

Application of abductive ILP to learning metabolic network inhibition from temporal data

Alireza Tamaddoni-Nezhad · Raphael Chaleil · Antonis Kakas · Stephen Muggleton

Received: 11 April 2005 / Revised: 6 March 2006 / Accepted: 31 March 2006 / Published online: 30 June 2006
Springer Science + Business Media, LLC 2006

Abstract In this paper we use a logic-based representation and a combination of Abduction and Induction to model inhibition in metabolic networks. In general, the integration of abduction and induction is required when the following two conditions hold. Firstly, the given background knowledge is incomplete. Secondly, the problem must require the learning of general rules in the circumstance in which the hypothesis language is disjoint from the observation language. Both these conditions hold in the application considered in this paper. Inhibition is very important from the therapeutic point of view since many substances designed to be used as drugs can have an inhibitory effect on other enzymes. Any system able to predict the inhibitory effect of substances on the metabolic network would therefore be very useful in assessing the potential harmful side-effects of drugs. In modelling the phenomenon of inhibition in metabolic networks, background knowledge is used which describes the network topology and functional classes of inhibitors and enzymes. This background knowledge, which represents the present state of understanding, is incomplete. In order to overcome this incompleteness hypotheses are considered which consist of a mixture of specific inhibitions of enzymes (ground facts) together with general (non-ground) rules which predict classes of enzymes likely to be inhibited by the toxin. The foreground examples are derived from *in vivo* experiments involving NMR analysis of time-varying metabolite concentrations in

Editor: Rui Camacho

A. Tamaddoni-Nezhad (✉) · S. Muggleton
Dept. of Computing, Imperial College London
e-mail: atn@doc.ic.ac.uk

S. Muggleton
e-mail: shm@doc.ic.ac.uk

R. Chaleil
Dept. of Biological Sciences, Imperial College London
e-mail: r.chaleil@imperial.ac.uk

A. Kakas
Dept. of Computer Science, University of Cyprus
e-mail: antonis@ucy.ac.cy

rat urine following injections of toxins. The model's performance is evaluated on training and test sets randomly generated from a real metabolic network. It is shown that even in the case where the hypotheses are restricted to be ground, the 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 is provided, non-ground hypotheses show a better predictive accuracy than ground hypotheses. The model is also evaluated in terms of the biological insight that it provides.

Keywords Abductive ILP · Bioinformatics · Modelling metabolic networks

1. Introduction

The combination of abduction and induction has recently been explored from a number of angles (Flach and Kakas, 2000). Moreover, theoretical issues related to the completeness of this form of reasoning have also been discussed by various authors (Yamamoto, 1997; Ito and Yamamoto, 1998; Inoue, 2001). Some implemented systems have been developed for combining abduction and induction (Muggleton and Bryant, 2000) and others have recently been proposed (Moyle, 2002; Ray, 2003). There have also recently been demonstrations of the application of abduction/induction systems in the area of Systems Biology (Zupan et al., 2003; King et al., 2004) though in these cases the generated hypotheses were ground. In general, the integration of abduction and induction is required when the following two conditions hold. Firstly, the given background knowledge is incomplete. Secondly, the problem must require the learning of general rules in the circumstance in which the hypothesis language is disjoint from the observation language.

The research reported in this paper is being conducted as part of the MetaLog project (MetaLog). This project aims at building of causal rules concerning the actions of toxins. These rules are built from Nuclear Magnetic Resonance (NMR) data in combination with background knowledge derived from the Kyoto Encyclopedia of Genes and Genomes (KEGG) (Ogata, 1999). The NMR spectra provide information concerning the flux of metabolite concentrations before, during and after administration of a toxin.

In this paper we use a logic-based representation and a combination of Abduction and Induction to model inhibition in metabolic networks. Inhibition is very important from the therapeutic point of view since many substances designed to be used as drugs can have an inhibitory effect on other enzymes. Any system able to predict the inhibitory effect of substances on the metabolic network would therefore be very useful in assessing the potential harmful side-effects of drugs. In our specific problem domain of modelling the phenomenon of inhibition, we use the frameworks of Abductive Logic Programming (ALP) (Kakas et al., 1993; Kakas and Denecker., 2002) and Inductive Logic Programming (ILP) (Muggleton, 1991; Muggleton and De Raedt, 1994). The purpose of abduction is to generate hypotheses (in the form of ground facts) about inhibited enzymes from the NMR observations of metabolite concentration. The purpose of the induction process is to learn general rules from these abducted hypotheses. These general rules describe the inhibition of enzymes in terms of a variety of properties, including chemical properties of the inhibitor and the functional class of enzymes. Part of the information about inhibition required by the induction process can be obtained from databases such as BRENDA.¹ 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.

¹ <http://www.brenda.uni-koeln.de/>.

In a previous paper (Tamaddoni-Nezhad et al., 2004) describing our initial investigation in this topic we modelled the effects of a single toxin, hydrazine. The previous model ignored the temporal variance of metabolite concentrations. By contrast, in this paper 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 have been made available to the project by our collaborators who studied these toxins as part of the COMET project (Lindon et al., 2003).

In our study, examples extracted from the NMR data consist of metabolite concentrations (up or 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 (from KEGG) and enzymes known to be inhibited by the toxin. This background knowledge, which represents the present state of understanding, is incomplete. In order to overcome this incompleteness hypotheses are considered which consist of a mixture of specific inhibitions of enzymes (ground facts) together with general rules (non-ground) which predict classes of enzymes likely to be inhibited by the toxin. Hypotheses about inhibition are built using Progol5.0 (Muggleton and Bryant, 2000). The predictive accuracy is assessed and compared for both ground and non-ground hypotheses. It is shown that even in the case where the hypotheses are restricted to be ground, the 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 is provided, non-ground hypotheses show a better predictive accuracy than ground hypotheses.

The rest of the paper is organised as follows. Section 2 introduces the underlying biological problem of our study. Section 3 gives the background to the task of logical modelling of scientific theories using abduction and induction. The experiments of learning ground and non-ground hypotheses for inhibition are then described in Section 4. Section 5 discusses related work and Section 6 concludes the paper.

2. Inhibition in metabolic networks

Metabolism is made by the processes used by organisms to transform elements from their environment into the essential building blocks for life. Moreover, metabolism provides a source of energy for cells and degrades toxic compounds in preparation for excretion. The graph of these interlinked chemical reactions is known as the *metabolic network* (Jeong et al., 2000; Ravasz, 2002; Alm and Arkin, 2003). The reactions that take place in this network need to be finely coordinated. Biochemical reactions are catalysed by highly specialised proteins known as enzymes. Without enzymes the rate of most reactions in living organisms would be too slow to sustain life. Enzymes act as keys controlling the activation of the different parts of the network and are therefore critical to its operation (Alberts et al., 1994).

