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

Virtual Site Visit – April 15, 2020

Martha Palmer, Jim Martin, Chris Heckman,
James Pustejovsky, Ross Beveridge,
Susan W. Brown

Computational Language and EducAtion Research
C L E A R

Our Team Now

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

Outline

- Current Architecture
- AIDA Explorer
- Current Progress and Future Plans
  - Cross-doc coref
  - Affine mappings for images
  - Multi-modal vector representations?
  - Annotating images and video
- Ontology Effort
- Covid-19

Current Architecture

Software Goals

- Combine multiple TA1’s
- Compact KB
- Streaming architecture
- AIF Compliance
- Metric Evaluation
- Collaboration with TA1 & TA2

Software Engineering - Ontology Objects

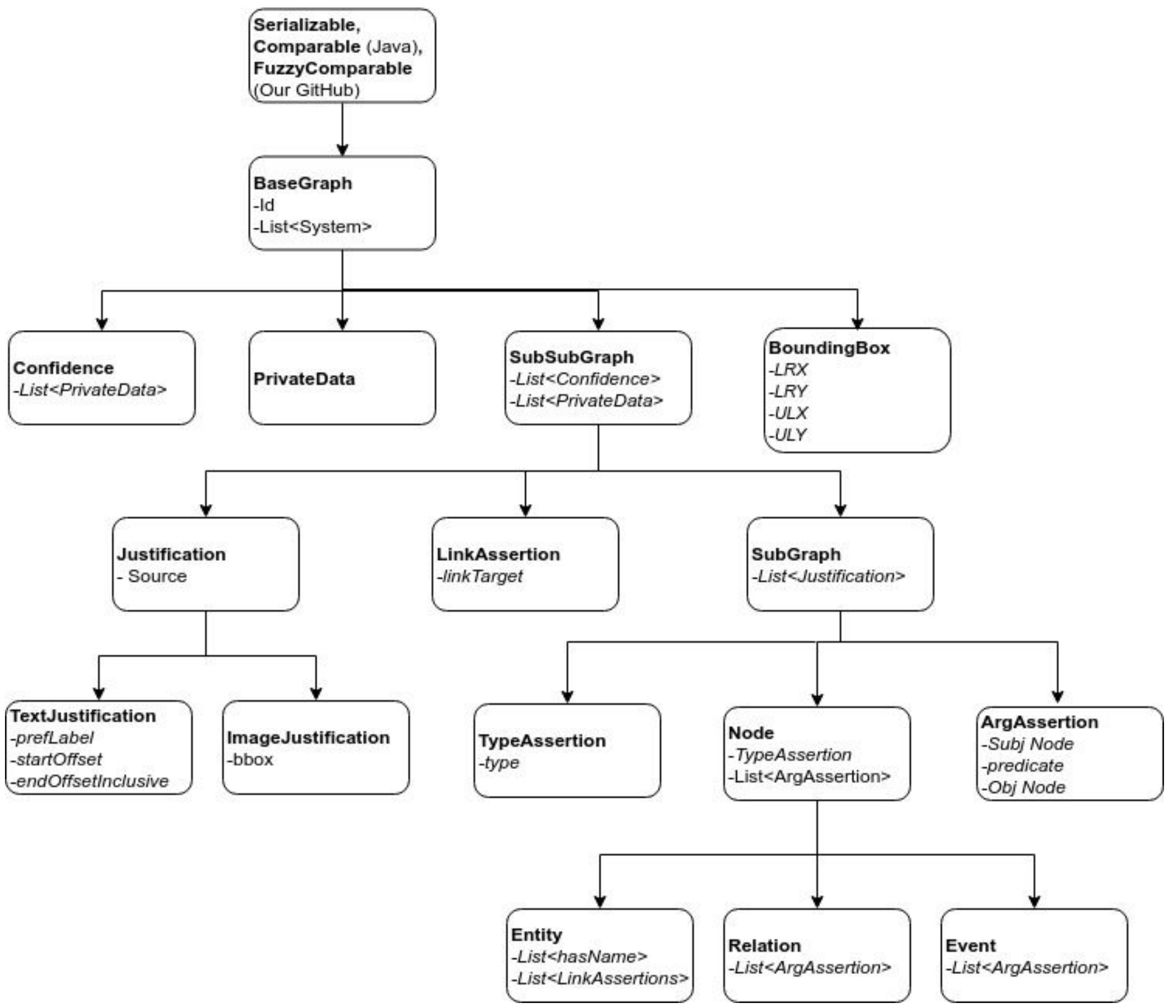

- Read()
- Compare()
- Merge()
- Write()

Diagram

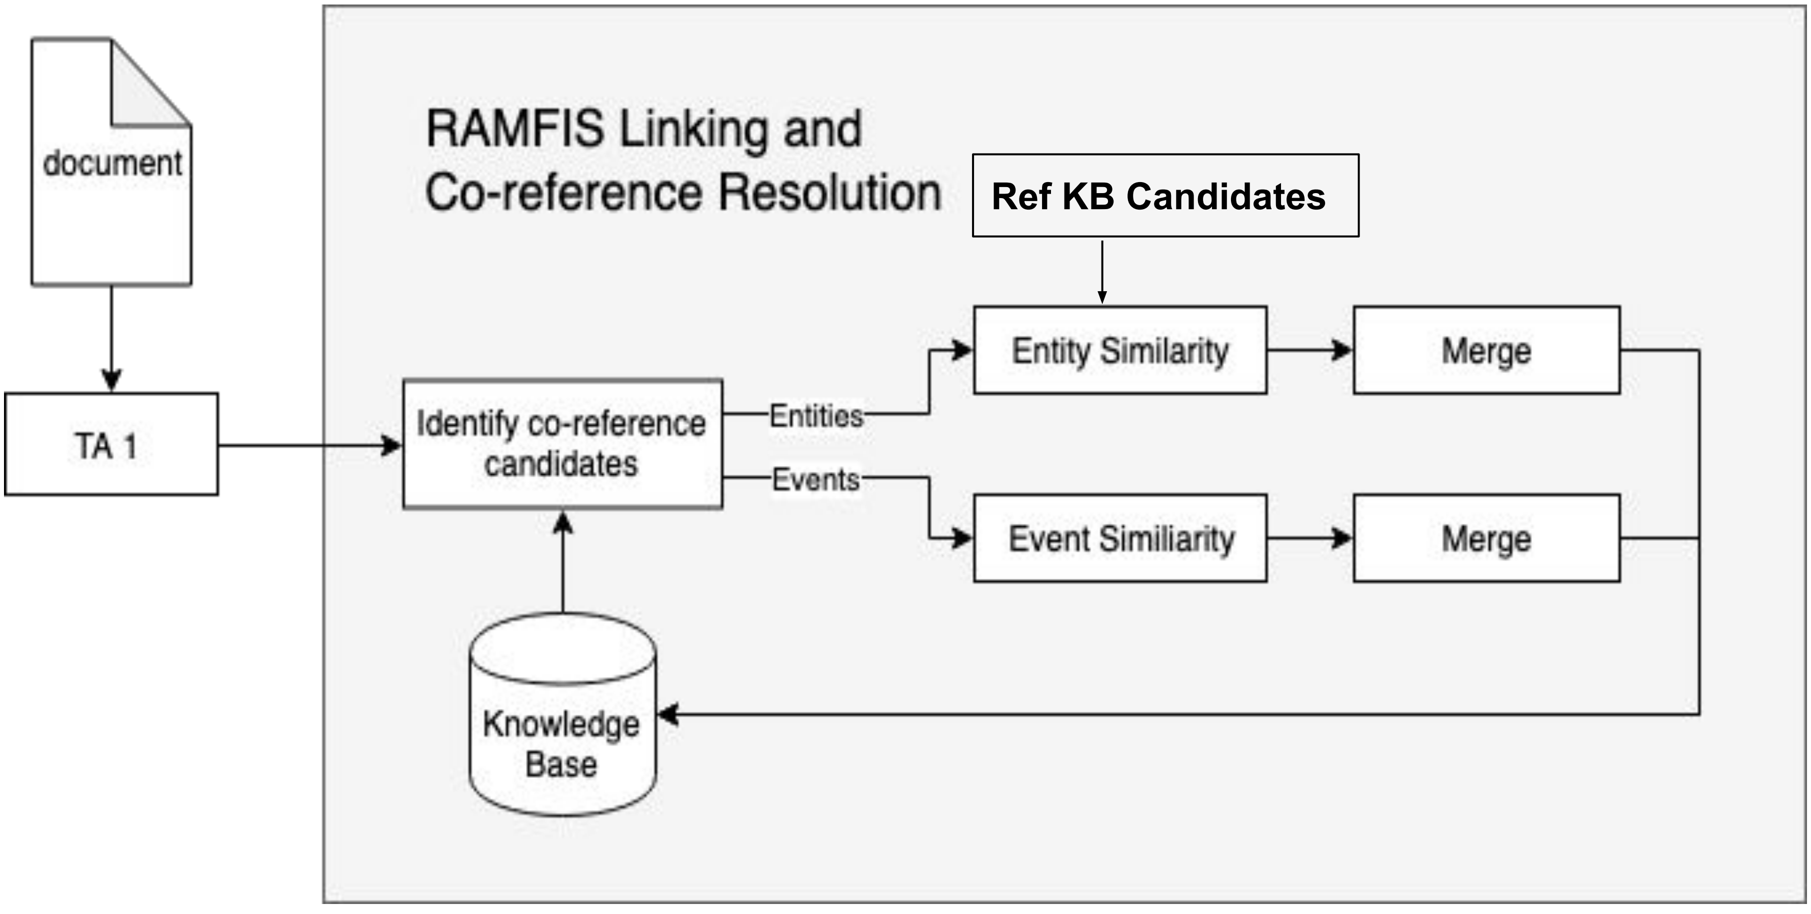
document
TA 1
RAMFIS Linking and
Co-reference Resolution
Ref KB Candidates
Identify co-reference candidates
Entities
Events
Entity Similarity
Event Similiarity
Merge
Merge
Knowledge Base

Similarity Criteria

Entities

- Type matching
- Fuzzy Name matching
- Justification overlap

Events

- Type matching
- Participant matching
- Justification overlap

Computational Language and EducAtion Research
CLEAR

Baseline coref scores on annotated datasets (cross-doc)

ECB* Data - scores for the common nodes

	Gold standard	TA1 output	Common	B^3 P	B^3 R	B^3 F1	MUC P	MUC R	MUC F1
Events	3437	5107	918	95.92	42.75	59.14	63.04	10.96	18.67
Entities	4268	8820	864	98.09	64.33	77.7	95.08	54.2	69.04
Combined	7705	13927	1782	95.75	57.05	71.5	54.71	10.96	18.26

RED** Data - B^3 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

*Event Coref Bank
** DEFT Richer Event Descriptions

Computational Language and EducAtion Research
CLEAR

Unsequestered data results

Entities

Method	B^3 Recall	B^3 Precision	B^3 F1	MUC Recall	MUC Precision	MUC F1
TA2 system /wo ref-kb linking	(2287 / 5108) 44.77%	(4969 / 5108) 97.29%	61.32%	(284 / 2867) 9.9%	(284 / 414) 68.59%	17.31%
TA2 system with ref-kb linking	(2634 / 5108) 51.56%	(4585 / 4838) 94.76%	**66.78%**	**(1447 / 2867) 50.47%**	(1447 / 2867) 50.47%	**64.22%**

Events

Method (Event Linking)	B^3 Recall	B^3 Precision	B^3 F1	MUC Recall	MUC Precision	MUC F1
TA2 system /wo ref-kb linking	(498 / 1259) 39.62%	(1138 / 1258) 90.49%	55.11%	(60 / 771) 7.78%	(60 / 101) 59.4%	13.76%
TA2 system with ref-kb linking	(502/1259) 39.88%	(1104 / 1254) 88.07%	54.9%	(148 / 771) 19.19%	(148 / 282) 52.48%	**28.11%**

Benefit of Ref-KB Linking - Cross Lingual Clusters

```
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Ades" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Gorad Adehsa"
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "ODS" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odesa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odess" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessa osh" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessae" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odesse" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessos" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessus" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odessza" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Odissos" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "Udessa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "ao de sa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "awdsa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "awdysa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "odesa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , "odessa" ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Горад Адэса ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одеса ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одесс ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одессæ ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одесса ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одесса ош ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одессе ,
"http://www.verbs.colorado/refkb/LDC2019E43/698740" , Одессы ,
```


Computational Language and EducAtion Research
CLEAR

Collaboration with TA1 and TA3

ISI

- Share results of event co-ref and similarity metrics

Texas

- Provide context vectors for entities
- Time
- Expanded entry points using Explorer

Work in Progress: Colorado

1. Event, entity and relation embeddings for improved co-ref
   - Learning through link prediction
   - BERT Context for Co-reference
2. Exploiting knowledge graph structure
   - Using graph neural networks to capture wider context
3. Scalable Approaches to Nearest Neighbor Search
   - Vector-based indexing of entities

Learning Embeddings with Link Prediction - TransE, CharTransE

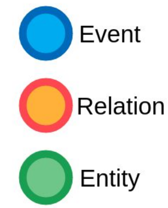

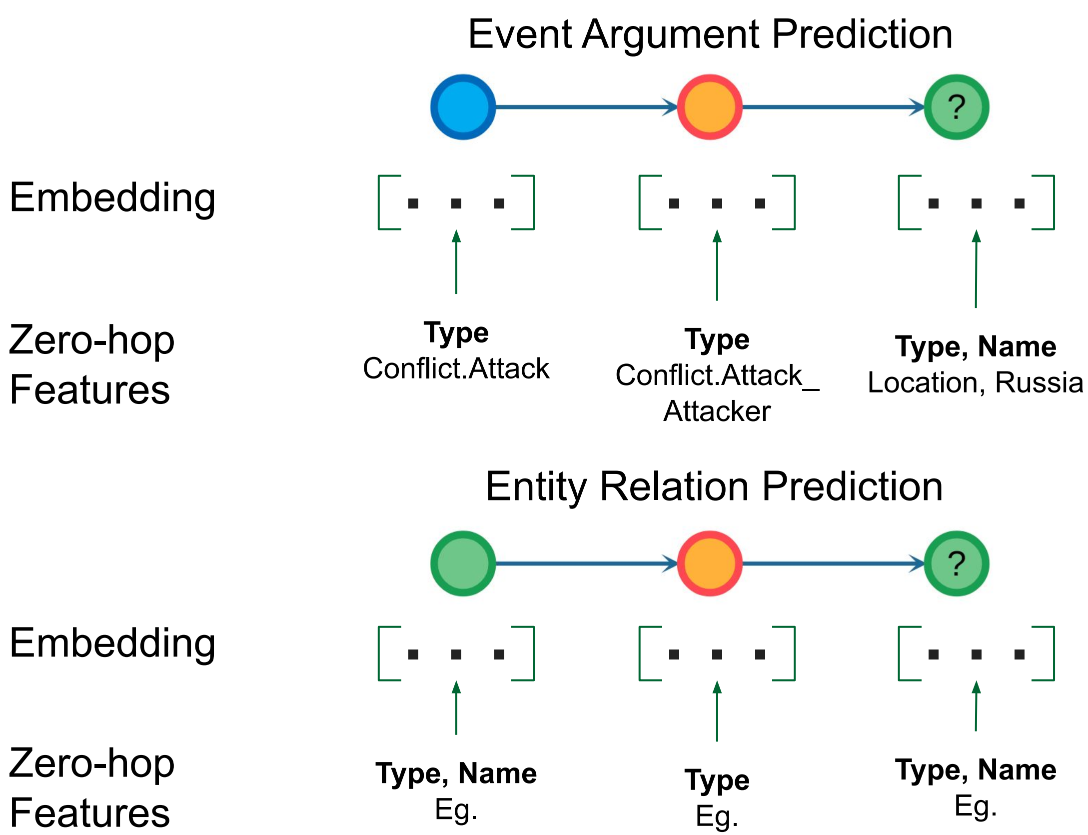

TransE: For each *(h, r, t)* ∈ S, sample *(h', r, t')* ∈ S'. Either corrupted tail, or head, or both.

Minimize Ranking Loss:

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

[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.

Preliminary results for Event Linking

ECB Data

	B^3 P	B^3 R	B^3 F1	MUC P	MUC R	MUC F1
CharTransE embs	60.43	**47.47**	53.17	53.14	**38.3**	**44.51**
TA2 events	95.92	42.75	59.14	63.04	10.96	18.67

Computational Language and EducAtion Research
CLEAR

Representations by Graph Aggregation - GraphTransE

Graph Neural Networks: GNNs use the graph structure and node features X_v to learn a representation vector of a node, h_v, or the entire graph, h_G

Representations are updated by aggregating k-hop neighborhood of a node as following:

$$a_v^{(k)} = \mathrm{AGGREGATE}^{(k)}\left(\left\{h_u^{(k-1)} : u \in \mathcal{N}(v)\right\}\right)$$

$$h_v^{(k)} = \mathrm{COMBINE}^{(k)}\left(h_v^{(k-1)}, a_v^{(k)}\right)$$

N(v) is the set of nodes adjacent to *v*

[2] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? CoRR, abs/1810.00826, 2018. URL http://arxiv.org/abs/1810.00826.

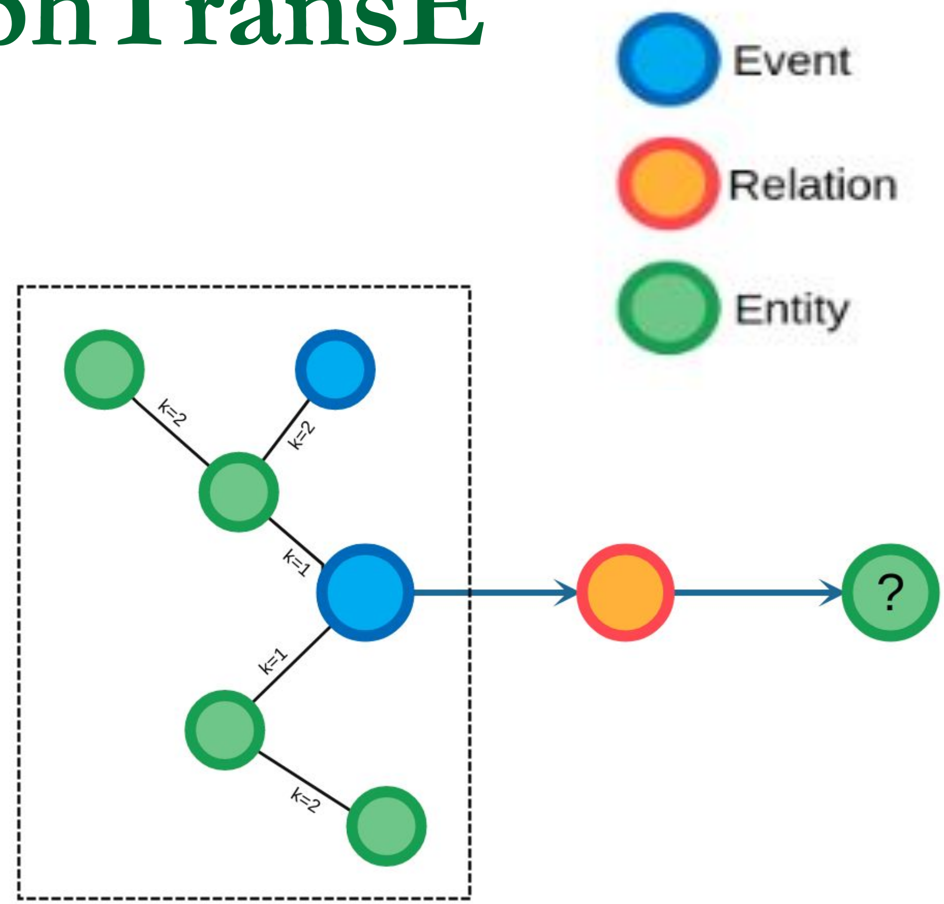

2-hop aggregation

GraphTransE - Composing Embeddings

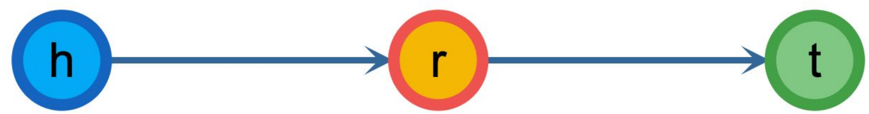

By the TransE architecture, we learn embeddings for (h, r, t) that follows **h + r ≈ t**

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

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

Clustering Techniques

HDBSCAN

1. Hierarchical Density-based spatial clustering of applications with noise
2. Non parametric

Incremental Clustering

1. Incrementally build the clusters by averaging the vector of the cluster upon merge
2. Pairwise comparisons of all the mentions are done
3. Threshold similarity = average similarity between coreferrent pairs

Results with Graph Embeddings

Method (Event Linking)	B^3 *F1*	*MUC F1*	*CEAFE F1*	***CONNL F1 (Average)***	*BLANC F1*	*# singleton clusters*	*# clusters size >= 2*
Single Cluster	5.76%	75.99%	0.08%	27.27%	2.8%	0	1
Same Subtype	32.91%	64.27%	14.3%	37.16%	61.57%	16	66
All Singletons	55.86%	0%	51.42%	35.76%	49.26%	1259	0
Random Inc	49.4%	36.94%	43.03%	43.12%	53.08%	511	117
TA2 system /wo ref-kb linking	55.11%	13.76%	52.39%	40.42%	50.1%	1110	48
TA2 system with ref-kb linking	54.9%	28.11%	54.3%	45.77%	50.44%	807	165
TransE HDBSCAN	49.66%	11.75%	46.76%	36.05%	49.64%	773	66
TransE Incremental	54.9%	1.4%	50.75%	35.68%	49.29%	1100	78
CharTransE HDBSCAN	39.34%	**64.36%**	45.35	51.85%	61.41%	110	68
CharTransE Incremental	32.94%	**64.27%**	14.31%	37.17%	61.58%	16	66
GraphTransE HDBSCAN	52.5%	53.7%	49.25	51.81%	55.94%	520	39
GraphTransE Incremental	**55.53%**	43.87%	**54.59%**	51.33%	57.87%	697	93

# events mentions	1259
# singleton clusters from anno	427
# clusters size >= 2 from anno	61

Computational Language and EducAtion Research
CLEAR

BERT Embeddings for Linking

- Create our own embeddings for the event and entity mentions for the LDC Unsequestered data
- Run similar experiments on these vectors
- This way we don't have to wait for the TA1 vectors to test the effectiveness of BERT

