A cognitive psychological model of linguistic intuitions: Polysemy and predicate order effects in copredication sentences

Christian Michel [*], Guido Löhr

Vrije Universiteit Amsterdam, Faculty of Humanities - Department of Philosophy, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

Abstract

Copredication is the phenomenon that two or more predicates can jointly apply to the same argument. In "The book is heavy and informative" the word *book* seems to refer to both a concrete physical object that can be heavy and an abstract content that can be informative. It has been observed that if the concrete sense of the nominal is triggered first, the copredication statement often sounds better compared to when the abstract sense is triggered first. However, the cognitive underpinnings of this effect are not well understood. In this theoretically oriented paper, we propose a predictive processing model of order effects aimed at advancing our understanding of that phenomenon. We also connect the debate regarding ordering preferences with an existing strand of research on ordering preferences in multi-adjective strings and the information structure of sentences.
© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Keywords: Predictive processing; Copredication; Order effects; Polysemy; Concreteness ambiguity

1. INTRODUCTION

A linguistic expression is polysemous if it has at least two senses that are intuitively related (Falkum, 2015; Löhr, 2021). Polysemy contrasts with homonymy, i.e., linguistic expressions that have two or more intuitively unrelated senses. The expression 'bank', for example, is a homonym. We can use it to express both the concept of a financial institution and the concept of a riverbank. The two senses are intuitively unrelated because there is little relevant overlap between our beliefs about the distinguishing or diagnostic features of financial institutions and riverbanks. However, 'bank' is arguably also polysemous. We can use the expression to talk about the building of the bank as well as its management, which are clearly related.

The different senses of polysemous words often vary in terms of the degree of abstractness, a phenomenon Löhr (2021) calls "concreteness ambiguity". For instance, 'school' can be used in the following senses: a building, an

* Corresponding author.
E-mail addresses: chris.michel08@gmail.com (C. Michel), loehrg@icloud.com (G. Löhr).

institution, or a group of people (students and teachers). These senses all differ in terms of the degree of concreteness or abstractness. A building is more concrete than an institution. Similarly, the word 'lunch' can denote both the concrete physical food as well as the more abstract process or event of consuming such food. 'Book' can refer to the physical object as well as its abstract informational content.

The phenomenon of concreteness ambiguity is relevant for studying the phenomenon of copredication. The term 'copredication' captures the phenomenon that two or more predicates can sometimes jointly apply to the same argument in a single sentence. This allows speakers to denote two or more distinct kinds of entity (often of different "semantic types" and degrees of concreteness) using a single nominal. To illustrate this phenomenon, consider the following examples:

(1) The manager entered[concrete_building] the bankrupt[abstract_institution] bank.
(2) The heavy[concrete_copy] book is informative[abstract_content].

In the case of (1), a single noun 'bank' is "copredicated" by 'entered' and 'bankrupt'. While the predicate 'entered' is intuitively taken to apply to the concrete building of the bank, the predicate 'bankrupt', in this context, is meant to apply to the more abstract financial institution whose existence conditions are independent of the concrete building that hosts it. In the case of (2), the physical copy of the book is said to be heavy while the predicate 'informative' requires an argument of a more abstract type given that only the content of the book can be informative in this context.

Many linguists and philosophers have argued that the phenomenon of copredication poses a challenge to standard referential semantics (cf. Chomsky, 2000; Collins, 2023; Pietroski, 2018). The challenge is to explain how the meaning of a sentence can be compositional if a single nominal is used to complement at the same time two predicates that require incompatible senses of the nominal. It is currently debated what the phenomenon of copredication tells us about the meaning of words like 'bank' and 'book' and how these terms relate to truth-conditional semantics (e.g., Chomsky, 2000; Gotham, 2017; Liebesman & Magidor, 2019; Pietroski, 2018). However, another important and arguably orthogonal question has only very recently found its way into the literature (cf. Ortega Andrés & Vicente, 2019): What are the cognitive mechanisms underpinning the processing of copredication statements?

One important step towards developing a psychology of copredication is to explain so-called predicate order effects. It has been observed that copredication sentences vary with respect to their acceptability if the order of the predicates is swapped (e.g., Asher, 2011; Asher & Denis, 2005; Brandtner, 2009; Murphy, 2021a; Ortega Andrés, 2020; Ortega Andrés & Vicente, 2019). Those predicate order effects seem to be systematic and related to different degrees of abstractness, thereby directly relying on a word's concreteness ambiguity. Consider the following two examples. (3) comes from Asher, 2011, and (4) from Murphy, 2021b:

(3a) The city has 500.000 inhabitants[concrete_area] and outlawed[abstract_institution] smoking in bars last year.
(3b) ? The city outlawed[abstract_institution] smoking in bars last year and has 500,000 inhabitants[concrete_area]

(4a) The White House is being repainted[concrete_object] and issued[abstract_institution] a statement concerning taxes.
(4b) ? The White House issued[abstract_institution] a statement concerning taxes and is being repainted[concrete_object].

The (a)-sentences have a concrete-abstract ordering of the senses of the nominal, while the (b)-sentences have an abstract-concrete ordering. The (a)-sentences with the concrete-abstract ordering sound more acceptable than the (b)-sentences (see Murphy, 2021a, for empirical evidence of this effect).

In this theoretical paper, we first review and reject a possible explanation of the predicate order effect (Section 2). We then (Section 3) summarize a recently proposed predictive processing (PP) model that we believe helps us understand how intuitions regarding the (in)felicity of copredication sentences are generated. We build on this model but expand it to explain a wider range of order effects (Section 4). The main proposal of this paper is that the order-sensitivity of acceptability intuitions is driven by different types of abstract expectation patterns in a hierarchically structured, probabilistic generative model entertained by the mind. We conclude by pointing to how further empirical research on order effects could be carried out within the PP paradigm.

2. THE CONCRETENESS EFFECT AND ORDER EFFECTS

According to the received psychological account of copredication, the "coactivation package" account, copredication is felicitous if "different sense interpretations are activated simultaneously and can be shifted between without additional processing costs" (Haber & Posio, 2020, p. 115; see also Murphy, 2021b, p. 29). Copredication generally works especially well for what is sometimes called "inherent polysemy", i.e., kinds of polysemy where the different senses are highly

related, in particular, when they are related logically (cf., Pustejovsky, 1995). For example, many things have both an informational content and a physical instantiation (films, books, articles, etc.). For these meanings, understanding the concept almost always requires understanding both the abstract and the concrete sense. However, Haber and Poesio's hypothesis raises the question of what explains the processing differences for copredication sentences that only differ in terms of the order of the predicates. In this section, we discuss one possible explanation of order effects.

Prima facie, the perhaps most obvious explanation for why we appear to prefer a concrete-abstract order might be that concrete concepts are generally easier to process than abstract concepts. Words like 'table' or 'dog' that refer to concrete physical entities that we can touch, hear, or see at least appear to be far easier to process than abstract words like 'democracy' or 'love' (Löhr, 2022). And, indeed, many studies suggest that *word forms* (participants are almost always presented with linguistic expressions - see Löhr, 2022) that are classified as abstract are processed more slowly in lexical decision- (e.g., James, 1975; Strain et al., 1995), recall- (e.g., Jefferies & Lambon Ralph, 2006; Nelson & Schreiber, 1992), comprehension- (e.g., Holmes & Langford, 1976; Schwanenflugel & Shoben, 1983; Hoffman et al., 2018) and production tasks (e.g., Goetz et al., 2007; Wiemer-Hastings & Xu, 2005). If one additionally endorses a "simple first" principle, i.e., we prefer to process first that which is easier to process, we get an explanation for order effects. Preference then boils down to a kind of processing advantage.

The argument that order effects can be explained in terms of a concreteness advantage has been put forth by Murphy (e.g., 2021a, 2021b). Murphy suggests that we prefer the concrete-abstract order because concrete concepts are semantically less complex than abstract concepts. He concludes that what is at work here is, more generally, a *simple-complex* order preference of the parser. Note that it is not the degree of abstraction directly that drives preferences in this view, but the associated semantic complexity. Abstract words are often semantically more complex (so there is some correlation between abstraction and complexity). Murphy (2021a, p. 123) calls this the *Principle of Incremental Semantic Complexity*, which he defines as "the number of cognitive modules, or "core knowledge systems (e.g., number sense, agents and their actions, basic geometric reasoning; Carey, 2009), a given sense draws upon." For instance, senses or concepts denoting physical objects have a low complexity because they only require the cognitive module responsible for conceptualizing geometrical and spatial information. Senses denoting an institution require additional modules like a theory of mind module. In general, the order preference for senses proposed by Murphy is: physical < information < event < institution (Murphy 2021a, p. 55, 2021b, p.23). This would not only explain the order effect but possibly even a concreteness advantage more generally. Let us now delve deeper into the proposal that abstract concepts are more difficult to process and hence drive order effects.

