

Completing causal networks by meta-level abduction

Katsumi Inoue · Andrei Doncescu ·
Hidetomo Nabeshima

Received: 26 July 2010 / Accepted: 8 March 2013 / Published online: 3 April 2013
© The Author(s) 2013

Abstract Meta-level abduction is a method to abduce missing rules in explaining observations. By representing rule structures of a problem in a form of causal networks, meta-level abduction infers missing links and unknown nodes from incomplete networks to complete paths for observations. We examine applicability of meta-level abduction on networks containing both positive and negative causal effects. Such networks appear in many domains including biology, in which inhibitory effects are important in several biological pathways. Reasoning in networks with inhibition involves nonmonotonic inference, which can be realized by making default assumptions in abduction. We show that meta-level abduction can consistently produce both positive and negative causal relations as well as invented nodes. Case studies of meta-level abduction are presented in p53 signaling networks, in which causal relations are abducted to suppress a tumor with a new protein and to stop DNA synthesis when damage has occurred. Effects of our method are also analyzed through experiments of completing networks randomly generated with both positive and negative links.

Keywords Abduction · Theory completion · Meta-reasoning · Causal network · Inhibition · Predicate invention · Default reasoning · Inductive logic programming · Systems biology

Editor: Paolo Frasconi.

This research was supported in part by the 2008-2011 JSPS Grant-in-Aid for Scientific Research (A) No. 20240016.

K. Inoue (✉)

National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan
e-mail: inoue@nii.ac.jp

A. Doncescu

LAAS-CNRS UPR 8001, Avenue du Colonel Roche, 31007 Toulouse, France
e-mail: andrei.doncescu@laas.fr

H. Nabeshima

Division of Medicine and Engineering Science, University of Yamanashi, 4-3-11 Takeda, Kofu, Yamanashi 400-8511, Japan
e-mail: nabesima@yamanashi.ac.jp

1 Introduction

Abduction and induction are both ampliative reasoning, and play essential roles in knowledge discovery and development in science and technology. Integration of abduction and induction has been discussed in such diverse aspects as implementation of inductive logic programming (ILP) systems using abductive methods (Yamamoto 2000; Inoue 2004; Ray et al. 2003; Ray and Inoue 2008; Kimber et al. 2009; Corapi et al. 2010) and “closing the loop” methodologies in scientific discovery (Flach and Kakas 2000; King et al. 2004; Ray et al. 2010). The use of prior or *background knowledge* in scientific applications has directed our attention to *theory completion* (Muggleton and Bryant 2000) rather than classical learning tasks such as concept learning and classification. There, *abduction* is mainly used to complete proofs of observations from *incomplete* background knowledge, while *induction* refers to generalization of abduced cases.

In scientific domains, background knowledge is often structured in a form of *networks*. In biology, a sequence of signalings or biochemical reactions constitutes a network called a *pathway*, which specifies a mechanism to explain how genes or cells carry out their functions. However, information of biological networks in public-domain databases is generally incomplete in that some details of reactions, intermediary genes/proteins or kinetic information are either omitted or undiscovered. To deal with incompleteness of pathways, we need to predict the status of relations which is consistent with the status of nodes (Tamaddoni-Nezhad et al. 2006; Yamamoto et al. 2010; Ray 2009; Inoue et al. 2009; Ray et al. 2010), or insert missing arcs between nodes to explain observations (Zupan et al. 2003; King et al. 2004; Tran and Baral 2009; Akutsu et al. 2009). These goals are characterized by abduction as theory completion, in which status of nodes or missing arcs are added to account for observations.

If a network is represented in a logical theory, inference on the network can be realized on a language in which each network element such as a node and a link is itself an entity of the language. Such a theory (or program) refers a language of the network, and is hence regarded as a *meta-theory* (or *meta-program*), while a theory (or program) representing networks is referred to as an *objective theory* (or *objective program*). Then, to perform abduction on networks, we need abduction on meta-theories, which is referred to as *meta-level abduction*. Note that *meta-reasoning* has been intensively investigated in logic programming, e.g., Kowalski (1990), Hill and Gallagher (1998), Costantini (2002), yet main inference considered there has been deduction.

Meta-level abduction was introduced in Inoue et al. (2010) as a method to discover unknown relations from incomplete networks. Given a network representing causal relations, called a *causal network*, missing links and nodes are abduced in the network to account for observations. The main objective in Inoue et al. (2010) is to provide a logical foundation and knowledge representation for *abducting rules*, which is an important abductive problem in ILP. It is notable that other abductive methods (Poole 1988; Corapi et al. 2010) need a predetermined set of candidate rules (called *abducible rules* or *abducibles*) and then select consistent combinations of abducibles to form explanations. In contrast, meta-level abduction does not need any such abducibles in advance. Meta-level abduction is implemented in SOLAR (Nabeshima et al. 2003, 2010), an automated deduction system for consequence finding, using a first-order representation for algebraic properties of causality and a full-clausal form of network information and constraints. Meta-level abduction by SOLAR is powerful enough to infer missing rules, missing facts, and unknown

causes involving *predicate invention* (Muggleton and Buntine 1988) in the form of existentially quantified hypotheses. Note that predicate invention had been intensively investigated in the initial stage of ILP research and has recently been revisited (Muggleton et al. 2012), since it should play an important role in discovery.

Meta-level abduction has been applied to discover physical skills in terms of *hidden rules* to explain given *empirical rules* in cello playing examples in Inoue et al. (2010), and a thorough experimental analysis with a variety of problem instances has been presented in Nabeshima et al. (2010). However, all those examples of meta-level abduction in Inoue et al. (2010) and Nabeshima et al. (2010) contain only one kind of causal effects, which are positive, and it was left open how to deal with *both positive and negative effects*. Then, we shall examine applicability of meta-level abduction to deal with networks expressing both positive and negative causal effects. Such networks are often used in biology, where *inhibitory effects* are essential in gene regulatory, signaling and metabolic networks.

In this paper, we present axioms for meta-level abduction to produce both positive and negative causal relations as well as newly invented nodes. We show two axiomatizations for such meta-level abduction. One is a set of *alternating axioms* which define relations of positive and negative causal effects in a double inductive manner. This axiom set reduces to the axioms for ordinary meta-level abduction defined in Inoue et al. (2010) when there is no negative causal link. The other axiomatization is a variant of the alternating axioms, but prefers negative links to positive ones if both are connected to the same node. In this case, reasoning in causal networks becomes *nonmonotonic*, and involves default assumptions in abduction.

Then, applications to p53 signal networks (Prives and Hall 1999; Meek 2009) are presented as case studies of our framework, in which meta-level abduction reproduces theories explaining how tumor suppressors work (Tran and Baral 2009) and how DNA synthesis stops (Shmulevich et al. 2002). Analysis of such abstract signaling networks, although simple, provides one of the most fundamental inference problems in computer-aided scientific research including Systems Biology: *Given an incomplete causal network, infer possible connections and functions of network entities to reach the target entities from the sources*. Meta-level abduction in this paper is crucial in this task for the following reasons. First, suggestion of possible additions in prior networks enables scientists to conduct hypothesis-driven experiments with those focused cases. If a suggested hypothesis is justified through a through set of experiments, the corresponding new links and/or nodes are considered to be discovered. In network completion, however, the larger the network becomes, the more abductive inference steps are required to get a hypothesis and the more candidate hypotheses are inferred. Then, it is hard for human scientists to consider all possibilities without losing any important ones. Therefore, automation of hypothesis enumeration is very important. Second, abduction in such network domains often involves a *goal* rather than an *observation*: A hypothesis is inferred to achieve the goal that has not been observed yet. For example, in drug design and pharmacology as well as therapeutic research, the effect of introduction of new entities and links to a known network is goal-oriented and the same hypotheses cannot be applied to other goals in general. This feature of goal-oriented abduction also exists in completing causal networks for improvement of physical techniques in musical performance (Inoue et al. 2010), in which specific skills are required for requested tasks.

Finally, scalability of meta-level abduction is analyzed through experiments of completing networks randomly generated with both positive and negative links. By varying the average degree of nodes to edges, from sparse to dense networks are generated with several

sizes. We will see that it is not easy to generate a large and dense network by keeping the consistency, since there are more chances to become inconsistent in such networks. Then the growth rate of hypotheses in the size of networks rather decreases in dense networks.

This is an extended version of the paper (Inoue et al. 2011), and contains several technical details such as an abductive procedure based on consequence finding, theoretical correctness for the proposed formalization and their proofs, detailed analysis of experiments on p53 pathways, and scalability issues. The rest of this paper is organized as follows. Section 2 offers the essential of meta-level abduction and its use for rule abduction. Section 3 then extends meta-level abduction to allow for two types of causal effects, in which positive and negative rules are called *triggers* and *inhibitors*, respectively, and investigate properties of two axiomatizations. Section 4 presents two case studies of meta-level abduction applied to completion of sub-networks in p53 signal networks. Section 5 shows experiments on causal networks, and analyzes scalability of our method in completing networks. Section 6 discusses related work, and Sect. 7 gives a summary and future work.

2 Meta-level abduction

This section revisits the framework for *meta-level abduction* (Inoue et al. 2010), and provides the correctness of rule abduction.

2.1 Causal networks

We suppose a background theory represented in a network structure called a *causal graph* or a *causal network*. A causal graph is a directed graph representing causal relations, which consists of a set of *nodes* and a set of (*directed*) *arcs* (or *links*).¹ Each node in a causal graph represents some event, fact or proposition. A *direct causal relation* corresponds to a directed arc, and a *causal chain* is represented by the reachability between two nodes. The interpretation of a “cause” here is kept rather informal, and just represents the connectivity, which may refer to a mathematical, physical, chemical, conceptual, epidemiological, structural, or statistical dependency (Pearl 2009). Similarly, a “direct cause” here simply represents the adjacent connectivity, while its effect is direct only relative to a certain level of abstraction.

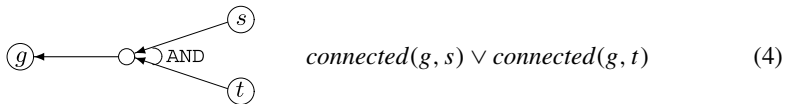

We then consider a first-order language to express causal networks. Each node is represented as a *proposition* or a (ground) atom in the language. When there is a direct causal relation from a node *s* to a node *g*, we define that *connected*(*g*, *s*) is true as in (1). Note that *connected*(*g*, *s*) only shows that *s* is one of possible causes of *g*, and thus the existence of another *connected*(*g*, *t*) (*s* ≠ *t*) means that *s* and *t* are alternative causes for *g*. Here, expression of causal relations is done at the *meta level* using the *meta-predicate* *connected*, while the *object level* refers to nodes in a causal network. An atom *connected*(*s*, *t*) at the meta level corresponds to a rule (*s* ← *t*) at the object level. The fact that a direct causal link cannot exist from *s* to *g* is represented in an (*integrity*) *constraint* of the form (2).

$$\textcircled{g} \leftarrow \textcircled{s} \quad \textit{connected}(g, s) \tag{1}$$

$$\textcircled{g} \leftarrow \textcircled{s} \quad \neg \textit{connected}(g, s) \tag{2}$$

¹Precisely speaking, our causal networks bring us more information than directed graphs, since negation, disjunctive effects and joint causes are all represented in a network.

A direct causal relation from s which has *nondeterministic effects* g and h , written as $(g \vee h \leftarrow s)$ at the object level, is represented in a disjunction of the form (3) at the meta level. On the other hand, the relation that “ g is *jointly* caused by s and t ”, written as $(g \leftarrow s \wedge t)$ at the object level, is expressed in a disjunction of the form (4) at the meta level, viz., $(g \leftarrow s \wedge t) \equiv (g \leftarrow s) \vee (g \leftarrow t)$.



There can be more than two atoms in a disjunction of the form (3) or (4). For example, $(g \leftarrow s \wedge t \wedge u)$ at the object level can be expressed as $connected(g, s) \vee connected(g, t) \vee connected(g, u)$. A complex relation of the form $(g \vee h \leftarrow s \wedge t)$ that has more than one node in both the left-hand and right-hand sides of the rule can be decomposed into two relations, $(s-t \leftarrow s \wedge t)$ and $(g \vee h \leftarrow s-t)$, where $s-t$ represents the intermediate complex. Then, any direct causal relation in a causal network can be represented by at most two disjunctions of atoms of the forms (3) and (4) at the meta level, using intermediate complexes.

In the above expression, (i) each atom at the object level is represented as a *term* at the meta level, and (ii) each rule at the object level is represented as a *fact* or a disjunction of facts at the meta level. The point (ii) can not only hold for rules given in the axioms, but can also be applied to express *inferred rules* at the meta level. Now, to express inferred rules, we introduce another meta-predicate *caused*. For object-level propositions g and s , we define that $caused(g, s)$ is true if there is a *causal chain* from s to g . Then, the causal chains are generally defined transitively in terms of *connected* as the axioms with variables:

$$caused(X, Y) \leftarrow connected(X, Y), \tag{5}$$

$$caused(X, Y) \leftarrow connected(X, Z) \wedge caused(Z, Y). \tag{6}$$

Here, the *caused/2* relation is recursively defined with the *connected/2* facts. The first-order expression with variables is thus used to represent that these axioms hold for all instances of them. Other algebraic properties as well as some particular constraints (e.g., $\neg caused(a, b)$) can also be defined if necessary. Variables in object-level expressions like $g(T)$ and $s(T)$ can be allowed in the meta-level expression like $connected(g(T), s(T))$, where the predicates g and s are treated as function symbols in the same way that Prolog can allow higher-order expressions. Here, an expression $g(T)$ can represent a set of nodes with the same property g with different values of the argument T such as time.


2.2 Rule abduction

Reasoning about causal networks is realized by deduction and abduction from the meta-level expression of causal networks together with the axioms for causal relations including (5) and (6). In *rule deduction*, we will later prove in Proposition 1 that, if a meta-level expression of the form $caused(g, s)$ for some facts g and s can be derived, it means that the rule $(g \leftarrow s)$ can be derived at the object level.

Similarly, we can realize *rule abduction* in the meta-level representation as follows. Suppose that a fact g is somehow caused by a fact s , which can also be regarded as an *input-output relation* that an output g is obtained given an input s . Here, g and s are called a *goal (fact)* (or a *target (fact)*) and a *source (fact)*, respectively. Setting an *observation* O as the causal chain $caused(g, s)$, we want to explain why or how O is caused. At the object level, O corresponds to the rule $(g \leftarrow s)$, which can be given as either a real observation (called an *empirical rule*) or a *virtual goal* to be achieved. An abductive task is then to find *hidden rules* that establish a connection from the source s to the goal g by filling the gaps in causal networks. As we will later see in Theorem 1, such an empirical rule can have more than one antecedent. For example, an object-level observation $(g \leftarrow s \wedge t)$ can be expressed as the meta-level formula $(caused(g, s) \vee caused(g, t))$ in the same way as (4).

Logically speaking, a *background theory* B consists of the meta-level expression of a causal network N and the axioms for causal relations at the meta level containing (5) and (6). When B is *incomplete*, there may be no path between g and s in B , that is, $caused(g, s)$ cannot be derived from B . Then, abduction infers an *explanation* (or *hypothesis*) H consisting of missing relations (links) and missing facts (nodes). This is realized by setting the *abducibles* Γ , the set of candidate literals to be assumed, as the atoms with the predicate *connected*: $\Gamma = \{connected(_, _)\}$. It is sometimes declared that there cannot exist any direct causal relation between the source and the goal, i.e., $\neg connected(g, s)$.

Formally, given a set O of formulas, a set H of instances of elements of Γ is an (*abductive*) *explanation* of O (with respect to B) if $B \cup H \models O$ and $B \cup H$ is consistent. A set of formulas can be interpreted as the conjunction of them. An explanation H of O is *minimal* if it does not imply any explanation of O that is not logically equivalent to H . Minimal explanations in meta-level abduction correspond to minimal additions in causal graphs, and are reasonable according to the principle of Occam’s razor. For example, suppose the observation $O = caused(g, s) \wedge caused(h, s)$, that is, the multiple causal chains between two goal facts g, h and the source fact s . Examples of minimal explanations of O containing two intermediate nodes are as follows.



