




Lund University
Cognitive Science

GEOMETRY OF MEANING

SEMANTICS BASED ON CONCEPTUAL SPACES

PETER GÄRDENFORS




The enigma of language learning

- After finishing high school a teenager understands about 60.000 words of her mother tongue
- From the age of two children learn to understand 8-10 new words per day
- A single example of a word is often sufficient to grasp its meaning
- We don't know how this process works
- My starting point: Word meanings have structures that speed up learning

1st paradigm: Symbolism

The computer as a metaphor for cognition

- The brain functions as a Turing machine
- The brain is an information-processing system
- The brain has a central processor and memory units
- The brain works on the basis of a (symbolic) code by which all information is represented





Symbolism

- The Turing machine as a symbol manipulator
- Information is captured in *predicates* (*variables* in a computer program) plus logical constructions
- Concepts are defined in terms of necessary and sufficient conditions using basic predicates
- $GM(x) := F(x) \ \& \ \exists (y) \ \exists (z)(C(y,x)\&C(z,y))$
- General problems: Where do the basic predicates come from? How can you *learn* the meaning of a predicate?





2nd paradigm: Connectionism

Cognitive processes can be modelled in artificial neural networks



Basic metaphor: Cognitive processes can be identified with the input and output activities of the neurons in the brain

Connectionism




- Based on (uninterpreted) inputs from receptors
- Distributed representations by dynamic connection weights
- *Problems:* What is represented in the network?
- Learning is in general very slow

3rd paradigm: Spatial models

Cognition can be modelled in topological and geometrical structures

- Conceptual spaces
- Information is organized in spatial structures
- Distances represent inverse similarity



Geometry of Meaning

- (1) What is semantics?
- (2) Conceptual spaces and semantic domains
- (3) Semantics as a meeting of minds
- (4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions
- (5) A cognitive theory of actions and events
- (6) Compositionality of meanings

Geometry of Meaning

(1) What is semantics?

(2) Conceptual spaces and semantic domains

(3) Semantics as a meeting of minds

(4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions

(5) A cognitive theory of actions and events

(6) Compositionality of meanings

Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?

Criteria for a semantic theory


- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?

Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?


Truth-functional semantics

- Realism vs conceptualism (vs nominalism)
- Truth-functional semantics builds on realism
- Extensional semantics



Truth-functional semantics


- Realism vs conceptualism
- Truth-functional semantics builds on realism
- Intensional semantics



Meanings are truth-conditions

Ontology

- Realism vs conceptualism
- Truth-functional semantics builds on realism
- Cognitive semantics builds on conceptualism



Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?

Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?

”Meanings ain’t in the head”

Putnam:

