Abduction and Question Answering *

Debra T. Burhans and Stuart C. Shapiro
Department of Computer Science and Engineering
University at Buffalo, The State University of New York

1 Introduction

Abduction can be viewed as self-questioning, where the search for an explanation is analogous to answering a “why” question. The question is of the form *Why P?*, where *P* is an observation or fact that has come to an agent’s attention. An abductive hypothesis is a possible answer to the question *Why P?* When abduction is viewed as a type of question answering, abductive hypotheses can be seen as a subset of *hypothetical* (or conditional) answers. These answers are comprised of two components: a hypothesis that is consistent with an agent’s knowledge, but whose validity can not be determined, and an associated specific or generic answer, whose correctness hinges on the validity of the associated hypothesis. Hypothetical answering and abductive reasoning are both ways of coping with incomplete information. The difference between the two lies in the way in which the hypothetical information is employed. In abduction, a reasoner chooses to believe an abductive hypothesis based on factors such as belonging to a set of abducibles and informativeness. In question answering (excepting the case of self-questioning), the answering agent is reasoning in service of a distinct questioner, and has no reason to adopt or consider the hypothesis associated with a hypothetical answer as an update to its knowledge base.

It is in the interest of an abductive reasoner to place higher value an abductive hypothesis that is more informative than other abductive hypotheses. This is true for the *questioner* also. The difference between abduction and question answering is that the answering agent does not customarily have access to the questioner’s knowledge base. This means that the informativeness of a hypothetical answer can not be determined by an answering agent, making it necessary to provide hypothetical answers at all possible levels of generality. It is generally not possible for the answering agent to measure the “informativeness” of an abductive hypothesis. This paper presents a discussion of the relationship between question answering and abduction.

2 Background

For the purposes of this discussion, question answering is assumed to take place in a resolution theorem prover. This is not a limitation on the ideas presented herein, rather, it simplifies some of the discussion. Further, the use of an answer literal in reasoning is assumed [Green, 1969b; Green, 1969a]. While the focus of the paper is on the the relationship between question answering and abduction, it is necessary to describe the answer classes specific, generic, and hypothetical in order to frame the problem. Specific and generic answers are not directly related to abduction, it is the class of hypothetical answers that is of particular interest.

2.1 Specific Answers

A specific or *extensional* answer is a witness that proves the truth of an existential hypothesis: it is a ground answer to a question. Specific answers are associated with proofs in theorem provers, which have often been employed as question answering frameworks [Green and Raphael, 1968; Green, 1969b; Green, 1969a]. A specific answer is associated with a clause containing only answer literals all of which are ground.

2.2 Generic Answers

A generic or *intensional* answer is generally associated with a rule that expresses general properties of classes or groups of objects that satisfactorily answer a question. Generic answers have been examined by researchers working in rule bases [Cholvy and Demolombe, 1986; Cholvy, 1990] as well as databases [Reiter, 1978a; Reiter, 1986; Imielinski, 1987; Motro, 1989] and cooperative answering [Lehnert, 1978; Kaplan, 1981; Webber, 1986]. Our definition of generic answers is based on variable sharing and requires the following definitions:

Variable Sharing Set of a Literal. The variable sharing set of a literal *l* with respect to a set of literals $\mathcal{C}$, is a set of literals $\mathcal{VS}(l)$ such that each literal in $\mathcal{VS}(l)$ is an element of $\mathcal{C}$, and shares at least one variable with *l*.

Closure of Variable Sharing of a Literal. The closure of variable sharing of a literal *l*, denoted $\mathcal{CVS}(l)$, is the fixpoint of the $\mathcal{VS}$ operator: $\mathcal{VS}(\mathcal{CVS}(l)) = \mathcal{CVS}(l)$. It is inductively defined as follows:

$\mathcal{CVS}_0 = \mathcal{VS}(l)$
$\mathcal{CVS}_1 = \mathcal{VS}(\mathcal{VS}(l))$

*To appear in *Proceedings of the IJCAI Workshop on Abductive Reasoning*, IJCAI-01, Seattle, Washington, August 2001

$\vdots$
$\mathcal{CVS}_n = \mathcal{VS}(\mathcal{VS}\ldots(\mathcal{VS}(l))\ldots)$
$\mathcal{CVS}_{n+1} = \mathcal{CVS}_n$
where $\mathcal{CVS}_n$ is the fixpoint. The maximum number of literals in this set is the number of literals in the clause $\mathcal{C}$, with respect to which $\mathcal{CVS}(l)$ is computed.

