Representation in Linguistics

Tim O’Donnell, May 2022

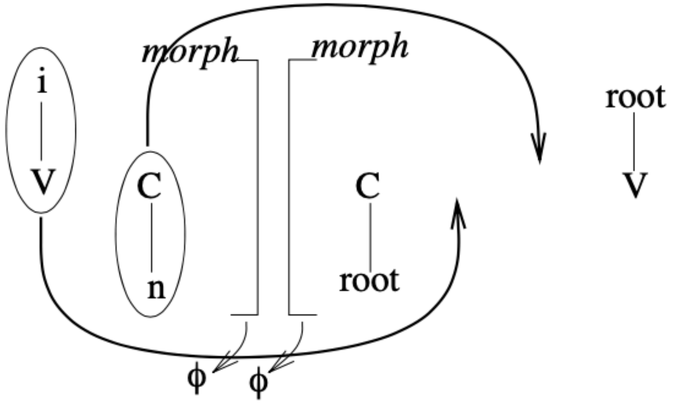
morph
morph
i
V
C
n
C
root
root
V
ϕ
ϕ

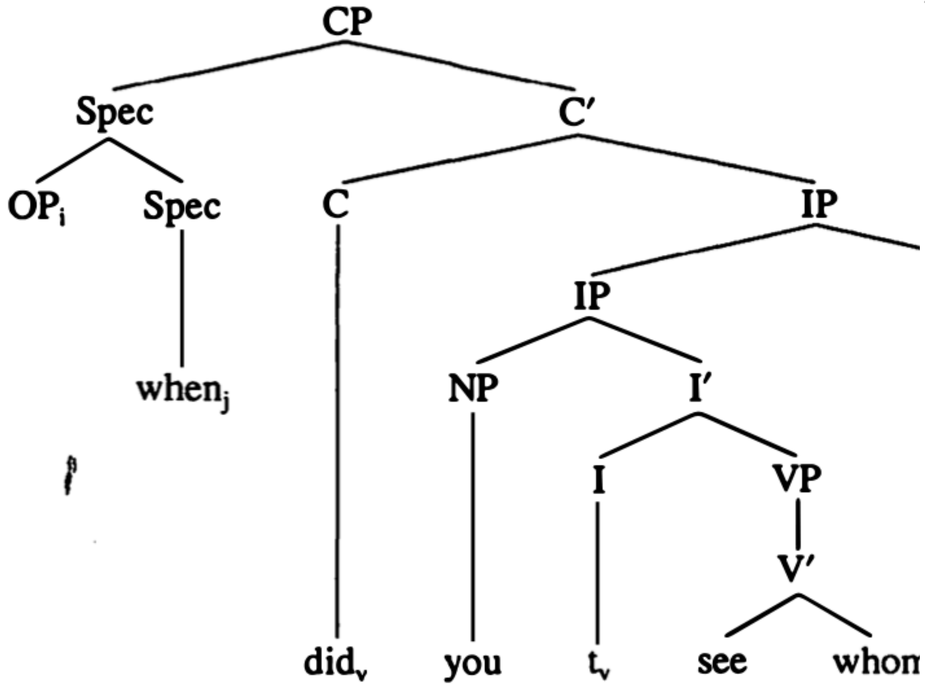
CP
Spec
C′
OP$_i$
Spec
C
IP
IP
when$_j$
NP
I′
I
VP
V′
did$_v$
you
t$_v$
see
whom

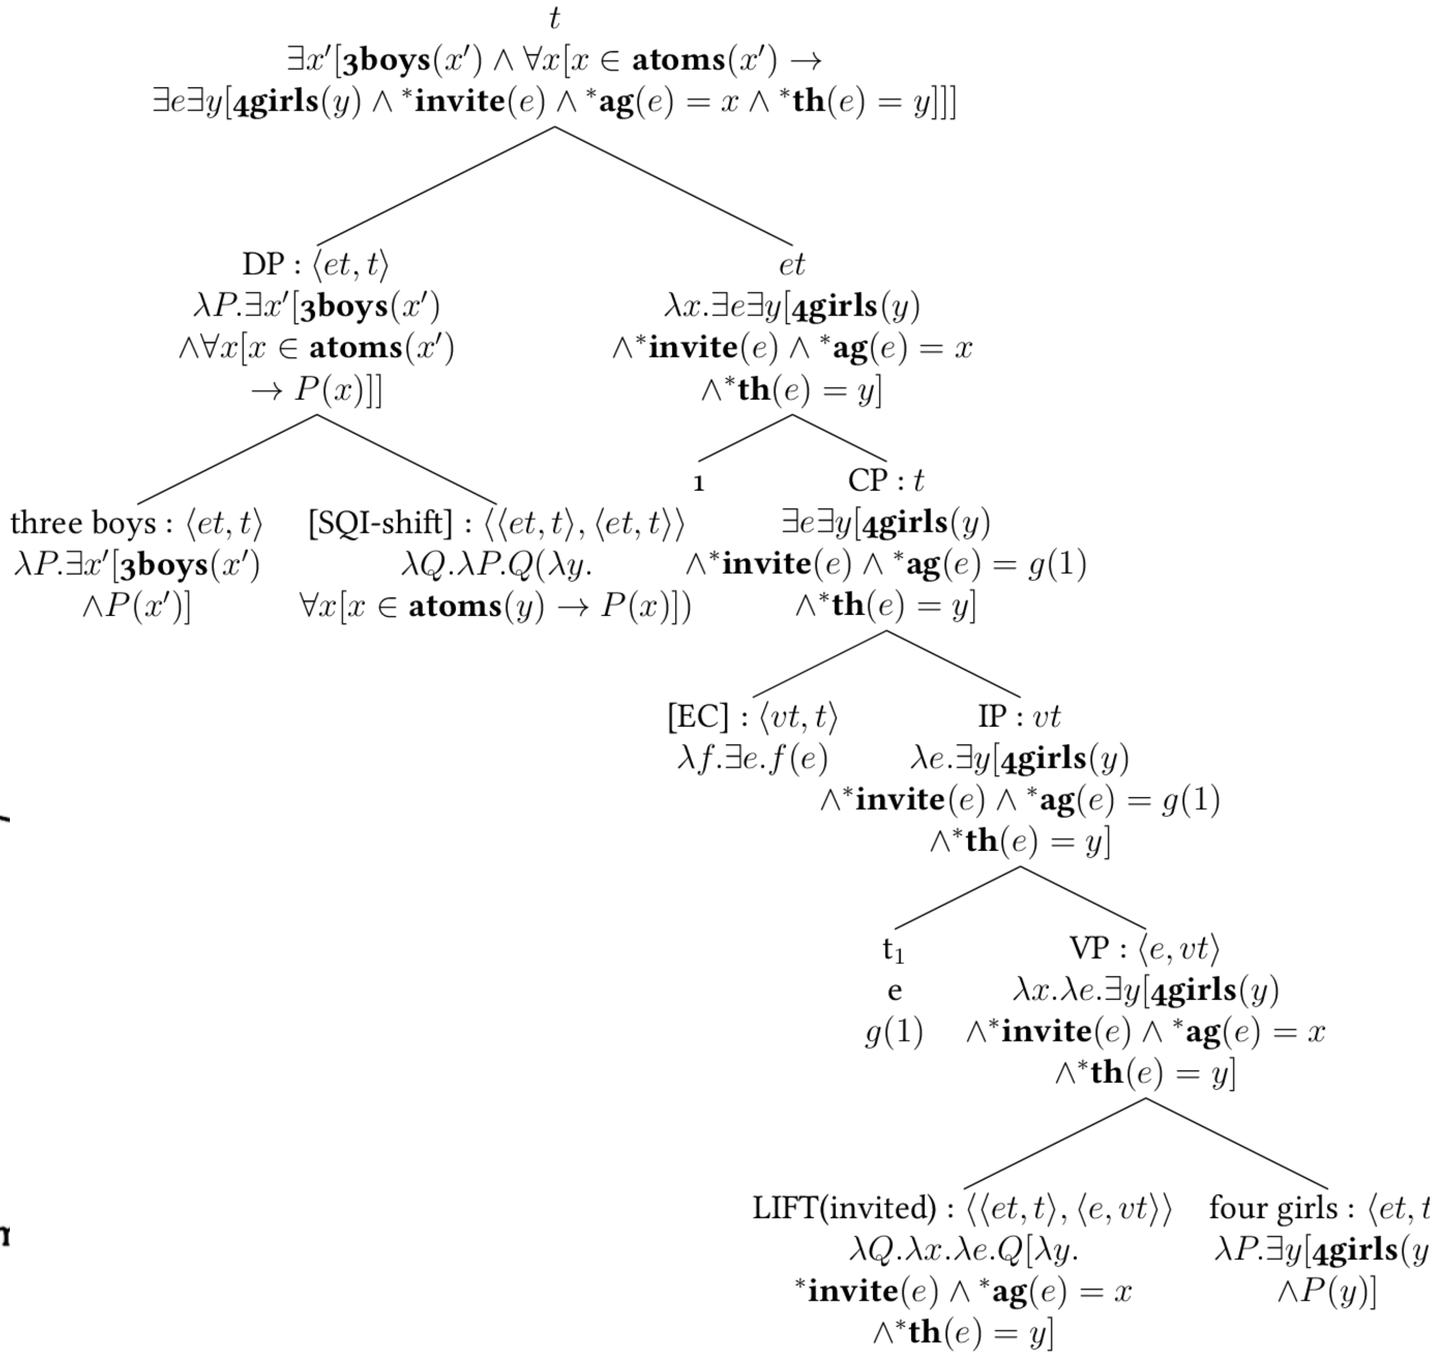
t
$\exists x'[\mathbf{3boys}(x') \wedge \forall x[x \in \mathbf{atoms}(x') \rightarrow \exists e \exists y[\mathbf{4girls}(y) \wedge {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = x \wedge {}^*\mathbf{th}(e) = y]]]$
DP : $\langle et, t \rangle$
$\lambda P.\exists x'[\mathbf{3boys}(x') \wedge \forall x[x \in \mathbf{atoms}(x') \rightarrow P(x)]]$
et
$\lambda x.\exists e \exists y[\mathbf{4girls}(y) \wedge {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = x \wedge {}^*\mathbf{th}(e) = y]$
three boys : $\langle et, t \rangle$
$\lambda P.\exists x'[\mathbf{3boys}(x') \wedge P(x')]$
[SQI-shift] : $\langle \langle et, t \rangle, \langle et, t \rangle \rangle$
$\lambda Q.\lambda P.Q(\lambda y. \forall x[x \in \mathbf{atoms}(y) \rightarrow P(x)])$
1
CP : t
$\exists e \exists y[\mathbf{4girls}(y) \wedge {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = g(1) \wedge {}^*\mathbf{th}(e) = y]$
[EC] : $\langle vt, t \rangle$
$\lambda f.\exists e.f(e)$
IP : vt
$\lambda e.\exists y[\mathbf{4girls}(y) \wedge {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = g(1) \wedge {}^*\mathbf{th}(e) = y]$
t$_1$
e
$g(1)$
VP : $\langle e, vt \rangle$
$\lambda x.\lambda e.\exists y[\mathbf{4girls}(y) \wedge {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = x \wedge {}^*\mathbf{th}(e) = y]$
LIFT(invited) : $\langle \langle et, t \rangle, \langle e, vt \rangle \rangle$
$\lambda Q.\lambda x.\lambda e.Q[\lambda y. {}^*\mathbf{invite}(e) \wedge {}^*\mathbf{ag}(e) = x \wedge {}^*\mathbf{th}(e) = y]$
four girls : $\langle et, t$
$\lambda P.\exists y[\mathbf{4girls}(y$
$\wedge P(y)]$

What are representations?

- Just “data structures” expressed in some formalism, notation, etc.
- We don’t have to think about this very much since we don’t ever worry about how these “representations” are “represented” anywhere else (e.g., in brains).
- Some covert commitments:
  - Speaker-illegibility (weak)
  - Theoretician-legibility (strong)
  - Interpretability: Systematic relationships (mappings) with other systems of the mind (or the outside world).

What are representations?

- Just “data structures” expressed in some formalism, notation, etc.
- We don’t have to think about this very much since we don’t ever worry about how these “representations” are “represented” anywhere else (e.g., in brains).
- Some covert commitments:
  - Speaker-illegibility (weak)
  - Theoretician-legibility (strong)
  - **Interpretability: Systematic relationships (mappings) with other systems of the mind (or the outside world).**

The Y-Model

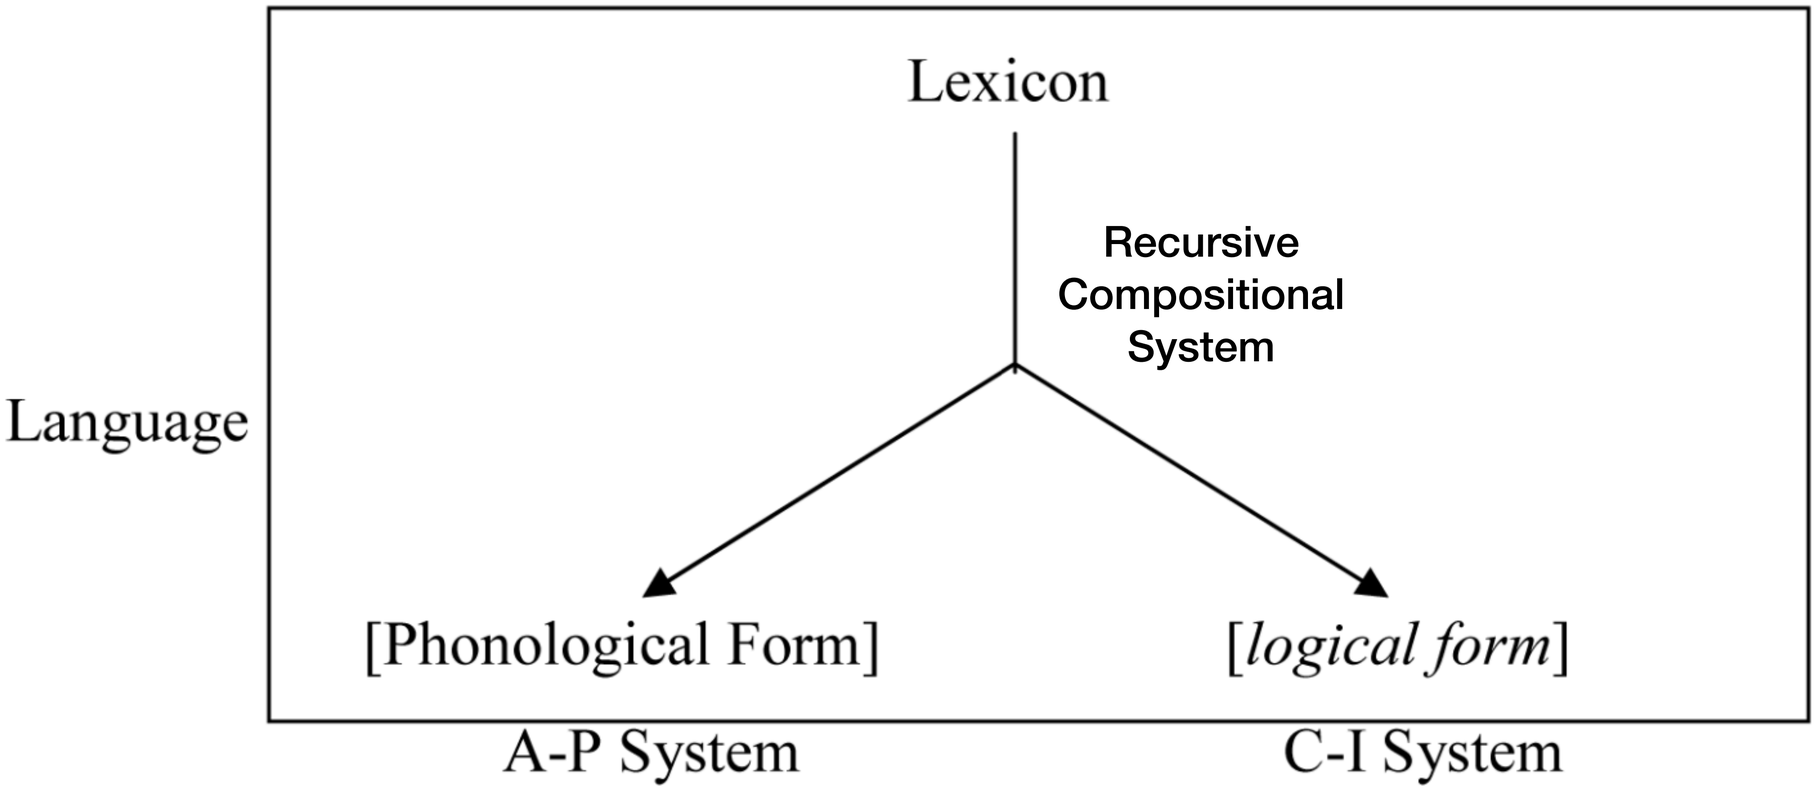
Language
Lexicon
Recursive
Compositional
System
[Phonological Form]
[*logical form*]
A-P System
C-I System

