Relational semantics

CS 690N, Spring 2017

Advanced Natural Language Processing

http://people.cs.umass.edu/~brenocon/anlp2017/

Brendan O'Connor

College of Information and Computer Sciences

University of Massachusetts Amherst

Why semantics?

Goal is to convert text into structured knowledge representations.
Some motivations:

- Automatically update databases of facts
- Infer new facts and relationships
- Answer complex questions, e.g.,
  what cheese-exporting countries are hereditary monarchies?
- Logic-check written arguments
- ...

[Slides: Jacob Eisenstein]

Why semantics?

Semantics is a stumbling block for NLP at all levels:

- I shot an elephant in my pajamas
- How to solve PP attachment question?
- Bilexical probabilities are just a noisy approximation

[Slides: Jacob Eisenstein]

Can your computer ever *really* understand you?

What does it really mean to understand language anyway?

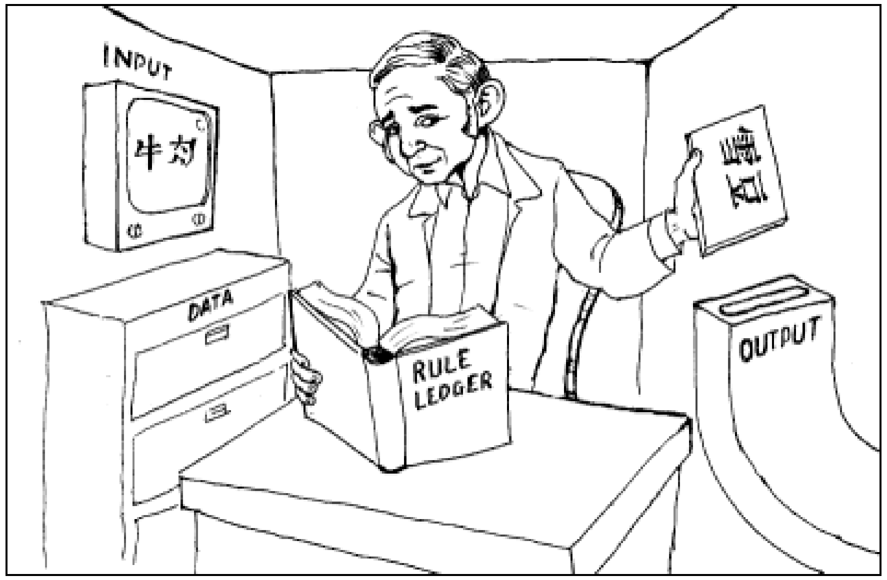

[Slides: Jacob Eisenstein]

Can your computer ever *really* understand you?

What does it really mean to understand language anyway?

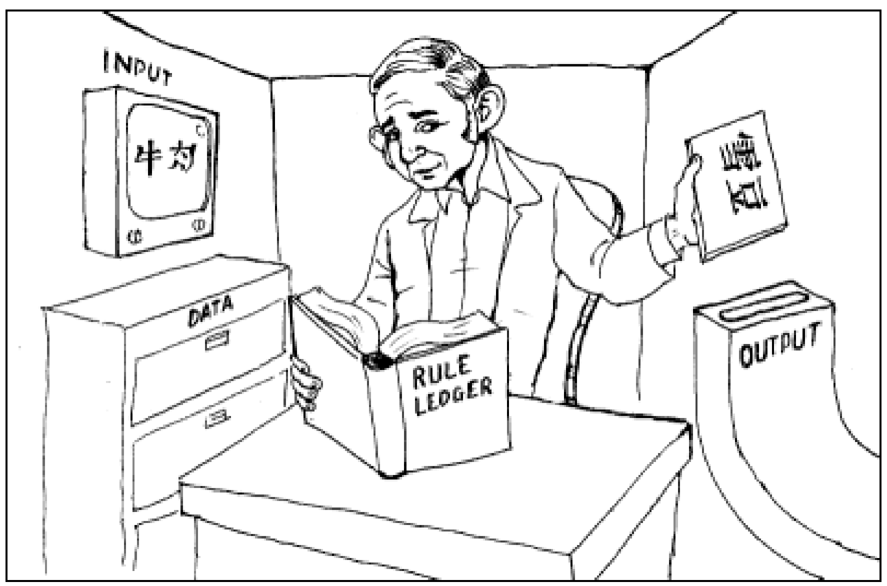

Some functional answers:

- Answer reading comprehension tests
- Determine whether a statement is true or false
- Choose the appropriate action
- Convert text to a *meaning representation*

[Slides: Jacob Eisenstein]

Language to Meaning

Information Extraction

Recover information about pre-specified relations and entities

More informative

Example Task

Relation Extraction

$$is\_a(OBAMA, PRESIDENT)$$

[Slides: ACL 2013 CCG tutorial]

Language to Meaning

Broad-coverage
Semantics

Focus on specific
phenomena (e.g., verb-
argument matching)

More informative

Example Task

Summarization


The New York Times

OBAMA

RACIAL BARRIER FALLS IN DECISIVE VICTORY

Democrats in Congres
Strengthen Grip

Obama wins
election. Big party
in Chicago.
Romney a bit
down, asks for
some tea.

[Slides: ACL 2013 CCG tutorial]

Language to Meaning

Semantic
Parsing

Recover **complete** meaning representation

More informative

Example Task

Database Query

What states border Texas?

Oklahoma
New Mexico
Arkansas
Louisiana

[Slides: ACL 2013 CCG tutorial]

Language to Meaning

Semantic Parsing

Recover **complete** meaning representation

More informative

Example Task

Instructing a Robot

at the chair,
turn right

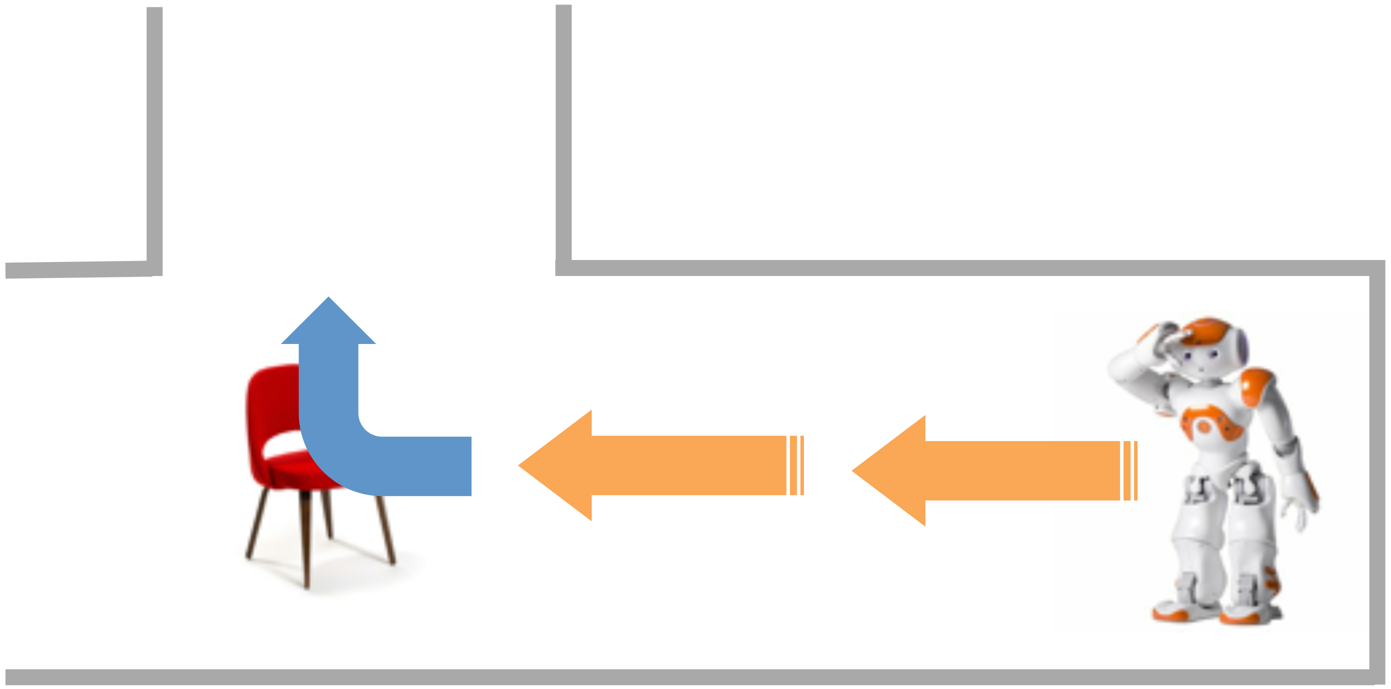

[Slides: ACL 2013 CCG tutorial]

Meaning

- Lexical semantics: individual words/phrases
  - KBs, embeddings, etc.
- Logical semantics
  - [e.g. questions as database queries ... theorem proving ...]
- Compositional semantics
- "Shallow" semantics: predicates, arguments
  - who did what to whom?
  - I bought a car from him <=> he sold me a car
- Practical examples: Information Extraction
- Major subtasks
  - Entities and coreference
    - I saw Bob, and **he** said hi
  - Time and Events

Desiderata for an MR

- Truth-conditional semantics
  - Every sentence is a logical statement (boolean, first order...)
  - Model-theoretic denotations: possible worlds (database states?) licensed by the sentence
  - Entailment and equivalence
- Non-ambiguity
- Expressiveness
- Maps to applications

Semantic parsing

- Semantic parsing: from NL to an MR
  - Typically "sem parse" applies to sentence-only analysis
- Lambda calculus: one common approach
  - Tie it to syntax: e.g. CFG extension (Montague-style semantics)
  - Current research: combinatory categorial grammar (CCG)

$$\mathrm{S} : \beta(\alpha) \rightarrow \mathrm{NP} : \alpha \quad \mathrm{VP} : \beta$$

$$\mathrm{VP} : \beta(\alpha) \rightarrow \mathrm{V} : \beta \quad \mathrm{NP} : \alpha$$

$$\begin{aligned} \textit{Abigail}, \mathrm{NP} &: \textsc{Abigail} \\ \textit{Max}, \mathrm{NP} &: \textsc{Max} \\ \textit{likes}, \mathrm{V} &: \lambda y.\lambda x.\textsc{like}(x, y) \end{aligned}$$

$$\begin{aligned} P &= \lambda y.\lambda x.\textsc{likes}(x, y)(\textsc{Max})(\textsc{Abigail}) \\ &= \lambda x.\textsc{likes}(x, \textsc{Abigail})(\textsc{Max}) \\ &= \textsc{likes}(\textsc{Max}, \textsc{Abigail}) \end{aligned}$$

Natural Language Understanding

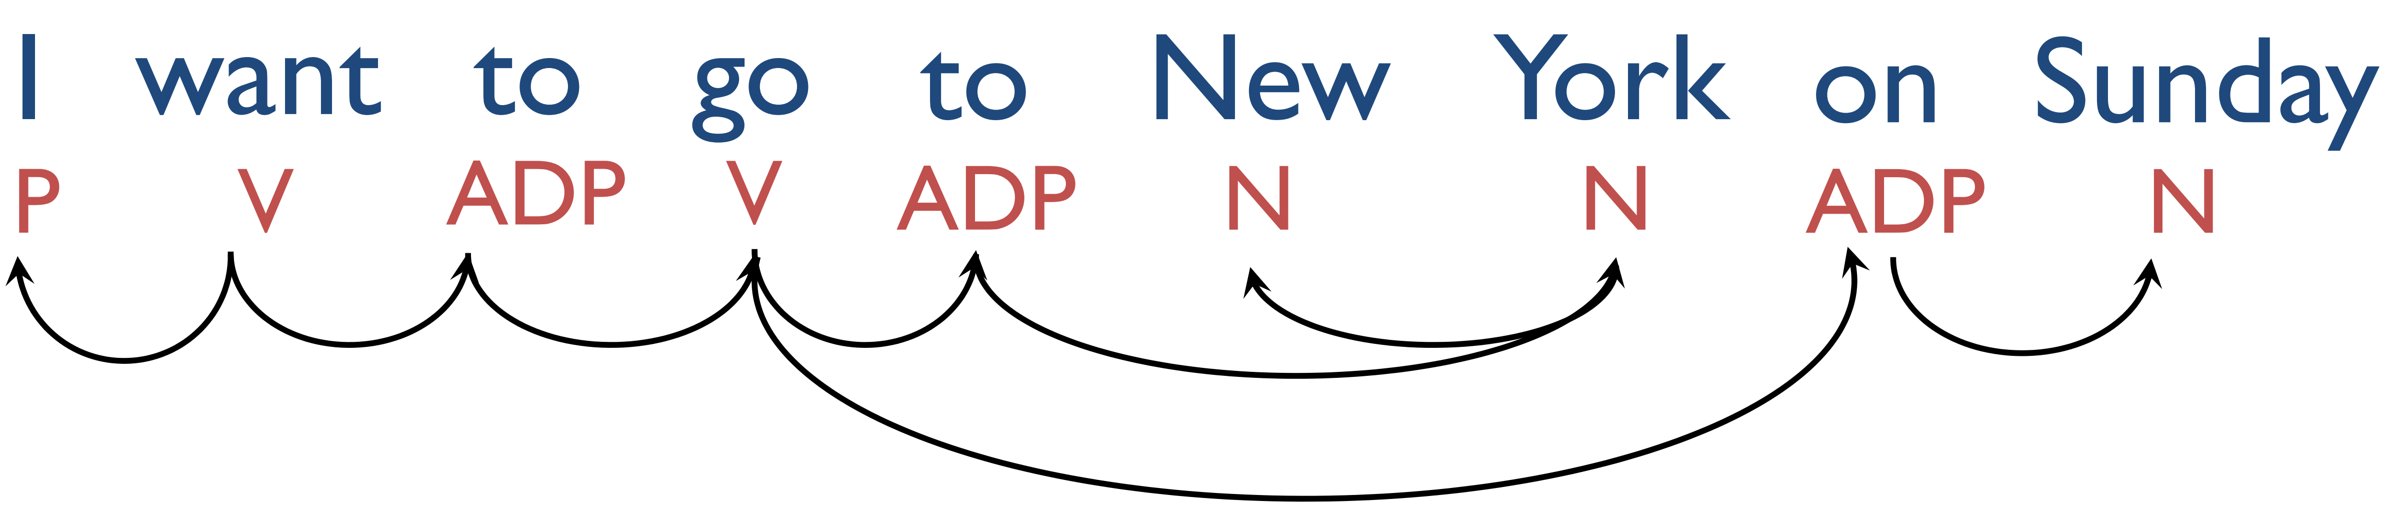

[Slides: Dipanjan Das]

