Negative Statements Considered Useful

Hiba Arnaout

max planck institut informatik

Myself: Hiba Arnaout

- PhD Candidate at Max Planck Institute for Informatics, Germany.

- MS of CS from the American University of Beirut, Lebanon.

- Research stays at BoschAI (21), The University of Edinburgh (20).

- Research interests:
  - *Former*: IR over KBs
  - *Current*: KB curation (negation)
  - *Next*: social computing (*let's talk!*)

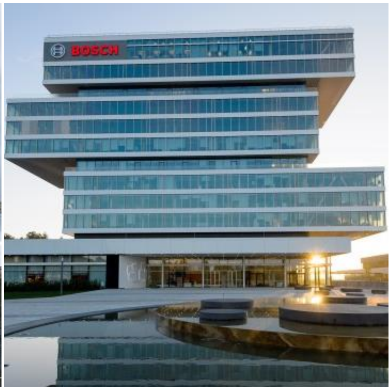

Outline

- **Knowledge Bases (KBs)**
  – *intro*

- **Effective Searching of KBs**
  – *1 slide*

- **The Case for Explicit Negation in KBs**
  – *remainder of presentation*

Outline

- **Knowledge Bases (KBs)**
  – *intro*
- **Effective Searching of KBs**
  – *1 slide*
- **The Case for Explicit Negation in KBs**
  – *remainder of presentation*

KBs not yet in action

Traditional search

KBs in action

Semantic search

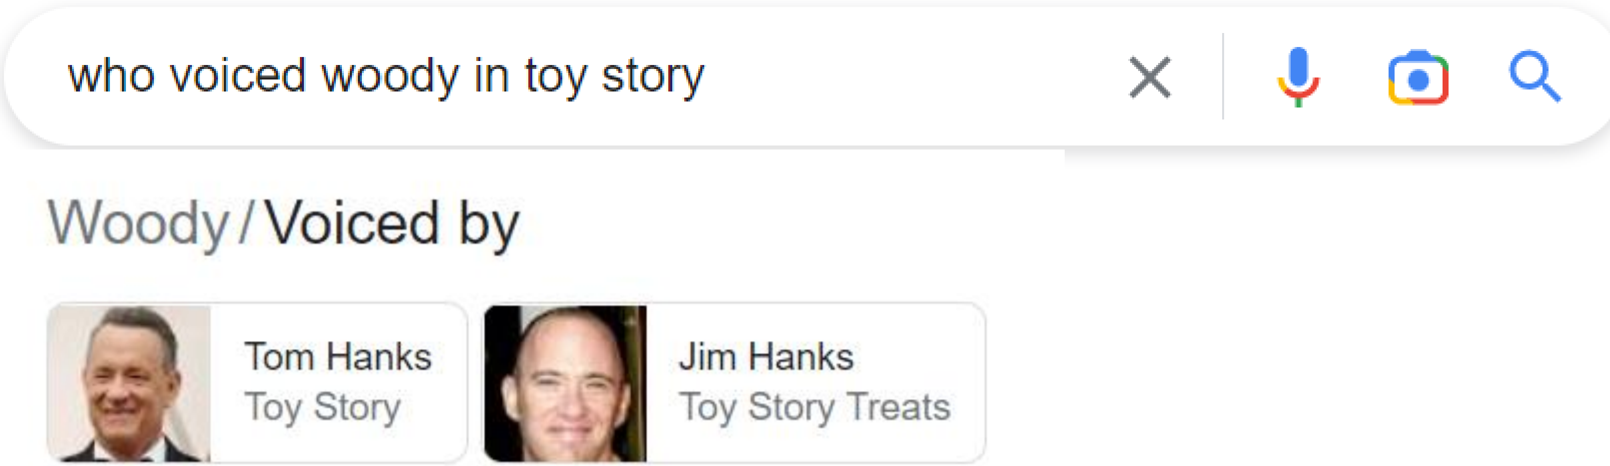

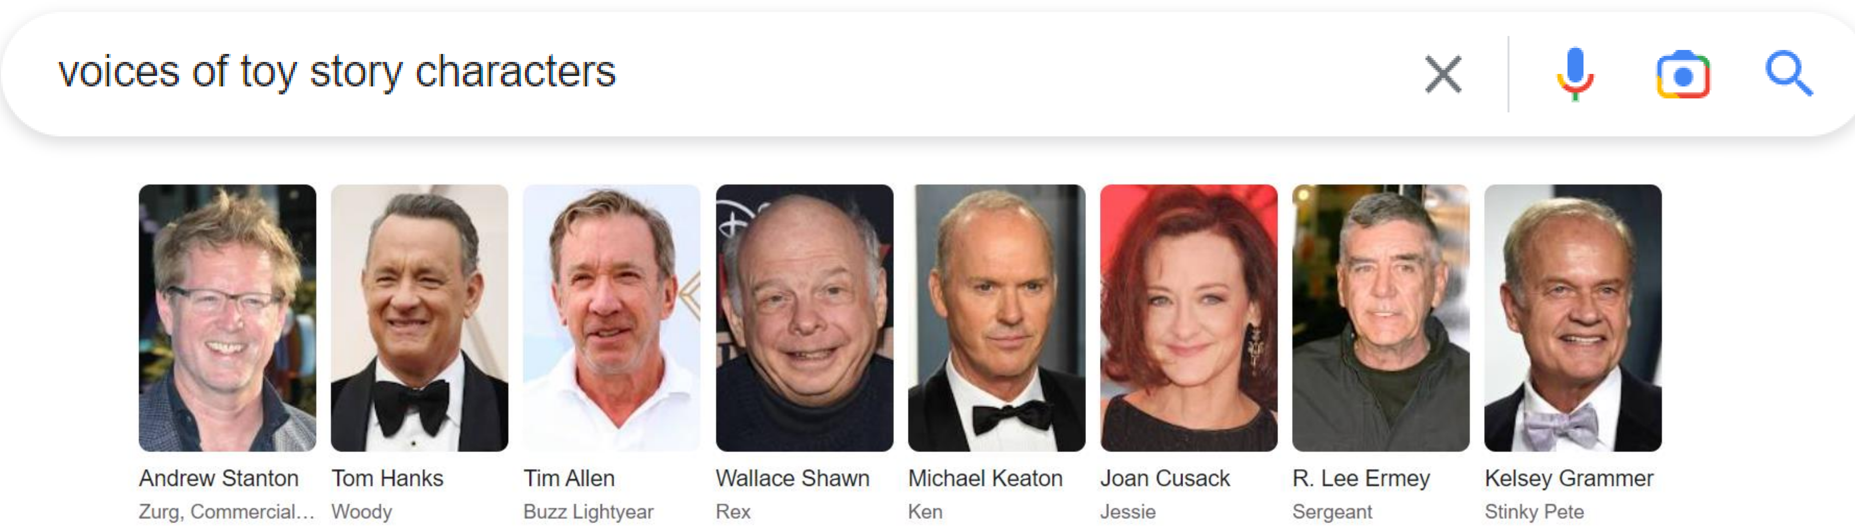

Knowledge Base (KB)

`(Subject, Predicate, Object)` triples about entities

```
(Tom Hanks, type, actor)
(Sheriff Woody Pride, type, fictional character)
(Sheriff Woody Pride, voice, Tom Hanks)
(Tom Hanks, birthName, ‘‘Thomas Jeffrey Hanks’’)
```

`Subject` entity
`Predicate` pre-defined relation
`Object` entity or literal

KBs are awesome

Impactful public and commercial projects:
Wikidata, YAGO, ConceptNet, Google, LinkedIn, ..

Applications:
- Cooking-chatbot need to return *items in a niçoise salad?*
- Query for *members of EU?*
- *Physician wants to speed-read on unfamiliar disease?*
- *Query Amazon's product graph to compare two smartphones?*

Colors indicate domain of KB

Each node is a public KB.

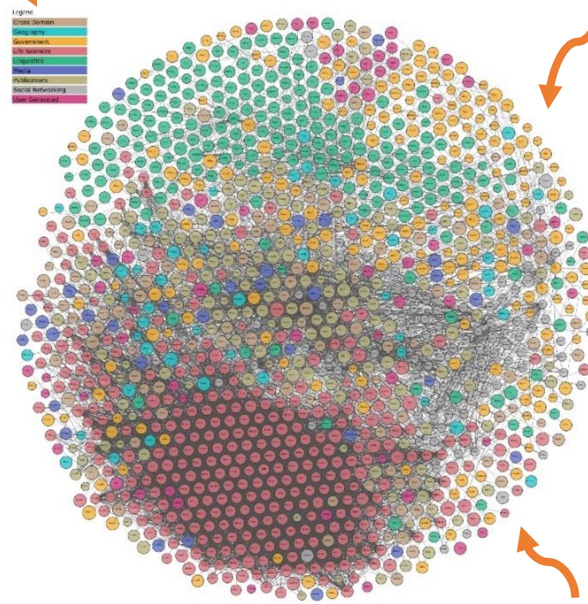

Links mean entities have equivalency in another KB.

Thousands of KBs
Billions of entities and statements

[Noy et al., Industry-Scale Knowledge Graphs: Lessons and Challenges, CACM'19]
[Weikum et al., Machine Knowledge: Creation and Curation of Comprehensive KBs, FnT'21]
https://lod-cloud.net

Outline

- **Knowledge Bases (KBs)**
  – *intro*
- **Effective Searching of KBs**
  – *1 slide*
- **The Case for Explicit Negation in KBs**
  – *remainder of presentation*

Effective querying of KBs

Semantic web problems using standard IR techniques:

- Ranking large result sets
- Extend triples with keywords
- Automatic query relaxation
- Result diversification
- Diversity-aware evaluation metric

Al Pacino movies?

```
<?m castMember Al Pacino>
```

?m in KB
Chinese Coffee
The Godfather
..

```
<?m starring Al Pacino>
<Al Pacino actedIn ?m>
<?m actor Al Pacino>
<?m actor Alfredo Pacino>
..
```

?m in KB
Empty list of results.

