Lund University
Cognitive Science

Peter Gärdenfors

Using conceptual spaces and event representations to construct semantics for artificial systems

GPT-3 – competent speaker?

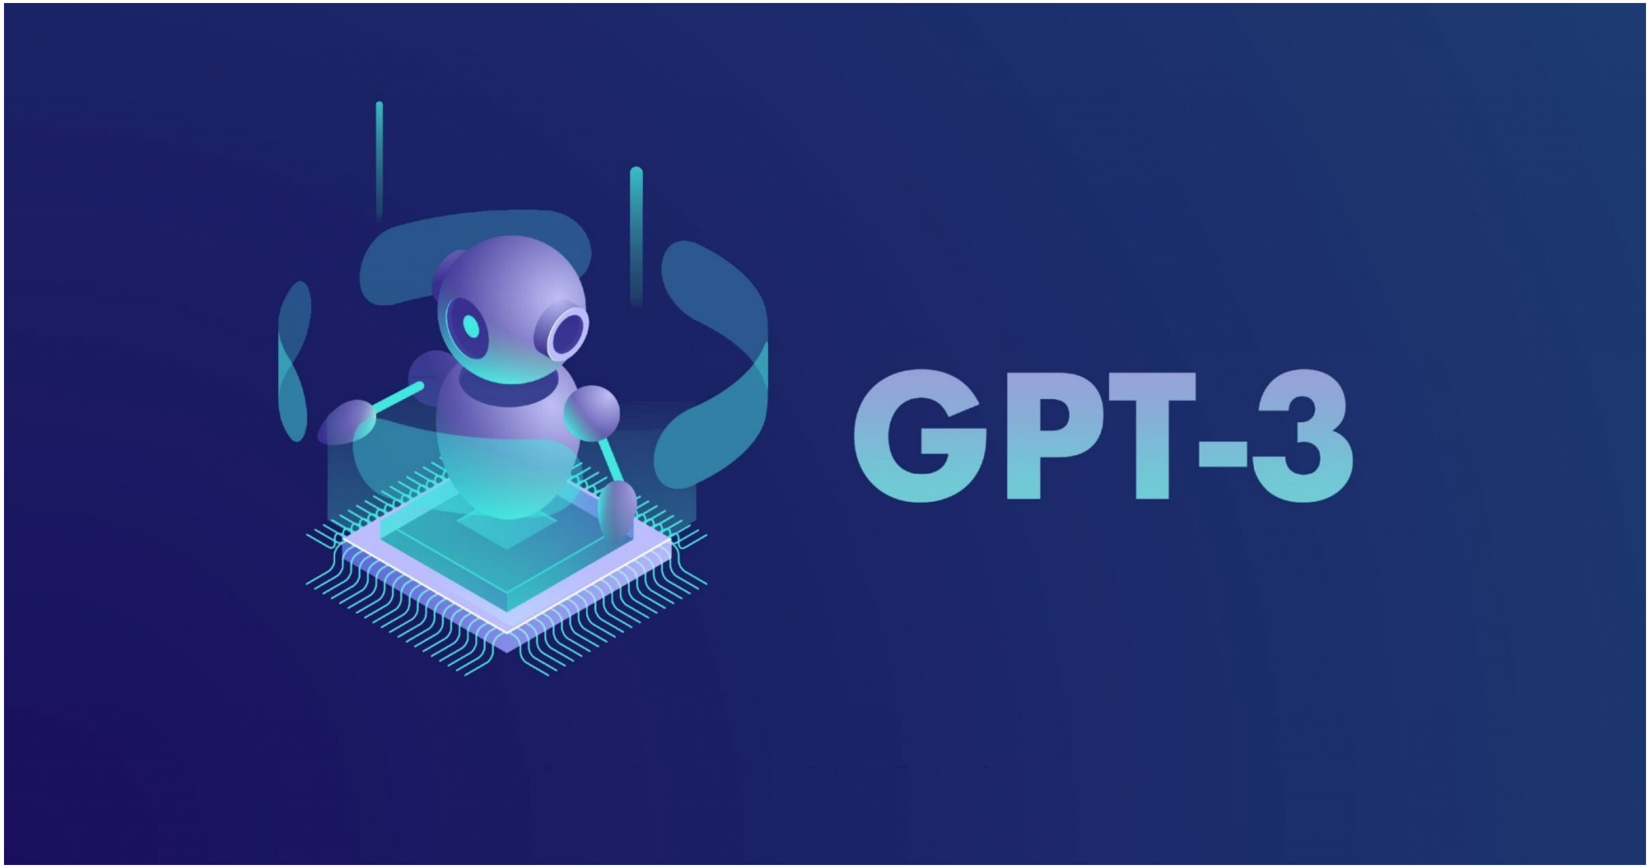

Main problem: GPT-3 is not *grounded* in the external world

Three levels of modelling in KR and ML

Symbolic models

Based on a given set of predicates with known denotation.

Representations based on logical and syntactic operations.

Problem: Where do the predicates come from?

Conceptual spaces

Based on a set of quality dimensions.

Representations based on topological and geometrical notions.

Problems: Where do the dimensions come from?

Connectionist models

Based on a (uninterpreted) inputs from receptors.

Distributed representations by dynamic connection weights.

Problems: What is represented in the network?

Learning is in general very slow.

