Normativity, Meaning Plasticity, and the Significance of Vector Space Semantics

Eleni Gregoromichelaki
Heinrich-Heine University and King's College London
elenigregor@gmail.com

Christine Howes
University of Gothenburg
christine.howes@gu.se

Arash Eshghi
Heriot-Watt University
a.eshghi@hw.ac.uk

Ruth Kempson
King's College London
ruth.kempson@kcl.ac.uk

Mehrnoosh Sadrzadeh
University College London
m.sadrzadeh@ucl.ac.uk

Julian Hough **Matthew Purver** **Gijs Wijnholds**
Queen Mary University of London
{j.hough,m.purver,g.j.wijnholds}@qmul.ac.uk

Abstract

This paper continues the discussion started in (Lücking et al., 2019), on the suitability or otherwise of Vector Space Semantics (VSS) as a model of semantics for NL in interaction.

1 Introduction

Lücking et al. (2019) argue that the *distributional hypothesis* (DH) cannot lead to a psychologically realistic modelling of natural language (NL) due to its inability to stand as an autonomous basis for semantics. Instead, they propose a model of the conceptual mechanisms underpinning NL interaction involving direct encoding of conceptual structures and processes on individual brains (Cooper, 2019). Problems of agent coordination are then resolved as meaning negotiation and learning based on game-theoretic modelling of symbolic signalling that presupposes mental states with hardwired discourse structure (DGBs).

We find many points of agreement with Lücking et al. (2019). However, we believe that not all versions of implementing DH/VSS fall under their criticism. Although in the past most operationalisations of DH have involved only word distributions, the recent multimodal trend involves not only textual but also image and even audio contexts (e.g. Kiela and Clark, 2017; Bruni et al., 2014). Indeed, from early on, such models have envisaged their extension to distributional representations that include situational contexts (see e.g. Landauer and Dumais, 1997, a.o.) and, in our view, at least the combination of Dynamic Syntax (DS; Kempson et al., 2001, 2016) and VSS (DS-VSS, Kempson et al., 2019; Sadrzadeh et al., 2018; Wijnholds et al., 2019) operates under assumptions resolving the issues the authors raise.

Instead of employing individualistic referential mechanisms, DS proposes that semantic content emerges in interaction rather than in the correspondence of representations in the brain to entities in the world (Gregoromichelaki, 2019; Gregoromichelaki et al., 2019). Hence, the structures manipulated by DS constitute complex, highly-structured predictive triggers (*affordances*) for further verbal/nonverbal actions. This idea has been computationally implemented in DS-TTR (Eshghi et al., 2017; Kalatzis et al., 2016; Eshghi and Lemon, 2014) where, in a Reinforcement Learning model, it is Record Types (RTs) of Type Theory with Records (TTR; Cooper, 2005; Cooper and Ginzburg, 2015) that are the triggers for further action (dialogue contexts): clusters of RTs are learned from interaction histories and, accordingly, a potential next response is chosen. But, under the same assumptions, the DS-VSS integration appears to be equally suitable for the same purpose, especially since it would appear to better capture the nondiscrete, gradient effects associated with such triggers.

On both the DS-TTR and DS-VSS views, normative semantic attributions do not concern facts about individuals but relational facts about characterisations of transactions of individuals with the sociomaterial environment. Meaning then arises in interaction, on the basis of affordances made available to agents by sociomaterial settings ('forms of life') that groups establish to direct their perceptual and action capacities. In concrete situations, agents selectively engage with multiple affordances available in such *affordance landscapes* (Rietveld et al., 2018; Bruineberg et al., 2018). These socially-established affordances constitute a general basis of normativity both for action/perception and NL meaning, in that individual agents can have partial or imperfect grasp of such potentials depending on their level of expertise. This is because individuals engage with affordances through the experience of *solicitations* (Dreyfus and Kelly, 2007): agents have abilities, dispositions, and concerns regarding their interactions which define the saliency of particular affordances in concrete situations; and individual abilities and values are acquired through histories of interactions in particular settings.

This, we suggest, is where the aptness of DH and VSS tools lies. In combination with DS, such models can be seen as implementing exemplar accounts of categorisation (Nosofsky, 2011) in that

the matrix representations record episodic memories of contexts of perception/action involving particular stimuli (here, *words*). Word forms in DS trigger sets of incremental actions and predictions; and past experiences with such stimuli-situation pairs is what is stored and retrieved in processing. Past co-occurrence, "similarity" relations, can then underpin associationist and probabilistic mechanisms of online selective attention (*affordance-competition*) that result in incrementally appropriate word retrieval (via activation facilitation) in production and contextualisation (narrowing-down or enrichment) in comprehension. Thus the significance of words emerges from joint (re)constructive acts during use: runtime operations over high-dimensional VS representations (e.g., context-aware analogy cf. Landauer, 2002) enable agents to engage with probabilistic distributions over fields of predictions of further opportunities for action thus grounding normativity in local exchanges. On the other hand, abstractions underpinning *explicit* normative judgements, e.g., truth-conditional judgements, reference, grammaticality etc, are phenomena definable only at a historical and group level of analysis, "bootstrapped" from more basic, domain-general psychological capacities, and do not play a fundamental grounding role in NL performance.

