Simulating Language
6: The evolution of compositionality

Simon Kirby

simon.kirby@ed.ac.uk

What's missing from our models so far?

- In all our models, both meanings and signals are *atomic*

- In reality (for all communicating species) both meanings and signals have internal structure
  - They have internal parts that can be recombined

- Does this matter at all?

How we leverage structure...

- What's the missing word?

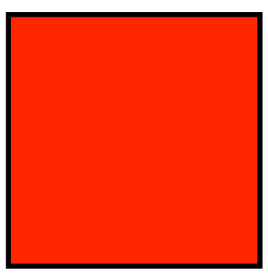

copiro

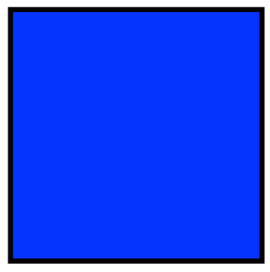

nelupi

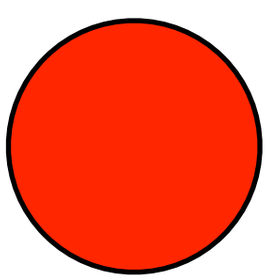

saroka

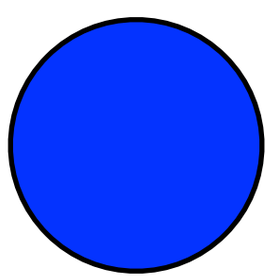

?????

How we leverage structure...

- What's the missing word?

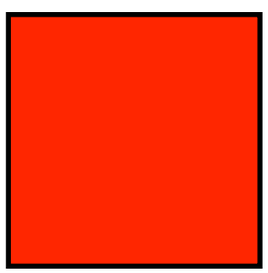

copiro

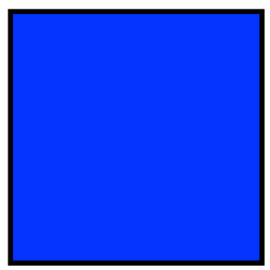

nelupi

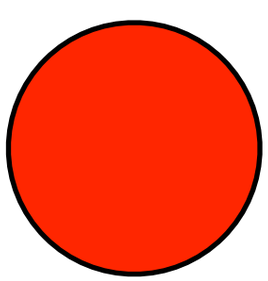

saroka

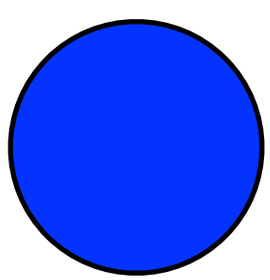

sotuga

How we leverage structure...

- What's the missing word?

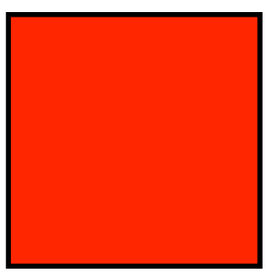

copiro

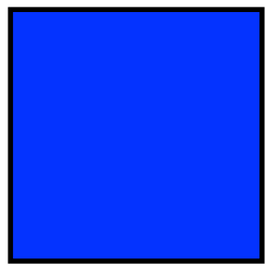

neluro

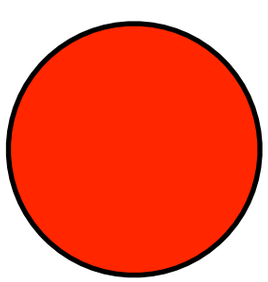

copida

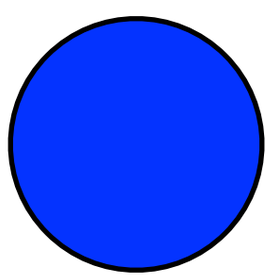

?????

How we leverage structure...

- What's the missing word?

