Vagueness in Intelligent Systems

Jonathan Lawry

Department of Engineering Mathematics, University of Bristol, UK

Ecsqaru 2013 Tutorial
Utrecht, July 7, 2013

What this Tutorial is not About

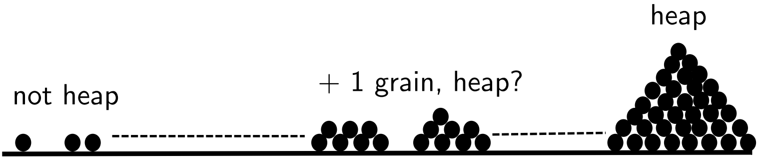

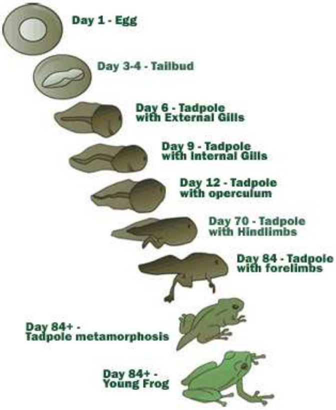

- The sorites paradox from the Greek for ‘heap’.
- Attributed to Eubulides of Miletus (4th century BCE).
- With many variations sorites susceptibility is often taken as a defining feature of vague predicates.

Is Vagueness Useful for Intelligent Systems?

it seems rather far fetched to conclude that we have simply tolerated a worldwide, several thousand year efficiency loss. [Barton Lipman]

- Vagueness pervades natural language yet it is frowned upon in the western scientific tradition.
- In science clarity and (semantic) precision are seen as being a fundamental prerequisite to progress.
- A hypothesis must be clearly formulated before it can be properly empirically tested.
- So why is vagueness so common in almost all aspects of language?
- So is it actually useful and if so how and why?

Some Possible Uses of Vagueness

- Kees Van Deemter (2009) has proposed a number of possible tasks in which vagueness may be useful, including 1 and 2:

1. *Future Contingencies:* Vagueness can mitigate the risk associated with making forecasts or promises.
2. *Search:* Order information about typicality is embedded in vague predicates. This has the potential to reduce search times.
3. *Consensus:* Vagueness enables agents to reach consensus while maintaining internal consistency.
4. *Language Learning:* The meaning of concepts emerge through a distributed process of interactions between agents leading to semantic uncertainty.
5. *Flexibility:* Small numbers of decision rules which can fire to an intermediate degree. Constraints and specifications which can naturally be relaxed or tightened.

What is Vagueness?

> *Today, vague predicates are standardly characterized by three main 'symptoms', namely as predicates that are sorities susceptible, that have borderline cases, and that have blurry boundaries. (Paul Égré)*

- Vagueness is a multifaceted phenomena and the different possible uses we have outlined exploit different aspects of vagueness.
- *Indeterminism:* Borderline cases may facilitate consensus as well as mitigating the risk of making promises or forecasts.
- *Semantic Uncertainty:* Explicit representation of uncertain about predicate definitions can aid learning and decision making.
- *Typicality:* Prototype based conceptual models may allow agents to effectively order candidate solutions during search.

Truth-gaps and Indeterminism

- Propositions can be *absolutely true* or *absolutely false* but there may be also be a *truth-gap*.
- Some propositions may be neither absolutely true nor absolutely false i.e. indeterminate or borderline
- For example, consider the proposition $p =$‘Ethel is short’.
- In this context *short* could be defined by two height thresholds $\underline{h} \leq \overline{h}$.
- Let Ethel’s height be h. Then p is absolutely true if $h \leq \underline{h}$, absolutely false of $h > \overline{h}$ and borderline otherwise.

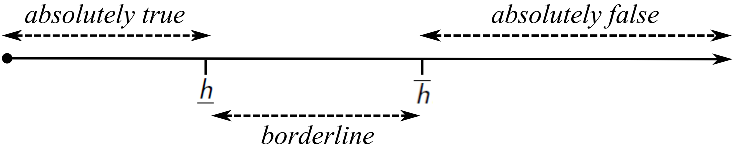

A Propositional Model of Indeterminism

- $\mathcal{L}$ is a language with propositional variables $\mathcal{P} = \{p_1, \ldots, p_n\}$, connectives $\neg, \wedge, \vee$ and sentences $S\mathcal{L}$.