[Arnaout and Elbassuoni, Effective Searching of RDF Knowledge Graphs, JWS'18]
[Arnaout and Elbassuoni, Result Diversity for RDF Search, IC3K'16]

Outline

- **Knowledge Bases (KBs)**
  – *intro*
- **Effective Searching of KBs**
  – *1 slide*
- **The Case for Explicit Negation in KBs**
  – *remainder of presentation*

But KBs are incomplete..

Is what they know true?

Most studies focus on precision
e.g., YAGO 95%

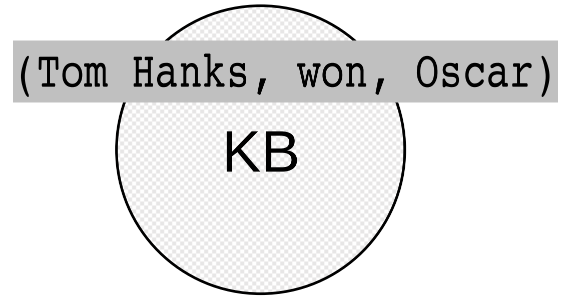

But, do they know what is true?
(recall/coverage/completeness)

Web-scale KBs

	Closed-world Assumption (CWA)	Open-world Assumption (OWA)
`(Tom Hanks, won, Oscar)?`	→ *Yes*	→ *Yes*
`(Tom Hanks, won, Nobel in Physics)?`	→ *No*	→ ***Maybe***

Enhance KB coverage, or?

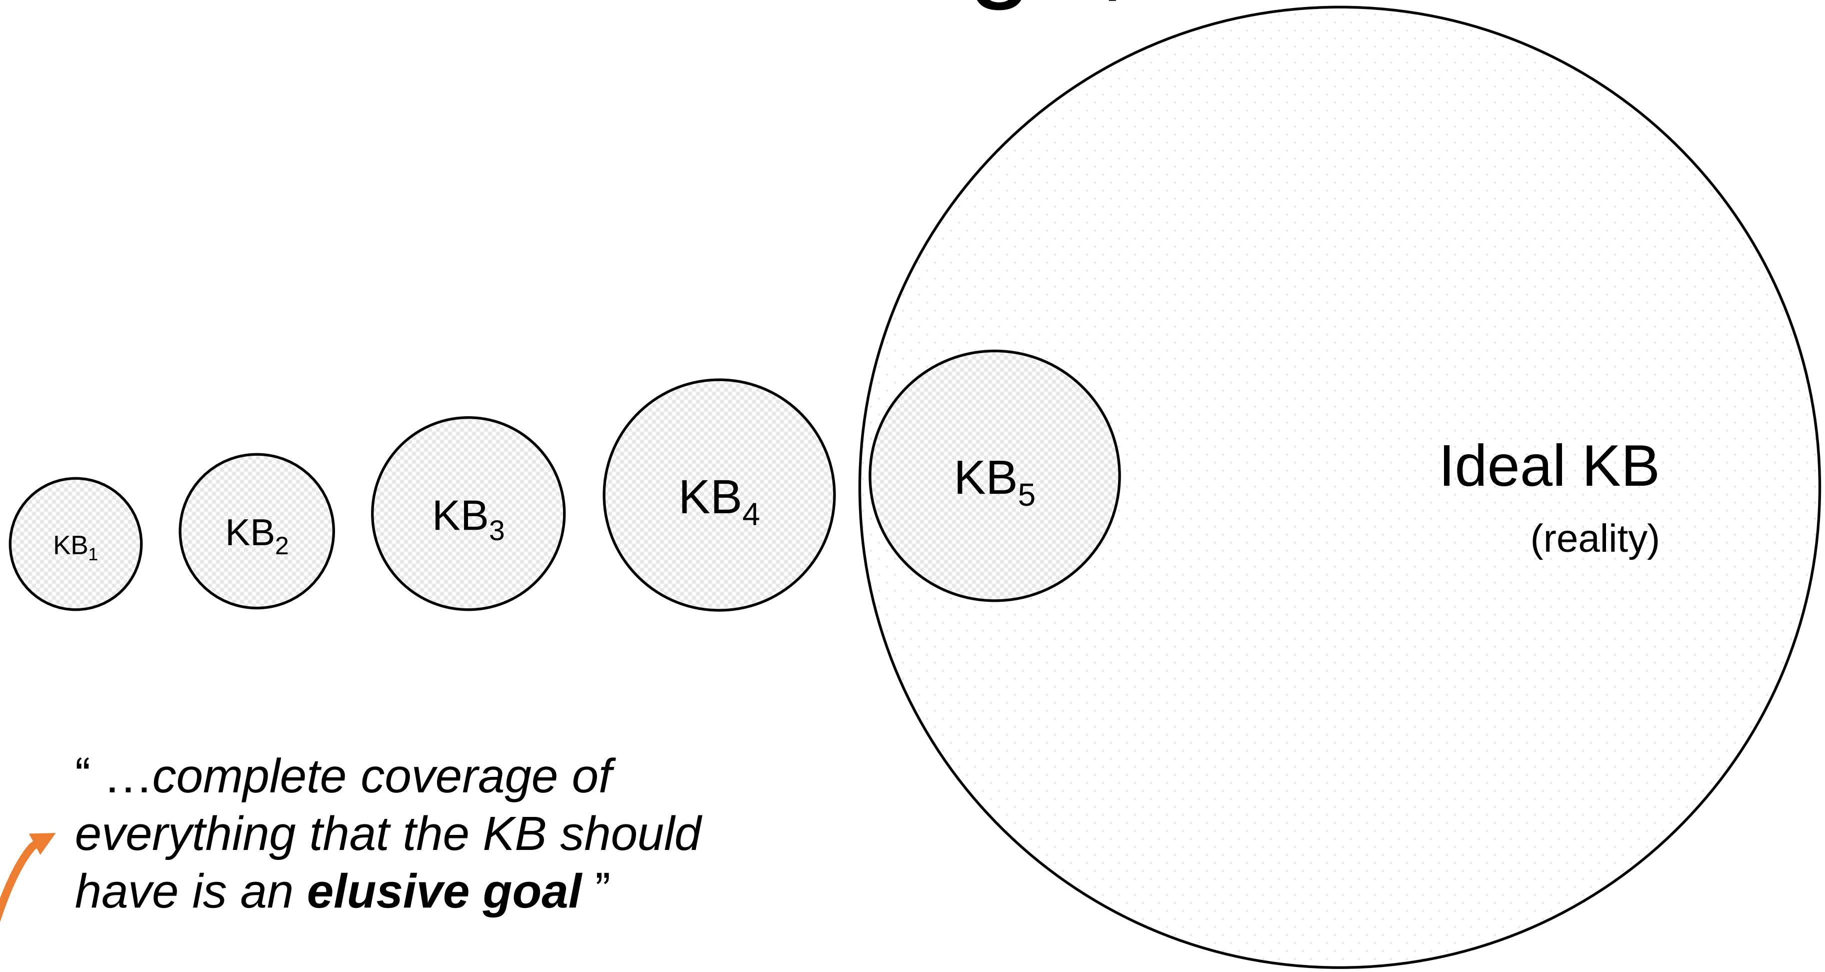

“ *…complete coverage of everything that the KB should have is an **elusive goal*** ”

Is the KB ~~complete~~?
useful

[Weikum et al., Machine Knowledge: Creation and Curation of Comprehensive KBs, FnT’21]

Important negative knowledge?

Proposal: Enrich KBs with *explicit* negations

- CWA (Yes, No)
- OWA (Yes, Maybe)
- PCWA: Partial-closed World Assumption
  (Yes, No, & Maybe)

Samples

`NOT(Tom Hanks, won, Nobel in Physics)` NOT IMPORTANT

`NOT(Stephen Hawking, won, Nobel in Physics)` IMPORTANT

`NOT(Peanut, isA, sport)` NOT IMPORTANT

`NOT(Peanut, isA, nut)` IMPORTANT *It's a legume*

Negations in KB: applications

- Allergic to pets, looking for a hotel that DOESN'T allow them?

→ *Ask a KB*

- I DON'T have an oven. What can I make for dinner?

→ *Ask a chatbot with access to KB*

- Celebrities who NEVER received an Oscar?