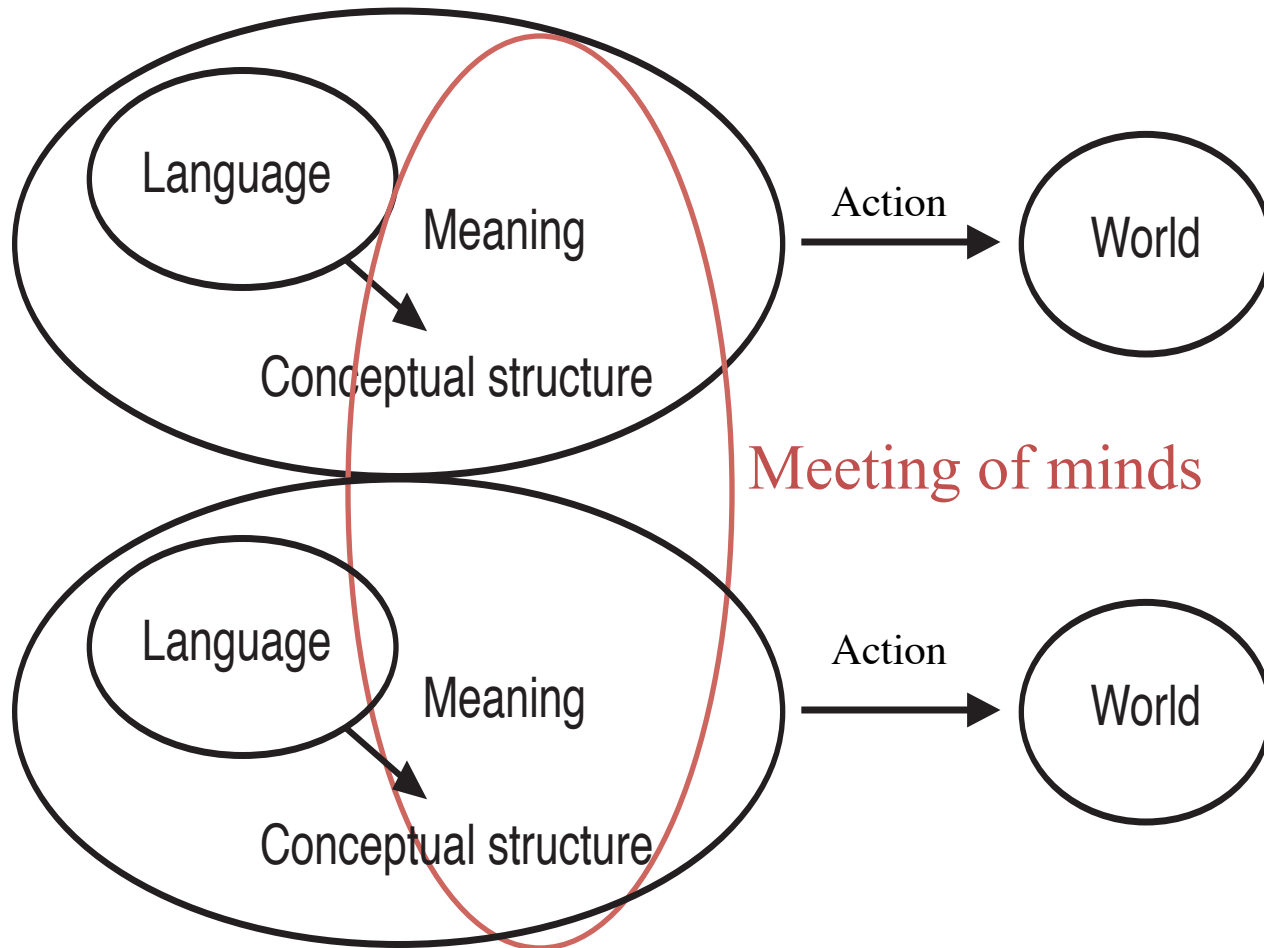
Suppose you are like me and cannot tell an elm from a beech tree. We still say that the extension of 'elm' in my idiolect is the same as the extension of 'elm' in anyone else's, viz., the set of all *elm* trees, and that the set of all beech trees is the extension of 'beech' in both of our idiolects. Thus 'elm' in my idiolect has a different extension from 'beech' in your idiolect (as it should). Is it really credible that this difference in extension is brought about by some difference in our *concepts*? My *concept* of an elm tree is exactly the same as my concept of a beech tree (I blush to confess). (This shows that the identification of meaning 'in the sense of intension' with *concept* cannot be correct, by the way). ... Cut the pie any way you like, meanings just ain't in the *head*!

Sharing mental representations results in an emergent semantics

- If everybody has their own mental space, how can we then talk about a representation being *the* meaning of an expression?
- Semantics is also a *product of communication* – meanings arise as a result of communicative interactions – meanings emerge in the heads
- Sharing of meaning puts *constraints* on individual meanings
- Socio-cognitive approach

Semantics as the meeting of minds

Mental structures (different for different individuals)



”Meanings are in the heads”

Putnam:

Suppose you are like me and cannot tell an elm from a beech tree. We still say that the extension of 'elm' in my idiolect is the same as the extension of 'elm' in anyone else's, viz., the set of all *elm* trees, and that the set of all beech trees is the extension of 'beech' in both of our idiolects. Thus 'elm' in my idiolect has a different extension from 'beech' in your idiolect (as it should). Is it really credible that this difference in extension is brought about by some difference in our *concepts*? My *concept* of an elm tree is exactly the same as my concept of a beech tree (I blush to confess). (This shows that the identification of meaning 'in the sense of intension' with *concept* cannot be correct, by the way). ... Cut the pie any way you like, meanings just ain't in the *head*!

Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? What is the relation between *action* processes and meaning?


How can we talk about what we see?

- The brain is multimodal – operates with different codes – vision, language, music ...
- How can we translate between the codes?
- We are good at it, but we do not know how it works.


Relations between perception and meaning

- In cognitive semantics the carriers of meaning are "image schemas"
- Relates to Gestalt psychology

Climb



Climber




Criteria for a semantic theory

- *Ontological*: What are meanings?
- *Semantic*: What is the relation between communicative acts and their meanings?
- *Epistemological*: How can meanings be learned?
- *Social*: What is the relation between individual speakers and their communal language?
- *Cognitive*: What is the relation between *perceptual* processes and meaning? *What is the relation between *action* processes and meaning?*

Relations between action and meaning

- Verbs normally express actions or results of actions
- Actions are represented cognitively *as force patterns*



Geometry of Meaning

(1) What is semantics?