H_1 and H_2 represent different connectivities, and we may want to enumerate different types of network structures that are missing in the original causal network. Here, H_1 corresponds to the four rules $\{(g \leftarrow \chi), (\chi \leftarrow s), (h \leftarrow \psi), (\psi \leftarrow s)\}$, hence rule abduction, i.e., abduction of rules, is realized. Moreover, these hypotheses contain existentially quantified variables, where χ and ψ are newly invented here. Those new terms can be regarded as either some existing nodes or new unknown nodes. Since new formulas can be produced at the object level, *predicate invention* (Muggleton and Buntine 1988) is partially realized in meta-level abduction.²

²Predicate invention in this form can also be regarded as a realization of *hidden object invention* discussed by Muggleton in Dietterich et al. (2008).

A hypothesis with a joint cause of the form (4) can be obtained by taking a disjunction of explanations of the form $connected(g, _)$ or by obtaining a *disjunctive answer* (Ray and Inoue 2007) for an observation containing a free variable X of the form $caused(g, X)$. Alternatively, this can be realized by adding a meta-level axiom:

$$connected(X, Y) \vee connected(X, Z) \leftarrow jointly_connected(X, Y, Z). \tag{7}$$

to the background theory B and the literals of the form $jointly_connected(_, _, _)$ to the abducibles Γ . Causes consisting of more than two joint links can be represented in a similar way. Atoms of the form $jointly_connected(_, _, _)$ together with axiom (7) can also be used to represent conjunctive causes of the form (4) in a causal network. That is, to express $(g \leftarrow s \wedge t)$ at the object level, the atom $jointly_connected(g, s, t)$ can be used instead of the formula $(connected(g, s) \vee connected(g, t))$ at the meta level.

The soundness and completeness of rule abduction in meta-level abduction can be derived as follows.³ For any meta-level theory N such that the predicate of any formula appearing in N is *connected* only, let $\lambda(N)$ be the object-level theory obtained by replacing every $connected(t_1, t_2)$ (t_1 and t_2 are terms) appearing in N with the formula $(t_1 \leftarrow t_2)$. We first show the correctness of meta-level deduction in the next proposition.

Proposition 1 *Suppose a meta-level theory N , which consists of disjunctions of facts of the form $connected(_, _)$. Let the background theory be $B = N \cup \{(5), (6)\}$. Then, $B \models (caused(g, s_1) \vee \dots \vee caused(g, s_n))$ if and only if $\lambda(N) \models (g \leftarrow s_1 \wedge \dots \wedge s_n)$.*

Proof We prove the proposition by induction on the depth of proof trees.⁴

Induction basis. It holds by the meaning of causal networks that $caused(g, s)$ is derived in a proof having depth 1 iff $connected(g, s) \in B$ (by (5)) iff $(g \leftarrow s) \in \lambda(N)$. Then, $(caused(g, s_1) \vee \dots \vee caused(g, s_n))$ is derived in a proof having depth 1 iff $(connected(g, s_1) \vee \dots \vee connected(g, s_n)) \in N$ iff $((g \leftarrow s_1) \vee \dots \vee (g \leftarrow s_n)) \in \lambda(N)$ iff $\lambda(N) \models (g \leftarrow s_1 \wedge \dots \wedge s_n)$.

Induction hypothesis. Suppose that the proposition holds for any formula $(caused(g, s_1) \vee \dots \vee caused(g, s_n))$ derived from B in a proof tree having depth d such that $d \leq k$.

Induction step. A formula $(caused(g, s_1) \vee \dots \vee caused(g, s_n))$ is derived in a proof tree having depth $k + 1$ iff $B \models ((connected(g, s_1) \vee \exists s'_1(connected(g, s'_1) \wedge caused(s'_1, s_1))) \vee \dots \vee (connected(g, s_n) \vee \exists s'_n(connected(g, s'_n) \wedge caused(s'_n, s_n))))$ (by (5) and (6)) such that $caused(s'_j, s_j)$ is derived in a proof tree having depth k for $j = 1, \dots, n$ iff $\lambda(N) \models (((g \leftarrow s_1) \vee \exists s'_1((g \leftarrow s'_1) \wedge (s'_1 \leftarrow s_1))) \vee \dots \vee ((g \leftarrow s_n) \vee \exists s'_n((g \leftarrow s'_n) \wedge (s'_n \leftarrow s_n))))$ (by the induction hypothesis) iff $\lambda(N) \models ((g \leftarrow s_1) \vee \dots \vee (g \leftarrow s_n))$ iff $\lambda(N) \models (g \leftarrow s_1 \wedge \dots \wedge s_n)$. □

Note in Proposition 1 that we do not need the λ -counterparts of axioms (5) and (6) in the object level. This logic reflects the assumption that transitivity of cause by \leftarrow holds for a chain in a causal network. Now we have the correctness of meta-level abduction.

³The properties presented here are generalizations of preliminary results in Inoue et al. (2010, Sect. 3.3), which proved the correctness when the observation is a fact, i.e., a single atom, of the form $caused(g, s)$.

⁴For example, SOL tableaux (Nabeshima et al. 2003, 2010) (Sect. 2.4) can be used as proof trees in deducing target consequences. Since SOL(AR) is complete for consequence-finding (Inoue 1992; Nabeshima et al. 2010), for any minimal clause C derived from a consistent axiom set, there is an SOL tableau producing C with a certain depth.

Theorem 1 Let N and B be the same theories as in Proposition 1. Suppose the observation $O = (\text{caused}(g, s_1) \vee \dots \vee \text{caused}(g, s_n))$, and let $\Gamma = \{\text{connected}(_, _)\}$. Then, H is an abductive explanation of $\text{caused}(g, s)$ with respect to B and Γ if and only if $\lambda(H)$ is a hypothesis satisfying that

$$\lambda(N) \cup \lambda(H) \models (g \leftarrow s_1 \wedge \dots \wedge s_n), \quad \text{and} \quad (8)$$

$$\lambda(N) \cup \lambda(H) \text{ is consistent.} \quad (9)$$

Proof The equivalence between the relation that $B \cup H \models O$ and the abductive derivation (8) holds by Proposition 1. The equivalence between the consistency of $B \cup H$ and the consistency (9) is obvious: $B \cup H$ is consistent because it contains no integrity constraint, and so is $\lambda(N) \cup \lambda(H)$. \square

2.3 Abduction of rules and facts

Besides the use in rule abduction, meta-level abduction can also be applied to *fact abduction* (Inoue et al. 2010), which has been focused on almost exclusively in research of abduction in AI.⁵ Abduction of facts at the object level can be formalized as *query answering* at the meta level. Given a goal of the form $\text{caused}(g, X)$, abduction of causes is computed by answer substitutions to the variable X . To this end, each abducible literal a at the object level is associated with the fact $\text{caused}(a, a)$ at the meta level. That is, an abducible can hold by assuming itself. Equivalently, the axiom for abducibles is expressed using the meta-predicate abd as:

$$\text{caused}(X, X) \leftarrow \text{abd}(X).$$

Then, each abducible a at the object level should be declared as $abd(a)$. Answer extraction for the query $\leftarrow \text{caused}(g, X)$ can be realized by giving the meta-level formula of the form:

$$\text{ans}(X) \leftarrow \text{caused}(g, X) \wedge \text{abd}(X). \quad (10)$$

Here, ans is the *answer predicate* (Iwanuma and Inoue 2002), and the variable X is used to collect only abducibles which cause g . An *integrity constraint* that two facts p and q cannot hold at the same time ($\leftarrow p \wedge q$) can be represented as:

$$\leftarrow \text{caused}(p, X) \wedge \text{caused}(q, Y) \wedge \text{abd}(X) \wedge \text{abd}(Y).$$

This makes any combination of abducibles that causes p and q incompatible. Such a set of incompatible abducibles is called a *nogood*. Finally, by combining rule abduction and fact abduction in the form of *conditional query answering* (Iwanuma and Inoue 2002), which extracts answers in a query with additional abduced conditions, meta-level abduction enables us to *abduce both rules and facts* (Inoue et al. 2010).

⁵In a philosophical work (Schurz 2008), fact abduction and rule abduction are classified as *factual abduction* and *law-abduction*, respectively. Our meta-level abduction also gives a realization of *2nd order existential abduction*, which is most important to produce new theories with new concepts (Schurz 2008).

2.4 Computation by consequence finding

All types of meta-level inferences in this section, including generation of existentially quantified hypotheses in meta-level abduction as well as conditional query answering to abduce rules and facts, can be realized by SOLAR (Nabeshima et al. 2003, 2010). SOLAR is a *consequence-finding* system based on SOL resolution (Inoue 1992) and the connection tableaux.

In SOLAR, the notion of production fields (Inoue 1992) is used to represent language biases for hypotheses. A *clause* is a disjunction of literals. A *production field* \mathcal{P} is a pair $\langle \mathbf{L}, \text{Cond} \rangle$, where \mathbf{L} is a set of literals and *Cond* is a certain condition. If *Cond* is *true* (empty), \mathcal{P} is denoted as $\langle \mathbf{L} \rangle$. A clause C *belongs to* $\mathcal{P} = \langle \mathbf{L}, \text{Cond} \rangle$ if every literal in C is an instance of a literal in \mathbf{L} and C satisfies *Cond*. Let Σ be a clausal theory. The set of consequences of Σ belonging to \mathcal{P} is denoted as $Th_{\mathcal{P}}(\Sigma)$. The *characteristic clauses* of Σ with respect to \mathcal{P} are defined as $\text{Carc}(\Sigma, \mathcal{P}) = \mu Th_{\mathcal{P}}(\Sigma)$, where μT denotes the set of clauses in T that are minimal with respect to subsumption. The *new characteristic clauses* of a clause C with respect to Σ and \mathcal{P} are defined as $\text{Newcarc}(\Sigma, C, \mathcal{P}) = \mu [Th_{\mathcal{P}}(\Sigma \cup \{C\}) \setminus Th_{\mathcal{P}}(\Sigma)]$.

Let B be a clausal theory (background theory) and O a set of literals (observations). Then, a set H of literals is obtained as an abductive explanation of O by *inverse entailment* (Inoue 1992):

$$B \cup \{\neg O\} \models \neg H, \quad (11)$$

where both $\neg O = \bigvee_{L \in O} \neg L$ and $\neg H = \bigvee_{L \in H} \neg L$ are clauses (because O and H are sets of literals and are interpreted as conjunctions of them). Similarly, the condition that $B \cup H$ is consistent is equivalent to $B \not\models \neg H$. Hence, for any hypothesis H , its negated form $\neg H$ is deductively obtained as a “new” consequence of $B \cup \{\neg O\}$ which is not an “old” consequence of B alone. Given the abducibles Γ , any literal in $\neg H$ is an instance of a literal in $\overline{\Gamma} = \{\neg L \mid L \in \Gamma\}$. Hence, the set of minimal explanations of O with respect to B and Γ is characterized as $\{H \mid \neg H \in \text{Newcarc}(B, \neg O, \langle \overline{\Gamma} \rangle)\}$, while the set of minimal nogoods with respect to B and Γ is $\{H \mid \neg H \in \text{Carc}(B, \langle \overline{\Gamma} \rangle)\}$.

SOLAR is complete for finding (new) characteristic clauses with respect to a given production field. SOLAR can thus be used to implement a *complete* abductive system for finding and enumerating minimal explanations from *full clausal theories* containing non-Horn clauses. A simple way to compute $\text{Newcarc}(\Sigma, C, \mathcal{P})$ in SOLAR is: (1) enumerate $\text{Carc}(\Sigma, \mathcal{P})$, and then (2) enumerate the *SOL tableau deductions* from $\Sigma \cup \{C\}$ with the top clause C (Nabeshima et al. 2010) by removing each produced clause subsumed by some clause in $\text{Carc}(\Sigma, \mathcal{P})$.

3 Reasoning about positive and negative causal effects

So far, meta-level abduction has been defined for causal networks without explicitly arguing the meaning of causes. Indeed, links in a causal network have been of one kind, and *connected*(g, s) at the meta level, i.e., ($g \leftarrow s$) at the object level, just represents that g directly depends on s somehow. However, mixing different types of causalities in one type of links often makes analysis of actual causes complicated (Pearl 2009). For example, suppose that increase of the amount of p decreases the amount of q and that increase of q causes increase of r . In this case, we cannot say that increase of p causes increase of r because q cannot mediate between p and r . For this problem, it is not appropriate to represent the

causalities as $(p \rightarrow q \rightarrow r)$ because transitivity does not hold. Instead, $(inc(p) \rightarrow dec(q))$ and $(inc(q) \rightarrow inc(r))$ would be more precise but we need more entities and relations between this kind: networks with two types of causalities, i.e., *positive* and *negative* causal effects. With this regard, from now on we can understand that each arc of the form *connected*(g, s) in Sect. 2 only represents positive effects.

We extend applicability of meta-level abduction to deal with networks expressing both positive and negative causal effects. Such networks are seen in biological domains, where *inhibition* effects negatively in gene regulatory, signaling and metabolic pathways. Now we consider two types of direct causal relations: *triggered* and *inhibited*. For two nodes g and t , the relation *triggered*(g, t) represents a positive cause such that t is a *trigger* of g , written as $g \leftarrow t$ in a causal network. On the other hand, the relation *inhibited*(g, s) represents a negative cause such that s is an *inhibitor* of g , written as $g \vdash s$ in a causal network.⁶ The meaning of these links will be given in two ways in Sects. 3.1 and 3.2.

As in Sect. 2.1, negation, disjunctive effects and conjunctive causes can be defined for *triggered* and *inhibited*, cf., (2), (3) and (4), and complex causal relations can be represented using those combinations and intermediate complexes. For instance, g is jointly triggered by t_1 and t_2 can be expressed as $triggered(g, t_1) \vee triggered(g, t_2)$.

The notion of causal chains is also divided into two types: the positive one (written *promoted*) and the negative one (written *suppressed*), respectively corresponding to *triggered* and *inhibited*. Now our task is to design the axioms for these two meta-predicates.

3.1 Alternating axioms for causality

Suppose first that there is no inhibitor in a causal network, that is, all links are positive. In this case, the axioms for *promoted* should coincide with (5) and (6):

$$promoted(X, Y) \leftarrow triggered(X, Y), \tag{12}$$

$$promoted(X, Y) \leftarrow triggered(X, Z) \wedge promoted(Z, Y). \tag{13}$$

Next, let us interpret the meaning of an inhibitor as a toggle switch of signals flowed in the inhibitor, just as an inverter in a logic circuit (Shmulevich et al. 2002). Then, in the presence of inhibitors, we need one more axiom (14), which blocks an adjacent inhibitor for X in order to promote X :

$$promoted(X, Y) \leftarrow inhibited(X, Z) \wedge suppressed(Z, Y). \tag{14}$$

As for the axioms of the negative causal chain *suppressed*, we can consider the following axioms (15), (16) and (17), which are the counterpart of positive ones (12), (13) and (14):

$$suppressed(X, Y) \leftarrow inhibited(X, Y), \tag{15}$$

$$suppressed(X, Y) \leftarrow inhibited(X, Z) \wedge promoted(Z, Y), \tag{16}$$

$$suppressed(X, Y) \leftarrow triggered(X, Z) \wedge suppressed(Z, Y). \tag{17}$$

⁶A trigger and an inhibitor are often called an *activator* and a *repressor*, respectively.

That is, a negative causal chain to X can be established if negative influence is propagated to X either directly by an adjacent inhibitor of an active (15) or promoted (16) item or indirectly by a trigger of suppressed (17). By this way, we can establish the connection between X and Y by mixing both positive and negative links.