→ Ask a QA system over KB

- This iPhone DOESN'T come with earphones.
- Negative sampling for ML problems, i.e., strong negatives.
- …

Discovering informative negations in open-world KBs: project timeline

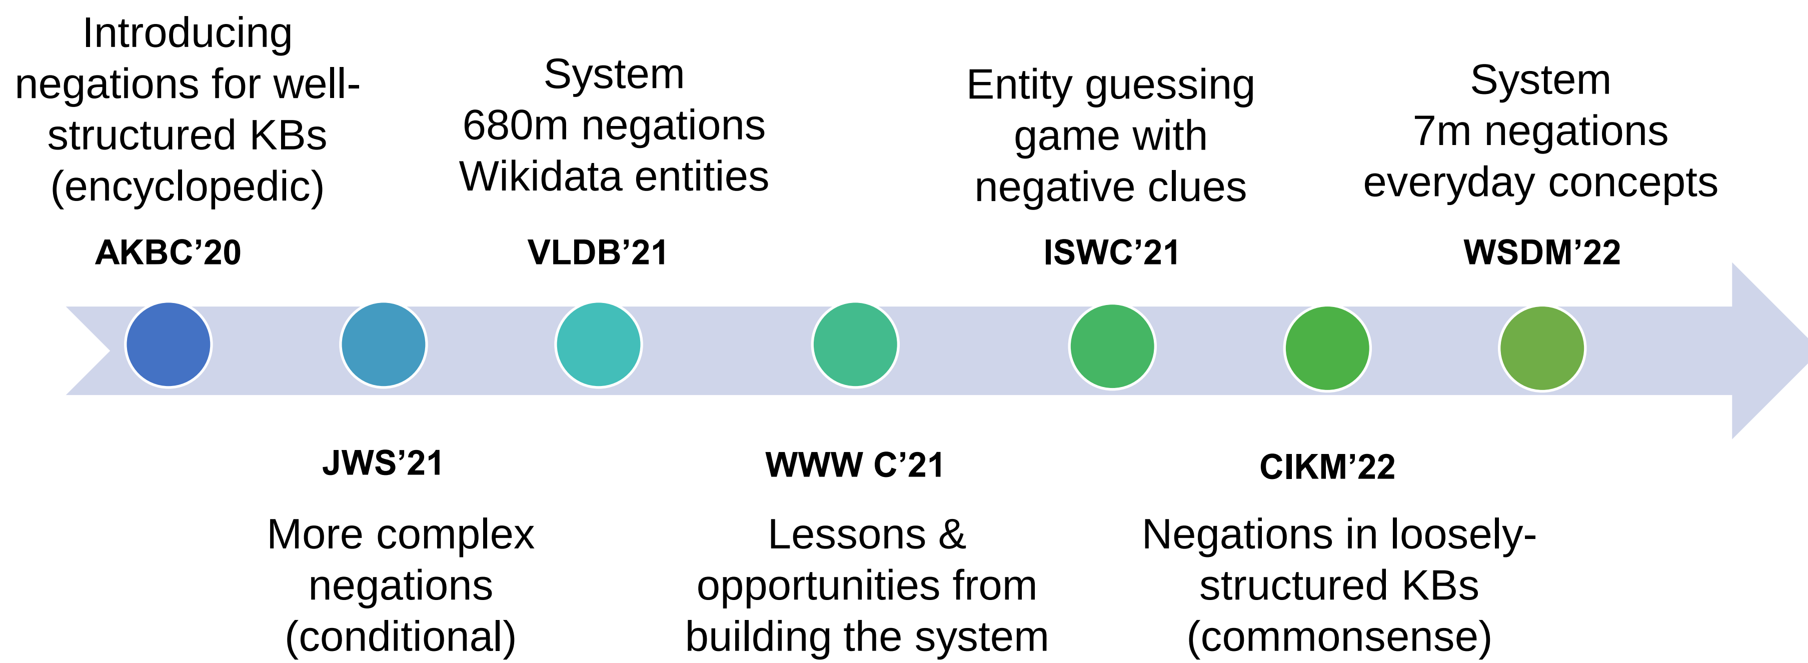

Discovering informative negation in open-world KBs: project timeline

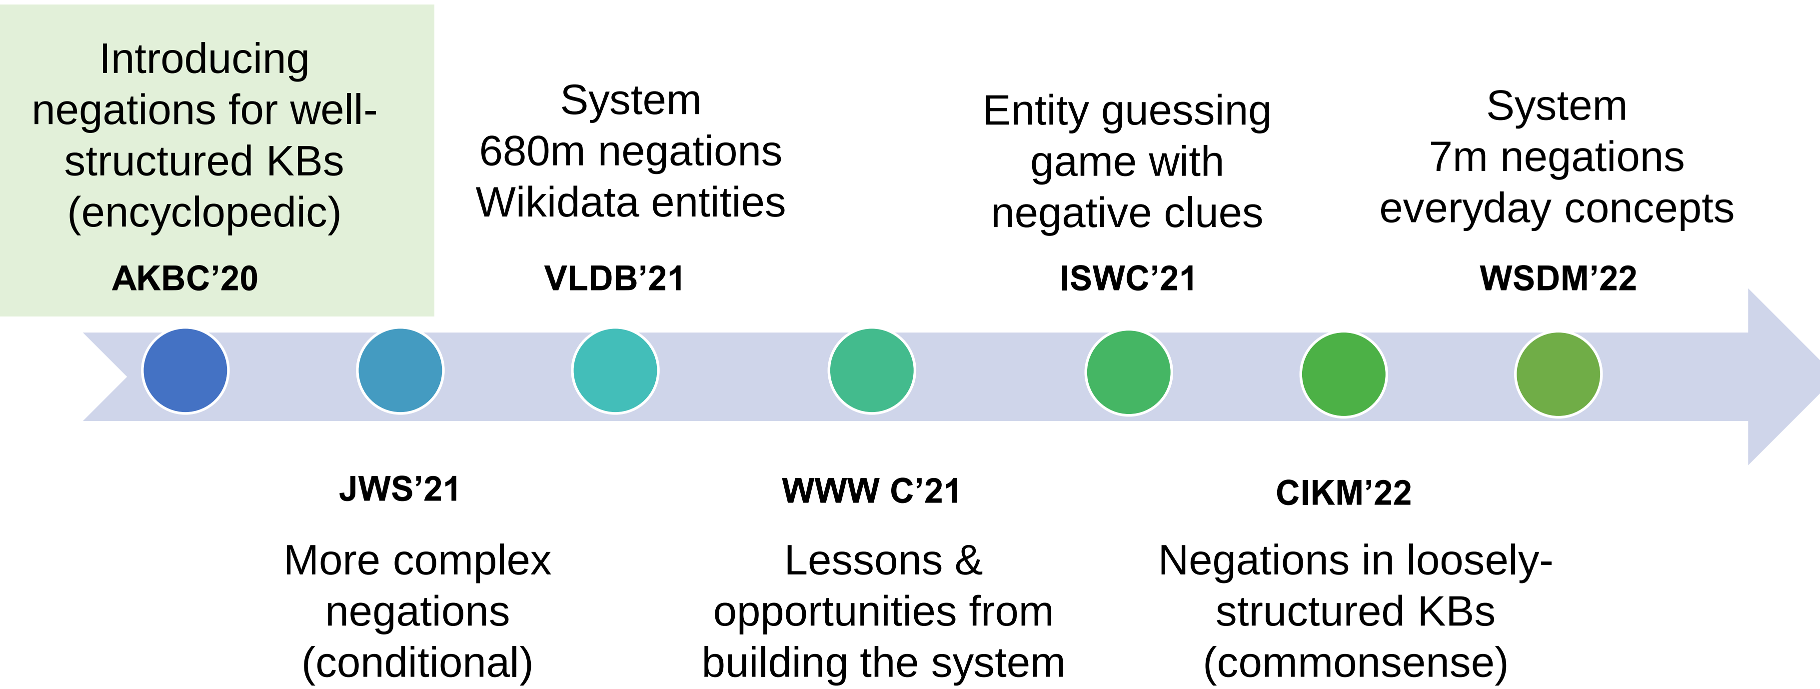

Discovering informative negation in open-world KBs: project timeline

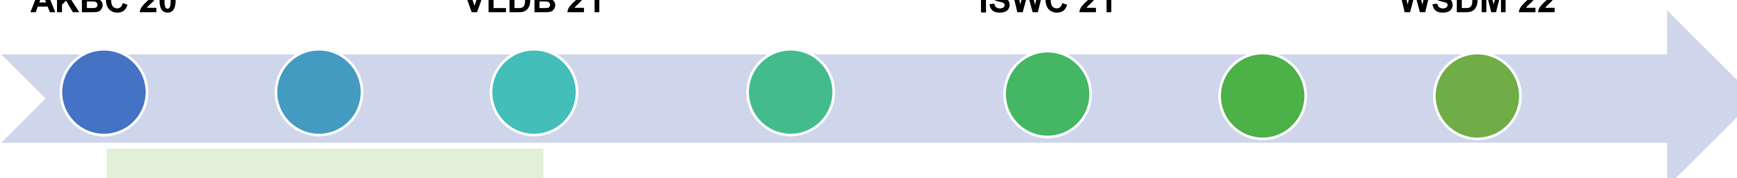

AKBC’20
Introducing negations for well-structured KBs (encyclopedic)

JWS’21
More complex negations (conditional)

VLDB’21
System
680m negations
Wikidata entities

WWW C’21
Lessons & opportunities from building the system

ISWC’21
Entity guessing game with negative clues

CIKM’22
Negations in loosely-structured KBs (commonsense)

WSDM’22
System
7m negations
everyday concepts

Discovering informative negation in open-world KBs: project timeline

AKBC’20
Introducing negations for well-structured KBs (encyclopedic)

JWS’21
More complex negations (conditional)

VLDB’21
System
680m negations
Wikidata entities

WWW C’21
Lessons & opportunities from building the system

ISWC’21
Entity guessing game with negative clues

CIKM’22
Negations in loosely-structured KBs (commonsense)

WSDM’22
System
7m negations
everyday concepts

Discovering informative negation in open-world KBs: project timeline

AKBC’20
Introducing negations for well-structured KBs (encyclopedic)

JWS’21
More complex negations (conditional)

VLDB’21
System
680m negations
Wikidata entities

WWW C’21
Lessons & opportunities from building the system

ISWC’21
Entity guessing game with negative clues

CIKM’22
Negations in loosely-structured KBs (commonsense)

WSDM’22
System
7m negations
everyday concepts

Discovering informative negation in open-world KBs: project timeline

- **AKBC'20** — Introducing negations for well-structured KBs (encyclopedic)
- **JWS'21** — More complex negations (conditional)
- **VLDB'21** — System 680m negations Wikidata entities
- **WWW C'21** — Lessons & opportunities from building the system
- **ISWC'21** — Entity guessing game with negative clues
- **CIKM'22** — Negations in loosely-structured KBs (commonsense)
- **WSDM'22** — System 7m negations everyday concepts

