(Open) Information Extraction: Where are we going?

SAPIENZA
UNIVERSITÀ DI ROMA

Claudio Delli Bovi
16/10/2015

Linguistic Computing Laboratory

About me

dellibovi@di.uniroma1.it

http://wwwusers.di.uniroma1.it/~dellibovi

bn:17381128n

First-year PhD student

LCL group @ Sapienza

Advisor: prof. Roberto Navigli

Focus (so far): (Open) Information Extraction

Outline

IE and OIE: some background

Outline

IE and OIE: some background

DefIE: OIE from textual definitions

Delli Bovi, Telesca, Navigli: **TACL** (to appear)

Outline

IE and OIE: some background

DefIE: OIE from textual definitions
Delli Bovi, Telesca, Navigli: **TACL** (to appear)

KBUnify: KB disambiguation and unification
Delli Bovi, Espinosa-Anke, Navigli: **EMNLP 2015**

Information Extraction

“A process of getting **structured** data from **unstructured** information in the text”

(Jurafsky and Martin, 2009)

“Identification of instances of a particular class of **relationships** in a natural language text, and the extraction of relevant **arguments** for that relationships”

(Grishman, 1997)

Information Extraction

Why?

Information Extraction

Why?

Machine Reading:

“*I hereby offer to bet anyone a lobster dinner that by 2015 we will have a computer program capable of automatically reading at least 80% of the factual content across the entire English speaking web, and placing those facts in a structured knowledge base.*”

(T. Mitchell. Reading the Web: A Breakthrough Goal for AI. AI Magazine, 2005)

Information Extraction

What?

Information Extraction

What?

Input:

- large corpus of unstructured text
- set of semantic relations
- labelled training data

Output:

- knowledge base of triples

⟨ **entity, relation, entity** ⟩

Information Extraction

What?

Input:

- large corpus of unstructured text
- set of semantic relations
- labelled training data

Output:

- knowledge base of triples

⟨ **entity, relation, entity** ⟩

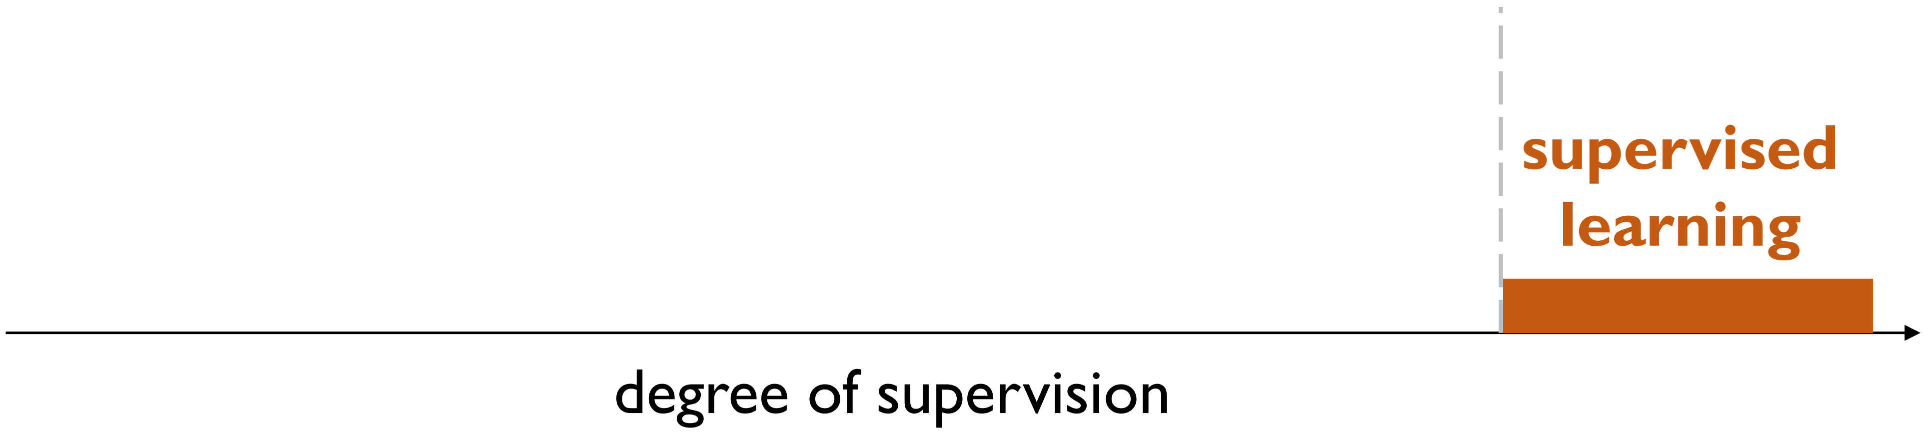

Information Extraction

What?

Input:

- large corpus of unstructured text
- set of semantic relations
- high-precision seeds/examples

Output:

- knowledge base of triples

⟨ **entity, relation, entity** ⟩

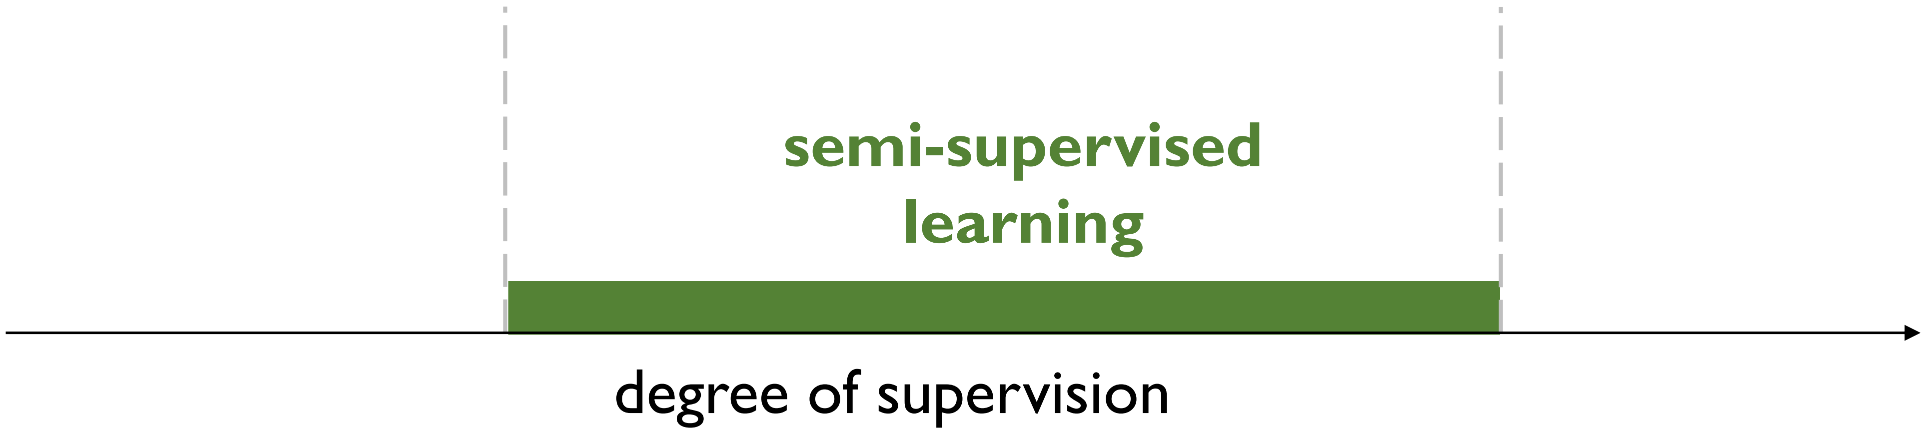

Information Extraction

What?

Input:

- large corpus of unstructured text

~~- set of semantic relations~~

Output:

- knowledge base of triples

⟨ **entity, relation, entity** ⟩

- set of semantic relations

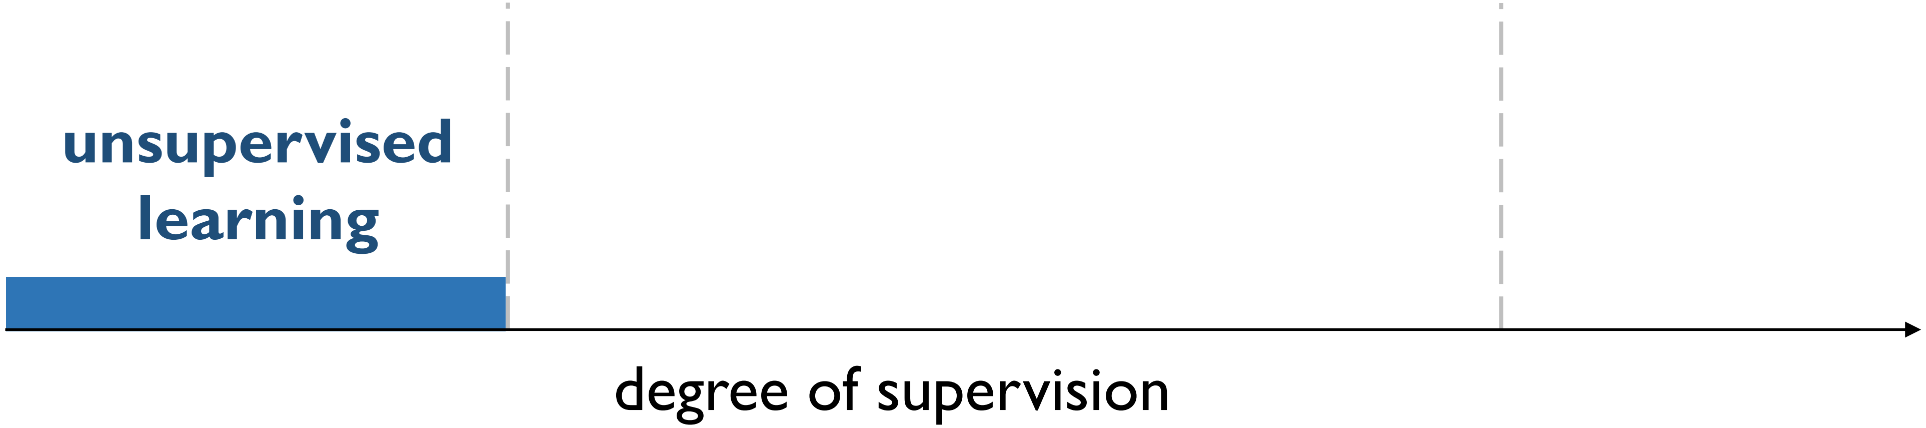

Information Extraction

How?

Information Extraction

How?

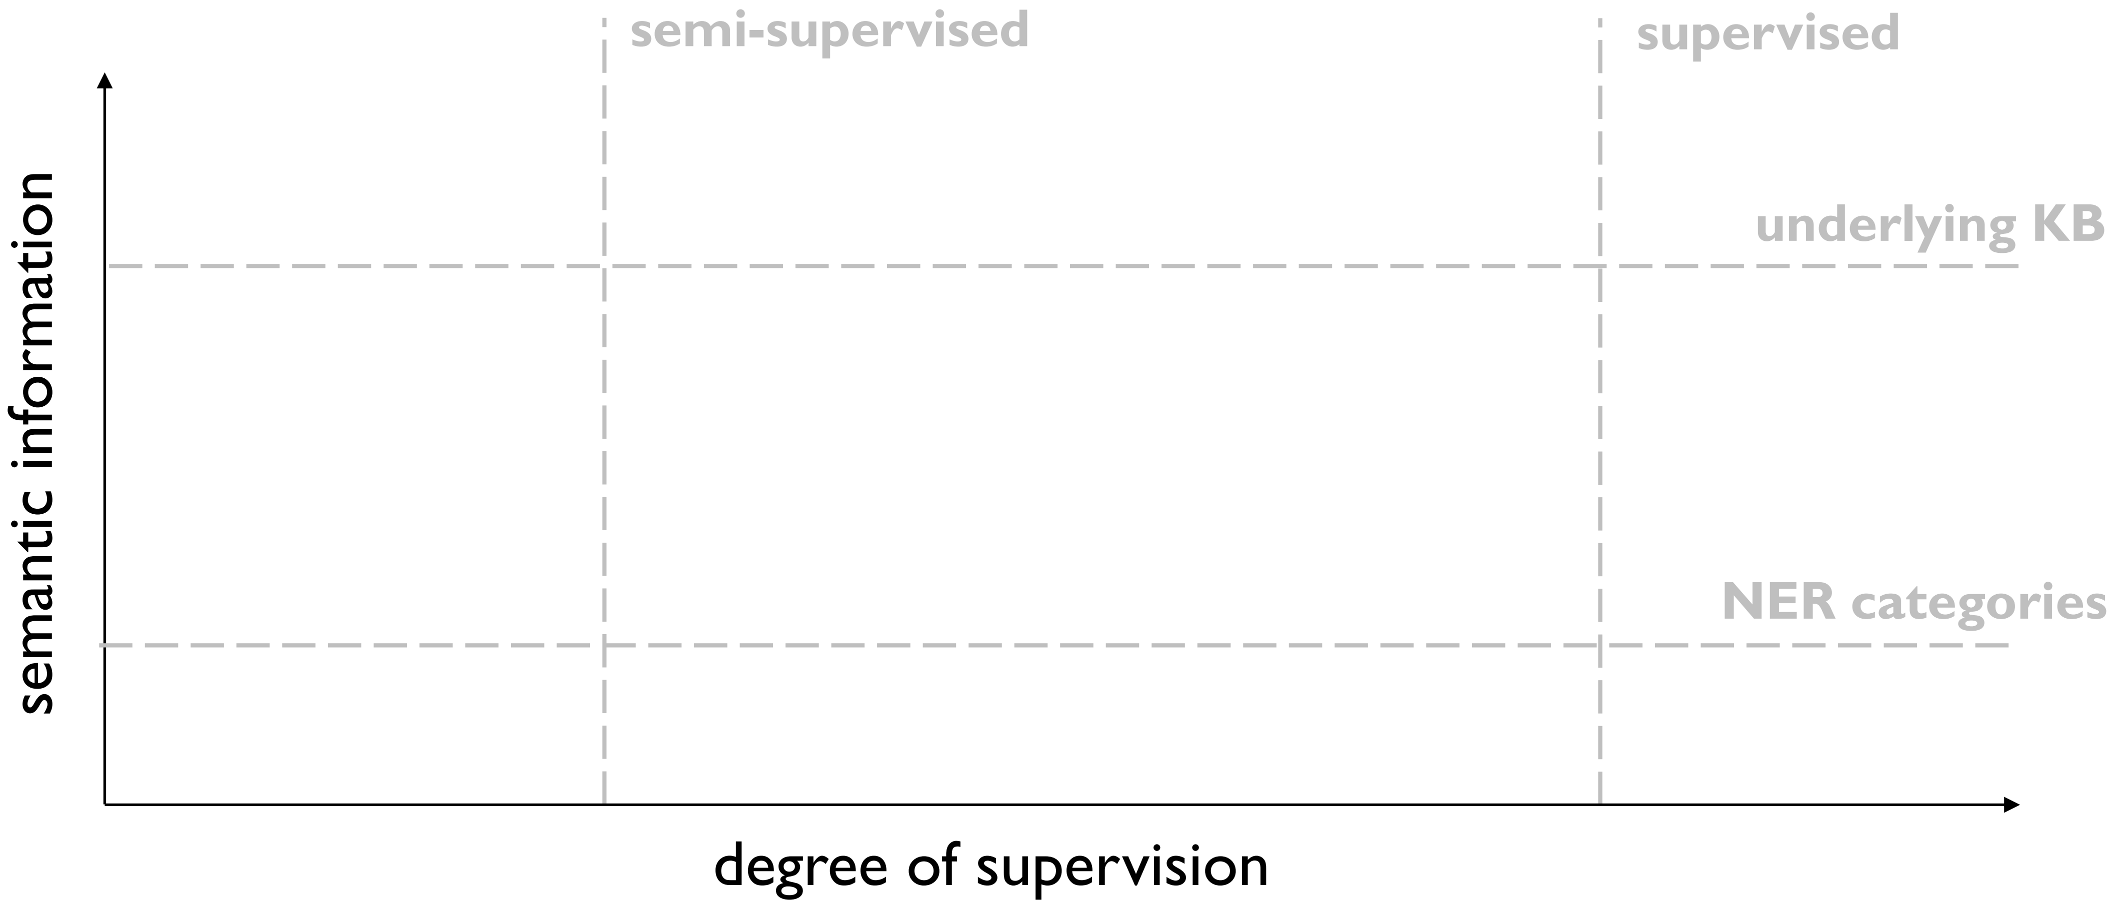

Information Extraction

How?

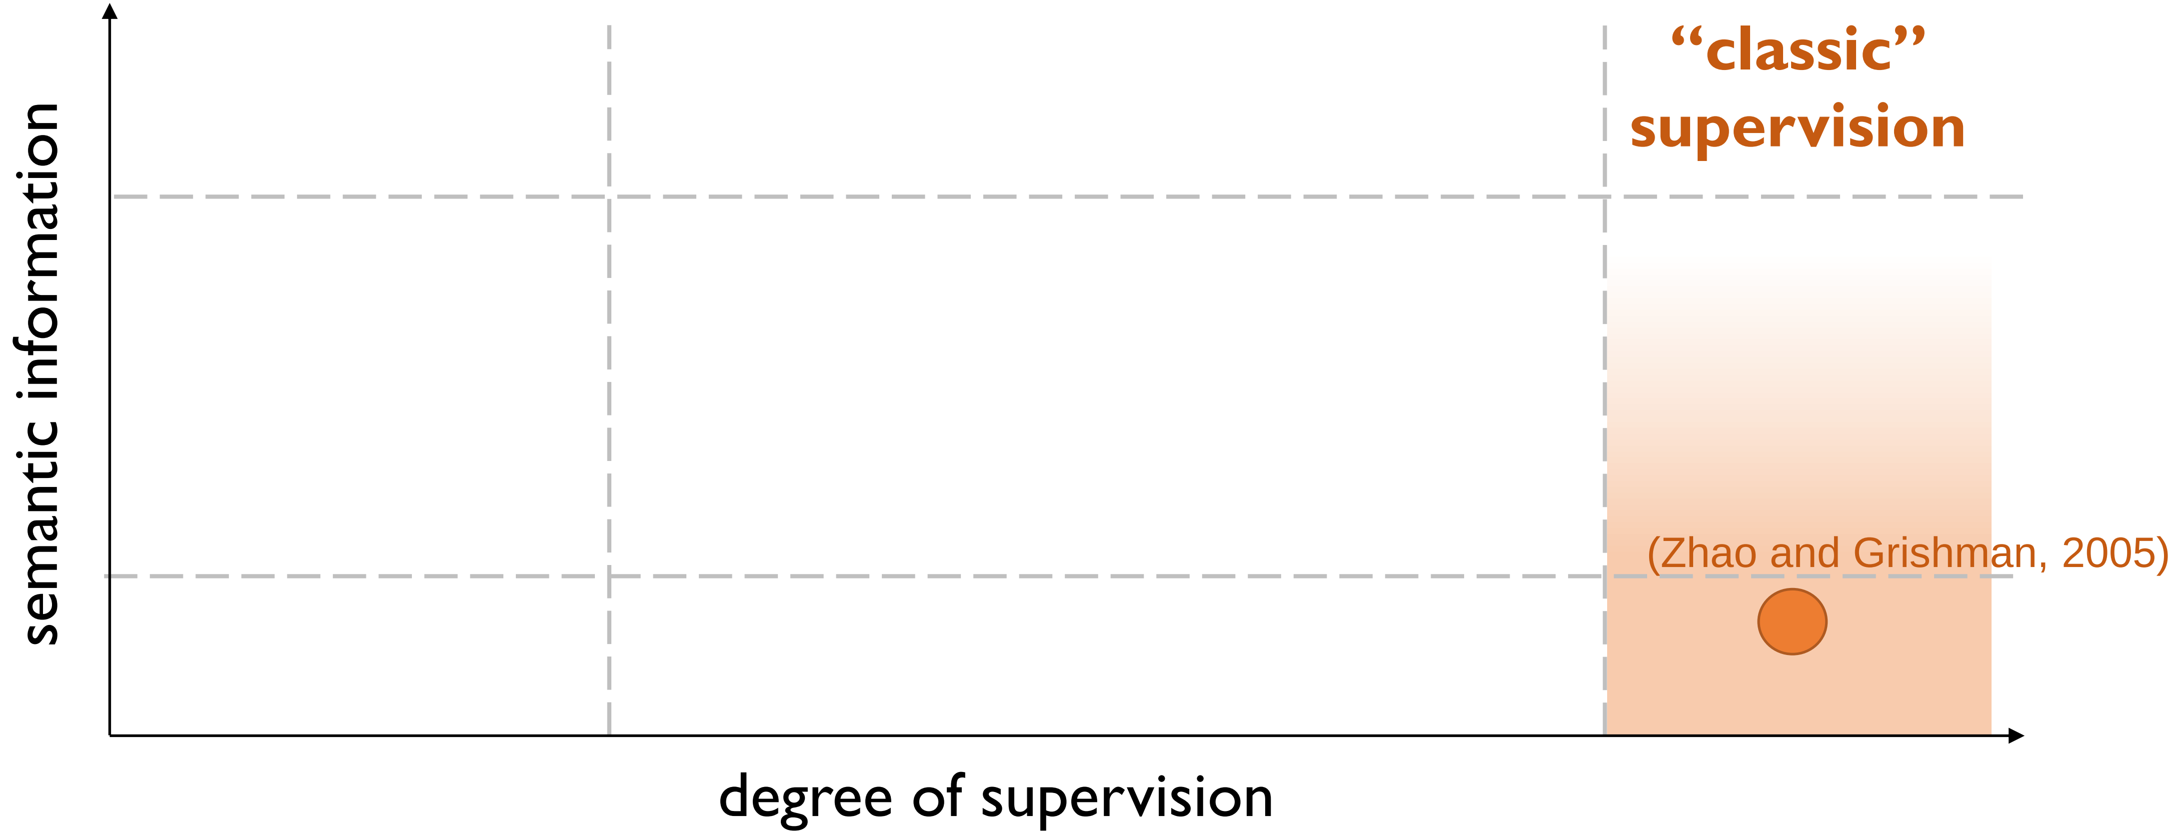

Information Extraction

How?

