NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Semantic Analysis

Semantic analysis is the process in Natural Language Processing (NLP) that determines the meaning of words, phrases, and sentences in a given context. It involves mapping the syntactic structure of a sentence to a formal representation of its meaning so that a computer can understand and reason about the information expressed in natural language.

Need for Semantic Representations

- Verifiability
- Unambiguous Representations
- Canonical Form
- Inference and Variables
- Expressiveness

Verifiability

Example: Does Maharani Serve Vegetarian Food?

Requires: Knowledge base

Representation: *Serves*(*Maharani,VegetarianFood*)

If the system finds a representation matching the input proposition in its knowledge base, it can return an affirmative answer. Otherwise, it must either say *No*, if its knowledge of local restaurants is complete, or say that it does not know if there is reason to believe that its knowledge is incomplete. This notion is known as **verifiability.**

Unambiguous Representations

Example: I wanna eat someplace that's close to ICSI.

The verb eat and someplace raises the ambiguity as they are not providing information on where to eat and what to eat.

Example: I want to eat Italian food.

The sentence have some vagueness although it is not ambiguous as there is some uncertainity regarding which Italian food he wants to eat.

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Canonical Form

Example:

Does Maharani have vegetarian dishes?

Do they have vegetarian food at Maharani?

Are vegetarian dishes served at Maharani?

Does Maharani serve vegetarian fare?

Since the answers given for each of these alternatives should be the same in all situations, we might say that they all mean the same thing, at least for the purposes of giving restaurant recommendations. In other words, at least in this domain, we can legitimately consider assigning the same meaning representation to the propositions underlying each of these requests. Taking such an approach would guarantee that our matching scheme for answering Yes-No questions will still work. The notion that inputs that mean the same thing should have the same meaning representation is known as the doctrine of **canonical form**. Canonical form does, of course, complicate the task of semantic analysis.

Inference and Variables

Example: Can vegetarians eat at Maharani?

Can we use canonical Form?

Here, it would be a mistake to invoke canonical form to force our system to assign the same representation to this request as for the previous examples. The fact that this request results in the same answer as the others arises not because they mean the same thing, but because there is a commonsense connection between what vegetarians eat and what vegetarian restaurants serve.

The term inference refer generically to a system's ability to draw valid conclusions based on the meaning representation of inputs and its store of background knowledge. It must be possible for the system to draw conclusions about the truth of propositions that are not explicitly represented in the knowledge base, but are nevertheless logically derivable from the propositions that are present.

Example: I'd like to find a restaurant where I can ger vegetarian food

This request does not make reference to any particular restaurant. The user is stating that they would like information about an unknown and unnamed entity that is a restaurant that serves

vegetarian food. Since this request does not mention any particular restaurant, the kind of simple matching based approach we have been advocating is not going to work. Rather, answering this request requires a more complex kind of matching that involves the use of variables.

Serves(*x,VegetarianFood*)

Expressiveness

Finally, to be useful a meaning representation scheme must be expressive enough to handle an extremely wide range of subject matter. The ideal situation, of course, would be to have a single meaning representation language that could adequately represent the meaning of any sensible natural language utterance.

First Order Logic (FOL)

First-Order Logic (FOL), also called predicate logic, is a formal system used to represent objects, properties, relationships, and rules in a domain so that computers can perform reasoning and inference. In Natural Language Processing, FOL is used to represent the meaning of sentences in a structured and machine-interpretable form.

Example: John Love Mary

Representation: Love(John,Mary)

John → subject (agent)

Love → relation/action

Mary → object

Components of FOL

- Constants
- Variables
- Predicates
- Functions
- Logical Connectives
- Quantifiers

Constants

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Constants represent **specific objects or entities** in the domain.

Examples: John, Mary, Paris, India, Computer

Example: John is a student.

Representation: Student(John)

John is a constant.

Variables

Variables represent **unknown or general objects**.

Examples: x, y, z

Example: Student(x)

Meaning: x is some student.

