Computational Fact Checking

A Content Management Perspective

Sylvie Cazalens[1] Philippe Lamarre[1] Julien Leblay[2] Ioana Manolescu[3] Xavier Tannier[4]

PVLDB Conference - August 2018

[1]Université de Lyon, INSA-Lyon – CNRS, LIRIS, France

[2]Artificial Intelligence Research Center – AIST, Japan

[3]Inria, LIX – CNRS and École Polytechnique, France

[4]Sorbonne Université – Inserm – LIMICS, France

Disclaimer

The slides presented during the WWW tutorial contained copyrighted materials.

In the present version, copyrighted materials have been replaced by blank frames.

What is this tutorial about?

- **About** how computer science can help *a posteriori* fact checking of claims:
  - Extracting claims from some discourse,
  - Searching for the facts the claims are based on,
  - Assessing the accuracy of the claim,
  - Providing perspective to claims
- **Not only about** fake news detection!
- **Not about** image and video fact checking

Companion paper in WWW 2018, "Journalism, Misinformation and Fact-Checking" track and at
`https://hal.archives-ouvertes.fr/hal-01722666`

Is fact-checking worth it?

"Some people will not be convinced"

No, they won't.

"Facts have a liberal bias" (Paul Krugman, Nobel prize in economics)

Source: https://www.nytimes.com/2017/12/08/opinion/facts-have-a-well-known-liberal-bias.html

Scientists and humanity scholars believe in a constructed, logical discourse, and believe humans yield to reason. Businesspeople know this is not true, in general. Businesspeople have thus an advantage in winning political competitions (George Lakoff, former Berkeley professor)

Source: https://georgelakoff.com/2016/11/22/a-minority-president-why-the-polls-failed-and-what-the-majority-can-do/

Conspiracy theory adepts have no problem believing two obviously contradicting theories [Wood et al., 2012]

Is fact-checking worth it? (continued)

We still think it is:

- For legal purposes
- As long as free, high-quality press remains
- Technology can help, if we get it right

Also, a source of many cool DB research problems!

Outline

News and journalism

Definition

news[1] *noun*

newly received or noteworthy information, especially about recent events

journalism *noun*

the activity of writing for newspapers, magazines, or websites or preparing news to be broadcast

- Journalists investigate, check the facts, explain, abiding by ethical principles including accuracy, objectivity, impartiality, and accountability
- Many countries have laws protecting freedom of the press, which also define the rights and responsibilities of news organizations.

[1] All definitions according to Oxford dictionary

Freedom of the press

In France

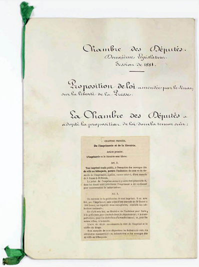
Chambre des Députés.
Deuxième Législature.
Session de 1881.

Proposition de loi amendée par le Sénat, sur la liberté de la Presse.

La Chambre des Députés a adopté la proposition de loi dont la teneur suit :

Chapitre premier.

De l'imprimerie et de la librairie.

Article premier.

L'imprimerie et la librairie sont libres.

Art. 2.

- The law dates back from 1881
- Born on the aftermath of insurrections in Paris ("Commune"), where defamation was widespread
- Already forbids publishing fake news causing "disturbance of the public sphere"

- France voted in July 2018 stricter regulations on news during elections

Regulating the news is one thing, but where to draw the line is another.

Free press is an essential ingredient of a democracy

To debate and express dissent

Romania, circa 1989:

Banner reads: “Ceaușescu re-elected at the 14th congress!”

He was in power since 1965.

Massive protests lead to approx 1000 dead.

No one convicted

To expose and explain how a society functions

Panama Papers, 2016:

Massive tax evasion offshore.

Known thanks to work by the International Consortium of Investigative Journalism (ICIJ).

Free press is an essential ingredient of a democracy

- To debate and express dissent.
- To confirm or refute public statements.
- To expose and explain how society functions.
- To keep the authorities accountable.

Daphne Caruana Galizia (1967-2017)

Long standing issues

Honest mistakes

- from incomplete and inaccurate sources
- from ambiguities of languages

Source: https://www.nytimes.com/2016/09/13/well/eat/how-the-sugar-industry-shifted-blame-to-fat.html

Long standing issues (continued)

Bias

On the part of journalists and reader

- from cultural, financial or political pressure
- as well as other social or psychological factors

Deception (including fake news)

- as old as journalism
- can take many forms
  (rumors, hoax, propaganda, satire, etc.)

