

# Probabilistic Logic Learning via Active Advice Seeking

Phillip Odom and Sriraam Natarajan  
Indiana University, Bloomington Indiana

## Abstract

Machine learning approaches that utilize human experts combine domain experience with data to generate novel knowledge. Unfortunately, most methods either provide only a limited form of communication with the human expert and/or are overly reliant on the human expert to specify their knowledge upfront. Thus, the expert is unable to understand what the system could learn without their involvement. Allowing the learning algorithm to query the human expert in the most useful areas takes full advantage of the data as well as the expert. We introduce *active advice-seeking* for relational domains. Relational logic allows for compact, but expressive interaction between the human expert and the learning algorithm. We demonstrate our algorithm empirically on several standard relational datasets.

## 1 Introduction

Probabilistic logic models (PLMs) [Getoor and Taskar, 2007; De Raedt *et al.*, 2008] combine the expressive power of first-order logic and the ability of probability theory to model noise and uncertainty. They have been inspired by databases [Friedman *et al.*, 1999; Heckerman *et al.*, 2004] and by logic [De Raedt *et al.*, 2007; Domingos and Lowd, 2009]. Given their expressivity, several powerful learning algorithms have been developed that allow for learning from interpretations [Domingos and Lowd, 2009; Natarajan *et al.*, 2008] and learning from entailment [Sato and Kameya, 1997; De Raedt *et al.*, 2007]. While efficient algorithms have been developed to learn the parameters of these models (either weights or probabilities), full model-learning (also called *structure learning* to denote learning of the logical structure) remains a challenging task. Recently, methods based on ensemble learning have been proposed that allow for efficient structure learning for PLMs [Natarajan *et al.*, 2015].

These methods essentially rely only on data. Given that the primary assumption is that data can be noisy, restricting humans to be mere labelers of the data, as is done in many popular approaches, seems wasteful. Recently, a formulation for incorporating prior knowledge as *preferences* over labels for the ensemble learning method was proposed [Odom *et*

*al.*, 2015]. The key idea was to explicitly trade-off between the label preferences suggested by the human expert and the posterior label distributions obtained from the data. It was demonstrated that advice was particularly useful where there was targeted noise - e.g., missing certain regions while segmenting, missing stop signs in some demonstrations etc.

While the framework of Odom *et al.* [2015] does not merely treat the given advice as “prior” knowledge, it assumes that all the advice is provided up-front before the learning takes place. Not only is this a potentially time consuming task for the experts, but it is also highly likely that they, not being experts in machine learning or probabilistic logic, would find it difficult to identify the domain knowledge that might be optimal for the learning algorithm. Hence, inspired by active learning [Settles, 2012], we propose *active advice-seeking* that aims to determine the regions of (relational/logical) feature space that is ideal for obtaining advice. For instance, will the accuracy of a model learned to predict heart attacks be higher if advice is given about the population who is overweight and has high blood pressure or about the population which smokes but exercises regularly? The answer is not clear but this is where active advice-seeking should be helpful. The goal of active advice-seeking is to lessen the responsibility of the expert both in terms of the effort that must be spent in specifying the advice, as well as the necessity that the expert understand the intricacies of the algorithm. It will automatically identify the regions of the feature space where the advice will be useful.

More precisely, the proposed algorithm presents a set of conjunctions of predicates as queries to the expert. The size of the set is pre-determined by a budget given by the expert (i.e., the algorithm and the expert agree in advance for the number of allowable queries). In order to compute the clause that should be queried, the algorithm learns a model from only the data to compute a score for each example, then it uses a regression clause learner to fit the scores. The best clause is presented to the expert who provides a preference over the labels. For instance, in a university domain, the body of the clause could be of the form  $prof(X) \wedge student(Y) \wedge paper(P,X) \wedge paper(P,Y)$ . The expert could then prefer the label to be  $advisedBy(X, Y)$ . Essentially, the system is asking the expert: what is your choice of label when a student and professor co-author the same paper? The expert replies saying, I prefer the student to be advised by the professor. Note that

this is a “soft” preference in that this preference may not always hold. This preference is then explicitly weighed against the data while learning the model.

