

Semantic Roles & Labeling

LING 571 — Deep Processing in NLP

November 16, 2022

Shane Steinert-Threlkeld

Announcements

- HW7:
 - Be careful with 0s (e.g. log, divide by 0)
 - Similarity vs. distance
 - Detailed readme! How did you treat sentence boundaries?

`scipy.spatial.distance.cosine(u, v, w=None)`

Compute the Cosine distance between 1-D arrays.

The Cosine distance between u and v , is defined as

$$1 - \frac{u \cdot v}{\|u\|_2 \|v\|_2}.$$

where $u \cdot v$ is the dot product of u and v .

Parameters:

u : (N,) array_like

Input array.

v : (N,) array_like

Input array.

w : (N,) array_like, optional

The weights for each value in u and v . Default is None, which gives each value a weight of 1.0

Returns:

cosine : double

The Cosine distance between vectors u and v .

Ambiguity of the Week



<https://twitter.com/venzann/status/1329337844175814656>

Ambiguity of the Week



Ambiguity of the Week



Semantic Roles

Semantic Analysis

- **Full, deep compositional semantics**
 - Creates full logical form
 - Links sentence meaning representation to logical world model representation
 - Powerful, expressive, AI-complete

Semantic Analysis

- **Full, deep compositional semantics**
 - Creates full logical form
 - Links sentence meaning representation to logical world model representation
 - Powerful, expressive, AI-complete
- **Domain-specific slot-filling:**
 - Common in dialog systems, IE tasks
 - Narrowly targeted to domain/task
 - e.g. ORIGIN_LOC, DESTINATION_LOC, AIRLINE, ...
 - Often pattern-matching
 - Low cost, but lacks generality, richness, etc

Semantic Role Labeling

- Typically want to know
 - *Who* did *what* to *whom*
 - ...*where*, *when*, and *how*

Semantic Role Labeling

- Typically want to know
 - *Who* did *what* to *whom*
 - ...*where*, *when*, and *how*
- Intermediate level:
 - Shallower than full deep composition
 - Abstracts away (somewhat) from surface form
 - Captures general predicate-argument structure info
 - Balance generality and specificity

Examples

Yesterday Tom chased Jerry

Yesterday Jerry was chased by Tom

Tom chased Jerry yesterday

Jerry was chased yesterday by Tom

Examples

Yesterday Tom chased Jerry

Yesterday Jerry was chased by Tom

Tom chased Jerry yesterday

Jerry was chased yesterday by Tom

- Semantic roles:
 - *Chaser*: Tom
 - *ChasedThing*: Jerry
 - *TimeOfChasing*: yesterday

Examples

Yesterday Tom chased Jerry

Yesterday Jerry was chased by Tom

Tom chased Jerry yesterday

Jerry was chased yesterday by Tom

- Semantic roles:
 - *Chaser*: Tom
 - *ChasedThing*: Jerry
 - *TimeOfChasing*: yesterday
- Same across all sentence forms

Full Event Semantics

- Neo-Davidsonian Style:
 - $\exists e \text{ } Chasing(e) \wedge \textcolor{green}{Chaser}(e, Tom) \wedge \textcolor{brown}{ChasedThing}(e, Jerry)$
 $\wedge \textcolor{teal}{TimeOfChasing}(e, Yesterday)$

Full Event Semantics

- Neo-Davidsonian Style:
 - $\exists e \text{ } Chasing(e) \wedge \textcolor{green}{Chaser}(e, Tom) \wedge \textcolor{brown}{ChasedThing}(e, Jerry)$
 $\wedge \textcolor{teal}{TimeOfChasing}(e, Yesterday)$
- Same across all examples

Full Event Semantics

- Neo-Davidsonian Style:
 - $\exists e \text{ Chasing}(e) \wedge \text{Chaser}(e, \text{Tom}) \wedge \text{ChasedThing}(e, \text{Jerry})$
 $\wedge \text{TimeOfChasing}(e, \text{Yesterday})$
- Same across all examples
- Roles: *Chaser*, *ChasedThing*, *TimeOfChasing*
 - Specific to verb “chase”
 - a.k.a. “Deep roles”

Main Idea

- Extract the semantic roles *without doing full semantic parsing*
- Easier problem, but still useful for many tasks
 - More data
 - Better models

Issues & Challenges

Issues & Challenges

- How many roles for a language?
 - Arbitrary!
 - Each verb's event structure determines sets of roles

Issues & Challenges

Issues & Challenges

- How can we acquire these roles?
 - Manual construction?
 - Some progress on automatic learning
 - Mostly successful on limited domains (ATIS, GeoQuery)

Issues & Challenges

Issues & Challenges

- Can we capture generalities across verbs/events?
 - Not really, each event/role is specific

Thematic Roles

Thematic Roles

- Solution to instantiating a specific role for every verb

Thematic Roles

- Solution to instantiating a specific role for every verb
- Attempt to capture commonality between roles