Discovering informative negation in open-world KBs: project timeline

AKBC'20
Introducing negations for well-structured KBs (encyclopedic)

JWS'21
More complex negations (conditional)

VLDB'21
System
680m negations
Wikidata entities

WWW C'21
Lessons & opportunities from building the system

ISWC'21
Entity guessing game with negative clues

CIKM'22
Negations in loosely-structured KBs (commonsense)

WSDM'22
System
7m negations
everyday concepts

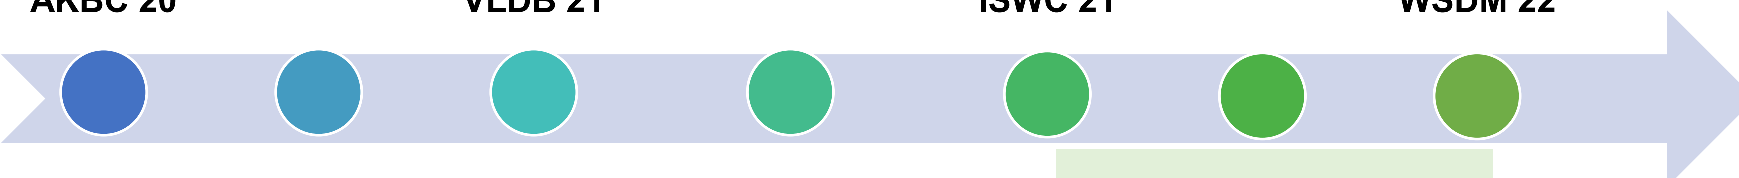

Discovering informative negation in open-world KBs: project timeline

Introducing negations for well-structured KBs (encyclopedic)
AKBC’20

JWS’21
More complex negations (conditional)

System
680m negations
Wikidata entities
VLDB’21

WWW C’21
Lessons & opportunities from building the system

Entity guessing game with negative clues
ISWC’21

CIKM’22
Negations in loosely-structured KBs (commonsense)

System
7m negations
everyday concepts
WSDM’22

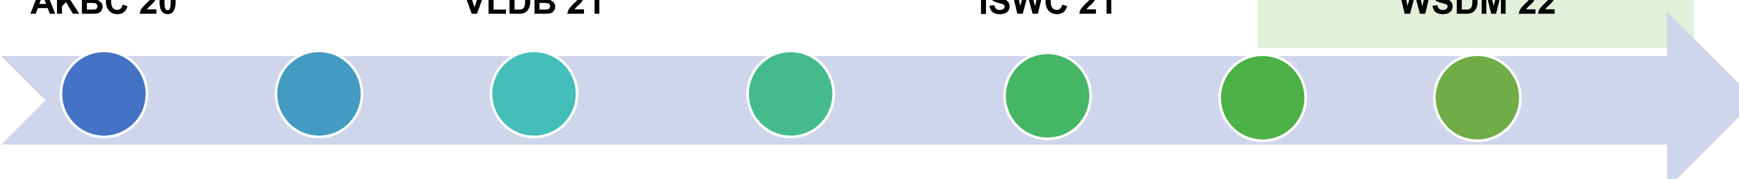

Discovering informative negation in open-world KBs: project timeline

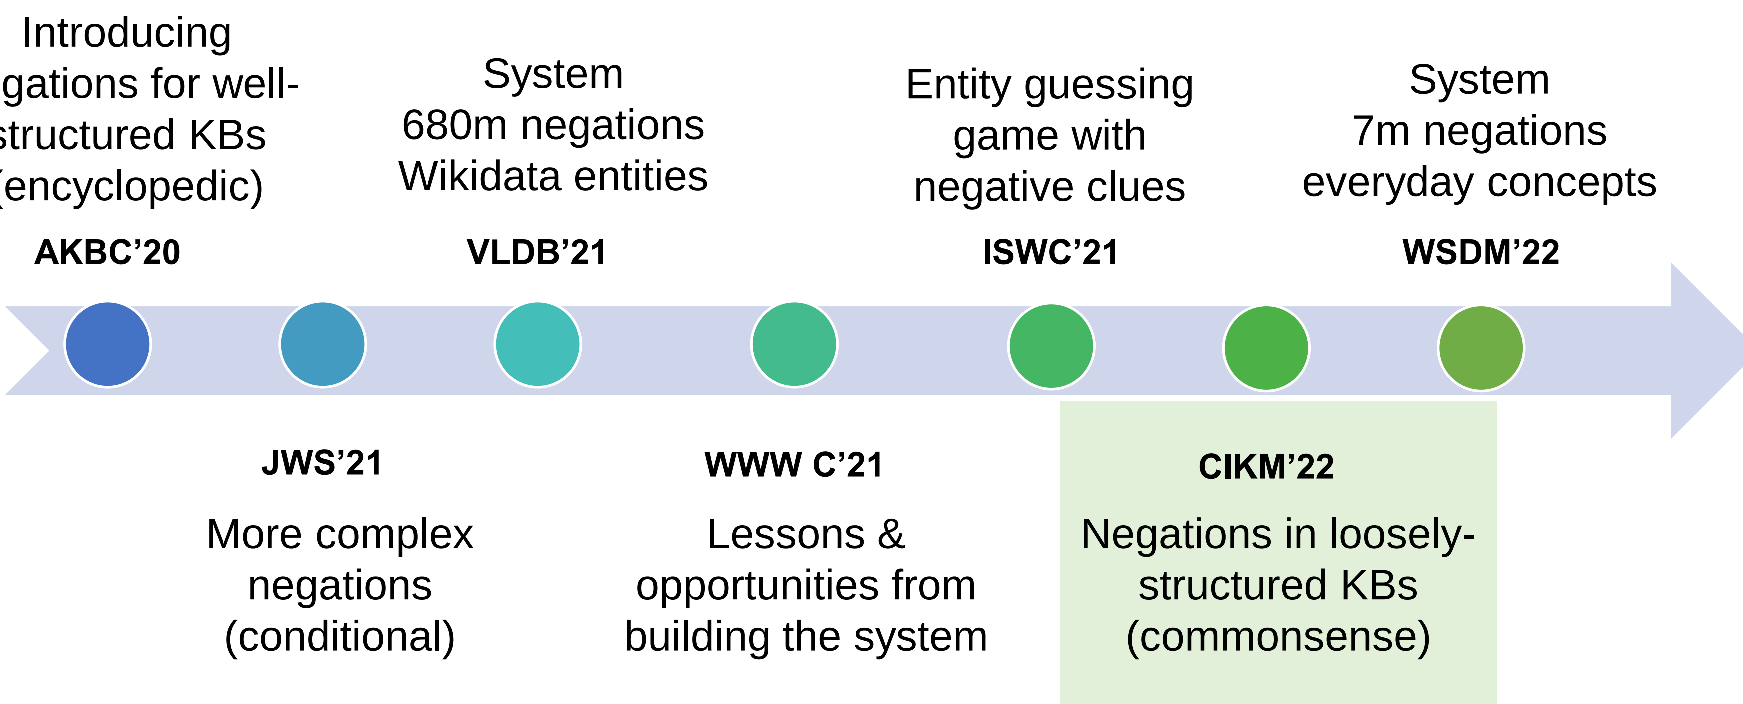

Commonsense knowledge (CSK)

Semi-structured knowledge about everyday concepts

- athlete, laptop, book,…

not to be confused with instances

~~*- Cristiano Ronaldo*~~*,* ~~*Dell XPS*~~*,* ~~*Harry Potter*~~ *…*

Long history in AI

- McCarthy (1959)
- Knowledge Bases: ConceptNet (1999) .., AscentKB (2021)

Commonsense Knowledge Bases (CSKB)

Store CSK as (`head`, `tail`) statements

- **Positive:**
  `(basketball, is a team sport)`

`Head` entity (concept)
`Tail` short phrase

- **Negative:**
  `NOT(basketball, is played on grass)`

Existing CSKBs focus on obtaining positives

ConceptNet < 2% negated statements

Enrich concepts in CSKBs with *informative* negations.

Existing solutions

- **Baseline: Closed-world Assumption**
  absent as negative → elephant has no eye, is not a song
  Many wrong or uninformative negatives

- **Text-based**: Mining query logs [CIKM'19]
  why can't elephants.. → jump, use computers, hide in trees
  Low informativeness

- **Corruption-based**: Negatives from positives [EMNLP'21]
  `(horse, is a pet) → NOT(horse rider, is a pet)`
  Type-inconsistency

- **LM-based**: Prompting LMs [ACL'20]
  Horses are not.. animal; they are mammals.
  Low accuracy

Existing solutions

- **Baseline: Closed-world Assumption**
  absent as negative → elephant has no eye, is not a song
  Many wrong or uninformative negatives

- **Text-based**: Mining query logs [CIKM'19]
  why can't elephants.. → jump, use computers, hide in trees
  Low informativeness

- **Corruption-based**: Negatives from positives [EMNLP'21]
  `(horse, is a pet) → NOT(horse rider, is a pet)`
  Type-inconsistency

- **LM-based**: Prompting LMs [ACL'20]
  Horses are not.. animal; they are mammals.
  Low accuracy

Existing solutions

- **Baseline: Closed-world Assumption**
  absent as negative → elephant has no eye, is not a song
  Many wrong or uninformative negatives

- **Text-based**: Mining query logs [CIKM'19]
  why can't elephants.. → jump, use computers, hide in trees
  Low informativeness

- **Corruption-based**: Negatives from positives [EMNLP'21]
  `(horse, is a pet) → NOT(horse rider, is a pet)`
  Type-inconsistency

- **LM-based**: Prompting LMs [ACL'20]
  Horses are not.. animal; they are mammals.
  Low accuracy

Existing solutions

- **Baseline: Closed-world Assumption**
  absent as negative → elephant has no eye, is not a song
  Many wrong or uninformative negatives

- **Text-based**: Mining query logs [CIKM'19]
  why can't elephants.. → jump, use computers, hide in trees
  Low informativeness

- **Corruption-based**: Negatives from positives [EMNLP'21]
  `(horse, is a pet) → NOT(horse rider, is a pet)`
  Type-inconsistency

- **LM-based**: Prompting LMs [ACL'20]
  Horses are not.. animal; they are mammals.
  Low accuracy

Existing solutions

- **Baseline: Closed-world Assumption**
  absent as negative → elephant has no eye, is not a song
  Many wrong or uninformative negatives

- **Text-based**: Mining query logs [CIKM'19]
  why can't elephants.. → jump, use computers, hide in trees
  Low informativeness

- **Corruption-based**: Negatives from positives [EMNLP'21]
  `(horse, is a pet) → NOT(horse rider, is a pet)`
  Type-inconsistency

- **LM-based**: Prompting LMs [ACL'20]
  Horses are not.. animals; they are mammals.
  Low accuracy

Our method: Uncommonsense

Discover informative negations about target concepts by exploiting positives about comparable concepts (e.g., type siblings).

