

PREDICTIVE PROCESSING AND THE SEMIOLOGICAL PRINCIPLE: COMMENTARY TO DUFFLEY

GUIDO LÖHR¹

¹<https://orcid.org/0000-0002-7028-3515>
University of Technology
Department of Philosophy
Eindhoven
Netherlands
loehrg@icloud.com

CHRISTIAN MICHEL²

²<https://orcid.org/0000-0001-9962-5403>
University of Edinburgh
Department of Philosophy
Edinburgh, Edinburgh
United Kingdom of Great Britain and Northern Ireland
s1443279@sms.ed.ac.uk

Article info

CDD: 410

Received: 12.08.2021; Revised: 28.09.2021; Accepted: 04.10.2021

<https://doi.org/10.1590/0100-6045.2022.V45N1.GM>

Keywords

Duffley

Meaning

Concepts

Predictive Processing

Polysemy

Abstract: The aim of this commentary is to underpin Duffley's notion of a stable mental content that corresponds to the literal word meaning with a computationally plausible cognitive theory. Our approach is to investigate what these stable contents could be according to the so-called *Predictive Processing* architecture. We argue that recent advances in cognitive science can make at least two contributions to the debate. First, they can provide some underpinning of Duffley's ideas of a stable linguistic meaning

associated with the sign. Second, they provide resources to understand how the semiological principle is compatible with a dynamic and flexible notion of "meaning".

1. Introduction¹

According to Patrick Duffley (2020, p. 3), “the primary task of linguistic semantics should be to work out an analysis of the mental content attached to the linguistic sign by the language community in which the latter is used.” Given that Duffley is a linguist interested in meaning, the topic of his *Linguistic Meaning Meets Linguistic Form* is supposed to answer the following question: what is the “mental content” of a given linguistic sign?

The core claim of Duffley’s book is that linguistic meaning needs to be stable, i.e., independent of context and world knowledge. Duffley emphasizes words among the linguistic and meaning bearing units, so we will focus in the following on words and their meanings. Linguistically encoded word meaning associated with a linguistic sign is fixed and the starting point for subsequent pragmatic processing rather than pragmatic processing being used to fix literal standing meaning. Thus, Duffley defends invariantism regarding concepts – however only in the psychological sense of the term ‘concept’ (c.f. Machery, 2009; Löhr, 2017; cf., Löhr, 2020 for different notions of ‘concept’). He does not endorse a strict distinction between semantics and pragmatics in the sense that pragmatics cannot penetrate the process of generating a propositional meaning (cf., Löhr, forthcoming, b, for a review of different views of this distinction). In the philosophical sense then, he

¹ Authors are listed in alphabetical order and contributed equally to the manuscript.

is a contextualist like Recanati (2010) or Carston (2008), who emphasize the continuity between semantics and pragmatics.

Duffley's motivation for the search for a stable mental content associated with our linguistic labels is the *semiological principle*. According to this principle, "there is a stable notional content associated with each linguistic sign" (p. 37). Duffley thinks that the semiological principle is violated by ignoring the stable semantic contribution of words. Duffley rejects the far-reaching softening of the semantic-pragmatic distinction that cognitive linguists like, e.g., Langacker (1987, p. 2) seem to endorse.

Why does Duffley think that linguistic processing requires a stable notional meaning in the first place? One possible reason, which Duffley has put forward elsewhere (Duffley, 2009), is related to the possibility of communication. "Publicity" is seen by many as a core desideratum for theories of concepts (e.g., Fodor, 1998; Prinz 2002).² Duffley argues that communication would be impossible if there were no stable meaning to which the sign is related and that is shared by all speakers in the language community. Another reason seems to be that Duffley believes that a linguistic sign simply cannot connect to a linguistic meaning if the latter is flexible and variable. The partner of the linguistic sign would be ill-defined and a moving target.³

² Note that we believe that stability and shareability/publicity are orthogonal. Alternatives, like Davidson (1986) show that we can share the same concept, without sharing stable concepts, because we could *converge* on the "same meaning" during a conversation.

³ We believe that an increasing body of evidence points to a rather holistic processing of linguistic material which undermines a semantic-pragmatic distinction, at least in the form of the sequential processing paradigm (first linguistic meaning, then pragmatic meaning).