One problem with reducing order effects to a concreteness effect is that the assumption that abstract concepts are generally more difficult to process is far from obvious (Löhr, 2023). In recent years, a more differentiated picture of the concreteness advantage has emerged. It has been shown, for example, that concreteness effects tend to disappear if the sentences are presented in an appropriate, "supportive" sentence context[1] (cf. Schwanenflugel & Shoben, 1983; Schwanenflugel & Stowe, 1989). So, while it may be that abstract concepts are more difficult to process than concrete concepts when presented out of context, this effect seems to disappear when they are presented as part of a sentence that provides a context that makes the abstract word predictable. More recently, Kousta et al. (2011) found that the concreteness effect may be reversed once word stimuli are controlled for various psycholinguistic variables. Specifically, abstract words with affective associations are processed faster than neutral concrete words. So, some concreteness effects that have been found might have been the result of poorly controlled studies. Therefore, we cannot take for granted that a straightforward, robust concreteness effect exists, and it is safe to say that the sentence context and other factors (like affective associations) are important factors that need to be controlled for.

Another consideration against the claim that the more abstract a concept is the more effortful it is to process is the following. Some scholars suggest that children can grasp more general features before specific instances are presented (e.g., Keil, 2021; Kemp, Perfors, & Tenenbaum, 2007). It is also well established that children do not always find more concrete concepts easier to acquire (e.g., Rosh, 1973; Lakoff, 2008). They prefer, for instance, 'dog' over the more concrete or specific 'terrier'. In this case, the more abstract concept of dog (versus terrier) seems not to be correlated with greater processing difficulty (see also Keil, 2021). Furthermore, the perception of complex objects or situations often involves first grasping an overall pattern before processing the concrete details (e.g., Barrett & Bar, 2009; Fillmore, 1975). Take an analogy from visual processing to create some intuition pump. There is no doubt that we are good at recognizing faces. However, the recognition of a face is a very complex task not least because we can do so under

---

[1] For instance, in the sentence "The final word of this sentence is *market*," "market" is embedded in a *neutral* context and in "The man bought two live lobsters at the seafood *market*," "market" is embedded in a *supportive* context (Schwanenflugel & Stower, 1989, p. 119).

many different conditions (light conditions, perspective, etc.). Faces are represented in the brain in higher-level neural layers that respond to more abstract and complex patterns (see, e.g., Freiwald & Tsao, 2010; Leibo et al., 2011; Riesenhuber & Poggio, 1999). The face of a specific person is a relatively abstract representation that captures invariances of all sorts and very different pixel patterns are categorized as the face of a specific person.

Note also that much of the idea that abstract concepts are more difficult to process depends on an underlying cognitive computational paradigm that one might endorse (see below). In particular, for models of cognition following hierarchical Bayesian approaches – which are getting more and more popular and influential (e.g., Tenenbaum et al., 2011) – there is no reason to assume that abstract concepts are more effortful to process. In those models, the learning of concepts and world knowledge amounts to learning, representing, and inferring concepts on many levels of abstraction in a hierarchical generative model. To bootstrap a model of the world, representations on each level must be constrained by more abstract representations on higher levels (so-called "hyper-priors"). Furthermore, inference happens on many levels in the hierarchical model at the same time.

Another interesting example is metaphor processing. According to one of the most supported psycholinguistic models of metaphor comprehension, a metaphor arises as an ad hoc abstraction process (Holyoak & Stamenković, 2018; Rodriquez Roderos, 2021). Combined with the finding that metaphorical processing can be as fast as literal processing (e.g., Gibbs, 1984; Glucksberg, 2003), this amounts to evidence that the processing of more abstract representations is not always more effortful. Thus, one might argue —quite to the contrary—that the processing of abstract concepts is less effortful because they are less detailed representations. The reason is that concrete concepts are often represented as detail-rich exemplars (e.g., Machery, 2009). In a recently proposed model of concepts that pulls together the current state of the neuroscience of concepts, the tokening of a more concrete concept often implies the activation of a much larger network of neuron assemblies (Dove, 2022; Michel, 2022) including in modality specific areas. Hence it might—on the contrary—be more effortful to represent and process concrete concepts in many instances.

If it is not the processing difficulty or effort of a concept that explains order effects, an alternative hypothesis might be that what is relevant are certain switching costs. It might be easier to *switch* from concrete to abstract compared to switching from abstract to concrete. However, empirical research on polysemy casts doubts on this explanation as well. One of the main findings in this literature is that the switch from one sense to another (unlike from one meaning to another in the case of homonymy) does not seem to coincide with an increased processing effort – whether or not these senses are abstract or concrete (Frazier & Rayner, 1990; Klepousniotou, et al., 2012; Eddington & Tokowicz, 2015; Quilty-Dunn, 2021).

Finally, a consideration that also speaks against the idea that abstract senses are per se more difficult to process are exceptions to the concrete-abstract bias, and, given appropriate contexts and contents, the bias does not manifest itself. Consider the following examples (both discussed in Murphy, 2021a; see also Asher & Denis, 2005, for (5a)):

(5a) The reproduction (of the painting) took place$^{abstract\_event}$ in that workshop and [was/is] eight feet tall$^{concrete\_painting}$.
(5b) The reproduction (of the painting) [was/is] eight feet tall$^{concrete\_painting}$ and took place$^{abstract\_event}$ in that workshop.
(6a) The entrance was being used to deliver food$^{abstract\_function}$ and was in need of repair$^{concrete\_object}$.
(6b) The entrance was in need of repair$^{concrete\_object}$ and was being used to deliver food$^{abstract\_function}$.

Sentences (5a) and (6a), which have an abstract-concrete ordering, are generally judged to be more acceptable than the corresponding variant sentences (5b) and (6b), with a concrete-abstract ordering.[2] It seems that if concrete concepts really were preferable in copredication sentences, we would likely not expect these exceptions unless, of course, one could show why the concreteness effect is overwritten in those cases.

3. A MODEL FOR LINGUISTIC INTUITIONS ABOUT COPREDICATION STATEMENTS

In what follows we provide a novel cognitive account for the observed concrete-abstract ordering preferences that does not rely on any concreteness effect (nor does it consider semantic complexity/processing effort per se as a factor that neuro-mechanistically explains order effects). Specifically, we want to provide an account of the circumstances under which the apparent concrete-abstract bias is overridden, i.e., an account of the *context-dependence* of ordering preferences. Our account also connects with research on the preference of the order of adjectives in multi-adjective strings; specifically, we build on the work of Scontras and colleagues (e.g., Scontras et al., 2017, 2019; Scontras, 2023). For that purpose, we argue that what drives acceptability intuitions about copredication sentences is the degree

---

[2] According to Murphy's study (2021a), (6b) has a slightly higher acceptance rating than (6a) (see p. 296).

to which certain background expectations are fulfilled or violated. Linguistic intuitions arise, so our proposal says, by a process of prediction error minimization in a complex hierarchy of expectations ("priors" as we will call them), which constitutes a generative model of the world, also including knowledge of language. We then apply this view specifically to copredication order effects. In the following, we first introduce the predictive processing (PP) framework, which is increasingly popular in cognitive science, and flesh out how it can model linguistic intuitions, specifically context-dependent ones.[3] In Section 4, we provide a novel account of order effects based on this PP model.

3.1. Predictive processing – A brief summary

Predictive processing (PP) is a relatively recent but very influential neuroscientific cognitive paradigm (see e.g., Clark, 2016, 2023; Hohwy, 2013; Sprevak, 2021a, 2021b, 2021c, 2021d, for more detailed introductions to the framework). According to PP, the brain tries to constantly anticipate or reconstruct its sensory input and minimize prediction errors. The brain is, therefore, highly proactive, not a passive processor of sensory input. Prediction error minimization is the single driving principle for brain activity. Some scholars (e.g., Friston, 2010; Hohwy, 2013) even claim that PP has the potential to be a great unified theory of brain, mind, and agency in uncertain environments.

The prediction error minimization principle has been motivated in different complementary ways. The first is by considerations of the efficiency of metabolically costly cognitive processing (evolution made it that we have energy efficient brains). It is more efficient to process an *error* signal than fully-fledged signals (very much like video compression: one only needs to store what is unexpected, i.e., what differs from frame to frame). The second is by considerations that have to do with the fact that we can learn (language, concepts, the mother's face, etc.) from very few examples (contrary to classical deep neural networks). To achieve this efficiency, we need an "inductive bias", i.e., some pre-existing knowledge/hypotheses/guesses that can themselves be learned or are evolutionarily endowed (e.g., Gershman, 2021). But inductive biases are nothing other than expectations/priors. The third is by theoretical biology (e.g., Friston, 2010): to survive, our body needs to keep its own and environmental parameters within a corridor of viable bounds. This implies that it needs to minimize "surprises", i.e., events and situations need to evolve as "expected" according to an internal model of the world.

