Handbook of Quantitative Studies of Science and Technology
A.F.J. van Raan (editor)

Chapter 12:

CO-CITATION BIBLIOMETRIC MODELING AS A TOOL FOR S&T POLICY AND R&D MANAGEMENT: ISSUES, APPLICATIONS, AND DEVELOPMENTS

J.J. Franklin and R. Johnston

Centre for Technology & Social Change,
University of Wollongong, P.O.Box 1144,
Wollongong, N.S.W.
Australia 2500

Summary

This chapter is a comprehensive analysis of one literature-based methodology for 'modeling' the intellectual organisation and content of scientific disciplines. The method, called 'co-citation bibliometric modeling', provides a detailed description of the international research front. It may describe new inter- or multi-disciplinary developments in science, identify the most rapidly evolving subdisciplinary topic areas, and characterise the research activity of nations and organisations. As a result, it has been or is being explored as an intelligence tool for use in science and technology policy by 6 national governments, and may be of interest to large, high technology-based corporations. The objective here is to provide:

1) a guide to the methodology and its limitations
2) a guide to uses and interpretations of the data as they have been explored in nationally sponsored studies
3) an analysis of major unresolved technical and policy relevant issues
4) an assessment of recent methodological improvements, and
5) an agenda for future investigations and applications of this information tool.

1. Introduction

Over the past few years there has been a steadily increasing demand for improved means of policy formulation and decision-making about research. Technology has become the dominant factor in determining international competitiveness, and nations and corporations are forced to compete in the global market for knowledge. As the rate of turnover in new scientific knowledge has increased, so has the cost, and risk, of engaging in leading-edge research. In addition, the gap, and even the distinction, between basic research and technological application has been significantly reduced, particularly for certain high-technology fields (Brown, 1985, pp.113-134). The resulting necessity of selectivity in funding has placed allocation decision mechanisms under greater scrutiny.

Thus science and technology (S&T) policy makers have come under increasing pressure to make the funding of research more systematic and accountable and to support more research that results in economically valuable technical applications. Major R&D-based corporations now share similar concerns to those of governments for maintaining scientific and technological advantage, performing more and more "strategic" research in an environment of increasingly scarce resources. Funders of research, whether in government or industry, are faced with a paradox: a need to be conservative in the allocation of R&D dollars and a drive to pursue high risk innovative opportunities that promise to payoff. Both motivations heighten the need for more reliable sources of intelligence about international research developments and improved techniques for aiding resource allocation decision-making.

A variety of quantitative approaches have been forwarded for these purposes. Perhaps the most methodologically sophisticated quantitative method that has been developed is co-citation bibliometric modeling. The methodology produces a detailed, descriptive "model" of the structure and content of the scientific research front. More than just a set of scientific-literature-based indicators, a bibliometric model provides a partial "mirror" of the 'actual' structure of science.

Bibliometrics were first used by information scientists and historians and sociologists of science as a way of analyzing the evolution of scientific knowledge. More recently, the method has been explored by S&T policy analysts and a few industrial R&D managers seeking to monitor developments in science. Bibliometric models are claimed to provide an empirically objective surrogate for international peer consensus about what is most important in current research. As a result, funding agencies in seven nations have undertaken exploration of bibliometric models for potential application to S&T policy objectives.

Given the potential power, and hence economic significance, of bibliometric models, it is not surprising that their development and evaluation has occurred in an environment made up of decidedly interested parties. In regard to its development, bibliometric modeling is in effect a new information technology undergoing refinement with the aim of capturing a market. There is competition in its development between two U.S. companies, the Center for Research Planning (CRP) and the Institute for Scientific Information (ISI).

However, it is the competition in the information technology market, between bibliometric models and other quantitative and semi-quantitative research planning and evaluation methods, and the concern of those subject to science policy decisions which have added controversy to the development and application of bibliometric modeling.

