3rd AAAI-Sponsored Knowledge Acquisition for Knowledge-Based Systems Workshop
Banff, Canada, November, 1988

A Methodology for Recognizing Consensus, Correspondence, Conflict and Contrast in a Knowledge Acquisition System

Mildred L. G. Shaw and Brian R Gaines
Knowledge Science Institute
Department of Computer Science
University of Calgary
Calgary, Alberta,
Canada T2N 1N4

ABSTRACT

One problem of eliciting knowledge from several experts is that experts may share only parts of their terminologies and conceptual systems. Experts may use the same term for different concepts, use different terms for the same concept, use the same term for the same concept, or use different terms and have different concepts. Moreover, clients who use an expert system have even less likelihood of sharing terms and concepts with the experts who produced it. This paper outlines a methodology for eliciting and recognizing such individual differences. KSS0 is a knowledge support system providing an integrated set of tools for knowledge acquisition and the rapid prototyping of knowledge-based systems. *Socio* is a tool in KSS0 that compares the conceptual systems of multiple experts in the same domain. It can be used to focus discussion between experts on those differences between them which require resolution, enabling them to classify them in terms of differing terminologies, levels of abstraction, disagreements, and so on. Socio promotes the full exploration of the conceptual framework of a domain of expertise by encouraging experts to operate in a "brain-storming" mode as a group, using differing viewpoints to develop a rich framework. It avoids social pressures forcing an invalid consensus by providing objective analysis of separately elicited conceptual systems.

INTRODUCTION

The elicitation of knowledge from one expert presents enough problems that elicitation from many might seem an unnecessary complication. However, a major reason for practical interest in expert systems is that in many domains, such as quality control, expertise is essentially distributed over many experts, and the purpose of the system is to bring it together (Hayes-Roth 1984). More fundamentally, the use of multiple experts in 'brain-storming' or synectics sessions has been found valuable in problem-solving, and is an attractive technique to prevent individuals experts 'blocking' and failing fully to explore and express their conceptual domains.

It is important to note in dealing with multiple experts that consensual agreement upon domain concepts is only one of many significant possibilities. Experts may legitimately have different terminologies for the same domain concept. They may describe it at different levels of abstraction. One expert may describe a concept in operational terms and another in descriptive terms. They may also legitimately use the same terminology for different domain concepts. They may be using the same term distinguished by different contexts. One expert may have a very different conceptual framework or strategy from another. These differences may be carried through to an expert system design that allows users to obtain advice based on different and mixed sources of expertise, and to ask for dissenting opinions (Boose & Bradshaw 1987).

Thus, in a knowledge acquisition system, it is important not to attempt to force a false consensus on a group of experts on the assumption that there is some 'correct' terminology and conceptual framework. However, it is also important to bring to light differences among the experts and make these clearly available for discussion. Some of them may reflect errors in elicitation, others differences in terminology, others differences in conceptual frameworks. In any event, the discussion of these differences is in itself a significant stage in the knowledge elicitation process.

The next section establishes a theoretical framework for knowledge elicitation from mutiple experts. The following sections describe a practical methodology based on this framework, a computer tool implementing this methodology, and results of using this tool to apply the methodology to practical knowledge elicitation.

CONCEPTUAL SYSTEMS IN A DOMAIN COMMUNITY

Figure 1 shows a domain in which an expert system could be constructed from the knowledge of existing experts. The two "experts" shown, Expert 1 and Expert 2, are best thought of as *roles* in the individuals involved. These expert roles are distinguished by involving conceptual systems which, when applied to problem solving in the domain, lead to recognition of the individuals playing the roles as beings "experts" in the domain. The individuals will also have other roles involving other conceptual systems appropriate to other domains in which they may, or may not, be experts.

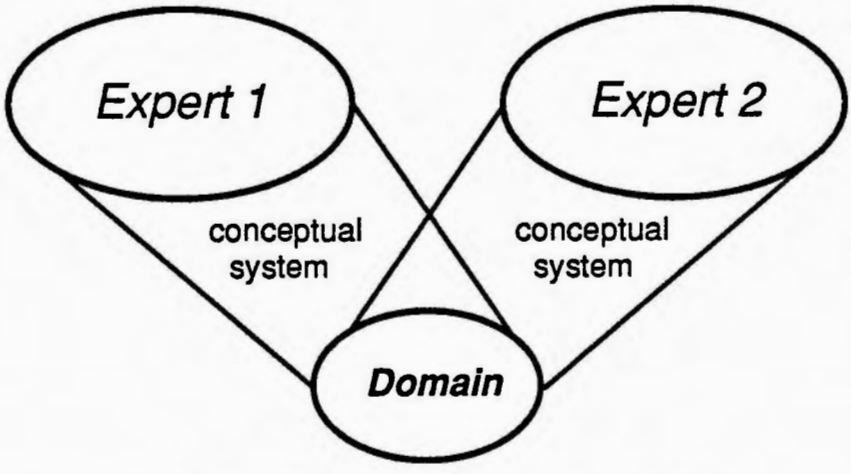

Fig.1 Experts acting in the same domain

The figure shows an overlap in the two conceptual systems indicating corresponding concepts. The experts will each have terminologies for expressing the concepts involved. The possible overlaps between concepts and terminologies leads to four relations which can occur between parts of the conceptual systems as shown in Figure 2:

- *Consensus* occurs when there is agreement in both the concepts and terminology used;
- *Correspondence* occurs when different terms are used for the same concepts;
- *Conflict* occurs when the the same terms are used for different concepts;
- *Contrast* occurs when different terms are used for different concepts.

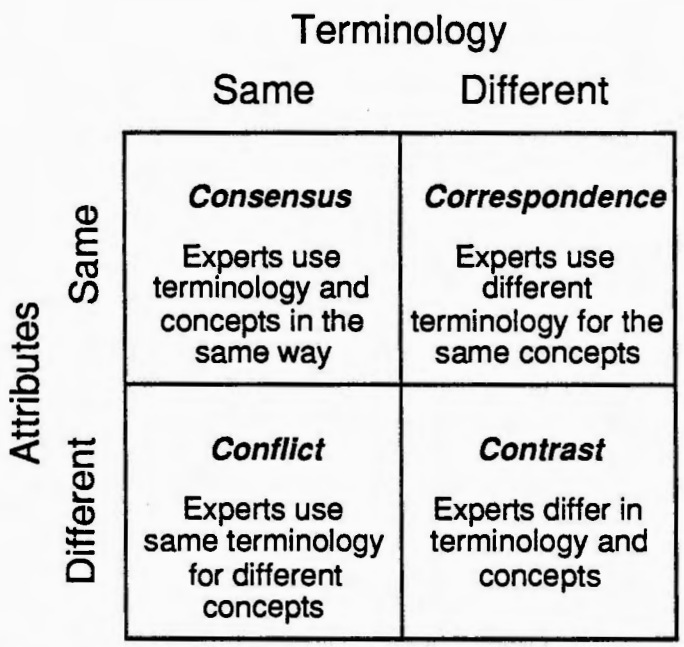

Fig.2 Consensus, correspondence, conflict and contrast among experts

The recognition of consensual concepts is important because it establishes a basis for communication using shared concepts and terminologies.

The recognition of corresponding concepts is important because it establishes a basis for mutual understanding of differing terms through the availability of common concepts.

The recognition of conflicting concepts is important because it establishes a basis for avoiding confusion over the labeling of differing concepts with same term.

The recognition of contrasting concepts is important because it establishes that there are aspects of the differing expertise about which communication and understanding may be very difficult, even though this should not lead to confusion. Such contrasts are more common than is generally realized. For example, it is possible to derive the same theorem in mathematics either by using an algebraic perspective, or a geometric one. There is nothing in common in these two approaches except the final result. It may still be possible to discuss the same domain using consensual and corresponding concepts that were not fundamental to the problem solving activities.

These considerations also apply to clients of the experts whose concepts and terminologies may be evaluated along these four dimensions in relation to those of the expert. The recognition of possible conflicts between the experts' and clients' use of terminology, and the provision of a variety of corresponding concepts, are major factors in the usability of an expert system.

In Figure 3, Client 1 is shown as having overlap only with Expert 1. He or she has no overlap whatsoever with any other individual in the figure, expert or user. Client 2 has overlap both with Client 3 and with the two experts. He or she is able to understand the concepts of Expert 2 and to a lesser extent those of Expert 1. Client 3 has no overlap with either expert, but has sufficient overlap with Client 2 that Client 2 can explain to him the advice given. In the development of an expert system, Client 1 may choose not to take advantage of Expert 2's expertise in the system. It would be particularly important to develop additional concepts and terminology to enable Client 3 to have adequate communication with, and understanding of, the system.