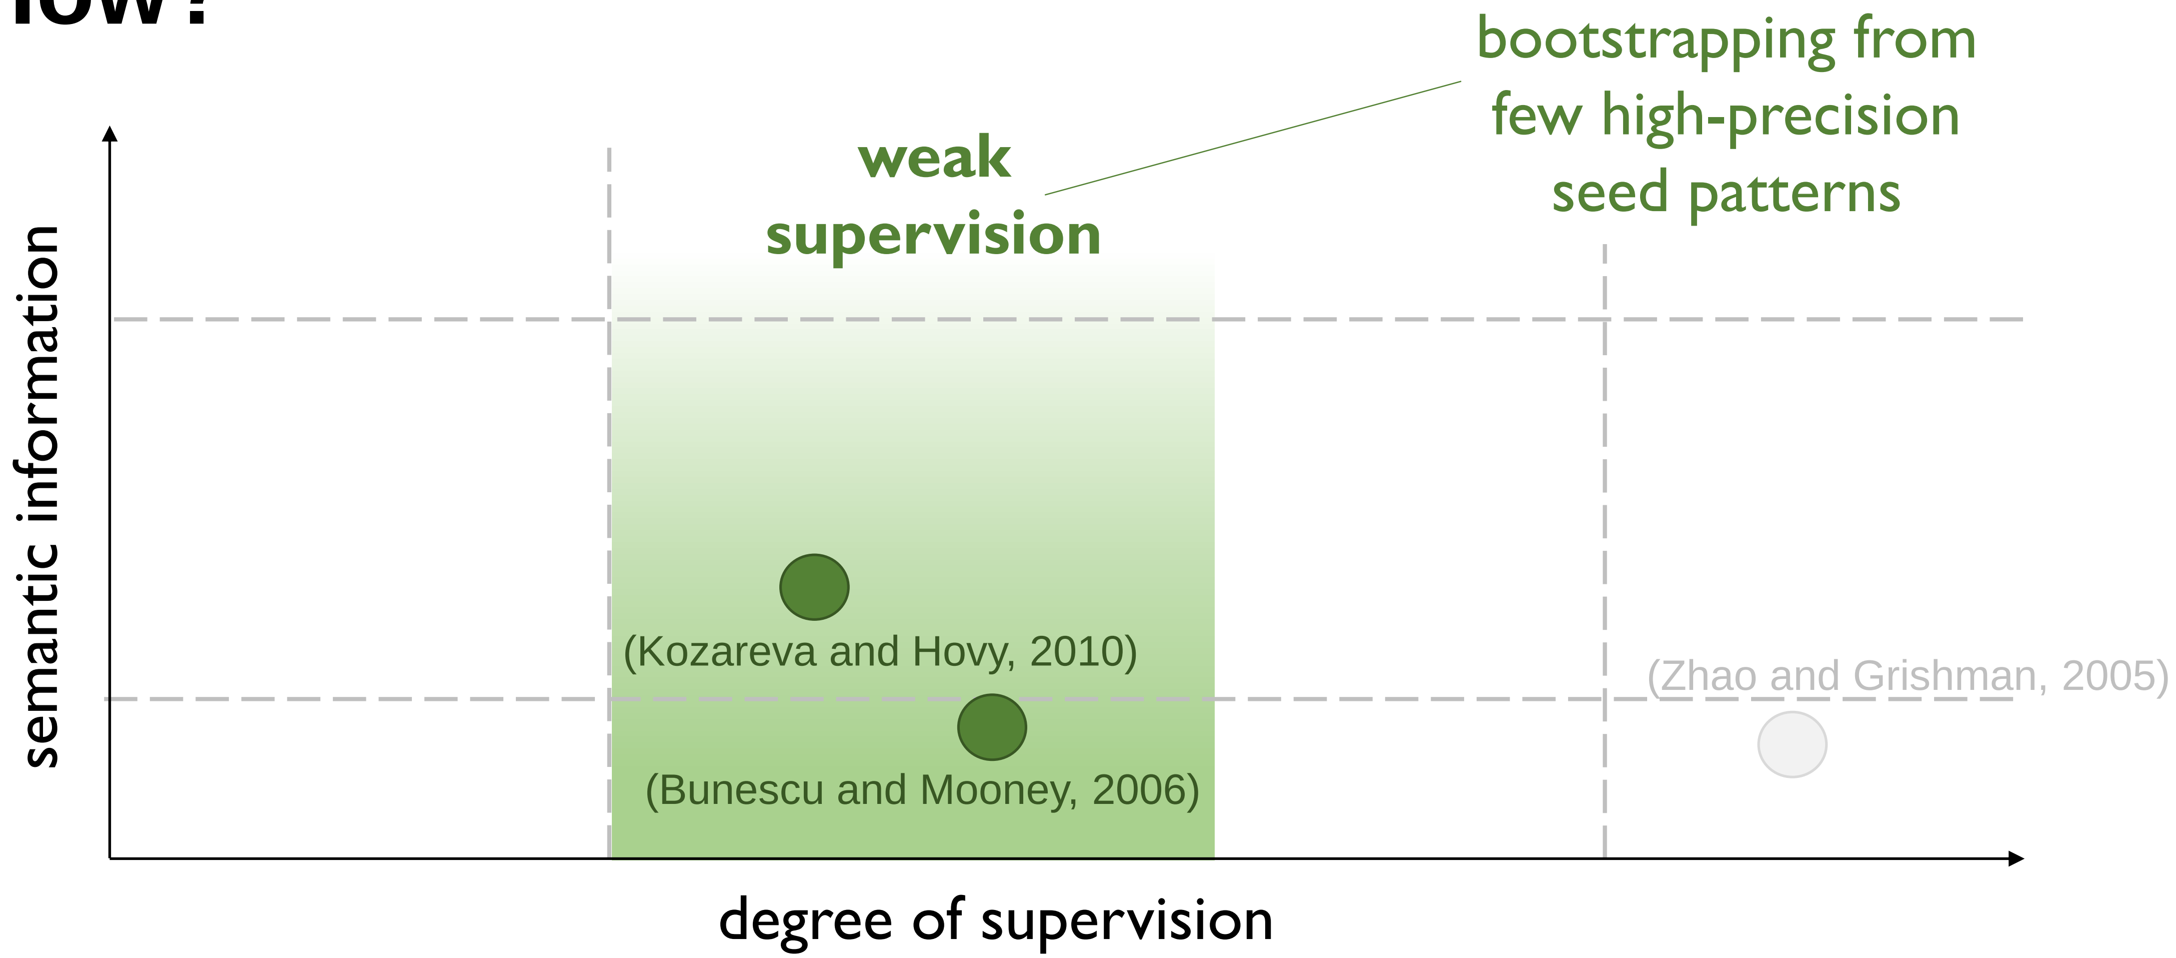

Information Extraction

How?

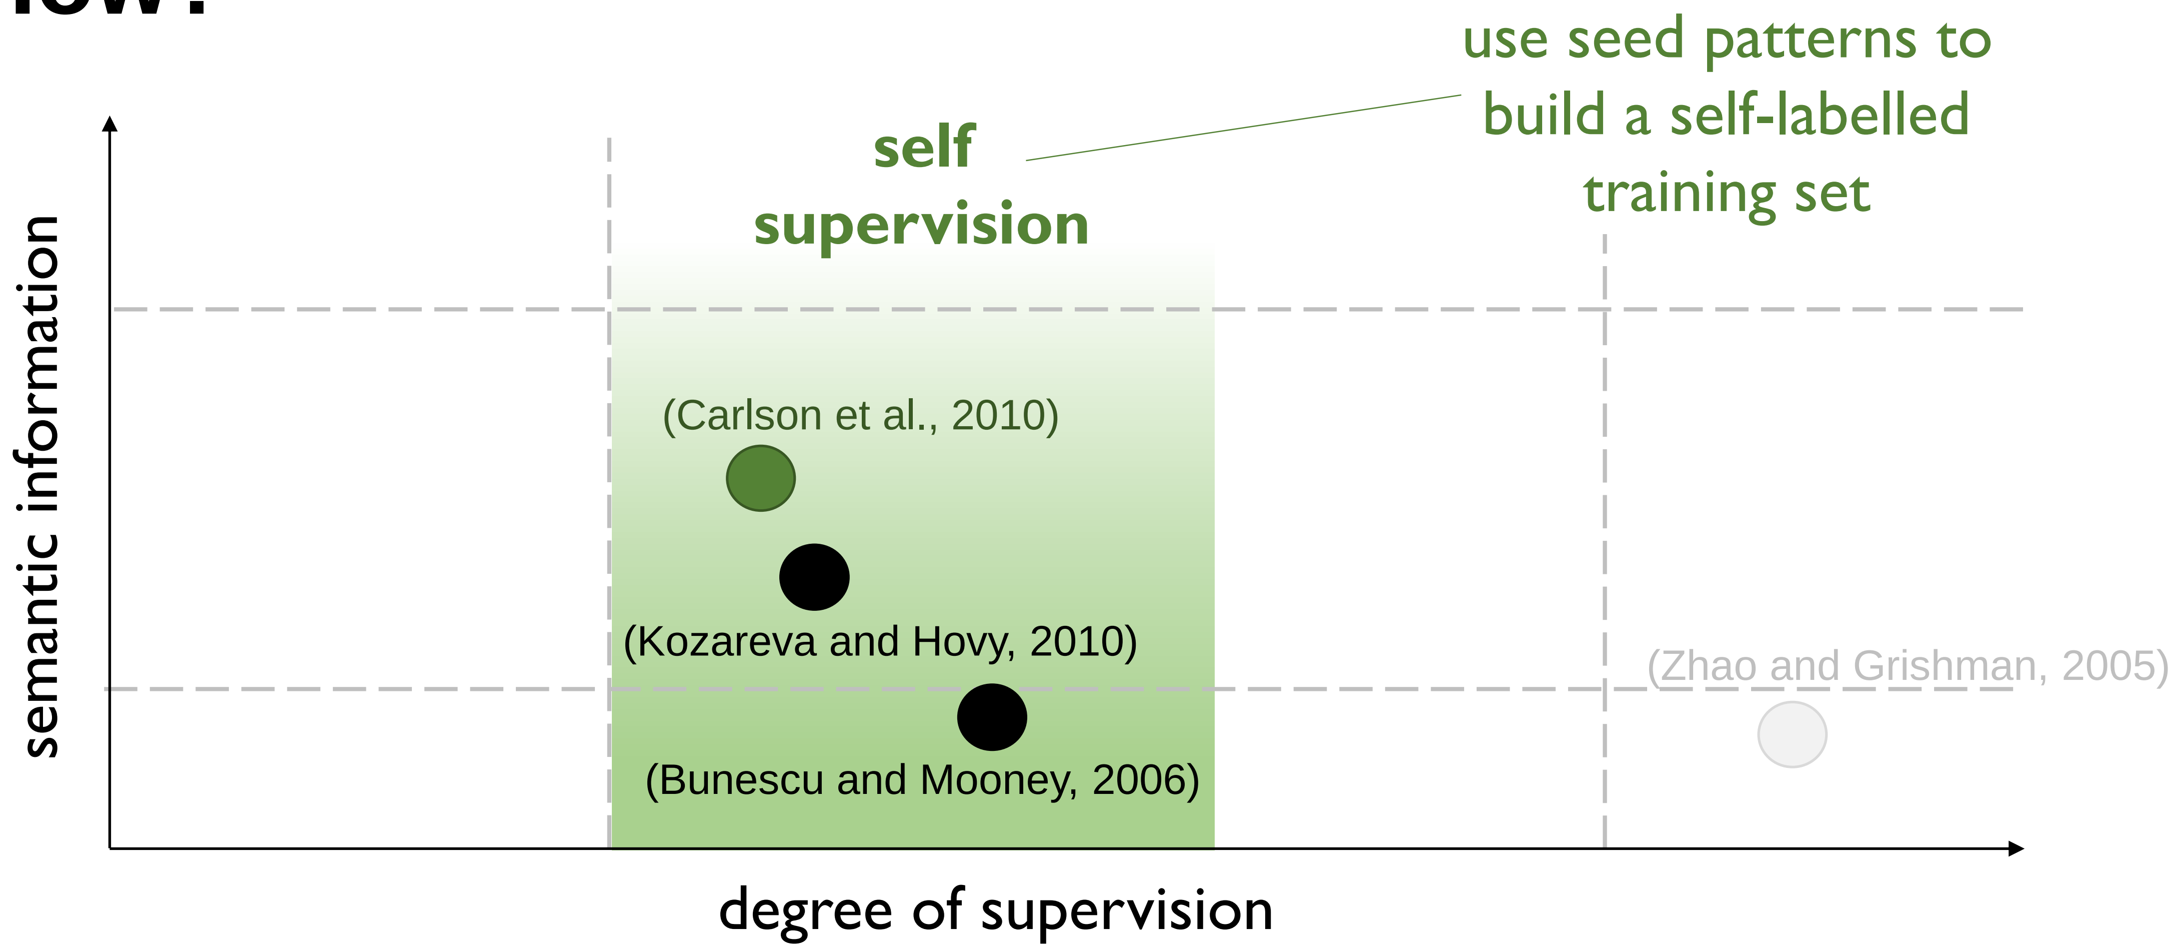

Information Extraction

How?

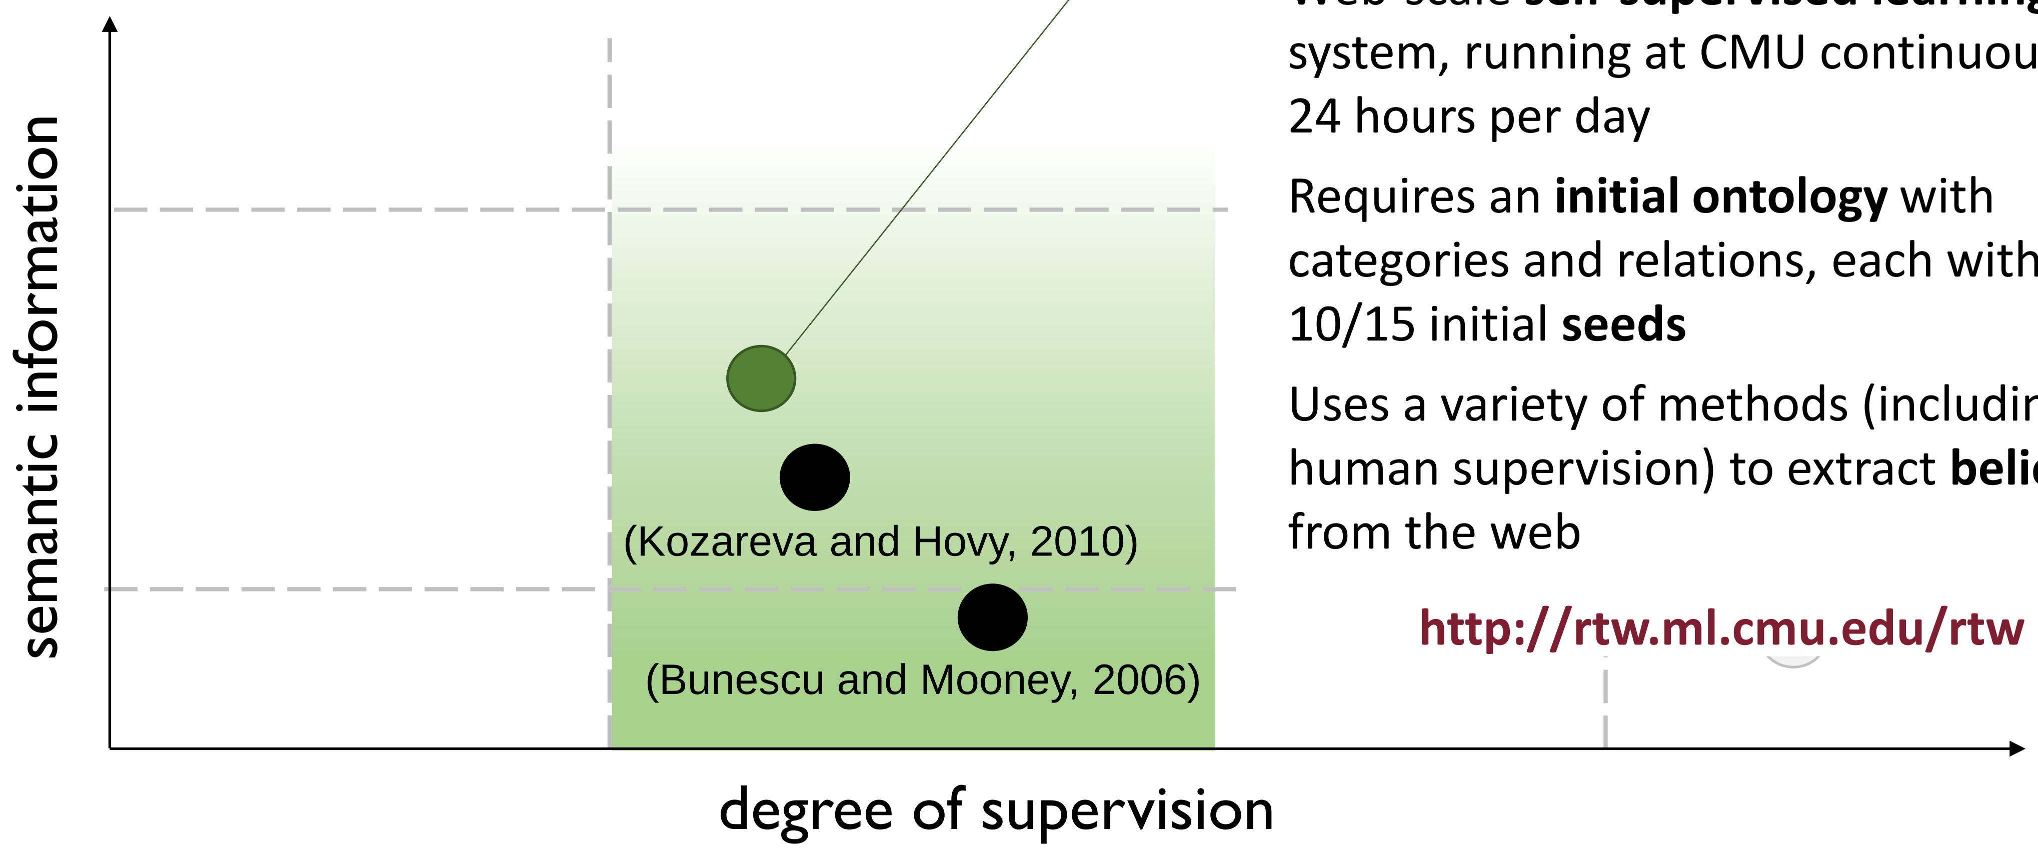

NELL – Never Ending Language Learning (Carlson et al., 2010)

Web-scale **self-supervised learning** system, running at CMU continuously 24 hours per day

Requires an **initial ontology** with categories and relations, each with 10/15 initial **seeds**

Uses a variety of methods (including human supervision) to extract **beliefs** from the web

http://rtw.ml.cmu.edu/rtw

Information Extraction

How?

