RAMFIS: Representations of vectors and Abstract Meanings for Information Synthesis – TA2

TAC 2019

Martha Palmer,
Rehan Ahmed, Cecilia Mauceri
University of Colorado, Boulder

Computational Language and EducAtion Research
C L E A R

Our Team

	KB/Ontology	Images and Video
Univ. Colorado	Martha Palmer (PI) Jim Martin, Susan Brown, **Rehan Ahmed**, Chris Koski, ….	Chris Heckman, **Cecilia Mauceri**,
Colo. State		Ross Beveridge, David White
Brandeis	James Pustejovsky, Peter Anick	James Pustejovsky Nikhil Krishnaswamy

How did we achieve highest frame recall score?

- Efficient AIF object manipulation
- Merge multiple TA1s
- Streaming clustering
- Simple linking metrics

How did we achieve highest frame recall score?

- **Efficient AIF object manipulation**
- Merge multiple TA1s
- Streaming clustering
- Simple linking metrics

AIF Objects (java)

- Read / Write
- Compare
- Merge

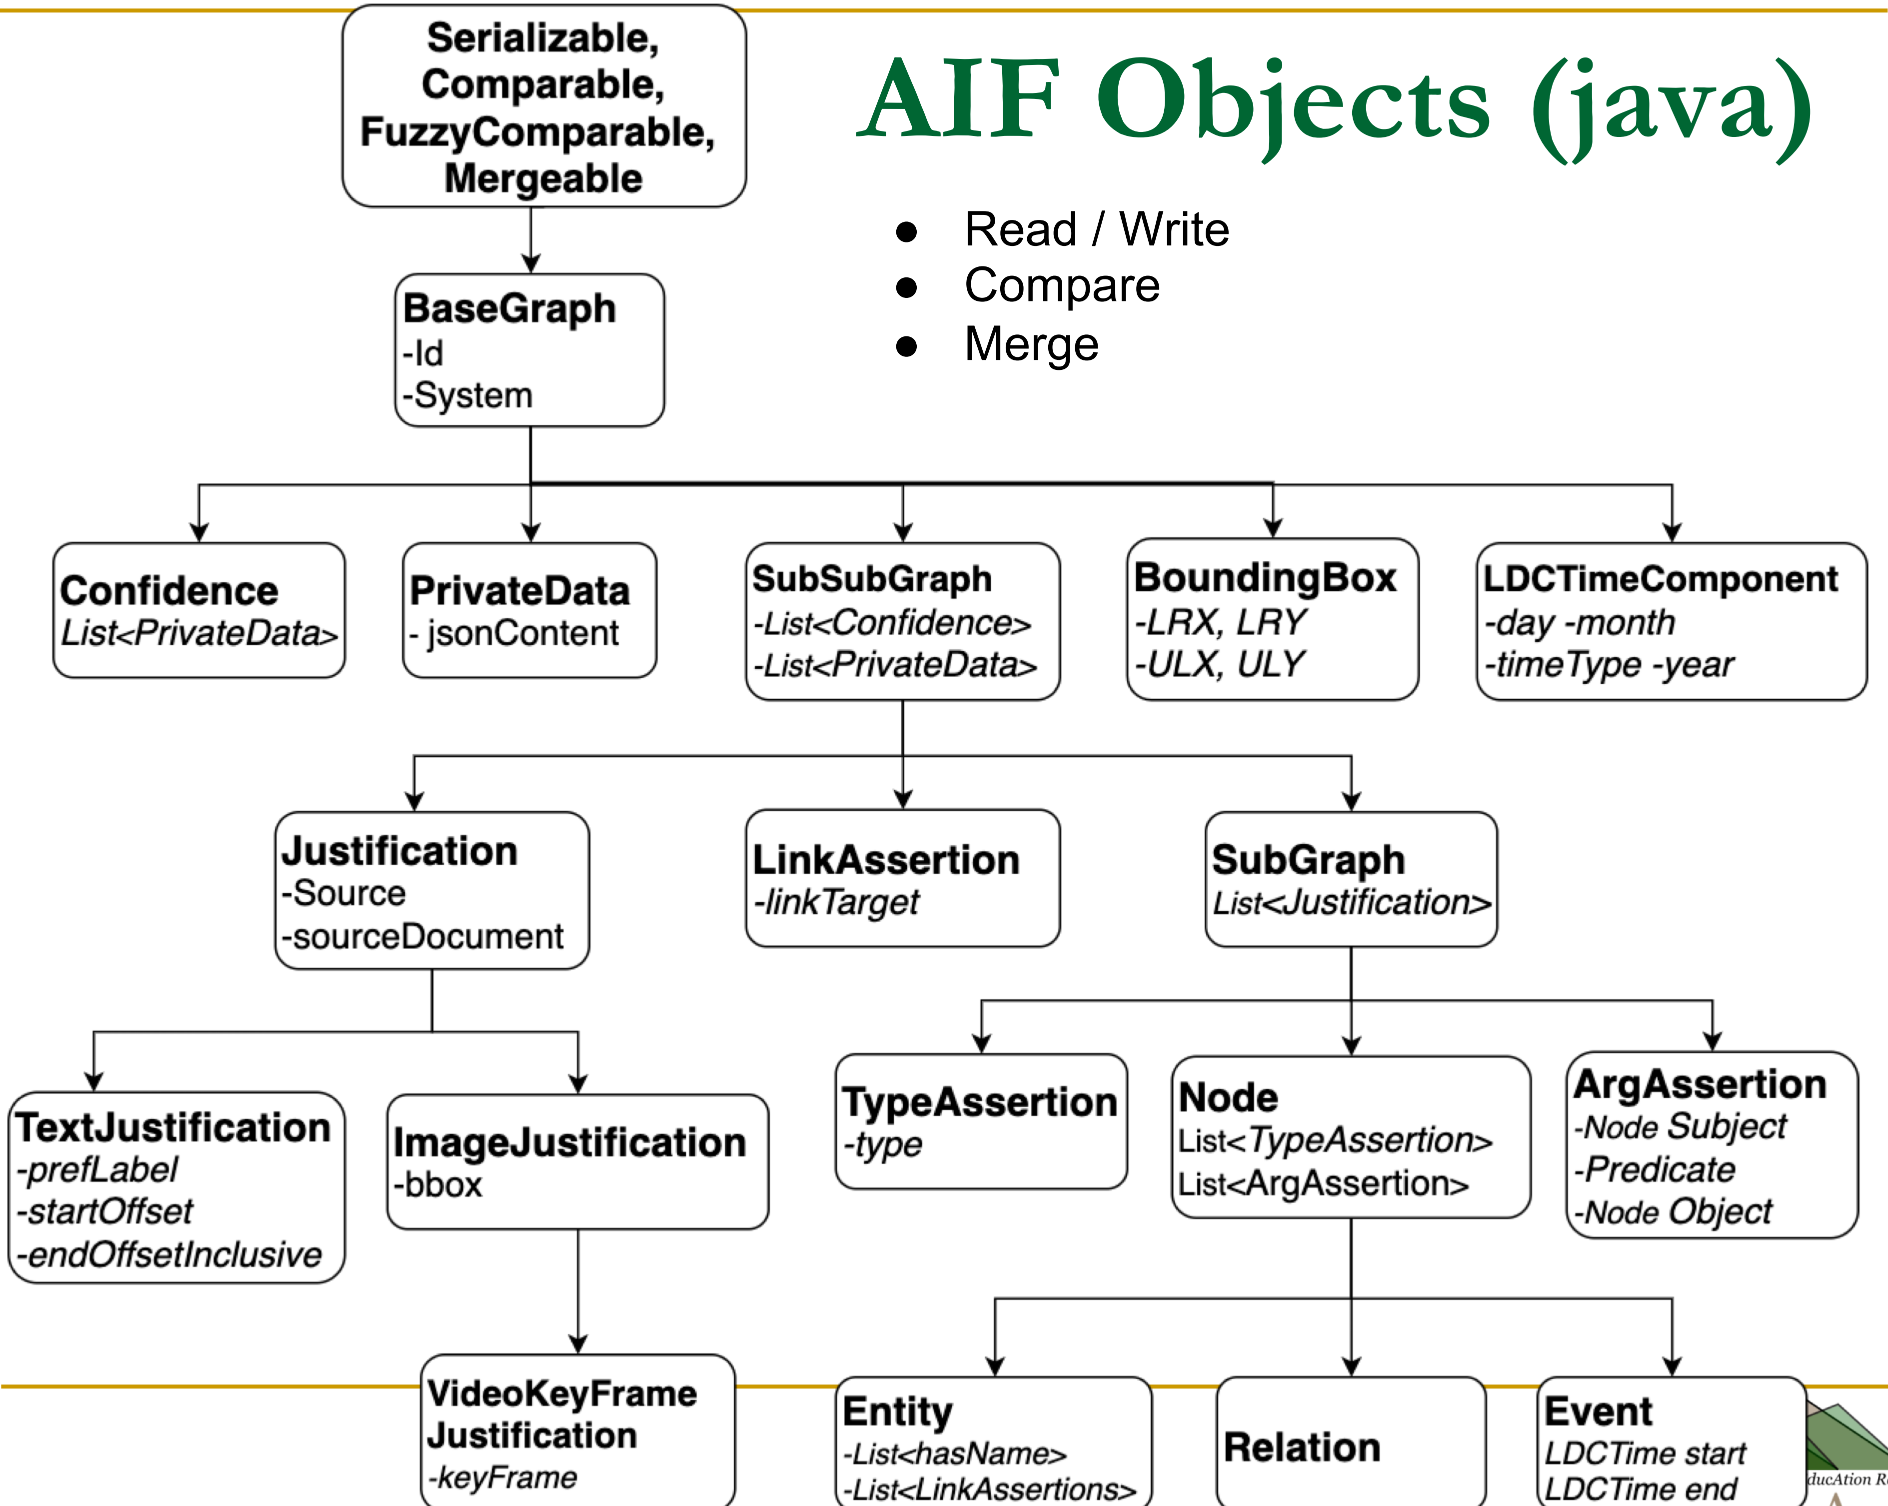

Software Engineering - Read/Write

- Read/Write Criteria
  - Distributed
  - Interfaces with many platforms

- Read

jena rdf4j GraphDB™

- Write
  - Efficient triples writer - AIF2Triples
  - The output can be split into smaller files (TA3 consumers liked this!)
  - Developed at Colorado

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node

Entity 1

List<hasName> :
["President Putin"]

List<Justification>:

Confidence: 0.9 (GAIA)

Entity 2

List<hasName> :
["Vladimir Putin"]

List<Justification>:

Confidence: 0.8 (BBN)

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node

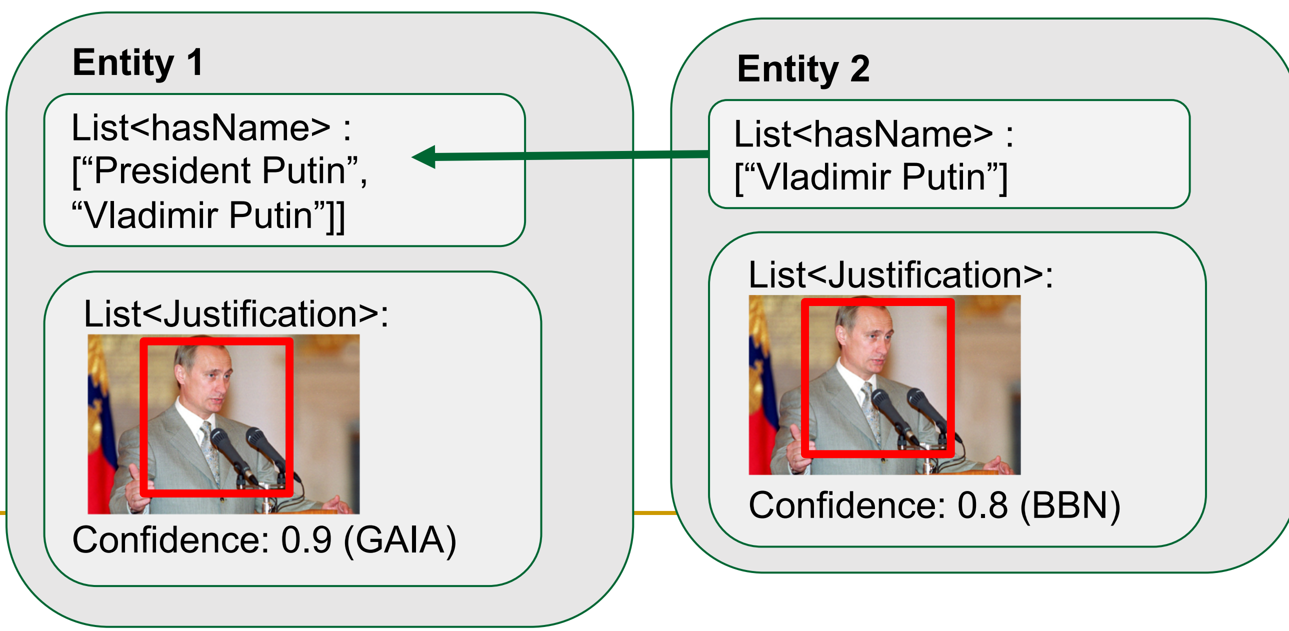

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node
  - Propagates through all sub-graphs

