Identifying civilians killed by police with distantly supervised entity-event extraction

Katherine A. Keith, Abram Handler, Michael Pinkham, Cara Magliozzi, Joshua McDuffie, and Brendan O'Connor

EMNLP 2017

College of Information and Computer Science
University of Massachusetts Amherst

Killings by police in the U.S.

July 17, 2014

Aug 9, 2014

July 5, 2016

July 6, 2016

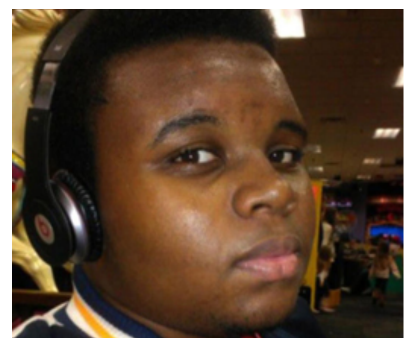

Eric Garner	New York, NY
Michael Brown	Ferguson, MO
Alton Sterling	Baton Rouge, LA
Philando Castile	Falcon Heights, MN

Data needed for policy making

Data needed for policy making

- Fatality Statistics?

Data needed for policy making

- Fatality Statistics?
- Racial disparity / discrimination?

Data needed for policy making

- Fatality Statistics?
- Racial disparity/discrimination?
- Most effective police departments/policing methods?

Data needed for policy making

- Fatality Statistics?
- Racial disparity / discrimination?
- Most effective police departments / policing methods?

DATA!

Issues in government data

Issues in government data

[Banks et al. 2015 (BJS/DOJ)]

Issues in government data

Number of U.S. police killings 2003-2009, 2011

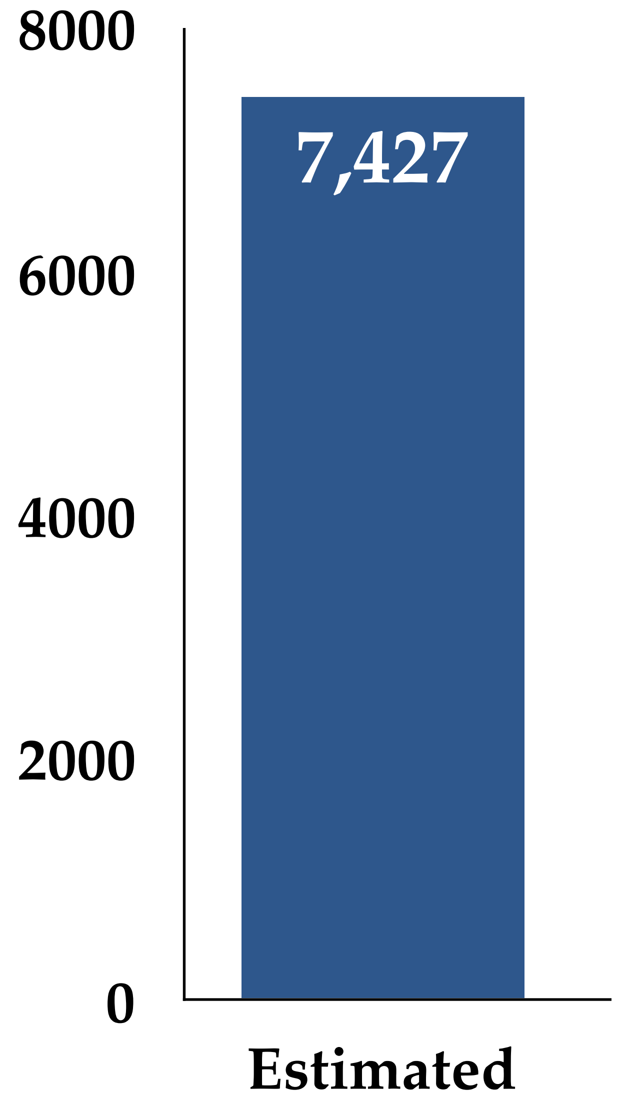

[Banks et al. 2015 (BJS/DOJ)]

Issues in government data

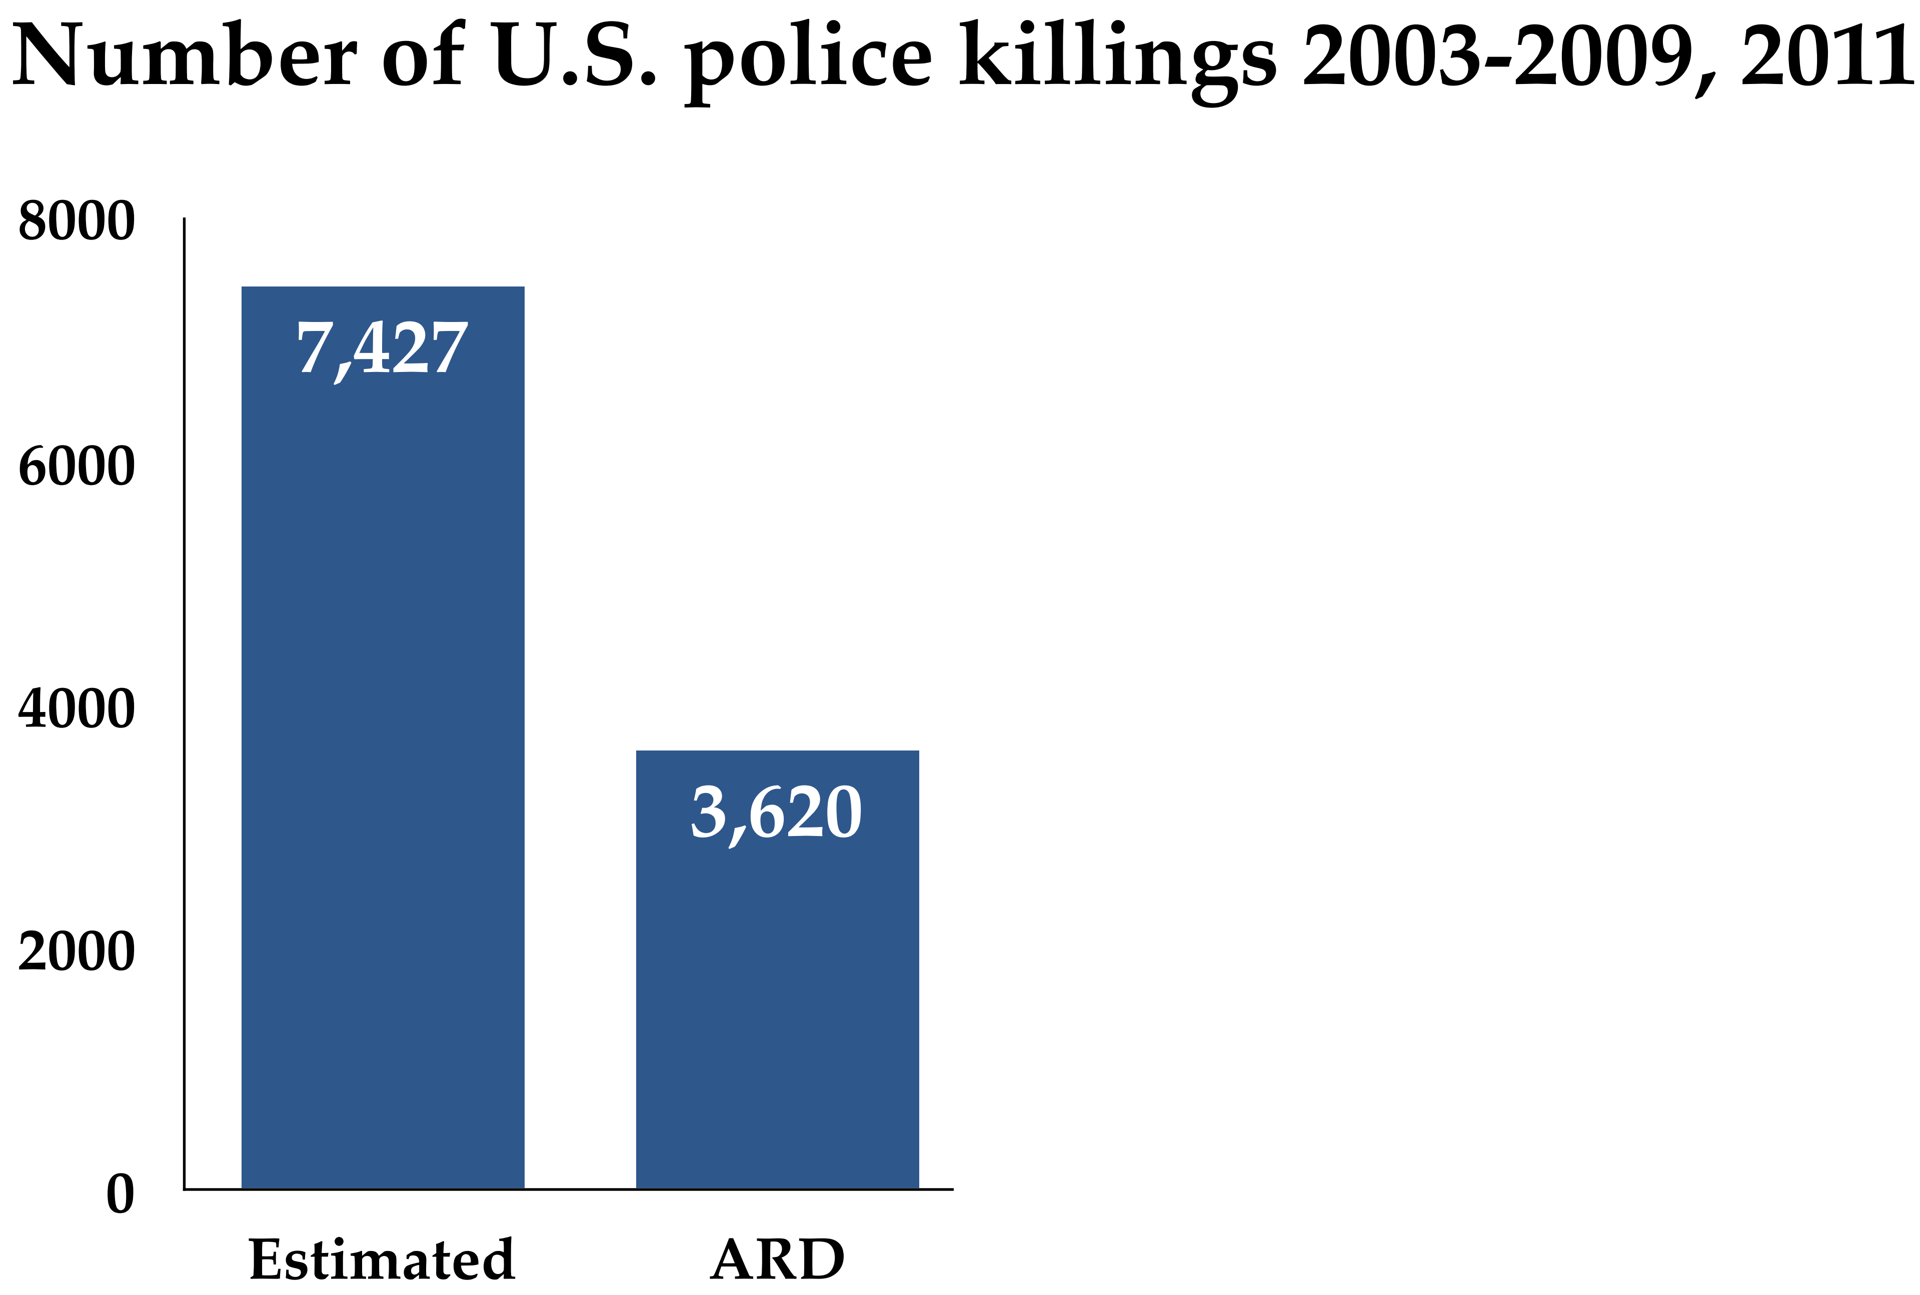

[Banks et al. 2015 (BJS/DOJ)]

Issues in government data

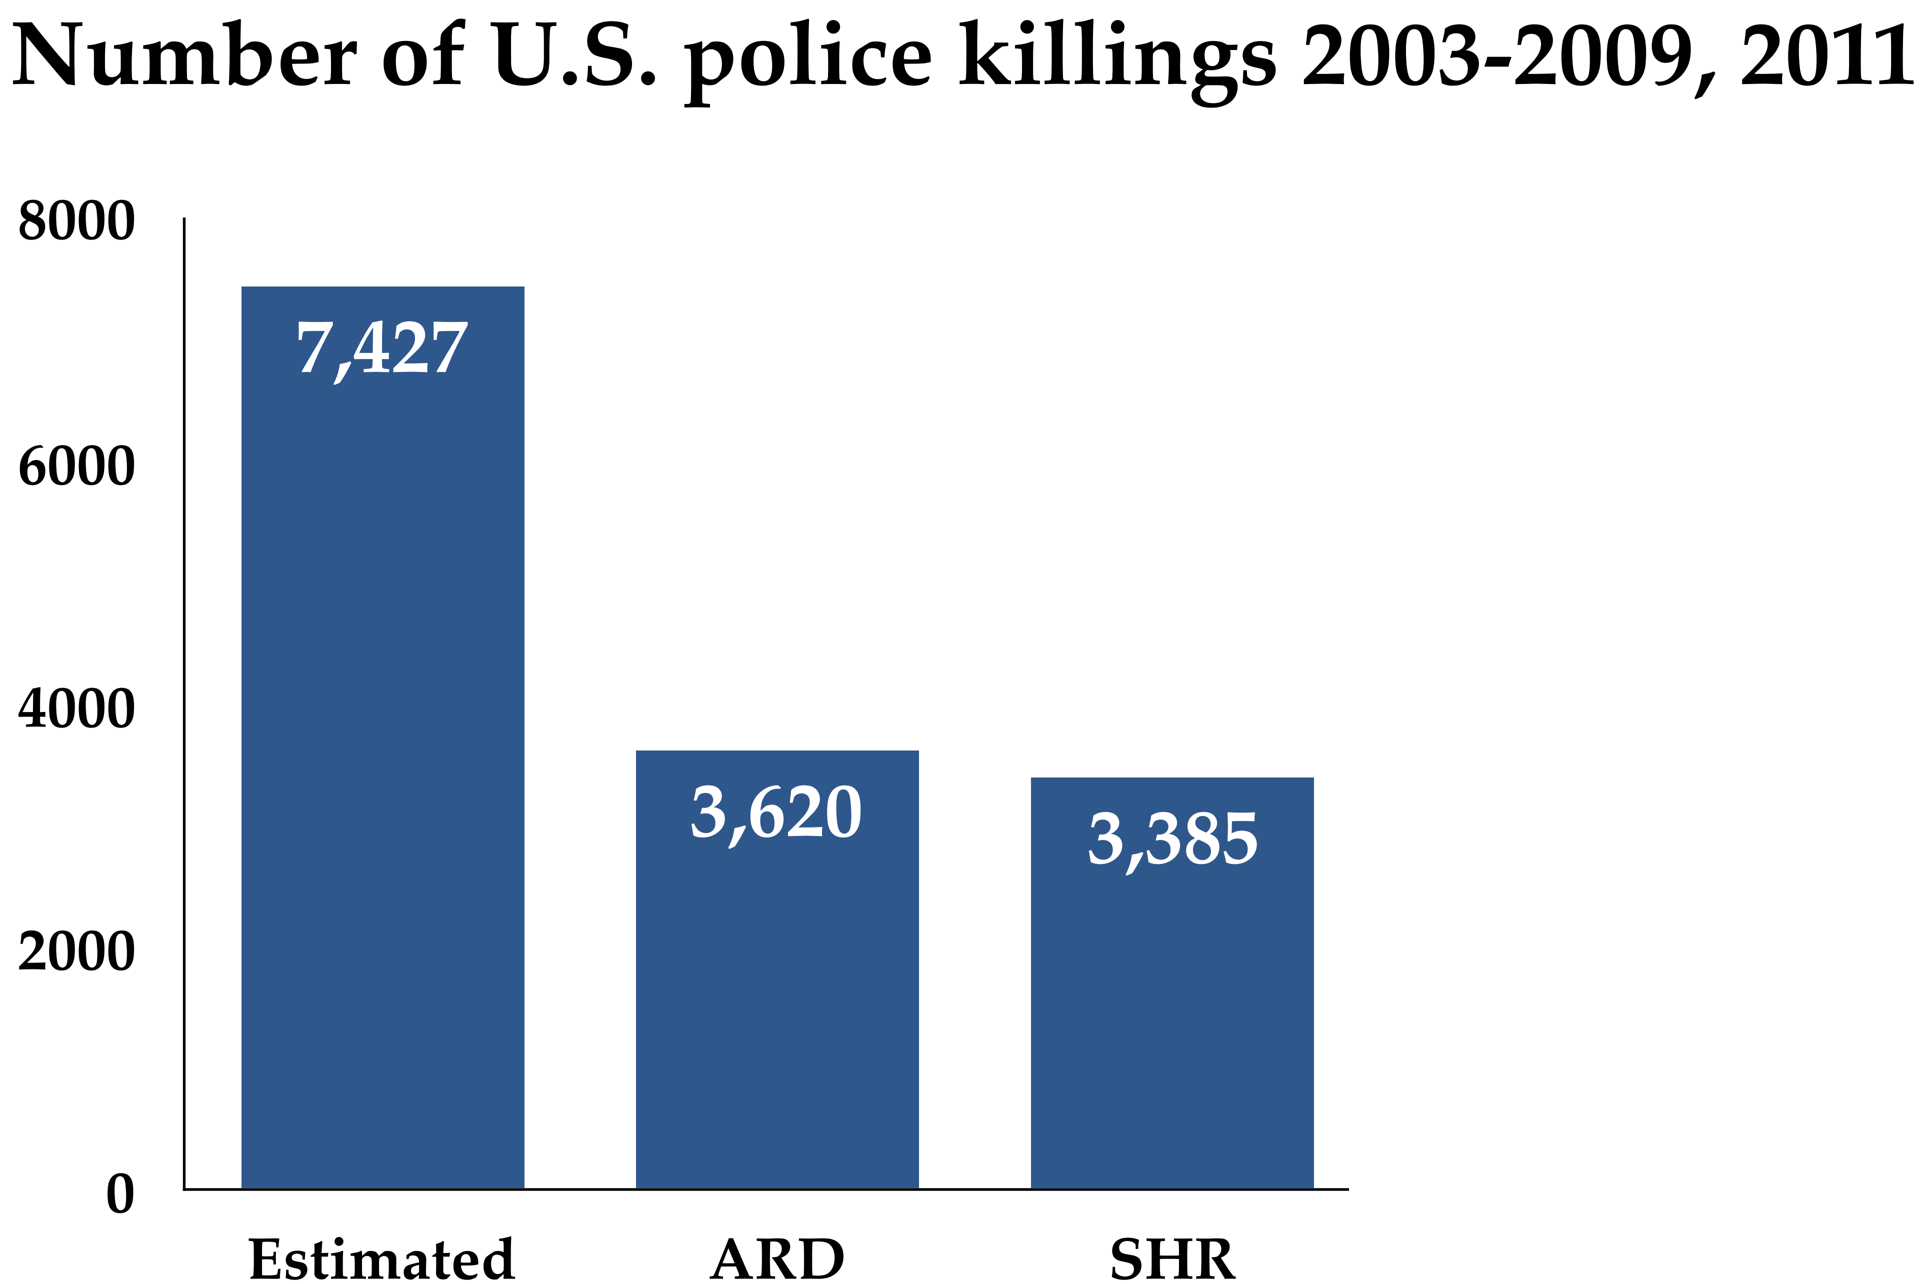

[Banks et al. 2015 (BJS/DOJ)]

Issues in government data

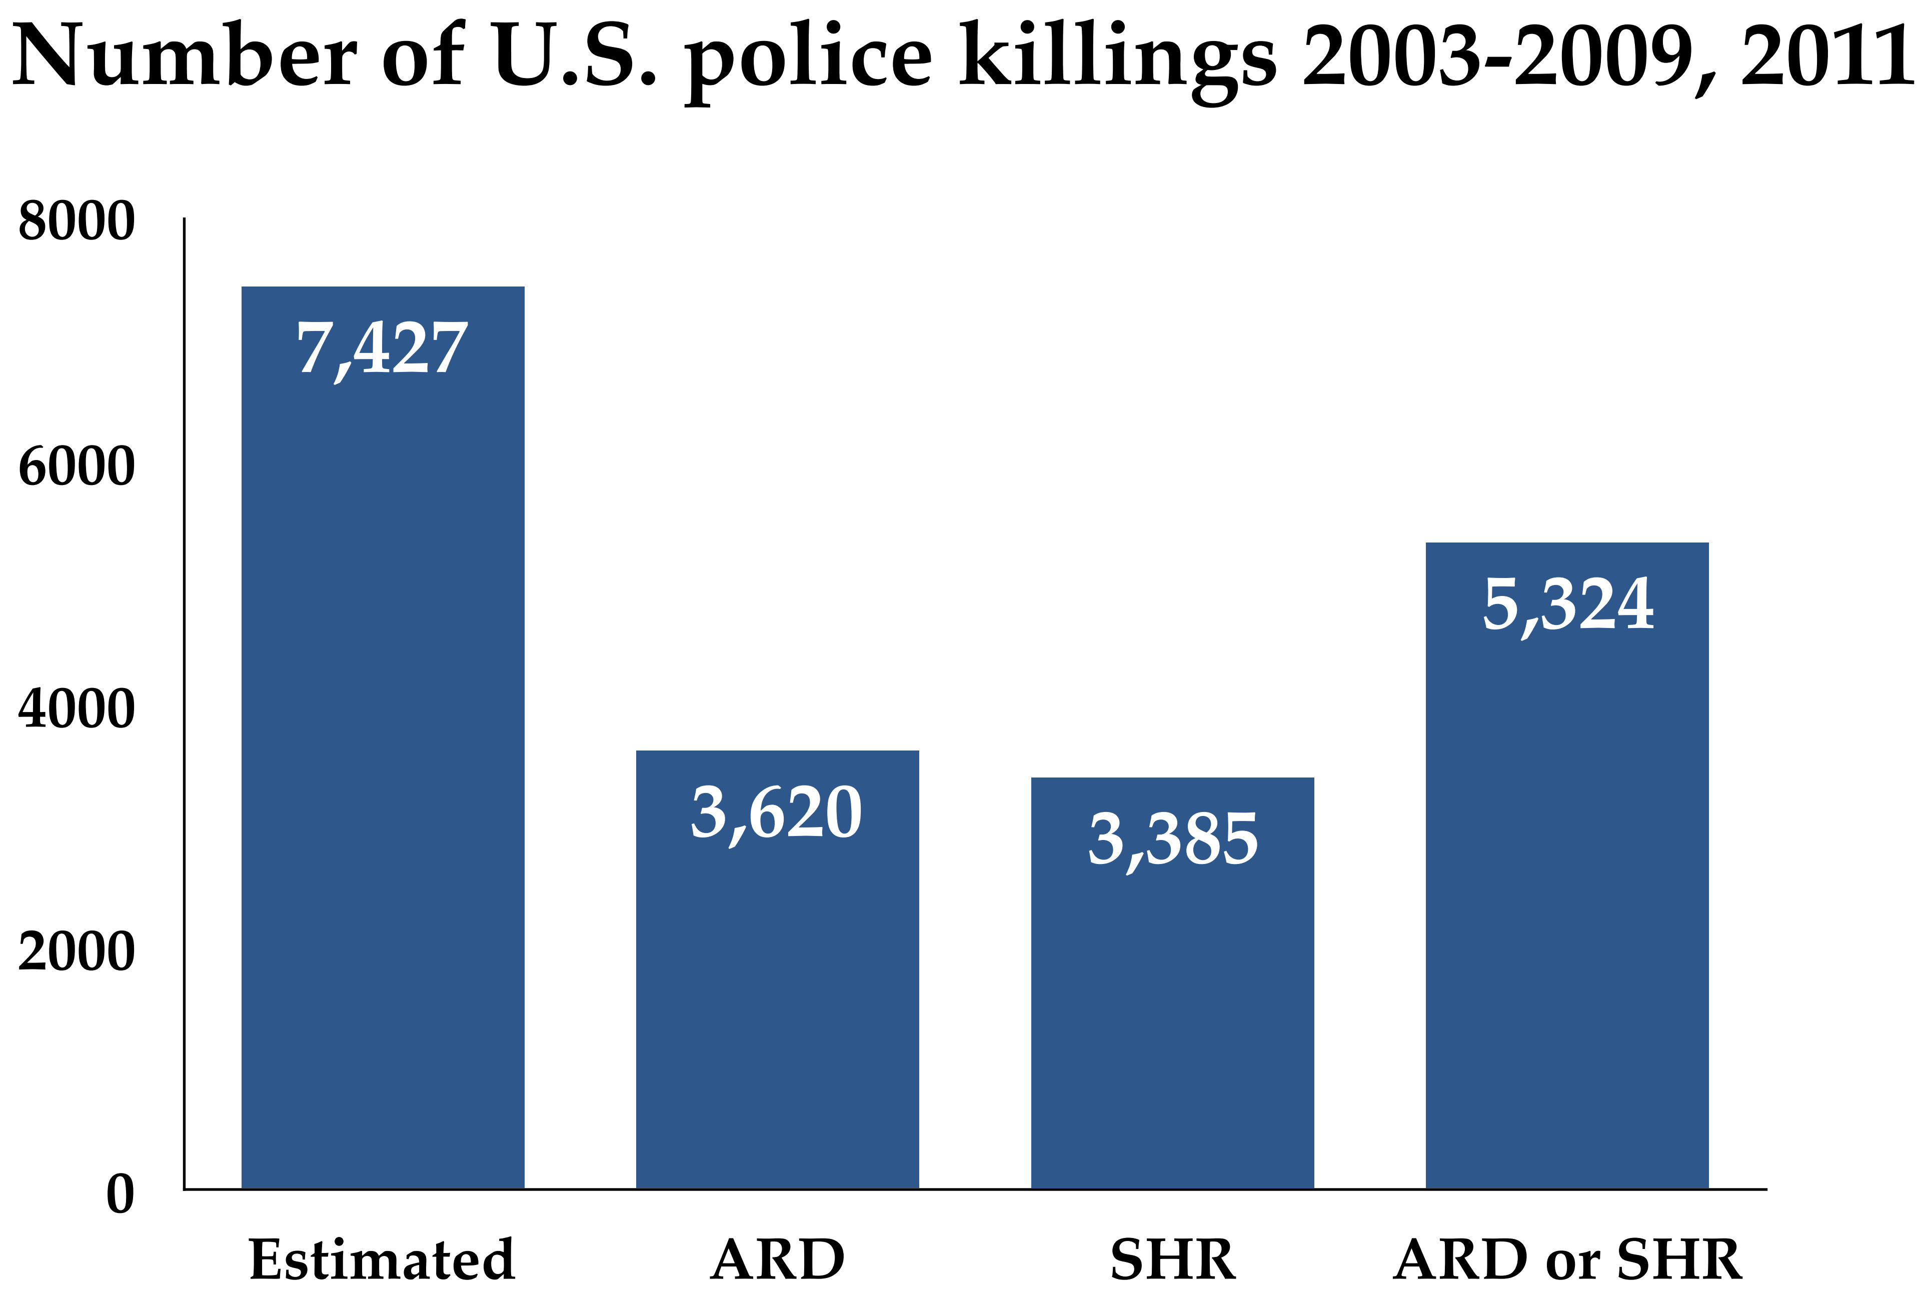

[Banks et al. 2015 (BJS/DOJ)]

