Meaning Representation in Natural Language Categories

Trevor Fountain and Mirella Lapata
Institute for Communicating and Collaborative Systems, University of Edinburgh

THE UNIVERSITY *of* EDINBURGH
informatics

What is Categorization?

Categorization is the process by which people group stimuli into categories and use those categories to reason about new stimuli they encounter.

Categories are generally represented according to either the **exemplar** or **prototype** approach.

Stimuli are generally represented using **feature norms**.

If we represent stimuli using **corpus features**, can we get similar performance on categorization tasks?

Does it matter whether we represent categories using **exemplars** or **prototypes**?

How can we compare different category and stimuli representations?

Category Representation

Multiple theories exist as to how people represent categories:

Classical

- *is_edible*
- *contains_seeds*
- *grows_above_ground*
- *part_of_a_plant*

A list of required features which all instances of FRUIT must possess

Prototype

A single prototypical FRUIT from which all instances are generated

Exemplar

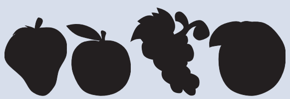

A number of stored instances of FRUIT to which new instances are likely to be similar

Meaning Representations

(a) Feature Norms

	has_4_legs	*used_for_eating*	*is_a_pet*
TABLE	12	9	0
DOG	14	0	15

(b) LSA

	Document 1	Document 2	Document 3
TABLE	0.02	0.98	-0.12
DOG	0.73	-0.02	0.01

(c) DV

	subj-of-walk	*subj-of-eat*	*obj-of-clean*
TABLE	0	3	28
DOG	36	48	19

(d) LDA

	Topic 1	Topic 2	Topic 3
TABLE	0.02	0.73	0.04
DOG	0.32	0.01	0.02

Semantic representations for TABLE and DOG using feature norms, Latent Semantic Analysis (LSA), Dependency Vectors (DV), and Latent Dirichlet Allocation (LDA).

Data

- Exemplars and feature norms were taken from McRae et. al (2005).
  - 541 exemplars in 41 categories
- Category labels and typicality ratings were collected via Amazon Mechanical Turk.
  - Mean reliability: 0.64, SD: 0.03

Models

$$sim_{w,c} = \sum_{x \in c} sim_{w,x}$$

(a) Exemplar

$$sim_{w,c} = sim_{w,c_{proto}}$$

(b) Prototype

In both models $sim_{w,x}$ is the cosine distance between stimuli.

Experiment

- Train four exemplar models (Nosofsky 1992) and four prototype models (Vanpaemel 2005), one per representation.
- For each model typicality $\approx sim(exemplar|category)$.
- Compare average correlation between model- and human-predicted typicality ratings.

Results

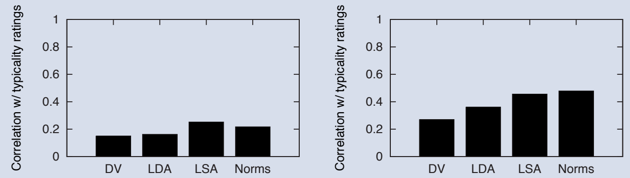

(a) Prototype (b) Exemplar

Average correlation for prototype (a) and exemplar (b) models between model- and human-predicted typicality rating, using various meaning representations.

Discussion

- Document co-occurrence (LSA) yields comparable performance vs. feature norms.
- Exemplar models overwhelmingly beat out prototype models.
- Low correlations even when using feature norms suggest that even humans have trouble with the task.
- Future work: explore specialized models for natural language categorization that are tailored to corpus-based meaning representations.

Bibliography

McRae, K., Cree, G. S., Seidenberg, M. S., and McNorgan, C. (2005). Semantic feature production norms for a large set of living and non-living things.

Nosofsky, R. M. (1992). Exemplars, prototypes, and similarity rules.

Vanpaemel, W., Storms, G., and Ons, B. (2005). A varying abstraction model for categorization.

sicsa*