All those considerations point to a view of the brain as a hierarchical prediction machine. *Hierarchical* means that we have a prediction cascade from general abstract patterns down to more concrete perceptual representations. Hierarchical models make brains very powerful because they allow for representing patterns on different scales of abstraction. Hence, they allow us to represent high level tacit principles, such as "hinge-beliefs", which we take for granted, like that there is an external world, complex patterns (our mother's face), situations, and causal inductive biases, but also representations along the lines of Gricean's conversational principles, specific perceptual patterns and so forth. All representations are predictions or expectations (or "priors", as we call them) for more concrete representations on a lower level. The brain's prediction minimization efforts approximate Bayesian inference. Bayesian approaches are increasingly popular as approaches to cognition. They build on Bayes' theorem, which describes how subjective "beliefs", i.e., states that represent the world in a way the agent relies on for further personal and sub-personal level cognition and action, should be updated based on new evidence.

Importantly, expectations of incoming stimuli can have different strengths This is realized as *precision weighting* of prediction errors in the PP framework (e.g., Clark, 2016, 2023). The more reliable an expectation is, the more it can force other conflicting expectations to be suppressed. Sometimes higher expectations are overwritten by reliable sensory evidence that is estimated by the system to be more accurate. Each situation involves expectations on multiple levels at the same time. Contextual representations also form part of that expectation hierarchy, which needs to establish an equilibrium with minimum conflict among all of them.

For instance, imagine you have been invited by a friend to a concert hall. Due to lack of interest in the concert itself (not in the company of your friend), you have not looked at the program. You are entering the audience hall. This situation is now represented as a contextual situational representation in your brain on a high level in the hierarchy which implies many other expectations. You are very surprised to suddenly see a clown on the stage. Your brain generates a large prediction error signal and tries to "explain" the scene, i.e., to adjust the expectation network. Subconsciously, you might entertain various hypotheses about the reliability of your (tacit) assumptions. Maybe my visual input is not reliable (am I drunk, or did I not look properly?) Or maybe I am confused and wrong to think I am at a classical concert (and

---

[3] We will not provide a detailed justification why the PP framework should be adopted in the first place, and we will not defend it against other frameworks here. Our motivation for applying the PP framework is that it is an independently motivated, increasingly popular framework with explanatory successes in many domains. Obviously, there are also detractors (e.g., Litwin & Milkowksi, 2020) and PP should still be seen as an emerging framework under development.

instead I am at a circus)? If all those hypotheses are discarded in a split second by your brain because of weak evidence, it now needs to adjust other expectations to get a coherent expectation network for the current situation. It turns out that the concert is a *contemporary* one and the musicians are dressed in costumes. Once you realize this, your expectation network settles down and re-establishes an equilibrium.

Expectations are also pervasive specifically when processing language. For instance, when you are reading, letters can prime (make expect) words and the other way round (e.g., Wheeler, 1970). Another example is that you don't expect sentences like "Can you change my tires" at a restaurant. Or, as another example, you can often easily finish the sentences of others, because you are able to anticipate what the other will say (e.g., Levinson, 2016). Empirical and theoretical support is mounting that indeed, predictive mechanisms play a crucial role in language processing (e.g., Heilbron et al., 2021; Pickering & Garrod, 2013; Pickering & Gambi, 2018).

For an interesting example of approaches to linguistics in the vicinity of PP consider so-called surprisal models. Hale (2001) proposed that a probabilistic version of the so-called Earley algorithm accounts for different linguistic phenomena like reading time predictions for garden path sentences. In surprise models, processing difficulty is accounted for by surprisal (inverse to the probability of encountering the next word). Levy (2008) follows Hale and proposes an "expectation-based syntactic comprehension" model to explain syntactic ambiguity resolution. According to this model, sentence comprehension proceeds incrementally. Evidence is being integrated step by rep and the distribution of preferred interpretation updated. Hale puts forward the "linking hypothesis" that processing difficulties (measured in reading time) are related to the surprisal of new information.

Surprise models—though they use probabilistic prediction as a paradigm—are quite limited.[4] For instance, Zarcone et al. (2016) criticizes surprise models, as they overlook attention and relevance, and endorse a predictive processing[5] approach in an explicit way. Leveraging the hierarchical structure, Zarcone at al. suggest that comprehenders have context-dependent expectations at various levels of abstraction. However, the authors focus on disentangling the notion of "linguistic salience", which—so they suggest—is being used in two senses, which can be cashed out nicely within a predictive processing approach. One notion of salience is what is top down available (which drives further expectation), and the other is where attention needs to focus on bottom-up sensory input. What is at work here is, of course, the precision weighting apparatus, which is in essence an attention mechanism orchestrating the balance of influence of previous knowledge versus new incoming sensory input.

It is important to stress that while it is generally recognized that prediction is—almost obviously—at work in language processing, not all accounts, like Hale's (2001) or Levy's (2008), are "PP accounts". The PP account of language processing goes much further, as Zarcone et al. have already suggested. Details do not matter here, but PP provides neuro-mechanistic details on an implementational level (e.g., Sprevak 2021d) and posits the representational structure in the form of a huge hierarchy with an abstraction gradient, just to mention some of the key commitments of PP that go beyond merely endorsing surprisal or prediction. In sum, the key idea we need from PP for the current purposes is that of an expectation hierarchy that needs to be brought into an equilibrium free of conflicting expectations (i.e., free of prediction errors). Those expectations span many levels of abstraction (and many are sub-personal and unconscious) and all of them operate simultaneously.

With this very brief sketch of the PP view of cognition and language processing in hand,[6] let's turn to an account of linguistic intuitions.

3.2. Linguistic intuitions about copredication sentences driven by a hierarchical web of expectations

According to predictive processing approaches to language processing, linguistic intuitions are driven by a system of dynamic cognitive expectations (e.g., Löhr & Michel, 2022). Our world knowledge (including knowledge of language) boils down to a large and complex expectation hierarchy constituted by "priors". Representations on one level constrain representations on lower levels and hence function as prior probabilities (e.g., Hohwy, 2013; Clark, 2016). The priors are dynamic, not fixed, in the sense that they are constantly updated—approximately by Bayesian inference—to posteriors which then function as new priors. In this model, plausibly, acceptability (or felicitousness) judgements are driven by intuitions that are determined by expectations on different levels of abstraction. Intuitions of oddness, in the PP approach, are understood as arising from monitoring prediction errors.

---

[4] A more detailed evaluation would exceed the scope of this paper but note that Hale (2001) explicitly endorses the Chomskian generative grammar paradigm, which might not be the most natural partner for PP (see footnote 8). Also, Earley algorithms only cover context-free grammars, which are a highly idealized model for natural language.

[5] They use the term "predictive *coding.*"

[6] We defer to the literature for further details about PP (e.g., Clark, 2016, Hohwy, 2013, Sprevak 2021a, 2021b, 2021c, 2021d).

A sentence (including a copredication sentence) is comprehended if we can"generate"[7] it from our model, such that it matches the bottom-up input of the perceived sentence. Generating a sentence from the mental model means mobilizing those priors on different levels that correspond to the variables that bring about the sentence causally. This implies that a selection of appropriate priors across the hierarchy form an error minimizing equilibrium. To give a toy example: the sentence "The apple is on the table" is comprehended if the model has generated this sentence using the priors it has available. Let's say on some higher (more abstract and complex) level there is a prior corresponding to a X_is_on_Y-schema,[8] on a lower (a less abstract level) an APPLE and TABLE prior. By filling in the slots X and Y (both of which are abstract concepts in themselves in the form of categories defining the class of instances that can be filled in), we generate a prior hierarchy that gives us the sentence meaning for "The apple is on the table". The brain's model is now in a situation where the visual system produces, bottom-up, the sentence "The apple is on the table", and the top-down prediction in the form of the sentence meaning matches the bottom-up input without any prediction error. In the case of the slightly erroneous sentence "The apple on table" the system is still able to comprehend the sentence. However, the best possible sentence the model can bring about is the correct "The apple is on the table". So there remains a permanent prediction error, which manifests itself in the form of a judgment of diminished acceptability i.e., a slight intuition of oddness.