Issues in government data

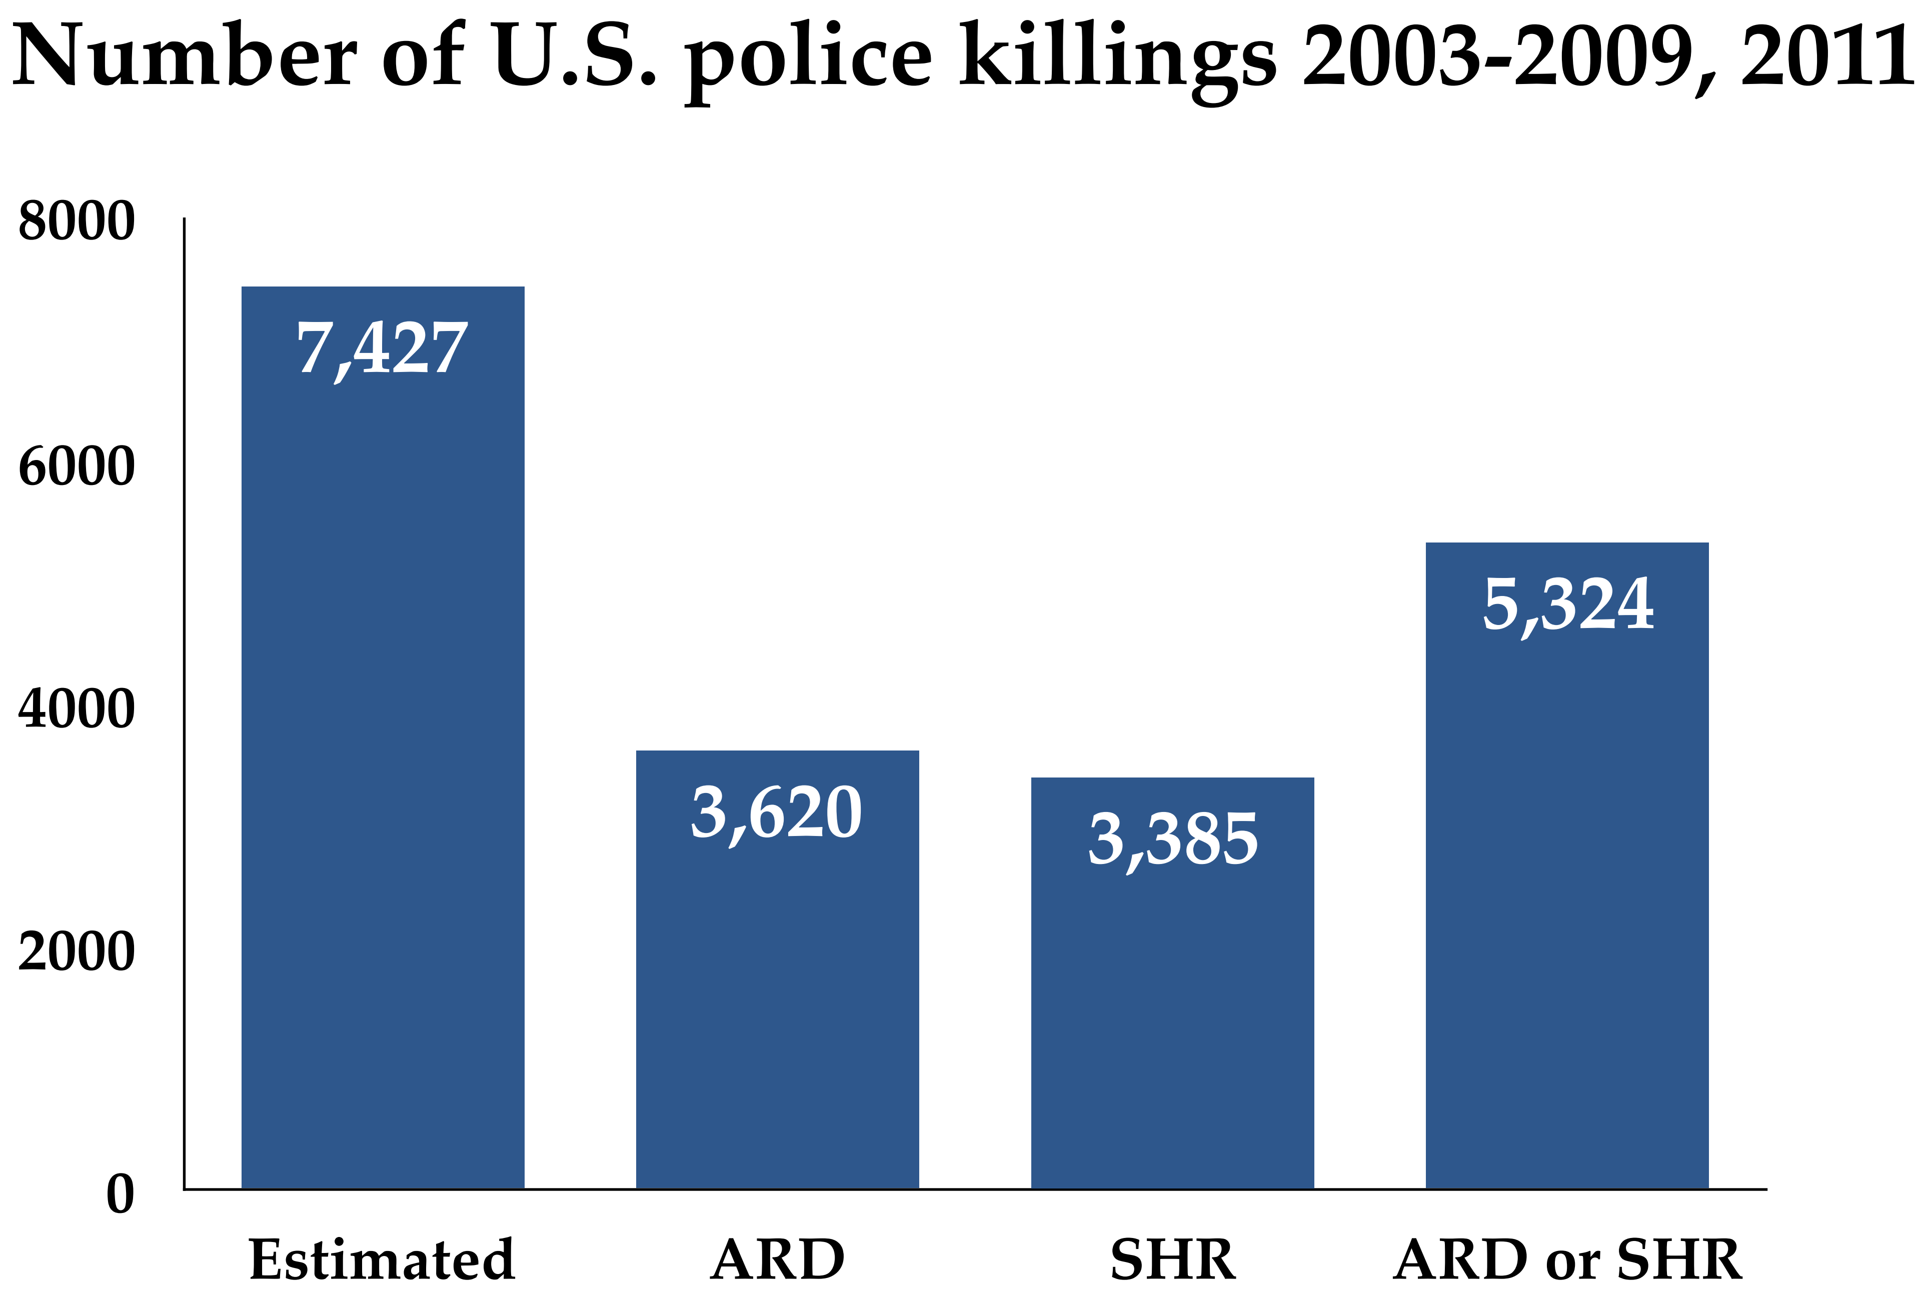

[Banks et al. 2015 (BJS/DOJ)]

Alternative data: media reports

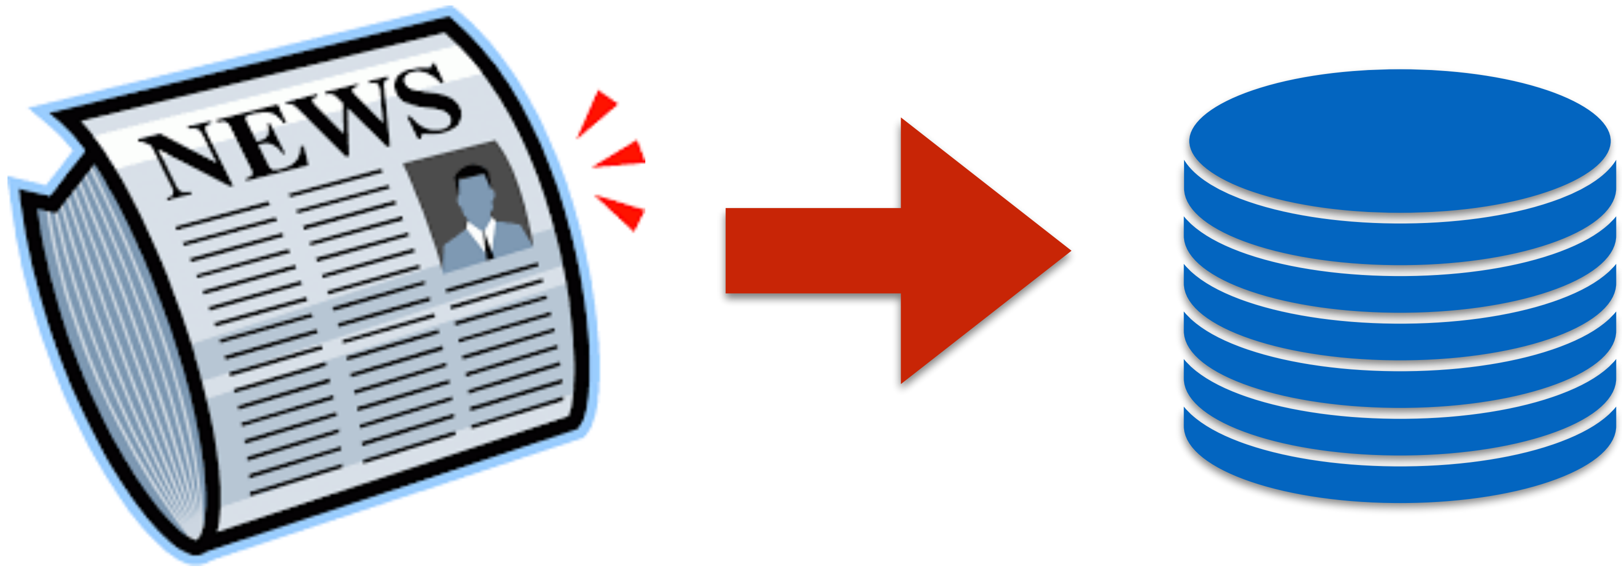

- Populate an **entity-event database** by manually reading news articles

Alternative data: media reports

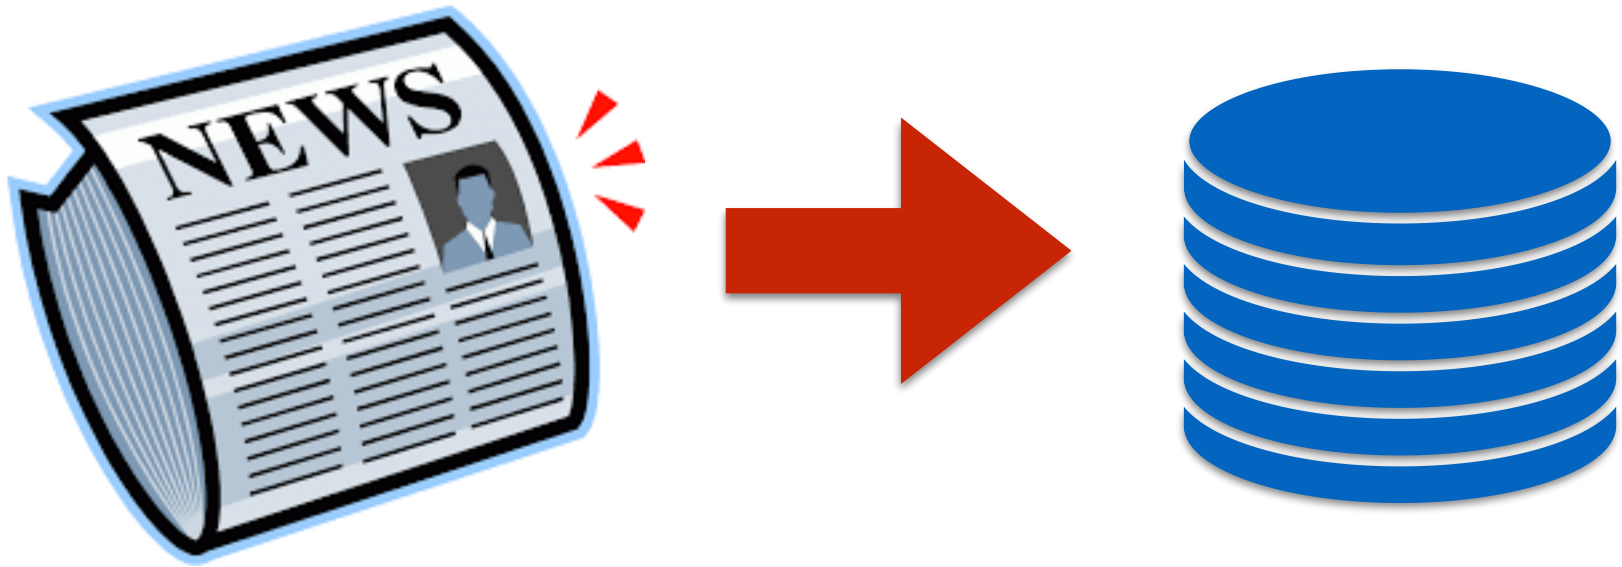

- Populate an **entity-event database** by manually reading news articles
- FatalEncounters.org, KilledByPolice.net, The Guardian, Washington Post…

Alternative data: media reports

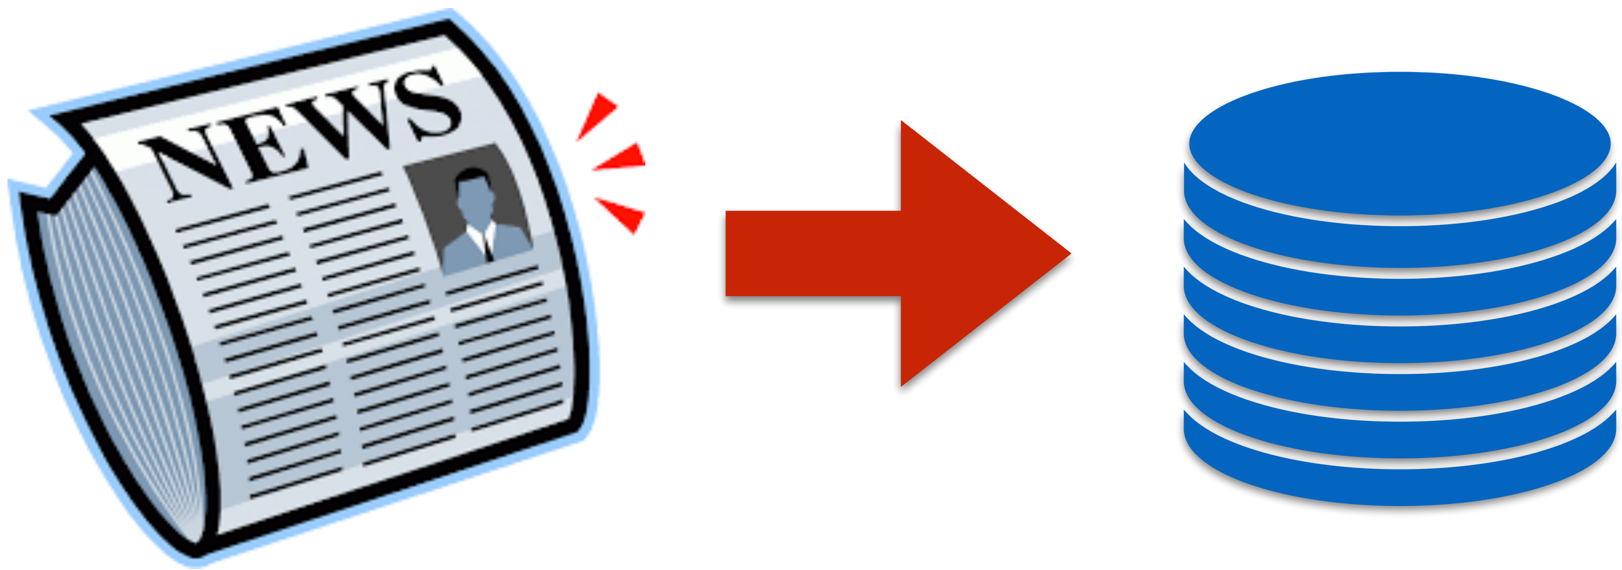

- Populate an **entity-event database** by manually reading news articles
- FatalEncounters.org, KilledByPolice.net, The Guardian, Washington Post…
- Fatal Encounters volunteers have read >2 million articles

Number of U.S. police killings 2003-2009, 2011

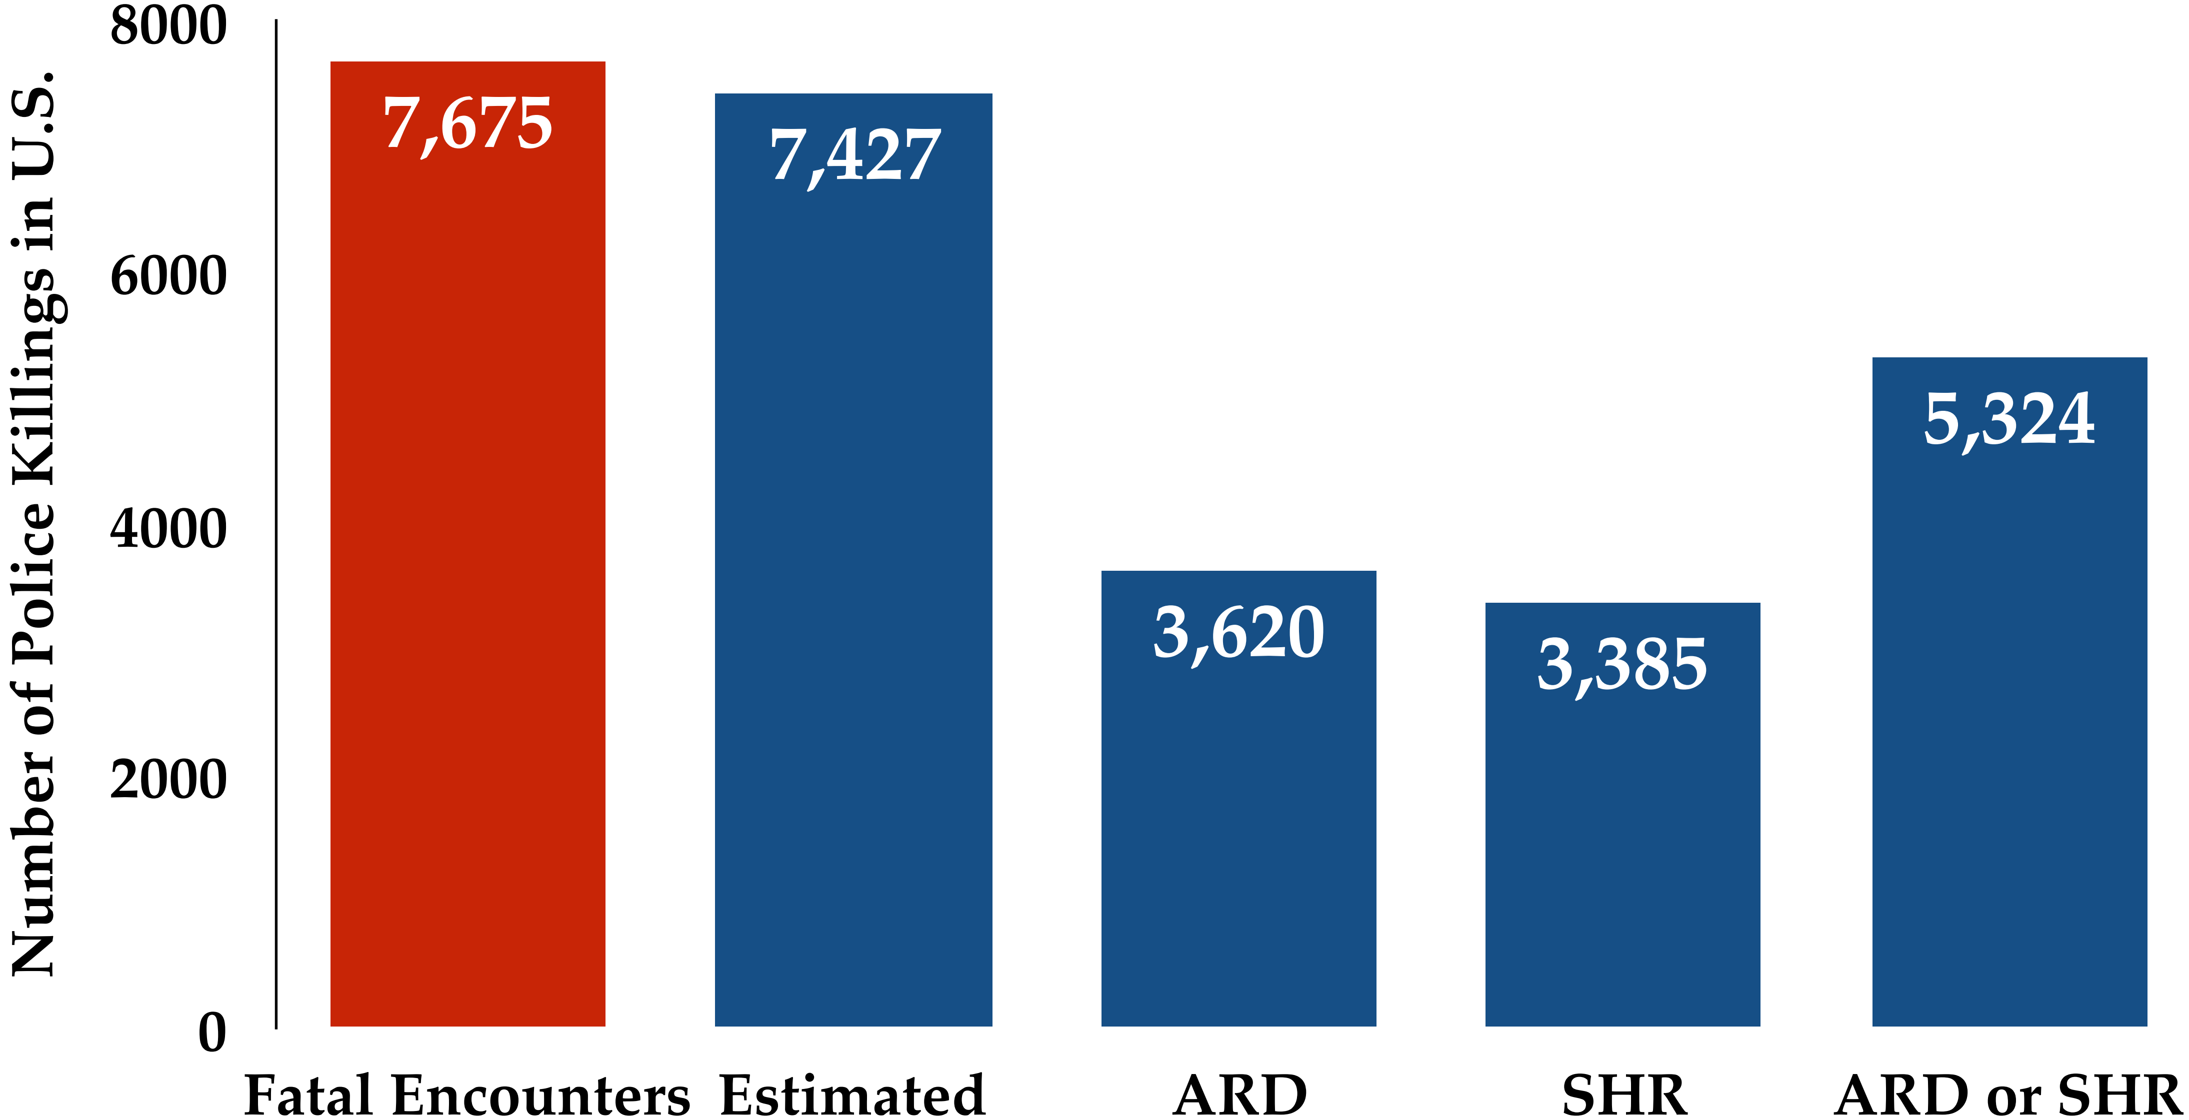

[Banks et al. 2015 (BJS/DOJ)]