(2) Conceptual spaces and semantic domains

(3) Semantics as a meeting of minds


(4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions

(5) A cognitive theory of actions and events


(6) Compositionality of meanings

Conceptual spaces

- Consists of a number of dimensions (colour, size, shape, weight, position ...)
- Dimensions have topological or geometric structures




Two linear quality dimensions




Conceptual spaces

- Consists of a number of dimensions (colour, size, shape, weight, position ...)
- Dimensions have topological or geometric structures
- Dimensions are sorted into *domains*
- Less distance means greater similarity of meaning
- Concepts are represented as *convex* regions of conceptual spaces




The color spindle



Conceptual spaces


- Consists of a number of dimensions (colour, size, shape, weight, position ...)
- Dimensions have topological or geometric structures
- Dimensions are sorted into *domains*
- Less distance means greater similarity of meaning
- Concepts are represented as *convex* regions of conceptual spaces




Why convexity?

- Makes learning more efficient

Learning from few examples




Learning from few examples



Why convexity?

- Makes learning more efficient
- Connects to prototype theory

Voronoi tessellation from prototypes



Cognitive economy: Once the space is given, you need only remember the prototypes – the borders can be calculated

Why convexity?

- Makes learning more efficient
- Connects to prototype theory
- Strong empirical support for colour categories (Jäger)

Properties vs. concepts

Properties: A *convex* region in a single domain

Concepts: A set of convex regions in a number of domains; together with (1) prominence values of the domains and (2) information about how the regions in different domains are correlated

Properties vs. concepts


Properties: A convex region in a single domain

Concepts: A set of convex regions in a number of domains; together with (1) prominence values of the domains and (2) information about how the regions in different domains are correlated

An example of an concept: "Apple"

| <i>Domain</i> | <i>Region</i> |
|---------------|---|
| Color | Red-green-yellow |
| Taste | Values for sweetness, sourness etc |
| Shape | "Round" region of shape space |
| Nutrition | Values for sugar, vitamin C, fibres etc |

