

Norm Participation Grounds Language

David Schlangen

Grundlagen der Computerlinguistik // Department Linguistik // Universität Potsdam

<http://clp.ling.uni-potsdam.de>

david.schlangen@uni-potsdam.de

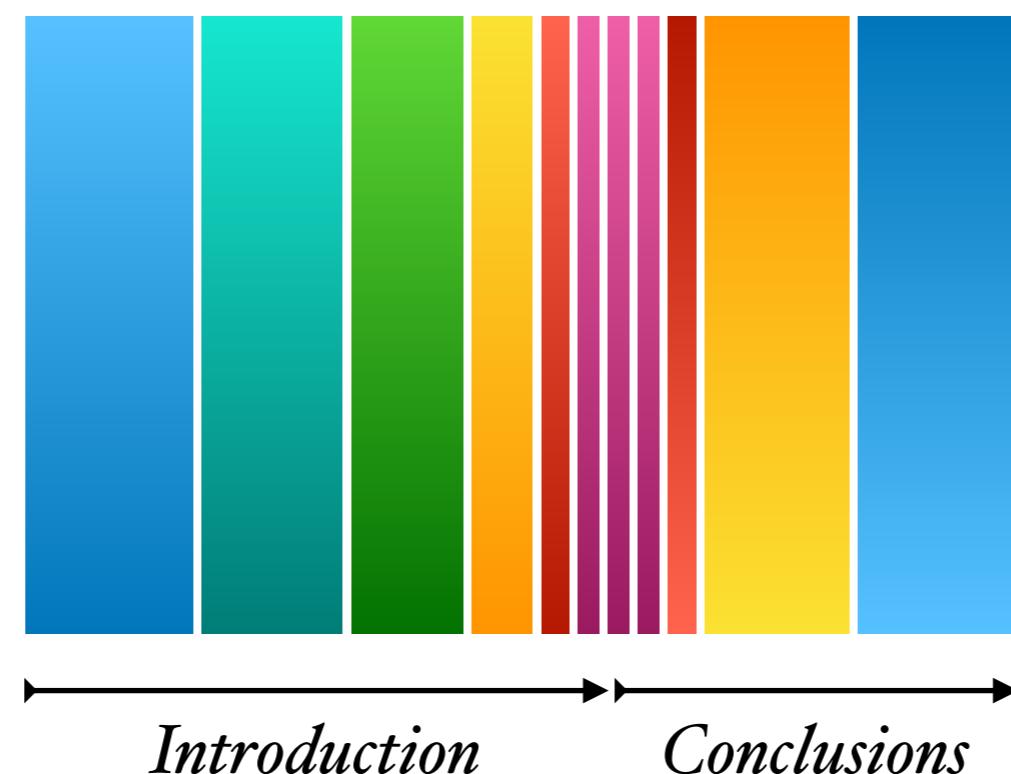
(dis)embodiment, CLASP Workshop

2022-09-16

These slides:

<https://clp.ling.uni-potsdam.de/talks>

Structure of this talk



Goals

- Negative Goal: Argue against “grounded language”.
- Positive Goals:
 - Argue for richer understanding of grounding, & emphasis on investigating & modelling certain dialogical skills.
 - Contribute to clearer understanding of conditions of use of NLP technology.

Resist!

- “Norm participation grounds language”

grounds language

Perception, Vision, Experience, ...

- Interpretability and Analysis of Models for NLP
- Language Grounding to Vision, Robotics and Beyond
- ...