On this view, then, individual agents' memories do not store transductions of perceptual input into symbolic conceptual representations (cf. Larsson, 2015). Instead, conceptual capacities are abilities to discriminate alternative responses to similar or dissimilar stimuli arrays (cues). Classical theories of learning, like reinforcement or discriminative learning (Rescorla and Wagner, 1972), can then be employed to model the constantly evolving fluid responsiveness to NL stimuli, even highly underspecified ones, like indexical pronouns, wh-elements, and names. For example, in learning the distinction between the English words *I*, *you*, and *he*/*she*, infants are initially expected to display inconsistencies and individual differences depending on their personal experience with input, as they will not be "attuned" sufficiently to the ambient social invariances that license the use of each form. Recorded episodes of experience with pronouns as cues for action initially will be too few and too restricted to enable development of speaker/addressee/non-participant discriminatory features to 'solicit' the affordances that characterise appropriate pronominal usage. But, in the face of discrepancies between their own predictions and actual experience, the infant will gradually come to discriminate salient aspects of the discourse environment to serve as cues for the choice of form. Such a shift only becomes possible, however, if there are options available, namely, sufficiently "similar" competing cues (in DS, *triggers* of actions) that occur in similar contexts (language games) like the various alternative forms of pronouns. Such triggers compete with each other on the basis of their predictive value regarding subsequent events (in DS, further opportunities for rewarding interaction or avoidance of undesirable consequences). Competition means that loss of associative strength by one cue results in reinforcement of the other(s) in the same category ensuing in an emergent systematic pattern of contrasts (Rescorla and Wagner, 1972). Moreover, given that lexical triggers are necessarily fewer than discourse situations/features, the same forms can come to acquire added triggering effects by the same process, i.e., contextual co-occurrence overlap and subsequent discrimination on the basis of prediction error: for example, *you* can come to include or not multiple addressees in multiparty dialogue, or acquire an impersonal use that might include the speaker when the combination of contextual features are sufficiently discriminative. Such cases cannot be handled easily by the model presented in Lücking et al. (2019) because these uses underdetermine, disregard, or eliminate the hard-wired distinctions postulated in their DGB-based modelling with arbitrary homonymies appearing as the only available solution.

Similarly, regarding proper names, it is storage of life episodes incidentally involving particular interlocutors that resolves the problem of "referential" uncertainty by means of relying solely on domain-general memory mechanisms rather than specific assumptions about conceptual/discourse structure. In fact, Gregoromichelaki et al. (2011) argue that "mindreading" effects can be accounted for exactly because of such co-occurrence mechanisms that employ names as cues for invoking past interactions with discourse participants to ground appropriate redeployment (Horton and Gerrig, 2005), rather than assuming explicit representations of common ground or metarepresentational reasoning.

Overall then, given DS-VSS modelling of both word meaning and syntax alike as (socio)-cognitive predictive and incremental mechanisms, compositional VSS employing tensors and tensor contraction provides a fruitful implementation of exemplar-based categorisation, thus modelling the emergence of NL polysemy, as well as 'ad hoc concept' enrichment/narrowing effects, which otherwise remain a mystery (Partee, 2018). Without such an extension of our theoretical vocabulary (Sadrzadeh et al., 2018), we believe that progressive achievement of NL acquisition, the emergent fluency of conversational exchange not only with familiars but in arbitrary multiparty exchanges, and the inexorability of NL change all threaten to continue to elude us (Kempson et al., 2018).

References

Bruineberg, J., E. Rietveld, T. Parr, L. van Maanen, and K. J. Friston (2018). Free-Energy Minimization in Joint Agent-Environment Systems: A Niche Construction Perspective. *Journal of Theoretical Biology 455*, 161–178.

Bruni, E., N. K. Tran, and M. Baroni (2014). Multimodal Distributional Semantics. *Journal of Artificial Intelligence Research 49*, 1–47.

Cooper, R. (2005). Records and record types in semantic theory. *Journal of Logic and Computation 15*(2), 99–112.

Cooper, R. (2019, June). Representing Types as Neural Events. *Journal of Logic, Language and Information 28*(2), 131–155.

Cooper, R. and J. Ginzburg (2015). Type theory with records for natural language semantics. *The Handbook of Contemporary Semantic Theory*, 375–407.