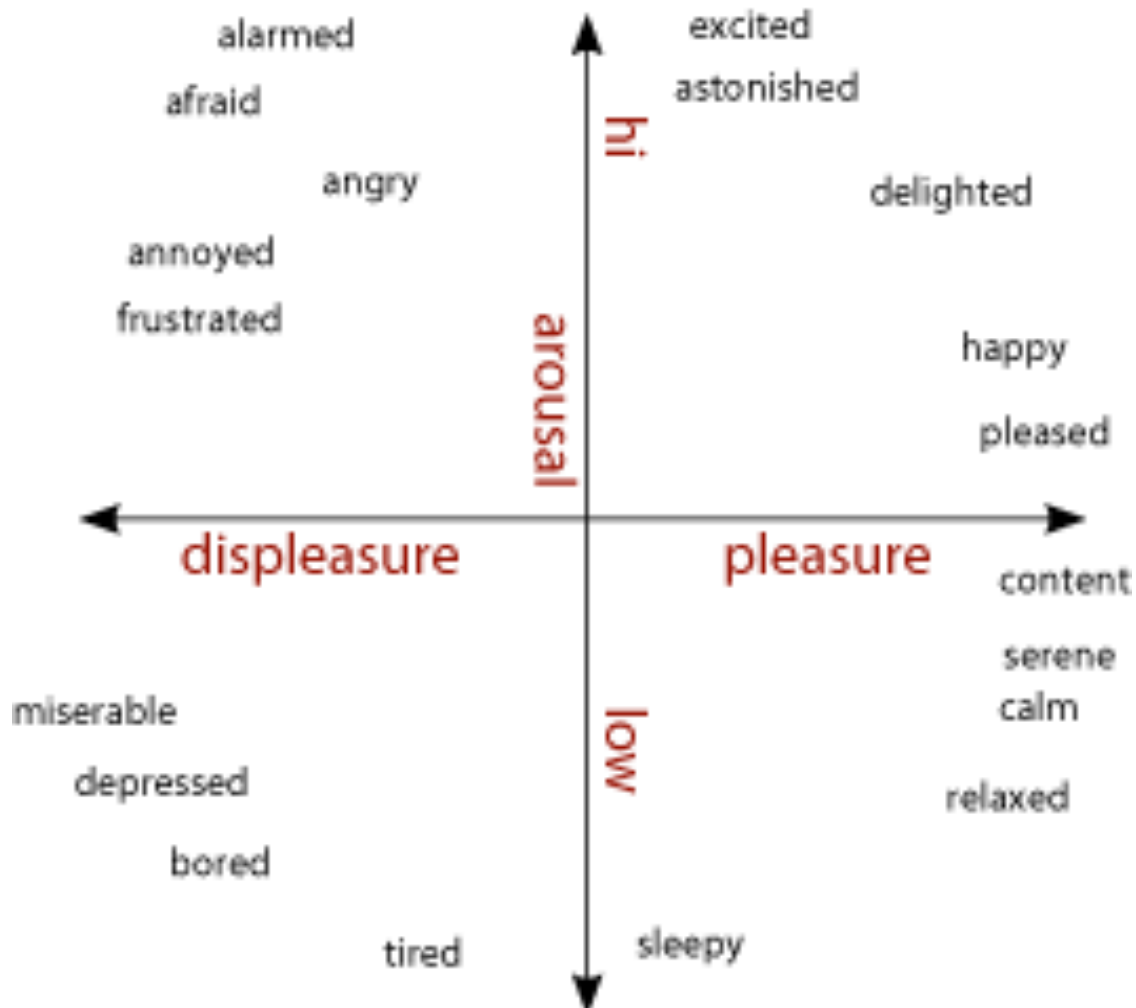
Shape space according to Marr



Sharing semantic domains

- Why is easy to explain to a four-year-old the meaning of the colour terms "chartreuse" and "mauve"
- ... but difficult to explain the monetary terms "inflation" and "mortgage"?
- Language understanding depends on *sharing* semantic domains

Emotional domain




Visuo-spatial domain

- Egocentric space
- Allocentric space

Action space

- Action identification depend on the forces that are exerted
- Represented as patterns of force vectors




Object category domain

- Domains for properties of objects
- Size, shape, weight, colour, taste ...
- Properties are convex regions of domains
- Categories are sets of properties
(+ correlations)

Goal space

- Locations in visuo-spatial domain transformed into goal space
- Extended by metaphorical mappings to more abstract goal spaces



Geometry of Meaning

- (1) What is semantics?
- (2) Conceptual spaces and semantic domains
- (3) Semantics as a meeting of minds
- (4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions
- (5) A cognitive theory of actions and events
- (6) Compositionality of meanings

Semantics as a meeting of minds

- Two kinds of processes

(1) Fast:

- Joint attention
- Common ground

(2) Slow:

- Coordinating meanings of words

Clark's common ground

- The participants in a conversation work together against a background of shared information
- As the discourse proceeds, the participants accumulate shared information by adding to it with each utterance
- Speakers design their utterances so that their addressees can readily identify what is to be added to that common ground


Why convexity?

- Makes learning more efficient
- Connects to prototype theory
- Strong empirical support for colour categories (Jäger)
- Makes it possible for minds to meet via communication

Colour naming game

- Communication game studied by Jäger and van Rooij
- Signaller and receiver have a common space for colours
- Signaller can choose between n messages


A language game with a few messages



Colour naming game

- Communication game studied by Jäger and van Rooij
- Signaller and receiver have a common space for colours
- Signaller can choose between n messages
- Signaller and receiver are rewarded for maximizing the similarity of the colours represented
- There exists a Nash equilibrium of the game that is a Voronoi tessellation (all regions convex)

Fixpoints in color mapping




The semantic reaction function

- *Expression function*: from communicative acts, e.g. utterings of words, onto mental states (represented by conceptual spaces).
- *Interpretation function*: from mental states to communicative acts

Language preserving neighbourhoods

This space
is discrete,
but combinatorial



Fixpoint semantics

- "Meeting of minds" \approx reaching agreement on a contract
- A semantics is a function that maps communicative expressions on conceptual spaces, and vice versa
- Minds meet when the representation-interpretation function mapping states of mind (represented by conceptual spaces) on states of mind via gestures or language finds a resting point – a *fixpoint*
- Related to equilibria in communication games
- Topological and geometric properties of conceptual spaces help generating fixpoints in communication activities
- Same mechanisms in fast and slow processes

The mathematical model

- States of mind of agents are points x in the product space of their individual mental representations C_i
- Similarity provides a metric structure to each C_i
- Additional assumptions about C_i : *convexity* and *compactness*
- If C_i are compact and convex, so is $C = \prod C_i$
- An representation-interpretation function $f: C \rightarrow C$
- It is assumed that f is *continuous*
- “Close enough” is “similar enough”. Hence continuity of f means that language can preserve similarity relations!

