

Ling 566

Oct 19, 2021

How the Grammar Works

Midterm feedback: Thank you!

- More polls
- More examples (coming right up!)
- More time between chapter covered in lecture & homework due (getting better?)
- Reading questions: can answer each other?
 - Yes, but still ask your own question too
- More explanation of linguistics terms
- Homework expectations unclear
- It can feel intimidating to ask questions

Overview

- What we're trying to do
- The pieces of our grammar
- Two extended examples
- Reflection on what we've done, what we still have to do
- Reading questions

What We're Trying To Do

- Objectives
 - Develop a theory of knowledge of language
 - Represent linguistic information explicitly enough to distinguish well-formed from ill-formed expressions
 - Be parsimonious, capturing linguistically significant generalizations.
- Why Formalize?
 - To formulate testable predictions
 - To check for consistency
 - To make it possible to get a computer to do it for us

Why does this matter to NLP?

- Understand how language works => better positioned to build technology that works with language
- For some applications, grammar engineering is a valuable component directly
- Grammar engineering can also support extremely detailed annotation

How We Construct Sentences

- The Components of Our Grammar
 - Grammar rules
 - Lexical entries
 - Principles
 - Type hierarchy (very preliminary, so far)
 - Initial symbol (S, for now)
- We combine constraints from these components.
 - Q: What says we have to combine them?

🌐 When poll is active, respond at pollev.com/emb

📱 Text **EMB** to **22333** once to join



W Syntax (so far) helps me:

understand other classes