We make the following key contributions: first, we introduce the notion of advice-seeking to the probabilistic logic model (PLM) community (and the general AI community). Second, we adapt a recent successful knowledge-based probabilistic logic learning algorithm to seek advice from the human expert. Third, we present the first relational algorithm that can go beyond data and interactively solicit input from the expert. Finally, we demonstrate using experiments, that such an approach is robust in learning from noisy data.

## 2 Background

Techniques for incorporating expert knowledge into learning is a key precondition for any active advice-seeking approach to be successful. We aim to introduce a broad learning paradigm that can use any method that incorporates prior knowledge. To that effect, we cover one advice-based framework which we will use to empirically validate our approach.

### 2.1 Advice-based PLMs

While there have been many knowledge-based systems developed for propositional models [Towell and Shavlik, 1994; Fung *et al.*, 2002; Torrey *et al.*, 2005; Le *et al.*, 2006; Kunapuli *et al.*, 2010], work on probabilistic logic models (PLMs) has not progressed as far. In PLMs, the expert is typically used to define some prior structure that can either be used as the complete structure or locally refined.

Recently, Odom *et al.* [2015] introduced a knowledge-based PLM method that learns seamlessly from data and any expert knowledge. While making use of Relational Function Gradient-Boosting (RFGB) to learn the structure and parameters of the model simultaneously [Natarajan *et al.*, 2012], they incorporate expert preferences which guide the structure and parameters to more robust models.

Extending previous work that considered knowledge as propositional Horn clauses [Towell and Shavlik, 1994; Fung *et al.*, 2002; Kunapuli *et al.*, 2013], they considered their advice as first-order logic Horn clauses. Thereby, allowing experts to give advice over different granularities of examples. The body of the clause specifies the examples over which the expert would like to give advice, while the head of the clause gives the preferred and avoided labels. For example, a cardiologist might suggest that patients whose close relatives had heart problems are more likely to have a heart problem.

Odom *et al.* [2015] incorporate this expert knowledge into RFGB [Natarajan *et al.*, 2012] which learns a series of relational regression trees [Blockeel, 1999]. Functional gradient-boosting aims to capture the error in the current model in a regression tree and then adds this regression tree to the model.

The gradients used by Odom *et al.* [2015] incorporate an additional term in the optimization function that pushes the model in the direction of the expert advice (represented by  $n_t$  and  $n_f$ -the number of advice which say that example  $x_i$

should be preferred/avoided)<sup>1</sup>

$$\Delta(x_i) = \alpha \cdot (I(y_i) - P(y_i; \psi)) + (1 - \alpha) \cdot [n_t(x_i) - n_f(x_i)]$$

While this approach has shown positive results in several difficult tasks, it still requires the expert to specify *all of the advice in advance*. Given a particular dataset, deciding the most useful advice is not a trivial problem. This problem is exacerbated by the fact that the expert has no expertise in machine learning. Active advice-seeking aims to alleviate this issue by querying the expert directly, using the training data as a guide to select the most useful queries. Previous work on active advice-seeking is limited to propositional queries in sequential decision making problems [Odom and Natarajan, 2016]. Grouping ground states into queries allowed the proposition algorithm to maximize the impact of the human expert. However, lifting advice to be relational is a more powerful and principled approach.

### 2.2 Active Learning

Active learning is a related research problem where the goal is to make use of an expert that can provide the labels of examples [Settles, 2012]. Pool-based active learning approaches assume a pool of unlabeled examples from which the learning algorithm should choose. In active advice-seeking, this pool of examples is the training set. While there are labels in the training set, it is assumed that either there are not sufficient training data (and thus there is missing knowledge) or the training data is noisy and so the labels should not be fully trusted. So while active learning aims for finding the labels of the examples, we are soliciting advice.

Most active learning methods repeat these steps:

1. Learn a model from training data
2. Compute uncertainty over unlabeled data
3. Select examples based on uncertainty and solicit label
4. Add labeled examples to training set

The process begins by learning a model with the current set of labeled data. This model is then used to compute some measure of uncertainty (this could be entropy, KL-divergence or other measures) that suggests how likely the model would correctly predict the unlabeled examples. Consider a simple, linear classifier with two possible unlabeled examples, one located close to the decision boundary with the other located far from the boundary. The example close to the decision boundary is more likely to effect the decision boundary and would be selected for labeling.

