Applications of Inductive Logic Programming

Ivan Bratko and Stephen Muggleton

Techniques of machine learning have been successfully applied to various problems [1, 12]. Most of these applications rely on attribute-based learning, exemplified by the induction of decision trees as in the program C4.5 [20]. Broadly speaking, attribute-based learning also includes such approaches to learning as neural networks and nearest neighbor techniques. The advantages of attribute-based learning are: relative simplicity, efficiency, and existence of effective techniques for handling noisy data. However, attribute-based learning is limited to non-relational descriptions of objects in the sense that the learned descriptions do *not* specify *relations* among the objects' parts. Attribute-based learning thus has two strong limitations:

A logic-based approach to machine learning is evaluated for its problem-solving merit in several industrial and academic situations.

- the background knowledge can be expressed in rather limited form, and
- the lack of relations makes the concept description language inappropriate for some domains.

Examples of such domains are presented in this article.

An attempt to overcome the limitations of attribute-based learning has led to recent development of a number of programs that learn at the level of first-order predicate logic. These include FOIL [19], Golem [17], and Progol [22], with Shapiro's program MIS as one of their early predecessors [21]. This has led to the inception of a new area of machine learning called *Inductive Logic Programming* [13]. For recent developments see [14]. ILP benefits from the solid theoretical framework provided by logic and logic programming.

The learning problem in ILP is normally stated as follows: given *background knowledge B*, expressed as a set of predicate definitions, positive exam-

Illustration: Rico Lins ©

ples E+ and negative examples E-, an ILP system will construct a predicate logic formula H such that:

- all the examples in E+ can be logically derived from B ∧ H, and
- no negative example in E- can be logically derived from B ∧ H.

This definition is similar to the general problem of inductive learning, but it insists on predicate logic representation of B and H. Typically for ILP systems B, H, E+ and E- will each be Prolog programs. B ∧ H is simply the Prolog program B extended by the Prolog program H. The use of Prolog throughout allows for a highly versatile representation language for all constituents of the problem. This versatility is reflected in the wide variety of ILP applications. ILP differs from other machine learning approaches owing to its insistence on a particular representation language. This has advantages in integrating techniques and theory with those inherited from the field of logic programming.

One of the main advantages of ILP over attribute-based learning is ILP's generality of representation for background knowledge. This enables the user to provide, in a more natural way, domain-specific background knowledge to be used in learning. The use of background knowledge enables the user both to develop a suitable problem representation and to introduce problem-specific constraints into the learning process. By contrast, attribute-based learners can typically accept background knowledge in rather limited form only. So in ILP, if the problem is to learn to distinguish cyclic from acyclic graphs, the graphs can be introduced by representing their edges as background knowledge. In addition, a recursive definition of the notion of a path within a graph can be provided. If the problem is to learn about properties of chemical compounds, the molecular structures can be introduced as background knowledge in terms of the atoms and bonds between them. If the task is to automatically construct a model of a physical system from the observed behaviors, complete mathematical apparatus that is considered relevant to the modeling domain is included in the background knowledge. Application of ILP involves development of a good representation of the examples together with relevant background knowledge. A general purpose ILP system is then applied.

The ILP framework can also be applied to automatic program synthesis from examples as follows. The existing, known predicates are introduced to a general ILP system as background knowledge. The target program is specified by examples of its input/output vectors. A common ILP exercise of this kind is the induction of the quick-sort program from examples, saying for instance that the list `[4,1,2]` sorts into `[1,2,4]`. Suitable background knowledge contains the definition of the predicates for list concatenation, and for partitioning of a list, with respect to some value, into the lists of "small" and "big" elements. Using this background knowledge and some ten examples and counterexamples, a typical ILP system will induce the known Prolog program for quick-sort in a few seconds of CPU time.

In the remainder of this article we describe selected applications of ILP. We mainly chose those applications that specifically benefit from the ILP's predicate logic descriptions, and from the background facility in ILP.

Finite Element Mesh Design

Finite element (FE) methods are used extensively by engineers and modeling scientists to analyze stresses in physical structures. Finite element methods require that the structures being modeled be partitioned into a finite number of elements, resulting in a finite element mesh (Figure 1). In order to design a numerical model of a physical structure it is necessary to decide the appropriate resolution of the mesh. Considerable expertise is required in choosing these resolution values. Too fine a mesh leads to unnecessary computational overheads when executing the model. Too coarse a mesh produces intolerable approximation errors.

