COMMUNICATION & COGNITIVE ARCHITECTURE

Week 1: Methodological Preliminaries

Daniel W. Harris

Intention Recognition

Intention Recognition

(Grice 1957, 1969)

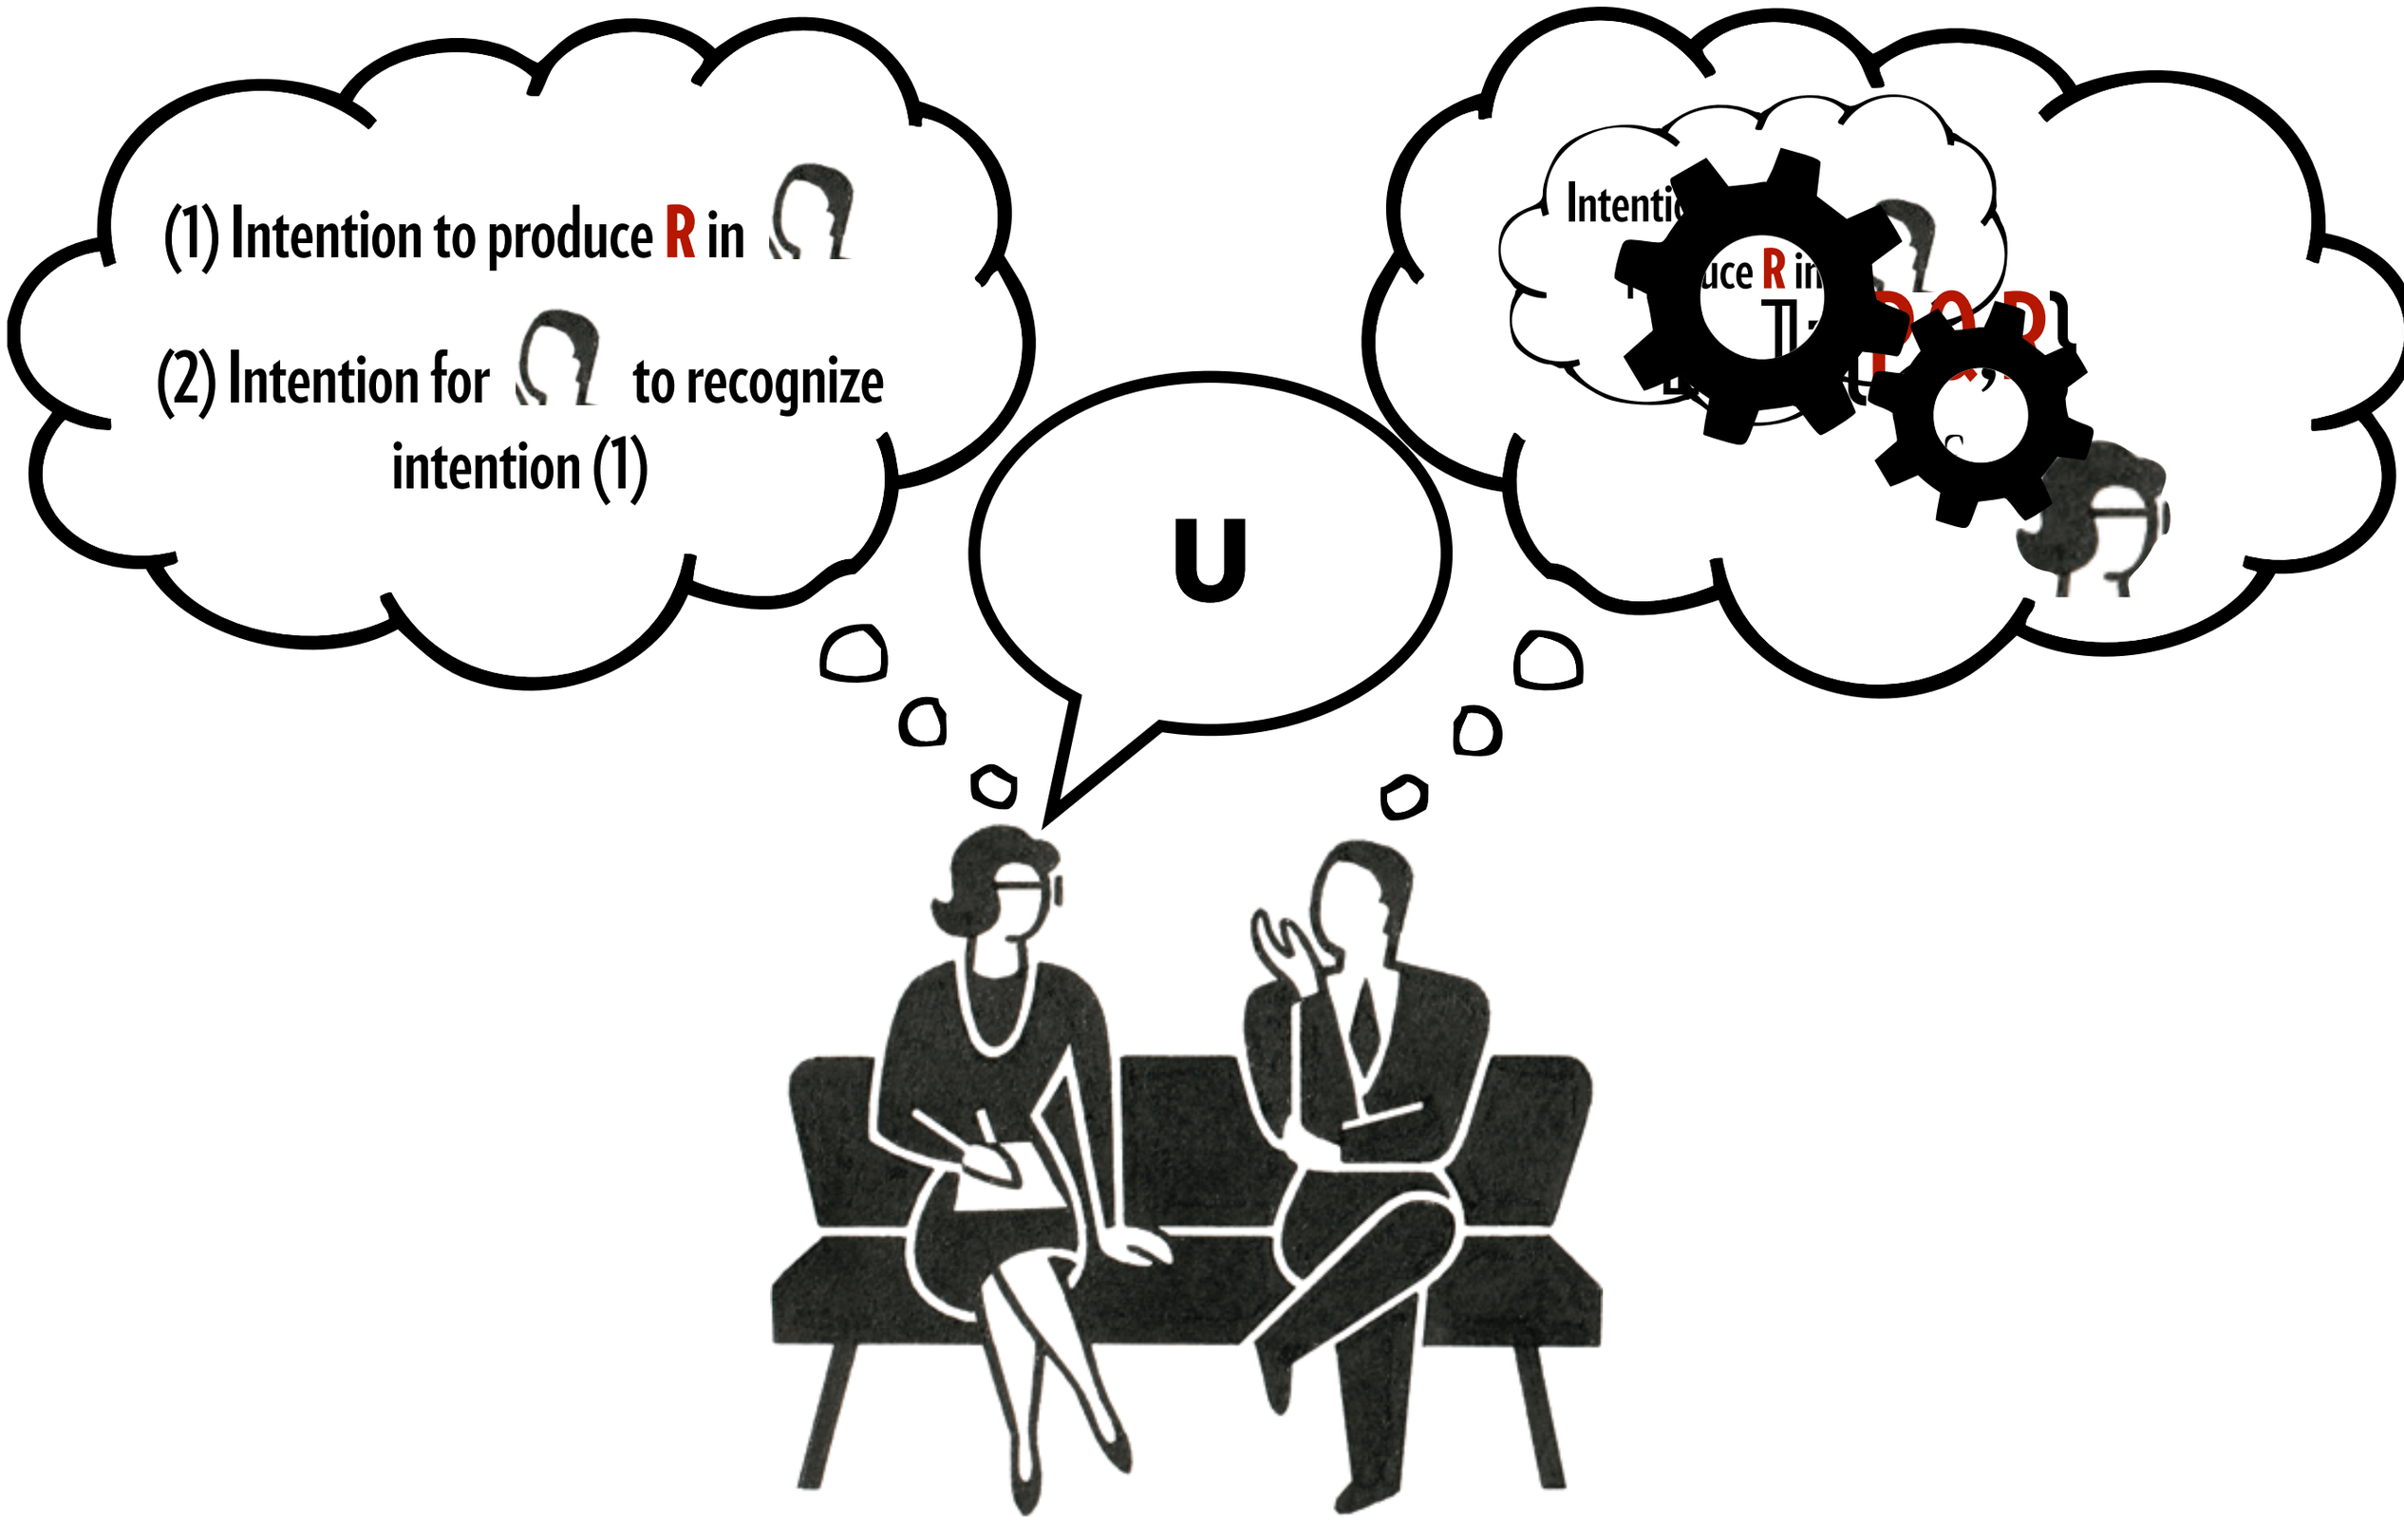

Intention Recognition and Communicative Acts

(Grice 1957, 1968; Strawson 1964; Schiffer 1972; Bach & Harnish 1979; Harris 2014, 2022)

(1) Intention to produce **R** in 

(2) Intention for  to recognize intention (1)

Assertoric Speech Act

R = a belief

Directive Speech Act

R = an intention

Conventional vs. Communicative Acts

(Grice 1968; Strawson 1964; Schiffer 1972; Bach & Harnish 1979; Harris 2029)

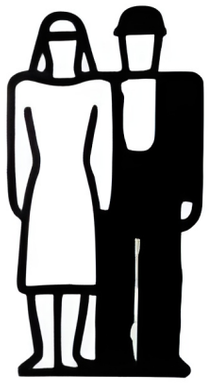

Illocutionary Acts

Conventional Acts

- Invocations of social conventions or rituals.
- E.g. The acts performed in weddings, court proceedings
- Must be performed relative to the jurisdiction of a specific social/ legal regime.
- Lots of cross-culture variability in goals and nature of acts
- Social kinds

Communicative Acts

- Attempts to communicate via intention recognition.
- E.g. assertions, questions, directives.
- Can be performed across social/ legal jurisdictions.
- Cross-cultural variability in the means of performance, but not in the nature of the acts themselves
- Natural kinds

Conceptual Analysis?

Utterer's Occasion Meaning

(Grice 1969)

"U meant something by uttering x" is true iff, for some audience A, U uttered x intending

(i) A to produce a particular response r;

(ii) A to think (recognize) that U intends (i);

(iii) A to fulfill (i) on the basis of fulfilling (ii).

Conceptual Analysis at Work

(Grice 1969)

"I do not think that one would want to say that U had meant something by throwing the banknote out of the window; that he had meant, for example, that A was to (should) go away."

Why Not Conceptual Analysis?

1. It hasn't produced consensus.

Grice (1957)
Strawson (1964)
Searle (1965)
Grice (1968)
Grice (1969)
Ziff (1969)
Searle (1969)
Schiffer (1972)
Harman (1974)
Loar (1976)
Bach & Harnish (1979)
McDowell (1980)
Grice (1982)
Loar (1982)
Schiffer (1982)
Blackburn (1984)
Sperber & Wilson (1986)
Recanati (1986)
Schiffer (1987)
Bach (1988)
Avramides (1989)
Thomason (1990)
Neale (1992)
Davis (1992)
Schiffer (1994)
Davis (2003)
Schiffer (2006)
Roberts (2012)
Scott-Phillips (2014)
Harris (2014)
Buchanan (2018)

•
•
•

Why Not Conceptual Analysis?

1. It hasn't produced consensus.

Why Not Conceptual Analysis?

1. It hasn't produced consensus.

2. All it *could* offer is a recapitulation of our folk theory of communication.

That wouldn't satisfy us in other empirical domains, and it shouldn't here.

Inter-Theoretic Reduction?

Explanation as Inter-Theoretic Reduction

(Nagel 1961)

Law of high-level theory

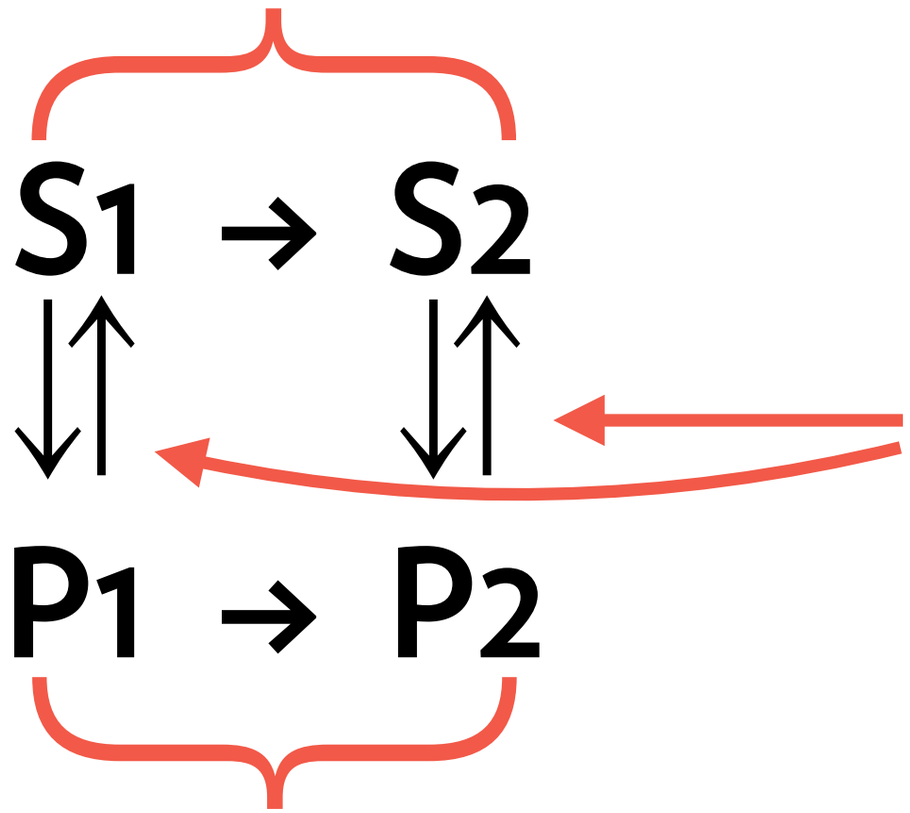

Law of lower-level theory

Bridge laws translating predicates of higher theory into predicates of lower theory.

Result:
lower-level laws entail higher-level laws.

Explication as Inter-Theoretic Reduction

S means something by an utterance

if and only if

S produces the utterance intending:

(1) to produce thereby a certain response R in a certain addressee A;

(2) that A recognize S's intention to produce R;

(3) that A's response be at least partly based on of her recognition of S's intention to produce it.

A bridge principle linking the vocabulary of pragmatics to the vocabulary of folk psychology.

Loar (1982): *Mind and Meaning*
Schiffer (1982): 'Intention-Based Semantics',

Explication as Inter-Theoretic Reduction

S means something by an utterance

if and only if

S produces the utterance intending:

(1) to produce thereby a certain response R in a certain addressee A;

(2) that A recognize S's intention to produce R;

(3) that A's response be at least partly based on of her recognition of S's intention to produce it.

This is an empirical hypothesis…

…whose content is ultimately cashed by this.

Loar (1982): *Mind and Meaning*
Schiffer (1982): 'Intention-Based Semantics',

Inter-Theoretic Reduction

What are the laws of pragmatics?

