NASSLLI 2022 · USC · 20-24 JUNE

SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

Daniel W. Harris
CUNY Graduate Center and Hunter College

What is (Compositional) Semantics?

A research program that we can find playing out in places like these:

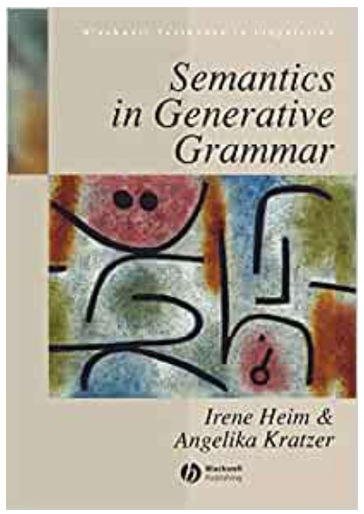

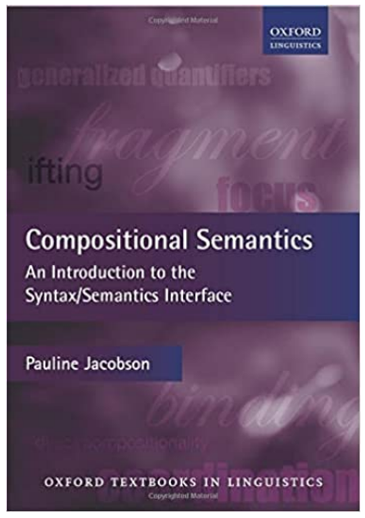

A presupposition of this course

Compositional semantics is a *progressive* research program that is in the process of converging on genuine knowledge of something.

Our guiding question

What is that something? What are we learning about when we're doing semantics?

Our guiding question

What is that something? What are we learning about when we're doing semantics?

The answer that we'll be exploring

When we do semantics, we're building a model of the proprietary database of a modular input-output system.

The Cognitive-Architectural Picture

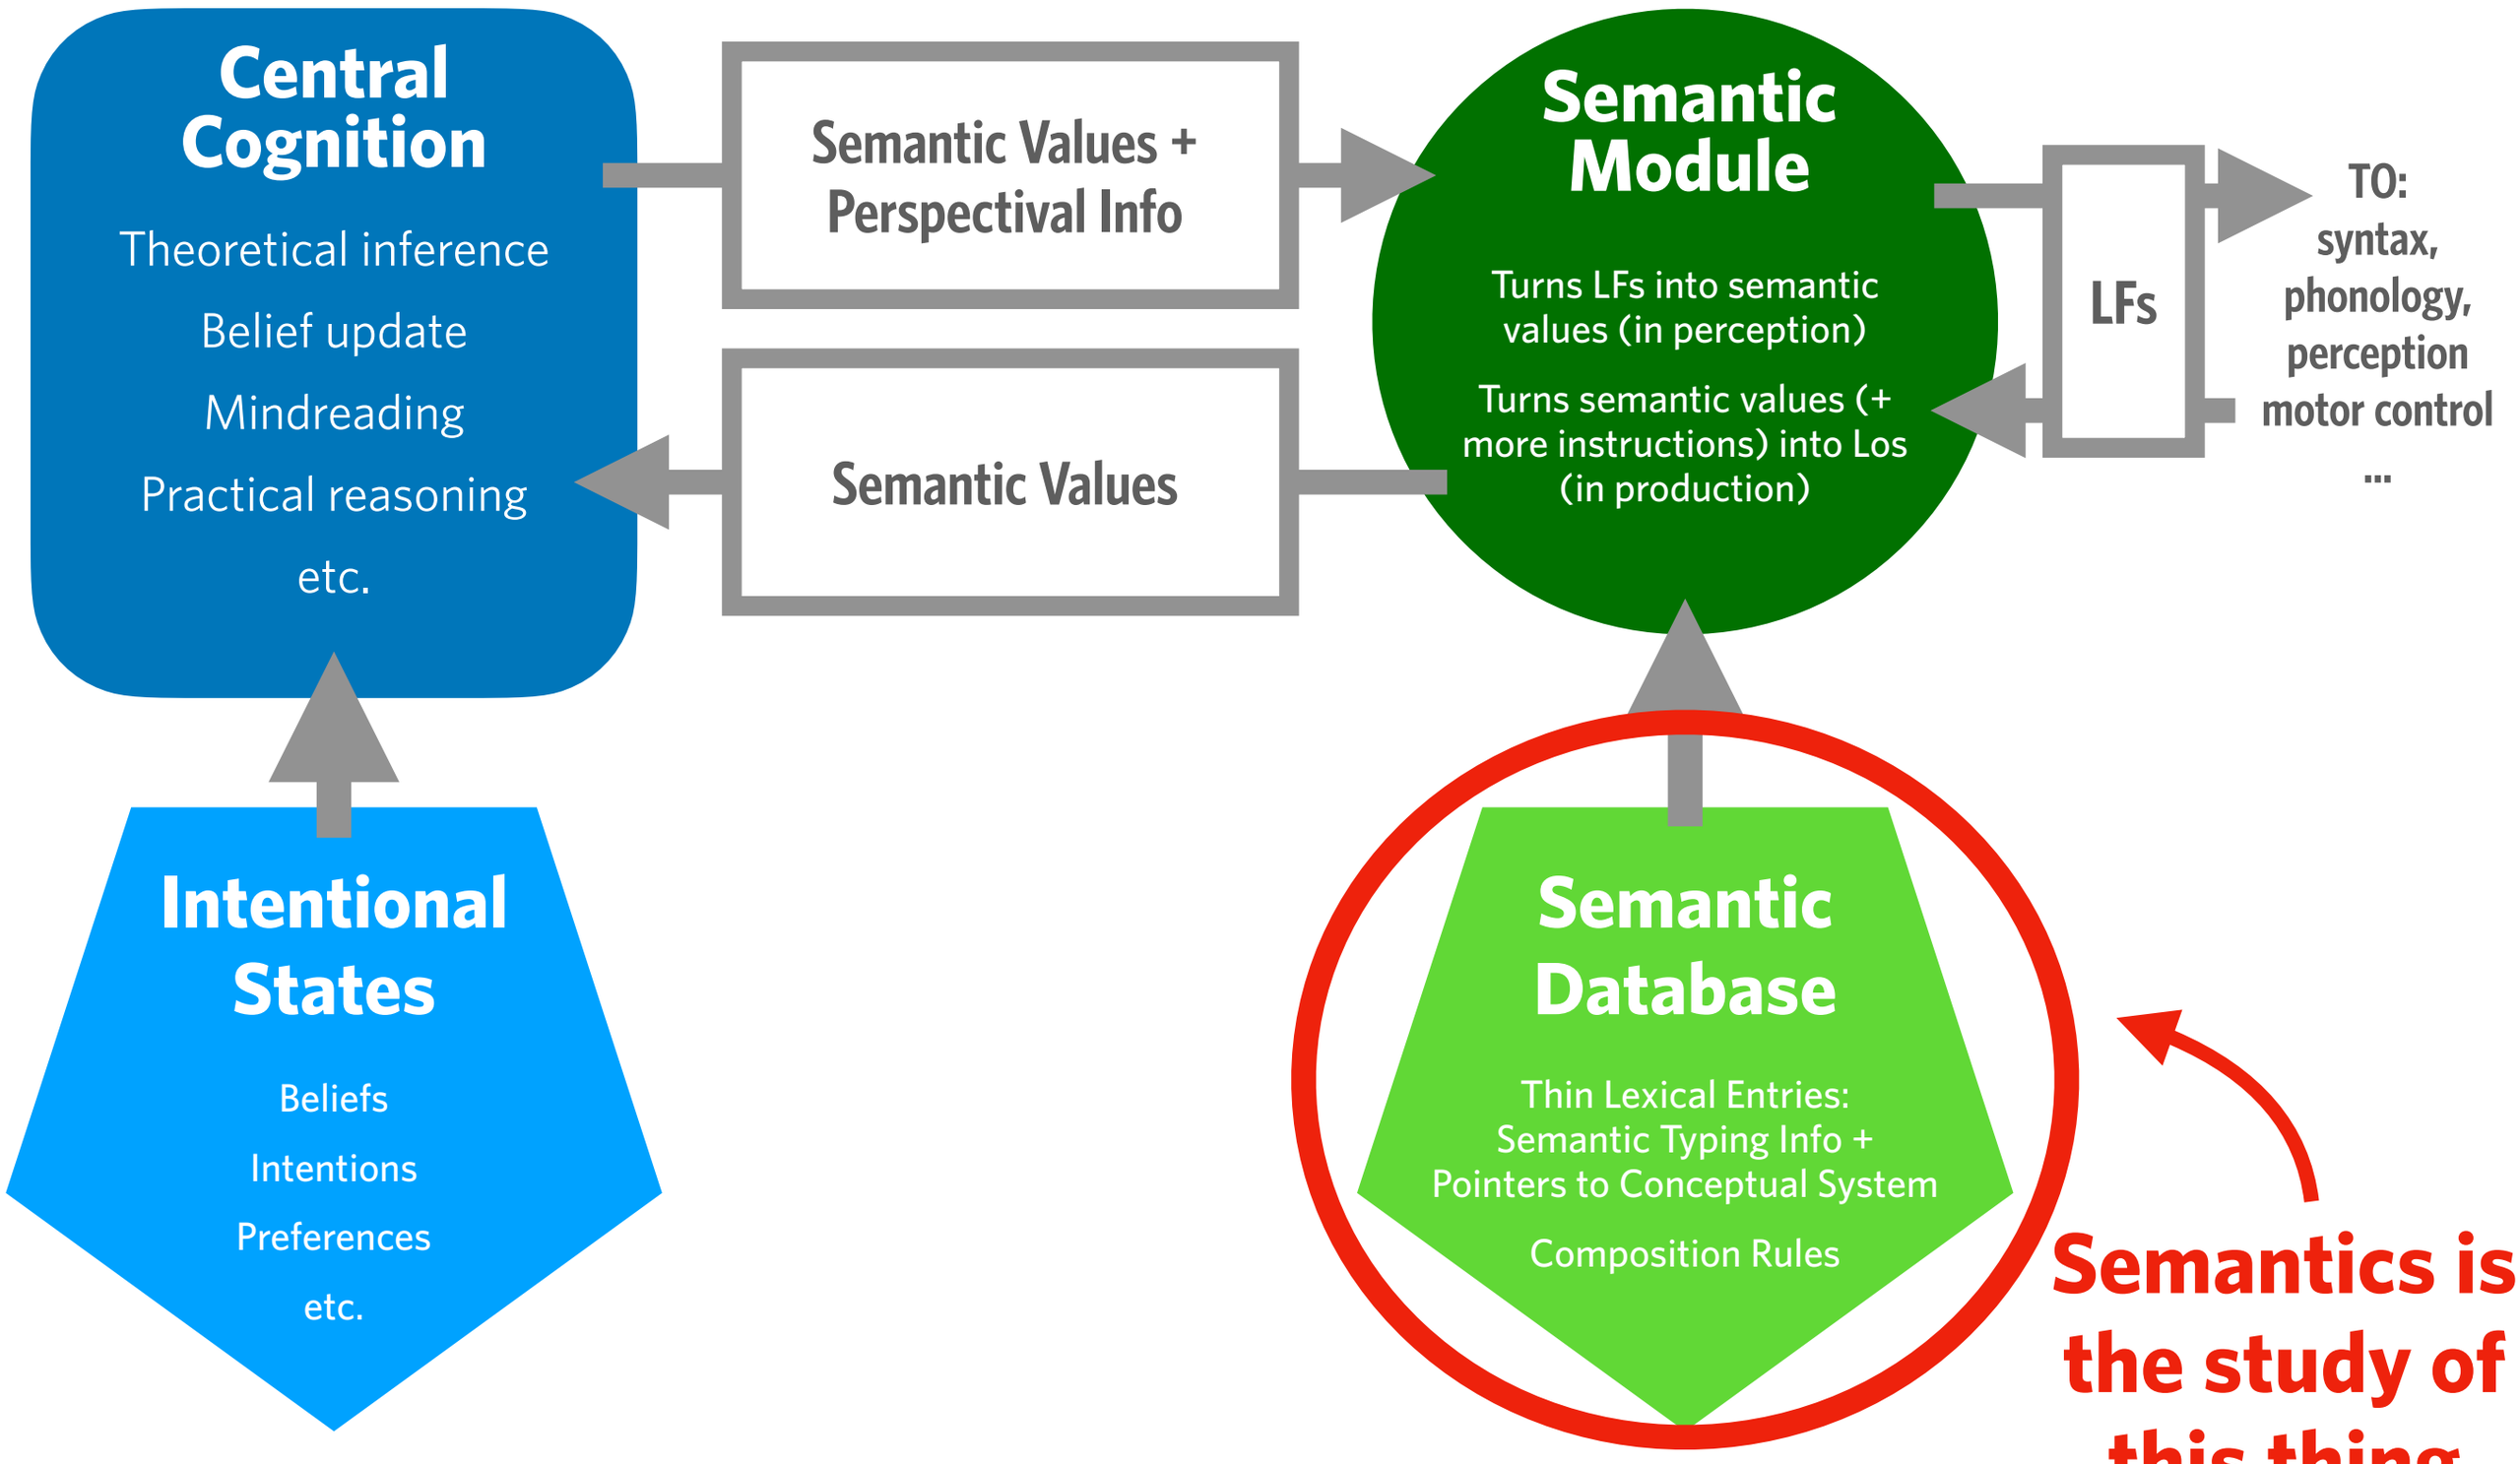

The axe I'll be grinding

This course will include material from some things that I have written (and some that I am still writing):

"Semantics without Semantic Content," *Mind & Language* 37, 2022. pp. 304–328.

"Speaker Reference and Cognitive Architecture," *Croatian Journal of Philosophy* 17 (3), 2017. pp. 319–349.

"Intention Recognition as the Mechanism of Human Communication," in *Sensations, Thoughts, Language: Essays in Honor of Brian Loar*, edited by Arthur Sullivan. Routledge, 2019. pp. 11–37.

Several (partially written) chapters of the manuscript of the book I have been working on, tentatively titled *Human Communication*.

But!

You know lots of things I don't. Please teach me!

NASSLLI 2022 · USC · 20–24 JUNE
SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

DAY 1:
SEMANTICS AND MODULARITY

Daniel W. Harris
CUNY Graduate Center and Hunter College

What is the subject matter of semantics?

What is the subject matter of semantics?

- Abstract objects (Montague, Katz)
- Patterns in token linguistic objects (Devitt)
- What an interpreter would have to know in order to be in a position to interpret other speakers of a language (Davidson)
- Conventions of public languages (Lewis)
- Facts about human psychology:
  - (Common) beliefs or knowledge (Lewis)
  - A specialized mental organ (Chomsky, Fodor, etc.)

If semantics is a progressive research program, shouldn't we know what it's about already?

Derek Ball (2014): "Semantics as Measurement"

It's actually somewhat common for a science to mature in its predictive and explanatory power even while scientists disagree about what they're studying.

Ball's example: in the 18th-Century, our understanding of temperature measurement slowly improved while a debate played about what was being measured: average kinetic energy or the quantity of a fluid called caloric.

The first-order and foundational debates played out in parallel, informing each other at various stages.

What is the subject matter of semantics? (Redux)

- Abstract objects (Montague, Katz)
- Patterns in token linguistic objects (Devitt)
- What an interpreter would have to know in order to be in a position to interpret other speakers of a language (Davidson)
- Conventions of public languages (Lewis)
- Facts about human psychology:
  - (Common) beliefs or knowledge (Lewis)
  - A specialized mental organ (Chomsky, Fodor, etc.)

Semantics as the study of (shared?) beliefs

David Lewis (1980): "Index, Context, and Content"

"A good grammar is one suited to play a certain role in a systematic restatement of our common knowledge about language. It is the detailed and parochial part — the part that would be different if we were Liars, or if we were Japanese."

Semantics as the study of (shared?) beliefs

David Lewis (1980): "Index, Context, and Content"

"A good grammar is one suited to play a certain role in a systematic restatement of our common knowledge about language. It is the detailed and parochial part — the part that would be different if we were Liars, or if we were Japanese."

...

"You might insist that a good grammar should be suited to fit into a psycholinguistic theory that goes beyond our common knowledge and explains the inner mechanisms that make our practice possible. There is nothing wrong in principle with this ambitious goal, but I doubt that it is worthwhile to pursue it in our present state of knowledge." (p.81)

Why isn't semantics the study of beliefs?

Daniel W. Harris (2022): "Semantics without Semantic Content"

In brief:

- There are good reasons to think that body of information that we're reverse engineering when we do semantics is not a body of beliefs.
- It isn't acquired in belief-like ways.
- It doesn't change in belief-like ways.
- It involves representations of things that ordinary speakers don't (and can't) have beliefs about.
- In cases where we know that there are contradictions between speakers' beliefs and the information in the semantic database, the beliefs don't change how speakers use language.

Modularity

THE MODULARITY OF MIND

Jerry A. Fodor

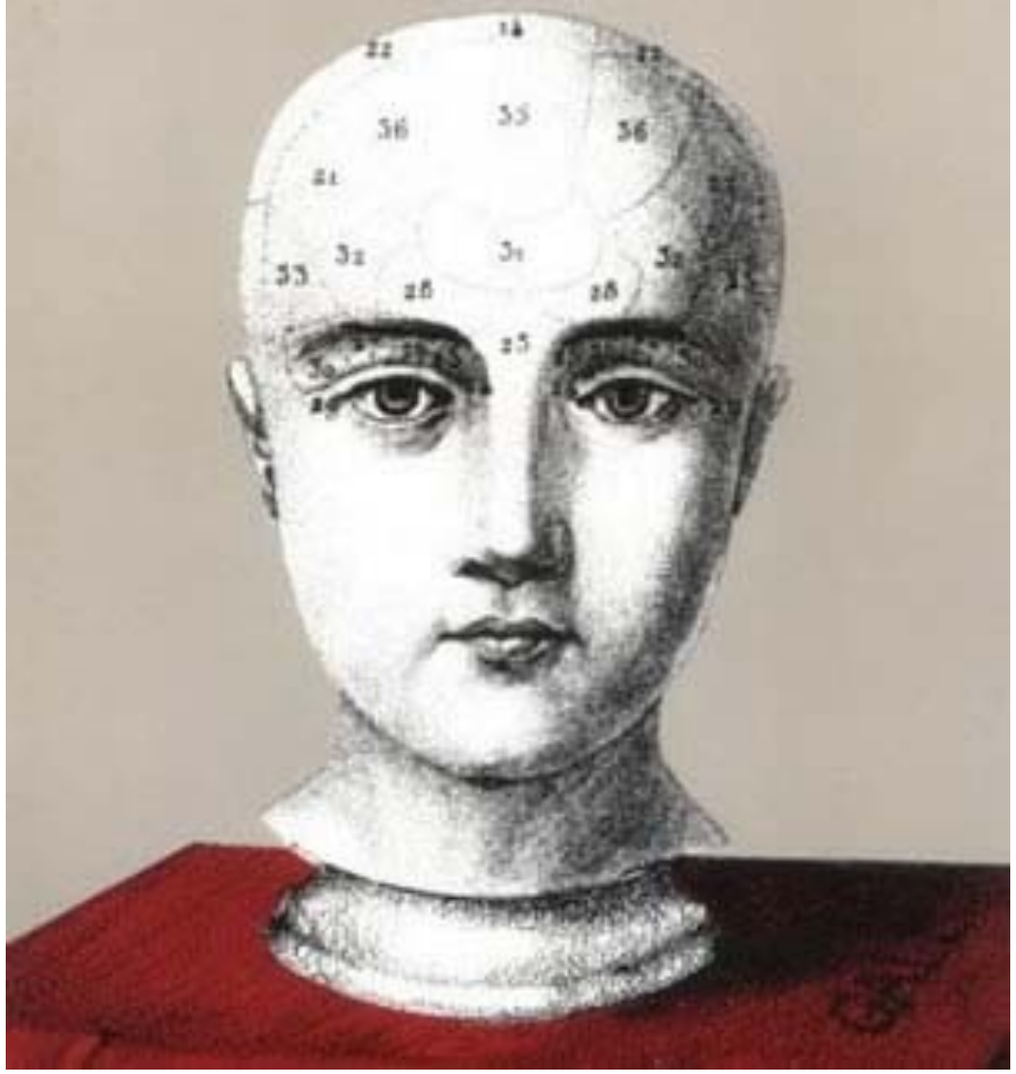

Sensory Transducers

Perception:
"Input Systems"

Auditory Perception

Visual Perception

Language:
"Input/Output System"

Central Cognition

- Theoretical inference
- Belief update
- Mindreading
- Practical reasoning
- (etc.)

Motor Control:
"Output System(s)"

Fodor's Picture of Cognitive Architecture

Characteristics of Modular Systems

- Proprietary database
- Proprietary representational format or vocabulary
- Domain specificity
- Central inaccessibility
- Informational encapsulation
- Mandatory operation
- Fast processing
- ‘Shallow’ outputs

- Amenability to computational modeling
- Characteristic and specific breakdown patterns
- Characteristic ontogenetic pace and sequencing (innate)
- Fixed neural architecture
- Distinctive evolutionary function

Characteristics of Modular Systems

- Proprietary database
- Proprietary representational format or vocabulary
- Domain specificity
- Central inaccessibility
- Informational encapsulation
- Mandatory operation
- Fast processing
- ‘Shallow’ outputs

- Amenability to computational modeling
- Characteristic and specific breakdown patterns
- Characteristic ontogenetic pace and sequencing (innate)
- Fixed neural architecture
- Distinctive evolutionary function

Proprietary Database

- A modular system has its own body of information/ rules to draw on.
- This “proprietary database” is distinct from the agent’s beliefs, which are available to central-cognitive processes.

Proprietary Database

Christopher Peacocke (1986):
"Explanation in Computational Psychology: Language, Perception and Level 1.5"

Marr distinguished three levels of description of a computational psychological process:

1. Computational: Describes which function is being computed and why.
2. Algorithmic: Describes the algorithm that computes the function.
3. Implementational: Describes how the algorithm is implemented by hardware.

Proprietary Database

Christopher Peacocke (1986):
"Explanation in Computational Psychology: Language, Perception and Level 1.5"

Marr distinguished three levels of description of a computational psychological process:

1. Computational: Describes which function is being computed and why.

Peacocke adds level 1.5, which describes "the information on which the algorithm draws"

2. Algorithmic: Describes the algorithm that computes the function.
3. Implementational: Describes how the algorithm is implemented by hardware.

The Cognitive-Architectural Picture

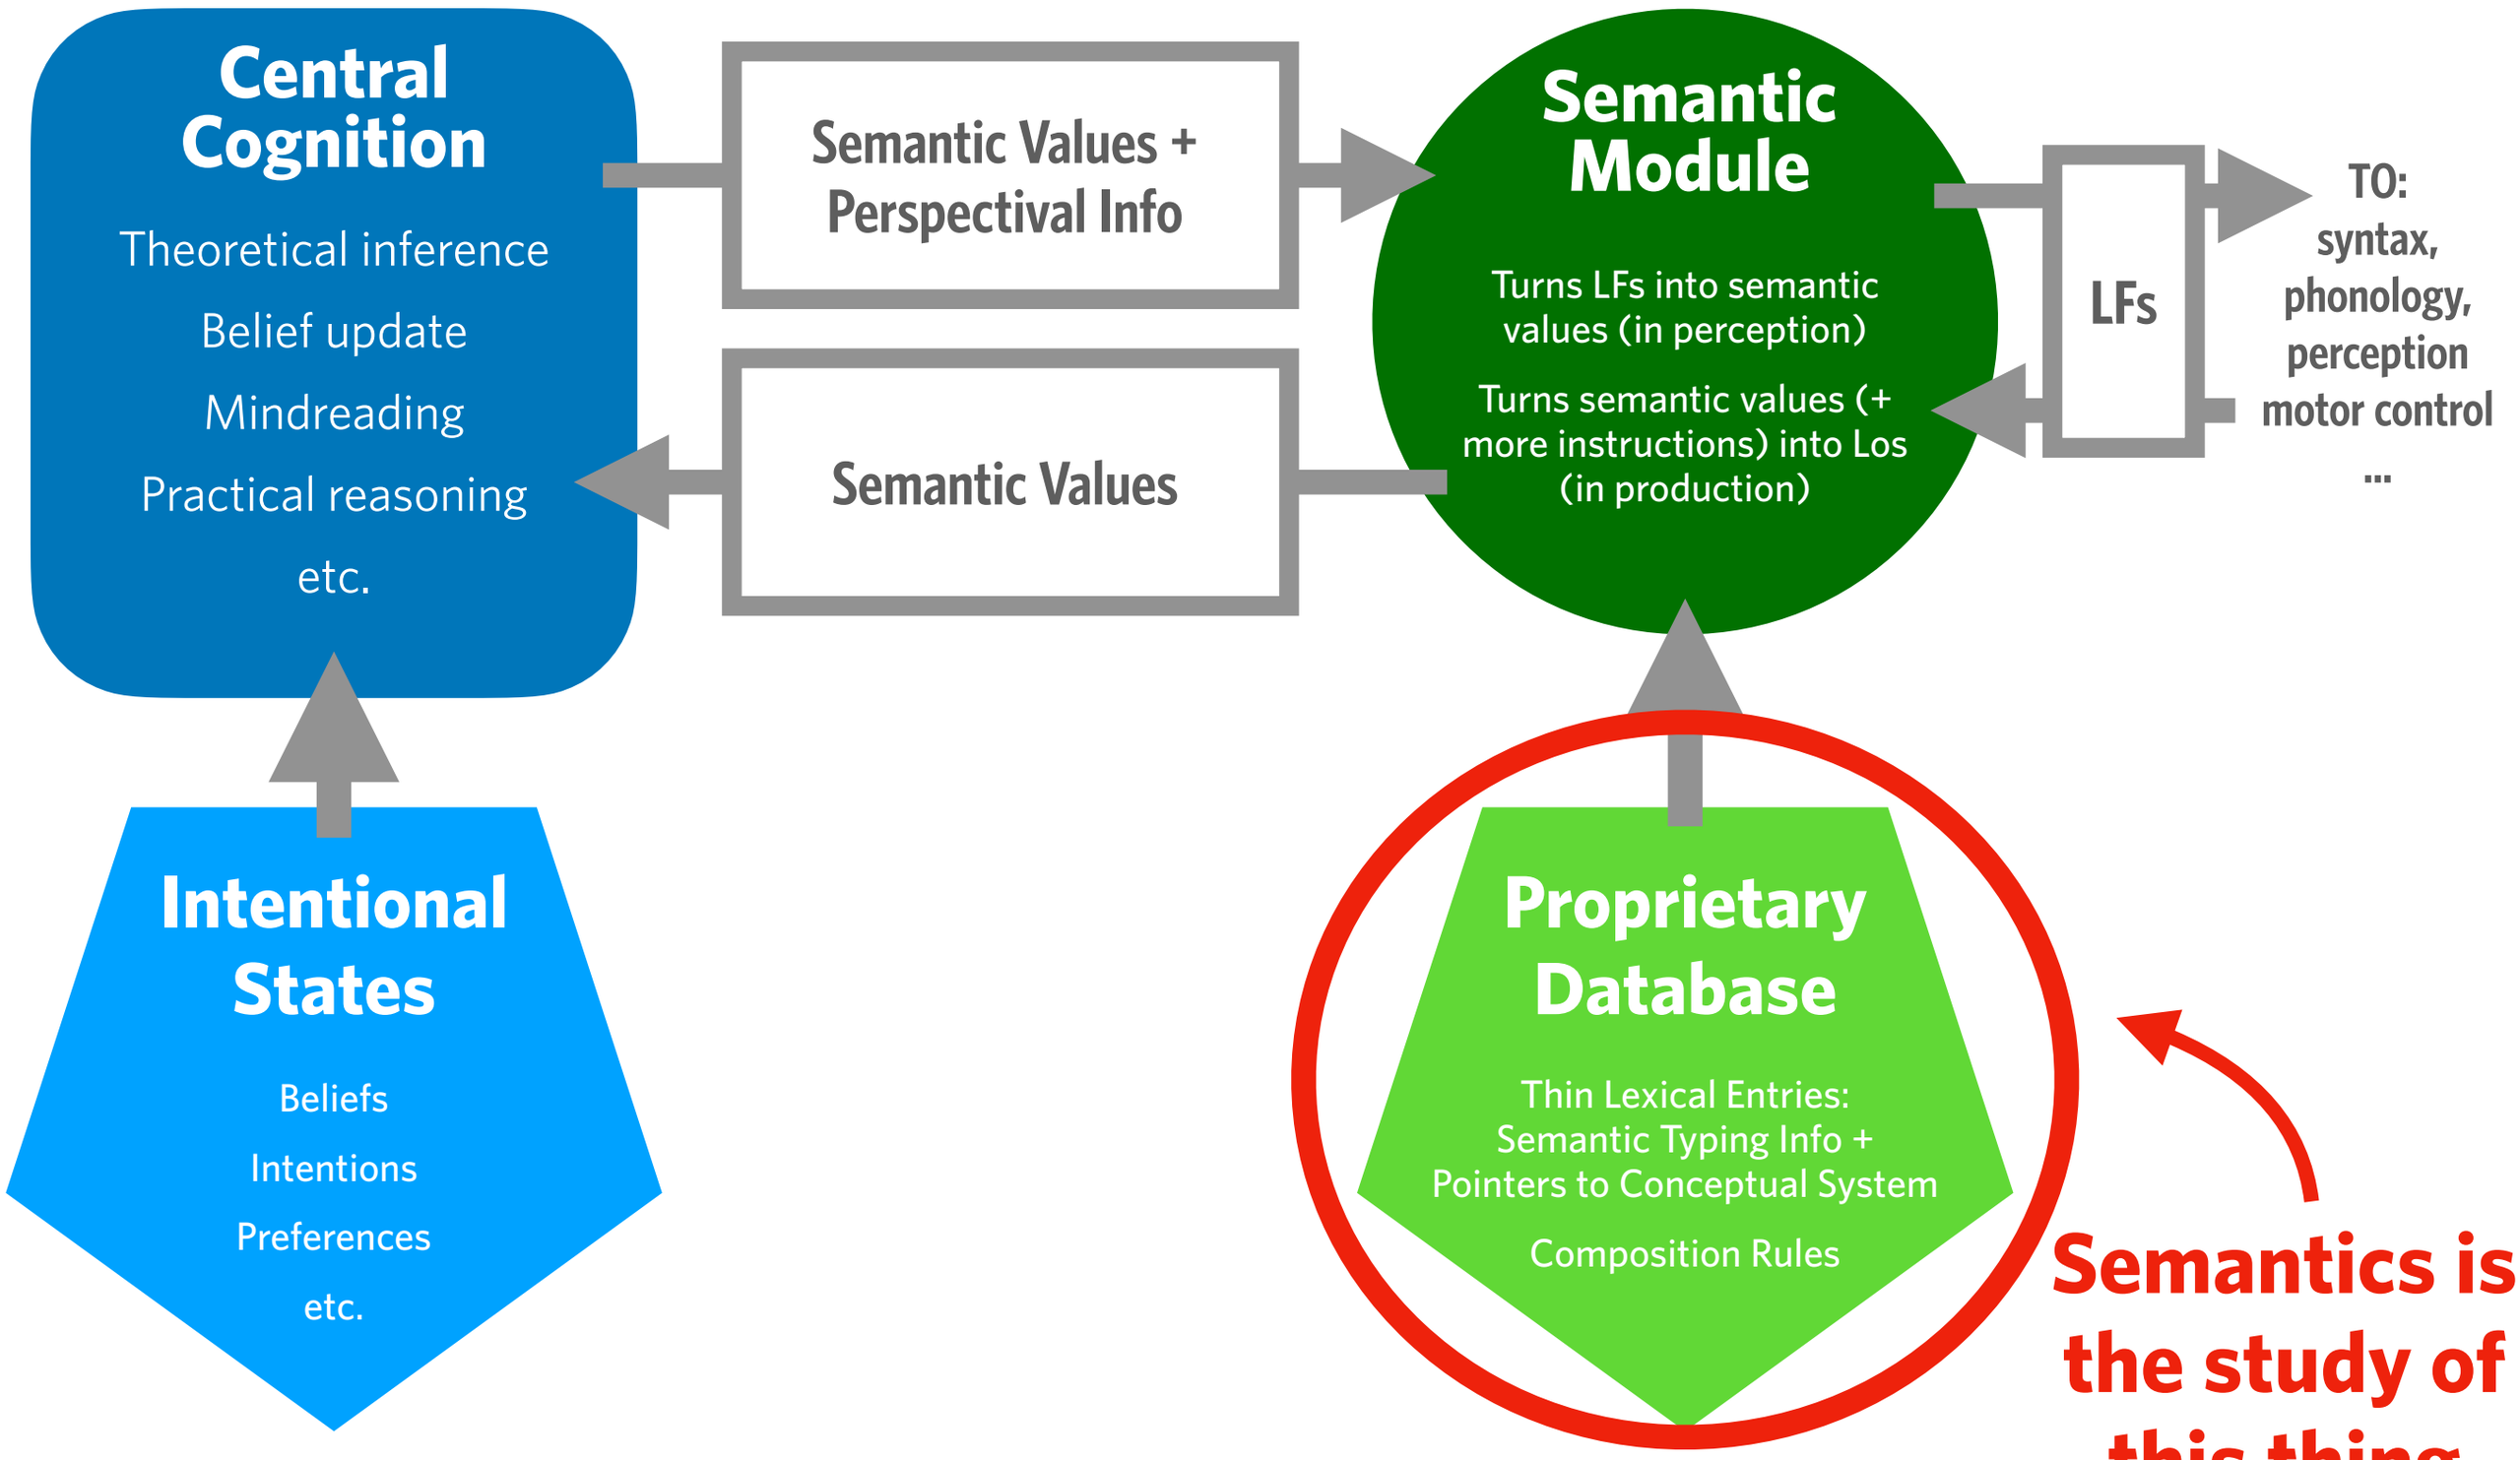

What's in the semantic database?

Heim & Kratzer (1998): Semantics in Generative Grammar

$$[\![he_i]\!]^c = g^c(i)$$

The content of "he," relative to an index i, is the value of the contextually given assignment for the index i.

$$[\![talk]\!]^c = \lambda x_e \,.\, x \text{ TALKS in } w$$

The content of "talk" is a function that maps a possible world and an individual to truth iff the individual talks at the world.

Functional Application

$$[\![\alpha\beta]\!]^c = \lambda w \,.\, [\![\alpha]\!]^c([\![\beta]\!]^c) \text{ or } [\![\beta]\!]^c([\![\alpha]\!]^c)$$

To derive the content of a compound expression, apply the function denoted by one to the argument denoted by the other.

Proprietary Representational Format/Vocabulary

Modular systems may have proprietary ways of *formatting* information.

Example representational formats:

- Language-like (discursive)
- Image-like (iconic)
- Map-like
- Diagram-like
- Neural net-like

Many have argued that perceptual and motor systems differ from central cognition in how their representations are formatted.

Differently formatted systems are limited in their interoperability.

Proprietary Representational Format/Vocabulary

- What is the format of whatever system does semantic processing?
    - If we take contemporary semantics seriously, it is a discursive format with the expressive power of (a fragment of) the typed lambda calculus.
- What about the format(s) of central cognition?
    - Extremely controversial, and there may be no single answer. But plausibly it is at least partially linguistic. (The Language of Thought Hypothesis)
- Does this mean that they are totally interoperable?
    - Not necessarily! It depends on their vocabulary.

The Semantic Module's "Vocabulary"

g^c
assignment function

i
numerical index

- These concepts play critical roles in semantic composition (on many views).
- But ordinary speakers do not possess them, in the usual sense.
- Non-semanticists can't have beliefs about them.
- Heterodox theorists who don't believe in them don't speak differently.

Domain Specificity

- Modular systems perform computations over limited subject matters.
- (This property tends to go along with a proprietary database and format/vocabulary.)
- E.g.: Perceptual input systems performs computations on information from specific sensory transducers.
- That's not true of language/semantics, which is multimodal and input/output.
- But semantics is always in the business of mapping LFs to semantic values.

Central Inaccessibility

- The proprietary database, operations, and intermediate representations of a modular system are off limits to central cognition.
  - E.g. we can't introspect the information on which visual perception draws, its intermediate outputs, or the computations it performs.
    - Same goes for semantic processing!
- This is why semantics is so hard! It takes painstaking reverse engineering.
- (The relationship to proprietary vocabulary should be clear.)

Central Inaccessibility of Semantic Information

$[\![\text{Dan}]\!]^c = \text{DAN}$

Names are devices of direct reference (Kaplan 1989)

$[\![\text{Dan}_i]\!]^c = g^c(i)$

Names are pronouns (Cumming 2008; Schoubye 2016)

$[\![\text{Dan}]\!]^c = \lambda x\ .\ x \text{ IS A DAN}$

Names are predicates (Fara 2015)

$[\![\text{Dan}]\!]^c = \lambda \Phi\ .\ \text{DAN IS } \Phi$

Names are quantifiers (Montague 1973)

- All live options in the semantics literature
- Deciding between them will take painstaking reverse engineering.
- Ordinary speakers' beliefs tell us *nothing* about which is correct.

Informational Encapsulation

- A module does its job in a way that is insensitive to the information and inner workings in central cognition.
  - The module can't reach into GC and grab beliefs to inform its operations.
  - GC can't reach into the module and interfere with its operations.

