APPLICATION OF SPREADING ACTIVATION TECHNIQUES IN INFORMATION RETRIEVAL

F. Crestani
Dipartimento di Elettronica e Informatica, Universitá di Padova, Padova, Italy.

Abstract.
This paper surveys the use of Spreading Activation techniques on Semantic Networks in Associative Information Retrieval. The major Spreading Activation models are presented and their applications to IR is surveyed. A number of works in this area are critically analyzed in order to study the relevance of Spreading Activation for associative IR.

Key words: spreading activation, information storage and retrieval, semantic networks, associative information retrieval, information processing, knowledge representation.

1. Introduction

This paper reviews the applications of Spreading Activation (SA) techniques on Semantic Networks in a very active research area: Information Retrieval (IR). The general motivation of this paper springs from past work in *associative retrieval*. The idea behind this form of information retrieval is that it is possible to retrieve relevant information by retrieving information that is "associated" with some information the user already retrieved and that is know it to be relevant. The associations between information can either be static and already existing at the time of the query session or dynamic and determined at run time. In the first case, associations among information items (document or parts of documents, extracted terms, index terms, concepts, etc.) are created before the query session, and they make use of semantic relationships between these items, such as for example thesaurus-like relationships among index terms, bibliographic citations among documents, or statistical similarity among documents or terms. In the last case, instead, the system determines associations between information items through interaction with the user, for example by retrieving documents that are similar to documents the user points out to be relevant (this particular technique is called "relevance feedback", as it it will explained in Section 2). Both these techniques have been explored for quite sometime by the IR community and there are various examples of working systems using them. In this paper we will concentrate on the first technique, which is the one most commonly called Associative Retrieval.

In Associative Retrieval associations among information items are often represented as a network, where information items are represented by nodes and associations by links connecting nodes. The heuristic rule consisting in retrieving information associated to information already assessed as relevant is often implemented by means of a technique called Spreading Activation. The purpose of this paper is to describe the different types of Spreading Activation used in the context of associative IR, and to report on past experience in developing and evaluating IR systems based on such technique. At present there is a decrease of interest in this area of research. This is mainly due to the fact that it is very time consuming to set up a network of associations among information items when the size of the document collection is very large. There is a current tendency in IR to experiment with larger and larger document collections. Many papers at the last ACM SIGIR Conferences (e.g. (Voorhees, 1994; Fujii and Croft, 1993)) and the TREC initiative (Harman, 1993) are evident examples of this trend. Most of the original work in associative IR was performed with small document collections, and often the associations among the information items were set up manually or semi-automatically. This, of course, becomes impossible when the document collection is very large. However, nowadays more and more computing power is becoming available and its cost is rapidly decreasing, making it possible to set up automatically associative networks for large document collections. We believe that this research area could return to be very active and that there is a lot to be learned from past experience. This survey paper tries to report about this past experience in a complete and critical way.

Another more personal motivation for the paper is related to the author's most recent research which brought him to develop a conceptual model for Associative IR (Crestani and van Rijsbergen, 1993). The model is based on a three levels network structure, and it is general enough to enable the modelling of any IR application. The model also enables the representation and use of application domain knowledge which could be used for the adaptation of the original user's query to the specific requirements of the application domain. The conceptual model can be implemented using various forms of network processing and three associative processing paradigms were considered: Spreading Activation on a Semantic Network, Neural Networks, and Inference Networks. With the intention of surveying and studying what has already been done in IR using these processing paradigms, this paper focus on Spreading Activation, reviewing this processing paradigm and surveying its applications to associative IR. This is part of the author's research activity ongoing at the Computing Science Department of the University of Glasgow. The main aim of the research is to develop a prototype of a interactive IR system based on a network representation structure, which makes use of application domain knowledge acquired by interactions with users. The SA processing framework and some of its applications to IR are analysed in this paper with this purpose in mind.

The paper is structured as follows: Section 2 briefly explain the IR problem and what has been done so far to solve it, Section 3 explain what a Semantic

Network is, and Section 4 describes the Spreading Activation technique and its variants, which is the most commonly used processing paradigm of Semantic Networks. The core of the paper lies in Section 5 that reports on a number of past experiences in using Spreading Activation in associative IR.

2. The Information Retrieval problem

Information Retrieval (for a good overview see (van Rijsbergen, 1979; Salton, 1989; Frakes and Baeza-Yates, 1992)) is a science that aims to store and allow fast access to a large amount of information. This information can be of any kind: textual, visual, or auditory. An *Information Retrieval System* (IRS) is a computing tool which stores this information to be retrieved for future use. Most actual IR systems store and enable the retrieval of only textual information or documents. To give a clue to the size of the task, it must be noticed that often the collections of documents an IRS has to deal with contain several thousands or even millions of documents.

A user accesses the IRS by submitting a query, the IRS then tries to retrieve all documents that "satisfy" the query. As opposed to database systems, and IRS does not provide an exact answer but produce a ranking of documents that appear to contain some relevant information. Queries and documents are usually expressed in natural language and to be processed by the IRS they are passed through a query and a document processors which breaks them into their constituents words. This process is called *indexing*, and in most modern IRS it is fully automatic. During indexing non-content-bearing words ("the", "but", "and", etc.) are discarded, and suffixes are removed, so that what remains to represent query and documents are two lists of terms. These terms are often weighted using some statistical weighting schema that is meant to capture the importance of a term in representing the document or the query informative content (for an old but still very useful survey on weighting schema see (Robertson and Sparck Jones, 1976)). Document indexing is performed off-line and document representation are stored in a "inverted file", that is a file that reports for each term a weighted list of documents in which the term appear. Query indexing is performed on-line, using the same procedure of document indexing so that the query representation can be compared using some similarity evaluation algorithms with the document representations. Good IR systems typically rank the matched documents so that those most likely to be relevant (those with the higher similarity with the query representation) are presented to the user first. Some retrieved documents will be relevant (with varying degree of relevancy) and some will, unfortunately, be irrelevant. The user appraises those ones that he considers relevant and feeds them through a process called *Relevance Feedback* (RF) which modifies the original query to produce a new improved query and as a consequence a new

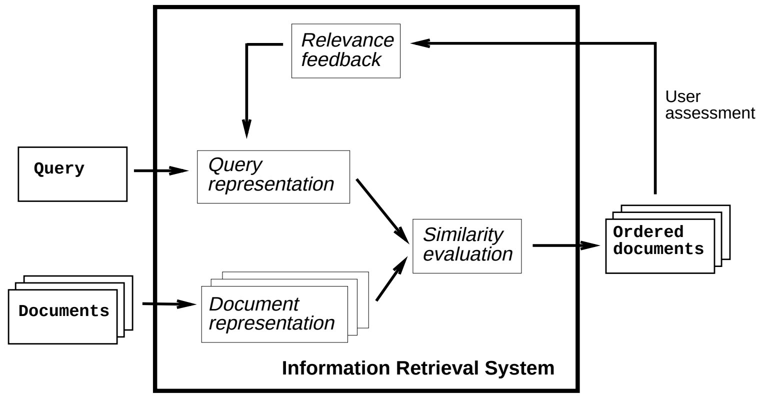

Figure 1. A classical Information Retrieval system.

ranking of documents. If the IR process is interactive this will go on until the user is happy of the resulting list of documents. An example of a such an IRS is depicted in Figure1.

In recent years big efforts have been devoted to the attempt to improve the performance of IR systems and research has explored many different directions trying to use with profits results achieved in other areas. An area which attracts much attention in IR is Artificial Intelligence, so much that a new branch of IR called *Intelligent IR* (IIR) arised from these studies. A few directions of the research in IIR are reported in (Croft, 1987; Jacobs, 1992). Among them, the use of knowledge based techniques in IR has received a particular attention. The aim is to use application domain knowledge in the indexing, in the similarity evaluation, or to enrich the query representation. This last use seems particularly promising because it would enable to retrieve documents indexed with terms not present in the query but relevant to the concepts expressed in the query. The most typical example of use of application domain knowledge in IR is depicted in Figure 2, where knowledge stored in a Knowledge Base is used to expand the original query formulation to include terms related to those originally used by the author.

