

Analysis and Refinement of Temporal Relation Aggregation

Taylor Cassidy^{1,2}, Heng Ji³

¹IBM T. J. Watson Research Center ²U.S. Army Research Lab

³Computer Science Department, Rensselaer Polytechnic Institute




Rensselaer

Outline



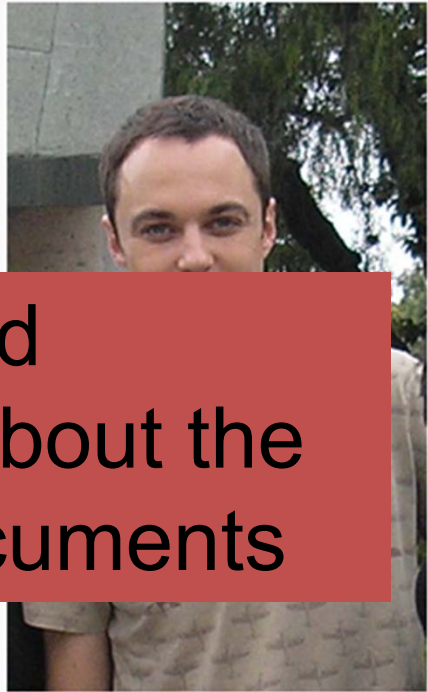
- Background: Adding time information to relations
- Aggregation and Classification Challenges
- Solutions
- Future work



Slot-filling



Jim Parsons



Parsons in 2008

Born	James Joseph Parsons March 24, 1973 (age 37) Houston, Texas, U.S.
Occupation	Actor
Years active	2000–present

Temporal Slot-filling: Gather and aggregate temporal information about the same relation across multiple documents

Jim Parsons, a graduate of **School Attended: University of Houston** School of Theater and Dance, won the Emmv on Sunday for Lead Actor in a Comedy Series for his work on The Big Bang Theory.

Mentions

Individuals

... **But when?**

What is “Temporal Information”?

- **Focus:** Information about a **relation** between two named entities expressible in terms of its **start and end time**
- **Simplifying Assumption:** relations are “interval shaped”
 - One occurrence per relation
 - No Gaps
- **Task:**
 - **Input:** mention of a relation between two named entities in a document.
 - **Given:** corpus of millions of newswire, weblog, and discussion forum documents
 - **Output:** aggregated temporal information about the relation that can be inferred (collectively) **from the entire corpus**




Grover Cleveland

Four-tuple Temporal Representation*



spouse(Jennifer Jones, Norton Simon) in 1971, in a ceremony on a yacht in the English Channel after a courtship of three weeks



*(Amigó et al., 2011)

Scoring Metric*

- Temporal Quality

- Let $S = \langle t^{(1)}, t^{(2)}, t^{(3)}, t^{(4)} \rangle$ be system output, $G = \langle g^{(1)}, g^{(2)}, g^{(3)}, g^{(4)} \rangle$ be gold standard

$$Q(S) = \frac{1}{4} \sum_i \frac{c}{c + |t^{(i)} - g^{(i)}|}$$

- An error of c time units produces a 0.5 score; scores produced with $c = 1$ year
- Each element in tuple is scored independently

- Overall Metric


$$P = \frac{\sum_{S^i \in C(S)} Q(S^i)}{M} \quad R = \frac{\sum_{S^i \in C(S)} Q(S^i)}{N}$$

- M : the number of system output tuples
- N : the number of gold standard tuples
- $C(S)$: the number of instances that have correct slot fills

