

LANGUAGE WITHOUT SHARED MEANING: PRELIMINARY CONSIDERATIONS ON THE EVOLUTIONARY BENEFITS OF SUBJECTIVITY

Alexandre Kabbach^{*1,2} and Aurélie Herbelot^{2,3}

*Corresponding Author: alexandre@kabbach.net

¹Department of Linguistics, University of Geneva, Geneva, Switzerland

²Center for Mind/Brain Sciences, University of Trento, Trento, Italy

³Department of Information Engineering and Computer Science, University of Trento, Trento, Italy

We propose a characterization of language that does not rest on the hypothesis that meaning is necessarily shared across interlocutors, since it is fundamentally grounded in the privacy and subjectivity of mental content. We first argue that the function of language is thus best characterized as the *coordination* rather than the *communication* between minds, since it is the constraints on *use*—rather than on meaning *itself*—that are negotiated between interlocutors during linguistic interactions. We then explore the evolutionary benefits of subjectivity and argue that it positively contributes to *adaptability* through: 1) *innovation*, as conceptual variability at the individual level increases the likelihood of the group of finding relevant conceptual innovations when exposed to environmental challenges; and 2) *transfer*, as subjective coordination allows for individual innovations to spread across a whole population, without individuals having to align their conceptual spaces and thereby lose the benefits of conceptual variability.

1. Introduction

Communication through language rests on mutual understanding: to understand you, I must figure out exactly what you *mean* by what you *say*. Successful communication therefore requires shared meaning and appeal to mental content. And yet, mental content remains inherently *private* and *idiosyncratic* in nature, which poses a major challenge to communication, for how can we indeed guarantee mutual understanding if the meaning speakers intend to communicate through language remains highly subjective and inaccessible to others (Pelletier, 2017, p.63)?

In this work, we propose to turn the problem on its head: rather than trying to reconcile subjectivity with communication, we explore what it would mean for the characterization of language to be structured around subjectivity, and what benefits subjectivity could have from an evolutionary perspective.

Our first contribution is to argue that the function of language that unfolds is actually best characterized as the *coordination* rather than the *communication* between minds—where communication and coordination are formally distinguished with respect to their requirement of shared meaning: strict for communication,

loose for coordination. Concretely, we argue that coordinating minds is best conceived as the *private satisfaction of shared constraints*, given that 1) speakers never have direct access to what their interlocutors actually mean, so that; 2) it is the constraints on *use*—rather than on meaning *itself*—that are actually negotiated between interlocutors during linguistic interactions.

Our second contribution is to formalize a clear hypothesis about the evolutionary benefits of subjectivity in improving adaptability through cumulative culture—that is, the ability of humans to innovate by creating new knowledge as well as to preserve existing knowledge by passing it on to other members of the group. We argue that subjectivity is advantageous in two ways: 1) for *innovation*, conceptual variability at the individual level increases the likelihood of the group of finding relevant conceptual innovations when exposed to environmental challenges; and 2) for *transfer*, subjective coordination allows for individual innovations to spread across a whole population, without individuals having to align their conceptual spaces and thereby lose the benefits of conceptual variability.

Through this work, we hope to contribute to the discussions on both the evolution and the function of language by decoupling the characterization of *communication success* from the characterization of *interpretation success*. Doing so allows us to treat the communicative function of language no longer as a theoretical prerequisite, “hardcoded” into our model of language, but as a mere hypothesis that can now be analyzed comparatively for explanatory adequacy.

2. Language, communication and the problem of subjectivity

2.1. *The code model of communication*

According to Sperber and Wilson (1986/1995, p.2) “From Aristotle through to modern semiotics, all theories of communication were based on [...] the code model” later formalized by Shannon and Weaver (1949), which characterizes communication as information transfer between a sender and a recipient who exchange messages encoded and decoded to and from (linguistic) signals through a potentially noisy channel.

The problem with such a characterization—as *pragmatics* has extensively shown—is that *there is more to language than what is said*, so that linguistic signals alone often do not suffice to unambiguously decode messages: they must be put into *context* (Sperber & Wilson, 1986/1995, ch.1 §2). But context here is to be understood not only as the linguistic (e.g. sentential or discourse) context under which communication takes place, but also as the non-linguistic context such as the set of background experiences, knowledge, beliefs, desires or assumptions grounding the interpretation of messages and the mapping of linguistic *form* to conceptual *meaning*. Recipients must understand not only *what is said*, but also what senders (or speakers) actually *mean* by what they say, and to do so they must identify their *intentions* (Grice, 1989).

