The Standpoint Theory of Vagueness

Brandon Bennett

School of Computing

University of Leeds

`brandon@comp.leeds.ac.uk`

LoMoReVI, Čejkovice — 17th September 2009

Overview

The presentation will be organised as follows:

- Examination of *vagueness* and the notion of ‘*standpoint*’.
- Basic *supervaluation semantics*.
- Relating *precisifications* to *possible worlds*.
- Semantic interpretation — truth dependent on reality and linguistic convention (à la Austin).
- Probability and Acceptability.
- The *sorites* paradox.

What is Vagueness?

- Vagueness is not Uncertainty.
- Vagueness is semantic (not epistemic) indeterminacy. (*pace* Williamson)
- *Vagueness is the phenomenon that arises because of lack of precise criteria for the applicability of natural terminology.*

Two Sources of Vagueness

Vagueness is actually a composite of (at least) two kinds of semantic indeterminacy:

Two Sources of Vagueness

Vagueness is actually a composite of (at least) two kinds of semantic indeterminacy:

- Conceptual Ambiguity — it is indeterminate what logical combination of objective properties constitute the meaning of a term. Definitions are often controversial (e.g. ‘murder’).

Two Sources of Vagueness

Vagueness is actually a composite of (at least) two kinds of semantic indeterminacy:

- Conceptual Ambiguity — it is indeterminate what logical combination of objective properties constitute the meaning of a term. Definitions are often controversial (e.g. ‘murder’).
- Threshold Vagueness — it is indeterminate to what degree a continuously varying property should hold in order to satisfy a vague qualitative predicate (e.g. tall).

See the example of ‘desert’ in the ClimateView applet.

The Notion of Standpoint

In making an assertion or a coherent series of assertions, one is taking a *standpoint* regarding the applicability of certain linguistic expressions to describing the world.

The Notion of Standpoint

In making an assertion or a coherent series of assertions, one is taking a *standpoint* regarding the applicability of certain linguistic expressions to describing the world.

The key proposal in this analysis is that the notion of ‘standpoint’ is central to the understanding of vagueness — and indeed other ‘non-classical’ aspects of ordinary language semantics.

What is a Standpoint?

A standpoint is a collection of commitments associated with an assertion or a sequence of assertions.

What is a Standpoint?

A standpoint is a collection of commitments associated with an assertion or a sequence of assertions.

Standpoints usually involve beliefs about the world.

What is a Standpoint?

A standpoint is a collection of commitments associated with an assertion or a sequence of assertions.

Standpoints usually involve beliefs about the world.

Taking a standpoint also involves a decision that a particular description is appropriate to the observed state of the world.

What is a Standpoint?

A standpoint is a collection of commitments associated with an assertion or a sequence of assertions.

Standpoints usually involve beliefs about the world.

Taking a standpoint also involves a decision that a particular description is appropriate to the observed state of the world.

Even when the relevant state of the world is clear, the choice of description may be controversial.

People and Standpoints

Vagueness is sometimes discussed in terms of different people having conflicting opinions about the use of a term.

People and Standpoints

Vagueness is sometimes discussed in terms of different people having conflicting opinions about the use of a term.

This is somewhat misleading since even a person thinking privately may be aware that an attribution is not clear cut.

Hence a person may change their standpoint.

People and Standpoints

Vagueness is sometimes discussed in terms of different people having conflicting opinions about the use of a term.

This is somewhat misleading since even a person thinking privately may be aware that an attribution is not clear cut.

Hence a person may change their standpoint.

Moreover this is not necessarily because they think they were mistaken. It can just be that they come to the view that a different standpoint might be more useful for communication purposes.

Different standpoints may be appropriate in different circumstances.

Consistency of Standpoints

Although vague terms can vary in meaning, an important feature of a standpoint is that it *ought* to be consistent.

Consistency of Standpoints

Although vague terms can vary in meaning, an important feature of a standpoint is that it *ought* to be consistent.

That is, while describing things in terms of a given standpoint terms should be applied according to uniform (albeit in some respects arbitrary) criteria.

If I say a person is tall then any person whose height is greater ought also to be judged tall — unless I change my standpoint.

Consistency of Standpoints

Although vague terms can vary in meaning, an important feature of a standpoint is that it *ought* to be consistent.

That is, while describing things in terms of a given standpoint terms should be applied according to uniform (albeit in some respects arbitrary) criteria.

If I say a person is tall then any person whose height is greater ought also to be judged tall — unless I change my standpoint.

Moreover, within a given standpoint constraints exist between the meanings of different terms — a person called ‘short’ cannot also be called ‘tall’.

Supervaluation Semantics

Supervaluation semantics enables a vague language to be logically interpreted by a set of possible precise interpretations (called *precisifications*).

This provides a very general framework within which vagueness can be analysed within a formal representation.

However, by itself it gives no analysis of the kinds of variability that occur in the meanings of natural language concepts.

Supervaluation Models

An possible world/precisificaiton model is a structure

$$\langle W, P, I, \mathcal{R}, \delta \rangle \ ,$$

where:

W is a set of possible worlds,
P is a set of precisifications,
I is a set of individual constants,
$\mathcal{R}$ is a set of relation symbols
and δ is a function from $W \times P \times \ \mathcal{R}$ to tuples of elements of I.

Semantic Definitions

I write $\langle w, p\rangle \Vdash \phi$ to mean that ϕ is true at $\langle w, p\rangle$.

For an atomic proposition $R(x_1, \ldots, x_n)$ I specify that

$$\langle w, p\rangle \Vdash R(x_1, \ldots, x_n) \text{ iff } \langle x_1, \ldots, x_n\rangle \in \delta(w, p, R)$$

The truth-conditions of the Boolean truth functions and quantifiers can then be specified in the usual way:

$$\begin{array}{lll}
\langle w, p\rangle \Vdash (\alpha \wedge \beta) & \text{iff} & \langle w, p\rangle \Vdash \alpha \text{ and } \langle w, p\rangle \Vdash \beta \\
\langle w, p\rangle \Vdash (\alpha \vee \beta) & \text{iff} & \langle w, p\rangle \Vdash \alpha \text{ or } \langle w, p\rangle \Vdash \beta \\
\langle w, p\rangle \Vdash \neg\phi & \text{iff} & \langle w, p\rangle \nVdash \phi \\
\langle w, p\rangle \Vdash \forall x[\phi(x)] & \text{iff} & (\forall i \in I)\{\langle w, p\rangle \Vdash \phi(i)\}
\end{array}$$

Necessity and Unequivocality

The possible world and precisification indices enable one to define two additional operators.

Necessity:

$$\langle w, p\rangle \Vdash \Box\phi \quad \text{iff} \quad (\forall u \in W)\{\langle u, p\rangle \Vdash \phi\}$$

Unequivocality:

$$\langle w, p\rangle \Vdash \mathbf{U}\phi \quad \text{iff} \quad (\forall q \in P)\{\langle w, q\rangle \Vdash \phi\}$$

Parameterised Precisifications

A useful idea is that the range of possible precisifications of a vague language can be parameterised by a (finite) number of values that determine the range of applicability of vague predicates.