The claim that it is possible to derive an accurate and objective picture of the dynamic, cognitive structure of science has been partially validated in early studies that compared bibliometric data to peer assessments. These and subsequent validations led to some inflated claims for the method, and, more significantly, to inflated expectations on the part of some policy makers. Sociologists and historians of science quickly pointed out the social and institutional complexity of research activity and the limitations of quantitative methods in representing these factors. Over interpretation of relatively simple S&T indicators and the prospect of misuse of more sophisticated methods then led to cautiousness on the part of policy makers. For the same reasons, components of the scientific community expressed concern, sometimes reacting defensively to the application of any quantitative decision-aid by 'outsiders', be they non-scientists or scientists in policy positions.

As a consequence of these controversies, if the full utility of bibliometric models (and of quantitative methods in general) is ever to be realized, then knowledge and application claims will have to be 'negotiated', in both cognitive and political terms, between the producers of the methods, the S&T policy community, and the scientific community. That negotiation must be based on a thorough knowledge of the capabilities, limitations, applications and interpretations of the methodology.

This paper seeks to advance the process. Sections 2 and 3 give a conceptual and methodological overview of bibliometric modeling. Section 4 raises and clarifies the major unresolved issues concerning validation, methodology and utilization for policy purposes. Section 5 serves to illustrate policy applications and interpretations of bibliometric model data, using recent national policy studies in Great Britain, the United States, Spain, Sweden and Australia as examples. That section also shows how difficulties and limitations encountered in those studies have led to important innovations in bibliometric modeling. Section 6 details the methodological innovations that have been made by both CRP and ISI, summarizing the state-of-the-art according to each, and compares the rival methodologies relative to issues of accuracy, objectivity and policy utility. The final section proposes potentially fruitful topics for future research. Many of the observations are drawn from experience in an ongoing program in Australia for the exploration of bibliometric models in which the authors are involved[1].

This paper is written with two audiences in mind: those who are largely unfamiliar with bibliometric modeling, and those who already understand the basic methodology and its applications. For those in the former group, Sections 2, 3, 5 and 7 will serve as an introductory "handbook" and the other sections may be of less initial interest. Sections 3, 4, and 6 are intended more for those in the second group.

2. Conceptual Overview of Bibliometric Models

Most simply, a co-citation bibliometric model is:

> A detailed representation of the structure and content of the international research front based on the strongly shared patterns of referencing among the current scientific literature papers.

A co-citation bibliometric model basically does two things. First, it defines the coherent research problem areas that appear to be the basic units of scientific activity. It does this by classifying and grouping current scientific literature papers by their common referencing to clusters of highly cited and highly co-cited previous papers (hence "co-citation bibliometric"). Second, it measures the interaction between the research problem areas defined in this way in order to portray the current research front as a hierarchical network of interacting specialty-, subfield- and field-sized research areas (hence "model").

The technique used for doing this is called "co-citation analysis". Two reference papers are said to be co-cited if they were cited together in the reference list of the same current paper (Small, 1973, pp. 265-269; Marshakova, 1973). If many current papers co-cite the same pair(s) of reference papers, then there is a high degree of referencing concensus about the connection between those older papers (even if they represent conflicting theories) (Small, 1974, p.40). This measure of scientific consensus is used to define co-citation clusters.

The basic unit of a bibliometric model is the co-citation cluster. Each cluster has two components: a set of highly cited and highly co-cited reference papers, called the "base literature" or "core" and a much larger set of current-year papers that referenced those, called the "current literature". It is the reference papers that are actually clustered. The resulting base literatures represent the cores of theories and methods around which the current research is organized. The current publications that reference a base literature share a common research focus, theoretical approach or methodological concern. The current literature therefore provides the most useful policy information, because it describes the current research front activity on the problem-area at the time period of the model. The clusters typically consist of between 25 and 150 current papers on average. A multidisciplinary, multinational model may define between 9,000 and 30,000 such clusters.

The bibliometric modeling method uses co-citation analysis to extract only the most strongly shared referencing patterns from the international scientific literature. These patterns represent an informal consensus among the publishing scientific community. This consensus supersedes the choices of individual scientists, and the perspective provided by a bibliometric model often transcends that of any single national or organizational peer group (Gilbert, 1977, pp.115-119; Garfield, et al., 1978, p. 184)[2].

For this reason, the models are considered representative of the international peer consensus about the most important research problems and developments in science.