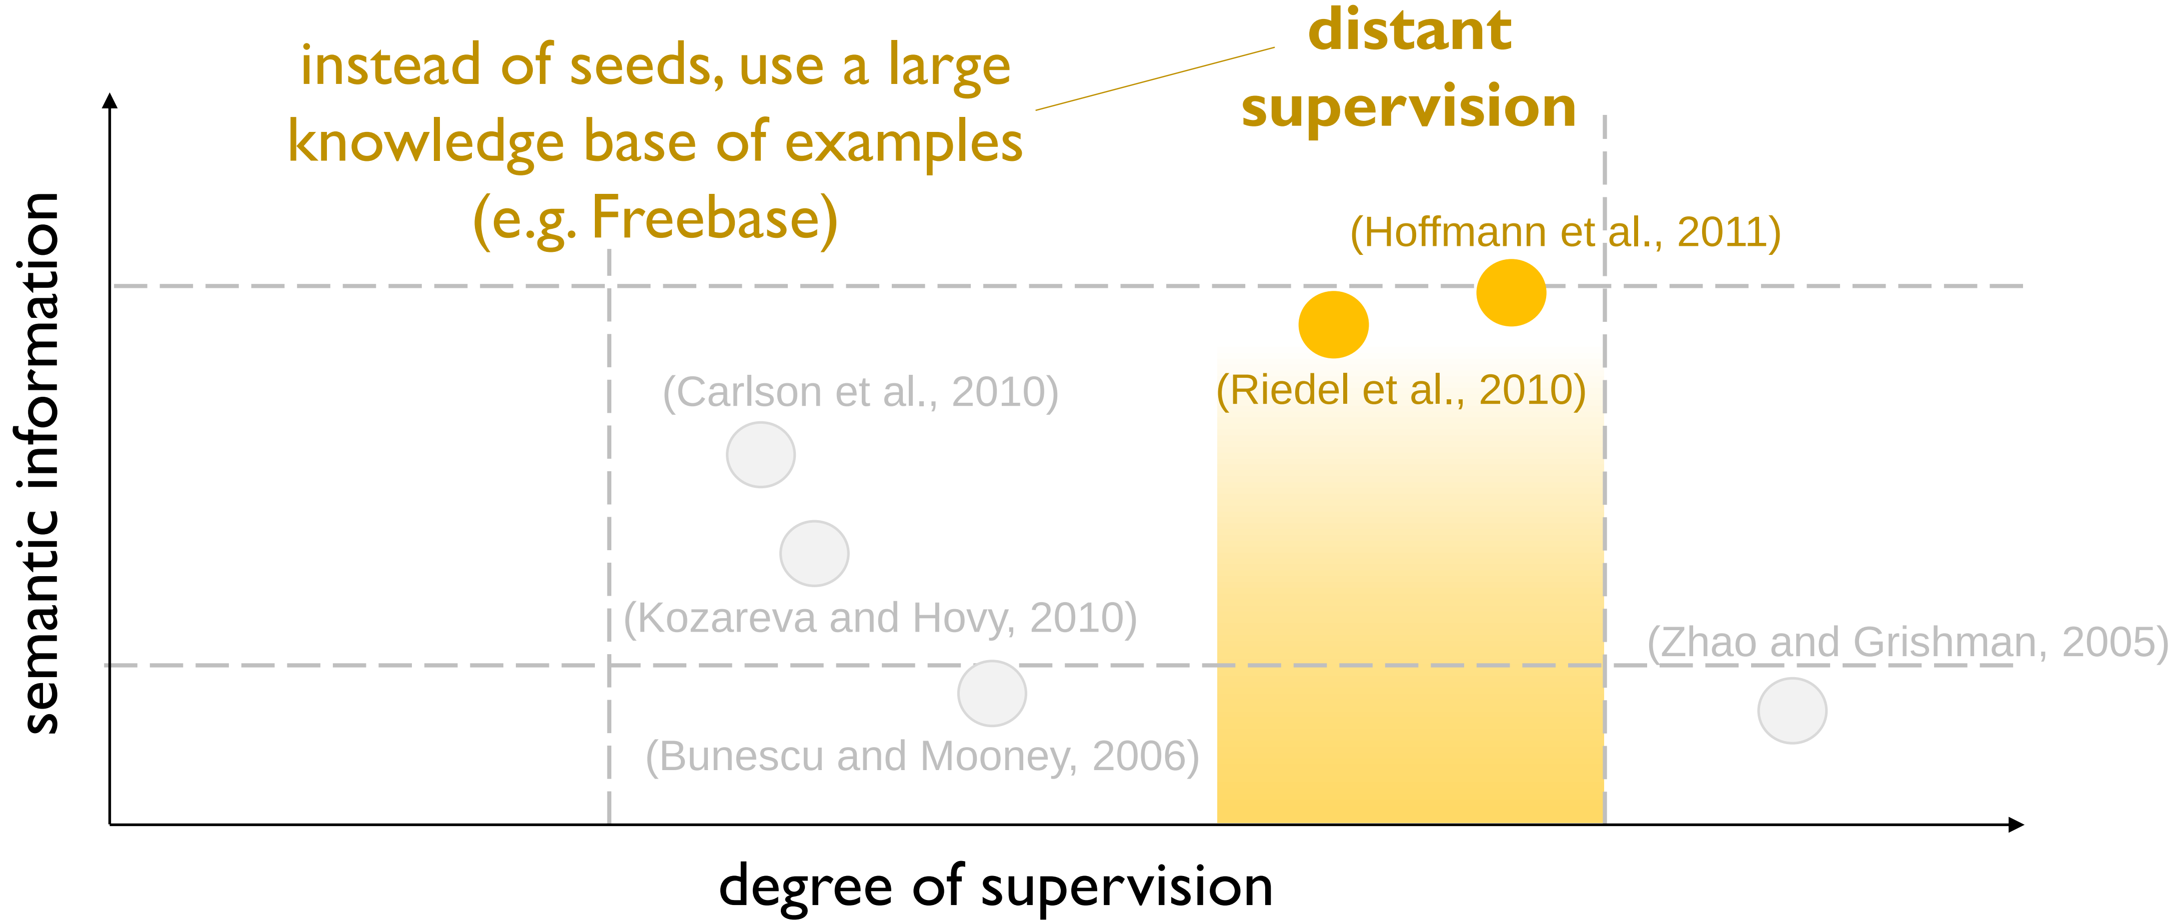

Information Extraction

How?

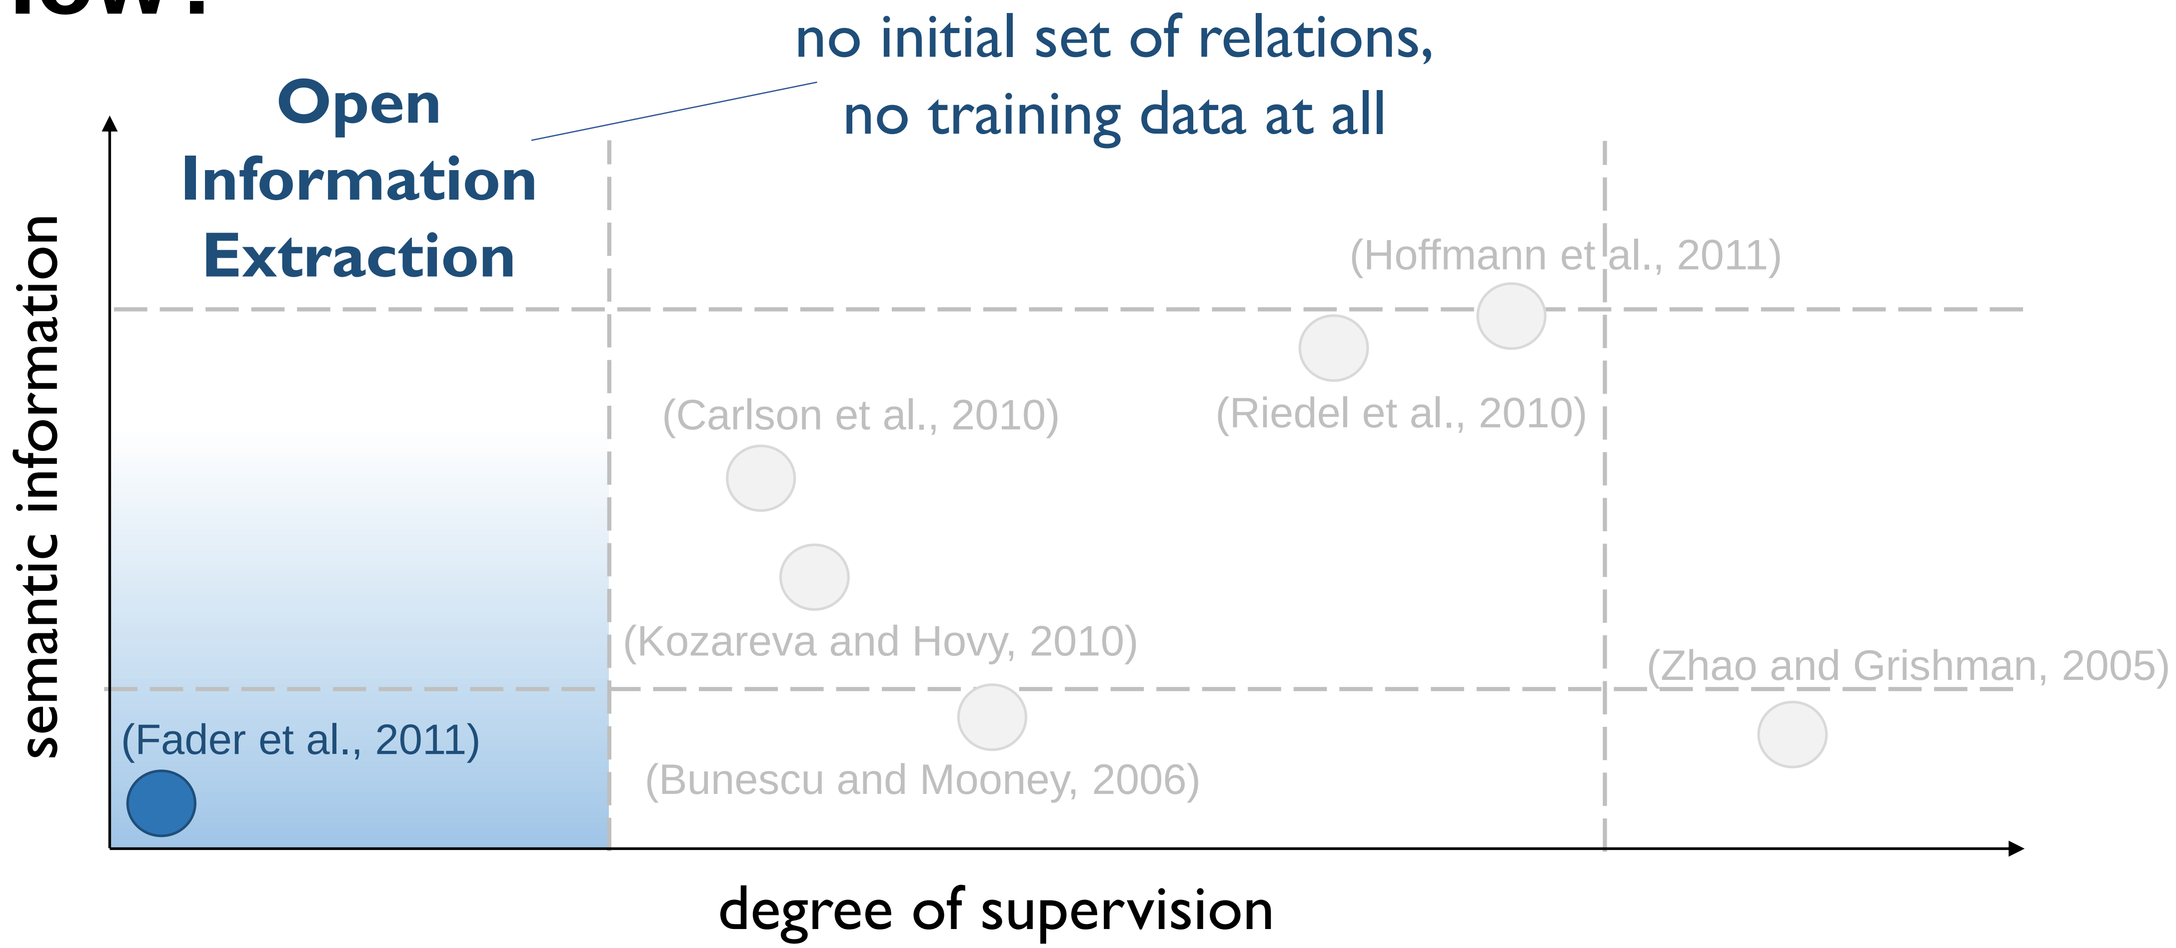

Information Extraction

How?

ReVerb (Fader et al., 2011)

For each sentence in the corpus:

- POS tagging and chunking
- Identify relation string as “well formed” sequence of words wrapping a **verb**
- Find the **nearest NPs** to the left and right
- Assign a **confidence** to the extraction

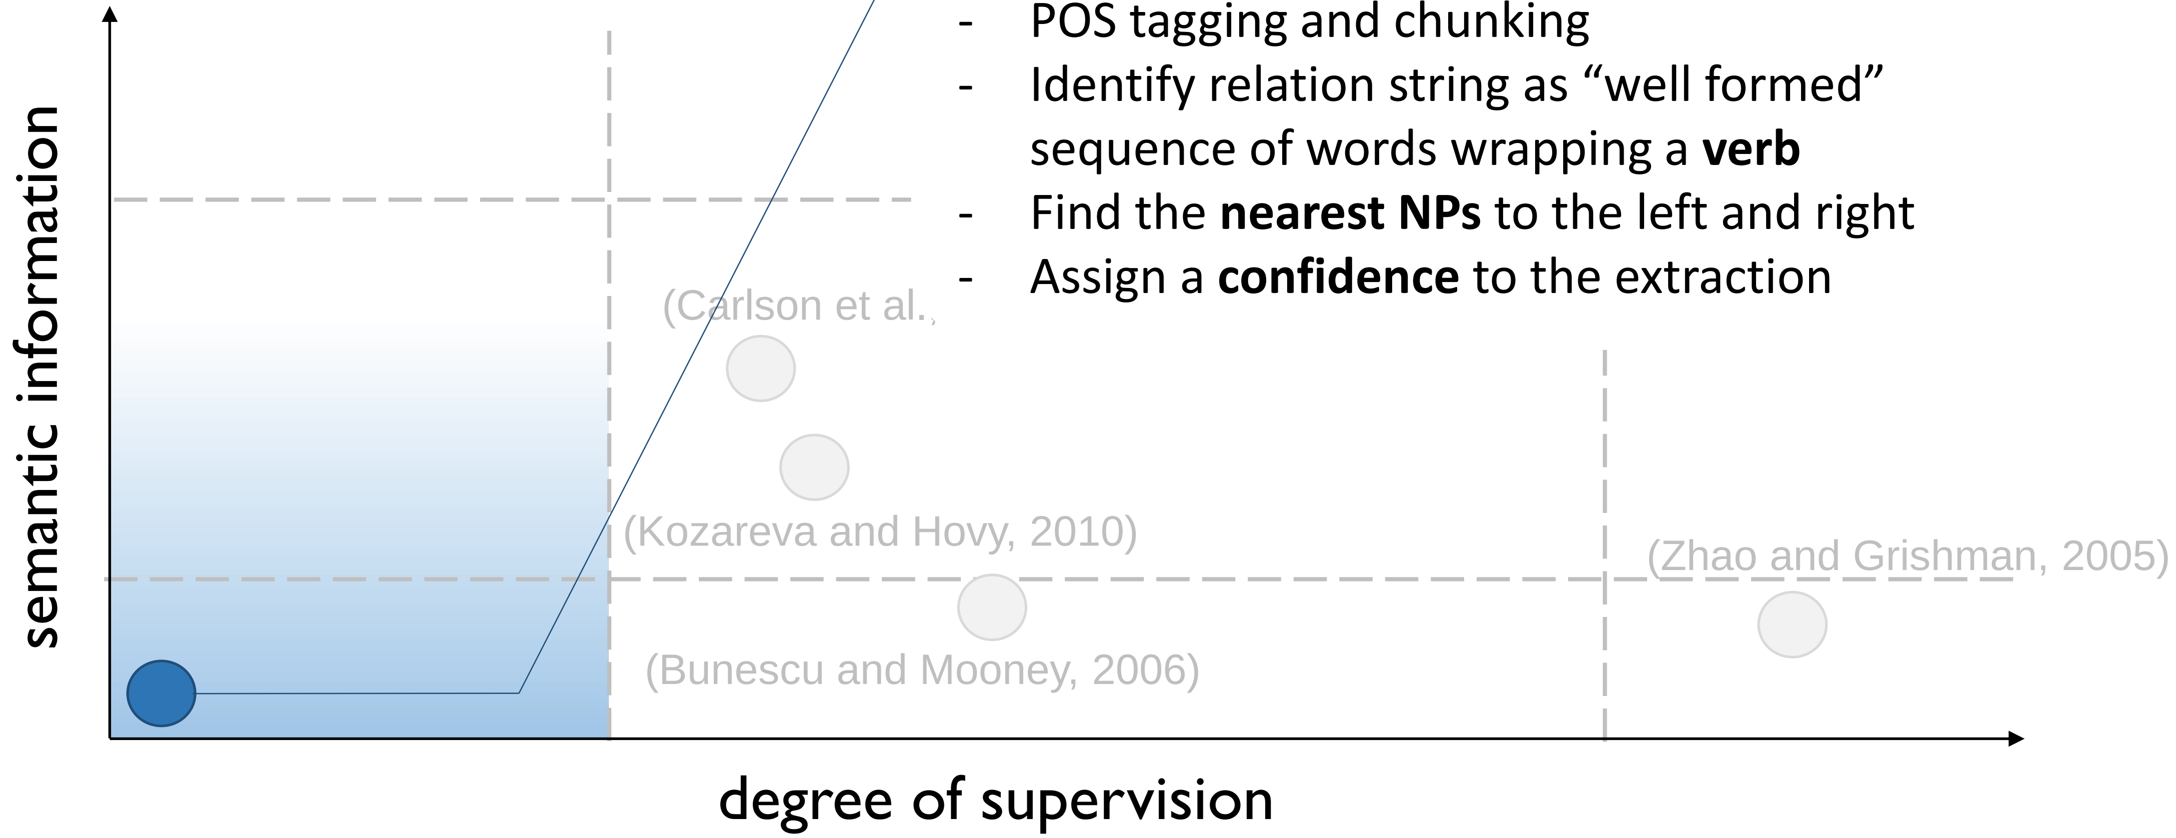

Information Extraction

How?

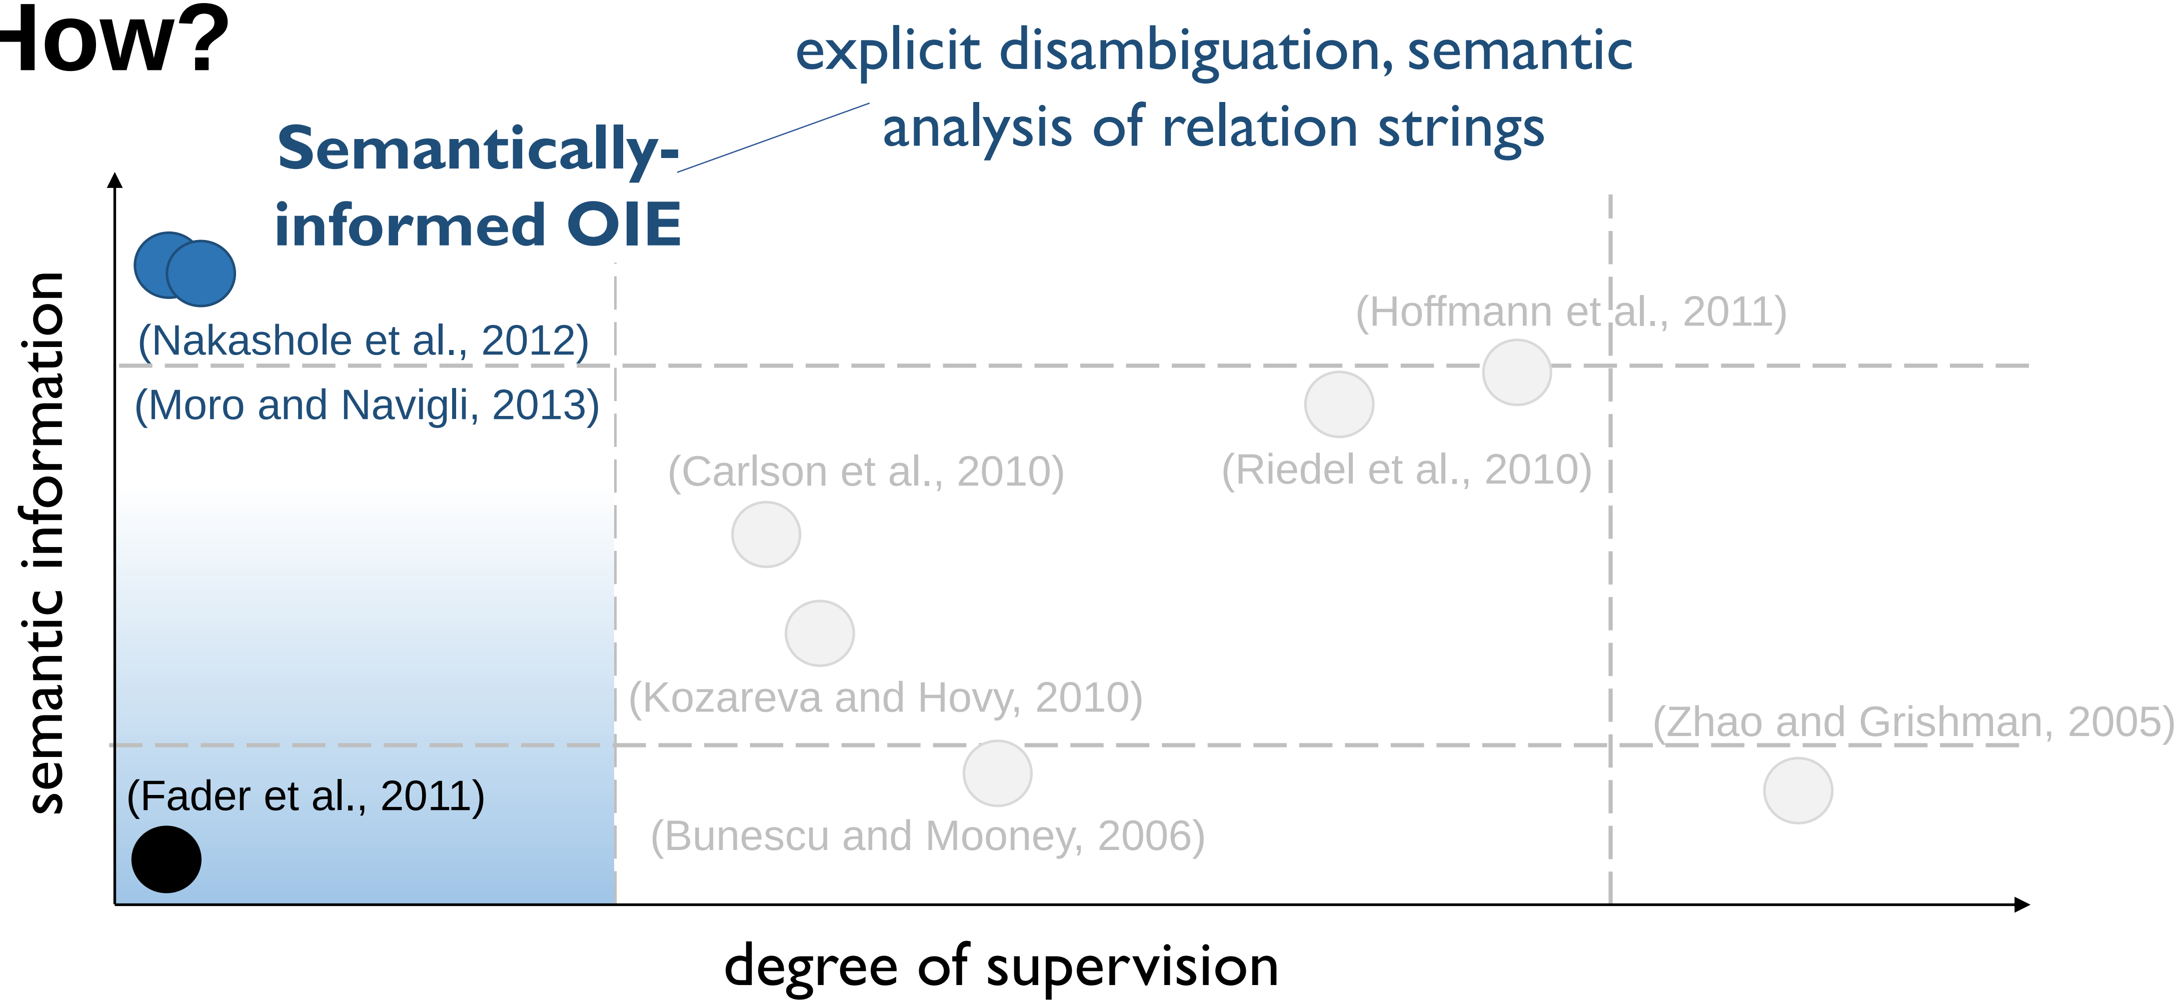

Information Extraction

How?

