A Logic for Approximate Reasoning*

Daphne Koller
Stanford University
Stanford, CA 94305
email: daphne@cs.stanford.edu

Joseph Y. Halpern
IBM Almaden Research Center
San Jose, CA 95120
mail: halpern@almaden.ibm.com

Abstract

We investigate the problem of reasoning with imprecise quantitative information. We give formal semantics to a notion of approximate observations, and define two types of entailment for a knowledge base with imprecise information: a cautious notion, which allows only completely justified conclusions, and a bold one, which allows jumping to conclusions. Both versions of the entailment relation are shown to be decidable. We investigate the behavior of the two alternatives on various examples, and show that the answers obtained are intuitively desirable. The behavior of these two entailment relations is completely characterized for a certain sublanguage, in terms of the logic of true equality. We demonstrate various properties of the full logic, and show how it applies to many situations of interest.

1 Introduction

In almost any situation involving quantitative information, some of the information is bound to be approximate and imprecise. Moreover, such imprecision can easily cause inconsistencies. Consider, for example, the following knowledge base KB:

> Bill is 1.8 meters tall.
> John is half a head taller than Bill.
> A head is 0.2 meters.

Although the information in this knowledge base is not intended to be completely precise, we might nevertheless want to conclude from it that John is 1.9 meters tall. It is clear that we want to view this conclusion as being only an approximation of the truth. In particular, if we later obtain the additional piece of information "John is 1.88 meters tall," we do not want to conclude that the resulting knowledge base KB' is inconsistent; rather, we view this as a problem due to inaccurate measurement. This shows that we cannot interpret "is" in approximate observations as true numeric equality, because we would end up deducing that the above knowledge base is inconsistent, thus enabling arbitrary conclusions.

The need for dealing with approximate information arises in many other contexts. We often want to say that a certain quantity (such as a probability) is very close to zero, without committing to a particular value. The technique of ϵ-semantics [Pea88] is based on this concept (see Section 6.3). When dealing with statistical information, we often use statements of the form "90% of birds fly;" however, we do not wish to infer that the number of birds is divisible by 10, as we could if we interpret this statement as "*precisely* 90% of birds fly." It is more appropriate to interpret it as "*approximately* 90% of birds fly." (See [GHK94] for a thorough discussion.) Problems relating to the intransitivity of the perceptual indistinguishability relation in human observations [SKLT89] can also be formulated and circumvented using approximate equality (see Section 6.2).

In this paper, we introduce a logic which enables us to deal with and reason about imprecise information and the inconsistencies that usually accompany it. Our logic extends standard real arithmetic with notions of approximate equality and inequality. We formalize approximate equality to allow some small but unspecified discrepancy between the values being compared.

Our main interest is in making deductions from knowledge bases, so we focus here on what we call *approximate entailment*, where we view "KB approximately entails φ" as meaning that we have reasonable justification for concluding φ given the knowledge base KB. For example, if we are interested in buying John a jacket, and we are given the knowledge base KB above, we would certainly think it justified to proceed under the assumption that John is about 1.9 meters tall. The problem becomes more difficult when we ask for inferences from the extended knowledge base KB' above. We present two different entailment relations that we

*This paper appears in *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR '92)*, 1992, pp. 153–164.

call *cautious entailment* and *bold entailment*. They differ in the degree to which they allow the agent to "leap to conclusions;" i.e., in the degree of default reasoning they incorporate. The knowledge base KB' cautiously entails that "John is approximately between 1.88 and 1.9 meters tall." Thus, given contradictory information, the cautious approach assumes the answer could be anywhere in between. On the other hand, the bold approach, given the same knowledge base, would be able to conclude that "John is approximately h meters tall" for each h between 1.88 and 1.9; any reasonable number can be used as a "guesstimate."

From this example, it is clear that cautious entailment is nonmonotonic: by adding additional information to the knowledge base KB, we lose the ability to deduce that "John is approximately 1.9 meters tall." On the other hand, bold entailment is usually monotonic in the sense that adding new data to the knowledge base does not force us to withdraw conclusions. From the knowledge base KB' we can still deduce that "John is approximately 1.9 meters tall." However, the bold logic is not a standard monotonic logic. Although we can deduce both that "John is approximately 1.9 meters tall" and that "John is approximately 1.88 meters tall," we cannot deduce their conjunction (see Section 5.2 for more details). At first, this might seem strange. But the intuition here is that, although we can work with any reasonable assumption about John's height, we do not want to work with contradictory assumptions simultaneously.

Both types of entailment can be reduced to the validity of a formula in the language of *real closed fields* [Tar51], and therefore are decidable. The decision procedure, however, does not give us much insight into the properties of entailment. To gain this insight, we consider several examples, and present general properties of our notion of approximate entailment. These show that approximate entailment agrees with our intuition in many situations. For example, we show that inferences made by either one of our entailment relations are always consistent with those obtained by taking approximate equality to be true equality. However, if the knowledge base is inconsistent with equality, as with KB' above, it entails only "reasonable" conclusions. We provide an elegant characterization of these entailment relations for a large sublanguage of our full language; in particular, the characterization justifies our choice of the names "bold" and "cautious." As a corollary to this characterization, we show that if our data is consistent even if approximate equality is treated as true equality, then we typically get *precisely* the conclusions that we get from true equality. Our characterization also shows that, for a large subclass of formulas, cautious entailment reduces to a variant of *preference semantics* [Sho87] (see Section 5.2).

While most of the paper focuses on issues concerning measurement, our approach is actually much more general. Given a notion of exact inference from a knowledge base with precise information, we can use our framework to extend it to a notion of approximate inference from a knowledge base of imprecise information. The notion of exact inference could well be probabilistic or nonmonotonic. In particular, we can apply these ideas to ϵ-semantics ([Pea88, GMP90]) and to the problem of computing degrees of belief from statistical information [GHK94] (see Section 6.3).

2 Syntax and Semantics

Since we want to focus on the basic issues arising from the problem of approximate numerical information, we restrict ourselves to considering a relatively simple framework where these issues arise. We begin with a core language $\mathcal{L}$, consisting of:

- the standard arithmetic operations of $+, -, \times, /$,
- the standard equality and inequality relations $=$ and $\leq$,
- a constant symbol d_r for each real number r (in our examples, we typically write, say, 0.1, rather than $d_{0.1}$),
- a countable collection $c_1, c_2, \ldots$ of uninterpreted constant symbols.

We form the set of *terms* by closing off the constants under $+, -, \times$ and $/$. The set $\mathcal{E}$ of *precise expressions* consists of formulas of the form $t = t'$ and $t \leq t'$ for terms t and t'. The language $\mathcal{L}$ is formed by closing off $\mathcal{E}$ under conjunction, disjunction, and negation.

In order to form the *approximate language* $\mathcal{L}^{\approx}$, we augment the language $\mathcal{L}$ with the *approximate equality and inequality* relations $\approx$ and $\preceq$. The set $\mathcal{A}$ of *approximate expressions* consists of formulas of the form $t \approx t'$ and $t \preceq t'$ for terms t and t'. The language $\mathcal{L}^{\approx}$ is formed by closing off $\mathcal{E} \cup \mathcal{A}$ under conjunction, disjunction, and negation.

We interpret $\mathcal{L}$ in the standard fashion. Terms are interpreted over the reals, with an additional undefined value ι (used to deal with the problem of division by zero). The symbols $+, -, \times, /, =, \leq$ receive their standard interpretation (extended to deal with ι), and the constant d_r is interpreted as the real number r.

