SAND2008-xxxx

SAND2008-2031C

Application of Inductive Learning in Reasoning Based Decision Aids

Michael Senglaub, PhD

Sandia National Laboratories
PO Box 5800 MS 1161
Albuquerque, NM 87185
mesengl@sandia.gov

Abstract

This article constitutes a "report from the field" on a decision aid / fusion system with a foundation in reasoning from a heuristic based perspective of philosophy. While the system employs multiple components of Peirce's 'method of scientific inquiry', the focus of the article is the inductive learning engine. Within a Peircean reasoning engine, induction is the work horse of reasoning functions. Induction is a form of reasoning in which a sample is used as a representation of a much larger population. It is an approximate form of reasoning. It serves as a hypothesis validation mechanism, as a reasoning algorithm and it supports learning and the modification of knowledge. The capability discussed in this note focuses on the extension of Finn's methodology to define a learning algorithm that supports a comprehensive Peircean reasoning engine.

Introduction

Development of fusion or decision support systems requires the integration of reasoning capabilities into the system solutions. This report documents the inductive learning sub-system of a decision aid we have developed at Sandia. The system architecture provides a foundation for an integrated reasoning system based on C.S. Peirce's method of 'scientific inquiry' which includes abduction, deduction and induction. Theories of human information processing can be found in the fields of cognitive psychology, neural-physiology or in philosophical theories. Modal logic and philosophy provide a heuristic assessment of the functions of human reasoning at an engineering level. The neural-physiological models are the 'physics' of human cognitive functionality, defining the electrical-chemical dynamics of the processes associated cognition. Engineering a solution to this information processing problem necessitates an understanding of the heuristics of reasoning. As in thermodynamics, the 'laws' of heat transfer between bodies are a heuristic representation of a molecular phenomena and are used in engineering solutions to the flow of energy in materials. In order to understand the physics of this energy transfer process requires an understanding of statistical physics.

Reasoning consists of the trained and/or ad-hoc process we use to solve problems. Within the context of our development, abduction provides the basis for hypothesis generation when new data/information enters our awareness. Deduction is the engine that supports the hypothesis selection process and induction becomes the hypothesis verification process. In addition, induction supports the functions of approximate reasoning and learning, which seems to be an innate human capability.

A brief description of the knowledge representation technology is provided in order to establish the foundation on which the mathematics of induction is constructed. The focus of the paper is the mathematics of inductive learning and the applications of induction as implemented in the Peircean decision aid.

Supporting Architecture

The application domain supports decision aid development in which data/information is convolved with knowledge to create a belief state, a theoretical interpretation of situational awareness, which forms the basis for making decisions. The mathematical foundations of formal concept analysis provides a foundation for developing sets of operators that define the functionalities of our reasoning system.

Formal concept analysis (FCA), developed by Ganter & Wille (Ref. 4,5), is based on ordered set theory and uses lattice theory as a rich technology for visualizing information and knowledge. FCA is based on the idea of a formal context, $\mathcal{K}_{FC}$, defined by a 'triple' as the one in equation 1.

$$K_{FC} = (G, M, I) \qquad \text{Eqn 1}$$

In this equation G and M are sets of objects and attributes respectively and I is a binary relation between the two sets. Within our problem domain we have refined the relation operator, I, to be a set of relations, each member corresponding to a specific predicate in the information domain. This permits us to assign some descriptive property to the binary relationship between the objects and attributes.

There is an operator, symbolized by $(\cdot)'$, which aids in the definition of formal concepts from the formal context.

$$\begin{aligned}(A)' &\equiv \{m \in M \mid (g,m) \in I, \forall g \in A\} \\ (B)' &\equiv \{g \in G \mid (g,m) \in I, \forall m \in B\}\end{aligned} \qquad \text{Eqn 2}$$

In equation 2, the operator action on the object set A produces the set of attributes common to objects within the

'A' set. Likewise, application of the operator on a set of attributes 'B' produces the set of objects which posses those attributes in common.

A construct from FCA that we will need later in the paper is the idea of a Begriff. Within FCA the 'context' constitutes a unit of knowledge / information and forms the basis of a lattice. A concept is a set of sets, (A,B), in which A is the set of objects and B is the set of attributes common to that set of objects. The prime operator defined above is the operator that defines the mapping between the two sets.

