Simulating Language
Lecture 12: Learning, culture, innateness

Kenny Smith

kenny.smith@ed.ac.uk

Approaching the end

9 **16/3-20/3**	Lecture 12 Learning, culture, innateness	*Catch-up lab 4pm-5pm only* *(optional)*	*Catch-up lab* *(optional)*
10 **23/3-27/3**	Lecture 13 Iterated Bayesian Learning: culture and innateness	*Catch-up lab 4pm-5pm only* *(optional)*	Lab 10 Extending Iterated Bayesian Learning
11 **30/3-3/4**	Feedback Meeting: Feedback on First Assignment	Lecture 14 The evolution of learning bias	Lecture 15 Human simulation

This week’s reading: Chater & Christiansen (2010)

- The classic view of acquisition
  - Language acquisition *should* be really hard or impossible, but children find it easy
  - They must have access to more than just data: innate, language-specific knowledge
  - An evolved language faculty?
- The cultural perspective (as advanced by Chater & Christiansen)
  - Languages are adapted to language learners
  - Language learners biases are *automatically the right ones*, because language will evolve to reflect those biases

Are these two perspectives really that different?

- Both posit a nice simple relationship between biases of learners and properties of language
- Differ in the details of why the fit between learners and languages exists
  - Classic account: learners have adapted to languages (?)
  - Cultural account: languages adapt to learners
- Differ in strength of bias?
  - Classic account: absolute constraints
  - Cultural accounts: weak biases, strong effects?
- Differ in hypothesised domain-specificity of biases
  - Classic account: must be domain-specific
  - Cultural account: biases likely to be domain-general

This lecture

- A case where culture really seems to **add something**
- Bottlenecks, generalisation, the evolution of syntax

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?

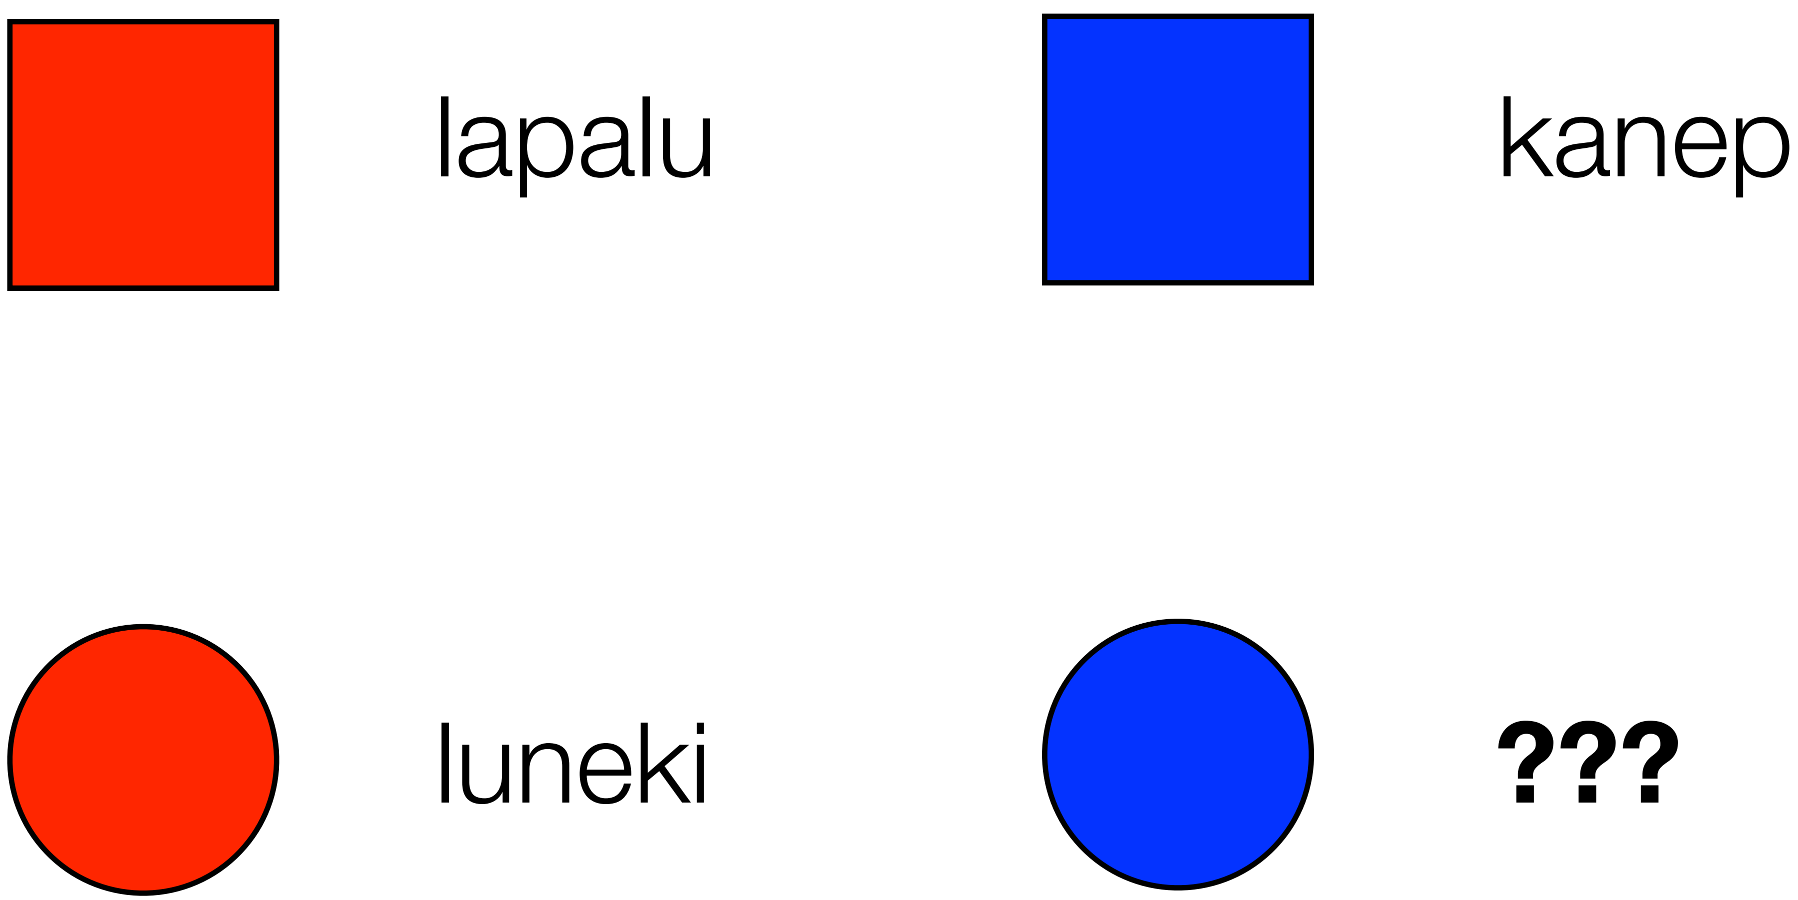

How we leverage structure...