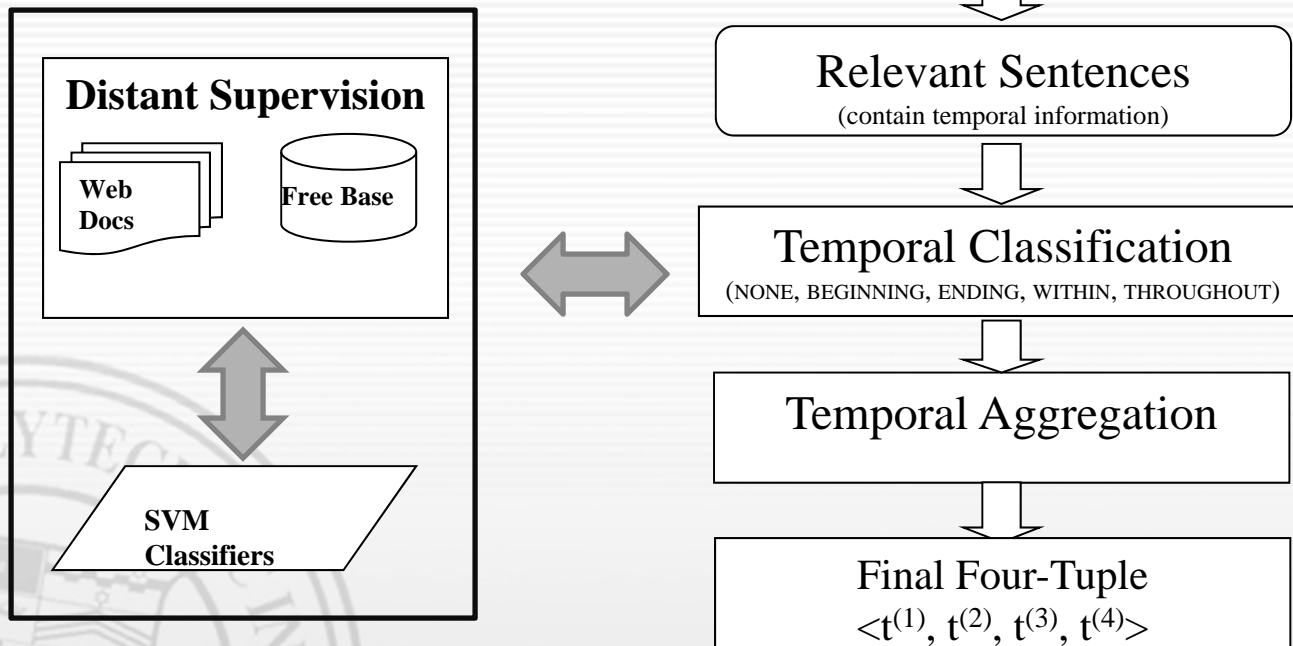
*(Ji, Grishman, & Dang, 2011)

Evaluation Examples

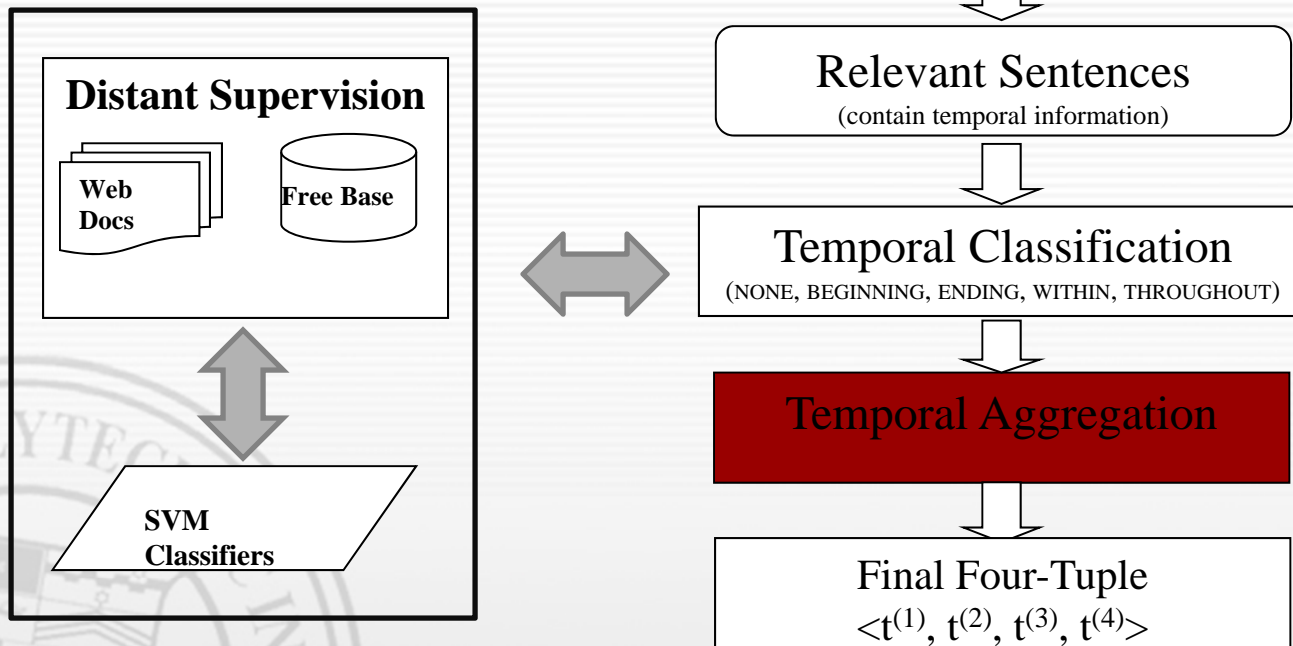
Infinite = 10000 -Infinite = 0
 $C_{\text{cons}} = C_{\text{vag}} = 5$