Co-citation analysis operates directly on the scientific literature data to reveal a structure that is inherent in those data, not imposed by the methodology. Bibliometric modeling therefore relies on criteria that are "internal" to scientific publication activity as opposed to being arbitrarily imposed according to criteria that are "external" to science. The co-citation clusters have been shown, on the basis of exploratory studies and validations against expert judgement, to represent with reasonable accuracy the 'actual' specialty units of which science is composed at the cognitive level (Small, et al., 1974; Small, 1977; Small, et al., 1979; Small, et al., 1980; Griffith, et al., 1972; Griffith, et al., 1974; Mullins, et al., 1977; Sullivan, et al., 1977; Garfield, et al., 1978; Coward, et al., 1984; Franklin, et al., April, 1986). For this reason, the co-citation clusters are often simply called "specialties"[3].

These, then, are the most important features of a bibliometric model: it extracts a primarily cognitive structure directly from the referencing consensus of the currently publishing scientists about the loci of the most important research problem-areas at a given time.

After the specialty clusters are defined, the interaction between them is used to define the structure of the current research front. Specialties are clustered on the basis of co-citation linkages between their base literatures in the same way that reference papers were clustered to form specialties. (The inter-specialty linkage strengths are by definition weaker than the intra-specialty linkages.) A matrix is generated that contains all the linkage strengths between all of the specialty clusters. The mathematical technique of multidimensional scaling is used to graphically represent the resulting set of relationships as a 'map', where clusters of highly interacting specialties are 'towns' and weaker linkages between clusters are 'roads' (de Solla Price, 1965; Small, et al., 1974; Kruskal, 1964; Shepard, 1962). Specialty clusters that share a related topic or similar intellectual focus automatically cluster together to form subfield-sized groupings, sometimes called "regions", and these in turn cluster to form field-sized groupings. Groups of related research areas are now commonly represented by "network diagrams", in which the proximity between them is inversely proportional to the lengths of the lines connecting them.

In effect, a bibliometric model is a current, detailed research classification system. It classifies specialty areas into larger subfields by a measure of the extent to which they draw on the same knowledge base. This structural feature of bibliometric models provides those concerned with the current state of science with an important alternative to standardized classification systems, which are designed more for adminstrative convenience than for reflecting the structure of the science. Without this classification, which is unavailable any other way, policy makers and R&D managers must rely on gross and outdated disciplinary categories for organizing research administration and funding, while the most important new research may be going on outside the domain of the old categories and therefore remain unnoticed (Cole, et al., 1978, p.220)[4].

Most importantly, bibliometric models place each research area within its intellectual and, to a more limited extent, social context. Research areas are not viewed in the isolation commonly enforced by unavoidable limitations of perspective. This means that it is possible to recognize intellectual bridges between previously unrelated research areas or new developments in science that had not been recognized. Validation exercises have shown that bibliometric models can identify inter-disciplinary developments in science that were not intuitively obvious to some technical experts, but were then recognized as meaningful by those specialists (Healey, et al., p.244; Coward, et al., 1984). This offers a considerable capability to S&T Policy and R&D Management, given the importance of new developments and inter-disciplinary research to innovation in science.

Because the bibliometric model describes a hierarchical structure for science (Small, et al., 1974, p.38), i.e., specialties within subfields within fields, the analyst can select the level of data aggregation to fit the policy question. Figure 1 gives an example drawn from a small, 1982 model designed to cover epitaxy and microlithography research within Solid State Physics (Coward, et al., 1984). Primary region 2 is shown to consist of 13 interacting subregions. One of those, subregion 2.3, contains 6 discrete but closely related specialty problem areas. The analytical focus can be shifted downward, if more detail is needed about individual specialties, organizations or researchers, or upward, if a broader perspective is appropriate.

