Fuzzy Logic & Vagueness

Phil Serchuk

Winter 2005
Honours Thesis

Department of Philosophy
University of Calgary
2500 University Drive NW
Calgary, Alberta T2N 1N4
Canada

Committee Members:
Jack MacIntosh
Nicole Wyatt
Richard Zach

Contents

1 Introduction **2**
1.1 Classical Logic and its Problems 3
1.1.1 Vagueness 4
1.1.2 What Vagueness is Not 5
1.2 Introduction to Many-valued Logics 7
1.2.1 3-valued Logic 8
1.2.2 Fuzzy Logic 9

2 The Sorites and Paradoxes of Vagueness **12**
2.1 The Sorites Paradox 13
2.2 Forms of the Sorites 16
2.3 Sorites-like Paradoxes 20
2.4 Vague Objects 23
2.4.1 Vague Objects and Gareth Evans 27
2.4.2 The Problem of the Many or the None 31

3 Vagueness in Practice **34**
3.1 Vagueness in Computer Science 36
3.2 Vagueness and the Law 37

4 Solutions to the Paradox **41**
4.1 Epistemicism 43
4.1.1 Alleged Empirical Evidence for Epistemicism 51
4.2 Fuzzy Logic 53

5 Conclusion **65**

1 Introduction

The phenomenon philosophers call ‘vagueness’ refers specifically to predicates that, in some sense, lack sharp boundaries. Predicates, for example, like ‘tall’ and ‘bald’ are considered vague because we are unable to determine the precise number of hairs that a man must have before he is ‘not bald’ or the exact height a man must be before he can be considered ‘tall’; predicates like ‘weighs 100 grams’ are, arguably, not vague because there is a precise number of grams that an object must weigh before we can say that it weighs 100 grams. Despite our inability to crisply define these words, humans have been able to regularly use these and other vague predicates without talking themselves into a contradiction. While the impact of philosophical vagueness has always been felt in other disciplines, the rise of Artificial Intelligence (AI) has cast problems of vagueness in a new light. For it was one thing to say, as ancient doctors did when they argued over the meaning of ‘medical experience’, that despite our inability to define ‘experience’ in terms of an exact number of observations we could still somehow gain ‘experience’; it is quite another to find a way of representing this intuitive fact in formal logic. As computational power increases applications in AI that were once thought impossible are suddenly being seriously researched; many of these systems are designed to interact with humans in natural language or mimic some aspect of human behaviour or intelligence. A key problem for researchers is the representation of vague predicates; as computer scientists begin to research vagueness they are beginning to find that its effects are as problematic for them as they are for philosophers. While classical logic is the built-in logic of every major programming language, an increasing number of AI researchers are turning to non-standard logics to represent vagueness and other linguistic phenomena.

In section 1, I give an overview of vagueness and distinguish it from seemingly related phenomena like uncertainty and ambiguity. I also discuss many-valued logics and give an overview of both a 3-valued and fuzzy logic. In section 2, I give a formal explanation of why the phenomenon of vagueness leads to a logical paradox and look at the different forms that paradoxes of vagueness can take. Theories of vague objects are also discussed along with a well-known argument by Gareth Evans that allegedly proves the impossibility of vague objects. In section 3, I discuss the impact of vagueness in two other disciplines, law and computer science, and outline some of the ways that these disciplines have come to deal with vagueness. I show that vagueness is not only a problem for philosophers and that a thorough account of vagueness requires that it be examined from the perspective of researchers in many other disciplines. In section 4, I discuss several philosophical methods of dealing with vagueness and focus on two theories in particular: epistemicism and fuzzy logic. Finally, I argue that while no one theory of vagueness is satisfactory, fuzzy logic is the only existing theory that can be readily implemented in the many disciplines where vagueness can be found; thus, until a philosophically satisfying theory is found, the ease with which fuzzy logic can be used in practice should be viewed as a major advantage when it is compared to competing theories of vagueness.

1.1 Classical Logic and its Problems

Classical logic adheres to the principle of bivalence, which says that for any sentence A, either A is true or A is false. This should not be confused with the law of the excluded middle, which says only that for any sentence A, $A \vee \neg A$ is true. Although the law of the excluded middle and the principle of bivalence are equivalent in classical logic, a logic that rejects bivalence does not necessarily reject the law of the excluded middle. This is because the law of the excluded middle only governs the validity of $A \vee \neg A$ and not, like bivalence, the truth values that A and $\neg A$ can take [Haa96, p. 67]. Although classical logic is subject to both bivalence and the law of the excluded middle, it seems as though there are many cases where our linguistic intuitions cannot be represented by only two truth values.

One example which has led some to question bivalence are future contingents; that is, statements about future events like 'the United States will go to war on April 9, 2013'. In 'De Interpretatione', Aristotle argues that the bivalence of future contingents entails the view that everything happens of necessity, for "it is necessary for there to be or not to be a sea battle tomorrow" [Ari00, 19a30-31] and, if 'there will be a sea battle tomorrow' is true now, then it "will not happen as chance has it" [Ari00, 18b24], but because it must occur. While Aristotle left the issue of future contingents unresolved, Łukasiewicz believed the problem could be solved by holding that future contingents may be neither true nor false and that they must possess a third truth value, 'indeterminate'. This truth value "represents 'the possible' and joins 'the true' and 'the false' as a third value" [Luk70, p. 53]. Self-referential sentences like 'this sentence is false' pose a similar problem: if the sentence is true then it must be false and if it is false then it must be true. It has been suggested that bivalence must be rejected in these cases as well so that the truth value ‘indeterminate’ can be used to represent these sentences; 3-valued logics that have been developed by Kleene and Łukasiewicz are examined in section 1.2.1.

1.1.1 Vagueness

Classical logic is generally seen as being unable to deal with the phenomenon of vague predicates. Vagueness fascinates so many people because it can be found in more than just logical or philosophical endeavors: its effects can be seen in everyday life as competent English speakers are unable to find boundaries to vague predicates they regularly use. A logician knows no more than a doctor or a cook when asked how many hairs are required to be on a man’s head for him to be considered bald. Everyone can agree that Willard, with one hair on his head, is bald and that Timothy, with 10,000 hairs, is not. But there will be little agreement over exactly how many hairs must be plucked from Timothy’s head until he may be considered bald. They may agree that plucking 9,900 hairs would be enough and that plucking 100 hairs would not, but they would be unable to find the minimum number of hairs that must be removed for Timothy to count as being bald. This is because for any number n that is under consideration, the logician can simply say “I agree that removing n hairs would be sufficient, but surely one hair does not make the difference between baldness and non-baldness. So, for Timothy’s sake, we should allow him to keep that extra strand of hair”. It is intuitively absurd to believe that anyone would identify a man with 1,000 hairs as being bald and a man standing next to him with 1,001 hairs as being non-bald, and it seems completely reasonable to assert that one hair cannot make the difference between baldness and non-baldness. So the group agrees with the logician’s claim and agrees to lower the number of hairs that must be removed by one, to $n-1$. Having agreed to two seemingly obvious premises, the group has unwittingly allowed a paradox to take hold. This is because the logician’s line of reasoning can be repeated over and over again until the group concludes that no hair would need to be removed for Timothy to qualify as being bald. A more detailed examination of the sorites paradox is given in sections 2.1 and 2.2.

These types of problems can be found throughout human discourse. It seems unknowable, if not outright indeterminate, whether or not a particular molecule is a part of Mt. Everest or a part of the ground next to Mt.

Everest; what times one can arrive at if instructed to meet 'around noon'; when 'a few' apples become 'many' apples; when a 'large' debt becomes a 'gigantic' debt. Arguably, even the identities of some electrons under certain circumstances is indeterminate. Vagueness can also infect quantifiers: although quantifiers like 'all' and 'none' are crisp in both natural language and first-order logic, some determiners may have a different connotative meaning in natural language than they would in first-order logic. Quantifiers like 'a few', 'several', and 'many' appear to be vague even though they are generally defined in first-order logic as being synonymous with 'more than one' [Phi99, p. 182]. Take the sentence 'many sophists studied Plato' where we define $\mathbf{X}$ as the set of all sophists, $\mathbf{Y}$ as the set of all those who have studied Plato, n as a constant integer, and p as a constant percentage. There are three possible definitions of 'many' [PtMW90, pp. 395–398], each of which yields a different sentence: $\|\mathbf{X} \cap \mathbf{Y}\| \geq n$, which would be equivalent to 'more than n sophists have studied Plato'; $(\|\mathbf{X} \cap \mathbf{Y}\| \div \|\mathbf{X}\|) \geq p$, which would be equivalent to 'more than p percent of sophists have studied Plato'; and $(\|\mathbf{X} \cap \mathbf{Y}\| \div \|\mathbf{X}\|) \geq (\|\mathbf{Y}\| \div \|\mathbf{D}\|)$, which would be equivalent to 'the percentage of sophists that studied Plato is greater than the percentage of the general population that has studied Plato'. While the different definitions are an example of ambiguity, the quantifiers themselves are vague and it is the vagueness of the quantifiers that is responsible for the ambiguity; if these predicates were crisply defined there would be no room for competing definitions.

1.1.2 What Vagueness is Not

It is important that vagueness be distinguished from other, arguably similar, linguistic phenomena. Ambiguity can be found throughout natural language. For example, a sentence of English which is formalized in first-order logic as $(A \vee \neg B) \wedge C$ has different truth conditions than a sentence formalized as $A \vee (\neg B \wedge C)$. Yet humans do not use brackets when speaking in natural language and we rely on other linguistic cues to determine the actual meaning of a sentence. A single word is also capable of having multiple meanings and the correct meaning may be determined, in whole or in part, by its pronunciation, spelling, or the context in which it is placed. The pronunciation of the Japanese word 'GOKAI' can allow it to mean 'misunderstanding', 'Go tournament', or 'five times' [Yok94, p. 142]. 'They're' and 'their' sound identical but have different meanings while 'there' sounds identical as well but has several context-dependent meanings. At the turn of the 20th century the term ‘horseless carriage’ was considered to be equivalent to ‘automobile’ yet nobody today would consider automobiles to be the equivalent of horseless carriages [Sow00, p. 351]; ‘cool’, ‘gay’, and ‘hip’ are all examples of words whose meanings have started to change or expand and may one day be considered inequivalent to their original meanings. ‘Light’ and ‘mad’ are examples of words whose ambiguity is more entrenched in the English language. In both cases, the meaning of the word must be inferred from the context in which it is used. A piece of furniture described as ‘light’ could be any colour but would probably weigh very little, while a photograph described as ‘light’ is most likely having its brightness described and not its weight. Ambiguity is seen as a separate phenomenon from vagueness because ambiguous statements can be resolved with additional information and clarifications, while vague sentences cannot. Not all authors share this view: in the introduction to section 4, I briefly examine supervaluations, a popular method of dealing with vagueness that is motivated in part by alleged similarities between ambiguity and vagueness [Fin75, p. 135].

Generalizations are frequently made in natural language that cannot be properly represented in first-order logic. Generalizations like ‘all birds can fly’ are intuitively acceptable to most humans even though penguins, newborn birds, and dead birds cannot fly [Sow00, p. 349]. Yet even after being presented with these exceptions few people would want to say that the generalization is completely false. A non-monotonic logic would be better suited to model this generalization, where we could represent the assumption that a given bird can fly in the absence of other information, such as its being a penguin. Although some generalizations only use vague predicates as less accurate versions of more specific predicates, there are some generalizations that require the presence of vague predicates. ‘Bill is elderly’ is a less accurate generalization of ‘Bill is 95 years old’, but the sentence ‘tall players usually have an advantage in basketball’ requires the vague predicate ‘tall’ and cannot be rewritten with a precise height [Wan96, p. 324]. Exceptions can be permanent, such as the exceptions to ‘all birds can fly’, or temporary, as is the case when the generalization ‘a car can drive from Calgary to Toronto’ fails because the car’s battery is dead [Sow00, p. 349]. Although it may be tempting to see these generalizations as borderline cases with borderline truth, the problem is not a feature of the language itself. ‘All birds can fly’ is problematic because it is false, not because there is anything vague about ‘all’, ‘birds’, or ‘fly’. Vagueness also differs from relativity: both an Eskimo and a Swede may be of above average height, but the average height of Swedes is different than the average height of Eskimos [Sai88, pp. 26–27]. It is important to note that there is nothing vague about ‘above average height’ and we could make the population to which it refers explicit so that the predicate is applied relative only to that population. Of course, were we to use ‘tall’ instead of ‘above average height’, context-sensitivity might be a valid consideration and there are cases where it seems as though vague predicates are relative; as a result, context-sensitivity is an important consideration in any study of vagueness.

Vagueness should also be distinguished from the phenomenon of uncertainty. ‘Uncertainty’ is a broad term that encompasses a wide range of unknown, imprecise, partial and inconsistent data. A knowledge base in Prolog represents information that it has no knowledge about as being false by default. This is different than truth-functional falsehood: if the knowledge base cannot derive the truth of P then it will report not only that P is false, but that $\neg P$ is false as well. Imprecision is another problem that plagues many knowledge bases and there are many potential sources of noisy data: information can be sensed incorrectly by a faulty sensor, improperly transmitted or interpreted as a result of human or computer error, or rounded off to a less precise value. Inconsistency can arise as a result of conflicting default information, conflicting recommendations, or inconsistent data that arrives in the system through user input or the merging of knowledge bases. Default values can conflict when incomplete generalizations are used in the reasoning process: a rule may state that Quakers are usually pacifists and Republicans are usually not, but without additional data on Republican Quakers there is no way for the system to determine which ‘usually’ to apply in a given case [Sow00, p. 367]. Conflicting recommendations can arise when several equally viable options are recommended for completing a single task. While these phenomena are similar to vagueness from the perspective of a computer scientist who must determine how this knowledge is to be represented, they are clearly distinct from vagueness and there is no reason to assume that a philosophical treatment of vagueness could successfully handle these phenomena.

1.2 Introduction to Many-valued Logics

Many-valued logics were originally developed to deal with some of the shortcomings of classical logic: Łukasiewicz tried to deal with the problem of future contingents by developing a 3-valued logic that allowed a third truth value, 'indeterminate', to be assigned to statements that did not seem to be either true or false. Since then, many-valued logics have been used by philosophers, mathematicians, and computer scientists for a variety of tasks, including the study of modal, probabilistic, and tense logics, the resolution of mathematical and philosophical paradoxes, and to deal with the indeterminacy of quantum mechanics [Res69, pp. 12–15]. Many-valued logics are as of much interest to computer scientists as they are to philosophers: switching theory, programming, and entire computer architectures can be based on many-valued logics [Rin77b, p. 2] and over the past few decades they have become of great use to AI researchers. Fuzzy logic is an infinitely-valued logic that is arguably the most commonly used many-valued logic in AI research. Fuzzy logic is usually used in these systems to deal with the vagueness and uncertainty that is encountered in both simple control and intelligent decision making tasks.

1.2.1 3-valued Logic

While the principle of bivalence entails classical logic's two truth values, 'true' and 'false', many-valued logics reject this principle by introducing new truth values. Although most many-valued logics consist of 3, 4, or infinitely many values, a many-valued logic can have any number of truth values. Many-valued logics are extensions of classical logic and always have 'true' and 'false' as truth values which behave, in relation to one another, as they would in classical logic [Got01, p. 5]. Because of this, the truth table for a 2-valued logic will always be a part of the truth table for a many-valued logic. This does not mean that classical logic is a subset of many-valued logics. For example, $A \vee \neg A$ is a tautology in classical logic that is not always a tautology in many-valued logics. In classical logic the principle of bivalence holds; that is, a sentence A can only have one of two possible truth values and its negation, $\neg A$, must have the truth value that A does not. So, for every sentence in first-order logic A, if A is true then $\neg A$ is false and if A is false then $\neg A$ is true. It then follows that $A \vee \neg A$ is a tautology because A can be any sentence in first-order logic and it will always be the case that either A or $\neg A$ is true. In a many-valued logic both A and $\neg A$ can have one of at least three possible truth values and there is no guarantee that either of them will be true; for example, if A is indeterminate in a 3-valued logic then $\neg A$ may also be indeterminate, as will $A \vee \neg A$.

The rows that correspond to the truth table for classical logic have been bolded on the truth table for Kleene's 3-valued logic that is found below. Truth values are represented numerically, where 0 represents 'false', 1 represents 'true', and $\frac{1}{2}$ represents an indeterminable truth value that is assigned to that which is possible and exists between 'the true' and 'the false'; that is, $\frac{1}{2}$ is truer than what is false but falser than what is true [Luk70, Res69]. This logic differs slightly from Łukasiewicz's 3-valued logic: Łukasiewicz sets the truth of $A \rightarrow_{Ł} B$ to 1, where $A = \frac{1}{2}$ and $B = \frac{1}{2}$, but Kleene sets the truth of $A \rightarrow_K B$ to be $\frac{1}{2}$.

$P \wedge Q$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$
$\mathbf{0}$	$\mathbf{0}$	0	$\mathbf{0}$
$\frac{1}{2}$	0	$\frac{1}{2}$	$\frac{1}{2}$
$\mathbf{1}$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$

$P \vee Q$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$
$\mathbf{0}$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
$\mathbf{1}$	$\mathbf{1}$	1	$\mathbf{1}$

$\neg P$	
$\mathbf{0}$	$\mathbf{1}$
$\frac{1}{2}$	$\frac{1}{2}$
$\mathbf{1}$	$\mathbf{0}$

$P \rightarrow_K Q$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$
$\mathbf{0}$	$\mathbf{1}$	1	$\mathbf{1}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
$\mathbf{1}$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$

$P \leftrightarrow_K Q$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$
$\mathbf{0}$	$\mathbf{1}$	$\frac{1}{2}$	$\mathbf{0}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
$\mathbf{1}$	$\mathbf{0}$	$\frac{1}{2}$	$\mathbf{1}$

Table 1: Truth Tables for Kleene's 3-valued Logic

As is the case in some other many-valued logics, the '$\wedge$' operator acts as a *min* function while the '$\vee$' operator acts a *max* function. While a three-valued logic can deal with future contingents by setting their truth value to $\frac{1}{2}$, it is unable to distinguish between cases of contradictory information and ignorance. Though it is not used to deal with vagueness, a 4-valued logic is given in [Bel77] that uses the values of 'none' and 'both' to make this distinction. More detailed information about these and other many-valued logics can be found in [Got01] and [Res69].

1.2.2 Fuzzy Logic

Infinitely-valued logics are many-valued logics with an infinite number of truth values, usually represented on the interval of real numbers $[0, 1]$. Although Łukasiewicz had worked on infinitely-valued logics as early as the 1920's, they had not been studied in computer science until 1965 when Lotfi A. Zadeh of the University of California published a paper on fuzzy sets, the backbone of fuzzy logic. Fuzzy logic is unique because it is a source of controversy not only in philosophy but in computer science as well, where it has been used in many successful applications. Yet despite its successes there are many compelling criticisms that fuzzy logic has yet to overcome. Both proponents and opponents of fuzzy logic tend to take extremist positions, a situation that makes fair accounts and criticisms of fuzzy logic difficult to come by. Engineers and computer scientists who use fuzzy logic vigorously defend their work and charge that their critics are conservatives who cannot see the fuzziness inherent to a given domain, a strong charge given the longstanding relationships many of the sciences have had with classical logic. Critics of fuzzy logic tend to take extreme positions as well: fuzzy logic has been described by Berkeley mathematician William Kahan as 'the cocaine of science' and as 'pornography' by Carnegie Mellon logician Dana Scott [Haa96, p. 230].

