Vagueness in Language: The Case Against Fuzzy Logic Revisited

ULI SAUERLAND[1]

Introduction

When I started interacting with logicians, I was surprised to learn that fuzzy logic is still a big and active field of basic research. This surprise stemmed from my experience with fuzzy logic in my own field, linguistic semantics: In semantics, fuzzy logic was explored in the analysis of vagueness in the early seventies by Lakoff [10], but has been regarded as unsuitable for the analysis of language meaning at least since the influential work of Kamp in 1975 [8], which I summarize below. Therefore, I held the belief that fuzzy logic, though it has been useful in technical applications—I once possessed a Japanese rice-cooker that was advertised to use fuzzy logic—, was not useful for the analysis of vagueness. As I have learned since, I was just ignorant of a big chunk of modern mathematical logic, and fuzzy logic is in fact a lively paradigm of research on vagueness. The lack of interaction between linguists and logicians that my experience illustrates seems to be a more general phenomenon: From the interaction with logicians, I have learned that many fuzzy logicians also do not seem to be aware of current work in linguistics on vagueness. The different attitudes towards fuzzy logic in linguistics and in logic are, of course, most likely rationally justified on the basis of the different goals of the two fields. However, historical accidents can also happen in the development of a field of science, in which case the different traditions in two fields may turn out to have no rational basis. So, which is it in the case of fuzzy logic and linguistics?

In this paper, I revisit the arguments against the use of fuzzy logic in linguistics (or more generally, against a truth-functional account of vagueness). In part, this is an exercise to explain to fuzzy logicians why linguists have shown little interest in their research paradigm. But, the paper contains more than this interdisciplinary service effort that I started out on: In fact, this seems an opportune time for revisiting the arguments against fuzzy logic in linguistics since three recent developments affect the argument. First, the formal apparatus of fuzzy logic has been made more general since the 1970s, specifically by Hájek [6], and this may make it possible to define operators in a way to make fuzzy logic more suitable for linguistic purposes. Secondly, recent research in philosophy has examined variations of fuzzy logic ([18, 19]). Since the goals of linguistic semantics

[1] I am very grateful to Sam Alxatib, Stephanie Solt, Hans Kamp, Chris Fermüller, and Libor Běhounek for comments that helped me to substantially improve this paper. I also thank the audience of the 2010 LoMoReVI conference at Čejkovice for inspiring discussion of the issues discussed here. This work was made possible by the financial support of the DFG grants SA 925/1 and SA 925/4, the latter within the ESF Eurocore LogiCCC project VAAG. All remaining errors are of course my responsibility.

seem sometimes closer to those of some branches of philosophy of language than they are to the goals of mathematical logic, fuzzy logic work in philosophy may mark the right time to reexamine fuzzy logic from a linguistic perspective as well. Finally, the reasoning used to exclude fuzzy logic in linguistics has been tied to the intuition that *p and not p* is a contradiction. However, this intuition seems dubious especially when *p* contains a vague predicate. For instance, one can easily think of circumstances where *'What I did was smart and not smart.'* or *'Bea is both tall and not tall.'* don't sound like senseless contradictions. In fact, some recent experimental work that I describe below has shown that contradictions of classical logic aren't always felt to be contradictory by speakers. So, it is important to see to what extent the argument against fuzzy logic depends on a specific stance on the semantics of contradictions. In sum then, there are three good reasons to take another look at fuzzy logic for linguistic purposes.

After recognizing that the argument against fuzzy logic in linguistics needs to be updated, I provide such an update in this paper. I conclude, however, that the conclusion reached by Kamp [8] and others still stands: For capturing vagueness in linguistic semantics, fuzzy logic or any other purely truth-functional theory doesn't provide a suitable mathematical formalism. As I already mentioned, this conclusion of course doesn't say anything about the intellectual merits of fuzzy logic in itself, but only shows that the interests of fuzzy logicians and linguists diverge: Linguists are interested in formulating specific formal models that can be used to account for linguistic phenomena. Logicians explore classes of formal models that have interesting mathematical properties and sometimes also have applications.

The paper is structured as follows. In the first section, I start off with summarizing work by Kamp [8] and others that effectively dissuaded linguists from looking at fuzzy logic for the last 35 years. As I show, the assumption that contradictions are contradictory has been central to this argument. In Section 2, I then examine the argument in the light of recent developments. Specifically I summarize recent experiments on the acceptability of contradictions. The new experimental study I report tests whether any truth-functional, fuzzy logic could provide an account of linguistic vagueness if the acceptability of contradictions is integrated into linguistic semantics. As I report, the study shows that a truth functional, fuzzy logic alone could not model vagueness in language successfully. In the conclusion, I mention formal theories of vagueness that seem better suited for the account of vagueness in language in light of the data in Section 2.

1 The classical argument

The argument that I summarize in this section, has been made by several people and I am actually not sure of its origin. Most influential in linguistics have been the presentations by Kamp [8] and Fine [4]. Kamp points out that the argument can be already found in work by Rescher [16]. But, Rescher's presentation seems to not have had the same effect as the later ones in dissuading linguists from applying fuzzy logic, since Lakoff still applied fuzzy logic in linguistics in 1973 [10], while I am not aware of any serious applications since then: There has been considerable work on vagueness in linguistics since 1973, just little of it makes use of fuzzy logic.[2] I summarize here mostly the work

[2]Fuzzy logic has been discussed by linguists more recently 1975, but only to dismiss its account for vagueness. There is also some work within fuzzy logic that addresses natural language (e.g. [15, 3], which I was

of Kamp [8] and to a lesser extent that of Fine [4].