All research areas are characterized by an array of statistical and demographic data that measure the size (i.e., total international activity), degree of coherency, location in the structure of the current research front ('mainstream' versus 'peripheral') and rate of development of research areas, and identify the active nations, organizations and researchers. Research areas are indexed by author names, the names of sponsoring or home organizations, and the significant words and phrases that appear in the titles of the current papers, allowing the model to be entered from a variety of vantage points and analytical objectives. The policy utility of these data is being investigated in the United States, Australia, Spain, and the Federal Republic of Germany, among other countries. These are some of the applications that have been investigated:

- characterizing concentrations of national research 'strength' and 'weaknesse' relative to national standards and to international patterns;

- assessing patterns of international collaborative research in fields of interest and selecting optimal partner-nations for cooperative research in designated subfields;

- profiling an industrial, academic or governmental research performing organization, pointing out its implicit (and often unrecognized) research priorities, potential 'holes' in the organization's research agenda, potential innovative opportunities, and the activity of competitors or potential collaborators;

- assessing the outcome of a funding scheme or research program;

- identifying potentially 'hot' research areas using an indicator of the rate of turnover of new knowledge in specialties and regions;

- tracking the development of a specialty or subfield over time using trend data from sequentially built models -- rapid changes may indicate microrevolutions in the science;

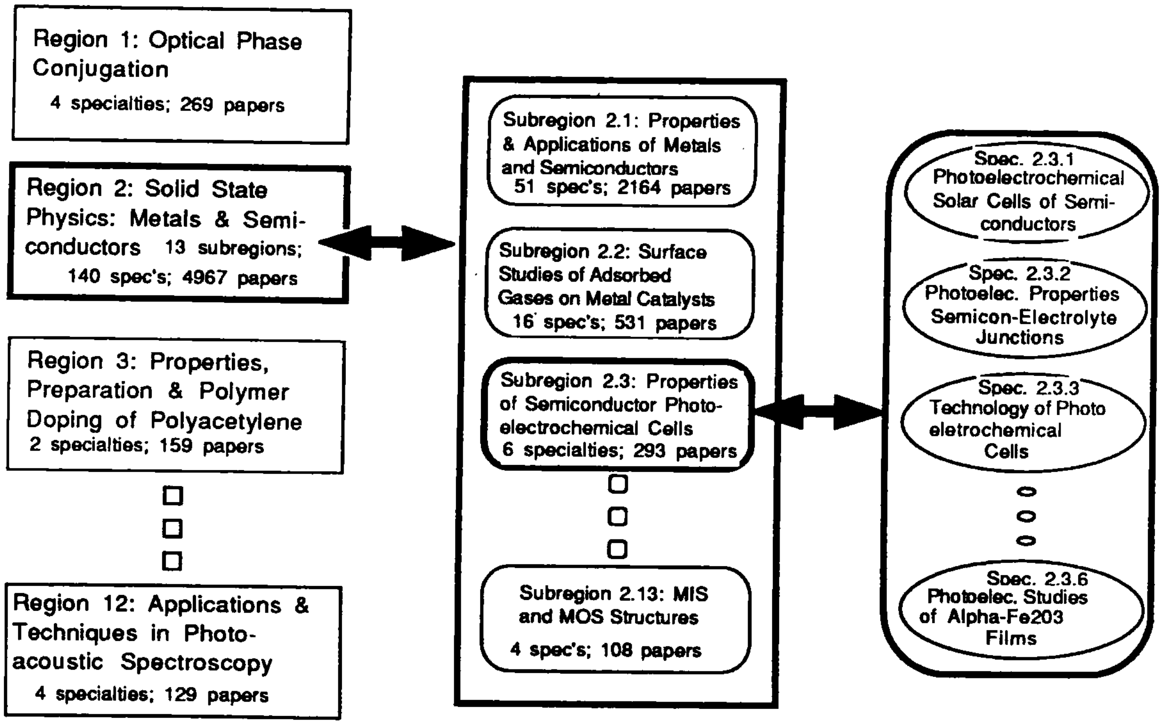

FIGURE 1

Hierarchy of Research Areas in a 1982 Model of Solid State Physics

- identifying where changes in national research activity were in line with or counter to the international patterns;

- profiling the research front activity in a subfield of policy interest, identifying research topics that may deserve special attention;

- selecting scientists for international review panels in subfields under investigation, and supplying background information and an international perspective to national peer panels involved in research planning and evaluation;

- identifying potential "strategic" research opportunities at the interface between science areas which have already been applied to technical problems (according to expert judgment or linkages between specialties and patent data) and those which appear to be 'untapped'.