While the merits of infinitely-valued logics were being debated by philosophers, computer scientists had began building actual systems that used fuzzy logic. Early fuzzy expert systems and controllers were developed in the 1970's and fuzzy logic's ability to use linguistic rules were being explored and put into practice: steam generators and cement kilns were among the earliest applications of fuzzy control. Decision support systems were also being developed and one of the first commercial decision support systems to use fuzzy logic was developed by the INFORM Corporation in 1986 [vA95, p. 279]. Over the 1970's and 1980's Japanese researchers had begun to embrace fuzzy logic and by 1989 the Japanese government partnered with 49 companies to found the Laboratory for International Fuzzy Engineering Research (LIFE) [Ter95, p. 1]. LIFE's six-year mandate was to develop new fuzzy technologies and applications in 'human-friendly' fields like robotics and expert systems. By 1990 fuzzy rice cookers, vacuum cleaners, refrigerators, and other electronics were yielding large profits for Japanese companies [LY99, p. 7]. The last part of the 1990's saw many more concrete applications of fuzzy logic, particularly in computer systems designed to help humans make difficult and complex decisions. In addition to traditional applications in business and engineering, these applications were developed for many disciplines across the natural and social sciences as researchers in these fields slowly began to see the phenomenon of vagueness as being inherent to certain aspects of their work.

Instead of using a finite set of truth values, fuzzy logic has an infinite number of truth degrees that can be assigned to sentences. This is seen by some to be more precise than a 3 or 4-valued logic. For example, Kleene's 3-valued logic represents all propositions that are neither true nor false as being 'indeterminate'. In fuzzy logic, the truth value of any proposition p will be in the interval $[0, 1]$ and $t(p)$ is used to represent the truth of p. As in classical logic, $t(p) = 0$ means p is completely false and $t(p) = 1$ means p is completely true. Fuzzy sets are ordinary sets that allow graded membership and are used to represent linguistic labels like 'old', 'tall', 'true' and 'false'. It is important to note that for any $p \in \Gamma$, where Γ is a fuzzy set, $t(p)$ only represents p's degree of membership in Γ and not the probability that p is a member of Γ. To claim the sentence 'Bill is tall' is true to degree 0.5 means that 'Bill is tall' is as true as it is false and not that there is a 50% chance that 'Bill is tall' is true and a 50% chance that it is false.

Fuzzy operators resemble the operators of other many-valued logics and are defined by the three operations we can perform on fuzzy sets:

$\Gamma' = \{p \in U : p \notin \Gamma\}$

$\Gamma \cup \Delta = \{p : p \in \Gamma \text{ or } p \in \Delta\}$

$\Gamma \cap \Delta = \{p : p \in \Gamma \text{ and } p \in \Delta\}$

Thus, negation can be defined as $t(\neg p) = 1 - t(p)$, conjunction as $t(p \wedge q) = min(t(p), t(q))$, and disjunction as $t(p \vee q) = max(t(p), t(q))$. Implication is more complicated as there are three potential constructions of fuzzy implication, each of which is equivalent to classical implication when its input values are restricted to 0 and 1 [BE02, p. 18]. The different implication operators are: $t(p \rightarrow_1 q) = t(\neg p \vee q)$, $t(p \rightarrow_2 q) = min(1, 1 - t(p) + t(q))$, and $t(p \rightarrow_3 q) = max(1 - t(p), t(p \wedge q))$. There are two different definitions of validity in fuzzy logic. The first preserves absolute truth by defining an argument as being valid if and only if its conclusion has a truth value of 1 whenever each of its premises has a truth value of 1; the second preserves degrees of truth by defining an argument as being valid if and only if the truth of its conclusion is greater than or equal to the truth of its weakest premise. A valid argument is sound if and only if its premises are true to degree 1. More information on fuzzy operators and the mathematical foundations of fuzzy logic can be found in [BE02] and [NW97].

2 The Sorites and Paradoxes of Vagueness

The phenomenon of vagueness that has been described thus far can be found throughout everyday life; for this reason, paradoxes of vagueness are of special concern to researchers from many disciplines. Typically, a paradox resulting from vagueness is referred to as a sorites or sorites-like paradox. The sorites paradox, also known as the paradox of the heap, is an attempt at formalizing our indecision over borderline cases. The sorites involves a heap of sand of sufficient size q and the true sentence P_q, 'q grains of sand constitute a heap of sand'. Now suppose one grain of sand is removed, so that the heap is now of size $q - 1$. Because it is odd to think that one grain of sand could make the difference between heapness and non-heapness, it seems as though P_{q-1} is still true. It is then reasonable to say that 'for all numbers n, if n grains of sand constitutes a heap of sand, then $n - 1$ grains constitute a heap of sand'. Yet we could repeat this process $q - 1$ times and reach the intuitively unacceptable conclusion that one grain of sand constitutes a heap; if we reject this conclusion, however, we are forced to conclude that there is some number n such that that n grains of sand constitute a heap and $n - 1$ do not. The reasoning of this paradox can be made explicit:

P(1): 100,000 grains of sand constitute a heap of sand
P(2): If 100,000 grains of sand constitute a heap of sand, then 99,999 grains of sand constitute a heap of sand as well
C(1): Therefore, 99,999 grains of sand constitute a heap of sand
P(3): If 99,999 grains of sand constitute a heap of sand, then 99,998 grains of sand constitute a heap of sand as well
C(2): Therefore, 99,998 grains of sand constitute a heap of sand
P(4): If 99,998 grains of sand constitute a heap of sand, then 99,997 grains of sand constitute a heap of sand as well
C(3): Therefore, 99,997 grains of sand constitute a heap of sand
⋮
P(100,000): If 2 grains of sand constitute a heap of sand, then 1 grain of sand constitutes a heap of sand as well
C(99,999): Therefore, 1 grain of sand constitutes a heap of sand

Note that the final conclusion is deduced by repeated application of modus ponens. If we were to model 'heap' in classical logic at least one premise from the above argument would need to be rejected. This would have the effect of imposing a strict bound on 'heap'. For example, we could reject P(99,000) and decide that 1,000 grains of sand constitute a heap of sand but 999 grains do not. Yet it is intuitively absurd to hold not only that one grain of sand makes the difference between heapness and non-heapness, but that it is precisely the thousandth grain of sand that makes the difference.

2.1 The Sorites Paradox

Problems of vagueness were first studied during the time of Aristotle, when Eubulides (born c. 400 B.C.E.) is claimed to have discovered the sorites paradox; in addition to the sorites, he is known to have discovered at least six other important paradoxes [Res01, pp. 77–78]. While these paradoxes were no doubt studied for their own sake, it is believed that Eubulides used the sorites to attack Aristotle's doctrine of the Mean [Mol69, p. 395]. While Eubulides doesn't explicitly attack the doctrine of the Mean, many ancient and contemporary commentators have inferred the connection based on the known animosity that existed between the two men: Eubulides was known to have slandered Aristotle a great deal and Eusebius, a bishop in the fourth century, claimed to have read a book by Eubulides that accused Aristotle of betraying Plato and spying for Philip of Macedon [Mol69, p. 394]. Yet not all modern commentators agree with this assessment: Barnes examines Eubulides's writings and Aristotle's alleged replies and concludes that Aristotle "either did not know the sorites or else kept his knowledge to himself . . . there is no reason to look for a philosophical debate between Eubulides and Aristotle . . . based upon soritical arguments" [Bar82, p. 41]. Although Eubulides is credited with the sorites paradox, he may have been inspired by Zeno of Elea's (c. 490–430 B.C.E.) millet seed paradox [Wil94, p. 9]. Zeno begins his presentation of the paradox by stating two facts: a single millet seed does not make a sound when it falls and a bushel of millet does. He then argues that because a millet seed is some fraction of the bushel, the millet seed must actually make a sound when it falls. Despite the similarities between the millet seed and sorites paradoxes, it does not appear as though Zeno was interested in the slippery slope of the paradox and it has been argued that he fell victim to a fallacy of composition [Sor03, p. 53]. Democritus (c. 460–370 B.C.E.) is also credited with a rudimentary version of the sorites that used a step-by-step argument to try to show that a cone either had many indentations or was really a cylinder [Sor03, pp. 53–54].

Diodorus Cronus (c. 315–285 B.C.E.) is credited with circulating Eubulides's paradoxes and, arguably, Diodorus's claims about motion and past tense statements can be connected to the sorites [Wil94, p. 10]. Chrysippus (c. 280–207 B.C.E.) is reported to have written several scrolls on the sorites, though they are all lost [Bob02, p. 218]. Williamson believes Chrysippus was committed to sharp cut-off points for vague predicates [Wil94, p. 13]; this belief is motivated, in part, by Chrysippus's claim that before a borderline case is reached in a sorites series one should fall silent and stop answering [Cic, 2.94]. Williamson's claim is disputed by Bobzien [Bob02], who argues that the surviving evidence shows that Chrysippus was hardly an epistemicist and, if anything, would have advocated for a non-epistemic solution to the sorites. Carneades (c. 214–129 B.C.E.) demonstrated the absurdity of the sorites by using it to show that gods do not exist; this theological version of the sorites is discussed in section 2.2. Carneades also attacked Chrysippus's suggestion of silence as being a temporary avoidance of the paradox that does nothing to further our understanding of it [Cic, 2.94]. Cicero (c. 106–43 B.C.E.) described the sorites as a "very vicious and captious style of arguing" [Sor03, p. 88] and believed that "nature has permitted us no knowledge of limits such as would enable us to determine, in any case, how far to go [in a sorites series]" [Cic, 2.94]. The sorites was well-known in ancient times and its reach extended beyond philosophical circles: the Latin poet Horace (c. 65–8 B.C.E.) mockingly uses the sorites to show that if any poet is old, then all the poets must be old; the Roman satirist Persius (c. 34–62 C.E.) claims to have found a limit to "Chrysippus's heap" by noting that the man who keeps doubling the asking price for his soul must eventually stop and choose a price [Bar82, p. 36].

Although much of Galen's (c. 129–216 C.E.) philosophy has been lost, he knew of the sorites and recognized its impact on terms that describe phases of human life, like 'boyhood' and 'adolescence' [Gal, 16.1–17.3]; while the relevance of describing these terms as vague may not have had much impact in Galen's time, the vagueness of these terms has implications for words like 'alive' and 'person' that are frequently used in contemporary moral philosophy. Galen's writings on the sorites reflect a centuries-old debate between two competing schools of medicine over the nature of medical experience, namely over whether a certain number of observations could constitute having experience. The Empirical school held that many observations would constitute having experience while the Logical school claimed that the sorites proved no amount of observations would be enough [Bur82, p. 319]. Responding to this, Empirical Doctors would claim that the sorites cannot be trusted because, “by the same reasoning, doubt and confusion enter into many other things which relate to the doings of men in spite of the fact that knowledge of these things is obvious and plain” [Gal, 16.1–17.3]. The Empirical Doctors claimed that the sorites offered a choice between accepting either an “argument” or “what is plain to the senses”; thus, they claimed, it could be easily rejected without further analysis [Bar82, p. 58]. Interestingly, the position of the Empirical school can be interpreted as advocating for a 3-valued approach to vagueness: one can interpret their claim that there was no precise answer to the sorites as being an appeal to indeterminacy; it is important to note, however, that the Empiricists had no desire to solve the sorites or create nonstandard logics and their reply was not intended to be philosophical in nature [Bar82, p. 59–64]. These types of arguments played an important role in ancient Greek medicine: at the time, medical knowledge was mostly speculative and medical theories were offered by doctors, sophists and natural philosophers [Llo79, pp. 89–98]. As a result, many ancient doctors required training in rhetoric to convince potential clients that they were more competent than competing doctors.

Though it could be found in logic textbooks after the middle ages [Wil94, p. 33], few philosophers studied the sorites between antiquity and the nineteenth century. Frege was one of the first philosophers in the nineteenth century to seriously examine vagueness and he concluded that it is simply a defect of natural language. Citing the law of the excluded middle as support for his argument, Frege claims that each concept “must have a sharp boundary” and every object clearly belongs or doesn’t belong to a given concept; concepts without such boundaries, he claims, are “wrongly termed a concept” [Fre03, p. 56]. Frege’s arguments entail one of two views: either vague predicates can eventually be sharpened and become legitimate concepts, or no sharpenings exist and they will be forever incoherent [Bur91, pp. 4–5]. Frege’s position stands in stark contrast to Wittgenstein, who held that vagueness was an essential feature of natural language. Wittgenstein believed vague concepts must have some meaning because in many cases they would serve us far better than precise alternatives [Hoo90, p. 64]. Peirce held that language would always be vague and believed that the phenomena we now separately describe as imprecision and unspecificity were a part of vagueness [Wil94, p. 47]. Unlike Frege and Peirce, Russell embraces vagueness and holds that all of natural language is vague; as evidence of this he gives examples of colour predicates, measurements, and proper names. Having established that these kinds of words are vague, Russell turns to the most precise words that he can find: logical terms and connectives. Russell argues that while logical terms can potentially have precise meanings, they can only operate on representations like words and symbols; because all of language is vague the logical terms will always operate on vague terms and thus be vague themselves. For Russell, language is where vagueness ends: he denies that vagueness can exist in the world, claiming that it "can only belong to a representation, of which language is an example" [Rus23, p. 35].

2.2 Forms of the Sorites

The word 'sorites paradox' need not only be applied to a 'heap', nor need it refer only to arguments of the form that were described in section 2. The paradox can be stated in a more compact and inductive form using any vague predicate:

P(1):	A man with one hair on his head is bald
P(2):	If someone with n hairs is bald, then someone with $n+1$ hairs is bald as well
C(1):	Someone with 10,000 hairs is bald

P(1):	A 120-year old person is elderly
P(2):	If someone n years-old is elderly, then someone $n-1$ years-old is elderly as well
C(1):	Someone 5 years-old is elderly.

These arguments can be shortened and the effect of mathematical induction highlighted by symbolizing the argument in predicate logic. Without any loss of generality, let $F_{100,000}$ be the true sentence '100,000 grains of sand constitute a heap of sand'. Using this conditional form of the sorites we can conclude that 1 grain of sand constitutes a heap of sand.

1	$F_{100,000}$	
2	$\forall n(F_n \supset F_{n-1})$	
3	$F_{100,000} \supset F_{99,999}$	$\forall$E, 2
4	$F_{99,999}$	$\Rightarrow$E, 1, 3
$\vdots$	$\vdots$	
$199,998$	F_2	$\Rightarrow$E, $\vdots$
$199,999$	$F_2 \supset F_1$	$\forall$E, 2
$200,000$	F_1	$\Rightarrow$E, $(199,998)$, $(199,999)$

Thus far, each version of the sorites has had the same logical structure and any method of dealing with the sorites would likely work for all of them. Yet the sorites need not be expressed only by conditionals; because $P \supset Q$ is logically equivalent to $\neg(P \wedge \neg Q)$ we can restate the inductive premise in the form $\forall n \neg(F_n \wedge \neg F_{n-1})$. Using this premise and without invoking DeMorgan's law, we can conclude that 1 grain of sand constitutes a heap of sand.

1	$F_{100,000}$	
2	$\forall n \neg(F_n \wedge \neg F_{n-1})$	
3	$\quad \neg F_{99,999}$	
4	$\quad F_{100,000} \wedge \neg F_{99,999}$	$\wedge$I, 1, 3
5	$\quad \neg(F_{100,000} \wedge \neg F_{99,999})$	$\forall$E, 2
6	$F_{99,999}$	$\neg$E, 3–5
$\vdots$	$\vdots$	
$399,994$	F_2	$\neg$E, $\vdots$
$399,995$	$\quad \neg F_1$	
$399,996$	$\quad F_2 \wedge \neg F_1$	$\wedge$I, $(399,994)$, $(399,995)$
$399,997$	$\quad \neg(F_2 \wedge \neg F_1)$	$\forall$E, 2
$399,998$	F_1	$\neg$E, $(399,995)$–$(399,997)$

In English, the inductive step of the above derivation would be 'it is not the case that n grains of sand constitute a heap of sand and $n-1$ grains of sand do not'. In [Wea00, p. 21], Weatherson argues that, for someone unfamiliar with the sorites paradox, an inductive premise of the form $\neg(P \wedge \neg Q)$ is more persuasive than one of the form $P \supset Q$. The plausibility of Weatherson's claim is discussed in section 4.2.

By DeMorgan's law, we know that $\neg(P \wedge \neg Q)$ is is logically equivalent to $\neg P \vee Q$ and yet another form of the sorites can be found using this equivalence, this time using $\forall n(\neg F_n \vee F_{n-1})$ as the inductive premise.

1	$F_{100,000}$	
2	$\forall n(\neg F_n \vee F_{n-1})$	
3	$\neg F_{100,000} \vee F_{99,999}$	$\forall$E, 2
4	$\neg\neg F_{100,000}$	$\neg\neg$I, 1
5	$F_{99,999}$	DS, 3, 4
$\vdots$	$\vdots$	
$299,996$	F_2	DS, $\vdots$
$299,997$	$\neg F_2 \vee F_1$	$\forall$E, 2
$299,998$	$\neg\neg F_2$	$\neg\neg$I, $(299,996)$
$299,999$	F_1	DS, $(299,997)$, $(299,998)$

In English, the inductive step of the above derivation would be 'either it is not the case that n grains of sand constitute a heap or $n-1$ grains of sand do'. The paradox seems far less persuasive when this disjunctive form of the premises is used. On this formulation, the premises do not seem to be intuitively true unless one already understands their logical equivalence to a more convincing form [Wea00, p. 21]. These equivalences pose a problem for some theories of vagueness as their account of why $P_n \supset P_{n-1}$ looks persuasive must also account for the persuasiveness, or lack of persuasiveness, of $\neg(P_n \wedge \neg P_{n-1})$ and $\neg P_n \vee P_{n-1}$. If a theory that evaluates vague sentences ascribes the same truth value to all three forms then it must give some account of why certain versions of the sorites are more compelling than others; if it ascribes different truth values, then it must explain its rejection of the classical equivalence and justify the truth values it gives to the different forms of the sorites.