Definition 2.1: A model v for $\mathcal{L}$ is a function assigning an element in $\mathbb{R} \cup \{\iota\}$ to each term, and a truth value to each formula, as follows:

- for each uninterpreted constant c_i, $v(c_i) \in \mathbb{R}$,
- for each constant d_r, $v(d_r) = r \in \mathbb{R}$,
- for each term $t \circ t'$, where $\circ \in \{+, -, \times, /\}$, we have $v(t \circ t') = v(t) \circ v(t')$ in the standard way, with the following exceptions:
  - if $v(t') = 0$, then $v(t/t') = \iota$,
  - if $v(t) = \iota$ or $v(t') = \iota$, then $v(t \circ t') = \iota$.
- if $\sim$ is one of $=$ or $\leq$, then $v(t \sim t')$ is *true* if both $v(t)$ and $v(t')$ are in $\mathbb{R}$ and $v(t) \sim v(t')$; otherwise

$v(t \sim t')$ is *false*,

- v is extended to Boolean combinations of precise expressions in the standard way. ∎

One might wonder why we take division to be a primitive operation in our language (at the cost of having to deal with ι), rather than defining it in terms of multiplication. The problem arises due to subtle interactions between division and the semantics of approximate equality, in such terms as a/b where $b \approx 0$. This is particularly relevant in such applications as ϵ-semantics. It turns out that the only way to handle such expressions appropriately is to allow division as a primitive operation (see Section 6.3).

To understand the interpretation of $\approx$ and $\preceq$, we first need to consider how we want to interpret a statement such as: "Bill is approximately 1.8 meters tall." We view such a statement as describing a measurement of Bill's height, taken with some unknown degree of inaccuracy. Thus, we take this to mean that Bill's height is within τ of 1.8 meters, for some unknown tolerance τ. In general, the tolerances for different measurements may be completely independent. We enforce this in our semantics by interpreting $\approx$ and $\preceq$ in a nonstandard, context-dependent manner; for each expression $e \in \mathcal{A}$, there is a tolerance $\tau(e)$ associated with e.[1]

Definition 2.2: The *tolerance function* τ is a function from $\mathcal{A}$ to $\mathbb{R}^+ = [0, \infty)$. Let $\mathcal{T}$ denote the set of tolerance functions. ∎

In this definition, we have chosen to allow 0 as a legal tolerance value; that is, the range of a tolerance function is $[0, \infty)$, not $(0, \infty)$. While this issue may seem minor, it has a number of side effects. For one thing, it allows us to state stronger theorems, with simpler proofs. But it also affects our ability to make certain inferences. We discuss this issue further in Section 5.1.

In order to relate the meaning of expressions in $\mathcal{L}^{\approx}$ to the semantics of $\mathcal{L}$, we need the following definition.

Definition 2.3: For a formula $\varphi \in \mathcal{L}^{\approx}$, and a fixed tolerance function τ, we define $\varphi[\tau] \in \mathcal{L}$ to be the same as φ, except that every approximate expression $t \approx t'$ is replaced by the expression $|t - t'| \leq d_{\tau(t \approx t')}$ and each expression $t \preceq t'$ is replaced by $(t - t') \leq d_{\tau(t \preceq t')}$. ∎

For example, if φ is $c_1 \approx c_2$, and τ is such that $\tau(c_1 \approx c_2) = 0.1$, then $\varphi[\tau] = |c_1 - c_2| \leq d_{0.1}$. Let τ_0 to be the tolerance function that assigns tolerance 0 to all expressions. Note that $\varphi[\tau_0]$ is precisely the result of interpreting all occurrences of "approximately equals" in φ as true equality. We say that φ is *consistent with equality* if $\varphi[\tau_0]$ is consistent.

Definition 2.4: An *augmented model* M for $\mathcal{L}^{\approx}$ is a pair (v, τ), where v is a model for $\mathcal{L}$ and τ is a tolerance function. For a formula $\varphi \in \mathcal{L}^{\approx}$, we define $M \models \varphi$ if $v \models \varphi[\tau]$. ∎

Note that $\tau(e)$ for expressions e that do not appear in a formula φ has no affect on the truth value of φ: if τ and τ' agree on all expressions that appear in φ, then for any v, we have $(v, \tau) \models \varphi$ iff $(v, \tau') \models \varphi$.

We define validity for $\mathcal{L}^{\approx}$ as usual: ψ is valid if $M \models \psi$ for all models M. The validity problem for $\mathcal{L}^{\approx}$ is of little interest, as the following example suggests.

Example 2.5: Let φ be $(c \approx 1) \Rightarrow (2 \times c \approx 2)$, and the model $M = (v, \tau)$, for $v(c) = 1.1, \tau(c \approx 1) = 0.1$, and $\tau(2 \times c \approx 2) = 0.15$. Then $M \not\models \varphi$, and therefore φ is not valid. ∎

In fact, as the following theorem shows, if we restrict attention to (Boolean combinations of) approximate expressions not involving division (which allows us to avoid all the complications of dealing with ι), the only valid formulas are those that are propositionally valid if we treat every approximate expression as a distinct primitive proposition.

Theorem 2.6: *Let $\varphi \in \mathcal{L}^{\approx}$ be a Boolean combination of approximate expressions not involving division. Let $\alpha(\varphi)$ be the propositional formula that results from replacing each approximate expression e in φ by a primitive proposition p_e. Then φ is valid over models of $\mathcal{L}^{\approx}$ iff $\alpha(\varphi)$ is propositionally valid.*

Thus, rather than considering validity, we concentrate on a different notion that we call *approximate entailment*.

3 Approximate Entailment

Given a knowledge base of approximate measurements, when should it entail a statement φ such as "John is approximately 1.9 meters tall"? We do not want to view φ as necessarily representing an actual measurement that was taken of John's height. Rather, we want it to be a useful working assumption. For example, if we are interested in buying John a suit, we may well be content with an approximate estimate of John's height. We do not expect formulas entailed by the knowledge base to be completely accurate. Moreover, we will rarely know exactly how accurate they are, since that depends on the accuracy of our initial measurements, which we do not typically know. However, we would like to have the property that the smaller the errors in the knowledge base, the smaller

[1] In some cases we may know that tolerances are related in some way. For example, we may know that different measurements were taken with the same measuring device, and therefore have the same maximum error. This issue was dealt with in [GHK94] by explicitly subscripting $\approx$ and $\preceq$ so that each $\approx_i$ denotes a different approximate equality relation. For the sake of simplicity, we have chosen not to extend the logic to express relationships between tolerances.

the errors in formulas entailed by the knowledge base. This is in the spirit of the standard ϵ-δ definition of limit.

Notice that it may not be possible to have all tolerances grow arbitrarily small simultaneously. For example, if our knowledge base consists of $(c \approx 1) \wedge (2c \approx 2.1)$, it is clear that both relevant tolerances cannot be arbitrarily small at the same time. We therefore introduce the concept of *minimal tolerance function*. Intuitively, this is one that chooses the smallest possible tolerances while still keeping the knowledge base consistent. We say that a tolerance function τ is *consistent with* KB if $KB[\tau]$ is satisfiable. We say that $\tau < \tau'$ for two tolerance functions τ, τ' if for all $e \in \mathcal{A}$, $\tau(e) \leq \tau'(e)$, and there exists some $e \in \mathcal{A}$ such that $\tau(e) < \tau'(e)$. We also define $||\tau||$ to be $\sup\{|\tau(e)| : e \in \mathcal{A}\}$. The tolerance function η is minimal for KB if it is in the closure of tolerance functions consistent with KB and there is no smaller tolerance function consistent with KB.

Definition 3.1: A tolerance function η is said to be *minimal for* KB if

1. for every $\epsilon > 0$, there exists a tolerance function τ consistent with KB such that $||\tau - \eta|| \leq \epsilon$,
2. there does not exist another tolerance function τ such that $\tau < \eta$ and τ is consistent with KB.

Let $\Omega(KB)$ be the set of tolerance functions minimal for KB. ∎

Recall that τ_0 is the tolerance function that assigns tolerance 0 to all expressions. It is easy to see that if KB is consistent with equality, then τ_0 is the unique minimal tolerance function for KB. Thus, the notion of minimal tolerance function becomes interesting only for knowledge bases that are inconsistent with equality.