The Y-Model

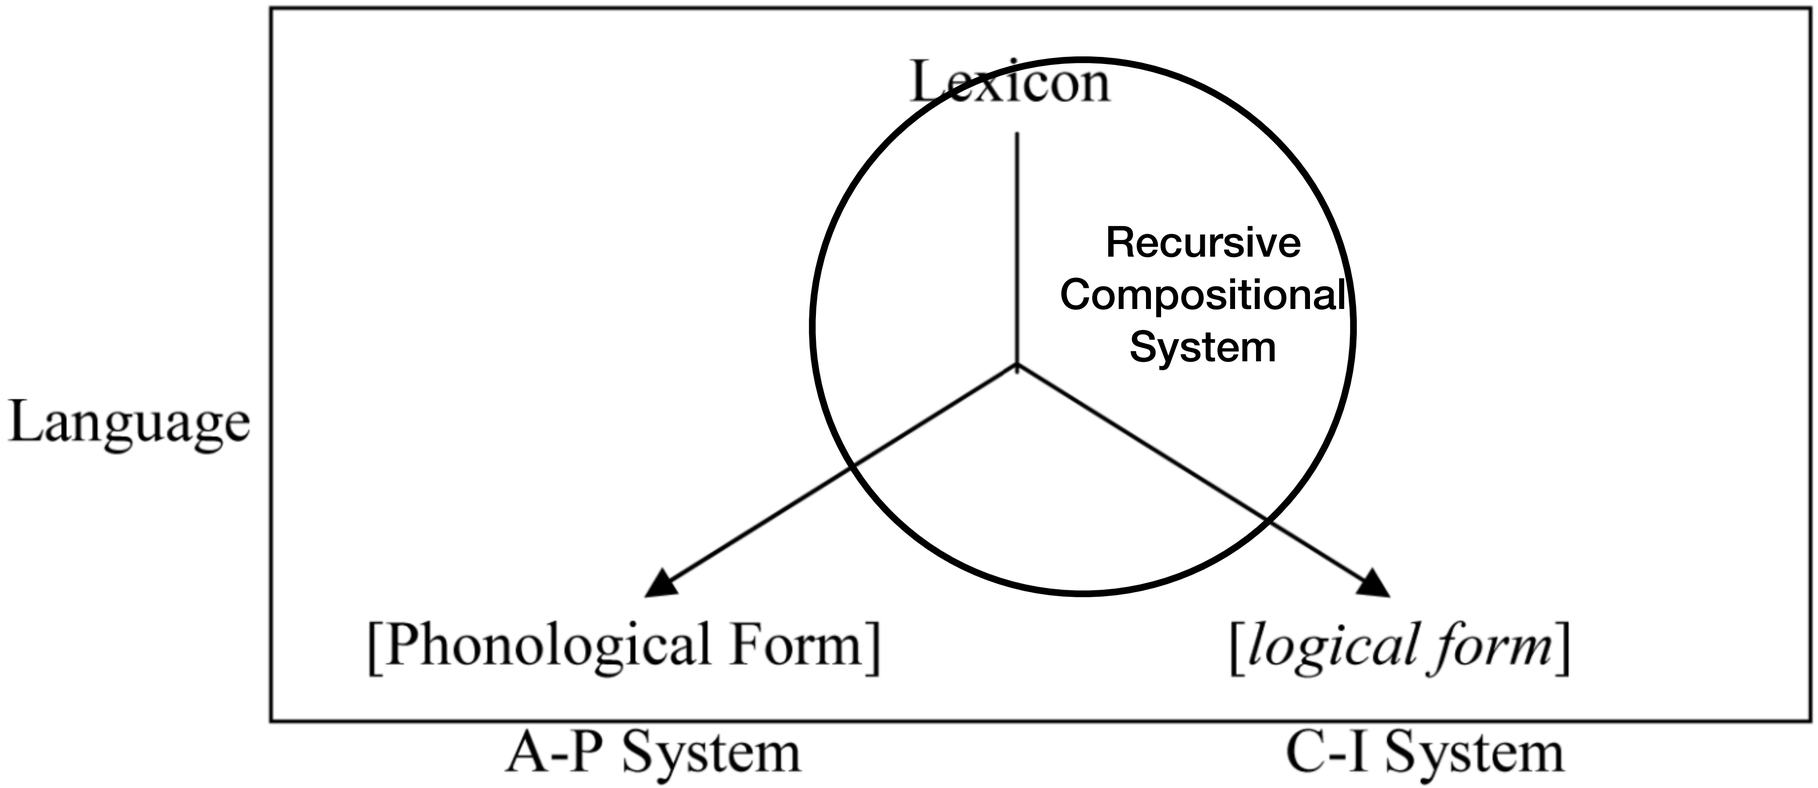
Language
Lexicon
Recursive
Compositional
System
[Phonological Form]
[*logical form*]
A-P System
C-I System

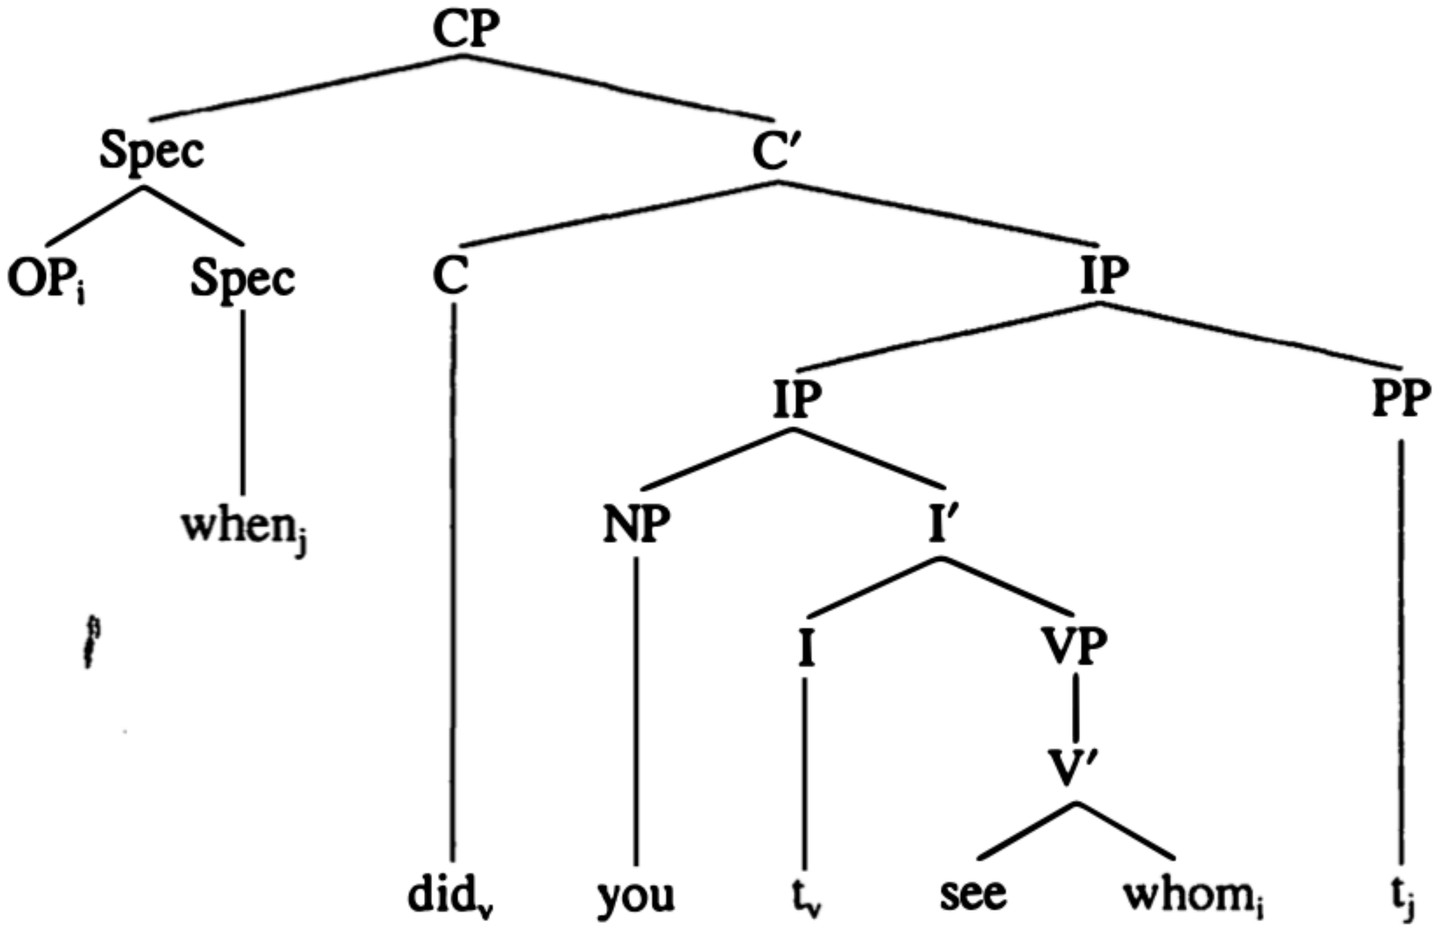
CP
Spec
C′
OP$_i$
Spec
C
IP
IP
PP
when$_j$
NP
I′
I
VP
V′
did$_v$
you
t$_v$
see
whom$_i$
t$_j$

The Y-Model

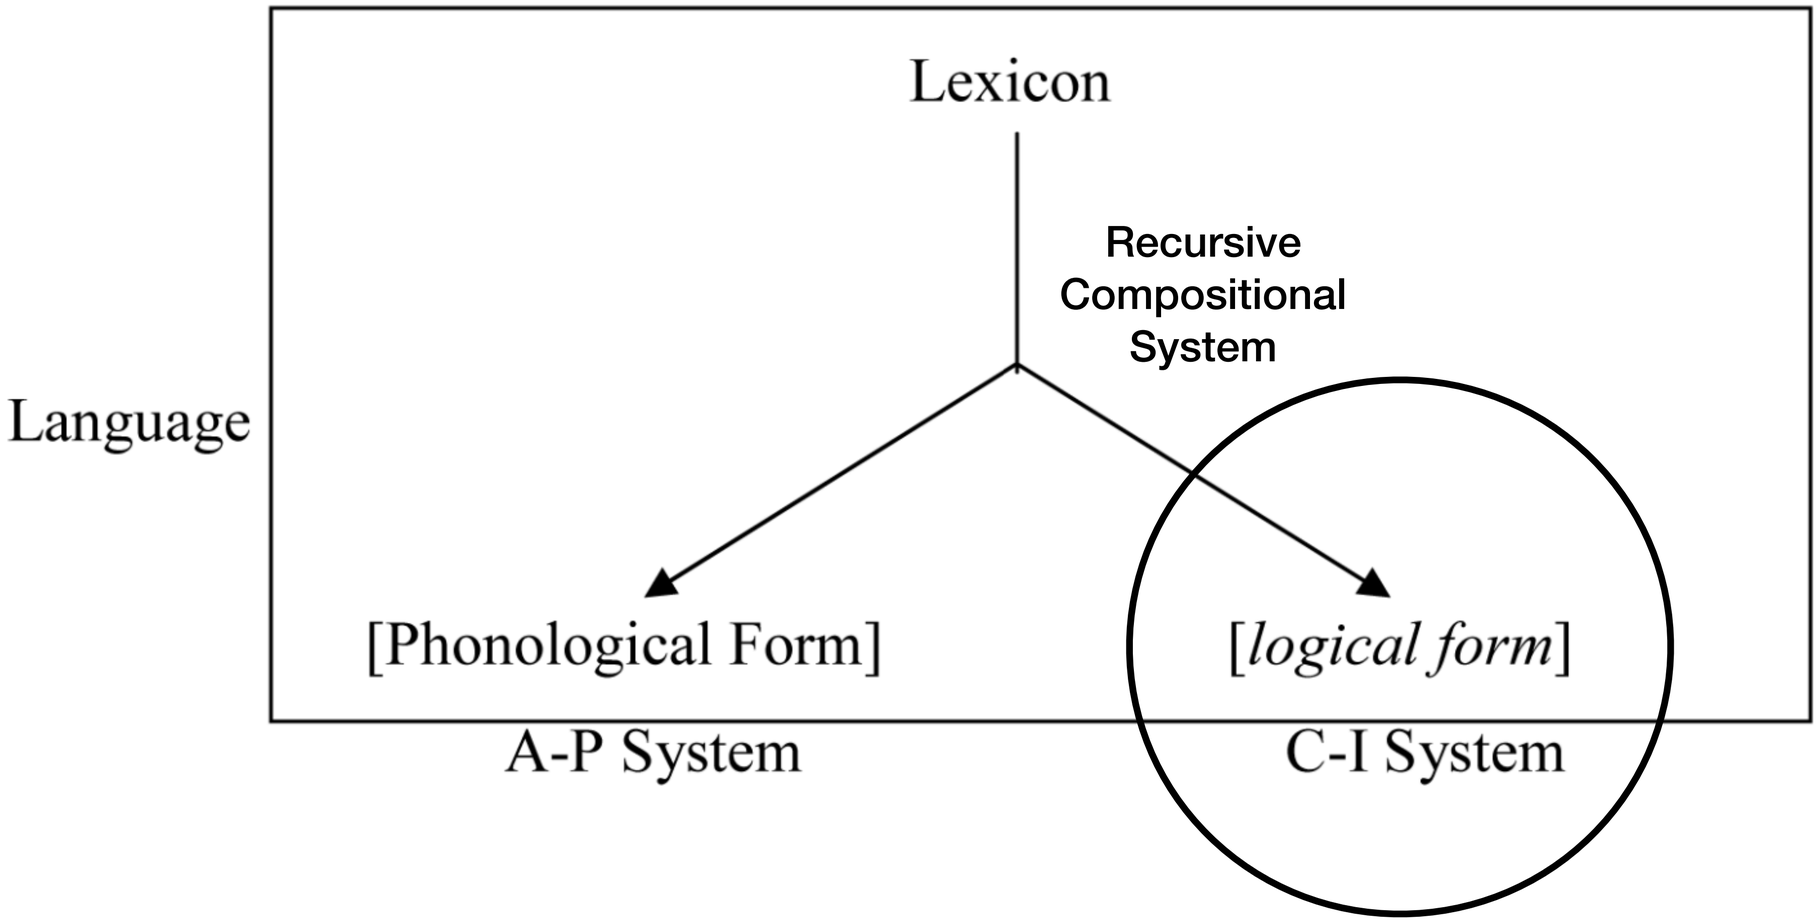
Lexicon
Recursive
Compositional
System
Language
[Phonological Form]
[logical form]
A-P System
C-I System