The logically equivalent forms of the sorites that were examined in this section have not, historically, been the only ways of formalizing the paradox. Although the theological sorites presented by Carneades proceeds in the same 'little by little' fashion as the logical sorites, modern readers will identify it as suffering from fallacies of ambiguity and equivocation. Carneades claimed to be able to use reason to show that gods do not exist; this argument was not intended to disprove the existence of gods but rather to prove that the Stoics explained nothing about them [Bar82, p. 45]. By assuming that Zeus is a god, Carneades derived the following sorites series: "if Zeus is a god, Poseidon his brother will be a god; and if Poseidon is a god, Achelous will be a god; and if Achelous, Neilos, and if Neilos, every river, and if every river, the streams would be gods ... but the streams are not gods; therefore, neither is Zeus. But if there were gods, Zeus would be a god. Therefore, the gods do not exist" [Emp97, 9.183]. It has been suggested that this non-paradoxical form of the sorites hinged on the principle that if X deserves T, and Y is similar enough to X, then Y deserves T: the theological sorites adheres to this principle because Zeus and Poseidon were related as brothers, Poseidon and Achelous both named masses of water, and Achelous and Neilos were both names of rivers [Bur82, p. 328]. However, unlike the logical sorites, the premises of the theological sorites can be disambiguated and a paradox can be avoided. The sentence 'if Poseidon is a god, Achelous will be a god' can be interpreted either as the false sentence 'if Poseidon (the god) is a god, then Achelous (the river) will be a god' or as the vacuous, non-paradoxically true sentence 'if Poseidon (the river) is a god, then Achelous (the river) will be a god'. In this case, the conclusion that 'Poseidon (the god) is a god' will not satisfy the antecedent of the conditional and the conclusion that the gods do not exist can be avoided.

2.3 Sorites-like Paradoxes

Although predicates vulnerable to sorites reasoning are always vague, vagueness can arguably apply to more than just predicates and it can be extended to include nouns and other concepts. Though these arguments do not all share the same form of the sorites, they each rely on inductive steps to move 'bit by bit' from a seemingly obvious truth to a seemingly obvious falsehood. This style of argument can be found in the puzzles presented by Zeno and Democritus and, as a result, they are frequently credited with having discovered a rudimentary version of the sorites. While predicates like 'bald' have a finite number of borderline areas, borderlines on a colour continuum are potentially infinite. The colour spectrum, and colour predicates, are frequently used to justify the use of fuzzy logic in dealing with vague predicates. Imagining a continuum of colours, there is some spot that is clearly red and some spot that is clearly orange. But between these two poles are potentially an infinite number of colours that seem to gradually and uniformly change in degree from red to orange; in other words, we can potentially assign a fuzzy degree from the set of real numbers to each colour that is between red and orange. Vagueness can also be found on a finite colour spectrum: for instance, given 100 colour patches that gradually shade from one colour into another and the assumption that if two colours are visually indistinguishable then they are the same colour, we can conclude that colour patches 1 and 100 are the same colour; nevertheless, a simple observation may tell us that patches 1 and 100 are really separate colours [Res01, p. 84].

Vagueness can also lead to moral paradoxes. If we grant that it is morally wrong to have an abortion on or after the mth minute of pregnancy, then it seems as though it will also be wrong to have an abortion at the $m - 1$th minute [Oli03, p. 168]. As was the case with heaps and balding men, we can continue this chain of reasoning until we reach the conclusion that abortion is never morally permissible. While one may accept or reject the morality of abortion on a variety of grounds, it seems misguided to allow sorites reasoning to decide morality. It is not difficult to see how a bit-by-bit sorites-style argument can be used to advocate for many moral issues: one can argue that 'if a bit of X is morally right then surely a bit more of X is morally right'. Yet if one accepts the sorites paradox as being an unproblematic, sound and valid argument worthy of basing moral reasoning on then one will be forced to accept the many other counterintuitive consequences of accepting the sorites in classical logic; one of these consequences is Unger's 'problem of the many or the none' that is discussed in section 2.4.2. Nonetheless, laws frequently attempt to govern morality. When this occurs, the arbitrariness of the sharp cut-off points selected become apparent: discrimination against minorities is generally illegal but affirmative action allows the majority to be discriminated against to meet certain quotas; these quotas are precise and it is strange to think that it is morally right to discriminate so long as it is to some precise value.

Claiming that a backwards induction can be reconstructed as a sorites series, Collins and Varzi [CV00] look at a take-it or leave-it game to show that rationality predicates are vague. The game runs as follows: a banker places \$2 on the table and gives player 1 the choice of taking the money or leaving it. If they leave the money, then the banker adds \$1 to the pot and gives player 2 the choice of taking the money. If player 2 leaves the money, then \$1 is added to the pot and player 1 now has the choice of taking the money. At any point, the player who takes the money gets the pot and the other player receives \$2 less. Play continues until the pot reaches \$101, at which point player 2 can take the money or give it to player 1; the player who doesn't get the \$101 receives \$99. Common sense says both players should wait until the last stage of the game; on the other hand, if they try to maximize their individual financial gain then a backward induction can occur and lead to a paradox. The paradox arises because on the last turn the only way for player 2 to maximize their payout is to take the \$101. But player 1 knows this, so it follows that player 1 should take the money on the second-last turn; this would give her \$100 instead of the \$99 she'd receive if she let player 2 take the \$101. However, player 2 knows it makes the most sense for player 1 to take the money on the second-last turn, so he will take \$99 on the third last turn; this will give player 1 only \$97. This chain of reasoning continues until it seems as though the only rational course of action is for player 1 to take \$2 on the first turn.

Prima facie, this kind of backward induction may seem similar to the surprise examination paradox, first conceived of by Quine as an executioner's paradox [Qui53] and published by O'Connor as an air-raid paradox [O'C48]. Each version of the paradox stems from the same set of premises: some event will occur next week, the event may occur on any day next week, and those involved will not know about the event until the day it occurs [Res01, p. 113]. Suppose the event is a surprise examination and a teacher has told their students about the surprise test using the wording of the premises above. From the first and second premises, it follows that Friday is the last day on which the test can occur. But if the test has not occurred by Thursday then the students will know on Thursday that the test must be on Friday, since Friday is the last possible day that the test could be given. This would contradict the third premise, that the students will not know about the exam until the day it is given. So the test cannot occur on Friday. It then follows that Thursday is the last day on which the test could be held. But if the test has not occurred by Wednesday then the students will know on Wednesday that the test must be on Thursday, since Thursday is now the last possible day that the test could be given. This chain of reasoning continues until we reach the conclusion that the teacher has lied and cannot possibly surprise the students with a test. While the surprise examination paradox hinges on backwards induction, it does not seem to have the same structure as the take-it or leave-it game and it has been suggested that the third premise of the surprise examination paradox suffers from the fallacy of equivocation [Res01, p. 113]. Yet no such explanation has been given for the take-it or leave-it game and a vulnerability to backwards induction is another possible symptom of vagueness.

2.4 Vague Objects

Most philosophers share Russell's view that vagueness can only be found in language and representations. Yet to its proponents, arguments for the existence of vague objects are intuitively obvious: Edgington describes a cloud whose water droplets get progressively less dense so there is some boundary where it seems uncertain whether a given droplet is a part of the cloud or simply next to it [Edg00, p. 29]. One cannot simply say that any adjoining droplet is a part of the cloud as this would allow us to say the whole earth is covered by one large cloud; yet it seems counterintuitive to claim that clouds are composed only of water droplets within a certain number of millimeters. Thus, a sorites series is formed:

P(1):	A water droplet 1 millimeter away from a cloud is a part of that cloud
P(2):	If a water droplet n millimeters away from a cloud is a part of that cloud then a water droplet $n+1$ millimeters away from that cloud is also a part of the cloud
C(1):	Every water droplet in the sky is a part of the same cloud

Proponents of vague objects see their thesis as being overwhelmingly intuitive: Copeland and Tye both describe the view that objects have precise boundaries as being common sense [Cop95, Tye95]. Yet critics of vague objects also see their position as being intuitive: McGee describes the thesis of vague objects as being "zany" and "wildly implausible" [McG97, pp. 141–143] and Dummett describes the notion of vague objects as being "not properly intelligible" [Dum75, p. 314]. Nonetheless, it is the proponents of vague objects who have mastered arguments from intuition as they have the luxury of positing clouds, mountains, cities and animals as examples of vague boundaries that can be found in everyday objects, as is the case in Edgington's cloud example. It has been claimed that the philosophical denial of vague objects stems from a belief that the physical world can be divided at the microscopic level into discrete objects that, in principle, could be sharply described; yet we could suppose that every piece of matter, including the underlying particles, must themselves decompose into smaller parts [Bur90, p. 279]. As we examine the underlying layers of an object, its bounds become fuzzier and we will find parts for which it is indeterminate whether or not they are a part of the object.

Vague objects are not limited to cases where molecules are being removed. There are molecules that are clearly a part of Mt. Everest and there are molecules that are clearly not a part of Mt. Everest, but between them there are molecules for which it is, arguably, indeterminate whether they are or are not a part of Mt. Everest [Tye90, p. 535]. To insist on precise boundaries is, *prima facie*, to claim that only one lump of molecules constitutes Mt. Everest and it is this lump that speakers of English refer to when they utter 'Mt. Everest'. Were we to refer to Mt. Everest plus the closest borderline molecule that is not a part of Mt. Everest then, while we may be referring to something, it is not Mt. Everest. While this conclusion is intuitively unacceptable, a rejection of precise boundaries leads to another unacceptable outcome: a sorites-style paradox.

P(1):	There is some atom at the base of Mt. Everest that is clearly a part of Mt. Everest
P(2):	If two atoms are adjacent to one another then they are a part of the same object
P(3):	There is a chain of atoms from Mt. Everest to Paris
C(1):	Paris is a part of Mt. Everest

The conclusion, that Paris is a part of Mt. Everest, is as absurd as claiming that one grain of sand constitutes a heap and those who argue for the existence of vague objects need to ensure that their preferred solution to the sorites can deal with these problems. Arguments for the existence of vague objects can also be constructed using the notion of object-parts. Morreau argues that when a whisker on the face of Tibbles the cat begins to loosen it becomes a questionable part of Tibbles, neither definitely a part nor definitely not a part of him [Mor02, p. 334]; when the whisker falls off, one wants to say that the Tibbles without the whisker is the same Tibbles as the one with the whisker. Objects are vague, according to Morreau, because they have questionable parts; these parts are gained and lost gradually such that there is no definite time when they change from being a part of an organism to not being a part of an organism [Mor02, p. 341].

The ship of Theseus paradox and the paradox of Sir John Cutler's stockings are both based on similar concerns about object-parts. The ship of Theseus paradox can be found in the writings of Plutarch where he describes a ship that, upon its return from Athens, was preserved for many years by replacing its decaying planks with new planks made from stronger timber; philosophers, he claimed, would ask whether the ship with the new planks was the same ship that returned many years earlier [Sor03, p. 132]. Hobbes modified this paradox by supposing that someone kept the old planks and eventually reassembled them [Sor03, p. 132]. Intuitively, the removal of one plank should not make a difference between a ship being the ship of Theseus and not being the ship of Theseus, yet it seems as though there is some point at which the reassembled ship becomes the ship of Theseus and the original ship becomes something else. This dilemma can be formulated as a sorites-style paradox:

P(1):	The ship of Theseus is composed of 5,000 planks that originate from the ship of Theseus
P(2):	Ship B is composed of 5,000 planks that originate from ship B
P(3):	If one plank from the ship of Theseus is exchanged for a plank from ship B, then the ship of Theseus and ship B are the same objects that they were before
C(1):	The ship of Theseus would be the same object if it were composed of 5,000 of its own planks as it would be if it were composed of 5,000 of ship B's planks

As the planks are switched either the object originally known as the ship of Theseus gradually becomes less identical with the actual ship of Theseus or there is some plank m for which it is the case that, were it to be removed, the ship in question could no longer be identified as the ship of Theseus. It also seems to be the case that as the reassembled ship is constructed either it gradually becomes more identical with the ship of Theseus or there is some plank l for which it is the case that, were it to be added, the ship in question would become the ship of Theseus. The paradox of Sir John Cutler's stockings is similar [Res01, p. 84]: over the years Sir John Cutler's favourite pair of stockings begin to wear out and, as they do, they are patched with a new piece of fabric. Eventually, every piece of the original stocking is replaced with a new piece of fabric and none of the original fabric remains. As was the case with the ship of Theseus, one can ask whether the remaining stockings are the same as Cutler's favourite pair and, if not, at what point they ceased to be the same thing. Because it is seemingly impossible to crisply determine how much of an original object can be replaced before it stops being that same object, Rescher argues, these are genuine examples of vagueness [Res01, p. 86]. They also seem to support the existence of vague objects: consider the scenario where a ship's planks are gradually switched with another ship and the borderline case where enough planks have been removed such that it is indeterminate whether the original ship, minus the planks, is identical to the original ship or the ship with which the planks were switched. Assuming we deny the existence of vague objects, Edgington claims, it follows that the borderline ship must be some third ship since it is neither the original ship nor the ship with which the planks were switched [Edg00, p. 32]. This follows because if it were identical to either the original ship or the ship with which the planks were switched then it would not be indeterminate whether it was the original ship or not. Yet this contradicts our strong intuition that there are only two ships and, if one accepts this argument, it seems as though we must reject the assumption that there are no vague objects.

A rejection of vague objects does not necessarily entail a rejection of the seemingly fuzzy boundaries of objects: those who accept the notion of vagueness can hold that it is the vagueness of the word 'Everest' that causes the apparent vagueness of the mountain; that is, words like 'Everest' are vague terms that denote precise objects [Alt01, p. 437]. McGee believes that this sort of explanation is "most natural" and uses a criticism of fuzzy logic to completely dismiss the notion of vague objects: because a fuzzy object requires a precise inner boundary such that all points within it are members of the object to degree 1, a precise outer boundary such that everything outside it are members of the object to degree 0, and a precise numerical function to link the two boundaries, it appears as though notions of fuzzy objects require even more precision, namely a precisely determined vague object [McG97, pp. 142–143]. McGee's argument hinges on the problem of higher-order vagueness, a well known-problem for nearly every theory of vagueness, including fuzzy logic, that is briefly discussed in section 4.2. Despite the convincing nature of McGee's argument, it seems unreasonable to altogether reject the notion of vague objects because of one problem with one proposed method of resolving vagueness.

Unlike McGee, Burgess believes the "more natural view" is that vague objects exist independently of us and that instead of being the source of vagueness, language simply determines which vague objects are referred to: it is doubtful, he writes, that "God [would] have use for more precise 'hills' when contemplating our interactions with [them]" [Bur90, p. 284]. Sainsbury notes that vague words like 'heap' do not require the existence of vague objects because every pile of grain has a definite number of grains [Sai88, p. 46]. While this view allows linguistic vagueness to be distanced from notions of vague objects, the argument only shows that the existence of vague words does not always entail the existence of vague objects: although a borderline heap of grain in the middle of an empty room appears to be a precise object, a borderline heap of grain laying in a grain-filled room would, to proponents of vague objects, constitute both a borderline heap and a vague object. For it would be indeterminate not only whether the number of grains in the pile constituted a heap, but whether certain grains were part of the heap, part of an adjoining pile, or part of the grain on the floor. So while 'heap' does not require the existence of vague objects, it seems as though there can be instances of 'heap' that cannot be resolved simply on the basis of the number of grains in the pile.

2.4.1 Vague Objects and Gareth Evans

In 1978 Gareth Evans published 'Can There Be Vague Objects?' [Eva78], a one-page article described by McGee as "[bearing] all the signs of having been written late one night and mailed off to *Analysis* the next morning" [McG97, p. 147]. In his article, Evans produces a proof that he claims establishes the impossibility of vague objects. Many commentators have found what appear to be small logical errors which could be rectified by adding certain assumptions. Yet the proof is accompanied only by a single paragraph that briefly introduces the concept of vague objects and none of these problems or their potential solutions are discussed by Evans. So while Evans's supporters tend to credit him with the required unwritten assumptions, "the problem [has become] one of trying to find the readings Evans had in mind" [Ove84, p. 97]; furthermore, these criticisms make it difficult to measure Evans's contribution to the literature: "there is little agreement regarding what is in dispute, what sorts of arguments might decide it, how Evans's argument addresses the problem, or what objections to that argument are relevant" [Hec98, p. 274]. Despite these uncertainties, nearly every study of vague objects discusses Evans's paper.

Evans's proof begins by assuming, for the sake of contradiction, that it is indeterminate whether a is identical to b. Because it is not, by assumption, indeterminate whether a is identical to a it follows by the contrapositive of

Leibniz's Law that a is not identical to b, since a has the property of being determinately identical to a while b does not. But the fact that a is not identical to b contradicts the assumption that it is indeterminate whether a is identical to b and, because this argument holds for any two objects a and b, it is claimed that identity statements can never be indeterminate. Note that although Evans does not give a schema of Leibniz's Law, it is generally assumed that he intends the classical reading of $a = b \supset (\Phi(a) \equiv \Phi(b))$. The proof has several problems with the semantics of its operators that, though seemingly small in scope, have forced others to try and infer what Evans 'really' meant. One common criticism of Evans's proof is the seemingly ambiguous use of his determinacy operator 'Δ': in some cases he uses it to mean that a sentence is determinately true while in others he uses it to mean that a sentence is either determinately true or determinately false; the former case is sometimes classified as 'it is determined that' and the latter as 'it is determined whether' [McG97, p. 147]. At the beginning of his proof, Evans assumes $\nabla(a = b)$, where '∇' is defined as an indeterminacy operator and the dual of 'Δ': in this case, '∇' is being used as part of the assumption that $a = b$ is indeterminate in truth value; it is not being used to claim that $a = b$ is not determinately true. Yet Evans's assumption that 'it is determinate that a is identical to a' is formulated as $\neg\nabla(a = a)$, where '$\neg\nabla$' is being used to assume that $a = a$ is determinately true; it is not being used to claim that $a = a$ is either determinately true or determinately false. The ambiguity arises from the fact that in one line of the proof 'Δ' is being used to refer to a sentence's truth value being determinate while, in another line, it is being used to refer to a sentence being determinately true. Yet Evans is clearly not intending for 'Δ' to refer to a sentence being determinately true: on this reading, if we apply duality to the key assumption $\neg\nabla(a = a)$, we can conclude $\neg\neg\Delta\neg(a = a)$, or $\Delta\neg(a = a)$; that is, it is determinate that a is not identical with a. Yet without this interpretation of 'Δ' Evans's proof fails; while McGee claims Evans can correct this ambiguity by simply adding an extra operator, Over notes that the resulting logic would no longer meet Evans's claim of being S5. Without S5 one of the theorems Evans relies on, $\Delta(\phi \supset \Psi) \supset (\Delta\phi \supset \Delta\Phi)$, would no longer hold and while one could change logics to overcome these semantic objections, it seems clear that Evans "did not have a precise logic clearly in mind" [Ove84, p. 98].