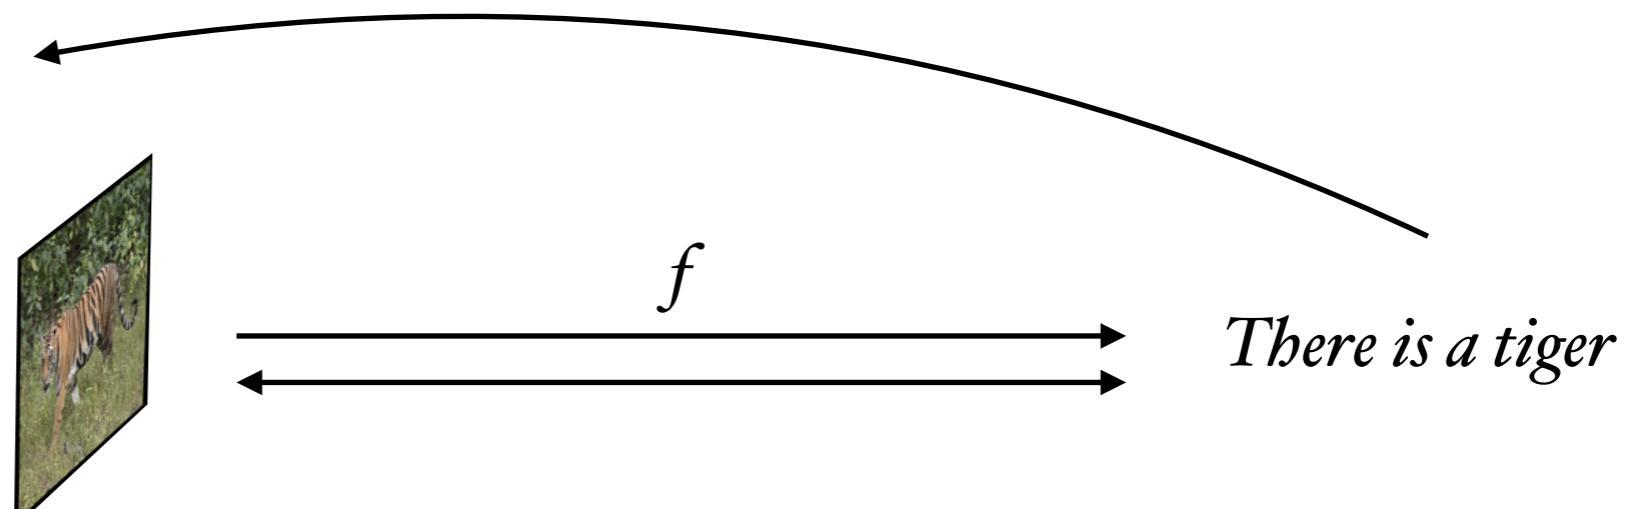
- Against accuracy!

