

Perceptions of the Business Knowledge Resource

April 2002

John L. Gordon: Applied Knowledge Research Institute, Blackburn Technology Management Centre, Challenge Way, Blackburn. BB1 5QB

Abstract:

A process identified here as Structural Knowledge Auditing, has evolved over five years and has now been applied successfully in a range of businesses and to a range of knowledge areas. The process involves collecting information about the structure of knowledge through interviews and then mapping this knowledge based on learning dependency to create a visualisation of a knowledge resource. Parameters are used to provide management information about the knowledge elicited. Values for these parameters are elicited during standard interviews. Information about the knowledge structure, the parameter values and other more general comment is used to compile a business report and recommendations concerning the development of the knowledge resource.

Ten sets of records of actual business projects have been used in this study to evaluate the validity of the approach and to consider generalised implications from this work. The study provides some evidence concerning the importance of the business knowledge resource, the availability of knowledge within a business, the views on knowledge acquisition adopted by various professional groups and suggests general approaches to the management of risk associated with the knowledge resource.

Final conclusions from this study claim that the methodology used leads to consistent and useful knowledge management support.

Acknowledgements:

Thanks to Martin Holden who has worked on each of the Structural Knowledge Audits that have contributed to this study.

1. Introduction:

The process known as Structural Knowledge Auditing (SKA) has now been used successfully in a number of organisations and applied to a range of knowledge categories or business sectors (Gordon 1999). The basis for this work is interviews with knowledgeable staff in order to discover and map a particular part of the business knowledge resource. Interviews are used to elicit the structural map of the knowledge area linked by 'learning dependency' and to elicit values for various parameters that are intended to provide useful information about the knowledge and the role that it plays in the business environment. The structure of the map and the parameter values, form the basis for an analysis of the knowledge resource leading to a more objective and explicit framework in which knowledge can be managed.

For each SKA, results (values and structure) are analysed collectively to generate recommendations for development of the knowledge resource. It is recognised however, that individual parameter values are subjective in that they have been derived by asking knowledgeable interviewees to assess the parameter and move a slider (0 .. 10) to qualify the parameter within the context of the knowledge area being studied. This paper considers all results and parameter values from 10 separate Structural Knowledge Audits. The intention is to examine the consistency with which values have been given and to consider implications of this overall analysis to the business knowledge resource in general.

1.1 The Range of SKAs undertaken

Table 1 provides a list of the SKAs that are to be considered in this study. The actual company has not been identified but the size of the company and the business area within which it operates is given.

	AUDIT	SIZE	START	END	TOPIC
1	Aerospace	Large	05:11:98	25:11:98	High Tech Fabrication
2	Prototyping Business	Small			Engineering Prototyping
3	Utility company	Large	12:03:99	26:03:99	Safety
4	Engine Manufacturer	Large	24:01:00	17:02:00	Business Winning
5	Business Consultant	Small	10:04:00	14:04:00	Consultancy Activities
6	Off-licence Retail Group	Small	09:05:00	06:06:00	Off-licence retail
7	Industrial Doors	Medium	07:06:00	05:07:00	Industrial Door Repair & Service
8	Computer Interface	Medium	20:10:00	14:11:00	Engineering Services
9	Computer Interface	Medium	18:12:00	06:02:01	Sales & marketing
10	Computer Interface	Medium	16:02:01	27:03:01	After Sales (installation, maintenance)
	Computer Interface	Medium	02:04:01	04:05:01	Cross Audit Analysis
11	Hotel	Medium	25:06:01	13:07:01	Hotel Operations

Table 1: Studies undertaken

It is suggested that the range of topics and the variety of businesses studied is sufficient for a useful analysis of common features and topic specific elements to be identified.

1.2 The scale of the source data

Each study from table 1 has involved senior staff with experience in their particular knowledge domain. Interviews were carried out by one experienced pair with a sound background in this area of work. All interviews were carried out within a business contract environment and non in an academic or learning environment. Each study undertaken lead to a business report and concrete recommendations concerning the knowledge resource.

	PEOPLE	INTERVIEWS
1	6	8
2	2	3
3	6	7
4	8	10
5	4	4
6	9	9
7	7	9
8	9	11
9	6	7
10	10	11
	4	1
11	7	7
Total	78	87

Table 2: Interview Data

Table 2 provides data for the individual (or pair) interviews. Each SKA also involved a group interview that is not shown in table 2. All of the people working on each project attended the group interview. Records for interviews 1 and 3 are approximate because proper records were not kept at the time. All interviews were about 2 hours duration.

Table 2 shows that approximately 78 people have been interviewed within this whole study and including 9 group interviews, approximately 96 interviews carried out.

Audit	Nodes	Arcs	Params	Pos Param	Arcs per Node
1	98	118	220	392	1.204
3	64	76	147	256	1.188
4	71	92	210	284	1.296
6	98	122	372	392	1.245
7	71	95	276	284	1.338
8	152	190	604	608	1.250
9	125	159	496	500	1.272
10	105	131	416	420	1.248
11	112	167	444	448	1.491
Total	896	1150	3185		
Average	99.56	127.78	353.89		1.281

Table 3: Map statistics for SKAs

Table 3 shows the number of knowledge nodes elicited for each study and in total (this excludes unrecorded data for audits 2 and 5). It also provides

similar information for arcs and parameters. In later work, the number of parameters elicited is 4 short of the maximum possible because there are 4 parameters elicited per node and the main or root knowledge node is left blank. Earlier studies did not elicit this information completely.

The arcs per node value is computed simply by dividing the number of arcs on the map by the number of nodes and is a crude measure of map complexity.

1.3 Details of the data captured

There are three main areas where data is captured during each SKA. These areas are:

- 1) The connectivity of the map
- 2) The parameters and their values
- 3) Comment from experts during interviews

The derivation of each knowledge node and the creation of map structure during interviews has been dealt with in earlier publications (Gordon 2000). Essentially a knowledge structure map contains boxes (nodes) that represent pieces of knowledge. Arcs connect from these parent nodes to child nodes that should already be known as a prerequisite of fully understanding the parent knowledge. When structure is developed in this way, it is possible for complex connectivity to develop as some knowledge is shown to be an essential prerequisite of several other knowledge areas.