Given the need for a fixed linguistic meaning, Duffley thinks that the semiological principle is inconsistent with much work in cognitive linguistics that attempts to explain a number of linguistic phenomena in terms of polysemy (Langacker, 1987; Lakoff, 1987 etc.). A word is polysemous, in Duffley's sense (see Löhr, forthcoming, a, for a review of different notions of polysemy), if the linguistic community associates many different but *related* meanings with a word form. It seems that polysemy might in some cases undermine the semiological principle because it suggests that each word is associated with a flexible cloud of mental contents as opposed to just one that is stable.

Duffley however allows for such clouds of mental content for *encyclopedic* words. Different parts of the complex knowledge network (which constitutes the linguistic meaning) might be activated on different use occasions. Take Duffley's example "Your cell phone is an elephant". Here only the notion of size and ungainliness are selected from the total linguistic meaning (which includes other features like elephants are big, mammals, have trunks, etc.) However, he insists that the encyclopedic knowledge structure must be stable enough to make communication possible.

So, while linguistic meaning is relatively stable, in Duffley's view, the meaning of words can be pragmatically modulated by selecting part of the encyclopedic knowledge structure. This makes his view compatible with the view that word meanings are dynamic and flexible, which is gaining increasing momentum in many disciplines (cf., Michel, 2020b).

However, Duffley suggests that there is a second, qualitatively distinct type of words, namely *monosemic* words, which are semantically much simpler compared to encyclopedic words. Examples are quantifiers like "any" or prepositions like "for". The meaning of those words is highly abstract and pragmatic modulation, so Duffley suggests,

works by enrichment rather than through the selection of parts of a rich encyclopedic knowledge structure. So, here he follows Ruhl (1989) and rejects polysemy of those words. For instance, in his case study of "for", Duffley shows that the many different senses (he deals with twelve) ascribed to this preposition should be seen as products by enrichment of a single highly abstract schema constituting its stable linguistic meaning.

The aim of this commentary is to underpin Duffley's notion of a stable mental content that corresponds to the literal word meaning with a computationally plausible cognitive theory. Our approach is to investigate what these stable contents could be according to the so-called *Predictive Processing* architecture.

We argue that recent advances in cognitive science can make at least two contributions to the debate. First, they can provide some underpinning of Duffley's ideas of a stable linguistic meaning associated with the sign. Second, they provide resources to understand how the semiological principle is compatible with a dynamic and flexible notion of "meaning".

Once we have spelled out a cognitive-computational model of "cognitive meaning", we can conclude that the semantic-pragmatic distinction really does not amount to very much anymore, beyond a merely terminological issue. The view that semantic and pragmatic processing are indeed difficult to disentangle on a cognitive level is supported by empirical findings that show that meaning processing is often immediate and holistic (e.g., Hagoort, & van Berkum 2007; Bašnáková, Weber, Petersson, van Berkum & Hagoort 2014). Furthermore, we suggest that a natural implication of the model proposed here is that - from a point of view of cognitive-computational representation and processing - a dichotomic distinction between monosemic and encyclopedic words might be unnecessary (though such a

distinction might still be useful from a point of view of linguistic analysis).

2. A cognitive-computational model for "word meaning"

More and more neuroscientists take the brain to entertain a hierarchical generative model (e.g., the so-called Predictive Processing (PP) framework: see Clark 2016; Hohwy 2013). The function of such a model is to constantly anticipate its sensorimotor input with minimal error. For that purpose, the brain's representations are organized hierarchically in network of nodes. "Predictions" (or hypotheses) are constantly generated on all levels, from lowest level perceptual representations (e.g., pixels in the retina) to the most abstract conceptual representations like complex situational or contextual patterns. Higher level representations serve as constraints for predictions on lower levels and lower levels feed prediction error signals upwards so that the model can be improved, and the predictions can be better next time.

If we map Langacker's "semantic space" (see, e.g., Langacker 1987, pp. 76-80) consisting of all the conceptualizations, i.e., both phonetic and semantic ones, on the many levels of schematization (or the equivalent construct-i-con of Construction Grammar) on such a hierarchical model, we get a highly interconnected multilayer system. Language comprehension and production according to this PP account can thus be viewed as a holistic process of finding an error minimal equilibrium in such an interlocked hierarchical system (e.g., Michel 2019; see also Pickering & Garrod for a model of language processing in the vicinity of Predictive Processing).