Natural Language Understanding

Shallow Semantics: ***Frames*** *and* ***Roles***

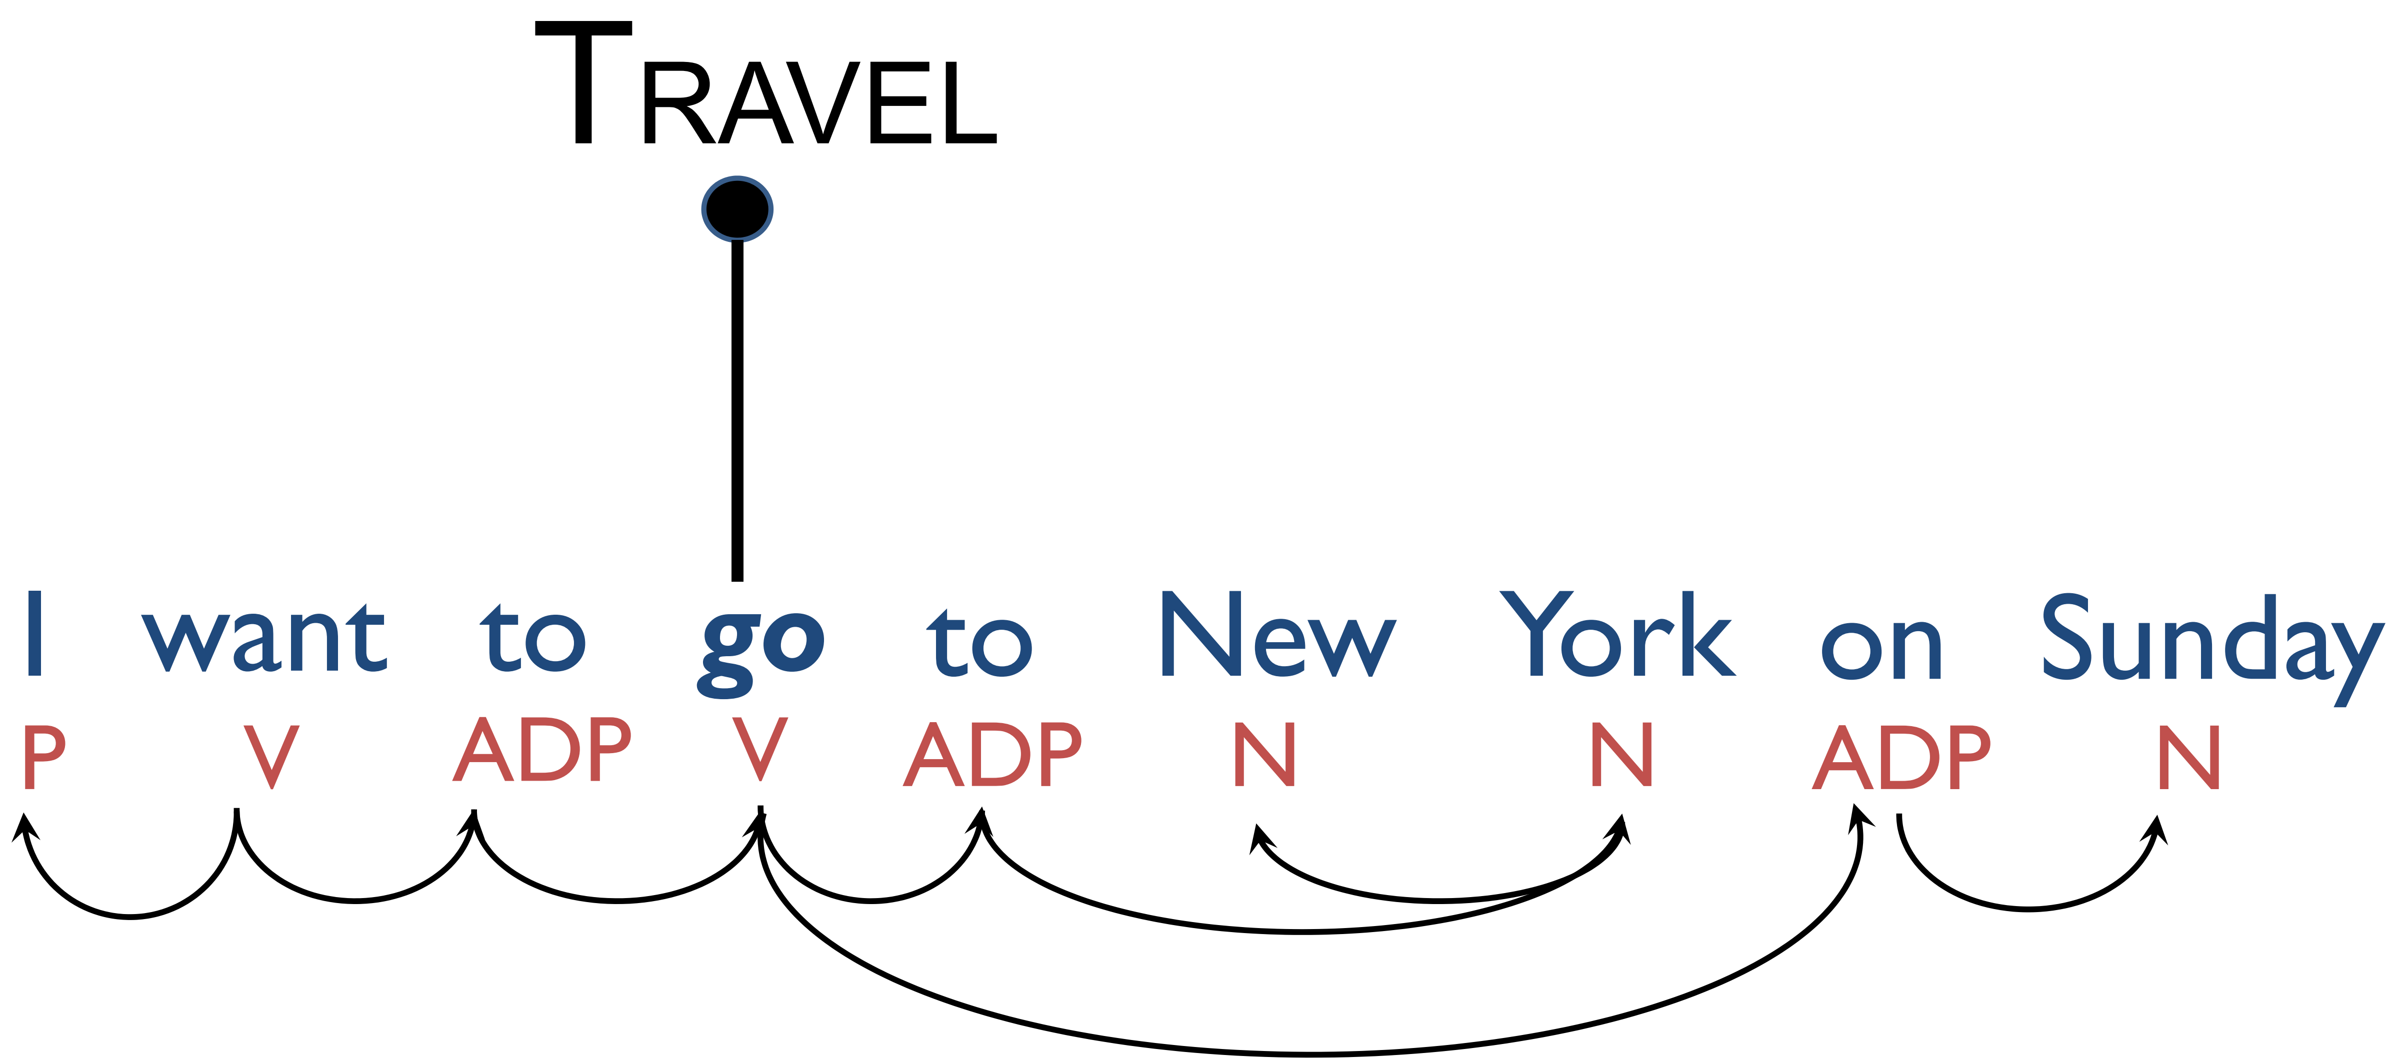

[Slides: Dipanjan Das]

