Modeling the Effects of Toxins in Metabolic Networks

Abductive and Inductive Reasoning for Learning Models of Inhibition in Biological Networks

BY ALIREZA TAMADDONI-NEZHAD, RAPHAEL CHALEIL, ANTONIS C. KAKAS, MICHAEL STERNBERG, JEREMY NICHOLSON, AND STEPHEN MUGGLETON

Abduction and induction are two forms of reasoning that have been widely used in machine learning. The combination of abduction and induction has recently been explored from a number of angles [1]. Moreover, theoretical issues related to the completeness of this form of reasoning have also been discussed by various authors [2]–[4]. Some implemented systems have been developed for combining abduction and induction [5] and others have recently been proposed [6], [7]. There have also recently been demonstrations of the application of abduction/induction systems in the area of systems biology [8], [9].

The research reported in this article is being conducted as part of the MetaLog project [10], which aims to build causal models of the actions of toxins from empirical data in the form of nuclear magnetic resonance (NMR) data, together with information on networks of known metabolic reactions from the Kyoto Encyclopedia of Genes and Genomes (KEGG) [11]. The NMR spectra provide information concerning the flux of metabolite concentrations before, during, and after administration of a toxin.

In an article [12] describing our initial investigation in this topic, we modeled the initial effects of a single toxin (hydrazine). The initial model ignored the temporal variance of metabolite concentrations. By contrast, in [13] we describe an extended study in which temporal variation is captured and the resulting model for hydrazine is contrasted with that of a second liver toxin (anit). The NMR data for hydrazine and anit were the first datasets which haves been made available to the project by our collaborators who studied these toxins as part of the COMET project [14]. The present article summarizes these studies and focuses on the main results, with less technical detail. In addition, this article contains more analysis of the significance of the results from a biological perspective. Interested readers are referred to [13] for more information on the use of abduction and induction in these studies.

In our study, examples extracted from the NMR data consist of metabolite concentrations (up-down regulation patterns extracted from NMR spectra of urine from rats dosed with the toxin) at 8 hours, 24 hours, 48 hours, 72 hours, and 96 hours after the injection of the toxin. Background knowledge consists of known metabolic networks and enzymes known to be inhibited by the toxin. This background knowledge, which represents the present state of understanding, is incomplete. For example, for many inhibitors the available data is not enough to generate any general rule. In order to overcome this incompleteness, hypotheses are considered that consist of a mixture of specific inhibitions of enzymes (ground facts) together with general rules that predict classes of enzymes likely to be inhibited by the toxin (nonground). Hypotheses about inhibition are built using Progol5.0 [5] and predictive accuracy is assessed for both the ground and the nonground cases. It is shown that even with the restriction to ground hypotheses, predictive accuracy increases with the number of training examples and in all cases exceeds the default (majority class). Experimental results also suggest that when sufficient training data are provided, nonground hypotheses show a better predictive accuracy than ground hypotheses. These results are also evaluated in terms of new biological insight provided by the ground hypotheses.

Inhibition in Metabolic Networks

The processes that sustain living systems are based on chemical (biochemical) reactions. These reactions provide the requirements of mass and energy for the cellular processes to take place. The complex set of interconnected reactions taking place in a given organism constitute its *metabolic network* [15], [16]. Not all reactions take place at the same time in this network, and they need to be finely coordinated. Biochemical reactions are sped up by highly specialized proteins, the enzymes. Enzymes are the most efficient catalyzers known, and most of the reactions taking place in living organisms would be too slow without them to sustain life. Enzymes control the activation of different parts of the network and are therefore the main element for coordination of the different parts of the metabolic network [17].

The assembly of full metabolic networks, made possible by data accumulated through years of research, is now stored and organized on metabolic databases and allows their study from a network perspective [18], [19]. Even with the help of this new systems biology approach to metabolism, we are still far apart from understanding many of its properties. One of the less understood phenomena, especially from a network perspective, is *inhibition*. Some chemical compounds can affect enzymes, impeding them to carry out their functions, and

0739-5175/07/$25.00©2007IEEE

hence affecting the normal flux in the metabolic network, which is in turn reflected in the accumulation or depletion of certain metabolites.

Inhibition is important from the therapeutic point of view since many substances designed to be used as drugs against some diseases can eventually have an inhibitory side effect on other enzymes. For example, Paracetamol is an inhibitor COX-3 cyclo-oxygenase [20], preventing the formation of arachidonic acid into prostaglandins, which are involved in pain and fever process. Paracetamol is transformed by the cytochrome P450 dependent enzymes producing N-acetyl-p-benzo-quinone imine (NAPQI), which is normally conjugated with glutathione for renal elimination. In case of overdose the glutathione production pathways are not able to produce enough glutathione and NAPQI, which is very reactive, accumulates in the liver leading to its failure [21]. Hence, any system able to predict the inhibitory effect of substances on the metabolic network would be useful in assessing the potential harmful side effects of drugs.

Preparation of Data and Background Knowledge

In this work we use experimental data on the accumulation and depletion of metabolites to model the inhibitory effect of a toxin such as hydrazine ($NH_2 - NH_2$)in the metabolic network of rats. Figure 1 shows the metabolic pathways subnetwork of interest also indicating with "up" and "down" arrows the observed effects of the hydrazine on the concentration of some of the metabolites involved. This subnetwork was manually built from the information contained in the KEGG metabolic database [11]. Starting from the set of chemical compounds for which there is information on up/down regulation after toxin treatment coming from the NMR experiments, we tried to construct the minimal network representing the biochemical links among them by taking the minimum pathway between each pair of compounds and collapsing all those pathways together through the shared chemical compounds. When there is more than one pathway of similar length (alternative pathways), all of them are included. Pathways involving "promiscuous" compounds (compounds involved in many chemical reactions) are excluded. KEGG contains a static representation of the metabolic network (reactions connecting metabolites); where the existence of a reaction is only conditioned by the existence of at least one gene coding for an enzyme catalyzing the reaction. NMR data provide information on the concentrations of metabolites and their changes with time. The NMR data used in this study represent variations of concentration of the metabolites (relative to their concentration before injection of hydrazine) that are measured at 8 hours, 24 hours, 48 hours, 72 hours, and 96 hours. The effect of toxin on the concentrations of chemical compounds is