Input: target concept

`dissertation`

Output: accurate & thematical negation

`NOT(dissertation, related to book tour)`
unlike similar piece of content
e.g., `autobiography, poetry book`

Overview

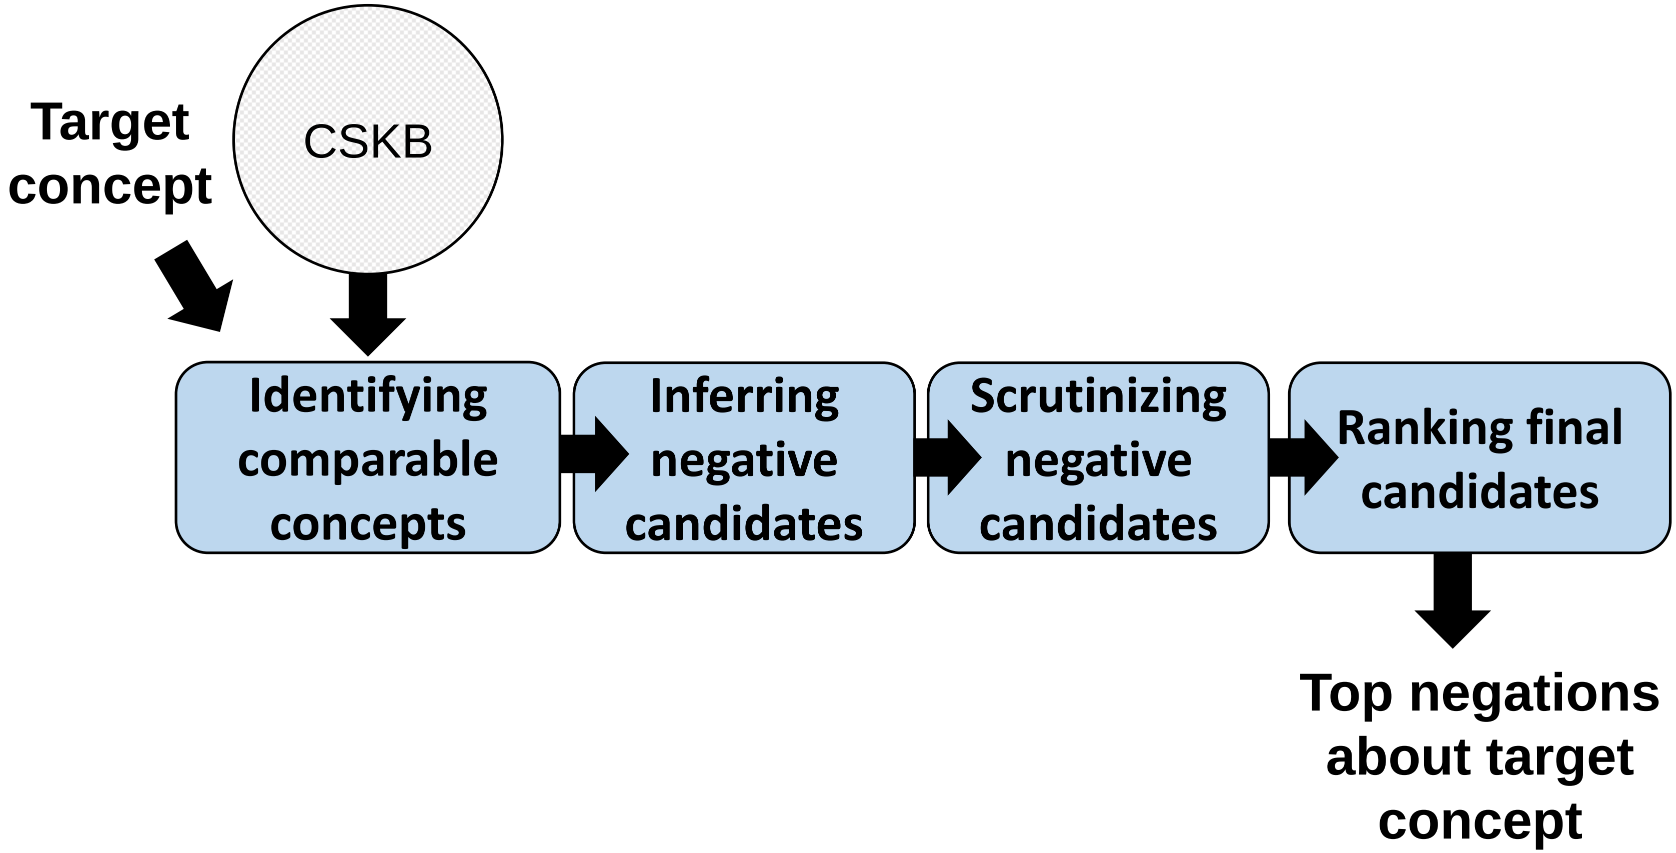
Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Overview