Figure 1 shows an actual knowledge structure map derived from a commercial study. The figure does not show any detail, it simply serves as a visual representation of the above description.

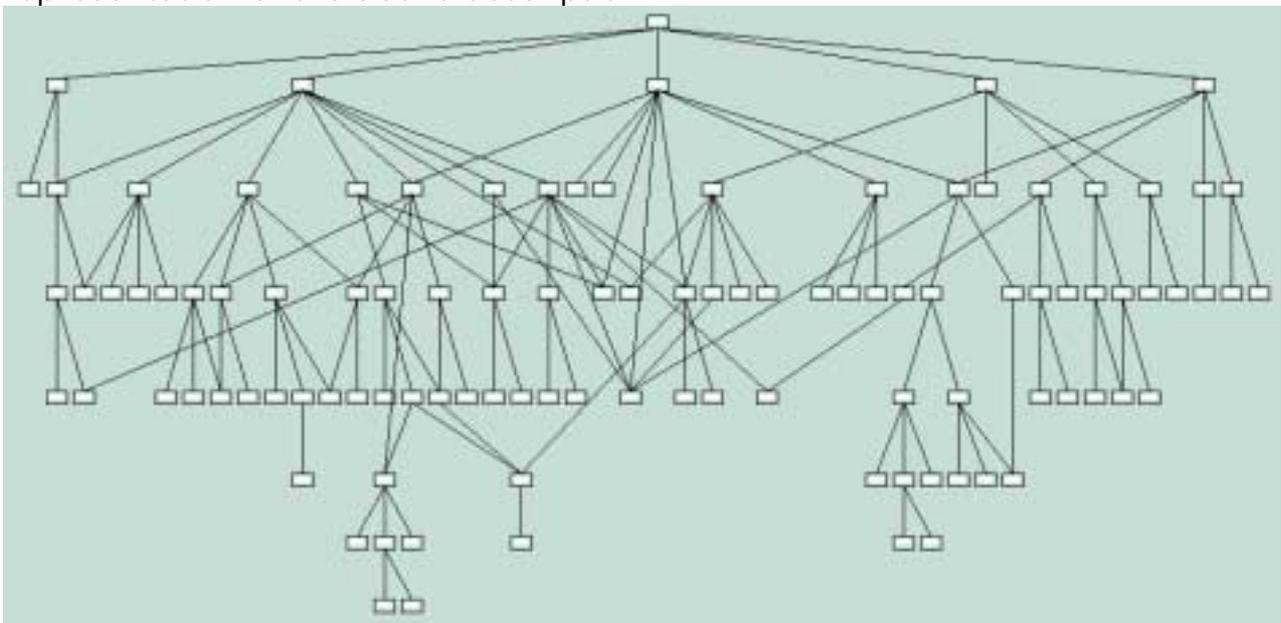


Figure 1: A Knowledge Structure Map

Earlier publications have also discussed the form and nature of the parameters elicited (Gordon 2000) but this information will be repeated here because it is central to the study being described in this work.

Parameter	Description	Low (0) value	High (10) value
Importance	Importance to the knowledge area being studied	Not important at all	Critical knowledge
Difficulty	Difficulty to replace the knowledge	Easy, readily available	Almost Impossible, very difficult
Study-Experience	How the knowledge is acquired	All Study	All Experience
Known By	How many people from this area have the knowledge	Less than 10%	All of them

Table 4: Description of parameters elicited

Table 4 identifies the 4 main parameters used in the studies along with descriptions and value assignment criteria.

During each interview, written records are also maintained concerning any key features that may influence the particular study within a knowledge area. These comments are reviewed along with the connectivity and parameter data during each audit. Although such comment is very useful in an individual audit it has less value in a comparative study.

2. Analysis of Results

Results from all audits will be presented in a way that aids a comparative study. Implications of these results will be considered in the next section.

2.1 Map Connectivity Data

A summary of map connectivity data is provided in table 3. The average size of a structural knowledge map is 100 nodes and the average number of arcs is 128. Comments from map users in business suggest that this is about the right size for a useful knowledge structure map. If the map was a great deal larger (more nodes) and the connectivity remained at a similar level then the map would become too difficult to use effectively. It has been suggested that several maps of this size would be better than one large map.

The average value for arcs per node, stated in table 3 is 1.281 (connectivity). When it is realised that a map often starts off at one root knowledge node and terminates with many leaf nodes, this level of connectivity can lead to quite a complex map. Certainly the map with the highest value (1.491) was a complex looking map that required quite a lot of effort to organise the layout to create a reasonable visualisation.

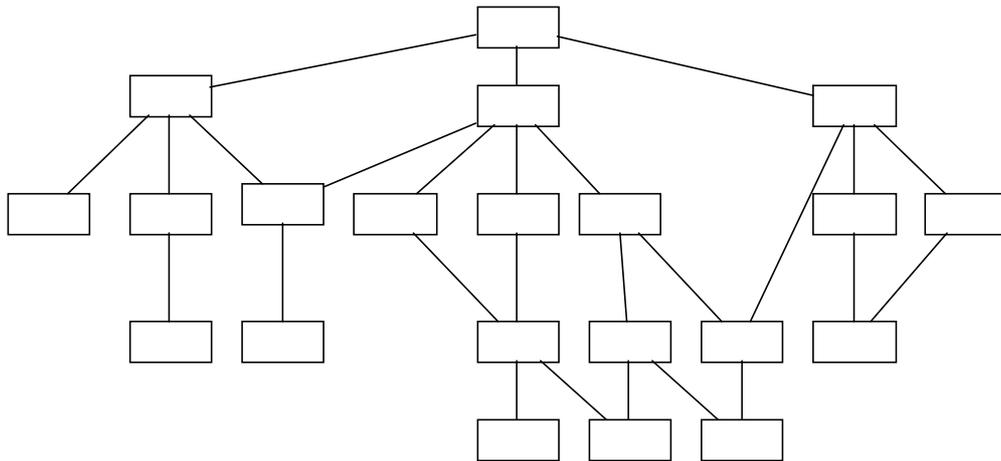


Figure 2: A Typical Knowledge Structure Map

The example map shown in figure 2 contains 21 knowledge nodes and 26 arcs. This gives a connectivity value of 1.238 (26/21). Just one more arc would yield a connectivity value of 1.286 and five more arcs would provide a connectivity value close to the maximum value in table 3.