Shallow Semantics: ***Frames*** *and* ***Roles***

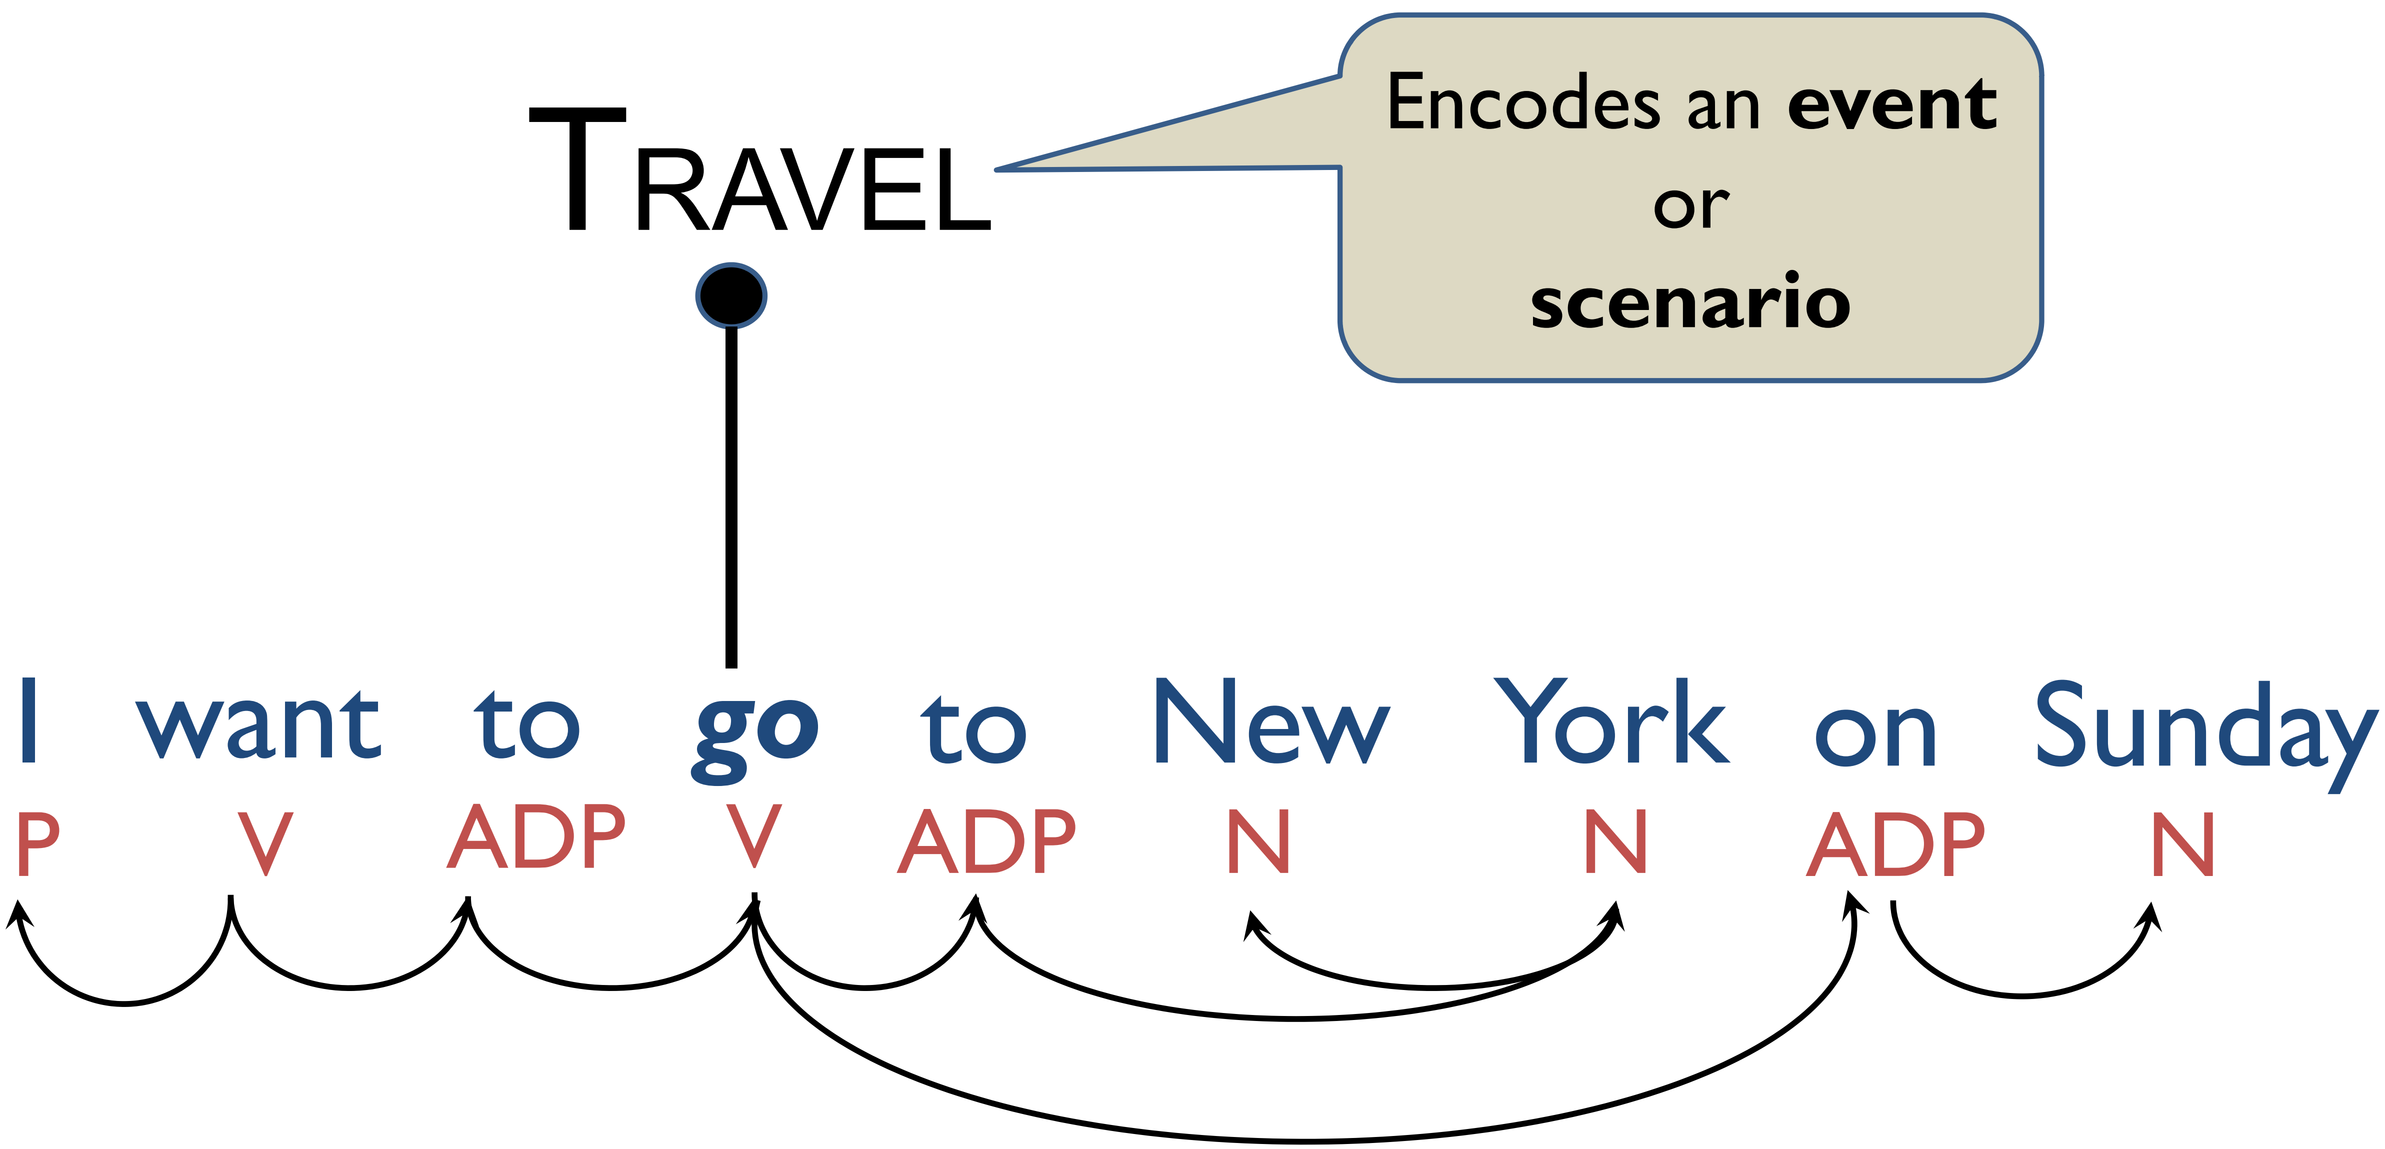

[Slides: Dipanjan Das]

Natural Language Understanding

Shallow Semantics: ***Frames** and **Roles***

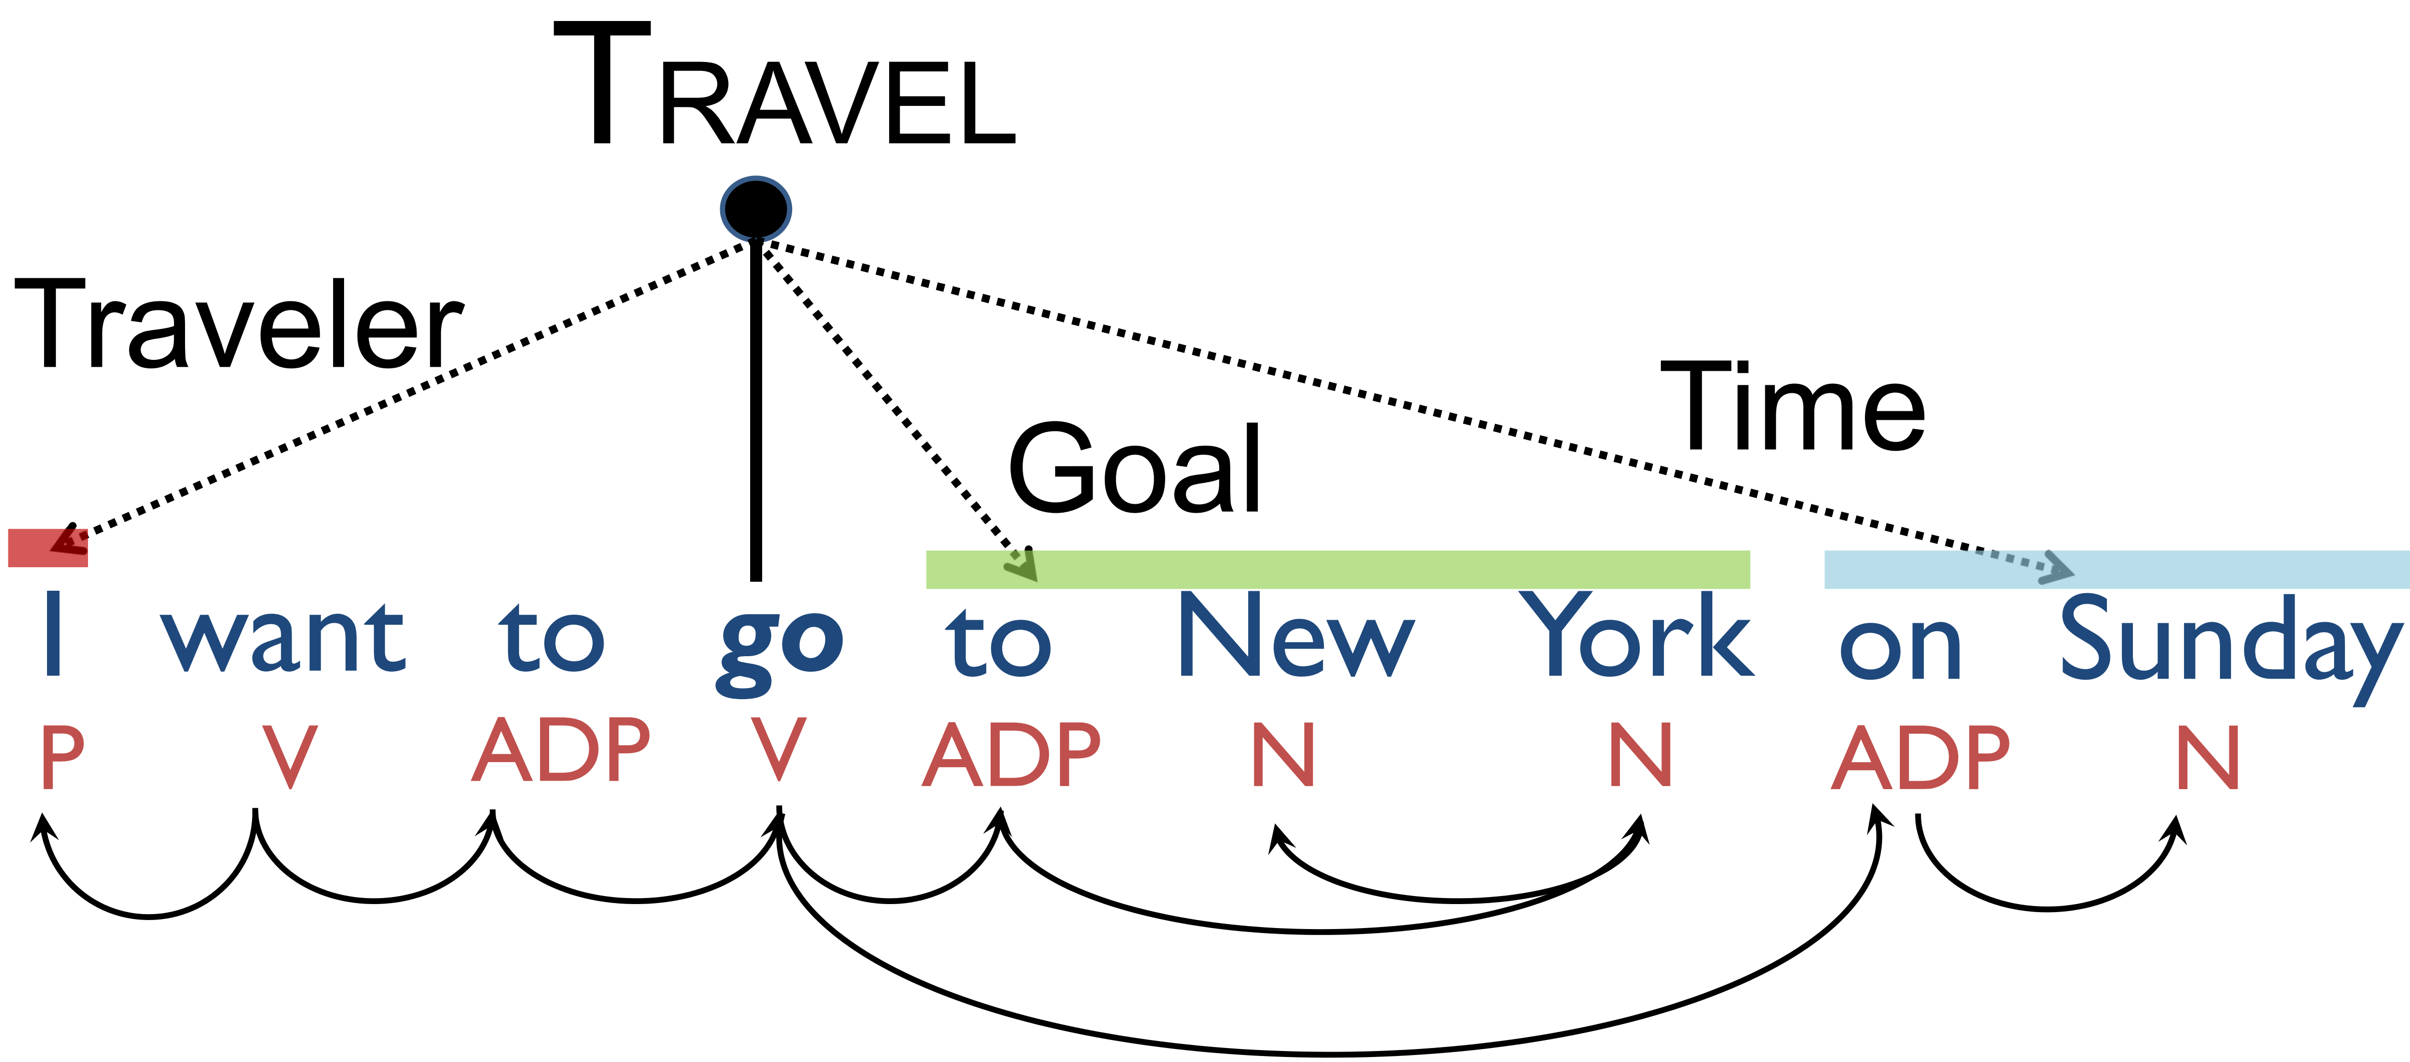

[Slides: Dipanjan Das]