Thematic Roles

- Describe common semantic roles of verbal arguments
 - e.g. subject of *break* is AGENT
 - AGENT: volitional cause
 - THEME: things affected by action

Thematic Roles

- Describe common semantic roles of verbal arguments
 - e.g. subject of *break* is AGENT
 - AGENT: volitional cause
 - THEME: things affected by action
- Enables generalization over surface order of arguments
 - John_{AGENT} broke the window_{THEME}
 - The rock_{INSTRUMENT} broke the window_{THEME}
 - The window_{THEME} was broken by John_{AGENT}

Thematic Roles

- Verbs take different roles
- The *break* verb could be formed as:

Thematic Roles

- Verbs take different roles
- The ***break*** verb could be formed as:
 - AGENT/Subject, THEME/Object

(John broke the window)

Thematic Roles

- Verbs take different roles
- The ***break*** verb could be formed as:
 - AGENT/Subject, THEME/Object *(John broke the window)*
 - AGENT/Subject, THEME/Object, INSTRUMENT/PP_{with} *(John broke the window with a rock)*

Thematic Roles

- Verbs take different roles
- The **break** verb could be formed as:
 - AGENT/Subject, THEME/Object *(John broke the window)*
 - AGENT/Subject, THEME/Object, INSTRUMENT/PP_{with} *(John broke the window with a rock)*
 - INSTRUMENT/Subject, THEME/Object *(The rock broke the window)*

Thematic Roles

- Verbs take different roles
- The **break** verb could be formed as:
 - AGENT/Subject, THEME/Object *(John broke the window)*
 - AGENT/Subject, THEME/Object, INSTRUMENT/PP_{with} *(John broke the window with a rock)*
 - INSTRUMENT/Subject, THEME/Object *(The rock broke the window)*
 - THEME/Subject *(The window was broken)*

Thematic Roles

- Thematic grid, Θ -grid, case frame
 - Set of thematic role arguments of verb
 - subject: AGENT; Object: THEME, or
 - subject: INSTR; Object: THEME

Thematic Roles

- Thematic grid, Θ -grid, case frame
 - Set of thematic role arguments of verb
 - subject: AGENT; Object: THEME, or
 - subject: INSTR; Object: THEME
- Verb/Diathesis Alternations
 - Verbs allow different surface realizations of roles
 - Doris_{AGENT} gave the book_{THEME} to Carv_{GOAL}
 - Doris_{AGENT} gave Carv_{GOAL} the book_{THEME}

Canonical Roles

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

Thematic Role Issues

- Hard to produce

Thematic Role Issues

- Hard to produce
- Standard set of roles
 - Fragmentation: Often need to make more specific
 - e.g. INSTRUMENTs can be subject or not

Thematic Role Issues

- Hard to produce
- Standard set of roles
- Fragmentation: Often need to make more specific
 - e.g. **INSTRUMENTS** can be subject or not

From Levin and Rappaport Hovav 2005:

- a. John broke the window with a rock.
- b. The rock broke the window.
- a. Swabha ate the banana with a fork.
- b. * The fork ate the banana.

Thematic Role Issues

- Hard to produce
- Standard set of roles
 - Fragmentation: Often need to make more specific
 - e.g. **INSTRUMENTS** can be subject or not
- Standard definition of roles
 - Most **AGENTS**: animate, volitional, sentient, causal
 - But not all... e.g.?

From Levin and Rappaport Hovav 2005:

- a. John broke the window with a rock.
- b. The rock broke the window.
- a. Swabha ate the banana with a fork.
- b. * The fork ate the banana.

Thematic Role Issues

- Hard to produce
- Standard set of roles
 - Fragmentation: Often need to make more specific
 - e.g. **INSTRUMENTS** can be subject or not
- Standard definition of roles
 - Most **AGENTS**: animate, volitional, sentient, causal
 - But not all... e.g.?

From Levin and Rappaport Hovav 2005:

- a. John broke the window with a rock.
- b. The rock broke the window.
- a. Swabha ate the banana with a fork.
- b. * The fork ate the banana.

[Google]_{Agent} found the answer.

Thematic Role Issues

- Strategies:
 - Generalized semantic roles: PROTO-AGENT/PROTO-PATIENT
 - Defined heuristically: PropBank
 - Semantic “proto”-roles: <http://decomp.io/projects/semantic-proto-roles/>
- Define roles specific to verbs/nouns: [FrameNet](#)

PropBank

- Sentences annotated with semantic roles
 - Penn and Chinese Treebank
 - Roles specific to verb sense
 - Numbered: Arg₀, Arg₁, Arg₂, ...
 - Arg₀: PROTO-AGENT; Arg₁: PROTO-PATIENT, etc

PropBank

- Arguments >1 are Verb-specific
 - e.g. *agree.01*
 - Arg₀: Agreeer
 - Arg₁: Proposition
 - Arg₂: Other entity agreeing
 - Ex1: [Arg₀ The group] agreed [Arg₁ it wouldn't make an offer]