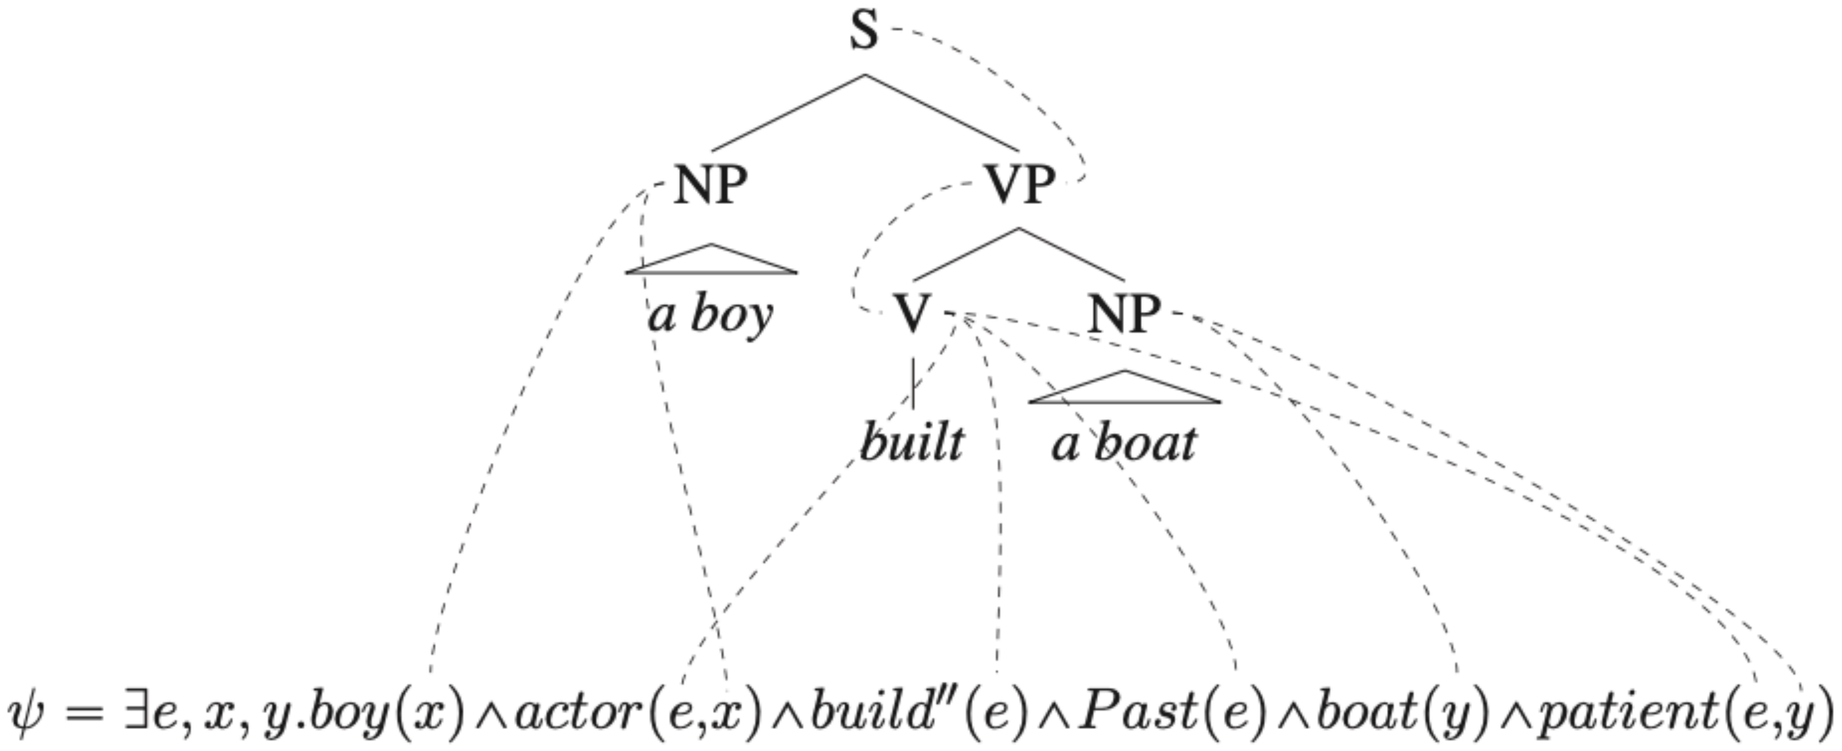
S
NP
VP
a boy
V
NP
built
a boat
$\psi = \exists e, x, y.boy(x) \wedge actor(e,x) \wedge build''(e) \wedge Past(e) \wedge boat(y) \wedge patient(e,y)$

Compositionality

- The meaning of a sentence is some function of the meaning of its parts and the way those parts are combined.
- We imagine that the problem is to map from some space of forms U to some space of meanings M, call the map $[\![\,]\!] : \mathrm{U} \rightarrow \mathrm{M}$
- Richard Montague (founder of formal semantics) suggested that compositionality in a mapping from a space (algebra) of forms to a space (algebra) meant that $[\![\,]\!]$ is a homomorphism (see Janssen, 1986).
  - $[\![F(x,y,z)]\!] = [\![F]\!]([\![x]\!], [\![y]\!], [\![z]\!])$.

A World Model (3DP3, Gothoskar et al 2021)

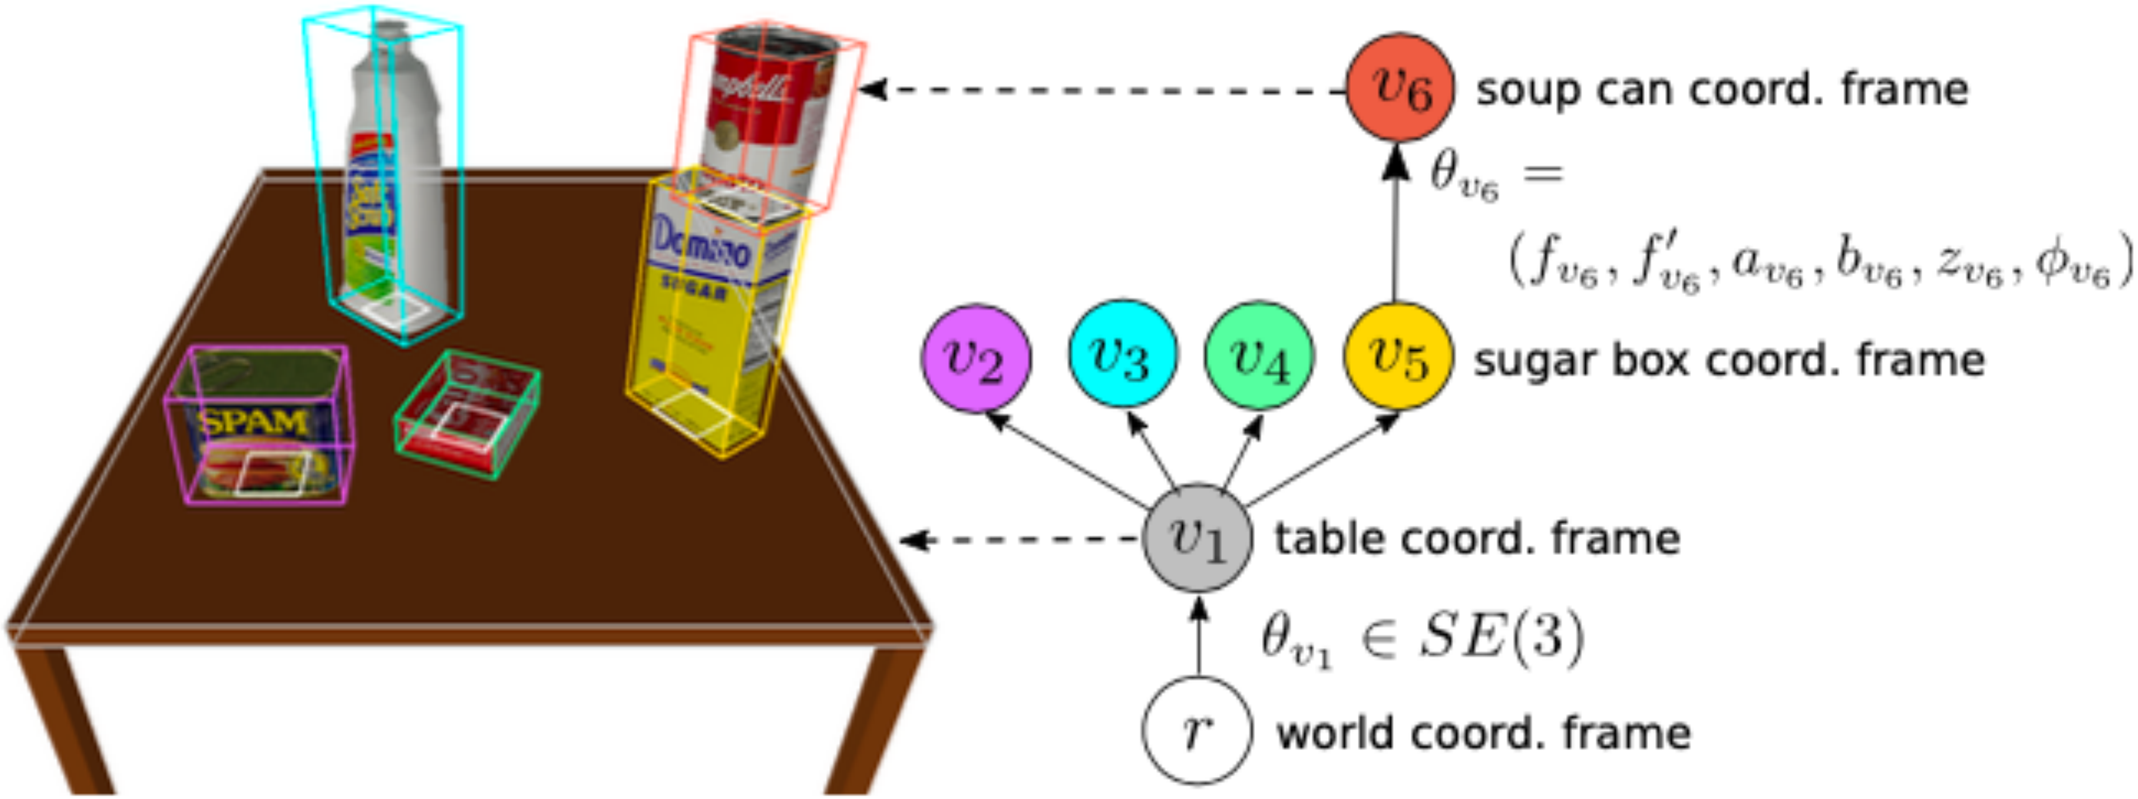

Scene

“find a can behind a box”

find

Sidebar: Views

- Active versus passive views of such structures
  - Derivation structures are instantiated in memory
  - Derivation structures are merely our conventions for representing the trace of a computational process
- “Derivational” versus “**Representational**” theories.
  - Step-by-step, algorithmic procedures for building structures, versus constraints on sets of structures stated in a meta-language
  - Proof-theoretic versus model-theoretic
  - Directed (generative) versus undirected (random field)
- Mentalist versus realist notions of the mappings

Interpretability as Compositionality

- The meaning of a sentence is some function of the meaning of its parts and the way those parts are combined.
- We imagine that the problem is to map from some space of forms U to some space of meanings M, call the map $[\![\,]\!] : \mathrm{U} \rightarrow \mathrm{M}$
- Richard Montague (founder of formal semantics) suggested that compositionality in a mapping from a space (algebra) of forms to a space (algebra) meant that $[\![\,]\!]$ is a homomorphism (see Janssen, 1986).
  - $[\![F(x,y,z)]\!] = [\![F]\!]([\![x]\!], [\![y]\!], [\![z]\!])$.
- This principle is too strict even for relatively simple examples in natural language…. **idioms**

Idioms

A “grammatical leak”

Tim rocked the cradle.

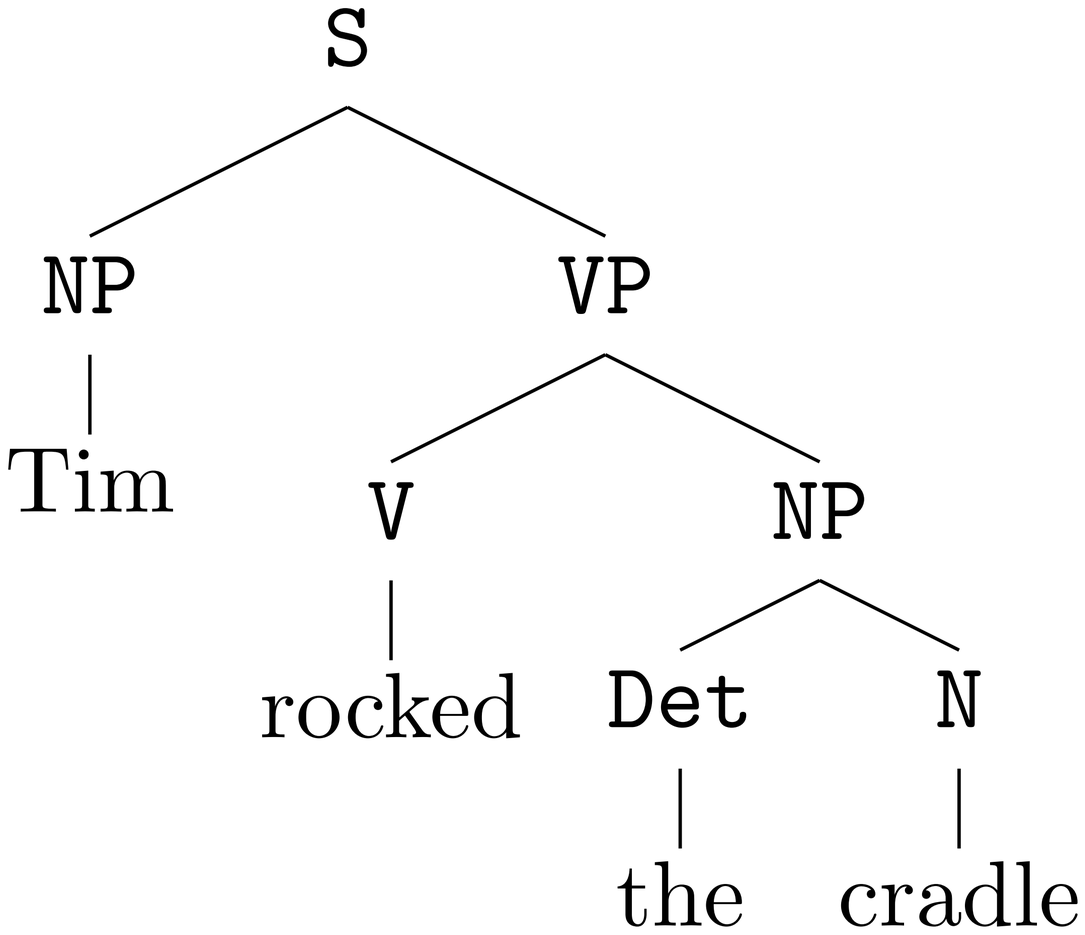

Idioms

A “grammatical leak”

Tim rocked the boat.

“Tim destabilized the situation.”

