Our present system would therefore be recognised as a 'fail-safe' selection approach, rather than an optimal one.

6. Conclusion

The current system has created a hypertext system for professional updating. Semantic links have been created through a knowledge-based analysis of the content of the chunks, providing specific labels to ensure that the reader selects an appropriate avenue for further exploration of the domain. This system has been independently trialled and evaluated with successful results as a learning system. This system is currently awaiting re-release after a further update of the content. The scope of this system both in terms of content and interface was restricted, but provided an effective context for assessment of the underlying design principles. The use of semantically typed links provides an opportunity for further support for learning to be provided without unnecessary restriction on systematic exploration of a domain of knowledge.

References

BIGGS (J. B.), COLLIS (K. F.)
1982, *Evaluating the Quality of Learning : The SOLO Taxonomy (Structure of the Observed Learning Outcome),* New York, Academic press.

BLENKINSOPP (A.), PAXTON (P.)
1989, *Symptoms in the Pharmacy,* Oxford, Blackwell Scientific.

BOBROW (W. G.), WINOGRAD (T.)
1977, "An Overview of KRL", *Cognitive Science,* 1, p. 3-46.

BRACHMAN (R. J.)
1979, "On the Epistemological Status of Semantic Networks", *in Associative Networks : Representation and Use of Knowledge by Computers,* N. V. Findler, ed., London, Academic Press.

BRATKO (I.), MOZETIC (I.), LAVRAC (N.)
1989, "KARDIO : A Study in Deep and Qualitative Knowledge for Expert Systems, London, MIT Press.

BRIGGS (J.), OATES (N.), TOMPSETT (C.)
1993, *Using Rules to Guide Learners through a Medical Database,* 20 (1).

BRUZA (P. D.)
1990, "Hyperindices : A Novey Aid for Searching in Hypermedia", p. 120-122, *in Hypertext : Concepts, Systems and Applications,* A. Rizk, N. Streitz, J. Andre, eds., Cambridge University Press.

CHELNOKOV (V. M.), ZEPHYROVA (V. L.)
1993, "Coherent Navigation in Hypertext Environments : The SMIsC Conception", *in Human-computer Interaction* (3rd international conference ; selected papers), L. J. Bass, J. Gornostaev, C. Unger, eds., Berlin, Springer-Verlag.

CLANCEY (J.)
1984, "Use of MYCINs Rules for Tutoring", p. 464-292, *in Rule-based Expert Systems,* B. G. Buchanan, E. H. Shortliffe, eds., Reading (Mass.), Addison-Wesley.

CLANCEY (W.)
1983, "The Epistemology of a Rule-based Expert System : A Framework for Explanation", *Artificial Intelligence,* 20 (2), 15 ff.

GÖDEL (K.)
1931, "Über formale unentschiedbare Sätze der Principia Mathematica und verwandter Systeme", *Monatshäfte für Mathematik und Physik,* 37, p. 349-360.

HALASZ (F.), SCHWARTZ (M.)
1994, *The Dexter Hypertext,* 37 (2), p. 30-39.

HAYES (P.)
1981, "The Logic of Frames", p. 451-458, *in Readings in Artificial Intelligence,* B. L. Webber, N. J. Nilsson, eds., Palo Alto (Ca.), Tioga Press.

JONASSEN (D. H.)
1993, "Effects of Semantically Structured Hypertext Bases on Users' Knowledge Structures", *in Hypertext : A Psychological Perspective,* C. McKnight, A. Dillon, J. Richardson, eds., Hemel Hempstead, Ellis Horwood.

LANDOW (G. P.)
1987, "Relationally Encoded Links and the Rhetoric of Hypertext", p. 331-344, *in Hypertext '87,* 1, J. B. Smith, F. Halasz, eds., Chapel Hill (N. Carolina), ACM Press.

MARSHALL (C. C.), HALSZ, (F. G.), ROGERS (R. A.) *et al.*
1991, "Aquanet : A Hypertext Tool to Hold your Knowledge in Place", *in Third ACM Conference on Hypertext,* San Antonio, Texas.

MARSHALL (C. C.), IRISH (P. M.)
1989, "Guided Tours and ON-Line Presentations : How Authors Make Existing Hypertext Intelligible for Readers", p. 15-26, *in Hypertext '89,* 1, F. Halasz, ed., Pittsburgh (Pennsylvania), ACM Press.

MINSKY (M. L.)
1975, *A Framework for Representing Knowledge,* New York, McGraw-Hill.

POPSELOV (D.)
1979, "Semiotic Models in Psychology and AI Systems", *in Human and Artificial Intelligence,* F. Klix, eds., North-Holland.

REMDE (J. R.), GOMEZ (L. M.), LANDAUER (T. K.)
1987, "Superbook : An Automated Approach Tool for Information Exploration-Hypertext", *in Hypertext '87,* Chapel Hill (N. Carolina), ACM Press.

SCHVANEVELDT (R. W.), DURSO (F. T.), GOLDSMITH (J. E.) *et al.*
1985, "Measuring the Structure of Expertise", *International Journal of Man-Machine Studies,* 23, p. 699-728.

TOMPSETT (C.)
1991, "Contextual Browsing in a Hypermedia Environment", *in NATO Invited International Workshop on Interactive Multimedia,* M. Giardina, ed., Quebec, New York, Springer Verlag.

1988, *Education Training and Knowledge-Base Design* (11).

TRILL (J.), DYKE (S.)
1991, *Computer Assisted Learning for Community Pharmacists : Second Trial-Evaluation Proceedings,* John Trill Associates.