3. Methodological Overview of Bibliometric Modeling

In 1974, ISI first began clustering parts of the SCI database using the co-citation methodology developed by H. Small (Small, et al., 1974). The modeling methodology developed at that time was used up until 1983-84, first by ISI and later by CRP. There are five generic steps to the approach, as illustrated by Figure 2. These steps are summarized below. They have been detailed at length elsewhere (Small, et al., 1974, p.20; Garfield, 1979, pp.98-147; Small, et al., 1983; Coward, et al., July 1984).

The first step is to select the origin database(s). Building a bibliometric model requires a citation-indexed database, one that indexes references to source documents and vica versa. All bibliometric models have been built primarily with data from ISI's Science Citation Index (SCI) and Social Science Citation Index) (SSCI) databases, which are the largest available multi-disciplinary citation-indexed databases[5]. The result is that all bibliometric models share the biases of these databases in favor of the English language and toward more established scientific journals[6].

The second step is to build the input database. The approach differs, depending on the type of model to be built: a "field model", "national model", or "global model". Field and national models were the first types to be made commercially available. Though ISI has been conducting annual cluster analyses of the entire SCI since 1974-5, it was not until 1983 that both ISI and CRP began building multidisciplinary, multinational (i.e., "global") models for distribution[7].

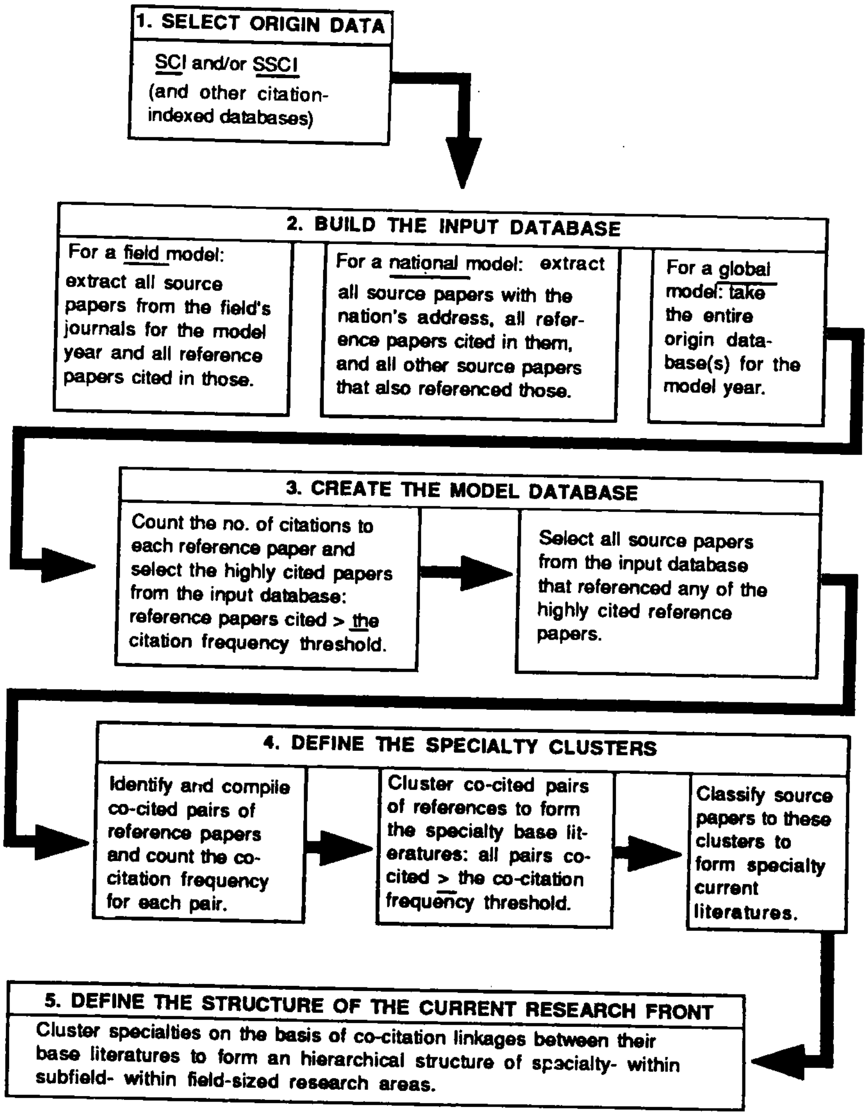

FIGURE 2

Model Construction: Summary of Generic Methodology

As the title implies, a field model is designed to cover the content of a single field as specified by the client. The "paradigm journals" for the field are identified by specialists, usually nominated by the client. They become the "core journal set." This set is expanded through an analysis of journal citation linkages (Narin, et al., 1972; Garfield, 1979, pp.148-239). The result is a "final journal set" that supposedly contains all articles of relevance to the field. The difficulty is in deciding where to limit the journal set, since the exclusion of a single journal may greatly alter the resulting definition of the field.

All current or "source" articles and all the references in those articles are extracted from the journal set for the year of the model. These references become the input data for the base literature component of the model database. All current papers in the origin database(s) that share a single reference with any of the journal set source papers are identified and compiled, and this combined set becomes the input to the current literature component of the model database[8].

A national model is a multidisciplinary model built from the perspective of a single nation. Instead of building the input database by starting with a journal set, the method starts by identifying all articles published in the model year that carry an address from the country of interest. All other current articles in the origin database(s) that shared a single reference with any of those national-origin papers are added to the current literature component. Only reference papers found in the national-origin papers are used as input to the base literature component of the model database. The same approach could be used to build a bibliometric model around the research activity of a single corporation.