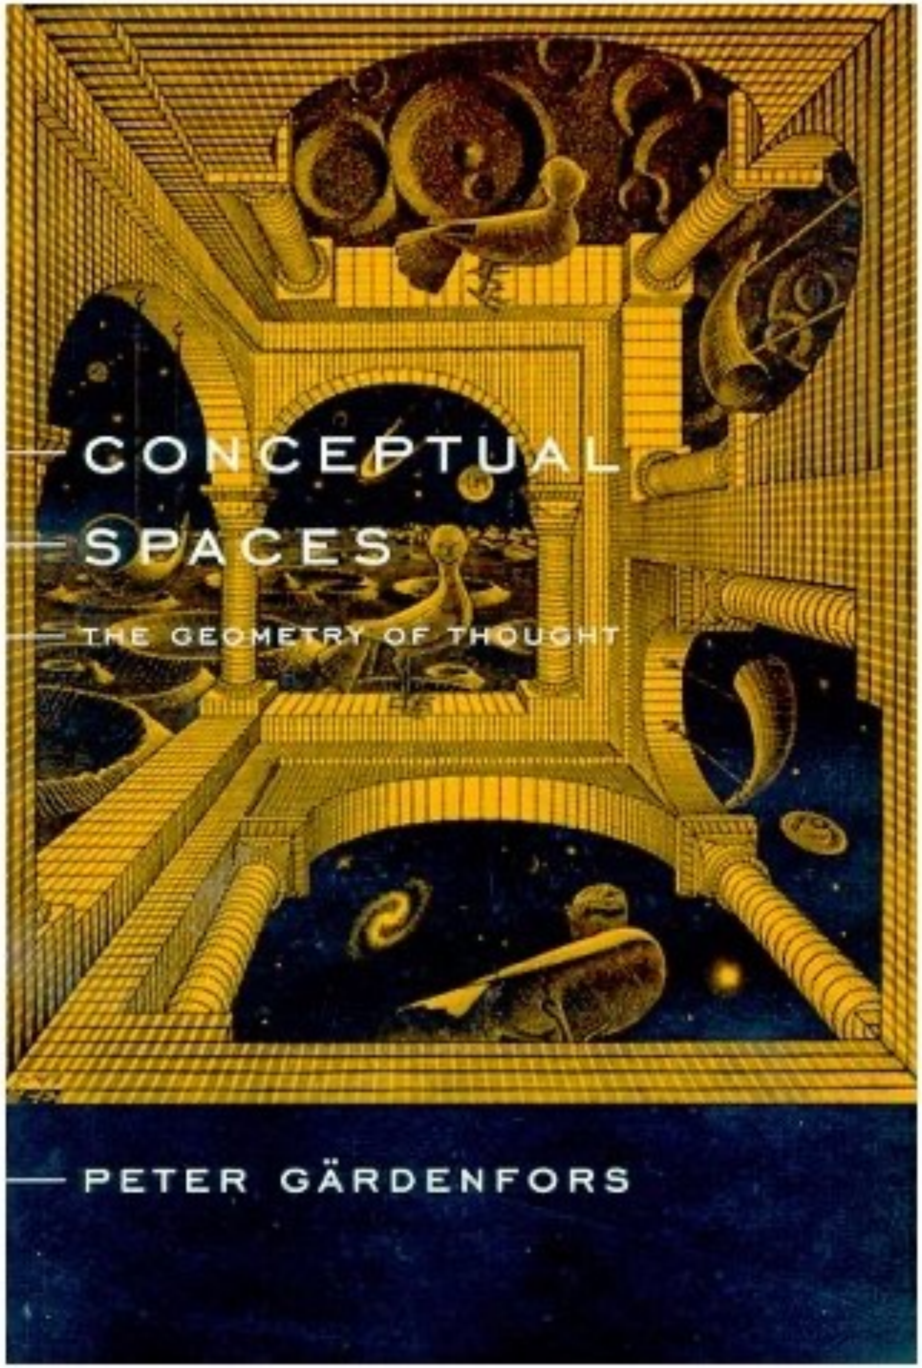
CONCEPTUAL
SPACES
THE GEOMETRY OF THOUGHT
PETER GÄRDENFORS

Conceptual spaces

- Information is organized by *quality dimensions*
- … that are sorted into *domains* (space, time, temperature, weight, color, shape …)
- Dimensions within domains are *integral*
- Domains are endowed with a *topology* or *metric*