Fact checking

Definition

fact-check *verb [with object]*

investigate (an issue) in order to verify the facts

Term in use since 1930 approx.

Source: Google N-gram viewer

Fact checking (Ye Good Ol' Days)

"The day I became a fact-checker at The New Yorker, I received one set of red pencils and one set of No. 2 pencils. [. . .] The red pencils were for underlining passages on page proofs of articles that might contain checkable facts [. . .] confirmed with the help of reference books from the magazine's library, including Merriam-Webster's Geographical Dictionary, the New Grove Dictionary of Music and Musicians and Burke's Peerage and Gentry."

Source:

nytimes.com/2010/08/22/magazine/22FOB-medium-t.html

Fact checking in the Internet era

- As the Internet took off, in the mid-90s, it gradually incorporated all other forms of media...
- ... allowing anyone publishing anything, while reaching a global audience.
- Gradually, journalists had to become more tech-savvy.

Fact checking has moved from **before** to **after** publication!

- A seminal article by [Cohen et al., 2011] gave birth to **computational journalism** as a discipline
- Since then, DB, IE, NLP, ML, KR communities have started work in the area

Context and problems

Q: "Is it true that in Moscow, Mercedes cars are being given to citizens?"
A: "Yes, but it is not Moscow but Leningrad, not Mercedes but Ladas, and not given to but stolen from."

Yeravan jokes, famous in the Eastern block during communism.

Outline

Fact-checking in the Internet era: what's new ?

The Web as the primary media

- Traditional news sources increasingly disseminate through the Web
- New outlets, e.g. so-called pure players, run 100% of their operations on the Web
- Social networks became major media outlets and conduit

Fact-checking in the Internet era: what's new ?

"Democratization" of authorship

- Non-media organizations (companies, government Web sites) and individuals gained access to large scale publishing means
- No editorial process or ethics required
- Line blurred between news producers and consumers.
- Sudden abundance of data (with varying quality/credibility)

Social networks have become a primary source for news

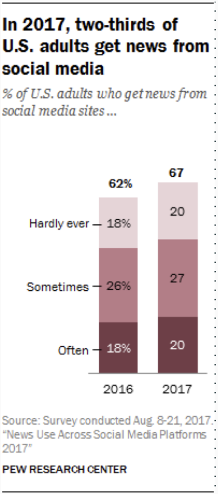

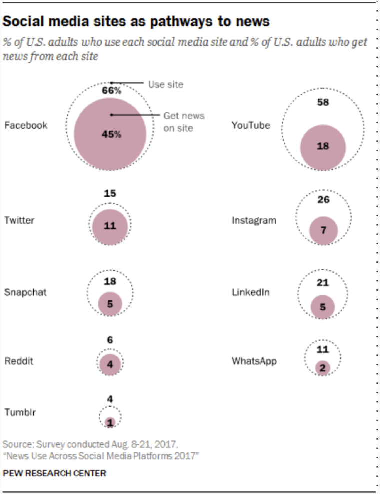

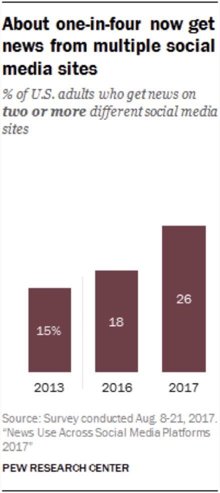

In some emerging countries, Internet "is" Facebook[2]

[2] qz.com/243086/facebook-is-creating-a-parallel-internet-in-emerging-markets/

Social media mishaps

Social networks are increasingly weaponized to spread dubious information

Example (Recent events)

- 2-3% of Facebook accounts are fake, 5% on Twitter[a]
- Twitter conducted a sweeping “bot purge” in February 2018[b]
- Russian meddling in the US presidential election
- Cambridge Analytica scandal vs. US elections and Brexit
- Insane man’s terror act in Germany wrongly connected to immigrants[c]

[a] nytimes.com/2017/11/03/technology/facebook-fake-accounts.html
[b] thedailybeast.com/inside-twitters-bot-purge
[c] lemonde.fr/international/article/2018/04/09/l-allemagne-sous-le-choc-apres-l-attaque-de-munster-et-l-attentat-dejoue-a-berlin_5282856_3210.html