In fuzzy logic, propositional expressions have a numerical value between 0 and 1 as their meaning, and negation and conjunction are interpreted as functions mapping a single such numerical value for negation or a pair of them for conjunction to another value between 0 and 1. Kamp [8, p. 131] raises the question what the truth value of conjunction should be if each conjunct has truth value $\frac{1}{2}$. I quote the relevant paragraph in full:

> What value would F($\wedge$) assign to the pair of arguments $(\frac{1}{2}, \frac{1}{2})$? It is plausible that the value should be $\leq \frac{1}{2}$. For how could a conjunction be true to a higher degree that one of its conjuncts? But which value $\leq \frac{1}{2}$? $\frac{1}{2}$ seems out because if $[\![\phi]\!]^{\mathscr{M}}_{M} = \frac{1}{2}$ then, if we accept our definition of $F(\neg)$, $[\![\neg\phi]\!]^{\mathscr{M}}_{M} = \frac{1}{2}$. So we would have $[\![\phi \wedge \neg\phi]\!]^{\mathscr{M}}_{M} = \frac{1}{2}$, which seems absurd. For how could a logical contradiction be true to *any* degree? However, if we stipulate that $(F(\wedge))(\frac{1}{2}, \frac{1}{2}) = 0$, we are stuck with the even less desirable consequence that if $[\![\phi]\!]^{\mathscr{M}}_{M} = \frac{1}{2}$, $[\![\phi \wedge \phi]\!]^{\mathscr{M}}_{M} = 0$. And if we choose any number between 0 and $\frac{1}{2}$, we get the wrong values for both $\phi \wedge \neg\phi$ and $\phi \wedge \phi$. [8, p. 131]

A few explanatory notes may be necessary to understand fully Kamp's presentation. Kamp assumes a two stage formalism that assigns to sentences of ordinary English their semantic value following Montague [13]. The semantic value is assumed by linguists to form the basis for semantic judgments of speakers like whether a sentence is true in particular scenario or whether one sentence entails another. Since Montague was a mathematician trained in model-theory, the semantic evaluation procedure bears many similarities to Tarskian model theory. In particular, the scenario in which a sentence is evaluated is represented as a set of parameters of the evaluation procedure. In Kamp's formal system these are the indices M and $\mathscr{M}$ that occur as super- and subscripts on the evaluation function $[\![—]\!]$. Since Kamp adopts Montague's two stage interpretation, the natural language expression *and* is first translated to the formal symbol $\wedge$, which is then interpreted model-theoretically. F is the function that assigns the value to logical constants in the second stage of the semantic evaluation procedure. For conjunction, the first question in the quotation is really about how the natural language expression *and* is interpreted, though $\wedge$ is also the symbol used also by logicians to represent conjunction in a logical system. Similarly, Kamp assumes that natural language negation (*not* or *it is not the case that*) is translated by the symbol $\neg$ in the intermediate representation, but here Kamp assumes that F($\neg$) maps p to $1 - p$. With these notes, Kamp's argument can be fully appreciated.

Kamp's argument is independent of the specific formalism he adopted, some of which like the two-stage interpretation procedure are no longer used in the field. In essence, the narrow point is that for sentences of the form *A and not A* we would want the semantic value 0, but for sentences of the form *A and A* we would want the semantic value of A itself. But, if there is an A with semantic value 0.5 and if the negation of A

unfortunately unable to access while writing this paper. But, note that these papers weren't published in linguistic journals, but in a fuzzy logic journal, so I ask to be forgiven for disregarding them. Furthermore, the proposal of [11] resembles fuzzy logic, but Lasersohn doesn't address standard logical expressions like negation in his account.

then also had semantic value 0.5, it is impossible to provide a semantic value for *and* that has the desired results. So, three concrete assumptions underlying the argument are that 1) there is a formula A with semantic value 0.5, 2) that for A, $\neg A$ also has semantic value 0.5, and 3) that $A \wedge \neg A$ has truth value 0.

As I already mentioned, Kamp points out that the argument is not original to his work, but can be already found in [16]. Furthermore, at the same time as Kamp, Fine [4] presents essentially the same arguments against multi-valued approaches to vagueness. Specifically, Fine (p. 270) also proclaims that *'Surely P & $-P$ is false even though P is indefinite.'*[3] This quote also makes clear that the contradictory nature of a sentence of the form A *and not* A is essentially self-evident for Fine just as for Kamp—a point, we will come back to in the next section.

Beyond the specifics of the presentation of the argument, there's also a broader question that the argument highlights: whether a truth-functional logical system is appropriate for capturing vagueness in linguistics. In a subsequent paragraph of his paper, Kamp summarizes his argument as follows: *'the truth value of a complex formula—say $\phi \wedge \psi$—should depend not just on the truth values of the components—i.e. ϕ and ψ—but also on certain aspects of these formulae which contribute to their truth values but cannot be unambiguously recaptured from them.'* This broader question is based on one of the central concerns of linguistic semantics: Compositionality. Linguists are very interested in the fact that we humans can combine the meanings of individual words to yield sentence meanings, and want to account for people's linguistic intuitions for sentences such as A *and* A or A *and not* A as a function of the semantic value of A. The broader question is whether this can be done truth-functionally if A contains a vague expression. After investigating critically the three concrete assumptions Kamp used in his presentation in the following, we will return to this broader question, and will present empirical evidence that truth-functionality does not obtain in language when vague predicates are involved.