Predicates

Predicates represent **properties of objects or relationships between objects**.

Examples: Student(x), Teacher(x), Loves(John, Mary), LivesIn(Ravi, India)

Types of predicates

- Unary
- Binary
- Ternary

1. Unary predicate (one argument)

- Student(John)
- Meaning: John is a student.

2. Binary predicate (two arguments)

- Likes(John, IceCream)
- Meaning: John likes ice cream.

3. Ternary predicate (three arguments)

- Give(Ravi, Book, Sita)

- Meaning: Ravi gave a book to Sita.

Functions

Functions return a **value related to an object**.

Example

- FatherOf(John)
- Meaning: The father of John.
- Example sentence
- The father of John is a doctor.
- Representation: Doctor(FatherOf(John))

Logical Connectives

Logical connectives combine logical statements.

Symbol	Meaning
∧	AND
∨	OR
¬	NOT
→	IMPLIES
↔	IF AND ONLY IF

Example using AND (∧)

- Sentence: John is a student and he is intelligent.
- Representation: Student(John) ∧ Intelligent(John)

Example using OR (∨)

- Sentence: John is a student or a teacher.
- Representation: Student(John) ∨ Teacher(John)

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Example using NOT (¬)

- Sentence: John is not a teacher.
- Representation: ¬Teacher(John)

Example using implication (→)

- Sentence: If someone is a student, then they study.
- Representation: Student(x) → Studies(x)
- Meaning: If x is a student, then x studies.

Quantifiers

Quantifiers express **how many objects satisfy a condition**.

Types:

- Universal Quantifier
- Existential Quantifier

Universal Quantifier (∀)

- Symbol: ∀
- Meaning: For all
- Example sentence: All humans are mortal.
- Representation: ∀x Human(x) → Mortal(x)
- Explanation: For every x, if x is human, then x is mortal.
- **Example:** All students study.
- Representation: ∀x Student(x) → Study(x)
- Meaning: Every student studies.

Existential Quantifier (∃)

- Symbol: ∃
- Meaning: There exists
- Example sentence: There is a student in the class.
- Representation: ∃x Student(x)
- Meaning: At least one student exists.
- **Example**
- Sentence: Some students like mathematics.

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

- Representation: ∃x Student(x) ∧ Likes(x, Mathematics)
- Meaning:There exists a student who likes mathematics.

Representing Events

- Events in NLP can also be represented using predicates.
- Example sentence
- John cut the cake with a knife.
- Representation
- Cut(John, Cake, Knife)
- This structure clearly shows **who performed the action and what objects were involved**.

Entity	Role
John	Agent
Cake	Theme
Knife	Instrument

Inference Using First Order Logic

FOL allows systems to **derive new knowledge from existing facts**.

Example

- Knowledge base
- ∀x Human(x) → Mortal(x)
  Human(Socrates)
- Inference: Mortal(Socrates)
- Explanation: Since all humans are mortal and Socrates is human, Socrates must be mortal.

Advantages of FOL

- Precise representation of meaning-Captures relationships between objects clearly.
- Supports reasoning and inference-Allows systems to derive new knowledge.

- Handles complex sentences-Can represent nested relationships and conditions.
- Used in knowledge bases and expert systems-Foundation for semantic analysis in NLP

Disadvantages of FOL

- Natural language meanings are sometimes too complex or ambiguous.
- Difficult to represent uncertainty or probability.
- Requires large knowledge bases.
- Not always efficient for large-scale NLP tasks.

Description Logic

Description Logics (DL) are a family of formal knowledge representation languages used to represent concepts (classes), relationships, and individuals in a domain. DL is mainly used in:

- Knowledge representation
- Ontology modeling
- Semantic Web
- Natural Language Processing
- It allows computers to describe structured knowledge and reason about it automatically.

Components of DL

- Concepts
- Roles
- Individuals

Concepts (Classes)

A **concept** represents a **set or category of objects** in a domain. Concepts are similar to **classes in object-oriented programming**.