Example 3.2: The "inconsistent" height knowledge base KB' from the introduction, written formally in our language, is the conjunction: $(c_B \approx 1.8) \wedge (c_J \approx c_B + c_H/2) \wedge (c_H \approx 0.2) \wedge (c_J \approx 1.88)$, where c_B denotes Bill's height, c_J denotes John's height, and c_H denotes the height of a head. Let τ be a tolerance function such that $\tau(e) = 0$ for any irrelevant expression e (not one of the four above), and let $\tau_1, \tau_2, \tau_3, \tau_4$ denote the values assigned by τ to the four expressions above. Let $\vec{\tau}$ denote $(\tau_1, \ldots, \tau_4)$. It is easy to see that KB' is consistent iff $\tau_4 \geq 0.02 - \tau_1 - \tau_2 - \tau_3/2$. Thus, if $\vec{\tau} = (0, 0, 0, 0.01)$, then τ is not a minimal tolerance function for KB' because it violates condition 1. On the other hand, if $\vec{\tau} = (0.03, 0, 0, 0.01)$, then τ is not minimal because it violates condition 2: there exists a smaller tolerance function consistent with KB' that assigns $(0.02, 0, 0, 0)$ to the relevant expressions. This last tolerance function is in fact minimal for KB', as is the one that assigns $(0.01, 0, 0.02, 0)$ to the relevant expressions. And generally,

$$\Omega(KB') = \{\tau : \tau_4 = 0.02 - \tau_1 - \tau_2 - \tau_3/2, \quad \tau(e) = 0 \text{ if } e \text{ is irrelevant}\} . ∎$$

We now give an example of a formula KB for which the set of tolerances consistent with KB is not closed, and some minimal tolerance function is not consistent with KB.

Example 3.3: Let KB be $(c_1 \times c_2 \approx 1) \wedge (c_1 \approx 0)$, and let $\tau_1 = \tau(c_1 \times c_2 \approx 1)$ and $\tau_2 = \tau(c_1 \approx 0)$. It is clear that for any value of τ_1 and any $\tau_2 > 0$, $KB[\tau]$ is consistent. Therefore, one minimal tolerance function for KB is τ_0, which is inconsistent with KB. Note that the tolerance function τ' for which $\tau'_1 = 1$ and $\tau'_2 = 0$ is also consistent with KB (and therefore fulfills condition 1). And, although $\tau_0 < \tau'$, there does not exist a tolerance function $\tau < \tau'$ which is consistent with KB. Therefore, τ' is also a minimal tolerance function, and $\Omega(KB) = \{\tau_0, \tau'\}$. ∎

Remark 3.4: Although a minimal tolerance function for KB is not necessarily consistent with KB, it is the case that KB is satisfiable (i.e., some tolerance function is consistent with KB) iff $\Omega(KB) \neq \emptyset$. From this point on, we consider only satisfiable KB's. ∎

Using the concept of minimal tolerance functions, we can now define entailment. Before we give the formal definitions, we give a little intuition. We would certainly like a knowledge base of the form $c \approx 1$ to entail, say $2c \approx 2$. Recall from Example 2.5 that $(c \approx 1) \Rightarrow (2c \approx 2)$ is not valid. However, given a tolerance τ_1 for $c \approx 1$, we can clearly find a tolerance τ_2 for $2c \approx 2$ to make it true, namely, $2\tau_1$. This is the key idea in our notion of entailment. Roughly speaking, we want it to be the case that KB entails φ if, given a tolerance function τ that makes KB true, we can find a tolerance function τ' that makes φ true. Clearly we want to put some constraints on τ' (for otherwise from $c \approx 1$ we could infer $c \approx 2$). We require that the closer τ is to a minimal tolerance function, the closer τ' is to τ_0. This corresponds to the intuition we described earlier, that the smaller the errors in the knowledge base, the smaller the errors in the conclusion. Note that we do not try to describe *how* τ' must go to τ_0 as a function of how τ gets small.

We can now define the first of our two notions of entailment.

Definition 3.5: We say that KB *cautiously entails* φ, written $KB \approx\!\!\models_c \varphi$, if for every minimal tolerance function $\eta \in \Omega(KB)$ there exists some function $f : \mathcal{T} \rightarrow \mathcal{T}$, and some $\epsilon > 0$ such that:

- for every tolerance function τ such that $||\tau - \eta|| < \epsilon$, we have $KB[\tau] \models \varphi[f(\tau)]$,
- for every sequence $(\tau^n)_{n=1}^{\infty}$ such that $\lim_{n\rightarrow\infty} \tau^n = \eta$, we have $\lim_{n\rightarrow\infty} f(\tau^n) = \tau_0$. ∎

The reader might wonder why we insist that $f(\tau^n)$ converge to τ_0, rather than to a minimal tolerance function. The reason is that, even if we have a knowledge base that is inconsistent with equality, we want the formulas entailed by the knowledge base to be consistent with equality, since we want our conclusions to be useful working assumptions. Thus, if we have a knowledge base KB such as $(c \approx 0) \wedge (c \approx 0.1)$, then we do not want to conclude KB, as we would be able to do if we just required that $f(\tau^n)$ approach a minimal tolerance function rather than τ_0. As we shall show, although KB does *not* cautiously entail KB, it does cautiously entail $0 \preceq c \preceq 0.1$, which seems more reasonable.

Bold entailment replaces most of the universal quantifiers in the definition of cautious entailment by existential quantifiers.

Definition 3.6: We say that KB *boldly entails* φ, written $KB \mathrel{|\!\!\approx_b} \varphi$, if either KB is unsatisfiable, or there exists some function $f : \mathcal{T} \rightarrow \mathcal{T}$, some minimal tolerance function $\eta \in \Omega(KB)$, and some decreasing sequence $(\tau^n)_{n=1}^{\infty}$ such that the following all hold:

- for all n, $KB[\tau^n] \models \varphi[f(\tau^n)]$,
- τ^n is consistent with KB for all n,
- $\lim_{n\rightarrow\infty} \tau^n = \eta$ and $\lim_{n\rightarrow\infty} f(\tau^n) = \tau_0$. ∎

Note that we restrict attention only to the tolerance functions τ^n in the first clause above, rather than requiring that this condition hold for all tolerance functions τ. The latter would also give us a reasonable notion of entailment. We chose our condition because it leads to a bolder notion of entailment; that is, it allows strictly more formulas to be entailed by a knowledge base. We have tried to make bold entailment as liberal as reasonably possible, while making cautious entailment as conservative as reasonably possible. Clearly other intermediate notions of entailment are possible.

As the names suggest, $\mathrel{|\!\!\approx_c} \subseteq \mathrel{|\!\!\approx_b}$ when viewed as relations on formulas. As Example 3.9 demonstrates, the containment is proper.

Proposition 3.7: *For $KB, \varphi \in \mathcal{L}^{\approx}$, if $KB \mathrel{|\!\!\approx_c} \varphi$ then $KB \mathrel{|\!\!\approx_b} \varphi$.*

We begin by showing a simple example of entailment.

