


InstaRead: An IDE for IE

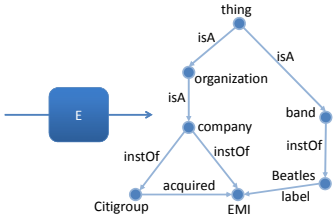
Short Presentation in CSE454

Raphael Hoffmann
January 17th, 2013

Computer Science & Engineering 


From Text To Knowledge


Citigroup has taken over EMI, the British music label of the Beatles and Radiohead, under a restructuring of its debt, EMI announced on Tuesday.



Human Effort

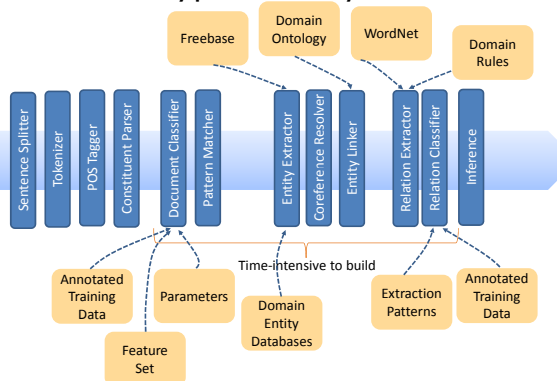
Evaluation time	#rels	#ann words	dev
ACE 2004 months	51	train 300K	12
MR IC 2010 months	15	train 115K	12
MR KBP 2011	16	8K	1.5hs

 team of annotators working in parallel

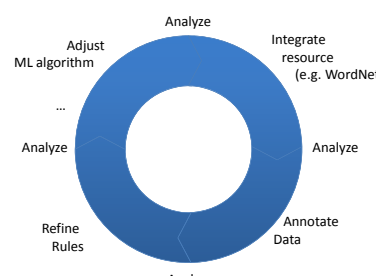
 teams of researchers working in parallel

Motivates research on learning with less supervision, but best-performing systems use more direct human input (eg. rules)

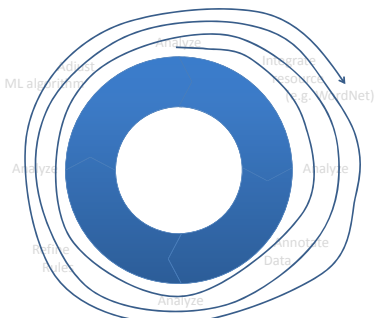
A Typical IE System

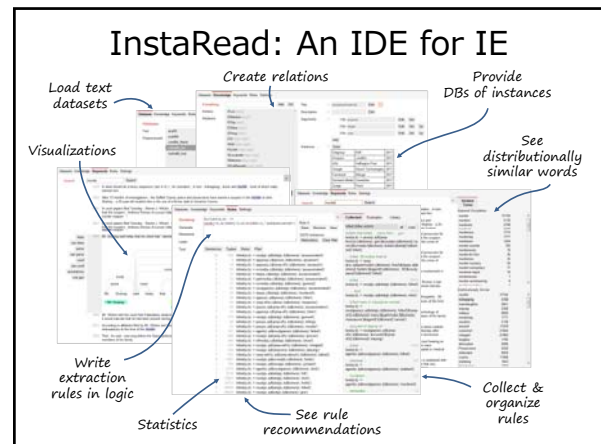
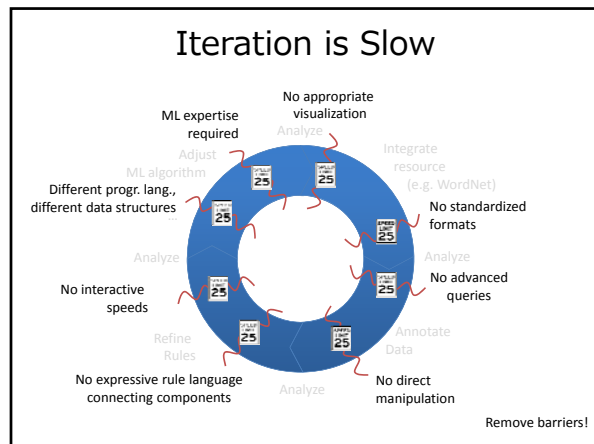


