CS 490: NATURAL LANGUAGE PROCESSING

Dan Goldwasser, Abulhair Saparov

Lecture 16: Compositional Semantics

RECAP

- We have discussed semantics.
  - I.e., the study of how meaning is represented in language.
- We discussed logical formalisms:
  - How to use logic to express the truth-functional meaning of sentences.
- Last lecture, we discussed lexical semantics,
  - Which focuses on the meaning of individual words,
  - And their relationships with each other.
- Today, we will discuss compositional semantics.
  - How is meaning of a larger phrase composed from the meanings of its constituent phrases?
- And how do we convert sentences into logical forms? (semantic parsing)

DO WE NEED A LOGICAL FORM?

- It's not obvious that human language processing involves converting natural language into logical form.
- Counterargument: Logical forms enable reasoning.
- But why not do reasoning in natural language?
  - I.e., natural language is the logical formalism.
- One potential roadblock: Ambiguity.
- Logical forms in a formal language are unambiguous.
  - Natural language is infamously ambiguous.

DO WE NEED A LOGICAL FORM?

- Consider the example:
  - '`All dogs chase a cat.`'
    - `∀d(dog(d) -> ∃c(cat(c) & chase(d,c)))`
    - `∃c(cat(c) & ∀d(dog(d) -> chase(d,c)))`
  - '`Sif and Fen are dogs.`'
    - `dog(sif) & dog(fen)`
  - '`Sif only chases Felix.`'
    - `chase(sif,felix) & ¬∃x(x≠felix & chase(sif,x))`
- If we take the second reading of '`All dogs chase a cat`', we can prove that '`Fen chases Felix.`'
- If we take the first reading, the proof is no longer valid.

DO HUMANS USE LOGICAL FORMS?

- Do humans "use logical forms"?
- There is some neuroscientific evidence that humans perform reasoning in a more abstract, modality-independent fashion.
- MacSweeney (2002) performed brain scans of 11 volunteers while they performed a reading comprehension task.
  - 4 deaf subjects who know British Sign Language (BSL).
  - 4 hearing subjects who know BSL.
  - 3 hearing subjects who don't know BSL.
  - Scanned subjects using fMRI (functional magnetic resonance imaging).

DO HUMANS USE LOGICAL FORMS?

