Scientometrics, Vol. 26. No. 1 (1993) 169-192

THE NEURAL NET OF NEURAL NETWORK RESEARCH

AN EXERCISE IN BIBLIOMETRIC MAPPING

A.F.J. VAN RAAN, R.J.W. TIJSSEN

Centre for Science and Technology Studies, University of Leiden, Wassenaarseweg 52, P.O. Box 9555 2300 RB Leiden (The Netherlands)

(Received January 6, 1992)

In this paper we discuss the limits and potentials of bibliometric mapping based on a specific co-word analysis. The subject area is neural network research. Our approach is a 'simulation' of expert assessment by offering the reader a narrative of the field which can be used as background information when 'reading' the bibliometric maps. The central issue in the applicability of bibliometric maps is whether these maps may supply 'additional intelligence' to users. In other words, whether such a bibliometric tool has an epistemological value, in the sense that it ecriches existing knowledge by supplying 'unexpected' relations between specific 'pieces' of knowledge ('synthetic value') or by supplying 'unexpected' problems ('creative value'). We argue that sophisticated bibliometric mapping techniques are indeed valuable for further exploration of these 'epistemological' potentials. In particular, these techniques may open new avenues to study science as a self-organizing system in the form of a 'neural network like' structure of which the bibliometric map is a first-order approximation. In that sense, this paper deals with the 'neural net of neural network research' as our bibliometric techniques in fact mimic a connectionistic approach.

Introduction

In this paper, we discuss the construction of maps of a specific research field by bibliometric (i.e., literature – based) methods. The subject area of this 'cartography' is neural network research. The applied bibliometric method is co-word analysis. In order to give the reader a first impression of the value and the usefulness of these 'science maps', we start this paper with a general overview, a *narrative* of neural network research. After this narrative, we discuss the construction of *bibliometric maps* of neural network research. The maps will be discussed with the narrative as a reference frame in order to investigate to what extent the bibliometric approach does reveal important features of the field. and whether it has the potential to yield additional information, for instance specific linkages between research specialities or emerging topics that are not or hardly covered by the narrative or the reviews on which the narrative has been based. Finally, we suggest that bibliometric maps

Elsevier, Amsterdam – Oxford – New York – Tokyo
Akadémiai Kiadó

themselves, as they are based on a connectionistic approach, reflect the self-organizing structure of science which may be conceived as a neural net of cognitive elements.

A narrative of neural network research

A *neural network* is system of a large number (typically more than 10^3) highly interconnected elements. These elements, *neurons*, have the capability to switch from an inactive to an active state when the total input (the sum of incoming signals from many other elements) exceeds a certain threshold value. This system thus resembles the biological neural system in which neurons emit ('fire') a voltage pulse (about 100 mV) when the sum of the inputs from their *synapses* (couplings between two neurons) exceeds a threshold. This pulse propogates itself through the network thus contributing to the input signal of other neurons. The use of pulses is a very efficient way to transport 'information' through the network. The strength of the linkage between two elements is, in principle, variable and in fact the whole operation of neural networks boils down to the creation of specific strengths of neuronal linkages. The process of learning, for example, can be seen as changing the strength of neuron connections. In artificial networks this can be used to give the system specific capabilities. For instance, in pattern recognition research one investigates which connection strengths are necessary to bring the network in a specific state. Neural networks therefore serve as models for cognitive and computational processes but, undoubtedly, they are extreme simplifications of biological systems, more particular the *brain*. Nevertheless, neural networks allow systematic, 'first approximation' studies on computational principles in biological systems, in particular *parallel computing* which differs significantly from *digital computing*. Psychologists use neural network models to understand *cognition*, i.e. cognitive states and processes in the human mind. Physicists are challenged by the *statistical mechanics* of neural networks, in particular the dynamical properties of neural networks as large and strongly coupled non-equilibrium systems. Another stimulation is the remarkable phenomenon that specific information is not localized in one place in the brain but distributed over many elements.

In 1943, *McCulloch* and *Pitts*[1] proposed networks of simple two-state elements to perform logic operations. A few years after the second worldwar, *Hebb*[2] introduced *connectionism* by presenting a model in which a neural system can learn through specific changes in the synaptic connections. In particular, the capability to perform simple recognition tasks (adaptive networks) was investigated in the sixties. Many neural network models for *adaptive* tasks such as *associative memory* and *pattern recognition* have been studied in the 1970's and the 1980's (*Grossberg*,[3] *Kohonen*[4, 5]).

An important development was the idea of *Little*[6] to explore the analogy between simple neural networks and magnetic systems. From then on, the interest of physicists in neural network research increased. A crucial further step was the work of *Hopfield*.[7, 8] He elaborated the analogy between long-time behaviour of neural networks with (symmetric) connections (in particular for associative memory), and statistical mechanical equilibrium properties of specific magnetic systems, *spin glasses*.

Why spin glasses? Physicists call a system (i.e., a 'collection' of very many molecules) a spin glass if it has randomly distributed ferromagnetic and anti-ferromagnetic interactions at a temperature below a specific value. The longe-range interactions in spin glasses resemble several important features of neural networks, in particular the phenomenon of coupling of one element to many others. Again, biological reality is quite different from physical systems like spin glasses. First, in a spin glass system the mutual interactions between the elements are known and based on random distributions, which means that statistical mechanics allows calculation of the macroscopic behaviour of the system. In the cortex, however, many neural connections are already fixed in a very specific way, both genetically and by learning. Therefore, in artificial neural networks one tries to find out how the connections between the elements must be, in order to bring the system into the desired mode of operation. A second important difference between the spin glass model and biological reality is that neurons in the cortex are not reciprocally symmetric. The dynamical properties of the biological neuronal system will therefore differ substantially from those of magnetic systems in statistical mechanical equilibrium (where pairwise interactions are strictly symmetric). In artificial neural networks, however, the statistical mechanical approach is useful, particularly in studies of associative memory where a symmetric architecture can be applied. For a review of recent developments and basic features of neural network theory as seen from a statistical mechanics viewpoint, we refer to *Sompolinsky*.[9]

In a simple network, information can only flow in the forward direction. This means mapping of the input neural-network layer state onto the output layer state. An early example of such a simple 'feedforward network' is the perceptron (*Rosenblatt*,[10] 1958). Interest in perceptrons declined in the sixties, and the limits of these one-layer networks contributed to the ascendency of arch-rival, *artificial* *intelligence* research (*Rappa* and *Debackere*,[11] *Valentine*[12]). Multi-layer variants of the perceptron however show a revival in recent years because of their usefulness in associative memory and pattern recognition tasks. Given the fixed configuration of the input layer, all neurons in subsequent layers 'compute' their states in parallel from the states of the preceding layer. This *parallel architecture* has an important linkage with research on *parallel processing* and, in particular, *parallel distributed processing*, one of the important current research topics in the network neural field. Pioneering work on parallel processing McCulloch-Pitts networks was done by *Von Neumann*[13] in the 1950's. Again, the simplicity of feedforward parallel computing as compared to the brain is obvious. In the cortex, feedback is an essential feature of the nervous system. These feedback loops are crucial for computational performance and, more specifically, for the control of animal and human behavioural patterns.