PropBank

- Resources:
 - Annotated sentences
 - Started w/Penn Treebank
 - Now: Google answerbank, SMS, webtext, etc
- Framesets:
 - Per-sense inventories of roles, examples
 - Span verbs, adjectives, nouns (e.g. event nouns)

PropBank

- proppbank.github.io
- Recent status:
 - 5940 verbs w/8121 framesets
 - 1880 adjectives w/2210 framesets
- Continued into [OntoNotes](#)
- [CoNLL 2005 and 2012 shared tasks]


AMR

- “Abstract Meaning Representation”
 - Sentence-level semantic representation
- Nodes: Concepts
 - English words; PropBank: predicates; or keywords (‘person’)
- Edges: Relations
 - PropBank thematic roles (ARG0-ARG5)
 - Others including ‘location,’ ‘name,’ ‘time,’ etc...
 - ~100 in total

AMR 2

- AMR Bank: (now) ~40K annotated sentences
- JAMR parser: 63% F-measure (2015)
 - Alignments between word spans & graph fragments
- Example: *“I saw Joe’s dog, which was running in the garden.”*

From [Liu et. al \(2015\)](#)



AMR 3

- Towards full semantic parsing
- “Deeper” than base PropBank, but:
 - No real quantification
 - No articles
 - No real vs. hypothetical events (e.g. “wants to go”)

FrameNet (Fillmore et al)

- Key insight:
 - Commonalities not just across different sentences w/same verb but across different verbs (and nouns and adjectives)

FrameNet (Fillmore et al)

- Key insight:
 - Commonalities not just across different sentences w/same verb but across different verbs (and nouns and adjectives)
- **PropBank**
 - [Arg0 Big Fruit Co.] increased [Arg1 the price of bananas].
 - [Arg1 The price of bananas] was increased by [Arg0 BFCo].
 - [Arg1 The price of bananas] increased [Arg2 5%].

FrameNet (Fillmore et al)

- Key insight:
 - Commonalities not just across different sentences w/same verb but across different verbs (and nouns and adjectives)
- **PropBank**
 - [Arg0 Big Fruit Co.] increased [Arg1 the price of bananas].
 - [Arg1 The price of bananas] was increased by [Arg0 BFCo].
 - [Arg1 The price of bananas] increased [Arg2 5%].
- **FrameNet**
 - [ATTRIBUTE The price] of [ITEM bananas] increased [DIFF 5%].
 - [ATTRIBUTE The price] of [ITEM bananas] rose [DIFF 5%].
 - There has been a [DIFF 5%] rise in [ATTRIBUTE the price] of [ITEM bananas].

FrameNet

- Semantic roles specific to frame
 - Frame: script-like structure, roles (frame elements)
 - e.g. CHANGE_POSITION_ON_SCALE: increase, rise
 - ATTRIBUTE; INITIAL_VALUE; FINAL_VALUE
- Core, non-core roles
- Relationships between frames, frame elements
 - Add causative: CAUSE_CHANGE_POSITION_ON_SCALE

Change of position on scale

VERBS:	<i>dwindle</i>	<i>move</i>	<i>soar</i>	<i>escalation</i>	<i>shift</i>
<i>advance</i>	<i>edge</i>	<i>mushroom</i>	<i>swell</i>	<i>explosion</i>	<i>tumble</i>
<i>climb</i>	<i>explode</i>	<i>plummet</i>	<i>swing</i>	<i>fall</i>	
<i>decline</i>	<i>fall</i>	<i>reach</i>	<i>triple</i>	<i>fluctuation</i>	ADVERBS:
<i>decrease</i>	<i>fluctuate</i>	<i>rise</i>	<i>tumble</i>	<i>gain</i>	<i>increasingly</i>
<i>diminish</i>	<i>gain</i>	<i>rocket</i>		<i>growth</i>	
<i>dip</i>	<i>grow</i>	<i>shift</i>	NOUNS:	<i>hike</i>	
<i>double</i>	<i>increase</i>	<i>skyrocket</i>	<i>decline</i>	<i>increase</i>	
<i>drop</i>	<i>jump</i>	<i>slide</i>	<i>decrease</i>	<i>rise</i>	

Core Roles

Core Roles

ATTRIBUTE	The ATTRIBUTE is a scalar property that the ITEM possesses.
DIFFERENCE	The distance by which an ITEM changes its position on the scale.
FINAL_STATE	A description that presents the ITEM's state after the change in the ATTRIBUTE's value as an independent predication.
FINAL_VALUE	The position on the scale where the ITEM ends up.
INITIAL_STATE	A description that presents the ITEM's state before the change in the ATTRIBUTE's value as an independent predication.
INITIAL_VALUE	The initial position on the scale from which the ITEM moves away.
ITEM	The entity that has a position on the scale.
VALUE_RANGE	A portion of the scale, typically identified by its end points, along which the values of the ATTRIBUTE fluctuate.