Natural Language Understanding

Shallow Semantics: ***Frames** and **Roles***

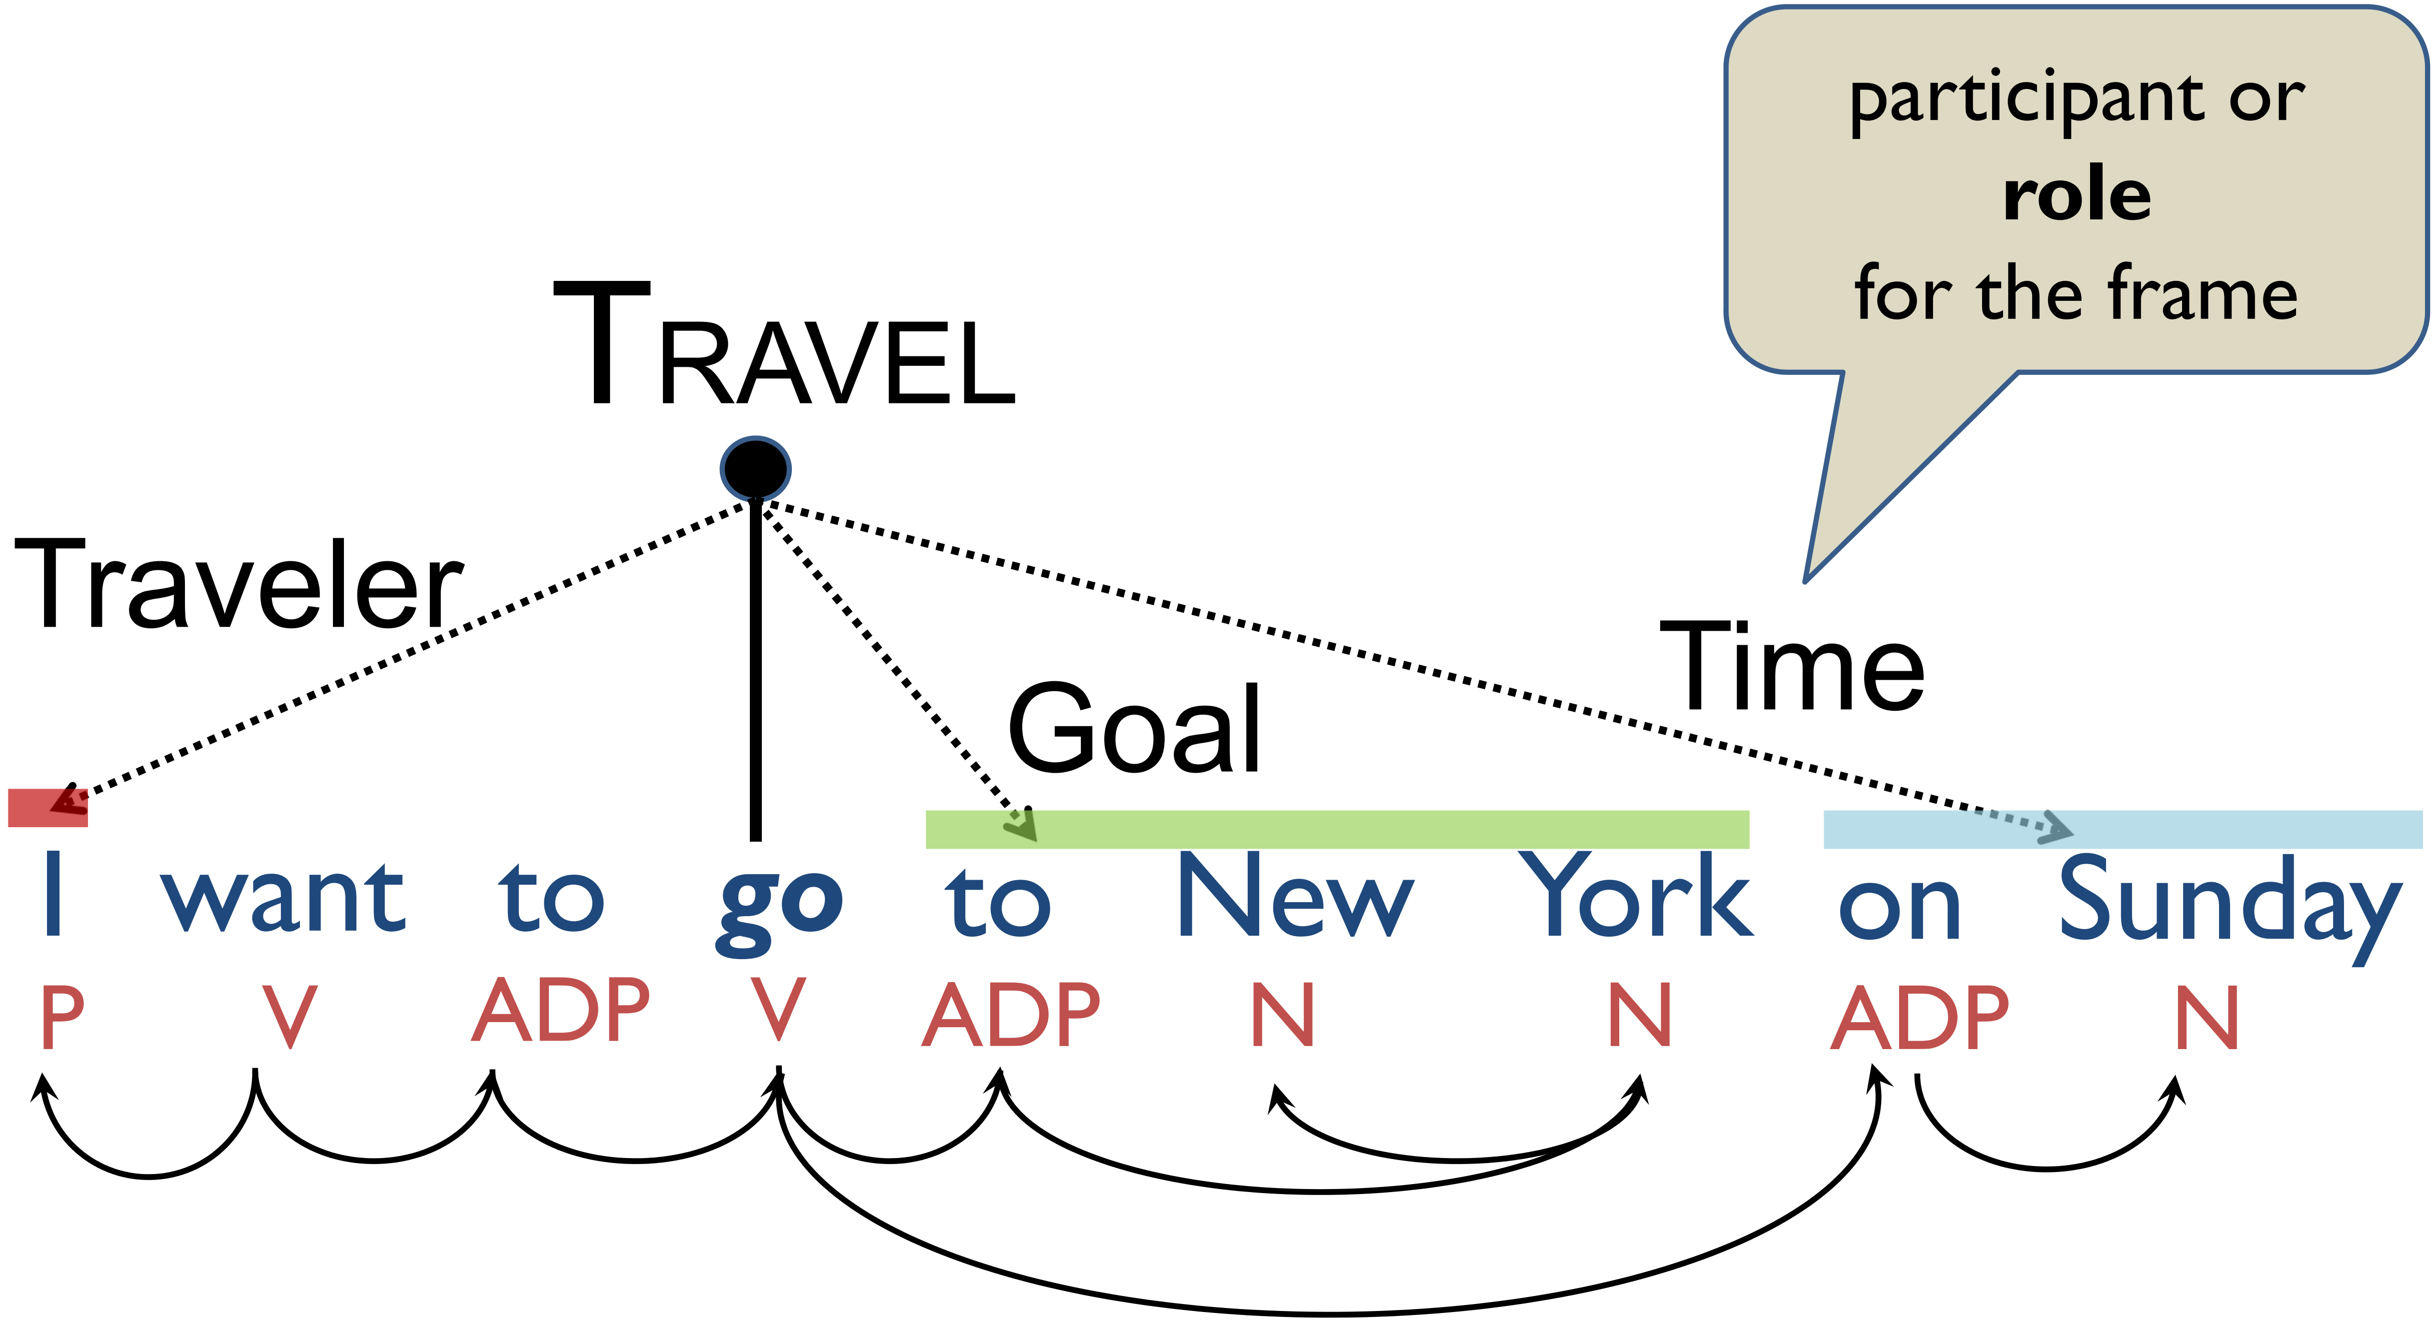

[Slides: Dipanjan Das]

Natural Language Understanding

Shallow Semantics: ***Frames** and **Roles***

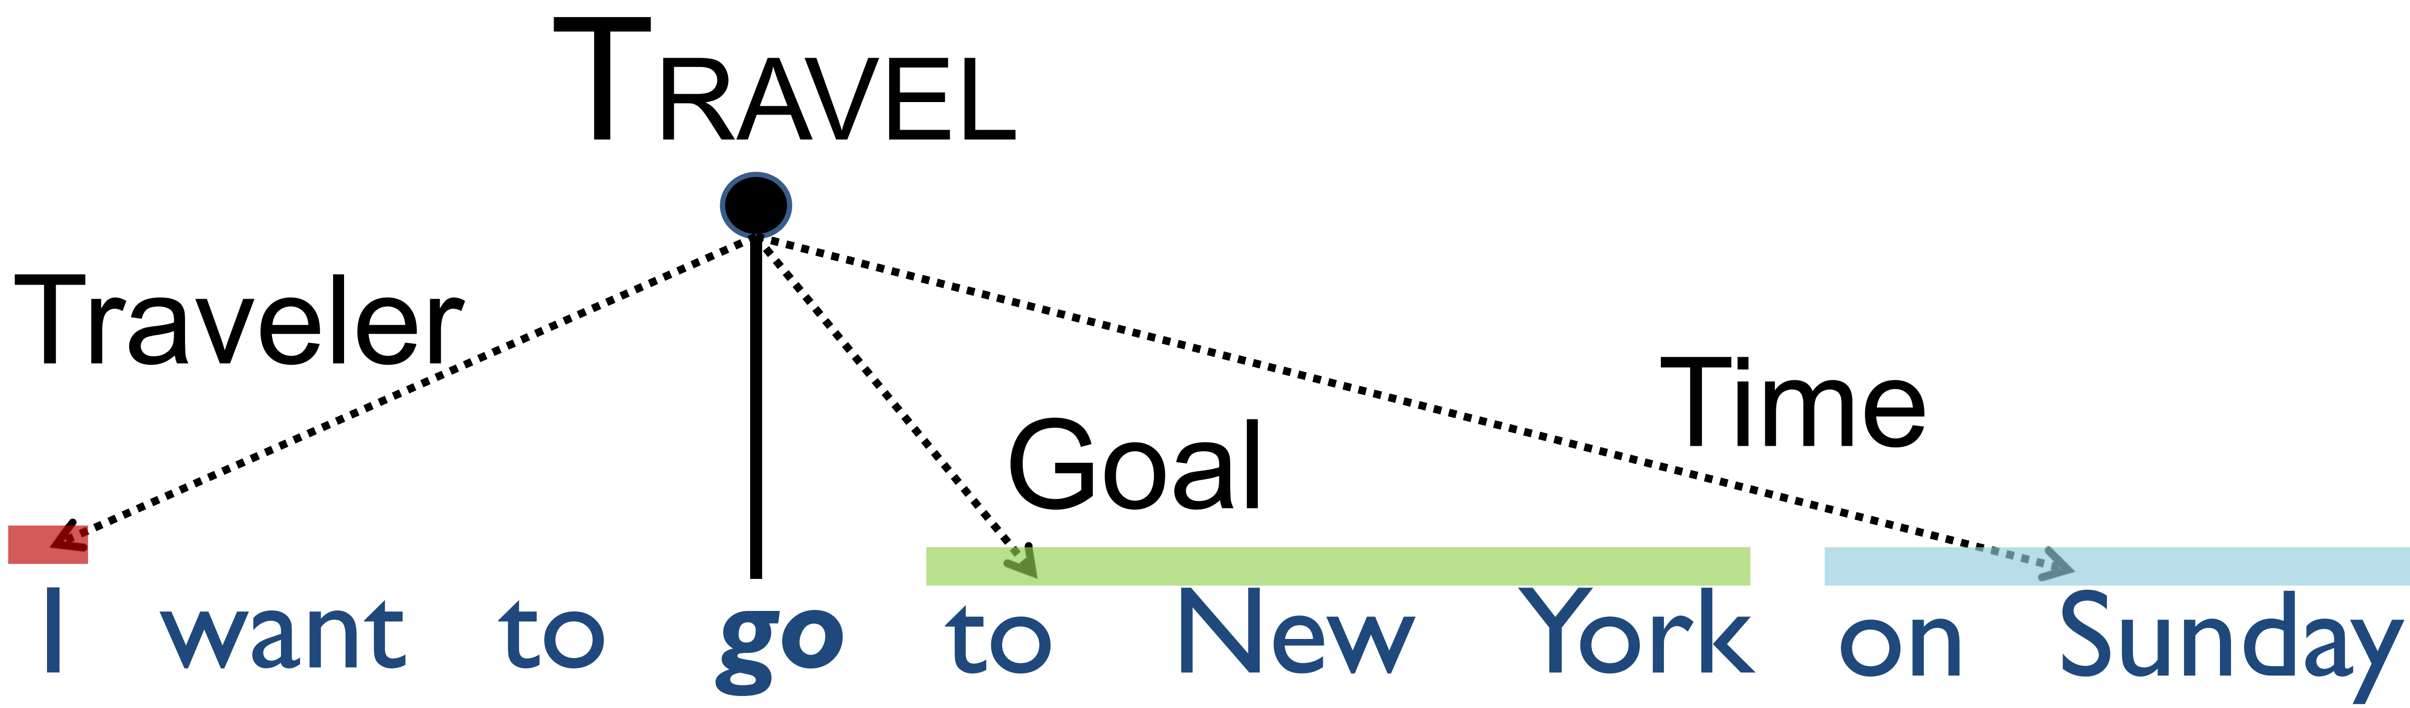

[Slides: Dipanjan Das]