understand what I'm getting
the computer to do

understand how to evaluate
NLP systems

not very much/not at all

by being interesting

🌐 When poll is active, respond at pollev.com/emb

📱 Text **EMB** to **22333** once to join



W In the future, I think syntax will help me:

understand other classes

understand what I'm getting
the computer to do

understand how to evaluate
NLP systems

not very much/not at all

by being interesting

An Example

A cat slept.


- Can we build this with our tools?
- Given the constraints our grammar puts on well-formed sentences, is this one?

Lexical Entry for *a*

$\langle a, \rangle$	$\left[\begin{array}{l} \text{word} \\ \\ \text{SYN} \\ \\ \text{SEM} \end{array} \right.$	$\left[\begin{array}{l} \text{HEAD} \\ \\ \text{VAL} \\ \\ \text{MODE} \\ \text{INDEX} \\ \text{RESTR} \end{array} \right.$	$\left[\begin{array}{l} \text{det} \\ \text{AGR} \\ \text{COUNT} \\ \\ \text{COMPS} \\ \text{SPR} \\ \text{MOD} \\ \\ \text{none} \\ j \\ \left\langle \left[\begin{array}{l} \text{RELN} \\ \text{BV} \end{array} \right] \right\rangle \end{array} \right.$	$\left[\begin{array}{l} \\ \\ 3sing \\ + \\ \\ \langle \rangle \\ \langle \rangle \\ \langle \rangle \\ \\ \\ a \\ j \end{array} \right.$	\rangle


- Is this a fully specified description?
- What features are unspecified?
- How many word structures can this entry license?

Lexical Entry for *cat*




- Which feature paths are abbreviated?
- Is this a fully specified description?
- What features are unspecified?
- How many word structures can this entry license?


Effect of Principles: the SHAC



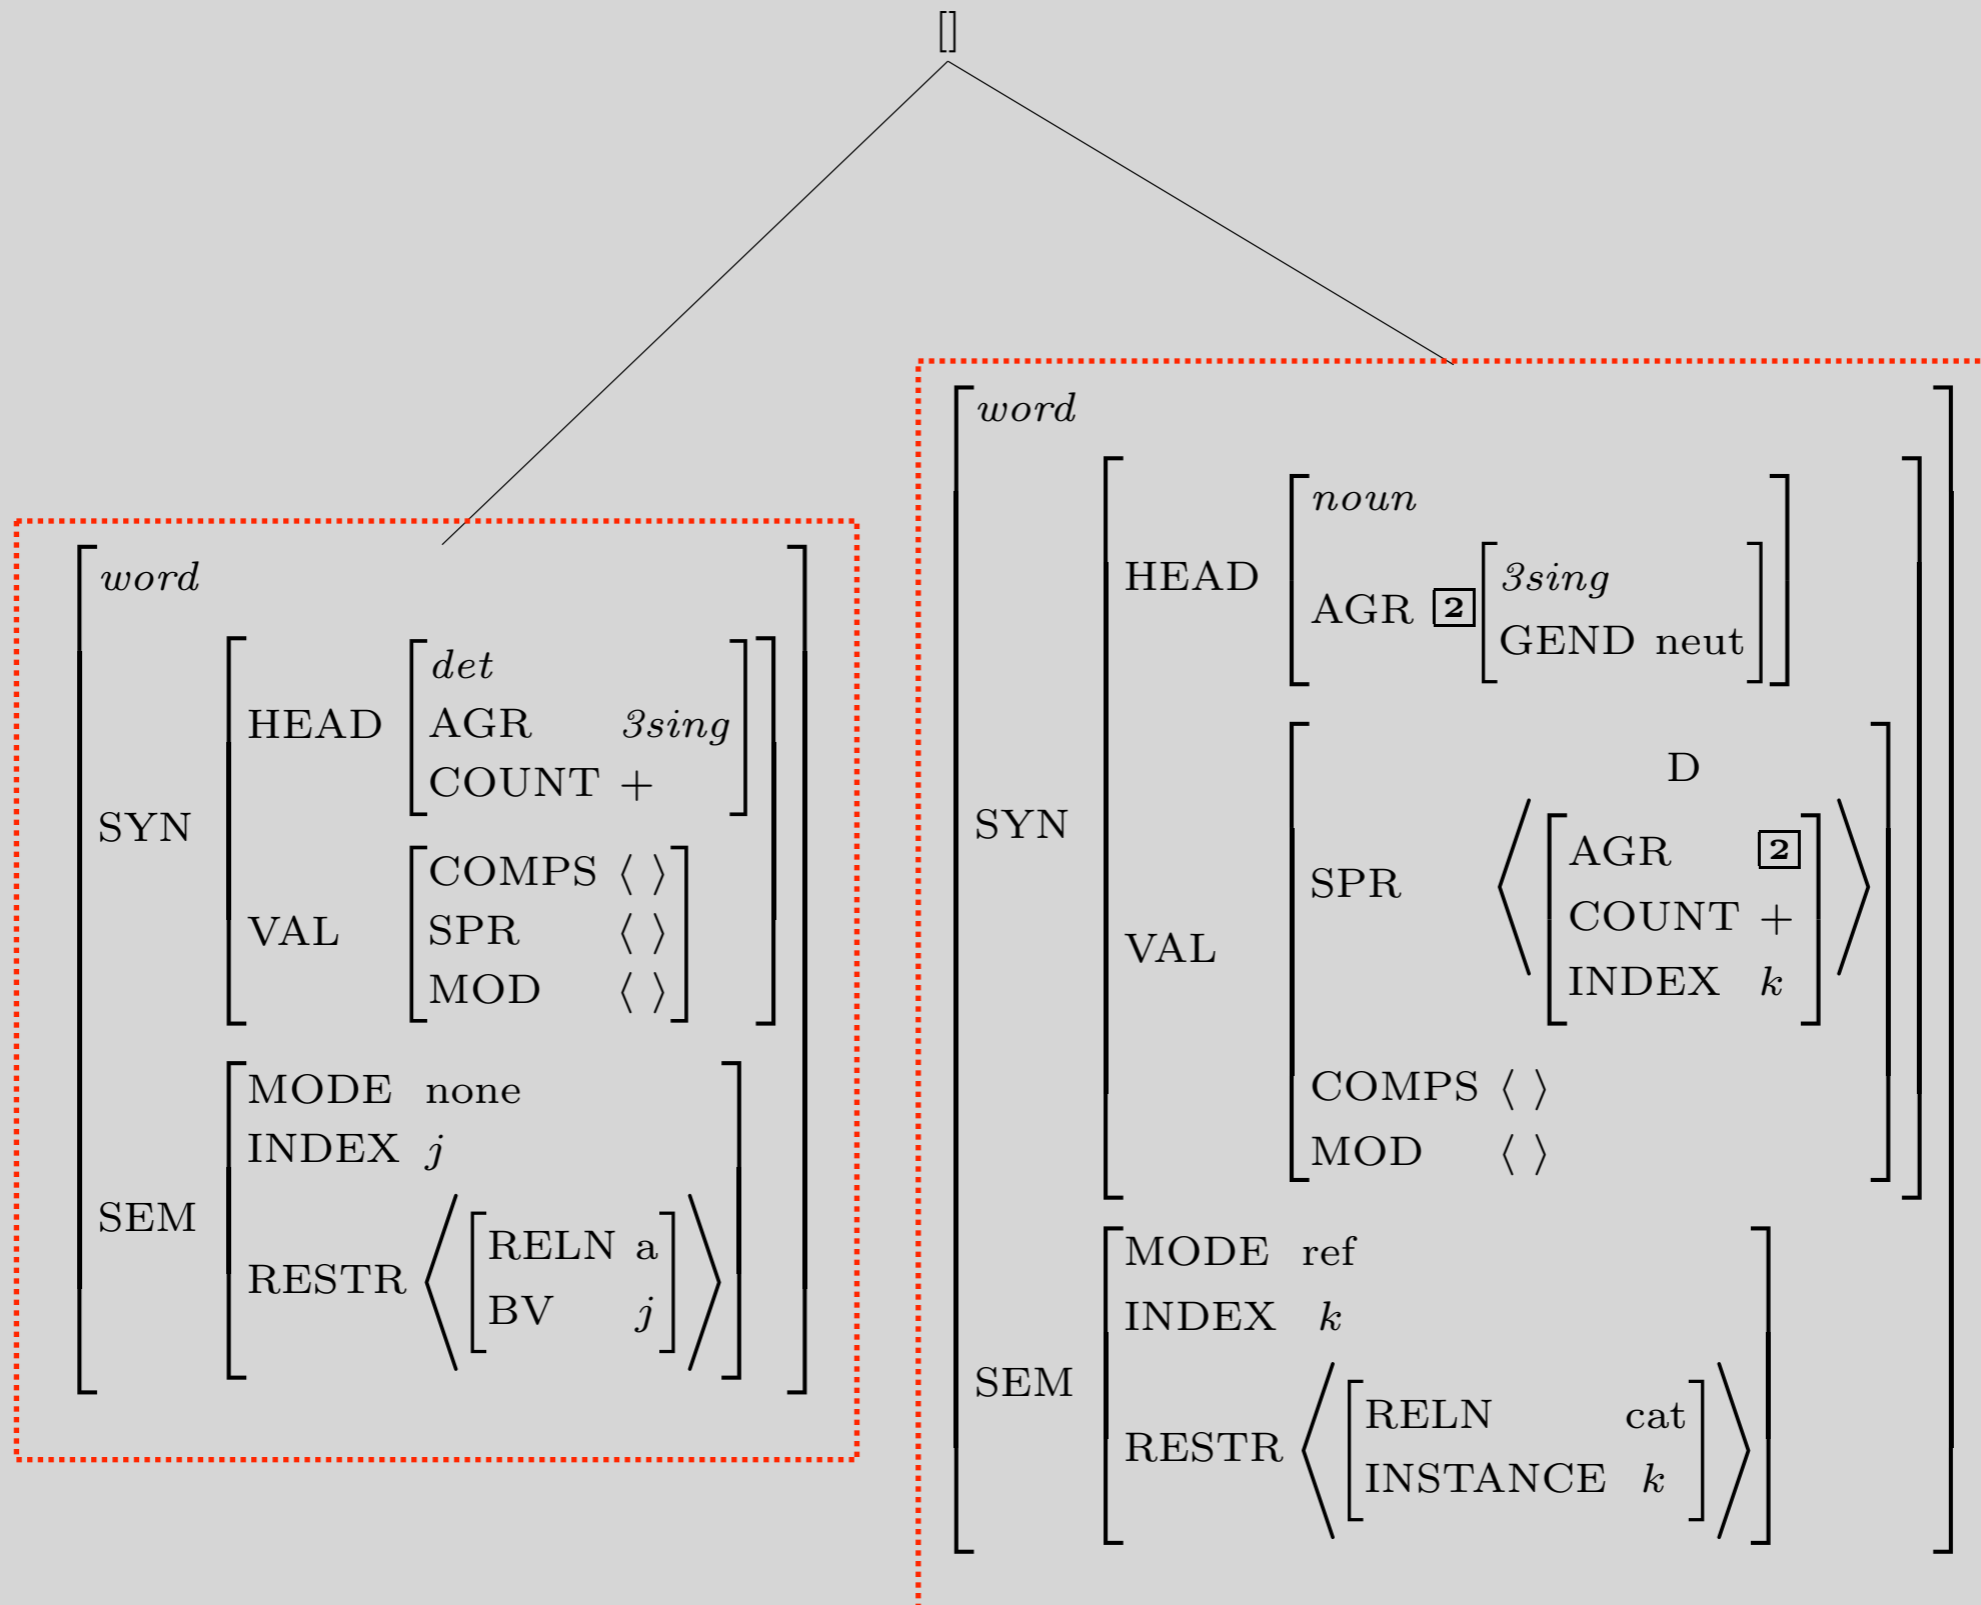
Description of Word Structures for *cat*




