Submitted to the Workshop for Qualitative Physics[*]

An Experimental Design Approach to Ambiguity Resolution[**]

Shankar A. Rajamoney
Gerald F. DeJong

Artificial Intelligence Research Group
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
1101 West Springfield Avenue
Urbana, IL 61801

Telephone: (217) 333-6446
Arpanet: shankar@uicsl.csl.uiuc.edu

Keywords

Experiment Design
Resolving Ambiguity
Dealing with Multiple Interpretations
Theory Formation and Refinement

Abstract

Ambiguity is a limiting obstacle to qualitative reasoning systems. In this paper, we present a new approach called *active ambiguity reduction* to deal with this problem. Active ambiguity reduction involves the purposeful alteration of the world to generate new information. This new information will be inconsistent with some of the ambiguous situations, thereby eliminating them from further consideration. Active ambiguity reduction may also be viewed as *experiment design*. This paper presents a theory of experiment design which is based on the principle of *refutation*. The theory describes three strategies for designing experiments - *elaboration*, *discrimination* and *transformation*. An experiment engine - an implementation of the theory - is also described. Importantly, the experiment engine is general and domain-independent and therefore readily integrable with existing qualitative reasoning systems.

---

[*] An extended version of this paper has been submitted to AAAI-87.

[**] This research was supported in part by a University of Illinois Cognitive Science/Artificial Intelligence Fellowship and in part by the Office of Naval Research under grant N-00014-86-K-0309.

1. Introduction

Qualitative reasoning systems encounter ambiguity [de Kleer79, Kuipers84] due to the qualitative nature of the reasoning and the local nature of the reasoning algorithms. Ambiguity is often a limiting obstacle to qualitative reasoning systems. Even simple world descriptions can lead to an unmanageable number of possibilities. Further, the ambiguity problem becomes dramatically worse with increased complexity of the world. Humans, not having the infallible bookkeeping abilities of computers, are overwhelmed by far fewer numbers of ambiguities. One human method to combat growing ambiguities is to actively interact with their environment. The interactions yield new data from the world which may eliminate ambiguities. Even when there are relatively few possible qualitative states, active pruning may be desirable. Humans often perform cheap tests to head off undesirable possibilities. They heft a snowball several times before throwing it or test the swimming pool temperature with a toe before diving in.

We propose a method called *active ambiguity reduction* as a partial solution to the problem of controlling ambiguities in qualitative reasoning systems. Active ambiguity reduction involves purposeful alteration of the world in such a way that observable behavior will dictate which of a number of ambiguities corresponds to reality. This process can be looked upon as conducting experiments in the world. The outcome of these experiments provide additional data which, provided the experiments are suitably chosen, will be inconsistent with a number of ambiguous qualitative scenarios. These ambiguities can then be dropped from further consideration.

Central to active ambiguity reduction is the issue of experimental design. The experimental design process must have several important features. First, it must be complete; if there is a way to tease apart an ambiguity the design system should find it. Second, it must be tolerant of unavailable data. Third, it should be efficient. Each experiment should evenly divide the ambiguities so that significant information is acquired regardless of the experiment's outcome. Fourth, it should be practical. Lighting a match is not a reasonable way to tell whether a nearby barrel contains water or gasoline.

This paper presents a theory of experiment design that is applicable to a wide variety of qualitative reasoning tasks - determining activity, prediction and envisionment, measurement interpretation, design, and fault diagnosis and troubleshooting. This theory includes three strategies for designing experiments - *elaboration*, *discrimination* and *transformation*. The theory and an experiment engine - an implementation of the theory - are described and their application to qualitative reasoning tasks like determining activity, measurement interpretation and prediction is described. The experiment engine is general and domain-independent. It is applicable to a number of existing qualitative reasoning systems. The requirements of the experiment engine are: 1) an inference engine that gives the predictions supported by a hypothesis for a given scenario and 2) domain-dependent information about the quantities that can be observed, measured and manipulated.

2. A Model of Experiment Design

The ideal experiment design system accepts a set of hypotheses, designs a series of experiments and returns the "correct" hypothesis i.e. a hypothesis that is consistent with reality. An external problem solver generates the hypotheses in the course of its problem solving and uses the experiment design system to find the correct hypothesis. For qualitative reasoning, the problem solver is the qualitative reasoning system and each hypothesis is a collection of assumptions that leads to ambiguity. The experiment design system interacts with an inference system which generates a (possibly incomplete) set of predictions that are supported by a hypothesis for a given scenario. The experiment system manipulates the hypotheses and predictions to find the correct hypothesis.