- What's the missing word?

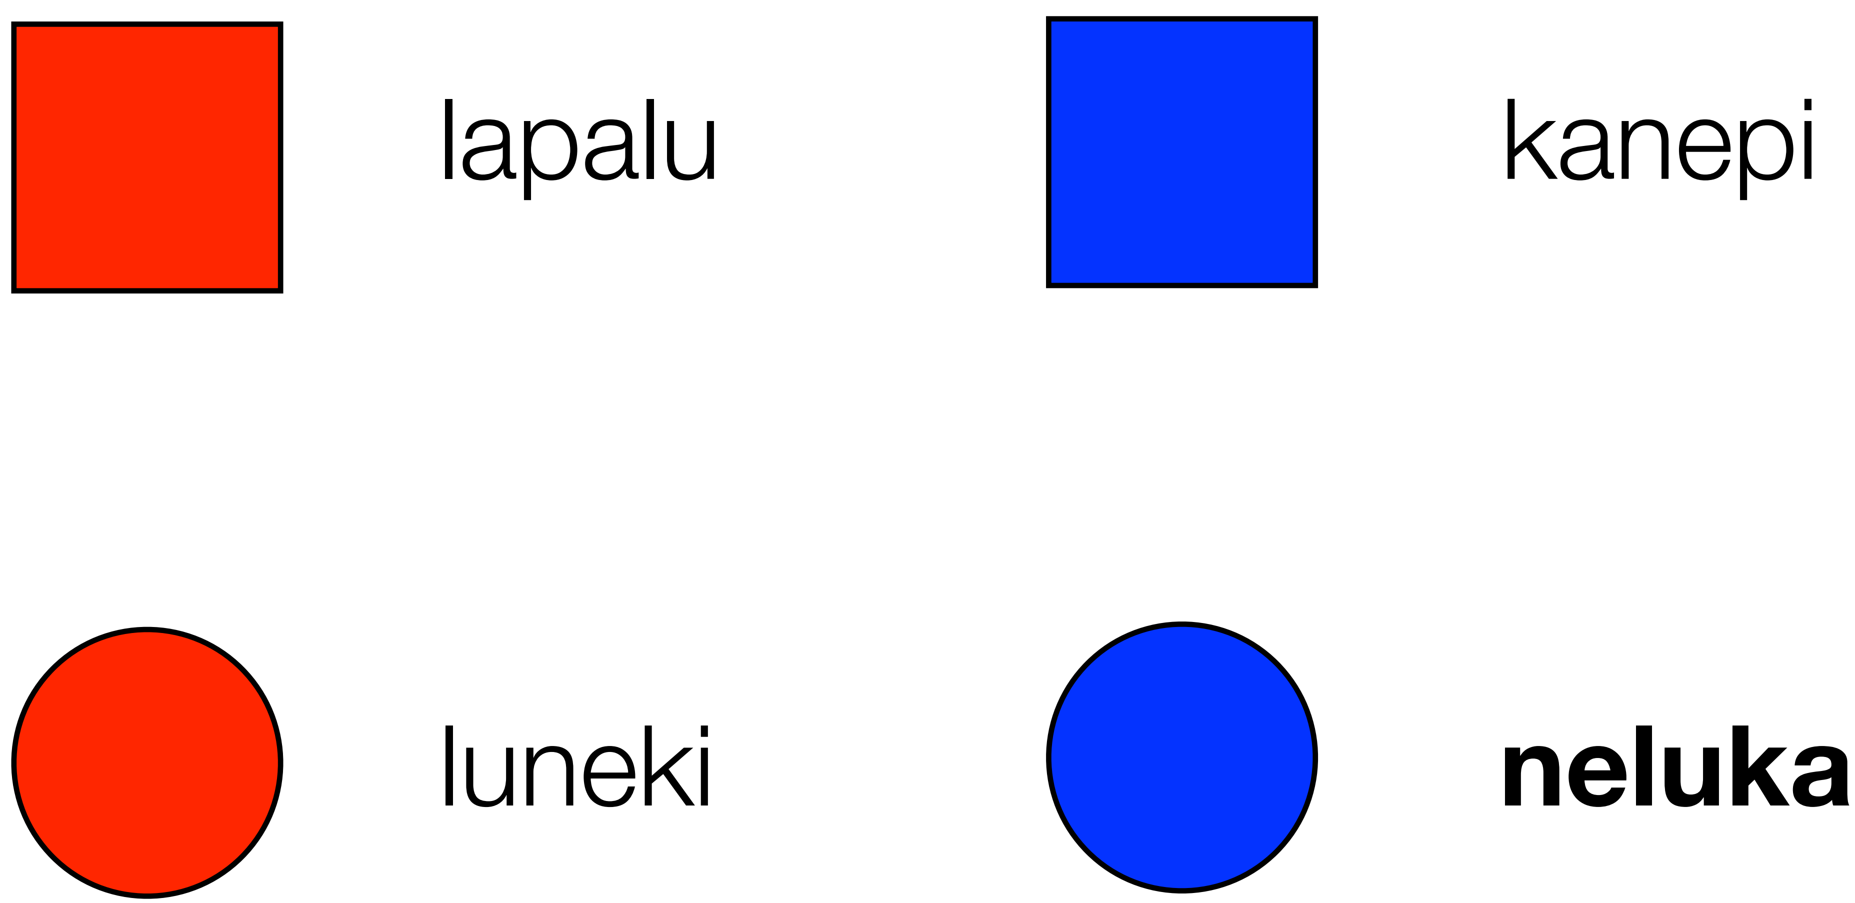

How we leverage structure...