Properties of Modular Systems

Informational Encapsulation

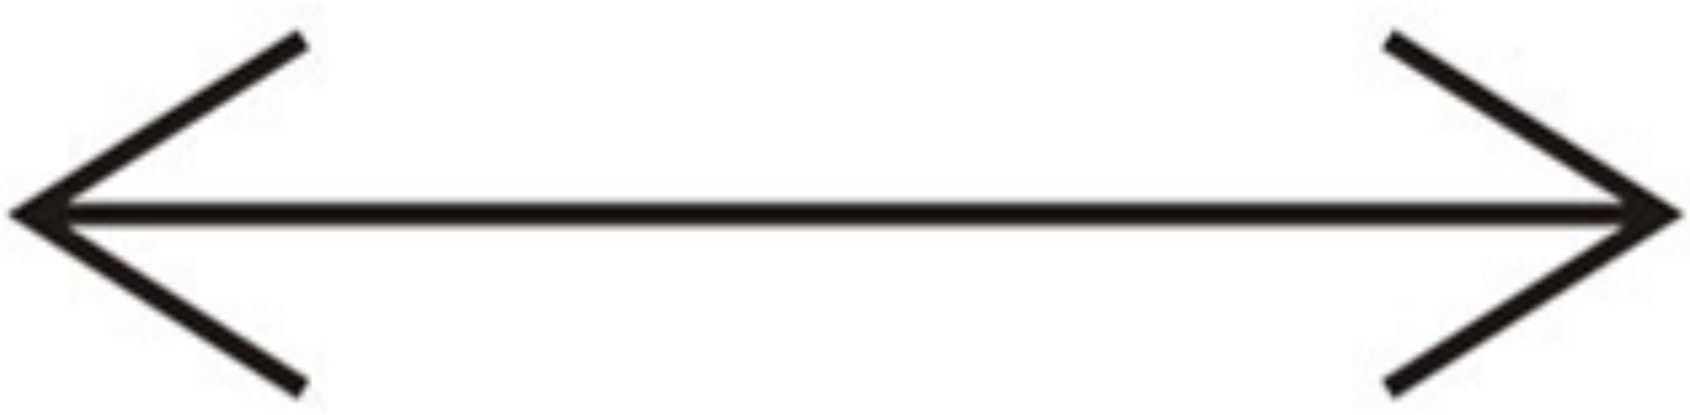

Properties of Modular Systems

Informational Encapsulation

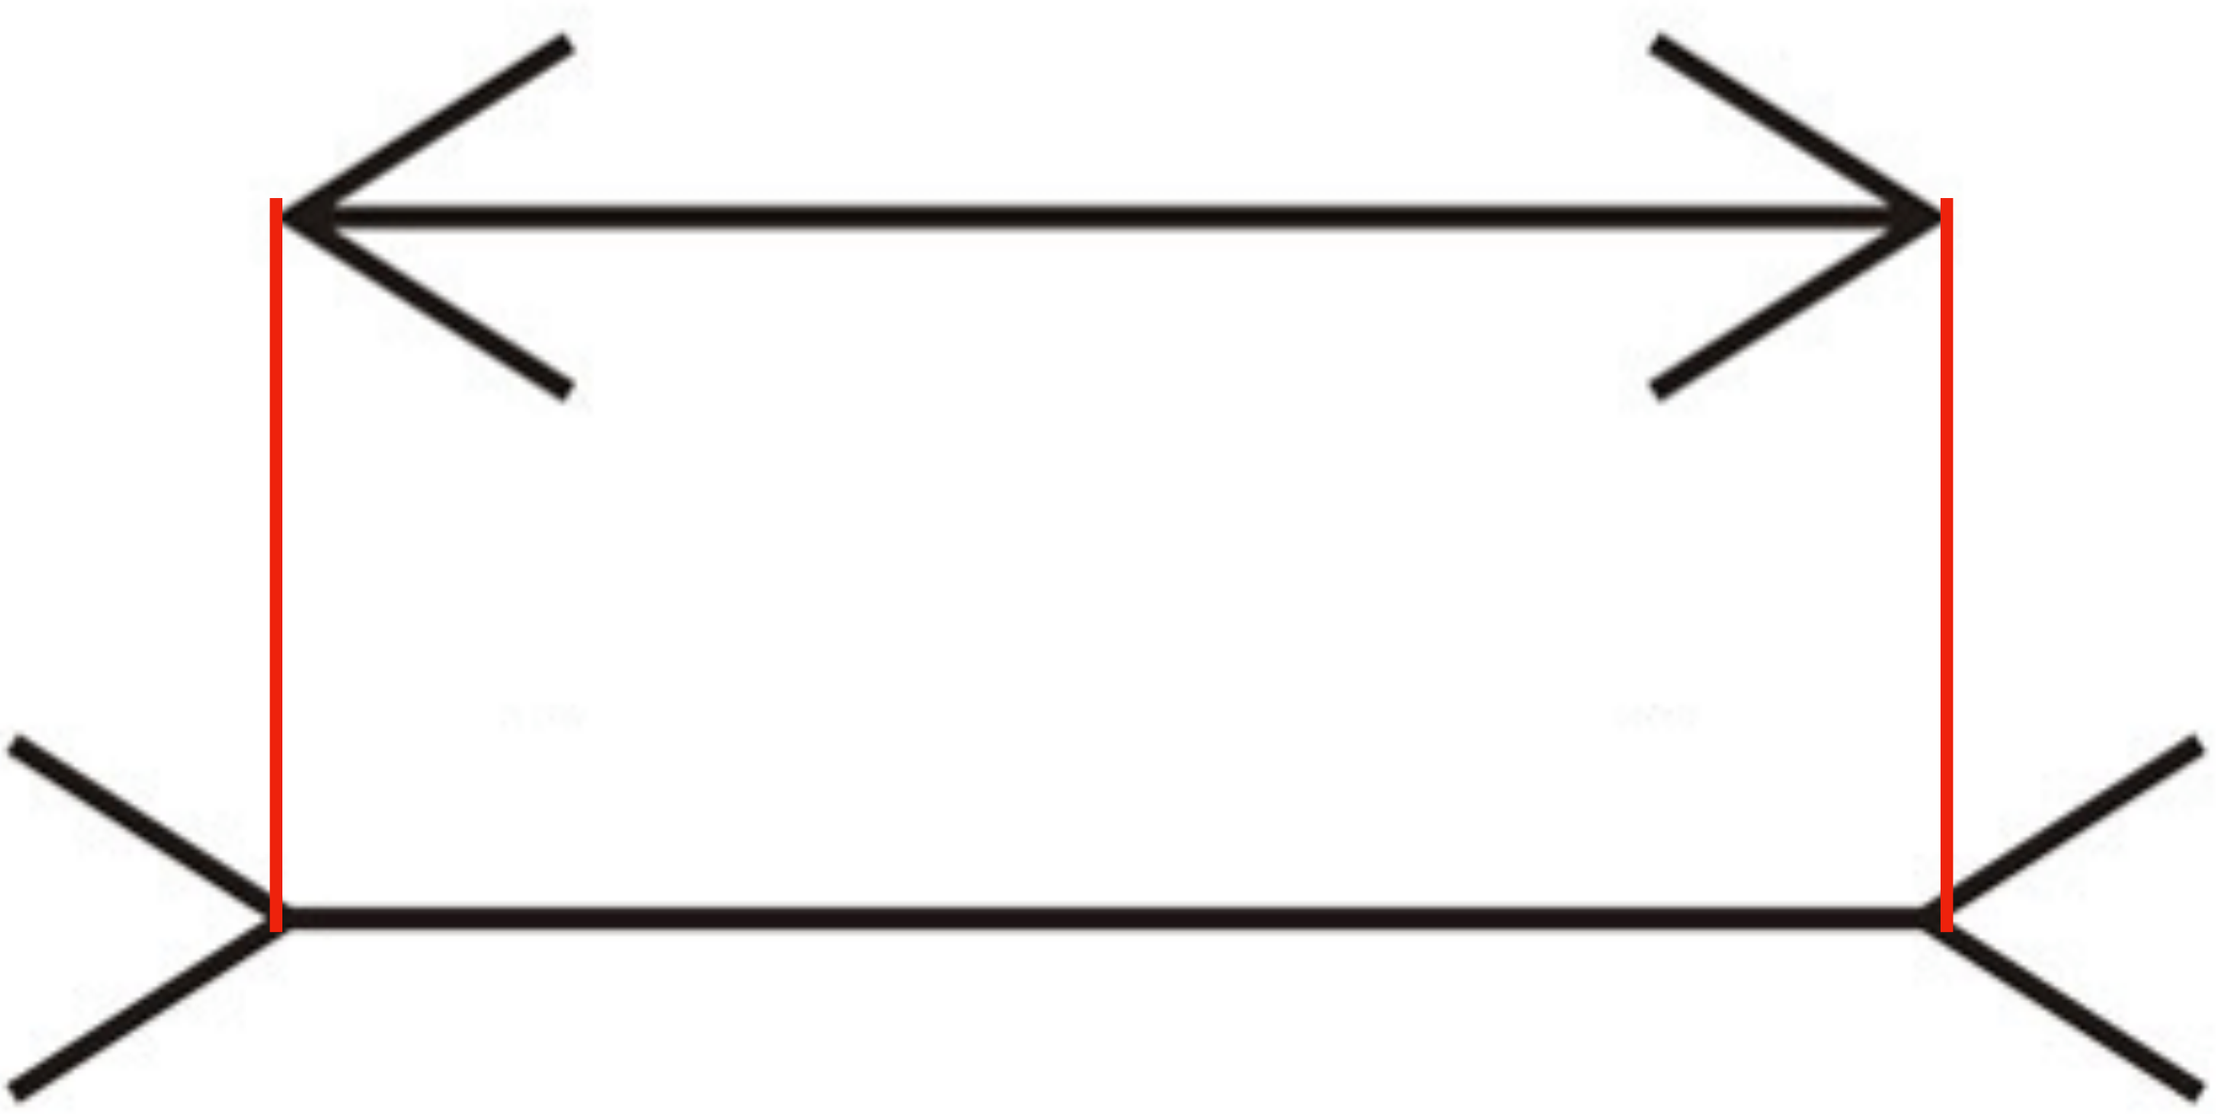

Properties of Modular Systems

Informational Encapsulation

Encapsulation of Semantics

$[\![\text{Dan}]\!]^c = \text{DAN}$

Names are devices of direct reference (Kaplan 1989)

$[\![\text{Dan}_i]\!]^c = g^c(i)$

Names are pronouns (Cumming 2008; Schoubye 2016)

$[\![\text{Dan}]\!]^c = \lambda x\ .\ x \text{ IS A DAN}$

Names are predicates (Fara 2015)

$[\![\text{Dan}]\!]^c = \lambda \Phi\ .\ \text{DAN IS } \Phi$

Names are quantifiers (Montague 1973)

- All live options in the semantics literature
- Theorists have passionate disagreements about which is correct.
- At least some of these theorists' beliefs conflict with the contents of their semantic databases.
- This does not affect how they use language *at all*.

Central Cognition

- Theoretical inference
- Belief update
- Mindreading
- Practical reasoning
- (etc.)

- Quinean
- Isotropic
- Domain General
- Not amenable to computational modeling

Fast Operation

- Modules do their jobs extremely quickly.
  - Subjectively, visual perception is instantaneous.
  - Likewise language perception.
  - (Language production seems more complicated. More on this on Friday.)

Mandatory Operation

Mandelbaum (2014): "The Automatic and the Ballistic"

- Modular processing is mandatory.
- Two kinds of mandatoriness:
  - Automatic: process can't help but start when it encounters a stimulus of the relevant kind.
  - Ballistic: process can't help but finish once it starts
- Linguistic processing normally has both properties:
- Possible exception to ballisticity: contextual information facilitates the parsing of garden-path sentences.
  Crain & Steedman (1985). "On not being led up the garden path: The use of context by the psychological parser"
- Not clear that there are any exceptions with semantics.

"Shallow" Outputs

- Because models have limited information and representational vocabularies, they have limited ability to represent "high-level" properties to the things they represent.
- In visual perception, modularists tend to deny that we represent complex theoretical kinds in perception.
- In the context of semantics, I will be arguing that semantic the outputs of semantic processing are shallower than is usually assumed.
- In particular: semantics doesn't represent anything that depends on contextual information (Days 2+3).

Amenability to computational modeling

- Because modular systems have well-bounded proprietary databases, we have had greater success creating models of them.
- By contrast, the open-endedness of central cognition makes it very difficult to model in a realistic way.
- Fodor's examples are visual perception and syntax.
- But of course things are also going well with semantics!

Some Objections to the Modularity of Semantics

1. The context-sensitivity objection

In standard semantic theories, semantic values of many expressions are contents that depend on context.

But if semantics is the study of a modular system, then semantic composition has to operate on context-independent representations.

Is it really possible to maintain that the usual way of doing semantics is a progressive research program if this basic presupposition about context sensitivity is false?

For the answer, come back tomorrow!

2. The polysemy objection

In fact, it's not just the meanings of the usual "context-sensitive" expressions that a modular system can't handle.

It's the meanings of all polysemous vocabulary—namely, pretty much every open-class lexical item.

To understand what someone says with one of these words, we need to use contextual information to choose a sense.

How could that be compatible with the modular theory????

For the answer, come back on Wednesday!

3. The thinking-in-language objection

The idea that semantics is (part of) an "input-output system" makes it sound as though we only use language for communicating with others.

But that isn't right! We also use language to think!

If language is an input-output system, how is that possible?

For the answer, come back on Thursday!

4. The top-down control objection

When we speak, we *seem* to have fine-grained control over which words we use.

Moreover, there is a lot of evidence that we use information about our addressees and the context to customize what we say.

Doesn't that require cognitive penetration of the semantic module? How is that compatible with the modular theory????

For the answer, come back on Friday!

Thanks for today!

NASSLLI 2022 · USC · 20–24 JUNE
SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

DAY 2:
CONTEXT SENSITIVITY

Daniel W. Harris
CUNY Graduate Center and Hunter College

1. The context-sensitivity objection

In standard semantic theories, semantic values of many expressions are contents that depend on context.

But if semantics is the study of a modular system, then semantic composition has to operate on context-independent representations.

Is it really possible to maintain that the usual way of doing semantics is a progressive research program if this basic presupposition about context sensitivity is false?

Content Semantics

Content Semantics

- Roughly: Any semantic theory according to which semantic values are contents, in roughly Kaplan's (1989) sense.
- The content of a (declarative) sentence is the proposition that one would assert by uttering it in a literal way.
- The contents of sub-sentential expressions are their contributions to the propositions expressed by sentences in which they appear: things like referents and properties.

A world-relativized truth condition

(States the conditions under which the sentence is true relative to a possible world, an assignment function, and a context of utterance.)

$$[\![\text{Ann smokes}]\!]^{w,g,c} = 1 \text{ iff Ann smokes at } w$$

$$[\![\text{Ann smokes}]\!]^{g,c} = \lambda w \,.\, \text{Ann smokes at } w$$

A proposition

A function that maps possible worlds to truth values. A popular way of formally modeling a piece of informational content—the kind of thing that a person can believe or assert.

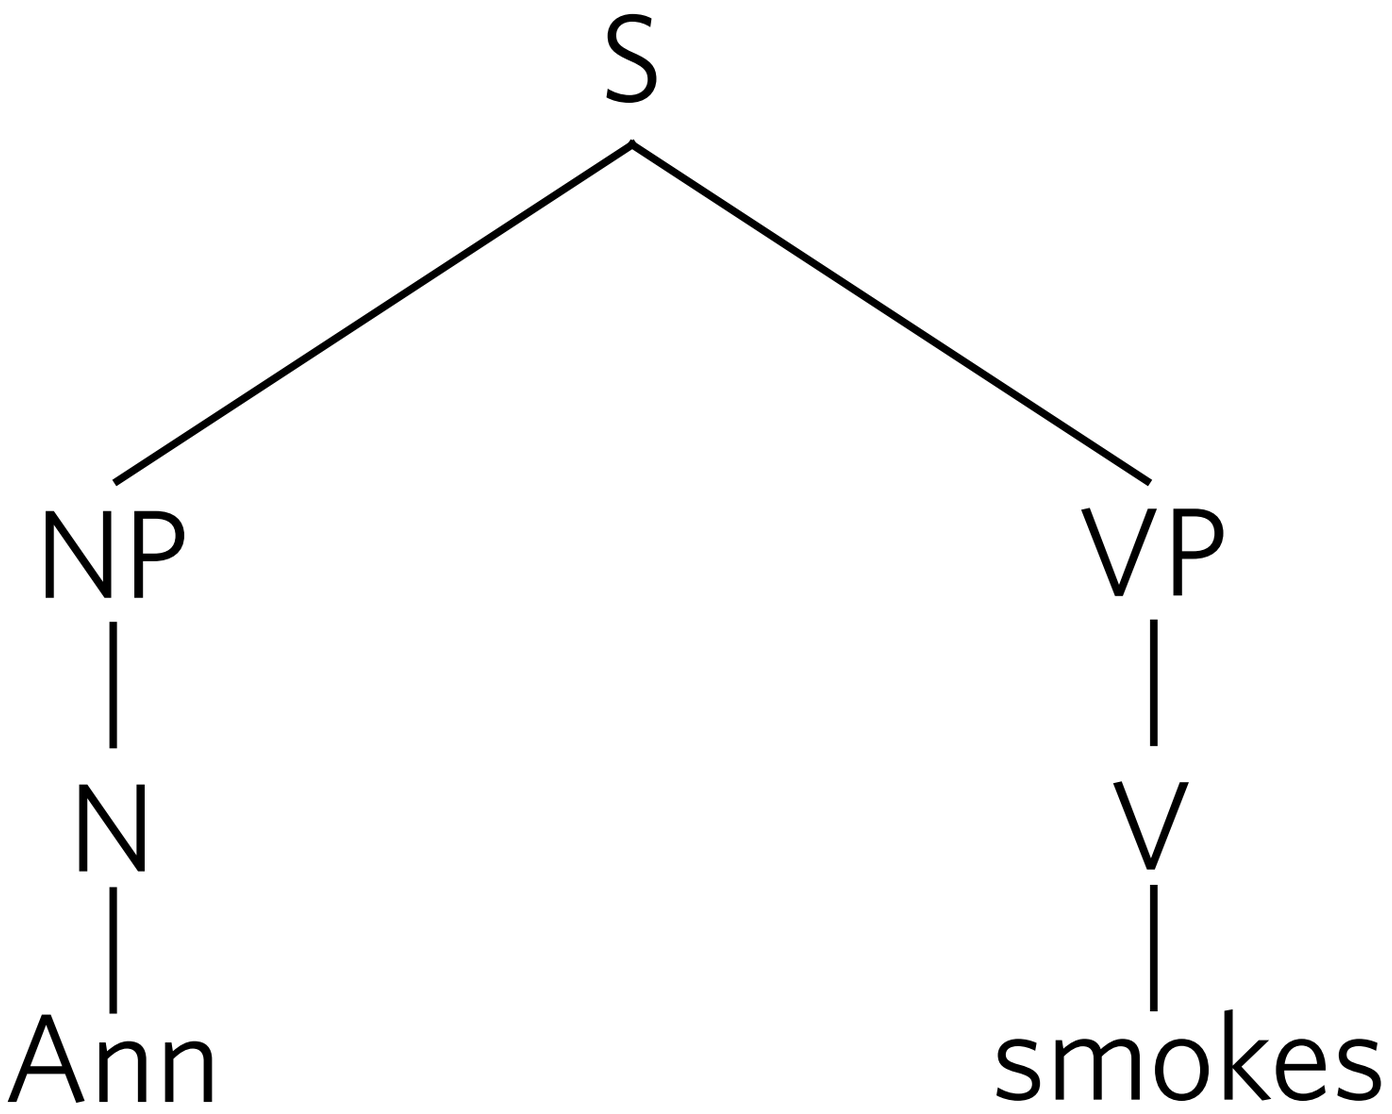

An (extremely simplified) LF

A phrase structure diagram which models the syntactic input to semantic processing during language perception.

$[\![\text{Ann smokes}]\!]^{w,g,c} = 1$ iff Ann smokes at w

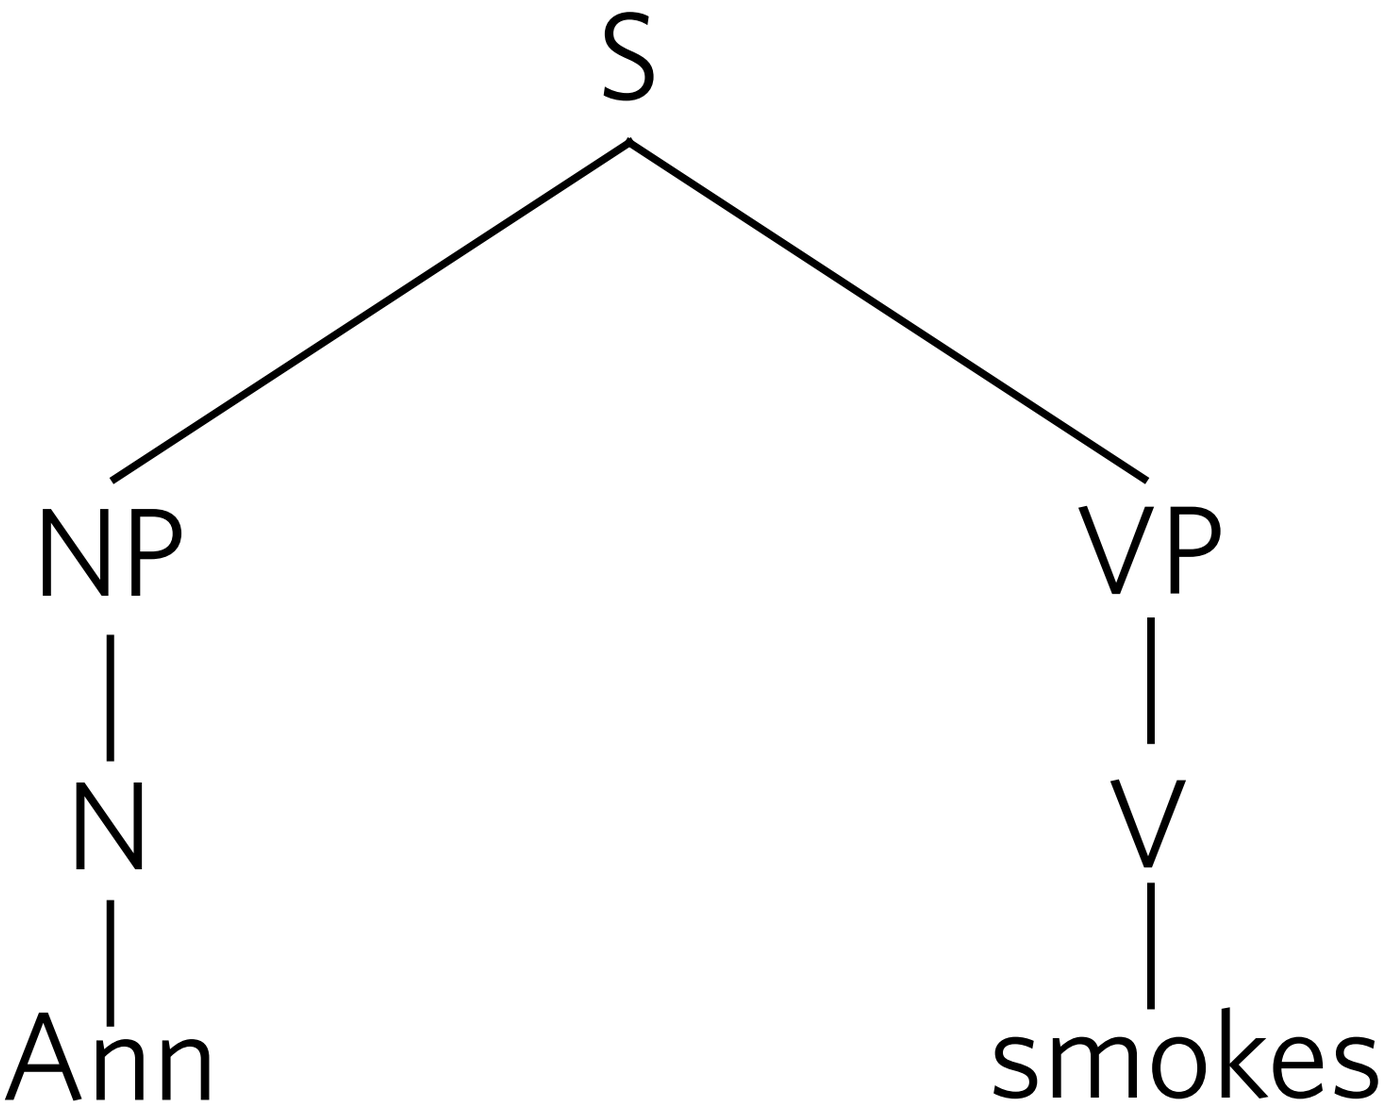

⟦Ann smokes⟧^{w,g,c} = 1 iff Ann smokes at w

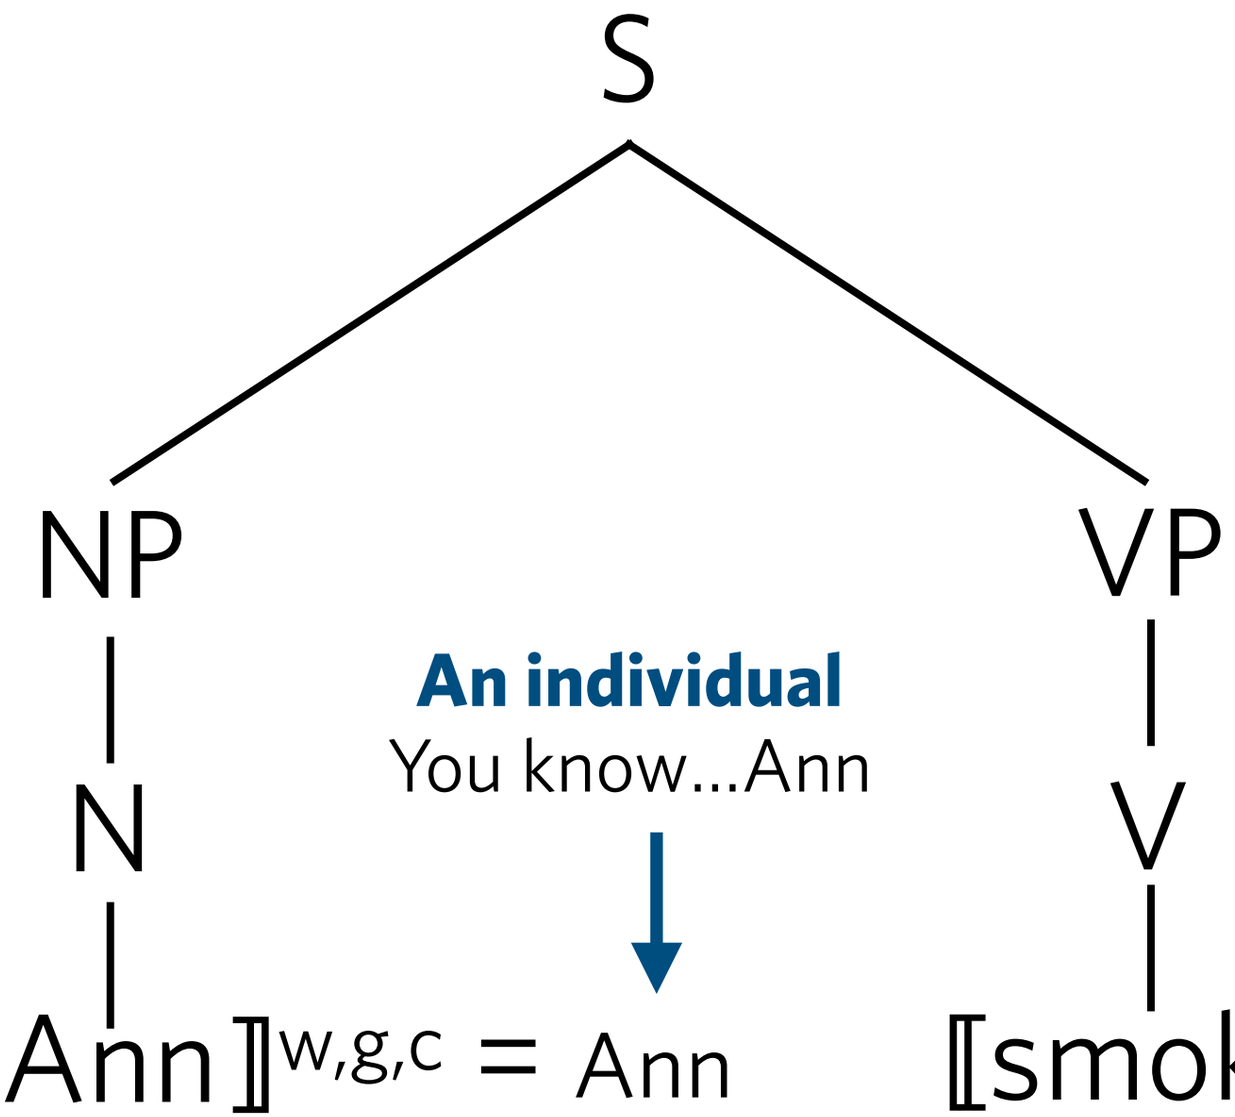

An individual
You know…Ann

A world-relativized property
Modeled as a function that takes individuals as inputs and maps them to truth (“1”) iff they smoke.

$[\![\text{Ann smokes}]\!]^{w,g,c}$ = 1 iff Ann smokes at w

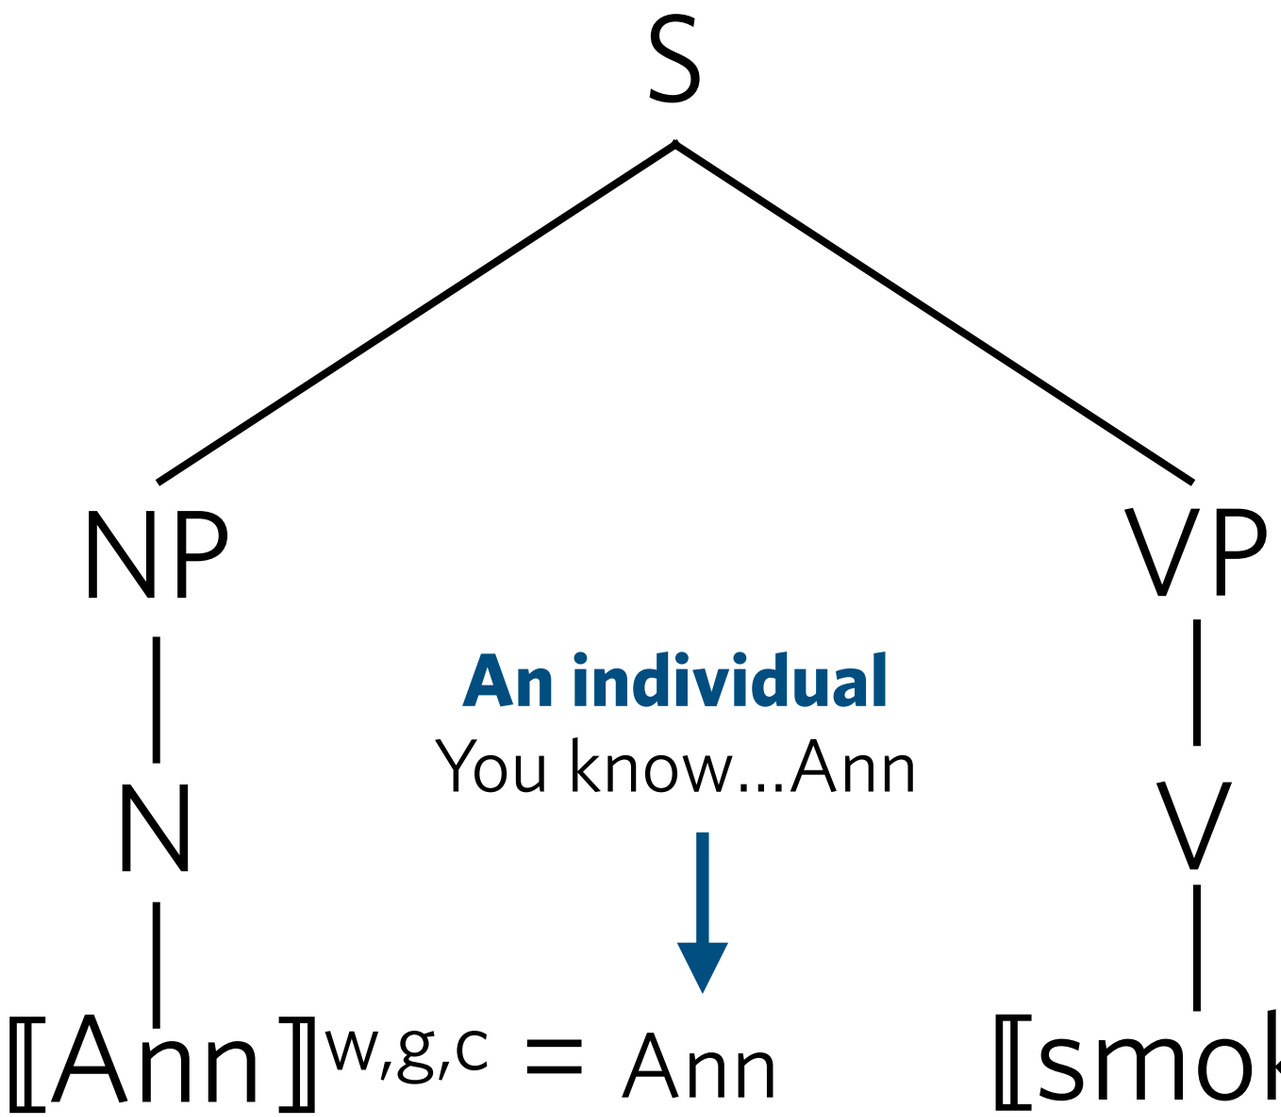

An individual
You know…Ann

A world-relativized property
Modeled as a function that takes individuals as inputs and maps them to truth ("1") iff they smoke.

Functional Application

$[\![\alpha\beta]\!]^{w,g,c} = [\![\alpha]\!]^{w,g,c}([\![\beta]\!]^{w,g,c})$ or $[\![\beta]\!]^{w,g,c}([\![\alpha]\!]^{w,g,c})$

To derive the content of a compound expression, apply the function denoted by one to the argument denoted by the other.

Context Sensitivity

Kaplan (1989): "Demonstratives"

An assignment function
A mapping from numerical indices on variables to entities.

g, c

A context
A formal model of the extralinguistic context of utterance

$$c = \langle \text{speaker}_c, \text{time}_c, \text{location}_c, \ldots \rangle$$

$$[\![\text{I}]\!]^{w,g,c} = \text{speaker}_c$$

$$[\![\text{now}]\!]^{w,g,c} = \text{time}_c$$

$$[\![\text{here}]\!]^{w,g,c} = \text{location}_c$$

Assignment Sensitivity

An assignment function
A mapping from numerical indices on variables to entities. → g,c ← **A context** A formal model of the extralinguistic context of utterance

$[\![he_i]\!]^{w,g,c} = g(i)$ (if g(i) is male; else undefined)

$[\![he_i \text{ smokes}]\!]^{w,g,c} = 1$ iff g(i) smokes at w

(as long as g(i) is male)

$[\![[\text{Every doctor}]_i \text{ denies that he}_i \text{ smokes}]\!]^{w,g,c} =$
1 iff Every doctor smokes at w

"roughly, what is meant by 'variable binding' is any semantic operation which removes (or reduces) assignment dependency" (Heim & Kratzer, p. 116).

How is the operative assignment function determined?

"...let us think of assignments as representing the contribution of the utterance situation. The physical and psychological circumstances that prevail when an LF is processed will (if the utterance is felicitous) determine an assignment to all the free variables occurring in this LF."

—Heim & Kratzer, p.243

"For free pronouns, the relevant assignment is given by the utterance context and represents the speaker's referential intentions. "

—Heim (2008): p.36

$$\mathbf{c} = \langle \text{speaker}_c, \text{time}_c, \text{location}_c, \mathbf{g_c} \ldots \rangle$$

Content Semantics

- The semantic value of an expression is its *content*.
  - Declarative sentences → propositions
  - referring expressions → referents
  - predicates → properties
- These contents are the things that get composed with each other.
- But since some expressions are context sensitive, the interpretation function depends on information about extralinguistic context.
- And so whatever system in our mind is doing the semantic computations needs that information too.

Constraint Semantics

- If semantics is the study of a modular system, then that system doesn't have access to information about extralinguistic context.
    - (Note that informational encapsulation is doing the work here.)
- So the semantic module can't compute the contents of context-sensitive expressions.
- So it can't compose those contents to get the contents of context-sensitive sentences.
- So, content semantics is incompatible with the idea that semantics is the study of a modular system.

Constraint Semantics

Constraint Semantics

- Roughly: Any semantic theory according to which semantic values are *partial* representations of contents.
- The meaning of a sentence tells us only what we could know about what a person would be saying with it if we didn't know anything about the context or their intentions, but assumed that they were speaking literally.
- Imagine finding a note that says, "he's here," without having any idea who *he* is or where *here* is.
- That's the sort of epistemic situation that your semantic module is in all the time.

Constraint Semantics

- A sentence's "semantic representation is a schema, which must be completed and integrated into an assumption about the speaker's informative intention" (Sperber & Wilson, Relevance, p.175)

- the meanings of declarative sentences "constrain without determining truth/ reference/satisfaction conditions" (Pietroski 2006, p. 34)

- "... the semantics of an expression gives the information that a competent speaker can glean from it independently of any context of utterance" (Bach 1987, p. 5)

- A sentence's semantic value is "a blueprint for (a template, a schematic or skeletal representation of) what someone will be taken to be saying when using [the sentence] to say something" (Neale 2005, p. 189).

See also work by Schiffer, Carston, Garcia-Carpintero, etc.

Implementing Constraint Semantics?

Some options that would require redoing a lot of stuff:

- Pietroski (2018), *Conjoining Meanings: Semantics without Truth Values*
- Variable-free semantics? (Jacobson 1999, 2014)
- Alternative semantics (Kratzer and Shimoyama 2002), but for all context-sensitive expressions?

These are all cool ideas that we should explore.

But do we *have to* revise the basics in order to have constraint semantics?

Constraints = Characters?

A context-relativized proposition

$$[\![\text{he}_i \text{ talks}]\!]^c = \lambda w.\ g_c(i) \text{ talks at } w$$

(as long as g(i) is male)

$$[\![\phi]\!] = \lambda c\ .\ [\![\phi]\!]^c$$

$$[\![\text{he}_i \text{ talks}]\!] = \lambda c\ .\ \lambda w.\ g_c(i) \text{ talks at } w$$

A context-insensitive character

A function that maps each context to the proposition that the sentence expresses in that context.

Constraints = Characters?

$[\![\text{he}_i \text{ talks}]\!] = \lambda c . \lambda w. \mathbf{g_c(i)}$ talks at w

The Semantic Module's "Vocabulary"

g^c
assignment function

i
numerical index

- These concepts play critical roles in semantic composition (on many views).
- But ordinary speakers do not possess them, in the usual sense.
- Non-semanticists can't have beliefs about them.
- Heterodox theorists who don't believe in them don't speak differently.

Implementing Constraint Semantics?

#Goals:

- Leave the system from Heim & Kratzer unchanged.
- Define an abstraction operation that turns their semantic values into constraints.
- Unlike characters, these constraints shouldn't be framed in terms of indices or assignments.
- That way, indices and assignments can be thought of as internal book-keeping devices of the semantic module.

Constraints, a first draft

(So far, we're ignoring variables different meanings.)

An index-relativized proposition

$$[\![\text{it}_1 \text{ stinks}]\!]^g = \lambda w \,.\, g(1) \text{ stinks at } w$$

$$[\![[\![\text{it}_1 \text{ stinks}]\!]]\!] = \lambda p_{st} \,.\, (\exists x) p = \lambda w \,.\, x \text{ stinks at } w$$

A constraint

A function that maps each proposition to "true" if and only if it is the sort of proposition that one could literally say with the sentence.

In English: The semantic value of "it stinks" is a property possessed by any proposition p such that, for some x, p is the proposition that x stinks.

Constraint Abstraction, first draft

$$[\![\text{it}_1 \text{ stinks}]\!]^g = \lambda w \,.\, g(1) \text{ stinks at } w$$

$$[\![[\![\text{it}_1 \text{ stinks}]\!]]\!] = \lambda p_{st} \,.\, (\exists x) p = \lambda w \,.\, x \text{ stinks at } w$$

Constraint Abstraction (preliminary version)

If α dominates unbound variables $v_i \ldots v_n$ and $[\![\alpha]\!]^g \in D_\tau$

Then $[\![\alpha]\!] = \lambda p_\tau . (\exists x^i) \ldots (\exists x^n) p = [\![\alpha]\!]^{g^{i \to x^i \ldots n \to x^n}}$

Constraint Abstraction, first draft

Constraint Abstraction (preliminary version)

If α dominates unbound variables $v_i ... v_n$ and $[\![\alpha]\!]^g \in D_\tau$

Then $[\![\alpha]\!] = \lambda p_\tau.(\exists x^i)...(\exists x^n)p = [\![\alpha]\!]^{g^{i \to x^i ... n \to x^n}}$

Two problems with this draft:

- It is not compositional (at least strictly speaking)
    - It needs info about syntactic constituents, and not just the semantic values of its daughters
- It doesn't distinguish the meanings of different variables.