For correctness

- interactive / dialectic
- (partially) holistic

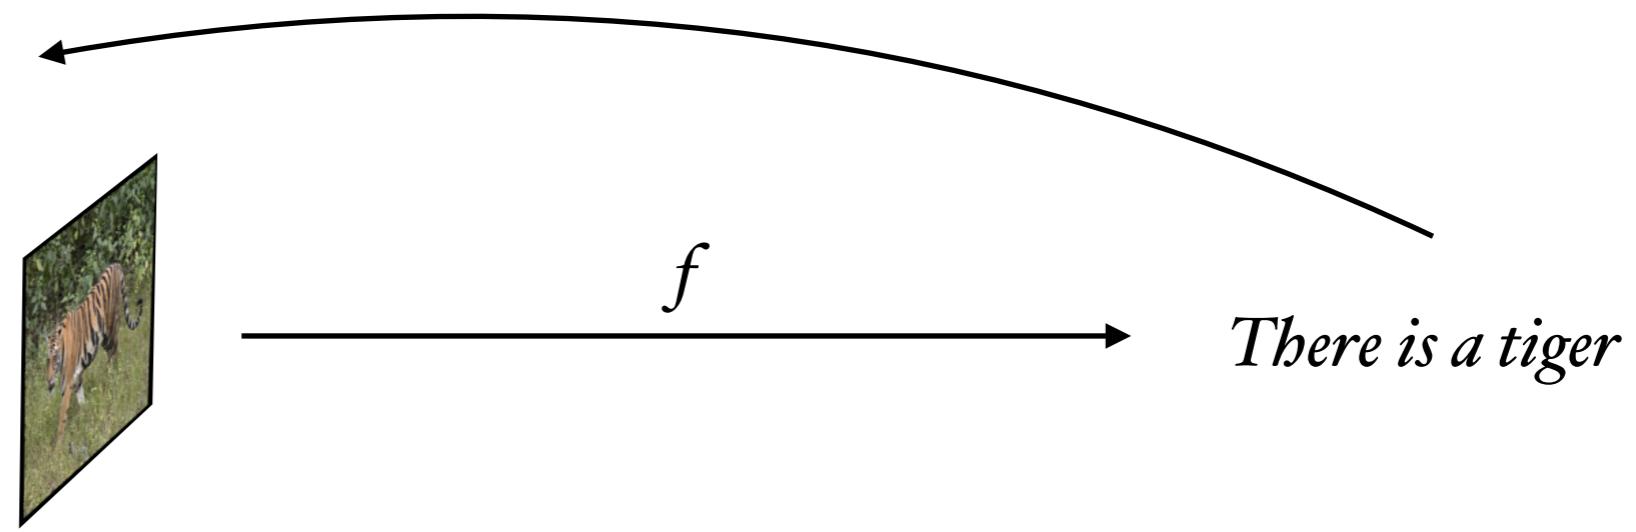
observation statements

Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



observation statements

Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

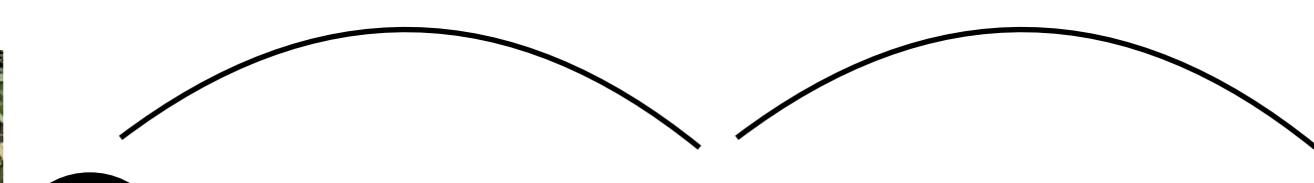


This is flawed, along two dimensions:

- Quantity; 3 basic moves in a language game (Sellars 1954)
 - Language-Entry
 - Language/Language
 - Language-Exit
- Quality: The link is *normative*.
 - Agent must be able to be mistaken about *applicability*.
 - The link must be *accepted* to exist.

commitments & entitlements

Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



epistemic
standing

There is a tiger

What makes this a justified assertion, a correct assertion, a correctly performed assertion;
what makes the asserter competent? What *entitles* asserter to make this assertion?

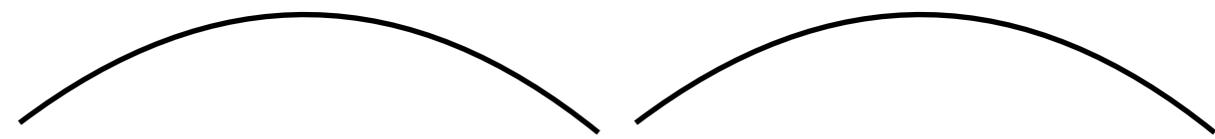
- *true statement*: SoA obtains, it's the truth-maker to the truth-bearer
- justified *assertion*:
 - epistemic conditions for forming reliable beliefs hold
 - there *is* a norm that justifies the forming of this belief
 - there *is* a norm that justifies the expression of this belief in this way

The norms that
entitle also *commit*!

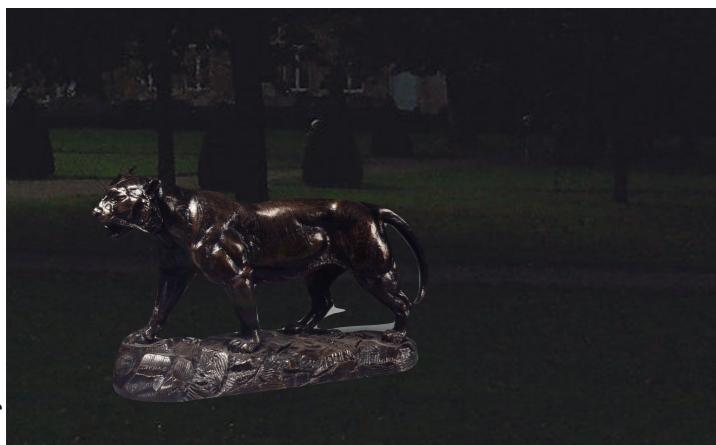
why beliefs?