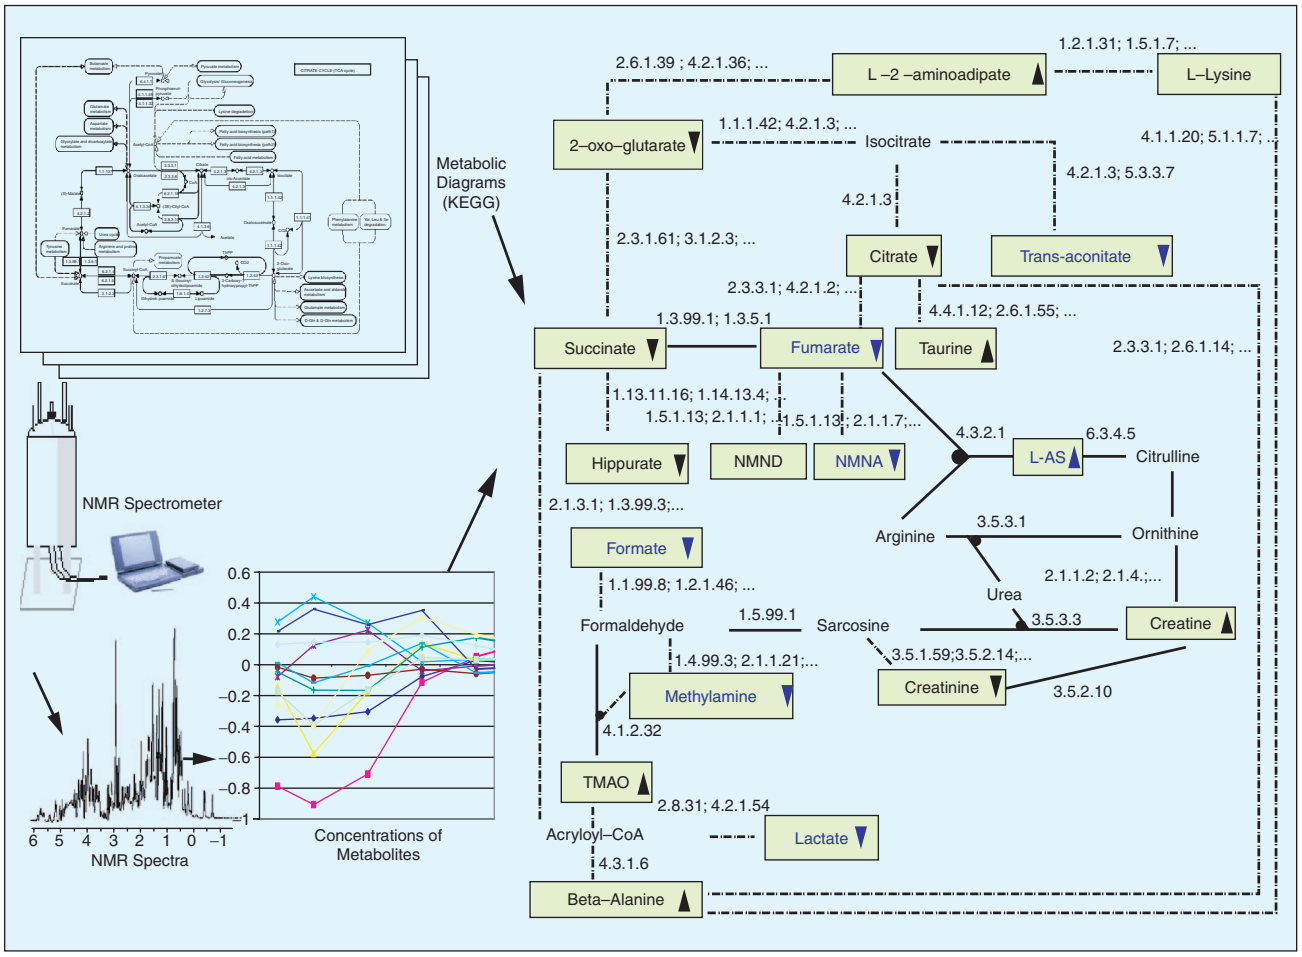

Fig. 1. A metabolic subnetwork involving metabolites affected by hydrazine. Information on up/down changes in metabolite concentrations from NMR spectra is combined with KEGG metabolic diagrams. The enzymes associated with a single reaction (solid line) or a linear pathway (dotted line) are shown as a single enzyme or a list of enzymes.

Inhibition is important from the therapeutic point of view since many substances designed to be used as drugs against some diseases can eventually have an inhibitory side effect on other enzymes.

coded in a binary way; i.e., only up/down changes (increasing/decreasing) in compound concentrations are incorporated in the model. In this subnetwork the relation between two compounds (edges in the network) can comprise a single chemical reaction (solid lines) or a linear pathway (dotted lines) of chemical reactions in the cases where the pathway between those compounds is composed by more than one reaction but not involving other compounds in the network (branching points). The directionality of the chemical reactions is not considered in this representation, and in fact it is left deliberately open. Although metabolic reactions flow in a certain direction under normal conditions, this may not be the case in "unusual" conditions like the one we are modeling here (inhibition). Inhibition of a given reaction causes the substrates to accumulate what may cause an upstream enzyme to start working backward in order to maintain its own substrate/product equilibrium.

The "one-to-many" relations (chemical reactions with more than one substrate or product) are indicated with a filled circle in Figure 1. The enzymes associated with the relations (single chemical reactions or linear pathways) are shown as a single enzyme or a list of enzymes.

Logical Modeling of Inhibition

Modelling a scientific domain is a continuous process of observing the phenomena, understanding these phenomena according to a currently chosen model, and using this understanding of an otherwise disperse collection of observations to improve the current general model of the domain. In this process of development of a scientific model one starts with a relatively simple model that becomes further improved and expanded as the process is iterated over. Any model of the phenomena at any stage of its development can be *incomplete* in its description. The task then is to use information given to us by experimental observations to improve and possibly complete this description. The development of our theories is then driven by the observations and the need for these theories to conform to the observations. Our approach will fall much in the same spirit of theories of scientific discovery [22], [23] in the sense that the development of a scientific theory is considered to be an incremental process of refinement guided strongly by the empirical observations.

Considering a logical approach to this problem of incremental development of a scientific model, philosophers of science have recognized the need to introduce new *synthetic* forms of reasoning, alongside with the analytical reasoning form of deduction. As early back as Aristotle we see two forms of synthetic logical reasoning: *abduction* and *induction*. In the 19th century, Charles Sanders Peirce [24] sets out clearly these three forms of syllogistic reasoning (deduction, abduction, and induction) and studies their respective role in the development of scientific theories. More recently, several authors (see, for example, [25], [26], and [1]) have studied abduction and induction from the perspective of artificial intelligence and cognitive science. In particular, the work in [1] is devoted to the problem of comparing these two forms of reasoning and investigating their possible unification or integration for the purposes of artificial intelligence.

Given a theory T that describes our current (incomplete) model of the scientific domain under investigation, and a set of (experimental) observations O, abduction and induction are employed in the process of assimilating in the current theory the new information contained in the observations. They both synthesize new knowledge H, thus extending the model T to $T \cup H$, according to the same formal specification of:

- $T \cup H| = O$, and
- $T \cup H$ is consistent

where $|=$ denotes the entailment relation of the formal logic used in the representation of our theory and consistency refers also to the corresponding notion in this logic. The particular choice of this underlying formal framework of logic is in general a matter that depends on the problem or phenomena that we are trying to model. In many cases this is taken to be first-order predicate calculus, as for example in the approach of theory completion in [5]. But other logics can be used; e.g., the nonmonotonic logics of logic programming with negation as failure or default logic when the modeling of our problem requires this level of expressibility. In many approaches of machine learning in artificial intelligence where we want to use automated forms of our logic, the choice of logic can also be driven by practical considerations of availability of effective computational models.

One way to distinguish the two forms of reasoning is to consider the extent to which we (a priori) allow the new knowledge H to complement the current theory T. Abduction is typically applied on a model in which we can separate two disjoint sets of predicates: the *observable* predicates and the *abducible* predicates. The basic assumption then is that our model has reached a sufficient level of comprehension of the domain such that all the incompleteness of the model can be isolated (under some working hypotheses) in its abducible predicates. The observable predicates are assumed to be completely defined in T; any incompleteness in their representation comes from the incompleteness in the abducible predicates.

