

TWO FRAMEWORKS FOR INTEGRATING KNOWLEDGE IN INDUCTION

Paul S. Rosenbloom
 Information Sciences Institute &
 Department of Computer Science
 University of Southern California
 4676 Admiralty Way
 Marina del Rey, CA 90292

Haym Hirsh
 Department of Computer Science
 Hill Center, Busch Campus
 Rutgers University
 New Brunswick, NJ 08903

William W. Cohen
 AT&T Bell Laboratories
 600 Mountain Avenue
 Room 3C-412A
 Murray Hill, NJ 07974

Benjamin D. Smith
 Information Sciences Institute &
 Department of Computer Science
 University of Southern California
 4676 Admiralty Way
 Marina del Rey, CA 90292

ABSTRACT

The use of knowledge in inductive learning is critical for improving the quality of the concept definitions generated, reducing the number of examples required in order to learn effective concept definitions, and reducing the computation needed to find good concept definitions. Relevant knowledge may come in many forms (such as examples, descriptions, advice, and constraints) and from many sources (such as books, teachers, databases, and scientific instruments). How to extract the relevant knowledge from this plethora of possibilities, and then to integrate it together so as to appropriately affect the induction process is perhaps the key issue at this point in inductive learning. Here we focus on the integration part of this problem; that is, how induction algorithms can, and do, utilize a range of extracted knowledge. Preliminary work on a transformational framework for defining knowledge-intensive inductive algorithms out of relatively knowledge-free algorithms is described, as is a more tentative problem-space framework that attempts to cover all induction algorithms within a single general approach. These frameworks help to organize what is known about current knowledge-intensive induction algorithms, and to point towards new algorithms.

INTRODUCTION

Inductive learning is a process whereby a definition of a concept is derived from a set of positive, and sometimes negative, examples of the concept. Key issues in inductive learning are the accuracy of the resulting definition – that is, the error rate it yields in classifying new examples – and the resources required to generate the definition (in terms of the number of examples and/or the amount of time and space needed).

The single most promising route towards reducing both the error rate and resource usage of inductive learning is to utilize whatever additional knowledge is available beyond the examples; that is, to convert induction from a weak to a strong method. However, to do this, the relevant knowledge must first be extracted from the sources in which it exists, such as books, teachers, databases, and scientific instruments. This extracted knowledge must then be integrated together for use by the induction process, in whatever form is appropriate – examples, descriptions, advice, constraints, or anything else. Here we focus on the integration task (extraction involves potentially everything from vision to natural language understanding, and more). Our goal is to begin the process of deriving principles for how knowledge-intensive induction algorithms both do now, and can in the future, provide such integration. The hope is that this will lead to both a useful descriptive framework for organizing existing approaches, as well as a prescriptive framework for generating new approaches that go beyond the existing ones.

We'll make this beginning by presenting two partial frameworks for knowledge integration in induction, along with implications drawn so far by applying them to four recently proposed knowledge-intensive induction algorithms. The focus here is specifically on knowledge integration for induction, rather than on the broader issue of knowledge integration in general, in the hope that the extra structure provided by the induction problem will lead to more powerful integration strategies than have been proposed for the general case. The more developed of the two knowledge-integration frameworks – and thus the one emphasized in this paper – is the *transformational* framework. It describes how knowledge-intensive induction algorithms are, and can be, derived by transforming traditional learning methods. The more tentative *problem-space* framework attempts to go beyond the transformational framework to the more difficult task of characterizing the fundamental components of all induction algorithms, whether knowledge-intensive or not. This framework is covered only briefly here as an intriguing possibility for the future.

THE TRANSFORMATIONAL FRAMEWORK

An induction algorithm can be viewed abstractly as a black box with one output port for the concept description and one or more input ports. A minimal induction algorithm has just one input port, for training data, with all other information being hardwired into the algorithm as a fixed *bias* [6]. The only way such an algorithm can use additional knowledge – other than by reprogramming – is to find some way of recoding the knowledge as pseudo-examples. For example, Quinlan describes how knowledge about type restrictions on the arguments of predicates could conceivably be used indirectly by the FOIL algorithm through the generation of pseudo-negative examples that cause FOIL to eliminate candidate hypotheses that would violate the type restrictions [8].