Some Non-Core Roles

DURATION	The length of time over which the change takes place.
SPEED	The rate of change of the VALUE.
GROUP	The GROUP in which an ITEM changes the value of an ATTRIBUTE in a specified way.

FrameNet

- Current status:
 - 1224 frames
 - 13686 lexical units (mostly verbs, nouns)
 - 10749 frame element relations
 - Annotations over:
 - Newswire (WSJ, AQUAINT)
 - American National Corpus
- Under active development
- Still relatively limited coverage

Semantic Role Labeling

Semantic Role Labeling

- Task of automatically assigning semantic roles for each argument

Typical Strategy

- Assign Parse to Input String
- Traverse parse to find all predicates
- For each predicate, examine each node and decide semantic role (if any)

Typical Strategy

function SEMANTICROLELABEL(*words*) **returns** **labeled tree**

parse ← PARSE(*words*)

for each *predicate* **in** *parse* **do**

for each *node* **in** *parse* **do**

featurevector ← EXTRACTFEATURES(*node*, *predicate*, *parse*)


CLASSIFYNODE(*node*, *featurevector*, *parse*)

J&M 3rd ed, [ch 20.6](#)


Semantic Role Labeling Features

- Governing predicate
- Phrase Type (NP, VP, etc)
- Headword of constituent
- Headword POS
- PATH from current node to predicate (NP↑S↓VP↓VBD)
- ...