Example 3.8: The consistent height knowledge base KB from the introduction, written formally in our language, is

$$(c_B \approx 1.8) \wedge (c_J \approx c_B + c_H/2) \wedge (c_H \approx 0.2) .$$

If we interpret $\approx$ as $=$, we can deduce that c_J, John's height, is 1.9 meters. As we might hope, using both cautious and bold entailment, KB entails that $c_J \approx 1.9$. Since $\mathrel{|\!\!\approx_c} \subseteq \mathrel{|\!\!\approx_b}$, it suffices to show this for cautious entailment. We proceed as follows: Since KB is consistent with equality, the only minimal tolerance function for KB is τ_0. Let τ_1, τ_2, τ_3 be the relevant coordinates of the tolerance function τ for KB. We choose $(f(\tau))(c_J \approx 1.9) = \tau_1 + \tau_2 + \tau_3/2$, and $f(\tau)(e) = 0$ for all other expressions $e \in \mathcal{E}$. Clearly, if $\lim_{n\rightarrow\infty} \tau^n = \tau_0$, then $\lim_{n\rightarrow\infty} f(\tau^n) = \tau_0$. Moreover, for any valuation v, if $|v(c_B) - 1.8| \leq \tau_1$, $|v(c_J) - v(c_B) - v(c_H)/2| \leq \tau_2$, and $|v(c_H) - 0.2| \leq \tau_3$, then $|v(c_J) - 1.9| \leq \tau_1 + \tau_3/2 + \tau_2 = f(\tau)(c_J \approx 1.9)$. Thus, if $(v, \tau) \models KB$, then $(v, f(\tau)) \models c_J \approx 1.9$. It follows that $KB \mathrel{|\!\!\approx_c} c_J \approx 1.9$, as desired. ∎

The following example helps explain the difference between bold and cautious entailment. Intuitively, cautious entailment allows no unjustified default assumptions about the relationships between the tolerances in the knowledge base, whereas bold entailment allows arbitrary assumptions about these relationships.

Example 3.9: Consider the the knowledge base KB' from Example 3.2. This knowledge base is clearly inconsistent with equality, so using true equality we can deduce anything. What can we deduce using approximate entailment? Recall from Example 3.2 that

$$\Omega(KB') = \{\tau : \tau_4 = 0.02 - \tau_1 - \tau_2 - \tau_3/2\} .$$

It is easy to see that for any model (v, τ) consistent with KB' such that $\tau \in \Omega(KB')$, we must have $1.88 \leq v(c_J) \leq 1.9$; moreover, every value in the range $[1.88, 1.9]$ is attained in one of these models. It follows that $KB' \mathrel{|\!\!\approx_c} 1.88 \preceq c_J \preceq 1.9$, and that we cannot get better bounds on c_J using cautious entailment. Thus, given contradictory information as to John's height, the cautious approach allows us to deduce only that the value of c_J is somewhere in the interval.

On the other hand, using the function f assigning $(f(\tau))(c_J \approx h) = \tau_4 - (0.02 - \tau_1 - \tau_2 - \tau_3/2)$ for any $1.88 \leq h \leq 1.9$ and zero elsewhere, we can deduce $KB' \mathrel{|\!\!\approx_b} c_J \approx 1.88$, $KB' \mathrel{|\!\!\approx_b} c_J \approx 1.89$, and in general, $KB' \mathrel{|\!\!\approx_b} c_J \approx h$ for every $h \in [1.88, 1.9]$. Thus, the bold approach allows us to deduce any value for John's height in the permissible range; we may use any reasonable working assumption for John's height. ∎

The two approximate entailment relations are defined for a particular language $\mathcal{L}$ and a semantics for it. Clearly the definitions make perfect sense for a far richer language, for example, one with first-order quantification and with interactions between tolerances. (We remark that the decision procedure in the next section also holds for this extended language, although our characterizations in Section 5 do not.) More interestingly, these relations can be extended to other semantics and other notions of satisfaction. The first clause in both of the approximate entailment definitions is based on the standard notion of satisfaction for precise formulas—$KB[\tau] \models \varphi[f(\tau)]$. We can replace the symbol $\models$ is this statement by a nonmonotonic notion, for example, or a probabilistic notion such as "holds with probability 1." This extension is explored further in Section 6.

4 A Decision Procedure for Approximate Entailment

In this section, we present decision procedures for the problems of deciding whether $KB \mathrel{|\!\!\approx_c} \varphi$ and whether $KB \mathrel{|\!\!\approx_b} \varphi$. Our decision procedures will be based on reducing these questions to the validity of certain formulas over the reals. We need first present the definition of a real closed field (see, for example, [Sho67]).

Definition 4.1: An *ordered field* is a field with a linear ordering $<$, where the field operations $+$ and $\times$ respect the ordering: that is, $x < y$ implies $x + z < y + z$, and if x, y are *positive* (where an element x is positive if $x > 0$) then so is $x \times y$. A *real closed field* is an ordered field where every positive element has a square root and every polynomial of odd degree has a root. ∎

Tarski [Tar51] showed that the theory of real closed fields coincides with the theory of the reals (under formulas containing only $+, \times, <, =$ and constants $0, 1, -1$). He also proved that the theory is decidable. Ben-Or, Kozen, and Reif [BKR86] extended this result to show that the complexity of the decision problem is exponential space.

When defining $\mathcal{L}^{\approx}$, we allowed a constant d_r for every real number r. Clearly, we cannot extend the decision procedure to formulas containing such constants. We therefore define a formula φ to be *rational* if for every constant d_r mentioned in φ, r is a rational number.

Theorem 4.2: *Given rational formulas $\varphi, KB \in \mathcal{L}^{\approx}$, we can in polynomial time construct formulas ψ_b and ψ_c over the vocabulary $\{0, 1, +, \times, <\}$ whose length[2] is linear in that of φ and KB, such that*

- *$KB \mathrel{|\!\!\approx_c} \varphi$ iff $\langle \mathbb{R}, 0, 1, +, \times \rangle \models \psi_c$,*
- *$KB \mathrel{|\!\!\approx_b} \varphi$ iff $\langle \mathbb{R}, 0, 1, +, \times \rangle \models \psi_b$.*

From the results of [Tar51] and [BKR86] mentioned above, we immediately get:

Corollary 4.3: *For KB, φ rational formulas, the problem of deciding whether $KB \mathrel{|\!\!\approx_c} \varphi$ (resp. $KB \mathrel{|\!\!\approx_b} \varphi$) is in exponential space.*

5 Properties of Entailment

Although we have a decision procedure for approximate entailment, it does not give us much insight into the properties of these relations. In this section, we explore these properties in greater depth. We begin by showing the connection between approximate entailment and standard entailment in $\mathcal{L}$ (the logic of true equality). This gives a complete characterization for a large fragment of our language. We then use this characterization in order to relate approximate entailment to standard nonmononotonic formalisms.

[2] The length of a rational formula φ is defined as the length of φ written as a string of symbols, where the length of d_q, where $q = a/b$ and a and b are integers, is the sum of the lengths of the binary representations of a and b.

5.1 Characterization

What kind of inferences can we make using our two notions of entailment? Most importantly, the inferences we make are always consistent with equality: they are always a subset of those we would obtain if we were to treat approximate equality as true equality. This is an important property; if it is consistent to interpret approximate equality as equality, then we do not want to conclude anything that would be inconsistent with this intepretation. One might also hope that for consistent knowledge bases, the converse also holds. This is in fact the case for bold entailment, and under certain conditions also for cautious entailment. Thus, we get all and only "reasonable" conclusions from knowledge bases consistent with equality. If this were the whole story, then there would be no need to introduce approximate equality at all. However, as we have already observed, some of the most interesting applications of approximate reasoning arise precisely when the knowledge base is inconsistent with equality. In this case, we do not want to be able to infer everything (as we could if we did not view equality as approximate). We will see, in fact, that the inferences that are lost are the "undesirable" ones.

Our goal is to characterize what we can infer in general. Roughly speaking, we want to show that a knowledge base KB cautiously entails φ if φ is true for *every* minimal tolerance function for KB; on the other hand, KB boldly entails φ if φ is true for *some* minimal tolerance function.

In order to relate entailment to the truth of formulas in $\mathcal{L}$, we first need a result that shows that entailment is, in some sense, continuous in the tolerance function.

Proposition 5.1: *Let KB and φ be two sentences in $\mathcal{L}^{\approx}$. Let τ^n be a decreasing sequence of tolerance functions such that $\lim_{n\to\infty} \tau^n = \eta$, and let f be a function such that $\lim_{n\to\infty} f(\tau^n) = \tau_0$. If, for every n, $KB[\tau^n] \models \varphi[f(\tau^n)]$, then $KB[\eta] \models \varphi[\tau_0]$.*

We can now give half of our desired characterization.

