KG⇔ML

Sixty years of knowledge graphs for language understanding

Tim Finin

University of Maryland, Baltimore County

KR2ML Workshop, 2019-09-19

Joint work with many colleagues and students

UMBC

http://bit.ly/kr2mlFinin

Overview

- How we got here
- Where are we anyway?
- Where are we going?
- Some recent work on
    - Using *knowledge graphs* to improve *machine learning*
    - Using *machine learning* to improve *knowledge graphs*

How we got here

"Knowledge graphs" of one kind or another have been used for more than 60 years for AI tasks, especially those involving language understanding

How we got here

An early example from 1955 representing ”dog bites cat”

```
       1           2          1           2
dog --> part of <-- teeth --> contact <-- cat
                                 ^
                                 |
                                much
```

Richens, R.H. and Booth, A.D. ‘Some methods of mechanised translation’, in Locke, W.N. and Booth, A.D. (Eds.) Machine translation of languages, pp. 24-46, 1955.

How we got here

Example from 1967 representing the food concept

FOOD: 1. That which living being has to take in to keep it living and for growth.
Things forming meals, especially other than drink

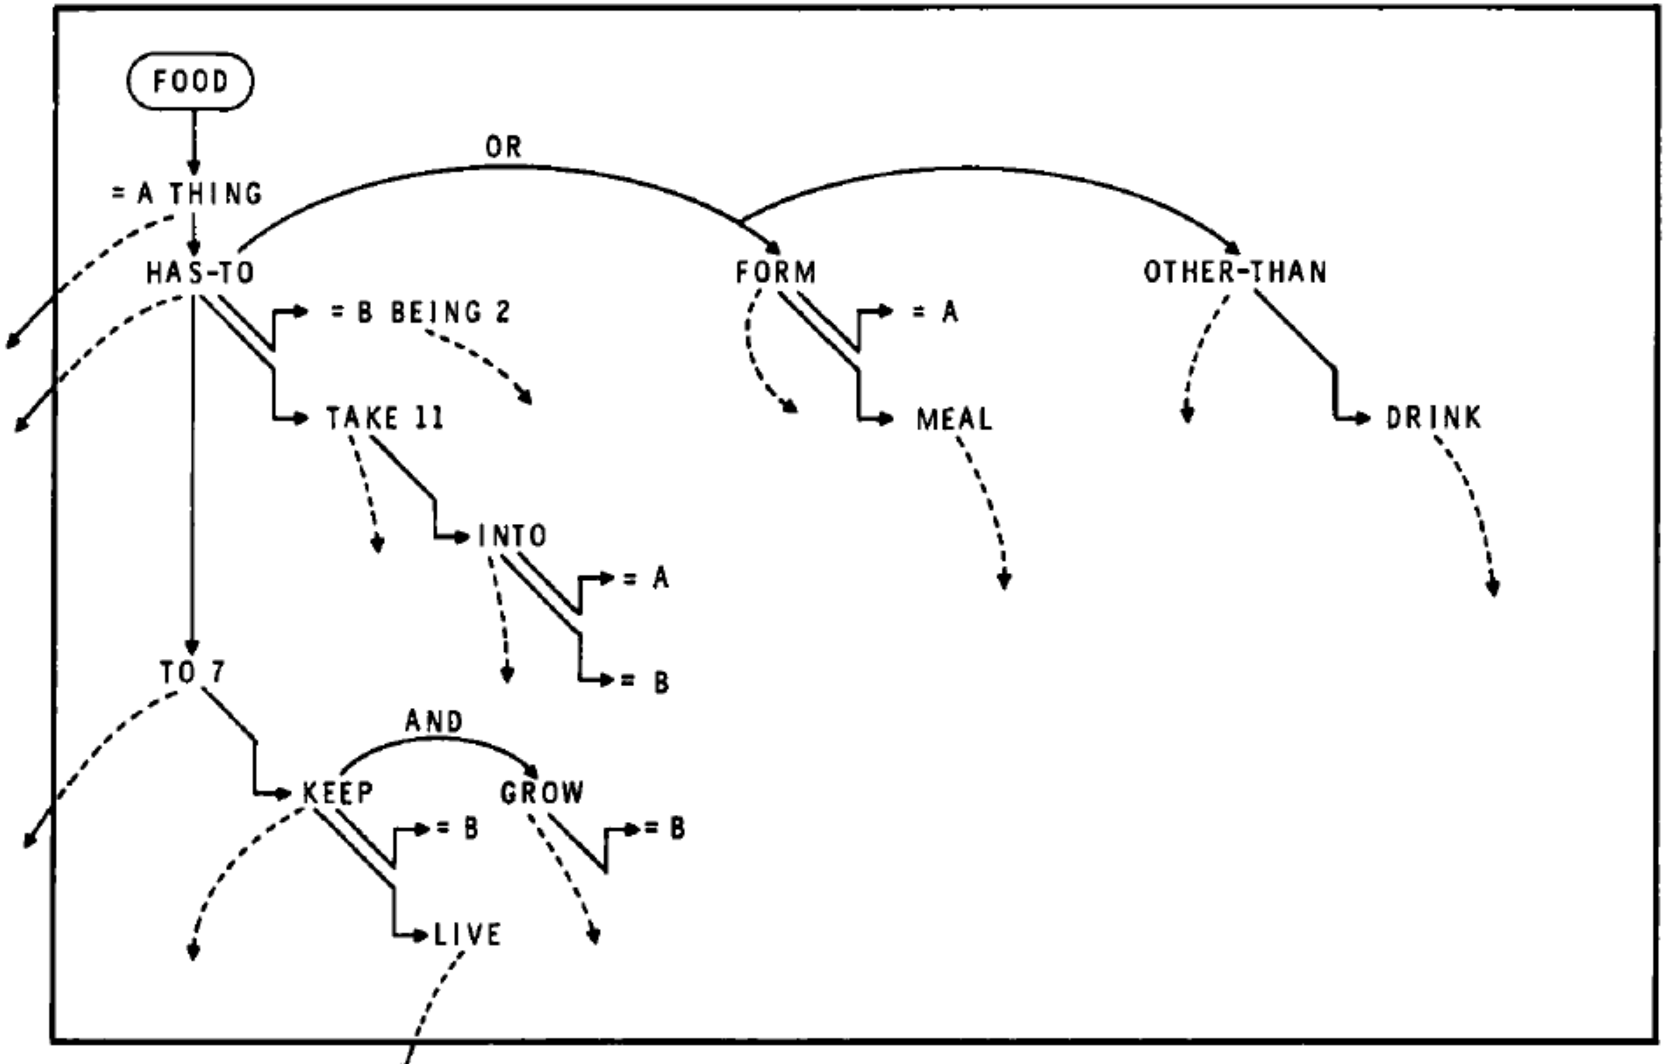

R. Quillian, Word concepts: A theory and simulation of some basic semantic capabilities, Behavioral Science, 12(5), 1967.

How we got here

I recall in the early 1970s using Lisp's symbol plists with *get, putprop* and *remprop* to create what we'd now call a *property graph* in support of AI problems

How we got here

Over the decades much important and useful knowledge representation work has been done in support of AI

These are just a few familiar examples

- Micro-planner
- Semantic networks + logic
- Minsky Frames
- Schank Scripts
- Object oriented systems
- What’s in a link?
- Logic programming
- KL-ONE
- Production systems
- Description Logic
- CYC
- Semantic Web
- OWL
- Linked Data
- Wikidata

Where are we, anyway?

After decades of slow but steady advances, know-ledge representation and AI have experienced a sea change that we can attribute to a convergence of multiple factors:

- Ubiquitous Internet and Web for sharing and accessing information
- Data availability, now that it’s easy to share
- Increased computing power
- Machine learning advances

How relevant are knowledge graphs today?

Where are we going?

KG ⇔ ML

It's hard to make predictions, especially about the future; for now, we're busy exploring two scenarios:

KG⇨ML: How can existing KGs support ML applications

- **#1** Wikidata gazetteers for Named Entity Recognition
- **#2** Better topic models with ontologies

KG⇦ML: How can ML enrich and improve KGs

- **#3** Inferring relations in KGs
- **#4** Detecting and repairing KG errors

Label Sparsity & Inconsistency in NER

#1

KG⇨ML

- Scale was a ten-week summer project at the Human Language Technology Center of Excellence at JHU (2019)
- Explored improving a TensorFlow NER system (Bi-LSTM-CRF & BERT) in English, Chinese & Russian
- Four core types (PER, ORG, GPE and LOC) and additional 12 finer grained types (e.g., government org, commercial org)
- Limited training data from Ontonotes and local annotations
- Exploited the Wikidata KG in several ways

Wikidata Knowledge Graph

- **Large knowledge graph** with ~700M statements about ~57M items
- **Fine-grained ontology**: ~2M types; ~5K properties
- **Multilingual**, strings tagged with language id
- Links to all of entity's **Wikimedia pages**
- Entities have a canonical **name** and **aliases** in one or more languages and multiple claims
- COE=Q64780099, with type *research institute*, name *Human Language Center of Excellence,* alias *HLTCOE,* and 12 properties

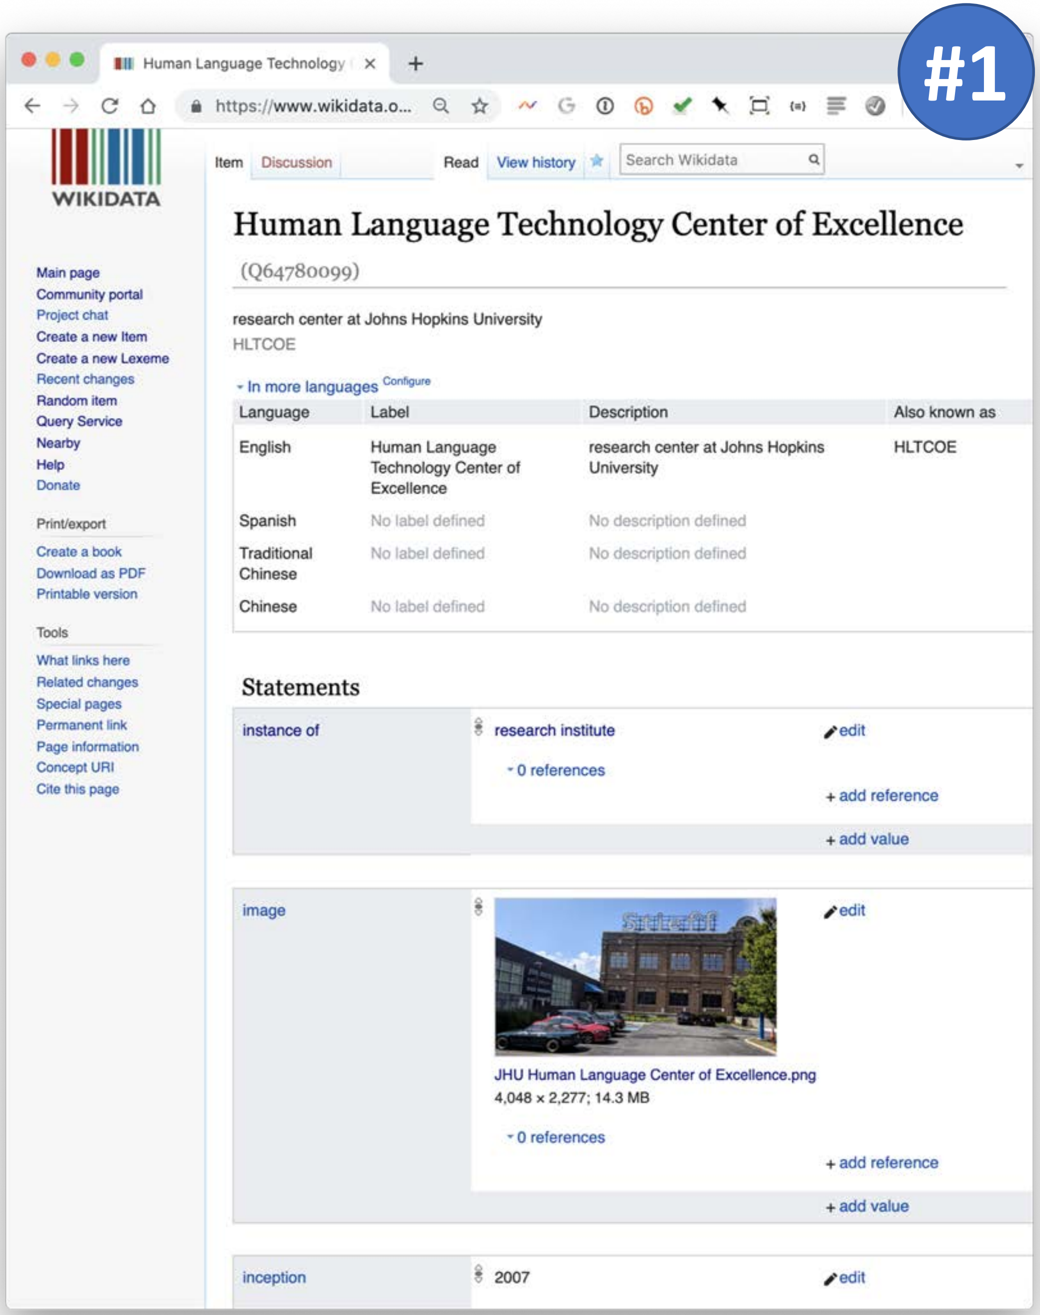

Q64780099

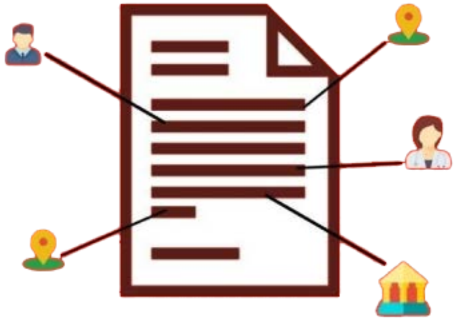

Three Use Cases Involving KGs

#1

#0 Preliminary: align NER types with Wikidata's huge type system

#1 Extract new training data from Wikidata & Wikipedia

#2 Create name gazetteers for each Scale type from Wikidata and use to construct new training data