Inferential processes are therefore required to properly decode messages, and decoding is probably best formalized as *interpretation*, given how messages are themselves best characterized in terms of mental content and as combinations of propositions and propositional attitudes (Sperber & Wilson, 1986/1995, p.57).

2.2. Subjectivity, similarity and communication success

Major problems then arise when trying to characterize communication success. Because mental content is inherently private and background experiences grounding conceptual knowledge necessarily idiosyncratic and thus highly subjective and speaker-specific (Pelletier, 2017, §6). Yet, the code model assumes that communication is successful when the message decoded by the recipient is *identical* to the one encoded by the sender. How can we guarantee, then, that senders and recipients do in fact understand one-another if the messages they map to and from identical linguistic signals are processed through *different* conceptual backgrounds?

The problem is well acknowledged¹ and usually resolved along two lines. First, by relaxing the identity constraint on messages to sheer similarity. Second, by emphasizing that agents in linguistic interactions are not passive but actively negotiate meaning and gradually align their conceptual representations through conversation (Clark, 1996; Pickering & Garrod, 2006). But what does it mean for two messages to be similar? And how much similarity is enough to guarantee communication success? In effect, the concept of similarity merely displaces the problem while still posing major theoretical and empirical challenges (Medin, Goldstone, & Gentner, 1993). Alignment-based approaches to communication are not without problems either, as they still remain committed to the code model's characterization of communication success, inasmuch as conceptual representations are expected to align across interlocutors during conversation and misalignment between speakers is only tolerated when it pertains to aspects irrelevant to the conversation at hand. Shared meaning is thus still very much of a requirement, and actually expected, given how speakers of the same (linguistic) community are assumed to share *common ground* (Kabbach & Herbelot, 2021, §2.3).

3. What exactly is the function of language?

3.1. The need for an alternative hypothesis

In any case, a characterization of communication success that rests on the hypothesis of shared meaning can only be an approximation at best, for speakers can never actually verify that they do share meaning in practice. As Sperber and Wilson (1986/1995, p.18) emphasize, “when human beings try to communicate with each other, they are aiming at something they can never, in fact, achieve”. The

¹And has been extensively discussed in the philosophy of language, notably through the question of meaning holism and instability. See (Jackman, 2020), especially §3.2 and references therein.

formalization of communication success in the code model is psychologically unrealistic, for it would actually require speakers to perform an infinite number of recursive mutual checks on their respective mental content to ensure their making use of the same context and guarantee their encoding and decoding identical messages (Sperber & Wilson, 1986/1995, ch1 §3).

Moreover, if all speakers can do during conversation is strive to *avoid conflict* rather than *enforce agreement*—given that overt disagreement is the only kind of information they actually have access to in practice—the question boils down to whether the *absence of (behavioral) conflict* really equates (*conceptual*) *agreement*, which is far from obvious. The structure of language and cognition indeed appears to be such that you and I can agree on the fact that *kittens are cute* without having to agree on what kittens are, or up to what age one can call a cat a kitten, on why we find kittens cute, or even on what constitutes a good prototype of cuteness. Many argue that, in such cases, conceptual variability between us does not really matter: language can remain vague and vagueness can even serve communication (van Deemter, 2010). But saying that conceptual variability does not matter does not mean that it is necessarily marginal. Regardless, it remains that language appears to be structurally *underdetermined* regarding concepts, in that my uttering the word *cat* does not even begin to convey the richness of my concept CAT. Therefore, our agreeing on the fact that *kittens are cute* remains what it is: an *appearance* of agreement between what are probably highly subjective and thus clearly distinct conceptual representations, compatible in this context only inasmuch as they do not lead to overt disagreement between us.

Considering subjectivity to be negligible overall in language processing is thus probably somewhat of an arbitrary choice—true *by definition* rather than by empirical observation—and the assumption of shared meaning may very well prove to be biased by our tendency to actually overestimate how much we share meaning with others (Martí, Piantadosi, & Kidd, 2019). Be that as it may, would it really make sense to talk about “communication” if it did not involve shared meaning in the first place? Or if the notion of mutual understanding was not grounded in a form of “sameness of concept” among interlocutors? The language-for-communication hypothesis still resists the relaxing of the shared meaning assumption, even if just for questions of clarity and consistency. If we are to call into question its characterization of success, we thus need more than a new model of communication: we need a new hypothesis about the function of language *itself*.