Normally some regions of the object require a denser mesh whereas in other regions a coarser mesh still suffices for good approximation. There is no known general method that would enable automatic determination of optimal, or even reasonably good meshes. However, many examples of successful meshes for particular objects have been accumulated in the practice of FE computations. These meshes can be used as sources of examples for learning about the construction of good meshes.

In general the mesh depends on the geometric properties of the object, the forces acting on it, and the relations between different components of the object. The mesh density in a region of the object depends also on the adjacent

Figure 1.
Partition of an object into a finite element mesh
(courtesy Bojan Dolsak)

regions. Because of these relational dependences, the ILP approach most naturally applies to the mesh design problem.

In the application of ILP to this problem [7], a structure to be partitioned is represented as (1) a set of edges, (2) the properties of the edges, and (3) relations among the edges. These properties and relations are represented as part of background knowledge by predicates, such as: short(Edge), loaded(Edge), two_side_fixed(Edge), neighbor_xy(Edge1,Edge2), etc. In experiments to learn a characterization of the density of a mesh in terms of these relations, ten meshes known to be numerically adequate have been used as sources of examples for learning. The target relation to be learned is: mesh(Edge,N) where Edge is the name of an edge in the structure, and N is the recommended number of finite elements along this edge. The available meshes comprise several hundreds of edges. Each edge is used as an example for learning, and typically some additional positive examples are derived from the meshes. The typical number of examples is between 300 and 600. Negative examples are generated by a form of closed-world assumption which gives rise to several thousands of negative examples.

Several relational learning algorithms were tried on this data including Golem [17], FOIL [19] and CLAUDIEN [6]. The resulting set of rules were of interest to expert users of the finite element methods. According to their comments, these rules reveal interesting relational dependences that the experts had not been aware of. Here we give an example of such a rule (in Prolog syntax as output by Golem, except that the variables were mnemonically renamed).

```
mesh( Edge, 7) :-
   usual_length( Edge),
   neighbor_xy( Edge, EdgeY),
   two_side_fixed( EdgeY),
   neighbor_zx( EdgeZ, Edge),
   not_loaded( EdgeZ).
```

This rule says that partitioning Edge into 7 elements is appropriate if Edge has "usual length", and has a neighbor EdgeY in the xy-plane so that EdgeY is fixed at both ends, and Edge has another neighbor EdgeZ in the xz-plane so that EdgeZ is not loaded. The following recursive rule was also generated by Golem.

```
mesh( Edge, N) :-
   equal( Edge, Edge2),
   mesh( Edge2, N).
```

This rule is interesting as it expresses, by recursion, a recurrent pattern in the structure. The rule observes that an edge's partition can be determined by looking for an edge of the same length and shape positioned symmetrically in the same object. In other words, this can be viewed as Golem's discovery that an edge may inherit a suitable partition from similar edges in the structure. Of course, for this rule to be computationally useful, at least some of the equivalent edges must have their partitions determined by some other rule.

The accuracy of the induced rule sets was investigated in detail in [7]. One method for estimating accuracy is cross-validation whereby a subset, say 90%, of all the available examples (edges) are used for learning, and the remaining examples are used for testing. Using this method, the test set accuracy of the rules induced by Golem was (on average) found to be as follows: the rules suggested correct partition of an edge into finite elements in 78% of all test cases, incorrect in 2% of the cases, and 20% of the test edges remained undecided (not covered by the induced clauses). Although the proportion of undecided edges here seems rather high, it is within an acceptable range for the practice of mesh design. Because of some general local consistency constraints used in mesh generators, many of the omissions can be automatically recovered by post-processing of the results generated by the ILP system.

Predicting the Mutagenicity of Chemical Compounds

The construction of new scientific knowledge from real-world data remains an active focus for machine learning. One such problem is the Structure/Activity Relationships (SAR) of chemical compounds. This forms the basis of rational drug design. One widely used method of SAR stems from the work of Hansch [10] and uses regression/discrimination to predict activity from molecular properties such as hydrophobicity, sigma effect, molar reflectivity and LUMO (the energy of the Lowest Unoccupied Molecular Orbital). This and many other traditional approaches are limited in their representation of molecular connectivity and structure. They take into account the global attributes of a molecule, but do not comprehensively consider the *structural relationships* in the molecule. Thus some possibly important information, comprised as patterns in the molecular structure, may remain unexploited.