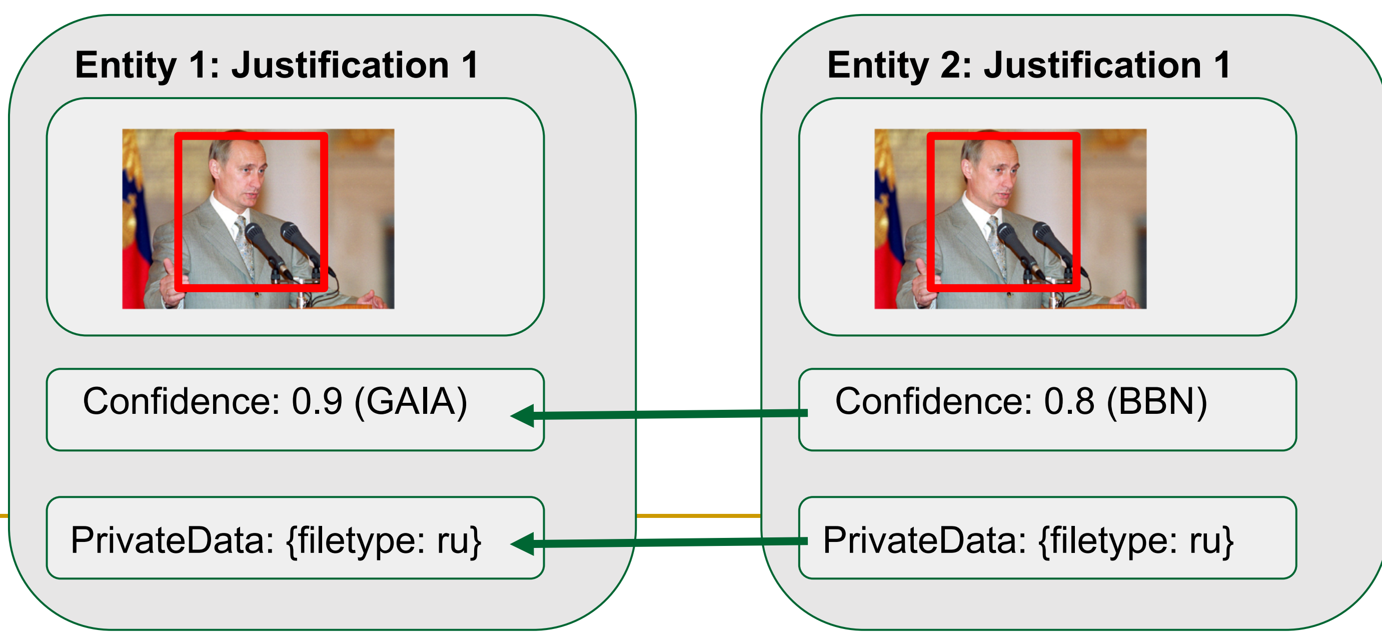

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node
  - Propagates through all sub-graphs

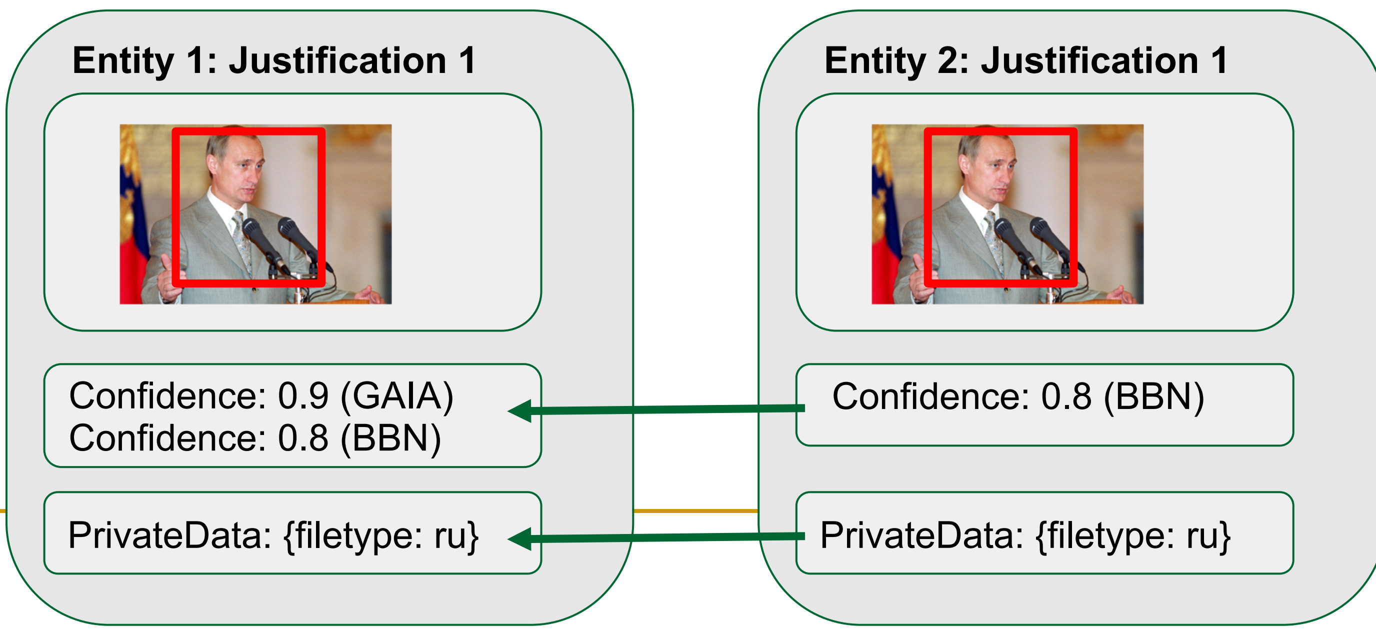

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node
  - Propagates through all sub-graphs

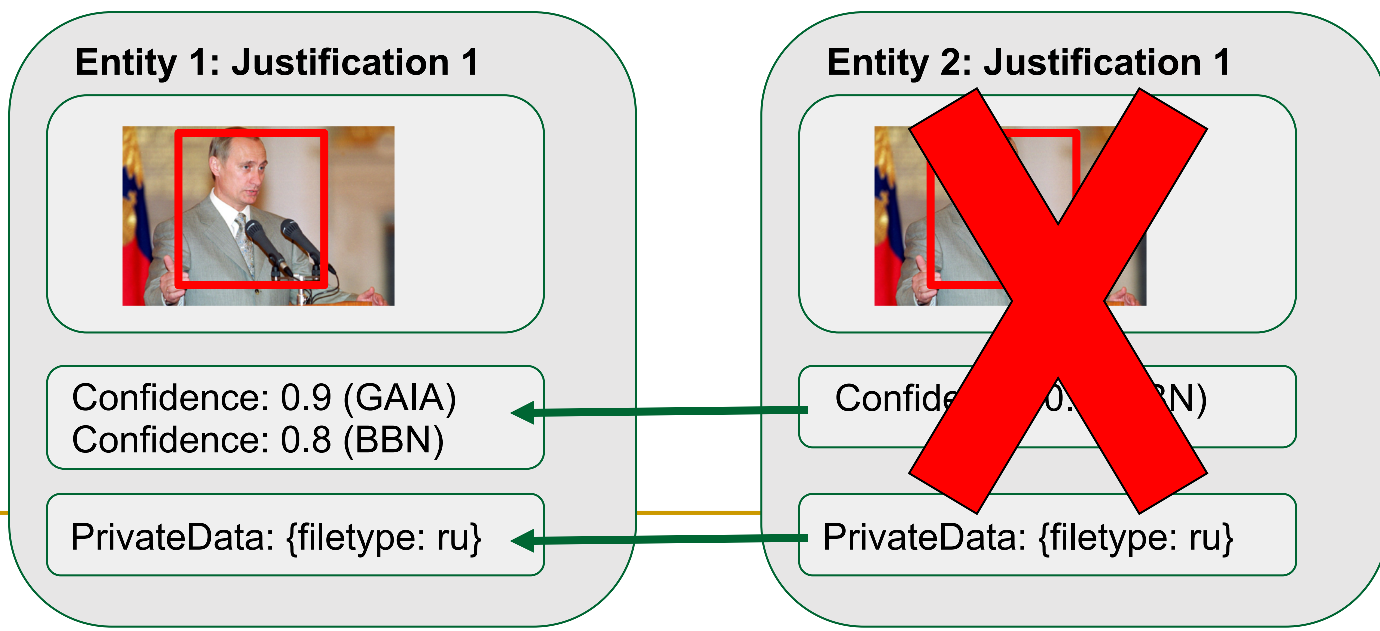

Computational Language and EducAtion Research
CLEAR

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node
  - Propagates through all sub-graphs

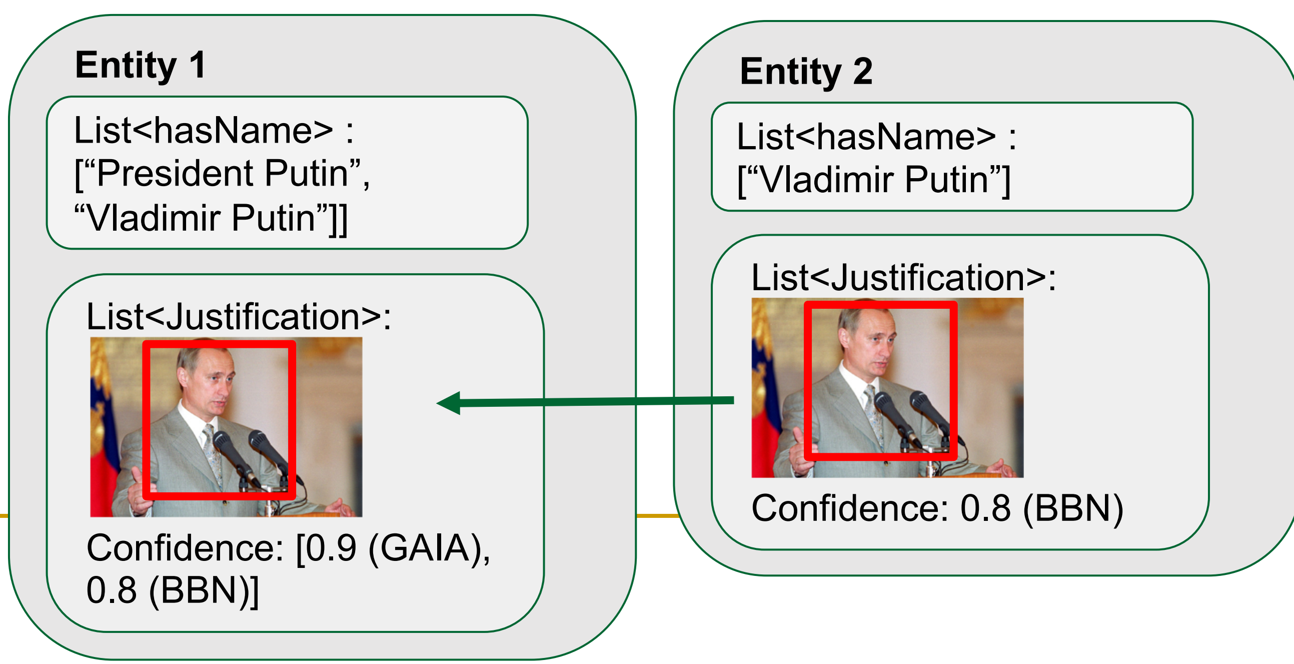

Computational Language and EducAtion Research
C L E A R

Software Engineering - Compare & Merge

- Each object has a comparison function (not just Entity, Event, Relation)
  - Merge duplicate justifications, private data, system information etc