The central fixpoint result

- Given a map $f: C \rightarrow C$, a *fixpoint* is a point $x^* \in C$ such that $f(x^*) = x^*$
- *Theorem* (Brouwer 1910): Every continuous map of a convex compact set on itself has at least one fixpoint (in practice many)
- In the Jäger & van Rooij game, the set of prototypes define the fixpoints
- Semantic interpretation: If individual meaning representations are “well-shaped” and language is plastic enough to preserve the spatial structure of concepts, there will be at least one equilibrium point representing a “meeting of minds”

Meanings are in the *heads*

Meanings emerge through the interaction between the members of a linguistic community

Language is a game with speech act moves where we try to coordinate (or negotiate) our meanings

A semantics for a language is ideally a fixpoint for the game

Fixpoints exist because our conceptual spaces have the right structure

Reality enters via the payoffs of communication. If meaning is not aligned with reality, then the communicators will suffer costs

Geometry of Meaning

- (1) What is semantics?
- (2) Conceptual spaces and semantic domains
- (3) Semantics as a meeting of minds
- (4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions
- (5) A cognitive theory of actions and events
- (6) Compositionality of meanings

Why are there word classes?

- Nouns, pronouns, adjectives, verbs, adverbs, prepositions, conjunctions and interjections
- No universal classification: Languages carve up the words in different ways
- In traditional linguistics, word classes are defined syntactically
- My aim is to give a cognitive/communicative motivation for the classes

Object categories

- Basis for *nouns*
- *Concepts* with information from several domains
- Which domains are included a category?
- Definitional generics: "A dog has four legs", "A cup has a colour", "*An idea has a colour"
- Domains are more or less salient (essential) for a category



Generic sentences

- *Definitional (4th level)*
- Madrigals are polyphonic
- A madrigal is polyphonic
- *Descriptive (3rd level)*
- Madrigals are popular
- *A madrigal is popular

Object category

- (i) a set of relevant domains (may be expanded over time)
- (ii) a set of convex regions in these domains (in some cases the region may be the entire domain)
- (iii) salience weights of the domains (roe vs. caviar)
- (iv) information about how the regions in different domains are correlated
- (v) information about part-whole relations

Change of weight of a dimension



Objects as categories

- Objects are represented as points in a conceptual space – a degenerate region
- The properties of an object are always consistent (no need for meaning postulates)
- Many points represent merely *possible* objects – a simple model of fictionality
- In addition to criteria for categories, a physical object is assumed to be spatiotemporally continuous (cf. object permanence)

Basic semantics of nouns and names

- The basic *communicative function* of a noun or a name is to identify a referent
- Nouns refer to concepts
- Names refer to objects
- Nouns do not always identify unique referents. Why does not everything have a name?
- Nouns are *cognitively economical*

Properties and adjectives

- Adjectives has two communicative functions:
 - (i) Help in identifying a referent: "Red book"
 - (ii) Informative (in predicative position). "The stove is hot"
- *Single domain hypothesis for adjectives*: The meaning of an adjective can be described as a convex region in a single domain
- No adjective means "green or orange"
- No adjective means "long and hot"

Evidence for the convexity criterion

- Jäger (2010) investigated data on color terminology from 110 languages
- 93.8% of the data could be explained by convex partitioning of the (incomplete) Munsell color chip space



Potential counterexamples

(i) *Ripe* strawberry = *red* and *sweet* strawberry

- Strong correlation generates dimension in product space

(ii) *Healthy* = not having pain, not having a high temperature, not having a high blood pressure, not having a high cholesterol level, etc

- *Healthy* is a region in a multidimensional "illness" domain

Kinship in a product space of generations

- Example of product space (from Joost Zwarts)



| | 0 | 1 | 2 | 3 |
|---|---------------------|------------|-------------|-------------------|
| 0 | I | mother | grandmother | great-grandmother |
| 1 | daughter | sister | aunt | (grand)aunt |
| 2 | granddaughter | niece | cousin | cousin |
| 3 | great-granddaughter | grandniece | cousin | cousin |

The development of domains

Linda Smith (1989): There is a *dimensionalization* of the knowledge system. ... Among the first words acquired by children are the names for basic categories – categories such as *dog* and *chair*, which seem well organized by overall similarities. Words that refer to superordinate categories (e.g., *animal*) are not well organized by overall similarity, and the words that refer to dimensional relations themselves (e.g., *red* or *tall*) appear to be understood relatively late

Establishments periods for domains

- *Establishment hypothesis*: If one word from a domain is learned during a certain establishment period, then other (common) words from the same domain should be learned during roughly the same period.