Typical Strategy




Typical Strategy



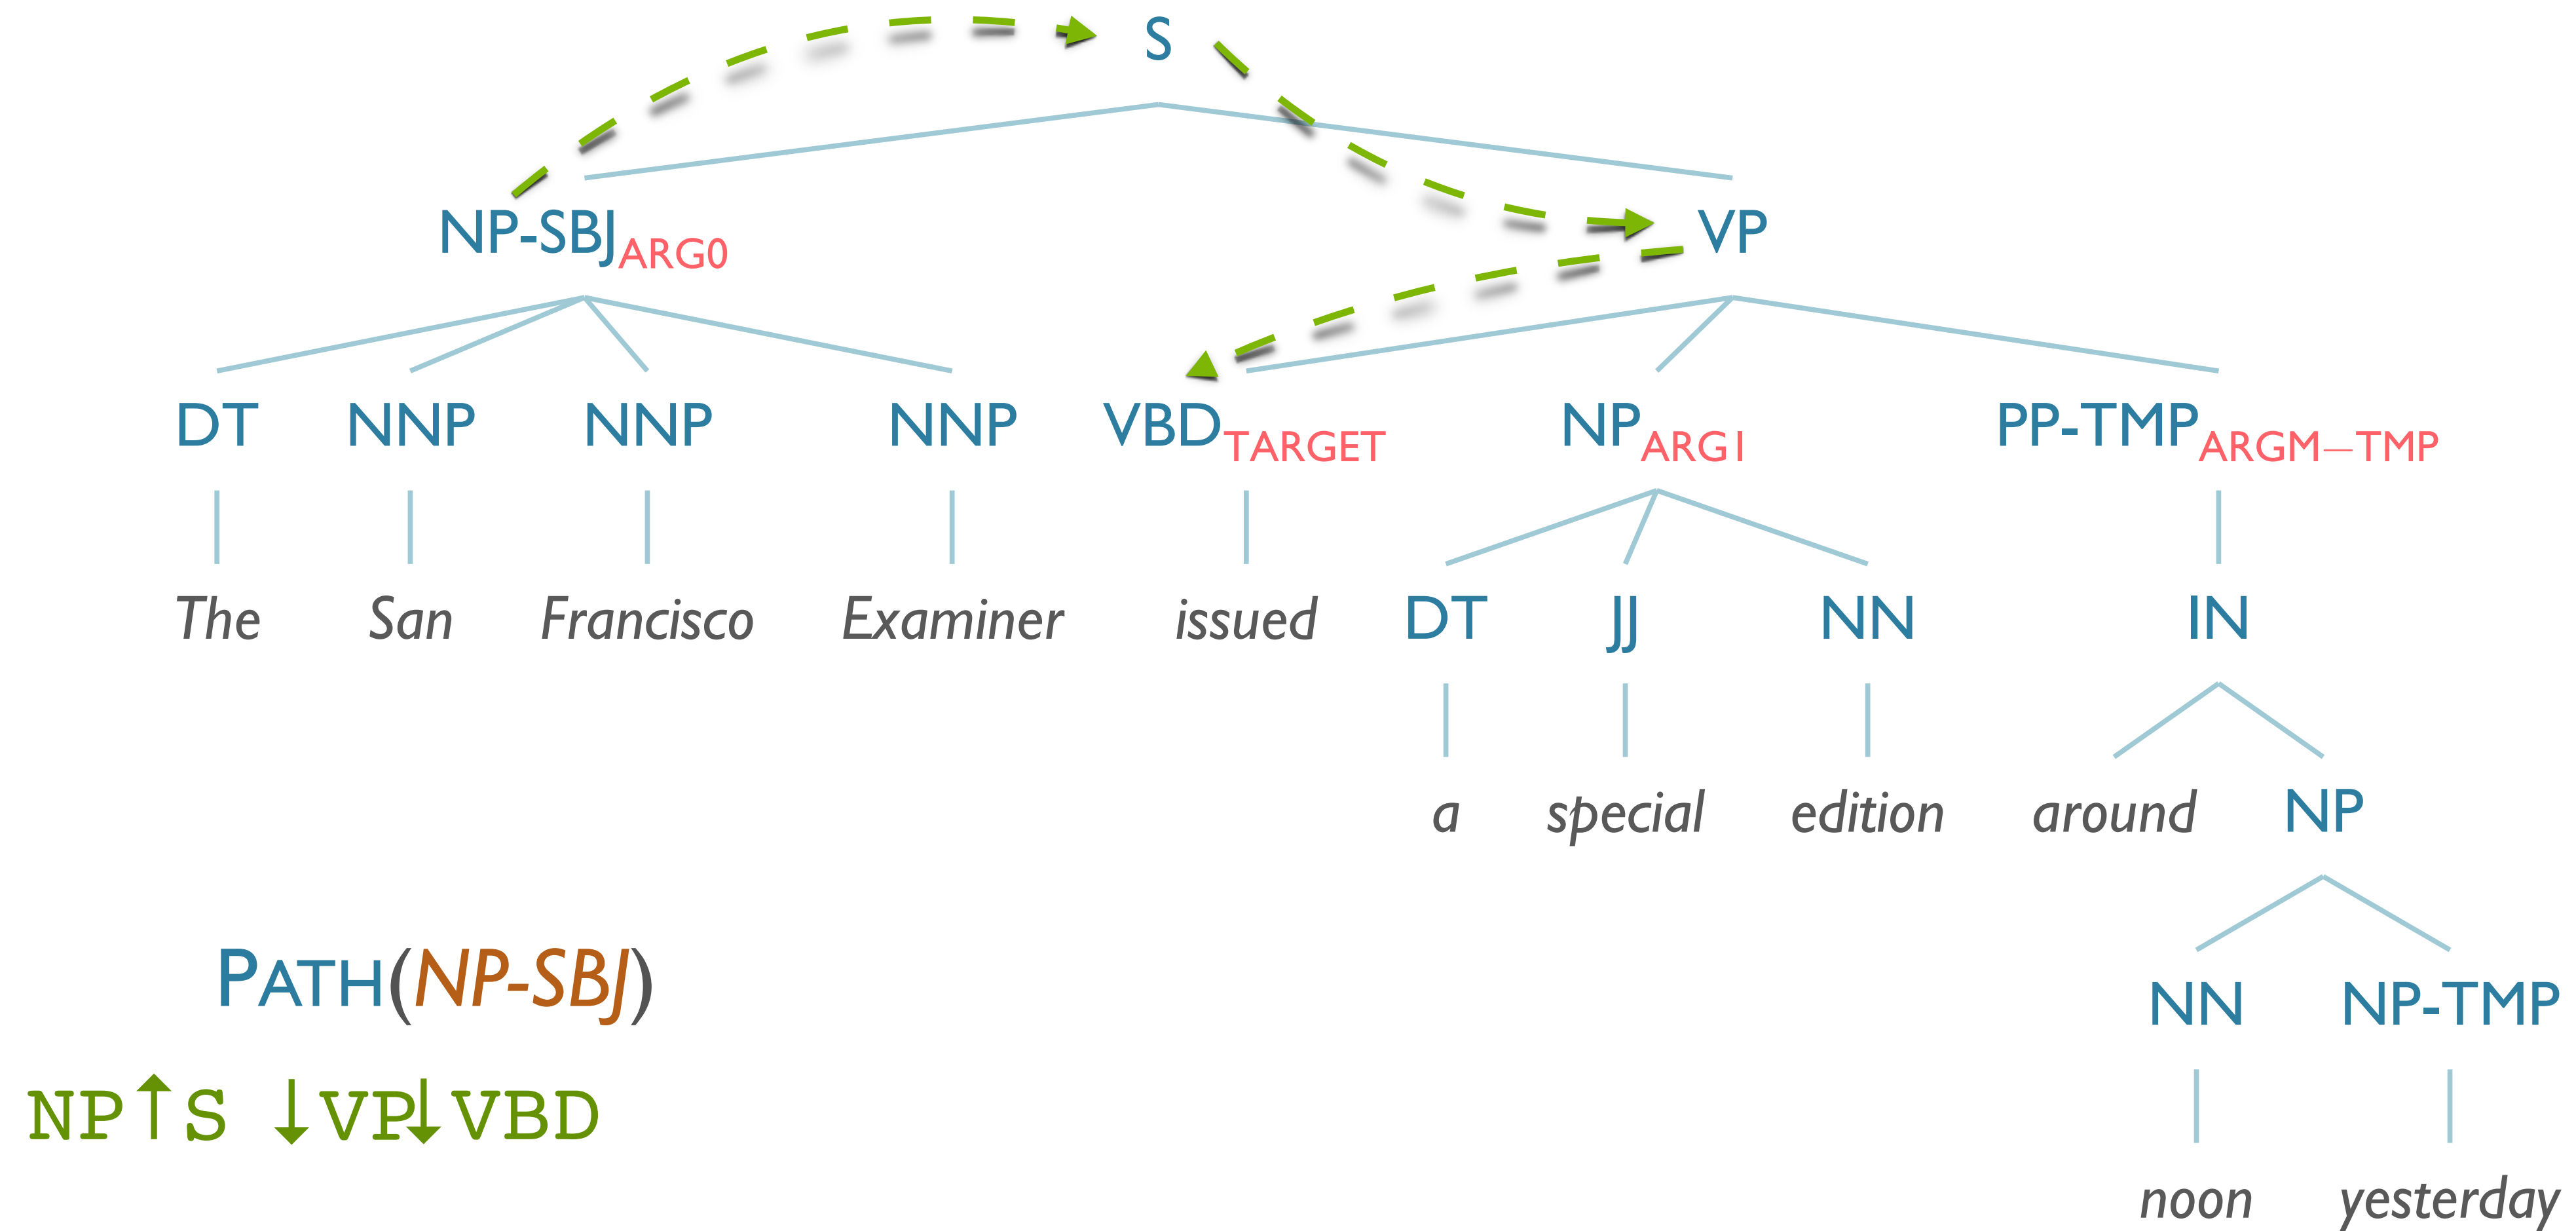
Typical Strategy



Typical Strategy



Typical Strategy




Some Semantic Role Labeling Applications

- Question answering:
 - Who did what to whom?
- Machine translation
 - Maintain agents/thematic roles through translation
- Dialogue systems