Among the many formalisms that could be used to represent application domain knowledge for IR applications, Semantic Networks or other network representation structure seem quite promising. The use of a network representation structure makes this approach particularly appealing for associative IR. The most used network knowledge representation structure is a Semantic Network, that uses Spreading Activation as its processing paradigm. In the following two section we will explain what a Semantic Network is and how Spreading Activation on a Semantic Network works.

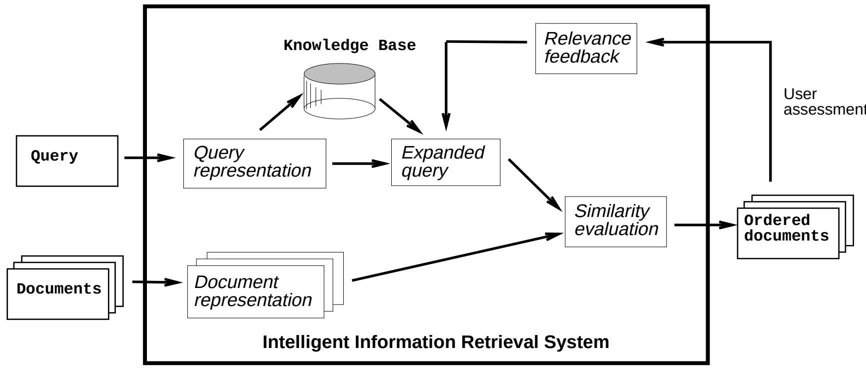

Figure 2. An Intelligent Information Retrieval system.

3. Semantic Networks

Since their introduction by Quillian in (Quillian, 1968), *Semantic Networks* have played a significant role in knowledge representation research. According to Quillian definition, Semantic Networks express knowledge in terms of concepts, their properties, and the hierarchical sub-superclass relationship between concepts. Each concept is represented by a node and the hierarchical relationship between concept is depicted by connecting appropriate concept nodes via "is-a" or "instance-of" links (Schiel, 1989). Nodes at the lowest level denote classes or categories of individuals while nodes at the higher levels denotes classes or categories of individuals. Concepts get more abstract as one moves up the is-a hierarchy. Properties are also represented by nodes, and the fact that a property applies to a concept is represented by connecting the property node and the concept node via an appropriate labelled link. Typically, a property is attached at the highest concept in the conceptual hierarchy to which the property applies, and if a property is attached to a node, it is assumed that it applies to all nodes that are descendants of that node. An example of a Semantic Network is depicted in Figure 3.

The term Semantic Networks has been used in a far more general sense in the knowledge representation literature than what has been described above, and for what concerns IR, researchers have often used the term Semantic Network to refer to an *Associative Network*. This is a generic network of information items in which information items are represented by nodes, and links express sometimes undefined and unlabeled associative relations among information items. In modern IR, where statistical techniques are used in the indexing phase to associate weights to index terms, the relationships among information items are sometimes weighted, so adding to the network a measure of the strength of associations. In Section 5 we will report a few examples of the kind of Associative Networks used in IR.

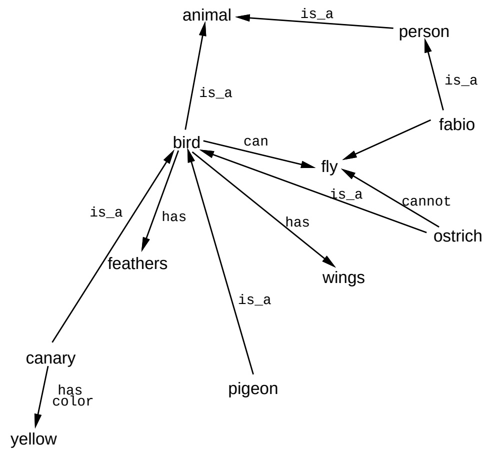

Figure 3. An example of a Semantic Network.

Semantic or Associative Networks are usually processed by means of a technique called Spreading Activation that will be explained in the next section.

4. Associative Retrieval using Spreading Activation

Historically speaking, SA was not the first associative processing paradigm to be used in IR. Studies on *Associative Retrieval* date as early as the 60s and had their origins in statistical studies of associations among terms and/or documents in a collection. The "associative linear retrieval model" is one of these earliest models based on the concept of associative retrieval. This model, in its basic idea (Salton, 1968), consists of expanding the original query using statistically determined term–term, term–document, and document–document associations. This technique is based on the assumption that there exists statistically determinable relations among terms, among documents, and among documents and terms. These associations can be represented in a similarity matrix. Quantitative evaluations of similarity between terms, for example, can be obtained by means of statistical analysis of term co-occurrence in documents. Associations between documents, based on a quantitative evaluation of their respective similarity, can be obtained evaluating similarities in the terms assignments to documents or by means of citations and other bibliographic indicators. There are many heavy assumptions on this model and more recent studies (Preece, 1981; Salton and Buckley, 1988) have lead to the conclusion that effective term expansion methods valid for a variety of different collections are difficult to generate. Moreover, IR systems based on this approach have shown a lack of consistent improvements in the effectiveness. This result can have various motivations. First, the similarities statistically derived from some pairs of documents, or some pairs of terms, may be valid only locally in the particular "environment" (or application domain) from which they are obtained. Second, most practical methods for computing the document associations are based on the assumption that the terms or the documents are originally uncorrelated, i.e. independent of each other. Such assumption is no more accepted in many of the new research directions of IR.

Recently, these models of associative retrieval has been revised using the so-called *Spreading Activation Model*, which is based on supposed mechanisms of human memory operations. Originated from psychological studies (see for example (Rumelhart and Norman, 1983)) it was first introduced in Computing Science in the area of Artificial Intelligence to provide a processing framework for Semantic Networks. Its use has been praised and criticised, but it is currently adopted in many different areas such as: Cognitive Science, Databases, Artificial Intelligence, Psychology, Biology, and lately to IR. The basic SA model has, however, been subject to various enhancements in order to make it more suitable to various application areas and the way it is used in IR is quite different from the original formulation in the area of psychology.

In the following three sections the SA model will be described in depth. Section 4.1 will describe the "pure" SA model, which consists in the sole use of the associative nature of a network representation as a search controlling structure. In Section 4.2 some more search controlling structures will be added in order to give preference to particular associations. Section 4.3 will show how the search controlling structure can be used in a interactive way using external feedback.

4.1. The pure Spreading Activation model

The SA model in its "pure" form is quite simple. It is made up of a conceptually simple processing technique on a network data structure.

The *network data structure* consists of nodes connected by links, as depicted in Figure 4. Nodes model objects or features of objects of the "real world" to be represented. They are usually labelled with the name of the objects they intend to represent. Links model relationships between nodes and they can be labelled and/or weighted. The connectivity pattern reflects the relationships between objects and/or features of objects of the "real world" to be represented. A links usually has a direction, a label, and/or a weight assigned according to a specific direction. This representation structure is very similar to a Semantic Network, but it is more general than the definition of Semantic Network given in Section 3. It could represent a Semantic Network, but also a more generic Associative Network.

The *processing technique* is defined by a sequence of iterations like the one schematically described in Figure 5. Each iteration is followed by another

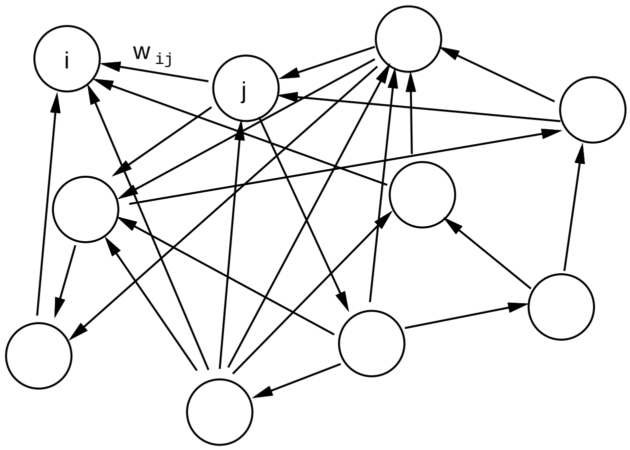

Figure 4. The network structure of a SA model.

iteration until halted by the user or by the triggering of some termination condition. An iteration consists of:

1. one or more pulses;
2. termination check.