- **Examples:** Person, Student, Teacher, Animal, Vehicle
- Example statement: Student ⊆ Person
- Meaning: All students are persons.
- **Example Representation**

- Sentence: Ravi is a student.
- Representation:Student(Ravi)
- Explanation: Ravi belongs to the concept **Student**.

Roles (Relationships)

A **role** represents a **relationship between two concepts or individuals**. Roles are similar to **relations in databases**.

- **Examples:** hasChild, teaches, owns, livesIn, drives
- **Example**
- Sentence: John teaches Ravi.
- Representation: teaches(John, Ravi)
- Here
- **John** → teacher
- **Ravi** → student
- **teaches** → role
- **Example with Concepts**
- Teacher ⊆ Person
  Student ⊆ Person
- Relationship: teaches(Teacher, Student)
- Meaning: Teachers teach students.

Individuals (Instances)

- Individuals represent **specific objects or entities in the domain**.
- **Examples:** John, Ravi, Sita, Paris, Car1
- **Example**
- Student(Ravi)
  Teacher(John)
- Meaning
- Ravi is a student
- John is a teacher

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

DL Knowledge Base

A DL knowledge base consists of two parts:

- **TBox (Terminological Box)**
- **ABox (Assertional Box)**

Terminological Box (TBox)

TBox contains **general knowledge about concepts and relationships**.

- Example
- Student ⊆ Person
  Teacher ⊆ Person
  Doctor ⊆ Person
- Meaning
- Students are persons
- Teachers are persons
- **Example:** Parent ≡ Person ⊓ ∃hasChild.Person
- Meaning: A parent is a person who has at least one child.

Assertional Knowledge Box (ABox)

ABox contains **specific facts about individuals**.

- Example
- Student(Ravi)
  Teacher(John)
  hasChild(John, Ravi)
- Meaning
- Ravi is a student
- John is a teacher
- John has a child Ravi

Concept Constructors in DL

DL provides operators to **create complex concepts**.

- Conjunction (AND)

- Disjunction (OR)
- Negation (NOT)
- Existential Restriction
- Universal Restriction

Conjunction (AND)

- Symbol: ⊓
- Example: Student ⊓ Athlete
- Meaning: Students who are athletes.
- Example sentence
- Ravi is a student and athlete.
- Representation
- (Student ⊓ Athlete)(Ravi)

Disjunction (OR)

- Symbol: ⊔
- Example: Student ⊔ Teacher
- Meaning: A person who is either student or teacher.

Negation (NOT)

- Symbol: ¬
- Example: ¬Student
- Meaning: Not a student.
- Example: ¬Student(John)
- Meaning: John is not a student.

Existential Restriction

- Symbol:∃
- Example: ∃hasChild.Person
- Meaning: A person who has at least one child.

Universal Restriction

- Symbol: ∀
- Example: ∀teaches.Student
- Meaning: A person who teaches only students.

Advantages of DL

- Structured knowledge representation
- Supports reasoning and inference
- Used in ontologies like **OWL**
- Helps build semantic knowledge bases
- Useful for NLP semantic understanding

Disadvantages of DL

- Cannot easily represent uncertainty
- Complex reasoning for large knowledge bases
- Limited ability for procedural knowledge

Syntax-Driven Semantic Analysis

Syntax-driven semantic analysis is the method of computing the meaning of a sentence based on its syntactic structure. The syntactic parse tree guides the semantic interpretation, and semantic rules are attached to grammar rules. Thus: Meaning of a sentence = Meaning of its parts + Syntactic structure. This approach is widely used in semantic parsing and compositional semantics.

Principle of Compositionability

The Principle of Compositionality states: The meaning of a sentence is determined by the meanings of its individual words and the way they are syntactically combined. This principle is fundamental in syntax-driven semantic analysis.