PATTY (Nakashole et al., 2012)

From patterns to **pattern synsets** (clusters of relation phrases that express the same relation)

Each pattern has **syntactic generalizations** and **semantic types** for its arguments:

<Person> 's [ADJ] voice * <Song>

Patterns are hierarchically organized in a **taxonomy**

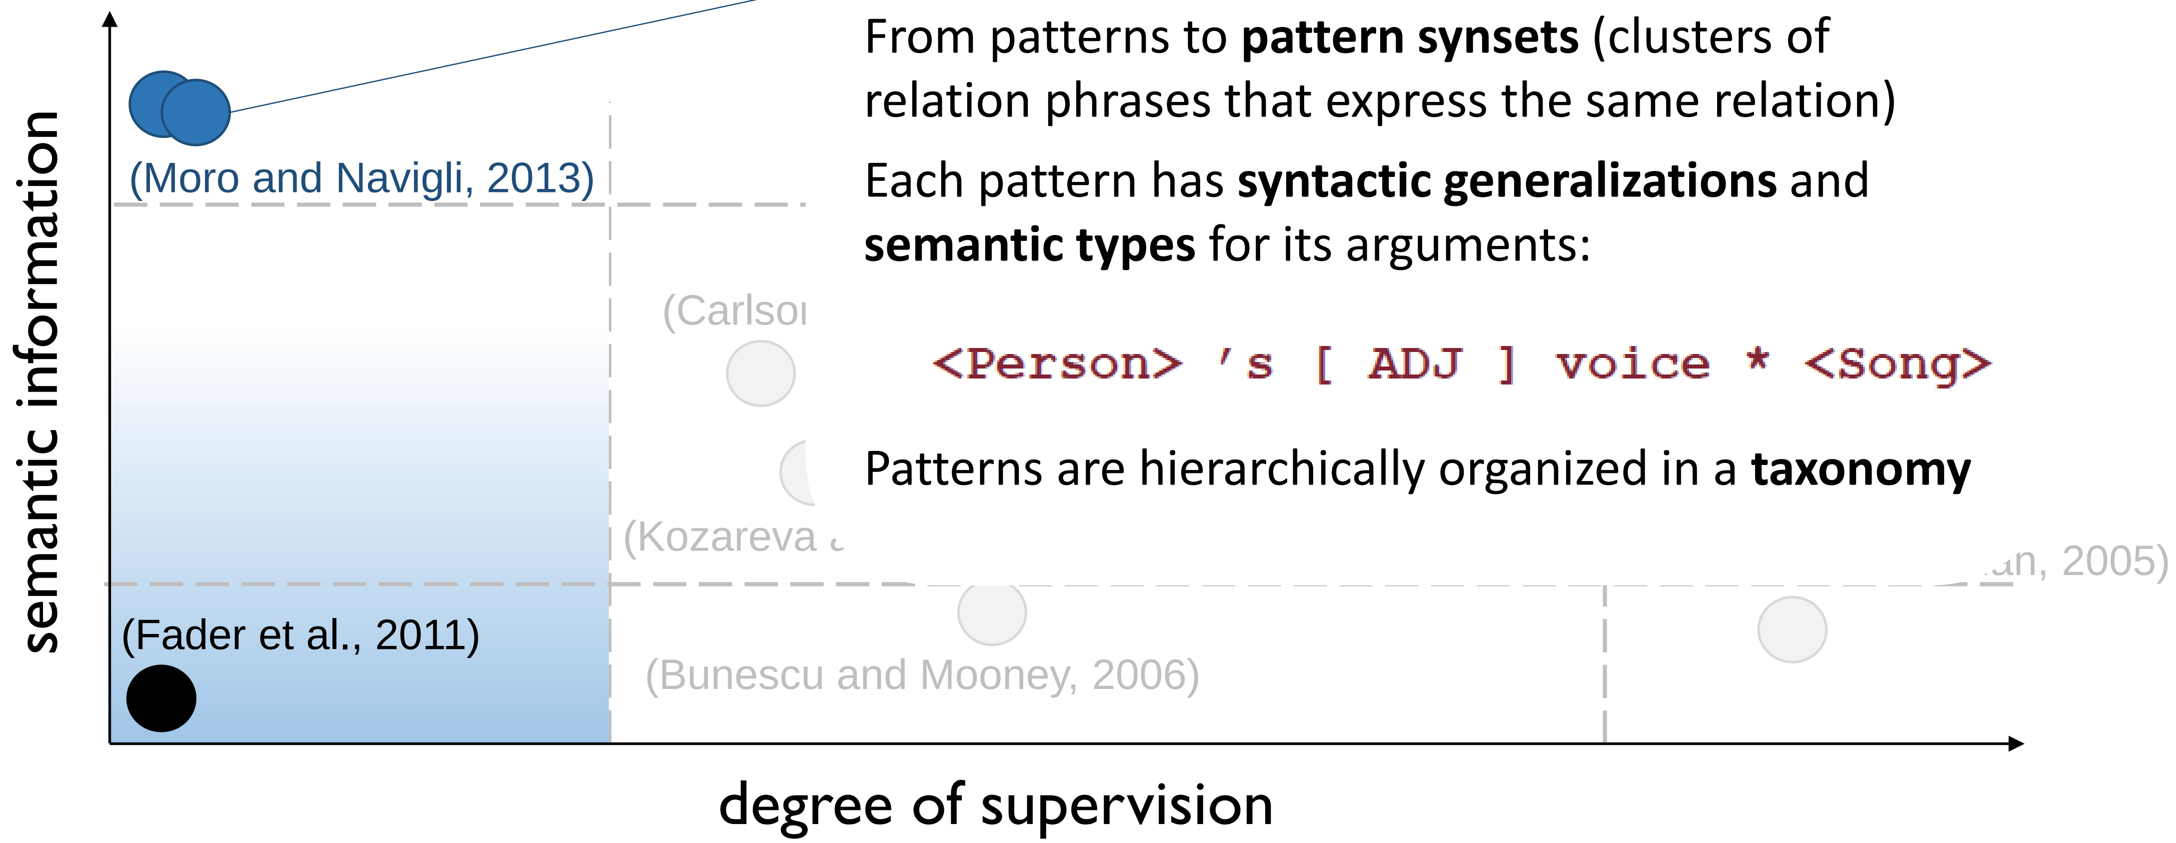

Information Extraction

How?

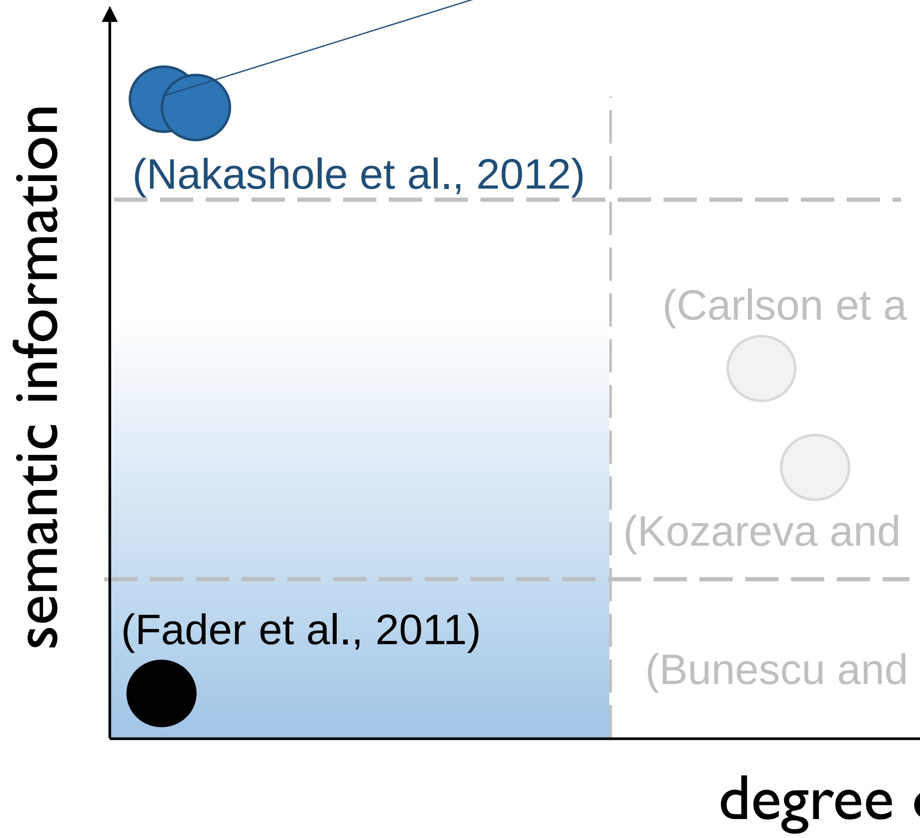

WiSeNet (Moro and Navigli, 2013)

Wikipedia-based **Se**mantic **Net**work: triples in the KB are determined by Wikipedia hyperlinks

Syntactically-grounded relational phrases with **Wikipedia categories** as semantic types

Relation synsets built using **soft clustering** techniques

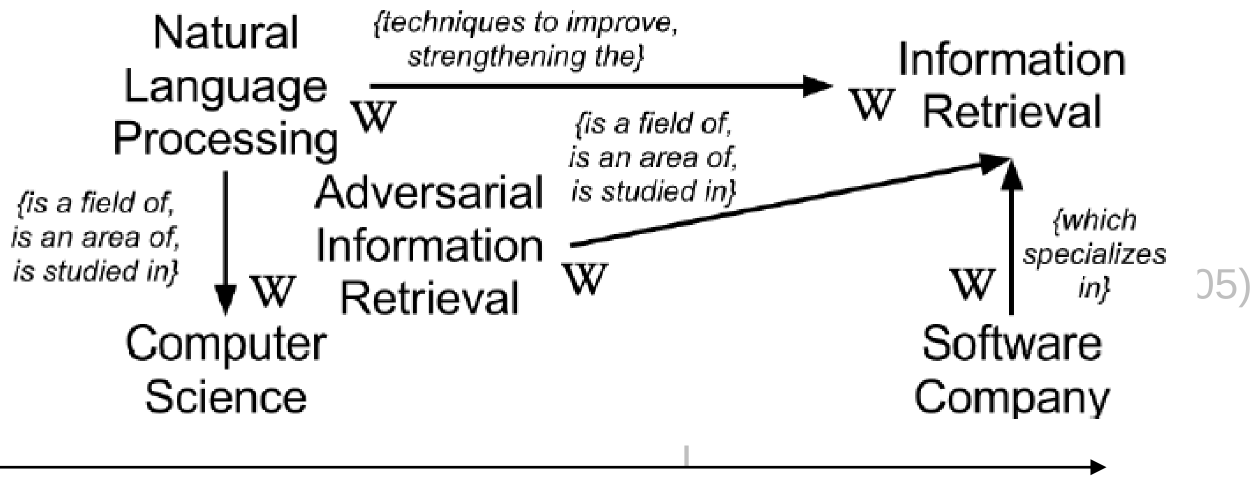

(Open) Information Extraction

OIE is great, but…

Sparsity: many relation phrases express the same relationship (e.g. synonyms, paraphrases)

Ambiguity: arguments (and relation phrases) are ambiguous!

Outline

IE and OIE: some background

DefIE: OIE from textual definitions

Claudio Delli Bovi, Luca Telesca and Roberto Navigli.
Large-Scale Information Extraction from Textual Definitions through Deep Syntactic and Semantic Analysis.
Transactions of the Association for Computational Linguistics (TACL), 2015.

SAPIENZA
UNIVERSITÀ DI ROMA

KBUnify: KB disambiguation and unification

Delli Bovi, Espinosa-Anke, Navigli: **EMNLP 2015**

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The idea:

instead of targeting massive and noisy corpora (like the web) and then trying to find a smart way to cope with the noise

target smaller but "denser" (and virtually noise-free) corpora of **definitional knowledge**.

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The idea:

instead of targeting massive and noisy corpora (like the web) and then trying to find a smart way to cope with the noise

target smaller but "denser" (and virtually noise-free) corpora of **definitional knowledge**.

Apply OIE techniques to extract as much information as possible!

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The tools:

- An underlying **inventory/knowledge base** (to which arguments and relation patterns will be connected)

- A **WSD/EL system** (to disambiguate concepts and entity mentions across the input text)

- A **syntactic parser** (to construct meaningful relation patterns and avoid sparsity)

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The tools:

- An underlying **inventory/knowledge base** (to which arguments and relation patterns will be connected)

- A **WSD/EL system** (to disambiguate concepts and entity mentions across the input text)

- A **syntactic parser** (to construct meaningful relation patterns and avoid sparsity)

http://babelnet.org

BabelNet

14 million entries

both **lexicographic** and **encyclopedic** knowledge

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The tools:

- An underlying **inventory/knowledge base** (to which arguments and relation patterns will be connected)

- A **WSD/EL system** (to disambiguate concepts and entity mentions across the input text)

- A **syntactic parser** (to construct meaningful relation patterns and avoid sparsity)

http://babelfy.org

Babelfy

unified graph-based approach to **EL** and **WSD**

unsupervised, based on **BabelNet**

DefIE: OIE from textual definitions

http://lcl.uniroma1.it/defie

The tools:

- An underlying **inventory/knowledge base** (to which arguments and relation patterns will be connected)

- A **WSD/EL system** (to disambiguate concepts and entity mentions across the input text)

- A **syntactic parser** (to construct meaningful relation patterns and avoid sparsity)

http://svn.ask.it.usyd.edu.au/trac/candc

C&C tools

log-linear parser and supertagger based on **CCG**

(theoretically) suited to **long-distance dependencies**

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

"Atom Heart Mother is the fifth album by English band Pink Floyd."

Textual definition d

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

Parsing

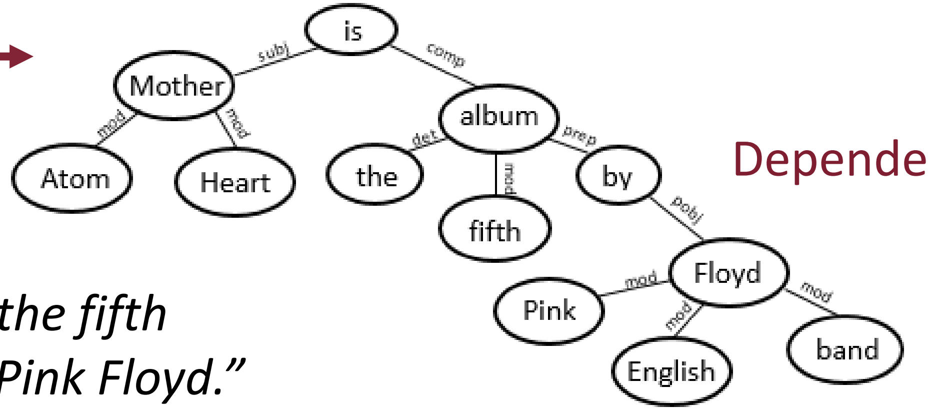

Dependency graph G_d

"Atom Heart Mother is the fifth album by English band Pink Floyd."

Disambiguation

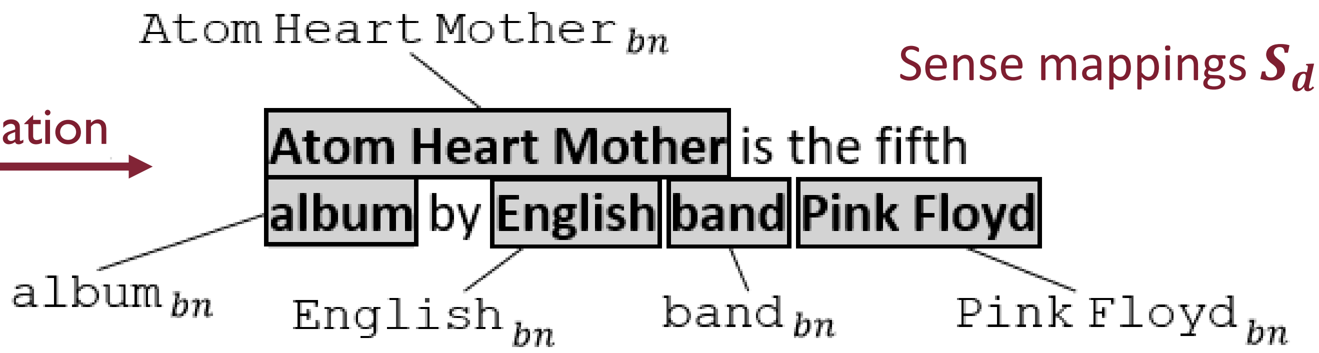

Sense mappings S_d

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

"Atom Heart Mother is the fifth album by English band Pink Floyd."