#3 Use tokens in name gazetteers as TensorFlow features

#1 WD types ⇔ Scale types

- Map Scale types to Wikidata's huge type system
- Some are simple:
  PER = wd:Q5 (human)
  AIR = wd:Q62447 (aerodrome)
- Others more complex:
  COMP = wd:Q3966 (computer_hardware) + wd:Q68 (computer) + wd:Q7397 (software)
  ORG = wd:Q163740 (nonprofit_organization) + 15 more
- Query Wikidata to collect names and aliases in English, Chinese and Russian for each type
- Filter/trim Wikipedia artifacts via regexs

```
aerodrome (Q62447) •34 ×11745 ↑↑↑↑
├──heliport (Q502074) •37 ×5120
├──military airbase (Q695850) •35 ×1245 ↑
│   ├──Royal Air Force (Q165862) •62 ×14
│   ├──??? (Q1428515) •3 ×2
│   ├──naval air station (Q6981985) •1 ×26 ↑
│   │   └──naval air auxiliary station (Q21601954) ×2
│   ├──Airbase 90 (Q10426742) •3
│   └──Airbase 60 (Q10426743) •3
├──highway strip (Q782667) •13 ×9 ↑
├──United States Coast Guard Air Stations (Q1104354) •4 ×1
├──airport (Q1248784) •130 ×24919 ↑
│   ├──alternate airport (Q392406) •8
│   ├──international airport (Q644371) •30 ×1206
│   ├──domestic airport (Q837800) •13 ×19
│   ├──special airport (Q1479818) •1 ×2
│   ├──regional airport (Q2138424) •6 ×2
│   ├──Airport of entry (Q4698971) •2
│   ├──airline hub (Q4811588) •36
│   ├──Executive airport (Q5419792) •1
│   ├──Relief airport (Q7311088) •2
│   ├──departure airport (Q15733670) ↑
│   ├──destination airport (Q15733672) ↑
│   ├──Non-towered airport (Q17144062) •2
│   ├──commercial airport (Q20977786)
│   ├──abandoned airport (Q20992031) ×1
│   ├──airport (Q21836433) ×15
│   ├──proposed airport/being built (Q44665966) ×91 ↑
│   └──general aviation airport (Q62782337)
├──mountain landing site (Q1497053) •3
├──glider airfield (Q2265915) •1 ×280
├──special airfield (Q2301048) •3 ×295
├──commercial airfield (Q2516330) •2 ×154
├──altiport (Q2840449) •6 ×8
├──seaplane base (Q3143713) •6 ×514
├──airstrip (Q3631092) •2 ×337
│   └──??? (Q1402469) •1 ×1
├──aéroclub (Q4438156) •4 ×3
├──Blue ice runway (Q4930096) •1 ×5
├──Floating airport (Q11558697) •2 ×1
└──balloonport (Q59388193) ×13
```

#1

#1 What type is Q1538727?

- Use case #1 requires finding an entity's Scale type, but wiki dumps only have immediate types
  - Q1538727 (**University of Angers**) is a Q3551775 (**University in France**)
- Which Scale type should it be, given it could be any of ~2M Wikidata types?
- SPARQL query on Wikidata for each Scale type for subtypes with ≥1 immediate instances (~20k total)
- Use result to make a compact & fast map from any WD entity to a Scale type
- Q1538727 is a [ORG, LOC]

```
university in France (Q3551775)
├──university (Q3918)
│   ├──higher education institution (Q38723)
│   │   └──educational institution (Q2385804)
│   │       ├──juridical person (Q155076)
│   │       │   ├──organization (Q43229)
│   │       │   │   ├──group of humans (Q16334295)
│   │       │   │   │   └──group of living things (Q16334298)
│   │       │   │   │       ├──structure (Q517966)
│   │       │   │   │       └──group of physical objects (Q61961344)
│   │       │   │   │           └──group (Q16887380)
│   │       │   │   │               └──object (Q488383)
│   │       │   │   │                   └──entity (Q35120)
│   │       │   │   └──agent (Q24229398)
│   │       │   │       └──individual (Q23958946)
│   │       │   │           ╘══entity (Q35120)
│   │       │   └──legal person (Q3778211)
│   │       │       ├──subject (Q830077)
│   │       │       │   ╘══entity (Q35120)
│   │       │       ╞══agent (Q24229398)
│   │       │       └──subject of a right (Q26720107)
│   │       │           ╞══subject (Q830077)
│   │       │           ╘══agent (Q24229398)
│   │       ├──institution (Q178706)
│   │       │   ╘══organization (Q43229)
│   │       ├──educational organization (Q5341295)
│   │       │   ╘══organization (Q43229)
│   │       └──facility (Q13226383)
│   │           ├──geographic location (Q2221906)
│   │           │   └──location (Q17334923)
│   │           │       └──geographic entity (Q27096213)
│   │           │           └──spatial entity (Q58416391)
│   │           │               └──spacio-temporal entity (Q58415929)
│   │           │                   ╘══entity (Q35120)
│   │           └──artificial geographic entity (Q27096235)
│   │               ├──artificial physical object (Q8205328)
│   │               │   ├──physical object (Q223557)
│   │               │   │   └──concrete object (Q4406616)
│   │               │   │       ╘══spacio-temporal entity (Q58415929)
│   │               │   └──artificial entity (Q16686448)
│   │               │       ╘══entity (Q35120)
│   │               ╘══geographic entity (Q27096213)
│   └──academic institution (Q4671277)
│       ╘══educational institution (Q2385804)
└──public scientific, cultural or professional establishment (Q3591586)
    └──public university (Q875538)
        ╞══university (Q3918)
        ╞══academic institution (Q4671277)
        └──public educational institution (Q23002037)
            ├──public institution (Q294163)
            │   ╞══juridical person (Q155076)
            │   ╘══institution (Q178706)
            ╘══educational institution (Q2385804)
```

#1

Use Case #2

#1

Replace names in existing NER training data with random names from appropriate **gazetteer** to get new training data

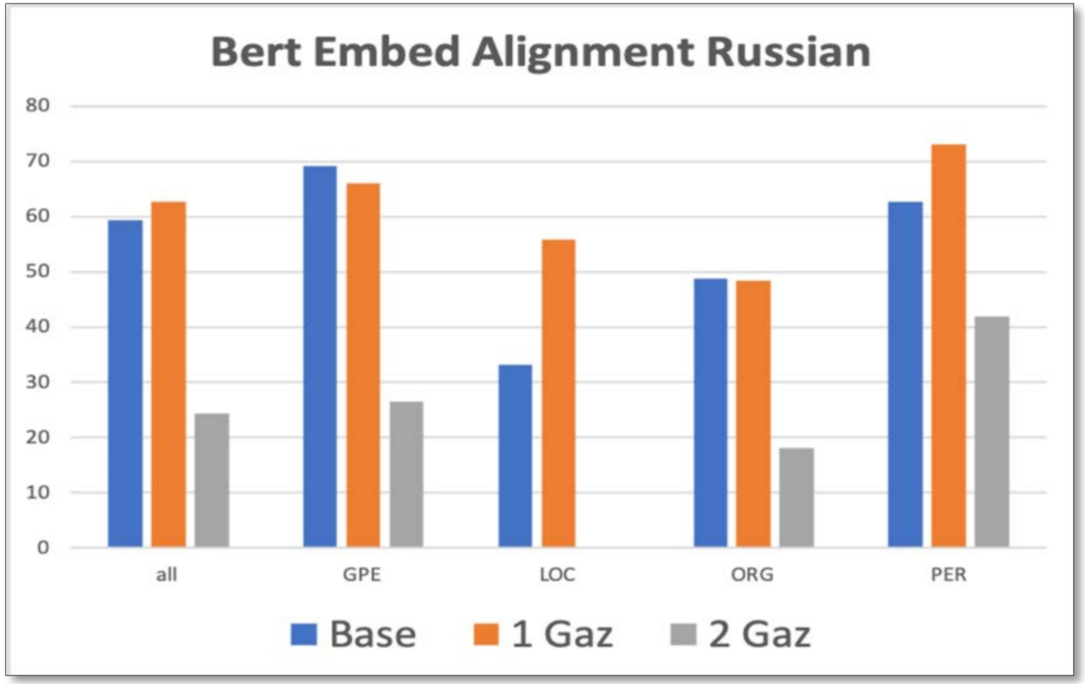

Shows promise, but needs some refinement for simple type systems

However	O
,	O
since	O
Disney	B-ORG
entered	O
Hong	B-GPE
Kong	I-GPE
,	O
the	B-ORG
Ocean	I-ORG
Park	I-ORG
,	O
sharing	O
the	O
same	O
city	O
as	O
Disney	B-ORG
,	O
has	O
felt	O
the	O
pressure	O
of	O
competition	O
.	

Original annotated training sentence

However	O
,	O
since	O
Colby	B-ORG
Village	I-ORG
Elementary	I-ORG
School	I-ORG
entered	O
Baltimore	B-GPE
County	I-GPE
,	O
Keds	B-ORG
Public	I-ORG
School	I-ORG
,	O
sharing	O
the	O
same	O
city	O
as	O
Faculty	B-ORG
of	I-ORG
Management,	I-ORG
University	I-ORG
of	I-ORG
Warsaw	I-ORG
,	O
has	O
felt	O
the	O
pressure	O
of	O
competition	O
.	O

New gazetteer training sentence

#1

Use Case #3

- Use tokens from gazetteers to add features for potential Scale types
- Produced modest but significant F1 improvements for English and Chinese
- Mixed results for Russian due to small gazetteer sizes

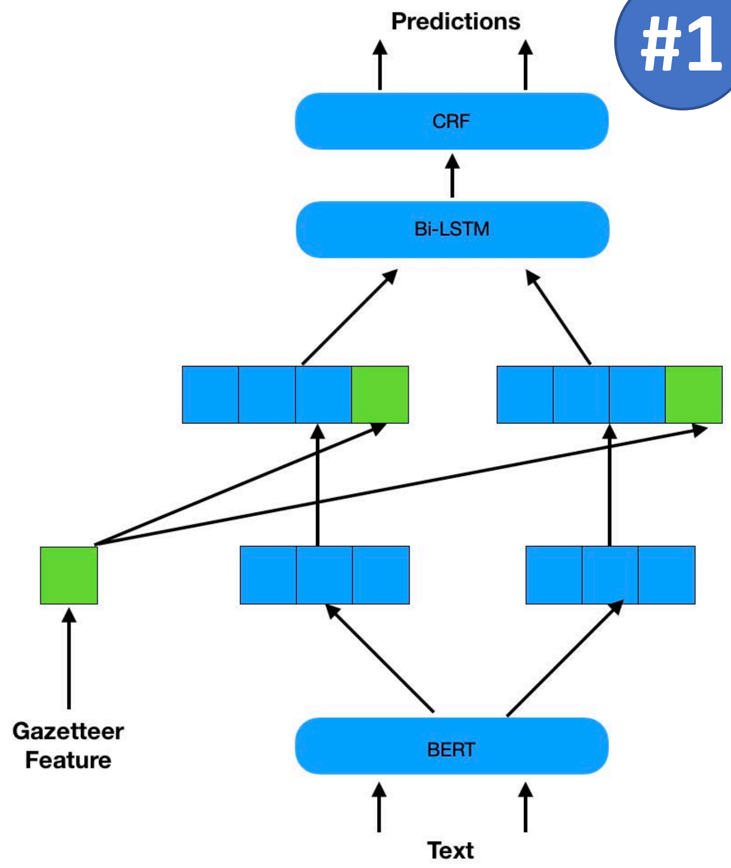

TF Bi-LSTM-CRF network w. BERT

Text	Jack	is	on	Hong	Kong	International	Airport	in	Lantau	Island	,	Hong	Kong
PER	B-PER	O	O	O	O	O	O	O	O	O	O	O	O
LOC	O	O	O	O	O	O	O	O	B-LOC	I-LOC	O	O	O
GPE	O	O	O	B-GPE	I-GPE	O	O	O	O	O	O	B-GPE	I-GPE
ORG	O	O	O	O	O	I-ORG	I-ORG	O	O	O	O	O	O

Ontology Grounded Topic Models

#2

KG⇨ML

- **Topic models** learn to represent a document's topics as a real-number vector, typically of length 100-300
- Often hard to know what each dimension/topic means
- Jennifer Sleeman developed techniques to enrich topic models with domain ontology concepts, and…
- Showed on two domains, climate science& cyber-security, that the technique:
  - **Improves model quality** using a standard perplexity metric
  - Makes models **easier to understand** by domain-experts

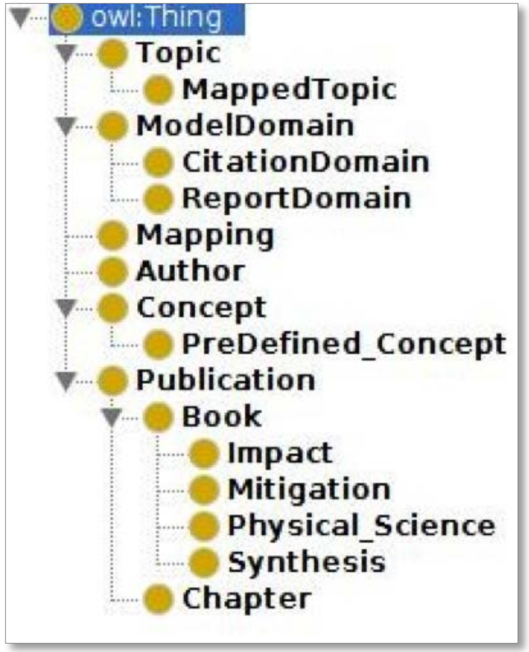

Scientific Influence in Climate Change Research

- Understand how scientific disciplines evolve & predict future directions by identifying and relating research and their influence
- Applied **cross-domain, dynamic topic modelling** jointly on (1) 25 years of IPCC Climate Change Reports and (2) the 200,000 papers they cite
- Explainable results via alignment with domain ontologies