Of particular interest in defining answer classes is the closure of variable sharing of the set of *answer literals*, which is the union of the closure of variable sharing of each of the answer literals. Consider the following clause:$\{R(x,z),\ P(y),\ Q(z), ANSWER(G(x))\}$. Given $l = ANSWER(G(x))$, and $\mathcal{C} = \{R(x,z),\ P(y),\ Q(z)\}$, the variable sharing of l is $\mathcal{VS}(ANSWER(G(x))) = \{R(x,z)\}$, and the closure of variable sharing of l is $\mathcal{CVS}(ANSWER(G(x))) = \{R(x,z),\ Q(z)\}$. This shows how the literals in a clause can be partitioned into three sets: answer literals, non-answer literals in the closure of variable sharing of the answer literals, and other non-answer literals.

A generic answer is associated with a clause that has the following properties: either the clause contains non-answer literals in addition to answer literals that are all in the closure of variable sharing of the answer literals, or there are only answer literals but they contain at least one variable (which we term *unconstrained generic answers*). The clause can be rewritten in rule form where the antecedent is the conjunction of the (negations of the) non-answer literals and the consequent is the disjunction of the question predicates, which are the arguments of the answer literals (with full predicate structure). Variables in the antecedent are universally quantified with wide scope over the entire rule, while non-shared variables in the consequent are locally, universally quantified.

2.3 Hypothetical Answers

Hypothetical, or *conditional* answers, are a less widely accepted class of answers in comparison with specific and generic answers. Hypothetical answers have received some attention [Wolstenholme, 1988; Burhans and Shapiro, 1999; Demolombe, 1992; Demolombe, 1997], in part due to their utility when reasoning takes place in the absence of a Closed World Assumption [Reiter, 1978b], as in abductive reasoning [Kakas *et al.*, 1998].

We identify hypothetical answers with clauses containing at least one non-answer literal that is *not* in the closure of variable sharing of the answer literals. Written in rule form, a hypothetical answer consists of two parts: one part is either a specific or generic answer, and the other is a condition that must hold in order for the accompanying answer to *be considered an answer*. The condition is termed the *hypothetical component* of the answer, and the rest of the answer is termed the *answer component*. Hypothetical answers are partitioned into specific and generic, based on the type of the answer component. The hallmark of the hypothetical answer is the way in which information belonging to the hypothetical component is identified, namely, literals *not* in the closure of variable sharing of the answer literals.

There are two important issues regarding this definition and previous work. First, hypothetical answers, and, in general, abductive hypotheses, have been identified with *all ground* rules. Some work on expanding the notion of abduction in order to support reasoning from "general observations" to general hypotheses has been done [Michalski, 1993]. The proposed definition imposes no limitations on the form of clauses other than the variable sharing property defined above. Second, the definitions provided for generic and hypothetical answers completely clarify the way in which clauses corresponding to rules should be classified: *rule* and *generic answer* are not the same thing.

3 Examples

A simple example from Rich and Knight [1991, page 192] is used as a first illustration of hypothetical answering. Consider the following knowledge base (KB):

1. Cats like to eat fish
2. Cats like to eat tuna
3. Calicos are cats
4. Herb is a tuna

and the question *What does Boots like to eat?* What sorts of answers arise according to the definitions given in the previous section? Note that the individual *Boots* is not mentioned in the KB: nothing whatsoever is known about Boots. Under ordinary circumstances a query about an unknown individual would fail. However, the following list of statements correspond to clauses that are generated as an answer search proceeds in a typical resolution refutation theorem prover:

1. If Boots is a calico then Boots likes to eat Herb.
2. If Boots is a cat then Boots likes to eat Herb.
3. If Boots is a calico then Boots likes to eat tuna.
4. If Boots is a calico then Boots likes to eat fish.
5. If Boots is a cat then Boots likes to eat tuna.
6. If Boots is a cat then Boots likes to eat fish.