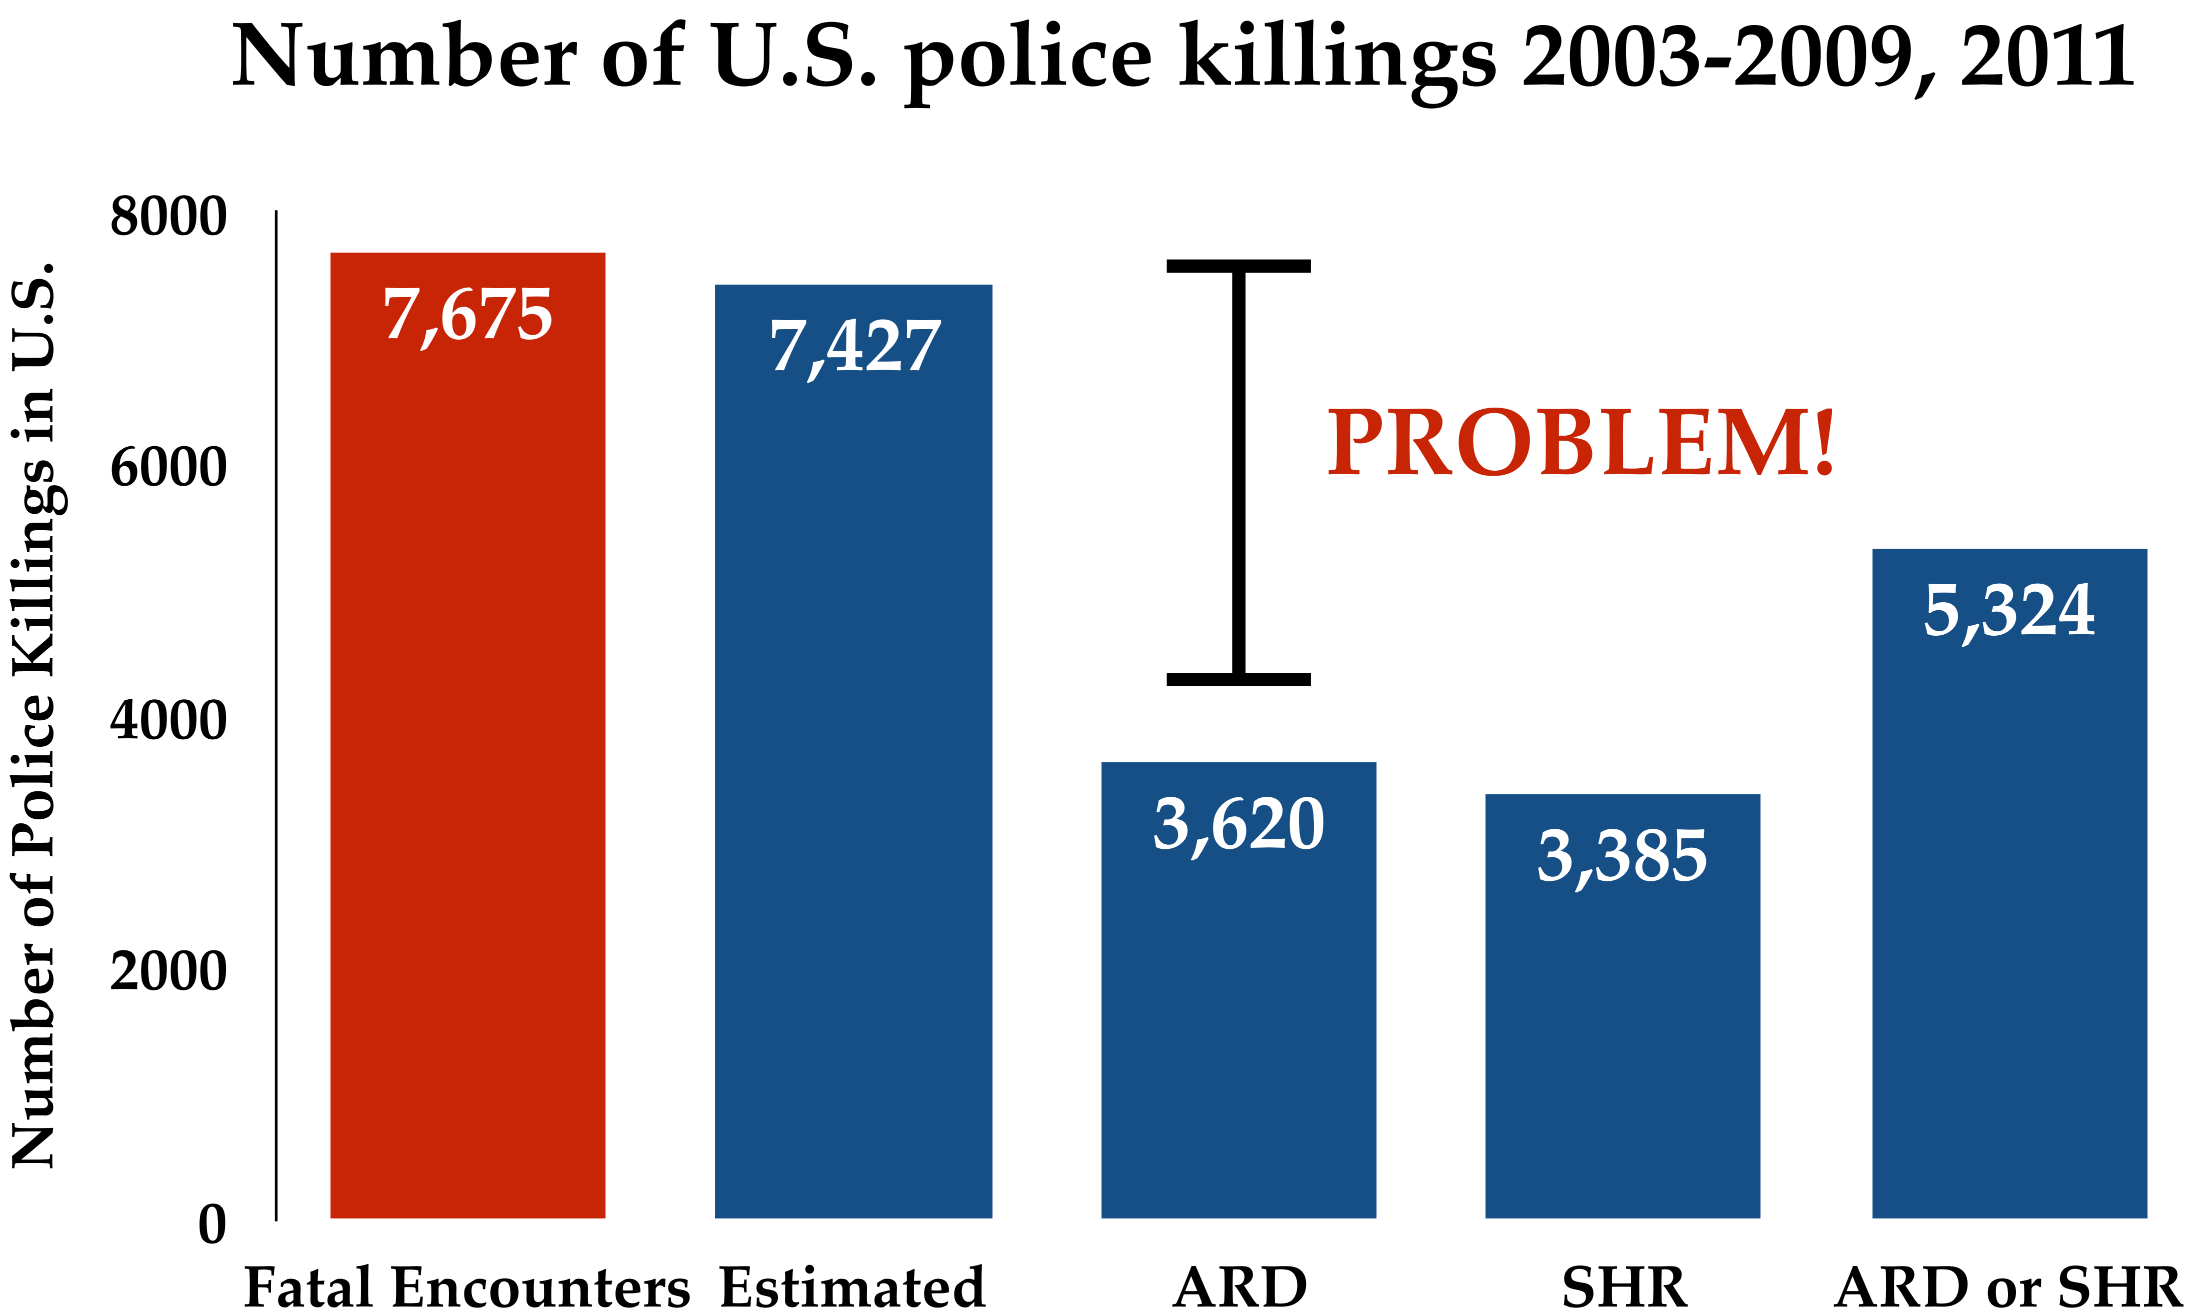

[Banks et al. 2015 (BJS/DOJ)]

Overview

Motivation:

Public data and government accountability

Overview

Motivation:

Public data and government accountability

Problems with existing approaches:

1. Manual updates are expensive
2. Continuous updates required

Overview

Motivation:

Public data and government accountability

Problems with existing approaches:

1. Manual updates are expensive
2. Continuous updates required

Goal:

Automatically update a police fatality database

Overview

Overview

sentence w/ entity

sentence w/ entity

sentence w/ entity

sentence w/ entity

Overview

sentence w/ entity 0

sentence w/ entity 1

sentence w/ entity 1

sentence w/ entity 0

Overview

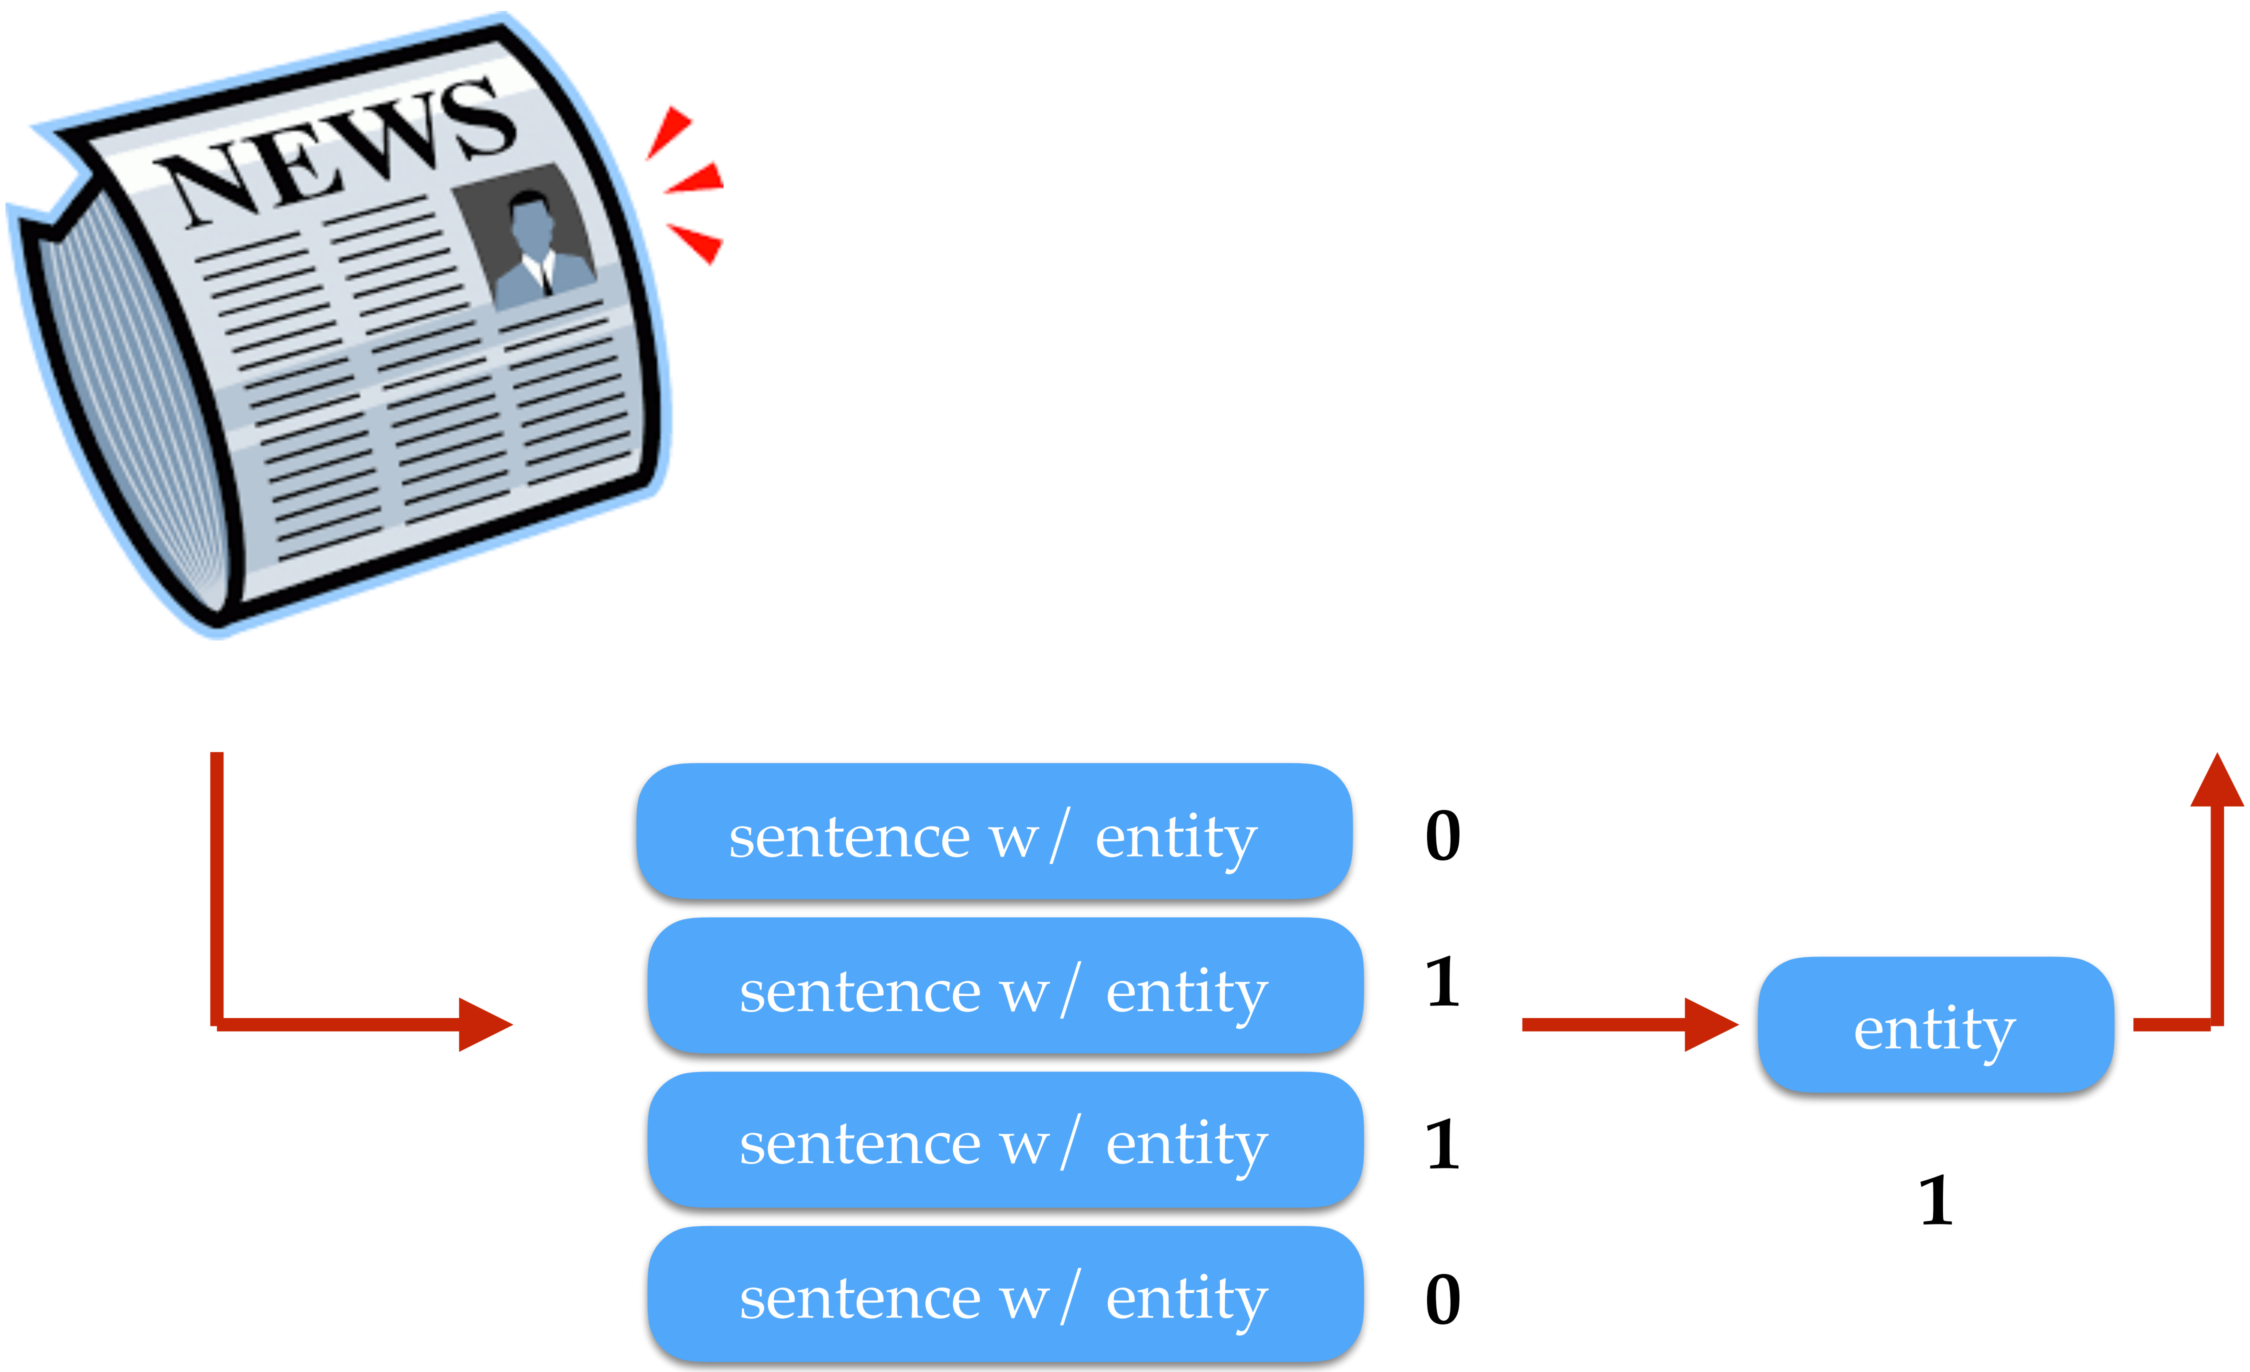

NEWS
sentence w/ entity
0
sentence w/ entity
1
sentence w/ entity
1
sentence w/ entity
0
entity
1

Overview

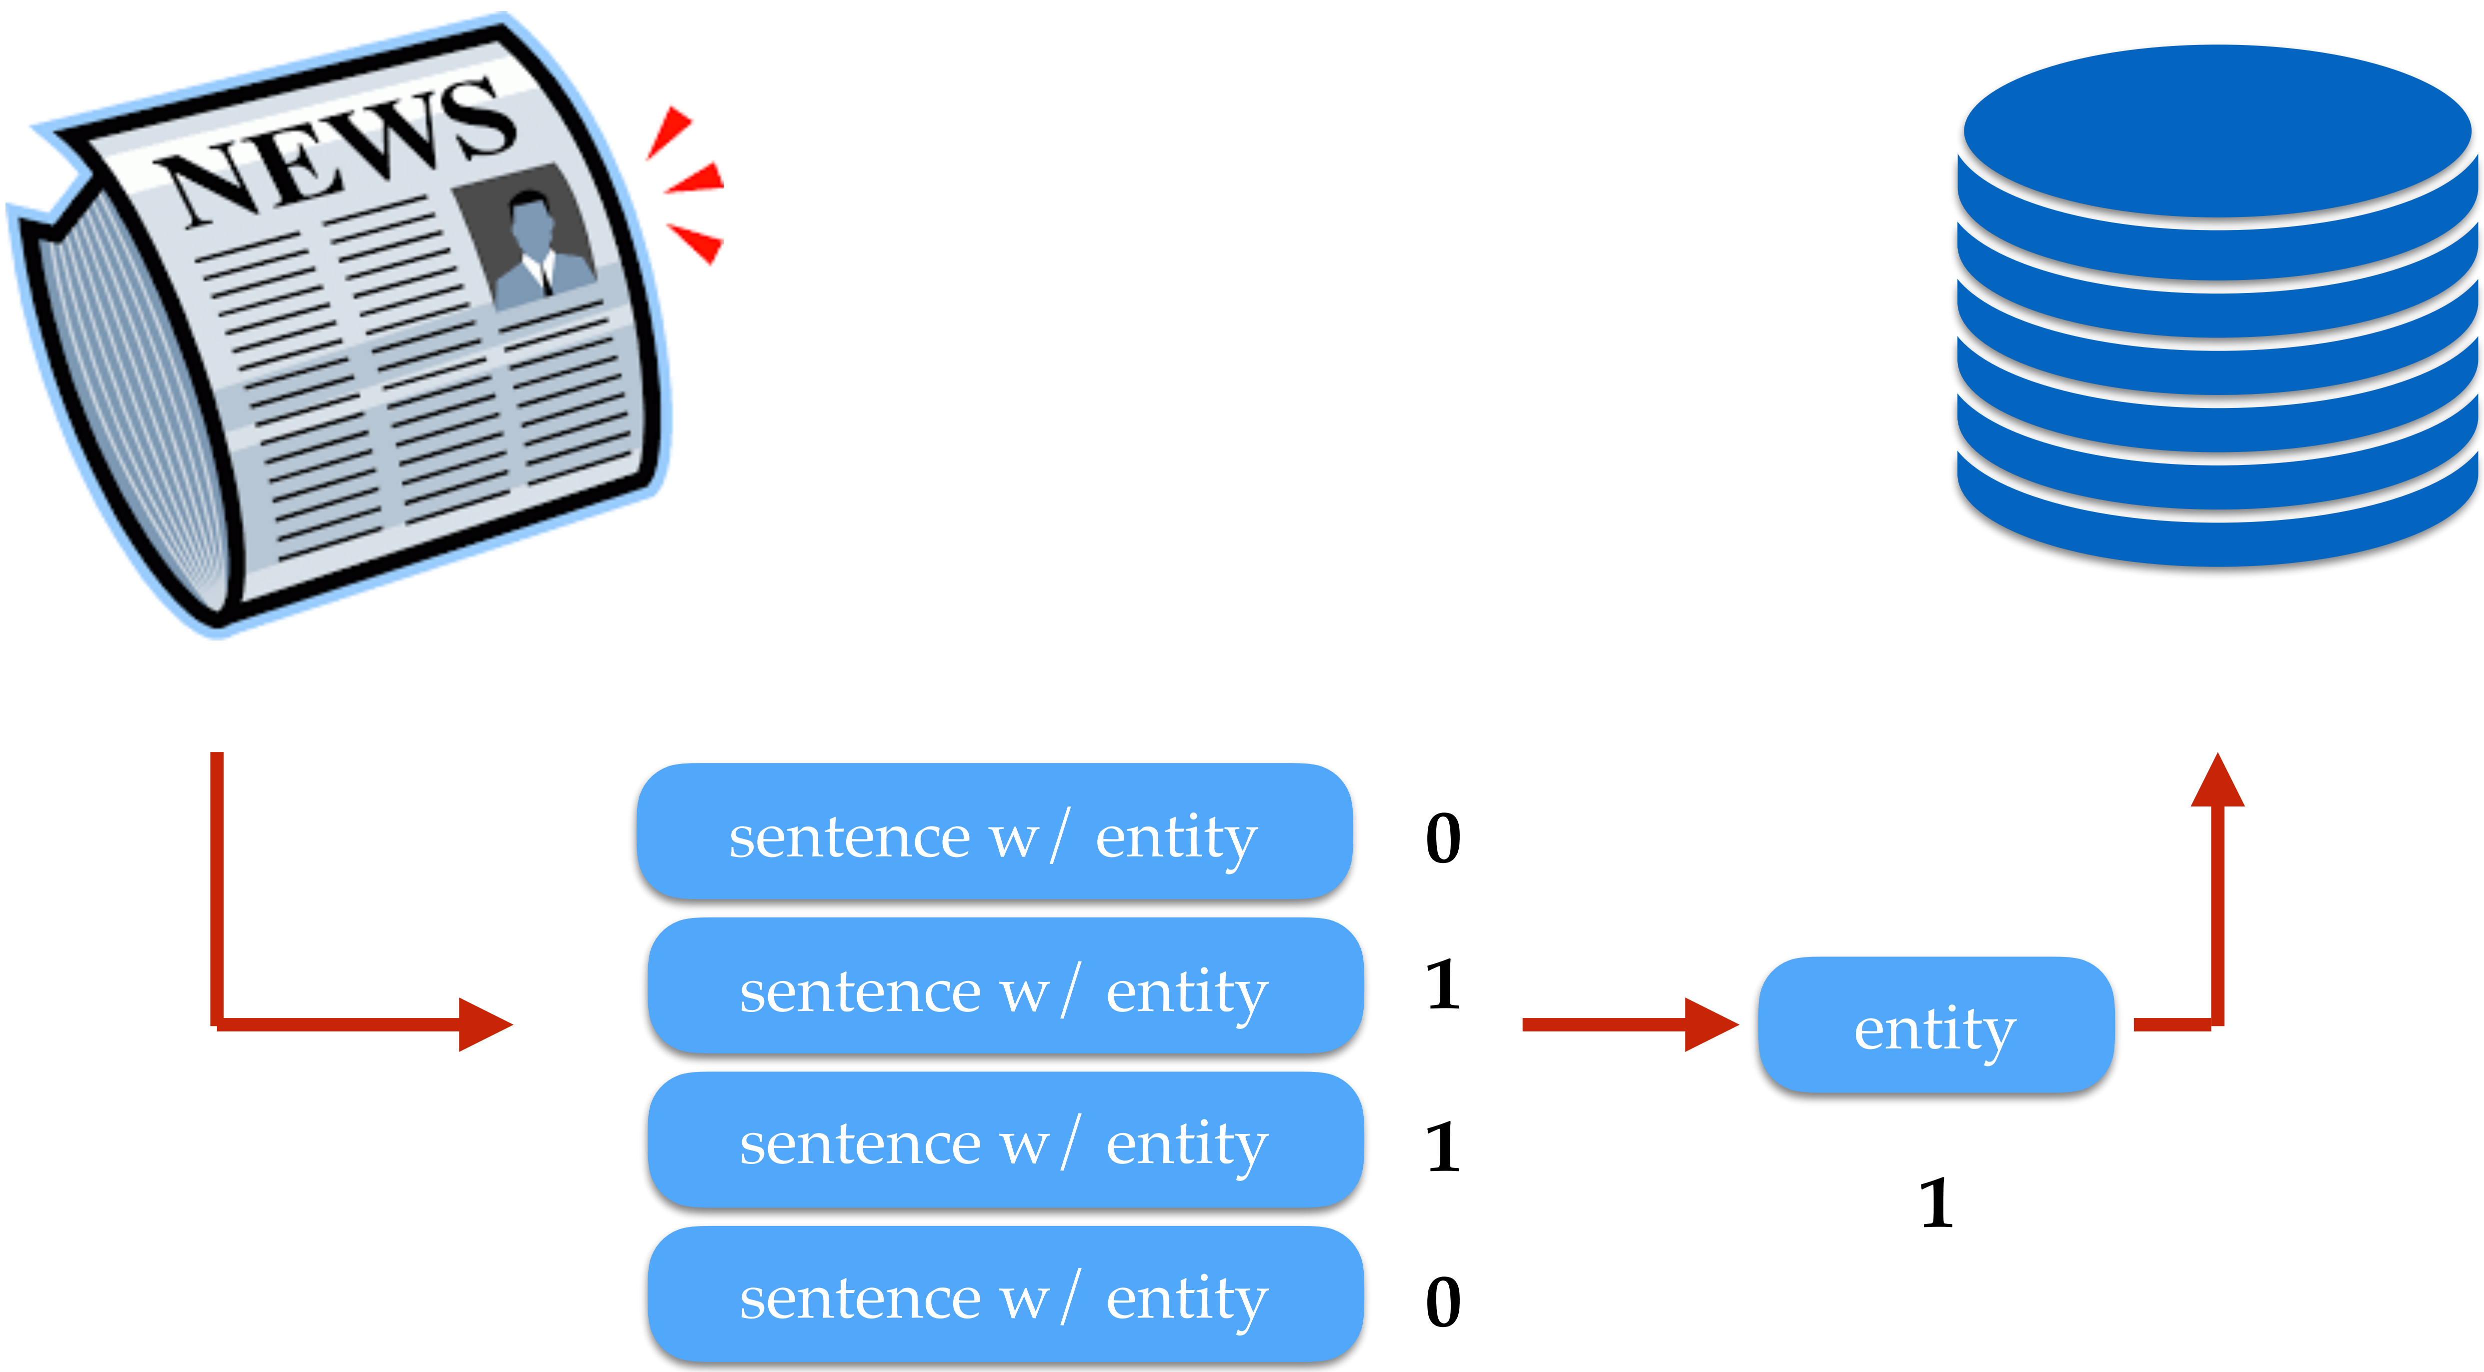
NEWS
sentence w/ entity
0
sentence w/ entity
1
sentence w/ entity
1
sentence w/ entity
0
entity
1