Nearest Neighbor DB Search

Challenge: Fast scalable approach for identifying co-reference candidates

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

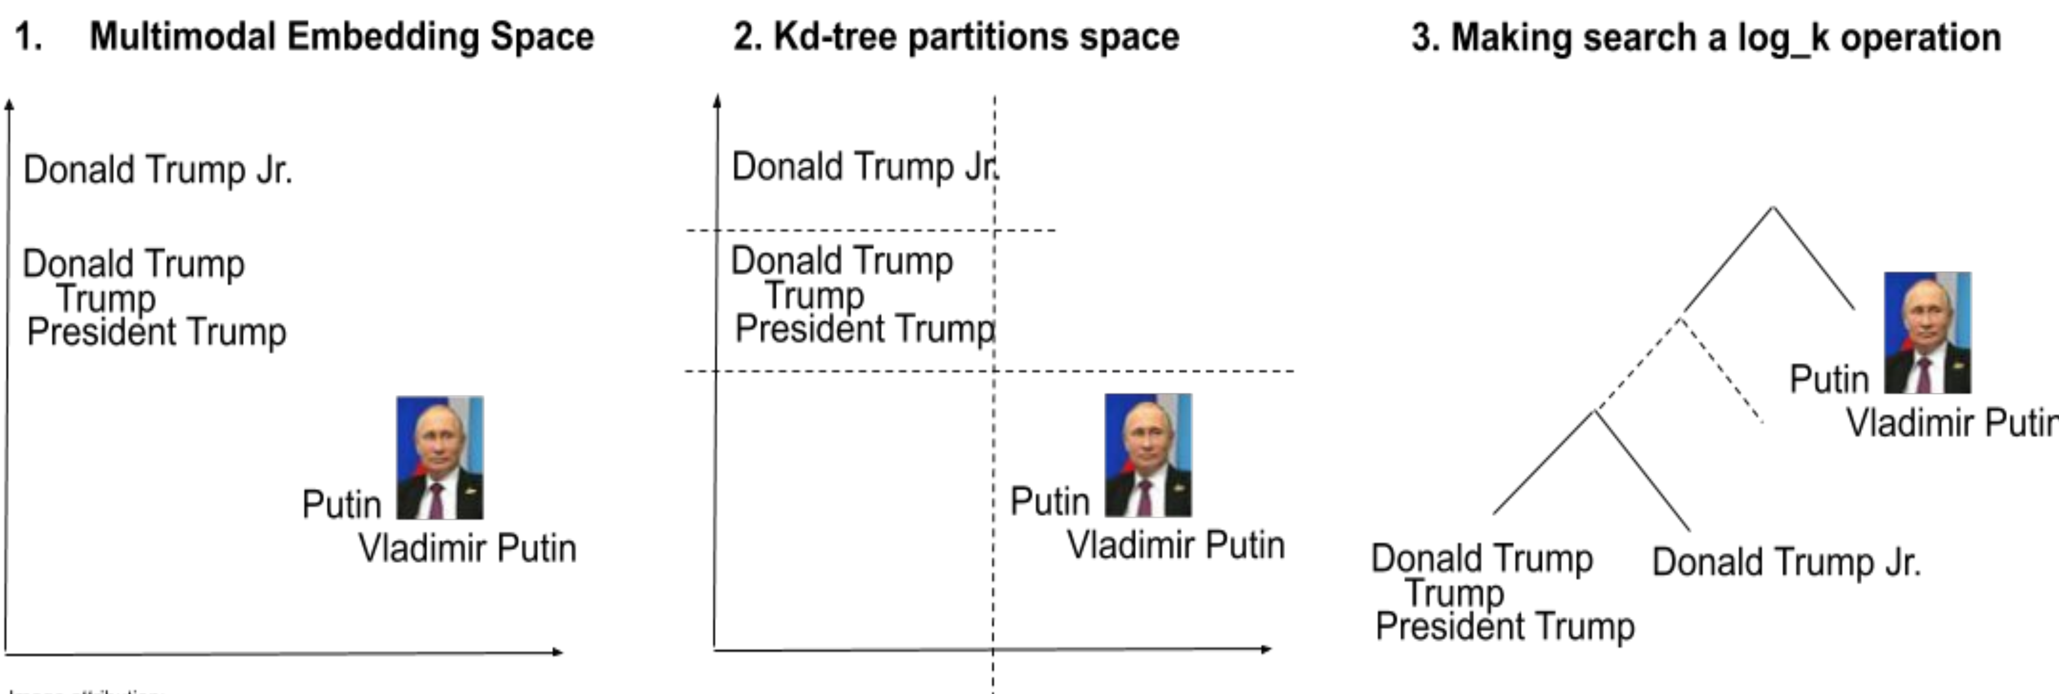

Image attribution:
Kremlin.ru [CC BY 4.0 (https://creativecommons.org/licenses/by/4.0)]

Computational Language and EducAtion Research
C L E A R

Brandeis TA2 Explorer

Brandeis University
Peter Anick, James Pustejovsky, Nikhil Krishnaswamy

Computational Language and EducAtion Research
CLEAR

Brandeis TA2 Explorer Goals

- Browse a TA2 knowledge graph without knowledge of underlying graph structure or query language
- Simple user interface for examining events, relations, and entities
- Lightweight back end optimized for most useful inspection/debugging needs

- V1: Browser for *events, relations* and *role fillers*
- V2: Extensions for examining *entities* and *coreference clusters*

Entities and coreference clusters

- An entity may be referred to in text by many names
- Preferred name within a coreference cluster = “handle”
- Coreference clusters can have mentions in multiple documents
- Mentions may fall into different (but usually hierarchically compatible) ontological categories
- Different coreference clusters may have the same handle
- Debugging:
  - Detecting incompatible members of a cluster
  - Finding independent entity clusters that should be joined

Two ways to explore entities

1. As fillers of events/relations

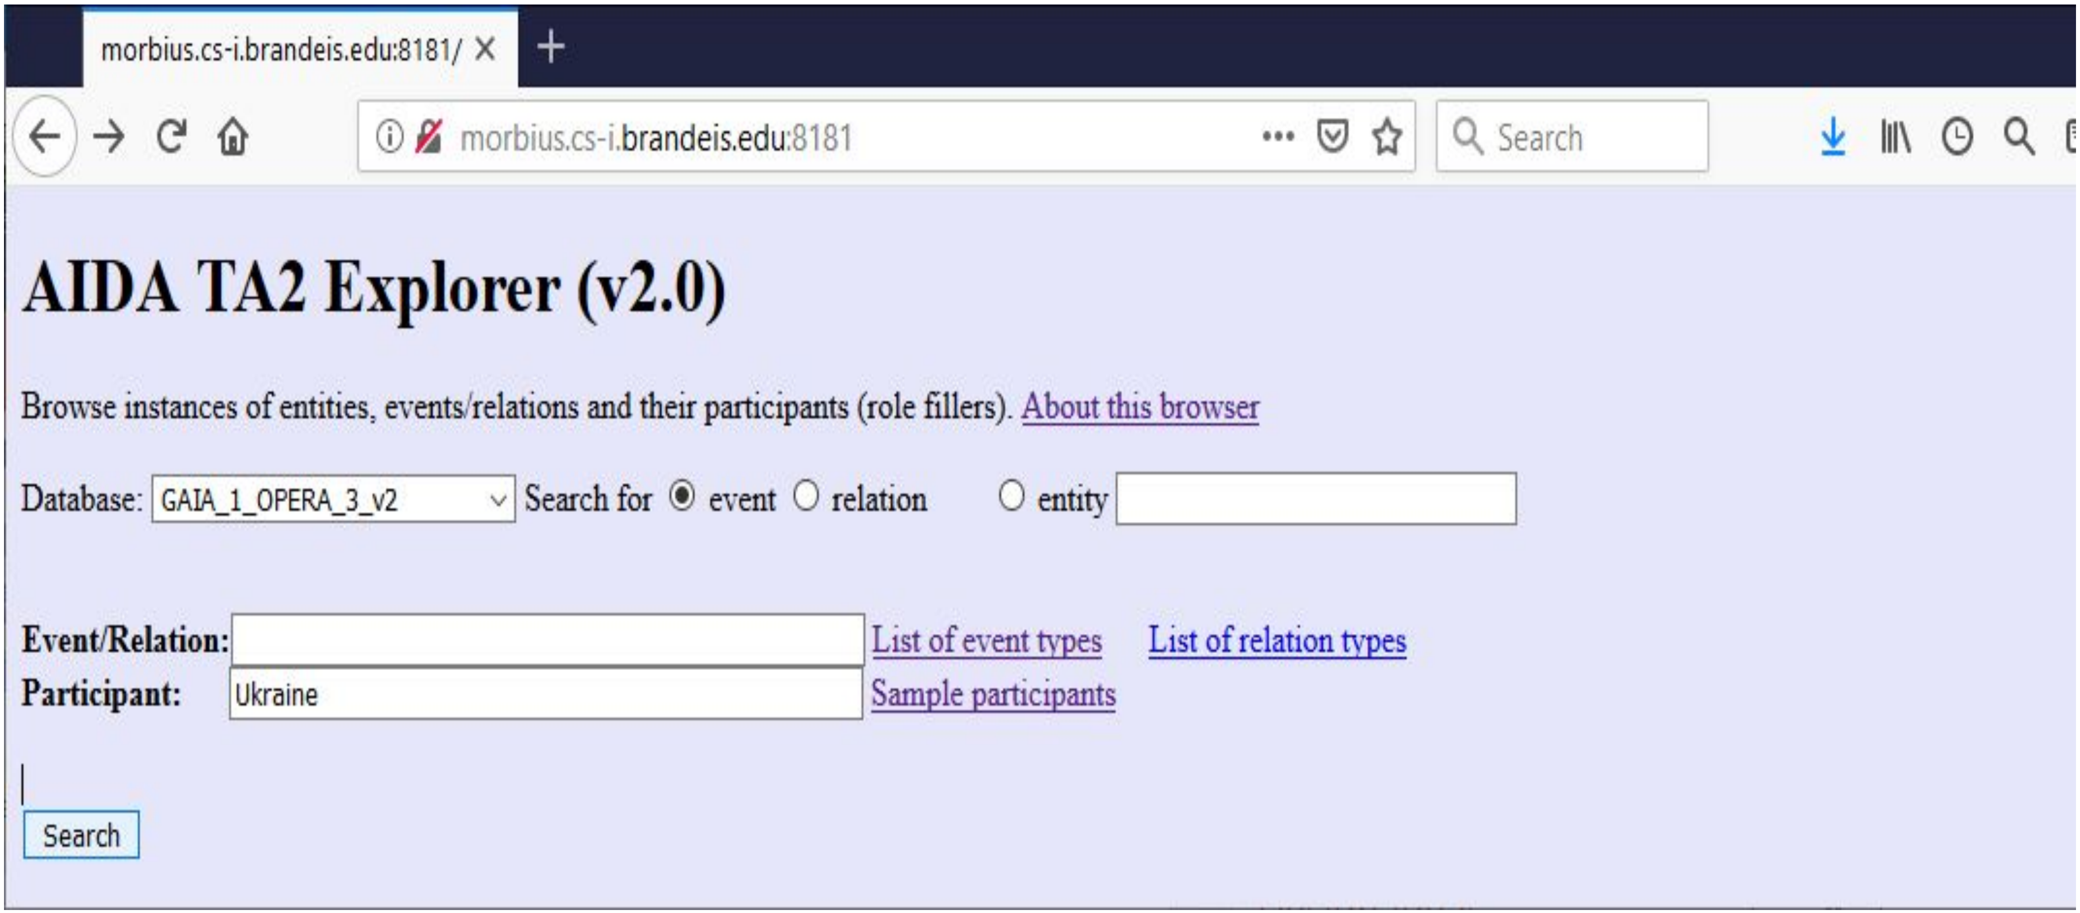

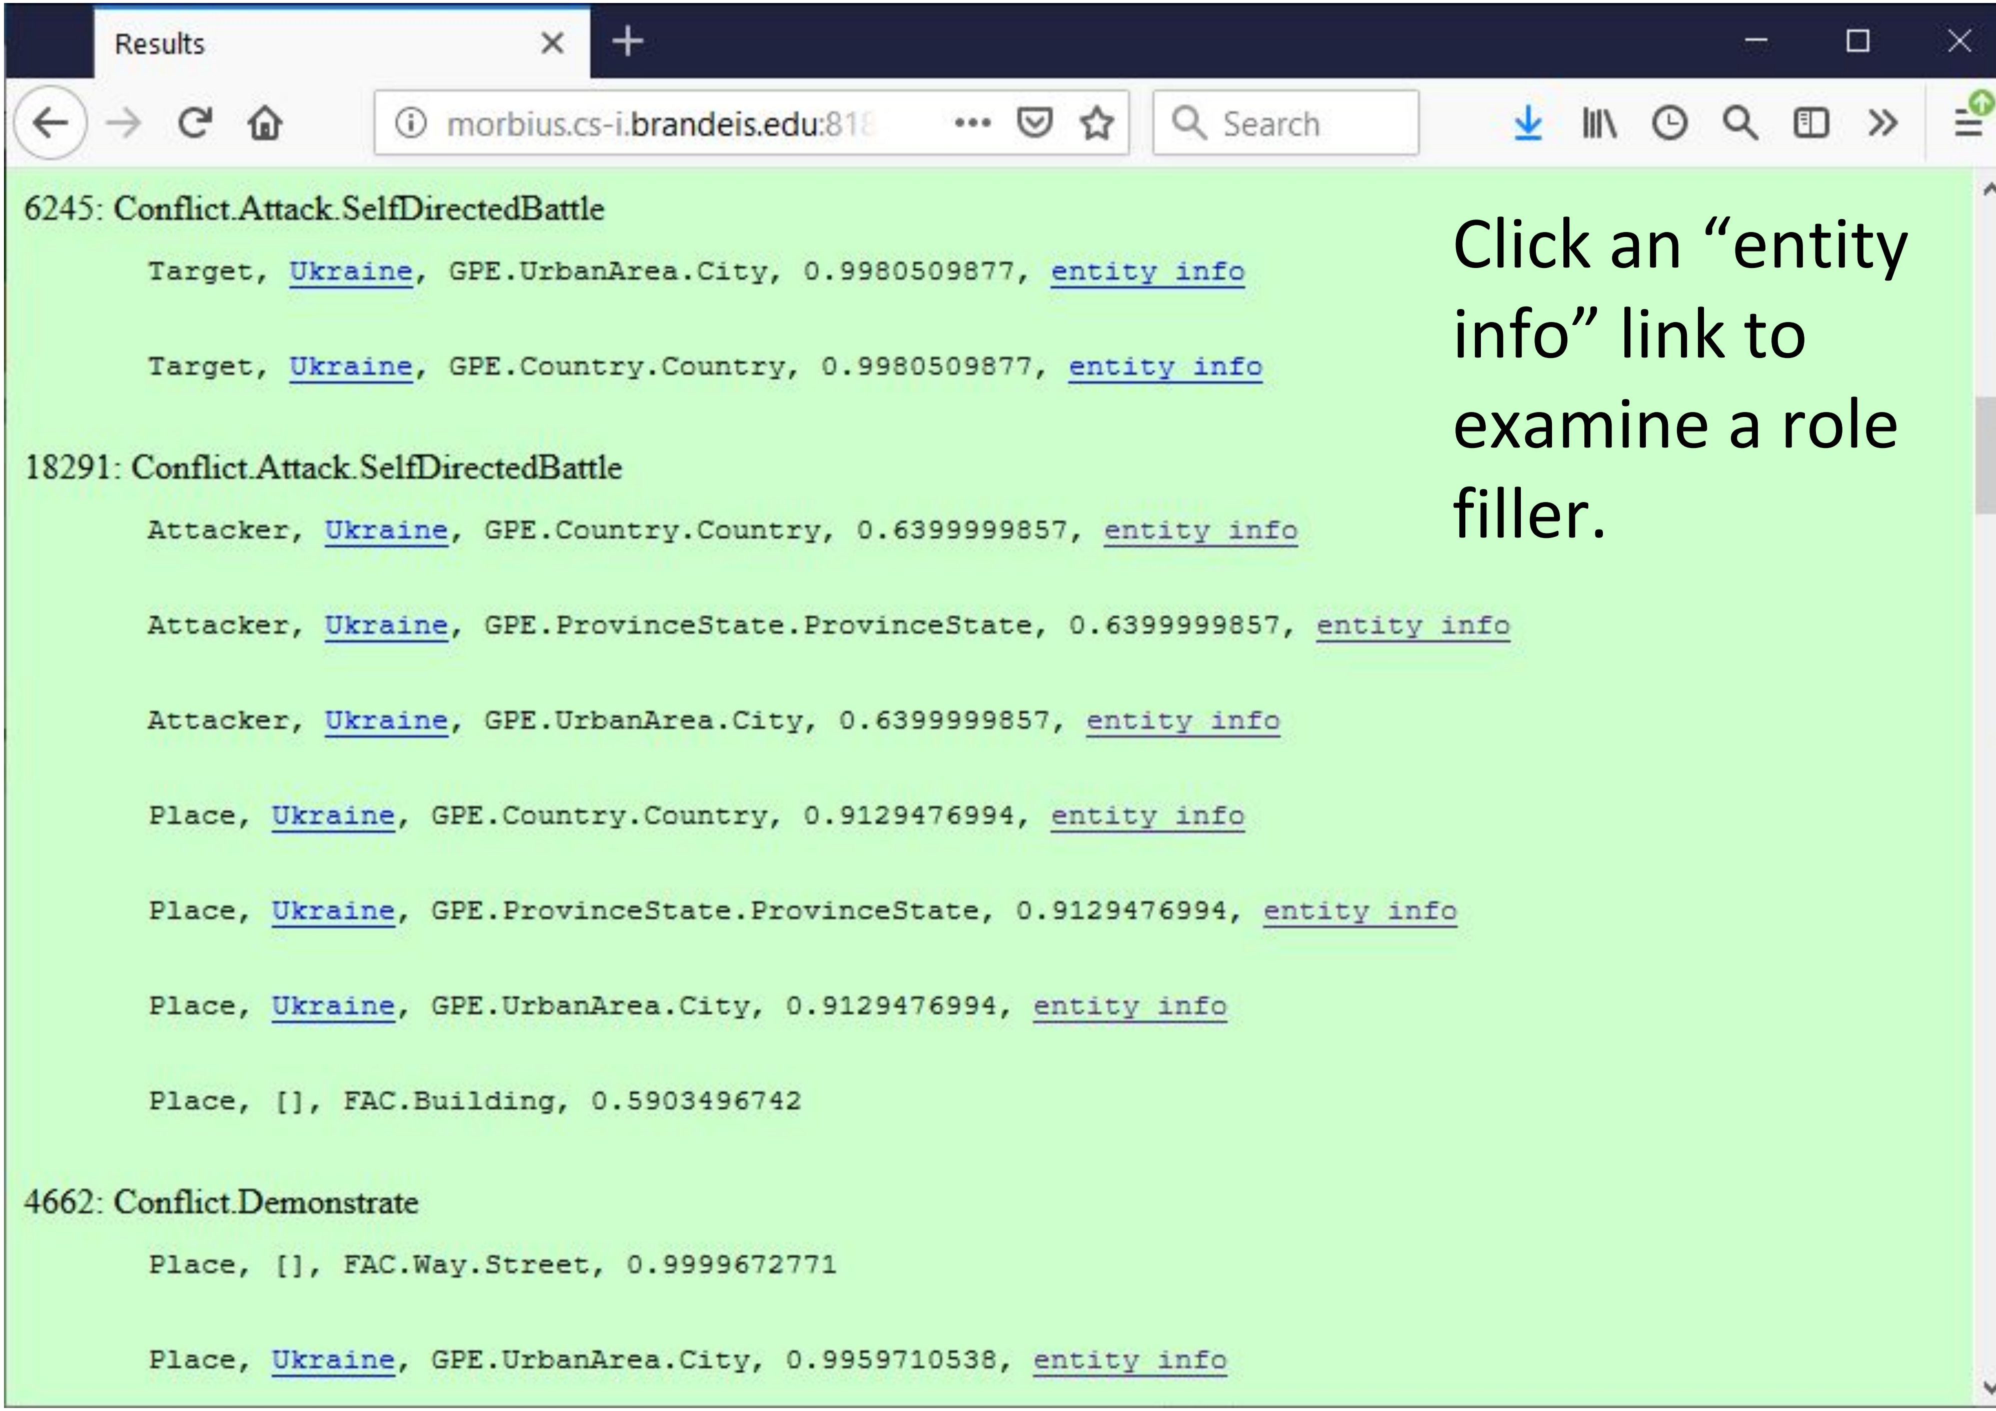

Click an “entity info” link to examine a role filler.