Parameterised Precisifications

A useful idea is that the range of possible precisifications of a vague language can be parameterised by a (finite) number of values that determine the range of applicability of vague predicates.

A *parameterised precisification* corresponds to a vector of parameter/value pairs:

```
P = [ pond_vs_lake_area_threshold=200,
      river_linearity_ratio=3,
      ...  ]
```

This fixes a precise interpretation of each vague concept.

Adjectives and Relevant Observables

In specifying a theory (an ontology) characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

Adjectives and Relevant Observables

In specifying a theory (an ontology) characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

To specify that height is positively relevant to tallness, one might write:

$$\mathsf{Rel}^{+}(\mathbf{height}, \mathsf{tall})$$

Adjectives and Relevant Observables

In specifying a theory (an ontology) characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

To specify that height is positively relevant to tallness, one might write:

$$\mathsf{Rel}^{+}(\mathbf{height}, \mathsf{tall})$$

The presence of such a relationship means that a suitable parameterised precisification semantics modelling the word 'tall' must include a threshold on height in relation to tallness.

Similar treatment applies to vague count nouns:

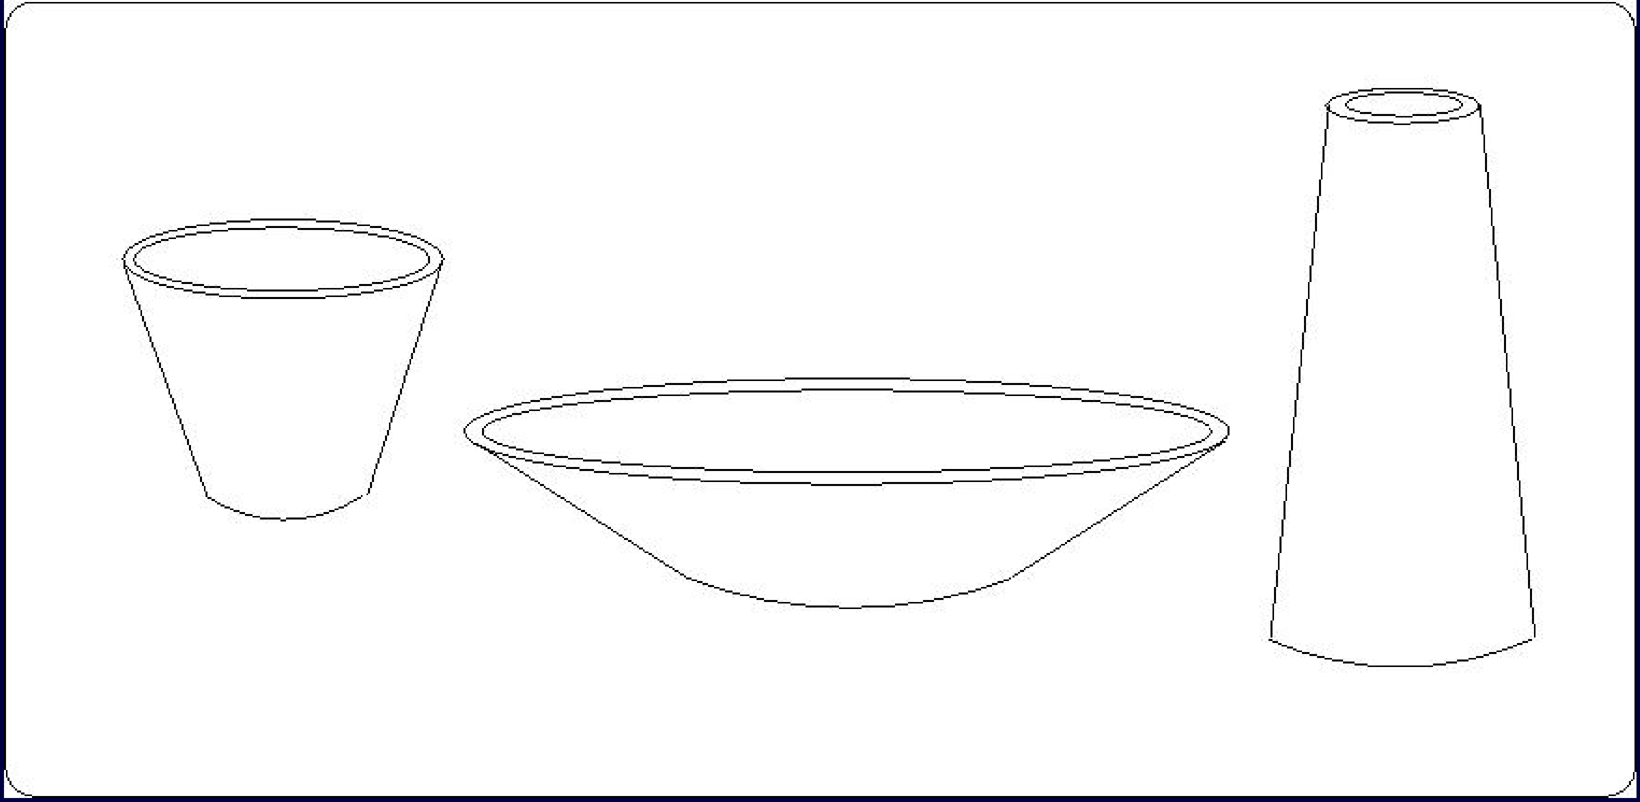

Example objects of the ‘vessel’ domain.

Measurement Structures — Models of the World

A *measurement structure* is a tuple

$$\mathcal{M} = \langle D, M, w \rangle \ ,$$

where

- D is a domain $\{\ldots, e_i, \ldots\}$ of entities;
- $M = \{\ldots, f_i, \ldots\}$ is a set of measurement function symbols;
- $w : M \times (D^n) \to \mathbb{R}$, were n, varies according to the arity of each function symbol $f \in M$.

$w(f, e_1, \ldots, e_n)$ gives the value of n-ary measurement function f, with respect to entities $\langle e_1, \ldots, e_n \rangle$, at world w.

Possible Worlds and Measurement Frames

Any given domain and set of measurement functions over that domain, determines a set of *possible worlds* as follows:

Worlds$(D, M) \ = \{ w \mid \langle D, M, w \rangle$ is a measurement structure $\}$.

A *measurement frame* is a structure that specifies all possible worlds determined by a given measurement structure:

$$\langle D, M, W \rangle \ ,$$

where $W =$ Worlds(D, M).