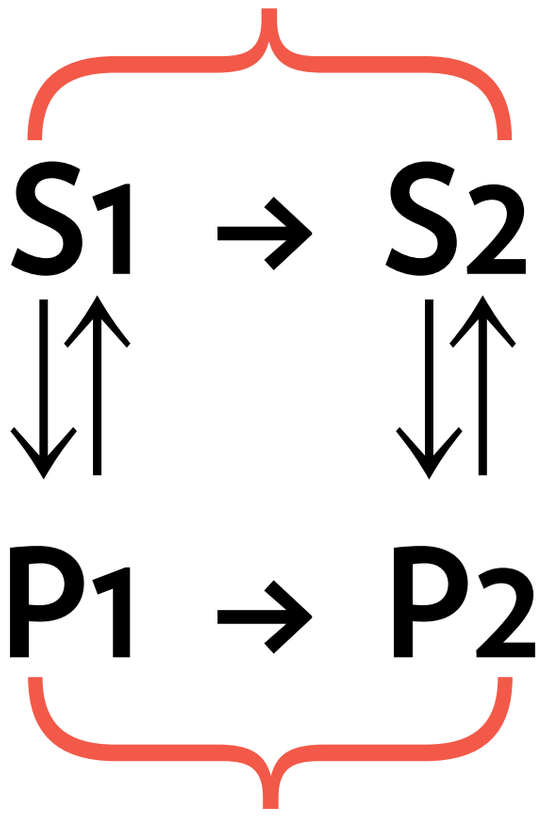

What are the laws of folk psychology, or, for that matter, of neuroscience?

The How vs. The Why

At best Grice taught us something about **how** we communicate.

But **why** do we do it that way?

Some Non-Gricean Theories of Communication

- What you mean by a signal is a matter of the signal type's **proper function** (Millikan, Skyrms)

- What you mean is a matter of the state of mind that you **express** (Green, Bar-On)

- What you mean is a matter of the public **commitment** that you undertake (Brandom, Geurts, Krifka)

- What you mean is a matter of **linguistic conventions** (Austin, Searle, Lepore & Stone, Stojnić)

Mechanistic Explanation

"It is false that [biology and psychology] are organized around laws. In particular, it is false that the usual form theoretical knowledge takes is a set of forward-looking law-like causal principles that directly describe real systems. Laws appear occasionally, but they are minor players, with none of the organizing role they play in physics."

—Peter Godfrey-Smith, "Reduction in Real Life"

"When we look at successful reductionist research programs in areas like biology, we do see an accumulation of information about how various biologically important processes occur. We now have a good understanding of processes like photosynthesis, respiration, protein synthesis, the transmission of signals in the brain, the action of muscles, the immune response, and so on. ... We are taking a high-level process or capacity, and explaining how it works in terms of lower-level mechanisms and entities."

—Peter Godfrey-Smith, "Reduction in Real Life"

"A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena."

—William Bechtel (2007): *Mental Mechanisms*


PIZZA
BY THE
SLICE!
YOUR LOGO HERE
HOT · FRESH · FAST
1
2
3
24/7 Pizza
BOX
WORLD'S SMALLEST PIZZERIA™
WORLD'S SMALLEST PIZZERIA

HOW IT WORKS

The base is precooked, toppings are added, and the pizza boxed. The pizza is then placed in cold storage, which holds a maximum of 70 pizzas at a regulated 37.5° Fahrenheit / 3° Celsius.

The robotic plate positions the selected pizza in front of the oven.

The mechanism lifts the lid, pushes the pizza and box in the oven, and then lifts the pizza out of the box. The convection oven heats the base and cooks the toppings in 3 minutes

The computer manages pizzas in the cold storage by optimizing their BBE (Best Before End) dates and cooking time.

Pizzas may have different parameters for temperature, cooking time, and browning power.

After cooking, the pizza is placed in its box and closed. It is delivered to the customer through the dispenser.

Mechanistic Explanation

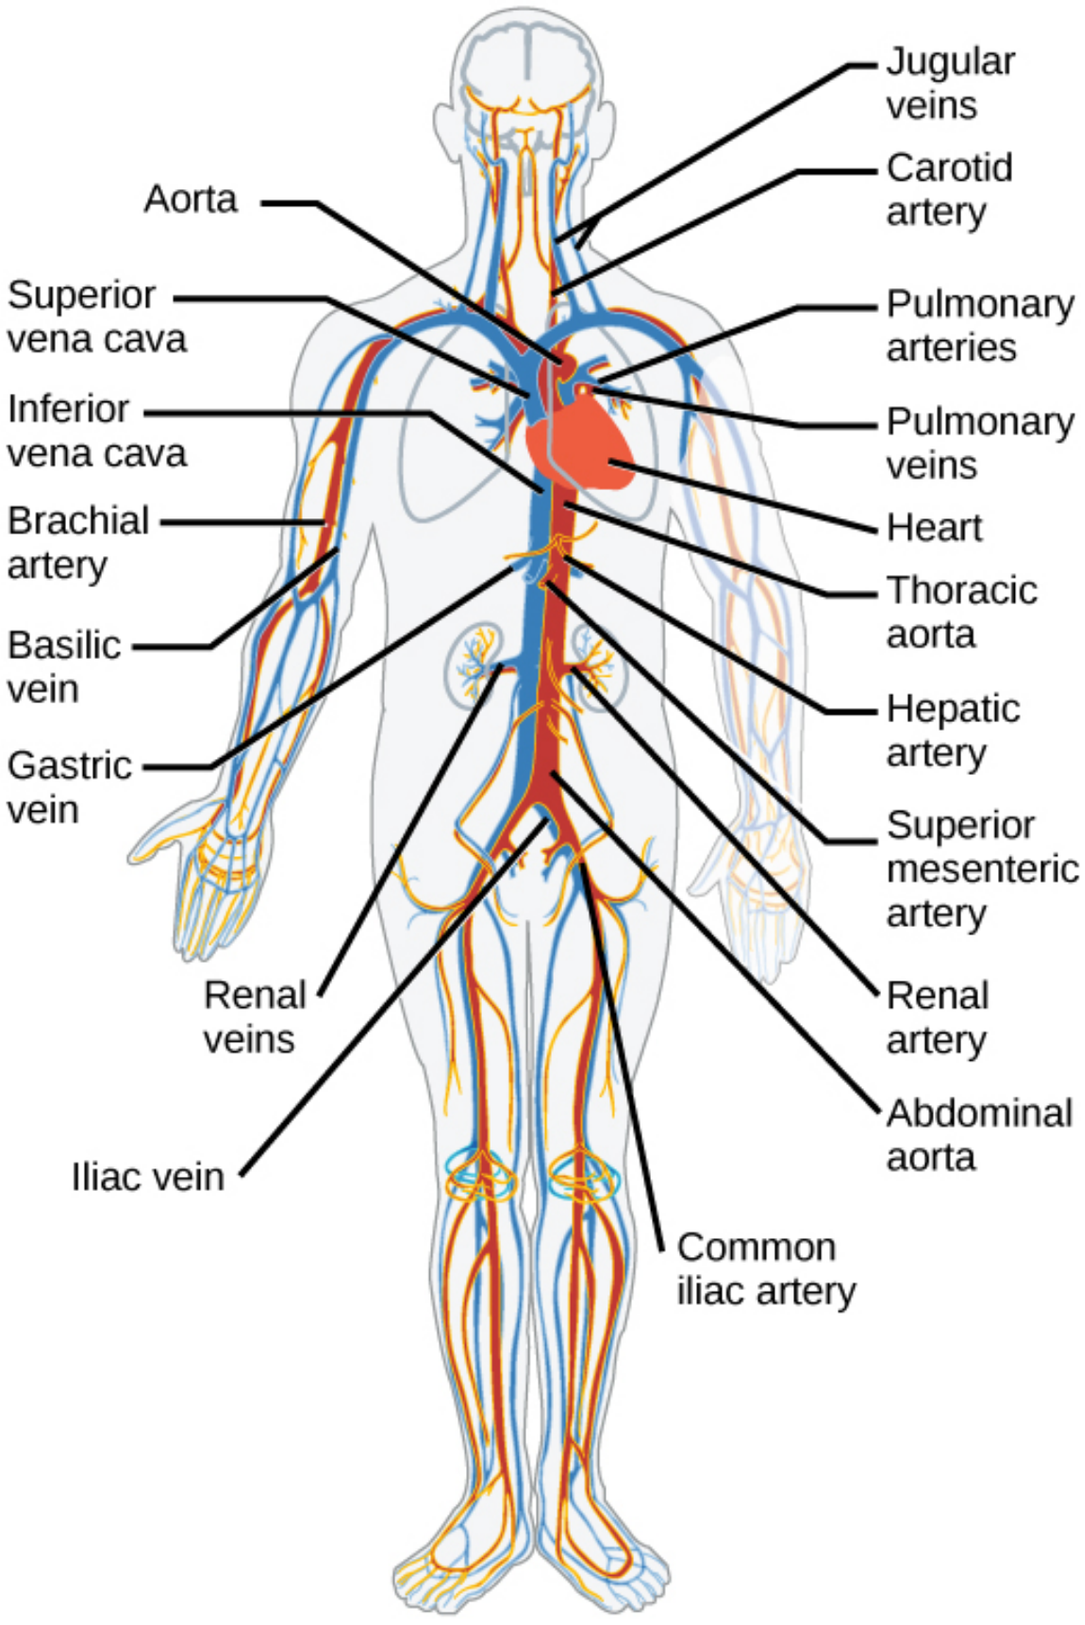

Mechanistic Explanation

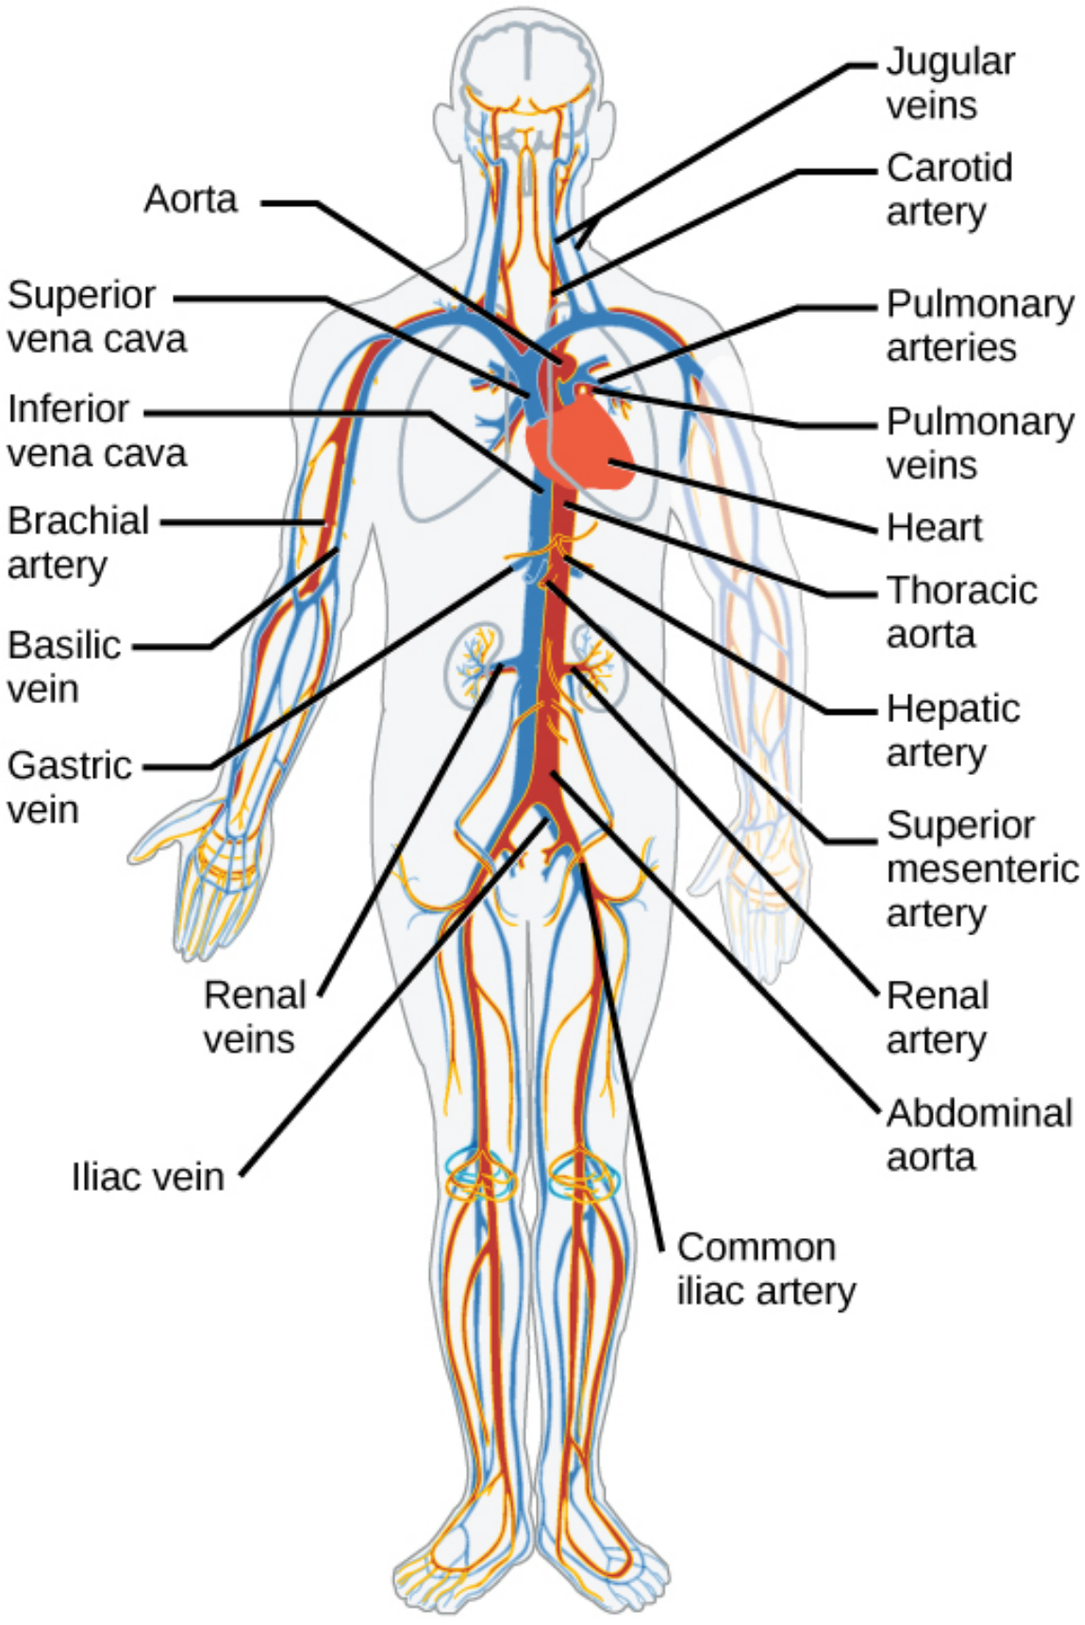

Some ways a heart can fail

<u>Constitutive causes:</u>

- Myocarditis
- Faulty valves

<u>Environmental causes:</u>

- Coronary artery disease
- High blood pressure

Mechanistic Explanation

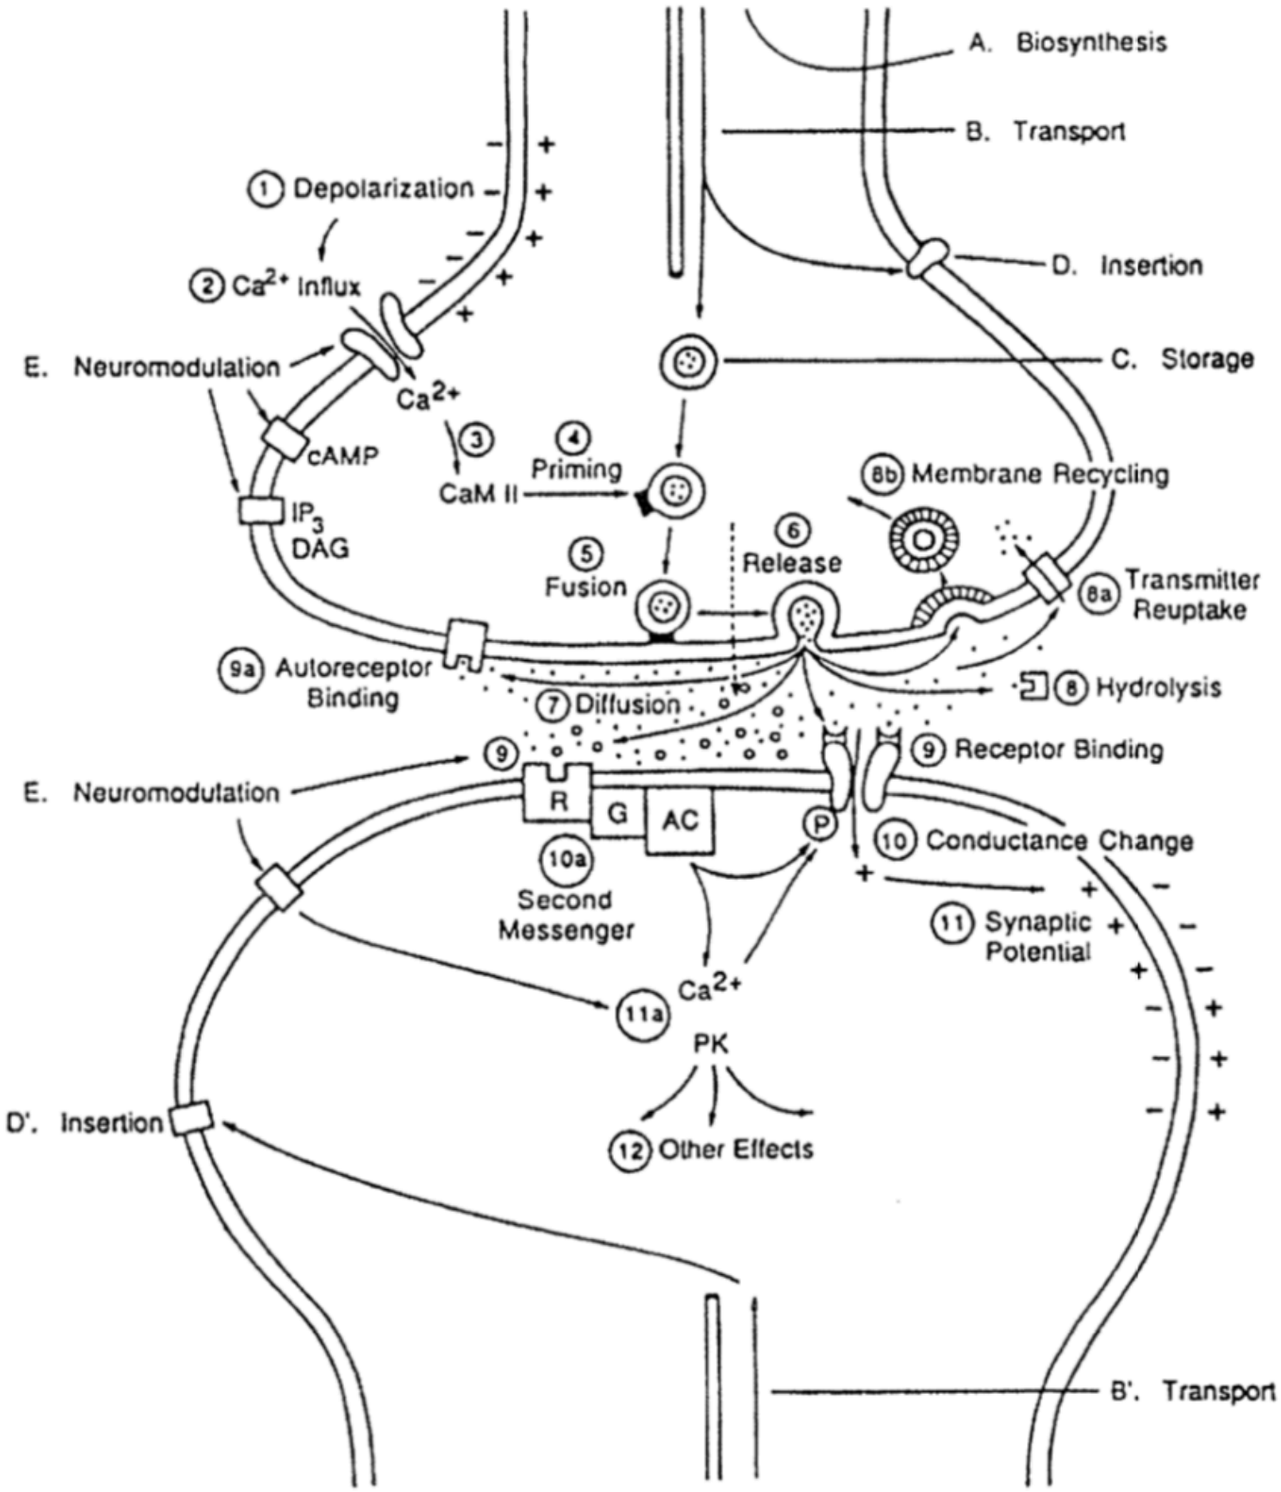

Machamer, Darden, and Craver (2001): 'Thinking about Mechanisms'
see also, Bechtel (2009): *Mental Mechanisms*