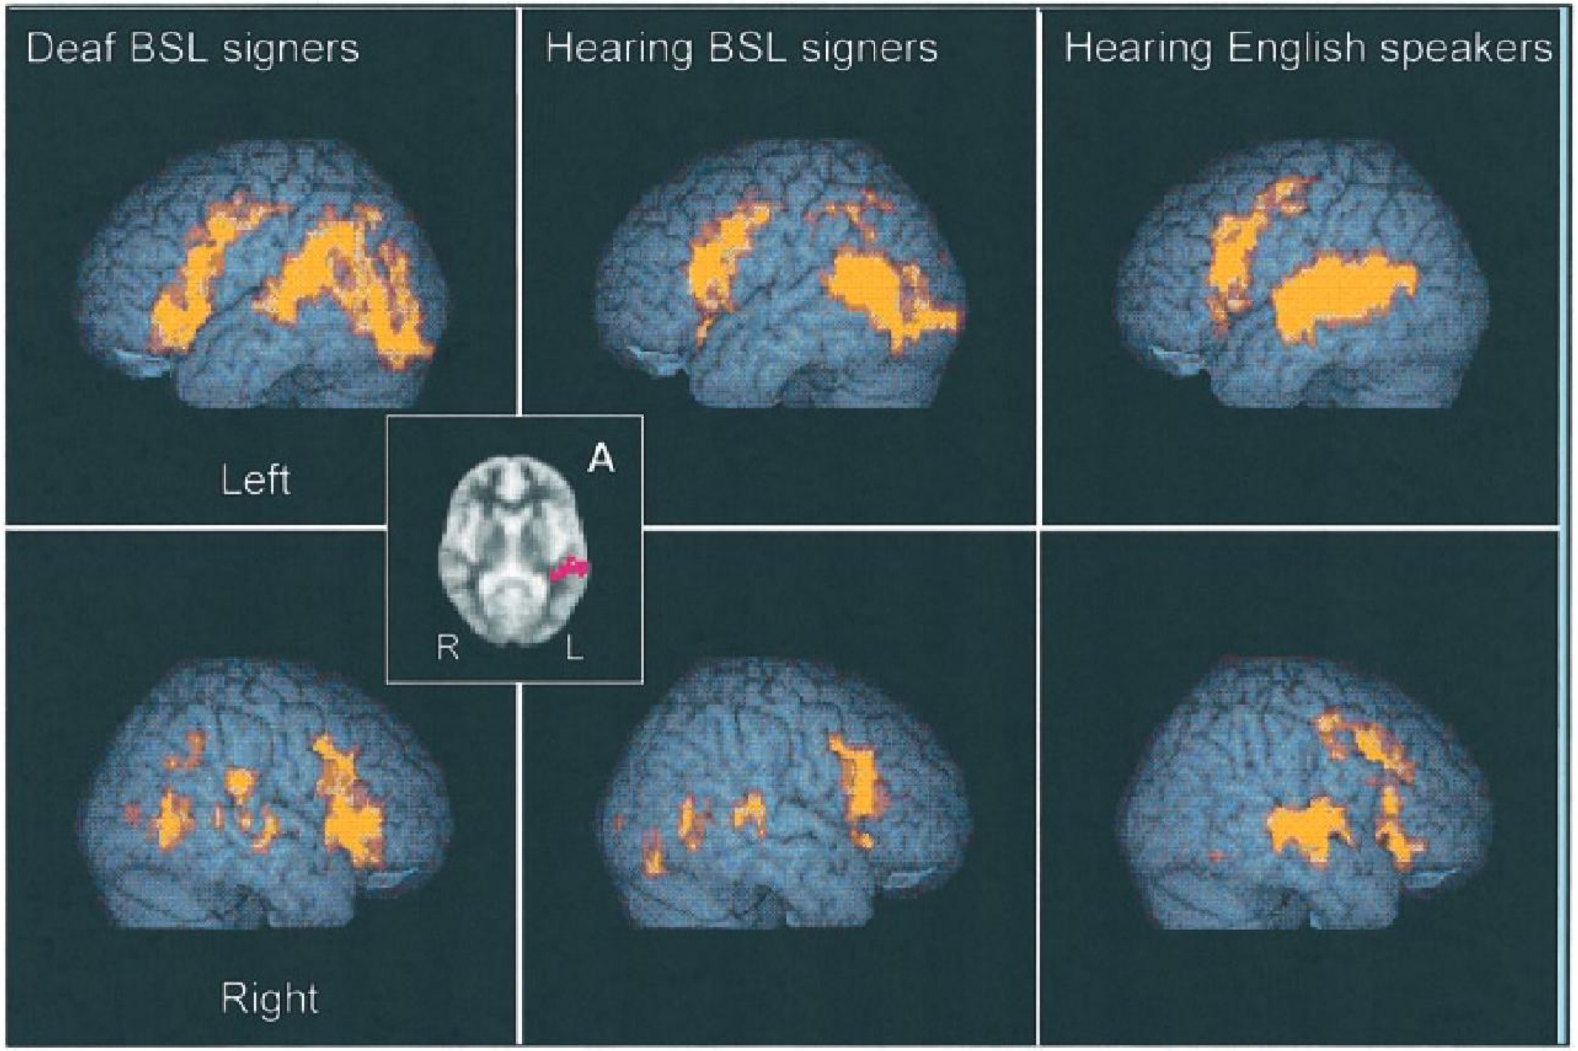

DO HUMANS USE LOGICAL FORMS?

- There are brain regions that are active across both deaf and hearing subjects.
- But maybe this is due to common syntactic processing across modalities?
  - Not likely since BSL and spoken English are very different grammatically.
  - BSL has OSV word order and nouns are head-initial (e.g., '`car blue`').

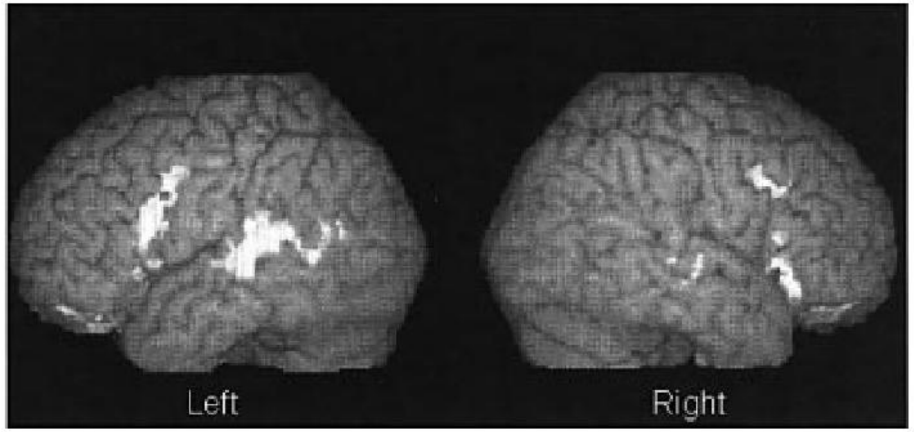

Fig. 2 Locations of common activation for audio-visual English (hearing) and BSL sentences (deaf). Activation up to 5 mm under the surface of the cortex is displayed.

DO HUMANS USE LOGICAL FORMS?

- Monti et al. (2012) used fMRI to localize which brain areas were active when subjects are given a language task vs a mathematical reasoning task.

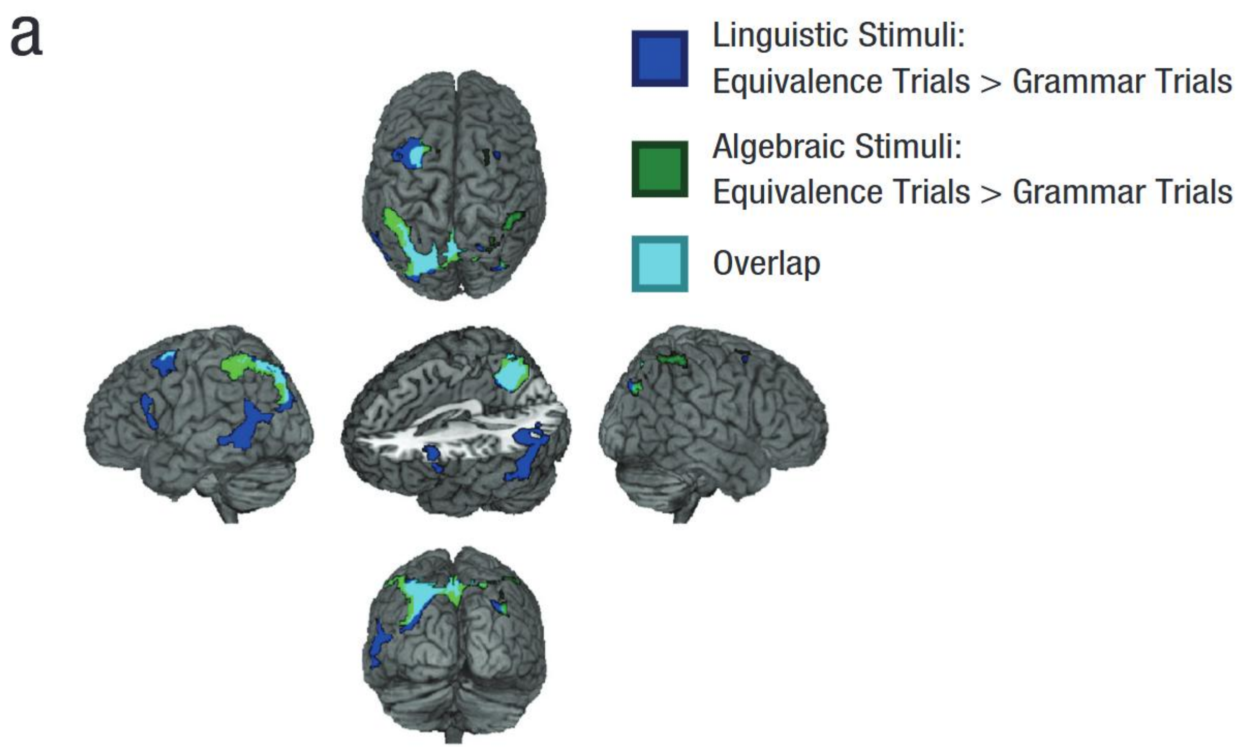

DO HUMANS USE LOGICAL FORMS?