- Merging is initiated by a Node
  - Propagates through all sub-graphs

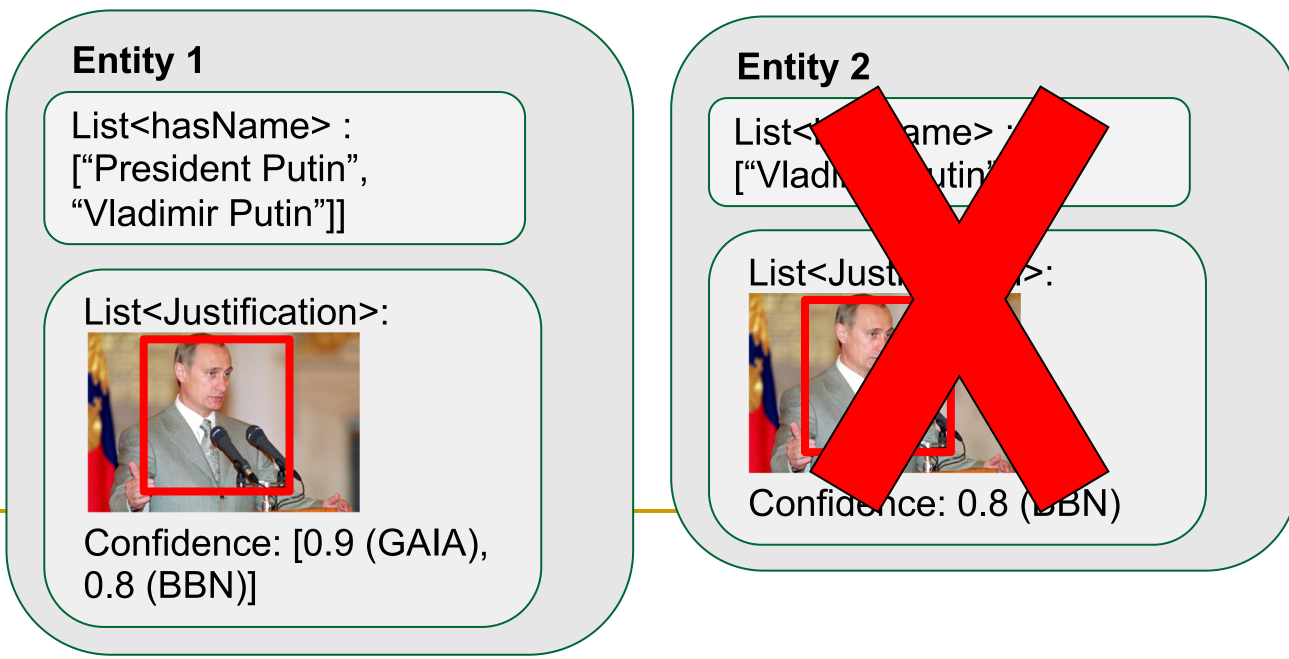

How did we achieve highest frame recall score?

- Efficient AIF object manipulation
- **Merge multiple TA1s**
- Streaming clustering
- Simple linking metrics

Benefits of Merging Multiple TA1

- Goal of AIDA to combine diverse data sources
- Additional coverage by using a diversity of models
- For example, increased coverage of reference KB links

Merging multiple TA1s

Merging the same source document across different TA1s

GAIA_1

HC0000A1T.ttl
HC0000AA3.ttl
HC0000AAP.ttl
HC0000AE1.ttl
…

OPERA_3

HC0000A1T.ttl
HC0000AA3.ttl
HC0000AAP.ttl
HC0000AE1.ttl
…

Merging multiple TA1s

Merging the same source document across different TA1s

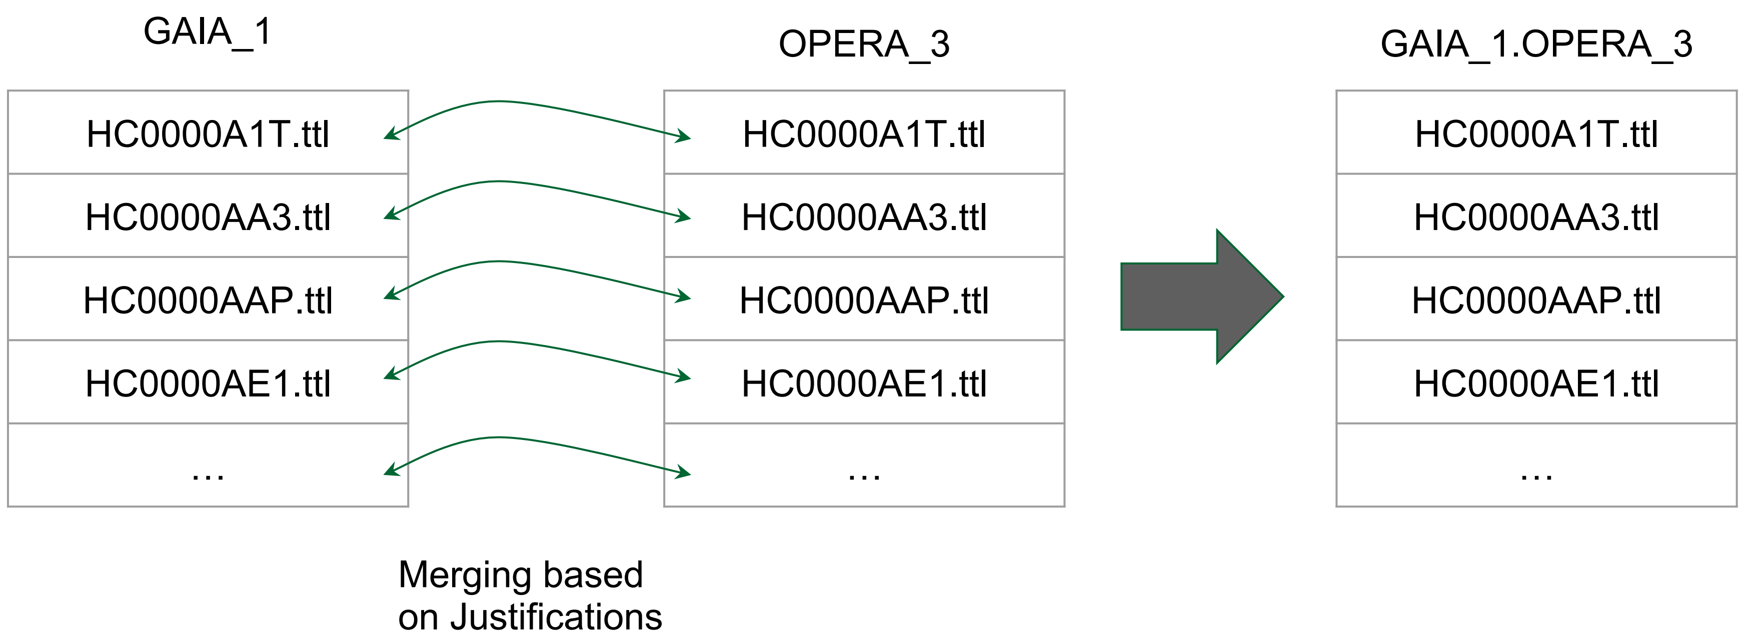

TAC 2019 Submissions

TA 1	Triples pre clustering	Triples post clustering
GAIA_1	31,987,759	30,324,882
GAIA_2	48,423,300	29,532,733
OPERA_3	23,290,306	12,665,445
GAIA_1 + Michigan_1	65,437,918	51,143,310
GAIA_1 + OPERA_3	45,787,436	35,134,812
GAIA_1 + JHU_5	60,421,533	55,194,984
…	…	…

OPERA_ADITI_V2

TAC 2019 Submissions

TA 1	Entities pre clustering	Entities post clustering	Events pre clustering	Events post clustering
BBN_1	270,168	232,785	107,050	89,836
GAIA_1	358,436	309,358	37,205	31,151
GAIA_2	459,044	310,437	34,127	23,743
OPERA_3	339,718	200,776	13,126	10,068
GAIA_1 + OPERA_3	587,977	458,931	43,526	36,800
GAIA_1 + JHU_5	758,978	690,166	85,393	75,820
…	…	…	…	…

How did we achieve highest frame recall score?

- Efficient AIF object manipulation
- Merge multiple TA1s
- **Streaming clustering**
- Simple linking metrics

Diagram

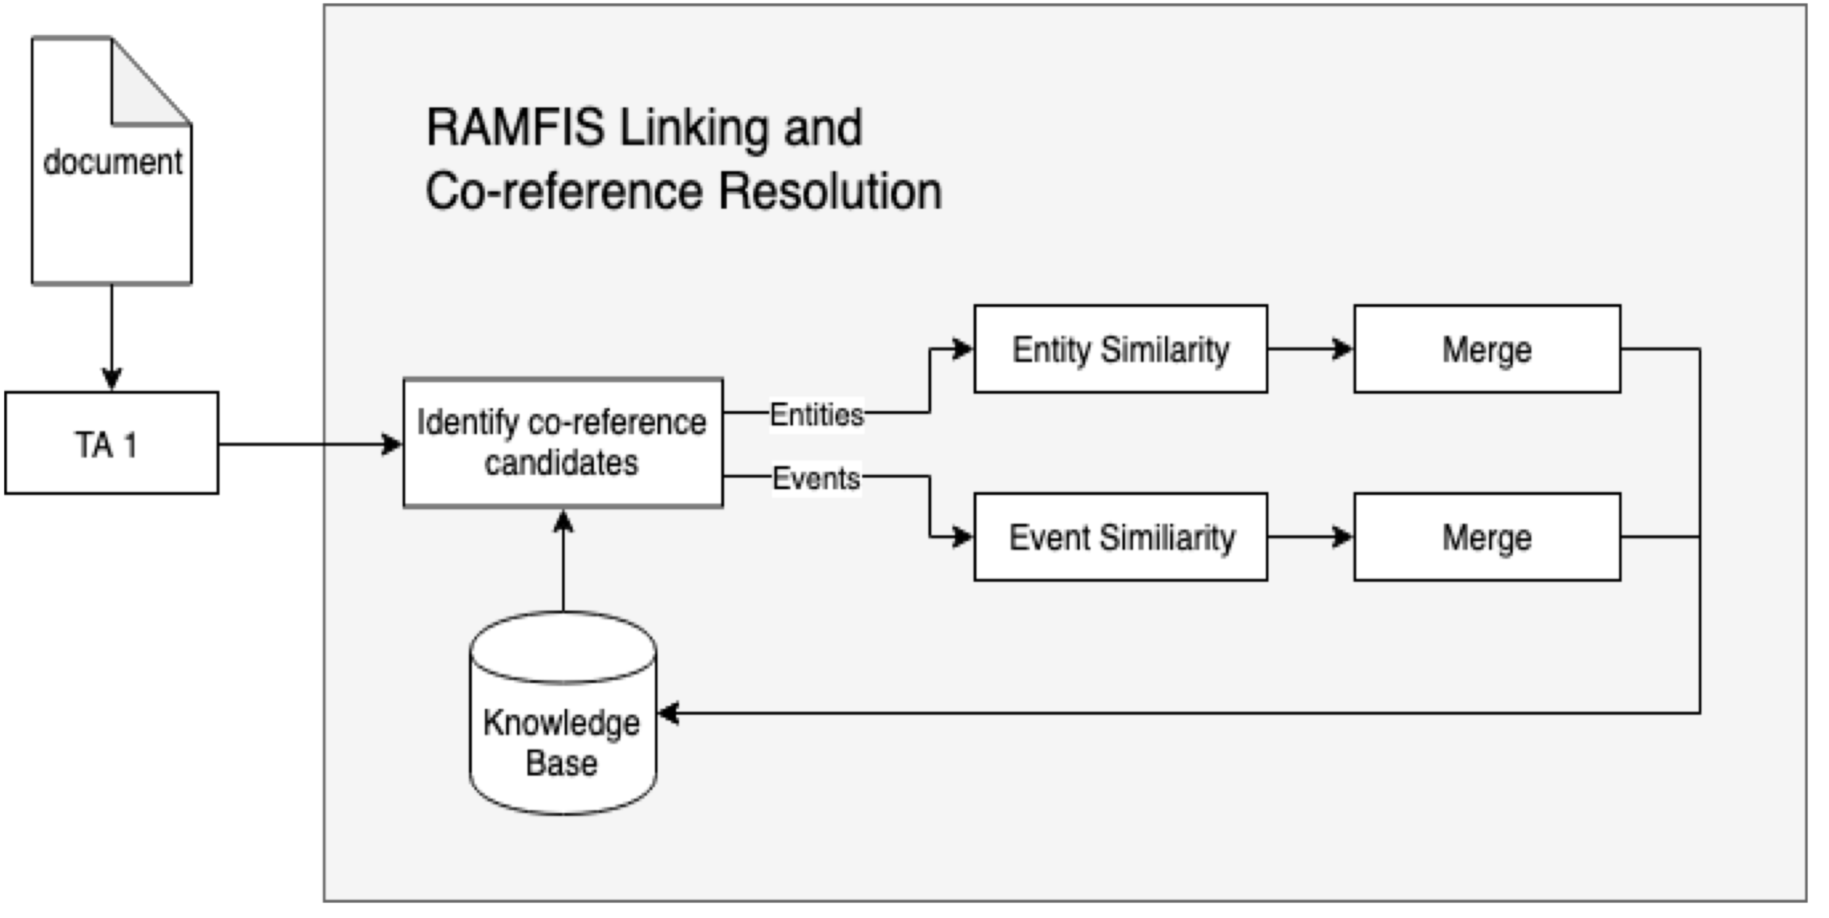
document
TA 1
RAMFIS Linking and
Co-reference Resolution
Identify co-reference candidates
Entities
Events
Entity Similarity
Event Similiarity
Merge
Merge
Knowledge Base