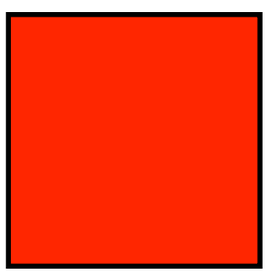

copiro

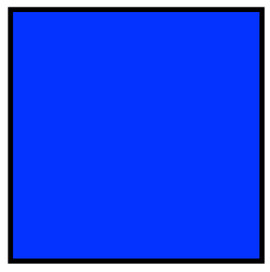

neluro

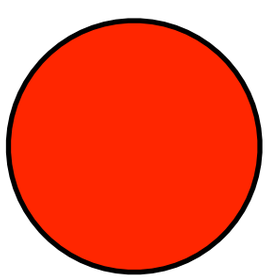

copida

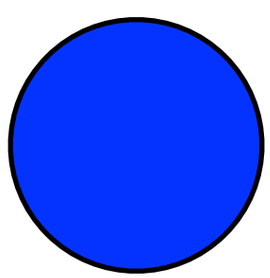

neluda

What's the difference?

- In the first example, the meanings and signals might as well have been unstructured/atomic
    - We were essentially seeing a vocabulary.
- In the second example, we relied on the fact that:
    - the meanings had internal structure (e.g. color and shape),
    - and the signals had internal structure (e.g. subsequences of syllables)
    - and the mapping utilises the structure in a way that allows us to **generalise**

Compositionality

- The crucial structure of the mapping is *compositionality*

Compositionality: the meaning of the whole is a function of the meaning of the parts and how they are put together.

- Arguably the most important feature of the syntax of human language
- Enables open-ended communication (more fundamentally than recursion)
- Strangely, it is rare and quite restricted in non-human animals, despite being a hugely beneficial trait!

Where does compositionality come from?

- Compositionally-structured meaning-signal mappings are adaptive, since they enable open-ended communication

- So... might suggest an explanation in terms of natural selection:

"Evolutionary theory offers clear criteria for when a trait should be attributed to natural selection: complex design for some function, and *the absence of alternative processes capable of explaining such complexity*. Human language meets these criteria." Pinker & Bloom (1990)

- But are there *alternative process*?

  And anyway, how exactly do properties of our innate endowment lead to observable properties of language (the adaptations they purport to explain)? This is **problem of linkage** again...

Iterated learning again

- To solve the problem of linkage, we need to turn again to the iterated learning model

- What happens if, instead of mappings between atomic meanings and signals, we allowed for meanings and signals with structure?

- Could we see a *cultural* rather than biological evolution of compositionality?

