

Editorial: Machine Learning and Discovery

Discovery as learning

In everyday language, the terms *learning* and *discovery* convey rather different meanings. The former suggests a gradual process, while the latter suggests a more rapid mental event, often involving some form of insight. Learning may lead to an unconscious change in knowledge, while one is always aware that a discovery has been made. The result of learning can be declarative or procedural, while the product of discovery is always declarative. Learning often involves a transfer of knowledge from teacher to student; in contrast, discovery involves acquiring knowledge from the environment without the aid of a tutor. Finally, all humans and most animals learn from experience, but we reserve the term *discovery* for the accomplishments of a select few. These boundaries are admittedly vague, but they exist nonetheless.

Despite the natural distinction between these concepts, the field of machine learning has always viewed discovery as one of its concerns. Undoubtedly, one reason for this interest is that, like learning, discovery often involves *induction* — the act of reasoning from specific facts or data to general rules or laws which provide a general characterization of those facts. Another reason (probably related to the first) is that, historically, researchers in machine discovery have also worked on learning problems and have applied related techniques to these tasks. Thus, contemporary machine learning researchers tend to view discovery as a difficult form of 'learning from observation' (Carbonell, Michalski, & Mitchell, 1983).

The historical role of discovery within machine learning, together with recent progress in automated discovery, suggested the need for a special treatment of that topic, and the current issue of *Machine Learning* is the result. The three papers in this issue are representative examples of current research in the area. Although they address different aspects of the discovery process, they also share some common underlying themes, which we consider below.

Laws and theories

The history of science suggests a natural distinction between two processes. In empirical discovery, one goes from observational or experimental data to some general laws that *summarize* or *describe* those data. In theory formation, one formulates some structural or process model that goes beyond description to *explain* one or more facts or laws. This dichotomy is probably better characterized as a continuum, but the distinction seems a useful one, with the early stages of a science focusing on empirical discovery and later stages turning to theory formation.

Most AI research has focused on the empirical side of discovery. This has included work on taxonomy formation and conceptual clustering (Michalski & Stepp, 1983), as well as research on the discovery of qualitative laws (Lenat, 1977) and quantitative laws (Langley, 1981). The empirical discovery task lets one focus on simple, general methods for finding regularities in data, and this is one of its attractions. Research on theory formation has been less common and has focused on specific domains, such as automatic programming (Amarel, 1986).

The papers in this issue represent a new stage of research on automated discovery, moving beyond the early efforts mentioned above. Both Kokar's COPER and Falkenhainer and Michalski's ABACUS employ knowledge of attributes' dimensions to constrain the search for laws, and this is an important form of theoretical bias. Rose and Langley's STAHLp moves from data stated as chemical reactions to componential models, which are themselves a simple form of theory. This suggests that we can expect two trends in future work — the use of theory-driven methods to constrain empirical discovery, and insights into the process of theory formation itself.

The importance of replication

Another feature shared by the systems described in this issue is that they all build upon earlier work. ABACUS refines the approach taken by Langley's (1981) BACON system and adds new methods as well, such as the ability to determine the conditions under which a law holds. The COPER program also includes abilities above and beyond those of BACON; for instance, it can determine the relevance of a variable even when its values are held constant. Similarly, STAHLp operates in the same domain as Zytkow and Simon's (1986) original STAHL model and uses many of the same heuristics. However, STAHLp also incorporates other heuristics that make it more robust than its precursor.

Moreover, the authors have tested their systems on many of the same tasks as used by their predecessors. Thus, Falkenhainer and Michalski show how ABACUS can rediscover the ideal gas law, conservation of momentum, and the law of falling bodies, all laws to which BACON had been applied. Kokar also tests his COPER system on the law of falling bodies, as well as a more complex example — Bernoulli's law of fluid flow. Finally, Rose and Langley have used their STAHLp

system to replicate *all* the chemical discoveries reported for Zytkow and Simon's STAHL program.

Replication is a central part of physics and chemistry, and we believe it also has an important role to play in the emerging science of machine learning. The papers in this issue constitute an important step in this direction, and we encourage other authors both to build upon the results of earlier research and to test their improved systems on the same tasks as their predecessors. In this way, we can ensure that our field is making progress, and we can measure the significance of that progress.

The need for synthesis

Another measure of progress involves the integration of previously separate ideas into a single coherent framework. One such approach would integrate methods for empirical discovery and methods for theory formation. A less obvious approach — but the one taken in all three papers — incorporates methods from outside machine discovery to improve the discovery process.

For example, Falkenhainer and Michalski employ a technique for learning from examples, another area of machine learning. Their ABACUS system applies the Aq algorithm (Michalski, 1983) to determine the conditions under which a given numeric law holds. Rose and Langley have borrowed a method from truth maintenance, an entirely different branch of artificial intelligence. Their STAHLp system employs a variant of de Kleer's (1984) assumption-based method for belief revision to modify incorrect chemical reactions. Kokar goes even farther afield, incorporating techniques from dimensional analysis to determine the completeness of a set of numeric terms and to identify the relevance of specific terms.

Again, we feel the papers in this issue set good examples for future work in machine discovery. Researchers should be willing to look beyond methods developed for their particular domain, and even beyond the confines of artificial intelligence. Machine learning is a science of heuristics and algorithms, but it holds no patent on useful mechanisms.

Motivations for machine discovery

As with machine learning in general, there are three basic motivations for studying the mechanisms of discovery: one may be interested in applications that require automating the discovery process, one may be concerned with the general conditions under which discovery can occur, or one may hope to model the historical details of scientific discovery.

The discovery papers in this issue represent all three of these goals. The ABACUS group is clearly concerned with general discovery techniques, but Falkenhainer and Michalski use their system as a nascent "scientist's aide" for

intelligent data analysis and discovery. Kokar's work seems to be motivated by the same dual concerns. Rose and Langley share the desire to understand the conditions for discovery, but they are more interested in modeling the early history of chemistry than in applications.

Of course, these different motives are not exclusive, and there is no reason that ideas developed in pursuit of one goal cannot be used to help achieve another. In fact, this is one of the strengths of our field — that researchers can draw on each others' results to achieve convergent methods despite divergent goals. We feel that the papers in this issue provide good examples of this synergy.

Pat Langley

University of California, Irvine
LANGLEY @ CIP.UCI.EDU

Ryszard S. Michalski

University of Illinois, Urbana
MICHALSK @ B.CS.UIUC.EDU

References

- Amarel, S. (1986). Program synthesis as a theory formation task: Problem representations and solution methods. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. II). Los Altos, CA: Morgan-Kaufmann.
- Carbonell, J.G., Michalski, R.S., & Mitchell, T.M. (1983). An overview of machine learning. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Palo Alto, CA: Tioga.
- de Kleer, J. (1984). Choices without backtracking. In *Proceedings of the Fourth National Conference on Artificial Intelligence* (pp. 79–85). Austin, TX: Morgan-Kaufmann.
- Langley, P. (1981). Data-driven discovery of physical laws. *Cognitive Science*, 5, 31–54.
- Lenat, D.B. (1977). Automated theory formation in mathematics. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (pp. 833–842). Cambridge, MA: Morgan-Kaufmann.
- Michalski, R.S. (1983). A theory and methodology of inductive learning. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Palo Alto, CA: Tioga.
- Michalski, R.S., & Stepp, R.E. (1983). Learning from observation: Conceptual clustering. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Palo Alto, CA: Tioga.
- Zytkow, J.M., & Simon, H.A. (1986). A theory of historical discovery: The construction of componential models. *Machine Learning*, 1, 107–136.