Charles James Sharp

<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



There is a tiger



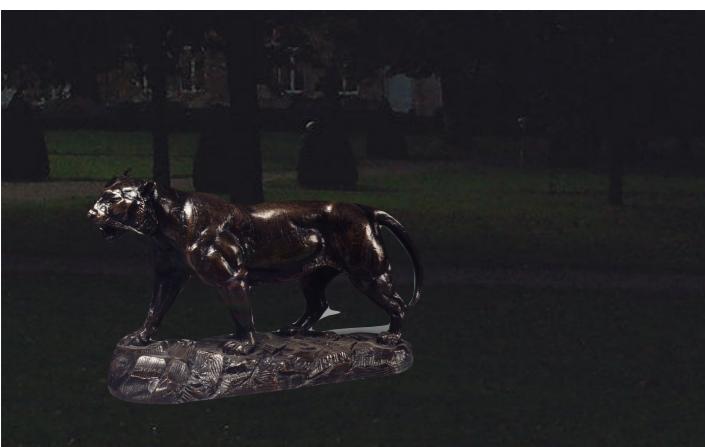
There is a tiger
Ah, no, that's my
new Barye copy.
*Oh, I really thought
there was a tiger!*



There is a tiger
WTF?

why inferences?

Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



A tiger is an animal
An animal is a living thing
material inferences / “common sense”

A sculpture is an artefact
A sculpture of something looks like that something

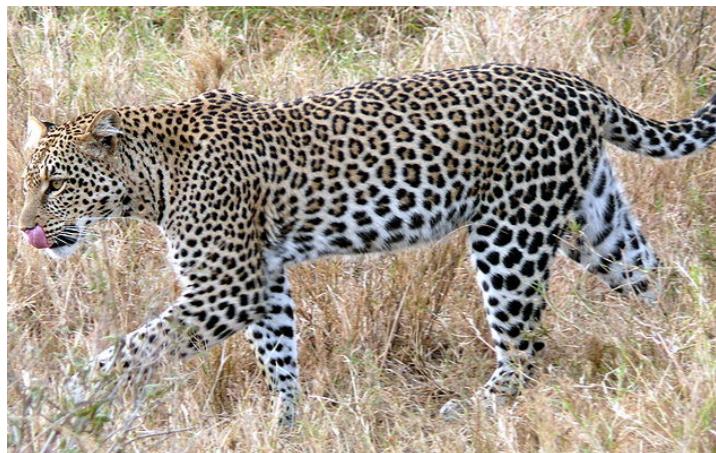
There is a tiger

There is a tiger

norm participation

Charles James Sharp

<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



Haplochromis (Wikipedia)

<https://creativecommons.org/licenses/by-sa/3.0/deed.en>

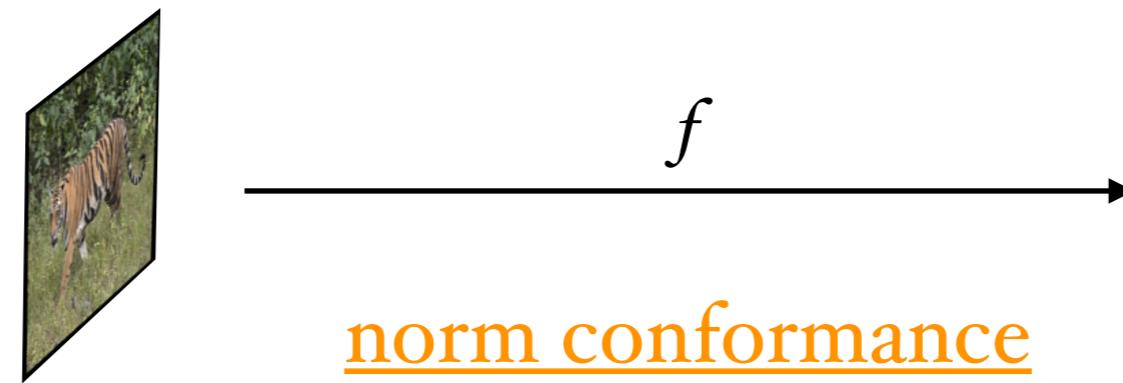
There is a /'ti:gə/
Yes, a /'taɪgə/

Tiger!
Yes, there's a tiger.