Scaling up SRL


Neural SRL

He et al 2017



Neural SRL

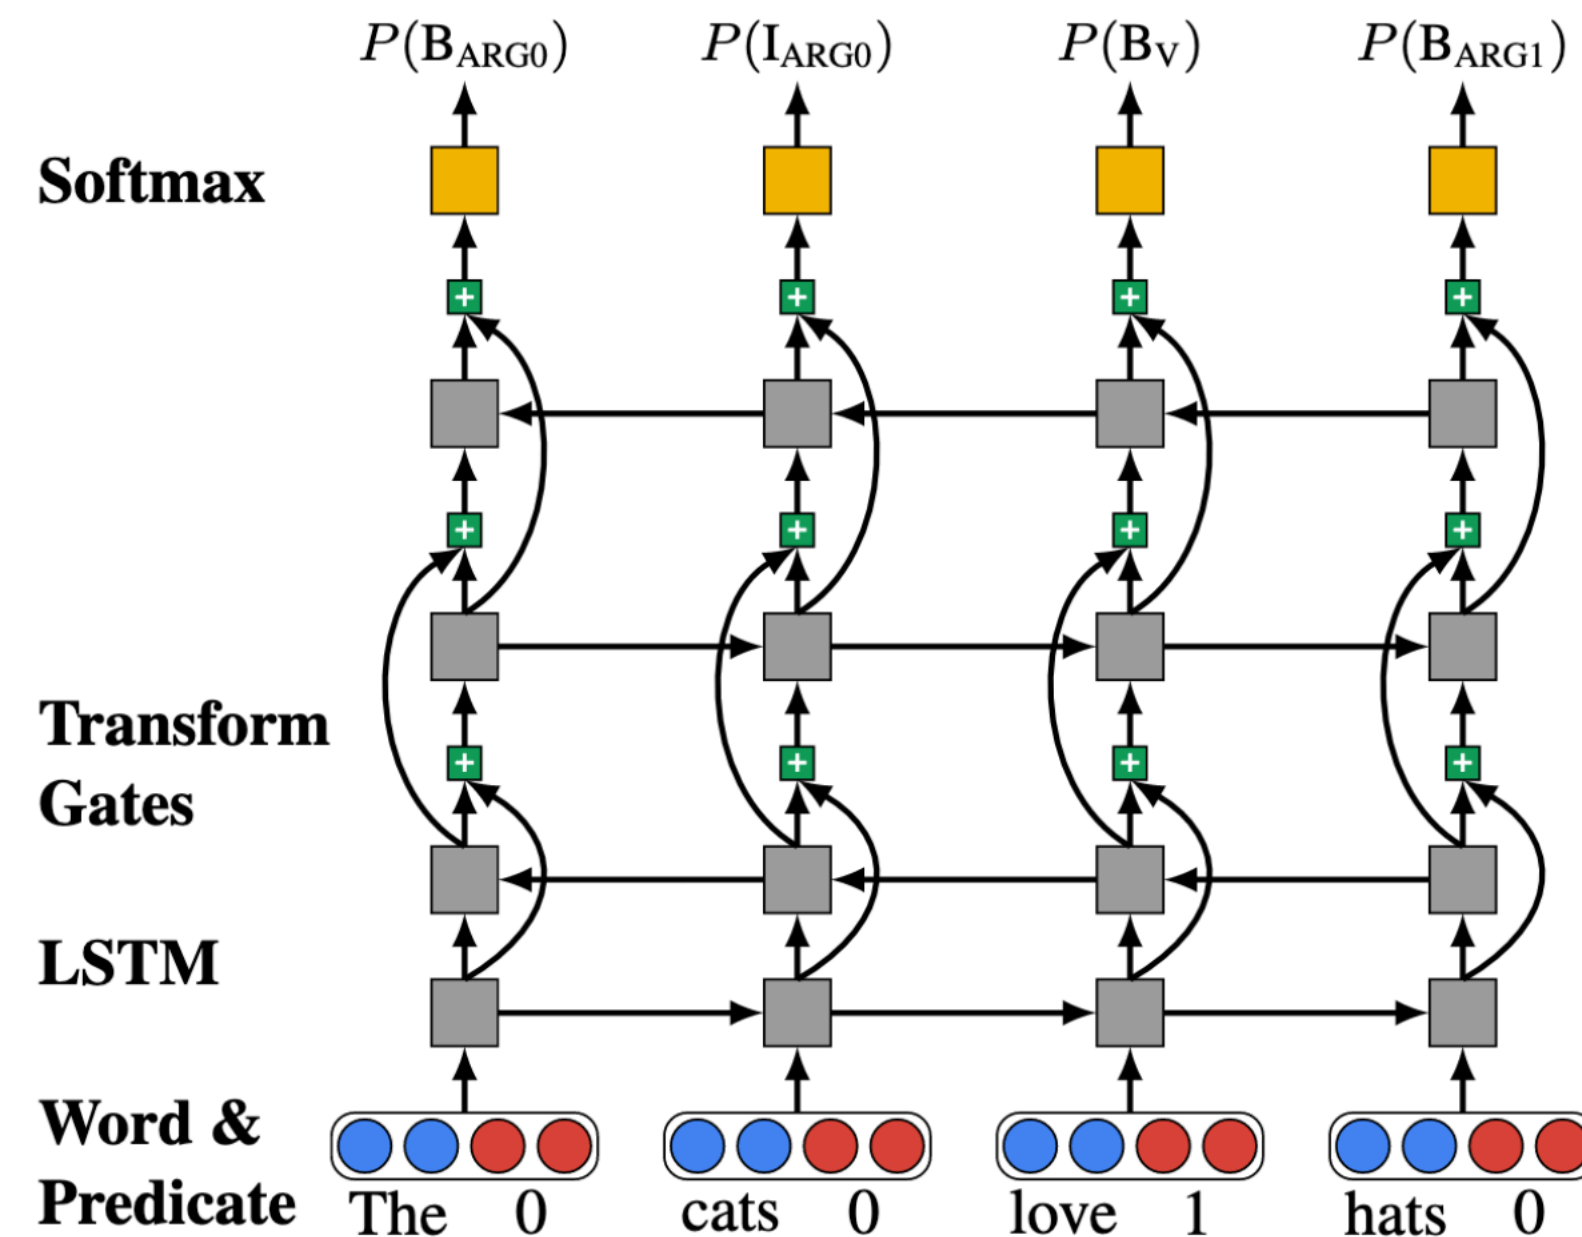
He et al 2017



No “detour” through syntactic parse

Neural SRL

He et al 2017



No “detour” through syntactic parse

Can be global or contextual
[contextual tends to improve]

QA-SRL

Question-Answer Driven Semantic Role Labeling: Using Natural Language to Annotate Natural Language

Luheng He Mike Lewis Luke Zettlemoyer

Computer Science & Engineering

University of Washington

Seattle, WA

{luheng,mlewis,lsz}@cs.washington.edu

Abstract

This paper introduces the task of question-answer driven semantic role labeling (QA-SRL), where question-answer pairs are used to represent predicate-argument structure. For example, the verb “introduce” in the previous sentence would be labeled with the questions “What is introduced?”, and “What introduces something?”, each paired with the phrase from the sentence that gives the correct answer. Posing the problem this way allows the questions themselves to define the set of possible roles, without the need for predefined frame or thematic role ontologies. It also allows for scalable data collection by annotators with very little training and no linguistic expertise. We gather data in two




Figure 1: QA-SRL annotations for a Wikipedia sentence.

(ARG0, ARG1, etc.). Existing task definitions can be complex and require significant linguistic expertise to understand,¹ causing challenges for data annotation and use in many target applications.