The collection of known metabolic networks, based on data accumulated through years of research, is now accessible through publicly-available metabolic databases, allowing their study from a network perspective (Papin, 2003; Alves et al., 2002). Even with the help of this new Systems Biology approach to metabolism, we are still far from understanding many of its properties. One of the less understood phenomena, especially from a network perspective, is *inhibition*. Some chemical compounds, known as inhibitors, can affect enzymes, impeding their function. This in turn affects the normal flux in the metabolic network, the result of which is reflected in the accumulation or depletion of certain metabolites.

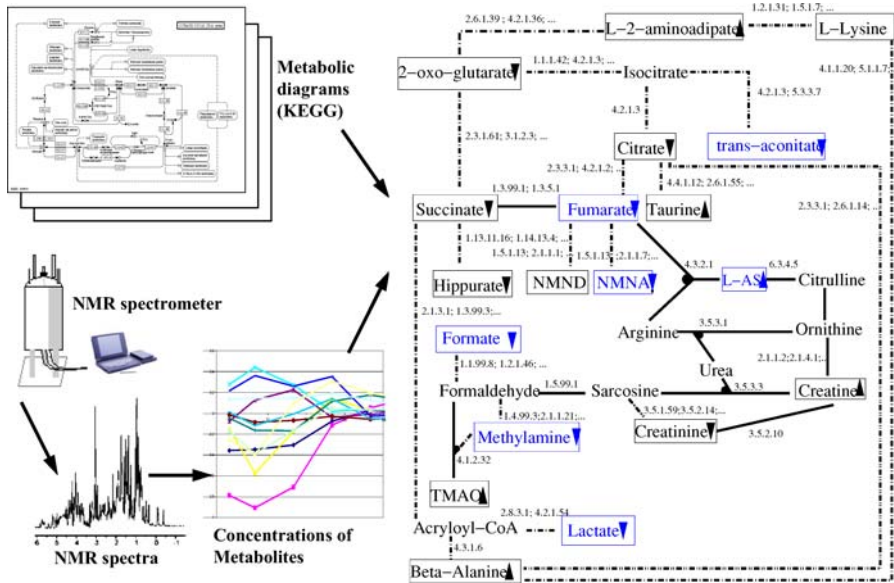


Fig. 1 A metabolic sub-network involving metabolites affected by hydrazine. Information on up/down changes in metabolite concentrations from NMR spectra is combined with KEGG metabolic diagrams. The enzymes associated with a single reaction (solid line) or a linear pathway (dotted line) are shown as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

Inhibition is important from a therapeutic point of view since many substances designed to be used as drugs can have an inhibitory effect on other enzymes. For example, Paracetamol is an inhibitor of COX-3 cyclo-oxygenase (Swierkosz et al., 2002), preventing the formation of arachidonic acid into prostaglandins, which are involved in the pain and fever process (Boutaud et al., 2002). 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 (Zimmerman and Maddrey, 1995). Any system able to predict the inhibitory effect of substances on the metabolic network would therefore be very useful in assessing the potential harmful side-effects of drugs.

In this work we use experimental data on the accumulation and depletion of metabolites to model the inhibitory effect of various toxins, including hydrazine ($\text{NH}_2\text{-NH}_2$), in the metabolic network of rats. Figure 1 shows the metabolic pathways sub-network of interest also indicating with “up” and “down” arrows, the observed effects of hydrazine on the concentration of some of the metabolites involved. This sub-network was manually built from the information contained in the KEGG metabolic database (Ogata, 1999). We started from the set of chemical compounds for which there is information from the Nuclear Magnetic Resonance (NMR) experiments on up- and down- regulation following toxin treatment. From this 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 reactions²) are excluded. KEGG contains a static³ 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 catalysing the reaction. NMR data provides 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) which are measured after 8 hours, 24 hours, 48 hours, 72 hours and 96 hours. The effect of a toxin on the concentrations of chemical compounds is coded in a binary way. Only up/down changes (increasing/decreasing) in compound concentrations are incorporated in the model.

In this sub-network 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 modelling here (inhibition). Inhibition of a given reaction causes the substrates to accumulate. This may cause an upstream enzyme to start working backwards 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. The enzymes associated with the relations (single chemical reactions or linear pathways) are shown as a single enzyme or a list of enzymes.

3. Logical modelling of scientific theories

Modelling a scientific domain is a continuous process of observing the phenomena, understanding these 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 which gets 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 (Popper, 1959; Hempel, 1965) 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 nineteenth century, Charles Sanders Peirce (Harstshorne, 1958; Peirce, 1957) sets out clearly these three forms of syllogistic reasoning—deduction, abduction and induction—and

² Selection of promiscuous compounds was done manually by considering those involved in more than 50 reactions.

³ In real life the enzymes are dynamically produced on demand when the reaction needs to occur.

studies their respective role in the development of scientific theories. More recently, several authors (see for example the books (Holland, 1989; Josephson and Josephson, 1994; Magnani, 2001; Flach and Kakas, 2000)) have studied abduction and induction from the perspective of Artificial Intelligence and Cognitive Science. In particular, the work in Flach and Kakas (2000) 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:

$$T \cup H \models O, \text{ and}$$

$$T \cup H \text{ is consistent.}$$

where \models 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 Muggleton and Bryant (2000). But other logics can be used, e.g. the non-monotonic logics of Logic Programming with Negation as Failure or Default Logic when the modelling 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 for these.

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, T , 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 T 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 sub-processes of the problem domain that we are modelling).

Having isolated the incompleteness of our model in the abducible predicates, these will form the basis of *abductive explanations* for understanding, according to the model, the specific observations that we have of our scientific domain. Abduction generates in these explanations (typically) *extensional* knowledge that refers only to the abducible predicates and that is specific to the particular state or scenario of the world pertaining to the observations explained and to the given model T from which they were generated. Adding an explanation to the theory then allows us to predict further observable information but this new predictive

power is restricted to come only through the already known rules in our theory that link the observable predicates to other predicates.

On the other hand, inductive inference generates *intentional knowledge* in the form of new general rules that now can provide, either directly or indirectly through the current theory T that they extend, new interrelationships between the predicates of our theory that can include the observable predicates and even in some cases new predicates. The inductive hypothesis thus introduces new, hitherto unknown, links between the relations that we are studying thus allowing new predictions on the observable predicates that would not have been possible before from the original theory under any abductive extension.

A *cycle of integration* of abduction and induction (Flach and Kakas, 2000) emerges that is suitable for our task of incremental scientific modelling. Abduction is first used to transform (and in some sense normalise) the observations to an extensional hypothesis on the abducible predicates. Then induction takes this as input and tries to generalize this extensional 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.

A simple example, adapted from Ray (2003), that illustrates this cycle of integration of abduction and induction is as follows. Suppose that our model contains the following rule and background facts:

sad(X) *if* *tired*(X), *poor*(X).
tired(*oli*), *tired*(*ale*), *tired*(*kr*),
academic(*oli*), *academic*(*ale*), *academic*(*kr*),
student(*oli*), *lecturer*(*ale*), *lecturer*(*kr*).

where our only observable predicate is *sad*/1.