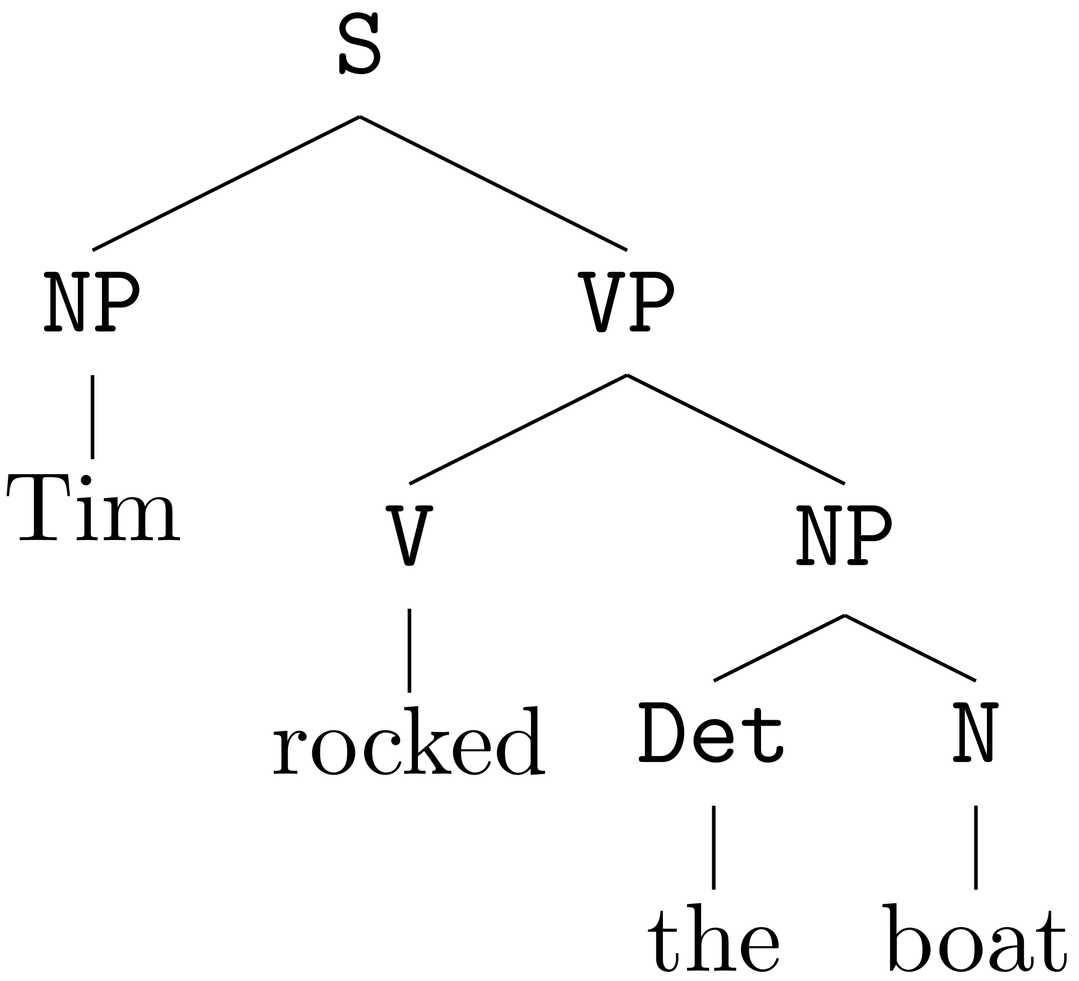

Idioms

A “grammatical leak”

Tim rocked the boat.

“Tim destabilized the situation.”

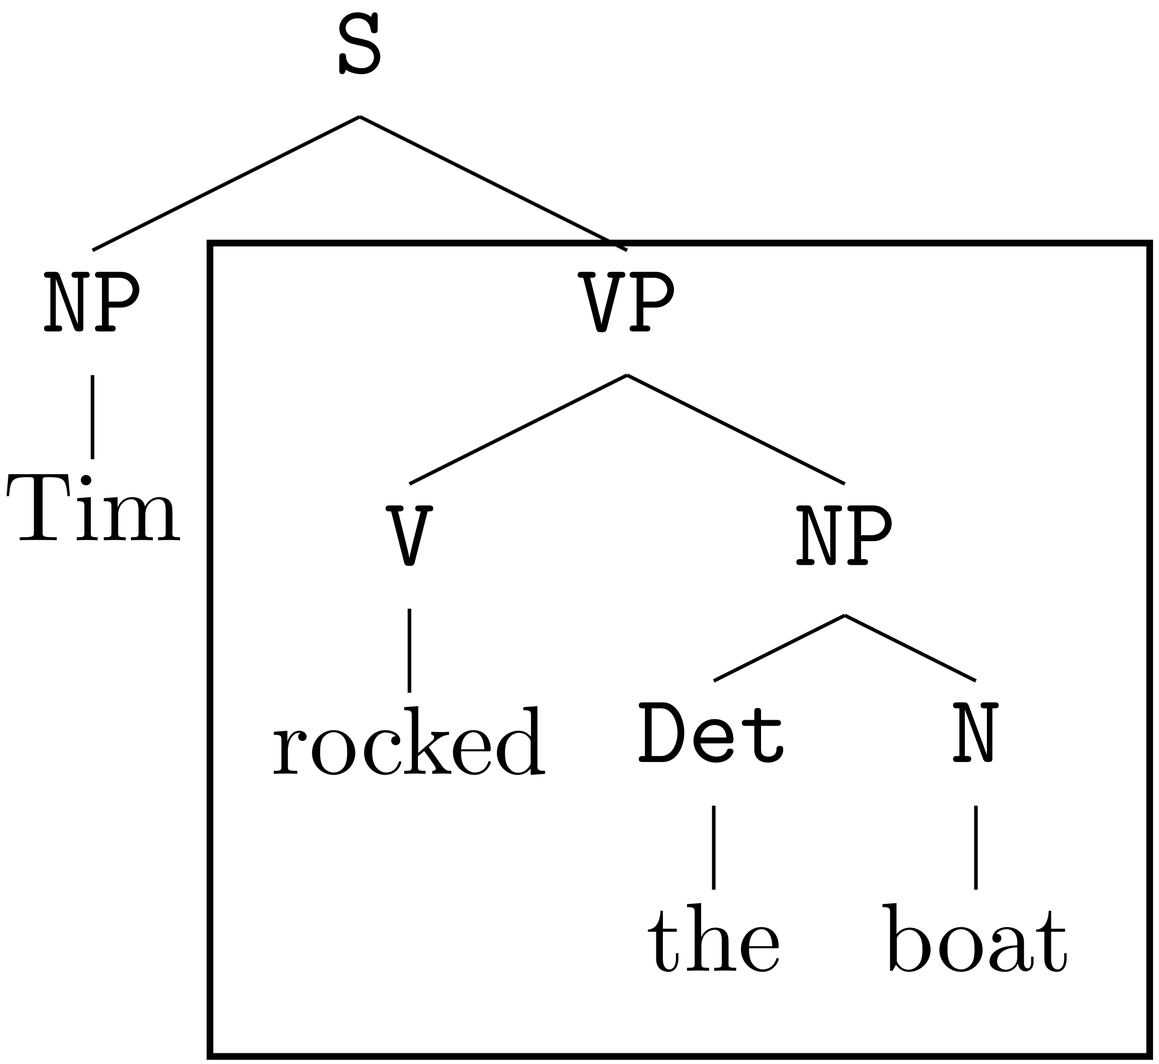

Idioms

A “grammatical leak”

Tim rocked the boat.

“Tim destabilized the situation.”

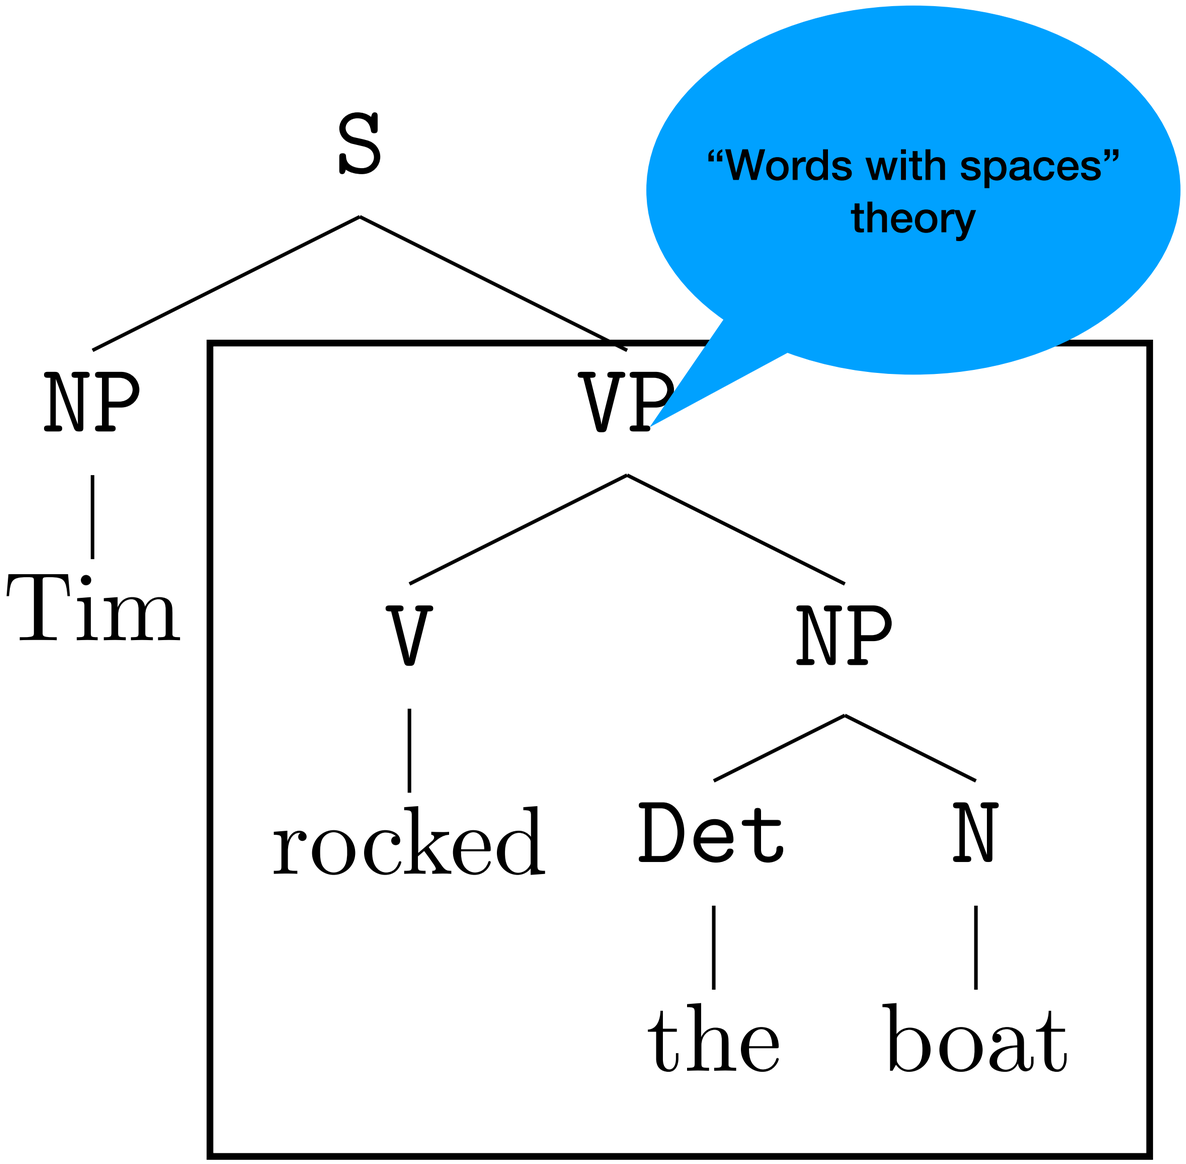

Idioms

A “grammatical leak”

- Two problems:

1. Individual words can be modified in many idioms

   Tim managed to rock an otherwise unrockable boat.

2. Nevertheless, the idiomatic meaning depends on the combination of words.

   Tim shook the canoe.

   Tim managed to shake an otherwise unshakeable canoe.

- Conclusion:
  - Individual words contribute some aspect of meaning.
  - But their combination also synergistically contributes some aspect of meaning.

Compositionality

- Different units and combinations of units seems to contribute different amounts of information about meaning: the is a contribution from *rock*, from *boat*, and **from the combination** *rock the boat*.
- The higher the proportion of information that comes from smaller units such as words, the more intuitively compositional the system is. *Rock the cradle* seems to mostly include information from *rock* and *cradle*.
- Maybe compositionality should be viewed not as a binary property of a system, but as a quantitative property:
  - **How locally** is meaning coded (on words, or on bigger units)?
- If we had such a definition, maybe we could put it into our objective functions…

A Simple Representional Model of Meaning

The classic Montague model

- Montague introduced the idea that meaning could be captured using model theory.
- Meaning spaces are just (structured) sets of situations (cf. possible worlds).
- In our case, the meaning spaces are just target interface systems of the mind.
- The meaning of an utterance is just a subset of possible situations in these interface spaces.
- Another perspective: an utterance is a probabilistic **conditioner** that chooses a particular subset of possible states of interfaced systems.
- Very week notion of representation in terms of probabilistic conditioning.