In this paper, we introduce a new question-answer driven SRL task formulation (QA-SRL)

the paper

QA-SRL

Question-Answer Driven Semantic Role Labeling: Using Natural Language to Annotate Natural Language

Luheng He Mike Lewis Luke Zettlemoyer

Computer Science & Engineering

University of Washington

Seattle, WA

{luheng,mlewis,lsz}@cs.washington.edu

Abstract

This paper introduces the task of question-answer driven semantic role labeling (QA-SRL), where question-answer pairs are used to represent predicate-argument structure. For example, the verb “introduce” in the previous sentence would be labeled with the questions “What is introduced?”, and “What introduces something?”, each paired with the phrase from the sentence that gives the correct answer.

Posing the problem this way allows the questions themselves to define the set of possible roles, without the need for predefined frame or thematic role ontologies. It also allows for scalable data collection by annotators with very little training and no linguistic expertise. We gather data in two

UCD *finished* the 2006 championship as Dublin champions ,
by *beating* St Vincents in the final .




Figure 1: QA-SRL annotations for a Wikipedia sentence.

(ARG0, ARG1, etc.). Existing task definitions can be complex and require significant linguistic expertise to understand,¹ causing challenges for data annotation and use in many target applications.

In this paper, we introduce a new question-answer driven SRL task formulation (QA-SRL)

the paper

QA-SRL

Editorial:
should've been /casserole/

Question-Answer Driven Semantic Role Labeling: Using Natural Language to Annotate Natural Language

Luheng He Mike Lewis Luke Zettlemoyer

Computer Science & Engineering

University of Washington

Seattle, WA

{luheng,mlewis,lsz}@cs.washington.edu

Abstract

This paper introduces the task of question-answer driven semantic role labeling (QA-SRL), where question-answer pairs are used to represent predicate-argument structure. For example, the verb “introduce” in the previous sentence would be labeled with the questions “What is introduced?”, and “What introduces something?”, each paired with the phrase from the sentence that gives the correct answer.

Posing the problem this way allows the questions themselves to define the set of possible roles, without the need for predefined frame or thematic role ontologies. It also allows for scalable data collection by annotators with very little training and no linguistic expertise. We gather data in two

UCD **finished** the 2006 championship as Dublin champions ,
by **beating** St Vincents in the final .

finished

Who finished something? - UCD
What did someone finish? - the 2006 championship
What did someone finish something as? - Dublin champions
How did someone finish something? - by beating St Vincents in the final

beating

Who beat someone? - UCD
When did someone beat someone? - in the final
Who did someone beat? - St Vincents

Figure 1: QA-SRL annotations for a Wikipedia sentence.

(ARG0, ARG1, etc.). Existing task definitions can be complex and require significant linguistic expertise to understand,¹ causing challenges for data annotation and use in many target applications.

In this paper, we introduce a new question-answer driven SRL task formulation (QA-SRL)

the paper

QA-SRL vs. PropBank

Sentence	CoNLL-2009		QA-SRL	
(1) Stock-fund managers , meantime , went into October with less cash on hand than they held earlier this year .	A0 AM-TMP	they year	Who had held something? When had someone held something? What had someone held? Where had someone held something?	Stock-fund managers / they earlier this year less cash on hand on hand
(2) Mr. Spielvogel added pointedly : “ The pressure on commissions did n’t begin with Al Achenbaum . ”	A0 A1 AM-MNR	Spielvogel did pointedly	Who added something? What was added? How was something added?	Mr. Spielvogel “ The pressure on commissions did n’t begin with Al Achenbaum . ” pointedly
(3) He claimed losses totaling \$ 42,455 – and the IRS denied them all .	A0 A1	IRS them	Who denied something? What was denied?	IRS losses / them
(4) The consumer - products and newsprint company said net rose to \$ 108.8 million , or \$ 1.35 a share , from \$ 90.5 million , or \$ 1.12 a share , a year ago .	A1 A3 A4	net \$/ago to	What rose? What did something rise from? What did something rise to? When did something rise?	net \$ 90.5 million , or \$ 1.12 a share \$ 108.8 million , or \$ 1.35 a share a year ago
(5) Mr. Agnew was vice president of the U.S. from 1969 until he resigned in 1973 .	A0 AM-TMP	he in	Who resigned from something? When did someone resign from something? What did someone resign from?	Mr. Agnew 1973 vice president of the U.S.

QA-SRL

- Much more info, including live data explorer:
 - <http://qasrl.org/>
- AI2 NLP Highlights podcast episode ft. Luke Zettlemoyer:
 - <https://soundcloud.com/nlp-highlights/96-question-answering-as-an-annotation-format-with-luke-zettlemoyer>
- For large-scale, “natural”/easy annotations, see also:
 - <http://decomp.io/projects/semantic-proto-roles/>
 - (And the other projects there beyond SRL as well)