Given the observations $O = \{sad(ale), sad(kr), not\ sad(oli)\}$ can we improve our model? The incompleteness of our model resides in the predicate *poor*. This is the only abducible predicate in our model. Using abduction we can explain the observations O via the explanation:

$E = \{poor(ale), poor(kr), not\ poor(oli)\}$.

Subsequently, treating this explanation as training data for inductive generalization we can generalize this to get the rule:

poor(X) *if* *lecturer*(X)

thus (partially) defining the abducible predicate *poor* when we extend our theory with this rule.

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 Moyle and Muggleton (1997), Muggleton and Bryant (2000); and Moyle (2000), 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* (Muggleton and Bryant, 2000) can be seen as a particular case of integration of abductive inference for constructing

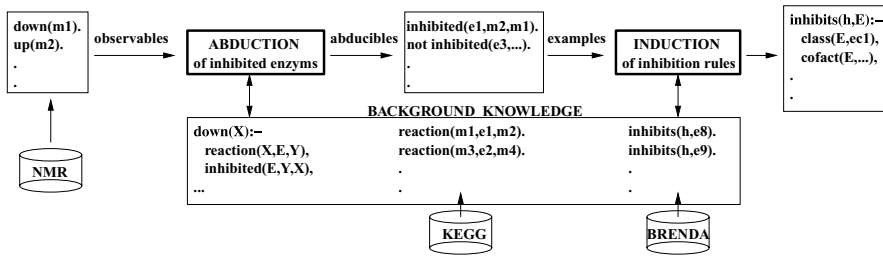


Fig. 2 An Abductive/Inductive framework for modelling inhibition. The purpose of the abduction process is to generate hypotheses about inhibited enzymes from the NMR observations of metabolite concentration. The purpose of the induction process is to learn from the abduced hypotheses, which are ground facts of inhibition, general rules about the inhibition of enzymes in terms of chemical properties of the inhibitor, functional class of enzymes etc

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 (King et al., 2004). In Moyle (2000) Theory Completion is realized in an ILP system, called ALECTO, which again integrates a phase of, *extraction-case abduction*, to transform each case of a training example to an abductive hypothesis with a phase of induction that generalizes these abductive hypotheses. It has been used (Moyle, 2002) to learn robot navigation control programs by completing the specific domain knowledge required, within a general theory of planning that the robot uses for its navigation.

In general, the integration of abduction and induction enhances the model development in several ways. It provides a better opportunity to test the correctness of the generated hypotheses as this can increase the scope of testing. The predictive accuracy (or degree of corroboration) of an abductive explanation can be used as a criterion to select amongst different possible explanations of the observations. But considering the abductive explanations as they are in their extensional form provides a relatively restricted scope for testing of the hypotheses. However, when we integrate this with induction we can get predictions outside the realm of the given observations and theory and thus increase the scope of testing of our hypotheses. In a *tight integration* of abduction and induction the choice of an explanation in the first abductive phase of the cycle is linked to the second phase of how well the explanation generalizes through induction. The abductive explanations and intentional hypothesis generated from them are thus evaluated together where the generalizability (measured by the predictive accuracy but also by other criteria such as the degree of compression) of the explanation feeds back in the choice of abductive explanation. Such frameworks of tight integration of abduction and induction with ILP include: Progol 5.0 (Muggleton and Bryant, 2000), ALECTO (Moyle, 2000), ACL (Kakas and Riguzzi, 2000), SOLDLR (Yamamoto and Fronhöfer, 2001), CF-Induction (Inoue, 2001) and HAIL (Ray, 2003). We will use Progol 5.0 to carry out the experiments in our study in this paper.

For our specific problem domain of modelling the phenomenon of inhibition the cycle of integration of abduction and induction is shown in Fig. 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 which models how the concentration of metabolites (e.g. up-down regulations) is related to inhibition of enzymes (see Section 3.2 for such a model). 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. 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.

3.1. Modelling in abductive logic programming

A framework that allows declarative representations of incomplete theories is that of Abductive Logic Programming (ALP) (Kakas et al., 1993; Kakas and Denecker., 2002). 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 formulae 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. Abducible predicates appear only in the conditions of the program rules with no definition in P . The integrity constraints, IC , represent *assertional knowledge* that we may have about our domain, augmenting the model in P , but without defining any predicates.

Given such an ALP theory the inference of abduction (i.e. of abductive explanation) is then specialized accordingly in the following way:

Definition 1. Given an abductive logic theory (P, A, IC) , an abductive explanation for an observation O , is a set, Δ , of ground abducible atoms on the predicates A such that:

- $P \cup \Delta \models_{LP} O$
- $P \cup \Delta \models_{LP} IC$.

where \models_{LP} denotes the logical entailment relation in Logic Programming⁴.

The abductive explanation Δ represents a hypothesis which when taken together with the model described in the theory T explains how a nonempty experimental observable O could hold. An abductive explanation partially completes the model as described in the theory T . When the observation O is a given training example in a learning problem we can see that such abductive explanations coincide with the concept of extraction-case abduction, as used in Kakas and Riguzzi (2000), Moyle and Muggleton (1997), and Moyle (2002), to transform each case of a training example to an abductive hypothesis.

The important role of the integrity constraints, IC , is to impose *validity* requirements on the abducible hypotheses Δ . They are modularly stated in the theory, separately from the basic model captured in P , and they are used to augment this with any partial information that we may have on the abducible predicates or other particular requirements that we may want the abductively generated explanations of our observations to have. In most practical cases the integrity constraints are of the form of clausal rules: $B_1 \wedge \dots \wedge B_n \rightarrow A_1 \vee \dots \vee A_k$ where A_1, \dots, A_k and B_1, \dots, B_n are positive literals. In these constraints, k can be possibly zero (we will then write the conclusion as **false**) in which case the constraint is a denial prohibiting any set of abducibles that would imply the conjunction B_1, \dots, B_n .

⁴ For example, when the program P contains no negation as failure then this entailment is given by the minimal Herbrand model of the program and the truth of formulae in this model.

An abductive explanation partially completes the model as described in the theory T . As mentioned above, this is a limited conservative extension of the model, due to the extensional nature of these explanations and their restricted scope of prediction to the specific situation(s) of the observations they explain. The subsequent generalization, through induction, of ground abductive explanations to intensional rules for the abducibles provides a greater level of completing the initial theory (model). The rules learned on the abducibles provide (partial) definitions of the abducible predicates, to be added in the program P , thus partially completing the initially incomplete model. We can then iterate the process of scientific analysis of the phenomena, i.e. gather new observations, explain them, generalize the explanations and test the inductive hypotheses thus generated, until we are satisfied from the level of completeness that our model has reached.