NeuroCognitive Mechanisms

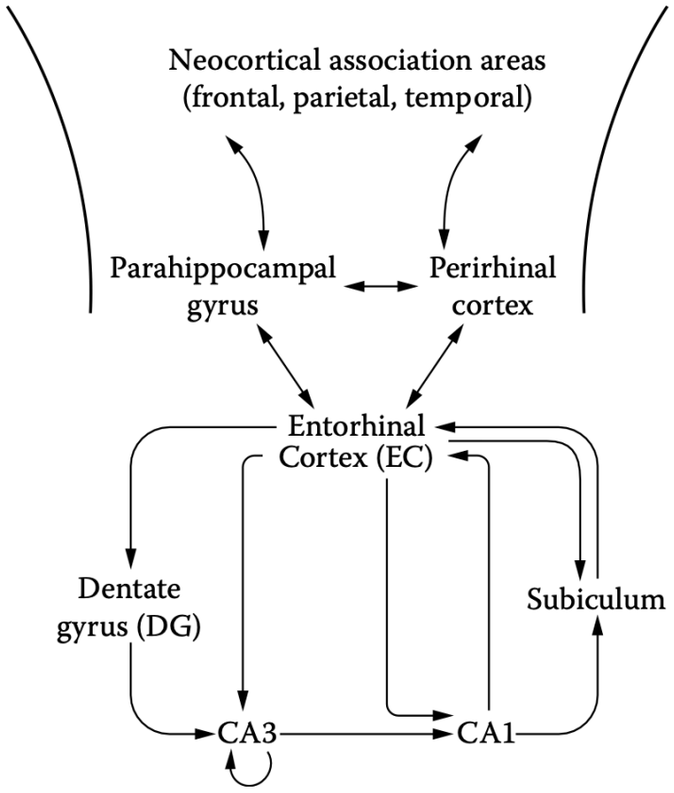

Figure 2.3 Schematic diagram of the hippocampal system. Informatio

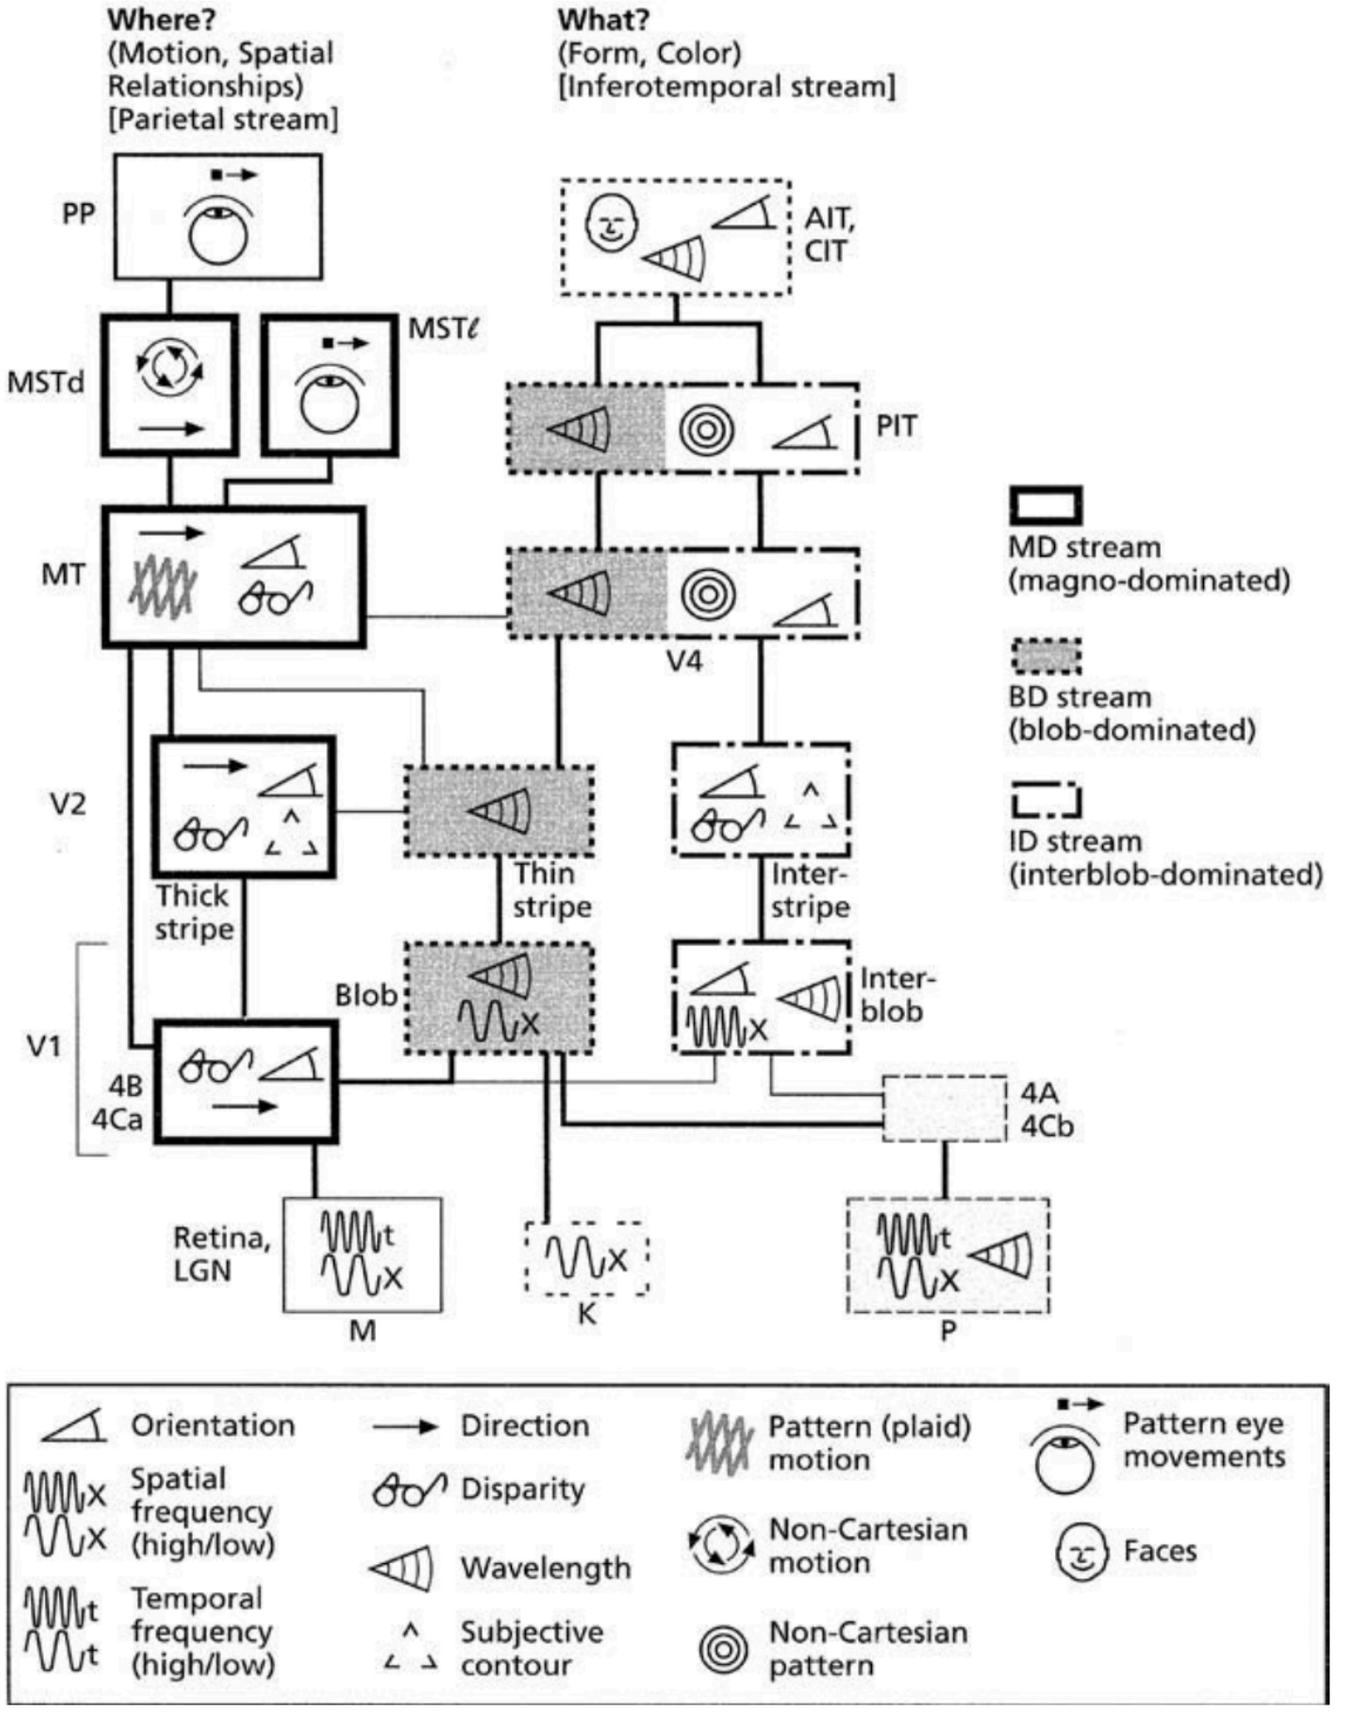

Figure 3.12 Van Essen and Gallant's iconic representation of the organization of multiple processing streams in the macaque, with boxes indicating different processing areas and icons representing the types of information to which cells in each area are responsive.

Bechtel (2007): *Mental Mechanisms*

~~Neuro~~Cognitive Mechanisms?

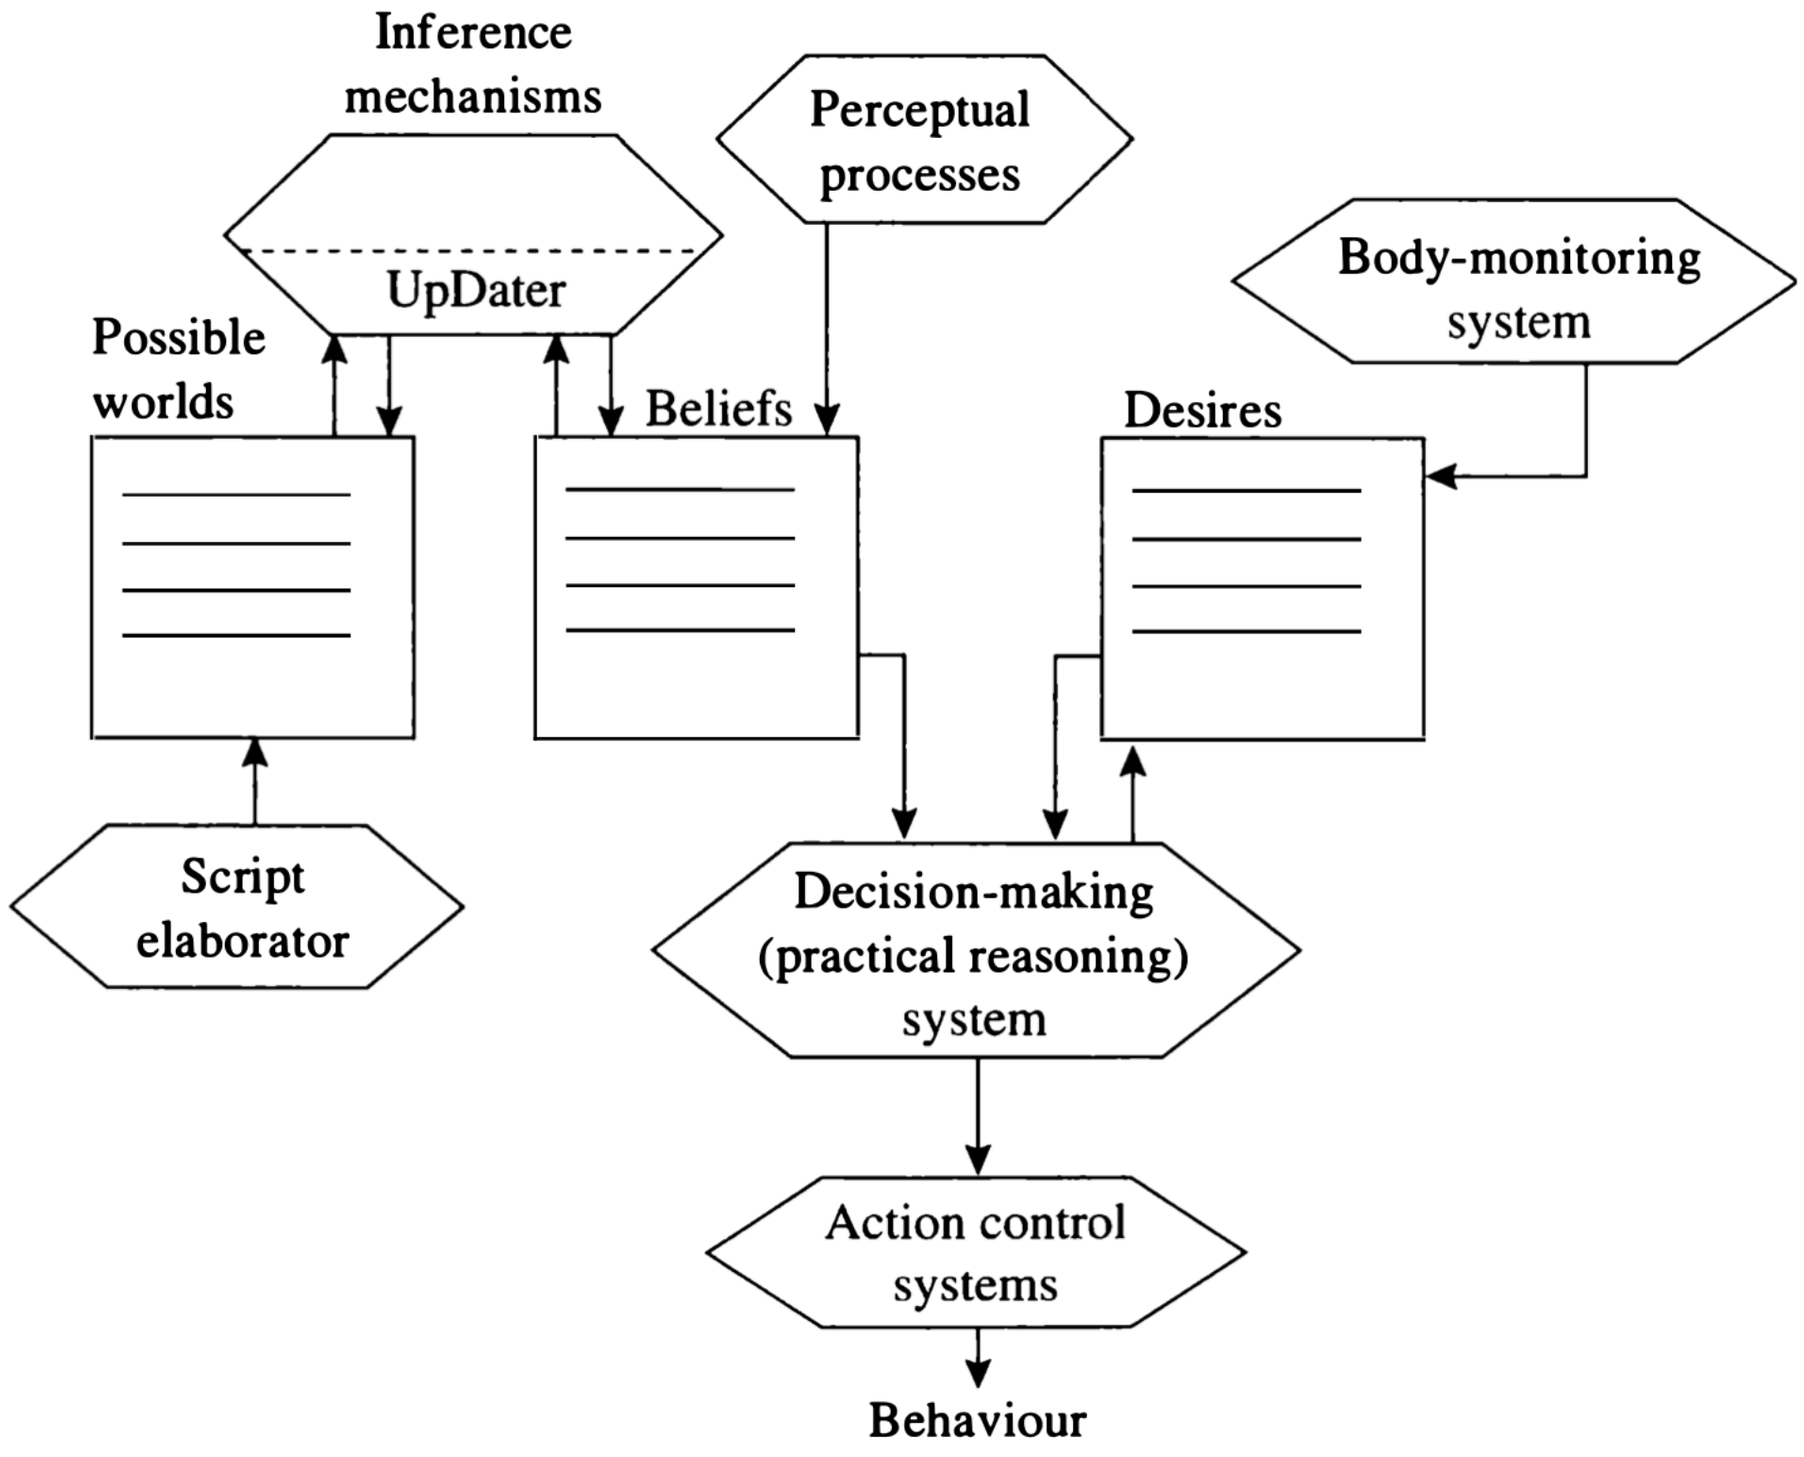

Fig. 2.2 The mental mechanisms posited by our theory of pretence

Nichols and Stich (2003): *Mindreading*

~~Neuro~~Cognitive Mechanisms?

