

Language as a Bottleneck of the Self

Richard Csaky, 2026.02.17

This essay treats language as a paradoxical cognitive technology: a compact, lossy code that nonetheless seems to enable the felt continuity of consciousness and deliberation. I propose that inner speech is primarily a communicative act—even when its audience is oneself—because producing language requires implicitly modeling an interpreter. That recursive “communicator–interceptor” loop yields a third-person stance on the self: we experience ourselves as an object because the brain simulates how another would parse our explanations. On this view, consciousness is not the executive source of action but an emergent side-effect of self-communication, which can sometimes influence behavior by feeding structured, compressed summaries back into the subconscious system that largely decides and acts. The model also clarifies familiar frictions between stated intentions and action, and suggests a resource trade-off: language can refine control and insight by forcing thought through an efficient channel, yet it can also inhibit performance by consuming neural capacity, explaining the phenomenology of flow, skill acquisition, and the difficulty of sustained narration. Finally, I sketch a continuum hypothesis: non-linguistic animals may possess forms of second-order volition without the same self-modeling phenomenology, while communicative species may exhibit graded consciousness and agency.

Note on LLMs: Abstract written by LLM, all other words are mine.

Introduction

(Clark et al., *The Extended Mind*, 1998) — *The major burden of the coupling between agents is carried by language. Without language, we might be much more akin to discrete Cartesian 'inner' minds, in which high-level cognition relies largely on internal resources. But the advent of language has allowed us to spread this burden into the world. Language, thus*

construed, is not a mirror of our inner states but a complement to them. It serves as a tool whose role is to extend cognition in ways that on-board devices cannot. Indeed, it may be that the intellectual explosion in recent evolutionary time is due as much to this linguistically-enabled extension of cognition as to any independent development in our inner cognitive resources. Perhaps language is an impediment to cognition but a great enabler of extension — what could then be possible if extension was enabled without relying on language?

Thesis

Language is strange. It is a compact, lossy representation of information, yet it also feels like the thing allowing me to be conscious — through language I can formulate alternatives and reason, as well as off-load, yet if language is just an afterthought of the subconscious then we arrive at a contradiction. Perhaps language really is just a means of communicating inner states and processes — including communicating them to myself. The inner dialogue thus construed makes sense as I probably need to continuously model how what I'm communicating would be interpreted by the other — therefore within myself I not only have to be the communicator, but the interceptor as well. This seems logical, without such a process how could I be able to predict the effect of my communication. The inner dialogue allows me to view myself as a third person precisely because the brain is trying to model a third person. Thus, the experience of consciousness is really a side-effect of this third-person view that I have upon myself, which I could not have without language and the modeling of the other. Perhaps it can even define or affect behavior in that through this third-person view my subconscious gains additional information. So the decision is still made by the subconscious but self-communication could influence it, just as processing language from others does. Andy Clark develops the “language as mind-transforming scaffolding” theme at book length: language isn't just expressing thought, it's a cognitive tool that helps build new forms of thought and self-control.

So the production of language and the inner monologue is aimed at explication and communication of behavior, and at the same time observing this *can* influence behavior, but not necessarily. This would also explain why sometimes people are not able to do what they say they want to do. I certainly don't need to explain my thoughts to myself, but it could be that forcing my thoughts through this lossy, but highly

information-efficient bottleneck—aimed at communication—provides additional insight.

The inner monologue is not really me, and yet even if behavior is produced by the subconscious it does not preclude free will. I am my subconscious, and just because I cannot fully communicate it to myself does not mean that I am not doing inner deliberation and I am not having second-order volitions. Indeed it is probably a physical impossibility to put into language the full processing of the subconscious (1, see below).

So perhaps animals too can have implicit second-order volitions, but without language they may lack the experience of being conscious, since they cannot model themselves as a third person and explain/communicate their own behavior/thinking to themselves, such as a truly third person would. Equally likely, this experience of consciousness and third-person view may be necessary for the subconscious to develop second-order volitions — it is hard to tell. But our closest cousins do use communication and some kind of primitive language, so it could simply be that they do have both free will and consciousness just a lower level — and indeed why shouldn't these be on a continuum.

(1) Putting my thoughts into language takes quite a lot of neural resources, explaining both why thinking (in language) too much can paralyze a person and preclude action, and why moments of flow or intense activity are associated—really enabled—by not employing the brain for self-communication. A person who is just starting to learn how to drive will certainly experience flow more often than an experienced driver, simply because the brain does not possess enough capacity / resources for additional self-narration. One can deliberate both with and without self-communication — meditation is an extreme example of the former and stream-of-thought writing as an extreme example of the latter. Normally we are somewhere in between, experiencing alternating periods with and without self-narration. Forcing oneself into sustaining either extreme is difficult. It is likely that there is a trade-off — the less resources I use to narrate my thoughts the more I have for the thoughts themselves, but the narration may provide additional insight by forcing my thoughts through this low-capacity channel that is aimed at communication. These two properties of language, high compression rate and optimized for transmitting information, are fundamental. Language is not easy either, I

often experience that I cannot put my thoughts into words — evidence for how unnatural and resource-intensive this process is.