Outline

1. Motivation and overview
2. Task and data
3. Model
4. Training
5. Evaluation
6. Conclusion

Example Dataset

Corpus

July 17, 2014

Aug 9, 2014

July 5, 2016

July 6, 2016

Database

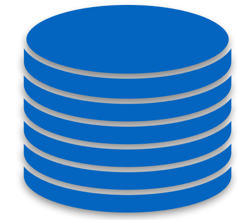

Eric Garner	New York, NY
Michael Brown	Ferguson, MO
Alton Sterling	Baton Rouge, LA
Philando Castile	Falcon Heights, MN

Example Dataset

Corpus

Database

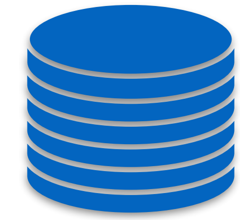

July 17, 2014	Eric Garner
Aug 9, 2014	Michael Brown
July 5, 2016	Alton Sterling
July 6, 2016	Philando Castile

Task: Database update

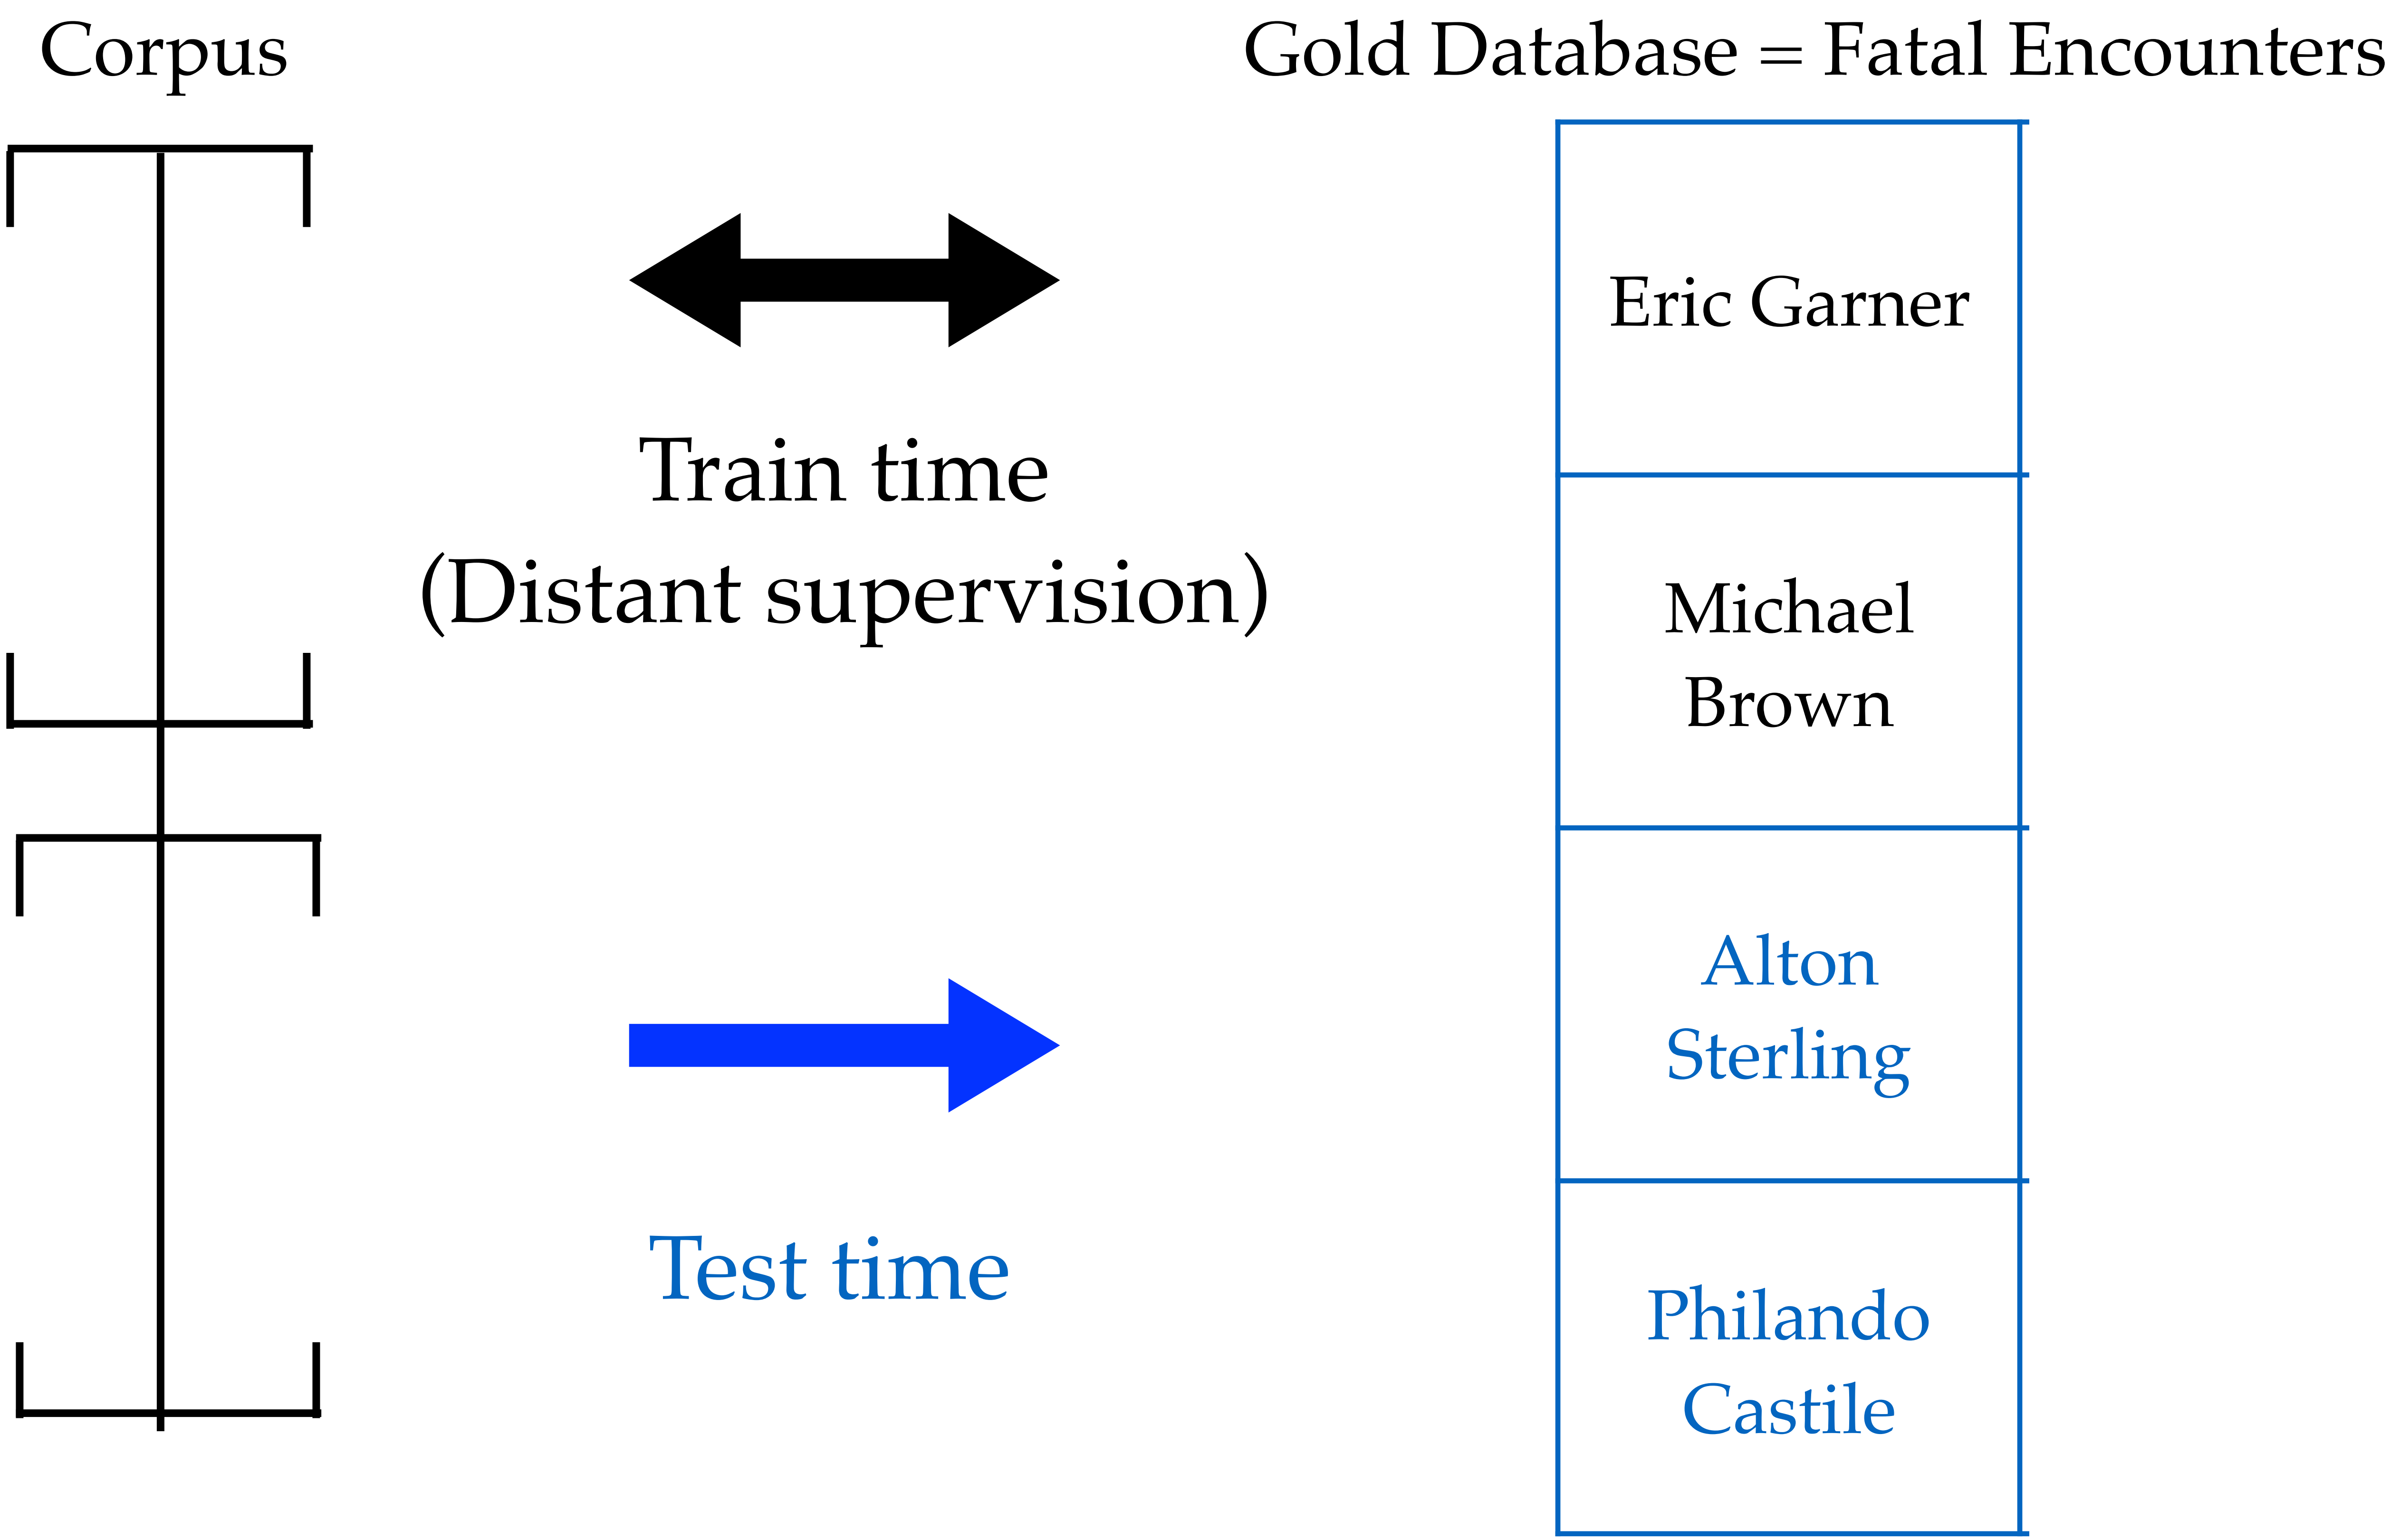

Collecting data

- Keyword-querying web scraper running throughout 2016
- Preprocessing: text extraction, deduplication, spaCy NER+parsing, name cleanups

Data

Knowledge base	Historical	Test
FE incident dates	Jan 2000 – Aug 2016	Sep 2016 – Dec 2016

News dataset	Train	Test
doc. dates	Jan 2016 – Aug 2016	Sep 2016 – Dec 2016

Data

Knowledge base	Historical	Test
FE incident dates	Jan 2000 – Aug 2016	Sep 2016 – Dec 2016

News dataset	Train	Test
doc. dates	Jan 2016 – Aug 2016	Sep 2016 – Dec 2016
total docs.	793,010	317,345

Data

Knowledge base	Historical	Test
FE incident dates	Jan 2000 – Aug 2016	Sep 2016 – Dec 2016
FE gold entities	17,219	452

News dataset	Train	Test
doc. dates	Jan 2016 – Aug 2016	Sep 2016 – Dec 2016
total docs.	793,010	317,345
total ments.	132,833	68,925
pos. ments.	11,274	6,132
total entities	49,203	24,550
pos. entities	916	258

Data

Knowledge base	Historical	Test
FE incident dates	Jan 2000 – Aug 2016	Sep 2016 – Dec 2016
FE gold entities	17,219	452

News dataset	Train	Test
doc. dates	Jan 2016 – Aug 2016	Sep 2016 – Dec 2016
total docs.	793,010	317,345
total ments.	132,833	68,925
pos. ments.	11,274	6,132
total entities	49,203	24,550
pos. entities	916	258

Data upper bound:
258 / 452 = 57% recall

Outline

1. Motivation and overview
2. Task and data
3. Model
4. Training
5. Evaluation
6. Conclusion

Test time

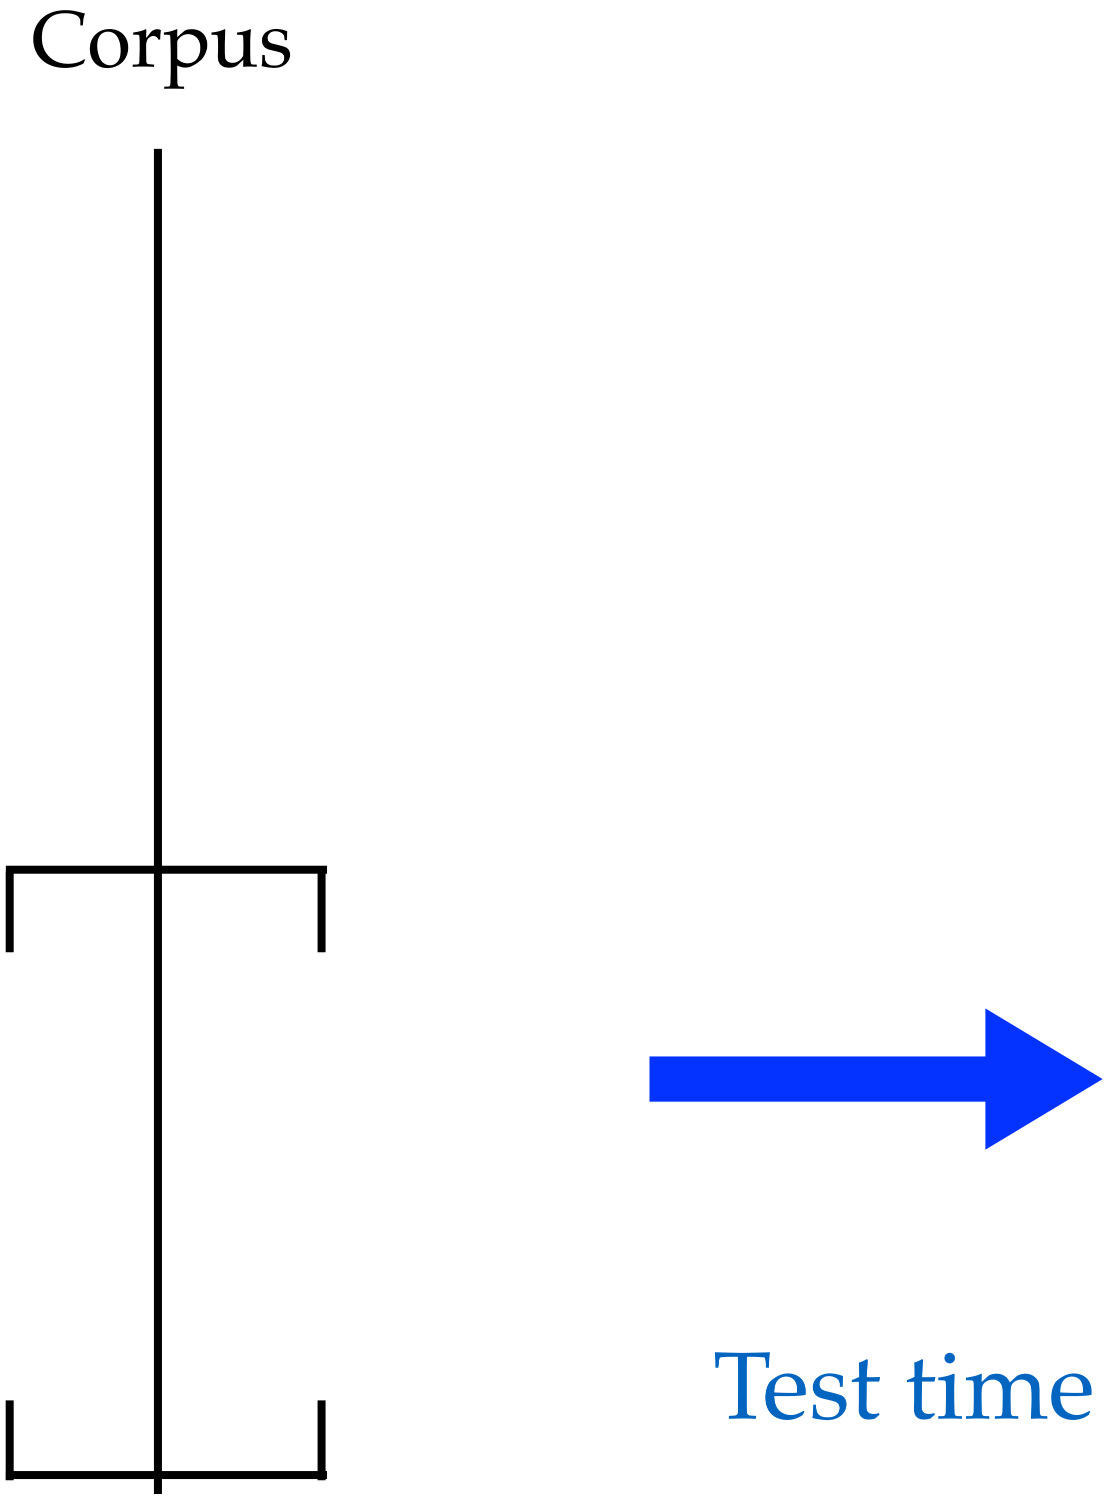

Test time

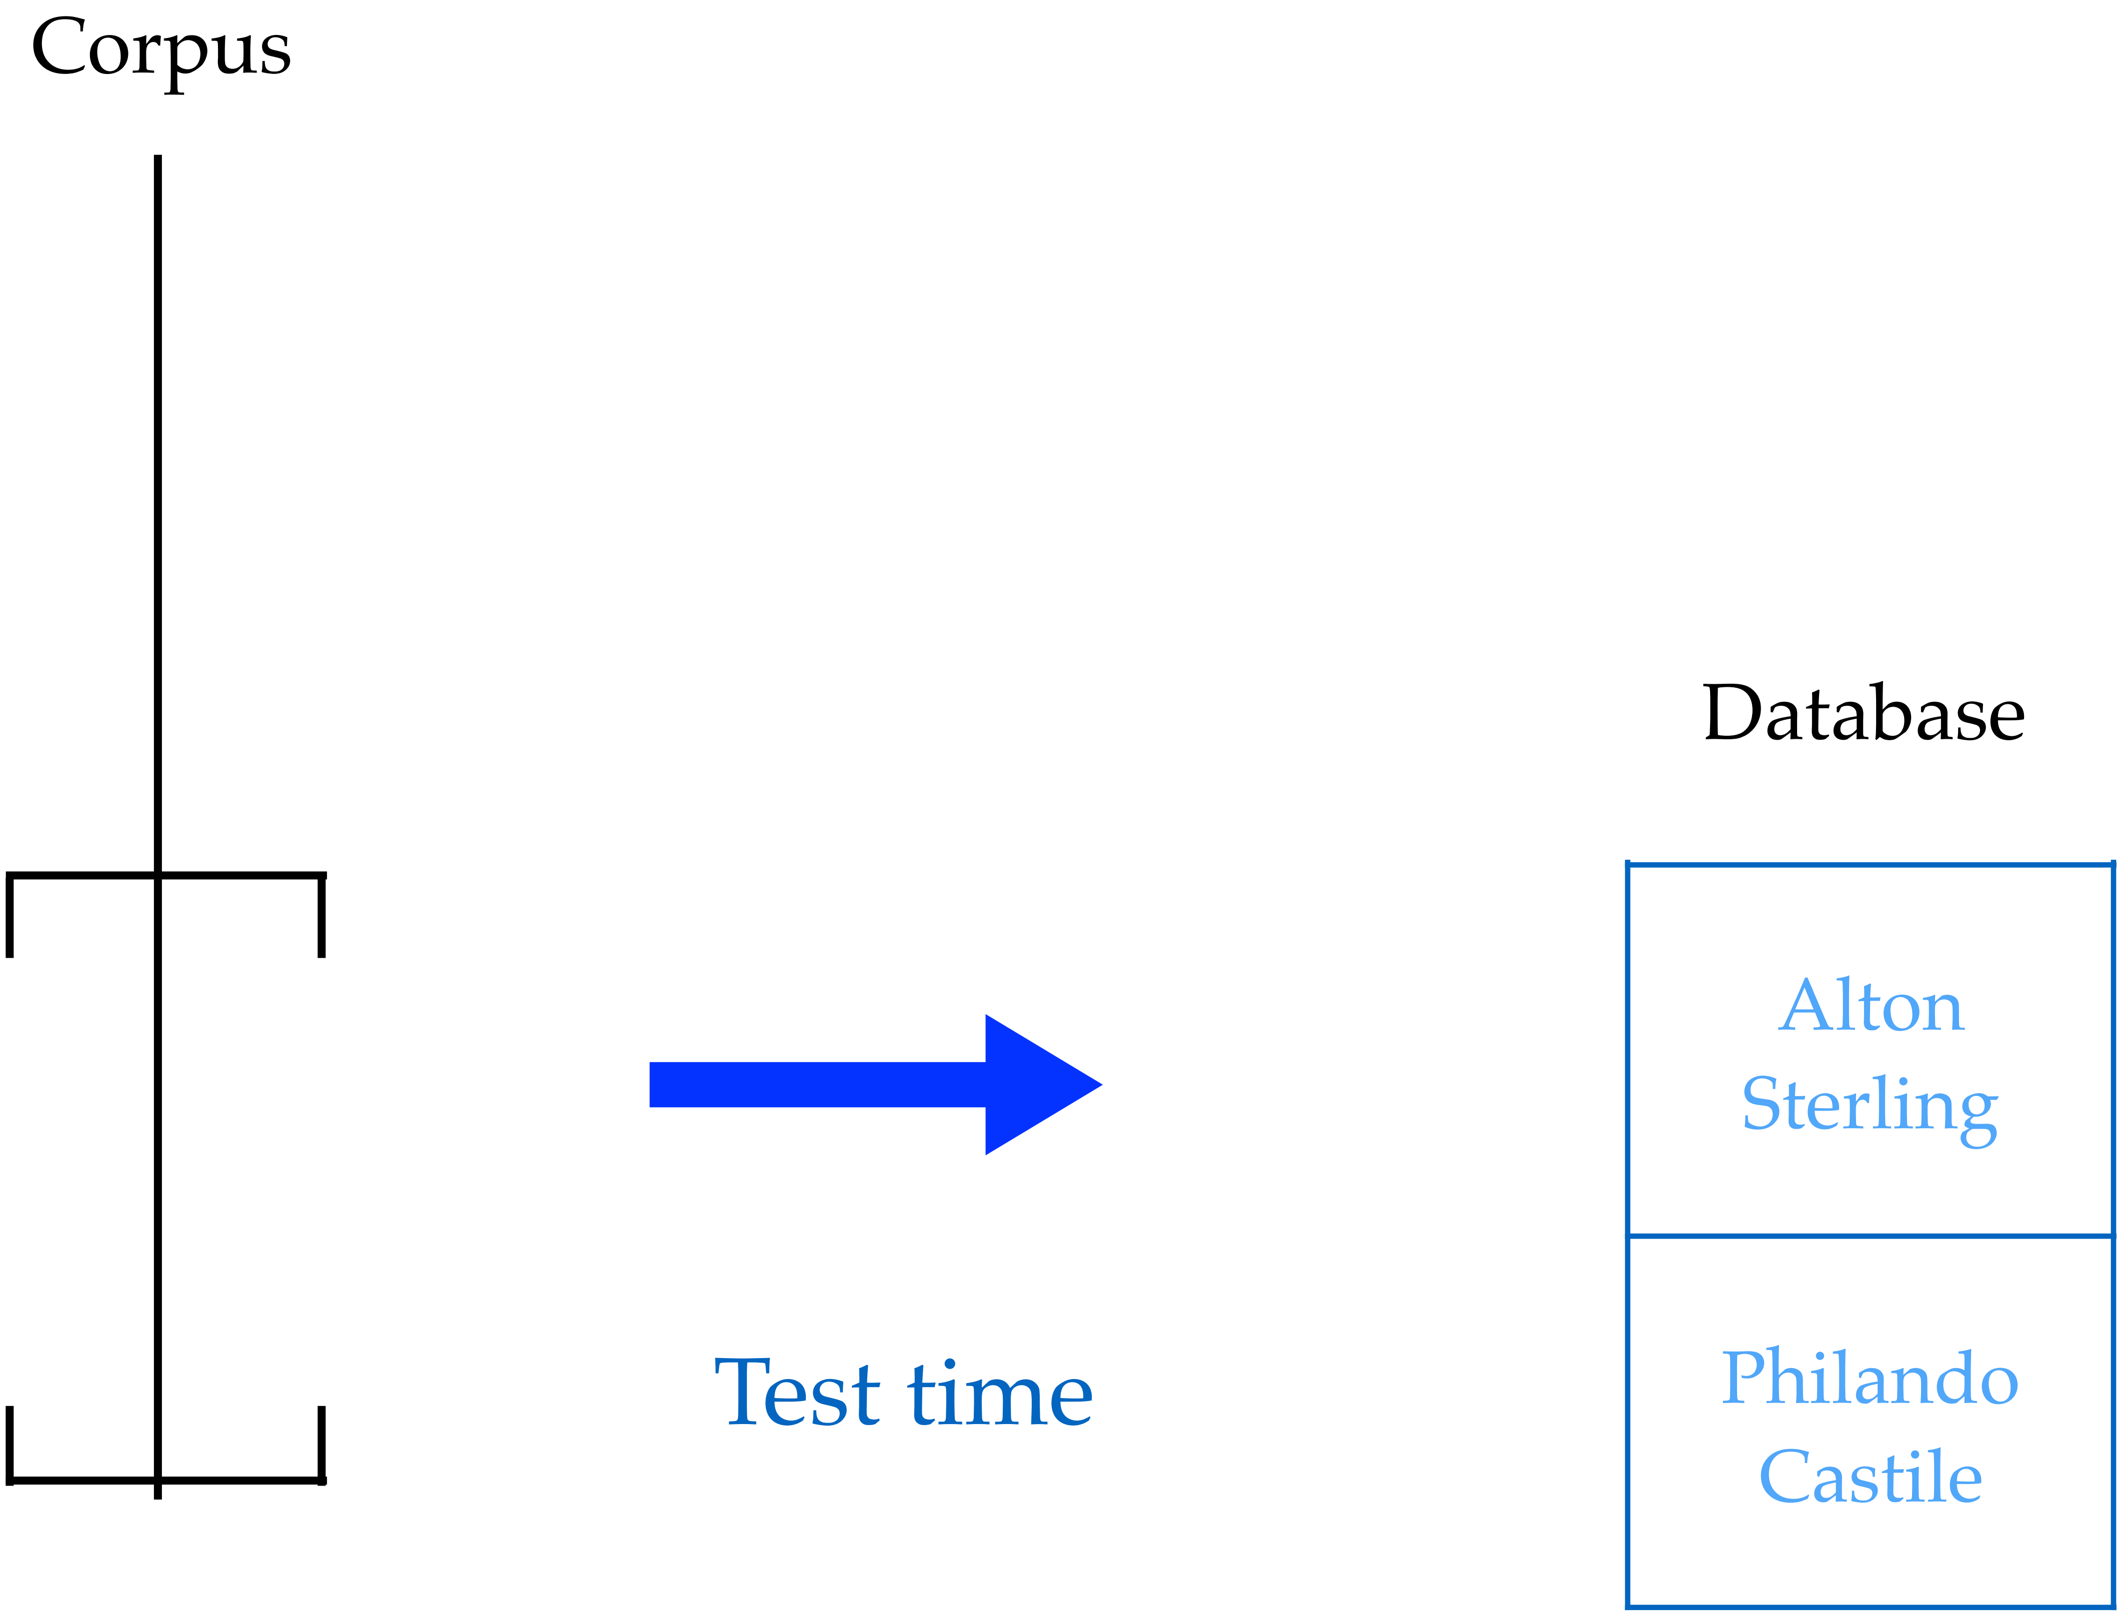
Corpus
Database
Test time
Alton Sterling
Philando Castile