Relevantly Similar Approaches:

- Homuncular functionalism (Dennett 1975)
- Functional analysis (Cummins 1975)
- Modular decomposition (Fodor 1983, Carruthers 2006)

In a slogan: Cognitive Architecture

Mechanistic Explanation of Human Communication?

Mechanistic Explanation of Human Communication?

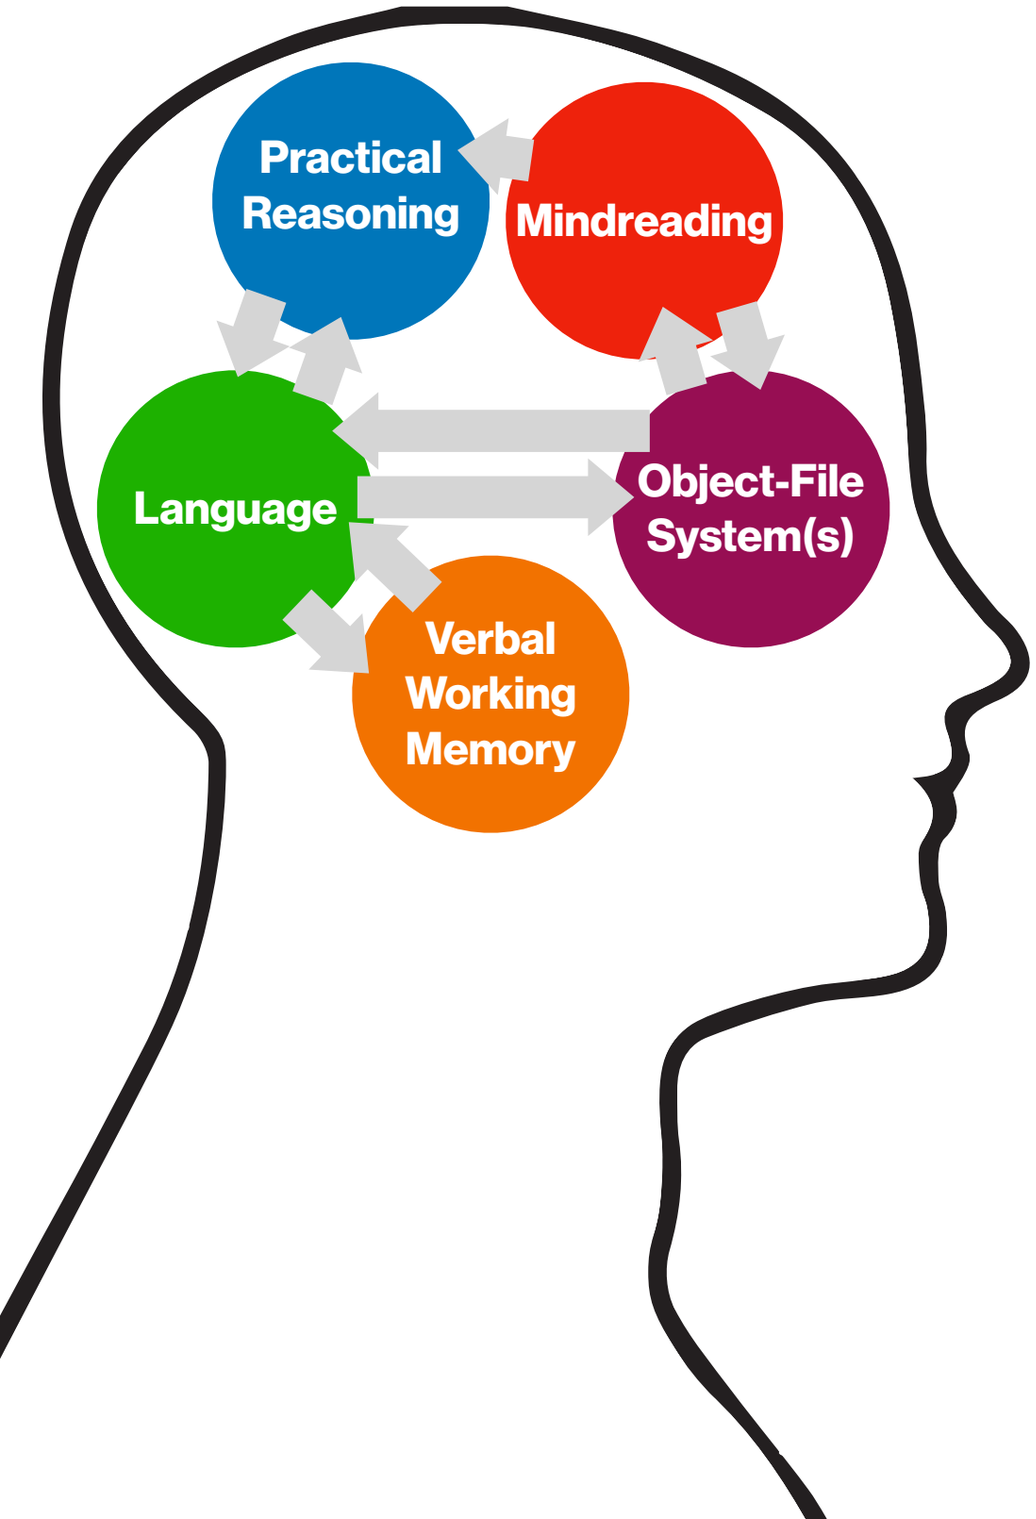

- What do we know about how these systems work?
- How do they interface with each other?
- What are their characteristic patterns of failure?
- What (scarce) resources do they rely on?
- How do they develop in kids?
- How did they evolve in our ancestors?

Mechanistic Explanation of Human Communication?

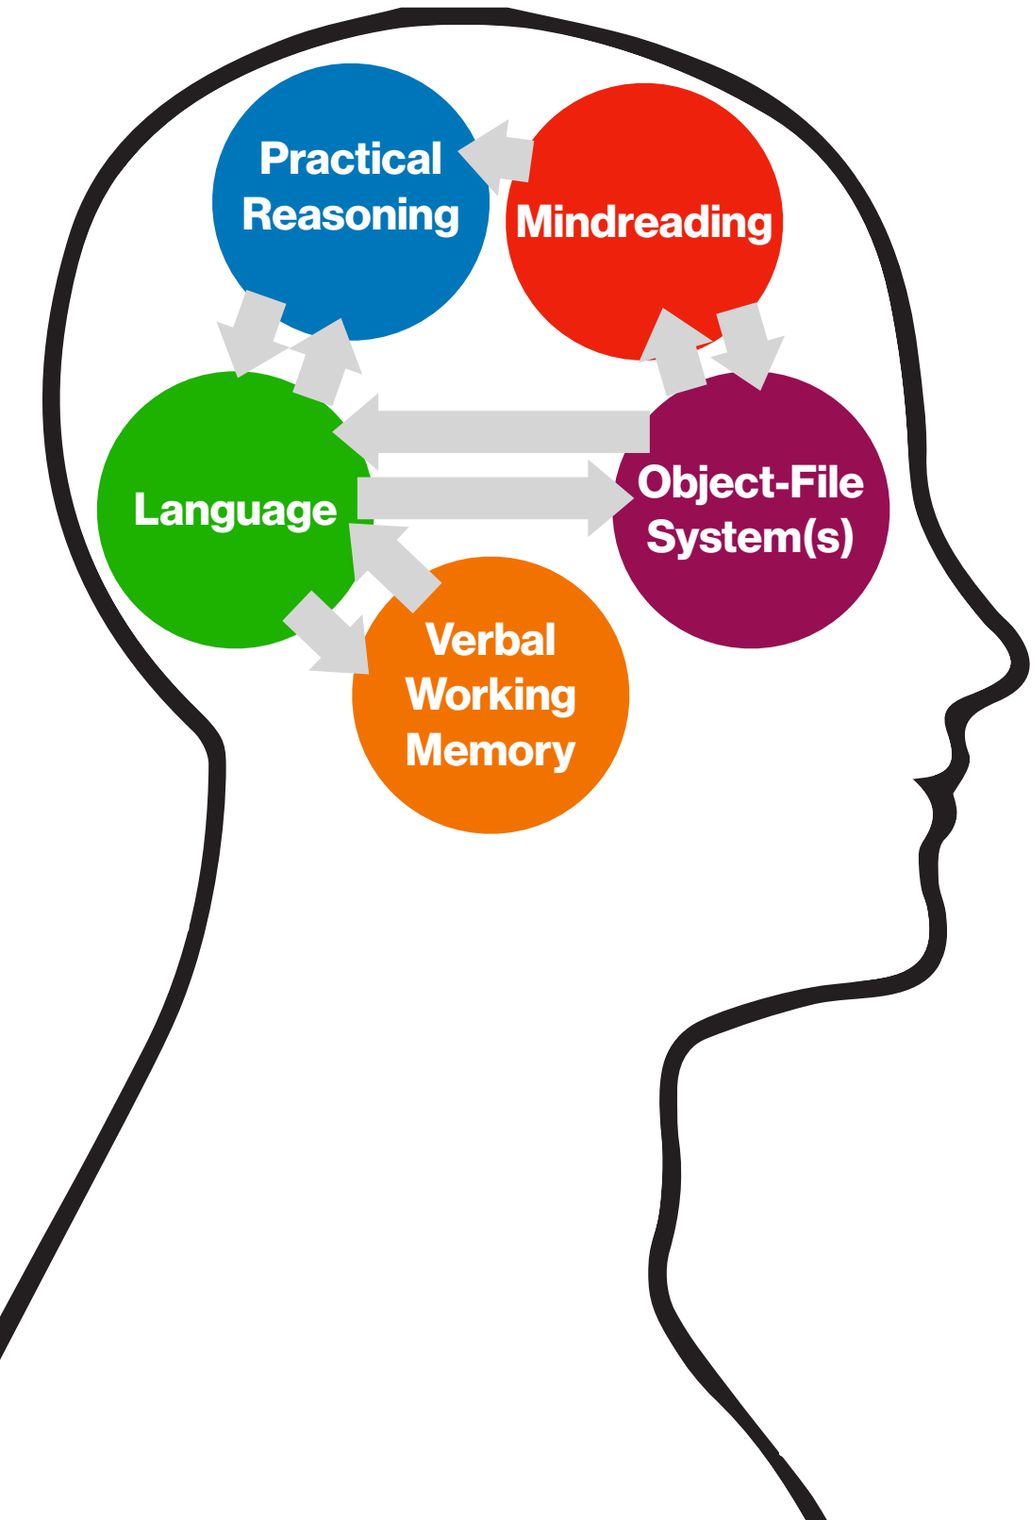

- Linguists, psychologists, biologists, anthropologists have had a lot to say about the individual components.
- But these people are mostly too siloed and busy doing their things to try to synthesize their insights.
- Philosophy opportunity! 🥳
- One worthy goal is to use knowledge from different fields to confirm and constrain our pursuits in others.
- Consilience!

Idealized Models

Idealized Models

"Model-based science features an "indirect" strategy for the representation and investigation of unknown systems. A model-builder first describes a hypothetical structure, usually a relatively simple one, and then considers similarity relations between this structure and the real-world "target" system that he is trying to understand."

—Peter Godfrey-Smith, "Reduction in Real Life"

Idealized Models

- Minimalist Idealization (Strevens, Weisberg)
- Special Cases Idealization (Godfrey-Smith)
- Galilean Idealization (McMullin, Weisberg)

Idealized Models

Minimalist idealizations eliminate all factors that fail to "make a difference" to the target phenomenon being studied.

Idealized Models

"In explaining Boyle's law, for example, theorists often introduce the assumption that gas molecules do not collide with each other. This assumption is false; collisions do occur in low-pressure gases. However, low-pressure gases behave as if there were no collisions. This means that collisions make no difference to the phenomenon and are not included in the canonical explanation. Theorists' explicit introduction of the no-collision assumption is a way of asserting that collisions are actually irrelevant and make no difference."

—Weisberg, 2007, 643

Idealized Models

- The legitimacy of a minimalist idealization is relative to a choice of target phenomena.
- E.g.: idealizing away from low-pressure collisions might be legit when studying the macroscopic behavior of gas.
- But it won't be legit if the goal is to study low-pressure collisions.
- Similarly: if our goal is to understand communication or communicative acts, it is not legit to idealize away from the forms of success that define them.

Idealized Models

Special cases models "directly describe one class of cases, which are simple and tractable, and use these as the basis for a more indirect understanding of the others. Understanding is achieved via similarity relations between the simple cases we have picked apart in detail, and the cloud of more complicated ones."

—Godfrey-Smith, 2009, 4

Idealized Models

- But there is no clear sense in which communication that involves context change is simpler than communication that does not involve context change.

- Indeed, cases of communication involving context change involves extra steps, over and above the simpler form of communication that is also present in non-context-involving cases.

- And mere similarity relations can't help us to understand no-context cases in terms of the context-involving ones.

Idealized Models

Galilean Idealization is "the practice of introducing distortions into theories with the goal of simplifying theories in order to make them computationally tractable."

—Weisberg, 2007, 640

Idealized Models

- Example: Galileo initially ignored air resistance in his model of projectile motion.
- Galilean idealizations are justified pragmatically: they're better than nothing, which is the alternative.
- They are provisional: our theory would be more predictive if we found a way to lift them.