3.2. Modelling inhibition in ALP

We will develop a model for analyzing (understanding and subsequently predicting) the effect of toxin substances on the concentration of metabolites. Our representation will use as the set of *observable* predicates the single predicate:

concentration(*Metabolite*, *Level*, *Time*)

expressing the fact that at some time, *Time*, a metabolite, *Metabolite*, has a certain level of concentration, *Level* which 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 Fig. 1. For example, the statement

reactionnode('1 – 2 – amino adipate', '2.6.1.39', '2 – oxo – glutarate')

expresses the fact that there is a direct path (reaction) between the metabolites 1 – 2 – amino adipate and 2 – oxo – glutarate 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 (re-balance) 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(*Enzyme*, *Metabolites1*, *Metabolites2*, *Time*)

capturing the hypothesis that at the time *Time* the reaction from *Metabolites1* to *Metabolites2* is inhibited by the toxin through an adverse effect on the enzyme, *Enzyme*, that normally catalyzes this reaction. For example,

inhibited('2.6.1.39', '1 – 2 – amino adipate', '2 – oxo – glutarate', 8)

expresses the abducible hypothesis that at time 8 the the reaction from $l - 2 - \text{aminoadipate}$ to $2 - \text{oxo} - \text{glutarate}$ via the enzyme 2.6.1.39 is inhibited by the toxin.

Hence the set of abducibles, A , in our ALP theory (P, A, IC) , contains the only predicate *inhibited*/4. 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 non-observable inhibitory effect of the toxin represented by this abducible predicate.

Given this vocabulary of predicates for our theory (P, A, IC) , we now need to provide the program rules in P and the integrity constraints IC that together model our problem domain. The rules in P describe an underlying mechanics of the effect of inhibition of a toxin by defining the observable *concentration*/3 predicate. This model is simple in the sense that it only describes at an appropriate high-level the possible inhibition effects of the toxin. The appropriateness of the abstraction is judged on two levels. On the one hand the model does not aim to model the low level chemical mechanism with which the inhibition occurs and hence it abstracts away from the complex details of the biochemical properties of the metabolic reactions. This also means that all the steps in the chains of reactions in the metabolic network are not needed provided that we do not change the topology of the network. On the other hand the model needs to be able to account for the connectivity of the reactions and the fact that the effects of inhibition can propagate in the network.

The underlying and simplifying working hypotheses of our model are the following:

- (1) the primary inhibitory effect of the toxin can be *localized* on the individual reactions of the metabolic pathways;
- (2) the underlying network of the metabolic pathways is correct and complete;
- (3) all the reactions of the metabolic pathways are a-priori equally likely to be affected by the toxin;
- (4) inhibition in one reaction is sufficient to cause change in the concentration of the metabolites;
- (5) the time delay needed for the toxin to begin its inhibitory effect is negligible.

The program P sets out general laws under which the effect of the toxin can increase or reduce the concentration of the metabolites. These are based on the simple idea that when a reaction from one metabolite X to another metabolite Y is inhibited then this causes a slowing down in the production of Y from X . Hence the concentration of Y decreases as less of Y is produced and at the same time the concentration of X increases as X is no longer consumed to produce Y . In other words, X starts to accumulate due to the inhibition of the reaction that transforms it to Y . The rules in P that capture these basic causal effects of inhibition are:

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

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

In addition, our simple model allows for knock on 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). In other words, the down regulation of a metabolite X can be caused by the inhibition of

some reaction upstream from X , that is not necessarily directly producing X , by repeatedly causing the down regulation of the intermediate metabolites from this reaction down to X . Analogously, a metabolite X can have increased concentration (i.e. concentration “up”) when another metabolite, Y , upstream from X has been cause to have increased concentration, an a knock on indirect effect of increased concentration from Y to X . The rules in P that capture this aspect of indirect effects are:

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

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

Both these rules contain the extra condition that the reaction linking the two metabolites is not inhibited since if the reaction from Y to X is inhibited this would block the knock on effect: in the first rule because the inhibition would cause as a direct effect (as we have described above) the “up” concentration of Y thus invalidating the condition that Y has “down” concentration and in the second rule because the inhibition would block the flow from Y to X even when the concentration of Y is “up”.

Note also that in all these rules of P we are assuming that the effect of the inhibition is present immediately with no delay and the effect remains at any time for which the reaction is inhibited. Indeed, over a period of time it is possible for the inhibition of some reaction to terminate and as a result the concentration of some of the metabolites to change over time from “up” to “down” or vice versa.

The above rules and working hypotheses give a relatively simple model that we have taken as a starting point for our study. 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 program P rules that express this. For example, the first rule above would be replaced by:

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

compensated(X,Enz,T):-
reactionnode(X,EnzI,Y),
different(EnzI,Enz),
increased(EnzI,Y,X,T).

where now the set of abducible predicates A includes also the predicate *increased(Enzyme, Metabolites1, Metabolites2, Time)* that captures the assumption that the flow of the reaction from *Metabolites1* to *Metabolites2* has increased at time T as a secondary effect of the presence of the toxin.

3.2.1. Validity requirements of the model

The abducible information of *inhibited/4* is required to satisfy several *validity requirements* captured in the integrity constraints *IC* of the model. The semantics of the Integrity Constraints as given by the requirement, $P \cup \Delta \models_{LP} IC$ for any abductive explanation Δ , gives them a role of filtering out any explanation that does not have the properties stated by them. They are not used for prediction, a task given solely to the program P and its extensions: they merely restrict the possible extensions of P . This allows the Integrity Constraints to be stated modularly in *IC*, separately from the program P and in many cases they can be changed without the need to reconsider the underlying model of P .

The integrity constraints typically involve general self-consistency requirements of the model such as:

$$\text{concentration}(X, \text{down}, T), \text{concentration}(X, \text{up}, T) \rightarrow \text{false}$$

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/4*, such as that a certain reaction is not reversible can be captured as a validity requirement. For example the integrity constraint:

$$\text{inhibited}(2.3.1.61, \text{'succinate'}, \text{'2-oxo-glutarate'}, T) \rightarrow \text{false}$$

prevents us from ever assuming that the reaction from ‘succinate’ to ‘2-oxo-glutarate’ is inhibited as we know that this reaction between these two metabolites in the pathway never operates in this direction.

Other such constraints can help us restrict further the form of the abductive explanations that we are looking for, essentially adding in this way extra working hypotheses to our model. We could, for example, be interested only in explanations whose inhibition effects are separated apart on the pathways network. This would be captured by an integrity constraint of the form:

$$\text{inhibited}(\text{Enz}, X, Y, T), \text{inhibited}(\text{Enz}1, Y1, Z, T), \text{close}(Y, Y1) \rightarrow \text{false}$$

where the auxiliary background predicate $\text{close}(Y, Y1)$ is defined in P and holds true iff the shortest distance between the two metabolite nodes Y and $Y1$ is smaller than a given minimum distance.