Test time

Corpus

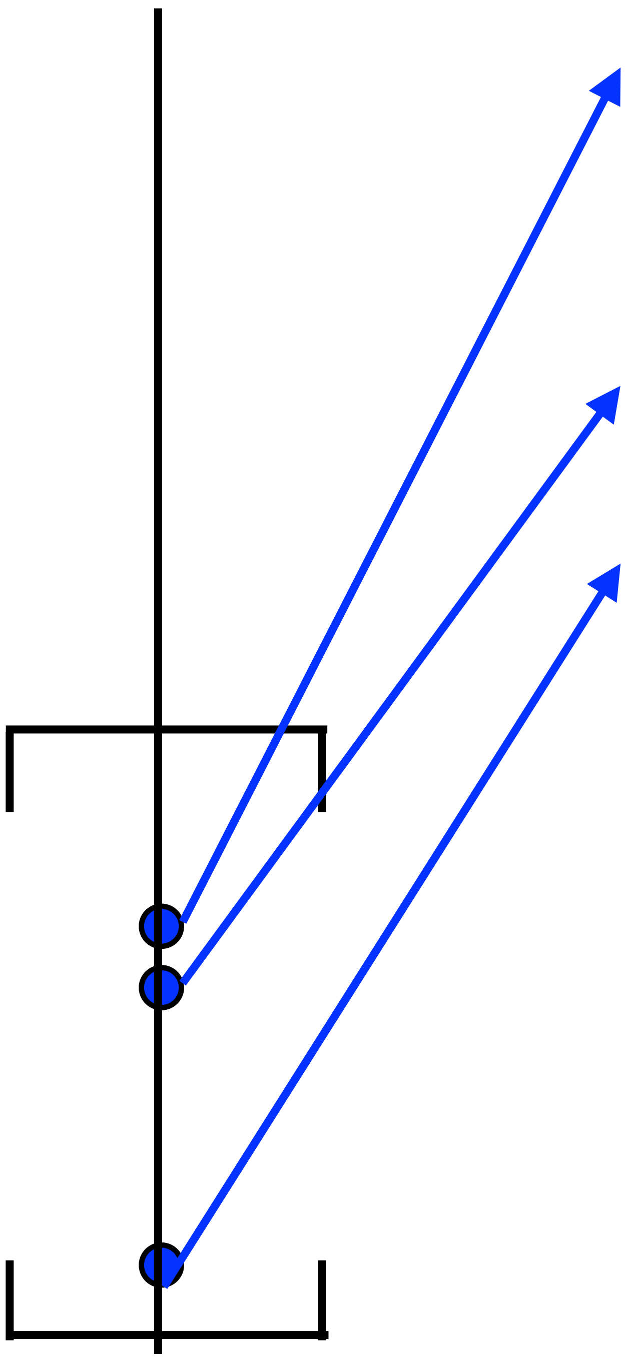

The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart

... the two officers involved in Tuesday's shooting of **Alton Sterling** ...

... **Alton Sterling** was a resident of Baton Rouge...

Test time

Corpus

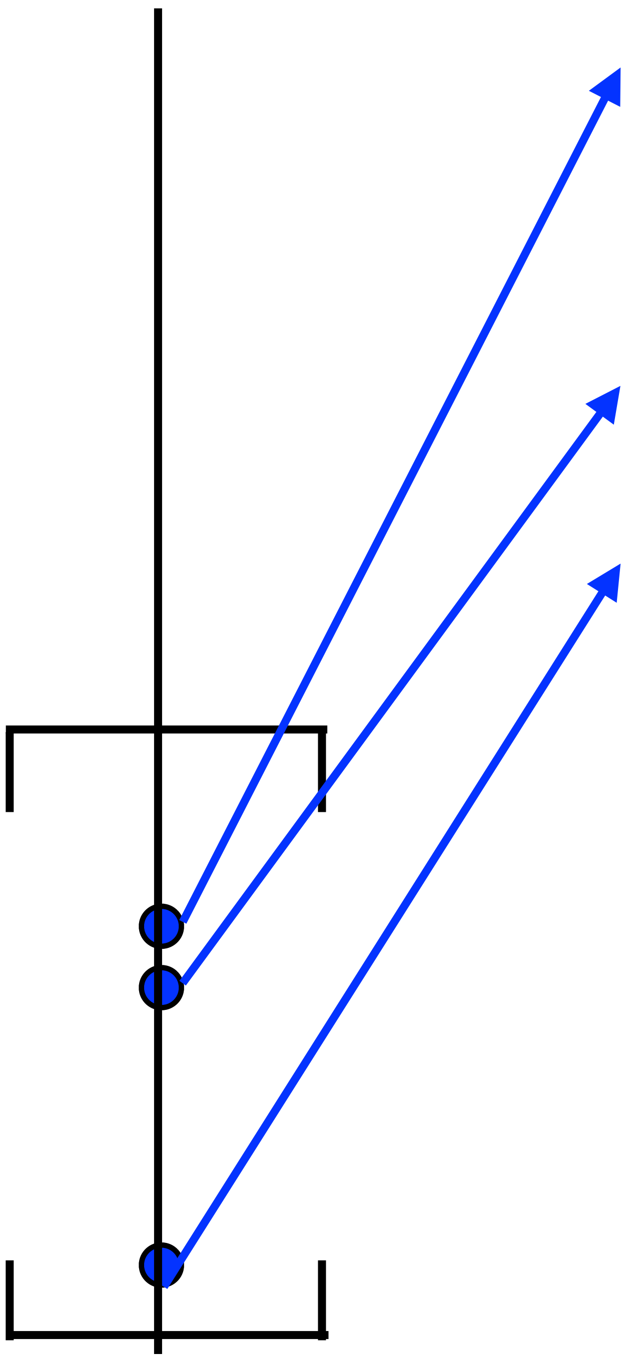

The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart

... the two officers involved in Tuesday's shooting of **Alton Sterling** ...

... **Alton Sterling** was a resident of Baton Rouge...

(1) predict: describes police fatality?

0.4

0.8

0.01

Test time

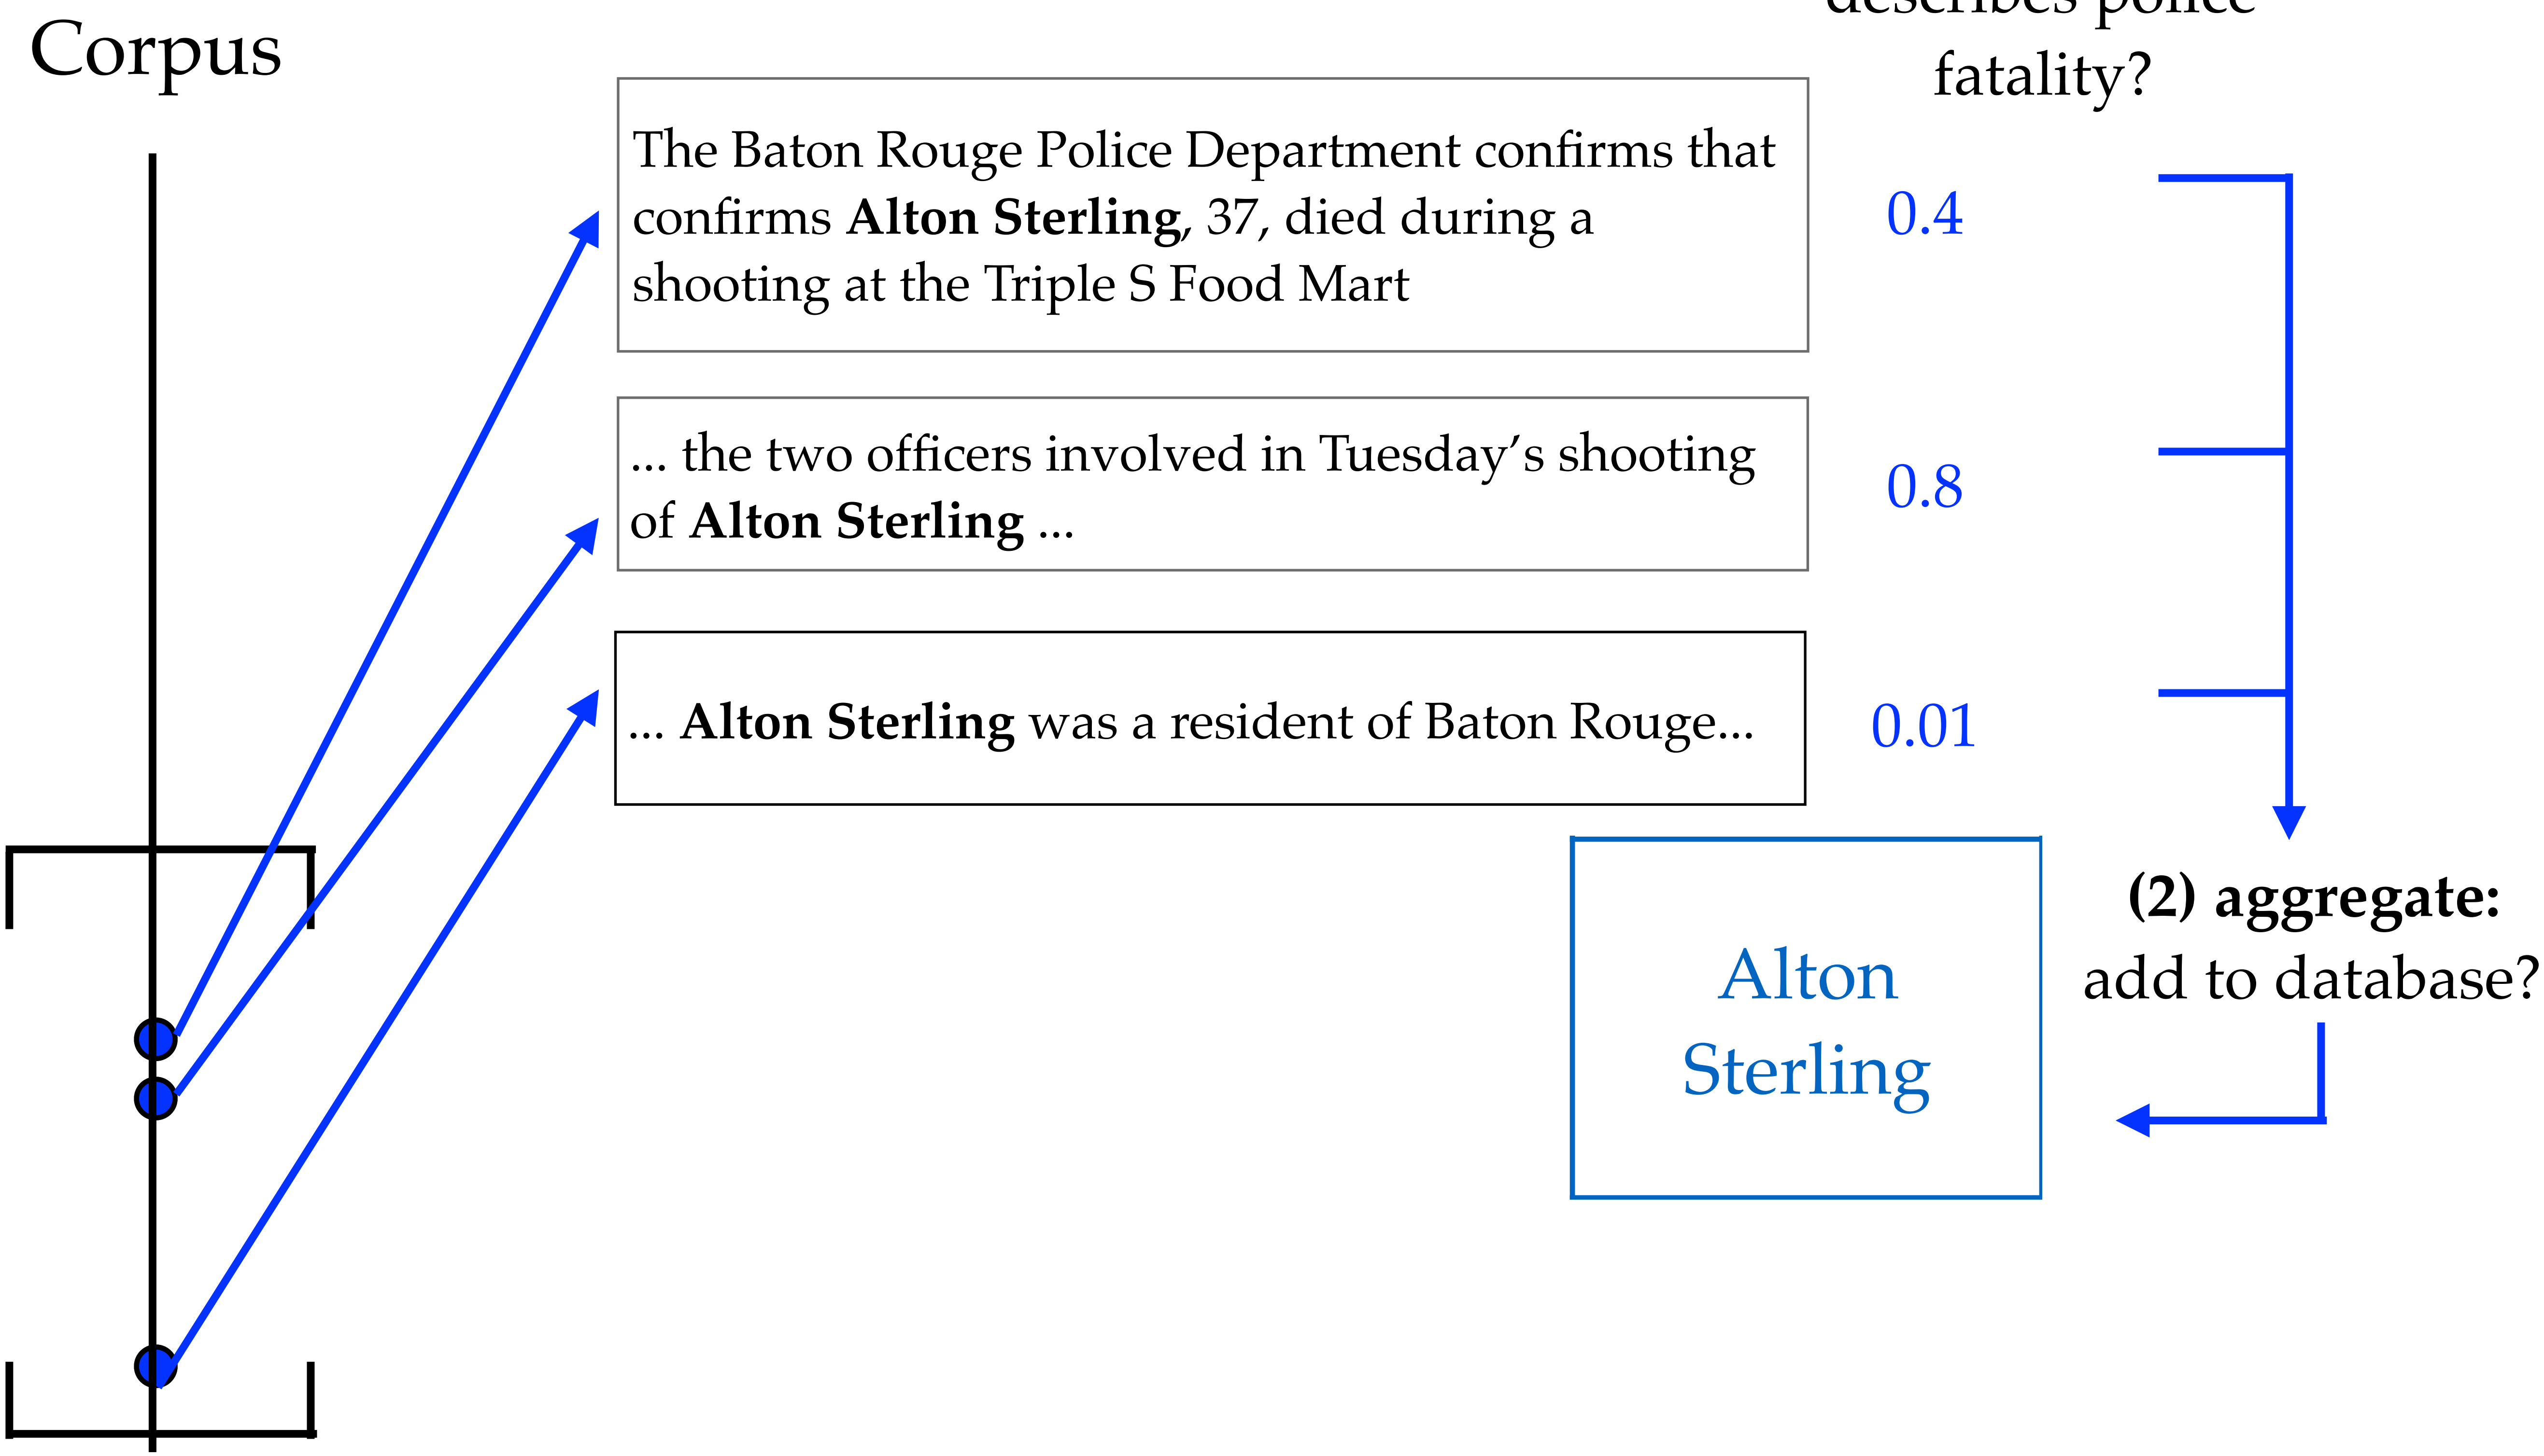
Corpus
(1) predict: describes police fatality?
The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart
0.4
... the two officers involved in Tuesday's shooting of **Alton Sterling** ...
0.8
... **Alton Sterling** was a resident of Baton Rouge...
0.01
(2) aggregate: add to database?
Alton Sterling

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

$$P(z_i = 1 | x_i) = \sigma(\theta^T f(x_i))$$

describes
police killing
event

sentence
text

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

$$P(z_i = 1 | x_i) = \sigma(\theta^T f(x_i))$$

describes police killing event

sentence text

e.g. logistic regression, convolutional neural network

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

1. **Feature-engineered logistic regression**
   - Syntactic dependency paths
   - N-grams
   - POS tags

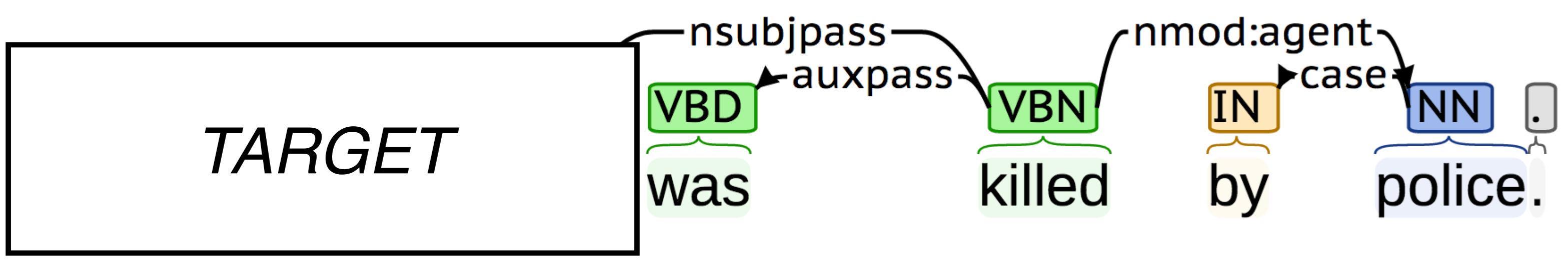

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

1. **Feature-engineered logistic regression**
   - Syntactic dependency paths
   - N-grams
   - POS tags

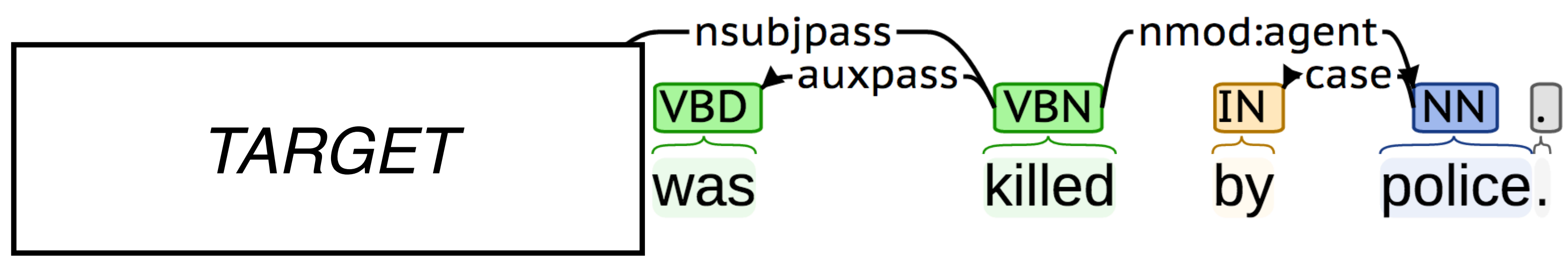

2. **Convolutional neural network**
   - *[Kim 2014]*
   - Used in other event detection work *[e.g. Nguyen and Grishman 2015]*

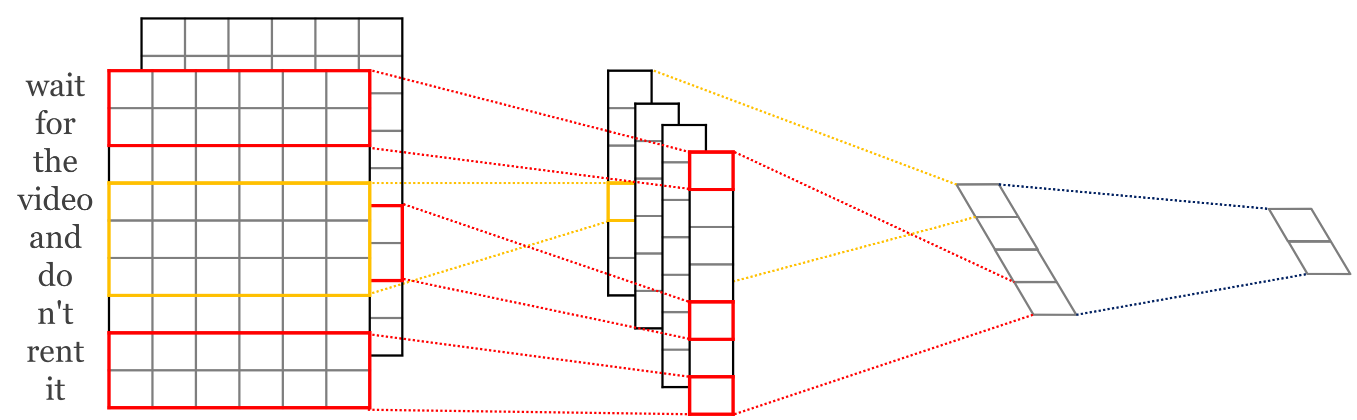

[Kim 2014]

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

	p(z | x)
The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart	0.4
... the two officers involved in Tuesday 's shooting of **Alton Sterling** ...	0.8
... **Alton Sterling** was a resident of Baton Rouge...	0.01

aggregate

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

p(z | x)

Sentence	p(z | x)
The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart	0.4
... the two officers involved in Tuesday 's shooting of **Alton Sterling** ...	0.8
... **Alton Sterling** was a resident of Baton Rouge...	0.01

aggregate

max	mean	median
.8	.403	.4

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

p(z | x)