In sum, linguistic intuitions about the acceptability and oddness of sentences in general (and copredication specifically) arise from violations of expectations, or in PP parlance, from prediction errors that cannot be suppressed by properly adjusting the expectations (for different reasons). Prediction errors are minimized on all levels in the hierarchy, however. Some errors on higher levels in the hierarchy are "brought to the attention" of the agent, such that clarifying action can be triggered. In the case of linguistic oddness, expectation errors signal a potential misunderstanding in communication. If the hearer's model cannot settle the errors by itself, the hearer is triggered by a feeling of oddness to "collect more evidence", for instance, by issuing a clarifying question to the speaker, or to reinterpret the sentence flexibly and charitably (by "overlooking" errors).

Intuitions about the felicitousness of copredication are then driven by our individual expectation networks which constitutes the language-model (as a submodel of the overall world-model). An analysis of felicitousness judgements should be carried out as an analysis of expectations on different levels of abstraction. Löhr & Michel (2022) have proposed a general model of copredication based on the predictive processing paradigm. It rests on the idea that words are structured as "information packages" precisely in the form of expectation hierarchies. A toy example for an expectation hierarchy is one with [BOOK] as a root node and [E-BOOK] and [PAPER-BOOK] as child nodes. [PAPER-BOOK] might in turn have features like [has weight], etc. as child nodes. When reading "book" we have (sub-personal) expectations about e-book and paper-book senses. When the paper-book sense is active (for instance because you are at a book fair), we have expectations about the features of paper-books, etc. The copredication sentence "The book is heavy and informative" is then felicitous, according to this account because we process the sentence only with the more abstract [BOOK] sense in mind, which allows us to expect both more specific senses (the physical object and information senses) evoked by the predicates. "The book is heavy and salty" is unexpected and sounds odd because the predicate *salty* is normally not associated with any of the predicates represented in the expectation package of 'book'.

Those word-associated expectation hierarchies in turn are embedded in different higher order structures of expectations or situational priors. For instance, if you are at an e-book fair, this situation is represented as a higher level situational prior, which primes the e-book sense when hearing "book" in a conversation. Those expectations external to the information package can explain why certain pairs of very similar sentences and in which the same two senses are invoked in the same order have different degrees of felicitousness. Consider the following sentences:

(10a) The newspaper has been attacked by the opposition[abstract_institution] and publicly burned by demonstrators[concrete_object].
(10b)?The newspaper was attacked by the opposition[abstract_institution] and fell off the table[concrete_object].

---

[7] "Prediction" in PP is always meant in this more general sense of "generation" (or "sampling") by a dynamically updating generative model, not the narrow sense of guessing some future event or state. See Pickering & Garrod (2013) for a model very much in line with PP, in which the comprehension and production system are identical. The underlying idea of a generative model used in language comprehension aligns with the "analysis by synthesis" principle (e.g., Neisser, 1967).
[8] Note that such schemas are reminiscent of "constructions" in construction grammar (or cognitive grammar) (e.g., Langacker, 1987; Goldberg, 2013). Indeed, as Michel (2023) argues, the PP paradigm can be seen as a cognitive computational underpinning for construction grammar, so PP is closer to construction grammar than to the Chomskian generative grammar paradigm. To elaborate on this would go beyond the scope of this paper, but we suggest that constructions are nothing other than expectations/priors.

(10) is an example of a specific case of such an asymmetry effect that can be found in Löhr & Michel (2022). They argue that (10a) is more felicitous because there is a higher-level prior (expectation), in the form of a familiar situation in which objects are burned in a protest directed against what those objects represent. This situational prior allows us to expect, when reading/hearing "newspaper", both the institution as well as the physical object sense. This prior, in the form of a metonymic construction, is not available in (10b). We will rely on this model to explain order effects in the next section.

4. MODELLING SENSE ORDER EFFECTS IN COPREDICATIONS

We suggest that there are a host of priors/expectations outside the nominal's package that have an influence on expected orderings of predicates. Those are represented as higher level priors and can explain the order-related asymmetries of acceptability judgements. The suggested account can accommodate the observed concrete-abstract bias and exceptions and spell out when the concrete-abstract order bias is overridden.

4.1. Explaining the concrete-abstract order bias

We suggest explaining the concrete–abstract preference observed by Murphy (2021a) by reference to insights from ordering preferences in multi-adjective strings (especially Scontras et al., 2017, 2019) and from information structure theory (e.g., Haviland & Clark, 1974).[9] The proposal is that in the absence of a rich context and sentence content— as is the case in most experiments by Murphy— by default we prefer an order from more objective to more subjective senses because this is the more reliable way to create common ground (the "Given" information) with the interlocutor. In richer contexts, and sentences with sufficient cues and information, this bias might easily be overridden. This idea follows Scontras et al.'s (e.g., 2017, 2019) explanation of the robust and cross-linguistic preference for a certain order of adjectives in multi-adjective strings combined with a noun. To use one of their extreme examples, while "the good smooth purple plastic chair" sounds ok, "the plastic good purple smooth chair" sounds clearly unacceptable. When classifying types of adjectives (adjectives describing material properties, colour, etc.), one can find the following preferred order: material > colour > shape > physical > age > value dimension (Scontras et al., 2017, p.56).[10]

How to explain this order preference has been the subject of debate and Scontras (2023) provides a very useful overview and discusses several predictors that have been proposed in adjective ordering research. Among them is the closeness of the adjective meaning to the meaning of the noun, the degree of meaning specificity, inherentness, definiteness of denotation, subjectivity, derivational syntactic proximity, noun likeness, frequency, length, lexical class, information gain (uncertainty reduction based in available evidence), pointwise mutual information (co-occurrence frequency), collocation frequency, and referential discrimination. The breadth of proposals is certainly impressive, and the fact that, as Scontras points out, there is "no winner" yet is intriguing. The specific phenomenon of ordering preferences seems complex and multifaceted, and is still not fully understood, even after over a century of research (it will become clear later that under the PP approach this somewhat "messy" picture and mixed evidence is to be expected). There are at least three complicating factors that we can identify. Firstly, order effects are heavily *context-dependent*. As Scontras (2023) points out, Ney (1983, p. 100) observed that "almost any order of adjectives seems to be possible depending on the intended meaning of the speaker or the situation in which the speaker frames an utterance." (cited in Scontras, 2023, p. 365). Secondly, there are subtle *cross-linguistic variations of the strength* of order effects. For instance, it matters whether a language has post- or pre-nominal adjectives. Rosales & Scontras (2019) found that order effects are weaker in Spanish then in English. Trainin & Shetreet (2021) found that in Hebrew order effects exist as well, but they are weaker than in English. Thirdly, different *factors might interact* in complex ways. For instance, Trainin & Shetreet (2023) have shown how order effects are weakened in Hebrew if lexical factors, like frequency/accessibility or length, interact.

---

[9] There are other important traditions in linguistics that deal with the grammar of adjectives and their ordering (see, e.g., Valois, 2006; Cinque, 2010) and which are located within the Chomskian generative grammar paradigm. We do not aim to propose a new theory of adjective orderings here (or discuss problems identified in this tradition), nor do we have space to discuss those existing traditions and how they relate to the here proposed account. We focus on predicate ordering and merely point to adjective order effects as an inspiration. Specifically, we do not want to argue at this stage that we can transfer our proposal related to predicate orderings to all phenomena related to the grammar of adjectives (and adverbs).

[10] This order is not exceptionless, of course. For instance, superlative adjectives or long adjectives do not follow this order.

We build on Scontras et al.'s approach based on the degree of subjectivity of the adjective, given that—compared to other accounts—it explains most robustly order effects,[11] and also seems very attractive from a conceptual point of view. Scontras et al.'s approach very plausibly starts from the question of what adjectives (or determiners in general) are doing for us. One function, they argue, is that they help us to determine what the intended referent is. The proposal of Scontras et al. (2019) is then that an order is preferred where adjectives of more objective properties are closer to the noun because they allow the interlocutors to align the referent described in more objective ways:

> More subjective adjectives introduce the potential for errors in alignment, as speakers and listeners might (faultlessly) disagree about category boundaries. When it comes to ordering preferences, speakers consolidate the less subjective, more useful content around the modified noun. The claim is that they do so in an attempt to aid the listener in establishing reference by minimizing errors in alignment. (Scontras et al., 2019, p. 10).

The crucial role of proper referent determination for communicative purposes has also recently been emphasized by Pickering & Garrod (2021) in their account of dialogue that builds on a *shared workspace* where the speaker and hearer need to align referents. So, we find Scontras et al.'s approach very well motivated.

The idea that order preference has to do with reliable referent determination can be applied to copredication. It would be surprising if Scontras et al.'s insight[12] had no bearing on the topic, especially given that some copredications have the form ["the" Adj1 Adj2 Noun], ["the" Adj1 "and" Adj2 Noun] or ["The" Noun "is" Adj1 "and" Adj2]. Consider:

(11a) The long^{process} and boring ^{content} lecture.
(11b) The boring ^{content} and long^{process} lecture.
(11c) The lecture is long^{process} and boring ^{content}.
(11d) The lecture is boring ^{content} and long^{process}.

