

Cost-sensitive Iterative Abductive Reasoning with Abstractions

Gianluca Torta¹ and Daniele Theseider Dupré² and Luca Anselma¹

1 Introduction

Several explanation and interpretation tasks, such as diagnosis, plan recognition and image interpretation, can be formalized as abductive reasoning. A number of approaches, including recent ones [1, 4], address the problem based on a task-independent representation of a domain which includes an ontology or taxonomy of hypotheses.

In this paper we adopt a similar representation, but we also deal with abduction as an iterative process where, like in model-based diagnosis, further observations are proposed to discriminate among candidate explanations; in addition, we take into account costs of observations and actions. In fact, discrimination also involves refining hypotheses, but this is performed down to an appropriate level which depends on the cost of actions (e.g. repair actions or therapy) to be taken based on the results of abduction, and on the cost of additional observations, which should be balanced with the benefits, in terms of more suitable actions, of better discrimination.

The presence of a domain representation with abstractions has a significant impact on this trade-off. In general, a better assessment of the situation at hand, based on additional observations, leads to a more focused action. However, the cost of observing the same phenomenon at different levels of abstraction may vary significantly; in fact, it could involve more or less costly medical or technical tests, or computationally complex image processing, possibly with additional costs due to the delay before taking an action.

Moreover, the knowledge base could have been designed independently of the explanation/action task (e.g. diagnosis and repair), and could therefore include a detailed description of the domain which is not necessary for the task; more generally, the convenience of a detailed discrimination may depend on the specific case at hand.

By explicitly considering abstractions in the iterative abduction process, we can often reduce the observation costs significantly, yet maintaining the ability to exploit detailed observations and knowledge when convenient (similar advantages have been shown in inductive classification with abstractions, e.g. [6]).

In the following, we first describe the knowledge we expect to be available. We then describe a basic iterative abduction loop and, finally, we concentrate on the criterion for selecting the next step in the loop: either performing a next observation at some level of detail, or stopping because the estimated most convenient choice is performing the action(s) associated with the current hypotheses.

2 Domain Representation

The basic elements of the domain model are a set of abducibles (atomic assumptions) $\mathcal{A} = \{A_1, \dots, A_n\}$ and a set of manifesta-

tions $\mathcal{M} = \{M_1, \dots, M_m\}$. Each abducible A_i is associated with an IS-A hierarchy $\Lambda(A_i)$ containing abstract values of A_i as well as their refinements at multiple levels; similarly, each manifestation M_j is associated with an IS-A hierarchy $\Lambda(M_j)$. We assume that the direct refinements v_1, \dots, v_q of a value V in a hierarchy (either $\Lambda(A_i)$ or $\Lambda(M_j)$) are mutually exclusive, and at most one of the leaf values in a hierarchy is true in each situation, i.e. we allow at most one instance for each abducible and observation; moreover, for each leaf value v of an abducible an a-priori probability $p(v)$ is given.

The hypotheses space $\mathcal{S}(\mathcal{A})$ for the abduction task is the set of all of the combinations γ of values drawn from one or more distinct hierarchies $\Lambda(A_i)$, while the manifestations space $\mathcal{S}(\mathcal{M})$ is the set of all of the combinations ω of values drawn from distinct hierarchies $\Lambda(M_j)$. The relationships between the abducibles and the manifestations are defined by the domain knowledge $\mathcal{K} \subseteq \mathcal{S}(\mathcal{A}) \times \mathcal{S}(\mathcal{M})$.

Given an instance of manifestations $\omega \in \mathcal{S}(\mathcal{M})$ and an instance of abducibles $\gamma \in \mathcal{S}(\mathcal{A})$, $(\gamma, \omega) \in \mathcal{K}$ means that ω is a possible observation set corresponding to the hypothesis set γ .

We associate costs with values of both abducibles and manifestations. Let $C \in \Lambda(A_i)$ be a value belonging to the IS-A hierarchy of A_i ; its cost $ac(C)$ is the cost of the action to be taken when A_i takes value C (e.g. a repair action if A_i represents a component and C denotes one of its fault modes). Let c_1, \dots, c_q be the children of C in $\Lambda(A_i)$, i.e. the possible refinements of value C . We assume that:

$$\max(\{ac(c_1), \dots, ac(c_q)\}) \leq ac(C) \leq \sum_{k=1}^q ac(c_k)$$

i.e. the action that we take for a value C of A_i costs no less than the most expensive action for its refinements and no more than taking the actions for all of such refinements. As for the manifestations, let $O \in \Lambda(M_j)$ be a value belonging to the IS-A hierarchy of M_j ; its cost $oc(O)$ is the cost of making the observation which refines value O into one of its children o_1, \dots, o_q in $\Lambda(O_j)$.