- Example
- Sentence: The dog chased the cat
- Word meanings: dog cat chase
- Semantic composition: chase(dog, cat)
- The meaning is obtained by combining the meanings of the words using syntactic structure.

Semantic Augmentations to CFG Rules

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Context-Free Grammar (CFG) describes sentence structure. Example rule: VP → V NP**.** To build meaning, CFG rules are **augmented with semantic rules**.

- **Example**
- Grammar rule: VP → V NP
- Semantic rule: VP.sem = V.sem(NP.sem)
- Meaning::The verb meaning is applied to the noun phrase meaning.
- **Example Sentence:** Mary likes chocolate
- Parse: VP → likes chocolate
- Semantic representation: like(Mary, chocolate)

Quantifier Scope Ambiguity

- Sentences with quantifiers may have **multiple interpretations depending on scope**.
- Example sentence: Every student read a book
- Two possible meanings:
- **Interpretation 1:** Each student read possibly different books.
- ∀x Student(x) → ∃y Book(y) ∧ Read(x,y)
- **Interpretation 2:** There is one book that all students read.
- ∃y Book(y) ∧ ∀x Student(x) → Read(x,y)
- Both interpretations are grammatically valid.

Unification Based approaches

A **unification-based approach** combines syntactic and semantic features using **feature structures**. These feature structures are unified during parsing.

- **Feature Structure Example**
- [CAT = NP
  SEM = John]
- **Example Grammar Rule:** S → NP VP
- Semantic unification: S.sem = VP.sem(NP.sem)
- **Example Sentence:** John runs
- Unification result: run(John)

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Semantic Attachments

Semantic attachments associate **semantic interpretation rules with grammar rules**. These rules allow computing meaning during parsing. We illustrate semantic attachments for common phrase types.

Sentences

- Grammar rule
- S → NP VP
- Semantic rule
- S.sem = VP.sem(NP.sem)
- Example
- John sleeps
- Representation
- sleep(John)

Noun and Genitive Noun Phrases

- **Noun Phrase**
- Example: the dog
- Semantic representation: dog(x)
- **Genitive Noun Phrase**
- Example: John's book
- Semantic Representation: book(y) ∧ possess(John,y)
- Meaning: The book belongs to John.

Adjective Phrases

- Adjectives modify nouns.
- Example: big dog
- Semantic representation: dog(x) ∧ big(x)
- Example
- red car
- Representation: car(x) ∧ red(x)

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Verb Phrases

Verb phrases represent **actions or events**.

- Example: Mary likes chocolate
- Representation: like(Mary, chocolate)
- Roles:

Entity	Role
Mary	Agent
Chocolate	Theme

Infinite Verb Phrases

- An **infinitive verb phrase** contains the word **to**.
- Example: John wants to eat pizza
- Representation: want(John, eat(John,pizza))
- Meaning: John desires the action of eating pizza.

Prepositional Phrases

- Prepositional phrases express **relationships such as location, instrument, or time**.
- Example: The book is on the table
- Representation: on(book, table)

Nominal Modifier Prepositional Phrases

- These modify **nouns**.
- Example: the book on the table
- Meaning: book(x) ∧ on(x, table)
- The phrase **on the table** modifies **book**.

Verb Phrase Modifier Prepositional Phrases

- These modify **verbs or actions**.

- Example: John ran in the park
- Representation: run(John) ∧ location(run, park)
- Meaning: The action occurred in the park.

Verb Argument Prepositional Phrases

- Some verbs require prepositional phrases as **arguments**.
- Example: John put the book on the table
- Representation: put(John, book, table)
- Here: on the table is a required argument of the verb **put**.