Development Cycle



Need Many Iterations





Key Ideas

1. User writes rules in simple, **expressive** language
Use First-Order Logic
2. User **instantly** sees extractions on lots of text
Use Database Indexing
3. User gets automatic rule **suggestions**
Use Bootstrapping, Learning

Demo

- Browser-based tool
- API
 - `cd /projects/pardosa/s1/raphaelh/github/readr/exp`
 - `source ../../init.sh`
 - `mvn exec:java -Dexec.mainClass=newexp.Materialize`

Project Ideas

In general: Create a new component & use InstaRead to explore, tie components together, and analyze

- Integrate exist. component (eg. Entity Linking, OpenIE)
- Mine rules from WordNet, FrameNet, VerbNet, Wiktionary
- Generate rule-suggestions through clustering
- Generate rule-suggestions through (better) bootstrapping
- Set rule (and thus extraction) confidence based on overlap with knowledge-base
- Develop rules/code to handle specific linguistic phenomenon (eg. time, modality, negation)
- Experiment with joint-inference of parsing + rules

More detailed slides (optional)


Rules are Crucial

For both hand-engineered and learned systems Goal: Enable experts to write quality rules *extremely quickly*

- Example

```

r(a, b) ← PER(a) ∧ smash(c, a)
           ∧ token(c, 'born')
           ∧ prep-in(a, b)
           ∧ LDC(b)
    
```



- Rules as patterns
- Rules as features
- Rules (or space of rules) typically supplied

Writing Rules in Logic

- FOL¹ is simple, expressive, extensible, widely used
- Introduce predicates for NER, dependencies, ...

```

PER ← smash born prep-in → LDC
    
```

```

r(a, b) ← PER(a) ∧ smash(c, a) ∧ token(c, 'born') ∧ prep-in(c, b) ∧ LDC(b)
    
```

- Rules are deterministic, execute in defined order

¹ We use the subset referred to as 'safe domain-relational calculus'

Rule Composition

- Define new predicates for similar substructure

```

killingsom('murder')
killingsom('assault/robbery')
killingsom('killing')
killingsom('wounding')

killingsom(c, b) ← prep-of(c, b) ∧ token(c, d) ∧ killingsom(d)
killingsom(c, b) ← m(c, b) ∧ token(c, d) ∧ killingsom(d)
killingsom(c, b) ← pass(c, b) ∧ token(c, d) ∧ killingsom(d)

killed(a, b) ← person(a) ∧ person(b) ∧ sub(pass(c, a) token(c, 'wounded') ∧
      prep-of(c, d) ∧ killingsom(d, b)
killed(a, b) ← person(a) ∧ person(b) ∧ prep-of(c, a) ∧ killingsom(c, b)
    
```

9 instead of 24 rules!

- More compact set of rules, better generalization

Enabling Instant Execution

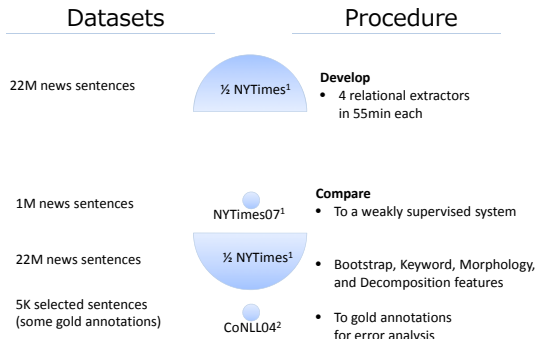
- Translation to SQL allows dynamic optimization

```

r(c) ← select sum(Low Harvey Oswald, a)
      person(a, p) ∧ smash(c, p) ∧
      token(c, b)
      ...
    
```

