Computational Fact Checking

Paolo Papotti

EURECOM
Sophia Antipolis

EDBT-INTENDED Summer School - 9th July 2022

Outline

- Motivation
- Pipeline
  - Worth checking → Evidence → Verification
- Data-driven verification
- Role of humans
- Ready for adoption?

Disclaimer

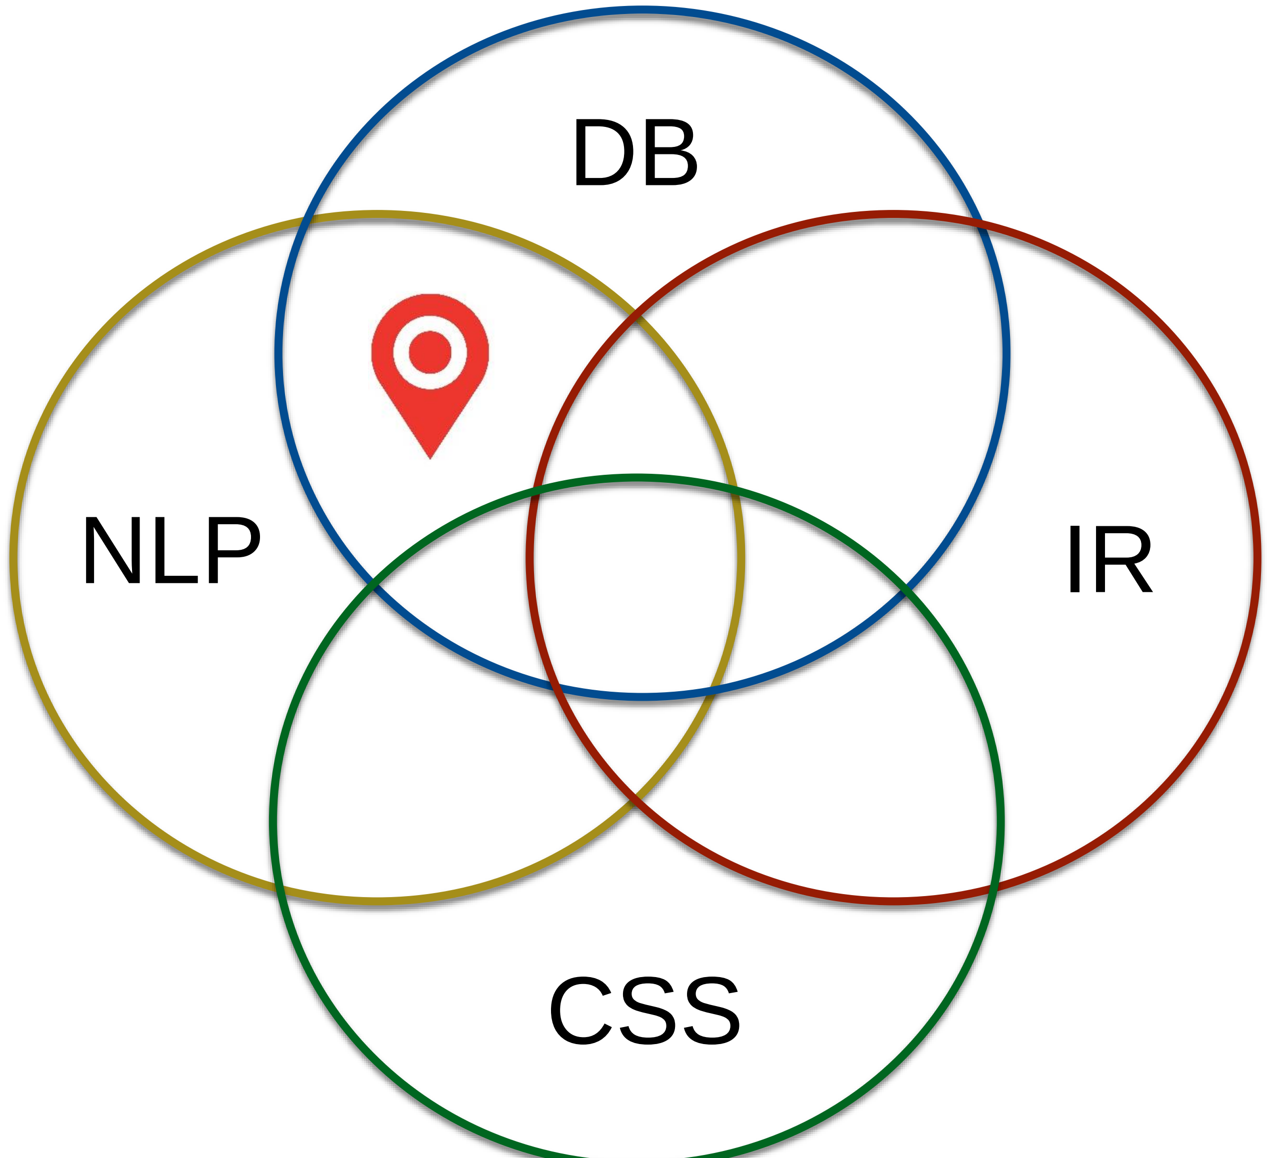

Fro
Ne

Cameron Harris i
electrical worker
Demczuk for The N

INDEP

Voices

Fake n
referen
what t

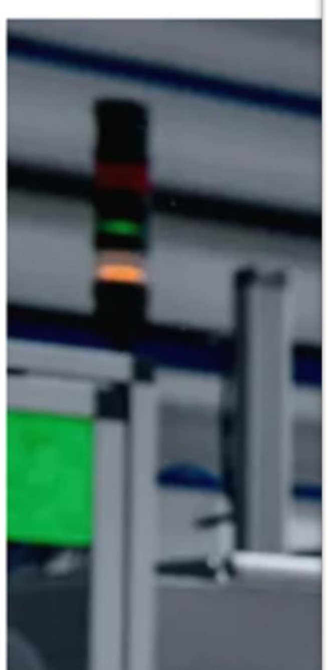

nytimes.con

Faceboo
Violence

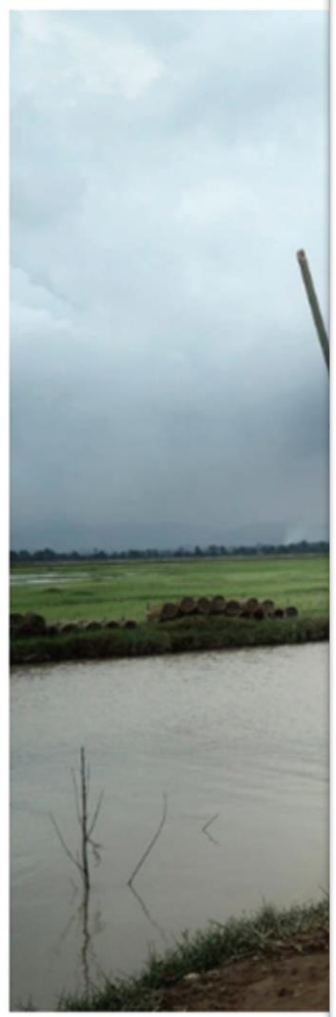

Rohingya refugees
Bangladesh, in 2017
keep its platform fr
Myanmar. Adam De

By Alexandra Steve

Nov. 6, 2018

ncbi.nlm.nih.gov/pmc/articles/PMC6657116/

Journal List › Glob Pediatr Health › v.6; 2019 › PMC6657116

Global Pediatric Health

Glob Pediatr Health. 2019; 6: 2333794X19862949.
Published online 2019 Jul 24.
doi: 10.1177/2333794X19862949

PMCID: PMC6657116
PMID: 31384629

Anti-Vaccine Decision-Making and Measles Resurgence in the United States

Olivia

Au

This

Intro

Meas
to otl
and J
psycl
State
resea
healt
for th

Epic

Durin
State

Glob Pediatr Health

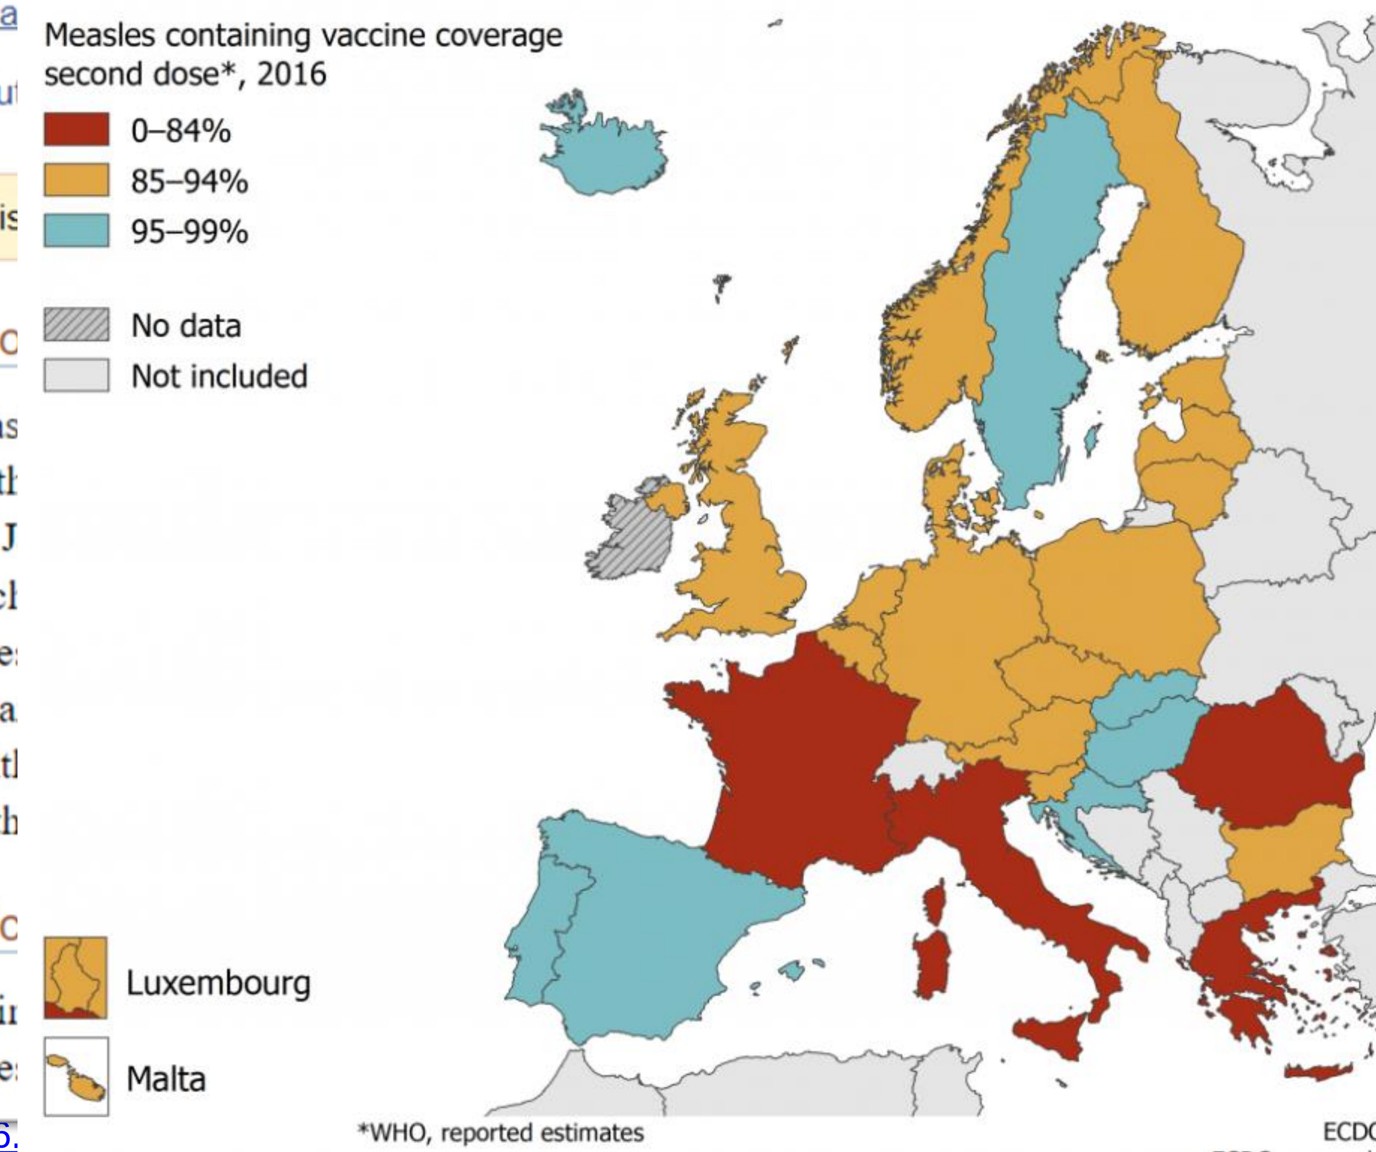

https://www.nytimes.com/2017/01/18/us/fa
news-hillary-clinton-cameron-harris.html?_r=0

gove-boris-johns
theresa-may-a7533806.

COVID-19 Infodemic

“The outbreak has been accompanied by a massive *infodemic* - an over-abundance of information – some accurate and some not – that makes it hard for people to find trustworthy sources and reliable guidance when they need it.”

World Health Organization

who.int/emergencies/diseases/novel-coronavirus-2019/... Incognito

To help alleviate suffering and save lives, WHO has been working night and day in five key ways:

1. Helping build countries' capacity to prepare and respond

2. Providing accurate information and fight the infodemic, together with numerous partners

3. Ensuring supplies of essential medical equipment for frontline health workers.

4. Training and mobilizing health workers.

5. Accelerating research and development.

https://www.who.int/emergencies/diseases/novel-coronavirus-2019/events-as-they-happen

Checking claims

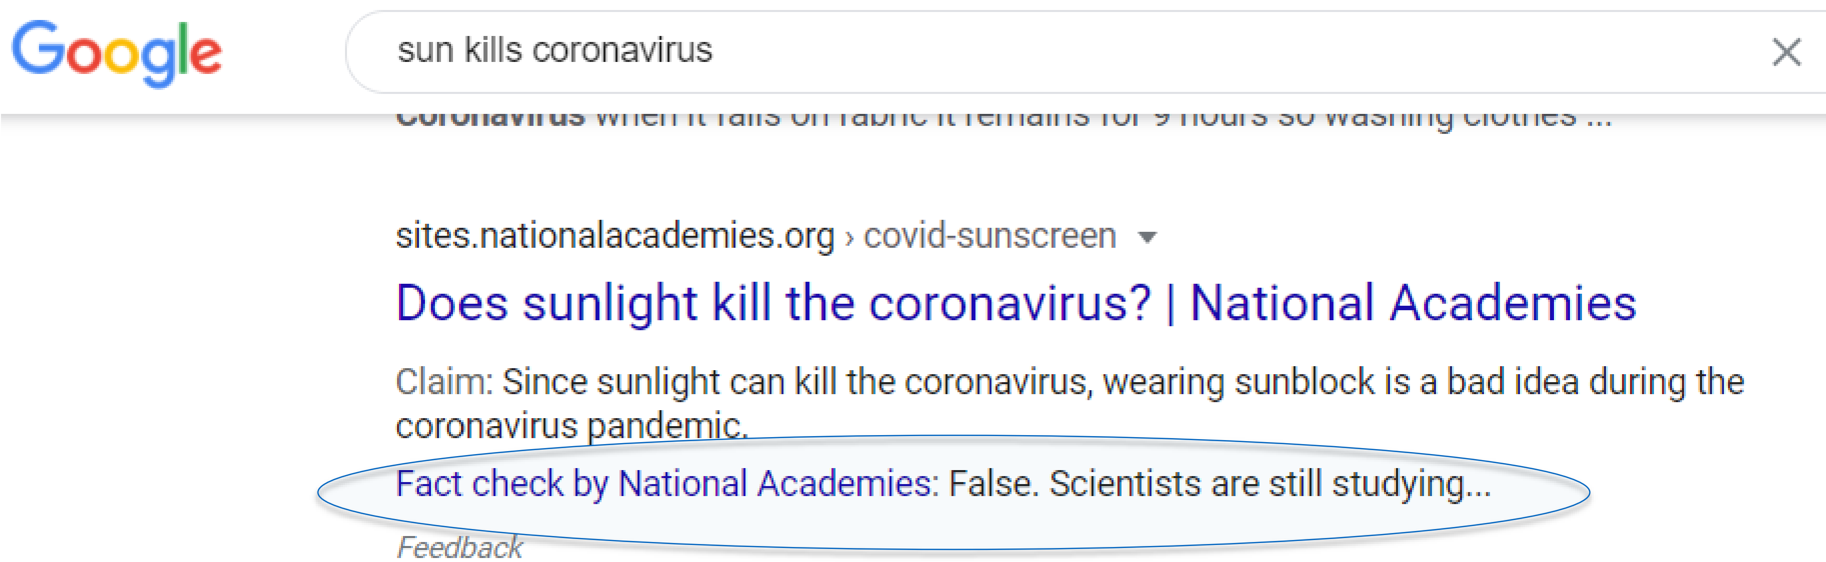

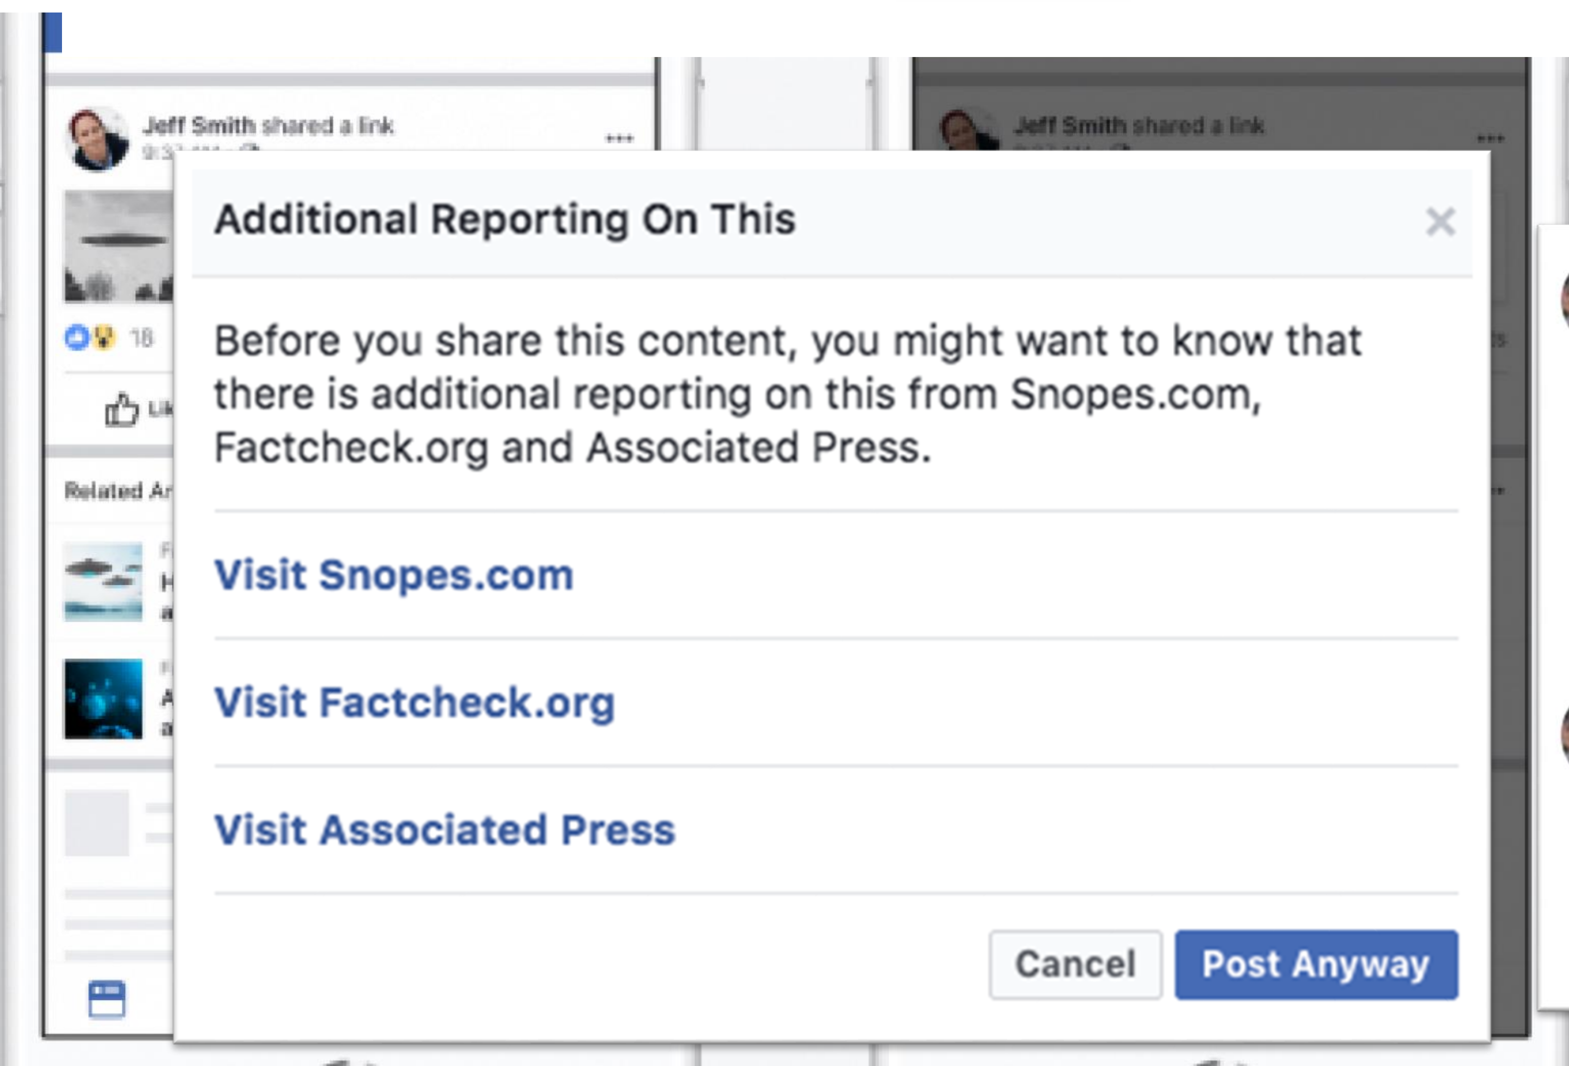

Fact checking entities

- Expert / Journalist
- Computational
- Crowdsourcing with end users

Expert based fact checking

BirdWatch Example
@birdwatchexample

Trump won the election by a landslide.

9:34 AM · Dec 17 2021 · Twitter Web App

96 Retweets **88** Quote Tweets **153** Likes

Fact-Check Nov 07

Claim: Donald Trump won the 2020 election, by a lot.

Verdict: Not Credible

Fact Checker: Lead Stories

Country: United States

Link:https://leadstories.com/hoax-alert/2020/11/fact-c ald-j-trump-on-twitter-I-won-by-a-lot.html

https://www.claimreviewproject.com/

Fact checking

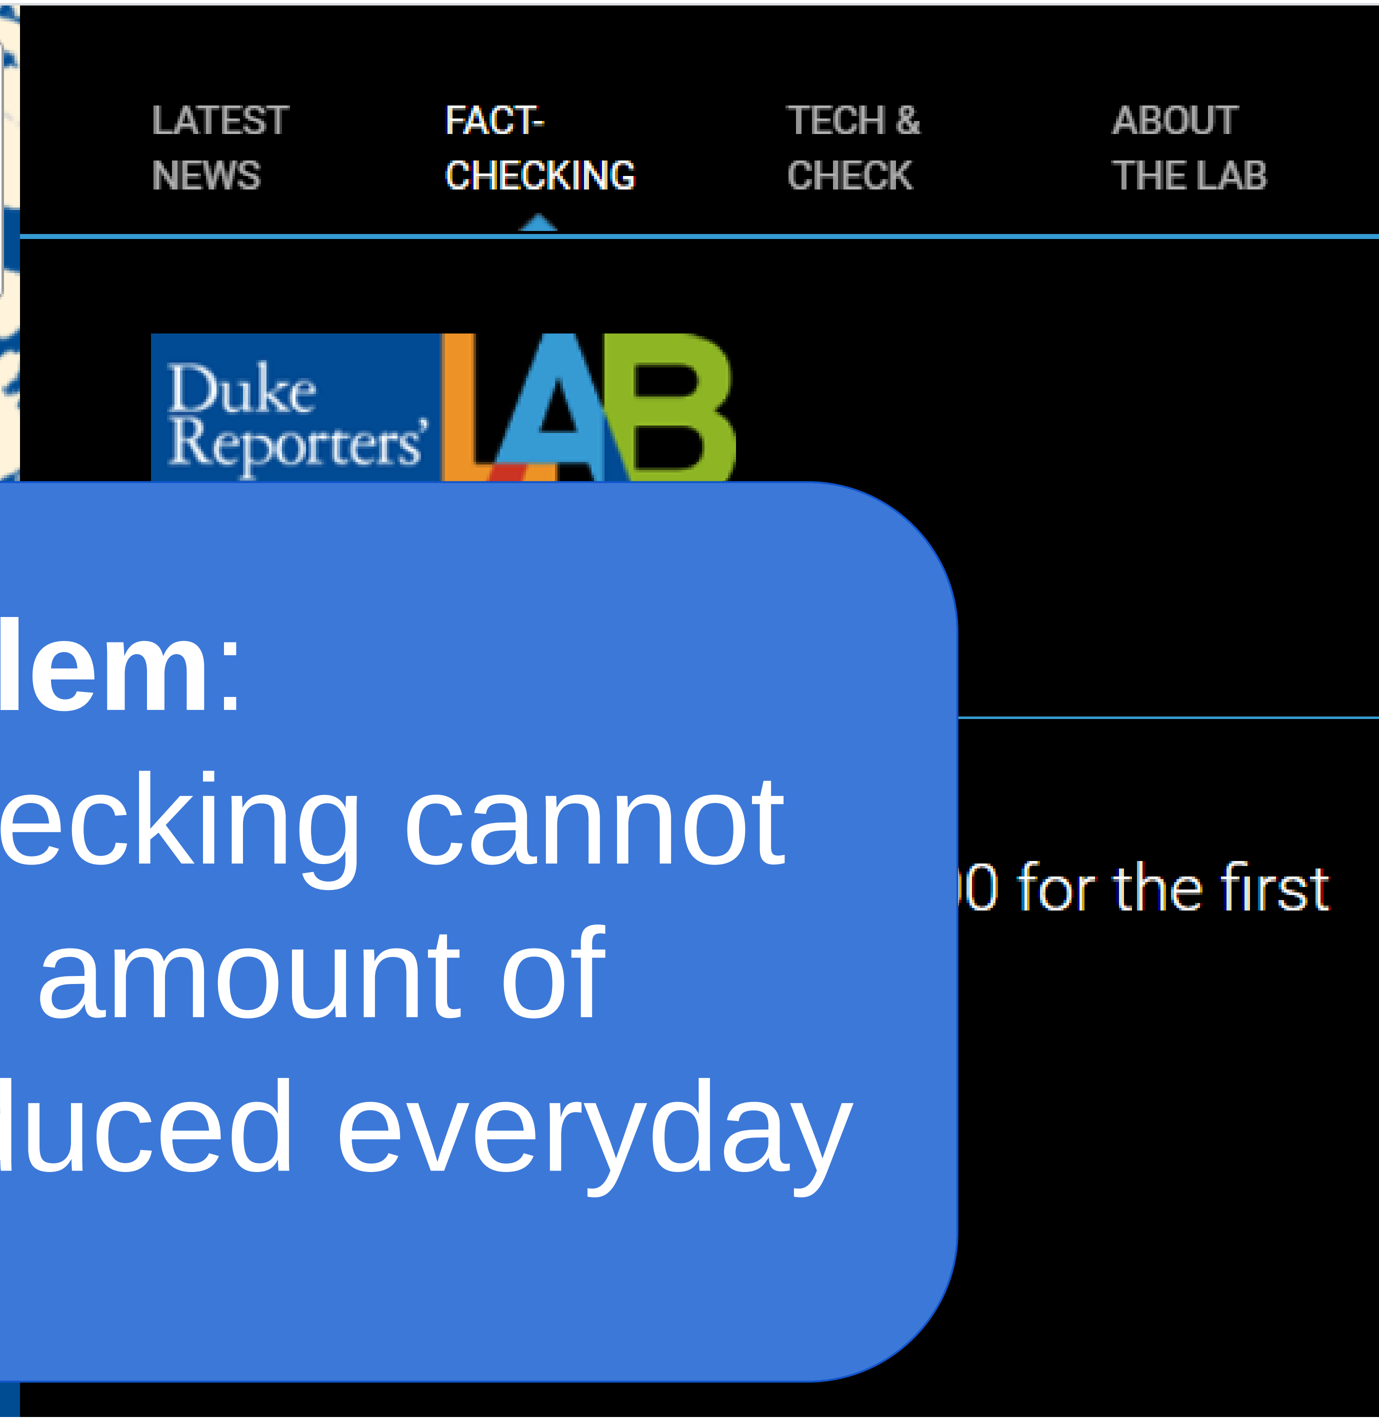

Problem:
Manual fact checking cannot scale to the amount of information produced everyday

https://www.poynter.org/coronavirusfactsalliance/

THE WALL STREET JOURNAL.

English Edition ▾ | May 25, 2020 | Print Edition | Video

Home World U.S. Politics Economy Business Tech Markets Opinion Life & Arts Real Estate WSJ. Magazine

VIRUS
es

FINDING A JOB | BACK TO WORK Q&A | GOING OUTSIDE SAFELY | CAN I GET SICK AGAIN? | CHILDCARE & SUMMER CAMPS | CAN I TRAVEL? | STATE REOPENINGS | TESTS & TREATMENTS | WHAT 6 FEET LOOKS LIKE | EN ESPAÑOL

TECH

Facebook's Fact Checkers Fight Surge in Fake Coronavirus Claims

Contractors battle bogus assertions about canine vaccines and free baby formula; 'We've maxxed out'

RECOMMENDED VIDE

1. Coronavirus Updat China Scraps GDP Target, Trump's Swing-State Tour

https://www.wsj.com/articles/facebooks-fact-checkers-fight-surge-in-fake-coronavirus-claims-11585580400

``**up to 22** days for the platform to downgrade and issue warning labels on **harmful misinformation content**''

https://secure.avaaz.org/campaign/en/facebook_coronavirus_misinformation/

Expert based fact checking

Advantages

- High credibility
- Can handle nuanced claims
- Can produce detailed evidence of fact checking

Disadvantages

- Not scalable: few claims per day
- Slow: average fact checking time >7 days
- Possible human bias
- Not easy for niche domains