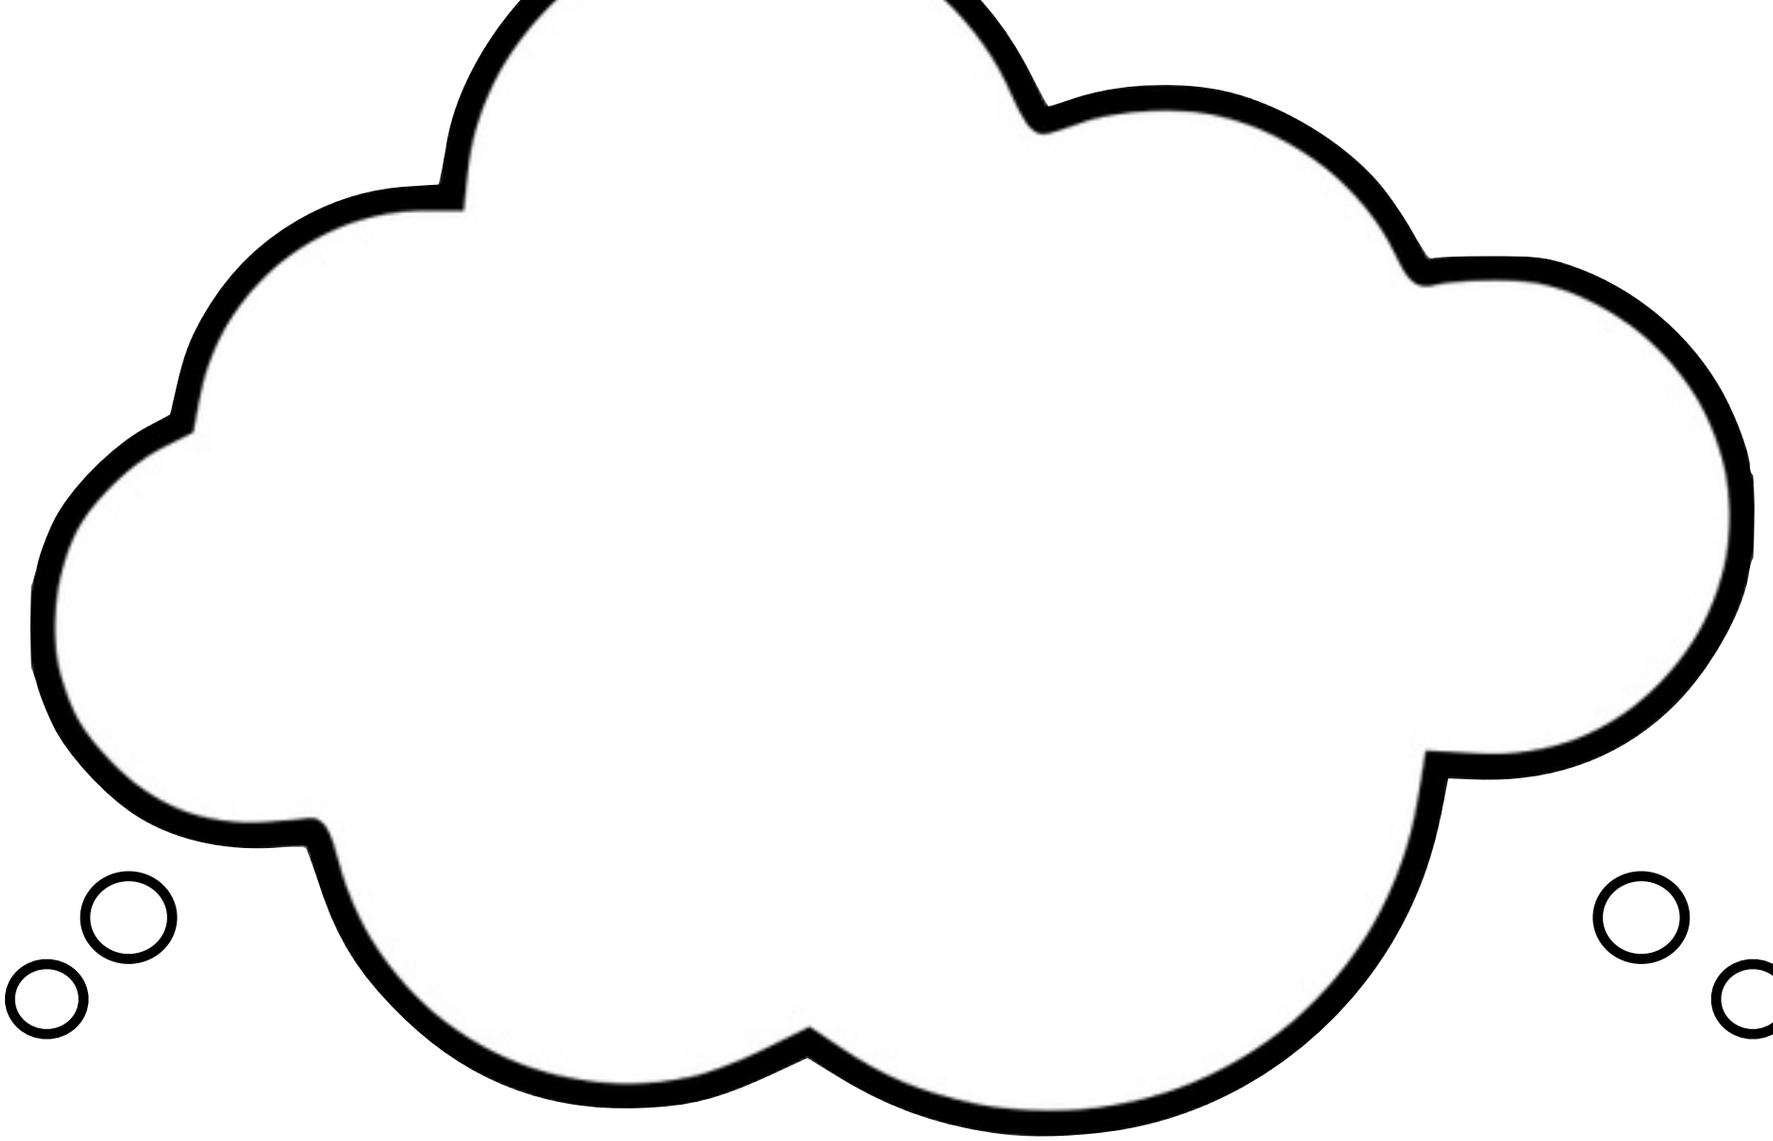

Dynamic Models of Conversation

Stalnaker (1978); Lewis (1979); etc etc

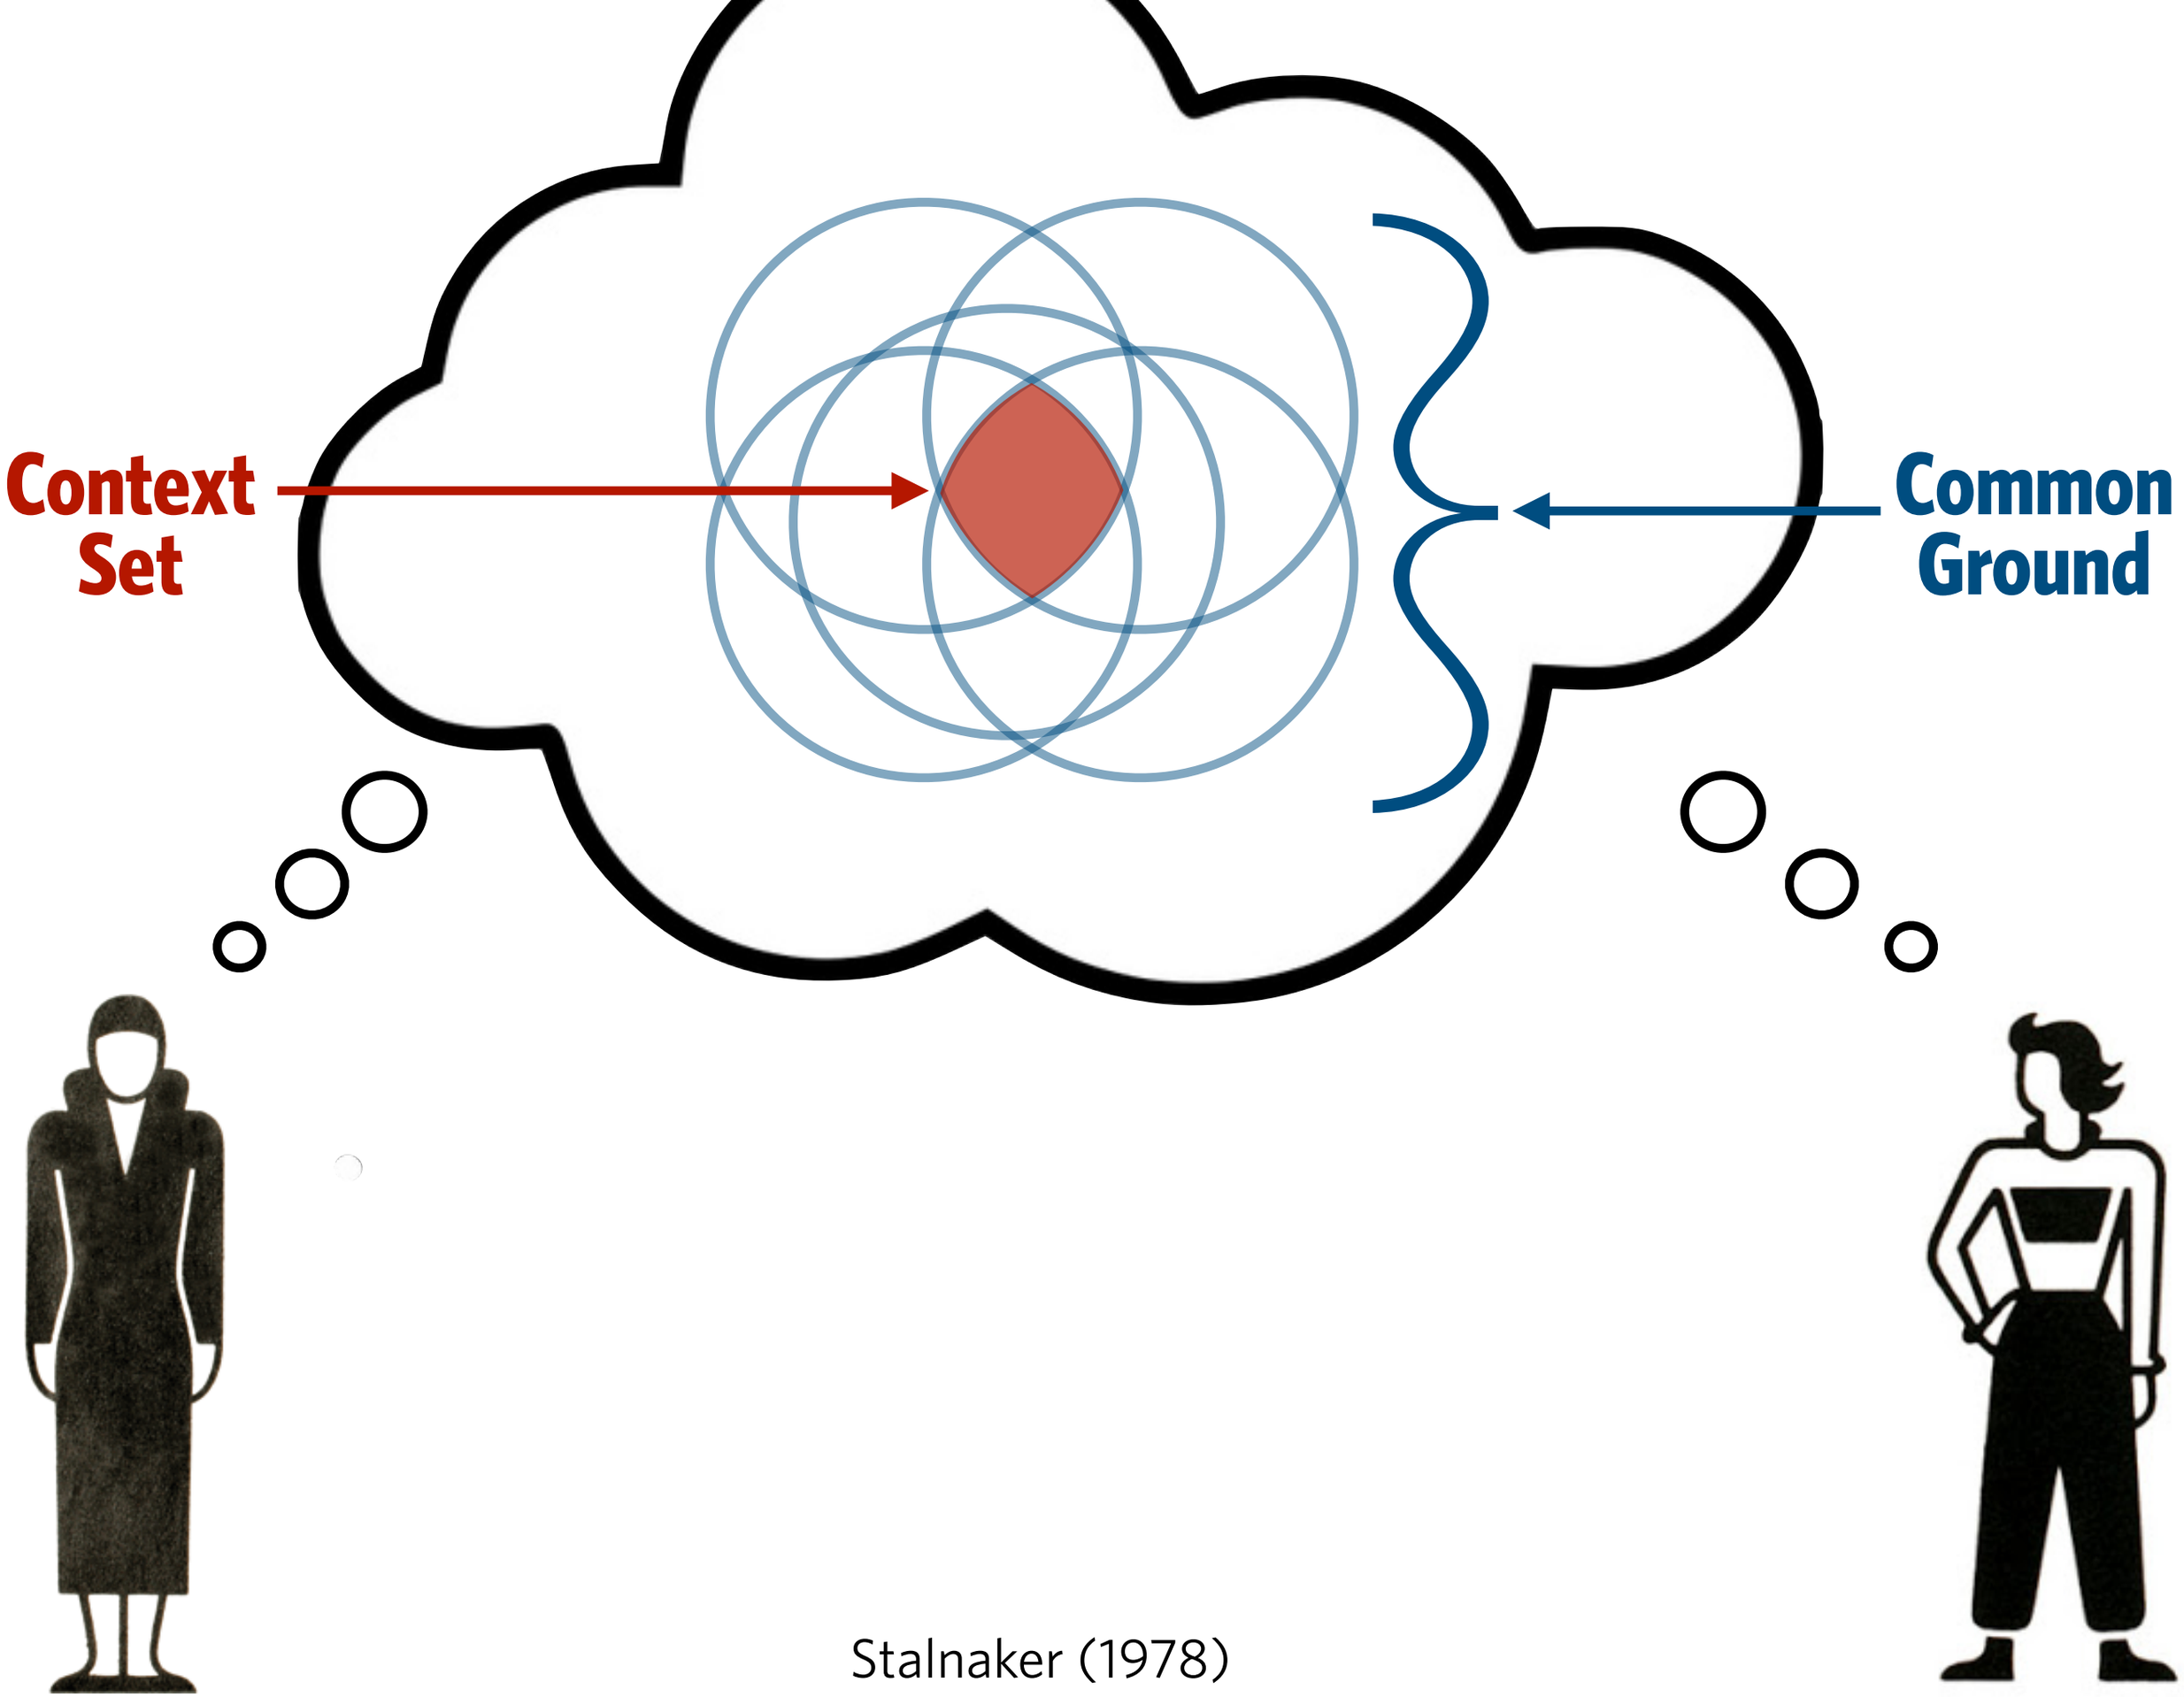

Stalnaker (1978)

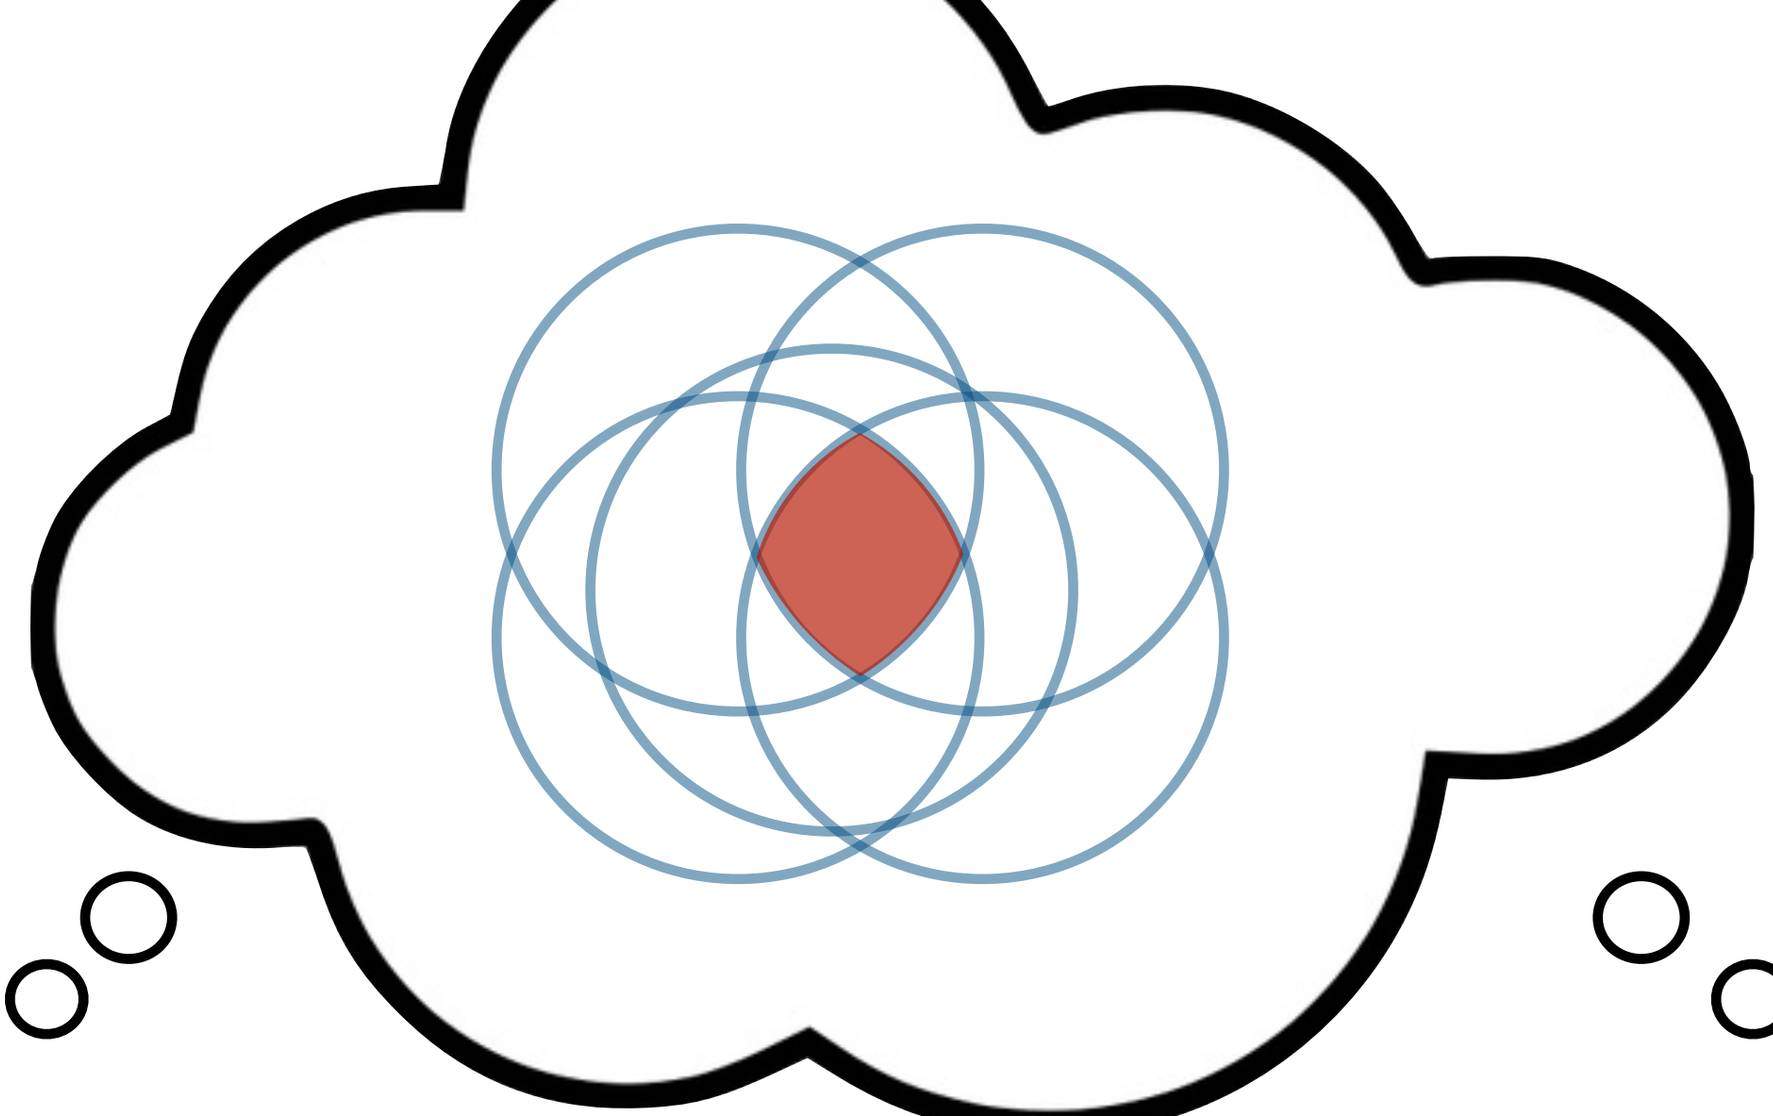

I'm reading a book by Wittgenstein.