There is a tiger
Well, it's a leopard.
Different fur, smaller.

intermediate summary

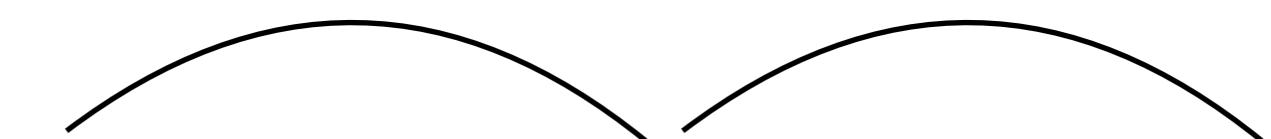
Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



norm conformance

- (Implicitly) assumes absolute authority of norms reflected in training data.
- Cannot make criteria explicit (= debatable).
- Machines cannot commit.

Charles James Sharp
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>



norm participation

There is a tiger
Really?

practical takeaways

- practical consequences:
 - let us try to understand better how these processes work: conceptually, empirically, formally, and computationally.
 - E.g., the relation between conversational grounding and norm grounding, their role in language acquisition, language maintenance (and change), meaning debates, etc.
 - *inter alia* Larsson & Myrendal 2017, Noble 2019, ...
 - Schlangen 2016, Attari et al. 2019; Rauhut et al. *forth.*
- let us be clear on what language-producing machines can and cannot be: instruments (with product liability), yes;
carriers of normative status, no

some Qs, and (fewer) As

- What if we increase accuracy to 100%?
 - Unlikely; but in any case, you'd still be measuring against whatever that one dataset implicitly encodes. Total authority (very few domains).
- Is that not just rule-conformance vs rule-following?
Reasons vs causes? Etc.
 - Maybe. For me, *norm* carries an additional force (& allows for implicitness), but that needs to be worked out in more detail...
 - Also, the *participation* (interaction, dialectic) part is less often discussed.
- Why not (Lewis-style) *convention*?
 - See above. The Lewis formulation doesn't seem to capture the right type of force. This is more than coordination?

some Qs, and (fewer) As

- Are you saying that inferences like “all tigers are animals” are purely convention, norms?
 - No. Being able to make that inference if you are a user of the word “tiger” is a normative expectation.
- Are you saying that all true inferences connected to a concept should by default (by norm) be known?
 - No. But some (currently thought to be true) ones should. Which those are can change over time. And expectations can differ. Hence, *participation*.
- Are you saying that all meaning making rests on norms? Where is the creativity?
 - Jein. There’s also norm *exploitation*, which bases ad-hoc inferences on norms and allows for creativity.

some Qs, and (fewer) As

- So, a normativist, Sellars-type conceptual role semantics. What's new?
 - Philosophers have, by and large, not bothered to look into the exact methods in which norms are created, enforced, adapted in and through dialogue. An invitation to do so.
- What does this have to do with AI technology, again?
 - It has to do with the use of AI technology. If you're claiming that the output is meaningful language, it has to follow the rules for meaning making. (Explainability, correctability, etc.)
 - (A different talk: How to ensure that the socio-technological system is set up in the right way for “language instruments”, as opposed to “language users”.)

some Qs, and (fewer) As

- What does it mean to be committed through language use?
 - It means being able to suffer the consequences of making a false statement, for example. Which can vary from being corrected to being punished.
 - If you sell a language-producing NLP product, you are getting committed by it.
- Is this not overly rationalistic? Few people want to (or can) debate meanings, or let debates be settled.
 - *If I have exhausted the justifications, I have reached bedrock and my spade is turned. Then I am inclined to say: 'This is simply what I do.'*
PU §217

some Qs, and (fewer) As

- So if I train a perfect „dialectic transformer” that does this kind of meta-linguistic interaction, does that then mean that I have built “an AI”?
 - You have something that plays our game of meaning making better. Whether we should let it play is another question.
 - (My hunch: It is not possible to learn meta-linguistic interaction from observational data.)
- How should I even evaluate my system if not via accuracy? Give me something practical.
 - Benchmarks that systematically test whole conceptual domains.

Goals

- Negative Goal: Argue against “grounded language”.
- Positive Goals:
 - Argue for richer understanding of grounding, & emphasis on investigating & modelling certain dialogical skills.
 - Contribute to clearer understanding of conditions of use of NLP technology.

Thank you.

Questions, Comments?