Following Scontras et al. (2017), the adjective ordering preference is expected to be long > boring, because a dimensional adjective ("long") scores higher in the subjectivity scale than a value adjective ("boring"). Indeed (11a) and (11c) sound more acceptable.[13]

Note that there might be a complication here due to our move from an adjective string without "and"-conjunction to one with such a conjunction (11a and 11b) and finally to the copredications (11c and 11d), which also have a conjunction. Rosales & Scontras (2019) found that conjunctions can neutralize adjective ordering preferences in languages like Spanish where an adjective can be chained *only* by a conjunction. Conjunctions interrupt hierarchical semantic composition (i.e., the adjective closer to the noun composes first with the noun). Rosales & Scontras could not detect significant ordering preferences in Spanish (which is also evidence for a hierarchical rather than serial composition of adjectives). They observe that, interestingly, in Indonesian, where conjunctions are optional, there *are* ordering preferences. How can we explain this? They suggest, plausibly, that "The regularity present in the conjunction-free cases bleeds over into the cases with conjunction" (Rosales & Scontras, 2019, p.10). We can flesh out this metaphorical explanation in more concrete terms within the PP framework. The order preference—acquired by the exposition to many nonconjunctive strings—is a prior that exerts its influence over cases with a conjunction. So, the order effects depend on what priors you actually have built into your language-model. The advantage of PP is that it can model this crosslinguistic diversity.

Consider another example discussed in Murphy (2021a).

(12a) John said that the *small$^{physical\_object}$* and *interesting$^{abstract\_content}$* book was on the table.
(12b) John said that the *interesting$^{abstract\_content}$* and *small $^{physical\_object}$* book was on the table.

(12a) is more acceptable than (12b). This can again be explained by the subjective-objective order preference.

---

[11] According to Scontras et al. (2017) subjectivity accounts for 85% of the variance in the ordering preferences for 26 adjectives from 7 semantic classes and for 61% for a larger group of 74 adjectives from 13 semantic classes.

[12] Or, more generally, it would be surprising that research on adjective ordering had no bearing on the topic of ordering preferences in copredication.

[13] Given the theoretical nature of this paper, we have not conducted a statistical, empirical study to support claims about acceptability judgments for some example sentences for which no acceptability data is available in the literature. However, for those examples we have carried out an informal and preliminary survey among English speakers that support those intuitions. Rather than closing the empirical file here, one of the main objectives of this paper is to inspire new empirical research that is focused on context effects and the proposed "expectation architecture."

The important point is that without much discourse context, i.e., for isolated sentences, we want to establish a reliable referent by using the more objective adjective. In the PP model, this is a default information structure in context- and content-poor situations. Subjectivity is plausibly correlated with abstraction, but what better explains the ordering bias in context-poor situations are expectations related to the *information structure*, i.e., to the question of what is given and what is new information. The role of determiners is to contribute to reference fixing and create a common ground before then moving to the new information.

We suggest that this idea can be generalized beyond the narrow phenomenon of adjective ordering to predicate ordering. A concrete–abstract sense order is preferred in context-poor conditions because something more concrete is better suited to creating common ground than something abstract. Once we have a richer context, which contributes to determining the common ground versus the new information, the default preferred order concrete-abstract might be easily overwritten.

Consider now a content-wise "enriched" version. The enrichment consists here of the qualifier 'incredibly', which can be interpreted as a cue for the new (interesting, surprising, etc.) information to which the hearer's attention should be drawn.

(11b) The boring [content] and long [physical_object] lecture.
(13) The boring [content] and **incredibly** long [physical_object] lecture.

We predict that (13) is now more acceptable than (11b). "Incredibly long" seems to be the interesting or new information and we expect it to be in second place.

There is a relation with metonymies here. Consider:

(14a) The White House is being repainted[concrete_object] and issued[abstract_institution] a statement concerning taxes.
(14b)?The White House issued[abstract_institution] a statement concerning taxes and is being repainted[concrete_object].

Here we again have a concrete–abstract order. We prefer to first single out the physical object (building) as a referent and then metonymically move to the more abstract sense (institution). In a way, this is a "double metonymy", as the White House is literally a building, and has this as its physical sense, but also the institutional sense. The objective-subjective order preference has a metonymic character as we want to establish a common ground that allows us to talk about more abstract/subjective properties. It is that *once* a concrete thing has been evoked, we have a default metonymic expectation for a more abstract sense. Facilitation of reference alignment in dialogue might also more generally explain the pervasiveness of all sorts of metonymies in language[14] (for an overview of the impressive variety of metonymies see, e.g., Littlemore, 2015). Note that in the case of (14) there might be an additional effect[15] at work (operating in the same direction) which makes (14b) especially infelicitous. This effect is related to expectation about the information structure of sentences. Repainting the White House is not especially interesting or important, while information about taxes is normally non-trivial and relevant. Therefore, the hearer wonders why the speaker mentions a trivial (though maybe a so far unknown) fact after an interesting one—see also Section 4.2 b.

The larger picture of the PP model is that the mind constantly tries to infer meaningful messages on a situational level, which are nothing more than complex and abstract higher-level patterns represented as priors in higher levels of the model hierarchy. Deriving situational patterns requires bringing under certain higher-level priors (i.e., to categorize) the pieces of information in the sentence to minimize violations of expectations. That the Given should come before the New is certainly an expectation that we reasonably have represented in our mind as a higher-level pattern/prior. Also, it is plausible that, given very little context and information, the most reasonable default assumption is that the Given is the more concrete (i.e., we have a "metonymic expectation" that manifests itself as a concrete-abstract order preference).

While we have argued that it is not necessarily more difficult to *process* abstract concepts when reading or listening (this is especially true for the PP model), it is often easier to *communicate* (and in consequence to *learn* them) by means of concrete/simple concepts because reference alignment is easier. The quite trivial reason is that concrete concepts correspond to priors that are closer to the sensory periphery, and we can then use ostension or perceivable properties to which we can direct joint attention to establish common ground involving those concepts. In a way, language as a

---

[14] Consider the classical metonymy in a restaurant setting: "The ham sandwich asked for the bill". Here, we suggest we see highly efficient "good enough processing" (Ferreira, et al., 2002) at work for referent determination in dialogue. Such a handy metonymy allows the hearer to understand perfectly well that "the ham sandwich" stands in for "the person who ordered the ham sandwich."
[15] Thanks to an anonymous reviewer for pointing this out.

social tool for communication works by channelling the speaker's intentions (that are represented high up in the hierarchy) in the form of cues through her sensory-motor periphery out into the world and back through the listener's sensorimotor periphery into higher levels of her mental model. Therefore, the concrete (i.e., sensorimotor related) is more privileged in communication, namely as a stepping-stone towards more abstract content.

4.2. Overriding the default concrete-abstract order bias

We suggest that multiple types of expectations, represented as hyper-priors in the mind's hierarchical PP model, can override the default concrete-abstract bias. We suggest that this bias arises only in sentences lacking a rich context and content. The concrete–abstract bias is a sort of default strategy to "keep going" in those cases. In other words, in the absence of a contextual hyper-prior, the current ("default") prior is not being updated to a new prior.[16]

The following list might not be exhaustive;[17] our aim is merely to make a start in providing a more comprehensive account of order effects of copredications and hopefully capture here at least the main factors that can explain most of the material discussed in the current literature. One might also be skeptical about whether a full systematic account can be given, as the mind might not represent expectations, biases, and associations in a very systematic way. Still, of course, we believe that many interesting generalizations can be made.

a) *Order-sensitive coherence expectations*

One sort of expectation is, of course, order-sensitive coherence expectations on the level of sentence (see also Murphy, 2021a). The coherence types relevant for order effects are causal and temporal coherence.[18] An example for (not very interesting) order effects of this type are: "She read and threw the newspaper" versus "She threw and read the newspaper".

b) *Expectations about the information structure*

However, there are other types of expectations, related to the "narrative", i.e., expectations about the content of the sentence and its context that do not fall under "coherence" in the more restricted sense. We suggest that those aspects are closely related to the information structure, which creates relevant expectations that influence order effects. This means that we have expectations, e.g., about the order of given and presupposed common ground information and new information, namely that the former should normally precede the latter. Those expectations are influenced by the many cues (e.g., grammatical and semantical) about information structure we can find in a sentence, but also by the context. We suggest that, as pointed out in the discussion of example (14), expectation about the information structure not only refers to the *novelty* but also to the *relevance* and *salience* of information.

c) *Formal structural primes*

Furthermore, there are other expectations created by more formal structural primes on sentence structure, stemming from an immediately previously processed sentence. In those cases, as we will see in a moment, we expect that the order in the current sentence is mirrored in some form.