Connectivity (as defined here) is not the only thing that makes maps look complex but it is a significant contributory factor. Other factors include knowledge nodes that are prerequisite of many parent nodes at several different levels in the hierarchy. One node at the lower right of figure 2 can be seen to have parents at more than one level above. Another factor is where child nodes have parent connections that originate from diverse sources that span the whole knowledge structure map. This happens regularly when fundamental or central knowledge areas are identified.

2.2 Size of maps

It should be noted that all of the studies featured in this work have taken about the same length of time in the data gathering phase. This is one explanation for the similar size of structure maps. It can also be seen however, from records of node growth (figure 3) that the rate of growth slows significantly during later interviews with more time spent on validation. This suggests that conducting more interviews over longer periods would not lead to significantly larger maps but probably to more accurate ones.

Validation is an important part of the interview procedure but it is also necessary to let each interviewee know what is currently on the map before they make a further contribution to it. This is part of the methodology for the elicitation of knowledge structure but it can be seen that it also has an effect on map size.

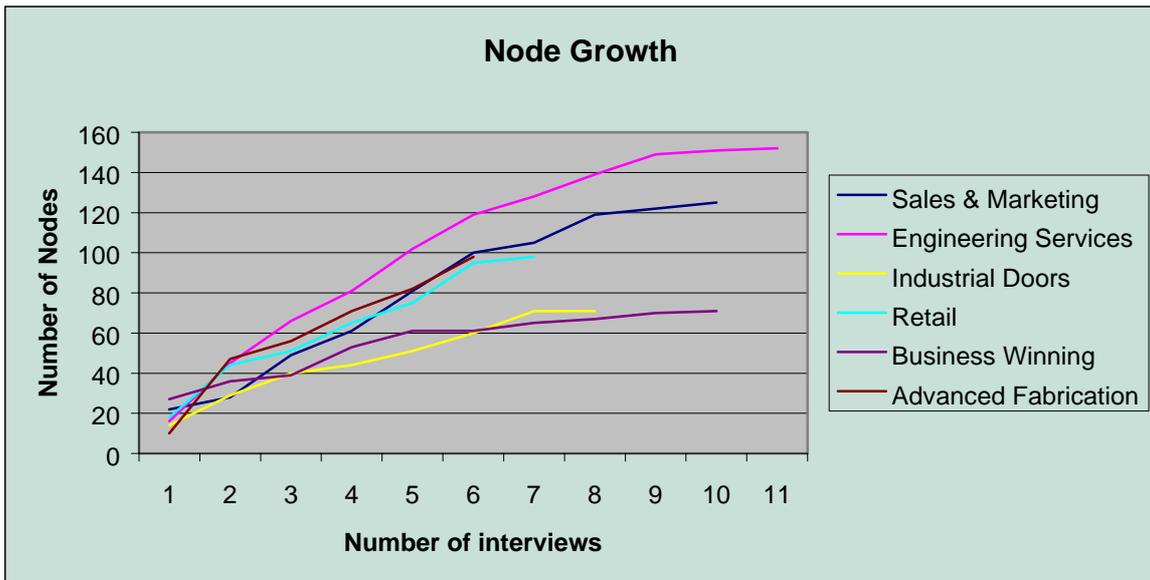


Figure 3: Record of node growth during interviews

2.3 Parameter Values supplied during interviews

The parameters that were elicited during interviews are introduced in section 1.3. In order to help interviewees assign parameters it was often necessary to list values already given so that they could see how the new value would reflect the relative position within the context of the more general knowledge area. In addition to this, some validation was carried out with different interviewees.

2.3.1 The Importance Parameter

Figure 4 shows all of the parameter values for importance for each SKA, plotted on the same axes. This has been done to estimate similarity and difference rather than to consider any particular study.

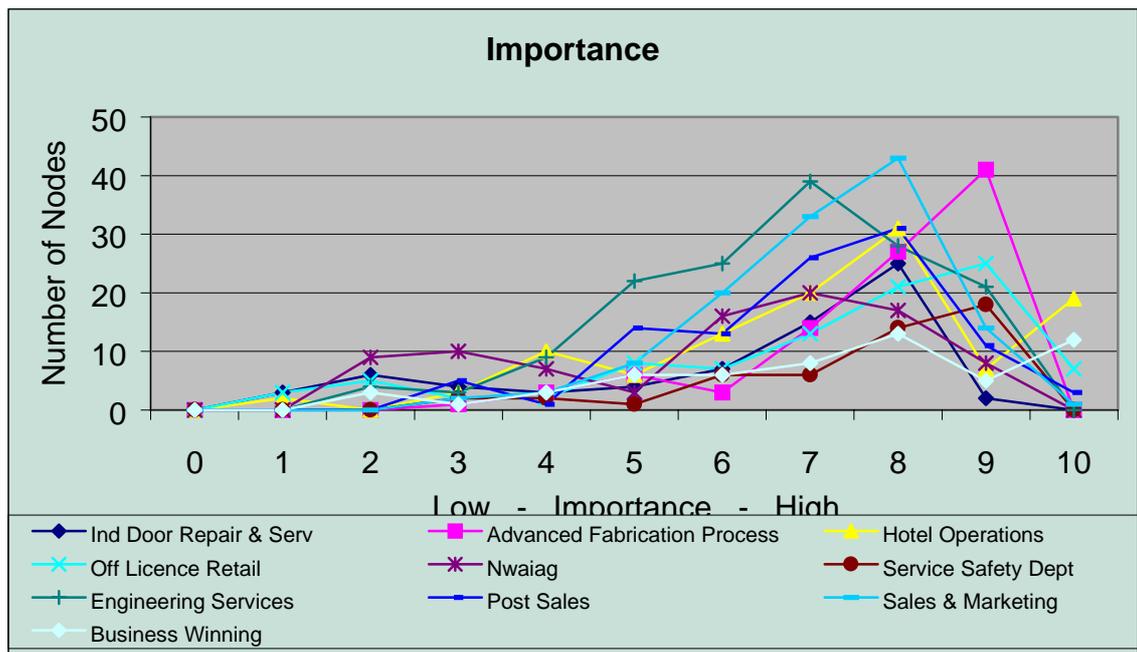


Figure 4: All importance parameter values

The importance values in figure 4, for all SKAs being studied, show considerable similarity. The average for all values given is shown in figure 4a. This figure seems to indicate that people generally view the knowledge resource as an important feature in their working lives. Little knowledge is seen as unimportant.