- Monti et al. (2012) used fMRI to localize which brain areas were active when subjects are given a language task vs a mathematical reasoning task.

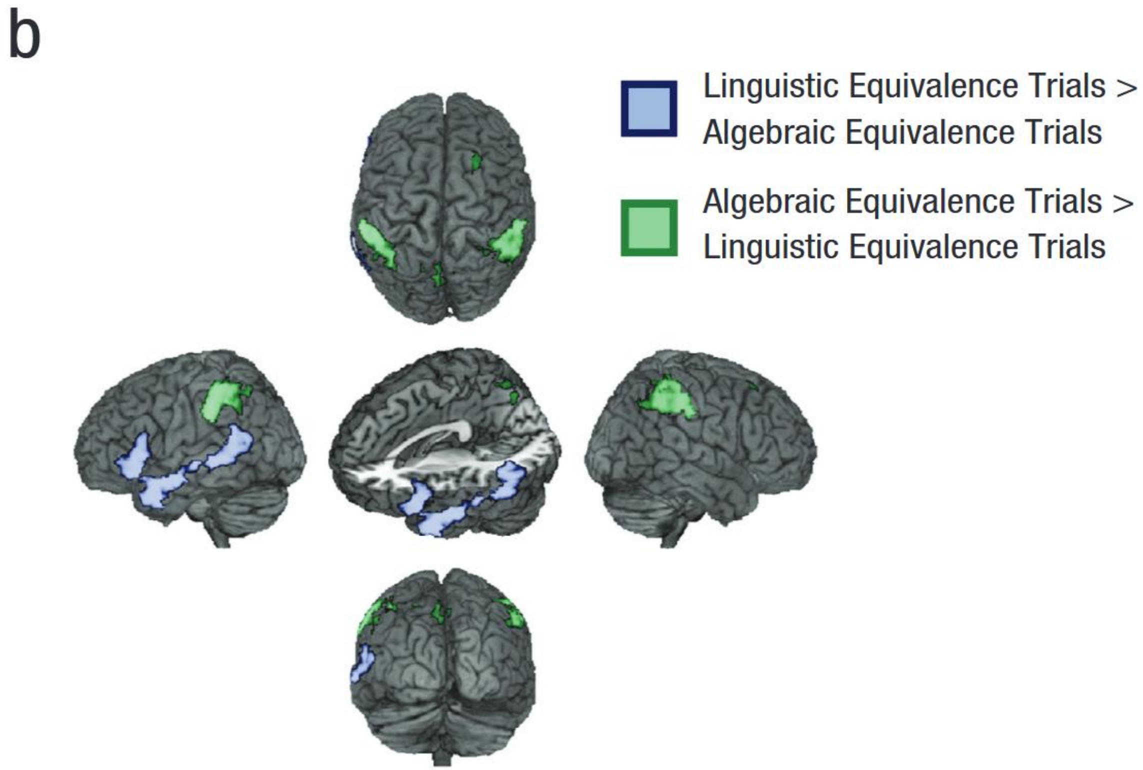

DO WE NEED LOGICAL FORMS?

- Monti et al. (2012) used fMRI to localize which brain areas were active when subjects are given a language task vs a mathematical reasoning task.
- Neuroscientific evidence supports the notion of a “language network” within the brain that is highly specialized for language processing.
- But this language network is not heavily involved in high-level reasoning.
  - E.g., mathematical reasoning.
- Logical forms are useful for other applications in NLP.
  - E.g., code generation
    - Text-to-SQL
    - etc...
  - Programs are logical forms!

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage:
    - If there are two sentences with different meanings, they should have different logical forms.
    - If there are multiple readings/interpretations of the same sentence, there should be a logical form for each reading.
    - We want to be able to capture all ambiguous interpretations (i.e., the formalism should accurately model semantic ambiguity) (see last two lectures)

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage
  - Language-independence:
    - This depends on the application.
    - If we wish to model translation from one natural language to another, Where the meaning is language-independent,
    - Ideally, two sentences in different languages but with the same meaning should map to the same logical form.

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage
  - Language-independence
  - Uniqueness:
    - Two sentences with the same meaning should map to the same LF.
    - Not always needed.
  - Amenable for reasoning:
    - If we want to reason over the logical forms,
    - It would help to have a set of inference rules with which we can construct proofs.

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage
  - Language-independence
  - Uniqueness
  - Amenable for reasoning
  - Easy to parse:
    - It shouldn't be too difficult to parse from natural language into LF.
    - If the logical form is more similar to the natural language utterance, semantic parsing is easier.
    - We will discuss semantic parsing in more detail later.

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage
  - Language-independence
  - Uniqueness
  - Amenable for reasoning
  - Easy to parse
  - Compositional:
    - The logical form of larger phrases should be composed of the logical forms of subphrases.

COMPOSITIONALITY

- Compositionality is also relevant in syntactic analysis.
  - The syntax tree of a larger phrase is composed of the syntax trees of its constituents.
  - Grammars with recursion enable syntactic compositionality.
  - E.g., context-free grammars.
- Semantic compositionality:
  - ‘a barn’ -> ∃b.barn(b)
  - ‘a dog’ -> ∃d.dog(d)
  - ‘in a barn’ -> ∃b(barn(b) & location(_,b))
  - ‘a dog in a barn’ -> ∃d(dog(d) & ∃b(barn(b) & location(d,b)))
- Exceptions: Idioms such as ‘break a leg’, ‘on the other hand’, etc.