Theorem 5.2: *Suppose $\varphi, KB \in \mathcal{L}^{\approx}$.*

- *If $KB \mathrel{|\!\!\approx_c} \varphi$ then for all $\eta \in \Omega(KB)$, we have $KB[\eta] \models \varphi[\tau_0]$.*
- *If $KB \mathrel{|\!\!\approx_b} \varphi$ then for some $\eta \in \Omega(KB)$, we have $KB[\eta] \models \varphi[\tau_0]$.*

Note that if $KB[\tau_0]$ is consistent (i.e., the knowledge base is consistent with equality), then $\Omega(KB) = \{\tau_0\}$. Thus, as a corollary to Theorem 5.2, we get that our inferences are always consistent with equality:

Corollary 5.3: *If $KB \mathrel{|\!\!\approx_b} \varphi$ (resp., $KB \mathrel{|\!\!\approx_c} \varphi$) then*

$KB[\tau_0] \models \varphi[\tau_0]$.

This result is important, because it says that we cannot conclude anything (using either notion of entailment) that we could not have concluded by viewing approximate equality as equality.

The converses to Theorem 5.2 and Corollary 5.3 do not hold in general, nor do we want them to. If $KB[\eta]$ is inconsistent, then $KB[\eta] \models \varphi[\tau_0]$ for all φ, but we do not want to conclude $KB \mathrel{|\!\!\approx_b} \varphi$ for all φ. In the case of bold entailment, such inconsistencies are the only problem.

Definition 5.4: We say KB is *min-consistent* if $KB[\eta]$ is consistent for all $\eta \in \Omega(KB)$. ∎

The knowledge-base in Example 3.3 is not min-consistent.

Theorem 5.5 : *If KB is min-consistent, then $KB \mathrel{|\!\!\approx_b} \varphi$ iff $KB[\eta] \models \varphi[\tau_0]$ for some $\eta \in \Omega(KB)$.*[3]

This result characterizes $\mathrel{|\!\!\approx_b}$ for min-consistent knowledge bases KB, and provides some justification for our calling this entailment "bold." It also lets us prove that, as long as our knowledge base is consistent with equality, then bold entailment lets us conclude precisely what we can conclude by viewing approximate equality as equality.

Corollary 5.6: *If $KB[\tau_0]$ is consistent, then $KB[\tau_0] \models \varphi[\tau_0]$ iff $KB \mathrel{|\!\!\approx_b} \varphi$.*

We can prove results analogous to Theorem 5.5 and Corollary 5.6 for cautious entailment. However, we must place some restrictions on formulas, as the following examples show.

The first example shows that we cannot deal with formulas that use precise equality.

Example 5.7: Let KB be $c \approx 0$ and φ be $c = 0$. The knowledge base is consistent with equality; therefore $\Omega(KB) = \{\tau_0\}$. From $KB[\tau_0]$ we can infer $\varphi[\tau_0]$ (the two are, in fact, equivalent). However, it is not the case, nor do we want it to be, that $KB \mathrel{|\!\!\approx_c} \varphi$. We point out that for the bold logic it is the case that $KB \mathrel{|\!\!\approx_b} \varphi$.[4] ∎

The second example shows that we cannot deal with formulas that mention division either. Here the problem is the possibility of division by zero.

Example 5.8: Let KB be $c \approx 0$ and φ be $\neg(1/c \succeq 0)$. Clearly KB is consistent with equality. Moreover, it is easy to see that $c = 0 \models \neg(1/c \geq 0)$: according to our semantics $1/0 = \iota$ and $\iota \geq 0$ is false. However, it is also easy to see that KB does not cautiously entail φ: for all τ with $\tau(c \approx 0) > 0$, there is no choice of τ' such that $(c \approx 0)[\tau] \models (\neg(1/c \succeq 0))[\tau']$.[5] ∎

Therefore, in order to obtain the desired results, we cannot allow precise equality and division in KB or φ. But, as the following three examples show, further restrictions are necessary as well.

Example 5.9: Let KB be $((c \approx 1) \wedge (c+1 \not\approx 2)) \vee (c \approx 0)$. Clearly, $\Omega(KB) = \{\tau_0\}$. From $KB[\tau_0]$ we can infer $c = 0$, and thus KB boldly entails $c \approx 0$. We might also hope that KB cautiously entails $c \approx 0$, but this is not the case. Let $\tau_1 = \tau(c \approx 1)$ and $\tau_2 = \tau(c+1 \approx 2)$. For any τ such that $\tau_1 > \tau_2$, we can easily see that there are models $(v, \tau) \models KB$ such that $v(c)$ is within τ_1 of 1. Thus, $KB \not\mathrel{|\!\!\approx_c} c \approx 0$. All we can deduce using cautious entailment is $KB \mathrel{|\!\!\approx_c} (c \approx 1) \vee (c \approx 0)$. Note that this is, in fact, a reasonable conclusion if we are being cautious in assuming relationships between different tolerances. ∎

Example 5.10: It is easy to see that $c \approx 0 \mathrel{|\!\!\approx_c} 2c \approx 0$; we simply define the function f in the definition of cautious entailment so that $f(\tau) = 2\tau$. Similarly, $c \approx 0 \mathrel{|\!\!\approx_c} d_r \times c \approx 0$ for any r. However, if c' is another uninterpreted constant in the language, then $c \approx 0 \not\mathrel{|\!\!\approx_c} c' \times c \approx 0$ (although $\Omega(c \approx 0) = \{\tau_0\}$ and $c = 0 \models c' \times c = 0$). The reason is that for every tolerance function τ such that $\tau(c \approx 0) > 0$ and every constant $B > 0$, there is some model (v, τ) such that $v(c' \times c) > B$. We cannot place an *a priori* bound on the tolerance required for $c' \times c \approx 0$ in terms of $\tau(c \approx 0)$. ∎

Example 5.11: Let KB be $(c \approx 0) \vee (c \approx 1)$ and φ be

$$\begin{array}{ll} [(c \approx 0) \wedge ((d \not\approx 0) \vee (d \approx c(c-1)))] & \vee \\ [(c \approx 1) \wedge ((d \approx 0) \vee (d \not\approx c(c-1)))] & . \end{array}$$

Under true equality, the second conjunct in each disjunct of φ is implied by the first, so that $KB[\tau_0] \models \varphi[\tau_0]$. Under approximate equality, on the other hand, the situation is very different: As we show in the full paper, $KB \not\mathrel{|\!\!\approx_c} \varphi$. To see why at an intuitive level, recall that the first clause in the definition of cautious entailment requires us to find a function f such that $KB[\tau] \models \varphi[f(\tau)]$, for every τ small enough. For any τ small enough, we can easily find v and v' such that $(v, \tau) \models c \approx 0$ and $(v', \tau) \models c \approx 1$. We must therefore have $(v, f(\tau)) \models (d \not\approx 0) \vee (d \approx c(c-1))$ and $(v', f(\tau)) \models (d \approx 0) \vee (d \not\approx c(c-1))$. However, we cannot define $f(\tau)(d \approx 0)$ and $f(\tau)(d \approx c(c-1))$ so that both implications hold. Essentially, the problem is that the two expressions $d \approx 0$ and $d \approx c(c-1)$

[3] If we were to require that the range of a tolerance function is $(0, \infty)$, so that τ_0 is not a a legal tolerance function, this theorem would not hold. Additional assumptions, similar to the well-behavedness assumptions below, would be necessary.

[4] In fact, it is easy to see that if KB is min-consistent, and $KB \mathrel{|\!\!\approx_b} \varphi$, then $KB \mathrel{|\!\!\approx_b} \varphi[\tau_0]$.

[5] This problem might seem to be an artifact of our particular semantics for division by zero. However, similar examples can be constructed for other choices.

appear both negated and unnegated in φ, inducing interactions between the two tolerances. ∎

These three examples essentially characterize the reasons why we do not get an analogue to Theorem 5.5 for cautious entailment, even for the restricted language. To make this precise, we need some definitions.