In practice, observable predicates describe the empirical observations of the domain that we are trying to model. The observations are represented by formulae that refer only to observable predicates [and possibly some background auxiliary predicates (see below)] typically by ground atomic facts on the observable predicates. The abducible predicates describe underlying (theoretical) relations in our model that are not observable directly but can, through the model T bring

about observable information. We also have *background* predicates that are auxiliary relations that help us link observable and abducible information (e.g., they describe experimental conditions or known subprocesses of the problem domain that we are modeling).

A *cycle of integration* of abduction and induction [27] emerges that is suitable for our task of incremental scientific modeling. Abduction is first used to transform (and in some sense normalize) the observations to an extensional hypothesis on the abducible predicates. Then induction takes this as input and tries to generalize this extentional information to general rules for the abducible predicates, now treating these as observable predicates for its own purposes. The cycle can then be repeated by adding the learned information on the abducibles back in the model as new partial information on the incomplete abducible predicates. This will affect the abductive explanations of new observations to be used again in a subsequent phase of induction. Hence, through this cycle of integration the abductive explanations of the observations are added to the theory not in the (simple) form that they have been generated but in a generalized form given by a process of induction on these.

The combination of abduction and induction has recently been studied and deployed in several ways within the context of inductive logic programming (ILP). In particular, a new form of ILP, called *theory completion* introduced in [28], [5], aims, as we have described above, to complete the current theory where the newly generated parts of the theory need not be in the form of clauses that refer directly to the predicates of the given training examples (observations). The realization of theory completion through *inverse entailment* [5] can be seen as a particular case of integration of abductive inference for constructing a "minimal" clause (called the bottom clause) and inductive inference to generalize this clause giving the new clause to be added to the theory. This is implemented in Progol 5.0 and applied to several problems including that of the discovery of the function of genes in a network of metabolic pathways [9].

For our specific problem domain of modeling the phenomenon of inhibition, the cycle of integration of abduction and induction is shown in Figure 2. The purpose of the abduction process is to generate hypotheses about inhibited enzymes from the NMR observations of metabolite concentration. For this we need to start with a theory that models how the concentration of metabolites (e.g., up-down regulations) is related to inhibition of enzymes. The purpose of the induction process is to learn from the abduced hypotheses, which are ground facts of inhibition, general rules about the inhibition of enzymes in terms of chemical properties of the inhibitor, functional class of enzymes, etc. Part of the information about inhibition required by the induction process can be obtained from databases such as BRENDA [29]. However, for many inhibitors the available data may not be enough to generate any general rule. The results of abduction, from the previous stage, then act as invaluable data for the induction process.

Abductive logic programming (ALP) [30], [31] is a framework that allows declarative representations of incomplete theories. In this framework a model or a theory T is described in terms of a triple (P, A, IC) consisting of a logic program P, a set of abducible predicates A, and a set of classical logic formulas IC, called the *integrity constraints* of the theory. The program P contains *definitional knowledge* representing the general laws about our problem domain through a complete definition of a set of *observable predicates* in terms of each other, background predicates (which are again assumed to be completely specified in P), and a set of abducible predicates that are open. Using this framework, we will develop a model for analyzing (understanding and subsequently predicting) the effect of toxin substances on the concentration of metabolites. We use as the set of *observables* the single predicate:

concentration (Metabolite, Level, Time)

expressing the fact that at some time, *Time*, a metabolite, *Metabolite*, a certain level of relative variation in concentration, *Level,* has been observed that in the simplest case can take the two values: *down* or *up*. In general, the concentration predication would contain a fourth argument, namely, the name of the toxin that we are examining, but we will assume here for simplicity that we are studying only one toxin at a time and hence we can factor this out. *Background* predicates such as:

reactionnode (Metabolites1,Enzymes,Metabolites2)

describe the topology of the network of the metabolic pathways as depicted in Figure 1. For example, the statement

reactionnode('l2aminoadipate','2.6.1.39','2 oxoglutarate')

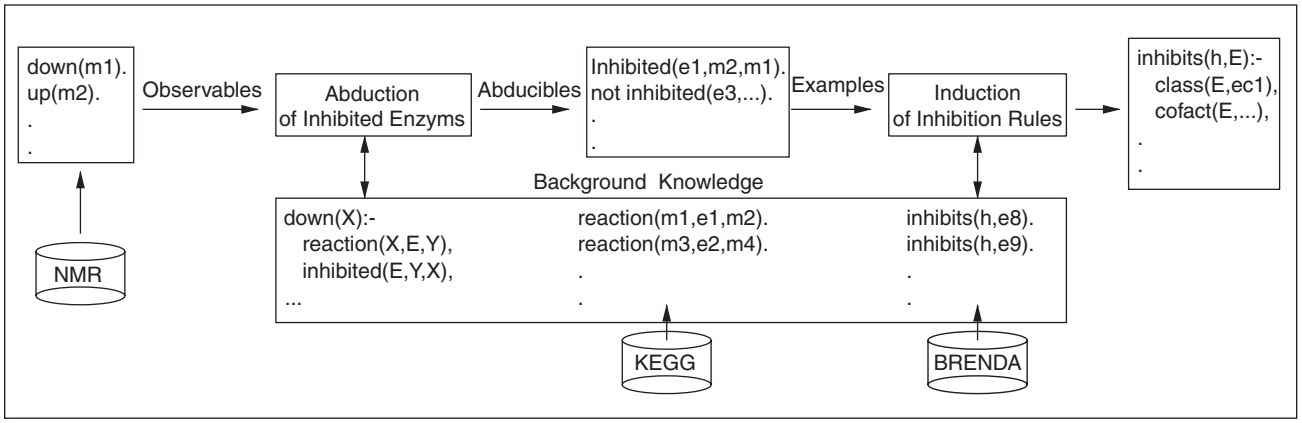

Fig. 2. An abductive/inductive framework for modeling inhibition.

Any system able to predict the inhibitory effect of substances on the metabolic network would be useful in assessing the potential harmful side effects of drugs.

expresses the fact that there is a direct path (reaction) between the metabolites *l2aminoadipate* and 2*oxoglutarate* catalyzed by the enzyme 2.6.1.39. More generally, we can have a set of metabolites on each side of the reaction and a set of different enzymes that can catalyze the reaction.

Note also that these reactions are in general reversible; i.e., they can occur in either direction and indeed the presence of a toxin could result in some reactions changing their direction in an attempt to compensate (rebalance) the effects of the toxin. The model also involves background biochemical data on enzymes and metabolites that would be used in the process of inductive generalization of the abduced hypotheses. The incompleteness of our model resides in the lack of knowledge of which metabolic reactions are adversely affected in the presence of the toxin. This is captured through the declaration of the *abducible* predicate:

inhibited(Enz,Metabolites1,Metabolites2, T)

capturing the hypothesis that at the time *T* the reaction from *Metabolites*1 to *Metabolites*2 is inhibited by the toxin through an adverse effect on the enzyme *Enz* that normally catalyzes this reaction. For example,

inhibited('4.1.2.32','methylamine',' tmao', 8)

expresses the abducible hypothesis that at time 8 the reaction from *methylamine* to *tmao* via the enzyme 4.1.2.32 is inhibited by the toxin.

Hence, the set of abducibles contains the only predicate *inhibited* and completing this would complete the given model. The experimental observations of increased or reduced metabolite concentration will be accounted for in terms of hypotheses on the underlying and nonobservable inhibitory effect of the toxin represented by this abducible predicate.