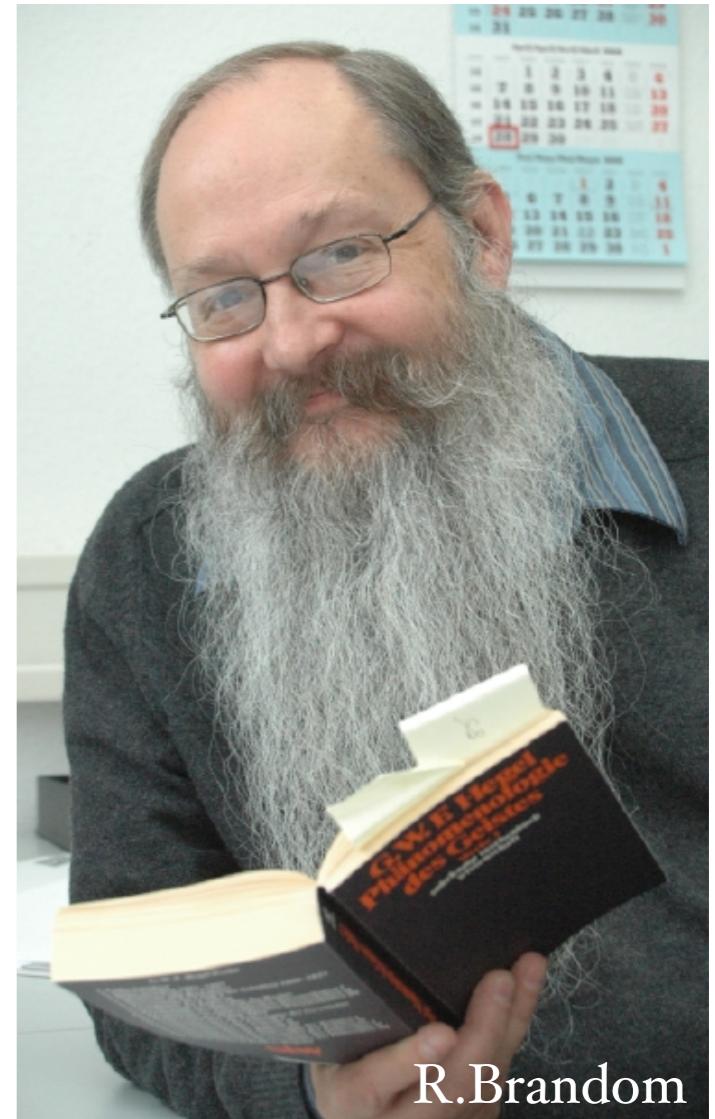
Appeal to authority



G.W.F. Hegel



W. Sellars



R. Brandom



Philosophische Bibliothek

G.W.F. Hegel
Phänomenologie des Geistes

SOME REFLECTIONS ON LANGUAGE GAMES

WILFRID SELLARS

1. It seems plausible to say that a language is a system of expressions the use of which is subject to certain rules. It would seem, then, that learning to use a language is learning to obey the rules for it. If this is what learning a language is taken as it stands, this thesis is subject to an obvious and devastating refutation. After formulating this refutation, I shall turn to the constructive task of attempting to formulate a thesis which is not subject to this refutation. This will involve certain distinctions the theoretical elaboration of which will, I believe, yield new insight into the psychology of language and of what might be called "norm conforming behavior" generally. The present paper contains an initial attempt along these lines.

2. The present thesis runs as follows:

"*L* is learning to obey the rules for *L* if and only if *L* is learning to obey the rules for *L*—in other words

—a sentence in a metalanguage.

Consequently, learning to obey the rules for *L* presupposes the ability to use the metalanguage (*ML*) in which the rules for *L* are formulated.

So that learning to use a language (*L*) presupposes having learned to use a language (*ML*). And so the thesis taken, having regard to the *ML* presupposition, has to learn to use a meta-language (*ML*) and so on.

But this is impossible (a vicious regress).

Therefore, the thesis is absurd and must be rejected.

