Dynamic–Bias Induction

Daniel Oblinger **Gerald DeJong**
Beckman Institute and Department of Computer Science
University of Illinois
Urbana, Il 61801
{oblinger, dejong}@cs.uiuc.edu

Abstract

A tradeoff exists between the range of learning tasks solved by an induction system, and its performance on those tasks. We propose dynamic–bias induction, an approach in which bias is dynamically constructed as a function of the learning task. This admits the possibility of a high performance inductive learner that applies to a wide range of learning tasks. We assess the benefits and limitations of dynamic–bias induction by comparing an implementation of the approach to two existing inductive learning systems.

1. INTRODUCTION

Learning performance is greatly influenced by the strength of the learner's bias. Strong bias reduces the complexity of induction by reducing the number of alternative hypotheses considered. Bias that applies to a narrow class of learning tasks is free to include a narrow space of inductive hypotheses. Bias associated with a narrow space of hypotheses is, by definition, strong bias, and thus results in high learning performance.

In addition to high learning performance, it is important for an inductive learner to apply to a broad class of learning tasks. To be an effective tool for use in a given domain, an induction system must be capable of solving a range of learning tasks within that domain. This requirement of broad applicability conflicts with the task specificity suggested as a means of improving learning performance. How can an inductive learning system simultaneously possess the improved learning performance of a task–specific bias while retaining the broad applicability of less task–specific bias? We propose dynamic–bias induction as a solution to this dilemma. In this solution bias is constructed dynamically as a function of the particular learning task at hand. This solution rests on a new form of information that we refer to as *relevance knowledge*. Relevance knowledge, when combined with the problem specification for a particular learning task, yields a strong, task–specific bias for that learning task. Given the description of a different learning task, the *same* relevance knowledge can yield a *different* bias specialized to that task.

This approach follows in spirit work done by Stuart Russell on deriving bias from "high–level regularities in the world" [Russell86, Russell87]. The particular class of regularities considered by a learner has a profound effect on the concepts that may be entertained and thus learned. The remainder of this abstract discusses requirements on dynamic–bias induction to motivate the particular high–level regularity we consider—relevance knowledge.

Additional requirements are needed to ensure that dynamic–bias induction is a practical tool for inductive learning. An obvious solution to the problem of broadly applicable, high performance induction is to encode a task–specific bias for each possible learning task. This set, when indexed by a particular learning task description, yields a strong bias encoded specifically for that task. Though such relevance knowledge solves the problem as stated, it is not a practical solution. Such a representation requires a bias for each possible learning task to be considered and encoded independently. Information expressed in the bias for one of these tasks cannot be applied to any of the other learning tasks.

The failure of the solution proposed above stems from its violation of additional constraints on relevance knowledge that ensure the practicality of dynamic–bias induction. We have identified three criteria for relevance knowledge. Relevance knowledge must be: compact, general, and available.

The solution proposed above fails because its relevance knowledge is not compact. In that solution, information regarding the bias for one learning task cannot be shared by other tasks, no matter how related the tasks are. Such sharing is important when biases for many learning tasks are derived from some common body of knowledge.

Relevance knowledge must be general. Though the goal of dynamic–bias induction is to capture task–specific bias, it is important that the relevance knowledge itself is independent of any particular learning task. The task dependance of the generated biases arises from the interaction of the particular problem specification and the general relevance knowledge. This independence from particular tasks is necessary if dynamic–bias induction is to be a practical approach to the problem of broadly applicable, high performance induction.

Relevance knowledge must be based on information that one can reasonably expect to be available. Our approach of providing the learner with specific knowledge about the task at hand is useless if the knowledge required is no more accessible to the human user than the definition of the target concept itself. Thus, relevance knowledge must be a form of knowledge that is both available prior to learning and detailed enough to yield strong, task–specific, bias.

2. MODEL OF DYNAMIC–BIAS INDUCTION

Figure 1 is a block diagram of our model of dynamic–bias induction. The shaded region corresponds to conventional (statically biased) induction (i.e., ID3 [Quinlan86], AQ11 [Dietterich82], FOIL [Quinlan90], version space induction, connectionist networks, etc.). The *training instances* in the figure are represented as instance attributes with their associated values, and a classification for each training instance. Instance attributes are referred to as domain features, and the classification of an instance is simply its value for the goal feature. The *concept definition* is an expression that predicts the goal features as a function of the remaining features. The *inductive bias* is a (possibly infinite) structured set of candidate target concept definitions. The unshaded portion of Figure 1 represents the enhancements of dynamic–bias induction. The *learning goal* is the name of the target concept to be induced. The *relevance knowledge* is some fixed source of domain knowledge used to dynamically construct biases.