A Precise Language
of Measurements and Thresholds

Let $\mathcal{L}(M, T, V)$ be the set of formulae of a first-order logical language whose non-logical symbols consist of:

- a finite set $M = \{f_1, \ldots, f_k\}$ of measurement function symbols,
- a finite set, $T = \{\mathrm{t}_1, \ldots \mathrm{t}_l\}$, of *threshold parameter* symbols,
- a denumerable set $V = \{\ldots, x_i, \ldots\}$ of variable symbols,
- inequality relations $<$, $\leq$.

Formulae of $\mathcal{L}(M, T, V)$

Every atomic formula of $\mathcal{L}(M, T, V)$ takes one of the forms:

A1. $f_j(x_1, \ldots, x_n) \leq f_k(y_1, \ldots, x_m)$

A2. $\mathsf{t}_i \leq \mathsf{t}_j$

A3. $\mathsf{t}_i < f_j(x_1, \ldots, x_n)$

A4. $f_j(x_1, \ldots, x_n) < \mathsf{t}_i$

Expressive Capabilities of $\mathcal{L}(M,T,V)$

Although the full language $\mathcal{L}(M,T,V)$ allows arbitrary logical combinations of the different kinds of atoms, certain restrictions on the atoms give rise distinct types of formulae.

Expressive Capabilities of $\mathcal{L}(M, T, V)$

Although the full language $\mathcal{L}(M, T, V)$ allows arbitrary logical combinations of the different kinds of atoms, certain restrictions on the atoms give rise distinct types of formulae.

Formulae containing only atoms of the form **A1** express constraints on the physical structure of the world.
Such formulae can be used to express a *physical theory*.

Expressive Capabilities of $\mathcal{L}(M,T,V)$

Although the full language $\mathcal{L}(M,T,V)$ allows arbitrary logical combinations of the different kinds of atoms, certain restrictions on the atoms give rise distinct types of formulae.

Formulae containing only atoms of the form **A1** express constraints on the physical structure of the world.
Such formulae can be used to express a *physical theory*.

Formulae of the form **A2** express ordering constraints between thresholds.
For example: t_short $\leq$ t_tall.
A set of such formulae will be called a *threshold constraint theory*.

Expressive Capabilities of $\mathcal{L}(M, T, V)$

Although the full language $\mathcal{L}(M, T, V)$ allows arbitrary logical combinations of the different kinds of atoms, certain restrictions on the atoms give rise distinct types of formulae.

Formulae containing only atoms of the form **A1** express constraints on the physical structure of the world.
Such formulae can be used to express a *physical theory*.

Formulae of the form **A2** express ordering constraints between thresholds.
For example: $\text{t\_short} \leq \text{t\_tall}$.
A set of such formulae will be called a *threshold constraint theory*.

Formulae of the form **A3** and **A4** express facts relating observables to threshold parameters.
E.g. $\mathbf{height}(x) \leq \text{t\_tall}$

Predicate Grounding

Each formula of $\mathcal{L}(M, T, V)$ defines a predicate of arity n, where n is the number of free variables in the formula. Hence, we can extend L by defining new predicate symbols by means of *predicate grounding* formulae of the form

(PG) $A x_1, \ldots, x_n [R(x_1, \ldots, x_n) \leftrightarrow \Phi(\mathsf{t}_1, \ldots, \mathsf{t}_m, x_1, \ldots, x_n)]$,

where $\Phi(\mathsf{t}_1, \ldots, \mathsf{t}_m, x_1, \ldots, x_n)]$ is any formula in $\mathcal{L}(M, T, V)$ incorporating parameters $\mathsf{t}_1, \ldots, \mathsf{t}_m$ and with free variables $x_1, \ldots, x_n$.

E.g. the predicate $\mathsf{Tall}(x)$ could be defined by:

$$\forall x [\mathsf{Tall}(x) \leftrightarrow (\mathsf{t\_tall} \leq \mathbf{height}(x))]$$

Extended Language with Predicates and Constants

Let $\mathcal{L}(M, T, V, R, N)$ be the language obtained by supplementing the vocabulary of $\mathcal{L}(M, T, V)$ by a set of predicate symbols R and a set of constant symbols N.

The set of atomic formulae is extended to include those of the form

A5. $R_i(\alpha_1, \ldots, \alpha_n)$,

where each $\alpha_i \in (V \cup N)$.

Predicate Grounding Theory

Given a language $\mathcal{L}(M, T, V, R, N)$, a *predicate grounding theory* θ for this language is a set of formulae of the form **PG**, containing one formula for each relation $R_i \in R$.

Thus, the predicate grounding theory defines every relation in R in terms of a formula of the sub-language $\mathcal{L}(M, T, V)$.

General Grounding Theories

A *general grounding theory* θ for $\mathcal{L}(M, T, V, R, N)$ consists of the conjunction of:

- a physical theory,
- a threshold constraint theory,
- and a predicate grounding theory.

General Grounding Theories

A *general grounding theory* θ for $\mathcal{L}(M, T, V, R, N)$ consists of the conjunction of:

- a physical theory,
- a threshold constraint theory,
- and a predicate grounding theory.

Let Θ be the set of all general grounding theories for $\mathcal{L}$.

Since each $\theta \in \Theta$ includes a predicate grounding theory, we can define a function,

$$\text{Def} : \Theta \times R \rightarrow \mathcal{L}$$

such that $\text{Def}(\theta, R)$ gives the definition of the relation R.
(It will have a number of free variables equal to the arity of R.)

Interpretation Function

The semantic interpretation function,

$$[\![\chi]\!]^{w,p,\theta}_{\mathfrak{M}}$$

gives the denotation of any formula or term χ of the language relative to:

- a given model $\mathfrak{M}$,
- a possible world $w \in W$,
- a precisification $p \in P$,
- and a general grounding theory θ.

Some Notation

The following notation is used to specify the interpretation function for the language:

- $\delta : (N \cup V) \rightarrow D$ is a general denotation function which gives a the denotation of both name and variable symbols, such that $\delta(\alpha) = \kappa(\alpha)$ if $\alpha \in N$ and $\delta(\alpha) = \xi(\alpha)$ if $\alpha \in V$
- $\mathfrak{M} \overset{x}{\sim} \mathfrak{M}'$ means that models $\mathfrak{M}$ and $\mathfrak{M}'$ are identical except for their variable assignment functions ξ and ξ'. And moreover, these assignment functions are identical, except that they may differ in the value assigned to the variable x.
- $\mathsf{Subst}([x_1 \Rightarrow \alpha_1, \ldots, x_n \Rightarrow \alpha_n], \phi)$ refers to the formula resulting from ϕ after replacing each variable x_i by α_i.