Definition

A valuation pair is a pair of binary functions $\vec{v} = (\underline{v}, \overline{v})$ where $\underline{v} : S\mathcal{L} \to \{0,1\}$, $\overline{v} : S\mathcal{L} \to \{0,1\}$ and $\underline{v} \leq \overline{v}$.
For sentence $\theta \in S\mathcal{L}$, $\underline{v}(\theta) = 1$ means that θ is absolutely true, and $\overline{v}(\theta) = 1$ means that θ is not absolutely false.

- We can also think of a valuation pair as a three-valued mapping with $\vec{v}(\theta)$ having possible values $\mathbf{t} = (1,1)$, $\mathbf{b} = (0,1)$ and $\mathbf{f} = (0,0)$.
- There are a number of ways of defining valuation pairs...

Two Types of Valuation Pairs

Definition

A Kleene valuation pair is valuation pair $(\underline{v}, \overline{v})$ *satisfying:*
$\forall \theta, \varphi \in S\mathcal{L}$

1. $\underline{v}(\neg\theta) = 1 - \overline{v}(\theta)$, $\overline{v}(\neg\theta) = 1 - \underline{v}(\theta)$
2. $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$ *and* $\overline{v}(\theta \wedge \varphi) = \min(\overline{v}(\theta), \overline{v}(\varphi))$
3. $\underline{v}(\theta \vee \varphi) = \max(\underline{v}(\theta), \underline{v}(\varphi))$ *and* $\overline{v}(\theta \vee \varphi) = \max(\overline{v}(\theta), \overline{v}(\varphi))$

- Let $\mathbb{V}_c$ denote the set of classical valuations on $\mathcal{L}$.

Definition

A supervaluation [Fine] pair is defined as follows: For $\Pi \subseteq \mathbb{V}_c$ *let*
$\forall \theta \in S\mathcal{L}$;

$$\underline{v}(\theta) = \min\{v(\theta) : v \in \Pi\} \text{ and } \overline{v}(\theta) = \max\{v(\theta) : v \in \Pi\}$$

- Let $\mathbb{V}_k$ and $\mathbb{V}_s$ denote the sets of Kleene and supervaluation pairs on $\mathcal{L}$, respectively.

Orthopairs and Kleene Valuations

- Kleene valuation pairs can be characterised by a pair of sets of propositional variables (P, N) where $P = \{p_i : \underline{v}(p_i) = 1\}$ and $N = \{p_i : \underline{v}(\neg p_i) = 1\}$
- Notice that since $\underline{v} \leq \overline{v}$ then $\underline{v}(p_i) = 1 \Rightarrow \overline{v}(p_i) = 1 \Rightarrow \underline{v}(\neg p_i) = 0$.
- Hence, $P \cap N = \emptyset$ i.e (P, N) is an *orthopair*.
- Example: Let $\mathcal{P} = \{p_1, p_2, p_3, p_4\}$ and $\vec{v} = (\{p_1\}, \{p_4\})$ then $\vec{v}(p_1) = \mathbf{t}$, $\vec{v}(p_4) = \mathbf{f}$, $\vec{v}(p_2) = \vec{v}(p_3) = \mathbf{b}$.
- If $p_i \in (P \cup N)^c$ then $\vec{v}(p_i) = \mathbf{b}$
- Hence, $|(P \cup N)^c|$ quantifies the vagueness of a particular valuation pair.
- (P_1, N_1) and (P_2, N_2) are *consistent* if $P_1 \cap N_2 = P_2 \cap N_1 = \emptyset$.

Semantic Precision: An Ordering on Precisification

- We now define *semantic precision* as a natural partial ordering on valuation pairs.

Definition

Given two valuation pairs $\vec{v}_1$ *and* $\vec{v}_2$, $\vec{v}_1 \preceq \vec{v}_2$ *if and only if* $\forall \theta \in S\mathcal{L}$, $\underline{v}_1(\theta) \leq \underline{v}_2(\theta)$ *and* $\overline{v}_1(\theta) \geq \overline{v}_2(\theta)$.

- $\vec{v}_1$ is less semantically precise than $\vec{v}_2$ if they disagree only for some set of sentences of $\mathcal{L}$, which being identified as either **t** or **f** by $\vec{v}_2$, are classified as being **b** by $\vec{v}_1$.

Theorem

- *For* $\vec{v}_1, \vec{v}_2 \in \mathbb{V}_k$, $\vec{v}_1 \preceq \vec{v}_2$ *iff* $P_1 \subseteq P_2$ *and* $N_1 \subseteq N_2$.
- *For* $\vec{v}_1, \vec{v}_2 \in \mathbb{V}_s$, $\vec{v}_1 \preceq \vec{v}_2$ *iff* $\Pi_1 \supseteq \Pi_2$.