UI for browsing entity clusters

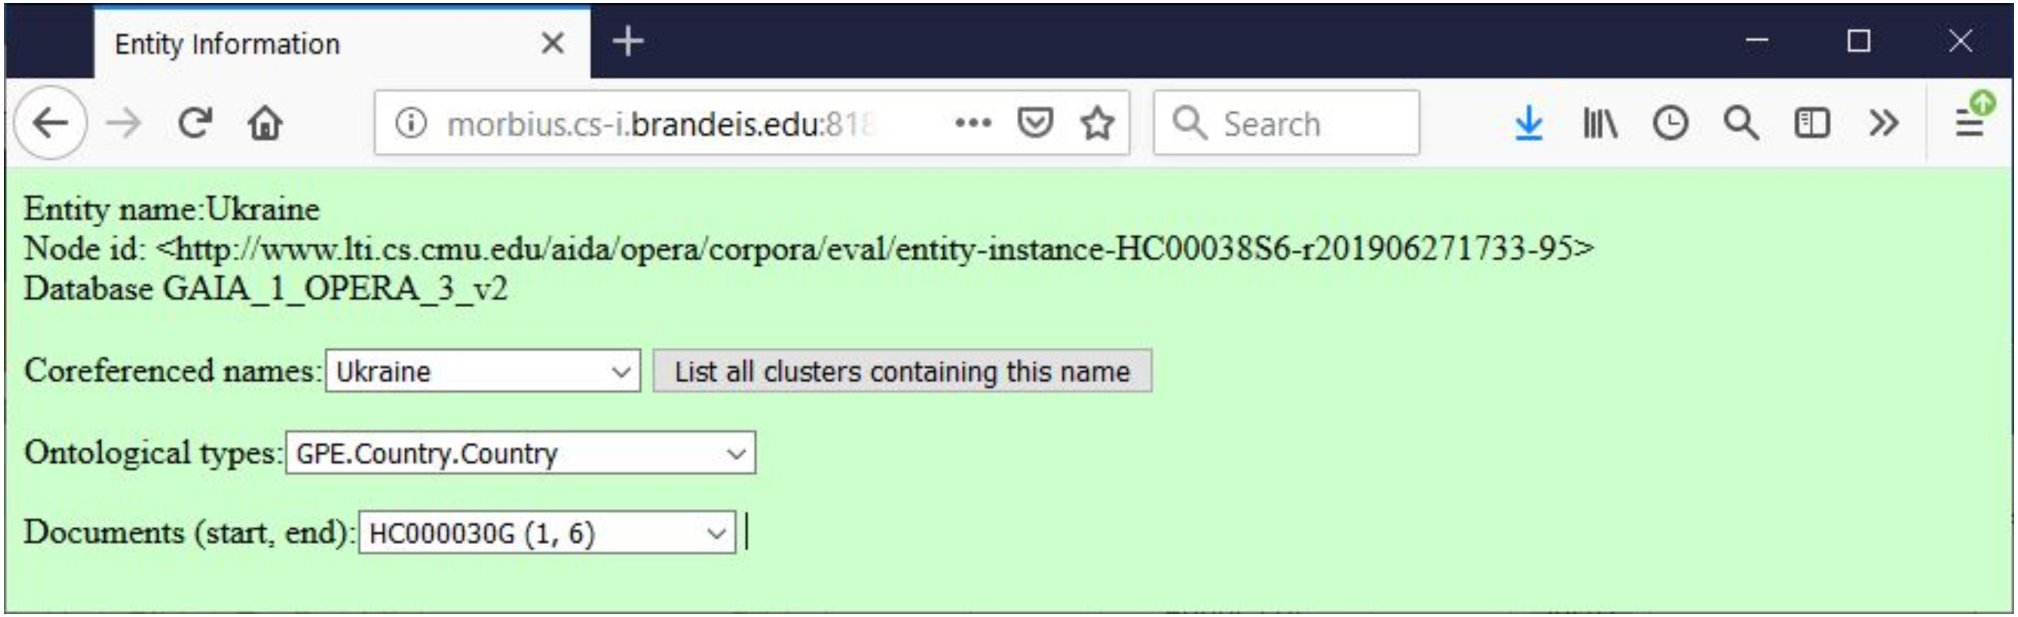

Alternative names used in the cluster

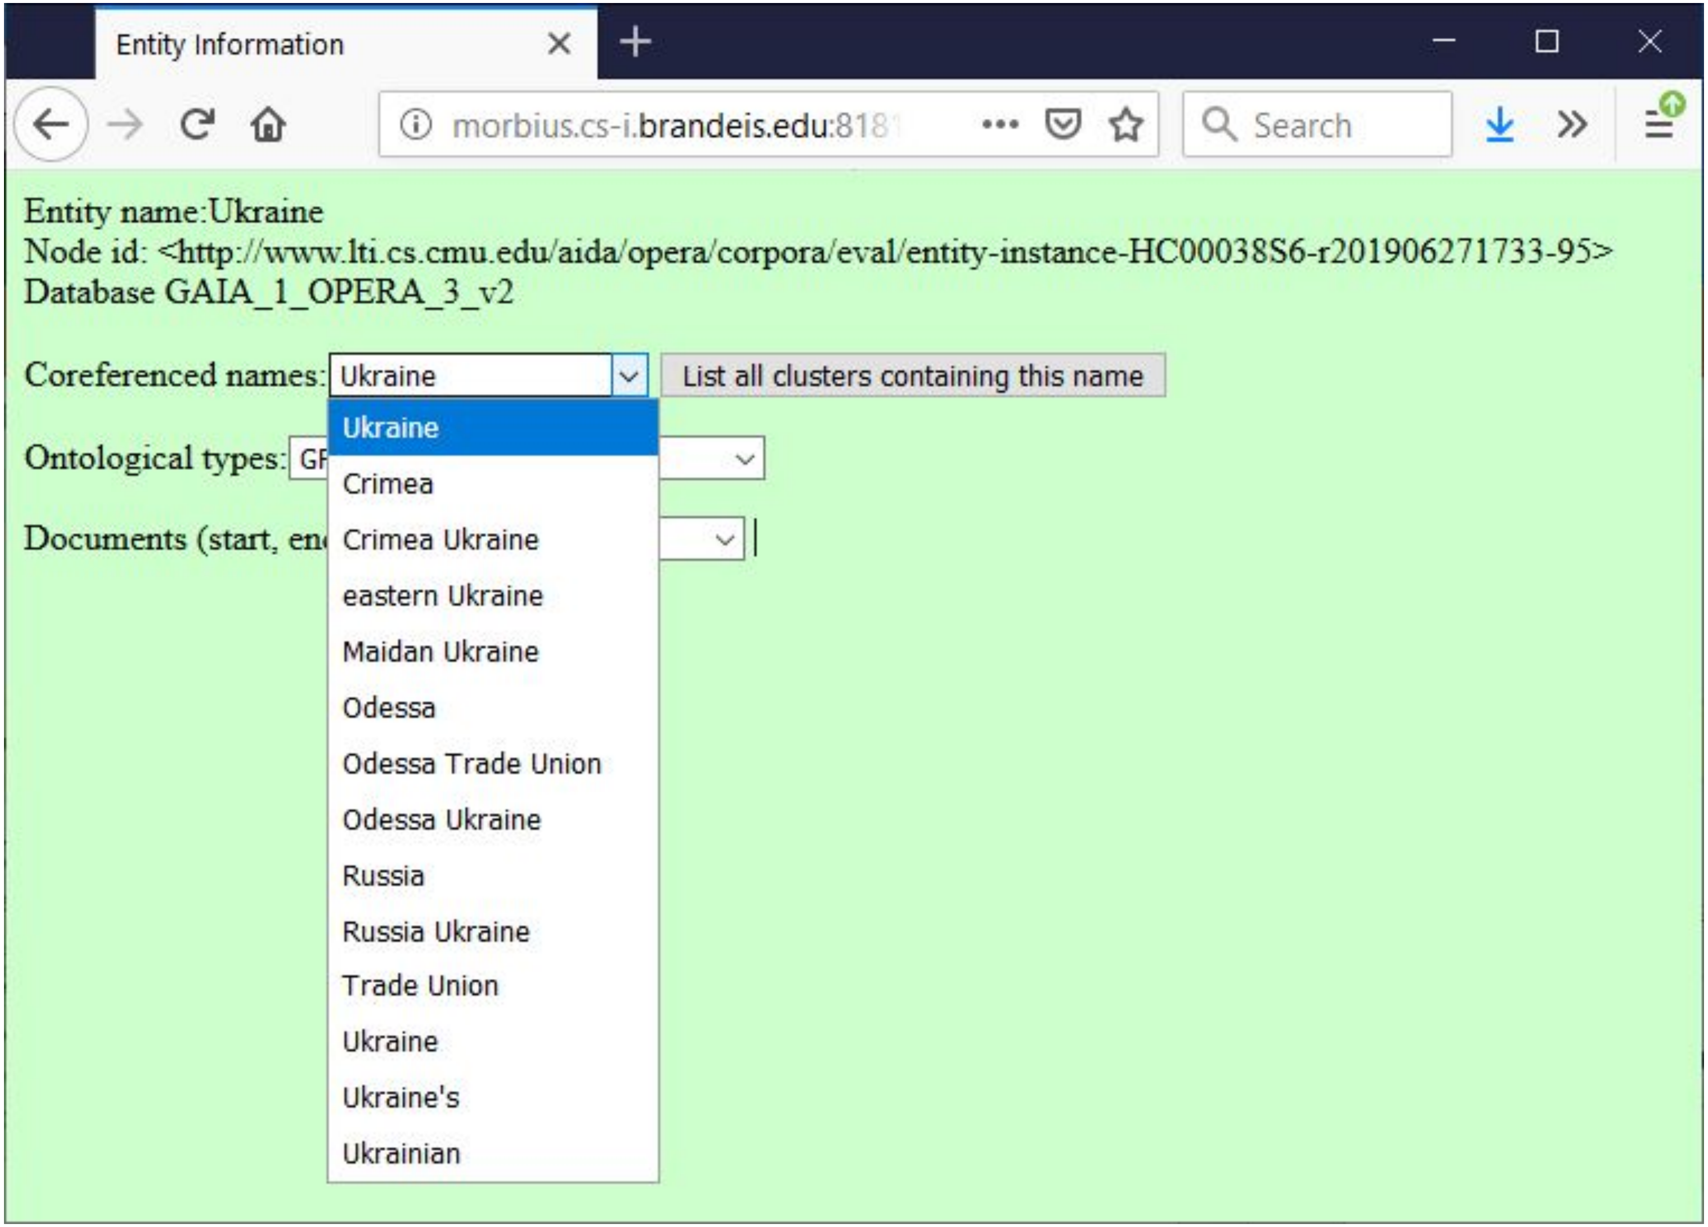

Alternative ontological types for mentions

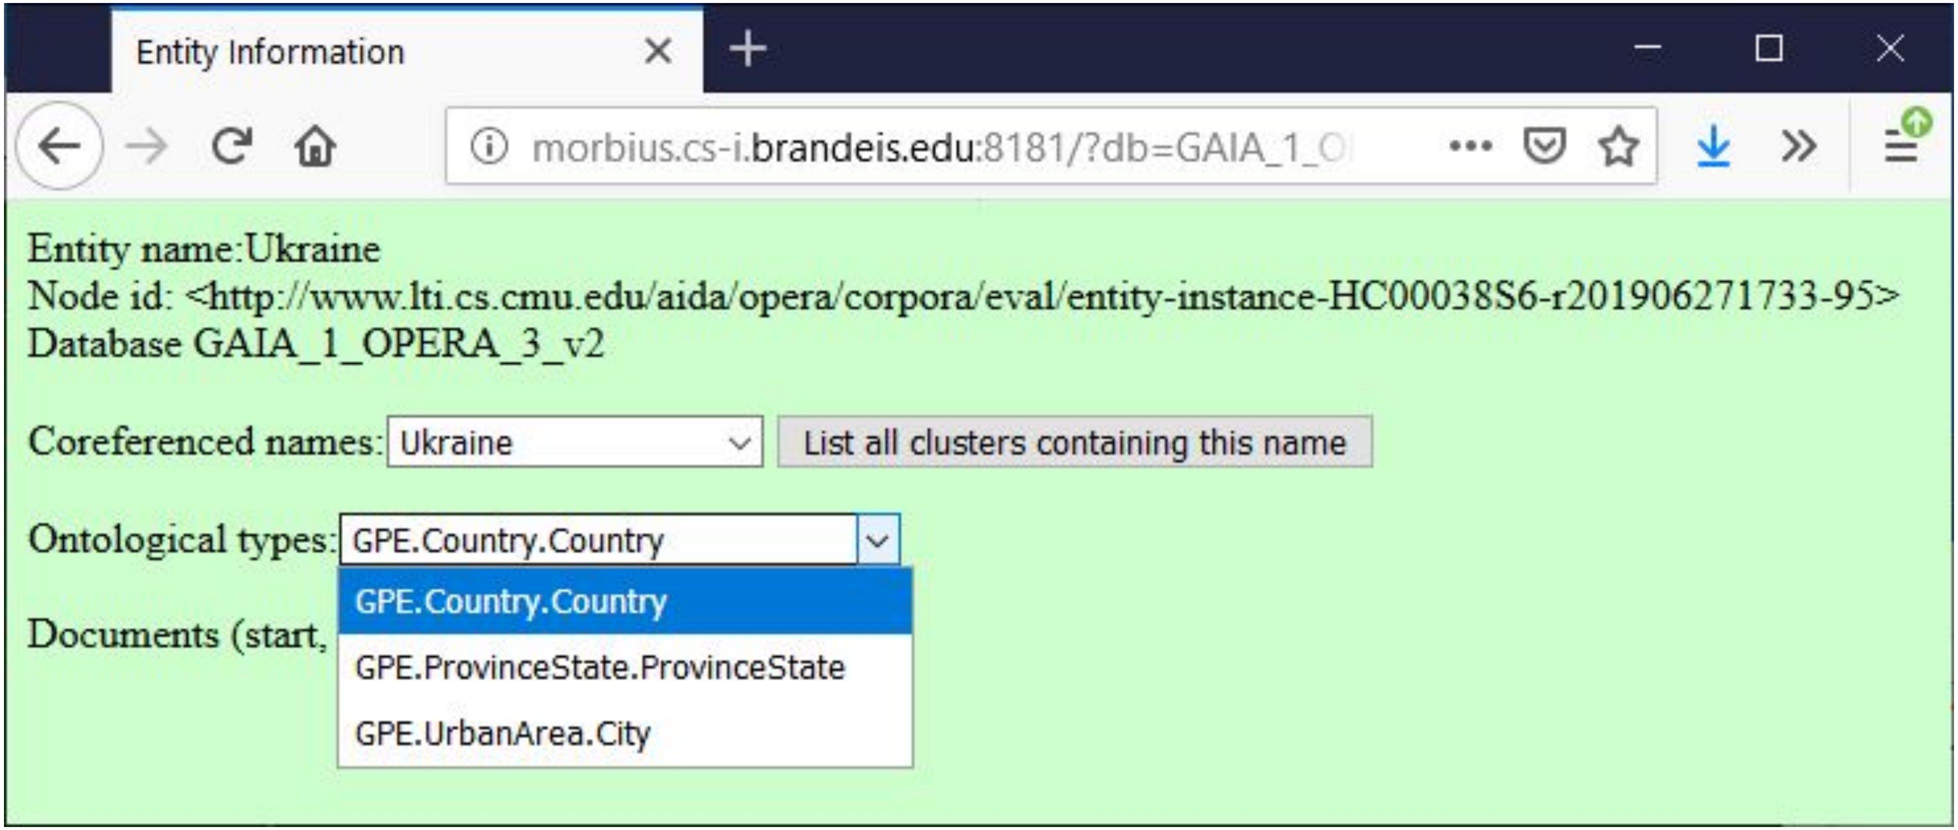

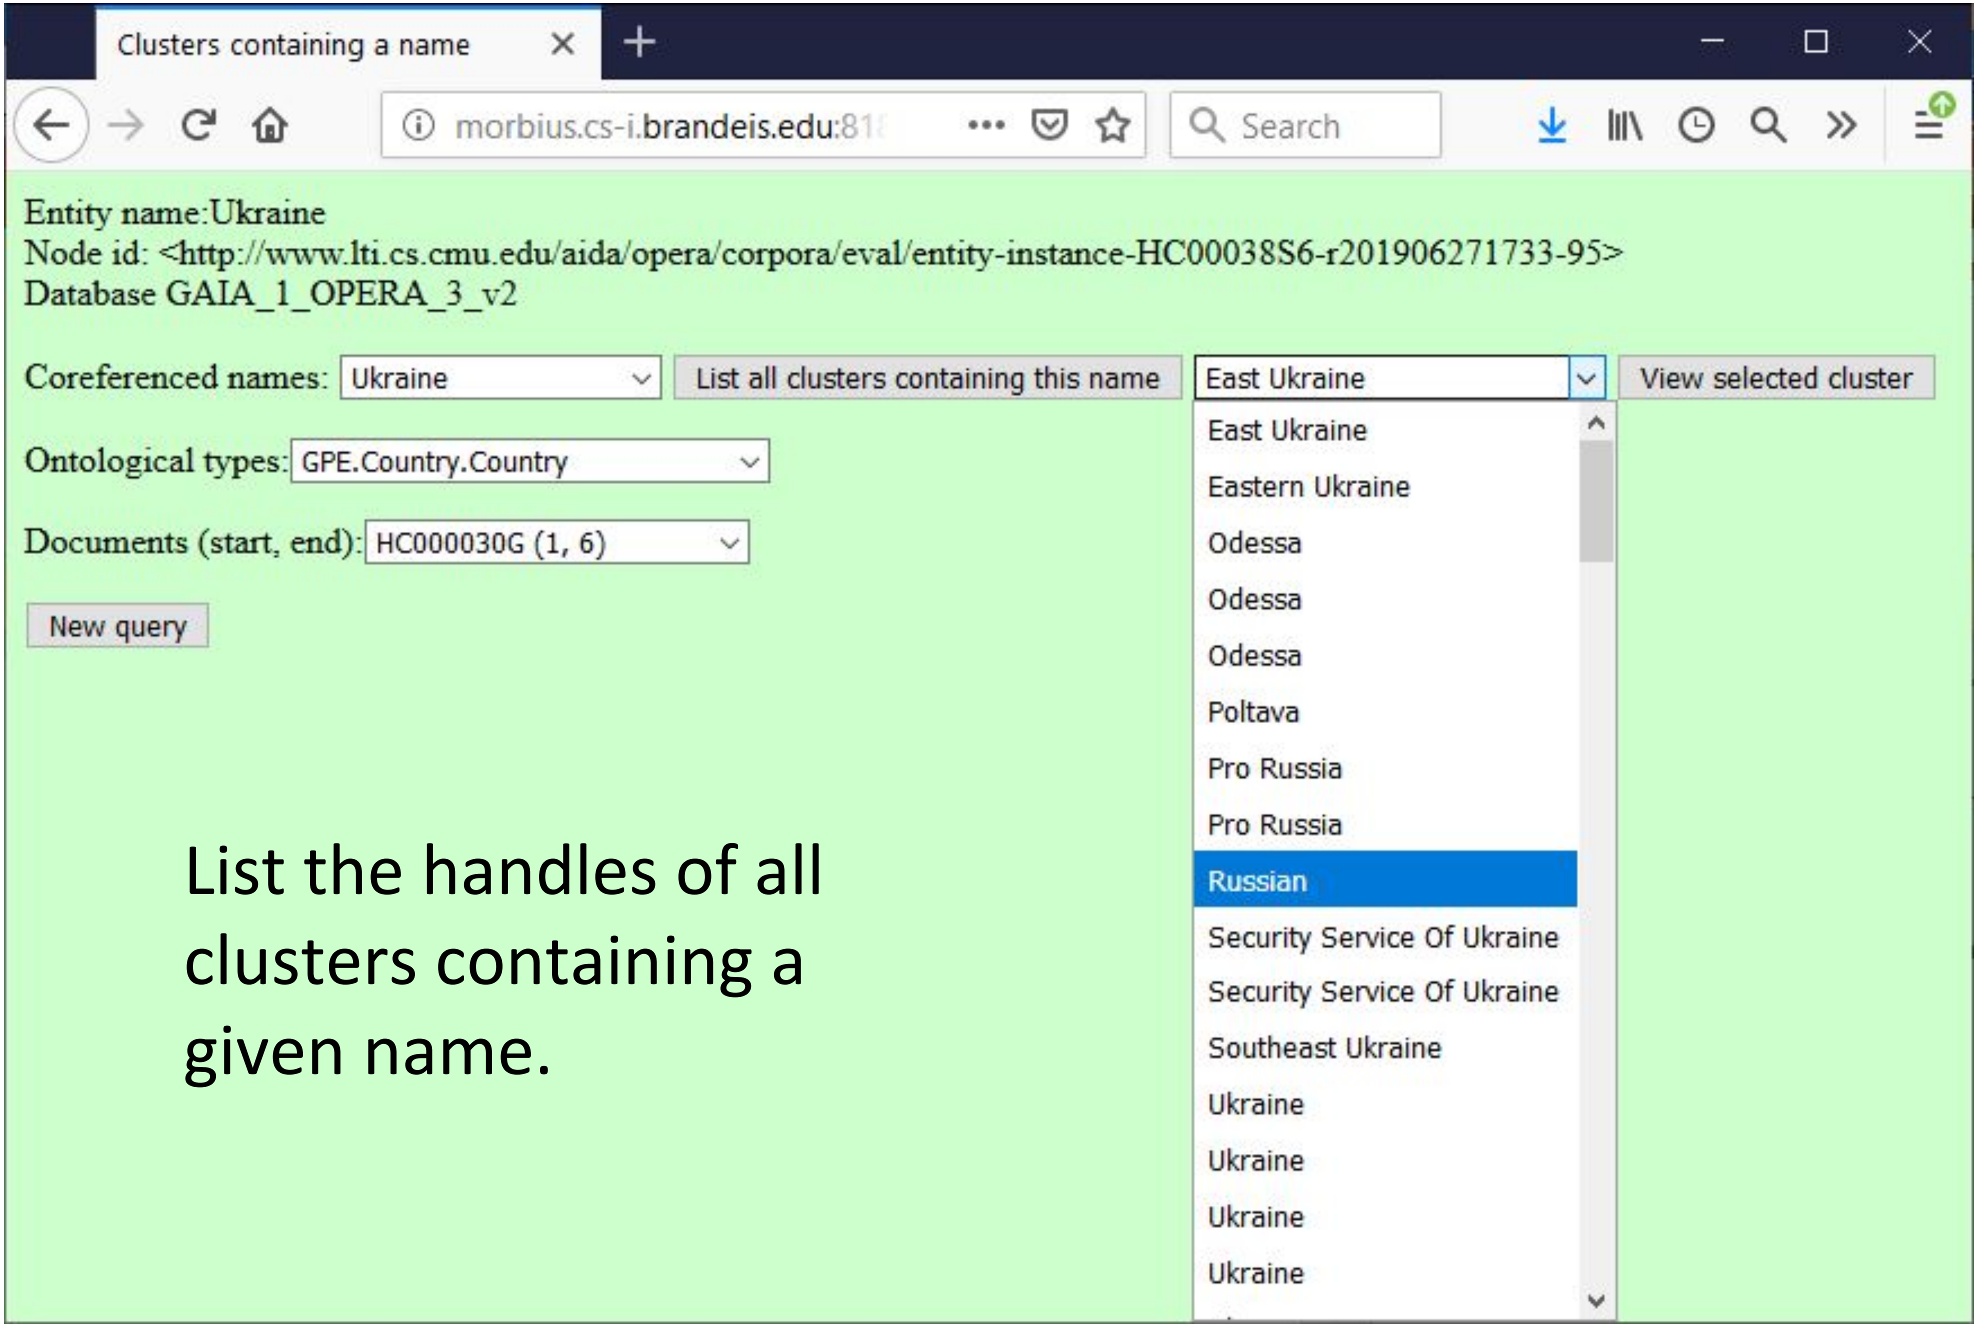

List the handles of all clusters containing a given name.

Select a cluster and view the set of named mentions within it.