3.2.2. Example explanations

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

$$\text{concentration}(\text{'2-oxo-glutarate'}, \text{down}, 8)$$

the following are some of its possible explanations

$$E_1 = \{\text{inhibited}(2.3.1.61, \text{'succinate'}, \text{'2-oxo-glutarate'}, 8)\}$$

$$E_2 = \{\text{inhibited}(2.6.1.39, \text{'l-2-aminoadipate'}, \text{'2-oxo-glutarate'}, 8)\}$$

$$E_3 = \{\text{inhibited}(1.1.1.42, \text{'isocitrate'}, \text{'2-oxo-glutarate'}, 8)\}$$

Combining this observation with the additional observation that

$$\text{concentration}(\text{'isocitrate'}, \text{down}, 8)$$

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('l - 2 - aminoadipate', 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-oxo-glutarate' down) and third observation (of 'l-2-aminoadipate' 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 'l-2-aminoadipate' to '2-oxo-glutarate'.

An alternative explanation would be

$$E'_2 = \{\textit{inhibited}(2.6.1.39, 'l - 2 - aminoadipate', '2 - oxo - glutarate', 8), \\ \textit{inhibited}(1.2.1.31, 'l - 2 - aminoadipate', 'l - lysine', 8)\}.$$

Applying a principle of *minimality* of explanations (Kakas et al., 1993) or more generally of *maximal compression* we would prefer the explanation E_2 over E'_2 .

3.2.3. Computing explanations by ALP and ILP systems

There are several systems (e.g. ACLP⁵, A-System⁶) for computing abductive explanations in ALP. Also some ILP systems, such as ALECTO and Progol 5, can compute abductive explanations as well as generalisations of these. Most ALP systems, unlike ILP systems, do not employ an automatic way of comparing different explanations at generation/search time and selecting from these those explanations that satisfy some criterium of compression or simplicity. On the other hand, ALP systems can operate on a richer representation language, e.g. that includes negation as failure. Hence although ILP systems can provide compact and minimal explanations ALP systems can provide explanations that have a more complete form. ALP explanations can contain extra hypotheses that are generated from ensuring that the integrity constraints are satisfied. Often these extra hypotheses are negative requiring that some abducible fact is false. Such hypotheses are left implicit in the explanations of ILP systems, where the integrity constraints are checked a-posteriori, not while the explanation is generated as in ALP systems.

Lets us illustrate the link and comparison between the ALP and ILP explanations by considering an example. Referring again to Fig. 1, a Progol 5 explanation for the two observations for metabolites *l - 2 - aminoadipate* and *succinate* would be:

$$E_{ILP} = \{\textit{inhibited}(2.6.1.39, \textit{true}, 'l - 2 - aminoadipate', '2 - oxo - glutarate', 8), \\ \textit{inhibited}(1.2.7.3, \textit{false}, '2 - oxo - glutarate', 'succinate', 8)\}.$$

This explanation does not carry any information on the rest of the network that is not directly connected with the observations. In particular, it does not give explicit information on how the integrity constraints (e.g. that the concentration of any metabolite in the network can not be predicted to be both “up” and “down”) are satisfied. The corresponding ALP explanation(s) have the form:

$$E_{ALP} = \{\textit{inhibited}(2.6.1.39, 'l - 2 - aminoadipate', '2 - oxo - glutarate', 8), \\ \textit{not inhibited}(1.2.7.3, '2 - oxo - glutarate', 'succinate', 8)\} \cup E_{Rest}.$$

⁵ <http://www.cs.ucy.ac.cy/aclp/>.

⁶ <http://www.cs.kuleuven.ac.be/~bertv/Assystem/>.

where E_{Rest} makes explicit further assumptions required for the satisfaction of the integrity constraints. In this example, if we are interested in the metabolite *isocitrate* then we could have two possibilities:

$$E_{Rest}^1 = \{not\ inhibited(1.1.1.42., '2 - oxo - glutarate', 'isocitrate', 8), \\ not\ inhibited(1.1.1.42., 'isocitrate', '2 - oxo - glutarate', 8)\}.$$

$$E_{Rest}^2 = \{not\ inhibited(1.1.1.42., '2 - oxo - glutarate', 'isocitrate', 8), \\ inhibited(1.1.1.42., 'isocitrate', '2 - oxo - glutarate', 8)\}.$$

These extra assumptions are left implicit in the ILP explanations as they have their emphasis on maximal compression. But the predictions that we get from the two types of ALP and ILP explanations are the same. Both types of explanations predict *concentration('2 - oxo - glutarate', down, 8)*. For *isocitrate* the first ALP explanation predicts this to have down concentration whereas the second one predicts this to have up concentration. The non-committal corresponding ILP explanation will also give these two possibilities of prediction depending on how we further assume the flow of the reaction between *2 - oxo - glutarate* and *isocitrate*.

4. 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 which 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 generalising the ground hypotheses. In particular, we test the following null hypothesis:

Null hypothesis: Generalising ground hypotheses, which are generated from the abductive model for inhibition, does not lead to increased predictive accuracy.

4.1. Materials and methods

In this experiment Progol 5.0⁷ is used to generate both ground and non-ground hypotheses. As a part of background knowledge, we use the relational representation of biochemical reactions involved in a metabolic pathway which is affected by the toxin. This information is extracted from KEGG as explained in Section 2. The observable data is up/down regulation of metabolites obtained from NMR spectra. The technique which has been used to transform raw time-series data is described in Crockford et al. (2005). The up/down regulations of metabolites at different time periods are then encoded as Prolog ground facts⁸. Background knowledge required for non-ground hypotheses can be obtained from databases such as BRENDA and LIGAND⁹ as discussed in Section 2. This background information can include information about enzyme classes, co-factors etc. In our experiments for learning non-ground hypotheses for hydrazine and ANIT, we include the possibility that a given chemical compound can be inhibiting a whole enzymatic class, since this situation is possible

⁷ Available from: <http://www.doc.ic.ac.uk/~shm/Software/progol5.0/>

⁸ Available from: <http://ilp.doc.ic.ac.uk/datasets/inhibition/>

⁹ <http://www.genome.ad.jp/ligand/>.

```

for  $t$  in (8hrs, 24hrs, 48hrs, 72hrs, 96hrs) do
  for  $j=1$  to  $n$  do
     $T_{s_{tj}}$  = a randomly selected example from  $E_t$ 
     $Tr_{tj} = E_t - T_{s_{tj}}$ 
    for  $i$  in (4,9,14,19) do
       $Tr_{tji} = i$  training example randomly sampled from  $Tr_{tj}$ 
    end
  end
end
for  $t$  in (8hrs, 24hrs, 48hrs, 72hrs, 96hrs) do
  for  $j=1$  to  $n$  do
    for  $i$  in (4,9,14,19) do
       $H_{tji}$  = learned hypotheses using the training set  $Tr_{tji}$ 
       $A_{tji}$  = predictive accuracy of  $H_{tji}$  on the test set  $T_{s_{tj}}$ 
    end
  end
end
for  $i$  in (4,9,14,19) do
  Plot average of  $A_{tji}$  versus  $i$  ( $j \in [1..n]$  and  $t \in [8, 24, 48, 72, 96]$ )

```

Fig. 3 Experimental method. E_t is the set of NMR data at time t and in this experiment $n = |E_t| = 20$. Variable $i = (4, 9, 14, 19)$ represents the varying size of a randomly chosen training set in a leave-one-out test strategy

in non-competitive 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 hrs to 96 hrs 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. The experimental method is detailed in Fig. 3.