Constraints, second draft

If all I know about your utterance is that you used “he smokes” literally, what can I know about what you have said? Roughly, that it is a proposition p with the following property: For some *male* individual x, p is the proposition that x smokes.

$[\![\text{he smokes}]\!] = \lambda p_{st} \,.\, (\exists x_e : x \text{ is male})(p = \lambda w \,.\, x \text{ smokes at } w)$

$[\![\text{she smokes}]\!] = \lambda p_{st} \,.\, (\exists x_e : x \text{ is female})(p = \lambda w \,.\, x \text{ smokes at } w)$

Constraint Properties

$$\mu(\text{he}) = \lambda x_e \,.\, x \text{ is male}$$

$$\mu(\text{she}) = \lambda x_e \,.\, x \text{ is female}$$

- These are specified for each variable in the lexicon
- They give the restrictions on what each variable can be used to say

A Multidimensional Semantics

The usual H&K semantic value goes here

For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$

A container for storing constraint properties for later use

A Multidimensional Semantics

The usual H&K semantic value goes here

For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$

A container for storing constraint properties for later use

For any variable v and numerical index i, $[\![v_i]\!]_2 = \{\langle i, \mu(v)\rangle\}$

For any non-variable lexical item α, $[\![\alpha]\!]_2 = \emptyset$.

A Multidimensional Semantics

For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$

Functional Application

$[\![\alpha\beta]\!]^{w,g} = \langle [\![\alpha\beta]\!]_1^{w,g}, [\![\alpha\beta]\!]_2 \rangle$, such that:

- $[\![\alpha\beta]\!]_1^{w,g} = [\![\alpha]\!]_1^{w,g}\left([\![\beta]\!]_1^{w,g}\right)$ or $[\![\beta]\!]_1^{w,g}\left([\![\alpha]\!]_1^{w,g}\right)$ (whichever is defined); and
- $[\![\alpha\beta]\!]_2 = [\![\alpha]\!]_2 \uplus [\![\beta]\!]_2$

The usual

An operation for combining and collating constraint properties

A Multidimensional Semantics

For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]_1^{w,g}, [\![\alpha]\!]_2 \rangle$

Functional Application

$[\![\alpha\beta]\!]^{w,g} = \langle [\![\alpha\beta]\!]_1^{w,g}, [\![\alpha\beta]\!]_2 \rangle$, such that:

- $[\![\alpha\beta]\!]_1^{w,g} = [\![\alpha]\!]_1^{w,g}\left([\![\beta]\!]_1^{w,g}\right)$ or $[\![\beta]\!]_1^{w,g}\left([\![\alpha]\!]_1^{w,g}\right)$ (whichever is defined); and
- $[\![\alpha\beta]\!]_2 = [\![\alpha]\!]_2 \uplus [\![\beta]\!]_2$

The usual

$\{\langle i,\phi\rangle\} \uplus \{\langle j,\psi\rangle\} = \{\langle i,\phi\rangle,\langle j,\psi\rangle\}$

$\{\langle i,\phi\rangle\} \uplus \{\langle i,\psi\rangle\} = \{\langle i,\phi\wedge\psi\rangle\}$

A Multidimensional Semantics

For every node α in every LF, $[\![\alpha]\!]^{w,g} = \langle [\![\alpha]\!]^{w,g}_1, [\![\alpha]\!]_2 \rangle$

Functional Application

$[\![\alpha\beta]\!]^{w,g} = \langle [\![\alpha\beta]\!]^{w,g}_1, [\![\alpha\beta]\!]_2 \rangle$, such that:

- $[\![\alpha\beta]\!]^{w,g}_1 = [\![\alpha]\!]^{w,g}_1 \left([\![\beta]\!]^{w,g}_1 \right)$ or $[\![\beta]\!]^{w,g}_1 \left([\![\alpha]\!]^{w,g}_1 \right)$ (whichever is defined); and
- $[\![\alpha\beta]\!]_2 = [\![\alpha]\!]_2 \uplus [\![\beta]\!]_2$

The usual

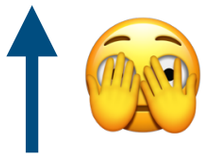

Definition of $\uplus$

For any numerical indices $x^1 \ldots x^n$, $y^1 \ldots y^m$ and functions $X^1 \ldots X^n$, $Y^1 \ldots Y^m$, if $\chi = \{\langle x^1, X^1 \rangle \ldots \langle x^n, X^n \rangle\}$ and $\gamma = \{\langle y^1, Y^1 \rangle \ldots \langle y^m, Y^m \rangle\}$, then $\chi \uplus \gamma$ is defined as follows:

For every $\langle x^i, X^i \rangle \in \chi$ and $\langle y^j, Y^j \rangle \in \gamma$:

- if $x^i = y^j$ and $\exists \tau : X^i, Y^j \in D_{\tau t}$ then $\langle x^i, \lambda y_\tau \,.\, X^i(y) \wedge Y^j(y) \rangle \in \chi \uplus \gamma$
- if $x^i \neq y^j$, then $\langle x^i, X^i \rangle \in \chi \uplus \gamma$ and $\langle y^j, Y^j \rangle \in \chi \uplus \gamma$

Nothing else is in $\chi \uplus \gamma$.

Constraint Abstraction, second draft

$$[\![\text{he}_1 \text{ smokes}]\!]^g = \begin{cases} [\![\text{he}_1 \text{ smokes}]\!]_1{}^g = \lambda w \,.\, g(1) \text{ smokes in } w \\ [\![\text{he}_1 \text{ smokes}]\!]_2 = \{\langle 1, \lambda x \,.\, x \text{ is male}\rangle\} \end{cases}$$

Constraint Abstraction (final version)

If $[\![\alpha]\!]_2 = \{\langle i, f^1\rangle \ldots \langle j, f^n\rangle\}$ and $[\![\alpha]\!]_1^g \in D_\tau$

Then $[\![[\![\alpha]\!]]\!] = \lambda p_\tau \,.\, (\exists x^1 : f^1(x^1)) \ldots (\exists x^n : f^n(x^n))(p = [\![\alpha]\!]_1^{g^{i \to x^1 \ldots j \to x^n}})$

$$[\![[\![\text{he smokes}]\!]]\!] = \lambda p_{st} \,.\, (\exists x_e : x \text{ is male})(p = \lambda w \,.\, x \text{ smokes at } w)$$

Indexicals

$\mu(\text{I}) = \lambda x_e$. x is the speaker of the utterance being interpreted

$\mu(\text{you}) = \lambda x_e$. x is the addressee of the utterance being interpreted[33]

$\mu(\text{today}) = \lambda x_e$. x is the day on which the utterance being interpreted was produced

$\mu(\text{now}) = \lambda x_e$. x is an interval of time during which the utterance being interpreted was produced

$\mu(\text{here}) = \lambda x_e$. x is a location in space within which the utterance being interpreted was produced[34]

$\mu(\text{this}) = \lambda x_e$. x is "proximal" from the perspective of the speaker of the utterance being interpreted

$\mu(\text{that}) = \lambda x_e$. x is "distal" from the perspective of the speaker of the utterance being interpreted

Some tentative conclusions

- Modularity → constraint semantics
- There are some cool and ambitious ways that we could pursue constraint semantics.
- But we don't *have to be* cool and ambitious to pursue it.
- We can non-destructively modify even a vanilla textbook semantic theory to get a (somewhat hacky) version of what we want.

Thanks for today!

NASSLLI 2022 · USC · 20–24 JUNE
SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

DAY 3:
POLYSEMY, WORD MEANINGS, AND CONCEPTS

Daniel W. Harris
CUNY Graduate Center and Hunter College

2. The polysemy objection

In fact, it's not just the meanings of the usual "context-sensitive" expressions that a modular system can't handle.

It's the meanings of all polysemous vocabulary—namely, pretty much every open-class lexical item.

To understand what someone says with one of these words, we need to use contextual information to choose a sense.

How could that be compatible with the modular theory??

For the answer, come back on Wednesday!

Word Meanings as Concepts

Ambiguity

Ambiguity

John is at the bank.

Ambiguity

John is at the bank.

Question:
How does a modular language system choose the right meaning?

Ambiguity

John is at the bank.

Hypothesis 1:
It picks one randomly, or using some heuristic, and tries again if the result doesn't make sense.

Compare: Structural Ambiguity

Rayner, K., Carlson, M., and Frazier, L. (1983). "The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences"

(49) The spy saw the cop with a revolver, but the cop didn't see him.

(50) The spy saw the cop with the binoculars, but the cop didn't see him.

Compare: Structural Ambiguity

Rayner, K., Carlson, M., and Frazier, L. (1983). "The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences"

???

(49) The spy saw the cop with a revolver, but the cop didn't see him.

(50) The spy saw the cop with the binoculars, but the cop didn't see him.

Compare: Structural Ambiguity

Rayner, K., Carlson, M., and Frazier, L. (1983). "The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences"

???

(49) The spy saw the cop with a revolver, but the cop didn't see him.

(50) The spy saw the cop with the binoculars, but the cop didn't see him.

An explanation: (49) violates the minimal attachment principle, which the parser uses as a heuristic when building phrase structures.

Same for lexical ambiguity?

Hirst, G. (1987). "Semantic interpretation and the resolution of ambiguity." Cambridge University Press.

???

(51) The astronomer married the star.

(52) The catcher filled the pitcher.[30]

???

Same for lexical ambiguity?

Hirst, G. (1987). "Semantic interpretation and the resolution of ambiguity." Cambridge University Press.

???

(51) The astronomer married the star.

(52) The catcher filled the pitcher.[30]

???

Two hypotheses:

Semantic priming causes us to choose the wrong meaning, and then we have to reanalyze.

We compose both meanings in parallel and then CG has to reject one, which can take extra time because of priming.

Polysemy

A word is polysemous if it has multiple *related* senses.

Usually distinguished from homonymy, in which a single word form has multiple unrelated senses.

bank·er[1] (băng′kər) *n.* **1.** One serving as an officer or owner of a bank. **2.** *Games.* The player in charge of the bank in some gambling games. **—bank′er·ly** *adj.*
bank·er[2] (băng′kər) *n.* One engaged in cod fishing off Newfoundland.

(American heritage dictionary.)

Polysemy

Jeff Bezos bought three **newspapers**.

Hot Take: All Ambiguity is the Same

$[\![\text{banker}^1]\!]^w = \lambda x$. x works at a bank in w

$[\![\text{banker}^2]\!]^w = \lambda x$. x plays the banker role in a board game in w

$[\![\text{banker}^3]\!]^w = \lambda x$. x works catching fish in the Grand Bank in w

Why not treat all ambiguity the same?

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

Table 1
Examples of polysemy in English.

Patterns and their Senses	Examples
Animal for Meat	The *chicken* walked on the grass/
(*chicken*, *turkey*, *fish*, etc.)	The *chicken* was well-salted
Material for Artifact	There is broken *glass* on the floor/
(*glass*, *tin, iron*, etc.)	She drank milk from the *glass*
Object for Representational Content	The *book* is very light to carry/
(*book*, *magazine*, *DVD*, etc.)	The *book* is very interesting
Container for Contents	She washed the *pot* after dinner/
(*pot*, *bowl*, *box*, etc.)	She stirred the *pot* with a spoon
Body Part for Object Part	He broke his *leg* last year/
(*leg*, *arm*, *back*, etc.)	That chair has a broken *leg*
Artist for Product	*Picasso* was born in 1881/
(*Picasso*, *Camus*, *Mozart*, etc.)	That museum has a *Picasso*
Place for Institution	The *White House* is being renovated/
(*White House*, *Wall Street*, *City Hall*, etc.)	The *White House* should make a decision
Place for Event	*Vietnam* shares a border with China/
(*Vietnam*, *Waterloo*, *Woodstock*, etc.)	He championed civil rights during *Vietnam*
Substance for Placing Substance at Goal	He bought some *butter* from the store/
(*butter*, *salt*, *water*, etc.)	He is going to *butter* the bread
Instrument for Action Involving Instrument	She has a red *shovel*/
(*shovel*, *hammer*, *rake*, etc.)	She is going to *shovel* the snow

Why not treat all ambiguity the same?

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

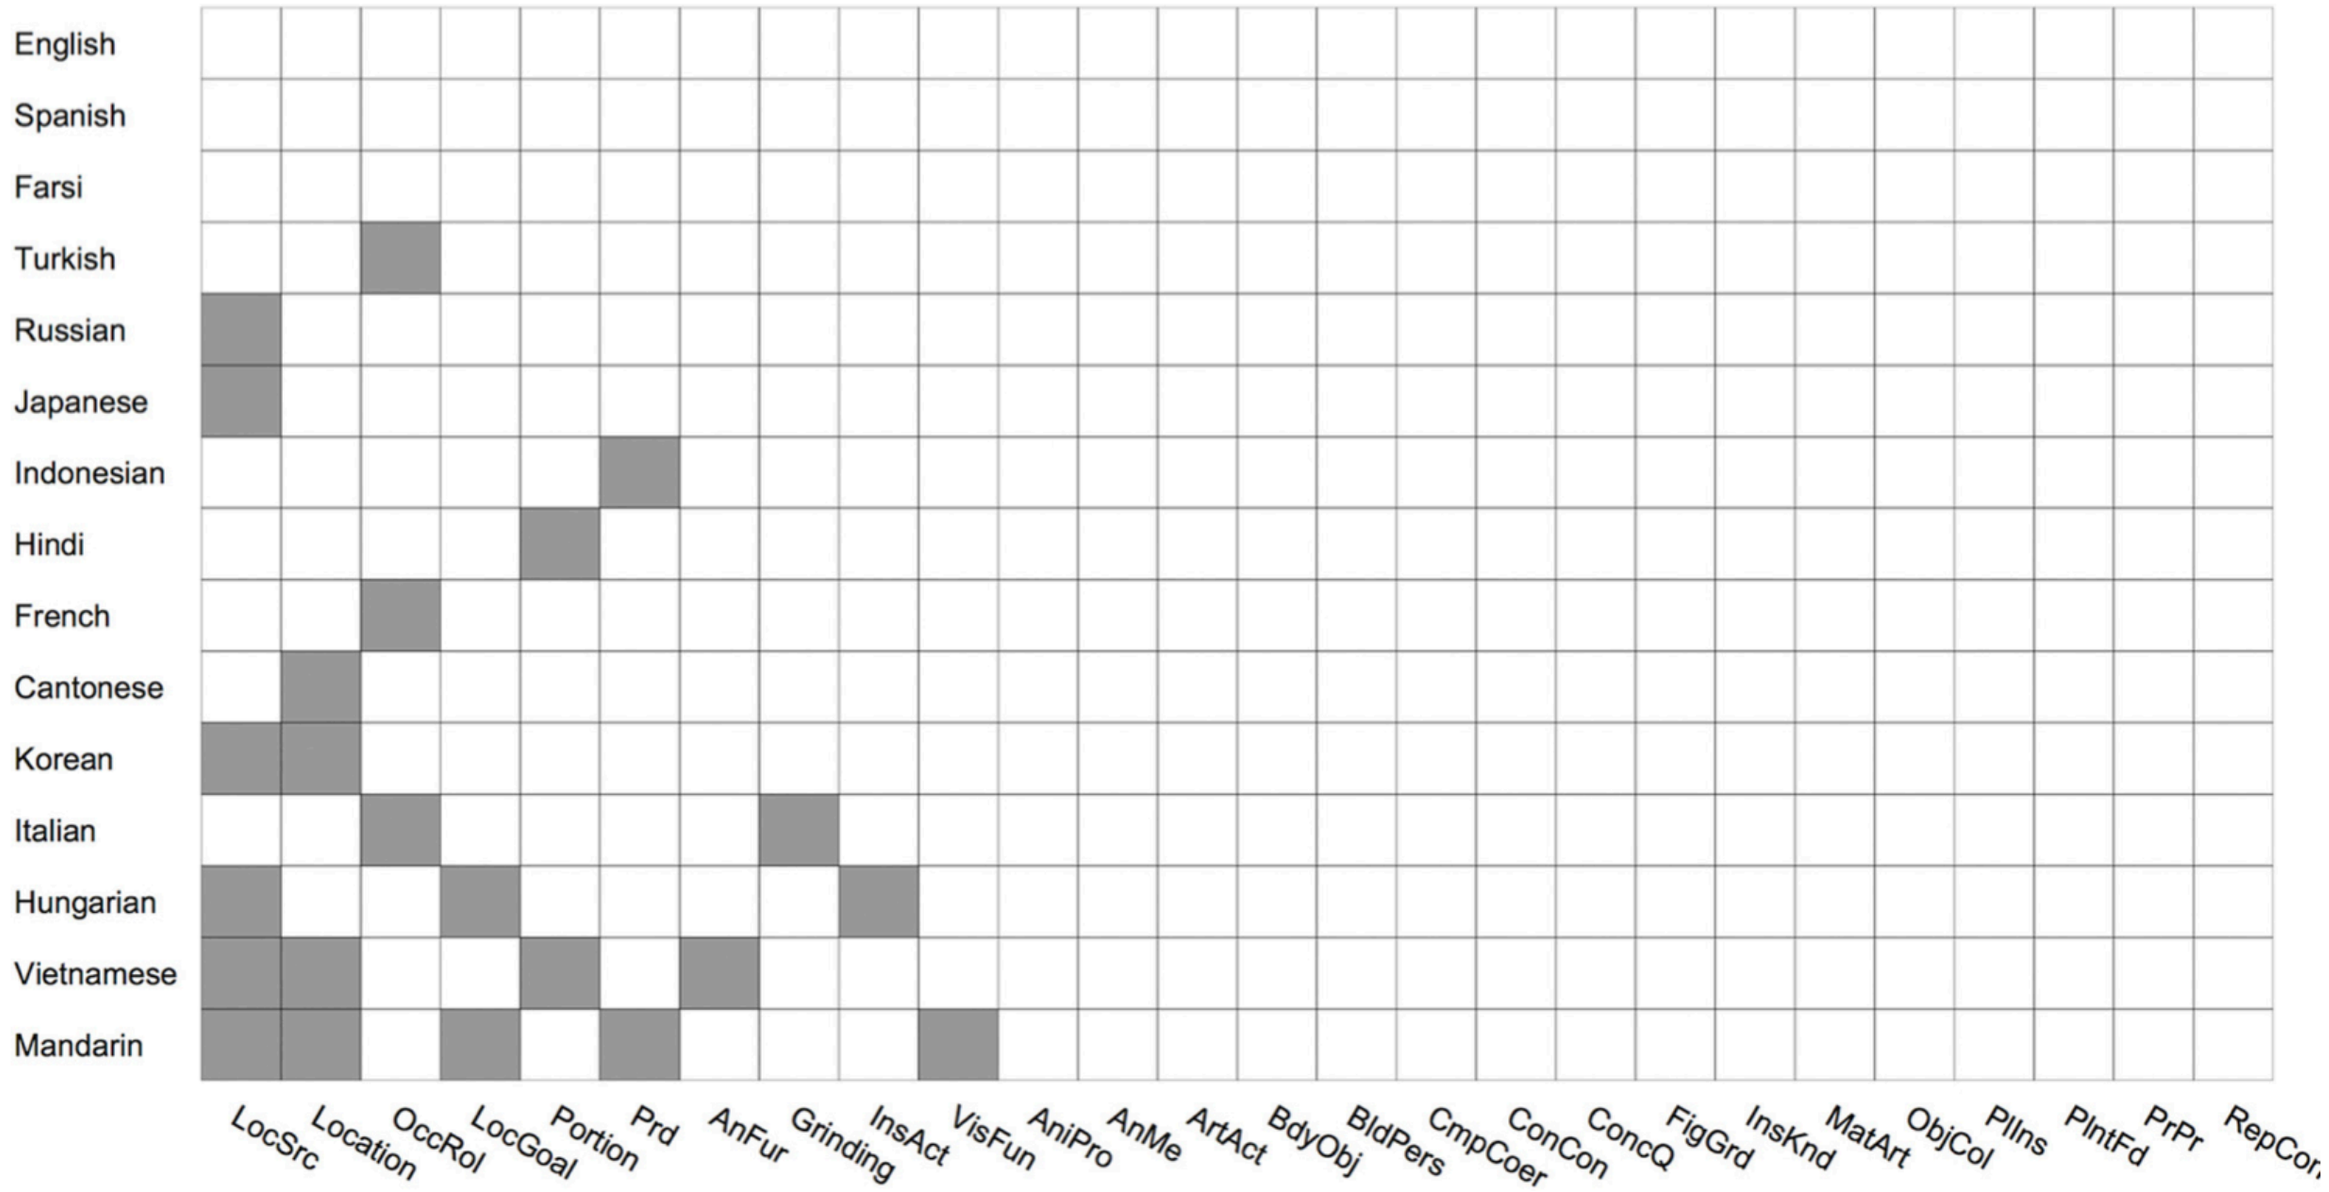

Fig. 2. Evidence for the presence of patterns across languages. A white box indicates evidence for a pattern in a language, and a gray box indicates no evidence for the pattern. The figure is ordered such that languages with evidence for more patterns are toward the top, and patterns that are attested across more languages are toward the right. The place for event pattern is excluded from this analysis (see footnote 5). Table 4 contains a legend for the pattern names.

Why not treat all ambiguity the same?

Liina Pylkkänen, Rodolfo Llinás, and Gregory L. Murphy (2006):
"The Representation of Polysemy: MEG Evidence," Journal of Cognitive Neuroscience

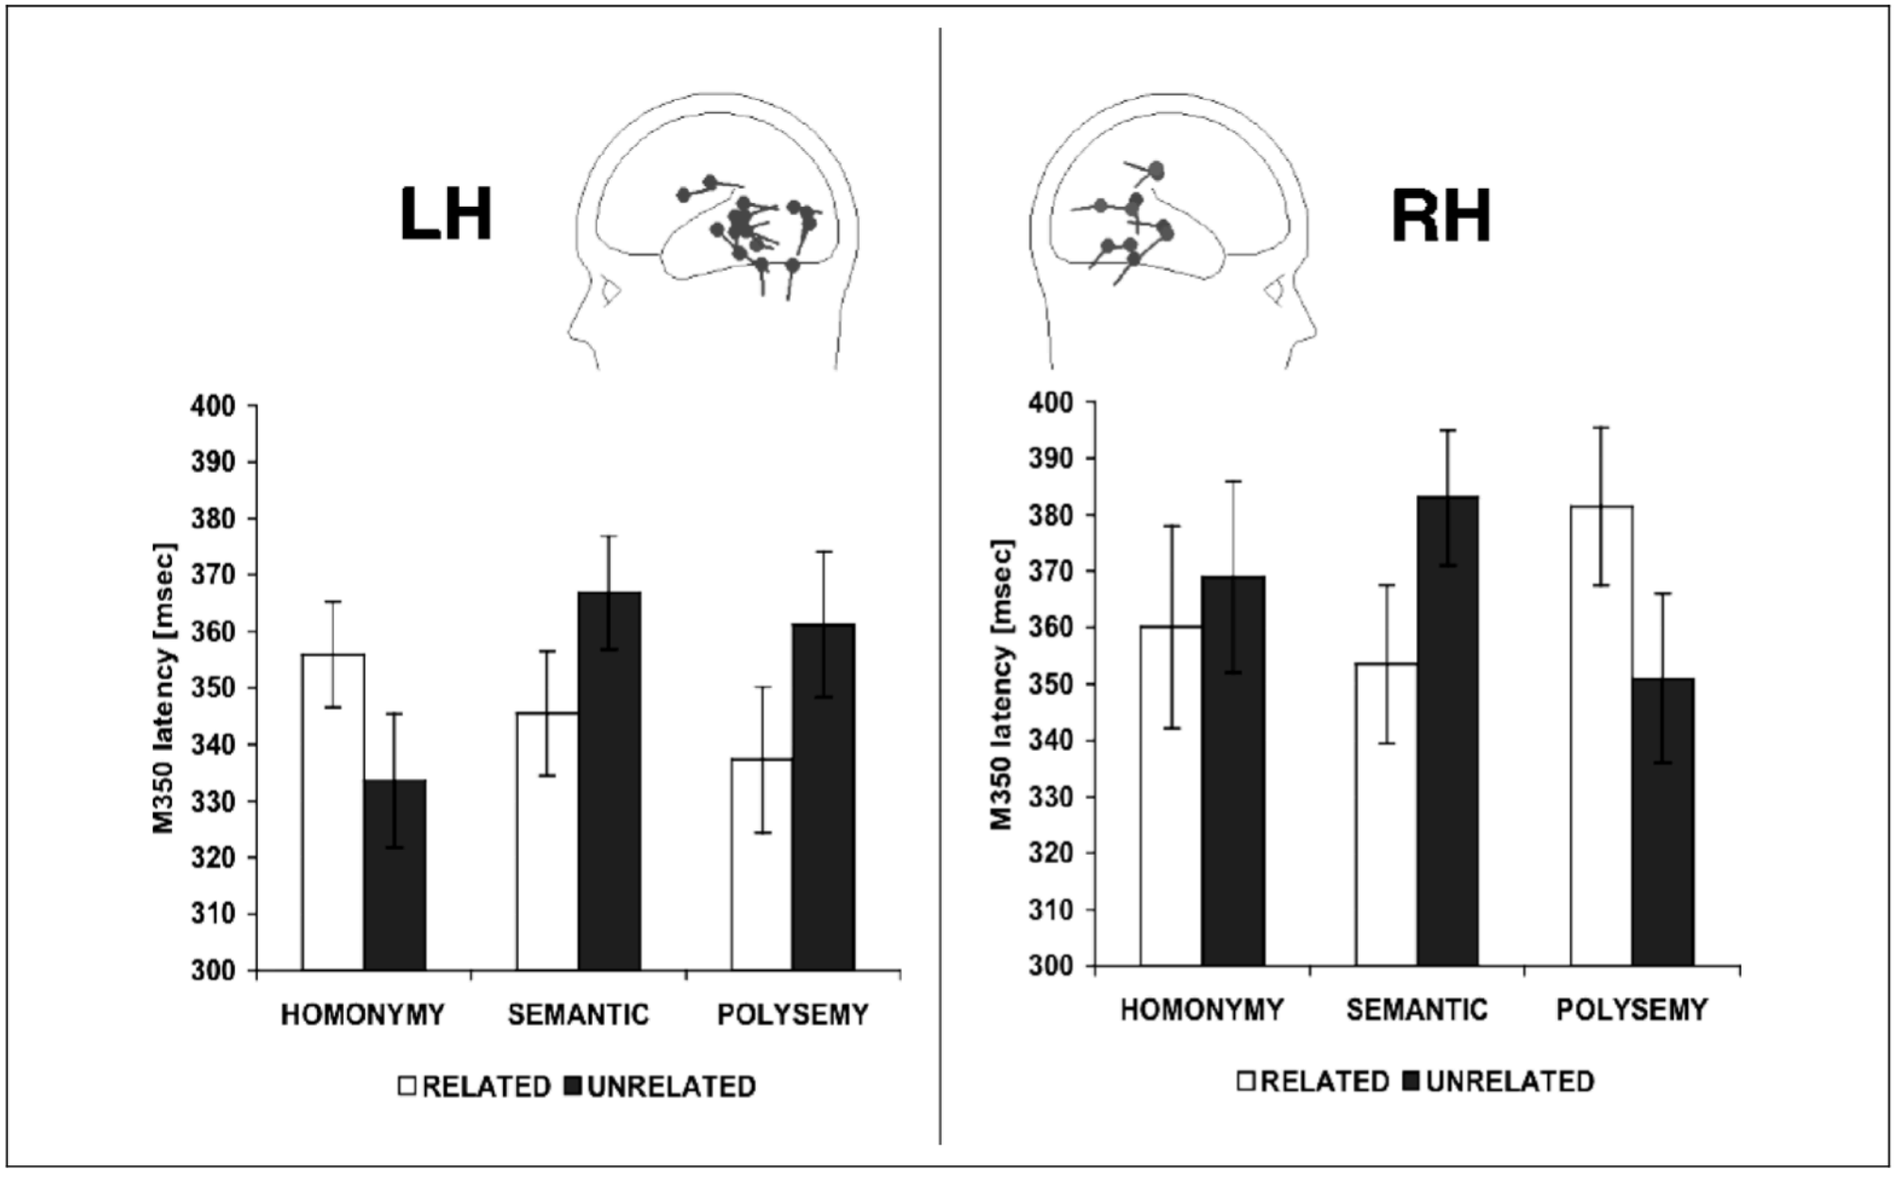

Figure 4. Effect of relatedness on M350 latencies. Error bars are *SD*s. The critical data are the priming effects (differences between related and unrelated conditions) for each word type.

Why not treat all ambiguity the same?

Chelsea M. Eddington and Natasha Tokowicz (2015): "How meaning similarity influences ambiguous word processing: the current state of the literature," *Psychonomic Bulletin & Review*, volume 22, pages 13–37

"It is likely that early research that showed an ambiguous word advantage in processing actually reflected an advantage for polysemous and not homonymous words."

Second Take:

Polysemy = 1 lexical item with two meanings
Homonymy = 2 lexical items

"Banker[1]" is now a variable whose meaning is given as a constraint property, which lists off the different properties that one can literally express by uttering "banker[1]".

$\mu(\text{banker}^1) = \lambda\phi_{et}$. [$\phi = \lambda w$. λx . x works at a bank in w] ∨
[$\phi = \lambda w$. λx . x plays the banker role in a board game in w]

$[\![\text{banker}^2]\!]^w = \lambda x$. x works catching fish in the Grand Bank in w

Second Take:

Polysemy = 1 lexical item with two meanings
Homonymy = 2 lexical items

$\mu(\text{banker}^1) = \lambda\phi_{et}$. [$\phi = \lambda w$. λx . x works at a bank in w] $\vee$
[$\phi = \lambda w$. λx . x plays the banker role in a board game in w]

$[\![\text{He}_i \text{ is a banker}^1_j]\!]^w = \lambda p_{st}$. $(\exists x : x \text{ is male})(\exists \phi : \phi = F \vee \phi = G)$ p = λw . x is ϕ in w

A property shared by every proposition p such that, for some male x and some property ϕ (one of the two senses of banker^1), p is the proposition that x is ϕ.

Against lists of meanings?

Charles Travis (1997): "Pragmatics"

(1) The leaves are green.

"Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood."

Against lists of meanings?

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

"Polysemy is important not only because it is ubiquitous, but also because it provides a source of linguistic creativity: to express new ideas, we needn't invent new words, but can instead extend existing words beyond their original meanings. In English, such creativity has yielded systematic patterns of senses: for instance, the same words are often used to label an animal or its meat (e.g., chicken, lamb, etc.), or a material and an artifact derived from that material (e.g., glass, tin, etc.)."

Against lists of meanings?

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

The book **took two years to write** and **has 400 pages**.

Third Take:

Polysemy = lexical item with abstract or structured meanings

$$\mu(\text{green}) = \lambda\phi_{et} \,.\, \phi = \Sigma\phi$$

When you hear someone say "green," your language module tells you that they're talking about a property of a certain kind Σ, but you have to figure out which one.

What is Σ?

It's whatever (possibly very complex, structured, theoretical) constraint has to be met by a property in order to be literally expressible with the word "green."

Cf. work by James Pustejovsky, Augustín Vicente, Deirdre Wilson, Robyn Carston, etc.

Third Take:

Polysemy = lexical item with abstract or structured meanings

$$\mu(\text{green}) = \lambda\phi_{et} \,.\, \phi = \Sigma\phi$$

Okay, but does this (possibly complex, structured, theoretical) information all have to be stored inside the semantic module?

It sounds like the module has a lot of world knowledge then!

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$

$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

Semantics is good at saying illuminating things about the **functional** aspects of meaning—the connective tissue that allows meanings to compose with each other.

But it doesn't tell us a lot about the **conceptual** parts of meanings, which tend to just get disquoted in our lexical semantic values.

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

The **functional** aspects of meaning behave like module-bound representations:

proprietary vocabulary

centrally inaccessible
(ask a lay person to explain the meaning of "every")

informationally encapsulated

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

By contrast, the **conceptual** aspects of meaning *don't* behave like module-bound representations:

Shared Vocabulary with CG

Centrally Accessible: We can articulate these aspects of word meanings, to some extent

Unencapsulated: We can learn new meanings on the fly and temporarily change how we use words (in their conceptual aspects)

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

Also notable:

The really funky kinds of polysemy tend to hang out in the **conceptual** parts of meaning.

Disclaimer: this isn't true of everything people call "polysemy"
(e.g. semantic-type flexibility), but those parts aren't funky!

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

"...there are two types of meaning at work. One is structural-functional. This is genuinely part of the language faculty, and so is within the domain of linguistic theory. It is generally abstract, and often yields to mathematical description. The other is core conceptual meaning, which tells us just which concepts (or properties, or whatever you prefer) our words express. These enter linguistic theory and the language faculty only through pointers, and are marked by the places our semantic theories fall back on disquotation."

Functional vs. Conceptual Meaning

Michael Glanzberg (2014): "Explanation and Partiality in Semantics"

$$[\![\text{tall}]\!] = \lambda w.\ \lambda x.\ \lambda d.\ \text{tall}(x,d) \text{ in } w$$

$$[\![\text{buy}]\!] = \lambda w.\ \lambda y.\ \lambda x.\ \lambda x.\ \text{buy}(e,x,y) \text{ in } w$$

$$[\![\text{every}]\!] = \lambda w.\ \lambda Y.\ \lambda X.\ X \subseteq Y$$

"I propose that what are in the lexicon corresponding to disquotation in our theories are simply pointers to other conceptual systems. If you like, they are pointers to *concepts* which are indicated by the non-quoted side of a disquotation clause."

Fourth Take:

Polysemy = lexical item with abstract or structured meanings

$$\mu(\text{green}) = \lambda\phi_{et} \,.\, \phi = \Sigma\phi$$

Hmm, so maybe this is a pointer to a concept, which is complex, structured, theoretical database entry out in central cognition?

Fourth Take:

Polysemy = lexical item with abstract or structured meanings

⟦green⟧ = λw. λx. x is green in w

Or maybe *this* is was the pointer all along.

Maybe the metalanguage itself is polysemous!

Concepts as Pointers

Jake Quilty-Dunn (2021):
"Polysemy and thought: Toward a generative theory of concepts"

- A lot of these issues about polysemy recur in the debate about concepts themselves!
- There are good reasons to think of concepts as atomistic representations (thought compositionality).
- There are also good reasons to think of concepts as richly structured databases of information (our use of concepts to categorize things).
- Let's have our cake and eat it too! Concepts are atomistic pointers to complex databases of information about their referents.

Cf. Pietroski?

Word meanings are concepts.

Word meanings systematically underdetermine concepts, and so compositionality is impossible!

Word meanings underdetermine concepts, but compositionality is possible with this one neat trick.

Word meanings are concepts.

