

Towards an Integrated Discovery System

Bernd Nordhausen

Pat Langley

Irvine Computational Intelligence Project

Department of Information & Computer Science

University of California, Irvine, CA 92717 USA

Abstract

Previous research on machine discovery has focused on limited parts of the empirical discovery task. In this paper we describe IDS, an integrated system that addresses both qualitative and quantitative discovery. The program represents its knowledge in terms of *qualitative schema**, which it discovers by interacting with a simulated physical environment. Once IDS has formulated a qualitative schema, it uses that schema to design experiments and to constrain the search for quantitative laws. We have carried out preliminary tests in the domain of heat phenomena. In this context the system has discovered both intrinsic properties, such as the melting point of substances, and numeric laws, such as the conservation of mass for objects going through a phase change.

I Introduction

In recent years, AI researchers have developed a number of systems that operate in the domain of scientific discovery. For instance, BACON [4] discovers numerical laws (e.g., the ideal gas law) and postulates intrinsic properties of object classes (e.g., atomic weight). ABACUS [2] is similar to BACON, but employs an improved search mechanism to find numeric laws in a more efficient manner. It also improves upon BACON by identifying qualitative preconditions on quantitative laws. GLAUBER [6] addresses a different aspect of empirical discovery - the formation of qualitative laws and object taxonomies.

Although each of these systems is successful at its task, each addresses only part of the overall problem of empirical discovery [5]. We are developing an integrated discovery system (IDS) that deals with a variety of empirical discovery tasks, including the formation of qualitative and numeric laws. Historically, qualitative discoveries have tended to lay the foundation for quantitative discoveries, but the latter can in turn lead to higher level qualitative discoveries. Our system operates in the same basic manner, first finding qualitative laws and then using them to aid in discovering quantitative relations.

IDS operates in a simulated world of simple physics and chemistry, thus overcoming one deficiency of previous dis-

covery systems. Previous systems were provided with data* and could not perform their own experiments. In contrast, IDS interacts with the simulated world through a set of *effectors* and *sensors*. Using an effector, the system can actively alter certain attributes of an object, e.g., by changing its location or heating it. Sensors let the program inspect certain attributes, such as the temperature and mass of an object. To carry out an experiment, the system applies effectors to a set of objects and uses its sensors to observe the manner in which those objects change over time.

In the following section, we introduce the representation that IDS employs to state qualitative laws. After this, we examine the mechanisms by which the system discovers qualitative laws and then consider how it uses the resulting schemas to aid its discovery of numeric laws. We close with some proposals for extending the system.

II Representing Qualitative Schemas

Before one can discover qualitative knowledge about the world, one must first have some way to represent that knowledge. Let us consider an example from the domain of heat phenomena. We might begin with a simple view of what happens when we heat an object, e.g., we expect the temperature of the object to increase. If we actually heat a solid, we will see that this occurs, but after some time we may also observe the appearance of a new liquid object. At this point the temperature increase stops and the mass of the liquid increases while the mass of the solid decreases. When the solid has disappeared, the temperature of the liquid begins to increase. This process continues until a new gaseous object appears. As before the mass of the gas increases while the mass of the liquid decreases, the temperature of both objects remains constant during this period. Finally, the liquid vanishes and the temperature of the gas increases, but so does its pressure.

IDS represents qualitative knowledge of this type in *qualitative schemas*. Our representation has been influenced by Forbus' [3] *qualitative process* (QP) theory, with qualitative schemas corresponding to envisionments in QP theory.

* Lenat's AM [7] is an exception, since it collected its own data and designs its own experiments. But the mathematical domain of AM allowed methods not easily extendable to "real world" domains.

III Inducing Qualitative Schemas

IDS begins with a simple qualitative schema for each of its effectors. For example, the initial schema of the heat effector consists of two states: $s0$, with one object and no active process, and $s1$, with an object touched by a heater and with the temperature of the object increasing. This represents IDS' initial knowledge of the results of applying the heat effector to an object.

The system carries out experiments to improve its schemas, which can be refined in several ways. First, if IDS encounters unfamiliar behavior, it adds a new state to the schema along with a link connecting it to the existing states. Second, the system may discover that an existing state can follow another known state; in this case it simply adds a new link connecting the states. Furthermore, any time new limit-points are found, the system adds quantity-conditions to the states.

