

What do humans know when they use language?

One of the core goals of the generative enterprise (and other cognitive approaches towards language) is to explore what humans know when they know a language. To achieve this goal, generative linguists typically explore properties of individual languages, the commonalities across languages, as well as the range and limits of observed variation. What has long been ignored in this approach are aspects of language restricted to language in use. Two assumptions that prevail the generative enterprise have conspired to this status quo:

- (1) Linguistic *competence* is the object of investigation and is to be distinguished from linguistic *performance*.
- (2) The *sentence* is the unit of grammatical analysis.

On this view, units of language (UoLs) that appear at the periphery of sentences and which serve to regulate linguistic interaction (*oh, huh, well...*) are not considered to fall within the purview of grammar. However, over the past two decades there is a growing consensus that the distribution of such UoLs can be successfully analyzed on the hypothesis that they occupy the very top of syntactic structures. If so, this has profound implications for the two assumptions above as summarized in (1') and (2').

- (1') UoLs that serve to regulate interaction should be considered part of linguistic competence.
- (2') Grammatical analysis should not be restricted to the classic notion of the sentence.

While there are several competing proposals to model the integration of aspects of language in use into our knowledge of language, in this series of workshops I focus on one such proposal: *the Interactional Spine Hypothesis* (Wiltschko 2021), according to which grammatical knowledge includes not only knowledge of how to construct the propositional aspects of language (*p-language*) but also its interactional aspects (*i-language*).

Once we acknowledge that (at least some) conditions on language use are part of our grammatical competence, new research questions emerge as well as the need for new methodologies which go beyond the exploration of sentences in isolation. During the three workshops, I explore some of these questions and methodologies.

First Workshop (May 19, 2025)

How do children acquire knowledge of language use?

Assuming that *i-language* is part of our linguistic competence raises the question regarding its acquisition. How is *i-language* acquired and how is it integrated with *p-language* during language development? These are the questions I pursue in the first workshop.

Based on two corpus-based case studies of the early use of *huh* and of backchannels, I demonstrate that interactional UoLs are among the first "words" to be acquired. This raises an interesting challenge for existing assumptions regarding language acquisition. Specifically, it is commonly assumed that syntactic structure matures in an upward fashion, i.e., that trees grow upwards (see Friedman et al. 2021 for a recent incarnation of this view). However, if *i-language* is associated with the very top of the tree, its early acquisition is mysterious.

In this workshop, I introduce an alternative: the *inward growing spine hypothesis* (Heim & Wiltschko, in print). Accordingly, syntactic acquisition proceeds in an inward fashion starting at both ends of the adult structure: categorizing structures at the bottom and the structure responsible for managing turn-taking at the top (Wiltschko's ResponseP). These structures are first linked via the layer of structure which – in adult-language – corresponds to CP. In subsequent stages, the intermediate structures responsible for grounding in interaction and for anchoring propositions are acquired. I will show that the inward growing spine hypothesis allows for a straightforward explanation of the acquisition path of *huh* and backchannels. Moreover, I argue that the observed patterns of acquisition support the view that *i-language* is part of our linguistic competence and that the sentence should not be considered the primary unit of analysis.

Second Workshop (May 20, 2025)

Does knowledge of language use affect the way we talk to ourselves?

Most frameworks that assume aspects of language in use to be part of syntactic structure minimally include a representation of the interlocutors. The Interactional Spine Hypothesis is no exception. Interactional roles are assumed to be introduced in dedicated *grounding* phrases, where speaker and addressee are represented as the holders of their individual *grounds* (a placeholder for their epistemic states). If grammatical knowledge includes representations of the (epistemic states of) speaker and (what the speaker assumes to be the epistemic state of their) addressee, then it is predicted that utterances in interaction are sensitive to the identity of the speaker and addressee, and the relation between them. This in turn requires a methodological shift in that the well-formedness of utterances must be judged relative to a particular context that takes the identity of the interlocutors into consideration. Much current research focusses on the use of honorifics and other addressee-oriented formality strategies to empirically explore the nature of these representations.

In this workshop I explore this question from a novel angle, namely self-talk. That is, if the addressee role is part of the grammatical representation of i-language, then the question arises as to what happens in situations when people talk to themselves, i.e., when there is no addressee. I demonstrate that self-talk provides us with striking evidence for the grammatical representation of the (epistemic state of the) addressee role.

Third Workshop (May 21, 2025)

Does knowledge of language use affect the way we express emotions?

When we use language in interaction, we do not only formulate our thoughts, we also express our emotions. Thus, the question arises as to how the expression of emotions is integrated into our knowledge of language use. It has long been established that the expression of emotions pervades all levels of language (from phonetic details affected by emotions to particular discourse strategies). However, what is conspicuously missing in the languages of the world – and thus arguably from our knowledge of language – are grammatical categories that are dedicated to the expression of emotions. This is not a trivial fact as one could easily imagine a language where speakers must specify how the propositional content of the utterance relates to their emotive state. But this does not seem to be the case, as I demonstrate in this workshop. I further discuss the consequences of this empirical finding for the question of how our knowledge of language relates to our emotions and ultimately for our understanding of human cognition more generally.

Public lecture (May 22, 2025)

Language makes us think AI knows stuff

It is well-known that AI frequently produces false information: output that appears plausible but is not factual. This is known as 'hallucinations' or 'bullshit'. Yet, AI pervades our lives including in domains where one might hope that factuality matters, e.g., medicine, warfare, law, and education. Thus, we find ourselves in a curious situation: Why do we place our trust in an intelligence that is not trustworthy? In this talk I explore this question from a linguistic angle. I argue that one reason for our trust in AI has to do with our unconscious knowledge of language.

I start by demonstrating that there is a grammar of certainty which is characterized by the absence of marking uncertainty. That is, when humans are certain about a fact, they use an unmarked declarative sentence (*It is raining*) to express this certainty. It is only when humans are not certain that sentences must be marked to indicate this uncertainty (*It might be raining; Apparently it is raining; I think it is raining*). I argue that this is an intrinsic characteristic of all human languages, and hence part of our unconscious linguistic knowledge. When someone utters *It is raining*, we are led to believe that they know that it is raining. Significantly, AI presents information with the use of such unmarked declaratives. In fact, it is trained to assert that it doesn't have beliefs or consciousness and hence will not say *I believe that ...* This leaves us with a situation where our knowledge of language leads us to interpret AI output as if AI knows stuff.