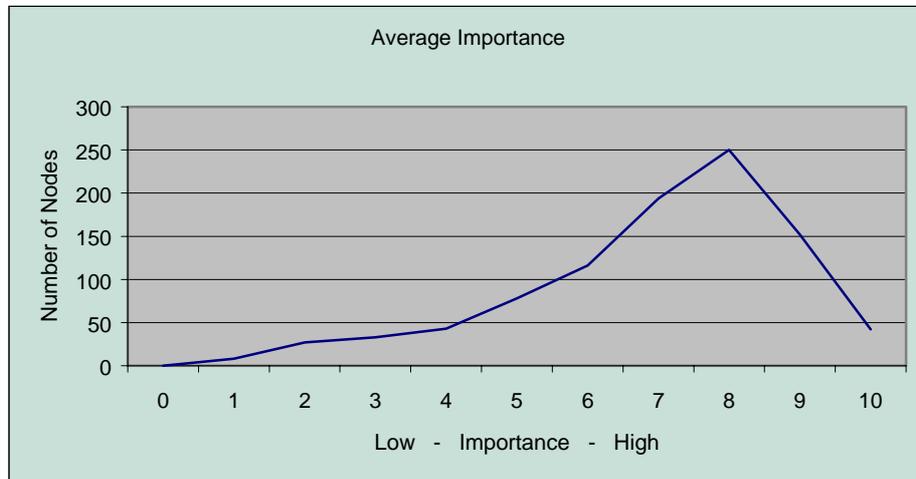


Figure 4a: Averages for all importance values over all SKAs

It could be argued that the results shown in figure 4 and 4a are what you would expect people to say about their knowledge. It is not possible to deny this, only to show that this is indeed what has been found.

2.3.2 The Difficulty to Replace Parameter

The results for the parameter that indicates how difficult to replace a piece of knowledge is thought to be are more scattered than those for importance.

Figure 5 is included for completeness but it is more difficult to identify a general trend.

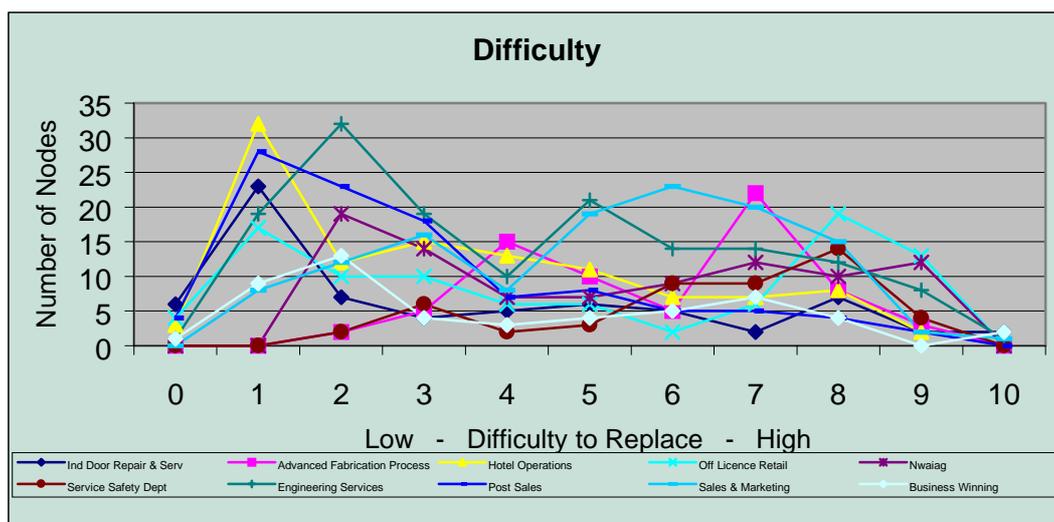


Figure 5: All Difficulty to Replace Parameter values

There is a weak indication within figure 5 that values for this parameter seem to have been given at either end of the scale (study or experience) rather

than in the middle (a mixture of both). A generalisation from closer inspection reveals that many of the engineering and management areas view their knowledge as easier to replace than areas such as sales and retail.

2.3.3 The Study – Experience Parameter

Figure 6 provides a result that is also inconclusive. In this case, closer inspection reveals that the engineering and management areas tend to view their knowledge more in the study category whilst the areas of sales, and service tend to view their knowledge more as acquired by experience. This probably reflects the way that the people interviewed have actually acquired their knowledge. As a commentator from the interview process it seems that those who are required to study professionally, seem to undervalue the subsequent progress they make with experience and those who have had much less need to study in any great detail, seem to miss reasonably clear indicators that show that much of the knowledge thought to be acquired through experience, could actually be gained through study.