What distinguish the pure SA model from other more complex models is the sequence of actions which composes the pulse. A pulse is made up of three phases:

1. preadjustment;
2. Spreading;
3. postadjustment.

In the preadjustment and postadjustement phases, which are optional, some form of activation decay can be applied to the active nodes. These phases are used to avoid retention of activation from previous pulses, enabling to control both activation of single nodes and the overall activation of the network. They implement a form of "loss of interest" in nodes that are not continually activated.

The spreading phase consists on a number of passages of activation weaves from one node to all other nodes connected to it. There are many ways of spreading the activation over a network (for a overview see (Preece, 1981)). In its more simple form, on a single unit level, SA consists first in the computation of the unit input using this formula:

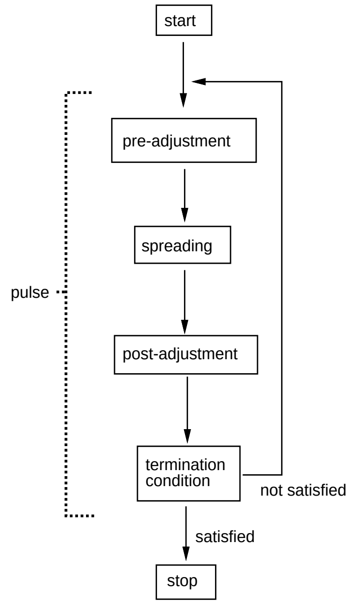

Figure 5. The pure SA model.

$$I_j = \sum_i O_k w_{ij}$$

where:
I_j is the total input of node j;
O_i is the output of unit i connected to node j;
w_{ij} is a weight associated to the link connecting node i to node j.

The input and the weight are usually real numbers, however their numerical type is determined by the specific requirements of the application to be modelled. In particular, they can be binary values (0 or 1), excitatory/inhibitory values (+1 or -1), or they can be real values indicating the strength of the relation between nodes. Usually the first two of these options are used in connection with networks with labelled links like for examples Semantic Networks, where the semantic value of the relation represented by the link determines, in the context of the application, the value to be associated to the link. The last option is mainly used for Associative Networks, where there is only one generic type of association that need to be weighted.

After a node has computed its input value, its output value must be determined. The numerical type of the output of a node is also determined by the requirements of the application. The two most used cases being the binary active/non-active type (0 or 1) and the real value type. In SA models there is

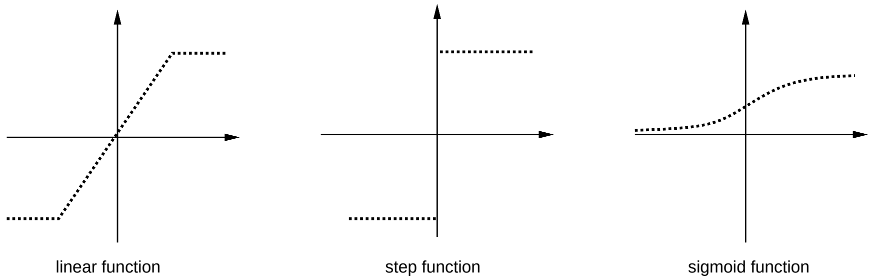

Figure 6. Some commonly used activation functions.

usually no distinction between "activation" or "output" of a unit. The activation level of a unit is its output value. This is usually computed as a function of the input value:

$$O_j = f(I_j)$$

There are many different functions that can be used in the evaluation of the output, some examples are depicted in Figure 6. The most commonly used function in pure SA models is the threshold function. It is used to determine if the node j has to be considered active or not. The application of the threshold function to the above formula in the case of binary value units gives:

$$O_j = \begin{cases} 0 I_j < k_j \\ 1 I_j > k_j \end{cases}$$

where k_j is the threshold value for unit j. The threshold value of the activation function is application dependent and can vary from node to node, therefore the notation k_j for the unit threshold has been used.

After the node has computed its output value, it fires it to all the nodes connected to it, usually sending the same value to each of them.

Pulse after pulse, the activation spreads over the network reaching nodes that are far from the initially activated ones. After a determined number of pulses has been fired a termination condition is checked. If the condition is verified than the SA process stops, otherwise it goes on for another series of pulses. SA is therefore iterative, consisting of a sequence of pulses and termination checks.

The result of the SA process is the activation level of nodes reached at termination time. The interpretation of the level of activation of each node depends on the application and, in particular, on the characteristics of the object being modelled by that node.

4.2. Constrained Spreading Activation

The pure SA model, however, presents some serious drawbacks:

– unless controlled carefully by means of the preadjustment and the postadjustment phases the activation ends up spreading all over the network;

– there is not a complete use of the information provided by the labels associated to the links, that is, there is no use of the semantics of the associations;

– it is difficult to implement some form of inference based on the semantics of associations.

These problems can find a solutions by taking into account in the processing technique the diverse significance of the relations among units. This can be achieved using the information provided by the labels on the links and by processing links in different ways according to their semantics. It is possible in this way to implement some form of heuristics, or to spread activation on the network according to some inference rules. A common way of implementing a processing technique which spreads the activation according to rules, is by means of constraints on the spreading. Here are some constraints commonly used in SA models:

distance constraint:the spreading of activation should cease when it reaches nodes that are far away in terms of links covered to reach them from the initially activated ones. This corresponds to the simple heuristic rule that the strength of the relation between two nodes decreases with their semantic distance. Relations can be classified according to their distance in term of links. Relations between two nodes directly connected are called first order relations. Relations between two nodes connected by means of an intermediate node are called second order relations, and so on. It is common to consider only first, second and, at most, third order relations, however this is application dependent.

fan-out constraint:the spreading of activation should cease at nodes with very high connectivity, or fan-out, that is at nodes connected to a very large number of other nodes. The purpose of this constraint is to avoid a too wide spreading which could derive from nodes with a very broad semantic meaning and therefore connected to many other nodes.

path constraint:activation should spread using preferential paths, reflecting application dependent inference rules. This can be modelled using the weights on links or, if links are labelled, diverting the activation flow to particular path while stopping it from following other less meaningful paths.

activation constraint:using the threshold function at a single node level, it is possible to control the spreading of the activation on the network. This can be achieved by changing the threshold value in relation to the total level of activation over the entire network at any single pulse. Moreover, it is possible to assign different threshold levels to each unit or set of units in relation to their meaning in the context of the application. Although this may cause a increase in the number of computations, it makes possible to implement various complex inference rules.

Referring to Figure 5 these constraints can be seen as acting during the preadjustment phase (this is the case for distance, fan-out, and path constraints) or during the postadjustment phase (mainly for activation level constraints). Therefore they can be considered as an enhancement of the pure SA model.

Another more practical advantage deriving from the fact that the activation does not spread over the entire network is that it permits a reduction of the computational effort of SA, because only a small portion of the units become active and send activation to other units.

A very good example of the use of constrained SA in IR is reported in (Kjeldsen and Cohen, 1987; Cohen and Kjeldsen, 1987). We will describe in more detail this model in Section 5.2, talking about a system called GRANT.