Most induction algorithms actually do provide some additional input ports that allow explicit provision of other types of information; that is, of knowledge beyond what is embodied in the examples. For example: the candidate elimination algorithm provides an input port for information concerning the partial ordering that defines the initial hypothesis space [4]; FOIL provides an input port for the set of relations that can be used in candidate hypotheses [8]; backpropagation provides input ports for learning-rule parameters, the network structure, and the initial connection weights [11]; and Bayesian learning algorithms provide an input port for prior probabilities [1]. Such ports expand the types of information that can be utilized at induction time, but still provide a very limited means for incorporating the full range of knowledge that may be available.

The transformational framework starts with the basic notion of black boxes and ports, as described above, and views knowledge-intensive induction algorithms as the composition of a core, usually knowledge-lean, algorithm plus a set of transformations. Although not all knowledge-intensive algorithms can be viewed in this fashion, when they can, the results can be quite informative. The four knowledge-intensive algorithms covered in the next section do all fit this framework, and are based on three distinct core algorithms – candidate elimination, FOIL, and backpropagation – and on two general types of transformations to these core algorithms:

- A *reformulation* transformation modifies the core algorithm so that its ports can handle a wider range of inputs, either by generalizing its existing ports or by adding new ones. A simple example of a reformulation is the modification of a decision-tree learner to allow it to accept a task-specific split criterion from an input port, rather than always using the same built-in criterion. A more sophisticated example of a reformulation is IVSM's derivation from the candidate elimination algorithm by converting its examples input port to take a more expressive class of inputs (i.e., version spaces) [3].
- A *preprocessor* transformation adds to the core algorithm a preprocessor that takes a form of input beyond what can be fed directly into the core algorithm's input ports, and translates this broader input into something that one or more of the core input ports can understand.

Quinlan's suggestion of using type constraints to create pseudo-negative examples is exactly a proposal for a preprocessor transformation. The preprocessor would have an input port that can accept type constraints, and would produce negative examples that can be fed into FOIL's existing input port. The combination of FOIL and this preprocessor thus defines a new learning algorithm that can take as input not only examples and relations, but also type constraints.

EXAMPLES

Much recent work on induction algorithms has focused on enhancing their ability to utilize additional knowledge during induction, and thus there are many learning systems we could consider. A full survey of such algorithms is beyond the scope of this paper, so we will focus here on just four recent knowledge-intensive algorithms, each of which provides the ability to utilize EBL-like domain theories, plus possibly some other forms of knowledge:

- IVSM has the ability to utilize EBL-like domain theories plus models of bounded inconsistency [3].
- FOCL has the ability to utilize (possibly partial) EBL-like domain theories plus constraints on predicate arguments [7].
- GRENDEL has the ability to specify the hypothesis space via a formal grammar – which can include an EBL-like domain theory – plus some simple ordering information [2].
- KBANN is a neural network algorithm that has the ability to utilize an EBL-like domain theory [13].

The remainder of this section considers these four systems in more detail.

IVSM is based on the candidate-elimination algorithm (CEA). It is derived by a reformulation of the CEA so that instead of basing its inner loop on the process of updating a version space with respect to a single example, it now updates the version space with respect to a second version space (by intersecting the two version spaces). This reformulation generalizes the CEA's examples input port so that it now accepts version spaces. In addition to this core reformulation transformation, IVSM also uses three distinct preprocessor transformations that are enabled by this reformulated input port. One preprocessor allows IVSM to emulate the CEA by taking examples and converting them into version spaces. A second preprocessor creates version spaces from combinations of examples and EBL-style domain theories. A third preprocessor creates version spaces from combinations of examples and a model of bounded inconsistency. When IVSM is combined with any one of these preprocessors, it actually yields a new induction algorithm: IVSM-CEA, IVSM-EBL, or IVSM-BI.

FOCL is based on FOIL. FOIL uses the information provided on its relations input port to determine what modifications to consider making to the current candidate hypothesis. Essentially, it considers adding the various relations – as instantiated with a mixture of old and new variables – to the current clause of the hypothesis, and uses an information-theoretic measure to determine which possibility is (locally) best. FOCL reformulates this possibility-generation strategy in two ways. First, it increases the set of possibilities by considering adding combinations of relations in a single step, rather than just individual relations. Second, it decreases the set of possibilities by eliminating those that violate constraints on the arguments of the relations (such as type and uniqueness restrictions). The first reformulation supports the addition of an input port for (possibly partial) EBL-like domain theories; the combinations of relations that occur in the condition sides of these rules form the basis for the relation combinations proposed in the reformulated algorithm. The second reformulation supports the addition of an input port for type and uniqueness constraints on the arguments of the relations that are proposed. This new port directly supports knowledge that FOIL would have needed a preprocessor to use.