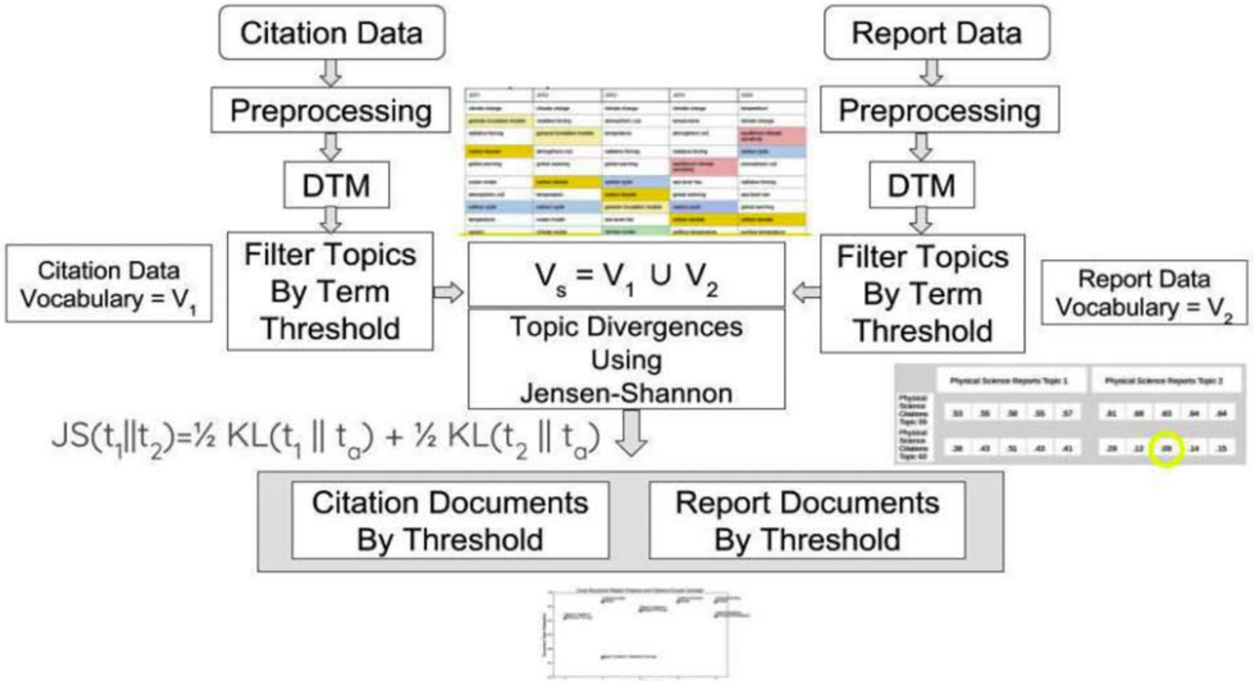

AR1	AR2	AR3	AR4	AR5
climate change	climate change	climate change	climate change	temperature
general circulation models	radiative forcing	atmospheric co2	temperature	climate change
radiative forcing	general circulation models	temperature	atmospheric co2	equilibrium climate sensitivity
carbon dioxide	atmospheric co2	radiative forcing	radiative forcing	carbon cycle
global warming	global warming	global warming	equilibrium climate sensitivity	atmospheric co2
ocean model	carbon dioxide	carbon cycle	sea level rise	radiative forcing
atmospheric co2	temperature	carbon dioxide	global warming	sea level rise
carbon cycle	carbon cycle	general circulation models	carbon cycle	global warming
temperature	ocean model	sea level rise	carbon dioxide	carbon dioxide
system	climate model	climate model	surface temperature	surface temperature

The dynamic topic model shows how significance of topics changes from AR1 report (1990) to AR5 report (2014)

Recognizing concept phrases

- Technical text includes key concepts phrases
- Hard to capture automatically using SOTA phrase extraction tools
- Solution: collect from glossaries & ontologies, e.g., Wikidata concepts
- Recognize instances in text & include these as extra tokens for topic model

"An intercomparison is undertaken of the tropical behavior of 17 coupled ocean-atmosphere models in which at least one component may be termed a general circulation model (GCM). The aim is to provide a taxonomy--a description and rough classification-of behavior across the ensemble of models, focusing on interannual variability. The temporal behavior of the sea surface temperature (SST) field along the equator is presented for each model, SST being chosen as the primary variable for intercomparison due to its crucial role in mediating the coupling and because it is a sensitive indicator of climate drift." -- Text from [Neelin, J. D., et al. "Tropical air-sea interaction in general circulation models." Climate Dynamics 7.2 (1992): 73-104.]

World Meteorological Organization
General Circulation Models
the atmosphere that
in the Earth
climate system
Assessment Report Working

Better

Ultraviolet Radiation
Forcing Mechanism
Fossil Fuel
Water Vapor
General Circulation Model
Fluorinated Gases

Sleeman, Halem & Finin, Ontology-Grounded Topic Modeling for Climate Science Research, Semantic Web for Social Good, ISWC, 2018

Adding concept phrases helps topic explainability

Word-based Topics	Concept-based Topics
change, ocean, level, global, model, mean, climate, figure, rise, surface,	temperature, anthropogenic, carbon dioxide, radiative forcing, sea level rise, greenhouse gases, snow, surface temperature, wind, global warming potential,
carbon, climate, change, emission, atmospheric, ocean, model, university, global, land,	carbon dioxide, carbon cycle, atmospheric co2, anthropogenic, temperature, land use, methane, fossil fuel, ppm, surface temperature,

Examples of top terms in two topics, with and without concepts

AR1	AR2	AR3	AR4	AR5
climate change	climate change	climate change	climate change	temperature
general circulation models	radiative forcing	atmospheric co2	temperature	climate change
radiative forcing	general circulation models	temperature	atmospheric co2	equilibrium climate sensitivity
carbon dioxide	atmospheric co2	radiative forcing	radiative forcing	carbon cycle
global warming	global warming	global warming	equilibrium climate sensitivity	atmospheric co2
ocean model	carbon dioxide	carbon cycle	sea level rise	radiative forcing
atmospheric co2	temperature	carbon dioxide	global warming	sea level rise
carbon cycle	carbon cycle	general circulation models	carbon cycle	global warming
temperature	ocean model	sea level rise	carbon dioxide	carbon dioxide
system	climate model	climate model	surface temperature	surface temperature

Our dynamic topic model shows changes in the significance of topics over 25 years, from the AR1 report to the AR5 report

Cybersecurity DTM

- Applied to cybersecurity documents

  16K Symantec malware reports (2000-2016), 4K arXiv papers on cryptology and security (1997-2017)
- Using 3,836 concept phrases extracted from Wikidata/Wikipedia
- Linked phrases (Denial of Service) to their acronyms (DOS)
- Shows influence between research (arXiv) and practice (Symantec)

With concept phrases	W/O concept phrases
quantum cryptography, phase, photon, cryptography, measurement, channel, system, eavesdropping, stage, polarization	quantum, state, communication, phase, cryptography, channel, eavesdropping, protocol, error, polarization
intrusion detection, universal, taxonomy, **intrusion detection system**, based, payload, classification, input, attack, alert	cell, network, intrusion, parameter, system, information, detection, method, space, approach

Examples of top terms in two topics with and without concepts

J. Sleeman, T. Finin, M.Halem and R. Bordawekar, Temporal Understanding of Cybersecurity Threats, under review, 2019.

Knowledge-Enriched Tensor Factorization

#3

KG⇦ML

- We developed a system to **predict likely relations** in a KG that outperforms RESCAL and similar systems
- Joint work with Ankur Padia, Frank Ferraro & Kostas Kalpakis
- Identifies relations *believed to hold* (**link verification**) rather than a *ranked list* of possible relations **(link ranking)**
- Computes & uses **relational similarity matrices** as features
- Evaluated on data from existing knowledge graphs, e.g., DBpedia and Freebase

Multi-relational Data as a Knowledge Graph

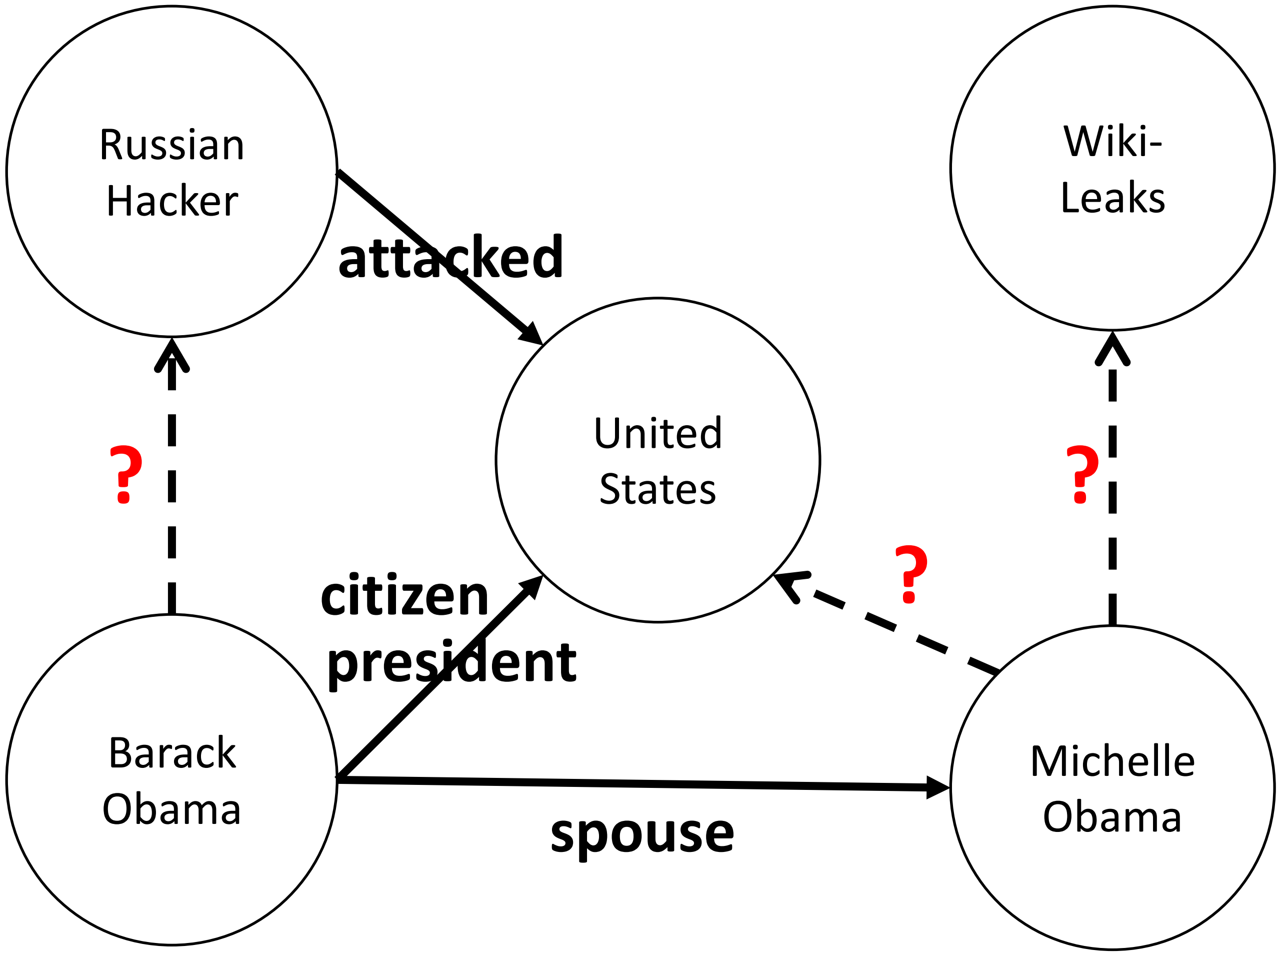

Belief 1: Russian_Hacker attacked United_States .

Belief 2: Barack_Obama Spouse Michelle_Obama .

Belief 3: Barack_Obama President United_States .

Belief 4: Barack_Obama citizen United_States .

Belief 5: Wikileaks isa non-profit_organization .

Learning Embeddings for Entities and Relations

Jointly learn Entity (E) and Relation (R) embeddings

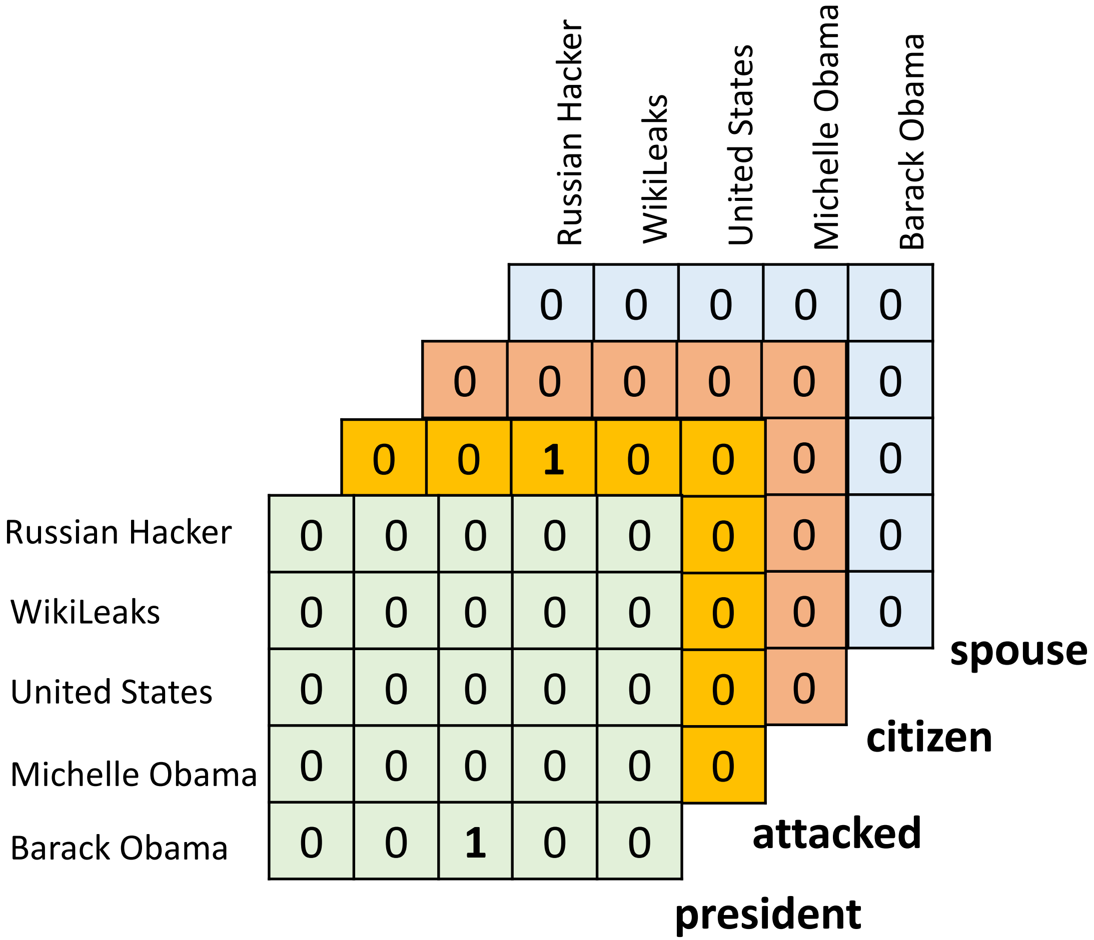

$$\chi \approx ERE^T$$

- χ – Data tensor of size *e x e x k*
- *E* – Shared entity matrix of size *e x p*
- *R* – Compact relation tensor of size *p x p x k*
- *p* – Latent dimension
- k – Number of relations
- e – Number of entities