Building a Phrase



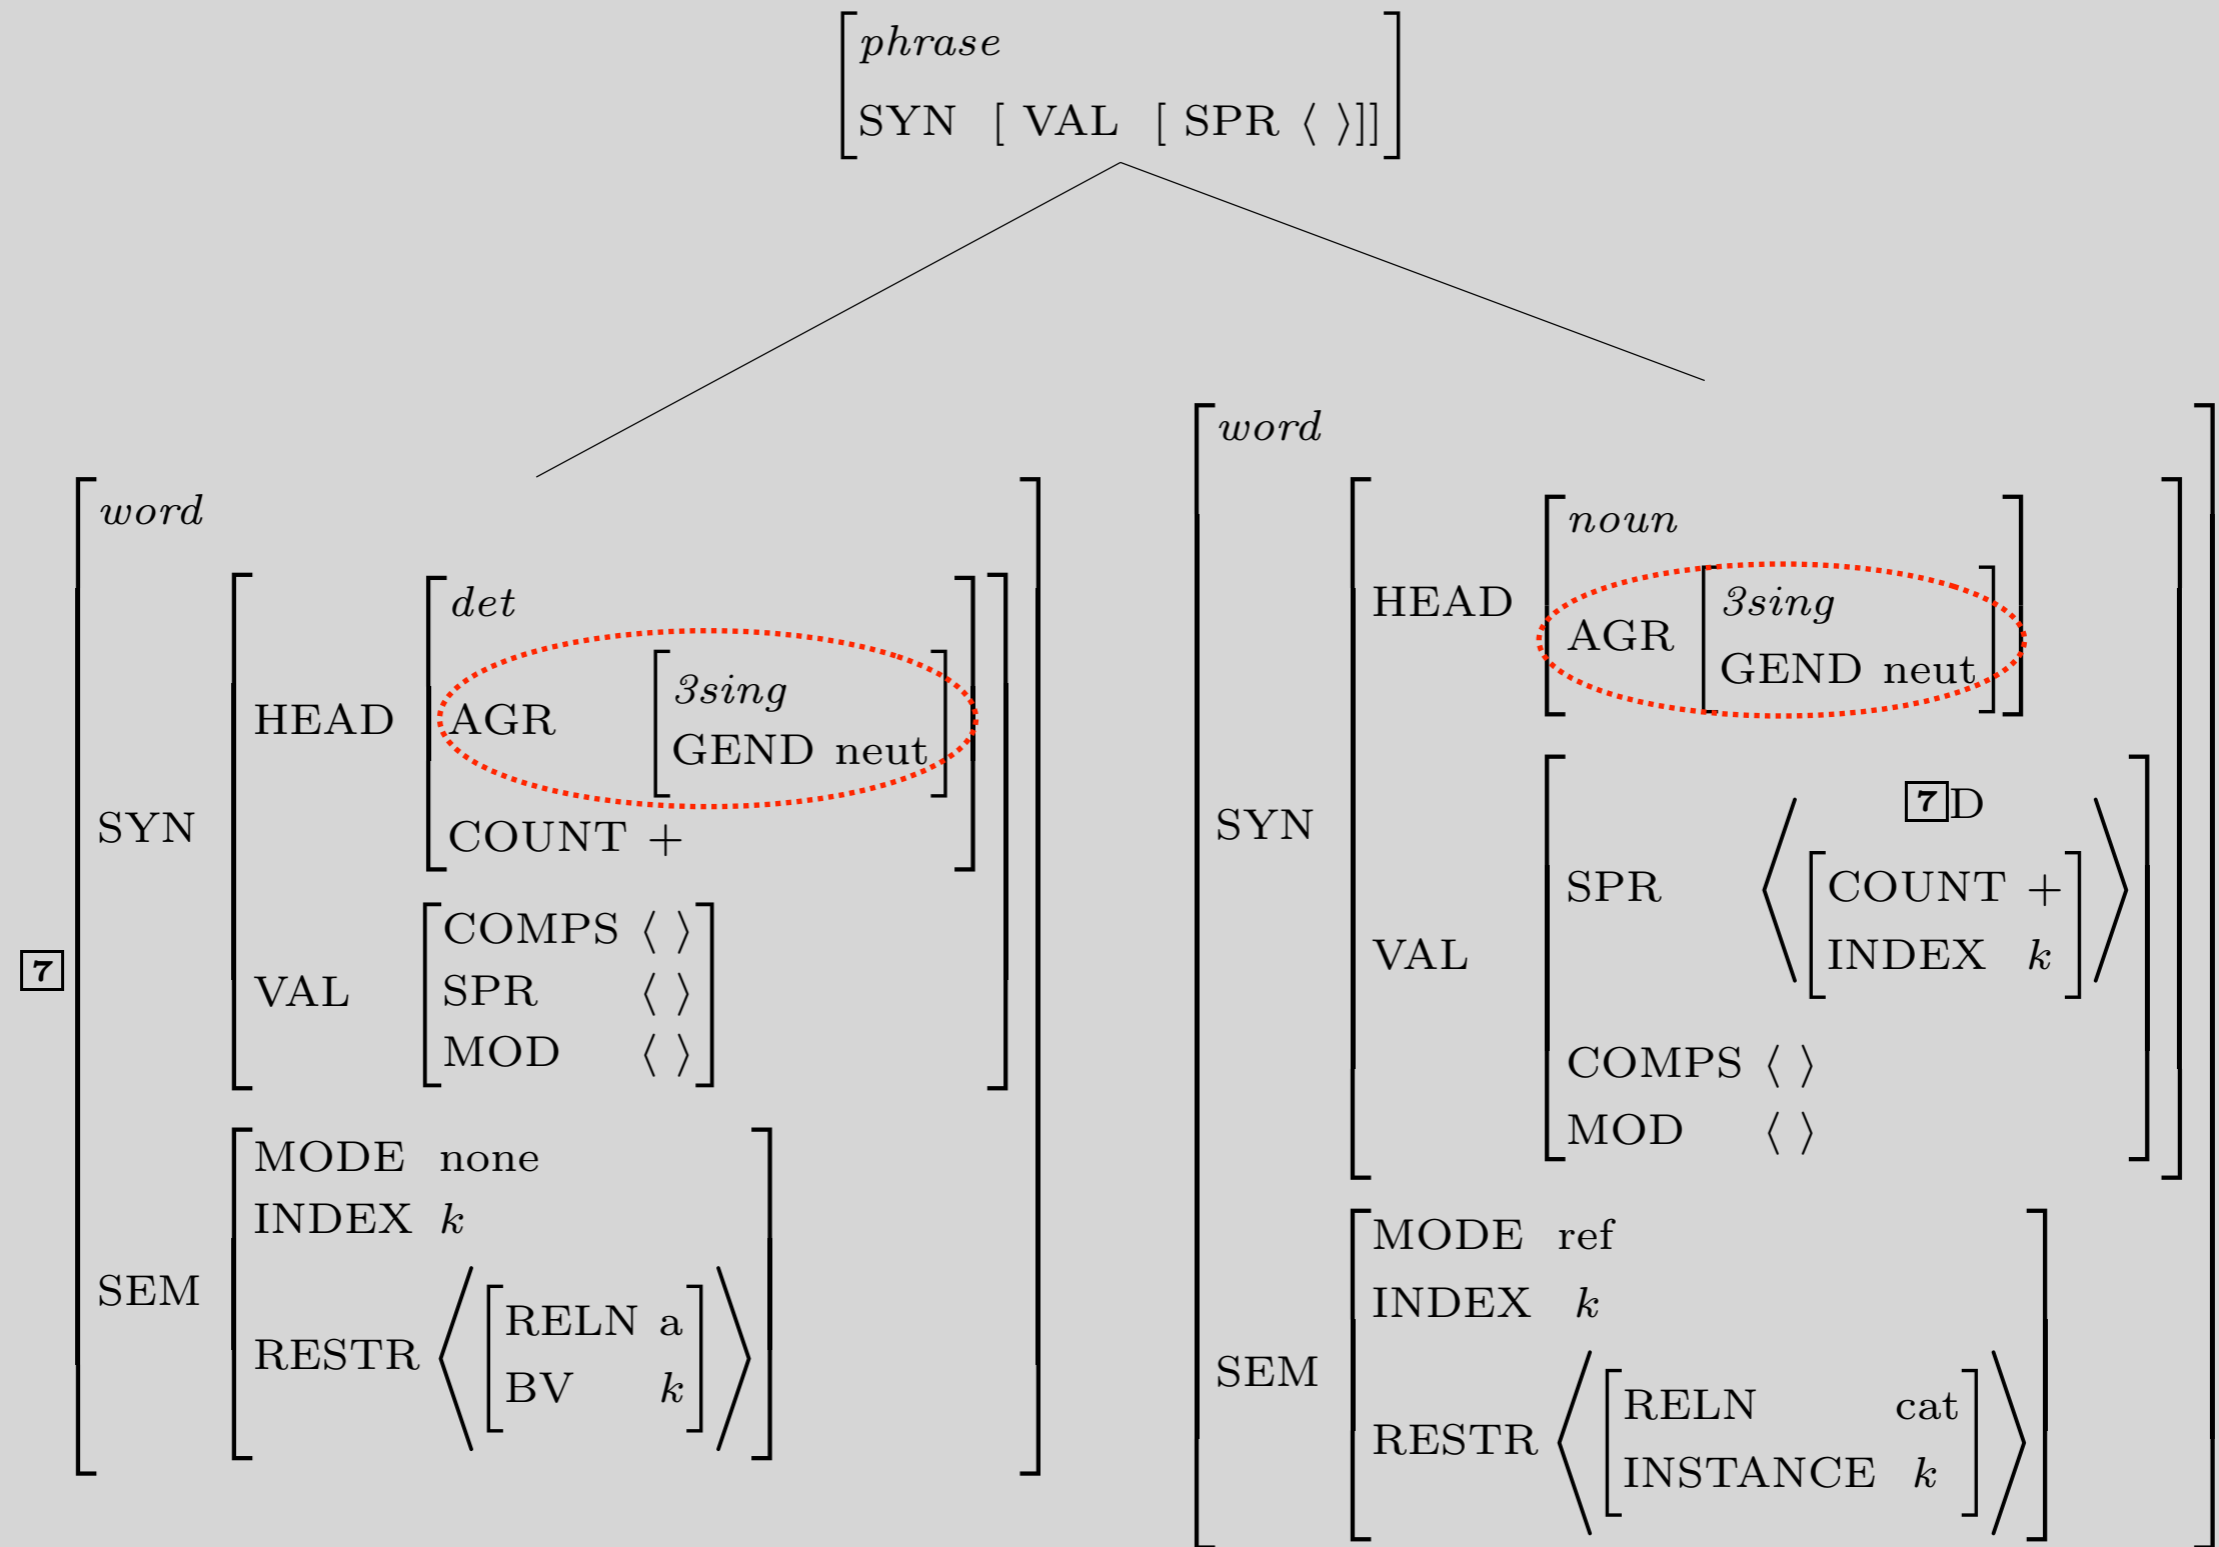
Constraints Contributed by Daughter Subtrees