Dreyfus, H. and S. D. Kelly (2007, March). Heterophenomenology: Heavy-Handed Sleight-of-Hand. *Phenomenology and the Cognitive Sciences 6*(1), 45–55.

Eshghi, A. and O. Lemon (2014). How domain-general can we be? Learning incremental dialogue systems without dialogue acts. In *Proceedings of Semdial 2014 (DialWatt)*.

Eshghi, A., I. Shalyminov, and O. Lemon (2017). Bootstrapping incremental dialogue systems from minimal data: the generalisation power of dialogue grammars. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Gregoromichelaki, E. (2019). Natural Languages as Distributed Action Systems. In *Third Dynamic Syntax Conference*, Valletta, Malta. May 2019.

Gregoromichelaki, E., C. Howes, and R. Kempson (2019). Actionism in syntax and semantics. In *Dialogue and Perception*. CLASP Papers in Computational linguistics. under review.

Gregoromichelaki, E., R. Kempson, M. Purver, G. J. Mills, R. Cann, W. Meyer-Viol, and P. G. T. Healey (2011). Incrementality and intention-recognition in utterance processing. *Dialogue and Discourse 2*(1), 199–233.

Horton, W. and R. Gerrig (2005). Conversational common ground and memory processes in language production. *Discourse Processes 40*(1), 1–35.

Kalatzis, D., A. Eshghi, and O. Lemon (2016). Bootstrapping incremental dialogue systems: using linguistic knowledge to learn from minimal data. In *Proceedings of the NIPS 2016 workshop on Learning Methods for Dialogue*, Barcelona.

Kempson, R., R. Cann, E. Gregoromichelaki, and S. Chatzikyriakidis (2016). Language as mechanisms for interaction. *Theoretical Linguistics 42*(3-4), 203–276.

Kempson, R., E. Gregoromichelaki, and C. Howes (2018). Language as mechanisms for interaction: Towards an evolutionary tale. In A. Silva, S. Staton, P. Sutton, and C. Umbach (Eds.), *Proceedings of the 12th TbiLLC Conference*, pp. 209–227.

Kempson, R., J. Hough, C. Howes, M. Purver, P. G. T. Healey, A. Eshghi, and E. Gregoromichelaki (2019). Why natural language models must be partial and shifting: a Dynamic Syntax with Vector Space Semantics perspective. In *Vector Semantics for Dialogue and Discourse workshop at IWCS 2019*.

Kempson, R., W. Meyer-Viol, and D. Gabbay (2001). *Dynamic Syntax: The Flow of Language Understanding*. Blackwell.

Kiela, D. and S. Clark (2017, December). Learning Neural Audio Embeddings for Grounding Semantics in Auditory Perception. *Journal of Artificial Intelligence Research 60*, 1003–1030.

Landauer, T. K. (2002, January). On the computational basis of learning and cognition: Arguments from LSA. In *Psychology of Learning and Motivation*, Volume 41, pp. 43–84. Academic Press.

Landauer, T. K. and S. T. Dumais (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review 104*(2), 211–240.

Larsson, S. (2015). Formal Semantics for Perceptual Classification. *Journal of Logic and Computation 25*(2), 335–369.

Lücking, A., R. Cooper, S. Larsson, and J. Ginzburg (2019). Distribution is not enough: going Firther. In *NLCS - Natural Language and Computer Science 6 workshop at IWCS 2019*.

Nosofsky, R. M. (2011). The Generalized Context Model: An Exemplar Model of Classification. In *Formal Approaches in Categorization*. Cambridge University Press.

Partee, B. H. (2018). *Changing notions of linguistic competence in the history of formal semantics*, pp. 172–196. Oxford University Press.

Rescorla, R. A. and A. R. Wagner (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy (Eds.), *Classical Conditioning II: Current Research and Theory 2*, pp. 64–99. Appleton-Century-Crofts.

Rietveld, E., D. Denys, and M. Van Westen (2018). Ecological-Enactive Cognition as engaging with a field of relevant affordances. In *The Oxford Handbook of 4E Cognition*, pp. 41. Oxford University Press.

Sadrzadeh, M., M. Purver, J. Hough, and R. Kempson (2018, November). Exploring Semantic Incrementality with Dynamic Syntax and Vector Space Semantics. In *Proceedings of the 22nd SemDial Workshop on the Semantics and Pragmatics of Dialogue (AixDial)*, Aix-en-Provence, pp. 122–131.

Wijnholds, G., M. Purver, M. Sadrzadeh, and R. Kempson (2019). Incremental Semantic Judgements. In *3rd Dynamic Syntax Conference*, Valletta, Malta. May, 2019.