The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart	0.4
... the two officers involved in Tuesday 's shooting of **Alton Sterling** ...	0.8
... **Alton Sterling** was a resident of Baton Rouge...	0.01

aggregate

noisy-or	max	mean	median
.881	.8	.403	.4

Noisy-Or

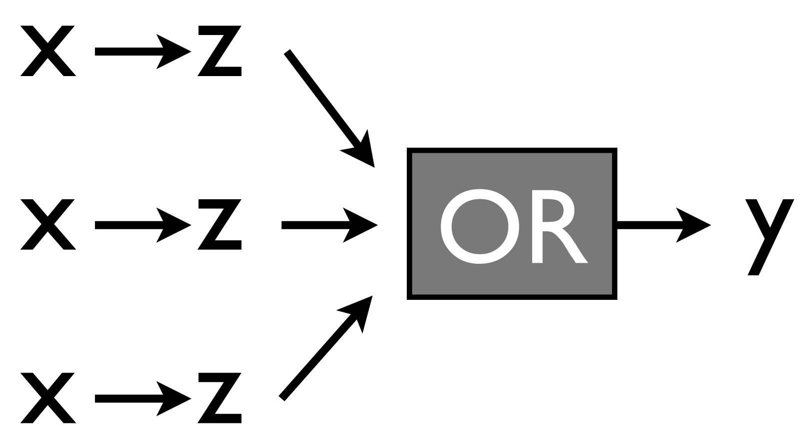

Noisy-Or

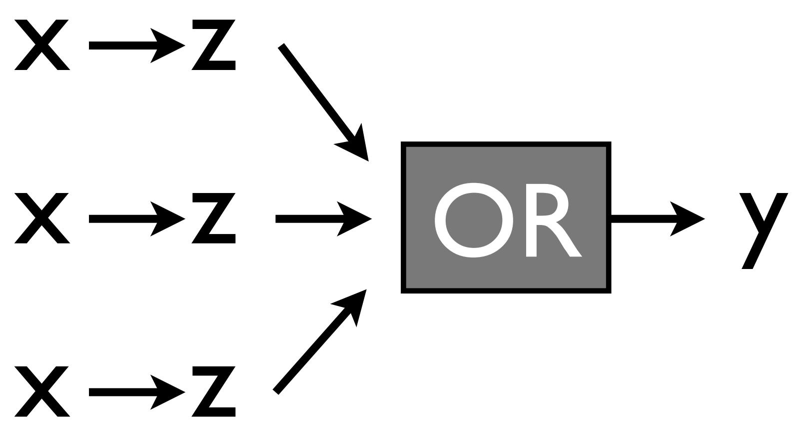

$$P(y_e = 1|x_{\mathcal{M}(e)}) = 1 - \prod_{i \in \mathcal{M}(e)} (1 - P(z_i = 1 \mid x_i))$$

entity label

set of sentences for given entity

sentence label

Model

(1) Predict sentence-level **event** assertions

(2) Aggregate **entity**-level predictions

	p(z | x)
The Baton Rouge Police Department confirms that confirms **Alton Sterling**, 37, died during a shooting at the Triple S Food Mart	0.4
... the two officers involved in Tuesday 's shooting of **Alton Sterling** ...	0.8
... **Alton Sterling** was a resident of Baton Rouge...	0.01

aggregate

noisy-or	max	mean	median
.881	.8	.403	.4

Outline

1. Motivation and overview
2. Task and data
3. Model
4. Training
5. Evaluation
6. Conclusion

Imputing training labels

Corpus

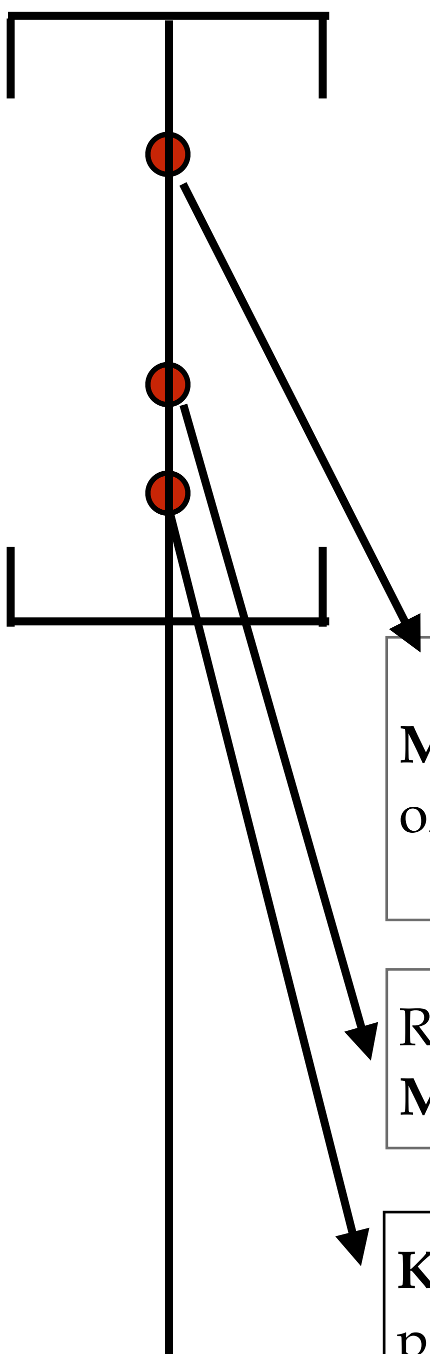

Michael Brown was killed by a white police officer in Ferguson , Mo.

Reporters interviewed the mother of **Michael Brown** last Sunday.

Katy Perry reacted on Twitter to the most recent police killing.

Database

Eric Garner
Michael Brown

Imputing training labels

Corpus

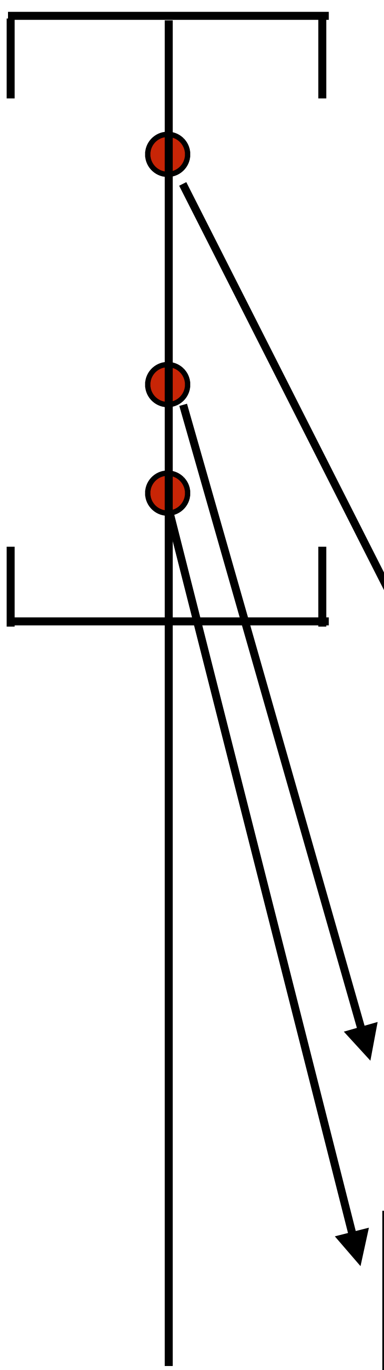

hand labeling is expensive
—> distant supervision

Michael Brown was killed by a white police officer in Ferguson , Mo.

Reporters interviewed the mother of **Michael Brown** last Sunday.

Katy Perry reacted on Twitter to the most recent police killing.

Database

Database
Eric Garner
Michael Brown

Imputing training labels

1. “Hard” labeling

2. “Soft” labeling

Imputing training labels

1. “Hard” labeling

Distant Supervision Assumption

[Mintz et al., 2009]