Consider again the heat example and the initial heat schema. IDS experiments by applying the heat effector to a block of ice. At first, the temperature of the ice increases, satisfying all conditions of state $s1$. Eventually, a new object (liquid water) appears; after this point the mass of this new object increases, while the mass of the ice decreases. IDS' heat schema does not yet contain a state for this behavior, so the system creates a new state (st) and adds it to the schema. This state has a heater and two objects, b and c . The process slot describes the qualitative behavior of the system - that the mass of object b decreases and the mass of object c increases. Since a new limit-point has been found, quantity-conditions are added to states $s1$ and $s2$. These conditions specify that the temperature of the object in state $s1$ is less than some limit-point C_1 and that the temperatures of the objects in state st are equal to C_1 .

After the ice disappears, state $s1$ again accurately describes the current behavior. When the temperature of the liquid reaches the limit-point C_1 , state st adequately describes the current behavior, so the system does not change the schema at this point. When the liquid disappears, IDS encounters unseen behavior; not only does the temperature of the object increase, but so does its pressure. Thus the system creates a new state ($s2$) and adds it to the schema. After further experimentation using different objects, IDS discovers that the object in $s2$ is always a gas, while the object in $s1$ is either a solid or a liquid. The object description for st is found in a similar way. This information is added, giving the final schema shown in Figure 1.

One can think of this schema-building process as a data-driven search through the space of possible schemas. In these terms, adding states and links make schemas more general, while augmenting the state description and adding quantity-conditions makes them more specific.

IV Discovering Quantitative Laws

Once IDS has formulated a qualitative schema, it uses that knowledge to constrain the search for numeric laws.**

**In addition, schemas provide a context for numeric laws. They describe not only the applicability of laws but also specify their pre-

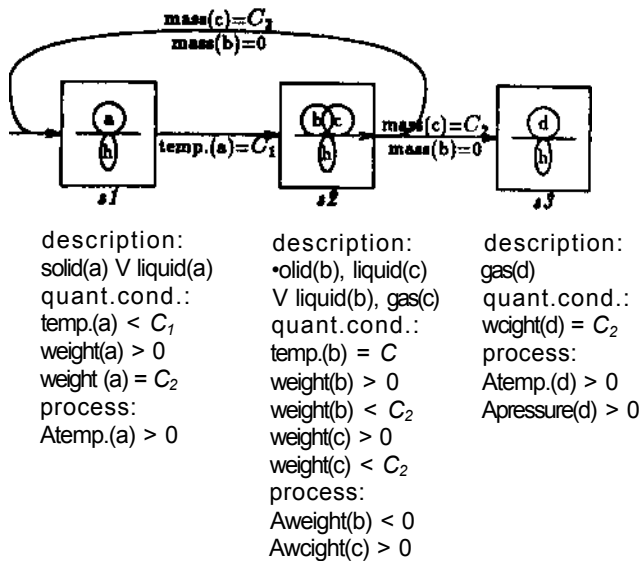


Figure 1: Qualitative Schema for heating an object

The schemas can be viewed as finite state diagrams that describe the behavior of objects over time. States correspond to intervals of time during which objects exhibit some constant behavior. Links specify connections between states, along with the conditions that must be satisfied to enter a successor state.

IDS represents each state as a frame with three slots. The *description* slot includes one or more classifications of the objects present in the state (e.g., *solid* or *acid*). This slot also includes structural descriptions (e.g. heater h touches object a , container a is connected to container b). The *quantity-conditions* slot contains statements about attributes of the objects in the state. These statements are expressed as equalities or inequalities between the quantities of attributes and *limit-points* (see below). The *process slot* is a list of zero or more *changes* that are occurring during the state. Like Forbus, we express a change in terms of the derivative of the changing attribute. For example, an increase in mass of object a is denoted $A_{mass}(a) > 0$.

A state ends only if the process reaches a limit-point, such as the melting-point, or if the agent intervenes, e.g., by turning off the heat. Limit-points are important because they are used in the quantity-conditions, and also because they form the basis for quantitative discoveries. Figure 1 presents a graphical illustration of a heat schema with the object description, the quantity-conditions, and the process for each state.

Although qualitative schemas are structurally similar to the envisionments of De Kleer [1] and Forbus, there is a major difference. Envisionments are *deduced* from structural or process descriptions, while qualitative schemas are *induced* from observations. In the following section we describe this discovery process.

Returning to our heat example, the system would use the schema in Figure 1 to run different experiments. The schema was discovered using a block of ice, so one experiment would examine the effect of varying the initial mass of the ice. Other experiments would vary the class of object used; for instance, IDS might see if the schema still holds when the heated object is hydrogen chloride or some other acid.