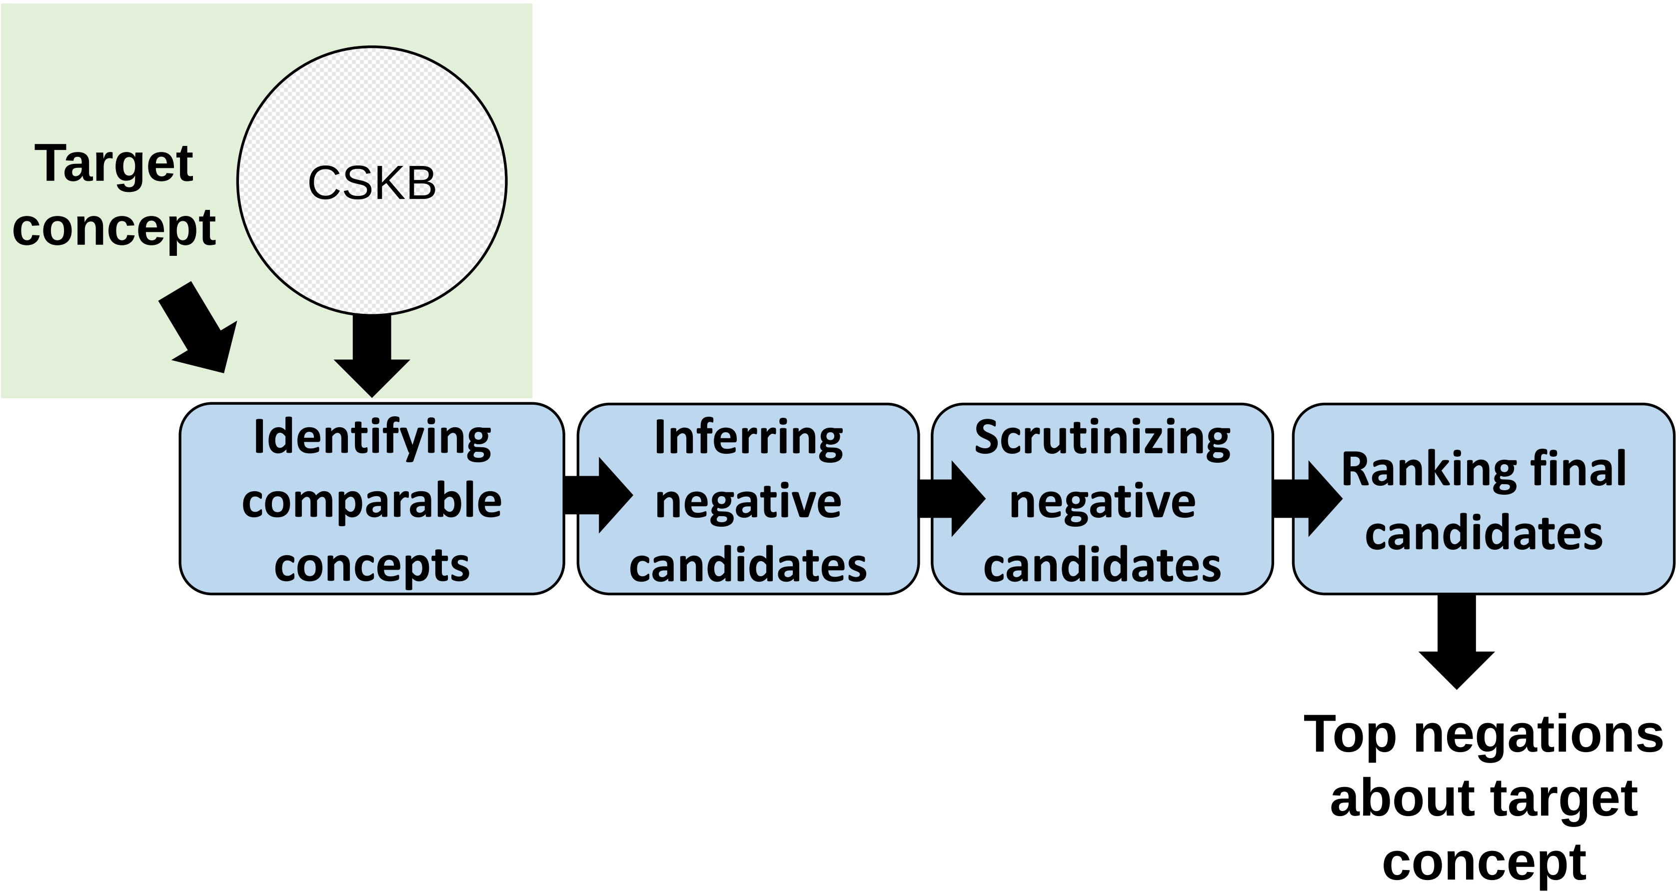

Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Overview

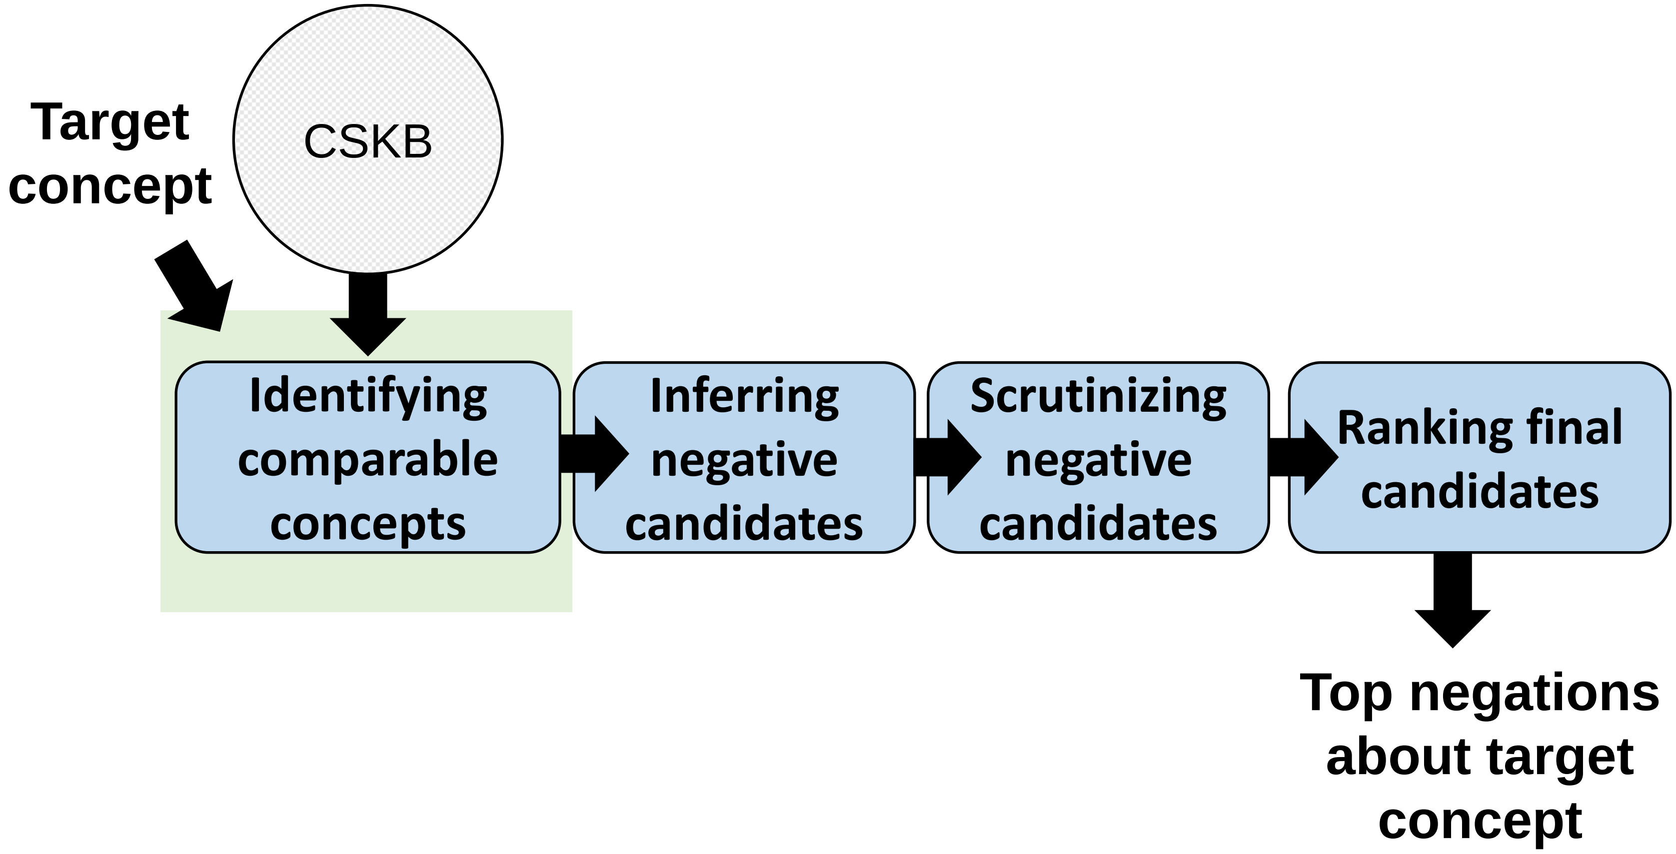
Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Overview

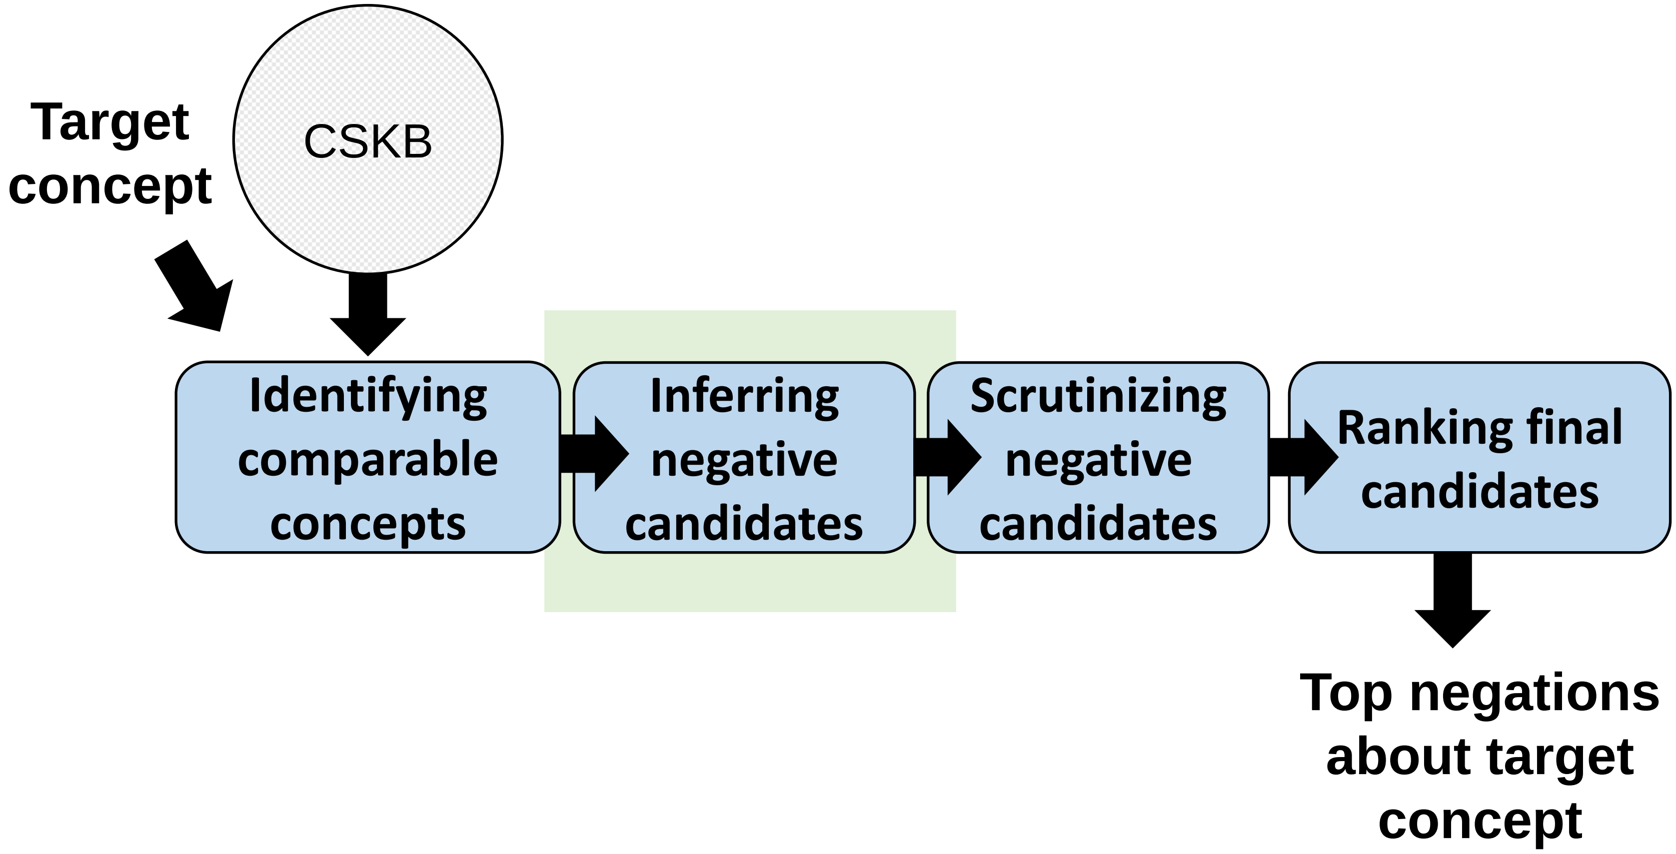
Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Overview