One nice property with the axiomatization by (12)–(17) is that all possible paths from a source to a goal, which is either positive or negative, can be obtained by meta-level abduction. We here prove this property step by step.⁷ Firstly, the correctness of the axiomatization by (12)–(17) with respect to deduction is shown as follows. Suppose a meta-level theory N , which represents a causal network with triggers and inhibitors, and a clause at the meta-level $O^+ = promoted(g, s)$ (resp. $O^- = suppressed(g, s)$), where g and s are nodes in N . When O^+ (resp. O^-) can be proved from $N \cup \{(12)–(17)\}$, a *proof* Π of O^+ (resp. O^-) is a sub-network of N such that Π consists of triggers and inhibitors that connect paths from s to g . A proof Π of O^+ (resp. O^-) is also called a *positive* (resp. *negative*) *causal chain* from s to g in N . Then, each proof of O^+ or O^- is a sequence of links, L_1, L_2, \dots, L_m , where m is the length of the chain and each L_j is either a trigger or an inhibitor in N such that the endpoint of L_1 is g , the start point of L_m is s , and the start point of L_j is the same as the endpoint of L_{j+1} for $j = 1, \dots, m - 1$.

Lemma 1 *Let N be a causal network, and s and g be nodes in N , which respectively represent a source and a goal. There exists a positive (resp. negative) causal chain from s to g in N if and only if there exists a path from s to g in N such that there are an even (resp. odd) number of occurrences of inhibitors.*

Proof We here prove the only-if direction of the lemma, but the if direction can be proved in a similar way. Let Π^+ (resp. Π^-) be a positive (resp. negative) causal chain from s to g in N . We prove the only-if direction of the lemma by a double induction on the number k of occurrences of inhibitors in Π^+ or Π^- .

Firstly, consider the case of $k = 0$. In this case, Π^+ consists of atoms with the predicate *triggered* only, each of which appears in either (12) or (13). Construction of Π^+ is indeed possible by Proposition 1 when we replace *triggered* with *connected* and axioms (12) and (13) with (5) and (6), respectively. On the other hand, if Π^- exists, it must contain at least one atom with *inhibited* by (15) or (16). Hence there is no Π^- in this case.

Secondly, consider the case of $k = 1$. If Π^+ exists, then it must contain an inhibitor *inhibited*(χ, ψ) for some χ and ψ in N from a rule of the form (14) after some chain of triggers via (13). Then *suppressed*(ψ, s) must be derived, but Π^+ cannot contain any other inhibitor by the assumption. By the case of $k = 0$, it is impossible to derive *suppressed*(ψ, s) without using any inhibitor. Hence, there is no Π^+ in this case. On the other hand, Π^- must contain either an inhibitor *inhibited*(χ, s) for some node χ in N from a rule of the form (15) or an inhibitor *inhibited*(χ, ψ) for some χ and ψ in N from a rule of the form (16) after some chain of triggers via (17). In the latter case, *promoted*(ψ, s) must be proved with no inhibitor, which is possible by the case of $k = 0$. In either case, Π^- exists in this case.

⁷A preliminary version of the completeness property presented in Inoue et al. (2011, Proposition 3.1) pre-supposes that the observation is a unit clause of the form *suppressed*(g, s) or *promoted*(g, s), and does not properly distinguish inhibitors contained in the axiom set and inhibitors abduced in an explanation.

Next, in the case of $k = m \geq 1$, suppose the two propositions: (I) for any positive causal chain Π^+ , the number m of occurrences of inhibitors in Π^+ is even; and (II) for any negative causal chain Π^- , the number m of occurrences of inhibitors in Π^- is odd.

Now, consider the case of $k = m + 1$. Then, Π^+ contains either a trigger from (13) or an inhibitor from (14). In the former case, the predicate of the other subgoal of (13) is *promoted*, but after some sequence of triggers, an inhibitor must be contained in Π^+ from (14). That is, at least one inhibitor is contained and the predicate of the rest of subgoals is *suppressed*, whose proof contains m inhibitors. In either case, by the induction hypothesis (II), m is odd so k is even. On the other hand, Π^- contains either an inhibitor from (16) or a trigger from (17). In the latter case, the predicate of the other subgoal of (17) is *suppressed*, but after some sequence of triggers, an inhibitor must be contained in Π^- from (16). That is, an inhibitor is encountered and the predicate of the other subgoals is *promoted*, whose proof contains m inhibitors. In either case, by the induction hypothesis (I), m is even, so k is odd. □

In axiomatization (12)–(17), the logical interpretations of $g \leftarrow s$ and $g \vdash s$ in causal networks are defined as follows. A trigger *triggered*(g, s) at the meta level is interpreted as $(g \leftarrow s) \wedge (\neg g \leftarrow \neg s)$ at the object level, which is now abbreviated as $(g \leftrightarrow s)$, then the truth value of s is preserved in the truth value of g with this trigger. On the other hand, an inhibitor *inhibited*(g, s) at the meta level can be interpreted as $(\neg g \leftarrow s) \wedge (g \leftarrow \neg s)$ at the object level, which is now abbreviated as $(\neg g \leftrightarrow s)$, and the truth value of s is reversed in the truth value of g with this inhibitor. For any meta-level clausal theory N such that the predicate of each clause appearing in N is either *triggered* or *inhibited*, let $\lambda(N)$ be the object-level theory obtained by replacing (i) every *triggered*(t_1, t_2) (t_1 and t_2 are terms) appearing in N with the formula $(t_1 \leftrightarrow t_2)$ and (ii) every *inhibited*(t_1, t_2) appearing in N with the formula $(\neg t_1 \leftrightarrow t_2)$.

Proposition 2 *Suppose a meta-level theory N , which consists of disjunctions of literals of the form *triggered*($_, _$) and disjunctions of literals of the form *inhibited*($_, _$). Let the background theory be $B = N \cup \{(12)–(17)\}$. Then, (1) $B \models (\text{promoted}(g, s_1) \vee \dots \vee \text{promoted}(g, s_n))$ if and only if $\lambda(N) \models (g \leftarrow s_1 \wedge \dots \wedge s_n)$; and (2) $B \models (\text{suppressed}(g, s_1) \vee \dots \vee \text{suppressed}(g, s_n))$ if and only if $\lambda(N) \models (\neg g \leftarrow s_1 \wedge \dots \wedge s_n)$.*

Proof Extending the proof of Proposition 1, we can prove that any (partial) proof Π_i of *promoted*(g_i, s_i) or *suppressed*(g_i, s_i) ($g_i \neq s_i$) from B contains m occurrences of inhibitors from N iff $\lambda(N) \models \underbrace{(\neg(\neg(\dots\neg(g_i))))}_{m} \leftarrow s_i$.⁸ By the result and the proof of Lemma 1, such a proof of *promoted*(g_i, s_i) (resp. *suppressed*(g_i, s_i)) ($g_i \neq s_i$) from B contains an even (resp. odd) number of occurrences of inhibitors from N . Then, (1) $B \models \text{promoted}(g_i, s_i)$ iff $\lambda(N) \models \underbrace{(\neg(\neg(\dots\neg(g_i))))}_{m:\text{even}} \leftarrow s_i$ iff $\lambda(N) \models (g_i \leftarrow s_i)$. On the other hand, (2) $B \models \text{suppressed}(g_i, s_i)$ iff $\lambda(N) \models \underbrace{(\neg(\neg(\dots\neg(g_i))))}_{m:\text{odd}} \leftarrow s_i$ iff $\lambda(N) \models (\neg g_i \leftarrow s_i)$. The proposition follows by applying these results to disjuncts of goal clauses. □

⁸To see this, if there is no inhibitor in Π_i , ($g_i \leftrightarrow s_i$) holds. If there is one inhibitor in Π_i , only one change of negation is introduced in a chain from s_i , which is then propagated to g_i , i.e., $(\neg g_i \leftrightarrow s_i)$. Similarly, an additional introduction of an inhibitor in Π_i adds one more negation at g_i .

Proposition 2 states that the axiomatization by (12)–(17) is sound and complete for enumerating positive and negative causal chains. Extending deductive inference to abductive inference, meta-level abduction on causal networks with positive and negative links can now be defined by letting the abducibles Γ be those atoms with the predicates *triggered* and/or *inhibited*: $\Gamma = \{\text{triggered}(_, _), \text{inhibited}(_, _)\}$, and observations or goals are given as (disjunctions of) literals either of the form *promoted*(g, s) or of the form *suppressed*(g, s). Hence, given positive and negative observations, we can abduce both positive and negative causes, and new nodes are produced whenever necessary.

Theorem 2 *Let N and B be the same theories as in Proposition 2. Let the abducibles be $\Gamma = \{\text{triggered}(_, _), \text{inhibited}(_, _)\}$. Then, H^+ is an abductive explanation of $O^+ = (\text{promoted}(g, s_1) \vee \dots \vee \text{promoted}(g, s_n))$ with respect to B and Γ if and only if $\lambda(H^+)$ is a hypothesis satisfying that*

$$\lambda(N) \cup \lambda(H^+) \models (g \leftarrow s_1 \wedge \dots \wedge s_n) \tag{18}$$

and $\lambda(N) \cup \lambda(H^+)$ is consistent. On the other hand, H^- is an abductive explanation of $O^- = (\text{suppressed}(g, s_1) \vee \dots \vee \text{suppressed}(g, s_n))$ with respect to B and Γ if and only if $\lambda(H^-)$ is a hypothesis satisfying that

$$\lambda(N) \cup \lambda(H^-) \models (\neg g \leftarrow s_1 \wedge \dots \wedge s_n) \tag{19}$$

and $\lambda(N) \cup \lambda(H^-)$ is consistent.

Proof The equivalence between the meta-level abductive entailment $B \cup H^+ \models O^+$ and the object-level abductive entailment (18) holds by Proposition 2(1). Similarly, the equivalence between the relation $B \cup H^- \models O^-$ and relation (19) holds by Proposition 2(2). The equivalence between the consistency of $B \cup H^+$ (resp. $B \cup H^-$) and the consistency of $\lambda(N) \cup \lambda(H^+)$ (resp. $\lambda(N) \cup \lambda(H^-)$) also holds: Since $B \cup H^+$ (resp. $B \cup H^-$) is consistent, so is $\lambda(N) \cup \lambda(H^+)$ (resp. $\lambda(N) \cup \lambda(H^-)$). \square

Theorem 2 shows the soundness and completeness of axiomatization (12)–(17) for meta-level abduction of positive and negative causal effects. Moreover, the consistency of the background theory $B = N \cup \{(12)–(17)\}$ in Theorem 2 also holds. However, both *promoted*(g, s) and *suppressed*(g, s) can be explained from this B for some N at the same time. This “semantic inconsistency” is, however, inevitable in this monotonic representation, since we can answer to any virtual query supposing a source s and a goal g . Then, to prevent derivations of promotion and suppression simultaneously for the same s and g , the following integrity constraint can be placed at the meta level.

$$\leftarrow \text{promoted}(X, Y) \wedge \text{suppressed}(X, Y). \tag{20}$$

The role of (20) is to derive (minimal) nogoods, i.e., a (minimal) set of incompatible instances of abducibles. Any abductive explanation must not include any nogood. For example, given the network

$$N_1 = \{\text{triggered}(g, t)\},$$

the minimal nogoods containing at most 2 abducibles are the following 13 clauses.⁹

$$\begin{aligned} & \{inhibited(g, t)\}, \{inhibited(t, g)\}, \{inhibited(\chi, \chi)\}^*, \\ & \{triggered(g, \chi), inhibited(\chi, t)\}, \{inhibited(g, \chi), triggered(\chi, t)\}, \\ & \{triggered(t, \chi), inhibited(g, \chi)\}, \{triggered(\chi, \psi), inhibited(\psi, \chi)\}^*, \\ & \{triggered(g, \chi), inhibited(t, \chi)\}, \{triggered(t, \chi), inhibited(\chi, g)\}, \\ & \{triggered(\chi, g), inhibited(\chi, t)\}, \{triggered(\chi, \psi), inhibited(\chi, \psi)\}^*, \\ & \{inhibited(\chi, g), triggered(\chi, t)\}, \{triggered(\chi, g), inhibited(\chi, t)\}, \end{aligned}$$

where χ and ψ are existentially quantified variables. A set with the symbol “*” does not refer constants g, t , and is also a minimal nogood of the empty network $N_0 = \emptyset$. This implies that any causal network including an instance of such a set with “*” becomes inconsistent. Hence, introduction of (20) into axiom set (12)–(17) can make a causal network inconsistent, that is, *the empty set \emptyset can become a nogood*. For example, given the network

$$N_2 = \{triggered(g, t), inhibited(g, s), triggered(s, t)\}, \quad (21)$$

$N_2 \cup \{(12)–(17)\}$ is consistent. However, $N_2 \cup \{(12)–(17), (20)\}$ is inconsistent, since both *promoted*(g, t) (by (12)) and *suppressed*(g, t) (by (16) and (12)) hold. Also, the p53 network (47), which will be given later in Sect. 4.1, becomes inconsistent if it is combined with this axiomatization $\{(12)–(17), (20)\}$. The (in)consistency of a theory with axioms (12)–(17) together with constraint (20) is characterized as follows.

Proposition 3 *Let N be the same as in Proposition 2. Let the background theory be $B = N \cup \{(12)–(17), (20)\}$. Then, B is consistent if and only if there exist no nodes g and s in N such that there are causal chains Π^+ and Π^- from s to g in N such that there are an even (resp. odd) number of occurrences of inhibitors in Π^+ (resp. Π^-).*


Proof If the condition is satisfied, Π^+ is a proof of *promoted*(g, s) and Π^- is a proof of *suppressed*(g, s). Then B becomes inconsistent by (20). Conversely, if B is inconsistent, both *promoted*(g, s) and *suppressed*(g, s) can be proved from B for some nodes g and s in N by (20). Let Π^+ be a proof of *promoted*(g, s), and Π^- a proof of *suppressed*(g, s). Then, Π^+ (resp. Π^-) contains an even (resp. odd) number of inhibitors by Lemma 1. \square

We next show a modification of this monotonic axiomatization to reduce such semantic inconsistencies using default assumptions.

⁹This enumeration is verified by SOLAR. Recall (Sect. 2.4) that a minimal nogood can be constructed as the negation of a characteristic clause of the background theory B , which is now the union of the causal network N at the meta-level and axioms (12)–(17) with constraint (20). Note that the number of (minimal) nogoods becomes infinite if we do not restrict the number of abducibles appearing in each nogood.

3.2 Axiomatization with default assumptions

Suppose two antagonistic direct causal relations appear simultaneously for the same node g as follows.



Our intuition on diagram (22) is as follows. (I) If the trigger t is present and the inhibitor s is not present, then g is triggered by t (and is not inhibited); (II) Else if s is present and t is not present, then g is inhibited by s (and is not triggered). These two cases are rather clear, but what happens for g if both t and s are somehow caused? Is it triggered or inhibited? The last axiomatization in Sect. 3.1 concludes that g is both promoted and suppressed by s , which leads to inconsistency under the existence of the integrity constraint of the form (20). However, often it is indicated that: (III) If both t and s are present, then g is inhibited by s (and is not triggered). Namely, *an inhibitor is preferred to a trigger*.¹⁰ This preference rule is often used in applications, particularly in the biological literature.¹¹

The last inference is *nonmonotonic*: a trigger of g works *if there is no inhibitor for g* , but if an inhibitor is added then the trigger stops working. We now show another axiomatization which reflects this principle of inhibitor preference. Depart from the monotonic axiomatization of causal chains (12)–(17), reasoning about networks is now made nonmonotonic and involves *default reasoning*. In the following new definitions of *promoted* and *suppressed*, we will associate an extra condition for each trigger to work.

$$promoted(X, Y) \leftarrow triggered(X, Y) \wedge no_inhibitor(X), \tag{23}$$

$$promoted(X, Y) \leftarrow triggered(X, Z) \wedge no_inhibitor(X) \wedge promoted(Z, Y), \tag{24}$$

$$promoted(X, Y) \leftarrow inhibited(X, Z) \wedge suppressed(Z, Y), \tag{25}$$

$$suppressed(X, Y) \leftarrow inhibited(X, Y), \tag{26}$$

$$suppressed(X, Y) \leftarrow inhibited(X, Z) \wedge promoted(Z, Y), \tag{27}$$

$$suppressed(X, Y) \leftarrow triggered(X, Z) \wedge no_inhibitor(X) \wedge suppressed(Z, Y), \tag{28}$$

$$\leftarrow promoted(X, Y) \wedge suppressed(X, Y). \tag{29}$$

The new axiom set for positive and negative causal chains (23)–(28) are the same as the monotonic version (12)–(17), except that each trigger to X ($triggered(X, _)$) must not be inhibited ($no_inhibitor(X)$) to give the positive effect to X in (23), (24) and (28).¹² Here, in-

¹⁰In some formalization of inhibitors (Leitgeb 2001; Fayruzov et al. 2010), an inhibitor is modeled to work against a trigger/inhibitor edge between nodes by blocking its activation/inhibition function. This results in the same effect as preferring inhibitors to triggers when an inhibitor works against activation of a trigger. Blocking inhibition by an inhibitor has the same effect as double inhibition in our formalization (see Remark 2).