A Simple Model of Meaning

The classic Montague model

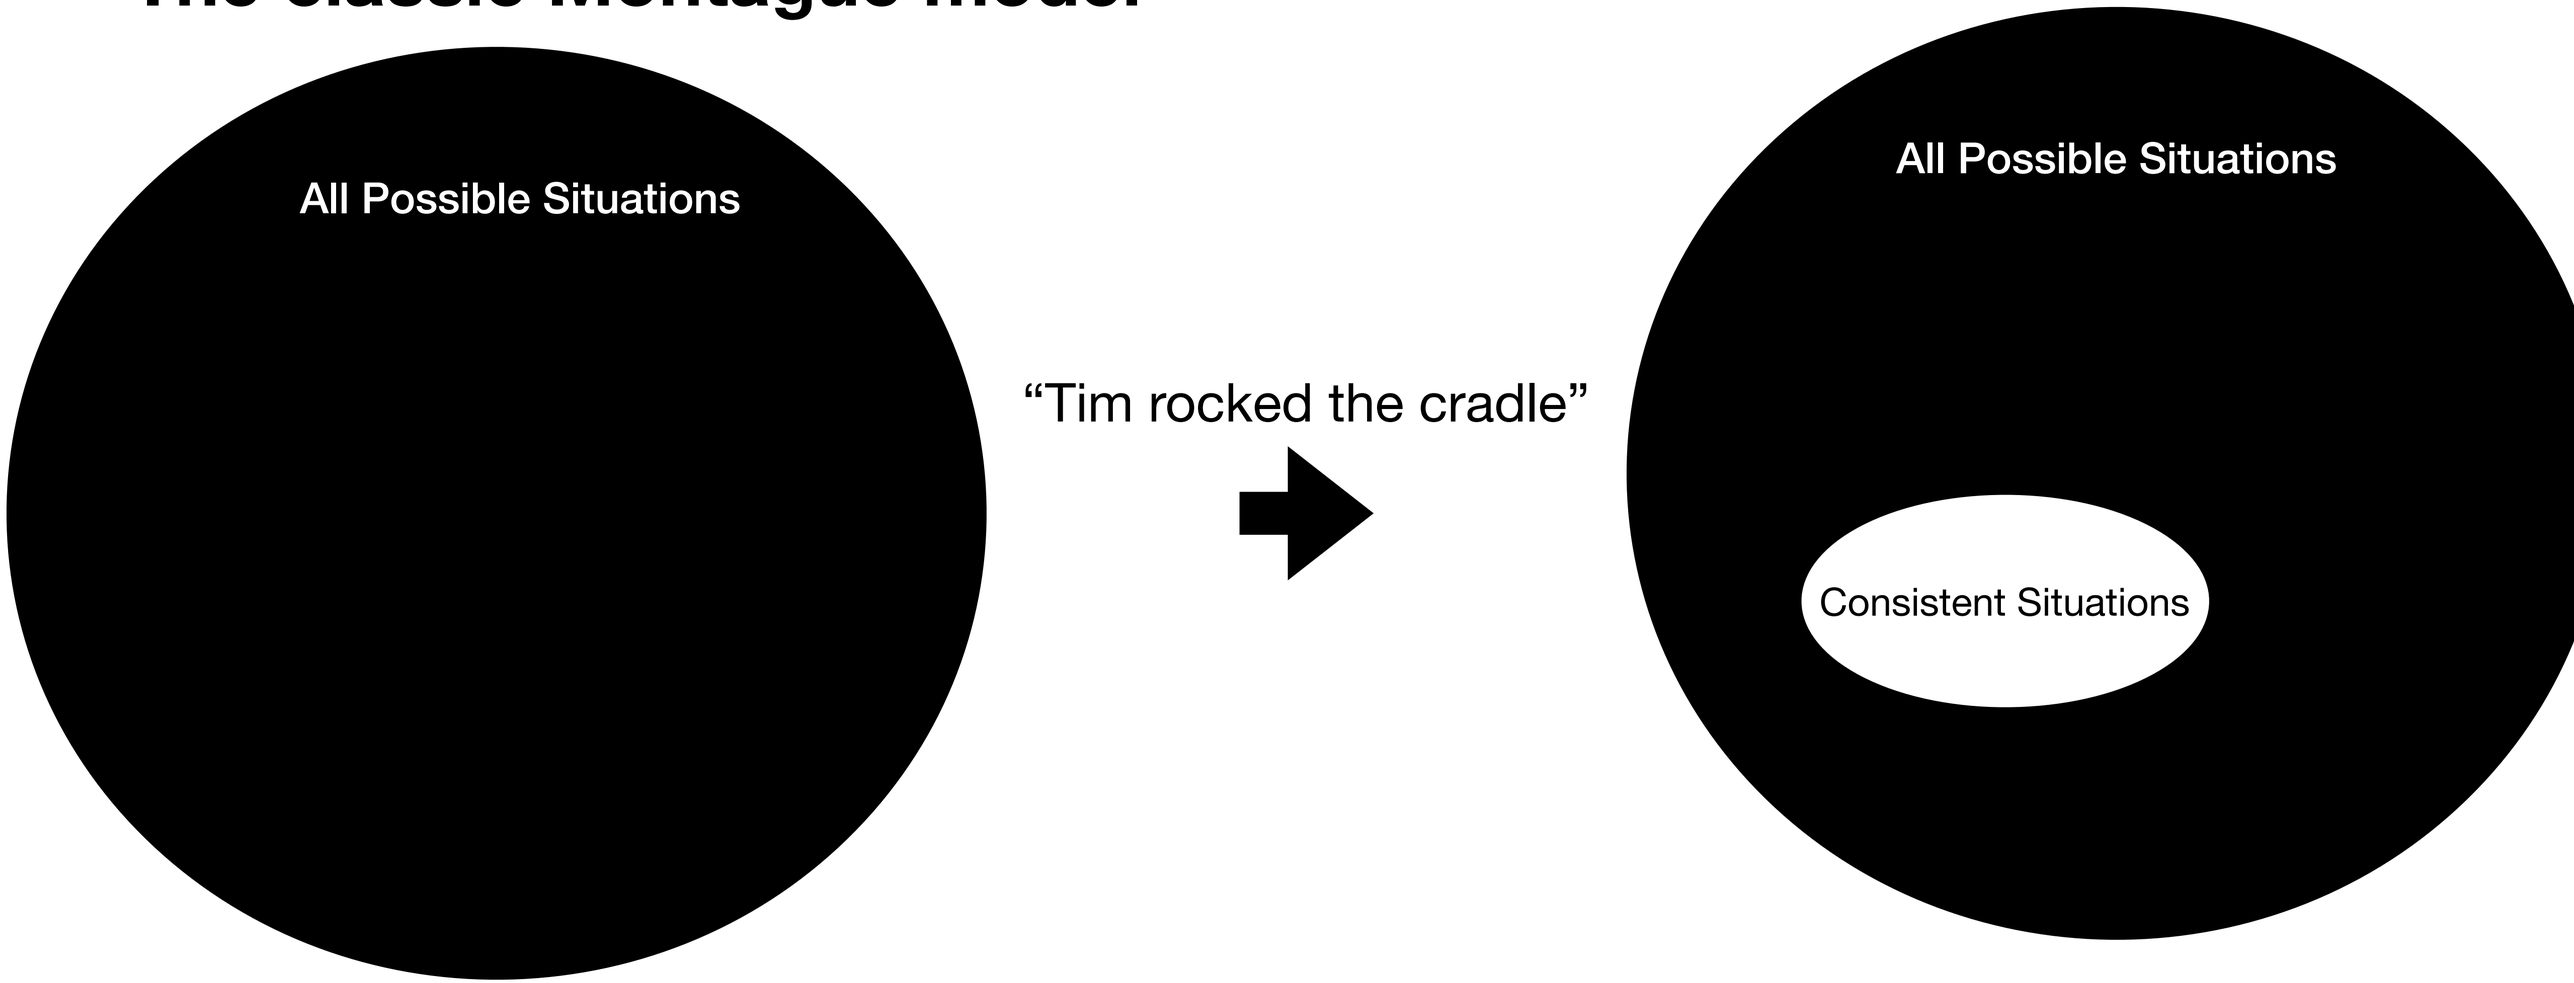

A Simple Model of Meaning

The classic Montague model

“Tim rocked the cradle”