Two linear quality dimensions

Time

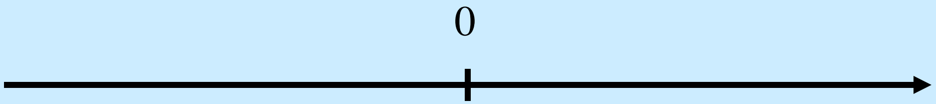

Weight

The color spindle

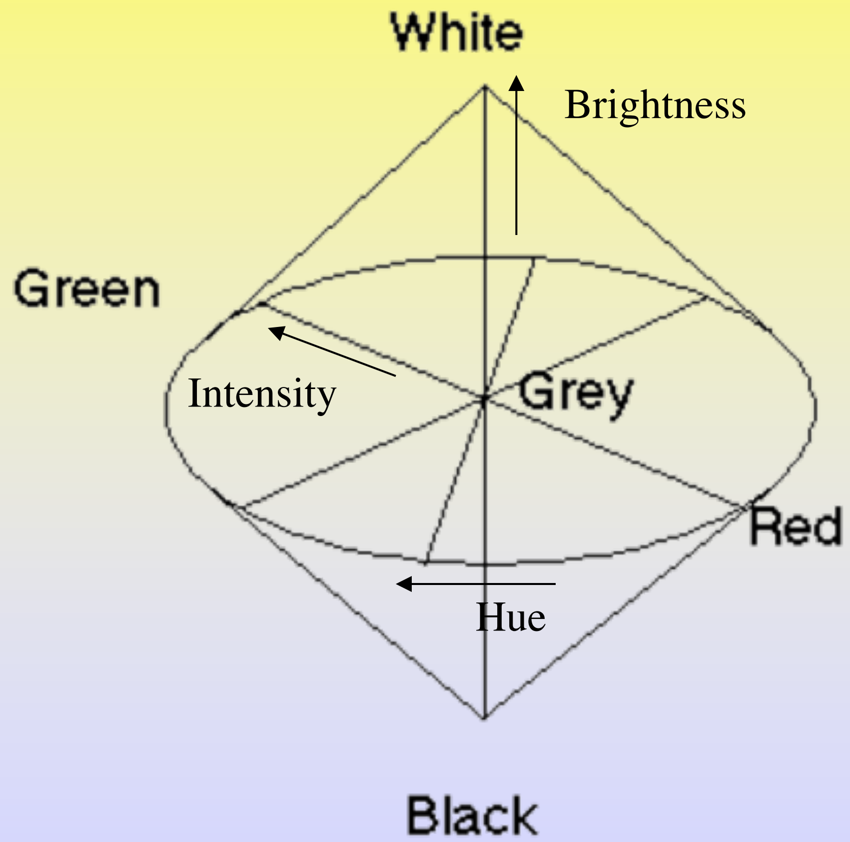

Conceptual spaces

- Information is organized by *quality dimensions*
- … that are sorted into *domains* (space, time, temperature, weight, color, shape …)
- Dimensions within domains are *integral*
- Domains are endowed with a *topology* or *metric*
- *Similarity* is represented by distance in a conceptual space

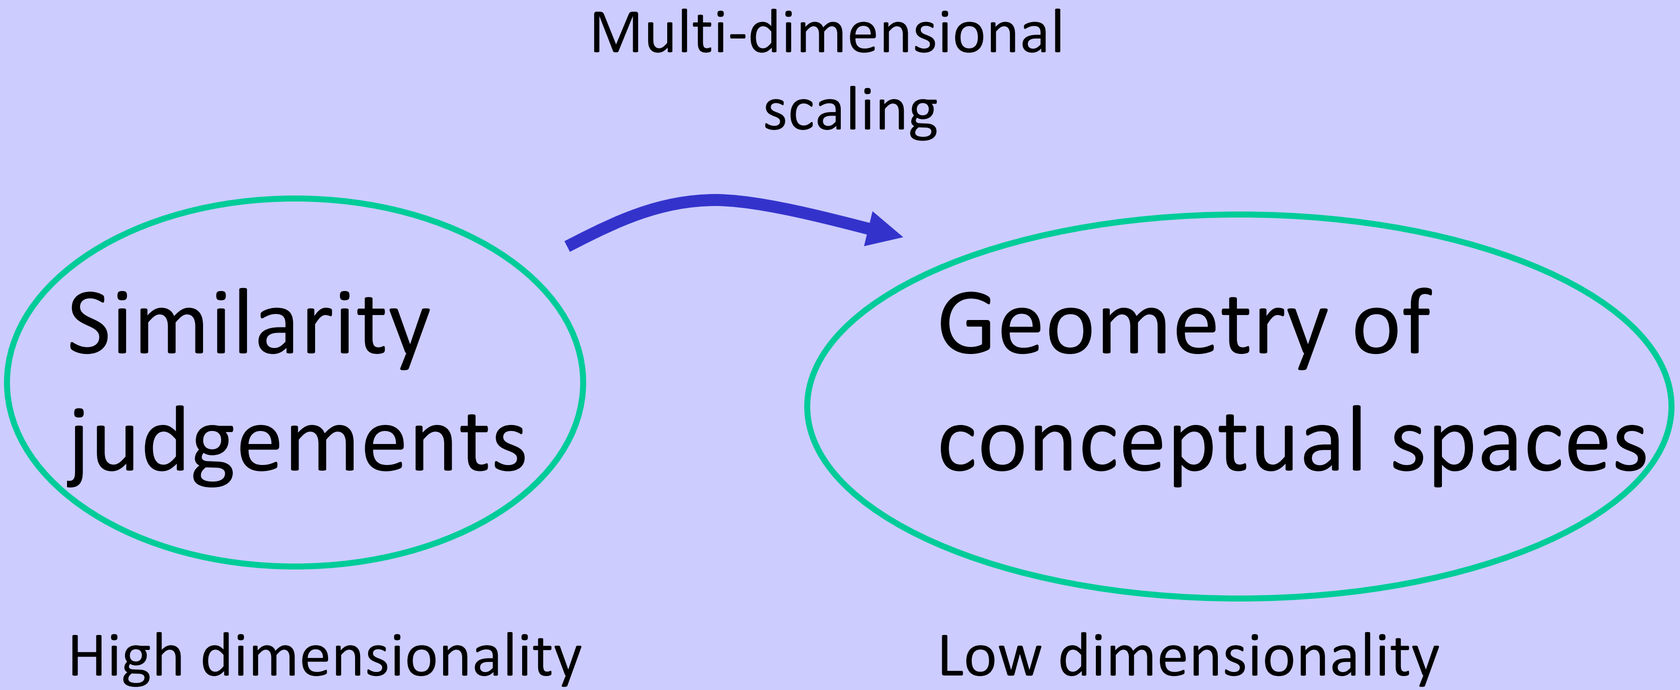
Multi-dimensional scaling
Similarity judgements
Geometry of conceptual spaces
High dimensionality
Low dimensionality