In this view, at least most nouns, verbs and adjectives are associated with "activation (or information) packages" (e.g., Ortega-Andrés & Vicente, 2019; Löhr & Michel, manuscript). Activation packages are sub-hierarchies of the overall model. The root node of an activation package represents the most abstract "sense" of a word, very much like Ruhl's (1989) monosemic meaning. All packages are tightly interconnected and together they form the overall generative model. Furthermore, a package's node structure might reach down to the sensorimotor periphery, i.e., the location of modality specific information, like color, sound, smell, etc.

Now, let's take Duffley's first case study of "for" (p. 37ff). This word would be associated with a root node that represents its most schematic sense. Below we would have more nodes that represent more specific senses over which the root node has abstracted. In the case of monosemic (as opposed to encyclopedic) words, the hierarchy is not very rich. Langacker's example of "tree", which Duffley refers to (p.169, figure 4.1), can be seen as a toy example of such an activation package. However, it is lacking depth and many more layers of nodes should be added and the structure should go down to the sensorimotor periphery.

In the PP view, which exact parts of a package are activated on a given occasion, however, is context dependent. A package can be activated *shallowly* (only the highest-level nodes) or *deeply* to different degrees (including lower-level nodes). Given the holistic processing of the PP model, no specific modulation of the package should be seen as privileged (see also Ludlow 2014). Because error minimization is approximate Bayesian inference, ultimately all processing is inference (and hence "pragmatic"). The distinction between semantics and pragmatics gets blurry.

There is, of course, a privileged *node*, namely the root node of the package, which is also content bearing. The root

node has the most abstract sense as mental content and is privileged because it has a link with a representation of the word sign (which is an information package itself with all of the phonological information - phonological in Langacker's broad sense, which considers phonological information itself as meaningful conceptualizations).

For the sake of terminology, we could perfectly follow Duffley and call the content represented by this root node "linguistic meaning". Different use occasions of the word might activate different parts of the associated information package. However, the link of the sign with the information package root node is stable as is (relatively) the whole information package. Only the context-dependent selected parts vary.

Our point is the following: Based on the best of our knowledge of the cognitive processing underlying language related actions, the processing of the information package, i.e., the pattern of which parts of it get activated, is holistic and *not* sequential (i.e., it is not the case that first the root node is processed, and only then other nodes). Moreover, it is context dependent. The context is represented as other higher order external nodes influencing the nodes of a given activation package. In other words, a word in a sentence is already located in a specific context. So according to the PP model, the sentence as a higher-level unit modulates (influences/constraints as a "prior") top down the meaning of the word potentially to a more specific sense (a lower node in the information package), just in virtue of forming part of a sentence (not yet in virtue of the sentence *use*).

With all that said, we might still agree with Duffley that a word in isolation (without any context in some idealized way of looking at it) just might "mean" its highly abstract "linguistic meaning" (the root node content) or the whole package associated with it. This is what the word sign "contributes" to the sentence meaning.

Take for example Duffley's example of "the bird is safe". The word "safe" might have the more specific senses "free of the risk of being harmed" and "free of the risk of being lost". Still, maybe the abstract (root node) meaning of the isolated (context free) "safe" is something like "free of the risk of something bad happening to it". In the sentence "The bird is safe" (without further context except the other words in the sentence) "safe" is being modulated into "free of the risk of being harmed". In this case, the word "bird" serves as a (linguistic) context for "safe" that modulates the generic "free of the risk of something bad happening to it" into "free of the risk of being harmed". The reason is that normally we don't expect a bird being lost, only that it is being harmed. In "The money is safe" the same word "safe" is being modulated into the other more specific sense, namely "free of the risk of being lost". The reason is that we normally do not expect that money is being harmed, only that it is lost.