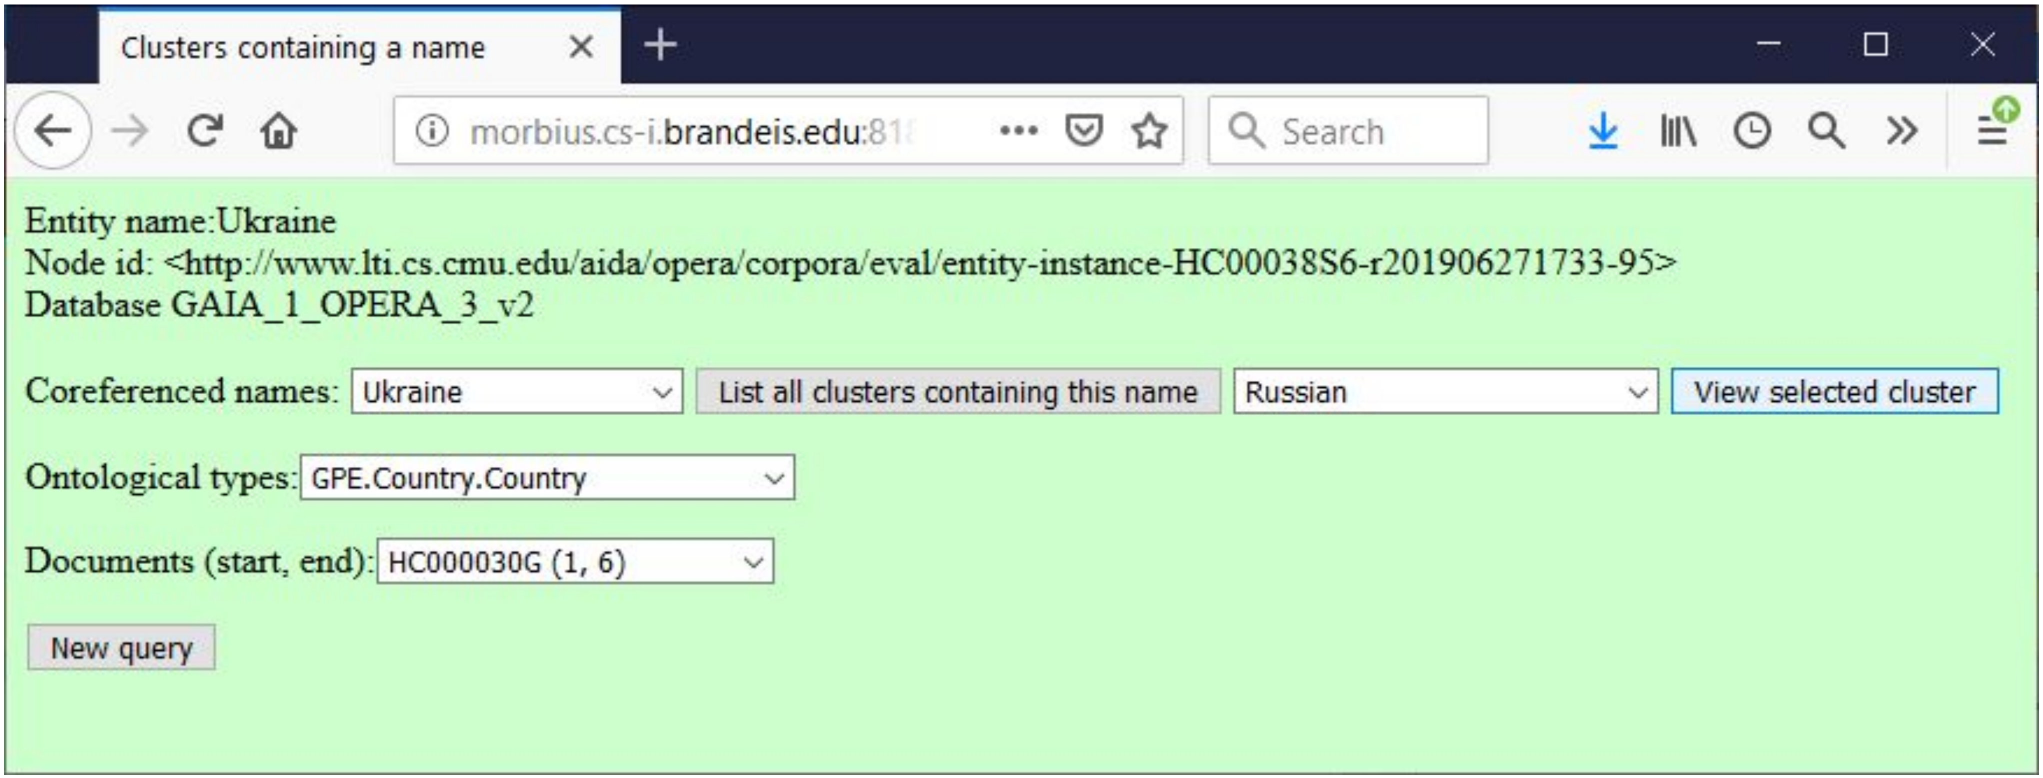

Coreferences across multiple documents

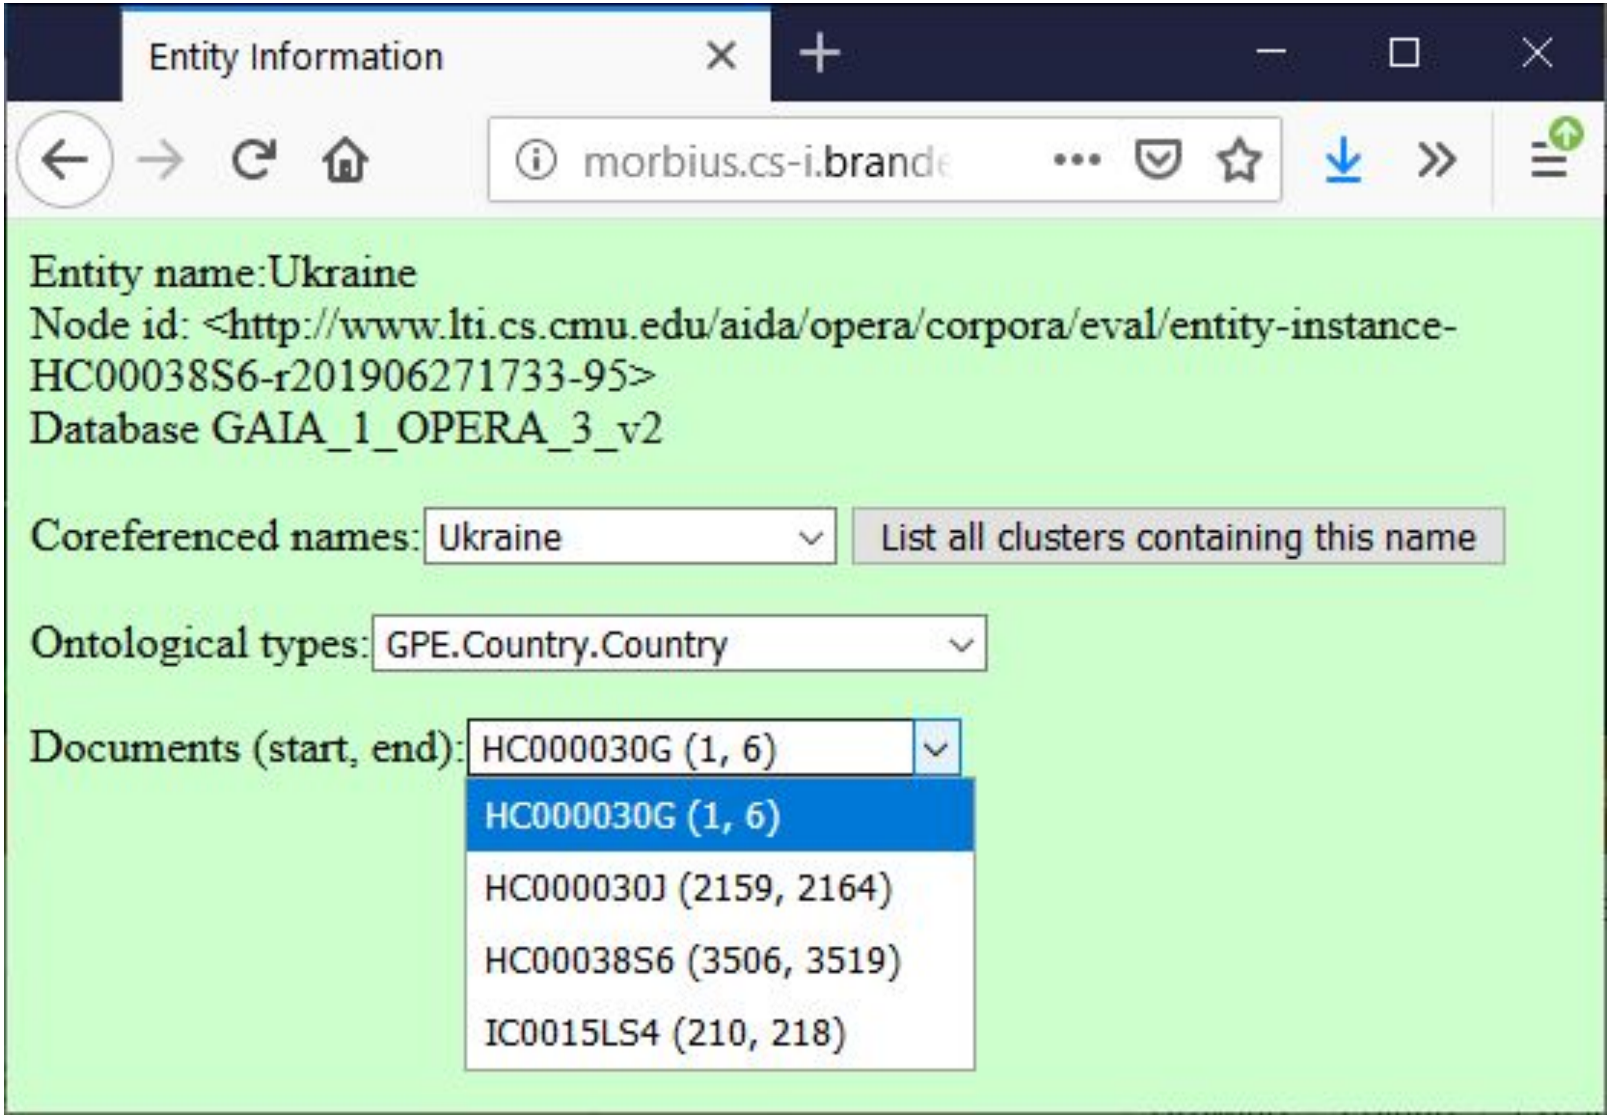

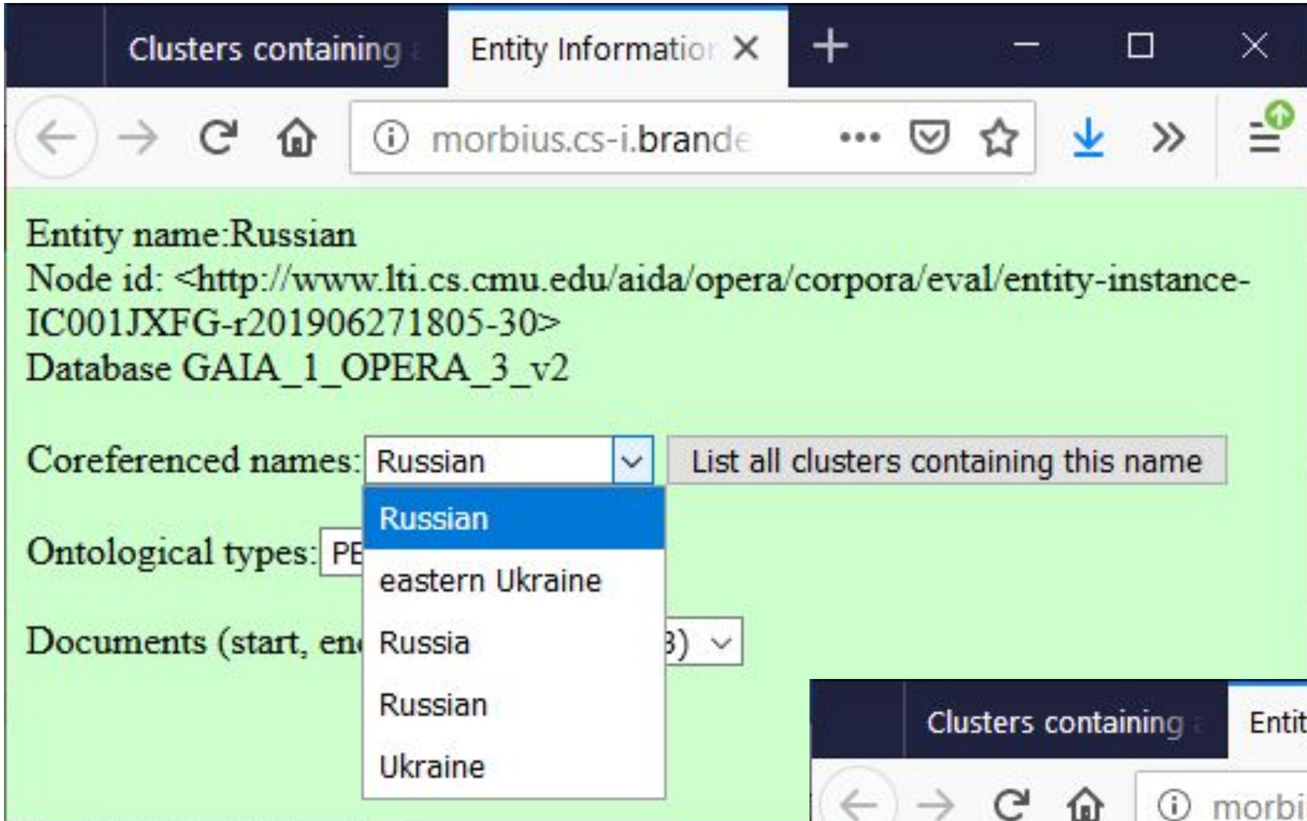

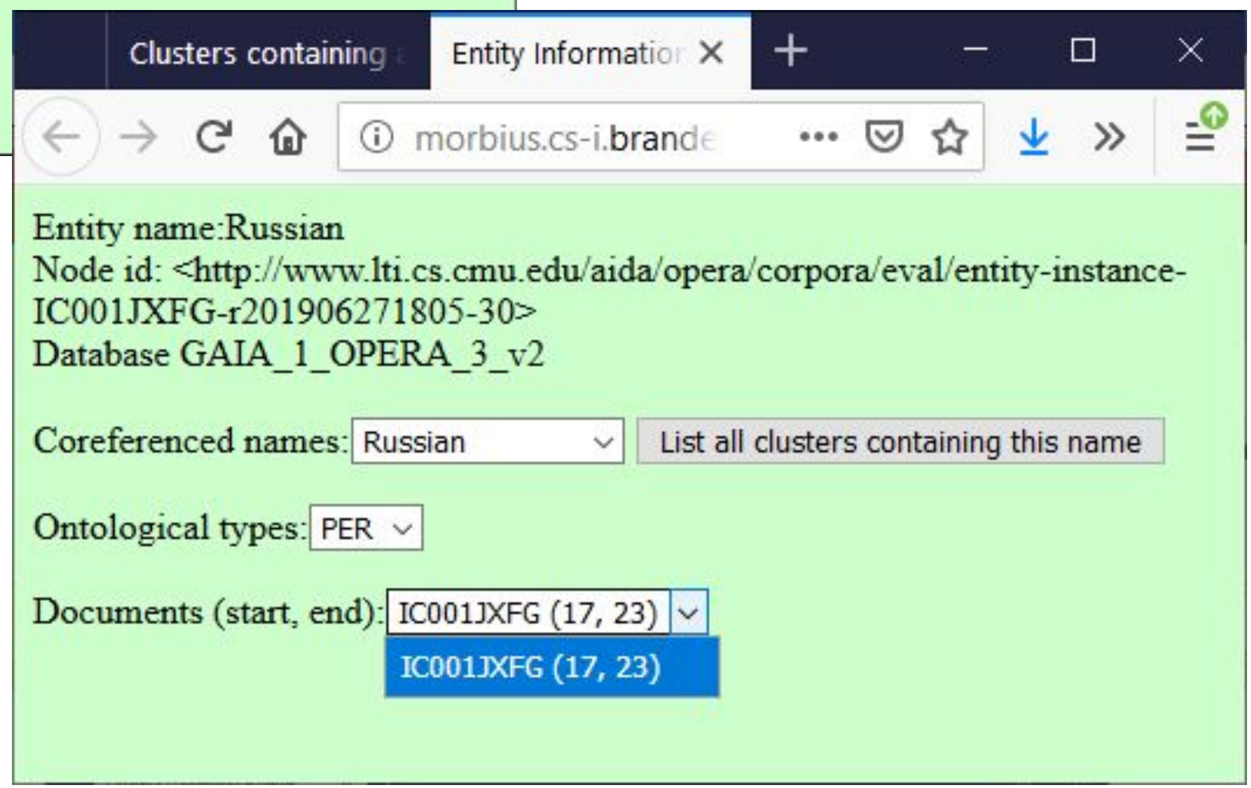

Russia-Ukraine coreferences within a single document may indicate a coreference error in TA1 output.