A brief history of fact checking initiatives

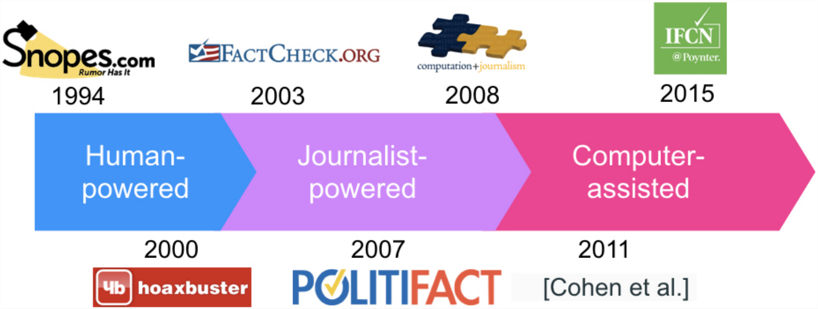

Fact checking sites today

Source: reporterslab.org/fact-checking (Duke Reporters Lab)

Types of information

World events are intertwined with longer-term social issues.

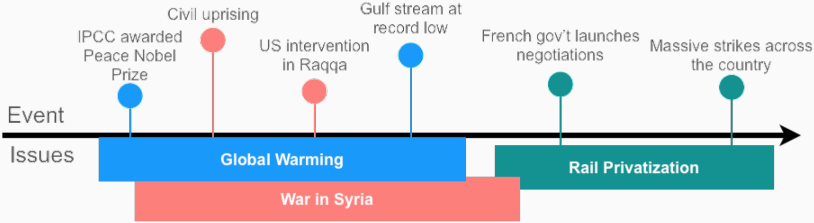

- Information is altered as it propagates across a social network (through bias, accumulated errors, and outright lies)
- Journalists must provide a short and high quality channel between the events on the public
- Stakeholders can be motivated, rely on rhetoric, persuasion
- Journalists must balance emotions, credibility and reasoning.

Computer science and journalism: how can we help?

1. **Data journalism**: journalistic work significantly or mainly based on (digital) data
2. **(Semi)-automated fact checking**: fact checking work where some tasks are delegated to software
   - **Our focus today**
   - Fact checking tasks will be detailed shortly
3. **Fake news detection**: software which estimates the level of falsehood of a piece of news
   - True, false, in-between...
   - May not use reference sources.

Not fact checking

Text classifier (true, fake) trained with major news agency content, resp. known fake content (often virulent style)

Source: https://towardsdatascience.com/i-trained-fake-news-detection-ai-with-95-accuracy-and-almost-went-crazy-d10589aa57c

Fact checking ingredients

To successfully check a claim, one needs to:

1. Lift the ambiguity
   - Vague statements lead to too many distinct interpretations, which one to check?
   - Clarify the context is which the claim is analysed (space, time...).
2. Ensure it is backed by sufficient references to sources.
   - Reliable reference sources give the background against which to check.
3. Validate the claim as consistent with the sources.
   - Some claims are crafted to mislead, i.e., look valid wrt a context or source that is irrelevant or flawed.

The need for transparency

The International Fact-Checking Network (IFCN) is sponsored by the Poynter Institute to "promote excellence in fact checking". Members commit to:

1. Non-partisanship and fairness.
2. Transparency of sources.
3. Transparency of function and organization.
4. Transparency of methodology.
5. Open and honest corrections.

Source: poynter.org/international-fact-checking-network-fact-checkers-code-principles

The limits of fact checking

- Confirmation bias: people are more likely to believe what fits their prior views.
  - Man-made part of the echo chamber.
  - Automated recommendation systems trap users in filter bubbles.

Yet:

- Filter bubbles and echo chambers are still being studied [Garrett, 2009, Garrett, 2016].
- Showing readers links to "related stories" reduces misperceptions more effectively [Bode and Vraga, 2015].

The limits of fact checking: Timing matters!

Emotionally engaging information, such as rumors and propaganda, spread faster than corrections on social networks [Shin et al., 2017].

- False news spread faster than true ones; most of the audience is reached in the first 24 hours [Vosoughi et al., 2018].
- If verification comes too late, false information has time to "stick" with audience[3].
- Backfire effect: defiance towards fact checkers may reinforce reader's perception if confronted directly [Nyhan and Reifler, 2010], near instant-correction making things worse [Garrett and Weeks, 2013].