This cycle accumulates the best examples to label at each step and has been shown to be effective especially in domains where there is a dearth of data available. However, labeling individual examples is not an effective use of human experts availability. Allowing expert’s to give advice results in the expert being able to select the ideal granularity of advice (over a single example or many examples). Active advice-seeking aims to effectively use human experts by providing clauses instead of ground examples. Not only does this provide a specific granularity of advice, but it also provides a compact description of the most uncertain examples.

<sup>1</sup>Note the difference to standard (only data) RFGB which optimizes  $(I(y_i) - P(y_i; \psi))$ .

### 3 Relational Active Advice-Seeking

The aim of relational active advice-seeking is to offload the task of selecting areas of the feature space to give targeted advice from the human expert to the learning algorithm. In relational models, experts are often asked to define the logical structure of the model with the parameters learned from data. However, it is important to be able to learn the full model (structure and parameters) especially in complex, real world domains. Experts can still provide valuable input about targeted areas of the feature space. The wide variety of potential expert advice complicates the advice-giving process and can lead the expert to give correct, but not relevant advice.

Previous work on advice-giving requires significant effort on the part of the expert to determine the relevant advice [Odom *et al.*, 2015; Kunapuli *et al.*, 2013]. If the expert provides exhaustive advice, the learning algorithm will be able to learn a robust classifier. However, the experts time is often limited and only a few queries can be answered. These queries should not be redundant, focusing on areas that are well covered by the data. Instead, they should focus on areas where the learning algorithm cannot distinguish the correct label or behavior. Thus, we extend relational advice-taking methods to active advice-seeking.

### 4 Problem Formulation

The overall goal of our algorithm is to identify regions of the feature space that the agent is most uncertain about and query the expert for advice on these regions. In the propositional case, this was handled by clustering examples based on the distribution over the labels and querying the expert over this cluster [Odom and Natarajan, 2016]. However, this heuristic may not suffice for relational tasks since there are significantly more negative examples than positives. Fortunately, the use of a rich representation such as first-order logic naturally allows us to query over the most uncertain regions of the feature space.

We represent the regions of feature space as conjunctions of predicates. Intuitively, this corresponds to grouping examples such that a particular condition is satisfied. More precisely, the goal of our algorithm is to select a set of conjunctions of first-order logic atoms about which to query the expert. These queries concisely describe the set of training examples to which the advice will apply. In order to select relevant areas of the feature space, the algorithm learns a clause (model) based on scores of the given examples. The goal of this learned model is to group similar examples based on their assigned score which measures the importance of that query. Queries have low scores if the algorithm is confident in its prediction, Otherwise, the query will receive a high score, making it more likely to be selected by the active advice-seeking algorithm. We explain the clause generation later in this section. We will now formally define advice:

**Definition 1** A set of advice ( $A$ ) is defined as a series of relational queries ( $Q_i$ ) and the experts corresponding response ( $R_i$ ), ie. ( $A = \langle (Q_1, R_1), (Q_2, R_2), \dots, (Q_n, R_n) \rangle$ ).

The algorithm solicits a sequence of queries that depend on the scoring function that will be discussed in detail later.

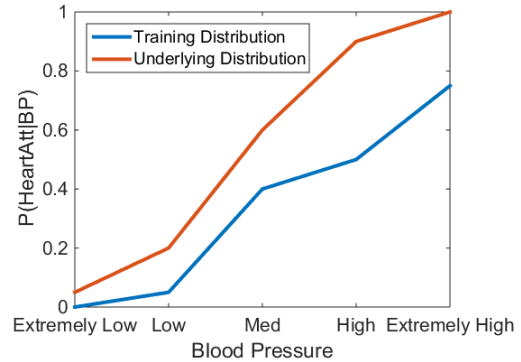


Figure 1: Example showing the distribution of heart attacks given blood pressure for an observed and underlying distribution.

The number of queries is dependent on the difficulty of the problem and the availability (query budget) of the expert.

**Definition 2** A Relational Query ( $Q$ ) is defined as a conjunction of literals ( $\wedge f_i$ ), which defines the set of examples to which the advice will be applied.  $Q$  will be shown to the human expert.

