# Common Sense Retrieval

**A. Julian Craddock**
University of B.C.
Vancouver, B.C., V6T 1W5
craddock@cs.ubc.ca

## Abstract

An important and readily available source of knowledge for common sense reasoning is partial descriptions of specific experiences. Knowledge bases (KBs) containing such information are called *episodic knowledge bases* (EKB). Aggregations of episodic knowledge provide common sense knowledge about the unobserved properties of 'new' experiences. Such knowledge is *retrieved* by applying statistics to a relevant subset of the EKB called the *reference class.*

I study a manner in which a corpus of experiences can be represented to allow common sense retrieval which is: 1. Flexible enough to allow the common sense reasoner to deal with 'new' experiences, and 2. In the simplest case, reduces to efficient database look-up. I define two first order dialects, $L$ and $QL$. $L$ is used to represent experiences in an episodic knowledge base. An extension, $QL$ is used for writing queries to EKBs [1].

## The problem

A corpus of declarative knowledge consisting of general concepts and their associated properties is frequently assumed adequate for common sense reasoning. At odds with this assumption I suggest that knowledge about *specific experiences* is necessary for flexible and efficient common sense reasoning. An experience is a common sense reasoner's (CSR) observation of its domain which can be described in terms of the properties of a set of observed objects. The descriptions of a CSR's experiences form its episodic knowledge base (EKB).

The problem addressed in this paper is to use knowledge about the specific experiences described in an EKB to make inferences about the unobserved properties of 'new' experiences; Inferences of the form 'Will the bird sitting on the lawn fly if I try to catch it?'. Such inferences are made by **retrieving common sense knowledge**: *General stable knowledge about the unobserved properties obtained by aggregating over a set of similar experiences in the EKB.* Thus, the CSR might infer that the bird sitting on the lawn flys if flying is part of the common sense knowledge about similar birds.

In this document I provide a first order framework appropriate for: 1. Describing experiences, and 2. Making inferences by retrieving common sense knowledge directly from an EKB. The framework differs from many existing knowledge representation techniques as it provides information about specific experiences. I judge the effectiveness of my model by its ability to flexibly and efficiently retrieve common sense knowledge.

## My solution

In order to solve the problem my model applies statistical knowledge obtained from old experiences to a new experience as follows:

**INPUT:** 1. A partial specification of the 'observed' properties of a new experience and some additional 'unobserved' properties, 2. An EKB of partial specifications of past experiences.

**SELECT:** a *reference class* [Kyburg, 1983][Kyburg, 1988] of experiences described in the EKB relevant to testing the hypothesis that the unobserved properties are also true of the new experience.

**MODIFY:** the criteria for membership in the reference class if there are no directly relevant experiences described in the EKB. This process provides the model with the flexibility necessary to reason about 'new' experiences: experiences for which there are no similar experiences already described in the EKB.

**AGGREGATE:** over the members of the reference class using statistical techniques to retrieve common sense knowledge about the unobserved properties.

**OUTPUT:** a conditional probability representing a measure of the support provided by the experiences in the EKB for the hypothesis.

[1]This research was supported by a NSERC of Canada postgraduate scholarship.

## Representing Experiences

In this section I briefly describe a first order language $L$ with the following properties:

1. Individual ground sentences of $L$ provide partial descriptions of specific experiences. An EKB is a set of such sentences.
2. A distinguished set of distinct ground terms allows the number of observations of objects described in an EKB sharing the same properties to be counted.

**Example 1** *$L$ allows a CSR to describe experiences such as 'two grad students drinking beer'. It also allows a CSR to count how many observations of 'two grad students drinking beer' it has described in its EKB'*

### Syntax

Objects are described using a distinguished set of distinct ground terms composed of *features*, *values* and *labels*. Individual features, values and labels are denoted by $f(i)$, $v(j)$ and $l(k)$ respectively, for some natural numbers $i$, $j$, and $k$. I sometimes write 'colour', 'size', 'name', ... for $f(i)$, $f(j)$, $f(k)$, ..., and 'red', 'large', 'julian', ... for $v(i)$, $v(j)$, $v(k)$ ....