1.1 Initial discussion

The argument by Kamp and Fine discussed in the previous section has been widely accepted in linguistics: As already mentioned, I am not aware of serious work on vagueness in linguistics in the last 35 years that is based on fuzzy logic. Nevertheless, I mentioned reasons for reexamining the argument at this point of time in the introduction already. From my perspective, the main motivation to reexamine the argument is evidence that speakers don't perceive contradictions of classical logic to be contradictory when it comes to borderline cases. The acceptance of borderline contradictions may be taken to argue that such contradictions don't have semantic value 0 as assumed by Kamp and Fine. In the following section, I investigate a truth-functional theory that assigns to borderline contradictions a semantic value greater than 0. However, this move is by no means necessary: It may also be that contradictions have semantic value 0, but that our intuitions are affected by factors other than the semantic value. Semantic values are after all only a theoretical concept within a complex semantic-pragmatic-psycholinguistic theory of human intuitions about sentence meaning. In the remainder of this section, I argue that if contradictions are taken to have semantic value 0, the argument Kamp and Fine presented still goes through even given developments in fuzzy logic since 1975.

[3] The &-sign represents conjunction and the $-$-sign negation in Fine's article.

First, one may ask whether the sentences used in the argument shouldn't just be excluded: The relevant sentences for the argument are of the form *A and not A* on the one hand and *A and A* on the other in the presentation of Kamp [8]. Such conjunctions can be excluded by design in technical applications of logic like my rice-cooker, but this cannot be done in linguistic semantics: They are both instantiated by grammatical sentences and native speakers of English do have judgments on their truth or falsity. In the case, of *A and A* the linguistic status perhaps warrants some further comment since utterances such as *It's raining and raining* or *I am bored and I am bored* are likely to be very rare. However, to explain the rarity of *A and A*-type utterances, we have to rely on the assumption that *A and A* has the same semantic value as *A*. Then it follows from widely accepted linguistic theories such as the pragmatics of Grice [5], that the extra effort of saying *A and A* rather than just *A* should not be undertaken by speakers unless there is a specific rhetorical effect speakers aim to accomplish by the repetition. So, Kamp is justified to assume that *A and A* has the same semantic value as *A*. In sum, a semantic theory that does not provide semantic values for *A and not A* as well as *A and A* would be incomplete from a linguistic perspective.

Second, one may ask whether *and* couldn't be ambiguous. Lexical ambiguity is after all a frequent phenomenon in natural language, and if there were two words *and* that could overcome Kamp's problem. In the case of *A and not A*, we could assume that we make use of the word and_1 that assigns to the pair $(0.5, 0.5)$ the semantic value 0. When we are looking at *A and A*, however, we make use of and_2 and we get 0.5 for the pair $(0.5, 0.5)$. But, this account would clearly be *ad hoc* and would need to be supplemented with an explanation of why and_1 cannot be made use of in *A and A*, while and_2 cannot be used in *A and not A*. In fact, for *A and B* probably a third and_3 would need to be postulated when B is neither A nor $\neg A$.

Third, fuzzy logic seems to have undergone tremendous progress. As an outsider, I cannot fully appreciate all the changes since the time of Kamp's and Fine's writing (1975). From my superficial understanding, it seems one change has been that different versions of the operators representing conjunction and negation are being considered in systematic ways. Since the argument as presented by Kamp and Fine relies on assumptions about the interpretation of conjunction and negation, it is worthwhile to look whether novel variations of fuzzy logic can overcome the problems set out by Kamp and Fine.

I use two conventions in the following three paragraphs: 1) I use the abbreviation V for the set of sentences with semantic value other than 0 and 1. 2) I use lower case letters as variables ranging over $[0, 1]$.

I furthermore strengthen one assumption concerning semantic value 0 in addition to the assumption that classical contradictions have truth value 0. Namely, I assume in addition that, if a complex formula F has truth value 0, then either a basic formula occurring in F has truth value 0 or 1 or F is a classical contradiction. So, for example $a \wedge \neg b$ may only have truth value 0 if either a is 0, b is 1, or a is equal to b. This assumption is necessary for the following argument because, for example, a fuzzy logic that assigns to $a \wedge b$ 1 if and only if $a = b = 1$ and 0 otherwise doesn't lead to linguistic problems of the type Kamp and Fine point out–it behaves just like classical logic with respect to classical contradictions of the type Kamp and Fine discuss. As Fermüller (p.c.) points out, though, the assumption I just introduced excludes two widely discussed types of fuzzy logic; namely, the Gödel-logic and product logic, which both define the negation of any positive truth value as 0 [7]. However, systems where *tall* is vague, but derived predicates such as *not tall* or *short* are not, are of lesser interest for linguistic purposes.