**Definition 3** An Expert Response ( $R$ ) is defined as a set of preferred labels ( $l+$ ), and a set of avoided label ( $l-$ ) given with respect to a relational query. Note that both  $l+$  /  $l-$  could be empty if the expert does not understand  $Q$  or if the query does not separate different classes.

If the expert is not satisfied with the query - possibly because the query does not properly delineate between labels - then the expert can provide no preferred or avoided labels. Such a query is not useful to the learning algorithm and squanders the time of the expert. The relational query and its accompanied response represent a single piece of advice that can be utilized by the knowledge-based learning algorithm. We now present an illustrative example before discussing the algorithm in detail.

#### 4.1 Illustrative Example

Consider the example of heart attack prediction given clinical information about the patients such as their blood pressure. The training set (e.g. one particular county in Wisconsin) might show patients with high (but not extreme) blood pressure having a relatively low incidence of heart attacks. This systematic difference could be attributed to local factors. The counties distribution and the true underlying distribution are shown in Figure 1.

Now consider soliciting advice about heart attacks and blood pressure from a cardiologist in California. Being unfamiliar with Wisconsin, the cardiologist might give broad, straight-forward advice. However, such knowledge might already follow from the training data. Examples of such advice include “extremely high blood pressure leads to heart attacks” and “heart attacks are not likely with low blood pressure”. While these pieces of advice are valid, they are not the most relevant advice for this particular learning problem.

If the algorithm had the ability to solicit advice, then it could direct the expert to give the most relevant advice at any point. Our proposed algorithm will identify areas in the data that are unclear and will instead query the expert automatically with “How likely are heart attacks when the blood pressure is high, but not extreme”. This is likely the most useful advice given the data. This approach not only benefits the learning algorithm, but reduces the burden on the expert who is only required to answer specific questions.

---

**Algorithm 1** Actively Seeking Advice for PLMs (ASAPlm)

---

```

function ASAPLM( $D, E, MaxQuery$ )
2:    $A = \emptyset$ 
    $M = \text{RFGB}(D)$   $\triangleright$  Model from Noisy Data
4:   for  $x_i \in D$  do  $\triangleright$  Compute Uncertainty per Example
        $R(x_i) = H(x_i)$ 
6:   end for
    $AQ = \text{LRC}(D, R)$   $\triangleright$  Learn Regression Clauses
8:   for  $i = 1$  to  $MaxQuery$  do  $\triangleright$  Query Expert
        $AQ_q = \text{MAXSCORE}(AQ)$ 
10:   $AQ = AQ - AQ_q$ 
        $Q = \text{QUERY}(E, AQ_q)$ 
12:   $A = A \cup Q$ 
   end for
14:   $M_F = \text{ADVLEARNER}(A, D)$   $\triangleright$  Learn with Advice
   return  $M_F$ 
16: end function

```

---

## 4.2 The Algorithm

Our proposed approach involves generating a set of queries, scoring those queries to rank them according to their usefulness, and finally soliciting the most useful queries to the human expert. The number of queries that can be requested depends on the problem (more difficult domains require more knowledge) and the availability of the human expert. The complete active advice seeking algorithm (ASAPlm) is shown in Algorithm 1. We will address each of these vital components in turn.

### Generating and Scoring Queries

Recall that in standard active learning, a model is learned from labeled data and using this model, some uncertainty measure is calculated to identify the most uncertain unlabeled example to query the expert. We take a similar approach with an important change. We learn a model using RFGB on the noisy data (line 3 of the algorithm) and compute the entropy over the examples given this model (lines 4-6). Following active learning, we define the score of an example as the entropy of the model’s prediction (line 5 of Algorithm 1), ie,

$$H(x_i) = \sum_{l \in \text{Labels}} P_l(y_i|x_i) \log(P_l(y_i|x_i))$$

where  $P(y_i|x_i)$  is learned using RFGB. Such uncertainty measures have performed extremely well in many active learning methods and similar results can be shown over relational data. The key difference is that the uncertainty is based on all of the training examples that satisfy the query. In our

empirical evaluation, we focus on entropy as our uncertainty measure. However, the framework is broad and allows for the selection of the most appropriate uncertainty function for the problem at hand.