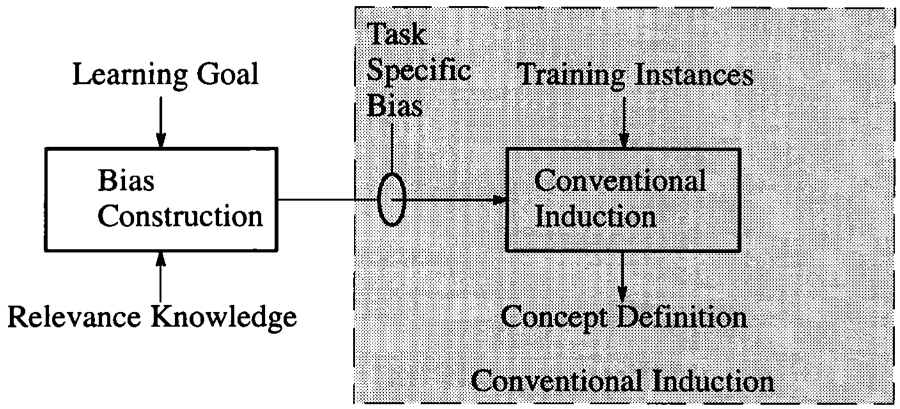

Figure 1: A model of dynamic–bias induction.

3. PROPERTIES OF RELEVANCE KNOWLEDGE

The use of relevance knowledge in our model of dynamic–bias induction, and the principles enumerated in the introduction both constrain relevance knowledge. We present these constraints as a set of properties on relevance knowledge. These properties are used in turn as the basis for the formalization of dynamic–bias induction.

According to our model in Figure 1, relevance knowledge gives rise to task–specific hypotheses. Many practical applications of induction are situated in domains where a wealth of knowledge regarding the domain features is available. Since the hypotheses are just particular combinations of domain features this knowledge can potentially provide strong guidance. We denote some of this knowledge using a relation over the domain features. Intuitively, this relation denotes the "relevance" of some features for the prediction of other features. As an example, consider the relevance of the feature *is–a–bird* for the prediction of the *can–fly* feature. Using the value of *is–a–bird* along with the values of other features like *is–dead*, and *has–wings*, the value of *can–fly* can be predicted, thus *is–a–bird* is relevant to *can–fly*. Since the aim of induction is the prediction of the target feature, the relevance of particular features for the prediction of other features can guide the task–specific choice of inductive bias. We refer to a declaration of the relevance of one feature to another as a *relevance link*. We use the '$\succ$' symbol to denote a relevance link. The relevant feature occurs on the left hand side, and the predicted feature is placed on the right. The relevance link above is written as: *is–a–bird*$\succ$*can–fly*. In general, $a \succ b$ iff there exists some concise concept definition of b that uses a.[1]

Like Russell's determination, the relevance link expresses generalities about the world used to constrain induction. Because the knowledge available prior to learning is likely to be incomplete/imprecise both the determination and relevance link are weaker than the corresponding deductive implication. Indeed the relevance link can be seen as a further step in the same direction, since it is weaker than the determination. If some predicate a determines a predicate b then it must be the case that a is relevant for the prediction of b, i.e. a is part of some concept definition for b. Russell's example determination "[a person's] Nationality determines [their primary] Language" implies the corresponding relevance link: "Nationality is relevant to Language," The converse, however, does not hold. Given some predicate a relevant to the prediction of b does not guarantee that a alone determines b. For example, *Gas–pedal–position* is relevant to determining a *car–velocity*, but it does not determine *car–velocity*. Many other features must also be included in any determination of the *car–velocity*, current gear, clutch, and break settings, gas level, road inclination, etc. Indeed it is precisely this ability to specify the relationship between a small number of concepts that facilitates the encoding of knowledge in our formalism.