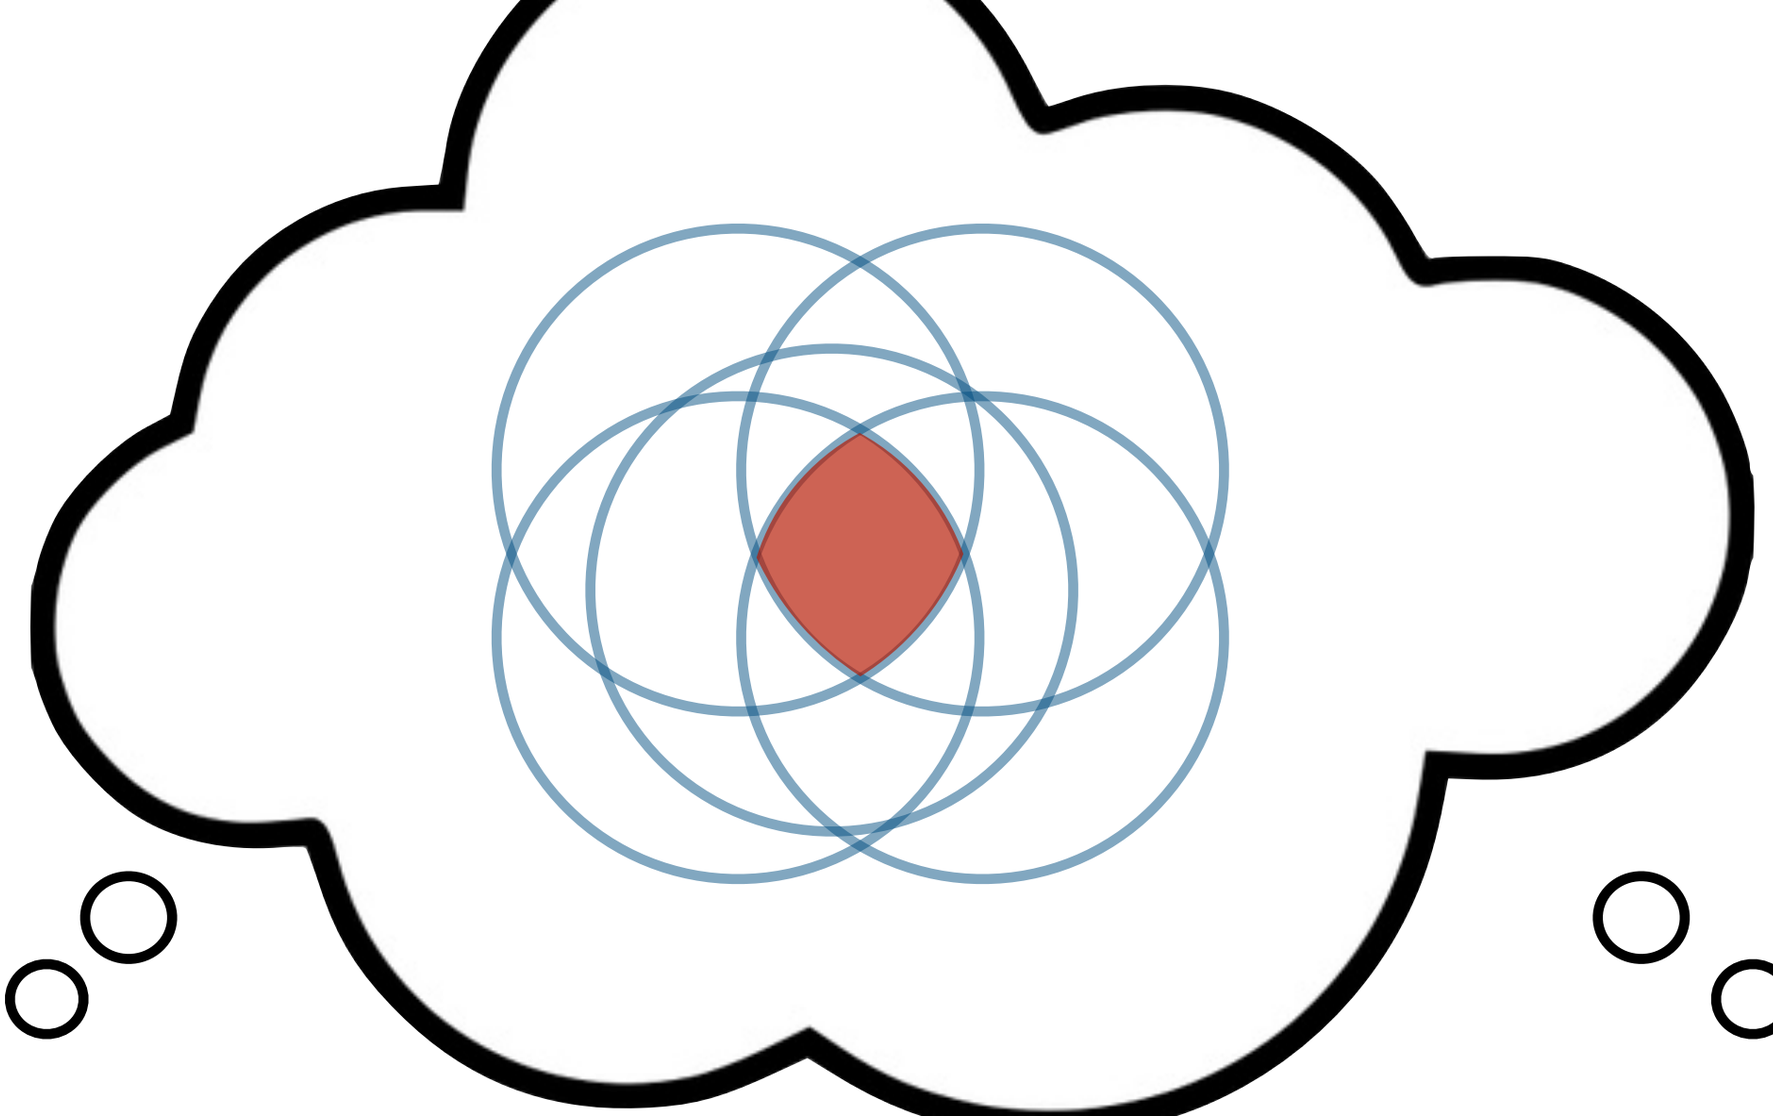

[I'm reading a book by Wittgenstein.] DECODING p

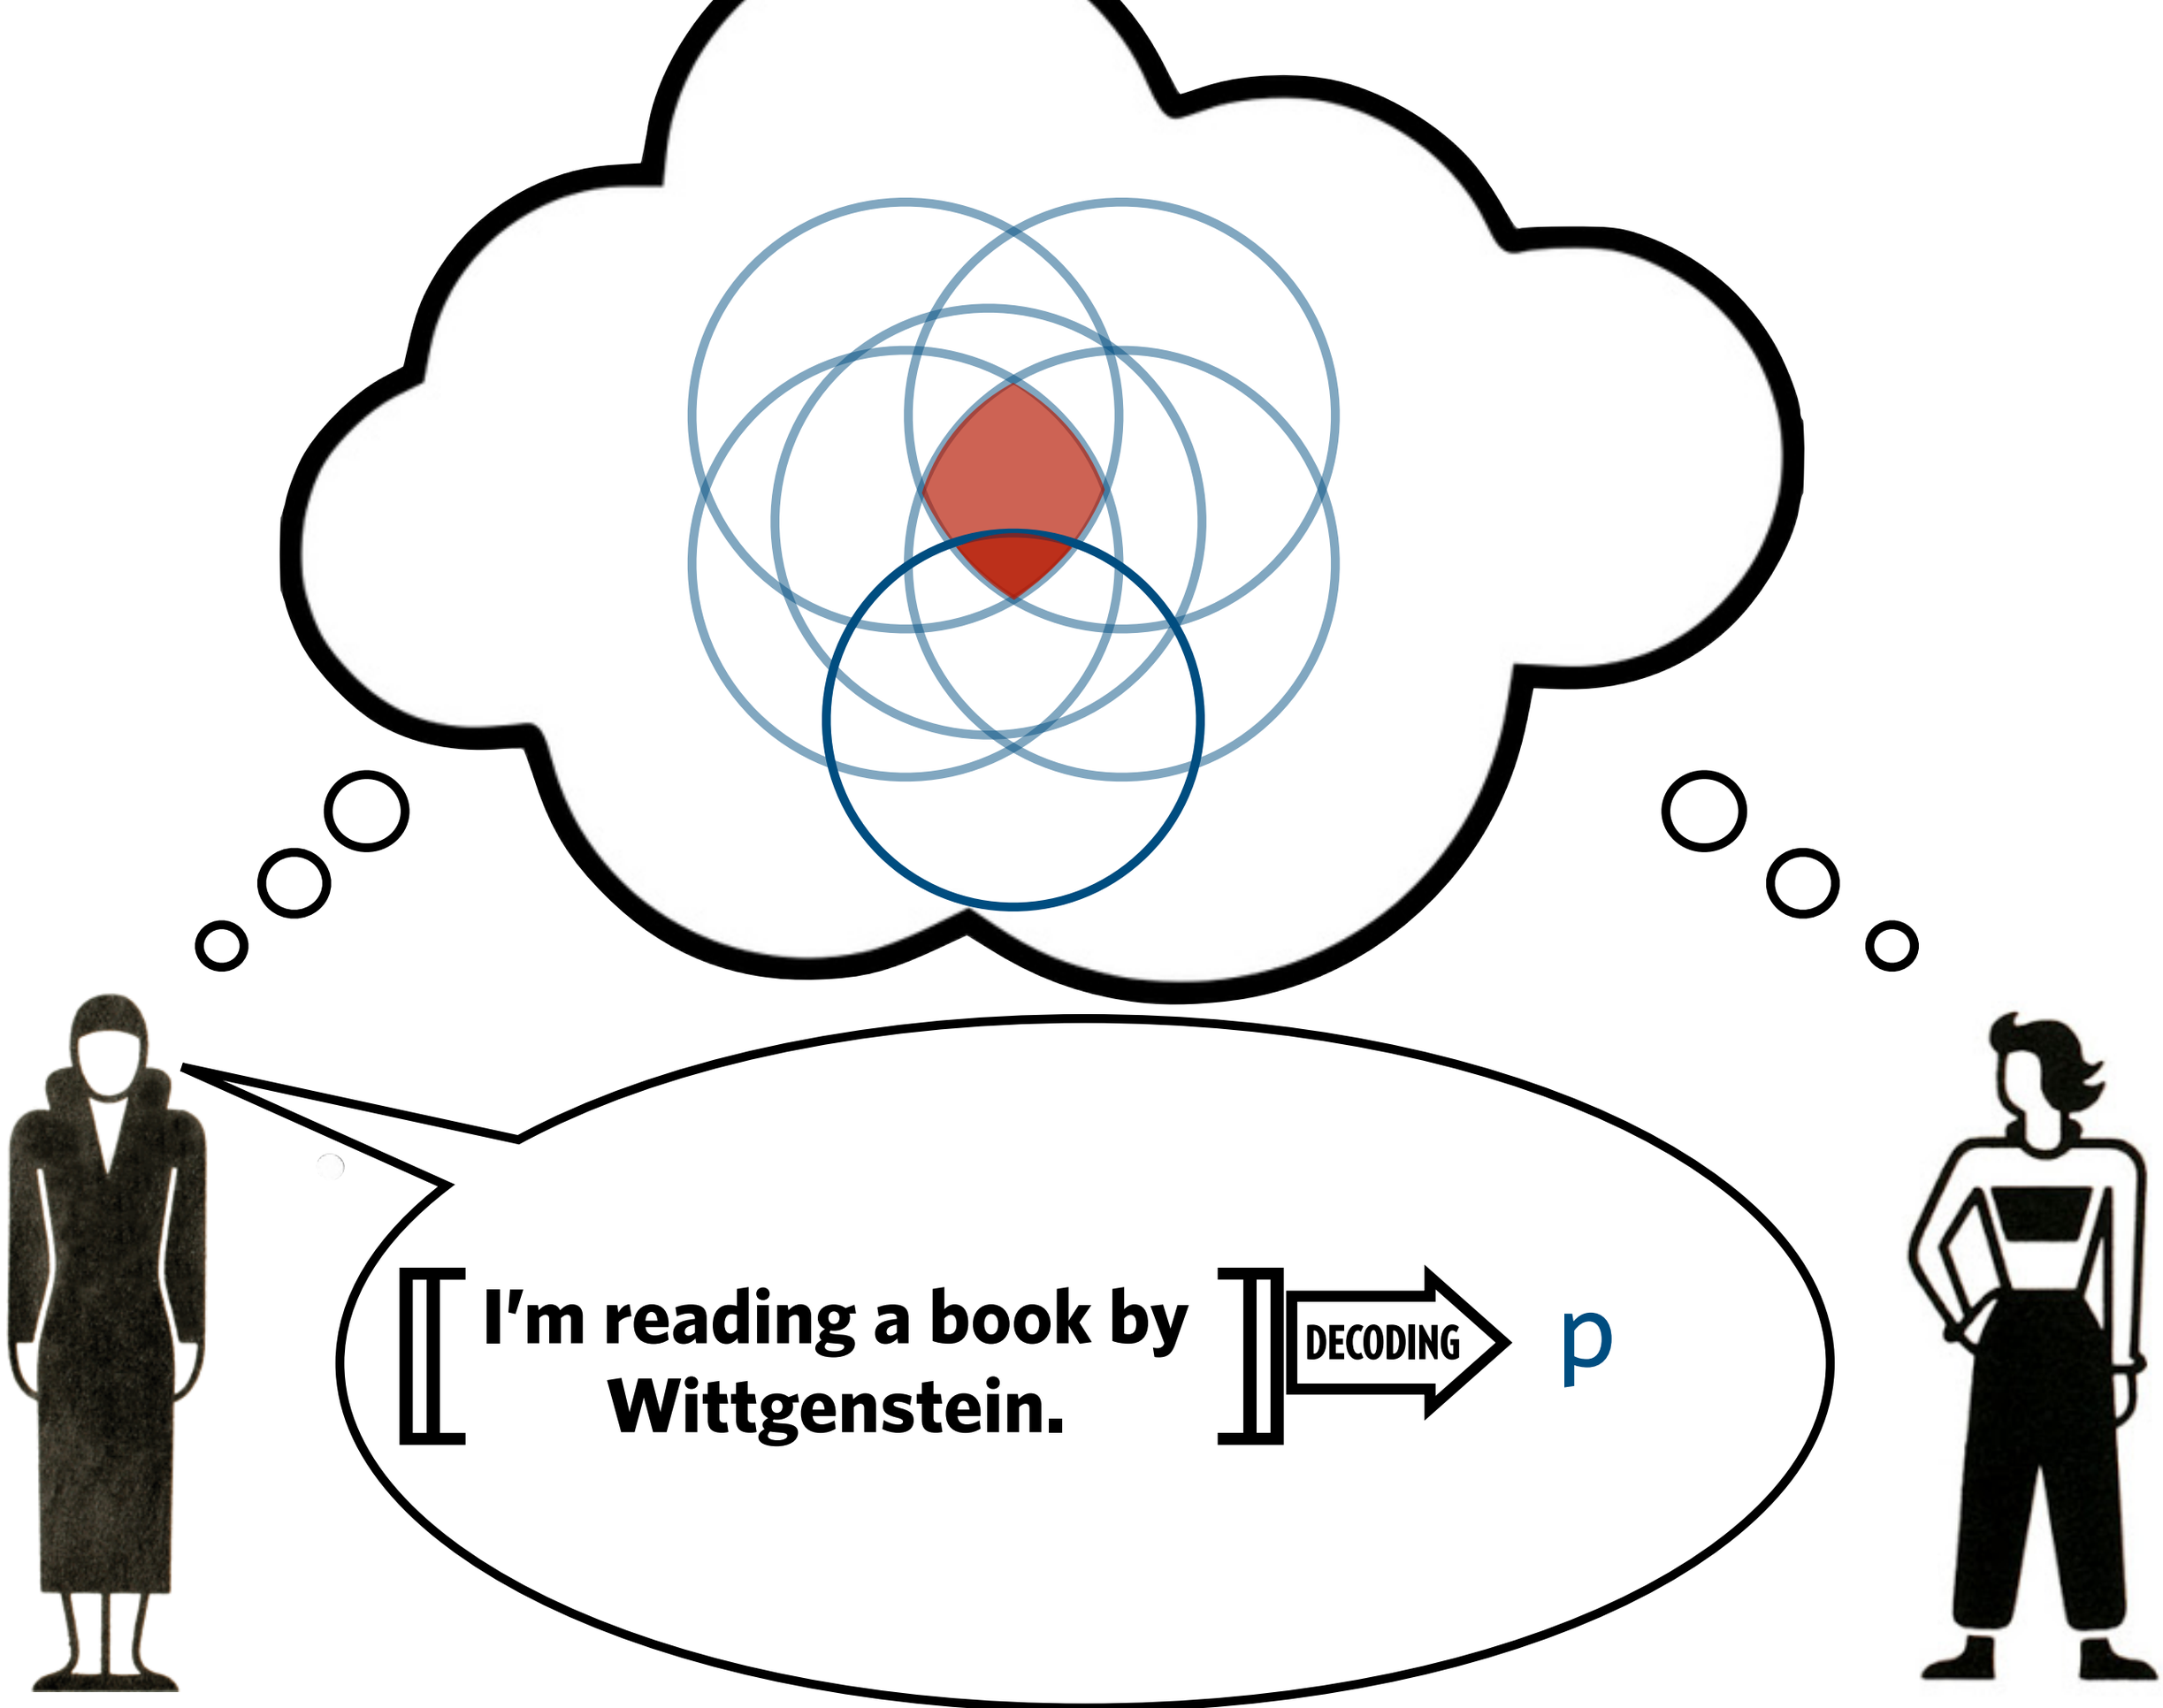
I'm reading a book by Wittgenstein.
DECODING
p