- What's the missing word?

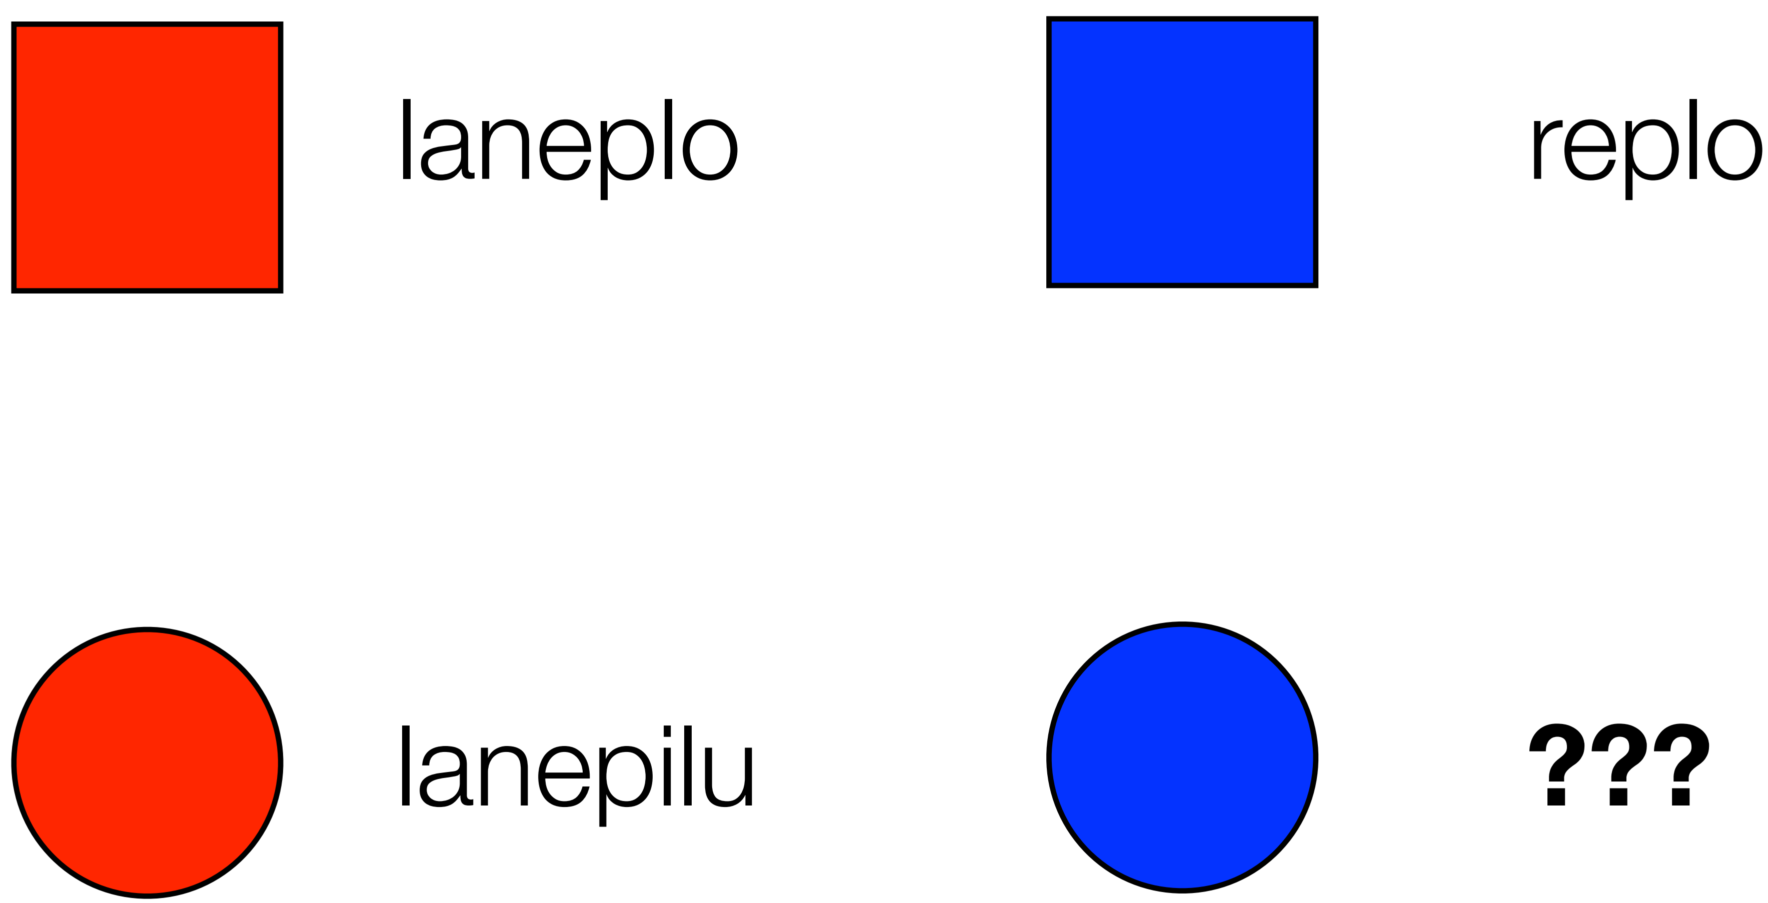

How we leverage structure…

- What’s the missing word?