Linking Candidates

PERSON: “Tr”

LOCATION: “Tr”

For all Entities of

- Same type
- Same name substring

Compare all pairs

Photo attributions:
Melania Trump - By Regine MahauxWeaver
Justin Trudeau - By Presidencia de la República Mexicana
Trump Tower - By Potro
Tribune Tower - By Luke Gordon

Linking Candidates

PERSON: “Tr”

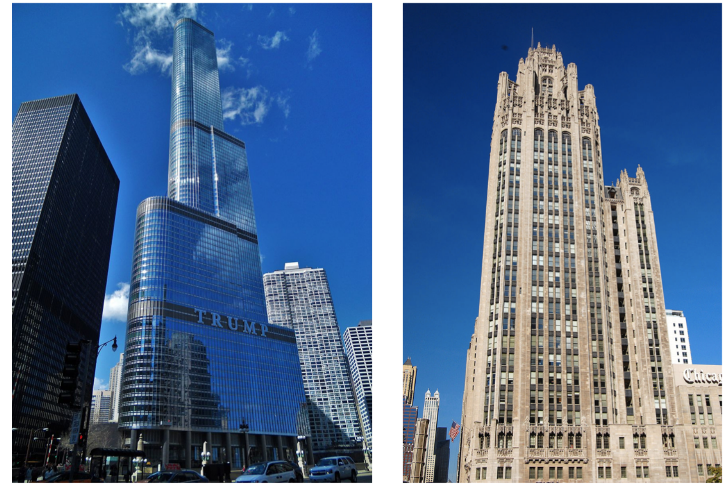

LOCATION: “Tr”

For all Entities of

- Same type
- Same name substring

Compare all pairs

Photo attributions:
Melania Trump - By Regine MahauxWeaver
Justin Trudeau - By Presidencia de la República Mexicana
Trump Tower - By Potro
Tribune Tower - By Luke Gordon

Computational Language and EducAtion Research
C L E A R

Linking Candidates

PROTEST
- Patient: Ukrainian Government

PROTEST
- Topic: Black Lives Matter

For all Event of

- Same type
- Same role label

Photo attributions:
Euromaidan Protests - By Mstyslav Chernov
Black Lives Matter Friday - By The All-Nite Images

How did we achieve highest frame recall score?

- Efficient AIF object manipulation
- Merge multiple TA1s
- Streaming clustering
- **Simple linking metrics**

Similarity Criteria

Entities

- Type matching
- Fuzzy Name matching
- Justification overlap

Events

- Type matching
- Participant matching
- Justification overlap

Similarity Criteria

Entities

- **Type matching** → PERSON, ORGANIZATION, GEOPOLITICAL ENTITY LOCATION ...
- Fuzzy Name matching
- Justification overlap

Events

- **Type matching** → ControlEvent MovementEvent ConflictEvent ..
- Participant matching
- Justification overlap

AIDA Ontology Types

Similarity Criteria

Entities

- Type matching
- **Fuzzy Name matching**
- Justification overlap

Events

- Type matching
- Participant matching
- Justification overlap

President Obama
Senator Obama
Obama ?
Mr. Obama ?
~~Michelle Obama~~
~~Mrs. Obama~~
Barack Obama
Barack H. Obama
Barack Hussein Obama
~~Barack Hussein Obama Sr.~~
Barack ?

Similarity Criteria

Entities

- Type matching
- **Fuzzy Name matching**
- Justification overlap

NYC
New York City
~~New York State~~
New York ?
NY ?
~~NYU~~
New York, New York

Events

- Type matching
- Participant matching
- Justification overlap

Computational Language and EducAtion Research
C L E A R

Similarity Criteria

Entities

- Type matching
- Fuzzy Name matching
- Justification overlap

Events

- Type matching
- **Participant matching**
- Justification overlap

PROTEST
- Patient: Entity 1
- Topic: Entity 2

PROTEST
- Patient: Entity 3
- Topic: Entity 2

PROTEST
- Patient: Entity 1

Similarity Criteria

Entities

- Type matching
- Fuzzy Name matching
- **Justification overlap**

Events

- Type matching
- Participant matching
- **Justification overlap**

ImageJustification Threshold

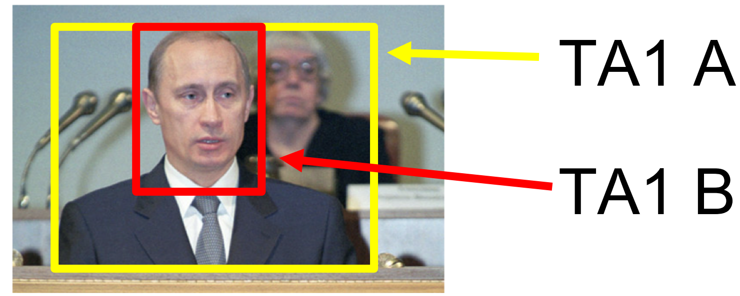

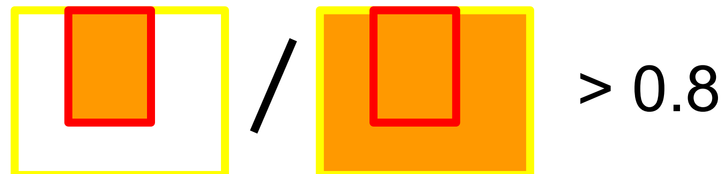

Intersection over union

TextJustification Threshold

… President Vladimir Putin …

Intersection over union > 0.8

Cross-Document Co-Reference Performance

Baseline coref scores on annotated datasets (cross-doc)

Event Coref Bank Data - scores for ∩

	Gold standard	TA1 output	∩	B^3 P	B^3 R	B^3 F1	MUC P	MUC R	MUC F1
Events	3437	5107	918	95.9	42.75	59.14	63.04	10.96	18.67
Entities	4268	8820	864	98.1	64.33	77.7	95.08	54.2	69.04
Both	7705	13927	1782	95.7	57.05	71.5	54.71	10.96	18.26

Baseline coref scores on annotated datasets (cross-doc)

DEFT Richer Event Descriptions BCUB score

	Precision	Recall	F1
Events	80.11	14.14	24.05
Entities	46.45	49.55	47.95
Combined	83.97	30.83	45.11

Room for improvement? Yes!

Graph Queries

	Prec(1a)	Recall(1a)	F1(1a)	Recall(1b)	**Frame Recall**
GAIA1_OPERA3	0.24	0.11	0.15	0.14	**0.05**

Zero-Hop Queries

	AP-B	AP-W	AP-T
GAIA1_OPERA3	0.0667	0.0667	0.0667

Future Work

- Linking using graph embeddings
- Nearest neighbor KB search
- Vector similarity
- Affine mapping between embedding vectors

Future Work

- **Linking using graph embeddings**
- Nearest neighbor KB search
- Vector similarity
- Affine mapping between embedding vectors

Event Linking by example (1)

A day after MH17 was shot down over Ukraine's warring eastern provinces on July 17, 2014, the United States government concluded from available evidence that the plane had been brought down by a Russian-made surface-to-air missile launched from rebel-held territory in eastern Ukraine. American officials said at the time that they believed the missile battery had most likely been provided by Russia to pro-Russian separatists.

Article: https://www.nytimes.com/2019/06/19/opinion/mh17-ukraine-russia-suspects.html

Event Linking - Building Knowledge Graph

A day after MH17 was shot down over Ukraine's warring eastern provinces on July 17, 2014, the United States government concluded from available evidence that the plane had been brought down by a Russian-made surface-to-air missile launched from rebel-held territory in eastern Ukraine. American officials said at the time that they believed the missile battery had most likely been provided by Russia to pro-Russian separatists.

Article: https://www.nytimes.com/2019/06/19/opinion/mh17-ukraine-russia-suspects.html

Computational Language and EducAtion Research

C L E A R

Event Linking - Knowledge Graph

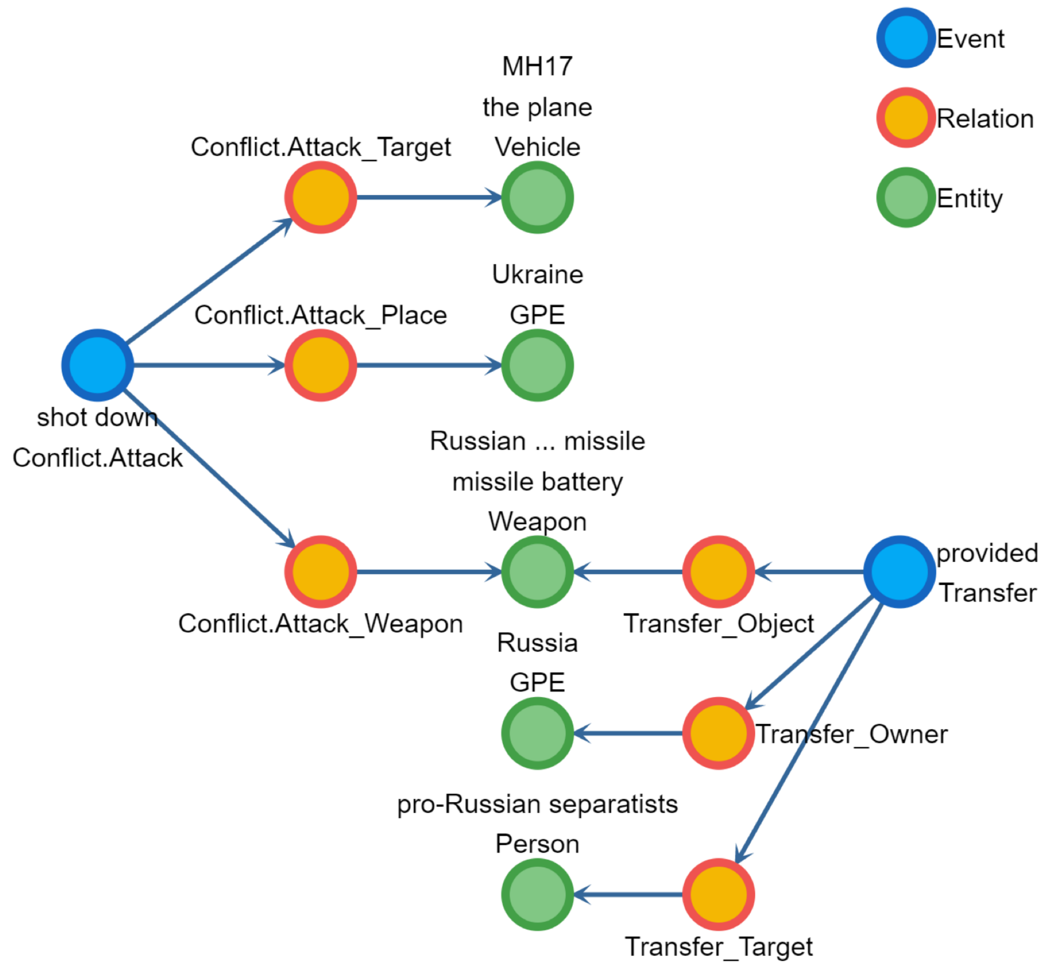

Event Linking as a graph problem

More specifically, a sub-graph isomorphism problem.