Relating Kleene and Supervaluation Pairs

- Consider the language $\mathcal{L}$ with $\mathcal{P} = \{p_1, p_2\}$
- Let $v_0, v_1, v_2, v_3 \in \mathbb{V}_c$ such that $v_0(\neg p_1 \wedge \neg p_2) = 1$, $v_1(p_1 \wedge \neg p_2) = 1$, $v_2(\neg p_1 \wedge p_2) = 1$, $v_3(p_1 \wedge p_2) = 1$

$\sqcap$

$\{v_1, v_2\} : \vec{v}(p_1 \wedge p_2) = \mathbf{f}, \vec{v}(p_1 \wedge \neg p_2) = \mathbf{b},$
$\vec{v}(\neg p_1 \wedge p_2) = \mathbf{b}, \vec{v}(\neg p_1 \wedge \neg p_2) = \mathbf{f}$

$\{v_0, v_3\} : \vec{v}(p_1 \wedge p_2) = \mathbf{b}, \vec{v}(p_1 \wedge \neg p_2) = \mathbf{f},$
$\vec{v}(\neg p_1 \wedge p_2) = \mathbf{f}, \vec{v}(\neg p_1 \wedge \neg p_2) = \mathbf{b}$

$\{v_1, v_2, v_0\}$

$\{v_1, v_2, v_3\}$

$\{v_0, v_3, v_1\}$

$\{v_0, v_3, v_2\}$

$\{v_0, v_1, v_2, v_3\} : \vec{v}(p_1 \wedge p_2) = \mathbf{b},$
$\vec{v}(p_1 \wedge \neg p_2) = \mathbf{b}, \vec{v}(\neg p_1 \wedge p_2) = \mathbf{b},$
$\vec{v}(\neg p_1 \wedge \neg p_2) = \mathbf{b}$

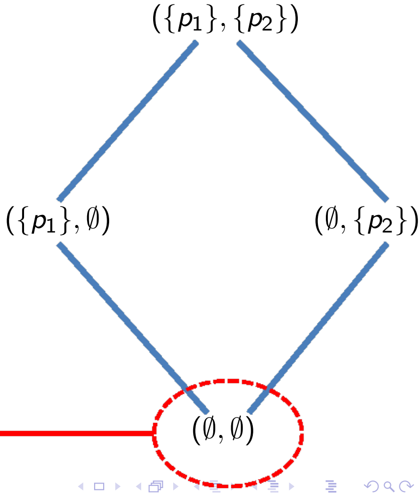

Penumbral Connections

- Penumbral connections are defined to be those 'logical relations [that] hold between indefinite sentences' [Fine].
- Given a sequence of heights $h_1 < h_2 < \ldots < h_n$ where only h_1 is absolutely not tall, and only h_n is absolutely tall, so that all other heights are borderline cases of tall.
- But if we learn that h_i is tall for any $i \in \{2, \ldots, n-1\}$ we immediately infer that h_{i+1} is tall, because $h_{i+1} > h_i$.
- Let $v_i \in \mathbb{V}_c$ be such that $v_i(\bigwedge_{j=1}^{i-1} \neg p_i \wedge \bigwedge_{j=i}^{n} p_i) = 1$ then take $\Pi = \{v_2, \ldots, v_n\}$.

Uncertainty and Vagueness

- Consider a combined model incorporating both indeterminism and epistemic uncertainty.
- In this context natural division of uncertainty types is:
- *Semantic Uncertainty*: This takes the form of uncertainty about what is the *correct* interpretation of $\mathcal{L}$. For example, an agent may be uncertain as to whether or not a proposition such as 'Ethel is short' is absolutely true or absolutely false even if they know Ethel's height precisely.
- *Possible Worlds Uncertainty*: This type of uncertainty arises from a lack of knowledge concerning the state of the world and in particular about the referents of sentences in $\mathcal{L}$. For example, an agent may not know Ethel's height precisely and hence be uncertain about the truth value of the proposition 'Ethel is short'.

Semantic Uncertainty and the Epistemic Stance

Usually it is uses of words, not words in themselves, that are properly called vague. [J. L. Austin]

- Recall the proposition 'Ethel is short' where short is defined in terms of lower and upper threshold values $\underline{h} \leq \overline{h}$.
- Here semantic uncertainty could take the form of a probability density function on $\{(\underline{h}, \overline{h}) \in \mathbb{R}^2 : 0 \leq \underline{h} \leq \overline{h}\}$
- Agents need to decide on assertions and update the conceptual models based their previous experience of communications with other agents.
- *The Epistemic Stance* [Lawry]: Individuals, when faced with such issues, find it useful as part of a pragmatic decision making and learning strategy to *assume* that there is a correct interpretation of $\mathcal{L}$.
- This a weaker form of the *epistemic theory of vagueness*. [Williamson].