Associative memory is the ability to retrieve stored information with help of other information, or from partial or corrupted sets of that earlier stored information: *content addressible memory*. In associative memory we have to distinguish between two main processes. First, the storage of information, which can be characterized as a learning process. Here the work of *Hebb*[2] played a pioneering role. According to *Hebb*, neurons active in a specific pattern will induce changes in their (synaptic) connections in such a way that the pattern will be reinforced. The system therefore is able to learn 'unsupervised', and here we notice the important link between associative memory and *self-organization*. We already mentioned that learning means adapting the neuronal connections. For many cases, the period after birth available for adapting the 'wires' is limited. Thus, the capability to learn get lost after some period of time. A striking example is the training of a child's 'lazy eye'. Retrieval is the second dynamical process. The associative nature of retrieval implies that the system must be able to handle many patterns of information simultaneously. Extremely important is the robustness of the neural network system against *noise* and failures of individual neurons and synapses.

The earlier mentioned work of *Hopfield* in the 1980's presents a first and detailed description of associative memory on the basis of the statistical mechanics of neural networks consisting of simple two-state neurons. Undoubtedly, the work of *Hopfield* marks the onset of a revival in neural network research. Robustness against noise (small amounts) is an important feature of the model. The symmetrical properties of the Hopfield model, however, are not compatible with the biological situation (symmetry of connections, symmetry of neuron activity values, and the fact that neurons equally exhibit excitatory and inhibitatory capacities). An earlier neural network model proposed by *Willshaw*[14] has a biologically more plausible structure. This model mimics the cortex with two types of neurons: excitatory ones to excite further neurons into the active state, and inhibitory neurons to inhibit the activity of other neurons. The crucial point is that information is encoded only in connections between excitatory neurons. *Willshaw* was inspired by the earlier mentioned intriguing phenomenon of the non-local, i.e. distributed character of information storage in the brain: each 'memory' is represented by a pattern distributed over many elements, and, consequently, each element is involved in representing many different memories. This phenomenon is, in fact, the 'natural solution' for robustness, i.e., resistance to local damages or inaccuracy. In the early 1960's the problem was tackled by using the *hologram* as a physical analogy. The hologram is indeed a remarkable form of distributed information storage. *Willshaw* pointed out the problems of the holographic approach and introduced his network that mimics the holographic principles in a simpler and non-optical way.

As mentioned earlier, the symmetry of synaptic connections in the neural network is a prerequisite for models based on equilibrium statistical mechanics. Reality, however, is asymmetric. An important part of recent theoretical work on neural networks focuses on this synaptic asymmetry. A wealth of new dynamical features presents itself. In particular, asymmetric networks can be used for making temporal associations.

In the past decade, much work has been done on learning algorithms. In terms of associative memory, learning is the organization of the states of neurons in a specific configuration that resembles a priori known configurations. In statistical mechanics terms, the stable configurations correspond to energy functions that must be minimized. These Boltzmann models indeed provide a useful framework for further theoretical investigations. Their usefulness to understand biological systems, however, is still questionable. Minimization of the energy function of specific neuron state configurations is strongly related to the problem of *optimization* and, in particular, combinatorial optimization. This problem is more complicated than retrieving memories as there is no 'exemplar' or 'clue' about the desired optimum. One has to find a minimum in the energy function starting with unbiased configurations. A well-known topic here is the solution of the *traveling salesman* problem.

'Dynamical' pattern recognition is essential for learning processes in living organisms. Reality is not a matter of fixed patterns but of time-dependent patterns: moving images, language and music recognition, etcetera. Learning a melody means that after the input of a few first sounds, the network generates ('recognizes') the rest of the melody. In other words, dynamical pattern recognition is strongly related to vital capabilities such as *hearing, speech recognition*, and *vision*.

In biological reality neurons have different functions. For instance, we have neurons to receive the input signals of the retina. These neurons have a sensoric function. Other neurons induced muscular activity, they have a motoric function. Sensoric and motoric neurons have quite different characteristics. However, sensoric input signals are coupled to output signals of a biological system, for example to make movements. There is causal relation between input and output, and the biological system thus operates in a closed-loop state. Biological learning means finding the correct transformation from changes in the sensoric input to the motoric response, a good example of the earlier mentioned unsupervised learning. This topic attracts much attention, in particular in the development of artificial autonomous systems such as *robots*.

Generally speaking, most of the artificial neural networks are as yet applied in pattern recognition tasks. Important examples are *signal* and *picture (image) processing* as well as *character recognition* (artificial reading). We may also expect a growth in application in more 'natural' tasks, such as sensoric *information processing* to execute complex movements (robotics), and in decision making in complex problems with incomplete information.

Interest in neural networks stems from application-oriented as well as from theoretical considerations. Neural networks suggest new architectures for computers and new learning methods. Undoubtedly, the major goal of neural network research is to understand the *nervous system*, and, eventually, the human brain. It is still an open question whether neural network based computers will perform better than our conventional digital computers. Mathematical calculation – a typical rule – based problem – is not the strongest side of neural networks. We certainly gained more insight into operation of large neural systems. But the impact of neural network research on *neurophysiology* is small. The simplified models are still far remote from biological reality. One does not even know whether one should consider the whole cortex (about 10^{11} neurons) as one enormous neural network, or that one should work with a system of many interconnected neural networks, as suggested by neurophysiological considerations. These considerations would suggest a system of about 10^6 networks of 10^5 neurons each. Maybe one single neuron is already a large network on itself. Therefore, it is clear that very recently the possible *fractal* structure of biological nervous systems is attracting more and more interest. Perhaps this approach will unravel other dynamical features of biological systems and characteristics of biological constraints.

Nevertheless, neural network research is now in a rapid evolution. Apart from the challange to explore the brain, neural networks have already become a new R & D area, devoted to such practical applications as mentioned above. Neural nets are now simulated in software and implemented on digital computers ('netware'). This represents a remarkable flirt of the original arch-rivals, neural networks and artificial intelligence research. Whether it will come to a marriage, will have to been seen. The next step will be the development of neural net-hardware. Semiconductor technology now allows for the relatively cheep production of microprocessors which can take the place of the simple two-state artificial neurons used so far.