Let us now illustrate those three expectation types with examples.

---

[16] We thank an anonymous reviewer for this explanation of the default interpretation.

[17] We thank an anonymous reviewer for pointing to one interesting phenomenon, too complex to cover here in detail, namely *heavy constituent shifting* (see, e.g., Wasow & Arnold, 2003, for an overview). In English, often, a syntactically complex noun phrase changes its canonical position like in "I gave Mary the valuable book that was extremely difficult to find" (which sounds more acceptable than "I gave the valuable book that was extremely difficult to find to Mary"). An interesting observation is that this shifting cannot be straightforwardly extrapolated to other languages (e.g., Japanese manifests a preference for placing first the heavy NP; see, e.g., Hawkins, 1994, Chapter 3). In general, other factors beyond syntactic complexity, like information structure, or language specific parsing peculiarities (e.g., left- versus right-branching) plausibly determine constituent ordering preferences. We believe that PP, has the potential to deal with this phenomenon, as it is associated more naturally with a use-based approach to language that can account for cross-linguistic diversity by positing variable complex systems of priors/expectations.

[18] Murphy (2021b, p.19) also mentions "extensional overlap" (i.e., featural commonalities between both predicates) as a type of coherence relation. While extensional overlap indeed influences acceptability judgements, it does not bear directly on explaining *order* effects, i.e., the asymmetry of judgements depending on the predicate order.

a) *Expectations related to order-sensitive coherence*

Consider the following example form Murphy (2021a), which turns out to be an exception to the concrete–abstract preference.

(15a) The entrance was in need of repair[object] and was being used to deliver food[function].
(15b) The entrance was being used to deliver food[function] and was in need of repair[object].

According to Murphy's data, (15b) with an abstract-concrete order is preferred over (15a).[19] We suggest that this is the case, because the first predication in (15b) sets up the context in which the functionality of the door is being made salient. This then motivates the need for repair because we can immediately infer that without such a repair the door's important food delivery functionality would be disrupted. This interpretation makes the sentence "coherent". Sentence (b) is structured as a simple "and"-conjunction and the coherence relation between the two predications is not explicit. The mind tries to make sense of the statement. In this case it can easily infer that "and" can be read elliptically as "... and **therefore**...".

(15a) contains the same pieces of information in a different order. The acceptability gap is not large but still significant. Here, "and" cannot be easily read elliptically. Rather, to make the sentence more explicitly coherent, we would need to change "and" for "because it ..." (or something like "... and this repair is important because the door was being used for food delivery").

A concern might arise here. (15a) and (15b) both contain the same information, and one could argue that both can lead to the same interpretation and there should be no difference in their acceptability. This means that, e.g., causal coherence per se of the two predicates cannot be the driver for the order asymmetry. An expectation and prediction error minimization approach can accommodate better the more subtle effects at work. In the case of a causal relation, if you have a cause (C) and an effect (E) we can accommodate both orderings C-E and E-C, namely: "C causes E" and "E is caused" by C by the above alluded method of "inferential repair." However, note that sentence (15a) requires a somewhat larger amount of inferential repair than (15b). This implies that (15a) has a higher prediction error between the perceived sentence (which corresponds to the bottom-up sensory input) and the repaired version (inferred interpretation). In fact, the very word "and" easily makes us expect "therefore" as "and therefore" is a frequently used expression. In (15a) one needs to delete "and" and add "because it." The higher prediction error manifests itself as a slightly lower degree of acceptability.

The idea here is then that there is a cost of repair, spelled out in PP terms as a prediction error, which drives acceptability judgements. The advantage of this way of looking at the phenomenon is that it promises to have a large scope of application. For instance, the idea can also be applied in the case of *zeugmatic* copredication-type sentences with order effects. Those are types of sentences with the form Predicate-Object 1-Object 2, i.e., a *verb* is used in two senses, mostly a literal and a metaphorical one. Here we clearly prefer a literal–metaphorical order, which corresponds to a concrete-abstract order.

(16a) The storm sunk my ship[literal] and my dreams[metaphorical].
(16b)?The storm sunk my dreams[metaphorical] and my ship[literal].

Here we suggest that the repair principle is also at work:

(17a) The storm sunk my ship and **therefore** my dreams.
(17b) The storm sunk my dreams **because it sunk** my ship.

The second sentence is more difficult to repair as more changes are necessary. The repaired sentence (17b) even sounds clumsy, because of the need for repeating "sunk".

b) *Expectations about the information structure*

In Section 4.1 we dealt with expectations related to information structure. We tend to interpret different linguistic cues as indications for what information is given and what information is new and should therefore be the focus of attention for the hearer.

---

[19] Murphy's data (2021a, p. 296) indicates a higher acceptability rating for (15b) compared to (15a).

(18a) John said that the *strangely shaped*[physical_object] and *boring*[content] book was on the table.
(18b) John said that the *boring*[content] and *strangely shaped*[physical_object] book was on the table.

In this example, (18b) sounds better than (18a). This violates the concrete–order bias. (18a), has a concrete–abstract order. "Strangely shaped" is more concrete and simpler than "boring" in the sense that it only involves the physical object module, while "boring" plausibly involves a folk-psychological module. But "strangely shaped" creates an expectation that requires further elaboration and explanation. "Strangely" serves as a cue that "strangely shaped" is the new information part in the information structure, the salient information to which attention should be given. So, when "boring" follows, we get a sort of "narrative disappointment" because it is not an elaboration, nor does it build on "strangely shaped". "Strangely shaped" is quite generic but more expectation-provoking such that we expect more information that explains in what sense the book is strangely shaped.

In (18b), "boring" does not create a further expectation, so "strangely shaped" can follow without any violation of expectations. Rather, the attention is now very much focused on what will follow in the next sentence because we expect some statement that elaborates on "strangely shaped". "Strangely shaped"—as the second predicate—creates a sort of "cliff-hanger" (which is different from a violation of expectations) which craves for a resolution.

c) *Other formal structural primes*

The following pair of sentences, taken from Murphy (2021a), clearly speaks for a simple-complex (or concrete-abstract) order preference. The statement (19a) has a concrete-abstract order and is more acceptable.[20]

(19a) The college has five floors[building] and starts at 8am[event/process].
(19b) The college starts at 8am[event/process] and has five floors[building].

According to our model, in (19a), "has five floors" activates an object/building sense of "college". The object sense, in the absence of any further context, establishes a safe (physical) referent, the Given. We suggest that this physical referent provides a "metonymic" expectation. In a metonymy something stands in for something else that is related in some form. Here the institutional building can stand in for the institution's activity. Therefore, once the physical school building sense has established a physical referent, we can expect a sense of school-activity. Because of this configuration of expectations, (19a) sounds better than (19b).

In (19b), the expression "starts at 8am" activates the college-activity sense, which does not make available a metonymic prior (there are no activity-for-building metonymies). We are left with a somewhat unexpected micro narrative. The fact that college starts at 8am generates some expectations that are not fulfilled by "has five floors", at least without any more specific context. In fact, the first part is about the starting point of a process, so it is reasonable to expect that more is being said about the process. But then there follows an apparently inconsequential description of the number of floors out of the blue. One wonders what the point is here.

So, there is a question about the coherence of (19b). The reason why (19b) sounds odd is that we do not understand the point of the second predicate given the first one. This coherence question should also arise in (19a), as it contains the same pieces of information. However, in (19a) the incoherence is overwritten because of the availability of the metonymic expectation.

However, the important point is that the concrete–abstract preference depends heavily on the context. The context makes available potentially different expectations. In the PP model, the context is also represented as higher-level priors. To see this, one can think of a context such that an *abstract–concrete* order might be preferred. Consider, for instance:

(20a) I have a broken leg and can't quickly get upstairs. My classroom is on the top floor of the school. The bus arrives at 7.45am at school. I have only time to get up 4 floors to arrive without delay if the school starts at 8am. Guess what, **my school indeed starts at 8am**[event/process] **and unfortunately has 5 floors**[building].
(20b) I have a broken leg and can't quickly get upstairs. My classroom is on the top floor of the school. The bus arrives at 7.45am at school. I have only time to get up 4 floors to arrive without delay if the school starts at 8am. Guess what, **my school unfortunately has 5 floors**[building] **and indeed starts at 8am**[event/process].

---

[20] In Murphy's study (2021a, p.296) (19a) obtained a higher acceptability rating than (19b).

In (20) the context of the "if"-conditional in the sentence before the (bold) target sentence evokes an expectation in the target sentence that first it is clarified whether the conditional is met or not. The condition is formulated with respect to the starting time of school. Therefore (20a) with an abstract-concrete (event/process - building) ordering sounds better than (20b).