The Begriff is the set of all concepts of a context. The set of concepts can be defined by an application of the 'prime' operator discussed in the previous sections. The first step in defining the Begriff is to define an attribute set consisting of the intent of a context. The power set of this attribute set provides the basis set on which the prime operator is applied to produce the extent of all concepts in a context.

$$
\begin{aligned}
S &= \{Intent\}_{formal\ context} \\
P_s &= \{p\} = \{\mathcal{P}(S)\} \\
\mathcal{B} &= \{\textstyle\sum_k \{(p_k)'\}, \{p_k\}\}
\end{aligned}
\qquad \text{Eqn 3}
$$

P_s represents the power set of the contexts intent, and $\mathcal{B}$ is the Begriff which is a partially ordered set of concepts without duplications that may result from the process of applying the 'prime' operator on each member of the power set.

A procedural approach to defining a Begriff can be found in Davey & Priestley's (Ref. 2) book. This approach relies on a process that uses a series of set intersections as the context is processed. The effect is the same while the equations above are a rigorous interpretation of the process described.

Induction

A summery paper by Evan Heit (Ref. 6) provides a nice analysis of induction from a psychological perspective. There work focused on the evaluation of inductive arguments of a form given in equation 4.

$$
\frac{Given\ a\ \ premise}{The\ conclusion\ must\ \ follow} \qquad \text{Eqn 4}
$$

In a deduction logic, if the premise is true the conclusion must be true, the inductive logic introduces an uncertainty. If the premise is true the conclusion is expected to be true. The psychological studies discussed in the review paper explored a number of factors associated with the argument and assessed the strength subjects placed on the conclusions.

The article describes conditions associated with good cases, sets of cases and properties for inductive reasoning. The examples they use are presumed to be single premise cases. As an example they discuss examples such as:

$$
\begin{aligned}
&\frac{Nearby\ houses\ were\ burgled}{My\ house\ will\ be\ burgled} \\
&vs. \\
&\frac{Houses\ 50\ miles\ away\ were\ burgled}{My\ house\ will\ be\ burgled}
\end{aligned}
\qquad \text{Eqn 5}
$$

The first case leads to a stronger induction than the later in the study. They provide similar cases for problems involving animals described to exist on an isolated island in order to attempt some form of knowledge constraint. It seems that, what has been missed is an implicit multi-premise induction being performed. this 'implicit' set of premises was not captured. Subjects bring prior knowledge into the problem and evaluate the strength of the inductive arguments based on this knowledge. Effectively things such as; " ...50 miles from here is a large city with a high level of crime, my neighborhood lies in a very safe area so the likelihood of my house being burgled is low".

Similarly the animal examples fail to take into account the fact that we possess varying levels of biological knowledge and recognize that premise and conclusion involving species from a similar order may possess common anatomical characteristics or susceptibilities. Their discussion of 'property', reflects a similar issue. Certain classes of predicate carry varying levels of prior knowledge or meaning and impart varying levels of strength to the inductive arguments. Their example uses the predicates 'thrives' versus 'secretes' one which is qualitative in nature while the second has a foundation in biology. The bottom line for me is the experiments discussed could not be considered single premise experiments with results drawn based on that supposition.

The discussion of sets of cases was of more interest in that it focused on the numbers and the diversity of the premises category. The greater the number of examples supporting a conclusion adds strength as does the greater the diversity of the examples (premises) in the inductive argument. It would seem that this is getting closer to the nature of induction and the inherent statistical character of the problem. Within the area of case sets no results dealing with counter examples was explored, This may be due to the single premise constraint imposed on the experiments. In this situation, the set of attributes and corresponding set of objects can represent instances that would result in both true and false conclusions.

Theoretical foundations

Within the domain of philosophy a significant body of research exists that addresses induction. The work used as the foundations for the Sandia decision aid is based on Peirce's modal of reasoning (Ref. 9, 10, 11) coupled to the work of V. Finn (Ref. 1, 3). Peirce defined induction '...as a form of reasoning from a sample to the whole sampled.' 'Induction is the mode of reasoning which adopts a conclusion as approximate.' Peirce indicated that there exists three kinds of induction all based on random samples. He seems to have used different terms for these types of induction but in general they consisted of a weak form, a strong form, and the 'gradual' form.