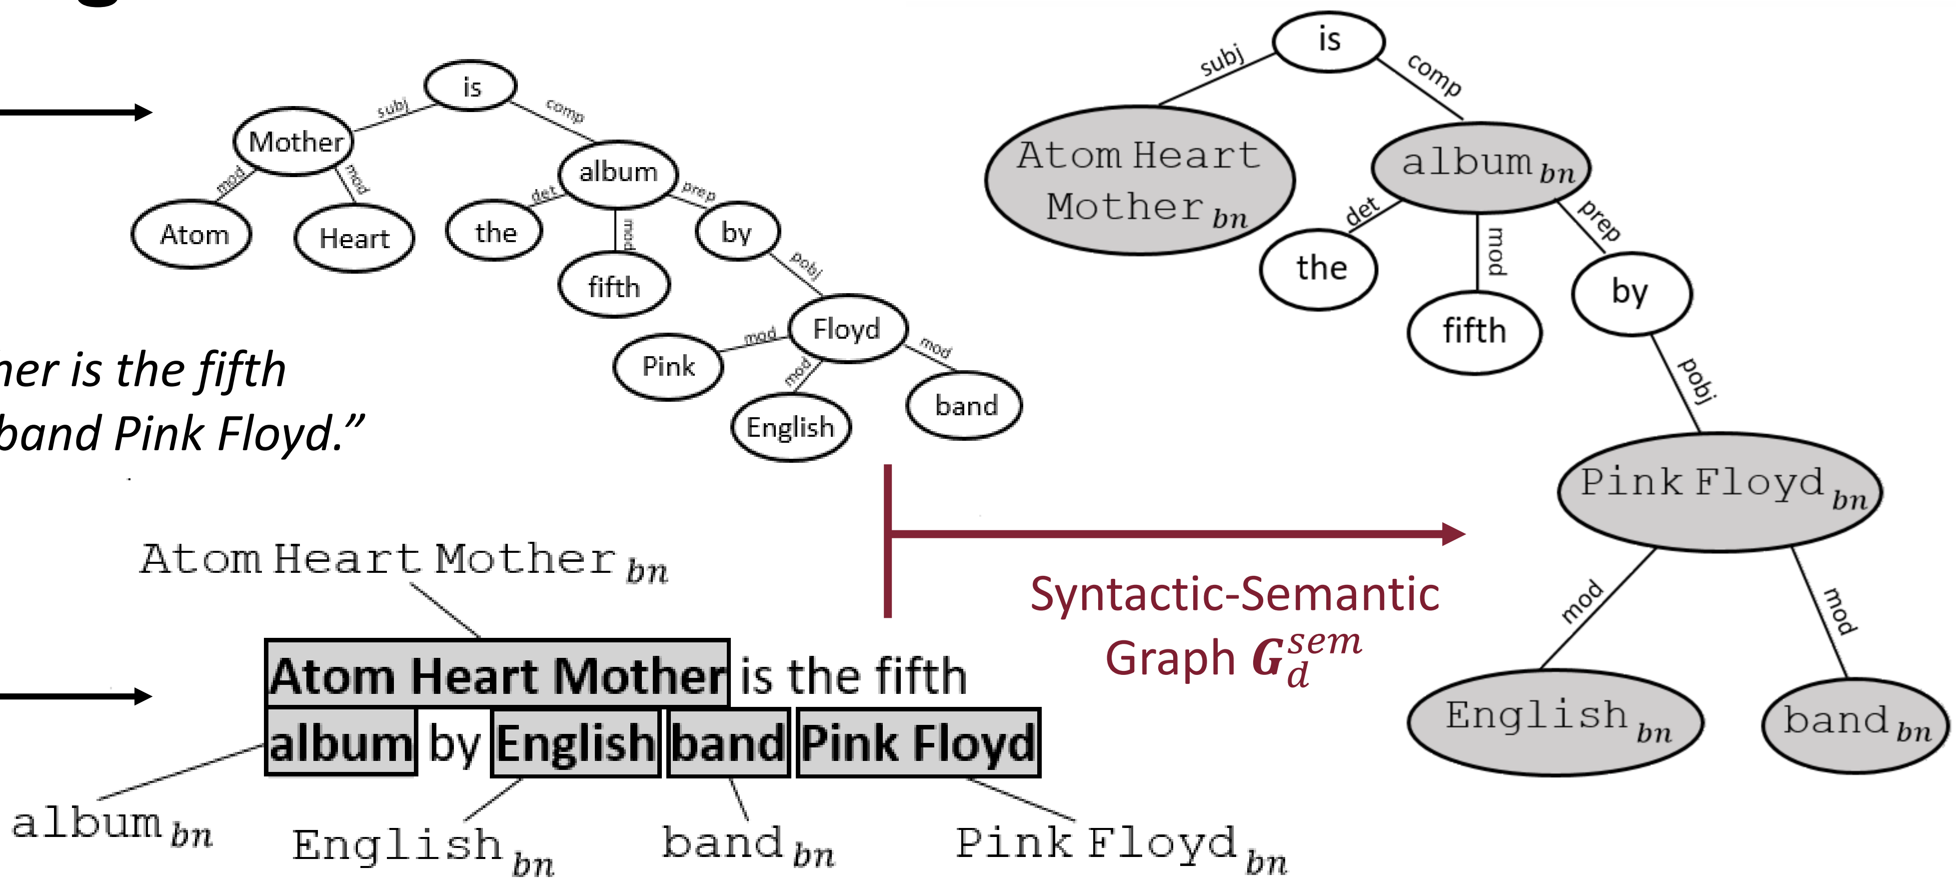

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

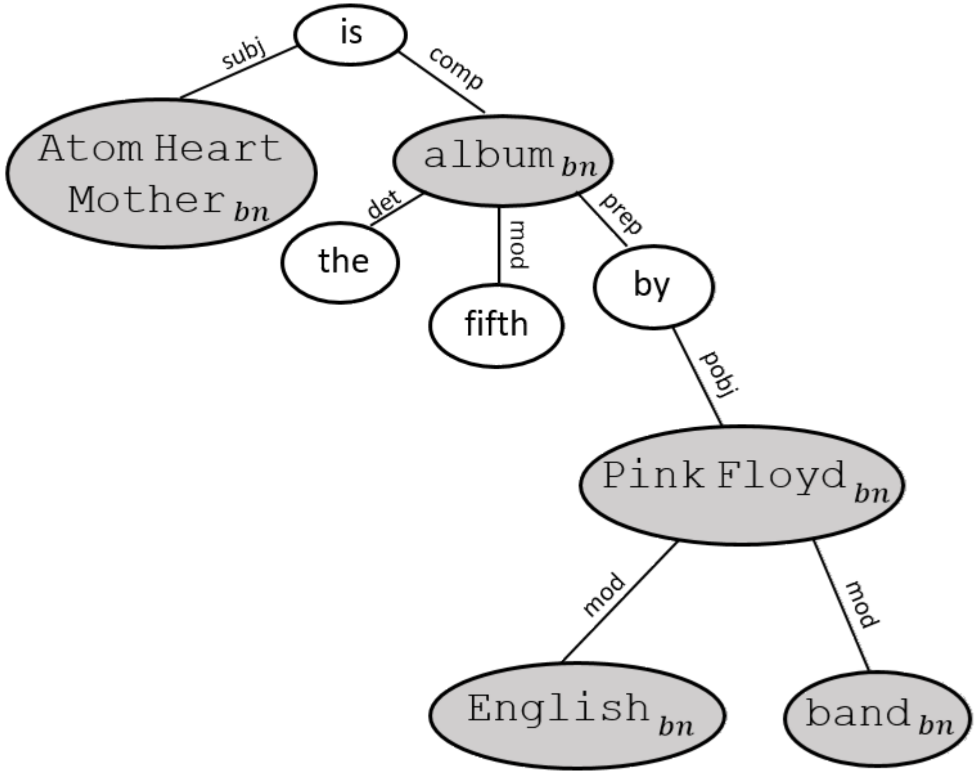

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

Extraction 1

$$\begin{cases} \mathrm{X} \rightarrow \texttt{is} \rightarrow \texttt{album}^1_{bn} \rightarrow \texttt{by} \rightarrow \mathrm{Y} \\ \mathrm{X} = \texttt{Atom Heart Mother}^1_{bn} \\ \mathrm{Y} = \texttt{Pink Floyd}^1_{bn} \end{cases}$$

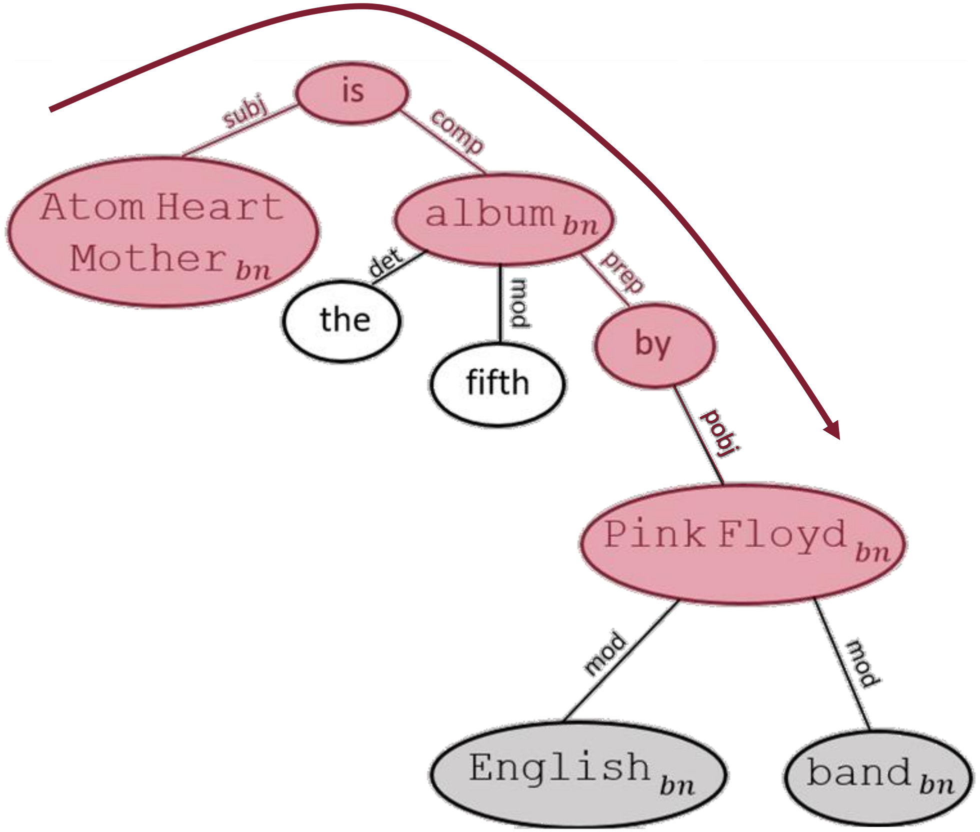

DefIE: How it works

http://lcl.uniroma1.it/defie

1. Extracting relation instances

Extraction 2

$$X \rightarrow \texttt{is} \rightarrow Y$$
$$X = \texttt{Atom Heart Mother}^1_{bn}$$
$$Y = \texttt{album}^1_{bn}$$

Extraction 1

$$X \rightarrow \texttt{is} \rightarrow \texttt{album}^1_{bn} \rightarrow \texttt{by} \rightarrow Y$$
$$X = \texttt{Atom Heart Mother}^1_{bn}$$
$$Y = \texttt{Pink Floyd}^1_{bn}$$

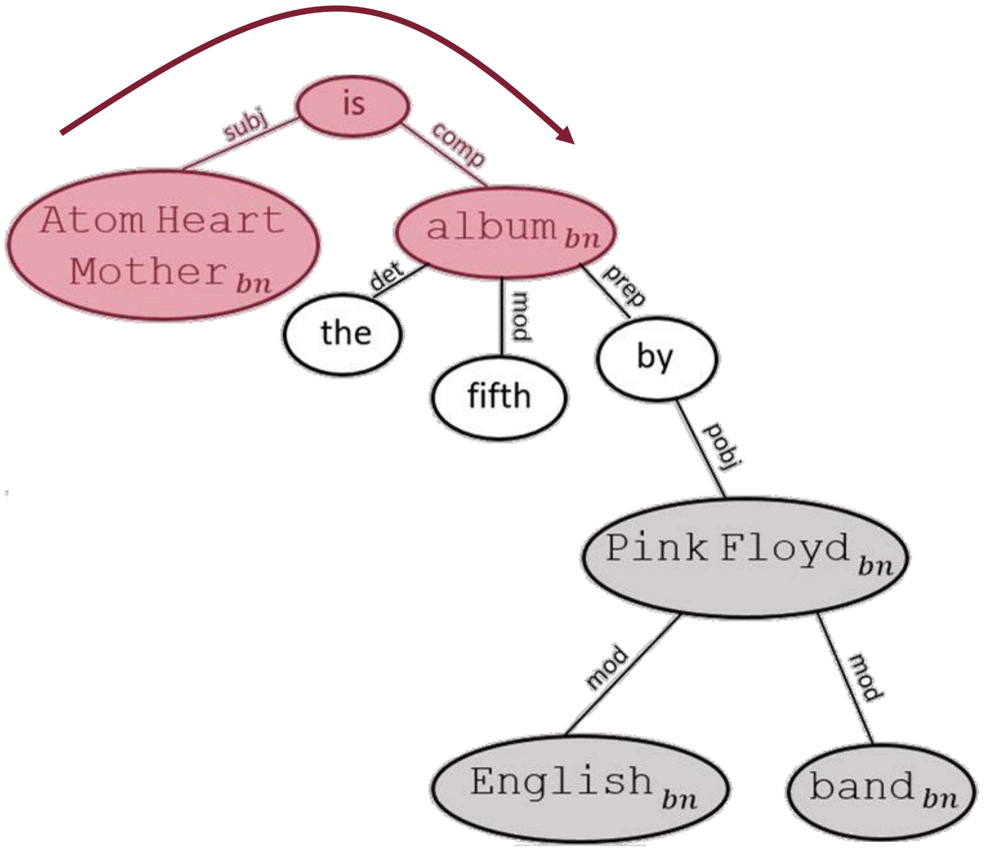

DefIE: How it works

http://lcl.uniroma1.it/defie

2. Relation typing and scoring

DefIE: How it works

http://lcl.uniroma1.it/defie

2. Relation typing and scoring

For each relation R:

Substitute each domain and range argument with its **hypernym** $\boldsymbol{h}$ (using the BabelNet taxonomy) and generate a **probability distribution over semantic types** for the two sets

DefIE: How it works

http://lcl.uniroma1.it/defie

2. Relation typing and scoring

For each relation R:

Substitute each domain and range argument with its **hypernym *h*** (using the BabelNet taxonomy) and generate a **probability distribution over semantic types** for the two sets

Compute the **entropy** of R as $H_R = -\sum_{i=1}^{n} p(h_i)\, log_2\, p(h_i)$

DefIE: How it works

http://lcl.uniroma1.it/defie

2. Relation typing and scoring

For each relation R:

Compute the **score** of R as

$$score(R) = \frac{|S_R|}{(H_R + 1)\, length(r)}$$

Total number of extracted instances for R

Domain and range entropy of R

Length of the relation pattern of R

DefIE: How it works

http://lcl.uniroma1.it/defie

2. Relation typing and scoring

Pattern	Score	Entropy
X *directed by* Y	4 025.80	1.74
X *known for* Y	2 590.70	3.65
X *is* $\texttt{election district}^{1}_{bn}$ *of* Y	110.49	0.83
X *is* $\texttt{composer}^{1}_{bn}$ *from* Y	39.92	2.08
X *is* $\texttt{street}^{1}_{bn}$ *named after* Y	1.91	2.24
X *is* $\texttt{village}^{2}_{bn}$ *founded in 1912 in* Y	0.91	0.18

DefIE: How it works

http://lcl.uniroma1.it/defie

3. Relation taxonomization

DefIE: How it works

http://lcl.uniroma1.it/defie

3. Relation taxonomization

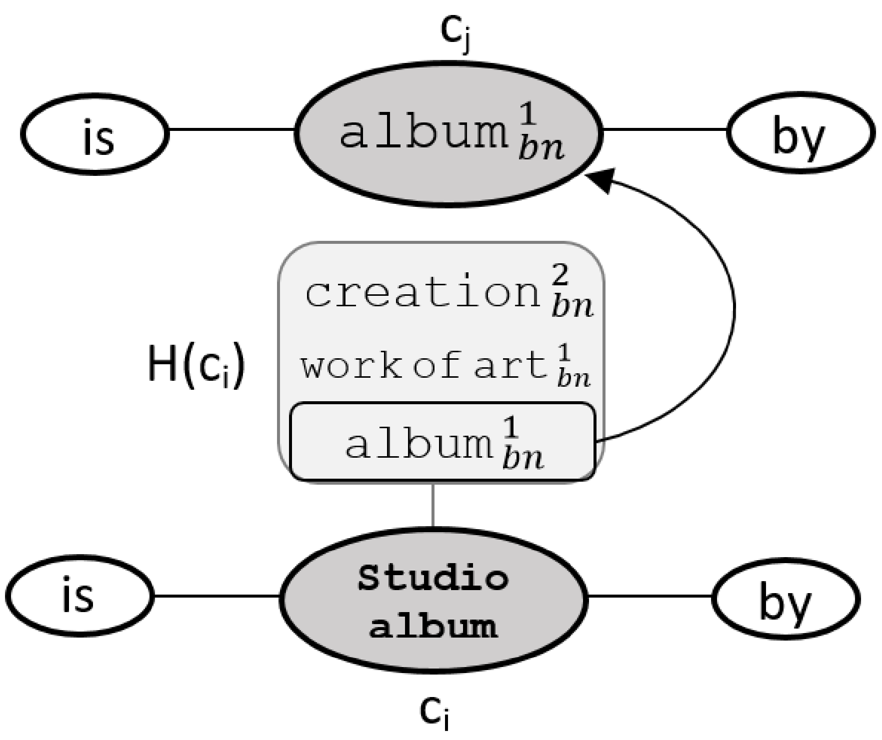

Hypernym Generalization

DefIE: How it works

http://lcl.uniroma1.it/defie

3. Relation taxonomization

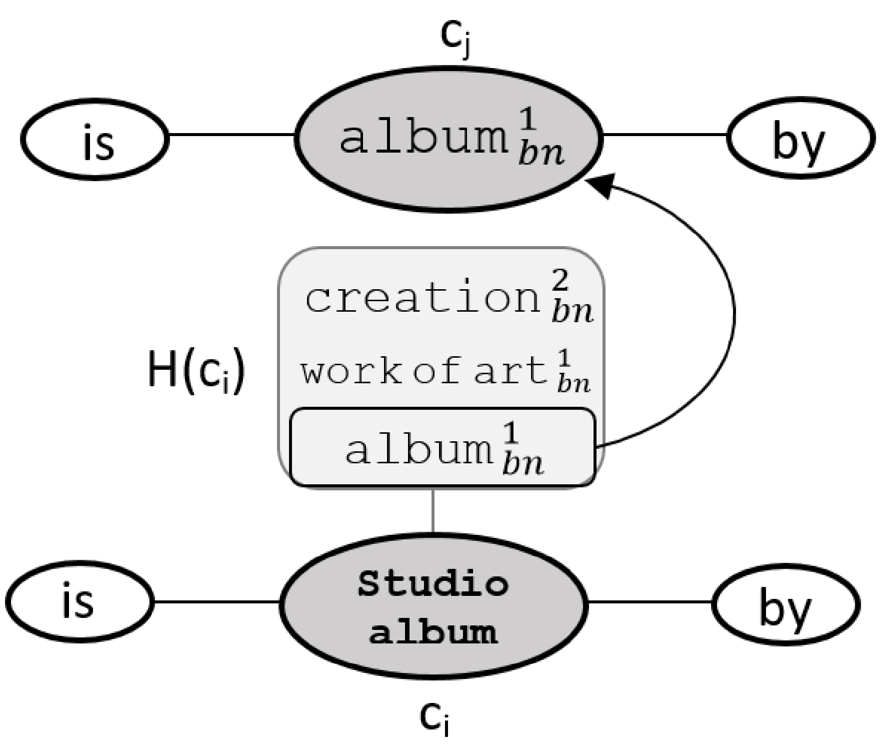

Hypernym Generalization

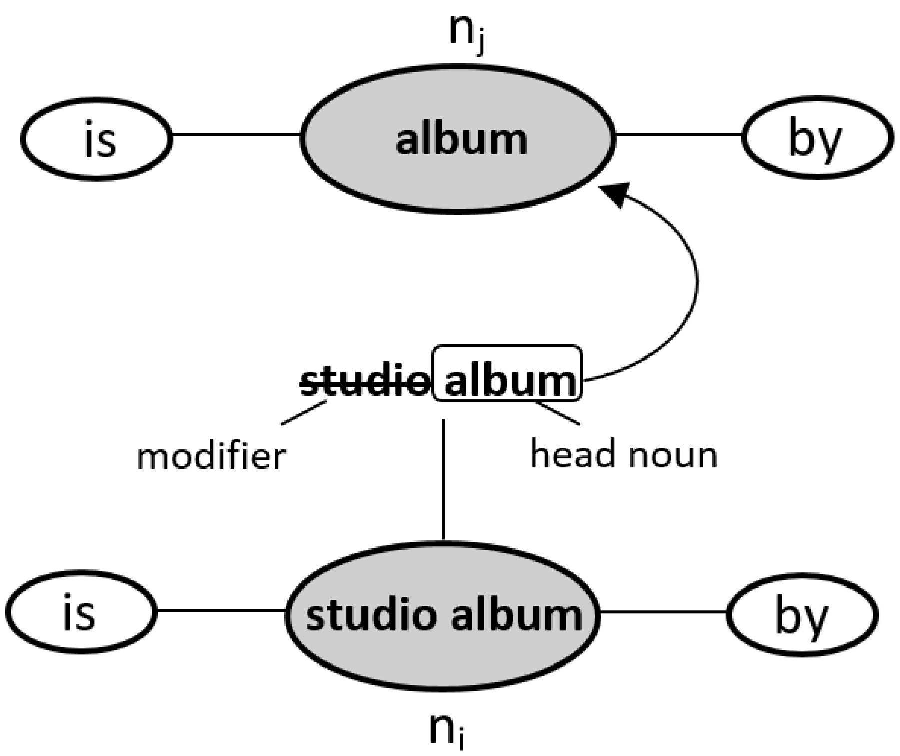

Substring Generalization

DefIE: Results

http://lcl.uniroma1.it/defie

Dataset:
whole set of English textual definitions in BabelNet 2.5

4 357 327 items from **5** different sources (Wikipedia, WordNet, Wikidata, Wiktionary, OmegaWiki)

DefIE: Results

http://lcl.uniroma1.it/defie

	DefIE	NELL	PATTY	ReVerb	WiSeNet
# Relations	255 881	298	**1 631 531**	664 746	245 935
Avg. extractions	81.68	**7 013.03**	9.68	22.16	9.24
# Extractions	**20 352 903**	2 089 883	15 802 946	14 728 268	2 271 807
# Entities	2 398 982	1 996 021	1 087 907	**3 327 425**	1 636 307
# Edges in the taxonomy	**44 412**	-	20 339	-	-

DefIE: Results

http://lcl.uniroma1.it/defie

Other evaluations:

- **Precision** and **coverage** of relations
- **Novelty** of information
- Quality of relation **taxonomization**
- Quality of **entity linking/disambiguation**
- **Impact** of definition sources

…

DefIE: Results

http://lcl.uniroma1.it/defie

Data and output soon available for download on the website!

Other evaluations:

- **Precision** and **coverage** of relations
- **Novelty** of information
- Quality of relation **taxonomization**
- Quality of **entity linking/disambiguation**
- **Impact** of definition sources

…

Outline

IE and OIE: some background

DefIE: OIE from textual definitions

Delli Bovi, Telesca, Navigli: **TACL** (to appear)

KBUnify: KB disambiguation and unification

Claudio Delli Bovi, Luis Espinosa-Anke and Roberto Navigli.
Knowledge Base Unification via Sense Embeddings and Disambiguation.
Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 726–736, Lisbon, Portugal, 17-21 September 2015.

