How Well Do Distributional Models Capture Different Types of Semantic Knowledge?

Dana Rubinstein, Effi Levi, Roy Schwartz and Ari Rappoport
Institute of Computer Science, The Hebrew University
{drubin80, efle, roys02, arir} @ cs.huji.ac.il

Abstract: In recent years, distributional models (DMs) have shown great success in representing lexical semantics. In this work we show that the extent to which DMs represent semantic knowledge is highly dependent on the type of knowledge. We pose the task of predicting properties of concrete nouns in a supervised setting, and compare between learning taxonomic properties (e.g., *animacy*) and attributive properties (eg., *size, color*). We employ four state-of-the-art DMs as sources of feature representation for this task, and show that they all yield poor results when tested on attributive properties, achieving no more than an average F-score of 0.37 in the binary property prediction task, compared to 0.73 on taxonomic properties. Our results suggest that the distributional hypothesis may not be equally applicable to all types of semantic information.

Distributional models

are vector representations of words, derived from statistics of words and their contexts in a large corpus.

Impressive results for:
- ★ word similarity
- ★ analogies
- ★ categorization
- ★ and more...

mouse = (0.003, 0.011, 0.082, 0.102, ...)

Research Question:

How well do *distributional models* capture semantic knowledge about *concepts* and their *properties*?

Can we use Distributional Models to predict that bananas are yellow?

Concepts are prototypical concrete objects.

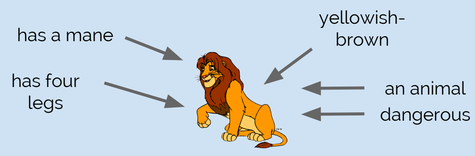

Properties are items of information that characterize the concepts.

The Task: Property Classification

We learn a classifier for each target property.
This property-classifier predicts for a given concept whether the property applies to it or not.

Each concept is represented by its corresponding vector in some Distributional Model.
We experiment with four state-of-the-art models:

- word2vec (Mikolov et al 2013)
- Distributional Memory (Baroni and Lenci 2010)
- Dependency word2vec (Levy and Goldberg 2014)
- GloVe (Pennington et al. 2014)

Formal Definitions

For each property p, we take concepts for which p applies to be positive instances, and concepts for which it does not as negative instances.

For example:
"is loud"

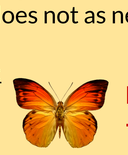

FALSE -1

TRUE +1

Notations:
$\mathcal{X}$ - the domain of concepts,
$\mathcal{Y}_p = \{1,-1\}$ - a binary label space.
$\phi: \mathcal{X} \rightarrow \mathbb{R}^n$ - a mapping from the concept domain to some Distributional Model space.
For each property p we learn a predictor

$$h_p: \phi(\mathcal{X}) \rightarrow \mathcal{Y}_p$$

★ Training was performed using linear SVM

The Data

We use the **McRae Feature Norms dataset** (McRae et al. 2005).
This data was collected in a set of experiments, where participants were presented with concepts and were asked to write down properties that describe them.

- ★ The data: 541 concepts and 2,526 different properties.
- ★ We focus only on 20 of the most frequent properties.
- ★ For each property, positive instances = concepts for which this property was elicited

Results

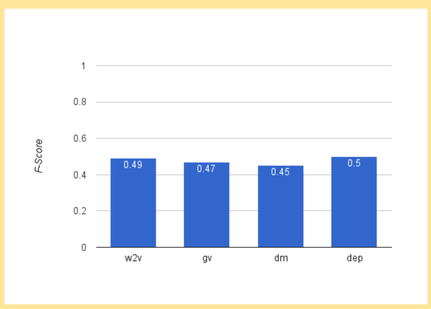

Average F-Scores over all properties for four different classifiers.

Each classifier was trained using a different state-of-the-art distributional model as feature representation for the data.

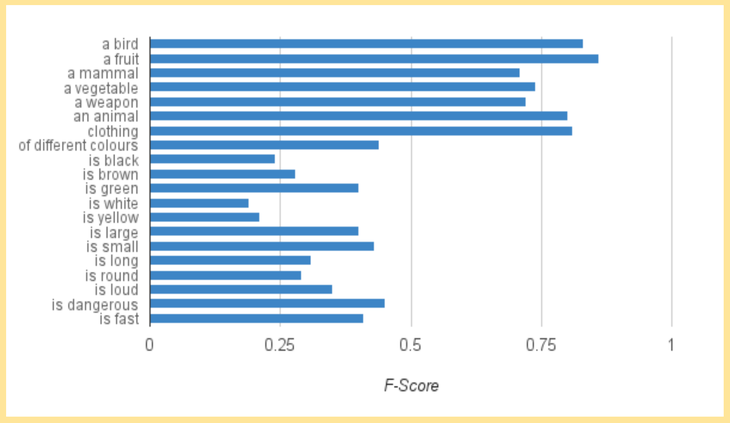

A deeper look into the performance of the **word2vec-based classifier.**

F-scores are presented for each property.

Note that several properties were learned much better than others.

Taxonomic Properties

- a bird
- a fruit
- a mammal
- a vegetable
- a weapon
- an animal
- clothing

w2v average F-score: 0.78

Attributive Properties

- is large
- is small
- is long
- is round
- is loud
- is dangerous
- is fast
- of different colors
- is black
- is brown
- is green
- is white
- is yellow

w2v average F-score: 0.34

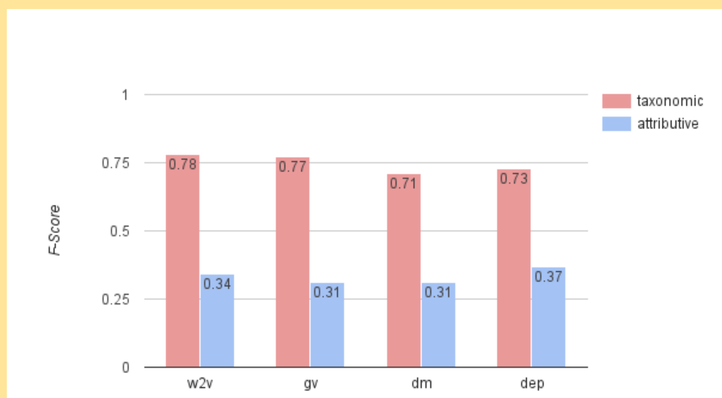

Average F-scores for taxonomic properties and attributive properties, shown for each of the classifiers.

All of the classifiers learn taxonomic properties better than attributive properties.

Results for a regression task show a similar trend (see the paper for details!)

Conclusion

In the context of learning semantic properties, state-of-the-art distributional models perform differently with respect to the type of property learned.

bananas are fruit

bananas are yellow

Future Work

- ★ Taxonomic properties as aggregates of attributive properties.
- ★ Fine grained property types (physically grounded vs. abstract)