GRENDDEL is also based on FOIL. The core transformation made in developing GRENDDEL is also a reformulation of FOIL's generation strategy for possible modifications to the current candidate hypothesis. However, GRENDDEL's reformulation is both more radical and more general. GRENDDEL generates possibilities by consulting generation rules specified in a context-free grammar. This supports broadening FOIL's input port from one that can take a list of relations to one that can handle a list of context-free grammar rules. A second smaller reformulation allows the processing of possibilities to be selectively deferred, and supports the addition of a second input port to specify this simple ordering information. The remainder of the GRENDDEL story is much like IVSM. The generalized input port facilitates the creation of a number of preprocessors that can accept a variety of types of input. This input is then translated into grammar rules that can be fed to this new port. These preprocessors allow GRENDDEL to accept the kinds of input utilized by (among others) EBL, FOIL, and FOCL, and thus to emulate these other algorithms. As with IVSM, there is a base GRENDDEL algorithm which takes grammars as inputs, and then there are a number of other induction algorithms that are based on GRENDDEL, such as GRENDDEL-EBL and GRENDDEL-FOIL, which are derived from it by adding specific preprocessors.

Finally, KBANN is based on backpropagation. It adds a preprocessor that takes as input an EBL-like domain theory, plus a list of environmental features not covered by the theory, and translates this knowledge into a form that can be fed into backpropagation's initial network topology and initial network weights ports. It leaves backpropagation's remaining input ports – such as its learning-rule parameters – intact.

ANALYZING INPUT PORTS IN THE TRANSFORMATIONAL FRAMEWORK

The transformational framework makes it possible to examine knowledge-intensive learners in more detail, by studying the set of input ports provided by the resulting algorithms, what kind of knowledge they can accept, and what key properties they possess (or fail to possess). Although we are still in the process of identifying what the key properties are for input ports, the list already includes at least two that seem critical.

- The *additivity* of an input port is determined by its ability to accept multiple independent fragments of knowledge at that port. Additivity is important because additive ports can serve directly as integrators for arbitrary amounts of knowledge of the types that they can accept. The prototypical example of an additive input port is the example port in standard induction algorithms. It can accept arbitrary amounts of new information, and combine it straightforwardly with whatever else the system knows. A classical example of a non-additive port is the learning-rate parameter in backpropagation. If more information is available, how should it be combined with what is already known? Must the old information simply be eliminated, and replaced by the new, or should the two values be averaged, or should something else happen?

For additive ports, the way in which inputs are combined usually depends on their interpretation. Examples can be viewed as constraints on the behavior of the concept being learned, so they are usually combined via an intersection operator. Other types of information might be combined via different operators, such as union or average.

- The *ease of use* of an input port is determined by how easy it is to express knowledge in the language provided by the port. Bayesian priors are a classic case of a difficult-to-use input port, with this difficulty most likely being the single biggest stumbling block in using Bayesian approaches to learning. Sometimes preprocessors can be added to make a port easier to use; however, the port's basic ease of use will still affect how easy it is to write the preprocessors. A good example of such a preprocessor for Bayesian priors is the use of the *minimum description length* principle, which, while it can be viewed as a Bayesian approach, replaces the task of assigning a prior probability to every concept with the arguably simpler

task of choosing an encoding scheme [9].

To illustrate these two properties of knowledge-intensive induction methods, as viewed from the transformational framework, we return to the four algorithms discussed above.

The core IVSM algorithm has two input ports, one for the partial-order information on which the version spaces are based and one for a collection of version spaces. The partial-order port is additive because it can handle an arbitrary number of elements plus ordering relations among them. It is also easy to use, but only for the narrow purpose of identifying (possibly parts of) candidate hypotheses and generality relationships among them. The version-space port is also an additive port – as with the traditional example input port, it can accept an unbounded set of inputs, and combine them (via version-space intersection) with what is already known. Its ease of use is intermediate between that provided by example ports at the low end (at least if they are being used for anything other than just examples) and languages like GRENDEL's grammars at the high end. When IVSM's preprocessors are considered, there are three new input ports, all of which are additive and relatively easy to use (for the restricted uses for which they are intended).