The goal of experimentation is to increase the system's knowledge about a situation. If everything that can be known about the situation is already known then experimentation will not help in finding the correct hypothesis. In such a case the situation is inherently ambiguous. However, in most situations, what is known is considerably less than what is predicted by the hypotheses. Therefore, designing experiments to find out what is not known about the situation can help in finding the correct hypothesis.

The underlying principle behind experimentation is the *refutation* of hypotheses. Experiments are designed with the specific purpose of finding information that is not compatible with the predictions of a hypothesis. When all hypotheses but one have been eliminated then the remaining hypothesis is the correct one. This principle of refutation is superior to the principle of *confirmation* of a hypothesis which involves verifying that every prediction supported by the hypothesis is correct. This is because some hypotheses may support an infinite set of predictions and therefore can never be completely confirmed. The strategies used to design experiments to refute hypotheses are described in detail in the next section.

A practical experiment design system has limitations in its ability to design all the appropriate experiments required to identify the correct hypothesis. Even after the complete space of experiments that can be designed by the system has been exhausted, there can be more than one hypothesis remaining that are consistent with the available information. Any or all of these hypotheses may be used for further problem solving since they are equivalent. However, for a given domain it may be possible to find criteria called *hypothesis selection criteria* for selecting a "best" hypothesis from the set of hypotheses left after experimentation. For example, such criteria may be based on the principle of Occam's razor that prefers "simple" hypotheses.

2.1. Strategies for Experiment Design

Three strategies for designing experiments to refute hypotheses - elaboration, discrimination and transformation - are described:

[a] Elaboration: In a given domain, there are usually some quantities that are readily observable or easily measurable. *Elaboration* involves observing or measuring the values of such quantities for a given scenario and refuting hypotheses that predict values for the quantity that are not compatible with the measured value. This strategy does not guarantee that the designed experiment will refute a hypothesis since all the hypotheses may support the observed value or may be agnostic about the value of the quantity. This strategy therefore involves a tradeoff between the ease of designing an experiment and the effectiveness of the experiment in refuting hypotheses.

[b] Discrimination: In elaboration the quantity to be measured is selected according to the ease with which it can be measured. This condition can be waived and elaboration can be used as the sole strategy for designing experiments. In this case, elaboration experiments result in a gradual increase of what is known about a scenario until every quantity that can be measured has been measured. However, this is very inefficient and will result in a large number of experiments that do not refute any hypothesis. In *discrimination* a quantity is selected only if its measurement will help in the refutation of hypotheses. Discrimination involves the measurement of a quantity which satisfies two criteria: 1) a number of different hypotheses should predict different values for the quantity 2) these values should be *discriminable* i.e. a measurement should be able to distinguish the different values. A quantity that satisfies these two criteria is called a *discriminant*. Once a discriminant has been selected an experiment to measure its value will result in the refutation of hypotheses that supported the other values. A discrimination experiment is guaranteed to refute hypotheses if the values of the discriminant can be grouped into sets of discriminable values and each set has at least one hypothesis that predicts a value from that set. Discrimination is more effective than elaboration since the experiments are directed towards the measurement of those quantities that will help in the refutation of hypotheses.

[c] Transformation: Elaboration and discrimination select quantities from a specified scenario for measurement. However, even after the space of measurable quantities for the scenario has been exhausted it may not be possible to identify the correct hypothesis. The ability to refute hypotheses is greatly enhanced if the experiment design system is endowed with the capacity to construct new scenarios or modify existing ones. This requires the use of domain-specific operators called *scenario transformation operators* that construct a scenario from a given scenario or propose changes to the given scenario. The ability to generate new scenarios helps in the refutation of hypotheses in two ways:

(1) The techniques of elaboration and discrimination can be applied to the new scenario. This results in the refutation of hypotheses that predicted the same value or indiscriminable values for a quantity in the original scenario and now predict different or discriminable values for the same quantity in the new scenario. This divergence of indiscriminable values can be used to select the appropriate transformations to be applied to the original scenario.