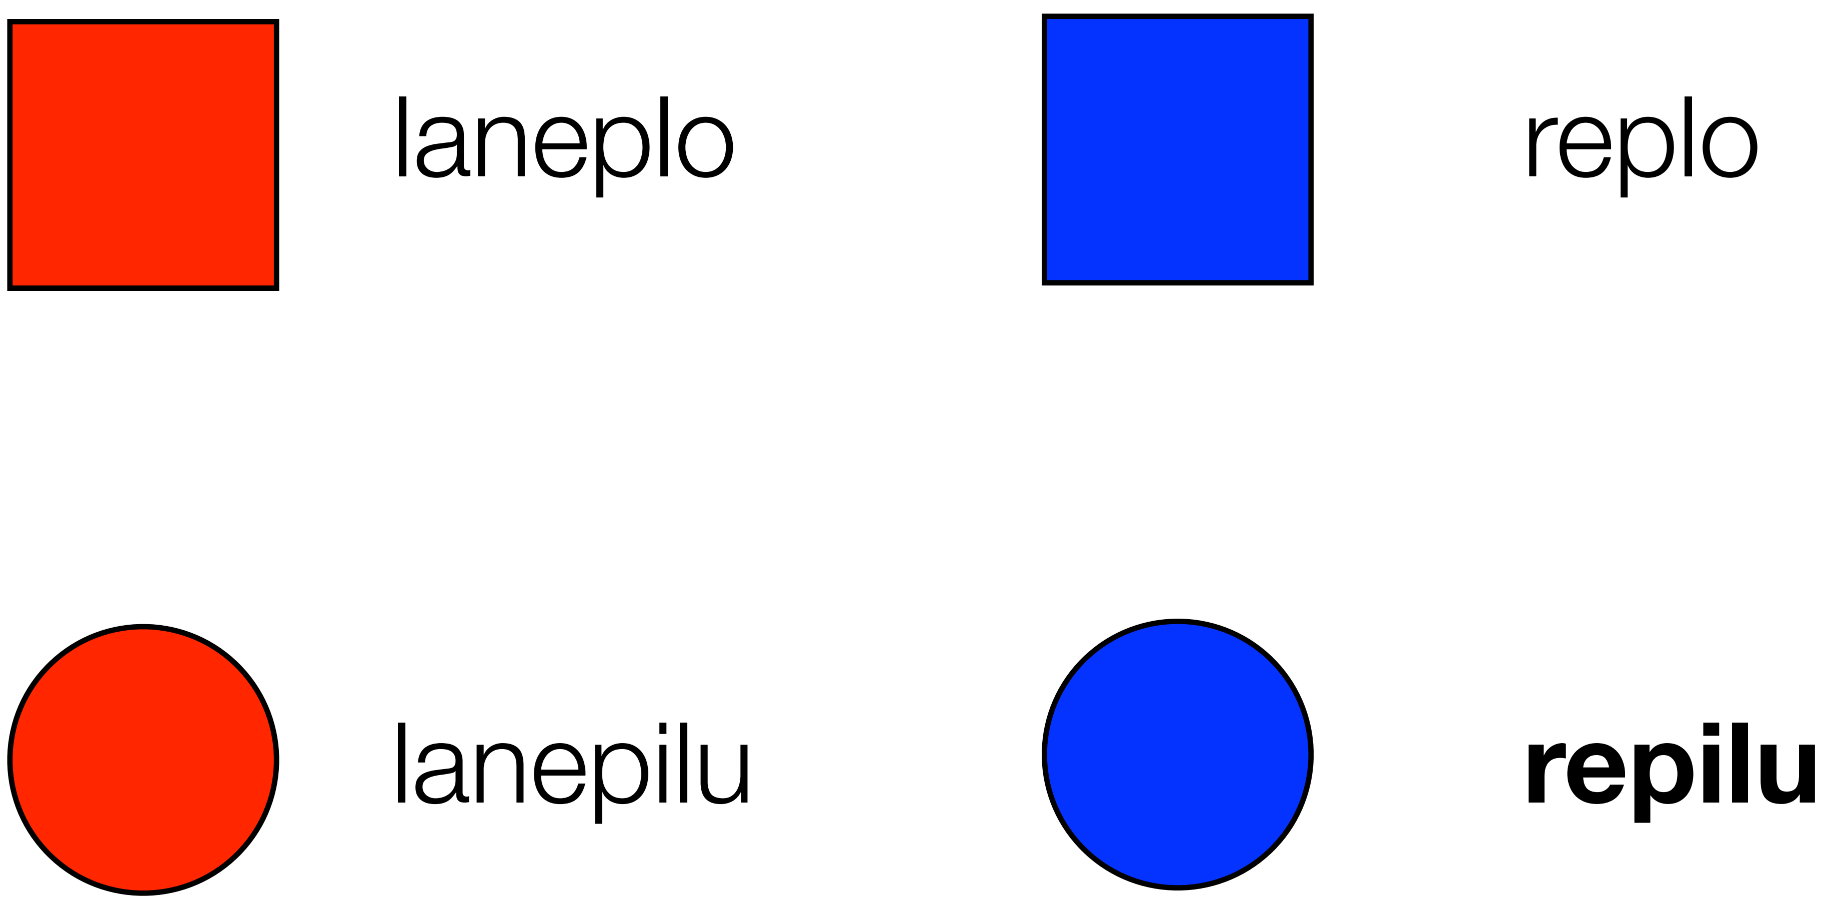

What’s the difference?

- 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 (almost) unique to humans, 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?