A bibliometric map of neural network research

Introductory remarks

In the foreign chapter we presented a *narrative* on a specific field in science, neural network research. It is a description of this research field based on opinions of experts in the field, given in, for instance, review articles. Parallel to the narrative, it is most interesting to explore what type of information, what picture of the field will emerge from a quantitative, *bibliometric*[15] approach of the same field. The rationale for such an approach is the following. Science is a complicated, heterogeneous system of activities characterized by many interrelated aspects. Nowadays, there is an enormous and ever increasing amount of information on scientific research, embodied in publications. It is a challenge to develop techniques for extracting well-structured patterns of information from such a rather 'amorphous' mass of data. This may reveal underlying and until now hidden features. A fruitful approach to solve this problem is the development of *bibliometric maps*. There are several important advantages of using such cartographical representations. A visualization of complex masses of data offer a more complete overview in less time. Furthermore, visual information is more easily remembered. Another very important point is, as indicated above, the reduction of information. It is a crucial problem to filter the significant features. As well shall see, mapping techniques developed in our group offer the possibilities to achieve such a data reduction. In other words, a *cartography of science* not only reformats the data into a specific visual representation, they also accomplish data reduction while retaining essential information. The next step is obvious. Maps are not only suitable for depicting a static structures. Time-series of maps enables a visualization of dynamic features of science, for instance the identification of important changes over time in the development of research fields, or shifts in emphasis of countries, research organizations, or research groups. Maps of science can be seen as tools for searching, identifying and analyzing structures of scientific activities as reflected by publications. They may point at merging fields of science, emerging new activities, and they offer insight into the position of countries, research organizations or institutes in a field of science. Maps aggregate data in a way no expert, with his or her background and perspective would be able to do. The cartographic approach is, so to say, independent of individual opinions. This is particularly advantageous in the case of broad and heterogeneous research fields. This does not mean, however, that maps can replace the opinions of experts. A thorough interpretation of science maps requires knowledge about the subject matter of the map, preferably from the 'users'.

Maps of science

In bibliometric analysis we may distinguish between *one-dimensional* and *two-dimensional* techniques. The one-dimensional techniques are based on *direct* counts (*occurrences*) of specific bibliographic items (e.g., publications and patents), or particular data-elements in these items, such as citations, keywords, addresses, etcetera. We call these techniques 'one-dimensional' as they are in principle represented by *lists* of numbers. There are, however, more possibilities, to make a bibliometric representation of the scientific endeavour. Whereas one-dimensional indicators are based on direct countings (i.e., the measurement of occurrences), the *two-dimensional* (relational) indicators are constructed from *co-occurrences* of specific information elements, such as the number of times keywords or citations are mentioned together in publications in a particular field of science. We already mentioned that publications carry different information elements such as author names, classification codes, keywords, references. In many cases, a publication is characterized by more than one author name, address, classification code, keyword, of reference. This means, that for each scientific field for instance all keywords of publications can be collected, and for each keyword in the compiled set one may analyse how many times that keywords occurs together (i.e., co-occurs) with any other keyword in publications involved (*Callon* et al.[16, 17]). This type of analysis thus yields an array, or *matrix*, of pair-wise keyword-relations, the 'word co-occurrence' or *co-word* matrix (hence we call this type of bibliometric analysis a 'two-dimensional' technique). Co-citation analysis is a similar technique. Here the co-occurence of references in scientific publications is the basis of a 'reference co-occurrence' or *co-citation* matrix.[18, 19, 20]

With help of special data-analytical techniques based on matrix-algebra, it is possible to convert the information in such a co-occurrence matrix into a spatial configuration of the elements (keywords, references) in a two-dimensional space. The best results are attained by a sophisticated combination of clustering-techniques and multi-dimensional scaling (for details we refer to *Tijssen*[21] and *Peters* and *Van Raan*[22, 23]). In this way, a 'map' is made of a scientific field, and, by applying the same techniques to patents (*Engelsman* and *Van Raan*[24]) of a combination of scientific and technological fields (mapping of R & D activities).

The advantage of co-word analysis is that it is applicable to fields where no citation-index is available, or where citation practices are such, that a citation-index does not cover sufficiently the research of those fields. This is especially the case for application-oriented research (*Van Raan* and *Tijssen*[25, 26]). Equally important, co-citation analysis has a much smaller 'retrieval rate' than co-word analysis. In other words: only a small part (e.g., 14%) of the publications upon which the co-citation analysis is based, is represented by the co-citation map. In co-word analysis, the majority of original publications is represented. For these striking differences we refer to *Braam* et al.[27, 28]

Co-word analysis is completely independent upon citation practices. Main caveats are: words may have other than purely descriptive purposes and their meaning is often context-dependent. The main advantage of co-word analysis is given by the nature of words: words are the foremost carrier of scientific concepts, their use is unavoidable and they cover an unlimited intellectual domain.

Finally, as words represent cognitive elements, a structure of interrelations between these cognitive elements may be conceived as a 'neural net'. In a later section we come back to this 'connectionistic approach', in relation with the concept of science as a self-organizing system.

The neural network research map

The main lines of our bibliometric mapping technique are as follows. First of all, we have to define our subject area. In this case, neural network research, we applied a 'two-step co-word analysis'.

For the time period 1985–1988 we collected with help of the INSPEC database relevant data of all publications characterized by 'neural network-' *or* 'neural net-' as a title-word *or* as a controlled term *or* as an uncontrolled term (the latter terms include keywords from abstracts of publications). The number of the thus analyzed 'neural network publications' was 1237. Of all these publications we made a frequency analysis of the uncontrolled terms (which can be regarded as the most 'author-related' keywords). We defined the fifty most frequent words (after corrections for synonymes etc.) as the 'neural network family words'. This is the first step of the word-analysis. The 'family words' involved are, for instance, 'artificial intelligence', 'associative memory', 'connectionism', 'content addressible memory', 'distributed processing', 'expert systems', 'Hebb', 'image processing', 'information storage', 'optical computing', 'nervous system', 'Von Neumann', 'parallel processing', 'pattern recognition', 'picture processing', 'robot', 'self-organization', 'spin glass', 'synapses', 'vision', etcetera. The complete set of these neural network family words can be found on the maps: they constitute the conceptual framework with which we construct the neural network maps.