Constraints Contributed by the Grammar Rule




A Constraint Involving the SHAC




Effects of the Valence Principle



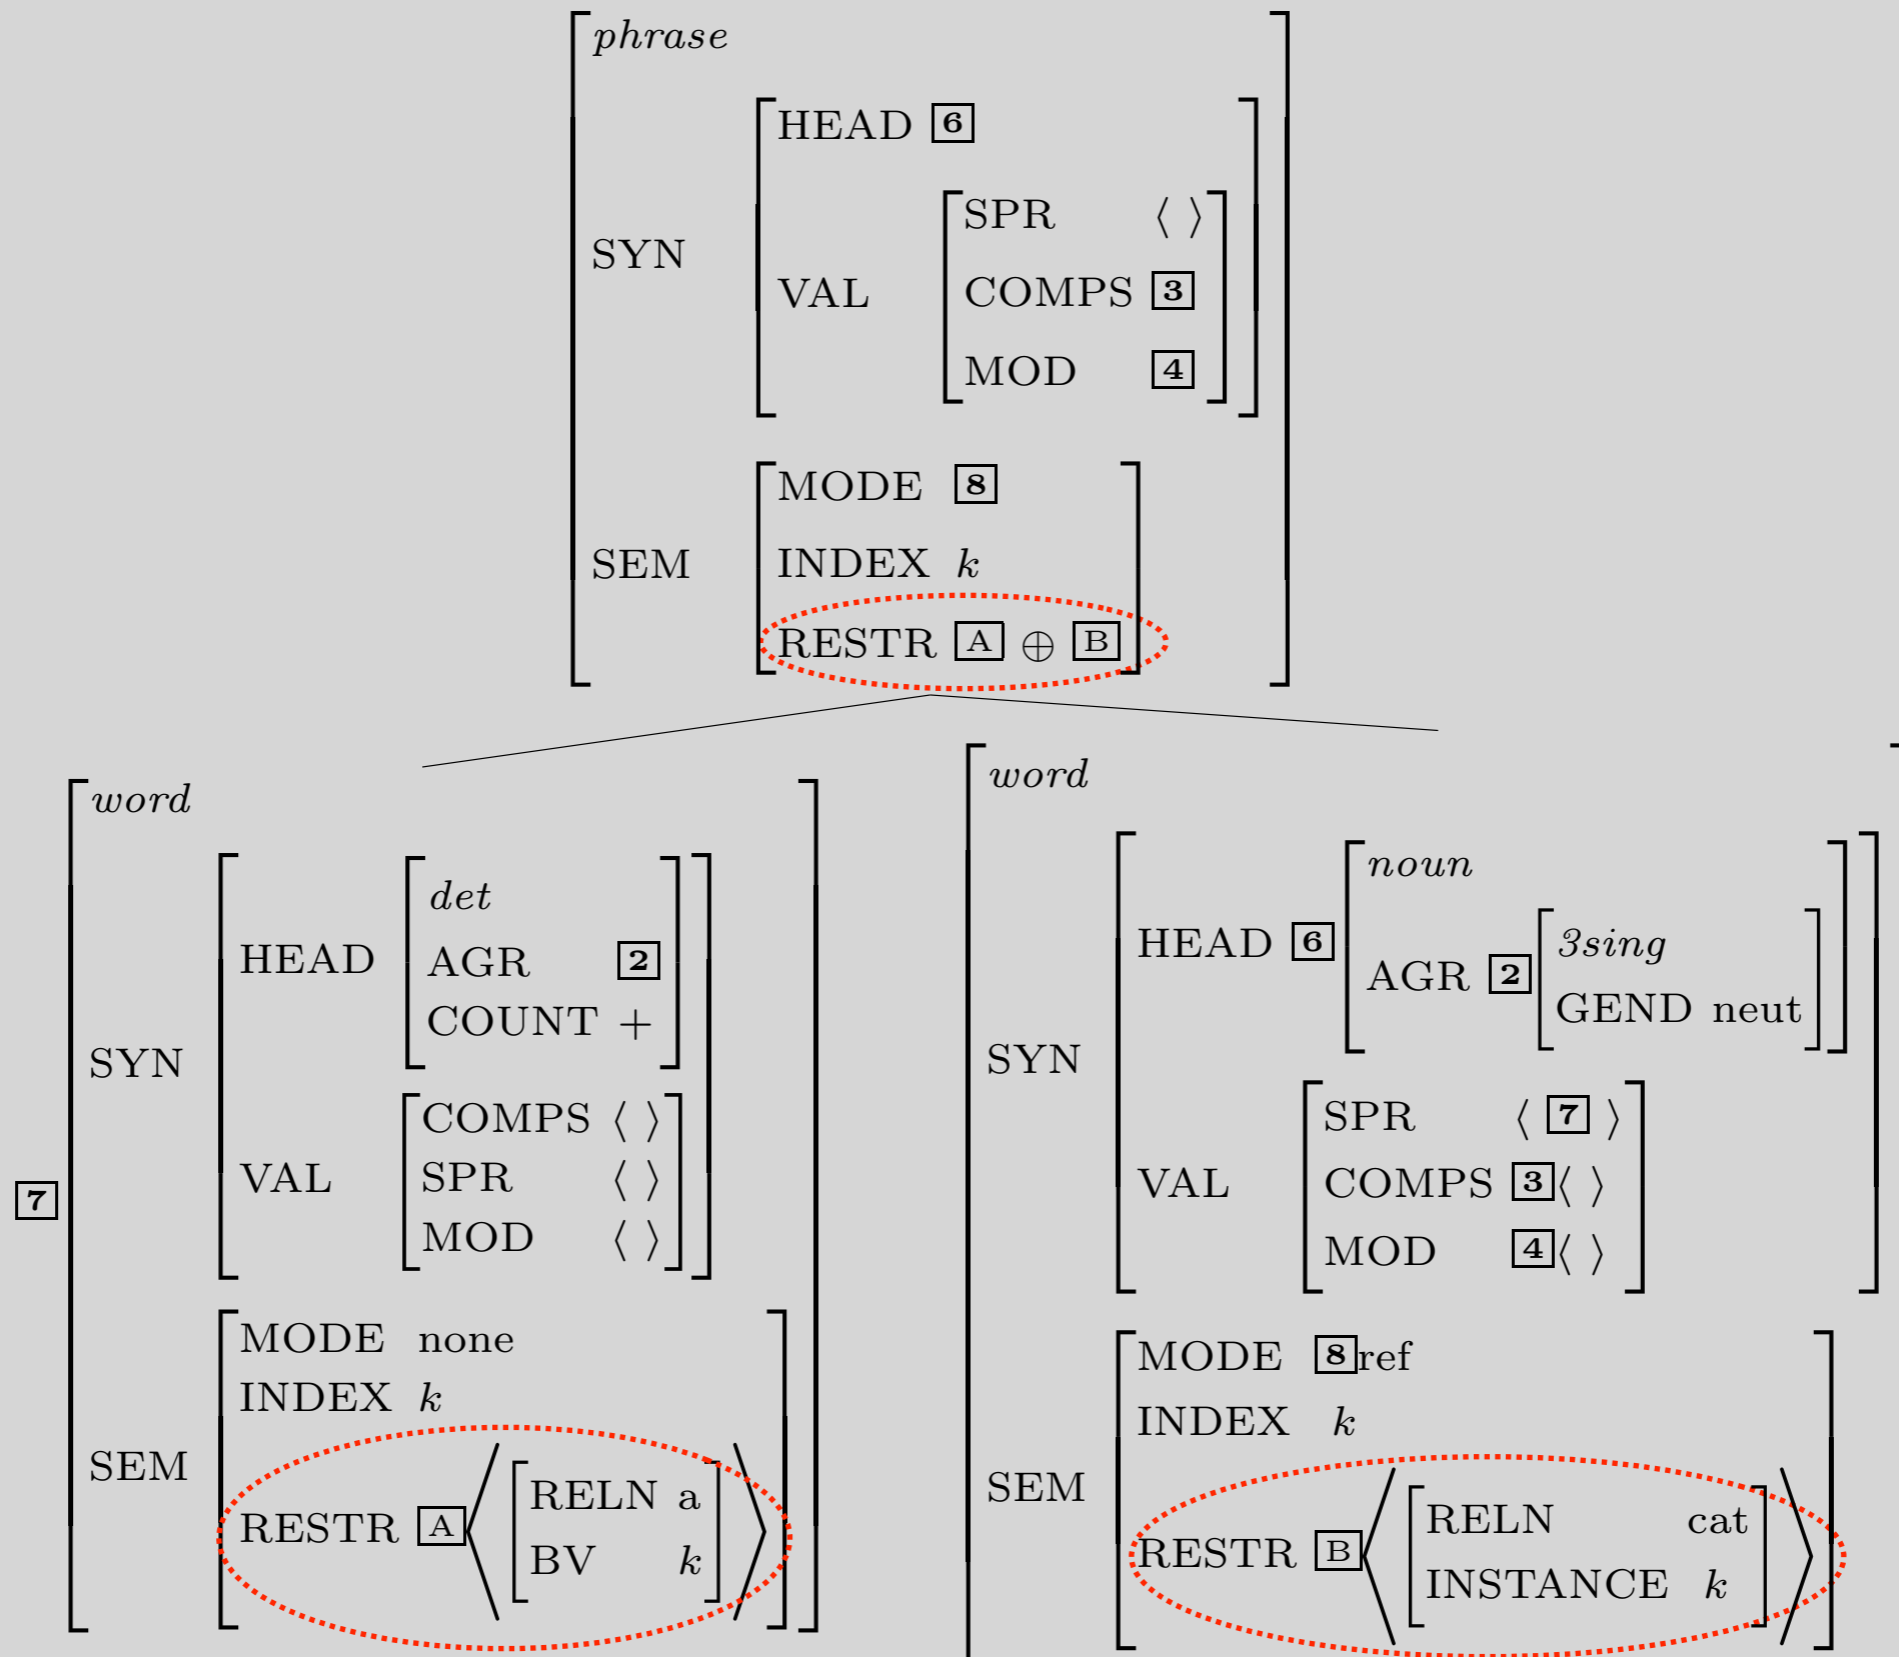
Effects of the Head Feature Principle




Effects of the Semantic Inheritance Principle



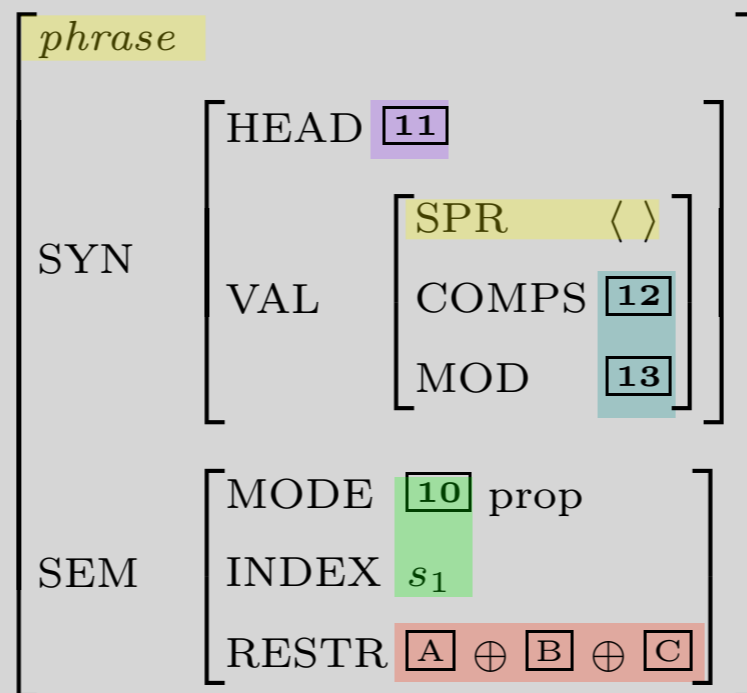
Effects of the Semantic Compositionality Principle



