Interacting Conceptual Spaces

Joe Bolt Bob Coecke Fabrizio Genovese Martha Lewis Dan Marsden Robin Piedeleu

September 29, 2016

1 Introduction

How should we represent concepts and how can they be composed to form new concepts, phrases and sentences? Conceptual spaces theory [Gärdenfors, 2004] gives a way of representing structured concepts, but does not provide a satisfactory account of how they compose. Categorical compositional distributional models [Coecke et al., 2010] provide a successful model of natural language, exploiting compositional structure in a principled fashion. We show how to combine these two accounts, endowing conceptual spaces with a compositional structure.

The categorical compositional distributional programme of [Coecke et al., 2010] successfully integrates two fundamental aspects of language meaning: firstly, the symbolic approach in which meanings of words compose to form larger units; and secondly, the distributional approach where word meanings are derived automatically from text corpora. These two approaches are unified by the key insight that each approach carries the same abstract structure, formalized using category theory.

The abstract framework of the categorical compositional scheme is actually broader in scope than natural language applications. It can be applied in other settings in which we wish to compose meanings in a principled manner, guided by structure. The outline of the general programme is as follows:

1. (a) Choose a compositional structure, such as a pregroup or combinatory categorial grammar.
   (b) Interpret this structure as a category, the **grammar category**.
2. (a) Choose or craft appropriate meaning or concept spaces, such as vector spaces, density matrices, or conceptual spaces.
   (b) Organize these spaces into a category, the **semantics category**, with the same abstract structure as the grammar category.
3. Interpret the compositional structure of the grammar category in the semantics category via a functor preserving the necessary structure.
4. This functor then maps type reductions in the grammar category onto algorithms for composing meanings in the semantics category.

2 The Category of Convex Relations

In [Bolt et al., 2016] we construct a new categorical setting for interpreting meanings in conceptual spaces which respects the important convex structure emphasized in the literature. We show that this category has the necessary abstract structure required by categorical compositional models. We construct convex spaces for interpreting a wide range of word types including nouns, adjectives, verbs, and relative pronouns.

Convex algebras are sets with a mixing operation for forming convex combinations. A convex relation is a binary relation between the underlying sets that respects the formation of convex combinations in a natural way. A **conceptual space** is then defined as one of these convex algebras.

Nouns are defined as convex subsets of a noun space N, and **sentences** as convex subsets of a sentence space S. Examples in a space encompassing colour, taste, and texture are:

$$\begin{aligned} banana &= \overbrace{[60, 95] \times [0.75, 1] \times [0.25, 1]}^{\text{colour (HSV)}} \times \overbrace{Conv(p_{sweet} \cup p_{bitter})}^{\text{taste}} \times \overbrace{[0.2, 0.5]}^{\text{texture}} \\ apple &= [0, 105] \times [0.75, 1] \times [0.5, 1] \times Conv(p_{sweet} \cup p_{sour}) \times [0.5, 0.8] \\ beer &= [40, 50] \times [0.85, 1] \times [0.1, 0.7] \times Conv(p_{sweet} \cup p_{sour} \cup p_{bitter}) \times [0, 0.01] \end{aligned}$$

Pictorially, we have:

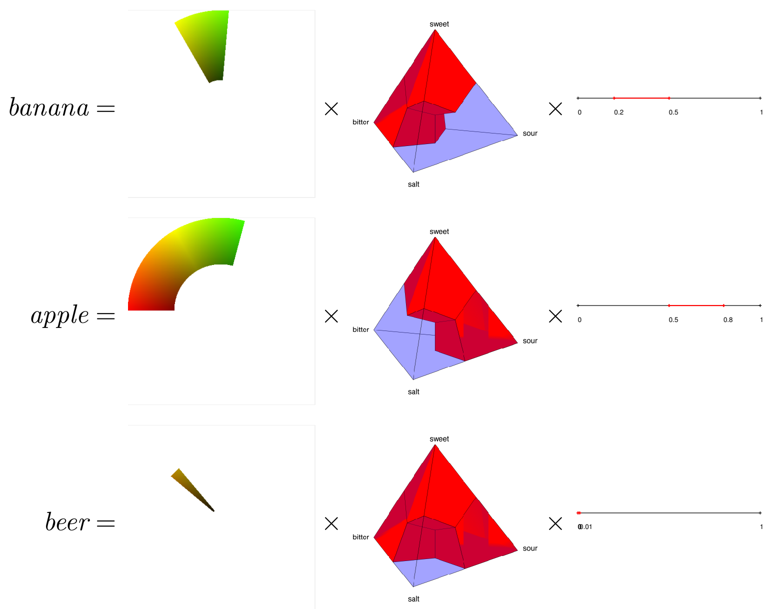

Adjectives are convex relations on the noun space, and **transitive verbs** are relations from two nouns to the sentence space.