We now need to provide the program rules and the integrity constraints of our model representation. The rules describe an underlying mechanics of the effect of inhibition of a toxin by defining the observable *concentration* predicate. This model is simple in the sense that it only describes at an appropriate high-level the possible inhibition effects of the toxin, abstracting away from the details of the complex biochemical reactions that occur. It sets out simple general laws under which the effect of the toxin can increase or reduce their concentration. Examples of these rules are:

```
concentration(X,down,T):-
    reactionnode(X,Enz,Y),
    inhibited(Enz,Y,X,T).
concentration(X,down,T):-
    reactionnode(X,Enz,Y),
    not inhibited(Enz,Y,X,T),
    concentration(Y,down,T).
```

The first rule expresses the fact that if a reaction producing metabolite *X* is inhibited at time *T,* then this will cause down concentration of this metabolite at this time. The second rule accounts for changes in the concentration through indirect effects where a metabolite *X* can have down concentration due to the fact that some other substrate metabolite *Y* that produces *X* was caused to have low concentration (even when the reaction is not currently inhibited). Increased concentration is modeled analogously with rules for "up" concentration. For example, we have:

```
concentration(X,up,T):-
    reactionnode(Y,Enz,X),
    inhibited(Enz,X,Y,T).
```

where the inhibition of the reaction from metabolite *X* to *Y* causes the concentration of *X* to go up as *X* is not (currently) consumed due to this inhibition. Note that for a representation that does not involve negation as failure, as we would need when using the Progol 5.0 system, we could use instead the abducible predicate *inhibited(Enz,Status,Y,X,T)* where *Status* would take the two values *true* and *false*.

The underlying and simplifying working hypotheses of this model are as follows:

- the primary effect of the toxin can be *localized* on the individual reactions of the metabolic pathways
- the underlying network of the metabolic pathways is correct and complete
- all the reactions of the metabolic pathways are a priori equally likely to be affected by the toxin
- inhibition in one reaction is sufficient to cause change in the concentration of the metabolites.

The above rules and working hypotheses give a relatively simple model, but this is sufficient as a starting point. In a more elaborate model we could relax the fourth underlying hypothesis of the model and allow, for example, the possibility that the down concentration effect on a metabolite, due to the inhibition of one reaction leading to it, to be compensated by some increased flow of another reaction that also leads to it. We would then have more elaborated rules that express this. For example, the first rule above would be replaced by:

```
concentration(X,down,T):-
        reactionnode(X,Enz,Y),
        inhibited(Enz,Y,X,T),
        not compensated(X,Enz,T).
compensated(X,Enz,T):-
```

```
    reactionnode(X,Enz1,Y),
    different(Enz1,Enz),
    increased(Enz1,Y,X,T).
```

where now the set of abducible predicates includes also the predicate *increased(Enz,M*1*,M2, T)* that captures the assumption that the flow of the reaction from *M*1 to *M*2 has increased at time *T* as a secondary effect of the presence of the toxin.

The abducible information of *inhibited* is required to satisfy several *validity requirements* captured in the integrity constraints of the model. These are stated modularly and separately from the program rules and can be changed without affecting the need to reconsider the underlying model. They typically involve general self-consistency requirements of the model such as:

```
false :-
    concentration(X,down,T),
    concentration(X,up,T).
```

expressing the fact that the model should not entail that the concentration of any metabolite is at the same time down and up. In addition, specific partial information that we may have on the abducible predicates *inhibited* (such as that a certain reaction cannot be inhibited by the toxin that we are examining) can be captured as a validity requirement.

Let us illustrate the use of our model and its possible development with an example. Given the pathways network in Figure 1 and the experimental observation that:

concentration('2oxoglutarate', down, 8)

the following are some of its possible explanations:

E_1 = *{inhibited('*2.3.1.61*','succinate','*2*oxoglutarate',* 8*)}*
E_2 = *{inhibited('*2.6.1.39*','l2aminoadipate','*2*oxoglutarate',* 8*)}*
E_3 = *{inhibited('*1.1.1.42*',' isocitrate','* 2*oxoglutarate',* 8*)}*.

Combining this observation with the additional observation that *concentration('isocitrate', down,* 8*)* makes the third explanation E_3 inconsistent, as this would imply that the concentration of *isocitrate* is up at time 8. Now if we further suppose that we have observed:

concentration('l2aminoadipate', up, 8)
then the above explanation E_2 is able to account for all three observations with no added hypotheses needed. The first observation (of *'*2*oxoglutarate'* down) and third observation (of *'l2aminoadipate'* up) are both accounted as direct effects of this inhibition while the second observation (of '*isocitrate'* down) is accounted for as an indirect knock on effect of the inhibition, assumed in E_2, of the upstream reaction from '*l2aminoadipate'* to '2*oxoglutarate'*. An alternative explanation would be
E'_2 = *{inhibited('*2.6.1.39*','l2aminoadipate','2oxoglutarate',* 8*),*
*inhibited('*1.2.1.31*','l2aminoadipate','l – lysine',* 8*)}*.

Applying a principle of *minimality* of explanations [35] or more generally of *maximal compression* we would prefer the explanation E_2 over E'_2.

Empirical Evaluation

The purpose of the experiments in this section is to empirically evaluate the inhibition model, described in the previous section, on real metabolic pathways and real NMR data. In this experiment we evaluate ground hypotheses that are generated using the inhibition model given observations about the change in the concentration of some metabolites. We also examine if we can improve the accuracy of the model by further generalizing the ground hypotheses.

In this experiment Progol 5.0 (available at http://www.doc.ic.ac.uk/~shm/Software/progol5.0/) is used to generate both ground and nonground hypotheses. As a part of background knowledge, we use the relational representation of biochemical reactions involved in a metabolic pathway that is affected by the toxin. This information is extracted from KEGG as explained above. The observable data are up/down regulations of metabolites obtained from NMR spectra. The technique that has been used to transform raw time-series data is described in [32]. The up/down regulations of metabolites at different time periods are then encoded as Prolog ground facts.

Background knowledge required for nonground hypotheses can be obtained from databases such as BRENDA [29], as discussed above. This background information can include information about enzyme classes, cofactors, etc. In our experiments for learning nonground hypotheses, we include the possibility that a given chemical compound can be inhibiting a whole enzymatic class, since this situation is possible in noncompetitive inhibition. For example, a very strong reducer or oxidant affecting many oxidoreductases (1.-.-.-). In our case, since the mechanism of inhibition of toxin is unknown, we leave this possibility open.

In this experiment we use up/down regulation of metabolites at 8 hours to 96 hours as training/test examples and apply a leave-one-out test strategy (randomly leave out one test example and use the rest as training data). The performance is then evaluated by varying the size of randomly chosen training sets. More details about this experimental method can be found in [13].