Fact checking entities

- Expert / Journalist
- Computational
- Crowdsourcing with end users

Computational fact checking

- The general problem:
  Given some content, assess if it is true or false

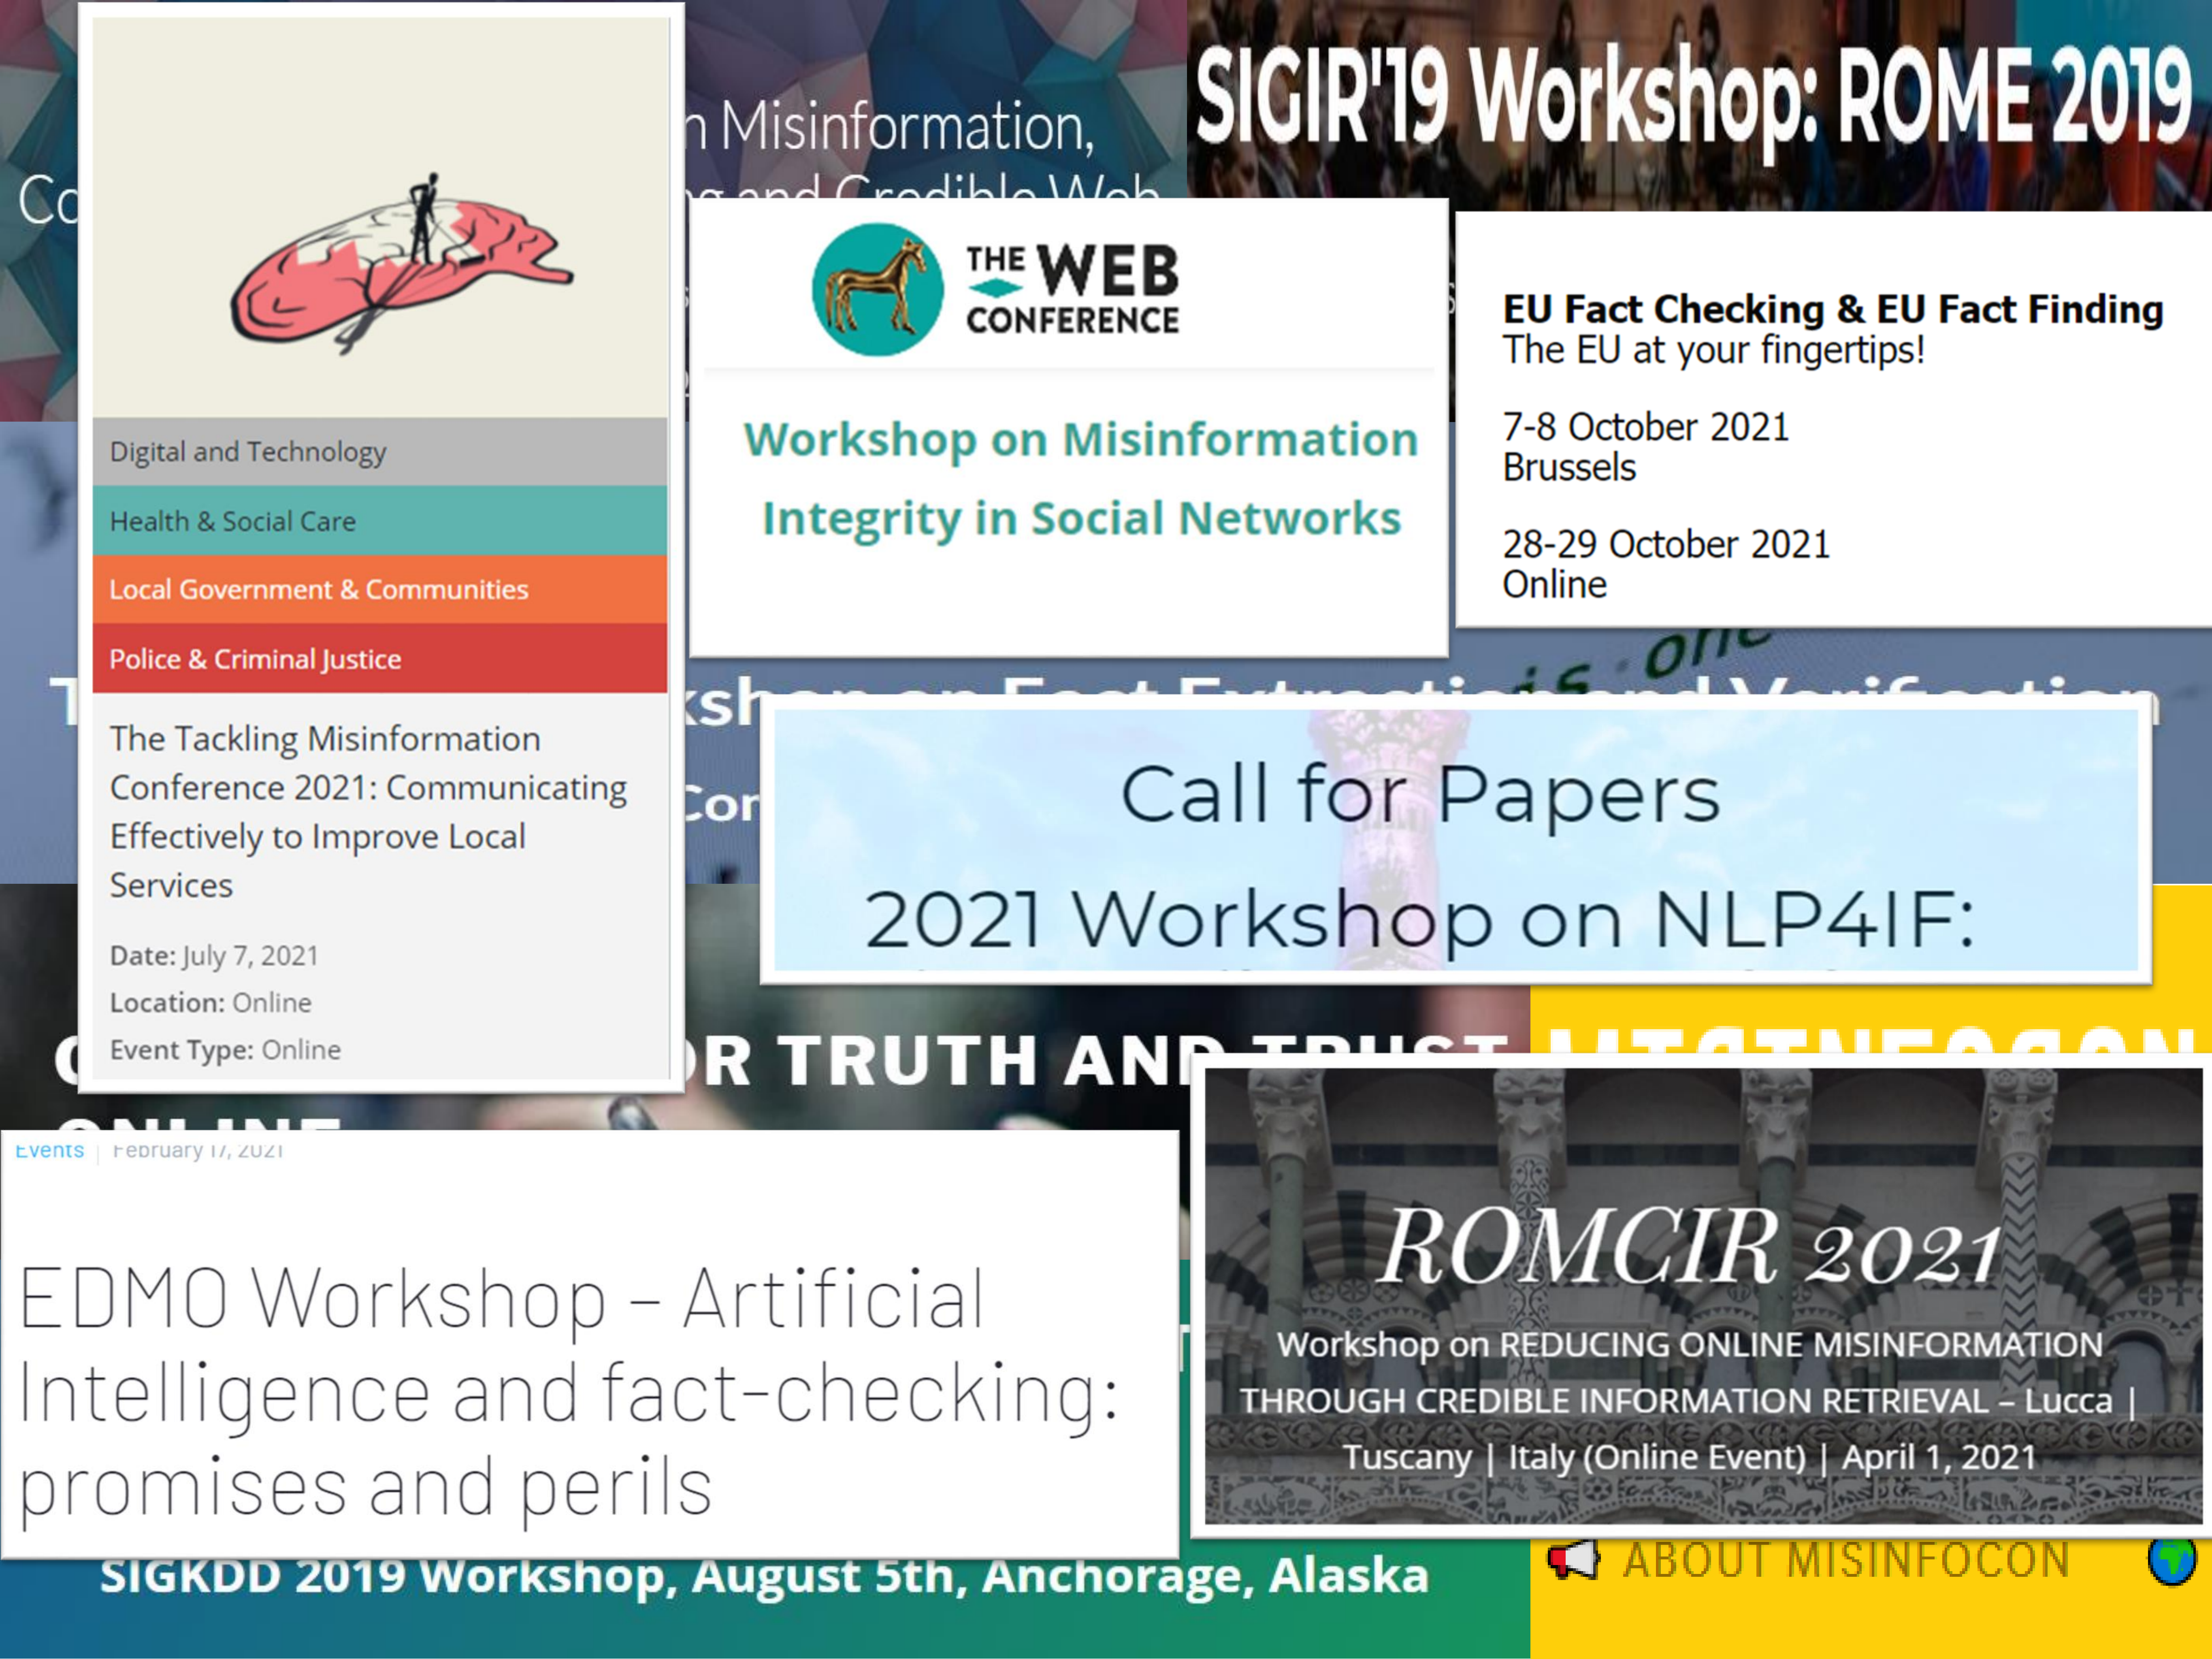
SIGIR'19 Workshop: ROME 2019
Digital and Technology
Health & Social Care
Local Government & Communities
Police & Criminal Justice
The Tackling Misinformation Conference 2021: Communicating Effectively to Improve Local Services
Date: July 7, 2021
Location: Online
Event Type: Online
THE WEB CONFERENCE
Workshop on Misinformation Integrity in Social Networks
EU Fact Checking & EU Fact Finding
The EU at your fingertips!
7-8 October 2021
Brussels
28-29 October 2021
Online
Call for Papers
2021 Workshop on NLP4IF:
EDMO Workshop – Artificial Intelligence and fact-checking: promises and perils
ROMCIR 2021
Workshop on REDUCING ONLINE MISINFORMATION THROUGH CREDIBLE INFORMATION RETRIEVAL – Lucca | Tuscany | Italy (Online Event) | April 1, 2021
SIGKDD 2019 Workshop, August 5th, Anchorage, Alaska
ABOUT MISINFOCON

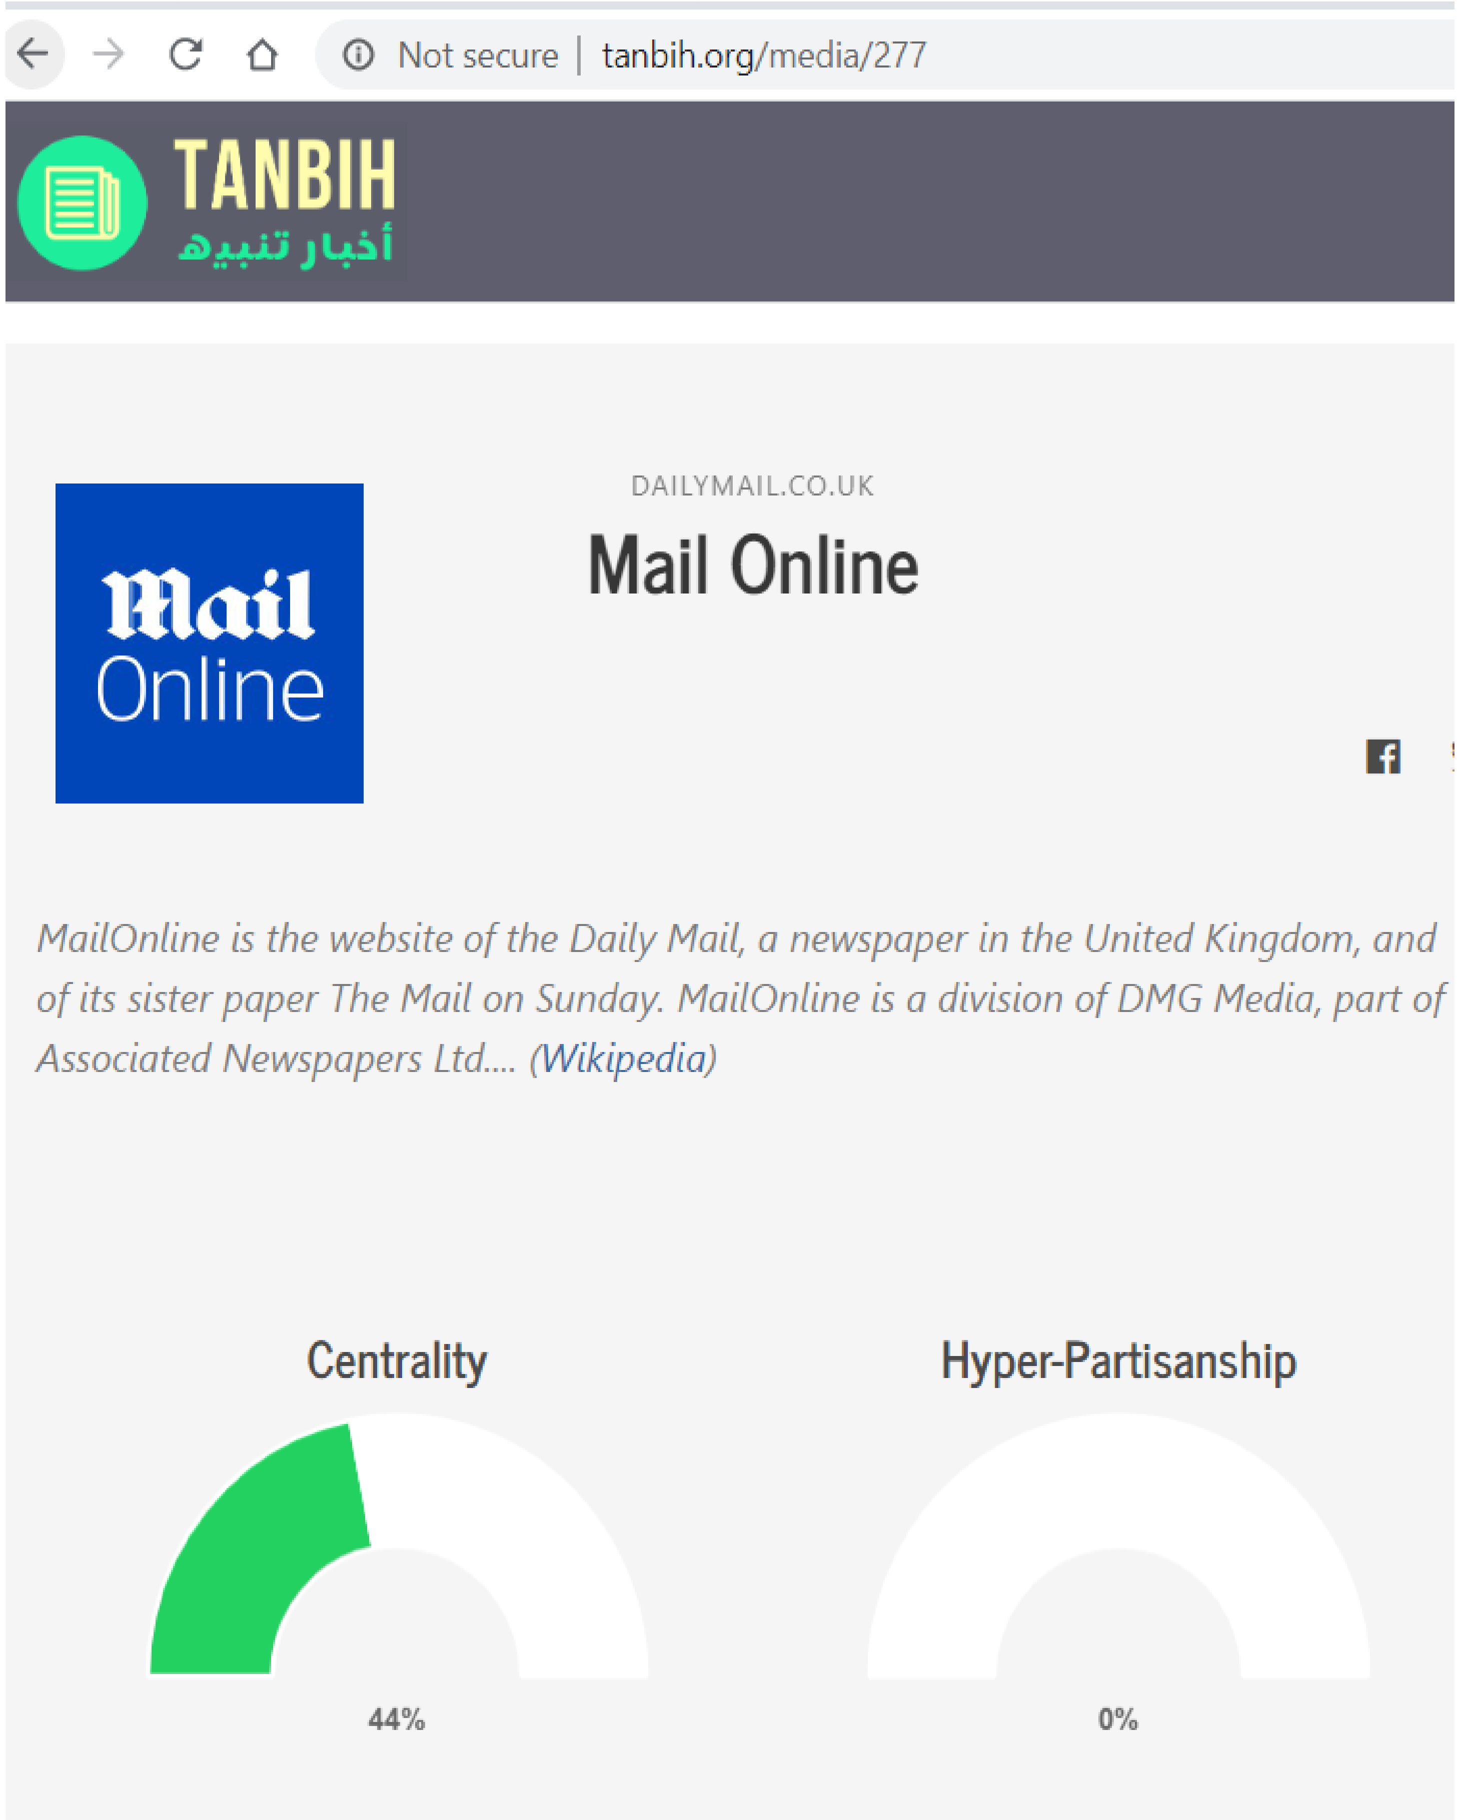

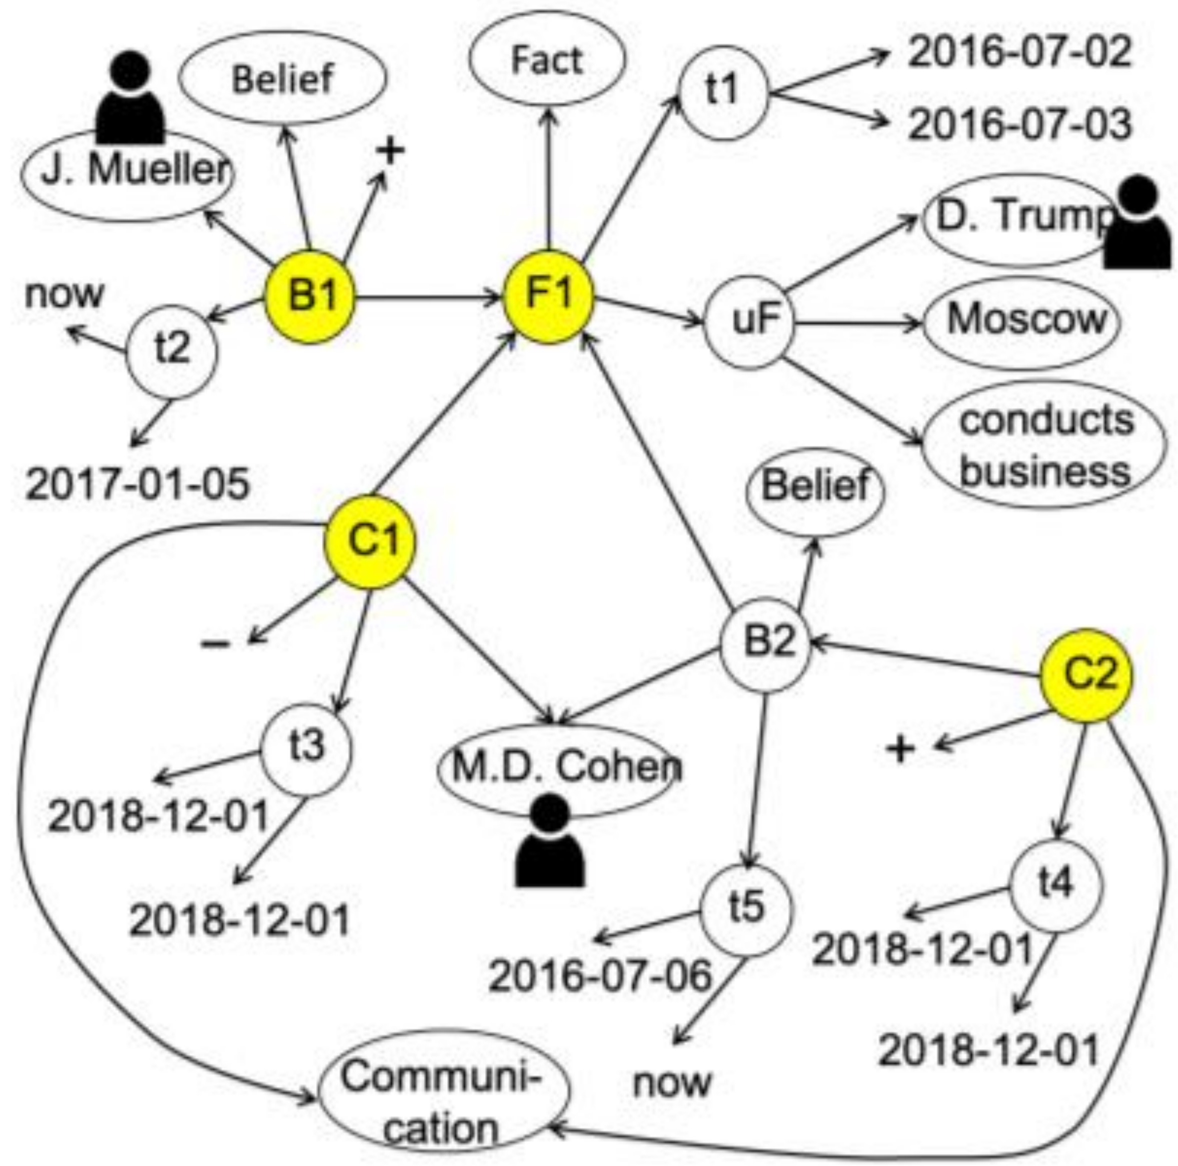

https://www.tanbih.org/

[Duroyon et al, 2019]

Computational pipeline

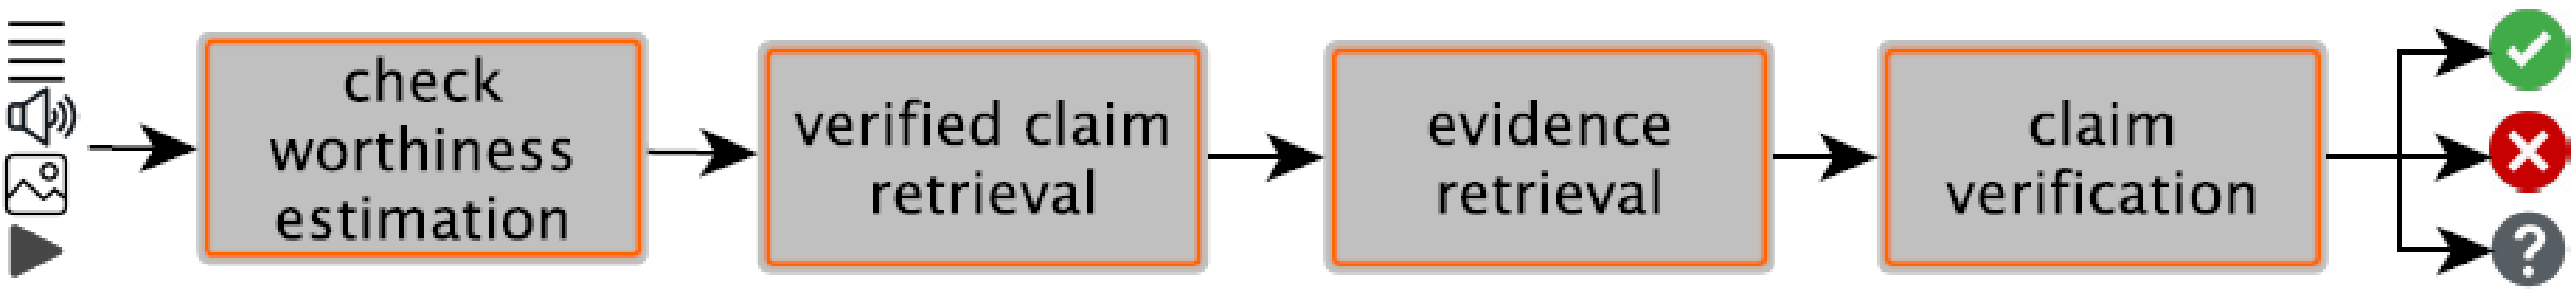

1. balance potential harm vs effort to check claim
2. quickly detect if viral claims have already been fact-checked (possibly in other language)
3. automatically retrieve relevant data from trusted sources
4. correct or incorrect
   but also partially correct, not enough evidence, out-of-context…

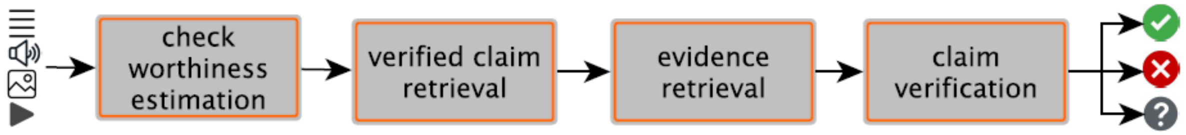

Find claims worth checking

- System produces a check-worthiness that provides fact-checkers ability to prioritize claims
- ClaimBuster [Hassan et al., 2017]
  - trained on manually annotated sentences (nonfactual, unimportant factual, checkworthy) features based on sentiment, named entities, part-of-speech tags, words, and claim length
- More models based on pre-trained transformers such as BERT/RoBERTa [Wright and Augenstein, 2020]

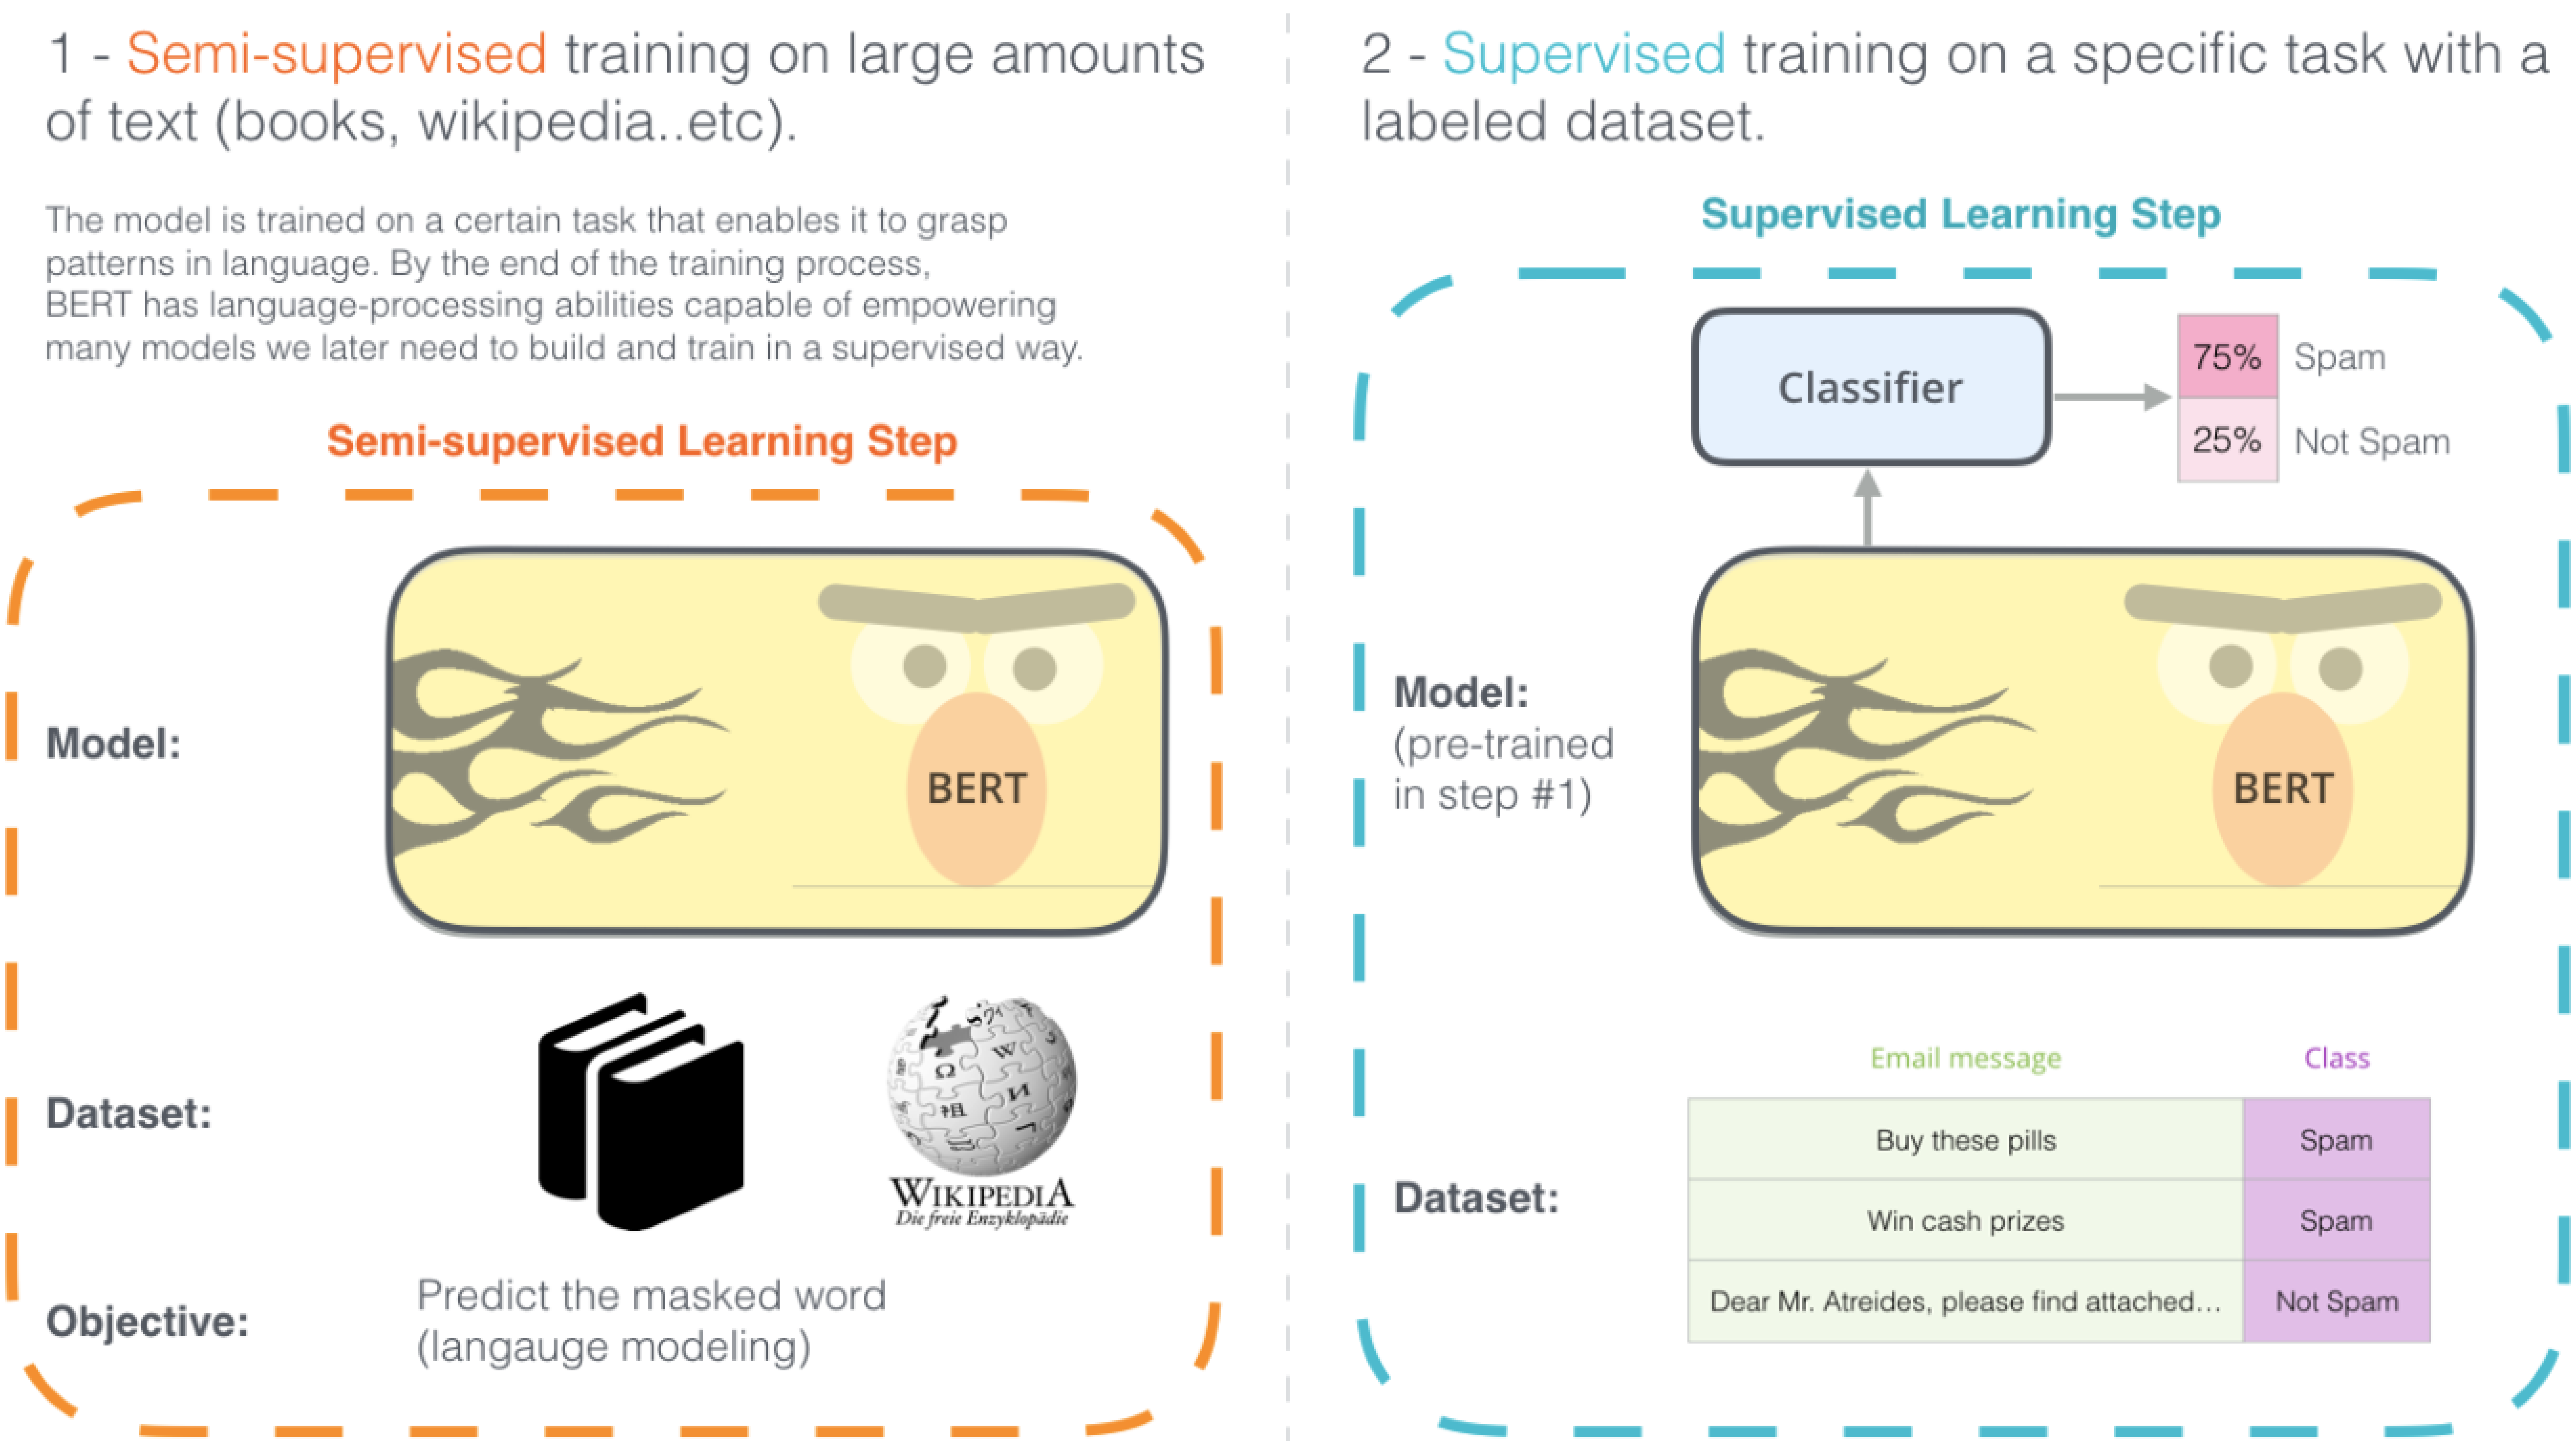

Fig. credits: https://jalammar.github.io/illustrated-bert/

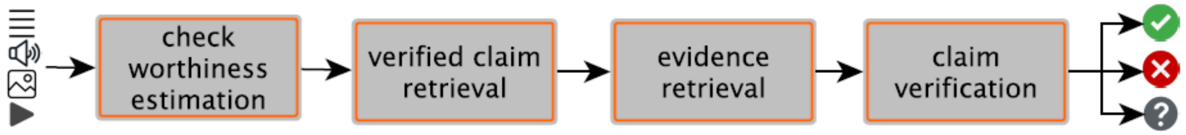

Detect previously checked claims

- Misleading claims repeated in different channels and languages, independently of existing fact-checks published with ClaimReview
- Once a claim established as misleading, the spread of its variants could be minimized with **rapid** detection
- [Shaar et al. 2020] matching with BERT and BM25 Datasets:
  - tweets, which are to be compared to claims in Snopes
  - political debates, to be matched to claims in PolitiFact