The strong form of induction consists of a sample or collection from a population in which it is possible to assess the proportion of the members of that population. The weak form deals with statements that could be disproved if a single counter example existed; e.g. 'liberals are intellectually bankrupt'. The third form of induction, gradual, is similar to the first form due to its quantitative nature. In this case an estimate of the population proportionality is made but each new sample acquired goes toward updating the proportions in the population.

Inductive Reasoning

Reasoning research in Russia has produced some very interesting results in which FCA is used as a knowledge representation technology and theories of abduction and induction have been developed. The work of V. Finn and V. Blinova(Ref. 1, 3) have used J.S. Mills (Ref. 8) canons as the guiding principles in developing a theoretical modal of inductive reasoning. Mills canons have been described by various researchers, as a set of inductive principles, a set of abductive principles and defined as principles describing causal reasoning. It seems that, in part due to ambiguities in definitions, that Mills canons have elements of each category.

Finn's (Ref. 1, 3, 7) work focused on the first of Mills five canons. His effort focused on an inductive 'learning' algorithm which may more accurately be characterized as an inductive reasoning algorithm.

Mill's (Ref. 8) canons consist of : the 'Method of agreement', 'Method of differences', 'Indirect method', 'Method of concomitant variation', and the 'Method of residues'. The description of the first canon used in Finn's algorithm was taken from Mill's *System of Logic*.

> The first canon: If two or more instances of the phenomenon under investigation have only one circumstance in common, the circumstance in which alone all the instances agree, is the cause (or effect) of the given phenomenon.

Finn's implementation of the first canon involves considering positive and negative examples of an objective that draw from a single set of attributes, and one or more unknown examples. The objective is to classify the unknowns according to the reference sets. The model develops a Begriff for the positive and negative example sets, then uses the concepts in each Begriff to classify the unknown examples as either a positive, negative or as an indeterminate.

The example in Blinova's paper consists of the positive and negative sets in equation 6.

$$\begin{array}{ll} g_1^+ = \{a,b,c\} & \\ g_2^+ = \{a,b,d\} & g_5^- = \{a,c,d\} \\ g_3^+ = \{a,b,e\} & g_6^- = \{b,c,d\} \\ g_4^+ = \{a,c,e\} & g_7^- = \{a,d,e\} \end{array} \qquad \text{Eqn 6}$$

The extra intent sets of the positive lattice consisted of {a,b}, {a,e}, {a,c}, and {a}. The extra intent sets of the negative lattice consisted of, {c,d}, {a,d}, and {d}.

The set of unknowns is defined in equation 7.

$$\begin{array}{l} g_8^\tau = \{a,b,c,e\} \\ g_9^\tau = \{c,d,e\} \\ g_{10}^\tau = \{a,b,c,d\} \end{array} \qquad \text{Eqn 7}$$

Finn's method classified the first case as a positive example, the second as a negative example and the third as indeterminate due to both positive and negative intents being subsets of the example intent.

The one aspect of Finn's method that I think needs modification is the heuristic imposed on the positive and negative lattices. Finn requires that there be '2' or more examples in a concept before it can be classified as a positive or negative concept in the respective lattices. From a theoretical perspective, imposing a heuristic goes against the grain. This constraint is better suited for implementation in the engineered solutions where we might wish to establish a bias toward false positive or false negative classifications.

Inductive Learning

An excellent foundation for inductive reasoning has been established by Finn's work with Mill's first canon. In our implementation we have opted for an inductive learning system as opposed to the reasoning implementation. The algorithm builds on Finn's work with minor modifications and extensions. In Finn's work he defines a positive and negative lattice to use in the classification process. We needed to learn in situations in which the context possessed an arbitrary number of goals. E.g. the terrorist incident database I used contained over 30 groups which represented the goal attribute. The algorithm developed is provided in the next paragraphs.