Prior Information as Relation Similarity

minimize
E, R

$$||\chi - ERE^T||^2 + \lambda_a||\mathrm{E}||^2 + \sum_1^k \lambda_r ||R_k||^2$$

\+ Use prior information

$$C_{ij}||R_i - R_j||^2$$

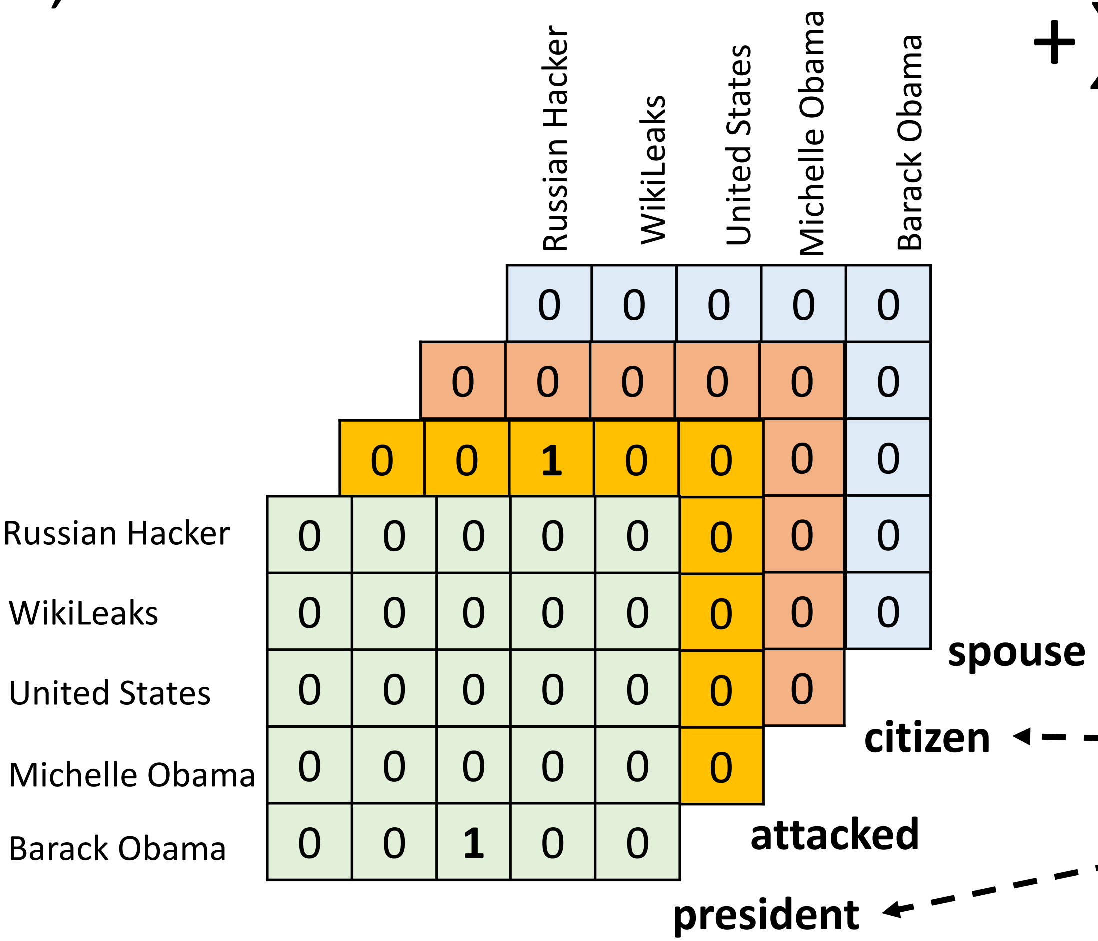

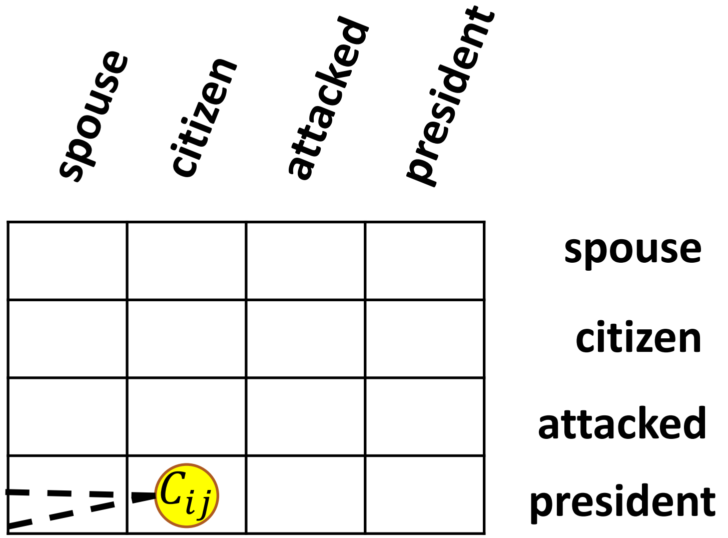

Prior Information as Relation Similarity

minimize E, R

$$||\chi - ERE^T||^2 + \lambda_a||\mathrm{E}||^2 + \sum_1^k \lambda_r ||R_k||^2$$

+ Use prior information

$$C_{ij}||R_i - R_j||^2$$

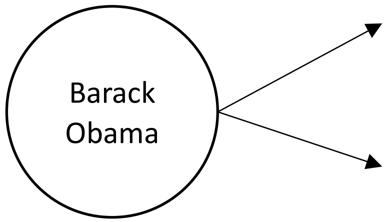

Agency

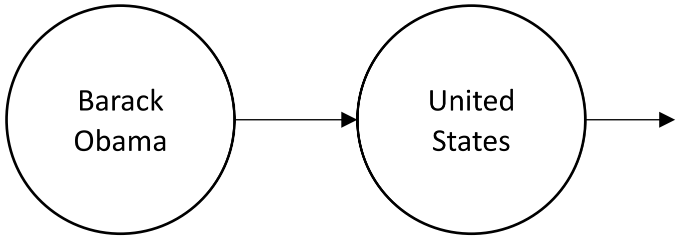

Transitivity

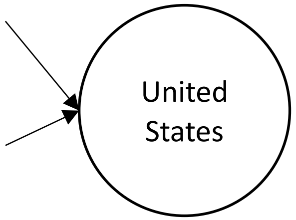

Patient

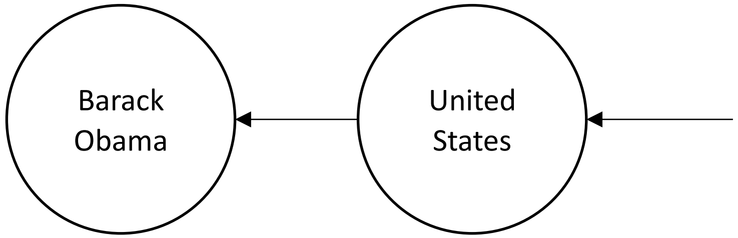

Reverse Transitivity

	spouse	citizen	attacked	president
spouse				
citizen				
attacked				
president		C_{ij}		

Relation similarity

Evaluated five metrics and found transitivity best

Symmetry: frequency relations have same entities as either subject or object

Agency: frequency relations have same subject

Patient: frequency relations have same object

Transitivity: frequency object of relation_1 is the subject of relation_2

Reverse Transitivity: frequency subject of relation_1 is the object of relation_2

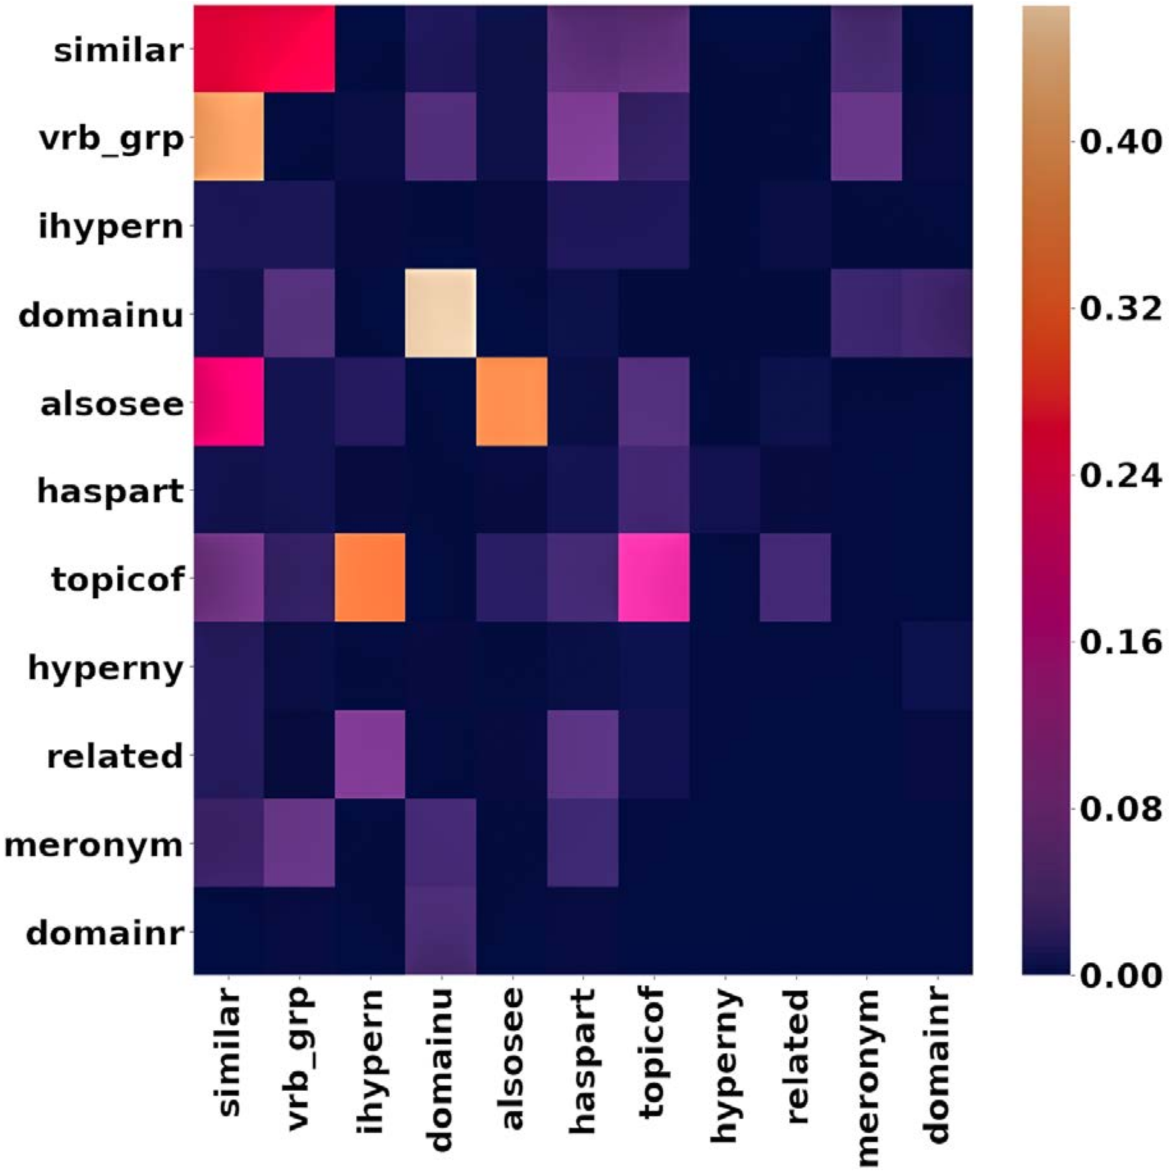

Heatmap for WIN18RR's similarity matrix using transitivity (relation names abbreviated)

Evaluation on eight datasets

Dataset	Domain	Entities /Nodes	Relations /Edges	Facts	Avg. Deg	Graph Density (= facts/entities^2)
Kinship	Social	104	26	10.7K	102.75	0.98798
UMLS	Medical	135	49	6.8K	50.01	0.37048
FB15-237	General	14.5K	237	310.1K	21.32	0.00147
DB10k	General	4.3K	140	10.0K	2.27	0.00052
FrameNet	Language	22.3K	16	62.3K	2.79	0.00013
WN18	Language	40.9K	18	151.4K	3.7	0.00009
FB13	General	81.1K	13	360.5K	4.45	0.00005
WN18RR	Language	40.9K	11	93.0K	2.27	0.00005

- Comparison with state-of-the-art tensor factorization methods and translation-based models
- Used Precision-Recall AUC evaluation metric
- Note that Kinship and UMLS are outliers w.r.t. graph density