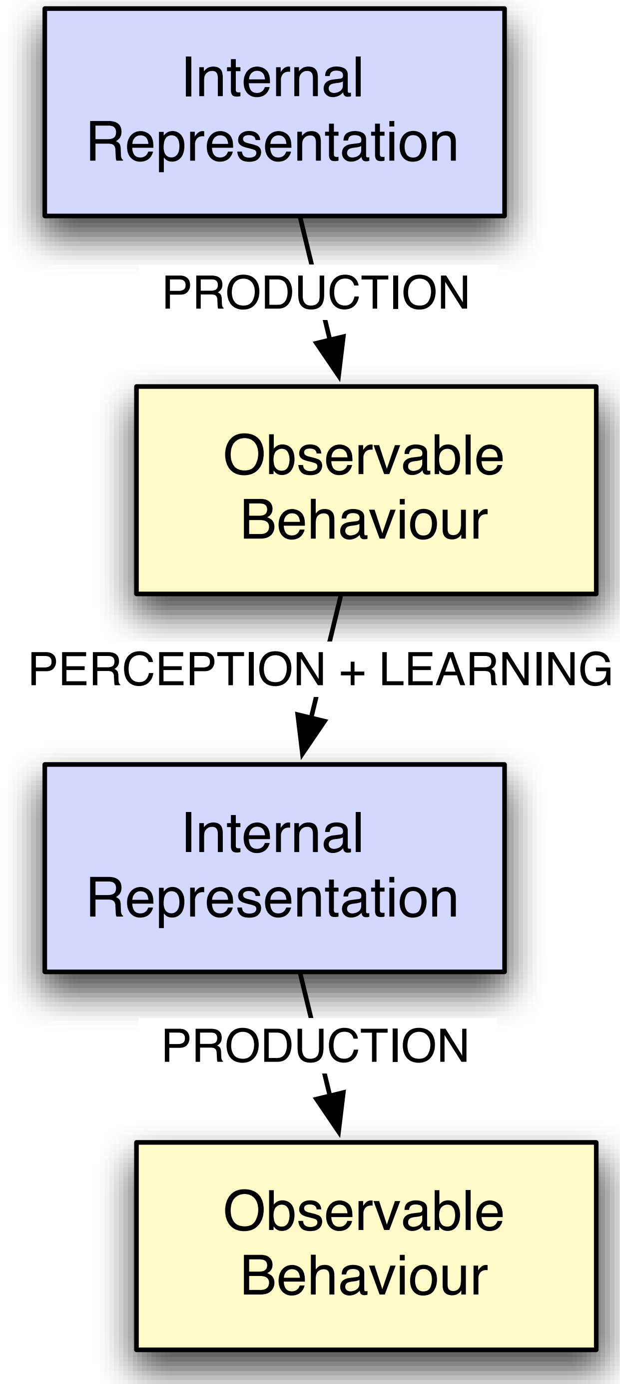

Holistic vs. Compositional

- It’s not the structure in meanings/signals that matters, but whether that structure is utilised by the mapping

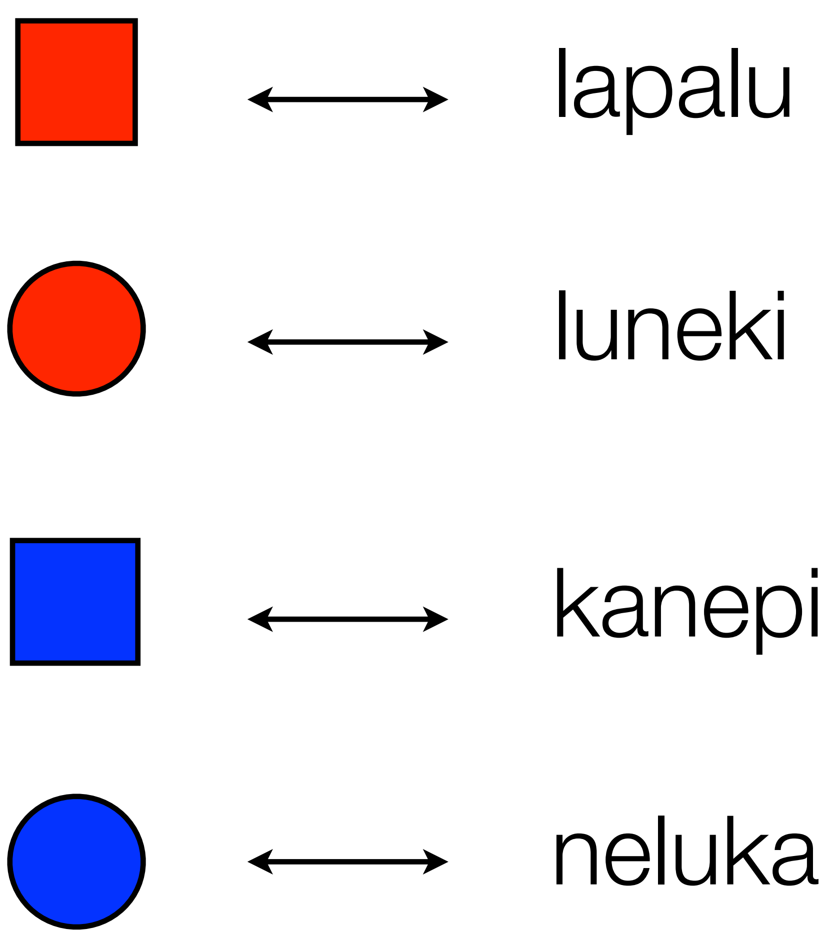

This mapping between meanings and signals does not preserve structure from one domain to the other. We call this a **holistic** language, and it’s equivalent to what we’ve been looking at in the course so far. It’s basically just a vocabulary.

Holistic vs. Compositional

- It’s not the structure in meanings/signals that matters, but whether that structure is utilised by the mapping

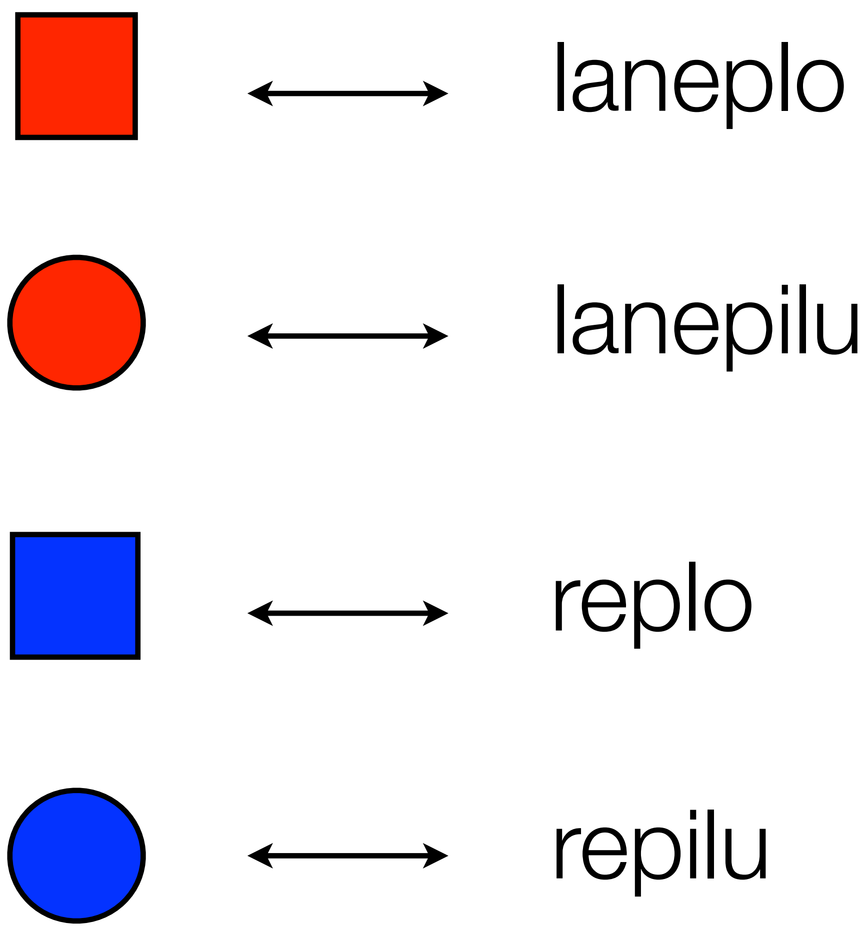

This mapping between meanings and signals *does* preserve structure from one domain to the other. We call this a **compositional** language. On a rudimentary level, it exhibits morphosyntactic properties. It enables generalisation to new meanings.