The algorithm begins by accepting the entire training set. In the example to follow a subset of the State Dept. terror incident database was used. The incidents covered a period from 1992 to 1998. The attributes describing the incident include; {date, location, target, tactic, result, organization}. Organization is the goal attribute. A Begriff is defined for each goal attribute in the training set. These Begriffs constitute the positive lattices defined in Blinova's paper.

$$B_k \equiv (A_k, B_k) | \; A_k \subseteq A; \; B_k = (A_k)' \quad \text{with} \quad k = goal\ attribute \qquad \text{Eqn 8}$$

The first adjustment occurs in the construction of the goal oriented Begriffs. Engineering factors are introduced to bias the ultimate algorithms toward false positives or false negatives. This is done by setting likelihood thresholds to some fraction of '1.0'. This translates into the number of examples associated with a concept in the lattice. We can require that for a positive example we require two or more examples or in the case of counter examples two or more examples to reflect a counter instance. The next step of the algorithm is to define the resultant classification Begriff after any counter examples are removed from the goal Begriff.

$$B_k^* \equiv B_k^{\sim} - \sum_{m \neq k} B_m^{\sim} \qquad \text{Eqn 9}$$

In expression 9, B* represents the positive classification Begriff for a goal attribute, in our test problem this might be the group HAMAS. The Begriff class $B^{\sim}$, represents the Begriffs containing the engineering bias adjustments. The resultant becomes the basis for the new classification context that will be constructed by the next steps.

Each concept remaining in the classification Begriff, B*, represents positive examples for identifying a specific objective. A concept, C = (A,B), represents a set of examples and the associated set of identifying attributes. The new object (the former goal attribute) consists of all remaining attributes after the operation in equation 9. This resultant set of attributes is the union of the intent of all concepts remaining in B*. After each new object is defined from the learning sets, they are added to a new context representing a new knowledge domain.

The next step involves constructing a new context Begriff and estimating a set of concept likelihoods that can be interpreted as a metric for assessing the appropriateness of sets of attributes being used as a defining attributes. We can readily define a likelihood of a set of attributes associated with a concept, what we are doing is assuming we can attribute that concept likelihood to each of the attributes comprising that concept. The mathematics of estimating the likelihood for each concept is defined in equation 10.

$$lh_{k,m} = SizeOf(\, C_{m,E} \,) \,/\, SizeOf(\, B_k^* \,) \qquad \text{Eqn 10}$$

In expression 10, m is a concept index in the classification Begriff k. SizeOf(.) is a function which determines a size metrics in a Begriff, basically the number of objects associated with the argument, either the Begriff or the concept in this case.

The final classification context captures the essence of the training set in which each goal attribute is an object in the constructed classification context with each concept assessed for likelihood as a discriminator for an object in the newly constructed context.

Engineering implementation

The Peircean decision aid incorporates 3 basic thresholds in the construction of knowledge that will be integrated into the knowledge base. Two of these thresholds apply to categorized or pre-defined knowledge and to inductively generated knowledge. Attributes that are defined with a probability of being a attribute associated with an object as well as real valued attributes that have gone through a fuzzification process can be pre-screened for inclusion as a descriptor. E.g. if the probability of being a descriptor is less than 0.25 it is removed from consideration as a descriptive attribute within a context.

The second type of threshold is associated with the inductive learning algorithms using the likelihood estimates in equation 10. This 2nd class is used to bias the system toward false positives or false negatives in the construction of the positive and negative Begriff's in the Finn inspired inductive learning model. In the case of the positive Begriff, the threshold sets a lower limit on which concepts can be considered good examples of the objective being classified. When a certain grouping of attributes does not emerge as a descriptor often it is effectively considered an outlier. In the case of the negative Begriff, we are looking for numerous occurrences of a particular attribute grouping to use as an indicator of a negative example of the objective goal. The criteria are structured such that a high threshold value for either the positive or negative thresholds effectively require greater number occurrences of either positive or negative examples.

Initial Application

The focus of this application is to provide decision support capabilities and / or augment the efforts of an intelligence analyst. The focus is the construction or assembly of knowledge which provides the basis for evaluation information collected through sensor and intelligence sources. The system enables the inclusion of modal logics in support of the various functions of the system, such as knowledge update and belief revision and potentially some form of disjunctive data filtering. Many of these linkages are 'zero order' at this point and can be tailored to support a specific application domain. For example, the requirements on the disjunctive logic required in an automated system are going to be more stringent than for an application supporting an intel analyst. Similar arguments apply to the modal logics associated with knowledge construction and revision.

In this paper the focus of the discussion has been the inductive learning engine that is integral to the system. The architecture to support all of the decision aid functions is displayed in Figure 1. The other major component of the system is the abductive hypothesis generation engine that supports the construction of the virtual belief state of the decision aid. This component is described in other documentation (Ref. 13).