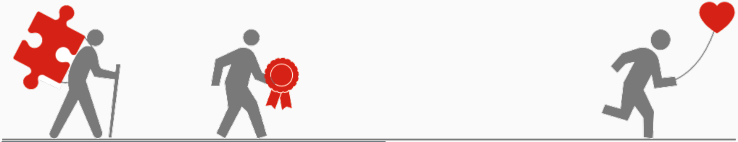

[3] jonathanstray.com/networked-propaganda-and-counter-propaganda

Outline

Misinformation and disinformation

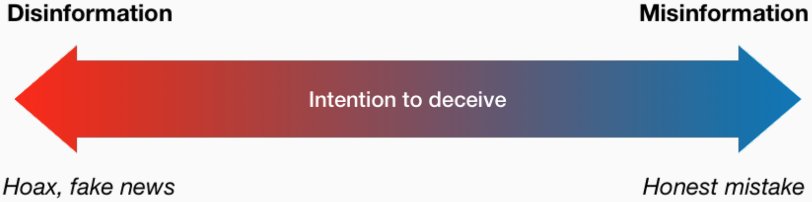

Rumors, myths, conspiracy theories

Widespread, often long-standing, misconceptions, on which one might base their judgment.

Hoax

Non-elaborate, unsubstantiated claim. Aimed at spreading virally.

Source: snopes.com/fact-check/david-hogg-on-campus-rumor-hoax/

Clickbait

Catchy title, poor content. Aimed mainly at attracting audience for financial gains.

Media hype

Catchy story, with some core element of truth, but vastly exaggerated.

Frequent in health, technology and science news.

Wrong context assignment

Genuine content (image, video, audio, or quote) planted in unrelated context to steer opinion.

Source:

https://abonnes.lemonde.fr/big-browser/article/2018/08/23/a-l-agence-france-presse-plongee-dans-le-service-fact-checking_5345538_4832693.html

Content doctoring

Realistic depiction of events that did not occur.

Source: wapo.com/news/the-intersect/wp/2018/03/25/a-fake-photo-of-emma-gonzalez-went-viral-on-the-far-right-where-parkland-teens-are-villains/

Incorrect factual claims

Claims with obvious interpretation and for which there exists reasonably relevant and accurate data.

> *"Our prisons are filled-up with foreigners."*
>
> BBC Question Time audience member, Oct. 20, 2016

Foreign citizens make up 9% of the general population and 12% of the prison population in England and Whales. [...] The number and proportion of foreign prisoners is falling: there were over 11,000 foreign prisoners in 2010.

Source: fullfact.org/immigration/foreigners-prison

Ambiguous or oversimplifying claims

The claim is open to multiple interpretations, some of which may be true, but not necessarily the most relevant one.
Typically requires in-depth analysis.

Source: politifact.com/punditfact/statements/2018/mar/28/blog-posting/was-ohio-student-suspended-staying-class-during-na/

Flavours of fake news [Rubin et al., 2015]

Even when there is intention to deceive, the purpose of the deception may vary a lot:

- coordinated and well-targeted information forgery
- simple lies that catch on to a large audience
- humor, satire, sarcasm

From characterization to detection [Shu et al., 2017]

Fake news is a news article that is intentionally and verifiably false.

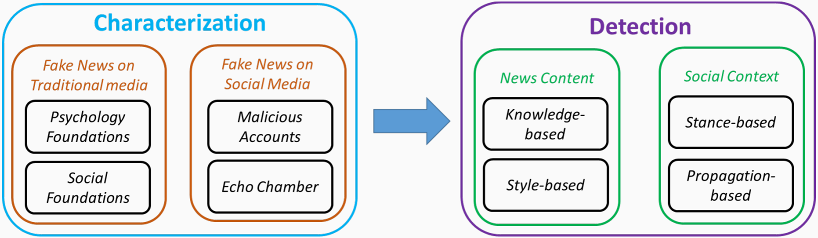

Outline

Correct but imprecise

- The French Railway Company (SNCF) went strike in spring 2018 protesting reform.
- Unions, the company, the government and other interest groups went in show of force.
- SNCF published a press release aledging the protest mobilization was rapidly falling.

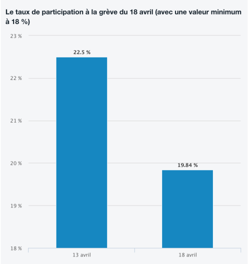

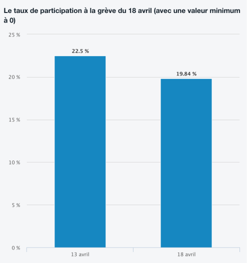

Source: lemonde.fr/les-decodeurs/article/2018/04/18/le-graphique-trompeur-de-la-direction-de-la-sncf-sur-le-taux-de-participation-a-la-greve_5287273_4355770.html

In-depth analysis

Detailed "forensic" study of past events for which reference sources were hard to obtain, witnesses speak late etc.