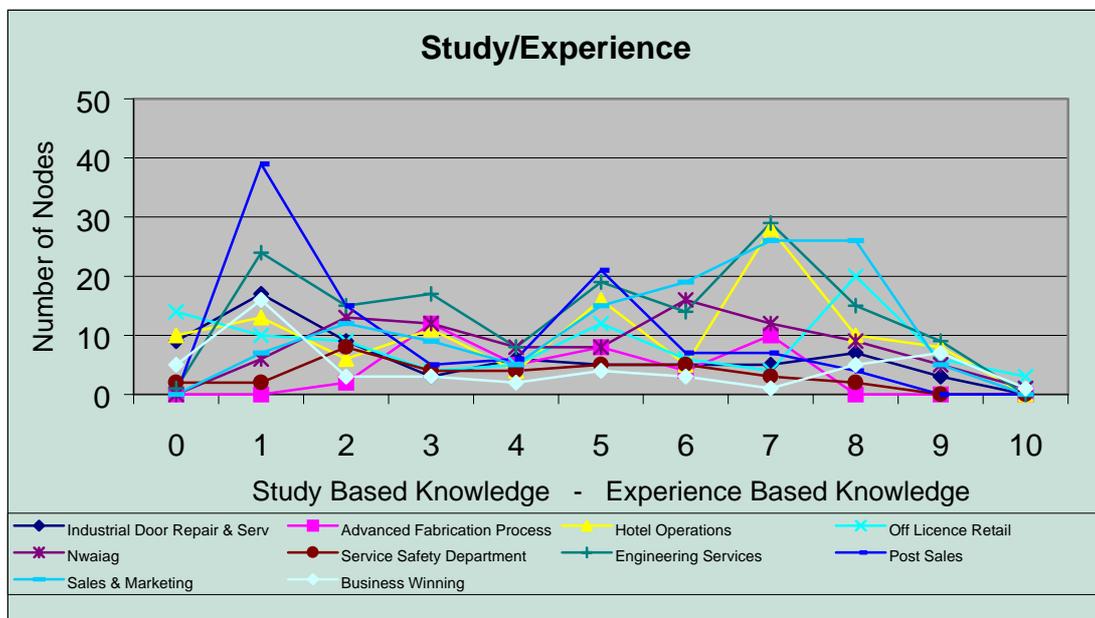


Figure 6: All study – Experience parameter values

2.3.4 The Knowledge By Parameter

A much clearer overall trend can be noticed in the known-by data. Figure 7 shows that the greater part of an organisational knowledge resource is known by a small percentage of the staff in a particular knowledge area. This is supported by the general trend shown in figure 7a.

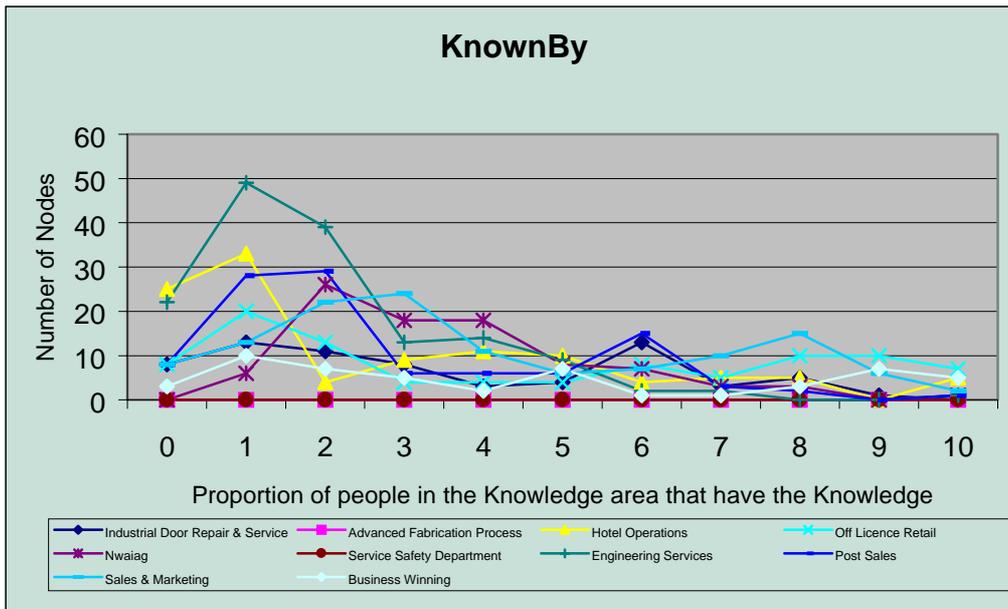


Figure 7: All Known-By parameter values

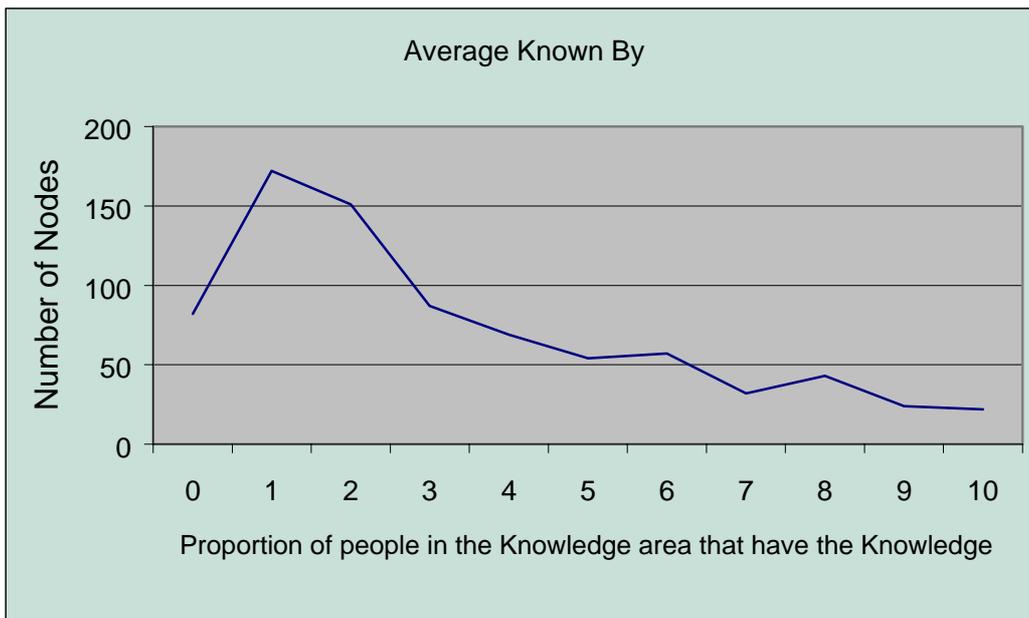


Figure 7a: Averages for all Known-By values over all SKAs

2.3.5 A combined parameter called Risk

The way 'risk' is computed has been described elsewhere (Gordon 2000) but it is knowledge that is most important *and* most difficult to replace *and* gained through experience *and* known by only a few staff. This combined parameter is intended to be an additional information source to help managers to target resources towards those parts of the knowledge asset that are in most urgent need of attention. Figures 8 and 8a show that there is a general trend of medium to high risk within an organisational knowledge resource. There are a few items that have very high risk and only a few that represent very little risk. Identifying and targeting resources at those items

with very high risk should proceed a more sustained management effort to deal with the majority of the knowledge resource.