Then these scores are used as regression values for the corresponding example and a set of weighted first-order-logic clauses are learned that can potentially group these examples (line 7, function LRC). These clauses are presented to the expert according to the learned weights. We learn these weighted clauses through an adaptation of RFGB where instead of learning  $P(y_i|x_i)$ , we want to learn a model for the uncertainty values of  $x_i$  (by fitting regression trees). The key intuition is that the regression trees find clauses that apply to examples with similar uncertainties. Note that unlike in discriminative learning where there are positive and negative examples, regression does not treat positive and negative examples differently. Every example has a uncertainty value and regression is just trying to fit those values. The learned clauses represent a set of possible queries from which the algorithm can select.

### Querying the Expert

After the queries have been generated and ranked, they can be used to solicit advice from the human expert. For a given relational query, the expert should supply the suggested preferred labels (should be considered more likely) and the avoided labels (should be considered less likely). Alternatively, the expert could decline to answer if the query is too general or incomprehensible. Declining is an indication that the active advice seeking algorithm is not selecting appropriate queries.

### Advice-based Learner

Given the advice, the final step is to utilize the advice-based learner to learn from both the training data as well as the expert advice. An ideal algorithm should trade-off between the sources of knowledge when they offer contradictory information. For the purposes of empirical validation, we utilize KB-PLL as our advice-based learner. It combines the target distribution of the training data and the distribution suggested by the advice to find a robust model (refer to section 2).

Overall, the proposed approach to active advice-seeking aims to effectively utilize the human expert by generating queries. These queries are targeted based on the perceived weaknesses in the training data. We now thoroughly investigate ASAPlm.

## 5 Experiments

Through our experiments, we aim to answer the following questions:

- Q1:** Does active advice-seeking result in more effective learning?
- Q2:** Is our algorithm robust to both random and systematic noise?
- Q3:** Is advice an efficient form of communication between algorithm and expert?

**METHODS:** We compare our method against two baselines. To evaluate our query generation method, we compare against learning with randomly generated queries (*Random Queries*).