COMPOSITIONALITY

- Compositionality is very useful for efficient parsing.
  - We can parse a sentence by first parsing smaller phrases.
  - Enables dynamic programming approaches.
- Compositionality is also important for generalization.
  - We can predict the syntactic/semantic structure of larger phrases/sentences based on its smaller constituent phrases.
  - Compositional generalization :
    - If a model correctly "understands" some phrases/sentences,

      How well does it "understand" larger phrases/sentences that contain those phrases?

COMPOSITIONAL GENERALIZATION

- Kim and Linzen (2020) introduced a dataset called COGS to test the semantic compositional generalization ability of NLP models.
- Some examples from COGS:

TRAINING

[[The girl]] = $\iota x.\, girl'(x)$, [[The cat]] = $\boldsymbol{\iota x.\, cat'(x)}$, [[The boy]] = $\iota x.\, boy'(x)$

[[The cat loves the girl]] = $love'(\boldsymbol{\iota x.\, cat(x)}, \iota x.\, girl'(x))$

[[The hedgehog sees the cat]] = $see'(\iota x.\, hedgehog'(x), \boldsymbol{\iota x.\, cat'(x)})$

GENERALIZATION

[[The boy loves the hedgehog]] = $love'(\iota x.\, boy'(x), \boldsymbol{\iota x.\, hedgehog(x)})$

COMPOSITIONAL GENERALIZATION

- Kim and Linzen (2020) introduced a dataset called COGS to test the semantic compositional generalization ability of NLP models.
- They evaluated transformers, unidirectional LSTMs, and bidirectional LSTMs.

Model	Dev.	Test	Gen.
Transformer	0.96	0.96	**0.35** (± 0.06)
LSTM (Bi)	0.99	0.99	0.16 (± 0.08)
LSTM (Uni)	0.99	0.99	0.32 (± 0.06)

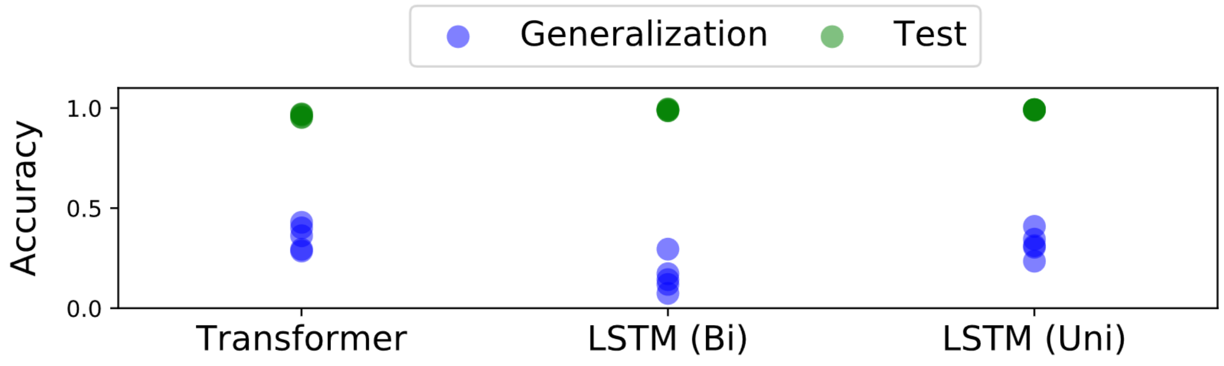

COMPOSITIONAL GENERALIZATION

- They considered two types of generalization:
  - Lexical generalization: A word is used in a previously-unseen context.

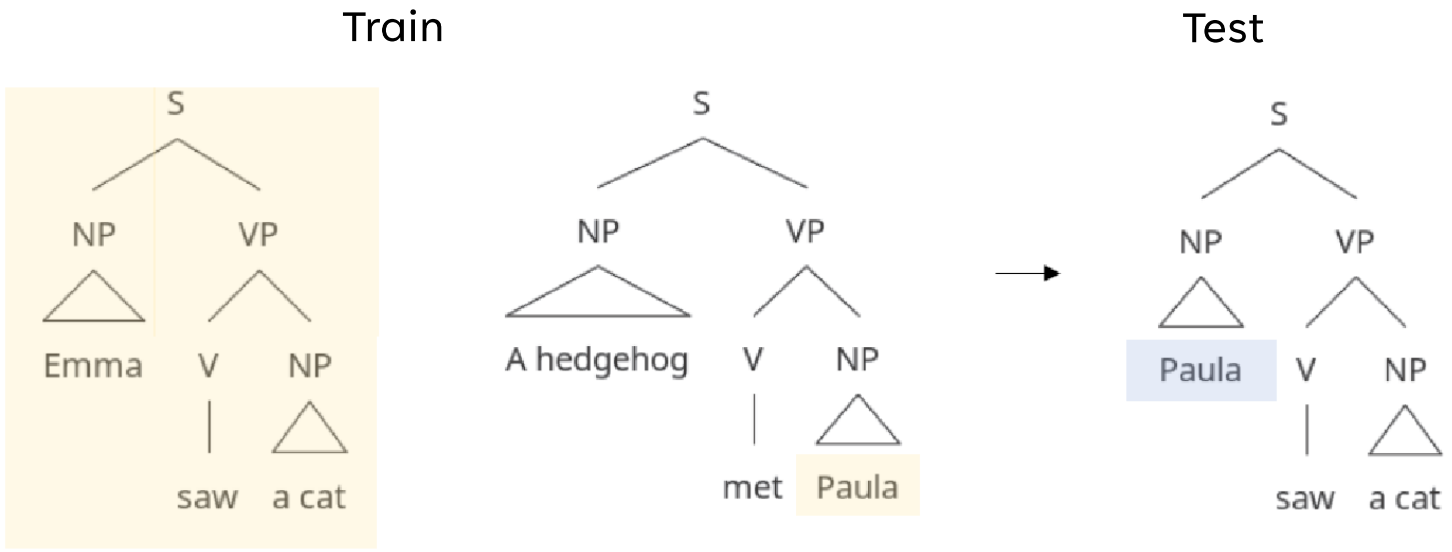

[Kim and Linzen, 2020]

COMPOSITIONAL GENERALIZATION

- They considered two types of generalization:
  - Lexical generalization: A word is used in a previously-unseen context.
  - Structural generalization: Familiar structures are combined into a novel larger structure.