SAPIENZA
UNIVERSITÀ DI ROMA

UNIVERSITAT
POMPEU FABRA

KB-Unify: Knowledge base unification via sense embeddings and disambiguation

http://lcl.uniroma1.it/kb-unify

The idea:

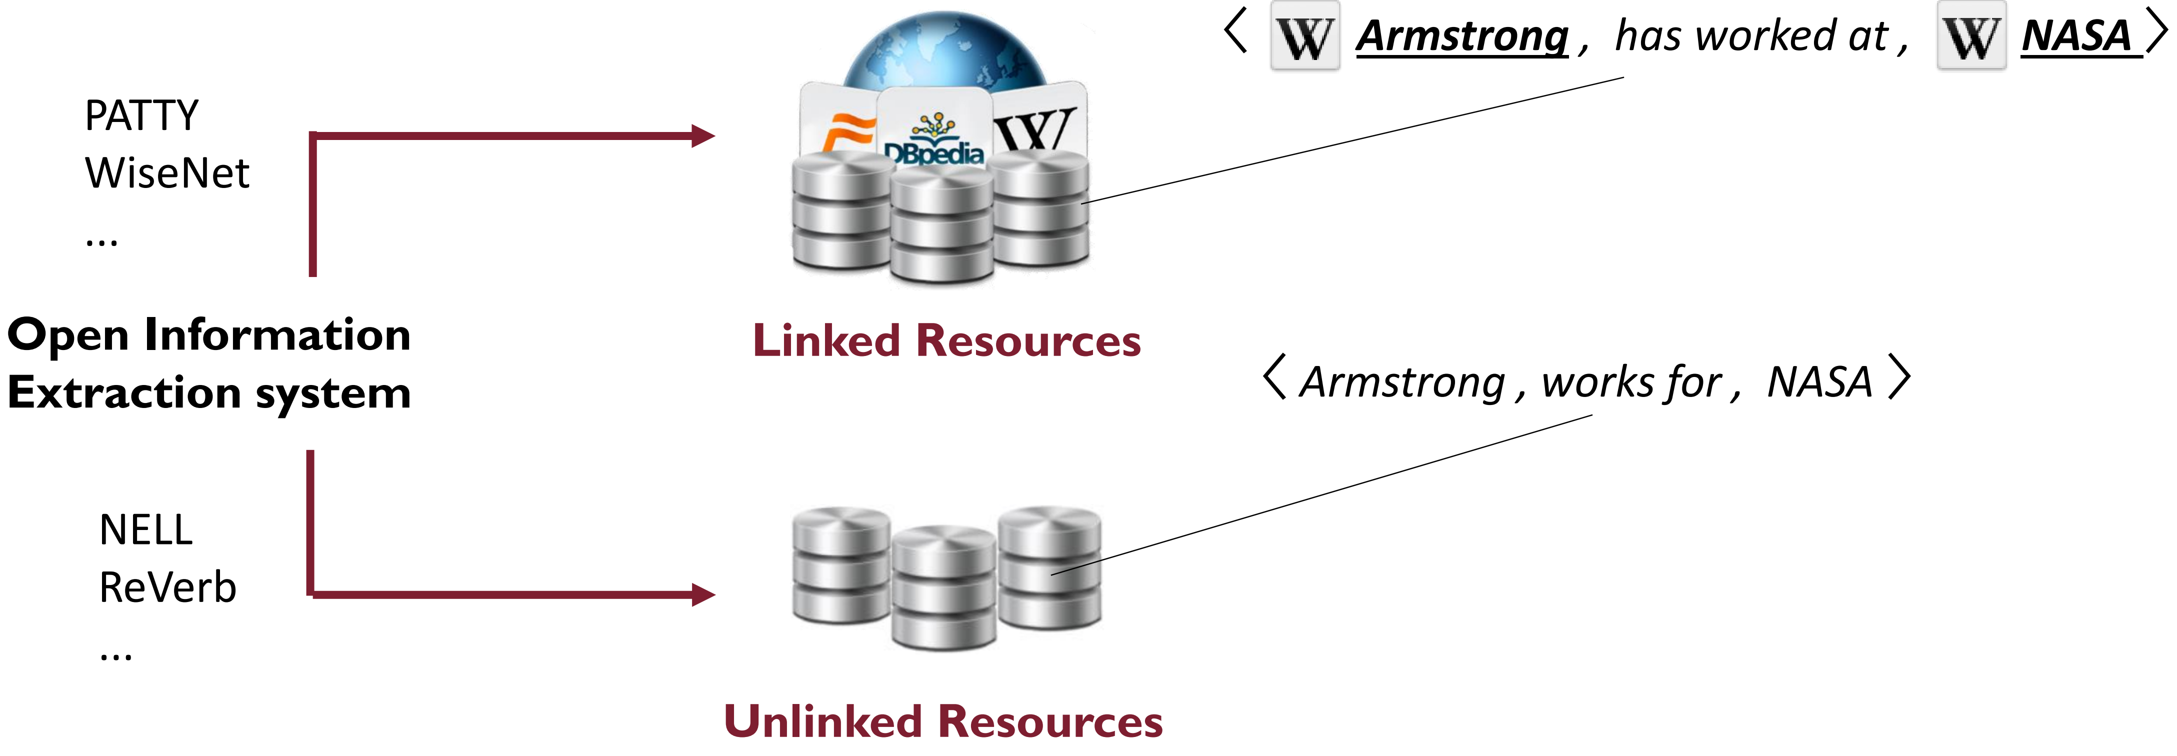

KB-Unify: Knowledge base unification via sense embeddings and disambiguation

http://lcl.uniroma1.it/kb-unify

The idea:

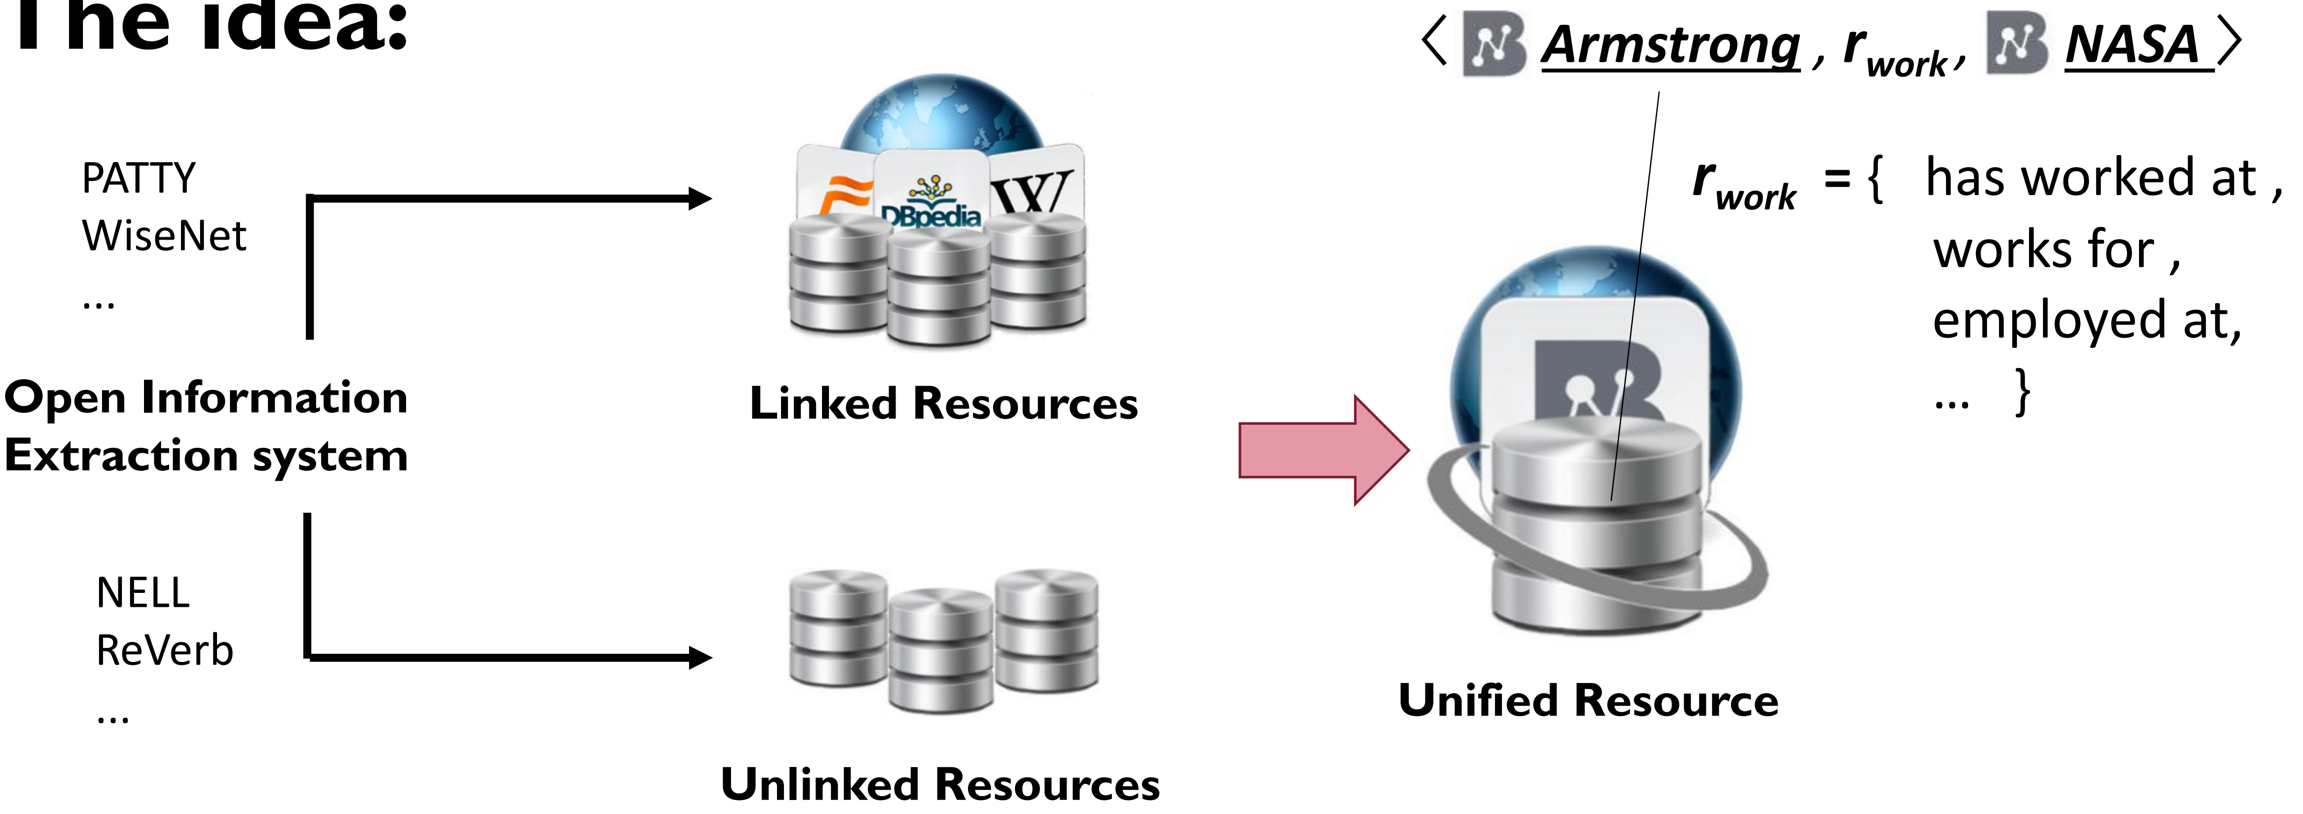

KB-Unify: Knowledge base unification via sense embeddings and disambiguation

http://lcl.uniroma1.it/kb-unify

The tools:

- A **WSD/EL system** (to disambiguate unlinked resources)

- A unified **sense inventory S** (to make the various resources "speak to each other")

- A unified **vector space $\mathbf{V_S}$** (to associate a vector with each item of **S**)

KB-Unify: Knowledge base unification via sense embeddings and disambiguation

http://lcl.uniroma1.it/kb-unify

The tools:

- A **WSD/EL system** (to disambiguate unlinked resources)

- A unified **sense inventory S** (to make the various resources "speak to each other")

- A unified **vector space** $\mathbf{V_S}$ (to associate a vector with each item of **S**)

Babelfy

Babelnet

KB-Unify: Knowledge base unification via sense embeddings and disambiguation

http://lcl.uniroma1.it/kb-unify

The tools:

- A **WSD/EL system** (to disambiguate unlinked resources)

- A unified **sense inventory S** (to make the various resources "speak to each other")

- A unified **vector space** $\mathbf{V}_S$ (to associate a vector with each item of **S**)

SensEmbed
(Iacobacci et al., 2015)

Sense-based embedding model

Popular word2vec architecture (**skip-gram**) trained on a **sense-annotated corpus**

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

A bird's-eye view

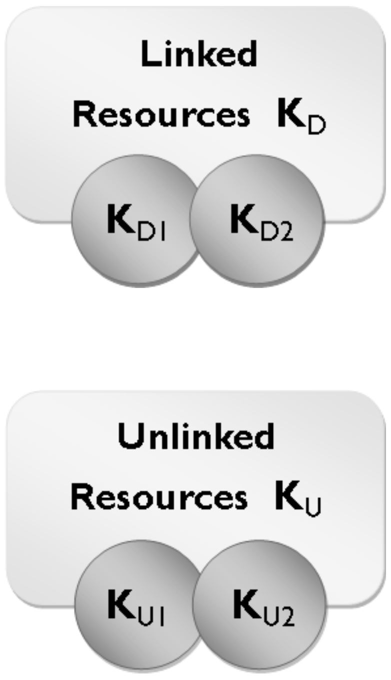

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

A bird's-eye view

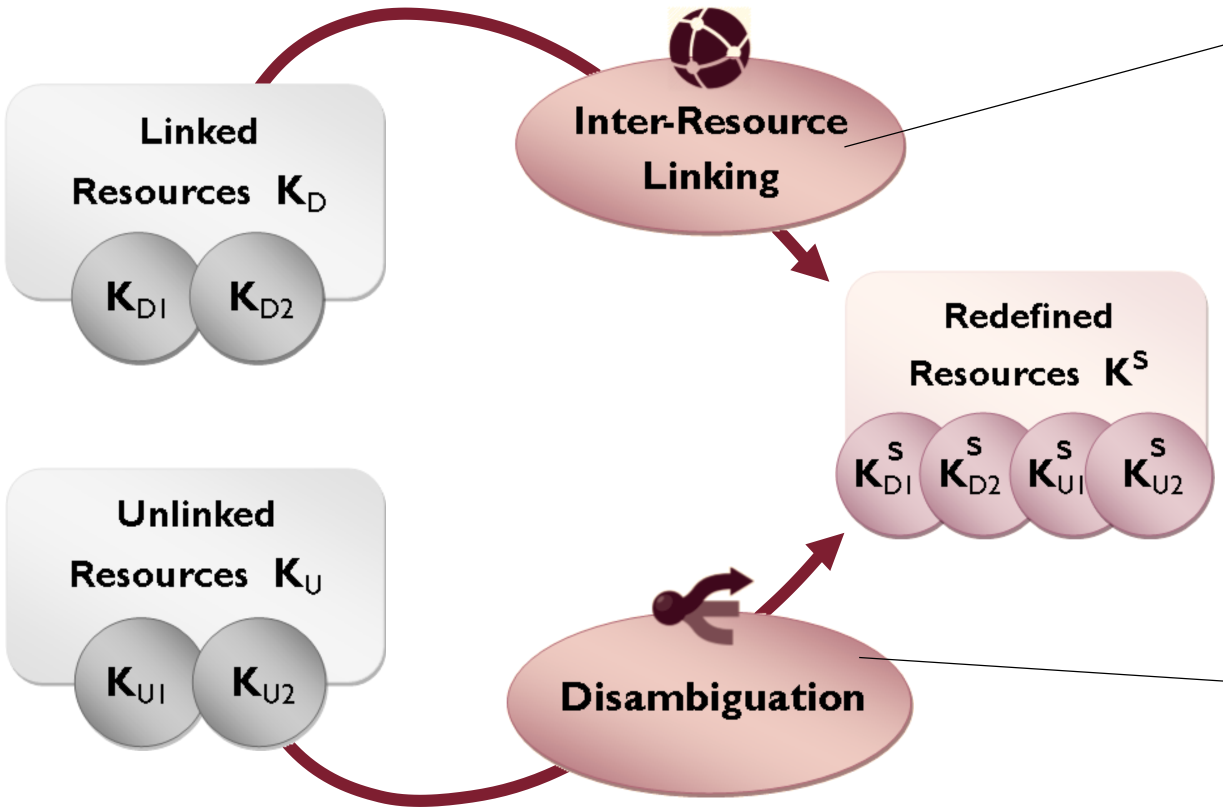

use **BabelNet mappings** to redefine each linked resource

disambiguate each unlinked resource using Babelnet as sense inventory (more on this later!)

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Disambiguation

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Disambiguation

Two basic intuitions:

1. Among all triples in target knowledge base, some of them (even if ambiguous) will be **easier to disambiguate**;

   e.g. ⟨ **Armstrong , works for , NASA** ⟩

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Disambiguation

Two basic intuitions:

1. Among all triples in target knowledge base, some of them (even if ambiguous) will be **easier to disambiguate**;

   e.g. ⟨ **Armstrong , works for , NASA** ⟩

2. In general, the disambiguation strategy should vary according to the **degree of specificity** of each relation.

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Disambiguation

Group the set of unlinked tripes by relation

For each relation r:

- Extract and disambiguate a subset of high-confidence **seed argument pairs** for r;
- Estimate the **specificity** of r by looking at the distribution of its disambiguated seeds in the vector space $\mathbf{V_S}$;
- Disambiguate the remaining argument pairs of r with Babelfy either **triple-by-triple** (if r is general) or **all at once** (if r is specific).

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Identifying seed argument pairs

⟨ **Armstrong ,**

works for ,

NASA ⟩

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Identifying seed argument pairs

$$\mathbf{v}_d = \{ \; v_d^1 \quad v_d^2 \quad v_d^3 \; \}$$

⟨ $\boxed{\text{Armstrong}}$,

works for ,

$\boxed{\text{NASA}}$ ⟩

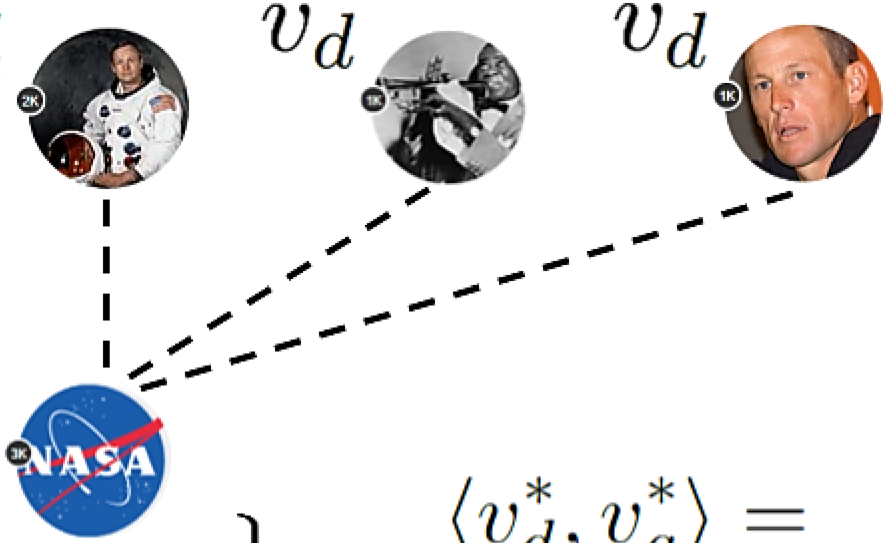

$$\mathbf{v}_g = \{ \; v_g^1 \; \}$$

$$\langle v_d^*, v_g^* \rangle = \underset{v_d \in \mathbf{v}_d,\, v_g \in \mathbf{v}_g}{\operatorname{argmax}} \frac{v_d \cdot v_g}{\|v_d\| \, \|v_g\|}$$

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Identifying seed argument pairs

$$\mathbf{v}_d = \{\ v_d^1 \quad v_d^2 \quad v_d^3\ \}$$

⟨ Armstrong ,

works for ,

NASA ⟩

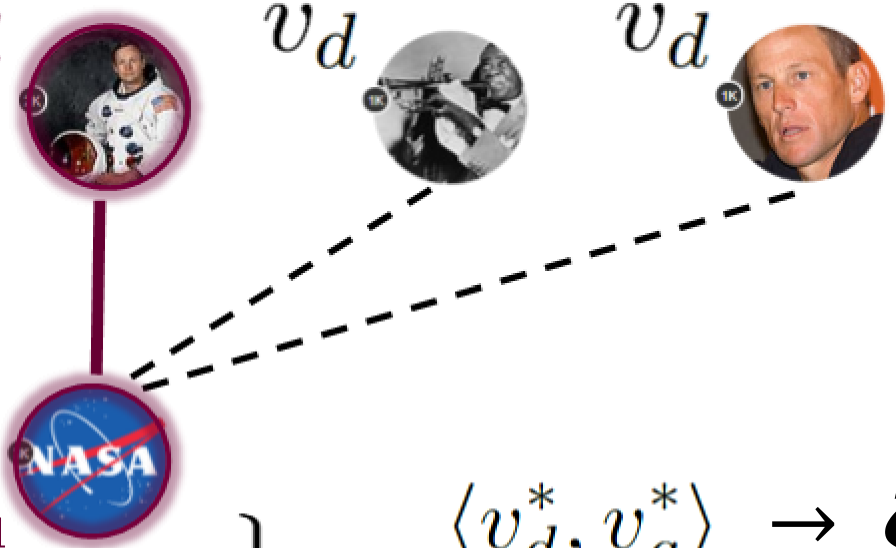

$$\mathbf{v}_g = \{\ v_g^1\ \}$$

$$\langle v_d^*, v_g^* \rangle \rightarrow \zeta_{dis}$$

Seed Disambiguation Confidence

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Ranking relations by specificity

$$\mu_k = \frac{1}{|\mathbf{v}_k|} \sum_{v \in \mathbf{v}_k} \frac{v}{||v||} \quad , \quad k \in \{D, G\}$$