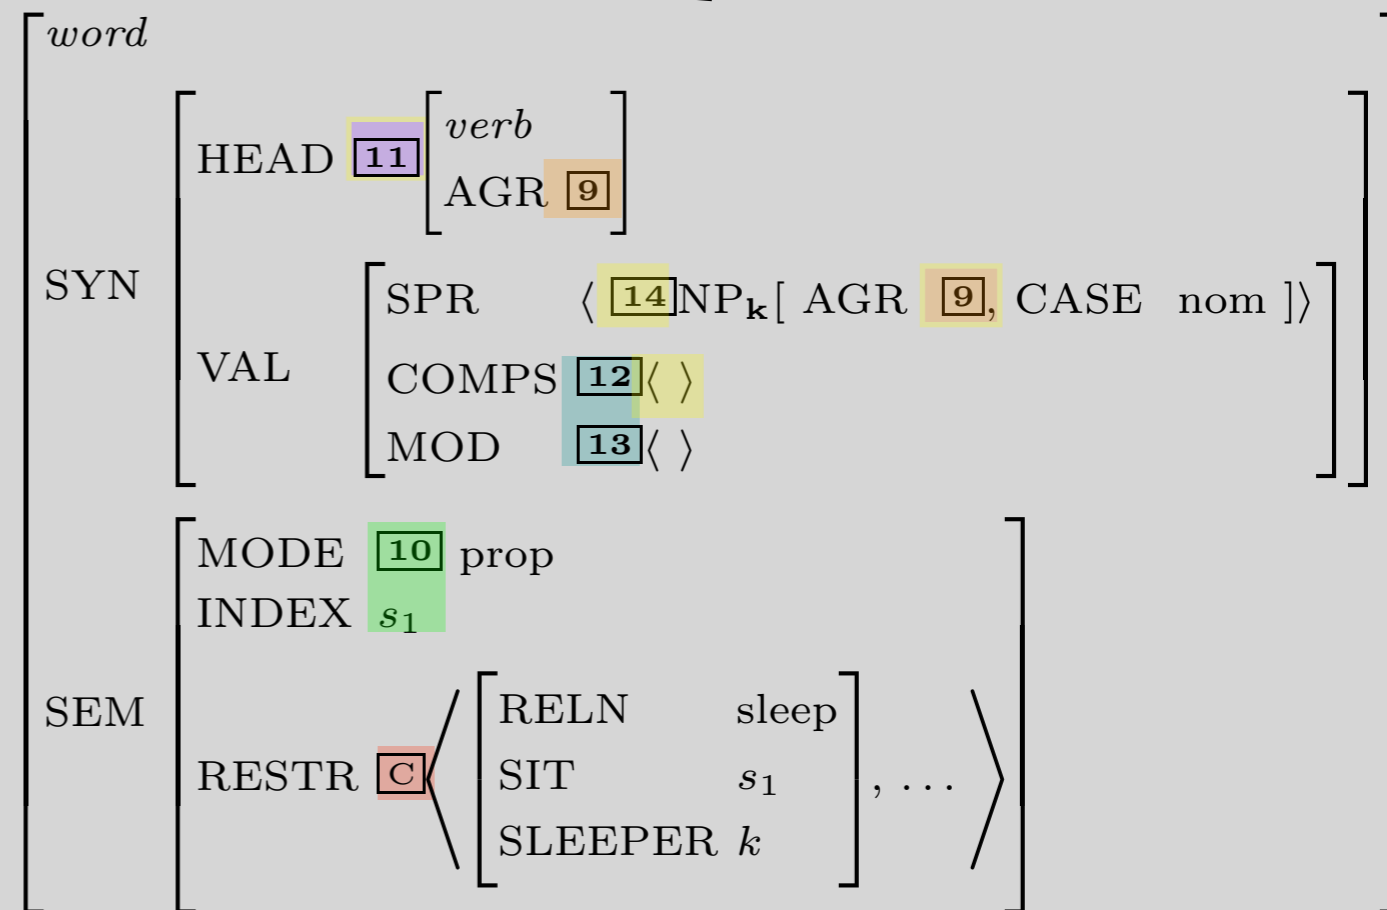


Lexical Entry for *slept*




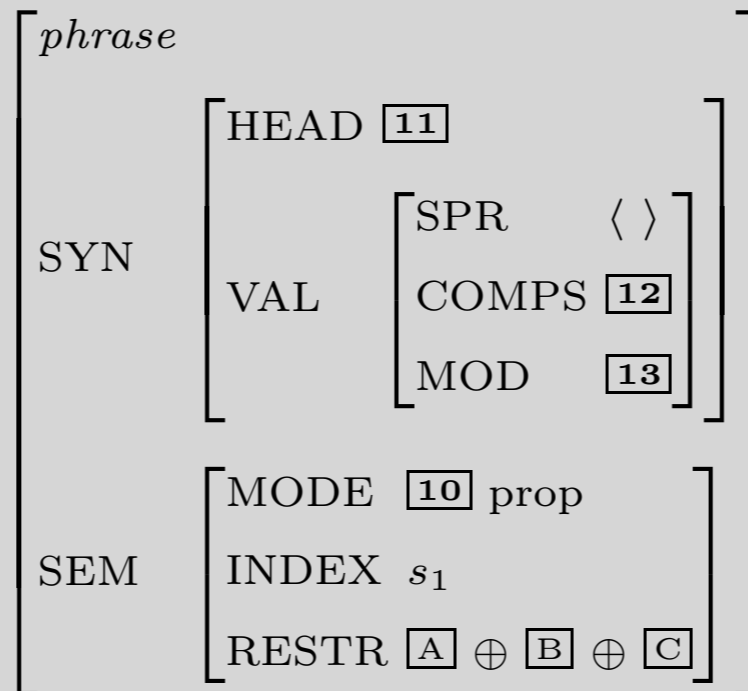
Another Head-Specifier Phrase





Key



Is this description fully specified?




Does the top node satisfy the initial symbol?




RESTR of the S node

$$\left\langle \begin{bmatrix} \text{RELN} & a \\ \text{BV} & k \end{bmatrix}, \begin{bmatrix} \text{RELN} & \text{cat} \\ \text{INST} & k \end{bmatrix}, \begin{bmatrix} \text{RELN} & \text{sleep} \\ \text{SIT} & s_1 \\ \text{SLEEPER} & k \end{bmatrix}, \dots \right\rangle$$