Each feature has a set of possible *values*. For example, the feature 'colour' might have the set of possible values 'red', 'green', 'yellow', and 'blue'. I assign values to features using n-ary relations. For example, I can describe objects which are 'red' 'cars' as follows: $R_0(x, red, colour) \wedge R_0(x, car, type)$. The result of assigning a value to a feature is called a primitive *property*. Complex properties are formed from primitive properties using the logical operators of $L$ which include the usual FOL operators.

In order for a CSR to count the number of observations of objects with the same property the EKB must be able to tell the observations apart. To allow this each object observed in an experience is denoted by an individual *label* which is unique to that object and that experience. For example, if a EKB contains only the ground sentence

1. $R_0(l(4), red, colour) \wedge R_0(l(5), red, colour)$

then I say that it knows of two observations of objects which are 'red' in 'colour'. If it contains the two ground sentences

1. $R_0(l(4), red, colour) \wedge R_0(l(5), red, colour)$
2. $R_0(l(36), car, type) \wedge R_0(l(36), red, colour)$

then I say that it knows of three observations of objects which are 'red' in 'colour' and of one observation of two objects which are both 'red' in colour'.

In addition to the usual FOL axioms I include axioms saying that there are distinct numbers and that distinct numbers are mapped to distinct individuals, features and values. This ensures that individual values, features, and labels are distinct, avoiding coreferentiality problems which the unique names assumption can not solve. I also add axioms to define the *complementation* operator and *exclusivity* relation:

• **The complementation operator '-'** describes set complementation, a restricted form of boolean negation [Craddock, 1992]. The usual axioms applicable to normal boolean negation apply to '-', with the addition of:

$$-R_{n-1}(l(i), \ldots, l(n), v(j), f(k)) \leftrightarrow (\exists y)(R_{n-1}(l(i), \ldots, l(n), v(y), f(k)) \wedge \neg(y = j))$$

and

$$-(\exists y)(R_{n-1}(l(i), \ldots, l(n), v(y), f(k)) \wedge \neg(y = j)) \leftrightarrow (\exists y)(R_{n-1}(l(i), \ldots, (n), v(y), f(k)) \wedge (y = j))$$

for some natural numbers $i$, ..., $n$, $j$ and $k$.

**Example 2** *The complement of 'red' with respect to 'colour' is the set of all possible values of 'colour' excluding 'red'. I write $-R_0(l(7), red, colour)$ if some individual has some colour other than red.*

• **The exclusive predicate** defines exclusive features. A feature $f(i)$ is exclusive if $E(f(i))$ is an axiom of $L$ such that:

$$(\forall x_1)\ldots(\forall x_n)(\forall y)(\forall y')(\forall z)$$
$$[(R_{n-1}(l(x_1), \ldots, l(x_n), l(y), f(z)) \wedge E(f(z)) \wedge \neg(y = y')) \rightarrow -(R_{n-1}(l(x_1), \ldots, l(x_n), l(y'), f(z)))]$$

**Example 3** *'Sex' is exclusive if it only makes sense to describe an individual by assigning a* single *value such as 'male' or 'female', but not both, to the feature 'sex'.*

**Example 4** *Suppose I wish to say that all features are exclusive. I write: $(\forall x)(E(x))$ in $L$ without having to use second order quantification.*

**Example 5** *For example,*

$$R_0(l(34), red, colour) \vee R_0(l(34), yellow, colour)$$

*specifies that $l(34)$ is either red or yellow in colour. If $E(colour)$ is an axiom of $L$ then*

$$(\exists y)(R(l(34), y, colour) \wedge \neg(y = blue))$$

*is a theorem.*

### An episodic KB

An **episodic KB** (EKB) is a closed set of axioms. These include: 1. All the axioms of $K$ (a subset of $L$ containing the theorems), and 2. A finite set of ground sentences written in $L$. Each element of the latter set is a partial description of an **experience**.