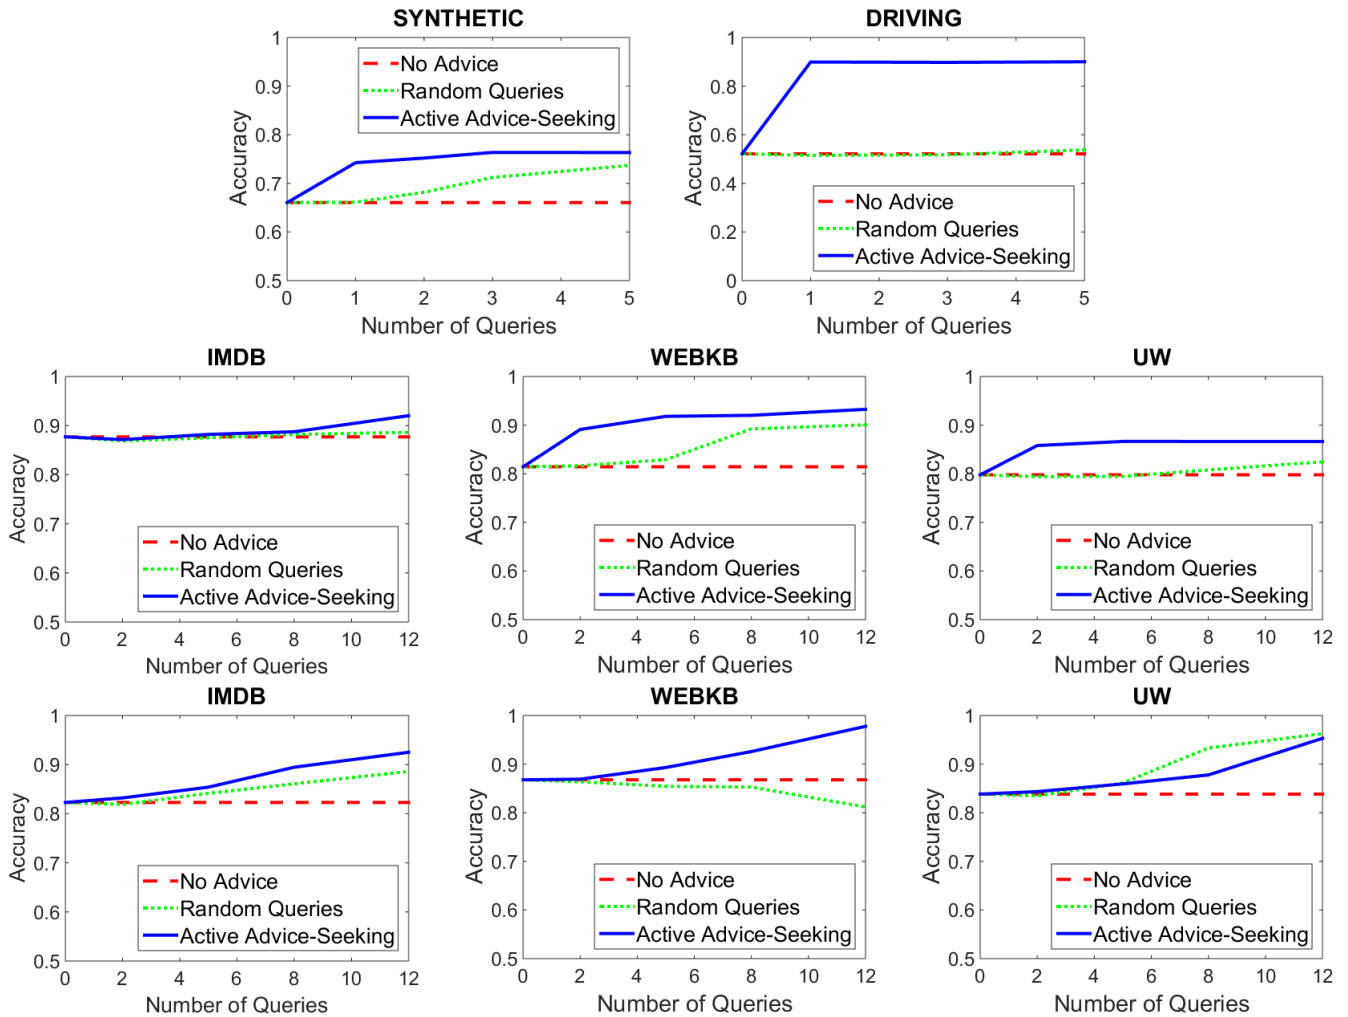


Figure 2: The learning curves from the various experimental domains. The top 5 results represent systematically noisy datasets, while the bottom 3 results represent randomly noisy datasets. Each learning curve shows accuracy as the number of queries to the expert increases. We compare Active Advice-Seeking to Random Queries and No Advice.

To evaluate the effectiveness of active advice-seeking, we compare against learning with no advice (*No Advice*). We also discuss the quality of the advice that is generated in each domain. Given our experience with the domains, we take the role of expert to answer the queries.

In all the experiments, we compare the accuracy of learned model. To show that our algorithm is capable of correcting noisy data, we added noise equal to 25% of the positive examples. Note that in the relational space the number of negative examples typically greatly outnumber the positive examples. This means that the impact of the noise is much less than 25%. To show that our algorithm is capable of correcting systematic noise, we label examples incorrectly in a targeted region of the feature space. The synthetic heart attack dataset and driving domain are tasks where systematic noise is natural. Heart problems effect different regions or ethnic groups in different ways. Many drivers consistently drive over the speed limit and roll through stop signs. For the remaining

datasets, imdb, webkb and uw, we have experimented with both systematic and noisy data. Each randomly noisy experimental domain has either 4 or 5 folds and we randomly add noise 5 times for each fold. Each systematically noisy experiment generated data for each fold or was repeated 5 times. For our relational advice-based learning algorithm, we use KBPLL [Odom *et al.*, 2015] with  $\alpha = 0.25$ .

**DOMAINS:** We have a variety of standard relational datasets as well as an imitation learning dataset focused on driving.

**IMDB:** This dataset is a movie database that consists of movies, actors, directors and their various genres. Our goal is to predict the worked under relationship (ie which actors worked on movies under a particular director). This dataset consist of 5 folds.