Now consider the following statements and their context:

(21a) All colleges that start at 8 am and have 5 floors will be closed on Monday. Guess what, **my college starts at 8 am[event/process] and has 5 floors[building]**!

(21b) All colleges that start at 8 am and have 5 floors will be closed on Monday. Guess what, **my college has 5 floors[building] and starts at 8am[event/process]**!

In (21a) the order in the previous statement parallels (and hence primes) the order in the second statement; therefore it sounds more acceptable than (21b).

Take another example to wrap up and bring the different cases of expectations together. Consider (Gotham, 2015, p. 140, also discussed in Murphy, 2021a, p. 44):

(22a) #The bank is FTSE-100 listed[institution] and used to be a police station[building].

(22b) The bank used to be a police station[building] and is FTSE-100 listed[institution].

Again, here we have a concrete–abstract (simple-complex) preference. The reason seems to be that there is no further context, and one might wonder what the point is of mentioning both predications together. Anyway, in this context- and content-poor case, the default concrete-abstract order seems to be preferred. However, in many other versions the b-type sentence sounds perfectly acceptable, for instance:

(23a) The bank is stock-exchange listed[institution] and one of the modern high-rise buildings in NY downtown[building].

(23b) The bank is one of the modern high-rise building in NY downtown[building] and stock-exchange listed[institution].

Here there is no clear preference for either (23a) or (23b). Plausibly this is because there is a connection between high rise buildings in NY downtown and bank headquarters that is not order-sensitive. Listed top banks are stereotypically associated with headquarters in fancy high rise buildings in NY downtown. So, it is easy to expect bank headquarters when we think of NY downtown buildings, and when we think of NY downtown buildings, we can easily think of bank headquarters.

5. CONCLUSION

We have provided an account for context-sensitive linguistic intuitions, specifically for the order-sensitive acceptability judgments of copredication sentences. We suggest that sentence processing is a holistic prediction error minimization process. Contextual factors contribute certain expectations and interact with the content and other prior expectations to produce the linguistic intuition about copredication sentences. We suggested that the observed simple-complex bias is a default strategy to deal with the lack of a richer content and context. We have drawn on research of ordering preferences in multi-adjective strings and generalized it to copredication. We suggested that the observed bias in context-poor conditions corresponds to a default expectation to establish a common-ground referent that is as reliable (objective) as possible. However, this bias is easily overridden once more context is available, for instance, in the form of order-sensitive causal patterns, structural formal primes from previous sentences, and previously established or assumed common ground.

The nature of this paper is theoretical insofar as it proposes to leverage an emerging and increasingly popular cognitive computational paradigm, Predictive Processing (PP), for empirical research into order effects. A first important point we wanted to make is that research into order effects needs to consider context effects and that the PP framework is especially suited to accommodate them. The PP framework more generally points towards new empirical approaches to understanding linguistic intuitions and order effects in the following way. In this paper we have brought to bear the main PP concepts on a high level of description, especially priors, prediction error minimization, and the hierarchical structure of priors. However, here has been no space to discuss in detail the rich toolbox that PP provides on the computational, algorithmic, and (neural) implementation levels (see Sprevak, 2021a, 2021b, 2021c, 2021d for an overview).

For instance, one interesting angle to further explore with a Bayesian approach is how ordering preferences change during an experiment.[21] We would expect, following the work of Zervakis & Mazuka (2013), that acceptance ratings are not stable during an experiment, given that Bayesian updating would influence them. Other —still speculative—lines of possible empirical research involve trying to localize the areas in the brain where certain hyper-priors are represented and then manipulate them (e.g., by non-invasive trans-magnetic stimulation techniques—TMS—by which the activation of certain brain regions can be temporarily manipulated) and observe how ordering preferences are influenced. Finally, we hope to have made a case that there might be benefits to joining forces for the research fields of adjective ordering and copredication order effects.

GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES

No use of such technology has been made.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Christian Michel: Writing – review & editing, Writing – original draft, Conceptualization. **Guido Löhr:** Writing – review & editing, Writing – original draft, Conceptualization.

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Asher, N., 2011. Lexical meaning in context: A web of words. Cambridge University Press.

Asher, N., Denis, P., 2005. Lexical ambiguity as type disjunction. In: Bouillon, P., Kanzaki, K. (Eds.) Proceedings of ᵗhe 3rd International Workshop on Generative Approaches to the Lexicon (G L2005), Geneva, Switzerland, pp. 10–17.

Barrett, L.F., Bar, M., 2009. See it with feeling: Affective predictions during object perception. Philos. Trans. Roy. Soc. London: Biol. Sci. 364 (1521), 1325–1334.

Brandtner, R., 2009. Events and their results: Pragmatic constraints on copredication. In: Talk given at Console XVIII, Universitat Autònoma de Barcelona, December, 19.

Carey, S., 2009. The Origin of Concepts. Oxford University Press, Oxford.

Chomsky, N., 2000. New Horizons in the Study of Language and Mind. Cambridge University Press.

Cinque, G., 2010. The Syntax of Adjectives: A Comparative Study. MIT press.

Clark, A., 2016. Surfing Uncertainty: Prediction, Action, and the Embodied Mind. Oxford University Press.

Clark, A., 2023. The Experience Machine: How Our Minds Predict and Shape Reality. Pantheon Books.

Collins, J., 2023. Copredication as Illusion. J. Semant., ffad014.

Dove, G., 2022. Abstract Concepts and the Embodied Mind: Rethinking Grounded Cognition. Oxford University Press.

Eddington, C.M., Tokowicz, N., 2015. How meaning similarity influences ambiguous word processing: The current state of the literature. Psychon. Bull. Rev. 22 (1), 13–37.

Falkum, I.L., 2015. The how and why of polysemy: a pragmatic account. Lingua 157, 83–99.

Ferreira, F., Bailey, K.G., Ferraro, V., 2002. Good-enough representations in language comprehension. Curr. Dir. Psychol. Sci. 11 (1), 11–15.

Fillmore, C.J., 1975. An alternative to checklist theories of meaning. In: Annual Meeting of the Berkeley Linguistics Society, Vol. 1, pp. 123–131.

Frazier, L., Rayner, K., 1990. Taking on semantic commitments: processing multiple meanings vs. multiple senses. J. Mem. Lang. 29, 181–200.

Freiwald, W.A., Tsao, D.Y., 2010. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. Science 330 (6005), 845–851.

Friston, K., 2010. The free-energy principle: A unified brain theory? Nat. Rev. Neurosci. 11 (2), 127–138.

---

[21] Thanks to an anonymous reviewer for suggesting this.

Gershman, S., 2021. What Makes Us Smart. In: What Makes Us Smart. Princeton University Press.

Gibbs Jr, R.W., 1984. Literal meaning and psychological theory. Cognit. Sci. 8 (3), 275–304.

Glucksberg, S., 2003. The psycholinguistics of metaphor. Trends Cogn. Sci. 7 (2), 92–96.

Goetz, E.T., Sadoski, M., Stricker, A.G., White, T.S., Wang, Z., 2007. The role of imagery in the production of written definitions. Read. Psychol. 28 (3), 241–256.

Goldberg, A., 2013. Constructionist approaches. In: Hoffmann, I.T., Trousdale, G. (Eds.), The Oxford Handbook of Construction Grammar. Oxford University Press, pp. 14–31.

Gotham, M., 2015. Copredication, quantification and individuation. University College London, PhD dissertation..

Gotham, M., 2017. Composing Criteria of Individuation in Copredication. J. Semant. 34 (2), 333–371.

Haber, J., Poesio, M., 2020. Assessing polyseme sense similarity through co-predication acceptability and contextualised embedding distance. In: Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics, pp. 114-124.

Hale, J., 2001. A probabilistic Early parser as a psycholinguistic model. In: Second meeting of the North American chapter of the Association for Computational Linguistics.

Haviland, S.E., Clark, H.H., 1974. What's new? Acquiring new information as a process in comprehension. J. Verbal Learn. Verbal Behav. 13, 512e521.

Hawkins, J.A., 1994. A Performance Theory of Order and Constituency. Cambridge University Press.

Heilbron, M., Armeni, K., Schoffelen, J. M., Hagoort, P., de Lange, F.P., 2021. A hierarchy of linguistic predictions during natural language comprehension. bioRxiv, 2020-12.

Hoffman, P., McClelland, J.L., Lambon Ralph, M.A., 2018. Concepts, control, and context: a connectionist account of normal and disordered semantic cognition. Psychol. Rev. 125 (3), 293.

Hohwy, J., 2013. The Predictive Mind. Oxford University Press.