The following is a simple example of an episodic KB, excluding the axioms of $K$:

**Example 6**

1) $R_0(l(0), red, colour) \wedge R_0(l(0), lrg, size)$
2) $R_0(l(1), Ph.D., deg.) \vee R_0(l(1), MSc., deg.)$
3) $(\exists y)(R_0(l(2), y, colour) \wedge \neg(y = red))$
4) $-(R_0(l(5), blue, colour) \wedge R_0(l(5), lrg, size))$
5) $R_0(l(6), red, colour) \wedge R_0(l(7), red, colour))$

⋮ ⋮

$n$) $R_0(l(34), Ned, name) \wedge R_0(l(34), red, colour)$

**Definition 1** *I write $EKB \vdash \alpha$ if the wff $\alpha$ is syntactically implied by the the axioms of the $EKB$.*

**Example 7** *Given the EKB in Example 3, I can write $EKB \vdash R_0(x, red,\ colour)$ and $EKB \not\vdash R_0(x, MSc., hasdegree)$.*

## Querying an episodic KB

In this section I extend $L$ to form a more expressive language $QL$ which allows the the $CSR$ to ask an $EKB$ for more detailed information about particular experiences. In particular I define a *probability term* $Prob(\alpha|\beta)_{ekb}$ such that $\alpha$ and $\beta$ are wff of $L$ and *ekb* denotes an episodic knowledge base as defined previously. Probability terms are interpreted as the conditional probability, with respect to an episodic knowledge base *ekb*, of $\alpha$ being true of a particular experience given that $\beta$ is true.

**Example 8** *Consider the query 'What is the conditional probability of a particular lawyer being rich?'. I can write this in $QL$ as $Prob(\ \alpha\ |\ \beta)_{ekb}$ such that $\beta$ is defined as $R_0(l(34), lawyer,\ occupation)$ and $\alpha$ is defined as $R_0(l(34), rich,\ financial\ \ status)$.*

In $QL$ conditional probabilities are calculated by *counting* the number of observations of objects with certain properties in the $EKB$. In order to count I include axioms for defining sets and set cardinalities in $QL$. For example, I define the *cardinality* of a wff $\alpha$ as the number of observations described in the $EKB$ of objects for which the property $\alpha$ is true.

**Definition 2** *Let $\alpha$ be a wff of $L$ such that $\langle x_1, \ldots, x_n\rangle$ is the n-tuple of labels denoting occurrences of individuals in $\alpha$. Let $\langle t_1, \ldots, t_n\rangle$ be a n-tuple of terms in $L$. $\alpha(x_1/t_1, \ldots, x_n/t_n)$ is the result of substituting each occurrence of $x_i$ by $t_i$.*

$$|\{\langle t_1, \ldots, t_n\rangle\ :\ EKB \vdash \alpha(x_1/t_1, \ldots, x_n/t_n)\}|$$

*is the cardinality of a wff $\alpha$, written $|\alpha|_{ekb}$, with respect to n'*

**Example 9** $|(R_0(l(987),\ red,\ colour) \wedge R_0(l(987),\ large,\ size))|_{ekb}$ *is the number of observations described in the EKB of objects with the property 'has colour red and size large'.*

The set of all observations described in the $EKB$ which are relevant to calculating the desired conditional probability is called the **reference class** of the probability term.

**Definition 3** *The reference class of $Prob(\alpha|\beta)_{ekb}$ is*

$$\{\langle t_1, \ldots, t_n\rangle\ :\ KB \vdash (\alpha \wedge \beta)(x_1/t_1, \ldots, x_n/t_n)\} \bigcup \{\langle t_1, \ldots, t_n\rangle\ :\ EKB \vdash (-\alpha \wedge \beta)(x_1/t_1, \ldots, x_n/t_n)\}$$

If the reference class of a probability term is not empty then I say that the probability term **succeeds**.

Probability terms with non-empty reference classes are interpreted as:

**Definition 4**

$$Prob(\alpha|\beta)_{ekb} \equiv_{df} \frac{|(\alpha \wedge \beta)|_{ekb}}{|(\alpha \wedge \beta)|_{ekb} + |(-\alpha \wedge \beta)|_{ekb}}$$

The denominator of each probability term is the cardinality of the reference class: a count of all the observations described in the EKB which are relevant to calculating the desired conditional probability.