Results: Link Verification Area Under Curve

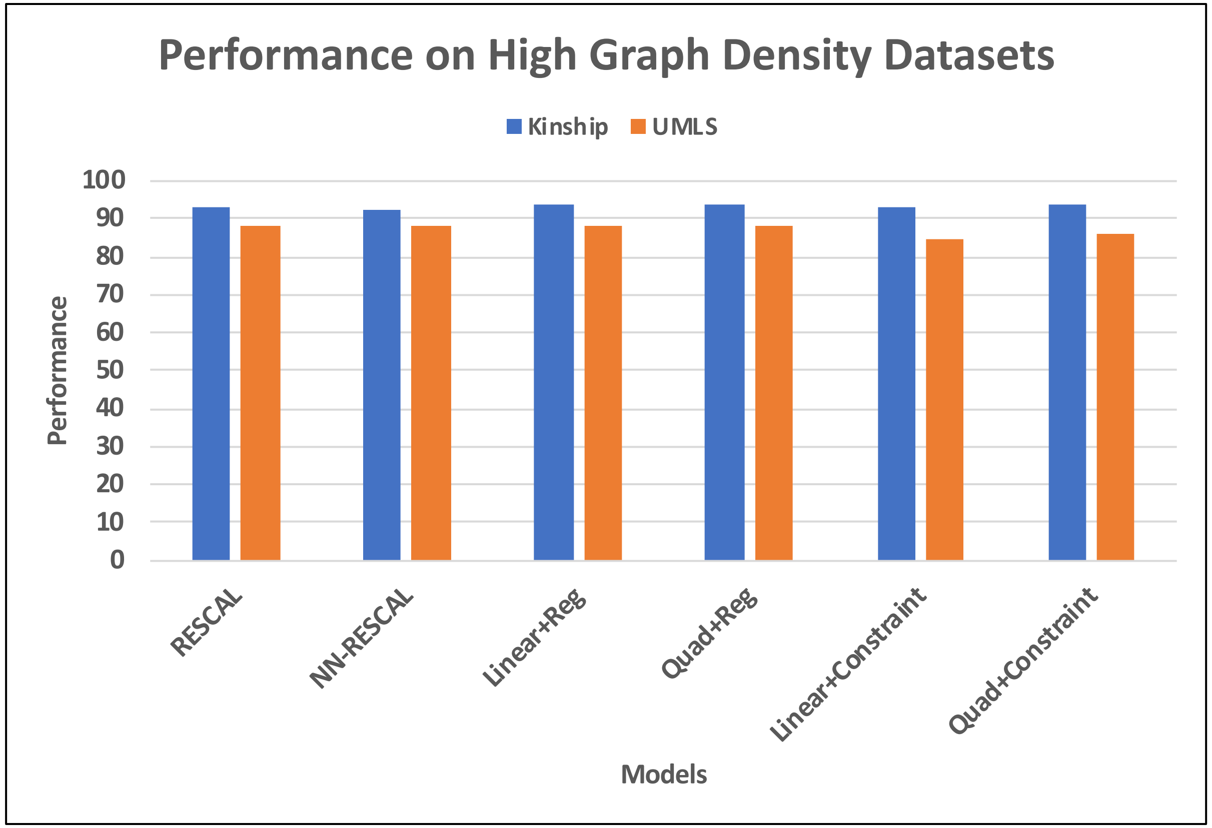

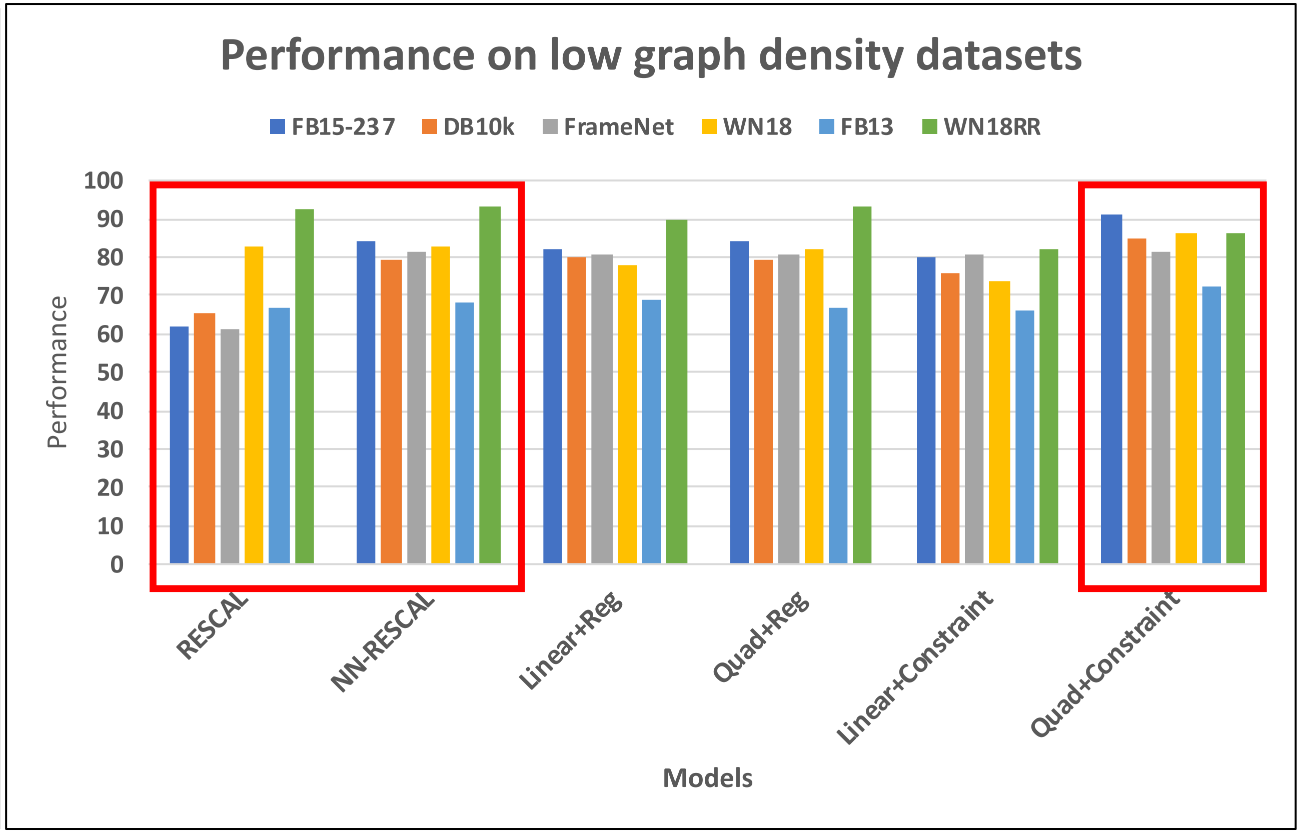

We evaluated four models and Quadratic+Constrained is best overall (statistically significant) for typical graphs with low density (e.g., Freebase, DBpedia, FrameNet)

Padia, Kalpakis, Ferraro and Finin, [Knowledge Graph Fact Prediction via Knowledge-Enriched Tensor Factorization](), Journal of Web Semantics, 2019.

Identifying & fixing IE errors

#4

KG⇦ML

- Dissertation work by Ankur Padia
- Many IE systems produce knowledge graph triples from text with provenance text, as in the TAC Knowledge Base Population tasks
- But error rate is high, e.g. F1 < 0.33
- We developed an independent system to
  - Identify triples that seem **inconsistent** w.r.t. their provenance, and
  - Jointly attempt to repair, if inconsistent
- Consistency vs. credibility

Belief learned by IE system:
`per:charges(Harry Reid, assault)`

Provenance identified by IE system:
Nevada's Harry Reid switches longtime stance to support assault weapon ban

Analysis output:

Is reading consistent:	Inconsistent
Suggested relation:	no repair

(a) An inconsistent reading with no correction.

Belief learned by IE system:
`per:cause_of_death(Edward Hardman, Typhoid fever)`

Provenance identified by IE system:
The Western Australian government agreed to offer the Government Geologist post to Hardman shortly before news of his death reached them . Early in April , he contracted typhoid fever , and died a few days later in a Dublin hospital on 6 April

Analysis output:

Is reading consistent:	Consistent
Suggested relation:	per:cause_of_death

(b) A consistent reading not requiring a correction. Notice the relation is unchanged.

Human assessment of TAC IE Errors

Category	Definition
Incorrect Relation	Relation not triggered or entailed
Subject missing	Entity is not mentioned in provenance
Object missing	Entity is not mentioned in provenance
Misc./ Format / Guidelines	Fact does not adhere to schema-specific guidelines and requirement