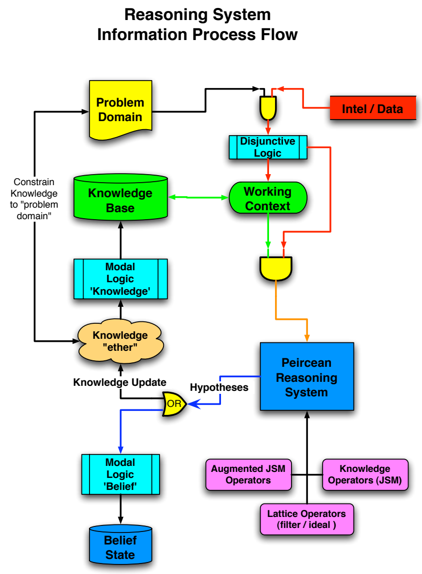

Figure 1. Decision support architecture.

Intel Analysis

One application domain we have been working supports an intel analyst. The rough scenario is an analyst is tasked with monitoring events for a possible terrorist attack in North America. Given this kind of problem there are many resources that may be utilized as pre-existing knowledge bases, such as a terrorist incident database.

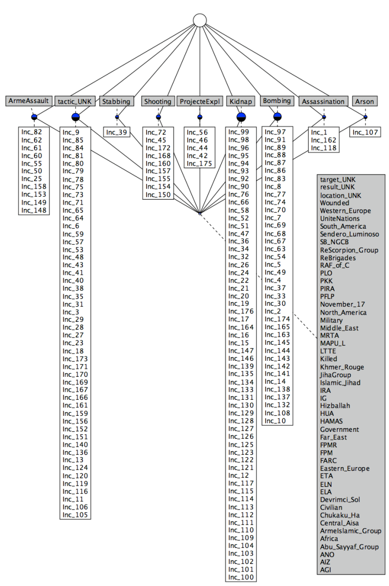

Figure 2. Raw terror incident data.

In this case a database from state department reports from 1992-1998 was used in an effort to understand possible patterns of behavior and tactical preferences by various groups. Raw data often conveys very little information to the analyst, Figure 2, but in many cases it is possible to process the information, to convert it to knowledge and find interesting and useful knowledge in the transformed data. The terror incidents were characterized by date, target, location, result, and the group responsible for the incident. The data was processed by the inductive learning engine in the Peicean Decision Aid (PDA) to construct the desired knowledge. The resultant knowledge base, in a lattice display, is depicted in Figure 3.

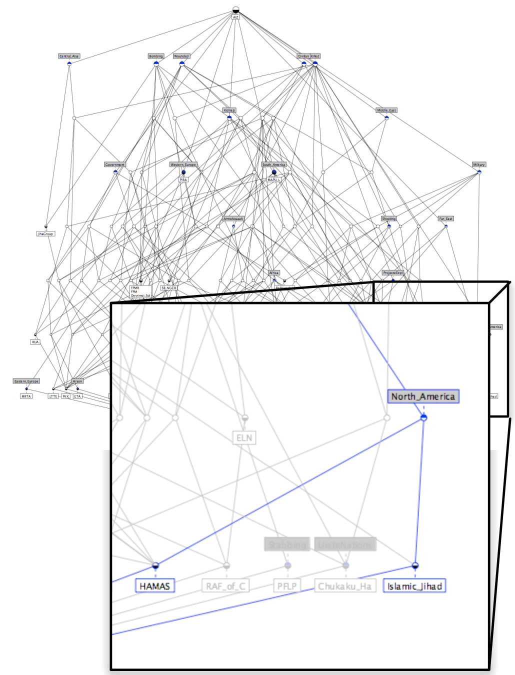

Figure 3. Terror incident database sample.

The highlighted section in the figure shows the result of a query concerning the area of operations. The knowledge base shows that 'HAMAS' and 'Islamic_Jihad' are the only groups operating in North America up to 1998. Similar queries show that HAMAS uses bombing as a tactic, impelling the analyst to key on information concerning the loss or theft of explosive material, as an example.

Like the case of the terror incident database, it is possible to process phone calls, bank transactions and other information bases to convert the information into knowledge permitting us to interpret data being collected by the analyst. Instead of seeing a series of transactions, we see linkages between banks known to launder money, the bank of a suspect, Confederate Bank, and a link to a new account / suspect at a Maryland bank, see Figure 4.