Domain/Range Centroids

$$\sigma_k^2 = \frac{1}{|\mathbf{v}_k|} \sum_{v \in \mathbf{v}_k} (1 - cos\,(v, \mu_k))^2$$

Domain/Range Variances

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Ranking relations by specificity

$$\mu_k = \frac{1}{|\mathbf{v}_k|} \sum_{v \in \mathbf{v}_k} \frac{v}{\|v\|} \quad , \quad k \in \{D, G\}$$

Domain/Range Centroids

$$\sigma_k^2 = \frac{1}{|\mathbf{v}_k|} \sum_{v \in \mathbf{v}_k} (1 - cos\,(v, \mu_k))^2$$

Domain/Range Variances

$$Gen(r) = \frac{\sigma_D^2 + \sigma_G^2}{2}$$

Specificity threshold: $\boldsymbol{\delta_{spec}}$

High $Gen(r)$ $(> \boldsymbol{\delta_{spec}})$

$\mathbf{v}_D$ 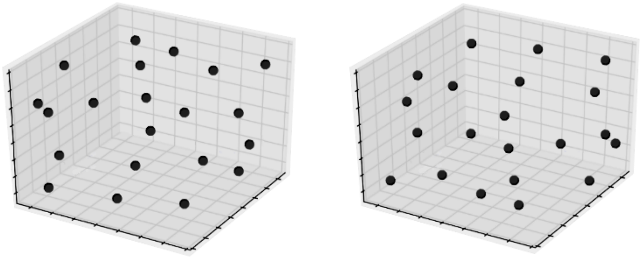 $\mathbf{v}_G$

Low $Gen(r)$ $(\leq \boldsymbol{\delta_{spec}})$

$\mathbf{v}_D$ 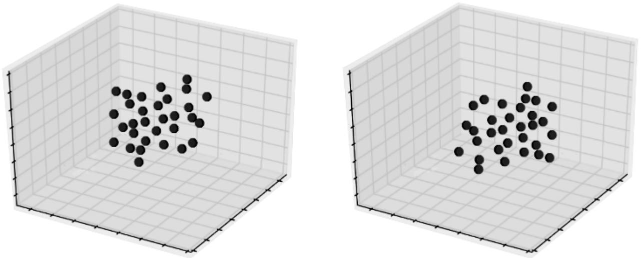 $\mathbf{v}_G$

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Disambiguation with Relation Context

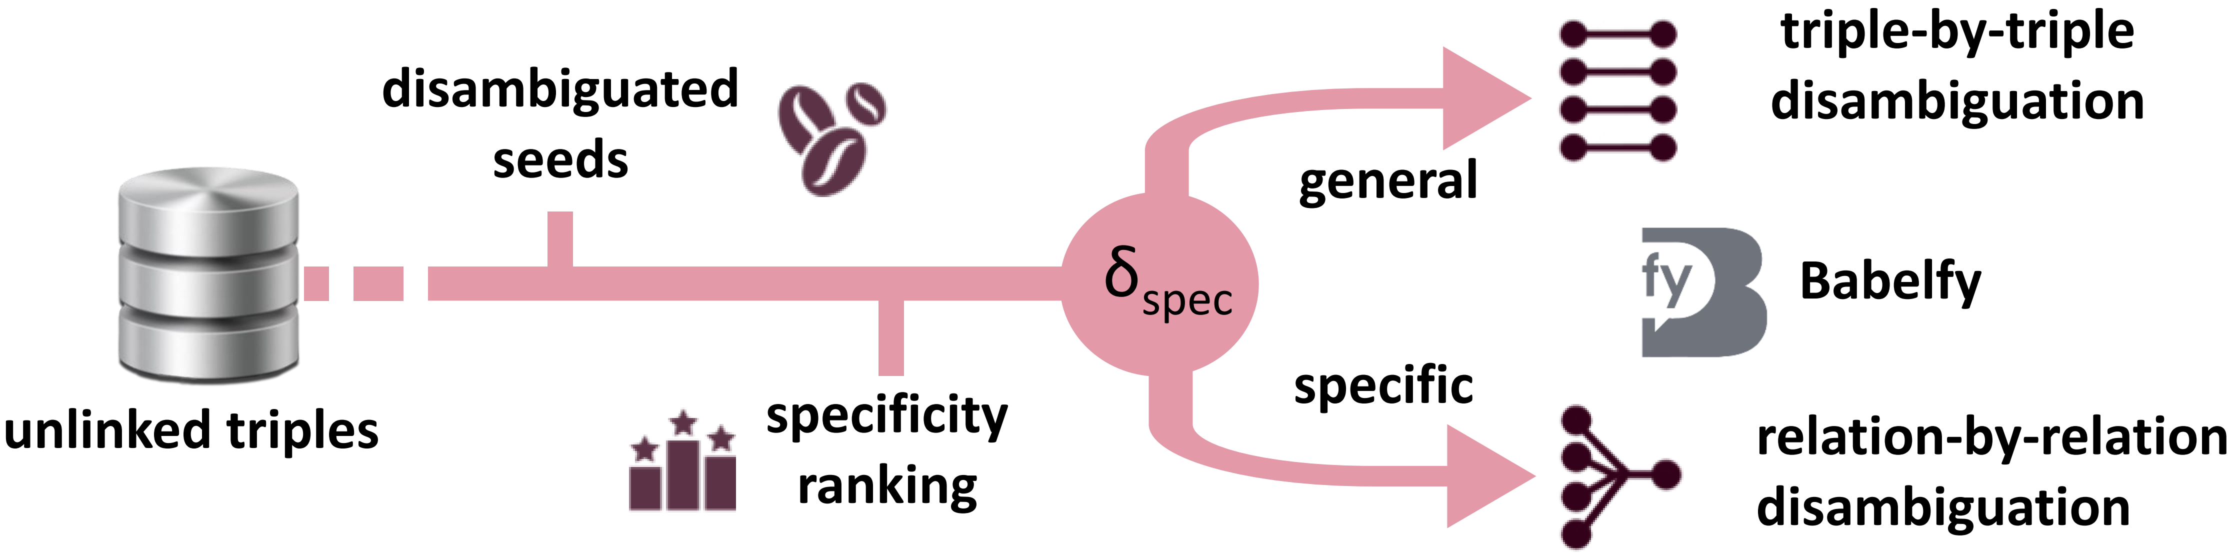

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

A bird's-eye view

represent each relation in the **unified vector space V_S** and compare them pairwise

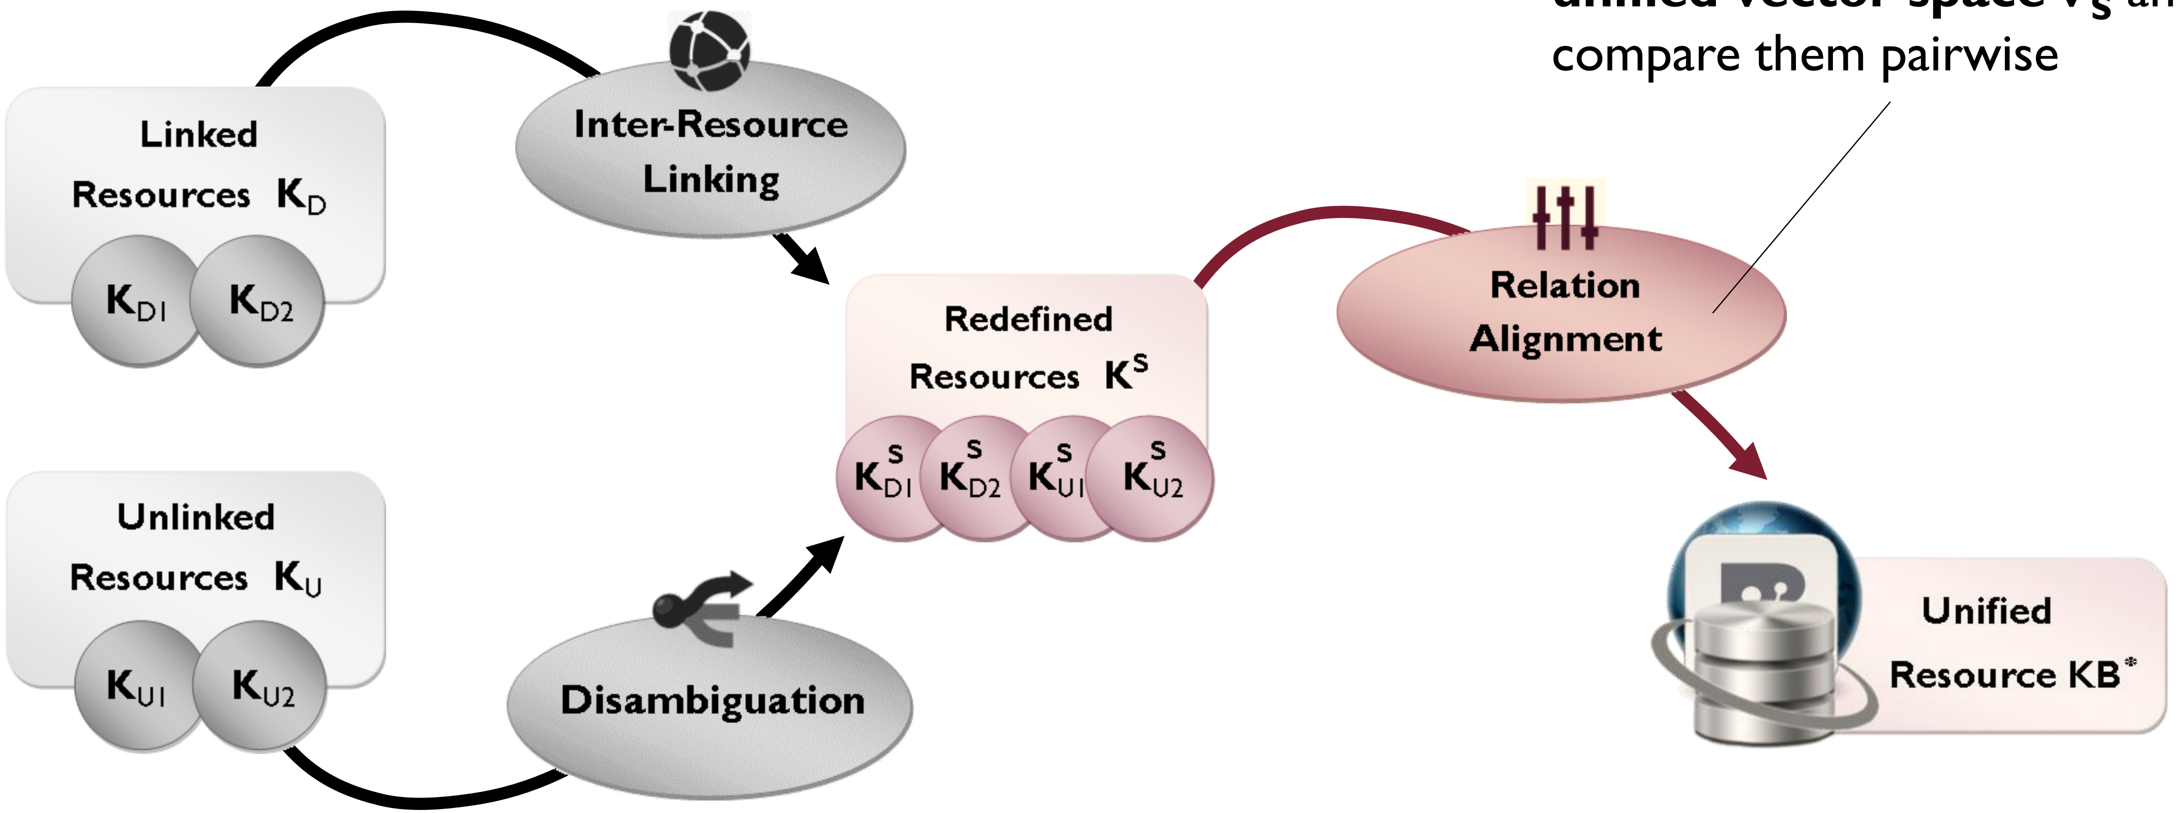

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Relation alignment

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Relation alignment

For each relation pair $\langle \boldsymbol{r_i}, \boldsymbol{r_j} \rangle$:

r_i	Relations	r_j
$\mu_D^{r_i}$ $\mu_G^{r_i}$	Centroids	$\mu_D^{r_j}$ $\mu_G^{r_j}$

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Relation alignment

For each relation pair $\langle \boldsymbol{r_i}, \boldsymbol{r_j} \rangle$:

r_i		**Relations**	r_j	
$\mu_D^{r_i}$	$\mu_G^{r_i}$	**Centroids**	$\mu_D^{r_j}$	$\mu_G^{r_j}$

Compare domain and range centroids pairwise:

$$\underset{k \in \{D,G\}}{s_k} = \frac{\mu_k^{r_i} \cdot \mu_k^{r_j}}{\|\mu_k^{r_i}\| \, \|\mu_k^{r_j}\|}$$

Relation Centroid Similarity

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Relation alignment

Fix a similarity threshold δ_{align}:

Domain Centroids

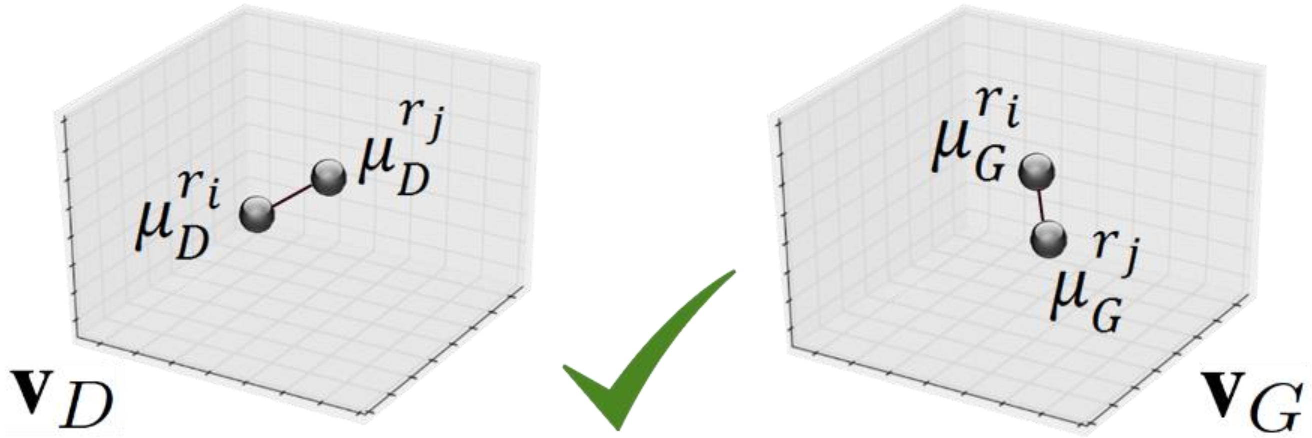

Range Centroids

$\frac{1}{2}(s_D + s_G) \geq \delta_{align}$? Align r_i and r_j and merge them in the same cluster

KB-Unify: How it works

http://lcl.uniroma1.it/kb-unify

Relation alignment

Fix a similarity threshold δ_{align}:

Domain Centroids

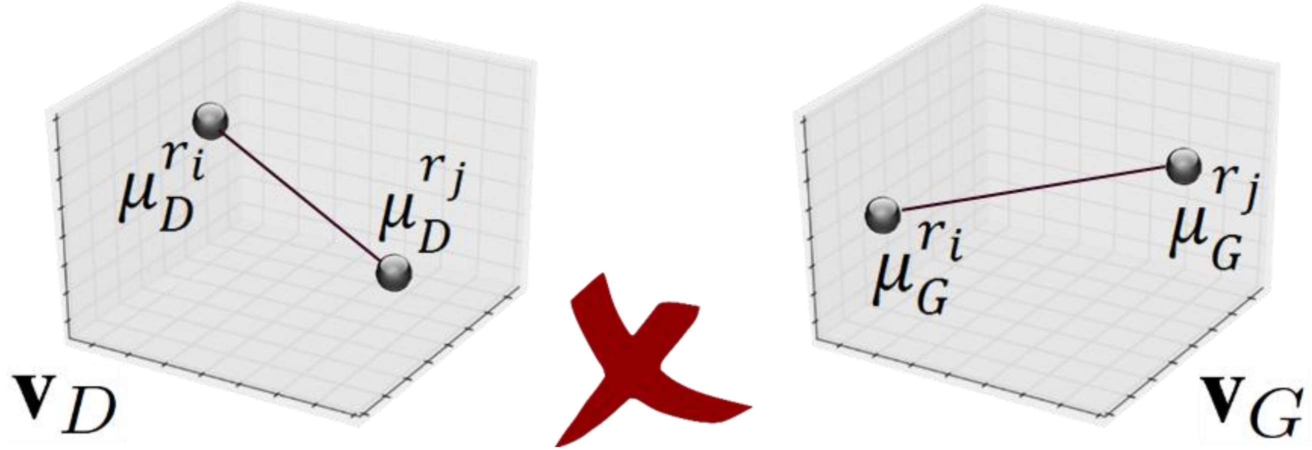

Range Centroids

$\frac{1}{2}(s_D + s_G) < \delta_{align}$? Leave r_i and r_j in separate clusters

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Evaluation

Experimental setup:

Linked Resources K_D : { **PATTY** , **WiSeNet** }

Unlinked Resources K_U : { **NELL** , **ReVerb** }

	PATTY	WiSeNet	NELL	ReVerb
relations	1,631,531	245,935	298	1,299,844
triples	15,802,946	2,271,807	2,245,050	14,728,268

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Disambiguation

Seed Precision:

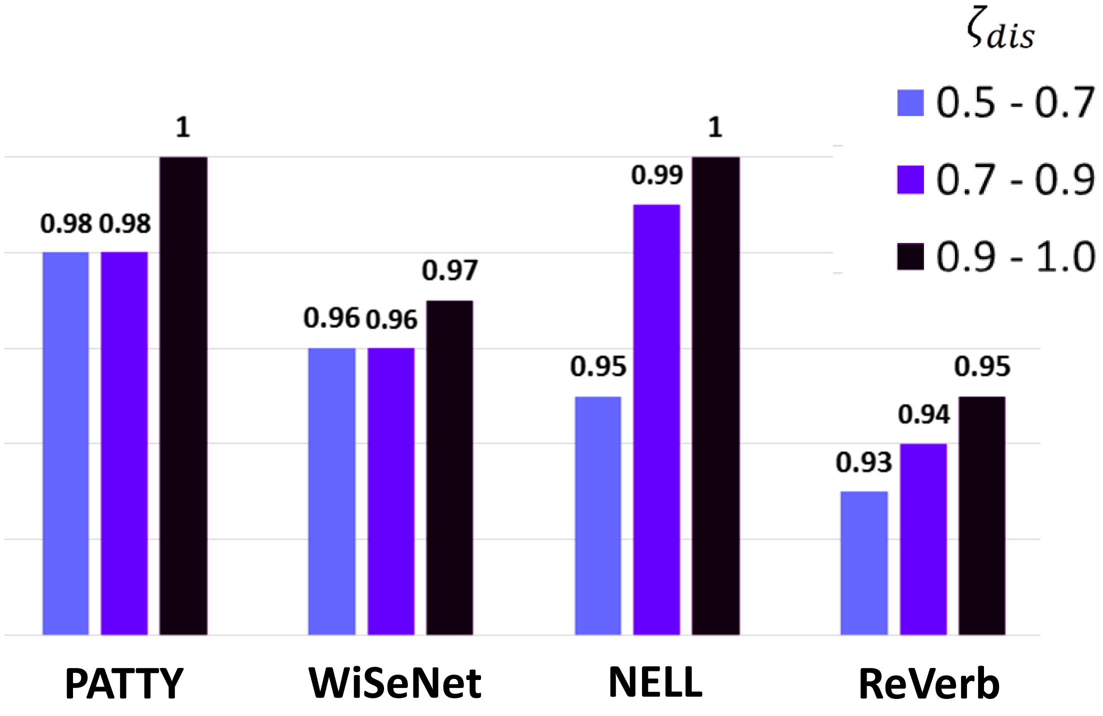

Coverage:

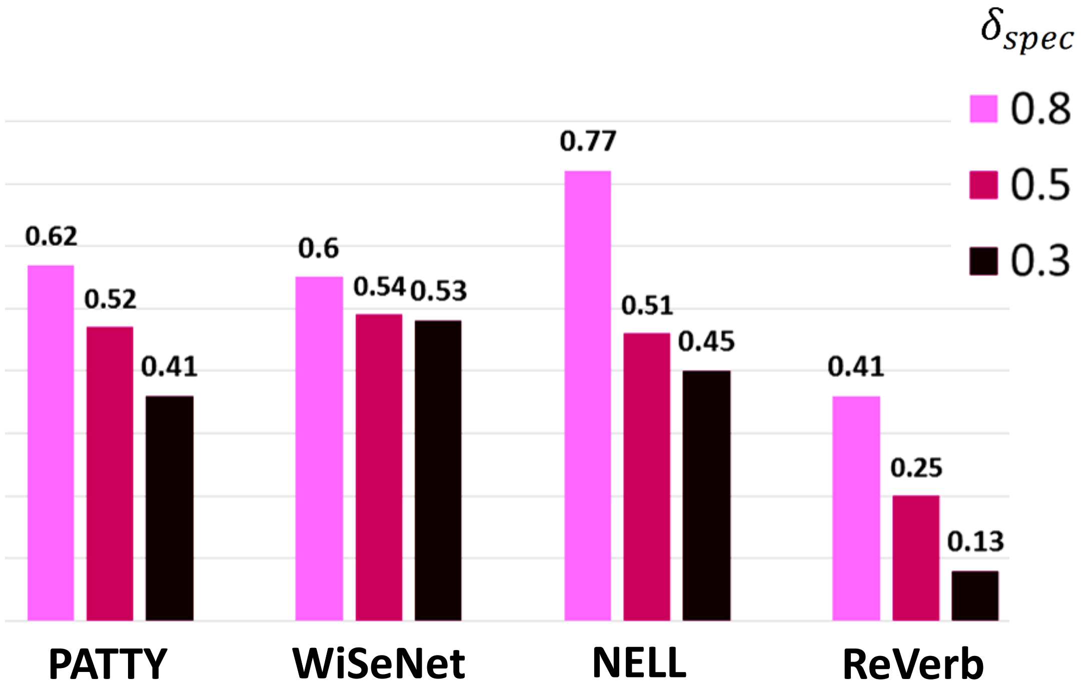

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Specificity ranking

For each ranked relation compute $\boldsymbol{Gen(r)}$ against the **average argument similarity** $\bar{s}$:

