Multitype Inference in Multistrategy Task-adaptive Learning: Dynamic Interlaced Hierarchies

Michael R. Hieb and Ryszard S. Michalski
Center for Artificial Intelligence
George Mason University, Fairfax, VA
hieb@aic.gmu.edu and michalski@aic.gmu.edu

Abstract

Research on multistrategy task-adaptive learning aims at integrating all basic inferential learning strategies—learning by deduction, induction and analogy. The implementation of such a learning system requires a knowledge representation that facilitates performing a multitype inference in a seamlessly integrated fashion. This paper presents an approach to implementing such multitype inference based on a novel knowledge representation, called *Dynamic Interlaced Hierarchies* (DIH). DIH integrates ideas from our research on cognitive modeling of human plausible reasoning, the Inferential Theory of Learning, and knowledge visualization. In DIH, knowledge is partitioned into a "static" part that represents relatively stable knowledge, and a "dynamic" part that represents knowledge that changes relatively frequently. The static part is organized into type, part, or *precedence* hierarchies, while the dynamic part consists of *traces* that link nodes of different hierarchies. By modifying traces in different ways, the system can perform different *knowledge transmutations* (patterns of inference), such as generalization, abstraction, similization, and their opposites, specialization, concretion and dissimilization, respectively.

Key words: multistrategy learning, inferential theory of learning, knowledge transmutation, generalization, abstraction, similization.

1. Introduction

The development of multistrategy learning systems requires a powerful and easily modifiable knowledge representation that facilitates multitype inference. This is particularly true in the case of *multistrategy task-adaptive learning* (MTL) systems that integrate a whole range of inferential strategies, such as empirical induction, abduction, deduction, plausible deduction, abstraction, and analogy (Michalski, 1990, 1991; Tecuci and Michalski, 1991; Tecuci, 1993). A MTL system adapts a strategy or a combination of strategies to the *learning task*, defined by the available input knowledge, the learner's background knowledge and the learning goal. A theoretical framework for the development of MTL systems has been presented in (Michalski, 1993).

This paper presents basic ideas underlying a knowledge representation proposed for the implementation of a MTL system and its use for implementing multitype inference. This representation, called *Dynamic Interlaced Hierarchies* (DIH), integrates ideas from our research on modeling human plausible inference, the Inferential Theory of Learning and the visualization of knowledge. DIH

encompasses many different forms of knowledge – facts, rules, dependencies, etc., and facilitates knowledge transmutations, described in the Inferential Theory of Learning (ITL) (Michalski, 1993). This paper shows how DIH supports several basic patterns of knowledge change (transmutations), such as generalization, abstraction, similization, and their opposites, specialization, concretion and dissimilization, respectively. These operations are performed on DIH traces, which correspond to well-formed predicate logic expressions associated with a degree of belief.

While our previous work has focused on the visualization of attribute-based representations for empirical induction (Wnek & Michalski, 1991), DIH allows the visualization of structural (attributional and relational) representations. The underlying assumption is that the syntactic structure for representing any knowledge should reflect as closely as possible the semantic relationships among the knowledge components, and facilitate knowledge modifications that correspond to the most frequently performed inferences. An early implementation of this idea was in the ADVISE system, which used three forms of knowledge representation: relational tables, networks and rules (Michalski et al., 1986).

The DIH approach assumes that a large part of human conceptual knowledge is organized into various hierarchies, primarily type, part and precedence hierarchies (see Section 3 for an explanation). Such hierarchies reflect frequently occurring relationships among knowledge components, and make it easy to perform basic forms of inference.

The initial idea for DIH stems from the core theory of human plausible reasoning (Collins & Michalski, 1989; Boehm-Davis, Dontas & Michalski, 1990). The theory presents a formal representation of various plausible inference patterns observed in human reasoning.

DIH is more fully described in (Hieb & Michalski, 1993).

2. Relevant Research

The core theory of Plausible Reasoning presents a system that formalizes various plausible inference patterns and "merit parameters" that affect the certainty of these inferences. This system combines structural aspects of reasoning (determined by knowledge structures) with parametric aspects that represent quantitative belief and other measures affecting the reasoning process.

Various components of the "Logic of Plausible Reasoning" have been implemented in several systems (Baker, Burstein & Collins, 1987; Dontas & Zemakova, 1988; Kelly, 1988). These implementations used various subsets of the inferences ("statement transforms") described in the core theory to investigate the parametric aspects of the theory. The implementations demonstrated how the core theory of plausible reasoning can be applied to various domains. DIH specifies a broader set of knowledge transmutations in a general and well-defined knowledge representation. These transmutations are part of a framework for both reasoning and learning.

The organization of concepts into various hierarchies has been proposed as a plausible structure for human semantic memory quite early (Collins & Quillian, 1972). The WordNet project at Princeton University, directed by George Miller, concerns the implementation of an electronic thesaurus using such a memory structure (Beckwith et

al., 1991). WordNet is a very large lexical database with approximately 50,000 different word forms. WordNet divides the lexicon into various categories including nouns, verbs, and modifiers (adjectives and adverbs). Significantly, the nouns are stored in topical hierarchies (both type and part hierarchies), lending support to the DIH representation. However, while WordNet can be used as a source of DIH hierarchies, it does not provide any inferential facilities.

Other relevant research includes the development of the Common Knowledge Representation Language (CKRL), done as part of an ESPRIT project (Morik, Causse & Boswell, 1991). CKRL offers a language in which knowledge can be exchanged between machine learning tools and it uses the set of *most* common representation structures and operators. While CKRL's representation for multistrategy learning seeks to integrate the various representations employed by several different learning programs for communication of knowledge between the machine learning tools, our aim is to develop a representation that facilitates an integration of learning and inference processes.

Semantic network knowledge representation systems, such as the KL-ONE family (Brachman et al., 1991), utilize a large network of relationships between concepts, intermixing different relationships. The hierarchies they use are tangled, in which a concept can have more than one parent. As a consequence, implementing knowledge transmutations, e.g., generalization, is not as easy as in DIH. DIH facilitates such transmutations because it uses only single-parent hierarchies, representing a structuring of a set of entities from a certain viewpoint. In DIH, a concept can belong to different hierarchies, reflecting the fact that a given concept (or object) can usually be classified from several different viewpoints.

The design of semantic networks is primarily oriented toward facilitating deductive inference, and is not usually concerned with knowledge visualization. The design of DIH is oriented toward facilitating multitype inference and providing a basis for the visual presentation of knowledge. DIH also utilizes a hierarchy of merit parameters to represent probabilistic factors associated with plausible reasoning.

3. Basic Components of DIH

The theory of plausible reasoning postulates that there are recurring patterns of human plausible inference. To adequately represent these patterns, one needs a proper knowledge representation. The DIH approach partitions knowledge into a "static" part and "dynamic" part. The static part represents knowledge that is relatively stable (such as established hierarchies of concepts), and a "dynamic" part that represents knowledge that changes relatively frequently (such as statements representing new observations or results of reasoning). The static part is organized into *type* hierarchies (TH), *part* hierarchies (PH) and *precedence* hierarchies. Precedence hierarchies include several subclasses, specifically, *measure* hierarchies (MH), *quantification* hierarchies (QH) and *schema* hierarchies (SH). The dynamic part consists of *traces* that represent knowledge involving concepts from different hierarchies. Each trace links nodes of two or more hierarchies and is assigned a degree of belief.

These hierarchies are composed of nodes representing abstract or physical entities, and links representing certain basic relationships among the entities, such as "type-of", "part-of" or "precedes". In the "pure" form, these

hierarchies are single parent, that is, no node can have more than one parent. The root node is assigned the name of the class of entities that are organized into the hierarchy from a given viewpoint.

A type (or generalization) hierarchy organizes concepts in a given class according to the "type-of" relation (also called a "generalization" or "kind-of" relation). For example, different types of "animals" can be organized into a "type" hierarchy.

A part hierarchy organizes entities according to a "part-of" relationship. For example, the world, viewed as a system of continents, geographical regions, countries, etc., can be organized into a part hierarchy. While properties of a parent node in the type hierarchy are inherited by children nodes, this does not necessarily hold for a part hierarchy. There are several different part relationships, which include part-component, part-member, part-location and part-substance (Winston, Chaffin and Herrmann, 1987).

To represent relationships among elements of ordered or partially ordered sets, a class of *precedence* hierarchies is introduced. Hierarchies in this class represent hierarchical structures of concepts ordered according to some precedence relation, such as "A precedes B", "A is greater than B", "A has higher rank than B", etc.

There are several subclasses of precedence hierarchies. One subclass is a *measure* hierarchy, in which leafs stand for values of some physical measurement, for example, weight, length, width, etc., and the parent nodes are symbolic labels characterizing ranges of these values, such as "low", "medium", "high", etc. Figure 1 shows a measure hierarchy of possible values of people's height. Dotted lines indicate a continuity of values between nodes. Arrows indicate the precedence order of the nodes. Another subclass hierarchy is a *belief* hierarchy, in which nodes represent degrees of an agent's beliefs in some knowledge represented by a trace.

Other subclasses of precedence hierarchies include a *rank* hierarchy and a *quantification* hierarchy. A rank hierarchy consists of values representing the "rank" of an entity in some structure, e.g., an administrative hierarchy or military hierarchy. A quantification hierarchy consists of nodes that represent different quantifiers for a set (An example is shown in Figure 2). A quantification hierarchy that is frequently used in commonsense reasoning includes such nodes as "one", "some" (corresponding to the existential quantifier), "most", and "all" (corresponding to the universal quantifier).

Each hierarchy has a heading that specifies its kind (TH, PH, MH, QH or SH) and the underlying concept (or viewpoint) used for the creation of the hierarchy. In addition, the type and part hierarchies also have a *top* node that in the type hierarchies stands for the class of

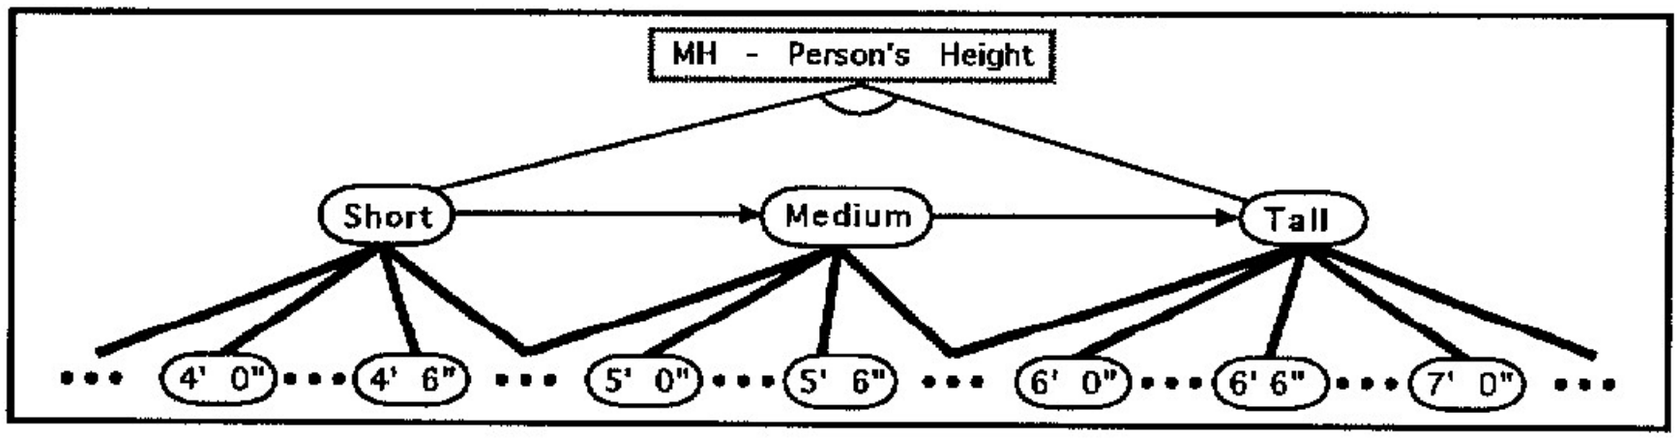

Figure 1: A measure hierarchy of values characterizing people's height.

all entities in the hierarchical structure, and in the part hierarchies for the complete object.

Schema hierarchies (or schema) are structures that indicate which hierarchies are connected in order to express multi-argument concepts or relationships. For example, the schema hierarchy for the concept of "physical-object" can be <shape, size>. This states that an attribute "shape" applies to any object that is a "physical-object" (a node in the "physical-object" hierarchy), and produces a shape value, which is a node in the "shape" hierarchy. The schema hierarchy for the concept of "giving" may be <giver, receiver, object, time> that states that this concept involves an agent that gives, an agent that receives, an object that is being given, and the time when the "giving" occurs. The agents, object and time are elements of their respective hierarchies.

DIH also makes a distinction between structural and parametric knowledge. The structural knowledge is represented by hierarchies and traces that link nodes of different hierarchies. Parametric knowledge consists of numeric quantities characterizing structural elements of knowledge. In DIH, this knowledge is represented via precedence hierarchies of *merit parameters*. The basic merit parameter is a belief measure that characterizes the "truth" relationship of a given component of knowledge representation (a trace), as estimated by the reasoning agent. Other merit parameters include the *forward* and *backward* strength of a dependency, *frequency*, *dominance*, etc. (Collins and Michalski, 1989; Michalski, 1993). In this paper, we will consider only one merit parameter, namely, the belief measure.

The theory of human plausible reasoning (Collins and Michalski, 1989) postulates that people rely primarily on the structural knowledge, and resort to parametric knowledge when the "structural" reasoning does not produce a unique result. They resist performing uncertain inferences based on only parametric knowledge, and they are not good at assigning a degree of certainty to a statement based only on the combination of the certainties of its constituents, without taking into consideration the meaning of the whole sentence. A reason for this may be that there does not exist a normative model for reasoning under uncertainty that is independent of the structural aspects of knowledge, i.e., its meaning. Plausible reasoning about a problem or question typically involves both structural and parametric knowledge components.