¹¹In some applications, it is possible to consider the preference for a trigger over an inhibitor. All theories explored in this section can be easily adapted in accordance with this reverse preference by just exchanging the roles of positive and negative causal relations.

¹²Instead of (26), ($suppressed(X, Y) \leftarrow inhibited(X, Y) \wedge no_inhibitor(Y)$) was used in an earlier version of this paper (http://ilp2010.dsi.unifi.it/pdf/ilp2010_submission_43.pdf). This axiom was only applied at an inhibitor to X to which $no_inhibitor(Y)$ is assumed. Although this is a possible formalization, it introduced an inelegant asymmetry between positive and negative influences.

clusion of $no_inhibitor(\chi)$ in association with $triggered(\chi, \psi)$ makes those three axioms *default rules* (Reiter 1980; Poole 1988). The literal of the form $no_inhibitor(_)$ is thus treated as a *default*, which can be assumed during inference unless contradiction occurs, and constraints can also be added to reject inconsistent cases with these assumptions. Then, a trigger $triggered(g, s)$ at the meta level is now interpreted as $(no_inhibitor(g) \rightarrow (g \leftrightarrow s))$ at the object level, that is, $(g \leftrightarrow s)$ is true if the default $no_inhibitor(g)$ is consistent with the union of the background theory B and a constructing hypothesis H . On the other hand, an inhibitor $inhibited(g, s)$ at the meta level can be interpreted as $(\neg g \leftrightarrow s)$ at the object level, hence the effect of an inhibitor does not need addition of a default. Finally, the integrity constraint of the form (29) is the same as (20) and prohibits the presence of both positive and negative causes between any pair of nodes X and Y .

Meta-level abduction under axiomatization (23)–(28) and the integrity constraint of the form (29) is defined in the same way as in Sect. 3.1. Abduction of *joint triggers* and *joint inhibitors* can also be realized in the same way as *jointly_connected* in (7) by adding the meta-level axioms

$$triggered(X, Y) \vee triggered(X, Z) \leftarrow jointly_triggered(X, Y, Z), \tag{30}$$

$$inhibited(X, Y) \vee inhibited(X, Z) \leftarrow jointly_inhibited(X, Y, Z). \tag{31}$$

to the background theory B and the literals of the form $jointly_triggered(_, _, _)$ and of the form $jointly_inhibited(_, _, _)$ to the abducibles Γ . Note that one of definitions (30), (31) of these meta-predicates can be omitted by simulating one by the other. For example, abduction of $jointly_inhibited(g, s, t)$ can be simulated by abduction of $inhibited(g, s-t) \wedge jointly_triggered(s-t, s, t)$, where $s-t$ is an intermediate complex.

Remark 1 Unlike axiom (7) of *jointly_connected* in Sect. 2.2, we can propagate negative effects from sources to goals through joint triggers and joint inhibitors by the axioms:

$$suppressed(G, X) \leftarrow jointly_triggered(G, A, B) \wedge suppressed(A, X), \tag{32}$$

$$suppressed(G, X) \leftarrow jointly_triggered(G, A, B) \wedge suppressed(B, X), \tag{33}$$

$$promoted(G, X) \leftarrow jointly_inhibited(G, A, B) \wedge suppressed(A, X), \tag{34}$$

$$promoted(G, X) \leftarrow jointly_inhibited(G, A, B) \wedge suppressed(B, X). \tag{35}$$

Axioms (32) and (33) show that suppression of one of the joint triggers causes suppression of the goal, while axioms (34) and (35) show that suppression of one of the joint inhibitors causes promotion of the goal. All these inverse effects can be caused by the suppression of one of the joint links regardless of the status of the other link.

As for default assumptions of the form $no_inhibitor(_)$, *default reasoning* can be implemented by assuming those literals whenever necessary during inference, and the consistency of such assumptions is checked each time they are added to the current set of abduced literals. This is a simple yet powerful method for default reasoning in the case of so-called *normal defaults (without prerequisites)* (Reiter 1980; Poole 1988). Hence, the abducibles Γ now also contain the literals of the form $no_inhibitor(_)$ along with $triggered(_, _)$ and $inhibited(_, _)$. For example, the abducibles allowing joint triggers can be defined as

$$\Gamma = \{triggered(_, _), inhibited(_, _), jointly_triggered(_, _, _), no_inhibitor(_)\}.$$

Note that addition of default literals through the abducibles $\{no_inhibitor(_)\}$ is necessary not only for abduction with defaults but for deduction involving defaults.

The soundness and completeness of meta-level abduction with the new axiomatization is guaranteed as in Theorem 2. Let M be a meta-level clausal theory such that the predicate of any literal appearing in M is either *triggered*, *inhibited*, or *no_inhibitor*. Now $\lambda^\dagger(M)$ is defined as the object-level theory obtained by replacing (i) every *triggered*(t_1, t_2) (t_1 and t_2 are terms) appearing in M with the formula $((t_1 \leftrightarrow t_2) \wedge t_1^*)$ (t_1^* is a new term uniquely associated with t_1), (ii) every *inhibited*(t_1, t_2) appearing in M with the formula $(\neg t_1 \leftrightarrow t_2)$, and (iii) every *no_inhibitor*(t) appearing in M with t^* . Note that t^* represents that t is not inhibited by default.

Theorem 3 *Suppose a meta-level theory N , which consists of disjunctions of literals of the form $triggered(_, _)$ and disjunctions of literals of the form $inhibited(_, _)$. Let the background theory be $B = N \cup \{(23)–(28)\}$. Let the abducibles be*

$$\Gamma = \{triggered(_, _), inhibited(_, _), no_inhibitor(_)\}.$$

Then, H^+ is an abductive explanation of $O^+ = (promoted(g, s_1) \vee \dots \vee promoted(g, s_n))$ with respect to B and Γ if and only if $\lambda^\dagger(H^+)$ is a hypothesis satisfying that

$$\lambda^\dagger(N) \cup \lambda^\dagger(H^+) \models (g \leftarrow s_1 \wedge \dots \wedge s_n) \tag{36}$$

and $\lambda^\dagger(N) \cup \lambda^\dagger(H^+)$ is consistent. On the other hand, H^- is an abductive explanation of $O^- = (suppressed(g, s_1) \vee \dots \vee suppressed(g, s_n))$ with respect to B and Γ if and only if $\lambda^\dagger(H^-)$ is a hypothesis satisfying that

$$\lambda^\dagger(N) \cup \lambda^\dagger(H^-) \models (\neg g \leftarrow s_1 \wedge \dots \wedge s_n) \tag{37}$$

and $\lambda^\dagger(N) \cup \lambda^\dagger(H^-)$ is consistent.

Proof Each explanation H obtained with the monotonic axioms (12)–(17) in Theorem 2 can also be obtained by incorporating defaults of the form *no_inhibitor*(t) in the corresponding explanation H^* with axioms (23)–(28), which are translated to atoms of the form t^* in $\lambda^\dagger(H^*)$. Then, H satisfies (18) (resp. (19)) iff H^* satisfies (36) (resp. (37)). \square

A background theory B can also include several knowledge about defaults, and here are some examples. First, when we are sure that there is no inhibitor for a node s , we can include the fact *no_inhibitor*(s) in the background theory B . For instance, *no_inhibitor*(s) can be declared to be true in B if s is a *terminal source node*:

$$no_inhibitor(X) \leftarrow source(X), \tag{38}$$

where each source s is declared as *source*(s) in B too. Such a terminal source node has the same effect as an abducible fact *abd*(s) in Sect. 2.3. We will see an example to use *source*($_$) and constraint (38) in Sect. 4.1. Second, a meta-level constraint:

$$\leftarrow no_inhibitor(X) \wedge inhibited(X, Y). \tag{39}$$

blocks to assume a default *no_inhibitor*(g) for any node g to which an inhibitor is connected. This constraint is natural in many cases, but is only optional because *inhibited*(g, k)

and $no_inhibitor(g)$ may coexist when, for example, there are multiple supports through derivations of $promoted(g, h_1), \dots, promoted(g, h_n)$ ($n \geq 2$) that share $no_inhibitor(g)$, yet only one $inhibited(g, k)$ exists.¹³

Remark 2 Introduction of (39) prunes a hypothesis containing $triggered(g, _)$ whenever an inhibitor $inhibited(g, _)$ exists. For example, for the causal network


$$N_3 = \{triggered(g, t), inhibited(g, s)\},$$

and the observation $O_2 = promoted(g, t)$, the set $E = \{no_inhibitor(g)\}$ can be obtained by (23) to explain O_2 but is not consistent with (39), so is not an explanation of O_2 . Instead, $F = \{inhibited(s, t)\}$ is an explanation of O_2 obtained by (25) and (26), i.e., $g \vdash s \vdash t$. Here, we can again see that an inhibitor is preferred to a trigger; for activation of g in the presence of an inhibitor of g , introduction of a new trigger of g is not enough, but *suppression of the existing inhibitor is really effective*. This is an example of how *double inhibitions* works as a promoter, which is actually seen in many biological systems, e.g., the p53 network (47) given in Sect. 4.1.

3.3 Consistency of nonmonotonic axiomatization

When constraint (29) is incorporated into the background theory B containing a causal network N and new axioms (23)–(28), the consistency of B is “more” guaranteed than in the case of Sect. 3.1. This is because, unlike axioms (12)–(17), the new axioms contain additional defaults of the form $no_inhibitor(_)$. For example, the causal network N_2 (21) is inconsistent with the previous axiom set $\{(12)–(17), (20)\}$, but $B_2 = N_2 \cup \{(23)–(28), (29)\}$ is now consistent. Still, both $promoted(g, t)$ and $suppressed(g, t)$ can be explained from B_2 and Γ , but their explanations are not the same: $\{no_inhibitor(g)\}$ explains the former, while $\{no_inhibitor(s)\}$ explains the latter. Again, the role of (29) is to identify each no-good to prune any incompatible combination of defaults and abducibles. That is, abducing literals with the predicates $triggered$ and $inhibited$ involves default assumptions of the form $no_inhibitor(_)$, and any inconsistent set of abducibles can be detected by subsumption checking with nogoods. In N_2 , $\{no_inhibitor(g), no_inhibitor(s)\}$ becomes a nogood.

However, there are some networks that can still be proved inconsistent under the new axiom set. Suppose the network



$$N_4 = \{inhibited(g, s), inhibited(g, t), inhibited(t, s)\}, \quad (40)$$

where the first inhibitor $inhibited(g, s)$ is a suppressor for g (by (26)) and the latter two consecutive inhibitors work as a promoter for g by (25) and (26). Since there is no trigger in N_4 , no default is involved for proving $promoted(g, s)$ and $suppressed(g, s)$, and neither goal can be preferred to the other. Hence, N_4 is proved inconsistent by (29).

The necessary and sufficient condition for the consistency of a meta-level background theory can be characterized by networks with multiple “inhibitor-only paths” as follows.

¹³This is a matter of majority when there are more triggers than inhibitors for the same node. In such a case, we could still prefer an inhibitor by introducing constraint (39), but alternatively might want to allow promotion by omitting this constraint. See also the discussion in Sect. 6.5.

Proposition 4 *Suppose a meta-level theory N , which is the same as in Theorem 3. Let the background theory be $B = N \cup \{(23)–(28), (29)\}$. Then, B is consistent if and only if there are no nodes g and s in N such that there exist a proof Π^+ of $\text{promoted}(g, s)$ and a proof Π^- of $\text{suppressed}(g, s)$ satisfying that (i) both Π^+ and Π^- have no triggers; (ii) there are an even number of occurrences of inhibitors in Π^+ ; and (iii) there are an odd number of occurrences of inhibitors in Π^- .*

Proof Suppose a proof Π^+ of $\text{promoted}(g, s)$ and a proof Π^- of $\text{suppressed}(g, s)$ satisfying the conditions (i), (ii) and (iii). Then, by (29), B becomes inconsistent. Conversely, if B is inconsistent, both $\text{promoted}(g, s)$ and $\text{suppressed}(g, s)$ can be proved from B by (29). However, those proofs Π^+ , Π^- do not use axioms (23), (24) and (28), since the literal of the form $\text{no_inhibitor}(_)$ cannot be proved. Hence, no trigger appears in either Π^+ or Π^- , and (i) holds. By Lemma 1, (ii) and (iii) hold. \square

It is easy to see that the causal network N_4 (40) is inconsistent with $\{(23)–(28), (29)\}$ by Proposition 4. Actually, N_4 is a minimal inconsistent network under axioms (23)–(28) and (29), since $\{\text{inhibited}(\gamma, \psi), \text{inhibited}(\gamma, \tau), \text{inhibited}(\tau, \psi)\}$ is a minimal nogood of $N_0 = \emptyset$ with those axioms. Note that the smallest inconsistent network is

$$N_5 = \{\text{inhibited}(g, g)\}. \tag{41}$$

In N_5 , a self-loop on g suppresses g if it passes the inhibitor an odd number of times, while it promotes g if it counts the inhibitor an even number of times. Remark 3 below suggests a minimal modification of the axiom set to guarantee the consistency, and Sect. 6.5 will discuss how to resolve such inconsistencies in general.

Since inhibitors are preferred to triggers, abduction of promotion is blocked by abduction of suppression if the explanation of the former is subsumed by that of the latter.

Proposition 5 *Let N , B , and g, s be the same as in Proposition 4. Suppose the abducibles $\Gamma = \{\text{triggered}(_, _), \text{inhibited}(_, _), \text{no_inhibitor}(_)\}$ and the observations $O^+ = \text{promoted}(g, s)$ and $O^- = \text{suppressed}(g, s)$. If E^+ is an explanation of O^+ with respect to B and Γ , then there is no explanation E^- of O^- with respect to B and Γ such that (i) all defaults of the form $\text{no_inhibitor}(_)$ in E^- are included in E^+ and (ii) all inhibitors in E^- are included in E^+ .*

Proof Suppose that E^+ is an explanation of O^+ . Assume further that an explanation E^- of O^- satisfying the conditions of the theorem exists. Since O^+ and O^- violate constraint (29), $E^+ \cup E^-$ is a nogood. The condition (i) implies that the triggers in E^- are also included in E^+ because abduction of a trigger always involves a default by axioms (23), (24) and (28). By this and the condition (ii), $E^- \subseteq E^+$ holds. Then, $E^+ \cup E^- = E^+$ holds, and thus E^+ is a nogood. This contradicts with the supposition that E^+ is an explanation of O^+ . \square

Proposition 5 implies that, whenever both $\text{promoted}(g, s)$ and $\text{suppressed}(g, s)$ are explained, their explanations must contain different triggers and inhibitors, and the union of the two explanations is a nogood. For example, for the network

$$N_6 = \{\text{triggered}(g, s), \text{inhibited}(g, s)\}. \tag{42}$$

$\{\text{no_inhibitor}(g)\}$ cannot be an explanation of $\text{promoted}(g, s)$ because it is a superset of \emptyset , an explanation of $\text{suppressed}(g, s)$, so that $\{\text{no_inhibitor}(g)\}$ becomes a nogood. This is a typical effect of preferring inhibitors.