Another Example




Head Features from Lexical Entries




Head Features from Lexical Entries, plus HFP




Valence Features: Lexicon, Rules, and the Valence Principle



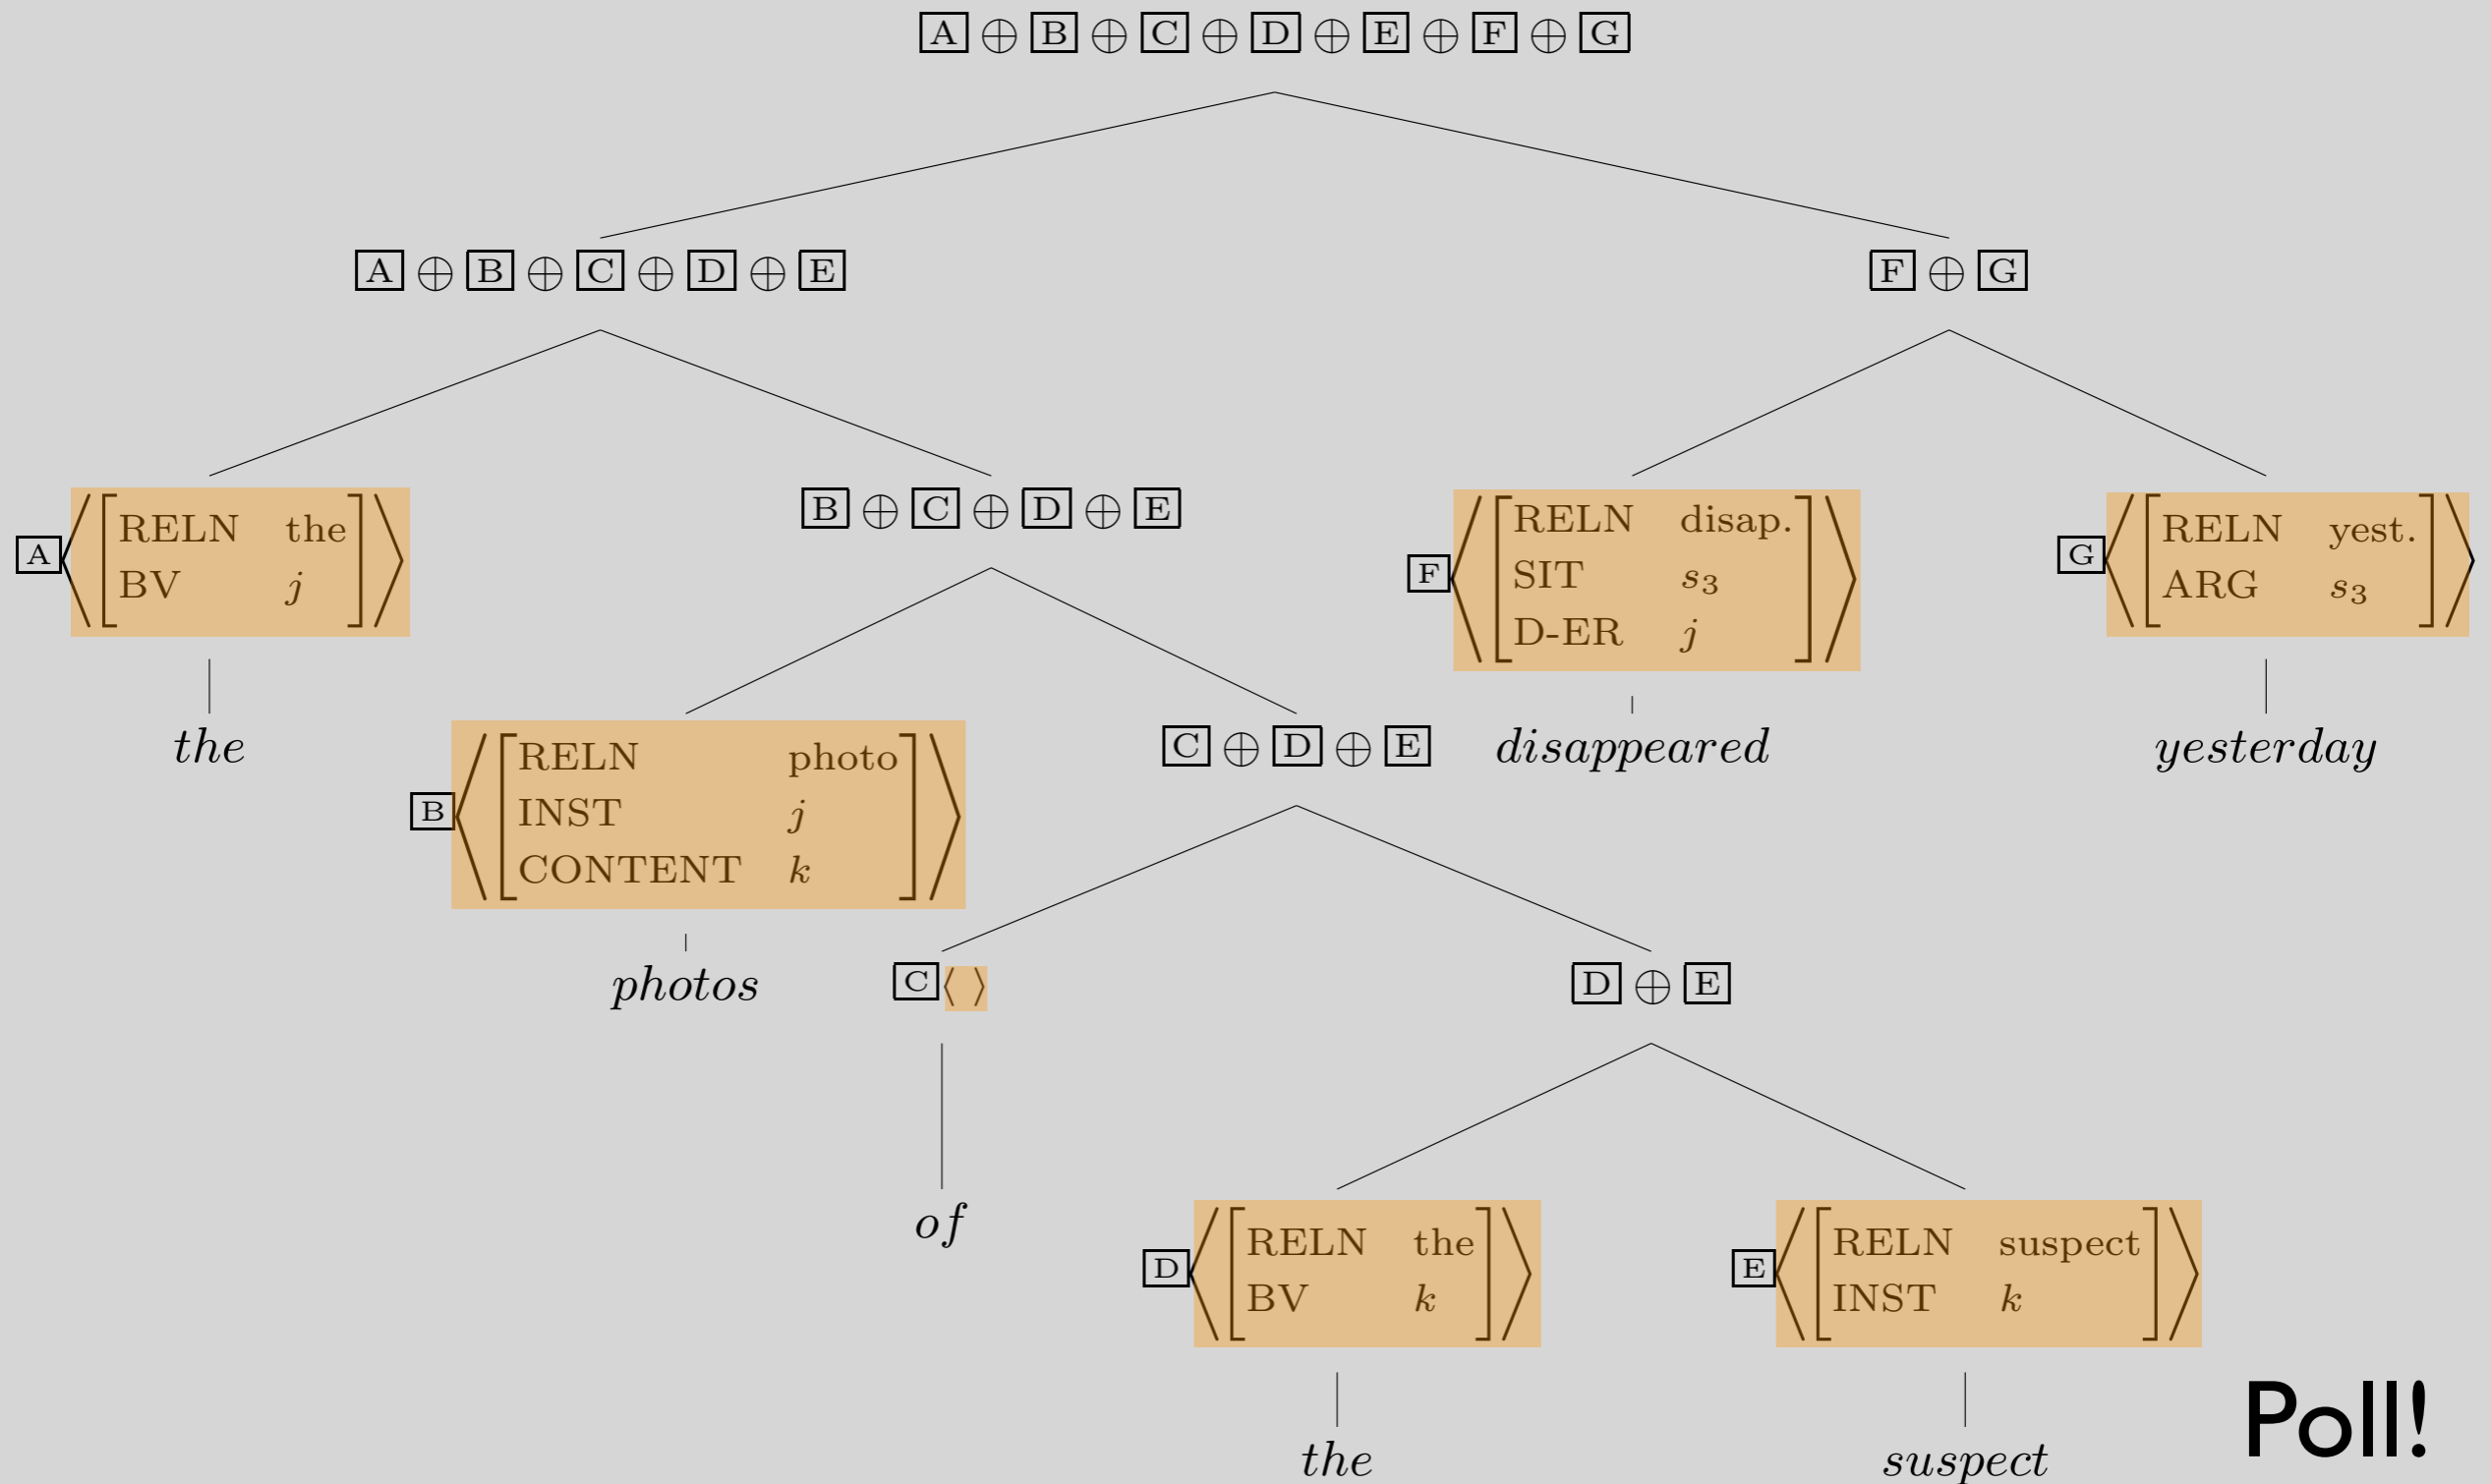
Required Identities: Grammar Rules



Two Semantic Features: the Lexicon & SIP




RESTR Values and the SCP




Poll!

An Ungrammatical Example



What's wrong with this sentence?

An Ungrammatical Example




What's wrong with this sentence?

So what?