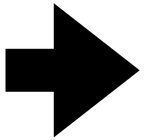

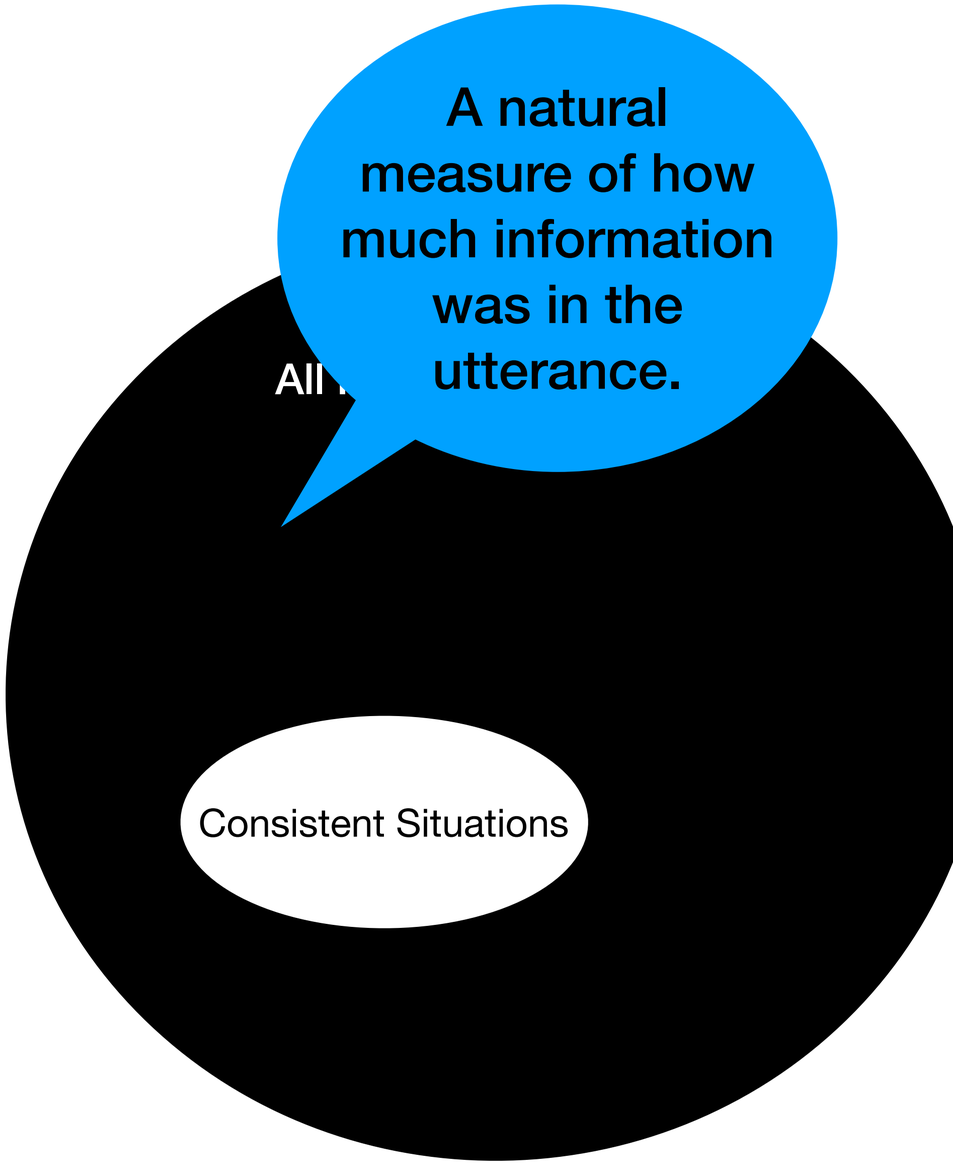

A Simple Model of Meaning

A probabilistic version

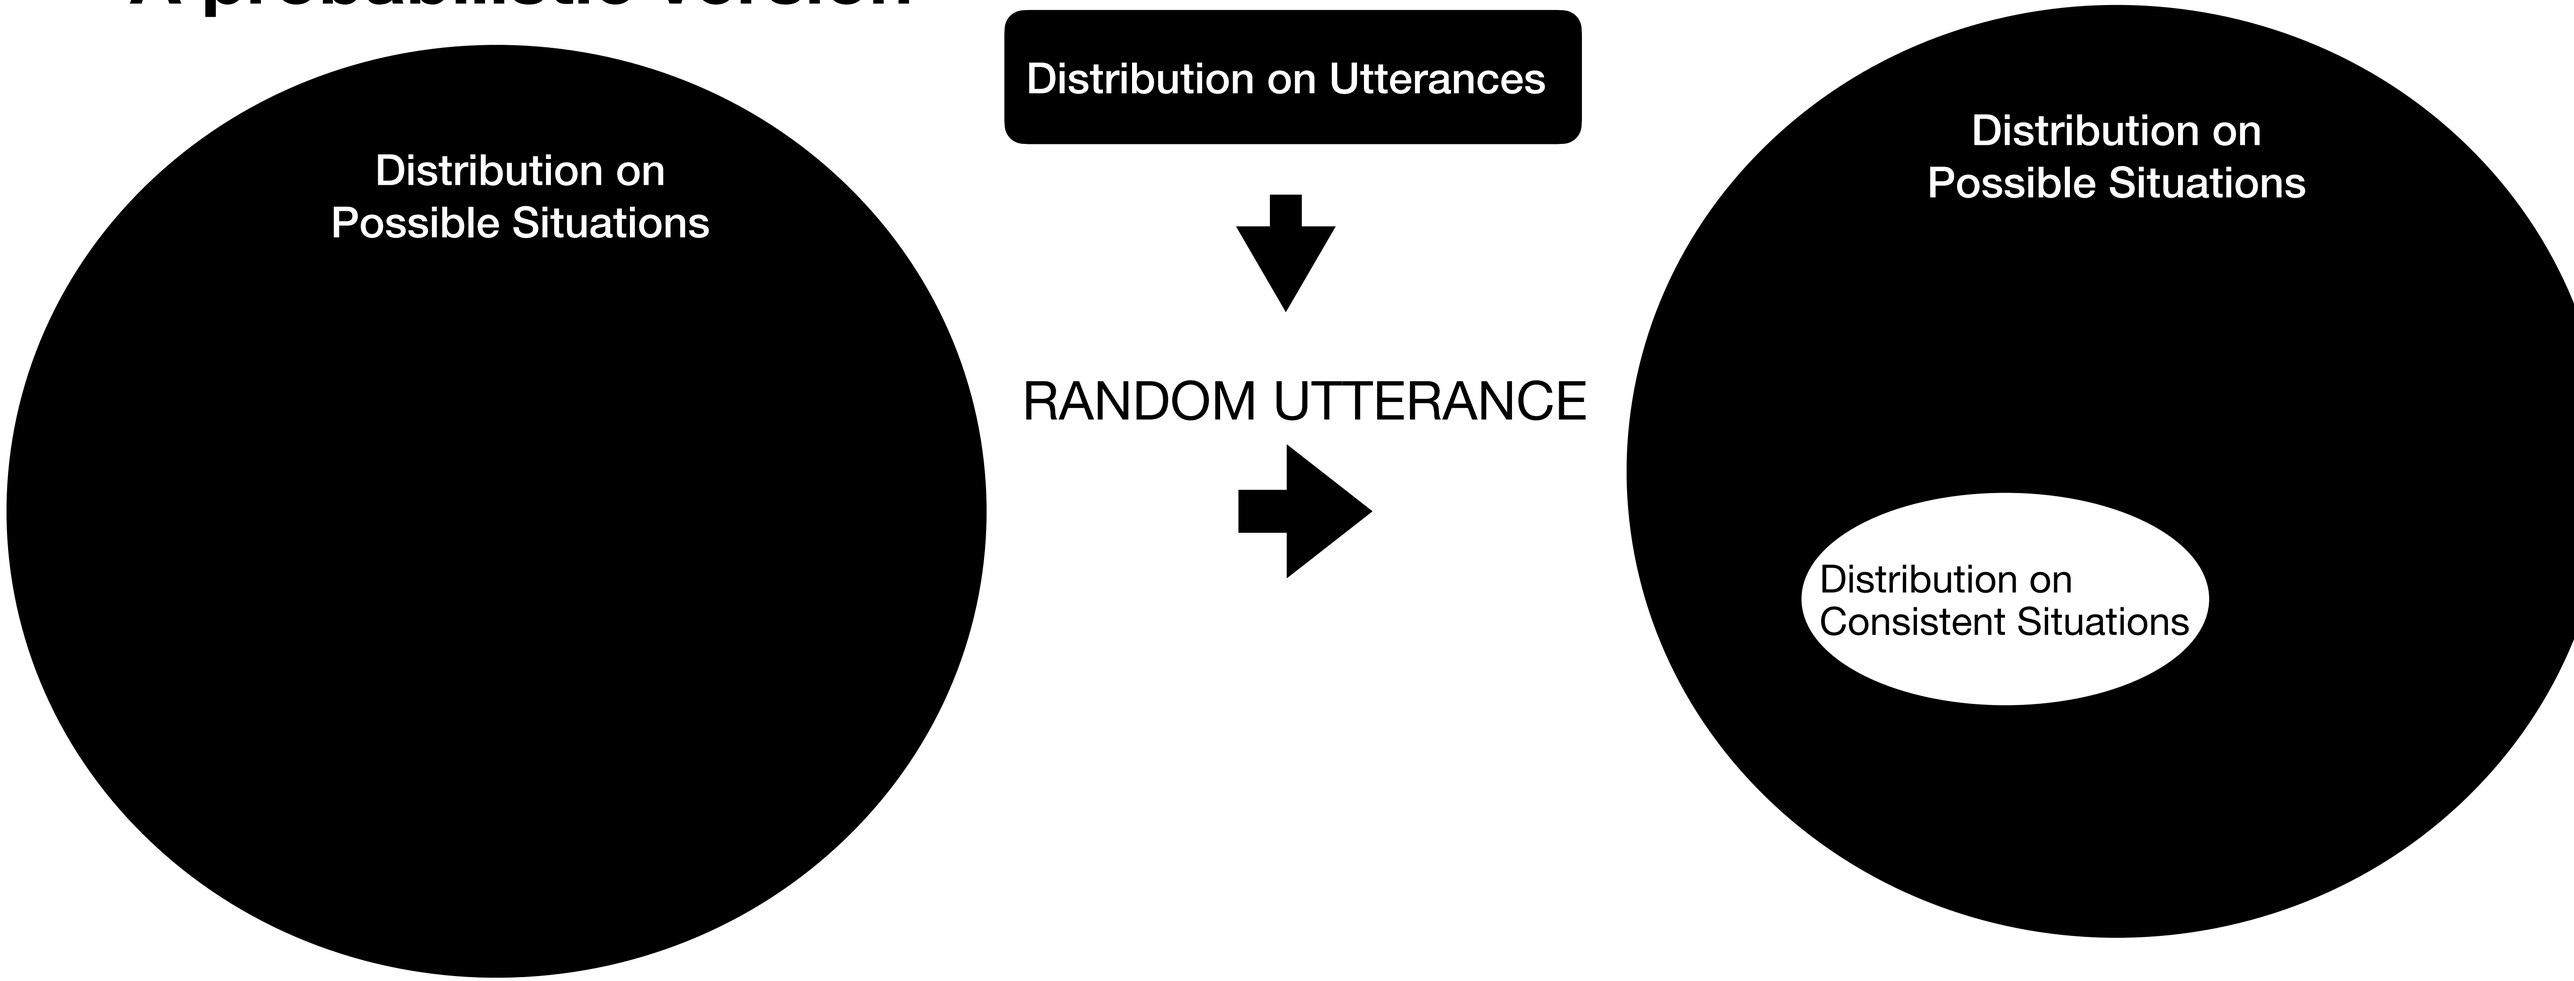

A Simple Model of Meaning

A probabilistic version

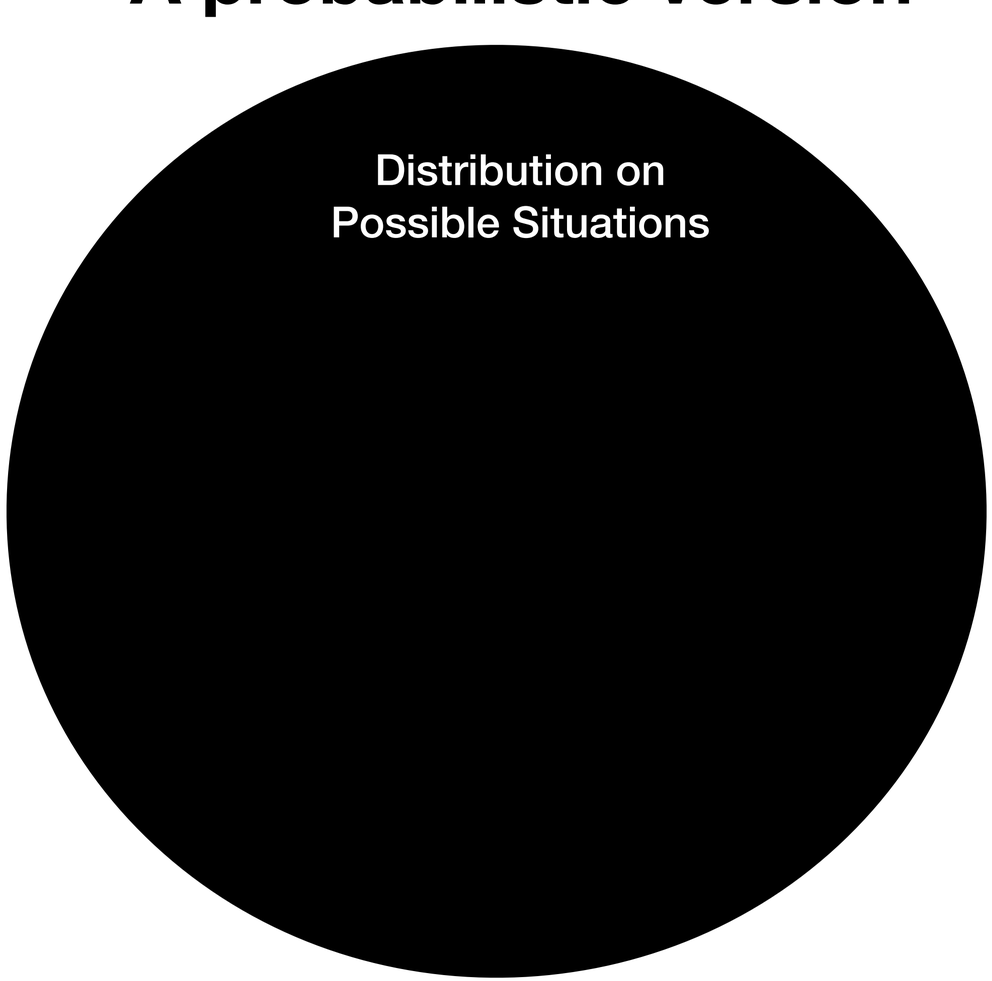

Distribution on Utterances

RANDOM UTTERANCE

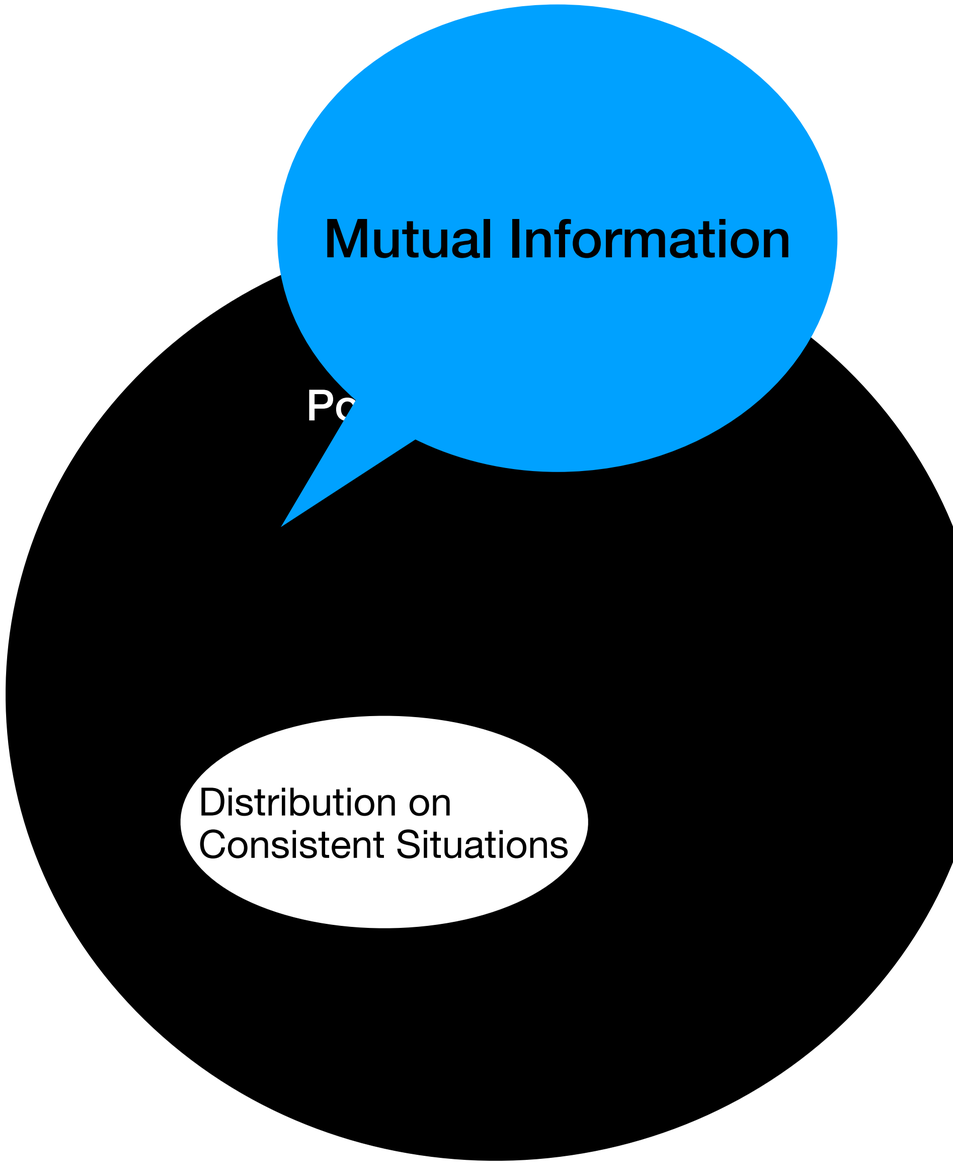

A Simple Model of Meaning

A probabilistic version

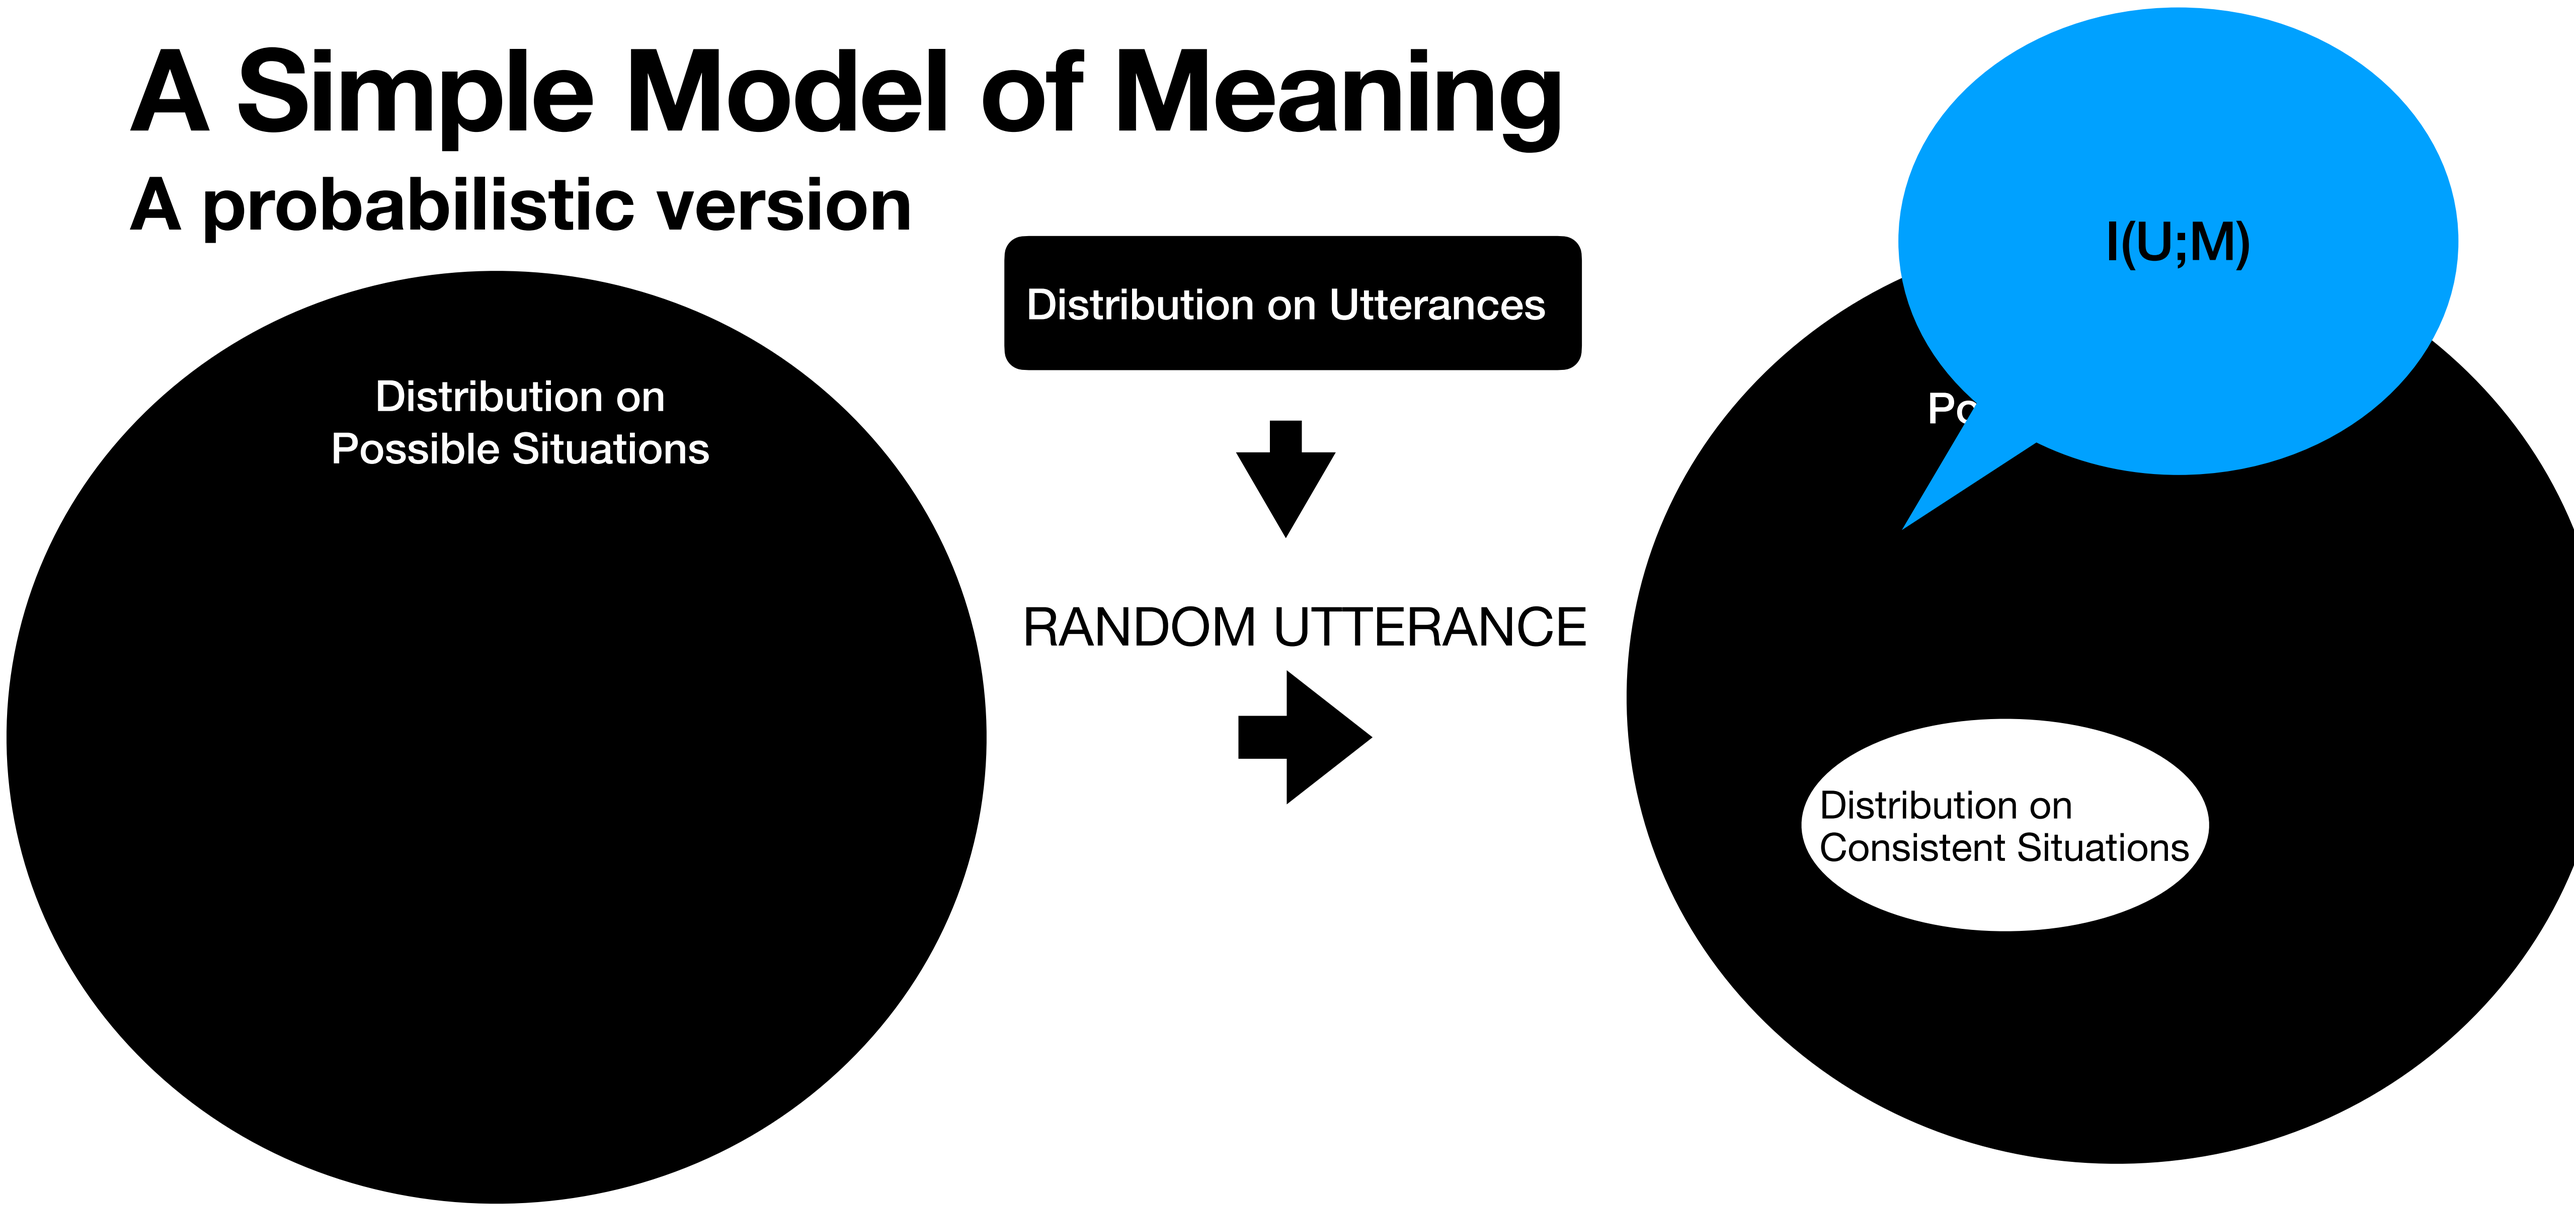

Decomposing the Mutual Information

Three kinds of information

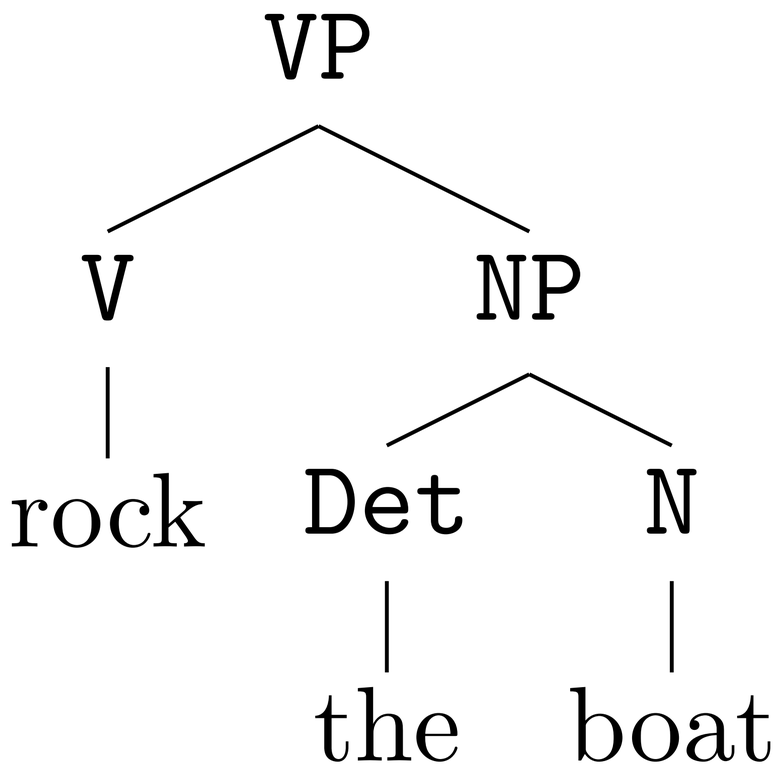

- Consider the mutual information between verb phrases consisting of a verb and direct object, and some meaning space:
  - $\mathrm{I(VP;M)}$
- Let's assume that our joint RV VP decomposes in to the set of RVs $\{\mathrm{V}, \mathrm{NP}\}$
- There are three kinds of ways that these two random variables can contribute to the meaning.

