Scientific Report for Xprag Application #x – Sam Alxatib

The meeting and the ensuing discussions were aimed at evaluating a new proposal regarding the assertability of vague propositions. Assertability, as will be seen shortly, is formalized as a turnstile. The justification for introducing the turnstile has to do with pragmatic principles, particularly the familiar maxims of quantity and quality (Grice 1975). It will seen first that adding the turnstile will make it possible to predict some recently discovered data regarding the psychology/use of vague predicates, and also suggest a solution to an often-cited deficiency of multi-valued systems. Discussions showed, however, that the idea is flawed. At this stage, the problems remain unresolved.

Multi-valued Logic

The language and the semantics of the logic under consideration are identical to those of basic multi-valued logics, the language being that of standard first-order logic, and the semantics – for presentational purposes – being those of basic fuzzy logic. We skip the definition of the language, but the semantics are defined below as a reminder. We define the valuation function v as follows:

(1) Let v be a function from wffs to the interval [0,1], then given a model $\mathcal{M}$

- (i) For any predicate letter P and term t, $v_{\mathcal{M}}(P(t)) = v_{\mathcal{M}}(t) \in v_{\mathcal{M}}(P)$
- (ii) $v_{\mathcal{M}}(\neg\phi) = 1 - v_{\mathcal{M}}(\phi)$
- (iii) $v_{\mathcal{M}}(\phi \vee \psi) = \max(v_{\mathcal{M}}(\phi), v_{\mathcal{M}}(\psi))$
- (iv) $v_{\mathcal{M}}(\phi \wedge \psi) = \min(v_{\mathcal{M}}(\phi), v_{\mathcal{M}}(\psi))$
- (v) $v_{\mathcal{M}}(\phi \rightarrow \psi) = 1$ if $v_{\mathcal{M}}(\psi) \geq v_{\mathcal{M}}(\phi)$

  $= 1 - (v_{\mathcal{M}}(\phi) - v_{\mathcal{M}}(\psi))$ otherwise.

The Turnstile

The turnstile $\models\!\!\!\equiv$, together with its auxiliary concepts, are defined in (T), (C), and (F) below.

(T) For a designated value d, formula Φ, and model $\mathcal{M}$, we define the assertability condition ($\overset{d}{\underset{\mathcal{M}}{\models\!\!\!\equiv}}$) as follows

$$\overset{d}{\underset{\mathcal{M}}{\models\!\!\!\equiv}} \Phi \quad \text{iff} \quad \frac{v_{\mathcal{M}}(\Phi) - {}^{\vee}\Phi}{{}^{\wedge}\Phi - {}^{\vee}\Phi} \geq d$$

(C) ${}^{\wedge}\Phi$ (to be read as the *ceiling of* Φ) $= i$ iff no model $\mathcal{M}$ is such that $v_{\mathcal{M}}(\Phi) > i$.

(F) ${}^{\vee}\Phi$ (to be read as the *floor of* Φ) $= i$ iff no model $\mathcal{M}$ is such that $v_{\mathcal{M}}(\Phi) < i$.