Remark 3 Proposition 4 captures the nature of inconsistent networks under axioms (23)–(28) and constraint (29), which minimizes the inconsistent cases only in the existence of two nodes connected by multiple inhibitor chains with odd and even numbers of occurrences of inhibitors. If we would further like to resolve inconsistencies in such cases, we could attach another default to axiom (25) as

$$\text{promoted}(X, Y) \leftarrow \text{inhibited}(X, Z) \wedge \text{suppressed}(Z, Y) \wedge \text{unsuppressed}(X), \quad (43)$$

and replace (25) by (43). Here, a literal of the form *unsuppressed*($_$) can be assumed whenever (43) is applied by putting it into the abducibles Γ . Then, any even number of occurrences of inhibitors must contain defaults of the form *unsuppressed*($_$), but an odd number of inhibitor occurrences may not need such a default. Hence, no network can be inconsistent under axioms (23), (24), (43), (26)–(28) and constraint (29). For example, in networks (40) and (41), the minimal nogood is obtained as $\{\text{unsuppressed}(g)\}$, which implies that only suppression is achieved on g in both cases, resulting in preferring inhibitors again.

However, this inconsistency resolution is not always welcome in real applications. A *negative feedback loop* like (41) is known to cause periodic oscillation between positive and negative states, and then it is harmful to completely ignore the possibility of promotion in such a case. This problem can only be solved by introducing time in the axioms, and will be discussed in Sect. 6.5. So far, we have a choice between (25) and (43), depending on our strategy to allow or disallow minimal inconsistencies in inhibitor-only paths.

3.4 Various inferences on causal networks

As in Sect. 2.3, axiomatization in Sects. 3.1 and 3.2 can be further combined with fact abduction to allow mixed forms of inferences. For example, to abduce a source node which promotes (resp. suppresses) some goal node g , clause (10) can be rewritten to clause (44) (resp. (45)):

$$\text{ans}(X) \leftarrow \text{promoted}(g, X) \wedge \text{source}(X), \quad (44)$$

$$\text{ans}(X) \leftarrow \text{suppressed}(g, X) \wedge \text{source}(X). \quad (45)$$

Here, a source node represented by X corresponds to an abducible with the predicate *abd*($_$), and both the predicates *promoted*($_$, $_$) and *suppressed*($_$, $_$) correspond to the predicate *caused*($_$, $_$) in (10). Then a source node at the object level is obtained as a ground term of X by answer extraction with these formulas. Combination of rule abduction and fact abduction can also be realized using consequence-finding techniques as in Sect. 2.3.

On the other hand, those nodes that can be promoted or suppressed by some source node s can be obtained by answer extraction for variable Y in the query of the form *promoted*(Y, s) or *suppressed*(Y, s). This last inference is called *prediction* rather than abduction, and can be realized by consequence finding too. Table 1 summarizes the correspondence between object-level and meta-level inferences. All types of meta-level inferences, possibly involving generation of existentially quantified hypotheses, can be realized by SOLAR. Recall that, in the context of inverse entailment (11), the negation of a goal is set as a *top clause*, and the negation of each abducible is given in a production field in SOLAR. In Table 1, “ $\leftarrow \text{caused}(_, _)$ ” in a “top clause” column is instantiated by either “ $\leftarrow \text{promoted}(_, _)$ ” or “ $\leftarrow \text{suppressed}(_, _)$ ”, and “*ans*($_$)” is an answer predicate to collect answer substitutions. In abducing object-level facts, a top clause can be further conditioned with the abducible literal “*abd*(X)” after *caused*(g, X) if the list of abducibles is given in the background theory. Like (44) and (45), abducibles are often represented as literals of the form *source*($_$).

Table 1 Correspondence between object-level inference and meta-level consequence finding

Object-level inference	Top clause in SOLAR	Production field
Rule verification	$\leftarrow \text{caused}(g, s)$	$\{\emptyset\}$
Fact abduction	$\text{ans}(X) \leftarrow \text{caused}(g, X)$	$\{\{\text{ans}(_)\}\}$
Fact prediction	$\text{ans}(X) \leftarrow \text{caused}(X, s)$	$\{\{\text{ans}(_)\}\}$
Rule prediction	unspecified	$\{\{\text{promoted}(_, _), \text{suppressed}(_, _)\}\}$
Rule abduction	$\leftarrow \text{caused}(g, s)$	$\{\{\neg \text{triggered}(_, _), \neg \text{inhibited}(_, _)\}\}$
Abducing rules and facts	$\text{ans}(X) \leftarrow \text{caused}(g, X)$	$\{\{\neg \text{triggered}(_, _), \neg \text{inhibited}(_, _), \text{ans}(_)\}\}$
Fact prediction + Rule abduction	$\text{ans}(X) \leftarrow \text{caused}(X, s)$	$\{\{\neg \text{triggered}(_, _), \neg \text{inhibited}(_, _), \text{ans}(_)\}\}$
Rule prediction + Rule abduction	unspecified	$\{\{\neg \text{triggered}(_, _), \neg \text{inhibited}(_, _), \text{promoted}(_, _), \text{suppressed}(_, _)\}\}$

In Table 1, “Rule verification” verifies if a given causal chain can be derived or not. “Fact prediction” computes *ramification* of a source s , i.e., to derive facts that can be caused by s . “Rule prediction” enumerates possible causal chains derivable from the given causal network, so a top clause is not provided in this case and characteristic clauses with respect to $\{\{\text{promoted}(_, _), \text{suppressed}(_, _)\}\}$ are just computed. The last three are mixed forms of inference with the conditional format; rules and facts are abduced or predicted under some completed rules by rule abduction. A complex example of abducing rules and facts will be shown in Sect. 4.1. The soundness and completeness of those meta-level inferences by SOLAR with respect to object-level inferences listed in Table 1 are guaranteed due to completeness of consequence finding in SOLAR (Inoue 1992; Nabeshima et al. 2010) and completeness of (conditional) answer extraction (Iwanuma and Inoue 2002; Inoue et al. 2006).

Abduction with default assumptions is also implemented in SOLAR (Inoue et al. 2006). Membership of a clause F in an *extension* of a default theory (Reiter 1980; Poole 1988)¹⁴ is guaranteed for each obtained consequence

$$F \leftarrow \text{no_inhibitor}(t_1) \wedge \dots \wedge \text{no_inhibitor}(t_m) \quad (46)$$


if $\{\text{no_inhibitor}(t_1), \dots, \text{no_inhibitor}(t_m)\}$ is not a nogood (Inoue et al. 2006, Theorem 4.5). In our meta-level abduction, F is a clause consisting of literals of such forms as $\neg \text{triggered}(_, _)$, $\neg \text{inhibited}(_, _)$ and $\text{ans}(_)$. Given a top clause C and a production field \mathcal{P} , SOLAR outputs the new characteristic clauses $\text{Newcarc}(B, C, \mathcal{P})$, and hence each new produced clause is always checked if it is not subsumed by any nogood in $\text{Carc}(B, \mathcal{P})$ (see Sect. 2.4). Hence, each clause F in a consequence of the form (46) belongs to some extension.

4 Case study: p53 signal networks

In this section, we will see that meta-level abduction can be well applied to completion of *signaling networks*. The importance of network completion in signaling networks has been

¹⁴Simply, a (*normal*) *default theory* is a pair (B, D) , where B and D are sets of first-order formulas representing a background theory and a set of defaults, respectively. An *extension* of a default theory (B, D) is the set of all consequences of $B \cup S$, where S is a maximal subset of D such that $B \cup S$ is consistent. Then, a formula F is considered as a rational consequence of (B, D) if F belongs to some extension of (B, D) .

Fig. 1 Causal network of the p53 pathway



recognized, since it is hard to observe activity levels and quantities of proteins in living organisms (Akutsu et al. 2009).¹⁵ Moreover, reporter proteins/genes are usually employed in signaling pathways, but designing and introducing reporter proteins are hard tasks. This is contrasted to the case of genetic networks, in which expression levels of most genes can be observed using DNA microarray/chip technologies.

As case studies of meta-level abduction, we consider two signaling networks containing the *p53 protein* (Prives and Hall 1999), but use it for different purposes. Although these networks are rather simple, they illustrate one of the most fundamental inference problems in Systems Biology: Given an incomplete biological network N , infer possible connections to promote or suppress the function associated with a node in N . We consider two target functions: suppression of tumors in cancer (Tran and Baral 2009) (Sect. 4.1) and switching DNA synthesis on and off (Shmulevich et al. 2002) (Sect. 4.2). Both problems are examples of *goal-oriented abduction*, in which additional links and nodes are abducted to realize given goals rather than to explain observed data.

4.1 Enumerating tumor suppressors

This subsection examines the p53 signal network presented in Tran and Baral (2009) by meta-level abduction. The p53 protein plays a central role as a tumor suppressor and is subjected to tight control through a complex mechanism involving several proteins (Meek 2009).

The p53 protein has the transactivator domain, which bounds to the promoters of target genes, then leads to protect the cell from cancer. The level and activity of p53 in the cell is influenced by its interactions with other proteins. Tumor suppression is enabled if the interacting partners of p53 do not inhibit the functionality of the transactivator domain. Mdm2 binds to the transactivator domain of p53, thus inhibiting the p53 from tumor suppression. UV (ultraviolet light) causes stress, which may induce the upregulation of p53. However, stress can also influence the growth of tumors.

These relations can be represented in solid lines of the causal network in Fig. 1. The corresponding formulas at the meta level can be simply represented by the clauses:

$$\begin{aligned}
 & triggered(cancer, uv), & & triggered(p53, uv), \\
 & inhibited(cancer, a), & & triggered(a, p53), \\
 & inhibited(a, b), & & jointly_triggered(b, p53, mdm2),
 \end{aligned}
 \tag{47}$$

¹⁵This feature of non-observability also exists in completing causal networks with respect to physical skills for cello playing (Inoue et al. 2010), in which internal causation in human bodies cannot be observed directly.

where a (“A” in Fig. 1) is the inhibitory domain of p53, and b (“B” in Fig. 1) is the complex p53-mdm2. The unit clause $jointly_triggered(b, p53, mdm2)$ can be replaced by the clause $(triggered(b, p53) \vee triggered(b, mdm2))$.

As in the setting of Tran and Baral (2009), we consider a tumor suppressor gene X (e.g., MdmX) such that mutants of X are highly susceptible to cancer. Suppose in some experiments that exposure of the cell to high level UV does not lead to cancer, given that the initial concentration of Mdm2 is high. Those initial conditions are represented as two facts,

$$source(uv), \quad source(mdm2), \tag{48}$$

that is, both UV and Mdm2 can be abducted whenever necessary.¹⁶ The meta-predicate $source$ thus behaves like the abducible predicate abd . Some other meta-level axioms can be introduced, e.g., $(no_inhibitor(S) \leftarrow source(S))$ (38). Supposing further that a high level of gene expression of the X protein is also observed, our objective in this experiment is to hypothesize about the various possible influences of X on the p53 pathway, thereby explaining how the cell can avoid cancer.

Our goal is now expressed as $\exists S(suppressed(cancer, S) \wedge source(S))$. Like (45) in Sect. 3.4, the top clause is then given in SOLAR as:

$$ans(S) \leftarrow suppressed(cancer, S) \wedge source(S). \tag{49}$$

The background theory B is now defined as the set consisting of the above causal network (47), (48), the top clause (49), the meta-axioms (23)–(28), the integrity constraint (29), the axioms for joint triggers (30), (32), (33), constraints for defaults (38), (39), two facts for defaults: $no_inhibitor(p53)$ and $no_inhibitor(x)$, and domain constraints for pruning such as $\neg inhibited(uv, Z)$ and $\neg inhibited(mdm2, Z)$. Let the abducibles be

$$\Gamma_x = \{triggered(_, _), inhibited(_, _), jointly_triggered(_, _, x)\},$$

expecting that a mutant of X bound to some Z is produced in suppressing the cancer from some source. Then the production field \mathcal{P} is set as:

$$\mathcal{P} = \{\overline{\Gamma}_x \cup \{ans(_), \neg no_inhibitor(_)\}, Cond\},$$

where $Cond$ is the length conditions such that each produced clause C must satisfy $|C \cap \{\neg triggered(_, _)\}| \leq 1$, $|C \cap \{\neg inhibited(_, _)\}| \leq 1$ and $|C \cap \{\neg jointly_triggered(_, _, x)\}| \leq 1$. Then, SOLAR produces the 26 new characteristic clauses in 10 seconds using a PC with Core 2 Duo 3.06 GHz and 4 GB RAM.

In these 26 consequences of SOLAR, the following two clauses are included:

$$ans(uv) \leftarrow triggered(x, uv) \wedge jointly_triggered(Y, p53, x) \wedge inhibited(b, Y) \wedge no_inhibitor(Y) \tag{50}$$

$$ans(uv) \vee ans(mdm2) \leftarrow triggered(x, uv) \wedge jointly_triggered(Y, mdm2, x) \wedge inhibited(b, Y) \wedge no_inhibitor(Y) \tag{51}$$

Both (50) and (51) give conditional answers. Consequence (50) represents a *definite answer* indicating that the p53- X complex has the unique source UV since both p53 and X are

¹⁶Each source is depicted with the symbol “↓” in the figure of a causal network.

caused by the same source UV. On the other hand, (51) represents a *disjunctive answer*: X is activated by UV but Mdm2 itself is assumed to be a source, hence the Mdm2- X complex has two sources. In fact, it takes more time to find consequence (51) than to find (50) in SOLAR. Those two formulas respectively correspond to the following hypotheses:

- (I) $triggered(x, uv) \wedge \exists Y (jointly_triggered(Y, p53, x) \wedge inhibited(b, Y))$,
 (II) $triggered(x, uv) \wedge \exists Y (jointly_triggered(Y, mdm2, x) \wedge inhibited(b, Y))$.

The variable Y in (I) or (II) represents a new complex synthesized from X and either p53 or Mdm2, respectively. That is, meta-level abduction generates a new mutant Y (“C” in Fig. 1) by combining X and either p53 or Mdm2, which then inhibits the existing complex b . Those two hypotheses are actually suggested in Tran and Baral (2009): (I) X directly influences p53 protein stability: UV causes stress then induces high expression of X , which then binds to p53, so p53 is stabilized and formation of Mdm2-p53 complex is prevented; (II) X is a negative regulator of Mdm2: UV causes stress then induces high expression of X , which then binds to Mdm2, which competes against inhibiting the Mdm2-p53 interaction (depicted in dashed lines in Fig. 1). In both cases, p53 (or “A”) can be functional as a tumor suppressor. In the biological viewpoint, however, the hypothesis (I) seems preferred because p53 has more chances to be bound to other proteins.

We again stress that all abduced links in hypotheses (I) and (II) are newly generated by meta-level abduction and that the new node Y is automatically invented during inference. In contrast, Tran and Baral (2009) consider an action theory for this p53 network, and prepare all possible ground candidate nodes and links as abducibles in advance. Hence, no new node is invented in the framework of Tran and Baral (2009).