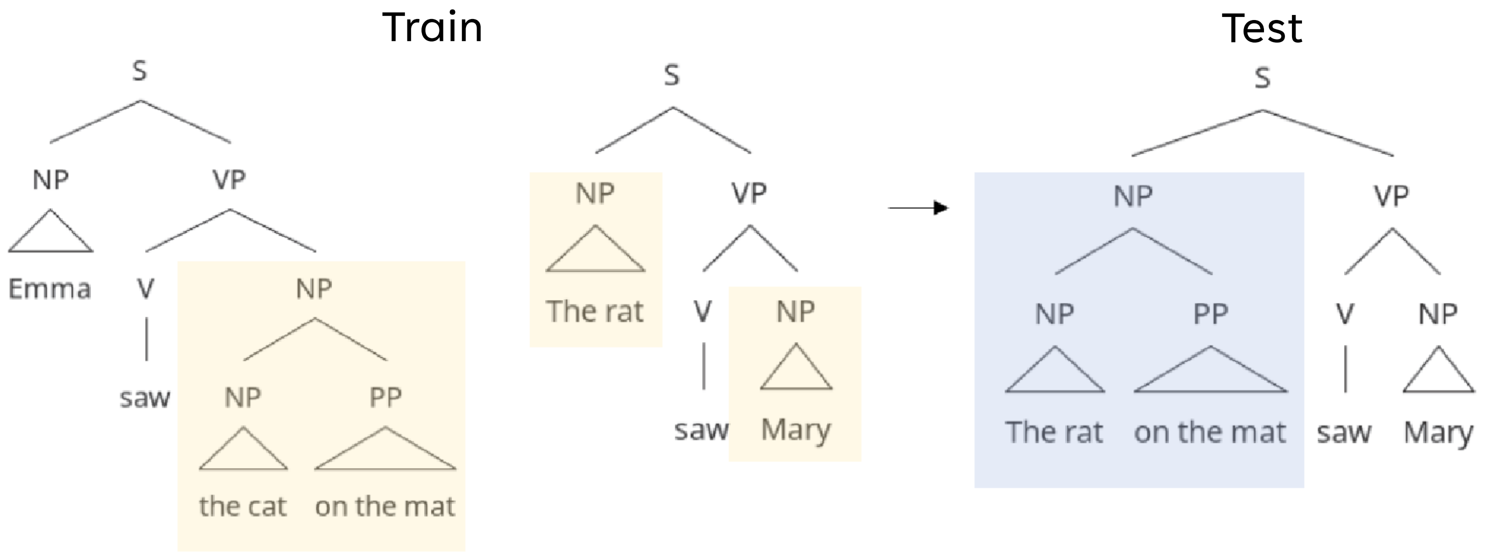

[Kim and Linzen, 2020]

COMPOSITIONAL GENERALIZATION

- They considered two types of generalization:
  - Lexical generalization: A word is used in a previously-unseen context.
  - Structural generalization: Familiar structures are combined into a novel larger structure.

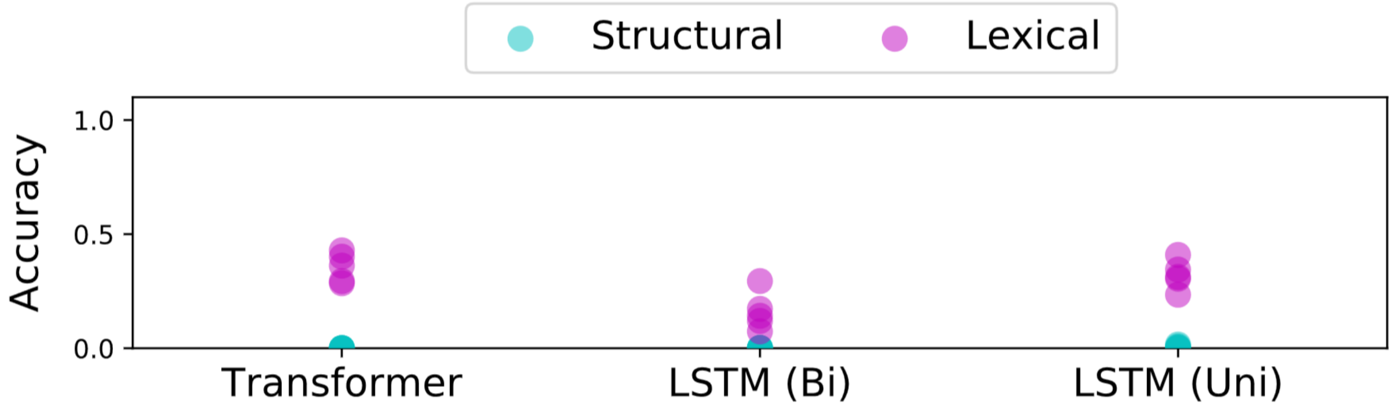

- They found that the tested models are incapable of structural generalization.
  - Newer models may do better?

[Kim and Linzen, 2020]

COMPOSITIONAL GENERALIZATION

- Compositionality is also a feature of reasoning.
- A proof can consist of simpler subproofs.
  - In order to find the full proof, an intelligent system must first find the subproofs,
  - Then combine them to produce the full proof.
- E.g., 2-hop question answering:

  ‘`Who won the Master’s Tournament the year Justin Bieber was born?`’

- To answer this, you must prove:
  - The year Justin Bieber was born is 1994.
  - Jose Maria Olazabal won the Master’s Tournament in 1994.