Belief Pairs

- In our current model epistemic uncertainty takes the form of uncertainty as to which is the *correct* valuation pair.
- Let $\mathbb{V}$ be a finite set (a certain class) of valuation pairs on $\mathcal{L}$.
- Let $w : \mathbb{V} \to [0,1]$ be a probability measure on $\mathbb{V}$ representing an agent's subjective belief.
- This naturally generates lower and upper measures as follows:
- $\underline{\mu}(\theta) = w(\{\vec{v} \in \mathbb{V} : \underline{v}(\theta) = 1\}) = w(\{\vec{v} \in \mathbb{V} : \vec{v}(\theta) = \mathbf{t}\})$ i.e. the agent's subjective probability that θ is absolutely true.
- $\overline{\mu}(\theta) = w(\{\vec{v} \in \mathbb{V} : \overline{v}(\theta) = 1\}) = w(\{\vec{v} \in \mathbb{V} : \vec{v}(\theta) \neq \mathbf{f}\})$ i.e. the agent's subjective probability that θ is not absolutely false.
- If $\mathbb{V} \subseteq \mathbb{V}_s$ then $(\underline{\mu}, \overline{\mu})$ are Dempster-Shafer belief and plausibility measures on $S\mathcal{L}$.
- If $\mathbb{V} \subseteq \mathbb{V}_k$ then $(\underline{\mu}, \overline{\mu})$ are Kleene belief pairs with some relationship to fuzzy and interval fuzzy truth degrees.

The Cautious Politician

I pledge to vote against any increase in fees in the next parliament and to pressure the government to introduce a fairer alternative.

- Consider a politician who is considering making a pledge at an upcoming general election.
- *Vague Pledge*: A *significant* reduction in the budget deficit, with only a *minor* increase in unemployment.
- *Crisp Pledge*: At least a 40% reduction in the budget deficit, with no more than a 2% increase in unemployment.
- Vague assertion: Lower initial reward but lower risk of failing to meet the pledge.
- Crisp assertion: Higher initial reward but higher risk of failing to meet the pledge.
- Can we develop a utility model for choosing between these assertions?

A Simple Utility Model

- Suppose that the agent must choose between a vague assertion θ and a crisp assertion φ.
- We assume $\underline{\mu}(\varphi) = \overline{\mu}(\varphi) = \mu(\varphi)$ and $\underline{\mu}(\theta) \leq \mu(\varphi) \leq \overline{\mu}(\theta)$.
- Let $x_1, x_2 \in \mathbb{R}^+$ be the (initial) rewards for asserting θ and φ respectively. (Simplify $x_1 = x_2 = x$)
- $y \in \mathbb{R}^+$ the reward for making an *absolutely true* assertion (the same for θ and φ).
- $-z$ for $z \in \mathbb{R}^+$ is the cost of making an assertion which is *absolutely false*.
- Assume *borderline* assertions are neutral with 0 reward or cost.
- $E(U_\theta) = x - z + \underline{\mu}(\theta)y + \overline{\mu}(\theta)z$,
  $E(U_\varphi) = x - z + \mu(\varphi)(y + z)$.
- $E(U_\theta) \geq E(U_\varphi)$ iff $\alpha \leq \frac{\overline{\mu}(\theta) - \mu(\varphi)}{\mu(\varphi) - \underline{\mu}(\theta)}$ where $\alpha = \frac{y}{z}$.

Interpretation of Utility Model

Some people are always critical of vague statements. I tend rather to be critical of precise statements; they are the only ones which can correctly be labeled 'wrong'.[Raymond Smullyan]

- We have an upper bound on the ratio of reward over cost for which $E(U_\theta) \geq E(U_\varphi)$.
- *The higher the cost of making an absolutely false assertion relative to the reward of making an absolutely true one, the more likely it is that the agent is better off making a vague rather than a crisp assertion.*
- If $\mu(\varphi) \leq \frac{\underline{\mu}(\theta)+\overline{\mu}(\theta)}{2}$ then $E(U_\theta) \geq E(U_\varphi)$ for all $\alpha \in [0,1]$ i.e. for $y \leq z$.
- A special case of this is when $\mu(\varphi) \leq 0.5$ and $\underline{\mu}(\theta) = 0$ and $\overline{\mu}(\theta) = 1$ i.e. the agent is *certain* that θ is *borderline*.

Emergent Consensus

"Then we are agreed nine to one that we will say our previous vote was unanimous!"