The assertability condition, and the notions of 'floor' and 'ceiling', are defined in terms of the valuation function $v_{\mathcal{M}}$. The ceiling (or floor) of a formula Φ is the highest (or lowest) degree of truth it can get: if it is impossible to find a model $\mathcal{M}$ in which $v_{\mathcal{M}}(\Phi)$ is greater than i, then i is said to be the ceiling of Φ. Likewise, if no model $\mathcal{M}$ is such that $v_{\mathcal{M}}(\Phi)$ is less than i, then i is the floor of Φ. Using ceilings and floors, we can define the *value-range* of a formula Φ as the interval $[{}^{\vee}\Phi, {}^{\wedge}\Phi]$. The value-range of an atomic proposition like $P(a)$, for example, is [0,1], because there are models in which a is a full member of P, and there are models in which a is a full non-member of P. But the value-range of a formula of the form $\phi \vee \neg\phi$ will be [0.5,1], since no model can give $\phi \vee \neg\phi$ a value less than 0.5. A formula Φ is said to be assertable relative to designated value d whenever the truth-value of Φ stretches d-far *within its value-range.* In the case of $P(a)$, this reduces to whether or not $v_{\mathcal{M}}(P(a)) \geq d$, for in this case the value-range is [0,1], and in order for the formula to be assertable, it has to be d-true within [0,1], i.e. d-true. But if Φ is $\phi \vee \neg\phi$, then the value-range shrinks to [0.5,1]. Say we set d at 0.5, and let $v_{\mathcal{M}}(\phi \vee \neg\phi) = 0.6$, then even though the truth-value of $\phi \vee \neg\phi$ is greater than d, if we focus only on the value-range, 0.6 is only 20%-true, which is lower than the 50% truth-requirement that we get from setting d at 0.5. With d at 0.5 (half-truth), $v_{\mathcal{M}}(\phi \vee \neg\phi)$ needs to be 0.75 in order for $\phi \vee \neg\phi$ to be assertable, for that marks the half-point in [0.5,1].

As mentioned in the introduction, the intuitive motivation for the turnstile is pragmatic. Recent findings by Ripley (2009), Alxatib and Pelletier (2010), and Sauerland (2010), show that contradictions are acceptable in borderline cases. The findings of Alxatib and Pelletier also suggest that tautological disjunctions like $\phi \vee \neg\phi$ are not. (Non-probabilistic) multi-valued logics make it possible for contradictions to take non-false truth-values, and make it possible for tautologous disjunctions to take non-true truth-values. The former reach their highest degree of truth, half-truth, whenever the conjuncts ϕ and $\neg\phi$ are themselves half-true. The latter reach their lowest degree of truth in the same case, namely when the two disjuncts are half-true. Since both the conjunction and the disjunction are half-true, the logic alone does not distinguish between the (un)acceptability of $\phi \wedge \neg\phi$ on the one hand, and $\phi \vee \neg\phi$ on the other, because each of ϕ, $\neg\phi$, $\phi \wedge \neg\phi$, and $\phi \vee \neg\phi$ will be half-true in the borderline range. But when we add to these considerations the intuition that contradictions (and tautologies) are uttered for informative/non-trivial reasons, we might follow up by adding that the truth of an expression is never intended to be evaluated relative to a value-range that the expression can never fall in, for if $\phi \vee \neg\phi$ were viewed in the full range of truth-values [0,1], there could never be a situation that falsifies it, and it would be uninformative of the speaker to utter an unfalsifiable proposition. Similarly, uttering $\phi \wedge \neg\phi$ cannot be considered cooperative if the proposition is known by the discourse participants to never reach truth. On a classical view, excluding the range of truth-values that are out of reach for $\phi \wedge \neg\phi$ and $\phi \vee \neg\phi$ would mean restricting one's attention to the range [0,0] for the former, and [1,1] for the latter. But on a multi-valued approach, the intervals will have non-zero length: [0,0.5] for contradictions and [0.5,1] for tautologies.

What results is a system that makes interesting predictions for borderline cases: we predict not only that contradictions are assertable, but also that their individual conjuncts are not. The reason is that while contradictions must be evaluated in the range [0,0.5], the conjuncts that comprise them may each be evaluated in the full range of truth-values. When $v_{\mathcal{M}}(P(a)) = 0.5$, $P(a)$ and $\neg P(a)$ will each be half-true relative to the range [0,1]. If the designated value is greater than 0.5, neither proposition will be assertable. But when they are conjoined together, the resulting formula will have a ceiling of 0.5, and will therefore be assertable if its value in $\mathcal{M}$ is 0.5. For the same reason, we also predict that $\neg(P(a) \vee \neg P(a))$ is assertable in the same situation. These results find empirical support in Alxatib and Pelletier (2010) and Ripley (2009).