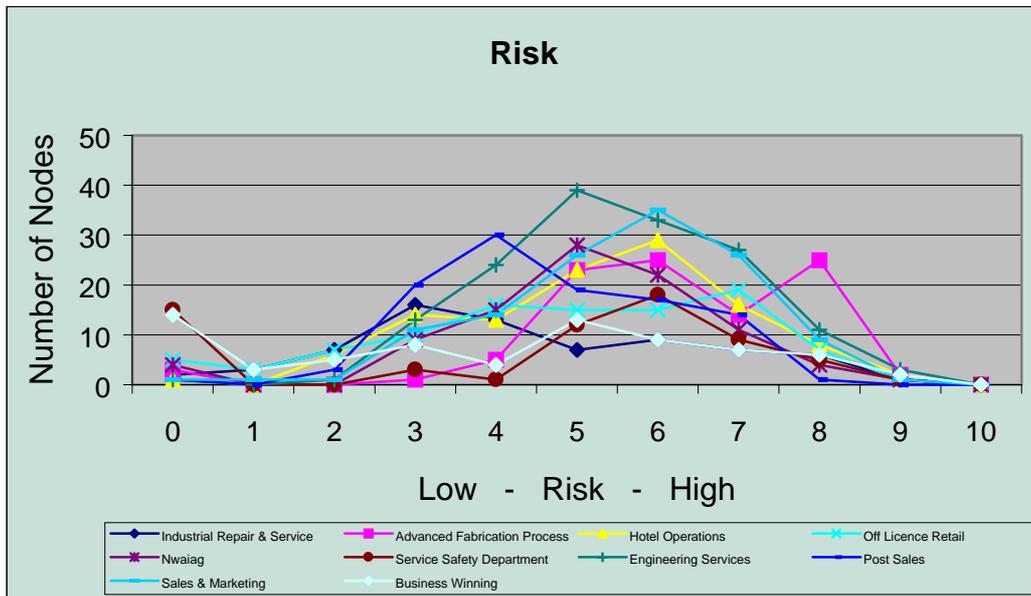


Figure 8: All Risk parameter values

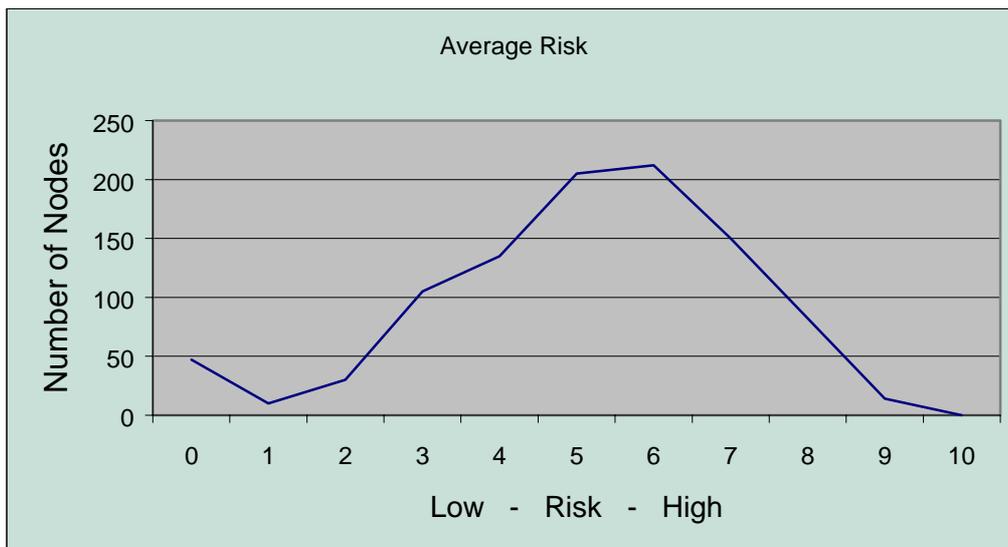


Figure 8a: Averages for all Risk values over all SKAs

2.4 Node Distribution

Table 3 shows the size of maps developed in actual projects in terms of the number of nodes elicited. Section 2.2 discusses possible reasons for the similar size of maps developed in separate projects. The SKA support tool 'SKAT' has an automatic layout function available that places nodes in the map at a particular level depending on the lowest level that a parent node exists. This means that if the automatic layout option is used then nodes on each map will be placed at a level that is consistent with a learning dependency hierarchy.

Figure 9 shows the node level distribution for all of the projects, actually, a few more projects than contained in the original list (project names have not been included). Node distribution is shown as a percentage of the total nodes for each map. This eliminates visualisation problems that would result from viewing the data for different size maps on the same axes.

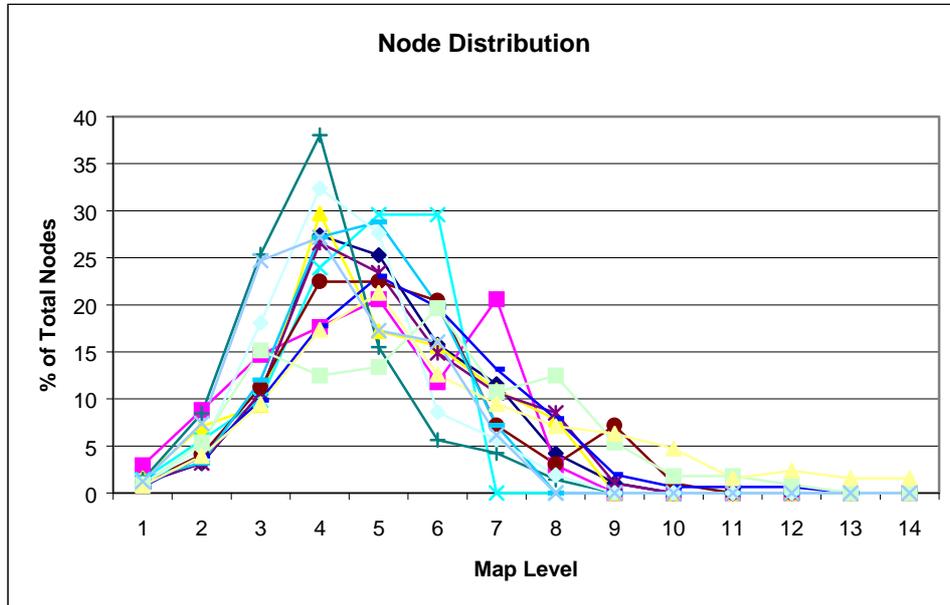


Figure 9: Node Level Distribution

The data from figure 1 provides a representative member of the graphs shown in figure 9. This map has only 8 levels however whilst figure 9 shows that some maps have 14 levels.