Geometry of Meaning


- (1) What is semantics?
- (2) Conceptual spaces and semantic domains
- (3) Semantics as a meeting of minds
- (4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions
- (5) A cognitive theory of actions and events
- (6) Compositionality of meanings


How is action space structured?

We know even less about the geometry and topology of action space than we know about shape space

Dynamic domains

Marr & Vaina
"Walk"





Gunnar Johansson's
patch-light technique
for analysing motion
perception

Humans are excellent at identifying actions



Kinematic specification of dynamics (Runesson)

The kinematics of a movement contains
sufficient information to identify the
underlying *dynamic force patterns*

Relations between action and meaning

- Verbs normally express actions or results of actions
- Actions are represented cognitively as *force patterns*
- Talmy's (1988) force dynamics is limited to interactions between two actors
- Embodied cognition focuses on the subject
- The role of forces in semantics has been underrated

Representational hypothesis

- The fundamental cognitive representation of an action is the *pattern of forces* that generates it
- Actions are more or less *similar* and show prototype effects
- An *action category* is a convex region in the space of force patterns

Similarities between actions

To identify the structure of the action space, *similarities* between actions should be investigated. This can be done with basically the same method as for similarities between objects.


Morphing actions to generate similarity space



Actor 1



Morph



Actor 2

Functional properties

Many objects are not categorized on the basis of their perceptual properties, but according to their functional properties

Paul Auster, *City of Glass*: "Not only is an umbrella a thing, it is a thing that performs a function – in other words, expresses the will of man. When you stop to think of it, every object is similar to the umbrella, in that it serves a function. A pencil is for writing, a shoe is for wearing, a car for driving.

Now my question is this. What happens when a thing no longer performs its functions? Is it still the thing, or has it become something else? When you rip the cloth off the umbrella, is the umbrella still an umbrella? You open the spokes, put them over your head walk out into the rain, and you get drenched. Is it possible to go on calling this object an umbrella? In general, people do. At the very limit, they will say the umbrella is broken. To me this is a serious error, the source all our troubles."

Main thesis

Functional properties correspond to convex regions in *action space*.

Example:

A chair affords
sitting actions
with support
for back for
one person




Events


- *Metaphysical* analyses (the ontology of events)
- *Cognitive* models of events that account for how events are mentally *represented*
- *Linguistic* expressions describing events or parts of events
- Cognitive and linguistic aspects are often confused
- E.g. symbolic representations: “break” ≈
[[X ACT_{<MANNER>}] CAUSE [Become [Y <*BROKEN*>]]]

A two-vector model of an event

- The *force vector* (pattern) acts on an patient
- From force space (generated by an action)
- The *result vector* describes the changes of the properties of the patient
- Changes in location or in category space
- Agent is not always necessary (fall, grow ...)




Example: pulling a sled



The force vector is the *cause*, the result vector the *effect*

More components of events (thematic roles)

- *Agent* represented in agent space that contains at least the force domain
- *Patient* represented in category space and physical space
- *Instruments*
- *Counterforces* exerted by the Patient
- *Intentions* of the Agent



Sentences refer to (parts of) events

Thesis: A sentence expresses a part of an event involving at least a force or a result vector and one entity

In analogy with the visual process, a sentence *focuses* on some parts of an event (“construal”)

“Victoria hits Oscar” vs “Oscar was hit by Victoria”

No simple relation between Agent/Patient and Subject/Object

Representing verb meanings


- *Thesis (single domain constraint)*: The meaning of a verb is a convex region of vectors that depends only on *a single domain*.
- Examples: *push* refers to the force vector of an event (and thus the force domain), *move* refers to changes in the spatial domain of the result vector and *heat* refers to changes in temperature.
- There are no verbs that mean ‘walk and burn’ (multiple domains) and there are no verbs that mean ‘crawl or run’ (not convex)

Predictions from the theory

- Explains *similarities* of verb meanings
- Explains *subcategories* of verbs
- *March, stride, strut, saunter, tread* etc are subregions of *walk*
- Explains metaphorical uses of verbs: "scything down a football player", "slicing a backhand"
- Predicts the division into manner and result verbs
- Manner verbs describe force vector (cause) and result verbs the result vector (effect)

Manner verb: "Push"

- Force applied to object
- Prototypically, push leads to change in position of object
- However, this change is not certain, due to counterforces
- Expectations can be tested with "but"



Result verbs

- Describes change in object:
"open", "grow", "paint"
- Do not determine the forces
that lead to the change
- Two basic kinds: change of
position ("move") and change
of properties ("shrink")



Perception and emotion verbs

- I hear the owl (result verb, "I" is patient)
- I listen to the owl (manner verb, "I" is agent)
- I persuade you, I scare you, I praise you, I blame you
- Apply to different domains of patient's emotional or cognitive space

Why either manner or result?