The ILP approach allows, however, the complete structural information to be taken into account. An ILP system *Progol* has been applied to the problem of identifying Ames test mutagenicity within a series of heteroaromatic nitro compounds [15, 22]. Hansch and coworkers have studied 230 compounds using classical regression [5]. For 188 compounds, they successfully obtained a linear regression function using hydrophobicity, LUMO and two handcrafted binary attributes indicative of some structural properties. This regression formula predicts high mutagenicity with very acceptable accuracy. However the remaining 42 compounds could not be successfully modeled by regression and no structural principles were proposed. This subset of 42 compounds will be

nitrofurazobe

4-nitropenta[cd]pyrene

High mutagenicity

6-nitro-7,8,9,10-tetrahydrobenzo[a]pyrene

4-nitroindole

Low mutagenicity

(A)

(B)

therefore referred to as "regression unfriendly." *Progol* was applied to this mutagenicity data using the split of the compounds into those with high mutagenicity and the rest as suggested by Hansch and coworkers. All the compounds were represented relationally in terms of atoms, bonds and their partial charge. This information was automatically generated by the modeling program QUANTA, and was represented as about 18300 Prolog facts (unit Horn clauses) for the entire set of 230 compounds. For the 188 compounds found to be amenable to regression, the additional Hansch attributes of LUMO and hydrophobicity were also provided. All this was supplied to *Progol* as background knowledge for learning. For these compounds, *Progol* constructed the following theory. A compound is highly mutagenic if it has (1) a LUMO value ≤ -1.937; or (2) a LUMO value ≤ -1.570 and a carbon atom merging six-membered aromatic rings; or (3) a LUMO value ≤ -1.176 and an aryl-aryl bond between benzene rings; or (4) an aliphatic carbon with partial charge ≤ -0.022. The theory has an estimated accuracy of 89%. This matches the accuracy of both the regression analysis of Hansch and coworkers, and a more recent effort using neural networks [24]. It should be noted, however, that *Progol*'s theory is easier to comprehend and was generated automatically, without access to any structural indicator variables handcrafted by experts specifically for this problem.

The advantage of ILP, however, became particular-ly clear on the remaining subset of the 42 "regression unfriendly" compounds. For these, *Progol* derived a single rule with an accuracy of 88% estimated from a leave-one-out validation (Figure 2). This is significant at $P < 0.001$. In contrast, linear regression and linear discrimination on the parameters used by Hansch and coworkers yield theories with accuracies estimated at 69% and 62%, which are no better than random guesses supplied with the default accuracy of 69%. Perhaps even more important than the predictive accuracy, *Progol*'s rule provides the new chemical insight that the presence of a five-membered aromatic carbon ring with a nitrogen atom linked by a single bond followed by a double bond indicates mutagenicity. *Progol* has therefore identified a new structural feature that is an alert for mutagenicity.

Some Other Applications of ILP
Biological classification of river water quality. River water quality can be monitored and assessed by observing various biological species present in the river. In particular, the riverbed macro-invertebrates are considered to be suitable indicators of the quality of water. Different species have different sensitivity to pollutants, and therefore the structure of the macro-invertebrate community in a river is well correlated with the degree and type of pollution. Dzeroski et al. [8] used ILP to analyze the relation between the samples of macro-invertebrates and the quality class of water. For learning, they used 292 field samples of

benthic communities taken from British Midlands rivers, classified by an expert river ecologist into five water-quality classes. They constructed a relational representation of these samples and used the ILP systems Golem [17] and CLAUDIEN [6] for inducing logic clauses from the data. The induced clauses were judged by experts to be intuitively appealing and largely consistent with their knowledge. In particular, the experts appreciated the symbolic explicitness of the generated descriptions. They considered this as a major advantage over neural net learning that was also applied to the same data.

Biomolecular modeling. ILP applications in biomolecular modeling aim to improve the understanding of the inter-relationships of chemical formula, three-dimensional structure, and function of molecules of biological importance. An overview of such applications of ILP can be found in [23]. These ILP applications involved applying Golem [17] to protein secondary structure prediction [14], prediction of ß-sheet topology and modeling the structure activity relationship (QSAR) of a series of drugs [11]. For secondary structure prediction Golem yielded predictive accuracies well in excess of any other contemporary approach. In the case of QSAR, predictive accuracies were not significantly higher than those produced by linear regression. However, in all three studies Golem discovered rules that provided insight into the stereochemistry of the system. Statistical techniques and neural networks are inferior in this aspect of comprehensibility, and are thus impaired for scientific discovery problems.