(2) A new technique called *differential discrimination* can be applied to the new scenario and the original scenario. Differential discrimination involves the measurement of a quantity that satisfies the following two criteria: 1) there are a number of hypotheses that predict the same value or indiscriminable values for the quantity in the original and transformed scenarios 2) however, the manner in which the value or indiscriminable values was reached under each hypothesis is different and discriminable. For example, the predictions in one scenario may be reached much faster or more of the predicted behavior may occur in the same time span as compared to the other scenario. Thus differential discrimination involves the discrimination of the second order behavior of a quantity across two scenarios. This strategy is efficient if transformations that lead to the discrimination of second-order behavior can be found. Then a single measurement in each scenario will refute a number of hypotheses that do not support the observed comparative behavior for the discriminant.

The problem of selecting the correct set of transformations to apply to a scenario can be viewed as a typical planning task. The initial state is the given scenario and the final state is a new scenario. The goal criterion to be satisfied by the new scenario is that it should have discriminable first-order or second-order behavior that leads to the refutation of one or more hypotheses. The plan is a sequence of transformations that converts the original scenario to a new scenario that satisfies the goal criterion.

These three strategies are best illustrated by an example from chemistry. In chemistry, we are often asked to determine an unknown salt. The process of determining the unknown salt is called *qualitative analysis* and is done by performing a series of experiments on samples of the unknown salt. Elaboration experiments correspond to the determination of the color, taste, smell, crystalline structure, physical state, litmus test reaction etc of the given salt. These experiments are very simple to perform and are invariably carried out to get rough idea of the likely candidates. Very rarely do these experiments conclusively determine the salt. Discrimination experiments are more difficult to design and perform. For example, we might need to determine whether a precipitate has formed using centrifugal precipitation or determine the solubility of the salt by measuring the concentration of the solution using titration. Transformation corresponds to converting the salt to another compound by chemical reactions. An example of transformation and elaboration is reacting the unknown salt with concentrated HCL and checking for a precipitate and identifying the chloride formed by examining the color of the precipitate. An example of transformation and discrimination is reacting the unknown salt with NH4Cl, NH4OH and (NH4)2S and performing the borax bead test if there is a precipitate to discriminate between CoS and NiS. Transformation and differential discrimination can involve comparing the speed with which the salt reacts with a given reagent. For example, the unknown salt is reacted with NH4Cl, NH4OH, (NH4)2CO3 and CaSO4 and if there is an immediate precipitate then Ba is present, if there is a tardy precipitate Sr is present and if there is no precipitate then Ca is present.

2.2. A Domain-Independent Experiment Engine

An experiment engine has been developed based on the model of experimentation described above. The inputs to the experiment engine are:

(1) A set of hypotheses.

(2) An inference engine that accepts a hypothesis and a scenario and returns a set of predictions supported by the hypothesis for the given scenario.

(3) Domain-dependent knowledge: There are three major sources of domain knowledge:

[a] A set of predicates that describe the quantities of the domain that can be measured or observed, the values that are discriminable, the parameters of a scenario that can be transformed etc. These predicates are required to design well-defined experiments that can be readily performed in the domain.

[b] A set of scenario transformation operators. These operators endow the experiment designer with the ability to construct new scenarios. A scenario transformation operator can change the quantities of some of the components of the scenario, the components of the scenario or the manner in which the components are organized. The current implementation supports only changes in the quantities of components.

[c] A set of hypothesis selection criteria. These criteria include domain specific knowledge and general principles like Occam's razor. These criteria are used to select the "best" hypothesis or hypotheses from the hypotheses remaining after experimentation.

The experiment engine uses elaboration, discrimination and transformation to design experiments. It refutes those hypotheses that support predictions that are not compatible with the results of the experiments. It finally returns the "best" hypothesis or hypotheses that are consistent with the available information.

Importantly, the experiment engine makes very few assumptions. Most qualitative reasoning systems [de Kleer84, Forbus84a, Kuipers84, Williams84] readily provide the type of inference engine that is required by the experiment design system. The domain-dependent knowledge required for designing the experiments is usually readily available. Since the experiment engine itself does not inspect the content of the hypotheses or predictions it is not limited by the domain-specific nature of the hypotheses or the predictions. Thus, the experiment engine is general and domain-independent and can be easily integrated with existing qualitative systems.

3. Application of Experiment Design to Qualitative Reasoning Tasks

Some of the qualitative reasoning tasks to which the experiment engine can be applied are:

Determining Activity: The primary reasoning task is to determine the physical changes in a given situation. If there are many processes imposing conflicting changes on a quantity then the qualitative nature of the reasoning will lead to ambiguity about the resulting change. For example, evaporation, condensation and flow of liquid out of a container can each change the level of the liquid in the container. Qualitative reasoning will not be able to establish whether the level is increasing, decreasing or constant unless it has further information to determine which process/es dominate/s. Experiments can be designed to measure such ambiguous quantities. In addition to resolving the ambiguity, the information from the experiments about the inequalities can help in further reasoning.