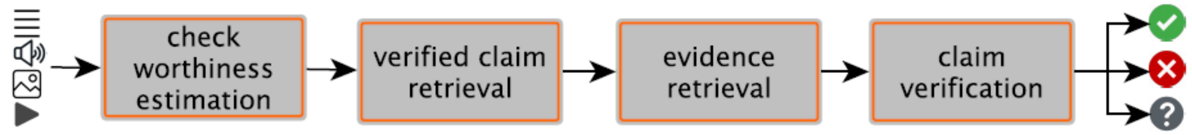

Evidence retrieval

- Find external evidence to help human/system checkers decide factuality of input claim
- Evidence lie in large text documents, audio-visual recordings and streams, or in multiple languages
- Given claim and a (closed) data collection (1) rank relevant objects (docs), or (2) extract specific pieces of evidence, e.g., a text snippet or a recording
  - FEVER dataset [Thorne et al., 2018] + BM25/NER/BERT

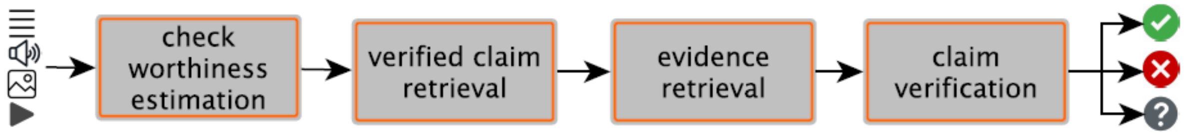

Automated verification

- Goal: internal tools presenting *evidence*, (*reasoning*) and *conclusion* regarding a claim, before the (human) fact-checker publishes article

- Non-explainable vs Explainable

  - + “General” claims, “Robust” to noise in data
    - Just a label as output

  - + Evidence of the arguments that support/refute
    - Stricter assumptions on claim and data

Standard “black box” verification

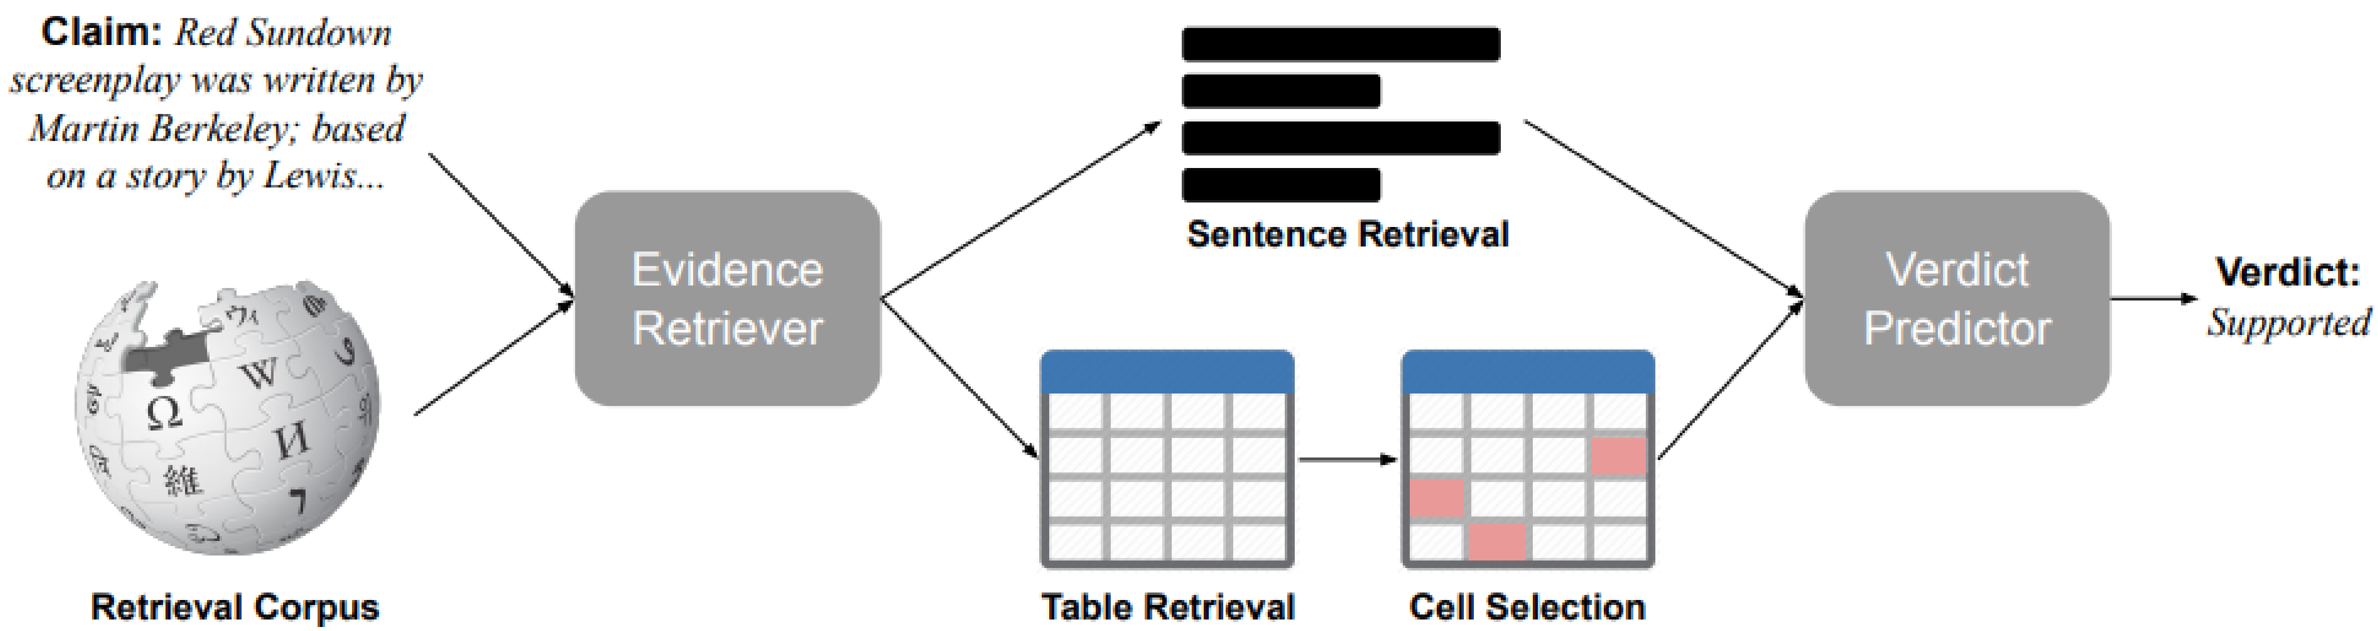

- Gather evidence + classifier (true/false/NEI)
- Can check “any” textual claim, can tolerate noise in corpus

[Thorne et al., 2018] [Aly et al, 2021]

Standard “black box” verification

Claim: Red Sundown screenplay was written by Martin Berkeley; based on a story by Lewis B. Patten, who often published under the names Lewis Ford, Lee Leighton and Joseph Wayne.

Evidence:

Page: `wiki/Red_Sundown`

$\mathbf{e}_1$(Introduction):

Red Sundown	
Directed by	Jack Arnold
Produced by	Albert Zugsmith
Screenplay by	Martin Berkeley
Based on	Lewis B. Patten
...	

Page: `wiki/Lewis_B._Patten`

$\mathbf{e}_2$(Introduction): He often published under the names Lewis Ford, Lee Leighton and Joseph Wayne.

Verdict: Supported

Well formed, factual, not ambiguous sentence A_C

High quality relevant sources A_S

[Aly et al, 2021]

Standard “black box” verification

Final Leaderboard

Rank	Team	FEVEROUS score	Accuracy	Evidence F1	Evidence Precision	Evidence Recall
1	Bust a move!	0.2701	0.5607	0.1308	0.0773	0.4258
2	Papelo	0.2592	0.5757	0.1187	0.0716	0.3460
3	NCU	0.2514	0.5229	0.1581	0.0991	0.3907
4	Z team	0.2251	0.4901	0.1312	0.0776	0.4264
5	EURECOM_Fever	0.2001	0.4779	0.1952	0.1373	0.3373

“Goal: internal tools presenting *evidence*, (*reasoning*) and *conclusion* regarding a claim, before the (human) fact-checker publishes article”

https://fever.ai/2021/task.html

Computational fact checking

- General problem: Given some content, assess if it is true or false
- Focus today:
  Given some **textual content** and some reference **data**, assess if the content is **true or false** and explain **why**
  - **≠ identifying check-worthy claims** [ClaimBuster]
    ≠ matching claims to existing checks [FullFact, ClaimReview]
    ≠ model trust of the sources [Tanbih]
    ≠ propagation, mitigation and intervention ["Combating Fake News: A Data Management and Mining Perspective" tutorial VLDB19]
- Effective, Scalable, **Interpretable**

Fact checking

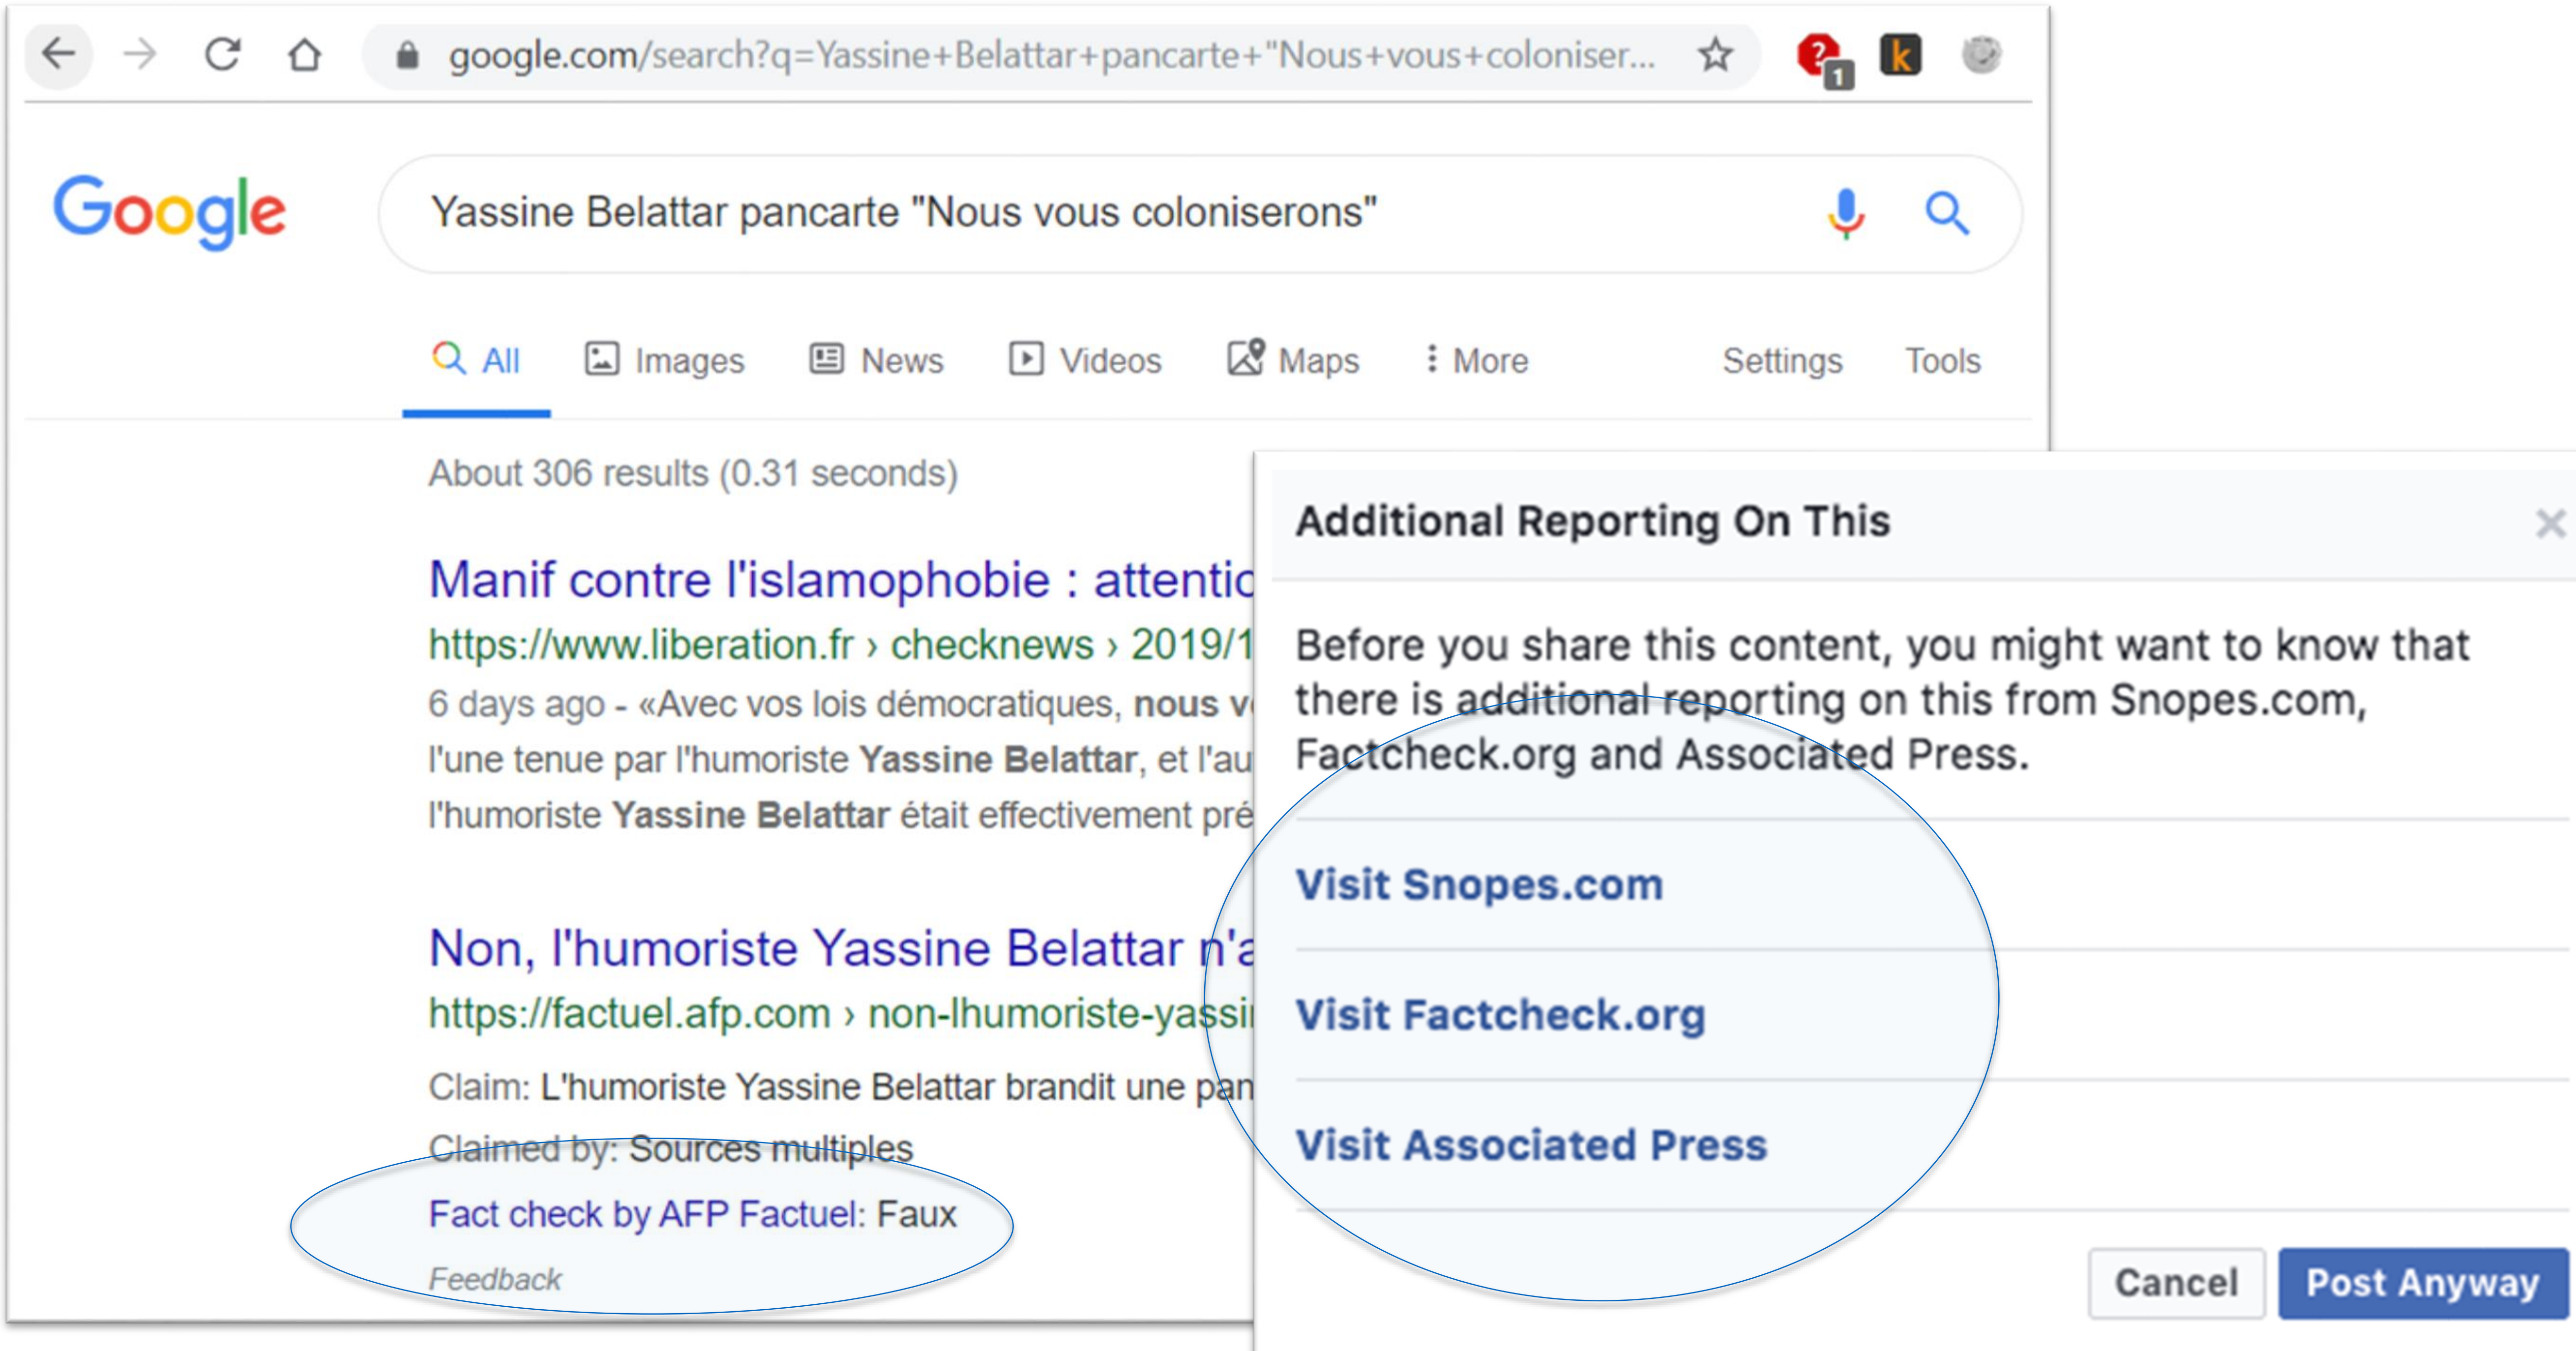

Computational fact checking

- Given some **textual content** and some reference **data**, assess if the content is **true or false** and explain **why**

  - Check(claim,data) =
    true/false label,
    confidence value,
    description of the subset of the data that implies the decision

- **Property** and statistical claims

Checking property claims

- Input textual claim:
  “Elon Musk is the founder of Chevrolet”
- Output: FALSE because
  “Chevrolet *was founded* in **1911** and Elon Musk *was born* in **1971**”

Checking property claims

Not in data. Count occurrences/text patterns/… on Web

…vrolet”

- Output: FALSE because
“Chevrolet *was founded* in **1911** and Elon Musk *was born* in **1971**”

Concepts of person, company, birth year, founding year, logical contradiction

Checking property claims

- Input textual claim:
  “Michael White’s alma mater is UT Austin”
- Output: FALSE because
  “M. White *works* at UT Austin and
  has Abilene Christian University and Yale
  Divinity School as *alma mater*”

Checking property claims

- Input
  - Textual claim: “Michael White’s alma mater is UT Austin” → *almaMater*(Michael White, UT Austin)

A_C

RDF Knowledge Graphs

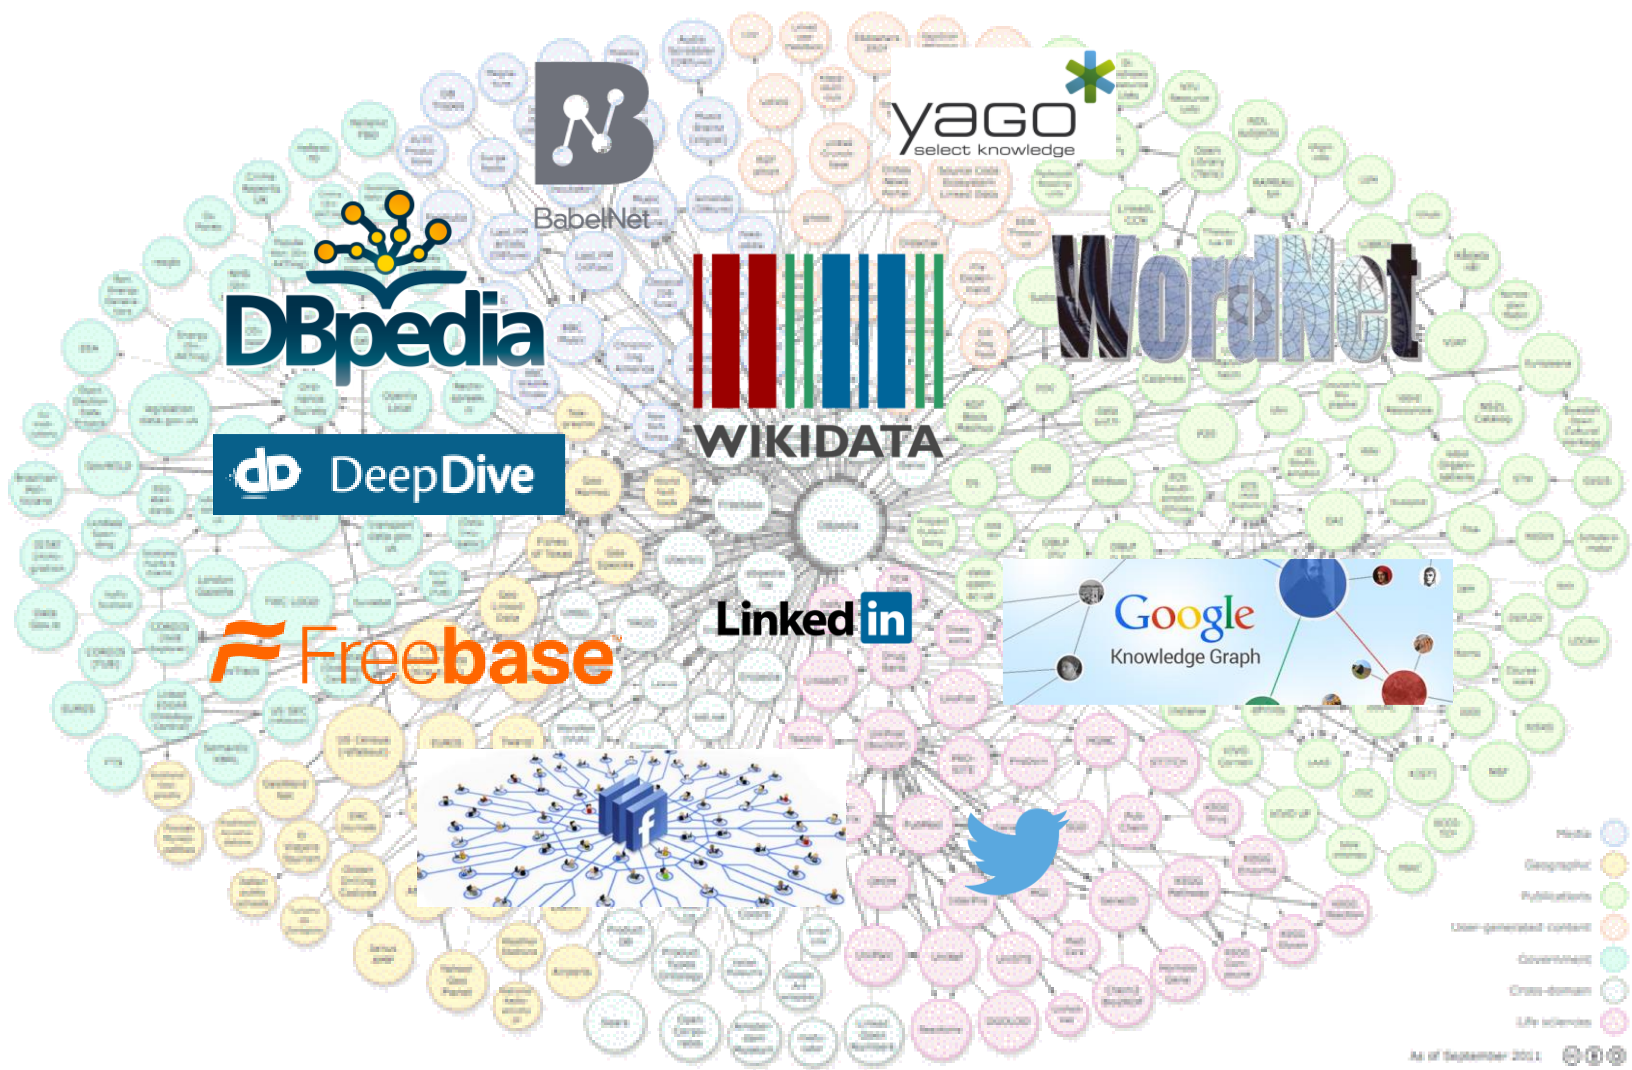

Enterprise: Walmart, Amazon, KPMG, …

RDF KGs

<Barack Obama> <spouse> <Michelle Obama> .

<Barack Obama> <birthDate> “1961-08-04” .

<Michelle Obama> <birthPlace> <Illinois> .

SUBJECT PREDICATE OBJECT

Name	# Entity types	# Entity instances	# Relation types
Knowledge Vault (KV)	1100	45M	4469
DeepDive [21]	4	2.7M	34
NELL [6]	271	5.19M	306
PROSPERA [20]	11	N/A	14
YAGO2 [16]	350,000	9.8M	100
Freebase [4]	1,500	40M	35,000
Knowledge Graph (KG)	1,500	570M	35,000

A_S

[Dong and Srivastava, 2015]

Computational fact checking

- Given some **textual content** and some reference **data**, assess if the content is **true or false** and explain **why**

  - Check(claim, *incomplete KG*) =
    true/false label,
    confidence value,
    description of the subset of the data that implies the decision

OWA

→ how to assess/explain if claim is **NOT** in KG?

Fact checking with KGs

- Structure-based: KL [1], KG-Miner [2], SFE [3]
- exploit the topological structure: fact modeled by predicate paths/proximity between subject and object

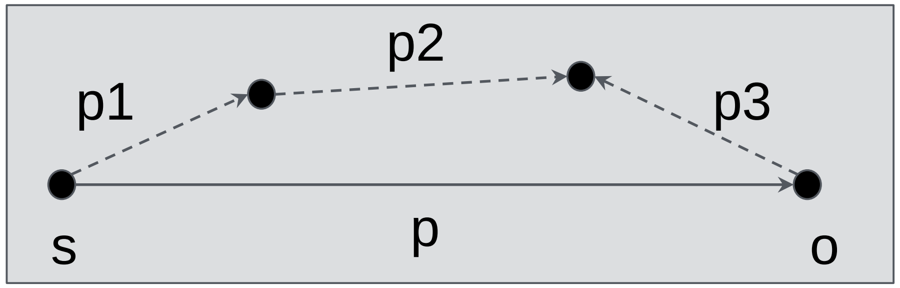

- Embeddings: TransE [4]
- relation in graph interpreted as a translation in a low dimensional vector space: check s+p=o

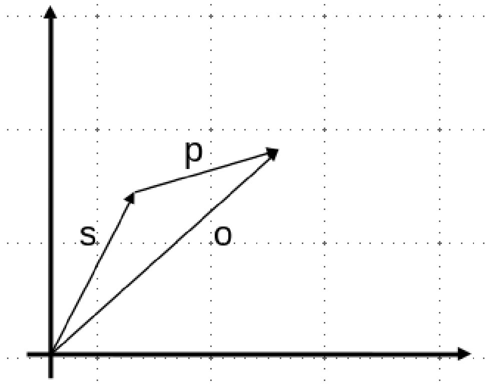

[1] Ciampaglia et al., PloS one, 2015 [2] Shi et al., AAAI, 2016
[3] Gardner et al., EMNLP, 2015 [4] Bordes et al., NIPS, 2013

Explaining the decision

- Output
  - “Michael White’s alma mater is UT Austin” is FALSE because
  - Evi 1: employer(Michael White, UT Austin)
    Evi 2: almaMater(Michael White, Abilene Christian Univ.)
    almaMater(Michael White, Yale Divinity School)

?

  - Approximate rules over the KG
    0.1: almaMater(x,y), almaMater(x,z), y<>z ⇒ ⊥
    0.3: almaMater(x,y), employer(x,y) ⇒ ⊥

Negative rules

$0.9\ \text{foundingYear}(x,z) \wedge \text{birthYear}(y,w) \wedge (z \leq w) \wedge \text{founder}(x,y) \Rightarrow \perp$

Atom = predicate from KG

[claim] founder(Chev., E.Musk)
[KG] foundingYear(Chev,1911)
[KG] birthYear(E.Musk,1971)

x=Chev., z=1911, …
At least one triple is incorrect

"negative" rules identify inconsistencies

$1.0\ \text{President}(x, \text{USA}) \wedge \text{PlaceOfBirth}(x,y) \wedge y \neq \text{USA} \Rightarrow \perp$

$0.9\ \text{isMarriedTo}(x,y) \wedge \text{hasChild}(x,y) \Rightarrow \perp$

[Ortona et al., 2018]

Horn rules discovery

0.7 hasChild(x,z) ∧ hasChild(y,z) ⇒ spouse(x,y)

child(Barack, Sasha) ∧ child(Michelle, Sasha)
⇒ spouse(Barack, Michelle)

"positive" rules address incomplete data

1.0 notableWork(y,x) ⇒ creator(x,y) (Wikidata)
0.8 hasChild(z,y) ∧ isMarriedTo(x,z) ⇒ hasChild(x,y) (Yago)

[Galárraga et al., 2013]
[Lajus et al, 2020]

RuleHub

http://rudik.eurecom.fr

Not secure | rudik.eurecom.fr:8080

RuleHub | Show rules | Add rules | Rule Administration | About | Login

Search Rules

List out rules base on following criteria. Users can vote the ***rule quality*** *and* ***rule confidence*** *by clicking the cell value.*

Knowledge Base DBpedia

Predicate http://dbpedia.org/ontology/spouse

Rule Type Positive

Human Confidence *from* 0.01 *to* 1.0

Apply

Select all | Deselect all | JSON Export | SHACL Export

Search:

	Type	Rule	Quality Evaluation	Human Confidence	Computed Confidence	Operation
☐	+	successor(subject,v0) & parent(v0,object) ⇒ spouse(subject,object)	5	0.45	0.28	Sample SPARQL Query
☐	+	predecessor(v0,object) & spouse(v1,v0) & successor(subject,v1) ⇒ spouse(subject,object)	5	0.15	0.25	Sample SPARQL Query

ExpClaim system

"Elon Musk is the founder of Chevrolet" → founder(Chevrolet, E.Musk)

Ahmadi et al. *Explainable Fact Checking with Probabilistic Answer Set Programming*. TTO 2019.

https://github.com/ppapotti/expclaim

Mine Logical Rules

founder(Chevrolet, E.Musk)

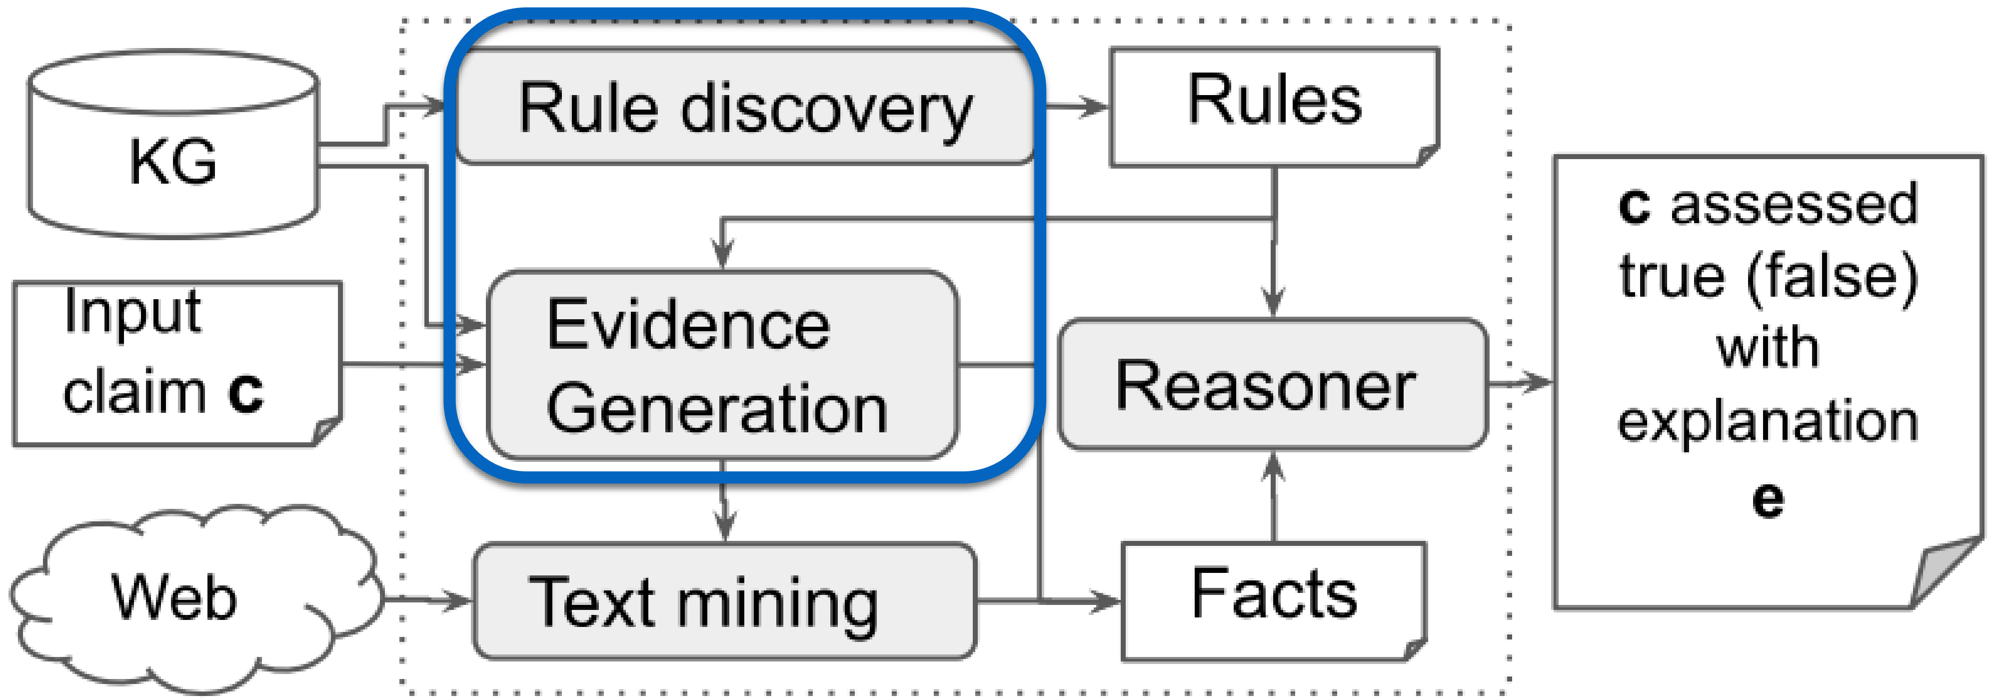

Gather rules and evidence (facts) from the KG

0.9 foundingYear(x,z) ∧ birthYear(y,w) ∧ (z≤w) ∧ founder(x,y) ⇒ ⊥

*[K] foundingYear(Chev,1911), [K] birthYear(E.Musk,1971), [**c**] founder(E.Musk,Chev)*

