Recognition vs. processing: exploring Phillips and Lewis (2013)

Seminar week #2: Non-information-density approaches to online processing, Winter 2015/2016

Asad Sayeed

Uni-Saarland

Today's agenda

- Fixing the time for the seminar.
- Presentation assignments.
- Phillips and Lewis (2013) – presentation and discussion.
  - Focus mostly on the methodological aspects.

So...

I have some questions for the audience. (Not again.)

What is language? (Yes, again. *sigh*)

Last time, we talked about competence and performance.

Generative linguistics: Chomsky and beyond.

- Divide language into (roughly speaking):
  - "Competence": the abstract knowledge of linguistic structure and the processes required to assemble it.
  - "Performance": the articulatory and perceptual vicissitudes involved in "producing" and "consuming" language.

But these are implicitly held to be specific to the human organism.

Which in turn leads to. . .

. . . a central question: what sort of relationship can we have between a "competence" theory and the observations of linguistic behaviour?

- We need a "linking theory".
- For a competence theory, the real question:
  - Words arrive in a sequence.
  - But the rules that define possible utterances (via formal representations) are not necessarily sequential.
  - Can we make a minimal inference between these facts?

Yes: exploit the derivation.

1960s and 1970s: the Derivational Theory of Complexity (DTC).

- Reliant on now-outdated notions of "deep structure" (DS) and "surface structure." (SS)
- Difficulty in processing $\approx$ derivational distance between DS and SS.

Yes: exploit the derivation.

1960s and 1970s: the Derivational Theory of Complexity (DTC).

- Reliant on now-outdated notions of "deep structure" (DS) and "surface structure." (SS)
- Difficulty in processing $\approx$ derivational distance between DS and SS.
- Alleged to have failed.
  - Some derivations at the time didn't correlate well to processing measures.

(Just to give a flavour of the DTC.)

A thing that probably worked: passive constructions.

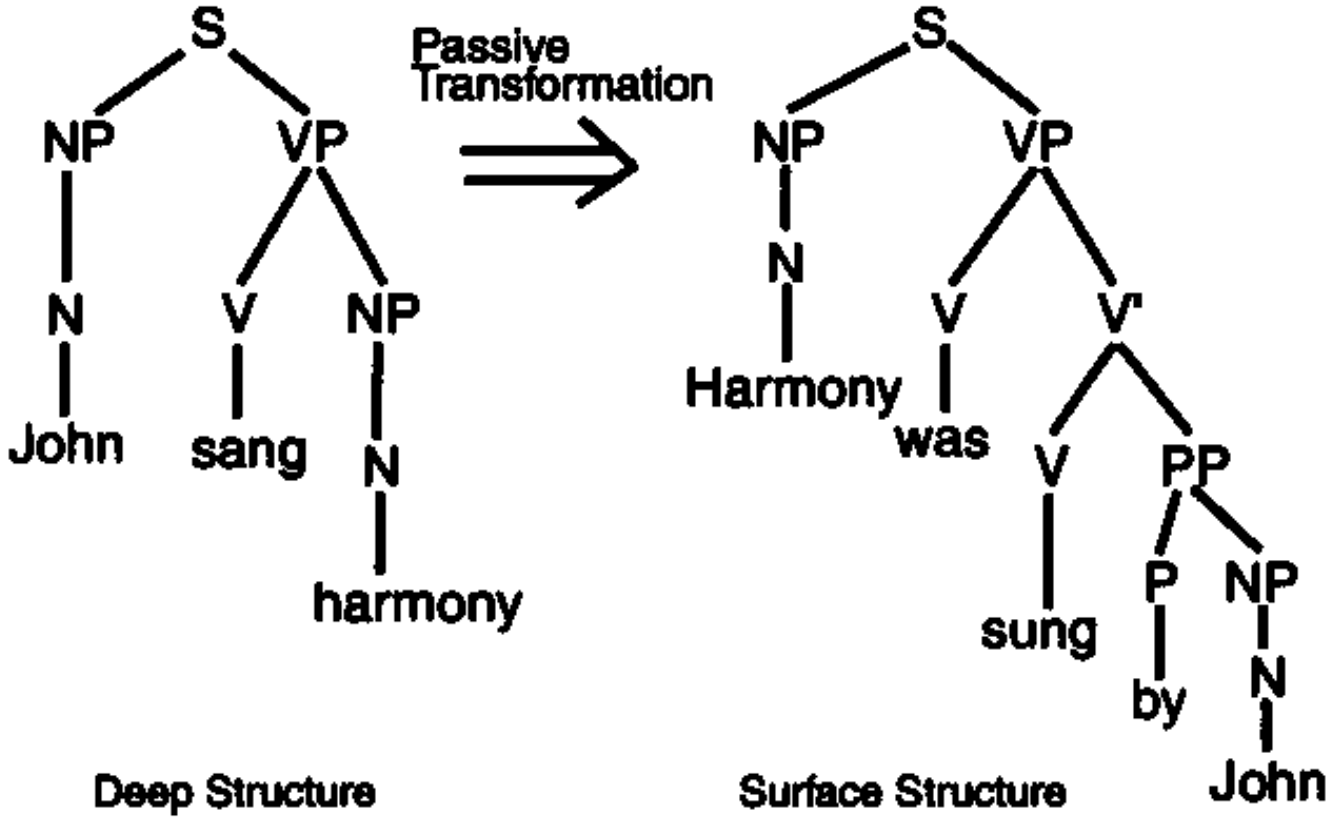

But what are we trying to accomplish here?