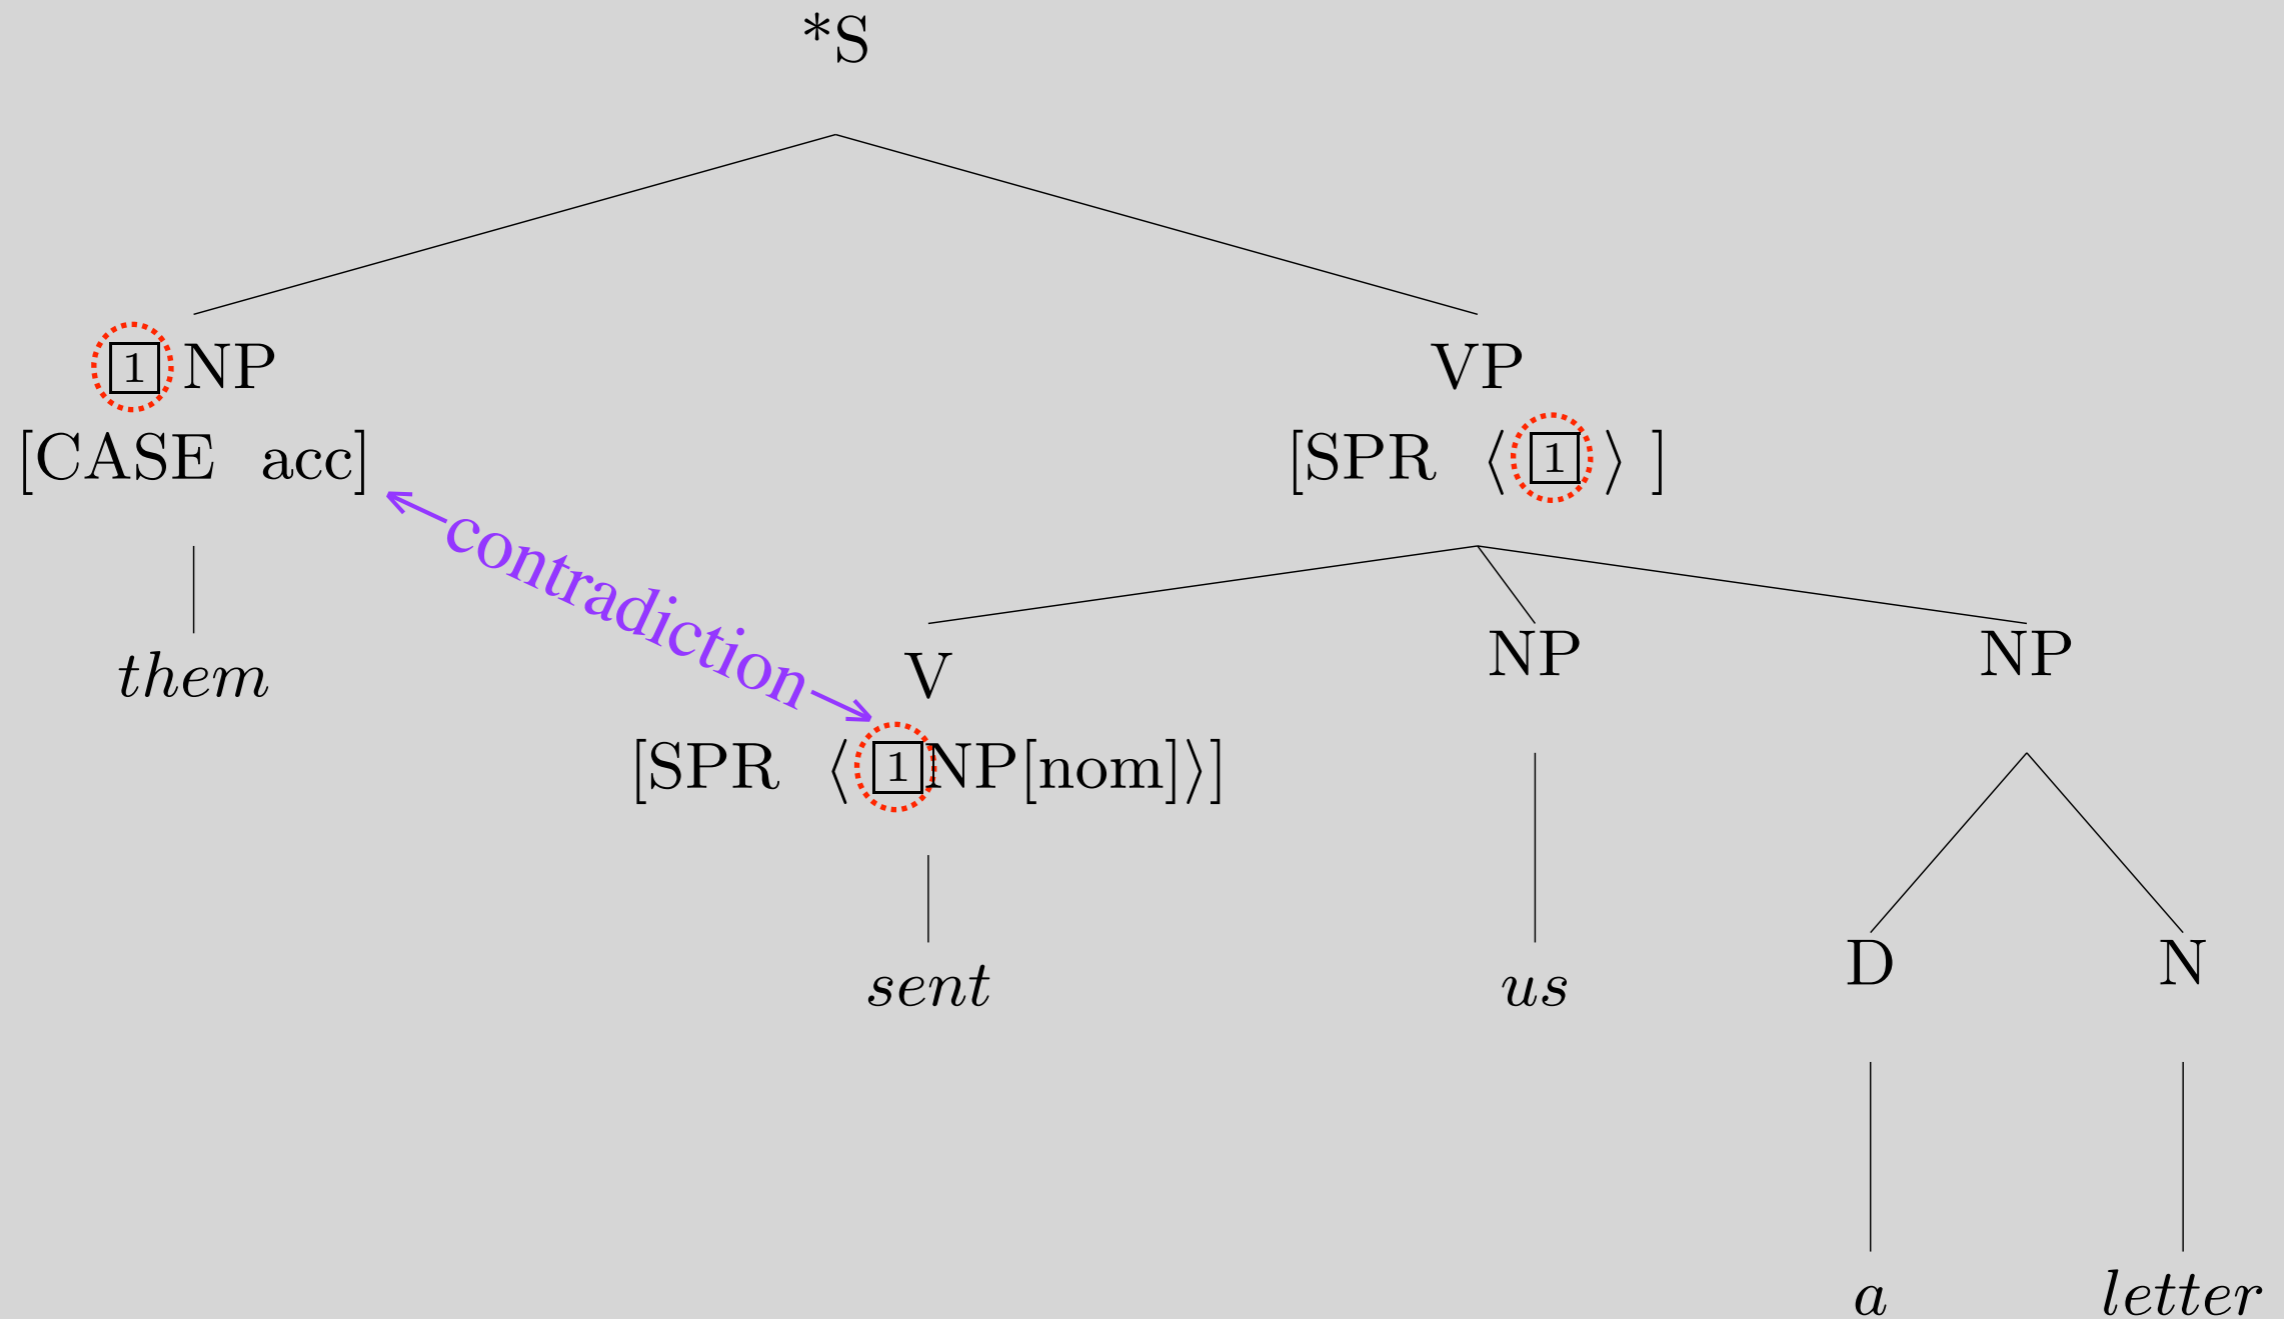
An Ungrammatical Example

The Valence Principle



An Ungrammatical Example

Head Specifier Rule



Exercise in Critical Thinking

- Our grammar has come a long way since Ch 2, as we've added ways of representing different kinds of information:
 - generalizations across categories
 - semantics
 - particular linguistic phenomena: valence, agreement, modification
- What else might we add? What facts about language are as yet unrepresented in our model?