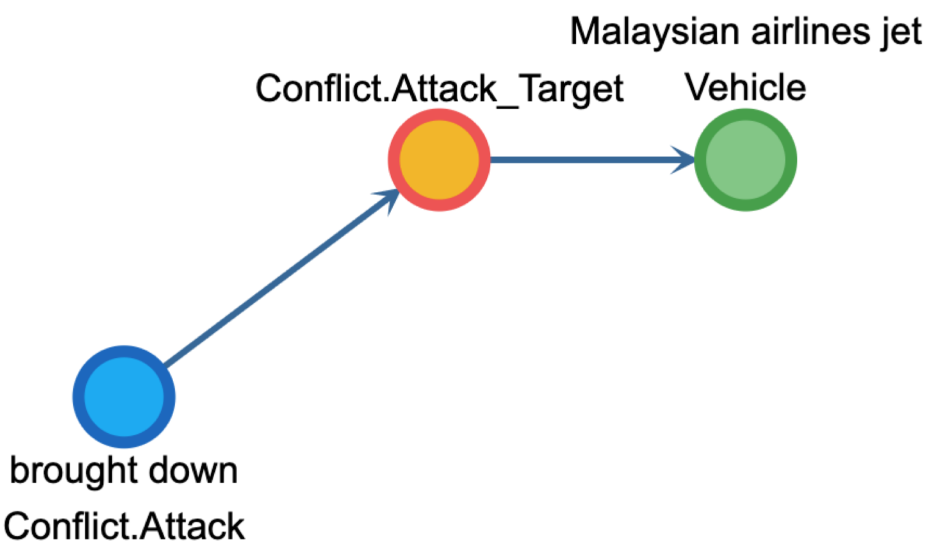

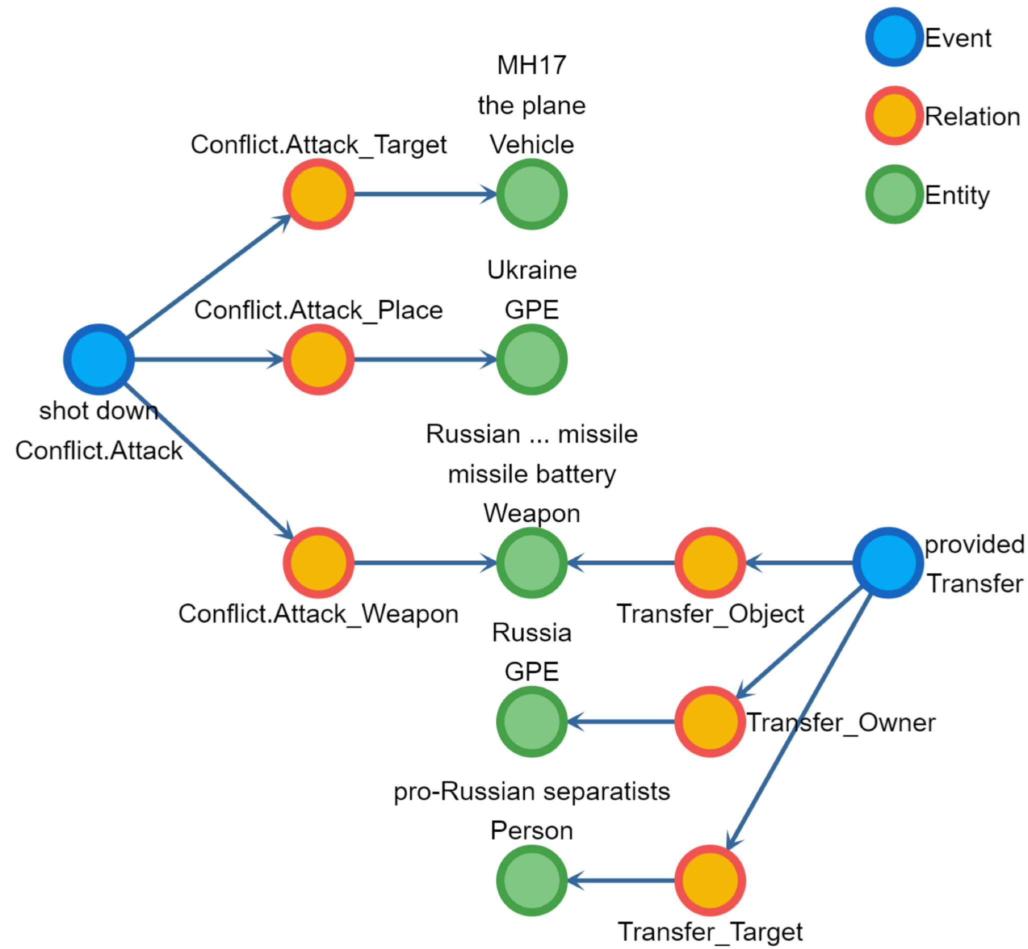

Event Linking as a graph problem

More specifically, a similarity based sub-graph isomorphism problem.

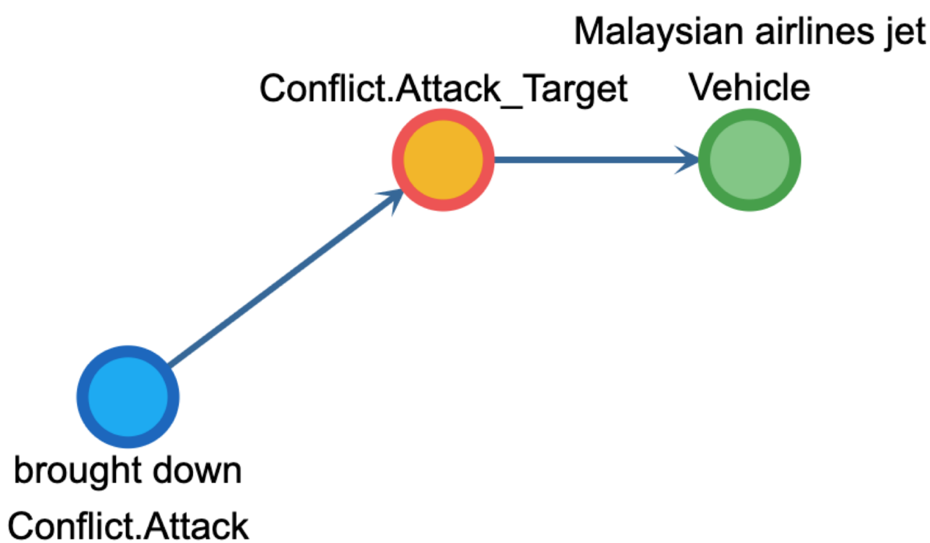

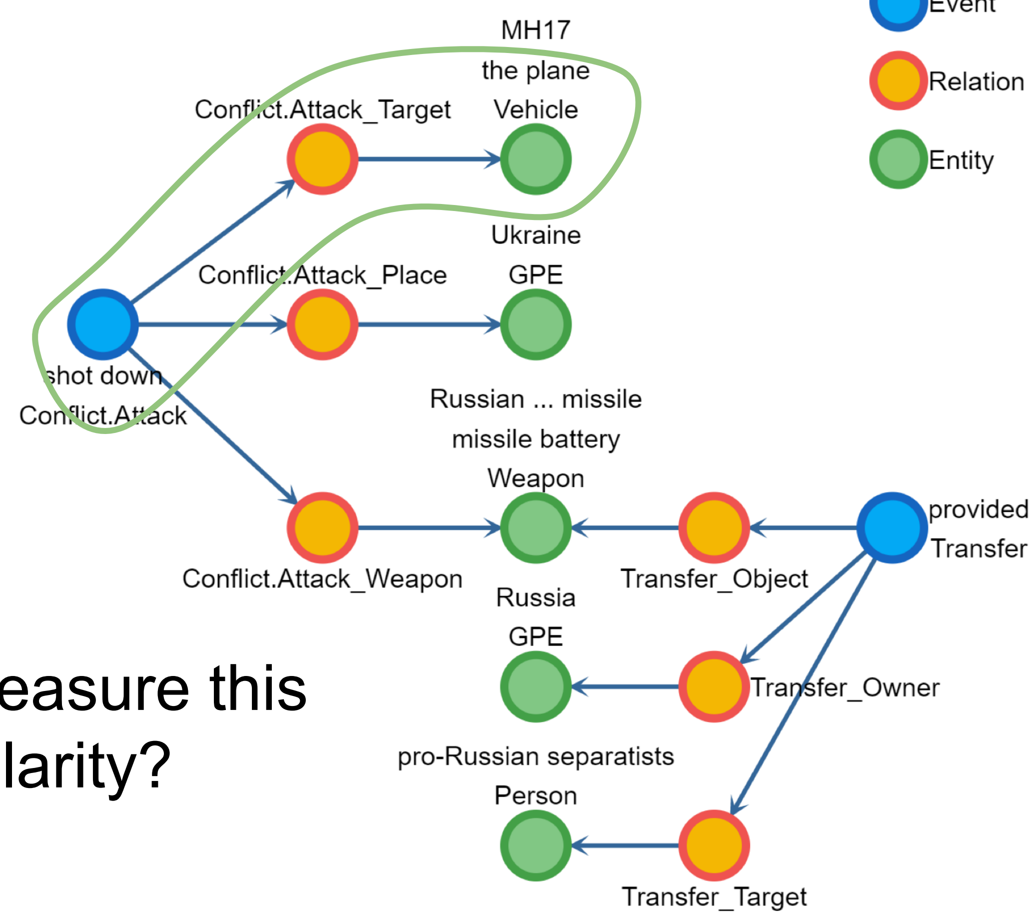

How do we measure this structural similarity?

Link Prediction - TransE (Bordes et al.)

“**Relationships as translations in the embedding space:** In this paper, we introduce TransE, an energy-based model for learning low-dimensional embeddings of entities. In TransE, relationships are represented as translations in the embedding space: if (h, l, t) holds, then the embedding of the tail entity t should be close to the embedding of the head entity h plus some vector that depends on the relationship”

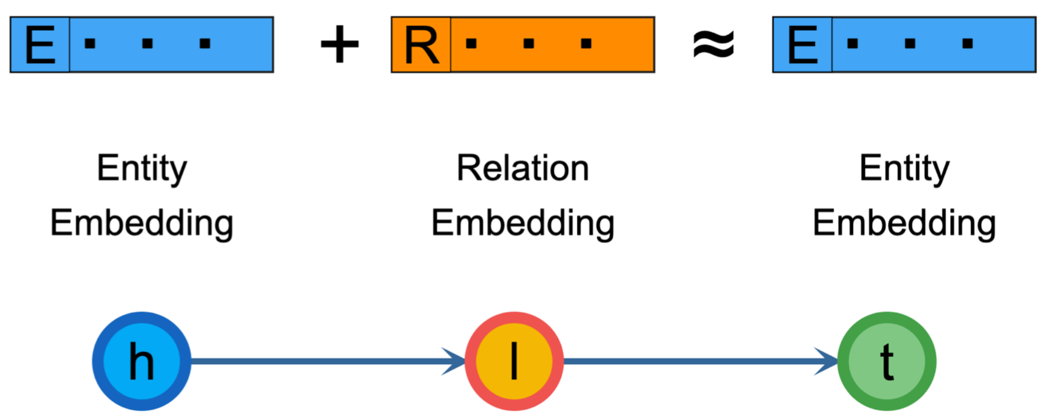

Learning Embeddings with Link Prediction

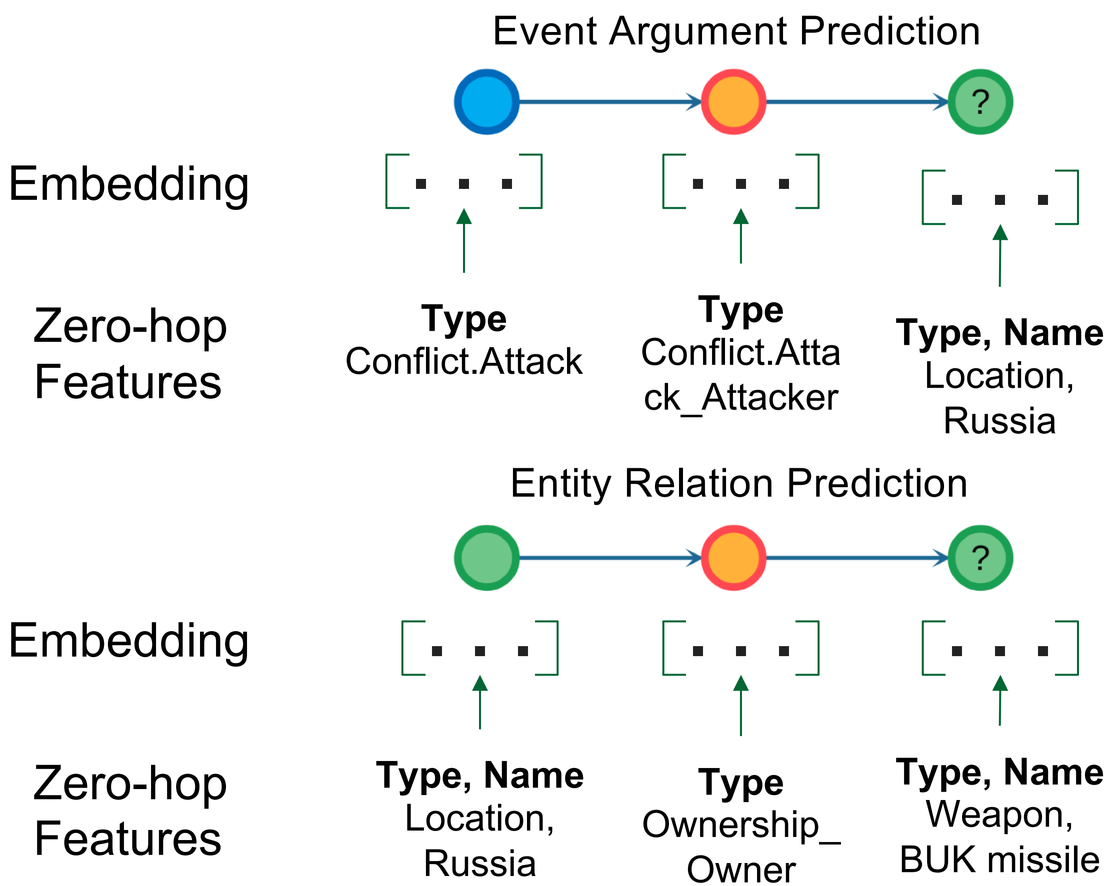

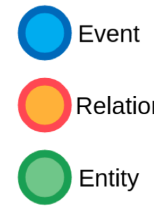

TransE: For each *(h, r, t)* ∈ S, sample *(h', r, t')* ∈ S'. Either corrupted tail, or head, or both.