As I understand from Hájek's non-technical summary [7], modern fuzzy logic generalizes the conjunction to be any binary operator $*$ on $[0,1]$ that is a *T-norm*, which is defined as a commutative, associative, non-decreasing, (i.e. $a \leq b$ and $c \leq d$ then $a * c \leq b * d$) function where 1 is the unit element (i.e. $1 * a = a$).[4] As we saw above, Kamp already leaves it open how to interpret conjunction *and*. His point is that there can be no single function resulting in the correct interpretation for both *A and not A* as well as *A and A*. For negation, however, Kamp only considers one possibility, namely the Łukasiewicz negation defined by $\neg a = 1 - a$. In Hájek's work, negation is assumed to be a any non-increasing (i.e. if $a \leq b$, then $\neg a \geq \neg b$) function with $\neg 0 = 1$ and $\neg 1 = 0$.[5] Obviously, the three concrete assumptions of Kamp's argument mentioned above are not guaranteed in such a more general system, specifically there maybe no A with $[\![A]\!] = 0.5$ or it may be that $\neg 0.5 \neq 0.5$. In the next two paragraphs, I show that Kamp's argument still goes through with this more general notion of negation.

First, assume that negation has a fix-point that is instantiated by a sentence A, i.e. there is a sentence A in V such that the semantic values of A and $\neg A$ are identical. Then it follows that $A \wedge \neg A$ must have the same semantic value as $A \wedge A$. But since only $A \wedge \neg A$ is a classical contradiction, this entailment contradicts our assumption that, of all complex sentences, all and only classical contradictions have semantic value 0. So, it cannot be that case that negation has a fix-point that is instantiated by a sentence. Note that this is already somewhat unexpected since negation represents some cognitive operation humans perform on the number scale. Intuitively it seems such an operation should be continuous and therefore should have a fix-point by Brouwer's theorem. Then the fix-point must not be instantiated by a sentence—i.e. be a gap in the space of possible truth-values. Concretely that entails that the mapping from a X to the semantic value *'A guy of height X is tall'* cannot be continuous. But, that runs counter an intuition underlying applications of fuzzy logic: that vagueness is captured by a continuous mapping of some perceptual quality (here: height) to truth values: intuitively, there shouldn't be any jumps in semantic value as height goes up.

Secondly, for argument's sake assume only that there is at least one pair A and B of logically independent (i.e. neither sentence entails the other and they also don't contradict each other) sentences in V with the same semantic value. This assumption is different from the previous one that A and $\neg A$ have the same truth value. Technically even the new assumption is not necessarily satisfied since the set of sentences of a language is only countably infinite while the set $[0,1]$ has a higher cardinality. Hence it may be technically possible to define a mapping from sentences to $[0,1]$ where no two sentences are mapped to the same value and Hájek's axioms are satisfied. However, linguistically

[4] In addition, continuity (or at least, continuity in the left argument) is often required of $*$. But this doesn't play a role in the following.

[5] Modern fuzzy logic often assumes additional constraints on negation, specifically concerning its interaction with conjunction. Since these are not important for the following argument though, I will not discuss them at this point.

the application of fuzzy logic wouldn't make any sense unless some generalization is achieved. And a generalization is only achieved if the initial assumption is satisfied: there are logically independent sentences A and B with the same truth value. But then it follows that $A \wedge \neg A$ must have a different truth value from $B \wedge \neg A$ since the former is a classical contradiction with semantic value 0, while the latter isn't, and therefore by assumption has a semantic value greater than 0.

In sum, this section showed that, as long as we maintain that of the complex sentences made up from basic ones that don't have truth value 0 nor 1, only the contradictions of classical logic have semantic value 0, fuzzy logic doesn't provide a useful model for vagueness in language. The argument was essentially already presented by Kamp, Fine and others more than 35 years ago as I summarized here. In the last few paragraphs, I provided some generalizations of this argument to cover more recent developments in fuzzy logic. As I have noted occasionally, my argument concerns linguistic utility not mere logical possibility: It is could be technically possible to define functions $*$ and $\neg$ that satisfy the axioms of fuzzy logic and don't lead to the contradictions noted above. But such functions if they exist would be linguistic monstrosities, and therefore I don't consider them seriously. However, there is still a serious candidate which I consider in the next section: a semantics based on fuzzy logic where classical contradictions do not necessarily have truth value 0.

2 The dialethic argument

The most controversial aspect of the argument summarized in Section 1 is the assumption that sentences *A and not A* should have semantic value 0. As I already quoted above, Kamp [8] and Fine [4] regarded it as evident that classical contradictions should have semantic value 0. Fine pronounces another semantic value to be *absurd*, while Kamp states that *surely* 0 must be the semantic value of a contradiction. Since the relevant sentences are contradictions in classical logic, the certainty of Kamp and Fine seems initially justified.

However, linguistic intuitions often stand in conflict to classical logic. And as noted above already, sentences such as *'What I did was smart and not smart.'* or *'Bea is both tall and not tall.'* sound not as non-sensical as classical logic would predict—following Ripley [17] I refer to such examples as *Borderline Contradictions*. In this section, I first summarize two pieces of experimental evidence from recent papers that show that people do find such borderline contradictions quite acceptable. This seems to undercut the argument against fuzzy logic as presented in the previous section (though see [9] for a defense). So, it raises the question whether actual linguistic intuitions on contradictions are in line with fuzzy logic. This question I address with a new experiment in the second subsection below. Specifically, I set up a comparison between $A \wedge \neg A$ and $B \wedge \neg A$, where B is logically unrelated to A, but has the same semantic value. Fuzzy logic predicts the contradiction $A \wedge \neg A$ and the conjunction $B \wedge \neg A$ to have the same semantic value if A and B have the same semantic value because of $\wedge$'s truth-functionality. But, my experimental results indicate that $A \wedge \neg A$ is significantly more acceptable than $B \wedge \neg A$. So I conclude that the amount of agreement to borderline contradictions is so high, that it cannot be accounted for directly by fuzzy logic (nor by classical logic).