Nodes of a hierarchy are elementary units of the DIH representation. Each node represents some real or abstract entity—a concept, an object, a process, etc. A given entity can be a node in multiple hierarchies, where each hierarchy structures a set of entities from a different viewpoint. The relevant viewpoint is determined by the context of the discourse.

As mentioned earlier, the basic structures in the DIH representation are hierarchies, nodes, traces and schema. Our research on DIH demonstrates that these structures provide a very natural environment for performing basic types of inference on statements. The subsequent sections show how these inferences are performed using the DIH representation.

4. DIH Traces

To describe the DIH knowledge representation, let us start by representing the following statement: "It is certain that some power plants in New York have mechanical

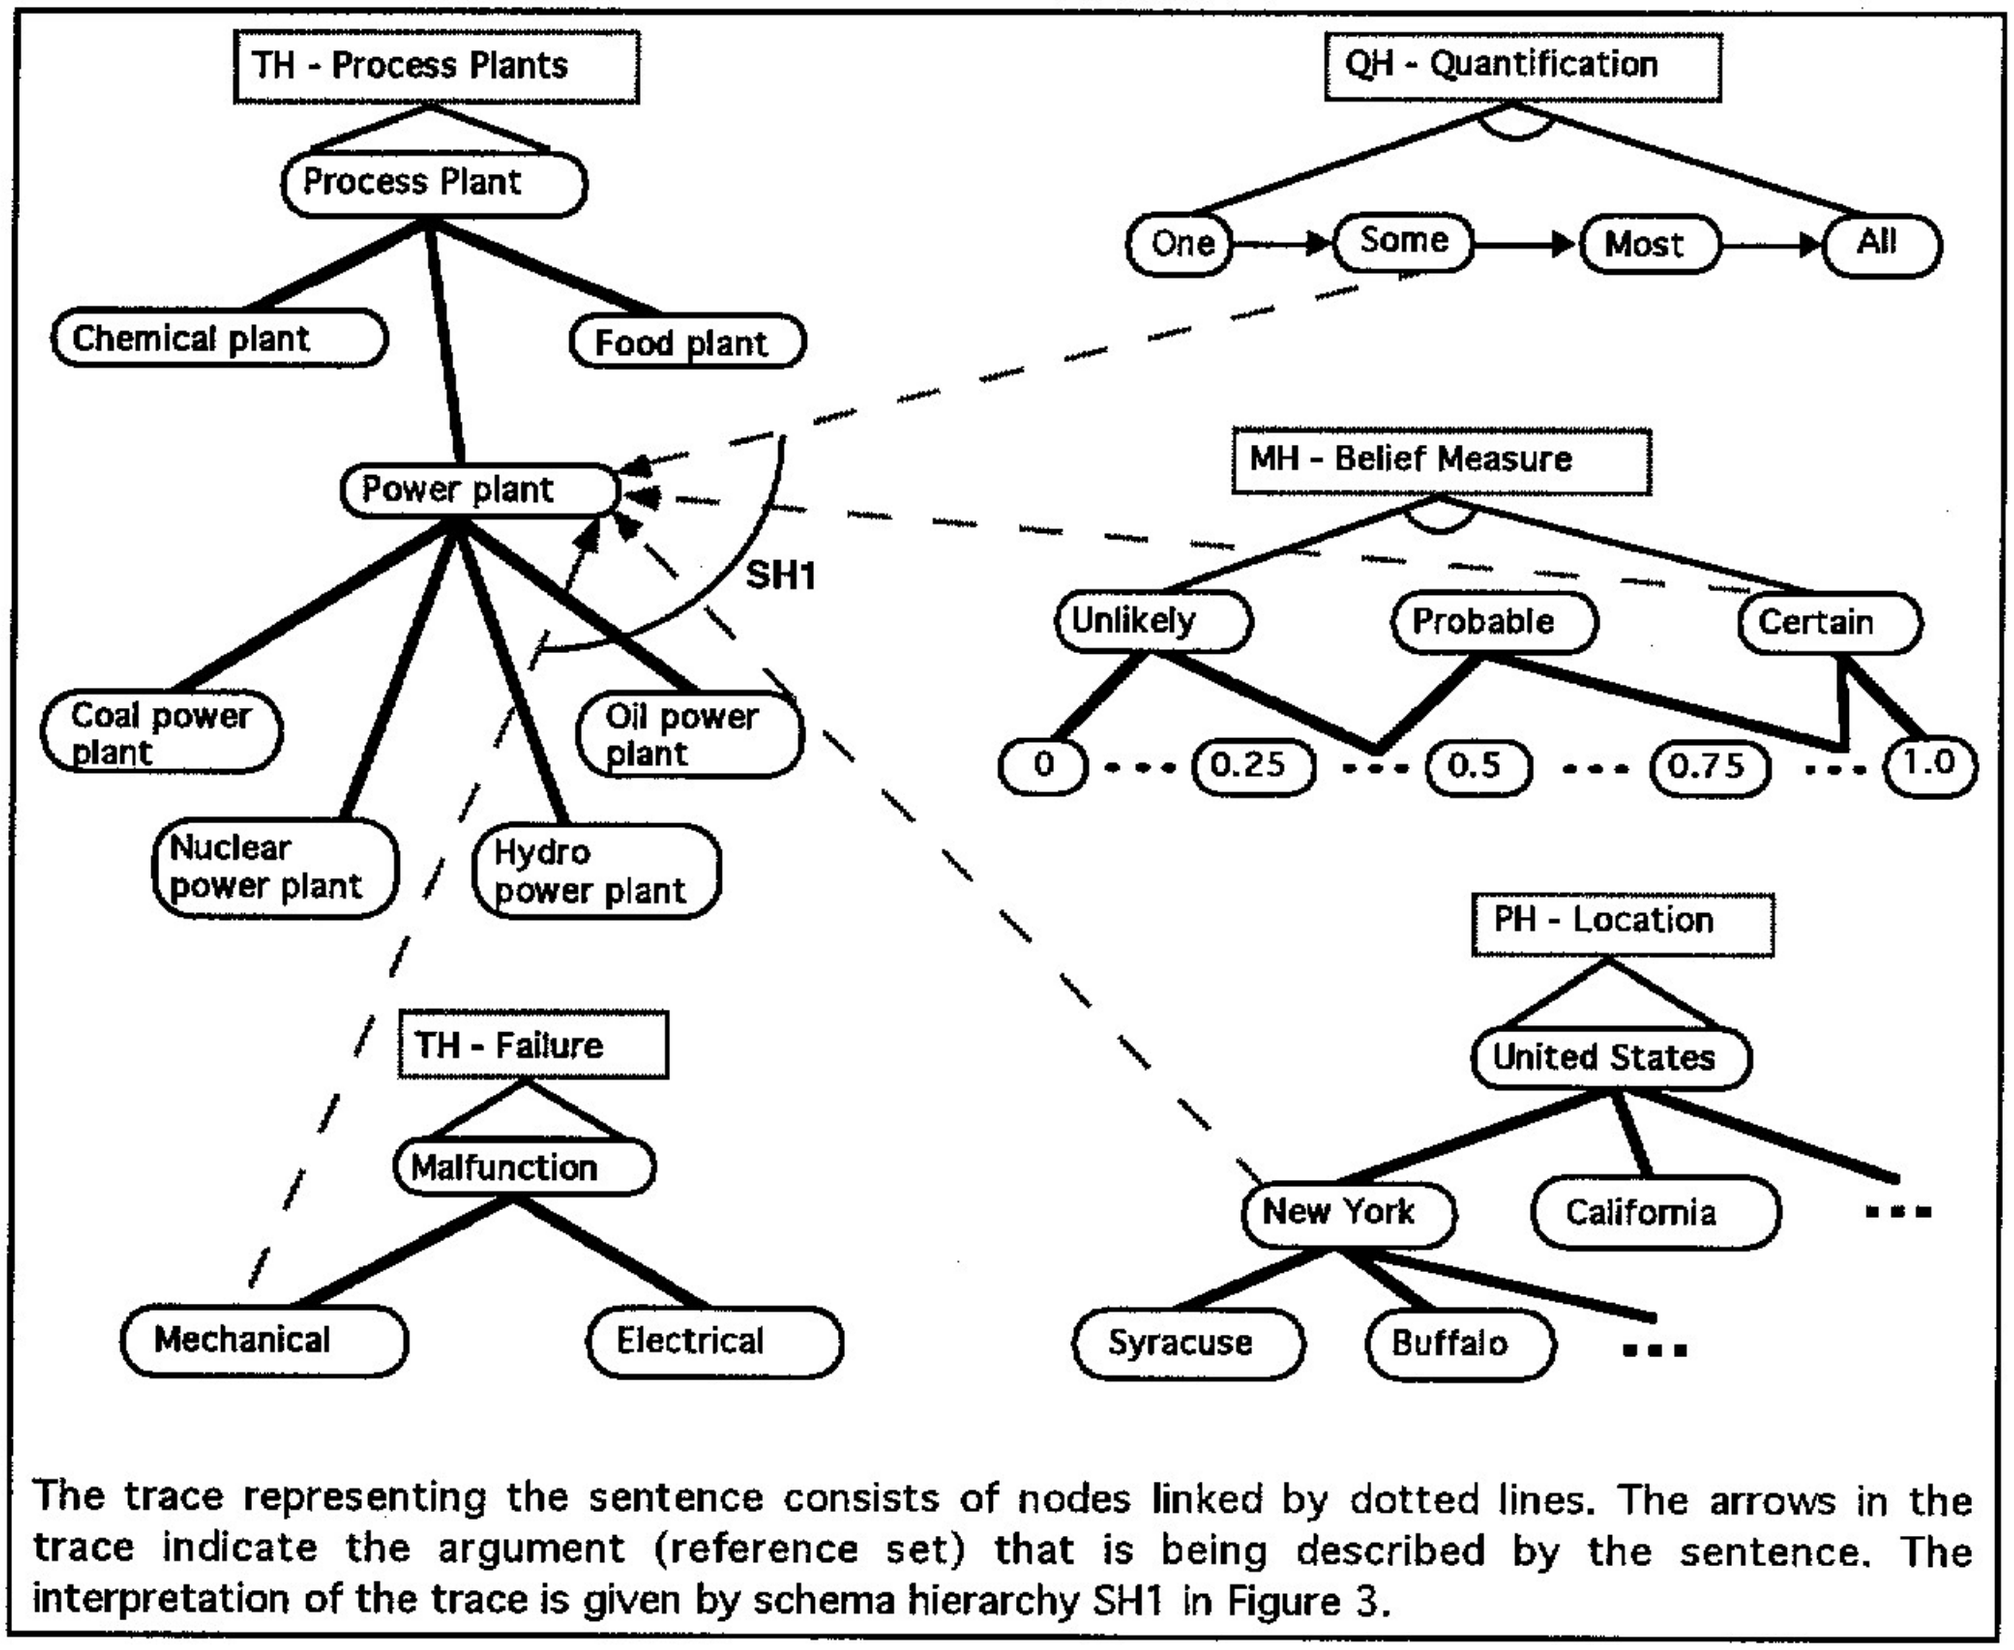

The trace representing the sentence consists of nodes linked by dotted lines. The arrows in the trace indicate the argument (reference set) that is being described by the sentence. The interpretation of the trace is given by schema hierarchy SH1 in Figure 3.

Figure 2: A DIH trace representing the sentence "It is certain that some power plants in New York have mechanical failures."

failures." Figure 2 presents this statement as a trace connecting nodes of five hierarchies: "Process plants" and "Failure", both type hierarchies; "Quantification", the quantification hierarchy; "Location", a part hierarchy; and "Belief measure" a measure hierarchy.

The interpretation of the trace is done on the basis of the schema hierarchy shown in Figure 3. The schema defines the universe of sentences that can be generated using concepts of these hierarchies, ordered according to the schema.

The convention for the direction of arrows in a trace is that they point from the nodes

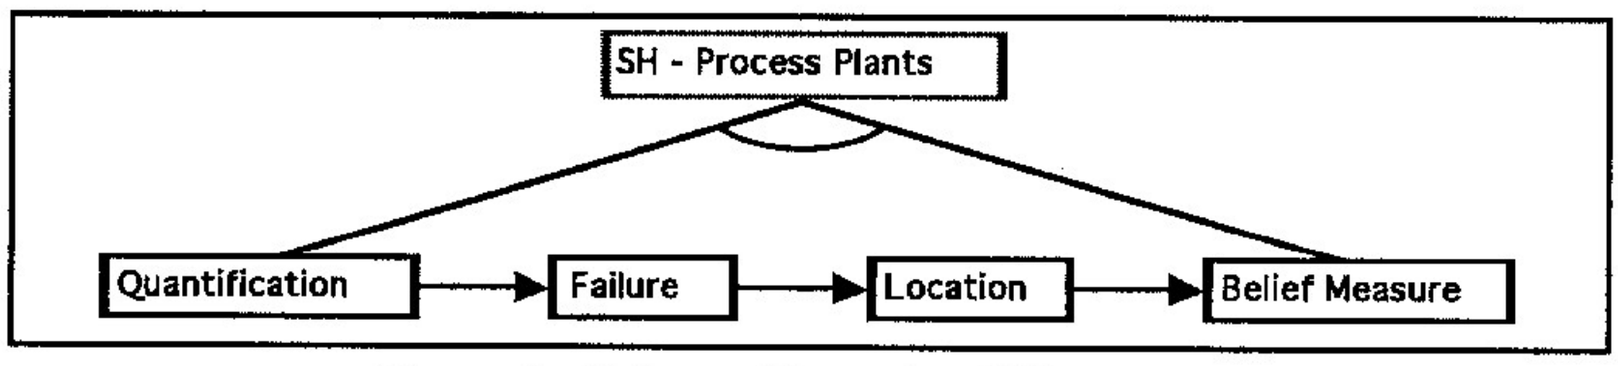

Figure 3: Schema hierarchy SH1.

denoting descriptive concepts to the *argument node* that stands for the set (or individual) being described, called a *reference set.* In this example, the set being described is "Power plant" in the hierarchy of Process Plants, thus the node representing it is the argument node. Other nodes linked by the trace represent descriptive concepts for the argument node. The belief measure takes values from a belief hierarchy, and refers to the entire trace rather than a single node, which is indicated by the schema.