**Example 10** *Consider the query 'What is the probability of a particular lawyer having a financial status other than rich?'. I can write this as $Prob(-\alpha|\beta)$, such that $\beta$ is defined as $R_0(l(987),\ lawyer,\ occupation)$ and $\alpha$ is defined as $R_0(l(987),\ rich,\ financial\,status)$. The numerator of the probability term is the cardinality of the set of all n-tuples of terms that satisfy $(-\alpha \wedge \beta)$. The reference class is the set of all n-tuples of terms that satisfy $(\alpha \wedge \beta)$, e.g. non-rich lawyers, and $(-\alpha \wedge \beta)$, e.g. non-rich lawyers, given the EKB. If $E(financial\ status)$ is an axiom, observations of 'lawyers' with financial statuses such as 'poor' and 'middle class' will be counted in $|(-\alpha \wedge \beta)|_{ekb}$*

## Modifying the reference class

In this section I discuss probability terms which **fail**:

**Definition 5** *A probability term $Prob(\alpha|\beta)_{ekb}$ fails if its reference class is the empty set.*

Probability terms fail because the knowledge in the knowledge base is incomplete: there are no experiences described in the EKB which are relevant to calculating the desired conditional probability. In this section I discuss two general mechanisms for identifying alternative, non-empty, reference classes using the episodic knowledge described in an EKB. I call these mechanisms *generalization* and *chaining*.

I argue that *generalization* and *chaining* are often more appropriate than techniques which depend upon the inclusion of *extra-logical* information [Asher and Morreau, 1991] in the KB, i.e., *assumptions of irrelevance* and *principles of defeasibility* [Touretzky *et al.*, 1991]. Such techniques often 'cover up' deficiencies in the KB's underlying representational strategy. Thus, I don't agree with suggestions that the complexity of these techniques is necessarily indicative of the complexity of the underlying problem (see for example [Touretzky *et al.*, 1991]).

Rather, I believe that the complexity is often an artifact of using an incorrect representational strategy. For example, knowledge representations such as taxonomic hierarchies, decision trees, and associative models of memory are founded on the basic premise that episodic knowledge should be *parsimoniously organized* on the basis of intuitively or statistically apparent structure. However, there are two problems with this premise:

1. Given a reasonable complex domain, the observed data may be too *sparse* to allow the identification of any useful structure.

2. Even in the presence of adequate data, straightforward Bayesian arguments show that only external information about the likely mix of queries is relevant to this determination [Schaffer, 1991].

The premise not only fails to address the problem of identifying structure in the absence of adequate data, but as a result of bias [Schaffer, 1991] in the representational strategy the data may be 'underfitted'. As a result the information necessary for answering *unforeseen* queries may be missing.

As a result, I argue that retrieval from an EKB is:

1. Appropriate in the presence of *sparse* data as no apriori structure needs to be identified.
2. Flexible, as episodic knowledge is structured in direct response to specific queries.

I now discuss how *generalization* and *chaining* apply the episodic knowledge in an KB directly to the problem of identifying a new reference class. Both mechanisms are motivated by two knowledge structures frequently found in the knowledge representation literature: 1. Knowledge hierarchies, and 2. Associative chains.

## Generalization

A probability term $Prob(\alpha|\beta)_{ekb}$ is **generalized** by relaxing the 'membership requirements' of its reference class. However, generalization through arbitrary syntactic manipulation of the wffs $\alpha$ and $\beta$ is inadequate. For example, the reference class of the query 'What is the probability that young lawyers are rich?' should not be arbitrarily relaxed to include observations of the financial status of 'young lawyers', 'dwarf elephants' and 'dead dogs'.

Instead, generalization should be constrained by the episodic knowledge in the EKB. In particular, I suggest that this knowledge allows us to

> "... ignore special characteristics of the ... event under consideration which are not known to be related to the property in question." [Kyburg, 1969]

without relying upon the inclusion of additional knowledge by a knowledge base designer.