The model which has been used for evaluating the hypotheses generated by Progol explicates the Closed World Assumption (CWA). In other words, we are working under the assumption that a reaction is not inhibited unless we have a fact which says otherwise:

```

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

```

The predictor which we have used in our experiments converts the three class problem which 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:

```

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

concentrationI(X,down,T):-
  not(concentrationI(X,up,T)).

```

4.2. Results and discussion

According to our model, there are many possible hypotheses which 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 time 8 hrs:

```

inhibited('2.6.1.39',true,'l-2-aminoadipate','2-oxo-glutarate',8).
  inhibited('2.3.1.61',false,'2-oxo-glutarate','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','mao',8).
  inhibited('4.3.1.6',true,'beta-alanine','acryloyl-coA',8).
  inhibited('4.3.2.1',true,'l-as','fumarate',8).

```

Examples of these ground hypotheses are illustrated in Fig. 4. 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, one of the hypothesised enzymes (i.e. EC2.6.1.39) was already 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 (Nicholls, 2001). Experimental evaluations *in vivo* are required to test this hypothesis.

The overall performance of ground and non-ground hypotheses are shown in Fig. 5. 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 non-ground rule is:

```

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

```

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

According to the comparison shown in Fig. 5, 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 small number of training examples, ground hypotheses (with closed world assumption) have a better predictive accuracy than non-ground hypotheses. These results suggest that for a small number of training examples (e.g. less than 45%) the induced non-ground 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%), non-ground hypotheses show a better performance than ground hypotheses (with a significance level of 0.01).

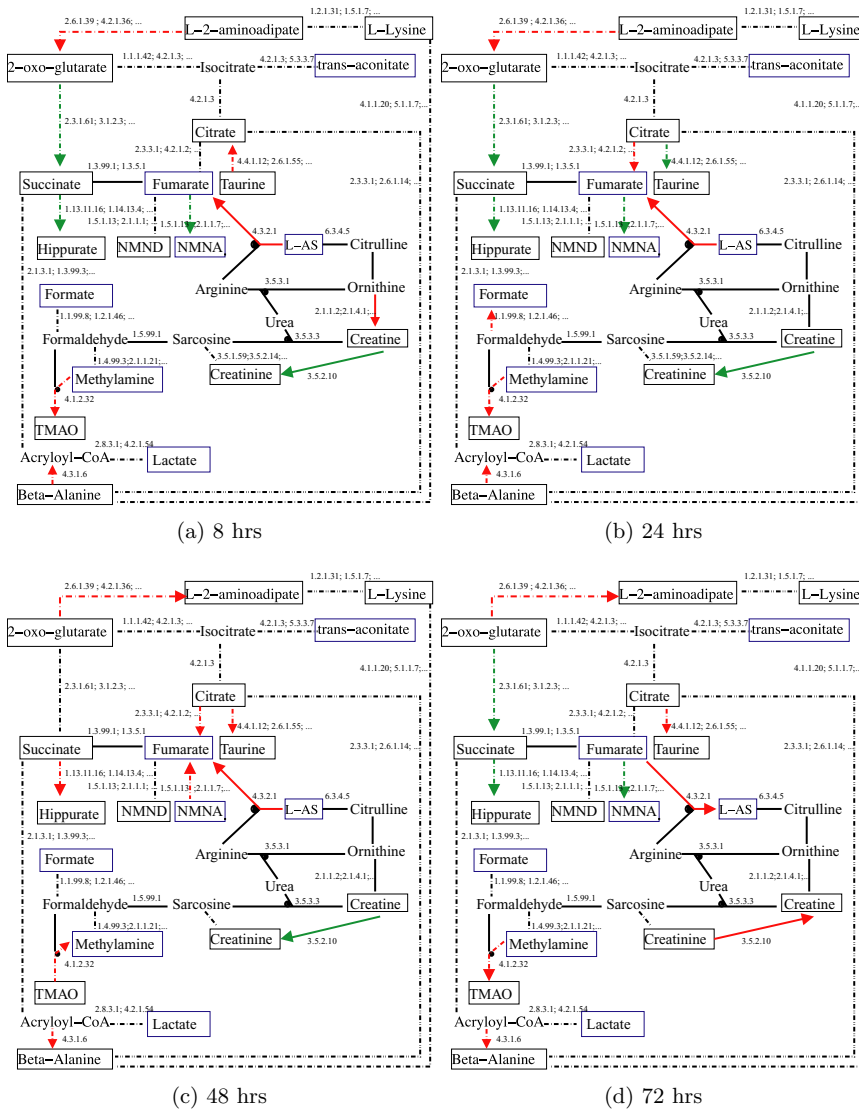



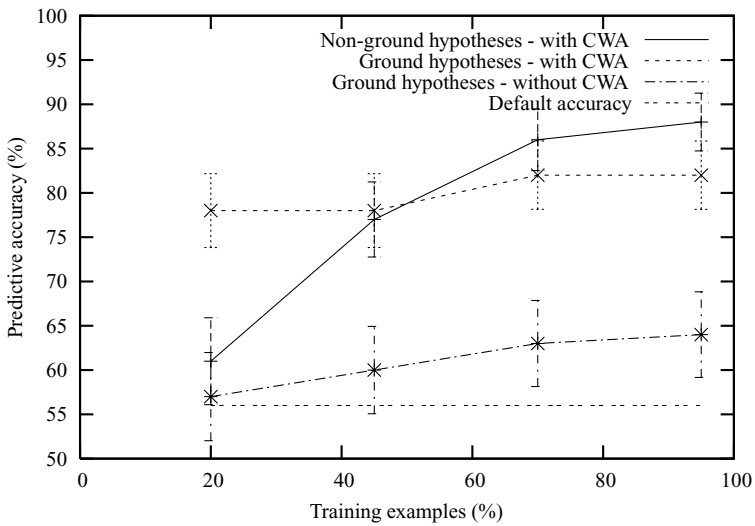
Fig. 4 Examples of ground hypotheses for Hydrazine at 8hrs, 24hrs, 48hrs and 72hrs. Red arrows correspond to *inhibited* and green arrows correspond to *not inhibited* hypotheses. The model suggests that some reactions remain inhibited through different time periods

5. Related work

The abduction technique which is used in this paper can be compared with the one in the robot scientist project (King et al., 2004) where Progol5.0 was used to generate ground hypotheses about the function of genes. Abduction has been also used within a system, called GenePath (Zupan et al., 2003), to find relations from experimental genetic data in order to facilitate the analysis of genetic networks. Similarly, in Papatheodorou (2005) abduction has been used to generate gene interactions and genetic pathways from microarray experimental data.