The model that has been used for evaluating the hypotheses generated by Progol explicates the closed-world assumption [33]. In other words, we are working under the assumption that a reaction is not inhibited unless we have a fact that says otherwise:

```
inhibited(Enz,false,X,Y):-
    reactionnode(Y,Enz,X),
    not(inhibited(Enz,true,_,_)).
```

The predictor that we have used in our experiments converts the three-class problem that we have ("up," "down," and "unknown") to a two-class prediction with "down" as the default class. For this purpose we use the following test predicate:

```
concentration1(X,up,T):-
    concentration(X,up,T),
    not(concentration(X,down,T)).
concentration1(X,down,T):-
    not(concentration1(X,up,T)).
```

According to our model, there are many possible hypotheses that can explain the up-regulation and down-regulation of the observed metabolites. However, Progol's search attempts to find the most compressive hypotheses. The following are examples of ground hypotheses returned by Progol for the inhibitory effect of hydrazine at 8 hours:

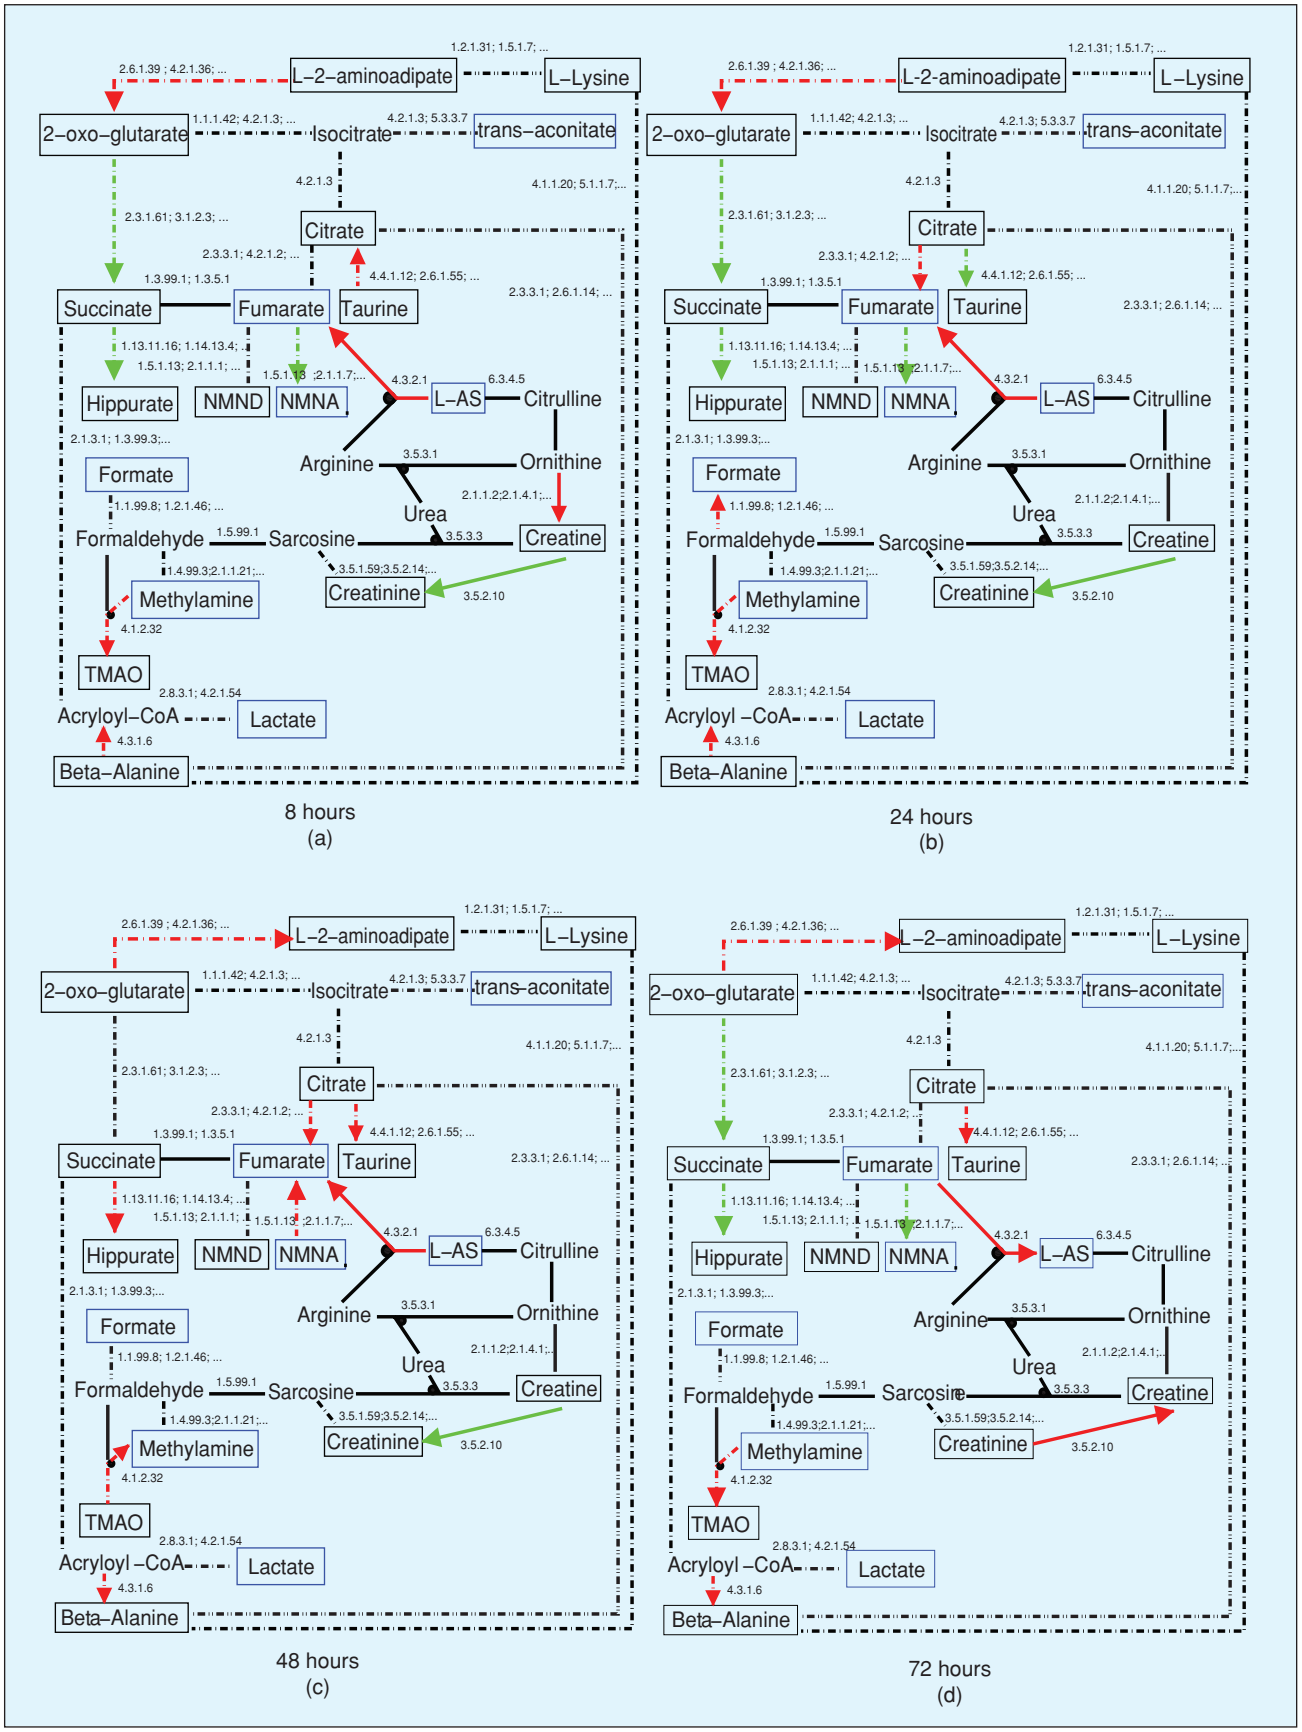

Fig. 3. Examples of ground hypotheses for Hydrazine at 8 hours, 24 hours, 48 hours, and 72 hours. Red arrows correspond to 'inhibited' and green arrows correspond to 'not inhibited' hypotheses. The model suggests that some reactions remain inhibited through different time periods.

inhibited('2.6.1.39',true,'l2aminoadipate','2oxoglutarate',8).
inhibited('2.3.1.61',false,'2oxoglutarate','succinate',8).
inhibited('1.13.11.16',false,'succinate','hippurate',8).
inhibited('2.6.1.-',true,'taurine','citrate',8).
inhibited('3.5.2.10',false,'creatine','creatinine',8).
inhibited('4.1.2.32',true,'methylamine','tmao',8).
inhibited('4.3.1.6',true,'beta-alanine','acryloyl-coA',8).
inhibited('4.3.2.1',true,'l-as','fumarate',8).