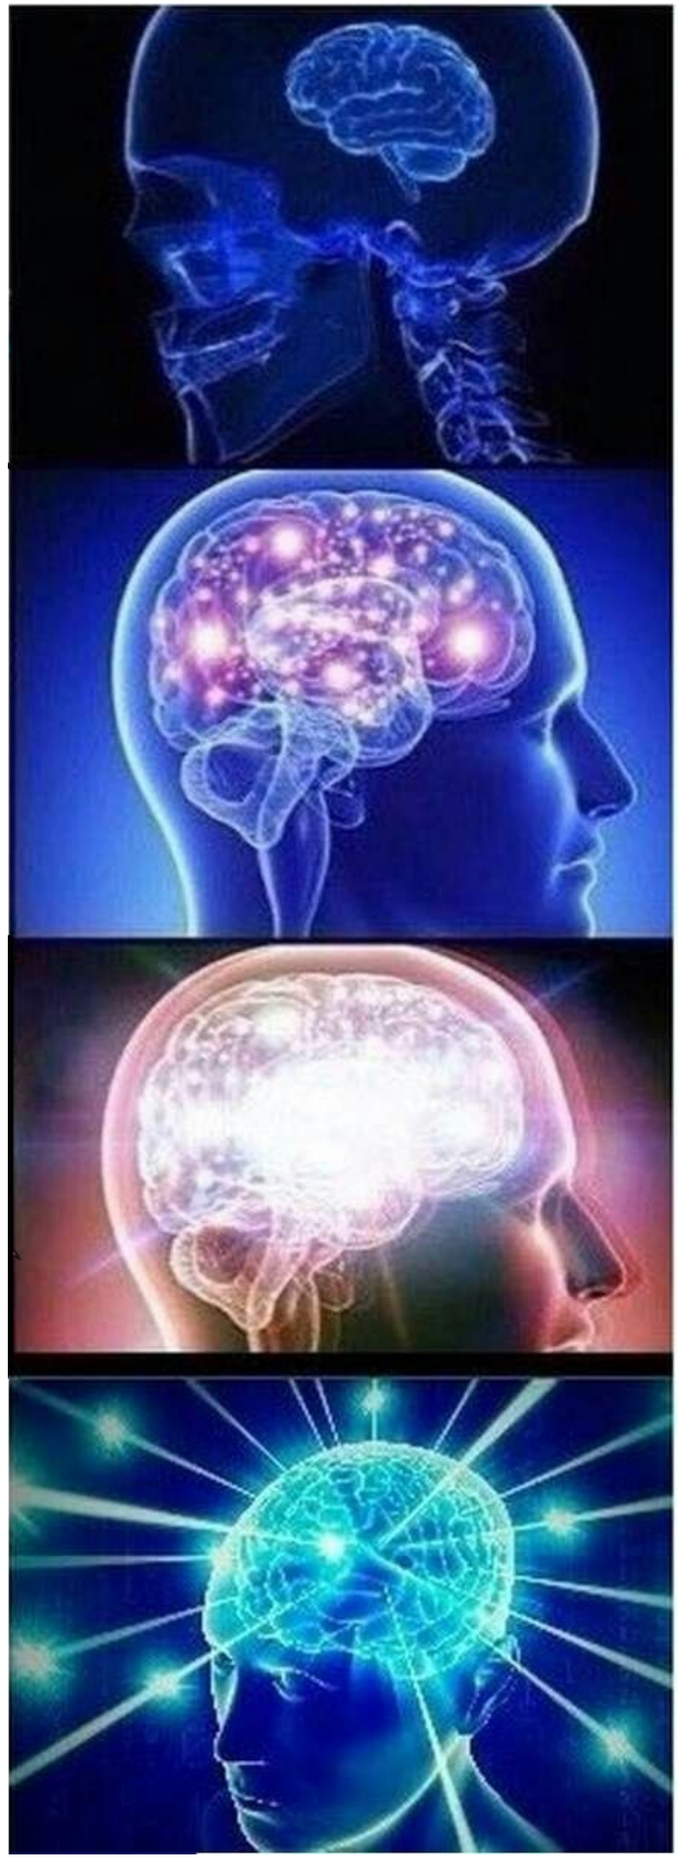

Conventionality of Polysemy

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

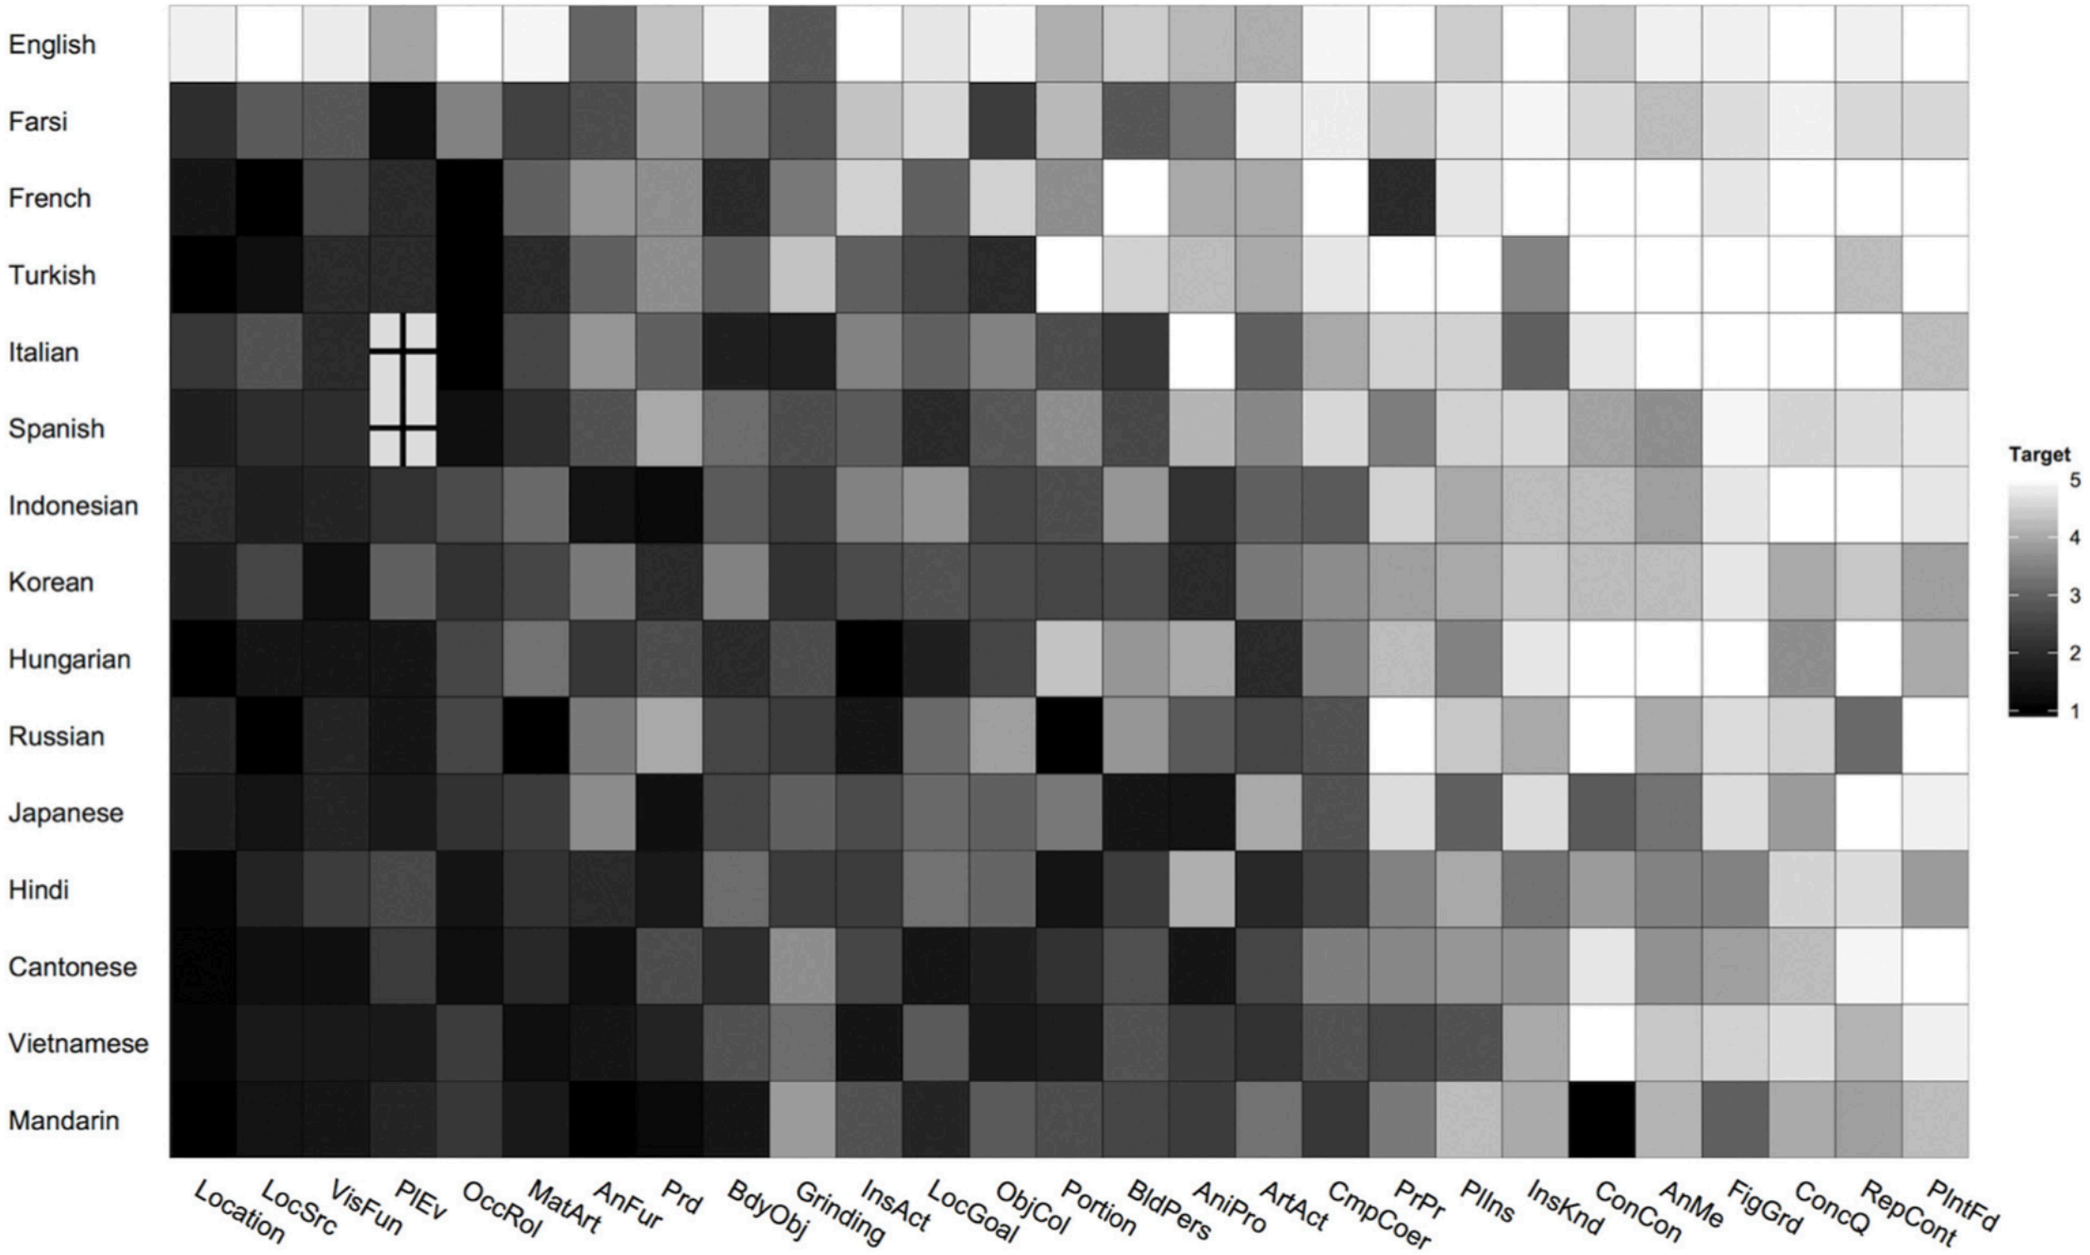

Fig. 3. The mean naturalness ratings of senses following each pattern, by language. Average rating is color-coded (white = most natural, black = least natural). White squares with crosses indicate missing data. The heatmap is ordered such that the languages with the highest average naturalness of senses (across patterns) are at the top, and patterns with the highest average naturalness of senses (across languages) are on the right.

Conventionality of Polysemy

Mahesh Srinivasan and Hugh Rabagliati (2015): "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy"

"...if polysemy is a direct reflection of conceptual structure, the same patterns and senses should be present across languages, but if polysemy corresponds to arbitrary lexicalized conventions, patterns and senses should be highly variable across languages. Our findings suggest that the structure of polysemy cannot be explained by either concepts or conventions on their own. Specifically, across 15 languages and 26 patterns of polysemy, we found very few instances where a language showed no evidence of having a particular pattern of polysemy (like the use of material words to label artifacts), which provides evidence against a conventions-only model. However, contrary to a concepts-only model, we found that many patterns are instantiated by different sets of senses across languages (e.g., glass labeling a drinking vessel, a mirror, etc.)."

(Extremely Tentative) Conclusions

- Polysemy is a head scratcher.
- We have some interesting options for how to plug it into compositional-semantic theories.
- Maybe the best option is not to change our theories at all, and let
- In some promising ways, facts about polysemy pattern with the boundary between the semantic module and central cognition.
    - E.g., it shows up as one of the aspects of meaning that doesn't behave in modular ways.
- But on the other hand, cross-linguistic variation suggests that polysemy is not an entirely extralinguistic

Thanks for today!

NASSLLI 2022 · USC · 20–24 JUNE
SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

DAY 4:
LANGUAGE IN THOUGHT AND AUDIENCE DESIGN

Daniel W. Harris
CUNY Graduate Center and Hunter College

Problem 1:

Non-Communicative Language Use


I'm good enough. I'm smart enough.

I guess I was wrong when I said I never promised anyone.

Inner Speech

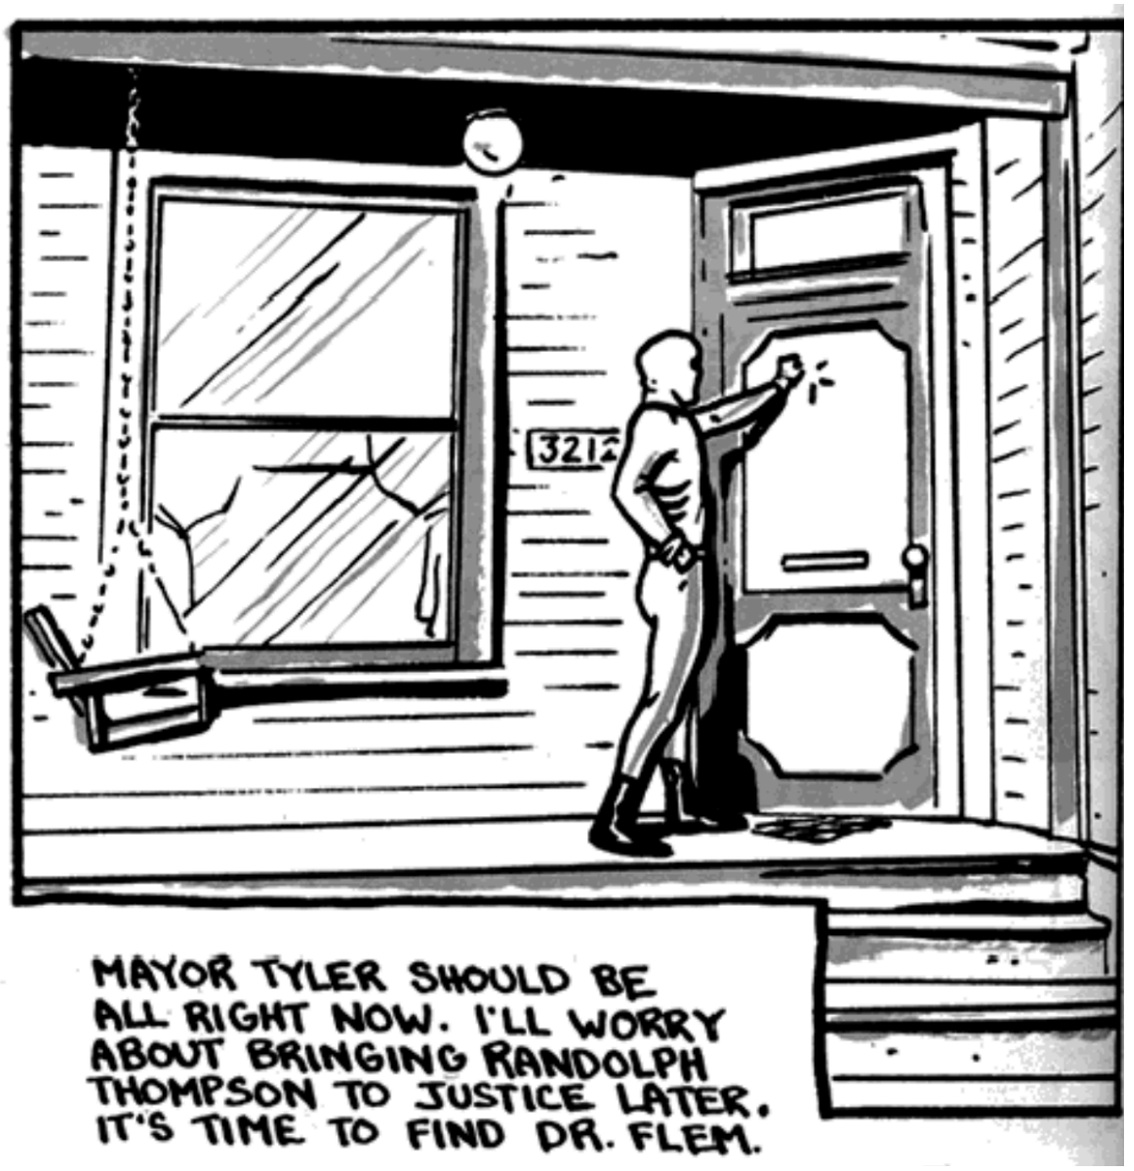

Inner Speech and Action Control

Luria, A. R. (1959). "The directive function of speech in development and dissolution," Part I. *Word*, 15, 341-352.

"What the child at first does with the help, and on the instructions, of the adult, he later begins to do by himself, supporting himself with his own speech; that speech as a form of communication with adults later becomes a means of organising the child's own behaviour and that function which was previously divided between two people later becomes an internal function of human behaviour"(341)

Inner Speech and Action Control

A Baddeley, D Chincotta, A Adlam (2001): "Working memory and the control of action: evidence from task switching"

- Gave subjects basic cognitive tasks to perform—e.g. alternating between simple additions (8+1) and simple subtractions (7–1).
- In the experimental condition, subjects were asked to repeat the days of the week (Monday, Tuesday,...) or the months of the year (January, February,...).
- This disrupts verbal working memory, and so (apparently) inner speech.
- Subjects do worse on switching tasks (among various other tasks) in this condition.
- Conclusion: Inner speech plays a role in the control of at least some actions.

Russian blues reveal effects of language on color discrimination

Jonathan Winawer*†‡, Nathan Witthoft*‡, Michael C. Frank*, Lisa Wu§, Alex R. Wade¶, and Lera Boroditsky‡

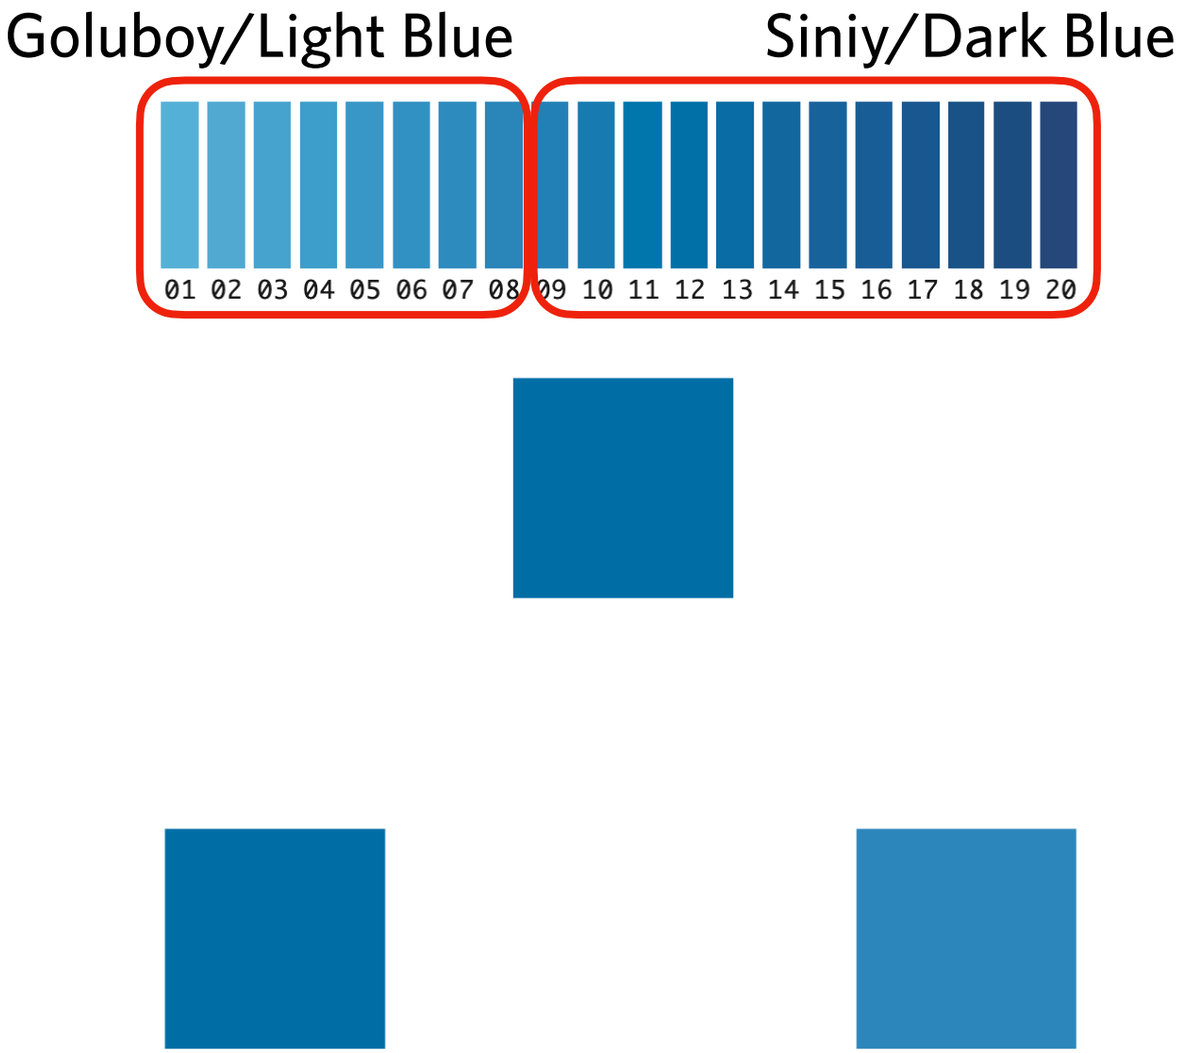

Fig. 1. The 20 blue colors used in this study are shown at the top of the figure. An example triad of color squares used in this study is shown at the bottom of the figure. Subjects were instructed to pick which one of the two bottom squares matched the color of the top square.

- Task: Say which two patches match.
- Russian speakers were faster when the colors were located near the light/dark boundary.
- The effect disappears when subjects' verbal working memories are occupied.
- So, a hypothesis: the difference is that two words take up more space in verbal working memory than one, or takes longer to read and write.

Problem 2:

Top-Down Control and Audience Design

Problem 2:
Top-Down Control and Audience Design

- It seems like we can make top-down, deliberate decisions about what to say and how to say it.
- These decisions seem at least sometimes to be informed by all sorts of beliefs, intentions, and mindreading.

Dad, what's that book?

This is one of my philosophy books.
It's by a philosopher named Ludwig Wittgenstein.
It's called *The Philosophical Investigations*.

Wittgen…stein?

Yes.

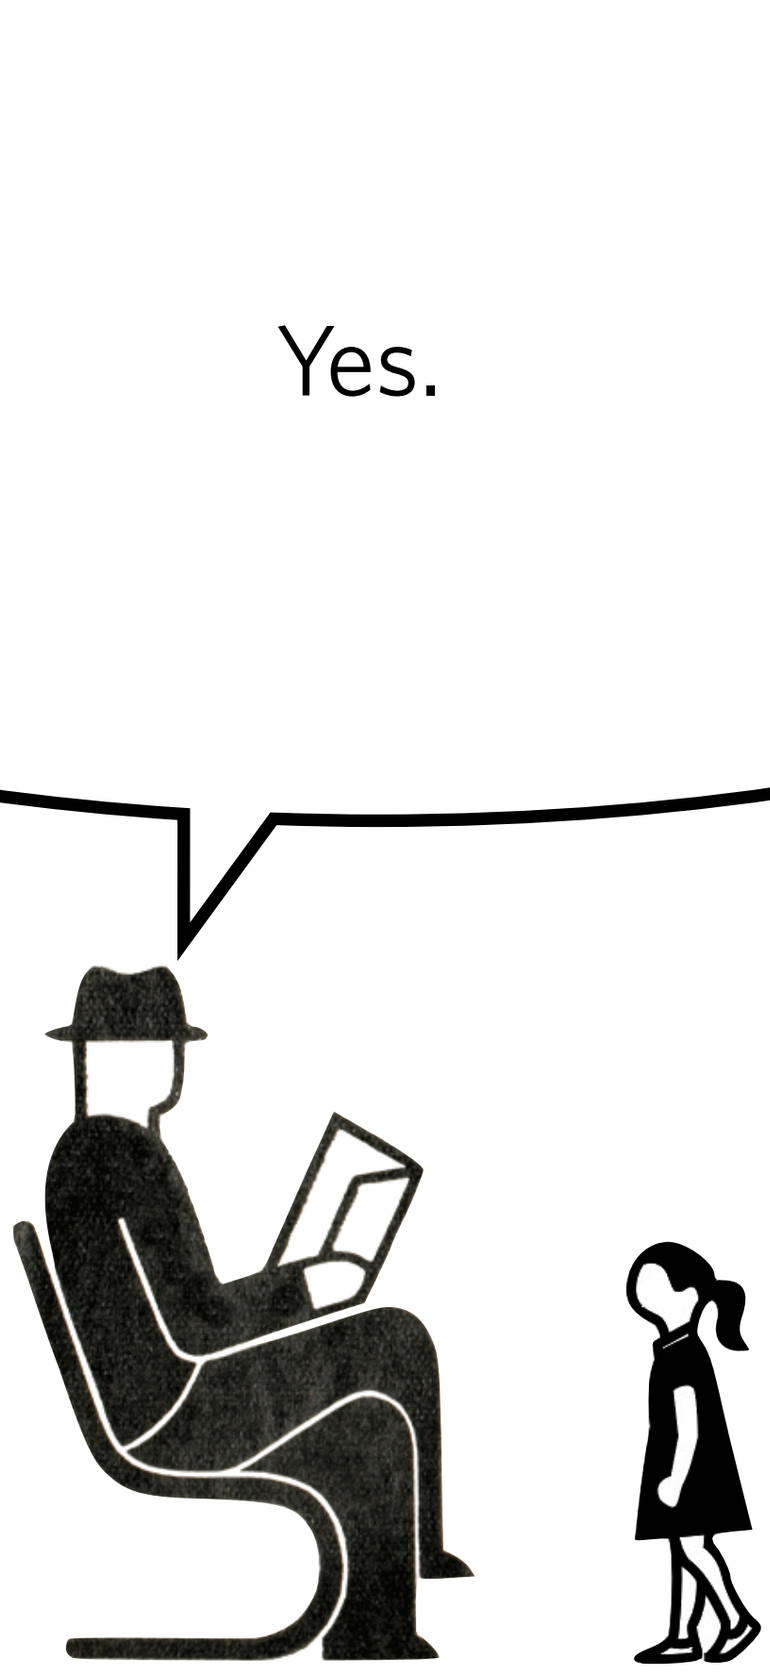

Do you want to know what it says?

Yes.

Well, here's one thing that it says:
In order to know what a rule tells us to do, we need help from other people.

Other people? Like our teachers?

Yes, or our friends, or our family. If they don't help us, we won't know what the rule means.

MESSAGE DESIGN

Oona doesn't know much about philosophy, and hasn't heard of Wittgenstein. So I will start with some very general information about the book to introduce her to a new topic.

This is one of my philosophy books. It's by a philosopher named Ludwig Wittgenstein. It's called *The Philosophical Investigations*.

The *Investigations*.

My colleague has read this book before,
and so all I need to do is to increment
some information she already has.

This is one of
my philosophy books.
It's by a philosopher
named Ludwig Wittgenstein.
It's called *The Philosophical Investigations*.

The *Investigations*.

Well, here's one thing that it says:
In order to know what a rule tells us to do,
we need help from other people.

SIGNAL DESIGN

I think the view is that following a
rule is an essentially social
practice.

Oona doesn't know what a "social practice" is, or what the word "essentially" means. So I will describe those ideas in simple terms.

Well, here's one thing that it says:
In order to know what a rule tells us to do, we need help from other people.

I think the view is that following a rule is an essentially social practice.

My colleague knows lots of philosophical terminology, and will be offended if I talk to her like she's a kid, so I will say "essentially social practice."

Well, here's one thing that it says: In order to know what a rule tells us to do, we need help from other people.

I think the view is that following a rule is an essentially social practice.

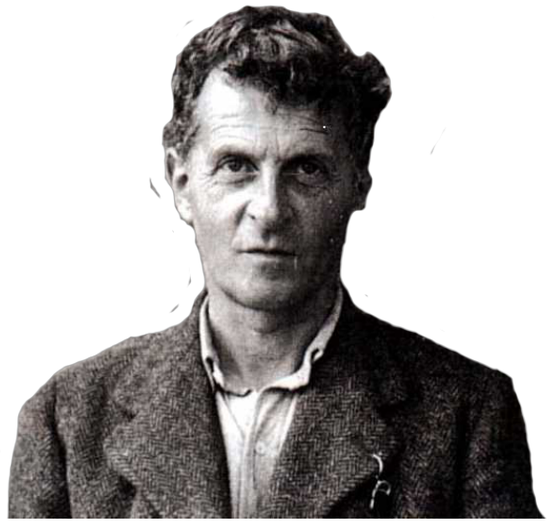

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

My addressee has never heard of Wittgenstein but knows what philosophers are.

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

My addressee has heard of *The Philosophical Investigations* but doesn't know (or doesn't remember) who wrote it.

A philosopher named 'Ludwig Wittgenstein'

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

My addressee remembers talking to me about a philosopher last week but doesn't know this is the same one.

A philosopher named 'Ludwig Wittgenstein'

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

My addressee knows who Ludwig Wittgenstein is, and by that name, but they might also know about some other Wittgensteins.

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

My addressee knows who Ludwig Wittgenstein is, and he is the most salient Wittgenstein for them.

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

LW is the most salient Ludwig with whom my addressee is on a first-name basis, and they also think that I am on a first-name basis with LW.

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

LW is currently very salient to my addressee, either because we have just been talking about him or for some other reason, and they can infer that I know this.

A philosopher named ‘Ludwig Wittgenstein’

The philosopher who wrote *The Philosophical Investigations*

The philosopher I was telling you about last week

Ludwig Wittgenstein

Wittgenstein

Ludwig

him

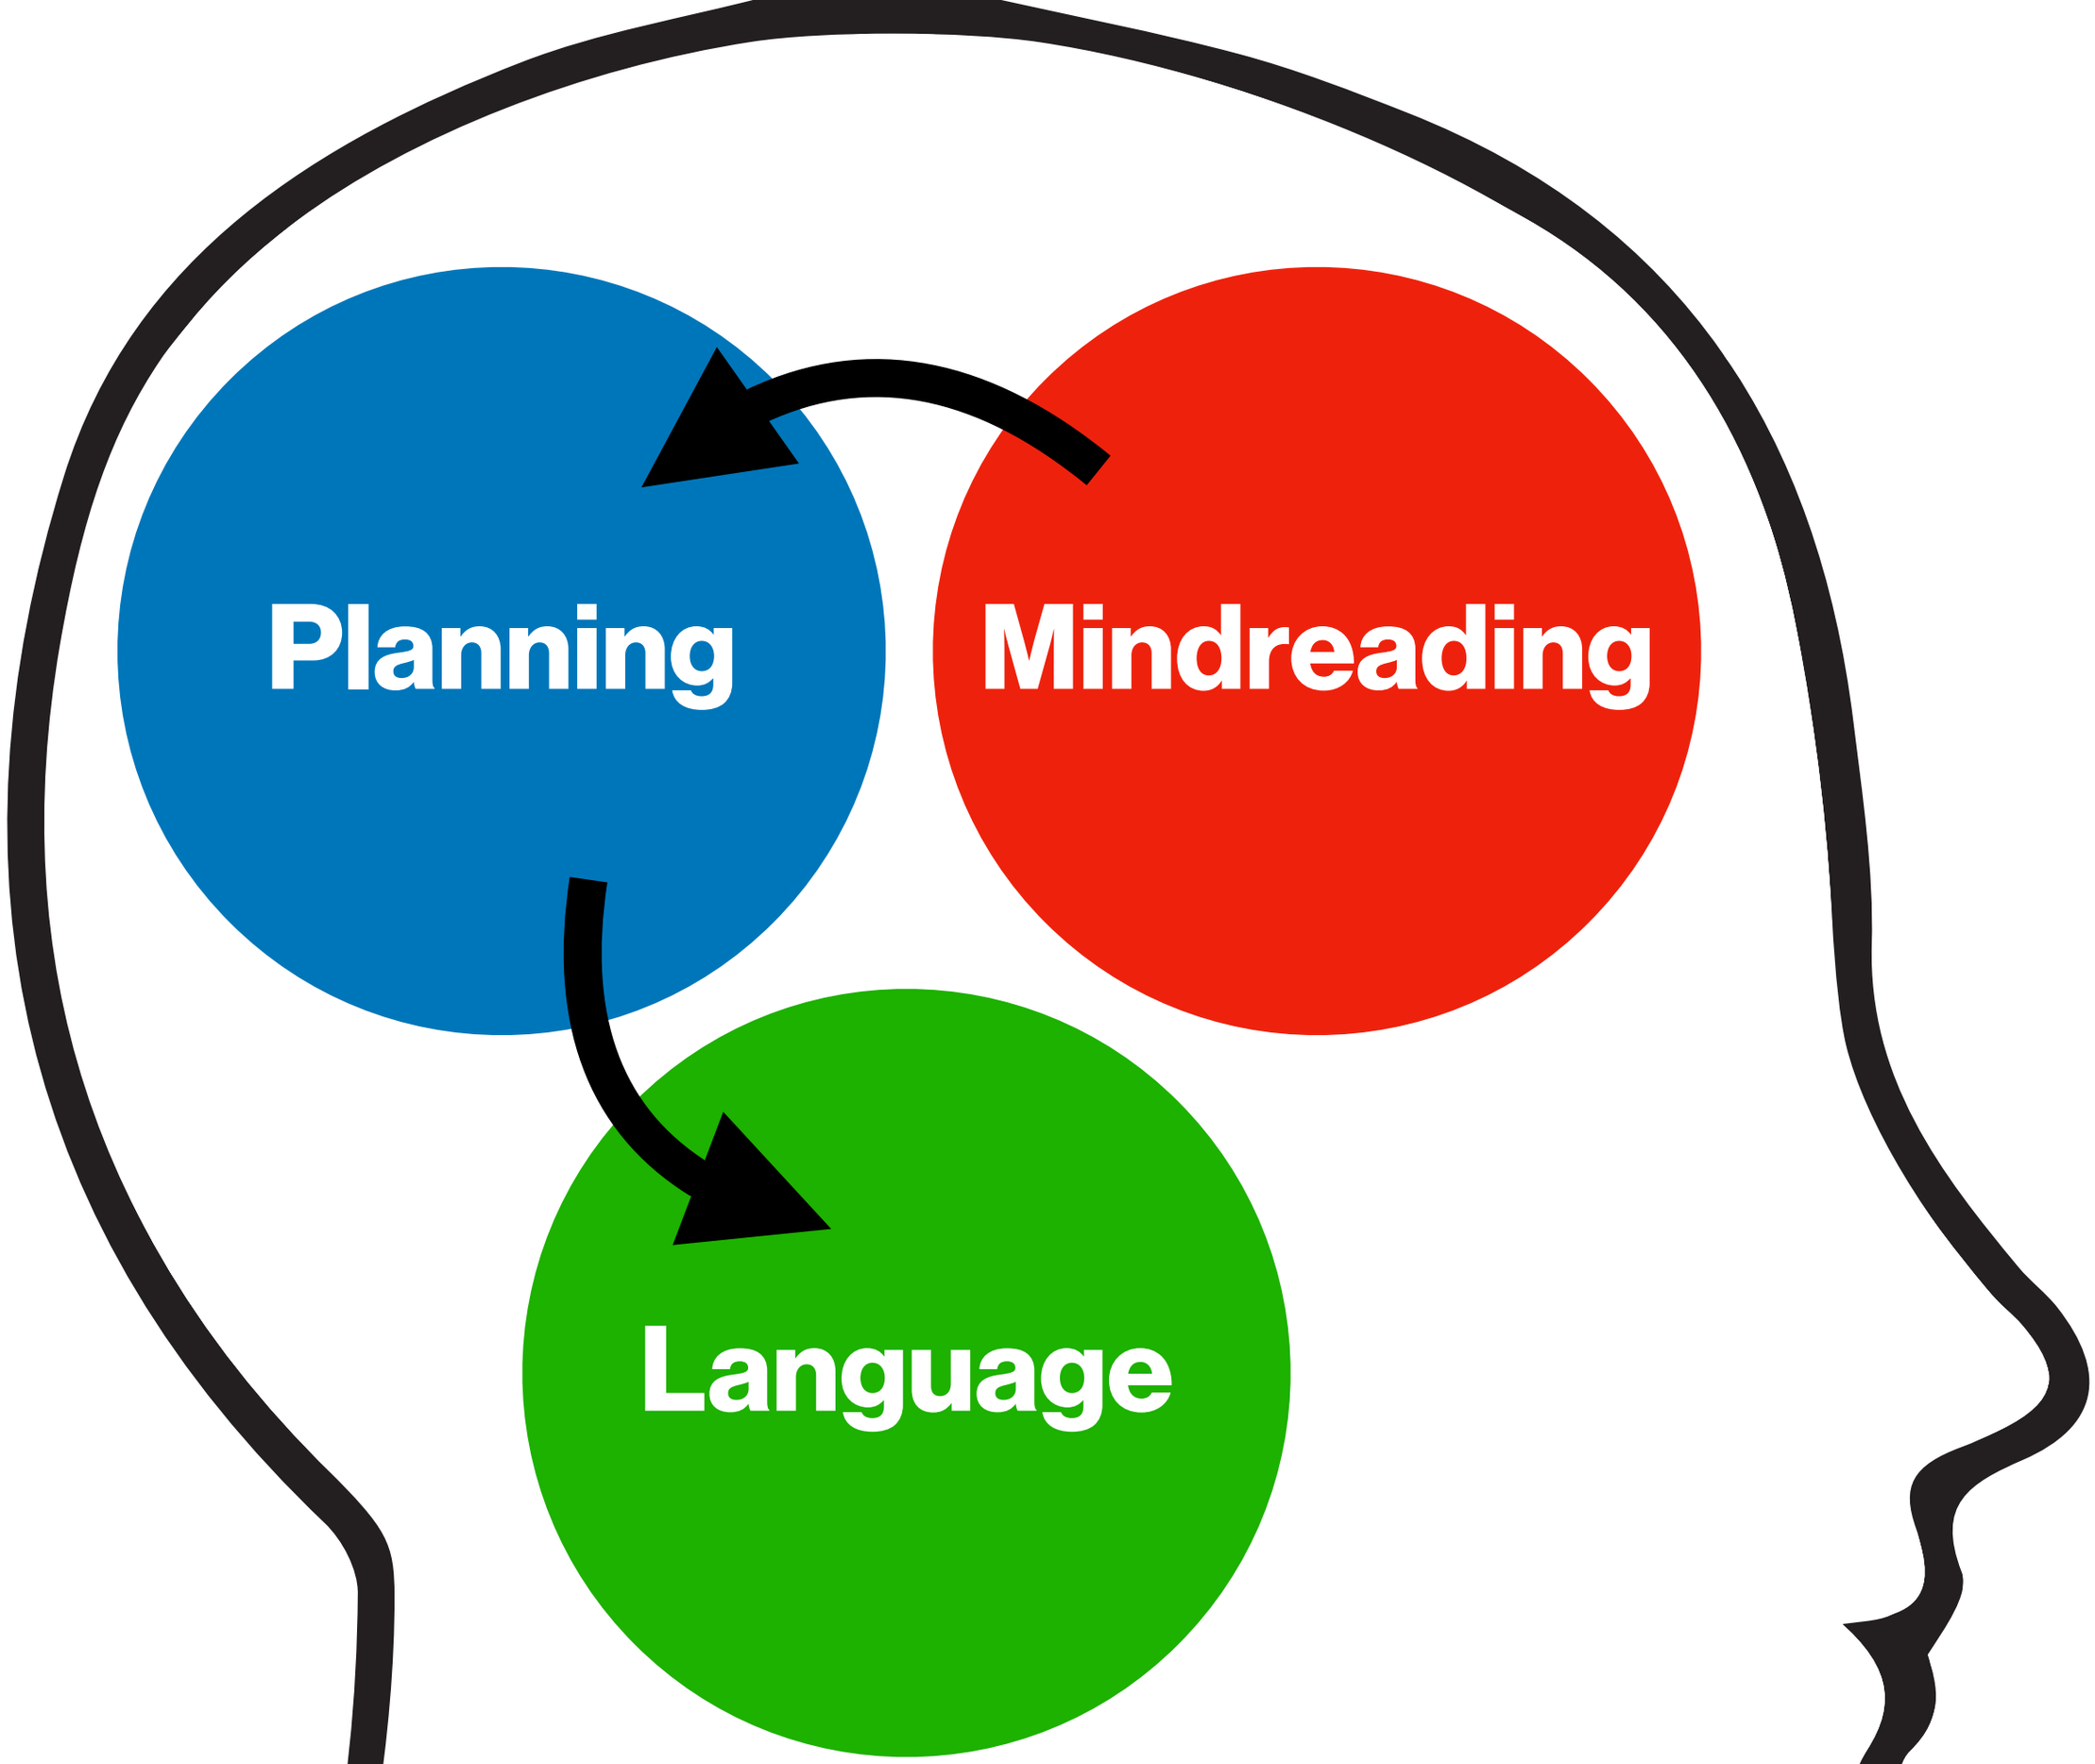
Planning
Mindreading
Language

Mindreading

Q: Do we really do all of this mindreading and practical reasoning in the course of a normal conversation?