(a) Hydrazine



(b) ANIT

Fig. 5 Performance of ground and non-ground hypotheses generated by Progol using a leave-one-out test strategy (as detailed in Fig. 3) for a) Hydrazine and b) ANIT

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 (Wellner, 1999; Moyle, 2002). In Moyle (2002) a framework of Theory Completion (Moyle, 2000), that also integrates abduction and induction, has been applied to this domain (i.e. learning robot navigation control theories) by completing Event Calculus theories (Moyle and Muggleton, 1997) for temporal reasoning.

Bayesian networks are among the most successful techniques which have been used for modelling biological networks. In particular, gene expression data has been widely modelled

using Bayes' net techniques (Friedman et al., 1998; Friedman et al., 2000; Imoto et al.) On the MetaLog project, Bayes' nets have also been used to model metabolic networks (Tamaddoni-Nezhad et al., 2003). A key advantage of the logical modelling approach in the present paper 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 modelling approach also produces explicit hypotheses concerning the inhibitory effects of toxins.

A number of classical mathematical approaches to metabolic pathway analysis and simulation exist. These can be divided into three main groups based around Biochemical Systems Theory (BST), Metabolic Control Analysis (MCA) and Flux Balance Analysis (FBA). BST and MCA are oriented toward dynamic simulation of cellular processes based on physico-chemical laws (Goodwin, 1963; Hess and Boiteux, 1968; Tyson and Othmer, 1978). However, progress towards the ultimate goal of complete simulation of cellular systems (Tyson and Othmer, 1978) has been impeded by the lack of kinetic information and attention in the last decade has been diverted to analysing the relative importance of metabolic events. FBA (Varma and Palsson, 1994; Edwards et al., 1999) unlike BST and MCA, does not require exact kinetic information to analyse the operative modes of metabolic systems. FBA, which includes the techniques of Elementary Flux Mode Analysis and Extreme Pathway Analysis, only requires stoichiometric parameters (the quantitative relationship between reactants and products in a chemical reaction). However, by contrast with the approach taken in the present paper, BST, MCA and FBA are not machine learning approaches, and most importantly do not incorporate techniques for extending the structure of the model based on empirical data.

6. Conclusions

In this paper we have studied how to use abduction and induction in scientific modelling and have concentrated 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 auto-regressive assumptions used in other temporal modelling approaches and demonstrates the strength and flexibility of the Abductive ILP approach for dealing with such problems. Our theories use a simplified model of temporal reasoning and hence there is scope for extending our approach by employing a more advanced time model such as that of the Event Calculus (as used in Moyle (2002)).

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 non-ground cases. These hypotheses were also evaluated in terms of biological insight provided. Experimental evaluations *in vivo* are required to test some of these hypotheses.

In the present study we used simple background knowledge concerning the class of enzymes to allow the construction of non-ground 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 (Muggleton et al., 2003).

Acknowledgements We would like to thank the anonymous reviewers for their comments. We also thank M. Sternberg, J. Nicholson, E. Holmes and F. Pazos 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 2003/4 during which this work had started.

References

- Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., & Watson, J.D. (1994). *Molecular Biology of the Cell*. Garland Publishing, Inc., third edition.
- Alm, E., & Arkin, A.P. (2003). Biological networks. *Curr. Opin. Struct. Biol.*, 13(2), 193–202.
- Alves, R., Chaleil, R.A., & Sternberg, M.J. (2002). Evolution of enzymes in metabolism: a network perspective. *Mol. Biol.*, 320(4), 751–70.
- Boutaud, O., Aronoff, D., Richardson, J., Marnett, L., & Oates, J. (2002). Determinants of the cellular specificity of acetaminophen as an inhibitor of prostaglandin H(2) synthases. *Proc Natl Acad Sci USA*, 99(10), 7130–7135.
- Crockford, D.J., Keun, H.C., Smith, L.M., Holmes, E., & Nicholson, J.K. (2005). Curve-fitting method for direct quantitation of compounds in complex biological mixtures using ¹H nmr: Application in metabonomic toxicology studies. *Anal. Chem.*, 77(14), 4556–4562.
- Edwards, J.S., Ramakrishna, R., Schilling, C.H. & Palsson, B.O. (1999). Metabolic flux balance analysis. In: Lee, S.Y. & Papoutsakis, E.T. editors, *Metabolic Engineering*. Marcel Deker.
- Flach, P.A., & Kakas, A.C., editors. (2000). *Abductive and Inductive Reasoning*. Pure and Applied Logic. Kluwer.
- Flach, P., & Kakas, A.C. (2000). Abductive and inductive reasoning: Background and issues. In: Flach, P.A., & Kakas, A.C., editors, *Abductive and Inductive Reasoning*, Pure and Applied Logic. Kluwer.
- Nir Friedman, Michal Linial, Iftach Nachman, & Dana Pe'er. (2000). Using bayesian networks to analyze expression data. *J. of Comp. Bio.*, 7, 601–620.
- Friedman, N., Murphy, K., & Russell, S. (1998). Learning the structure of dynamic probabilistic networks. In: *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference (UAI-1998)*, (pp. 139–147), San Francisco, CA. Morgan Kaufmann Publishers.
- Goodwin, B.C. (1963). *Oscillatory organization in cells, a dynamic theory of cellular control processes*. Academic Press, New York.
- Harstshorne, C., Weiss, P., & Burks, A., editors. (1958). *Collected Papers of Charles Sanders Peirce*. Harvard University Press.
- Hempel, C.G. (1965). *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*. Free Press, New York.
- Hess, B., & Boiteux, A. (1968). Oscillatory organization in cells, a dynamic theory of cellular control processes. *Hoppe-Seylers Zeitschrift für Physiologische Chemie*, 349, 1567–1574.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1989). *Induction: Processes of Inference, Learning and Discovery*. MIT Press.
- Imoto, S., Goto, T., & Miyano, S. Estimation of genetic networks and functional structures between genes by using bayesian networks and nonparametric regression. In: *Proceeding of Pacific Symp. on Biocomp.*
- Inoue, K. (2001). Induction, abduction and consequence-finding. In: *Proceedings of the International Workshop on Inductive Logic Programming (ILP01)*, (pp. 65–79). Springer-Verlag.
- Inoue, K. (2001). Inverse entailment for full clausal theories. In: *LICS-2001 Workshop on Logic and Learning*.
- Ito, K., & Yamamoto, A., (1998). Finding hypotheses from examples by computing the least generalisation of bottom clauses. In: *Proceedings of Discovery Science '98*, (pp. 303–314). Springer.
- Jeong, H., Tombor, B., Albert, R., Oltvai, Z.N., & Barabasi, A.L. (2000). The large-scale organization of metabolic networks. *Nature*, 407(6804), 651–654.
- Josephson, J.R., & Josephson, S.G., editors. (1994). *Abductive Inference: Computation, Philosophy, Technology*. Cambridge University Press.
- Kakas, A.C., & Riguzzi, F. (2000). Abductive concept learning. *New Generation Computing*, 18, 243–294.
- Kakas, A.C., & Denecker, M. (2002). Abduction in logic programming. In: Kakas, A.C., & Sadri, F. editors, *Computational Logic: Logic Programming and Beyond. Part I*, number 2407, (pp. 402–436).
- Kakas, A.C., Kowalski, R.A., & Toni, F. (1993). Abductive Logic Programming. *Journal of Logic and Computation*, 2(6), 719–770.
- King, R.D., Whelan, K.E., Jones, F.M., Reiser, P.K.G., Bryant, C.H., Muggleton, S.H., Kell, D.B., & Oliver, S.G. (2004). Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427, 247–252.