The second step in our procedure is to define publication sets for each of the family words (as an uncontrolled term) *separately* (i.e., a set of publications characterized by 'artificial intelligence', by 'associative memory', etc.) and to analyze per set the frequency of each *other* family word. Thus the original set of 1237 neural network publications has been extended with the sets of publications characterized by the family words. In this way, the number of total publications involved, is extended with more than an order of magnitude. For instance, the number of artificial intelligence publications in the period 1985–1988 is 5343, expert systems 9270, parallel processing 4747, image processing 2494, information storage 2163, spin glass 1518, pattern recognition 3601, picture processing 4983, and for the topic vision 6568 publications are involved. For each of the above sets characterised by one specific family-word we determined the frequency of each *other* family-word. For example, we established for all 'artificial intelligence' publications (as defined above) how many times these publications are also characterized by 'associative memory', 'connectionism', etc. (including 'neural net-'). Many publications will be related to more than one of these topics, so the number of publications involved in our total set will certainly be not the sum of all these individual counts. We found that the number of publications in our total set is about 20,000.

In this way, we were able to compose a 50×50 co-word matrix which can be described as a specific structure of neural network research 'interwoven' in its direct 'environment'. More specifically: it is a bibliometric map of the *texture* of neural network research and related fields. If we would have composed a co-word matrix of the same 50 family words for their co-occurrences *within* the first-step set (i.e., publications characterized by the word neural network) then the map would rather picture the structure of neural network research itself. For instance, it would indicate the relation of artificial intelligence and self-organization *as far as* this relation is embedded in publications that *also* deal with neural network research (i.e., publications would be involved characterized by at least the following keywords: neural network *and* artificial intelligence *and* self-organization).

Now, in our two-step approach, the 'family-words' originally generated by neural network publications (first step) are used to define the research fields related to these words and the occurrence of each possible word-pair is determined for all the sets separately. So, our map shows the relation between artificial intelligence and self-organization as it emerges from publications characterized by the keywords artificial intelligence *and* self-organization, without necessarily neural network as a keyword.

Results and discussion

Figures 1 and 2 present our ('two-step') co-word maps of neural network research and its 'environment', for the periods 1981–1984 and 1985–1988. The selected words in the map and the links between them clearly show that research on neural networks is characterized by a high degree of interdisciplinarity. It is assumed that linkages between specific words or concepts reflect cognitive links between the relevant research specialties. A purely textual description of a field, as given in the narrative, makes it virtually impossible to offer a *synthetic* view of all the relations between the topics in neural network research, and to provide a comprehensive structural overview. The maps, however, show this network of linkages between important (sub)fields and disciplines within neural network research, such as biology, cognitive phychology, computer science, and physics. The overall structure is quite clear. We notice the remarkable distinction between neural network research and the artificial intelligence research. We come back to this important point furtheron. Research on vision and pattern recognition functions more or less as a bridge between the two fields, particularly in the first period.

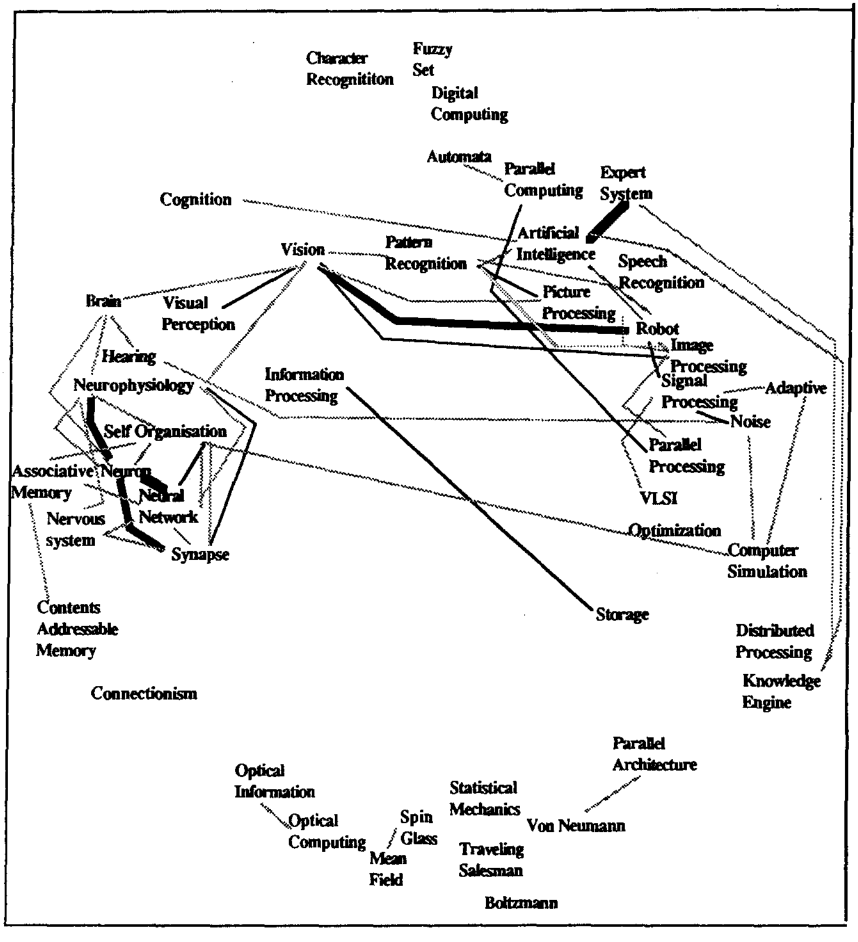

Strength of linkages (Jaccard index): $J > .099$

$.050 < J \leq .099$

$.015 < J \leq .050$

$J \leq .015$ not shown

Fig. 1. Co-occurrence structure of keywords related to the keyword "Neural Network", 1981 – 1984