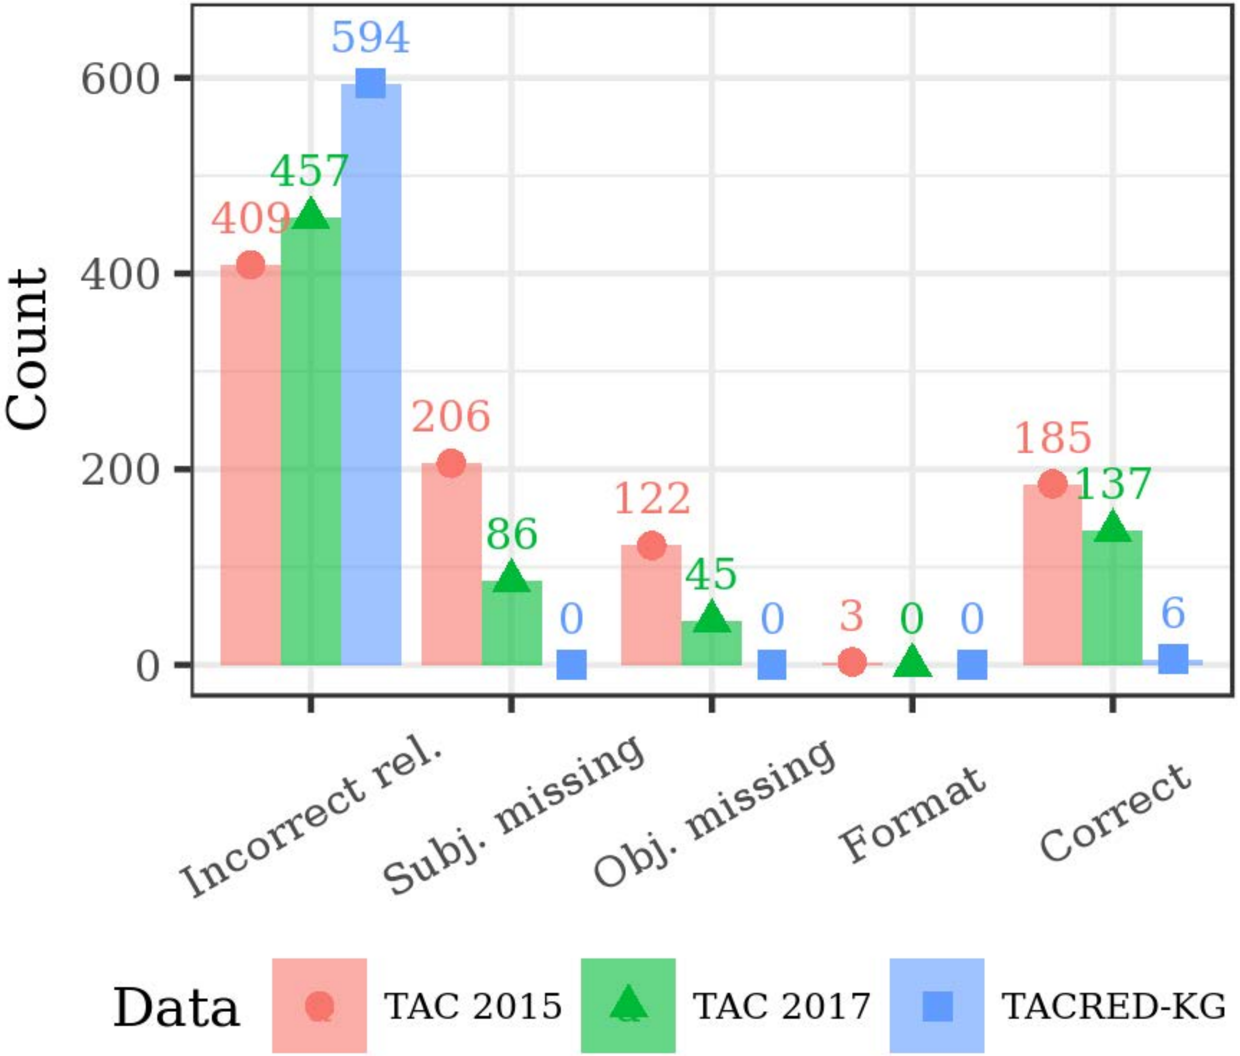

Approach: MLP Architecture

Representation

Belief and Provenance Combination

Feature Learning and Classification

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 1: Prepare Input

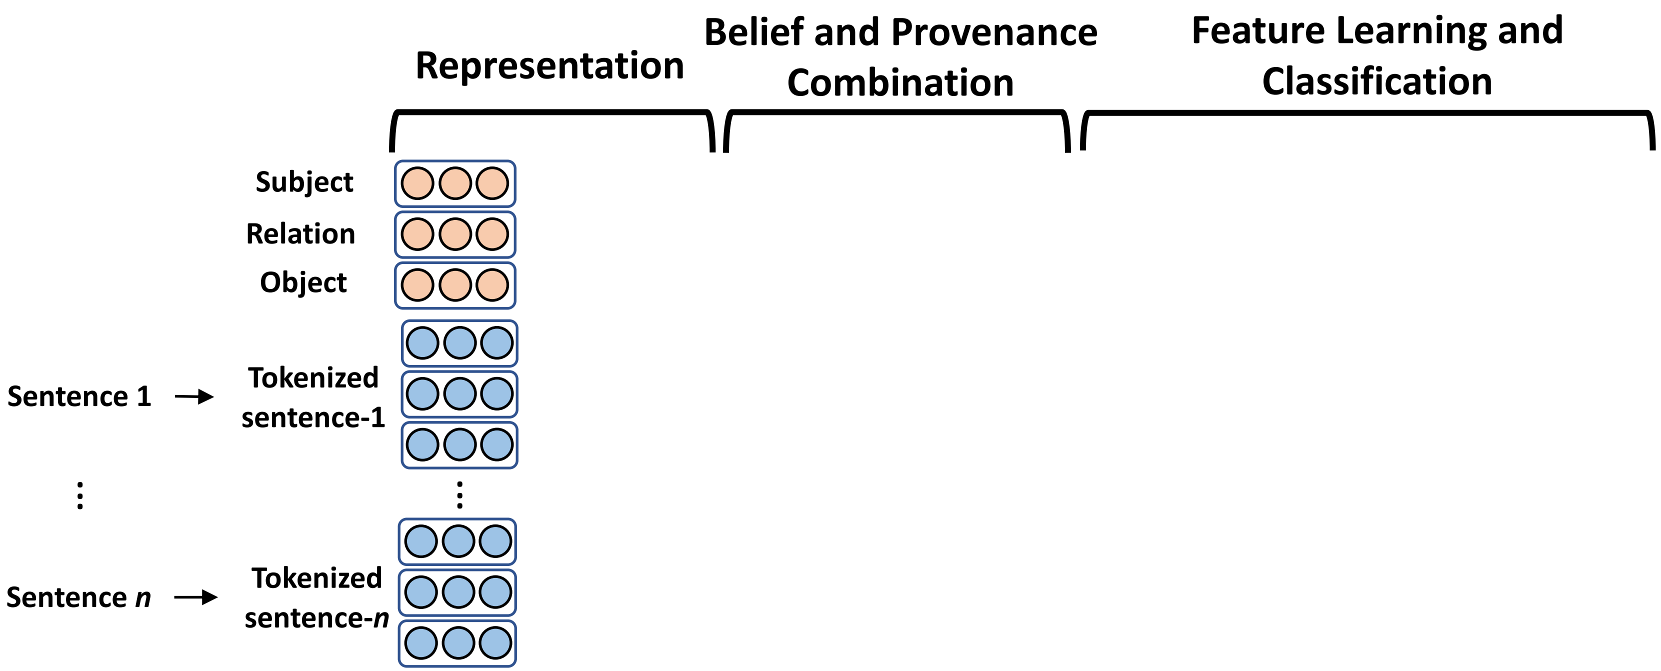

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 2: Obtain Representation

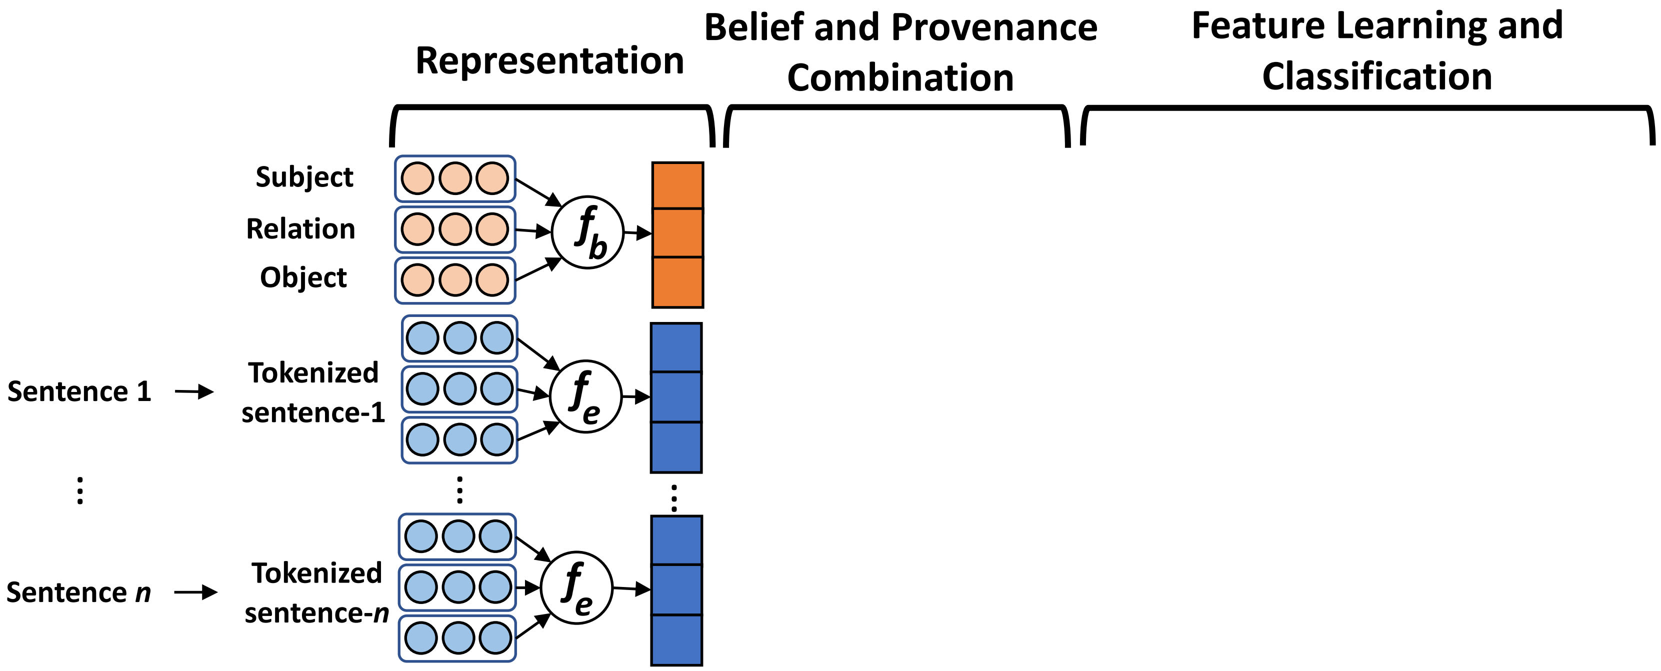

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 3: Apply Attention/Combination