Using the formalism of the annotated predicate logic (Michalski, 1983), this trace can be interpreted as: "(Some)x, [type(x) = Power plant) & [location(x) = New York] & [failure(x) = mechanical]: Belief = 1.0." This statement is a quantified conjunction of several *elementary* statements. An elementary statement expresses one property of the reference node (set), for example, "Location(Power plant) = New York."

In a formal expression of an elementary statement, the reference set ("Power plant") is called an *argument,* the predicate ("Location") is called a *descriptor*, and the value of the descriptor ("New York") is called the *referent.* Thus, an elementary statement is formally expressed in the form "descriptor(argument) = referent".

In Figure 2, the square boxes contain the heading of the hierarchy. The concept specified in the heading is the general descriptor for the hierarchy. The nodes in the hierarchy are possible values of this descriptor.

The schema hierarchy, SH1, in Figure 3 is used for the interpretation of the trace represented in Figure 2. The heading indicates the type of hierarchy (SH: Schema Hierarchy) and the reference set of the trace. Since the schema hierarchy is a precedence hierarchy, a valid interpretation of the schema requires each of the descriptors in order. Thus the first element of the trace must be from the quantification hierarchy, the second from the failure hierarchy, the third from the location hierarchy and the last from the hierarchy of belief measures. This schema hierarchy is also utilized for examples in Section 4.

Adding knowledge to the DIH representation is done by creating hierarchies and specifying traces that express statements involving nodes of different hierarchies. To allow proper interpretation of a trace, the schema is also specified by indicating relevant descriptors and their order.

DIH allows one to represent complex forms of knowledge, involving different kinds of quantifiers, multi-argument predicates, different types of logical operations on them, and to associate degrees of belief with individual statements. A more complete description of the DIH representation system is given in (Hieb & Michalski, 1993).

5. Multitype Inference in DIH

The core theory of plausible reasoning introduced in (Collins & Michalski, 1989) gives four knowledge transmutation operators (also called transforms) – generalization, specialization, similization and dissimilization. The Inferential Theory of Learning (Michalski, 1993) specifies several additional operators, of which abstraction and concretion are incorporated into DIH. (In (Collins and Michalski, 1989), the abstraction and concretion transmutations were called referent generalization and referent specialization, respectively.)

Transmutation	Symbol	Relevant Hierarchies	Inference Type
Argument Generalization	AGen	Type, Part	Deductive
Argument Specialization	ASpec	Type, Part	Inductive
Quantification Generalization	QGen	Quantification	Inductive
Quantification Specialization	QSpec	Quantification	Deductive
Abstraction	Abs	Type, Part, Precedence	Deductive
Concretion	Con	Type, Part, Precedence	Inductive
Argument Similization	ASim	Type, Part	Analogical
Argument Dissimilization	ADis	Type, Part	Analogical
Referent Similization	RSim	Type, Part, Precedence	Analogical
Referent Dissimilization	RDis	Type, Part, Precedence	Analogical

Table 1: Basic knowledge generation transmutations.

Generalization (specialization) transmutations extend (contract) the reference set. They are done either by argument generalization (specialization) or by quantification generalization (specialization). Argument generalization is accomplished by moving above the node representing the reference set in a type hierarchy. Quantification generalization is accomplished by moving up the quantification hierarchy.

Abstraction (concretion) transmutations decrease (increase) the amount of information about the reference set. A way to accomplish such a transmutation is by moving above the node in the type or part hierarchy that corresponds to a value of some descriptor in the sentence represented by the trace.

Similization (dissimilization) transmutation is done by replacing a node corresponding to the reference set (argument) or a descriptor value (referent) by a node at the same level of hierarchy, which corresponds to a similar (dissimilar) concept within the context of the given hierarchy. In the case of dissimilization, the resulting trace is linked with a negation node, because the generated inference is a negation of the original sentence (Michalski, 1993).

These transmutations can be given a simple conceptual interpretation, if one assumes that nodes at each level of hierarchy are ordered by the relation of similarity, that is, nodes that correspond to similar concepts (in the context of the given hierarchy) are located near each other, and nodes that correspond to dissimilar concepts are placed far away from each other. Such an arrangement is natural for precedence hierarchies. In sum, similization and dissimilization transmutations are performed by sideways node movements, while generalization (specialization) and abstraction (concretion) are performed by upward (downward) node movements.

Table 1 lists all the above knowledge transmutations, specifying their abbreviated name, the relevant hierarchies, and the underlying inference type. The relevant hierarchies are the kinds of hierarchies for which the transmutations are valid. The various kinds of part hierarchies are not shown, but are distinguished in DIH. Additional constraints are necessary in some kinds of part hierarchies to maintain the validity of the transmutation.

Figure 4 presents a schematic diagram illustrating how knowledge transmutations

modify a trace. A dotted line represents a link in a trace. An arrow means that the trace is moving to a new node in the indicated direction by performing the indicated transmutation. The quantification transmutations operate over the entire trace, rather than on a single node, as do the transformations involving the merit parameters.

One form of generalization transmutation moves a node in the quantification hierarchy upward, another form moves a node (argument) in the type hierarchy upward. The "+" indicates a strengthening of a merit parameter, or the movement of the link to a node that is "higher" in the particular merit parameter measure hierarchy. The "–" indicates a weakening of the merit parameter, or the movement of the link down in the hierarchy.

Moving a node in a trace in a manner that corresponds to a deductive inference (Table 1) produces a new trace (statement) with the same truth status as the original trace. In the case of node movement that corresponds to inductive or analogical inference, the smaller the node movement ("perturbation"), the more plausible the resulting inference.

The Argument Generalization transmutation represents a deductive inference. The abstraction operation is also deductive. In contrast, Argument Specialization, Quantification generalization and Concretion are inductive, because they produce traces (statements) that logically entail the original traces (statements).

The above transmutations can be usually done in a number of different ways, by moving to different alternative nodes. The plausibility of the generated statements depends on additional merit parameters, such as dominance, typicality, multiplicity, similarity, frequency, etc. (Collins and Michalski, 1989). These issues will be the subject of future research.

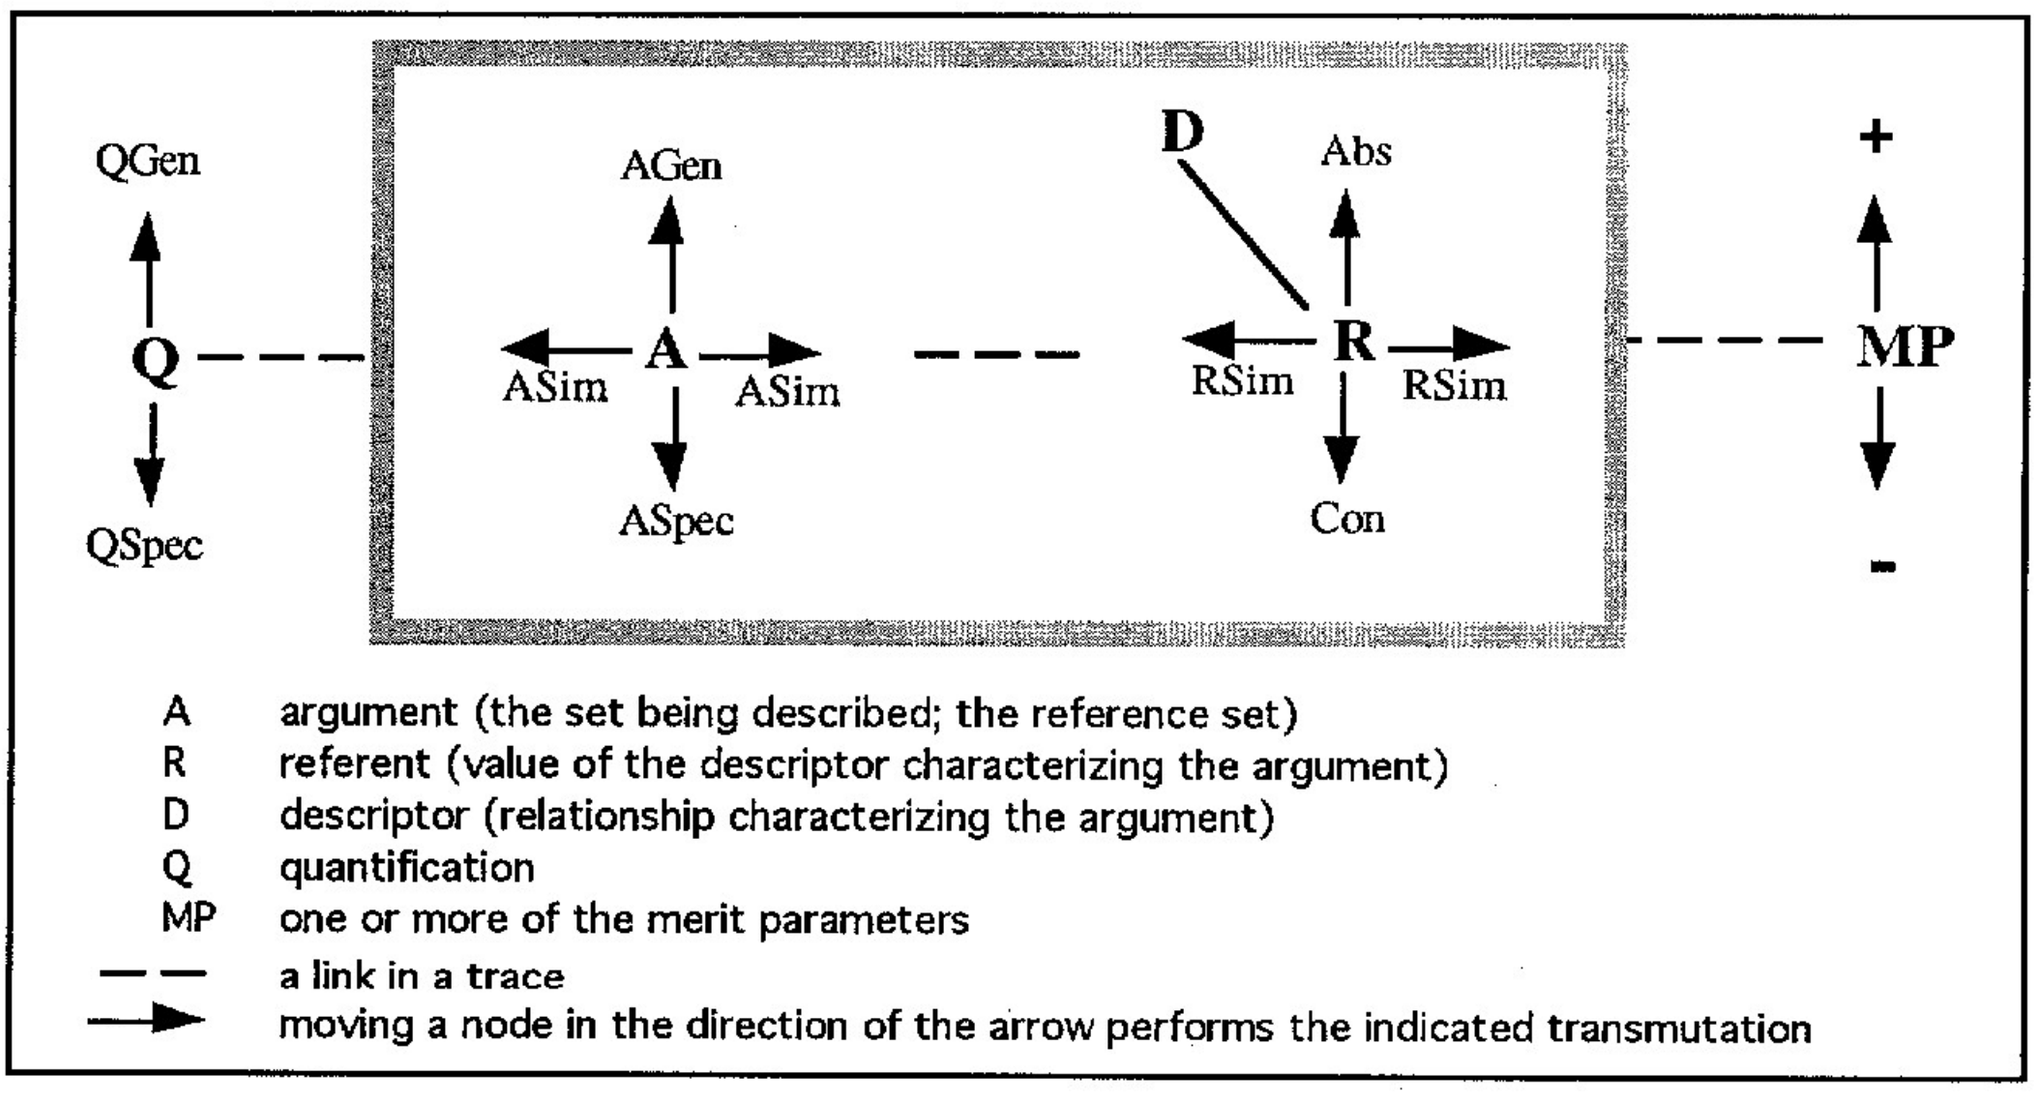

Figure 4: Diagram of knowledge transmutations in DIH.