Concepts are *convex* regions in conceptual spaces

The color spindle

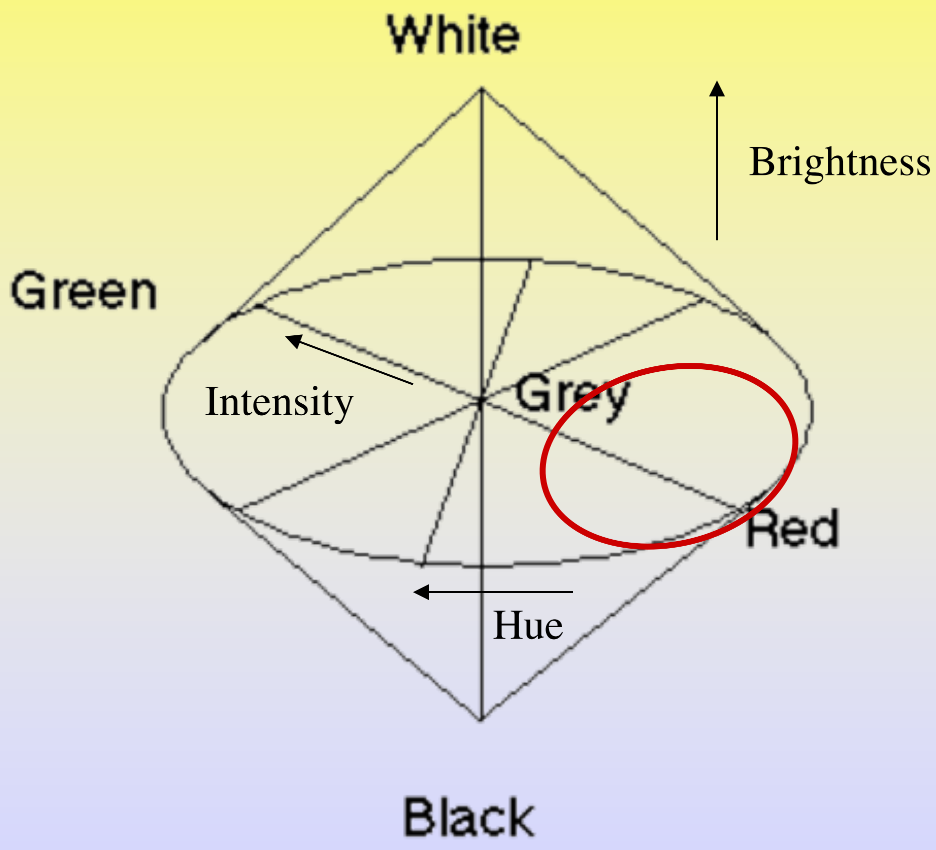

Why convexity?

- Handles concepts with vague borders
- Connects to prototype theory
- A robin is more typical bird than an ostrich
- An apple is more typical fruit than a kiwi
- Children learn prototypical meanings first

Categorization in conceptual spaces

Voronoi tessellations around prototype objects divides conceptual spaces into categories based on the nearest neighbour rule, i.e. each object is associated with the prototype closest to it.