Inducing program invariants with ILP. In formally proving the correctness of procedural programs, one needs to find suitable conditions that always hold at given points in the program. Such a precondition has to be sufficiently strong to imply the postcondition of the program. Of particular interest is the problem of finding suitable conditions that are true inside program loops, called loop invariants. In general, the construction of loop invariants is considered difficult, and is usually done simply by guessing. Bratko and Grobelnik [3] explored the idea that ILP techniques can be used for automatically constructing loop invariants. A program that is to be proved correct can be executed, and the resulting execution traces can be used as learning examples for an ILP system. The states of the program variables at a given point in the program represent positive examples for the condition associated with that point in the program. Negative examples can be generated by employing a kind of "controlled closed-world assumption." Suitable loop invariants were straightforwardly induced for simple programs that are used in typical correctness proof exercises in [3]. The automatic induction of an invariant for a parallel program was also demonstrated. The scaling up of this approach to larger programs has not been investigated yet.

Data refinement in program design. In program construction from higher order specification, functions in the specification language (higher level) are to be implemented in the target language (lower level). Thereby abstract data types at the higher level are to be refined into concrete data types at the target language level. For example, sets can be reified into lists. In [3] this refinement problem is formulated in the ILP framework. As an illustration, the general ILP program Markus [9] was used to implement by induction the set union operation from abstract, high-level specification.

Innovative design from first principles. Bratko [2] formulated an approach to innovative design as an ILP problem. The design process is viewed as the process of structuring available elementary components in such a way that they together realize some specified target behavior. The approach addresses the design from "first principles" in the sense that the functional behavior of an artifact is derived from the *physics* of the elementary components available to the designer. The approach involves: specification of the target artifact by examples of its intended behavior, *qualitative physics* definition of the behavior of the elementary components available, and ILP as the mechanism for conceptually constructing the device. As an illustration, the Markus program [9] was applied to constructing simple electric circuits from examples of their intended behavior and the qualitative physics of some simple electrical components.

Qualitative system identification. A fundamental problem in the theory of dynamic systems is system identification. This can be defined as follows: given examples of the behavior of a dynamic system, find a model that explains these examples. Motivated by the hypothesis that it should be easier to learn qualitative than quantitative models, Bratko et al. [4] formulated the qualitative identification problem as an ILP problem. In their work, models are sets of Qualitative Differential Equations (QDEs) that constrain the values of the system variables. A Prolog implementation of QDE constraints normally used in qualitative physics is provided as background knowledge for learning. Example behaviors of the modeled system are used as positive training examples, while negative examples are generated as near misses. Models of simple dynamical systems have been induced using general ILP systems.

Conclusion

ILP has been applied to difficult, industrially or scientifically relevant and not yet satisfactorily solved problems. In the major applications described (mesh design, mutagenicity, river water quality), the results obtained with ILP using real industrial or environmental data are better than with any other known approach, with or without ML. Mesh design and mutagenicity are good examples of problems where relational background knowledge is most natural, and converting this into attribute-value form would be, if possible, awkward at best. In many of these applications, the users—domain specialists—are becoming increasingly interested in the understandability of the induced concept descriptions. This

helps them to obtain new insights in their problem domains. The representational flexibility in ILP helps the issue of understandability.

In all applications, general purpose ILP systems were used. Accordingly, a typical ILP application amounts to designing a good relational representation of the problem, including the definition of relevant background knowledge. A major strength of ILP systems, compared with other machine learning approaches, is that they accept background knowledge in forms as general as Prolog programs. This often completely changes the nature of representational engineering, which is essential part of machine learning applications as observed by Langley and Simon [12].

On the other hand, a major obstacle to more effective use of ILP at present is the relative inefficiency of the existing ILP systems, and their rather limited facilities for handling numerical data. Therefore, in problems for which attribute-value representation is adequate, attribute-based learning is more practical simply for efficiency reasons.