4.3. Spreading Activation with feedback

A further enhancement of the pure SA model can be obtained by means of feedback from an external source. In this case, an external evaluation of the activation level of some units or of the entire network provides some form of constraint that would be difficult or impossible to implement in the form of automatic rules. This external feedback can arrive from another process or can be provided by the user of the system. The user evaluates the activation level reached by some nodes and modify it according to his requirements. This may result in a following spreading of activation adjusted by the user feedback. Moreover, it is also possible to enable the user to indicate some particular spreading path so that activation can follows directions given by the user which can differ from those specified by path constraints. From this point of view SA can adapt itself to the specific user's needs.

This model is particularly useful in application where there would be too many inference rules to be represented in the form of constraints and where it is necessary to provide an external control by means of a user's evaluation of the results achieved by the SA. In the context of IR, it would be possible to use this model with user provided relevance feedback. The use of feedback from user can either be made in the preadjustment phase, so that the user directs the spreading of activation of the pulse, or in the postadjustment phase, enabling the user to evaluate the result of the spreading and direct the following pulse accordingly. Some examples of the use of this model are reported in the next Section.

5. The use of Spreading Activation in Information Retrieval

SA techniques used in IR are based on the existence of maps specifying relations between terms or between documents, as the case may be. Nodes correspond to terms, documents, articles, journals, subject classifications, authors, and so forth. There is no homogeneity in the network. A node can represent anything. Links indicate the association of a node with another node, as, for instance, an author with a document he/she wrote or a document with a document it cites. An example of a fragment of a document collection representation is shown in Figure 7. Specific link types include term occurrence, document publication, term assignment by indexing, document authorship, document assignment to classification, document citation, and so forth. The set of node types and link types is determined by the data available and by the purpose of the application. The representation structure is therefore application specific and the same structure cannot be applied to different applications. Some examples of differing network representations will be shown in the following of this section. It is also important to note that relationships could actually be expressed by a pairs of links. Authorship, for instance, can be represented by both "authored by" and "is author of" links. Both links in such pairs connect the same two nodes, but their source and destination roles are reversed. Specific processing rules may inhibit activation in either directions, use them in different ways, or associate different weights with the different directions.

Given such a representation structure the network activation starts by placing a specified activation level at some starting node. These nodes are usually identified by the initial query formulation or by documents or terms retrieved by an earlier search operation, the second option being often used in systems with relevance feedback. The activation first reaches nodes located closest to the starting nodes, then spreads through the network using links. The activation level of a node is computed using one of the functions specified above in Section 4.1. The process ends when some termination condition is reached. The activation level of documents at the end of the spreading process is used to compute the relevance level of each document.

Most of the SA techniques used in IR systems differs from the pure SA models in several respects:

– the activation level of a node reached by the spreading of activation is determined by the starting activation level and the type of nodes and links traversed before reaching it;

– distance constraints are usually imposed by stopping or degrading the activity at some specified distance from the original node;

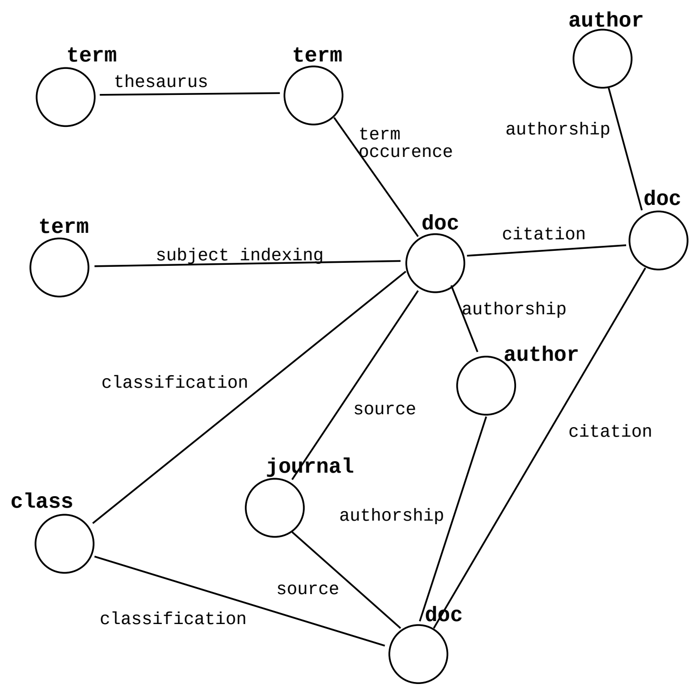

Figure 7. An example of document collection network representation.

– nodes with a large branching ratio (fan-out), that is, nodes connected to many other nodes, may receive special treatment in the activation process in order to avoid or boost a large spreading of activation on the network;

– the activation process follows specified rules, that try to mimic some sort of inference in the process of associative retrieval.

It has been demonstrated in (Preece, 1981) that better retrieval results can be obtained by a SA process which uses some of the above mentioned characteristics.

Much of the effectiveness of the process is, however, crucially dependent on the availability of a representative network. The problem of building a network which effectively represents the useful relations (in terms of the IR's aims) has always been the critical point of many of the attempts to use SA in IR. These networks are very difficult to build, to maintain and keep up to date. Their construction requires in depth application domain knowledge that only experts in the application domain can provide. Furthermore, their construction is a very expensive and time consuming process, that results almost impossible for large collections and/or collections spanning over a large application domain. This problem, often present in applications of SA to

IR, is the main reason for the increasing interest in techniques for automatically build network representations and, in particular, for the application of machine learning techniques in IR. In this paper we will not address this problem. The application of machine learning techniques to IR is a large area of work where very interesting results are being achieved, and limitation in space do not allow us to even briefly report about such work. However in Section 5.7 we will report about a few experiments aimed at building automatically, via statistical techniques, the necessary network representation upon which SA techniques can be used.

A complete SA system that makes use of diverse link types and of spreading rules with distance and fan-out constraints, has never been implemented with ordinary document collections. Various prototype systems are presented in the IR literature but no commercial system based on a SA model actually exists. Some of the characteristics of these prototypes will be briefly presented in the flowing part of this section. The purpose of this review is not to report in details about such prototypes, but only to show advantages and problems related to the practical application of SA to IR.

5.1. The first steps: Preece's and Shoval's works

S.E. Preece's work reported in details in (Preece, 1981) can be considered one of the first attempts to use associative search by SA in IR. In his PhD thesis, he examined in depth the SA approach to associative retrieval. He argued that most of the classical approaches to IR could be explained in terms of different SA processing techniques on a network representation of the document collection. This division between data structure and processing technique can be seen as a first attempt to conceptual modelling IR applications. Combining different network data structures with different processing techniques he showed how it is possible to implement the Boolean model, the Vector Space model and use various forms of weighting for associative retrieval. Moreover, he showed how, using relevance feedback, SA can be used for automatic classification and indexing, and for concept building. It is particularly in these two applications that the approach shows much of its potential. Certainly those ideas come from the appearance of the first papers on Neural Networks and it is easy to see the intention of expanding the SA paradigm in that direction. However, the computing power Preece had the possibility to use in the late 70s could not enable him to see the use of SA together with a machine learning paradigm. Indeed, all the experiments he performed to test the model were made using a collection of very small size, which could not be compared with the size of collections in (even then) "real-world" applications. His thesis can be considered a good survey of earlier works in SA. The use of a common formalism in describing different SA techniques over a common data structure put order in many previous attempts. Moreover, from the analysis and comparison of previous works many new ideas came out. Some of this ideas have been used by other researcher much later.

A major drawback of this work stand in the use of manually build Semantic Networks to represent the document collection. This is a drawback common to almost all the other works presented in this section. The effort in representing all the relationships between documents, terms, authors, etc. in a document collection is such that even for very small collections it results un-practicable. Thus, has Preece pointed out in his thesis, an automatic way of construction of this document collection network representation must be adopted.