Natural Language Understanding

Shallow Semantics: ***Frames*** *and* ***Roles***

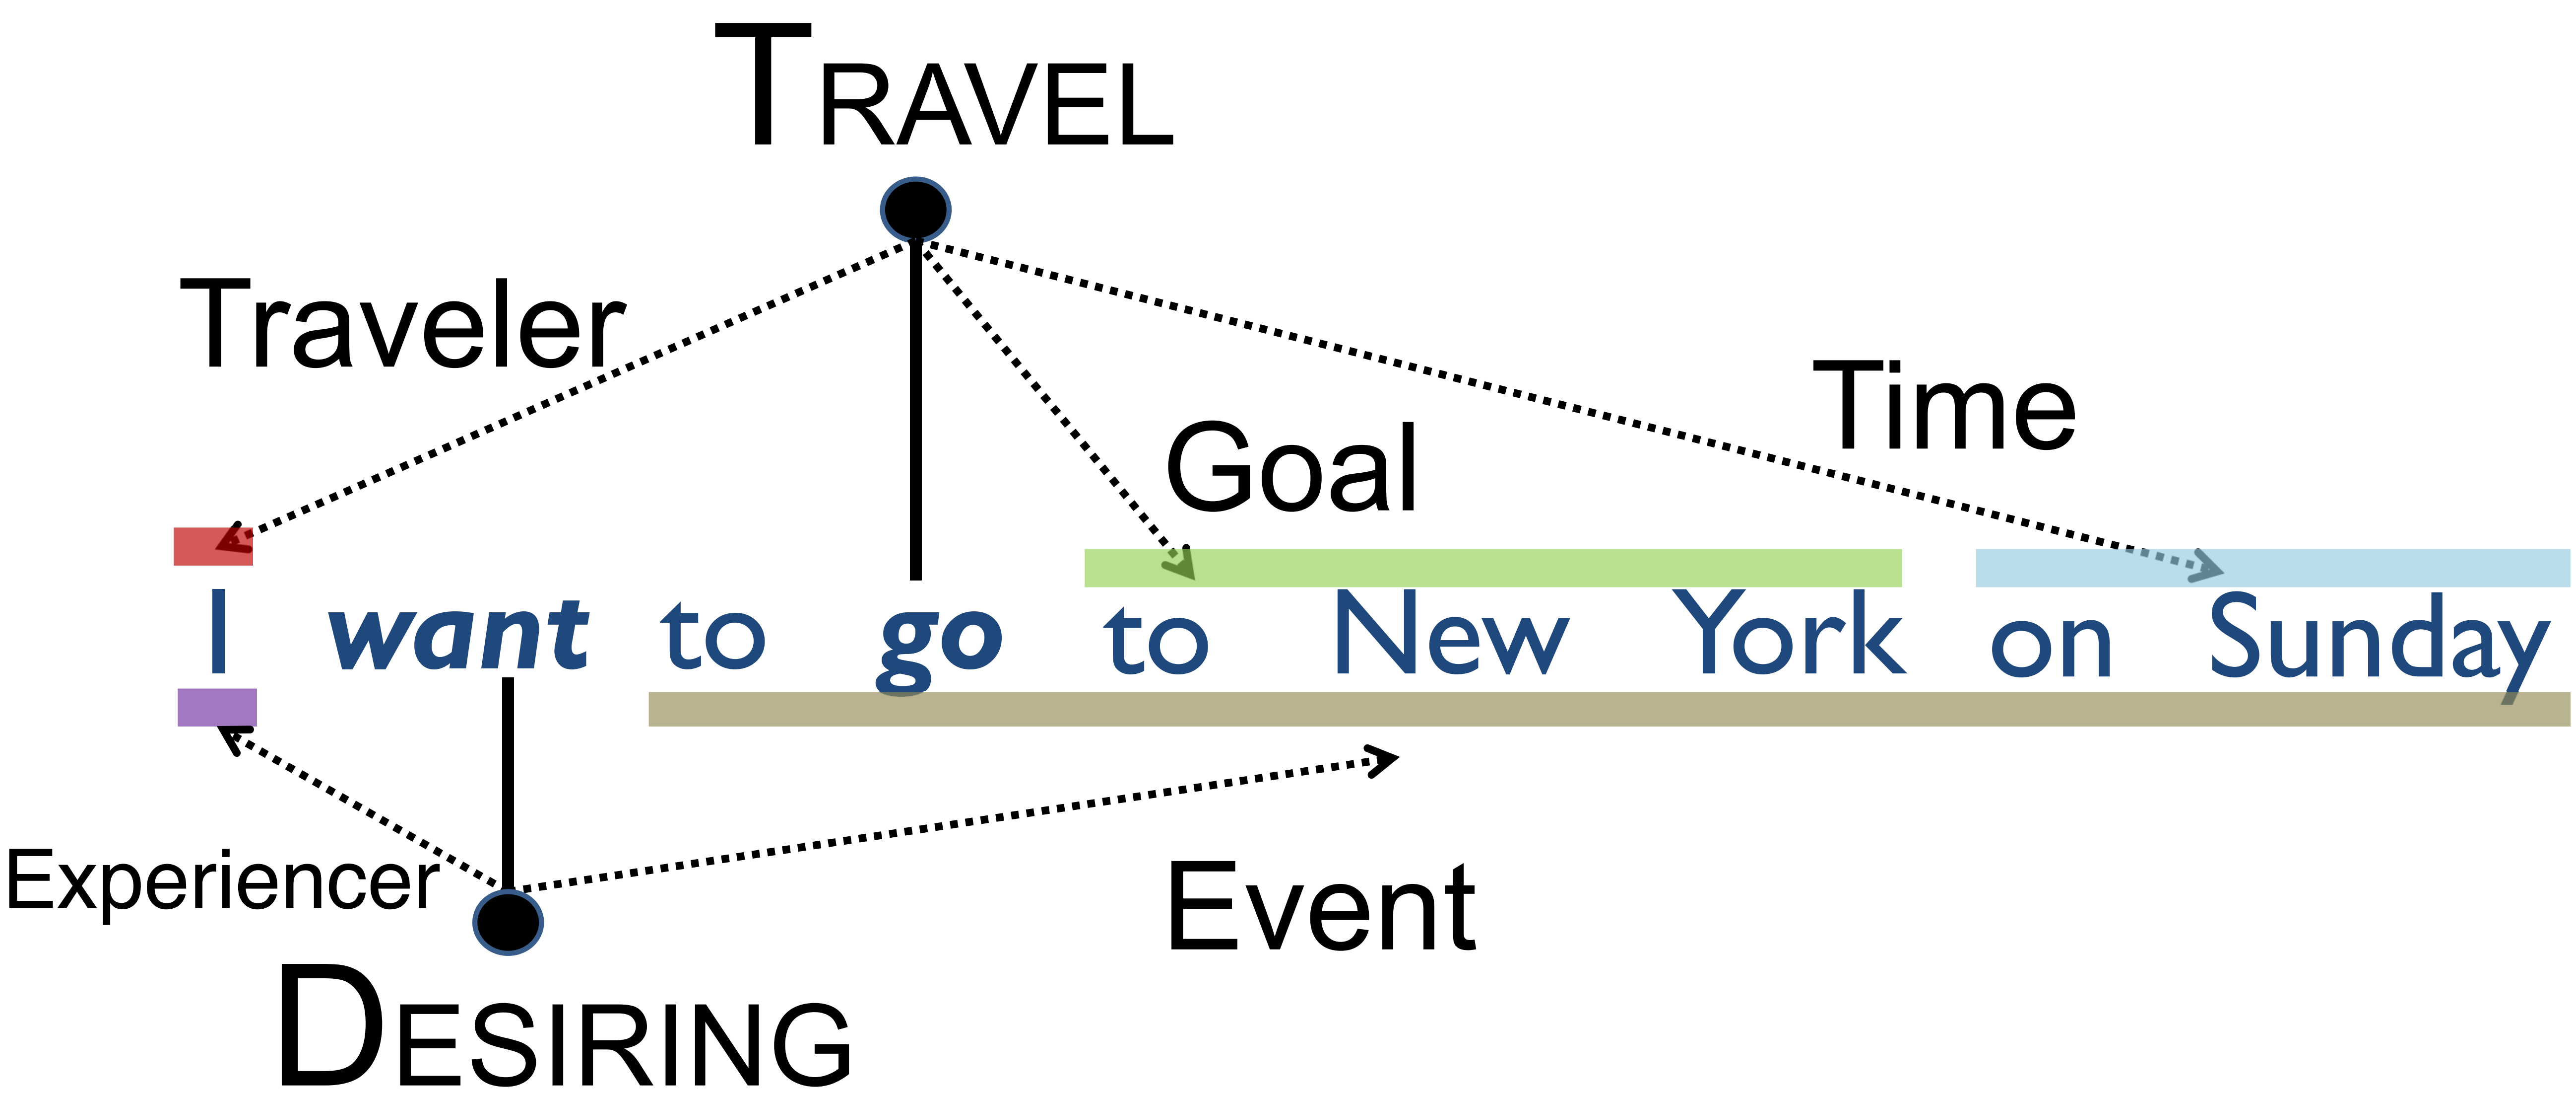

[Slides: Dipanjan Das]

The entity-relation paradigm

http://thelousylinguist.blogspot.com/2017/03/using-ibm-watson-knowledge-studio-to.html