According to the criteria in the introduction, the representation of relevance knowledge must be compact. We achieve this compactness by sharing information about the bias (relevant features) for different learning tasks within a *domain*, where a domain is simply a group of related learning tasks. Relevance knowledge expresses systematic properties at the domain level, *not* the task level. The relevance of some feature α for the predication of another feature β is domain level knowledge, since it potentially impacts many tasks within that domain. For each learning task with goal feature δ, where $\beta \succ \delta$, $\alpha \succ \beta$ impacts that task since, $\alpha \succ \beta$, and $\beta \succ \delta$ implies $\alpha \succ \delta$. This transitivity of the relevance relation permits the compact representation of relevant features for tasks within a domain by encoding them as a network of domain–specific (but task–independent) relevance knowledge.

Dynamic–bias induction uses specialized domain knowledge as a source of strong, task–specific bias for induction. The next properties of relevance knowledge enable it to express two commonly available types of specialized knowledge.

One commonly available type of specialized knowledge indicates whether some feature is positively or negatively correlated with another feature. A feature is positively correlated with another fea-

1. A concept definition c is a "concise" concept definition for b if there does not exists some definition c' for b where the predicates in c' are a proper subset of those in c.

ture if positive values for the first are associated with positive values for the second. So, *is–a–bird* is positively correlated with *can–fly*. A feature is negatively correlated with another, if positive values for the first are associated with negative values for the second. So, *has–hair* is negatively correlated with *can–fly*. The relevance link denotes a positive correlation between its left and right hand sides. A negative correlation is denoted by negating the left hand side of the relevance link; the example above is written: *¬has–hair≻can–fly*.

Another common type of specialized knowledge organizes features into conceptual units. These groups of features are used for predicting other features. The expression *is–speeding∧police–present* is a conceptual unit that is relevant to the feature *traffic–citation–received*. Such groups of features are treated as a unit when combined with other features. Treating groups of features as a unit greatly reduces the number of inductive hypotheses since the combined features are treated like a single feature.

We have a number of constraints on relevance knowledge derived from our model of dynamic–bias induction. We summarize these constraints as a set of properties of relevance knowledge:

1) Relevance knowledge is a *relation* over domain features.

2) Relevance knowledge captures *domain–level* knowledge in a transitive relation.

3) Relevance knowledge can express *correlation* information between features.

4) Relevance knowledge can express *groups* of features to be treated as a unit.

4. RESULTS

In the extended version of this paper [Oblinger94] we use the properties enumerated here to motivate our syntactic encoding of relevance knowledge, and the formal specification of the bias (set of inductive hypotheses) derived from this knowledge. The use of dynamic–bias induction is demonstrated by encoding a simple domain as a set of relevance links. A pair of experiments are performed using learning tasks within this domain. These experiments empirically measure the effectiveness of dynamic–bias induction. In the first experiment, we show the benefit of dynamically constructing bias by comparing our approach to a statically biased induction system. For this experiment FOIL [Quinlan90] is given an optimal static bias for a class of problems. Its performance is compared with an implementation of dynamic–bias induction, given relevance knowledge for the same task. In the second experiment our approach compares quite favorably to FOCL [Pizzani91], another induction system that dynamically derives its bias from a background knowledge. This comparison provides some sense for the constraint relevance knowledge places on induction, relative to other forms of knowledge also used to dynamically derive bias for induction.

References

[Dietterich82] T. G. Dietterich, B. London, K. Clarkson and G. Dromney, "Learning and Inductive Inference,"in *The Handbook of Artificial Intelligence, Vol. III*, P. R. Cohen & E. A. Feigenbaum (ed.), William Kaufman, Inc., Los Altos, CA, 1982.

[Pazzani91] M. J. Pazzani, "A Knowledge–Intensive Approach to Learning Relational Concepts," *Proceedings of the Eight International Workshop on Machine Learning*, Chicago, Ill, 1991, pp. 432–436.

[Oblinger94] D. Oblinger, G. DeJong, "Dynamic Bias Induction", Technical Report UIUC–BI–AI–94–03, Beckman Institute, Urbana, Ill, 1994

[Quinlan86] J. R. Quinlan, "Induction of Decision Trees," *Machine Learning 1*, 1 (1986), pp. 81–106.

[Quinlan90] J. R. Quinlan, "Learning Logical Definitions From Relations," *Machine Learning 5*, (1990), pp. 239–266.

[Russell86] S. Russell, "Preliminary Steps Towards Automation of Induction", *Proceedings, Fifth National Conference on Artificial Intelligence*, (1986), pp. 477–484

[Russell87] S. Russell, "A Declarative Approach to Bias in Concept Learning", *Proceedings, Sixth National Conference on Artificial Intelligence*, (1987), pp. 505–510