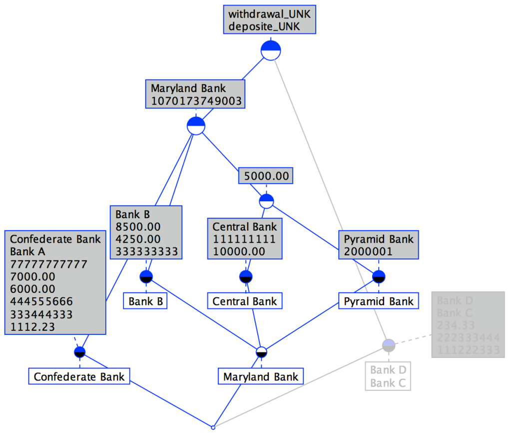

Figure 4. Knowledge associated with bank transactions.

Forensics analysis

A second decision support area explored involved a problem of nuclear forensics. In this problem we have databases of assayed nuclear material as reference samples. The problem involves determining the origin of an unknown nuclear material. The complexity of the problem can be rather daunting. The material carries its entire history in its chemical composition. The mining, enrichment, fabrication, operational history, and its reprocessing affects sample compositions.

This problem was appealing because it had the potential of breaking the algorithms as implemented within the PDA system. The sample data consisted of over 100 real valued attributes which were fuzzified into over 500 qualitative attributes for describing the materials in the database.

Figure 5 captures some of the complexity of the information contained in this knowledge base. The lattice does show that we can uniquely define the source of interdicted material. The bottom row of the lattice has separate instances for each reactor which is why we can make that assertion. Additional assessments will be made on expanded data sets to explore the robustness of the algorithms and the data bases.