Now if we add yet another layer of context, say a scene with a cat threatening a bird in a cage, the sentence might be modulated into the *message* meaning "The bird is safe *from the cat*". We can now add further layers of context like a specific situation embedded in a more general situation. Imagine a father calming down his daughter who is concerned that the cat might devour the bird. The overall message inferred might then be *Don't worry, trust your Daddy who cares for you: you need not be concerned with the bird being eaten by the cat.*

Now in the PP framework all of those layers of "meaning" can be mapped onto the hierarchy. Linguistic processing does not happen in a linear fashion from the most abstract root node of the word to the highest-level meta-message involving the word in a sequential pragmatic process. Rather, the whole system tries to come into an error minimizing equilibrium state holistically on all levels simultaneously.

We now get a picture of a stable linguistic word meaning, but at the same time it is clear that a word as an activation package is highly flexible and constantly being context-dependently modulated (i.e., a subnetwork of nodes is being selected), without there being some (cognitively) privileged modulation. The root node is privileged and has a special role only from the point of view of an outside linguistic description with isolated context-free words. In such a situation the word is "mentioned" but not "used". At the moment where the word is used, it is inevitably a use in a context and hence it will be modulated. In other words, "linguistic meaning" as the ideal stable notional non-modulated content of the root node is a cognitive idealization. Any cognitive process happens in some context (some other connected nodes are firing or not). This does not mean however, that such a notion of "linguistic meaning" is not useful for linguistic description and analysis, or that it is not useful to give a special name ("linguistic meaning") to the root node content or the whole activation package.

3. Embodiment of word meaning

We would now like to touch upon Duffley's critique on embodiment. Duffley claims that many abstract words are not embodied. It all depends, of course, on what you mean by "embodiment" and "abstract" (cf., Löhr, forthcoming, a, for an attempt to define the notion of abstract concepts). A source of confusion might be to think that embodied linguistic cognition is committed to the idea that every word corresponds to co-activations of a concrete, low level sensory or motor representation (cf., Löhr, 2019). Indeed, words like "knowledge" and "for" do not seem to correspond to any concrete imagery or ways of interacting

with the world. Still, we can separate philosophical issues of linguistic meaning from questions of how we apply and represent the conditions of application of such words cognitive-psychologically (Löhr, 2020). The latter could be embodied even if the former is externalized (Löhr, 2021). Imagine for instance voting scenes when hearing the word "democracy". Even if such event representations fail as a sufficient philosophical analysis of the meaning of "democracy" we may still rely on such representations to decide whether or not to apply the word.

We believe that Duffley could perfectly take on board a notion of embodiment spelled out in terms of the model proposed here (cf., Löhr, 2019, Löhr & Michel, manuscript). An activation package tree bottoms out at the sensorimotor level. The hypothesis is that even highly abstract, monosemic words like "any" should be seen as embodied (note that Langacker famously endorsed that abstract concepts like "time" are represented in spatial - and hence sensorimotor - terms). To the objection (Duffley, in personal conversation) that "any" is a merely "mental process of random selection of a referent from a reference mass", the embodied cognitivist can respond that a "process of selection" is an action, and hence its conceptualization is in (complex) sensorimotor terms.

So, according to our proposed cognitive model, there might be no reason to posit two qualitatively distinct types of representation, those corresponding to monosemic and to encyclopedic words. Both can in principle be implemented cognitively as activation packages, with abstracted root nodes. "Elephant" is an activation package with an abstract root node that - on its own - has a very unspecific content that is compatible with all of the more specific encyclopedic knowledge that is represented in lower-level nodes (Eliasmith 2013 or Thagard 2019 call this a "semantic pointer" architecture). Such root nodes arise from a complex

abstraction and convolution process of sensorimotor experiences; or indirectly (e.g., in case of learning a definition) by grasping *other* words that *are* sensorimotor grounded (see also Michel 2020a).

Depending on the context, sometimes the activation package can be processed deeply down to the sensorimotor level. Then vivid and specific imagery could be activated upon hearing an abstract word (e.g., a voting scene for "democracy"). Such imagery and other sensorimotor information is simply part of the package, awaiting optional activation.

Finally, note that the predictions made by the PP-based model cannot be separated from perception-action because of the holistic way the PP brain works. We have a model with interconnected representations on all levels in the hierarchy. Prediction making takes place all the time and in the whole model simultaneously and perception, action and cognition are tightly integrated.