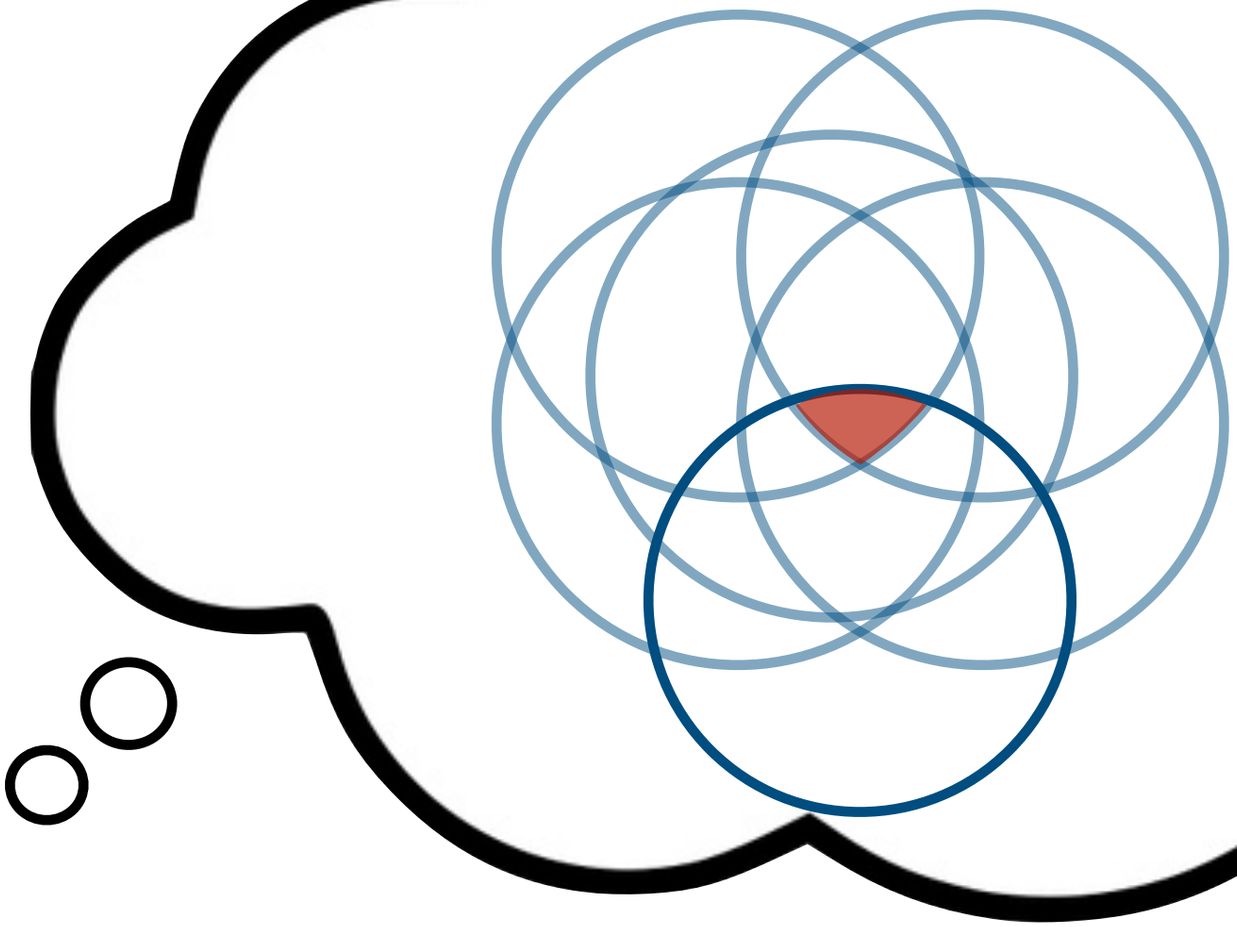

What was his big idea?

[[What was his big idea?]] DECODING q

Idealized Models

- It is plausible to think of dynamic models of conversation as partaking in Galilean idealization.
- The basic idea: these models idealize away from the processes in communication that we don't know how to computationally model: mindreading, planning, non-intersective belief update, etc.
- These processes all involve non-monotonic reasoning (a.k.a. non-demonstrative inference, abduction, inference to the best explanation, central processing, etc.)
- Non-monotonic reasoning is a black box: 

IDEALIZATION

- By idealizing away from the parts of communication that rely on non-monotonic reasoning, we elide the parts that we don't know how to computationally model.
- This allows us to treat conversation as a rule-governed, deterministic process that behaves according to scrutable principles.
- This, it seems to me, is the best thing that dynamic models have going for them.

IDEALIZATION

- But! Galilean idealizations must eventually be lifted.
- And their legitimacy is relative to our explanatory goals.
- E.g.: Galileo's model of projectile motion would not have been good for studying air resistance.
- Similarly: since this model elides many aspects of human communication, it is not a good model for studying those things.
- This is not to say that it hasn't been useful in lots of ways!

Idealized Models

So we now have a sketch of how scientific work proceeds in the case of early stages of reductionist work on complex systems. In that situation, the currency of scientific work is often models of important processes; models of possible mechanisms, possible dependency structures, that might in time give us an account of the real mechanisms. Once we say it like this, it becomes apparent that this is what a large proportion of work in the cognitive sciences is concerned with today—models of learning, models of numerical cognition, models of the processing of syntax.

—Peter Godfrey-Smith, "Reduction in Real Life"

A Preview

1. Intention Recognition and its Psychological Underpinnings

(Sept 14)

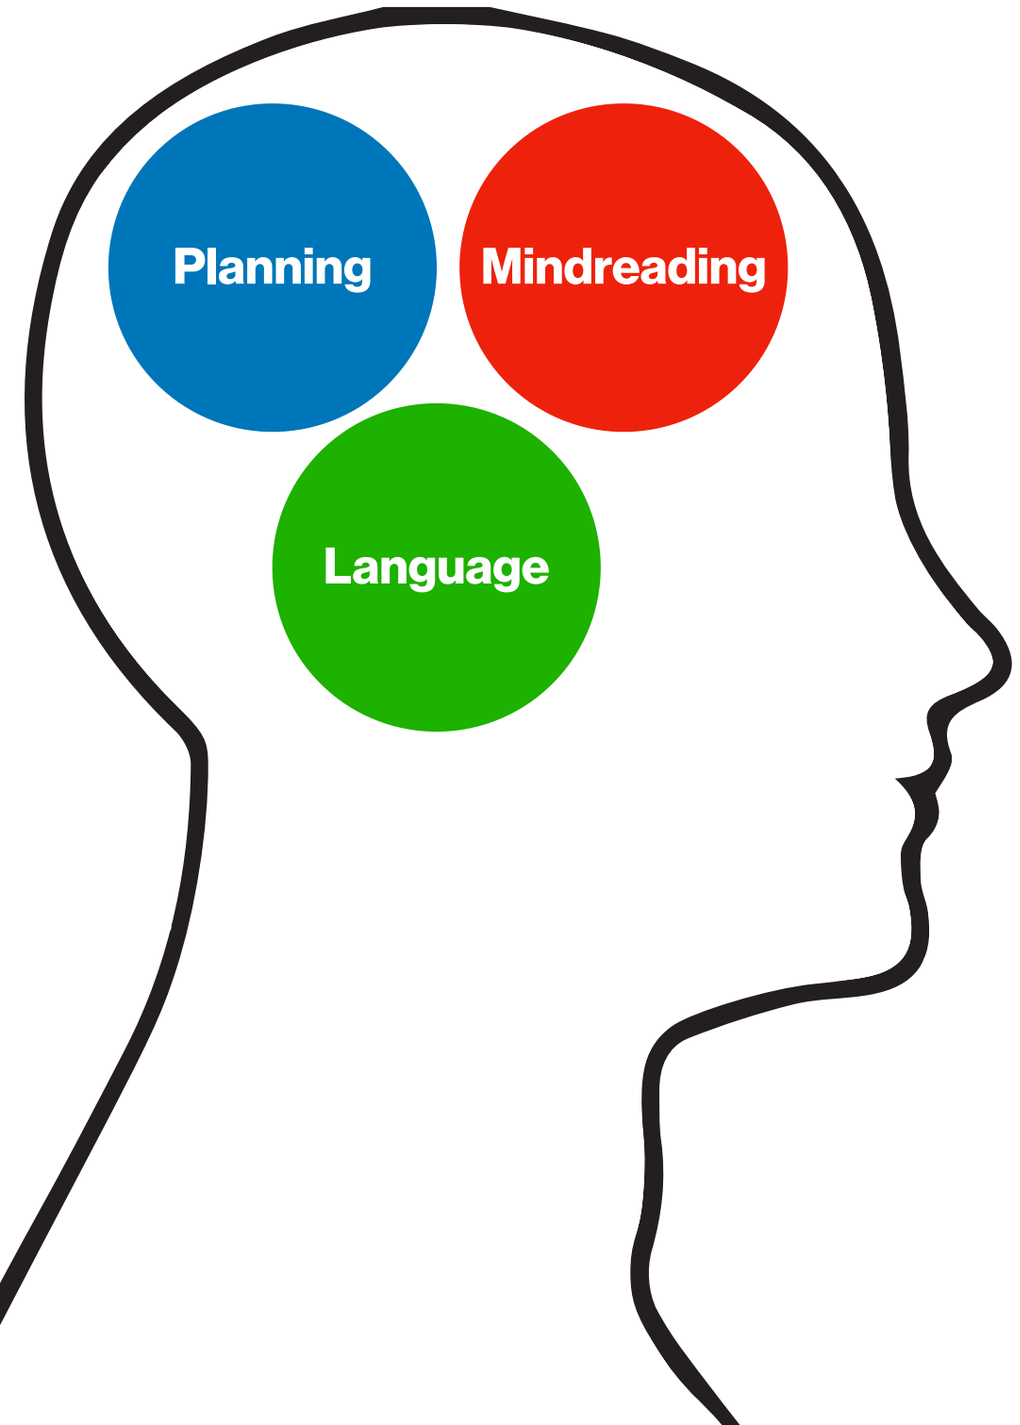

Our capacity for intention recognition emerges from our capacities for mindreading, planning, and language.

Whereas language is a Fodor-modular system, planning and mindreading are not.

2. Designing Communicative Acts

(Sept 21)

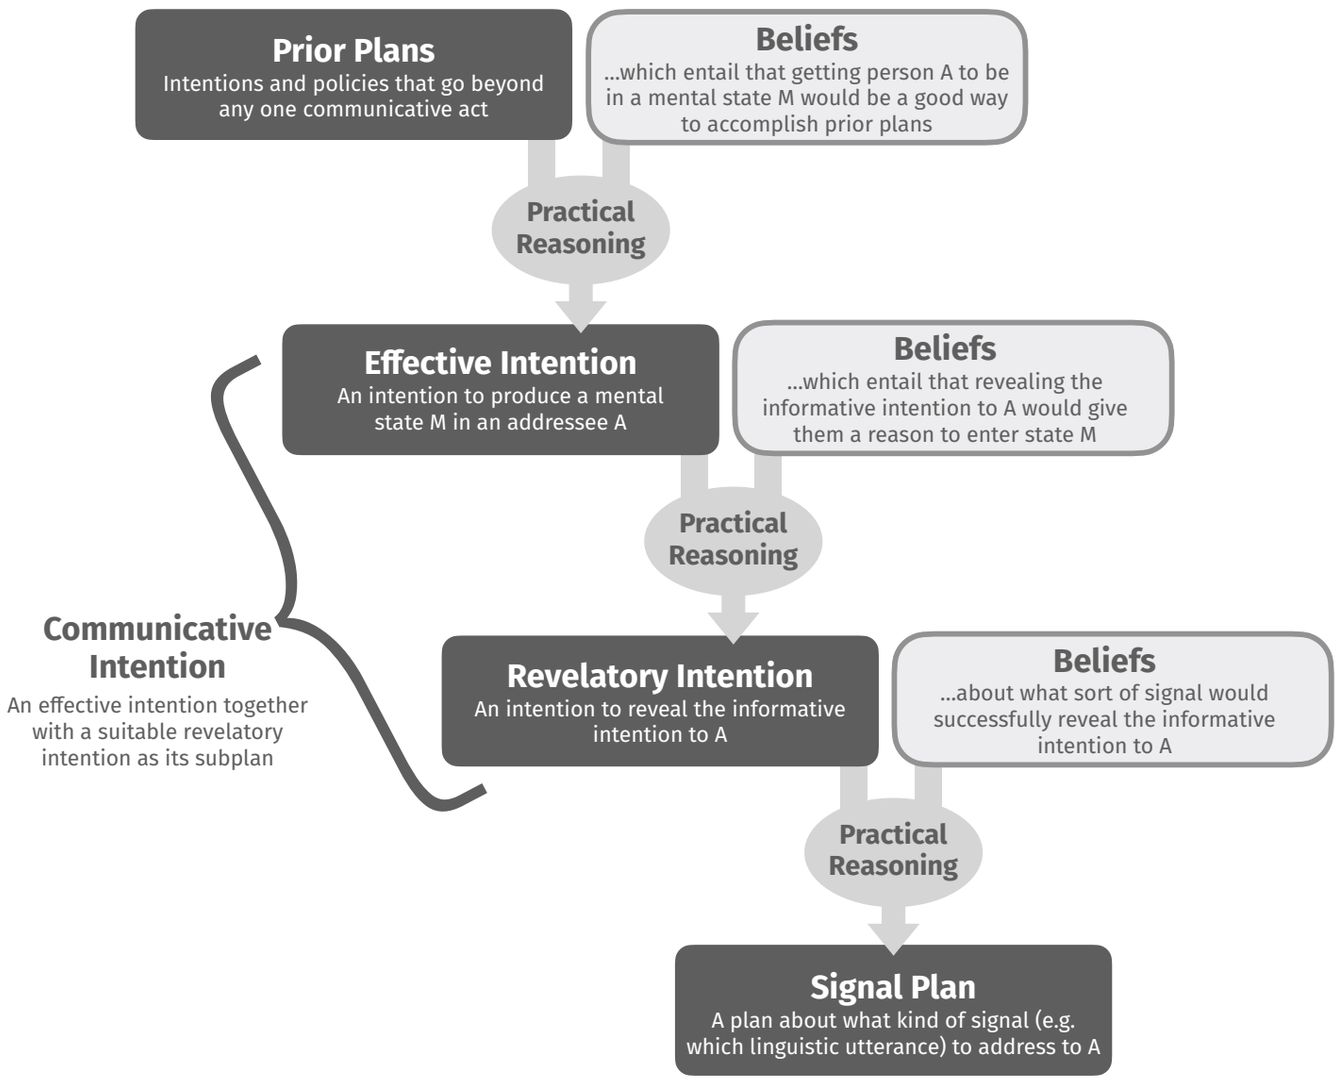

We form communicative intentions as part of the process of designing communicative acts for our addressees.

This design process makes human communication much more powerful and efficient, and this is why we bother with communicative intentions.