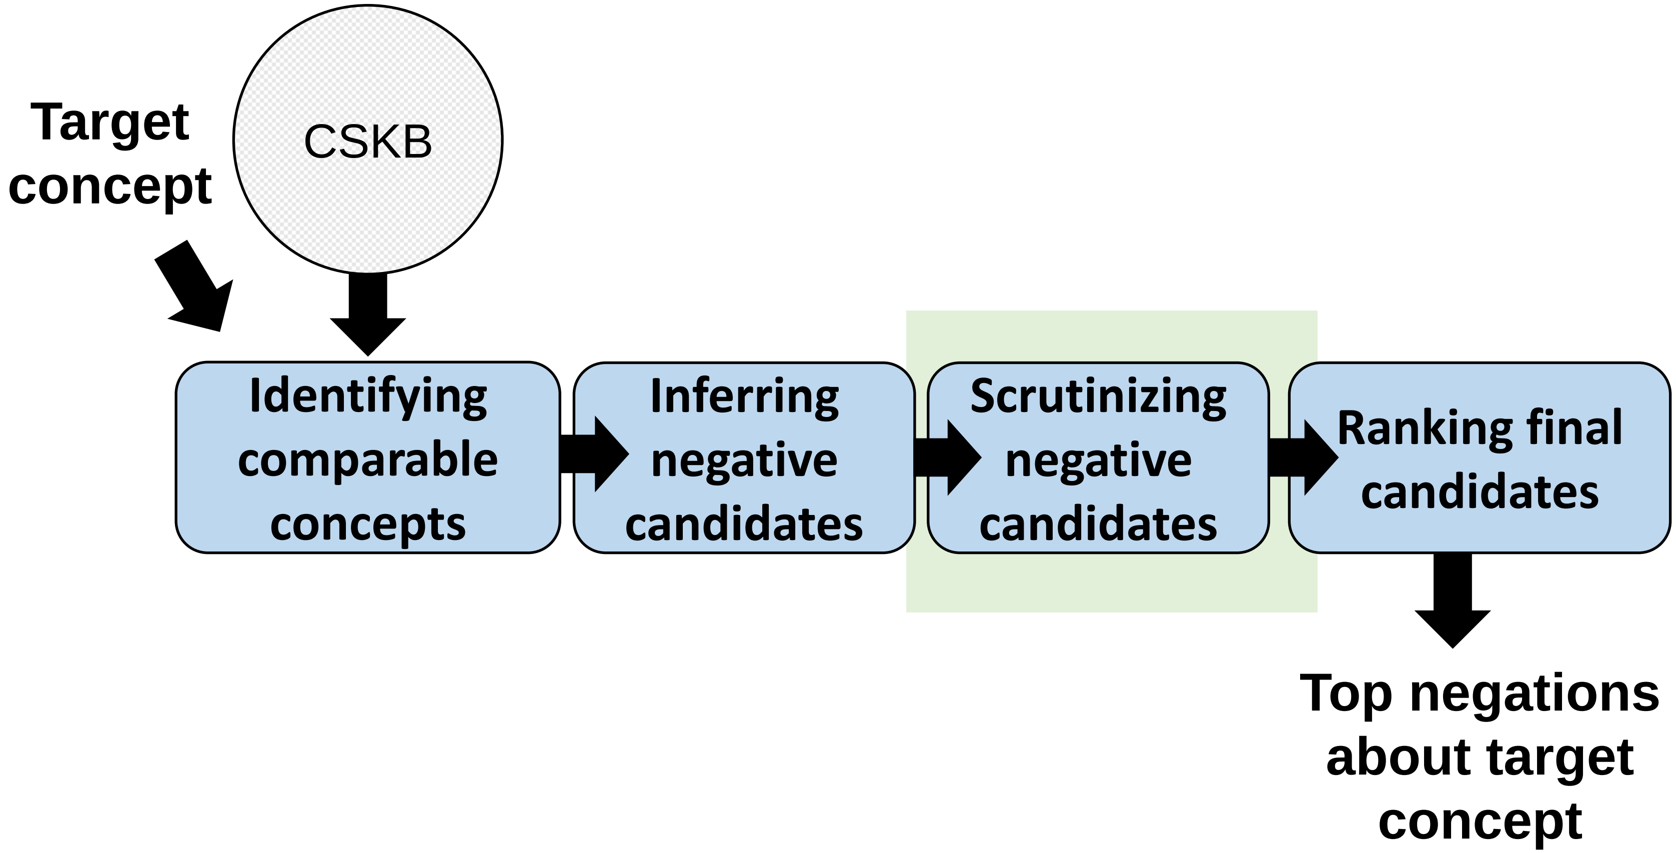

Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Overview

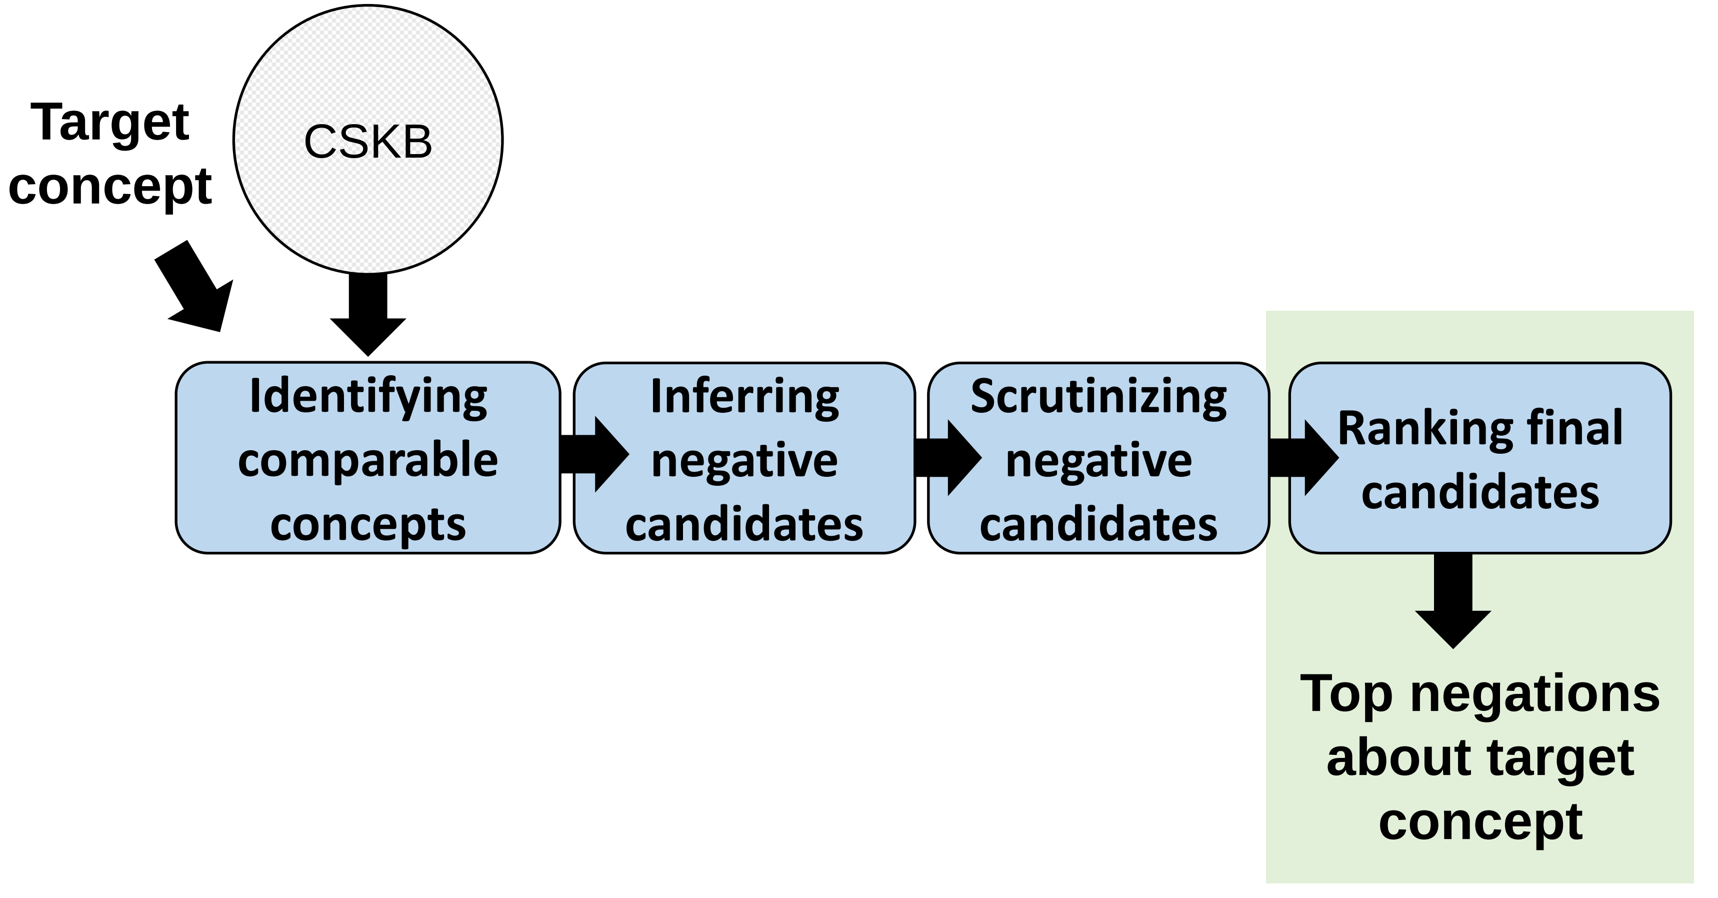

Target concept
CSKB
Identifying comparable concepts
Inferring negative candidates
Scrutinizing negative candidates
Ranking final candidates
Top negations about target concept