Semantic objections aside, Evans's argument has been heavily criticized for two major assumptions: firstly, that any account of vague objects requires one to accept the notion of indeterminate identity and, secondly, that Evans's conclusion actually does contradict a theory of vague objects. Evans begins with the assumption that it is indeterminate whether $a = b$ and then derives a contradiction; thus, even if one accepts Evans's proof one could only claim that there is no such thing as indeterminate identity. Although Evans believes that this is enough to reject vague objects, he offers no account of why vague objects require indeterminate identity. Copeland criticizes Evans's proof but gives the following example to support the assumption that vague objects require indeterminate identity: suppose a seventeenth century British sailor names a natural harbour along the coastline of what is now Western Australia as 'New Devon' and that, a year earlier, a French sailor had named roughly the same spot as 'Nouvelle Provence' [Cop95, p. 83]; Copeland believes that 'New Devon = Nouvelle Provence' must be indeterminate since it seems neither true nor false. Yet Copeland's claim is rejected by many authors: Morreau gives an account of vague objects that do not have indeterminate identities; that is, they can be identical to themselves or different from other objects but are indefinitely identical to nothing [Mor02, p. 357]. Burns notes that indefinite identity is not required for vague boundaries: when a and b denote the same vague object there is no reason to deny the truth of $a = b$ and when they do not denote the same vague object then there is no reason to deny the truth of $\neg a = b$ [Bur91, p. 15].

Because few proponents of vague objects believe that classical logic alone can suitably handle vagueness, it should come as no surprise that one common response to Evans is to claim that the classical form of Leibniz's Law is invalid in domains with vague objects. One such claim has to do with Evans's use of the contrapositive of Leibniz's Law, an unproblematic move in classical logic, according to which a statement and its contrapositive are logically equivalent. If one allows for the possibility of a many-valued logic or a 2-valued logic with truth-gaps, however, then this principle does not always hold: $A \vdash \neg\nabla A$ is valid but its contrapositive, $\nabla A \vdash \neg A$, is not [PW96, p. 327]. A different criticism has to do with Evans's use of classical logic: Copeland notes that if we grant Evans's initial assumption, that it is indeterminate whether a is identical to b, then one should be willing to consider using fuzzy logic [Cop95, p. 84]. Copeland then gives a fuzzy logic equivalent of Leibniz's Law and shows that it is unable to generate the same consequences as its classical counterpart. For example, as the antecedent of Leibniz's Law decreases in truth the truth of the consequent will decrease as well. This is problematic for when Φ contains the '∇' operator the value of $\Phi(a) \equiv \Phi(b)$ can quickly become 0. For example, when Leibniz's Law is formulated as

$a = b \supset (\lambda x[\nabla(x = a)]a \equiv \lambda x[\nabla(x = a)]b)$ and the value of $\lambda x[\nabla(x = a)]b$ is 1 the value of Leibniz's Law will be non-integral, thus giving the alleged contradiction in Evans's proof only partial truth [Cop95, p. 87]. Fuzzy logic is also used in [Pri98], where Priest uses a fuzzy identity relation to claim that Evans's use of the substitutivity of identicals is inappropriate.

While Evans's argument loses its force in non-classical logics, many authors have avoided discussing that contentious issue, choosing instead to focus on the proof's alleged contradiction. Tye grants the applicability of Leibniz's Law and argues that Evans's conclusion does not threaten vague objects because one can have definite identity statements that are vague. Suppose a and b denote similar vague objects made of slightly different chunks of matter: $a = b$ will be vague because both a and b are vague but $a = b$ is clearly false [Tye90, p. 556]. Heck claims that we must distinguish, in the case of non-classical logics, between the unsatisfiability of a formula and the truth of its negation; thus, Heck claims, Evans only shows the unsatisfiability of $\Delta a = b$ and does not threaten a conception of vague objects without definite identity [Hec98]. Others reject Evans's conclusion as being vulnerable to counterexamples. Using Frege's claim that "it is not clear when to declare a saddle to be in the middle of one mountain and when between two mountains" [Qui60, p. 126], Burgess cites the following example as *prima facie* evidence of indeterminate identity:

> let us suppose that a community ... has named one mountain 'Aphla' and one mountain 'Ateb'. Unbeknownst to them, the only visible peak of the eminence they call 'Aphla' and the only visible peak of the eminence they call 'Ateb' are separated by a saddle in the manner envisaged by Frege ... there is no question here of anything being a borderline case of a mountain; 'Aphla' seems to denote one and only one mountain, as does 'Ateb'. What appears to be indeterminate is whether Aphla is Ateb [Bur90, p. 269].

Another counterexample is presented by Sainsbury, who notes that the Ship of Theseus paradox seems to show that identity is governed by the vague principle that "replacing some, but not too many, parts of an artifact does not destroy it, but leaves you with the very same artifact" [Sai88, p. 48]. The domain of quantum mechanics offers yet another counterexample to Evans's proof. Imagine an atom that captures a free electron to form a negative ion which, in turn, releases an electron and reverts to its neutral state [Low94, p. 110]. On the standard interpretation of quantum mechanics we cannot follow the trajectory of an electron or track it in any way and, because there is clearly an electron before capture and after re-emission, it seems indeterminate whether the electron emitted by the ion is identical with the electron captured by the atom [Mai00, p. 263]. The view of quantum mechanics Lowe subscribes to holds that the electrons are in an 'entangled' state such that the number of electrons present is determinate but their identities are not [Low94, p. 110]; thus, Lowe claims, it is well known that this is not a case of epistemic indeterminacy since there is no determinate fact of the matter as to whether the two electrons are identical.

2.4.2 The Problem of the Many or the None

It was in Unger's 1979 paper 'There Are No Ordinary Things' [Ung79] that what would come to be known as 'the problem of the none' was first proposed. Because one grain of sand does not constitute a heap of sand and the addition of one grain of sand cannot make the difference between there being or not being a heap, Unger concludes that there are no such things as heaps. He goes on to extend this argument to ordinary things such as "pieces of furniture, rocks and stones, plants and ordinary stars, and even lakes and mountains" [Ung79, p. 119]. Unger calls this the direct argument; realizing that one may object to it, Unger gives several indirect arguments for his conclusion, each of which is based on the sorites paradox. The first indirect argument is called the 'sorites of decomposition'. Assuming that there is at least one stone, that it consists of a finite number of atoms and that one atom cannot make the difference between there being and not being a stone, one can conclude that a stone has an infinite number of atoms. This argument can be represented in a sorites series:

P(1): There is at least one stone
P(2): Every stone consists of a finite number of atoms
P(3): If a stone is comprised of n atoms and one atom is removed then a stone will remain, comprised of $n-1$ atoms

---

C(1): An infinite number of atoms can be removed from a stone

Note that **C(1)** follows from the premises because there is no limit imposed on the number of times **P(3)** can be applied. Thus, because the conclusion contradicts the premise that there is at least one stone and because any ordinary thing can be used in place of 'stone', Unger concludes that there are no ordinary things. Of course, one can instead choose to reject this conclusion by rejecting **P(3)**; as is the case with linguistic vagueness, one must now explain why there is some number of atoms n such that n atoms form a stone and $n - 1$ do not. Unger's next indirect argument is called the 'sorites of slicing and grinding' and involves a table that is slowly cut into fifths; if one does not identify which fifth constitutes the change from tablehood to non-tablehood, then one must conclude that there are no such things as tables. However, if one chooses some fifth that marks this change, "we may imaginatively divide the problematic fifth itself into, say, fifths. And then we may slice and grind again ... down to ten twenty-fifths ... we may perform again and again the same partitioning procedure until any resister will find absurd his own putative sensitivity" [Ung79, p. 134]. Unger goes on to restate his argument with both a 'sorites of cutting and separating' and a 'sorites of accumulation'.

In 1980, Unger published a follow-up paper that gave readers an alternate explanation for the phenomenon described in his previous paper: instead of using the sorites to conclude that there were no objects, one could use an account of vague objects to claim that there were many objects. 'The Problem of the Many' [Ung80] relies on a conception of vague objects because it assumes that ordinary objects, like clouds, lack crisp definitions and boundaries. Consider the group of water droplets that compose what appears to be one cloud in the sky: because none of the water droplets are themselves clouds, it must be the case that some combination of water droplets is what forms a cloud. The 'problem of the many' takes hold because there are millions of potential combinations of water droplets that could form what we believe is a single cloud and no one combination is any better than the others. It would be arbitrary to pick one combination and say that it is the cloud since no one combination has a greater claim to cloudhood than any other; thus, concludes Unger, it must be the case that there are many clouds. As was the case with 'the problem of the many' the argument can be extended to other ordinary things, though Unger concedes that there may be cases where there is a non-arbitrary reason for choosing one combination over another [Ung80, p. 416]. 'The problem of the many or the none' is a problem for proponents of vague objects because Unger's arguments fail if one rejects the notion of vague objects. It is only because we assume there is no precise number or arrangement of atoms required to constitute ordinary objects that we accept Unger's premises; were we to hold that 'stones' were composed of n atoms arranged in a certain precise configuration, then 'the problem of the many or the none' would be defused.

3 Vagueness in Practice

Although vagueness can be found throughout natural language, problems with vague predicates are sometimes avoided in everyday life by imposing artificial limitations on vague predicates. For example, governments define an adult as a person who is at least 18 years of age and a minor as someone 17 years and 364 days of age or less [Nis04, p. 51]. These artificial boundaries introduce many conceptual problems as there does not seem to be anything about a given person that makes them a 'minor' one day and an 'adult' the next. The definition is arbitrary and its underlying concepts are counterintuitive, yet without a crisp definition it is difficult to see how a legal system could function. Furthermore, these limitations can only work in certain contexts: a computer system could not properly interact with humans in natural language if it precisely defined predicates like 'bald', 'tall', and 'rich', though presumably it could properly interact using the precise definitions of 'minor' and 'adult'. There have been many attempts to duplicate the successful use of vague predicates by natural language speakers in computer systems using nonstandard logics; classical logic is rejected because of the sharpness imposed by bivalence. Although formal accounts of vagueness have traditionally been confined to the philosophical domain, the scope of the problem should give pause to anyone wishing to discuss the problem in purely abstract terms. Every philosophical account of vagueness has a broad range of implications in many other disciplines and before subscribing to a particular theory one should carefully examine these consequences.

Besides computer science and legal applications, applications of vagueness have been found in many unlikely areas. In [San94], Sangalli suggests democracy would be better served if fuzzy ballots were used; that is, instead of crisply choosing one candidate or position votes could be partially cast, theoretically using values from an infinite continuum. Defending his proposal, Sangalli simply says that most people do not take black or white positions on issues and if people could vote exactly how they felt the outcome of a close election or referendum might be different. It is doubtful Sangalli is seriously advocating electoral reform; rather, he seems to be suggesting that the world is far less precise than advocates of precise boundaries would have us believe. Wang's paradox shows that even mathematics is susceptible to vagueness: because 0 is a 'small' number, and any number following a 'small' number is itself 'small', every number must be 'small' [Dum75, p. 303]. While Dummett dismantles Wang's paradox by concluding that the use of vague predicates is incoherent [Dum75, p. 319], not everyone agrees that vagueness and mathematics are incompatible with one another. In [Ben00a], Bendegem claims that it would have been possible for a vague mathematics to develop instead of our traditional, precise mathematics. He then proves several theorems using such a math, including “equality implies rough equality”, “there is a large number k such that for all small numbers n, $k + n$ is roughly equal to n” and “small numbers have few prime numbers” [Ben00a, pp. 27-29].

Perhaps the most unique application of vagueness has been its use in defining the world view of a small group of computer scientists and engineers. This group refers to vagueness as ‘fuzziness’ and exclusively uses fuzzy logic to deal with the fuzziness they claim to find in the world. Although the origins of this kind of thinking can be traced to the Japanese fuzzy boom, the view gained popularity in North America in 1993 when Bart Kosko published ‘Fuzzy Thinking’ [Kos93] and proposed a fuzzy view of the world. In both this and his follow-up book, ‘Heaven in a Chip: Fuzzy Visions of Society and Science in the Digital Age’ [Kos99], Kosko rejected probability theory and classical logic and argued that nearly every aspect of life could be measured in degrees of fuzziness including parking a car, gambling, tax forms, property, war, living and dying [Kos93, Kos99]. His work is a subject of much controversy: while many in the fuzzy logic community seem to view it as a vindication of their work, critics dismiss Kosko’s arguments by pointing to his apparent extremism. In 2000, Vladimir Dimitrov coined the term ‘fuzziology’ to describe what he calls the fuzziness of human knowledge, its sources and dynamics; to Dimitrov, Fuzziologists embrace the unknown and do not impose rules on it, avoid the need for certainty and common sense ideas, and do not fight with the dynamics of life [D^+02, p. 17]. Although these views will not be studied here it is important to distinguish between several important philosophical concepts that these ideologies threaten to blur. Philosophically, and for the purposes of this paper, ‘vagueness’ refers strictly to the phenomenon that occurs when predicates satisfy one of the forms of the sorites discussed in section 2.2. ‘Fuzzy logic’ refers strictly to the many-valued logic that was discussed in section 1.2.2 and the arguments put forth in support of fuzzy logic are given only in the context of a solution to the sorites paradox and the phenomenon of vagueness.

3.1 Vagueness in Computer Science

The natural language capabilities of humans has prevented vagueness from posing a substantial practical problem in many domains; as technology begins to replace human operators and experts, however, computer scientists are finding themselves faced with the challenge of developing systems that can understand and reason with vague predicates. Vagueness is primarily a problem for computer systems that receive some input in natural language, a process that typically occurs when a system is designed to interact with humans or to reason on information obtained from humans. Expert systems are designed to replicate human reasoning in fields that have traditionally required a human expert. Human experts fill an expert system's knowledge base with facts and inference rules which are then used by the system to make decisions. If human knowledge could easily be expressed in the language of first-order logic then the construction of expert systems would be a simple task. Unfortunately, expert knowledge is rarely precise; instead, vague predicates are frequently used by experts to describe vague elements that are inherent to their domain. Vagueness plays a crucial role in these domains and cannot simply be replaced by precise alternatives. Vague predicates allow complex and confusing quantitative information to be quickly conveyed in an intelligible manner, an important consideration in many expert systems. For example, intelligibility was ranked as the most important factor in automated fraud detection by Lloyds/TSB [Ben00b, p. 98]. This is because the consequences of false accusations are severe and every suspect claim flagged by the system must be reviewed by a human expert who needs to understand why a particular claim was flagged.

Many medical expert systems require patient input that must be given in natural language. Though one could attempt to count the number of times one sneezes or Kleenexes used over the course of a day, it would be nearly impossible for one to input many of their symptoms in a quantitative manner and it is much more likely that one would simply say that their sinuses were 'somewhat' congested. One could attempt to remove vagueness by adding crisp rules to the system that appeared to take vagueness into account. Instead of a rule 'if ϕ's sinuses are congested then α' rules could be added that say 'if ϕ is mildly congested then α' and 'if ϕ is somewhat congested then β'. While this may be practically appealing, it is not theoretically satisfying. Instead of one crisp boundary between ϕ having sinus congestion and ϕ not having sinus congestion there are now two crisp boundaries and vagueness is still present in the rules. The Health Management System (HMS) developed at the Omron Corporation uses fuzzy logic to model vagueness; the HMS creates personalized programs designed to meet the needs of over 10,000 individuals and uses fuzzy logic to model linguistic variables so that both expert knowledge and patient input can be properly represented [Isa95, p. 65]. A fuzzy controller is presented by [GNZ93] that automates dialysis; this controller requires a nonstandard logic because there is currently no precise mathematical model of dialysis [GNZ93, p. 147]. The management of traffic, both of automobiles and people, is another popular application for controllers that use vague linguistic terms. The fuzzy controller presented in [NN00] coordinates traffic signals and directs traffic using expert knowledge obtained from a police officer. An elevator controller presented in [KKLKS98] uses vague terms like 'up-peak' and 'lunch' to define different time periods; each term represents a different set of operating conditions that are used to optimize elevator performance in each time period.

There is no precise definition for emotional words or creative expressions and it seems as though they may be vulnerable to the sorites paradox. Emotions do not crisply change and it is not the case that there is a precise number of unhappy events that are required to change a person's mood from happy to unhappy; a small unhappy event is not enough to make the difference between happiness and sadness, yet after a sufficient quantity of unhappy events there is a perceptible change in attitude. This is an important consideration for researchers who are attempting to mimic emotional responses in robots designed to interact with the public. The ability to accurately simulate emotion on a human-like face is an important step in making computer agents more accessible to the general public, particularly in recreational and educational domains where a sudden and unjustified change from happy to sad or angry could be frightening to some and ultimately undermine the robot's effectiveness [AM03, p. 3918]. A fuzzy logic system is used to model emotions on a a robot face in [AM03] and on a 3-D model in [BHPN01]; in each case, the authors use fuzzy sets to represent emotions and facilitate gradual changes.

3.2 Vagueness and the Law

Vagueness can be found throughout the legal domain: while some laws are written with precise language and intended to be applied under an exact set of circumstances, others are left intentionally vague for social or political purposes. While these types of laws may seem precise, or at least precise enough to pass into law, their interpretations are rarely clear as the situations humans find themselves in can rarely be quantitatively described. Edgington writes that social workers face this difficulty on a daily basis as they must constantly "draw the line between the acceptable and the unacceptable ... [they] come in for severe criticism for removing children from their parents' care unnecessarily and for not removing children from their parents' care when it would have been better had they done so" [Edg01, p. 371]. Even when lines are precisely drawn extenuating circumstances are usually taken into account and many people, including judges and juries, may have intuitive difficulties accepting a quantitative definition that doing something to one precise extent is illegal but doing it to some lesser extent is not. Yet if one accepts the notion of legal indeterminacy then any account of vagueness in the law must be reconciled with notions of legal integrity and the rule of law. In [End00], Endicott argues that vagueness is an essential feature of legal systems that makes the rule of law unattainable: theories that reject legal indeterminacy must either take an epistemic view, that vague predicates actually have sharp boundaries that are necessarily hidden from us, or hold that linguistic indeterminacy is eliminated by some special mechanism, such as a principle preventing the plaintiff from succeeding in indeterminate situations [End00, pp. 58–59]. This type of mechanism is hardly satisfactory, for if it is indeterminate whether a law ϕ applies in case β then it seems as though there will be cases where it is indeterminate whether it is indeterminate whether a law ϕ applies in case β.

Even when written with the best linguistic intentions, vague predicates and ambiguous words find their way into laws. A British law that allows police to order the organizers of raves to shut down their sound equipment applies to "a gathering ... at which amplified music is played during the night ... by reason of its loudness and duration and the time at which it is played, is likely to cause serious distress" [End00, p. 57]. Suppose one million rave organizers are brought in court on charges of disobeying the anti-rave law. The cases are identical to each other with the exception that the raves could be ordered in a series such that the volume played at concert r_m is imperceptibly lower than the volume played at concert r_{m+1} [End00, p. 161]. In each case, the judge will need to decide whether the volume level constituted "amplified music ... [causing] serious distress"; in other words, the judge must either decide to convict or acquit all of the organizers or choose some rave r_n such that the music at r_n was loud enough to cause serious distress and the imperceptibly quieter music at r_{n-1} was not. While the law allows graded penalties it is committed to judicial bivalence [End00, p. 73]: violators of the same crime can be assessed different penalties but they must be found either guilty or not guilty; there is no partial guilt.