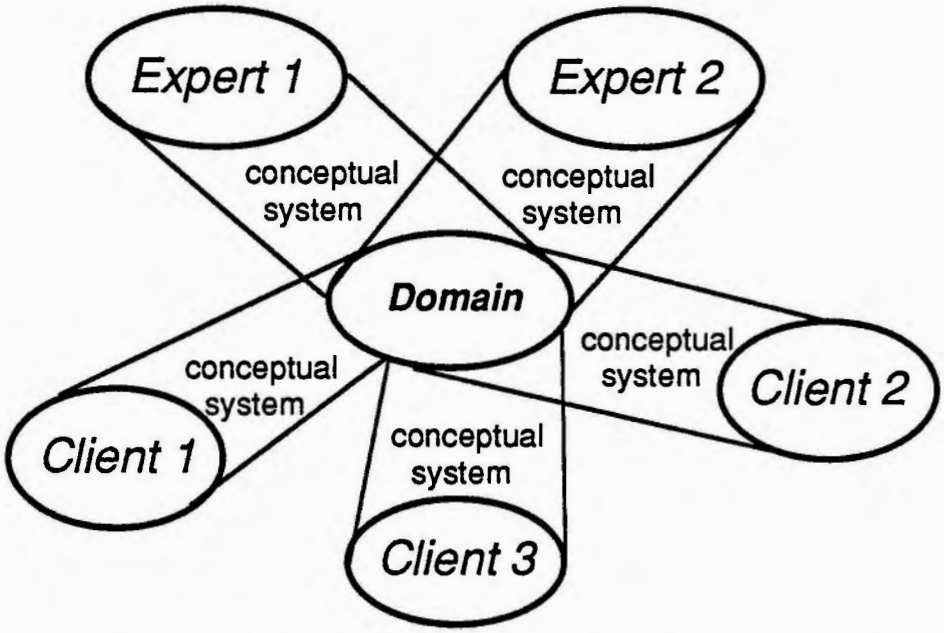

Fig.3 Experts and clients operating in the same domain

These are some of the possible relations which may occur between the conceptual systems of users and experts. The model is one of *roles* within an individual using a particular *conceptual system* and expressing it in a particular way to enter into an externally perceived and valued *domain community*. This model of coherent intellectual processes, each with its own individuality, each using its own conceptual system and its own terminology to enter into the domain community provides a cogent picture of the rationale behind, and the conflicts involved in, the development of communication and understanding in the community and in expert systems playing a role in that community.

DERIVING CONCEPTUAL SYSTEMS

Personal Construct Psychology (Kelly 1955, Shaw 1980) has been widely used as a basis for developing methods of knowledge acquisition from experts in a given domain (Shaw & Gaines 1983, Boose 1984, Diederich, Ruhmann & May 1987). Personal Construct Psychology is a psychology of the individual and many of its applications emphasize the idiosynchratic nature of conceptual systems (Mancuso & Shaw 1988). However, it is also a psychology of the individual interacting with the world and embedded in society (Shaw 1985). It encompasses shared concepts and conceptual systems, and their generation through experience. In particular it encompasses those systems that are so widely shared and so significant that they are construed as *knowledge*, and treated as having an existence virtually independent of their carriers.

Entity-attribute, or repertory grid, methodologies based on Personal Construct Psychology allow a significant part of the conceptual systems of experts to be elicited through interactive interviewing techniques such as those of KSS0 (Gaines 1987) and Aquinas (Boose and Bradshaw 1987). The resultant data structure is one in which the terms for entities and attributes in a domain have been specified by the expert, together with the values of those entities along the dimensions of the attributes. The entities are usually concrete items in the domain whose nature, definition and names can be agreed by experts and clients. The attributes reflect individual conceptual systems and may be used, and labeled, idiosynchratically.

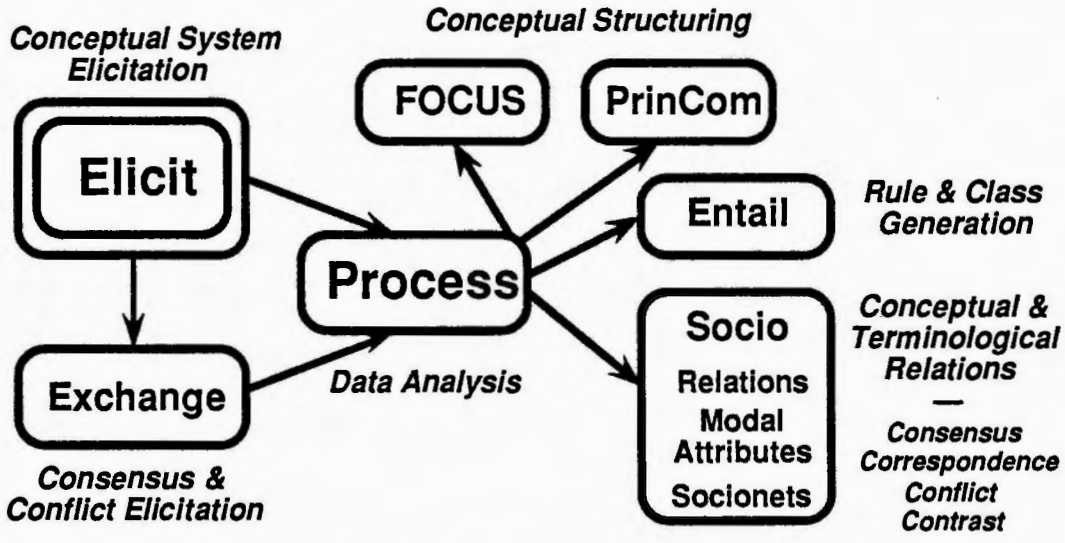

Fig.4 Some tools in Knowledge Support System Zero

Figure 4 shows the main tools in KSS0 relevant to the issues discussed in this paper:

- *Elicit* accepts specifications of entities within a domain and provides an interactive graphical elicitation environment within which the experts can distinguish entities to derive their attributes. The resultant conceptual system is continuously analyzed to provide feedback prompting the expert to enter further entities and attributes.
- *Exchange* allows the terms in the conceptual system derived from one expert to be used by another in order to determine whether the two experts have consensus or conflict in their use of terminology and concepts.
- *Process* gives access to various cluster analytic methods for the analysis and display of the conceptual systems elicited: FOCUS shows the system as a hierarchical structure; PrinCom as a spatial map; and Entail as a logical structure of classes and rules.
- *Socio* compares elicited and exchanged grids in a variety of ways to determine consensus, correspondence, conflict and contrast, and its methodology is described in detail in the next section.

Figures 5 and 6 give examples of the interactive graphical elicitation of attributes using KSS0 in a study of the consistency of expertise across experts, and across time, in a small group of geographers specializing in mapping techniques and their application to geological exploration (Shaw & Woodward 1987). Figure 5 shows some of the mapping techniques used as entities through a rating screen from the program, *Elicit*, in which a geographer is rating the entity, *trend surface analysis*, on the attribute, *doesn't incorporate geologic model—incorporates geologic model*. Figure 6 shows some of the characteristics derived as attributes through a match screen also from *Elicit* in which a geographer is being asked to distinguish *punctual kriging* from *universal kriging*. Figure 7 shows an elicited entity-attribute grid.

Exchange methodologies were developed for the measurement of understanding and agreement between either two individuals, two roles or on two occasions (Shaw 1980). To do this two people, possibly experts with differing points of view, each elicit a grid in an area of common knowledge or experience. Each may choose his own entities independently of the other, and elicit and rate his or her attributes quite separately. Each then can *Exchange* his or her grid, that is use the other's entities and attributes but fill in his or her own the rating values. For example, in terms of Figure 5, the exchanging expert would see the terms *doesn't incorporate geologic model—incorporates geologic model*, but the entities would all be to the left and have to be dragged to the scale with no knowledge of where the other expert had previously placed them.