3. Planning Conversations Together

(Sept 28)

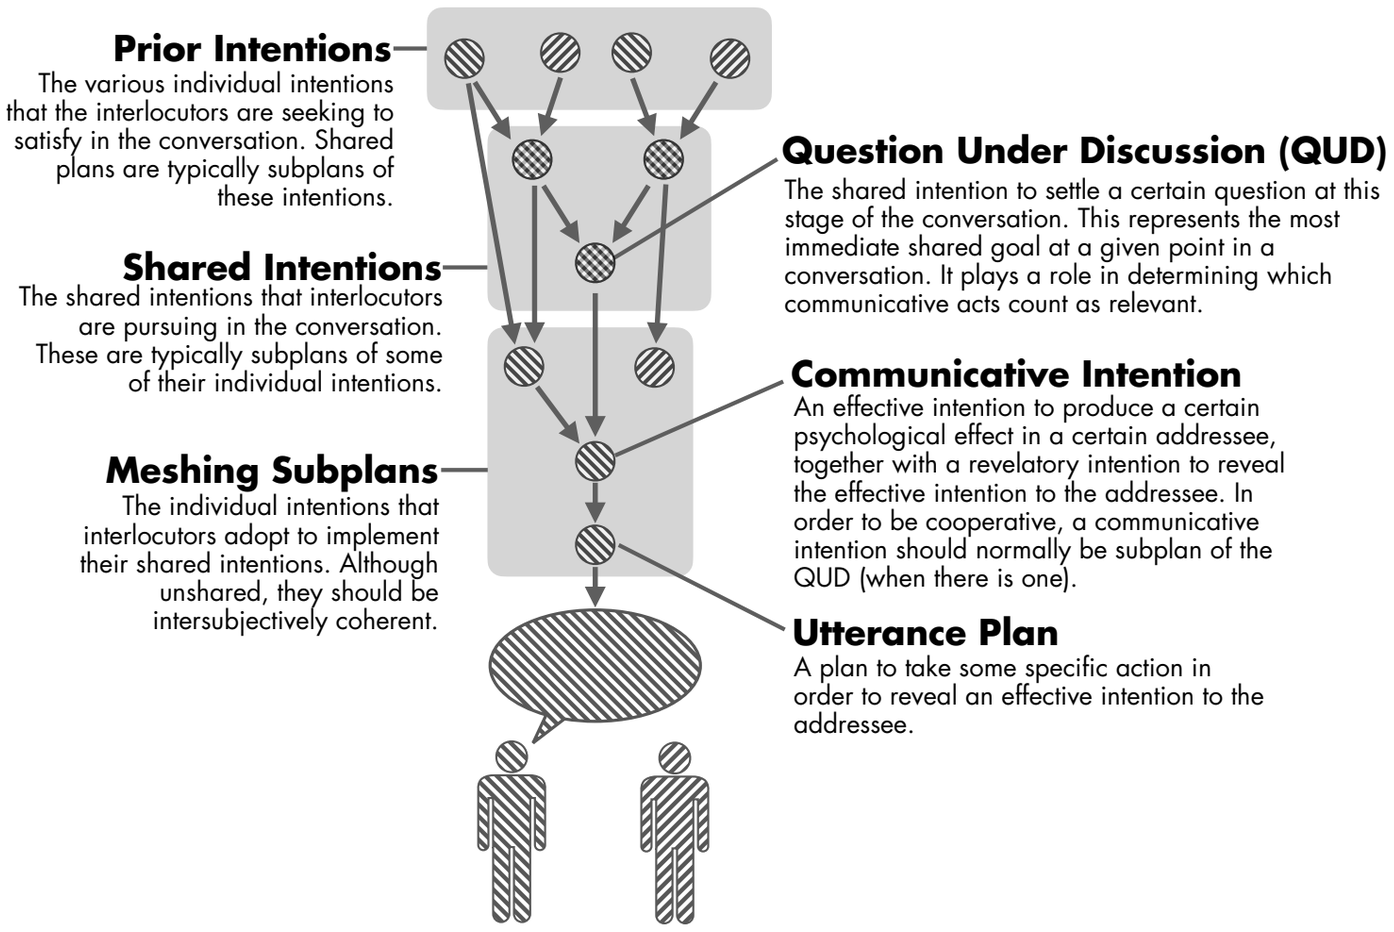

We reveal our intentions to our addressees to allow them to coordinate their communicative efforts with ours.

More generally, we reap many benefits from treating communication as a shared, cooperative activity that is governed by shared plans.

4. Mindreading in Human Communication

(Oct 5)

We do a lot more mindreading when communicating with others than we are conscious of.

This mindreading is cognitively resource intensive, but we do it anyway when we can, because the benefits of intention recognition are worth it.

5. Background Information and Common Ground

(Oct 26)

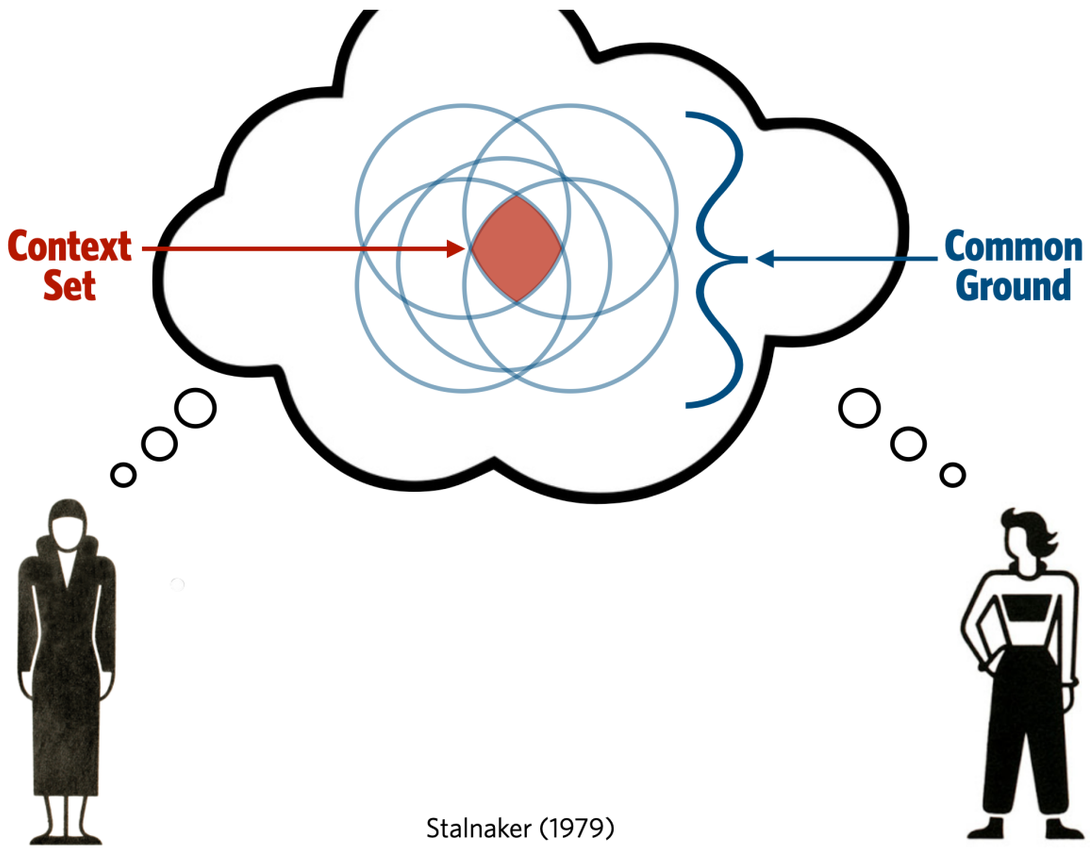

Stalnaker (1979)

The most influential models of conversation tell us that we are constantly drawing on and updating the "common ground." This is an idealization.

Really, we constantly reason about each others' mental states, about their representations of our states (etc), and also when it is worthwhile to do this reasoning.

6. Discourse Referents as Object Files

(Nov 2)

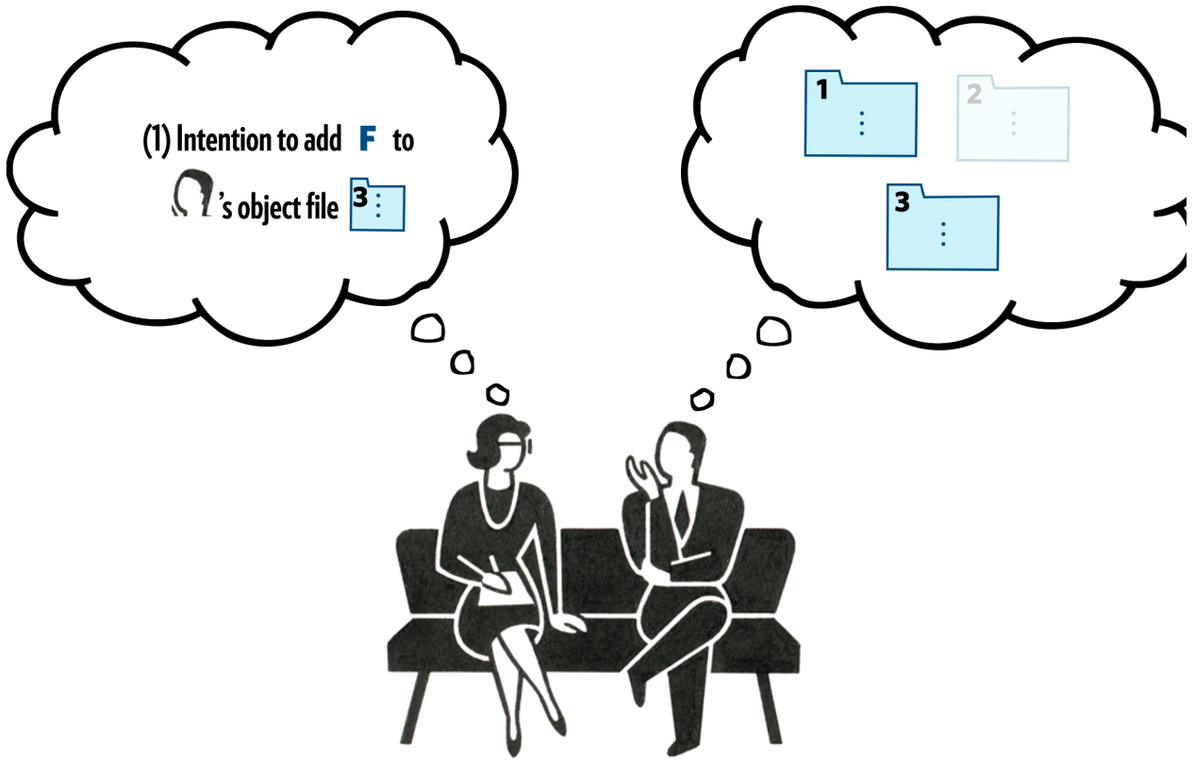

Influential models of conversation say that we track what we're talking about through a conversation using "discourse referents."

A promising hypothesis is that the subject matter of these models are the same mental representations that psychologists call "object files" and philosophers of mind call "mental files."

7. Natural-Language Semantics and Cognitive Architecture

(Nov 9)

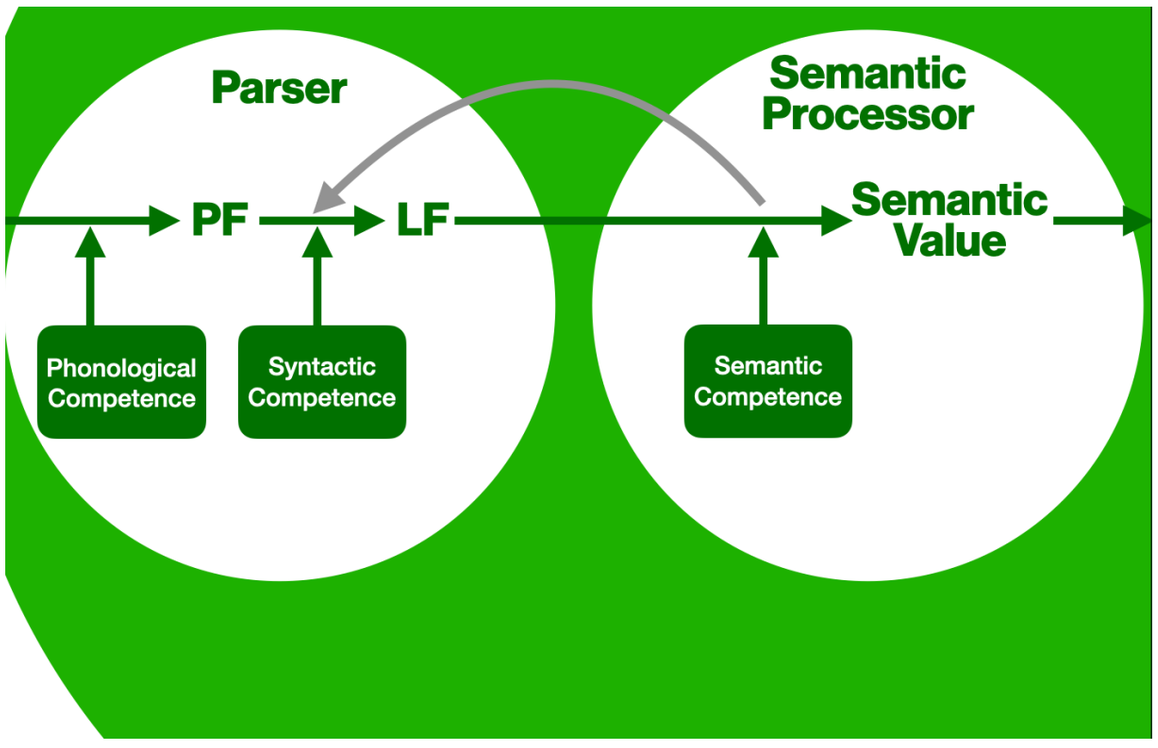

The best explanation of the successes of natural-language semantics is that semanticists are studying a modular input-output system, of the kind first posited by Fodor.

This places some constraints on how we do semantics, but we can non-destructively modify existing theories.

8. Polysemy, Language, and Thought

(Nov 16)

Jeff Bezos bought three **newspapers**.

Nearly every word in natural language is polysemous.

To explain this, we need to find the right division of labor between the grammar and the conceptual system.

We also have some interesting options for how a semantic theory should represent polysemous meanings.

9. How to do Things with Verbal Working Memory

(Dec 7)

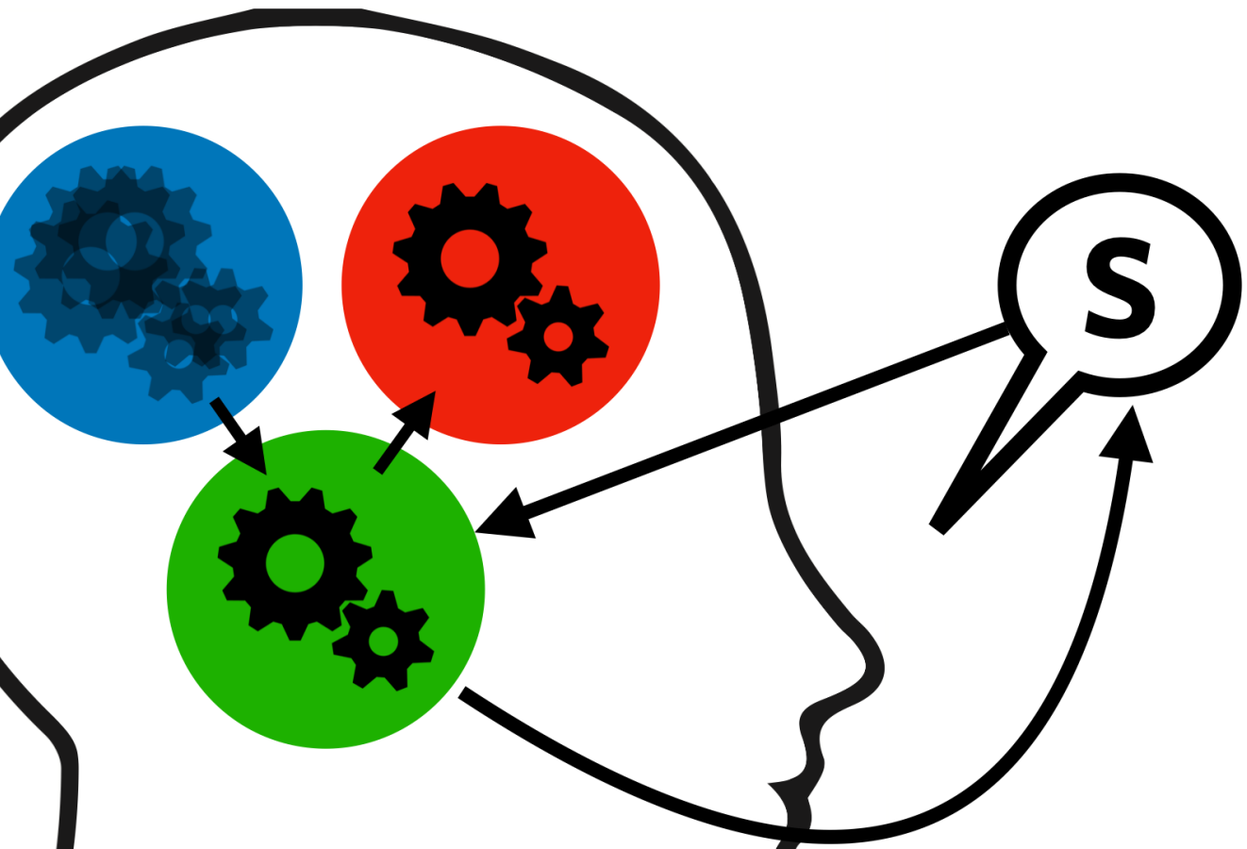

The theory sketched in earlier chapters, plus some straightforward assumptions about verbal working memory, can explain lots of stuff! E.g.:

- Non-communicative language use
- How we can use language to clarify our thoughts
- Neo-whorfian influences of language on thought

10. The Evolution of Language

(Makeup Class: Dec 14?)

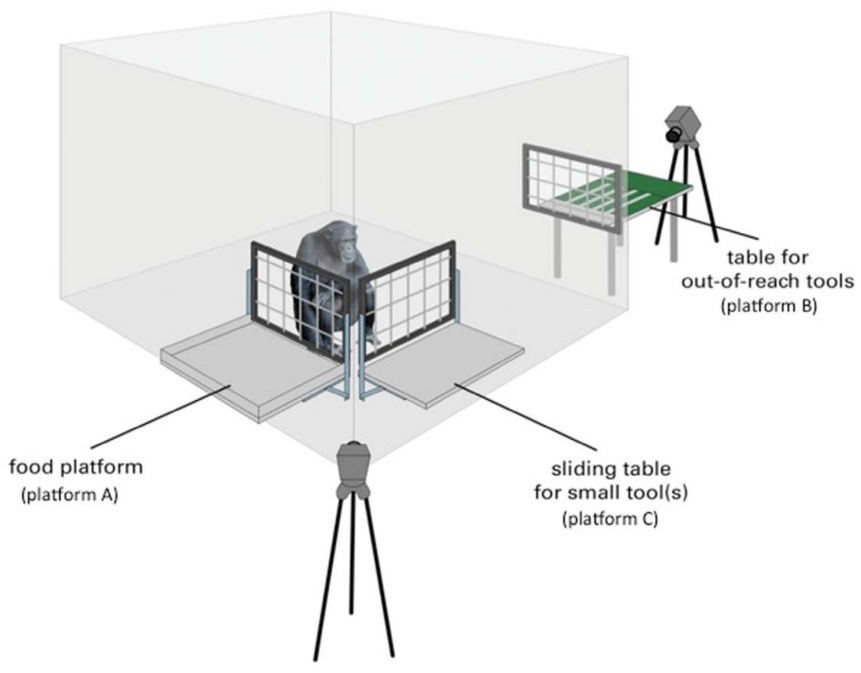

Ordas, Schumacher, and Call (2012): “Sequential Tool Use in Great Apes”

Language-use depends on our capacity for intention recognition.

But it doesn’t follow that our ancestors had to become intention recognizers before they could evolve language.

It is plausible that other great apes have homologues of most of the cognitive ingredients in our capacity for intention recognition, but they aren’t as developed or integrated as ours.

LFG!