Acknowledgments

References

1. Bratko, I. Applications of machine learning: Towards knowledge synthesis. *New Generation Computing 11*, (1993), 343–360.
2. Bratko, I. Innovative design as learning from examples. In *Proceedings of the International Conference on Design to Manufacture in Modern Industries,* Bled, Slovenia, June 1993.
3. Bratko, I., Grobelnik, M. Inductive learning applied to program construction and verification. In *AI Techniques for Information Processing*, J. Cuena, Ed. North-Holland 1993. Also In *Proceedings of the ILP'93 Workshop,* Bled, Slovenia, April 1993.
4. Bratko, I., Muggleton, S., Varvsek, A. Learning qualitative models of dynamic systems. In *Inductive Logic Programming,* S. Muggleton, Ed. Academic Press, London, 1992.
5. Debnath, A.K., Lopez de Compadre, R.L., Debnath, G., Schusterman, A.J. and Hansch, C. *J. Medicinal Chemistry 34*, (1991), 786–797.
6. DeRaedt, L. and Bruynooghe, M. A theory of clausal discovery. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence.* Morgan Kaufmann, San Mateo, CA, 1993, pp. 1058–1063.
7. Dolvsak, B, Bratko, I., and Jezernik, A. Finite-element mesh design: An engineering domain for ILP application. In *Proceedings of the Fourth International Workshop on Inductive Logic Programming ILP-94,* Bad Honnef/Bonn, 1994.
8. Dvzeroski, S., DeHaspe, L., Ruck, B.M., and Walley, W.J. Classification of river water quality data using machine learning. In *Procedings of the Fifth International Conference on the Development and Application of Computer Techniques to Environmental Studies (ENVIROSOFT'94),* 1994.
9. Grobelnik, M. Markus: An optimized model inference system. In *Proceedings of the Logic Approaches to Machine Learning Workshop,* Vienna, August 1992.
10. Hansch, C., Malong, P.P., Fujita, T. and Muir, M. *Nature 194,* (1962), 178–180.
11. King, R.D., Muggleton, S., Lewis, R.A. and Sternberg, M.J.E. Drug design by machine learning: The use of inductive logic programming to model the structure-activity relationship of trimethoprim analogues binding to dihydrofolate reductase. In *Proceedings of the National Academy of Sciences 89,* (1992), 11322–11326.
12. Langley, P., Simon, H. Applications of machine learning and rule induction. *Commun. ACM 38,* 11(Nov. 1995).
13. Muggleton, S. Inductive logic programming. *New Generation Computing 8,* 4 (1991), 295–318.
14. Muggleton, S., Ed. *Inductive Logic Programming.* Academic Press, London, 1992.
15. Muggleton S. Bayesian inductive logic programming. In *Proceedings of the Eleventh International Conference on Machine Learning,* W. Cohen and H. Hirsh, Eds. (1994), pp. 371–379.
16. Muggleton, S., DeRaedt L. Inductive logic programming: Theory and methods. *J. Logic Programming 19,* 20 (1994), 629–679.
17. Muggleton, S. and Feng, C. Efficient induction of logic programs. In *Proceedings of the First Conference on Algorithmic Learning Theory,* Arikawa, S., Goto, S., Ohsuga, S. and Yokomori, T., Eds. Japanese Society for Artificial Intelligence, Tokyo, 1990, pp. 368–381.
18. Muggleton, S., King, R.D. and Sternberg, M.J.E. Protein secondary structure prediction using logic. *Prot. Eng. 5,* 7 (1992), 647–657.
19. Quinlan, J.R. Learning logical definitions from relations. *Machine Learning 5,* (1990), 239–266.
20. Quinlan, J.R. *C4.5: Programs for Machine Learning.* Morgan Kaufmann, San Francisco, 1993.
21. Shapiro, E. *Algorithmic Program Debugging.* MIT Press, Cambridge, MA, 1983.
22. Srinivasan, A., Muggleton, S.H., King, R.D., Sternberg, M.J.E. Mutagenesis: ILP experiments in a non-determinate biological domain. In *Proceedings of the Fourth International Workshop on Inductive Logic Programming ILP-94,* Bad Honnef/Bonn, 1994.
23. Sternberg M., King R., Lewis R., Muggleton S. Application of machine learning to structural molecular biology. *Philosophical Transactions of the Royal Society B, 344* (1994) 365–731.
24. Villemin, D., Cherqaoui, D. and Cense J.M. *J. Chim. Phys. 90* (1993), 1505–1519.

About the Authors:

IVAN BRATKO is a professor in the Faculty of Electrical Engineering and Computer Science, Ljubljana University, Slovenia. Current research interests in machine learning have been in learning from noisy data, combining learning and qualitative reasoning, and various applications of machine learning and Inductive Logic Programming, including medicine and control of dynamic systems. **Author's Present Address:** Department of EE/CS, Ljubljana University, Trzaska 25 61000 Ljubljana, Slovenia; email: Ivan.Bratko@fer.uni-lj.si

STEPHEN MUGGLETON is EPSRC Advanced Research Fellow at Oxford University Computing Laboratory. Current research interests include Inductive Logic Programming, learnability theory, Bayesian statistics, and computational biology. **Author's Present Address:** Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD UK; email: Stephen.Muggleton@comlab.oxford.ac.uk