Minimize Ranking Loss:

$$\sum_{(h,r,t)\in S} \sum_{(h',r,t')\in S'_{(h,r,t)}} \left[\gamma + d(\mathbf{h}+\mathbf{r},\mathbf{t}) - d(\mathbf{h}'+\mathbf{r},\mathbf{t}')\right]_+$$

[1] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In Advances in Neural Information Processing Systems, pages 2787–2795, 2013.

Composing Embeddings

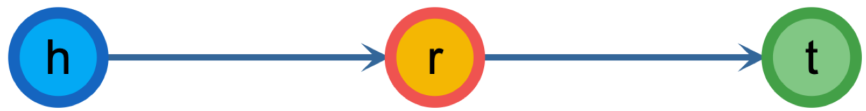

By the TransE architecture, we learn embeddings for (h, r, t) that follows
h + **r** ≈ **t**

Therefore, to compose the embeddings of h (head) and t (tail) that explicitly accounts for the context of the triple we can follow:

Given (h, r, t) ∈ KG:
- Composition(tail) = (**h + r) + t**
- Composition(head) = **h + (t – r)** (since, **h ≈ t – r**)

Composing Embeddings - ECB Example

Document 1 event

Police apprehended Jackson at about 2:30 a.m. and booked him for the misdemeanour before his release , making for a long night with a playoff looming on Sunday at Pittsburgh against the Steelers

Document 2 event

Chargers receiver Vincent Jackson was arrested on suspicion of drunk driving on Tuesday morning five days before a key NFL playoff game

Composing Embeddings - using Blender's AIDA parser

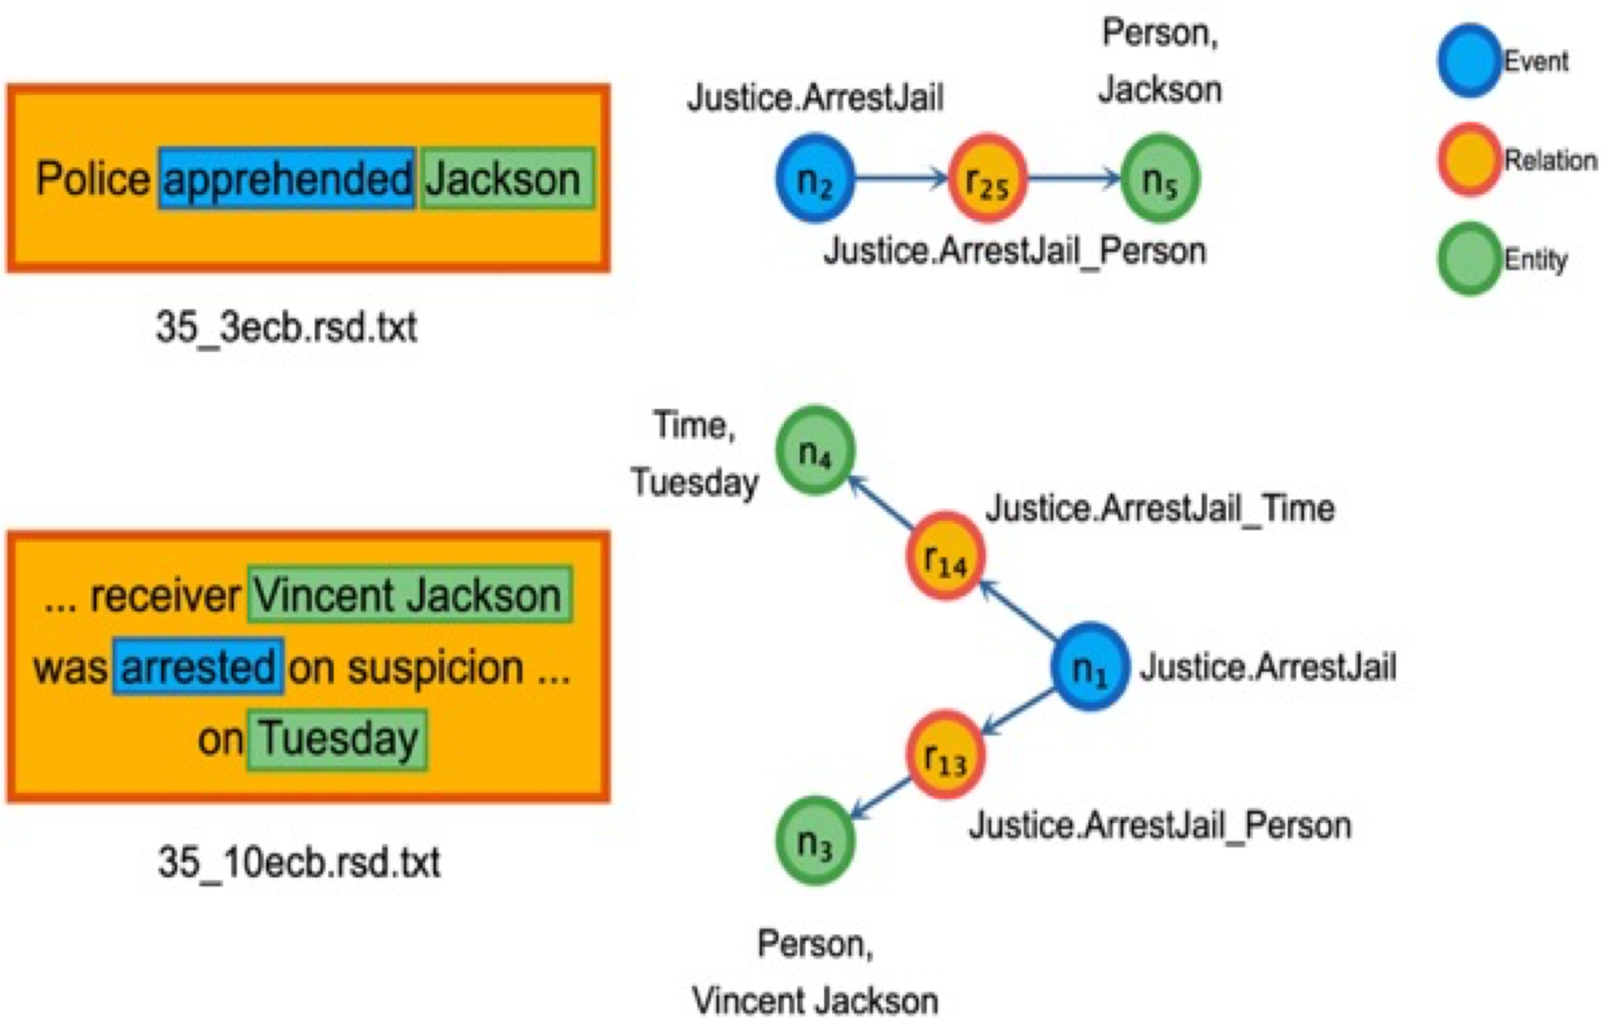

Composing Embeddings - Similarity

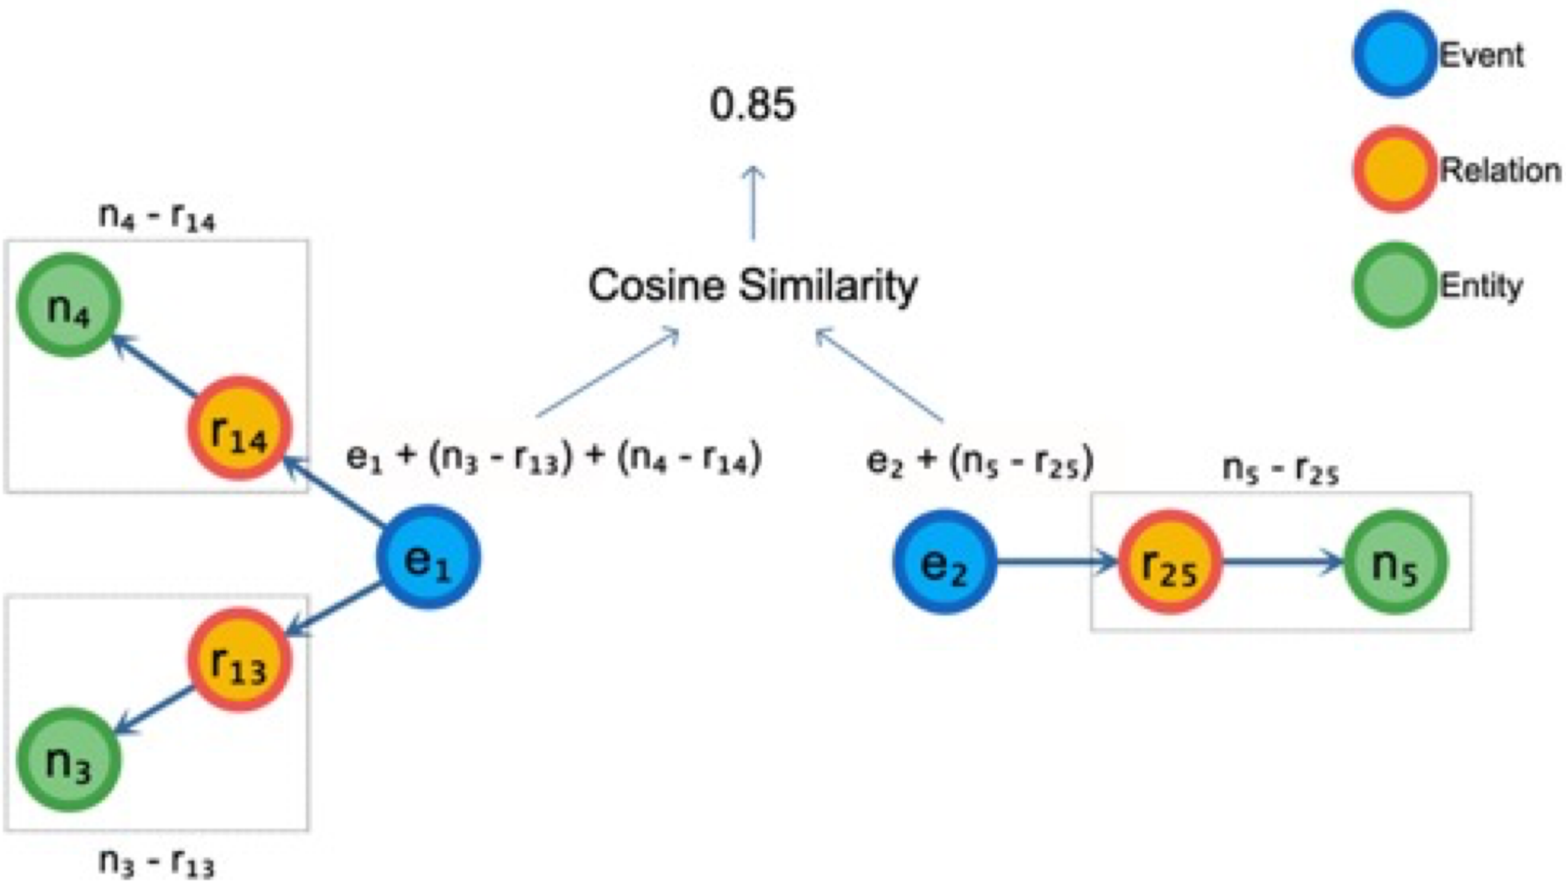

Computational Language and EducAtion Research
CLEAR

Preliminary results for Event Linking on ECB corpus

Method	**BCUB Recall**	**BCUB Precision**	**BCUB F1**	**MUC Recall**	**MUC Precision**	**MUC F1**
TA2 system only	(377 / 886) 42.53%	(852.8 / 886) 96.25%	58.99%	(54 / 529) 10.2%	(54 / 86) 62.79%	17.56%
Graph Embeddings (CC)	(548 / 886) 61.83%	(390 / 886) 44%	51.41%	(270 / 529) **51.03%**	(270 / 512) 52.73%	**51.87%**
Graph Embeddings + TA2 system	(430 / 886) 48.54%	(550 / 886) 62.08%	**54.48%**	(200 / 529) 37.8%	(200 / 412) 48.5%	**42.5%**