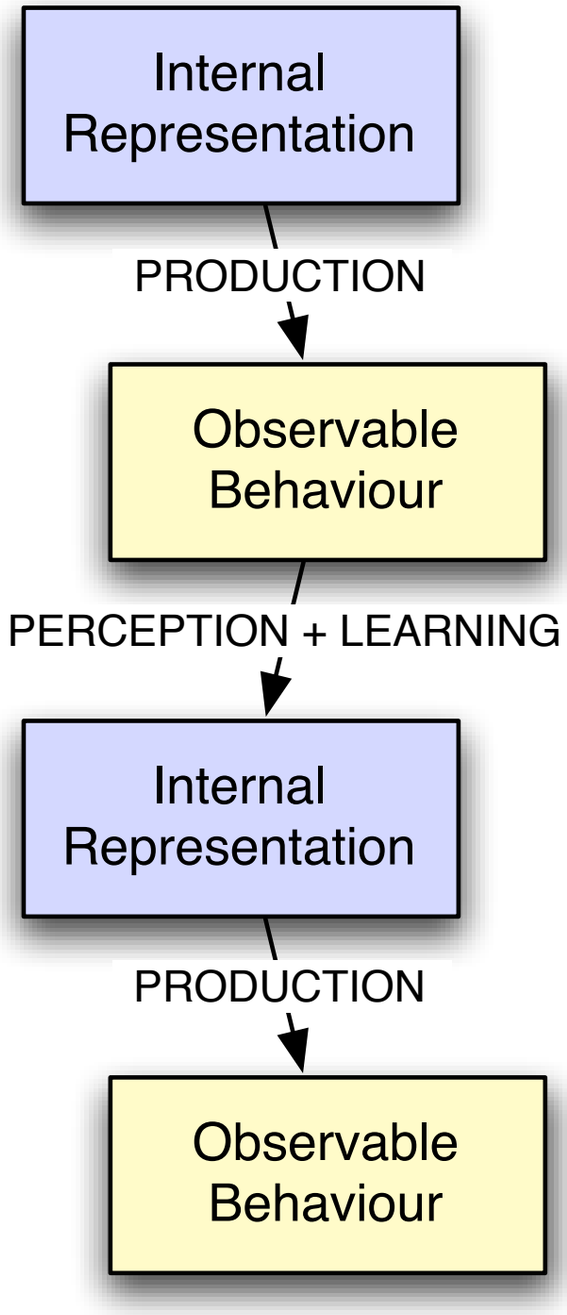

Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. Cognition, 141, 87-102.

The simplest possible model?

- What's the simplest setup that would still allow us to compare **compositional** and non-compositional (**holistic**) languages?

- Signals: two syllable words, with two possible syllables

  baba, *baki*, *kiba*, *kiki*

  ['aa', 'ab', 'ba', 'bb']

- Meanings: two features, with two possible "values" on each feature

  square+red, circle+red, square+blue, circle+blue

  ['02', '12', '03', '13']

Some grammars

$$
\begin{aligned}
S:02 &\rightarrow aa \\
S:03 &\rightarrow ab \\
S:12 &\rightarrow bb \\
S:13 &\rightarrow ba
\end{aligned}
$$

Holistic

$$
\begin{aligned}
S &\rightarrow A\ B \\
A:0 &\rightarrow a \\
A:1 &\rightarrow b \\
B:2 &\rightarrow a \\
B:3 &\rightarrow b
\end{aligned}
$$

Compositional

$$S:\{02,03,12,13\} \rightarrow aa$$

Degenerate

A very general prior

- Occam's razor: **simpler solutions are more likely than complex ones**

- Faced with different theories of the world (or data), we should prefer the simpler ones

- We can actually measure simplicity by looking at how much it takes to encode (roughly, write down) our grammars:

$$S : \{02,03,12,13\} \rightarrow aa$$

$$\begin{aligned} S &\rightarrow A\ B \\ A:0 &\rightarrow a \\ A:1 &\rightarrow b \\ B:2 &\rightarrow a \\ B:3 &\rightarrow b \end{aligned}$$

$$\begin{aligned} S:02 &\rightarrow aa \\ S:03 &\rightarrow ab \\ S:12 &\rightarrow bb \\ S:13 &\rightarrow ba \end{aligned}$$

More complex

Learning

$$P(h|d) \propto P(d|h)P(h)$$

Posterior: learners pick grammars based on their probability given the sentences they see

Likelihood: favour grammars that predict the data well

Prior: favour simple grammars

What happens when we iterate in a chain?