W What grammatical phenomena are we missing?

Overview

- What we're trying to do
- The pieces of our grammar
- Two extended examples
- Reflection on what we've done, what we still have to do
- Reading questions

Reading Questions

- I'm a little unsure about why the COMPS list remains empty so often. If an N takes a D as a specifier (SPR), wouldn't that D require an N as a complement (COMPS)?
- Why do Ns and the Ds they combine with have the same INDEX?

(65)

\langle dog ,	SYN	HEAD	$\left[\begin{array}{l} \textit{noun} \\ \textit{AGR} \quad \textit{3sing} \end{array} \right]$
	VAL	SPR	$\langle \textit{DP}_i \rangle$
		COMPS	$\langle \rangle$
		MOD	$\langle \rangle$
	SEM	MODE	ref
		INDEX	i
		RESTR	$\left\langle \left[\begin{array}{l} \textit{RELN} \quad \textit{dog} \\ \textit{INST} \quad \quad \quad \textit{i} \end{array} \right] \right\rangle$

(70)

$\langle a, \rangle$	<i>word</i>								
	SYN	HEAD	<i>det</i>	AGR	<i>3sing</i>	COUNT	+		
		VAL	COMPS		$\langle \rangle$	SPR	$\langle \rangle$		
			MOD		$\langle \rangle$				
	SEM	MODE	none						
		INDEX	<i>i</i>						
		RESTR			\langle	$\left[\begin{array}{l} \text{RELN} \\ \text{BV} \end{array} \right]$	\rangle	exist	\rangle
								<i>i</i>	

Reading Questions

- How do we know when to leave SPR, COMPS and MOD empty in a lexical entry or phrase and when not to? It feels like every lexical entry should be as complete as possible if we are to say they are entries used to build sentences.
- A similar issue applies to SEM features.
- How do we know which features ought to be left empty for which lexical items? I don't fully understand the difference between which features are essential, for instance, between an isolated lexical entry and a lexical item's feature structure when it appears in a tree.

Reading Questions

- The difference between (12) and (23) (the lexical entries for "sent" and "send") makes me wonder which goes first? (a) Is it the sentence itself that determines the lexical entries of words? (b) Or is it lexical entries of words that define well-formed sentences?

Reading Questions

- If a word has two completely different meanings and roles in different contexts then do we create two separate lexical entries for it or do we underspecify things in single lexical entry such that it can be used in all the cases? For example, reading has two meanings - action of reading text and a town in England. My intuition is that we should create two separate lexical entries but wanted to confirm.

Reading Questions

- Can you please re-articulate the difference between a lexical entry and a lexical tree (word structure)?
- I can see that a lexical entry "gives rise to" (licenses?) the lexical tree, but how do we decide what feat-struct to include vs. leave underspecified in either one? For example, on page 175, the lexical tree (13) for they has included AGRs, CASEs, and tagging which are non-existent in its lexical entry (12).

Reading Questions

- When is it okay for two different sentences, that don't even have all the same words, to have the same semantics?
- How do we tell which words are semantically empty?
- Are prepositions almost always devoid of semantics?

Reading Questions

- When does a word have empty RESTR like 'to' in this chapter, and are there other tests for determining that? In the example given, the claim is that We sent Lee two letters has the same meaning as We send two letters to Lee, but to me it seems like We sent Lee two letters only has one interpretation where Lee is the sendee whereas We send two letters to Lee has the two interpretations described in the text. It seems then like the two sentences aren't exactly the same, but I'm also not sure the semantic difference comes from the word 'to'. Are there other arguments for 'to' or other words having empty RESTR?

Reading Questions

- What is the difference between the RESTR list, comma separated and the sum symbol separated?
- What are some other ways to represent the RESTR of *us*?

Reading Questions

- I'm curious where we can find the information of the single-ness of "a letter" in the RESTR list in the example tree (10). The feature for "a" has RELN "exist." However, if it is "two letters," how does the feature "exist" show the "two-ness" of the phrase? Wouldn't it be better to add a NUM feature in SEM to every noun phrase to specify the quantity (like the constant feature in HW2)?

Reading Questions

- I think my confusion here stems from not intuitively understanding what RESTR signifies in SEM-cat. Is it the truth conditions that must be met for the phrase to be valid? If so, where do we indicate that we're concerned that a "letter" exists? Does "the" do the heavy lifting of confirming existence?

Reading Questions

- I am still confused about the INDEX value in SEM, sometimes it refers to the SIT in RESTR such as in (23) and it corresponds to INST in (24), and in (8), the INDEX is the same for letter and its SPR.

Reading Questions

- "Since we have not imposed any constraint requiring that semantic roles be realized syntactically, this does not present any technical problem. And having an ADDRESSEE role for the noun letter, even when no addressee is mentioned, seems quite intuitive." I didn't quite get this - I thought restrictions were meant to be binding? That if the word needs an addressee, it has to be there and mentioned in RESTR? Agreed "letter" doesn't need an addressee all the time. So can't we put brackets around ADDRESSEE to indicate that it is optional, like we do in the COMPS/SPR lists?

Reading Questions

- I am a bit confused on why letter has an ADDRESSEE feature. This is assuming something about the letter that we do not know, as a letter does not technically have to be addressed to anyone. If the sentence was "They gave us a letter" instead of "They sent us a letter," would letter still have the addressee feature? You could give someone a letter that was addressed to someone else, like giving a letter to the postman. It just seems to me that our grammar is generating information, rather than explaining what is/is not syntactically and semantically valid.

Reading Questions

- I'm still a little bit confused by the two RESTR values of Lee even after reading the footnote on page 191... Would you walk us through this in class?
- In 6.2.1, according to the lexical entry, the word letter takes an optional PP complement which semantically represents the addressee, e.g. "letter to Kim". I was wondering why such PPs are treated as complements rather than modifiers considering they can always be omitted.

We send two letters to Lee

(19) $\left[\begin{array}{ll} \text{RELN} & \text{group} \\ \text{INST} & i \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{speaker} \\ \text{INST} & l \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{member} \\ \text{SET} & i \\ \text{ELEMENT} & l \end{array} \right],$

$\left[\begin{array}{ll} \text{RELN} & \text{send} \\ \text{SIT} & s_7 \\ \text{SENDER} & i \\ \text{SENDEE} & j \\ \text{SENT} & k \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{two} \\ \text{BV} & k \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{letter} \\ \text{INST} & k \\ \text{ADDRESSEE} & m \end{array} \right],$

$\left[\begin{array}{ll} \text{RELN} & \text{name} \\ \text{NAME} & \text{Lee} \\ \text{NAMED} & j \end{array} \right]$

We send two letters to Lee

(20) $\left[\begin{array}{ll} \text{RELN} & \text{group} \\ \text{INST} & i \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{speaker} \\ \text{INST} & l \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{member} \\ \text{SET} & i \\ \text{ELEMENT} & l \end{array} \right],$

$\left[\begin{array}{ll} \text{RELN} & \text{send} \\ \text{SIT} & s_7 \\ \text{SENDER} & i \\ \text{SENDEE} & j \\ \text{SENT} & k \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{two} \\ \text{BV} & k \end{array} \right], \left[\begin{array}{ll} \text{RELN} & \text{letter} \\ \text{INST} & k \\ \text{ADDRESSEE} & m \end{array} \right],$

$\left[\begin{array}{ll} \text{RELN} & \text{name} \\ \text{NAME} & \text{Lee} \\ \text{NAMED} & m \end{array} \right]$

Reading Questions

- In chapter 6, why was the CASE feature designed to be a sub-feature of HEAD instead of being a sub-feature of AGR? Is there any specific reason for that? When I was doing homework (Problem 8, Chapter 4), I chose to put CASE under the AGR intuitively, so that the SHAC rule will hold. (E.g. languages like German, the determiner would have the same case as the word it specifies).

Reading Questions

- Does every treebank based on this grammar have a consistent manner for creating the restrictions list? For example, would the word 'us' always have the same restrictions list for every sentence it occurs in in a particular treebank?
- This is more of a general question: Is a particular theory of grammar judged by how complicated (in terms of human annotation effort) it is to create a treebank based on the grammar proposed by the theory?

Reading Questions

- When building the trees (say, the one on Page 172 or 178), why don't we need to use tags to demonstrate the identity of MODE and INDEX, guaranteed by the Semantic Inheritance Principle, just as we do to the other rules/principles? Here the mother and head daughter share the same MODE ref and INDEX k, and that is explicitly stated in the text below, but why do we need not tag them to show this?

Reading Questions

- Why do we use lettered indexes for RESTR? This requires that we keep referring back to the lexicon but I had assumed that you wouldn't need to do that with the tree. Is this just to save space or is there another reason?

Reading Questions

- In comparing the examples on pages 169 and 170, I see some differences. I understand, for instance, the CASE gets filled in because of the context. But why is COMPS<> list empty now? I'd think the word doesn't take an optional PP, in this case, doesn't mean it can't. So I was expecting there to be an optional PP (i.e. that still holds). Is the purpose of lexical entries suppose to show what *can be done* with the word or what *is done* with the word?