Idioms and Compositionability

- Idioms are expressions whose meanings **cannot be derived from individual word meanings**.
- Example
- kick the bucket
- Literal meaning
- kick(bucket)
- Actual meaning
- die

Challenges with Idioms

- They violate the **principle of compositionality**.
- Meaning must be stored as a **single lexical unit**.
- **Example**
- Sentence
- He kicked the bucket
- Interpretation
- die(he)

Lexical Semantics

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Lexical semantics is the branch of linguistics and NLP that studies the meaning of words and the relationships between word meanings. In natural language processing, understanding the meaning of words is essential for tasks such as:

- Machine translation
- Information retrieval
- Question answering
- Text summarization
- Word sense disambiguation

Word Sense

word sense is one of the possible meanings of a word. Many words are **ambiguous** because they have multiple senses.

- **Example Word:** bank
- Possible meanings: Financial institution, Edge of a river
- **Example Sentences**
- Sentence 1: She deposited money in the bank.
- Meaning: bank = financial institution
- Sentence 2: They sat on the bank of the river.
- Meaning: bank = river edge
- Thus, the correct sense depends on the **context**.

Types of Word Senses

- Polysemy
- Homonymy
- Homophones
- Homographs
- Metonymy

Polysemy

NLP UNIT-4: Requirements for representation, First-Order Logic, Description Logics – Syntax-Driven Semantic analysis, Semantic attachments – Word Senses, Relations between Senses, Thematic Roles, selectional restrictions

Polysemy occurs when a word has **multiple related meanings**. The meanings share a **common semantic origin**.

- **Example Word:** paper
- Possible meanings: Writing material, Academic article, Newspaper
- **Example Sentences**
- She wrote her notes on paper.
- He published a research paper.
- I read the morning paper.

All senses relate to **written information**.

Homonymy

Homonymy occurs when two words share the **same spelling or pronunciation but have unrelated meanings**.

- **Example Word:** bat
- Meaning 1: flying mammal
- Meaning 2: sports equipment
- **Example Sentences**
- The bat flew out of the cave.
- He hit the ball with the bat.

Homophones

Homophones are words that **sound the same but have different meanings and spellings**.

- **Example**
- flower – flour
- Examples
- She bought a flower.
- The cake needs flour.

Homographs

Homographs are words that have the **same spelling but different meanings and sometimes different pronunciations**.

- **Example**
- lead
- Meaning 1: metal
- Meaning 2: to guide
- Examples
- The pipes are made of lead.
- She will lead the team.

Metonymy

Metonymy occurs when a word is used to refer to something **closely associated with it**, rather than the thing itself. In metonymy, one entity **stands for another related entity**.

- **Example**
- Sentence: The White House announced new policies.
- Meaning: White House → U.S. government administration
- **Example**
- Hollywood produces many movies.
- Meaning: Hollywood → American film industry
- Metonymy is common in everyday language and must be handled by NLP systems.

Relationships Between Word Senses

Lexical semantics studies **relationships between meanings of words**. These relationships help organize vocabulary in **semantic networks**.

Types of Relationships:

- Synonomy
- Antonymy
- Hypernymy
- Hyponymy
- Meronymy
- Holonymy

- Semantic Fields

Synonomy

- **Synonyms** are words with **similar meanings**
- **Examples**
- big – large
  car – automobile
- Example sentences
- This is a big building.
- This is a large building.

Antonymy

- **Antonyms** are words with **opposite meanings**.
- **Examples**
- hot – cold
  happy – sad
- Example
- The coffee is hot.
- The ice cream is cold.

Hypernymy (Is-A Relationship)

- A **hypernym** is a **more general term representing a category**.
- **Example**
- animal → dog
- Meaning
- A dog is a type of animal.

Hyponymy

- A **hyponym** is a **more specific word within a category**.
- **Example**
- dog → animal
- Example hierarchy

- vehicle
  ├── car
  ├── bus
  └── truck

Meronymy (Part-Whole Relationship)

- **Meronymy** represents a **part-of relationship**
- **Examples**
- wheel → car
  finger → hand
  page → book
  engine → car
- Example sentence
- A wheel is part of a car.

Holonymy

- Holonymy represents the **whole-to-part relationship**.
- Example
- car → wheel
- Meaning
- A car has wheels.