2.1 Experimental evidence for dialetheism

In this section, I summarize two recent experiments, one by Alxatib and Pelletier [1] and the other by Ripley [17], that show that many speakers indeed regard sentences of the form *A and not A* as informative and accept them under the right set of circumstances.

Alxatib and Pelletier's experiment Alxatib and Pelletier [1] primarily set out to argue against the account of truth value gaps by Bonini *et al.* [2]; however, their result is also of interest for our purposes. Their result confirms the intuition that borderline contradictions are most acceptable for borderline cases.

Alxatib and Pelletier [1] made subjects evaluate sentences as descriptions of a picture. They presented subjects a picture of five men drawn in the style of an American police line-up. Behind the men, a height scale is visible so in addition to their relative height, also the absolute heights are known. The five men are $5'4''$, $5'7''$, $5'11''$, $6'2''$, and $6'6''$ tall. Alxatib and Pelletier asked their subjects to judge sentences as true or false given the picture, and offered also a third, *can't tell* response choice. In addition to the sentences *'X is tall'* and *'X is not tall'*, they also included the sentence (1), which is of primary interest to us:

(1) *X* is tall and not tall.

Sentence (1) is, of course, a contradiction in classical logic. However, Alxatib and Pelletier report that 44.7% of their subjects judge (1) to be true for the borderline case, the $5'11''$ man, and only 40.8% judge (1) to be false in that case. The amount of agreement is lower for all the other heights.

It is noteworthy that the subjects in Alxatib and Pelletier's study regard sentence (1) to be true more frequently than either of its conjuncts: This is clearly the case for *'X is not tall'*, which is judged true by 25.0% and false by 67.1% of subjects for the $5'11''$ tall man. *'X is tall'* is judged true by 46.1% and false by 44.7% for the $5'11''$ tall man. So, *X is tall* is judged true more frequently than (1), but it is also judged false more frequently. By comparing the percentage of agreement among those subjects indicating either agreement or disagreement (i.e. not the *can't tell*-option), Alxatib and Pelletier argue that (1) at 52.3% is more frequently judged true than *'X is tall'* at 50.8%.

Ripley's experiment Ripley [17] reports an experiment he conducted on contradictory sentences. Ripley's experiment uses graded acceptability judgments rather than just a binary true/false-choice. Ripley also confirms the intuition that borderline contradictions are most acceptable for borderline cases, and furthermore shows that this intuition holds across the following four different ways to express a borderline contradiction.

(2) a. The circle is near the square and it is not near the square.
b. The circle is and is not near the square.
c. The circle neither is near the square nor is not near the square.
d. The circle neither is nor isn't near the square.

Participants were assigned to one of these four conditions, so that each subject would only see one of the sentences in (2). They were then asked to indicate for seven different pictures whether they agreed with the sentence. Specifically, they were asked to indicate agreement on a scale from 1 to 7 (disagree to agree).

Ripley classifies subjects by distinct response patterns. The most frequent pattern Ripley refers to as the hump pattern, and describes it as follows: The maximum degree of agreement is indicated for one of the middle pictures, and agreement slopes down monotonically to both sides. This pattern is exhibited by more than 50% of the subjects in Ripley's experiment and corroborates the Alxatib and Pelletier finding just reported that borderline contradictions are most acceptable for borderline cases. Ripley describes furthermore that only about 10% of subjects indicate uniform disagreement to the contradictory statement in all conditions. This shows that Kamp's and Fine's intuition that classical contradictions should *surely* be judged wrong are not reflected directly in most speakers' intuitions, but only in the judgments of a small minority.

Moreover, Ripley reports the average amount of agreement: The maximum amount of agreement is 4.1, slightly above the midpoint of Ripley's 1 to 7 scale. The picture with this amount of agreement is one of the middle pictures, and the average amount of agreement also exhibits the hump patterns that the majority of individuals exhibit.

Finally, Ripley compares the four different conditions in (2) that he used in his experiment. Ripley reports that in his study there was no significant difference between the versions a and c without ellipsis and the b and d versions with ellipsis. This result, as Ripley points out, is important in evaluating proposals that trace back contradiction tolerance to an ambiguity in the vague term involved like the one of Kamp and Partee [9]. A second comparison, Ripley reports is between the conjunctive version a and b and the disjunctive ones in c and d. In this comparison, as well, the differences are not significant.

Discussion Both summarized papers confirm the intuition that borderline contradictions are generally found to be quite acceptable. The highest degree of agreement is found for the borderline cases. We could probably still attribute this agreement to classical contradictions entirely to pragmatics (*pace* Ripley [17]), but *prima facie* it provides support for a non-classical logic where at least some contradictions are acceptable. If we accept that, the argument against fuzzy logic presented in Section 1 doesn't go through any longer.