We can associate an action cost also with any instance $\gamma = \{C_1, \dots, C_r\} \in \mathcal{S}(\mathcal{A})$ of abducibles simply as $ac(\gamma) = \sum_{i=1}^r ac(C_i)$, i.e. we assume that independent actions are taken for each of the abducibles values that appear in γ . With a slightly more complex computation we can also associate an action cost with a set of instances $\Gamma = \{\gamma_1, \dots, \gamma_s\}$ representing the cumulative action cost if Γ is the final set of explanations. For each abducible A_i s.t. (a value of) A_i appears at least in one $\gamma \in \Gamma$, we compute a new hierarchy $\Lambda(A_i, \Gamma)$ by considering the portion of $\Lambda(A_i)$ up to the least upper bound $LUB(A_i, \Gamma)$ that covers all of the values of A_i that appear in Γ and by further removing from such a sub-tree all of the values that do not appear in Γ .

In this way, it may happen that the cost $ac(C)$ of a value $C \in \Lambda(A_i, \Gamma)$ is larger than the sum of the costs $ac(c_k)$ of its children,

¹ Università di Torino, Italy, email: {torta,anselma}@di.unito.it

² Università del Piemonte Orientale, Italy, email: dtd@mfn.unipmn.it

since not all of the children of C defined in $\Lambda(A_i)$ need to appear in $\Lambda(A_i, \Gamma)$. We therefore update (with a bottom-up computation) the ac costs in $\Lambda(A_i, \Gamma)$ to new costs ac^* in order to reestablish this property. The action cost of Γ is then computed just as:

$$ac(\Gamma) = \sum_{A_i \in \Gamma} ac^*(LUB(A_i, \Gamma))$$

3 Iterative Abduction

We rely on the following generic loop for iterative explanation:

Input is a set of values $\omega_I = \{O_1, \dots, O_t\}$ representing the initial observations, i.e. the values of a set of manifestations $\{M^1, \dots, M^t\} \subseteq \mathcal{M}$.

Generate a set Γ of candidates (i.e. explanations of ω_I).

```

loop
   $O := \text{NextStep}(\Gamma)$ ;
  if  $O = \text{STOP}$  then exit
  else
    perform observation to refine  $O$  into one of its children  $o_k$ ;
     $\Gamma := \text{Update}(\Gamma, o_k)$ 
end

```

end

That is, we assume that one or more initial observations are given; that there is a way to generate candidate explanations based on them (see below), and to update candidates based on additional observations; and we proceed with selecting and performing one observation at a time, which, of course, is in general suboptimal, as in [3, 2].

In this paper we aim at providing a general approach to the selection of the next step; we do not provide a general approach to candidate generation and update which could involve a mix of abduction and consistency reasoning; its formulation would depend on the way \mathcal{K} is represented. With hierarchies of abducibles, moreover, abstract as well as detailed assumptions may take part in explanations; a general criterion which is suitable in this setting is the preference for *least presumptive* explanations [5], which generalize minimal (wrt set inclusion) explanations: an explanation that (also based on the IS-A hierarchy) implies another explanation is not least presumptive. In the following we assume that the candidates computed at each iteration represent the least presumptive explanations of the observations collected so far.

4 Choosing the Next Step

Let Γ be the current candidate set and let OBS be the set of possible observations (including refinements of previous observations). In order to decide whether to stop or to proceed with a new observation $O \in OBS$, we select the minimum among:

- the action cost $ac(\Gamma)$ associated with Γ
- for each $O \in OBS$, the estimated cost $c(O)$, which is the sum of the cost $oc(O)$ of observing O and the expected cost of the candidate set after observing O , i.e.:

$$c(O) = oc(O) + \sum_{k=1}^q p(o_k) \cdot c(\Gamma_k)$$

where $\Gamma_1, \dots, \Gamma_q$ are the possible candidate sets that would result by observing O and getting values o_1, \dots, o_q respectively, $p(o_k)$ is the probability of getting value o_k (computed based on current candidates Γ as in [3, 2]) and $c(\Gamma_k)$ is the estimated cost of Γ_k as detailed in the following.