Two ways to explore entities

1. As fillers of events/relations
2. Directly by name

Then select a cluster by its handle

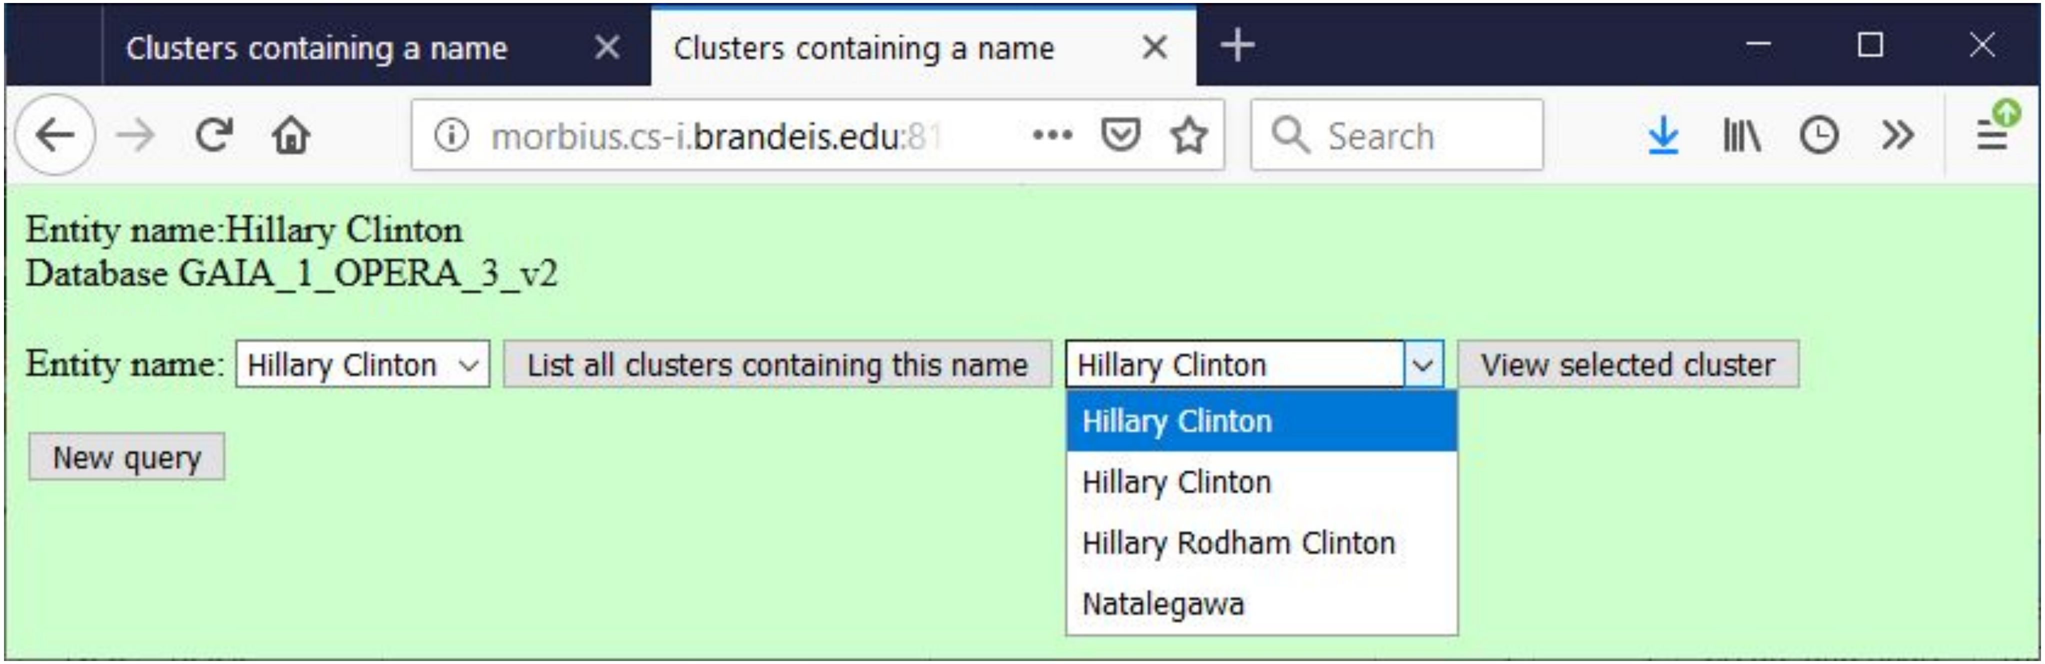

Mention names in cluster …d5a6

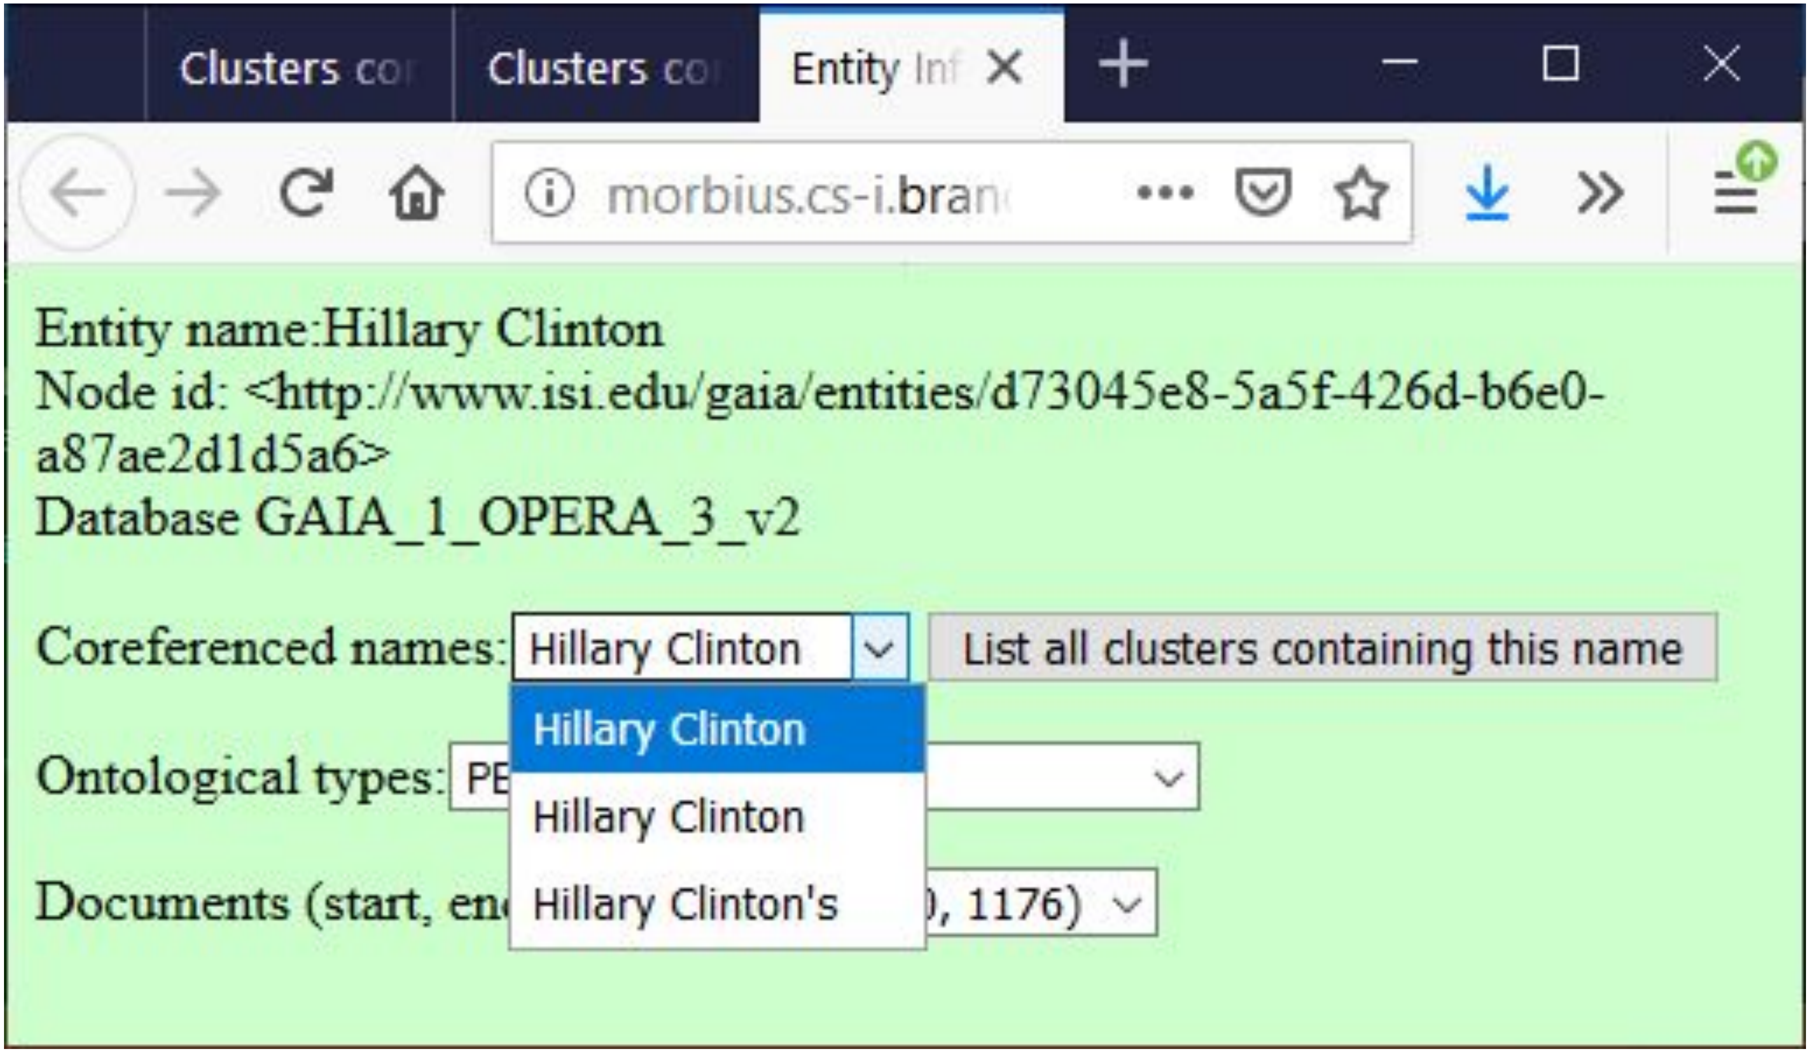

Cross-document coreferences appear in cluster …d5a6

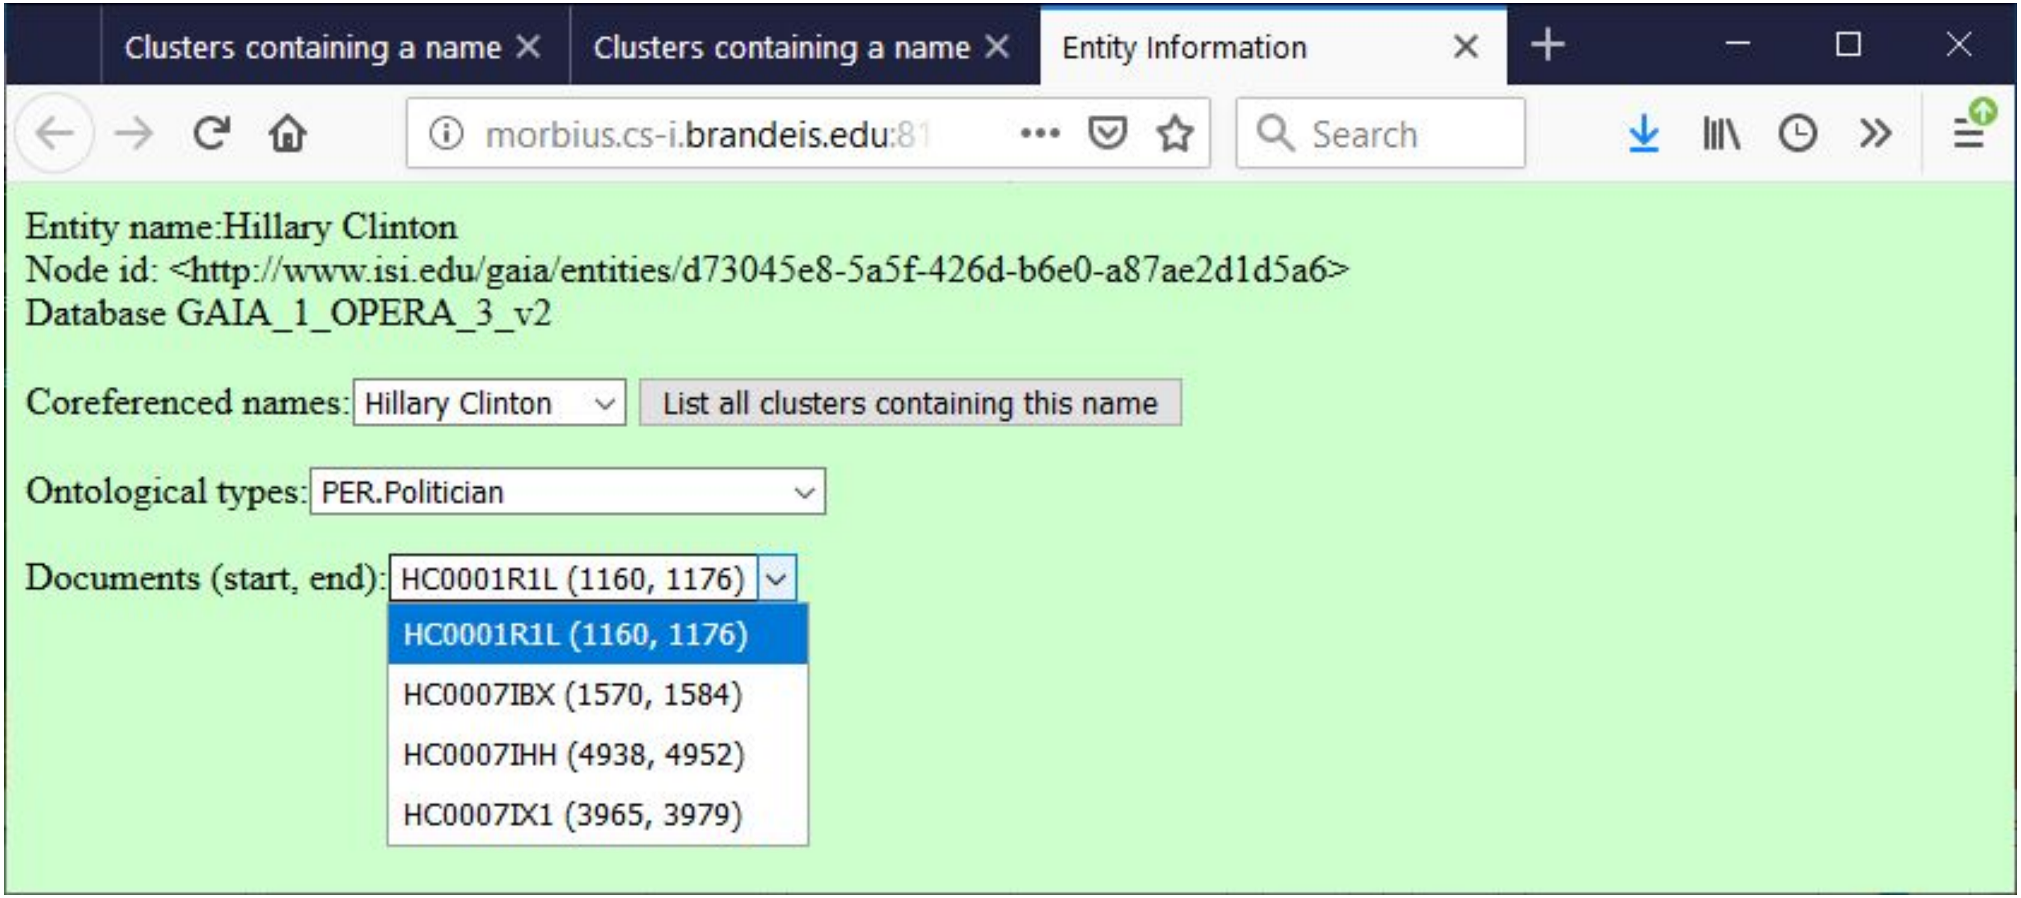

Conflicting mentions in cluster …22-75

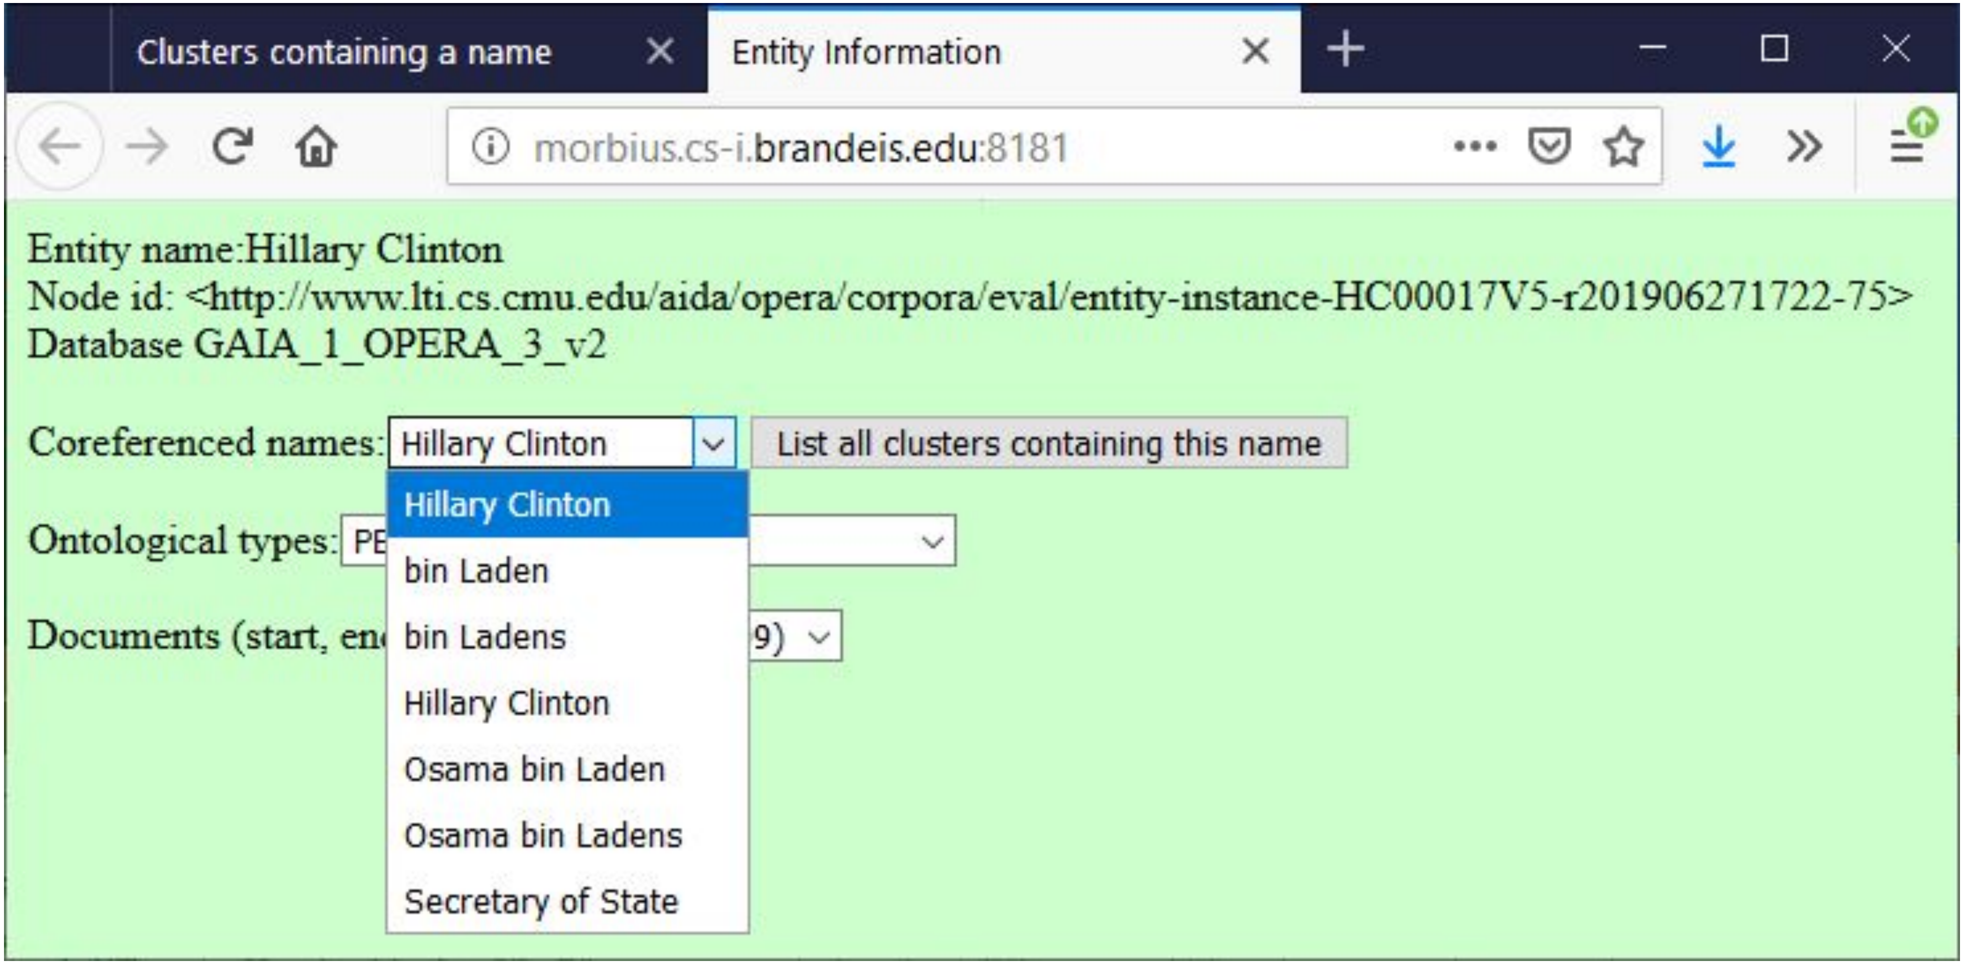

Conflicting mentions in cluster …09-52, traceable to parsing error

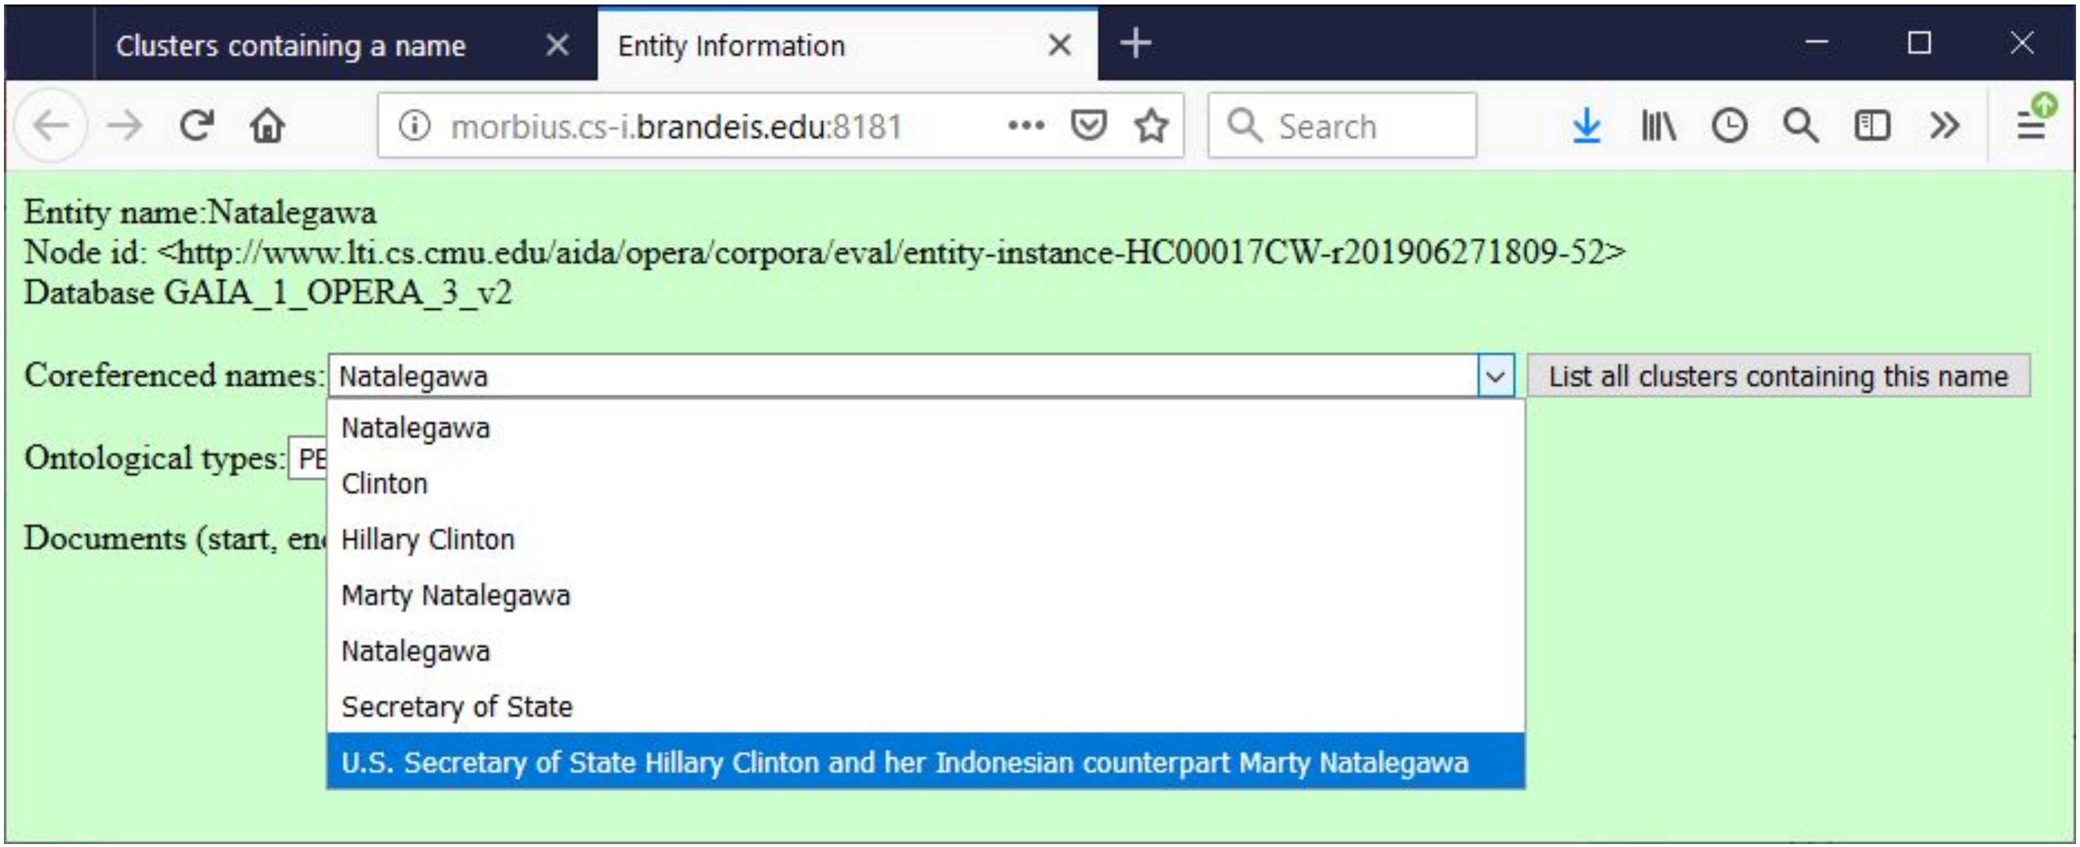

Conflicting mentions in cluster …09-52, traceable to parsing error *within a single document*

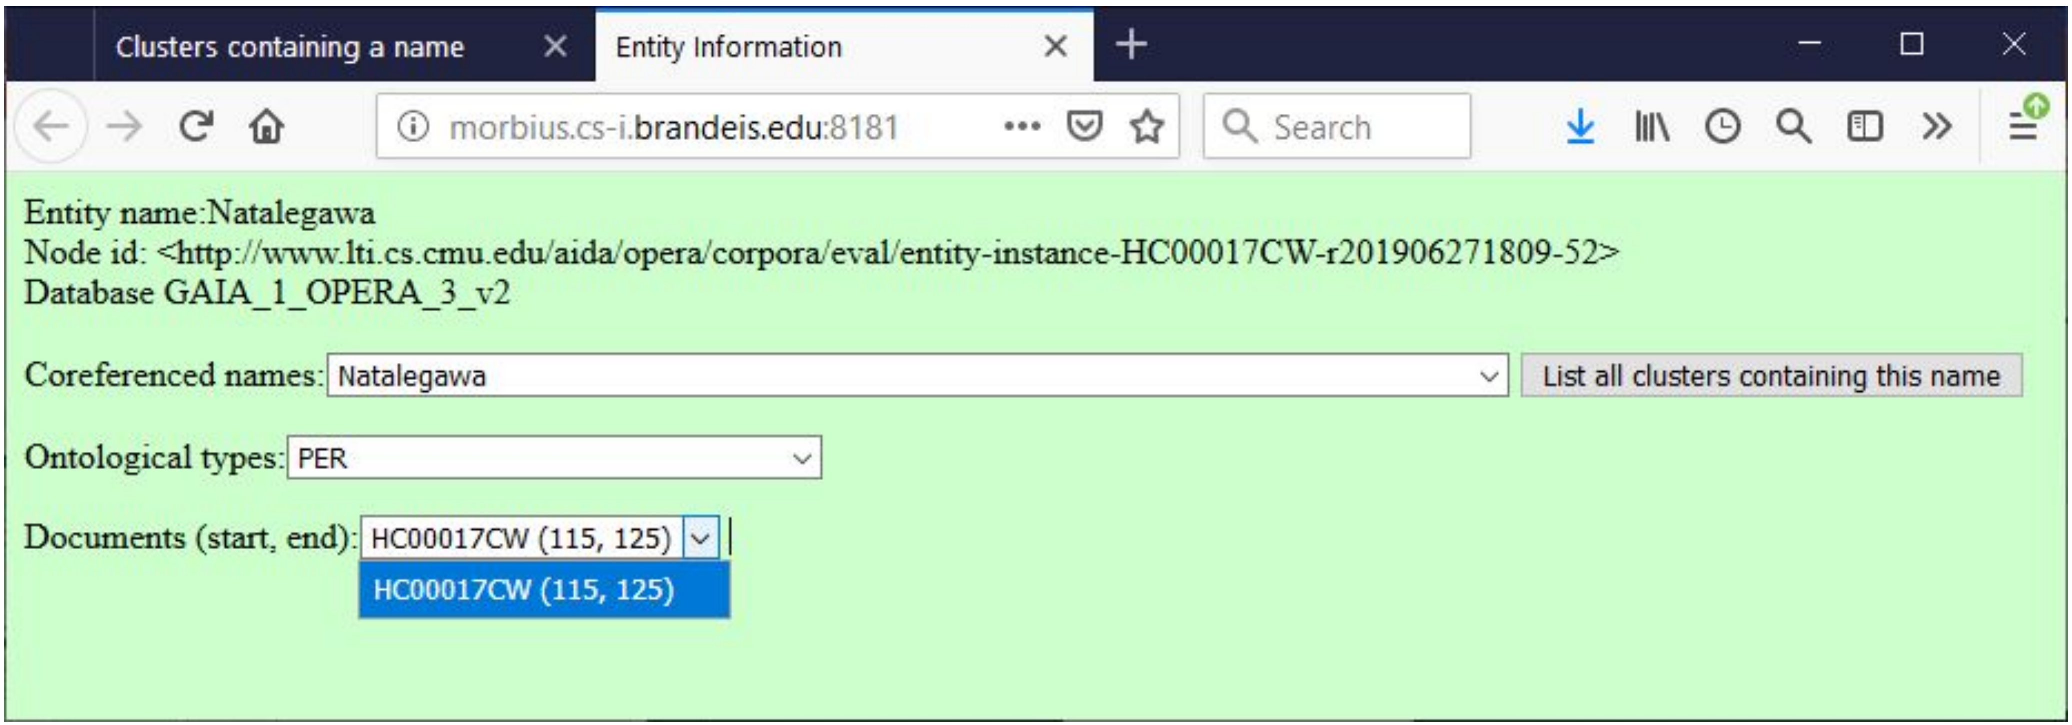

Mentions in cluster …16-29. Potential for cross-document coreference with cluster …d5a6