Approach Overview



Approach Overview



(Ji et al., 2013 KIS)

Intermediate Relations

Int. Rel	four-tuple
BEGINNING	$\langle s, e, s, \infty \rangle$
ENDING	$\langle -\infty, e, s, e \rangle$
BEG_AND_END	$\langle s, e, s, e \rangle$
WITHIN	$\langle -\infty, e, s, \infty \rangle$
THROUGHOUT	$\langle -\infty, s, e, \infty \rangle$
BEFORE_START	$\langle e, \infty, e, \infty \rangle$
AFTER_END	$\langle -\infty, s, -\infty, s \rangle$
BEFORE_START*	$\langle s, \infty, s, \infty \rangle$
AFTER_END*	$\langle -\infty, e, -\infty, e \rangle$
NONE	$\langle -\infty, \infty, -\infty, \infty \rangle$

Jones married Norton **Simon**, in **1971**, ...

$\langle 1971-01-01, 1971-12-31, 1971-01-01, \infty \rangle$

$\langle spouse(Jones, Simon), 1971 \rangle$ **classified as** BEGINNING

+ *Aggregation Operator*

A spokeswoman for the **Norton Simon** Museum says **Jones** died **Thursday** ...

$\langle -\infty, 2009-12-11, -\infty, 2009-12-11 \rangle$

$\langle spouse(Jones, Simon), 2009-12-11 \rangle$ **classified as** AFTER_END

|| *Final Four-tuple*

$\langle 1971-01-01, 1971-12-31, 1971-01-01, 2009-12-11 \rangle$

*(Ji, Grishman, & Dang, 2011)

How to Aggregate?

- 1. Gather instances from the corpus in which a relation **r** co-occurs with a time expression **γ**.
- 2. label each classification instance **<r, γ>** with an intermediate relation (e.g. WITHIN)
- 3. Aggregate all **<r, γ>_i** to yield one final four-tuple for r
- A four-tuple is “consistent” iff.
 - **1. $t^{(1)} \leq t^{(2)}$**
 - **2. $t^{(3)} \leq t^{(4)}$**
 - **3. $t^{(1)} \leq t^{(4)}$**

How to Aggregate?

- **Simple Aggregation:**

Max-Constrain*: $T_{\text{final}} = \langle \max t^{(1)}, \min t^{(2)}, \max t^{(3)}, \min t^{(4)} \rangle$

- **Take the latest $t^{(1)}$ and $t^{(3)}$ from all four-tuples, and the earliest $t^{(2)}$ and $t^{(4)}$ from all four-tuples.**

- **Problem:** systems yield many incorrect four-tuples, resulting in a poor final four-tuple under max-constrain
- **Previous work's solution:** order four-tuples by classifier confidence**

* (Ji, Grishman, Dang, 2011) ** (Ji et al., 2013 KIS)

How to Aggregate four-tuples?

- Order four-tuples by confidence.
- Apply max-constrain to the top two.
- Accept the resulting four-tuple as the new #1 only if it is consistent
- Repeat

Iterative Pair-Wise Max-constrain*

T_0 and T_1 are inconsistent –
don't aggregate!

T_0 < $-\infty$, 1999-12-31, 1999-01-01, 1999-12-31 >

*classifier
confidence*

99%

T_1 < $-\infty$, 2009-07-31, 2009-07-01, ∞ >

71%

T_2 < $-\infty$, 1998-12-31, 1998-01-01, ∞ >

65%

T_0 combines with T_2
to yield a more
informative,
consistent four-tuple


< $-\infty$, 1998-12-31, 1999-01-01, 1999-12-31 >

*(Ji et al., 2013 KIS)

How to Aggregate four-tuples?