Properties of experiences can be 'ignored' by expanding feature values:

**Definition 6** *The value $v(i)$ of a feature $f(j)$ in $R_{n-1}(x_1, \ldots, x_n, v(i), f(j))$ is expanded by replacing $i$ with an existentially quantified variable $y$ to get $(\exists y)(R_{n-1}(x_1, \ldots, x_n, v(y), f(j)))$.*

**Example 11** *$Prob(\alpha|\beta')_{ekb}$ is a generalization of $Prob(\alpha|\beta)_{ekb}$ if $\beta'$ is the result of expanding one or more property predicates in $\beta$.*

All the possible generalizations obtained by expanding feature values can be ordered in a partial lattice:

**Lemma 1** *Let $Prob(\alpha|\beta)_{ekb}$ be a probability term. The set of all possible generalizations of $Prob(\alpha|\beta)_{ekb}$ with non empty reference classes formed by expanding the property predicates of $\beta$ forms a partial lattice.*

In [Craddock, 1992] I argue that the most relevant generalization of a probability term is the minimal element in the lattice of generalizations. This method adopts a principle similar to the *principle of specificity* described in the *non-monotonic logic* literature, (see for example, Poole [Poole, 1990]). If there are several equally minimal elements, the one obtained by expanding the least statistically relevant feature is chosen, in accordance with Kyburg's [Kyburg, 1988] conditions of statistical epistemological relevance.

**Example 12** *Suppose $Prob(R_0(x,\ flys,\ moves) \mid R_0(x,\ red,\ colour) \wedge R_0(x,\ bird,\ type))$ fails. There are two minimal elements in the lattice of generalizations:*

$$(1)\quad Prob\left(R_0(x, flys, moves) \,\middle|\, \begin{array}{c} R_0(x, y, colour) \\ \wedge R_0(x, bird, type) \end{array}\right)$$

$$(2)\quad Prob\left(R_0(x, flys, moves) \,\middle|\, \begin{array}{c} R_0(x, red, colour) \\ \wedge R_0(x, y, type) \end{array}\right)$$

*Let $abs(r^p_{\langle x,y\rangle})$ be the absolute value of the correlation coefficient between two variables $x$ and $y$ given some property $p$. Suppose,*

$$abs(r^{type=bird}_{\langle moves,colour\rangle}) < abs(r^{colour=yellow}_{\langle moves,type\rangle)})$$

*then generalization (1) is the appropriate generalization of the original query.*

Given a probability term $Prob(\alpha|\beta)_{ekb}$, episodic knowledge contained in the EKB can be analyzed using well understood statistical techniques for measuring the association between features. These techniques can be used to identify the primitive properties described in $\beta$ which are most relevant to predicting the truth of $\alpha$.

## Chaining

As a result of the partial nature of descriptions of experiences in probability terms and EKBs there are cases in which every generalization of a probability term will have an empty reference classes. An intuitive solution is to **chain** rather than generalize the original probability term.

**Example 13** *Suppose the episodic knowledge contained in a particular EKB was collected using two experiments:*

1. *Experiment 1 recorded the observed properties of objects using only the features 'Virus' and 'Syndrome', and*
2. *Experiment 2 recorded the observed properties of objects using only the features 'Syndrome' and 'Cancer'.*

*Now, let $\alpha$ be defined as $R_0(l(6),\ skin,\ cancer)$, $\beta$ as $R_0(l(6),\ HIV^+,\ virus)$, and $\gamma$ as $R_0(l(6),\ AIDS,\ syndrome)$. The probability term $Prob(\alpha \mid \beta)_{ekb}$ fails. Furthermore, as the partial lattice of generalizations with non-empty reference classes is also empty, generalization also fails.*

*However, suppose that the conditional probability $Prob(\gamma \mid \beta)_{ekb}$ is high, e.g. If you are 'HIV$^+$' then you have 'AIDS'. The reference class of the original probability term 'overlaps' the reference class of $Prob(\alpha \mid \gamma)_{ekb}$, e.g. If you have 'AIDS' then you have 'Skin cancer', and thus provides an estimate of the desired conditional probability (assuming 'noisy-or' relationships [Pearl, 1988]). In this example, retrieval has 'chained' from the property 'HIV$^+$' to 'AIDS' to 'Skin cancer' in a manner similar to reasoning in associative models of human memory.*

Intuitively, if one set of observations overlaps another set then they are *similar* and related in some way. Depending on the strength of this similarity, predictions made from the new reference class should be representative of predictions made from the original reference class. The validity of this assumption depends upon how *similar* the new reference class is to the old and how well this similarity can be measured. For example, the metric for measuring similarity presented in the previous example is only one of many possibilities.