LOGICAL FORMALISMS

- What makes a good logical formalism/meaning representation?
  - Coverage
  - Language-independence
  - Uniqueness
  - Amenable for reasoning
  - Easy to parse
  - Compositional
  - Human-readable:
    - Logical forms should be easy to understand by humans.
    - Ideally, they shouldn't require extensive special training (e.g., a university course on formal semantics).

SEMANTIC PARSING

- Semantic parsing is the task of converting natural language to logical form.
  - NLP models that are trained to convert natural language into logical form (e.g., code) are effectively performing semantic parsing.
- Consider the following example:
  - We want to parse '`Sif chases Felix`' into `chase(sif,felix)`
  - We can model the syntax of the natural language with a grammar.
  - Logic (and any formal language) can easily be described with a CFG.

SEMANTIC PARSING

- A simple CFG for English:

```
S -> N VP          N -> 'Sif'
VP -> V N          N -> 'Felix'
V -> 'chases'
```

- CFG for first-order logic:

```
S -> S '&' S       S -> P '(' T ')'
S -> S '|' S       S -> P '(' T ',' T ')'
S -> S '=>' S      P -> 'chases'
S -> '(' S ')'     T -> V
S -> '¬' S         V -> 'x'
S -> '∀' V S       T -> 'sif'
S -> '∃' V S       T -> 'felix'
```

SYNCHRONOUS GRAMMARS

- Combine these grammars to model both English and FOL simultaneously?

S -> <N_1 V N_2, V ‘(’ N_1 ‘,’ N_2 ‘)’>
N -> <‘Sif’, ‘sif’>
N -> <‘Felix’, ‘felix’>
V -> <‘chases’, ‘chases’>

- This is a synchronous context free grammar (SCFG).
- We derive/parse the sentence and logical form simultaneously:

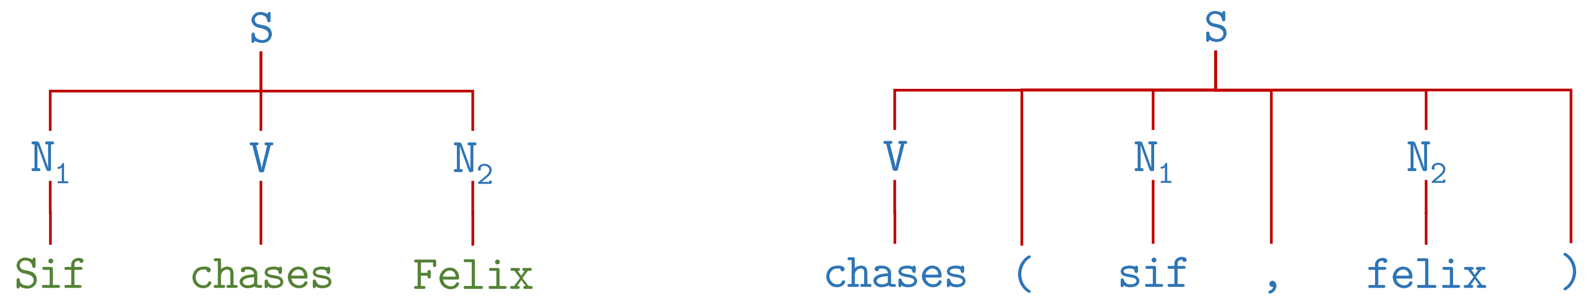

PARSING WITH SYNCHRONOUS GRAMMARS

- In semantic parsing, however, we only have ‘`Sif chases Felix`’.
  - How do we obtain the logical form?
- Consider the grammar:

```
S -> <N_1 V N_2, V ‘(’ N_1 ‘,’ N_2 ‘)’>
N -> <‘Sif’, ‘sif’>
N -> <‘Felix’, ‘felix’>
V -> <‘chases’, ‘chases’>
```

- Focus on just the natural language part of the grammar:

```
S -> N_1 V N_2
N -> ‘Sif’
N -> ‘Felix’
V -> ‘chases’
```

PARSING WITH SYNCHRONOUS GRAMMARS

- Use any CFG parsing method to parse '`Sif chases Felix`'.
  - E.g., Earley parsing

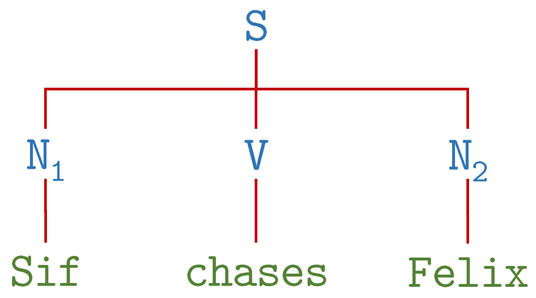

- Then we can reconstruct the derivation tree for the logical form by inspecting each rule in the above tree.
  - For each rule, we look at the right-hand side to determine how to construct the logical form.

PARSING WITH SYNCHRONOUS GRAMMARS

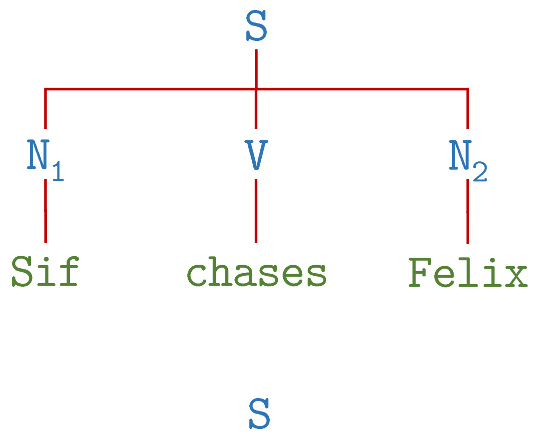


```
S -> <N1 V N2, V '(' N1 ',' N2 ')'>
N -> <'Sif', 'sif'>
N -> <'Felix', 'felix'>
V -> <'chases', 'chases'>
```