The work of *P. Shoval* (Shoval, 1981), developed in parallel to Preece's one, is an attempt to implement interactive query expansion using SA on a Semantic Network. Again this approach is quite simple compared to the most recent ones, however we must consider it a seminal work whose directions are still followed by current research. The knowledge base employed by Shoval is a Semantic Network based upon a thesaurus. The link types are the those used by a common thesaurus, expressing hierarchical relationships, synonymous relationships, and general relatedness. Moreover, two more typer of links were added: generator links, which combines source/generic words to multi-words concepts (e.g. "information" and "systems" are linked to "information system"), and model links, which are used to extend the knowledge about concept by linking a word with a component of its meaning (e.g. "business" is linked to "organizational area"). The processing technique is a form of SA with feedback, where the feedback comes from the user via an interactive process. The SA technique is as follows:

1. The system accepts user terms and immediately expand them along the Semantic Network links.

2. The new terms are matched against each others. If an intersection is found it is first verified that the match also involves additional entry terms, and if it does it is suggested to the user who is asked to judge whether or not the suggested term is in the right direction.

3. If the user rejects that term the search in that direction stops, otherwise the term is further expanded, only after marking it so that it cannot be retrieved in the next wave. The strategy is a "best first" one, and only the "front line" terms are kept, since they are supposed to capture the meaning of their parent terms.

4. The process continues as in 2 as long intersections are found.

5. Finally, the system collects the suggested terms and shows them to the user who for each term has the choice of accepting it (and in this case it can be used to pursue that expanding direction), rejecting it (in which case the system will

backtrack to find alternative terms to expand), or asking why it was suggested (in which case the system prints out the path followed to reach it from the starting terms).

This system presents very interesting characteristics, such as the almost automatic construction of the network representation structure, as long a thesaurus is available, and the fact that this knowledge is not domain dependent, since it could be based on a generic thesaurus. These aspects makes this system very powerful even compared with more recent ones. The SA technique, however, is quite simplistic since it requires a constant feedback from the user. Although this could be useful sometimes, in this case too much intervention is required from the user who may not want "to have to" give so much feedback. Shoval proposes also an alternative search strategy that does not require so much feedback from the user, since it assumes that all intersections found at phase 2 are accepted. This however is very similar to a form of pure SA technique and it is likely to produce a large spreading of activation over the entire network, with the effect that the user will have to look through a long list of suggested terms at phase 5.

5.2. A SUCCESSFUL SYSTEM: GRANT

P.R. Cohen and R. Kjeldsen's GRANT system (Kjeldsen and Cohen, 1987; Cohen and Kjeldsen, 1987) is one of the first systems to use constrained SA in IR. In GRANT, knowledge about research proposal and potential funding agencies is organised using a Semantic Network. Research topics and agencies are connected using a wide variety of association links to form a dense network. A query expresses one or more research topics, or one or more funding agencies. The search is carried out by constrained SA on the network, making full usage of almost all the types of constraints described in Section 4.2. In particular, it makes large usage of path constraints in the form of "path endorsement".

From an heuristic point of view, GRANT can be considered as an inference system that applies repeatedly a single inference schema:

$$IF\ x\ AND\ R(x,y)\ \rightarrow\ y$$

where $R(x,y)$ is a path connecting the two nodes x and y which could be of one or more links. In the particular application for which GRANT was designed, i.e. finding founding agencies for research proposals, this is equivalent to an inference rule of the form: "if a founding agency is interested in topic x and there is a relation between topic x and topic y than the founding agency is likely to be interested in the related topic y". The path endorsement process gives preference (positive endorsement) to some paths and it enables to avoid (with a negative endorsement) some misleading paths. The evaluation mechanism of the paths enable to rank the retrieved nodes.

The use of constraints on the spreading of the activation over the network and of rules to "endorse" some particular paths enable the system to achieve very interesting results. In fact, the authors demonstrated that the use of constrained SA for the application they were considering gives reasonable values of recall and precision[1]. This values, for the application under consideration, were found to be better than those provided by simple keyword search. This technique has been demonstrated to be particularly good for "difficult cases", that is for cases that could have been difficult even for a human expert, though sometimes it provided misleading results for "simple cases".

Developing a system like GRANT involves, first of all, a significant amount of knowledge engineering to construct the Semantic Network. This work consists of an in depth analysis of the domain in which the system will operate in order to determine the appropriate concepts and relationships to build in the network, and the preferences to give to paths of activation spreading over it. One of the major limitations of GRANT, from a SA point of view, is the difficulty of adjusting the parameters of the path endorsement. Without a proper tuning of these parameters a considerable high fallout rate[2] has to be expected.

5.3. A TESTBED FOR IR EXPERIMENTATION: I^3R

The declared purpose of *W.B. Croft, T.J. Lucia, J. Crigean and P. Willet* in designing I^3R (among the many papers about this system, see (Croft et al., 1988; Croft et al., 1989)) was mainly devoted to study the possibility of retrieving documents by "plausible inference". The use of SA was therefore only incidental. In I^3R, it was chosen to implement plausible inference as a form of constrained SA, taking GRANT as an example.

I^3R is designed to act as a search intermediary. It accomplishes its task using domain knowledge to refine query descriptions, initiating the appropriate search strategies, assisting the users in evaluating the output, and reformulating queries. In its initial version (Croft and R.H.Thompson, 1987) the domain knowledge was represented using an AND/OR tree of concept frames, while documents were represented by means of single term descriptors. The system used the domain knowledge to infer concepts that are related to those mentioned in the query. The inference mechanism used a form of "propagation of certainty" on the concept frames. In later versions of I^3R the knowledge

[1] These are two well known effectiveness measures in IR. Recall is the proportion of all documents in the collections that are relevant to a query and that are actually retrieved, while precision is the proportion of the retrieved set of documents that is relevant to the query.

[2] The fallout rate is the ratio between number of non relevant documents retrieved and the total number of non relevant documents in the collection. This is another well know performance measure for IR systems.

representation structure was refined to a sort of Semantic Network, of the kind depicted in Figure 8. Looking at that figure, we must remember that it is not necessary to distinguish among concepts, terms, and documents in the network structure. The important fact is that they are nodes and they behave in exactly the same way as the nodes described in Section 4.1.

Several processing techniques have been used on such a representation structure, that could also be used for browsing. In particular in (Croft et al., 1988; Croft et al., 1989) the following specific form of constrained SA was used:

1. the starting points of the SA are the top-ranked documents from a probabilistic search;

2. initially links connecting document nearest neighbours and document citations are used for spreading activation; these links represent the strongest plausible relationships between documents;

3. in the remaining cycles of activation only nearest neighbours links are used; citation relationships are interesting only in relation to the starting documents;

4. weights on links are used in the evaluation of the node's activation level; they are specified as "credibility" values associated to inference rules representing the existence relationships between the two nodes;

5. documents that have been used as part of an activation path are not used again if they are reactivated.

The authors implemented a retrieval paradigm called "multiple sources of evidence" using these constraints on the basic SA model. This paradigm is the central point of the research using I^3R. It springs from the intuition that a document is more likely to be relevant if its relevance is supported by many different clues.

The experimental results showed the possibility of improving the performance of a generic IR system based on the sole use of nearest neighbour and citation information. However, the magnitude of the improvement varies from one collection to another, therefore showing the difficulty of using SA in operational IR systems.

This piece of research can be considered one of the best attempts to combine constrained SA model with the most sounded IR probabilistic techniques. We believe that it is in this direction that the use of SA in IR will find its bets results.