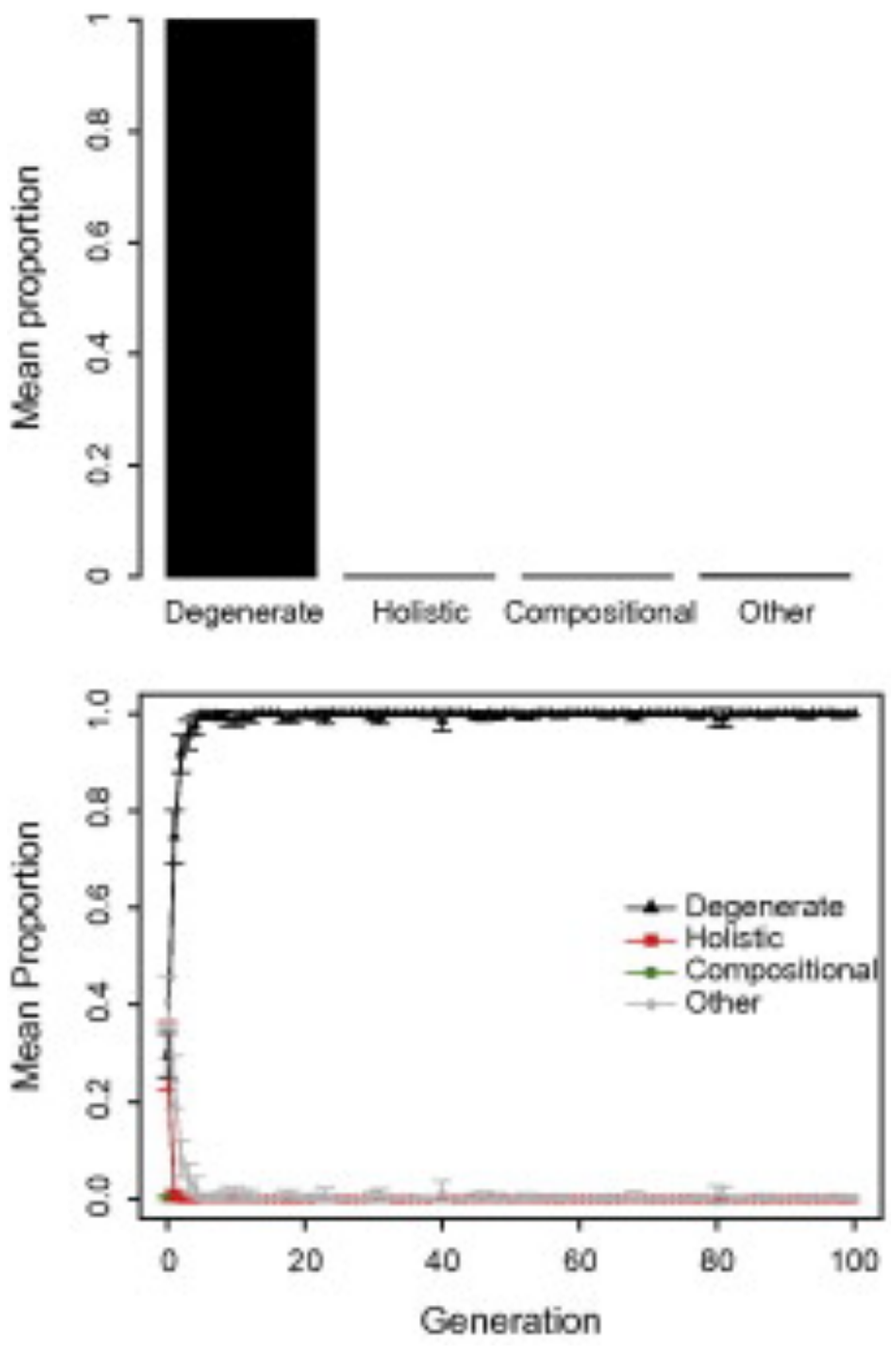

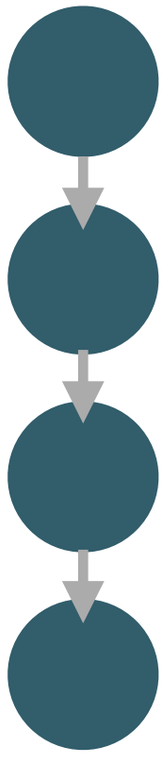

$$S : \{02,03,12,13\} \rightarrow aa$$

Very **learnable** (i.e. simple), but **inexpressive**

Communication

- Language adapts to the learner. So simplest possible language emerges, but it's useless for communication!

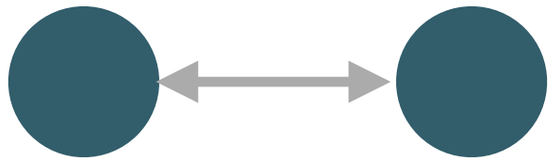

- An alternative model: two agents interact with each other and learn from their interactions.

- Use the simple "rational" speaker that we implemented before.

What happens when a pair interact?