PARSING WITH SYNCHRONOUS GRAMMARS

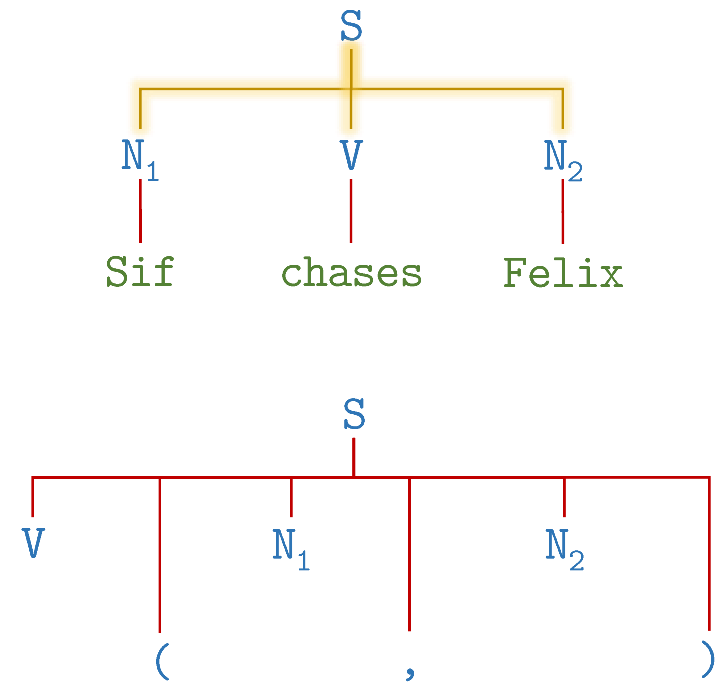

S -> <N_1 V N_2, V ‘(’ N_1 ‘,’ N_2 ‘)’>
N -> <‘Sif’, ‘sif’>
N -> <‘Felix’, ‘felix’>
V -> <‘chases’, ‘chases’>

PARSING WITH SYNCHRONOUS GRAMMARS

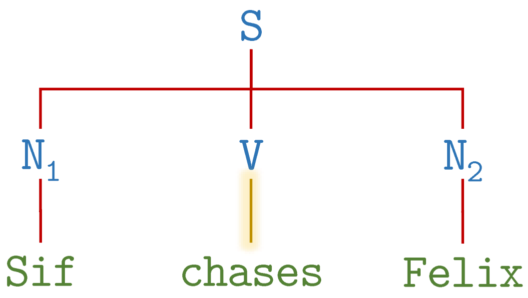

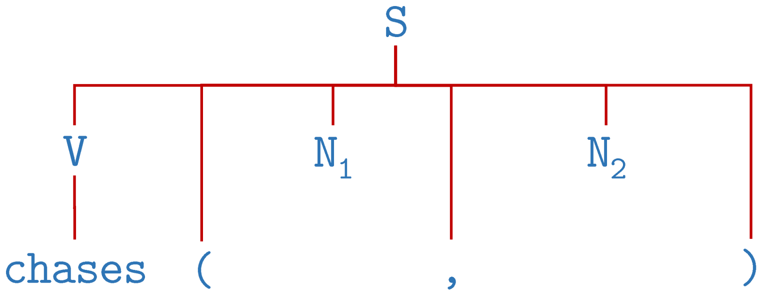

S -> <N_1 V N_2, V '(' N_1 ',' N_2 ')'>

N -> <'Sif', 'sif'>

N -> <'Felix', 'felix'>

V -> <'chases', 'chases'>

PARSING WITH SYNCHRONOUS GRAMMARS

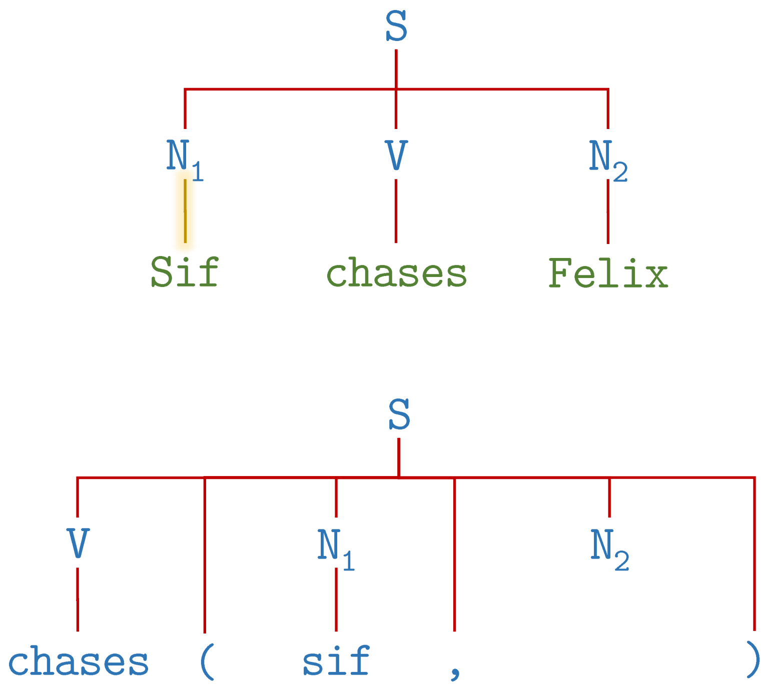


```
S -> <N1 V N2, V ‘(’ N1 ‘,’ N2 ‘)’>
N -> <‘Sif’, ‘sif’>
N -> <‘Felix’, ‘felix’>
V -> <‘chases’, ‘chases’>
```