5.4. The evaluation stage: the Cornell University experience

G. Salton and C. Buckley from Cornell University described, in a very interesting article (Salton and Buckley, 1988), an evaluative comparison of some

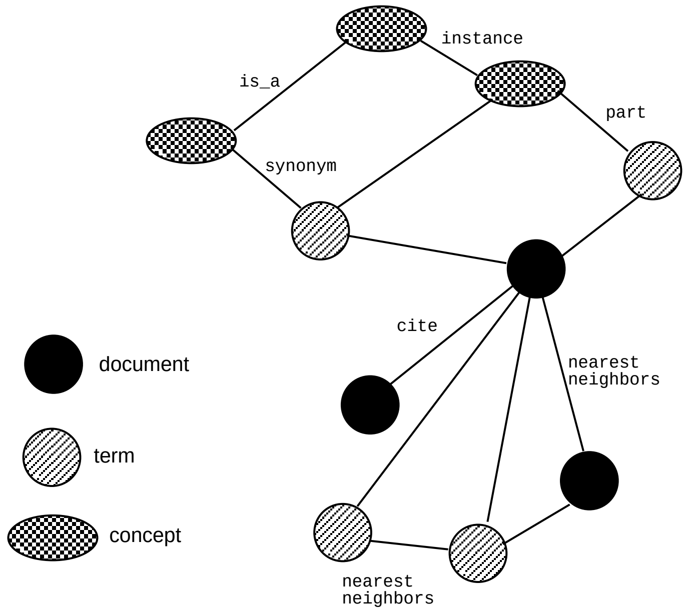

Figure 8. I[3]R network representation structure.

SA models with the vector space model. In particular, the effectiveness of one SA model is evaluated and compared to that of a linear associative model based on vector processing in the vector space model. The SA model used in this evaluation is a refined version of the pure SA model, allowing only two links to be covered from the initially activated node to the final one. The network structure used by the model is also quite simple as can be seen from the example depicted in Figure 9. On such a simple network structure the authors used a pure SA model enhanced by means of a normalisation factor on the activation function. Weights on links are determined using term frequency, which measures the frequency of occurrence of a term in a given document or query. This SA model is compared to the vector space model that benefits from years of experimentations, and where normalisation factors were carefully set in order to obtain the best performance. Moreover, the vector space model makes also use of the inverse document frequency, which measures the relative importance of a term in a collection as an inverse function of the its posting frequency.

The evaluation used several test document collections and were performed using several slightly different weighting schemas. As in most of the papers coming from the Cornell group, the evaluation is very accurate and can be taken as an example of very precise IR system evaluation.

However, despite the big effort in the comparative evaluation of their proposed SA model with the vector processing model, it must be said that

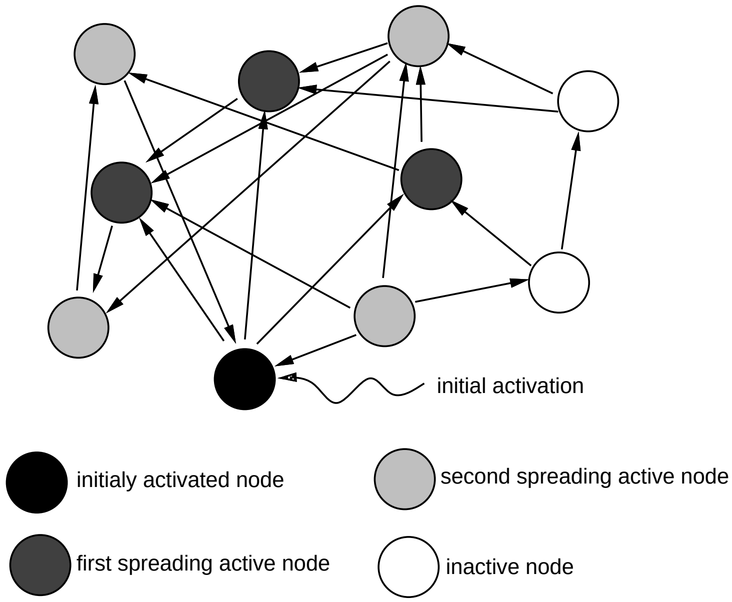

Figure 9. Network used by Salton's evaluations.

the model they used is conceptually very similar to the vector space model and can be considered a very simplified version of SA. It is not surprising that the evaluation were more favourable to the vector processing model than to the SA model. Evaluations were conducted on a ground more favourable to vector processing model than to the SA one, and the conclusions reached by the authors for the particular SA model they used should not be extended to SA models using more complex control structure like those presented in Sections 4.2 and 4.3.

5.5. Spreading Activation and Connectionism

H. Kimoto and T. Iwadera's AIRS (Associated Information Retrieval System) (Kimoto and Iwadera, 1990) is a IR prototype system that incorporates the SA processing technique in a dynamic thesaurus. The basic concept and distinctive feature of AIRS is that it determines the user's interest from user's sample of relevant documents to produce "term information". This term information is used to construct a dynamic thesaurus that generates, at retrieval time, associated keywords.

The dynamic thesaurus is represented as a network where nodes represent terms and links represent semantic relationships between terms. The network is obtained combining information from a static thesaurus with term information. Term information is obtained by ranking terms according to their frequency and location in a set of relevant documents provided by the user. In particular, terms co-occurrence is use to build new links between term nodes in the dynamic thesaurus. The network structure is quite peculiar and it differs from the classical connectionist network architecture. In fact, there are not weights on links but on nodes. These weights are obtained from term information.

The retrieval of documents relevant to a query is performed with a process similar to a constrained SA. The set of terms used in the query is expanded adding associated terms according to their weights and links. Constrains are set in order to limit the distance of the spreading of the activation on the network and only terms with weights over a predefined threshold are used as associated search terms.

Several experiments performed by the authors, although with a very small collection of documents, showed that the performance of their model are better than those obtained using a static thesaurus. However there are a few drawbacks in their approach. The first is a conceptual one: the authors claim they use a connectionist approach. A connectionist approach, based on the use of a Neural Network would be quite different. What distinguish Neural Network from SA (Rumelhart et al., 1986) is mainly the presence of a non linear activation function and, in particular, the presence of a learning procedure that is used to modify the weights on links so that the spreading of the activation over the network reflects some desired pattern. In AIRS there is no activation function for the nodes and there is no learning procedure. If this is to be considered as a connectionist approach to IR, then it is certainly a very poor one. Much research has been devoted to the application of Neural Networks in IR, as reported for example in (Crestani, 1991). Kimoto and Iwadera's approach is certainly much better as a SA one than a connectionist one. Another point is that the model is not much different from classical statistical models, and therefore similar performance to other approaches (like for example the Cornel one) has to be expected. Most approaches based on statistics use co-occurrence information for query expansion. Finally, the complex typology of relations among terms provided by a classical thesaurus is not used at all. All kinds of relations among terms are represented using a single type of connection, resembling more an associative network than a thesaurus.

5.6. THE MARKER PASSING MODEL AND SCISOR

S. Fahlman's NETL (Fahlman et al., 1981) can be considered the first attempt at encoding Semantic Networks as a massively parallel network of simple processing elements. The NETL system consist of a central (serial) computer connected to a large number of node and link elements, each of which is a hardware element. The underlying processing model is the so called *Marker Passing Model*. This model is based on the classical network structure, where a concept is represented as a node, and an association between concepts is represented by a link. A node communicates with another node by propagating a small number of simple messages called "markers" along its link. NETL uses marker passing to perform simple inferences based on set intersection and transitive closure operations. The intersection operation locates items that shares a set of properties whereas the transitive closure operation handles inheritance as well as closure of relations like part-of. These operations are performed in parallel and allows the system to conduct a very fast search.

A limitation of systems based on the Marker Passing Model and of NETL in particular is that communication between network elements is a small number of discrete markers that are essentially boolean conditions and a network element can only detect the presence or absence of a marker in the input. This all or none nature makes these systems incapable of supporting "best match" or "partial match" operations. This is not the place to report the limitations of NETL, that are discussed at length in (Fahlman et al., 1981), but there has been a series of attempts to use modified versions of the Marker Passing Model to IR. In the following we will report about the most interesting ones.

P. Jacobs and L. Rau (Jacobs and Rau, 1993) and in particular Rau (Rau, 1987), attempts to use a form of "constrained" Marker Passing Model, very similar to Constrained SA, to understand and retrieve documents. *Rau's SCISOR* is a question-answering system that partially parse, understand, and answer questions about fact learned from short newspaper stories in the domain of corporate takeovers. The architecture of the system is quite complex, since it is composed of several tools such as TRUMP, a parser and semantic interpreter of the natural language input, KING, a natural language generator for the natural language output, FLUSH, a knowledge acquisition tool used to acquire knowledge from short newspaper stories, and so forth. We will not enter into the details of such a system (there is a number of papers about this system, the most complete one being the above cited one), but we will only analyse the knowledge representation structure and its processing technique.