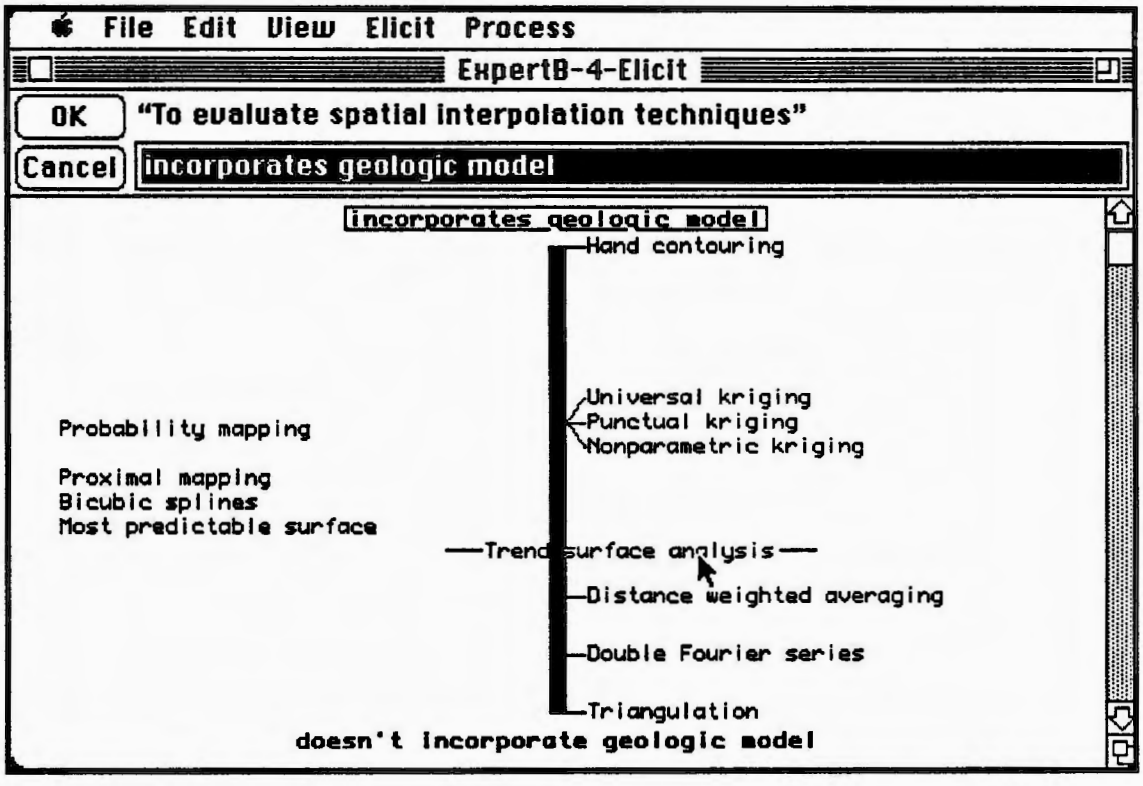

Fig. 5 KSS0 attribute elicitation screen showing entities being rated

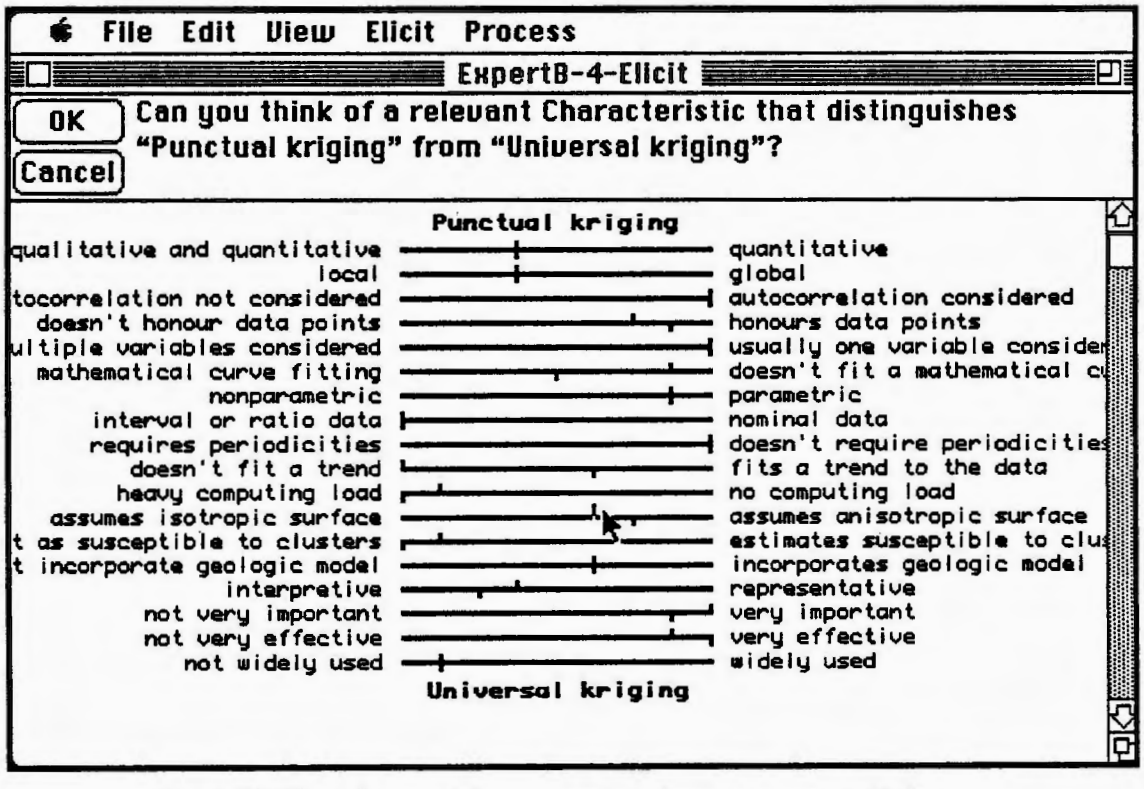

Fig. 6 KSS0 entity match screen showing ratings on attributes

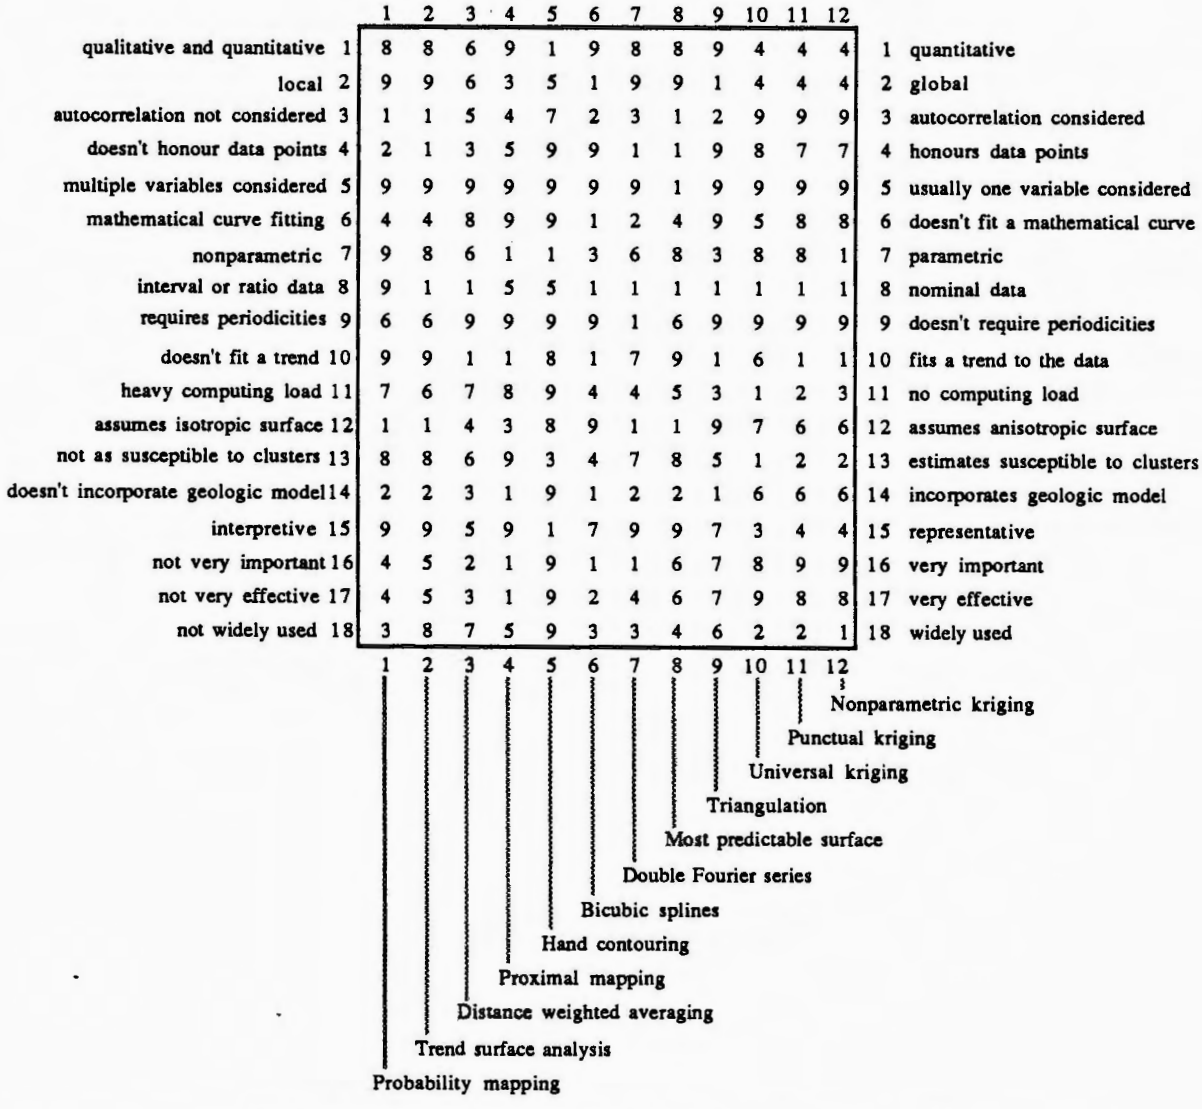

Fig.7 Conceptual system in entity-attribute grid

DERIVING RELATIONS BETWEEN CONCEPTUAL SYSTEMS

The *Socio* analysis in KSS0 allows members of a community to explore their agreement and understanding with other members, and to make overt the knowledge network involved (Shaw 1980, 1981, 1988). It is an extension of techniques such as SOCIOGRIDS (Shaw 1980) for deriving socionets and mode constructs from groups of individuals construing the same class of entities. The objective of Socio is to take different conceptual systems in the same domain and compare them for their structure, showing the similarities and differences. It may be regarded as the implementation of a simple form of analogical reasoning. Figure 8 shows the basis of operation of Socio—consider one set of data as being the base class defined by its entities, their attributes and values, and consider variant classes:

- *Entity-Attribute Compare*: The conceptual system at the upper left has the same entities and attributes but possibly differing values. Socio analyzes the matches between the entities and the attributes in this and the base class according to the values, and shows those entities and

attributes that are similar and those which are different. A typical application is to see whether experts agree on the definitions of attributes by asking them to separately fill in the values for a grid exchanged between them.

- *Entity Compare*: The conceptual system at the upper right has different entities but the same attributes. Socio analyzes the matches between the entities in this and the base class, and for each entity in the base class shows the closest matching entity in the other class. A typical application is to see whether experts are using different terminologies for the same entities by asking them separately to define entities and fill in the values for a domain defined through a class with agreed attributes.
- *Attribute Compare*: The conceptual system at the lower left has different attributes but the same entities. Socio analyzes the matches between the attributes in this and the base class, and for each attribute in the base class shows the closest matching attribute in the other class. A typical application is to see whether experts are using different terminologies for the same attributes by asking them to separately define attributes and fill in the values for a sub-domain defined through a class with agreed entities.

If, as in the lower right, neither entities nor attributes are common then no comparison is possible.

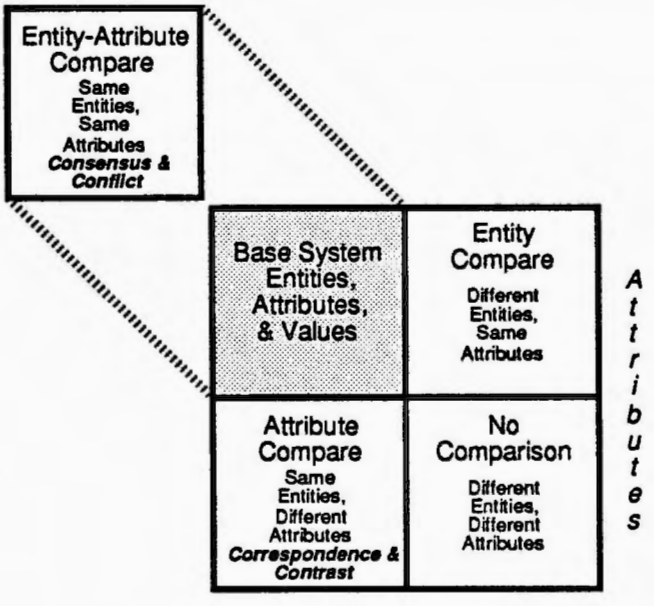

Fig.8 Possible comparisons between base and related systems

When a number of classes representing the same sub-domain are available, Socio also provides two other forms of analysis:

- *Mode Entities and Attributes*: Socio attempts to derive "modal" entities or attributes that reflect a consensus among the experts. It does this by extracting those which occur as highly matched entities or attributes across the majority of conceptual systems. A typical application is to reach consensus on critical concepts that are associated with a rich vocabulary at differing levels of abstraction.