If $ac(\Gamma)$ is the minimum among the costs, we stop; otherwise we observe the O with the smallest $c(O)$.

Let $\Gamma_k = \{\gamma_1, \dots, \gamma_s\}$ be one of the candidate sets involved in the above formula (note that each candidate γ_i may contain ground as well as abstract causes) and $ac(\Gamma_k)$ be its action cost, i.e. the cost of stopping at Γ_k , which must be compared with the estimated cost of acting after a further discrimination and refinement.

In principle, this estimation step would require to simulate all the possible observation sequences and outcomes and, for each of them, to assess the point where it is convenient, on average, to stop and perform the actions; in order to avoid such an intractable search, we assume that the abductive process will continue as follows: first, one of the $\gamma_i \in \Gamma_k$ will be isolated; then, γ_i is refined level by level, up to a point where performing an action is estimated to be convenient. Therefore the estimated cost of Γ_k is:

$$c(\Gamma_k) = \min(ac(\Gamma_k), ic(\Gamma_k) + rac(\Gamma_k))$$

where $ic(\Gamma_k)$ is the estimated cost of isolating a single $\gamma_i \in \Gamma_k$ and $rac(\Gamma_k)$ is the estimated additional refinement and action cost once some γ_i has been isolated.

In this proposal, we estimate the cost $ic(\Gamma_k)$ as follows:

$$ic(\Gamma_k) = \sum_{i=1}^s -p(\gamma_i) \cdot \log(p(\gamma_i)) \cdot \overline{oc}(\gamma_i)$$

where $-\log(p(\gamma_i))$ is the estimated number of observations needed for isolating γ_i [3] and $\overline{oc}(\gamma_i)$ is an estimate of the cost of a single observation³. The cost $rac(\Gamma_k)$ of refining its members $\gamma_i = \{C_{i,1}, \dots, C_{i,r_i}\}$ until an action is taken is estimated by:

$$rac(\Gamma_k) = \sum_{i=1}^s \left(p(\gamma_i) \cdot \sum_{j=1}^{r_i} c(C_{i,j}) \right)$$

where $c(C_{i,j})$ is the estimated cost associated with $C_{i,j}$.

In case action costs do not depend on the current context, each cost $c(C_{i,j})$ can be pre-computed offline with a bottom-up visit of the taxonomies of the causes. In this proposal we have adopted a formula similar to the one for $c(\Gamma_k)$, i.e.:

$$c(C_{i,j}) = \min(ac(C_{i,j}), ic(C_{i,j}) + rac(C_{i,j}))$$

where $ic(C_{i,j})$ is the estimated cost of isolating a single child of $C_{i,j}$ in the hierarchy and $rac(C_{i,j})$ is the estimated additional refinement and action cost once some child of $C_{i,j}$ has been isolated⁴.

REFERENCES

- [1] Ph. Besnard, M.-O. Cordier, and Y. Moinard, ‘Ontology-based inference for causal explanation’, in *Knowledge Science, Engineering and Management, 2nd Int. Conf., LNCS 4798*, pp. 153–164, (2007).
- [2] L. Console, D. Theseider Dupré, and P. Torasso, ‘Introducing test theory into abductive diagnosis’, in *Proc. 10th Int. Work. on Expert Systems and Their Applications*, pp. 111–124, Avignon, (1990).
- [3] J. de Kleer and B.C. Williams, ‘Diagnosing multiple faults’, *Artificial Intelligence*, **32**(1), 97–130, (1987).
- [4] B. Neumann and R. Möller, ‘On scene interpretation with description logics’, in *Cognitive Vision Systems*, 247–275, Springer, (2006).
- [5] D. Poole, ‘Explanation, prediction: an architecture for default, abductive reasoning’, *Computational Intelligence*, **5**, 97–110, (1989).
- [6] J. Zhang, A. Silvescu, and V. Honavar, ‘Ontology-driven induction of decision trees at multiple levels of abstraction’, *LNCS*, **2371**, 316–323, (2002).

³ We have defined \overline{oc} as a function of γ_i to possibly take into account the level of detail of observations related with γ_i

⁴ Note that when $C_{i,j}$ is a leaf of the hierarchy, $c(C_{i,j})$ is the action cost $ac(C_{i,j})$.