Decomposing the Mutual Information

Three kinds of information

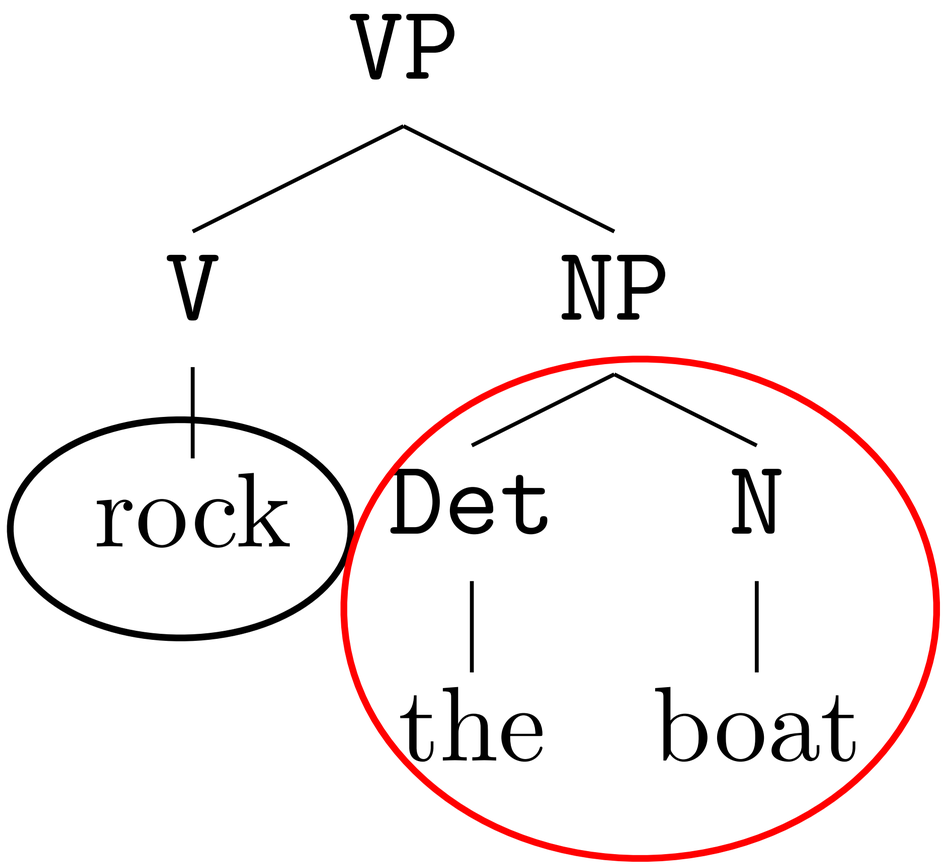

- **Unique information**:
  - I({V};M) and I({NP};M)

Decomposing the Mutual Information

Three kinds of information

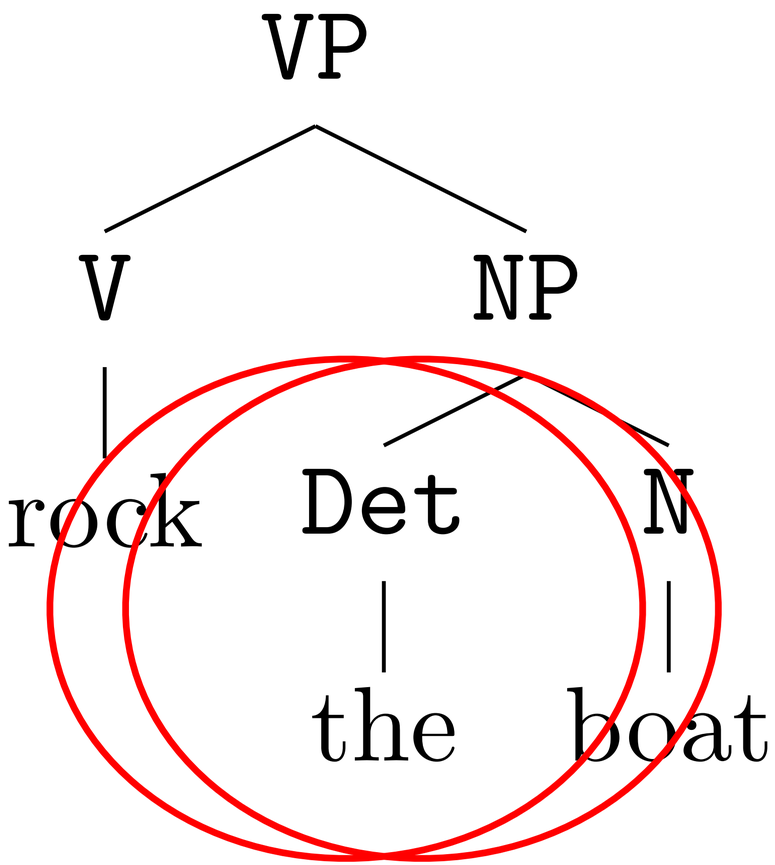

- **Unique information**:
  - I({V};M) and I({NP};M)
- **Redundant information:**
  - I({V}{NP};M)

Decomposing the Mutual Information

Three kinds of information

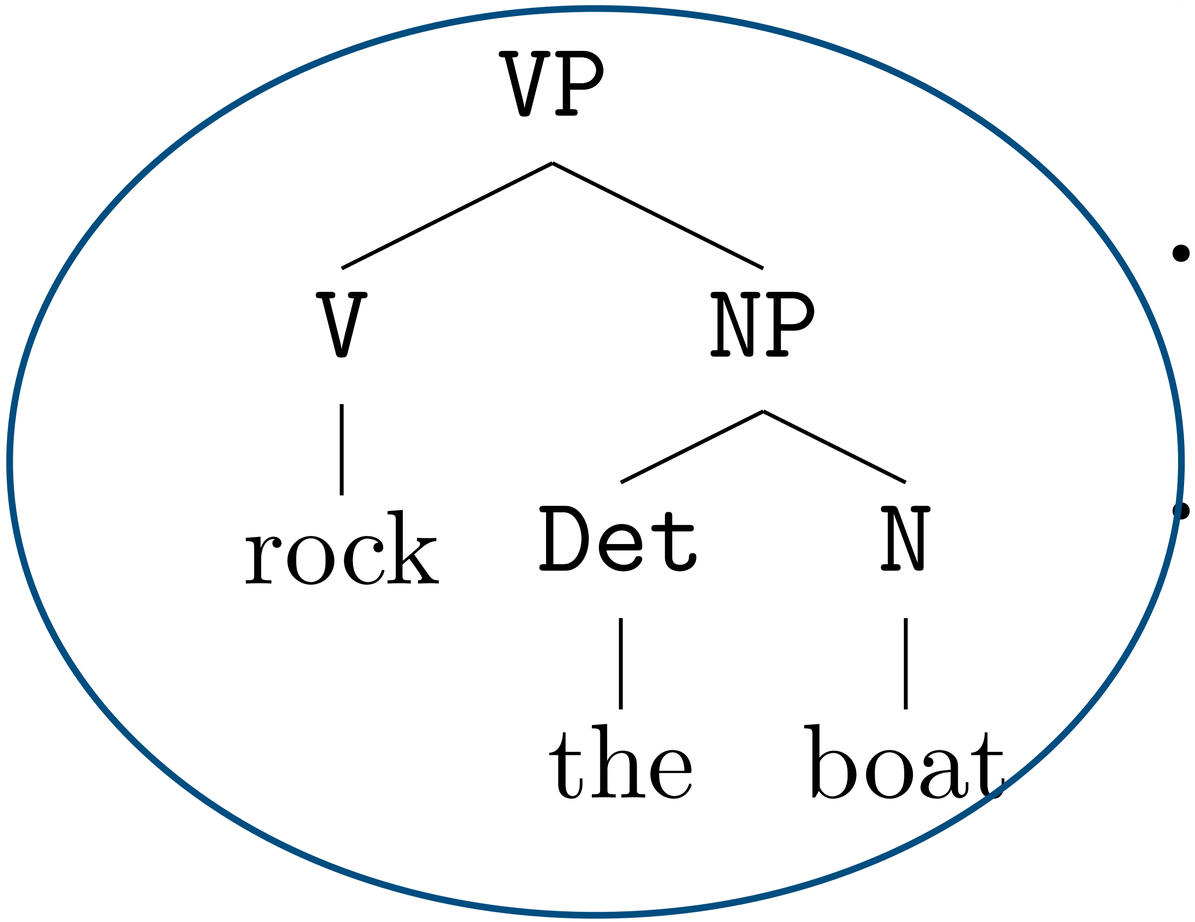

- **Unique information**:
  - I({V};M) and I({NP};M)
- **Redundant information:**
  - I({V}{NP};M)
- **Synergistic information:**
  - I({V, NP};M)

Decomposing the Mutual Information

Three kinds of information

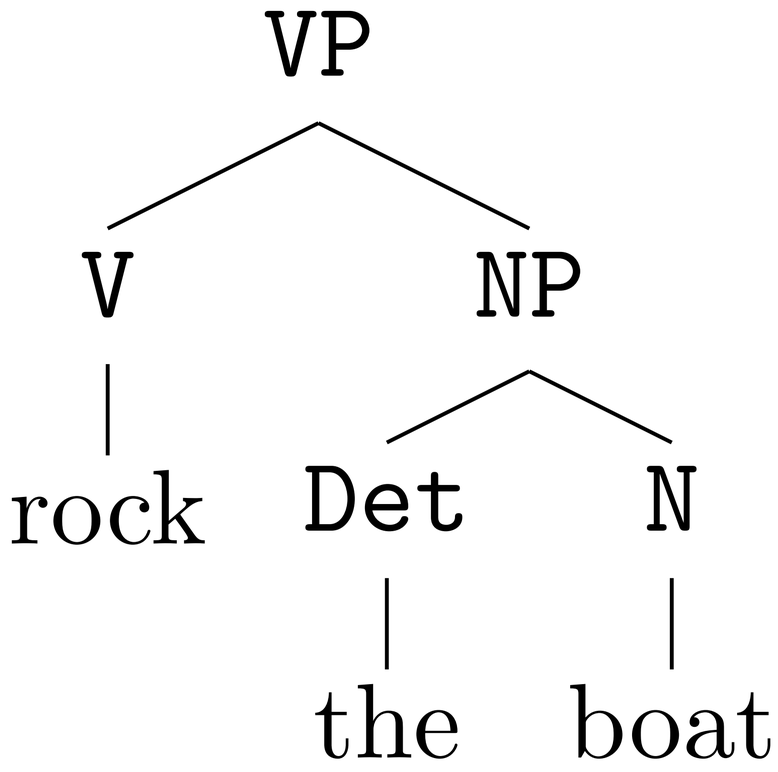

- **Decomposition**
  - I(V,NP;M) =
    I({V};M) + I({NP};M) + I({V}{NP};M)
    +I({V, NP};M)
  - I(NP;M) = I({NP};M) + I({V}{NP};M)
  - I(V;M) = I({V};M) + I({V}{NP};M)