3.2. *The language-for-coordination hypothesis*

The primary contribution of our work is to propose a radical change of perspective on the role subjectivity is assumed to play in language processing (in general) and in interpretation (in particular). In considering that subjectivity is structural rather than marginal, we argue that the function of language is actually best characterized

as the coordination rather than the communication between minds.² Concretely, we distinguish coordination from communication with respect to their requirement of shared meaning: strict for communication, loose for coordination. In effect, coordination dispenses itself from having to characterize success in terms of necessary conceptual alignment between speakers. Indeed, interpretation can be thought of as a problem of *constraint satisfaction*, but one that admits multiple solutions and that is resolved independently by speakers taken individually, so that they need not converge to identical solutions.

Let us take a concrete example to illustrate our point: it is common in the scientific literature to find counter-arguments to the subjectivity of meaning revolving around the problem of concept-to-word mapping. The argument goes more or less as follows: if you call a *dog* what I call a *cat*, how can the two of us communicate with one-another? Yet, things need not be so binary: there is ample leeway between my saying *cat* when you say *dog* and you and I having *identical* concepts of CAT and DOG. The example is of course limited—what are identical or similar concepts in the first place?—but it gives us an intuitive sense of the claim: the mapping between words and concepts need not be one-to-one, so that different speakers can actually associate different concepts to identical words.

In fact, the process of *making sense* of linguistic signals could perfectly accommodate the pervasiveness of subjectivity specifically because it remains private and because meaning is never negotiated directly across speakers. You and I certainly do have to agree that what we have in front of us is referred to as a *cat* and not a *dog*, but we will never negotiate our conceptual representations of CAT and DOG directly: only the constraints they have to satisfy externally. And since the words *cat* and *dog* are structurally underdetermined with respect to our concepts of CAT and DOG, the problem of constraint satisfaction so defined may very well admit multiple solutions that will be produced privately by each of us, within the subjectivity of our respective conceptual spaces.³

²Our use of the term “coordination” is somewhat arbitrary and primarily motivated by the desire to clearly depart from the term “communication” without resorting to a neologism. Nonetheless, it also rests on the observation that, to coordinate, people need not *do the same thing*. The examples of coordination provided by Clark (1996, p.3) for instance—such as *waltzing* or *playing music*—give a sense of the intuition. From not *doing* the same thing to not *thinking* the same thing, there is a bridge between behavior and mind that we will take the liberty to cross here.

³We neither provide nor commit to any particular characterization of the notion of “constraint”, though we do commit to the view that concepts are mental entities. In the above example, constraints could for instance be considered “referential”, insofar as you and I have to agree that the word *cat* should be used to refer to the animal we have in front of us. But in a different conversation, I may try to convince you that *cats are adorable creatures*, and constraints on usage of *cats* or *adorable* will be formalized differently. Anyhow, the point made by the language-for-coordination hypothesis is that aligning usage does not require aligning concepts. The intuition is plain: we need not agree for the same reasons. “Making sense” only requires that we find at least *one way* to agree, that is, one way to accommodate our subjective concepts to the situation at hand.

4. Coordination and subjectivity from an evolutionary perspective

4.1. *Language-for-coordination and explanatory adequacy*

Our proposition is not without empirical support. Kabbach and Herbelot (2021), for instance, question the standard view that communication succeeds because the impact of subjectivity is negligible (i.e. the view that speakers of the same linguistic community share significant *common ground* despite having private mental content (Clark, 1996)). They specifically investigate whether common ground emerges from linguistic interactions in the first place, and show using computational models of lexical meaning that aligning different models on *parts* of their semantic spaces does not necessarily lead to increased *overall* similarity between them. Worse, they show that models often manage to improve superficial alignment by actually resorting to idiosyncratic rather than commonly shared aspects of their conceptual spaces, providing thereby a computational characterization of the distinction between agreeing and compatible semantic representations.

Such work, of course, is no definite proof of our argument. But it stands in the long tradition of computational work that, without guaranteeing what *is*, can at least give us some intuition about what *could be*: in this case, superficial behavioral alignment does not have to equate deep conceptual agreement. To us, the role of subjectivity in language should no longer be downplayed: in practice, conceptual similarity across people is extremely hard to validate experimentally (Kabbach & Herbelot, 2021, §2.2) and the fact of the matter is that “different subjects give individually different results on the many tasks about meaning that have been administered over the decades in cognitive psychology” (Pelletier, 2017, p.74). Those empirical data have to be accounted for one way or another.