Source: https://www.mediapart.fr/journal/france/040418/nicolas-sarkozy-bien-servi-les-interets-de-kadhafi-voici-les-preuves

Promise verification

Validating past claims made about the future

Source:

politifact.com/truth-o-meter/promises/obameter/subjects/politifacts-top-promises/

Reversal tracking

Checking a personality's position or stance on a specific issue over time.

Source: politifact.com/texas/statements/2018/jan/10/beto-orourke/beto-orourke-flip-flops-requiring-public-service-y/

State of the art

Outline

FactCheck.org

- In-depth analysis of a claim
- First politics, now science, health etc.
- Particularly active during US elections

> “*We are not going to let our campaign be dictated by fact-checkers.*”
>
> Neil Newhouse, pollster for Republican nominee Mitt Romney

Politifact

Truth-o-meter

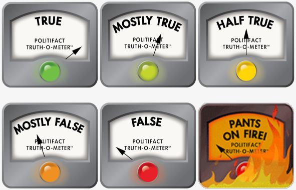

Flip-o-meter

- In-depth political claim analysis
- Simple classifaction for checked claims
- Position reversals
- DB published through an API

Les Décodeurs team of “Le Monde”

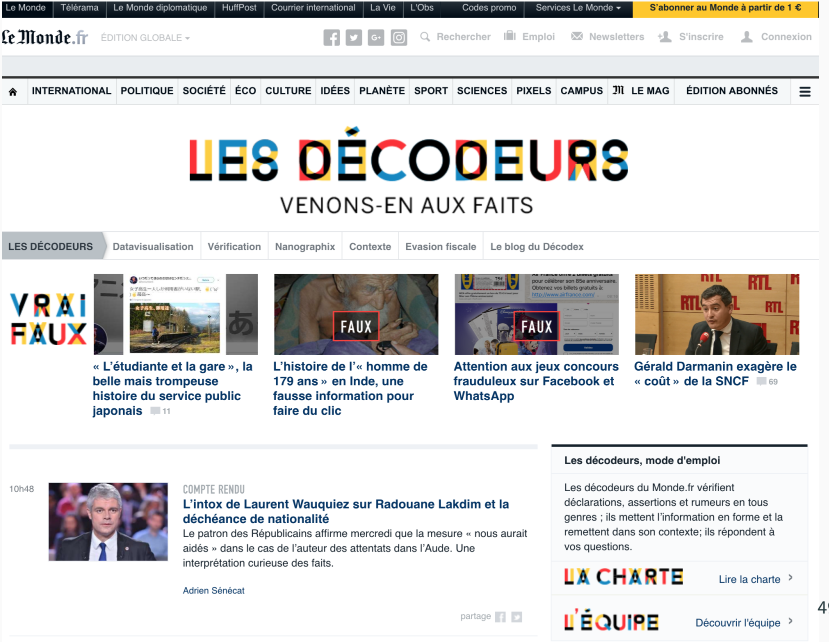
Le Monde | Télérama | Le Monde diplomatique | HuffPost | Courrier international | La Vie | L'Obs | Codes promo | Services Le Monde | S'abonner au Monde à partir de 1 €

Le Monde.fr ÉDITION GLOBALE

Rechercher | Emploi | Newsletters | S'inscrire | Connexion

INTERNATIONAL | POLITIQUE | SOCIÉTÉ | ÉCO | CULTURE | IDÉES | PLANÈTE | SPORT | SCIENCES | PIXELS | CAMPUS | LE MAG | ÉDITION ABONNÉS

LES DÉCODEURS

VENONS-EN AUX FAITS

LES DÉCODEURS | Datavisualisation | Vérification | Nanographix | Contexte | Evasion fiscale | Le blog du Décodex

VRAI FAUX

« L'étudiante et la gare », la belle mais trompeuse histoire du service public japonais 11

FAUX

L'histoire de l'« homme de 179 ans » en Inde, une fausse information pour faire du clic

FAUX

Attention aux jeux concours frauduleux sur Facebook et WhatsApp

Gérald Darmanin exagère le « coût » de la SNCF 69

10h48

COMPTE RENDU

L'intox de Laurent Wauquiez sur Radouane Lakdim et la déchéance de nationalité

Le patron des Républicains affirme mercredi que la mesure « nous aurait aidés » dans le cas de l'auteur des attentats dans l'Aude. Une interprétation curieuse des faits.