SCISOR manipulates conceptual structures represented in the KODIAK knowledge representation language. KODIAC can be seen as a hybrid frame and Semantic Network based language. Knowledge can be stored in SCISOR in one of these three forms: specific (or episodic), abstract, or semantic. The distinction between these forms of knowledge, whose definitions are quite obvious from their names, is quite subtle and is almost a continuum. The resulting structure is very similar to a Semantic Network, where these levels of abstraction are merged together to form a single network structure. Upon this network structure another level of organisation is superimposed: groups of related specific or abstract concepts (i.e. nodes) are linked together through a common node, called TAG. This TAG allows the system to detect whether two or more concepts appear in the same event or episode. TAGs are used in the process of retrieval, which is achieved using a “priming” or constrained spreading activation. As concepts of the representation structure are instantiated in the system by a query, instances of these nodes that are related via category membership links are marked (primed or activated) and their TAG activation values are increased. When a certain subset of the concepts in a episode is marked and the TAG activation level exceeds a threshold, the entire episode is put into the system’s short term memory buffer, where a constraint limits the number of episodes that can be stored at the same time. This process correspond to “spontaneous” retrieval. A match filter is then used to refine the retrieved episodes. Note that retrieval occurs as a side effect of understanding, since the instantiation of new concepts causes waves of activation to spread through the network.

Leaving aside the fact that SCISOR is not really an IR system, but mainly a question-answering system, it has a few limitations we would like to point out. An important limitation of SCISOR is that it is not clear what the system can and cannot answer. In fact, SCISOR seems only capable of answering questions about information “explicitly” stored in the knowledge base. Any information that could be reconstructed from information stored in the knowledge base is simply not available, although the content-addressable nature of KODIAC allows some limited form of inference to access “derivable” information. However the line between what is explicitly stored in the knowledge base and what can be figured out is not sharp and a user could be potentially left with the doubt of weather the system does not have the information searched or simply it cannot be retrieved because it cannot be derived. A second limitation of SCISOR is in its complex knowledge representation structure which is quite difficult to build and to keep up to date. Despite the use of a sophisticated NLP tool like TRUMP, there is always the need of manual intervention for encoding information contained in the stories. The complexity of the representation structure makes it necessary to have the encoding performed by an expert, whom will have to constantly check the consistency of the entire knowledge base.

Moldovan too, in (Kim and Moldovan, 1993; Chung and Moldovan, 1994), attempts to use the Marker Passing Model for fast classification and retrieval. However this attempt is only partially concerned with IR, since it is was developed mainly for the task of “message understanding and classification”. Although reporting about this work is outside the purpose of this paper, we considered important to cite his work since it has some similarity with constrained SA.

5.7. Spreading Activation on Hypertexts

An area that is becoming increasingly important is the one related to the use of hypertext systems in IR. The use of hypertexts in IR is directed towards enabling the user to search for document not only by querying a document base, but also by browsing it. However, in oder to provide such a tool to final users it is necessary to build up the hypertext from the document collection, possibly using information already present in the documents. The size of the collections normaly used in IR makes it impossible to build by hand the hypertext representing the collection, therefore a large research effort has been devoted to finding techniques for the automatic construction of hypertext to be used in IR. We will not enter into this subject in this paper, since it is outside our scope, however it is important to note that there are a few research works devoted to the use of SA in automatically build hypertext for IR.

One very interesting such work is the one by *J. Savoy* that is reported in (Savoy, 1992). Savoy assumes the existence of a hypertext of documents and proposes a methodology for constructing automatically a Bayesian tree of documents terms. This tree, whose nodes are terms extracted automatically from the collection, represents the probabilistic dependency relationships existing between terms, and can be browsed by the user using a interface that enables him to have a “fish eye view” of relationships connecting terms in the tree. The user, traversing this tree, marks the terms he considers most usefull to express his information need. The user can express his belief about the importance of a term using a scale of possible beliefs, such as: very pertinent, pertinent, no assessment, irrelevant, very irrelevant. As the user marks terms other terms becomes increasingly relevant or irrelevant following the rules of belief propagation on Bayesian Networks (Pearl, 1988). Once this process of expanding the set of terms initially indicated by the user has ended, it causes a new process to start over the network of documents. This process consists of a form of constrained SA where the activity level of a document is evaluated using a function of statistically determined term weights (the usual indexing weights) and parameters indicating the importance of terms relative to the user’s interests (determined using the beliefs in the Bayesian tree). The spreading of activation on the hypertext is controlled by a set of rules like the following:

- Documents nodes are classified into classes (e.g. research papers, surveys, short articles, etc.) and a variable is attached to documents to identify the class. This is considered when evaluating the activation level of a document, so that the user can select the type of document he is most interested in and spread the activation only on this type of documents.

- In the first wave of SA links can be followed in any direction, while for successive waves hierarchical links are only followed towards specialisation.

– The number of waves of activation is limited to 5, though this is only a default value that can be change by the user. Indeed the best results were achieved using only 2 waves.

– During each wave the system does not allow reactivation of a document if that document and the document from which activation comes have a common parent.

Savoy's system, through the above described combination of a belief propagation over the indexing space and SA over the document space, achieves positive results, although it is not possible to compare its performance with traditional IR systems. A major drawback of this system is that, though it provides an automatic construction of the Bayesian tree for the indexing space, it still requires a manual construction of the hypertext. This can be a very heavy burden if the collection is large.

At the Department of Electronics and Informatics of the University of Padova (Padua) another attempt of using SA in IR is under study by *M. Agosti et al.*. The starting point is to model a set of raw IR data by means of a conceptual schema. In IR the term conceptual schema refers to a conceptual structure describing semantic relationships among IR data, i.e. among the different objects (documents, index terms, concepts, etc.) taking part in an IR application (Agosti et al., 1990). A conceptual schema of a specific IR application provides the user with a frame of reference in the query formulation process and can be very useful if the user is allowed to browse it. In the IR field the idea of using conceptual models is new. Most IR applications have an "ad hoc" data model. In the proposed conceptual model the application of the classification mechanism to the IR data implies working with three different levels of abstraction: documents, index terms, and concepts, as depicted in Figure 10. Links connect nodes to express a semantic relation between them. There are links connecting objects of the same type (on the same level) and links connecting objects of different type. For example a link connecting two index terms indicates that the two terms occur quite often together. A link connecting an index term with a document indicates that the document has been indexed using that term or that that term occurs in the document. A link between an index term and a concept indicates that the concept can be expressed using that index term. Links on the document level represent bibliographical citations or similarity between documents.