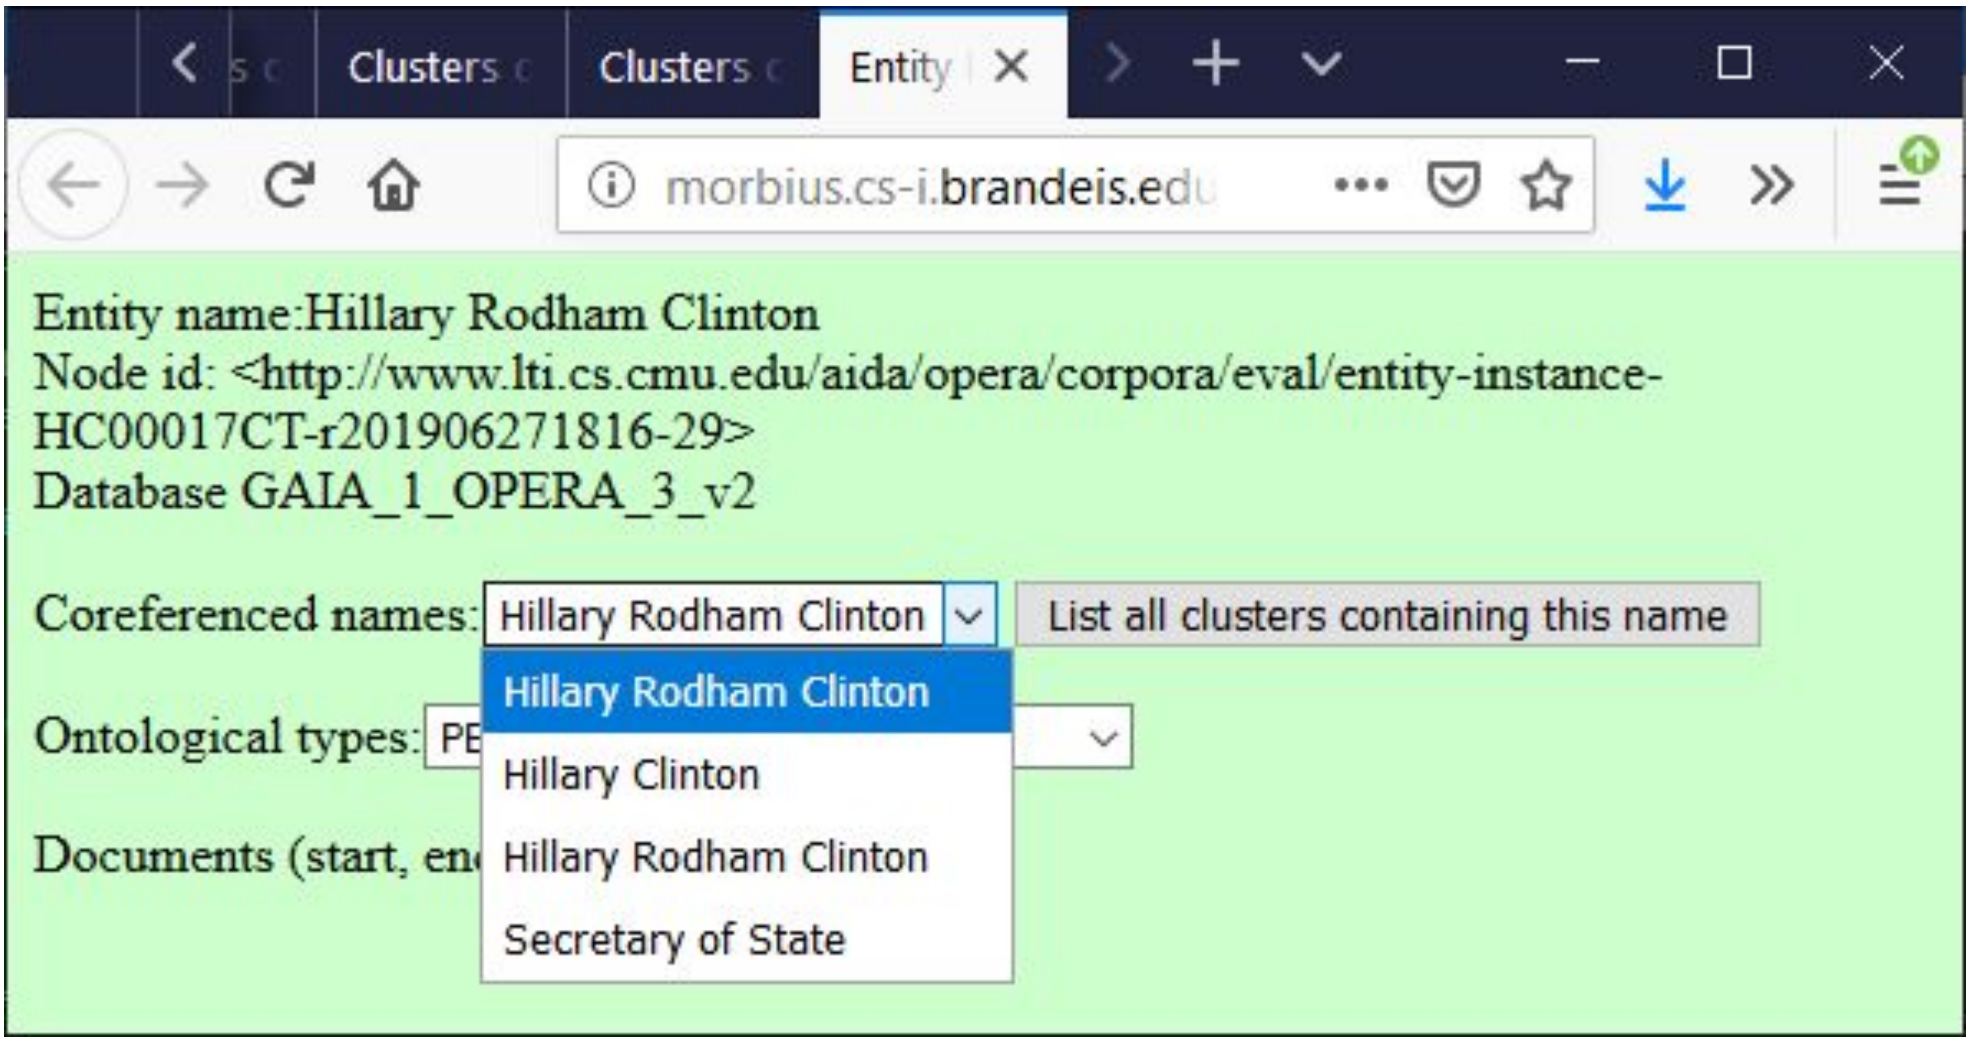

Status and next steps

- Usable by TA2 performers
  - For each knowledge graph, TA2 runs 5 sparql queries and sends results to Brandeis as zip archive.
  - Output of queries is processed into an SQL database with 8 tables.
  - Data accessible for browsing via web UI.
- Share prototype with Next Century (developing a more expansive model)
- Extend coverage to event coreference
- Support TA3 performer queries over TA2 events/relations (e.g., events with slots filled by specified ontological types)

CSU – Update April 15, 2020

RAMFIS: Representations of vectors and Abstract Meanings for Information Synthesis

Ross Beveridge
David White

Glimpse Ahead in our Talk

- We have now demonstrated how to discover mappings to support inferences between AIDA relevant feature (embedding) spaces.
- Both BBN and RPI have been most helpful and their data is key to our demonstrations.
- Our experiment shows practicality of discovering mapping and utility of the mapping to perform co-reference between two TA1s data.
- Note: the encoding mapping is revealed without direct access to either TA1's internal system.

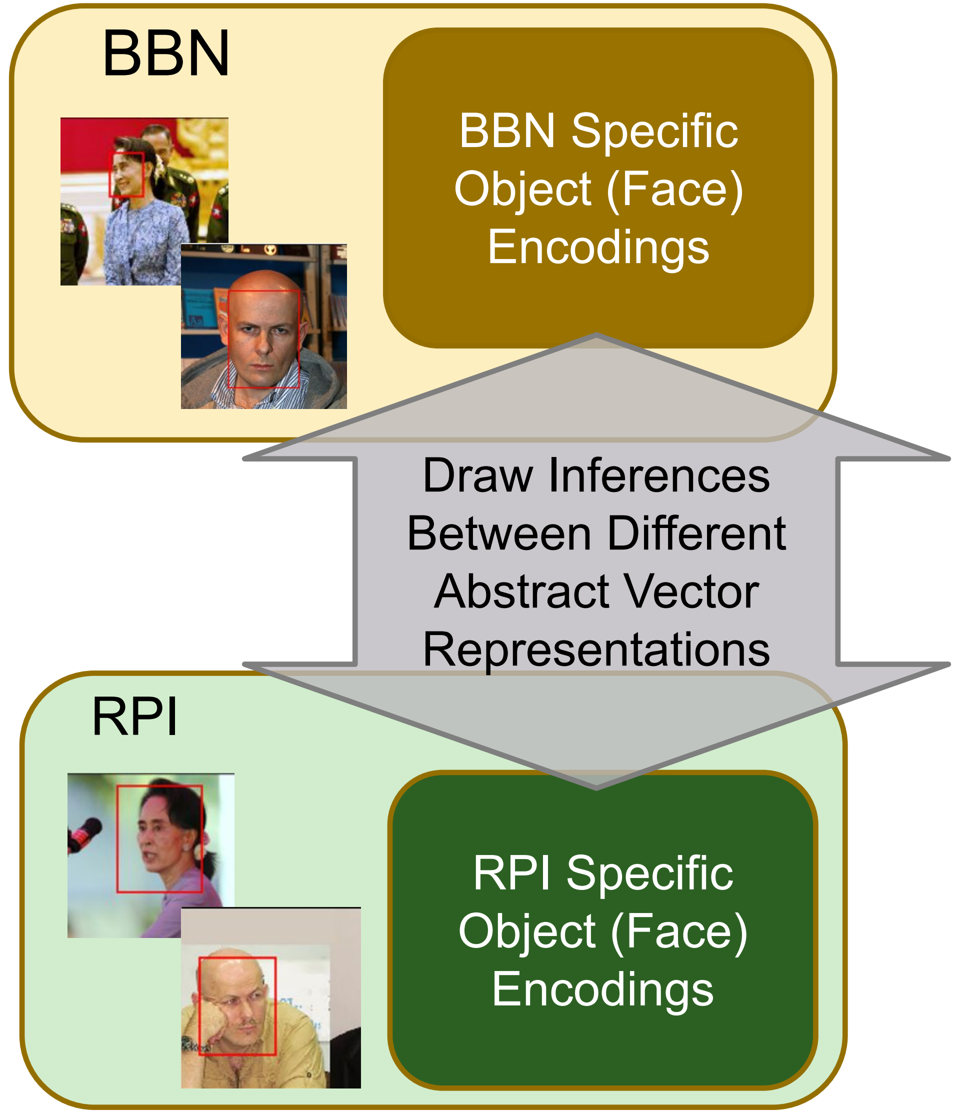

But First: Encoding in General

Essentially all modern object recognition systems encode the visual appearance of an instance of an object class in a highly abstracted feature vector.

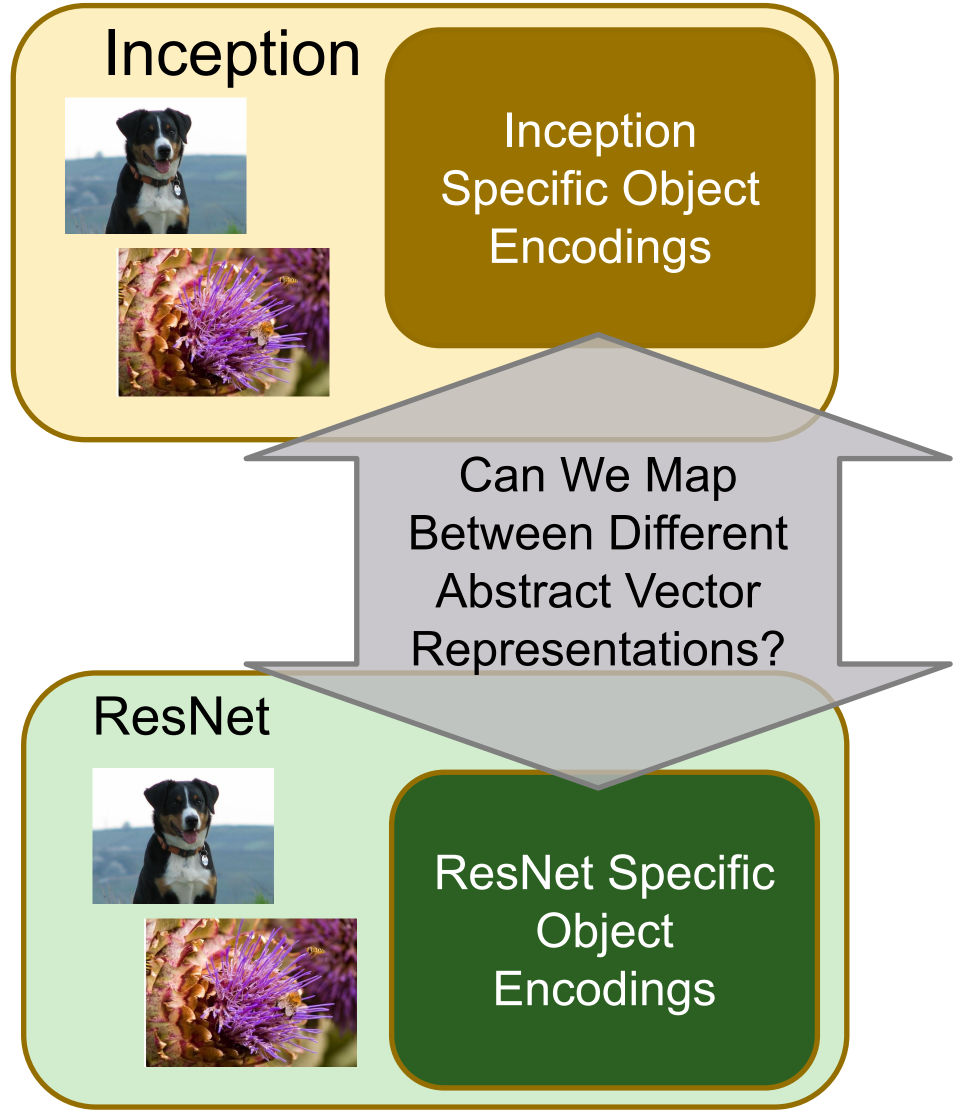

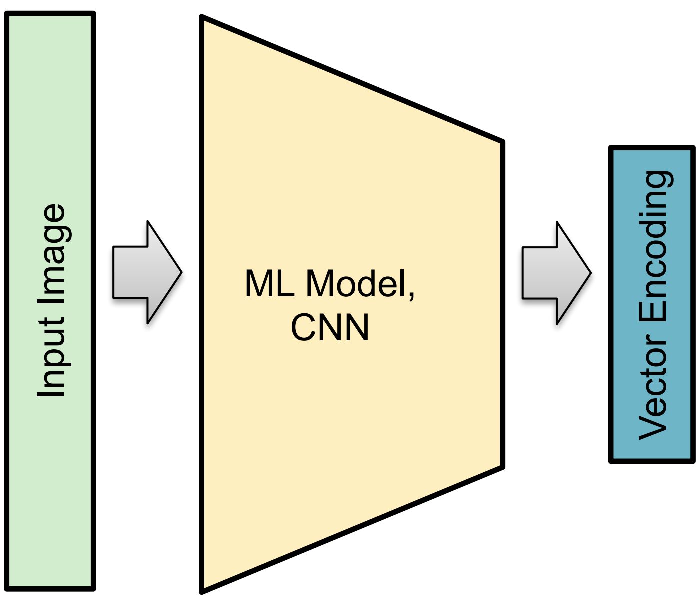

Trivial (Identity) Mapping

We have a run a more comprehensive and larger scale set of experiments with10 distinct and common ways of encoding ImageNet object classes. Here is the `null hypothesis`, i.e. feed features from one system to another without alteration.

10 CNN Backend Classifiers

10 CNN Feature Extractors

Source (feature extractor) \ Target (classifier)	Inception V1	Inception V2	MobileNet V2 1.4 224	ResNet V1 152	ResNet V2 152	Inception V3	Inception V4	Inception ResNet V2	NASNet Large	PNASNet Large
Inception V1	71.06%	0.10%	0.07%	0.08%	0.13%	0.12%	0.09%	0.08%	0.10%	0.07%
Inception V2	0.11%	73.94%	0.12%	0.14%	0.08%	0.14%	0.09%	0.10%	0.10%	0.10%
MobileNet V2 1.4 224	0.08%	0.09%	74.60%	0.08%	0.11%	0.11%	0.09%	0.12%	0.10%	0.11%
ResNet V1 152	0.09%	0.08%	0.13%	78.78%	0.22%	0.12%	0.11%	0.09%	0.11%	0.08%
ResNet V2 152	0.13%	0.15%	0.08%	0.18%	78.70%	0.07%	0.10%	0.10%	0.11%	0.13%
Inception V3	0.15%	0.12%	0.09%	0.27%	0.04%	78.88%	0.11%	0.11%	0.06%	0.08%
Inception V4	0.05%	0.10%	0.10%	0.10%	0.15%	0.07%	80.39%	0.12%	0.10%	0.10%
Inception ResNet V2	0.08%	0.11%	0.06%	0.10%	0.07%	0.12%	0.15%	81.16%	0.06%	0.12%
NASNet Large	0.03%	0.16%	0.09%	0.10%	0.09%	0.13%	0.09%	0.11%	82.66%	0.08%
PNASNet Large	0.09%	0.10%	0.09%	0.09%	0.10%	0.14%	0.04%	0.08%	0.06%	82.94%

Without alteration features from one are meaningless to another!

Affine Mappings

Latest result showing recognition rates using mapped encodings between all combinations of the 10 CNNs studied.
More red means greater drop. No drop is more than 12%

	Target (classifier)									
Source (feature extractor)	Inception V1	Inception V2	MobileNet V2 1.4 224	ResNet V1 152	ResNet V2 152	Inception V3	Inception V4	Inception ResNet V2	NASNet Large	PNASNet Large
Inception V1	71.06% 0.0%	62.82% -11.6%	68.17% -4.1%	68.45% -3.7%	68.31% -3.9%	67.36% -5.21%	66.77% -6.04%	65.86% -7.32%	64.97% -8.58%	65.40% -7.97%
Inception V2	69.91% -5.44%	73.94% 0.0%	72.65% -1.74%	72.73% -1.63%	72.66% -1.73%	72.34% -2.17%	72.04% -2.57%	71.55% -3.24%	71.01% -3.97%	71.24% -3.65%
MobileNet V2 1.4 224	68.17% -8.62%	66.27% -11.16%	74.60% 0.0%	71.57% -4.06%	71.25% -4.49%	70.56% -5.42%	69.93% -6.25%	69.43% -6.93%	69.16% -7.30%	69.72% -6.54%
ResNet V1 152	73.98% -6.09%	73.42% -6.80%	76.94% -2.33%	78.78% 0.0%	77.04% -2.21%	76.47% -2.93%	76.08% -3.42%	75.86% -3.70%	75.26% -4.46%	75.12% -4.65%
ResNet V2 152	74.79% -4.97%	74.42% -5.45%	77.19% -1.92%	77.79% -1.16%	78.70% 0.0%	76.75% -2.48%	76.35% -2.98%	76.10% -3.31%	75.38% -4.22%	75.44% -4.15%
Inception V3	75.57% -4.19%	75.55% -4.22%	77.91% -1.22%	77.9% -1.24%	77.62% -1.59%	78.88% 0.0%	77.63% -1.58%	77.50% -1.75%	77.12% -2.23%	77.17% -2.17%
Inception V4	78.29% -2.61%	78.49% -2.37%	79.84% -0.69%	79.75% -0.79%	79.57% -1.01%	79.79% -0.74%	80.39% 0.0%	79.68% -0.88%	79.64% -0.93%	79.61% -0.97%
Inception ResNet V2	79.63% -1.89%	79.63% -1.89%	80.71% -0.55%	80.80% -0.44%	80.52% -0.78%	80.69% -0.57%	80.82% -0.41%	81.16% 0.0%	80.77% -0.48%	80.69% -0.58%
NASNet Large	81.00% -2.01%	81.30% -1.65%	82.43% 0.28%	82.32% -0.41%	82.25% -0.50%	82.42% -0.29%	82.61% -0.06%	82.58% -0.10%	82.66% 0.0%	82.65% -0.01%
PNASNet Large	81.16% -2.15%	81.40% -1.86%	82.62% -0.39%	82.52% -0.52%	82.46% -0.59%	82.72% -0.27%	82.80% -0.18%	82.77% -0.20%	82.84% -0.13%	82.94% 0.0%

There always exists an approximate affine mapping between systems !

Faces: Are Vectors Anonymous?

- For biometric applications there is considerable interest in:
  - What it means to encode a face?
  - What happens if an encoding vector is stolen?
  - Specifically, can identity be reconstructed from a stolen vector?
  - Can a stolen vector help with impersonation on a different vendor's system?
- Keep in mind these types of questions are closely related to what we are doing with different TA1 provided encodings.
- Also, whether the actual identity, i.e. name, associated with a vector is revealed is a separate issue from whether vectors across vendors can be meaningfully compared!

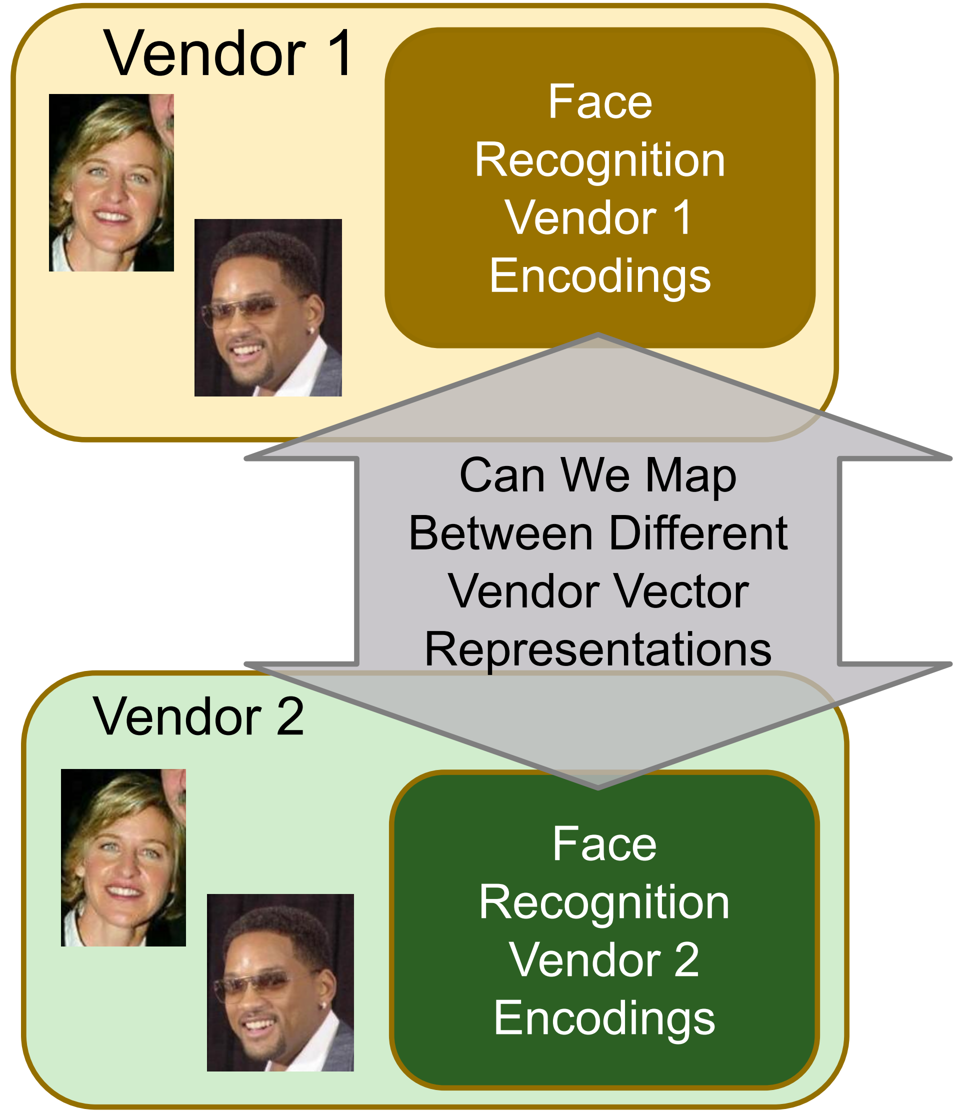

Computational Language and EducAtion Research
CLEAR

Co-Reference Experiment

- Picked 2 TA1 Performers: BBN and RPI.
- Established a Dataset for Experiment.
- Infer the mapping between encodings:
  - From Identity labeled samples.
  - Co-located (same image) samples.
- Part 1: ROC for known co-located samples.
- Part 2: Nearest-neighbor associations between BBN and RPI.