- In many decision making and negotiation scenarios intelligent agents need to reach a common shared position or viewpoint about some set of propositions.
- One route to such a consensus is for each agent to adopt a more vague interpretation of underlying predicates so as to soften directly conflicting opinions.
- Truth-gaps enable agents to reach consensus by weakening their viewpoints while maintaining internal consistency.

Consensus by Weakening

If you can't be kind, at least be vague. [Judith Martin]

- Suppose two agents a_1 and a_2 need to reach agreement about the proposition $p =$'UK inflation is currently low'.
- f =the actual level of UK inflation
- Agent a_i defines *low* in terms of lower and upper threshold $\underline{f}_i \leq \overline{f}_i$.

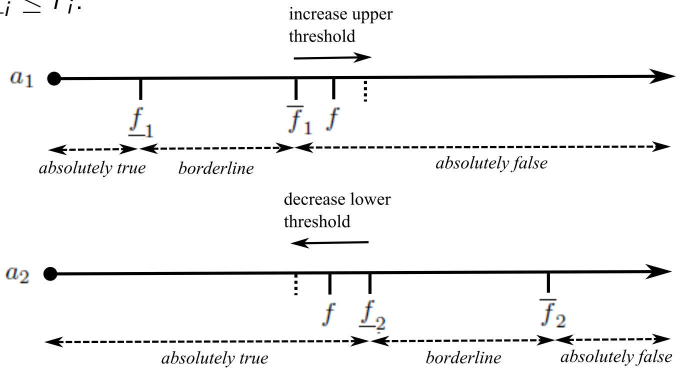

A Natural Consensus Operator

- We use the orthopair notation to define a consensus combination operator in the Kleene valuation pair framework.

Definition

Given valuation pairs $\vec{v}_1$ *and* $\vec{v}_2$ *with orthopairs* (P_1, N_1) *and* (P_2, N_2) *then* $\vec{v}_1 \odot \vec{v}_2$ *is defined by the orthopair* $((P_1 - N_2) \cup (P_2 - N_1), (N_1 - P_2) \cup (N_2 - P_1))$.

- For literals we have that:

$$\underline{v_1 \odot v_2}(l) = \max(\min(\underline{v_1}(l), \overline{v_2}(l)), \min(\underline{v_2}(l), \overline{v_1}(l)))$$
$$\overline{v_1 \odot v_2}(l) = \min(\max(\overline{v_1}(l), \underline{v_2}(l)), \max(\overline{v_2}(l), \underline{v_1}(l)))$$

$\odot$	**t**	**b**	**f**
t	**t**	**t**	**b**
b	**t**	**b**	**f**
f	**b**	**f**	**f**

A Multi-Agent Simulation Study on Consensus

- Language size n propositional variables, k agents.
- Some useful metrics are:
- Inconsistency: $I(\vec{v_1}, \vec{v_2}) = \frac{|P_1 \cap N_2| + |P_2 \cap N_1|}{n}$
- Vagueness: $V(\vec{v}) = \frac{|(P \cup N)^c|}{n}$
- We now set a consistency threshold $\gamma \in [0, 1]$ on belief combination, so that valuations $\vec{v_1}$ and $\vec{v_2}$ can only be combined using $\odot$ if $I(\vec{v_1}, \vec{v_2}) \leq \gamma$.
- At each discrete time-step two agents characterised by $\vec{v_1}$ and $\vec{v_2}$ are selected at random from the population.
- If $I(\vec{v_1}, \vec{v_2}) \leq \gamma$ then both $\vec{v_1}$ and $\vec{v_2}$ are replaced by $\vec{v_1} \odot \vec{v_2}$, and otherwise are left unchanged.
- Each experiment (i.e. specific combination of parameter values) was run 100 times and results shown are averages across these runs.

Distinct Valuations at Steady-State

- As the inconsistency threshold tends to one the number of distinct valuations (or viewpoints) also tends to one.

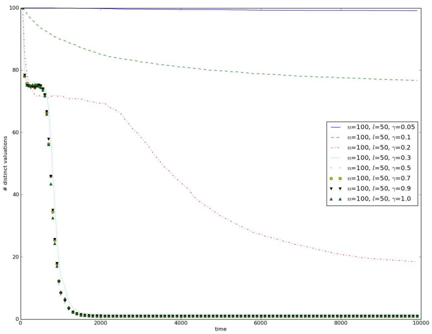

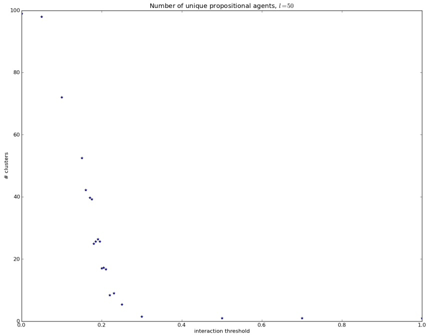

Vagueness at Steady-State

- As the inconsistency threshold tends to one then the vagueness level tends to zero.
- Intermediate inconsistency threshold levels result in a number of distinct vague viewpoints at steady state.