Computational Language and EducAtion Research
CLEAR

Future Work

- Linking using graph embeddings
- **Nearest neighbor KB search**
- **Vector similarity**
- Affine mapping between embedding vectors

Nearest Neighbor DB Search

Challenge: Fast scalable approach for identifying co-reference candidates

Solution: Vector representation of DB entries stored in kd-tree

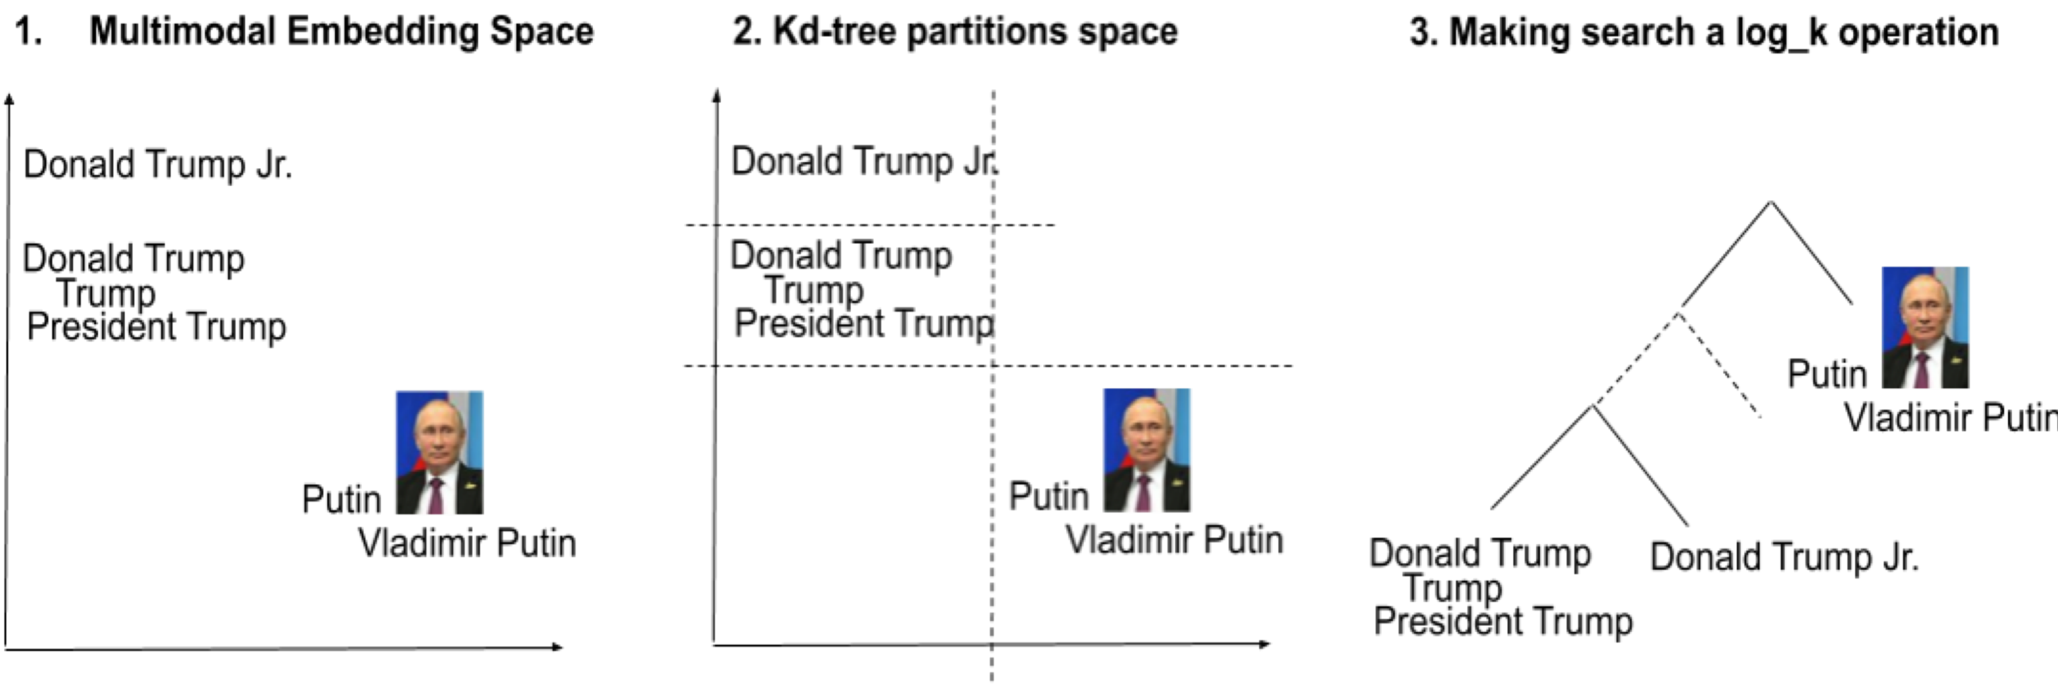

Computational Language and EducAtion Research
CLEAR

Future Work

- Linking using graph embeddings
- Nearest neighbor KB search
- Vector similarity
- **Affine mapping between embedding vectors**