Concerning other 24 explanations obtained by SOLAR, 14 solutions promote a by inhibiting b , and 10 inhibit cancer directly. Although direct inhibition of cancer in the latter case is feasible, we consider that revival of p53 as a tumor suppressor by promoting a is more important, since it relies on *double inhibition* of a through inhibition of b , which is biologically more admissible. Then, in such 14 explanations, 6 solutions have the definite answer $ans(uv)$, 6 have the definite answer $ans(mdm2)$, and 2 have the disjunctive answer $ans(uv) \vee ans(mdm2)$. In these solutions, there are simple ones that never produce complexes with X , e.g.,

$$ans(uv) \leftarrow inhibited(b, x) \wedge triggered(x, uv),$$

but they are not very useful, since *inhibition is usually achieved by a mutant of some protein combined with another protein* in biological systems. One of the disjunctive answers has the same abducibles as in (50) except that X is triggered by Mdm2 instead of UV:

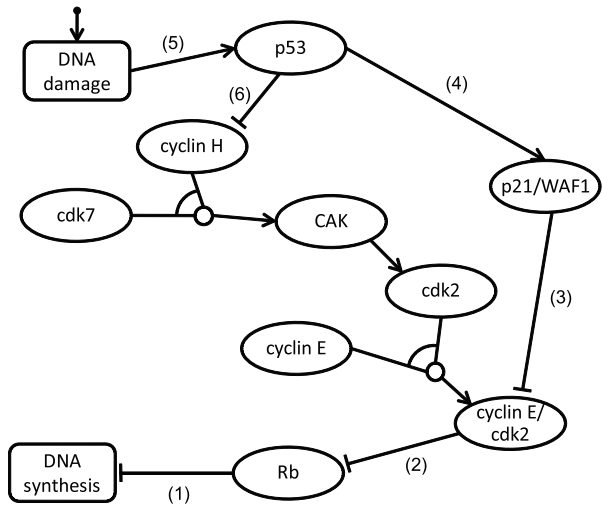
$$ans(uv) \vee ans(mdm2) \leftarrow triggered(x, mdm2) \wedge jointly_triggered(Y, p53, x) \wedge inhibited(b, Y) \wedge no_inhibitor(Y).$$

One of other interesting solutions is

$$ans(uv) \leftarrow jointly_triggered(b, Z, x) \wedge inhibited(Z, p53),$$

which represents that suppression of b is realized by a complex of X and a new protein Z that is inhibited by p53. In this explanation, the additional axiom (32) works well.

The p53 regulatory network includes a complex array of upstream regulators and downstream effectors. There are versatile functions of p53 in the suppression of tumors; deriving

Fig. 2 Causal network of the CDK pathway

the cell to apoptosis and ensuring the normal cell growth as a guardian of life. The stabilization of p53 seems to be the result of inhibition by Mdm2 and Mdm4, while activators such as HIPK2 and DYRK2 increase the p53 response. This complex system of activation/inhibition shows that p53 is a bottleneck of many regulatory mechanisms. The results obtained in this section are important in the sense that the activation/inhibition mechanism of p53 is linked to some proteins that might not have been found out yet. Meta-level abduction is thus crucial for this discovery task, and inferred hypotheses can suggest to scientists necessary experiments with gene knockout mice as minimally as possible.

4.2 Recovering links in CDK networks

The next case study is completion in the switch network of *cyclin-dependent kinases* (CDKs) (Shmulevich et al. 2002). CDKs are kinase enzymes that modify other proteins by chemically adding phosphate groups to them (phosphorylation), and are involved in the regulation of the cell cycle. A CDK is activated by association with a cyclin, forming a cyclin-CDK complex (Fig. 2).

The cdk2/cyclin E complex inactivates the retinoblastoma (Rb) protein, which then releases cells into the synthesis phase. Cdk2/cyclin E is regulated by both the positive switch called CAK (cdk activating kinase) and the negative switch p21/WAF1. p21/WAF1 is activated by p53, but p53 can also inhibit cyclin H, which is a source of the positive regulator of cdk2/cyclin E. The negative regulation from p53 works as a defensive system in the cells: when DNA damage occurs, it triggers p53, which then turns on the negative regulation to stop DNA synthesis, so that the damage should be repaired before DNA replication to prevent passing damaged genetic materials onto the next generation.

For the CDK network represented in Fig. 2 with *source*(dna_damage), several experimental problems are designed by removing some links from Fig. 2. Then, meta-level abduction is applied to verify how those removed links can be recovered and whether other interesting links and unknown missing links can be inferred in explaining the observation

suppressed(dna_synthesis, dna_damage).

Table 2 Results of recovering links in the CDK pathway

Removed links	<i>LCond</i>	#H	Time [sec]	Minimal recovered links		
				p21-mediated	cdk2 activate	cyclin E inhibit
(1)	$len \leq 1$	4	10.7	(1)	–	–
(1), (2)	$len \leq 2$	9	18.8	(1g), (2g)	–	–
(1), (3)	$len \leq 2$	17	18.5	(1), (3)	(1), +	(1), (7)
(2), (4)	$len \leq 2$	24	20.3	(2), (4)	(2), +	(2), (7)
(4), (6)	$len \leq 2$	22	22.3	(4)	(6), +	(7)
(1), (2), (3), (4), (5)	$len \leq 5$	392	138.4	(1g), (2g), (3g), (4), (5)	(1g), (2g), (5), +	(1g), (2g), (5), (7)

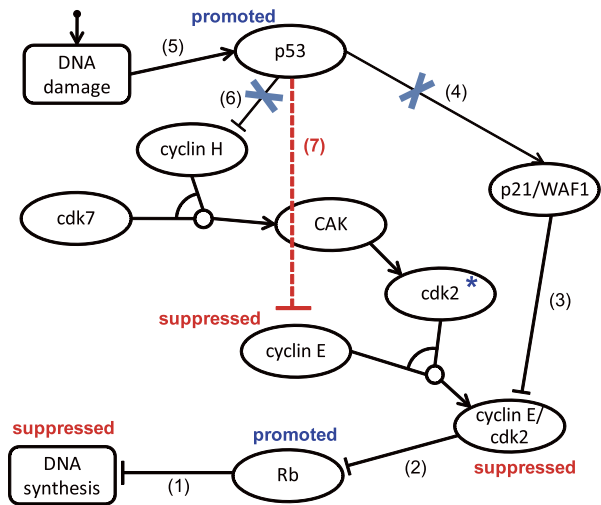
The objective of this experiment is to show how meta-level abduction can be well applied to complete incomplete networks. Recovery of removed links is a good testbed for this purpose because the existing natural system can be considered as an ideal solution. Yet, looking at other hypotheses, we understand that the same functions can be realized in different ways.

Table 2 shows a part of experimental results. All experiments are done in the environment of Mac mini with Core 2 Duo 1.83 GHz and 2 GB RAM. The maximum search depth of SOLAR is set to 5 for computing the characteristic clauses and 10 for computing the new characteristic clauses. The table shows 6 problems, each of which is given as a network obtained by removing the links shown in the table. In the table, the production field for each problem is defined as follows:

$$\mathcal{P} = \{ \{ \neg triggered(_, _), \neg inhibited(_, _), \neg no_inhibitor(_) \}, LCond \wedge Occ-Cond \},$$

where *LCond* is a condition on the maximum length *len* of each consequence (described in the table) and *Occ-Cond* = $|C \cap \{ \neg no_inhibitor(_) \}| \leq 1$, that is, the number of occurrences of (the negation of) defaults in each consequence *C* is at most 1. The “#H” column shows the number of new characteristic clauses (minimal explanations) obtained by SOLAR. “Time” is computation time to obtain all minimal hypotheses and to check their consistencies.

The results of recovering removed links are shown in the table. The symbol “+” means the default *no_inhibitor*(cdk2), and link (7) is *inhibited*(cyclin_e, p53). Two nodes are declared to have no inhibitors in the axioms so that they do not need to be included in each explanation: *no_inhibitor*(p53) and *no_inhibitor*(p21_WAF1). Here, among the number #H of explanations, we pick up three important groups of solutions depending on suppression paths from dna_damage to dna_synthesis. The first solution group goes through the right path via p21/WAF1, i.e., (5)-(4)-(3)-(2)-(1). The second one uses two cyclin-cdk joint triggers via inhibition (6) of cyclin H by p53. The third one generates the new link (7) which inhibits cyclin E by p53. For instance, when links {(1), (3)} are removed, there exists the explanation in the first group that contains exactly the same links as the removed ones among the 17 hypotheses. For a link (N) for N = 1, 2, 3, the recovered link (Ng) means that a more general hypothesis than (N) is obtained. For example, when {(1), (2)}, i.e., *inhibited*(dna_synthesis, rb) \wedge *inhibited*(rb, cyclin_e/cdk2) is removed, {(1g), (2g)}, which is $\exists X (inhibited(dna_synthesis, X) \wedge inhibited(X, cyclin_e/cdk2))$, is recovered. Under axiomatization (23)–(28), explanations in the second group always contain the default *no_inhibitor*(cdk2) (+). By this reason, when links {(1)} or {(1), (2)} are removed, no minimal explanation via the second path is obtained because any explanation is always subsumed by an explanation in the first group without the default. We

Fig. 3 p53-cyclin E-mediated pathway

can observe that two consecutive links are replaced by general ones with existentially quantified variables at recovery, e.g., $\{(1), (2)\}$, but links that are not connected are recovered as they are in the first group, e.g., $\{(1), (3)\}$. Notice that only either (4) or (6) is recovered at the first or second group, respectively, when links $\{(4), (6)\}$ are removed. Actually, the negative regulation by p53 has the two paths to suppress cdk2/cyclin E via (4) and (6), which correspond to the first and second paths, respectively. This shows that the biological network in Fig. 2 is *robust*, and the system can still work well even if one of the two paths is cut. In other words, if these two are cut simultaneously, only one recovery is logically sufficient to realize this function, and SOLAR outputs such minimally sufficient hypotheses in this case. Interestingly, we can also find other recoveries by connecting nodes that have not been originally connected. The third solution group realizes such a case by newly imposing inhibitor (7) (Fig. 3).

A well-known effect of p53 expression is a block in the cell division cycle, and induction of p53 leads to cell growth arrest and apoptosis. Although the function of p53 to induce a G1 arrest by binding with p21 has been well defined, the control of cyclin E by p53 has not been established yet. In our experiment, we have applied abductive inference on an incomplete p53 network related to its activation by DNA damage, and one of the results obtained in Fig. 3 clearly shows that p53 degrades the cyclin E protein by finding (7) and consequently imposes cycle arrest. This is a biologically important function verified by our method. A possible explanation here is based on *analogical reasoning* such that the inhibition of cyclin H from p53 in the same network can suggest a possible inhibition of cyclin E from p53. This analogy could be regarded as a form of inductive reasoning because it strives to provide understanding of what is likely to be true. We note that link (7) is automatically generated in our experiment without giving such an analogical bias.¹⁷ In the literature, examples of this function can be found in genetic analysis of *Drosophila melanogaster* (Mandal et al. 2005), which has shown that disruption of the mitochondrial electron transport chain activates a G1–S checkpoint as a result of control of cyclin E by p53, and in a recent evidence (Mandal

¹⁷We can also supply such a bias as $\text{inhibited}(X, p53) \wedge \text{cyclin}(X)$ in our framework by giving facts $\text{cyclin}(\text{cyclin_h})$ and $\text{cyclin}(\text{cyclin_e})$.

et al. 2010) that p53 modulates the activity of the ubiquitin-proteasome system to degrade cyclin E and thereby imposes cell cycle arrest. Our results hence explain those results from a logical point of view.

In this subsection, we have verified that previously known solutions, either well-known or just reported in the literature, can be generated by our abductive method. In other words, meta-level abduction has the ability to recover or reproduce such solutions *from an incomplete network*. The experiments show that such known solutions can be found among many other solutions, so recovery rate would be low if our task were to recover exactly the same removed links. From the biological viewpoint, however, this phenomenon is not surprising and can be explained as follows. The biological mechanism is so complex that any configuration of a biological network cannot be optimally organized *per se*. Once a biological function is lost in a network for some reason, it is not naturally recovered as it was. Actually, there are many possibilities to realize the same function, since the inverse problem does not have a unique solution in general. Moreover, biological systems are *robust* and have much *functional redundancy*, and survival of individuals and species preservation are of overriding importance for them (Wagner 2005). In other words, a biological system is *resilient*, that is, it has the ability to endure and successfully recover from perturbations such as genetic and environmental changes. Having multiple stable states (solutions) in a biological system, transition to a new stable state occurs in the face of changes so that they can survive and evolve. Our results having many alternative solutions can thus be understood as a logical account for an example of biological resilience. Now, *the purpose of abduction here is not recovery of particular links but recovery of the function itself*, i.e., suppression of DNA synthesis. Hence, it is essential for the abductive system to guarantee that important solutions are never lost. Under this completeness condition, it is better to restrict the number of possible solutions as many as possible, and we will discuss this important issue in Sect. 7. Finally, biologists often decompose networks into smaller sub-networks to make biological analysis easier, and thus even those simplified p53 networks in this section are useful for considering all possible connectivities that are hard to find for human experts.

5 Experiments of network completion

This section addresses the issue of scalability of meta-level abduction, which is important when applications involve large networks. To this end, we show experiments of network completion by meta-level abduction on networks randomly generated with some parameters varied, and see scalability of the method.

For experiments, networks are generated randomly, each of which consists of the specified numbers of nodes and edges including both positive and negative links. We consider 70 combinations of the numbers of nodes and edges, whose ratios are among 3 : 1, 2 : 1, 1 : 1, 1 : 2, 1 : 3, 1 : 4 and 1 : 5 (7 ratios), each varies from 10 nodes to 100 nodes (10 cases). Note that each ratio determines the average *degree* (the number of adjacent nodes connected to a node), but it does not mean that the degree is always fixed to every node. The consistency of each network instance together with axioms (23)–(28) and constraints (29, 39) is checked when the network is generated; if a randomly generated network is proved inconsistent by SOLAR, that is, the empty clause is produced as the unique characteristic clause, the next network is randomly generated until a consistent one is found. By this way, 30 consistent random network instances are generated for each test case. We set the timeout for each run of SOLAR on one network instance to 1,800 seconds. Then, problems with ratio 1 : 5 and 100 nodes (500 edges) cannot be solved within the time limit. We also do not allow bidirectional links between any pair of nodes like $triggered(a, b) \wedge inhibited(b, a)$, so that there




Fig. 4 Total computation time of network completion

is at most one link, which is either a trigger or an inhibitor, is allowed between two nodes. Then, there are ${}_{10}C_2 = 45$ edges for any complete graph with 10 nodes, and we cannot generate any network for problems with ratio 1 : 5 and 10 nodes (50 edges). Hence, we test $7 \times 10 - 2 = 68$ cases, and $68 \times 30 = 2,040$ network instances are generated.

The observation for each network instance is given as *suppressed*($n - 1, 0$) where n is the number of nodes in the network, supposing that 0 is the source node and $n - 1$ is the goal node. We limit the maximum length of hypotheses to 3, which is considered reasonable in many real applications. The maximum search depth of SOLAR in running each network instance is set to 5 for computing the characteristic clauses and 8 for computing the new characteristic clauses. Running environment is Core 2 Duo 1.83 GHz with 2 GB RAM. Figure 4 is the graph of results with all 7 nodes-edges ratios on the average computation time to compute all explanations for each test case with 30 network instances. Each time includes computation of generating all possible hypotheses and their consistency checking. Figure 5 is the graph of results on the average number of explanations generated by network completion for each test case. The ratio of m nodes to n edges is denoted as “ $m : n$ ” in these graphs.

From Figs. 4 and 5, it is observed that, as the number of nodes (and edges) increases, both the CPU time and the number of hypotheses grow in all cases of ratios. As the degree of a network increases, completion of intermediate nodes in the network occurs more frequently and then time and hypotheses also increase. In particular, when new nodes are created, possible new connections between the source and the goal become larger. However, as the ratio of nodes to edges becomes smaller, e.g., 1 : 4 and 1 : 5, that is, *as the average degree of a network increases, the growth rates of hypotheses decrease and tend to converge*. This is a remarkable result in this experiment, and this dynamics is explained as follows. Under the limit of inference depth and maximum length of consequences, there is no room for further addition of new edges into such a “dense” network by keeping the consistency, so the number of hypotheses does not increase too much.