MAP Reasoning

“Elon Musk is the founder of Chevrolet”

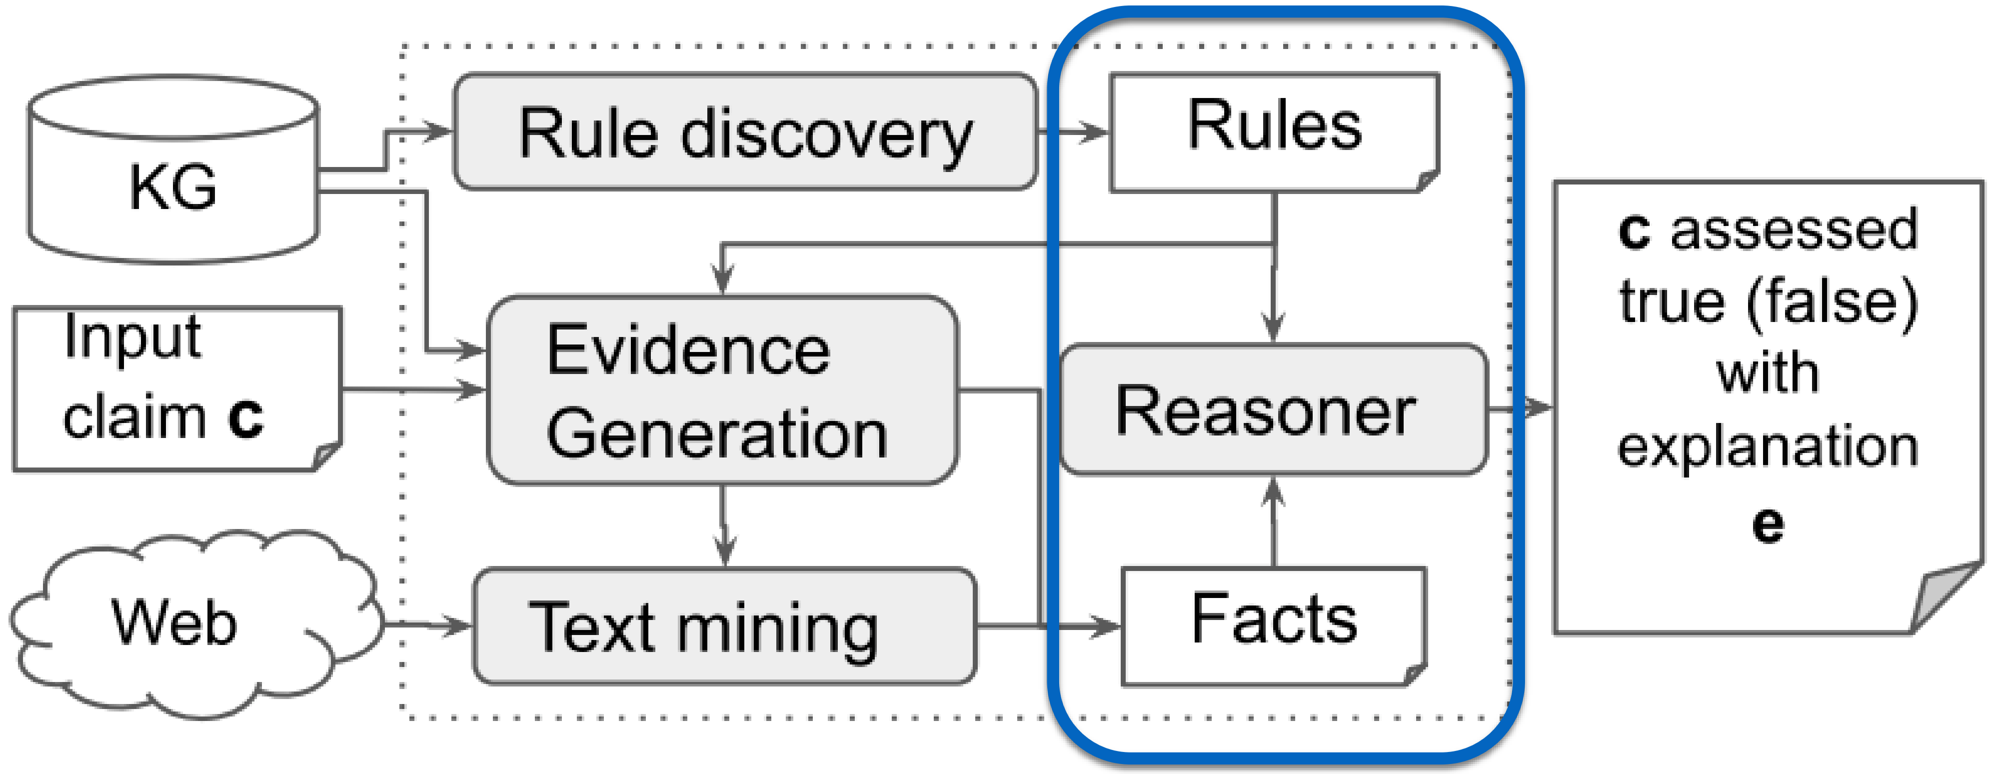

A probabilistic extension of answer set programs with the concept of **weighted rules** derived from Markov Logic

[Lee et al., 2016]

```
FALSE : almaMater(Michael White, UT Austin)
    ← employer(Michael White, UT Austin)
    ← occupation(Michael White, UT Austin)
    ← almaMater(Michael White, Abilene
       Christian Univ.), almaMater(Michael
        White, Yale Divinity School)
```

- Baselines: text miner (CredE), link prediction (KGM), hard rules (ASP)

	almaMat.	*deathPl.*	*spouse*	*vicePres.*
CREDE	.41(.03)	.59(.06)	.44(.07)	.36(.15)
KGM	.73(.08)	.68(.01)	.86(.01)	**.81**(.03)
ASP	.70(.06)	.01(.01)	.31(.08)	.18(.16)
MAP	**.88**(.14)	.75(.15)	**.87**(.11)	.66(.22)
MAP+W	**.88**(.09)	**.83**(.11)	.86(.10)	.68(.18)

Computational fact checking

- Given some **textual content** and some reference **data**, assess if the content is **true or false** and explain **why**

  - Check(claim,data) =
    true/false label,
    confidence value,
    description of the subset of the data that implies the decision

- Property and **statistical** claims

Checking statistical claims

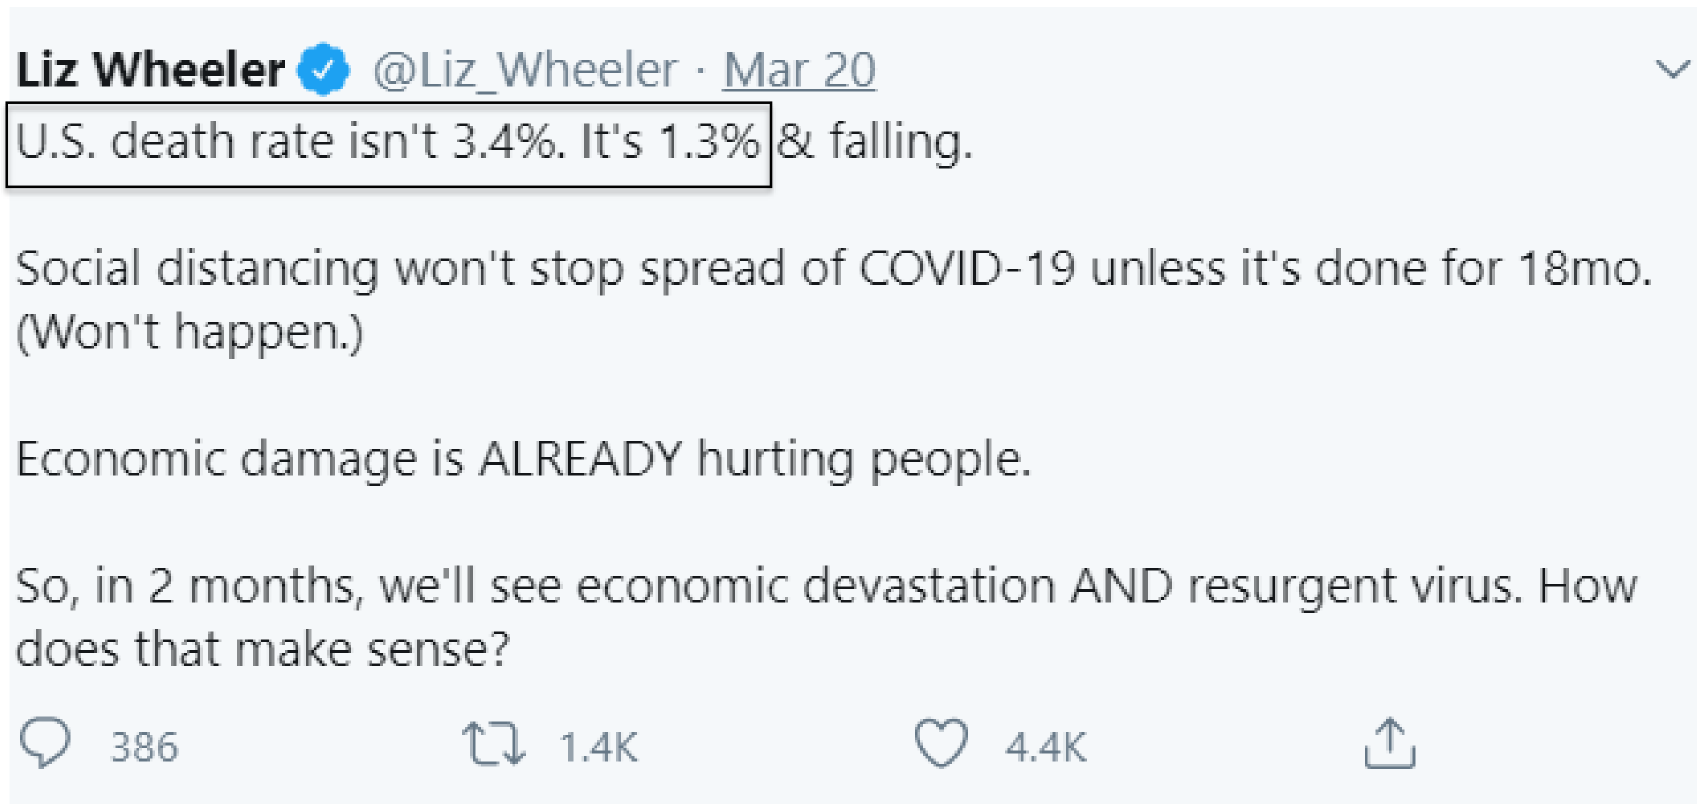

Country	deaths	cases
IT	100	10000
FR	110	9900
USA	536	190000

- Rules (FDs, DCs) could be mined over relational data
  In most cases values are there
  Less rich schema: check a few claims (deaths<cases)
- What about *derived values* such as “death rate”?

Checking statistical claims

Liz Wheeler @Liz_Wheeler · Mar 20

U.S. death rate isn't 3.4%. It's 1.3% & falling.

Social distancing won't stop spread of COVID-19 unless

(Won't happen.)

Economic damage is ALREADY hurting people.

So, in 2 months, we'll see economic devastation AND re

does that make sense?

386 1.4K 4.4K

- Output: FALSE

https://coronacheck.eurecom.fr /en,

How did we verify your statement? ×

Claim: 0.01300000000000(

Infered formula: a/b

Infered formula value: 0.(

List of variables:

Variable name	
Table	
Month	March
Country	USA
Value	5367

Variable name	
Table	
Month	
Country	USA
Value	188172

Online since March 2020

>20k claims checked

Problem model

- Scenario: textual claim **C**, and a set of relational tables **D** (e.g., https://github.com/CSSEGISandData/COVID-19) A_S

- A **claim** describes the comparison **op**(<,=,<>,>) between the value of query **q** and a parameter **p**, when **q** is executed on **D**. A claim is *correct* if there if there exists **q(D) op p** for it A_C
  - “The number of new cases increased **100%** in March in France” → *explicit* claim
  - “The number of deaths increased in March in France” → *general* claim

- If there is no **q(D)** s.t. **op p**, then the claim is false (CWA)

- Challenges: (i) number of possible queries is huge, (ii) spurious match

Checking statistical claims

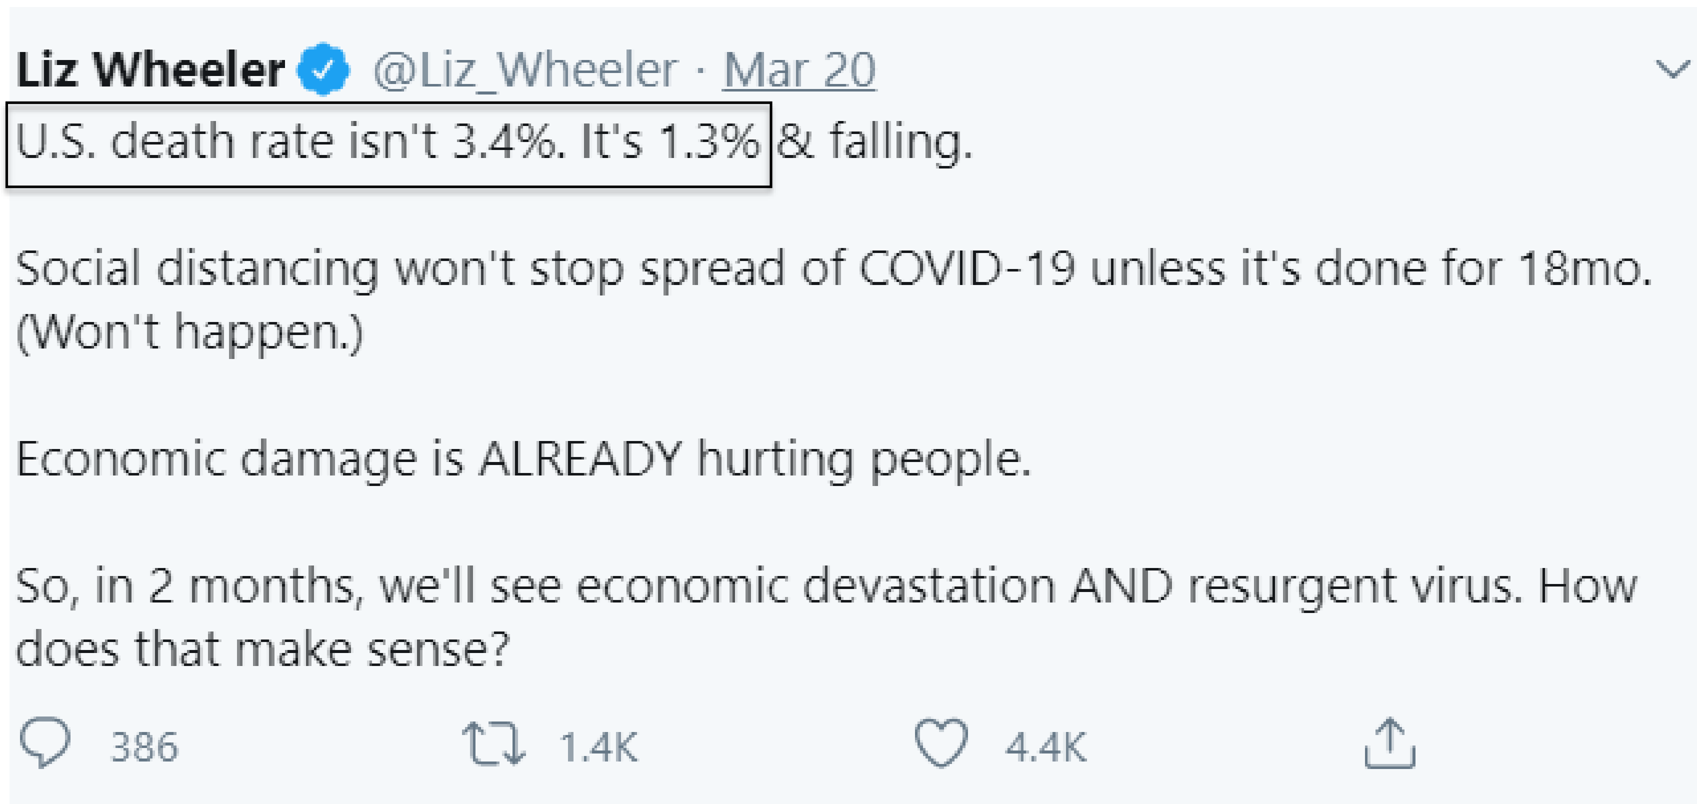

Country	deaths	cases
IT	100	10000
FR	130	9900
USA	536	190000

1. Lots of true queries:
   look up: 100, 10000, 110….
   ratio: 100/130, 130/9900, 9900/130 …
2. 1.3% may come from 130/100!

From text to query elements to data driven check

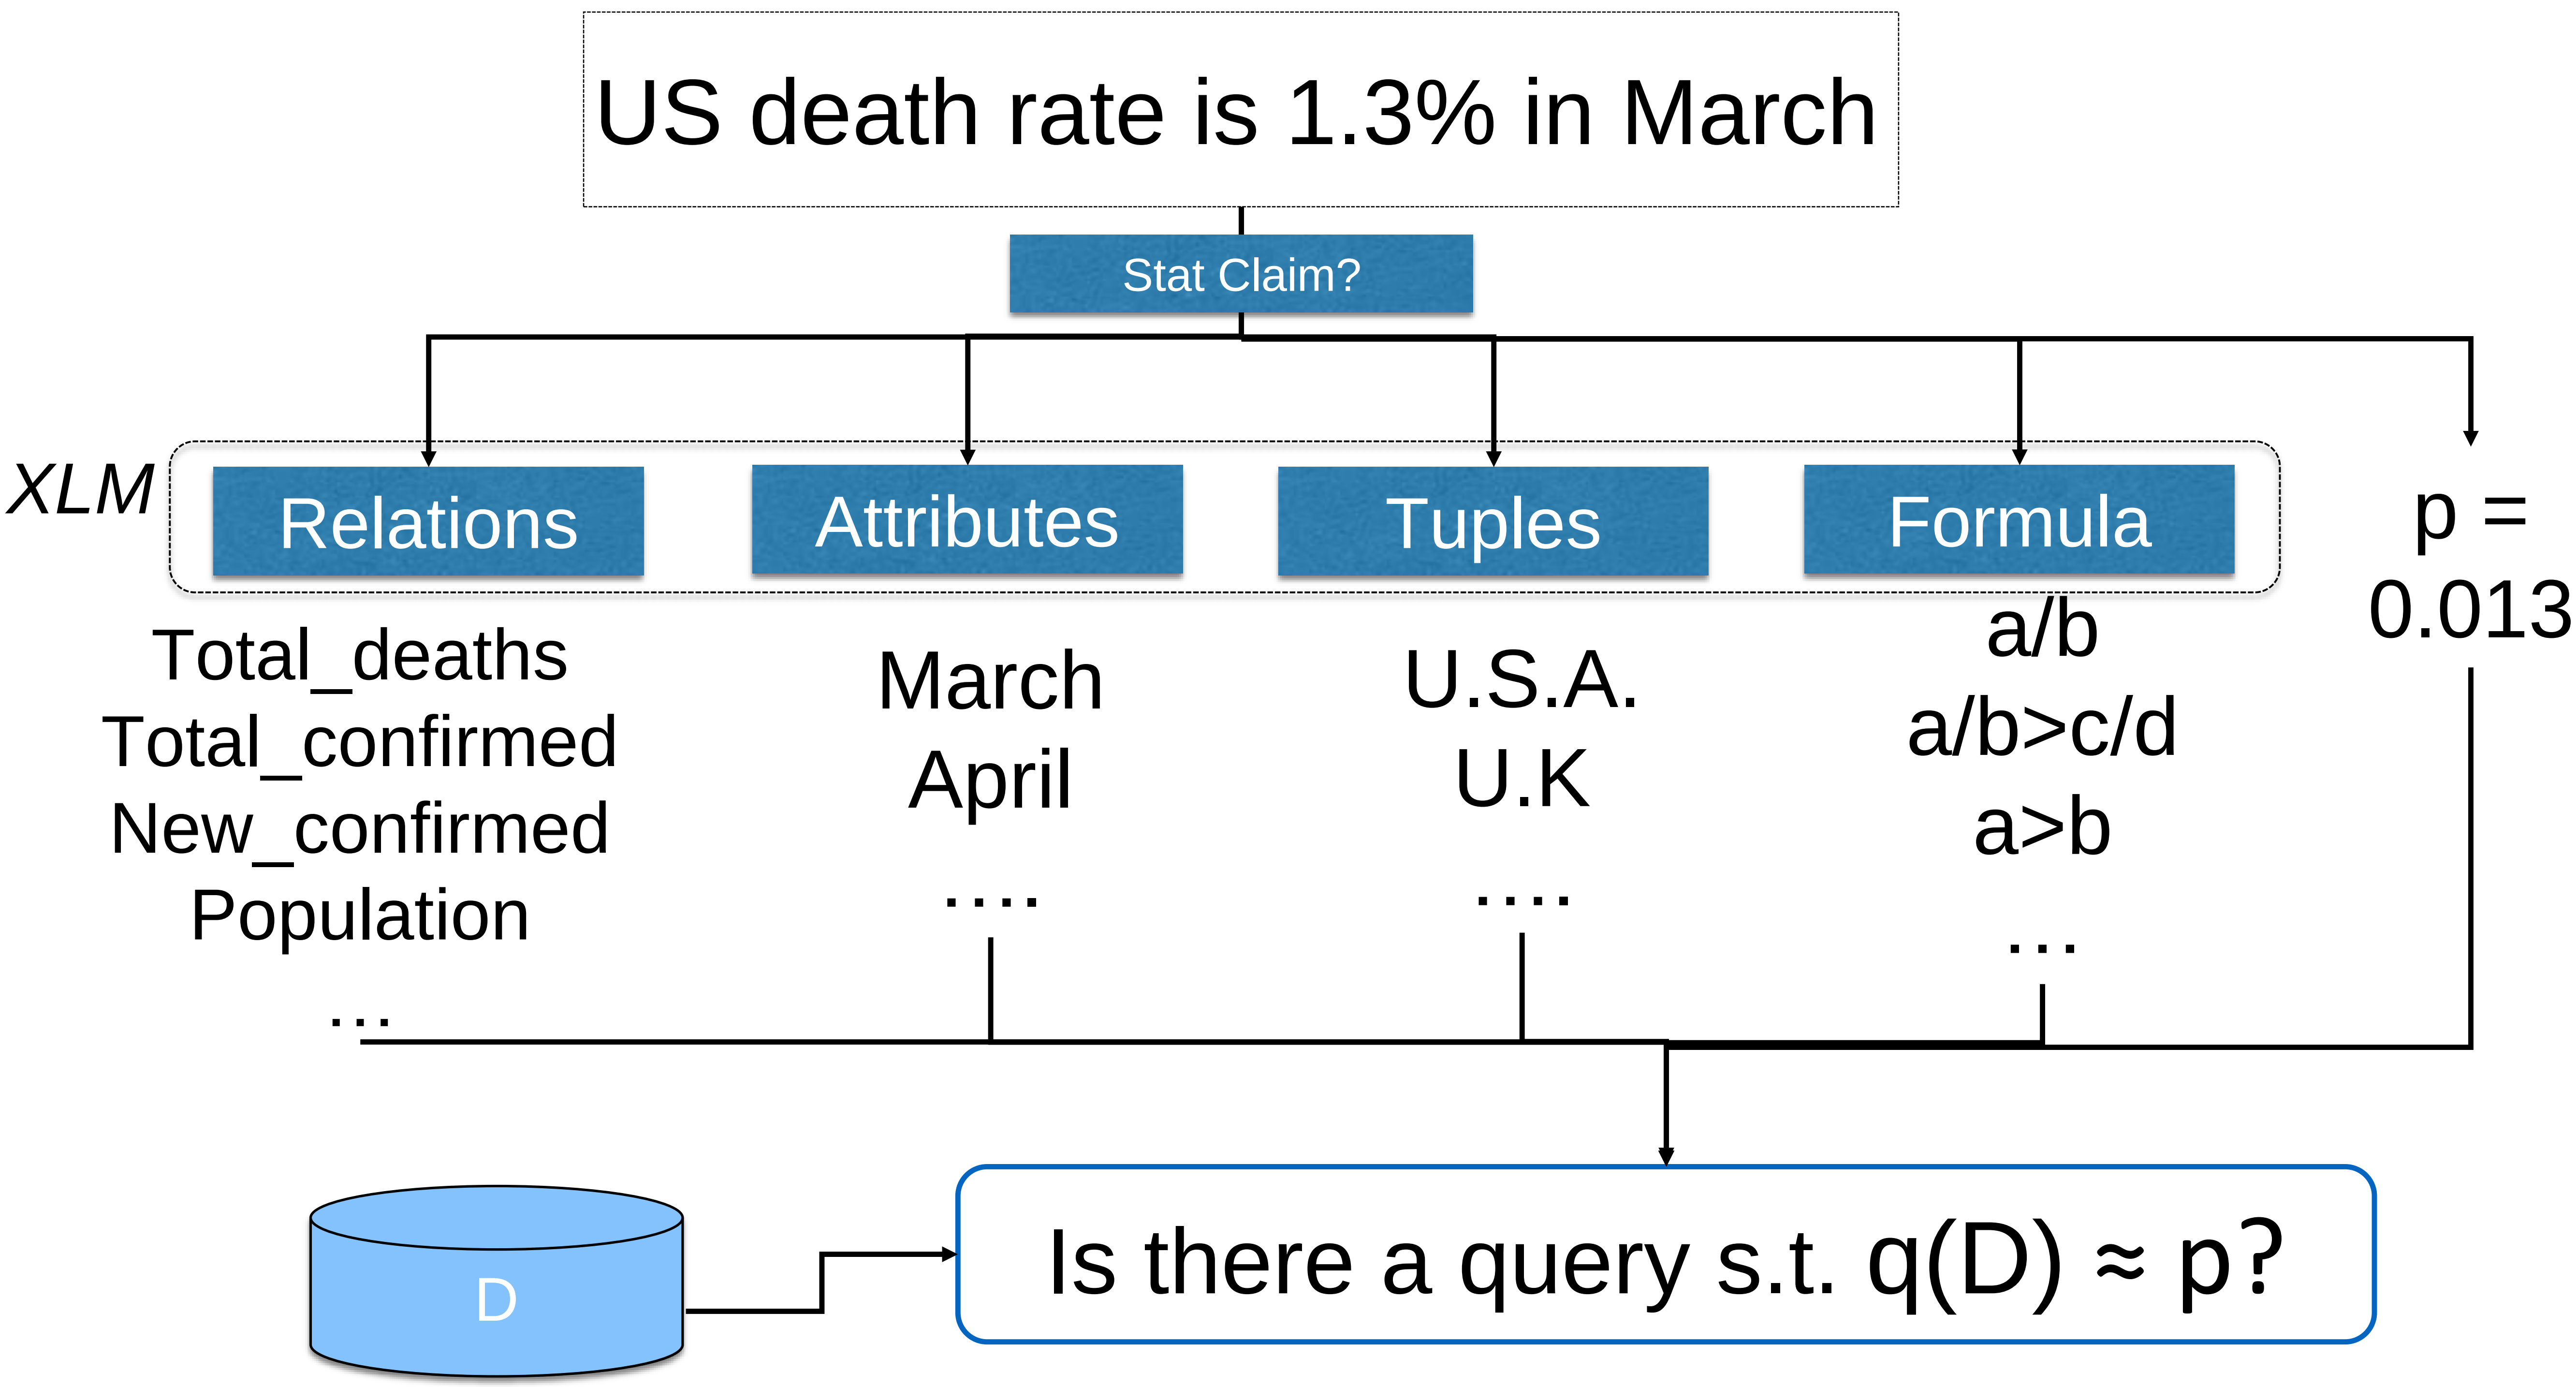

User feedback

- “Yesterday the number of new cases in France decreased **a lot**”
  New_confirmed,
  14-15 Sept,
  France,
  (b-a)/a < **?**

- “In **Europe**, the number of new cases is decreasing this month’
  New_confirmed,
  Aug-Sept,
  tuple **?**
  (b-a)/a < 0

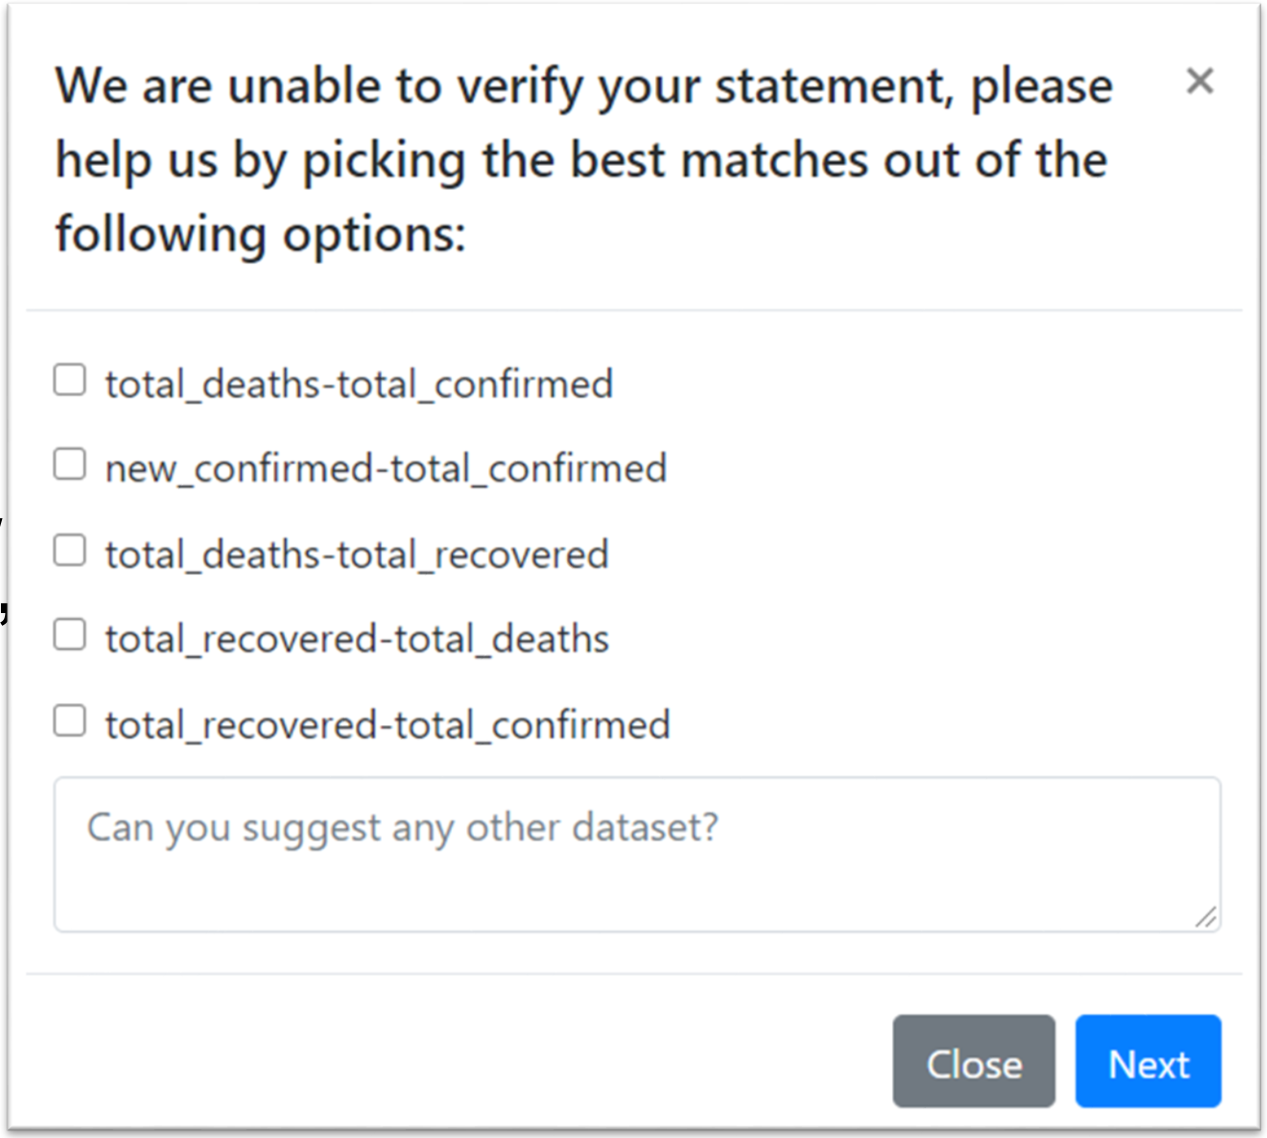

Scrutinizer

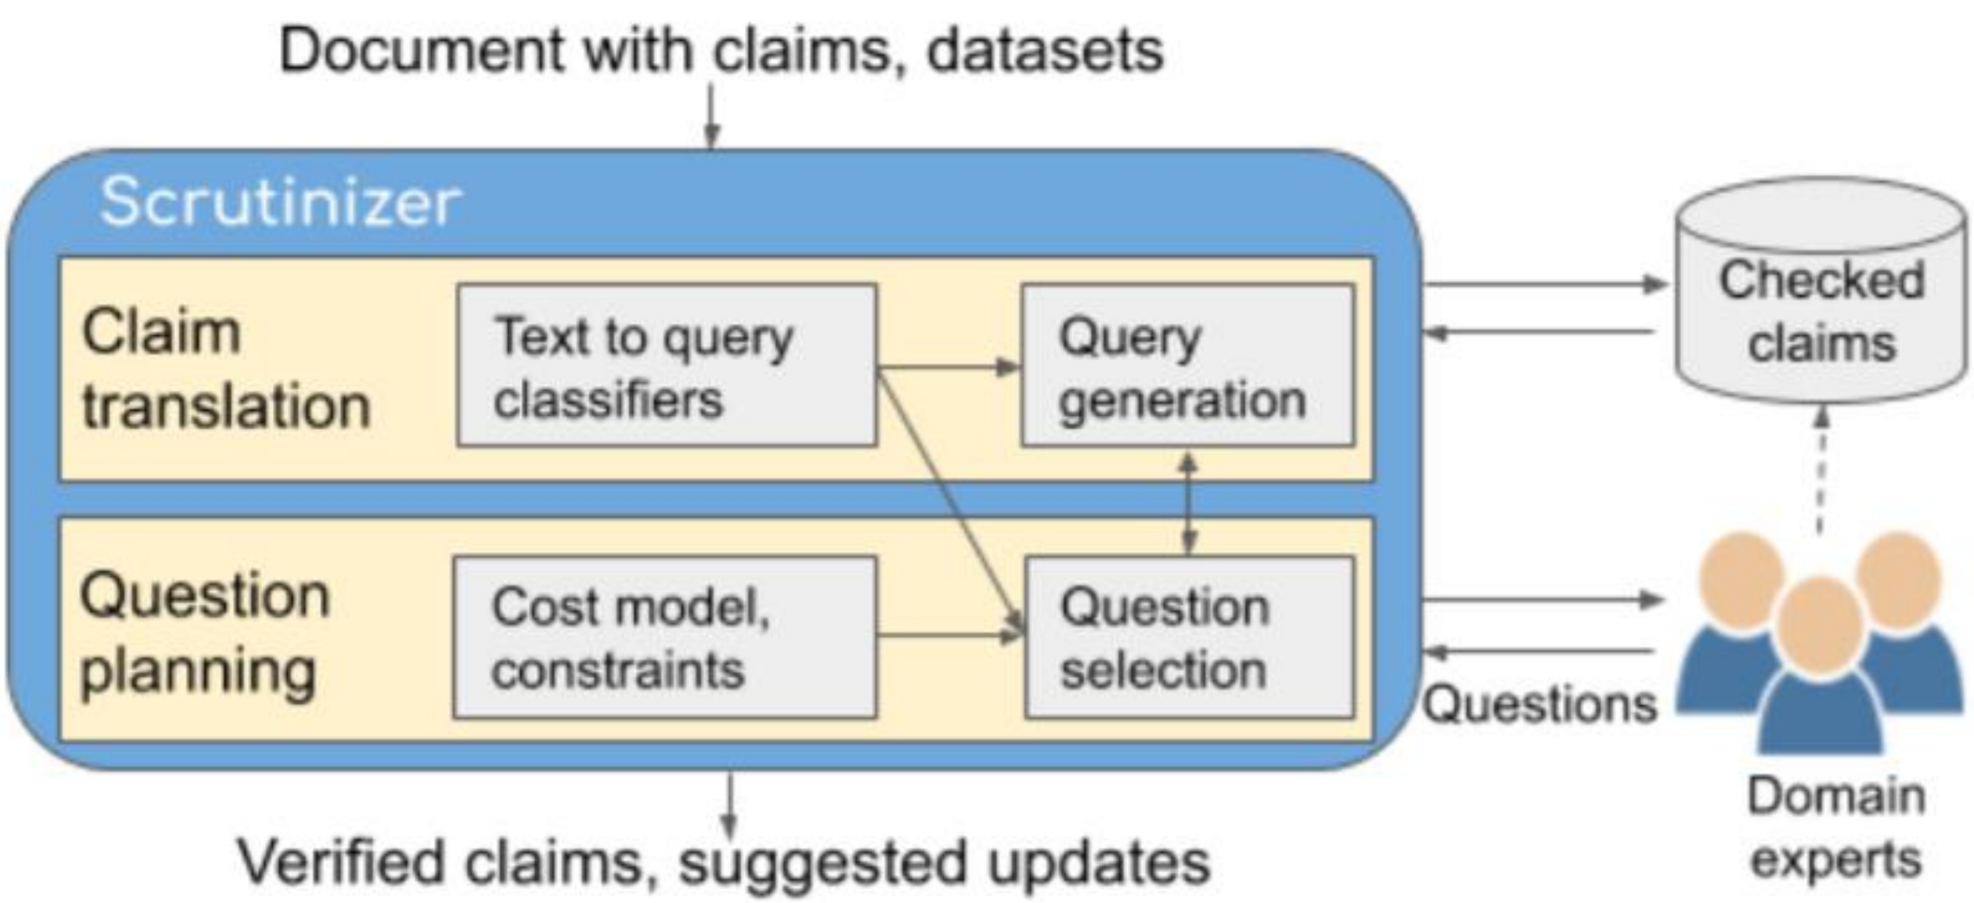

Scrutinizer: A Mixed-Initiative Approach to Large-Scale, Data-Driven Claim Verification. Karagiannis et al. – PVLDB 2020 https://github.com/geokaragiannis/statchecker

International Energy Agency

Within gross output sectors there is a shift in regional shares of total world production for individual industries. India gains a larger share of world nonmetallic mineral production, in part because of rapid growth in its domestic construction industry. India's share of world steel production also increases from 2012 to 2040, while China's share remains relatively flat. In the paper industry, China's share of world production increases as growth is slower in other regions, including in Russia and in other non-OECD Europe and Eurasia. For basic chemicals, the Middle East, China, and India have increasing production shares throughout the projection period. Finally, most non-OECD countries increase their shares of total world agricultural sector production through 2040, and most OECD countries have decreasing shares. Within gross output sectors there is a shift in regional shares of total world production for individual industries. India gains a larger share of world nonmetallic mineral production, in part because of rapid growth in its domestic construction industry. India's share of world steel production also increases from 2012 to 2040, while China's share remains relatively flat. In the paper industry, China's share of world production increases as growth is slower in other regions, including in Russia and in other non-OECD Europe and Eurasia. For basic chemicals, the Middle East, China, and India have increasing production shares throughout the projection period. Finally, most non-OECD countries increase their shares of total world agricultural sector production through 2040, and most OECD countries have decreasing shares.