Interpretation Function

Most clauses of the interpretation function are standard and do not depend on the indices w, p, θ:

- $[\![\neg\phi]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$ if $[\![\phi]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{f}$, otherwise $= \mathbf{f}$;
- $[\![\phi \wedge \psi]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$ if $[\![\phi]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$ and $[\![\psi]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$, otherwise $= \mathbf{f}$;
- $[\![\forall x[\psi]]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$ if $[\![\psi]\!]^{w,p,\theta}_{\mathfrak{M}'}$ for all $\mathfrak{M}'$ such that $\mathfrak{M} \overset{x}{\sim} \mathfrak{M}'$, otherwise $= \mathbf{f}$;
- $[\![\mathsf{n}]\!]^{w,p,\theta}_{\mathfrak{M}} = \kappa(\mathsf{n})$;
- $[\![\tau_i \leq \tau_j]\!]^{w,p,\theta}_{\mathfrak{M}} = \mathbf{t}$ if $[\![\tau_i]\!]^{w,p,\theta}_{\mathfrak{M}}$ is less than or equal to $[\![\tau_j]\!]^{w,p,\theta}_{\mathfrak{M}}$, otherwise $= \mathbf{f}$.

Interpretation Dependence on the Possible World

The value of measurement functions does of course depend on the possible world in which the measurement is made; and hence, their interpretation is relative to the w index:

- $[\![f(\alpha_1, \ldots, \alpha_n)]\!]_{\mathfrak{M}}^{w,p,\theta} = w(f, \delta(\alpha_1), \ldots, \delta(\alpha_n))$;

Interpretations Involving Vagueness and Ambiguity

The following clause takes account of vagueness by interpreting threshold parameters in accordance with the precisification index:

(I-thresh) $[\![\mathsf{t}]\!]^{w,p,\theta}_{\mathfrak{M}} = p(\mathsf{t})$

The next clause handles conceptual ambiguity by interpreting predicates according to the grounding theory index:

(I-pred) $[\![R(\alpha_1, \ldots, \alpha_n)]\!]^{w,p,\theta}_{\mathfrak{M}} =$

$$[\![\mathsf{Subst}([x_1 \Rightarrow \alpha_1, \ldots, x_n \Rightarrow \alpha_n], \mathsf{Def}(\theta, R))]\!]^{w,p,\theta}_{\mathfrak{M}}$$

The Satisfaction Relation

On the basis of the interpretation function, a semantic satisfaction relation can be defined by

$$\mathfrak{M}, \langle w, p \rangle \Vdash_{\theta} \phi \quad \text{iff} \quad [\![\phi]\!]_{\mathfrak{M}}^{w,p,\theta} = \mathbf{t}\ .$$

This says that formula ϕ is true in model $\mathfrak{M}$, at world w and precisification p, with predicate grounding theory θ.

Interpretation Sets

The *interpretation set* of a proposition relative to a model $\mathfrak{M}$ and grounding theory θ is given by:

$$[\![\phi]\!]^{\theta}_{\mathfrak{M}} = \{\langle w, p \rangle \mid (\mathfrak{M}, \langle w, p \rangle \Vdash_{\theta} \theta \wedge \phi)\}$$

Interpretation Sets

The *interpretation set* of a proposition relative to a model $\mathfrak{M}$ and grounding theory θ is given by:

$$[\![\phi]\!]^{\theta}_{\mathfrak{M}} = \{\langle w, p\rangle \mid (\mathfrak{M}, \langle w, p\rangle \Vdash_{\theta} \theta \wedge \phi)\}$$

Note that we require that the conjunction $\theta \wedge \phi$ is entailed, rather than simply ϕ.

This is because the satisfaction relation $\Vdash_{\theta}$ only enforces the predicate grounding part of the general grounding theory θ.

Adding the θ conjunct ensures that the interpretation set only includes possible worlds and precisifications that also satisfy the physical theory and threshold constraints associated with θ.

A Formal Model of Standpoint

A *standpoint* is modelled by a tuple,

$$\langle B, A, \theta, \Gamma \rangle \ ,$$

where:

- $B \subseteq W$ is the agent's *belief set* — i.e. the set of possible worlds that are compatible with the agent's beliefs,
- $A \subseteq P$ is the set of precisifications that the agent considers to be *admissible*,
- θ is the general grounding theory that characterises the predicate definitions, physical theory and threshold constraints employed by the agent.
- $\Gamma \subseteq \mathcal{L}$ is a set of propositional judgements held to be true by the agent.

Standpoint Frames

A *standpoint frame* characterises the epistemic state and linguistic conventions of an agent, but not the actual judgement that (s)he holds.

Thus, it is a structure

$$\langle B, A, \theta \rangle \ ,$$

where B, A, θ satisfy the same stipulations as for a standpoint.

Standpoint Relative Propositional Attitudes

By comparing the interpretation set of a proposition ϕ to an agent's standpoint we can define various attitudes that the agent might have towards the proposition:

- $\mathfrak{M}, \langle B, A, \theta \rangle \Vdash \mathsf{Definitely}(\phi)$ iff $(B \times A) \subseteq [\![\phi]\!]^{\theta}_{\mathfrak{M}}$.
- $\mathfrak{M}, \langle B, A, \theta \rangle \Vdash \mathsf{CouldSay}(\phi)$ iff
  $(B \times \{p\}) \subseteq [\![\phi]\!]^{\theta}_{\mathfrak{M}}$, for some $p \in A$.
- $\mathfrak{M}, \langle B, A, \theta \rangle \Vdash \mathsf{CouldBe}(\phi)$ iff
  $(\{w\} \times A) \subseteq [\![\phi]\!]^{\theta}_{\mathfrak{M}}$, for some $w \in B$.
- $\mathfrak{M}, \langle B, A, \theta \rangle \Vdash \mathsf{CouldBeSay}(\phi)$ iff
  $\langle w, p \rangle \in [\![\phi]\!]^{\theta}_{\mathfrak{M}}$, for some $w \in B$ and some $p \in A$.

Stable and Coherent Standpoints

The formal definition of standpoint that I have given is perhaps too general, in that it does not constrain the judgement set to be consistent with the standpoint frame.

We can take account of this by restricting attention to more reasonable standpoints:

$$\mathfrak{M} \Vdash \mathsf{Stable}(\langle B, A, \theta, \Gamma \rangle) \quad \text{iff} \quad \mathfrak{M}, \langle B, A, \theta \rangle \Vdash \ \mathsf{Definitely}(\Gamma)$$

Stable and Coherent Standpoints

The formal definition of standpoint that I have given is perhaps too general, in that it does not constrain the judgement set to be consistent with the standpoint frame.

We can take account of this by restricting attention to more reasonable standpoints:

$$\mathfrak{M} \Vdash \text{Stable}(\langle B, A, \theta, \Gamma\rangle) \quad \text{iff} \quad \mathfrak{M}, \langle B, A, \theta\rangle \Vdash \ \text{Definitely}(\Gamma)$$

A weaker restriction could also be useful for some purposes:

$$\mathfrak{M} \Vdash \text{Coherent}(\langle B, A, \theta, \Gamma\rangle) \quad \text{iff} \quad \mathfrak{M}, \langle B, A, \theta\rangle \Vdash \ \text{CouldBeSay}(\Gamma)$$

Fuzzy Logic

So far I have not mentioned the most popular approach to vagueness employed in AI: *fuzzy logic*.