While obviously vague predicates like 'tall' can be avoided when drafting laws, other words are not so easily recognized as being vague. Although laws including these words can be sharpened, there will always be borderline cases where even the precise definition of something like a vehicle will be called into question and be put up to scrutiny. Philipps gives an example where using fuzzy logic to create a legal definition would classify a car as being entirely in the set of vehicles and a toy car as not at all being in the set of vehicles, while a skateboard would be classified as both a vehicle and a piece of equipment for pedestrians and have partial membership in each set [PS99, p. 123]. Although laws that precisely define vague predicates like 'vehicle' may be easier to enforce, their arbitrariness harms the integrity of the law. One ancient Jewish law held that although one or more chicken eggs may lead to defilement, 'one egg minus one sesame seed does not defile', while another law stated that a fledgling found within fifty cubits of an aviary is considered to be property of the aviary's owner [Sha02, p. 120]. When Rabbi Yrmiya asked if a fledgling with one leg inside the fifty cubits and the other outside of it would be considered the owner's property, he was thrown out of his college for attempting to be provocative [Sha02, p. 121]. An intuitive reply by a contemporary reader would most likely, and with good reason, question the strictness of these laws as in the former case an arbitrary boundary is employed to prevent a sorites series from forming while in the latter an object is found to have uncertain boundaries. The Talmud has crisply defined a vague term such that the definition of one egg requires that each and every one of its sesame seeds be present and accounted for.

Another legal area where vagueness can be found, and where fuzzy expert systems have been proposed, is in traffic law. In [Bor99], Borgulya presents an expert system that uses fuzzy sets to represent important facts that cannot be expressed quantitatively, such as the reason for the defendant's violation of the law and the severity of the injuries sustained. Vague linguistic terms that would normally be used to describe the reason for a law's violation, such as carelessness or flagrant disregard, can be represented in this way. Extenuating circumstances can also be evaluated using this approach, and criteria such as technical defects in the vehicle can be weighted and taken into consideration when the system produces a verdict. These examples of vagueness in practice are meant to show that vagueness is a major problem in many disciplines. As a result, there is great incentive for philosophers to work together with professionals in other disciplines to develop practical and theoretically satisfying methods of dealing with vagueness.

4 Solutions to the Paradox

Solutions to the sorites have typically only been examined from a philosophical perspective. While it is necessary for any good theory of vagueness to pass philosophical scrutiny, a good account of vagueness should also be useful when applied in the many disciplines where vagueness can be found. This section will give a detailed analysis of two proposed methods of dealing with vagueness: the epistemic view, which holds that classical logic is completely compatible with what appears to be vagueness, and fuzzy logic, which introduces an infinite number of truth values to deal with the phenomenon. Both the epistemic and fuzzy account of vagueness can be seen as extreme solutions and several moderate theories have also been proposed, including one requiring a 3-valued logic and one that is known as supervaluationism. Each theory of vagueness has a unique set of strengths and offers a different view of vagueness; yet discussions of these theories usually focus not on their strengths but on their many weaknesses. Tye nicely summarizes the current state of affairs: "whichever way we turn in our attempts to understand vagueness—whether we move to the right and become arch conservatives in the manner of the epistemic theories, or we shift to the left and embrace the liberal chic of alternative logics—we quickly become enmeshed in difficulties" [Tye95, p. 18].

Like epistemicists, supervaluationists believe that classical logic can be reconciled with vagueness; unlike epistemicists, supervaluationists claim classical semantics must be rejected to accomplish the task. Defending supervaluationism, Fine [Fin75] claims that vagueness and ambiguity are closely related: because we can easily disambiguate ambiguous sentences, Fine claims, we can similarly resolve vague sentences. One way of determining the truth value of a given ambiguous sentence is as follows: if each disambiguation of the ambiguous sentence is true then the ambiguous sentence is also true; if each disambiguation of the ambiguous sentence is false then the ambiguous sentence is also false. If different disambiguations lead to different truth values then the ambiguous sentence is neither true nor false. Fine believes this method can be expanded to vague sentences using the notion of 'precisifications', which is the term supervaluationists use to refer to sharpenings of vague predicates. For example, one precisification of 'tall' may draw the borderline at $1.8m$ while another would draw the borderline at $1.6m$. According to the supervaluationist, vague statements are only true when they are supertrue and only false when they are superfalse [Wil94, p. 144]: statements are supertrue if they are true on all admissible precisifications and superfalse if they are false on all admissible precisifications. Statements that are true on some admissible precisifications and false on others are neither true nor false. Dealing with the sorites paradox, supervaluationists believe that the inductive step of a sorites series is superfalse [Bur91, p. 71] because on each precisification there is some n that marks the borderline between heapness and non-heapness. Supervaluationism cannot treat connectives truth-functionally: assume Bob is a borderline case of a bald man and let P be the sentence 'Bob is bald'. Because every precisification of 'bald' guarantees the truth of one of the disjuncts, $P \vee \neg P$ must be supertrue [Tye94, p. 192]. But neither P nor $\neg P$ can be supertrue, because if either of them were then Bob would not be borderline bald. So both 'Bob is bald' and 'Bob is not bald' may be neither true nor false even though their disjunction will be supertrue. This holds for any P, since $P \vee \neg P$ will always be a tautology but neither P nor $\neg P$ is necessarily supertrue or superfalse. Whether this is a benefit or a liability for supervaluationists is debatable; Tye [Tye94] believes this example is enough to reject supervaluationism while Williamson [Wil94, p. 147] argues that it shows the superiority of supervaluations to many-valued approaches. More information about supervaluationism and the alleged link between vagueness and ambiguity can be found in [Fin75], [Sai88], [Tye89], [Wil94] and [Sor98a].

The first person to use a 3-valued logic to represent vague sentences was the Swedish logician Sören Halldén; Halldén believed borderline vague sentences were 'meaningless' in the sense that they were neither true nor false [Wil94, p. 103]. 3-valued logics are used to advocate both truth gap and truth glut theories of vagueness: on the truth gap theory, borderline vague sentences are neither true nor false while on the truth glut theory borderline vague sentences are both true and false. A major criticism of the 3-valued approach is that there is no good way of distinguishing between cases of true vague sentences and borderline vague sentences. For example, there is no way of determining which n is such that a man with n hairs on his head is bald and a man with $n+1$ hairs is indeterminately bald. This problem, known as higher-order vagueness, can be found in every theory of vagueness. Yet it is particularly problematic for the 3-valued account because it so closely mirrors the implausible distinction that classical logic requires between a man with n hairs being bald and a man with $n+1$ hairs being not bald. Another criticism is that the 3-valued approach gives us no way of ranking truth. For example, in 3-valued logic the sentence 'a man with 100 hairs on his head is bald' is just as indeterminate as the sentence 'a man with 500 hairs on his head is bald', yet it intuitively seems as though there is something truer about the first sentence. For more information on using a 3-valued logic to represent vagueness, see [Tye90] and [Tye95]. While supervaluations and 3-valued theories present an interesting middle ground, the remainder of this section will focus the epistemic and fuzzy conceptions of vagueness.

4.1 Epistemicism

If it weren't for the strong defense given by its two main proponents, Timothy Williamson and Roy Sorensen, it is doubtful anyone would take the epistemic account of vagueness seriously. This theory holds that although every vague predicate has sharp boundaries, it is necessarily unknowable whether a particular borderline instance of a vague predicate is either true or false. The counterintuitive nature of this theory is well understood by its advocates: Sorensen describes it as his "strangest belief" [Sor01, p. 1] and Williamson describes the theory as "incredible", admitting that many people "may think that [they] cannot conceive how a vague statement could be true or false in an unclear case" [Wil94, p. 3]. Yet epistemicism may not be any stranger than alternate explanations of vagueness: many accounts of vagueness require a rejection of bivalence, a move that entails a rejection of "classical truth-conditional semantics . . . yet classical semantics and logic are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains" [Wil94, p. 186]. Epistemicists typically support their theory by "emphasizing the problems facing non-classical alternatives" [KS96, p. 18] and most of their critics agree that it offers the simplest explanation of vagueness because it allows "the linguistic phenomena which we associate with vagueness to sit perfectly comfortably alongside classical logic and semantics" [Wri95, p. 134].

Another major tenet of epistemicism is that a rejection of bivalence is incoherent: by assuming that bivalence fails we assume that for some sentence P, $\neg(t(\ulcorner P \urcorner) \vee t(\ulcorner \neg P \urcorner))$ holds; but because Tarski's disquotational schema for truth tells us that $t(\ulcorner P \urcorner) \equiv P$ and $t(\ulcorner \neg P \urcorner) \equiv \neg P$ must be true as well, we can substitute the appropriate terms and derive $\neg(P \vee \neg P)$ which, by DeMorgan's law, yields the contradiction $\neg P \wedge \neg\neg P$ [Wil94, pp. 188–189]. This proof can be criticized on many grounds: in some logics DeMorgan's laws do not hold while in others $P \wedge \neg P$ is not a straightforward contradiction; fuzzy logic allows, for example, $P \wedge \neg P$ to be only partially false. One can also criticize Williamson's use of Tarski's schema. Pelletier and Stainton point out that "if a sentence 'P' is true if and only if P, and if P and $\neg\neg P$ always have the same truth value, then of course there is bivalence" [PS03, p. 374]; in fuzzy logic the Tarski biconditionals will be partially true when P is partially true or partially false.

While the epistemicist's claim that vague predicates have hidden boundaries is difficult for many to accept, perhaps the theory's biggest challenge lies in explaining why these boundaries are unknowable to us all. Williamson tries to connect these two claims by introducing the idea of omniscient speakers who know everything relevant to the application of a given predicate in borderline cases; when asked to proceed down a sorites series until a borderline is reached Williamson claims that each speaker will necessarily stop at the same point, for if they did not then the speakers would not be omniscient as one would have knowledge that the others did not [Wil94, p. 200]. This example is intended to carry the epistemic view that our indecision over the applicability of vague predicates in borderline cases stems from ignorance: were we to have all the relevant information, as omniscient beings would, then we would know where the borderline lies. Yet Williamson's account of omniscient speakers begs the question by assuming that borderline cases can be decided solely on the basis of the speaker's knowledge and not, for example, on the basis of their opinions. Williamson's example places the onus on his critics to explain why, if vagueness is not an epistemic phenomenon, omniscient beings would give different answers to the question of where the borderline lies. Yet it should not come as a surprise that omniscient beings may not know the solutions to logical paradoxes: few people would claim that an omniscient speaker knows the truth value of 'this sentence is false' or the number of sides on a round square. Without Williamson's seemingly question-begging assumption that borderline vague cases are decided purely on the basis of knowledge, an account of omniscient speakers is completely compatible with the view that crisp borderlines do not exist. Descending along a sorites series, the sum of speaker s_1's knowledge may lead them to the opinion that the last true application of predicate F is F_n while that same knowledge may lead speaker s_2 to the opinion that the last true application of F is F_{n+5}; in the absence of epistemic principles, there is no reason to suppose that the summed knowledge of an omniscient being deductively entails knowledge of a vague predicate's borderline.

Most epistemicists offer general explanations as to why the precise borderlines are unknowable. For example, some have claimed that ignorance is the default state of knowledge and that no further explanation is warranted [KS96, p. 19]. Williamson, however, provides a more thorough explanation: he likens our ignorance in borderline cases to the kind of ignorance faced when we are presented with inexact knowledge; because we can never have reliable knowledge about borderline cases, he claims, whenever we think we have such knowledge it will be accidental or based on luck. Because these lucky guesses do not constitute reliable knowledge, Williamson claims that our knowledge is inexact and "our beliefs [will be] reliable only if we leave a margin for error"; this principle says that "'A' is true in all cases similar to cases in which 'it is known that A' is true" [Wil94, pp. 226-227]. Wright summarizes Williamson's 'margin of error' principle with the following example: "to know where the sharp boundary falls in a spectrum of patches ... [one] would have to know of some patch j that it is red while at the same time knowing that j', its neighbour, is orange. But [one] cannot know that an item, x, is red unless [one's] impression that x is red is flanked on both sides by patches that are red" [Wri95, p. 149]. On the epistemic view, the claim that j' is the first orange colour patch is just a lucky guess because j' is indistinguishable from j; were j' identical to j we would still have formed the same belief, though in this case it would have been an unlucky guess. One seemingly unfortunate consequence of this view is that we can never reliably know whether two colours are truly identical [Gol00, p. 171]. Yet this contradicts the assumption implicit in every sorites argument that there is some arbitrarily high or low number n such that it is determinately true and knowable that F_n holds. For example, it seems intuitively obvious that we know there is some height n such that T_n is definitely true. So it will be knowable that T_{n+1} is flanked on both sides by reliable truths, namely T_n and T_{n+2}. If we accept this, then we cannot apply the 'margin of error' principle to T_{n+1} and there will be little stopping us from proceeding down a sorites series; we cannot say that the principle starts applying at T_{n-1} or any other specific n because we have no way of knowing when our knowledge is no longer reliable. Yet to apply the 'margin of error' principle to T_{n+1} and deny that we reliably know its truth is absurd: there is nothing unreliable about claiming we know a man nine feet in height is clearly and determinately tall.

While Williamson's account of epistemicism has become the standard, there are other arguments that have been put forth for an epistemic theory of vagueness. Kuczynski argues for sharp cut-offs by claiming that if one is indeterminately bald then they cannot be bald, since being indeterminately bald "is a way of being outside the class of bald things"; therefore, he claims, it must be incoherent to say that something is "neither inside nor outside the class of bald things" [Kuc03, p. 154]. While Kuczynski acknowledges that some may find the assumption that something not in the class of bald things is non-bald question-begging, he dismisses this concern by claiming that it entails a redefinition of the logical operator 'not'; part of its meaning, he claims, is that if something is not an X then it must be a non-X. Yet it is not clear why his opponents would find this conclusion troublesome: every nonstandard logic redefines 'not' and while Kuczynski may reject such logics, it is difficult to take the assumption that something can be indefinitely true seriously solely in the framework of classical logic. Horwich defends an epistemic view by arguing that the claim that vague predicates lack sharp boundaries is a theoretical one that is difficult to justify; instead, the phenomenon of vagueness is best described by saying a vague predicate is one where we are unwilling to apply the predicate or its negation [Hor97, p. 930]. Horwich attacks Williamson's view that borderlines are unknowable on the grounds that it is "attributed to a special form of unreliability—an external failure—and not, as it should be, to the internal difficulty in making a judgement" [Hor97, p. 931].

The plausibility of the epistemic conception of vagueness can be further reduced by examining the notion of context-dependency: epistemicists specifically dismiss context-dependent factors as being irrelevant when determining the applicability of vague predicates. One should be careful to distinguish between the claim that context, in some cases, plays a role in vagueness and the contextualist or interest-relative conception of vagueness. Though the latter is not discussed here, it is outlined by Graff in [Gra00] and criticized by Stanley in [Sta03]. Williamson claims that vagueness and context-dependence are distinct phenomena; he justifies this distinction by claiming that precise words like 'now' depend on context while the apparent vagueness of vague predicates cannot be removed even if the context is fixed [Wil94, p. 215]. Yet Williamson gives no examples of vague predicates that do not depend on context and he gives no account of what it means to fix a given context. While it would be unproblematic for many theories of vagueness to hold that a vague predicate P is context-dependent and the vague predicate Q is not, the epistemicist is forced to reject context-dependency in all cases. Furthermore, one need not hold that all of vagueness results from context-dependency to hold that, in some cases, the vagueness of a given predicate is the result of context-dependent factors. Consider the context of 'tall': it could refer to tall men, tall bald men, tall bald rich men, tall bald rich Canadian men, tall bald rich Canadian men living in Brazil, and so on; while it is arguable whether contexts can be sharpened *ad infinitum*, it should be clear that fixing a given context is no easy task.

To Williamson, generalizations like 'every man with m hairs is bald and every man with $m+1$ hairs is not' express necessary truths. This is yet another way that the epistemic thesis is counterintuitive: on the epistemic view, the number of hairs required for a man with a small head to be considered 'bald' is the same number of hairs that men with average, large, or irregularly shaped heads require to become 'bald'. Yet the implications are far worse than this: Tom Thumb may look as though his head is covered in hair yet be biologically incapable of growing $m+1$ hairs. The epistemicist's demand for necessary truths requires that all such men are necessarily bald, despite their appearance to the contrary. Now consider a giant with $m+1$ hairs: though he is necessarily 'not bald', his hair may cover one one-thousandth of his scalp and it may be the case that every competent speaker of English would classify him as being 'bald'. The epistemicist may now claim that baldness cannot be expressed simply in terms of the raw number of one's hairs; the necessary truth in question may be 'every man with a hair density of n is bald and $n+1$ is not'. Yet this is also insufficient: suppose the head of the giant is covered in hair that is only 0.5mm long; although competent speakers of English would classify him as 'bald', his hair density could easily be greater than $n+1$. The epistemicist can keep building these context-dependent factors into potentially precise definitions of vague predicates but counterexamples will likely always be found. Thus, the epistemic view requires one to hold either that we are intuitively wrong about the truth of every counterexample, or that the necessary truths that form a predicate's definition either depend on irrelevant facts or have an unknowable form. If the truths depend on irrelevant facts then the context-dependent factors won't matter because, the epistemicist would claim, there is some other principle that decides the truth of 'bald'. If the truths have an unknowable form then both the value of cut-offs and the measurements or units that the values refer to are unknowable; that is, not only is the value m, or values $m_1 \ldots m_n$, unknowable, but it is also unknowable whether the m refers to a number of hairs, a hair density, a head shape, a hair density and ethnic background, and so on. None of these alternatives seem palatable and it seems as though the sharp boundaries that the epistemic thesis demands are just the beginning of the counterintuitive consequences the thesis is ultimately committed to.

Technical arguments aside, there seem to be many counterexamples to the epistemic conception of vagueness. In the 1950's the U.S. supreme court ordered segregated school systems to be integrated 'with all deliberate speed'; a phrase shown by internal court documents to have been intentionally chosen to avoid giving school districts precise deadlines [Tap95, p. 198]. The epistemicist needs to either deny the vagueness of 'with all deliberate speed', perhaps claiming that it is meaningless or unintentionally synonymous with a precise time limit, or explain why despite the Justices's best efforts a precise deadline was unknowingly imposed. Purportedly vague predicates can be constructed to further motivate the impossibility of sharp cut-offs. Keefe cites the example of shutters in southern France that are painted similar shades of blue; it seems implausible to say that, either when the first shutter was painted this colour or as more and more shutters were painted similar colours, the predicate 'shutter-blue' was given fixed boundaries [KS96, p. 21]. In [Dor03], Dorr describes a scenario involving two speakers hunting for fruit. The speakers communicate using only a hoot and a yelp: hoots indicate the presence of 'more' fruit and yelps indicate the presence of 'less' fruit. Dorr examines several different cases of this scenario and concludes that even when speakers have all the requisite knowledge they are still unable to find the sharp boundaries that the epistemicist requires. While these are strong objections to the epistemic account of vagueness, one must be careful when crafting analogous cases: attempting to further trap the epistemicist, Tappenden imagines a scenario where the word 'brownrate' is introduced and defined as applying if integration was completed within one year and not applying if integration was not completed within five years, with further sharpenings to be determined as cases arise [Tap95, p. 198]. Yet it is questionable whether this is a genuine example of vagueness: although there is a gap between proceeding at 'brownrate' and 'non-brownrate', the gaps are precisely drawn and one could try and say that instead of being an instance of vagueness, this is simply a case of improper definition. Tappenden claims that this objection holds the notion of 'proper definition' to too high a standard: English speakers are fully aware of when they can and cannot use 'brownrate' and competent speakers would have no trouble understanding its meaning. One can, however, claim that 'brownrate' is an instance of a precise predicate that, in some cases, neither applies nor fails to apply. The notion of something neither applying nor failing to apply is compatible with both classical logic and natural language: "for example, a horse neither applies nor *fails-to-apply* to any polygon" [LR02, p. 422] and few, if any, philosophers claim that classical logic must be abandoned in these cases.