6. Visualizing DIH-Based Inference

This section illustrates several basic transmutations through a series of self-explanatory examples. These examples involve the same original statement, represented as a trace in Figure 2. Given the original statement, these transmutations generate new statements illustrated by DIH traces in Figures 5 through 12.

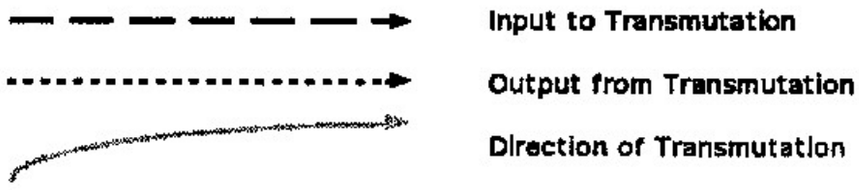

The legend above is used for interpreting the following figures. The input statement is the same as that of Figure 2, without the belief measure hierarchy. All of the examples are interpreted according to the schema SH1 shown in Figure 3.

There are two referents in the input statement. The resulting statements (output) show the results of the given transmutation assuming that there are no merit parameters that assist in the specialization or concretion and that the similization operator finds a single "most similar" node using the descriptors given. The Background Knowledge (BK) is the learner's prior knowledge that is relevant to the learning process.

7. MTL-DIH System

The research on DIH aims at developing a representation that will facilitate all basic inferential strategies and knowledge transmutations to be implemented in the multistrategy task-adaptive learning system (MTL-DIH).

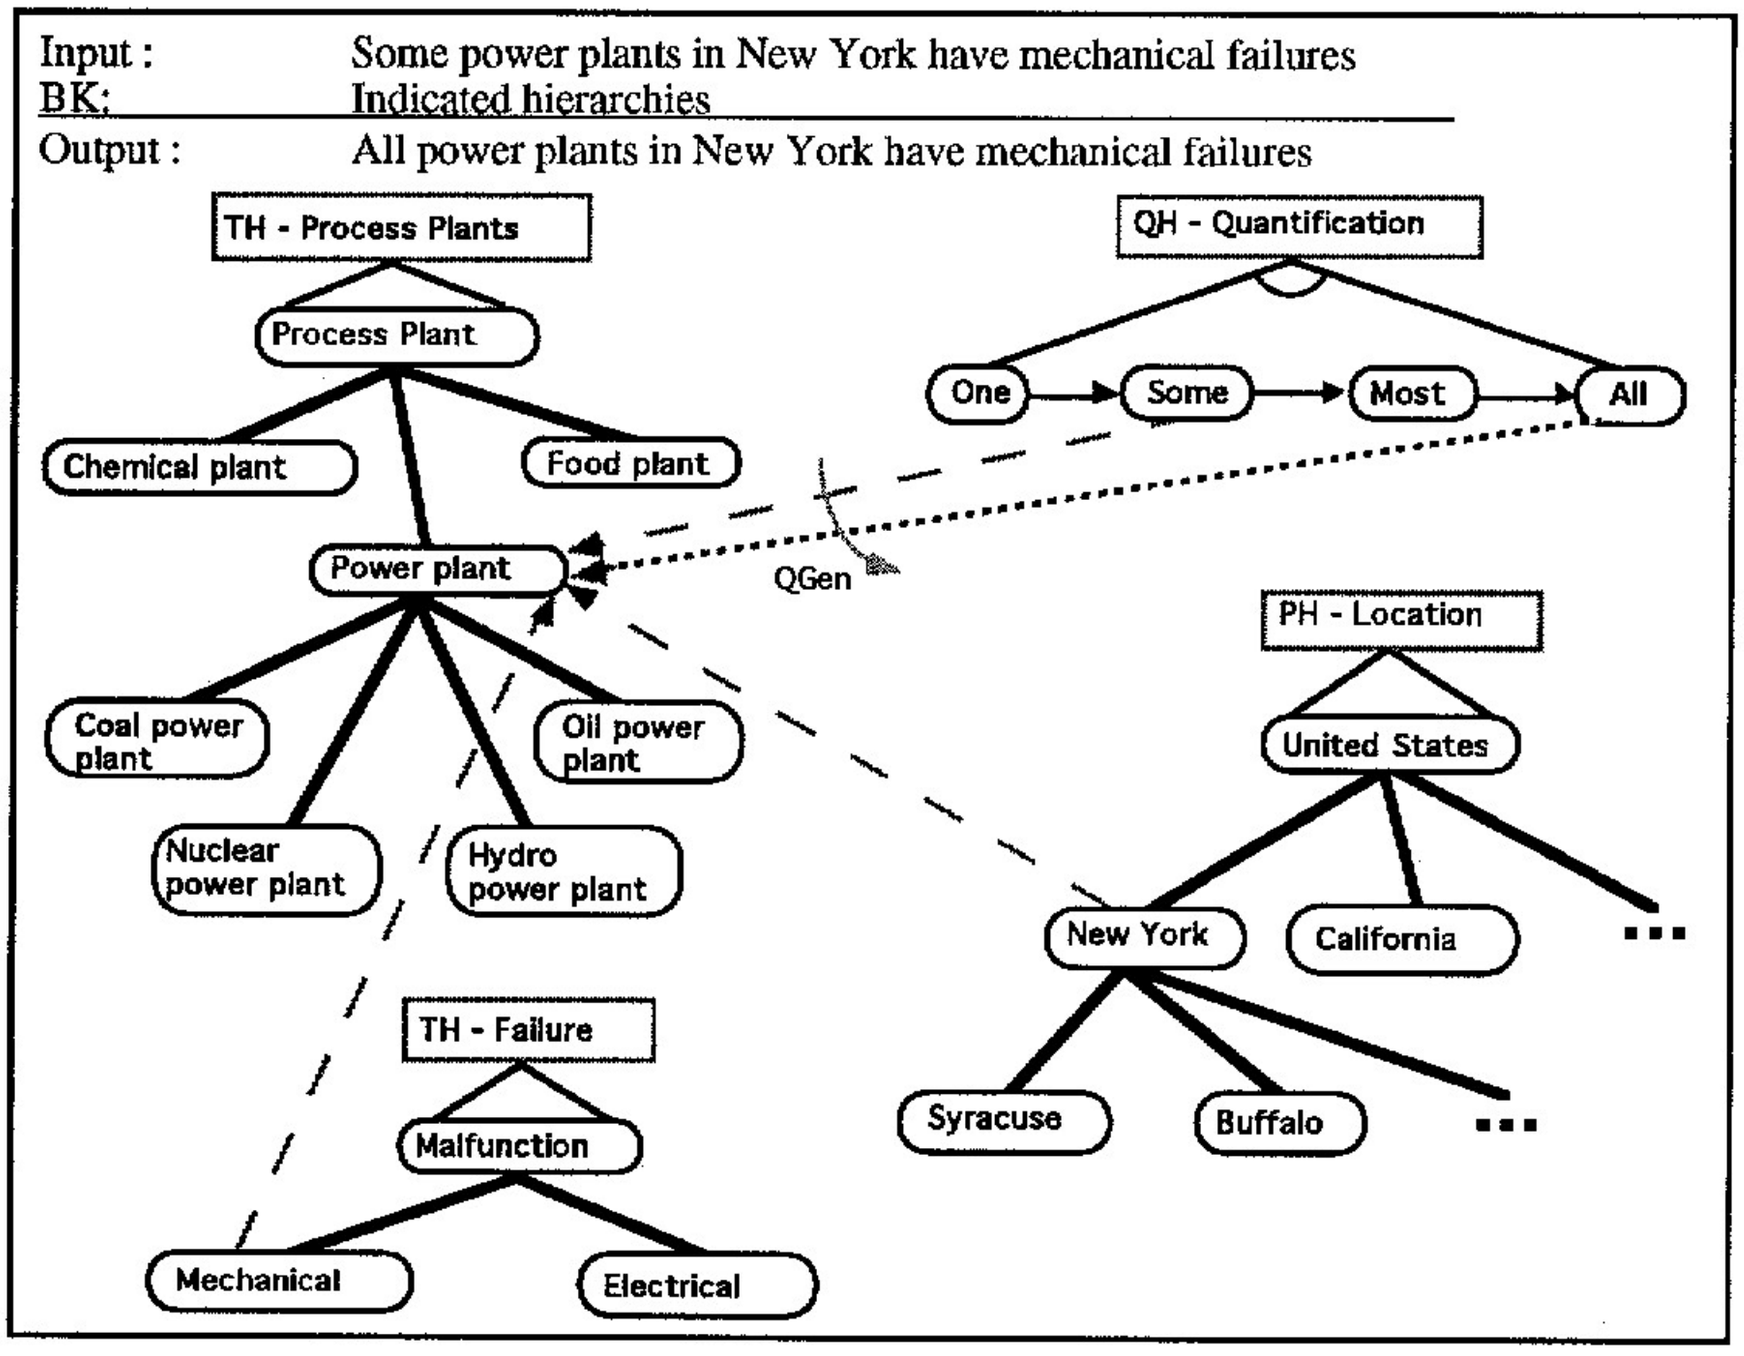

Figure 5: Inductive generalization based on quantification.