In fuzzy logic the absolute truth values of classical logic are replaced by numerical values akin to probabilities, which are supposed to represent degrees of truth.

For example, the fact that I am not particularly tall might be represented by:

$$\mu(\mathsf{Tall}(\mathsf{brandon})) = 0.4$$

where the fuzzy interpretation function μ gives the degree of truth of a proposition.

Fuzzy Logic

So far I have not mentioned the most popular approach to vagueness employed in AI: *fuzzy logic*.

In fuzzy logic the absolute truth values of classical logic are replaced by numerical values akin to probabilities, which are supposed to represent degrees of truth.

For example, the fact that I am not particularly tall might be represented by:

$$\mu(\text{Tall}(\text{brandon})) = 0.4$$

where the fuzzy interpretation function μ gives the degree of truth of a proposition.

But, what does this number mean?

Vagueness and Probability

Fuzzy logicians tend to take the numerical truth values as primitive. But if pressed they tend to say something like:

> *The fuzzy truth value of a proposition is a measure of the likelihood that a person would assent to that proposition.*

The standpoint semantics enables us to use a rather different approach to modelling this assent likelihood.

Probability Distribution Over Precisifications

In the semantics defined above, the space of precisifications of a vague language was modelled simply as a set of possible assignments to threshold parameters.

However, it is clear that some thresholds are much more appropriate than others.

For instance, a suitable threshold for tallness (of a human male) might be around 6”.

We can model this by introducing a probability distribution over the set of precisifications.

Probabilistic Standpoint Frames

A *probabilistic standpoint frame* is a structure:

$$\langle B, \mathsf{Acc}, \theta \rangle \ ,$$

where $\mathsf{Acc} : P \rightarrow [0 \ldots 1]$,
such that Acc is a probability distribution over P.

A Measure of Acceptability

The probability distribution over P, encapsulated in the function Acc enables us to define the following measure of the *acceptability* of a formula to an agent:

$$\mathsf{Acc}_{\mathfrak{M}}(\langle B, \mathsf{Acc}, \theta\rangle, \phi) \;=\; \Sigma\{\mathsf{Acc}(p) \mid\; (\exists w \in B)[\mathfrak{M}, \langle w, p\rangle \Vdash_{\theta} \phi]\}$$

Adding Probability Over the Belief Set

To go one step further, we could also replace the belief set of an agent probability distribution over possible worlds, corresponding to the likelihood that an agent assigns to possible states of the world.

A standpoint frame would then be a structure

$$\langle \mathsf{BelProb}, \mathsf{Acc}, \theta \rangle \ ,$$

And the acceptability measure would become:

$$\mathsf{Acc}_{\mathfrak{M}}(\langle \mathsf{BelProb}, \mathsf{Acc}, \theta \rangle, \phi) \ = \ \Sigma\{\mathsf{BelProb}(w) \cdot \mathsf{Acc}(p) \mid \ \mathfrak{M}, \langle w, p \rangle \Vdash_{\theta} \phi\}$$

Acceptability of ‘Tall’

Ascription of a vague predicate, such as ‘tall’ is interpreted as equivalent to a constraint-like proposition asserting that threshold for tallness must lie below the height of the person who is judged to be tall.

- $\mathsf{Tall}(x) \implies \tau(\mathsf{Tall}) \in [0 \ldots \mathbf{height}(x))$
- $\neg\mathsf{Tall}(x) \implies \tau(\mathsf{Tall}) \in (\mathbf{height}(x) \ldots \infty)$

The *accepttability* of $\mathsf{Tall}(x)$ might be evaluated as the probability that the threshold τ(Tall) lies in the range $[0 \ldots \mathbf{height}(x))$.