- *Socionets*: Socio derives a socionet showing the degree to which each expert is able to make the same distinctions as another expert, even if they use different terminology. A typical application is to validate the knowledge acquisition process by determining whether the structure derived conforms with known relations between the experts.

CONSENSUS, CORRESPONDENCE, CONFLICT AND CONTRAST

Using the concepts developed above it is possible to develop a complete methodology for eliciting and analyzing consensus, correspondence, conflict and contrast in a group of experts, and implement this as an automatic process using the tools in KSS0. The methodology has three phases shown in Figures 9 through 11.

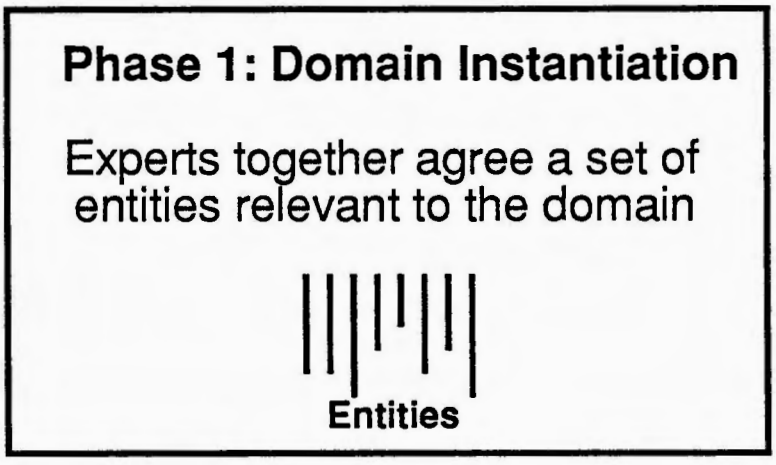

Fig.9 Co^4 Methodology Phase 1

Phase 1: Domain Discussion and Instantiation

In phase 1 a group of experts comes to an agreement over a set of entities which instantiate the relevant domain. This is the initial phase of any entity-attribute methodology, whether used with individuals or groups. However, with individuals the elicitation techniques may be used to elicit more entities as the exploration of the conceptual domain proceeds. When comparing multiple experts it is important that a set of entities is established at the start of the comparative study, and that the experts mutually agree on the definitions of these.

A convenient way to generate this set of entities is have each expert individually use Elicit to enter his or her conceptual system for a domain, and then extract the elicited entities from all the grids for discussion and consolidation by the group. This has the advantage that the experts gain some experience in the use of the KSS0 tools and can take advantage of the full elicitation facilities.

Phase 2: Conceptualization and Feedback

In phase 2 each expert individually elicits attributes and values for the agreed entities. The resultant conceptual systems will have the same entities but different attributes and can be analyzed by the *Attribute Compare* component of Socio as shown in Figure 8. This takes each attribute in one grid and determines the best matching attribute in the other grid, if there is one. The result is a mapping from the attributes in one expert's grid to those in another's as shown in Figure 10.

In evaluating this mapping we are not particularly interested in the terminology used but rather whether one expert has an attribute that can be used to make the same distinctions between the entities as does the other expert, regardless of whether these distinctions are called by the same terms. If such a *correspondence* occurs then the experts have a basis for mutual understanding of the underlying concept.

If an attribute in one system has no matching attribute in the other then it stands in *contrast* to all the other expert's attributes, and it may be very difficult for the other expert to understand the use of this term.

Note that the arrows in Figure 10 need not be symmetric. Attribute A in G1 may be best matched by attribute B in G2, but attribute B in G2 may be best matched by attribute C in G1. The only constraint is that if an attribute has an incoming arrow then it will have an outgoing arrow.

Thus, the second phase provides the basis for analysis of correspondence and contrast as already discussed.