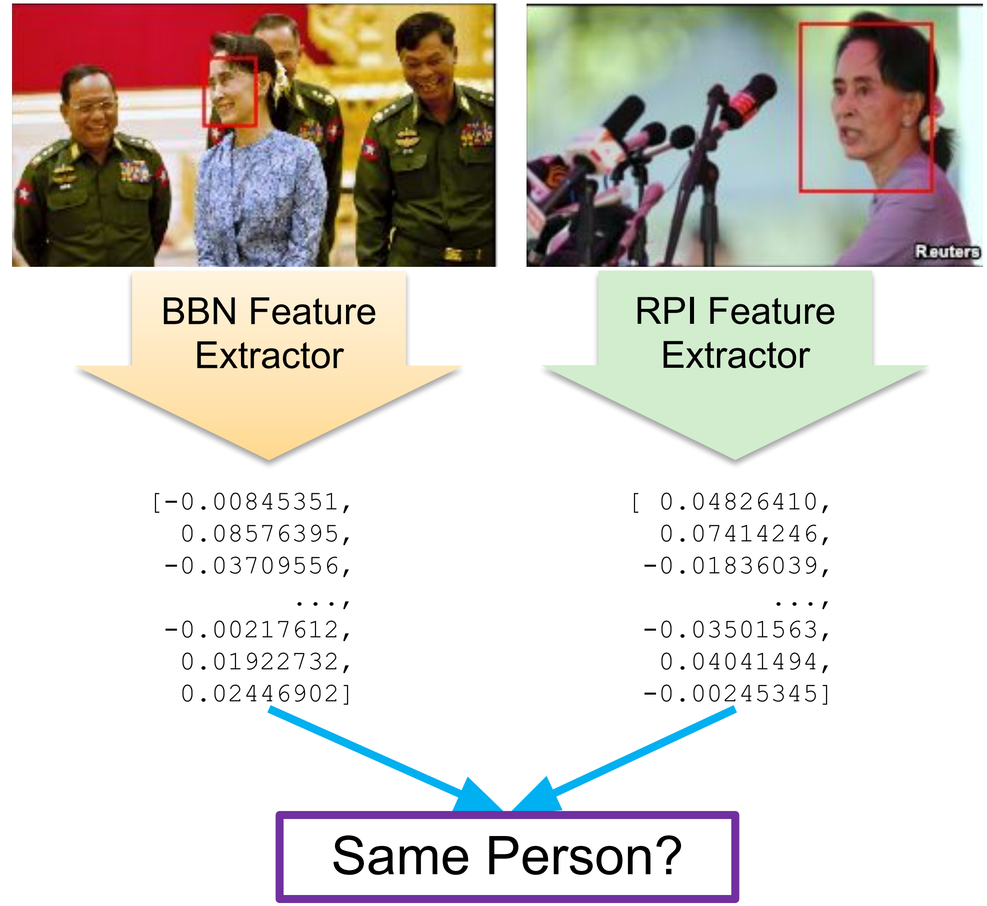

Part 2 is most interesting, showing that it is entirely possible to carry out joint analysis across different TA1 documents through the co-reference linkage built upon the discovered mapping between encodings.

Computational Language and EducAtion Research
CLEAR

Data Setup Details

From BBN:

- 295 identities, multiple embeddings each
- ~4k labeled embeddings total
- ~340k unlabeled embeddings from M18 corpus
- Each unlabeled embedding includes a bounding box and document ID

From RPI:

- 367 identities, 1 embedding each
- 367 labeled embeddings total
- ~50k unlabeled embeddings from M18 corpus
- Each unlabeled embedding includes a bounding box and document ID

What is Common Between BBN and RPI data:

- 67 labelled identities in common (BBN aggregated per-identity)
- ~7k spatially co-located unlabeled embeddings

Mapping from 67 People

- Our previous finding used 1.3 million training samples
  - Many fewer are likely needed
- Inferring a 512 x 512 affine mapping from only 67 example pairs is not possible
- We tried and, not surprisingly, the computation proved unstable
- Punchline – **we need an identity label-free path to discover the mapping**

Histograms of R^2 scores of models fit on 100 random partitions

BBN → RPI

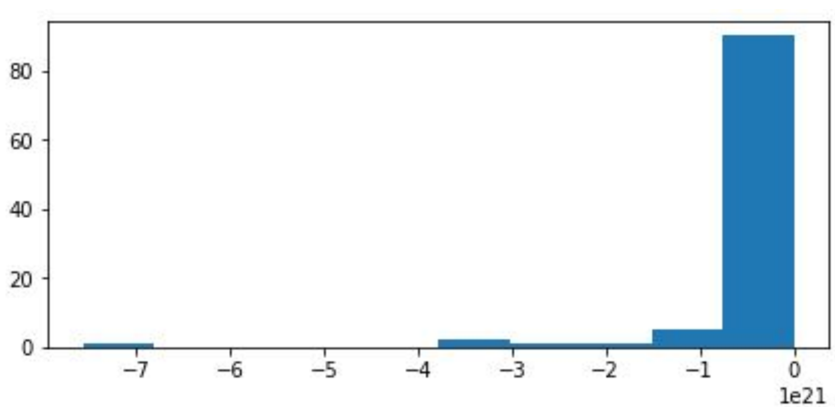

RPI → BBN

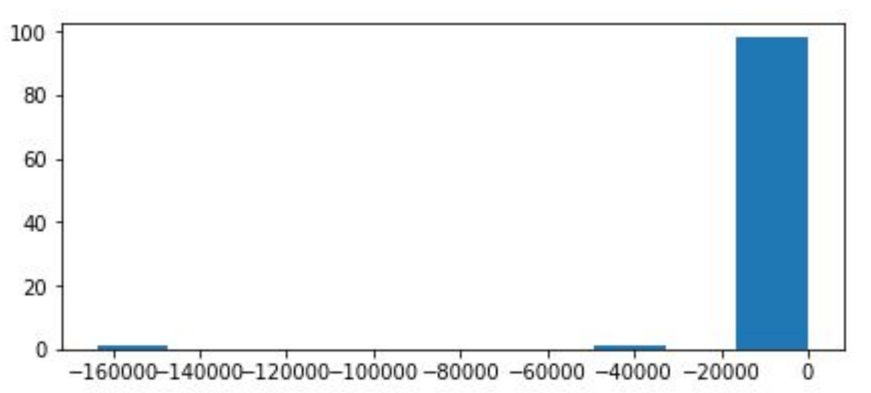

Mapping from Co-location

- Both BBN and RPI provide a bounding box and document identifier
- First, we cross-referenced all embeddings via document ID
- Then, we iteratively matched bounding boxes with the greatest Intersection Over Union (IOU) until no overlapping bounding boxes remained.
  - Essentially, a simple greedy co-location matching algorithm
  - Note: our process currently excludes embeddings from video frames

BBN Face Detection

RPI Face Detection

HC0002RJH.jpg

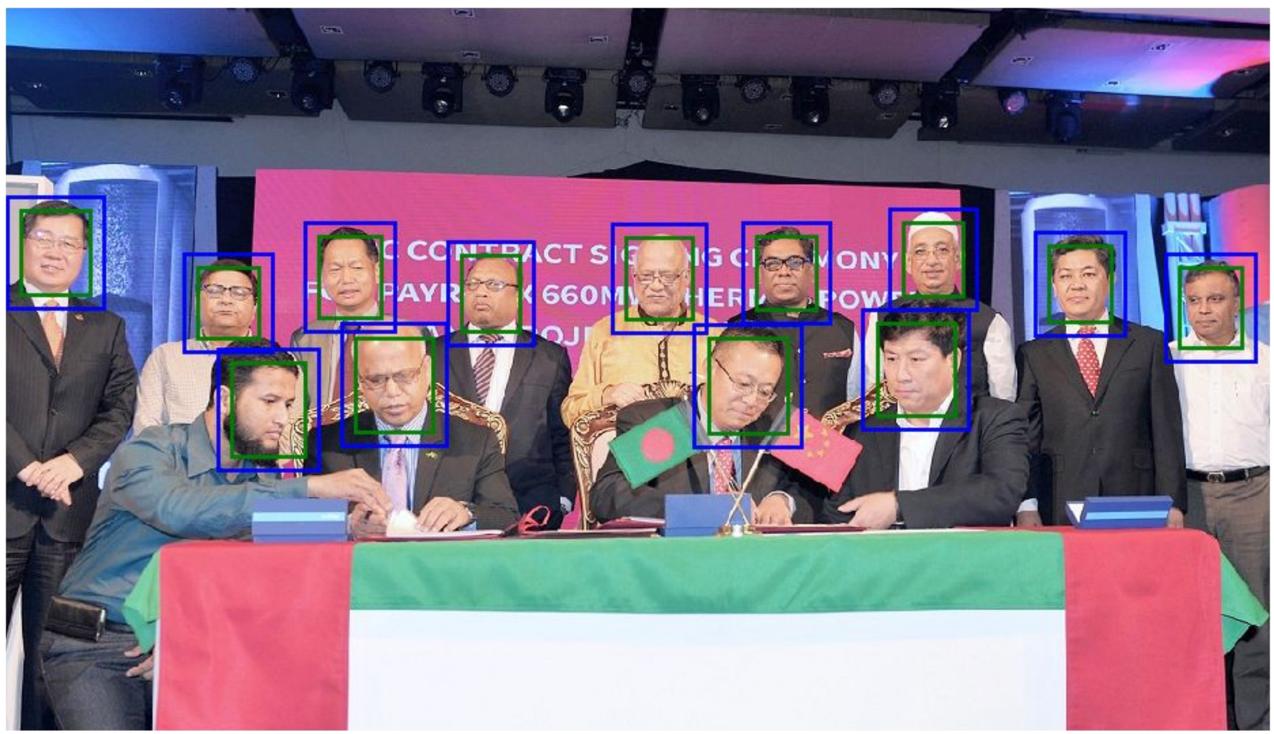

The M Matrix from Co-location

- Ordinary least squares regression on 7k paired embeddings
- Essentially, find M minimizing L2-norm between y and Mx
  - Where y and x are corresponding pairs of the same face instance

$$x, y \in \mathbb{R}^{512}, \text{ so } M \in \mathbb{R}^{512 \times 512}$$

- Using ~7k pairs and 6-fold cross-validation
  - ~6k training pairs, ~1k testing pairs
  - ~6k x 512 = ~3 million constraints for 262,144 parameters in mapping
- After finding M, the reserved testing pairs are used to evaluate the quality of the mapping (ROC curve on next slide).

Part 1: ROC for Affine Map

- For cross-fold tests cases the recognition ROC is shown using the affine mapping determined from the training samples.
- In this test, we expect often the same image of the same face are paired.
- However, the ROC is based upon comparing embeddings and so are measuring what we care about, namely are vectors comparable after mapping.

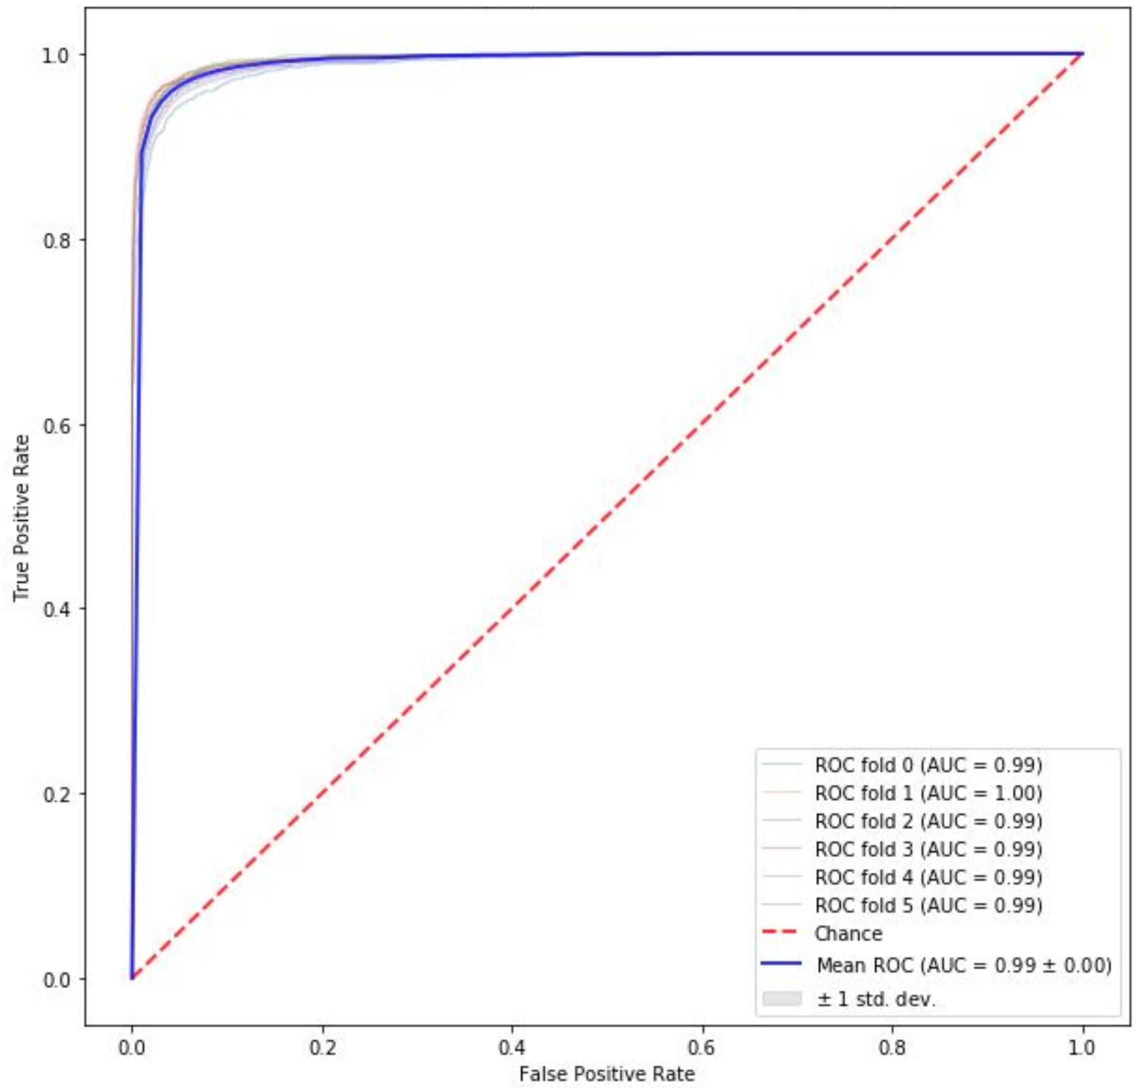

ROC curve for pairwise L2-distance classification applied to mapped feature vectors

Part 1: ROC No Mapping

- This is a bit of surprise!
- The same experiment as before but without using any mapping.
- At a False Positive Rate of 0.05 the True Positive Rate is 0.97 with mapping and 0.78 without.
- It appears that BBN and RPI embeddings, even without a mapping, are similar.

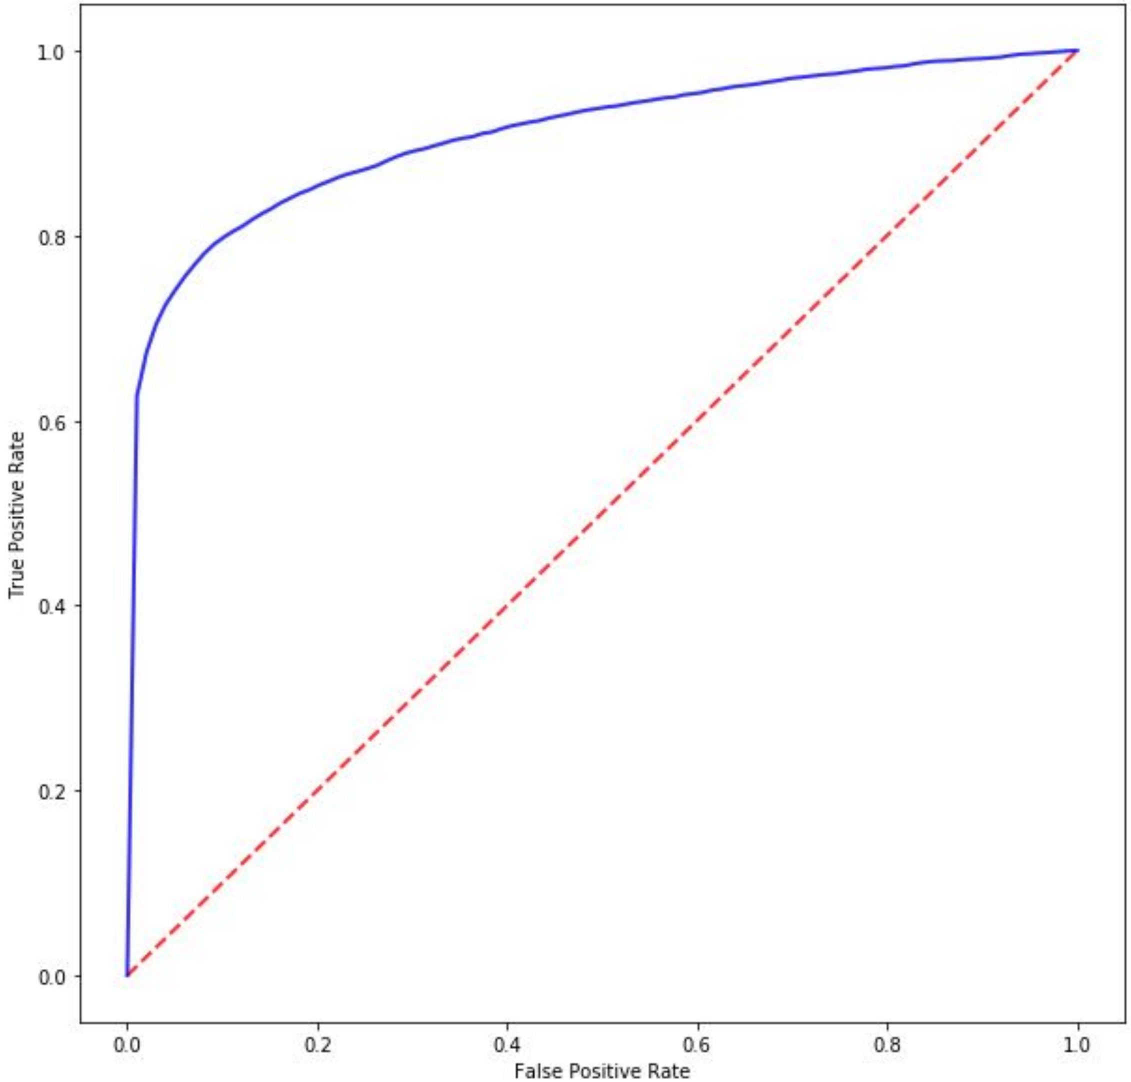

ROC curve for pairwise L2-distance classification (**w/o mapping**)

Computational Language and EducAtion Research
CLEAR

Part 2: Co-Reference

- Month 9 Pilot Eval is now un-sequestered and this includes images.
- Using the mapping found by co-location, we converted ALL 340k BBN embeddings into RPI space
- Then we compare each one to ALL 50k RPI embeddings
- Sort pairs based on match strength, excluding image pairs which correspond to locations in the same or very similar images
  - This step was done using ORB features and FLANN

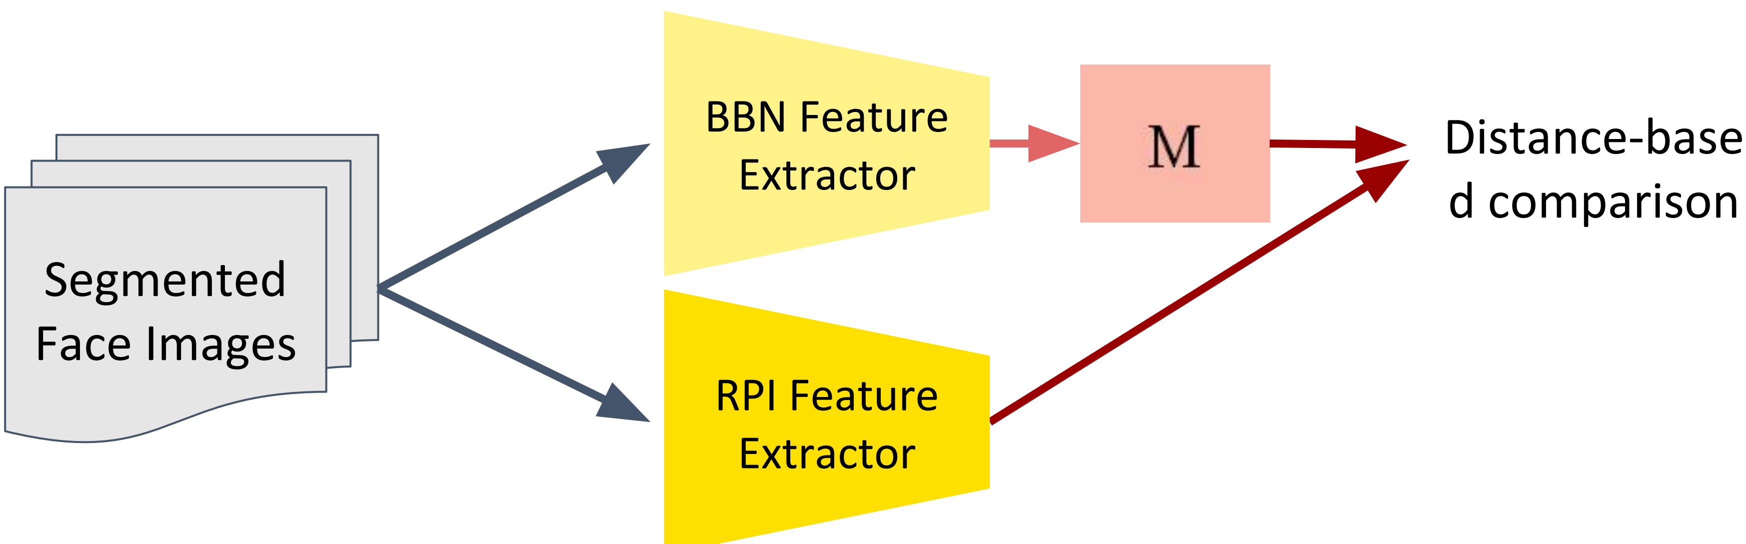

Computational Language and EducAtion Research
CLEAR

Part 2: Aung San Suu Kyi

- Here is one example - incumbent State Counsellor of Myanmar
- The pairing BBN to RPI was flagged because of embedding similarity
  - No name is provided by either BBN or RPI
- However, the association was automatic
  - We after the fact mapped this unknown paired finding with “Aung San Suu Kyi”

BBN embedding, mapped

RPI embedding

Part 2: Oles Alekseevich Buzina

- Ukrainian journalist and writer, murdered near his home in Kiev (unsolved)
- The BBN to RPI association was again found automatically
- We added the name later

BBN embedding, mapped

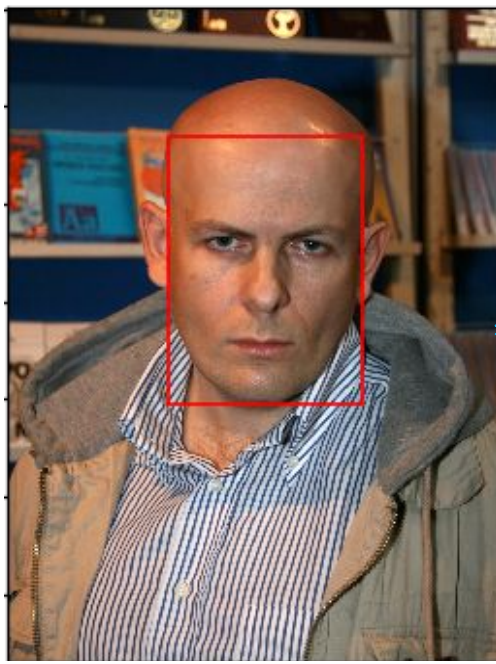

RPI embeddings

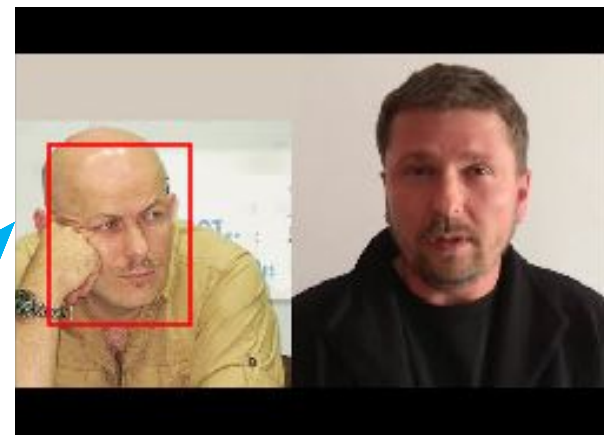

Computational Language and EducAtion Research
CLEAR

Next Steps

- These pilot studies are promising.
- Large-scale studies using CASIA-Webface, VGGFace2, and Labelled Faces in the Wild are underway now. Because of the scale of these experiments and known ground truth the results will go a long way to quantifying the reliability and utility of the mappings we are discovering.
- In the context of AIDA, to experiment with a broader range of embedding mapping tasks more data is needed.
- We should be braced for commercial biometrics systems working hard to obscure identity in their systems. In other words, now we have seen how easily we've mapped between BBN and RPI embeddings, it raises the broader question of if/when embedding obscuration may become a goal for some.
- Following upon the previous point aimed specifically at faces, there is an interesting larger scale issue about feature representations associated with image understanding and their use in context such as AIDA. Tensions will exist between commonality of representation (see what has happened with BBN and RPI) and secondary reasons for obscuration of mappings.