Simulating the transition to syntax

- Starting in the late ’90s, there were a series of simulations using different techniques to try and understand the cultural evolution of syntactic structure (see e.g. Kirby & Hurford 2002)
- Tricky requirement:
    - we need a learning model that is capable of detecting, and using, syntactic structure when it is there in the data,
    - but we don’t want to simply impose syntactic structure from the outset.
- We need a learner that is happy with either **holistic** or **compositional** languages

Example model (Kirby, 2002)

- Meanings are simple predicate logic expressions. e.g.:
  loves(mary, john)
  thinks(mary, likes(john, heather)

- There are 5 different individuals, 5 simple predicates, and 5 predicates of propositional attitude in the agents' world

- Signals are simply strings of random characters from the alphabet. e.g.:
  agjds
  gfhiyjilkq
  marylovesjohn

Learning

- Agents attempt to induce a simple grammar that covers the meaning-signal pairs that they hear

- Fundamental principle: Learning is compression

  - Learners try and fit the data heard, but also generalise by compressing their grammar (cf. Occam's Razor)

  - Learning is a trade-off between fit to data and generalisation

Two steps to learning

- Incorporation (for each utterance heard)

S/**loves(john, mary)** → johnlovesmary

- Generalisation (whenever possible, within certain heuristic constraints)

S/**loves(peter, mary)** → peterlovesmary
S/**loves(john, mary)** → johnlovesmary

↯

S/**loves(*x*, mary)** → C/*x* lovesmary
C/**john** → john
C/**peter** → peter

A diffusion chain

1. Start with one learner and one adult speaker neither of which have grammars.
2. Choose a meaning at random.
3. Get speaker to produce signal for that meaning (may need to “invent” random string).
4. Give meaning-signal pair to learner.
5. Repeat 2-4 one hundred and fifty times.
6. Delete speaker.
7. Make learner be the new speaker.
8. Introduce a new learner (with no initial grammar)
9. Repeat 2-8 thousands of times.

Results 1: initial stages

- Initially, speakers have no language, so “invent” random strings of characters

- A *protolanguage* emerges for some meanings, but no structure. These are *holistic* expressions:

`ldg`	“Mary admires John”
`xkq`	“Mary loves John”
`gj`	“Mary admires Gavin”
`axk`	“John admires Gavin”
`gb`	“John knows that Mary knows that John admires Gavin”

Big complex grammar
but low expressivity

S/loves(john,mary) $\rightarrow$ `sdx`

S/admires(mary,gavin) $\rightarrow$ `gj`

S/admires(john,gavin) $\rightarrow$ `axk`

S/admires(gavin,heather) $\rightarrow$ `nui`

S/loves(john,heather) $\rightarrow$ `my`

S/loves(mary,john) $\rightarrow$ `xkq`

S/admires(mary,john) $\rightarrow$ `ldg`

S/thinks(john,loves(mary,gavin)) $\rightarrow$ `fi`

S/thinks(heather,loves(heather,gavin)) $\rightarrow$ `ad`

S/thinks(john,admires(heather,gavin)) $\rightarrow$ `xuy`

S/knows(gavin,loves(gavin,mary)) $\rightarrow$ `k`

S/knows(gavin,loves(john,mary)) $\rightarrow$ `ysw`

S/thinks(mary,knows(gavin,loves(heather,john))) $\rightarrow$ `pq`

S/thinks(mary,knows(heather,loves(heather,john))) $\rightarrow$ `rr`

S/knows(john,knows(mary,admires(mary,john))) $\rightarrow$ `lr`

... (plus another 101 rules)

Results 2: many generations later...

gjhftejm	“Mary admires John”
gjhftejwp	“Mary loves John”
gjqpftejm	“Mary admires Gavin”
gjqpfhm	“John admires Gavin”
ihuitejugjqpfhm	“John knows that Mary knows that John admires Gavin”

Results 2: many generations later...

gj h f tej m
John Mary admires
"Mary admires John"

gj h f tej wp
John Mary loves
"Mary loves John"

gj qp f tej m
Gavin Mary admires
"Mary admires Gavin"

gj qp f h m
Gavin John admires
"John admires Gavin"

i h u i tej u gj qp f h m
John knows Mary knows Gavin John admires
"John knows that Mary knows that John admires Gavin"

$$S/p(x,y) \rightarrow \texttt{gj}\ A/y\ \texttt{f}\ A/x\ \ B/p$$
$$S/p(x,q) \rightarrow \texttt{i}\ A/x\ \ D/p\ \ S/q$$
$$A/\text{heather} \rightarrow \texttt{dl}$$
$$A/\text{mary} \rightarrow \texttt{tej}$$
$$A/\text{pete} \rightarrow \texttt{n}$$
$$A/\text{gavin} \rightarrow \texttt{qp}$$
$$A/\text{john} \rightarrow \texttt{h}$$
$$B/\text{detests} \rightarrow \texttt{b}$$
$$B/\text{loves} \rightarrow \texttt{wp}$$
$$B/\text{hates} \rightarrow \texttt{c}$$
$$B/\text{likes} \rightarrow \texttt{e}$$
$$B/\text{admires} \rightarrow \texttt{m}$$
$$D/\text{believes} \rightarrow \texttt{g}$$
$$D/\text{knows} \rightarrow \texttt{u}$$
$$D/\text{decides} \rightarrow \texttt{ipr}$$
$$D/\text{says} \rightarrow \texttt{p}$$
$$D/\text{thinks} \rightarrow \texttt{m}$$