- Strong support from linguistic analyses
- Connection between force vector and result vector not direct (counterforces etc)
- Makes it difficult to *learn* the mapping



Possible counterexample: "climb"

- Seems to involve both upward motion (result) and clambering (manner)
- Can be analyzed as manner where the force vector has an upward direction

(1) Oscar climbed the mountain.

(2) Oscar climbed down the mountain.

(3) Oscar climbed along the rope.

(4) The train climbed the mountain.

(5) ?The train climbed down the mountain.

(6) The snail climbed up the side of the tank.

State verbs

- E.g. be, sleep, hate
- No change involved (identity vector in property space)
- No force applied
- Result verbs

Intentional verbs

- Agents can exert forces intentionally
- Intentional = desired change in *goal space*
- *Kill* vs *murder*
- *Wink* = *blink* + intention to “awaken the attention of or convey private intimation to person” (*The Oxford Concise Dictionary*)
- Summarize *two* events: The fulfilment of a goal (change in goal space) and the physical action (change in physical property)

Adverbs (modifying verb vectors)



- Analogous to how adjectives modify nouns
- "Walk backward" (*direction* of vector)
- "Push strongly" (*magnitude* of vector)
- Apply to one domain (?)

Prepositions

- The core meaning of a preposition depends only on a *single domain*
- Spatial domain (*above*), force domain (*on*), time domain (*during*)
- Complicating factor: Prepositions have many metaphorical uses



The geometry of prepositions

- Space described by *polar* coordinates
- Results in a different kind of convexity





Locational prepositions


Inside and outside





Near and far



Front, back, left of, right of



Above, below and between




Directional prepositions

- Goal prepositions: *to, into, onto, towards*
- Source prepositions: *from, out of, off, away from*
- Route prepositions: *through, over, along, around, across*
- Analysed in terms of *paths*
- Convexity of paths can be defined

Some prepositions depend on forces

- *In, on, against, over, under ...*




Some prepositions depend on forces

- *In, on, against, over, under ...*



Some prepositions depend on forces

- *In, on, against, over, under ...*



Demonstratives

- *This, that, here, and there*
- Form a word class on their own
- *This* and *that* refer category space, while *here* and *there* refer to physical space
- Communicative role: To create joint attention
- First in combination with pointing
- Then in relation to common ground
- Pointing to the inner world

Quantifiers

- Flagship of truth-functional semantics
- Modelled in Montague grammar
- Here more similar to Lorenzen/Hintikka's game-theoretic semantics
- Langacker (2003): "Any cat" and "each cat" involve *fictive objects* (points in conceptual spaces)
- Two basic kinds: *proportional* (all, most, some) and *representative instance* (every, each, any)

Quantifiers

- Proportional quantifiers: multiplex-mass interchange
- *All, some* and *most* denote proportions
- Representative instance quantifiers: construed as *fictive individuals*
- *Any*: Arbitrary selection of one of the members (e.g. the individual called Anycat)
- *Each*: Sequential examination, one at a time
- *Every*: Simultaneous examination
- *A*: Every/each/any cat is lazy. *B*: *Which one?

Cognitive grounding of linguistic categories

- Object concepts *Nouns*
- Properties *Adjectives*
- Domain relations *Prepositions*
- Force & result vectors *Verbs*
- Modifying vectors (?) *Adverbs*
- Present entities *Demonstratives*
- Events *Sentences*

Cognitive grounding of linguistic categories

- Object concepts *Nouns*
 - Properties *Adjectives*
 - Domain relations *Prepositions*
 - Force & result vectors *Verbs*
 - Modifying vectors (?) *Adverbs*
 - Present entities *Demonstratives*
 - Events *Sentences*
-
- Single domain reference + convexity!