- Each predicate maps to a fragment of SQL
- Intensional and extensional predicates
- Indices, caching, SSDs important

Experimental Setup



¹ LDC2008T19, ² Roth & Yih, 2004

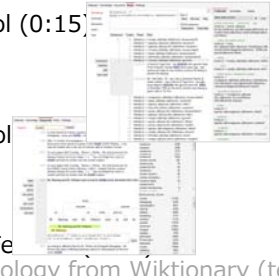
Comparison to Weakly Supervised Extraction

		attendedSchool	founded	killed	married
Rules	Precision	1.00	.91	.90	.90
	#extractions	1,411	997	189	4,694
Weakly supervised	Precision	0	.71	N/A	.50
	#extractions	5	14	N/A	2

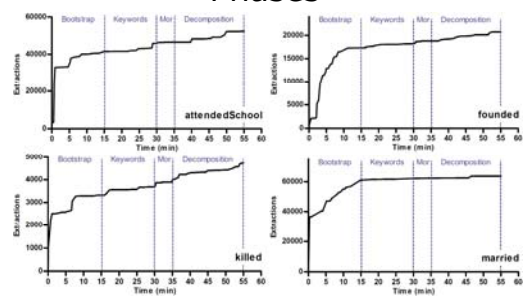
- Precision consistently at least 90%
- Works well, even when weak supervision fails

Development Phases

1. Bootstrap Tool (0:15)
2. Keywords Tool (0:15)
3. Morphology Fe
mined morphology from Wiktionary (tense etc.)
4. Decomposition Feature (0:20)
enable chaining of rules



Comparison of Development Phases



- Bootstrap initially effective, but recall limited
- Decomposition effective for some relations

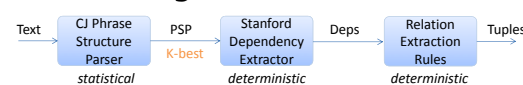
Error Analysis

- On CoNLL04 'killed' wrt gold annotations: Re .34, Pr .98

False Negatives due to		False Positives due to	
Preprocessing (missing predictions)		Preprocessing (wrong predictions)	
NER	30	NER	0
Dependencies	25	Dependencies	2
Co-references	7	Co-references	0
Rules (missing predictions)		Rules (wrong predictions)	
Lexical items	89	Lexical items	0
Syntactic variation	17	Syntactic variation	0
Reasoning chain	10	Reasoning chain	0

- 36% of all errors are due to incorrect pre-processing

Joint Parsing & Relation Extraction



- If top parse is wrong ...
 - In 35% of cases: correct at top-2
 - In 50% of cases: correct among top-5
 - In 90% of cases: correct among top-50
- Heuristic: Select first parse among top-k with most extractions

	Recall	Precision	F1
Pipeline	.342	.978	.507
Joint	.412	.973	.58

18/19 flipped to correct parse
Almost no drop in precision

Runtime Performance

	avg	median	max
#SQL queries per relation (55min)	55	54	66
#join tables per SQL query	4.3	4	11
Execution time (s) per SQL query	1.5	.074	52.5

- On 22M sentences, 3.7B rows in 75 tables, 140GB
- Most queries execute in 74ms or less
- Outliers due to Bootstrap tool
 - Aggregation over millions of rows motivates streaming or sampling

Logic

- Often Horn clauses are sufficient, but sometimes we need \neg, \forall, \exists
- Example 1: *founded* relation

Michael Dell built his first company in a dorm-room.

Mr. Harris built Dell into a formidable competitor to IBM.

$\neg(\text{hd} : \text{prop-into}(c, d))$

Desired conjunct

- Example 2: Integration of co-reference

$\neg(\text{hd} : \dots)$

"nearest noun-phrase which satisfies ..."