Holmes, V.T., Langford, J., 1976. Comprehension and recall of abstract and concrete sentences. J. Verbal Learn. Verbal Behav. 15 (5), 559–566.

Holyoak, K.J., Stamenković, D., 2018. Metaphor comprehension: a critical review of theories and evidence. Psychol. Bull. 144 (6), 641–671.

James, C.T., 1975. The role of semantic information in lexical decisions. J. Exp. Psychol. Hum. Percept. Perform. 1 (2), 130.

Jefferies, E., Lambon Ralph, M.A., 2006. Semantic impairment in stroke aphasia versus semantic dementia: a case-series comparison. Brain 129 (8), 2132–2147.

Keil, F., 2021. Why are concrete to abstract models of conceptual change so appealing but so wrong? 74th Annual Meeting of the Society for Philosophy and Psychology.

Kemp, C., Perfors, A., Tenenbaum, J.B., 2007. Learning overhypotheses with hierarchical Bayesian models. Dev. Sci. 10 (3), 307–321.

Klepousniotou, E., Pike, G.B., Steinhauer, K., Gracco, V., 2012. Not all ambiguous words are created equal: An EEG investigation of homonymy and polysemy. Brain Lang. 123, 11–21.

Kousta, S.T., Vigliocco, G., Vinson, D.P., Andrews, M., Del Campo, E., 2011. The representation of abstract words: why emotion matters. J. Exp. Psychol. Gen. 140 (1), 14.

Lakoff, G., 2008. Women, Fire, and Dangerous Things: What Categories Reveal About the Mind. University of Chicago Press.

Langacker, R.W., 1987. Foundations of Cognitive Grammar. Stanford University Press.

Leibo, J. Z., Mutch, J., & Poggio, T., 2011. Why the brain separates face recognition from object recognition. In: Advances in neural information processing systems, 24.

Levinson, S.C., 2016. Turn-taking in human communication–origins and implications for language processing. Trends Cogn. Sci. 20 (1), 6–14.

Levy, R., 2008. Expectation-based syntactic comprehension. Cognition 106 (3), 1126–1177.

Liebesman, D., Magidor, O., 2019. Copredication, counting, and criteria of individuation: a response to Gotham. J. Semant. 36 (3), 549–561.

Littlemore, J., 2015. Metonymy. Cambridge University Press.

Litwin, P., Miłkowski, M., 2020. Unification by fiat: arrested development of predictive processing. Cognit. Sci. 44 (7), e12867.

Löhr, G., 2021. Does polysemy support radical contextualism? On the relation between minimalism, contextualism and polysemy. Inquiry 67 (2).

Löhr, G., 2022. What are abstract concepts? On lexical ambiguity and concreteness ratings. Rev. Philos. Psychol. 13 (3), 549–566.

Löhr, G., 2023. Does the mind care about whether a word is abstract or concrete? Why concreteness is probably not a natural kind. Mind Lang., 1–20. https://doi.org/10.1111/mila.12473.

Löhr, G., Michel, C., 2022. Copredication in context: A predictive processing approach. Cognit. Sci. 46, e13138.

Machery, E., 2009. Doing Without Concepts. Oxford University Press.

Michel, C., 2022. A hybrid account of concepts within the predictive processing paradigm. Rev. Philos. Psychol. 1–27. https://doi.org/10.1007/s13164-022-00648-8.

Michel, C., 2023. Scaling up predictive processing to language with construction grammar. Philos. Psychol. 36 (3), 553–579.

Murphy, E., 2021a. Linguistic representation and processing of copredication. PhD Dissertation. University College London.

Murphy, E., 2021b. Predicate order and coherence in copredication. Inquiry, 1–37.

Neisser, U., 1967. Cognitive Psychology. Appleton-Century-Crofts, New York.

Nelson, D.L., Schreiber, T.A., 1992. Word concreteness and word structure as independent determinants of recall. J. Mem. Lang. 31 (2), 237–260.

Ney, J.W., 1983. Optionality and choice in the selection of order of adjectives in English. General Linguistics 23 (2), 94–128.

Ortega Andrés, M., Vicente, A., 2019. Polysemy and co-predication. Glossa: A J. General Linguist. 4 (1).

Ortega Andrés, M., 2020. Polysemous nouns and copredication. PhD Dissertation. Universidad del País Vasco.

Pickering, M.J., Gambi, C., 2018. Predicting while comprehending language: a theory and review. Psychol. Bull. 144 (10), 1002.

Pickering, M.J., Garrod, S., 2013. An integrated theory of language production and comprehension. Behav. Brain Sci. 36 (04), 329–347.

Pickering, M.J., Garrod, S., 2021. Language Use and Social Interaction. Cambridge University Press, Understanding dialogue.

Pietroski, P.M., 2018. Conjoining Meanings: Semantics Without Truth Values. Oxford University Press.

Pustejovsky, J., 1995. The Generative Lexicon. MIT Press.

Quilty-Dunn, J., 2021. Polysemy and thought: Toward a generative theory of concepts. Mind Lang. 36 (1), 158–185.

Riesenhuber, M., Poggio, T., 1999. Hierarchical models of object recognition in cortex. Nat. Neurosci. 2, 1019–1025.

Rodriquez Roderos, C., 2021. The Processing of Non-Nominal Metaphors. PhD Dissertation, Humboldt-Universität zu Berlin.

Rosales Jr, C.M., Scontras, G., 2019. On the role of conjunction in adjective ordering preferences. Proc. Linguist. Soc. Am. 4 (1), 32.

Rosch, E.H., 1973. Natural categories. Cogn. Psychol. 4 (3), 328–350.

Schwanenflugel, P.J., Shoben, E.J., 1983. Differential context effects in the comprehension of abstract and concrete verbal materials. J. Exp. Psychol. Learn. Mem. Cogn. 9 (1), 82–102.

Schwanenflugel, P.J., Stowe, R.W., 1989. Context availability and the processing of abstract and concrete words in sentences. Read. Res. Q. 24 (1), 114–126.

Scontras, G., 2023. Adjective ordering across languages. Annu. Rev. Linguist. 9, 357–376.

Scontras, G., Degen, J., Goodman, N.D., 2017. Subjectivity predicts adjective ordering preferences. Open Mind 1 (1), 53–66.

Scontras, G., Degen, J., Goodman, N.D., 2019. On the grammatical source of adjective ordering preferences. Sem. Pragmat. 12, 7.

Sprevak, M., 2021a. Predictive coding I: Introduction. PhilSci-Archive http://philsci-archive.pitt.edu/id/eprint/19365.

Sprevak, M., 2021b. Predictive coding II: The computational level. PhilSci-Archive http://philsci-archive.pitt.edu/id/eprint/19366.

Sprevak, M., 2021c. Predictive coding III: The algorithmic level. PhilSci-Archive http://philsci-archive.pitt.edu/id/eprint/19488.

Sprevak, M., 2021d. Predictive coding IV: The implementation level. PhilSci-Archive http://philsci-archive.pitt.edu/id/eprint/19669.

Strain, E., Patterson, K., Seidenberg, M.S., 1995. Semantic effects in single-word naming. J. Exp. Psychol. Learn. Mem. Cogn. 21 (5), 1140.

Tenenbaum, J.B., Kemp, C., Griffiths, T.L., Goodman, N.D., 2011. How to grow a mind: Statistics, structure, and abstraction. Science 331 (6022), 1279–1285.

Trainin, N., Shetreet, E., 2021. It's a dotted blue big star: on adjective ordering in a post-nominal language. Lang. Cogn. Neurosci. 36 (3), 320–341.

Trainin, N., Shetreet, E., 2023. Subjectivity predicts adjective ordering preferences in Hebrew, but lexical factors matter too. Lang. Cogn. Neurosci., 1–7.

Valois, D., 2006. Adjectives: Order within DP and Attributive APs. In: Everaert, M., van Riemsdijk, H. (Eds.) The Blackwell Companion to Syntax, Volume I. Blackwell Publishing.

Wasow, T., Arnold, J., 2003. Post-verbal constituent ordering in English. In: Rohdenburg, G., Mondorf, B. (Eds.), Determinants of grammatical variation in English, De Gruyter, pp. 119-154.

Wheeler, D.D., 1970. Processes in word recognition. Cogn. Psychol. 1 (1), 59–85.

Wiemer-Hastings, K., Xu, X., 2005. Content differences for abstract and concrete concepts. Cognit. Sci. 29 (5), 719–736.

Zarcone, A., Van Schijndel, M., Vogels, J., Demberg, V., 2016. Salience and attention in surprisal-based accounts of language processing. Front. Psychol. 7, 844.

Zervakis, J., Mazuka, R., 2013. Effect of repeated evaluation and repeated exposure on acceptability ratings of sentences. J. Psycholinguist. Res. 42, 505–525.