2. “Soft” labeling

(1)"Hard" labeling

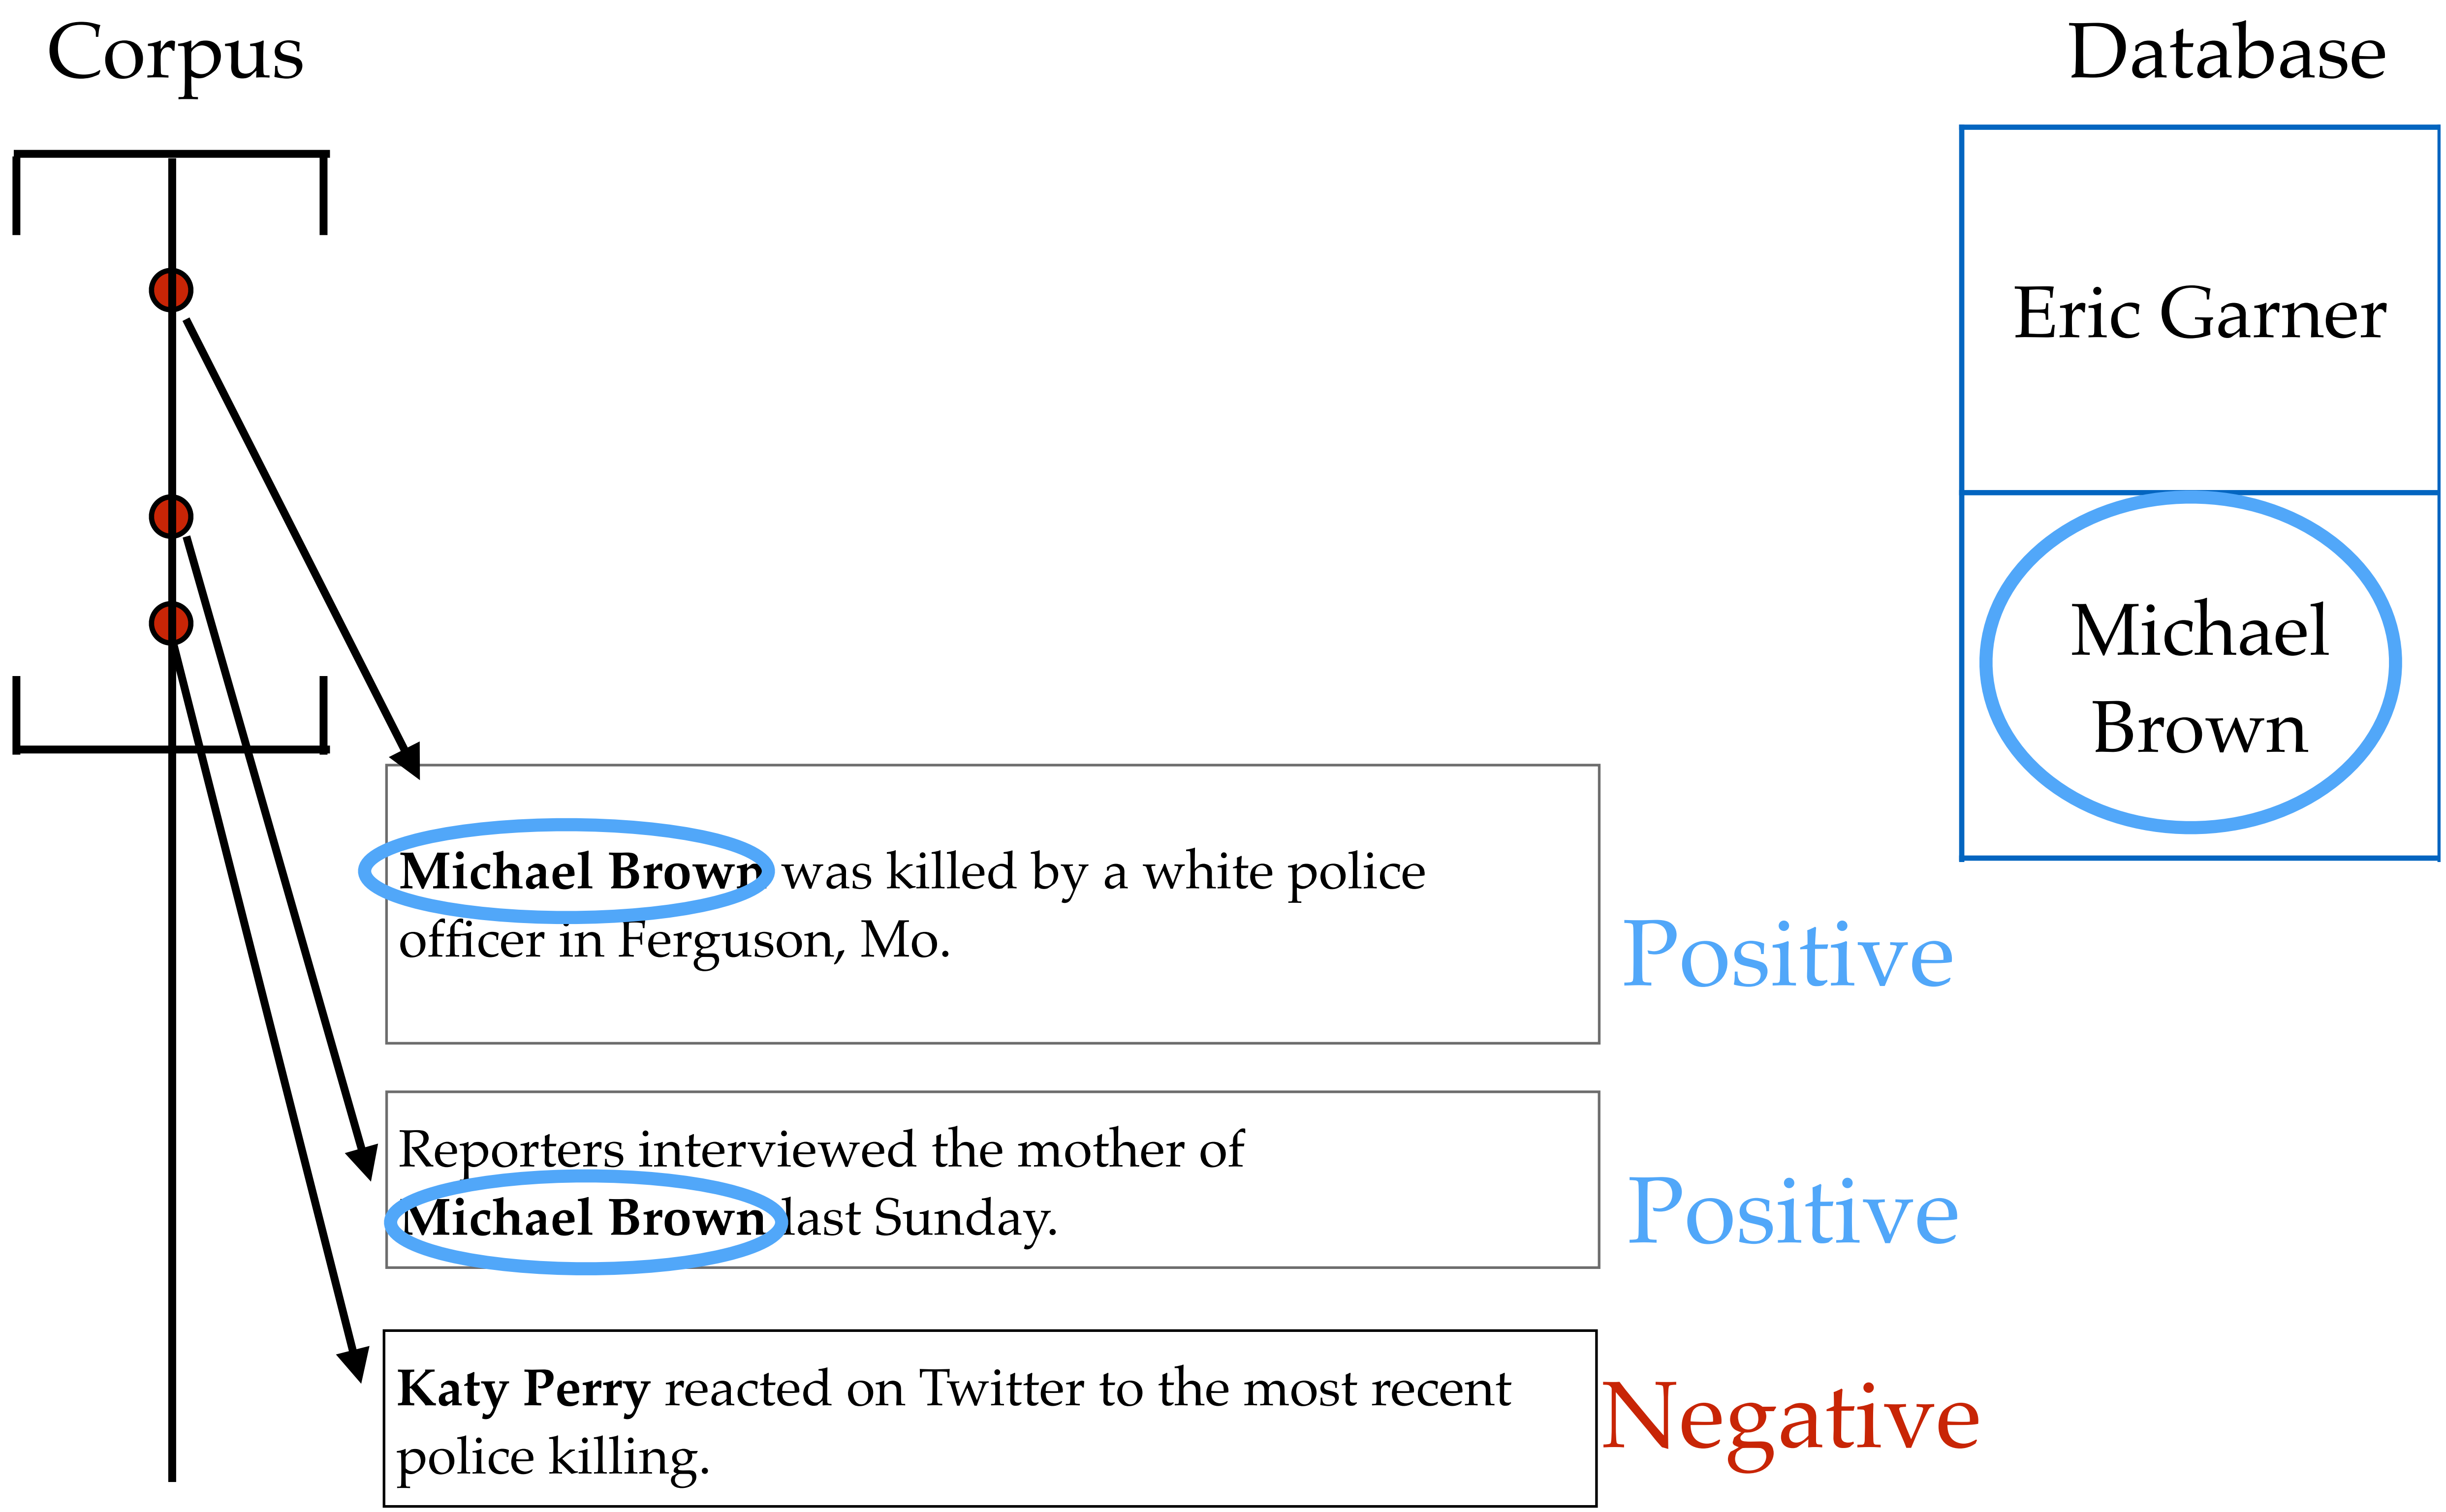

(1)"Hard" labeling

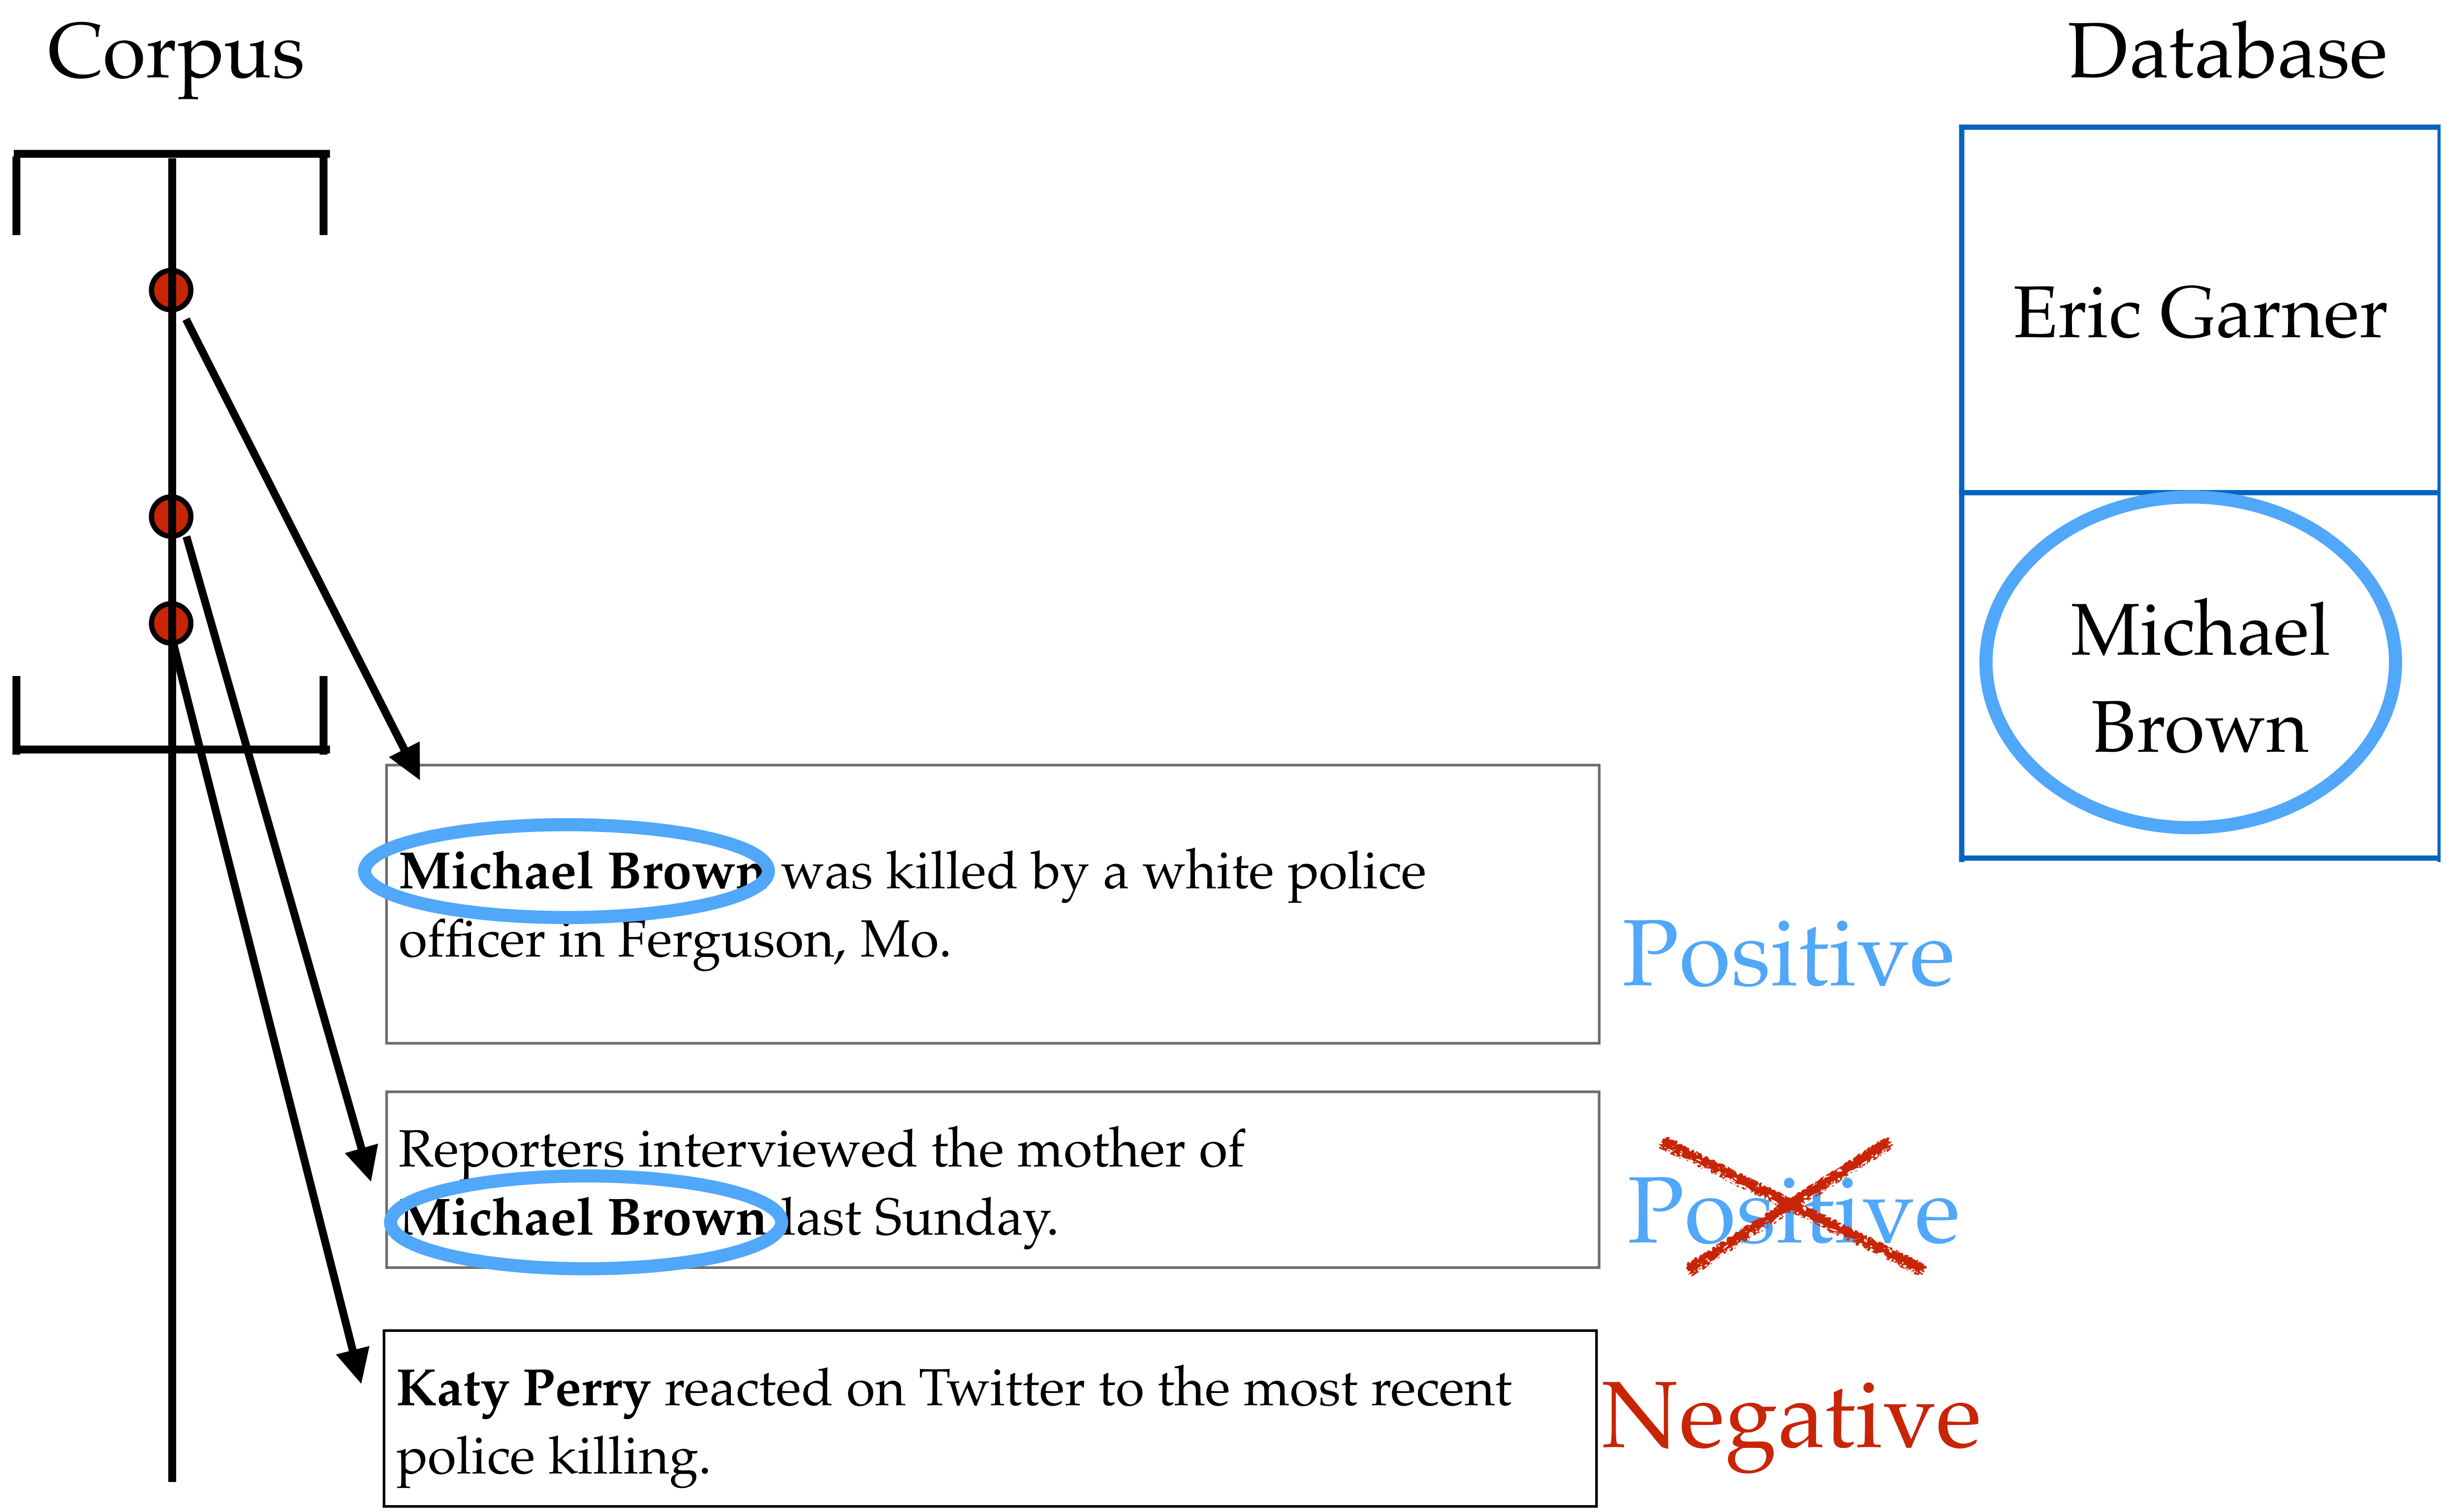

(1)"Hard" labeling

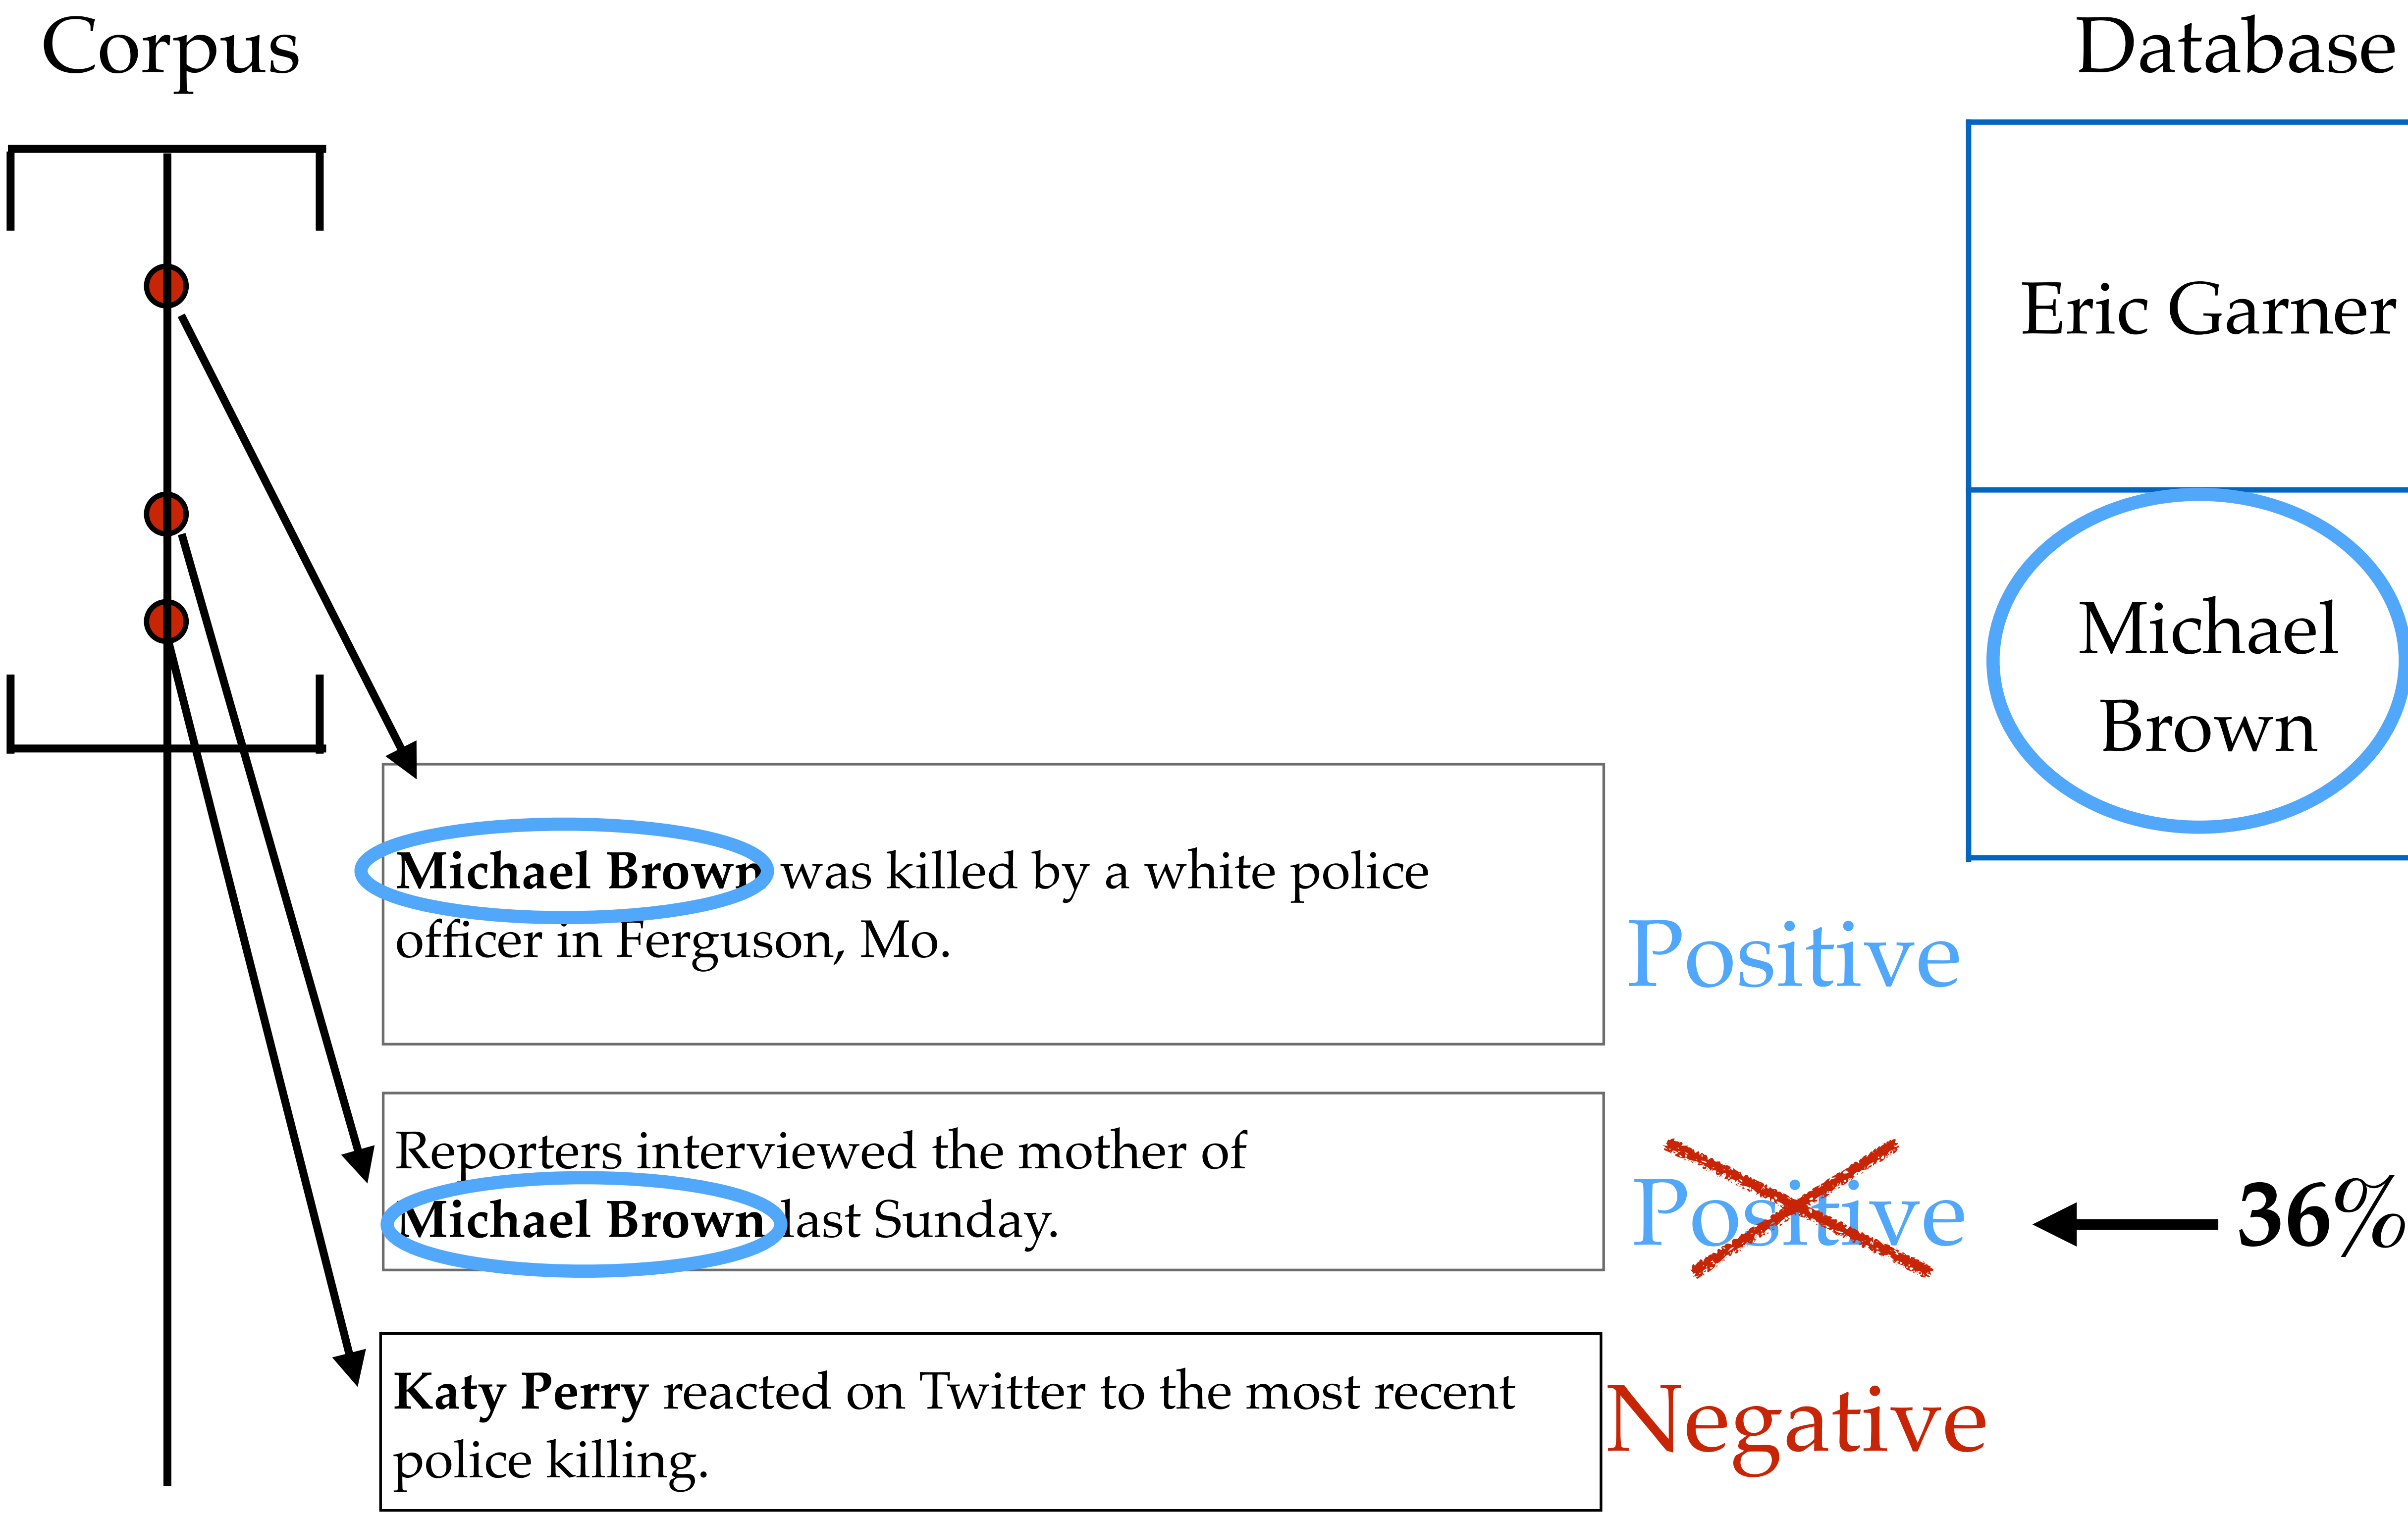

Imputing training labels

1. “Hard” labeling

Distant Supervision Assumption

[Mintz et al., 2009]

2. “Soft” labeling

“At least one” assumption

[Bunescu and Mooney 2007]

(2)"Soft" labeling

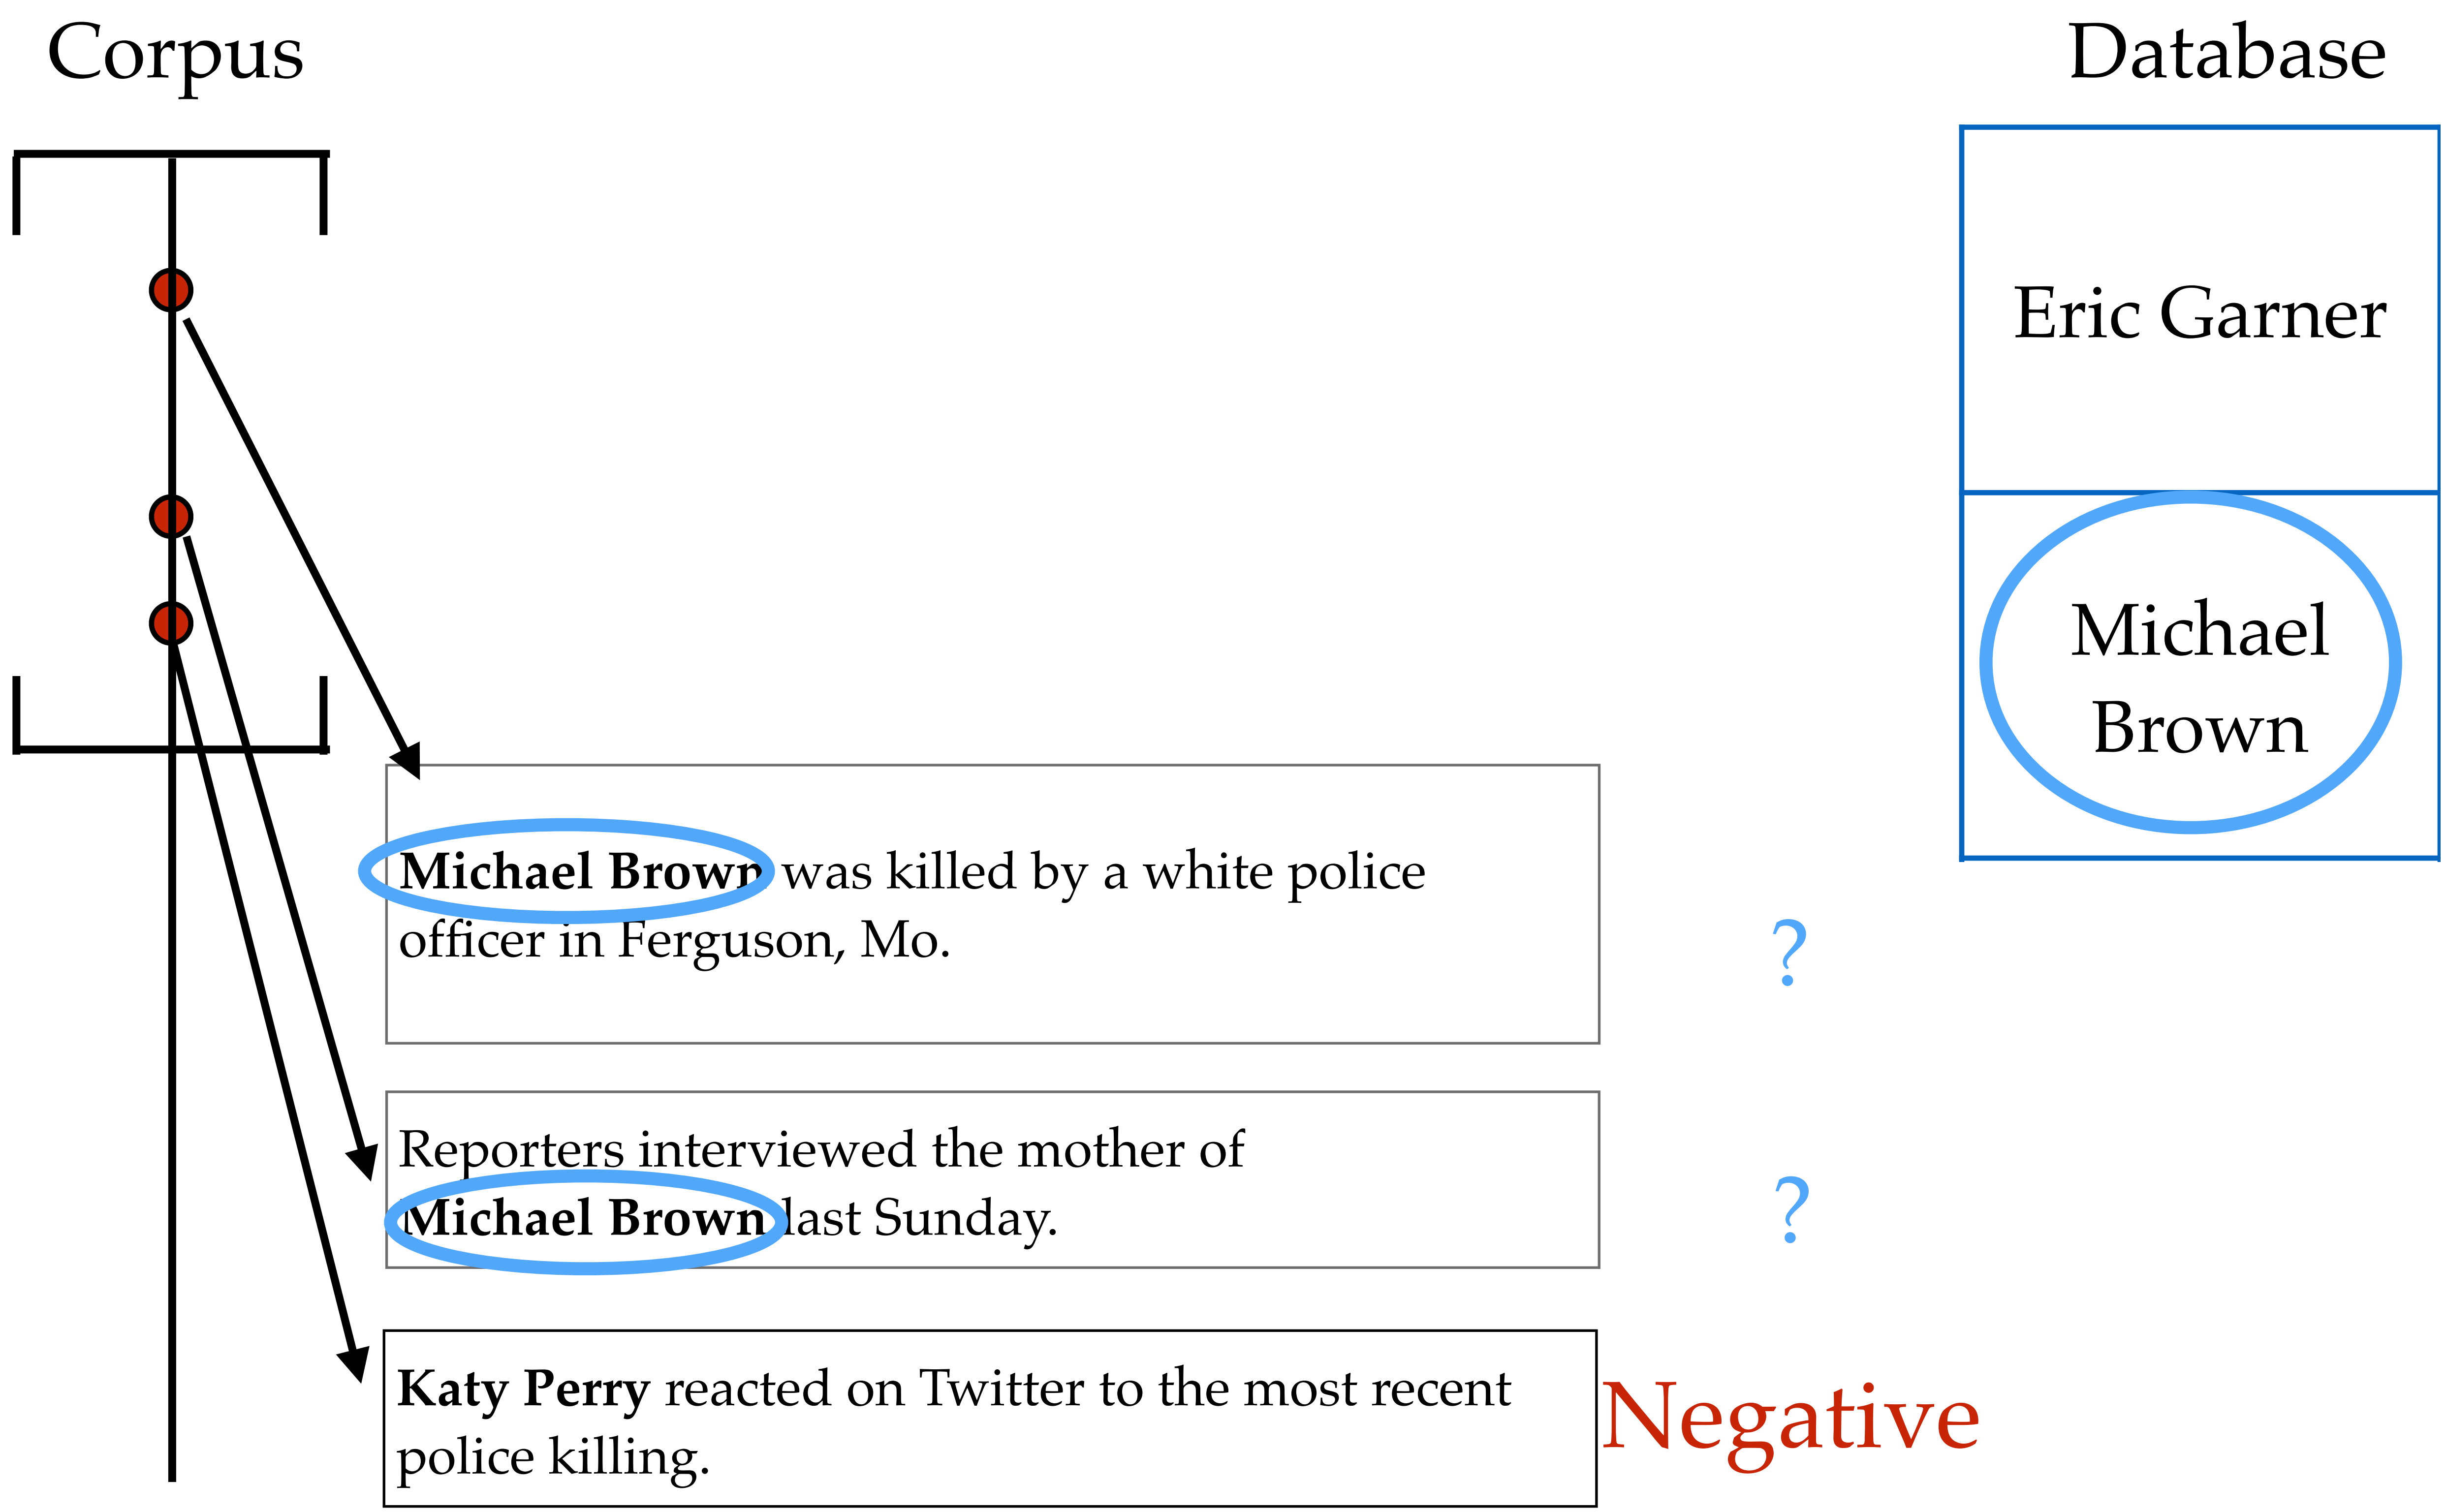

(2) “Soft” labeling

EM Training *[Dempster et al. 1977]*

(2) “Soft” labeling

EM Training *[Dempster et al. 1977]*

Initialize with hard distant labels

(2) “Soft” labeling

EM Training *[Dempster et al. 1977]*

Initialize with hard distant labels

E-Step:

Marginal posterior probability for each z_i

$$q(z_i = 1) = \frac{P(z_i = 1, y_{e_i} = 1 | x_{\mathcal{M}(e_i)})}{P(y_{e_i} = 1 | x_{\mathcal{M}(e_i)})}$$

probability sentence i is a police fatality event

entity label

set of all sentences for the given entity

(2) "Soft" labeling

EM Training *[Dempster et al. 1977]*

Initialize with hard distant labels

E-Step:

Marginal posterior probability for each z_i

$$q(z_i = 1) = \frac{P(z_i = 1, y_{e_i} = 1 | x_{\mathcal{M}(e_i)})}{P(y_{e_i} = 1 | x_{\mathcal{M}(e_i)})}$$

probability sentence i is a police fatality event

entity label

set of all sentences for the given entity

M-Step:

$$\max_{\theta} \sum_i \sum_{z \in \{0,1\}} q(z_i = z) \log P_\theta(z_i = z \mid x_i).$$

classifier parameters

Outline

1. Motivation and overview
2. Task and data
3. Model
4. Training
5. Evaluation
6. Conclusion

Model results

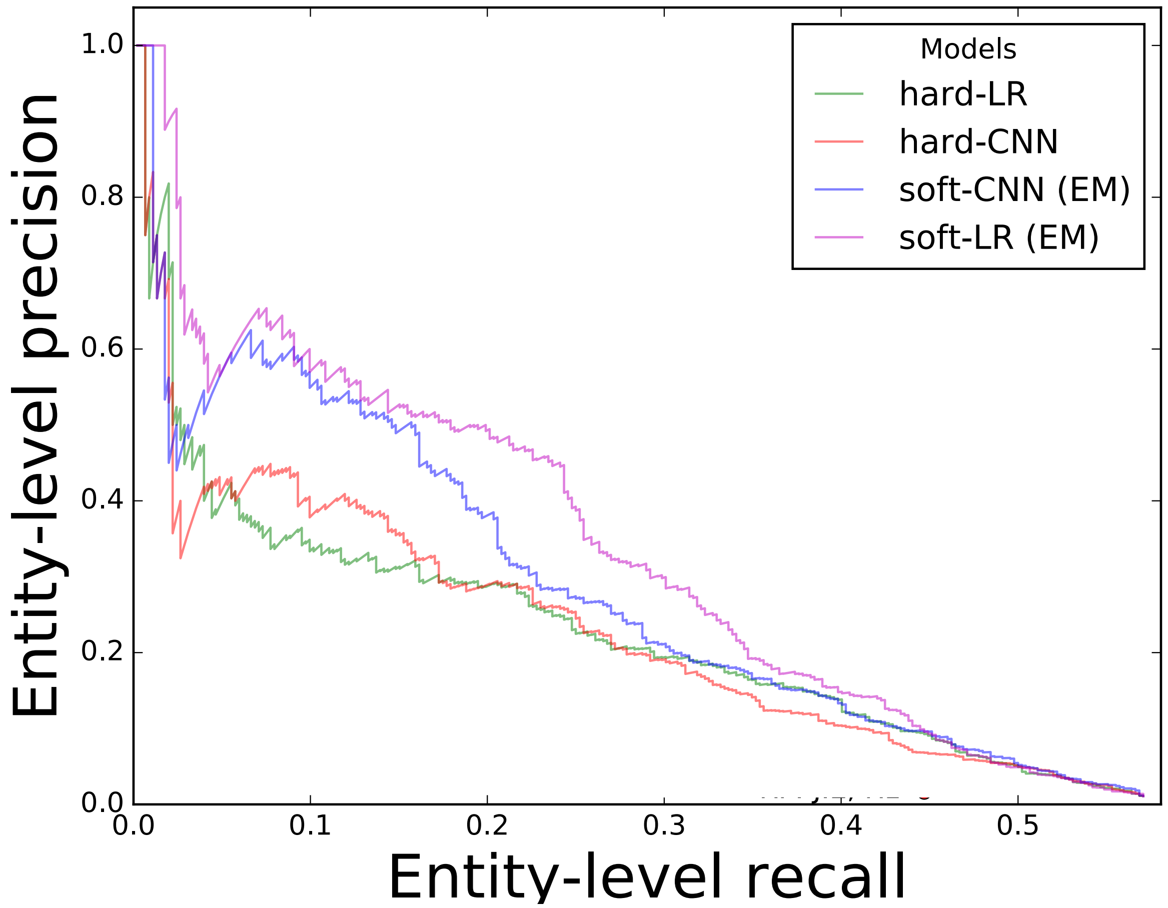

Model results

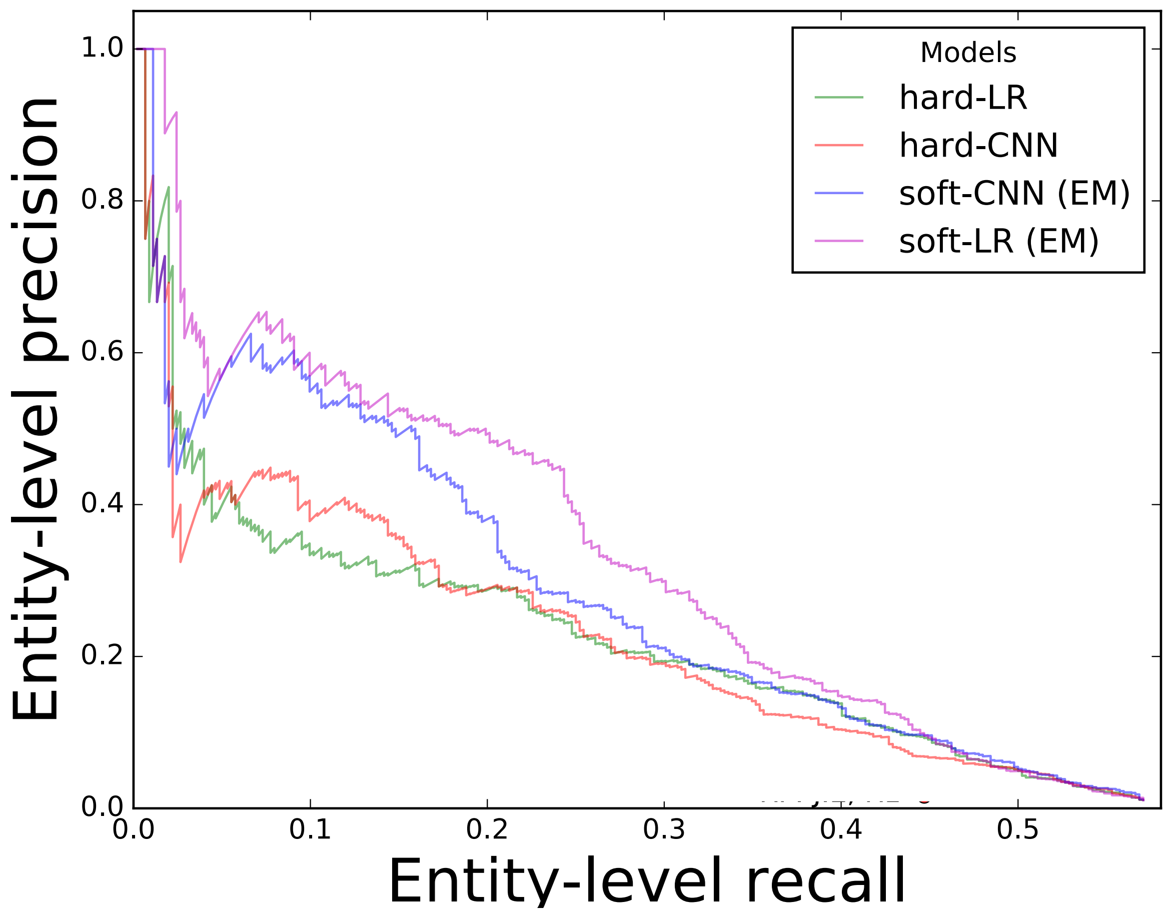

Model	AUPRC	F1
Data upper bound	0.57	0.73

Model results

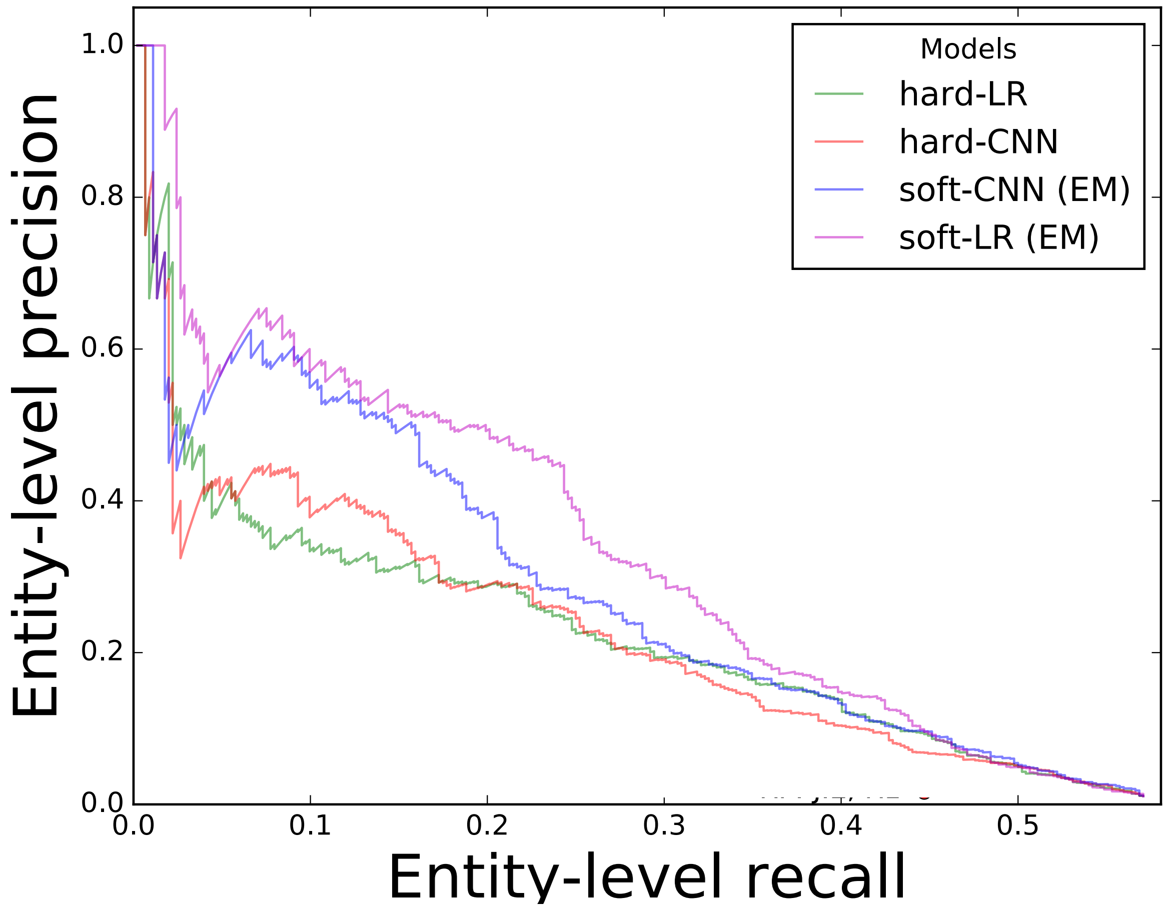

Model	AUPRC	F1
hard-LR, dep. feats.	0.117	0.229
hard-LR, n-gram feats.	0.134	0.257
hard-LR, all feats.	0.142	0.266
Data upper bound	0.57	0.73

Model results

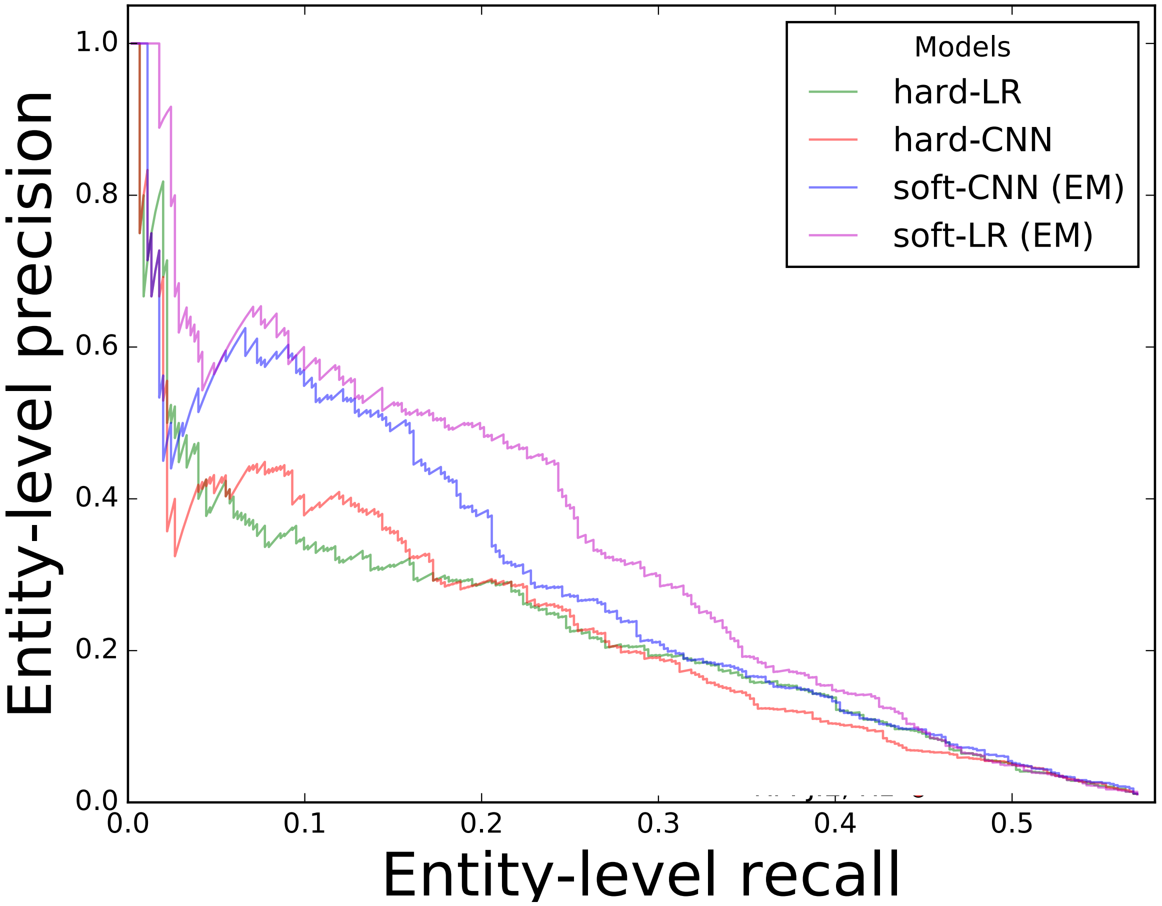

Model	AUPRC	F1
hard-LR, dep. feats.	0.117	0.229
hard-LR, n-gram feats.	0.134	0.257
hard-LR, all feats.	0.142	0.266
hard-CNN	0.130	0.252

Data upper bound	0.57	0.73

Model results

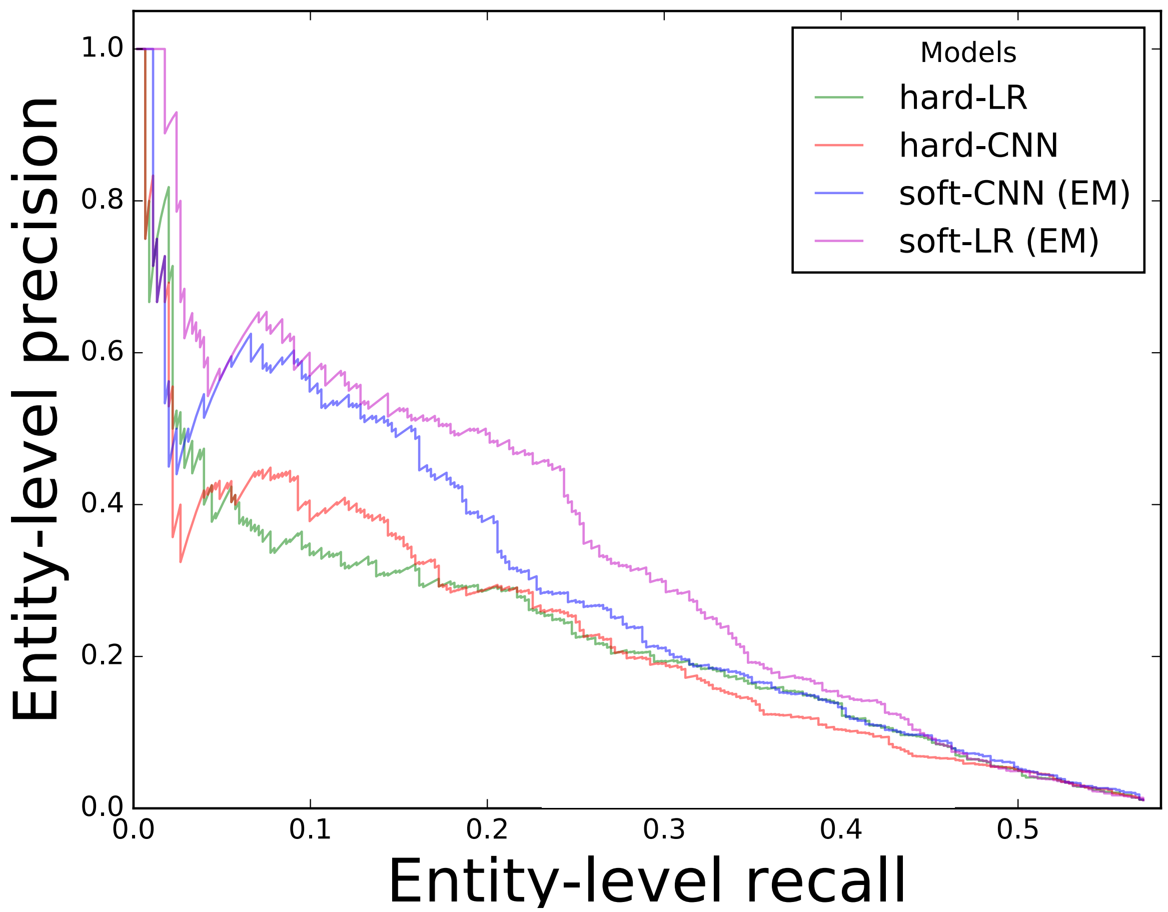

Model	AUPRC	F1
hard-LR, dep. feats.	0.117	0.229
hard-LR, n-gram feats.	0.134	0.257
hard-LR, all feats.	0.142	0.266
hard-CNN	0.130	0.252
soft-CNN (EM)	0.164	0.267
soft-LR (EM)	**0.193**	**0.316**
Data upper bound	0.57	0.73

Off-the-shelf event extractors

Off-the-shelf event extractors

SEMAFOR
(trained for FrameNet)
[Das et al. 2014]

Off-the-shelf event extractors

SEMAFOR
(trained for FrameNet)
[Das et al. 2014]

RPI-JIE
(trained for ACE)
[Li and Ji 2014]

Off-the-shelf event extractors

SEMAFOR
(trained for FrameNet)
[Das et al. 2014]

RPI-JIE
(trained for ACE)
[Li and Ji 2014]

Used in gun violence
database pipeline
[Pavlick and Callison-Burch 2016]

Off-the-shelf event extractors

	Rule	Prec.	Recall	F1
SEMAFOR (trained for FrameNet) *[Das et al. 2014]*				
RPI-JIE (trained for ACE) *[Li and Ji 2014]*				

Off-the-shelf event extractors

	Rule	Prec.	Recall	F1
SEMAFOR (trained for FrameNet) *[Das et al. 2014]*	R1	0.011	0.436	0.022
RPI-JIE (trained for ACE) *[Li and Ji 2014]*	R1	0.016	0.447	0.030

R1: killing event

Off-the-shelf event extractors

	Rule	Prec.	Recall	F1
SEMAFOR (trained for FrameNet) *[Das et al. 2014]*	R1	0.011	0.436	0.022
	R2	0.031	0.162	0.051
RPI-JIE (trained for ACE) *[Li and Ji 2014]*	R1	0.016	0.447	0.030
	R2	0.044	0.327	0.078

R1: killing event

R2: R1 and patient = entity

Off-the-shelf event extractors

	Rule	Prec.	Recall	F1
SEMAFOR (trained for FrameNet) *[Das et al. 2014]*	R1	0.011	0.436	0.022
	R2	0.031	0.162	0.051
	R3	0.098	0.009	0.016
RPI-JIE (trained for ACE) *[Li and Ji 2014]*	R1	0.016	0.447	0.030
	R2	0.044	0.327	0.078
	R3	0.172	0.168	**0.170**

R1: killing event

R2: R1 and patient = entity

R3: R2 and agent = police

Off-the-shelf event extractors

	Rule	Prec.	Recall	F1
SEMAFOR (trained for FrameNet) *[Das et al. 2014]*	R1	0.011	0.436	0.022
	R2	0.031	0.162	0.051
	R3	0.098	0.009	0.016
RPI-JIE (trained for ACE) *[Li and Ji 2014]*	R1	0.016	0.447	0.030
	R2	0.044	0.327	0.078
	R3	0.172	0.168	**0.170**
soft-LR (EM)				**0.316**

R1: killing event

R2: R1 and patient = entity

R3: R2 and agent = police

Top entities at test time

rank	name	positive	analysis
1	Keith Scott	true	
2	Terence Crutcher	true	
3	Alfred Olango	true	
4	Deborah Danner	true	
5	Carnell Snell	true	
6	Kajuan Raye	true	
7	Terrence Sterling	true	
8	Francisco Serna	true	
9	Sam DuBose	false	name mismatch
10	Michael Vance	true	
11	Tyre King	true	
12	Joshua Beal	true	
13	Trayvon Martin	false	killed, not by police
14	Mark Duggan	false	non-US
15	Kirk Figueroa	true	
16	Anis Amri	false	non-US
17	Logan Clarke	false	shot not killed
18	Craig McDougall	false	non-US
19	Frank Clark	true	
20	Benjamin Marconi	false	name of officer

Top entities at test time

rank	name	positive	analysis
1			
2			
3			
4			
5			
6			
7			
8			
9	Sam DuBose	false	name mismatch
10			
11			
12			
13	Trayvon Martin	false	killed, not by police
14	Mark Duggan	false	non-US
15			
16	Anis Amri	false	non-US
17	Logan Clarke	false	shot not killed
18	Craig McDougall	false	non-US
19			
20	Benjamin Marconi	false	name of officer

Outline

1. Motivation and overview
2. Task and data
3. Model
4. Training
5. Evaluation
6. Conclusion

Goal: database update

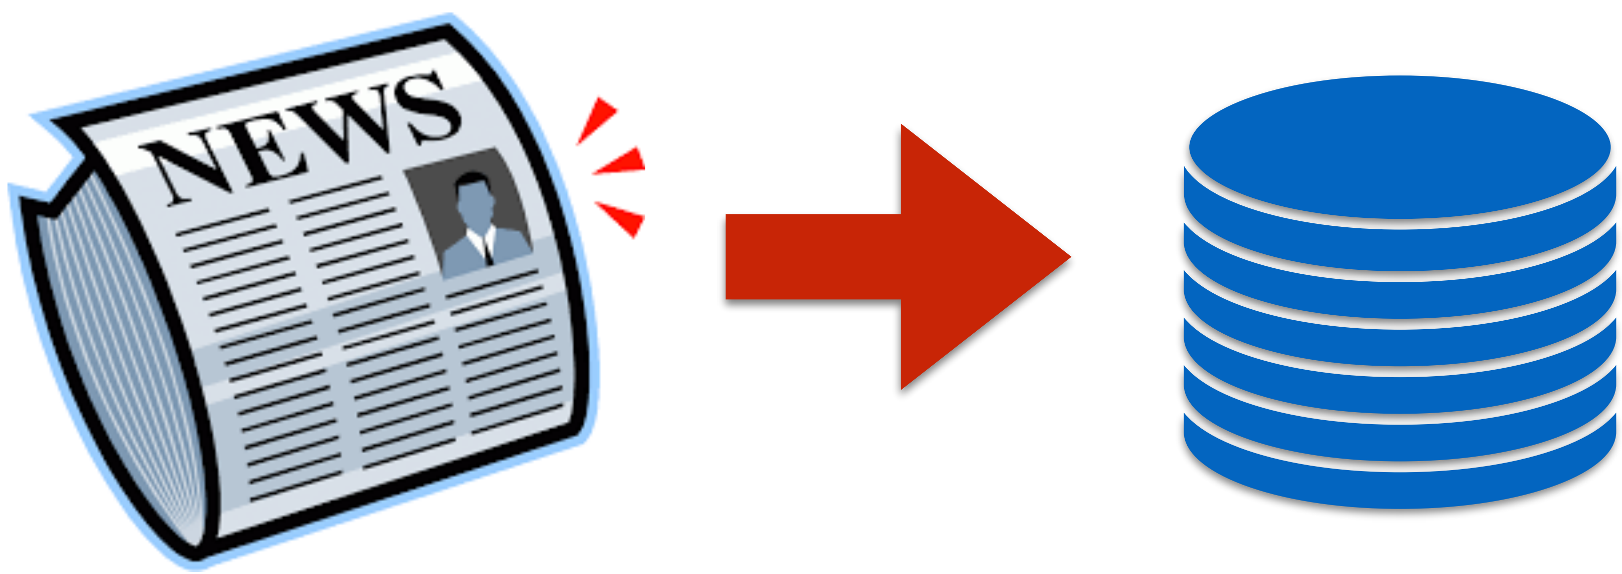

Sample Output

(1) **Walter Scott**

- A group prayer is held on April 12 , 2015 at the site where **Walter Scott** was killed by a North Charleston police officer in North Charleston , South Carolina View photos A group prayer is held on April 12 , 2015 at the site where **Walter Scott** was killed by a North Charleston police officer in North Charleston , South Carolina (AFP Photo/JOE RAEDLE) (BUTTON)
  dl date 2016-12-06 Doc 2173194_36_4 pred=0.998
- The shooting happened just months after **Walter Scott** , an unarmed black man , was killed by white police officer Michael Slager when he fled a traffic stop in North Charleston .
  dl date 2016-12-16 Doc 2203135_323_0 pred=0.991
- A man walks past the lot where **Walter Scott** was killed by a North Charleston police officer Saturday after a traffic stop in North Charleston , S.C. , Thursday , April 9 , 2015 .
  dl date 2016-12-06 Doc 2172211_194_0 pred=0.99

(2) **Keith Scott**

- News of the jury 's failure to reach a verdict came just a few days after a prosecutor in Charlotte , N.C. , announced no charges would be filed against a police officer in the September shooting of **Keith Scott** , an African American man whose death inspired violent protests in North Carolina .
  dl date 2016-12-02 Doc 2163436_27_0 pred=0.97
- Nation/World Keith Lamont Scott , pictured at right in a photo released by his family , was fatally shot by police in Charlotte , North Carolina on Sept. 20 , 2016 .
  dl date 2016-12-02 Doc 2163074_100_0 pred=0.951
- People march in Charlotte , N.C. , on Sept. 23 to protest the fatal police shooting of Keith Lamont Scott .
  dl date 2016-12-20 Doc 2213883_298_0 pred=0.947

(3) **Alton Sterling**

- Hundreds of miles away , protesters marched outside a convenience store in Baton Rouge , Louisiana , where **Alton Sterling** was fatally shot Tuesday while police tackled him in a parking lot .
  dl date 2016-12-29 Doc 2241447_83_0 pred=0.995
- [rtsh3xr.jpg?quality=80&strip=all&w=50] Ieshia L. Evans , a demonstrator protesting the shooting death of **Alton Sterling** is detained by law enforcement near the headquarters of the Baton Rouge Police Department in Baton Rouge , Louisiana , on July 9 .
  dl date 2016-12-27 Doc 2234040_59_0 pred=0.995
- old **Alton Sterling** , a black man killed by white Baton Rouge officers after a confrontation at a convenience store .
  dl date 2016-12-27 Doc 2235302_71_0 pred=0.995

Future Work

- Other model architectures (e.g. LSTMs)
- Other domains for database update problem
- Extract additional event information
- Build interactive interface for practitioners

Contributions

- Distant supervision approach much cheaper
- Public data for the social good
- New NLP task, released data publicly
- Progress towards fully-automatic system

Thanks!

Code and data:

http://slanglab.cs.umass.edu/PoliceKillingsExtraction/

Acknowledgements:

- Amazon Web Services (AWS) Cloud Credits for Research program.
- D. Brian Burghart for advice on police fatalities tracking