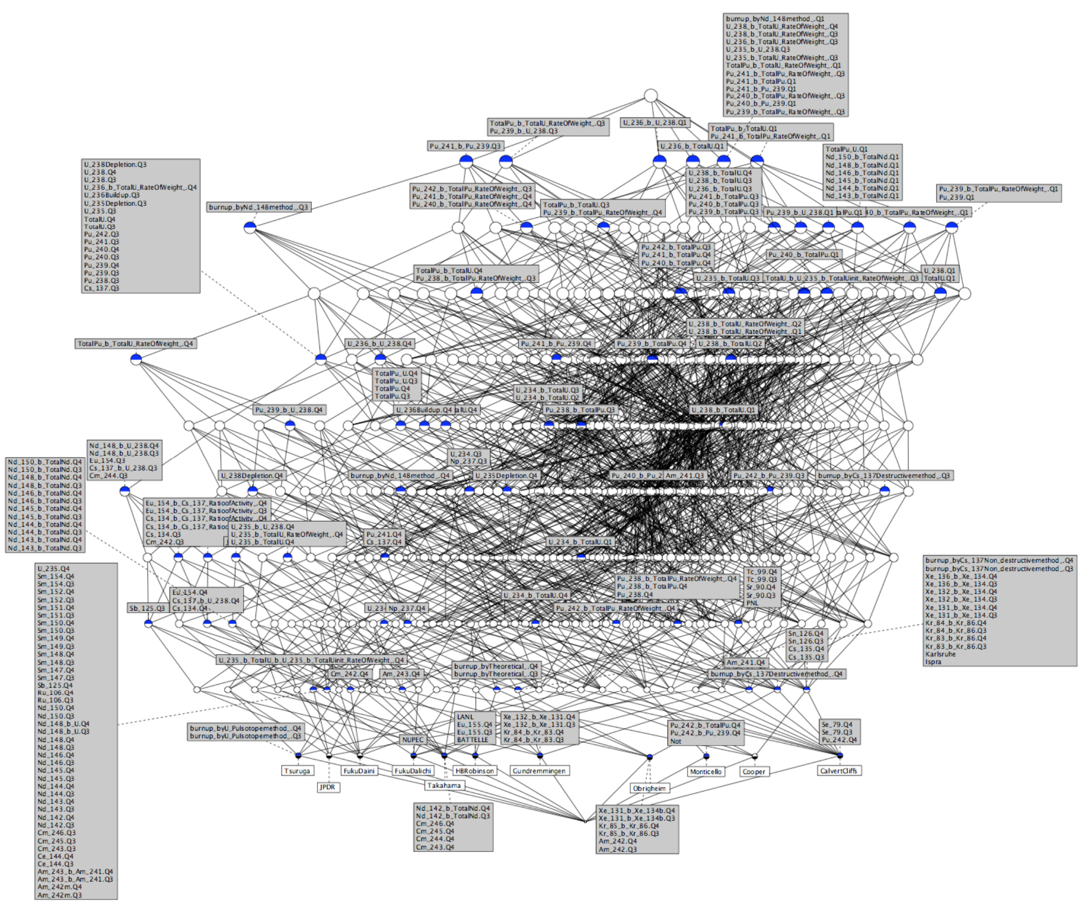

Figure 5. Example of a forensic lattice. (for demonstration of structural complexity)

Preliminary Qualitative Evaluation

Using the theories and technologies of Peircean reasoning provides the analyst with knowledge that can assist their daily activities as opposed to adding to their cognitive load. What we have is a capability that can process massive amounts of information that is likely to over whelm a decision maker faced with classical decision support technologies such as an air traffic control system. In these types of system raw data is presented to the decision maker who must reason or internally fuse the information provided. The inductive learning engine performs a function that produces knowledge that can be used to evaluate data that enters a decision makers field of perception.

The structure of the solution developed is a very 'transparent' system. This transparency was part of the initial requirements to enable confidence to be developed by a user of the system. The abductive component is easily validated by examination of the virtual belief cache that is constructed. In this cache, the hypothesis support is provided along with implicit support and unresolved information, permitting a decision maker to validate in real time, conclusions generated. Validation of the inductive learning engine is a continuing activity. We are continually searching for more difficult and complex problems to employ in the overall system. The validation approach is to use the variability of the engineered solution in a design of experiment approach to assess the limits of the knowledge constructed by the inductive engine and evaluate based on the abductive solutions generated based on data presented to the system.

What has been produced in this effort is a robust flexible decision support functionality that has its roots in reasoning, knowledge representation and logic theory. The system is a hybrid solution using these technologies in a manner in which the best technology is matched to function. We have applied the integrated solution to a number problems utilizing differing process methodologies and have attempted to break the system by going well beyond toy problems. While the solution is a 70 - 80% solution we already have seen the system produce solutions that produced varying levels of surprise, in terms of insights provided in complex analysis domains.

References

1. V. G. Blinova, D. A. Dobrynin, V. K. Finn, S. O. Kuznetsov, e. al., "Toxicology analysis by means of the JSM-method," in Bioinformatics, 2003.

2. B. A. Davey and H. A. Priestley, "Introduction to Lattices and Order," 2002, pp. 310.

3. V. K. Finn, T. Gergely, and S. O. Kuznetsov, "Plausible reasoning for open problem domains," in Intelligent Control, 1996.

4. B. Ganter, "Attribute exploration with background knowledge," Theoretical Computer Science, 1999.

5. B. Ganter, R. Wille, T. H. Darmstadt, and F. Mathematik, "Applied Lattice Theory: Formal Concept Analysis," in cl.inf.tu-dresden.de, 1997.

6. E. Heit, "Properties of inductive reasoning," in Psychonomic Bulletin & Review, 2000.

7. S. O. Kuznetsov, "Machine Learning on the Basis of Formal Concept Analysis," in Automation and Remote Control, 2001.

8. J. S Mill, "A system of logic, ratiocinative and inductive," 1856.

9. E. T. Nozawa, "Peircean semeiotic: a new engineering paradigm for automatic and adaptive intelligent systems design," Information Fusion, 2000.

10. C. S. Peirce, "The Essential Peirce: Selected Philosophical Writings. Volume 2 (1893-1913) - Page 361," 1998, pp. 640.

11. C.S Peirce, J. Hoopes (ed.), "Peirce on Signs: Writings on Semiotic - Page 189," 1991, pp. 284.

12. C. S Peirce, K.L. Ketner (ed.), "Reasoning and the Logic of Things: The Cambridge Conferences Lectures of 1898," 1992, pp. 297.

13. M.E. Senglaub, "Automated Decision Support in a Complex Information Space", 13th ICCRTS, June, 2008.