Answers 1 and 2 are *specific hypothetical answers* because the associated answer component is specific (all ground). Similarly answers 3-6 are *generic hypothetical answers* because the associated answer component is generic. There are only two distinct hypothetical components across the six answers: *Boots is a cat* and *Boots is a calico*. Because all calicos are cats, these hypotheses are related taxonomically: if Boots is a calico then, based on the information in the KB, Boots is a cat. The important point here is that this taxonomic relation may not exist in the mind of the questioner: a questioner might know that Boots is a calico and nothing about the relationship between calicos and cats. In this case (which is just one example of this phenomenon) it would clearly be detrimental (as well as arbitrary) to deny "answer" status to an answer with hypothetical component *Boots is a calico* simply because it is taxonomically subsumed by *Boots is a cat*. It is, however, appropriate to disregard *syntactically subsumed* clauses when selecting from among a set of answers. This contrasts with the situation in abduction, where taxonomic subsumption is important: it reflects the generality, thus the informativeness, of a particular abductive hypothesis. The taxonomy is part of the reasoner's knowledge, whereas the taxonomy of the questioner is not necessarily part of the answering agent's knowledge.

When taxonomically related answers are given, it is incumbent upon the answering agent to convey not only the answers but their taxonomic relationship. To enable this we define an *answer set* that embodies taxonomic information. Note that this information is *not* requested as part of a question, yet is an essential component of a cooperative answer.

Not all hypothetical components of answers are ground. The following example shows a universally quantified hypothetical component with an unconstrained generic answer as a consequent. Given a knowledge base with one rule:

1. ∀ x ((dog x) ⇒ (barks x))

and the question *Does everything bark?*, the answer generated by our theorem prover is, in clause form, {(¬(DOG S11)), (ANSWER (BARKS S11))}. The corresponding answer in rule form is (DOG S11) ⇒ (BARKS S11). As shown by Luckham and Nilsson [1971], Skolem constants and functions can be replaced by universally quantified variables in clauses *outside* the context of an ongoing proof. When should this replacement take place: should it be *before* or *after* the answer literals are divided into the hypothetical, and generic or specific answer component of a hypothetical answer? If it takes place beforehand, it seems to indicate that S11 in both literals should be replaced with the same variable, thus yielding a generic answer. This would conflict with our definition of the closure of variable sharing, which stipulates that the literal (DOG S11), an all-ground literal, which does not share variables with the answer literal (ANSWER (BARKS S11)), should appear in the hypothetical position of the answer. A consideration of the original question, and the semantics of both the question and answer, leads to clarification on this issue. The original question asked whether *everything* barks. If S11 is replaced by the same variable in both clauses, thus becoming a shared, universally quantified variable in a generic answer, the answer would be $\forall\ x\ [(dog\ x)\ \Rightarrow\ (barks\ x)]$. The meaning of this rule is that if something is a dog then it barks, it makes no claim about whether or not the arbitrary individual barks. On the other hand, by first separating the literals into the two components of the answer, and replacing S11 by a universally quantified variable with *local scope* in the two answer components, the answer would be $\forall\ x\ (dog\ x)\ \Rightarrow\ \forall\ y\ (barks\ y)$. The meaning of this rule is that if everything is a dog then everything barks. This is clearly the right answer. Thus, not only must translation from Skolem functions and constants to variables take place outside the scope of reasoning, but it must be delayed until the literals in a clause have been properly separated into the hypothetical and answer components.

The hypothetical component of an answer describes conditions on the associated answer, whether that answer is generic or specific. In contrast to generic answers, which describe *conditions on the individuals* that satisfy the question, hypotheticals described *conditions on the answer itself*. These conditions are analogous to abductive hypotheses.