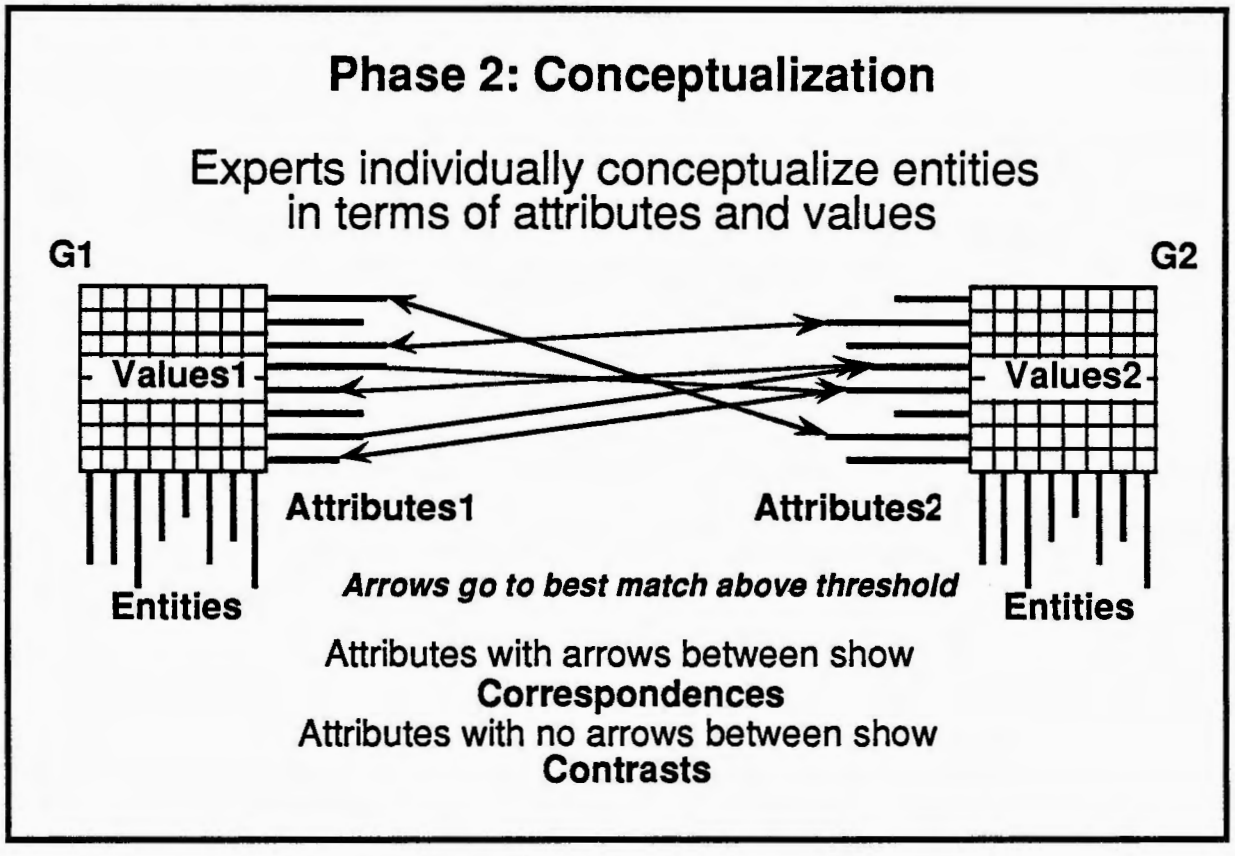

Fig.10 Co^4 Methodology Phase 2

Phase 3: Exchange and Compare

In phase 3 each expert individually exchanges elicited conceptual systems with every other expert, and fills in the values for the agreed entities on the attributes used by the other experts. The resultant conceptual systems will have the same entities and attributes and can be analyzed by the *Entity-Attribute Compare* component of Socio as shown in Figure 8. This takes each attribute in one grid and determines whether it matches the corresponding attribute in the other grid.

The result is a map showing *consensus* when attributes with the same labels are used in the same way and *conflict* when they are not as shown in Figure 11. Thus, the third phase provides the basis for analysis of consensus and conflict as already discussed.

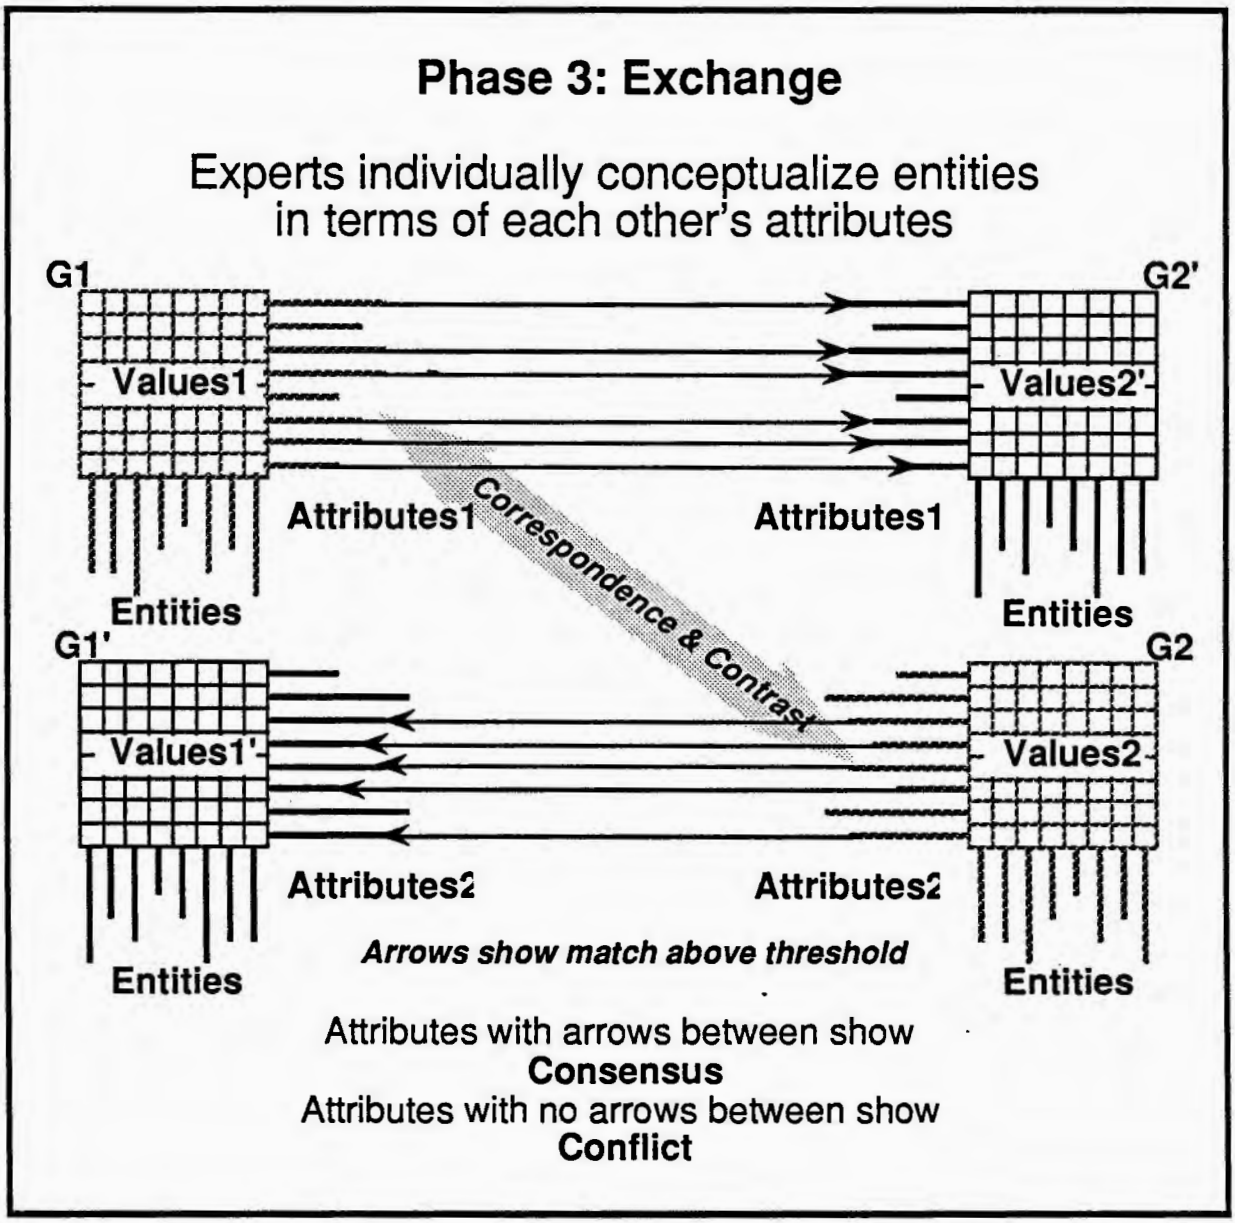

Fig.11 Co^4 Methodology Phase 3

Figure 11 also shows the correspondence and contrast relations analyzed in phase 2 as a relations between two of the grids used in phase 3. Thus, for two experts, four grids obtained by one elicitation and one exchange each, are sufficient to classify the relations between attributes in terms of consensus, correspondence, conflict and contrast. The methodology scales up linearly for each expert, so that n experts will be involved in n elicitations, one base elicitation and n-1 exchanges.