Another welcome prediction of this approach is that we no longer assign $\phi \wedge \phi$ and $\phi \wedge \neg\phi$ the same status in borderline cases. Fuzzy logic has been criticized for assigning the same truth-value to these conjunctions in the borderline range – when $v_{\mathcal{M}}(\phi) = 0.5$, both conjuncts will have the value 0.5. The proposed approach makes different predictions: $\phi \wedge \phi$ will be assertable whenever ϕ itself is assertable, because adding ϕ in a conjunction will have no effect on the value-range of the formula, but conjoining ϕ with its negation will shrink the value-range to [0,0.5], as was already demonstrated, making the contradiction more assertable than any of ϕ, $\neg\phi$, $\phi \wedge \phi$, and $\neg\phi \wedge \neg\phi$.

Moreover, we predict a difference between the formulae $P(a) \wedge \neg P(a)$ on the one hand, and $P(a) \wedge \neg Q(a)$ on the other. In fuzzy logic, these two formulae are assigned the same truth-value whenever $P(a)$ and $Q(a)$ are both half-true. But if $P(a)$ and $Q(a)$ are logically-independent, the value-range of the second formula will be [0,1], rather than [0,0.5]. So, given a designated value $d > 0.5$, we predict $P(a) \wedge \neg Q(a)$ to be unassertable, but we predict the contradiction $P(a) \wedge \neg P(a)$ to be assertable. This result was tested and verified empirically by Sauerland (2010).

Problems

The system in its current formulation is in need of major revisions, owing to two problem with the assertability of disjunctions, and a problem with the assertability of conditionals.

The first problem was noted in our first meeting by Uli Sauerland. Take a model $\mathcal{M}$ such that $v_{\mathcal{M}}(P(a)) = 0.5$ and $v_{\mathcal{M}}(Q(a)) = 0$ (as an example, say that a is borderline case of tallness, and a is not rich at all). The conditions on the turnstile give the prediction that $P(a) \wedge \neg P(a)$ is assertable, but, shockingly, that $(P(a) \wedge \neg P(a)) \vee Q(a)$ is not. In fact, the problem generalizes to every case $(\phi \wedge \neg\phi) \vee \psi$, where ψ is logically-independent from ϕ. Because ψ is logically-independent, it will give the entire formula a floor of 0 and a ceiling of 1 again, and it will therefore follow that whenever this formula is assertable, its assertability will never be due to the contradictory first disjunct (except when the designated value d is less than or equal to 0.5).

The second problem is that tautological disjunctions, already discussed above, will actually turn out to be too stringent. Consider the case where $v_{\mathcal{M}}(P(a)) = 0.8$, and where $d = 0.7$. Clearly, $P(a)$ will be assertable in this case. But even though $P(a)$ is assertable, $P(a) \vee \neg P(a)$ will not be. This is because the floor of the disjunction, 0.5, will be subtracted from its value in $\mathcal{M}$, giving us 0.3, and when this is divided by the difference between the formula's ceiling and floor, we get the value 0.3/0.5=0.6, which is lower than d.

The third problem is that a conditional like $\phi \to \phi$ will satisfy the condition of assertability; the floor and ceiling of $\phi \to \phi$ are both 1, so the denominator of the assertability fraction will be zero. But the numerator will be zero as well because the floor of the formula is equal to its value in any given model. So $\phi \to \phi$ is assertable iff $0/0 \geq d$, which cannot be evaluated. This problem holds if we follow the definition given on pg. 1 for the semantics of fuzzy conditionals. We could check other definitions, but the more general goal is to look for a semantics that allows neither tautologies nor contradictions (i.e. no formulae whose value is 1 or 0 in every model), for it is exactly these formulae that cause the fraction to have an uninterpretable value. Alternatively, we could impose a further restriction on assertability, namely, that the floor and ceiling of the given formula may not be equal, and in the cases where they are equal, the formula is not assertable because of its hopeless uninformativity.