The entity-relation paradigm

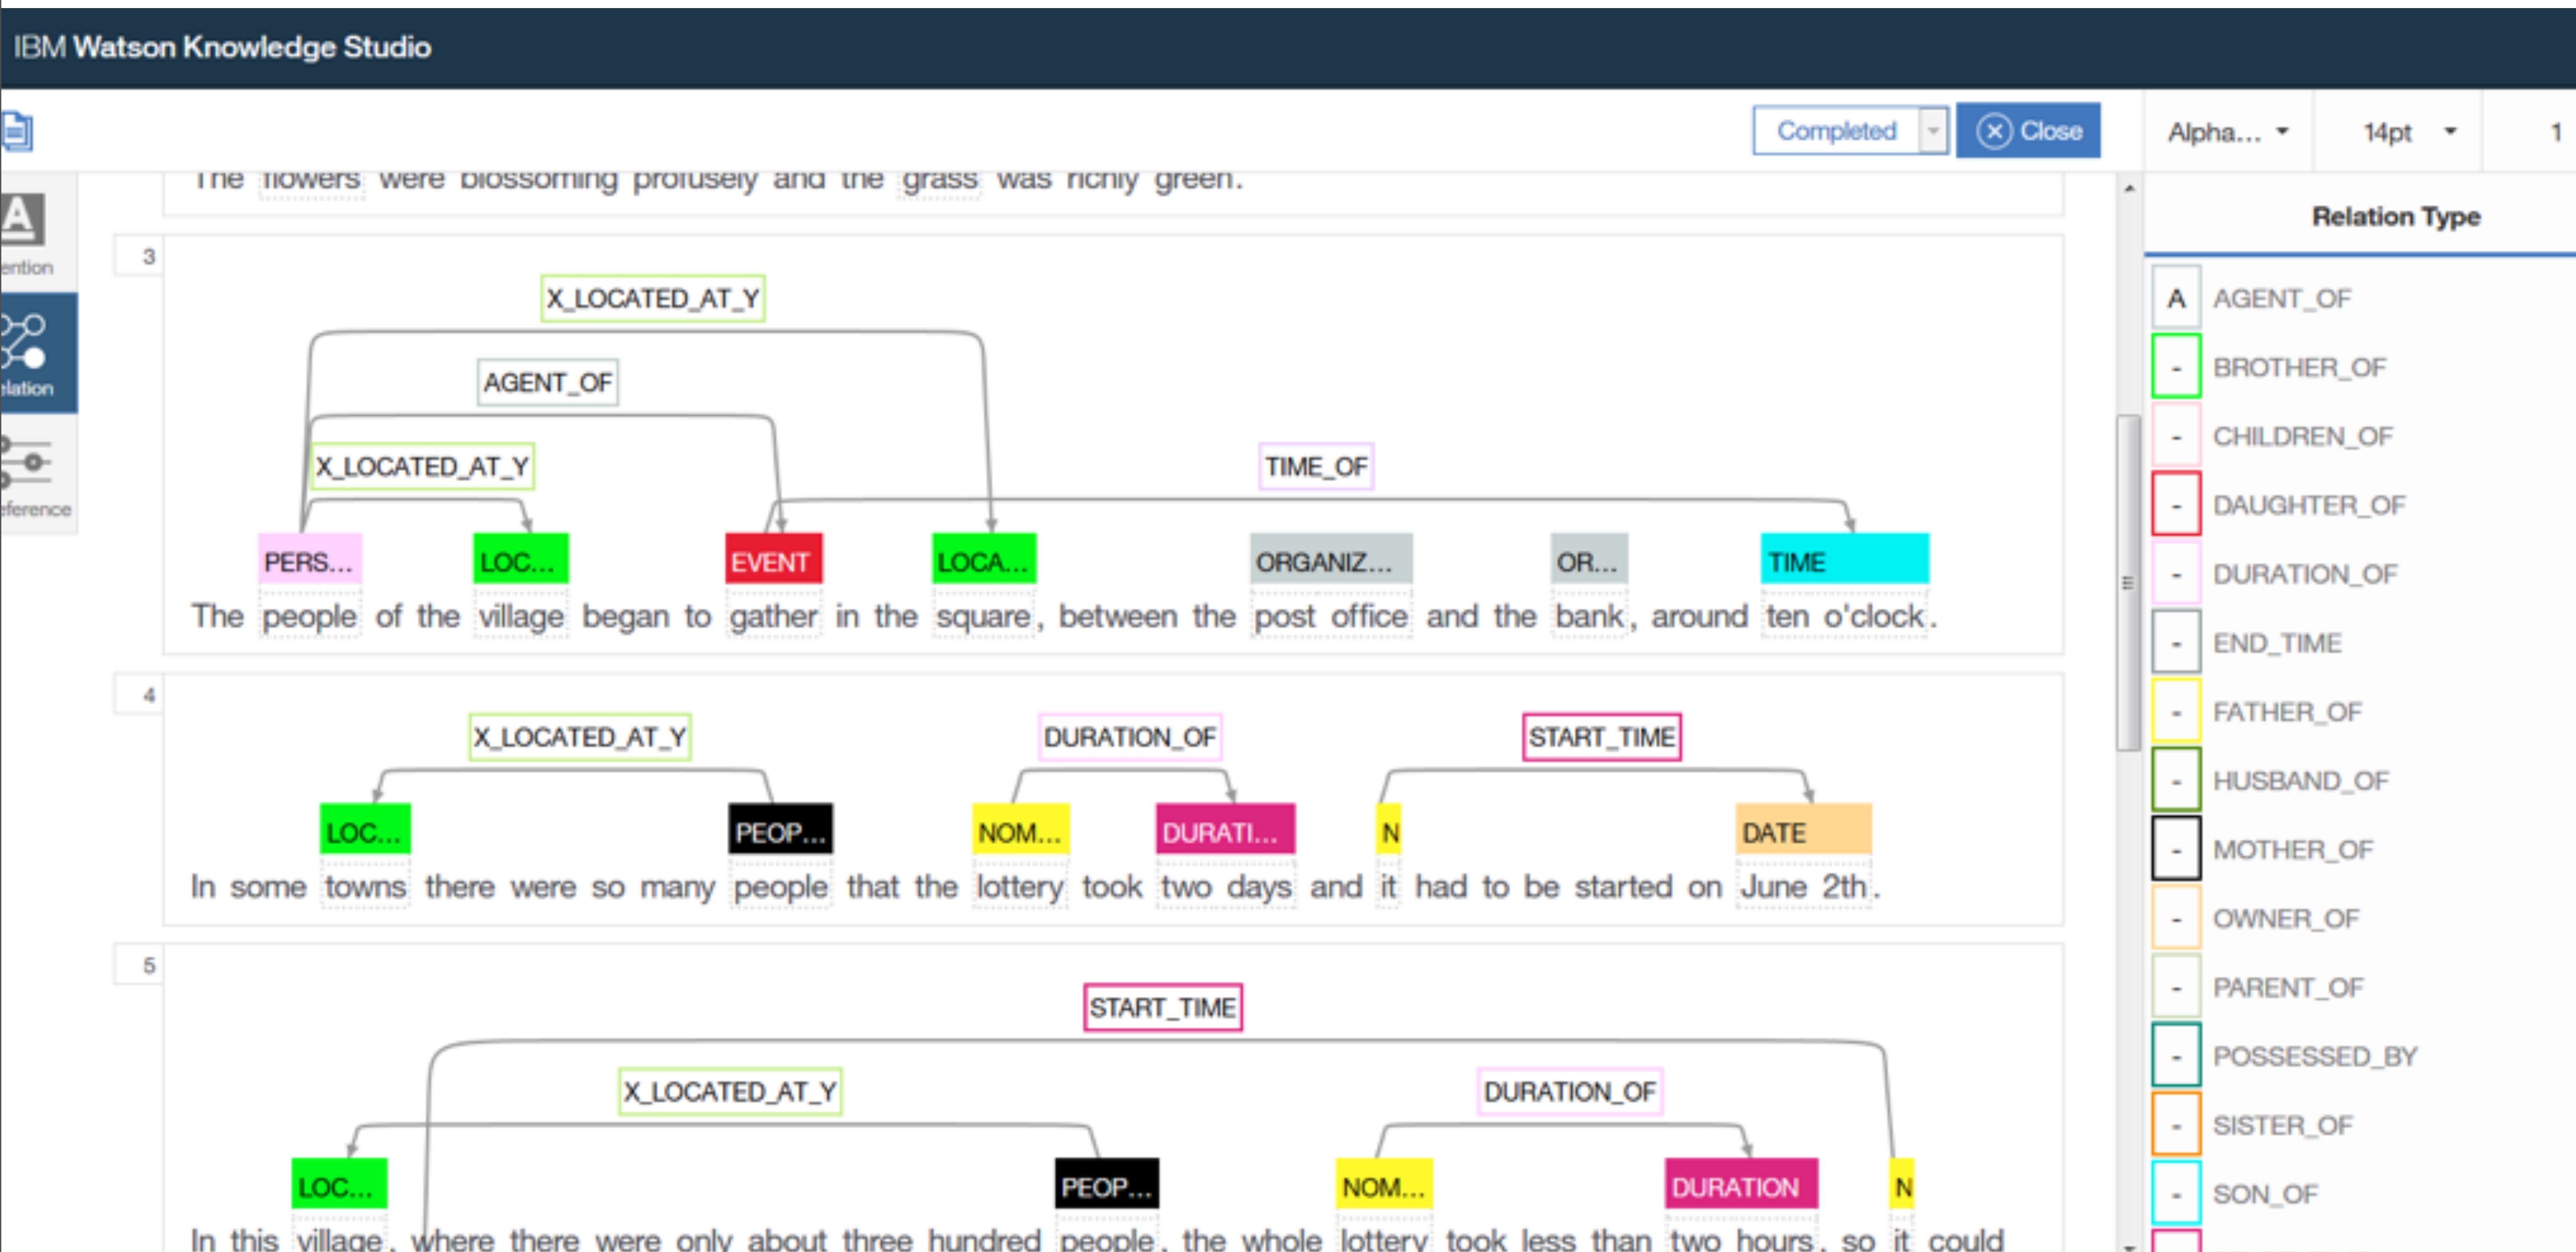

Abstract Meaning Representation

(Broad-coverage MR)

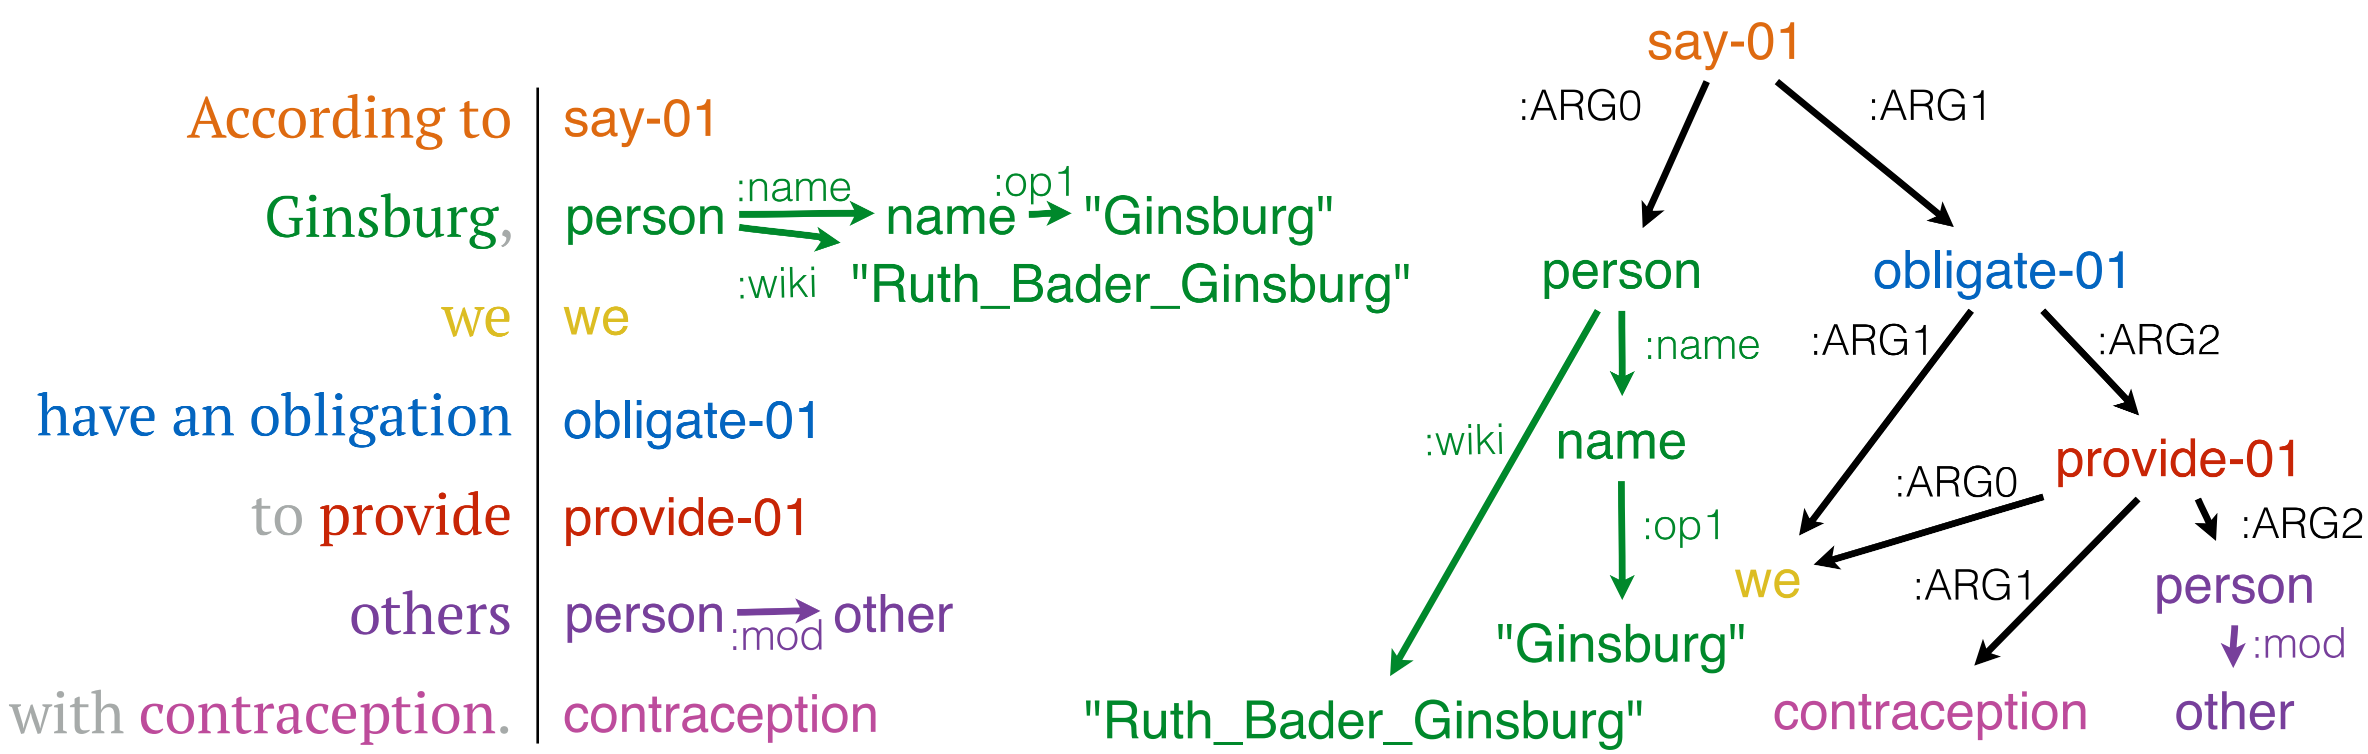

- Annotations reflect a *neo-Davidsonian* logical representation
  - Broad event/predicate classes (“frames”)
  - No deeper sharing of frames across lexical items (buy vs. sell)

https://github.com/nschneid/amr-tutorial/tree/master/slides

Event analysis in intl. relations

(Narrow-coverage MR)

- Analyze time-series of friendly vs. hostile country-country interactions, coded from newswire text
  - Manual coding (~1960's): hire people to read thousands of articles (inconsistencies!)
  - Machine coding (KEDS) -- rule-based S-V-O or S-V-PP extraction *[Phil Schrodt (1993, 1994... 2011)]*
- Various current efforts: ICEWS, OEDA, etc.

Event analysis in intl. relations

(Narrow-coverage MR)

EXAMPLES OF WEIS EVENT CODES

11. REJECT

111 Turn down proposal; reject protest demand; threat
112 Refuse; oppose; refuse to allow

12. ACCUSE

121 Charge, criticize, blame, disapprove
122 Denounce, denigrate, abuse

13. PROTEST

131 Make complaint (not formal)
132 Make formal complaint or protest

17. THREATEN

171 Threat without specific negative sanctions
172 Threat with specific nonmilitary negative sanctions
173 Threat with force specified
174 Ultimatum: threat with negative sanctions and time

18. DEMONSTRATE

181 Non-military demonstration; walk out on
182 Armed force mobilization, exercise and/or display

Table 2
WEIS Coding of 1990 Iraq-Kuwait Crisis

Date	Source	Target	WEIS Code	Type of Action
900717	IRQ	KUW	121	CHARGE
900717	IRQ	UAE	121	CHARGE
900723	IRQ	KUW	122	DENOUNCE
900724	IRQ	ARB	150	DEMAND
900724	IRQ	OPC	150	DEMAND
900725	IRQ	EGY	054	ASSURE
900727	IRQ	KUW	160	WARN
900731	IRQ	KUW	182	MOBILIZATION
900801	KUW	IRQ	112	REFUSE
900802	IRQ	KUW	223	MILITARY FORCE

Event analysis in intl. relations

(Narrow-coverage MR)