Geometry of Meaning

- (1) What is semantics?
- (2) Conceptual spaces and semantic domains
- (3) Semantics as a meeting of minds
- (4) Cognitive grounding of word classes: nouns, adjectives, verbs and prepositions
- (5) A cognitive theory of actions and events
- (6) Compositionality of meanings

Compositions of meanings

Three kinds:

(1) Direct


(2) Modifier-head

(3) Metaphorical

- Traditionally compositions are defined for words
- Here, domains and functions are composed

Direct composition

- "Blue rectangle" created as a region in product space of colour space and shape space (category space)
- Structural properties of conceptual spaces are preserved under such composition
- Components can be recovered by projections on dimensions




Modifier-head composition

- Many combinations cannot be handled as direct products:
- "Tall squirrel"
- "Honey bee"
- "Criminal lawyer"
- "Leberkäse"

Modifier-head composition

Concepts are sensitive to context

Hot bath water is not a subcategory
of "hot water"




The effect of contrast classes

red: of the colour of fresh blood, rubies, human lips, the tongue, maple leaves in the autumn, post-office pillar boxes in Gt. Brit.


Advanced Learner's Dictionary of Current English.

- Red book
- Red wine
- Red hair
- Red skin
- Red snapper
- Redwood

The embedded skin color space




Radial projections



“Stone lion”


The only domains that are common to “stone” and “lion” are *shape*, *material* and *color*. The values for the shape of “lion” and the value for the material and the color of “stone” determine the meaning of “stone lion”.

The mechanism of metaphor



(a)

A bumpy road



(b)

A bumpy relationship

Taking stock

- A semantic theory that satisfies the six criteria
- An analysis of domains based on conceptual spaces
- The role of semantic domains in language learning
- A model of meaning created by the meeting of minds
- A model of actions and event based on conceptual spaces
- A semantics of verbs
- The single domain constraint for adjectives, verbs and prepositions (and maybe other word classes)
- Cognitive/communicative grounding of word classes
- A geometric theory of compositionality

Computational issues

- Symbolic methods involve rule following based on recursive procedures on tree-like structures
- Conceptual spaces are based on vectorial representations. Calculations involve matrix multiplications (coordinate transformations), inner products, Voronoi tessellations, etc

The Semantic Web is not semantic

- At best it is ontological
- Taxonomy + inference rules
- Berners-Lee, Hendler and Lassila:
“Fortunately, a large majority of the information we want to express is along the lines of ‘a hex-head bolt is a type of machine bolt’.”
- Unfortunately, this is not true

The Semantic Web is not semantic

- OWL (and other web languages) "adds more vocabulary for describing properties and classes, among other relations between classes, cardinality, equality, richer typing of properties, characteristics of properties ... " (McGuinness and Hamelen2004)
- This is exactly what is to be expected of a language with the expressivity of first order logic where concepts (properties and relations) are defined in terms of *sets* of objects
- The symbols of a web ontology language are not *grounded* outside the system

Conceptual spaces as a tool for the Semantic Web

- *Conceptual hierarchies* (taxonomies) will emerge from the inclusion structure of the concept regions of a domain (e.g. via a Voronoi tessellation)
- *Identities of concepts* corresponds to identities of regions
- *Identities of names* given by identities of points (more complicated when vectors are incomplete)
- *Property characteristics* (e.g. transitivity and symmetry) are implicit in the topology or geometry of the domains
- Explicit inferences based on additional rules become superfluous

CSML

Conceptual Space Markup Language

- Adams and Raubal 2009

```
<csml:ContrastClass csml:ID="#Tall"  
  csml:domainID="#Size">  
  <csml:aVector>  
    <cn>-1.0</cn>  
    <cn>0.0</cn>  
  </csml:aVector>  
  <csml:qVector>  
    <ci>#Height</ci>  
    <ci>#Width</ci>  
  </csml:qVector>  
  <csml:ccMin> -0.7 </csml:ccMin>  
  <csml:ccMax> 1.0 </csml:ccMax>  
</csml:ContrastClass>
```

The role of attention in language production

(Ongoing work with Peter Ford Dominey in the WYSIWYD project)



The role of attention in language production

- Suppose you want to express that the iCub gives the ball to Peter
- "Peter is given the ball by the iCub"
- "The ball is given to Peter by the iCub"
- "*It* is given by the iCub to *him*"
- "Peter receives the ball from the robot"
- Etc

The role of attention in language production

- General rule: The *subject* of the sentence expresses the focus of attention of the speaker
- "Peter admires the iCub" vs. "The iCub is admired by Peter"
- The choice of verb can focus on *action* or *result*
- "The iCub pushed the table" vs. "The iCub moved the table"