Definition 5.12: The constant c is said to be *bounded by KB* if KB implies $d_r \preceq c \preceq d_{r'}$ for some constants r and r'. ∎

Notice that the constant c' is not bounded by the KB $c \approx 0$ in Example 5.10.

Definition 5.13: Let φ be a formula in $\mathcal{L}^{\approx}$. We say that an expression $e \in \mathcal{A}$ *appears positively* (resp., *appears negatively*) in φ if there is an instance of e which is in the scope of an even (resp. odd) number of negations. We say that φ is *strictly independent* if there is no expression e that appears both positively and negatively in φ. ∎

Note that the formula φ in Example 5.11 is not strictly independent. Strict independence might seem, at first glance, to be a harsh restriction. But this is not the case. Consider, for example, $(c \approx 0) \wedge (c \not\approx 0)$. This formula is not strictly independent, but the almost identical formula $(c \approx 0) \wedge (c + 0 \not\approx 0)$ is. In general, it is simple to transform any φ to a strictly independent formula φ', that, apart from interactions between different tolerances, is equivalent.

Definition 5.14: The pair KB, φ is said to be *well-behaved* if

- neither KB nor φ contain any division operations or precise equality expressions,
- KB is negation free,
- all the constants in φ and KB are bounded by KB,
- φ is strictly independent. ∎

How reasonable is the assumption of well-behavedness? Since we are mainly interested in knowledge bases with approximate information, not allowing precise equality in this context does not seem unduly restrictive. Not allowing division is, of course, a nontrivial restriction, but still seems to cover many interesting examples. As we have seen, strict independence is a very mild restriction. Although for general knowledge bases we want negations, in this case we are reasoning only about quantitative relations among measured quantities and numerical constants. It seems reasonable to expect that the information we have in such a knowledge base would be positive. Finally, although the assumption that all constants that appear be bounded may seem restrictive, note that the bounds can be arbitrarily large. In practice, we often do have *some* bounds on the size of constants, perhaps not very precise. For example, if we are talking about heights of people, we surely have a lower bound of 0 and an upper bound of 3 meters. It should not hurt to add such bounds to the knowledge base. We therefore believe that in practice, knowledge bases will often be well-behaved (or can easily be made so). As we now show, well-behavedness is sufficient to guarantee that we avoid the problems in the examples above, as well as other difficulties.

Assuming well-behavedness, we can now prove that if an assertion holds at some minimal point, then it also holds for any sequence tending to that point.

Proposition 5.15: *Let KB, φ be well-behaved, and let η be a tolerance function such that $KB[\eta]$ is satisfiable. If $KB[\eta] \models \varphi[\tau_0]$, then there exists a function f and some $\epsilon > 0$ such that for all τ such that $\|\tau - \eta\| < \epsilon$, we have $KB[\tau] \models \varphi[f(\tau)]$, and for all sequences τ^n such that $\lim_{n\to\infty} \tau^n = \eta$, we have $\lim_{n\to\infty} f(\tau^n) = \tau_0$.*

We can now prove the following analogues to Theorem 5.5 and Corollary 5.6.

Theorem 5.16: *If the pair KB, φ is well-behaved, then $KB \mid\!\approx_c \varphi$ iff $KB[\eta] \models \varphi[\tau_0]$ for all $\eta \in \Omega(KB)$.*

Corollary 5.17: *If KB, φ are well-behaved and $KB[\tau_0]$ is consistent, then $KB \mid\!\approx_c \varphi$ iff $KB[\tau_0] \models \varphi[\tau_0]$.*

We can also show that the assumptions of Theorem 5.16 are stronger than those of Theorem 5.5:

Lemma 5.18: *If KB is bounded and negation-free then KB is min-consistent.*

We do not have an elegant characterization of bold entailment in the case where KB is not min-consistent, nor for cautious entailment in the case where φ and KB are not well-behaved. As our various examples show, we still get reasonable entailments even when these conditions are not met.

5.2 Preference semantics

Our characterization theorems emphasize the importance of minimal tolerance functions. Tolerance functions consistent with a knowledge base KB are possible combinations of measurement errors that could have led to the formation of KB. Since we prefer to believe that the errors made were as small as possible, we can view the ordering $<$ on tolerance functions as defining a preference relation on tolerance functions, in the spirit of [Sho87]. Therefore, for the fragment of our language for which the characterization theorems (Theorems 5.5 and 5.16) hold, approximate entailment reduces to reasoning in the preferred models. Using these results, we can now show that for well-behaved formulas, cautious entailment is closely related to Shoham's notion of *preferential entailment* [Sho87].

Definition 5.19: For two augmented models $M = (v, \tau)$ and $M' = (v', \tau')$, $M < M'$ if $\tau < \tau'$. The augmented model M is a *preferred model* of KB if $M \models KB$ and there is no other augmented model M' such that $M' < M$ and $M' \models KB$. KB *preferentially entails* φ, written $KB \models_< \varphi$, if for any preferred model M of KB, $M \models \varphi[\tau_0]$.[6] ∎

Theorem 5.20: *For well-behaved KB, φ the following are true:*

- *An augmented model $M = (v, \eta)$ of KB is a preferred model of KB iff $\eta \in \Omega(KB)$.*
- *$KB \approx\!\!\!|_c\, \varphi$ iff $KB \models_< \varphi$.*

We can view minimal tolerance functions as frames of mind, and models v for $\mathcal{L}$ as possible worlds. In the frame of mind corresponding to the minimal tolerance function η, the agent believes φ if $KB[\eta] \models \varphi[\tau_0]$. In the cautious approach, the agent believes only deductions made in all frames of mind. In the bold approach, the agent believes deductions made in any frame of mind. Note, however, that the agent can believe φ in one frame of mind and ψ in another, while not believing $\varphi \wedge \psi$ in any frame of mind. Therefore, it is possible that $KB \,|\!\!\approx_b \varphi$, $KB \,|\!\!\approx_b \psi$, while $KB \,|\!\!\not\approx_b \varphi \wedge \psi$.

Using the terminology of [KLM90], what this discussion has shown is that bold entailment is not closed under the *And* rule. We have also shown that neither bold nor cautious entailment is closed under the *Reflexivity* rule: it is not necessarily the case that $KB \,|\!\!\approx_b KB$ or that $KB \,|\!\!\approx_c KB$. Our characterization theorems suggest why this should be so: formulas on the left-hand side of $|\!\!\approx_b$ or $|\!\!\approx_c$ are evaluated with respect to minimal tolerance functions; formulas on the right-hand side are evaluated with respect to τ_0. Therefore, for KB inconsistent with equality, KB will not entail itself (nor would we want it to). Several examples have demonstrated that cautious entailment is nonmonotonic: adding information to the knowledge base can cause inferences to be lost. Interestingly, bold entailment is monotonic for a large fragment of the language. Both $|\!\!\approx_b$ and $|\!\!\approx_c$ are closed under most of the other rules suggested in [KLM90] (under certain restrictions such as strict independence). We discuss this issue in more detail in the full paper.

6 Applications

6.1 Error propagation

So far, we have only looked at very simple examples, with two or three quantities of interest, and very few interactions between them. In real-world situations, there will be many different quantities, each of which will be relevant in a variety of computations. This can cause complex interactions, as shown in the following simple example.