4 Hypothetical Answers and Abduction

As the foregoing discussion demonstrates, there are a number of similarities between hypothetical answering and abduction. There are also important differences. The observations or facts comprising the input to the abductive reasoning process are *true*. The status of the propositional content of a question with regard to truth or falsity is, initially, *unknown*. Whether or not the question has a "yes" answer, indicating the truth of the proposition put by the question, the answering agent may not have an interest in adding the proposition to its knowledge base. The existence of multiple explanations in abduction is a problem: often these explanations are *competing* in the sense that if one is true the others are not. The existence of multiple hypothetical answers is *not* a problem. All such answers are implied by the knowledge base, and their status is not provisional. The answering agent does not have the necessary information to discriminate between hypothetical answers: that information is privy to the questioner. Table 1 shows a comparison between abductive reasoning and the generation of hypothetical answers to a question.

As shown in Table 1, abduction begins with observable facts and reasons towards provisional hypotheses. The observable facts are known to be true and are added to the knowledge base of the abductive reasoner. The reasoner then searches for possible explanations for the observations. In contrast, question answering begins with a proposition that is neither known to be true nor known to be false to a reasoner. The goal of reasoning is to determine whether the proposition is true. In the absence of the ability to do that, the reasoner may produce *hypothetical answers*.

At the end of the reasoning process, abduction involves accepting as true hypotheses that are neither provably true nor provably false. At the end of the reasoning process performed by a question answerer, the hypothetical component of an answer to a question is neither assumed to be true or false by an answering agent. Rather, the hypothetical answer is added to a set of answers from which an answering agent may draw in order to generate responses. The answering agent may *choose* to offer the hypothetical component of an answer, which *formally* corresponds to an abductive hypothesis, back to the questioner as a follow-up question, but clearly the purpose and effect is different.

Bromberger [1992] provides a clear explication of what constitutes an *explanation*: abductive reasoning is the search for explanations. The need for explanations arises in the face of a *p-predicament*. A *p-predicament* is a situation in which nothing an intelligent agent knows accounts for the truth or lack thereof of a proposition *p*. The search for an explanation that alleviates a *p-predicament* is often initiated by asking a "why-question". Not all why-questions are requests for explanations, but when they are Bromberger identifies an "inner question" that is presumed true. For example, the why-question *Why is the sky blue?* has as its inner question *Is the sky blue?* Inner questions in this context function as presuppositions.

The analogy between Bromberger's observations and our discussion of abduction and question answering is clear. In response to the (observation/fact) *the sky is blue*, an abduc-

Table 1: A Comparison between Abduction and Hypothetical Answers

	Abduction	*Hypothetical Q/A*
Start State of Reasoning Agent KB	$\mathcal{KB}$	$\mathcal{KB}$
Information "Trigger" for Reasoner Example:	Fact/Observation The sky is blue	Question Is the sky blue?
Representation of Information	$\mathcal{P}$	$\mathcal{P}$ (propositional content of question)
KB of Reasoning Agent	$\mathcal{KB}' = \mathcal{KB} \cup \mathcal{P}$	$\mathcal{KB}$
Reasoning Proceeds With:	$\mathcal{KB}'$ (or $\mathcal{KB} \cup \neg\mathcal{P}$ depending on how reasoner operates)	$\mathcal{KB} \cup \neg\mathcal{P}$ (refutation proof)
Purpose of Reasoning	Find an explanation for $\mathcal{P}$	Determined the validity of $\mathcal{P}$
Conclusion of Reasoning Example:	$\mathcal{KB} \vdash \mathcal{H} \rightarrow \mathcal{P}$ If it is not raining then the sky is blue	$\mathcal{KB} \vdash \mathcal{H} \rightarrow \mathcal{P}$ If it is not raining then the sky is blue
Interpretation of Results	$\mathcal{H}$ is an abductive hypothesis (assume that $\mathcal{H}$ is in the set of abducibles)	$\mathcal{H} \rightarrow \mathcal{P}$ is a hypothetical answer $\mathcal{H}$ is a new question that is unanswerable by this reasoner
Action following Reasoning	Adopt $\mathcal{H}$ as a defeasible hypothesis	Return $\mathcal{H} \rightarrow \mathcal{P}$ as a hypothetical answer
KB of Reasoning Agent Action:	$\mathcal{KB}'' = \mathcal{KB}' \cup \mathcal{H}$ Assume: It is not raining	$\mathcal{KB}$ New answer: If it is not raining then the sky is blue