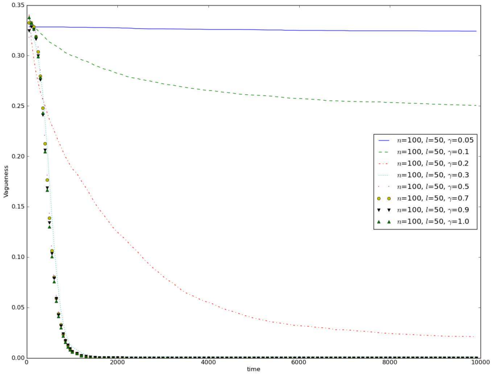

- Suppose an agent is only uncertain about the correct level of semantic precision for interpretation of $\mathcal{L}$.
- Formally; suppose w is non-zero only on a sequence $\vec{v}_1 \preceq \vec{v}_2 \preceq \ldots \preceq \vec{v}_k$
- If $\vec{v}_i \in \mathbb{V}_s$ then $\underline{\mu}$ and $\overline{\mu}$ are partially compositional *necessity* and *possibility* measures on $S\mathcal{L}$ respectively.
- If If $\vec{v}_i \in \mathbb{V}_k$ then $\underline{\mu}$ and $\overline{\mu}$ are fully compositional *lower and upper fuzzy truth-degrees* satisfying [Lawry, Gonzalez]: $\forall \theta, \varphi \in S\mathcal{L}$

$$\underline{\mu}(\neg\theta) = 1 - \overline{\mu}(\theta),\ \overline{\mu}(\neg\theta) = 1 - \underline{\mu}(\theta)$$
$$\underline{\mu}(\theta \wedge \varphi) = \min(\underline{\mu}(\theta), \underline{\mu}(\varphi)),\ \overline{\mu}(\theta \wedge \varphi) = \min(\overline{\mu}(\theta), \overline{\mu}(\varphi))$$
$$\underline{\mu}(\theta \vee \varphi) = \max(\underline{\mu}(\theta), \underline{\mu}(\varphi)),\ \overline{\mu}(\theta \vee \varphi) = \max(\overline{\mu}(\theta), \overline{\mu}(\varphi))$$

Vague Predicates, Typicality and Blurry Boundaries

- Consider a language $\mathcal{L} = \{L_1, \ldots, L_n\}$ with connectives $\wedge, \vee, \neg$ and single variable x; L_i are unary predicates (labels) for describing elements of an underlying metric space (Ω, d).
- Let $F\mathcal{L}$ denote the formula of $\mathcal{L}$ obtained by recursively applying the connectives to $L_i(x) : i = 1, \ldots, n$.
- $L_i = (P_i, \epsilon_i, \delta_i)$; $P_i \subseteq \Omega$ is a set of prototypes for L_i, ϵ_i is a positive real-valued random variable and $\epsilon_i \sim \delta_i$ [Lawry, Tang].
- For $x \in \Omega$, $L_i(x)$ is true provided x is *sufficiently close* to P_i i.e. $d(x, P_i) \leq \epsilon_i$

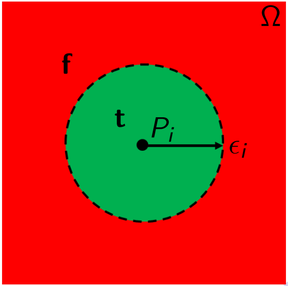

Random Set Neighbourhoods

- Predicate L_i is characterised by the following consonant random set:

$$\mathcal{N}_{L_i}^{\epsilon_i} = \{x \in \Omega : d(x, P_i) \leq \epsilon_i\}$$

- So that: $L_i(x)$ is true if and only if $x \in \mathcal{N}_{L_i}^{\epsilon_i}$
- Then neighbourhoods for formula in $F\mathcal{L}$ can be defined recursively according to the rules: For $\theta, \varphi \in F\mathcal{L}$;

$$\mathcal{N}_{\neg\theta}^{\vec{\epsilon}} = (\mathcal{N}_{\theta}^{\vec{\epsilon}})^c,\ \mathcal{N}_{\theta\wedge\varphi}^{\vec{\epsilon}} = \mathcal{N}_{\theta}^{\vec{\epsilon}} \cap \mathcal{N}_{\varphi}^{\vec{\epsilon}},\ \mathcal{N}_{\theta\vee\varphi}^{\vec{\epsilon}} = \mathcal{N}_{\theta}^{\vec{\epsilon}} \cup \mathcal{N}_{\varphi}^{\vec{\epsilon}}$$