Step1: identify comparable concepts

Target concept: `dissertation`

We need a similarity measure to collect *type siblings*:

- **Hypernymy relations**
  other piece of content
  `ad, article, podcast, poetry book..`
  Lacks ranking

- **Entity embeddings**
  compute closest neighbors cosine similarity
  `article, university, poetry book, scholar..`
  Lacks type consistency

A combination: <u>rank</u> using embeddings and <u>retrain</u> co-hyponyms: `article, poetry book, ..`

[Hertling et al., WebIsALOD: Providing Hypernymy Relations extracted from the Web as LOD ISWC'17]
[Yamada et al., Wikipedia2Vec: Toolkit for Learning .. Embeddings of Words and Entities .., EMNLP'20]

Step2: infer candidate negations

CSKB

Target:
`dissertation`

(dissertation, academic writing)
(dissertation, is evaluated)

Siblings:
`poetry book, article, autobiography, comic strip`

(autobiography|article|poetry book, in bookstore)
(autobiography|poetry book, is personal)
(autobiography, has chapters)
(article, is reviewed)

MINUS

NOT(dissertation, in bookstore)
NOT(dissertation, is personal)
NOT(dissertation, has chapters)
NOT(dissertation, is reviewed)

Step3: scrutinize candidates

Target:
`dissertation`

```
NOT(in bookstore)
NOT(is personal)
NOT(has chapters)
NOT(is reviewed)
```

We need plausibility checks to eliminate *false negatives*:

1. **Internal**
   compute semantic similarity between candidates and what the KB *already* knows
   similarity (`is reviewed`, `is evaluated`) > thresh.

2. **External**
   in case of missing positives, LM as source of knowledge (zero-shot, masking subject of phrase)
   `[MASK] has chapters -> dissertation`

[Reimers et al., Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, EMNLP'19]

Step3: scrutinize candidates

Target: dissertation

`NOT(in bookstore)`
`NOT(is personal)`
~~`NOT(has chapters)`~~
~~`NOT(is reviewed)`~~

We need plausibility checks to eliminate *false negatives*:

1. **Internal**
   compute semantic similarity between candidates and what the KB *already* knows
   similarity (`is reviewed`, `is evaluated`) > thresh.

2. **External**
   in case of missing positives, LM as source of knowledge (zero-shot, masking subject of phrase)
   `[MASK] has chapters -> dissertation`

[Reimers et al., Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, EMNLP’19]

Step4: quantify informativeness

Target:
`dissertation`

`NOT(in bookstore)`
`NOT(is personal)`

Real example:
thousands!

Simple statistical measure: relative frequency

How <u>common</u> is this phrase <u>amongst</u> type siblings?

score(`dissertation, NOT(in bookstore)`) = ¾

[`negation`] , [sc]
[provenance]

Output

`NOT(dissertation, found in bookstore)`, 0.75
Unlike similar <u>piece of content</u>
such as `poetry book`, `article`, & `autobiography`

Intrinsic evaluation: setup

- **CSKB**

  Ascent++ (total = 2 million statements)

  200 target concepts (*randomly sampled*)
  `elephant, acne, marriage proposal,..`

- **Metrics**

  Plausibility (inverse: false negatives)
  crowd (*is the statement truly negative?*)

  Informativeness
  crowd (*is the statement interesting?*)

  Recall
  ground-truth (Conceptnet's 14k negated statements)

Intrinsic evaluation: baselines

1. **Closed-world baseline**
   absent = negative, *no ranking*

2. **GPT-3^{neg}**
   prompt with negative keywords (zero-shot), *ranking* = *probabilities*

3. **Quasimodo^{neg}**
   Quasimodo KB's negations, *ranking* = *confidence*

4. **NegatER**
   corruption-based, *ranking* = *LM's classification scores*

Intrinsic evaluation: results

Method	False Negatives	Informative-ness	Negations for `elephant`
CW-baseline	**0.07**	0.07	`NOT(perform comedy)`
Quasimodo^{neg}	0.61	0.32	`NOT(survive)`
GPT-3^{neg}	0.63	0.30	`NOT(in the zoo)`
NegatER	0.26	0.29	`NOT(interested)`
Uncommonsense	0.25	**0.50**	`NOT(carnivore)`

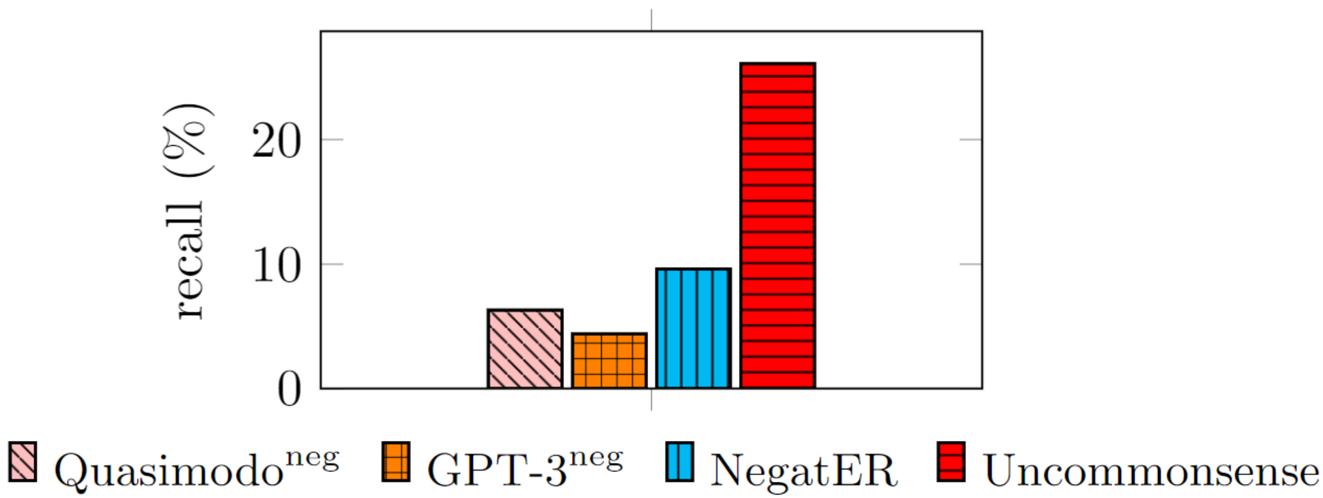

Intrinsic evaluation: informativeness per domain

Method	**Animal**	**Food**	**Activity**	**Social**	**Object**	**Other**
Quasimodo^{neg}	0.41	0.44	0.24	**0.39**	0.20	0.24
GPT-3^{neg}	0.14	0.46	0.44	0.17	0.22	0.23
NegatER	0.13	0.14	0.17	0.26	0.15	0.18
Uncommonsense	**0.67**	**0.55**	**0.49**	**0.39**	**0.42**	**0.45**
Sample concept	`lynx`	`cake`	`yoga`	`wedding`	`tripod`	`belief`

Inside uncommonsense: ablation

Configuration	**Plausibility**	**Informativeness**
w/o comparable concepts	0.81	0.26 (-24%)
w/o plausibility checks	0.51 (-24%)	0.38
w/o ranking	0.61	0.29 (-21%)
Complete	0.75	0.50

Use case: negative trivia

- Crowdsourcing task
- Provenance-extended negations **+32% in informativeness**

Target concept	**Negation**	**Explanation**
`gorilla`	`NOT(territorial)`	unlike similar wild animal such as (`tiger`, `lion`, ..)
`vinegar`	`NOT(has iron)`	unlike similar acidic food such as (`tomato`, `lemon`, ..)

Use case: KB completion

- Task: identifying *unseen* positive statements.
- Method: fine-tune LM as classifier of candidates (True/False), LM=BERT, CSKB=ConceptNet.
- Crucial ingredient: negative samples for training the model.
- Instead of *trivial* ..
  Use ***strong*** negatives, generated by uncommonsense.

  +4% in accuracy

Code for fine-tuning BERT: https://github.com/tsafavi/NegatER

Use case: multiple-choice QA

- Dataset: CommonsenseQA.
- Automatic elimination of ***unlikely*** answers.
  +18% in helpful eliminations
  (i.e., removal of an incorrect answer)

Which of the following is a territorial animal?

A. Lion
~~B. Elephant~~
~~C. Gorilla~~
~~D. Bear~~

Uncommonsense

[Talmor et al., A Question Answering Challenge Targeting Commonsense Knowledge, NAACL'19]

Summary

- Take away message
  - Current KBs lack negative knowledge
  - Explicit negation increases useability, e.g., QA
  - Statistical inference methods to generate accurate & sal negations.
- Challenges/Opportunities
  - Plausibility-informativeness tradeoff
  - Maintenance, e.g., real-world changes
  - Long tail entities, e.g. less famous people
  - Universal social negations (new MPI project: Cultural CSK)

Thank you!
Questions?

- Acknowledgments 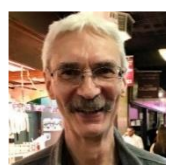 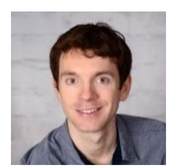  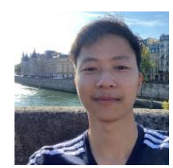

Contact: hibaarnaout.com or SCAN ME