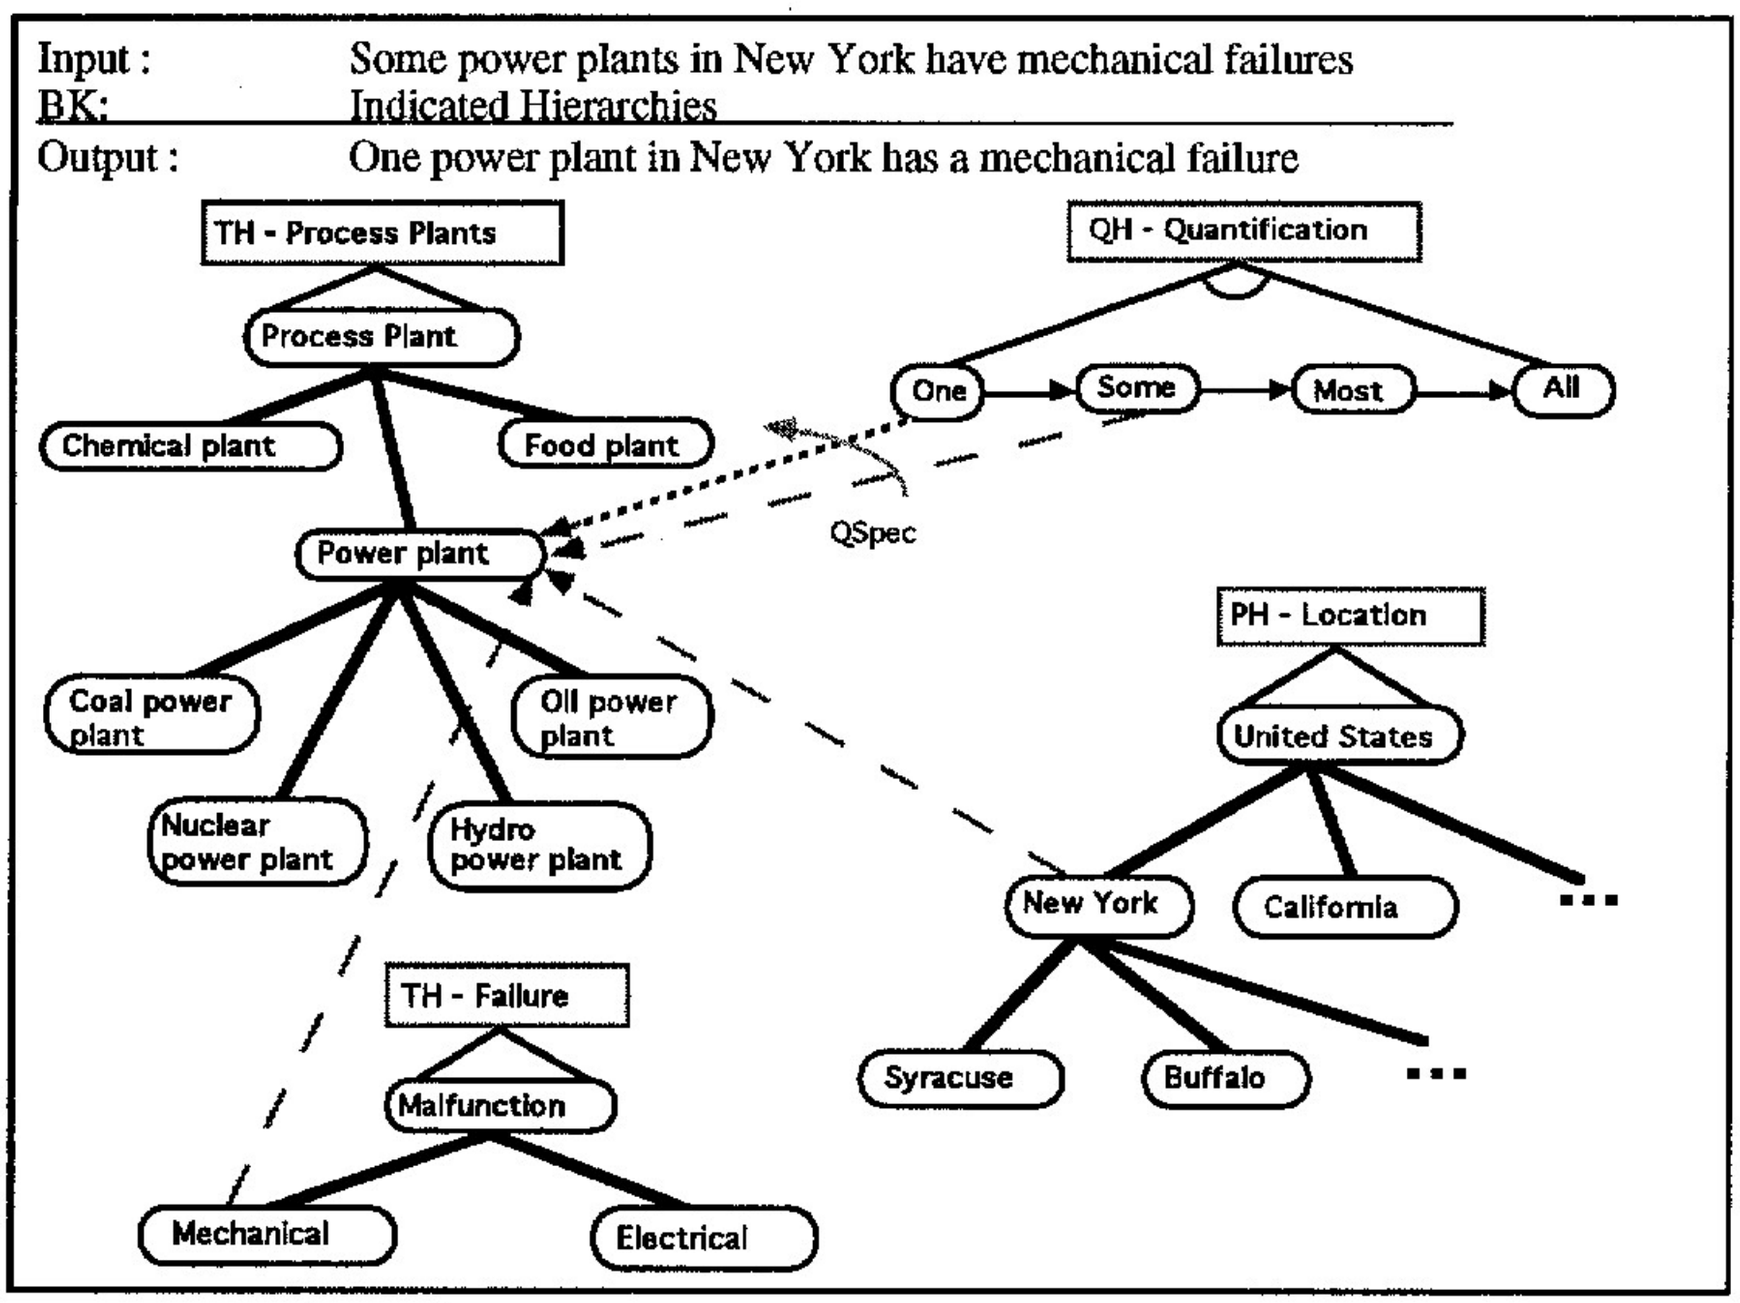

Figure 6: Deductive specialization based on quantification.

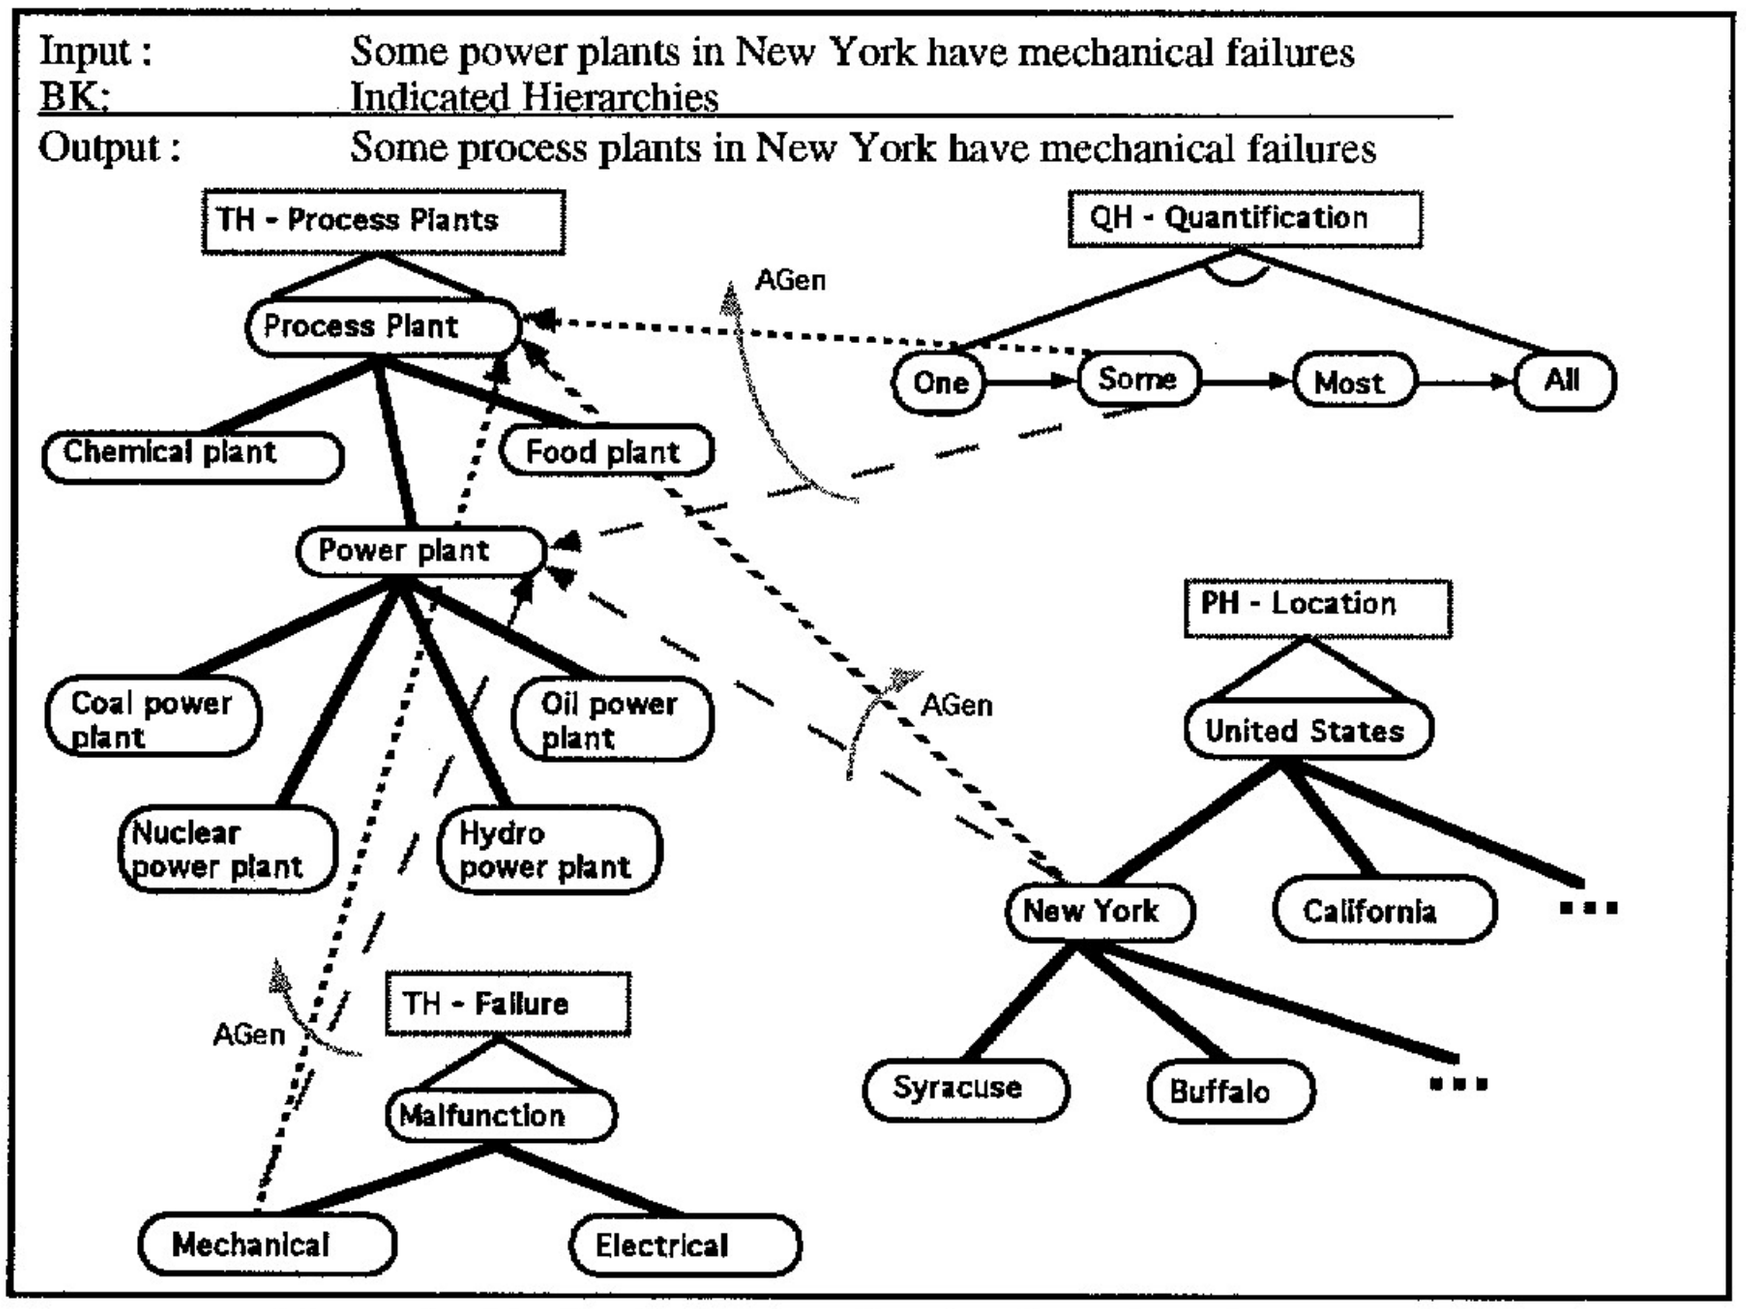

Figure 7: Deductive generalization based on the argument.