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Specificity ranking

For each ranked relation compute $\boldsymbol{Gen(r)}$ against the **average argument similarity** $\bar{s}$:

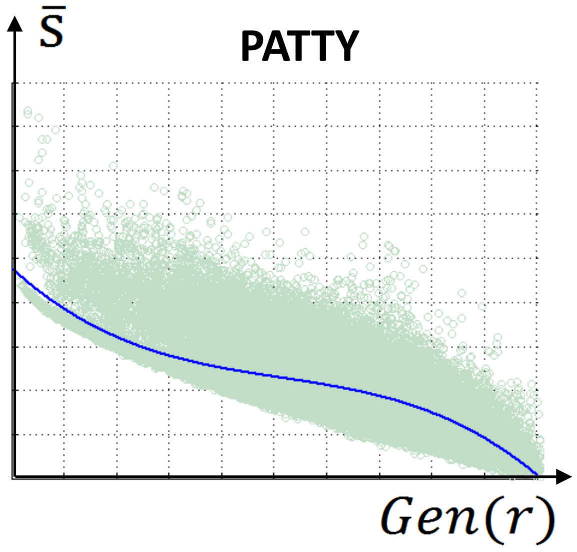

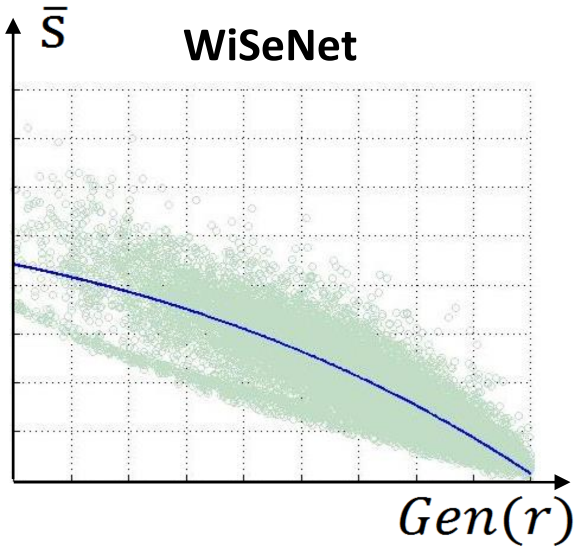

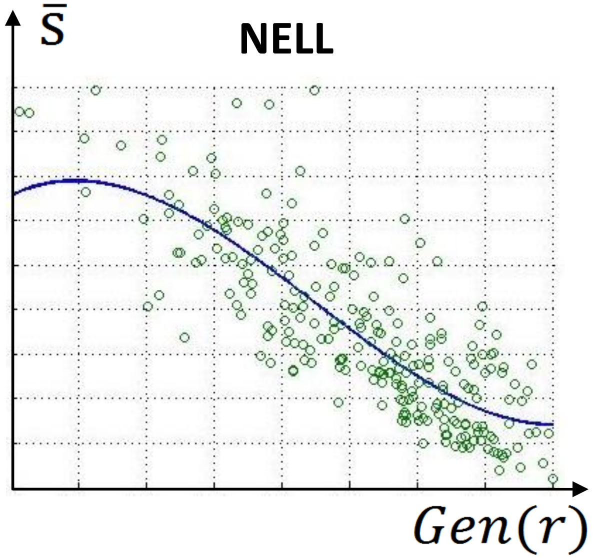

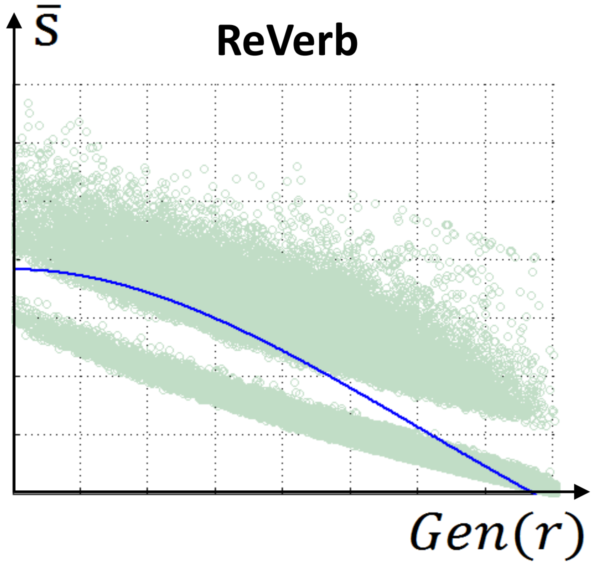

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Specificity ranking

For each ranked relation compute $Gen(r)$ against the **average argument similarity** $\bar{s}$:

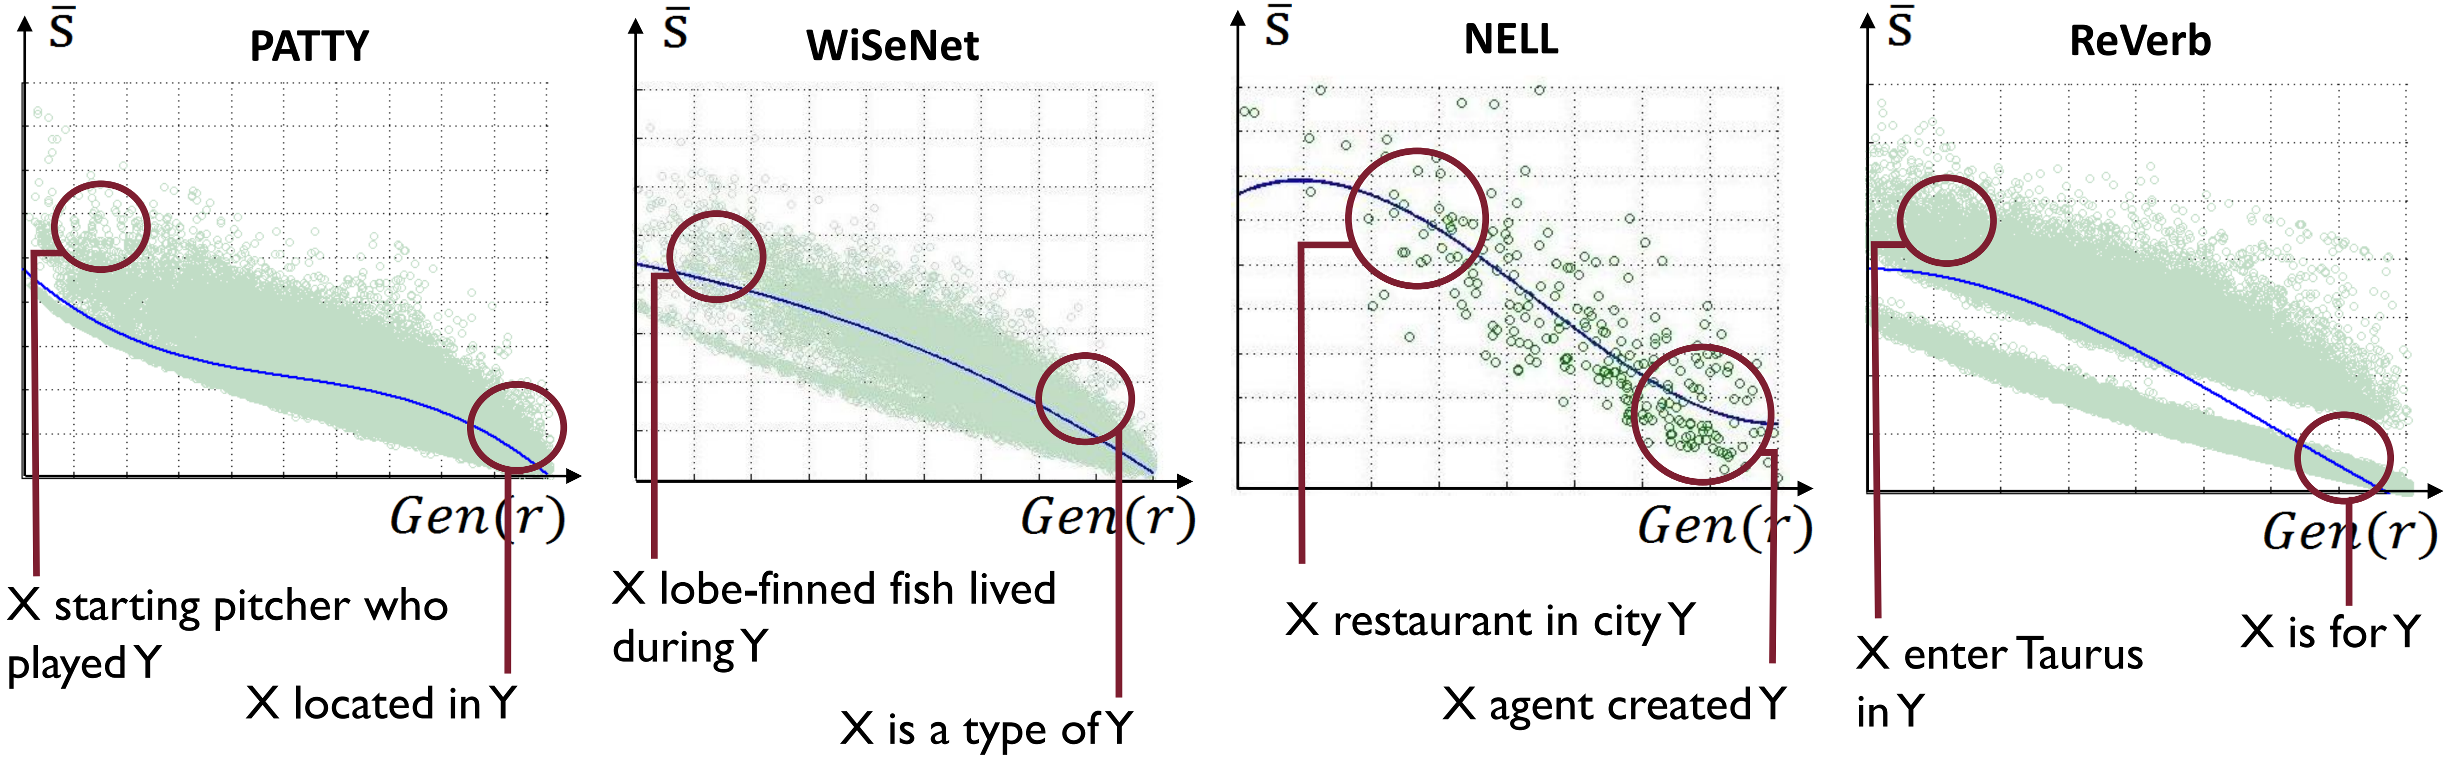

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Cross-resource relation alignment

Samples of **150 candidate alignments** for different alignment thresholds δ_{align} manually evaluated (in terms of **paraphrasing**) by two human judges

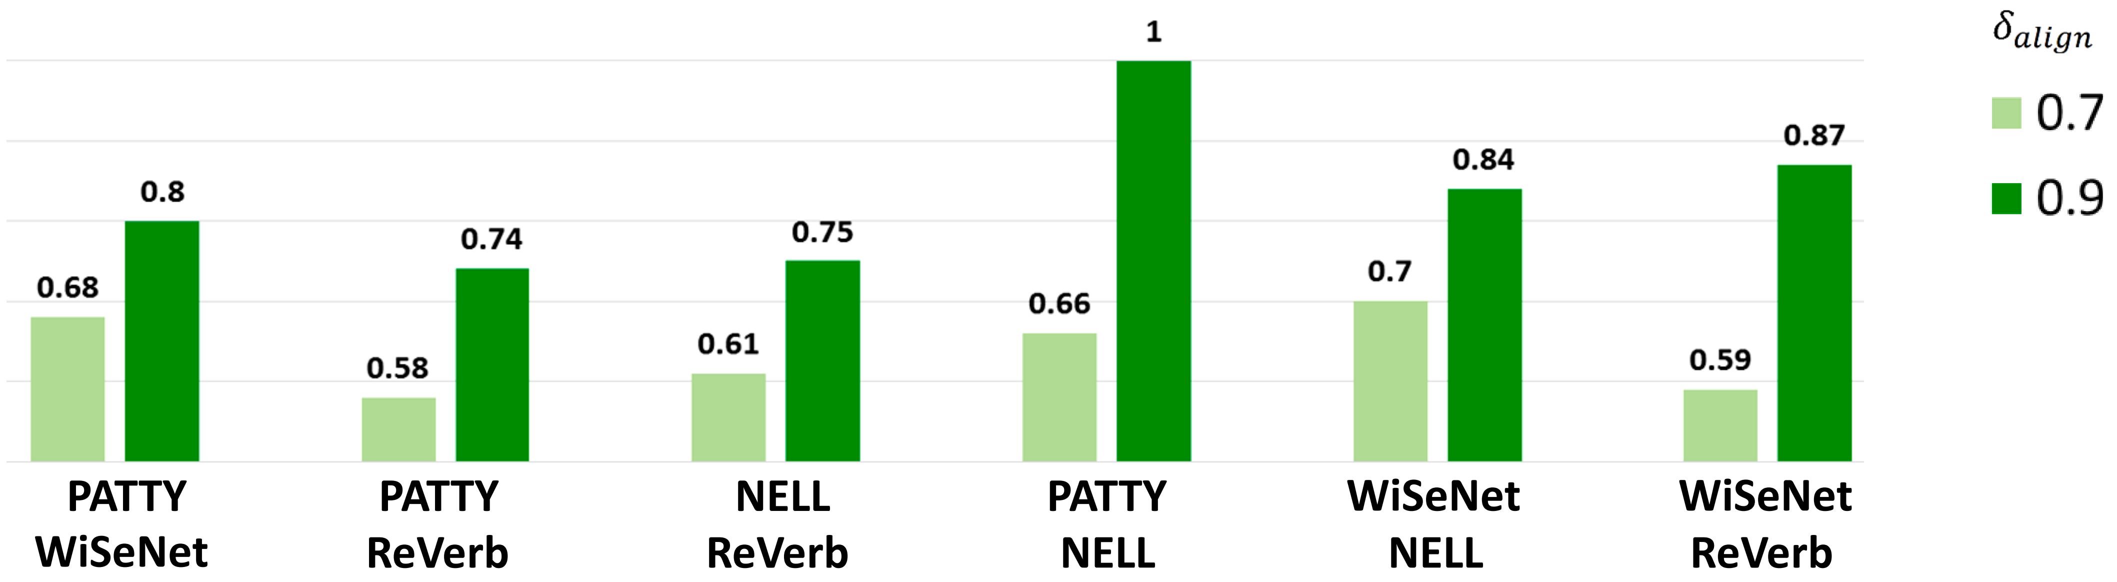

KB-Unify: Experiments

http://lcl.uniroma1.it/kb-unify

Cross-resource relation alignment

Some examples:

PATTY-WISENET		ζ_{align}
portrayed	's character	0.84
debuted in	first appeared in	0.86

PATTY-REVERB		ζ_{align}
language in	is spoken in	0.81
mostly known for	plays the role of	0.70

NELL-REVERB		ζ_{align}
bookwriter	is a novel by	0.88
personleadscity	is the mayor of	0.60

NELL-PATTY		ζ_{align}
worksfor	was hired by	0.72
riveremptiesintoriver	tributary of	0.89

NELL-WISENET		ζ_{align}
animaleatfood	feeds on	0.72
teamhomestadium	play their home games at	0.88

REVERB-WISENET		ζ_{align}
has a selection of	offers	0.82
had grown up in	was born and raised in	0.85

Wrap up and Conclusion

Wrap up and Conclusion

DefIE: A full-fledged OIE pipeline targeted to textual definitions, with explicit semantic characterization of both arguments and relation patterns

Wrap up and Conclusion

DefIE: A full-fledged OIE pipeline targeted to textual definitions, with explicit semantic characterization of both arguments and relation patterns

KB-Unify: An approach to knowledge base disambiguation and unification based on a shared sense inventory and a sense-based vector space model

Wrap up and Conclusion

Take-home message(s):

Web-scale OIE is absolutely great, but…

1. **Definitional knowledge is important**: sometimes it is worth it to just step back and analyze from where valuable information is extracted (**quality vs. quantity**)

2. **Making sense of the output is important:** semantic analysis can be used to let different OIE outputs "speak to each other" and benefit from mutual enrichment

Thanks!

¡Gracias!

Gràcies!

Grazie!

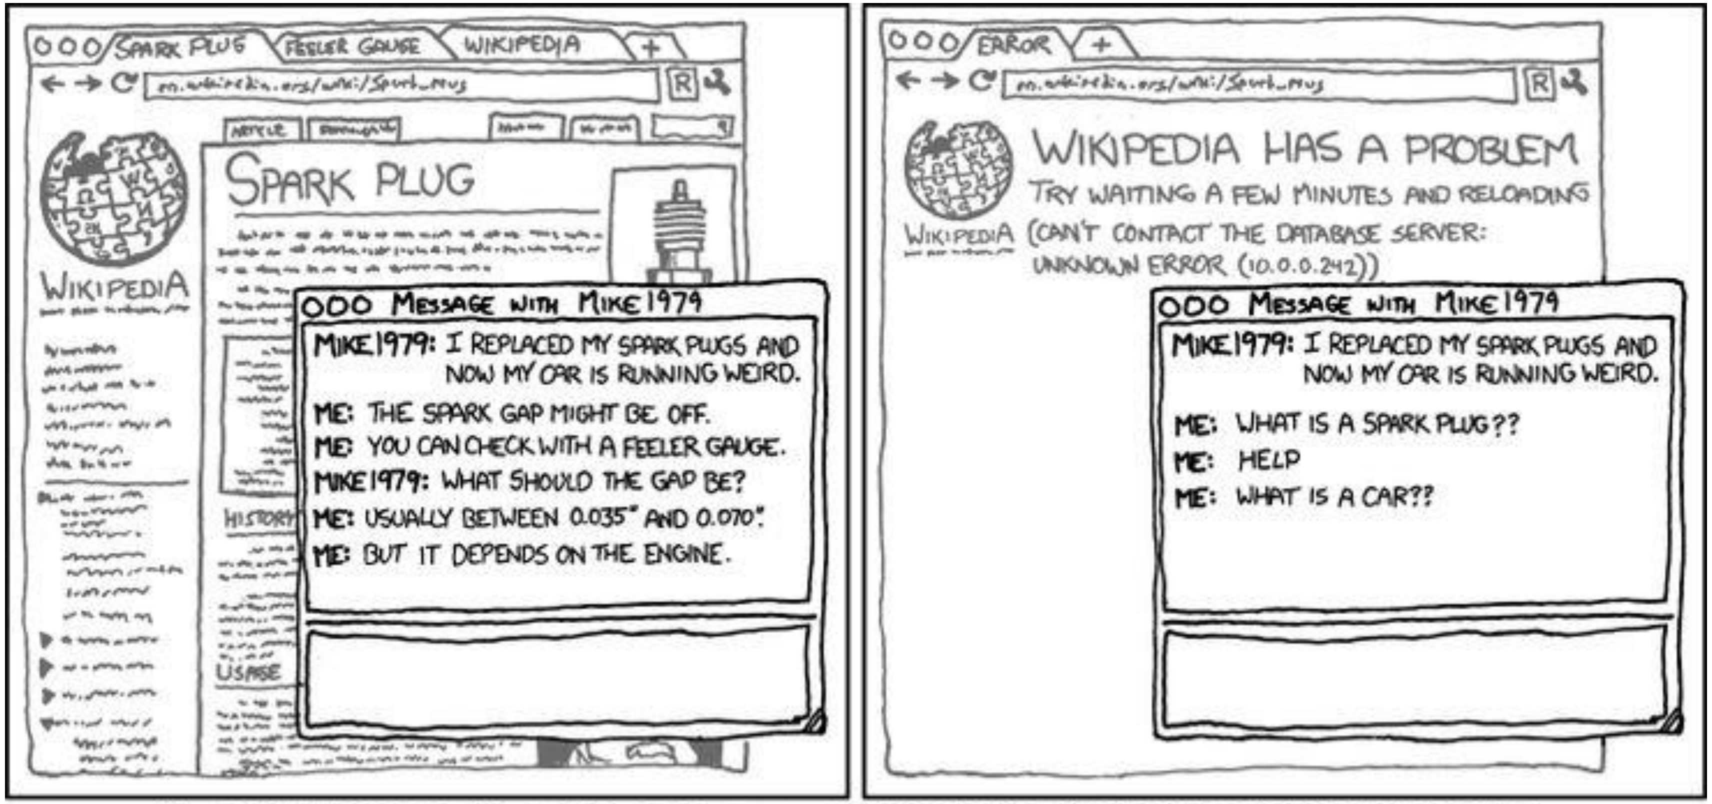

WHEN WIKIPEDIA HAS A SERVER OUTAGE, MY APPARENT IQ DROPS BY ABOUT 30 POINTS.