Want to resolve mismatch between derivational order and linear order.

But what are we trying to accomplish here?

Want to resolve mismatch between derivational order and linear order.

But, uh, if there is a conflict, isn't linear order the "real" one?

But what are we trying to accomplish here?

Want to resolve mismatch between derivational order and linear order.

But, uh, if there is a conflict, isn't linear order the "real" one?

- After all, it's what people DO.
- And what is language without people to "DO" language?

But what are we trying to accomplish here?

Want to resolve mismatch between derivational order and linear order.

But, uh, if there is a conflict, isn't linear order the "real" one?

- After all, it's what people DO.
- And what is language without people to "DO" language?

Then when "linear" and "derivational" diverge...

. . . shouldn't the "most" linear theory "win"?

Not so fast! Now we finally get to Phillips and Lewis (2013).

- The question is: What is a linguistic theory *for*?

. . . shouldn't the "most" linear theory "win"?

Not so fast! Now we finally get to Phillips and Lewis (2013).

- The question is: What is a linguistic theory *for*?
- Approach from generative linguistics:
  - What does it mean for a speaker to "know" a language — what do they "know"?
  - It's only an **assumption** that what they "know" is how to put a sentence together "left-to-right".

. . . shouldn't the "most" linear theory "win"?

Not so fast! Now we finally get to Phillips and Lewis (2013).

- The question is: What is a linguistic theory *for*?
- Approach from generative linguistics:
  - What does it mean for a speaker to "know" a language — what do they "know"?
  - It's only an **assumption** that what they "know" is how to put a sentence together "left-to-right".

There's a more basic question. . .

What is a linguistic object?

Well, we know it when we see it.

An implicit claim: language as set.

- What humans acquire as children: ability to distinguish between strings.
  - "Grammatical" utterance: belongs to the language you learned.
  - "Ungrammatical" utterance: does not belong to the language you learned.
- That sentences (and the criterion to decide language membership) are highly complex. . .

. . . does not permit us to assume incrementality.

So what *is* the relationship of processing to grammar?

. . . does not permit us to assume incrementality.

So what *is* the relationship of processing to grammar? Philips and Lewis give us three possible positions:

- **Literalism** – formal derivations are temporally related somehow to the actual structure building systems.
  - What a lot of "beginners" expect.
  - So yeah, we don't have little trees floating in our head, but something is happening that is "tree-like".
  - Not common in literature.

. . . does not permit us to assume incrementality.

So what *is* the relationship of processing to grammar? Philips and Lewis give us three possible positions:

- **Formalism** – formal derivations are related to actual structure-building systems, but not temporally.
  - Not a common position at all.
  - Cognitive system can construct derivations, use them in comprehension. (strategically?)

. . . does not permit us to assume incrementality.

So what *is* the relationship of processing to grammar? Philips and Lewis give us three possible positions:

- **Extensionalism** – grammar is just an abstraction, representing "all and only" well-formed sentences of language.
  - Most practicing generative linguists/syntacticians assume this.
  - Limits the accountability of linguistic theory.
  - Actual mental left-to-right construction process? Not a "must have".

But there is a problem.

Extension leaves the grammatical system as a purely abstract object.

But there is a problem.

Extension leaves the grammatical system as a purely abstract object.

- Thin empirical basis – acceptability judgements only.

But there is a problem.

Extension leaves the grammatical system as a purely abstract object.

- Thin empirical basis – acceptability judgements only.
- What constrains derivations?
  - Generative linguists often appeal to notions of parsimony, "efficiency", "economy".
  - But where do these come from if not cognitive restrictions?

But there is a problem.

Extension leaves the grammatical system as a purely abstract object.

- Thin empirical basis – acceptability judgements only.
- What constrains derivations?
  - Generative linguists often appeal to notions of parsimony, "efficiency", "economy".
  - But where do these come from if not cognitive restrictions?
- How would children learn this "partitioning function" if it had no measurable relationship to performative considerations?

On the other hand, it's liberating.

Sort of.

- One could just argue that these are different levels of representation.
  - Levels of analysis – computational, algorithmic, and implementational.
  - Well-established in cognitive science (Marr's levels, 1982)
- Grammar merely belongs to the computational level, explaining "real" behaviour is someone else's job.
- Would we demand that everything be explained in terms of neurons?

So maybe it's just good strategy.

Phillips and Lewis: "strategic extensionalism".

- Set-partitioning of sentence into grammatical and ungrammatical: an "interim" goal.
- But ultimately we want to organize the explanation in a manner that connects to the cognitive mechanism.