A small map of the knowledge area of 'negotiation' discussed later in section 3.1, was constructed as a focused addition to a larger study. Even this map however with only 24 nodes has roughly the same node distribution suggested in figure 9.

The average of the graphs from figure 9 is shown in figure 10 along with the average number of leaf nodes at each level

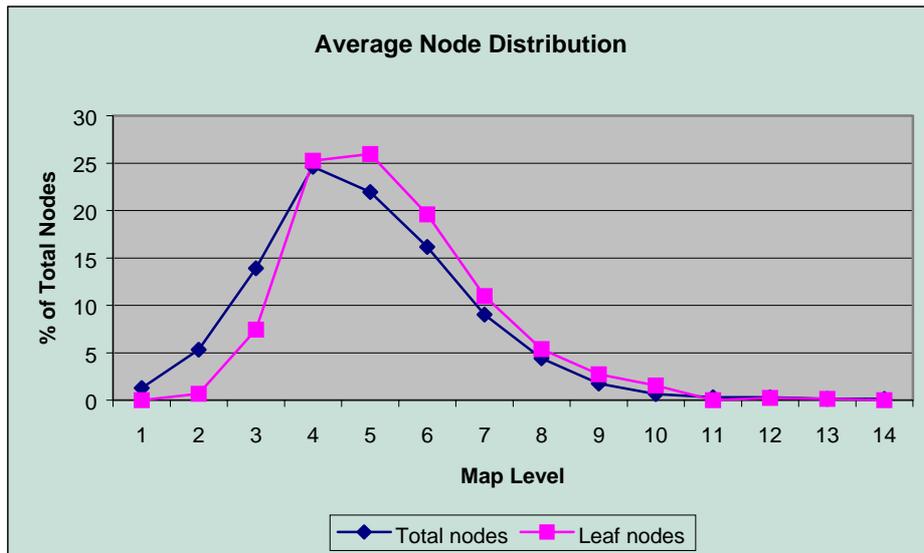


Figure 10: Average nodes and leaf nodes at each level

It can be seen that the number of leaf nodes at each level on maps follows a similar pattern to the total nodes but there is a level lag for the profile. This suggests that map exploration is not complete at any level. Exploration is controlled by the relevance of the knowledge resource and business issues and the time and size constraints discussed.

3. Implications of findings

The information considered in section 2 has the same value as data from controlled interviews. It is essentially opinion but that opinion has been focused by the interview methodology and the context within which the interviews were carried out. All of the information concerns the organisational knowledge resource and within this context it has value to anyone interested in the application of human knowledge.

3.1 Map Connectivity

The Structural Knowledge Auditing method has identified a particular type of knowledge structure map. Justification for the way the map is derived and organised has been provided in earlier work. Table 3 lists the quantity of data that is being considered during this study and it is argued that this data and the consistency in its derivation represents adequate justification for a comparative study.

Sections 2.1 and 2.2 discuss map complexity and size. Records of map growth during interviews and comments from business people viewing completed maps suggest that a knowledge structure map that contains about 100 knowledge nodes is a useful size. Smaller maps have been derived to investigate particular fundamental knowledge areas that have featured on organisational maps but have not been expanded. These maps are useful for the exploration of important knowledge areas and seem to be appropriately used within the context of a larger map. One area that featured on several organisational maps was the area of negotiation. This is not an exclusively

knowledge area but it does have a knowledge component. The map derived for negotiation contained 24 nodes, 27 arcs (low complexity) and contained 92 parameter values and was probably similar (my estimation of similarity) to the example shown in figure 2. Knowledge area maps from the lower range of node numbers seem to be less successful within the business context. The map featuring in this study with the highest number of nodes was well received but its complexity was less than the average value given in table 3. The map shown with the highest complexity value and with 112 nodes was the most difficult to read.

This discussion suggests that the target size and complexity for a knowledge area map to be well received would contain about 120 nodes and have a complexity of about 1.28. Clearly this cannot be engineered during interviews and it would not be correct to try to do it. The point is however, that these factors can affect the usefulness of a knowledge structure map.

Map or Graph connectivity is not only a problem here. It is a problem in the general area of graph research. The World Wide Web also presents significant connectivity problems. Carrière and Kazman (Carrière 1997) consider connectivity issues related to searching the web that are very similar to those faced in this work. The paper considers issues of visualising structures that have similar shapes to those met in Knowledge Structure Maps and they also seek to take advantage of node connectivity for searching purposes. The 'hot spots' identified, could have implications to knowledge structure map analysis.

3.2 Map Layout

There is a reasonably substantial research base for automatic map and graph layout (Giuseppe 1994, Ioannis 1998) and also useful alternatives for information visualisation (Kimelman 1994). Map (or graph) complexity is a key factor in useful visualisation. There are a number of different research approaches to the visualisation of complex maps. Map layout is an important contributor to map usefulness. When a knowledge structure map has been fully elicited it needs to be laid out in an acceptable way, taking account of map readability, complexity and knowledge content and relationships. During elicitation, maps also need to be reorganised from time to time so that validation can take place more efficiently.

Several automated and semi automated functions are available for map layout design. Work spanning 2.5 years has led to the conclusion that map layout needs to be carried out by an experienced person with directed automated assistance where and when required.

There is also a real need to investigate alternative visualisation options that can be used appropriately either as whole or partial visualisations for map areas. There is no reason why visualisations cannot be mixed for optimal viewing of a complex map. This work has not yet been carried out within the context of SKA and its specific needs.

One worthy point though is that people seem to like a large printed (plotted) version of the map to hang on the wall. Computer versions offer manipulation, resizing and progressive disclosure amongst other things but people still like a big printed map to hang on the wall. They can stand around this and if they like and bang their fingers on certain areas to make a point in a human way.