Prediction: Prediction involves determining how the scenario evolves with time. If there is ambiguity in determining activity or if the initial scenario was incompletely specified then prediction will result in an *envisionment* [de Kleer77, Forbus84b] - a set of possible paths into which the scenario can develop. Envisionment can lead to a potentially very large number of possible situations. Experiment design can help in curbing the combinatorial explosion by resolving ambiguities and generating information to eliminate possible situations. For example, if a mixture of alcohol and water is heated and if the system does not know whether the boiling point of alcohol is greater than, less than or equal to that of water then it will generate situations in which alcohol boils first, water boils first and both boil together. An experiment that determines the boiling points of alcohol and water or one that examines the liquid remaining in the container or the vapor produced by boiling will help in determining which situation actually occurs. If this were a small part of a bigger computation, for example, understanding the functioning of a complex distillation factory, the information obtained by performing such experiments will considerably reduce the computational complexity.

Measurement Interpretation: Measurement interpretation [Forbus83, Forbus84b, Forbus86] involves determining the changes taking place from a set of observations and an incomplete description of an initial scenario. For example, if the system is given a process vocabulary that includes evaporation, absorption and flow of liquids and is shown a scenario in which the level of the liquid in a container is decreasing then measurement interpretation determines which processes are acting even if information about the preconditions is missing. Typically measurement interpretation results in a set of alternate interpretations. In the previous example, liquid might be evaporating, being absorbed, be flowing out of the container or a combination of these. Experiment design can determine the correct interpretation.

4. Discussion

We have outlined a method to cope with certain spurious ambiguities which arise in the normal process of qualitative reasoning. The approach, called *Active Ambiguity Reduction*, would be invoked by a qualitative reasoning system confronted either with a) an unmanageably large number of ambiguous situations as might occur in constructing a full qualitative envisionment [de Kleer79, Forbus84b], or b) some ambiguity entails an intolerable feature so that it is important for the system to discover whether this qualitative state actually corresponds to reality. When invoked, the system identifies measurements which would serve to disambiguate among the postulated alternatives. The system then proposes experiments by which these measurements can be obtained. A major portion of the research contribution addresses the problem of *experimental design*.

The current research is complementary to research in the area of fault diagnosis [Davis82, de Kleer86a, Genesereth84]. For example, the diagnostic approach described in [de Kleer86a] produces a set of candidates each of which explains the observed differences in behavior between the model and the artifact. An active ambiguity reduction module might use this set of candidates as input hypotheses. Experiments would then be designed to identify the actual fault/s in the model. An ATMS [de Kleer86b] could also be used to organize the hypotheses that are input to the AAR experiment engine. The assumptions or collections of assumptions underlying a contradiction would then form the hypotheses to be tested.

Active ambiguity reduction might also be used to complement Forbus' notion of measurement interpretation [Forbus86]. Forbus describes how constraints from envisionment can restrict the total number of interpretations for a given set of measurements. In general, this approach results in a number of interpretations each of which is consistent with the available measurements. The set of interpretations might be given to the AAR module as ambiguities resulting in the design of experiments to generate the additional measurements required to further reduce the set of interpretations.

In [de Kleer79] de Kleer describes two approaches to resolving ambiguities - 1) using quantitative information and 2) using teleological arguments. The AAR approach operates within the framework of qualitative reasoning. However, teleological information can also be used if it is input as either one of the predictions supported by the hypotheses or as a hypothesis selection criterion.

5. Theory Formation and Refinement

Qualitative reasoning, and QP theory in particular, rely on an accurate description of all the processes of a domain. If the theory is flawed - for example, a process is missing, a precondition is incorrect, a limit point is unknown (a new landmark value [Kuipers85] is needed), or an influence is missing then discrepancies may arise between predictions of the model and observations of reality. Continuing research includes the construction of a model of theory refinement [Mitchell86, Rajamoney87] based on Forbus' QP theory [Forbus84b]. The resulting system has a similar motivation to STAHLP [Rose86] but is experimentally oriented and performs in the framework of qualiative reasoning.