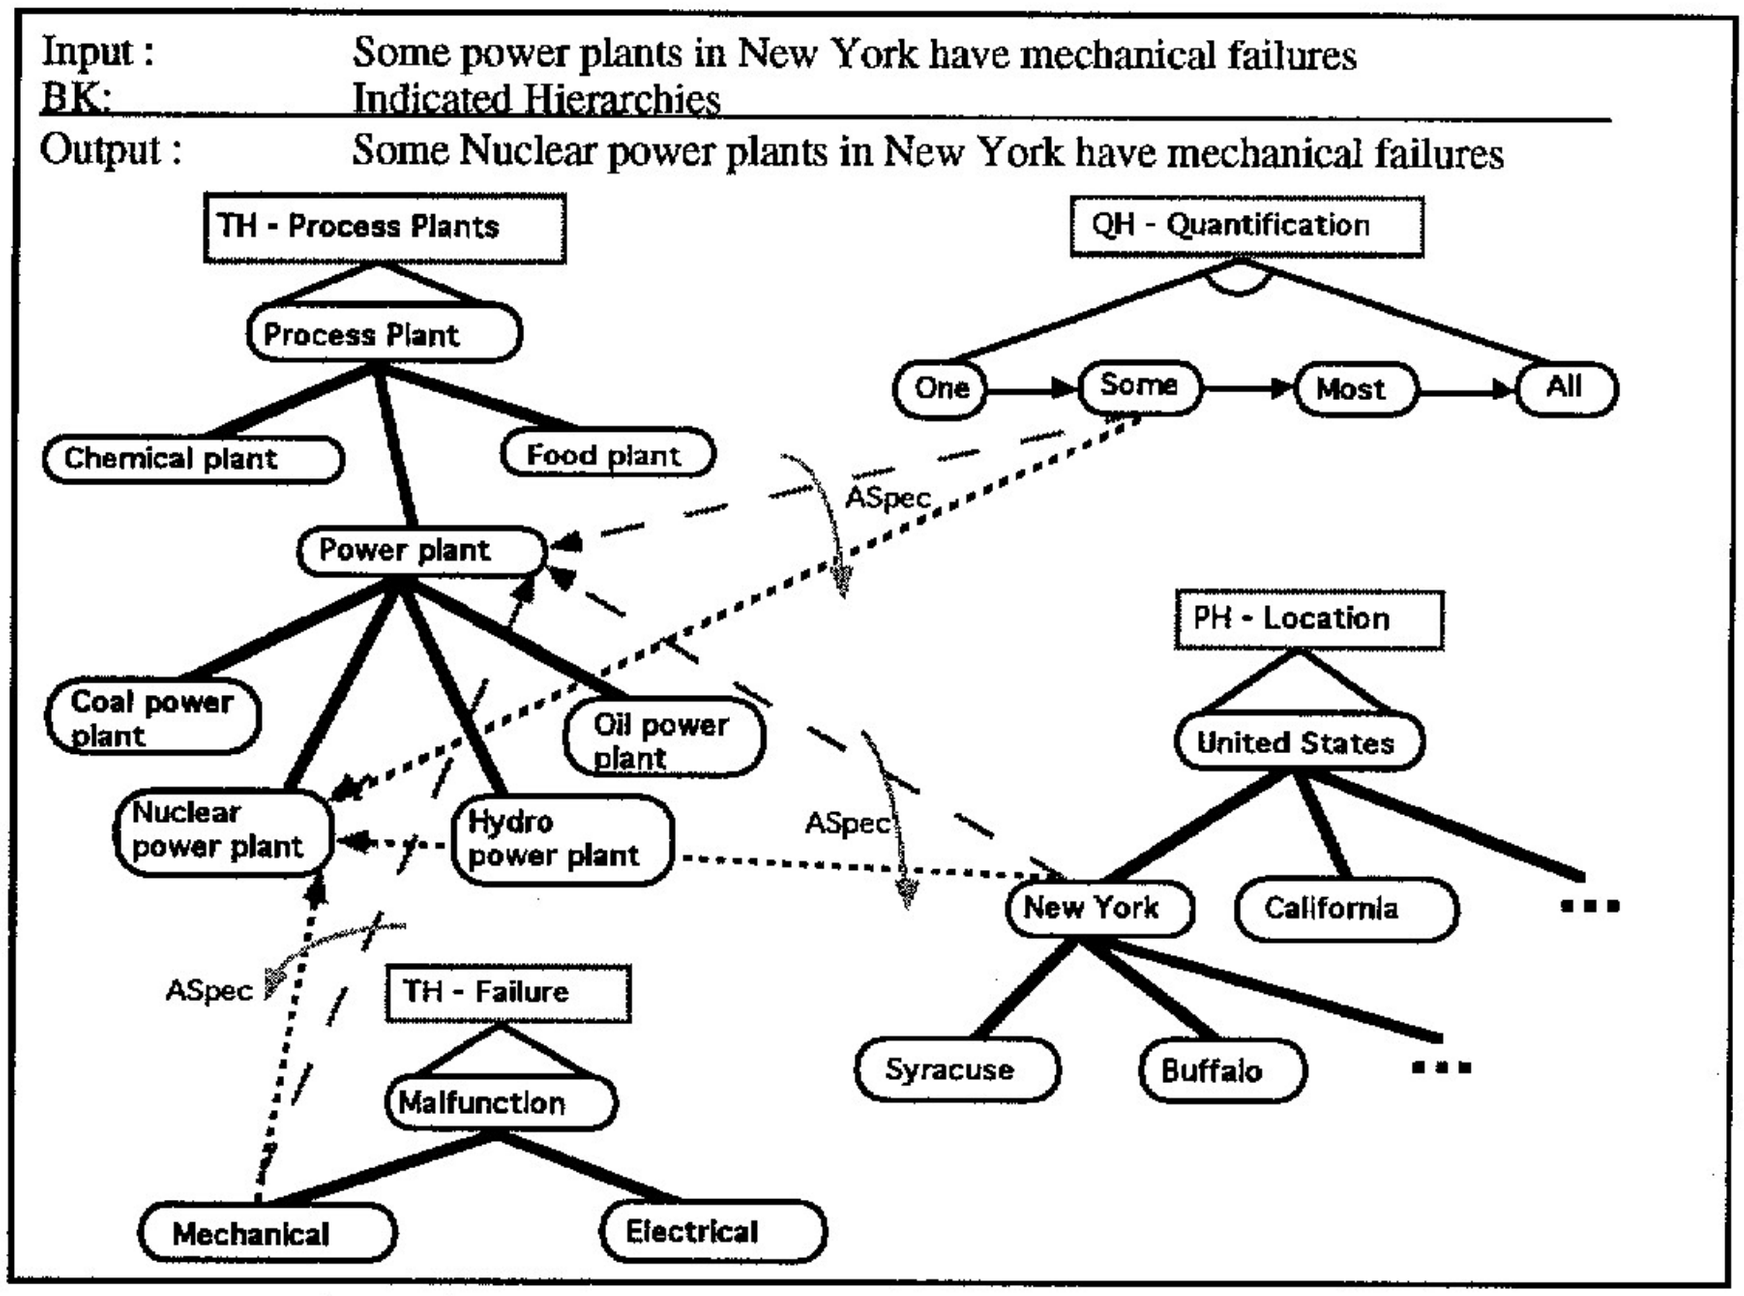

Figure 8: Inductive specialization based on the argument.

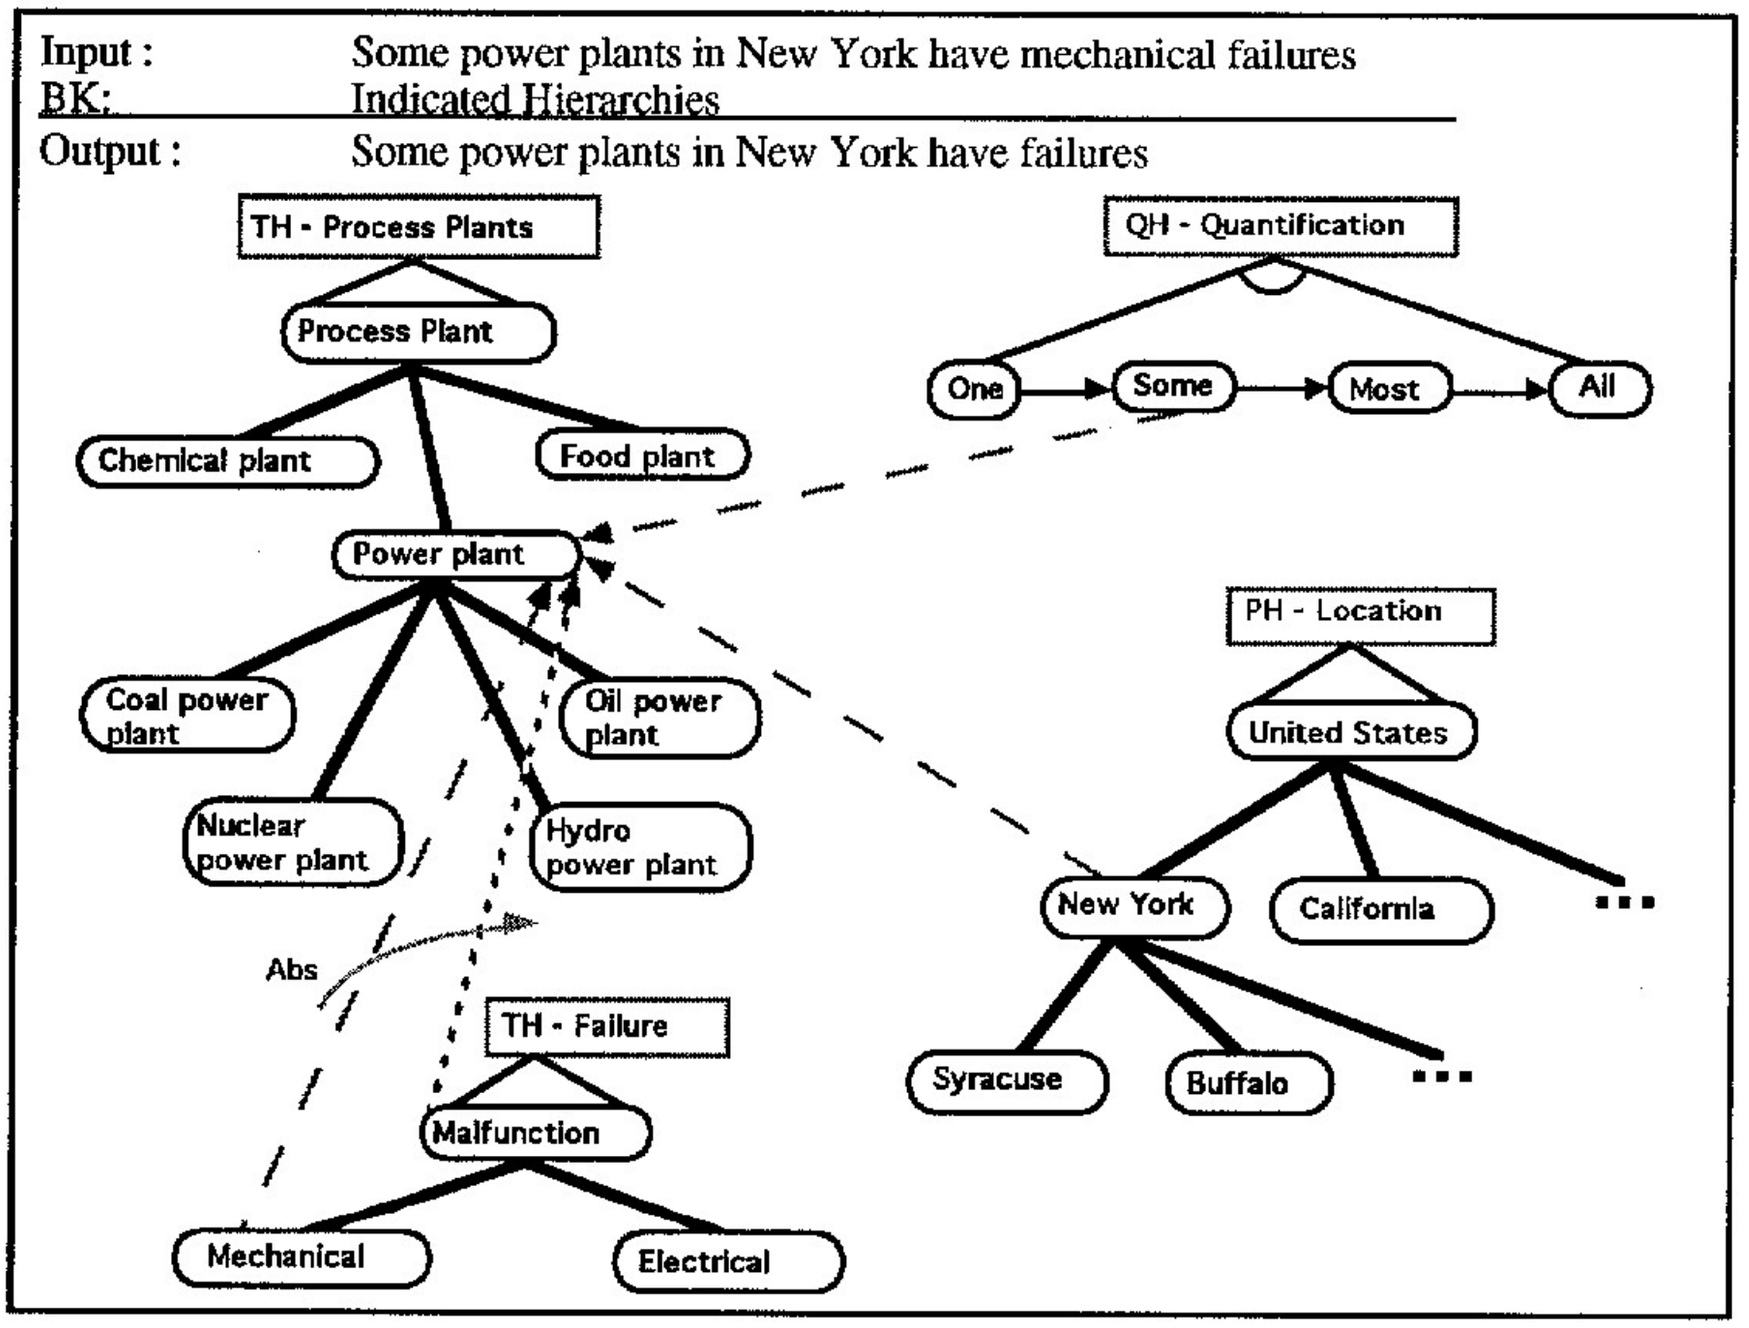

Figure 9: Abstraction transmutation.