A: Yes!

...sometimes!

...it's complicated!

The Director Task

Keysar, Barr, and Horton (1998): "The Egocentric Basis of Language Use: Insights From a Processing Approach,"

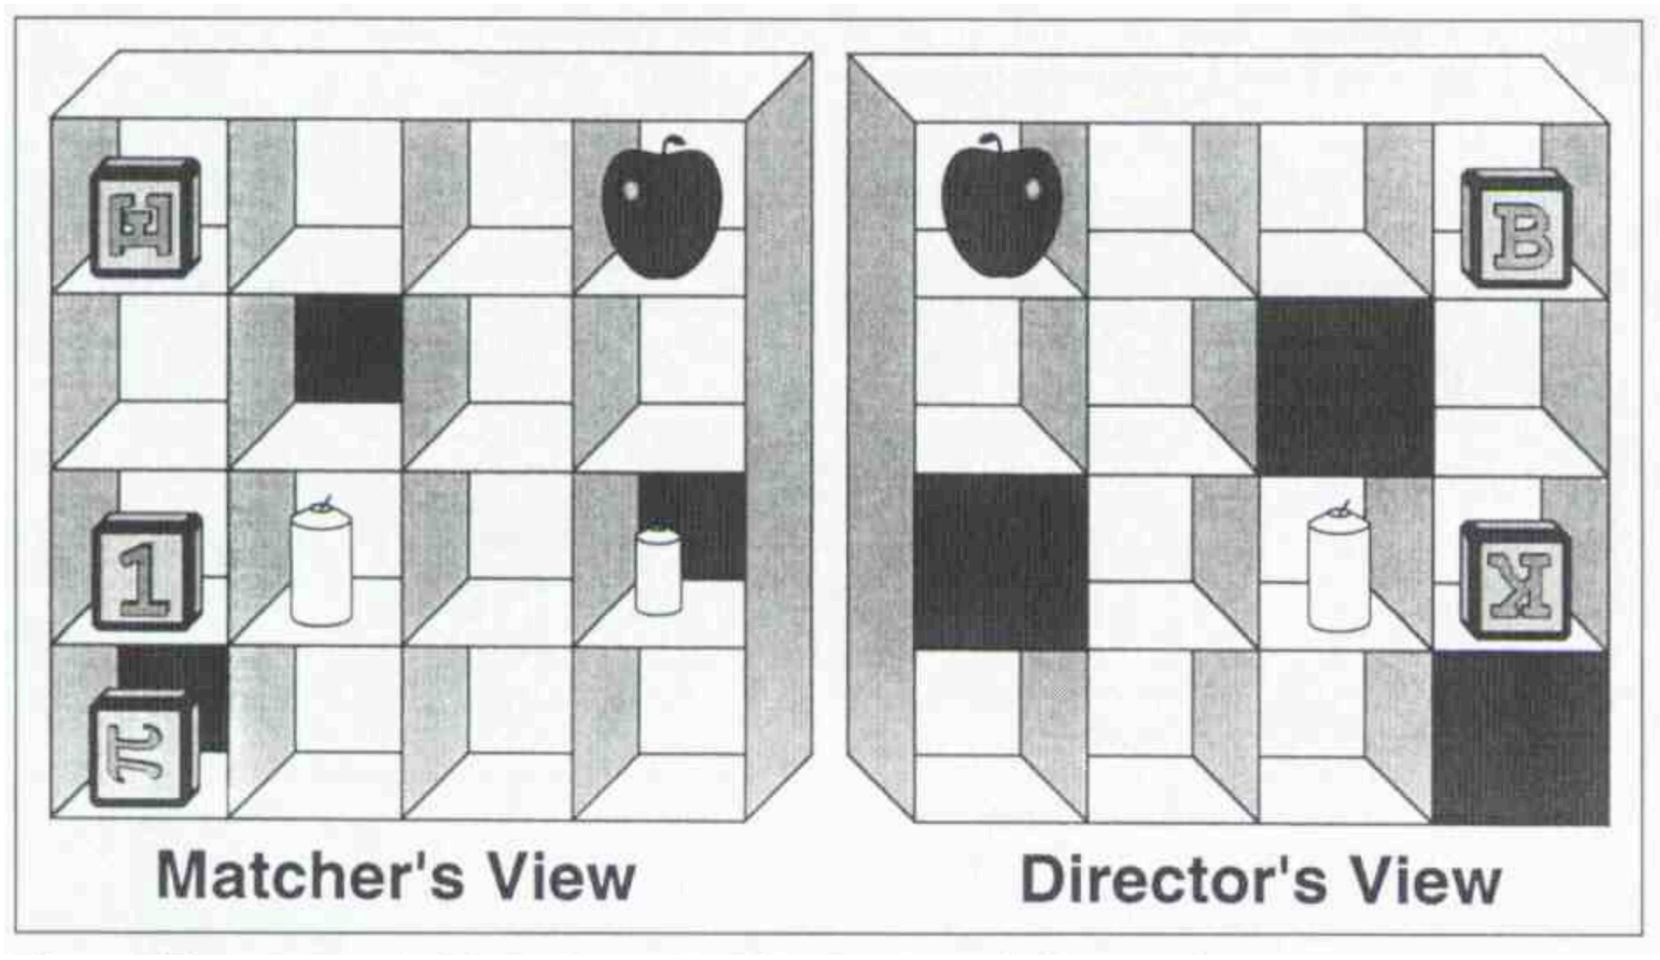

Director's instructions to Matcher:
"Put **the bottom block** below the apple."

If the Matcher moves the block marked ⊏, then they have reasoned "egocentrically"—i.e., failed to account for the Director's perspective.

The Anchor-and-Adjust Model

Speakers and hearers are often sensitive to others' perspectives.

But not always. Some patterns:

- cognitive load → more egocentric (Keysar 2008)
- Verbal-working-memory deficit → more egocentric (Lin et al 2010)
- Happier → more egocentric (Converse et al 2008)
- Younger children → more egocentric (Keysar 2008)
- Eye tracking studies: everyone is at least partly egocentric at first (Keysar et al 1998)

Theory:
We anchor to our own perspective, and adjust away if we have enough cognitive resources.
(Keysar 2007; Barr 2014; Epley et al 2004)

The Resource-Rational Model

- Eye-tracking studies: Subjects consider both their own and others' perspectives, even early in processing (Nadig & Sedivy 2002; Heller et al 2008, etc.)

- Speakers compensate for uncertainty about addressees' perspective by using more informative descriptions (Hawkins et al 2021)

- Subjects who encounter egocentric interlocutors repeatedly learn to invest more effort in later interactions (Hawkins et al 2021)

Theory:

We reason not only about others' states of mind, but also about how likely they are to be thinking about our states of mind, and about how much effort will be worth putting into this reasoning. (Hawkins et al 2021)

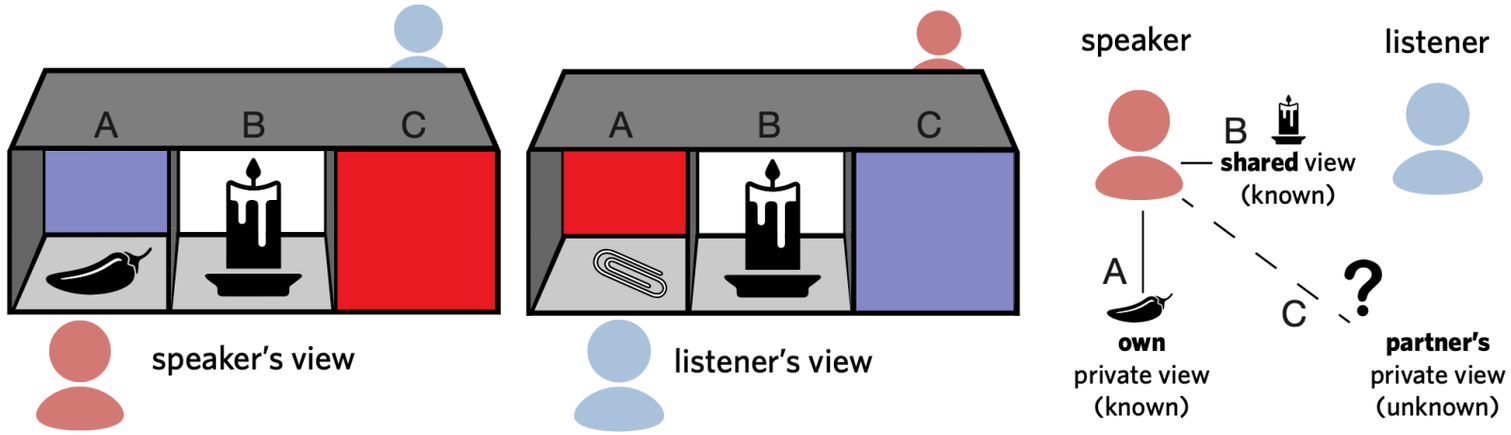

NASSLLI 2022 · USC · 20–24 JUNE
SEMANTICS AS THE STUDY OF A MODULAR SYSTEM

DAY 5:
VERBAL WORKING MEMORY

Daniel W. Harris
CUNY Graduate Center and Hunter College

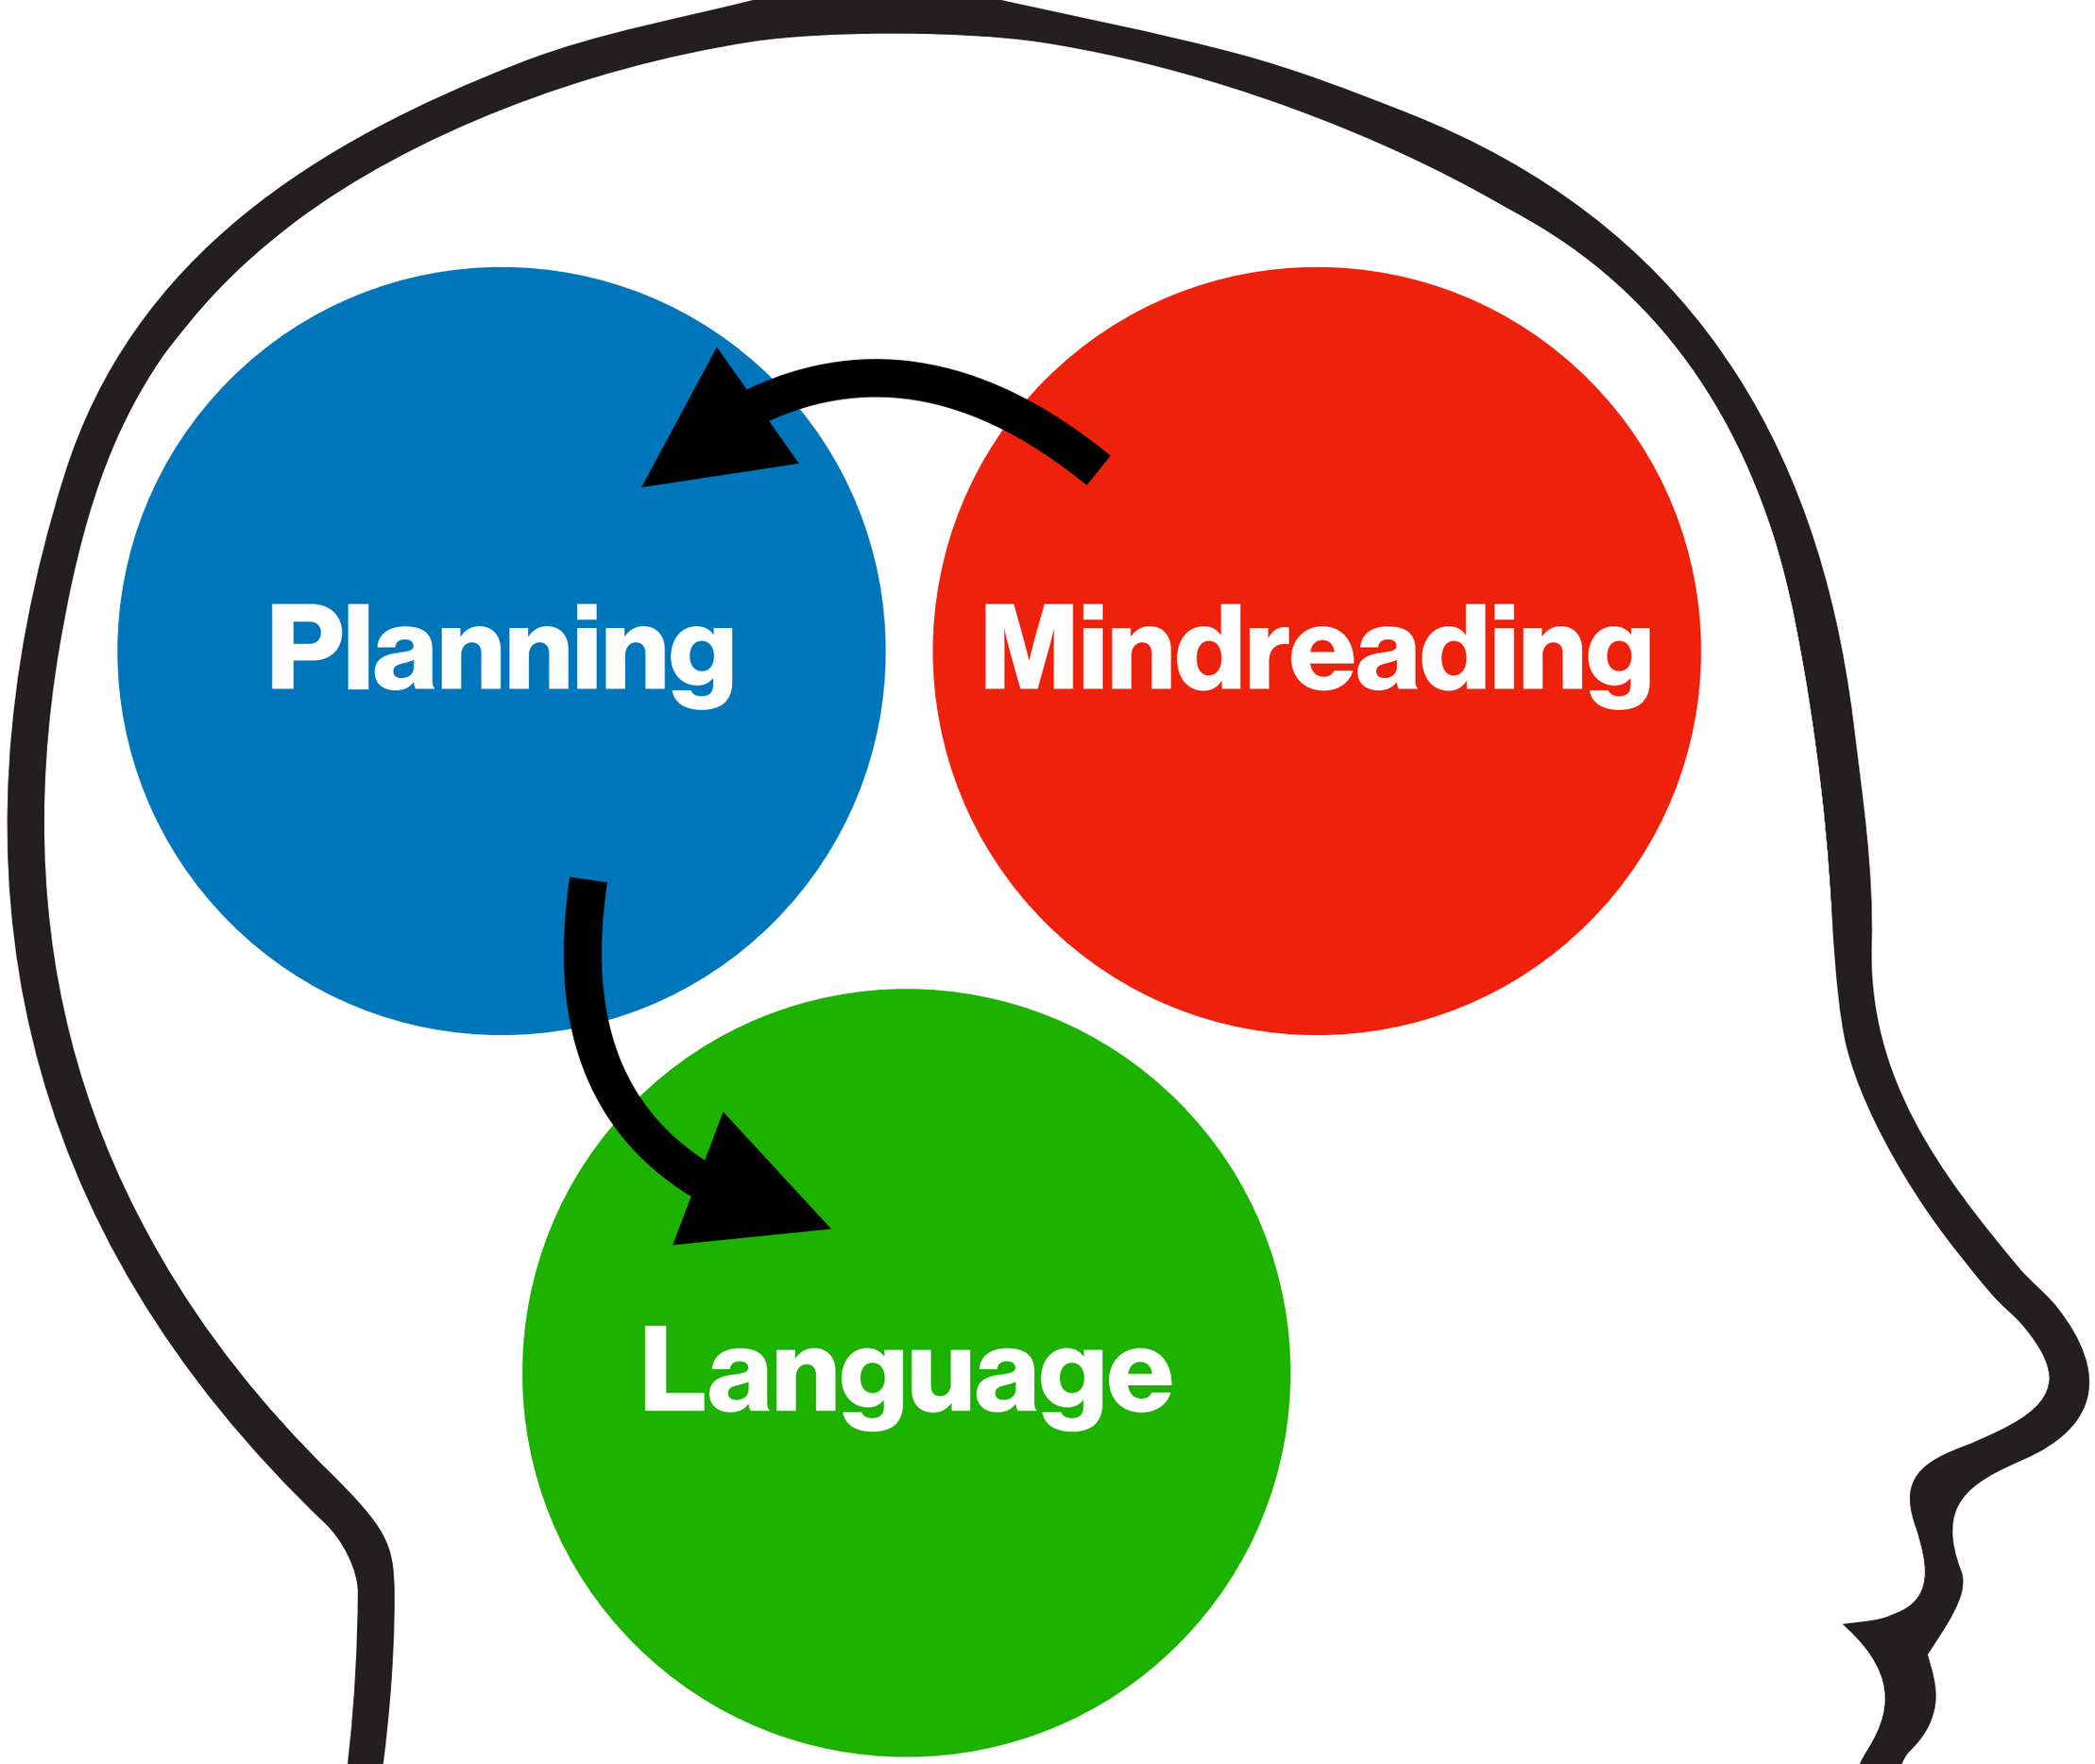
Planning
Mindreading
Language

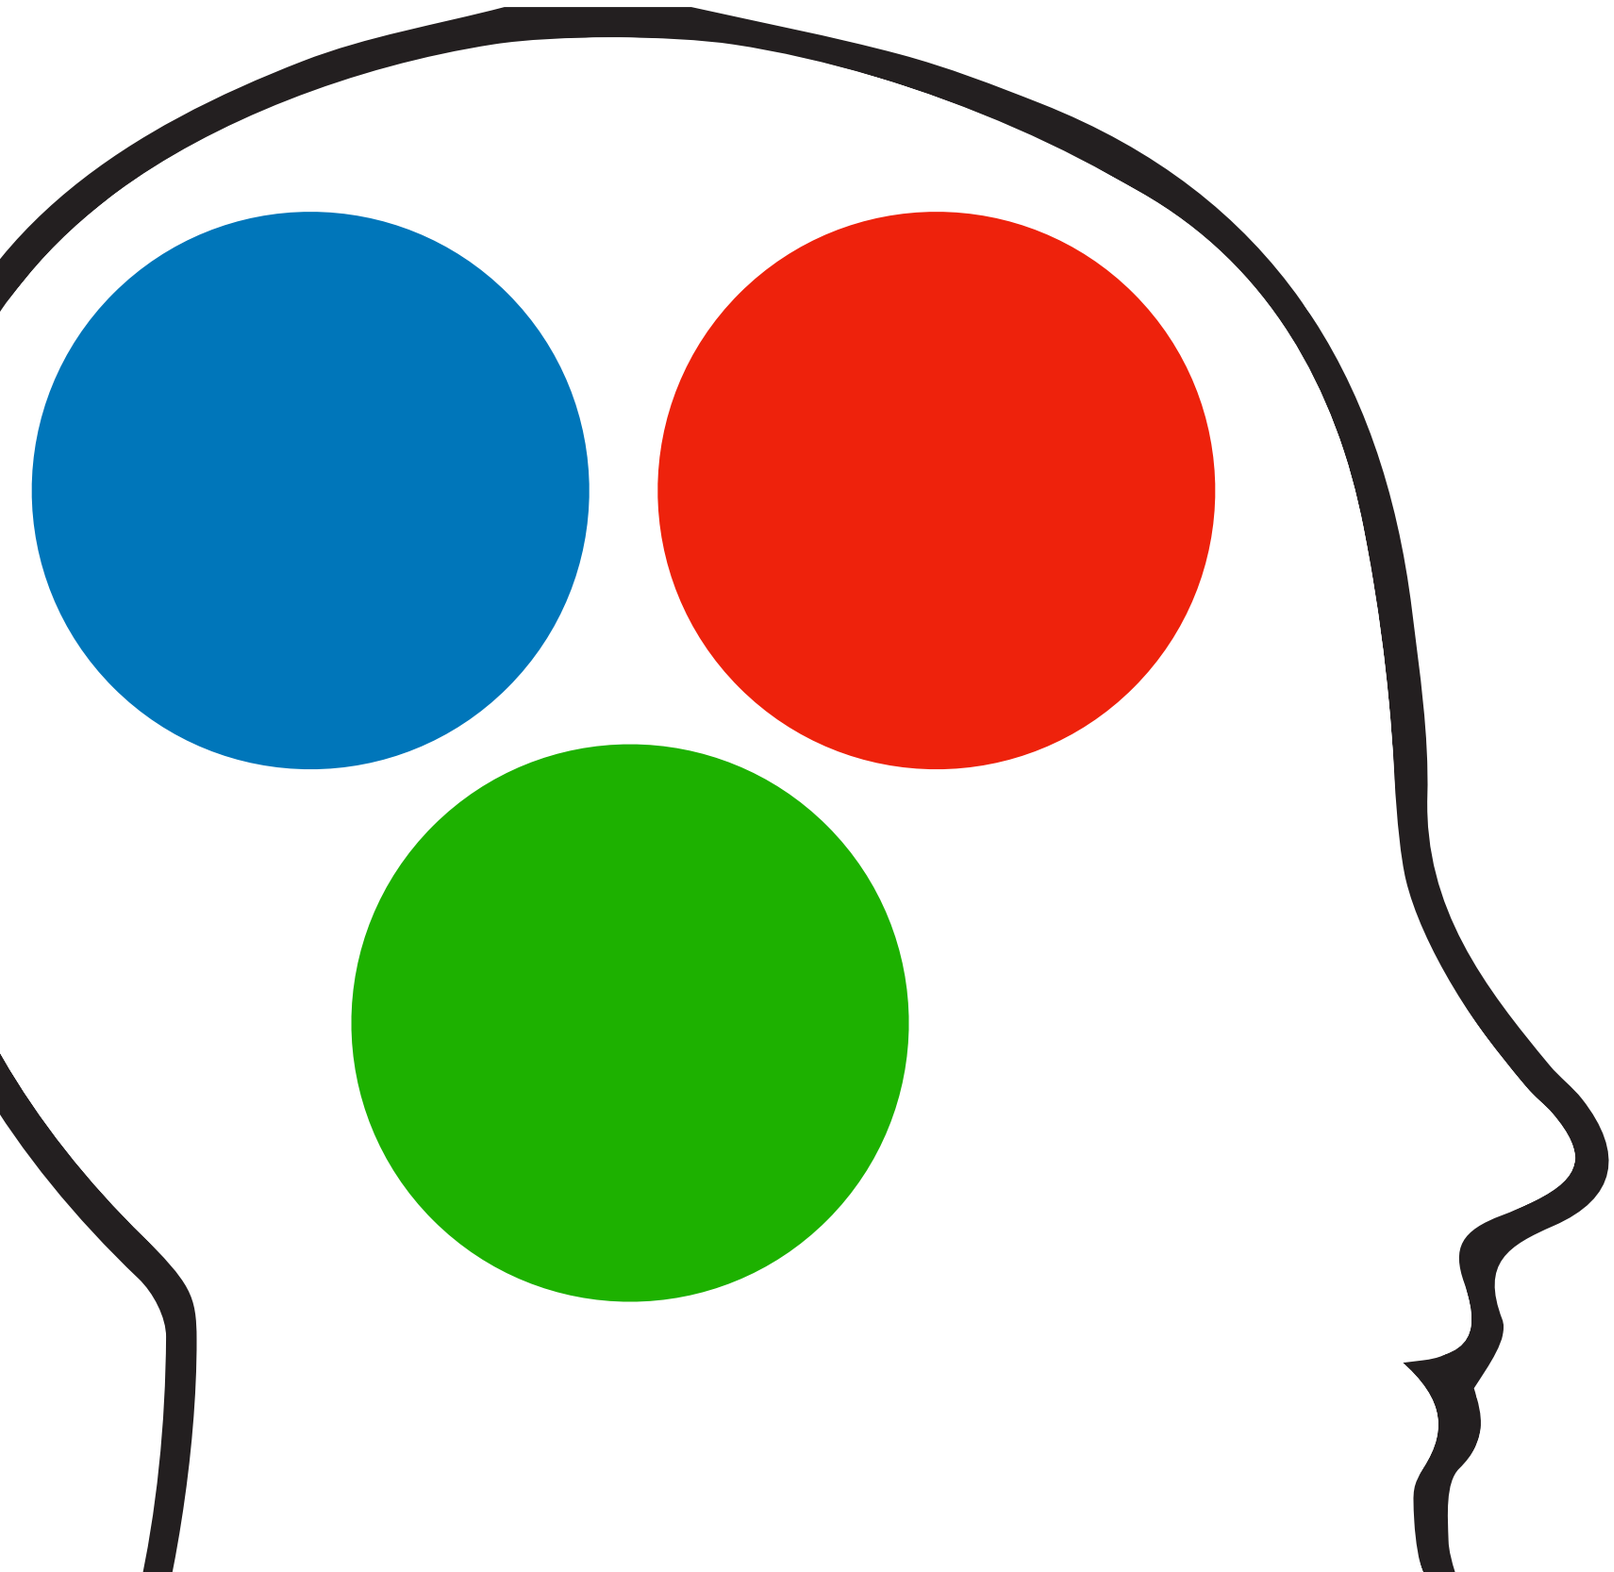

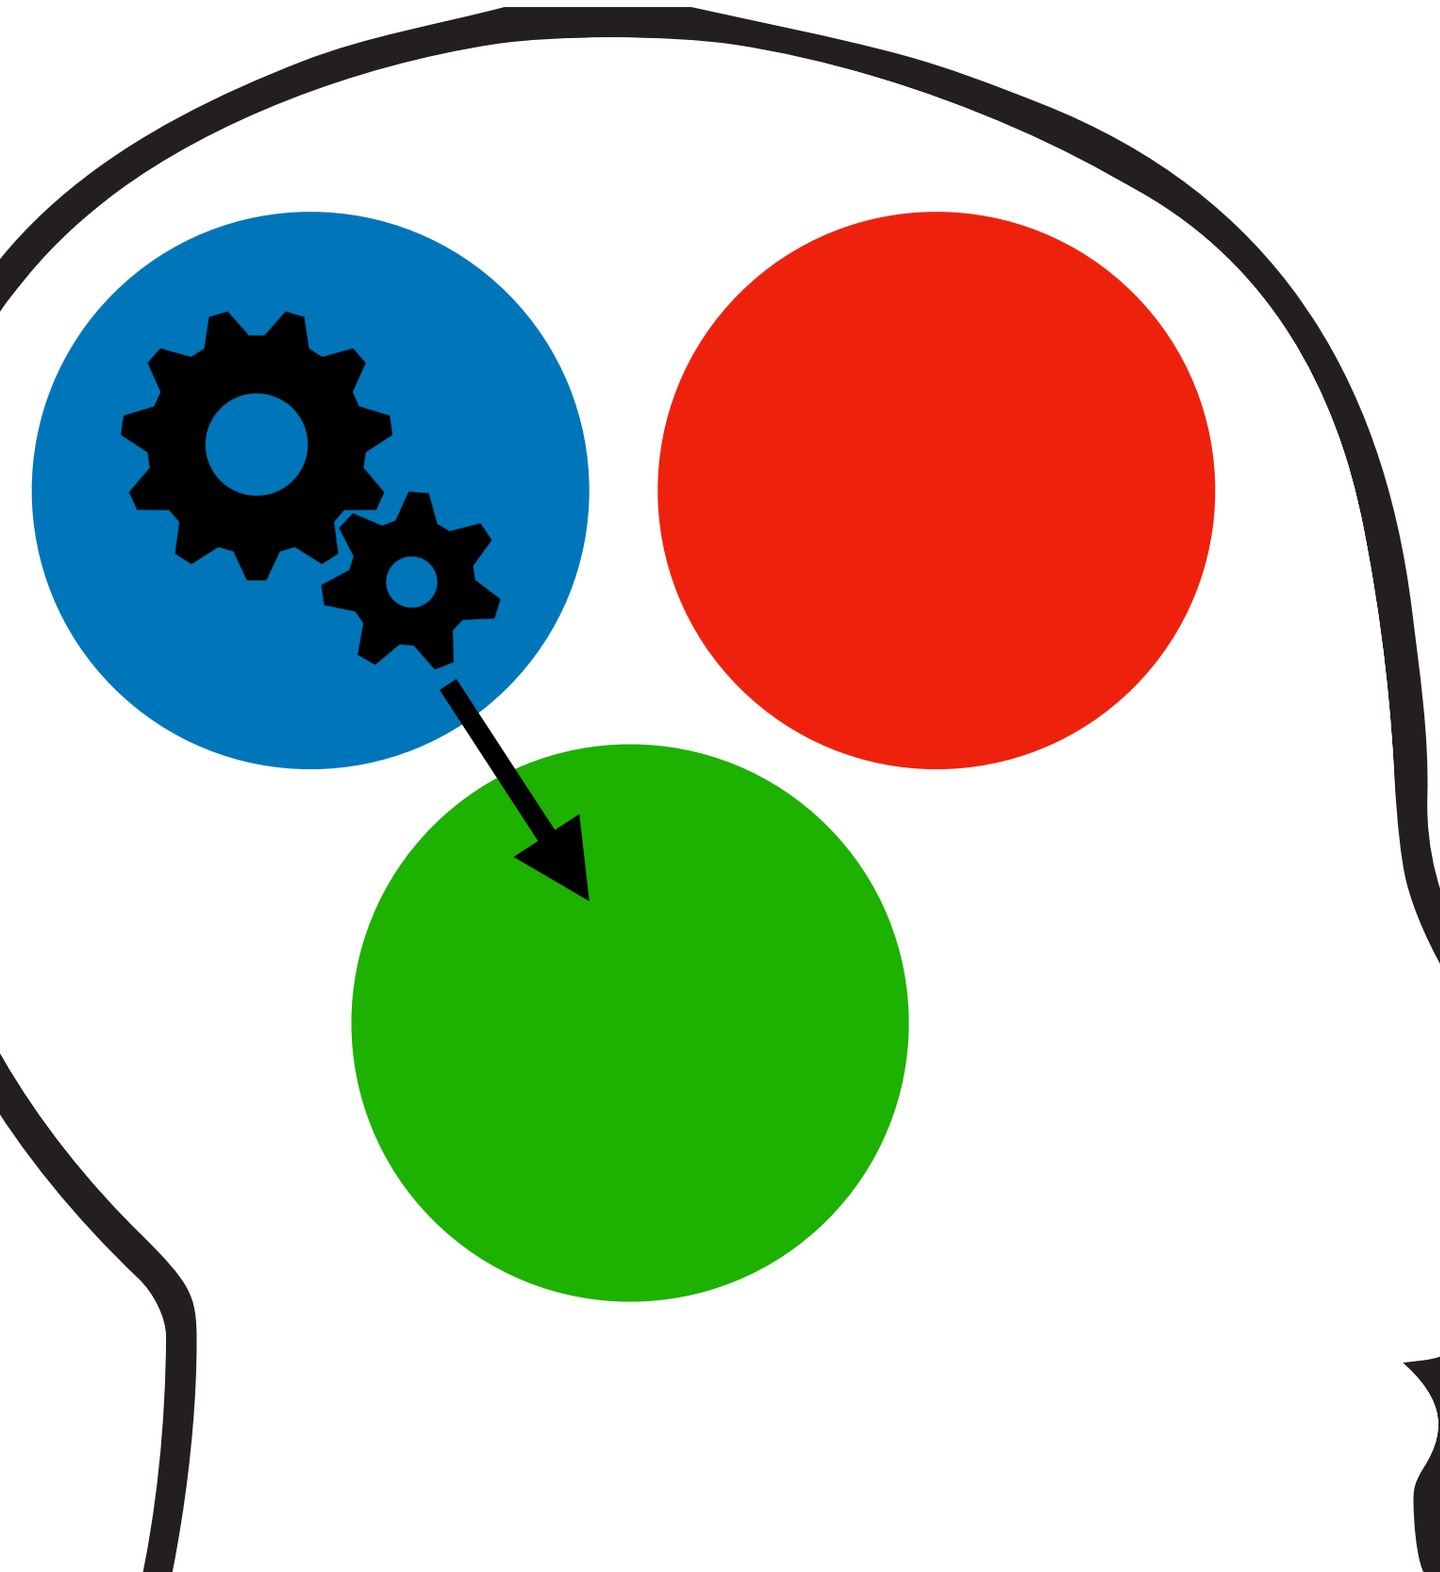

The planning system sends instruction to the language system.