Chaining can be defined recursively as follows:

**Definition 7** *$Prob(\alpha|\gamma)_{ekb}$ is a chain on $Prob(\alpha \mid \beta)_{ekb}$ iff:*

1. $\{\langle t_1,\ldots,t_n\rangle : EKB \vdash \gamma(t_1/x_1,\ldots,t_m/x_m)\} - \{\langle t_1,\ldots,t_m\rangle : KB \vdash \beta(t_1/x_1,\ldots,t_m/x_m)\} \neq \emptyset$, *i.e. overlap exists, and/or*
2. *$\exists\delta$ such that $Prob(\alpha|\gamma)_{ekb}$ is a chain on $Prob(\alpha|\delta)_{ekb}$ and $Prob(\alpha|\delta)_{ekb}$ is a chain on $Prob(\alpha|\beta)_{ekb}$.*

As with generalization a partial lattice of possible chainings can be defined. However, as there is a potentially large number of ways to chain a probability term the lattice may be very large. In Craddock [Craddock, 1992] I discuss several classes of heuristics for choosing the most appropriate chain. Some of these are concerned with the length or degree of the chain and are similar to those addressed by Touretzky et. al. [Touretzky *et al.*, 1991]. Others are concerned with assessing the *similarity* of reference classes using techniques found in machine learning [Aha *et al.*, 1991]. In [Craddock, 1992] I examine the applicability of these heuristics using data obtained from the "UCI repository of machine learning data bases and domain theories.".

## Relationship to other work

I assume that knowledge about specific past experiences is necessary for common sense reasoning. In [Craddock, 1992] I show that this assumption is supported by the psychological literature on human *episodic memory*. Of particular interest are results showing that humans make predictions using reference classes containing as few as one previous experience. Using such knowledge my model is able to treat common sense retrieval as an example of the reference class problem [Kyburg, 1983] [Kyburg, 1988]:

> The problem of identifying statistical knowledge about past experiences which is epistemically relevant to making an inductive inference about the unobserved properties of a new experience.

I extend this problem to include the modification of empty reference classes. The extension provides the flexibility necessary to deal with incomplete knowledge.

My model assumes that a good basis on which to retrieve common sense knowledge is something like an associative database, where retrieval is merely look-up or pattern matching (the model is relevant to databases as statistics are computed from individuals and not from classes of unknown cardinality). This assumption forms the basis for *efficient* common sense reasoning in many models of reasoning, i.e., [Levesque, 1986], [Levesque, 1989],[Etherington *et al.*, 1989], [Frisch, 1987], [Davis, 1990], [Davis, 1987], and [Stanfill and D., 1986]. However, unlike many of these models and other machine learning techniques, i.e., [Aha *et al.*, 1991], I reason directly with partial descriptions of experiences in probability terms and EKBs.

My model does not depend upon pre-defined common sense knowledge supplied in the form of defaults or preference orderings as in many representations of common sense knowledge, i.e., [Etherington, 1987], [Boutilier, 1991], [Poole, 1990], [Konolige, 1987]. Although such techniques are formally well understood they can be criticized as depending upon the inclusion of additional 'extra-logical' knowledge [Asher and Morreau, 1991]. Nor does my model depend upon the episodic knowledge being 'pre-structured' into general concepts and relations. Instead, episodic knowledge is manipulated in direct response to specific queries.

My model can use techniques similar to those used by machine learning in order to measure relevance and similarity. However, unlike machine learning my model does not apply the techniques to improve the predictability of properties from a *static* reference class; It applys the techniques in order to identify a *new* reference class from which the conditional probability of the properties can be predicted. As suggested by Kyburg [Kyburg, 1988], considerations as to the appropriate application of these techniques mirror considerations appropriate to the formalization of nonmonotonic logics.

My model uses simple techniques to aggregate over sets of incomplete descriptions of experiences, obtaining common sense knowledge directly from an EKB. Furthermore, episodic knowledge is used directly to respond to the problem of incomplete knowledge. I argue that this approach is: 1. Potentially efficient as it reduces to database lookup in the simplest case, and 2. More *flexible* than existing models in dealing with experiences with which the common sense reasoner has no 'relevant' prior experience.