Small, simple grammar
infinite expressivity

Quantitative results: languages evolve

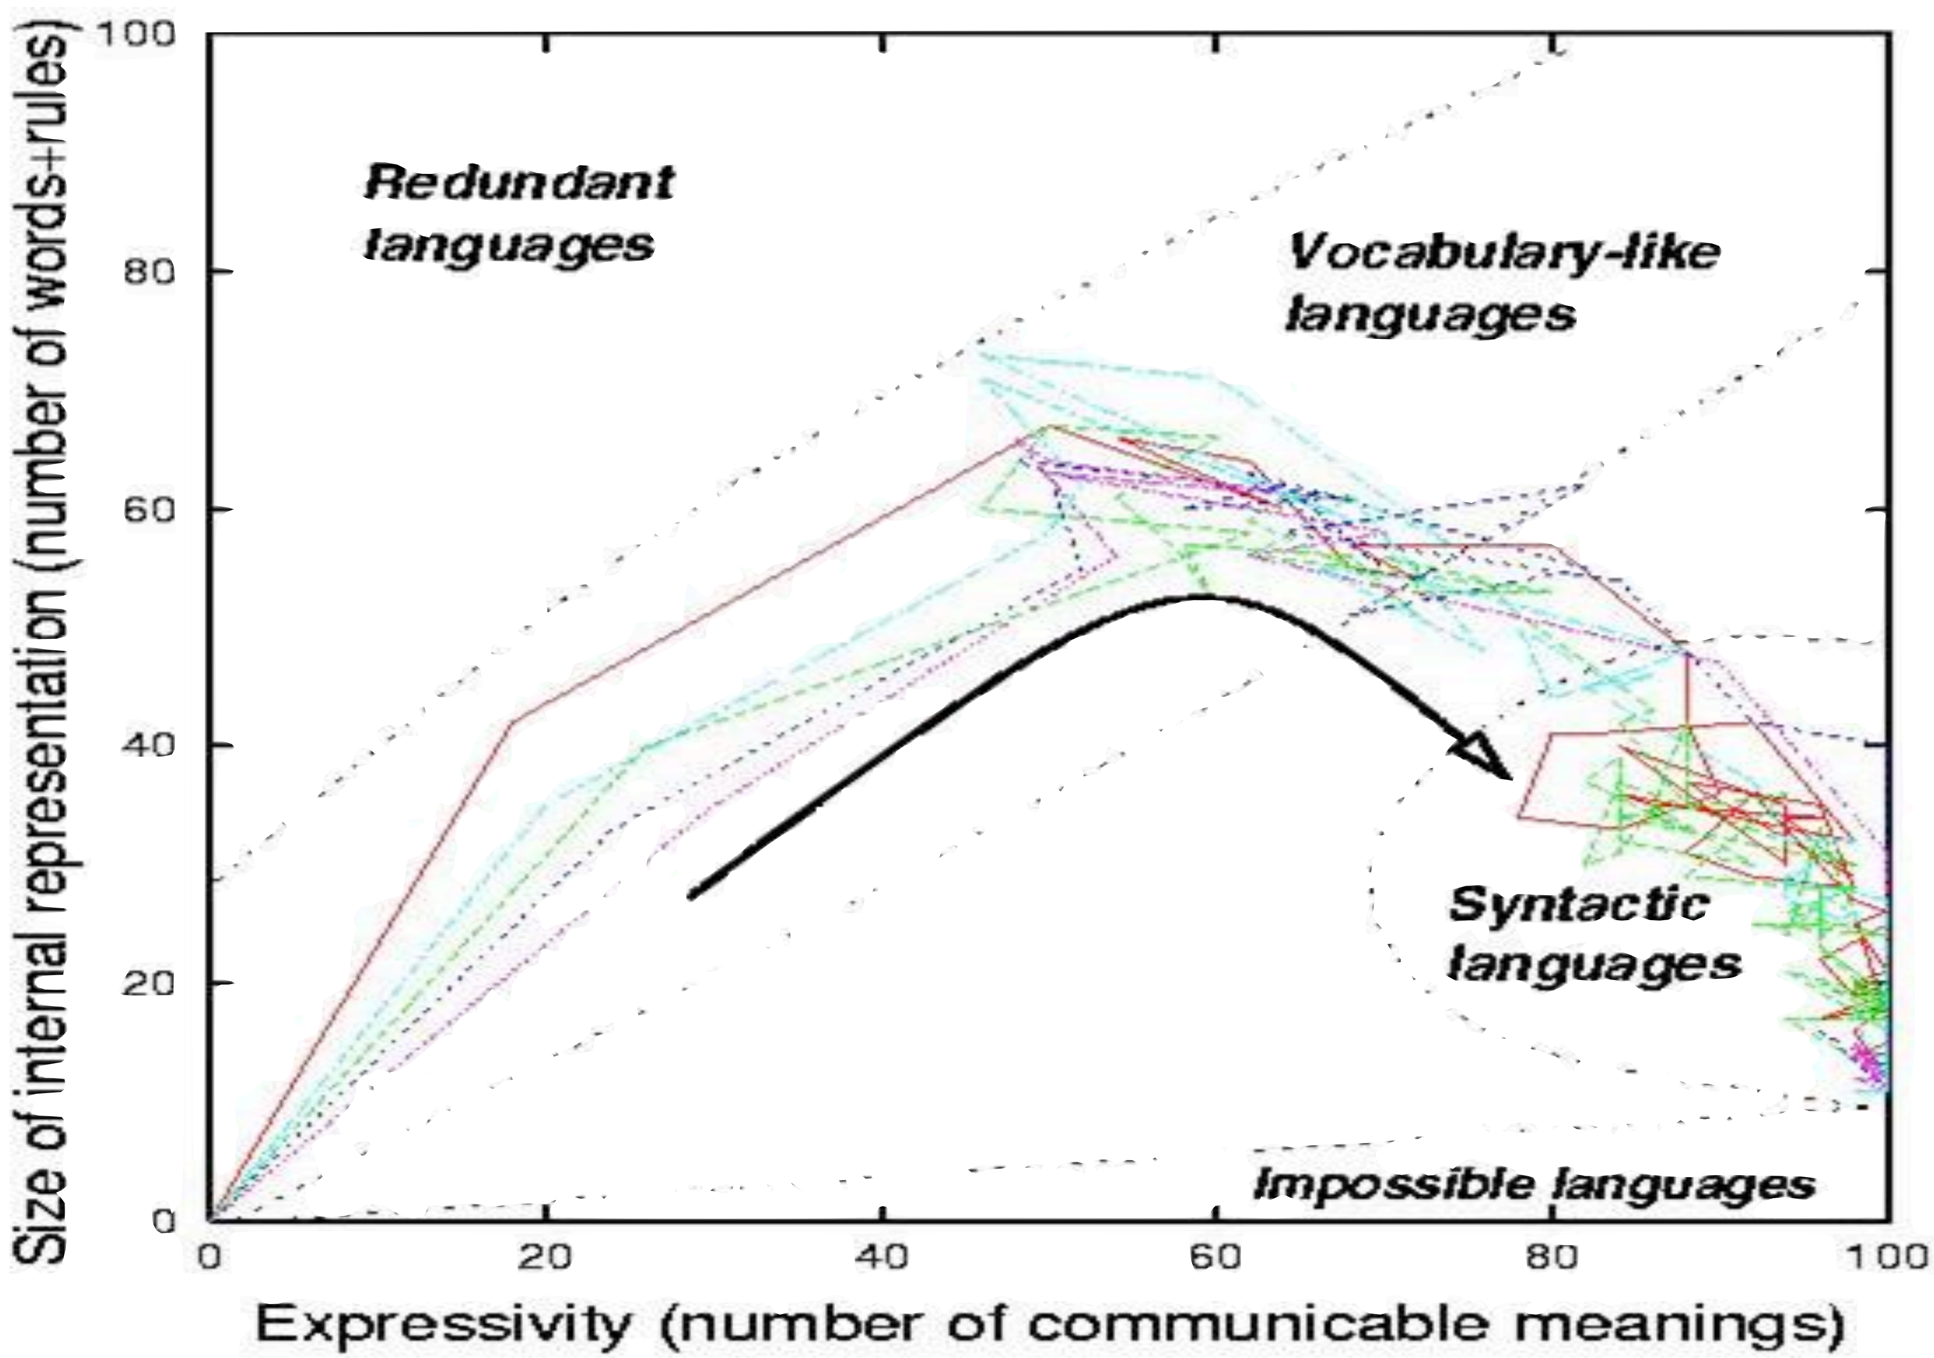

What's going on?

- There's no biological evolution in this iterated learning model
- There isn't even any communication or notion of function in model at all.
- So, why are structured languages evolving?
- **Languages themselves are evolving to the conditions of the iterated learning process in order that they are learnable.**
- The agents never see all the meanings…
- Only languages that are *generalisable* from limited exposure are stable.

Language has to fit through a narrow *bottleneck*

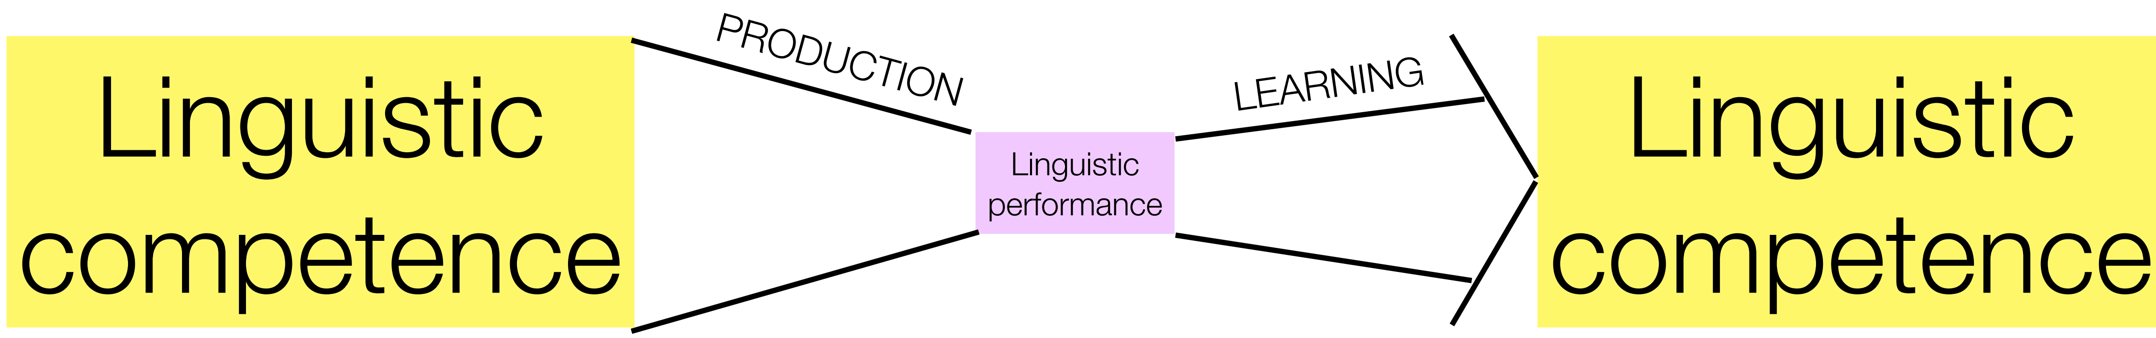

- This has profound implications for the structure of language
- Language becomes generalisable from a limited subset of utterances:
    - When meanings are structured, signals become structured
    - *Generalisable* equates to *compositional* in this case
- Syntax is an adaptive response **by language** (arising from cultural evolution) to the problem of getting through this bottleneck

From simulations to experiments

- Since running these simulations (and many more like them) there have been criticisms that this process is implausible

- Is it really likely that random ‘mistakes’ could lead us from a holistic protolanguage to a compositional syntax? Do these learning algorithms really reflect the human language learning biases?

- I’ll come back to this in the final lecture on the course, when I talk about **human simulation.**

- **Next week’s lecture:** ruining all this with Bayes (and then fixing it again)

Readings for this lecture

- Kirby, S. (2002). Learning, bottlenecks and the evolution of recursive syntax. In E. J. Briscoe (Ed.), *Linguistic Evolution through Language Acquisition: Formal and Computational Models* (pp. 173-204). Cambridge: Cambridge University Press.

- Kirby, S. & Hurford, J. (2002) The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the Evolution of Language* (pp. 121-148). London: Springer Verlag.