Example 6.1: Consider a robot operating in a blocks world, whose primitive actions are: grasp (g), ungrasp (u), and move arm (m). Suppose that the robot has estimates on how long these tasks take. Such estimates are clearly useful in the context of planning. Let $c_g \approx 2$, $c_u \approx 1.5$, and $c_m \approx 5$ be the estimates for the time taken by these three actions, respectively. Assume that a particular plan r (raise block) requires one grasp and one move action, so that the robot estimate that $c_r \approx c_g + c_m$, whereas plan l (lower block), requires one move and one ungrasp, so that $c_l \approx c_m + c_u$. The robot can deduce that $KB \,|\!\!\approx_c c_r \approx 7 \wedge c_l \approx 6.5$. Now suppose that the robot executes plan r, measures the time it actually takes, and discovers that it takes 7.5 seconds. The robot would like to use this measurement as a new approximation for how long this plan usually takes. It therefore adds $c_r \approx 7.5$ to KB, obtaining KB'. The robot can now conclude that $KB' \,|\!\!\approx_b c_g \approx 2.5 \wedge c_l \approx 6.5$; this corresponds to the case that the mistaken estimate was for action g, and therefore the time for plan l is unaffected. Alternatively, the robot can conclude $KB' \,|\!\!\approx_b (c_g \approx 2) \wedge (c_m \approx 5.5) \wedge (c_l \approx 7)$, corresponding to the case that the mistaken estimate was for action m. Yet another alternative is that $KB' \,|\!\!\approx_b (c_g \approx 2) \wedge (c_m \approx 5) \wedge (c_l \approx 6.5)$, corresponding to the case that the cause of the discrepancy was overhead in plan r. Any intermediate assignment of errors is also possible. Note that each of these alternatives corresponds to a scenario that "explains" the discrepancy between the estimated time for the plan and the actual time for the plan by assuming that the estimates were as correct as possible, and making the minimal change required to account for the discrepancy. If the robot is not willing to leap to any of these conclusions, then it could use cautious entailment, and obtain ranges for the estimated time for each action and plan. ∎

We see that the process of considering the different causes for the discrepancy, and deducing from those how the times for different plans could be affected, is done automatically by approximate entailment. Essentially, errors are propagated back to their possible sources, and then forward to their logical conclusions. This type of reasoning is useful in many other applications. For example, it arises in complex numerical computations, where each subroutine can introduce errors (such as rounding errors), which then propagate in many ways, affecting more than one result.

6.2 Measurement theory

The problems of inexact measurement and numerical inaccuracies have been extensively investigated in the field of measurement theory. While there are many

[6] Shoham's notation for M preferred to M' is $M' < M$. We choose to represent preference as $M < M'$ in order to maintain consistency with the notation for tolerance functions. Moreover, this is not quite Shoham's definition of preferential entailment. The exact analogue of Shoham's definition would have $M \models \varphi$, not $M \models \varphi[\tau_0]$.

points of commonality between our approach and measurement theory, there are also some significant differences. Measurement theory investigates the issue from an axiomatic standpoint. Their measurement data typically contains relative observations about the objects being measured. For example, if we are measuring the heights of people, we may observe that John is taller than Bill. The general theory attempts to find axioms guaranteeing that numbers that "satisfy" the observations in an appropriate sense can be assigned to the objects. Thus, they start from axioms, rather than models, as we do.

When dealing with inexact measurement, the problems encountered typically involve *intransitivities*. Consider the following classic example [LR57]:

Example 6.2: Let c_n denote a standard cup of coffee that contains n granules of sugar. The agent cannot differentiate between c_n and c_{n+1} by taste; therefore, its KB will contain $c_n \approx c_{n+1}$ for every n. However, for some m, the agent will be able to tell that c_m is sweeter than c_1, so KB will contain $c_1 \prec c_m$. ∎

Intuitively, from the point of view of measurement theory, this problem arises because each measurable quantity c_i has a "true value," and an interval around it that the agent cannot differentiate from the true value. It is shown [SKLT89] that if the $\prec$ relation satisfies certain axioms, then we can find an assignment v and a *threshold function* δ such that taking $c_i \approx c_j$ iff $v(c_i) \in [v(c_j) - \delta(c_j), v(c_j) + \delta(c_j)]$ is consistent with all the observations. Such a pair (v, δ) is called a *threshold representation*.

This shows another key difference between our approach and measurement theory. Measurement theory associates the "tolerance" δ with a quantity such as c_i, whereas we associate a tolerance with a measurement such as $c_i \approx c_j$, viewing each measurement as having its own (independent) uncertainty. However, we can capture the measurement theory notion of tolerance in our approach if we put additional constraints on our tolerance function τ. In particular, we could require for each variable c_i, the tolerances of all comparisons involving c_i are the same; i.e., for all j and k, we could require that $\tau(c_k \approx c_i) = \tau(c_j \approx c_i)$. Thus, a threshold representation corresponds to an augmented model with some additional constraints on the tolerance function, so axioms that guarantee the existence of a threshold representation also guarantee the existence of such an augmented model.

6.3 Probabilistic entailment

Another important type of numerical information is probabilistic knowledge. An agent may frequently use probabilities to deal with its uncertainty about the truth of various sentences. For example, in [Nil86], Nilsson suggests a framework for probabilistic logic, which, for the case of propositional logic, is essentially as follows.

Consider a finite propositional language over the propositions $p_1, \ldots, p_k$. There are $K = 2^k$ truth assignments, or worlds, for this language; let us denote them by $w_1, \ldots, w_K$. Given a probability distribution π over the K worlds, we define the probability of a propositional sentence α to be

$$\pi(\alpha) = \sum_{w_i \models \alpha} \pi(w_i) \ . \qquad (1)$$

The agent's knowledge base consists of a set of constraints on the probabilities of different sentences; for example, $\pi((p_1 \vee p_3) \wedge \neg p_6) \leq 0.4$. Using *probabilistic entailment*, the agent deduces constraints on the probability of a sentence α from the probabilistic constraints in the knowledge base.

In our framework, we define the constant c_i to be the probability $\pi(w_i)$. Any constraint on the probability of sentences can be replaced by a constraint referring only to the constants c_i, using Equation (1). Thus, the problem of probabilistic entailment can be expressed in $\mathcal{L}$. But what happens if the numbers appearing in the probabilistic constraints are not known to be precise? As Nilsson points out, "just as it is possible to assign inconsistent *true-false* truth values to sentences, it is also possible to assign them inconsistent probabilities." He suggests the heuristic, based on a geometric interpretation, of moving to a "nearby" consistent probability distribution. This is, in fact, the effect of using approximate equality (instead of true equality) to represent the constraints, and using our framework for approximate entailment.

The same framework for assigning probabilities to propositional sentences is also used in Pearl's ϵ-semantics [Pea88]. The goal of ϵ-semantics is to provide probabilistic semantics for default reasoning—defaults of the type "birds fly" are interpreted as meaning "almost all birds fly," and are given semantics using the conditional probability of the "fly" given "bird" (where "fly" and "bird" are propositions in the language). However, the statement $\pi(\mathit{fly} \wedge \mathit{bird})/\pi(\mathit{bird}) = 1$ does not accurately represent the meaning of the sentence "almost all birds fly," since it is inconsistent with the existence of non-flying birds. Pearl's solution to the problem is essentially equivalent to representing that sentence as $\pi(\mathit{fly} \wedge \mathit{bird})/\pi(\mathit{bird}) \approx 1$ (using the transformation of π into constants c_i as described above).[7] Using this representation, Pearl defines a notion of *ϵ-entailment*, written $\models_\epsilon$. As we now show, our notion of cautious entailment agrees with ϵ-entailment, except when the knowledge base is declared inconsistent by the ϵ-semantics approach. In this case, cautious entailment can tolerate the inconsistency and still provide reasonable answers.

[7] We cannot represent this formula as $\pi(\mathit{fly} \wedge \mathit{bird}) \approx \pi(\mathit{bird})$. Assume, for example, that we also knew that most things are not birds, that is, $\pi(\mathit{bird}) \approx 0$. Then, for any τ such that $\tau(\pi(\mathit{bird}) \approx 0) \leq \tau(\pi(\mathit{fly} \wedge \mathit{bird}) \approx \pi(\mathit{bird}))$, the assertion $\pi(\mathit{fly} \wedge \mathit{bird}) \approx \pi(\mathit{bird})$ would be trivially true. This is not the case for the original assertion above.