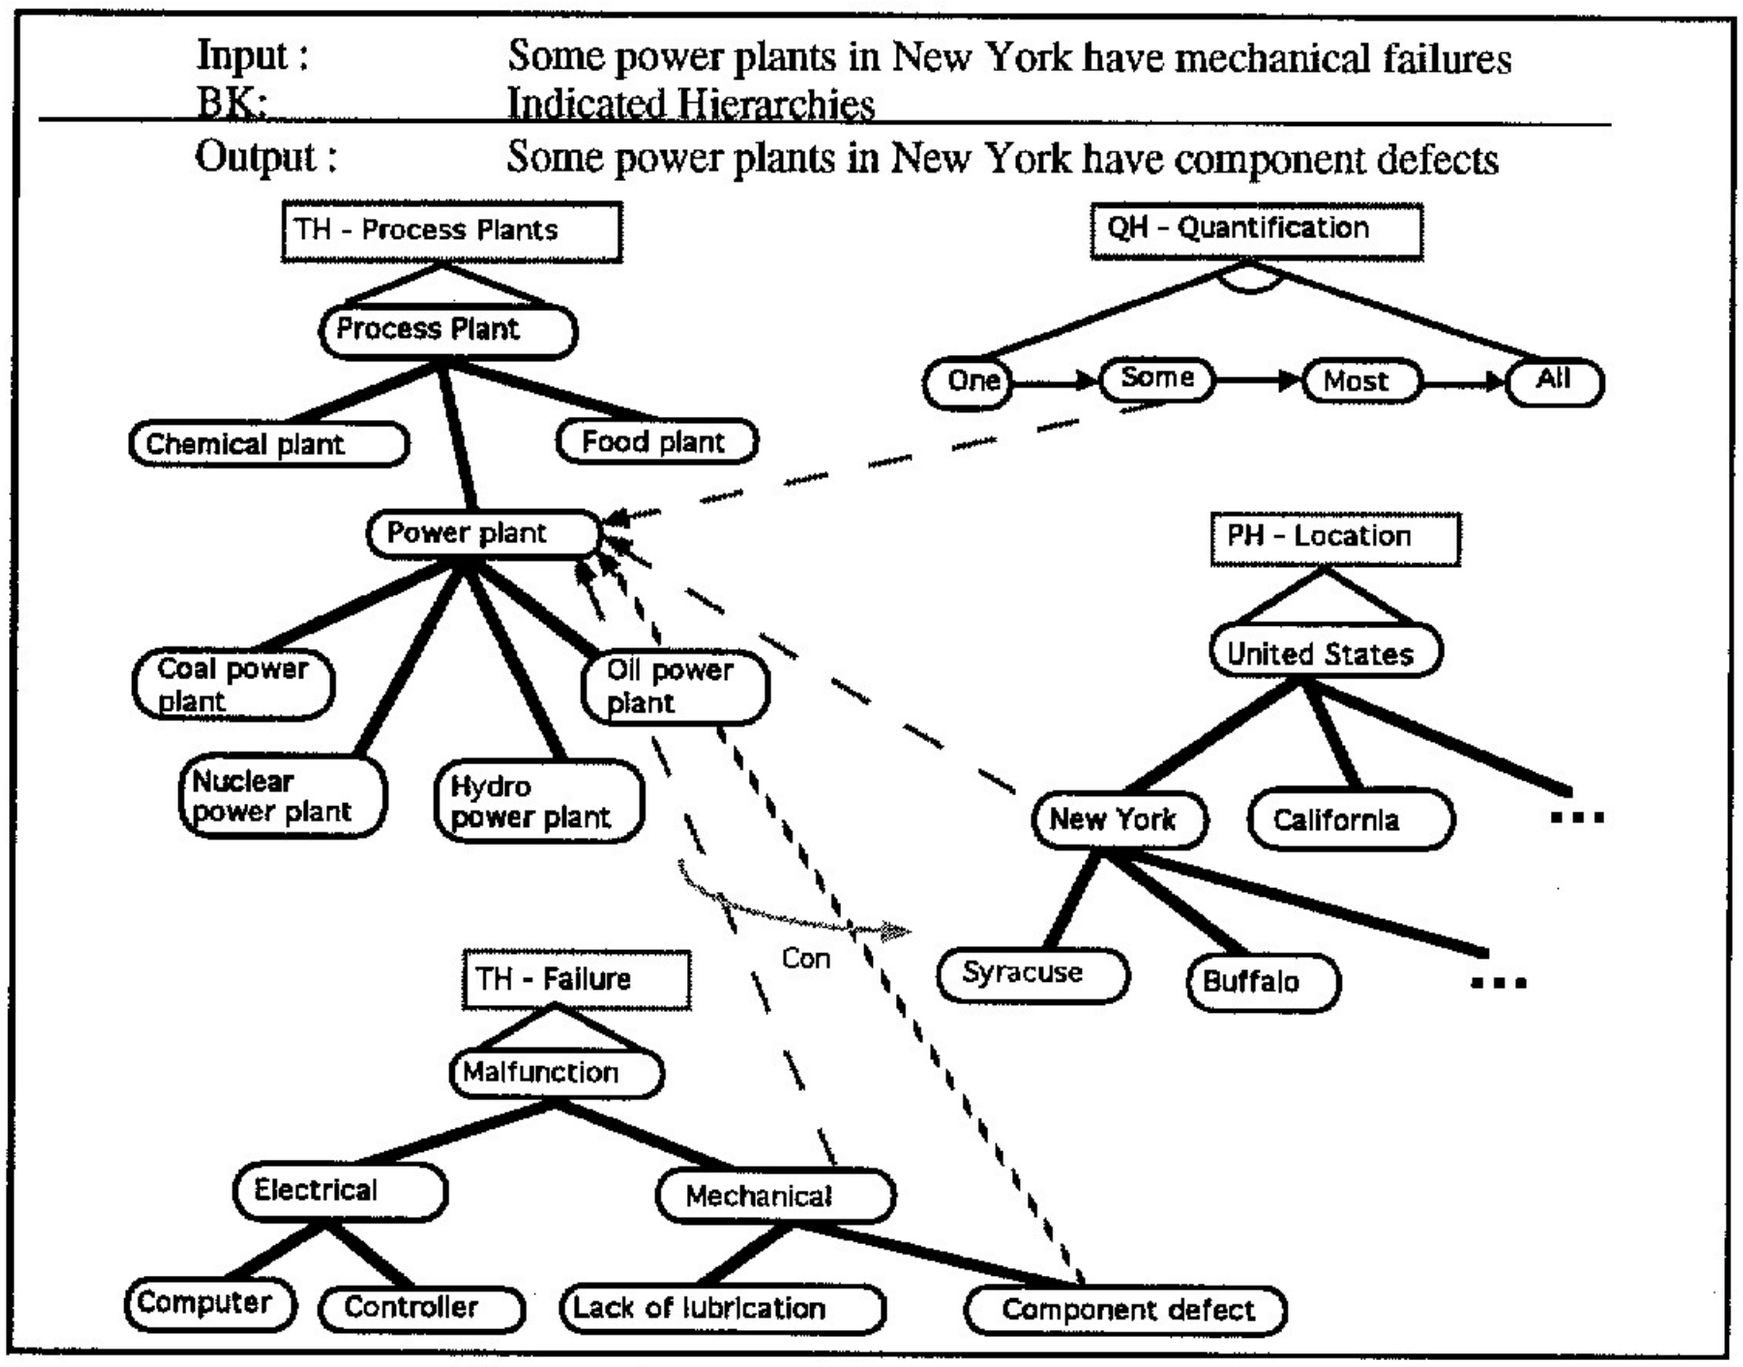

Figure 10: Concretion transmutation.

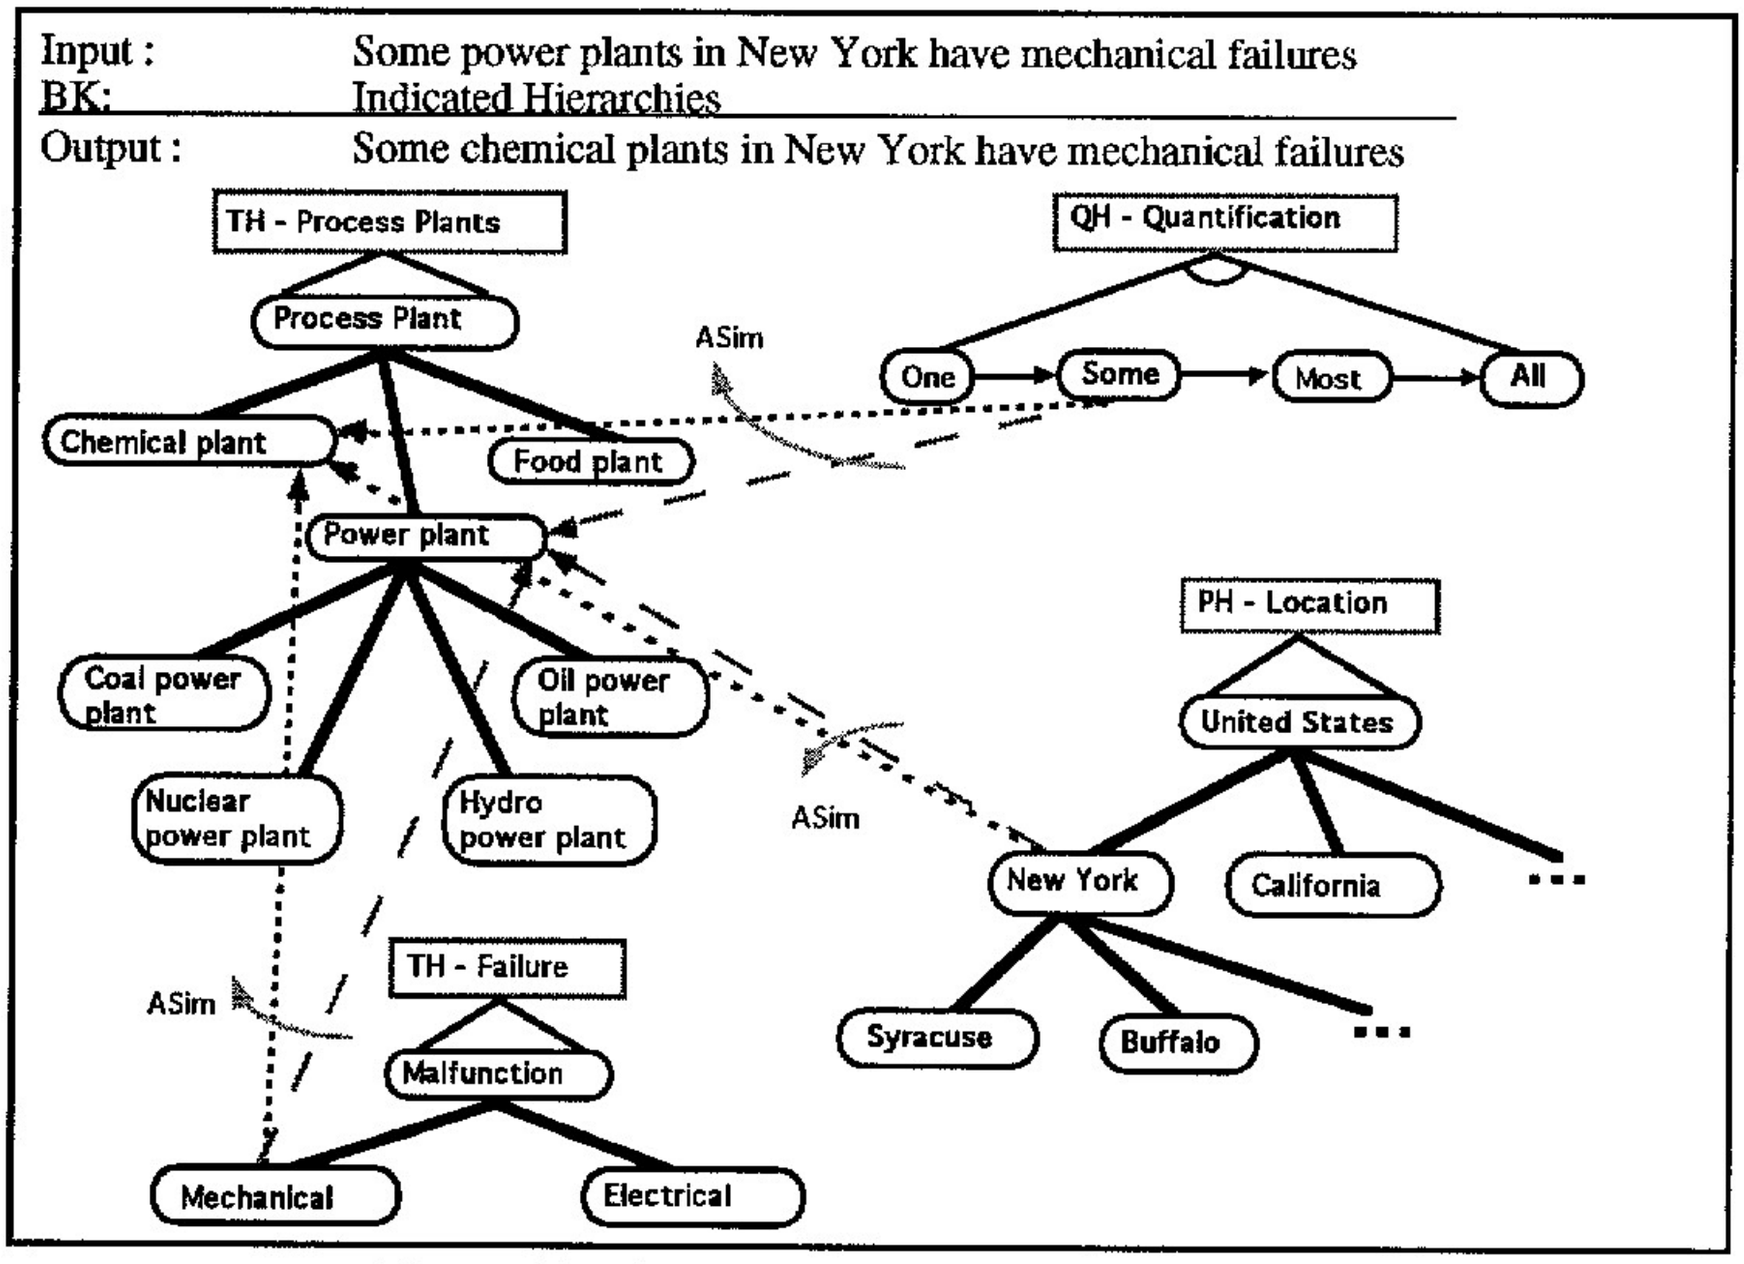

Figure 11: Argument similization transmutation.