It has been recognized that the complexity of *enumeration problems* may not be defined in terms of functions of the size of the input because the number of solutions can be exponential in the input size (T'kindt et al. 2007). For example, Johnson et al. (1988) defines that an enumeration problem can be solved in polynomial total time if an algorithm exists with




Fig. 5 The number of hypotheses generated by network completion




Fig. 6 Average time required to find one hypothesis in network completion

running time bounded by a polynomial function of the combined size of the input and the output. With this regard, Fig. 6 shows the graph of average time to generate one hypothesis, which is total running time divided by the number of generated hypotheses. The figure indicates that time to compute each hypothesis in enumeration does not grow exponentially with increase of input network sizes for all cases of networks. In Fig. 6, when there are fewer nodes, e.g., 10–30, the running time is short and then there are some overheads of file read and start-up of Java Virtual Machine. On the other hand, when many nodes exist (90–100), time for subsumption checking (consistency checking) increases due to a large number of hypotheses and then average time to compute one hypothesis also increases, but this happens only when the ratio of nodes to edges becomes smaller (1 : 3–1 : 5). Hence, it is possible to argue that enumeration scalability can be observed in this experiment of network completion.

Finally, it should be noted that no domain-dependent knowledge other than an initial network is incorporated in each run of this experiment. In real-world applications, however, we expect that the number of hypotheses can be decreased in general by incorporating more background knowledge and constraints and by specifying more restricted production fields.

6 Discussion and related work

6.1 Rule abduction

The method of *rule abduction* by means of meta-level abduction was firstly introduced in Inoue et al. (2010), and this paper has extended it to deal with positive and negative causal links. Although few works on rule abduction exist previously, they focus on positive effects only unlike our work. Moreover, the patterns of *abducible rules* must be determined in advance. The framework of abductive systems with such abducible rules were firstly considered in Poole (1988), which associates a unique name with each ground instance of a rule schema that is predetermined as possible hypotheses, and those names are treated as abducible atoms in abductive reasoning. This is a convenient method when we know exact patterns of rules as strong biases. However, it is impractical or impossible to prepare all patterns of rules in advance in order to abduce missing rules, and predicate invention is basically impossible by this method.

Meta-reasoning has been discussed intensively in logic programming (Kowalski 1990; Hill and Gallagher 1998; Costantini 2002), but rule abduction has never been considered in the literature of meta-reasoning. While the importance of both abduction and meta-reasoning is discussed in Kowalski (1990), their combination has not been discussed explicitly. Actually it is possible to embed abductive procedures into meta-programming to perform abduction, but this is distinguished from abduction at the meta level. Christiansen (2000) computes both abduction and induction in a unified system of meta-programming, and uses separate forms of reasoning in their actual computation followed by the system. In Christiansen (2000), the *demo* predicate is used to generate parts of programs that are necessary to derive the goal based on techniques of constraint logic programming.

As a predecessor of this work, applications of SOLAR to complete networks are discussed in Ray and Inoue (2007), but negative causes and effects are not handled and joint (positive) causes are not considered there. CF-induction (Inoue 2004) can induce explanatory rules, and its applications to completion of causal rules and abduction of network status in metabolic pathways are shown in Yamamoto et al. (2010). CF-induction can directly induce first-order full clausal theories at the object level, but predicate invention and hypothesis enumeration at the object level are not easier than those by meta-level abduction. This is because predicate invention is realized by *inverse resolution* (Muggleton and Buntine 1988) and the search space for hypothesis enumeration is huge in general.

6.2 Combining abduction and induction

Abduction is also used to induce rules in several ILP systems other than CF-induction (Inoue 2004). Earlier works are summarized in several papers in Flach and Kakas (2000), see, e.g., Mooney (2000), Sakama (2000). Most such previous systems use abduction to compute inductive hypotheses efficiently, and such integration is useful in theory refinement. More recently, those works (Yamamoto 2000; Ray et al. 2003; Kimber et al. 2009; Corapi et al. 2010) use procedures for *abductive logic programming* (ALP), which are different from

SOLAR that is used for full clausal abductive theories as in our meta-level abduction. Progol (Muggleton 1995) is the first ILP system based on inverse entailment, and Muggleton and Bryant (2000) improves it to realize theory completion. ILP systems in Yamamoto (2000), Ray et al. (2003), Kimber et al. (2009) extend Progol's applicability to larger classes of induced logic programs, which are yet limited to a set of Horn clauses. Such extensions do not necessarily depend on ALP, and Ray and Inoue (2008) uses SOLAR instead of ALP to deal with non-Horn clauses. These systems abduce part of hypotheses from the background theory together with the negation of an observation. In Corapi et al. (2010), an ILP problem is fully translated into an equivalent ALP problem, and then an ALP procedure is used to abduce a hypothesis for an observation from the background theory and an inductive bias called a *top theory*. Their system TAL is an extension of TopLog (Muggleton et al. 2008) to allow for multiple clause hypotheses with negation-as-failure (or default negation). To abduce rules in ALP, TAL uses the naming method as in Poole (1988), thus depends on a strong bias to induce rules, while our meta-level abduction defines the abducibles as meta-predicates only.

There are other types of works on combining ALP and ILP. Dimopoulos and Kakas (1996) propose a general framework in which a new abductive logic program is learned, given a previous abductive program as background knowledge and training examples. Their framework is followed by Kakas and Riguzzi (1997), Lamma et al. (2000). On the other hand, Inoue and Haneda learn an abductive program from a non-abductive logic program with positive and negative examples, and new abducibles are acquired there (Inoue and Haneda 2000). The covering relation in Dimopoulos and Kakas (1996), Kakas and Riguzzi (1997) and Lamma et al. (2000) is defined based on *abductive entailment*, while Inoue and Haneda (2000) did not change the entailment relation from the ordinary one in extended logic programs.

6.3 Abduction and induction in causal theories

Some attempts in ALP, ILP and AI have contributed to abduction and induction in causal theories. In these works, a causal theory is formally represented in a more specialized way as a set of individual rules between causes and their effects. The *event calculus* (Kowalski and Sergot 1986) is a meta-theory for reasoning about time and action in the framework of logic programming. *Abductive event calculus* is an abductive extension of event calculus (Eshghi 1988), and can thus be regarded as a kind of meta-level abduction. Abductive event calculus has been extended for applications to planning, e.g., Shanahan (2000), but has never been used for abducing causal theories.

Moyle (2003) uses a theory completion technique of Muggleton and Bryant (2000) and Yamamoto (2000) to learn a causal theory in the form of logic programs based on event calculus, given positive examples of input-output relations. In this work, a complete initial state is required as an input and a complete set of *narrative* facts is computed in advance, and thus observations handled in our work cannot be explained. Otero (2005) considers causal theories represented in logic programs in the case of incomplete narratives. These previous works need either *frame axioms* or *inertia rules* in logic programs. The former causes the frame problem and the latter requires induction in nonmonotonic logic programs. Inoue et al. (2005) induce causal theories represented in an *action language* given an incomplete action description and observations, but requires an algorithm to learn finite automata to compute hypotheses, which may search the space of possible permutations of actions. Unlike our meta-level abduction, all these previous works on induction of action theories do not consider invention of new events or objects.

Tran and Baral (2009) use an action language which formalizes notions such as *causal rules* of the form “X causes Y”, *trigger rules* of the form “X triggers Y”, and *inhibition rules* of the form “X inhibits Y” to model cell biochemistry, and apply it to hypothesize about signaling networks presented in Sect. 4.1. As discussed in Sect. 4.1, all ground candidate nodes and links to be added are prepared as abducible causal/trigger/inhibition rules in advance in Tran and Baral (2009). In contrast, all abduced causal relations as well as new nodes are automatically generated in our meta-level abduction.

6.4 Completion of biological networks

There are several works on completing biological networks. Notably, work on Robot Scientist (King et al. 2004) adopts abduction to complete biochemical pathways. The abduction mechanism of Robot Scientist detailed in Reiser et al. (2001) suggests the use of SOLDR resolution (Yamamoto 2000), which is a version of SOL resolution (Inoue 1992) restricted for Horn theories. In Reiser et al. (2001), a *reaction* is defined as a pair of sets of compounds representing the substrates and the products of the reaction, and a *metabolic graph* is defined in such a way that each node is given as a set of compounds available by sequences of reactions. This representation is used to deal with biochemical networks represented as hyper-graphs and to jointly propagate causes of the parent nodes to those of their child nodes. However, since each node does not correspond to one compound but represents a set of compounds, the number of nodes in a metabolic graph becomes large. In contrast, we adopt causal networks, which are simpler than metabolic graphs, and the meta-predicate *jointly_connected* or a disjunction of *connected* literals is used to represent a direct multi-causal relationship by (7) or (4). Moreover, our causal network deals with inhibition, while Reiser et al. (2001) does not take negative effects into account.

Metabolic pathways are updated by several reasons. As discussed in Sect. 1, known metabolic pathways are often incomplete so that they should be completed (King et al. 2004). New experimental data often reveal that previously known pathways are incorrect so that they should be revised as opposed to merely extended (Ray et al. 2010). Moreover, biological systems are robust to perturbations and can survive in critical situations like loss of some important genes or links by finding bypasses in pathways (Ishii et al. 2007) and by creating novel reactions that are not normally used (Nakahigashi et al. 2009). In computational models, metabolic pathways are completed in Schaub and Thiele (2009) and are revised in Ray et al. (2010) using *answer set programming* (ASP). These works do not invent new nodes, and the work (Schaub and Thiele 2009) does not consider inhibition. Furthermore, in revising pathways in Ray et al. (2010), deletion is realized through addition of atom δ_R representing retractability of rule R , where negation-as-failure $not\delta_R$ appears in the body of each retractable rule R represented in ASP. This kind of deletion is only conceivable under the assumption that we can declare retractable links that are subject to change in advance. However, real deletion of connections between metabolites is generally impossible *in vivo* once they have been established. Instead of deleting links from a network, our method *adds* new links and nodes by abduction, yet the switch between presence (activation) and absence (inhibition) of a node can be controlled by our method due to the nonmonotonic axiomatization (23)–(28). Our abductive method is then more realistic, since we can simulate the biological effect of double inhibition by inhibiting inhibitors in the p53 network in Sect. 4.1. Abduction in metabolic pathways with inhibition is also considered in Tamaddoni-Nezhad et al. (2006), although the problem setting in Tamaddoni-Nezhad et al. (2006) is to infer the state of each reaction, which is different from network completion and no new links/nodes are abduced there. Moreover, this kind of qualitative abstraction is a very difficult task, and

indeed the axiomatization in Tamaddoni-Nezhad et al. (2006) can be inconsistent with some networks and observed data, as is discussed in Ray (2009).

For *gene regulatory networks*, Gat-Viks and Shamir (2003) determine a class of regulation functions, by which regulators determine transcription, and analyze their complexity. Zupan et al. (2003) construct networks from mutant data using abduction, but use experts' heuristic rules for construction. These works are different from ours in that they use more specific methods at the object level to reconstruct networks, depending on the problem domains. On the other hand, completion of *signaling networks* is analyzed in a general way by Akutsu et al. (2009), in which unknown Boolean functions are guessed in a *Boolean network* (Kauffman 1993) whose network topology is known. This contrasts with our setting that a network is incomplete and its topology is not fixed.

6.5 Extension of network representation

Besides completion of biological networks, logical inference on networks has been considered with several different formalizations in the literature. Leitgeb (2001) defines *inhibition nets* as a special case of artificial neural networks to represent inhibitory connections between nodes and excitatory connections. Inhibition nets can represent some nonmonotonic logics and can simulate logic programs. In other words, inhibitors can be used to realize nonmonotonic reasoning, and can replace negation-as-failure in logic programs to some extent. On the other hand, Fayruzov et al. (2010) formalize inference on gene regulatory networks using ASP by utilizing nonmonotonic behavior with negation-as-failure. In Fayruzov et al. (2010), each particular inference problem is translated into an equivalent computational problem in ASP at the meta level. Inoue (2011) shows that logic programs with negation-as-failure can simulate Boolean networks (Kauffman 1993) and vice versa. These works support a claim that inhibition and negation-as-failure can simulate each other, and that inference on networks with inhibitors is nonmonotonic.

Although we have shown two axiomatizations in Sect. 3, other formalizations can be still considered for controlling inference in different ways. In fact, we have noticed that there are at least two arguable cases using the axiomatization in this paper. The first one is a competing case that there are multiple triggers and one or fewer inhibitors for a node g . Then, instead of preferring the unique inhibitor to all triggers as in Sect. 3.2, we could use the *majority function* or assign *probabilities* to them to determine the value of g . Setting a *threshold* on a node is also useful, so that a node is activated if the sum of input signals for the node exceeds the threshold. Such a mechanism is implemented using ASP in Fayruzov et al. (2010). Alternatively, guess of a Boolean function for updating the truth value of g in the framework of Boolean networks is another interesting direction in this case, as shown in Akutsu et al. (2009). Network N_4 (40), which is inconsistent with (23)–(28) and (29), can also be handled in such a way. The use of *paraconsistent logic* (Damásio and Pereira 1998) would also help to infer some meaningful information from part of knowledge that are not related to inconsistency.

The second important case is a *negative feedback loop*, in which a node g depends negatively on itself. The smallest example of negative feedback loops is $N_5 = \{\textit{inhibited}(g, g)\}$ (41), which is inconsistent by Proposition 4. Interestingly, if we interpret N_5 as a Boolean network, then any stable state (called a *point attractor*) of N_5 is characterized by a model of the Boolean equation $(\neg g \equiv g)$ (Inoue 2011). Since this is unsatisfiable, N_5 has no stable state. Actually, a negative loop is the source of periodic oscillation (called a *cycle attractor*) in Boolean networks. Hence, it is admissible to return no explanation for an observation in a network containing negative feedback loops, as long as we are concerned with solutions

in a static or equilibrium state. On the other hand, to represent *dynamic behaviors* in our framework, we need to introduce *time* for causal axioms. For example, axiom (25) can be reexpressed as

$$\textit{promoted}(X, Y, T + 1) \leftarrow \textit{inhibited}(X, Z) \wedge \textit{suppressed}(Z, Y, T).$$

by incorporating the time argument. Lejay et al. (2011) have examined such an extended axiomatization in an application of meta-level abduction to treatment of hypertension, but a formal work is needed to explore this extension.

7 Conclusion

The method of meta-level abduction (Inoue et al. 2010) has been extensively explored in this paper. In particular, we have allowed representation of positive and negative effects in causal networks, and deductive and abductive methods have been investigated for reasoning about networks. Two axiomatizations for causal networks have been presented: one treats positive and negative effects equally and the other prefers inhibitors to triggers. The latter involves default reasoning on triggers. With this extension, nonmonotonic reasoning in causal networks and completion of positive and negative links are now possible by meta-level abduction.

Besides reasoning with positive and negative causal effects discussed in this paper, meta-level abduction has several advantages in abduction. For example, multiple observations are explained at once, full clausal theories can be allowed for background knowledge, both rules and facts can be abduced at the object level, and predicate invention is partially realized as existentially quantified hypotheses.

Problem solving with meta-level abduction consists of three steps: (I) design of meta-level axioms; (II) representation of domain knowledge at the meta level; and (III) restriction of the search space to treat large knowledge. This work supposes an incomplete network for Step (II), and hence representation of a problem is rather tractable. On the other hand, we have made great effort on Step (I), by gradually extending the positive causal networks (Sect. 2) to alternating axioms (Sect. 3.1) and then to nonmonotonic formalization (Sect. 3.2). This formalization process was based on trial and error by running SOLAR many times, and case studies in Sect. 4 were used as testbeds for this purpose. Step (III) is important not only for more efficient computation of hypotheses but for making the user of meta-level abduction to select appropriate hypotheses more easily. This goal is achieved by introducing more constraints, and it is important to explore more useful methods for inducing such constraints. In Lejay et al. (2011), a general method to generate constraints is proposed to prune many useless solutions by assuming that given link information does not change so that a trigger cannot become an inhibitor and vice versa. For example, if $\textit{triggered}(g, t)$ is true in the given network, then $\neg\textit{inhibited}(g, t)$ is assumed automatically. If we further assume non-existence of negative feedbacks, the constraint $\neg\textit{inhibited}(t, g)$ is further added there. These two additional constraints are, however, still weak because they can be also obtained as characteristic clauses for network N_1 in Sect. 3.1. Then, a stronger constraint $\neg\textit{triggered}(t, g)$ is introduced by assuming that the direction of a link is not changed, but is often too restrictive and cannot be applied in a complicated case. Domain-dependent constraints are more powerful, and we have seen in Sect. 4.1 which hypotheses are more important than others, and have shown that analogical biases are useful in Sect. 4.2. Incorporation of these guides would be a promising way to focus on particular

solutions. Thus, discovering a more general and useful method for constraint generation is left as a future work.