## Discussion

The model presented in this paper is a principled way of incorporating specific knowledge about experiences into common sense reasoning. It provides :

- An expressive and well understood framework for talking about experiences - knowledge which computers can readily obtain about the real world.
- A potentially flexible and efficient mechanism for retrieving common sense knowledge: 1. Directly from past experiences, and 2. In the presence of sparse data.
- A mechanism for mapping from inductive techniques for retrieving 'statistical' common sense knowledge to well understood deductive techniques for reasoning with it (see for example, Bacchus [Bacchus, 1990]).
- A well defined, formal framework, for talking about reference classes. In particular, the framework allows:
  1. The definition of a reference class of past experiences.
  2. The modification of an empty reference class.
  3. The retrieval of information from a reference class.

Common sense retrieval is not intended as a replacement for the powerful deductive techniques found in the common sense reasoning literature. Rather, I show how retrieval can be approached in a flexible and efficient manner. I believe that the inefficiency and inflexibility of many existing models of common sense reasoning is a direct result of not explicitly representing knowledge about specific past experiences.

There are two obvious problems with the model:

1. Its reliance on large sets of data. This is becoming more acceptable given recent advances in computational hardware. In [Craddock, 1992] I describe data reduction in which only distinct experiences - experiences with the same form - are represented as separate entries in the KB.
2. The methods for modifying the reference class briefly discussed in this paper are computationally expensive. However, I point out that it is reasonable to assume that a CSR will retrieve what it knows quickly and retrieve what it doesn't know much more slowly.

On going work is currently addressing these problems.

### Acknowledgements

I would like to acknowledge the help of P. Turney and D. Poole.

## References

Aha, D.; Kibler, D.; and Albert, M. 1991. Instance-based learning algorithms. *Machine Learning* 6(1):37–66.

Asher, N. and Morreau, M. 1991. Commonsense entailment: A modal theory of nonmonotonic reasoning. In *Proceedings of 12th IJCAI.* 387–392.

Bacchus, F. 1990. *Representing and reasoning with probabilistic knowledge: a logical appropach to probabilities.* MIT Press.

Boutilier, C. 1991. Inaccessible worlds and irrelevance: Preliminary report. In *Proceedings of 12th IJCAI.* 413–418.

Craddock, J. 1992. Access and retrieval of information from an episodic knowledge representation. *University of British Columbia* Forthcoming Ph.D. Thesis.

Davis, L. 1987. *Genetic algorithms and simulated annealing.* Morgan Kaufmann.

Davis, E. 1990. Partial information and vivid representations. In *Representation in mental models.*

Etherington, D.; Borgida, A.; Brachman, R.; and Kautz, H. 1989. Vivid knowledge and tractable reasoning: Preliminary report. In *International Joint Conference on AI.* 1146–1152.

Etherington, D. 1987. A semantics for default logic. In *Proceedings of 8th IJCAI.* 495–498.

Frisch, A. 1987. Knowledge retrieval as specialised inference. Technical Report 214, University of Rochester.

Konolige, K. 1987. On the relation between default theories and autoepistemic logic. In *Proceedings of 8th IJCAI.* 394–401.

Kyburg, H.E. 1969. *Probability theory.* Englewood Cliffs, N.J.: Prentice-Hall.

Kyburg, H. E. 1983. The reference class. *Philosophy of Science* 50:374–397.

Kyburg, H. 1988. Epistemological relevance and statistical knowledge. Technical Report 251, University of Rochester.

Levesque, H. 1986. Making believers out of computers. *Artificial Intelligence* 30:81–108.

Levesque, H. 1989. Logic and the complexity of reasoning. Technical Report KRR-TR-89-2, University of Toronto.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference.* Morgan Kaufmann.

Poole, D. 1990. Dialectics and specificity: conditioning in logic-based hypothetical reasoning (preliminary report). In *CSCSI.*

Schaffer, C. 1991. Overfitting avoidance as bias. In *Proceedings of the IJCAI Workshop on Evaluating and Changing Representation in Machine Learning,* Sydney, Australia.

Stanfill, C. and D., Waltz 1986. Toward memory-based reasoning. *Communciations of the ACM* 29(12):1213–1228.

Touretzky, D.S.; Thomason, R.H.; and Horty, J.F. 1991. A skeptics menagerie: Conflictors, preemptors, reinstaters, and zombies in non-monotonic inheritence. In *Proceedings of 12th IJCAI.* 478–483.