Industrial sector delivered energy consumption varies by region, according to differences in industrial gross output, energy intensity (measured as energy consumed per unit of gross output), and the composition of industries. Enterprises are able to reduce energy consumption in a number of ways, including improving industrial sector processes to reduce energy waste and recover energy lost (often process heat), increasing the use of cogeneration, and recycling materials and fuel inputs to reduce costs and improve efficiency. In terms of industrial fuel use, natural gas and electricity are the fastest-growing forms of industrial energy use in the OECD region (Figure 7-2), with each energy source increasing by about 0.7%/year from 2012 to 2040. Consumption of liquids, coal, and renewable energy in the OECD industrial sector grows more slowly, averaging 0.4%/year (liquids), 0.3%/year (renewable energy), and 0.2%/year (coal). As a result, from 2012 to 2040, the natural gas and electricity shares increase from 28.7% to 30.3% (natural gas) and from 14.9% to 15.6% (electricity). As in OECD, the fastest-growing forms of industrialsector energy consumption in the non-OECD region from 2012 to 2040 are natural gas (2.2%/year) and electricity (1.6%/year). Non-OECD consumption of renewable energy also expands rapidly, by an average of 1.7%/year from 2012 to 2040, while consumption of liquid fuels and coal increases by 1.4%/year and 0.8%/year, respectively (Figure 7-3). Most of the world growth in industrial sector energy use occurs in the emerging non-OECD economies.

The strong rates of growth in industrial sector consumption of electricity and natural gas in both the OECD and non-OECD regions are attributable to increases in the other industrials group of nonenergy-intensive manufacturing (see Table 7-1). Although the other industrials are not energy intensive, they do make up approximately 30% and 36% of total OECD and non-OECD industrial sector delivered energy consumption, respectively. Moreover, the manufacture of bulk chemicals in the non-OECD region expands rapidly in the Reference case. Because nonenergy-intensive manufacturing—including both other industrials and other chemicals—relies heavily on electricity and natural gas, consumption of both energy sources shows strong growth in the OECD and non-OECD industrial sectors, in comparison with most other energy sources.

Biomass currently provides most of the renewable energy (excluding hydroelectricity) consumed in the industrial sector and continues to do so throughout the projection, largely because of its role in providing byproduct energy to the pulp and paper industry. OECD countries typically have either flat or declining growth in the pulp and paper industry, resulting in the slower growth

Figure 7-2. OECD industrial sector delivered energy consumption by energy source, 2012–40 (quadrillion Btu)

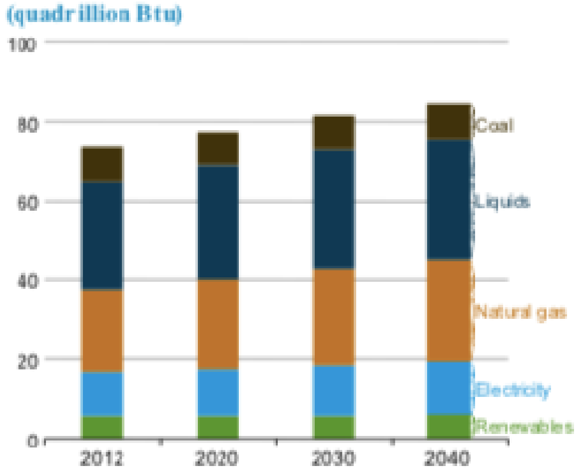

Figure 7-3. Non-OECD industrial sector delivered energy consumption by energy source, 2012–40 (quadrillion Btu)

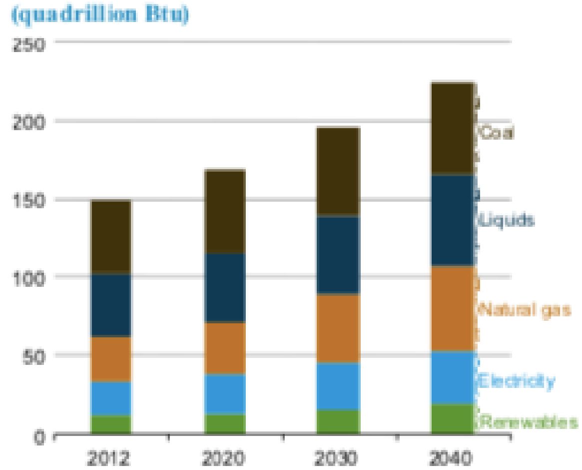

of renewables compared with total industrial sector energy consumption. In most of the non-OECD countries, the pulp and paper industry grows significantly, with corresponding growth in industrial sector renewable energy use.

Coal becomes a less important source for industrial energy consumption in both the OECD and non-OECD regions in the IEO2016 Reference case. In the OECD industrial sector, coal use increases by 0.2%/year from 2012 to 2040, and its share of total delivered industrial energy consumption declines slightly, from 12% in 2012 to 11% in 2040. Similarly, in the non-OECD industrial sector, the coal share of total industrial delivered energy falls from 32% in 2012 to 27% in 2040. The iron and steel industry is the largest consumer of coal in the industrial sector, and as regions shift from coal-fired furnaces to electric arc furnaces, coal use for iron and steel production declines. In addition, several significant industrial manufacturing countries, including the United States and China, are initiating policies to reduce greenhouse gas (GHG) emissions from their industrial sectors by switching to electricity and natural gas and by improving energy efficiency in industries that produce large amounts of GHG emissions.

Energy-intensive industries

The following industries are considered to be energy-intensive: food, pulp and paper, basic chemicals, refining, iron and steel, nonferrous metals (primarily aluminum), and nonmetallic minerals (primarily cement). Together, they account for about half of all industrial sector delivered energy use. In 2012, OECD energy-intensive industries accounted for about 54% of the region's total industrial sector energy consumption, and non-OECD energy-intensive industries accounted for about 51% of the industrial sector total. Consequently, the quantity and fuel mix of future industrial sector delivered energy consumption will be determined largely by the overall levels of energy consumption in those seven industries. In addition, the same industries emit large quantities of carbon dioxide (CO_2), related to both their energy consumption (combustion emissions) and their production processes (process emissions). Figure 7-4 and Figure 7-5 show energy consumption shares of the energy-intensive industries compared with all industrial sector energy consumption (including feedstock consumption) in 2012 and 2040 for the OECD and non-OECD, respectively. The energy consumption shares of the energy-intensive industries are shown as percentages of total delivered energy consumption in the OECD and non-OECD industrial sectors.

Increases in energy efficiency and changes in industrial gross output affect the growth of industrial sector energy consumption. Anticipated energy efficiency improvements in the industrial sector temper the growth of industrial energy demand, particulary for the energy-intensive industries. Recycling is a key contributor to industrial energy efficiency improvements, especially in the pulp and paper, iron and steel, and nonferrous metals industries *(see box on page 120)*.

Among the energy-intensive industries, the largest consumer of delivered energy is the basic chemicals industry, which in 2012 accounted for about 19% of total delivered energy consumption in the OECD industrial sector and about 14% in the non-OECD industrial sector. In both regions, the basic chemicals share of industrial energy use in the IEO2016 Reference case rises to about 20% in 2040 (Figure 7-4 and Figure 7-5). The chemicals industry in general uses petrochemical feedstocks, which are included in its energy use. In 2012, petrochemical feedstocks accounted for roughly 60% of the energy consumed in the chemicals sector (which includes both energy-intensive basic chemicals and nonenergy-intensive other chemicals). Intemediate petrochemical products (or building blocks), which go into products such as plastics, require a fixed amount of hydrocarbon feedstock as input. For any given amount of chemical output, depending on the fundamental chemical process of production, a fixed amount of feedstock is required, which greatly reduces opportunities for decreasing fuel consumption in the absence of any major shifts toward recycling and bio-based chemicals.

Figure 7-4. Energy-intensive industry shares of total OECD industrial sector energy consumption, 2012 and 2040 (percent of total)

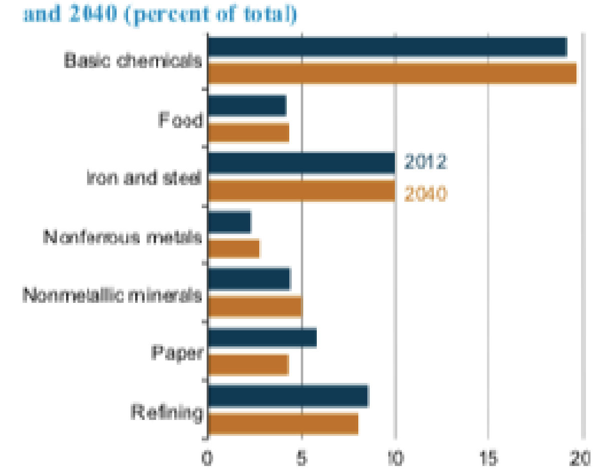

Figure 7-5. Energy-intensive industry shares of total non-OECD industrial sector energy consumption, 2012 and 2040 (percent of total)

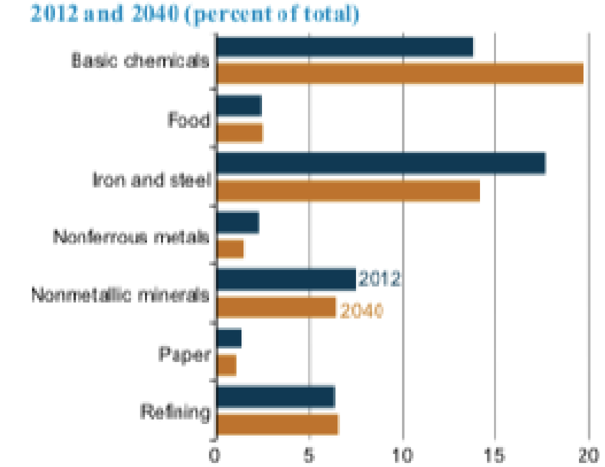

Claims Marked Up

Within gross output sectors there is a shift in regional shares of total world production for individual industries. India gains a larger share of world nonmetallic mineral production, in part because of rapid growth in its domestic construction industry. India's share of world steel production also increases from 2012 to 2040, while China's share remains relatively flat. In the paper industry, China's share of world production increases as growth is slower in other regions, including in Russia and in other non-OECD Europe and Eurasia. For basic chemicals, the Middle East, China, and India have increasing production shares throughout the projection period. Finally, most non-OECD countries increase their shares of total world agricultural sector production through 2040, and most OECD countries have decreasing shares. Within gross output sectors there is a shift in regional shares of total world production for individual industries. India gains a larger share of world nonmetallic mineral production, in part because of rapid growth in its domestic construction industry. India's share of world steel production also increases from 2012 to 2040, while China's share remains relatively flat. In the paper industry, China's share of world production increases as growth is slower in other regions, including in Russia and in other non-OECD Europe and Eurasia. For basic chemicals, the Middle East, China, and India have increasing production shares throughout the projection period. Finally, most non-OECD countries increase their shares of total world agricultural sector production through 2040, and most OECD countries have decreasing shares.

Industrial sector delivered energy consumption varies by region, according to differences in industrial gross output, energy intensity (measured as energy consumed per unit of gross output), and the composition of industries. Enterprises are able to reduce energy consumption in a number of ways, including improving industrial sector processes to reduce energy waste and recover energy lost (often process heat), increasing the use of cogeneration, and recycling materials and fuel inputs to reduce costs and improve efficiency. In terms of industrial fuel use, natural gas and electricity are the fastest-growing forms of industrial energy use in the OECD region (Figure 7-2), with each energy source increasing by about 0.7%/year from 2012 to 2040. Consumption of liquids, coal, and renewable energy in the OECD industrial sector grows more slowly, averaging 0.4%/year (liquids), 0.3%/year (renewable energy), and 0.2%/year (coal). As a result, from 2012 to 2040, the natural gas and electricity shares increase from 28.7% to 30.3% (natural gas) and from 14.9% to 15.6% (electricity). As in OECD, the fastest-growing forms of industrial sector energy consumption in the non-OECD region from 2012 to 2040 are natural gas (2.2%/year) and electricity (1.6%/year). Non-OECD consumption of renewable energy also expands rapidly, by an average of 1.7%/year from 2012 to 2040, while consumption of liquid fuels and coal increases by 1.4%/year and 0.8%/year, respectively (Figure 7-3). Most of the world growth in industrial sector energy use occurs in the emerging non-OECD economies.

The strong rates of growth in industrial sector consumption of electricity and natural gas in both the OECD and non-OECD regions are attributable to increases in the other industrials group of nonenergy-intensive manufacturing (see Table 7-1). Although the other industrials are not energy intensive, they do make up approximately 30% and 36% of total OECD and non-OECD industrial sector delivered energy consumption, respectively. Moreover, the manufacture of bulk chemicals in the non-OECD region expands rapidly in the Reference case. Because nonenergy-intensive manufacturing—including both other industrials and other chemicals—relies heavily on electricity and natural gas, consumption of both energy sources shows strong growth in the OECD and non-OECD industrial sectors, in comparison with most other energy sources.

Biomass currently provides most of the renewable energy (excluding hydroelectricity) consumed in the industrial sector and continues to do so throughout the projection, largely because of its role in providing byproduct energy to the pulp and paper industry. OECD countries typically have either flat or declining growth in the pulp and paper industry, resulting in the slower growth

Figure 7-2. OECD industrial sector delivered energy consumption by energy source, 2012–40 (quadrillion Btu)

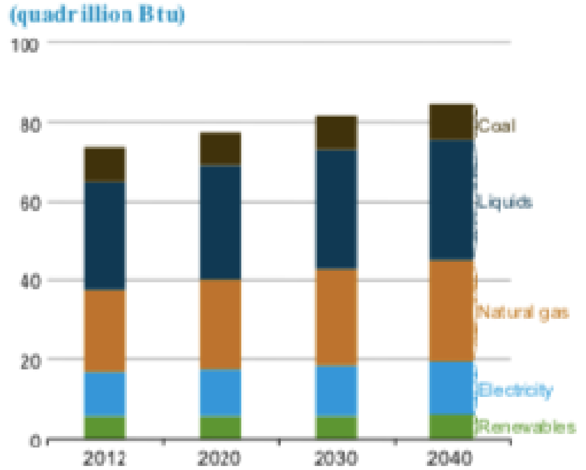

Figure 7-3. Non-OECD industrial sector delivered energy consumption by energy source, 2012–40 (quadrillion Btu)

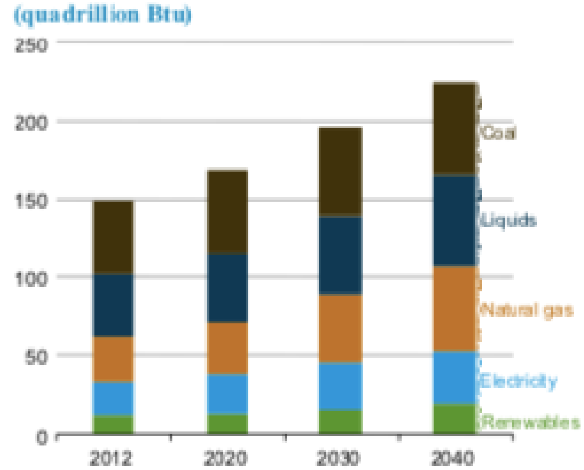

of renewables compared with total industrial sector energy consumption. In most of the non-OECD countries the pulp and paper industry grows significantly, with corresponding growth in industrial sector renewable energy use.

Coal becomes a less important source for industrial energy consumption in both the OECD and non-OECD regions in the IEO2016 Reference case. In the OECD industrial sector, coal use increases by 0.2%/year from 2012 to 2040, and its share of total delivered industrial energy consumption declines slightly, from [illegible]. Similarly, in the non-OECD industrial sector, the coal share of total industrial delivered energy falls from 32% in 2012 to 27% in 2040. The iron and steel industry is the largest consumer of coal in the industrial sector, and as regions shift from coal-fired furnaces to electric arc furnaces, coal use for iron and steel production declines. In addition, several significant industrial manufacturing countries, including the United States and China, are initiating policies to reduce greenhouse gas (GHG) emissions from their industrial sectors by switching to electricity and natural gas and by improving energy efficiency in industries that produce large amounts of GHG emissions.

Energy-intensive industries

The following industries are considered to be energy-intensive: food, pulp and paper, basic chemicals, refining, iron and steel, nonferrous metals (primarily aluminum), and nonmetallic minerals (primarily cement). Together, they account for about half of all industrial sector delivered energy use. In 2012, OECD energy-intensive industries accounted for about 54% of the region's total industrial sector energy consumption, and non-OECD energy-intensive industries accounted for about 51% of the industrial sector total. Consequently, the quantity and fuel mix of future industrial sector delivered energy consumption will be determined largely by the overall levels of energy consumption in those seven industries. In addition, the same industries emit large quantities of carbon dioxide (CO2), related to both their energy consumption (combustion emissions) and their production processes (process emissions). Figure 7-4 and Figure 7-5 show energy consumption shares of the energy-intensive industries compared with all industrial sector energy consumption (including feedstock consumption) in 2012 and 2040 for the OECD and non-OECD, respectively. The energy consumption shares of the energy-intensive industries are shown as percentages of total delivered energy consumption in the OECD and non-OECD industrial sectors.

Increases in energy efficiency and changes in industrial gross output affect the growth of industrial sector energy consumption. Anticipated energy efficiency improvements in the industrial sector temper the growth of industrial energy demand, particularly for the energy-intensive industries. Recycling is a key contributor to industrial energy efficiency improvements, especially in the pulp and paper, iron and steel, and nonferrous metals industries *(see box on page 120)*.

Among the energy-intensive industries, the largest consumer of delivered energy is the basic chemicals industry, which in 2012 accounted for about 19% of total delivered energy consumption in the OECD industrial sector and about 14% in the non-OECD industrial sector. In both regions, the basic chemicals share of industrial energy use in the IEO2016 Reference case rises to about 20% in 2040 (Figure 7-4 and Figure 7-5). The chemicals industry in general uses petrochemical feedstocks, which are included in its energy use. In 2012, petrochemical feedstocks accounted for roughly 60% of the energy consumed in the chemicals sector (which includes both energy-intensive basic chemicals and nonenergy-intensive other chemicals). Intermediate petrochemical products (or building blocks), which go into products such as plastics, require a fixed amount of hydrocarbon feedstock as input. For any given amount of chemical output, depending on the fundamental chemical process of production, a fixed amount of feedstock is required, which greatly reduces opportunities for decreasing fuel consumption in the absence of any major shifts toward recycling and bio-based chemicals.

Figure 7-4. Energy-intensive industry shares of total OECD industrial sector energy consumption, 2012 and 2040 (percent of total)

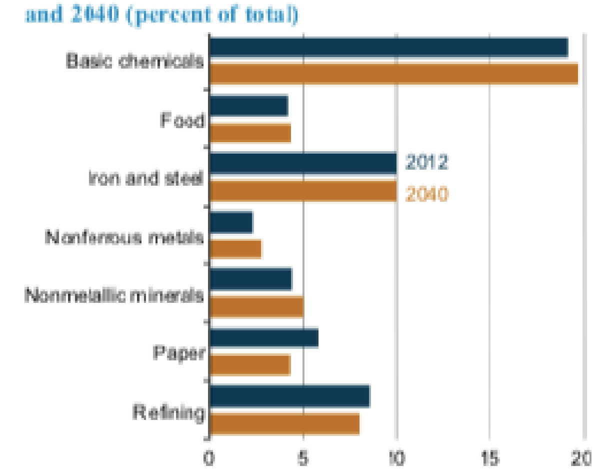

Figure 7-5. Energy-intensive industry shares of total non-OECD industrial sector energy consumption, 2012 and 2040 (percent of total)

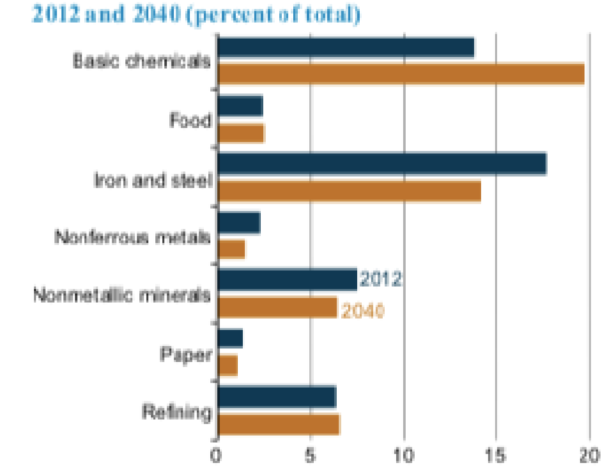

Problem

- The pace of oil demand growth slows, and all of the 11.5 million barrels per day (mb/d) increase between 2017 and 2040 takes place in developing economies. Demand growth is consistently strong in the Middle East and India, particularly for trucks and petrochemical feedstocks. But it is China that becomes the world's biggest oil consumer and, by 2040, the largest net oil importer in history.

A_C

Index	2017	2018	...	2030	2040
PGElecDemand	22 209	22 793	...	29 349	35 526
PGINCoal	2 390	2 412	..	2 341	2 353
TFCelec	21 465	22 040	...	28 566	34 790
...	...	...	...	...	...

A_S

Claim Translation

- Large number of claims in batch
- Classifier accuracy is lower than CoronaCheck, bigger space for datasets, attributes, formulas
- We have access to the checkers

In 2017, global electricity demand grew by 3%.