Decomposing the Mutual Information

Three kinds of information

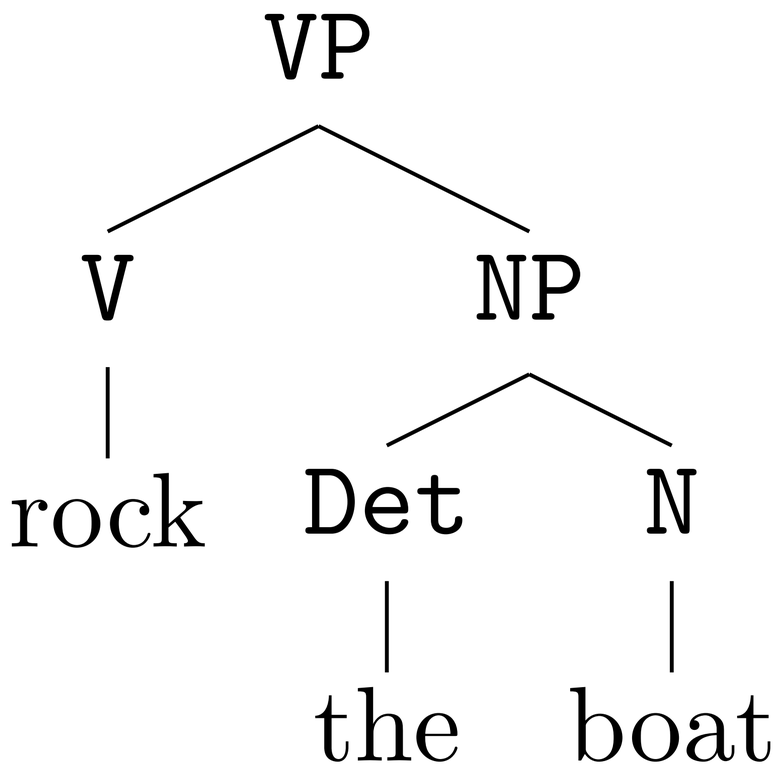

- **Partial Information Decomposition: Beer and Williams, 2010.**
  - I(V,NP;M) = I({V};M) + I({NP};M) + I({V}{NP};M) +I({V, NP};M)
  - I(NP;M) = I({NP};M) + I({V}{NP};M)
  - I(V;M) = I({V};M) + I({V}{NP};M)

Decomposing the Mutual Infor

Three kinds of information

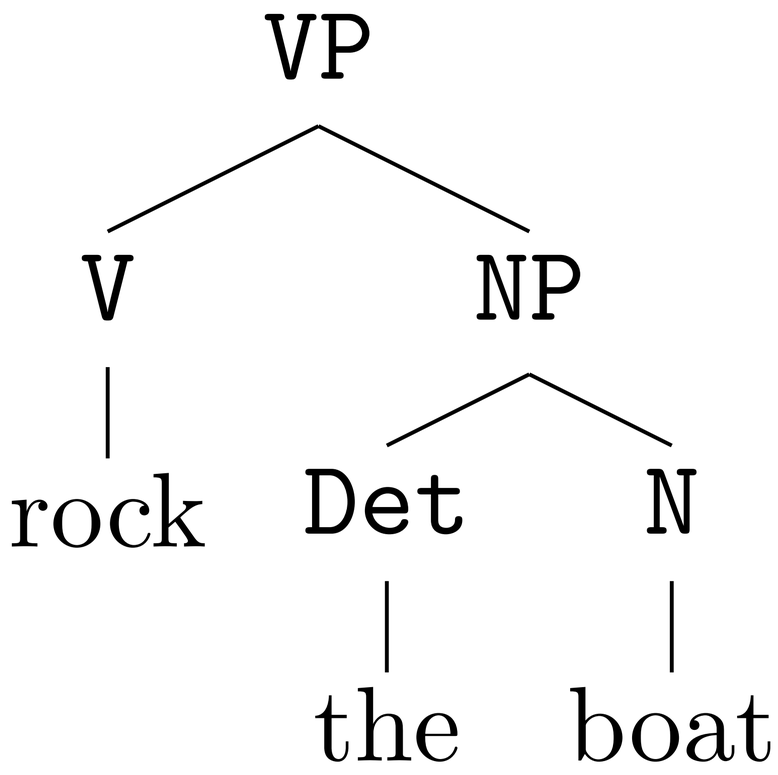

- **Partial Info**
  Beer and Wi
  - I(V,NP;M) = I({V};M) + I({NP};M) + I({V}{NP};M) +I({V, NP};M)
  - I(NP;M) = I({NP};M) + I({V}{NP};M)
  - I(V;M) = I({V};M) + I({V}{NP};M)

Need to provide a definition for either redundancy or unique information.

Decomposing the Mutual Infor

Three kinds of information

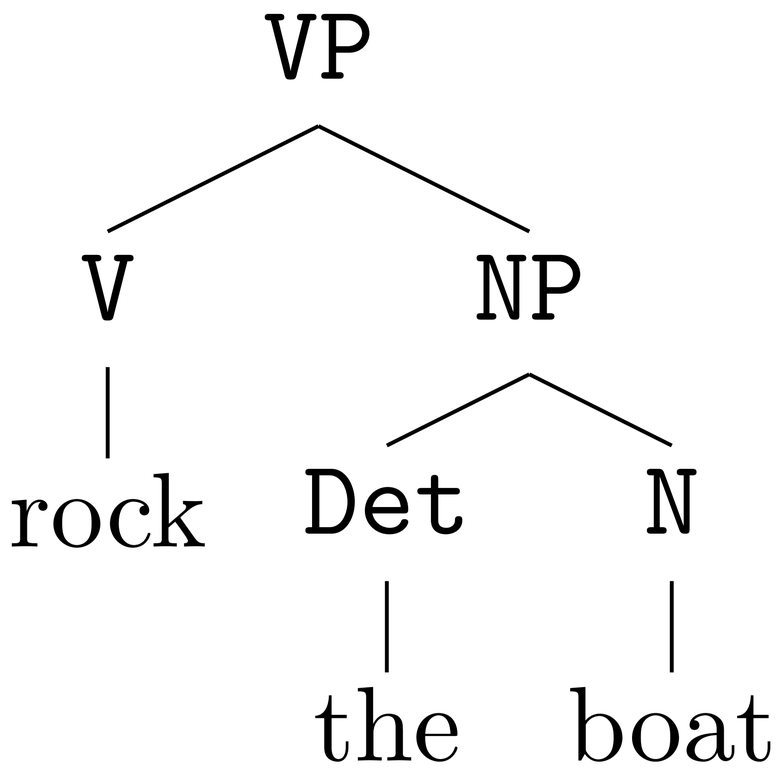

- **Partial Info**
  Beer and Wi
  - I(V,NP;M) =
    I({V};M) + I({
    +I({V, NP};M)
  - I(NP;M) = I({N
  - I(V;M) = I({V};M)

Need to provide a definition for either redundancy or unique information

Bertschinger et al. 2014

Characterizing morphological systems using partial information decomposition

Michaela Socolof and Timothy J. O’Donnell

The relation between meaning and form in morphology (word structure)

In language we have meanings that we want to express, and linguistic forms that we use to express them

Example: In English, we communicate the meaning of [plural] with *-s*

Meaning	**Form**
APPLE-**PL**	apple-**s**

The relation between meaning and form in morphology (word structure)

There are a variety of ways that meanings can correspond to linguistic forms, e.g.:

One-to-one correspondence

Meaning	**Form**
APPLE-**PL**	apple-**s**

Many-to-one correspondence

Meaning	**Form**
WALK-**3rd-SG**	walk-**s**

Agglutinative vs fusional

Different relations between meaning and form can be seen clearly in morphology

Well-known distinction between agglutinative and fusional languages

- Highly agglutinative = a single morpheme expresses a single unit of meaning (highly compositional)
- Highly fusional = a single morpheme bundles together multiple units of meaning (less compositional)

Agglutinative vs fusional

Agglutinative: Hungarian

Meaning	Form
girl-SG-DAT	lány-∅-nak
girl-PL-DAT	lány-ok-nak
girl-SG-TERM	lány-∅-ig
girl-PL-TERM	lány-ok-ig

Fusional: Russian

Meaning	Form
cat-SG-DAT	кот-у
cat-PL-DAT	кот-ам
cat-SG-GEN	кот-а
cat-PL-GEN	кот-ов

Methods

24 languages from UniMorph, noun paradigms

Example paradigm, Latin:

Meaning	Form
gazela-NOM-SG	gazela
gazela-NOM-PL	gazelae
gazela-GEN-SG	gazelae
gazela-GEN-PL	gazelārum
gazela-DAT-SG	gazelae
gazela-DAT-PL	gazelīs
gazela-ACC-SG	gazelam
gazela-ACC-PL	gazelās
gazela-ABL-SG	gazelā
gazela-ABL-PL	gazelīs
gazela-VOC-SG	gazela
gazela-VOC-PL	gazelae

Methods

We have to define our source and target variables

Elements of the meaning are our sources, elements of the form are our targets

Meaning variables: CASE, NUMBER, DEFINITENESS (in relevant languages)

Form variables: character slots in the suffix, aligned with LingPy (List and Forkel 2021)

M_1	M_2	M_3	F_1	F_2	F_3	F_4
КОТ	GEN	SG	КОТ	а	-	-
КОТ	DAT	PL	КОТ	а	м	-
КОТ	INS	PL	КОТ	а	м	и

Table 6: Random variable structure for three word forms in Russian

Results

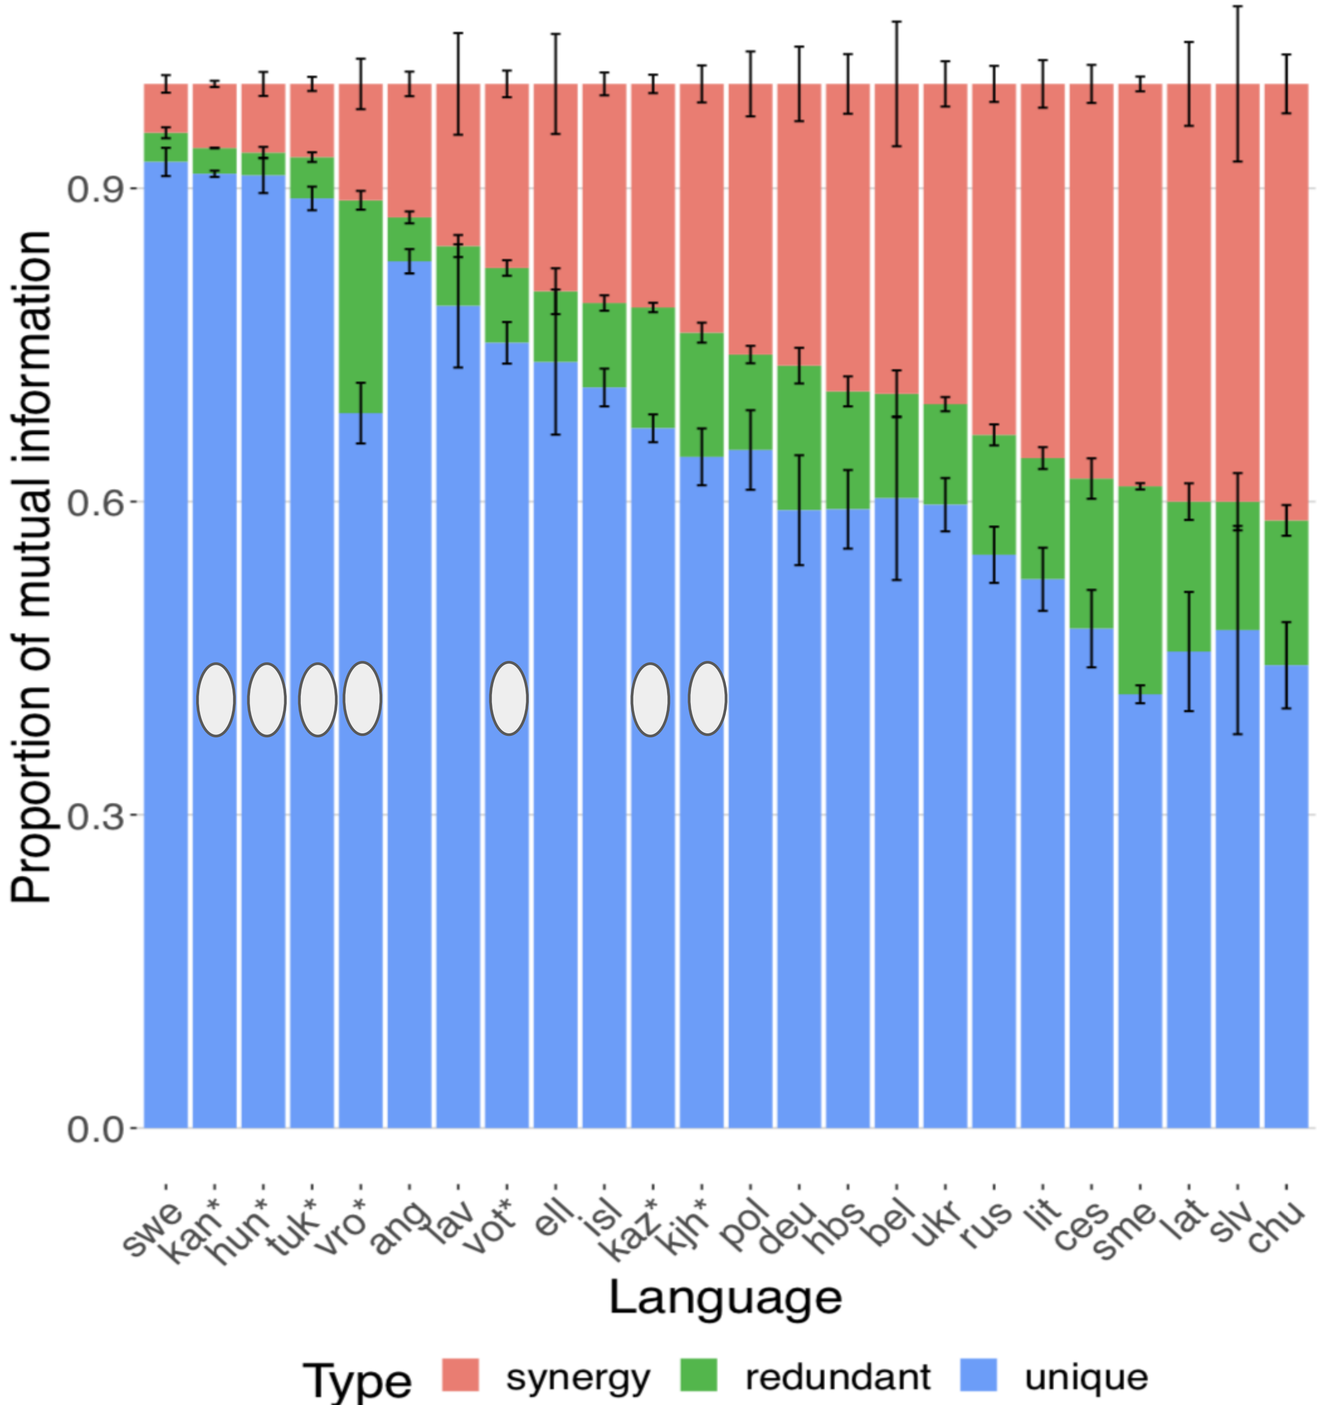
Proportion of mutual information
0.0
0.3
0.6
0.9
swe
kan*
hun*
tuk*
vro*
ang
lav
vot*
ell
isl
kaz*
kjh*
pol
deu
hbs
bel
ukr
rus
lit
ces
sme
lat
slv
chu
Language
Type
synergy
redundant
unique

Quantitative formalization of Compositionality

- The proportion of unique information in a system is a measure of compositionality.
- There are differentiable versions of these quantities.
- Requires some (measure on a) meaning space.

Conclusion

- Showed several different studies looking at classical linguistic concepts.
- More to say about definitions of productivity and the exact nature of linguistic generalization.
- Much more work to be done to use this definition of compositionality.
- Would love to collaborate on developing large-scale versions of these ideas.

Thank you!