tive reasoner effectively asks itself the question *Why is the sky blue?* In response to a question whose propositional content is *the sky is blue*, a question answerer attempts to determine the validity of the proposition. An abductive reasoner is designed to make *assumptions* that are not entailed by its knowledge, while a question answerer is expected to respond with information that is true according to its knowledge.

Abductive reasoning provides a way to handle the failure to prove a proposition $\mathcal{P}$. Each failed branch of the proof tree in a resolution theorem prover corresponds to a non-empty set of clauses. The clauses on such a branch do entail $\mathcal{P}$, the problem is that such clauses are neither known to be true, nor are they known to be false. Above it was mentioned that this is precisely the status of the propositional content of a question. A questioner will take a proposition whose truth is unknown and *pose* it as a question to another agent, whereas an abductive reasoner will simply make an assumption about its truth. In the context of question answering, an answering agent may in turn derive a hypothetical answer with its associated "question" (hypothetical component), namely, a proposition whose truth is unknown. In this context the hypothetical component of an answer makes an ideal follow-up question that the answering agent may (if circumstances permit) pose to the questioner. Alternatively, in response to a hypothetical answer a *questioner* can try to get yet another answering agent to verify the hypothetical component of the answer. Like abduction, question answering provides a way to handle the failure to prove a proposition $\mathcal{P}$.

While abductive hypotheses are generally associated with failed branches of a proof tree [Cox and Pietrzykowski, 1987], other criteria are also used to narrow the selection of abductive hypotheses:

- The hypothesis should convey a *cause*, it should not explain an effect in terms of another effect [Kakas *et al.*, 1998]. Cox and Pietrzykowski refer to this property as *basic* [1987].
- The hypothesis should not be subsumed by any other hypotheses.

In order to enforce these requirements and to limit the number of possible abductive hypotheses a system may infer, abductive reasoning systems incorporate a preferred, predefined set of "abducibles": predicates that generally do not appear as the consequence of any rules. In order for a clause to be considered as a possible abductive hypothesis it must match one of these predefined predicates. Abducibles are "needed to distinguish abductive explanations from inductive generalizations" [Kakas *et al.*, 1998, page 237].

Note that *generic answers* are generalizations, whereas hypothetical answers are explanations. The desire for obtaining a "minimal", non-redundant, basic explanation for an observation determines what a system will consider to be an abducible hypothesis. The question answering task involves no such requirement: in fact, it is important in question answering to be able to provide answers at *many levels of specificity*, and generalizations are important types of answers. This is because the questioner's knowledge may not include the same concepts as those of the answering agent, particularly in terms of specificity/generality of concepts. Thus, no *a priori* assumptions about what will comprise a satisfactory answer to questioner may be made in the absence of a large amount of knowledge about the questioner's knowledge which is generally impractical. For question answering, not only the clauses at the *ends* of the failed branches of a proof tree are interesting, but all intermediate clauses along the branches comprise answers, excepting the root, which is simply the negation of the original question. Our scheme for distinguishing generic and hypothetical answers cleanly shows which of these clauses are associated with potential abductive hypotheses.

4.1 Circumstances under which Hypothetical Answers Arise

In a back-chaining reasoner, hypothetical answers are discovered "along the way" as the search for specific and generic answers proceeds. Most hypothetical answers function merely as resolution intermediates and are ultimately subsumed: when a hypothetical answer is subsumed by a specific or generic answer, its associated answer component becomes an answer. When a hypothetical answer is *not* subsumed, however, it stands on its own as both relevant and informative. It is relevant because it is connected to the question by a chain of reasoning. It is informative in that it provides the questioner with new information, namely, a question for which the reasoner has no answer. It is less informative than specific and generic answers in that it neither gives a satisfying instance for a question nor a general description of a satisfying instance, but it does provide information about what needs to be known in order for an answer to be forthcoming.