Does this mean that now we can adopt fuzzy logic for language interpretation? While this is mostly an open question, there is some indication in the data above already that indicated that borderline contradictions are actually still a problem for fuzzy logic. Not because their semantic value is 0 as in the classical argument, but because now their semantic value is higher than predicted by fuzzy logic. Specifically, I noted that agreement to $A \wedge \neg A$ is greater than agreement to either A and $\neg A$ in Alxatib and Pelletier's results, though only marginally so. At the same time, I would not expect agreement to a conjunction $A \wedge B$ to be higher than that of either conjunct (i.e. clear non-monotonicity) unless A and B stand in some logical or pragmatic relation to each other. In the following section, I report on a new experiment that confirms this intuition.

2.2 Dialetheism and fuzzy logic: an experimental test

My main goal in this section is to examine, whether fuzzy logic is suitable for linguistic purposes if dialetheism is accepted as the basis for data like those reported in the previous section. The experimental results by Alxatib and Pelletier [1] and Ripley [17] show that the argument against fuzzy logic for linguistic semantics by Kamp [8] and Fine [4]

pair	A/B	sentence
1	A	A $5'10''$-guy is tall.
	B	A guy with \$100,000 is rich.
2	A	A car driving at 70 mph is fast.
	B	A 250 page book is thick.
3	A	A 83 degree Fahrenheit day is hot.
	B	A town 45 miles away is far.
4	A	A 2 hour flight is long.
	B	A 50 year old guy is old.
5	A	A 3280-foot mountain is high.
	B	A 10 day vacation is long.

Figure 1. Basic Borderline Statements.

is less clear-cut that it was previously made out to be. Given that there is some experimental support for dialetheism that option needs to be considered, and Kamp and Fine failed to do so.

A crucial prediction fuzzy logic makes is truth-functionality of conjunction. I.e. if a and b have the same truth value, $a \wedge \neg a$ and $b \wedge \neg a$ have the same truth value too. However, intuitively this is not how borderline contradictions seem to behave in language: If we are talking about a $5'10''$-guy with kind-of-blondish hair, I find *'He's a tall blond guy.'* much more difficult to swallow than *'He's both tall and not tall.'* Here *tall blond* may receive a non-intersective interpretation, but imagine being in New York and then compare *'Boston is large and far away.'* with *'Boston is large and not large.'*—the latter seems more acceptable to me. I now describe an experiment I conducted to confirm this intuition experimentally. Specifically, my goal was to investigate the case where A, B, $\neg A$, and $\neg B$ are agreed to at the 50% level. My expectation was that the borderline contradictions $A \wedge \neg A$ and $B \wedge \neg B$ would be more acceptable than the crossed conjunctions $A \wedge \neg B$ and $B \wedge \neg A$.

Pre-test To confirm the intuitions just mentioned experimentally, the main problem is that the ideal test would require two sentences judged true to exactly the same degree. But it is not realistic to find such sentences. I decided to settle for sentences in the 40% to 60% range. I decided to start with using a greater set of sentences than just two, and only use those in the analysis that satisfy the identity of truth value requirement. So, we need several statements at the borderline. To determine suitable borderline values, I conducted a pre-test. The test was conducted over the internet using the Amazon MTurk platform. 50 subjects participated and each received 10 cents for the participation. The subjects read a short instruction, part of which told them people are often not strictly logical. Then subjects were asked to answer twelve questions such as *How tall is a guy who is tall and not tall? (in inches)*. It was assumed that the median response would represent a borderline value for the respective predicate. The sentences shown in Figure 1 were constructed from the result as basic borderline propositions. For example, the median response to the question mentioned above was 70. From this the basic borderline proposition in the first line (i.e. pair 1, sentence A) of Figure 1 was derived: *'A $5'10''$*

ba-	A	A $5'10''$-guy is tall.
sic	B	A guy with \$100,000 is rich.
nega-	$\neg A$	A $5'10''$-guy isn't tall.
tions	$\neg B$	A guy with \$100,000 isn't rich.
contra-	$A \wedge \neg B$	A $5'10''$-guy is tall and a guy with \$100,000 isn't rich.
dictions	$B \wedge \neg A$	A guy with \$100,000 is rich and a $5'10''$-guy isn't tall.
conjunc-	$A \wedge \neg A$	A $5'10''$-guy is tall and a $5'10''$-guy isn't tall.
tions	$B \wedge \neg B$	A guy with \$100,000 is rich and a guy with \$100,000 isn't rich.

Figure 2. Sample conditions used in the main experiment (pair 1).

(i.e. 70 inches) guy is tall.'. In Figure 1, the propositions are furthermore grouped into pairs of A and B, which play a role in the construction of the main experimental items as described in the following section.

Subjects, method, and materials quad The main experiment also was conducted over the internet using the Amazon MTurk platform. 100 subjects participated, and the only selection criterion used was that their HIT-acceptance rate as reported by the Amazon system be greater than 50%.[6] Subjects' gender, native language, and country of residence were reported, but not used as exclusion criteria.

The experiment used eight conditions per pair of items, which are illustrated in Figure 2. The example conditions in 2 are drawn from the pair 1 of borderline values of Figure 1. The other conditions were created in the same manner out of each of the subsequent pairs of lines of Figure 1. The conditions were created to test both the agreement to individual conjuncts, as well as to the conjunctions. In particular, both the borderline assertions A and B and their negations $\neg A$ and $\neg B$ were tested because the comparison of the conjunctions would be relevant for my goals only if all four of these received truth value 0.5. The crucial comparison then is between the two classical contradictions *AnA* and *BnB* on the one hand and the non-contradictions *AnB* and *BnA* on the other. Note that the conditions were designed so as to make sure that all conditions would have a negation in the second conjunct. This was done so that all four of these conditions would have comparable grammatical and logical complexity.