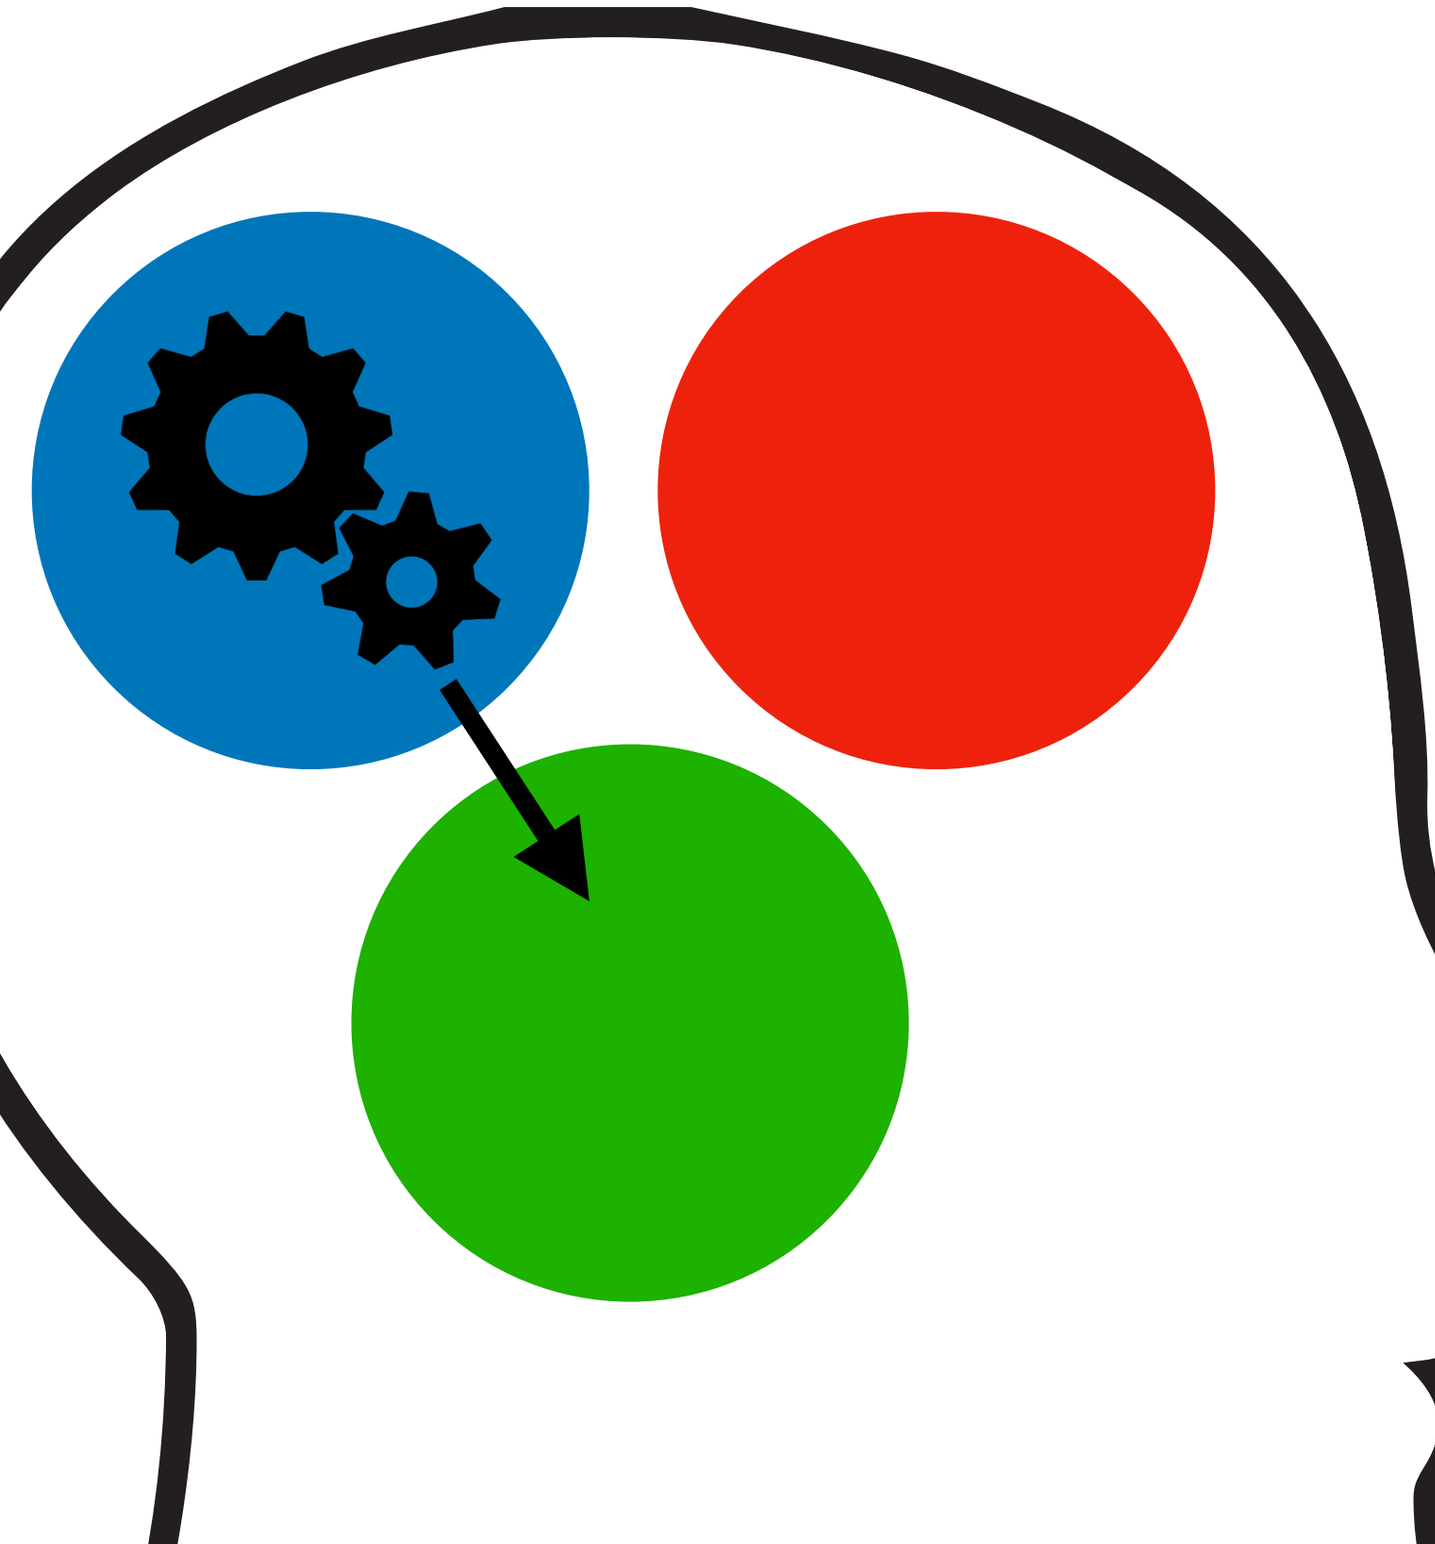

The planning system sends instruction to the language system.

These instructions are a subplan of S's communicative intention(s).

Planning

PRIOR INTENTION

Intention to communicate that Dan is talking.

SUBPLAN

Instruction for language system to encode

λp . (∃x : x is the speaker) p = λw . x is talking at w

Planning

PRIOR INTENTION

Intention to communicate that Dan is talking.

SUBPLAN

Instruction for language system to encode

$\lambda p . (\exists x : x \text{ is the speaker}) \; p = \lambda w . x \text{ is talking at } w$

Vocal Rehearsal

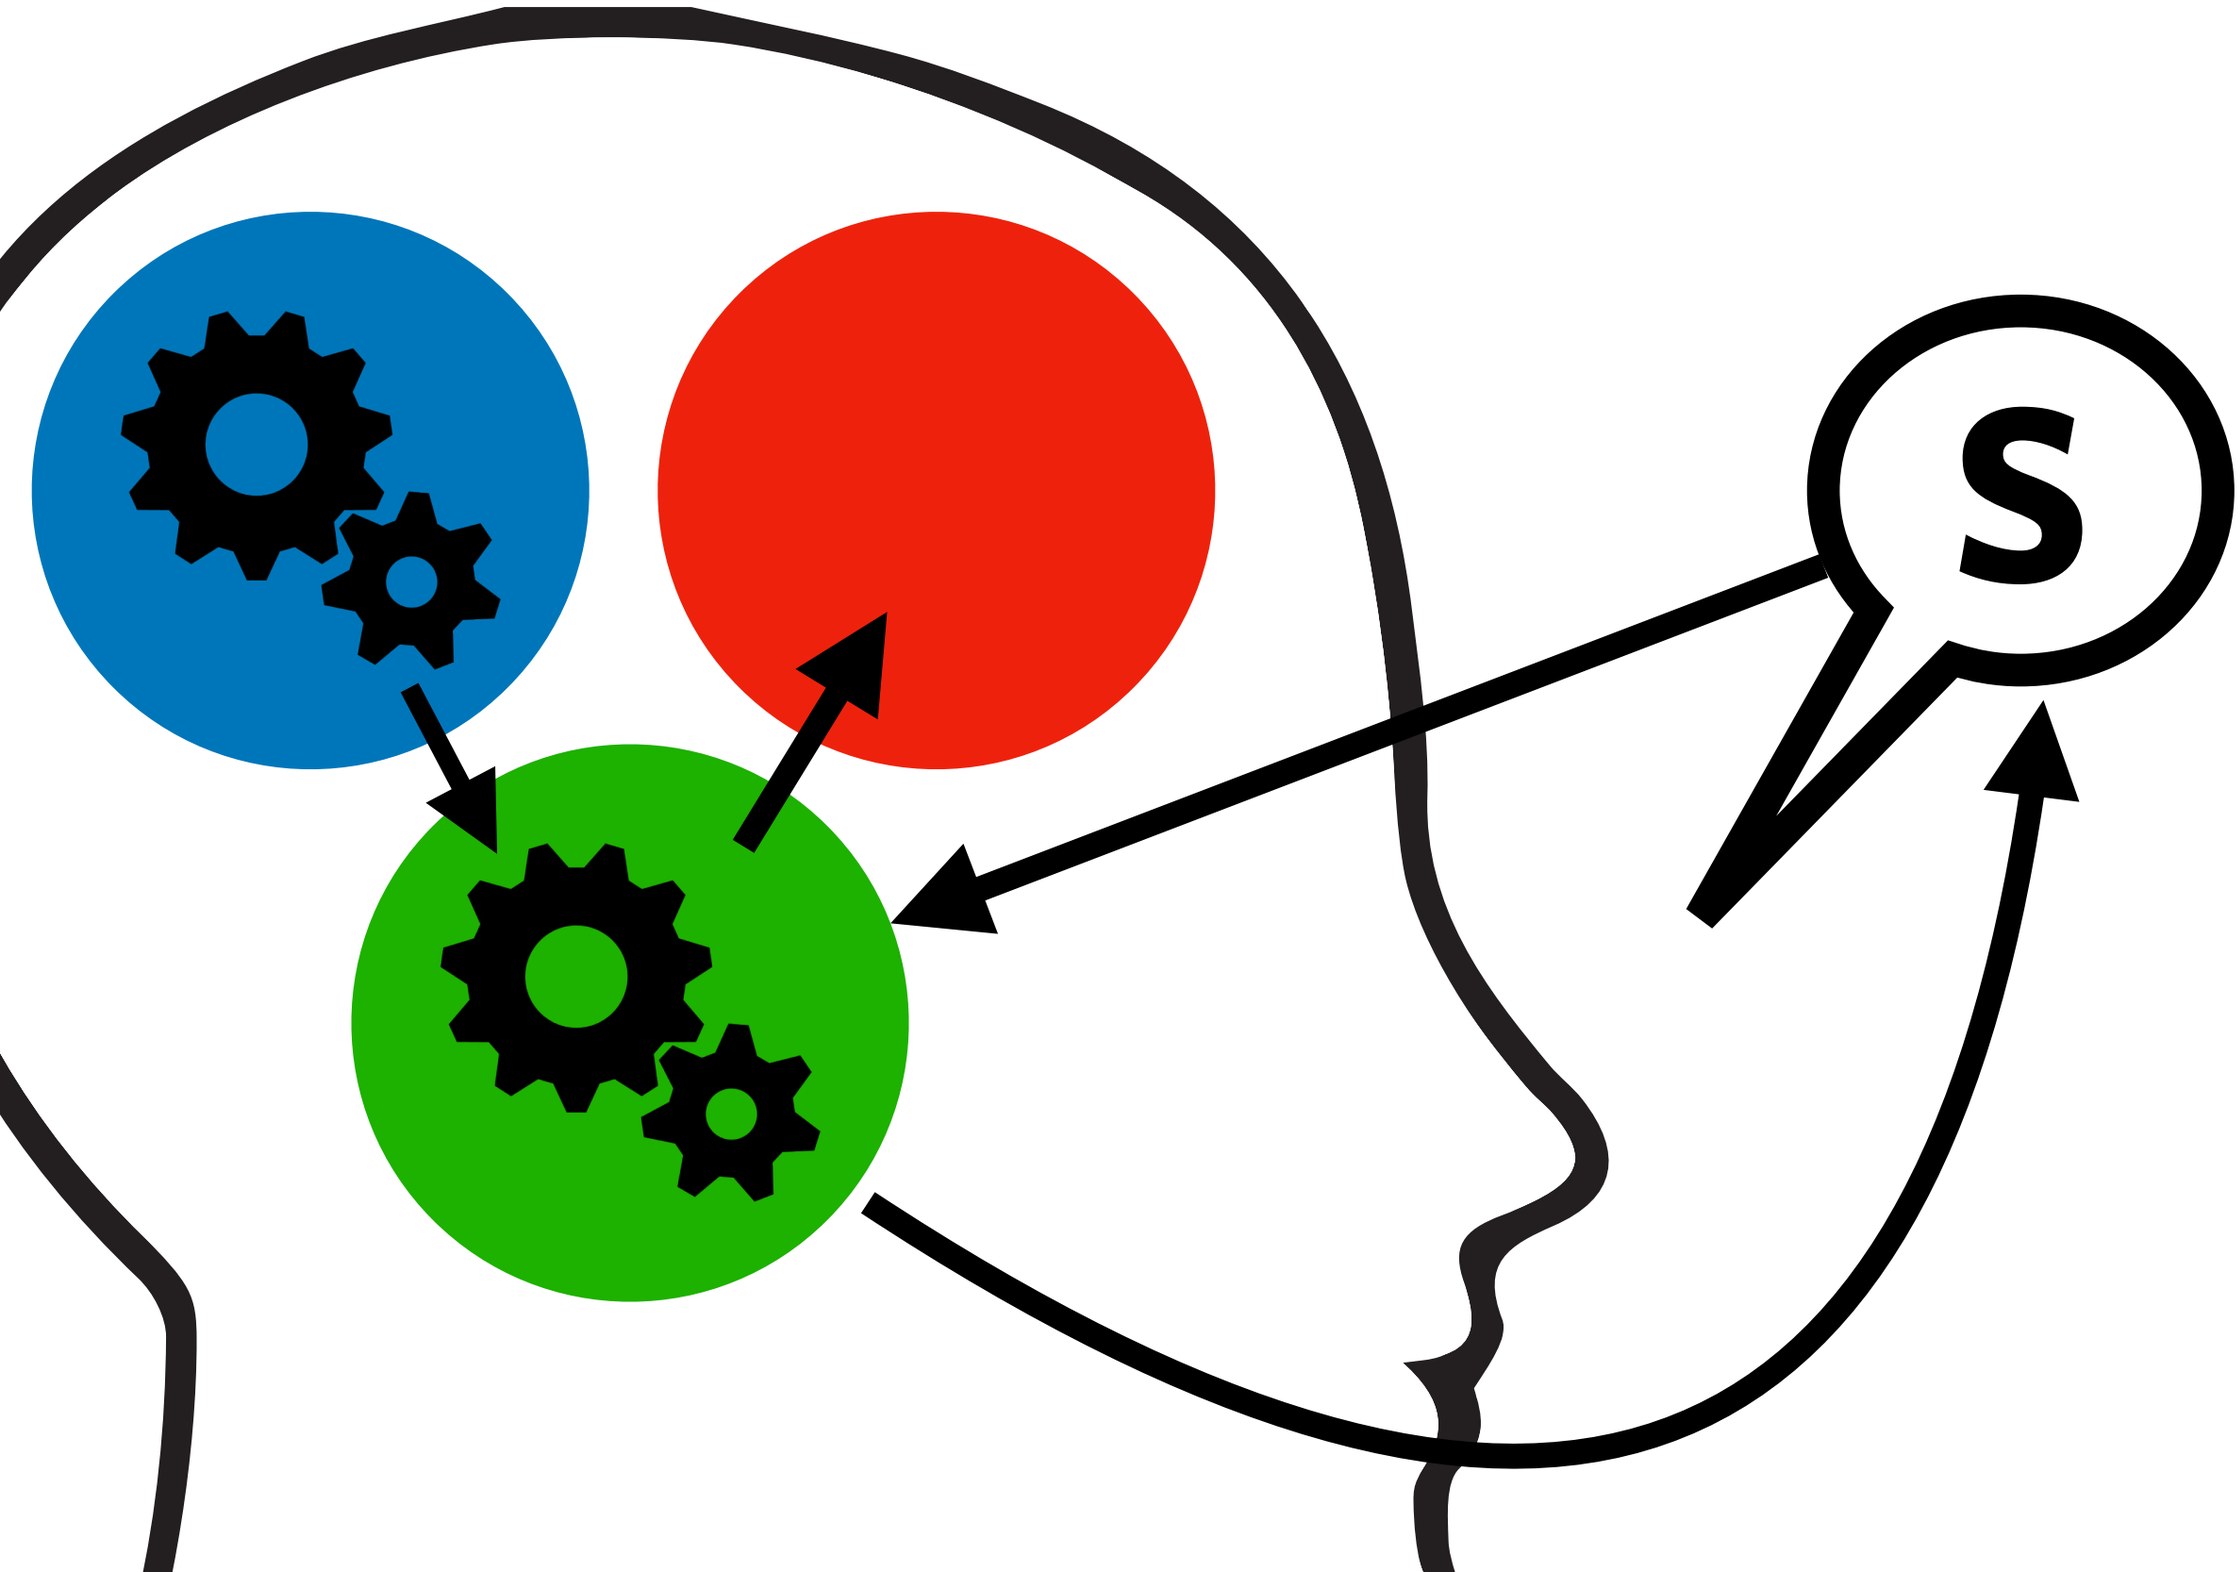

Subvocal Rehearsal

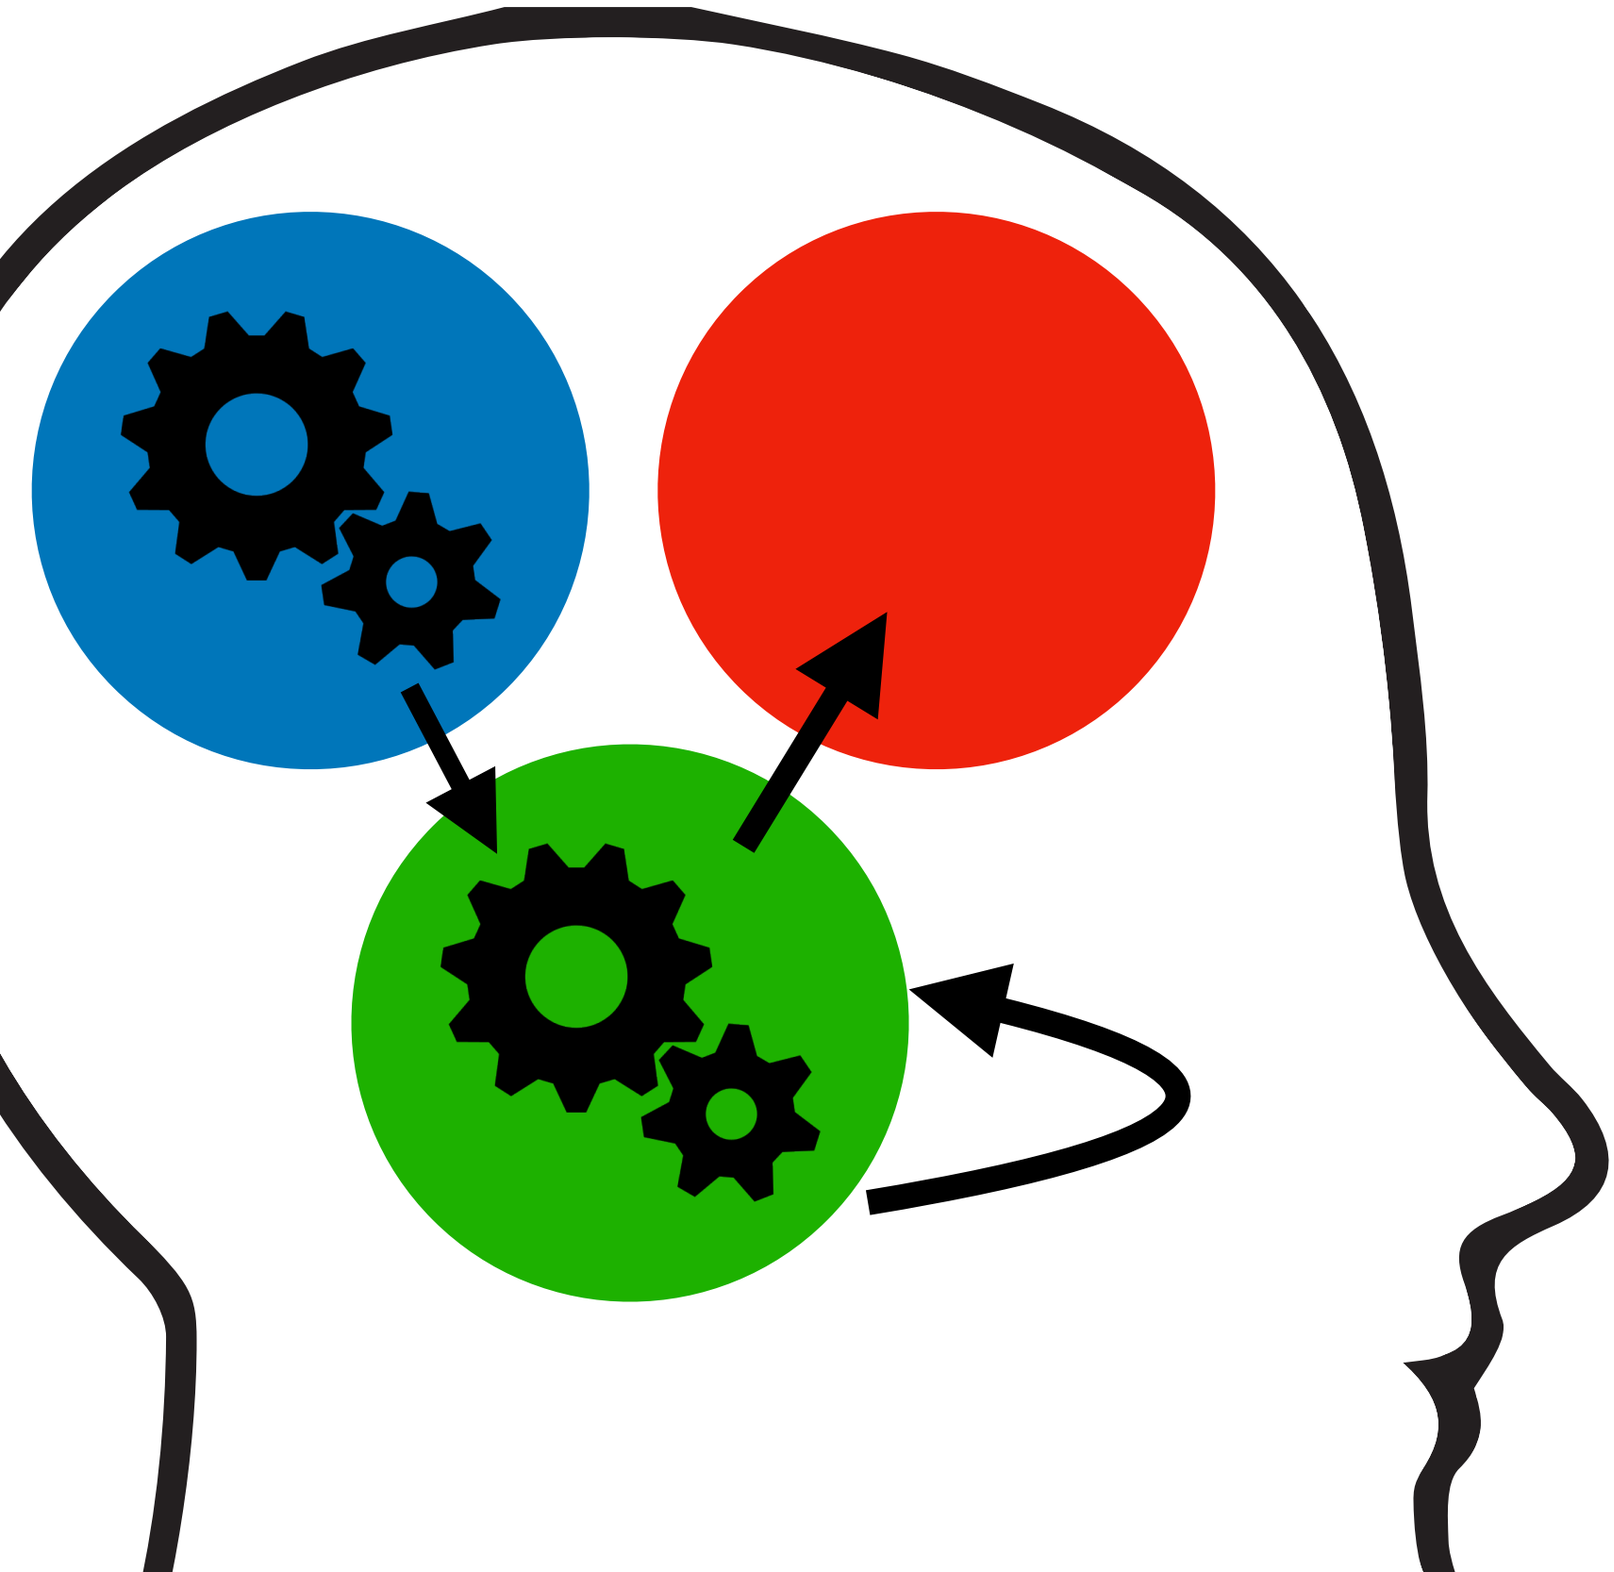

Working Memory

Baddeley (2000)