In a previously shown example, it was seen that *all* answers are hypothetical when there is a complete lack of information about individuals mentioned in a question. Hypotheses about properties of such individuals must be made in order to provide answers. The only situation under which non-hypothetical answers can be found regarding a completely unknown individual is in the presence of *unconstrained generic information*. For example, if a knowledge based contains a universally quantified formula such as $\forall\, x\ (P\ x)$, then $(P\ a)$ can be inferred for any individual a.

Hypothetical answers will be present when there is *disjunctive* information in a knowledge base. A disjunction, where neither disjunct is known to be true or known to be false, represents a lack of information. If $\mathcal{A} \vee \mathcal{B}$ is known, and neither $\mathcal{A}$ nor $\mathcal{B}$ is known to be true, then hypothetical answers will arise. Such circumstances also lead to the discovery of abductive hypotheses in an abductive reasoner.

4.2 The Hypothetical Component of an Answer as a Question

An important characteristic of the hypothetical answer is that the hypothetical component is itself a *a good question*. If the question corresponding to the hypothetical component can be answered, the associated *answer component* of the hypothetical answer is immediately recognized as either a specific or generic answer. The fact that the hypothetical component is a "detachable" question arises from the way we have defined the form of hypothetical answers: because there is no variable sharing between the hypothetical and answer components of the answer, all quantifiers in the hypothetical component may be *locally scoped*. This means the hypothetical component can be effectively detached from the rest of the answer. Similarly, an abductive hypothesis can serve as a question to ask another agent. In the absence of a verification, the abductive reasoner may still choose to believe the hypothesis based on other information it has to support that belief.

4.3 Reasoning in Abduction and Question Answering

It is clearly not always practical, in the face of lack of information, to generate *all* possible hypothetical answers. One of the ways in which abduction can be limited is by defining a set of abducible predicates. This is not proposed for the case of hypothetical answering, rather, if hypothetical answers are accumulating, it is practical to immediately return some of them to the questioner: provided the question answering situation allows for interactive information exchange between questioner and answering agent. Hypothetical answers and abductive hypotheses may be more quickly available in a top-down reasoner than proof results: it may be beneficial for a reasoner to be abductive in the interest of conserving time and other resources.

5 Summary

In summary, hypothetical answers provide a way of giving a questioner information in the face of incomplete knowledge of an answering agent. Abductive hypotheses represent a subset of the the hypothetical components of hypothetical answers. Some of the differences between hypothetical answers and what are traditionally considered as abductive hypotheses are the following:

- Hypothetical answers whose only difference is the level of generality are *all* interesting.
- The hypothetical component of an answer may contain existentially or universally quantified, locally scoped variables.
- Hypothetical answers have a purely syntactic definition that is domain-independent and does not rely on a predefined set of predicates.
- The purpose of hypothetical answers is not the adoption of a hypothesis that is neither provably true nor provably

false by the reasoner that produces it, rather, their purpose is to provide at least some information to a distinct questioner.

References

[Bromberger, 1992] Sylvain Bromberger. *On What We Know We Don't Know*. The University of Chicago Press; Center for the Study of Language and Information, Chicago; Stanford, 1992.

[Burhans and Shapiro, 1999] Debra T. Burhans and Stuart C. Shapiro. Finding Answers with a Resolution Theorem Prover. In *AAAI-F99 Symposium on Question Answering*, Cambridge MA, 1999. AAAI Press.

[Cholvy and Demolombe, 1986] L. Cholvy and R. Demolombe. Querying a Rule Base. In *Proceedings of the First International Workshop on Expert Database Systems*, pages 365–371, 1986.

[Cholvy, 1990] Laurence Cholvy. Answering Queries Addressed to a Rule Base. *Revue d'Intelligence Artificielle*, 4(1), 1990.

[Cox and Pietrzykowski, 1987] P.T. Cox and T. Pietrzykowski. General diagnosis by abductive inference. In *Proceedings of the IEEE Symposium on Logic Programming*, pages 183–189. IEEE Press, 1987.