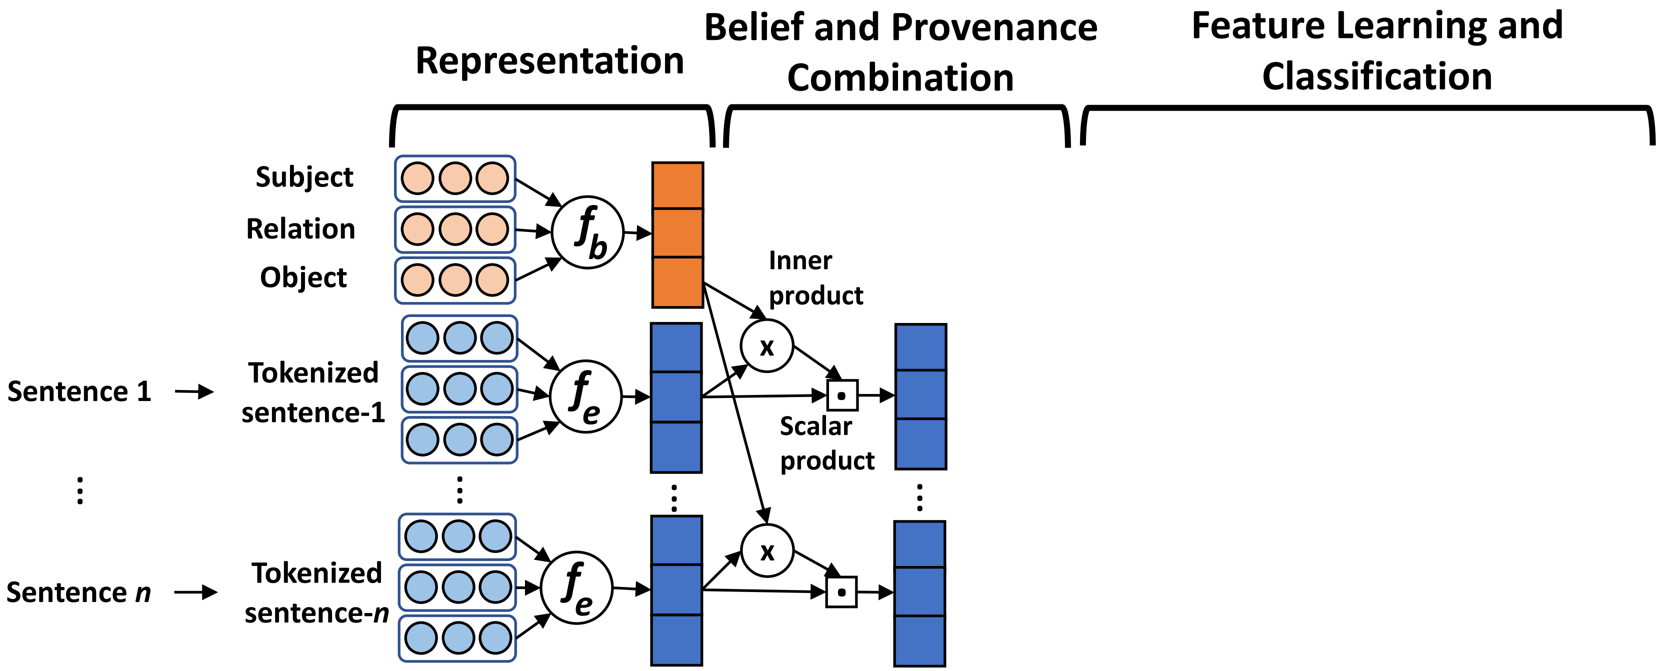

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 4: Aggregate

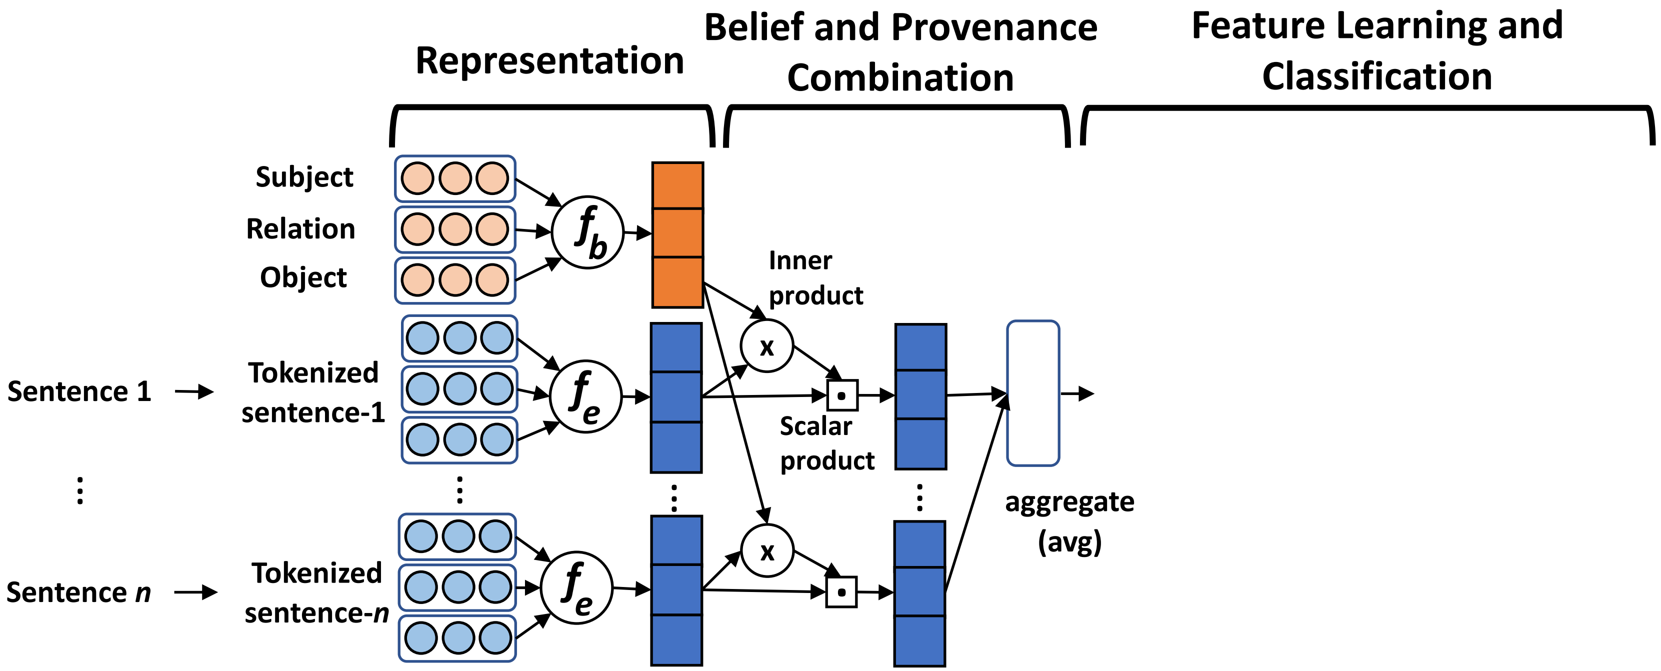

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 5: Learn Abstract Features

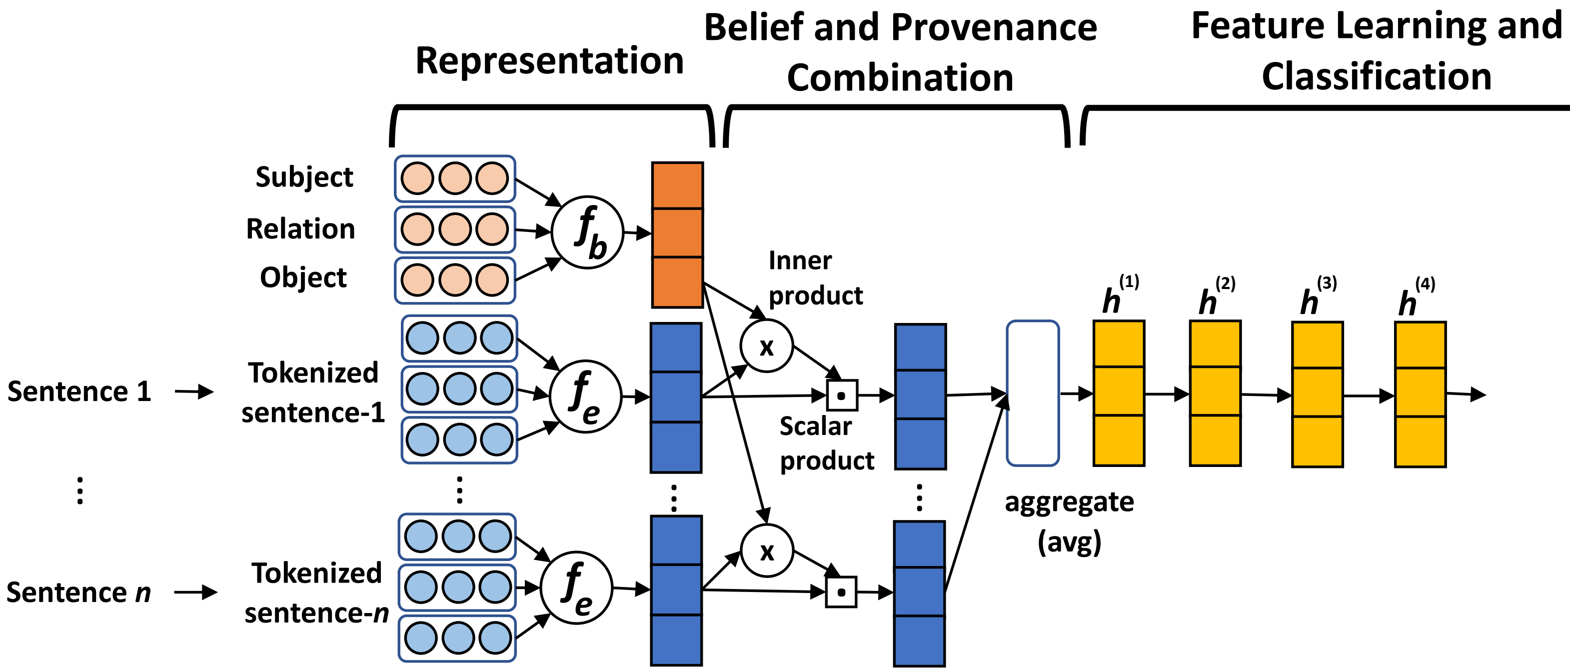

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 6: Add Skip Connection

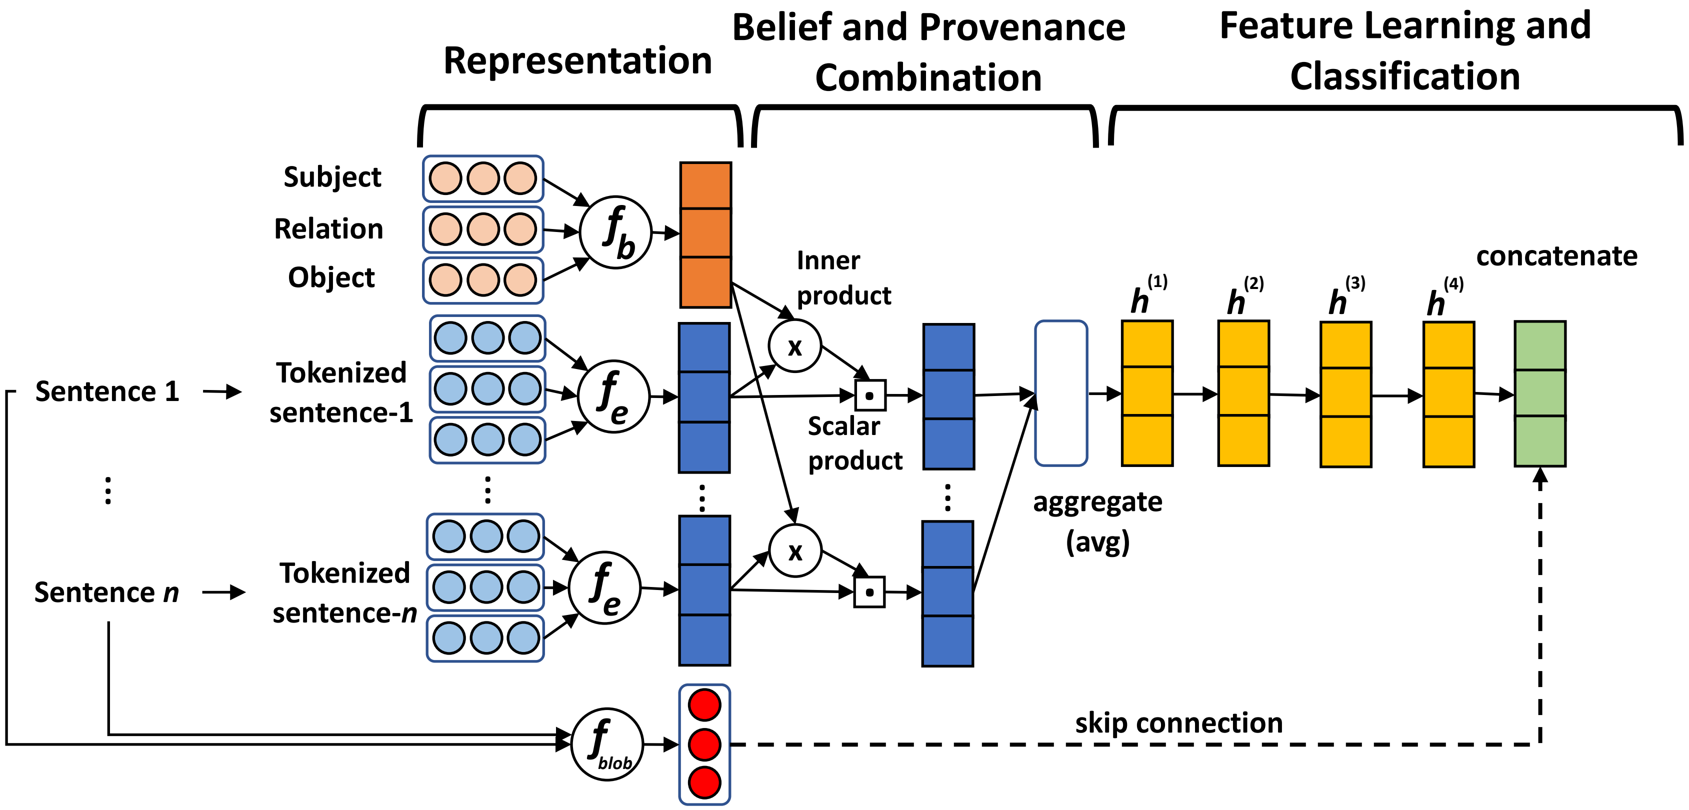

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Step 7: Jointly Classify

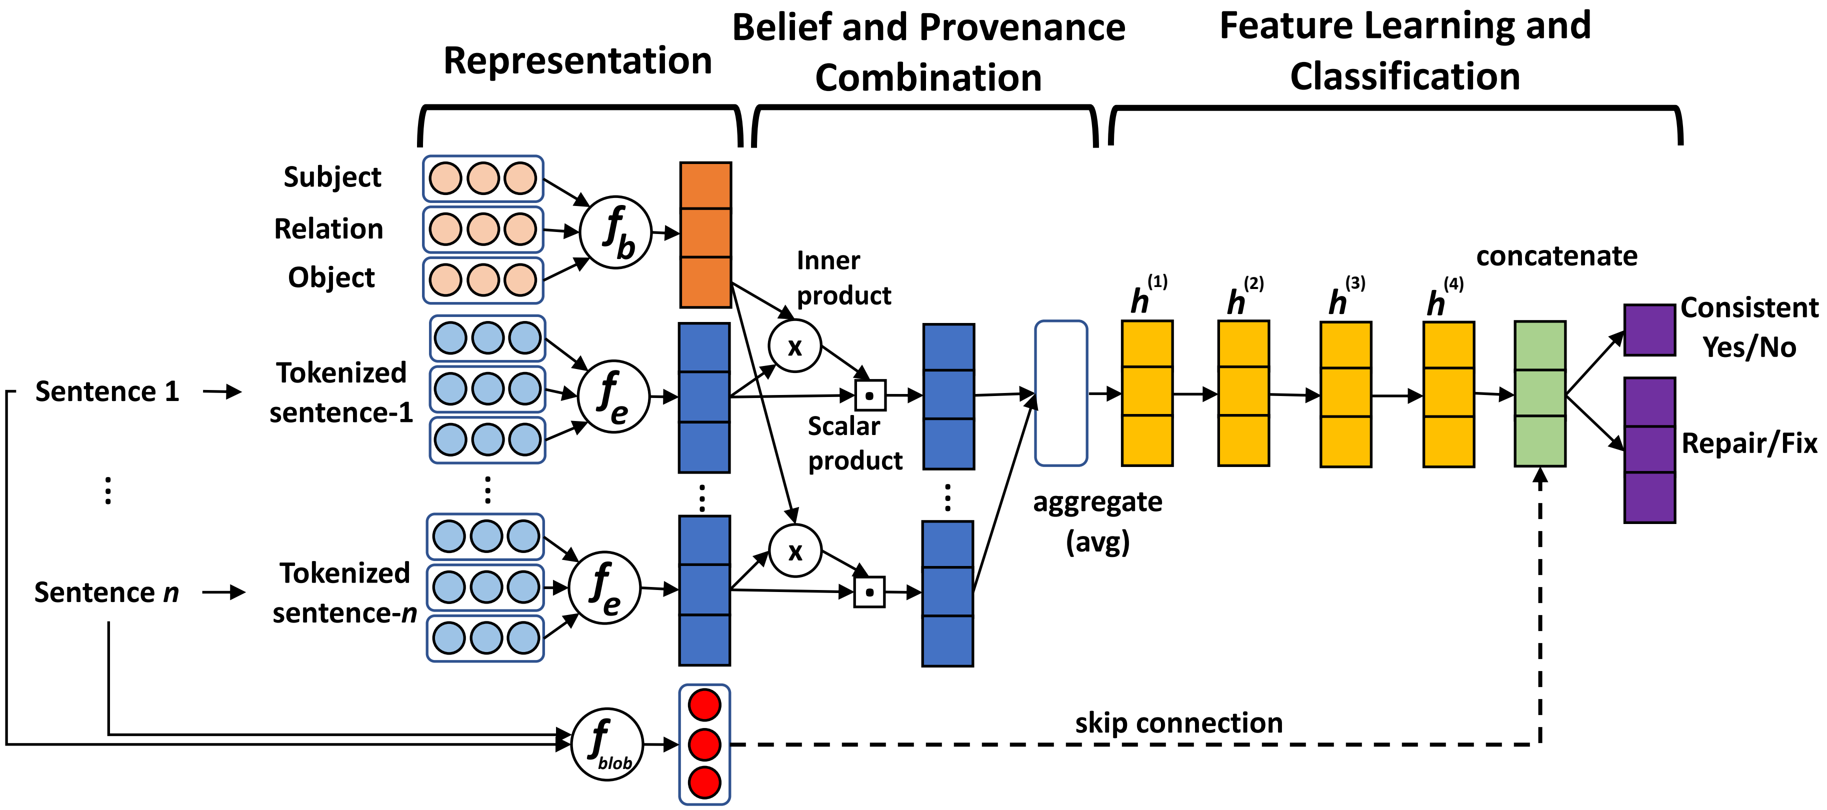

Given a belief & set of provenance sentences, we jointly determine their consistency and a repair if deemed inconsistent

Experiments and Questions

- Are existing credibility models sufficient: DeClarE (Popat et al., 2018) and LSTM-text (Rashkin et al., 2017?
- What are important features and representations?
- How important is word order?
- Evaluation using three datasets
  - TAC KBP 2015 & 2017; output from IE systems + gold standard
  - TACRED relation extraction dataset
  - 27% of 34k beliefs judged consistent – TAC 2015
  - 36% of 57k beliefs judge consistent – TAC 2017
  - <1% of 106k belief are consistent – TACRED-KG

Using existing credibility models for Consistency?

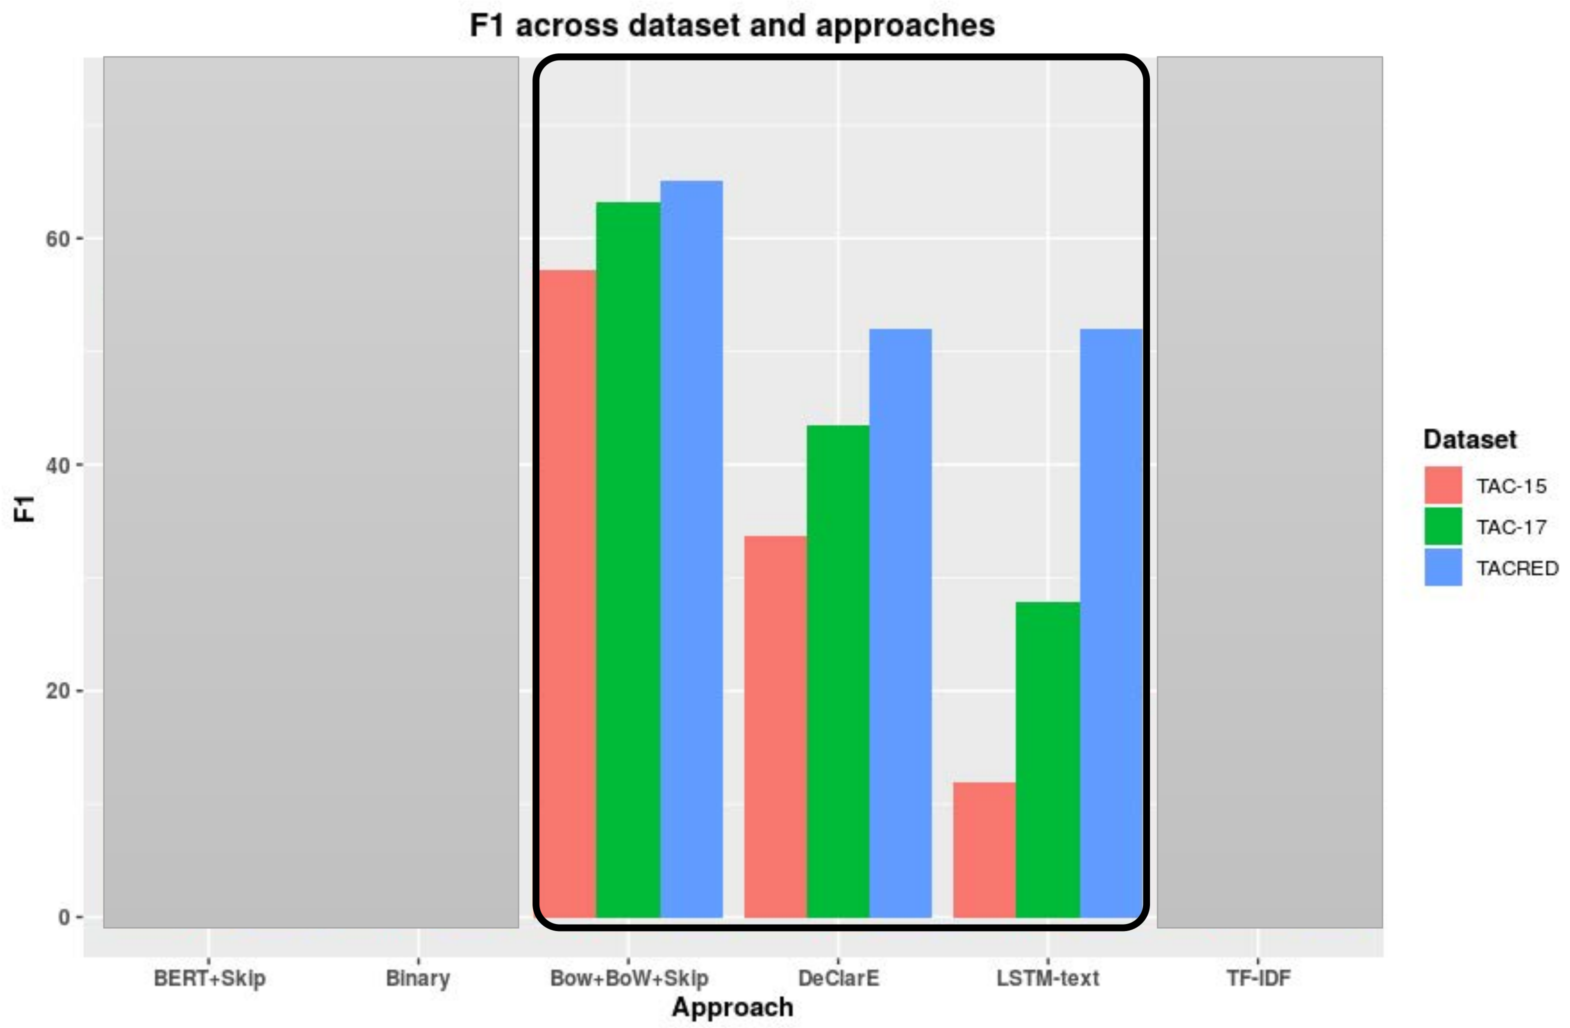

Credibility model do n ot work well for the consistency task

Using existing credibility models for Repair?