Semantic Fields

A **semantic field** is a group of words that are **related in meaning and belong to the same conceptual domain**. Words within a semantic field share a **common conceptual area**.

- **Example: Color Field**
- red
  blue
  green
  yellow
- All belong to the **color semantic field**.
- **Example: Transportation Field**
- car
  bus
  train

airplane
bicycle

- All belong to the **transportation semantic field**.
- **Example: Emotion Field**
- happy
  sad
  angry
  excited
  afraid
- These words describe **emotions**.
- Semantic fields help NLP systems organize vocabulary into **meaningful conceptual groups**.

Semantic (Thematic Roles)

Consider these two sentences.

- Sasha broke the window.
- Pat opened the door.

What can you infer?

The subjects Sasha and Pat, what we might call the breaker of the window-breaking event and the opener of the door-opening event, have something in common. They are both volitional actors, often animate, and they have direct causal responsibility for their events.

Importance of Semantic (Thematic) Roles

Semantic roles are important in NLP tasks such as:

- Question answering
- Information extraction
- Machine translation
- Text summarization
- Dialogue systems

Thematic roles

Thematic roles are a way to capture this semantic commonality between breakers and openers. We say that the subjects of both these verbs are agents. Thus,is the thematic role that represents an abstract idea such as volitional causation. Similarly, the direct objects of

both these verbs, the BrokenThing and OpenedThing,are both prototypically inanimate objects that are affected in some way by the action. The semantic role for these participants is theme.

Thematic Role	Definition
AGENT	The volitional causer of an event
EXPERIENCER	The experiencer of an event
FORCE	The non-volitional causer of the event
THEME	The participant most directly affected by an event
RESULT	The end product of an event
CONTENT	The proposition or content of a propositional event
INSTRUMENT	An instrument used in an event
BENEFICIARY	The beneficiary of an event
SOURCE	The origin of the object of a transfer event
GOAL	The destination of an object of a transfer event

Thematic Role	Example
AGENT	*The waiter* spilled the soup.
EXPERIENCER	*John* has a headache.
FORCE	*The wind* blows debris from the mall into our yards.
THEME	Only after Benjamin Franklin broke *the ice...*
RESULT	The city built a *regulation-size baseball diamond...*
CONTENT	Mona asked *"You met Mary Ann at a supermarket?"*
INSTRUMENT	He poached catfish, stunning them *with a shocking device...*
BENEFICIARY	Whenever Ann Callahan makes hotel reservations *for her boss...*
SOURCE	I flew in *from Boston.*
GOAL	I drove *to Portland.*

Diathesis (Verb) Alterations

Diathesis alternations are changes in the syntactic structure of a sentence that preserve the same underlying meaning. In other words, the same event can be expressed using different syntactic patterns.

- ohn broke the window.
- AGENT THEME
- John broke the window with a rock.
- AGENT THEME INSTRUMENT
- The rock broke the window.
- INSTRUMENT THEME

- The window broke.
- THEME
- The window was broken by John.
- THEME AGENT

The set of thematic role arguments taken by a verb is thematic grid often called the thematic grid, θ-grid, or case frame. We can see that there are case frame (among others) the following possibilities for the realization of these arguments of break:

- AGENT/Subject, THEME/Object
- AGENT/Subject, THEME/Object, INSTRUMENT/PPwith
- INSTRUMENT/Subject, THEME/Object
- THEME/Subject

Problems with Thematic Roles

Although thematic roles are useful, they have several limitations.

- **Lack of Clear Definitions** Some roles are difficult to define precisely.

  Example: Mary opened the door with a key. Is **key** an instrument or a cause? Different interpretations are possible.

- **Too Many Possible Roles** Many roles exist, and it is difficult to create a **fixed universal set**. Examples include:
- Agent
- Experiencer
- Beneficiary
- Stimulus
- Cause
- Instrument

  The boundaries between these roles are often unclear.

- **Verb-Specific Differences**

  Different verbs may require **different role interpretations**.