- **Problem:** systems yield many incorrect four-tuples
- **Previous work's solution:** order four-tuples by classifier confidence
- **Remaining Problems:**
 - Many incorrect four-tuples have high classifier confidence.
 - In particular: many relations are classified as holding at the document creation time
 - It only takes one "bad apple" to derail the process and corrupt the final four-tuple.



Addressing Aggregation Challenges Rensselaer

- Suppose T_1 is **incorrect**, but was assigned high confidence...

T_0 $\langle -\infty, 1999-12-31, 1999-01-01, 1999-12-31 \rangle$ 71%

T_1 $\langle -\infty, 2009-07-31, 2009-07-01, \infty \rangle$ 99%

T_2 $\langle -\infty, 1998-12-31, 1998-01-01, \infty \rangle$ 65%



“Bad Apple” effect

- T_1 is incorrect, but has high confidence



T_1	$\langle -\infty, 2009-07-31, 2009-07-01, \infty \rangle$	99%
T_0	$\langle -\infty, 1999-12-31, 1999-01-01, 1999-12-31 \rangle$	71%
T_2	$\langle -\infty, 1998-12-31, 1998-01-01, \infty \rangle$	65%

$\langle -\infty, 1998-12-31, 2009-07-01, \infty \rangle$

- Only four-tuples consistent with the incorrect T_1 can contribute to the final four-tuple

How to Aggregate?

- **Problem:** systems yield many incorrect four-tuples
- **Previous work's solution:** order four-tuples by classifier confidence
- **Remaining Problems:**
 - Many incorrect four-tuples have high classifier confidence
 - In particular: many relations are classified as holding at the document creation time
 - It only takes one “bad apple” to derail the process and corrupt the final four-tuple.
- **Solutions:**
 - (1) Fix some common errors using rules
 - Use background knowledge to (2) eliminate or fix bad four-tuples and (3) add additional relevant four-tuples

(1) Post Relational States




- Verb Phrases use tense and aspect to help locate their predication time
 - tense is anaphoric (Higginbotham, 2006; Partee, 1984) and “binds” to a time
- NP’s lack these clues (in English).
 - We rely on context, world knowledge, reasoning, etc., to locate the properties NP’s denote.
- **An Exception:** modifiers like *former*, *then-*, and *ex-* introduce a *post-state* (Tonhauser, 2002)
- **Solution:** apply simple rules to correctly relate NP’s modified by *former*, *then-*, and *ex-* to the Document Creation Time

(2) Entity Existence Constraints



- **Common cause:** deceased person's relations are mentioned in a newspaper article. Systems treat these as true at the time of publication.
- **Solution:** Obtain **entity lifespan information** from Wikipedia and convert WITHIN labels to AFTER_END or BEFORE_START for out-of-lifespan time-expressions.



(2) Entity Existence Constraints



Document Creation Time (DCT): **2010-08-16**

The **London** home of composer George Frideric Handel is holding an exhibition about its other famous resident – **Jimi Hendrix**

Baseline

~~<resident(Hendrix, London), 2010-08-16> classified as WITHIN~~

Using Background Knowledge

Death

1970-09-18

[*lifespan*]

BEFORE

[D
C
T]

1947-11-27

1970-09-18

2010-08-16

<resident(Hendrix, London), 2010-08-16> classified as AFTER_END*

This process eliminates many four-tuples that are incorrect, yet have a high classifier confidence!