**WEBKB:** This dataset is a university dataset that consists of webpages and their hyperlinks. Our goal in this domain is to predict which webpage belongs to a faculty member based on the webpages and their linking structure. This dataset has 5

Domain	Query Generated
Driving	What if there is a car in the left lane?
Synthetic	What if a person has medium to high blood pressure?
IMDB	Do female actors work under people in crime movies?
WEBKB	What is the title of students working on projects?
UW	What is the relationship between students and TA's?

Table 1: The top queries generated in each domain for the systematically noisy datasets. Experts respond to these queries by providing  $l + /l-$  in each domain.

folders.

**UW:** This dataset is a university dataset that consists of professors, students, courses, and publications each having various relationships and features. Our goal is to predict the advisedby relationship. This dataset has 4 folds.

**SYNTHETIC:** The goal of the synthetic dataset, from the illustrative example, is to predict heart attacks given the blood pressure. There is a systematic difference (see Figure 1) between the training set and the testing set. This dataset was generated 5 independent times.

**DRIVING:** The driving domain focuses on navigating down a 5-lane highway, avoiding the other cars on the road [Judah *et al.*, 2014]. The possible actions are to stay in the current lane or change lanes to the left or right. The size of the training set and testing set are 100 trajectories consisting of 10000 total training examples.

### 5.1 Systematic Noise

The results with systematic noise (Figure 2, top 2 rows) are shown for the synthetic and driving domains as well as each of the standard relational datasets. Together they show the power of our proposed approach when dealing with systematic noise. In most datasets, the algorithm is capable of selecting useful queries immediately, providing significant impact. Random queries demonstrate gradual performance gains in the synthetic and webkb domains, but fail to have a positive effect on the other domains. While random queries do not cause performance to degrade, they have an extremely difficult time isolating systematic noise especially when there are more features. A key reason there is very little change in these domains is that the queries generated were ambiguous and useful only for a few examples. For instance, a common query in the driving domain is “What action should I take if there is a car both to my left, right, AND in front”. While this is a possible scenario, it is not likely in this dataset and there is no obvious advice to give for these states. Alternatively, the queries generated from the active advice-seeking algorithm select more relevant and overall useful queries. Thus, **Q1** is answered affirmatively that our proposed approach is able to learn effectively in the presence of systematic noise.

### 5.2 Random Noise

The standard relational domains (Figure 2, bottom row) are used to show that even when noise is random, our proposed method can still generate high-quality queries to the expert. Random noise should be more difficult for our algorithm, as there may not be specific regions of the feature space that need attention. However, across all three domains, our pro-

posed approach achieves consistent success, generating performance gains with each query. In contrast, randomly generated queries can yield positive performance (as in imdb or uw), or actually result in a model that is worse than relying on the data (as in webkb). It may seem counter intuitive for advice to be harmful. However, consider the query “Is a student advised by a professor”. While it may seem that the advice should be that student are advised by professors, there are many student and many professors. Therefore, such an advice could result in many false positives as a student is not advised by most professors. Thus, our proposed approach is robust to random noise as well as systematic noise (**Q2**).

### 5.3 Quality of Advice

The preceding empirical results show that our proposed approach is able to generate relevant queries that yield significantly higher accuracy in nearly all of the domains for both systematic and noisy experiments. However, the interpretability of the queries is vital as the experts need to easily comprehend the queries in order to give the proper advice. Table 1 shows the top query generated for each domain (systematic noise). In the driving domain, the query asks what action to take when there is a car in the left lane. The expert response would be to stay in the current lane. As another example, in the uw domain, the query asks about the relationship between students and TA’s. While TA’s might help teach student, the advice would say that TA’s cannot advise students. The best queries are heavily influenced by the noise in the training set. Overall, the queries are concise (as shown in Table 1) and effective (as shown in the empirical validation). Thus, advice is an efficient form of communication (**Q3**).

## 6 Conclusion

We presented the first advice-seeking framework for PLMs. Our method, inspired by active learning, queries the expert with sub-spaces of the feature space where advice can be provided as preferences over labels. The key insight is that the learning algorithm can better query the expert based on the uncertainty in the data as compared to the expert providing all advice pieces in advance. Our experimental results across standard data sets proved that such a method is effective in soliciting useful advice. Evaluating on larger data sets such as electronic health records is an important future direction, as is exploring the different measures of uncertainty for grouping the different examples. Learning from multiple experts by weighing them explicitly is another direction that we will explore. Finally, performing user studies on more sophisticated test beds is an interesting research direction.