  Example

- John broke the window.
- vs.
- John saw the window.
- In the first case John is **Agent**, in the second he is **Experiencer**.

- **Ambiguity**
- Example
- The door opened with the key.
- Is the key an **instrument** or a **cause**?
- This ambiguity creates problems for semantic role labeling.

Selection Restrictions

A selectional restriction is a semantic type constraint that a verb imposes on the kind of concepts that are allowed to fill its argument roles.

- I want to eat someplace nearby.
- I want to eat Malaysian food.
- The restaurant serves green-lipped mussels.
- Which airlines serve Denver?
- What can you infer from above sentences
- Sentence1: the sensible interpretation, eat is intransitive and the phrase someplace nearby is an adjunct that gives the location of the eating event.
- Sentence2: In the nonsensical speaker-as-Godzilla interpretation, eat is transitive and the phrase someplace nearby is the direct object and the THEME of the eating, like the NP Malaysian food.
- Sentence3: illustrates the offering-food sense of serve, which ordinarily restricts its THEME to be some kind of food Example.
- Sentence4: illustrates the provides a commercial service to sense of serve, which constrains its THEME to be some type of appropriate location.

Representing Selection Restrictions

The neo-Davidsonian representation of an event consists of a single variable that stands for the event, a predicate denoting the kind of event, and variables and relations for the event

roles. Ignoring the issue of the λ-structures and using thematic roles rather than deep event roles, the semantic contribution of a verb like eat might look like the following:

$$\exists e,x,y\ Eating(e) \wedge Agent(e,x) \wedge Theme(e,y)$$

With this representation, all we know about y, the filler of the THEME role, is that it is associated with an Eating event through the Theme relation.

To stipulate the selectional restriction that y must be something edible, we simply add a new term to that effect:

$$\exists e,x,y\ Eating(e) \wedge Agent(e,x) \wedge Theme(e,y) \wedge Isa(y,EdibleThing)$$

When a phrase like ate a hamburger is encountered, a semantic analyzer can form the following kind of representation:

$$\exists e,x,y\ Eating(e) \wedge Eater(e,x) \wedge Theme(e,y) \wedge Isa(y,EdibleThing) \wedge Isa(y,Hamburger)$$

This representation is perfectly reasonable since the membership of y in the category Hamburger is consistent with its membership in the category EdibleThing, assuming a reasonable set of facts in the knowledge base.

A more practical approach is to state selectional restrictions in terms ofWordNet synsets rather than as logical concepts. Each predicate simply specifies a WordNet synset as the selectional restriction on each of its arguments. A meaning representation is well-formed if the role filler word is a hyponym (subordinate) of this synset. For our ate a hamburger example, for instance, we could set the selectional restriction on the THEME role of the verb eat to the synset food, nutrients, glossed as any substance that can be metabolized by an animal to give energy and build tissue. Again, the filler of a role need not match the restriction synset exactly; it just needs to have the synset as one of its superordinates. We can apply this approach to the THEME roles of the verbs imagine, lift, and diagonalize, discussed earlier. Let us restrict imagine's THEME to the synset fentity, lift's THEME to physical entity, and diagonalize to matrix. This arrangement correctly permits imagine a hamburger and lift a hamburger, while also correctly ruling out diagonalize a hamburger.

Selection Preferences

In the earliest implementations, selectional restrictions were considered strict constraints on the kind of arguments a predicate could take. For example, the verb eat might require that its THEME argument be [+FOOD]. Early word sense disambiguation systems used this idea to rule out senses that violated the selectional restrictions of their governing predicates. Very quickly, however, it became clear that these selectional restrictions were better represented as

preferences rather than strict constraints. For example, selectional restriction violations (like inedible arguments of eat) often occur in well-formed sentences, for example because they are negated or because selectional restrictions are overstated.

Examples:

- But it fell apart in 1931, perhaps because people realized you can't eat gold for lunch if you're hungry.
- In his two championship trials, Mr. Kulkarni ate glass on an empty stomach, accompanied only by water and tea.