Examples of ground hypotheses are illustrated in Figure 3. In this figure, red arrows correspond to 'inhibited' and green arrows correspond to 'not inhibited' hypotheses. As shown in this figure, the model suggests that some reactions remain inhibited through different time periods. According to the domain experts who evaluated these results, one of these enzymes (i.e., EC2.6.1.39) was known to be inhibited by hydrazine. Another hypothesis suggested by the model agrees with the speculations about the inhibition of enzyme EC4.3.2.1 by hydrazine [34]. Experimental evaluations in vivo are required to test this hypothesis.

The overall performance of ground and nonground hypotheses are shown in Figure 4. In this graph, the vertical axis shows the predictive accuracy and the horizontal axis shows the number of training examples. According to this graph, we have a better predictive accuracy when we use the closed-world assumption compared to the accuracy when we do not use this assumption. The reason for this is that the closed-world assumption allows the rules of the model (as represented in Progol) to apply in more cases than without the assumption. These graphs also show that in all cases the overall accuracy is above the default accuracy (a model that simply guesses the majority class) and increases with the number of training examples.

In this experiment, Progol also attempted to generate general rules for inhibition, effectively trying to generalize from the ground facts in the abductive explanations. An example of such a nonground rule is:

inhibited(Enz, true, M1, M2) :-
reactionnode(M2,Enz,M1),
class(Enz,'aminotransferase')

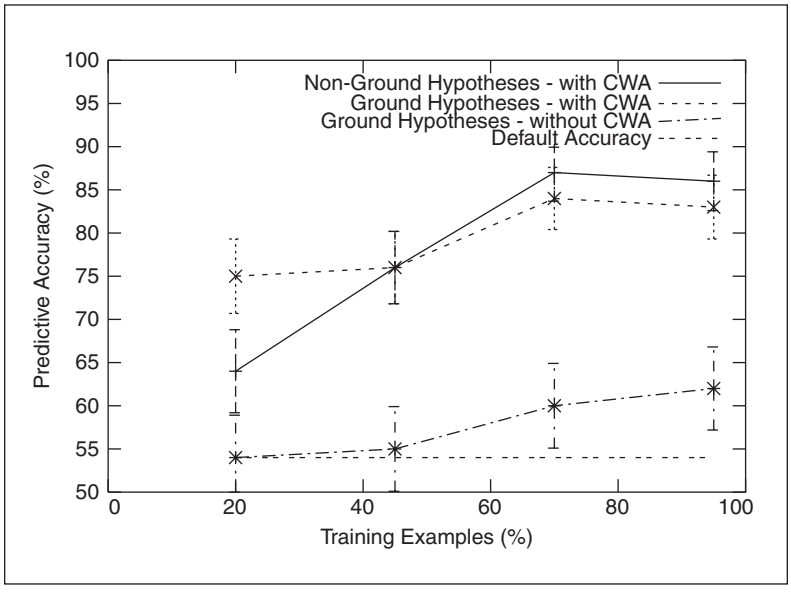

Fig. 4. Performance of ground and nonground hypotheses generated by Progol using a leave-one-out test strategy.

expressing the information that reactions that are catalyzed by enzymes in the enzymatic class *'aminotransferase'* are inhibited by the toxin.

According to the comparison shown in Figure 4, it is instructive to accept these (seemingly overgeneral) rules into our model and examine the effect of this generalization on the predictive accuracy of the model compared with the case where only ground explanations are allowed. This figure shows that for a small number of training examples, ground hypotheses (with closed-world assumption) have a better predictive accuracy than nonground hypotheses. These results suggest that for a small number of training examples (e.g., less than 45%) the induced nonground hypotheses are either too general or overfitted the training data and therefore lead to a lower predictive accuracy than the ground hypotheses. However, when more training examples are provided (i.e., more than 70%), nonground hypotheses show a better performance than ground hypotheses.

Related Work

The abduction technique that is used in this article can be compared with the one in the robot scientist project [9] where Progol5.0 was used to generate ground hypotheses about the function of genes. Abduction has been also used within a system called GenePath [8] to find relations from experimental genetic data in order to facilitate the analysis of genetic networks. Similarly, in [35] abduction has been used to generate gene interactions and genetic pathways from microarray experimental data. Combinations of abduction and induction have been also used for learning robot planners by completing the specific domain knowledge required, within a general theory of planning that the robot uses for its navigation [36], [6].

Bayesian networks are among the most successful techniques that have been used for modeling biological networks. In particular, gene expression data have been widely modeled using Bayes net techniques [37], [38]. On the MetaLog project, Bayes nets have also been used to model metabolic networks [39]. A key advantage of the logical modeling approach in this article compared with the Bayes net approach is the ability to incorporate background knowledge of existing known biochemical pathways together with information on enzyme classes and reaction chemistry. The logical modeling approach also produces explicit hypotheses concerning the inhibitory effects of toxins.

Our approach can also help address an important chellenge in 'top-down' systems biology, namely how to describe and represent holistically the disregulation of a metabolic system with multiple cell types where homeostatic function is dispersed in space and time.

Conclusions

We have studied how to use abduction and induction in scientific modeling concentrating on the problem of inhibition of metabolic pathways. Our work has demonstrated the feasibility of a process of scientific model development through an integrated use of abduction and induction. This is to our knowledge the first

time that abduction and induction have been used together in building life-science models from empirical data. We also address the problem of extreme disparities of scale between the temporal measurements underlying the experiment and the model, respectively. This involves avoiding standard autoregressive assumptions used in other temporal modeling approaches and demonstrates the strength and flexibility of the abductive ILP approach for dealing with such problems.

In this study, hypotheses about inhibition were built using the ILP system Progol5.0 and predictive accuracy was assessed for both the ground and the nonground cases. These hypotheses were also evaluated in terms of biological insight provided. Experimental evaluations in vivo are required to test some of these hypotheses.

Maximization of drug efficacy and safety are major issues in the pharmaceutical industry and understanding the mechanistic interactions of drugs with their desired (pharmacological) and undesired (toxic) targets is of great scientific, medical, and indeed economic importance. Our new approach can give new insights into the metabolic network responses of man and animals to drugs at the system level and therefore should prove to be a valuable tool in drug discovery and development. Moreover, because there are many network commonalities between animal models and humans, it should be possible to create novel predictive models of drug toxicity that cross species boundaries and that are applicable to man which historically has been a major challenge in pharmaceutical research. Our approach is made more attractive still by the fact that non- or minimally invasive metabolic metrics (from urine or plasma) can be used to describe intact system function that will assist drug safety evaluations in phase I-IV clinical trials and potentially in the population at large in future molecular epidemiology studies.