So maybe it's just good strategy.

Phillips and Lewis: "strategic extensionalism".

- Set-partitioning of sentence into grammatical and ungrammatical: an "interim" goal.
- But ultimately we want to organize the explanation in a manner that connects to the cognitive mechanism.

But "principled extensionalism" still not nonsense: "what" and "how" might just be separate questions.

Implementation independence is crucial.

Implementation independence:

- The same system can be implemented in different ways without changing the abstract system.
- We don't have to change the grammatical theory just because we find that people parse a particular way.

Implementation dependence:

- "only ever realized one way" at lower level.

Implementation independence is crucial.

Implementation independence:

- The same system can be implemented in different ways without changing the abstract system.
- We don't have to change the grammatical theory just because we find that people parse a particular way.

Implementation dependence:

- "only ever realized one way" at lower level.

Common implicit assumption: generative grammar is independent.

But that's an empirical question.

Independence is also a common belief in AI:

- We still can't simulate human "wetware" accurately.
- But so what? A formal description of human capability (e.g. language) should suffice.

Aside from the fact that no one has ever built this machine, it begs the question.

We want to understand the human system.

Even if we can "port" human language to another "platform":

- It "arose" in the human system.
- It's structure may still be dependent on neurobiological and performative considerations.
- We can't just assume implementation independence.

So how do we assess implementation-dependence?

We need to know whether speakers can use multiple ways to construct the same interpretations.

So how do we assess implementation-dependence?

We need to know whether speakers can use multiple ways to construct the same interpretations.

- There's not much evidence that speakers have those multiple ways.
  - Speaking and understanding have common goals.
  - Large amount of evidence for incremental construction of representations.
  - Reanalysis: the "human parser" seems to go back and repair errors, e.g. garden paths.

But if we accept implementation-dependence. . .

. . . then what do we need for a theory of "real-time" grammar?

But if we accept implementation-dependence. . .

. . . then what do we need for a theory of "real-time" grammar?Phillips and Lewis:

- First, recognize that the "real-time grammar" is only part of the story.
  - There's of course the whole articulatory mechanism.
  - Interaction with intention, communicative goals, etc.
  - Just because you have a real-time grammar, doesn't mean that language "works".

But if we accept implementation-dependence. . .

. . . then what do we need for a theory of "real-time" grammar?Phillips and Lewis:

- The grammar does not "live" in a "perfect" real-life parsing machine.
  - This is where the "grimy mirror" I talked about last time comes in.
  - Noisy environments, distractions produce "incorrect" analyses.
  - Listener may not make "full use of the input" – but constructs well-formed representations.

But if we accept implementation-dependence. . .

. . . then what do we need for a theory of "real-time" grammar?Phillips and Lewis:

- The question of ambiguity: there can be multiple well-formed analyses — doesn't mean that the grammar is not "real-time".
- Reductionism: "processing" accounts vs. "grammar" accounts
  - Attempting to take things out of the grammar – perfectly reasonable given evidence.
  - e.g. memory accounts, even UID. (remember my "weak UID hypothesis"?)

But lastly...

Phillips and Lewis decide not to commit to strict incrementality.

- Need only be "roughly" left-to-right.
- "Growing evidence that comprehenders often build structural positions in their parses before encountering the words..."
- (ie, predictive parsing? Prediction at what level...)

And then Phillips and Lewis go through an analysis of various evidence.

Which wasn't my main purpose in assigning this paper.

(Although of course we can discuss it.)

So what was my point, then?

Well, going back to last week's lecture:

So what was my point, then?

Well, going back to last week's lecture:

- Communicative-efficency approaches: have a much more specific idea of what the grammar is *for*.

So what was my point, then?

Well, going back to last week's lecture:

- Communicative-efficency approaches: have a much more specific idea of what the grammar is *for*.
  - Namely: an efficient solution to the linguistic "organism" accomplishing communicative goals.

So what was my point, then?

Well, going back to last week's lecture:

- Communicative-efficency approaches: have a much more specific idea of what the grammar is *for*.
  - Namely: an efficient solution to the linguistic "organism" accomplishing communicative goals.
  - But that contains an implicit idea of the relationship of the grammar to the machine...

But if you DON'T make that assumption. . .

. . . then it's *not* completely obvious that the grammar *should* have a direct relationship to the process of parsing.

But if you DON'T make that assumption. . .

. . . then it's *not* completely obvious that the grammar *should* have a direct relationship to the process of parsing.

- And if it's *not* obvious, then all the issues brought up by Phillips and Lewis make sense.

But if you DON'T make that assumption. . .

. . . then it's *not* completely obvious that the grammar *should* have a direct relationship to the process of parsing.

- And if it's *not* obvious, then all the issues brought up by Phillips and Lewis make sense.
- You can't just assume that the grammar has *anything* to do with the parser – it requires experimental evidence.

And it happens that some of that evidence is not necessarily available from studying language-as-communication.

As we'll see in the rest of the seminar.