What should the sentence space for food and drink be like? We give a very simple example where the events are either positive or negative, and surprising or unsurprising. We therefore use a sentence space of 2-tuples. The first element of the tuple states whether the sentence is positive (1) or negative (0) and the second states whether it is surprising (1) or unsurprising (0). The convex structure on this space is the convex algebra on a join semilattice induced by element-wise max. We have four points in the space: positive, surprising (1, 1); positive, unsurprising (1, 0); negative, surprising (0, 1); and negative, unsurprising (0, 0). Sentence meanings are convex subsets of this space.

3 Concepts in Interaction

Finally, our categorical description of conceptual spaces allows us to use grammatical structure to compose meanings. The composition is effected by the relation $\epsilon_N : N \times N \to \{*\}$, where (a, b) maps to $*$ if $a = b$. All notation used in this section is fully described in [Bolt et al., 2016].

The application of $yellow_{adj}$ to *banana* works as follows.

$$\begin{aligned} yellow\ banana &= (1_N \times \epsilon_N)(yellow_{adj} \times banana) \\ &= (1_N \times \epsilon_N)\{(\overrightarrow{x}, \overrightarrow{x}) | x_{colour} \in yellow\} \\ &\quad \times ([60, 95] \times [0.75, 1] \times [0.25, 1] \times Conv(p_{sweet} \cup p_{bitter}) \times [0.2, 0.5]) \\ &= [60, 75] \times [0.75, 1] \times [0.25, 1] \times Conv(p_{sweet} \cup p_{bitter}) \times [0.2, 0.5] \end{aligned}$$

Notice, in the last line, how the hue element of the banana has altered from $[60, 95]$ to $[60, 75]$. The same calculation gives us:

$$soft\ apple = [0, 105] \times [0.75, 1] \times [0.5, 1] \times Conv(p_{sweet} \cup p_{sour}) \times [0.4, 0.6]$$

Using the definition of *taste* that we gave, we find that although sweet bananas are good:

$$\begin{aligned} bananas\ taste\ sweet &= (\epsilon_N \times 1_S \times \epsilon_N)(bananas \times taste \times sweet) \\ &= (\epsilon_N \times 1_S)(banana \times (green\ banana \times \{(1, 1)\} \cup yellow\ banana \times \{(1, 0)\}) \\ &= \{(1, 1), (1, 0)\} = positive \end{aligned}$$

sweet beer is not so desirable:

$$\begin{aligned} beer\ tastes\ sweet &= (\epsilon_N \times 1_S \times \epsilon_N)(beer \times taste \times sweet) = \{(0, 1)\} \\ &= negative\ and\ surprising \end{aligned}$$

Relative Pronouns The compositional semantics we use can also deal with relative pronouns, i.e. words such as 'which'. This uses additional compositional structure. $\mu_N : N \times N \to N$ maps matching pairs down to the singleton, i.e. $(a, a) \mapsto a$, and $\iota_S : S \to \{*\}$ is a discarding map. As an example, we can form the noun phrase *Fruit which tastes bitter*. This has the structure:

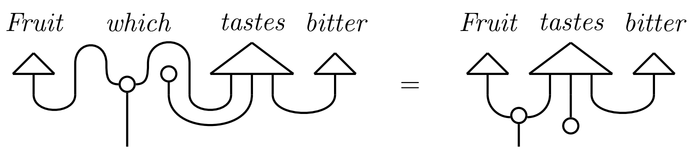

In our example, we find that *Fruit which tastes bitter* = *green banana*:

$$\begin{aligned} &Fruit\ which\ tastes\ bitter \\ &\quad = (\mu_N \times \iota_S \times \epsilon_N)(Conv(bananas \cup apples) \times taste \times bitter) \\ &\quad = (\mu_N \times \iota_S)(Conv(bananas \cup apples) \times (green\ banana \times \{(0, 0)\})) \\ &\quad = \mu_N(Conv(bananas \cup apples) \times (green\ banana)) = green\ banana \end{aligned}$$

References

[Bolt et al., 2016] Bolt, J., Coecke, B., Genovese, F., Lewis, M., Marsden, D., and Piedeleu, R. (2016). Interacting conceptual spaces. In Workshop on *Semantic Spaces at the Intersection of NLP, Physics and Cognitive Science*.

[Coecke et al., 2010] Coecke, B., Sadrzadeh, M., and Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 36:345–384.

[Gärdenfors, 2004] Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought.* The MIT Press.