## References

- [Blockeel, 1999] Hendrik Blockeel. Top-down induction of first order logical decision trees. *AI Communications*, 12(1-2), 1999.
- [De Raedt *et al.*, 2007] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: A probabilistic prolog and tis application in link discovery. In *IJCAI*, 2007.
- [De Raedt *et al.*, 2008] Luc De Raedt, Paolo Frasconi, Kristian Kersting, and Stephen Muggleton. *Probabilistic inductive logic programming*. Springer, 2008.
- [Domingos and Lowd, 2009] Pedro Domingos and Daniel Lowd. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan & Claypool, 2009.
- [Friedman *et al.*, 1999] Nir Friedman, Lise Getoor, Daphne Koller, and Avi Pfeffer. Learning probabilistic relational models. In *IJCAI*, 1999.
- [Fung *et al.*, 2002] Glenn Fung, Olvi L. Mangasarian, and Jude W. Shavlik. Knowledge-Based support vector machine classifiers. In *NIPS*, pages 01–09, 2002.
- [Getoor and Taskar, 2007] Lise Getoor and Ben Taskar. *Introduction to statistical relational learning*. Cambridge: MIT Press, 2007.
- [Heckerman *et al.*, 2004] David Heckerman, Christopher Meek, and Daphne Koller. Probabilistic entity-relationship models, prms, and plate models. In *ICML*, 2004.
- [Judah *et al.*, 2014] Kshitij Judah, Alan Fern, Prasad Tadepalli, and Robby Goetschalckx. Imitation learning with demonstrations and shaping rewards. In *AAAI*, 2014.
- [Kunapuli *et al.*, 2010] Gautam Kunapuli, Kristin P. Bennett, Amina Shabbeer, Richard Maclin, and Jude W. Shavlik. Online knowledge-based support vector machines. In *ECML*, pages 145–161, 2010.
- [Kunapuli *et al.*, 2013] Gautam Kunapuli, Phillip Odom, Jude Shavlik, and Sriraam Natarajan. Guiding autonomous agents to better behaviors through human advice. In *ICDM*, 2013.
- [Le *et al.*, 2006] Quoc V. Le, Alex J. Smola, and Thomas Gärtner. Simpler knowledge-based support vector machines. In *ICML*, pages 521–528, 2006.
- [Natarajan *et al.*, 2008] Sriraam Natarajan, Prasad Tadepalli, Thomas Dietterich, and Alan Fern. Learning first-order probabilistic models with combining rules. *Annals of Mathematics and AI*, 54(1), 2008.
- [Natarajan *et al.*, 2012] Sriraam Natarajan, Tushar Khot, Kristian Kersting, Burnd Gutmann, and Jude Shavlik. Gradient-based boosting for statistical relational learning: The relational dependency network case. *Machine Learning*, 86(1), 2012.
- [Natarajan *et al.*, 2015] Sriraam Natarajan, Kristian Kersting, Tushar Khot, and Jude Shavlik. *Boosted Statistical Relational Learners: From Benchmarks to Data-Driven Medicine*. Springer, 2015.
- [Odom and Natarajan, 2016] Phillip Odom and Sriraam Natarajan. Active advice seeking for inverse reinforcement learning. In *AAMAS*, 2016.
- [Odom *et al.*, 2015] Phillip Odom, Tushar Khot, Reid Porter, and Sriraam Natarajan. Knowledge-based probabilistic logic learning. In *AAAI*, 2015.
- [Sato and Kameya, 1997] Taisuke Sato and Yoshitaka Kameya. Prism: A symbolic statistical modeling language. In *IJCAI*, 1997.
- [Settles, 2012] Burr Settles. *Active Learning*. Morgan & Claypool, 2012.
- [Torrey *et al.*, 2005] Lisa Torrey, Trevor Walker, Jude Shavlik, and Richard Maclin. Using advice to transfer knowledge acquired in one reinforcement learning task to another. In *ECML*, 2005.
- [Towell and Shavlik, 1994] Geoffrey Towell and Jude Shavlik. Knowledge-based artificial neural networks. *Artificial Intelligence*, 69:119–165, 1994.