One idea that is directly suggested by this analysis of IVSM is that there is no reason its three distinct preprocessors couldn't all be used simultaneously. Because they all output version spaces, and the version-space port is additive, it should be possible to intermingle information based on examples, domain theories, and bounded inconsistency (thus effectively creating a new algorithm that subsumes the three existing ones).

FOCL's three input ports – for examples, (possibly partial) domain theories, and argument constraints – are all additive, as they can all accept arbitrary amounts of knowledge of their chosen input types. They are also all easy to use for their intended purposes, but difficult to use for other purposes.

GRENDEL's three input ports accept examples, grammars, and ordering information (information about what portions of the hypothesis space should be tried first). Regarding ease of use, the example port has the standard properties; the ordering port is similar to a Bayesian-priors port but likely to be somewhat easier to use because it is much less demanding; and the grammar port is relatively easy to use for most purposes. The example and ordering ports are both additive; however the grammar port is only semi-additive, in that the grammars are closed under union, but not under intersection. Thus the additivity of the grammar port depends on the way in which grammars are used. If a grammar is used as a suggestion as to which hypotheses are most likely – as when grammars are used to encode a domain theory – then grammars can be easily combined with a union operator. However, when grammars are used as constraints on the hypothesis space, it is impossible to generate a separate grammar for each constraint and then integrate the constraints by intersecting the grammars (as IVSM would intersect its version spaces).

KBANN's three input ports accept examples, domain theories, and environmental features. The examples port is much like any other examples port – it is additive and easy to use for its intended purpose (but difficult to use for other purposes). The domain theory port is additive and easy to use. The environmental-features port is like the examples port, being additive and easy to use for very limited purposes.

IMPLICATIONS OF THE TRANSFORMATIONAL FRAMEWORK

Pulling back up now from these detailed analyses to look at the picture more globally, several general implications can be discerned. The first implication is that multiple pieces of knowledge can be combined in three distinct fashions. The first approach feeds the knowledge into multiple of the core algorithm's input ports, and depends on the structure of the core algorithm to perform the integration. For example, KBANN integrates a domain theory with examples by feeding the domain theory to the core

network topology and weight ports, while feeding examples directly to the core examples port. The core algorithm – that is, backpropagation – then combines this knowledge during its normal processing. The second approach utilizes a multi-ported preprocessor that integrates the knowledge provided to its input ports in the process of generating input for the core algorithm. One example is GRENDDEL-FOCL's emulation of FOCL via a preprocessor that combines knowledge from all of FOCL's input ports (except for the examples port) in the process of converting this knowledge into a single grammar for use by GRENDDEL. A second example is IVSM-EBL's use of a preprocessor to integrate knowledge from its examples and domain-theory ports in the process of generating version spaces for the core IVSM algorithm. The third integration approach is to utilize an additive port that can integrate across multiple pieces of knowledge sent to a single port. IVSM is a good example of this, as its version-space port is an effective additive input port.

The second implication is that the insights underlying different knowledge-intensive algorithms can often be transferred or combined in useful ways. In cases where two knowledge-intensive algorithms are based on the same core algorithm, and where they have transformed the core algorithm in different ways, it should be possible to combine many of the transformations without a great deal of difficulty. For example, GRENDDEL's generalization of FOIL's relations port to accept grammars could be combined with FOCL's techniques for pruning hypotheses using typing constraints. It would be an interesting question to see whether this approach would have any advantages over using a preprocessor to incorporate all of FOCL's knowledge into a GRENDDEL grammar, as in GRENDDEL-FOCL.

In cases where the core algorithms are different, transfer of a more abstract sort can still occur. For example, IVSM's additivity based on version-space intersection leads to asking whether GRENDDEL's grammars could support a comparable operation: the answer is no, since context-free grammars are not closed under intersection. This also suggests the new research topic of modifying GRENDDEL so that it is more additive. For example, since the intersection of a context-free language and a regular language is a context-free language, it might be possible to create a new version of GRENDDEL that has an additive port for regular languages in addition to the existing (non-additive) port for context-free languages.