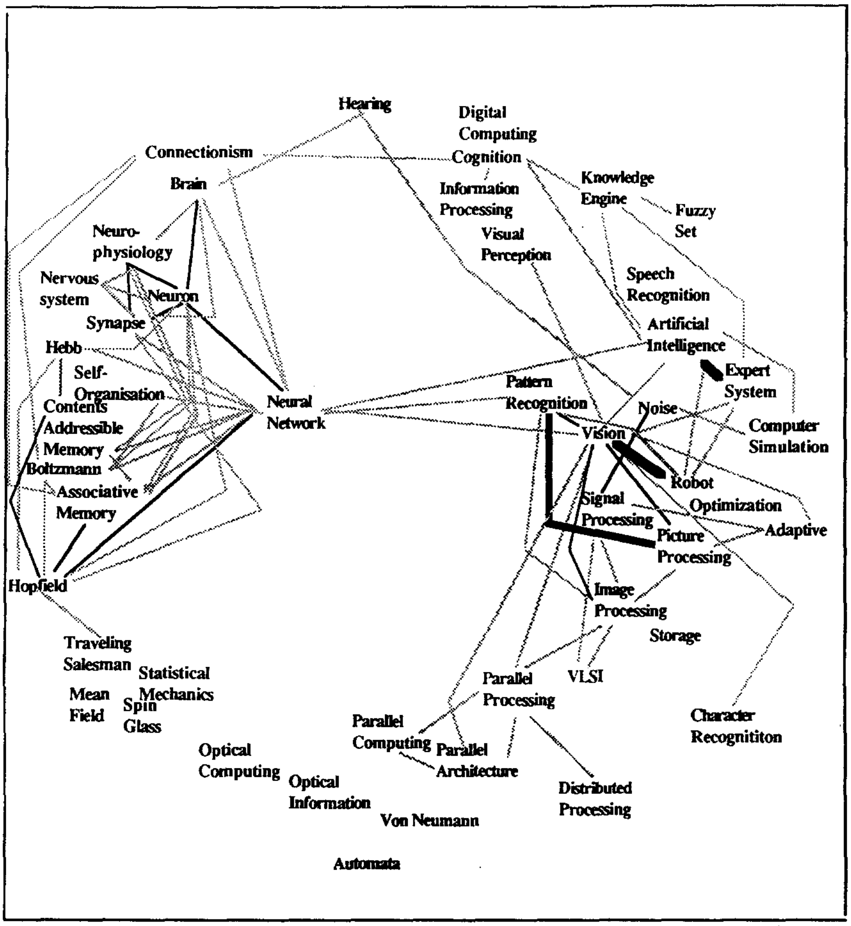

Strength of linkages (Jaccard index):

$J > .099$	
$.050 < J \leq .099$	
$.015 < J \leq .050$	
$J \leq .015$	not shown

Fig. 2. Co-occurrence structure of keywords related to the keyword "Neural Network", 1985–1988

After this general impression, we will discuss the maps in more detail. We focus on the most recent map, 1985–1988 (Fig. 2), as it gives the best approximation of the field at this moment. Moreover, because of the strong growth of the field, the numbers of publications involved in the 1981–1984 map are considerably smaller. The narrative should be used as a reference frame, as background information to 'read' the map. It simulates, as it were, an 'expert use' of our bibliometric map. We already noticed a striking feature: there are two main clusters, one around neural networks, and one around artificial intelligence and vision. At the end of the sixties, the original field (neural networks and artificial intelligence) splitted up by a dramatically emerging arch-rivalry (*Minsky* and *Papert*,[29] *Papert*,[30] *Rappa* and *Debackere*,[11] *Valentine*[12]). This historical event is still characterizing the whole area, and this divergence between the fields of neural network research and artificial intelligence research couldn't be pictured better. But there is much more to see than just two clusters representing the arch-rival fields. Therefore, we 'read' the map by analysing cluster by cluster. Special attention will be paid to topics (research specialties) having three or more linkages to other topics with a linkage strength (Jaccard index) $\geq$ 0.15. Two linkages would mean just a position in a 'linear' chain, three linkages is the simplest node-structure. The three-or-more-linkages topics are regarded to have a central position in the map, or at least in the cluster they apparently belong to.

We start our analysis at the central left side of the map, where we see the *neural network cluster*. Remarkably, the most central concept 'neural network' takes a rather eccentric position in its 'own' cluster. This can be explained by the linkages this word has with most of the words in the whole map, making 'neural network' indeed central to the whole structure. Some topics play a self-evident central role: brain, neurophysiology, neuron, synapse. They are related to the biological basis of neural network research. They clearly take positions close together: the upper part of the neural network cluster. It is the 'natural' subcluster. We find connectionism and *Hebb* as two psychological research topics (of course, *Hebb* stands for his specific approach in associative memories, see in the section *Narrative* closely related to the biological basis. Hebb indeed functions in a middle position between the natural (biological) upper part and the lower, 'artificial' part. In this artificial neural network we find most of the key-topics. Some are already very established and 'old', but still indicate current and very important topics: associative memory, contents addressible memory; and others indicate the more recent developments in close relation with the 'old' concepts: Hopfield (model), Bolzmann (models), and self-organization. With help of these central words we are able to identify further research topics closely related to but not shown on the map for two reasons: (1) these *additional words* did not reach the word-occurrence or, most probably, the word co-occurrence thresholds, and (2) lack of space prevents us to lower the thresholds (the map would become virtually unreadable). We, however, have to experiment further with these important technical details as a number of these 'additional' words are indeed very important and should figure on the map. There are two approaches to tackle this problem. These additional words can be added as labels to specific key-words on the map, or they can be used for a 'local mapping', i.e., a mapping of a specific part of the larger map, as a kind of local enlargement. Work in our group is in progress to develop this type of mapping (*Peters* and *Van Raan*[23]). The first approach, the addition of 'labels', is shown in Fig. 3.