Figure 1
USA Actions Towards USSR, 1948-1978

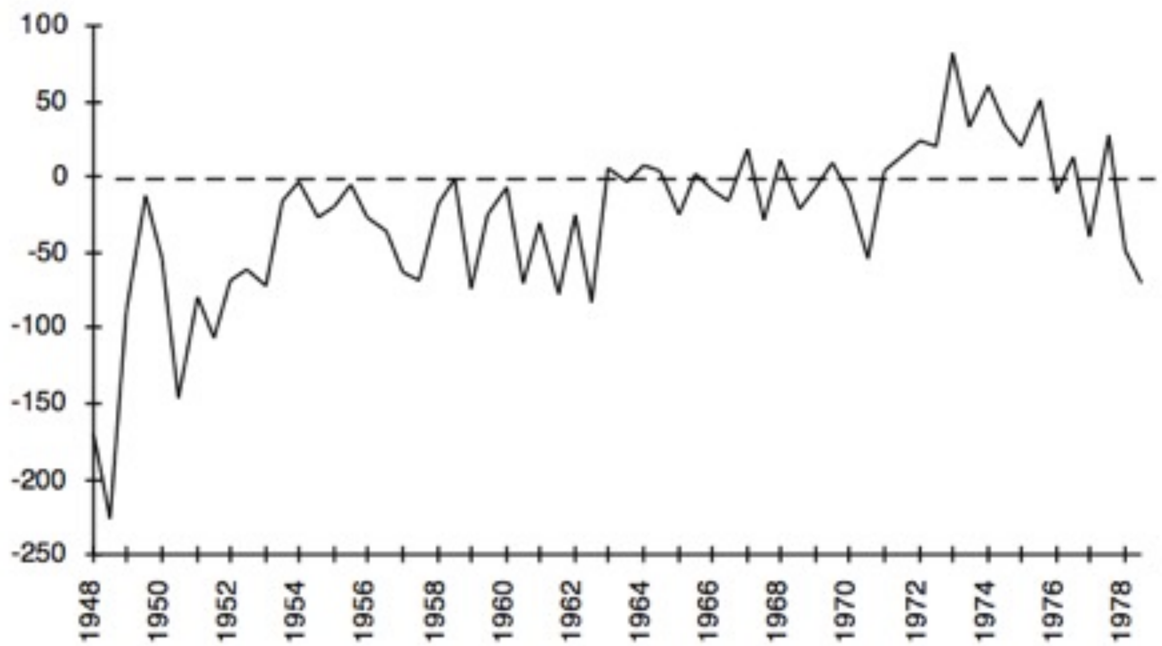

Figure 2
Israel-Palestinian interactions, 1982-1992

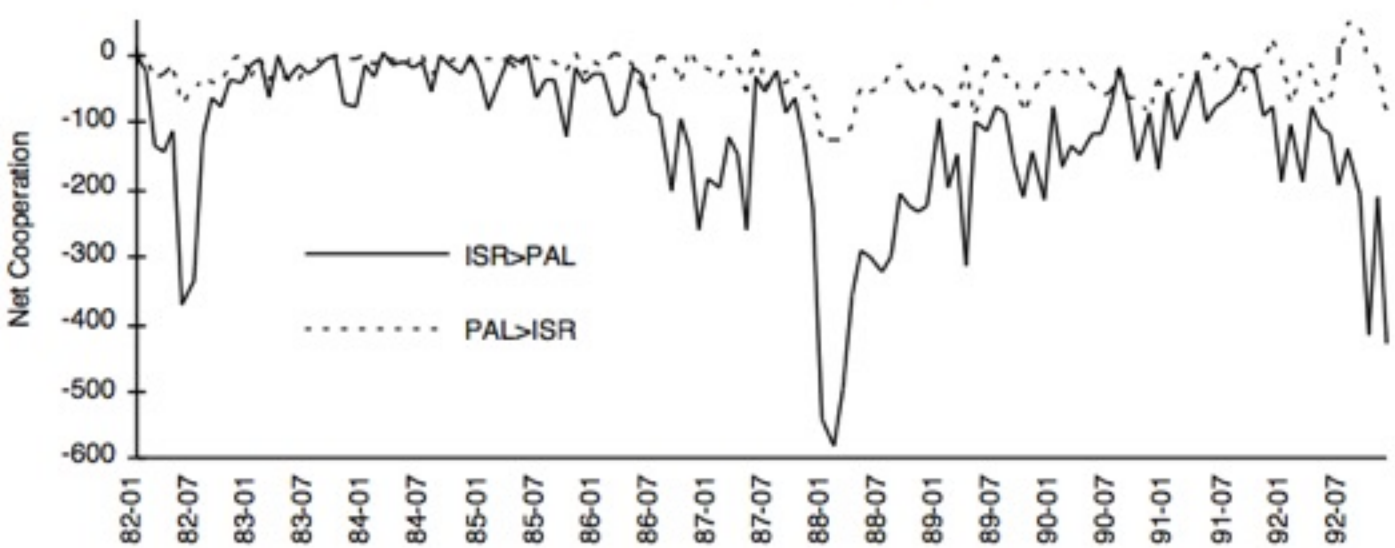

Crosscorrelation of Arms Transfers and International Cooperation from Recepient to Supplier

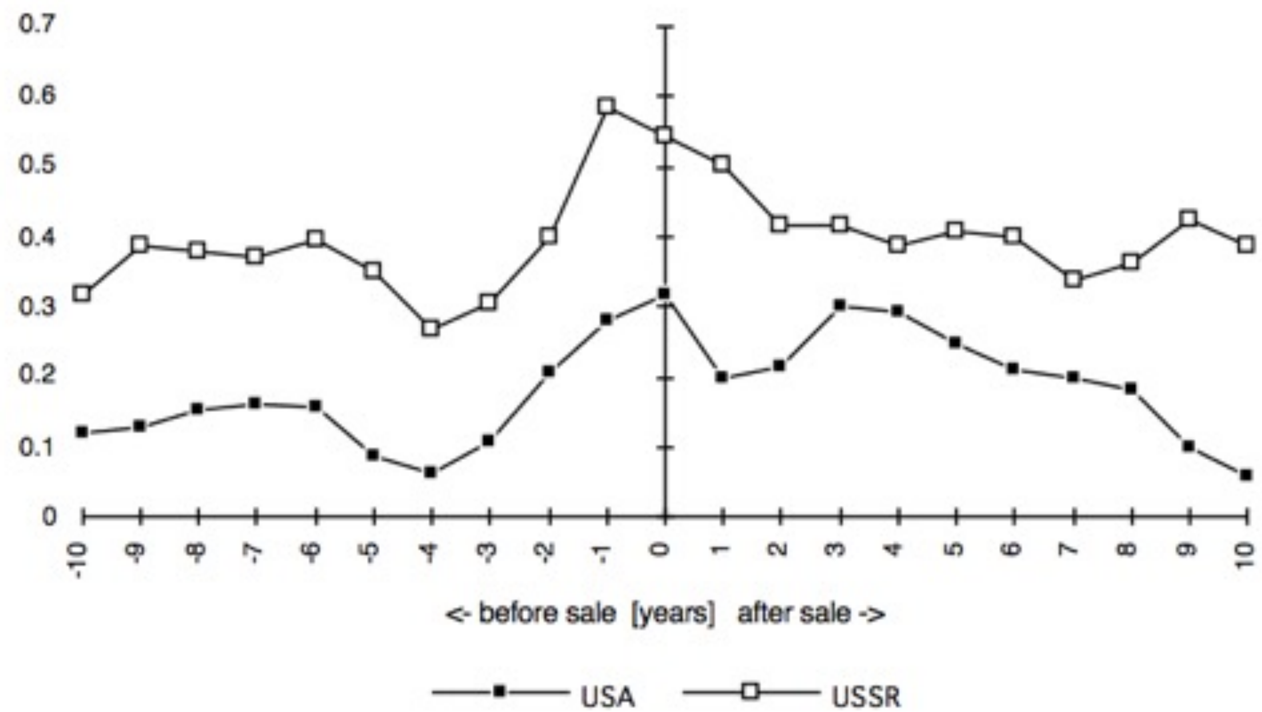

(These graphs are from manual coding; IE evaluations in Schrodt and Gerner 1994, King and Lowe 2001)

Message Understanding Conferences (MUC)

- Bakeoff format: shared task, dataset, hidden test set for competitive evaluation
- Different domains – involving specific events
  - (1987) MUC-1: Fleet operations
  - (1991-2) MUC-3, 4: Terrorist activities in Latin America
  - (1993-7) Corporate Joint Ventures, Microelectronic production, Negotiation of Labor Disputes, Airplane crashes, and Rocket/Missile Launches
- ACE (1999-2008) – Automated Content Extraction

MUC Template-Filling IE

<u>Input</u>: text

San Salvador, 19 Apr 89 (ACAN-EFE) – [TEXT] Salvadoran President-elect Alfredo Cristiani condemned the terrorist killing of Attorney General Roberto Garcia Alvarado and accused the Farabundo Marti National Liberation Front (FMLN) of the crime.

...

Garcia Alvarado, 56, was killed when a bomb placed by urban guerrillas on his vehicle exploded as it came to a halt at an intersection in downtown San Salvador.

...

Vice President-elect Francisco Merino said that when the attorney general's car stopped at a light on a street in downtown San Salvador, an individual placed a bomb on the roof of the armored vehicle.

...

According to the police and Garcia Alvarado's driver, who escaped unscathed, the attorney general was traveling with two bodyguards. One of them was injured.

<u>Output</u>: extract an event record ("terrorist attack") with the following attributes:

Incident: Date
Incident: Location
Incident: Type
Perpetrator: Individual ID
Perpetrator: Organization ID
Perpetrator: Organization Confidence
Physical Target: Description
Physical Target: Effect
Human Target: Name
Human Target: Description

MUC Template-Filling IE

Input: text

San Salvador, 19 Apr 89 (ACAN-EFE) – [TEXT] Salvadoran President-elect Alfredo Cristiani condemned the terrorist killing of Attorney General Roberto Garcia Alvarado and accused the Farabundo Marti National Liberation Front (FMLN) of the crime.

...

Garcia Alvarado, 56, was killed when a bomb placed by urban guerrillas on his vehicle exploded as it came to a halt at an intersection in downtown San Salvador.

...

Vice President-elect Francisco Merino said that when the attorney general's car stopped at a light on a street in downtown San Salvador, an individual placed a bomb on the roof of the armored vehicle.

...

According to the police and Garcia Alvarado's driver, who escaped unscathed, the attorney general was traveling with two bodyguards. One of them was injured.

Output: extract an event record ("terrorist attack") with the following attributes:

Incident: Date	- 19 Apr 89
Incident: Location	El Salvador: San Salvador (city)
Incident: Type	Bombing
Perpetrator: Individual ID	"urban guerrillas"
Perpetrator: Organization ID	"FMLN"
Perpetrator: Organization Confidence	Suspected or Accused by Authorities: "FMLN"
Physical Target: Description	"vehicle"
Physical Target: Effect	Some Damage: "vehicle"
Human Target: Name	"Roberto Garcia Alvarado"
Human Target: Description	"attorney general": "Roberto Garcia Alvarado" "driver" "bodyguards"

Pipeline
(finite-state transducers)

Text

1. Complex Words
2. Basic Phrases
3. Complex Phrases
4. Domain Events
5. Merging Structures

Structure

<u>Syntax steps</u>

Names, multiwords...

NPs, verb groups, phrase structure...

<u>Domain-specific semantics</u>

Event Patterns

<Company/ies> <Set-up> <Joint-Venture> with <Company/ies>

↓

Relationship:	TIE-UP
Entities:	"Bridgestone Sports Co."
	"a local concern"
	"a Japanese trading house"
Joint Venture Company:	–
Activity:	–
Amount:	–

<Produce> <Product>

↓

Activity:	PRODUCTION
Company:	–
Product:	"golf clubs"
Start Date:	–

Bridgestone Sports Co. said Friday it has set up a joint venture in Taiwan with a local concern and a Japanese trading house to produce golf clubs to be shipped to Japan.

The joint venture, Bridgestone Sports Taiwan Co., capitalized at 20 million new Taiwan dollars, will start production in January 1990 with production of 20,000 iron and "metal wood" clubs a month.

(4/5) Domain Events

(5/5) Merge Structures

Activity:	PRODUCTION
Company:	–
Product:	“golf clubs”
Start Date:	–

Activity:	PRODUCTION
Company:	“Bridgestone Sports Taiwan Co.”
Product:	–
Start Date:	DURING: January 1990

Relationship:	TIE-UP
Entities:	“Bridgestone Sports Co.” “a local concern” “a Japanese trading house”
Joint Venture Company:	–
Activity:	–
Amount:	–

Relationship:	TIE-UP
Entities:	–
Joint Venture Company:	“Bridgestone Sports Taiwan Co.”
Activity:	–
Amount:	NT$20000000

Bridgestone Sports Co. said Friday it has set up a **joint venture** in Taiwan with **a local concern** and **a Japanese trading house** to **produce golf clubs** to be shipped to Japan.

The joint venture, **Bridgestone Sports Taiwan Co.**, **capitalized** at **20 million new Taiwan dollars**, will start **production** in **January 1990** with production of 20,000 iron and “metal wood” clubs a month.

(4/5) Domain Events
(5/5) Merge Structures

Activity:	PRODUCTION
Company:	–
Product:	"golf clubs"
Start Date:	–

Activity:	PRODUCTION
Company:	"Bridgestone Sports Taiwan Co."
Product:	–
Start Date:	DURING: January 1990

Relationship:	TIE-UP
Entities:	"Bridgestone Sports Co." "a local concern" "a Japanese trading house"
Joint Venture Company:	–
Activity:	–
Amount:	–

Relationship:	TIE-UP
Entities:	–
Joint Venture Company:	"Bridgestone Sports Taiwan Co."
Activity:	–
Amount:	NT$20000000

(4/5) Domain Events

(5/5) Merge Structures

Activity:	PRODUCTION
Company:	–
Product:	“golf clubs”
Start Date:	–

Activity:	PRODUCTION
Company:	“Bridgestone Sports Taiwan Co.”
Product:	–
Start Date:	DURING: January 1990

Relationship:	TIE-UP
Entities:	“Bridgestone Sports Co.”
	“a local concern”
	“a Japanese trading house”
Joint Venture Company:	–
Activity:	–
Amount:	–

Relationship:	TIE-UP
Entities:	–
Joint Venture Company:	“Bridgestone Sports Taiwan Co.”
Activity:	–
Amount:	NT$20000000

Decide identity coreference through name-matching and type compatibility; if arguments are coreferent, merge events

Activity:	PRODUCTION
Company:	“Bridgestone Sports Taiwan Co.”
Product:	“iron and ‘metal wood’ clubs”
Start Date:	DURING: January 1990

(4/5) Domain Events
(5/5) Merge Structures

Decide identity coreference through name-matching and type compatibility; if arguments are coreferent, merge events

Activity:	PRODUCTION
Company:	–
Product:	"golf clubs"
Start Date:	–

Activity:	PRODUCTION
Company:	"Bridgestone Sports Taiwan Co."
Product:	–
Start Date:	DURING: January 1990

Activity:	PRODUCTION
Company:	"Bridgestone Sports Taiwan Co."
Product:	"iron and 'metal wood' clubs"
Start Date:	DURING: January 1990

Relationship:	TIE-UP
Entities:	"Bridgestone Sports Co." "a local concern" "a Japanese trading house"
Joint Venture Company:	–
Activity:	–
Amount:	–

Relationship:	TIE-UP
Entities:	–
Joint Venture Company:	"Bridgestone Sports Taiwan Co."
Activity:	–
Amount:	NT$20000000

Relationship:	TIE-UP
Entities:	"Bridgestone Sports Co." "a local concern" "a Japanese trading house"
Joint Venture Company:	"Bridgestone Sports Taiwan Co."
Activity:	–
Amount:	NT$20000000

Empirical Rule-based NLP

- Originally FASTUS was just a preprocessor for a more complex system. It was too slow, they threw it out -- deadline pressure
- Hours vs Minutes runtime on development set -- much faster development iterations

Empirical Rule-based NLP

- Originally FASTUS was just a preprocessor for a more complex system. It was too slow, they threw it out -- deadline pressure
- Hours vs Minutes runtime on development set -- much faster development iterations

January: Designed FASTUS
Jan-May: Development
May 6: First test of the FASTUS system on a blind test set of 100 terrorist reports, which had been withheld as a fair test, and we obtained a score of **8% recall and 42% precision**.
At that point we began a fairly intensive effort to hill-climb on all 1300 development texts then available, doing periodic runs on the fair test to monitor our progress. This effort culminated in a score of **44% recall and 57% precision** in the wee hours of **June 1**, when we decided to run the official test. The rate of progress was rapid enough that even a few hours of work could be shown to have a noticeable impact on the score. Our scarcest resource was time, and our supply of it was eventually exhausted well before the point of diminishing returns.
We were thus able, **in three and a half weeks, to increase the system's F-score by 36.2 points, from 13.5 to 49.7**.