The third implication is that additional effort would be usefully spent looking at how the two general classes of transformations could be applied to further aspects of existing algorithms, both those considered here as well as others.

BEYOND THE TRANSFORMATIONAL FRAMEWORK

The transformational framework is somewhat unsatisfying for several reasons: it does not apply to all knowledge-intensive learners; it does not apply to knowledge-weak learners (which actually do achieve some forms of knowledge-integration even in simply being able to accept varying numbers of examples and learn from them); and it doesn't say much about how to merge the insights across knowledge-intensive algorithms that have different core algorithms. Our continuing work attempts to go beyond the transformational framework by developing a *problem space* framework that attempts to identify the core functionalities that underly all induction algorithms, and then to understand how all of the knowledge utilized by a learner – examples, domain theories, etc. – is integrated together via its mapping on to these functionalities. In terms of the transformational framework, the goal here can be expressed as finding a single black box and set of input ports that conceptually lie at the heart of all induction algorithms.

The problem-space framework is organized around the concept of the space of candidate hypotheses, thereby continuing the existing line of analyses that have viewed induction as search [12; 5; 10]. In this framework the role of knowledge is first off to specify, constrain, and order the elements – that is, the states – of this space. In the four algorithms we have focused on here, specification of the states in the space occurs rather directly via GRENDDEL's grammar port, FOCL's relations port, and IVSM's partial-

order port. Constraints on the set of states considered are provided by IVSM's version spaces and FOCL's argument constraints. Ordering information about the states is provided by GRENDEL's ordering port and GRENDEL's and FOCL's examples ports (though rather indirectly, through their information-theoretic measures). However, none of the four systems allows direct statement of all three types of knowledge. GRENDEL comes the closest, though it requires all constraints to be stated indirectly in terms of what can be generated via the grammar. KBANN is the furthest away, as it cannot accept direct statement of any of these types of knowledge. It does however accept some such information indirectly; for example, its domain theory (plus information about additional domain features) indirectly determines what can and cannot be in the hypothesis space, by determining the network topology.

The remaining use of knowledge in this framework is to provide method-specific knowledge about how to search the space of hypotheses. IVSM is at one extreme, in that it makes no use of such knowledge – it always maintains a representation of all hypotheses that are consistent with all of the knowledge available so far. FOCL, GRENDEL, and KBANN all utilize greedy search algorithms. FOCL uses its relation and domain-theory ports to generate candidate changes at each step, its argument-constraint port to eliminate candidates, and its examples port to order the candidates (via its information-theoretic measure). GRENDEL uses the detailed structure of its grammar rules to generate the candidates at each step – two grammars that generate the same terminal language could lead to different greedy searches if they are specified in terms of different sets of rules. It also uses the information from its ordering port as a first cut at ordering the candidates, and then its examples port to complete the ordering (again via its information-theoretic measure). KBANN uses its examples port to determine the direction in which to descend the gradient in its greedy search (via backpropagation) and its learning-rate port to determine the size of the steps taken in that direction.

Although the problem-space framework is still in a very preliminary stage of development, one insight already revealed by this analysis is that, though all four of the knowledge-intensive algorithms studied here use EBL-like domain theories, they use them in three qualitatively different ways. Two of the algorithms – KBANN and IVSM – trust their domain theories enough to use them to directly affect the space of candidate hypotheses, though they do this in different fashions. KBANN uses the domain theory to specify the initial space of candidate hypotheses (that is, the network structure). In contrast, IVSM uses the domain theory (along with examples) to constrain the space of candidate hypotheses that was earlier generated from information provided to its partial-order port. FOCL distrusts its domain theory sufficiently to allow it to affect only the search strategy; that is, the domain theory is used only to order the search for a hypothesis, and never to prune the space. It thus gets less constraint from its domain theory, but is also able to recover more gracefully if the theory is wrong. GRENDEL's treatment of the domain theory depends on how the domain theory has been converted into a grammar; GRENDEL can employ either a KBANN-like strategy, in which the theory determines the search space, or a FOCL-like strategy, in which the theory orders the search space. In GRENDEL-FOCL, the variant of GRENDEL discussed above, the domain theory orders the search space.

As work continues on the problem-space framework, the insights derived from it should (hopefully) get both broader and deeper.