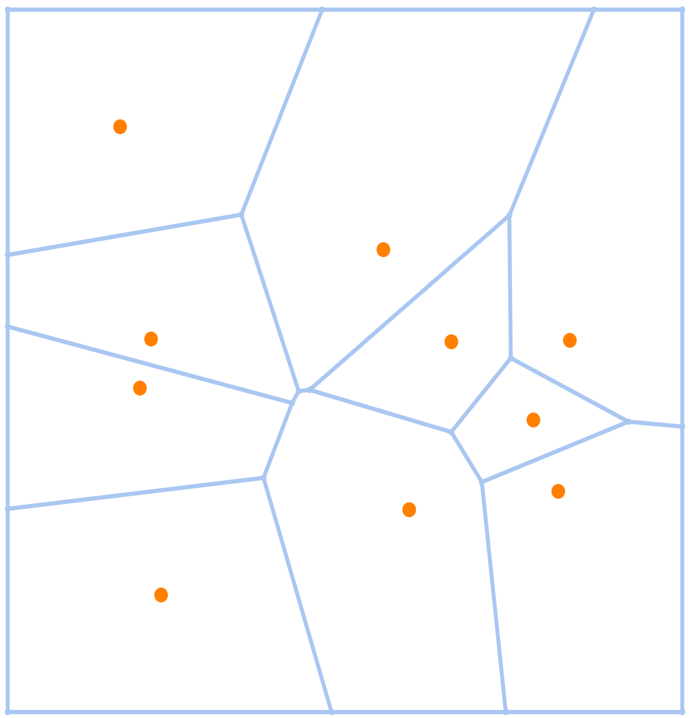

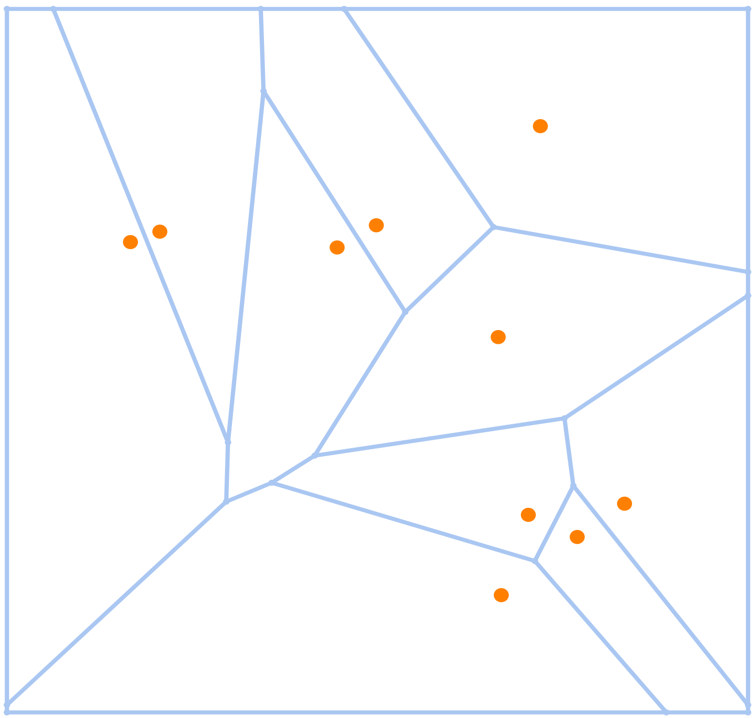

Why convexity?

- Handles concepts with vague borders
- Connects to prototype theory
- Makes learning more efficient

Learning from few examples

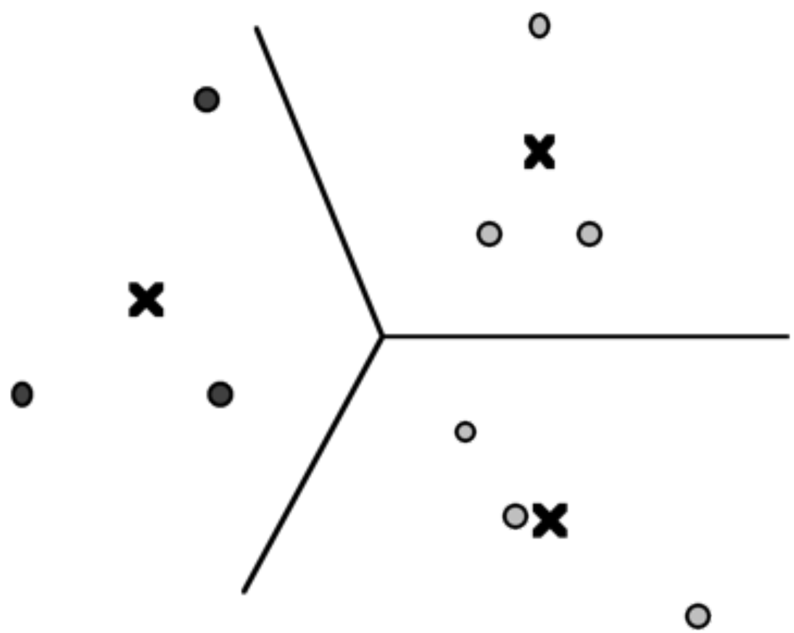

Learning from few examples

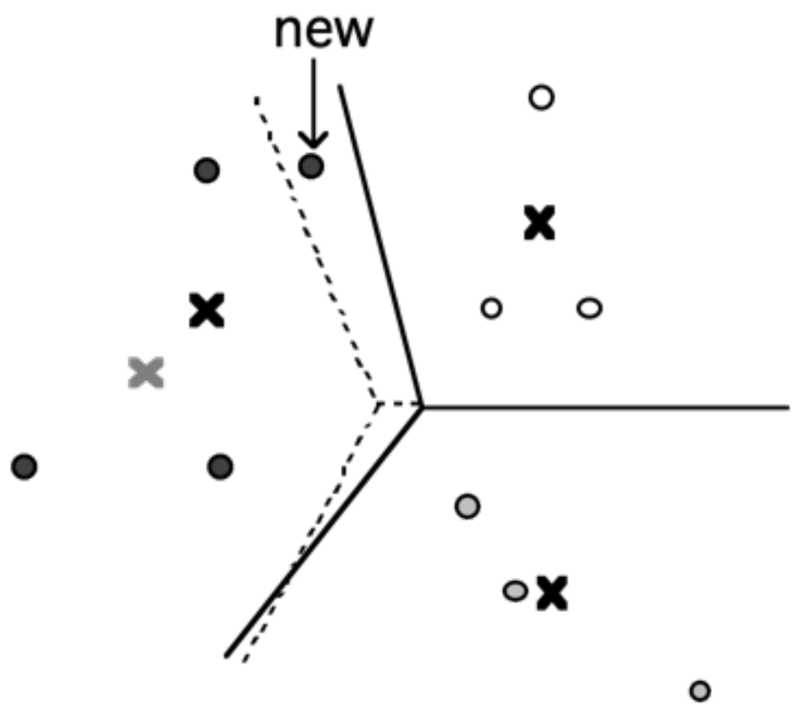

Properties vs. Object categories

Properties: A convex region in a single domain

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

Properties and adjectives

- *Thesis:* (single domain constraint for adjectives) The meaning of an adjective can be described as a convex region in a single domain
- Examples: *heavy* (weight), *hot* (temperature), *bitter* (taste), *round* (shape)
- No adjective means ”green or orange”
- No adjective means ”long and hot”