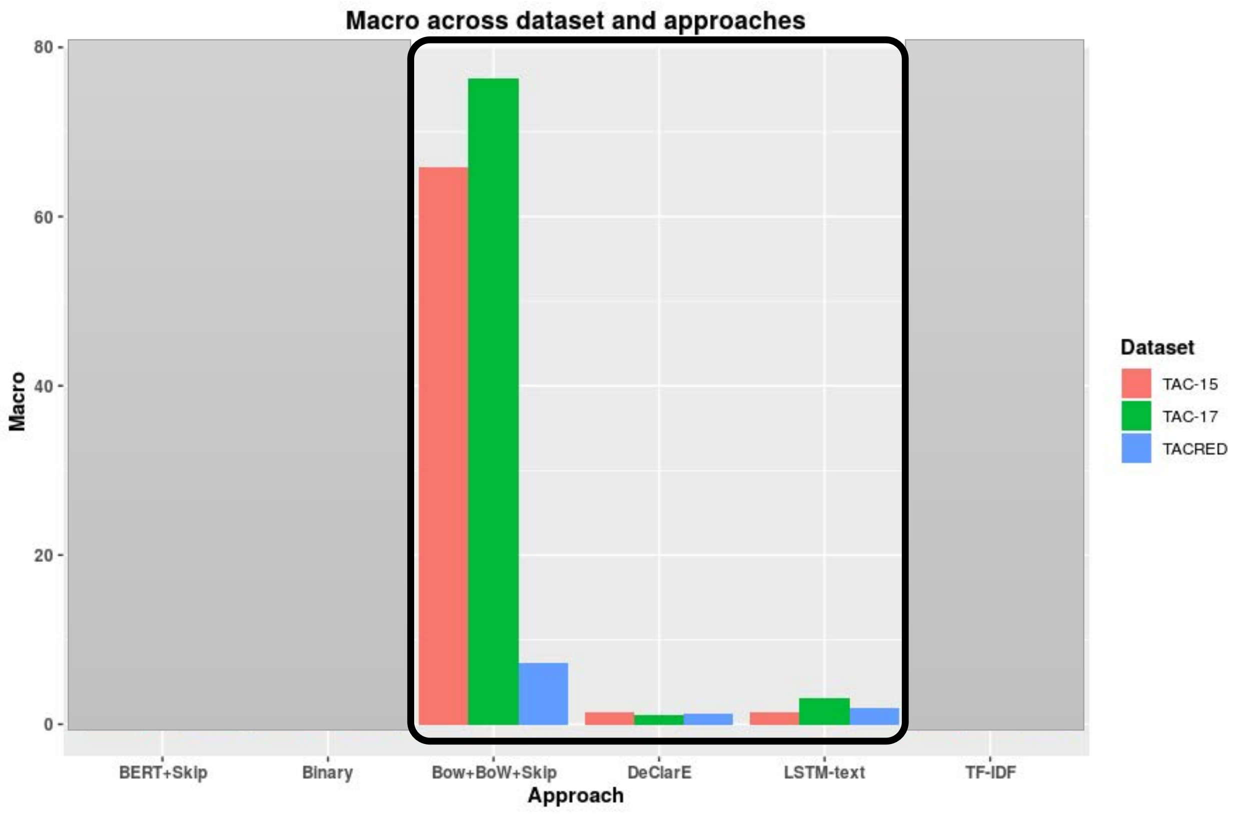

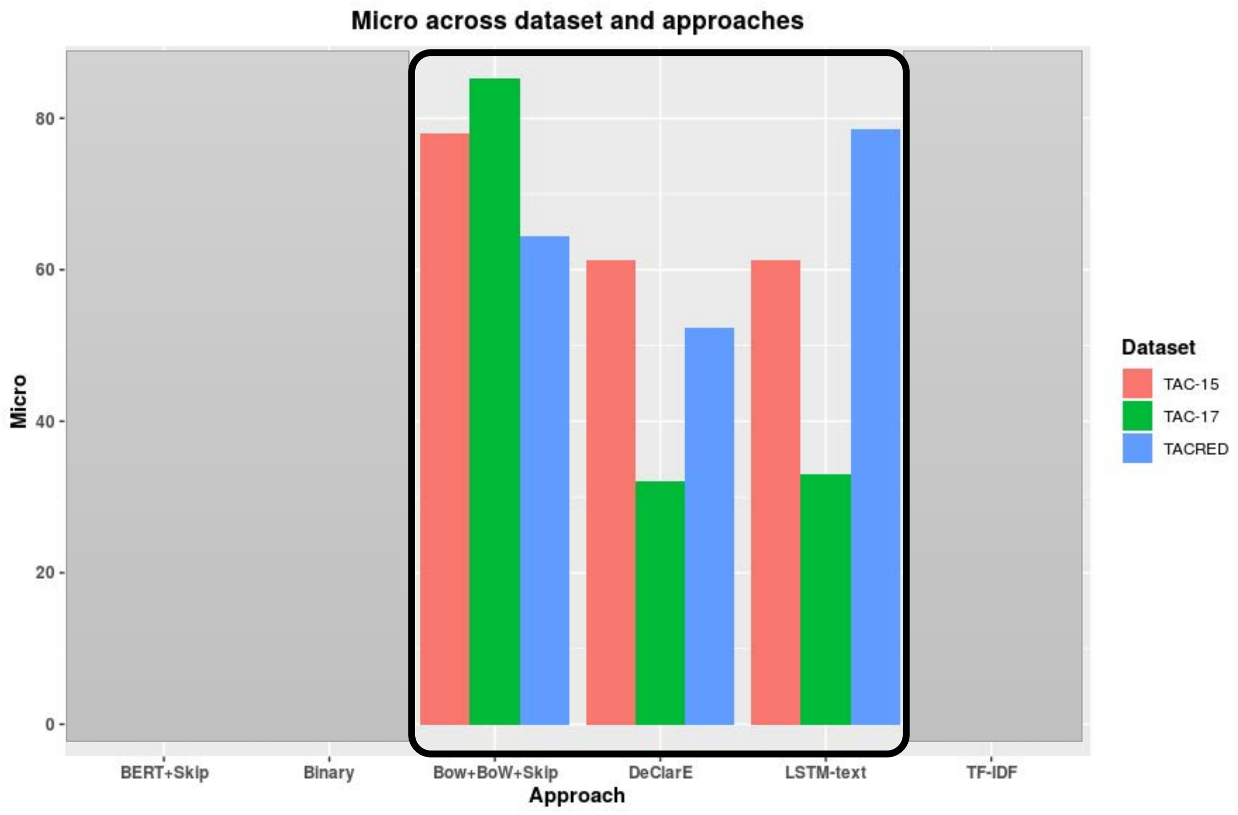

Nor do credibility models work well for the repair task

What Representations are Effective?

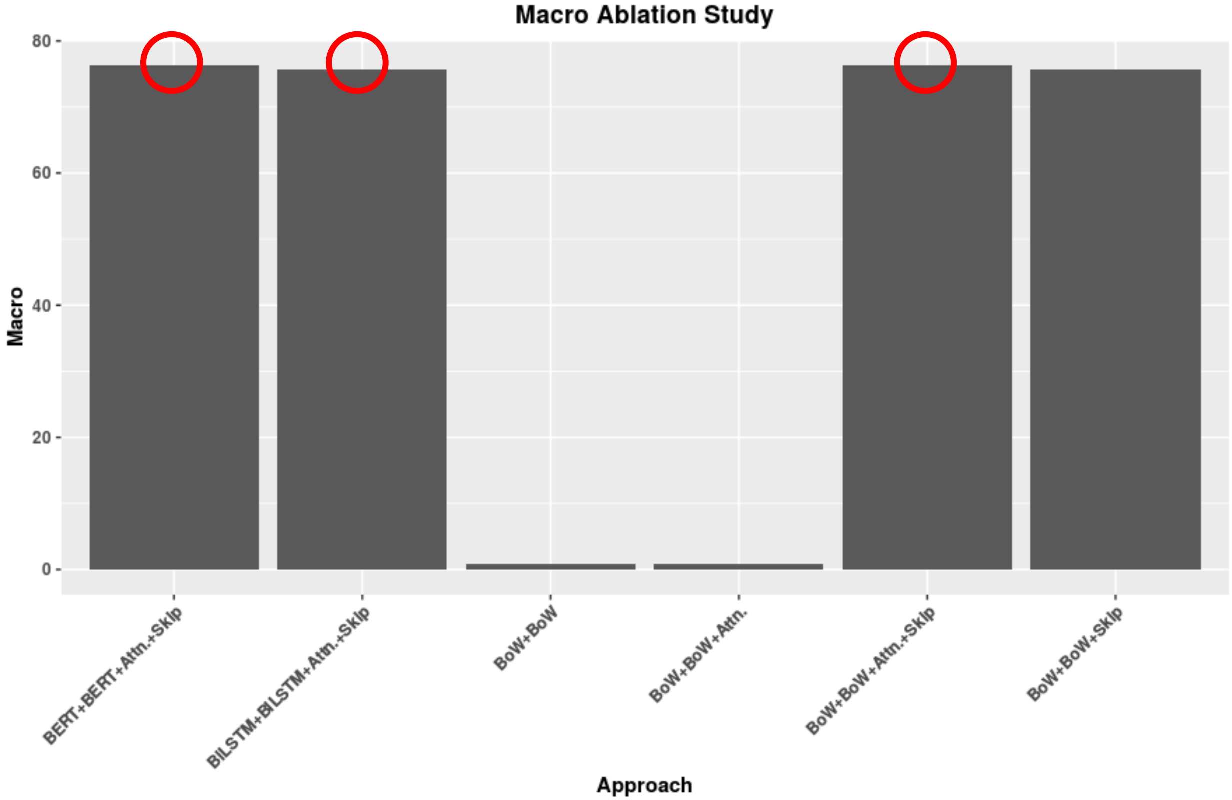

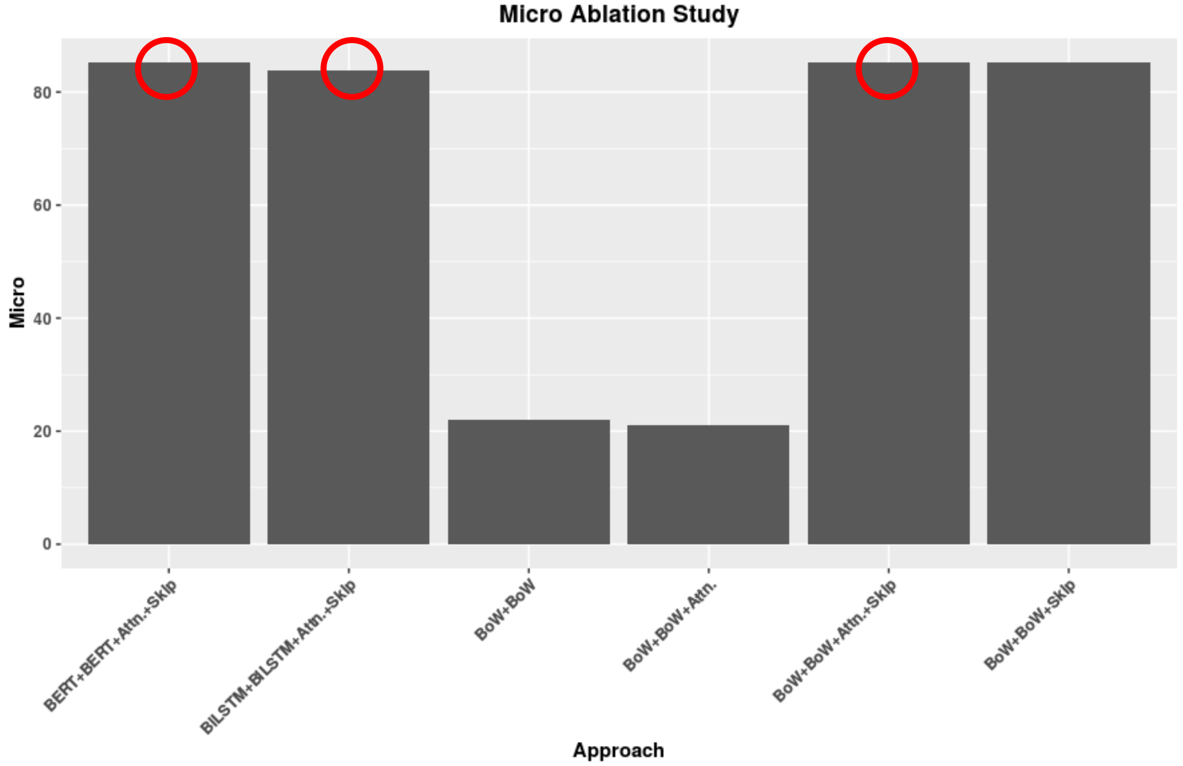

BoW with Attn. + Skip performs similar to BERT and Bi-LSTM

KGcleaner Conclusions

- Measuring **Consistency,** not **Credibility**, is often desired
- Most of the mistakes encountered while learning can be classified as *relation not entailed*
- IE system errors are systematic and are lexical in nature
- Simple model can outperform a more expressive SOTA one
- Language composability can be compromised
- Skip connection increases model performance slightly
- Data balancing using weighted cost function helped address imbalanced data

Conclusions

Knowledge graphs & machine learning have a symbiotic relationship

- Machine learning techniques can create, augment and improve knowledge graphs
- Knowledge graphs can provide data and features for machine learning systems to learn from

Both aspects are important and being actively pursued by the research community

- There are many important and hard problems that remain