3. Now consider this a simple and straightforward way of preserving the essential claim of the thesis while freeing it from the refutation. It consists in substituting the phrase "learning to *comply* with the rules for *L*" in place of "learning to obey the rules for *L*". The thesis is then that the use of *A* in circumstances *C'* is to be equated simply with "doing *A* when the circumstances *C'*—regardless of how one comes to do it. [It is granted that "complying with" is not the same as "obeying".] The thesis is then that the use of *A* in circumstances *C'* is an instance of an act of stipulation. A person who has the habit of doing *A* in *C* would then be conforming to the above rule even though the idea that he was to do *A* in *C* had never occurred to him. He would be conforming to the rule even though he did *A* in *C* without having the idea that he was to do *A* in *C* or without having the idea that he was to do *A* in *C* or without having the idea that he was to do *A* in *C*.

4. The approach we are considering, after proposing the above definition of "complying with a rule", argues that whereas *obeying* rules involves using the language in which the rules are formulated, *complying* with rules does not, so that whereas the thesis put in terms of obeying rules leads to a vicious regress, it *cannot* do so once the above substitution is made. Learning to use a language (*L*) no longer entails having learned to use the metalanguage (*ML*) nor does

20

LANGUAGE AS THOUGHT AND AS COMMUNICATION

WILFRID SELLARS

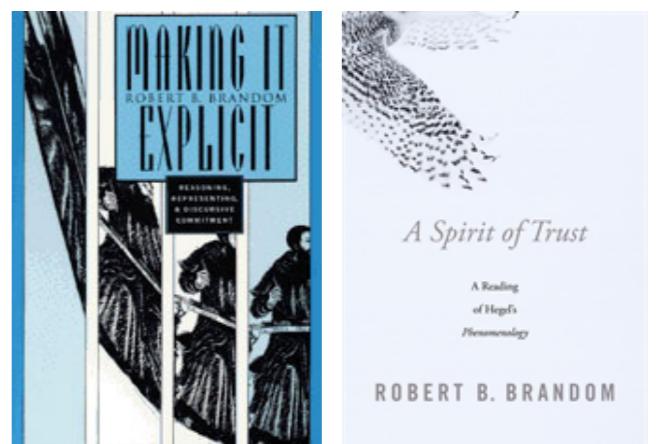
My aim in this paper is to throw light from several directions on the intimate connections which exist between natural thinking and the linguistic behavior which is said to "express" it. The point which I shall ultimately define and defend enough behaviorism in its methodological orientation is not, initial appearances to the contrary, behavioristic in its substantive contentions. It can, nevertheless, be characterized as an attempt to give a naturalistic interpretation of the intentionality of conceptual acts.

The early sections (I-IV) stress the essentially rule-governed character of linguistic behavior. I argue that a proper understanding of the nature and status of linguistic rules is a *modus quod* *modus* of a correct interpretation of the nature of which linguistic behavior can be said to be ruled and not merely to *express* conceptual activity. The second, and larger part of the paper (Sections VII-L) is devoted to exploring the sense (or senses) in which language can be said to "express" thought. A distinction is drawn between three different contexts in which the verb occurs. It is argued that they involve radically different meanings which, if confused, obfuscate the relation between language as conceptual act and language as means of communication, and preclude the possibility of an adequate philosophy of language.

1

There are many interesting questions about the exact meaning or meanings of the term "rule" in non-philosophical contexts. What, for example, is the difference between a "rule" and a "principle"? Are principles simple "rules" in that they are not special applications of more general rules? Or are the two terms different in some other way? Or are principles rules for choosing rules? Is the principle of induction, for example, a higher order rule for choosing law-like statements, themselves construed as extra-logical rules of inference? Though these ques-

506



A Spirit of Trust
A Reading
of Hegel's
Phenomenology

ROBERT B. BRANDON

<https://sites.pitt.edu/~rbrandom/>

norm participation



I think it's a hairy woodpecker.

Hm. It's a woodpecker, clearly. But I think it's a downy woodpecker. Its beak is pretty short.

Yeah, but I've seen hairy woodpeckers with such short beaks. But I agree it seems pretty small overall.

Let's ask Mary!

norm participation



This is humanity's best invention in the last 5 years!

Murder!! Bill Gates!!
Remote control!!!@@

More on accuracy

- If you have a series like 2 4 6 8 10 12 14 ...
 - you measure *accuracy* by generating more from the sequence, and measuring how many generated items are correct
 - you measure *correctness* by well, in this case, an ability to formulate the rule explicitly