(3) Inter-Relation Dependencies

Input relation: *title(Chudinov, Prime Minister)*

Extract inter-relation dependencies from Wikipedia

title(Chudinov, Prime Minister)
BEFORE
title(Usenov, Prime Minister)



Hand code interval based reasoning rules for four-tuples

Retrieve documents relevant to dependent relations

Apply inference rules

“Prime Minister Daniar Usenov said at a press conference **Monday** that the government will uphold its current economic measure”



title(Chudinov, Prime Minister)
AFTER_END* @ Monday

Classification Results

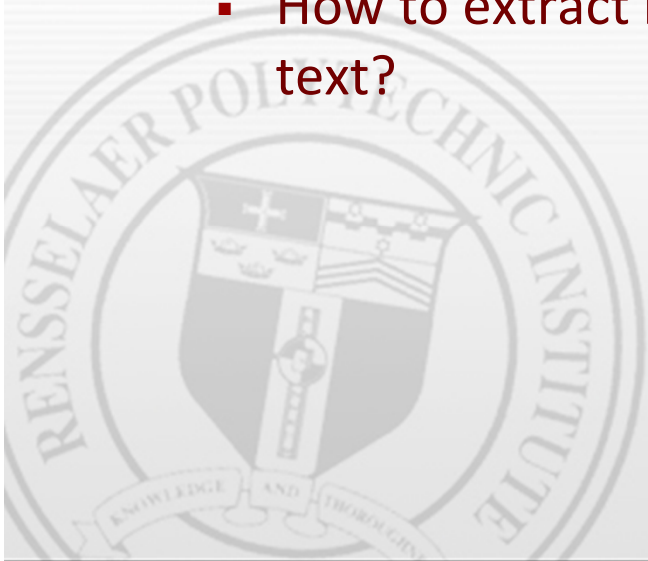
Settings	P	R	F
Baseline System	.337	.294	.314
Baseline + (1)	.341	.298	.318
Baseline + (1) + (2)	.353	.309	.329
Baseline + (1) + (2) + (3)	.360	.315	.336

- Rule based approach to handling temporal nominal modifiers outperforms statistical methods
- Background knowledge about entity existence helps eliminate inaccurate temporal information
- Background knowledge about relation ordering facilitates addition of mutually beneficial temporal information

Remaining Analysis & Challenges



- Title relation attributed via Noun Phrase
 - e.g., (former) President Smith
 - Under what circumstances does relation hold at DCT? At implicit and explicit predication times related to surrounding verbs?
- Relation ordering information
 - Under what circumstances does temporal information about one relation inform that of another?
 - How to extract relation ordering information automatically from text?



World Knowledge



“O’Donnell ... suggested **Wednesday** that the Obama administration - particularly **Vice President Joe Biden**, who **represented** Delaware in the Senate for **decades** - was behind them.”

- *Is title (Joe Biden, Vice President) true at*
 - *the predication time of *represented*?*
 - *... at Document Creation Time?*
- How do we know?
 - “Represent Delaware” not consistent with “Vice President” (of the US)

“Hidden *former*”



“Isakov is the latest of several top officials from the country’s caretaker government to resign so they can prepare for elections. Others include **First Deputy Prime Minister Almazbek Atambayev, ...**”

title(Atambayev, First Deputy Prime Minister) relation not true at DCT
It is clear that Atambayev already has resigned

In November 2000, **Chinese President Jiang Zemin** paid a state visit to Laos, the first visit to Laos by a Chinese president

title(Jiang Zemin, President) not true at DCT

This sentence is from a list of historical visits by Chinese Presidents to Laos

To what time does "*former*" apply?



Secretary of State Hillary Rodham Clinton says **former Philippines President Corazon Aquino** “**helped** bring democracy back” to her country after years of authoritarian rule”

Is *title(Corazon Aquino, President)* true at the predication time of *helped*?
...at Document Creation Time?

Former US President Bill Clinton and US journalists Euna Lee and Laura Ling **returned** Wednesday from North Korea, one day after North Korea’s leader Kim Jong-Il pardoned the two women

Is *title(Bill Clinton, President)* true at the time of *returned*?
... at Document Creation Time?

Future work: rigorously define elements of context that govern temporal interpretation of NP predicates

Extracting Relation Ordering From Text Rensselaer (This Work)

Honorable
Elena Kagan

11th Dean of Harvard Law School

Preceded by [Robert Clark](#)

Succeeded by [Martha Minow](#)



title(Minow Dean) and
title(Kagan, Dean)
are temporally linked

- In this work we assumed precedence relations were given



Extracting Relation Ordering From Text (**Future Work**)



“Martha Minow, who succeeded Kagan as Harvard Law School dean, ...”



title(Minow Dean) and
title(Kagan, Dean)
are temporally linked

Future work: Extract Precedence (and other temporal) relations from raw text



Thank You