The result is a nationally skewed global model. The view of global science is skewed in that if the target nation was not active in a particular subfield (i.e., no scientists from that nation published a paper in that subfield during the model year that also made it into the model), then the model is unlikely to contain any information about that subfield[9].

The third step is to create the model database. Only highly cited reference papers from the input database and the current papers that referenced those qualify for potential inclusion in the bibliometric model. The number of citations received by each reference paper is counted, this count is normalized in one of several ways to compensate for differences in publication and citation practices between fields, and a citation frequency threshold is set. (The techniques used for normalizing citation frequencies and determining the citation threshold are discussed in Sections 4 and 6).
Only reference papers cited above the normalized citation frequency

threshold are retained for the model database.

The purpose of the citation threshold, and the co-citation threshold discussed below, is to separate the meaningful referencing 'signals' from the 'noise'. Less than 5% of the hundreds of thousands of references in the input database eventually qualify for inclusion in the model. It has been shown that roughly 35% of all documents are never cited anywhere, and only 1% are ever cited 6 or more times (de Solla Price, 1965, pp.510-515). From the perspective of bibliometric modeling, the large majority of published research represents transitory statistical 'noise' from which the meaningful patterns of referencing consensus must be extracted using co-citation analysis.

The fourth step is to use these patterns to define the specialty clusters. This is done by first clustering the highly cited reference papers using co-citation analysis to form the base literature cores of the specialties, then classifying the current papers to those clusters by their references to them. All co-cited pairs of references, called reference-reference pairs (R-R pairs), are identified, the co-citation frequency for each R-R pair is counted, those counts are normalized by a mean of the citation frequencies of the individual references in the pairs, and a co-citation frequency threshold is set. Only R-R pairs co-cited above the threshold level are retained. A co-citation threshold of 10 means that the two references must have been cited together in the reference lists of 10 or more different current papers. Some R-R pairs are co-cited as many as 100 times.

The highly cited and highly co-cited R-R pairs are then clustered to form the specialty base literatures. All single pairs or pairs of pairs that form above the set co-citation threshold become specialty base literature cores. A "single-link" clustering algorithm is used: all R-R pairs that qualify and that share a single reference are linked (Sneath, et al., 1973, pp.188-308; Garfield, et al., 1978, p.187). For example, the pairs R1-R2, R2-R3, R3-R1, and R4-R2 will all cluster together to form the core for a single specialty. This sometimes results in the "chaining" of a very large number of R-R pairs and the formation of what are called "macro-clusters" or "macro-specialties." In contrast, some R-R pairs remain uncluster with other pairs.