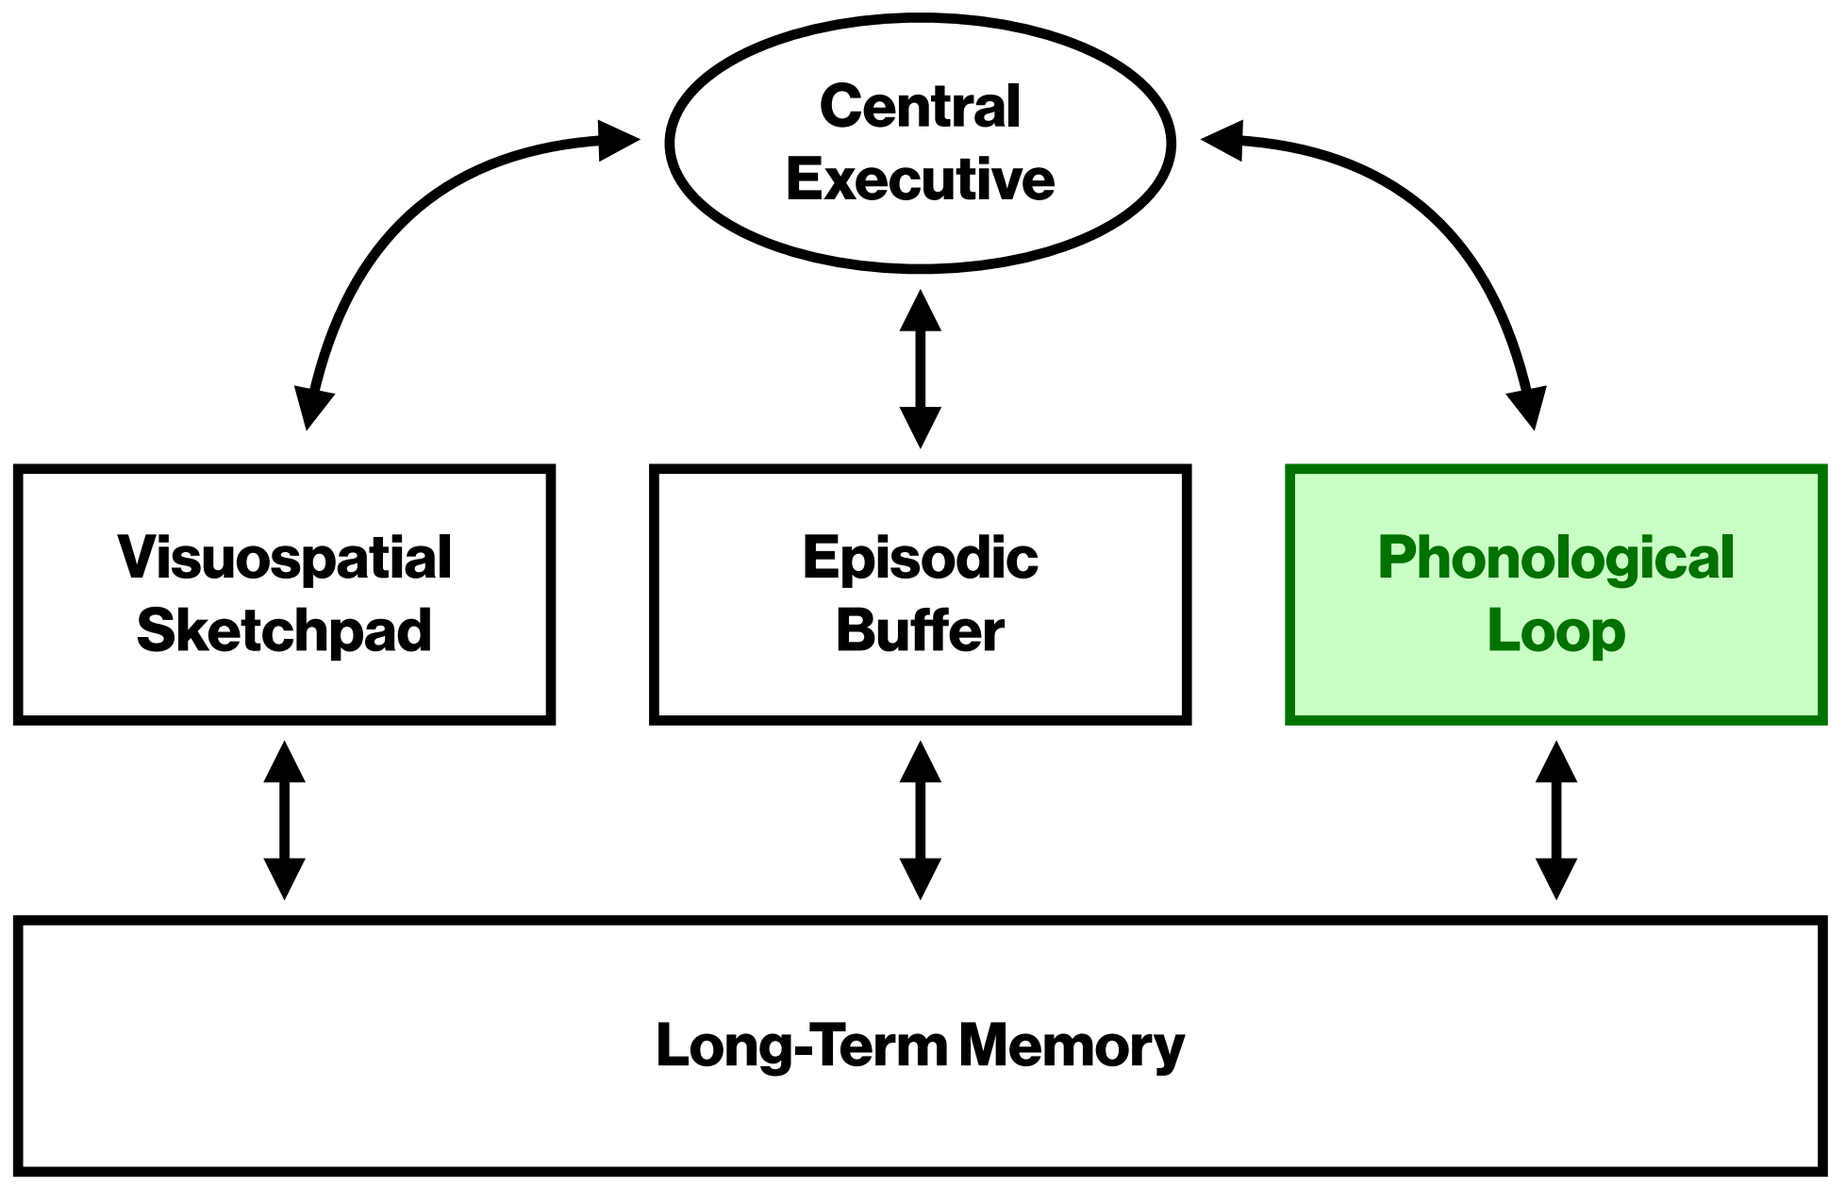

Phonological Loop

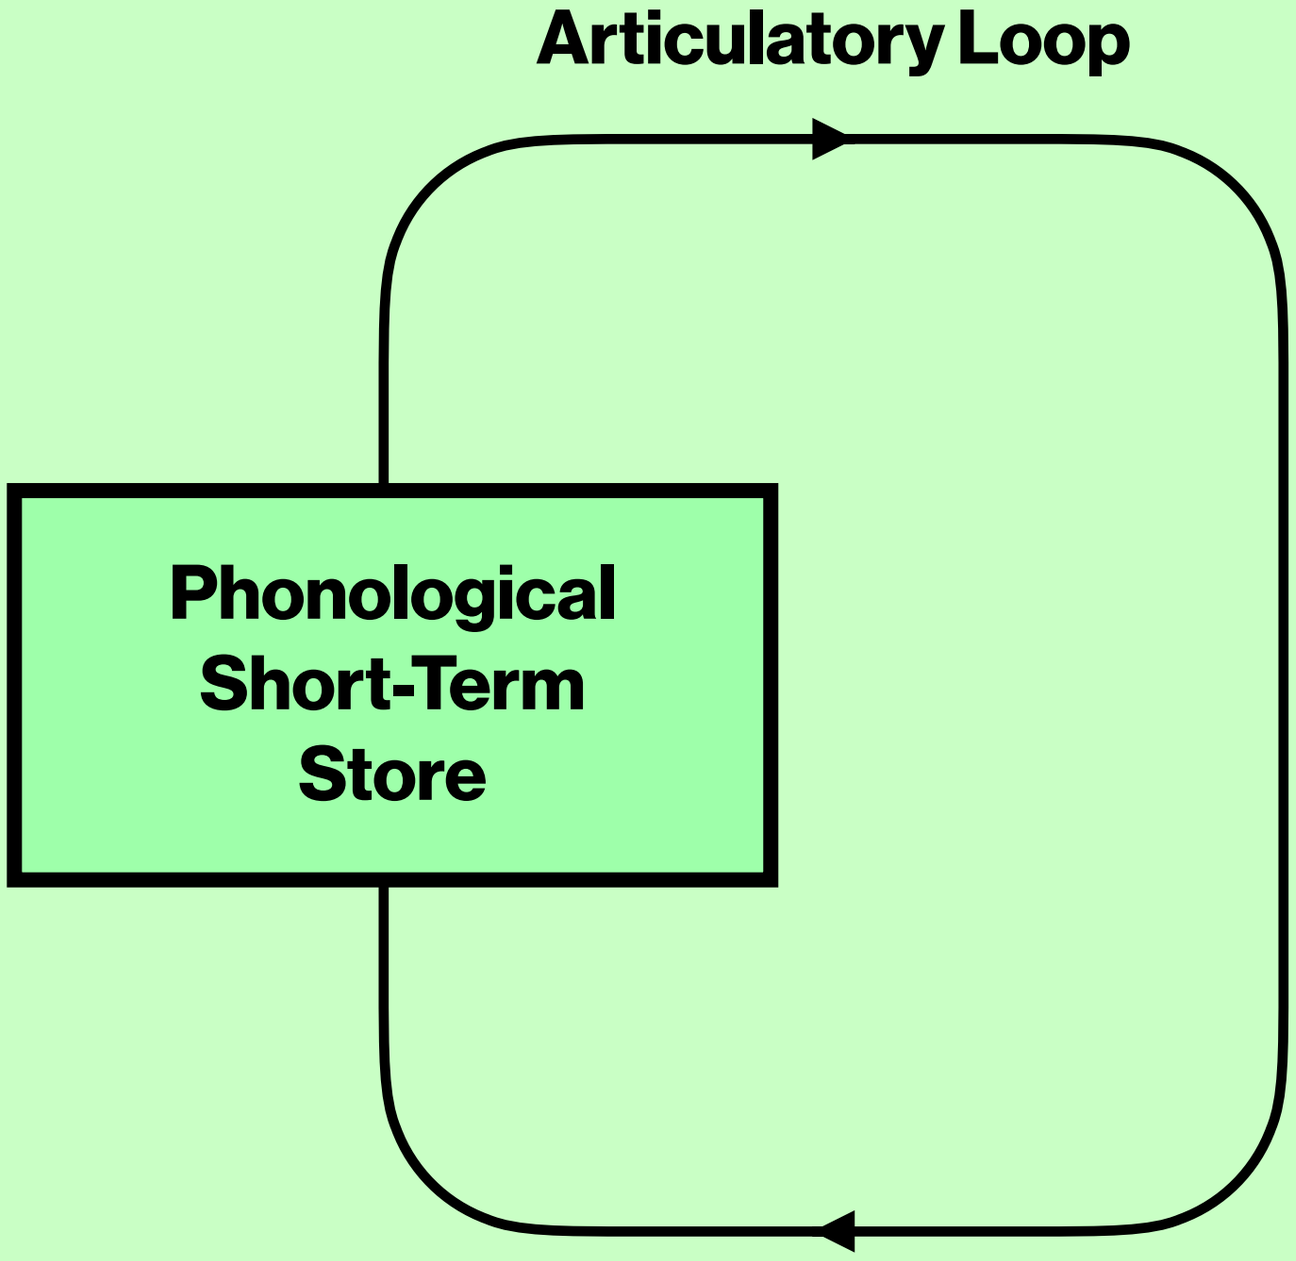

Phonological Loop

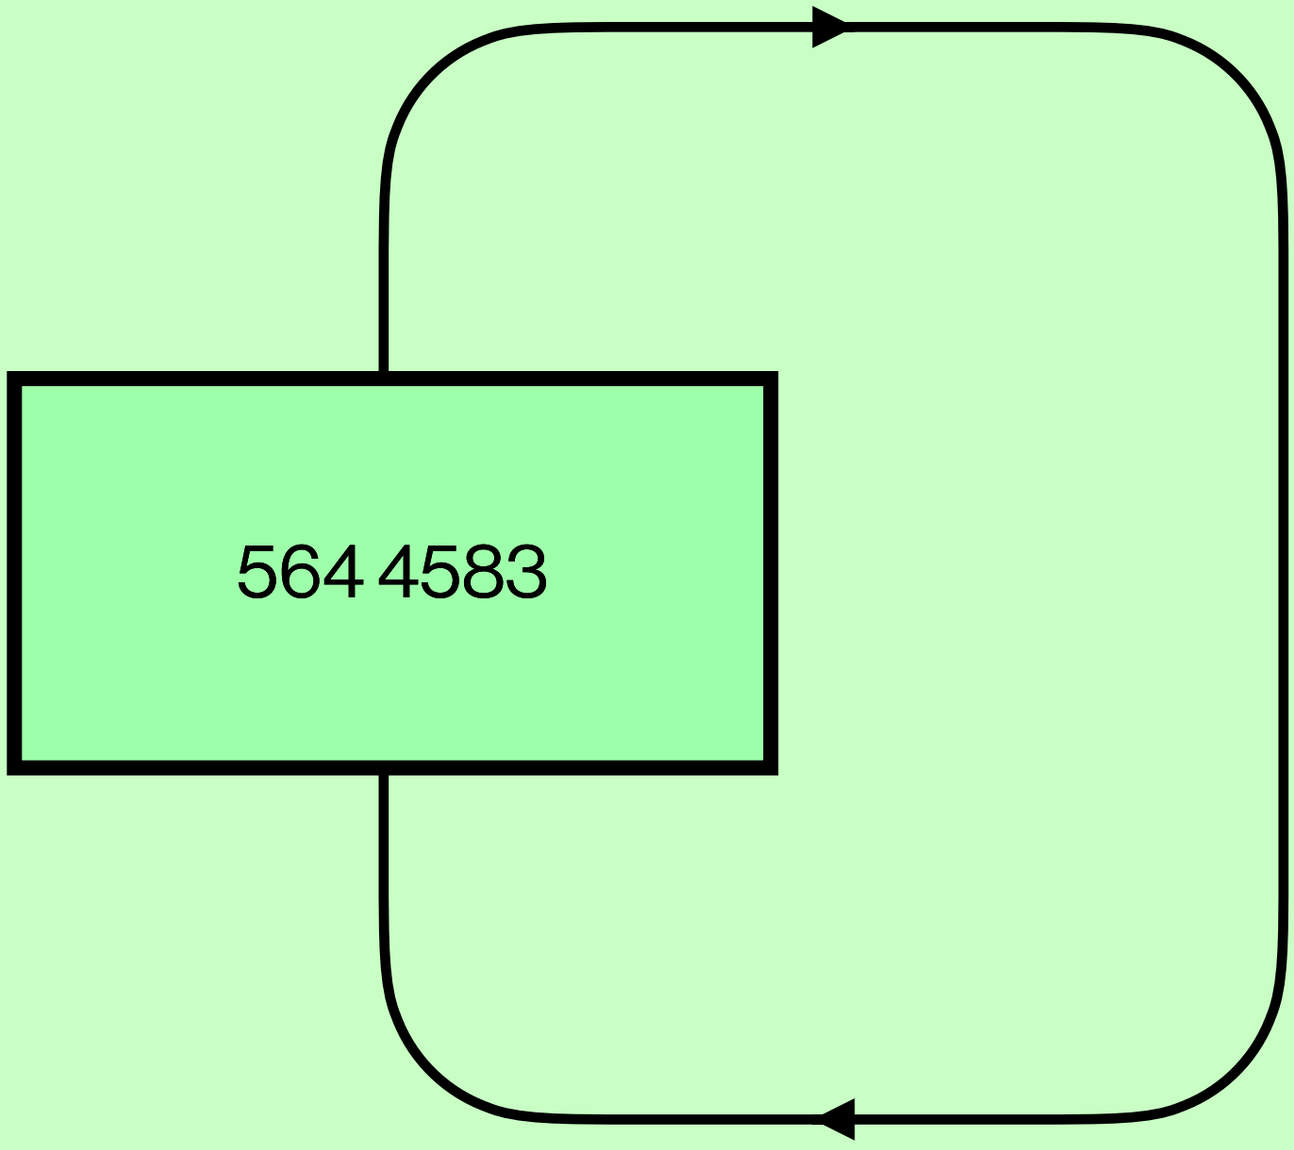

Phonological Loop

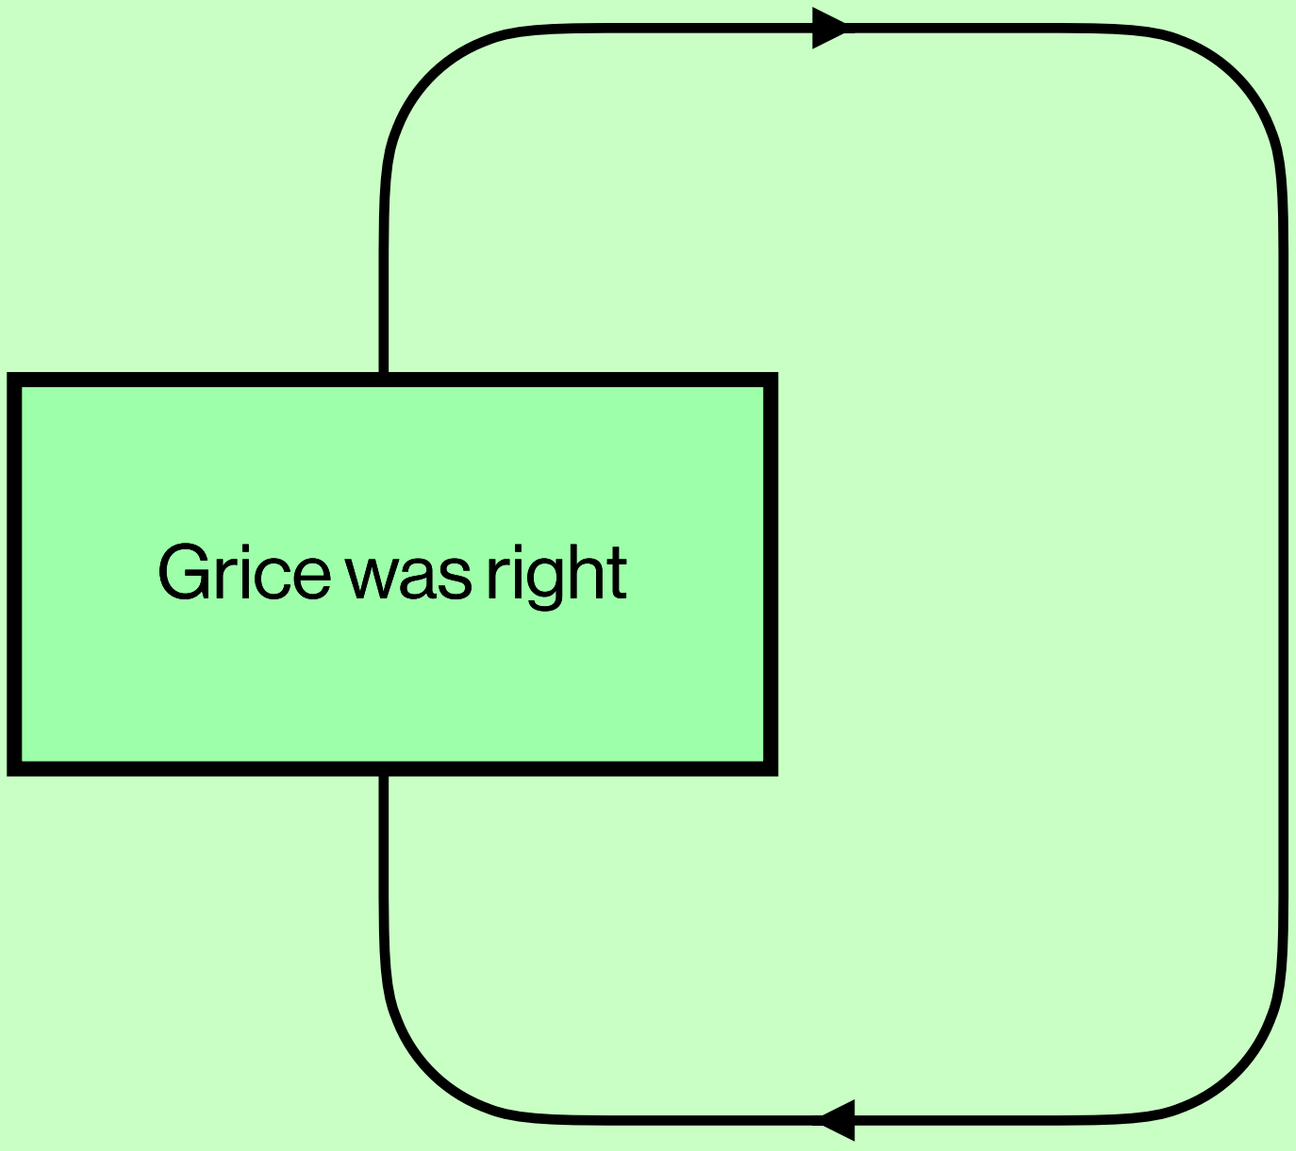

Vol. 63, No. 2 March, 1956

THE PSYCHOLOGICAL REVIEW

THE MAGICAL NUMBER SEVEN, PLUS OR MINUS TWO: SOME LIMITS ON OUR CAPACITY FOR PROCESSING INFORMATION [1]

GEORGE A. MILLER

Harvard University

My problem is that I have been persecuted by an integer. For seven years this number has followed me around, has intruded in my most private data, and has assaulted me from the pages of our most public journals. This number as-

judgment. Historical accident, however, has decreed that they should have another name. We now call them experiments on the capacity of people to transmit information. Since these experiments would not have been done

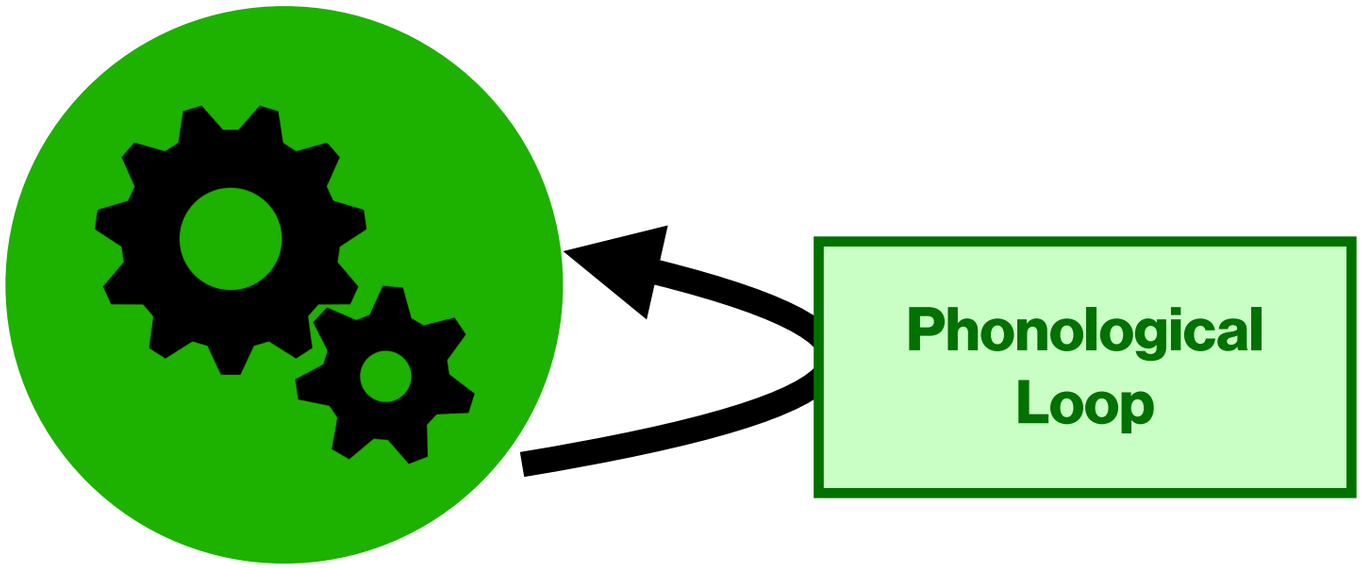
Phonological
Loop

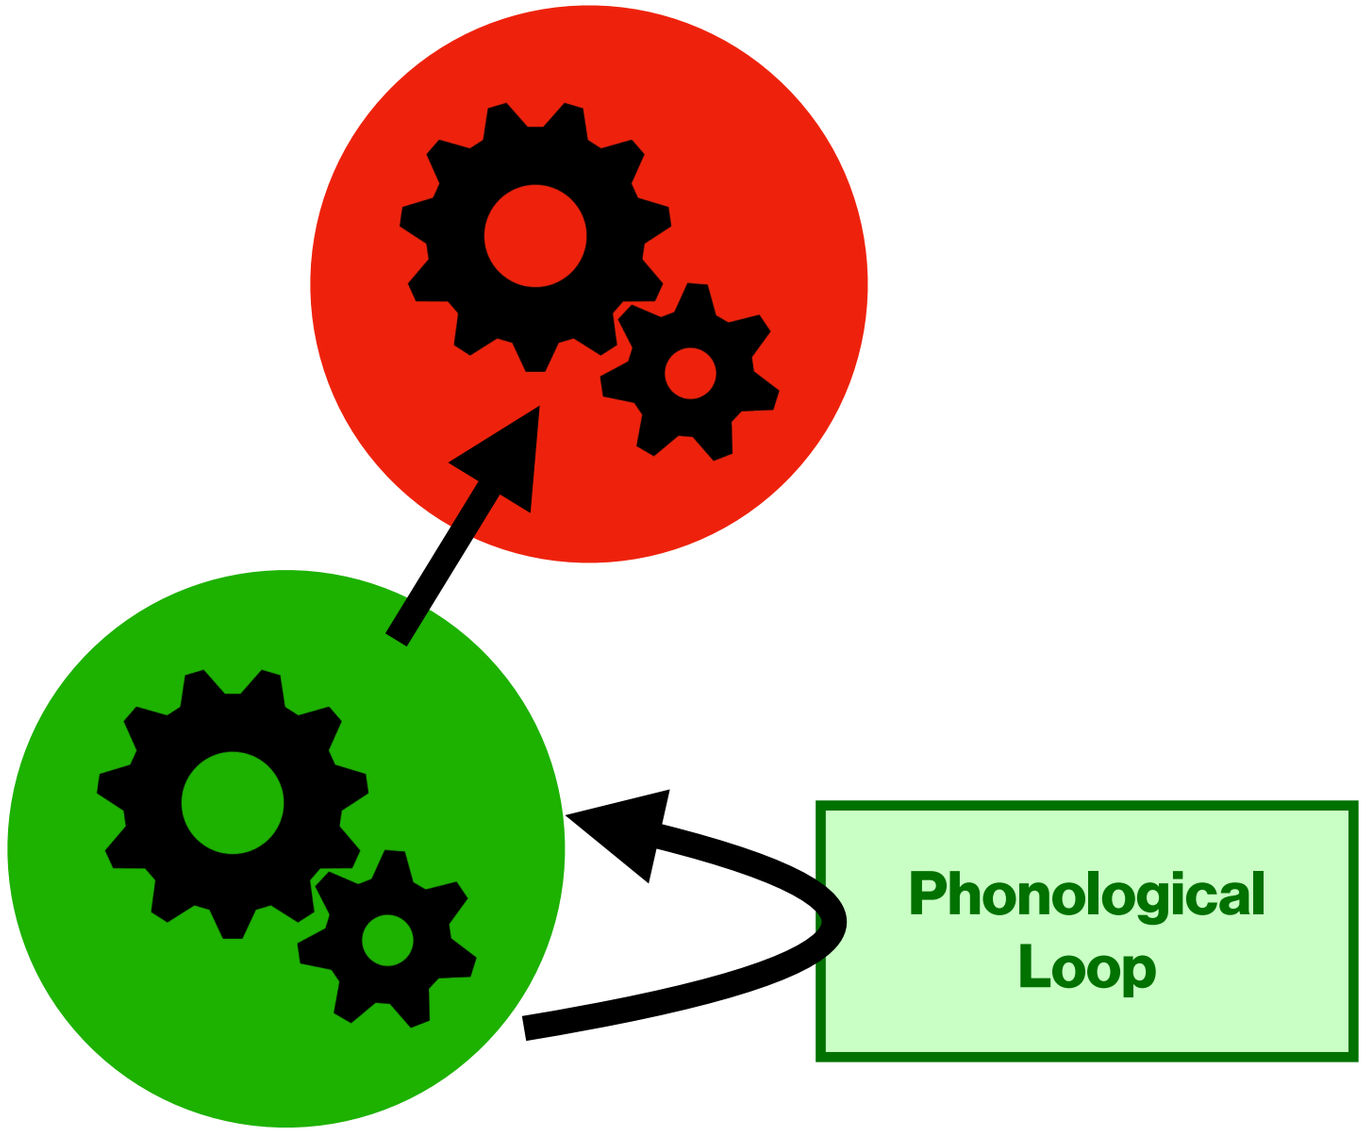
Phonological
Loop

Subvocal Rehearsal

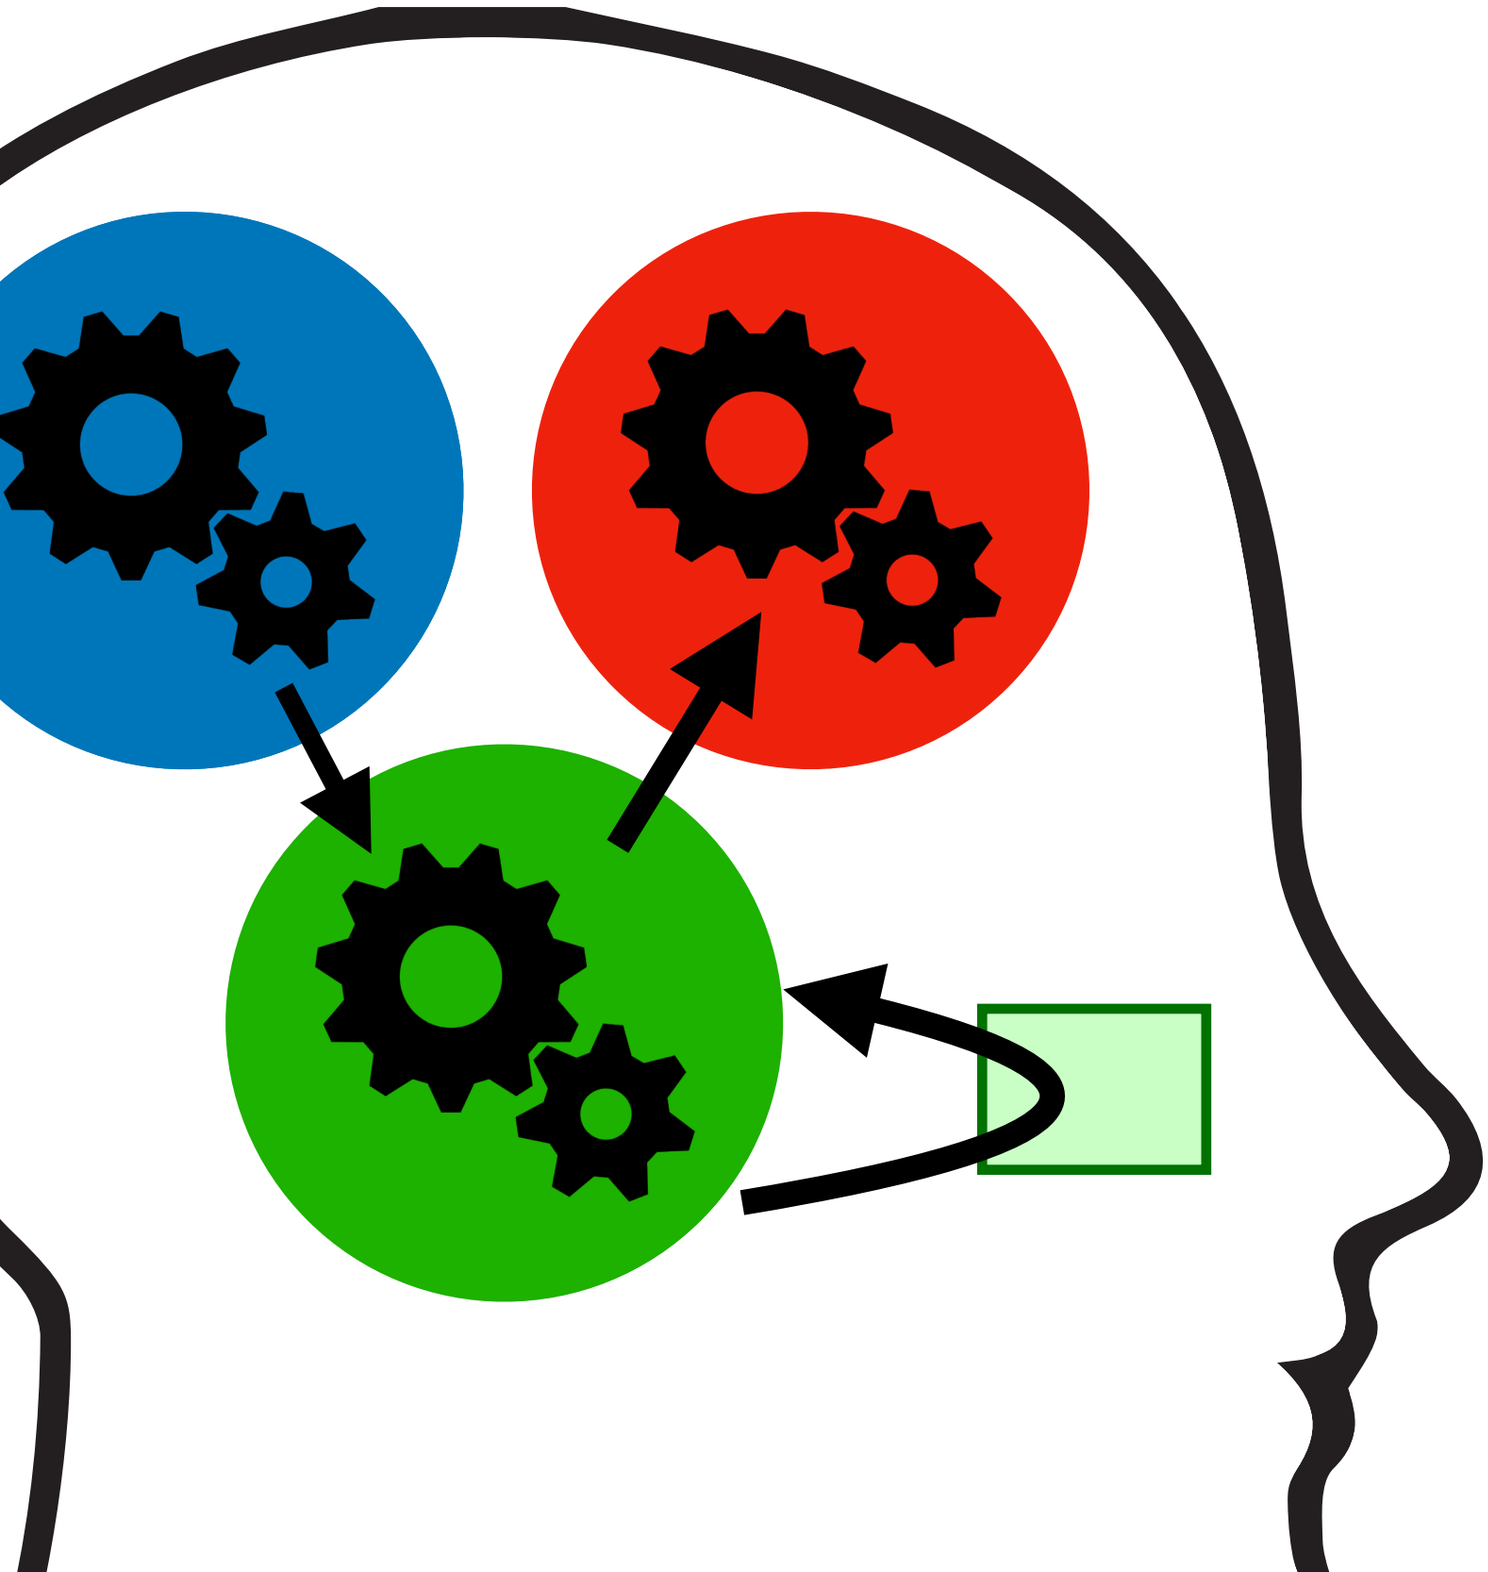

- Verbal-working-memory deficit → more egocentric speech and interpretation (Lin et al 2010)
- cognitive load → more egocentric speech and interpretation (Keysar 2008)
- cognitive load → faster speech, but more errors (Ivanova and Ferreira 2008)

Cf. work on language production and audience design by V. Ferreira (2009)

Say the next thing!

Subvocally repeating instructions enhances performance on some tasks.

(Baddeley, Chincotta, & Adlam 2001; Emerson and Miyake 2003; Saeki and Saito 2004)

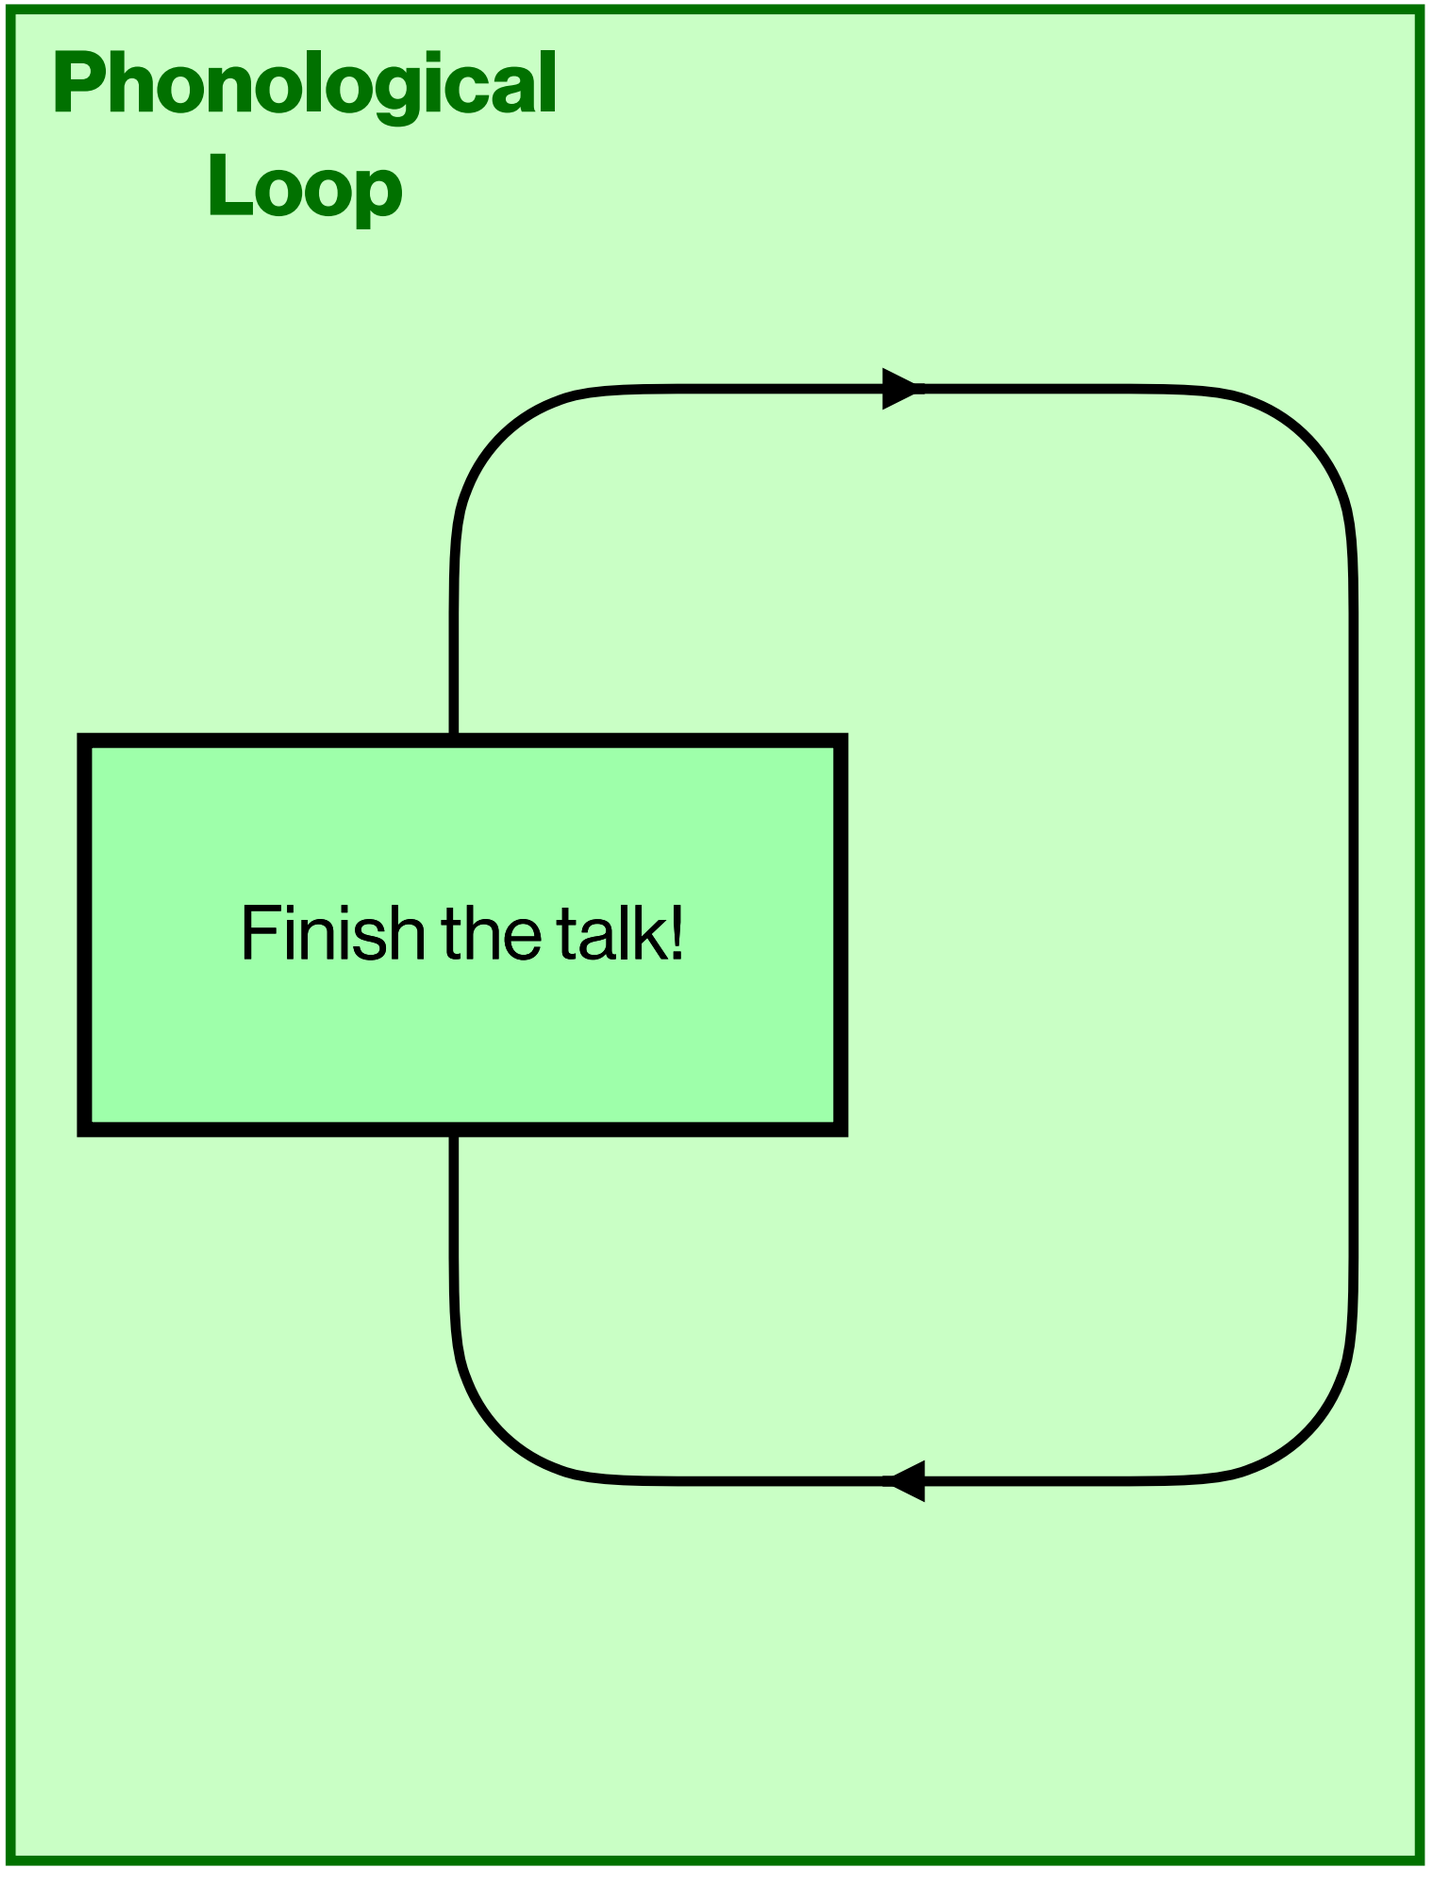

Russian blues reveal effects of language on color discrimination

Jonathan Winawer*†‡, Nathan Witthoft*‡, Michael C. Frank*, Lisa Wu§, Alex R. Wade¶, and Lera Boroditsky‡

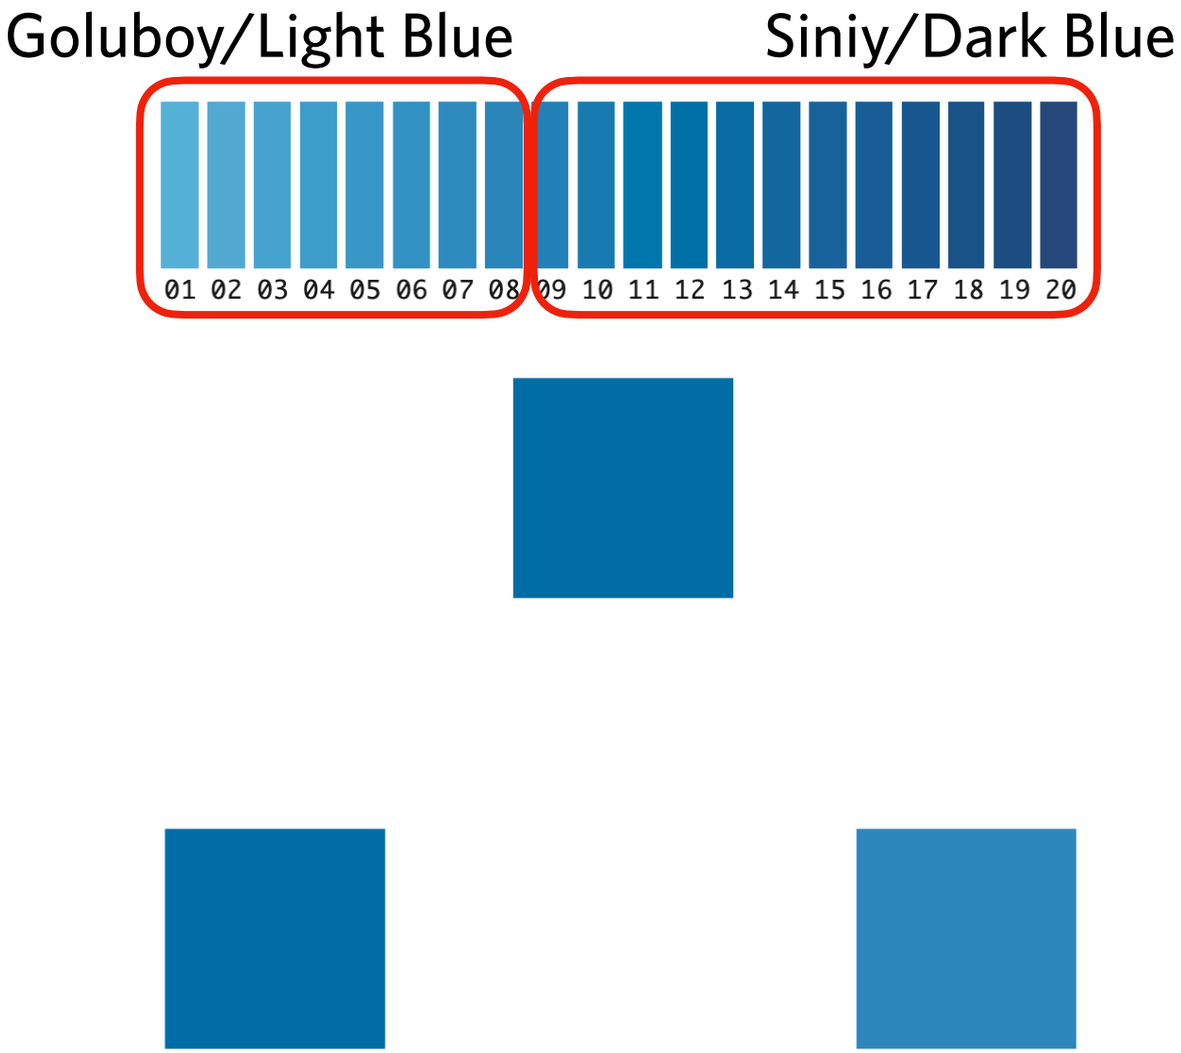

Fig. 1. The 20 blue colors used in this study are shown at the top of the figure. An example triad of color squares used in this study is shown at the bottom of the figure. Subjects were instructed to pick which one of the two bottom squares matched the color of the top square.

- Task: Say which two patches match.
- Russian speakers were faster when the colors were located near the light/dark boundary.
- The effect disappears when subjects' verbal working memories are occupied.
- So, a hypothesis: the difference is that two words take up more space in verbal working memory than one, or takes longer to read and write.

Russian Speaker's Phonological Loop

Goluboy or Siniy?

English Speaker's Phonological Loop

Light blue
or dark blue?

Clarifying Thoughts?

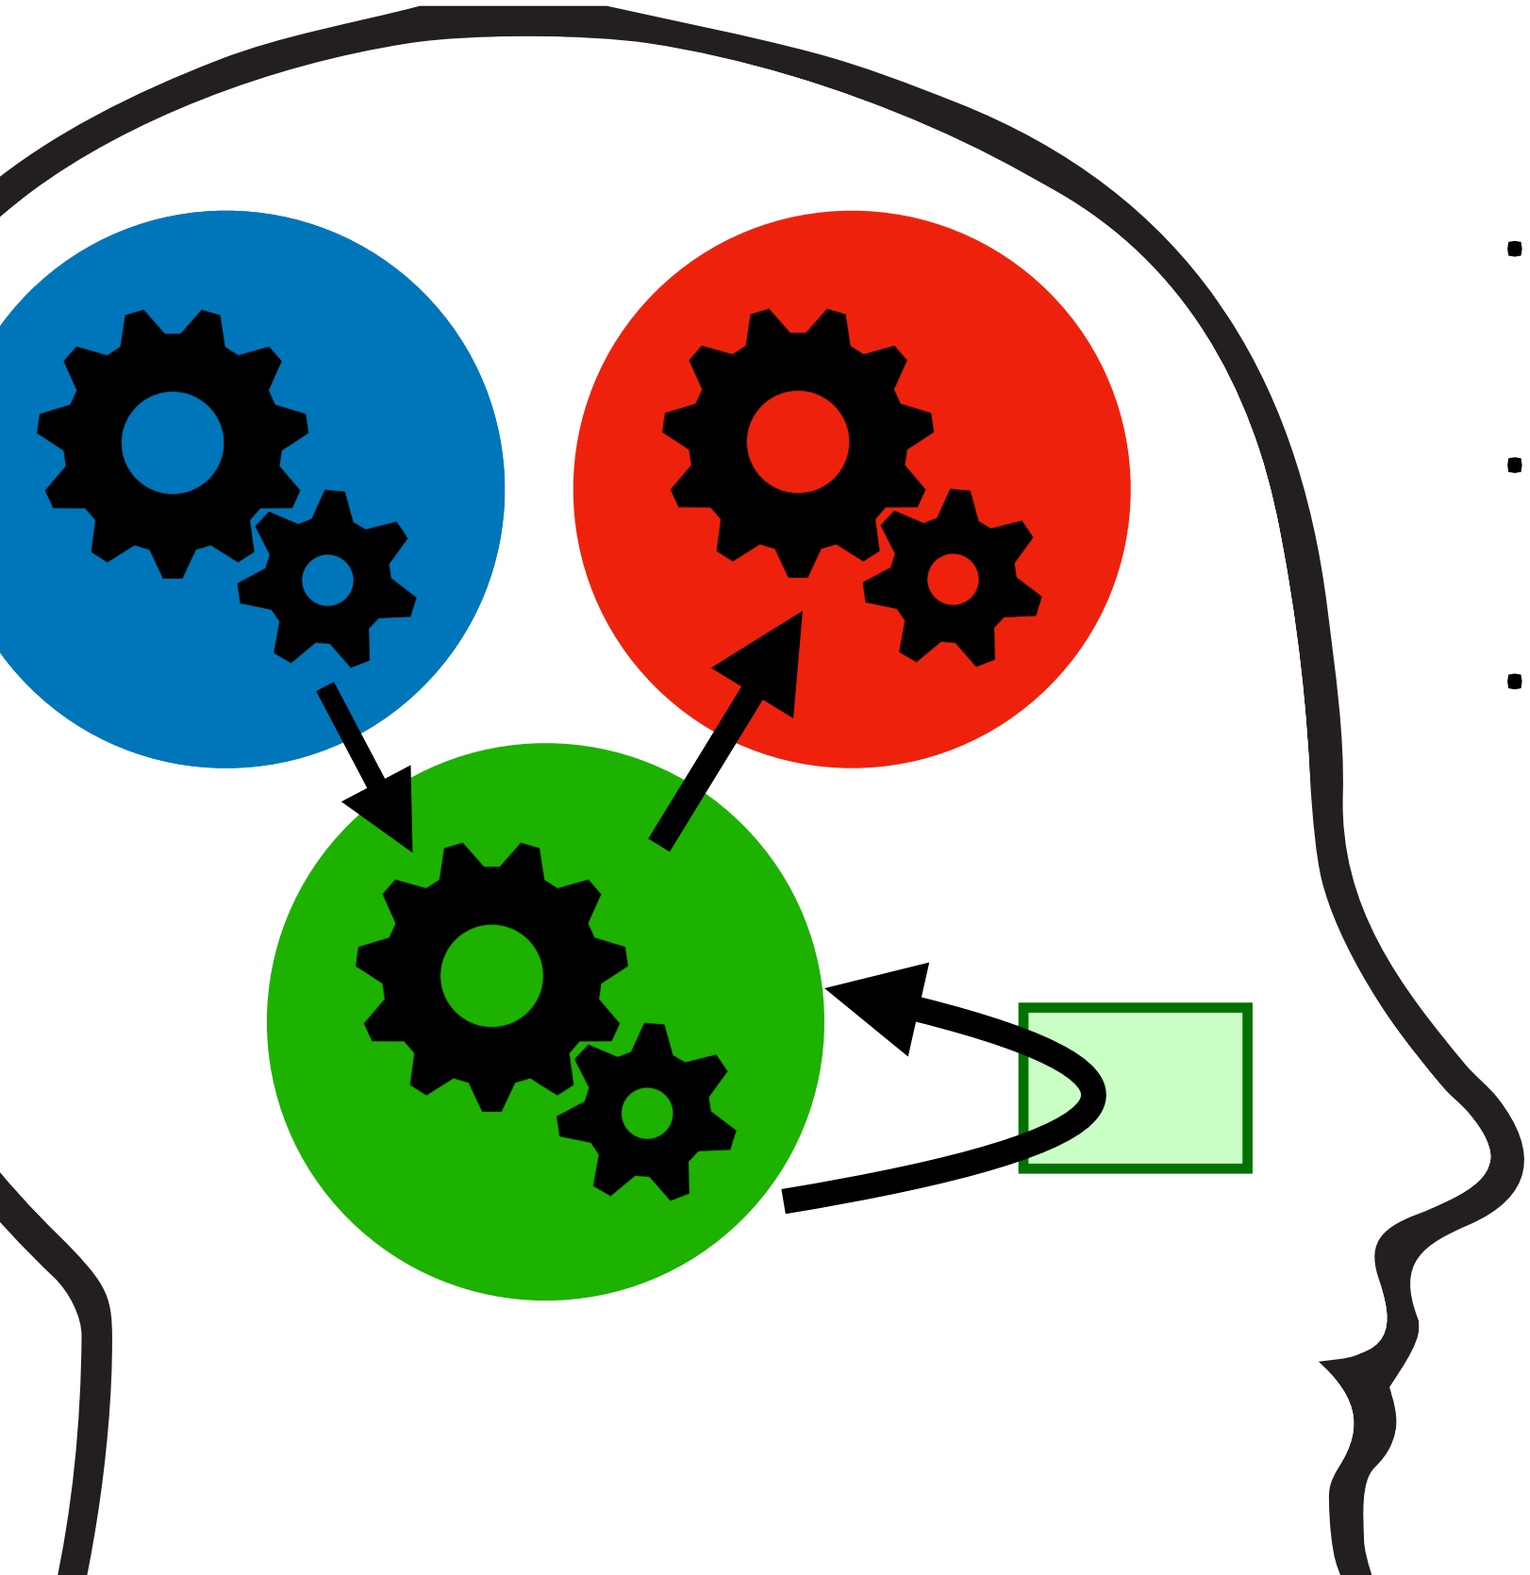

- It *seems* like we sometimes speak in order to clarify our thoughts.
- Writing is a good example of this for many people.
- Is this even compatible with Grice's view that what we say is determined by our intentions?