Five different questionnaires were created from the experimental items. This was done to make sure that each subject would see each basic item A or B only once in the experiment, either on its own or as part of a more complex sentence. Recall that there were five pairs (A, B) of the basic borderline assertions created from Figure 1. So, then for example one questionnaire may have contained A and B of pair 1, $\neg A$ and $\neg B$ of pair 2, $A \wedge \neg B$ of pair 3, $B \wedge \neg A$ of pair 4, and $A \wedge \neg A$ and $B \wedge \neg B$ of pair 5. Each questionnaire therefore contained eight experimental items. In addition, twelve filler sentences were added to each questionnaire. The order of sentences was randomized individually in each questionnaire.

The subjects were divided into five groups of twenty. Each group was assigned one of the questionnaires. Subjects were instructed that the questionnaire was not a math

[6]Work conducted within the MTurk platform is always reviewed by the requesters and then either accepted or rejected. The HIT-acceptance rate indicates how often an individual's work has been accepted in the past.

test, but one of how true they felt the sentences are. Furthermore, they were given some examples of such intermediate truth, namely 95% for *'The Earth is round'* since it is not perfectly round, 50% for *'Mammals are furry'*. Subjects were then asked to indicate for each of the test items how true they felt it was.

Results and Discussion All values given by the subjects were in the 0 to 100 range, even though there were no restrictions on the input implemented in the html-code—i.e. a subject could have entered any alphanumeric string. So, that all responses were in range shows that subjects had paid attention to the instructions. Some subjects skipped single questions, but only 4 of 2000 data points were missing.

The average relative agreement to A, B, $\neg A$, and $\neg B$ was never exactly 50%. However, it had been decided to consider those items further where the relative agreement for each of A, B, $\neg A$, and $\neg B$ was between 40% and 60%. Only pair 1 of the conditions satisfied this requirement. The means for eight conditions in the four excluded conditions are shown in Figure 3.[7] The data in Figure 4 shows the average relative agreement (mean) and the standard error (SE) for each of the eight conditions of pair 1:

pair	A	$\neg A$	B	$\neg B$	$A \wedge \neg A$	$B \wedge \neg B$	$A \wedge \neg B$	$B \wedge \neg A$
2	54.6	47.1	**34.8**	59.9	52.5	42.3	53.1	42.4
3	47.9	**22.1**	59.8	**70.3**	51.6	54.9	50.6	46.4
4	**30.3**	**18.7**	54.6	69.8	27.5	72.0	43.1	56.2
5	**63.5**	**34.8**	55.2	49.9	45.6	49.3	60.8	43.8

Figure 3. Mean agreement in % for excluded pairs (violations marked by boldface).

	borderline				contradiction		conjunction	
	A	$\neg A$	B	$\neg B$	$A \wedge \neg A$	$B \wedge \neg B$	$A \wedge \neg B$	$B \wedge \neg A$
mean	45	42	45.75	47.25	48.15	46.5	43.2	25.65
SE	(6.7)	(6.0)	(7.0)	(7.0)	(4.3)	(6.1)	(4.3)	(5.4)

Figure 4. Mean agreement in % and standard error (SE) for pair 1.

The agreement to $A \wedge \neg A$ and $B \wedge \neg B$ is lower than we expected given that the basic propositions were borderline. To compare borderline propositions, A, B, $\neg A$, and $\neg B$, with the classical contradictions, $A \wedge \neg A$ and $B \wedge \neg B$, a Welch two sample t-test was computed. The difference was shown to not be significant at the $p < 0.05$ level ($t(103.4) = -0.50, p = 0.62$). This is, however, in line with the findings Ripley [17], who attributes the low agreement to borderline contradictions to cultural influence. Another reason may have been the greater grammatical and logical complexity of the contradiction items compared to the basic borderline propositions.

The more telling comparison is that between borderline contradictions and non-contradictory conjunctions. We find 47.3% agreement to the contradictions, $A \wedge \neg A$

[7] As Libor Běhounek (p.c.) points out, the comparison of sentence A of pair 2 and sentence B of pair 4 indicates that truth functionality more generally doesn't hold for graded truth judgments: Though the basic formulas are judged the same in the mean, the derived formulas ¬A and ¬B respectively are judged differently (with the difference likely to be significant, though I haven't tested this).

and $B \wedge \neg B$, compared to only 34.4% agreement to the non-contradictory conjunctions, $A \wedge \neg B$ and $B \wedge \neg A$. A Welch two sample t-test was computed to compare the two conditions. The test showed that subjects agreed to the borderline contradictions significantly more frequently than to the non-contradictory conjunctions at the $p < 0.05$ level ($t(77.0) = 2.60, p = 0.011$). The result shows an account of agreement levels based on fuzzy logic would not make the right prediction for conjunctions of borderline cases.

The experimental result therefore disconfirms the predictions of an account of speakers' graded truth judgments based solely on fuzzy logic.

3 Conclusions

The main point of this paper argues that fuzzy logic cannot provide a complete model for linguistic semantics of vagueness: Though speakers have no problems assigning intermediate truth values to sentences, the truth values assigned to complex sentences aren't systematic in a way that a fuzzy logic would predict. Therefore, adopting a fuzzy logic instead of a classical one doesn't provide a better approach to vagueness. In both cases, the account of truth-value judgments for complex sentences containing vague predicates cannot be that of the logical system, but needs to be relegated to pragmatics.