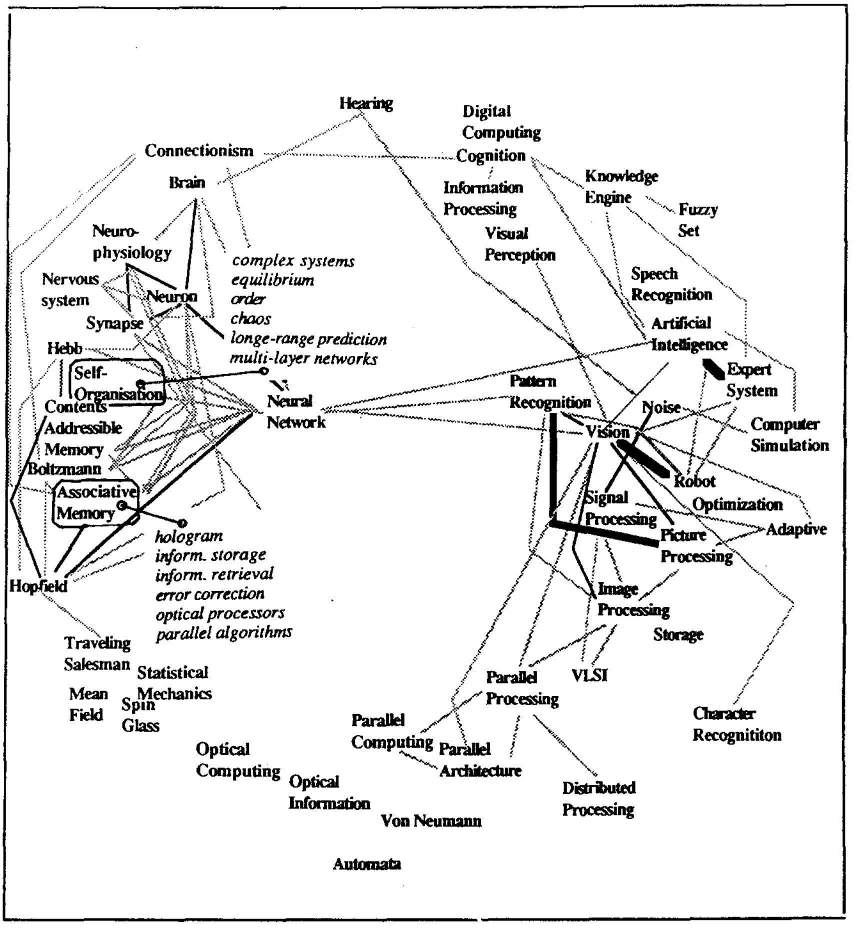

Fig. 3. Same as Fig. 2, but "labels" of additional keywords attached to main keywords (examples)

Returning to our central topics in the lower, 'artificial' part of the neural network cluster, we find the following *additional words* that increase in frequency as compared to an earlier four-year period (Fig. 3):

- related to associative memory: hologram, information storage and retrieval, error correction, optical processors, parallel algorithms;
- related to content addressable memory: fault tolerance, associative storage, high-speed retrieval;
- related to Hebb: high-order recursive neural network;
- related to self-organization: complex systems, equilibrium, order and chaos, longe-range prediction, multi-layered networks;
- related to Boltzmann: simulated annealing, combinatorial optimization, parallel optical systems, backward error-propagation network, distributed memory architectures, fast non-volatile programmable spatial light modula, Hinton machines, hidden units;
- related to Hopfield: stored patterns, high-temperature superconductors, numerical simulations;
- related to the most central topic, neural networks: computer vision systems, back-propagation, feature extraction, feedback, neural computing, parallel computing, artificial neurons.

There is little doubt that these additional words together with the main central words indicate topics of crucial importance in current neural network research. We may compare the above information with the main lines as pictured by the narrative, and find striking resemblances, but also important new information in terms of further topics or clear emphases. We mention for instance the evident role of error (fault) tolerance, the importance of feedback-, back(ward)-, multi-layer, and thus non-linear systems, the influence of recent statistical mechanical concepts such as simulated annealing, complex systems and equilibria, and order & chaos.

Below the *neural network cluster*, and nicely linking up with the just discussed physical, 'artificial' lower part of this cluster, we first find, going counter-clockwise, a very important part of physics. As there are no strong direct linkages with the main cluster, we may regard such a word aggregation as a kind of general 'knowledge supply' cluster. It is, in fact, the *statistical mechanics cluster* with spin glass research, the traveling salesman problem, and the more generic terms statistical mechanics and mean field. It is clear that the position of this statistical mechanics cluster, its main concepts and its relation with the neural network cluster (Boltzmann, in particular the linkage with Hopfield) as visible on the map, agrees very well with the information as given by the narrative. The next 'word aggregation' is again an important part of physics: optics. But its emphasis is on computing capabilities, thus linking up quite naturally with a further cluster around the central concepts parallel processing and parallel architecture (optical data processing is a parallel technique par excellence). In the *optics cluster* important additional concepts (with increased activity as compared to the earlier period 1981–1985) are:

- related to optical computing: non-linear optical properties, integrated opto-electronics, semiconductors, spatial light modulaters;
- related to optical information: non-linear optics, holography, optical fibres, (semiconductor) lasers, liquid-crystal light valve.

The *parallel processing cluster* includes distributed processing, which indicates a very important relation as discussed in the narrative. Furthermore, the work of *Von Neumann* and (cellular) automata are positioned in the direct vicinity of the core of the parallel processing cluster. We notice that this parallel processing cluster takes an almost opposite position to digital computing, which nicely illustrates the fundamental difference between both types of computing.

For this parallel cluster we find the following important additional words:

- related to automata: cellular automata, decidability;
- related to parallel computing: shared memory;
- related to parallel architecture: computerized picture processing, computer vision, systolic array, real-time control, transputers;
- related to parallel processing: vector processor;
- related to distributed processing: reliability, real-time systems.

The parallel processing cluster has several connections with a following, very nearby cluster dealing with visual techniques of *computer vision*. Central concepts in this large and important application-oriented conglomerate of research topics are: image processing (acting as a bridge to the parallel processing), the practically synonyme picture processing, signal processing, pattern recognition, robot, adaptive, and the most central concept of vision (also connected directly with parallel processing). This computer vision cluster clearly focuses on automated pattern recognition, in particular for industrial robots with manipulative tasks such as position control and assembling. In this cluster we also find VLSI (strongly related to integrated circuit (IC-) technology and computer-aided design), storage (e.g., optical discs), optimization (strongly related to operations research, control systems, computer-aided design), and, in a rather remote position, character recognition (e.g., recognition of Chinese characters, hand-written characters and numerals): both pattern recognition and the more generic concept of vision are linked up with neural networks. For this computer vision cluster we find the following important additional words:

- related to image/picture processing and: remote sensing, edge detection;
- related to pattern recognition: edge detection;
- related to robot: computer vision;
- related to adaptive: robustness;
- related to storage: reliability;
- related to optimization: combinatorial optimization, reliability;
- related to character recognition: segmentation, computer vision.

There are two direct linkages (i.e., linkages with a Jaccard index $\geq$ 15) of the computer vision cluster with our main neural network cluster. In this cluster, pattern recognition is positioned most closely to neural networks. These linkages indicate that research on computer vision and pattern recognition tasks are strongly related to recent developments in neural network research. As we know from the narrative, this relation is undoubtedly one of the most characteristic features of current neural network research.

The computer vision cluster on its turn is positioned closely to what we call the *artificial intelligence cluster*. Expert systems is the central topic very closely related to artificial intelligence. This cluster seems to be somewhat long-drawn, via knowledge engine to cognition as central concepts. Artificial intelligence is the domain of research on knowledge representation, logic programming, software engineering (PROLOG, LISP) and thus, typically, rule-based systems. In relation to this, natural language processing is also an important research topic in artificial intelligence. This explains the close position of speech recognition on our map.

Additional concepts of increasing importance related to artificial intelligence and expert systems are reasoning and user interfaces.