Planning

PRIOR INTENTION

Intention to practice saying that .

Grice was right

SUBPLAN

Instruction for language system to say that

Grice was right

Clarifying Thoughts

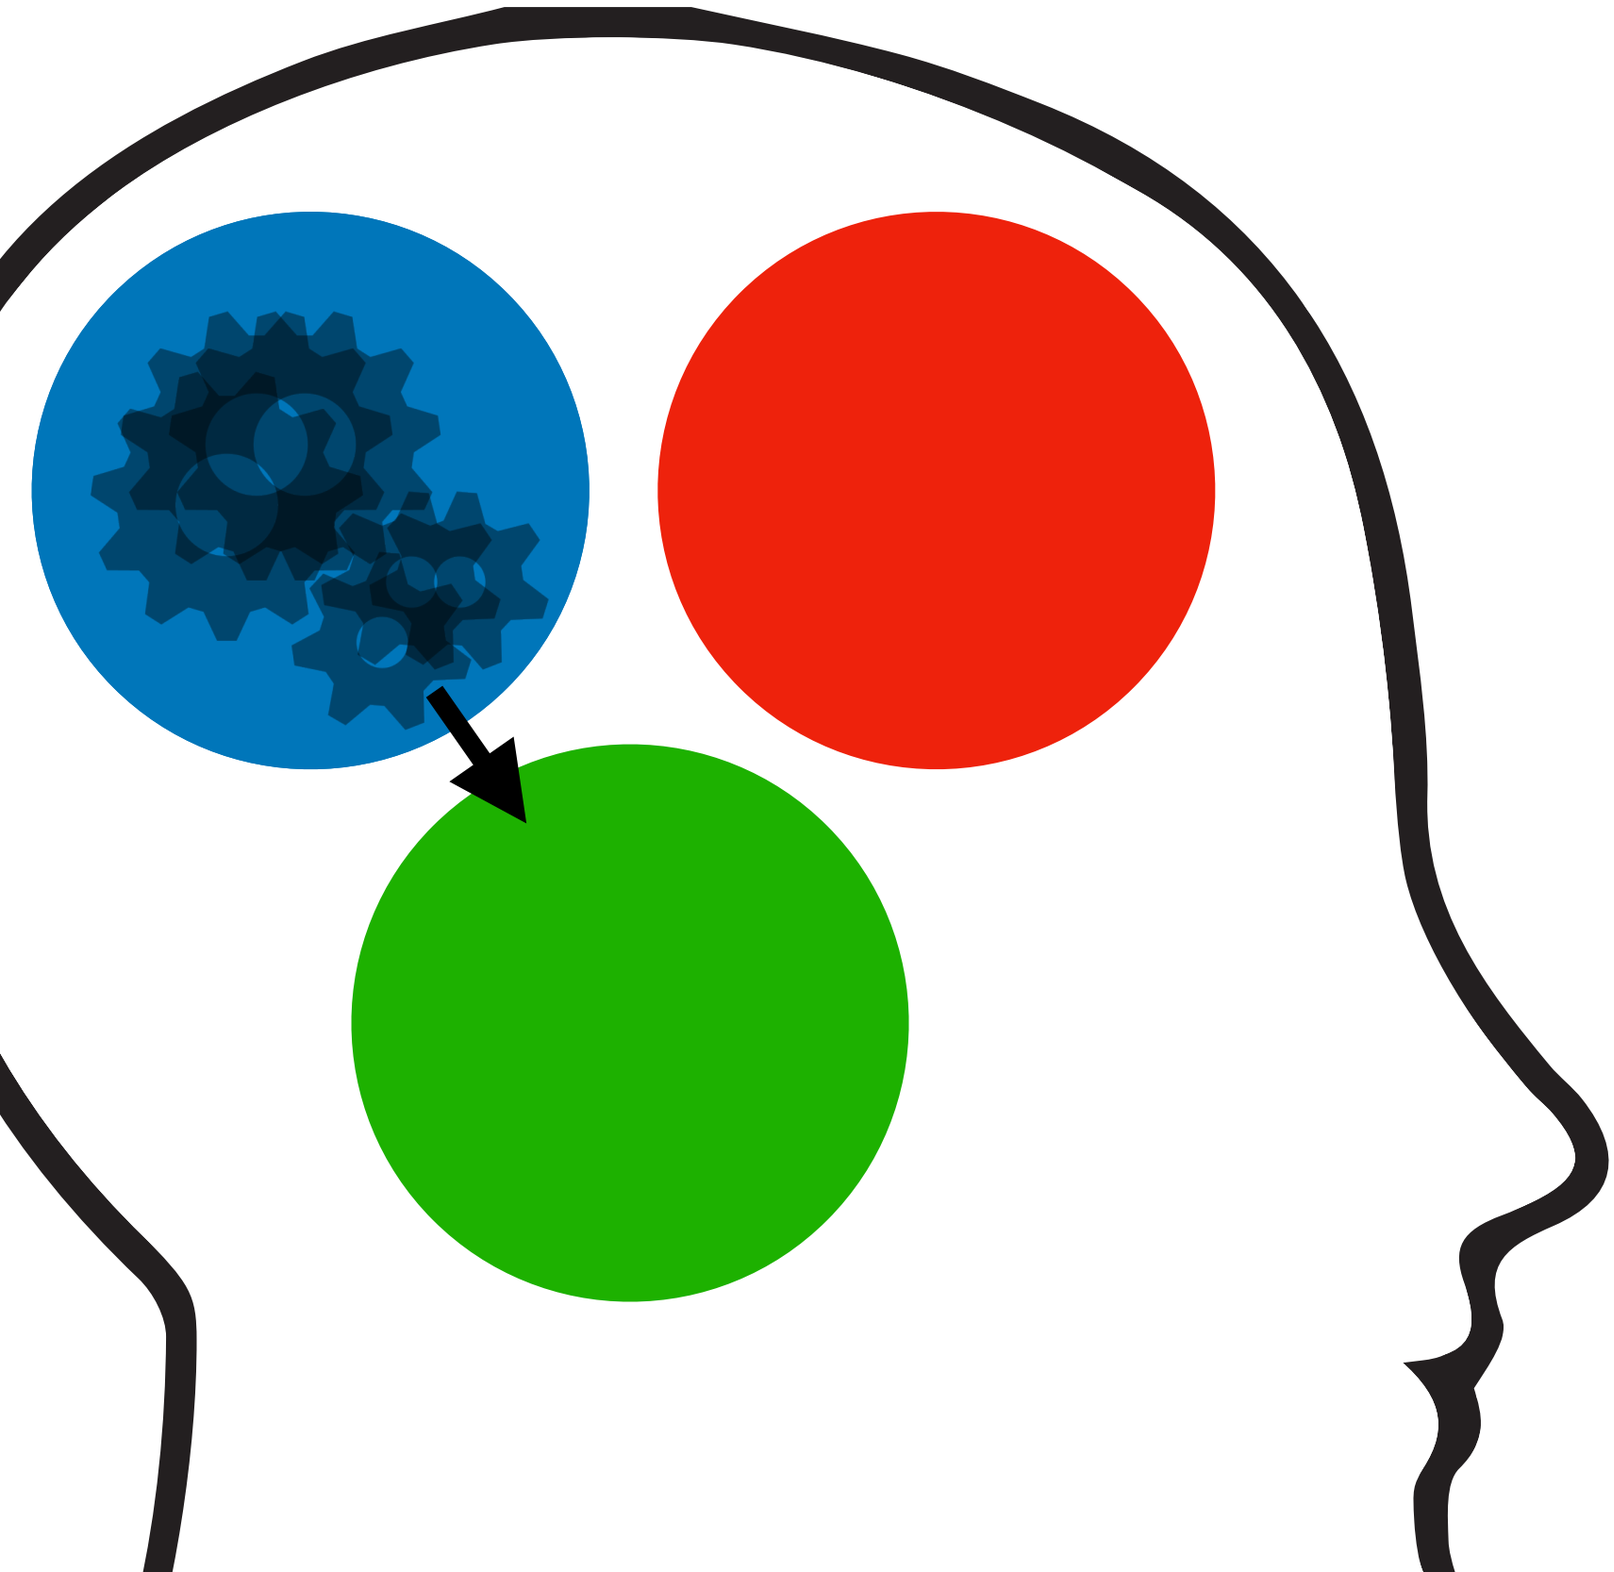

Clarifying Thoughts

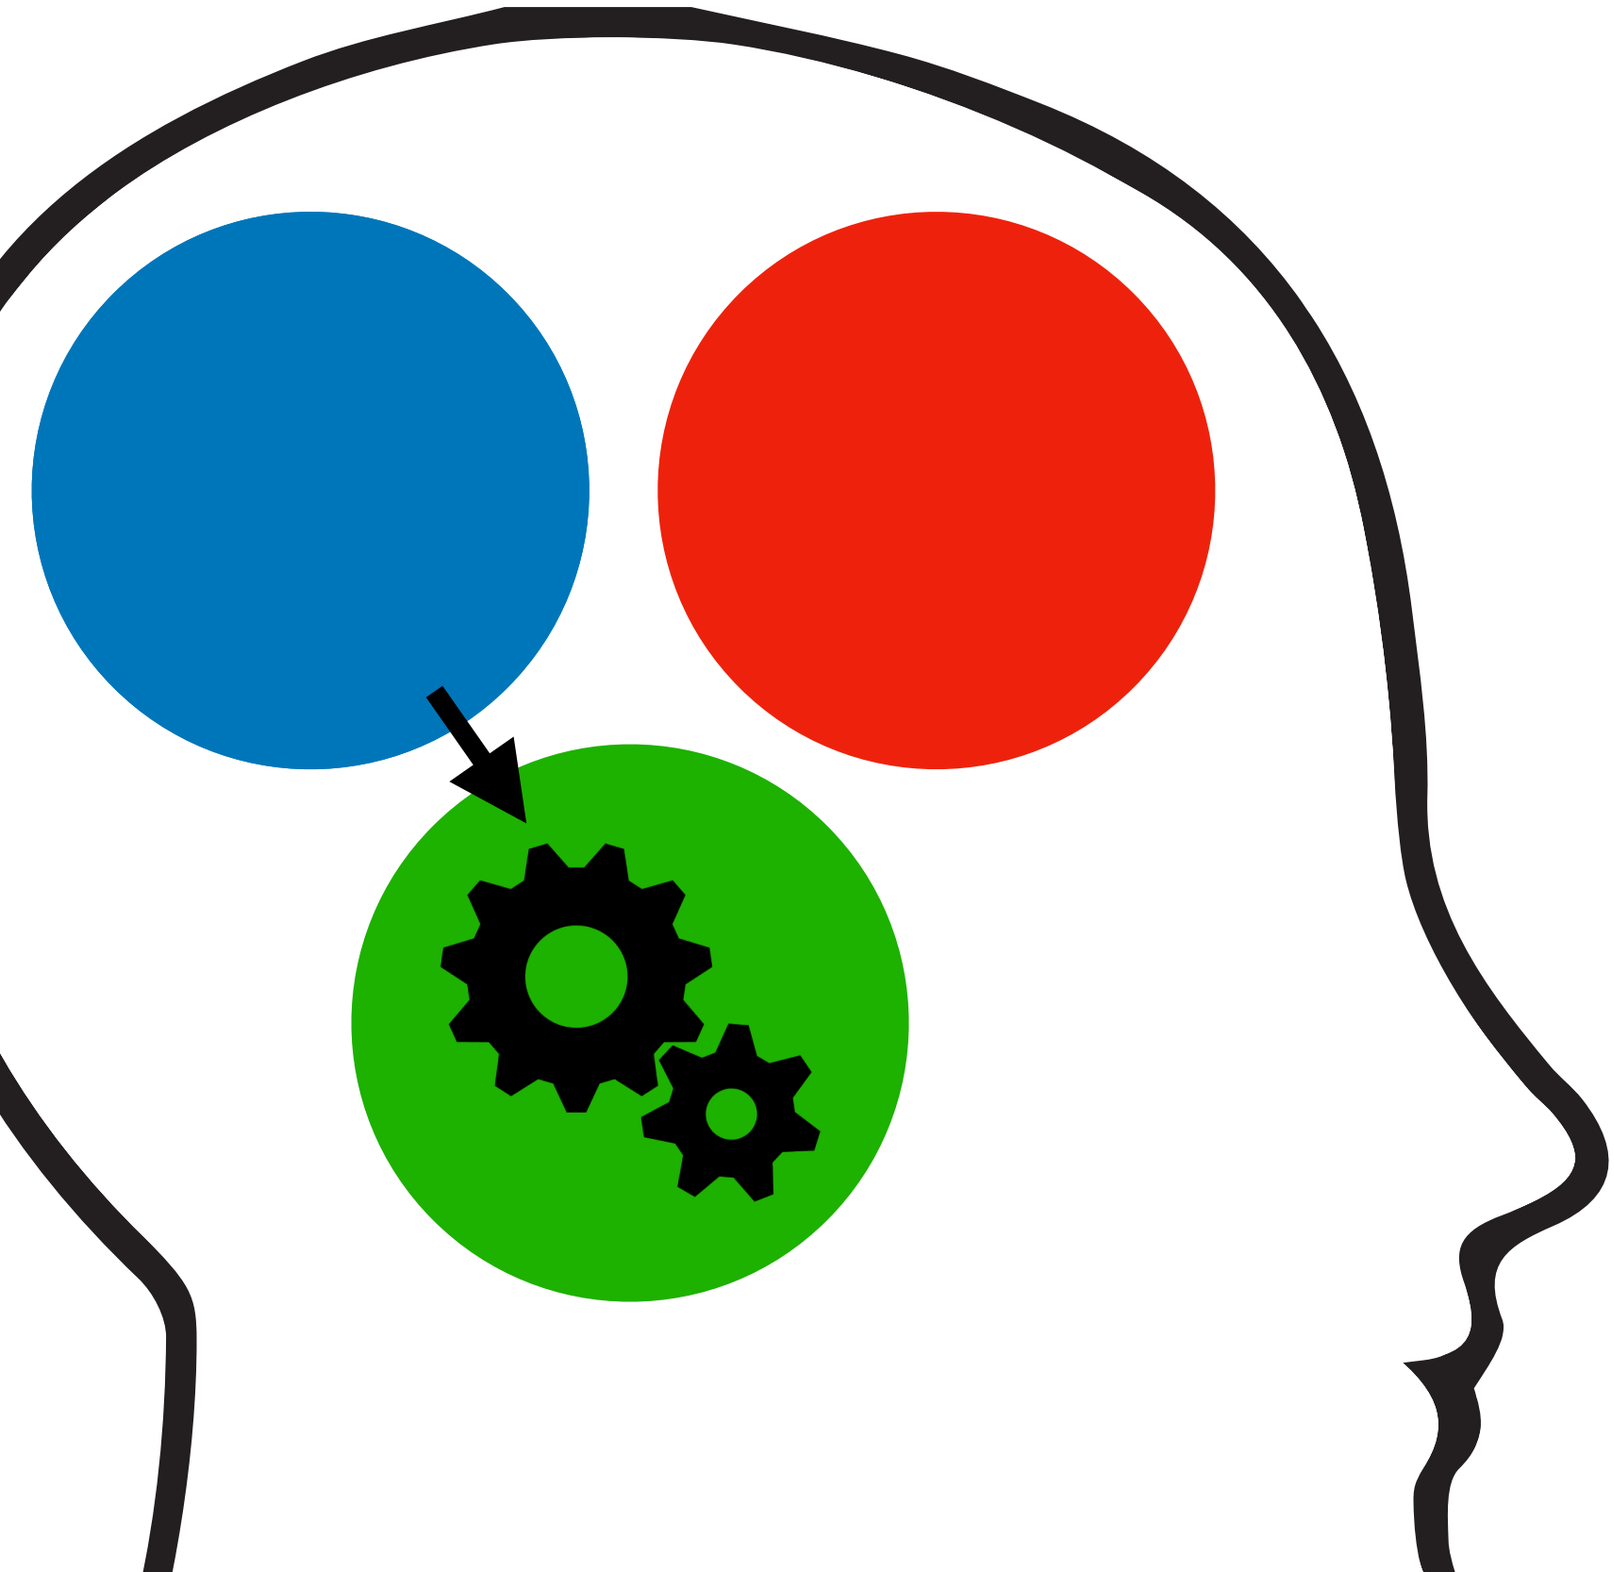

Clarifying Thoughts

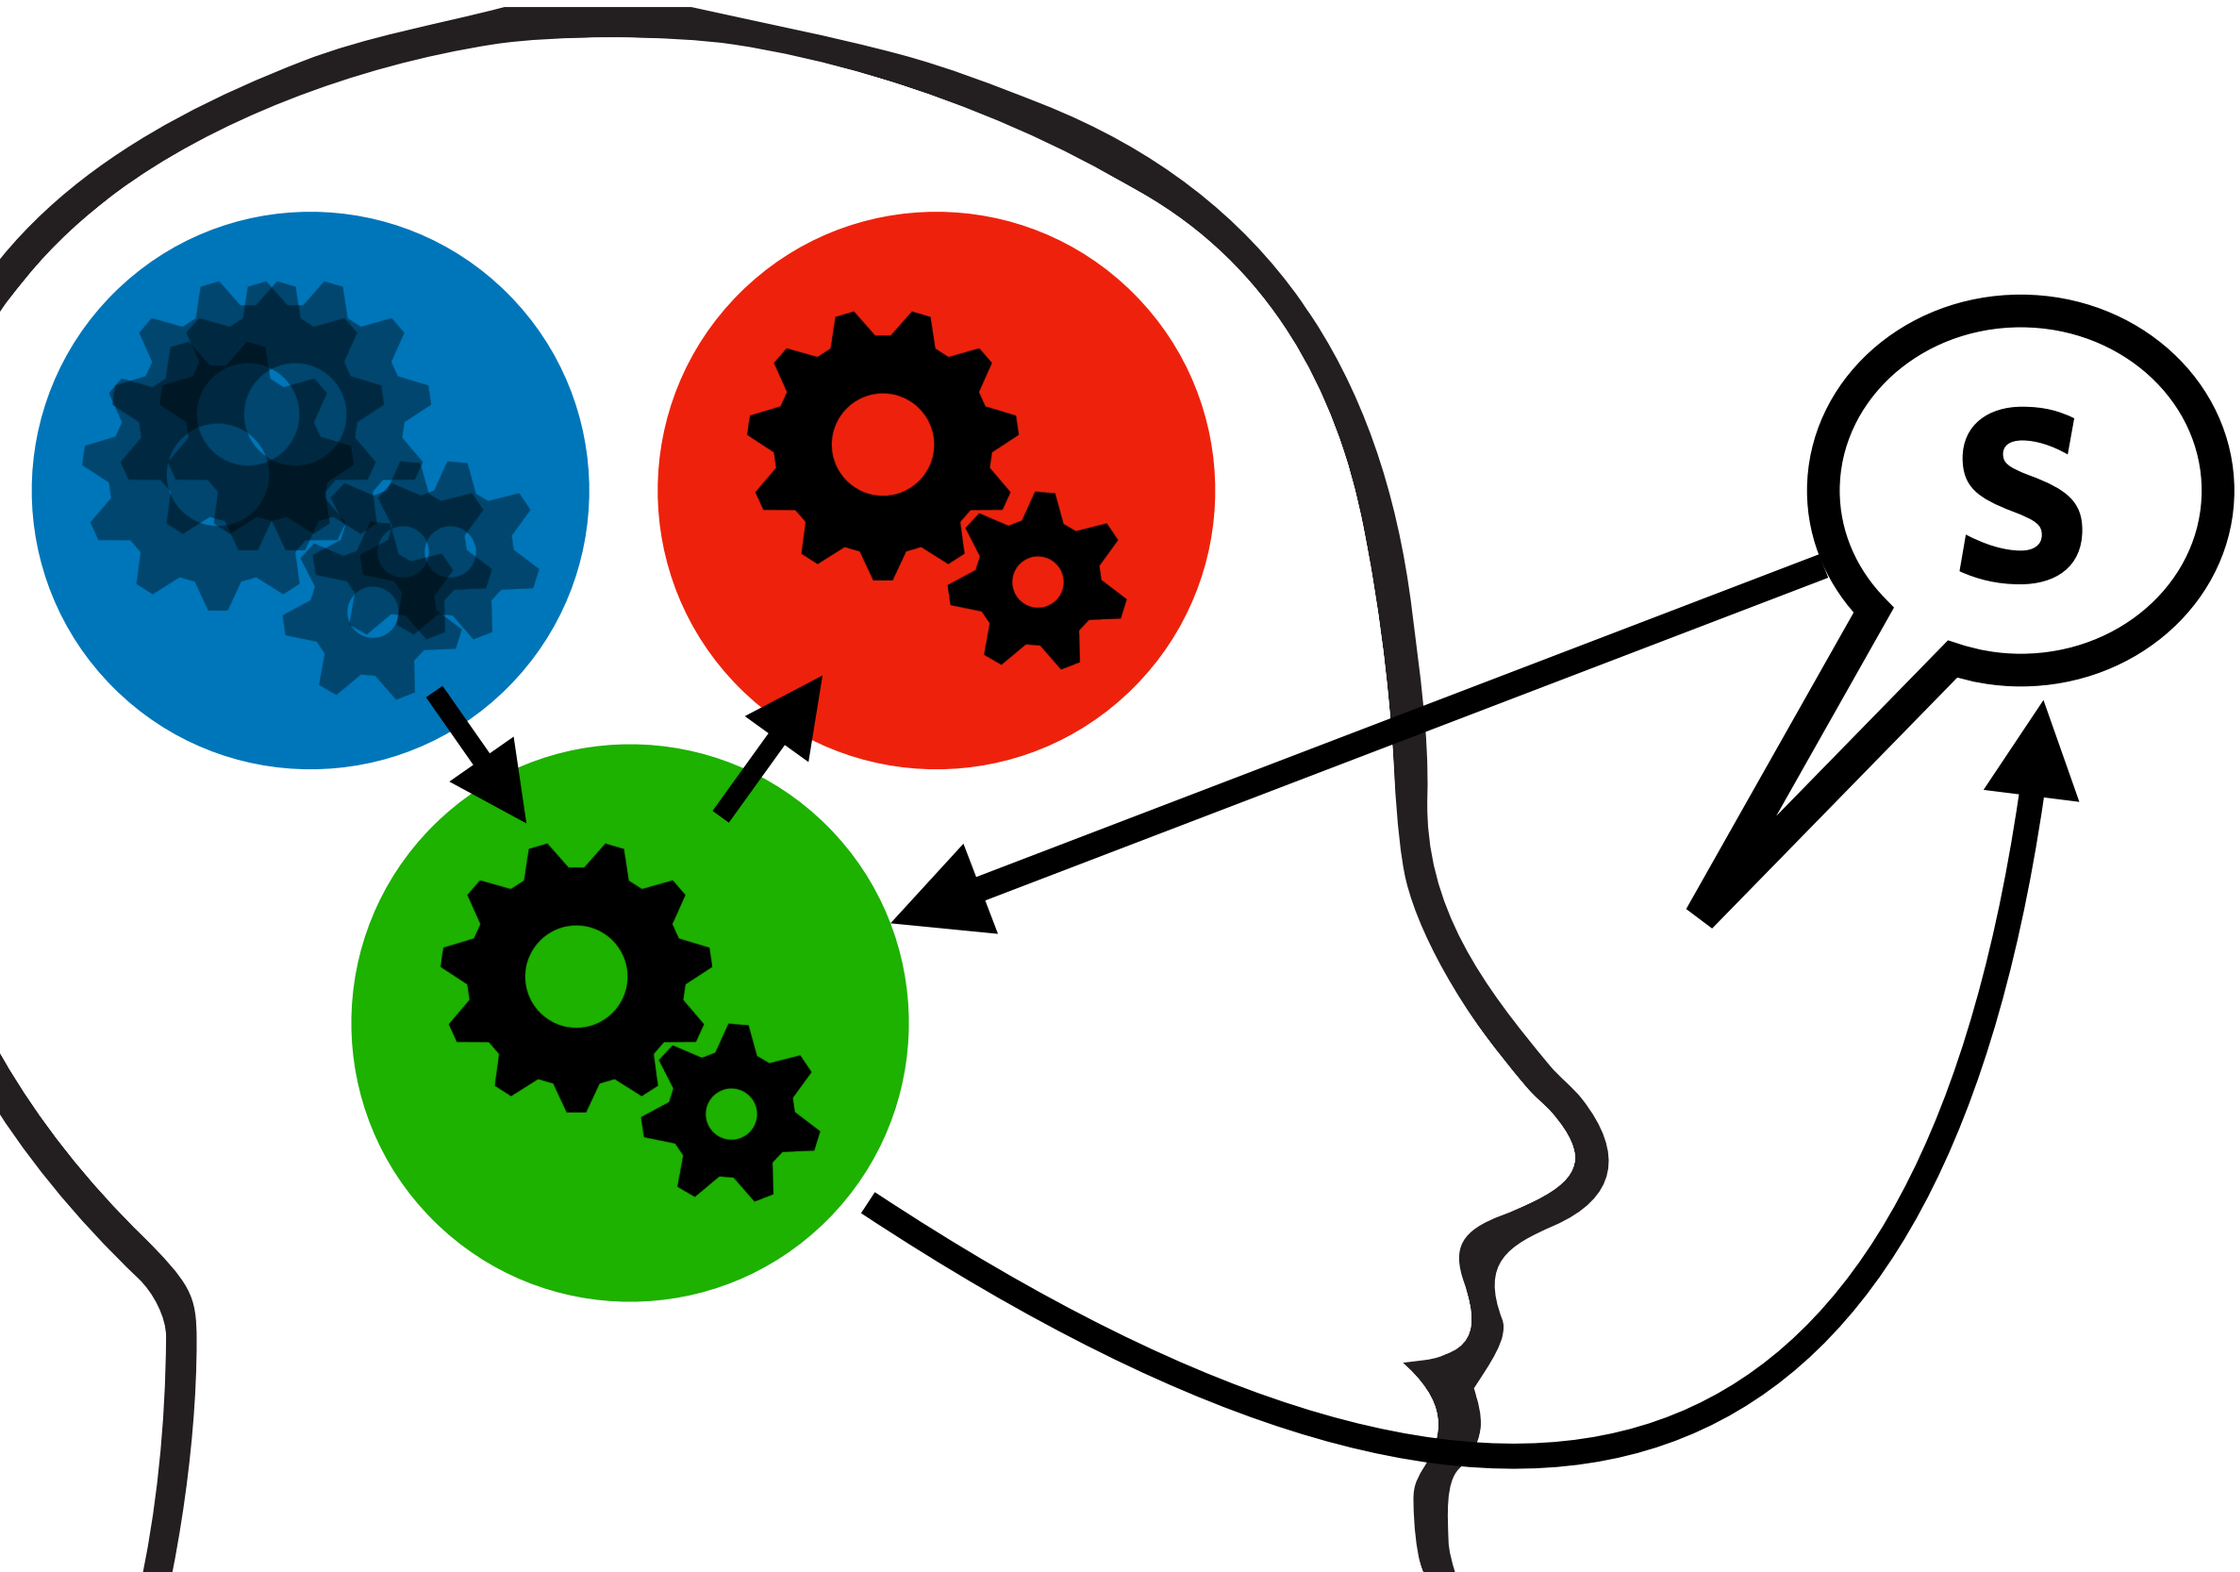

Clarifying Thoughts

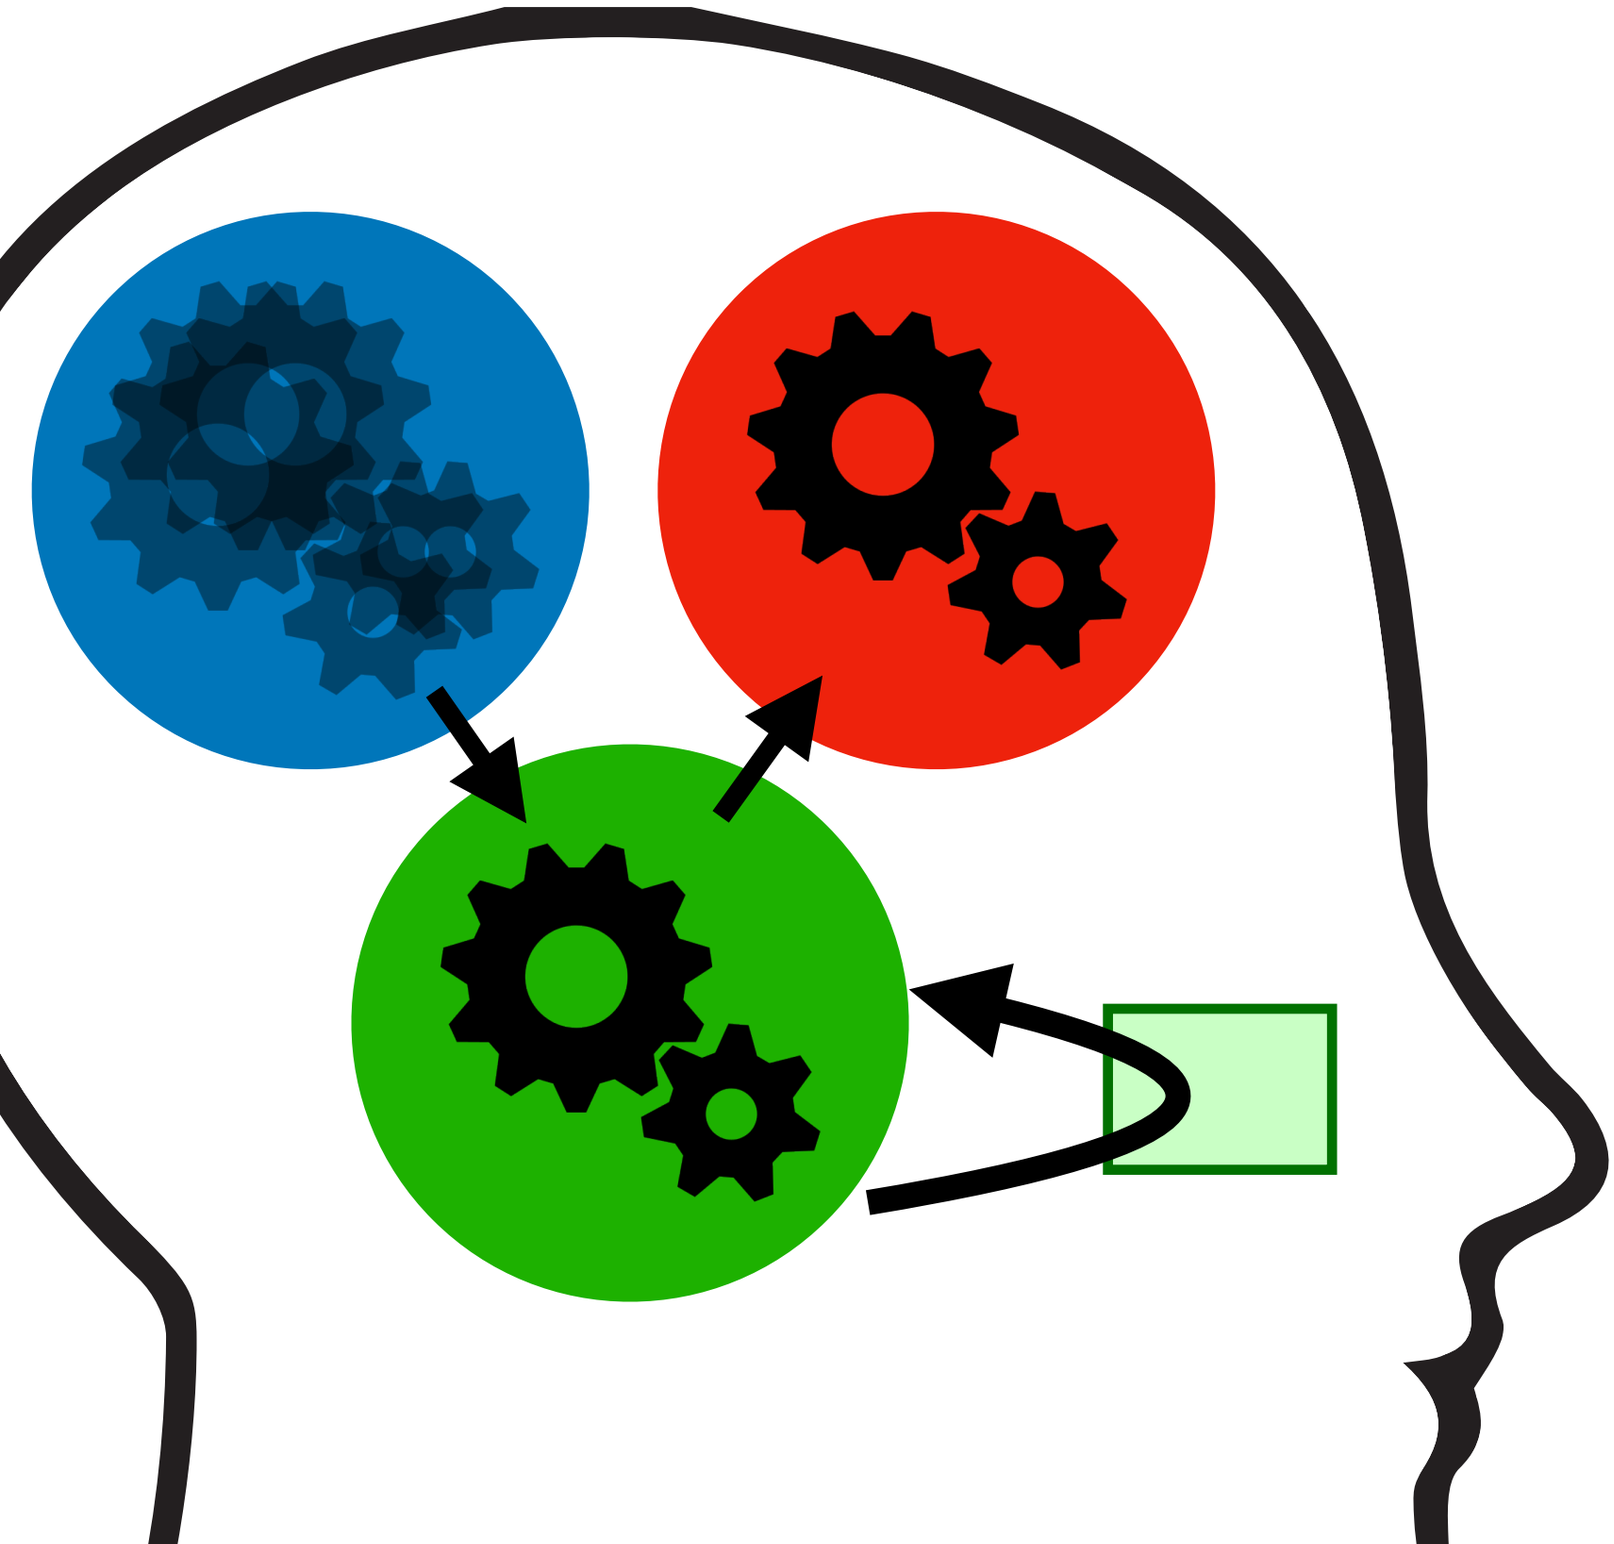

Tentative Conclusions

- In principle, we could use a modular language system for all kinds of things:
  - Short-term memory
  - Making Info available to other parts of the mind
  - Clarifying thoughts
  - Testing outputs before speaking
- There is at least some evidence that we actually do all of this.
- This view also holds some promise of explaining some documented forms of linguistic relativism.

Course Upshots

- We have good reasons to think that semantics is the study of a modular system.
- There are interesting and plausible strategies for handling some of the main challenges to this view:
  - Context sensitivity → thin meanings
  - Polysemy → division of labor with conceptual system
  - Top down control and thinking in language → working memory
- All of this is speculative!
- But hopefully it is also generative.