[Demolombe, 1992] Robert Demolombe. A strategy for the computation of conditional answers. In Bernd Neumann, editor, *Proceedings of the 10th European Conference on Artificial Intelligence, ECAI 92*, pages 134–138. John Wiley and Sons, 1992.

[Demolombe, 1997] Robert Demolombe. Uncertainty in intelligent databases. In Amihai Motro and Philippe Smets, editors, *Uncertainty Management in Information Systems*, pages 89–126. Kluwer Academic Publishers, 1997.

[Green and Raphael, 1968] C. Cordell Green and Bertram Raphael. The user of theorem-proving techniques in question-answering systems. In Sr. Richard B. Blue, editor, *Proceedings of 23rd National Conference*, pages 169–181, Princeton, NJ, 1968. Association for Computing Machinery, Brandon/Systems Press, Inc.

[Green, 1969a] Cordell Green. Applications of theorem proving to problem solving. In Donald E. Walker and Lewis M. Norton, editors, *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 219–239. IJCAI, IJCAI, May 1969.

[Green, 1969b] Cordell Green. Theorem-proving by resolution as a basis for question-answering systems. In D. Michie and B. Melzer, editors, *Machine Intelligence 4*, pages 183–205. Edinburgh University Press, 1969.

[Imielinski, 1987] Tomasz Imielinski. Intelligent Query Answering in Rule Based Systems. *Journal of Logic Programming*, 4(3):229–257, September 1987.

[Kakas *et al.*, 1998] A. C. Kakas, R. A. Kowalski, and F. Toni. The role of abduction in logic programming. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Handbook of Logic in AI and Logic Programming 5*, pages 235–324. Oxford University Press, 1998.

[Kaplan, 1981] S. Jerrold Kaplan. Appropriate responses to inappropriate questions. In Aravind K. Joshi, Bonnie J. Weber, and Ivan A. Sag, editors, *Elements of Discourse Understanding*, pages 127–144. Cambridge University Press, Cambridge, 1981.

[Lehnert, 1978] Wendy G. Lehnert. *The Process of Question Answering*. Lawrence Erlbaum, 1978.

[Luckham and Nilsson, 1971] David Luckham and Nils J. Nilsson. Extracting information from resolution proof trees. *Artificial Intelligence*, 1971.

[Michalski, 1993] R. S. Michalski. Inferential theory of learning as a conceptual basis for multistrategy learning. *Machine Learning*, 11:111–151, 1993.

[Motro, 1989] Amihai Motro. Using integrity constraints to provide intensional answers to relational queries. In *Proceedings of VLDB89, the 15th International Conference on Very Large Databases*, pages 237–246, Amsterdam The Netherlands, August 1989. VLDB89.

[Reiter, 1978a] Raymond Reiter. Deductive question-answering in relational data bases. In Herve Gallaire and Jack Minker, editors, *Logic and Databases*, pages 149–177. Plenum Press, New York NY, 1978.

[Reiter, 1978b] Raymond Reiter. On closed world data bases. In Herve Gallaire and Jack Minker, editors, *Logic and Databases*, pages 55–76. Plenum Press, New York NY, 1978.

[Reiter, 1986] Raymond Reiter. A sound and sometimes complete query evaluation for relational databases with null values. *Journal of the ACM*, 33(2):349–370, 1986.

[Rich and Knight, 1991] Elaine Rich and Kevin Knight. *Artificial Intelligence*. McGraw Hill, 2 edition, 1991.

[Webber, 1986] Bonnie Lynn Webber. Questions, Answers and Responses: Interacting with Knowledge-Base Systems. In Michael L. Brodie and John Mylopoulos, editors, *On Knowledge Base Management Systems: Integrating Artificial Intelligence and Database Technologies*, chapter 26, pages 365–402. Springer-Verlag, 1986.

[Wolstenholme, 1988] Dave Wolstenholme. *Research and Development in Expert Systems IV*, chapter Saying "I Don't Know" and Conditional Answers, pages 115–125. The British Computer Society Workshop Series. Cambridge University Press, Cambridge, 1988.