These three phases result in the experts' conceptual systems having become overt and inter-related. They lead naturally to later phases in which classes, objects and rules can be developed incorporating consensual, corresponding, and some of the contrasting attributes as kernel knowledge and, possibly, the conflicting and remaining contrasting attributes as 'other opinions.'

EXAMPLES OF THE METHODOLOGY IN ACTION

This section gives an example of the methodology in action using data from the study of geographers specializing in mapping techniques previously cited (Shaw & Woodward 1987).

Figure 12 shows an entity-attribute comparison from the geographic study in which expert B has filled in the values for a class defined by entities and attributes elicited from expert A. The print out shows the matches sorted with best first. The cumulative percentage is given of the number of attributes with matches greater than the value shown. In the list of attributes at the top, it can be seen that there is consensus on *interval data—nominal data*, but conflict on *requires no model—requires a model* and *linear interpolation—nonlinear interpolation*. There is clear consensus on the top four attributes, clear conflict on the lower five, and uncertainty about the remaining three.

In the list of entities at the bottom, it can be seen that there is close agreement on *vector trend analysis*, but high disagreement on *probability mapping*. This output can be used to focus a discussion between the experts on why they differ in their views of *probability mapping* and the classification of mapping techniques in terms of *linear* or *nonlinear interpolation*. For example, the first attribute, *requires no model—requires a model* shows high disagreement, and looking further into the elicited data it can be seen that expert A thinks that *probability mapping requires a model*, whereas expert B thinks that *probability mapping requires NO model*. On inquiring into this, the explanation given was couched in terms of what one actually means by the term "model", indicating the conflicting use of terminology.

```
G1:G2 33.3% over 80.0 (ExpertB attribute-consistency-with ExpertA)
  1:   8.3% ≥ 90.9   A2: Interval data - Nominal data
  2:          81.8   A4: Global - Local
  3:          81.8   A5: Intuitive - Mathematical
  4:  33.3% ≥ 81.8   A6: Requires spatial search - Does not require spatial search
  5:  41.7% ≥ 75.0  A10: Difficult to understand - Easily understood
  6:          72.7   A3: Non-polynomial - Polynomial
  7:  58.3% ≥ 72.7   A7: Discontinuous - Continuous
  8:  66.7% ≥ 59.1  A12: Does not consider non-spatial attributes - Considers non-spatial
  9:  75.0% ≥ 56.8  A11: Few points - Many points
 10:  83.3% ≥ 50.0   A8: Does not honour data - Honours data
 11:          47.7   A1: Requires no model - Requires model
 12: 100.0% ≥ 47.7   A9: Linear interpolation - Non-linear interpolation

G1:G2 8.3% over 80.0 (ExpertB entity-consistency-with ExpertA)
  1:   9.1% ≥ 83.3  E11: Vector trend surface analysis
  2:  18.2% ≥ 77.1   E2: Trend surface analysis
  3:          72.9   E4: Distance weighted averaging
  4:  36.4% ≥ 72.9   E7: Bicubic splines
  5:          70.8   E3: Kriging
  6:          70.8   E5: Proximal mapping
  7:          70.8   E8: Double Fourier series
  8:  72.7% ≥ 70.8   E9: Most predictable surface
  9:  81.8% ≥ 60.4  E10: Negative exponential surface
 10:  90.9% ≥ 56.2   E6: Hand contouring
 11: 100.0% ≥ 43.8   E1: Probability mapping
```

Fig.12 Entity-attribute comparison of expert B with expert A

Figure 12 can be redrawn as a *difference grid* where rating values (in this case 1 to 5) for expert B's ratings of expert A's entities on his attributes are subtracted from expert A's similar rating values respectively. Figure 13 shows this with the entities and attributes about which they agree the most in the top right corner, shown by no difference or a difference of only 1; and those with most disagreement towards the bottom left, shown by the maximum difference of 4 or a large difference of 3. Hence from this difference grid, the consensus and conflicts can easily be identified and discussed by the experts.

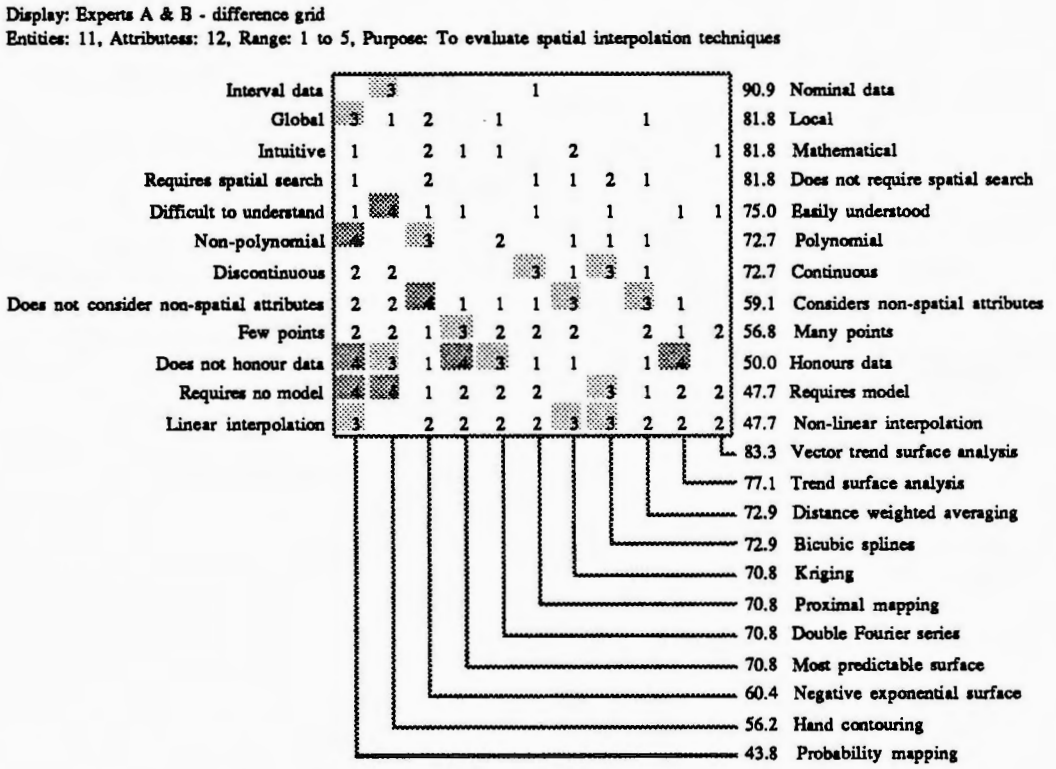

Fig.13 The difference grid for experts A and B

```
G1<:G2 62.5% over 80.0 (ExpertA attribute-construed-by ExpertB)
  1:   6.2% ≥ 88.5    G1A2: Local - Global
                      G2A2: local - global
  2:  12.5% ≥ 87.5    G1A3: Low level data - High level data
                      G2A8:   nominal data - interval or ratio data
  3:          86.5    G1A1: Does not honour data points - Honours data points
                      G2A4:  doesn't honour data points - honours data points
  4:          86.5    G1A7: Short distance autocorrelation - Long distance autocorrelation
                      G2A2:                              local - global
  5:  31.2% ≥ 86.5    G1A9: New geographical technique - Old geographical technique
                     G2A18:              not widely used - widely used
  6:  37.5% ≥ 85.4   G1A16: Not widely used - Widely used
                     G2A18: not widely used - widely used
  7:  43.8% ≥ 83.3    G1A5: Discontinuous - Continuous
                      G2A2:          local - global
  8:  50.0% ≥ 82.3    G1A4: Mathematically complex - Mathematically simple
                     G2A11:   heavy computing load - no computing load
  9:          81.2   G1A10:   Hard to adapt to multivariate - Easy to adapt to multivariate
                      G2A5: usually one variable considered - multiple variables considered
 10:  62.5% ≥ 81.2   G1A12:   Does not require spatial search - Requires spatial search
                     G2A13: estimates susceptible to clusters - not as susceptible to clusters
 11:          79.2    G1A6: Does not require a priori model - Requires a priori model
                      G2A2:                                  local - global
 12:  75.0% ≥ 79.2   G1A11:      Few points - Many points
                     G2A18: not widely used - widely used
 13:          76.0   G1A13:          Does not use polynomial - Uses polynomial
                      G2A6: doesn't fit a mathematical curve - mathematical curve fitting
 14:  87.5% ≥ 76.0   G1A15: Not very effective - Very effective
                     G2A17: not very effective - very effective
 15:  93.8% ≥ 72.9   G1A14: Not very important - Very important
                     G2A18:    not widely used - widely used
 16: 100.0% ≥ 71.9    G1A8:        Models the stationarity - Assumes stationarity
                      G2A3: autocorrelation not considered - autocorrelation considered
```