- Note that $\vec{\epsilon} = (\epsilon_1, \ldots, \epsilon_i)$ appears above because the neighbourhood for a general formula θ may be defined in terms of any or all of these n thresholds.
- E.g. for $\theta = \bigwedge_{i=1}^{n} L_i$,

$$\mathcal{N}_{\theta}^{\vec{\epsilon}} = \{x \in \Omega : d(x, P_i) \leq \epsilon_i, i = 1, \ldots, n\}$$

Probability, Similarity and Membership Functions

- Given this framework we should also take account of the semantic uncertainty concerning the threshold parameters $\epsilon_i : i = 1, \ldots, n$.
- For a given $x \in \Omega$ we can determine the probability that $L_i(x)$ is true, according to:

$$\mu_{L_i}(x) = \delta_i(\{\epsilon_i : \epsilon_i \geq d(x, P_i)\}) = \Delta_i(d(x, P_i))$$
$$\text{where } \Delta_i(t) = \delta_i(\{\epsilon_i : \epsilon_i \geq t\})$$

- $\mu_{L_i} : \Omega \rightarrow [0, 1]$ is a membership function (single point coverage) of the random set defining L_i in Ω.
- We can also think in terms of similarity to prototypes.
- $S_i(x, y) = \Delta_i(d(x, y))$ is a *similarity measure* and $\mu_{L_i}(x) = S_i(x, P_i)$

Membership for Compound Formula

- For a compound formula $\theta \in F\mathcal{L}$ calculating the probability that $x \in \mathcal{N}_{\theta}^{\vec{\epsilon}}$ in general requires a joint distribution on $(\epsilon_1, \ldots, \epsilon_n)$.
- Here we consider just two special cases:
- *Total Independence:* All thresholds are independent of all other thresholds. Appropriate when the different predicates relate to different independent characteristics of the elements in Ω e.g. *tall* and *rich*. In this case:

$$\mu_{L_i \wedge L_j}(x) = \mu_{L_i}(x) \times \mu_{L_j}(x)$$

- *Total Dependence:* All thresholds are rescalings of a shared underlying threshold. $\epsilon_i = f_i(\epsilon) : i = 1, \ldots, n$ where f_i is an increasing function. Appropriate when predicates all refer to strongly related features or characteristics e.g. colour, taste etc. In this case:

$$\mu_{L_i \wedge L_j}(x) = \min(\mu_{L_i}(x), \mu_{L_j}(x))$$

Example: Colour Categories

- Let $\Omega = [0,1]^3$ (normalised rgb space) with metric $d(x,y) = \|x - y\|$.
- Predicates are *green* $= ((0,1,0), \epsilon, \delta)$ and *blue* $= ((0,0,1), \epsilon, \delta)$
- We consider two cases: 1) δ is the uniform distribution on $[0, 0.9]$ and 2) δ is a normalised Gaussian with mode 0 and $\sigma = 0.55$.

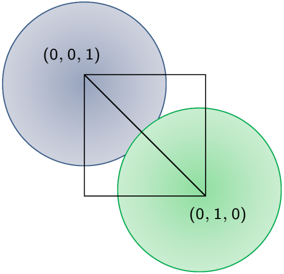

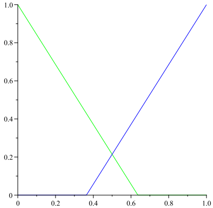

uniform

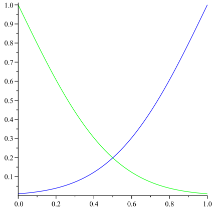

gaussian

A Predicate Model Incorporating Truth-Gaps

- So far our underlying truth-model for vague predicates has been Boolean.
- The approach can naturally be extended to incorporate explicit borderlines.
- Take $L_i = (P_i, \underline{\epsilon_i}, \overline{\epsilon_i}, \delta_i)$ where $\underline{\epsilon_i} \leq \overline{\epsilon_i}$ and $(\underline{\epsilon_i}, \overline{\epsilon_i}) \sim \delta_i$

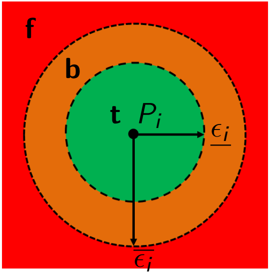

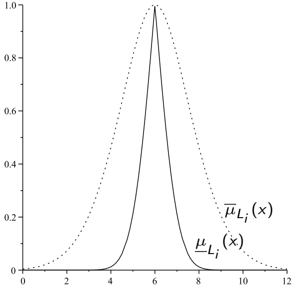

Language Games and Category Evolution

- Luc Steels (1995) has argued that to be more realistic language models must take into account the evolutionary nature of language learning.
- Language games model language learning through pairwise interactions between agents.