CONCLUSION

We have begun the process of understanding knowledge-intensive induction algorithms by presenting a transformational framework for creating knowledge-intensive methods from knowledge-weak methods, using the framework to analyze four recent algorithms, and deriving from these analyses general implications about the integration of knowledge in induction. We also described a more preliminary problem-space framework that attempts to identify the core functionalities of any learning method and

how various learning methods are created by mapping out how knowledge sources can be used to define these functionalities.

Beyond what has already been described, one fundamental insight revealed by these two frameworks, and the analyses they yield of existing knowledge-intensive learners, is a path towards simple yet powerful knowledge-intensive induction algorithms. First, additive ports need to be developed that provide broad languages for the basic functionalities of specifying, constraining and ordering hypothesis spaces. Ideally, such ports and languages should combine, for example, the best aspects of IVSM's version spaces and GRENDEL's grammars, yet still cover all of these basic functionalities. Second, comparable ports need to be developed to allow knowledge to be used in whatever search method is chosen. For greedy methods, this tends to be knowledge about proposing, constraining, and ordering the options at each step. Third, a range of preprocessors need to be created that can translate a wide variety of forms of knowledge into these ports. Ultimately this leads to a direct concern about knowledge extraction, as the preprocessors get closer and closer to the prime sources of knowledge (such as books), and thus raises a variety of additional issues about how and when knowledge is extracted. Ultimately the hope is to complete these two frameworks, fuse them into a single more comprehensive framework, analyze the full space of existing knowledge-intensive induction algorithms, and use the resulting insights to build one or more new algorithms that go significantly beyond the existing ones.

ACKNOWLEDGMENTS

This project is partially supported by the National Aeronautics and Space Administration (NASA) under cooperative agreement number NCC 2-538.

REFERENCES

1. Buntine, W. Classifiers: A theoretical and empirical study. 12th International Joint Conference on Artificial Intelligence, Sydney, 1991, pp. 638-644.
2. Cohen, W. W. Compiling prior knowledge into an explicit bias. *Machine Learning: Proceedings of the Ninth International Workshop*, Aberdeen, 1992, pp. 102-110.
3. Hirsh, H.. *Incremental Version Space Merging: A General Framework for Concept Learning*. Kluwer Academic Publishers, Boston, MA, 1990.
4. Mitchell, T. M. Version spaces: A candidate elimination approach to rule learning. *Proceedings of the 5th International Joint Conference on Artificial Intelligence, IJCAI*, Cambridge, MA, 1977, pp. 305-310.
5. Mitchell, T. M. An Analysis of Generalization as a Search Problem. *Proceedings of IJCAI-79*, 1979.
6. Mitchell, T. M. The Need for Biases in Learning Generalizations. Tech. Rept. CBM-TR-117, Department of Computer Science, Rutgers University, 1980. (Reproduced in Shavlik, J. W. & Dietterich, T. G. (1990). *Readings in Machine Learning*. San Mateo, CA: Morgan Kaufmann.)
7. Pazzani, M. J., & Kibler, D. "The utility of knowledge in inductive learning." *Machine Learning* 9 (1992), 57-94.
8. Quinlan, J. R. "Learning logical definitions from relations." *Machine Learning* 5 (1990), 239-266.
9. Quinlan, J. R. & Rivest, R. L. "Inferring decision trees using the minimum description length principle." *Information and Computation* 80 (1989), 227-248.
10. Rosenbloom, P. S. Beyond generalization as search: Towards a unified framework for the acquisition of new knowledge. *Proceedings of the AAAI Symposium on Explanation-Based Learning*, AAAI, Stanford, CA, 1988, pp. 17-21.
11. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*, D. E. Rumelhart, J. L. McClelland, and the PDP Research Group, Eds., Bradford Books/MIT Press, Cambridge, MA, 1986, pp. 318-362.
12. Simon, H. A., & Lea, G. Problem solving and rule induction: A unified view. In *Knowledge and Cognition*, L. W. Gregg, Ed., Erlbaum, Potomac, MD, 1974, pp. 105-127.
13. Towell, G. G., Shavlik, J. W., & Noordewier, M. O. Refinement of approximate domain theories by knowledge-based neural networks. *Proceedings of the Eighth National Conference on Artificial Intelligence*, AAAI, Boston, MA, 1990, pp. 861-866.