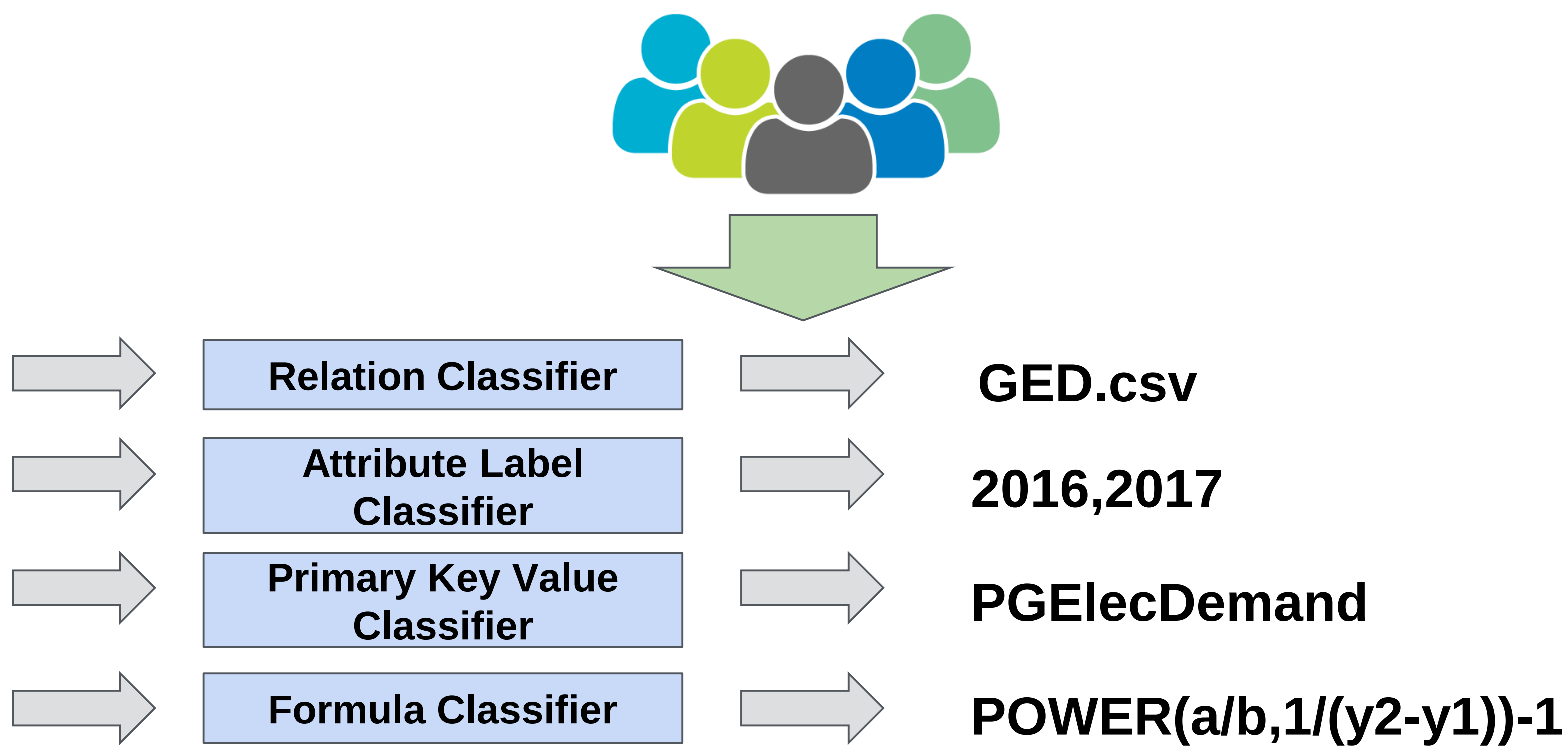

Scaling it up

- Thousands of claims
- Multiple experts in the process
- Hundreds of datasets

Crowdsourcing Workflow

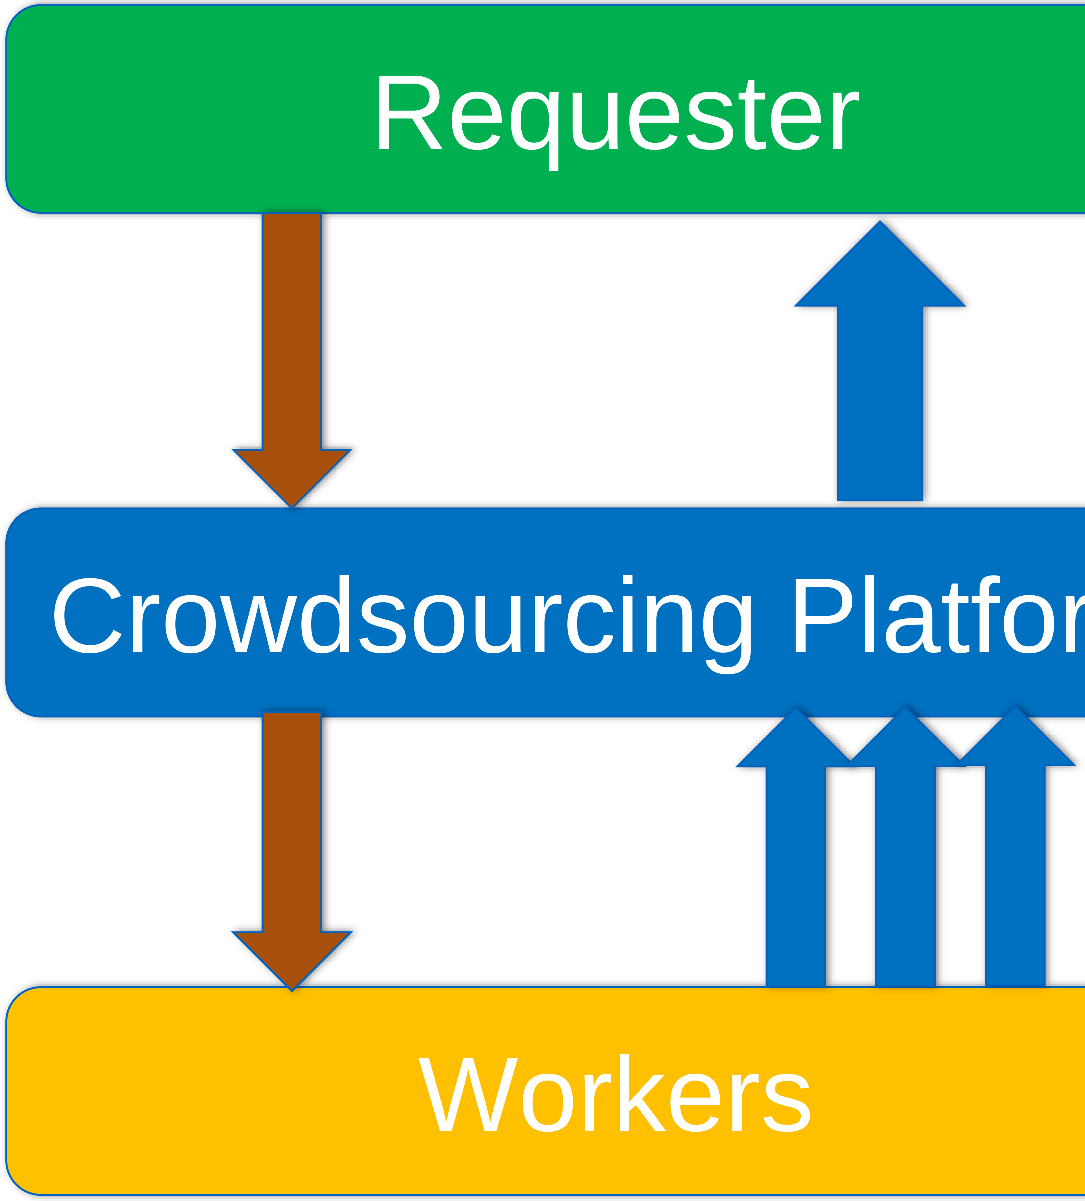

Needs optimization

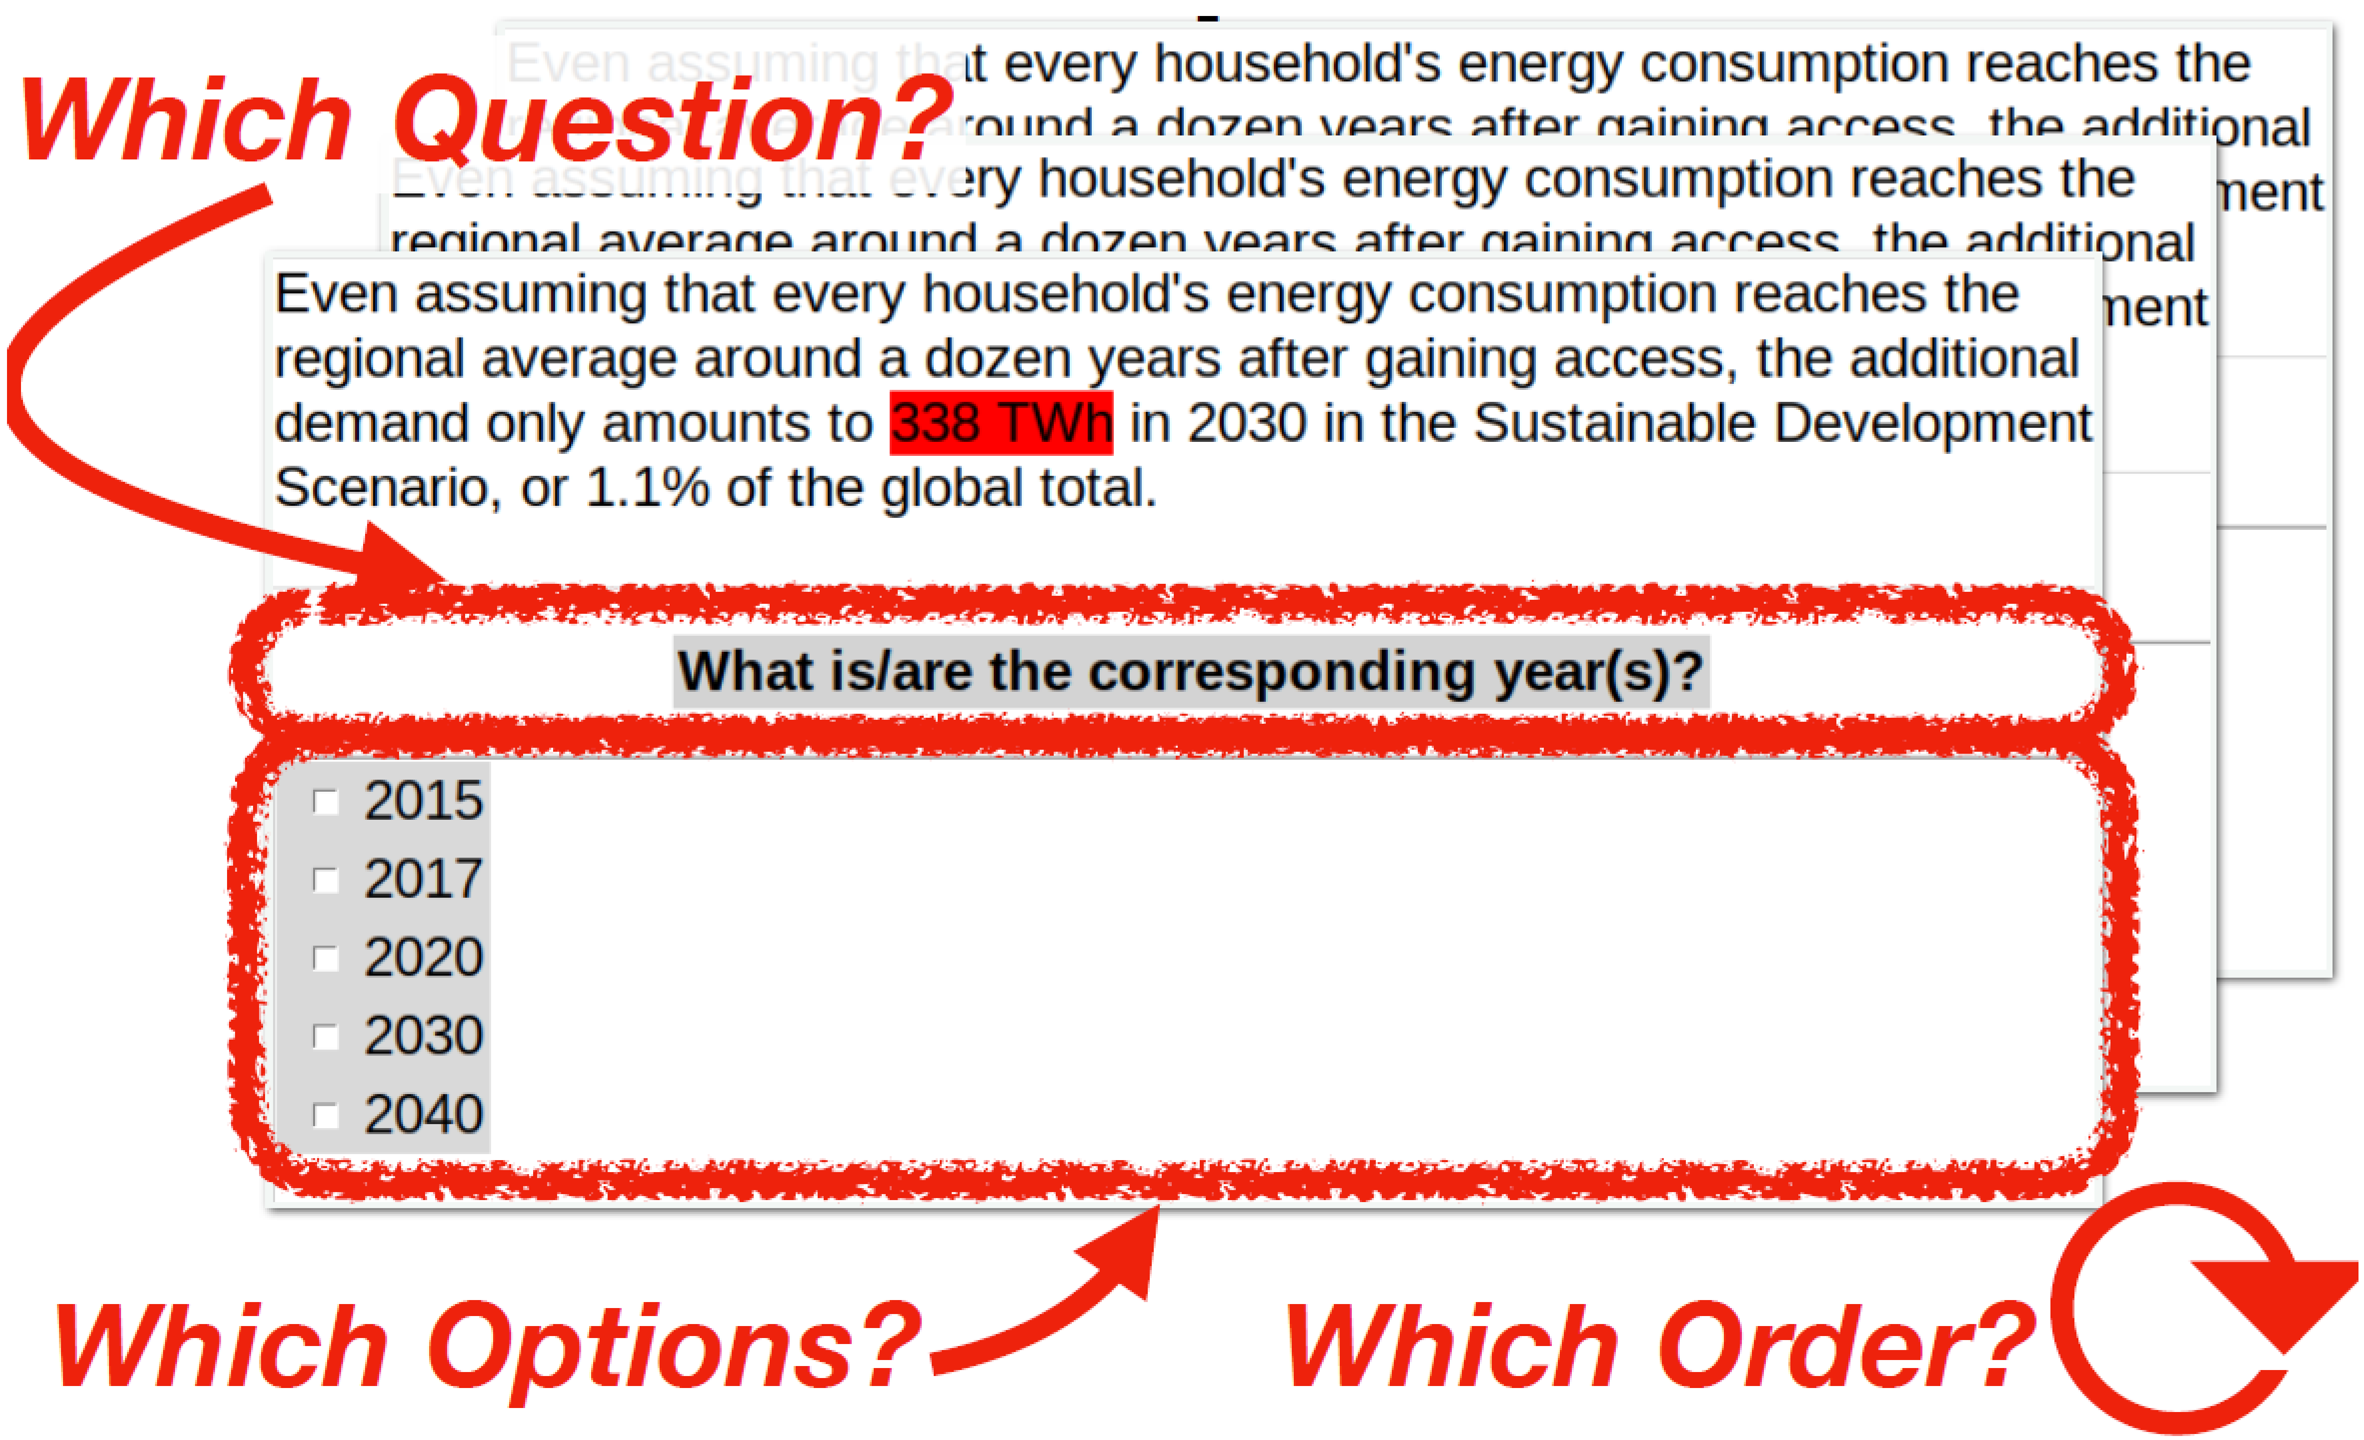

Verify Cheapest Claims First?

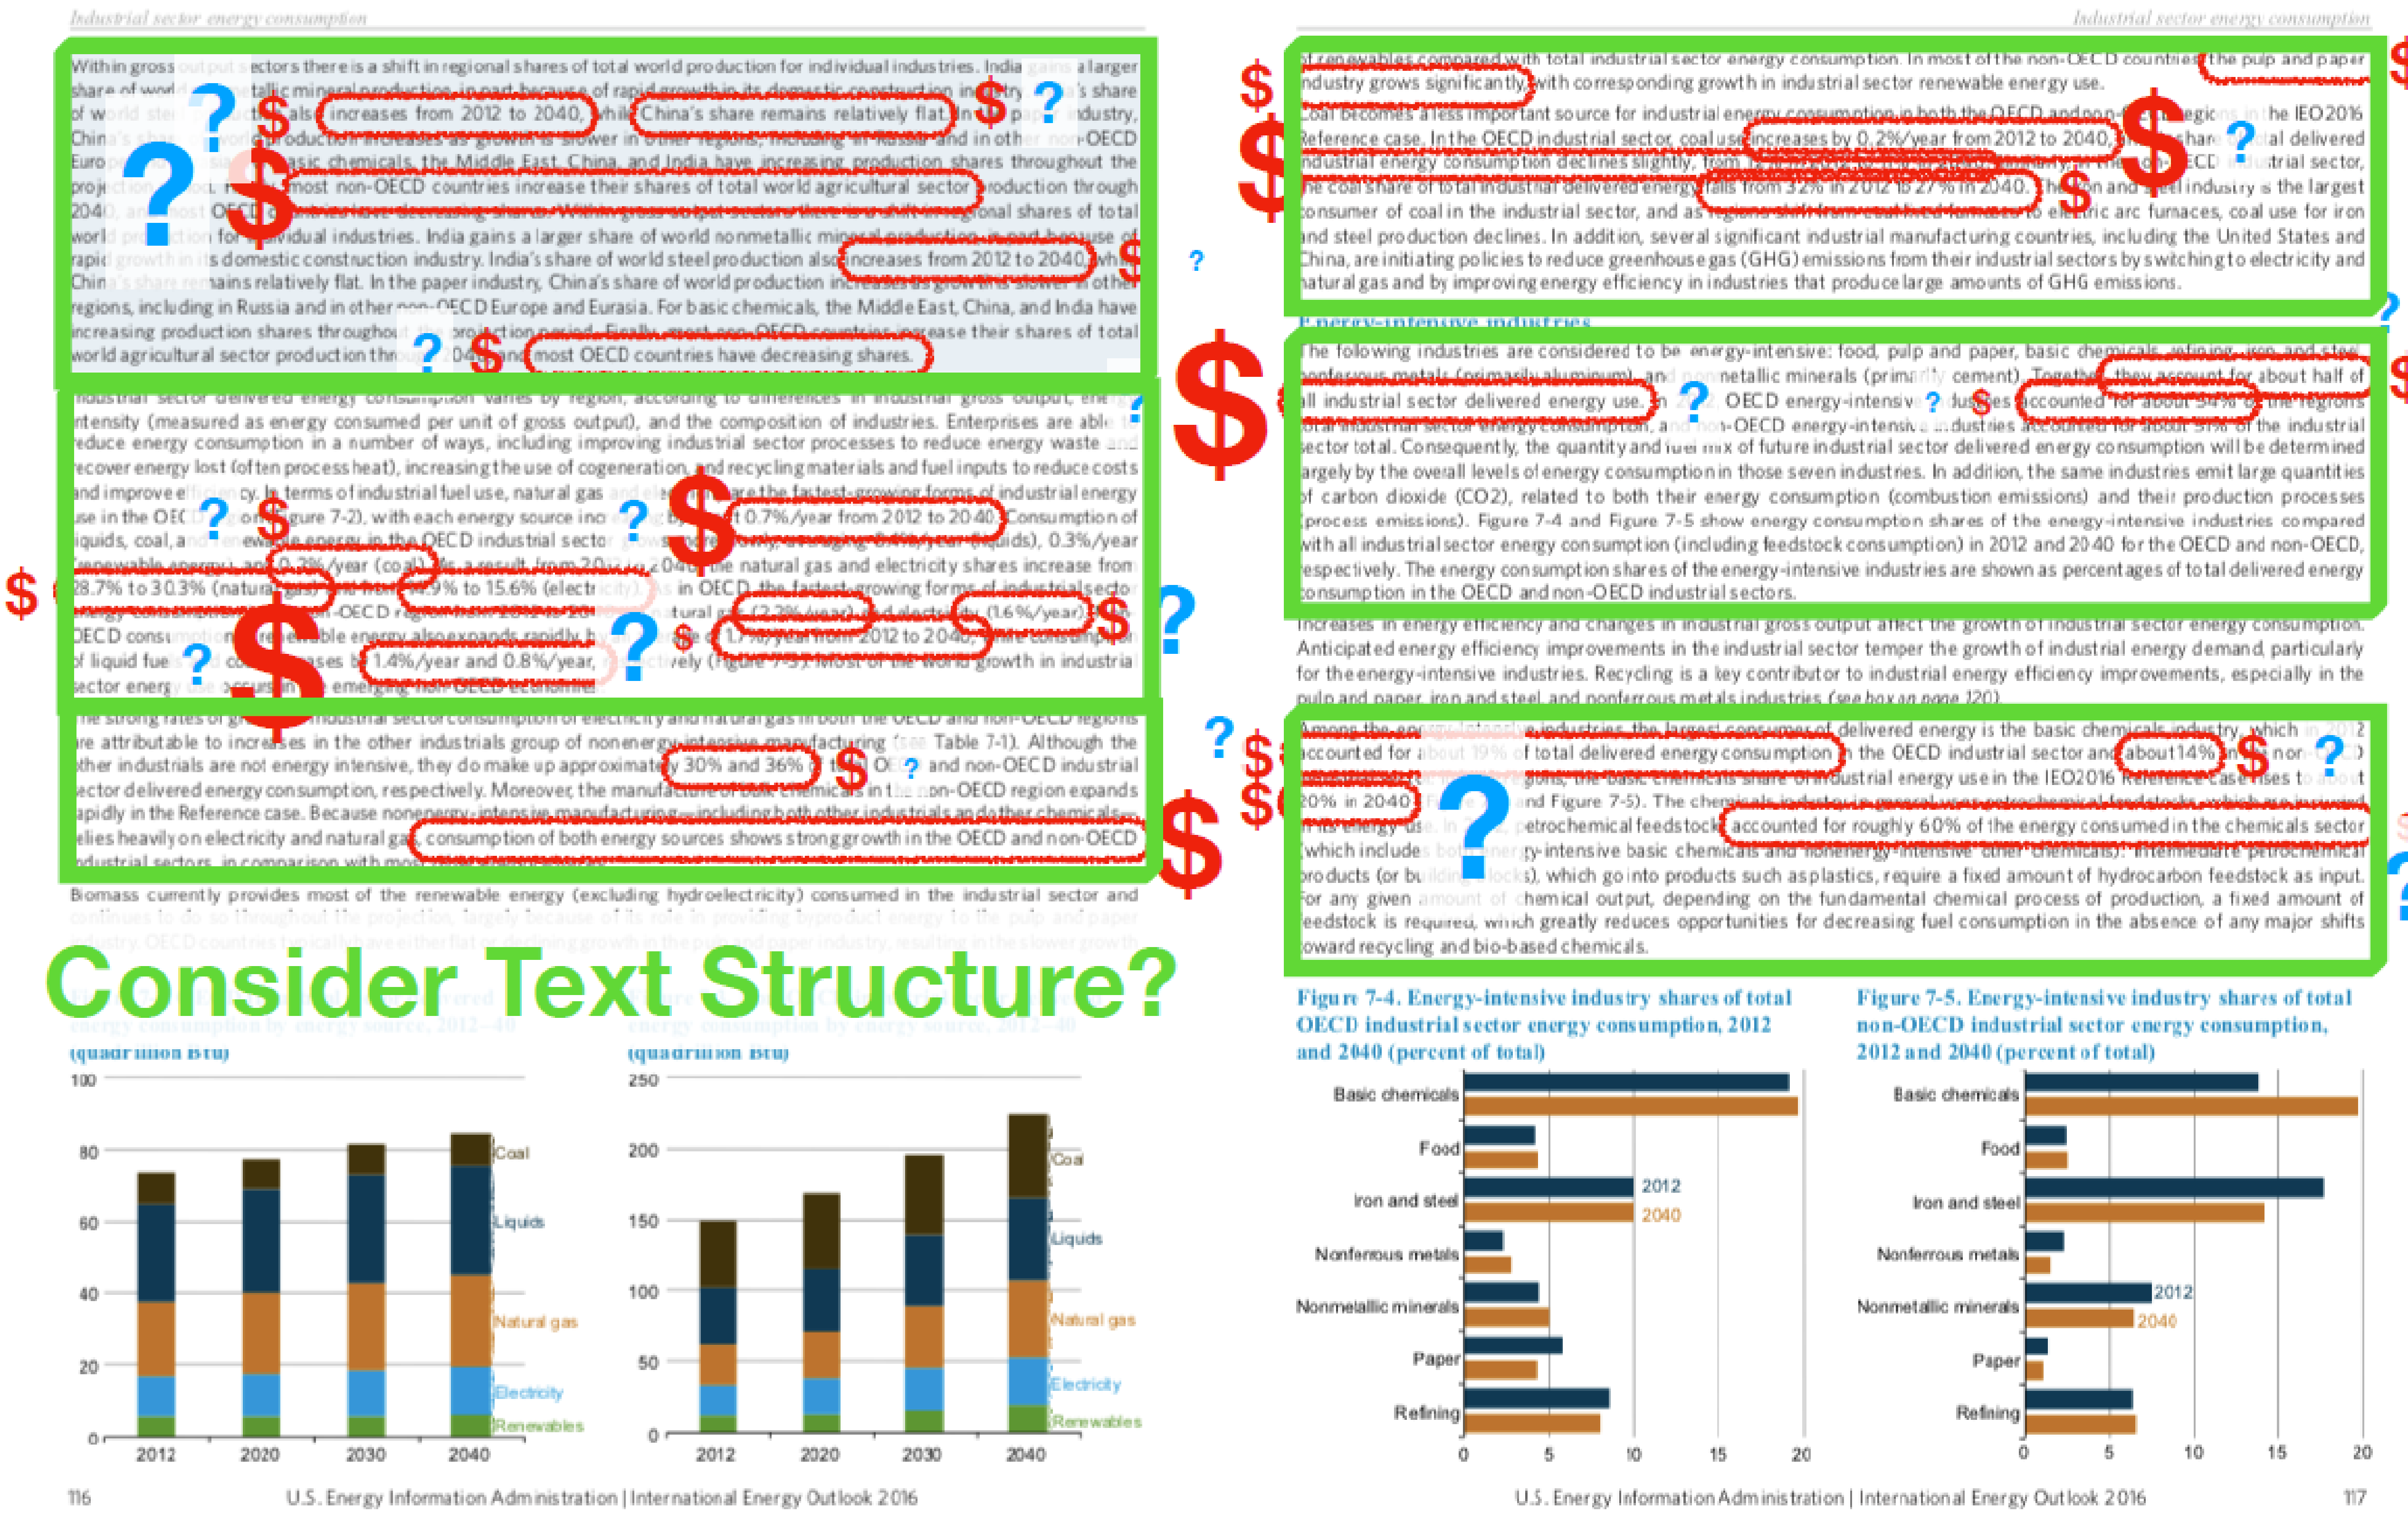

Consider Text Structure?

Verify Interesting Claims First?

Energy claims evaluation

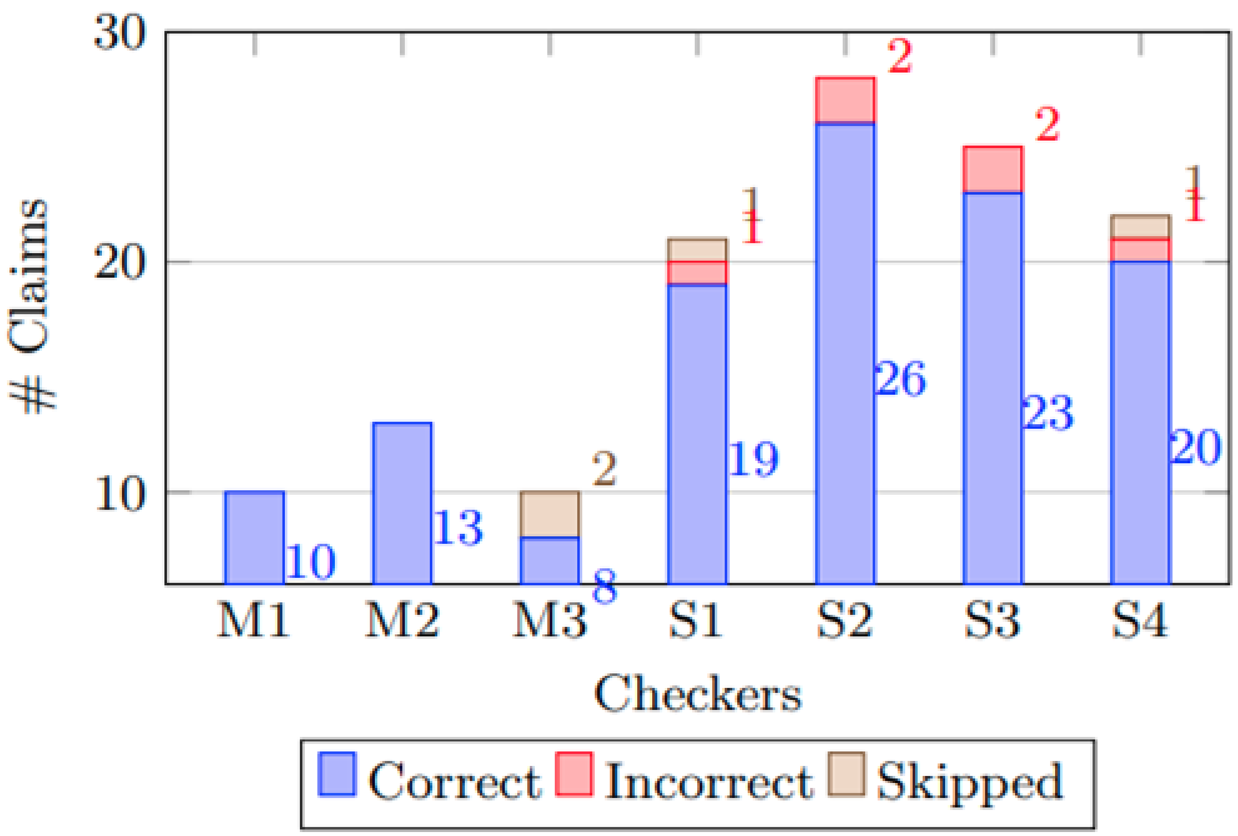

Figure 5: Number of claims verified in 20 minutes by checkers with the manual process (M1–M3) and with our system (S1–S4).

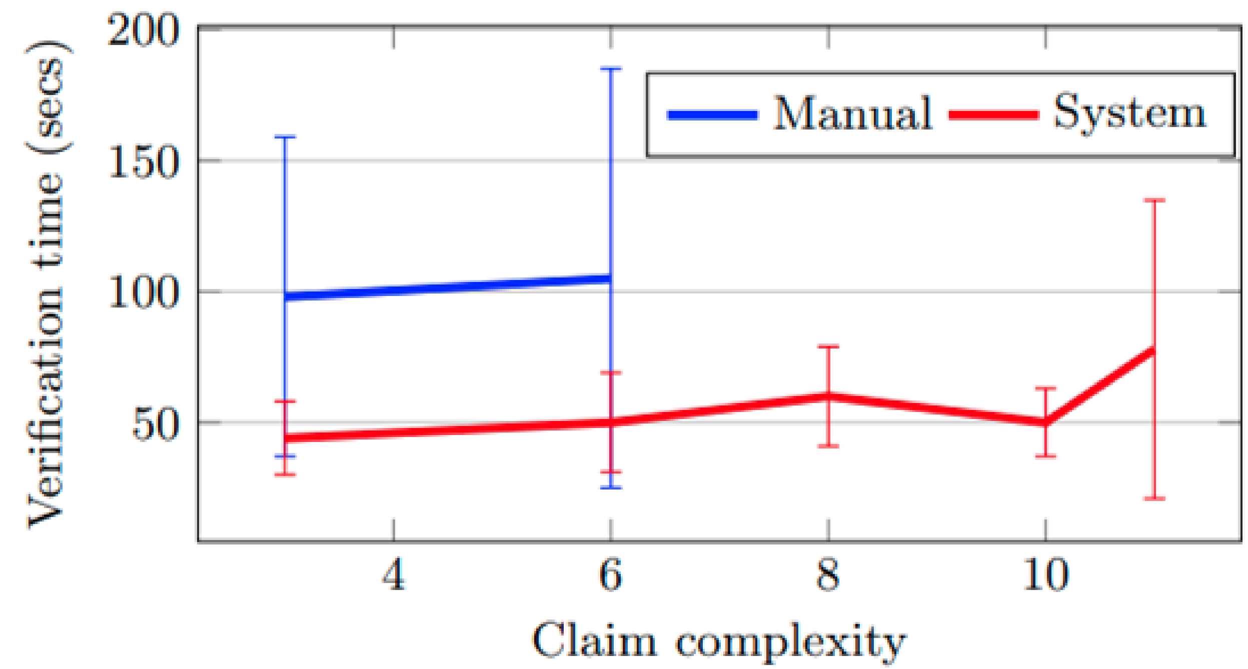

$$\mathrm{CAGR} = \left(\frac{V_{\mathrm{final}}}{V_{\mathrm{begin}}}\right)^{1/t} - 1$$

Compound Annual Growth Rate

[Cao et al WebDB 18, Chen et al ICLR 20, Herzig et al ACL 20]

Experiments

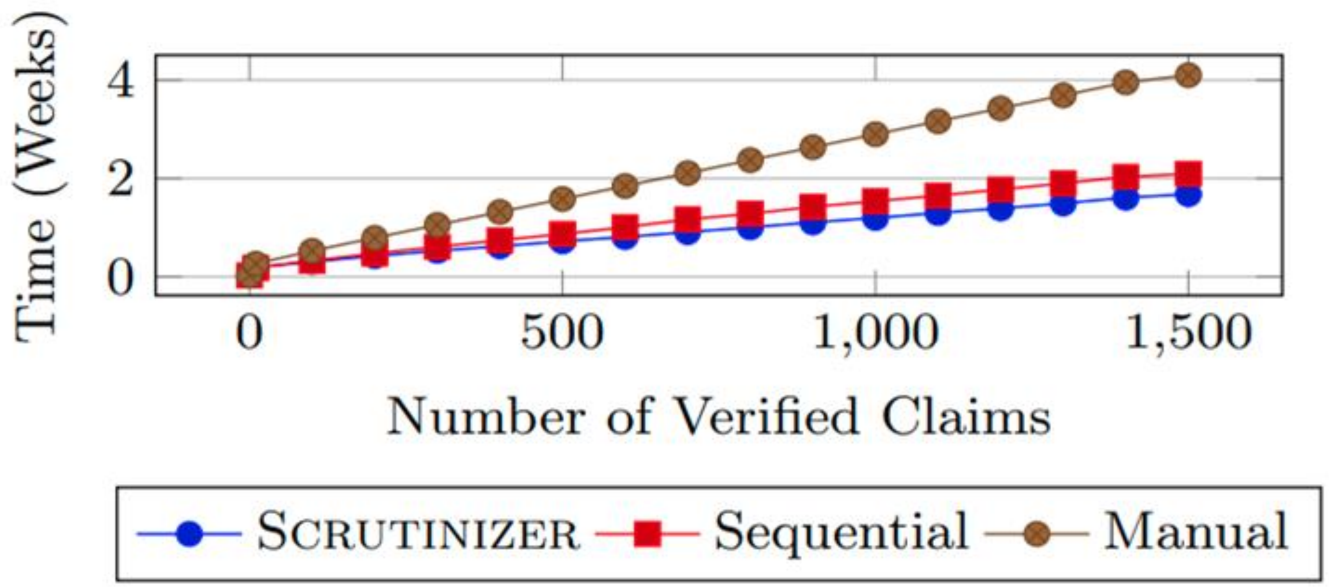

Figure 6: Accumulated verification time over verification period.

Table 4: Summary of simulation results.

	Manual	Sequential	Scrut.
Time (Weeks)	4.1	2.1	1.7
% Savings	-	49%	59%
Avg. Accuracy	-	40%	47%
Max Accuracy	-	46%	53%
Comp. (Mins)	-	14	28

Fact checking entities

- Expert / Journalist
- Computational
- **Crowdsourcing with end users**

Crowdsourced fact checking

WikiTRIBUNE
Evidence-based Journalism

Fiskkit
A better way to discuss the news
Paste article link here

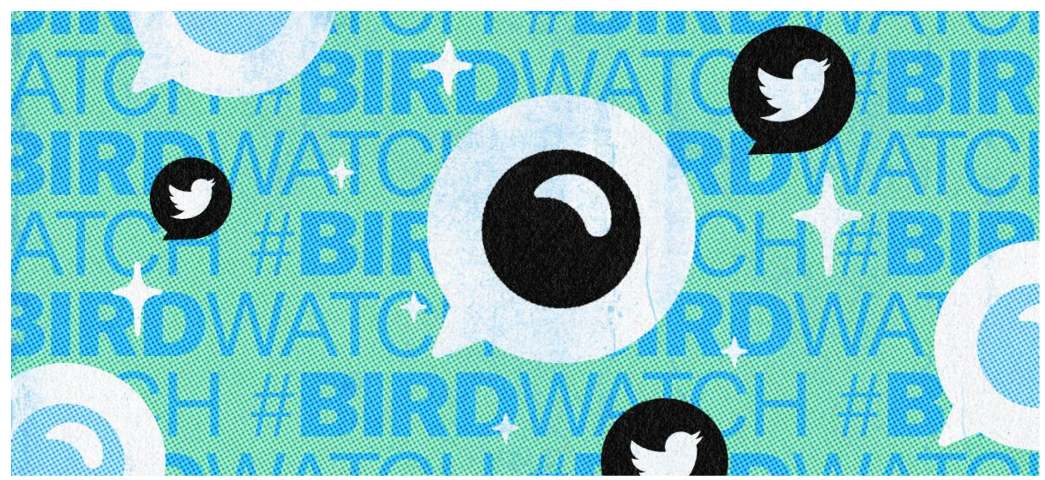

TextThresher

Very few thriving projects

Crowdsourced fact checking

Advantages

- Leverage large number of users
- High scalability
- Faster response

Disadvantages

- Lower credibility
- Risk of manipulation by partisans
- Need to be aware of human biases
- Imbalance in checking popular vs important topics

Birdwatch (Twitter)

- (US only) Community-Based Fact-Checking
- “Qualified” users fact-check tweets → notes
- and rate other notes → ratings

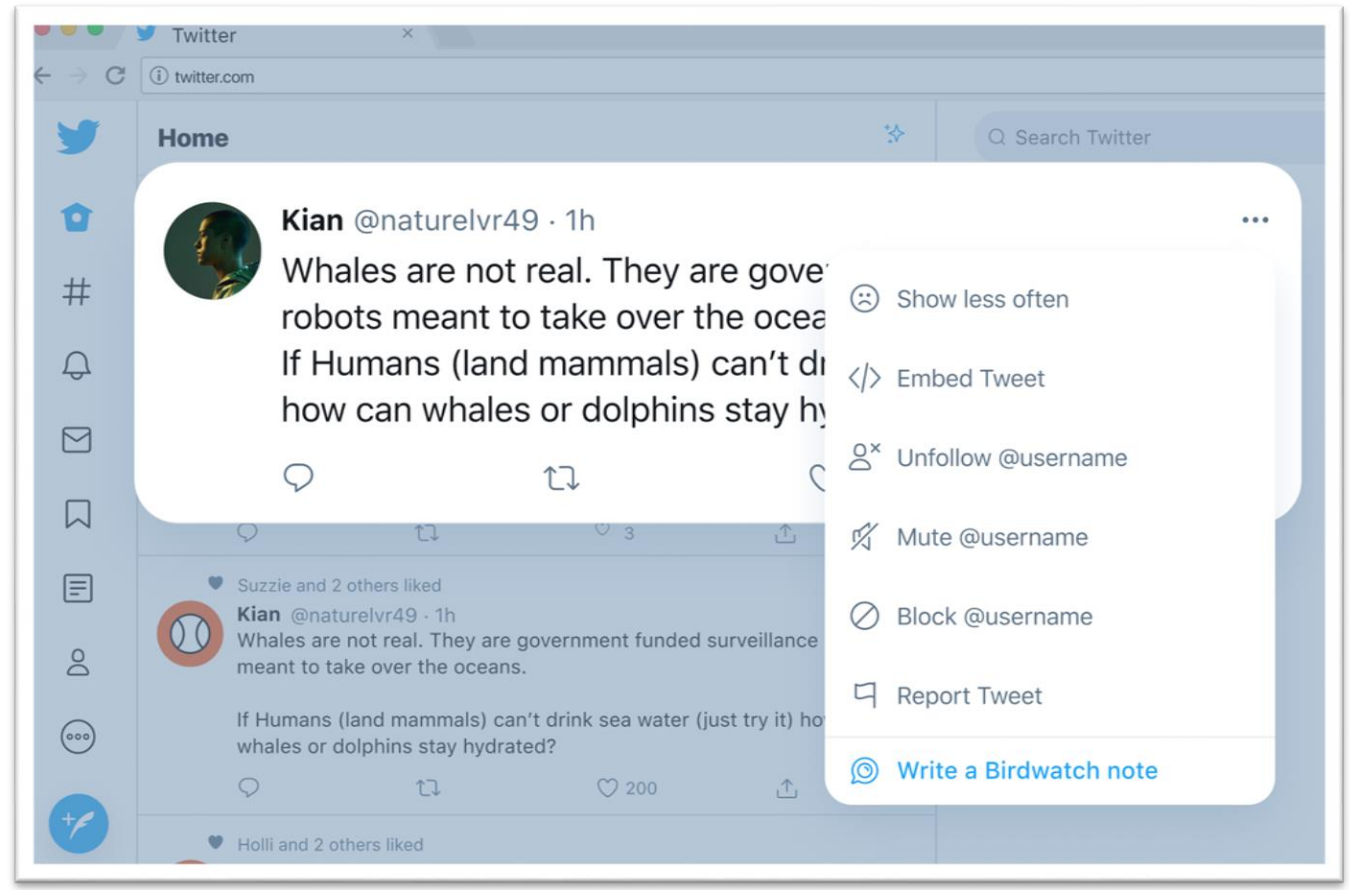

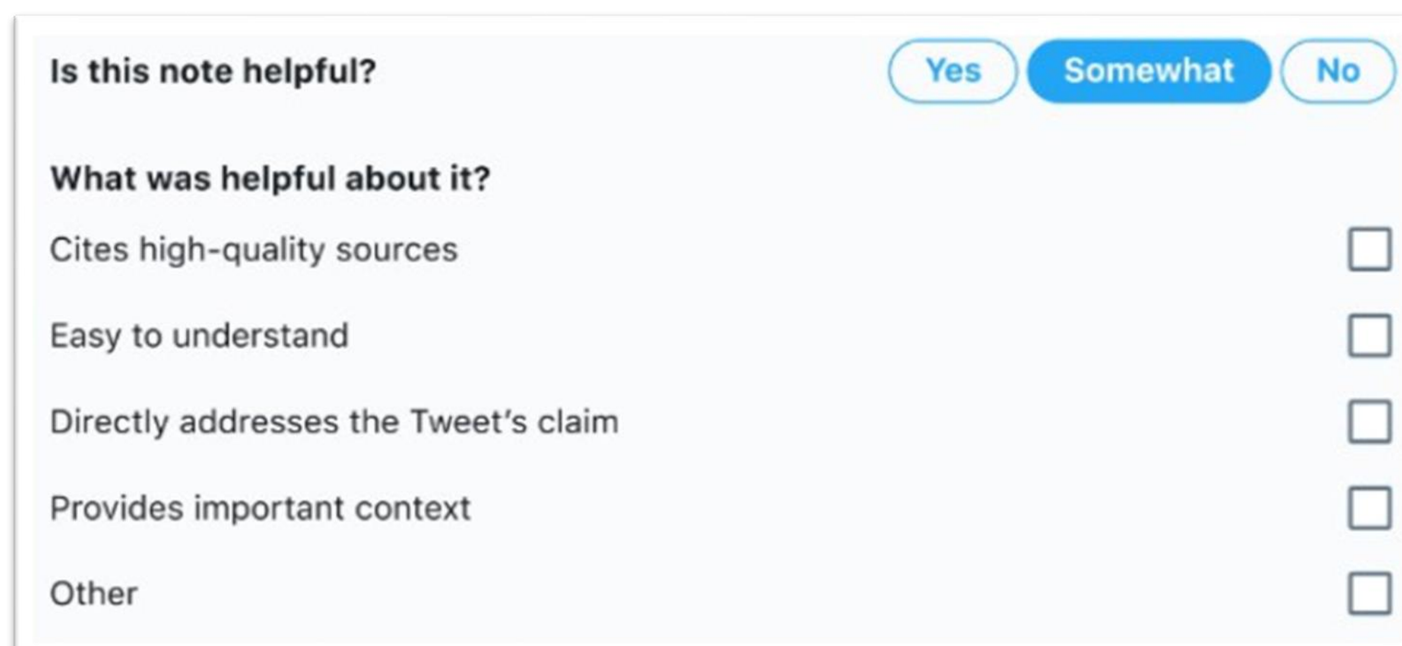

https://twitter.github.io/birdwatch/

How is Birdwatch doing?