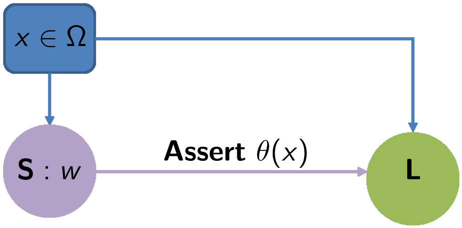

Update?
Yes if $\mu_\theta(x) < w$
No otherwise

Updating Conceptual Models

Assert $L_i(x)$: $\lambda \in [0,1], \alpha \geq 1$

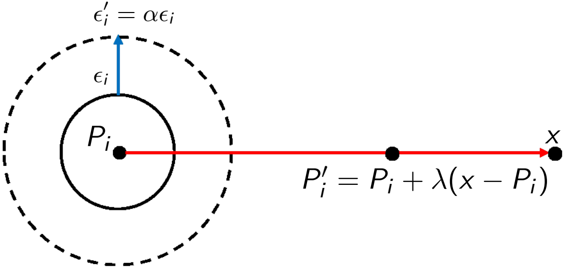

Assert $\neg L_i(x)$: $\lambda \in [0,1], \alpha \leq 1$

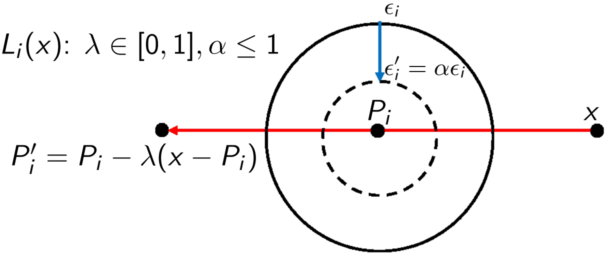

Measuring Communication Success

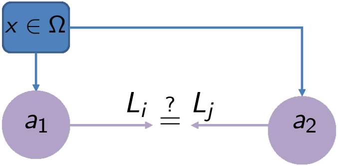

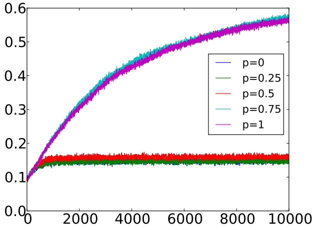

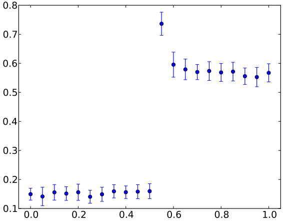

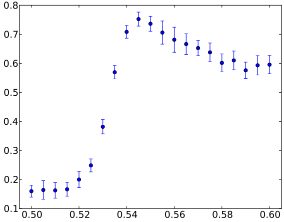

Flexible Specifications

- There is a growing need to formally verify the design of autonomous systems against a predefined specification.
- *Example:* All trajectories must not pass close to either $\mathbf{O_1}$ nor $\mathbf{O_2}$ and must terminate close to $\mathbf{T}$.

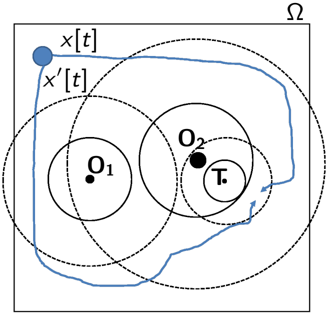

- This specification can be systematically weakened by admitting *borderline* cases.

Conclusions: Enriched Conceptual Models

Man, unlike anything organic or inorganic in the universe, grows beyond his work, walks up the stairs of his concepts, emerges ahead of his accomplishments. [John Steinbeck]

- Vagueness is a multi-faceted phenomenon with different aspects useful in different contexts.
- In all cases we are really referring to *enriched conceptual models*.
- Boolean representations of concepts are flat!
- All the examples we have considered exploit additional structure in predicate definitions.
- But there is usually a computational cost.
- Vagueness $\not\Rightarrow$ cheap back of an envelope calculation!