Theorem 6.3: *Let Δ be a set of defaults and α a sentence in the language of ϵ-semantics. Let KB and φ be their respective analogues in $\mathcal{L}^{\approx}$, using the transformation described above. Then if $\Omega(KB) = \{\tau_0\}$, then $\Delta \models_\epsilon \alpha$ iff $KB \mathrel{|\!\approx_c} \varphi$.*

Note that this equivalence holds only if the knowledge base is consistent arbitrarily close to τ_0. If this is not the case, then ϵ-semantics would declare the knowledge base to be inconsistent, allowing arbitrary deductions. Consider, for example, the following variant (from [Pea88]) of the famous *lottery paradox* [Kyb61].

Example 6.4: A large number of people buy tickets for a lottery that will have a single winner. Let c_i represent the probability that person i will win. The probability that any one person will win is very low. Therefore, we might choose to represent our knowledge in ϵ-semantics using the statement $c_i \approx 0$. However, we know that one person will certainly win, leading to the conclusion $c_1 + c_2 + \ldots + c_N = 1$, where N is the number of people participating in the lottery. The resulting knowledge base is inconsistent in the ϵ-semantics framework. However, using our approach, this inconsistency is avoided, and the knowledge base can be used for making inferences. ∎

Both Nilsson and Pearl have suggested specific nonmonotonic variants of their basic logics. Our formalism easily extends to encompass these proposals. Both proposals use the same basic idea: instead of looking at all probability distributions over worlds consistent with the constraints, one should look at one particular "special" one—the probability distribution satisfying the constraints that has *maximum entropy* (see [GMP90] for more details). This leads to nonmonotonic notions of inferences: we leap to the conclusions sanctioned by the distribution of maximum entropy consistent with our information, although they may not be sanctioned by other distributions consistent with our information. We can easily extend our framework to deal with nonmonotonic inferences, and thus capture these nonmonotonic approaches. As before, let $KB, \varphi \in \mathcal{L}^{\approx}$ represent the knowledge base and desired conclusion in our framework. Recall that the first condition in the definition of cautious entailment states that, for every τ, $KB[\tau] \models \varphi[f(\tau)]$. Until now, we used $\models$ to denote the standard notion of entailment. Instead, we can replace $\models$ by $\models_{ME}$, where $\models_{ME}$ allows any inferences which hold in the model v of $KB[\tau]$ having maximum entropy (where we now view v as a probability assignment, so that talking about its entropy makes sense), leaving the remainder of the definition unchanged. The key point here is that the choice of inference rule is completely orthogonal to our treatment of approximate equality. Our approach can be applied to any notion of inference rule, to convert a logic for reasoning about equality to one for reasoning about approximate equality.

A final example of the generality of this framework uses yet another language and inference mechanism. In [GHK94], we present a technique which deduces *degrees of belief* from a first-order knowledge base augmented with statistical information about the domain. That is, given such a knowledge base KB and a formula φ, we define the notion of the degree of belief in φ given KB, denoted $\Pr_\infty(\varphi|KB)$ (see [GHK94] for details). The statistical information has the form "the proportion of flyers among birds is 90%;" this type of information is usually based on some sort of statistical sampling, and is therefore only approximate. Moreover, as pointed out in [GHK94], if we were to interpret the statistical statement above as being precisely true, we would deduce that the number of birds is a multiple of 10, an inference which is surely undesirable. Therefore, approximate equality rather than precise equality is used in [GHK94]. However, there is no analogue to (cautious or bold) entailment. Rather, the approach of [GHK94] can be viewed as using an analogue to validity (for an appropriate nonstandard notion of $\models$). Since, as we have seen, very few interesting deductions regarding approximate inference can be made using validity, not much can be deduced if we have approximate equality in both φ and KB. As an example of this phenomenon, suppose KB is "10% of birds are yellow, 20% of birds are green, and no birds are both yellow and green," and φ is "30% of birds are yellow or green." Using the approach of [GHK94], we cannot deduce $\Pr_\infty(\varphi|KB) = 1$. Moreover, this approach could not deal with inconsistent numerical information in KB. By using the approach outlined in this paper, one could define both a bold and a cautious version of $\Pr_\infty$, and deal with these issues in a satisfactory way.

7 Conclusions

We have presented a logic for approximate reasoning, and defined two notions of approximate entailment used to make default deductions from an imprecise knowledge base. One might ask why we should bother designing a new logic, rather than using say, a variant of *relevance logic* [AB75], *fuzzy logic* [Zad75], or any one of a number of nonmonotonic logics. Each of these logics has some properties that we view as desirable in our setting. The use of relevance logic would block the deduction of arbitrary formulas from an inconsistent knowledge base. Fuzzy logic would allow us to express the notion of approximate equality (although in a way that is very different from that captured by our semantics). Nonmonotonic logics allow the type of nonmonotonic behavior we want in the height knowledge base mentioned above. However, since these logics were not designed specifically to handle approximate measurement, none of them can capture all the situations in which we are interested. Threshold representations and semantics based on intervals can be used to directly express approximate quantities. In fact, Parikh's [Par83] theory of *vague reals* is a logical framework for doing this. However, as we explained in Section 6.2, approximate quantities differ from approximate measurements. Moreover, Parikh's framework requires observations to be given in terms of ranges rather than exact numbers; this is not always feasible.

Our logic has many interesting and intuitive properties. We concentrated on demonstrating these properties for a particular language (the language of arithmetic), and for the classical notion of entailment. However, we also showed variants of our logic for probabilistic and even nonmonotonic logics. We view this logic as providing a general and coherent framework for dealing with approximate information and the numerical inconsistencies that usually accompany it.

Acknowledgements

The authors thank Adam Grove for numerous helpful discussions, and Ron Fagin and Moshe Vardi for useful comments on a draft of the paper. This research was sponsored in part by the Air Force Office of Scientific Research (AFSC), under Contract F49620-91-C-0080.

References

[AB75] A. Anderson and N. D. Belnap. *Entailment: The Logic of Relevance and Necessity.* Princeton University Press, Princeton, N.J., 1975.

[BKR86] M. Ben-Or, D. Kozen, and J. H. Reif. The complexity of elementary algebra and geometry. *Journal of Computer and System Sciences*, 32(1):251–264, 1986.

[GHK94] A. J. Grove, J. Y. Halpern, and D. Koller. Random worlds and maximum entropy. *Journal of A.I. Research*, 2:33–88, 1994.

[GMP90] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. In *Proc. National Conference on Artificial Intelligence (AAAI '90)*, pages 646–652, 1990.

[KLM90] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[Kyb61] H. E. Kyburg, Jr. *Probability and the Logic of Rational Belief.* Wesleyan University Press, Middletown, Conn., 1961.

[LR57] R. D. Luce and H. Raiffa. *Games and Decisions.* Wiley, New York, 1957.

[Nil86] N. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1986.

[Par83] R. Parikh. The problem of vague predicates. In R. S. Cohen and M. W. Wartofsky, editors, *Language, Logic, and Method*, pages 241–261. D. Reidel Publishing Company, Dordrecht, Netherlands, 1983.

[Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, San Francisco, Calif., 1988.

[Sho67] J. R. Shoenfield. *Mathematical Logic.* Addison-Wesley, Reading, Mass., 1967.

[Sho87] Y. Shoham. A semantical approach to nonmonotonic logics. In *Proc. 2nd IEEE Symp. on Logic in Computer Science*, pages 275–279, 1987. Reprinted in M. L. Ginsberg (Ed.), *Readings in Nonmonotonic Reasoning*, Morgan Kaufman, San Francisco, Calif., 1987, pp. 227–250.

[SKLT89] P. Suppes, D. M. Krantz, R. D. Luce, and A. Tversky. *Foundations of Measurement*, volume 2. Academic Press, Inc., 1989.

[Tar51] A. Tarski. *A Decision Method for Elementary Algebra and Geometry.* Univ. of California Press, 2nd edition, 1951.

[Zad75] L. A. Zadeh. Fuzzy logics and approximate reasoning. *Synthese*, 30:407–428, 1975.