1. How are check-worthy claims selected by the crowd?

2. What sources are used to support a checking decision and how reliable are they?

3. Are crowd workers able to reliably assess the veracity of a tweet?

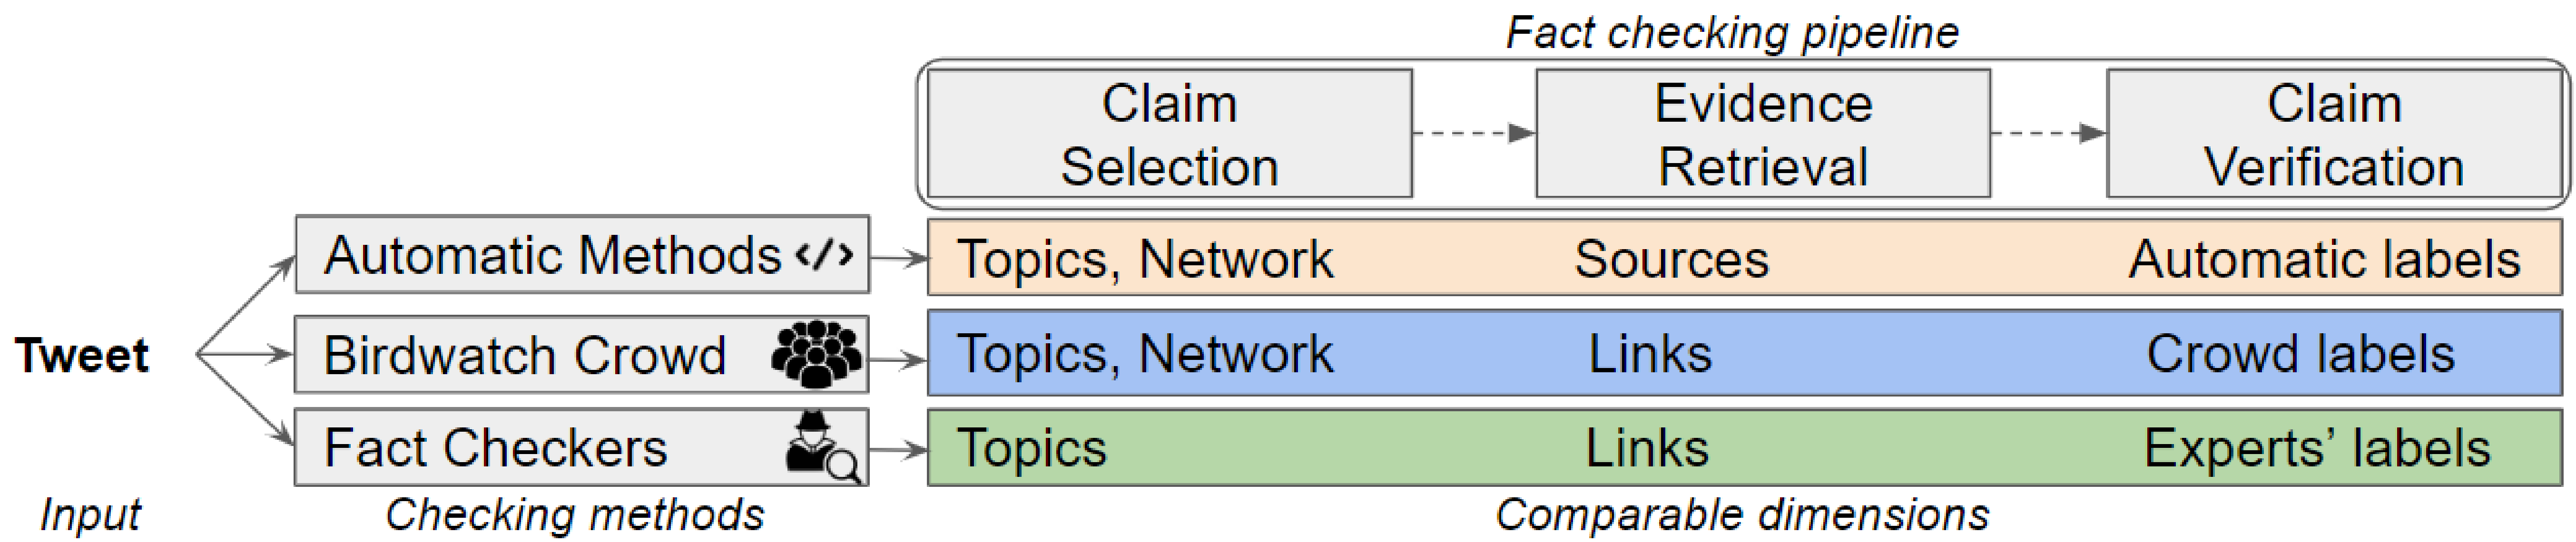

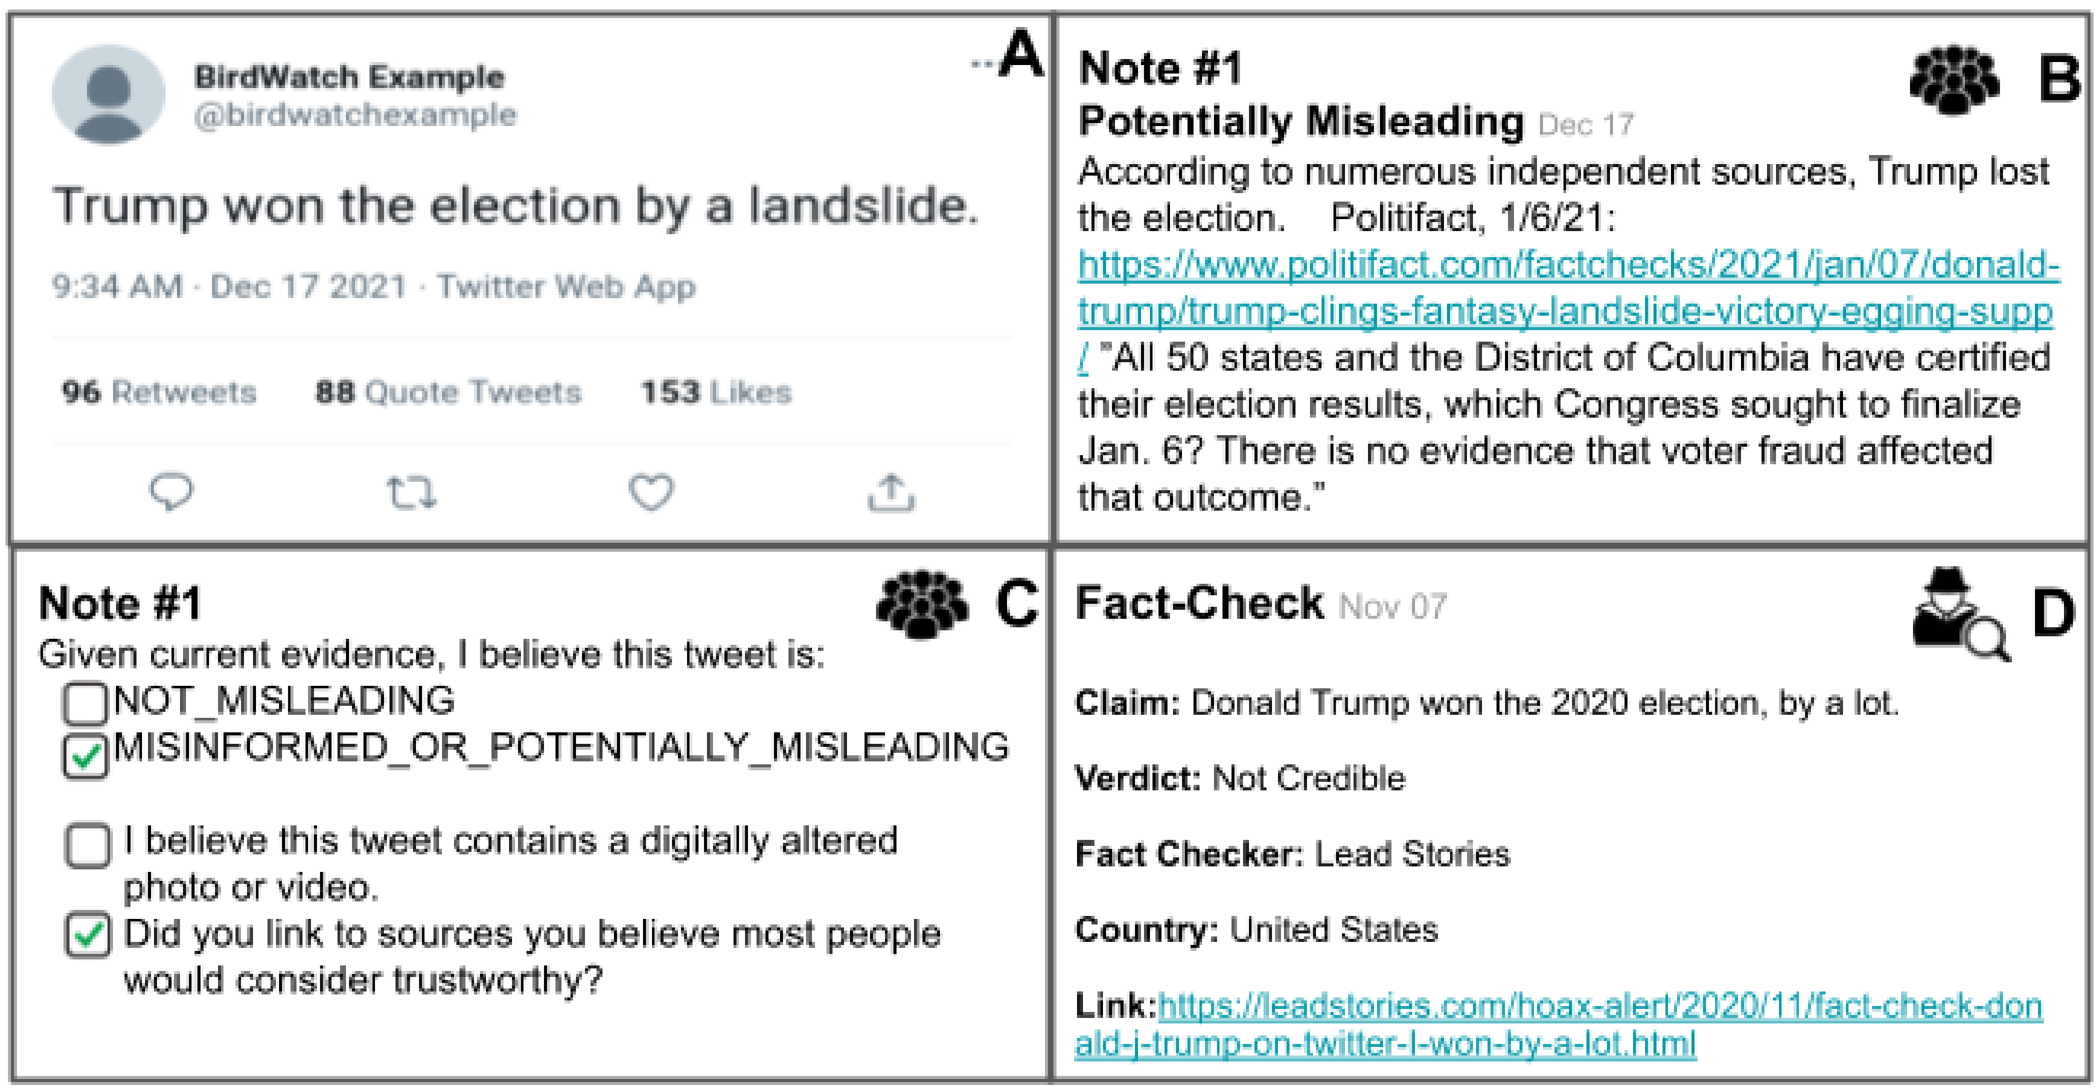

[M. Saeed et al., 2022]

Dataset

- Birdwatch (BW) up to Sept 2021: 16k notes, 87k ratings
- ClaimReview (CR) fact-checks: 77k
- Manual match by authors (500) and crowd (5.5k)
  - 2208 tweets (3043 notes) matching CR checks
    - https://birdwatch.eurecom.fr/

How are check-worthy claims selected by the crowd?

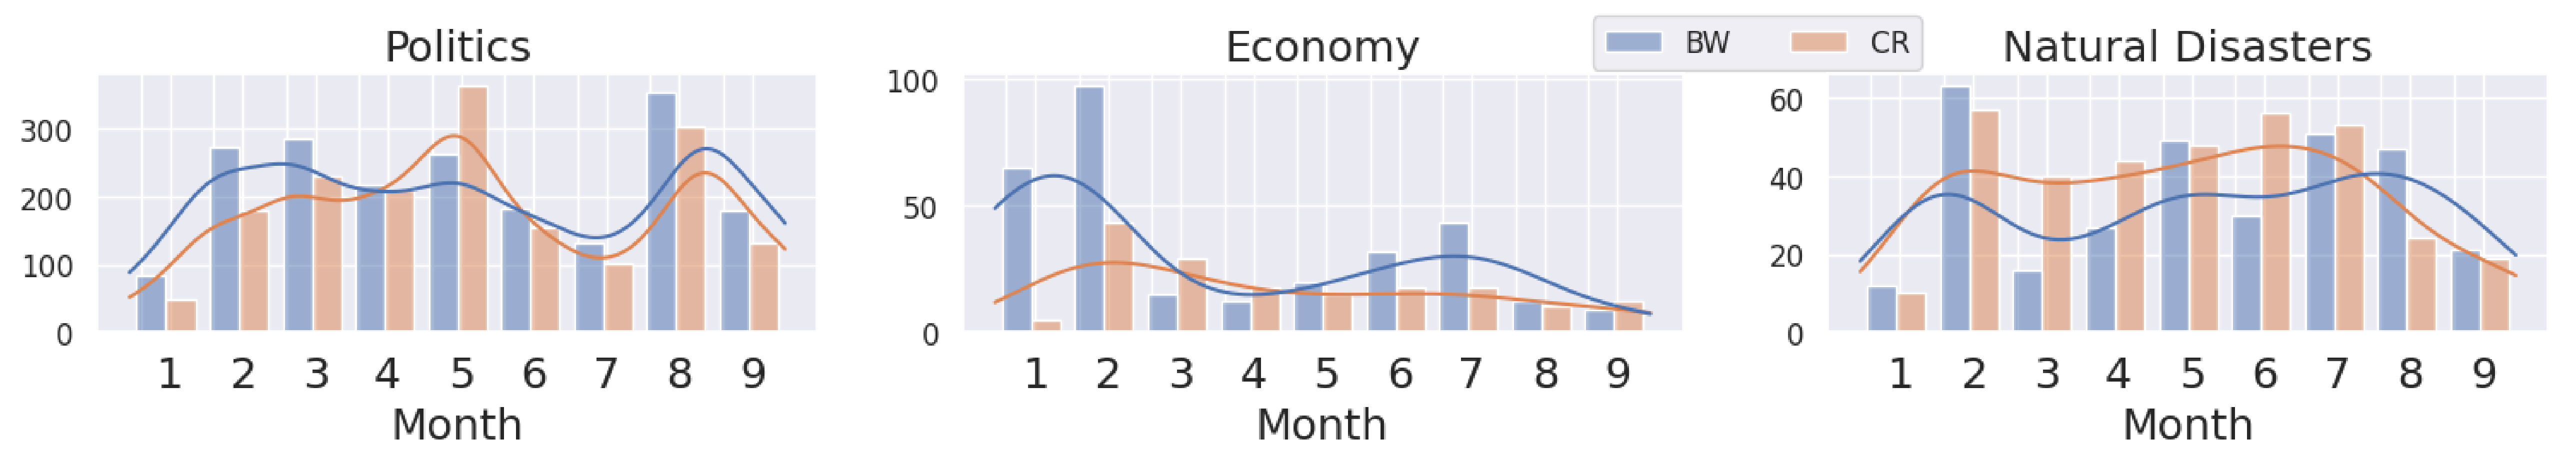

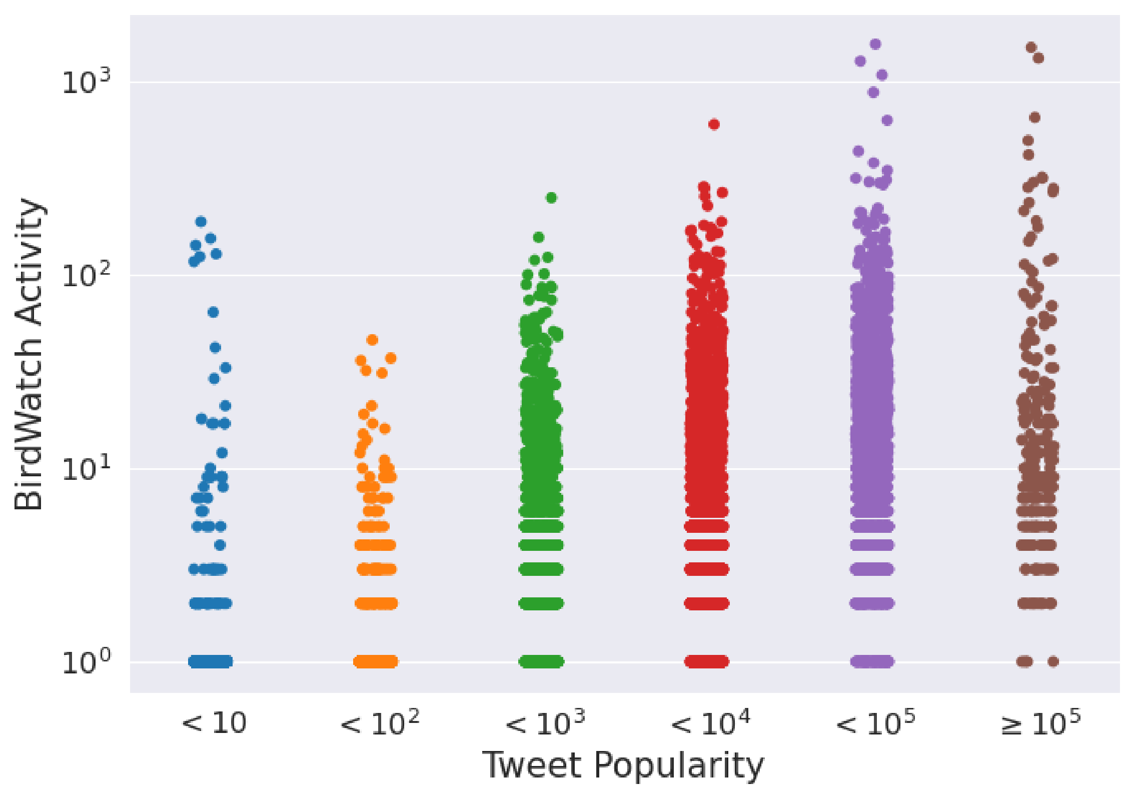

What sources support a decision and how reliable are they?

- BW: 13k links covering 2k domains
  Skew: half links from 29 domains

- CR: 77k links covering only 73 domains of fact-checking groups and journalists

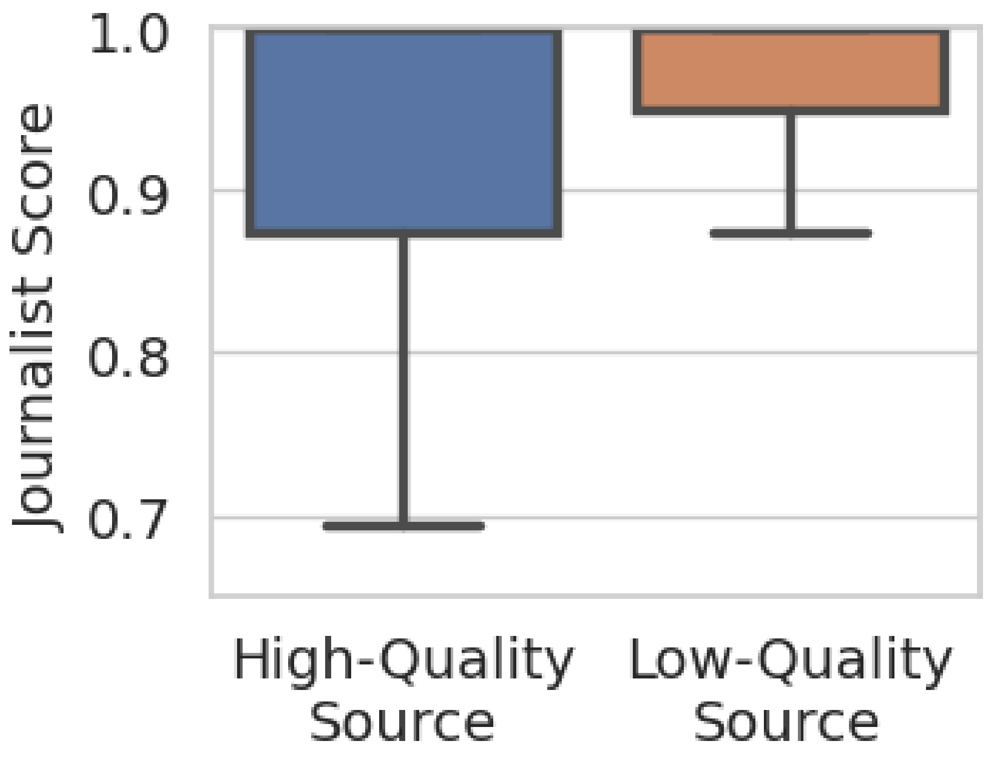

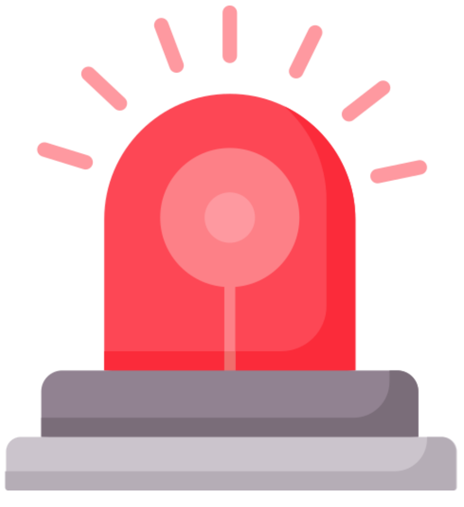

Are crowd workers able to assess the veracity of a tweet?

		BirdWatch				
		Notes		Tweets		
		MM	**NM**	**MM**	**NM**	**Tie**
Claim Review	**credible**	209	**25**	126	**9**	9
	mostly_credible	56	**14**	44	**7**	5
	not_credible	**1983**	184	**1476**	62	55
	not_verifiable	300	25	225	8	9
	uncertain	225	22	156	8	9

- majority of BW labels match the CR ones

Ready for adoption?

[IJCAI'21 survey https://arxiv.org/pdf/2103.07769.pdf]

Automated Fact-Checking for Assisting Human Fact-Checkers

Preslav Nakov[1*], **David Corney**[2], **Maram Hasanain**[3], **Firoj Alam**[1], **Tamer Elsayed**[3], **Alberto Barrón-Cedeño**[4], **Paolo Papotti**[5], **Shaden Shaar**[1], **Giovanni Da San Martino**[6]
[1]Qatar Computing Research Institute, HBKU, Qatar, [2]Full Fact, UK, [3]Qatar University, Qatar, [4]Università di Bologna, Italy, [5]EURECOM, France, [6]University of Padova, Italy
{pnakov, fialam, sshaar}@hbku.edu.qa, david.corney@fullfact.org, a.barron@unibo.it, {maram.hasanain, telsayed}@qu.edu.qa, papotti@eurecom.fr, dasan@math.unipd.it

Abstract

The reporting and analysis of current events around the globe has expanded from professional, editor-lead journalism all the way to citizen journalism. Politicians and other key players enjoy direct access to their audiences through social media, bypassing the filters of official cables or traditional media. However, the multiple advantages of free speech and direct communication are dimmed by the misuse of the media to spread inaccurate or misleading claims. These phenomena have led to the modern incarnation of the *fact-checker* — a professional whose main aim is to examine claims using

Such organizations are also potential beneficiaries of or leaders in automated fact-checking research. As misinformation became a major concern globally, tech companies, national and international agencies began work in the area. Recently, several international initiatives have also emerged such as the *Credibility Coalition*[5] and *EUfactcheck*,[6]. Along side some tools have also been made available such as Google Factcheck[7] and Hoaxy[8]. Moreover, fact-checking is a common task in settings that go beyond online misinformation, as the verification of content accuracy is a priority for many organizations [Karagiannis *et al.*, 2020].

A large body of research is devoted to developing automatic systems for fact-checking [Vo and Lee, 2018; Shu

Research directions

- Applicability & Usability

A_C
- More expressive power
  [Mori et al, 2022] [Aly et al, 2021]

A_C
- Ambiguity in NL
  [Veltri et al, 2022][Wenzel, 2019]

- Explanations in NL
  [Atanasova et al, 2020]
  [Kotonya and Toni, 2020]

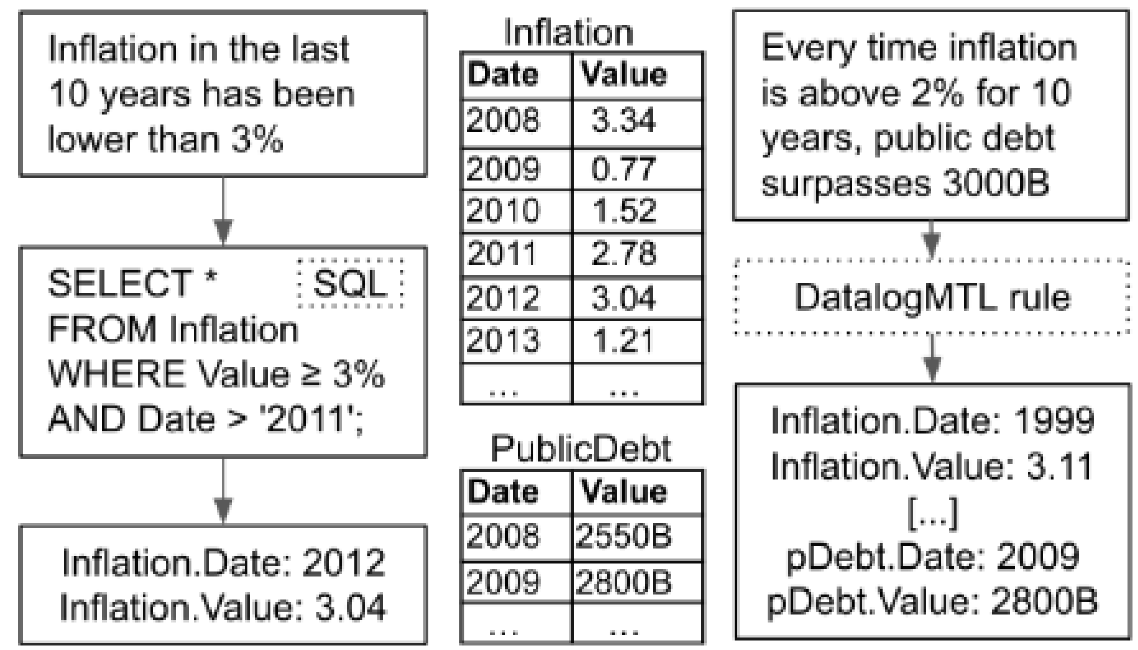

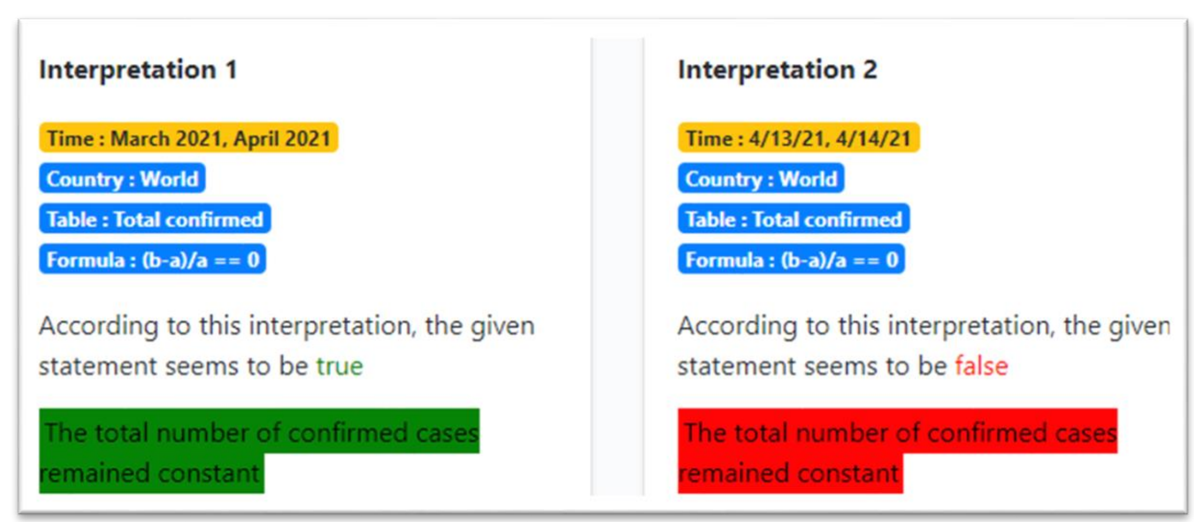

Research directions

- Automatic methods poor with NL text (*Twitter slang)* and at matching *stale claims* → *crowd is good at this!* A_C
- Crowdsourcing and aut. methods fail with subtle textual claims → *experts are good at this!* A_C
- Experts are slow and rely on few sources → *crowd and aut. methods scale!* A_S
- Aut. method X is good on claims of type Y given clean data Z…. A_C A_S
- How to combine these approaches and methods? Collaborative solutions [Qu et al 2021]

Research directions

- Automatic tools fact check in real-time all content, provide evidence, possible labels
- Platform users provide a first line of defense
- Experts provide quality checks in every step
- *Assuming we model trust and cost for all actors*
    - Design novel hybrid human-machine solutions that coordinate this joint effort!

(more) Research directions

- Multi language
- System bias
- Multimodality
- Explainability
- Real time
- …

[redacted]
Yesterd

A few mins ago I wa
"information you ha
clicked the pop-up
If this is their latest
pretty pathetic, like
stupid pop-up and
vaccines!

Additional reporti

Sometimes, different news
story. You can take a look
sites.

Content that you shared

SCIENCE FE
False: Cu
do not ca

Learn more about how fa

Like Share

119

[redacted]
5 h ·

Tired of fact checkers censoring posts claiming it's partly false when it's mostly true?! Follow the below mentioned cheat code.

Go to settings>privacy>blocking>search for fact check/fact checkers and block them all. Since I have done this none of my stuff have been censored... I keep checking every once in a while since more popped up I have blocked like 50 so far.
A way around censorship. Spread the knowledge.

12:35

for being so concerned about
one's post. Give me a break! You
to trust you?

46° 87% 08:57

ional Reporting

nal Reporting
ntent from
lnews.com

ifferent news sources
ng perspectives on the
. You can check out
rting from independent
hecking sites.

you shared has been
by Science Feedback.