Fig.14 Attribute comparison of expert B with expert A

Figure 14 shows an attribute comparison from the geographic study in which expert B has specified attributes and filled in the values for a class defined by entities elicited from expert A. The print out shows the matches sorted with best first. The cumulative percentage is given of those with matches greater than the value shown, and the attribute from B which best matches each from A is shown beneath it. It can be seen that:

- The first or highest match which accounts for 6.2% of the attributes has a level of 88.5 out of a possible 100 if they were identical. That is, both experts are using the attribute *local-global* in the same way. This is an example of *correspondence*. However, it is not an interesting one because we can see that the terms used are the same and it is effectively arising from a *consensus*.
- The first and second matches together account for 12.5% of all attributes, and they are matched over the level of 87.5. For the second attribute, when Expert A uses the term *low-level-data—high-level data,* expert B is using the term *nominal data—interval or ratio data.* This shows a difference in terminology which can be interpreted as their *levels of abstraction* being different in their construing of this topic. This illustrates the two experts having attributes in *correspondence*.
- The third match again shows a *correspondence* that can be interpreted as arising from *consensus*, with both experts using the attribute *does not honour data points—honours data points* in the same way. Note that this showed up as *conflict* in Figure 12, probably reflecting that, in this study, the methodology was more complex and the exchange grids from which Figure 12 derives were elicited and discussed *before* those from which Figure 14 derives.

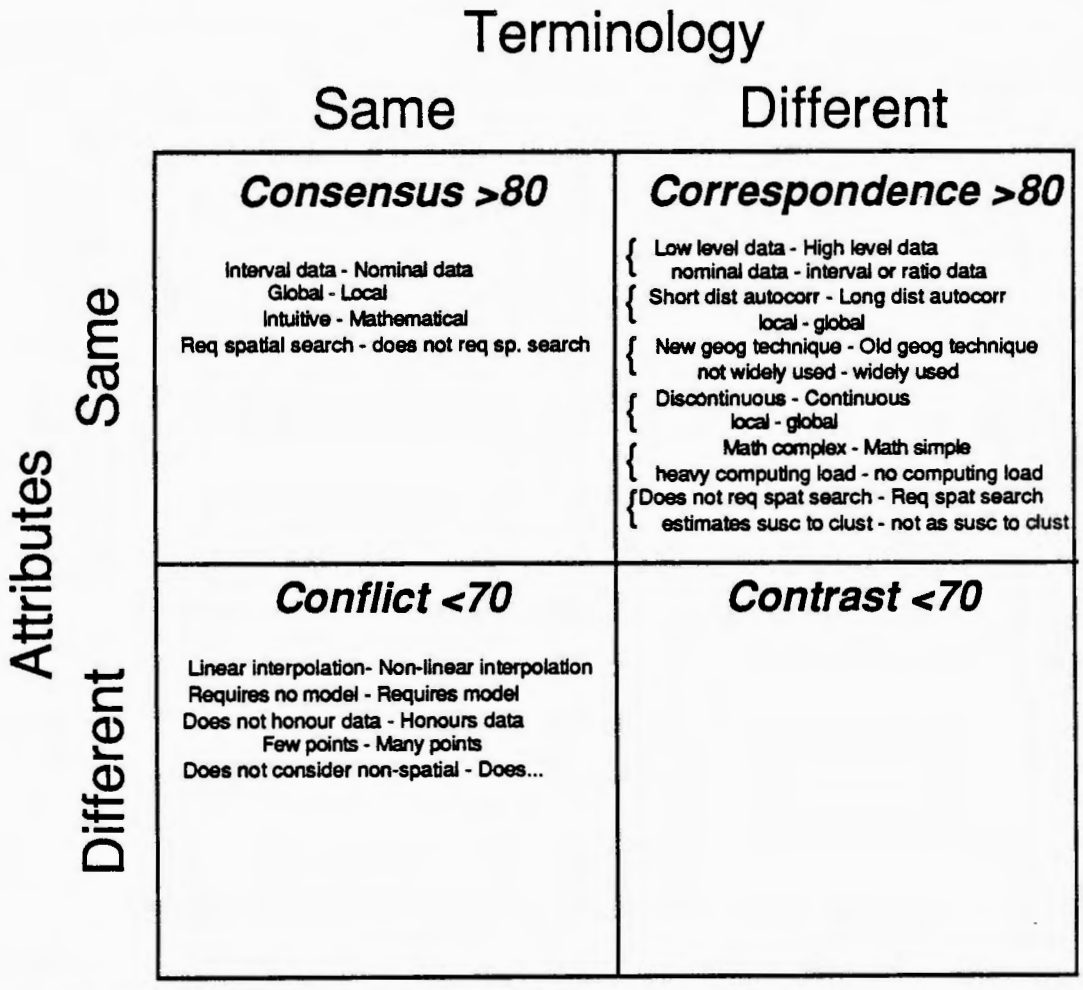

Fig.15 Consensus, correspondence, conflict and contrast from Figs.12 and 14

- The fourth match shows a *correspondence* between *short-distance autocorrelation—long-distance autocorrelation* and *local—global*. Notice that *local—global* used by expert B was also used in the first match indicating that expert A has two attributes *short-distance autocorrelation—long-distance autocorrelation* and *local—global* which are used similarly to each other but with different terminology, whereas expert B has only one. In fact, looking on to the seventh and eleventh matches, it can be seen that expert A has two more attributes *discontinuous—continuous* and *does not require a priori model—requires a priori model* which correspond to expert B's single attribute *local—global*. This shows a differences in richness of concepts not necessarily making new distinctions in the class so far defined by the entities.
- The eighth match, still over the level of 82, again shows *correspondence*. It shows the attribute *heavy computing load—no computing load* is being used by expert B to correspond to *mathematically complex—mathematically simple* used by expert A. We can interpret this as a difference in terminology corresponding to a correlation in the real-world.

Figure 15 shows each of the attributes from Figures 12 and 14 put into the appropriate quadrant of Figure 2. There is no significant example of *contrast* in this data, possibly because the two experts work very closely together. Such examples do arise in a full analysis of the three experts in the original study.

```
Mode Attribute 1: 13 attributes in 3 grids at 80.0
  G1A1:   Does not honour data points - Honours data points
  G1A2:                        Global - Local
  G1A5:                    Continuous - Discontinuous
  G1A7: Long distance autocorrelation - Short distance autocorrelation
  G2A2:                        global - local
  G2A4:    doesn't honour data points - honours data points
  G2A7:                    parametric - nonparametric
  G2A9:       requires periodicities - doesn't require periodicities
 G2A10:      fits a trend to the data - doesn't fit a trend
 G2A12:     assumes isotropic surface - assumes anisotropic surface
  G3A1:                        global - local
  G3A6:                   periodicity - non-periodicity
  G3A7:                   very smooth - non-smooth data

Mode Attribute 2: 8 attributes in 3 grids at 80.0
 G1A12:   Does not require spatial search - Requires spatial search
  G2A3:    autocorrelation not considered - autocorrelation considered
 G2A13: estimates susceptible to clusters - not as susceptible to clusters
 G2A15:                    representative - interpretive
  G3A3:                data restrictions - no data restrictions
  G3A8:                low computer cost - high computer cost
  G3A9:               no error estimates - error estimates
 G3A13:       not so effective technique - very effective technique

Mode Attribute 3: 6 attributes in 3 grids at 80.0
  G1A4:     Mathematically complex - Mathematically simple
  G1A9: New geographical technique - Old geographical technique
 G1A16:           Not widely used - Widely used
 G2A11:       heavy computing load - no computing load
 G2A18:            not widely used - widely used
 G3A11:        non-linear surface - linear surface

Mode Attribute 4: 4 attributes in 3 grids at 80.0
  G1A3: Low level data - High level data
  G2A8:   nominal data - interval or ratio data
  G3A2:        nominal - interval
  G3A5: non-continuous - continuous

Mode Attribute 5: 2 attributes in 2 grids at 80.0
 G1A10:   Hard to adapt to multivariate - Easy to adapt to multivariate
  G2A5: usually one variable considered - multiple variables considered
```

Fig.16 Mode attributes from three experts

Figure 16 shows a *mode attributes* analysis of data from the three geographers which extends the clusters noted above. These mode attributes are essentially inter-conceptual-system clusters of corresponding attributes from the three experts. The first may be interpreted as centering around *Global-Local* encompassing a variety of concepts related to this; the second as around autocorrelation techniques and their consequences; the third as around the complexity of the technique; the fourth around the type of data; and the fifth around the number of variables considered.