Most of the data used in discovering numeric laws are not directly observable, but are gathered in the form of limit-points and state durations. This information is recorded as attribute-value pairs during the matching of a schema to an experimental run. Thus, the system records the values of the limit-point C_1 for different objects and uses these attribute-value pairs as data in its search for numeric laws. Like BACON, the system formulates a quantitative law upon finding some numeric term with a constant value.

IDS discovers two basically different forms of numeric laws. First it finds numeric terms that are constant for all objects of a given class. Langley et al. [4] have called such terms *intrinsic properties*. For example, the system notices that all instances of the class of ice have the same value for the limit-point C_1 . Thus it stores an intrinsic value for the property C_1 and associates this value with the ice class. In fact, this value corresponds to the melting point of water. IDS also discovers that the zero mass is a critical value for all objects, since this is the point when object appear and disappear. This can be viewed as an intrinsic value associated with all objects.

IDS also discovers numeric laws that relate the attributes of different objects within the same instance of a schema. For example, the system notices that the masses of the solid, the liquid, and the gas within the same instance of the heat schema are always equal. Based on this regularity, it postulates a conservation law stating that the mass of an object remains constant as it goes through a phase change.

V Concluding Remarks

In this paper we have described IDS, a system that integrates the process of qualitative and quantitative discovery. We have focused on a single example - involving heat phenomena - to illustrate the acquisition of qualitative schemas and their role in discovering numeric laws. However, the qualitative schema representation and IDS' discovery methods are general enough to cover a wide range of physical and chemical phenomena. For instance, the system has also induced a schema that describes simple chemical reactions and another that describes Black's law of specific heat. We have also used qualitative schemas to represent the fluid-flow of two connected containers filled with liquids [3] and the osmosis of two liquids with different concentrations [8], though IDS has not yet generated this knowledge itself.

We are extending the discovery system on several fronts. Our next step is to incorporate a more robust search mechanism and post-conditions.

anism, such as those used in BACON and ABACUS, to support the discovery of more complex numeric laws. In addition, we must currently supply the system with a concept hierarchy, and are actively extending the system to construct taxonomies on its own initiative. In forming these taxonomies, the next version of IDS will use symbolic attributes, numeric attributes, and information derived from qualitative schemas. As the capabilities of IDS grow, so will the need for an improved agenda mechanism [6] that directs not only the discovery process but also the design of experiments. Even though the IDS project is still in an early phase, it has already led to promising results that have improved our understanding of the complex process of scientific discovery.

Acknowledgments

This work was supported in part by Contract N00014-84-K-0345 from the Information Sciences Division, Office of Naval Research.

We would like to thank Randy Jones and Don Rose for their help on this work along with others from the UCI machine learning group who gave us valuable comments on earlier drafts.

References

- [1] De Kleer, J. "Qualitative and Quantitative Knowledge in Classical Mechanics," Technical Report 325, MIT AI Lab, Cambridge, MA, 1975.
- [2] Falkenhainer, B. C. and R. S. Michalski, "Integrating Quantitative and Qualitative Discovery: The ABACUS System." *Machine Learning* 1:4 (1986) 367-422.
- [3] Forbus, K. D. "Qualitative Process Theory." In D. G. Bobrow (Ed.), *Qualitative Reasoning about Physical Systems*. Cambridge, MA: MIT Press, 1985, pp. 85-168.
- [4] Langley, P., J., G. L. Bradshaw, and H. A. Simon "Rediscovering Chemistry with the BACON System." In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (Eds.), *Machine learning: An Artificial Intelligence Approach*. Palo Alto, CA: Tioga Publishing Company, 1983, pp. 307-329.
- [5] Langley, P. and Nordhausen, B. "A Framework for Empirical Discovery." In *Proceedings of the International Meeting on Advances in Learning*. Les Arcs, France, July, 1986.
- [6] Langley, P., J. Zytkow, H. A. Simon, and G. L. Bradshaw "The Search for Regularity: Four Aspects of Scientific Discovery" In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach, Vol. II*. Los Altos, CA: Morgan Kaufmann, 1986, pp. 425-269.
- [7] Lenat, D. B. "Automated theory formation in mathematics." In *Proc. IJCAI*. Cambridge, MA, August, 1977, pp. 833-841.
- [8] Rajamoney, S., G. DeJong, and B. Falting "Towards a model of conceptual knowledge acquisition through directed experimentation." In *Proc. IJCAI-85*. Los Angeles, CA, August, 1985, pp. 688-690.