In the current study we used simple background knowledge concerning the class of enzymes to allow the construction of nonground hypotheses. Despite this limited use of background knowledge, we achieved an increase in predictive accuracy over the case in which hypothesis were restricted to be ground. In future work we hope to extend the representation to include structural descriptions of the reactions involved in a style similar to that described in [40].

Acknowledgments

We would like to thank F. Pazos, C. Caulcott, N. Cooper, C. Rawlings, P. Bang, and E. Holmes for their useful discussions and advice and D. Crockford and T. Ebbels for preparing the NMR data. This work was supported by the DTI project "MetaLog— Integrated Machine Learning of Metabolic Networks applied to Predictive Toxicology." The third author is grateful to the Department of Computing at Imperial College, London, for hosting him in 2003/2004 during which this work had started.

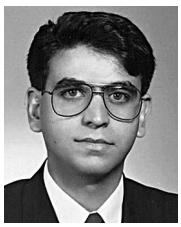

Alireza Tamaddoni-Nezhad is a research associate in the Department of Computing and the Centre for Integrative Systems Biology at Imperial College London. His main research interest is machine learning, in particular inductive logic programming (ILP) and applications to bioinformatics. His current research includes machine learning of metabolic networks using abduction and induction. During his Ph.D. study at Imperial College (2001-) he has developed a genetic-based search for the ILP system CProgol. He also has an M.Sc. in computer engineering from Tehran Polytechnic (1996) and a B.Sc. in computer engineering from Sharif University of Technology (1993).

Raphael Chaleil is a research associate in the Department of Biological Sciences at Imperial College London. He has a Licence in biochemistry and molecular and cellular biology from the University of Tours in France (1997) and an M.Sc. in molecular modeling and bioinformatics from Birkbeck College London (2000). During his Ph.D. study at Imperial College (2000-) he has worked on computational studies of metabolism from a network approach.

Antonis C. Kakas is a professor at the Computer Science Department of the University of Cyprus. He obtained his Ph.D. in theoretical physics from Imperial College London in 1984. In 1989 he started working in computational logic and artificial intelligence. His main research interests are abduction and argumentation, with specific interest in the integration of abductive, inductive, and constraint logic programming and applications in the areas of computational bioscience, agent argumentative deliberation, and the theory of actions and change. He has coedited the book *Abduction and Induction: Essays on their Relation and Integration* and a special issue of the *Journal of Logic Programming* on abductive logic programming.

Michael Sternberg holds the Chair of Structural Bioinformatics in the Department of Biological Sciences at Imperial College, and is also the Director of the Imperial College Centre for Bioinformatics. He obtained his first degree in Cambridge in natural sciences specializing in theoretical physics. He obtained the M.Sc. in Computing at Imperial College. His D.Phil. research in Oxford in Lord Phillip's laboratory was on the analysis and prediction of protein structure and this area remained the focus for his research. Subsequent posts included: postdoctoral research in Oxford; a lectureship in the Department of Crystallography, Birkbeck; and the head of the Biolomolecular Modelling Laboratory, Imperial Cancer Research Fund (1988-2001).

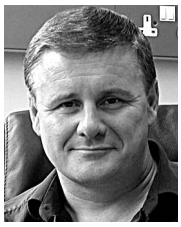

Jeremy Nicholson is a Professor of Biological Chemistry (1998-) at Imperial College where he also heads the Department of Biomolecular Medicine. He obtained his B.Sc from Liverpool University in 1977, and his Ph.D. from London University in 1980 in biochemistry. He was with Birkbeck College, London, from 1985 to 1998. He is the author of more than 500 scientific papers, patents, and articles. He received the 1992 Royal Society of Chemistry Silver Medal for Analytical

Science, the 1997 Royal Society of Chemistry Gold Medal for Analytical Chemistry, the 1994 Chromatographic Society Jubilee Silver Medal, the 2002 Pfizer Prize for Chemical and Medicinal Technologies, the 2003 Royal Society of Chemistry medal for Chemical Biology, and the 2006 Pfizer Global Research and Development prize for Chemistry. He holds visiting and honorary professorships in several countries and is on the editorial board of ten international science journals. He is also a consultant to numerous pharmaceutical companies in the UK, Europe, and the United States and is a founder director of Metabometrix.

Stephen Muggleton holds the EPSRC chair of bioinformatics and is director of the Imperial College Computational Bioinformatics Centre (2001-) and director of modeling for the Imperial College Centre for Integrated Systems Biology. His career has concentrated on the development of theory, implementations, and applications of machine learning, particularly in the field of inductive logic programming. Over the last decade he has collaborated increasingly with biological colleagues, in particular Prof. Mike Sternberg, on applications of machine learning to biological prediction tasks. These tasks have included the determination of protein structure, the activity of drugs and toxins, and the assignment of gene function. Previous posts were as professor of machine learning at the Computer Science Department, University of York (1997–2001); reader in machine learning and research fellow at Wolfson College Oxford (1993–1997); EPSRC advanced research fellow (1993–1997); visiting associate professor (Fujitsu Chair) at the University of Tokyo; and EPSRC postdoctoral fellow and Turing Institute Fellow (1987–1992). He has a Ph.D. in artificial intelligence from Edinburgh University (1986) and a B.Sc. in computer science from Edinburgh University (1983). Professional positions include fellow of the American Association for Artificial Intelligence (2002-), editor-in-chief of the *Machine Intelligence* series; panel member for the DTI Functional Genomics initiative (2002-), and the BBSRC EBI Committee (2004-).

Address for Correspondence: A. Tamaddoni-Nezhad, Department of Computing, Imperial College London. E-mail: atn@doc.ic.ac.uk.