These five mode attributes can be interpreted as indicating stereotypical lines of reasoning most used by these experts. This output can be used as a basis for discussion among the experts on whether these conceptual clusters should be split because they confound different concepts expressed in apparently corresponding attributes, or retained as being the same concept expressed in different terminologies. Once this form of analysis has been discussed by the group it is readily edited and extended.

```
Attribute Links
 G3<:G2    71.4% over   80.0 (ExpertC attribute-construed-by ExpertB)
 G1<:G2    62.5% over   80.0 (ExpertA attribute-construed-by ExpertB)
 G2<:G1    61.1% over   80.0 (ExpertB attribute-construed-by ExpertA)
 G2<:G3    44.4% over   80.0 (ExpertB attribute-construed-by ExpertC)
 G3<:G1    42.9% over   80.0 (ExpertC attribute-construed-by ExpertA)
 G1<:G3    31.3% over   80.0 (ExpertA attribute-construed-by ExpertC)
```

Fig.17 Socionet analysis of three experts

Figure 17 shows a *socionets* analysis of the same data based on a set of comparisons like that shown in full in Figure 14. It can be seen from the first two links that the conceptual system of expert B encompasses the majority of attributes used by experts A and C. This indicates that expert B has a deeper knowledge of the topic than either expert A or expert C. That of expert A encompasses the majority of the attributes of B, shown in the third link. However, that of expert C does not encompass many of those of A and B, and that of A does not encompass many of those of C indicating that C has a different point of view, or a background of different experiences from those of A and B. This could then be explored in detail with expert C.

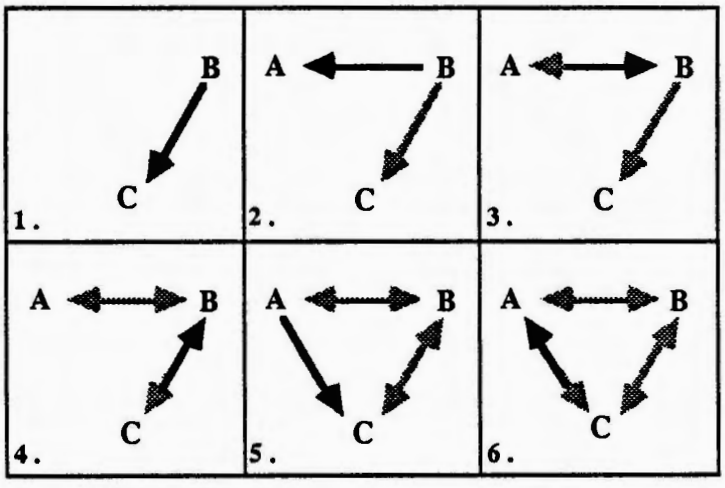

Fig.18 Socionet of relations between experts' conceptual systems

Figure 18 shows this information expressed in the socionet links of relations between the experts' conceptual systems for the domain. Each new link is shown by a black arrow as it is added into the sequence.

SUMMARY OF METHODOLOGY

Phase 1: Domain Discussion and Instantiation

1. The problem is discussed with the experts and a domain identified.
2. The experts are introduced to the use of KSS0. At this time a specific task is agreed upon and a purpose for the grid elicitation developed.
3. The first elicitation consists of the elicitation of a grid by each expert individually. This gives each expert experience in using KSS0 and provides an initial set of entities and attributes particular to each expert.
4. These grids are processed using FOCUS and PrinCom, and the results reported and explained to each expert respectively.
5. The experts meet together to discuss each set of entities, review and revise the purpose or context of the elicitations, and agree on a common set of entities which all understand.

Phase 2: Conceptualization and Feedback

6. This common entity set provides the basis for the next phase of data collection which consists of the elicitation of a grid by each expert using the set of common entities and a common rating scale. This elicitation is done separately for each expert.
7. These grids are also processed on FOCUS and PrinCom, and the results reported and explained to each expert respectively.
8. This data is then processed on Socio, and attributes of correspondence and contrast identified.

Phase 3: Exchange and Compare

8. The next phase of data collection consists of the experts exchanging their first grids (entities and attributes) with each other expert in the group to produce other grids.
10. This data is then processed on Socio, and attributes of consensus and conflict identified.
11. From the Socio results a mode grid is produced to identify the content of the shared lines of thought emanating from the distinctions made by all the experts together.
12. From the Socio results the socionet sequence is produced to identify the subgroups of experts who think and act in similar ways, and any who are different in their thinking from the main lines of thought in the group.

Phase 4: Rule Derivation and Validation

13. This data is then processed using Entail (Gaines & Shaw 1986) to produce rules, classes and/or objects with which to prime an expert system shell.
14. The final data collection phase consists of giving each expert a selected list of entailments generated from this grid. A set of examples from four separate levels of significance of entailment are then selected and randomly presented to the expert who is asked to rate each rule on a four point scale from correct to incorrect. This data is used to check that the set of rules produced by Entail is an accurate one.
15. The rules are tagged with any information resulting from these analyses, such as the name of the expert, his contribution to the group of experts, and his relationship with the other experts in the group with respect to the sub-domain under consideration.

CONCLUSIONS

The methodology described in this paper provides facilities for revealing the similarities and differences in the concept systems of different experts, or the same experts at different times, construing a domain defined through common entities or attributes. It can be used to focus discussion between experts on those differences between them which require resolution, enabling them to classify them in terms of differing terminologies, levels of abstraction, disagreements, and so on.

Note that the derivation of consensual, corresponding, conflicting and contrasting attributes is completely algorithmic, based solely on the data in the grids. This derivation is done by a computer program, not a knowledge engineer, and its basis can be demonstrated clearly to the experts through computer output such as the difference grid of Figure 13. Thus, there are no opinons being expressed about the correctness of the use of the attributes and terminology, that the differences highlighted are 'right' or 'wrong.' It is open to the experts to consider, discuss and explain these differences, changing or retaining them as they wish. Conflicts can be retained in the final system if desired by tagging classes, objects and rules with the sources from which they derive.

Note also that the methodology described applies equally to the relations between the conceptual structures of experts and representative clients for the system. The differences in conceptual systems shown in Figure 3 and the problems that may arise from them may be made overt, and the methods for overcoming these for the different classes of clients, may be analyzed in detail using the derivation of consensual, corresponding, conflicting and contrasting attributes from experts and clients.

The methodology promotes the full exploration of the conceptual framework of a domain of expertise by encouraging experts to operate in a "brain-storming" mode as a group, using differing viewpoints to develop a rich framework. It avoids social pressures forcing an invalid consensus by providing objective analysis of separately elicited conceptual systems.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. The KSS0 system was made available by the Centre for Person Computer Studies. We are grateful to many colleagues at the knowledge acquisition workshops for discussions which have highlighted many of the issues raised in this paper.

REFERENCES

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**, 27-33. California: American Association for Artificial Intelligence.

Boose, J.H. & Bradshaw, J.M. (1987) Expertise transfer and complex problems: using AQUINAS as a knowledge acquisition workbench for knowledge-based systems. **International Journal of Man-Machine Studies(26)**, 3-28.

Diederich, J., Ruhmann, I. & May, M. (1987) KRITON: a knowledge-acquisition tool for expert systems. **International Journal of Man-Machine Studies(26)**, 29-40.

Gaines, B.R. (1987) Rapid prototyping for expert systems. Oliff, M. (Ed). **Proceedings of International Conference on Expert Systems and the Leading Edge in Productions, Planning and Control.** pp. 213-241. University of South Carolina.

Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. **Fuzzy Sets and Systems**, 8(3), 315-328 (April).

Hayes-Roth, F. (1984). The industrialization of knowledge engineering. Reitman, W., Ed. **Artificial Intelligence Applications for Business**. pp. 159-177. Norwood, New Jersey: Ablex.

Kelly, G.A. (1955). **The Psychology of Personal Constructs**. New York: Norton.

Mancuso, J.C. & Shaw, M.L.G., Eds. (1988). **Cognition and Personal Structure: Computer Access and Analysis**. New York: Praeger Press.

Shaw, M.L.G. (1980). **On Becoming a Personal Scientist**. London: Academic Press.

Shaw, M.L.G., Ed. (1981). **Recent Advances in Personal Construct Technology**. London: Academic Press.

Shaw, M.L.G. (1985). Communities of knowledge. Epting, F. & Landfield, A.W., Eds. **Anticipating Personal Construct Psychology**. pp. 25-35. Lincoln: University of Nebraska Press.

Shaw, M.L.G. (1988). An interactive knowledge-based system for group problem-solving. **IEEE System, Man and Cybernetics**, (in press).

Shaw, M.L.G. & Gaines, B.R. (1983). A computer aid to knowledge engineering. **Proceedings of British Computer Society Conference on Expert Systems**, 263-271 (December). Cambridge.

Shaw, M.L.G. & Woodward, J.B. (1987) Validation of a knowledge support system. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.18-0-18-15. Banff (October).