For the most part, I discussed the semantics of borderline contradictions such as *'She is tall and she is not tall.'* . Specifically I showed experimentally that the truth-functionality assumed by fuzzy logic fails. This finding holds regardless of whether we take the classical position that borderline contradictions are really contradictory (i.e. semantic value 0) or the dialethetist position where the semantic value of borderline contradictions is assumed to be the one speakers' intuitions actually assign to it, i.e. some truth value greater than 0. The specific problems the two positions give rise to are different though: For the classical position, the problem is simply that, if two logically unrelated A and B have the same semantic value, $A \wedge \neg A$ is less true than $B \wedge \neg A$ is judged to be. For the dialethetist position, the problem is the data I show above: When two logically unrelated A and B that have essentially the same semantic value, $A \wedge \neg A$ is judged more true than $B \wedge \neg A$ is. So, either way we cannot derive the semantic value of $A \wedge \neg A$ by applying a T-norm * to the semantic values of A and $\neg A$ that we would also apply to in evaluating $B \wedge \neg A$. This cannot be captured by a truth-functional logic, including a fuzzy one.

The results have implications for the account of vagueness in language beyond being problematic for classical and fuzzy logic. Fuzzy logic specifically cannot account for the data discussed because it builds vagueness directly into the logical composition. The data discussed speak instead for a separation of vagueness and logical composition, i.e. systems with the following two components: 1) a parameter of evaluation relating to vagueness and 2) a not fully truth-functional logic. Such parametric systems include epistemic approaches as well as sub- and super-valuation, and even ones such as fuzzy pluri-valuationism of Smith [19] where their non-vague logic is a fuzzy logic. The parametric approaches differ with respect to the nature of the paramater of evaluation and also with respect to the nature of the non-vague logic. The semantic value of a sentence is then determined from the results of evaluating the sentence relative to one or more parameter settings—in the case of vague expressions, typically more than one parameter setting would need to be considered. Specifically, if we have a probabilistic measure $\int$ on the parameter space P, truth judgments on sentences could be related to $\int_{p \in P} f(p) \partial p$ as is done, for instance, by Lassiter [12].

While our main result doesn't tell us much about the nature of the parameters, the experimental results discussed in Section 2—specifically the new finding showing that borderline contradictions are of higher acceptability than conjunctions of two unrelated propositions—tell us something about the non-vague logic used. Namely, none of these results are straightforwardly predicted if the logic used is completely classical. I think it is still possible to maintain a classical logic and allow the evaluation parameters to be shifted quite freely in the course of the evaluation of a sentence. But, a non-classical logic together with a parameter for vagueness provides a more straightforward account of the data.

BIBLIOGRAPHY

[1] Sam Alxatib and Jeffrey Pelletier. On the psychology of truth-gaps. In Nouwen et al. [14], pages 13–36.
[2] Nicolao Bonini, Daniel Osherson, Riccardo Viale, and Tim Williamson. On the psychology of vague predicates. *Mind and Language*, 14:377–393, 1999.
[3] Antonín Dvořák and Vilém Novák. Formal theories and linguistic descriptions. *Fuzzy Sets and Systems*, 143(1):169–188, 2004.
[4] Kit Fine. Vagueness, truth, and logic. *Synthese*, 30:265–300, 1975.
[5] Paul Grice. *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass., 1989.
[6] Petr Hájek. *Metamathematics of Fuzzy Logic*. Springer, 1998.
[7] Petr Hájek. Fuzzy logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2009 edition, 2009.
[8] Hans Kamp. Two theories of adjectives. In Edward Keenan, editor, *Formal Semantics of Natural Languages*, pages 123–155. Cambridge University Press, 1975.
[9] Hans Kamp and Barbara H. Partee. Prototype theory and compositionality. *Cognition*, 57:129–191, 1995.
[10] George Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2:458–508, 1973.
[11] Peter Lasersohn. Pragmatic halos. *Language*, 75:522–551, 1999.
[12] Daniel Lassiter. Vagueness as probabilistic linguistic knowledge. In Nouwen et al. [14].
[13] Richard Montague. The proper treatment of quantification in ordinary English. In R.H. Thomason, editor, *Formal Philosophy: Selected Papers of Richard Montague*, pages 247–270. Yale University Press, New Haven, Conn., 1974[1970].
[14] Rick Nouwen, Robert van Rooij, Uli Sauerland, and Hans-Christian Schmitz, editors. *Vagueness in Communication*. Springer, Heidelberg, Germany, 2011.
[15] Vilem Novák. Antonyms and linguistic quantifiers in fuzzy logic. *Fuzzy Sets and Systems*, 124(3): 335–351, 2001.
[16] Nicholas Rescher. *Many-Valued Logic*. McGraw-Hill, New York, N.Y., 1969.
[17] David J. Ripley. Contradictions at the borders. In Nouwen et al. [14], pages 169–188.
[18] Stephen Schiffer. Vagueness and partial belief. *Philosophical Issues*, pages 220–257, 2000.
[19] Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, USA, 2009.

Uli Sauerland
Zentrum für Allgemeine Sprachwissenschaft (Center for General Linguistics)
Schützenstr. 18
10117 Berlin, Germany
Email: uli@alum.mit.edu