A methodology for the automatic construction of this schema was proposed in (Agosti and Crestani, 1993). Starting from a raw set of documents and a thesaurus like structure among concepts statistical and IR techniques are used to determine nodes and links. The resulting network schema can be browsed by means of a hypertext tool. The hypertext tool can be used for simple browsing or for query formulation, enabling the user to build up a query by moving through the IR data on different levels of abstraction and picking up

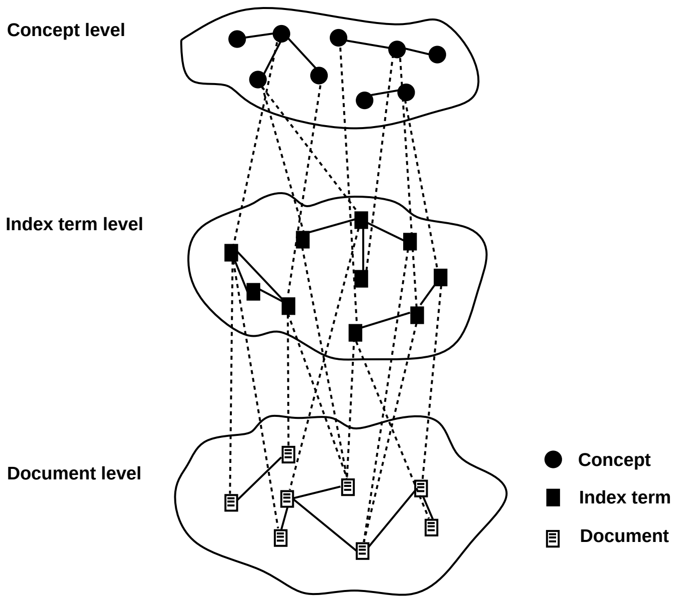

Figure 10. A three levels conceptual network.

the entities that better represent his information need (Agosti et al., 1995). After the user has built up a query an automatic procedure making use of the different semantics associated to links and node type spreads activation over the network to concepts, index terms, or documents that are closely related to those used by the user in the query formulation. A set of constraints take into account the node and the link type to control the spreading over the network. Active documents nodes are put in a retrieval list for the user to browse. He can provide some feedback to the system by marking the document nodes in the retrieval list that he considers relevant. In this way he assess if the spreading has been successful or not. This process in similar to the Relevance Feedback technique used in advanced IR system. He can then start a new spreading of activation and continue its search in a iterative and interactive process.

This work is currently at a prototypical stage. We will report in more details in a future paper on the architecture and on the evaluation of our prototype.

6. Conclusions

We reported on various approaches to the use of Spreading Activation in IR. The above presented exemplification of applications of SA to IR can be considered quite exhaustive, and although a few other attempts to use SA in IR can be found in literature, they are quite similar to those described here.

The field has proved itself to be very interesting and capable of providing good results. We expect to see much more research devoted to the subject when the problem of building network schemas of IR applications will find easier solutions by means of automatic techniques. A few research goups are currently working in that direction.

Acknowledgements

At the time of writing the author was a visiting Research Fellow at the Department of Computing Science of the University of Glasgow, Scotland. This work has been partly supported by the Italian National Research Council (CNR) under the project "Sistemi informatici e calcolo parallelo - P5: Linea di Ricerca Coordinata MULTIDATA".

References

Agosti, M. and Crestani, F. (1993). A methodology for the automatic construction of a Hypertext for Information Retrieval. In *Proceedings of the ACM Symposium on Applied Computing*, pages 745–753, Indianapolis, USA.

Agosti, M., Crestani, F., Gradenigo, G., and Mattiello, P. (1990). An approach to conceptual modelling of IR auxiliary data. In *Proceedings of IEEE International Conference on Computer and Comunications*, Scottsdale, Arizona, USA.

Agosti, M., Melucci, M., and Crestani, F. (1995). Automatic authoring and construction of hypertext for Information Retrieval. *ACM Multimedia Systems*, 3(1):15–24.

Chung, M. and Moldovan, D. (1994). Applying parallel processing to natural-language processing. *IEEE Expert*, pages 36–44.

Cohen, P. and Kjeldsen, R. (1987). Information Retrieval by constrained spreading activation on Sematic Networks. *Information Processing & Management*, 23(4):255–268.

Crestani, F. (1991). A survey on the application of Neural Networks's supervised learning procedures in Information Retrieval. Rapporto Tecnico CNR 5/85, Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo - P5: Linea di Ricerca Coordinata Multidata.

Crestani, F. and van Rijsbergen, C. (1993). Modelling Adaptive Information Retrieval. Depertimental Research Report IR-93-2, Department of Computing Science, University of Glasgow, Glasgow, Scotland.

Croft, W. (1987). Approaches to Intelligent Information Retrieval. *Information Processing & Management*, 23(4):249:254.

Croft, W., Lucia, T., and Cohen, P. (1988). Retrieving documents by plausible inference: a preliminary study. In *Proceedings of ACM SIGIR*, Grenoble, France.

Croft, W., Lucia, T., Crigean, J., and Willet, P. (1989). Retrieving documents by plausible inference: an experimental study. *Information Processing & Management*, 25(6):599–614.

Croft, W. and R.H.Thompson (1987). I^3R: a new approach to the design of Document Retrieval Systems. *Journal of the American Society for Information Science*, 38(6):389–404.

Fahlman, S., Touretzky, D., and van Roggen, W. (1981). Cancellation in a parallel semantic network. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 257–263.

Frakes, W. and Baeza-Yates, R., editors (1992). *Information Retrieval: data structures and algorithms*. Prentice Hall, Englewood Cliffs, New Jersey, USA.

Fujii, H. and Croft, W. (1993). A comparison of indexing techniques for japanese text retrieval. In *Proceedings of ACM SIGIR*, pages 237–246, Pittsburgh, PA, USA.

Harman, D. (1993). Overview of the first TREC conference. In *Proceedings of ACM SIGIR*, pages 36–47, Pittsburgh, PA, USA.
Jacobs, P., editor (1992). *Text-Based Intelligent Systems*. Lawrence Erlbaum Associates, Hillsdale, new Jersey, USA.
Jacobs, P. and Rau, L. (1993). Innovations in text interpretation. *Artificial Intelligence*, 63:143–191.
Kim, J. and Moldovan, D. (1993). Classification and retrieval of knowledge on a parallel marker-passing architecture. *IEEE Transactions on Knowledge and Data Engineering*, 5(5):753–761.
Kimoto, H. and Iwadera, T. (1990). Construction of a dynamic thesaurus and its use for Associative Information Retrieval. In *Proceedings of ACM SIGIR*, Brussels, Belgium.
Kjeldsen, R. and Cohen, P. (1987). The evolution and performance of the GRANT system. *IEEE Expert*, summer:73–79.
Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Mateo, California.
Preece, S. (1981). *A spreading activation model for Information Retrieval*. PhD thesis, University of Illinois, Urbana-Champaign, USA.
Quillian, R. (1968). Semantic memory. In Minsky, M., editor, *Semantic Information Processing*, pages 216–270. The MIT Press, Cambridge, MA, USA.
Rau, L. (1987). Knowledge organization and access in a conceptual information system. *Information Processing & Management*, 23(4):269–283.
Robertson, S. and Sparck Jones, K. (1976). Relevance weighting of search terms. *Journal of the American Society for Information Science*, 27:129–146.
Rumelhart, D., McClelland, J., and PDP Research Group (1986). *Parallel Distributed Processing: exploration in the microstructure of cognition*. MIT Press, Cambridge.
Rumelhart, D. and Norman, D. (1983). Representation in memory. Technical report, Department of Psychology and Institute of Cognitive Science, UCSD La Jolla, USA.
Salton, G. (1968). *Automatic information organization and retrieval*. Mc Graw Hill, New York.
Salton, G. (1989). *Automatic Text Processing*. Addison-Wesley.
Salton, G. and Buckley, C. (1988). On the use of spreading activation methods in automatic Information Retrieval. In *Proceedings of ACM SIGIR*, Grenoble, France.
Savoy, J. (1992). Bayesian inference networks and spreading activation in hypertext systems. *Information Processing & Management*, 28(3):389–406.
Schiel, U. (1989). Abstraction in Semantic Networks: axiom schemata for generalization, aggregation and grouping. *SIGART Newsletters*, 107:25–26.
Shoval, P. (1981). Expert/consultation system for a retrieval data-base with semantic network of concepts. In *Proceedings of ACM SIGIR*, pages 145–149.
van Rijsbergen, C. (1979). *Information Retrieval*. Butterworths, London, second edition.
Voorhees, E. (1994). Query expansion using lexical-semantic relations. In *Proceedings of ACM SIGIR*, pages 61–69, Dublin, Ireland.

Address for correspondence: Fabio Crestani, Dipartimento di Elettronica e Informatica, Universitá di Padova, I-35131 Padova, Italy, email: fabio@dei.unipd.it.