As for the first two problems, no good solution is in sight. In our later meetings we attempted to change the definition of assertability to a recursive definition. The basic case is the case for (unnegated) atomic formulae, and here we use a variant of the definition we already have. But as soon as we tackle the the recursive step involving negation, we run into problems: saying that $\neg\phi$ is assertable iff ϕ is itself not assertable will not do, because we have evidence that $\phi \wedge \neg\phi$ is assertable in borderline cases, and by this recursive step we will require ϕ to be assertable and to be unassertable at the same time, i.e. that the value of ϕ be greater than and also less than the given designated threshold, which is clearly impossible. So the assertability of a negated proposition must be defined differently. An alternative that comes to mind is

$$\mathrel{\underset{d}{\overset{\mathcal{M}}{\models}}} \neg\phi \text{ iff it is not the case that } \mathrel{\underset{1-d}{\overset{\mathcal{M}}{\models}}} \phi$$

This definition brings negation closer to what is often refered to as "choice" or "strong" negation: $\neg P(a)$ is assertable just in case the value of $P(a)$ is lower than $1-d$. If $d = 0.7$, then in order for $\neg P(a)$ to be assertable, the membership of a in P has to be lower than 0.3; if its degree of membership is any higher than that, it no longer qualifies as a "non-member" of P. In other words, if a's membership in P is greater than 0.3, and our designated value is 0.7, a will be too much of a "P" for us to assert $\neg P(a)$.

Now if a is borderline-P, then $P(a)$ will not be assertable, and $\neg P(a)$ will not be assertable either. The conjunction, in the recursive step, will have to be evaluated relative to its value range. I refrain from formalizing this idea here, but once we find the value-range of the conjunction, [0,0.5], we say that it is assertable iff its individual conjuncts are assertable relative to [0,0.5]. For borderline cases, this condition is satisfied.

Disjunctions, however, will complicate the picture yet again. Both the problematic cases will be correctly accounted for if we say that

$$\mathrel{\underset{d}{\overset{\mathcal{M}}{\models}}} \phi \vee \psi \text{ iff } \mathrel{\underset{d}{\overset{\mathcal{M}}{\models}}} \phi \text{ or } \mathrel{\underset{d}{\overset{\mathcal{M}}{\models}}} \psi$$

In the case of $(\phi \wedge \neg\phi) \vee \psi$, the assertability of the formula will depend on whether either disjunct is assertable, i.e. on whether $\phi \wedge \neg\phi$ is assertable or ψ is assertable. Here it will not matter if ψ is false; if ϕ is borderline, $\phi \wedge \neg\phi$ will be assertable because of the condition on conjunctions. We also predict that $\phi \vee \neg\phi$ is assertable only when ϕ or its negation satisfies their own assertability conditions, that is, that whenever ϕ (or $\neg\phi$) is assertable, $\phi \vee \neg\phi$ will be assertable as well. This solves both the disjunction-related problems described above.

But now consider $\neg(\phi \vee \neg\phi)$. Our treatment of negation as strong negation was motivated by the attested acceptability of $\phi \wedge \neg\phi$ in the borderline range. We also have evidence that $\neg(\phi \vee \neg\phi)$ is acceptable in the very same range (a borderline-tall individual is "neither tall nor not tall"). But under the current formulation, the outermost negation in this formula will be interpreted strongly, making $\neg(\phi \vee \neg\phi)$ assertable with respect to value d just in case $\phi \vee \neg\phi$ is *not* assertable with respect to $1-d$. This means that both ϕ and$\neg\phi$ must not be assertable relative to $1-d$. If ϕ is $P(a)$, the condition on the assertability of $\neg(P(a) \vee \neg P(a))$ will be that $a \in P < 1-d$ and $a \in P \geq d$, which is possible only when $d < 0.5$. The problem can be understood intuitively as saying that the outermost negation should be treated as weak negation: when it is said that a is neither tall nor not tall, what is meant that we cannot assert that a is tall, and we cannot assert that a not tall. In other words, "a is neither tall nor not tall" is assertable just in case "a is either tall or not tall" is not assertable, and this will not result from our strong condition on the assertability of negation.

This is the picture that we have so far, and we are continuing our search for the correct formulation of the system.