One very important topic that we could not describe much in this paper is evaluation of logically possible hypotheses obtained in meta-level abduction. To evaluate hypotheses, some statistical methods have been proposed. For example, hypotheses can be ranked according to frequencies of literals appearing in them and corresponding paths (Inoue et al. 2009), and can be given their scores according to their fitness with observed data (Gat-Viks and Shamir 2003). In general, obtained hypotheses can be used to explain more than given observations. For example, cross-validation can be used to evaluate hypotheses as long as observed data can be used for both training and testing, although case studies in Sect. 4 are not of this kind. Often, within a certain maximum number of alternative solutions, providing a set of *diverse solutions* rather than a few most promising solutions is rather a good idea. To this end, we need more domain-dependent or general constraints and heuristics to restrict the search space and the number of possible solutions, so the future work mentioned above is also important here.

This paper has shown applications of meta-level abduction to signaling pathways, but we expect that the proposed method can also be applied to abduction in metabolic pathways with inhibition (Tamaddoni-Nezhad et al. 2006; Yamamoto et al. 2010; Ray et al. 2010) as well as transcription networks (Gat-Viks and Shamir 2003). Another interesting application is to understanding and modeling of robustness of biological systems during severe and unexpected environments (Wagner 2005; Ishii et al. 2007; Nakahigashi et al. 2009). We hope that meta-level abduction could contribute to a future breakthrough in sciences including biology and medicine by helping scientists to discover important missing and unknown networks.

Acknowledgements The authors would like to thank many people who had discussions on a wide range of topics related to this work: Tomoya Baba and Kenji Nakahigashi for their comments from viewpoints of biologists, Domenico Corapi, Koichi Furukawa, Bob Kowalski and Stephen Muggleton for discussions on integration of meta-reasoning and abduction, Jean-Charles Faye and Olivier Sordet from L'Institut Claudius Regaud for their verification of our results on completed p53 pathways, Koji Iwanuma and Oliver Ray for their contribution on conditional query answering and disjunctive answers, Hiroshi Maruyama and colleagues of the Systems Resilience project for discussions on robustness, Philip Reiser for his early attempt to complete biological networks by SOLAR, and the reviewers of this paper for their many comments and suggestions.

References

- Akutsu, T., Tamura, T., & Horimoto, K. (2009). Completing networks using observed data. In *LNAI: Vol. 5809. Proceedings of ALT'09* (pp. 126–140). Berlin: Springer.
- Christiansen, H. (2000). Abduction and induction combined in a metalogic framework. In *Abduction and induction—essays on their relation and integration* (pp. 195–211).
- Corapi, D., Russo, A., & Lupu, E. C. (2010). Inductive logic programming as abductive search. In *Technical communications of ICLP'10* (pp. 54–63).
- Costantini, S. (2002). Meta-reasoning: a survey. In A. C. Kakas & F. Sadri (Eds.), *LNAI: Vol. 2408. Computational logic: logic programming and beyond, essays in honour of Robert A. Kowalski, Part II* (pp. 253–288). Berlin: Springer.
- Damásio, C. V., & Pereira, L. M. (1998). A survey of paraconsistent semantics for logic programs. In D. M. Gabbay & P. Smets (Eds.), *Handbook of defeasible reasoning and uncertainty management systems* (Vol. 2, pp. 241–320). Berlin: Springer.
- Dietterich, T., Domingos, P., Getoor, L., Muggleton, S., & Tadepalli, P. (2008). Structured machine learning: the next ten years. *Machine Learning*, 73, 3–23.
- Dimopoulos, Y., & Kakas, A. (1996). Abduction and inductive learning. In L. De Raedt (Ed.), *Advances in inductive logic programming* (pp. 144–171). Amsterdam: IOS Press.

- Eshghi, K. (1988). Abductive planning with event calculus. In *Proceedings of the 5th international conference on logic programming* (pp. 562–579).
- Fayruzov, T., Janssen, J., Cornelis, C., Vermeir, D., & De Cock, M. (2010). Extending Boolean regulatory network models with answer set programming. In *Proceedings of 2010 IEEE international conference on bioinformatics and biomedicine* (pp. 207–212).
- Flach, P. A. & Kakas, A. C. (Eds.) (2000). *Abduction and induction—essays on their relation and integration*. Norwell: Kluwer Academic.
- Gat-Viks, I., & Shamir, R. (2003). Chain functions and scoring functions in genetic networks. *Bioinformatics*, 19(1), i108–i117.
- Hill, P. M., & Gallagher, J. (1998). Meta-programming in logic programming. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of logic in artificial intelligence and logic programming* (Vol. 5, pp. 421–497). London: Oxford University Press.
- Inoue, K. (1992). Linear resolution for consequence finding. *Artificial Intelligence*, 56, 301–353.
- Inoue, K. (2004). Induction as consequence finding. *Machine Learning*, 55, 109–135.
- Inoue, K. (2011). Logic programming for Boolean networks. In *Proceedings of IJCAI-11* (pp. 924–930).
- Inoue, K., & Haneda, H. (2000). Learning abductive and nonmonotonic logic programs. In *Abduction and induction—essays on their relation and integration* (pp. 213–231).
- Inoue, K., Bando, H., & Nabeshima, H. (2005). Inducing causal laws by regular inference. In *LNAI: Vol. 3625. Proceedings of ILP'05* (pp. 154–171). Berlin: Springer.
- Inoue, K., Iwanuma, K., & Nabeshima, H. (2006). Consequence finding and computing answers with defaults. *Journal of Intelligent Information Systems*, 26, 41–58.
- Inoue, K., Sato, T., Ishihata, M., Kameya, Y., & Nabeshima, H. (2009). Evaluating abductive hypotheses using an EM algorithm on BDDs. In *Proceedings of IJCAI-09* (pp. 810–815).
- Inoue, K., Furukawa, K., Kobayashi, I., & Nabeshima, H. (2010). Discovering rules by meta-level abduction. In L. De Raedt (Ed.), *LNAI: Vol. 5989. Inductive logic programming: revised papers from the 19th international conference (ILP'09)* (pp. 49–64). Berlin: Springer.
- Inoue, K., Doncescu, A., & Nabeshima, H. (2011). Hypothesizing about networks with positive and negative effects by meta-level abduction. In P. Frasconi & F. A. Lisi (Eds.), *LNAI: Vol. 6489. Inductive logic programming: revised papers from the 20th international conference (ILP'10)* (pp. 114–129). Berlin: Springer.
- Ishii, N., Nakahigashi, K., Baba, T., et al. (2007). Multiple high-throughput analyses monitor the response of *E. coli* to perturbations. *Science*, 316, 593–597.
- Iwanuma, K., & Inoue, K. (2002). Minimal answer computation and SOL. In *LNAI: Vol. 2424. Proceedings of JELIA'02* (pp. 245–257). Berlin: Springer.
- Johnson, D., Yannakakis, M., & Papadimitriou, C. (1988). On generating all maximal independent sets. *Information Processing Letters*, 27(3), 119–123.
- Kakas, A. C., & Riguzzi, F. (1997). Learning with abduction. In *LNAI: Vol. 1297. Proceedings of the 7th international workshop on inductive logic programming* (pp. 181–188). Berlin: Springer.
- Kauffman, S. A. (1993). *The origins of order: self-organization and selection in evolution*. London: Oxford University Press.
- Kimber, T., Broda, K., & Russo, A. (2009). Induction on failure: learning connected Horn theories. In *LNAI: Vol. 5753. Proceedings of LPNMR'09* (pp. 169–181). Berlin: Springer.
- King, R. D., Whelan, K. E., Jones, F. M., Reiser, P. G. K., 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.
- Kowalski, R. A. (1990). Problems and promises of computational logic. In J. W. Lloyd (Ed.), *Computational logic* (pp. 1–36).
- Kowalski, R. A., & Sergot, M. J. (1986). A logic-based calculus of events. *New Generation Computing*, 4(1), 67–95.
- Lamma, E., Mello, P., Milano, M., Riguzzi, F., Esposito, F., Ferilli, S., & Semeraro, G. (2000). Cooperation of abduction and induction in logic programming. In *Abduction and induction—essays on their relation and integration* (pp. 233–252).
- Leitgeb, H. (2001). Nonmonotonic reasoning by inhibition nets. *Artificial Intelligence*, 128, 161–201.
- Lejay, G., Inoue, K., & Doncescu, A. (2011). Application of meta-level abduction for the treatment of hypertension using SOLAR. In *Proceedings of the AINA-2011 workshop on bioinformatics and life science modeling and computing* (pp. 495–500). Los Alamitos: IEEE Comput. Soc.
- Mandal, S., Guptan, P., Owusu-Ansah, E., & Banerjee, U. (2005). Mitochondrial regulation of cell cycle progression during development as revealed by the tenured mutation in *Drosophila*. *Developmental Cell*, 9, 843–854.
- Mandal, S., Freije, W. A., Guptan, P., & Banerjee, U. (2010). Metabolic control of G1–S transition: cyclin E degradation by p53-induced activation of the ubiquitin-proteasome system. *The Journal of Cell Biology*, 188(4), 473–479.

- Meek, D. W. (2009). Tumour suppression by p53: a role for the DNA damage response? *Nature Reviews. Cancer*, 9, 714–723.
- Mooney, R. (2000). Integrating abduction and induction in machine learning. In *Abduction and induction—essays on their relation and integration* (pp. 181–191).
- Moyle, S. (2003). Using theory completion to learn a robot navigation control programs. In *LNAI: Vol. 2583. Proceedings of ILP'02* (pp. 182–197). Berlin: Springer.
- Muggleton, S. (1995). Inverse entailment and Progol. *New Generation Computing*, 13(3/4), 245–286.
- Muggleton, S., & Bryant, C. (2000). Theory completion and inverse entailment. In *LNAI: Vol. 1866. Proceedings of ILP 2000* (pp. 130–146). Berlin: Springer.
- Muggleton, S., & Buntine, W. (1988). Machine invention of first-order predicate by inverting resolution. In *Proceedings of the 5th international workshop on machine learning* (pp. 339–351). San Mateo: Morgan Kaufmann.
- Muggleton, S., Santos, J. C. A., & Tamaddoni-Nezhad, A. (2008). Toplog: ILP using a logic program declarative bias. In *LCNS: Vol. 5366. Proceedings of ICLP'08* (pp. 687–692). Berlin: Springer.
- Muggleton, S. H., De Raedt, L., Poole, D., Bratko, I., Flach, P., Inoue, K., & Srinivasan, A. (2012). ILP turns 20—Biography and future challenges. *Machine Learning*, 86(1), 3–23.
- Nabeshima, H., Iwanuma, K., & Inoue, K. (2003). SOLAR: a consequence finding system for advanced reasoning. In *LNAI: Vol. 2796. Proceedings of TABLEAUX'03* (pp. 257–263). Berlin: Springer.
- Nabeshima, H., Iwanuma, K., Inoue, K., & Ray, O. (2010). SOLAR: An automated deduction system for consequence finding. *AI Communications*, 23(2–3), 183–203.
- Nakahigashi, K., Toya, Y., Ishii, N., Soga, T., Hasegawa, M., Watanabe, H., Takai, Y., Honma, M., Mori, H., & Tomita, M. (2009). Systematic phenome analysis of Escherichia coli multiple-knockout mutants reveals hidden reactions in central carbon metabolism. *Molecular Systems Biology*, 5, 306.
- Otero, R. P. (2005). Induction of the indirect effects of actions by monotonic methods. In *LNAI: Vol. 3625. Proceedings of ILP'05* (pp. 279–294). Berlin: Springer.
- Pearl, J. (2009). *Causality: models, reasoning, and inference* (2nd ed.). Cambridge: Cambridge University Press.
- Poole, D. (1988). A logical framework for default reasoning. *Artificial Intelligence*, 36, 27–47.
- Prives, C., & Hall, P. A. (1999). The p53 pathway. *Journal of Pathology*, 187, 112–126.
- Ray, O. (2009). Towards a rational approach for the logical modelling of inhibition in metabolic networks. In *Proceedings of the 2009 IEEE international workshop on bioinformatics and life science modeling and computing* (pp. 18–923).
- Ray, O., & Inoue, K. (2007). A consequence finding approach for full clausal abduction. In *LNAI: Vol. 4755. Proceedings of the 10th international conference on discovery science* (pp. 173–184). Berlin: Springer.
- Ray, O., & Inoue, K. (2008). Mode-directed inverse entailment for full clausal theories. In *LNAI: Vol. 4894. Proceedings of ILP'07* (pp. 225–238). Berlin: Springer.
- Ray, O., Broda, K., & Russo, A. (2003). Hybrid abductive inductive learning: a generalisation of Progol. In *LNAI: Vol. 2835. Proceedings of ILP'03* (pp. 311–328). Berlin: Springer.
- Ray, O., Whelan, K., & King, R. (2010). Logic-based steady-state analysis and revision of metabolic networks with inhibition. In *Proceedings of the 3rd international workshop on intelligent informatics in biology and medicine* (pp. 661–666).
- Reiser, P. G. K., King, R. D., Kell, D. B., Muggleton, S. H., Bryant, C. H., & Oliver, S. G. (2001). Developing a logical model of yeast metabolism. *Electronic Transactions in Artificial Intelligence*, 5-B2(024), 223–244.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–132.
- Sakama, C. (2000). Abductive generalization and specialization. In *Abduction and induction—essays on their relation and integration* (pp. 253–265).
- Schaub, T., & Thiele, S. (2009). Metabolic network expansion with answer set programming. In *LNAI: Vol. 5649. Proceedings of ICLP'09* (pp. 312–326). Berlin: Springer.
- Schurz, G. (2008). Patterns of abduction. *Synthese*, 164(2), 201–234.
- Shanahan, M. (2000). An abductive event calculus planner. *The Journal of Logic Programming*, 44, 207–239.
- Shmulevich, I., Dougherty, E. R., Kim, S., & Zhang, W. (2002). Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics*, 18(2), 261–274.
- Tamaddoni-Nezhad, A., Chaleil, R., Kakas, A., & Muggleton, S. (2006). Application of abductive ILP to learning metabolic network inhibition from temporal data. *Machine Learning*, 65, 209–230.
- T'kindt, V., Bouibede-Hocine, K., & Esswein, C. (2007). Counting and enumeration complexity with application to multicriteria scheduling. *Annals of Operations Research*, 153, 215–234.
- Tran, N., & Baral, C. (2009). Hypothesizing about signaling networks. *Journal of Applied Logic*, 7(3), 253–274.
- Wagner, A. (2005). *Robustness and evolvability in living systems*. Princeton: Princeton University Press.

- Yamamoto, A. (2000). Using abduction for induction based on bottom generalization. In *Abduction and induction—essays on their relation and integration* (pp. 267–280).
- Yamamoto, Y., Inoue, K., & Doncescu, A. (2010). Integrating abduction and induction in biological inference using CF-induction. In H. Lodhi & S. Muggleton (Eds.), *Elements of computational systems biology* (pp. 213–234). New York: Wiley.
- Zupan, B., Demsar, J., Bratko, I., 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.