The almost opposite position of the artificial intelligence cluster as compared to the neural network cluster, reflects the *arch-rival relation* between these two fields as discussed earlier. Nevertheless, there is a direct linkage between artificial intelligence and neural network research visible on our map. Partially, this will be due to the simple fact that when two things are arch-rivals, they often will be mentioned together. But as we discussed in the narrative, recent developments (neural network inplementation in digital computing) do build a bridge between the two (former?) arch-rivals.

We also notice a linkage between artificial intelligence and vision, artificial intelligence and robots, and between expert systems and robots. These linkages enhance the relation of the computer vision cluster with the artificial intelligence cluster. This connection is further enhanced by the remarkable position of visual perception (linked up with vision) in the upper, cognition-oriented part of the artificial intelligence cluster.

We notice the remarkable position of noise, just between the computer vision cluster and the artificial intelligence cluster. Noise, of course, is a very generic term, extremely important in all systems with interconnected elements. Notice the linkage of noise with the computer vision cluster through signal processing, (e.g., 'signal-to-noise' ratio) and the linkage with the artificial intelligence cluster through computer simulation. These linkages indicate the role of noise in, for instance, (in)stabilities in (digital) simulations of complex systems. The linkage between noise and hearing is a trivial one, but should nevertheless appear, as is indeed the case.

Moving to the upper part of the artificial intelligence cluster, we already noticed the central position of knowledge engines. This topic is on the one side connected with artificial intelligence and expert systems, which is obviously the more application-oriented side (the use of artificially intelligent techniques for problem solving, fault diagnosis, decision support systems). On the other side there is a connection with cogition, the more basic side (knowledge acquisition, knowledge elicitation, reasoning). The linkage with fuzzy set relates to problems of uncertainty, probability, and 'approximate reasoning' in formal logic (thus called 'fuzzy logic').

Cognition is the other central concept in this part of the cluster. It links the artificial intelligence work with psychology, which is clearly visible on the map by the linkage between cognition and connectionism. Cognition represents the 'human' or 'natural' side of the subjects involved in artificial intelligence: reasoning, problem solving, decision making. But it also represents such other human capabilities as learning and perception. This explains the close positions of the keywords information processing and visual perception. Learning and perception are the cognitive concepts which cannot be investigated successfully by the typically rule-based artificial intelligence approaches alone. Here neural network research, by its connectionistic and associative character, offers more promising developments. Therefore, it is almost unavoidable that we indeed find a direct linkage on our map between the two archrivals, artificial intelligence and neural networks. With cognition the circle of clusters is closed: we are back to the main cluster of our map, neural network research. Here connectionism is the central concept. For this topic we find many additional topics of increasing importance, all indicating the central character of connectionism, and also the vitality of this well-established concept. As additional keywords we find: computer vision, knowledge representation, massively parallel networks, natural language processing, distributed representations, genetic algorithms, linguistics, machine learning.

In order to get an impression of *dynamical* features of neural network research, we have to analyse changes over time. Therefore, we compare our 1985–1988 map with the map for the earlier period 1981–1984 (Fig. 1). We concentrate on the main differences between the two maps. Comparing the neural network cluster with its predecessor in the period 1981–1984, we immediately see striking differences: the absence of the Hopfield model, and the shift of the Boltzmann model from a rather remote position in the statistical mechanics cluster to the heart of the neural network cluster. Moreover, the whole statistical mechanics cluster moved considerably towards the neural network cluster, as if it has been dragged in by the Boltzmann model. Connectionism changed its earlier position near content addressable memory to a more central position in the upper part of the neural network cluster, linked up with cognition. As we already mentioned, comparison of the 1981–1984 map with the 1985–1988 map also shows that neural network left its position in the middle of its own cluster and moved to a position more central to the whole area.

The optics cluster hardly changes its (relative) position. But for the parallel processing cluster, the computer vision cluster, and the artificial intelligence clusters considerable rearrangements has taken place. Parallel and distributed processing techniques (together with automata) joint, near Von Neumann. The computer vision related keywords have also clustered, and the same is the case for the artificial intelligence related words. The shift of visual perception and vision from a position near the 'biological' side of the neural cluster towards a clustering with computerized image and signal processing shows the advancing automation of visual techniques.

Connectionism and distributed processing of information were important research topics already in the 1940's and 1950's, and even before that. But it is the interaction of these concepts with physics, mathematics, biology and, in particular, technology (which is, in many cases, again the result of physics research, think of semiconductor technology) which open new ways in both further scientific development and in technological application.

From practical applications to epistemological values

Maps of science are not only useful descriptive tools, but may also serve as an interface between 'objective', literature-based data and 'subjective' accounts of experts. To this end, we used Figs 1 and 2 as input for an international questionnaire distributed amongst researchers active in various parts of the neural network research field. The main goal of this survey was twofold: (1) to assess the validity of these maps, as compared to the views of experts on (interdisciplinary) relations between research topics, (2) to obtain data on current trends in interdisciplinary relations, and which developments are to be expected in the near future. Regarding the first issue, the responses of the experts indicated that our maps have, overall, a fair degree of resemblance with their own views on the intellectual structure of the field. It was found, however, that the background of the experts strongly influenced their view on the field as a whole, as well as their judgement of the map's accuracy. A full account of the results is presented in *Tijssen*.[31] The maps also proved to be very useful as a means to elicit detailed information regarding the second issue. The information from seven experts who completed the questionnaire in detail, gives a fuzzy picture of current trends and future developments. If we consider the scope and the complex, multidisciplinary nature of neural network research, it is, of course, no surprise that the comments from the experts are very diverse, focussing on different subject areas and varying in range and level of detail. The only feature their opinions have in common is that natural sciences and life sciences are integrating in neural network research (particularly physics and neurophysiology), and that this process is likely to continue in the years to come.

The general objective of our research is to explore the limits and potentials of bibliometric mapping based a specific, 'two-step' co-word analysis. In particular, we aim at the *epistemological potential* of bibliometric mapping, i.e., its value as a means of *advancing* knowledge in addition to the knowledge it is based upon.