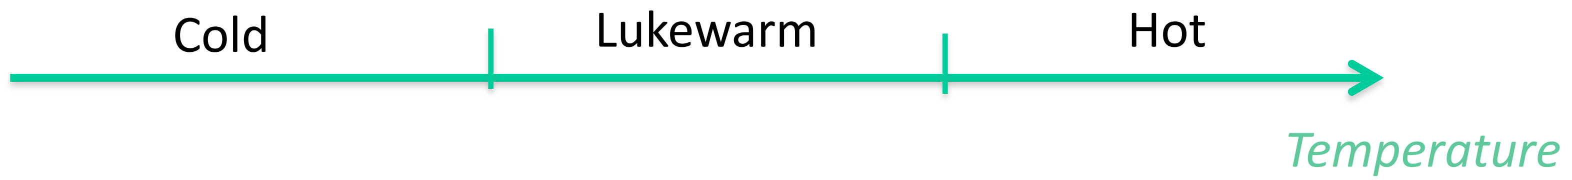

Properties vs. Object categories

Properties: A convex region in a single domain

Object categories: A number of convex regions in different domains; together with (1) salience values of the domains and (2) information about how the regions in different domains are correlated

Objects categories are expressed by *nouns*

An example of an object category: ”Apple”

Domain	*Region*
Fruit	**Values for skin, flesh and seed type**
Color	**Red-green-yellow**
Taste	**Values for sweetness, sourness etc**
Shape	**"Round" region of shape space**
Nutrition	**Values for sugar, vitamin C, fibres etc**

Cognitive grounding of word classes

- Object categories → *Nouns*
- Properties → *Adjectives*
- Dynamic properties → *Verbs*
- Spatial relations → *Prepositions*

Identifying actions

Representational hypothesis for actions

- 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 (Paul Hemeren's Ph.D. thesis)
- An *action concept* is a convex region in the space of force patterns
- Action/force ≈ shape/space

Representing verb meanings

- *Thesis:* (single domain constraint for verbs) The meaning of a verb root is a *convex* region of vectors that depends only on *a single domain*.
- Manner verbs for forces, result verbs for changes
- 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)

A two-vector model of an event

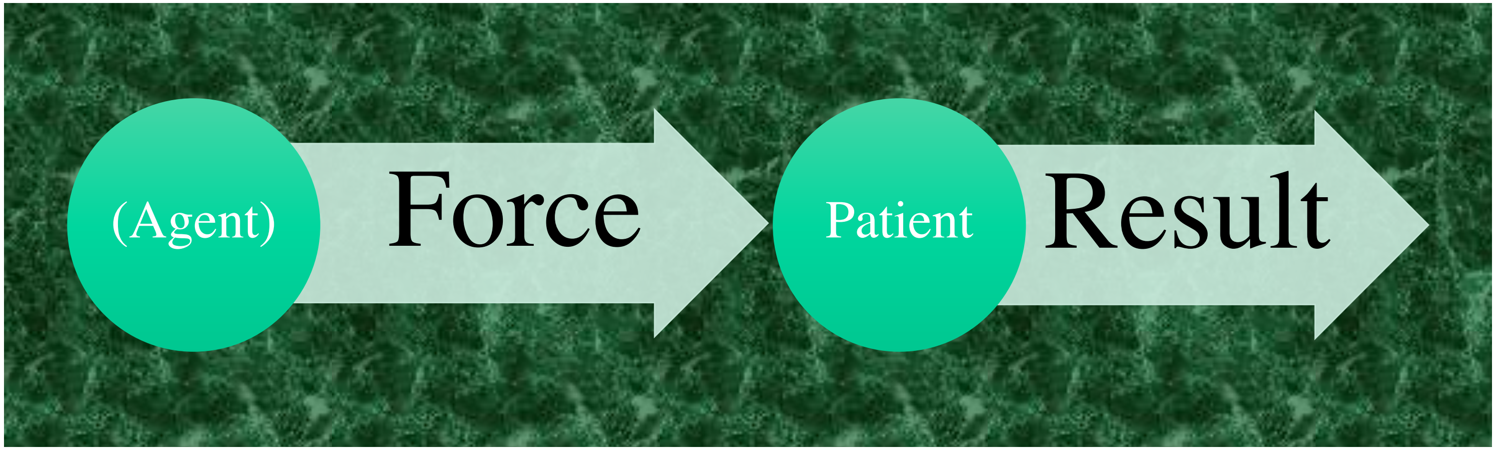

- The *force vector* (pattern) acts on an patient
- From force space (often generated by an action)
- The *result vector* describes the changes of the properties of the patient (can be null)

Sentences refer to (aspects of) events

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

Patients and agents ➔NPs / Vectors ➔VPs

In analogy with the visual process, a sentence *focuses* on some parts of an event