Calculating Acceptability of Conjunctions

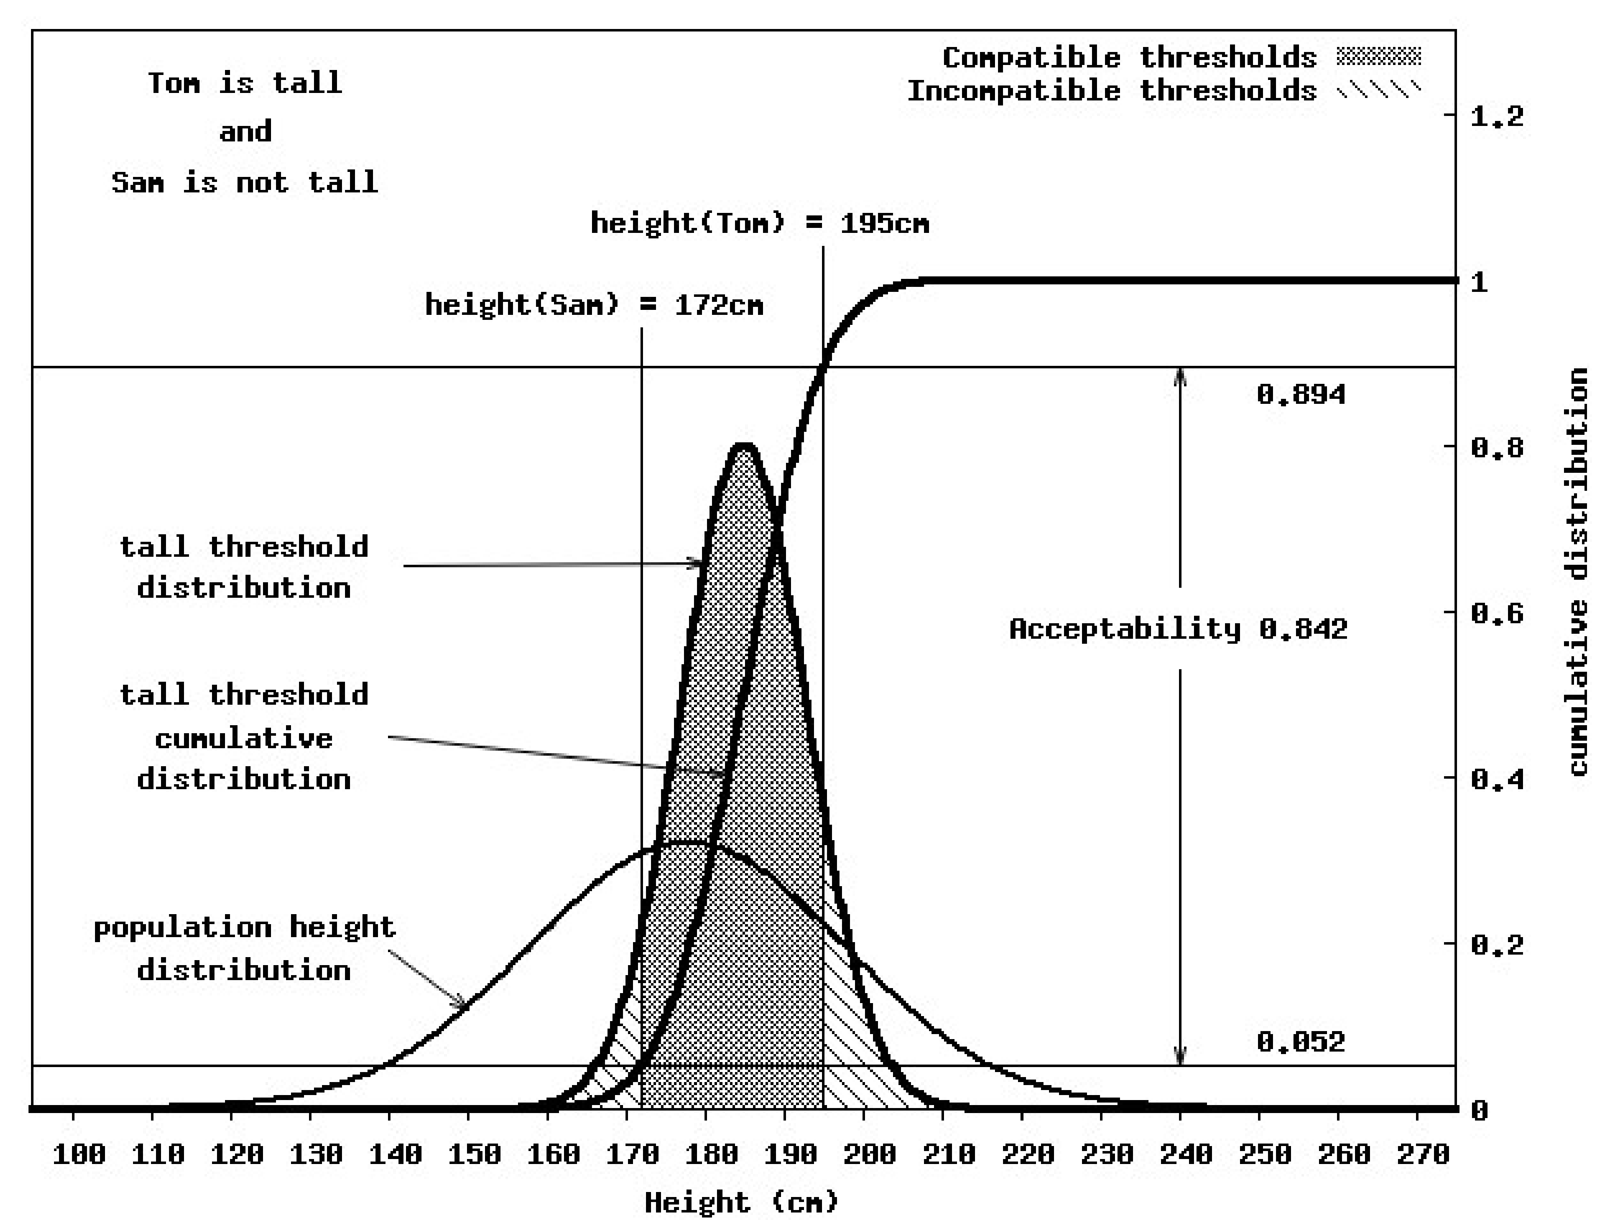

The Sorites Paradox

The acceptability measures just defined enable one to give a new account of how the *sorites paradox* arises.

The Sorites Paradox

The acceptability measures just defined enable one to give a new account of how the *sorites paradox* arises.

Note that propositions with high acceptability must be true over a wide range of possible precisifications.

The Sorites Paradox

The acceptability measures just defined enable one to give a new account of how the *sorites paradox* arises.

Note that propositions with high acceptability must be true over a wide range of possible precisifications.

Suppose I am confronted with a line of people, such that the first person is clearly not tall, but each successive person is slightly taller than the previous one. I am then asked to make a divide between the tall and non-tall people.

The Sorites Paradox

The acceptability measures just defined enable one to give a new account of how the *sorites paradox* arises.

Note that propositions with high acceptability must be true over a wide range of possible precisifications.

Suppose I am confronted with a line of people, such that the first person is clearly not tall, but each successive person is slightly taller than the previous one. I am then asked to make a divide between the tall and non-tall people.

Wherever I make the divide there will be a very small range of precisifications which satisfy my classification.

The Sorites Paradox

The acceptability measures just defined enable one to give a new account of how the *sorites paradox* arises.

Note that propositions with high acceptability must be true over a wide range of possible precisifications.

Suppose I am confronted with a line of people, such that the first person is clearly not tall, but each successive person is slightly taller than the previous one. I am then asked to make a divide between the tall and non-tall people.

Wherever I make the divide there will be a very small range of precisifications which satisfy my classification.

Hence, the acceptability of my claim must be very low.

Conclusions

I have presented a general framework for modelling the semantics of vagueness.

The semantics takes account of both conceptual ambiguity (via the grounding theory index θ) and threshold vagueness (via the precisification index p).

Adding probabilities over the space of precisifications provides a mechanism for a statistical treatment of vagueness that is very different from that provided by fuzzy logic.

The ‘acceptability’ measure I have defined also seems to give some insight into the nature of the sorites paradox.

Old and Extra Slides Follow

How do we chose our words?

The cognitive and judgemental processes involved in making a decision about how to describe the world are complex, and the details go well beyond the scope of this analysis.

Here I assume that the agents involved are honest, and that their primary concern is to choose words which they believe will effectively communicate meaning. That is, other agents will agree with the choice of words.

Representing Judgements

Every standpoint is associated with a set of judgements.

I represent a judgement as a relationship between a *standpoint* and a *proposition*.

Thus, a judgement is represented by a formula of the form:

$$\text{Judge}(s, \phi)$$

Commitment

Judgements imply more than just the words with which they are stated. Propositions have consequences.

I say that a person is *committed* to all consequences (whether they know them or not) of their judgements.

To indicate that the standpoint s is associated with a commitment to proposition ϕ, I write:

$$\mathrm{Commit}(s, \phi)$$

Commitment Axioms

Every judgement implies a commitment:

$$\text{Judge}(s, \phi) \;\rightarrow\; \text{Commit}(s, \phi)$$

Commitment Axioms

Every judgement implies a commitment:

$$\text{Judge}(s, \phi) \;\rightarrow\; \text{Commit}(s, \phi)$$

Logical implications of commitments are also commitments:

$$(\text{Commit}(s, \alpha) \wedge \text{Commit}(s, (\alpha \rightarrow \beta)) \;\models\; \text{Commit}(s, \beta)$$

Commitment Axioms

Every judgement implies a commitment:

$$\text{Judge}(s,\phi) \;\rightarrow\; \text{Commit}(s,\phi)$$

Logical implications of commitments are also commitments:

$$(\text{Commit}(s,\alpha) \wedge \text{Commit}(s,(\alpha \rightarrow \beta)) \;\models\; \text{Commit}(s,\beta)$$

One is always committed to what is logically true:

$$(\models_{CL} \phi) \;\models\; \text{Commit}(s,\phi)$$

Commitment to a Background Theory

One is always committed to some theory, Θ, of necessary semantic constraints:

$$\models \ \text{Commit}(s, \Theta)$$

Commitment to a Background Theory

One is always committed to some theory, Θ, of necessary semantic constraints:

$$\models \text{Commit}(s, \Theta)$$

Θ may be thought of as an *ontology* describing necessary relationships between the meanings of terms.