Just as they deny genuine linguistic vagueness, epistemicists deny the existence of vague objects. While one can try to apply Williamson's epistemic principles to arguments against vague objects, Sorensen gives a different kind of argument in [Sor98b]. Sorensen tries to prove that vague objects cannot exist by proving that no object can exist without its boundaries. Sorensen presents several different methods of destroying objects and, in each case, comes to the same conclusion: once part of an object's boundary disappears, the object itself must cease to exist; because objects cannot exist without their boundaries, Sorensen claims vague objects cannot exist. Sorensen's account of boundaries seems immune to objections from fuzzy logicians: by holding that some point of an object is a part of it to degree 0.5 we are simply creating two new boundaries, the area between 1 and 0.5 and between 0.5 and 0, that are still subject to Sorensen's argument [Sor98b, p. 11]. Markosian notes, however, that there are two potential definitions of 'boundary' which Sorensen seems to use interchangeably: an outer boundary definition, where an object's boundary consists of its outer boundary points, and a vague object definition, which is based on a fuzzy set of an object's determinate and indeterminate points [Mar00]. On this view, Sorensen's argument fails because there is no one use of 'boundary' that makes the argument consistent. One advantage of applying Williamson's epistemic theory to vague objects is that it offers a potential, though counterintuitive, solution to Unger's 'problem of the many or the none': for every object there is some unknowable but precise quantity of atoms that must be present to qualify as that particular object. Thus, the epistemicist holds, the induction step of Unger's argument fails for the same reason the induction step in the arguments of linguistically vague predicates fail: because it is not the case that 'if an object is composed of n atoms and one atom is removed then that same kind of object will remain, comprised of $n-1$ atoms'. For example, an object with n atoms may qualify as a 'rock' but that same group of atoms, minus one atom, would not.

Because epistemicists deny the existence of vagueness, they fail to distinguish between cases of apparent linguistic vagueness and cases of apparent vague objects. Yet the epistemic theory cannot simply be transfered from one phenomenon to another: what is referred to as linguistic vagueness occurs on a spectrum such that one end represents property P while the other end represents property $\neg P$ and a borderline's place on the spectrum is determined by some facts, be it a necessary truth or a context-dependent combination of factors. To the epistemicist, the relationship between men on the spectrum of bald men is linear: if we are confident calling a certain man ‘bald’ then every man with less hair than him will be ‘bald’ as well and if we are confident calling a certain man ‘not bald’ then every man with more hair than him will be ‘not bald’ as well. Yet it is difficult to see how the epistemicist could claim that this relationship holds for seemingly vague objects: imagine a small table composed of some number of atoms. While we could slowly add atoms so that the table becomes so disfigured that it definitely no longer qualifies as a table, we could also add atoms in a certain pattern so that the table is simply enlarged or beautified. The generalization of the form ‘every object that is table-like and is composed of n atoms is a table and every object that is table-like and is composed of $n-1$ atoms is not’ cannot be found in this case because a linear spectrum does not exist: atoms can be added, potentially infinitely, to create larger tables; atoms can be removed, one by one or all at once, to create a smaller table or simply remove decorations from a larger one. While one may criticize the above generalization as being an incomplete definition of tablehood, no account of what criteria should be used to judge tablehood is given by the epistemicist and it is difficult to see what kind of generalization could be used in such a case. Furthermore, the epistemicist cannot account for considerations like the placement of the table or the size of the individuals for which it was designed as these are context-dependent factors that the epistemicist claims play no role in generalizations. Even in the same context atoms can be added or subtracted in different ways and, without creating a substantial and arbitrary account of tablehood, it doesn’t seem as though tablehood can be represented on a linear spectrum.

The criticisms of epistemicism that have been presented in this section are representative of the general feeling of disbelief that accompanies the epistemic view. While proponents of this view acknowledge its implausibility, the ideological divide between epistemicists and their critics goes beyond mere implausibility. To Tye, epistemicists have wrongly construed what it means to be a borderline case: “our ordinary concept of a borderline case is the concept of a case that is neither definitely in nor definitely out, [where] ‘definitely’ here does not seem to mean ‘known’ or ‘knowably’” [Tye95, p. 18]. Using the example of tiles “qualitatively identical except in color, and so alike in color for it to be a matter of indifference which we choose for any purpose”, Edgington more colourfully describes the epistemic view as “[depriving the difference between truth and falsity of its point] . . . [and serving] no purpose beyond a fetish for precision” [Edg01, p. 373]. Ludwig and Ray feel that epistemicism solves little as it manages only to “trade a puzzle for a mystery, namely, how we could not in principle discover boundary lines" [LR02, pp. 440–441]. Originally, vagueness was seen as a paradox because we allowed our intuitions to guide us when describing the sorites: it is in virtue of our intuitions that we believe its premises to be true and its conclusion false. It is not unreasonable to hold that, in the absence of overwhelming evidence to the contrary, any theory of vagueness should conform to the same intuitions that led us to label it a paradox in the first place. Even if one can overcome the *prima facie* implausibility of the epistemicist's sharp boundaries, many more implausible and counterintuitive premises are required to make the epistemic thesis compatible with vague objects and contextual factors.

4.1.1 Alleged Empirical Evidence for Epistemicism

Although proponents of the epistemic conception of vagueness readily admit that their theory is counterintuitive, some researchers have claimed to find empirical evidence to support such a view. In [BOVW99], Williamson et al. conducted the following survey: participants were divided into two groups, 'truth-judgers' and 'falsity-judgers'; while both groups were shown the same group of predicates, 'truth-judgers' were asked to give the smallest number n such that F_n was true while 'falsity-judgers' were asked to give the largest number n such that F_n was false. Each question was of the following form, with the appropriate truth condition and predicate substituted [BOVW99, p. 5]:

> When is it false to say that a man is 'tall'? Of course, the adjective 'tall' is false of very small men and true of very big men. We're interested in your view of the matter. Please indicate the greatest height that in your opinion makes it false to say that a man is 'tall'.
> It is false to say that a man is 'tall' if his height is less than or equal to ___ centimeters

Their results showed a significant gap between the values given by 'truth-judgers' and 'falsity-judgers'. For example, in the second iteration of their survey, the mean value of the minimum n needed for 'old' to apply to a given person was 76.59 years, while the maximum n needed for 'old' to not apply to a given person was 62.27 years. Although the authors admit that these results are consistent with truth gap and fuzzy logic approaches, they dismiss these theories on philosophical grounds and proceed to advance an epistemic conception of vagueness. After testing participants's responses to vague predicates, two different groups of participants were asked similar questions about average heights and ages. Like vague predicates, these predicates have boundaries that are unknown to most people; unlike vague predicates, however, they clearly do have sharp boundaries. Thus, the researchers claimed, if participants gave similar answers to these questions it would support the hypothesis that "estimates of an acknowledged, but unknown, boundary are generated in a manner similar to estimates of the true and false regions of continua associated with vague predicates" [BOVW99, p. 15]. In this second round of studies the participants did give similar answers, leading the researchers to conclude that a "psychological interpretation of vagueness rests on the assumption of a sharp but unknown boundary" [BOVW99, p. 16].

While empirical data supporting a particular theory of vagueness could make its philosophical consequences more palatable, there are several problems with the study conducted by Williamson et al. While in the question's preamble participants are asked to "indicate the greatest height" that they believe "makes it false to say that a man is 'tall'", participants are asked to write a number n in the blank space such that, for all people with a height less than or equal to n, they believe it would be false to say that a man is 'tall'. While a user following the preamble's directions will fill in the survey correctly, someone following the directions preceding the input space may not. For example, it is false to say that a man is 'tall' if his height is less than or equal to 10 centimeters but it is not the case that 10 centimeters is the greatest height that makes it false to say that a man is 'tall'; while participants did not give these kinds of answers, the 'less than or equal to' condition allowed them to input a value that may not have been what they considered to be the greatest allowable. I duplicated one part of this study to see whether or not this ambiguity affected the results. As was the case in Williamson's study, participants were divided into two groups: 'truth-judgers' and 'falsity-judgers'. Each group was asked about the same four predicates: 'tall', 'old', 'long', and 'far away'; each of these predicates was used in one of Williamson's studies. After answering all four questions participants were then asked, for each of the predicates they gave an answer to, if the number they gave was in fact the absolute greatest, or least, where the corresponding predicate would apply or fail to apply. Of the 286 respondents whose responses were used to generate the preliminary data I use in this paper, 59.4% of respondents indicated that they did not originally give such a number to at least one of the questions.

There are other problems with the study: on the view advanced by the authors, vague predicates have unknown boundaries that are allegedly real while predicates like 'average height' have unknown boundaries that are clearly real. Yet it is not unreasonable to assume that many people might think that they know, with some reasonable degree of accuracy, what these averages are; this is problematic for the authors's interpretation because it introduces an additional variable, namely whether or not the participants believed they had a concrete answer. In the case of vagueness it is doubtful that many people hold this belief: preliminary data from my own study shows that, when asked to give a crisp boundary to a vague predicate, 14% of respondents gave some sort of comment indicating their perceived inability to do so. Note that these comments were completely unsolicited: the wording of the questions posed to participants in this section of my study was identical to the wording used by Williamson et al. Although these participants did enter a value, comments like "it's all proportionate to the context of those the person being called 'tall' is being compared to. I'm afraid I have no answer, sorry. I just can't give any truthful answer" and "I think that there is more to the questions than the quantifying numbers" show that many participants felt as though they were in no position to decide on a specific height that would make a man 'tall'.

4.2 Fuzzy Logic

Perhaps the greatest advantage of fuzzy logic is its compatibility with plausible assumptions about vague predicates. In fuzzy logic, each vague predicate is represented by a fuzzy set and objects in the domain are members of that set to varying degrees. For example, if 'Susan' is a member of the set 'tall' to degree 0.6 then the sentence 'Susan is tall' will be true to degree 0.6. Proponents of fuzzy logic claim that one of its main advantages is that it can handle a broad range of phenomena: future contingents and self-referential sentences can be given a truth value of 0.5 and treated as they would be in a 3-valued logic, while a theory of vague objects can be incorporated by using fuzzy sets to represent vague objects. For example, the peak of Mt. Everest may be a member of the set 'Mt. Everest' to degree 1 while some mound of dirt near its base may be a member of the set to degree 0.5. Degrees of truth are determined by membership functions which assign a degree of membership to each of the elements in the domain. Partial membership allows fuzzy sets to represent more than just truth conditions: when represented by fuzzy sets, vague predicates act much differently than they would in classical logic.

In classical logic 'tall' would be crisply defined, for example, as a description of one who is at least 1.75 meters tall. Anyone with a height of less than 1.75 meters would be considered 'not tall', regardless of their actual height. Suppose that ϕ is 1.75 meters tall and α is 1.74 meters tall; intuitively, it seems as though either ϕ and α are both 'tall' or both 'not tall'. By letting the fuzzy set Δ represent the vague predicate 'tall', ϕ and α can both have membership in Δ and each be considered tall to a different degree. One possible membership function for Δ is:

if $height(x) < 1.4m$ then $tall(x) = 0$
if $1.4m \leq height(x) \leq 2.1m$ then $tall(x) = \frac{height(x)-1.4m}{0.7m}$
if $height(x) > 2.1m$ then $tall(x) = 1$

In this case, ϕ would be a member of Δ to degree 0.5 while α would be a member of Δ to degree 0.49, symbolizing our intuition that the vague predicate 'tall' is more applicable to ϕ than it is to α. Membership functions are not standardized; they are developed for a specific applications and must be manually tuned through trial and error. Suppose we want to use a fuzzy set to represent Tom Thumb's conception of 'tall'; furthermore, assume that Tom considers himself to be only somewhat 'tall', believes that anyone around his height is equally 'tall', and is precisely $0.4m$ in height. One possible membership function for the set is:

if $height(x) < 0.1m$ then $tall(x) = 0$
if $0.2m \leq height(x) \leq 0.38m$ then $tall(x) = \frac{height(x)-0.2m}{0.36m}$
if $0.38m \leq height(x) \leq 0.42m$ then $tall(x) = \frac{1m}{2m}$
if $0.42m \leq height(x) \leq 0.6m$ then $tall(x) = \frac{height(x)-0.24m}{0.36m}$
if $height(x) > 0.6m$ then $tall(x) = 1$

This particular membership function represents Tom Thumb's usage of 'tall'. While one would not usually want to design a membership function based on one person's use of a particular word, this example highlights fuzzy logic's ability to give an intuitively appealing account of how vague predicates act.

There are several ways that fuzzy logic can handle the sorites paradox. Consider the following membership function for F, a vague predicate:

if $n < 10,000$ then $t(F_n) = 0$
if $10,000 \leq n \leq 90,000$ then $t(F_n) = \frac{n-10,000}{80,000}$
if $n > 90,000$ then $t(F_n) = 1$

Consider a typical sorites series where P_1 is of the form $F_{100,000}$, each P_i is of the form $F_n \rightarrow_i F_{n-1}$:

$\mathbf{P}_1$: $F_{100,000}$
$\mathbf{P}_2$: $F_{100,000} \rightarrow_i F_{99,999}$
$\mathbf{C}_1$: $F_{99,999}$
$\vdots$
$\mathbf{P}_{50,002}$: $F_{50,000} \rightarrow_i F_{49,999}$
$\mathbf{C}_{50,001}$: $F_{49,999}$
$\vdots$
$\mathbf{P}_{100,000}$: $F_2 \rightarrow_i F_1$
$\mathbf{C}_{99,999}$: F_1

Given the membership function above, $\mathrm{P}_2 \ldots \mathrm{P}_{10,002}$ have a truth value of 1 and and $\mathrm{P}_{90,002} \ldots \mathrm{P}_{100,000}$ have a truth value of 0. While the truth values of $\mathrm{P}_{10,003} \ldots \mathrm{P}_{90,001}$ will depend on which definition of fuzzy implication is used, every definition takes each premise to be true to some degree that's less than 1. Note that instead of giving a single truth value to the generalization $\forall n(F_n \rightarrow_i F_{n-1})$ we are now giving different truth values to each assumption of the form $F_n \rightarrow_i F_{n-1}$.

Depending on which definition of the fuzzy implication operator is used, these truth values may lead to counterintuitive consequences. Consider the first and third definitions of fuzzy implication, $t(p \rightarrow_1 q) = t(\neg p \vee q)$ and $t(p \rightarrow_3 q) = max(1 - t(p), t(p \wedge q))$: on each of these definitions, $t(\mathrm{P}_2) = 1$, $t(\mathrm{P}_{50,002}) = 0.5$ and $t(\mathrm{P}_{100,000}) = 1$. Using this definition of truth values, sorites reasoning fails because the truth value of each conditional decreases as it nears F's borderline. Yet these definitions of fuzzy implication fail to explain why the sorites looks so convincing: originally, we were unable to reject any one conditional in a sorites series because each one seemed equally true or false. Because, on this view, P_2 and $\mathrm{P}_{100,000}$ are supposed to be twice as 'true' as $\mathrm{P}_{50,002}$, it seems as though both of these definitions are inadequate.

On the second definition of fuzzy implication, $t(p \rightarrow_2 q) = min(1, 1 - t(p) + t(q))$, the results are far more intuitive: $t(\mathrm{P}_2) = 1$, $t(\mathrm{P}_{50,002}) = 0.9999875$ and $t(\mathrm{P}_{100,000}) = 1$. On this view, P_2 and $\mathrm{P}_{100,000}$ are only slightly 'truer' than $\mathrm{P}_{50,002}$ and a better account of sorites reasoning can be given: if we take fuzzy validity to be defined in terms of preservation of truth degrees then the sorites fails because modus ponens are no longer considered to be a valid form of reasoning. If we take fuzzy validity to be defined in terms of preservation of absolute truth, then the sorites looks convincing because it is a valid argument where each premise is nearly indistinguishable in truth value; it fails because its premises, namely $\mathrm{P}_{10,003} \ldots \mathrm{P}_{90,001}$, are slightly less than completely true [Pao03, p. 365]. Nonetheless, one may still object to the fact that the different conditionals in a sorites argument are given different truth values and Paoli claims that "there can be no way to get a uniform assignment of truth degrees to the conditional [premises] of a sorites argument if we stick to Łukasiewicz's infinite-valued logic" [Pao03, p. 366]. Responding to this type of claim, Tye notes that one could try arguing that, for example, "it is never wholly true that one is tall" [Tye95, p. 12]; in this case, even sentences like $\mathrm{F}_{100,000}$ that seem completely true will always have a truth value less than 1. If one extends this view to hold that it is never wholly true that one is 'not tall' then, using the second implication operator, each step along the sorites series will have the same truth value since neither the antecedent nor the consequent will ever have a value of 0 or 1. Yet Tye notes that this reply will likely fail as one would be forced to conclude that the world's tallest man was never 'tall' enough to be 'definitely tall' and this man will be always 'not tall' to some small degree.