- Lindon, J.C., Nicholson, J.K., Holmes, E., & et al. (2003) Contemporary issues in toxicology. the role of metabonomics in toxicology and its evaluation by the comet project. *Tox Appl Pharm*, 187, 137–146.
- Magnani, L. (2001). *Abduction, Reason and Science*. Kluwer Academic/Plenum Publishers.
- MetaLog: <http://www.doc.ic.ac.uk/bioinformatics/metalog/>.
- Moyle, S.A. (2000). *An investigation into Theory Completion techniques in Inductive Logic Programming*. PhD thesis, Oxford University Computing Laboratory, University of Oxford.
- Moyle, S., & Muggleton, S.H. (1997). Learning programs in the event calculus. In: *Proceedings of the Seventh Inductive Logic Programming Workshop (ILP97)*, LNAI 1297, (pp. 205–212). Springer-Verlag.
- Moyle, S. (2002). Using theory completion to learn a robot navigation control program. In: *Proceedings of the 12th International Conference on Inductive Logic Programming*, (pp. 182–197). Springer-Verlag.
- Muggleton, S.H., & Bryant, C.H. (2000). Theory completion using inverse entailment. In: *Proc. of the 10th International Workshop on Inductive Logic Programming (ILP-00)*, (pp. 130–146), Berlin. Springer-Verlag.
- Muggleton, S.H., Tamaddoni-Nezhad, A., & Watanabe, H. (2003). Induction of enzyme classes from biological databases. In: *Proceedings of the 13th International Conference on Inductive Logic Programming*, (pp. 269–280). Springer-Verlag.
- Muggleton, S., & De Raedt, L. (1994). Inductive logic programming: Theory and methods. *Journal of Logic Programming*, 19(20), 629–679.
- Muggleton, S. (1991). Inductive logic programming. *New Generation Computing*, 8(4), 295–318.
- Nicholls, A.W., Holmes, E., Lindon, J.C., Farrant, R.D., Haselden, J.N., Damment, S.J.P., Waterfield, C.J., & Nicholson, J.K. (2001). Metabonomic investigations into hydrazine toxicity in the rat. *Chemical Research in Toxicology*, 14(8), 975–987.
- Ogata, H., Goto, S., Sato, K., Fujibuchi, W., Bono, H., & Kanehisa, M. (1999). KEGG: Kyoto Encyclopedia of Genes and Genomes. *Nucleic Acids Res.*, 27(1), 29–34.
- Papathodorou, I., Kakas, A., & Sergot, M. (2005). Inference of gene relations from microarray data by abduction. Technical Report 2005/3, Department of Computing, Imperial College London.
- Papin, J.A., Price, N.D., Wiback, S.J., Fell, D.A., & Palsson, B.O. (2003). Metabolic pathways in the post-genome era. *Trends Biochem. Sci.*, 28(5), 250–8.
- Peirce, C.S. (1957). *Essays in the Philosophy of Science*. Liberal Arts Press.
- Popper, K. (1959). *The Logic of Scientific Discovery*. Basic Books, New York.
- Ravasz, E., Somera, A.L., Mongru, D.A., Oltvai, Z.N., & Barabasi, A.L. (2002). Hierarchical organization of modularity in metabolic networks. *Science*, 297(5586), 1551–5.
- Ray, O., Broda, K., & A. Russo. (2003). Hybrid Abductive Inductive Learning: a Generalisation of Progol. In: *13th International Conference on Inductive Logic Programming*, volume 2835 of *LNAI*, (pp. 311–328). Springer Verlag.
- Świerkosz, T.A., Jordan, L., McBride, M., McGough, K., Devlin, J., & Botting, R. M. (2002). Actions of paracetamol on cyclooxygenases in tissue and cell homogenates of mouse and rabbit. *Med Sci Monit*, 8(12), 496–503.
- Tamaddoni-Nezhad, A., Kakas, A., Muggleton, S.H., & Pazos, F. (2004). Modelling inhibition in metabolic pathways through abduction and induction. In: *Proceedings of the 14th International Conference on Inductive Logic Programming*, (pp. 305–322). Springer-Verlag.
- Tamaddoni-Nezhad, A., Muggleton, S., & Bang, J. (2003). A bayesian model for metabolic pathways. In: *International Joint Conference on Artificial Intelligence (IJCAI03) Workshop on Learning Statistical Models from Relational Data*, (pp. 50–57). IJCAI.
- Tyson, J.J., & Othmer, H.G. (1978). The dynamics of feedback control circuits in biochemical pathways. *Progress in Theoretical Biology*, 5, 1–62.
- Varma, A., & Palsson, B.O. (1994). Metabolic flux balancing: Basic concepts, scientific and practical use. *BioTechnology*, 12, 994–998.
- Wellner, B.R. (1999). An abductive-inductive learning framework for logic-based agents. Master's thesis, Imperial College of Science Technology and Medicine.
- Yamamoto, A., & Fronhöfer, B. (2001). Finding hypotheses by generalizing residues hypotheses. In: *Proceedings of the International Workshop on Inductive Logic Programming (ILP01)*, Springer-Verlag (pp. 107–118).
- Yamamoto, A. (1997). Which hypotheses can be found with inverse entailment? In: *Proceedings of the Seventh International Workshop on Inductive Logic Programming*, Berlin. LNAI 1297 (pp. 296–308).
- Zimmerman, H. J., & Maddrey, W. C. (1995). Acetaminophen (paracetamol) hepatotoxicity with regular intake of alcohol: analysis of instances of therapeutic misadventure. *Hepatology*, 22(3), 767–773.
- Zupan, B., Bratko, I., Demsar, J., Juvan, P., Halter, J.A., Kuspa, A., & Shaulsky, G. (2003). Genepath: a system for automated construction of genetic networks from mutant data. *Bioinformatics*, 19(3), 383–389.