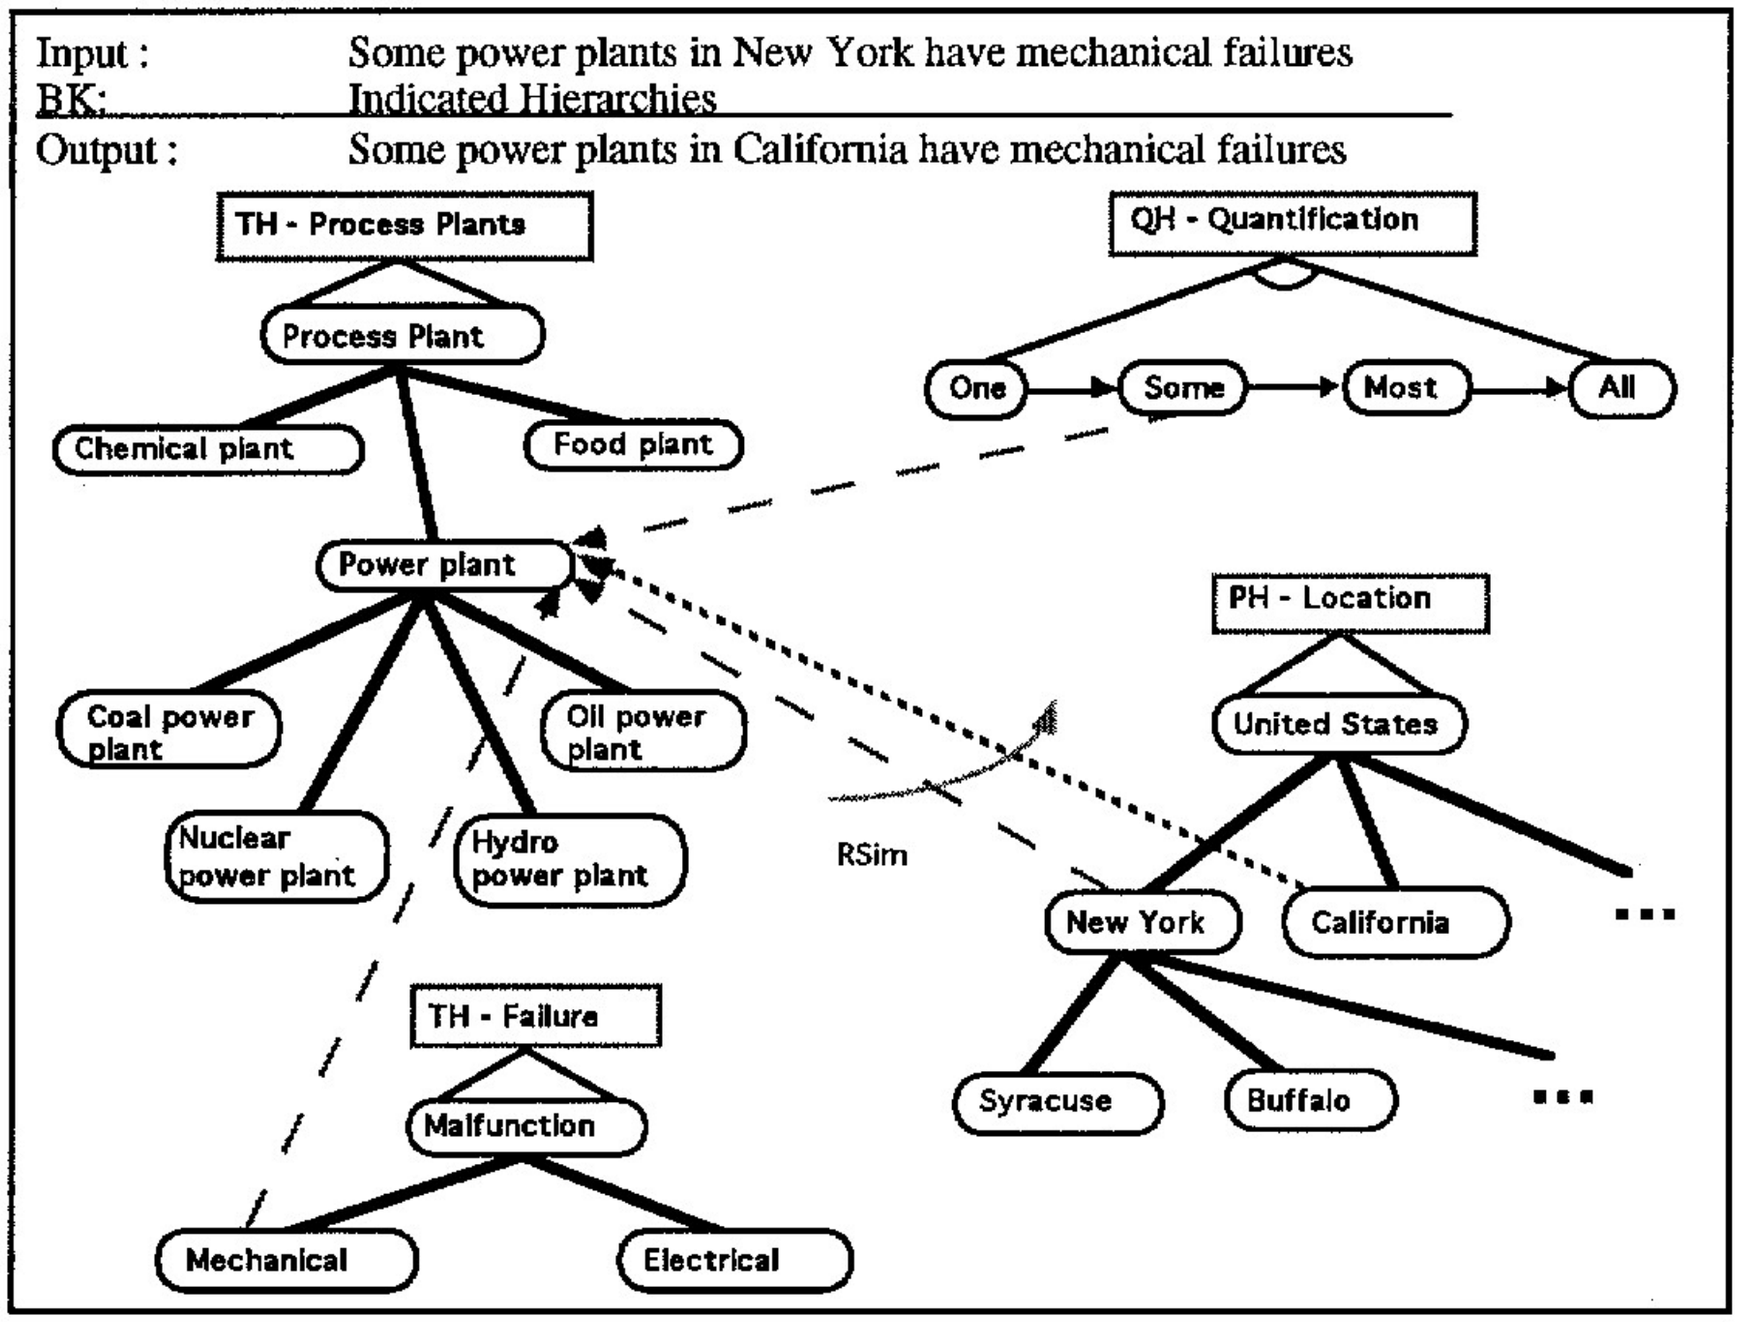

Figure 12: Referent similization transmutation.

Although issues related to the implementation of an MTL system are beyond the scope of this paper, we will briefly outline the basic ideas. We have been pursuing two approaches, MTL-JT, which builds a plausible justification tree to "understand" a user's input (Tecuci, 1993), and a second one, MTL-DIH, based on DIH.

In the MTL-DIH approach, a learning strategy is determined by analyzing the learning task. This analysis relates the input information to the learner's background knowledge and the learning goal. The input information to the system is assumed to be given in the form of logic statements. It can be concept examples, concept descriptions, rule examples, rules or a combination of the above. The system re-represents the input as a trace, or set of traces. Background knowledge is the part of the learner's prior knowledge that is relevant to the input and the learning goal.

The learning goal specifies criteria characterizing knowledge to be learned. There are different kinds of learning goals, such as to predict new information, to explain the input, to classify a fact or concept instance, to create an abstract description from an operational one or conversely, to create a problem solution or a plan. It is assumed that the learning goal is determined by a teacher or by the control module of the system.

The learning process involves determining the type of relationship between the given input and the background knowledge, and performing a sequence of knowledge transmutations, involving input and background knowledge, to produce knowledge satisfying the learning goal.

8. Summary and Future Research

The DIH knowledge representation presented serves as the basis for implementing multistrategy task-adaptive learning. It builds upon ideas of the Inferential Theory of Learning and the core theory of plausible reasoning. Although it is closely related to the semantic network representation, it represents a significantly different approach, and contains many new ideas that make it particularly useful for representing multitype inference. These include the idea of dividing the knowledge representation into a static part and a dynamic part, the organization of knowledge in which basic forms of inference can be performed via simple trace perturbations, and the introduction of various precedence hierarchies, such as the schema hierarchy, the measure hierarchy, and the quantification hierarchy.

The primary purpose of this paper was to demonstrate how DIH supports several basic knowledge generation transmutations, specifically, generalization, specialization, abstraction, concretion, similization and dissimilization. The first version of DIH has been implemented in Smalltalk, and used as a tool for investigating the interactive display and modification of traces in hierarchies. The visual display of inference is particularly useful in situations that involve traces connecting only a few hierarchies (that is, representing short sentences). To facilitate knowledge visualization, the system has an option to present traces with only a limited number of neighboring nodes, rather then connecting complete hierarchies.

In DIH, the more knowledge structures there are in background knowledge, the easier it is to assimilate new knowledge, or to plausibly explain input statements. DIH is an efficient, representation, because most knowledge modifications consist of forming or changing traces, without affecting the established hierarchies.

Many issues remain to be addressed in future research. Among these issues are the representation of more complex forms of knowledge—mutual implications, various types of dependencies, temporal and spatial knowledge, and the development of methods for determining the affect of merit parameters on the reasoning process.

Acknowledgments

The authors thank Eric Bloedorn, Mark Burnstein, David Hille and Ken Kaufman for their thoughtful comments on this paper.

This research was conducted in the Center for Artificial Intelligence at George Mason University. The Center's research is supported in part by the National Science Foundation under grant No. IRI-9020266, in part by the Advanced Research Projects Agency under the grant No. N00014-91-J-1854, administered by the Office of Naval Research and the grant No. F49620-92-J-0549, administered by the Air Force Office of Scientific Research, and in part by the Office of Naval Research under grant No. N00014-91-J-1351.

References

Baker, M., Burstein, M.H., and Collins, A.M., "Implementing a Model of Human Plausible Reasoning," *Proceedings of the Tenth International Joint Conference of Artificial Intelligence*, pp. 185-188, Los Altos, CA: Morgan Kaufman, 1987.

Brachman, R.J., McGuinness, D.L., Patel-Schneider, P.F., Resnick, L.A. and Borgida, A., "Living with CLASSIC: When and How to Use a KL-ONE-Like Language," *Principles of Semantic Networks - Explorations in the Representation of Knowledge*, J.F. Sowa (ed.), Morgan Kaufmann Publishers, 1991.

Beckwith, R., Fellbaum, C., Gross, D., and Miller, G.A., "WordNet: A Lexical Database Organized on Psycholingistic Principles," *Using On-line Resources to Build a Lexicon*, U. Zernick (Ed.), Hillsdale, NJ: Erlbaum, 1991.

Boehm-Davis, D., Dontas, K., and Michalski, R.S., "A Validation and Exploration of the Collins-Michalski Theory of Plausible Reasoning," *Reports of the Machine Learning and Inference Laboratory*, MLI 90-5, Center for Artificial Intelligence, George Mason University, Fairfax, VA, 1990.

Collins, A., and Michalski, R.S., "The Logic of Plausible Reasoning: A Core Theory," *Cognitive Science,* Vol. 13, pp. 1-49, 1989.

Collins, A., and Quillian, M.R., "How To Make A Language User," *Organization of Memory,* E. Tulving and W. Donaldson (eds.), New York: Academic, 1972.

Dontas, K., and Zemakova, M., "APPLAUSE: An Implementation of the Collins-Michalski Theory of Plausible Reasoning," *Proceedings of the Third International Symposium on Methodologies for Intelligent Systems*, Torino, Italy, 1988.

Hieb, M.R. and Michalski, R.S., "A Knowledge Representation System Based on Dynamic Interlaced Hierarchies: Basic Ideas and Examples," *Reports of the Machine Learning and Inference Laboratory*, MLI 93-5, Center for Artificial Intelligence, George Mason University, Fairfax, VA, 1993.

Kelly, J., *PRS: A System for Plausible Reasoning*, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, 1988.

Michalski, R. S., "A Theory and Methodology of Inductive Learning," Chapter in the book, *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. Carbonell and T. Mitchell (Eds.), TIOGA Publishing Co., Palo Alto, 1983, pp. 83-134.

Michalski, R.S., "Towards a Unified Theory of Learning: Multistrategy Task-adaptive Learning," *Reports of the Machine Learning and Inference Laboratory*, MLI 90-1, Center for Artificial Intelligence, George Mason University, Fairfax, VA, 1990.

Michalski, R.S., "Inferential Learning Theory as a Basis for Multistrategy Task-Adaptive Learning," *First International Workshop on Multistrategy Learning*, pp. 3-18, Harpers Ferry, West Virginia, 1991.

Michalski, R.S., "Inferential Theory of Learning: Developing Foundations for Multistrategy Learning," in *Machine Learning: A Multistrategy Approach, Volume 4,* R.S. Michalski & G. Tecuci (Eds.), Morgan Kaufmann Publishers, 1993.

Michalski R. S., Baskin A. B., Uhrik C., Channic T., Borodkin S.,. Boulanger A. G, Rodewald L., and Reinke, "A Technical Description of the ADVISE.1 Meta-Expert System that Integrates Multiple Knowledge Representations and Learning Capabilities," ISG 86-8, Department of Computer Science, University of Illinois, Urbana, 1986.

Morik, K., Causse, K., and Boswell, R., "A Common Knowledge Representation Integrating Learning Tools," *First International Workshop on Multistrategy Learning*, pp. 81-96, Harpers Ferry, West Virginia, 1991.

Tecuci, G. & Michalski, R.S., "A Method for Multistrategy Task-adaptive Learning Based on Plausible Justifications," in Birnbaum, L., & Collins, G. (Eds.) *Machine Learning: Proceedings of the Eighth International Workshop*, San Mateo, CA, Morgan Kaufmann, 1991.

Tecuci, G., "An Inference-Based Framework for Multistrategy Learning" in *Machine Learning: A Multistrategy Approach, Volume 4,* R.S. Michalski & G. Tecuci (Eds.), Morgan Kaufmann Publishers, 1993.

Winston, M.E., Chaffin, R., and Herrmann, D., "A Taxonomy of Part-Whole Relations," *Cognitive Science,* Vol. 11, pp. 417-444, 1987.

Wnek, J. and Michalski, R.S., "An Experimental Comparison of Symbolic and Subsymbolic Learning Paradigms: Phase I - Learning Logic-style Concepts," *First International Workshop on Multistrategy Learning*, pp. 324-339, Harpers Ferry, West Virginia, 1991.