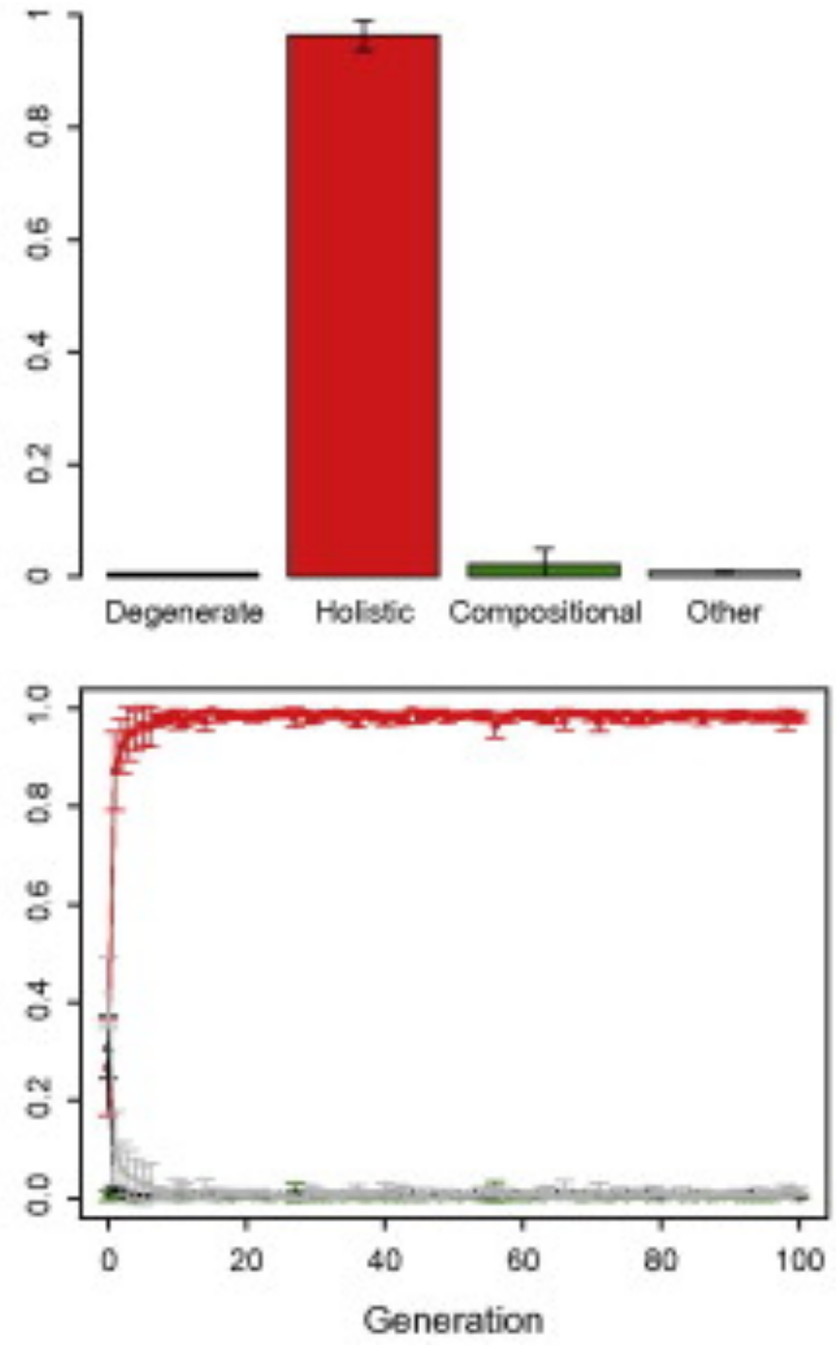

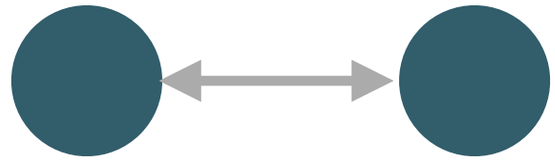

$$S : 02 \rightarrow aa$$
$$S : 03 \rightarrow ab$$
$$S : 12 \rightarrow bb$$
$$S : 13 \rightarrow ba$$

Expressive, but not very **learnable** (i.e. complex)

OK, what about both iteration *and* interaction?