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

No simple relation between Agent/Patient and Subject/Object

Sentence comprehension

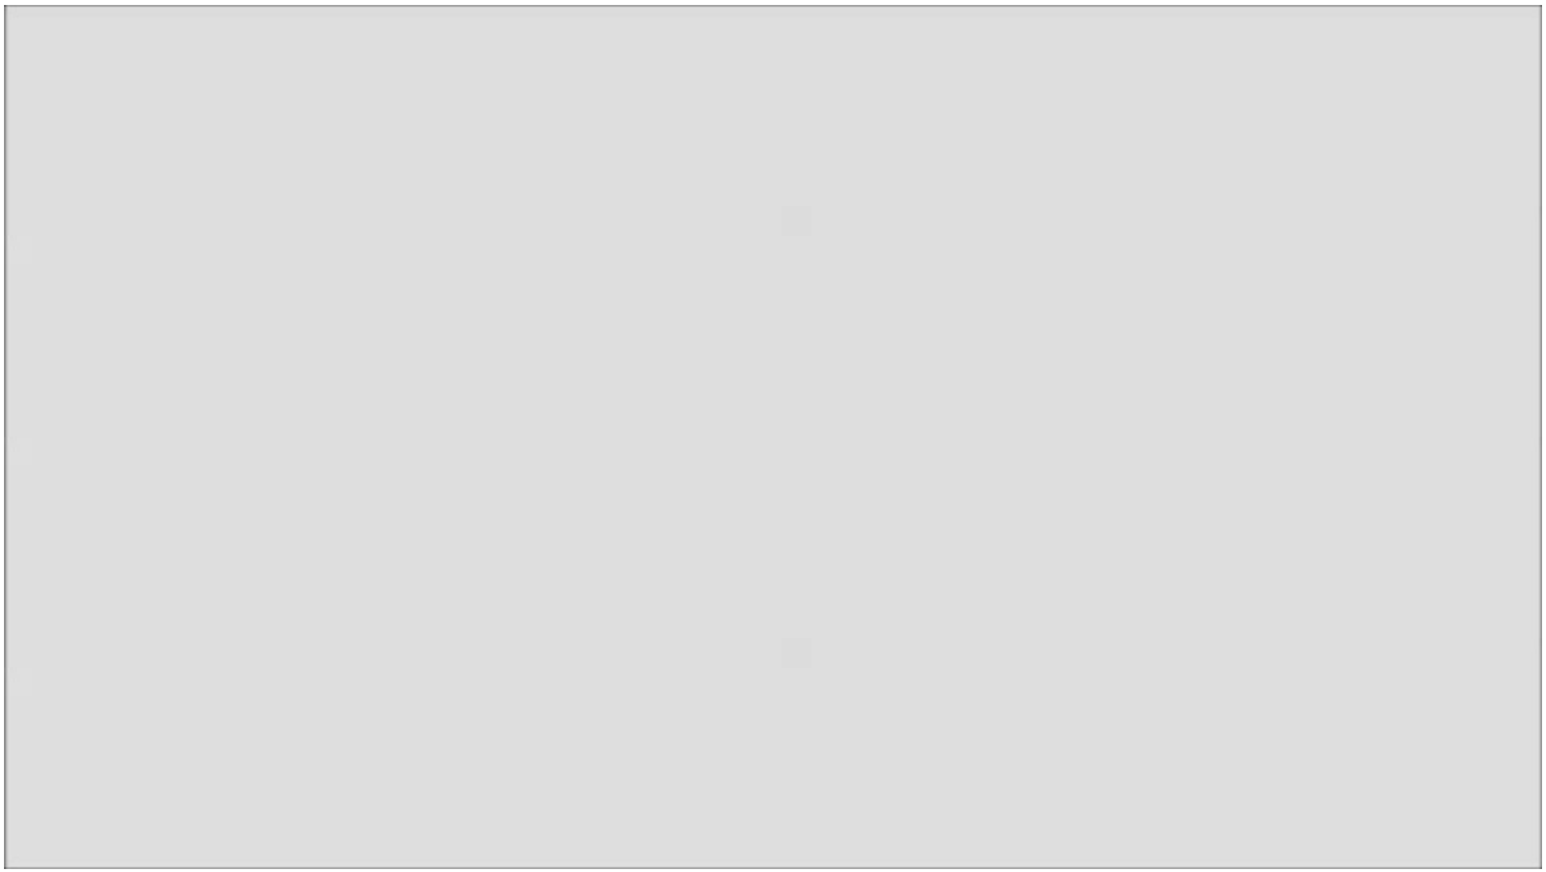

LUND action recognition system

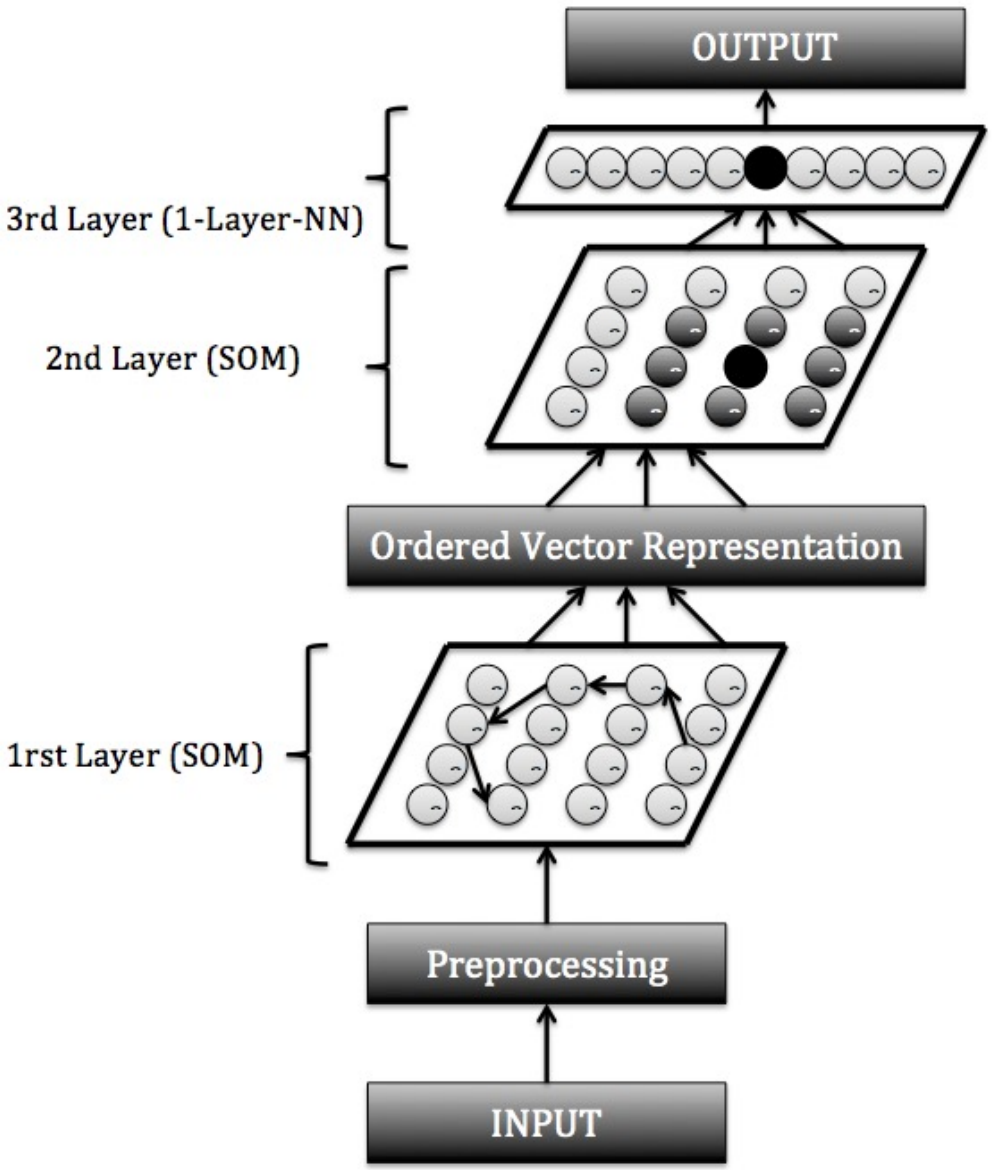

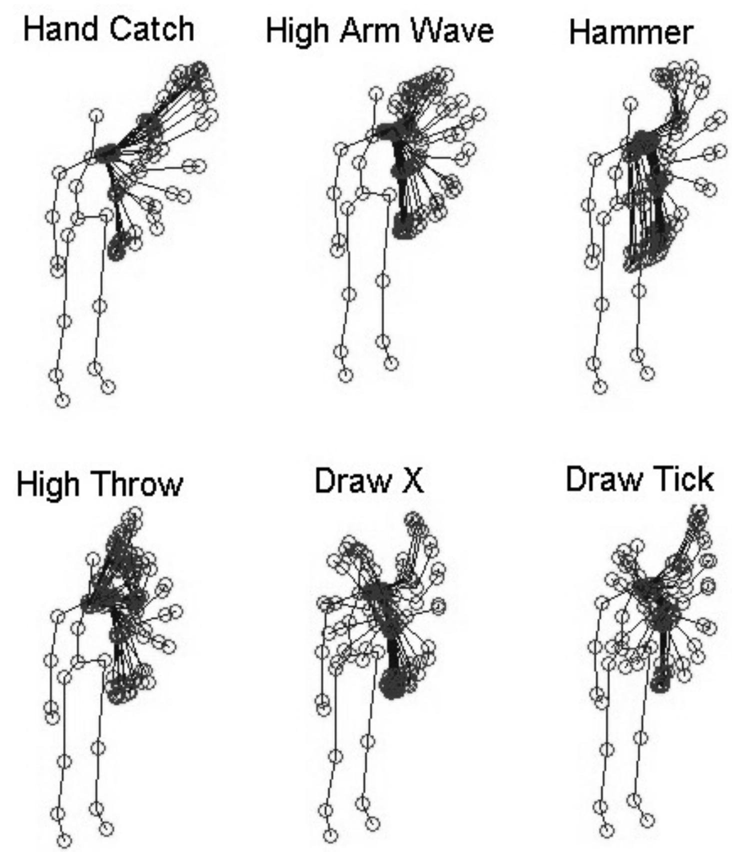

Real time action identification

From Zahra Gharaee's Ph.D. thesis

Semantic *structures* of word classes

• *Adjectives*	Regions of domains
• *Nouns*	Regions of products of domains
• *Verbs*	Vectors of domains
• *Prepositions*	Regions and vectors with landmarks
• *Adverbs*	Multipliers of vectors
• *Demonstratives*	Regions of perceptual domains
• *Quantifiers*	Properties of sets

Take home messages for conceptual engineers

- Organize information in separable domains
- Determine a distance measure
- Find prototypes
- Learn correlations between domains