Image Encoding

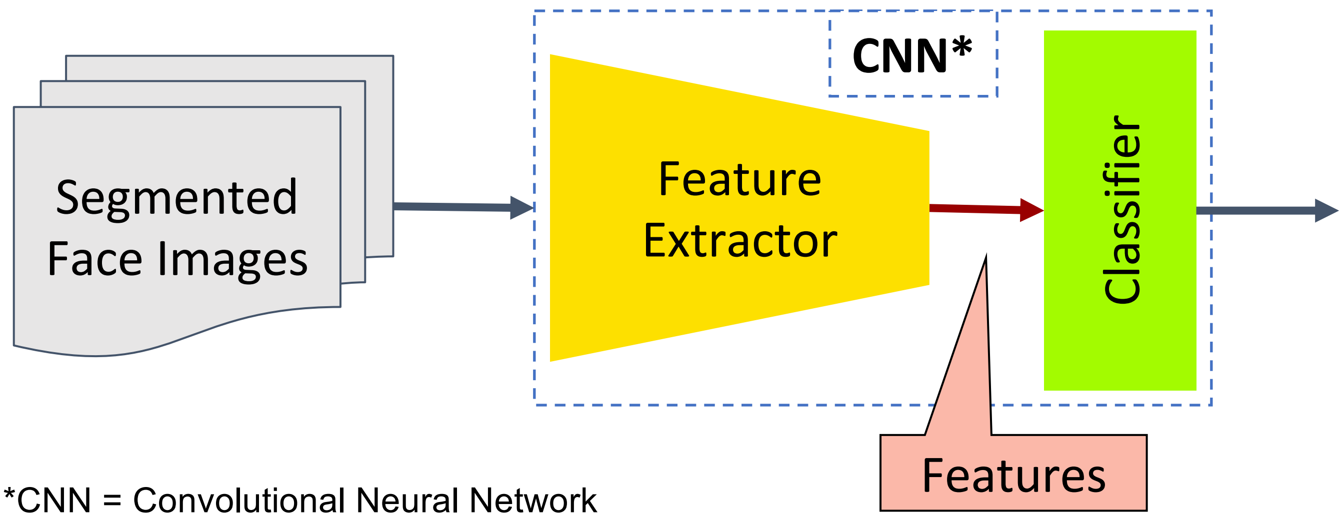

*CNN = Convolutional Neural Network

 ≈  ≈ 

Computational Language and EducAtion Research
CLEAR

Image Encoding

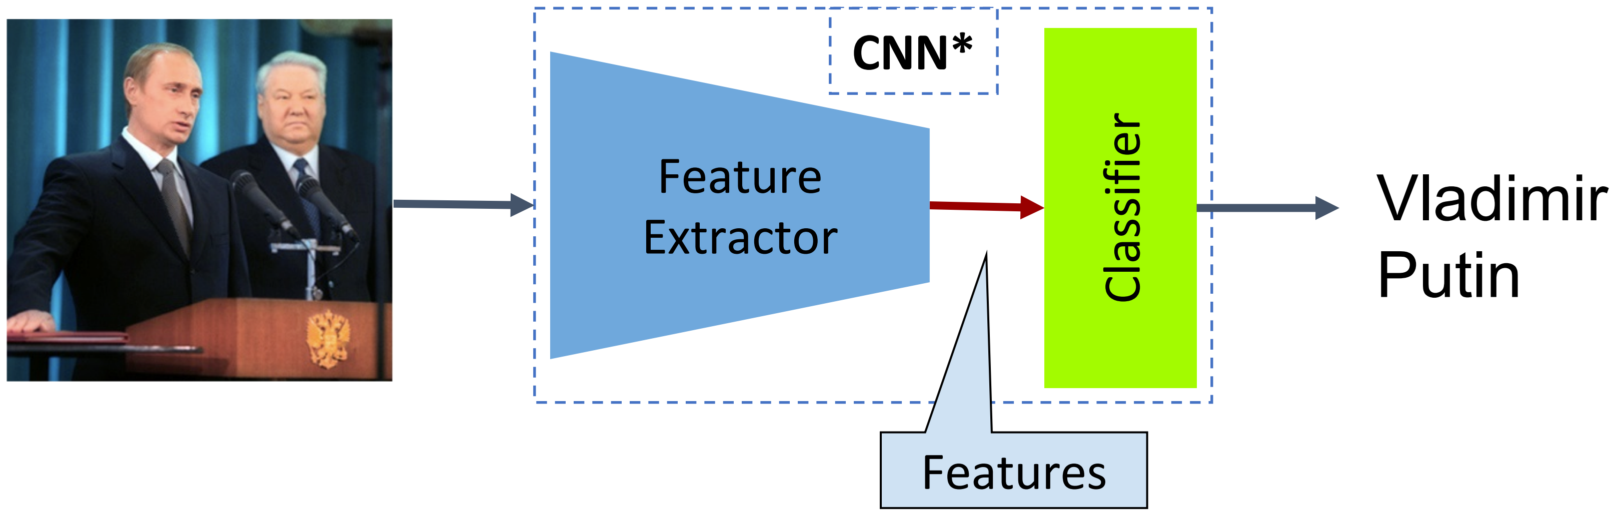

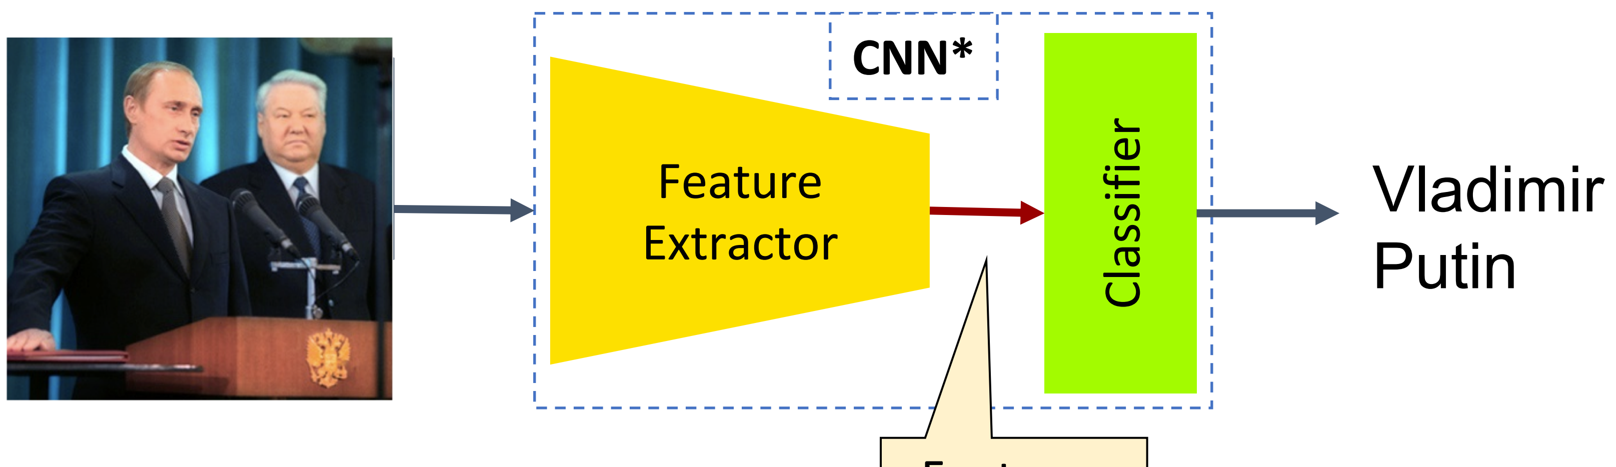

*CNN = Convolutional Neural Network

Image Encoding

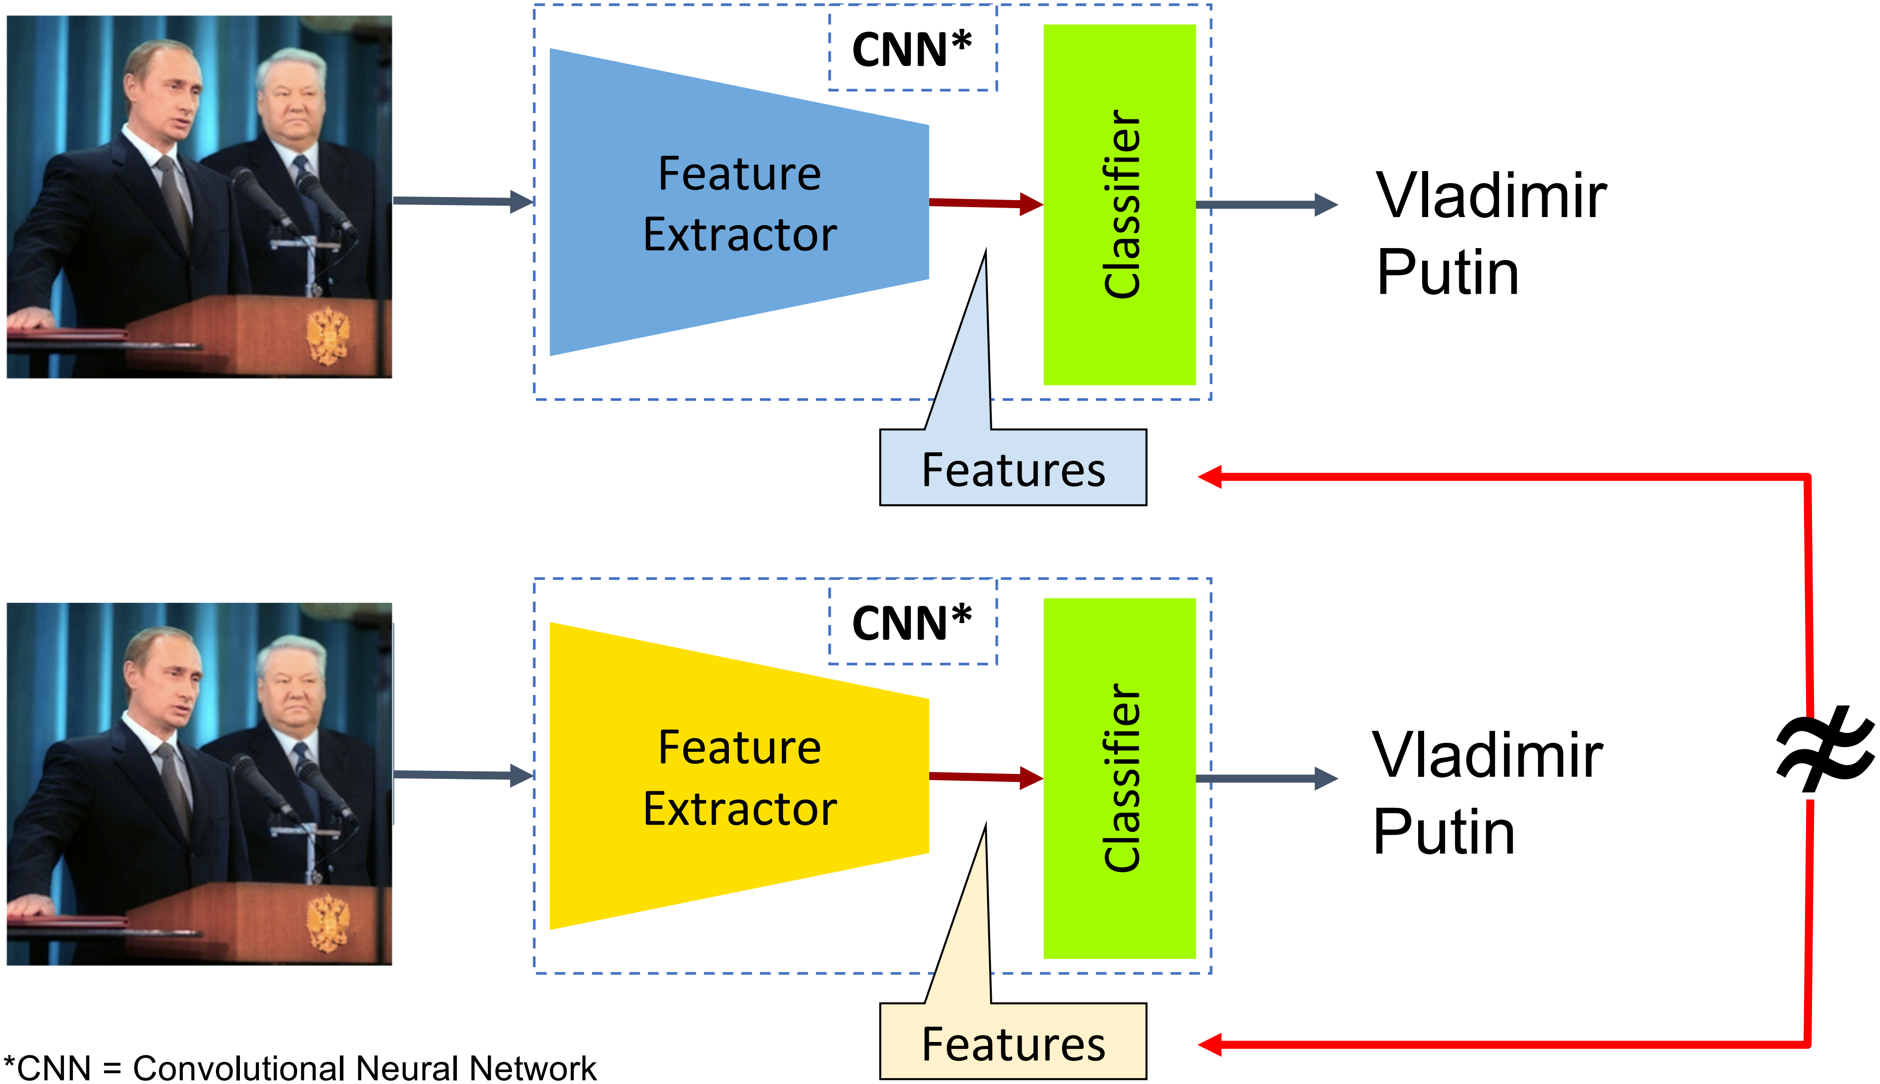

*CNN = Convolutional Neural Network

We establish a mapping between these two features

$$A(x) = M_{B \to A} B(x)$$

Face features

We establish a mapping between these two features

$$A(x) = M_{B \rightarrow A} B(x)$$

Affine Map

Solving for the Affine Mapping

A(x)

B(x)

Minimize the euclidean distance between

$$\mathbf{A(x)}\ \text{and}\ \mathbf{M_{B \rightarrow A} B(x)}$$

Solving for the Affine Mapping

A(x)

B(x)

...

Minimize the euclidean distance between

$$\mathbf{A(x)} \text{ and } \mathbf{M_{B \rightarrow A}B(x)}$$

Cross-TA1 linking with diverse CNN models produces 99% accuracy

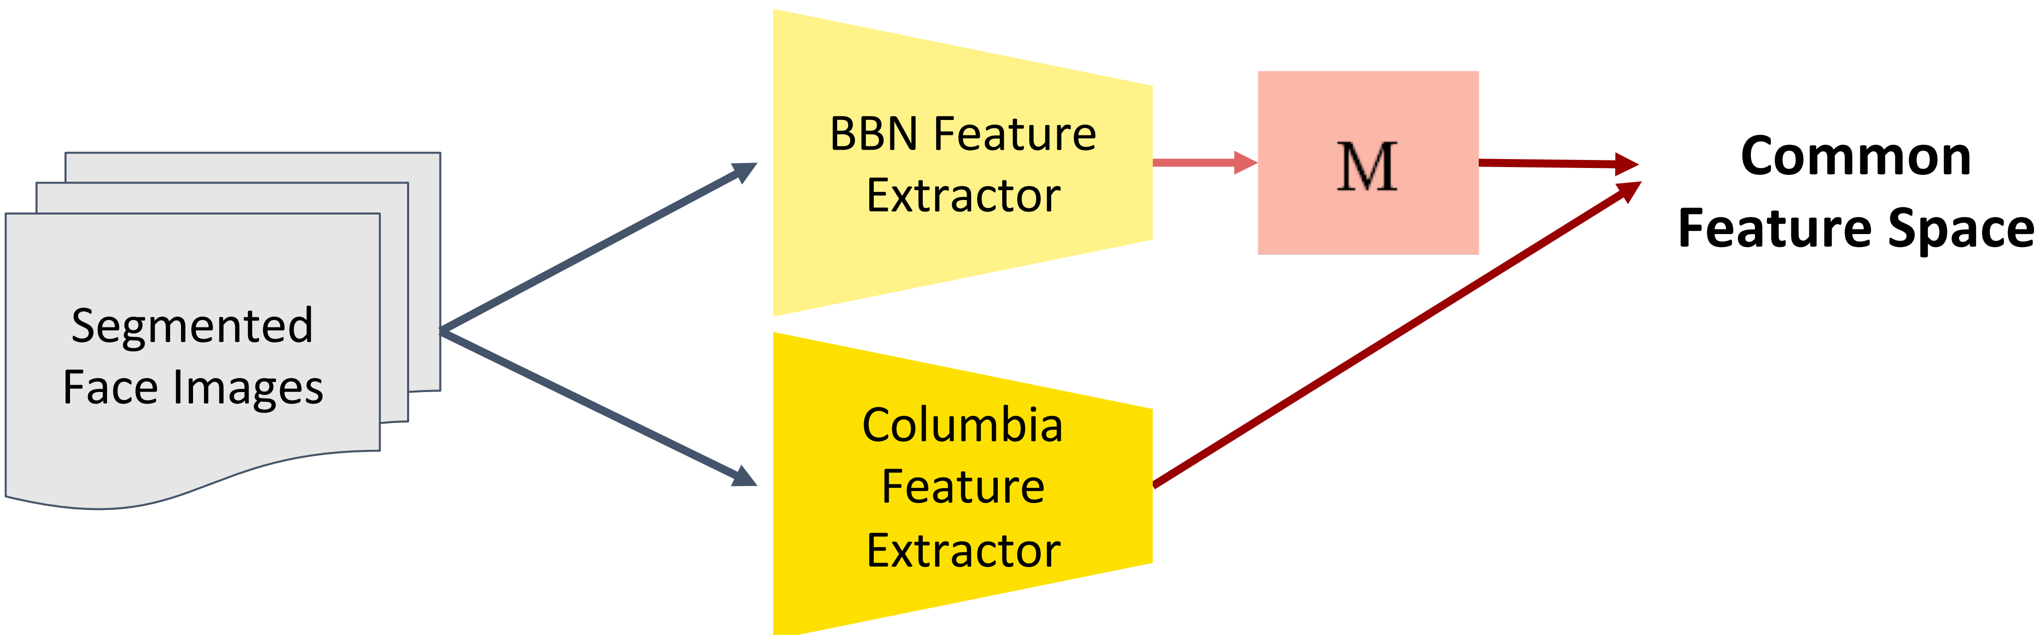

BBN: generated from FaceNet trained on CASIA-WebFace;
Columbia: generated from FaceNet trained on VGGFace2;

Summary

High frame recall is achieved using

- Efficient object manipulation
- Input from multiple TA1s
- Simple linking metrics
- Streaming clustering

Paths to improvement

- Graph embeddings
- Multimodal nearest neighbor KB search
- Affine mapping between vector spaces