Commitment to a Background Theory

One is always committed to some theory, Θ, of necessary semantic constraints:

$$\models \text{Commit}(s, \Theta)$$

Θ may be thought of as an *ontology* describing necessary relationships between the meanings of terms.

In some circumstances one might modify Θ; but this is a more radical change of conceptualisation than a simple shift of standpoint.

Inconsistent Standpoints

Just because standpoints *ought* to be consistent does not mean they always are.

Inconsistent Standpoints

Just because standpoints *ought* to be consistent does not mean they always are.

Within my notation it is easy to identify inconsistent standpoints:

$$\text{Commit}(s, \phi) \wedge \text{Commit}(s, \neg\phi) \models \neg\text{Consistent}(s)$$

Inconsistent Standpoints

Just because standpoints *ought* to be consistent does not mean they always are.

Within my notation it is easy to identify inconsistent standpoints:

$$\text{Commit}(s, \phi) \wedge \text{Commit}(s, \neg\phi) \models \neg\text{Consistent}(s)$$

If a rational agent becomes aware that their standpoint is inconsistent they ought to change it.

The Sorites Paradox

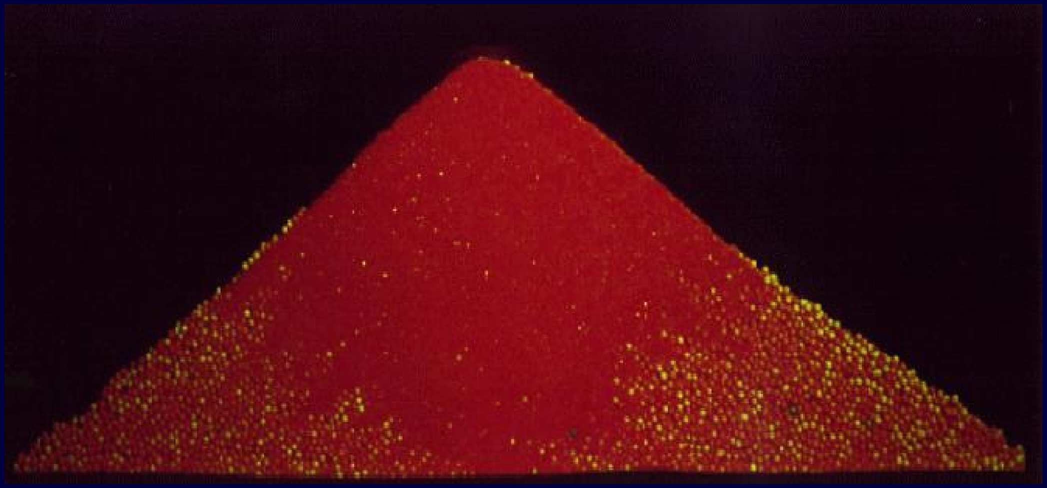

How many grains make a *heap*?

Logical Structure of the Paradox

Assuming we have various collections of grains organised in as heap-like an arrangement as possible.

The following argument *seems* to be valid:

1. $\forall x[(\text{num\_grains}(x) = 10{,}000{,}000) \rightarrow \text{heap}(x)]$
2. $\forall x \forall y[(\text{heap}(x) \wedge (\text{num\_grains}(y) = \text{num\_grains}(x) - 1)) \rightarrow \text{heap}(y)]$
3. $\therefore \quad \forall x[(\text{num\_grains}(x) = 0) \rightarrow \text{heap}(x)]$

Obligations

The solution proposed in this theory is that implications based on insignificant or indiscriminable differences are of a weaker kind than those based on purely logical inference.

I call this kind of implication an *obligation*.

Obligations

The solution proposed in this theory is that implications based on insignificant or indiscriminable differences are of a weaker kind than those based on purely logical inference.

I call this kind of implication an *obligation*.

The induction step of the sorites now becomes:

- Commit(s, x is a heap) $\wedge$
  Commit(s, x contains one fewer grain than y)
  $\models$ Obligation(s, y is a heap)

Obligations

The solution proposed in this theory is that implications based on insignificant or indiscriminable differences are of a weaker kind than those based on purely logical inference.

I call this kind of implication an *obligation*.

The induction step of the sorites now becomes:

- Commit$(s, x$ is a heap$)$ $\wedge$
  Commit$(s, x$ contains one fewer grain than $y)$
  $\models$ Obligation$(s, y$ is a heap$)$

As long as one is not committed to the universalised form of the second premiss, the false conclusion does not arise.

Obligation Inference Rules

Every commitment implies an obligation:

$$\text{Commit}(s, \phi) \models \text{Obligation}(s, \phi)$$

Obligations are closed under implication:

$$\text{Obligation}(s, \alpha) \wedge \text{Obligation}(\alpha \rightarrow \beta) \models \text{Obligation}(s, \beta)$$

Obligation Inference Rules

Every commitment implies an obligation:

$$\text{Commit}(s, \phi) \models \text{Obligation}(s, \phi)$$

Obligations are closed under implication:

$$\text{Obligation}(s, \alpha) \wedge \text{Obligation}(\alpha \rightarrow \beta) \models \text{Obligation}(s, \beta)$$

We can identify *incoherent* obligations as follows:

$$\text{Obligation}(s, \phi) \wedge \text{Obligation}(s, \neg\phi) \models \neg\text{Coherent}(s)$$

Diagnosis of the Paradox

The inferences that give rise to slippery slope sorites arguments are a matter of ‘obligation’ rather than strict logical deduction.

Diagnosis of the Paradox

The inferences that give rise to slippery slope sorites arguments are a matter of ‘obligation’ rather than strict logical deduction.

Vague concepts are such that certain (unusual) situations cannot be fully described fully in terms of these concepts, without incoherent obligations arising.

Diagnosis of the Paradox

The inferences that give rise to slippery slope sorites arguments are a matter of ‘obligation’ rather than strict logical deduction.

Vague concepts are such that certain (unusual) situations cannot be fully described fully in terms of these concepts, without incoherent obligations arising.

In most circumstances we are not pushed into exhaustive classification of sets of very similar samples. So, on the whole we can take a coherent standpoint on how to describe the situation at hand.

Semantic Definitions

I write $\langle w, p\rangle \Vdash \phi$ to mean that ϕ is true at $\langle w, p\rangle$.