4. Conclusion

Duffley's aim is to show how a single stable mental content associated with a linguistic form can go a long way in explaining a number of linguistic phenomena that many cognitive linguists take to be evidence for polysemy. We have tried to underpin his idea with a cognitive model that assumes that the mind is predicting rather than merely reacting. We argued that Duffley's notion of a stable word meaning can be identified with the root node of neurally organized activation packages, or the whole package.

However, a picture of linguistic processing arises that is holistic and where words are highly flexible and dynamic entities that are constantly modulated within a context of

representations on all levels of abstraction (sentences, situations, situational contexts, etc). Words, when used - not when theorized about as isolated context-free entities - are always used in a context which drives the modulation of the package without there being any privileged modulation.

In this way, we can reconcile Duffley's concern for the need of a stable linguistic meaning in order to save the semiological principle, on the one hand, and the tendency within the cognitive linguistics camp to consider the boundary between semantics and pragmatics as blurred, on the other.

References

- Bašnáková, J., Weber, K., Petersson, K. M., van Berkum, J., & Hagoort, P. (2014). Beyond the Language Given: The Neural Correlates of Inferring Speaker Meaning. *Cerebral Cortex*, 24(10), 2572–2578.
- Carston, R. (2008). *Thoughts and utterances: The pragmatics of explicit communication*. John Wiley & Sons.
- Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Davidson, D. (1986). A Nice Derangement of Epitaphs. In E. Lepore (Ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (pp. 433-446). Blackwell.
- Duffley P. (2009). Does the Rubber Ever Hit the Road? What Human Language Shows about the Relation Between Mind and Reality. <https://youtu.be/g2vtsdgDeC0>
- _____. (2020). *Linguistic meaning meets linguistic form*. Oxford University Press, USA.

- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.
- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford University Press.
- Hagoort, P., & van Berkum, J. (2007). Beyond the sentence given. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 801–811.
- Hohwy, J. (2013). *The predictive mind* (First edition). Oxford University Press.
- Langacker, R. W. (1987). *Foundations of cognitive grammar*. Stanford, California: Stanford University Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. University of Chicago press.
- Ludlow, P. (2014). *Living words: Meaning underdetermination and the dynamic lexicon* (First edition). Oxford University Press.
- Löhr, G. (2017). Abstract concepts, compositionality, and the contextualism-invariantism debate. *Philosophical Psychology*, 30(6), 689-710.^[1]_[SEP]
- _____. (2019). Embodied cognition and abstract concepts: Do concept empiricists leave anything out? *Philosophical Psychology*. 32(2), 161–185.^[1]_[SEP]
- _____. (2020). Concepts and categorization: do philosophers and psychologists theorize about different things? *Synthese*. 197(5), 2171-2191. Online, 2018.
- _____. (2021). Social constructionism, concept acquisition and the mismatch problem. *Synthese* 198, 2659–2673. Online 2019.

- _____. (forthcoming, a). What are abstract concepts? On lexical ambiguity and concreteness ratings. *Review of Philosophy and Psychology*.
<https://doi.org/10.1007/s13164-021-00542-9>
- _____. (forthcoming, b). Does polysemy support radical contextualism? On the relation between minimalism, contextualism and polysemy. *Inquiry*.
<https://doi.org/10.1080/0020174X.2020.1868329>
- _____. & Michel (manuscript). Copredication in context: A Predictive Processing account.
- Michel, C. (2019). The Liar Paradox in the predictive mind. *Pragmatics & Cognition*, 26(2-3), 239-266.
- _____. (2020a). Overcoming the modal/amodal dichotomy of concepts. *Phenomenology and the Cognitive Sciences* 20, 655–677.
<https://doi.org/10.1007/s11097-020-09678-y>
- _____. (2020b). Concept contextualism through the lens of Predictive Processing. *Philosophical Psychology* 33 (4), 624-647
- Ortega-Andrés, M., & Vicente, A. (2019). Polysemy and copredication. *Glossa: a journal of general linguistics*, 4(1).
- Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, 36(04), 329–347. <https://doi.org/10.1017/S0140525X12001495>
- Prinz, J. J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. MIT Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Clarendon Press.

- Ruhl, C. (1989). *On monosemy: A study in linguistic semantics*. SUNY Press.
- Thagard, P. (2019). *Brain-mind: From neurons to consciousness and creativity (Treatise on mind and society)*. Oxford University Press.