This surplus value may be found in 'synthetic' or 'creative' elements. The first type of knowledge growth is related to the discovery of new relations between specific pieces of knowledge, the latter type is related to the discovery of new problems which demand priority in solution. A rather obvious prerequisite of such an 'epistemological tool' is its capacity to indicate *existing* concepts, 'pieces of knowledge', problems, etc., that are essential in the field concerned. We think that sophisticated bibliometric mapping techniques are indeed tools which sufficiently fulfil the above prerequisite. A next step will be a further exploration of the epistemological potential as defined above. The challenging point here is that bibliometric maps may be regarded as *cognitive patterns* resembling stored information in neural nets. In other words, a bibliometric map represents, in first approximation, the *self-organizing* character of scientific activities in the form of a *neural network-like* structure (and thus our neural network map could be regarded as the 'neural net of neural network research'). This paradigmatic metaphor (also developed independently by *Ziman*[32]) closely links with our earlier empirical evidences (recently published in *Nature*[33]) that science can be regarded as a self-organizing system. Further empirical work is necessary to elaborate this exotic approach. In particular, it is challenging to investigate how the different neural net *layers*, as discussed by *Ziman*, can be distinguished – at least partly – by bibliometric methods.

*

The authors are grateful to their colleagues *Renger de Bruin* and *Harrie Peters* for their expert help in database searching and data collection. Part of this research was financed by the Netherlands Ministry of Education and Science, The Hague, and part by Elsevier Science Publishers, Amsterdam.

References

1. W.S. McCulloch, W.A. Pitts, A logical calculus of the ideas immanent in nervous activity, *Bull Math. Biophys.*, 5 (1943) 115–133.
2. D.O. Hebb, *The Organization of Behavior*, New York, John Wiley, 1949.
3. S. Grossberg, *Neural Networks*, 1 (1988) 63.
4. T. Kohonen, *Self-Organization and Associative Memory*, Berlin, Springer-Verlag, 1984.
5. T. Kohonen, *Neural Networks*, 1 (1988) 3.
6. W.A. Little, *Math. Biosci.*, 19 (1974) 101.
7. J.J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA*, 79 (1982) 2554–2558.
8. J.J. Hopfield, D.W. Tank, *Science* 233 (1986) 625.
9. H. Sompolinsky, Statistical mechanics of neural networks, *Physics Today*, (December 1988) 70–80.

10. F. ROSENBLATT, The perceptron: a probabilistic model for information storage and organization in the brain, *Psychological Review*, 65 (1958) 386–408.
11. M.A. RAPPA, K. DEBACKERE, The emergence of a new technology: the case of neural networks. Working Paper WP3031-89-BPS, Alfred P. Sloan School of Management, Massachusetts Institute of Technology, Cambridge (MA), 1989.
12. E. VALENTINE, Neurals nets: From Hartley and Hebb to Hinton, *J. Math. Psychology*, 33 (1989) 348–357.
13. J. VON NEUMANN, Probabilistic logics and the synthesis of reliable organisms from unreliable components. In: C.E. SHANNON, J. MCCARTHY (Eds.) *Automata Studies*, Princeton N.J.: Princeton University Press, 1956, pp. 93–98.
14. D.J. WILLSHAW, O.P. BUNEMAN, H.C. LONGUET-HIGGINGS, Non-holographic associative memory, *Nature*, 222 (1969) 960–962.
15. A.F.J. VAN RAAN (Ed.) *Handbook of Quantitative Studies of Science and Technology*, Amsterdam, Elsevier Science Publishers, 1988.
16. M. CALLON, J.-P. COURTIAL, W.A. TURNER, S. BAUIN, From translations to problematic networks: an introduction to co-word analysis, *Social Science Information*, 22 (1983) 191–235.
17. M. CALLON, J. LAW, A. RIP, *Mapping the Dynamics of Science and Technology*, London, MacMillan Press Ltd, 1986.
18. H. SMALL, Co-citation in the scientific literature: a new measure of the relationship between publications, *Journal of the American Society for Information Science*, 24 (1973) 265–269.
19. H. SMALL, E. SWEENEY, Clustering the science citation index using co-citations. I: A comparison of methods, *Scientometrics*, 7 (1985) 393–404.
20. H. SMALL, E. SWEENEY, E. GREENLEE, Clustering the science citation index using co-citations. II: Mapping science, *Scientometrics*, 8 (1985) 321–340.
21. R.J.W. TIJSSEN, A quantitative assessment of interdisciplinary structures in sciences and technology: co-classification analysis of energy research, *Research Policy* (to be published).
22. H.P.F. PETERS, A.F.J. VAN RAAN, (1991) Co-word based science maps. Part I: Direct multidimensional scaling maps in chemical engineering (to be published).
23. H.P.F. PETERS, A.F.J. VAN RAAN, Co-word based science maps. Part II: Clustering and multidimensional scaling combination maps in chemical engineering (to be published).
24. E.C. ENGELSMAN, A.F.J. VAN RAAN, Mapping of Technology. A First Exploration of Knowledge Diffusion Amongst Fields of Science. Research Report to the Ministry of Economic Affairs, CWTS 91–01, 1991. An adapted version will be published in the Policy Studies on Technology & Economy (BTE) Series Ministry of Economic Affairs, The Hague.
25. A.F.J. VAN RAAN, R.J.W. TIJSSEN, An Overview of Quantitative Science and Technology Indicators Based on Bibliometric Methods. Technology/Economy Programme (TEP), OECD, report 27769, Paris, 1990.
26. A.F.J. VAN RAAN, R.J.W. TIJSSEN, Numerical methods for information on aspects of science: Scientometric analysis and mapping. In: CH. OPPENHEIM, J.-M. GRIFFITHS, CH.L. CITROEN (Eds), *Perspectives in Information Management*, Vol. 2, Bowker-Saur/Butterworth-Heinemann, London (1990) p. 203–228.
27. R.R. BRAAM, H.F. MOED, A.F.J. VAN RAAN, Mapping of science by combining co-citation and word analysis. I: Structural aspects, *Journal of the American Society for Information Science*, 42 (1991) 233–251.
28. R.R. BRAAM, H.F. MOED, A.F.J. VAN RAAN, Mapping of science by combining co-citation and word analysis. II: Dynamical aspects, *Journal of the American Society for Information Science*, 42 (1991) 252–266.
29. M. MINSKY, S. PAPERT, *Perceptrons: An Introduction to Computational Geometry*, Cambridge (MA), MIT Press, 1969.
30. S. PAPERT, One AI or many? *Daedalus*, 117 (1988) 1–14.

31. R.J.W. TIJSSEN, Internal and external representations of science: comparing bibliometric maps with expert data. Report to the Netherlands Advisory Council for Science Policy, The Hague. Report CWTS 91–04, Centre for Science and Technology Studies, Leiden. 1991.
32. J. ZIMAN, The neural net model of innovation, *Science & Public Policy*, 18 (1991) 65–75.
33. A.F.J. VAN RAAN, Fractal dimension of co-citations, *Nature*, 347 (1990) 626.