For an atomic proposition $R(x_1, \ldots, x_n)$ I specify that

$$\langle w, p\rangle \Vdash R(x_1, \ldots, x_n) \ \text{ iff } \ \langle x_1, \ldots, x_n\rangle \in \delta(w, p, R)$$

The truth-conditions of the Boolean truth functions and quantifiers can then be specified in the usual way:

$$\begin{array}{lcl}
\langle w, p\rangle \Vdash (\alpha \wedge \beta) & \text{iff} & \langle w, p\rangle \Vdash \alpha \ \text{ and } \ \langle w, p\rangle \Vdash \beta \\
\langle w, p\rangle \Vdash (\alpha \vee \beta) & \text{iff} & \langle w, p\rangle \Vdash \alpha \ \text{ or } \ \langle w, p\rangle \Vdash \beta \\
\langle w, p\rangle \Vdash \neg\phi & \text{iff} & \langle w, p\rangle \nVdash \phi \\
\langle w, p\rangle \Vdash \forall x[\phi(x)] & \text{iff} & (\forall i \in I)\{\langle w, p\rangle \Vdash \phi(i)\}
\end{array}$$

Necessity and Unequivocality

The possible world and precisification indices enable one to define two additional operators.

Necessity:

$$\langle w, p\rangle \Vdash \Box\phi \quad \text{iff} \quad (\forall u \in W)\{\langle u, p\rangle \Vdash \phi\}$$

Unequivocality:

$$\langle w, p\rangle \Vdash \mathbf{U}\phi \quad \text{iff} \quad (\forall q \in P)\{\langle w, q\rangle \Vdash \phi\}$$

Parameterised Precisifications

The key idea is that the range of possible precisifications of a vague language can be parameterised by a (finite) number of values that determine the range of applicability of vague predicates.

Parameterised Precisifications

The key idea is that the range of possible precisifications of a vague language can be parameterised by a (finite) number of values that determine the range of applicability of vague predicates.

A *parameterised precisification* corresponds to a vector of parameter/value pairs:

```
P = [ pond_vs_lake_area_threshold=200,
      river_linearity_ratio=3,
      ...  ]
```

This fixes a precise interpretation of each vague concept.

Parameterised Precisifications

The key idea is that the range of possible precisifications of a vague language can be parameterised by a (finite) number of values that determine the range of applicability of vague predicates.

A *parameterised precisification* corresponds to a vector of parameter/value pairs:

```
P = [ pond_vs_lake_area_threshold=200,
      river_linearity_ratio=3,
      ...  ]
```

This fixes a precise interpretation of each vague concept.

(Note: In my earlier work standpoints were identified directly with parameterised precisifications.)

Adjectives and Relevant Observables

In specifying an ontology characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

Adjectives and Relevant Observables

In specifying an ontology characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

To specify that height is positively relevant to tallness, one might write:

$$\mathsf{Rel}^{+}(\mathbf{height}, \mathsf{tall})$$

Adjectives and Relevant Observables

In specifying an ontology characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

To specify that height is positively relevant to tallness, one might write:

$$\mathsf{Rel}^{+}(\mathbf{height}, \mathsf{tall})$$

The presence of such a relationship means that a suitable parameterised precisification semantics modelling the word ‘tall’ must include a threshold on height in relation to tallness.

Adjectives and Relevant Observables

In specifying an ontology characterising the meanings of vague terms, a key relationship that one needs to identify is which *observables* are relevant to a particular vague adjective.

To specify that height is positively relevant to tallness, one might write:

$$\mathsf{Rel}^{+}(\mathbf{height}, \mathsf{tall})$$

The presence of such a relationship means that a suitable parameterised precisification semantics modelling the word ‘tall’ must include a threshold on height in relation to tallness.

My thesis is that vague count nouns inherit their vagueness by being defined in terms of vague adjectives.

Relating Standpoints to Parametrised Supervaluation Semantics

If we take the set of propositions associated with a standpoint s, we can determine the set of indices that are compatible with s:

$$\mathcal{P}_s \;=\; \{\langle w, p\rangle \mid \forall\phi[\mathsf{Judge}(s,\phi) \rightarrow (\langle w, p\rangle \Vdash \phi)]\;\}$$

Relating Standpoints to Parametrised Supervaluation Semantics

If we take the set of propositions associated with a standpoint s, we can determine the set of indices that are compatible with s:

$$\mathcal{P}_s \;=\; \{\langle w, p\rangle \mid \forall\phi[\mathsf{Judge}(s,\phi) \rightarrow (\langle w, p\rangle \Vdash \phi)]\;\}$$

Similarly, given a particular world w, we can identify the set $\mathcal{P}_{ws}$ of precisifications that are compatible with standpoint s at that world:

$$\mathcal{P}_{ws} \;=\; \{p \mid \forall\phi[\mathsf{Judge}(s,\phi) \rightarrow (\langle w, p\rangle \Vdash \phi)]\;\}$$

Truth Relative to a Standpoint

The set $\mathcal{P}_{ws}$ can be used to determine the truth of an arbitrary proposition ϕ relative to standpoint s at possible world w:

- $[\![\phi]\!]^s_w = \mathbf{t}$ iff $\langle w, p \rangle \Vdash \phi$ for all $p \in \mathcal{P}_{ws}$
- $[\![\phi]\!]^s_w = \mathbf{f}$ iff $\langle w, p \rangle \Vdash \neg\phi$ for all $p \in \mathcal{P}_{ws}$
- $[\![\phi]\!]^s_w = \mathbf{i}$ iff $([\![\phi]\!]^s_w \neq \mathbf{t}$ and $[\![\phi]\!]^s_w \neq \mathbf{f})$

Truth Relative to a Standpoint

The set $\mathcal{P}_{ws}$ can be used to determine the truth of an arbitrary proposition ϕ relative to standpoint s at possible world w:

- $[\![\phi]\!]^s_w = \mathbf{t}$ iff $\langle w, p \rangle \Vdash \phi$ for all $p \in \mathcal{P}_{ws}$
- $[\![\phi]\!]^s_w = \mathbf{f}$ iff $\langle w, p \rangle \Vdash \neg\phi$ for all $p \in \mathcal{P}_{ws}$
- $[\![\phi]\!]^s_w = \mathbf{i}$ iff $([\![\phi]\!]^s_w \neq \mathbf{t}$ and $[\![\phi]\!]^s_w \neq \mathbf{f})$

Thus, those propositions whose truth is not determined by a particular standpoint at a given world are assigned the *indefinite* truth value, **i**.

Conclusions

I have attempted to shed light on the semantics of vague terminology by means of a logical analysis of the concept of a *standpoint* held in relation to a linguistic assertion.

This seems to lead to an explanatory account of certain linguistic phenomena.

Although the theory is somewhat complex, I believe it is sufficiently well-defined to be implemented within computational information systems.

Much further work remains to be done.