Selection Associations

One of the most influential has been the selectional association model of Resnik. Resnik defines the idea of selectional preference strength as the general strength amount of information that a predicate tells us about the semantic class of its arguments. For example, the verb eat tells us a lot about the semantic class of its direct objects, since they tend to be edible. The verb be, by contrast, tells us less about its direct objects. The selectional preference strength can be defined by the difference in information between two distributions: the distribution of expected semantic classes P(c) (how likely is it that a direct object will fall into class c) and the distribution of expected semantic classes for the particular verb P(c|v) (how likely is it that the direct object of the specific verb v will fall into semantic class c). The greater the difference between these distributions, the more information the verb is giving us about possible objects.

The Kullback-Leibler or KL divergence D(P||Q) expresses the difference between two probability distributions P and Q

$$D(P||Q) = \sum_x P(x) \log \frac{P(x)}{Q(x)}$$

The selectional preference $S_R(v)$ uses the KL divergence to express how much information, in bits, the verb v expresses about the possible semantic class of its argument.

$$\begin{aligned} S_R(v) &= D(P(c|v)||P(c)) \\ &= \sum_c P(c|v) \log \frac{P(c|v)}{P(c)} \end{aligned}$$

Resnik then defines the selectional association of a particular class and verb as the relative contribution of that class to the general selectional preference of the verb:

$$A_R(v,c) = \frac{1}{S_R(v)} P(c|v) \log \frac{P(c|v)}{P(c)}$$

Verb	Direct Object Semantic Class	Assoc	Direct Object Semantic Class	Assoc
read	WRITING	6.80	ACTIVITY	-.20
write	WRITING	7.26	COMMERCE	0
see	ENTITY	5.79	METHOD	-0.01

Selection Preference Via Conditional Probability

An alternative to using selectional association between a verb and the WordNet class of its arguments is to use the conditional probability of an argument word given a predicate verb, directly modeling the strength of association of one verb (predicate) with one noun (argument). The conditional probability model can be computed by parsing a very large corpus (billions of words), and computing co-occurrence counts: how often a given verb occurs with a given noun in a given relation.

The conditional probability of an argument noun given a verb for a particular relation P(n|v, r) can then be used as a selectional preference metric for that pair of words

$$P(n|v,r) = \begin{cases} \frac{C(n,v,r)}{C(v,r)} & \text{if } C(n,v,r) > 0 \\ 0 & \text{otherwise} \end{cases}$$

The inverse probability P(v|n, r) was found to have better performance in some cases

$$P(v|n,r) = \begin{cases} \frac{C(n,v,r)}{C(n,r)} & \text{if } C(n,v,r) > 0 \\ 0 & \text{otherwise} \end{cases}$$

An even simpler approach is to use the simple log co-occurrence frequency of the predicate with the argument logcount(v,n, r) instead of conditional probability as this seems to do better for extracting preferences for syntactic subjects rather than objects.

Evaluating Selection Preferences

One way to evaluate models of selectional preferences is to use pseudowords. A pseudoword is an artificial word created by concatenating a test word in some context (say banana) with a

confounder word (say door) to create banana-door). The task of the system is to identify which of the two words is the original word.

To evaluate a selectional preference model (for example on the relationship between a verb and a direct object) we take a test corpus and select all verb tokens. For each verb token (say drive) we select the direct object (e.g., car), concatenated with a confounder word that is its nearest neighbor, the noun with the frequency closest to the original (say house), to make car/house). We then use the selectional preference model to choose which of car and house are more preferred objects of drive, and compute how often the model chooses the correct original object (e.g., car).

Another evaluation metric is to get human preferences for a test set of verb argument pairs, and have them rate their degree of plausibility. This is usually done by using magnitude estimation, a technique from psychophysics, in which subjects rate the plausibility of an argument proportional to a modulus item. A selectional preference model can then be evaluated by its correlation with the human preferences.