Science Feedback Fact-Checker
**False: US measles outbreak due
to unvaccinated Americans who**

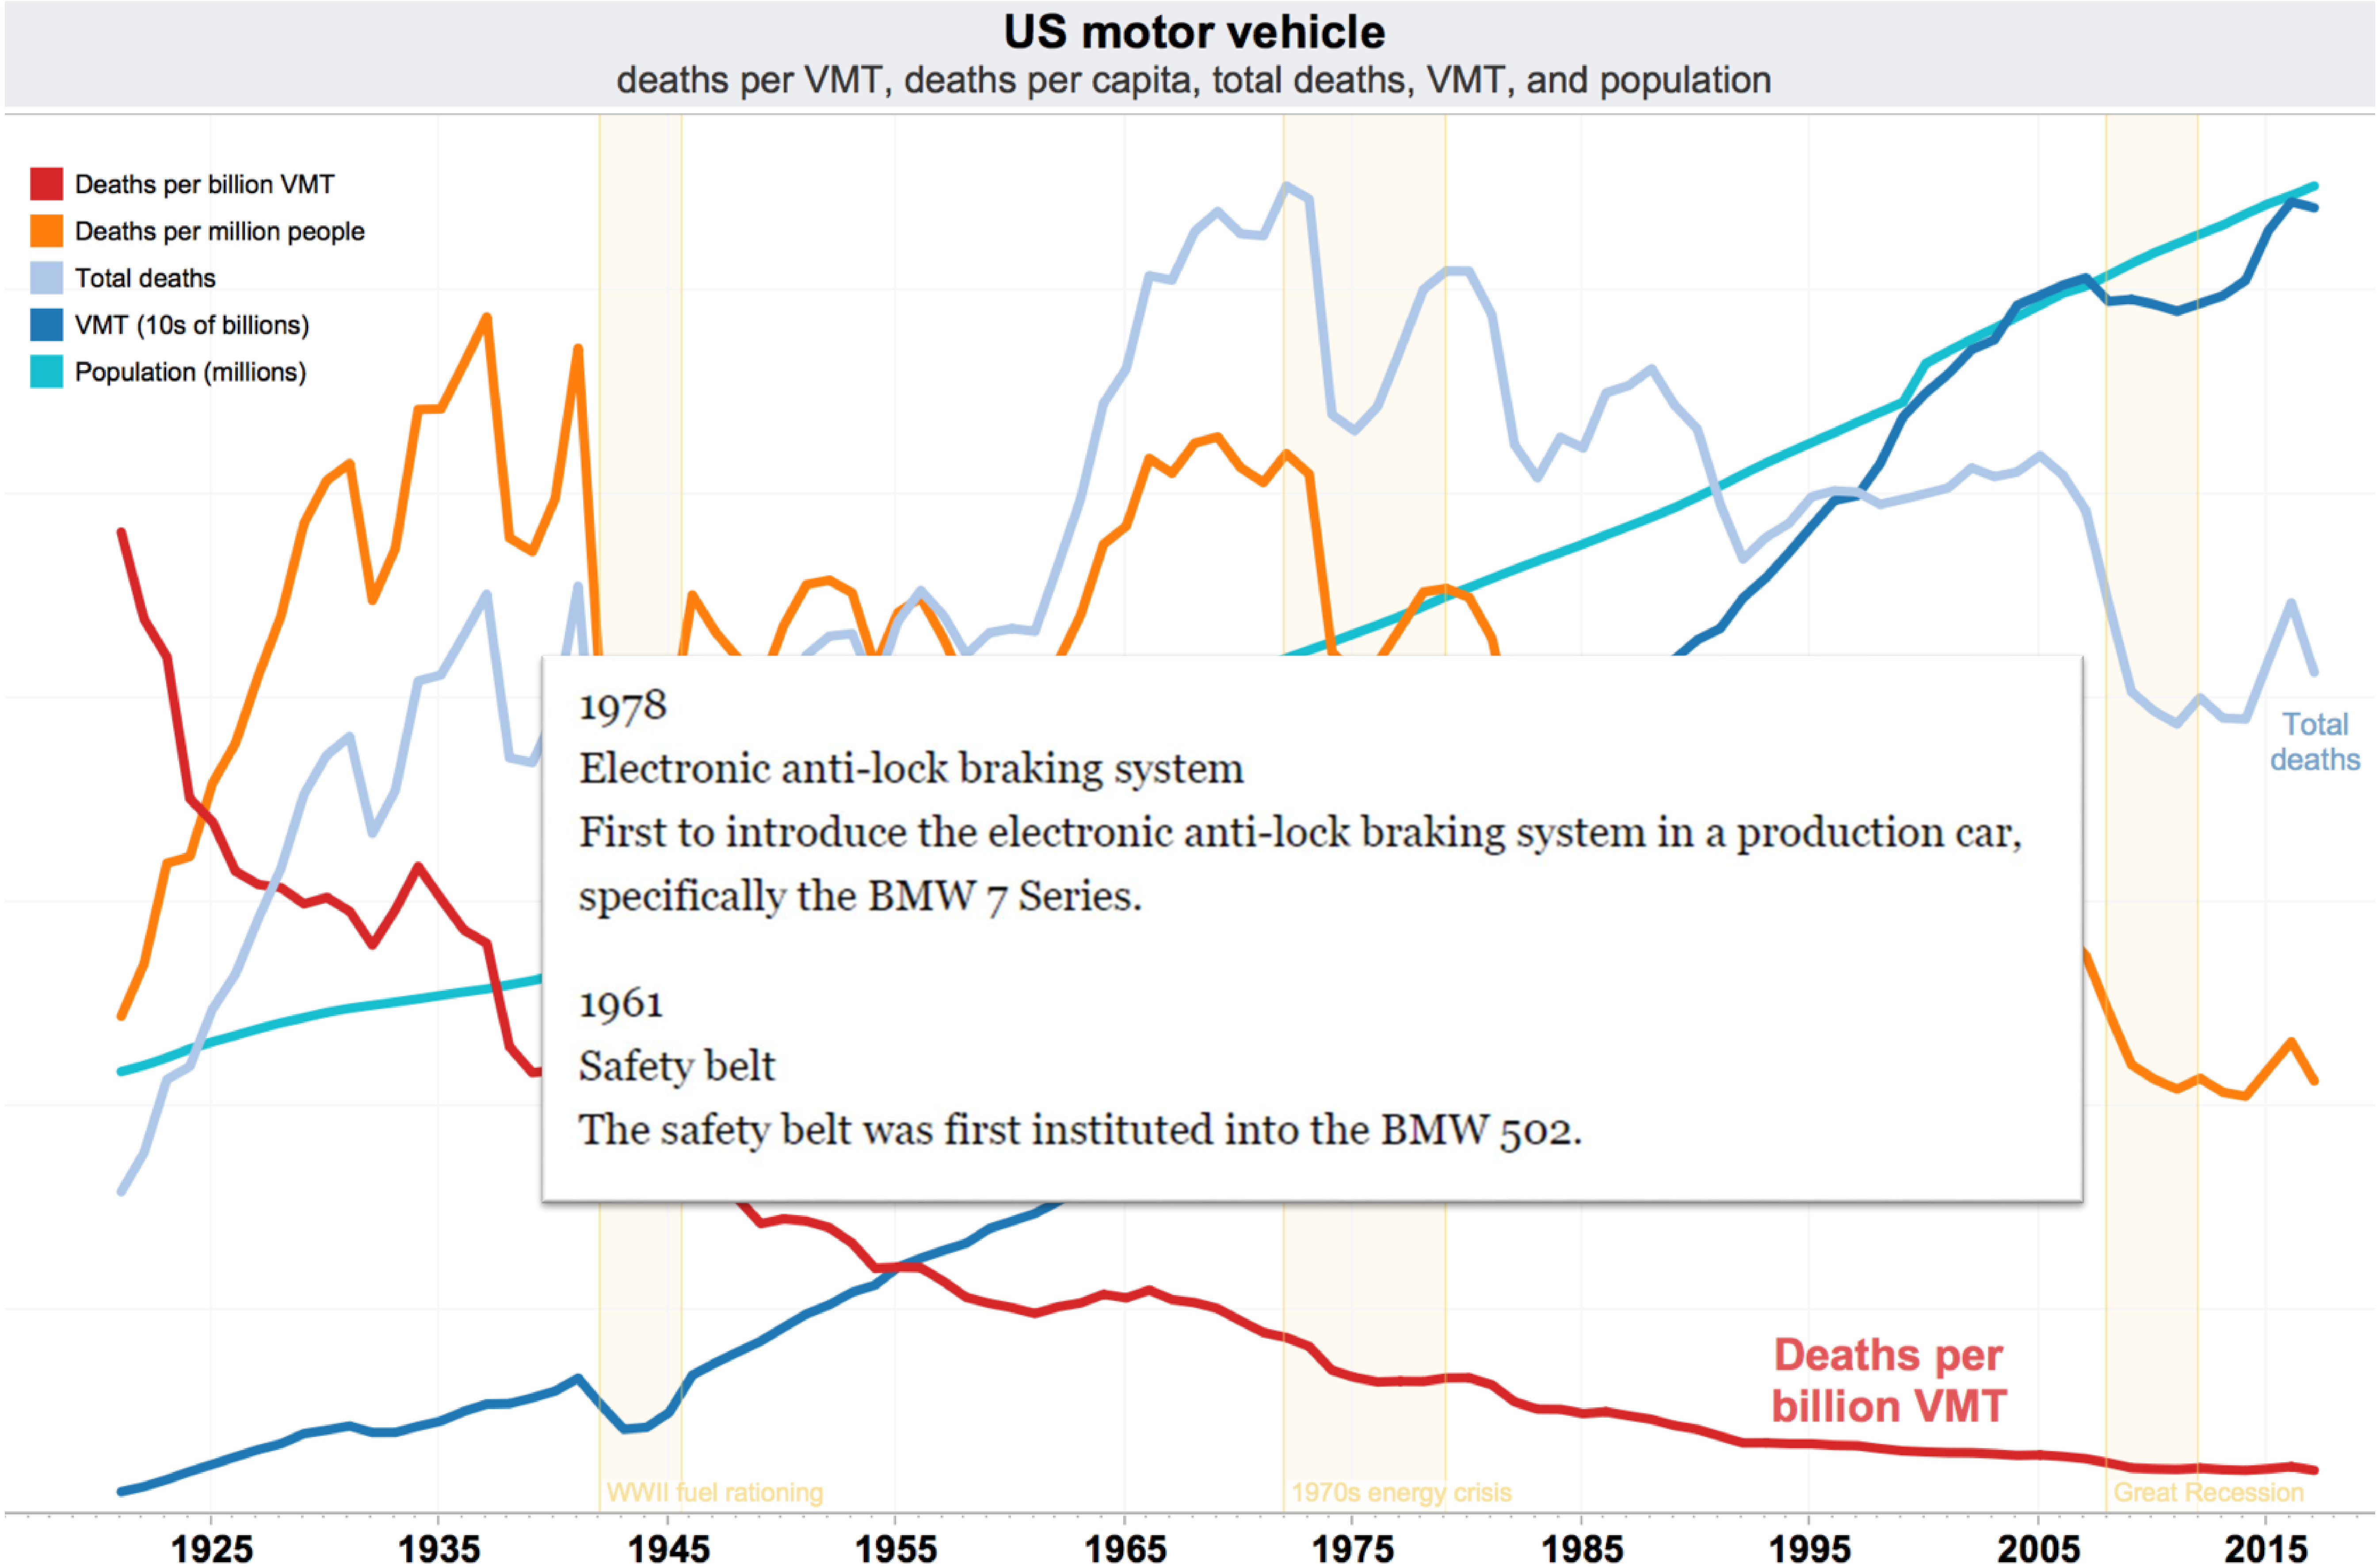

https://en.wikipedia.org/wiki/Motor_vehicle_fatality_rate_in_U.S._by_year

Actors and actions to reduce car deaths

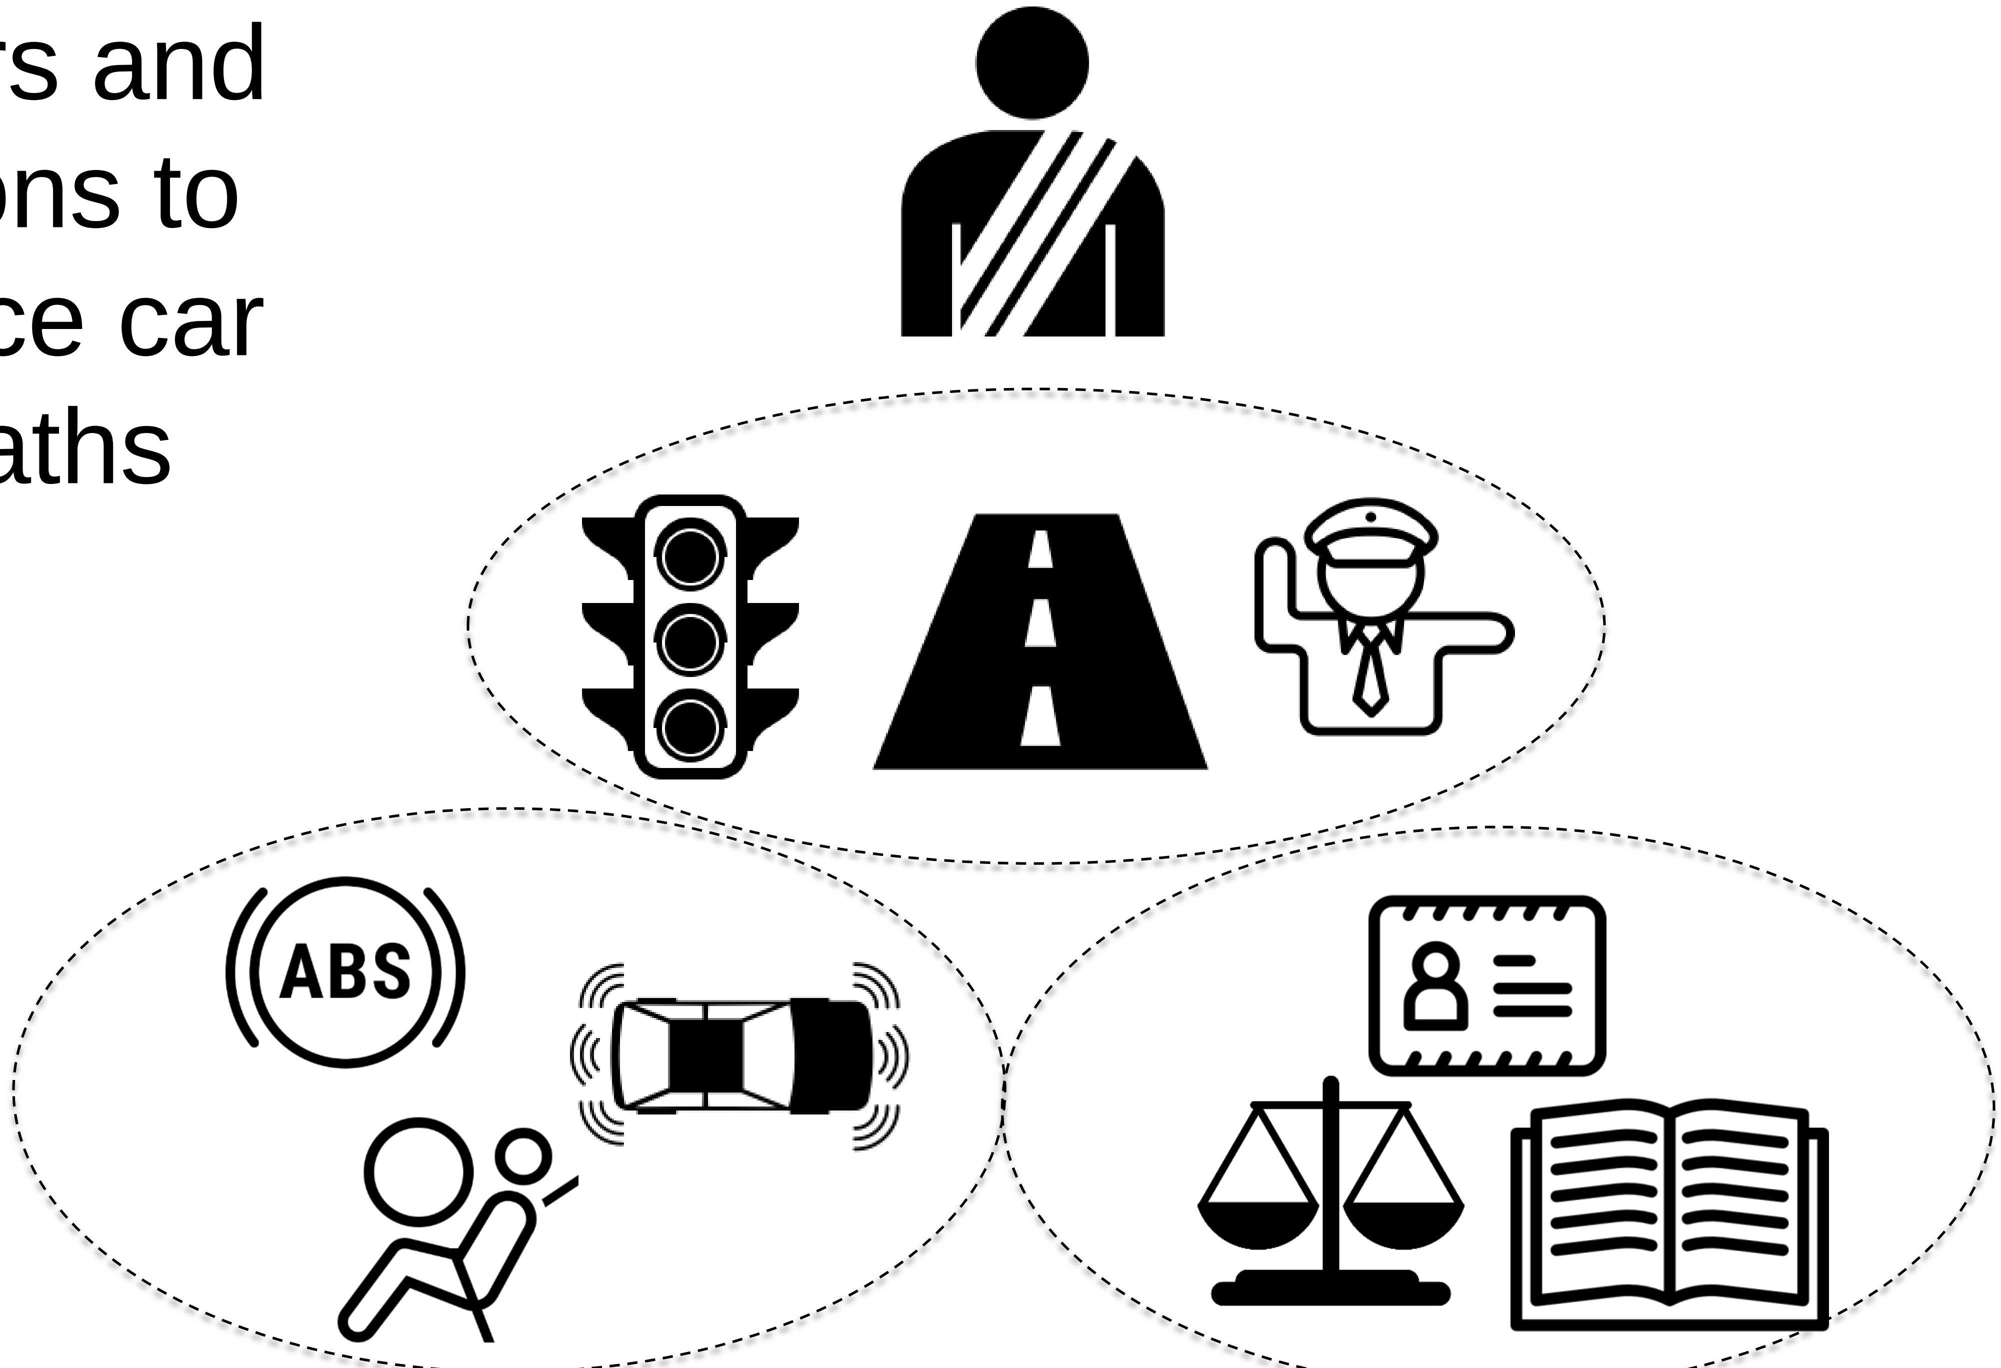

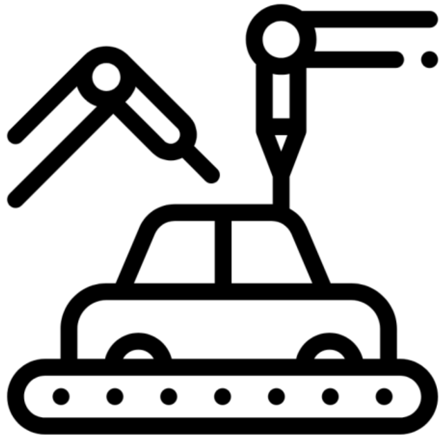

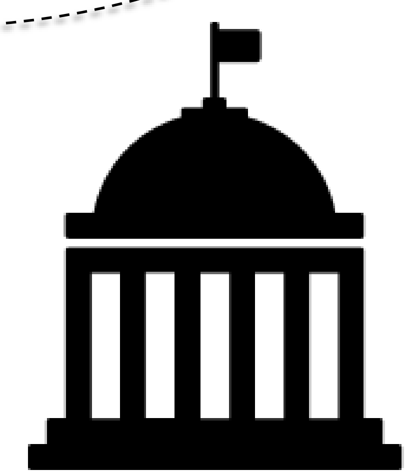

Actors and actions to *reduce misinformation*

2021
EU: Digital Services Act
UK: Online Safety Act

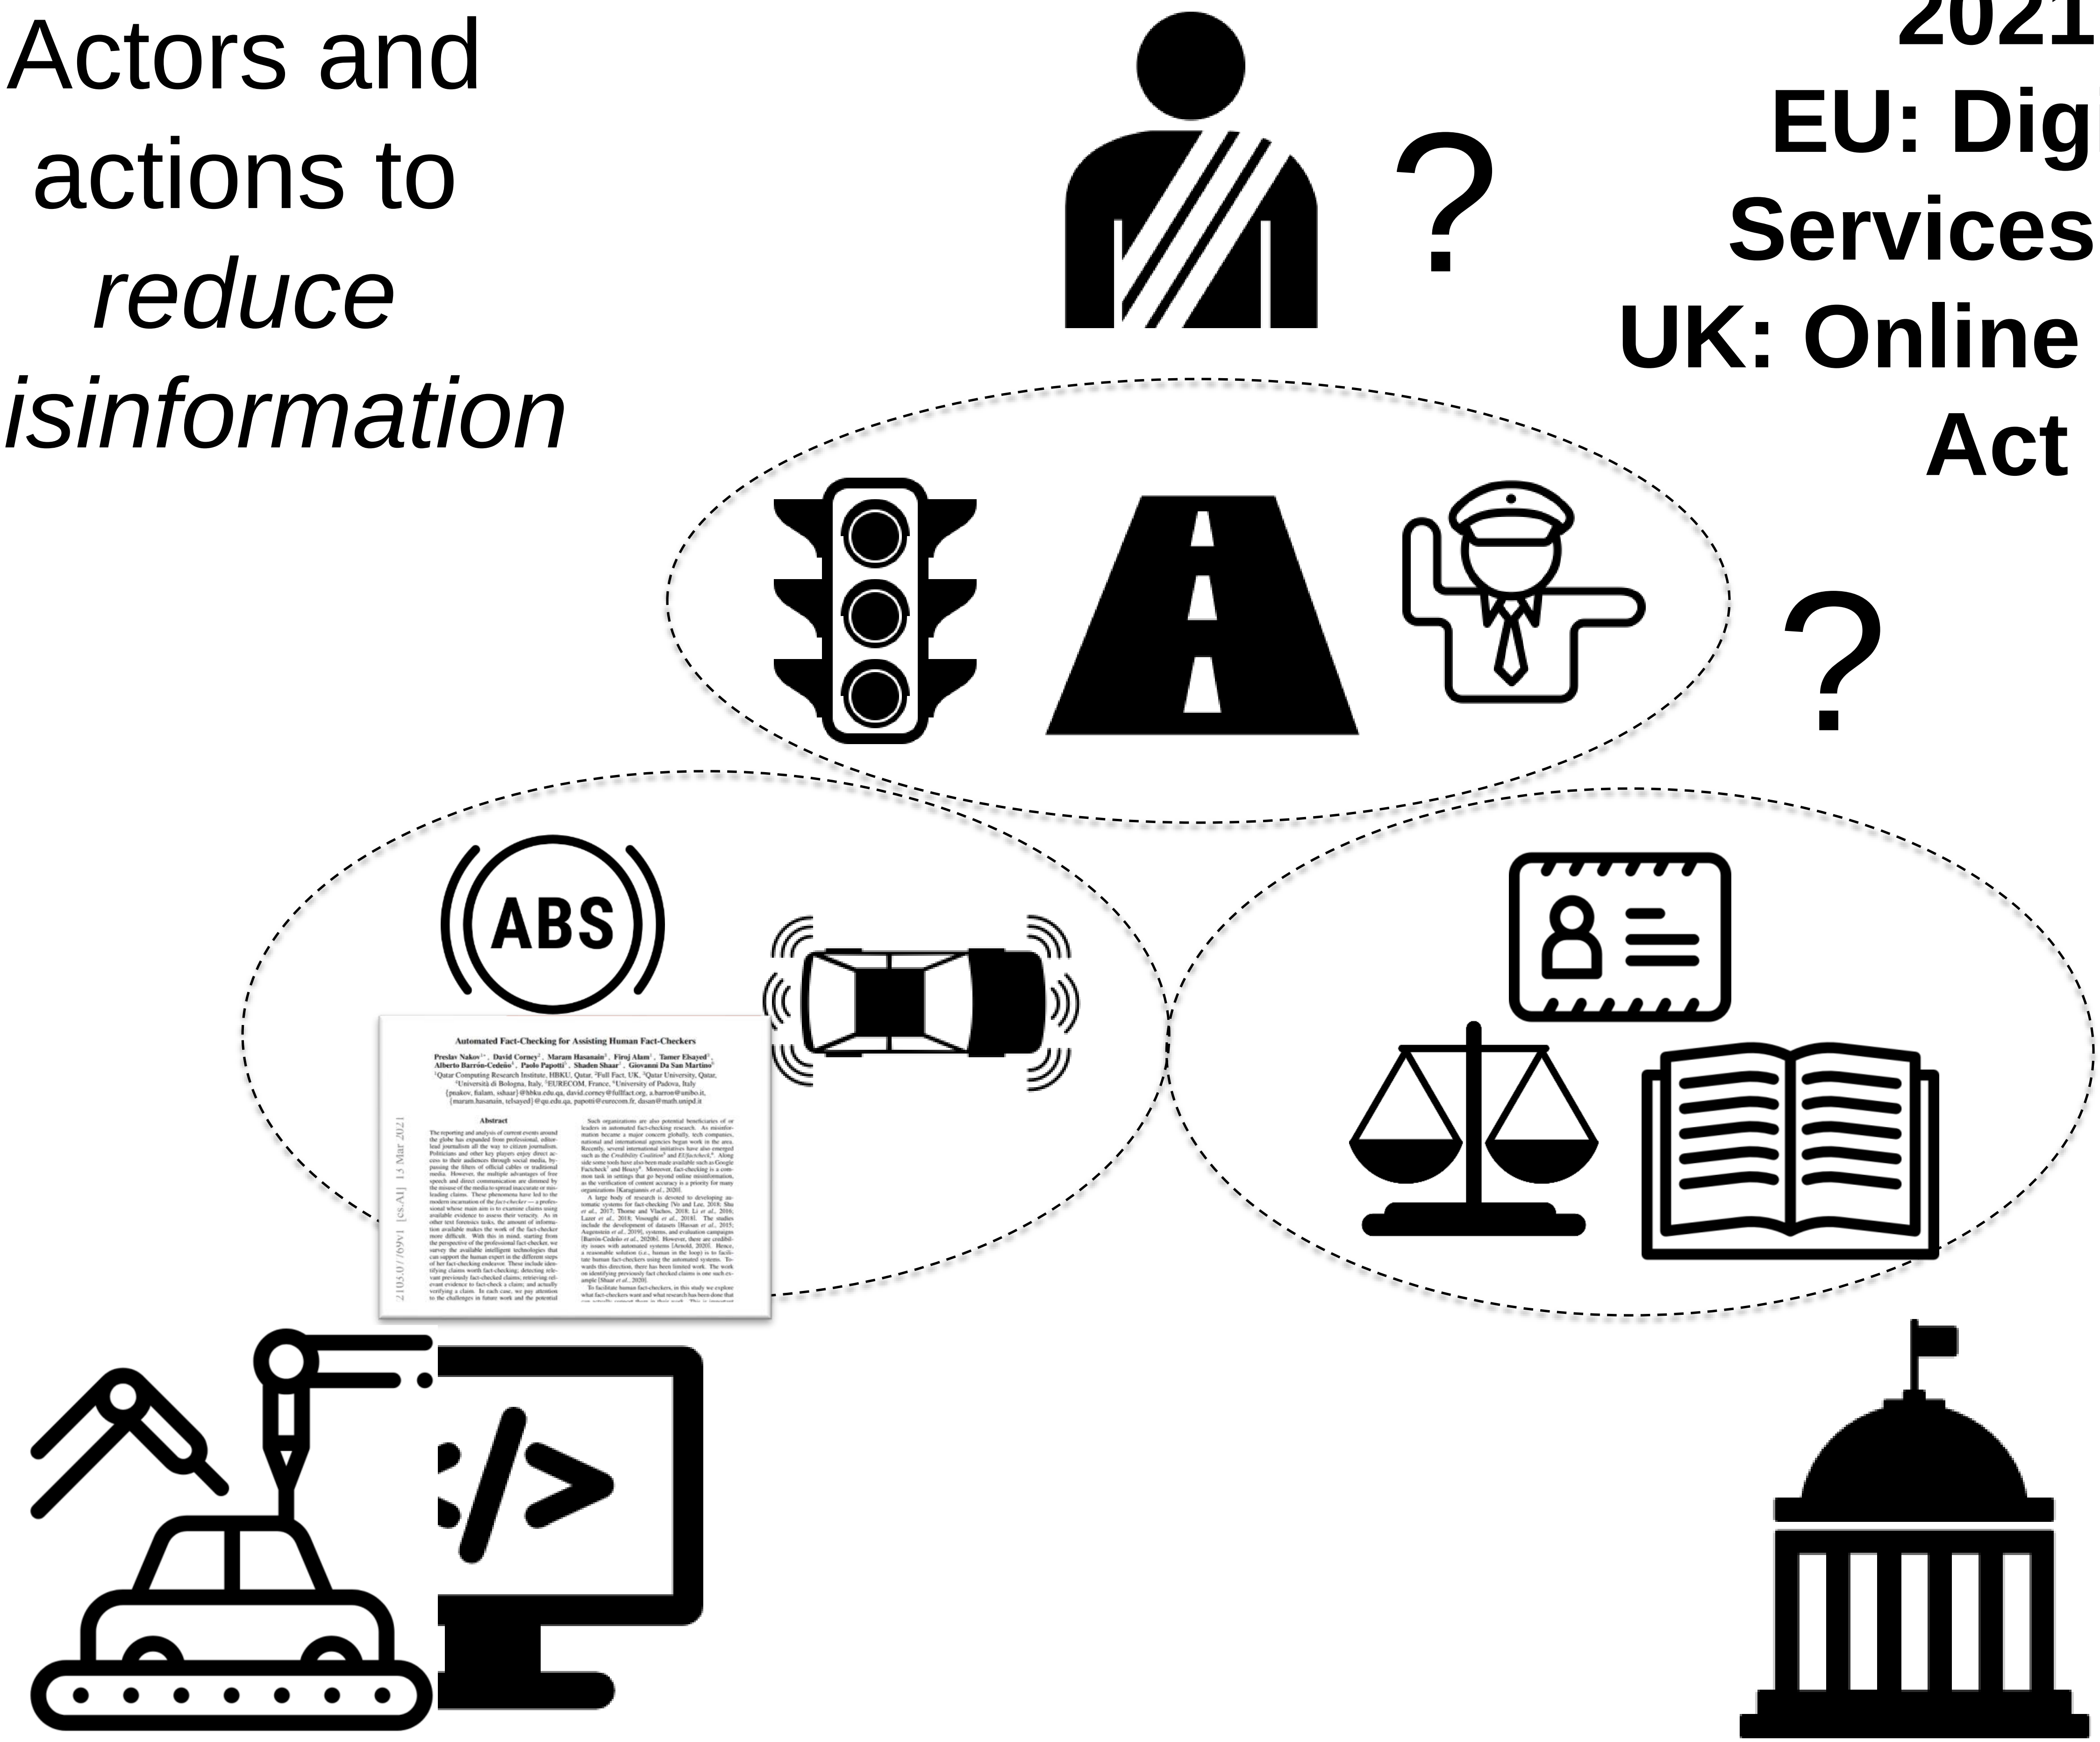

- Data driven fact checking is happening
- More research and engineering is needed

- But it's not only about tech
  - Regulation + Education

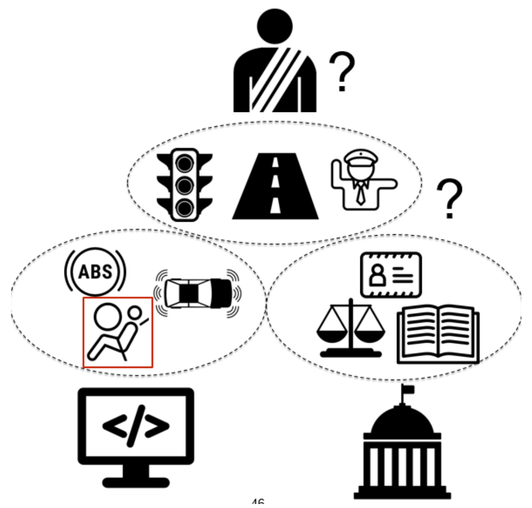

http://www.eurecom.fr/~papotti/

@paolopapotti

References 1

- Naser Ahmadi, Joohyung Lee, Paolo Papotti, and Mohammed Saeed. Explainable fact checking with probabilistic answer set programming. In TTO, 2019.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, Arpit Mittal: FEVEROUS: Fact Extraction and VERification Over Unstructured and Structured information. NeurIPS Datasets and Benchmarks 2021
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, Isabelle Augenstein: Generating Fact Checking Explanations. ACL 2020: 7352-7364
- Bordes, Antoine, et al. "Translating embeddings for modeling multi-relational data." Advances in neural information processing systems. 2013.
- T.D. Cao, I. Manolescu, and X. Tannier. Searching for truth in a database of statistics. In WebDB, pages 4:1–4:6, 2018.
- W. Chen, H. Wang, J. Chen, Y. Zhang, H. Wang, S. Li, X. Zhou, and W. Y. Wang. Tabfact: A large-scale dataset for table-based fact verification. In ICLR, 2020
- Ciampaglia, Giovanni Luca, et al. "Computational fact checking from knowledge networks." PloS one 10.6 (2015): e0128193.
- Xin Luna Dong, Divesh Srivastava: Knowledge Curation and Knowledge Fusion: Challenges, Models and Applications. SIGMOD Conference 2015: 2063-2066
- Jonathan Lajus, Luis Galárraga, Fabian M. Suchanek: Fast and Exact Rule Mining with AMIE 3. ESWC 2020: 36-52
- Joohyung Lee, Yi Wang: Weighted Rules under the Stable Model Semantics. KR 2016: 145-154
- Ludivine Duroyon, François Goasdoué, Ioana Manolescu: A Linked Data Model for Facts, Statements and Beliefs. WWW (Companion Volume) 2019: 988-993

References 2

- Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, Fabian M. Suchanek: AMIE: association rule mining under incomplete evidence in ontological knowledge bases. WWW 2013: 413-422
- Matt Gardner and Tom M Mitchell. 2015. Efficient and Expressive Knowledge Base Completion Using Subgraph Feature Extraction.. In EMNLP. 1488–1498.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. Toward automated fact-checking: Detecting check-worthy factual claims by ClaimBuster. In SIGKDD, pages 1803–1812, 2017.
- J. Herzig, P. K. Nowak, T. Muller, F. Piccinno, and J. M. Eisenschlos. TAPAS: weakly supervised table parsing via pre-training. In ACL, 2020
- Georgios Karagiannis, Mohammed Saeed, Paolo Papotti, and Immanuel Trummer.Scrutinizer: A mixed-initiative approach to large-scale, data-driven claim verification. VLDB, 13(11):2508–2521, 020.
- Neema Kotonya, Francesca Toni: Explainable Automated Fact-Checking: A Survey. COLING 2020: 5430-5443
- Marco Mori, Paolo Papotti, Luigi Bellomarini and Oliver Giudice. Neural machine Translation for Fact-Checking Temporal Claims. FEVER 2022
- Preslav Nakov, David P. A. Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, Giovanni Da San Martino: Automated Fact-Checking for Assisting Human Fact-Checkers. IJCAI 2021: 4551-4558
- Stefano Ortona, Venkata Vamsikrishna Meduri, Paolo Papotti: Robust Discovery of Positive and Negative Rules in Knowledge Bases. ICDE 2018: 1168-1179

References 3

- Yunke Qu, Kevin Roitero, Stefano Mizzaro, Damiano Spina, and Gianluca Demartini. Human-in-the-Loop Systems for Truthfulness: A Study of Human and Machine Confidence. In: Proceedings of the 3rd Conference for Truth and Trust Online (TTO 2021).
- Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. That is a known lie: Detecting previously fact-checked claims. In ACL, pages 3607–3618, 2020.
- Shi, Baoxu, and Tim Weninger. "ProjE: Embedding projection for knowledge graph completion." Thirty-First AAAI Conference on Artificial Intelligence. 2017.
- Saeed et al. Crowdsourced Fact-Checking at Twitter: How Does the Crowd Compare With Experts? Under review
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. The fact extraction and VERification (FEVER) shared task. In FEVER, pages 1–9, 2018.
- Enzo Veltri, Donatello Santoro, Gilbert Badaro, Mohammed Saeed, Paolo Papotti: Pythia: Unsupervised Generation of Ambiguous Textual Claims from Relational Data. SIGMOD Conference 2022: 2409-2412
- Y. Wu, P. K. Agarwal, C. Li, J. Yang, and C. Yu. Toward computational fact-checking. PVLDB, 7(7):589–600, 2014.
- Andrea Wenzel. To Verify or to Disengage: Coping with “Fake News” and Ambiguity. International Journal of Communication 13(2019)
- Dustin Wright and Isabelle Augenstein. Claim check-worthiness detection as positive unlabelled learning. In Findings of EMNLP, pages 476–488, 2020.