Annotating Images and Video

Questions raised at November PI meeting:

- Need exhaustive annotation for TA1s?
  - Cannot measure precision and recall
- How to exhaustively annotate video and images?
- Is precision/recall important for TA3s?

Solutions

- Use existing benchmarks when available (object detection, pedestrian tracking, etc.)
- Focus annotation on
  - Cross-document coreference resolution
    - Location
    - Time
    - Image-Text coreference
  - Unusual ontology types such as Events
    - Spatial relationships
  - Hypothesis generation

Example Annotations: Coreference

Location: Altamira Square, Caracas

Use external knowledge to train localization system.

Example Annotations: Event Detection

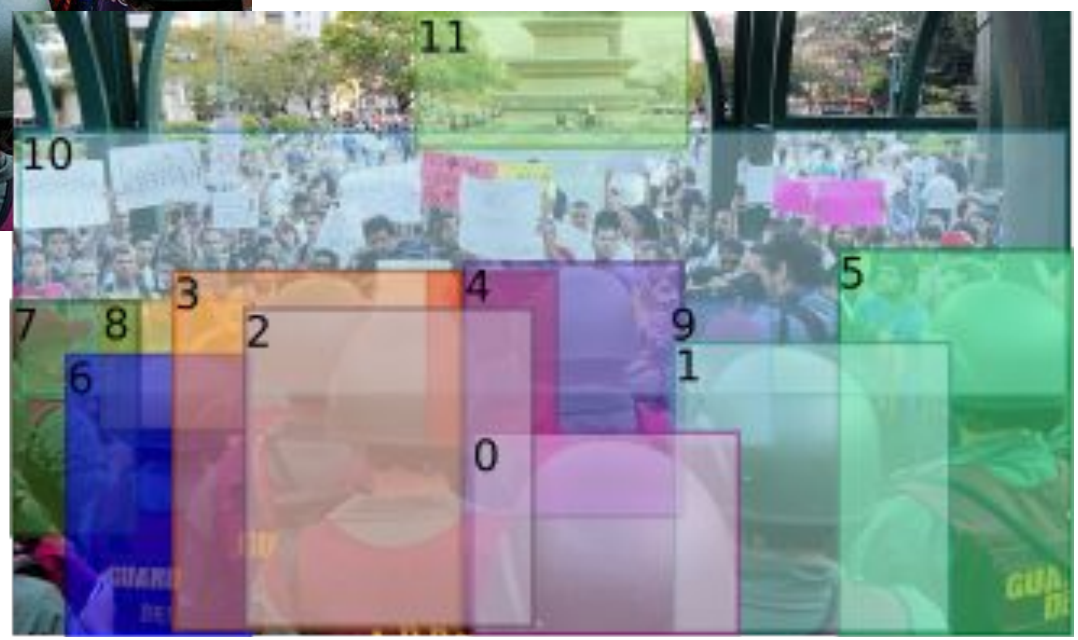

Spatial Relationships:
- Relative Depth
- Orientation

Computational Language and EducAtion Research
C L E A R

Example Annotations: Event Detection

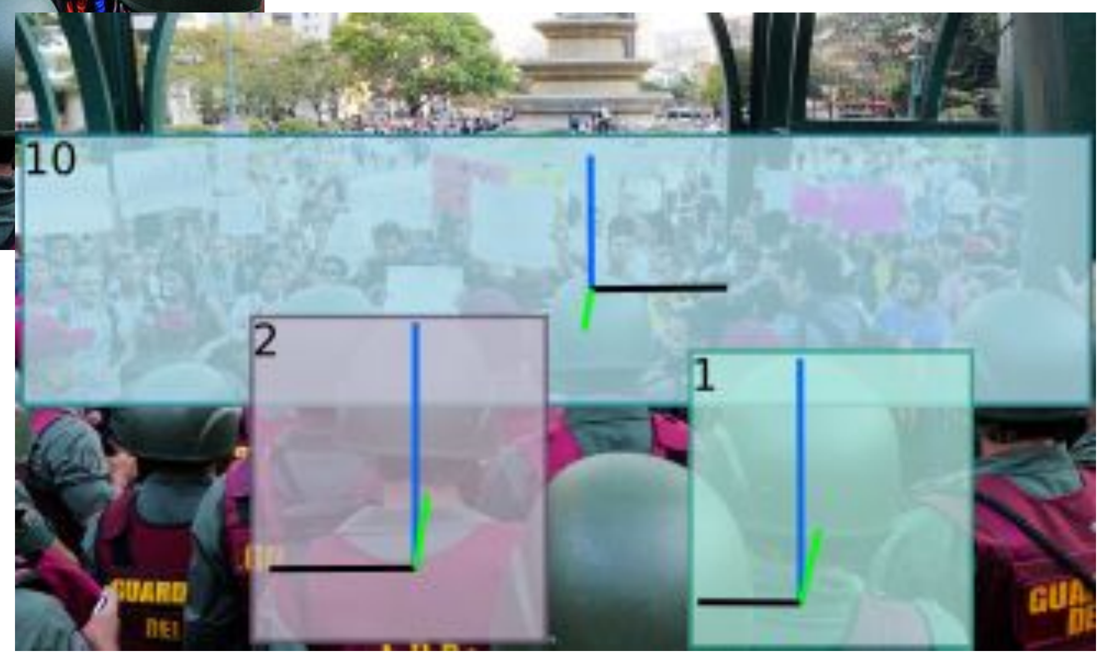

Spatial Relationships:

- Relative Depth
- Orientation

Example Annotations: Event Inference

Spatial Relationships:
- Relative Depth
- Orientation

Vox ML Annotation Task

Caption	Focus Activity(ies)	Objects/entities in scene	Spatial Relations	Depicted Activities	Potential Activities	Changes in Circumstance
A man and woman drinking together at a restaurant	Man drinking from glass	Man, woman, glasses, cups, bottles, sunglasses	Man holding glass, woman beside man, bottle in front of man, bottles behind woman, ...	**glass**: man drinking from glass, man holding glass; **cup**: woman holding cup	**cup**: woman drinking from cup; **glass**: man setting down glass; **bottle**: man/woman drinking from bottle...	**Woman drinking from cup:** to drink from the cup, the woman must lift it to her mouth; ...

Ontology Effort –

Past, Present & Future

Colorado: AIDA Working Ontology

Collaboration with LDC on Annotation Ontology

1. Changes to existing types and how they are annotated
2. New entity or event types related to those changes
3. New types for the Venezuela scenario
4. Cross-over from AIDA ontology to KAIROS ontology

Allowing Events as Arguments to Events

In two limited contexts:

Justice events: Types that fill the existing Crime role now can be existing annotated events or a new GenericCrime event type.

Some Contact events: A new Topic role allows annotated events as fillers

New GenericCrime event

Crime changed from an entity type to an event

We defined arguments and allowable “fillers” for those arguments

Crime’s subtypes were eliminated:

- BehavioralCrime
- FinancialCrime
- PoliticalCrime

Event types with new Topic role

- CommandOrder
- CommitmentPromiseExpressIntent
- Negotiate
- Prevarication
- RequestAdvise
- ThreatenCoerce

Still No Topic role

- Collaborate
- Discussion
- FuneralVigil
- MediaStatement
- PublicStatement

Either because the role seemed inappropriate or because including it would make annotation too difficult

New entity type for Topic role

- Most topics have no established annotation ontology type, e.g., “dogs” or “to read a book”.
- We proposed InformationObject, which stands for any kind of proposition in discourse.
- LDC considered it too open ended for exhaustive annotation - countered with a subtype

InformationObject.TopicFiller

- An ad-hoc category
- An InformationObject that is a filler for the Topic role in
  - CommandOrder,ThreatenCoerce,
  - CommitmentPromiseExpressIntent,
  - Negotiate, Prevarication, RequestAdvise,

i.e., events that have no other annotatable ontology type

- Not a coherent category in the real world

New types for Venezuela scenario

17 new types across entities, events and relations

Coverage for:

- Disease outbreaks
- Coups
- Drone control
- Hoaxes/Fraud

New entity types

Medical/Health Condition or Issue (MHI)

MHI.Disease

MHI.SymptomPresentation

PER.ProfessionalPosition.MedicalPersonnel
(extended and renamed Paramedic)

VEH.Aircraft.Drone

New event types: Disease scenario

Life.Injure.IllnessDegradation

Medical.Intervention

Disaster.DiseaseOutbreak

ArtifactExistence.Shortage

Extension of Life.Injure and Life.Die to include Agent role and MedicalHealthIssue role

New event types: Govt. upheaval

Conflict.Coup

Government.Convene

New event types: Drone explosions

ArtifactExistence.ArtifactFailure

Inspection.TargetAimAt

Annotation post-processing

Elizabeth Spaulding & Susan Brown

Computational Language and EducAtion Research

C L E A R

Problems with LDC's annotation schemes

- Generic person type PER has subtypes which are more like roles
  - PER.Ambassador
- Goal: change these PER subtypes to relations
  - PER.Ambassador → PER "has_role" Ambassador

- "Prevent" subtypes name an event which doesn't happen
  - Mvmt.PreventEntry, Vote.PreventVote
- Goal: create new "Prevent" type
  - Mvmt.PreventEntry → Prevent.PreventEntry
  - Vote.PreventVote → Prevent.PreventVote

Annotation files

- arg_mentions.tab
- rel_mentions.tab
- rel_slots.tab
- evt_mentions.tab
- evt_slots.tab

arg_mentions.tab

root_uid	argmention_id	text_string	description	type	subtype	subsubtype
IC0015LNI	EMIC0015LNI.000748	anti-junta	anti-junta	sid	ideological	ideological
IC0015LNI	EMIC0015LNI.000761	anti-junta activists	anti-junta activists	per	protester	unspecified

Computational Language and EducAtion Research
C L E A R

rel_mentions.tab

root_uid	relationmention_id	text_string	description	type	subtype
IC0015LNI	RMIC0015LNI.000024	anti-junta	Anti junta activists	generalaffiliation	memberoriginreligionethnicity

Each relation must refer to a specific string in the text

rel_slots.tab

root_uid	relationmention_id	slot_type	argmention_id
IC0015LNI	RMIC0015LNI.000024	rel014arg02entityorfiller	EMIC0015LNI.000748
IC0015LNI	RMIC0015LNI.000024	rel014arg01person	EMIC0015LNI.000761

Computational Language and EducAtion Research
C L E A R

Annotation changes

- How do we structure our changes without disrupting the format of the annotation files?
  - For the PER subtype issue, the solution would not be simple mapping - we may have to create new relation mentions?
- Mvmt.PreventEntry → Prevent.PreventEntry may be an easier change
  - We wouldn't have to create new mentions

Overlap between KAIROS/AIDA Ontologies

Martha Palmer
University of Colorado

Computational Language and EducAtion Research
CLEAR

AIDA Ontology background

- AIDA Program Ontology
  - Used sporadically by some team members
- AIDA Annotation Ontology
  - The primary source of ERE for AIDA
  - Very fine-grained because of TA2 constraints
-
- Open Question - Will AIDA performers be ok with changes to Program Ontology?
  - So far, yes

KAIROS Ontology effort

- Focus primarily on defining Event Primitives
- Reusing many AIDA event types,
  - often at a more coarse-grained level,
  - Ex. Transaction-Exchange
- Creating new event types as well
  - Necessitates new entity types for slot fillers
  - Trying to borrow from AIDA
  - Highlighting issues with AIDA event taxonomy

Major Areas for initial primitives

- Construction
- Transactions
- Conflict
- Contact/Communication
- Movement
- Life events
- Health

AIDA Event oddities

- Resistance to Events as Event arguments,
  - ex., Crime
  -
- Emphasis on Relations between Events, such as Causal and Temporal

Label	Illness, sickness
Description	a disease or period of sickness affecting the body or mind
Parents/Domain	Life

Slot Role	Slot constraints
Target/Victim	PER,
Means (cause)	Event, WEA (biological),
Agent (in case of WEA)	PER, ORG, GPE
Disease	MedicalCondition

Temporal	Start/End= point , Duration = range

(*) ”The FBI concluded that the killers were victims of mental illness” ✓ ✓
(*) “early release on compassionate grounds of prisoners with terminal illnesses” ✓ ✓

Not a common event in IED

AIDA Event oddities

- Resistance to Events as Event arguments,
  - ex., Crime
  -
- Emphasis on Relations between Events, such as Causal and Temporal
- LDC refusal to annotate Cause as an argument slot - In contrast w/ preference for annotating PER.Protestor rather than as a RELATION

Computational Language and EducAtion Research
C L E A R

AIDA Entity oddities

FAC - GeographicalArea - Border Checkpoint

LOC - GeographicPoint - Address

AIDA Entity oddities - resolved

FAC - GeographicalArea - Border
Checkpoint

(could inherit from FAC & area)

LOC - GeographicPoint - Address

(could inherit from LOC & point)

Cross-program Ontology - Goals

- Priorities for multimodal information are key
- Rational entity ontology
  - appropriate properties as Relations
- Rational Event ontology
  - broad coverage
- Principled approach to multiple inheritance
  - Curation via cross-linguistic metonymy

Multiple Inheritance curated via Cross-lingual Metonymy

James Pustejovsky
Brandeis University

Computational Language and EducAtion Research
C L E A R

Multiple Inheritance curated via Cross-lingual Metonymy

- Metonymic types (below) justify multiple inheritance
- Functional types (*president, pilot, driver, protester, bomber*) do not.

INFO · PHYSOBJ	article, book, email
CONTAINER · CONTENT	bottle, bucket, spoon
PRODUCER · PRODUCT	Honda, Apple
ORG · (INFO · PHYSOBJ)	newspaper, magazine
ORG · LOC · HUMANGROUP	university, city
GOVORG · CAPITALLOC	Moscow, London
EVENT · INFO	lecture, exam
EVENT · HUMANGROUP	class
ANIMAL · FOOD	chicken, lamb, fish

Data and Evidence for Cross-lingual Metonymy

- Peters, Wim. "Metonymy as a cross-lingual phenomenon." In Proceedings of the ACL 2003 workshop on Lexicon and figurative language-Volume 14, pp. 1-9. Association for Computational Linguistics, 2003.

- Srinivasan, Mahesh, and Hugh Rabagliati. "How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy." Lingua 157 (2015): 124-152.
  English Cantonese Farsi French Hindi Hungarian Indonesian Italian Japanese Korean Mandarin Russian Spanish Turkish Vietnamese

- Zhu, Huichun, and Barbara C. Malt. "Cross-linguistic evidence for cognitive foundations of polysemy." In Proceedings of the Annual Meeting of the Cognitive Science Society, vol. 36, no. 36. 2014.
  English Chinese

- Alonso, Hector Martinez, Bolette Sandford Pedersen, and Nuria Bel. "Annotation of regular polysemy and underspecification." In Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp. 725-730. 2013.
  Danish Spanish English

- Pustejovsky, James, Anna Rumshisky, Jessica L. Moszkowicz, and Olga Batiukova. "GLML: Annotating argument selection and coercion." In Proceedings of the Eighth International Conference on Computational Semantics, pp. 169-180. Association for Computational Linguistics, 2009.
  English Spanish Italian

- Boleda, Gemma, Sabine Schulte im Walde, and Toni Badia. "Modeling regular polysemy: A study on the semantic classification of catalan adjectives." Computational Linguistics 38, no. 3 (2012): 575-616.
  Catalan

- Martinez Alonso, Hector. "Annotation of regular polysemy: an empirical assessment of the underspecified sense." PhD diss., Universitat Pompeu Fabra, 2013.
  English Danish Spanish

- Babarinde, Olusanmi. "Lexical Ambiguity In Yoruba: Its Implications For Second Language Learners". Journal of Languages, Linguistics and Literary Studies (JOLLS) Volume 5. June 2018
  Yoruba

Dot Types and Polysemy

- The lunch$_{\text{EVENT}}$ lasted two hours. But it$_{\text{ENTITY}}$ was delicious
- The book$_{\text{ENTITY}}$ cost \$25.00. And it$_{\text{PROPOSITION}}$ is confusing!

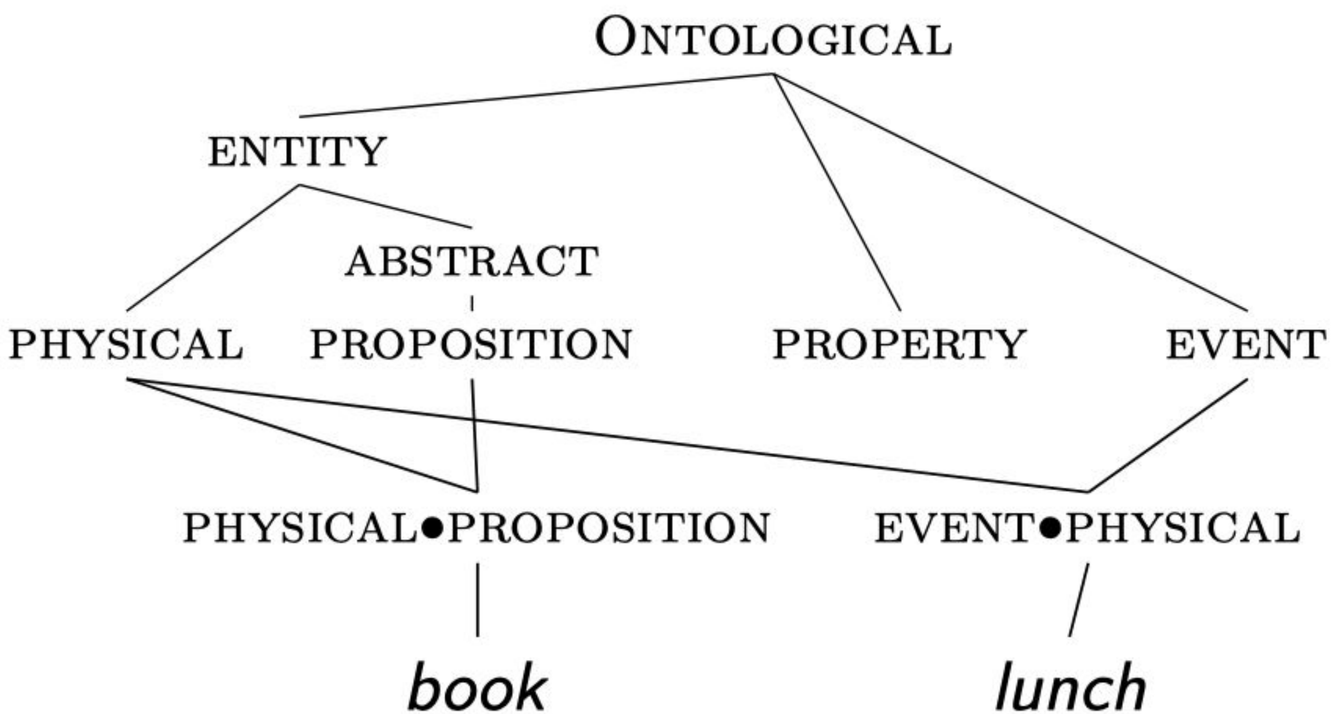

Computational Language and EducAtion Research
C L E A R

Cross-linguistic Examples

- Hindi

Letter (PhysObj*Info):
I received your **letter** = Mujhē tumhārā **patr** milā
(मुझे तुम्हारा **पत्र** मिला)
I liked your **letter** = Mujhē tumhārā **patr** pasand āyā
(मुझे तुम्हारा **पत्र** पसंद आया)

Lecture (Event*Info):
I liked your lecture = Mujhē tumhārā **vyākhyān** pasand āyā
(मुझे तुम्हारा **व्याख्यान** पसंद आया)
I learned from the lecture = Mainē **vyākhyān** sē sīkhā
(मैंने **व्याख्यान** से सीखा)
I began a new lecture = Mainē ēk nayā **vyākhyān** śurū kiyā
(मैंने एक **व्याख्यान** शुरू किया)
I finished the lecture = Mainē **vyākhyān** khatm kar diyā
(मैंने **व्याख्यान** खत्म कर दिया)

Cross-linguistic Examples

- German, Italian

Container*Content - German

Tim drank another glass. = Tim trank noch ein **Glas**.

Tim bought a handblown glass. = Tim kaufte ein mundgeblasenes **Glas**.

Information*Physobj - Italian

He grabbed the book I was handing to him. = Afferro il **libro** che gli stavo porgendo.

It is impossible to summarize this book. = E impossibile riassumere questo **libro**.

AIDA Data Model for Metonymy

- Ontological types should apply cross-linguistically
- Dot types (metonymies) have cross-linguistic justification as inheriting from multiple superordinates
- If we can determine which dot types reflect metonymy in a given language, then we can link metonymic behavior to the ontology types and guide transfer learning across language that way

COVID-19 - Brandeis

Demo of Semantic Visualization over

- Heng Ji Blender CORD-19 data
- Brandeis/Harvard/SIFT/IHMC Covid-19 Dataset