3.3 Implications of Parameter Values

In section 2.3 it was concluded that:

- 1) Staff view knowledge as an important resource (even if senior managers may not)
- 2) The greater part of an organisational resource is known by relatively few staff
- 3) People who have to study extensively to gain expertise in their knowledge area tend to consider that much of the knowledge they have must be gained through study even though it is often clear that they have really acquired it through experience.
- 4) People who have not studied extensively for their area of work tend to feel that the larger part of their knowledge domain must be acquired through experience even though it is often clear that much of the knowledge is very procedural and could easily be passed on through study.
- 5) Very generally, engineers and managers seem to think that their knowledge is easier to replace and sales and service staff tend to think that their knowledge is more difficult to replace.

It seems clear that several of the perceptions above are incorrect. This could be explained by a general tendency for people who for some reason invest significant time in knowledge acquisition realising that they don't actually know very much. People who acquire skills over time without explicit study agendas possibly seem to feel that this is the only way to acquire their domain knowledge. It is possible that the way that people actually acquire knowledge influences their opinion about appropriate methods of knowledge acquisition. People who study more seem to think that their knowledge is easier to replace.

The implication of (2) above, is the one addressed by the more recent approach to multi-skilling. SKA provides, within a target knowledge domain, some indication as to the level of multi-skilling that exists. What is not clear is how much this can reasonably be improved. It is very unlikely that everyone in an organisation can possess all organisational knowledge, thus representing the ultimate in multi-skilling. The other extreme is that each person knows only that knowledge that they apply daily (or possibly not even this). Multi-skilling is seen as a more efficient way to utilise staff. It is possible that efforts in the multi-skilling area will move the peak of the graph in figure 7 / 7a to the right.

A different approach for higher risk knowledge is to use intelligent systems (computer or AI methods) to store organisational knowledge. The graphs do not show the knowledge that may be retained in computer systems. It is unlikely that the interview technique used to derive these parameters would

be successful in estimating the knowledge available in computer systems. One of the main reasons for this is that computers contain a lot of information but currently very little knowledge.

3.4 Combined Parameters – Risk

Figures 8 and 8a show a similarity between studies and a useful average. Taking things further and calculating standard deviations etc would be to push the credibility of this data a little too far. From a knowledge management perspective, the figures indicate that there is a significant need within organisations for a sort of explicit knowledge management organised by managers and targeted at identifiable knowledge areas. It is also suggested in section 2.3.5 that urgent effort to target a few high risk areas can be followed by a more permanent programme of knowledge resource development. This tends to suggest that a standardised model for this sort of explicit knowledge management based on SKA could prove useful for many organisations and business types.

3.5 Node distribution similarity

It is unrealistic to believe that the similarity in node distribution as defined in section 2.4, represents some fundamental property of knowledge structure. Certainly the data gathered cannot be used to infer this. There must be a reason for the similarity however since figure 9 shows data from 13 separate maps and it was stated that 'negotiation' has a similar distribution.

The suggested reason for this similarity is that both interviewees and interviewers remained focused, at least in part, on the main knowledge study area whilst each map was developed. This idea is reinforced by the methodology that asks parameter related questions that continually relate back to the main knowledge topic.

The study of negotiation is interesting here because it was carried out as a separate investigation to review more about a node on another map where negotiation was a leaf node. The suggestion is that if negotiation had been explored more fully as part of the original work that it would have been less well developed than in the separate study.

These observations seem to support conclusions drawn in section 3.1 that relate to optimum map size. The word optimum is being used here in a 'useful to manager the knowledge resource' context. The average leaf node distribution shown in figure 10 is a feature of map size constraints.

4. Conclusion

The purpose of this study has been to investigate the perceptions of knowledge holders (and investigators to some extent) when knowledge is studied in a controlled way using learning dependency and parameter elicitation as a structural framework. Section 1 stated that the context for this investigation is to consider these implications to an organisational knowledge resource. This study is not exhaustive but it is representative. For instance

there is more data available on map connectivity than presented here but the data that is presented is representative of the whole.

Looking back through this study the main points that have been uncovered are:

- a) There seems to be an optimum map size that people (managers) are happy with.
- b) There seems to be a map complexity level beyond which the map would be less useful.
- c) There appears to be consistency in map exploration that reinforces an optimum map size and justifies a restricted length project set.
- d) Some regularity in parameter elicitation can be seen as validation for an otherwise very subjective value elicitation strategy.
- e) Regularity in parameter elicitation also suggests that there are commonly held views between knowledge holders in some very different situations (section 3.3).
- f) These methods can be used to provide managers with an overall understanding of the knowledge resource that is held by staff. Section 3.4 is just one example of this.
- g) Node distribution analysis provides additional evidence to the claim for consistency and credibility in the application of the methodology.
- h) Node distribution also appears to support claims for optimum map size for studies.
- i) People's perception of the knowledge resource seems to be influenced by the knowledge acquisition method.

All of these claims relate to the way that knowledge holders have revealed information about knowledge within the context of a particular study objective. The list above uses words such as 'suggests' and 'appears' rather than 'shows' or 'is' deliberately. The reader may feel that even these words are too strong or may feel that they are understatements.

The author feels justified in claiming that the methodology used in this work is rigorous enough to identify real patterns and strong enough to justify claims about an organisational knowledge resource. The method helps to organise and validate the information elicited from knowledge holders so that their perceptions of the knowledge resource are properly represented. In turn, the visualisation and analysis can lead to useful data on which to build a management lead knowledge resource development and protection programme.

References

Carrière J, Kazman R: 1997: WebQuery: Searching and Visualizing the Web through Connectivity: in Sixth International WWW Conference 1997

Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis Tollis 1994: Annotated Bibliography on Graph Drawing Algorithms in *Computational Geometry: Theory and Applications* 4:235-282

Gordon J.L. 2000: Creating Knowledge Structure Maps to support Explicit Knowledge Management: Applications and Innovations in Intelligent Systems VIII, Eds, Macintosh, Moulton & Coenen: Pages 33-48: Springer

Gordon J.L. 1999: Creating Knowledge Maps by Exploiting Dependent Relationships: Knowledge Based Systems, Vol13 (April 2000) pages 71 - 79, Elsevier Science

Ioannis Tollis, Giuseppe Di Battista, Peter Eades, Roberto Tamassia 1998: Graph Drawing, Prentice Hall. (ISBN 0-13-301615-3)

Kimelman D. Leban B, Rothy T, Zernik D 1994: Reduction of Visual Complexity in Dynamic Graphs, DIMACS Graph Drawing'94