In [Wea00], Weatherson criticizes the ability of fuzzy logic to deal with classical equivalences of the conditional form of a sorites series's inductive step. Weatherson claims fuzzy logic yields counterintuitive results when $\neg(P \wedge \neg Q)$ or $\neg P \vee Q$ is used in place of $P \rightarrow_i Q$. Recall that in fuzzy logic $t(p \wedge q) = min(p, q)$ and $t(p \vee q) = max(p, q)$. Thus, on the second definition of fuzzy implication,

$$\begin{aligned}
t(F_{100,000} \rightarrow_2 F_{99,999}) &= t(\neg F_{100,000} \vee F_{99,999}) &= \\
&= t(\neg(F_{100,000} \wedge \neg F_{99,999})) = 1 \\
t(F_2 \rightarrow_2 F_1) &= t(\neg F_2 \vee F_1) &= \\
&= t(\neg(F_2 \wedge \neg F_1)) = 1, \text{ but} \\
t(F_{50,000} \rightarrow_2 F_{49,999}) &= 0.9999875 \neq t(\neg(F_{50,000} \wedge \neg F_{49,999})) &= \\
&= t(\neg F_{50,000} \vee F_{49,999}) = 0.5
\end{aligned}$$

As n approaches F's borderline the truth value of the conjunctive and disjunctive versions of the premises will approach 0.5 while the truth value of the conditional will always be the constant $1 - \epsilon$, a higher degree of truth. Weatherson claims:

> there's no explanation in [fuzzy logic] for why all premises of the form of [the conjunction] look persuasive. This is quite bad, because if anything [the conjunction] is *more* plausible than [the conditional]. Consider the following thought experiment. You are trying to get a group of (typically non-responsive) undergraduates to appreciate the force of the Sorites paradox. If they don't feel the force of [the conditional], what do you use to persuade them? My first instinct is to appeal to something like [the conjunction]. And if that doesn't work, I start appealing to theoretical considerations about how our use of *tall* couldn't possibly pick a boundary between *a* and *a'*. As far as I can tell, I find [the conditional] plausible *because* I find [the conjunction] plausible, and my first reaction when teaching would be to assume my students would feel likewise ... I can't imagine trying to defend [the conjunction] to an undergraduate by appeal to something like [the conditional], nor can I imagine that the reason I find [the conjunction] intuitively plausible is that I'm tacitly inferring it from [the conditional] [Wea00, p. 21]

Weatherson's criticism hinges on the claim that while the conjunctive form of the premises is more convincing than conditional, fuzzy logic gives the conditional a higher truth value than the conjunction. Yet Weatherson's claim about our intuitions is wrong: of the 248 respondents whose responses were used to generate the preliminary data from the study I conducted, 25% of people found the conditional form of the inductive step most convincing, compared to 12.1% for its conjunctive form and 8.8% for its disjunctive form. While these results cannot be used to infer specific truth degrees, it seems as though the higher truth value the conditional receives is actually a benefit, and not a liability, for fuzzy logic. While one intuition is satisfied by fuzzy logic, another is not: although 54.8% of participants surveyed claimed the disjunctive version was least persuasive, compared to 22.6% for the conjunctive version, fuzzy logic gives these forms equivalent truth values. Despite the few counterintuitive results that have been presented thus far, fuzzy logic is still of great practical value: the most persuasive form of the inductive step, the conditional, is given a higher truth value than its conjunctive and disjunctive counterparts and the difference in truth value between $F_{10,001} \rightarrow_i F_{10,000}$ and $F_{10,000} \rightarrow_i F_{9,999}$ is negligible when applied in practice.

Another common criticism of fuzzy logic is its seemingly arbitrary membership functions. For example, there is no evidence to support the notion that $F_{10,001}$ is true to the specific degree of 0.0000125. Perhaps, critics claim, the sentence '$F_{10,001}$ is true to degree 0.0000125' is itself only true to degree ϕ. Yet this is not where the questioning ends as ϕ is an arbitrary degree as well; thus, perhaps the sentence 'the sentence '$F_{10,001}$ is true to degree 0.0000125' is true to degree ϕ' is itself only true to degree Φ, and so on, seemingly *ad infinitum*. This is the problem of higher-order vagueness: given a vague sentence P, it seems impossible to claim that the sentence 'P is vague' is precise in any way as there is no quantitative definition for vagueness; thus, the sentence 'P is vague' must itself be vague. Whether or not higher-order vagueness continues *ad infinitum* depends on how one conceives of its generation, an issue that is discussed by Burns [Bur91, p. 77–80]. Nonetheless, higher-order vagueness has proved to be a difficult problem for fuzzy logicians to overcome and while its implications for the theoretical uses of fuzzy logic are great, its practical drawbacks are minimal. For example, although Williamson is able to give a detailed account of the fuzzy logician's inability to construct an appropriate meta-language [Wil94, p. 127–131] many computer systems have still successfully been built using fuzzy logic and, while the presence of higher-order vagueness may diminish the effectiveness of these systems, they seem to be far more effective than their classical counterparts.

Membership functions, typically derived by trial and error, have been described as being intrinsically vague [CPF95, p. 2] and when expert knowledge is not available there is no standard method of generating membership functions [TT01, p. 161]. Neural networks and genetic algorithms are two of the automated methods that have been suggested to overcome this problem and a good overview of the ways these methods can be combined with fuzzy logic is given by Tettamanzi and Tomassini [TT01]. These methods have been applied by many researchers: Herrmann presents a hybrid fuzzy-neural system in [Her95] that uses a neural network to extract fuzzy rules from learned data, Takagi and M. Lee use neural networks to automatically generate nonlinear membership functions in [TL93] and Hong and C. Lee give a learning algorithm that uses the results from training examples to generate membership functions and inference rules in [HL98]. Genetic algorithms have also been suggested by several researchers for determining and optimizing membership functions in fuzzy systems and several examples of this method can be found in [Kar91], [LT93], and [JMP$^+$02]. In a genetic algorithm, each solution is encoded as a gene and checked against a fitness measure to determine which gene will be used to generate the next generation of solutions. The next generation is then generated by applying genetic operators like crossover and mutation to the fittest member and the accuracy of membership functions will generally improve with each new generation.

Fuzzy logic is frequently criticized for the way its truth-functionality deals with correlations. In his infamous attack on fuzzy logic, Charles Elkan gives the example of the fuzzy rule

$$\mathit{watermelon}(x) \leftrightarrow_i (\mathit{redinside}(x) \wedge \mathit{greenoutside}(x)),$$

where m has a red inside with truth degree 0.5 and a green outside with degree 0.8 [Elk94, p. 5]. Using the fuzzy *min* operator, $\mathit{watermelon}(m)$ would be true with degree 0.5, yet the mutually reinforcing pieces of evidence imply a much greater degree of truth. Elkan's watermelon example has been criticized for being an incorrect formulation of a rule and several alternate rules have been proposed. Garcia proposes that

$$(\mathit{redinside}(x) \rightarrow_i \mathit{watermelon}(x)) \wedge (\mathit{greenoutside}(x) \rightarrow_i \mathit{watermelon}(x))$$

be used to indicate that the two predicates both contribute to the implication that x is a watermelon [Gar94, p. 23]. The linguistic rule 'if x is a red inside and x is a green outside then x is a watermelon is very true' is proposed by Jamshidi and Vadiee to represent the mutually reinforcing qualities of x's having a green outside and a red inside [JV94, p. 36]. Yet Elkan properly responds that background information, in this case the mutually reinforcing evidence that watermelons have red insides and green outsides, may not always be known and it is unreasonable to believe that correct reasoning should only occur when all of the background information is made explicit.

While Elkan's watermelon example may not be as conceptually challenging as he would like, an example presented in [Ad94] exposes one of fuzzy logic's major weaknesses. In this example, we are supposed to imagine that the plane John Doe was traveling on has crashed in a remote location and no information is available as to whether or not anyone has survived. It is then suggested that it would be reasonable to claim both that $t(\mathit{alive}(\mathit{JohnDoe})) = 0.5$ and $t(\mathit{dead}(\mathit{JohnDoe})) = 0.5$. $\mathit{alive}(\mathit{JohnDoe}) \wedge \mathit{dead}(\mathit{JohnDoe})$ would thus be true to degree 0.5 which is, of course, an absurd conclusion that contradicts our belief that such a statement should have a truth of degree 0. One potential reply is that the initial assumption $t(alive(JohnDoe)) = 0.5$ applies the semantics of probability theory to fuzzy logic: $t(alive(JohnDoe)) = 0.5$ means that John Doe is a member of the set 'alive' to degree 0.5, not that the probability of John Doe's being alive is 0.5. Entemann notes that while it is possible someone in John Doe's condition is partially alive and partially dead, that is $t(alive(JohnDoe) \wedge dead(JohnDoe)) = 0.5$, such a claim requires knowledge that the scenario explicitly rules out by saying that no additional information is available [Ent02, p. 71]. Thus, to derive the original premise that $t(dead(JohnDoe)) = 0.5$, one must accept that 'alive' and 'dead' are fuzzy predicates and the scenario must be modified to include some information about John Doe's condition, namely that he is severely injured and dying.

These replies are largely unsatisfactory as they avoid the thrust of the original objection: fuzzy logic has no method of representing ignorance. Suppose n is some arbitrarily large number for which there is no information indicating whether or not it is prime and let P be the sentence 'n is prime'. It seems as though the only option one has in fuzzy logic is to assign a value of 0.5 to both P and $\neg P$: though it may be true that there is a subjective probability of 50% that either P is true or $\neg P$ is true, this assignment is not intended to be probabilistic. Were we to assign a value greater than 0.5 to P or $\neg P$ it would imply that we had some information about the truth or falsity of P. Furthermore, one cannot simply claim that in some way P is both true and false as mathematical truths are one of the few things that nearly everyone agrees is crisply defined: numbers greater than 1 are either completely prime or completely composite and there is no way of claiming that P is partially true or partially false. This example has the same vulnerabilities as the John Doe example: the contradiction $P \wedge \neg P$ is true to degree 0.5, despite our intuitons that it should be true to degree 0; the disjunction $P \vee \neg P$ is only true to degree 0.5, despite our intuitions that it should be true to degree 1.

Advocates of classical logic see these criticisms as being reason enough to quickly dismiss alternate logics: for example, while Haack admits that "vague sentences may not be bivalent", she claims that "a division of vague sentences into three classes—true, false and neither, is liable to give results as counterintuitive as those consequent on the use of a bivalent logic"; thus, she claims, "though replacing vague by precise expressions may lead to uncertainty due to inadequacies of measuring techniques, this uncertainty does not threaten bivalence ... it seems most economical ... to regard classical logic as an idealisation of which arguments in ordinary discourse fall short, but to which they can be approximated" [Haa96, pp. 124–125]. While Haack does not endorse epistemicism, it is difficult to see how one can argue that classical logic satisfactorily approximates our use of vague predicates without appealing to some similar account of vagueness. Neither Haack nor the other proponents of classical logic have given a non-epistemic account of this alleged approximation. Furthermore, it is difficult to see where there is room for approximation in classical logic: the point of classifying arguments as being deductively valid is so that we can have confidence in the validity of an argument based on its form and not its content. Were we to allow this notion of approximation in a 2-valued logic we would be forced to use one truth value to represent both approximate and non-approximate truth. For example, 'true' would describe a sentence that is either true or is an approximation of a sentence that could be true. It is precisely because it is unacceptable to approximate arguments in a 2-valued system that fuzzy logic has been proposed to deal with vague predicates and the resulting effects of approximation. In a 2-valued logic, the effects of incorrect approximations can be catastrophic. Were we to incorrectly approximate the false sentence P as being true in a large knowledge base the amount of revisions required would be far greater than simply changing P's truth value: sound arguments with P as a conclusion would become unsound, sets of premises including P would become unsound, and the truth values of complex sentences that make up P could switch. Furthermore, the consequences of incorrect approximations are always the same, regardless of whether the approximation of P met some high degree of accuracy or whether P was a total guess; that is, the result of incorrectly classifying a borderline tall man as 'tall' is the same as incorrectly classifying Tom Thumb as being 'tall'.

While it may seem as though fuzzy logic reduces the need for precision, Tye claims that fuzzy logic merely replaces vagueness with "the most refined and incredible precision ... [so that] the result is a commitment to precise dividing lines that is not only unbelievable but also thoroughly contrary to vagueness" [Tye94, p. 191]. Tye is referring to what proponents of fuzzy logic consider to be one of its greatest strengths and what critics claim is one of its greatest weaknesses: intuitive membership functions. Because there is no formal method for generating membership functions, each application of fuzzy logic has a different tuning of membership functions. This flexibility makes fuzzy logic useful in domains where mathematical models are not available and allows membership functions to fit the needs of a particular domain; it also highlights the seemingly arbitrary nature of membership functions. Although classical logic can be criticized for imposing arbitrarily crisp boundaries on vague terms, critics claim it is difficult to see why it is more desirable to have a fuzzy membership function impose an arbitrarily fuzzy boundary. So, while it may be reasonable to classify a particular borderline tall man as being 'tall' to degree 0.5, it may be just as reasonable to classify the man as being 'tall' to degree 0.49 or 0.51. Although 'tall' is being represented by a fuzzy set, objects still have membership in that set to a crisp degree and there is rarely a good justification for assigning one crisp value over another. While this is another difficulty with fuzzy logic that seems to cripple its ability to theoretically deal with vagueness, it has little effect in practical applications where "the difference between clear truth and almost truth—between 1 and 0.99—is an insignificant difference upon which, normally, nothing hangs" [Edg01, p. 375].

Although the fuzzy truth value is just as arbitrary as the classical one, fuzzy logic allows us to minimize the impact of arbitrary decisions. If we incorrectly classify a borderline tall man as being tall to degree 0.5 when he is in fact true to degree 0.49 then, instead of being completely wrong, as we would be if we incorrectly represented him as being determinately non-tall in classical logic, we will have only misclassified the person by a hundredth of a truth degree. Furthermore, the effects of these incorrect classifications are minimized as the validity of an argument, truth of a set of premises or the truth value of dependent complex sentences would only change by a small degree. Consider the case where we originally approximate P as being true to degree 0.6 and Q as being true to degree 0.5; originally, $P \wedge Q$ is true to degree 0.5. Yet if we later discover that our approximation of Q was inaccurate and Q is, in fact, only true to degree 0.49 then the truth of $P \wedge Q$ will only decrease by 0.01, to 0.49. In classical logic, where both P and Q are originally approximated as being true, $P \wedge Q$ will be classified as being completely true. Yet if we discover that Q was incorrectly approximated then the truth of $P \wedge Q$ will change to 0 and be completely false, a major change resulting from a minor error.

Criticisms like Tye's are rooted in the belief that the classical approach of claiming a sentence P is 'true' is less precise than claiming 'P is true to degree ϕ' and, because the nature of vagueness seems to preclude precision, we should reject any approach that requires even more precision than classical logic. These criticisms are misplaced because they rely on an ambiguous use of 'precision': in classical logic, vague sentences are said to be 'imprecise' not because they lack detail but because their truth values are inaccurate and do not adequately represent their actual truth value; sentences in fuzzy logic are 'precise' not because they are intended to be an accurate representation of truth but because they are intended to give more detail about a sentence's truth. For example, the sentences 'P is true', 'P is true to degree 1' and 'P is true to degree 0.8' may each be inaccurate; in fact, they may be equally inaccurate if P is actually true to degree 0.9. Yet both 'P is true to degree 1' and 'P is true to degree 0.8' give more information about the truth of P than 'P is true' as they do not obscure our belief that P may be only partially true. Furthermore, one need not be committed to the precise dividing lines Tye attributes to fuzzy logic to believe that 'P is true to degree ϕ' is not only more informative but more accurate than 'P is true'. Without considering dividing lines it still seems as though 'Bob, who is 7 feet in height, is tall' is truer in at least some way than 'Bill, who is 5'8 in height, is tall'; while the fuzzy representation might be incorrect, it is clearly both more informative and accurate to claim that the former sentence is true to degree 1 and the latter is true to some arbitrarily smaller degree than it is to claim that both sentences are equally true. This notion is captured by Weatherson in [Wea00], where many-valued logics are rejected in favour of a nonlinear 'truer than' relation. Without appealing to a linear ordering, additional truth values or degrees of truth, Weatherson claims that the 'truer than' relation allows us to represent our intuitions that certain vague sentences are 'truer' than others. Another comparative logic is given by Paoli: in it, truth degrees need not be totally ordered, exist without a top bound, and may be either positive or negative [Pao03]. While comparative logics seem promising, it is only recently that they have been used to handle vagueness and they are limited in scope; Weatherson's logic does not allow us to claim that 'Bob, who is 7 feet in height, is tall' is 'truer' than 'Sally, who has $100, is rich'.

There are few philosophers who claim that fuzzy logic can accurately model vagueness on a theoretical level. Edgington describes fuzzy logic as being "a precise model of an imprecise phenomenon" that is only of "instrumental value" [Edg01, p. 375]. Weatherson claims that "without a compelling argument . . . for why we must suffer to be content with useful fictions like [fuzzy logic], it seems . . . at best that [fuzzy logic] is a useful waystation on the road to truth" [Wea00, p. 3]. This criticism need not be nearly as harsh as it seems: the epistemic and supervaluationist approaches are equally susceptible to criticisms and, in practice, neither of these views is nearly as useful a 'waystation' as fuzzy logic. It would be impossible to implement either of these approaches in a practical application such as an expert system and, even if they could be implemented, they would be far less effective than a fuzzy implementation. No matter how forcefully it is asked, an epistemic robot greeter will be forced to fall silent when queried about whether a given borderline tall man is 'tall'; "according to [the] staunch epistemicist ... [one does] not and cannot know where the border is [and has] no business venturing an opinion on the state ... of the fellows near that border" [Sha03, p. 57]. In practice, however, human greeters have no trouble giving useful answers to these kinds of questions. The supervaluational robot greeter will claim that a given borderline tall man is neither 'tall' nor 'not tall' and, once again, the human greeter seems far more capable than her robotic replacement. The fuzzy robot, however, reports exactly as the human greeter would: presented with a borderline tall man, the fuzzy robot reports that it has seen a borderline 'tall' man; given a slightly taller man, the robot reports that it has seen a man that is 'somewhat tall', and so on. Although when modeling vague predicates fuzzy logic is no more theoretically satisfying than classical logic or competing theories of vagueness, it seems as though arbitrarily deciding partial membership in a fuzzy set is more intuitively acceptable than arbitrarily deciding full membership in a classical set. Both intuitively and in practice, the fuzzy approach is an improvement over the way classical logic and other proposed solutions to the sorites deal with vague predicates.

5 Conclusion

Unlike many other philosophical paradoxes, vagueness and its consequences are wide ranging and can be found throughout human knowledge and natural language; its presence in non-philosophical domains is being increasingly recognized today and was first recognized over two millennia ago. The sheer quantity of domains in which vagueness can be found necessitates that any solution to the sorites paradox be examined from many different perspectives. While this requirement may be waived in favour of a theoretically satisfying account of vagueness, philosophers have yet to find such an account; in the absence of such a theory, the practical applicability of any theory should be taken into consideration. The epistemic conception of vagueness was rejected as being both philosophically unsound and intuitively implausible; yet despite its many counterintuitive claims and consequences, epistemicism and classical logic are seen by some philosophers and a majority of computer scientists as being preferable alternatives to fuzzy logic. While fuzzy logic fares no better than classical logic at modeling our theoretic conception of vagueness, the practical advantages of using it to model vagueness are numerous. Fuzzy logic offers an intuitive explanation of why the sorites paradox fails and partially matches our intuitions about which forms of the sorites are most persuasive. There have been many successful applications of fuzzy logic in computer science and fuzzy logic has made positive contributions to the way in which computer scientists deal with vagueness by helping them develop systems that can better model human knowledge and communicate in natural language. Strong theoretical objections can be given to every philosophical account of vagueness and neither the supervaluational nor epistemic account of vagueness has led to any successful applications. Yet despite the impressive practical successes of fuzzy logic, neither philosophers nor computer scientists should be satisfied with the theoretical results it yields. While fuzzy logic may simply be a waystation on the road to a satisfying theory of vagueness, its intuitively plausible results and many successful applications are clear evidence that when compared the existing theories of vagueness, fuzzy logic is, by far, the most useful waystation available.

References

[Ad94] Yianni Attikiouzel and Christopher deSilva. A better path to duplicating human reasoning. *IEEE Expert*, 9(4):14, August 1994.