Once the base literature cores are clustered, current papers in the model database are assigned to specialty clusters. A current paper needs only one reference to a base literature to be considered part of the current research in that specialty area. One result is that some current papers are assigned to more than one specialty[10].

A second is that some current papers can be seen to have many more references to a base literature than others, i.e., they are more central to the focus of that specialty cluster.

Through this process, some number of distinct but interacting specialty clusters of varying sizes are formed. Each consists of two components, a base literature of reference papers and a current literature of source papers.

The fifth step is to build the model by clustering the specialties to define the structure of the current research front. This is done by effectively lowering the co-citation threshold above which R-R pairs must qualify in order to cluster. Inter-specialty co-citation linkages between different base literatures pull related specialties together. Some specialties will be isolated or linked to only a few other specialties, while some specialties will be linked to many other specialties. For a typical model built using the original methodology, the result is a very large, central core of highly interactive "mainstream" specialties surrounded by some number of smaller "satellite" clusters of interactive specialties, themselves surrounded by many "peripheral" specialties and groupings of two to four specialties (Small, et al., 1979, p.455). The central core is then broken up by reclustering it at a slightly higher co-citation level. The result is a hierarchical view of the structure of the current research front in which specialties are seen to belong to subfields and subfields to fields.

An optional sixth step is the definition of research "trends" or cluster "strings". This requires having two or more models built for sequential time-periods using comparable methodologies. Specialty clusters involved in trends are those whose base literatures continued completely or partially intact from one time-period to the next (Coward, 1980; Small, et al., 1983; Coward, et al., 1984). The degree of "stability" of the base literature of the specialty over time is taken as a measure of the continuity in the scientific peer consensus about the intellectual focus of that research area (Small, 1977, p.143). The success of the method depends on the comparability of clustering methods used for the sequential models. If the methods are not comparable, then it is difficult to tell whether shifts are due to differences in method or actual developments in the science.

4. The Major Unresolved Issues

Three types of unresolved issues exist concerning bibliometric models: 1) the issue of the validity of the models as representations of scientific activity; 2) issues concerning the objectivity, accuracy and efficiency of the model building methodology; and, 3) issues related to the policy utility of bibliometric model data.

The three types are obviously related. The useability of any quantitative method depends on its validity. Validity is partially an issue of accuracy and objectivity, but demonstration, recognition and acceptance of validity, as well as utility, is a social and political process. Explorations of the validity and utility of bibliometric models have taken place in a highly politically charged environment. There has been a tendency for both expectations for and criticisms of bibliometric models to be inflated.

In the first place, the commercialization of bibliometric models has occurred at a time when policy makers have come under increasing pressure to make research funding decision making more systematic, productive and accountable. Under pressure from higher and more politicized levels of decision making to provide unequivocal decisions, some policy analysts have pressed S&T indicators to be decision algorithms. This has led to justified criticism of the use of quantitative methods and cast particular suspicion on more sophisticated methods such as bibliometric modeling.

Second, attempts to establish the credibility of bibliometric models have taken place at a time when the notion of the "sovereignty" of science has been forced to give way to the economic and social requirements of public support for R&D. This, the debate over the efficacy of the peer review system (for example: Cole, et al., 1981; Porter, et al., 1985), and the potential for misuse of quantitative methods have led the peer community to express serious reservations about the use of any quantitative methods by non-scientists or scientists in policy positions.

Third, the arena has been politicized by competition between producers and users of different quantitative methods. Some have found it necessary to defend their use of certain methods by attacking the use of other methods. The front for disagreement about which methods are most valid and useful has generally formed along national borders, with analysts in one nation criticising methods being explored in or exported from other nations[11]. This has served only to confuse validation efforts and to slow development of a better understanding of the complementary capabilities of different quantitative methodologies.

