Modelling metaphor with attribute-based semantics

Luana Bulat, Stephen Clark and Ekaterina Shutova
Computer Laboratory, University of Cambridge, United Kingdom

Conceptual metaphor theory

- Lakoff and Johnson (1980): metaphor can be explained through the presence of systematic associations between two concepts or domains.

TIME is MONEY

target domain *source domain*

- These associations allow us to project knowledge and inferences from the source domain onto the target domain, for example, we can reason about TIME in terms of the properties of MONEY.

You're wasting my time
This gadget will save you hours
How do you spend your time?

Research question

- We hypothesise that semantic models based on cognitively motivated properties will provide a better means of capturing and generalising metaphorical mechanisms.
- Can we identify metaphors more accurately by modelling semantics using (a proxy for) human conceptual knowledge?

Feature norms

- Also referred to as property norms or attributes.
- Represent human conceptual knowledge.
- Datasets constructed by asking participants to identify the most important attributes of a concept (McRae et al., 2005; Devereaux et al., 2013)

BANANA	CELLO	SHOES
is_yellow, 29	a_musical_instrument, 26	made_of_leather, 24
a_fruit, 25	has_strings, 16	has_heels, 15
is_edible, 13	made_of_wood, 16	worn_on_feet, 13
is_soft, 12	found_in_orchestras, 13	has_laces, 13
grows_on_trees, 11	is_large, 13	worn_for_protection, 11
eaten_by_peeling, 10	requires_a_bow, 9	has_high_heels, 10

Attribute-based representations

- Problem: attributes have only been collected for a few hundred words (they are expensive and hard to collect).
- If we are to use them in large-scale applications we need to be able to learn attribute representations for any word.
- Solution: predict attribute-based representation from linguistic one. (Fagarasan et al., 2015)

1. Treat the McRae dataset as a bag-of-properties (dimensions: properties, counts: production frequencies).
2. Induce a cross-modal map between linguistic and property-based representations using partial least squares regression.

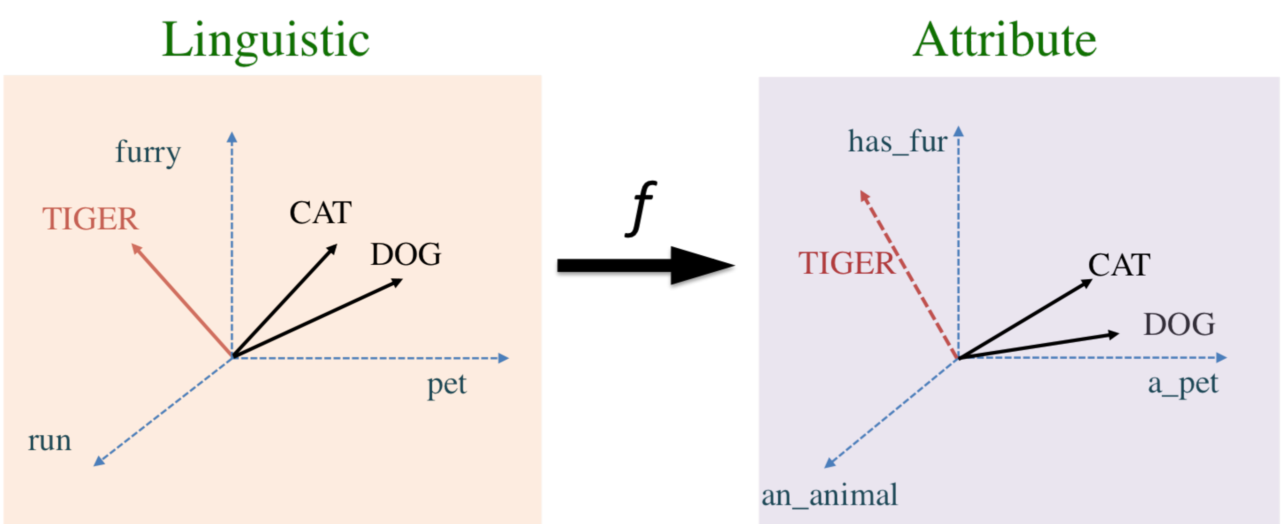

3. Use the learned weights to predict attribute vectors for new words from their linguistic representations.

Learning linguistic representations

We construct two types of linguistic representations:

- EMBED: 100-dimensional word embeddings learnt from Wikipedia using the standard log-linear skip-gram model with negative sampling of Mikolov et al. (2013).
- SVD: 100-dimensional vectors obtained by applying SVD to sparse count-based distributional vectors (Wikipedia, contexts as top 10K most frequent lemmatised words, counts re-weighed using PPMI).

Learning attribute-based representations

We construct two types of attribute representations.

- ATTR-EMBED: induce property-norm representations from EMBED
- ATTR-SVD: induce property-norm representations from SVD

Task: metaphor classification

- We compare the performance of SVD, EMBED, ATTR-SVD and ATTR-EMBED on a metaphor classification task, in order to test our hypothesis as to whether attribute-based semantic representations provide better concept generalisations for metaphor modelling than the widely-used dense linguistic representations.

Metaphorical	Literal
black humor	black dress
filthy mind	filthy garment
young moon	young boy
ripe age	ripe banana
shallow argument	shallow grave
stormy applause	stormy sea

- We use the Tsvetkov et al. (2014) dataset of adjective-noun pairs manually annotated for metaphoricity.
  1. Training: 884 literal and 884 metaphorical pairs
  2. Test: 100 literal and 100 metaphorical pairs

Method

- We use an SVM classifier.
- The input to the classifier is the concatenation of the L2-normalised adjective and noun vectors.
- We evaluated the performance of our classifier on the Tsvetkov test set in terms of precision, recall and F1-score.

Results

- Both types of attribute-based vectors outperform their dense counterparts, which lends support to our hypothesis that property norms offer a suitable level of generalisation of the source and target domains.
- Our hypothesis is that attribute-based methods perform better because the attribute-based dimensions are cognitively motivated and represent cognitively salient properties for concept distinctiveness.

Vectors	P	R	F1
EMBED	0.84	0.65	0.73
ATTR-EMBED	0.85	0.71	0.77
SVD	0.86	0.64	0.73
ATTR-SVD	0.74	0.77	0.75

Qualitative analysis

- Advantage of modelling semantics using attributes is the interpretability of features: every dimension in the space has a fixed interpretation (e.g. is_round, a_bird).
- We can gain insight into how the attributes of metaphorical expressions differ from those of the literal ones.
- In *literal expressions*: adjective and noun in literal expression share a lot of properties. In *metaphorical expressions*: highest ranked properties for noun are ranked low for the adjective and vice-versa.
- This is consistent throughout the test set, showing that the components of literal expressions share many more features than the components of the metaphorical ones.

Conclusion

- We present the first method that uses large-scale attribute-based semantic representations for metaphor identification.
- Our results demonstrate that attribute-based representations provide a suitable level of generalisation for capturing metaphorical mechanisms.

EPSRC
The Leverhulme Trust
UNIVERSITY OF CAMBRIDGE