[ADKP04] Bas Aarts, David Denison, Evelien Keizer, and Gergana Popova, editors. *Fuzzy Grammar: A Reader*. Oxford University Press, 2004.

[Alt01] Torin Alter. Vague names and vague objects. *Dialogue*, XL:435–41, 2001.

[AM03] S. Ansari and H. Mobahi. Fuzzy perception, emotion and expression for interactive robots. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics (SMCC04)*, volume 4, pages 3918–23, Washington, D.C., USA, October 2003. IEEE Press.

[Ari00] Aristotle. De interpretatione. In S. Marc Cohen, Patricia Curd, and C. D. C. Reeve (trans.), editors, *Readings in Ancient Greek Philosophy: From Thales to Aristotle*. Hackett Publishing Company, 2000.

[Bar82] Jonathan Barnes. Medicine, experience and logic. In Jonathan Barnes, Jacques Brunschwig, Myles Burnyeat, and Malcolm Schofield, editors, *Science and Speculation: Studies in Hellenistic Theory and Practice*, pages 24–68. Cambridge University Press, 1982.

[BE02] James J. Buckley and Esfandiar Eslami. *An Introduction to Fuzzy Logic and Fuzzy Sets*. Advances in Soft Computing. Physica-Verlag, 2002.

[Bea03] JC Beall, editor. *Liars and Heaps*. Oxford University Press: New Essays on Paradox, 2003.

[Bel77] Nuel D. Belnap. A useful four-valued logic. In J. Michael Dunn and George Epstein, editors, *Modern Uses of Multiple-Valued Logic*, volume 2, pages 8–37. D. Reidel Publishing Company, 1977.

[Ben00a] Jean Paul Van Bendegem. Alternative mathematics: The vague way. *Synthese*, 25:19–31, 2000.

[Ben00b] Peter J. Bentley. Evolving fuzzy detectives: An investigation into the evolution of fuzzy rules. In *[DSO$^+$00]*, pages 89–106. 2000.

[BHPN01] The Duy Bui, D. Heylen, M. Poel, and A. Nijholt. Generation of facial expressions from emotion using a fuzzy rule based system. In M. Stumptner, D. Corbett, and M. Brooks, editors, *14th Australian Joint Conference on Artificial Intelligence (AI 2001), Adelaide, Australia*, volume 2256 of *Lecture Notes in Artificial Intelligence*, pages 83–94, Berlin, 2001. Springer.

[Bob02] Susanne Bobzien. Chrysippus and the epistemic theory of vagueness. *Proceedings of the Aristotelian Society*, 102:217–38, 2002.

[Bor99] Istvan Borgulya. Two examples of decision support in law. *Artificial Intelligence and Law*, 7:303–21, 1999.

[BOVW99] Nicolao Bonini, Daniel Osherson, Riccardo Viale, and Timothy Williamson. On the psychology of vague predicates. *Mind & Language*, 14(4):377–93, December 1999.

[Bur82] M. F. Burnyeat. Gods and heaps. In Malcolm Schofield and Martha Craven Nussbaum, editors, *Language and Logos: Studies in Ancient Greek Philosophy Presented to G. E. L. Owen*. Cambridge University Press, 1982.

[Bur90] J. A. Burgess. Vague objects and indefinite identity. *Philosophical Studies*, 59:263–87, 1990.

[Bur91] Linda Claire Burns. *Vagueness: An Investigation into Natural Languages and the Sorites Paradox*. Kluwer Academic Publishers, 1991.

[Cic] Cicero. Academica. In *[KS96]*, pages 59–60.

[Cop95] B. Jack Copeland. On vague objects, fuzzy logic and fractal boundaries. In *[Hor95]*, pages 83–96. 1995.

[CPF95] Luis M.M. Custódio and Carlos A. Pinto-Ferreira. Agreement: A logical approach to approximate reasoning. In *Progress in Artificial Intelligence, 7th Portuguese Conference on Artificial Intelligence, EPIA '95*, volume 990 of *Lecture Notes in Computer Science*, Funchal, Madeira Island, Portugal, October 1995. Springer.

[CV00] John Collins and Achille C. Varzi. Unsharpenable vagueness. In *[Hil00]*, pages 1–10. 2000.

[D+02] Vladimir Dimitrov et al. Fuzzy logic: A framework for the new millennium. volume 81 of *Studies in Fuzziness and Soft Computing*, chapter "Understanding Society", pages 9–143. Physica-Verlag, 2002.

[Dor03] Cian Dorr. Vagueness without ignorance. In John Hawthorne and Dean Zimmerman, editors, *Language and Philosophical Linguistics*, volume 17 of *Philosophical Perspectives*, pages 83–113. Blackwell Publishing, 2003.

[DSO+00] Yasuhiko Dote, Yukinori Suzuki, Seppo Ovaska, Takeshi Furuhashi, and Rajkumar Roy, editors. *Soft Computing in Industrial Applications*. Springer-Verlag, Great Britain, 2000.

[Dum75] M. A. E. Dummett. Wang's paradox. *Synthese*, pages 325–65, 1975.

[Edg00] Dorothy Edgington. Indeterminacy de re. In *[Hil00]*, pages 27–44. 2000.

[Edg01] Dorothy Edgington. The philosophical problem of vagueness. *Legal Theory*, 7:371–78, 2001.

[Elk94] Charles Elkan. The paradoxical success of fuzzy logic. *IEEE Expert*, 9(4):3–8, August 1994.

[Emp97] Sextus Empiricus. Adversus mathematicos. In Brad Inwood and L. P. Gerson (trans.), editors, *Hellenistic Philosophy: Introductory Readings*, page 373. Hackett Publishing Company, 1997.

[End00] Timothy A. O. Endicott. *Vagueness in Law*. Oxford University Press, 2000.

[Ent02] Carl Entemann. Fuzzy logic: Misconceptions and clarifications. *Artificial Intelligence Review*, 17(1):65–84, 2002.

[Eva78] Gareth Evans. Can there be vague objects? *Analysis*, 38(4):208, 1978.

[FHK+02] Takeshi Furuhashi, Frank Hoffmann, Mario Köppen, Seppo Ovaska, and Rajkumar Roy, editors. *Soft Computing and Industry: Recent Applications*. Springer, 2002.

[Fin75] Kit Fine. Vagueness, truth and logic. In *[KS96]*. 1975.

[Fre03] Gottlob Frege. Concepts. In *[ADKP04]*, page 33. 1903.

[Gal] Galen. On medical experience. In *[KS96]*.

[Gar94] Oscar N. Garcia. Known concerns about fuzzy logic. *IEEE Expert*, 9(4):22–24, August 1994.

[GNZ93] S. Giove, M. Nordio, and A. Zorat. An adaptive fuzzy control module for automatic dialysis. In *[KS93]*, pages 146–56, 1993.

[Gol00] Laurence Goldstein. How to boil a live frog. *Analysis*, 60(2):170–78, 2000.

[Got01] Siegfried Gottwald. *A Treatise on Many-Valued Logics*. Studies in Logic and Computation. Research Studies Press Ltd., 2001.

[Gra00] Delia Graff. Shifting sands: An interest-relative theory of vagueness. In *[Hil00]*, pages 45–81. 2000.

[Haa96] Susan Haack. *Deviant Logic, Fuzzy Logic: Beyond the Formalism*. The University of Chicago Press, 1996.

[Hec98] Richard G. Heck, Jr. That there might be vague objects (so far as concerns logic). *The Monist*, 81(2):274–97, April 1998.

[Her95] Christoph Herrmann. A hybrid fuzzy-neural expert system for diagnosis. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 1, pages 494–502, 1995.

[Hil00] Christopher S. Hill, editor. *Vagueness*, volume 28 of *Philosophical Topics*. The University of Arkansas Press, Fayetteville, Arksansas, 2000.

[HL98] Tzung-Pei Hong and Chai-Ying Lee. Learning fuzzy knowledge from training examples. In Georges Gardarin, James C. French, Niki Pissinou, Kia Makki, and Luc Bouganim, editors, *Proceedings of the 1998 ACM CIKM International Conference on Information and Knowledge Management, Bethesda, Maryland, USA, November 3-7, 1998*, pages 161–66. ACM, 1998.

[Hoo90] Christopher Hookway. Vagueness, logic and interpretation. In David Bell and Neil Cooper, editors, *The Analytic Tradition: Meaning, Thought and Knowledge*, pages 61–82. Basil Blackwell, 1990.

[Hor95] Terry Horgan, editor. *Spindel Conference: Vagueness*, volume XXXIII of *The Southern Journal of Philosophy Supplement*. University of Memphis, 1995.

[Hor97] Paul Horwich. The nature of vagueness. *Philosophy and Phenomenological Research*, 57(4):929–35, 1997.

[Isa95] Satoru Isaka. Fuzzy logic applications at omron corporation. In *[LYZ95]*, pages 55–67. 1995.

[JMP+02] R. Joarder, S. Mondal, B. Podder, D. Pratihar, and Mohit. Prediction of power requirement in turning using a ga-fuzzy approach. In *[FHK+02]*, pages 167–78. 2002.

[JV94] Mohammad Jamshidi and Nader Vadiee. The promising future of fuzzy logic. *IEEE Expert*, 9(4):36–38, August 1994.

[Kar91] C. L. Karr. Applying genetics to fuzzy logic. *AI Expert*, 6(2):26–33, 1991.

[KKLKS98] C. Kim, J. Kim, H. Lee-Kwang, and K. Seong. Design and implementation of a fuzzy elevator group control system. In *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, volume 28, pages 277–87. IEEE, May 1998.

[Kos93] Bart Kosko. *Fuzzy Thinking*. Hyperion, New York, 1993.

[Kos99] Bart Kosko. *Heaven in a Chip: Fuzzy Visions of Society and Science in the Digital Age*. Harmony Books, USA, 1999.

[KS93] E.P. Klement and W. Slany, editors. *Fuzzy Logic in Artificial Intelligence, 8th Austrian Artificial Intelligence Conference, FLAI '93, Linza, Austria, June 1993*, volume 695 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 1993.

[KS96] Rosanna Keefe and Peter Smith, editors. *Vagueness: A Reader*. MIT Press, 1996.

[Kuc03] John-Michael Kuczynski. Is there non-epistemic vagueness? *Indian Philosophical Quarterly*, XXX(2), 2003.

[Llo79] G. E. R. Lloyd. *Magic, Reason and Experience: Studies in the Origin and Development of Greek Science*. Cambridge University Press, 1979.

[Low94] Jonathan Lowe. Vague identity and quantum indeterminacy. *Analysis*, 54:110–14, 1994.

[LR02] Kirk Ludwig and Greg Ray. Vagueness and the sorites paradox. In James E. Tomberlin, editor, *Language and Mind*, volume 16 of *Philosophical Perspectives*, pages 419–61. Blackwell Publishing, 2002.

[LT93] Michael A. Lee and Hideyuki Takagi. Integrating design stages of fuzzy systems using genetic algorithms. In *Second IEEE International Conference on Fuzzy Systems*, volume 1, pages 612–17. IEEE, 1993.

[Luk70] Jan Lukasiewicz. Philosophical remarks on many-valued systems of propositional logic. In O. Wojtasiewicz, editor, *Selected Works*, pages 153–78. North-Holland, Amsterdam, 1970.

[LY99] Reza Langari and John Yen. *Fuzzy Logic: Intelligence, Control, and Information*. Prentice-Hall Inc., 1999.

[LYZ95] Reza Langari, John Yen, and Lotfi A. Zadeh, editors. *Industrial Applications of Fuzzy Logic and Intelligent Systems*. IEEE Press, 1995.

[Mai00] Anna Maidens. Are electrons vague objects? In Niall Shanks and Robert B. Gardner, editors, *Logic, Probability and Science*, Poznan Studies in the Philosophy of the Sciences and Humanities, pages 261–77. Rodopi, 2000.

[Mar00] Ned Markosian. Sorensen's argument against vague objects. *Philosophical Studies*, 97:1–9, 2000.

[McG97] Vann McGee. "kilimanjaro". In Ali A. Kazmi, editor, *Meaning and Reference*, volume 23 of *Canadian Journal of Philosophy Supplementary Volume*, pages 141–63. University of Calgary Press, 1997.

[Mol69] Jon Moline. Aristotle, eubulides and the sorites. *Mind*, 78(311):393–407, 1969.

[Mor02] Michael Morreau. What vague objects are like. *The Journal of Philosophy*, 99(7):333–61, 2002.

[Nis04] Vesa A. Niskanen. *Soft Computing Methods in Human Sciences*, volume 134 of *Studies in Fuzziness and Soft Computing*. Springer, Berlin, 2004.

[NN00] Riku Nevala and Jarkko Niittymaki. Fuzzy logic for two-phase traffic signal control for coordinated one-way streets. In *[DSO+00]*, pages 107–16. 2000.

[NW97] Hung T. Nguyen and Elbert A. Walker. *A First Course in Fuzzy Logic*. CRC Press, 1997.

[O'C48] D.J. O'Connor. Pragmatic paradoxes. *Mind*, 57(227):358–59, 1948.

[Oli03] Doris Olin. *Paradox*. Central Problems of Philosophy. McGill-Queen's University Press, 2003.

[Ove84] D. E. Over. Vague objects and identity. *Analysis*, 49(2):97–9, June 1984.

[Pao03] Francesco Paoli. A really fuzzy approach to the sorites paradox. *Synthese*, 134:363–87, 2003.

[Phi99] Lothar Philipps. Approximate syllogisms - on the logic of everyday life. *Artificial Intelligence and Law*, 7(2):227–34, September 1999.

[Pri98] Graham Priest. Fuzzy identity and local validity. *The Monist*, 81(2):331–43, April 1998.

[PS99] Lothar Philipps and Giovanni Sartor. Introduction: From legal theories to neural networks and fuzzy reasoning. *Artificial Intelligence and Law*, 7:115–28, 1999.

[PS03] Francis Jeffry Pelletier and Robert J. Stainton. On 'the denial of bivalence is absurd'. *Australasian Journal of Philosophy*, 81(3):369–82, 2003.

[PtMW90] Barbara H. Partee, Alice ter Meulen, and Robert E. Wall. *Mathematical Methods in Linguistics*, volume 30 of *Studies in Linguistics and Philosophy*. Kluwer Academic Publishers, 1990.

[PW96] Terence Parsons and Peter Woodruff. Worldy indeterminacy of identity. In *[KS96]*, pages 321–37. 1996.

[Qui53] W. V. Quine. On a so-called paradox. *Mind*, 62(245):65–7, 1953.

[Qui60] W. V. Quine. *Word and Object*. M.I.T. Press, 1960.

[Ral94] Anca L. Ralescu, editor. *Applied Research in Fuzzy Technology: Three Years of Research at the Laboratory for International Fuzzy Engineering (LIFE), Yokohama, Japan*. Kluwer Academic Publishers, Boston, 1994.

[Res69] Nicholas Rescher. *Many-valued Logic*. McGraw-Hill, 1969.

[Res01] Nicholas Rescher. *Paradoxes: Their Roots, Range, and Resolution*. Open Court, 2001.

[Rin77a] David C. Rine, editor. *Computer Science and Multiple-Valued Logic: Theory and Applications.* North-Holland Publishing Company, 1977.

[Rin77b] David C. Rine. An introduction to multiple-valued logic. In *[Rin77a]*, pages 2–16. 1977.

[Rus23] Bertrand Russell. Vagueness. In *[ADKP04]*, pages 35–40. 1923.

[Sai88] R.M. Sainsbury. *Paradoxes.* Cambridge University Press, 1988.

[San94] Arturo Sangalli. Vote, vote, vote for fuzzy logic. *New Scientist*, page 51, 12 November 1994.

[Sha02] Ron A. Shapira. Fuzzy logic and its application to legal reasoning - a comment to professor zadeh. In Marilyn MacCrimmon and Peter Tillers, editors, *The Dynamics of Judicial Proof: Computation, Logic, and Common Sense*, volume 94 of *Studies in Fuzziness and Soft Computing*, pages 119–33. Physica-Verlag, 2002.

[Sha03] Stewart Shapiro. Vagueness and conversation. In *[Bea03]*. 2003.

[Sor98a] Roy Sorensen. Ambiguity, discretion, and the sorites. *The Monist*, 81(2):215–33, 1998.

[Sor98b] Roy Sorensen. Sharp boundaries for blobs. *Philosophical Studies*, 91:275–95, 1998.

[Sor01] Roy Sorensen. *Vagueness and Contradiction.* Oxford University Press, 2001.

[Sor03] Roy Sorensen. *A Brief History of the Paradox: Philosophy and the Labyrinths of the Mind.* Oxford University Press, 2003.

[Sow00] John F. Sowa. *Knowledge Representation: Logical, Philosophical, and Computational Foundations.* Brooks/Cole, 2000.

[Sta03] Jason Stanley. Context, interest-relativity, and the sorites. *Analysis*, 63:269–80, 2003.

[Tap95] Jamie Tappenden. Some remarks on vagueness and a dynamic conception of language. In *[Hor95]*, pages 193–201. 1995.

[Ter95] T. Terano. Fuzzy engineering - its progress at life and future prospects. In *Proceedings of 1995 IEEE International Conference on Fuzzy Systems*, volume 1, pages 1–8, Yokohama, Japan, 1995. International Joint Conference of the Fourth IEEE International Conference on Fuzzy Systems and The Second International Fuzzy Engineering Symposium.

[TL93] Hideyuki Takagi and Michael Lee. Neural networks and genetic algorithm approaches to auto-design of fuzzy systems. In *[KS93]*, pages 68–79, 1993.

[TT01] A. Tettamanzi and M. Tomassini. *Soft Computing: Integrating Evolutionary, Neural, and Fuzzy Systems*. Springer, 2001.

[Tye89] Michael Tye. Supervaluationism and the law of excluded middle. *Analysis*, 49:141–3, 1989.

[Tye90] Michael Tye. Vague objects. *Mind*, 99(396):535–57, 1990.

[Tye94] Michael Tye. Sorites paradoxes and the semantics of vagueness. *Philosophical Perspectives*, 8:189–206, 1994.

[Tye95] Michael Tye. Vagueness: Welcome to the quicksand. In *[Hor95]*, pages 1–22. 1995.

[Ung79] Peter Unger. There are no ordinary things. *Synthese*, 41:117–54, 1979.

[Ung80] Peter Unger. The problem of the many. *Midwest Studies in Philosophy*, 5:411–67, 1980.

[vA95] Constantin von Altrock. Fuzzy logic applications in europe. In *[LYZ95]*, pages 277–310. 1995.

[Wan96] Pei Wang. The interpretation of fuzziness. *IEEE Transactions on Systems, Man and Cybernetics - Part B*, 26(2):321–26, April 1996.

[Wea00] Brian Weatherson. True, truer, truest. *Philosophical Studies*, 200? forthcoming.

[Wil94] Timothy Williamson. *Vagueness*. Routledge, 1994.

[Wri95] Crispin Wright. The epistemic conception of vagueness. In *[Hor95]*, pages 133–59. 1995.

[Yok94] Toshihiko Yokogawa. Inference function for understanding linguistic instructions. In *[Ral94]*, pages 139–213. 1994.