References

[1] P.A. Flach and A.C. Kakas, Eds., *Abductive and Inductive Reasoning, Pure and Applied Logic*. Norwell, MA: Kluwer, 2000.
[2] A. Yamamoto, "Which hypotheses can be found with inverse entailment?" in *Proc. 7th Int. Workshop Inductive Logic Programming*, Berlin, 1997, pp. 296–308.
[3] K. Ito and A. Yamamoto, "Finding hypotheses from examples by computing the least generalization of bottom clauses," in *Proc. Discovery Science '98*, Berlin, vol. 1532, LNAI, Springer-Verlag, pp. 303–314.
[4] K. Inoue, "Induction, abduction and consequence-finding," in *Proc. Int. Workshop Inductive Logic Programming,* vol. 2157, LNAI, Springer-Verlag, 2001, pp. 65–79.
[5] S.H. Muggleton and C.H. Bryant, "Theory completion using inverse entailment," in *Proc. 10th Int. Workshop Inductive Logic Programming*, Berlin, vol. 1866, LNAI, Springer-Verlag, 2000, pp. 130–146.
[6] S. Moyle, "Using theory completion to learn a robot navigation control program," in *Proc. 12th Int. Conf. Inductive Logic Programming*, vol. 2583, LNAI, Springer-Verlag, 2002, pp. 182–197.
[7] O. Ray, K. Broda, and A. Russo, "Hybrid abductive inductive learning: A generalisation of Progol," in *Proc. 13th Int. Conf. Inductive Logic Programming*, 2835, LNAI, Springer-Verlag, 2003, pp. 311–328.
[8] B. Zupan, I. Bratko, J. Demsar, P. Juvan, J.A Halter, A. Kuspa, and G. Shaulsky, "GenePath: A system for automated construction of genetic networks from mutant data," *Bioinformatics*, vol. 19, no. 3, pp. 383–389, 2003.
[9] R.D. King, K.E. Whelan, F.M. Jones, P.K.G. Reiser, C.H. Bryant, S.H. Muggleton, D.B. Kell, and S.G. Oliver, "Functional genomic hypothesis generation and experimentation by a robot scientist," *Nature*, vol. 427, pp. 247–252, 2004.
[10] MetaLog [Online]. Available: http://www.doc.ic.ac.uk/bioinformatics/metalog/.
[11] H. Ogata, S. Goto, K. Sato, W. Fujibuchi, H. Bono, and M. Kanehisa, "KEGG: Kyoto encyclopedia of genes and genomes," *Nucleic Acids Res.*, vol. 27, no. 1, pp. 29–34, 1999.
[12] A. Tamaddoni-Nezhad, A. Kakas, S.H. Muggleton, and F. Pazos, "Modelling inhibition in metabolic pathways through abduction and induction," in *Proc. 14th Int. Conf. Inductive Logic Programming*, vol. 3195, LNAI, Springer-Verlag, 2004, pp. 305–322.
[13] A. Tamadonni-Nezhad, R. Chaleil, A.Kakas, and S.H. Muggleton, "Application of abductive ILP to learning metabolic network inhibition from temporal data," *Machine Learning.*, vol. 64, pp. 209–230, 2006.
[14] J.C. Lindon, J.K. Nicholson, E. Holmes, et al., "Contemporary issues in toxicology. The role of metabonomics in toxicology and its evaluation by the COMET project," *Tox. Appl. Pharm.*, vol. 187, pp. 137–146, 2003.
[15] E. Ravasz, A.L. Somera, D.A. Mongru, Z.N. Oltvai, and A.L. Barabasi, "Hierarchical organization of modularity in metabolic networks," *Science*, vol. 297, no. 5586, pp. 1551–1555, 2002.
[16] E. Alm and A.P. Arkin, "Biological networks," *Curr. Opin. Struct. Biol.*, vol. 13, no. 2, pp. 193–202, 2003.
[17] B. Alberts, D. Bray, J. Lewis, M. Raff, K Roberts, and J.D. Watson, *Molecular Biology of the Cell*, 3rd ed. New York: Garland, 1994.
[18] J.A. Papin, N.D. Price, S.J. Wiback, D.A. Fell, and B.O. Palsson, "Metabolic pathways in the post-genome era," *Trends Biochem. Sci.*, vol. 28, no. 5, pp. 250–258, May 2003.
[19] R. Alves, R.A. Chaleil, and M.J. Sternberg, "Evolution of enzymes in metabolism: A network perspective," *Mol. Biol.*, vol. 320, no. 4, pp. 751–770, 2002.
[20] T.A. Swierkosz, L. Jordan, M. McBride, K. McGough, J. Devlin, and R.M. Botting, "Actions of paracetamol on cyclooxygenases in tissue and cell homogenates of mouse and rabbit," *Med. Sci. Monit.*, vol. 8, no. 12, pp. 496–503, Dec. 2002.
[21] H.J. Zimmerman and W.C. Maddrey, "Acetaminophen (paracetamol) hepatotoxicity with regular intake of alcohol: analysis of instances of therapeutic misadventure," *Hepatology*, vol. 22, no. 3, pp. 767–773, Sept. 1995.
[22] K. Popper, *The Logic of Scientific Discovery*. New York: Basic Books, 1959.
[23] C.G. Hempel, *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press, 1965.
[24] C.S. Peirce, *Essays in the Philosophy of Science*. Liberal Arts Press, 1957.
[25] J.R. Josephson and S.G. Josephson, Eds., *Abductive Inference: Computation, Philosophy, Technology*. Cambridge, UK: Cambridge Univ. Press, 1994.
[26] L. Magnani, *Abduction, Reason and Science*. Norwell, MA: Kluwer, 2001.
[27] P. Flach and A.C. Kakas, "Abductive and inductive reasoning: Background and issues," in *Abductive and Inductive Reasoning, Pure and Applied Logic*, P.A. Flach and A.C. Kakas, Eds. Norwell, MA: Kluwer, 2000.
[28] S. Moyle and S.H. Muggleton, "Learning programs in the event calculus," in *Proc. 7th Inductive Logic Programming Workshop*, vol. 1297, LNAI, Springer-Verlag, 1997, pp. 205–212.
[29] BRENDA Webpage [Online]. Available: http://www.brenda.uni-koeln.de/.
[30] A.C. Kakas, R.A. Kowalski, and F. Toni, "Abductive logic programming," *J. Logic Computat.*, vol. 2, no. 6, pp. 719–770, 1993.
[31] A.C. Kakas and M. Denecker, "Abduction in logic programming," in *Computational Logic: Logic Programming and Beyond. Part I*, A.C. Kakas and F. Sadri, Eds., 2002, pp. 402–436.
[32] D.J. Crockford, H.C. Keun, L.M. Smith, E. Holmes, and J.K. Nicholson, "Curve-fitting method for direct quantitation of compounds in complex biological mixtures using 1h NMR: Application in metabonomic toxicology studies," *Anal. Chem.*, vol. 77, no. 14, pp. 4556–4562, 2005.
[33] K.L. Clark, "Negation as failure," in *Readings in Nonmonotonic Reasoning*, M.L. Ginsberg, Ed. Los Altos, CA: Kaufmann, 1987, pp. 311–325.
[34] A.W. Nicholls, E. Holmes, J.C. Lindon, R.D. Farrant, J.N. Haselden, S.J.P. Damment, C.J. Waterfield, and J.K. Nicholson, "Metabonomic investigations into hydrazine toxicity in the rat," *Chem. Res. Toxicol.,* vol. 14, no. 8, pp. 975–987, 2001.
[35] I. Papatheodorou, A. Kakas, and M. Sergot, "Inference of gene relations from microarray data by abduction," Dept. Computing, Imperial College London, London, UK, Tech. Rep. 2005/3, 2005.
[36] B.R. Wellner, "An abductive-inductive learning framework for logic-based agents," master's thesis, Imperial College of Science Technology and Medicine, London, UK, 1999.
[37] N. Friedman, M. Linial, I. Nachman, and D. Pe'er, "Using Bayesian networks to analyze expression data," *J. Comp. Bio.*, vol. 7, no. 3-4, pp. 601–620, 2000.
[38] S. Imoto, T. Goto, and S. Miyano, "Estimation of genetic networks and functional structures between genes by using Bayesian networks and nonparametric regression," in *Proc. Pacific Symp. Biocomputing,* 2002, pp. 175-186.
[39] A. Tamaddoni-Nezhad, S. Muggleton, and J. Bang, "A Bayesian model for metabolic pathways," in *Proc. Int. Joint Conf. Artificial Intelligence Workshop Learning Statistical Models from Relational Data*, Acapulco, Mexico, 2003, pp. 50–57, 2003.
[40] S.H. Muggleton, A. Tamaddoni-Nezhad, and H. Watanabe, "Induction of enzyme classes from biological databases," in *Proc. 13th Int. Conf. Inductive Logic Programming*, vol. 2835, LNAI, Springer-Verlag, 2003, pp. 269–280.