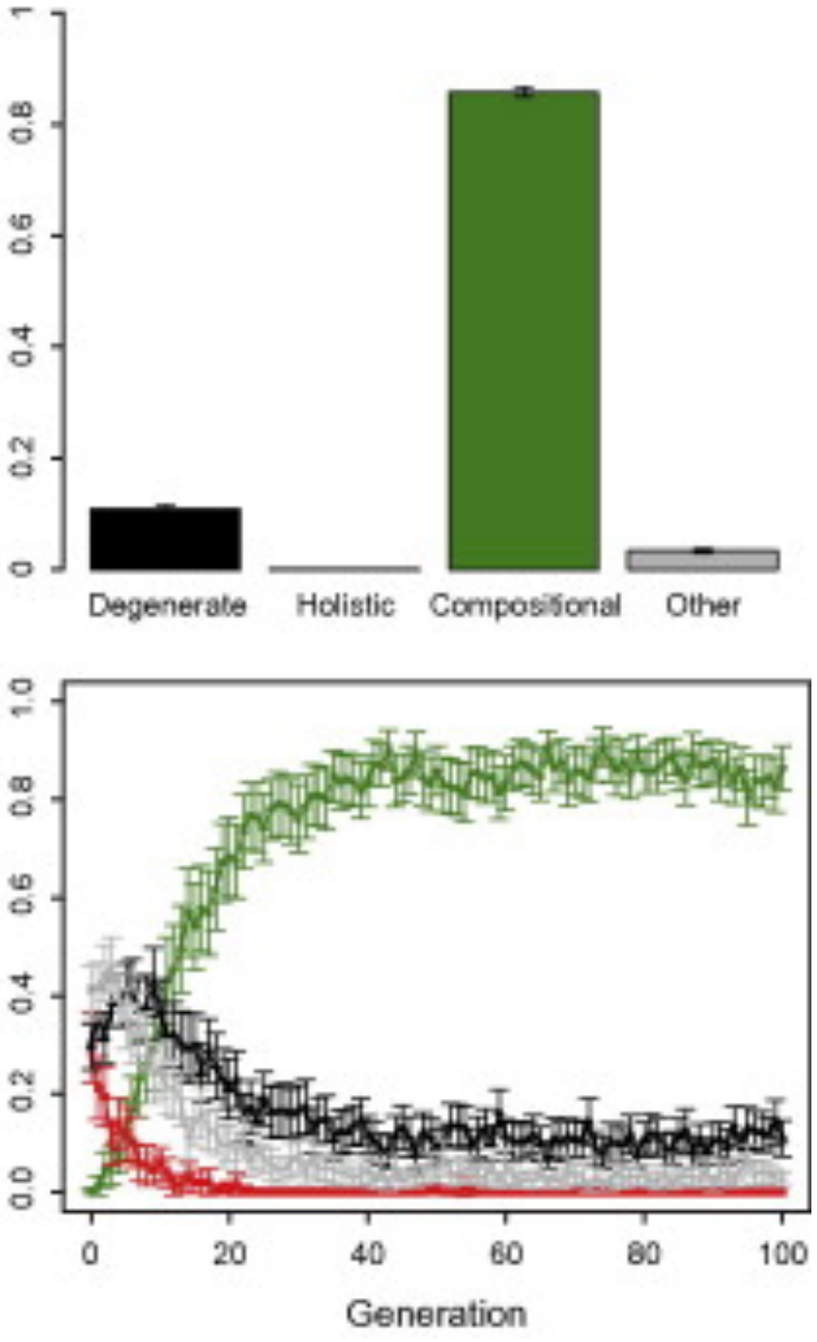

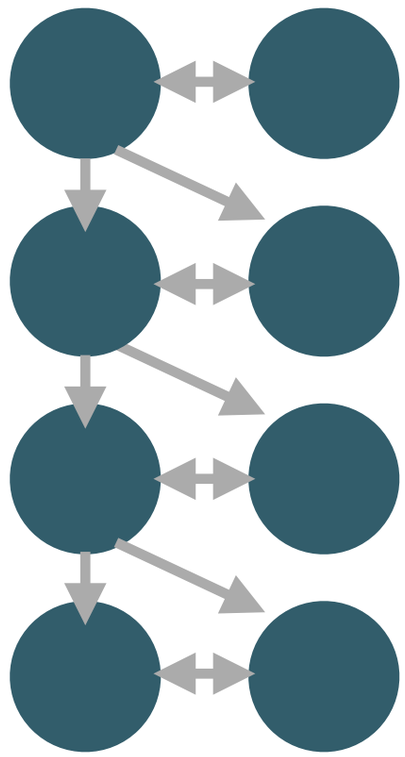

$$
\begin{aligned}
S &\rightarrow A\ B \\
A:0 &\rightarrow a \\
A:1 &\rightarrow b \\
B:2 &\rightarrow a \\
B:3 &\rightarrow b
\end{aligned}
$$

Expressive, and fairly **learnable** (i.e. reasonably simple)

How confident can we be in this result?

- This is an interesting result, but how realistic is it?
- Kirby et al (2015) recreate the simulation in the experiment lab
- Participants come into the lab and learn a miniature holistic language, then use it to communicate with another participant
- New pairs of participants learn from the behaviour of the previous pair
- New learners + communication -> compositional languages
  New learners + no communication -> degenerate languages
  No new learners + communication -> holistic languages

Language has to fit through a narrow *bottleneck*