4.2. *Adaptability: the evolutionary benefits of subjectivity*

Interestingly, support for the language-for-communication hypothesis may come from considerations about the emergence of language and asking what benefits subjectivity could have from an evolutionary perspective. One way to concretely tackle the problem is to consider the impact of cumulative culture on adaptability. According to Mithen (1996), human societies distinguish themselves from apes by a striking technological gap, giving our species a major survival advantage as it enables us to adapt much more easily to environmental changes. The superiority of those technological abilities have been said to arise from cumulative culture, i.e., from our ability to innovate by creating new knowledge as well as to preserve existing knowledge by passing it on to other members of the group (Mesoudi, 2011). The question that arises, then, is how to provide a functional characterization of the cognitive processes involved in cumulative culture, and more specifically of the role language plays in both (conceptual) *innovation* and *transfer*.

One interesting take on the matter is the study of Toya and Hashimoto (2018) which investigates the evolutionary benefits of recursion. Drawing a parallel be-

tween mental operations and action sequences involved in toolmaking, they show how specific types of recursive operations can lead to improved tool manufacturing strategies and to the production of more and more diverse tools—ultimately leading to improved population fitness, better adaptability and increased survival capabilities. Yet, if their study paves the way for an empirical investigation of the question that interests us here—demonstrating how increased conceptual diversity can positively contribute to innovation and adaptability—it tells us little about *transfer* and how a specific innovation may spread across a whole population. More specifically, will members of a group necessarily transmit *whole* production strategies or rather focus on enabling all members to produce the desired outcome irrespective of the production method used? Which approach will prove better able to entrench innovations and ultimately improve adaptability? This is probably where we can best foresee the possible contribution of the language-for-coordination hypothesis: if what matters is *what* you make and not *how* you make it, relaxing the constraint on shared meaning could prove decisive. As Toya and Hashimoto (2018) indeed detail, a single linearized sequence may correspond to different hierarchical structures (see Figure 1) so that if the point is to match the sequence rather than a particular structure (the *what* versus the *how*) diversity in the *how*—what they call “diversity of production”—can actually prove beneficial. This directly echoes considerations of §3.2 regarding the underdeterminacy of language vis-a-vis concepts and the benefits of subjectivity for coordination.




Figure 1.: Different hierarchical structures corresponding to different combinations of recursive operations generating the linearized sequence *abcd*.

We can thus now formulate a clear hypothesis about the possible evolutionary benefits of subjectivity: by enabling agents to coordinate *without* having to align their respective conceptual spaces, language would allow for the spread of a particular conceptual innovation within a group, without losing the benefits for innovation of maintaining distinct conceptual spaces across individuals.

5. Conclusion

Is the function of language best characterized in terms of communication or coordination? Inasmuch as communication entails that meaning must be shared across interlocutors when coordination does not, we argue for the latter. Indeed, considering the major evolutionary benefits that subjectivity brings to cumulative culture and adaptability, we argue that we should reconsider the place of subjectivity

as being central to any account of human language and cognition, rather than a negligible byproduct of marginally different background experiences grounding conceptual knowledge.

References

- Clark, H. H. (1996). *Using Language*. Cambridge, UK: Cambridge University Press.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Jackman, H. (2020). Meaning Holism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2020 ed.). Metaphysics Research Lab, Stanford University.
- Kabbach, A., & Herbelot, A. (2021). Avoiding Conflict: When Speaker Coordination Does Not Require Conceptual Agreement. *Frontiers in Artificial Intelligence*, 3, 95.
- Martí, L., Piantadosi, S. T., & Kidd, C. (2019). Same Words, Same Context, Different Meanings: People are unaware their own concepts are not always shared. In *Proceedings of CogSci* (pp. 2296–2302). Montréal.
- Medin, D., Goldstone, R., & Gentner, D. (1993). Respects for Similarity. *Psychological Review*, 100(2), 254–278.
- Mesoudi, A. (2011). *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. Chicago and London: University of Chicago Press.
- Mithen, S. J. (1996). *The Prehistory of the Mind: a Search for the Origins of Art, Religion and Science*. London: Thames and Hudson.
- Pelletier, F. J. (2017). Compositionality and Concepts—A Perspective from Formal Semantics and Philosophy of Language. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and Concepts in Linguistics and Psychology* (pp. 31–94). Cham: Springer International Publishing.
- Pickering, M. J., & Garrod, S. (2006). Alignment as the Basis for Successful Communication. *Research on Language and Computation*, 4(2), 203–228.
- Shannon, C. E., & Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.
- Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and Cognition* (Second ed.). Oxford: Basil Blackwell.
- Toya, G., & Hashimoto, T. (2018). Recursive Combination Has Adaptability in Diversifiability of Production and Material Culture. *Frontiers in Psychology*, 9, 1512.
- van Deemter, K. (2010). *Not Exactly: In Praise of Vagueness*. Oxford: Oxford University Press.