In general there will be many different changes to the theory that remove the anomaly. Each change may result in a distinct qualitative theory. The resulting theories form a set of ambiguous hypotheses. We are extending the ADEPT system [Rajamoney85, Rajamoney86] to integrate our notions of active ambiguity reduction to experimentally determine which of the various hypotheses correspond to reality.

The ADEPT system deals with the incomplete and inconsistent theory problems. The approach used by ADEPT involves:

(1) Monitoring the execution of the system's plans.

(2) Detection of contradictions if the system's predictions are not compatible with the observations.

(3) Hypothesizing reasons which could resolve the contradiction.

(4) Designing experiments to test each of the hypotheses.

(5) Incorporating the information obtained by the experiments into the domain theory.

ADEPT uses five classes of experiments to discriminate among hypotheses, perform measurements, find dependencies among parameters, classify objects based on their behavior with respect to a property and define new properties of objects based on their behavior in a situation. These experiments are used to gather relevant information to identify the correct hypothesis.

The extensions to ADEPT that we are presently exploring involve using Forbus' ATMI (Across Time Measurement Interpretation) module to generate interpretations of observations made by the system and using the AAR module to prune the interpretations. If the AAR module refutes all the interpretations or if no consistent interpretation is available from the ATMI module then the theory refinement module hypothesizes changes to the current theory that can interpret the observations. The ramifications of these hypothesized theories for different scenarios are obtained from Forbus' QPE (Qualitative Process Engine) and are used by the AAR and the experiment design module to find the theory that best accounts for all the generated information.

References

[Davis82] R. Davis, H. Shrobe, W. Hamscher, K. Wieckert, M. Shirley and S. Polit, "Diagnosis Based on Description of Structure and Function," *Proceedings of the National Conference on Artificial Intelligence*, Pittsburgh, PA, August, 1982, pp. 137-142.

[de Kleer77] J. de Kleer, "Multiple Representations of Knowledge in a Mechanics Problem-Solver," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, MA, August 1977, pp. 299-304.

[de Kleer79] J. de Kleer, "The Origin and Resolution of Ambiguities in Causal Arguments," *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, Tokyo, Japan, 1979.

[de Kleer84] J. de Kleer and J. S. Brown, "A Qualitative Physics Based on Confluences," *Artificial Intelligence 24*, (1984), .

[de Kleer86a] J. de Kleer and B. C. Williams, "Reasoning About Multiple Faults," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 1986.

[de Kleer86b] J. de Kleer, "An Assumption-Based Truth Maintenance System," *Artificial Intelligence 28*, (1986), pp. 127-162.

[Forbus83] K. Forbus, "Measurement Interpretation in Qualitative Proces Theory," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, Germany, 1983.

[Forbus84a] K. Forbus, "Qualitative Process Theory," *Artificial Intelligence 24*, (1984), .

[Forbus84b] K. D. Forbus, "Qualitative Process Theory," Technical Report 789, Ph.D. Thesis, MIT AI Lab, Cambridge, MA, August 1984.

[Forbus86] K. Forbus, *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 1986.

[Genesereth84] M. R. Genesereth, "The Use of Design Descriptions in Automated Diagnosis," *Artificial Intelligence 24*, (1984), .

[Kuipers84] B. Kuipers, "Commonsense Reasoning About Causality: Deriving Behavior from Structure," *Artificial Intelligence 24*, (1984), .

[Kuipers85] B. Kuipers, "The Limits of Qualitative Simulation," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, 1985.

[Mitchell86] T. M. Mitchell, R. Keller and S. Kedar-Cabelli, "Explanation-Based Generalization: A Unifying View," *Machine Learning 1*, 1 (January 1986), pp. 47-80.

[Rajamoney85] S. Rajamoney, G. F. DeJong and B. Faltings, "Towards a Model of Conceptual Knowledge Acquisition through Directed Experimentation," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985. (Also appears as Working Paper 68, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.)

[Rajamoney86] S. A. Rajamoney, "Automated Design of Experiments for Refining Theories," M. S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, May 1986. (Also appears as Technical Report T-2213, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.)

[Rajamoney87] S. A. Rajamoney and G. F. DeJong, "The Classification, Detection and Handling of Imperfect Theory Problems," Submitted to the *Tenth International Conference on Artificial Intelligence*, Milan, Italy, Aug 1987.

[Rose86] D. Rose and P. Langley, "STAHLp: Belief Revision in Scientific Discovery," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 1986.

[Williams84] B. Williams. "Qualitative Analysis of MOS Circuits." *Artificial Intelligence 24*. (1984). .