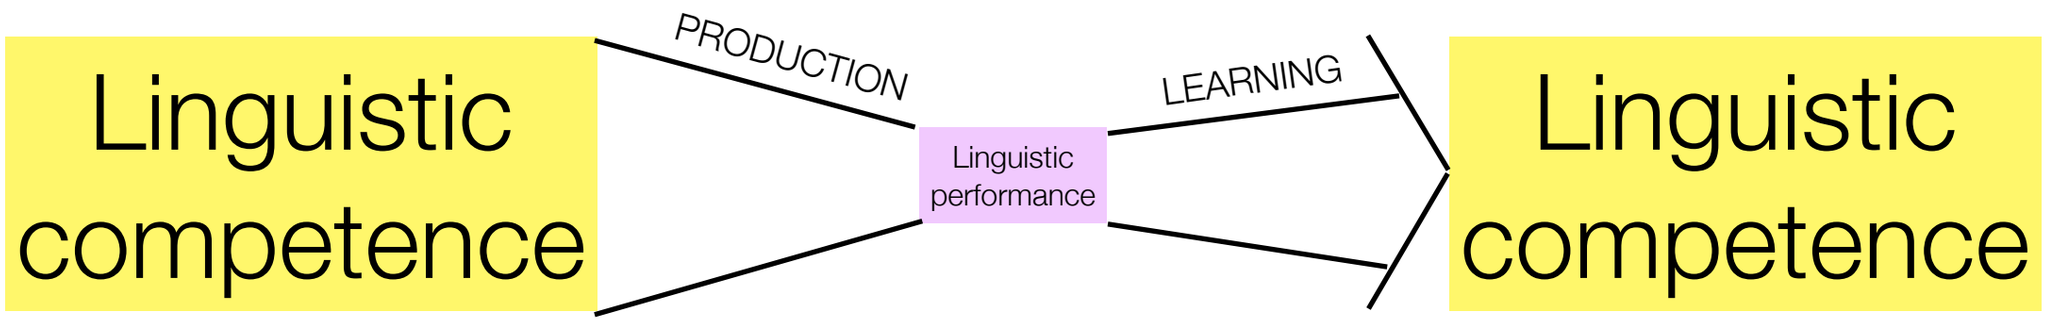

- This has profound implications for the structure of language

- Only languages that are *generalisable* from limited exposure are stable if they are transmitted to new learners

- Only languages that are unambiguous are stable if they are used by speakers who avoid ambiguity

- Compositional syntax is an adaptive response **by language** (arising from cultural evolution) to the problem of getting through this bottleneck

Up next

- Labs: a replication of the model in Kirby et al (2015)

- Coming next… we've been assuming particular prior biases throughout this course, but where do they come from?
  - Next lecture: learning how to learn

- Final weeks of the course: how **biological evolution** can shape learning and culture, and how this finally answers some fundamental questions about whether language is innate