PARSING WITH SYNCHRONOUS GRAMMARS

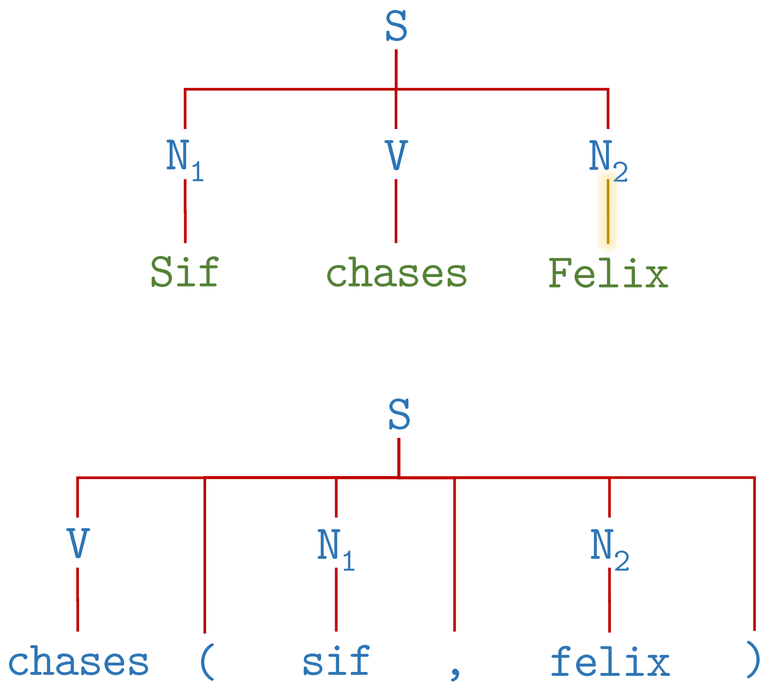


```
S -> <N1 V N2, V ‘(’ N1 ‘,’ N2 ‘)’>
N -> <‘Sif’, ‘sif’>
N -> <‘Felix’, ‘felix’>
V -> <‘chases’, ‘chases’>
```

PARSING WITH SYNCHRONOUS GRAMMARS

- Note that, when we reconstruct the logical form, there may be more than one matching rule.
- Consider the slightly modified grammar:

```
S -> <N1 V N2, V '(' N1 ',' N2 ')'>
N -> <'Sif', 'sif'>
N -> <'Felix', 'felix_lee'>
N -> <'Felix', 'felix_mendelssohn'>
V -> <'chases', 'chases'>
```

- When we parse 'Sif chases Felix' and inspect the rule N -> 'Felix', there are two matching rules.
  - We can choose either rule to produce a valid logical form.
- There are two valid logical forms:
  - chases(sif,felix_lee) and chases(sif,felix_mendelssohn)
  - Example of semantic ambiguity.
- SCFG can capture both syntactic and semantic ambiguity.

PRAGMATICS

PRAGMATICS

- Our discussion on semantics has focused on the meaning of individual sentences.
- In reality, sentences almost never appear in isolation.
  - Natural language sentences have context.
  - Surrounding sentences can change the meaning of other sentences.
  - E.g., ‘`I was so focused, I didn’t realize how much time had passed. Time flies like an arrow.`’
  - vs ‘`The wizard studied flies that possessed the ability to travel through time. He held an arrow in his hand and the flies were attracted to it. Time flies like an arrow.`’

PRAGMATICS

- The sentence ‘`Time flies like an arrow`’ has multiple ambiguous interpretations.
  - Each interpretation can be represented with a distinct logical form.
- The surrounding context helps us to determine which interpretation is most likely.
- Consider another example:
  - ‘`Everyone looked up.`’
  - ‘`During gym class, Alice kicked the ball high into the air. Everyone looked up.`’
  - ‘`The supervillain turned on the machine, and the sky began to change color everywhere. Everyone looked up.`’

PRAGMATICS

- ‘`Everyone`’ in ‘`Everyone looked up`’ is an example of underspecification.
  - We didn’t specify whether it’s ‘`Everyone in the gym class`’ or ‘`Everyone outside in the world`’.
  - We were able to infer that information from the context.
- Why do we leave out this information from sentences?

PRAGMATICS

- We can think of human language as a product of optimization:
  - Humans want to communicate information to each other.
  - But we want to minimize the energy cost of the communication.
    - Talking a lot is tiring.
  - We want to maximize information transfer while minimizing energy cost.
- Speakers want to minimize the number of words they produce.
- Listeners want to make sure they have accurately received the information.
  - But accurate information transfer typically requires more words.
- Thus, natural language strikes a balance between using enough words to convey the correct information, but no more than that.
  - So we avoid repeating information if it can be inferred from context.

GRICE'S MAXIMS

- Paul Grice (1975) posited that human languages follow four 'maxims'.

1. Maxim of quantity:
   - Be informative
   - Make your contribution as informative as required, but no more.
2. Maxim of quality:
   - Be truthful
   - If you are unsure or don't know, say so.
3. Maxim of relevance:
   - Be relevant

GRICE'S MAXIMS

- Paul Grice (1975) posited that human languages follow four 'maxims'.

4. Maxim of manner:
    - Try to ensure the listener understands you.
    - Avoid language that's difficult to understand for the listener.
    - Avoid ambiguity
        - Some ambiguity is okay, if the correct interpretation is obvious to the listener from context.
        - But if not, the listener might ask you to clarify, which increases the cost of communication.
    - Try to consider what the listener knows and doesn't know.

GRICE’S MAXIMS

- These maxims are not axioms or theorems.
  - They are not universally true.
- E.g., What if your goal is deception?
  - You wouldn’t want to follow the maxim of quality.
- If you want to change the subject, then you shouldn’t follow the maxim of relevance.

PRAGMATICS IN NLP

- If an NLP model is to accurately perform tasks that involve context, such as long-document question answering, it needs to be able to correctly model how context affects the meaning of sentences.
  - And correctly infer the full meaning of underspecified sentences.

- E.g.,
  - '...During gym class, Alice kicked the ball high into the air. Everyone looked up and saw a falcon grab the ball... Jay didn't come to school that day since he wasn't feeling well...'

    (the story could be much longer)
  - 'Did Jay see the falcon?'

QUESTIONS?