These factors have speeded as well as pressurized the development of bibliometric models (and other quantitative and semi-quantitative approaches). Development of the methodology has been partially a process of trial-and-error. Each new CRP bibliometric model has incorporate innovations that resulted from previous studies. The two companies that build bibliometric models, ISI and CRP, now offer two very different types of models, though the basic application of the co-citation clustering method is essentially the same. Changes in the modeling method and the proliferation of types of models have both benefited and confused the consumers. Clients have heard the announcement of an important new innovation just as the old product was being delivered. The initial purchasers of models have been left in a position not unlike that of the first purchasers of personal computers, and subsequent buyers have been unclear about what exactly they were getting. Part of the objective of this paper is to clarify those developments.

The rest of this section attempts to raise and clarify the major unresolved issues.

4.1 Concerning Validity

The validity of bibliometric models has been critically examined from two perspectives. It has primarily been considered as a question of the accuracy with which the models match the 'actual' structure of science. The only authoritative baseline concerning this is the judgment of scientific peers. Secondly, the validity of the specialty clusters has been examined relative to definitions of "specialty" in the Sociology and History of Science.

Thus far validation oriented bibliometric model studies could be considered generally successful, particularly given the politically charged environment in which they have been conducted. The models have been subjected to the most extensive and detailed validation exercises of any quantitative method. Validations relative to peer judgment have not been wholly unambiguous, but they have been largely positive and convincing. Early studies by both information scientists and sociologists of science successfully demonstrated that the specialty clusters accurately locate/represent the specialties of science. The two most systematic and thorough validation exercises within policy contexts, the British ABRC study and another CRP study performed for the Swedish government, both produced evidence in favor of the validity of the method (Coward, et al., 1984; Healey, et al., 1985; Franklin, et al., 1986).

The debate about the validity of bibliometric model specialty clusters as representations of the 'actual' specialties of science is a old one (see Chubin, 1985). The debate originally resided within the Sociology of Science, but the commercialization of bibliometric modeling moved the debate into the realm of Science Policy Studies. We contend that the debate represents a confusion on three different levels: 1) a general confusion about the relationship between the social and cognitive aspects of scientific research activity; 2) a confusion about what bibliometric model specialties actually represent; and, 3) a confusion between the value perspectives of sociologists of science and S&T policy analysts.

The earliest bibliometric studies tended to equate the social and cognitive aspects of scientific research activity (Price, 1965; Crane, 1972). This assumption has haunted the development and validation of bibliometric modeling. Initial demonstrations of an ability to accurately locate the specialties of science were interpreted as a claim that the specialty clusters describe all aspects of the scientific specialty, social and cognitive, and all research related to each specialty. Sociologists and historians of science reacted on the basis of this interpretation, pointing out the social complexity and institutional control of research activity (Whitley, 1977; Edge, 1979), while bibliometricians (as they were called) adopted the basic theoretical and methodological foundation provided by early research, because it was strongly supported by empirical evidence. Both sides were right.

A co-citation cluster does not describe all aspects of the scientific specialty. First, the models are built from scientific literature data, and publication in juried periodicals is perhaps the most formal and restricted mode of communication in science. Secondly, the modeling methodology is specifically designed to filter out the research about which there is lower levels of referencing consensus. What the earliest co-citation cluster studies showed was that the method classifies papers on the basis of shared intellectual interest and thereby defines specialty clusters with coherent cognitive foci (Small, 1973; Small, 1974; Small, et al., 1974; Griffith, et al., 1974; Small, 1978). Subsequent studies observed that specialty clusters do not describe all research related to a problem area, but that they do accurately locate and characterize those problem areas (Sullivan, et al., 1977). The research of the past five years supports this interpretation of the bibliometric model specialty: they are partial but accurate indicators of the 'actual' specialties of science.

Sociologists of science and science policy analysts respond to the partialness or selectiveness of bibliometric models from two very different value perspectives. Sociologists and historians of science focus on the social and institutional aspects of science. The utility of bibliometric models for S&T Policy rests on the fact that they transcend the 'local', social perspective and reflect the international context within which science operates at the cognitive level. Publication in juried periodicals may be the most formal and restricted mode of communication in science, but it is through this mechanism that science appears to legitimate its findings, thereby filtering out the less significant and the unsubstantiated[12].

Such a filter may be desirable from the perspective of S&T Policy,