- Current work in supervised event extraction (feature-based, neural network...)
- ACE entity/event dataset: ~dozen event types and mention-level annotations

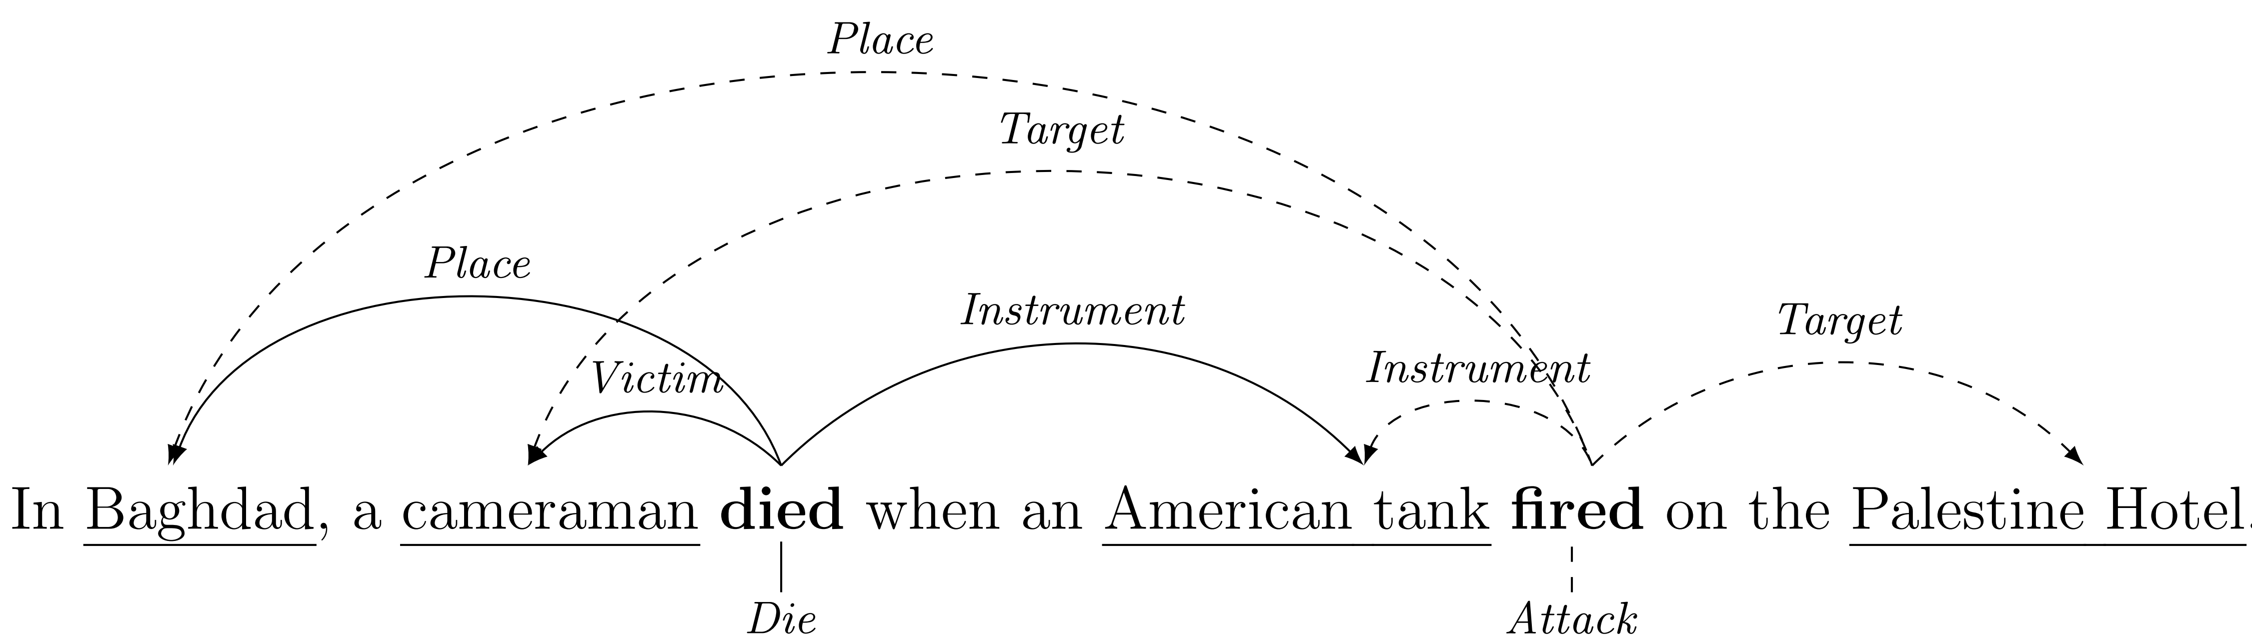

[Li et al. 2013]

Category	Type	Feature Description
Trigger	Lexical	*1.* unigrams/bigrams of the current and context words within the window of size 2 *2.* unigrams/bigrams of part-of-speech tags of the current and context words within the window of size 2 *3.* lemma and synonyms of the current token *4.* base form of the current token extracted from Nomlex (Macleod et al., 1998) *5.* Brown clusters that are learned from ACE English corpus (Brown et al., 1992; Miller et al., 2004; Sun et al., 2011). We used the clusters with prefixes of length 13, 16 and 20 for each token.
	Syntactic	*6.* dependent and governor words of the current token *7.* dependency types associated the current token *8.* whether the current token is a modifier of job title *9.* whether the current token is a non-referential pronoun
	Entity Information	*10.* unigrams/bigrams normalized by entity types *11.* dependency features normalized by entity types *12.* nearest entity type and string in the sentence/clause
Argument	Basic	*1.* context words of the entity mention *2.* trigger word and subtype *3.* entity type, subtype and entity role if it is a geo-political entity mention *4.* entity mention head, and head of any other name mention from co-reference chain *5.* lexical distance between the argument candidate and the trigger *6.* the relative position between the argument candidate and the trigger: {before, after, overlap, or separated by punctuation} *7.* whether it is the nearest argument candidate with the same type *8.* whether it is the only mention of the same entity type in the sentence
	Syntactic	*9.* dependency path between the argument candidate and the trigger *10.* path from the argument candidate and the trigger in constituent parse tree *11.* length of the path between the argument candidate and the trigger in dependency graph *12.* common root node and its depth of the argument candidate and parse tree *13.* whether the argument candidate and the trigger appear in the same clause

Table 1: Local features.

[Li et al. 2013]

Category	Type	Feature Description
Trigger	Lexical	*1.* unigrams/bigrams of the current and context words within the window of size 2 *2.* unigrams/bigrams of part-of-speech tags of the current and context words within the window of size 2 *3.* lemma and synonyms of the current token *4.* base form of the current token extracted from Nomlex (Macleod et al., 1998) *5.* Brown clusters that are learned from ACE English corpus (Brown et al., 1992; Miller et al., 2004; Sun et al., 2011). We used the clusters with prefixes of length 13, 16 and 20 for each token.
	Syntactic	*6.* dependent and governor words of the current token *7.* dependency types associated the current token *8.* whether the current token is a modifier of job title *9.* whether the current token is a non-referential pronoun
	Entity Information	*10.* unigrams/bigrams normalized by entity types *11.* dependency features normalized by entity types *12.* nearest entity type and string in the sentence/clause
Argument	Basic	*1.* context words of the entity mention *2.* trigger word and subtype *3.* entity type, subtype and entity role if it is a geo-political entity mention *4.* entity mention head, and head of any other name mention from co-reference chain *5.* lexical distance between the argument candidate and the trigger *6.* the relative position between the argument candidate and the trigger: {before, after, overlap, or separated by punctuation} *7.* whether it is the nearest argument candidate with the same type *8.* whether it is the only mention of the same entity type in the sentence
	Syntactic	*9.* dependency path between the argument candidate and the trigger *10.* path from the argument candidate and the trigger in constituent parse tree *11.* length of the path between the argument candidate and the trigger in dependency graph *12.* common root node and its depth of the argument candidate and parse tree *13.* whether the argument candidate and the trigger appear in the same clause

Table 1: Local features.

:)

[Li et al. 2013]

Category	Feature Description
Trigger	*1.* bigram of trigger types occur in the same sentence or the same clause *2.* binary feature indicating whether synonyms in the same sentence have the same trigger label *3.* context and dependency paths between two triggers conjuncted with their types
Argument	*4.* context and dependency features about two argument candidates which share the same role within the same event mention *5.* features about one argument candidate which plays as arguments in two event mentions in the same sentence *6.* features about two arguments of an event mention which are overlapping *7.* the number of arguments with each role type of an event mention conjuncted with the event subtype *8.* the pairs of time arguments within an event mention conjuncted with the event subtype

Table 2: Global features.

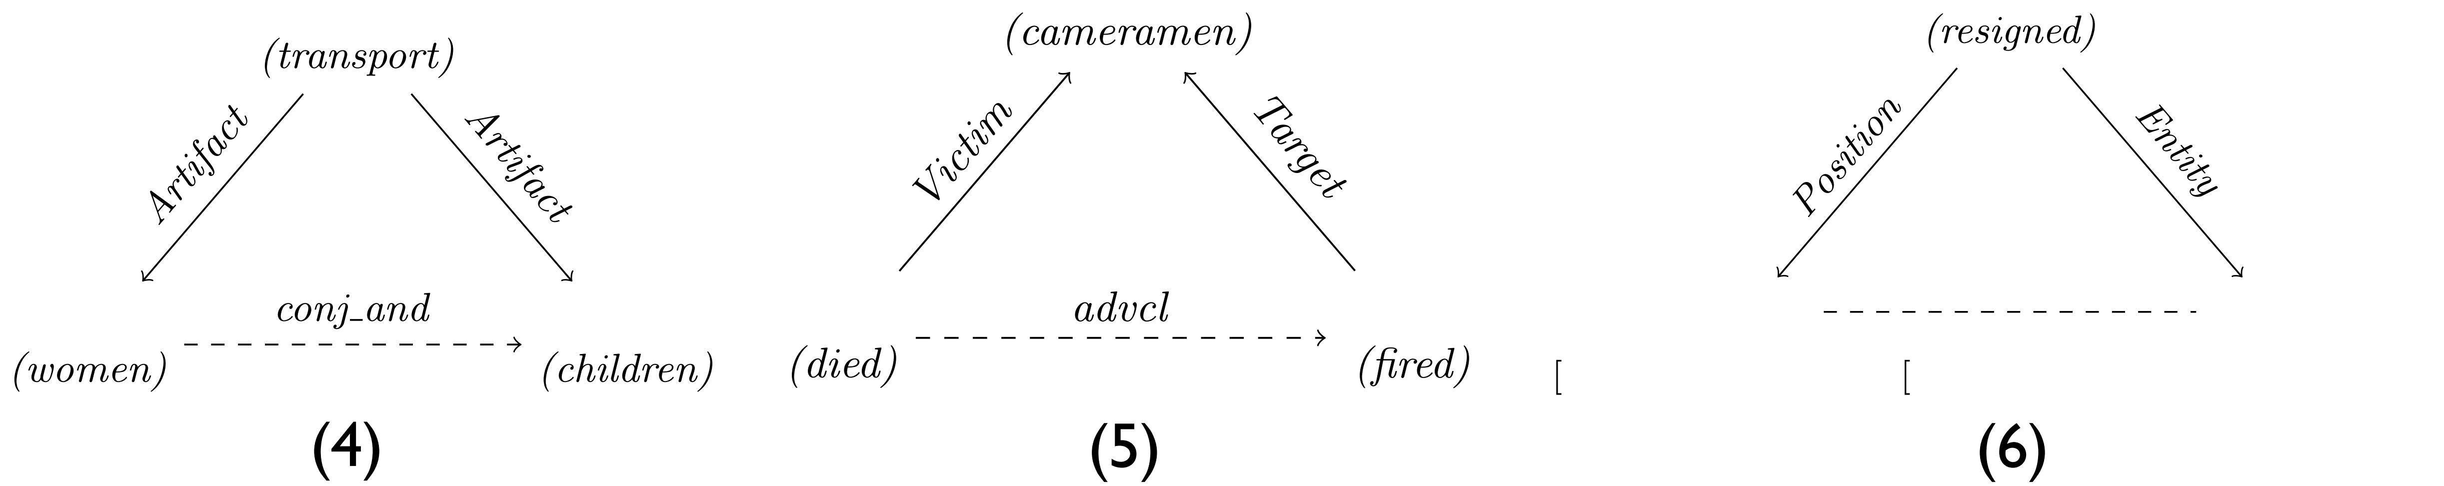

[Li et al. 2013]