Adrien Sénécat

partage

Les décodeurs, mode d'emploi

Les décodeurs du Monde.fr vérifient déclarations, assertions et rumeurs en tous genres ; ils mettent l'information en forme et la remettent dans son contexte; ils répondent à vos questions.

LA CHARTE — Lire la charte ›

L'ÉQUIPE — Découvrir l'équipe ›

Fact-checking blogs of main media

- France: liberation.fr, afp.fr
- In the US: TruthTeller from the Washington Post closed in 2014 circa.
- In Italy: major media not interested, prominent fact-checkers moved to US media school

Crosscheck from First Draft News

- Supported by Google News Initiave
- Relies on volunteers
- Trains the public to critical thinking and news analysis

Outline

Fact checking pipeline

Definition (Fact checking [Babakar and Moy, 2016])

Defined as a four-stage process where

(i) media sources are monitored,

(ii) claims are spotted,

(iii) claims are checked,

(iv) fact checking analysis results are created and published.

Fact checking from a content management perspective

Overview of **tasks**, **inputs** and **outputs**

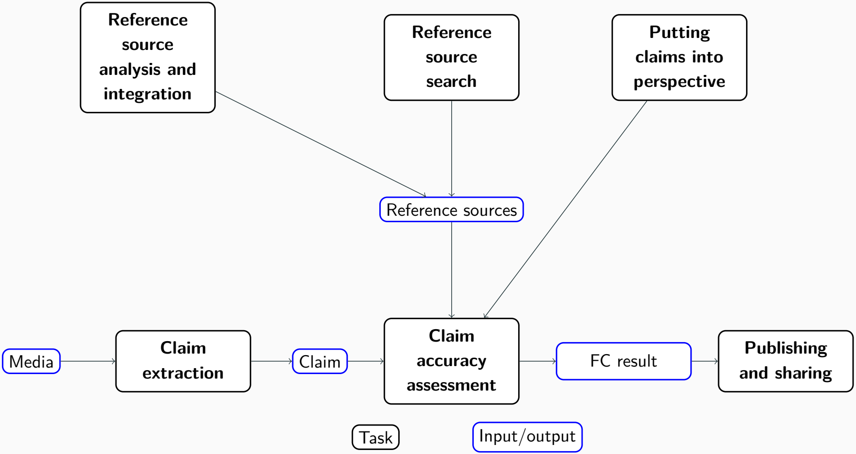

Fact checking from a content management perspective: partial reference list

More references available in companion paper [Cazalens et al., 2018]

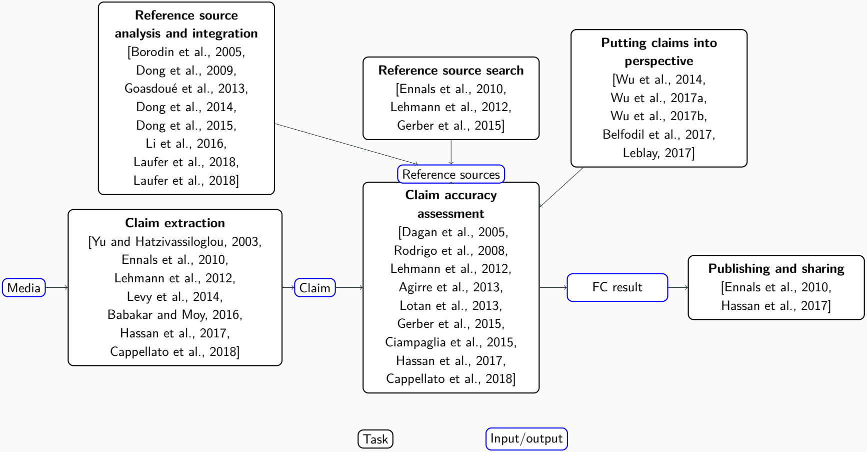

Which reference data source to use?

- **Fixed** (known in advance):
  - ClaimBuster [Hassan et al., 2017]
  - DisputeFinder (PolitiFact API) [Ennals et al., 2010]
  - FullFact (internal DB of manually checked claims) [Babakar and Moy, 2016]
  - TruthTeller (claims manually checked by Factcheck.org);
  - The Décodex plug-in developed by Le Monde also leverages their past fact checking analyses
- **Web search**:
  - DeFacto [Lehmann et al., 2012, Gerber et al., 2015], ClaimBuster [Hassan et al., 2017], CLEF CheckThat 2018 winners [Cappellato et al., 2018]

Professional journalism is very picky on source quality.

Which reference data source to use? (continued)

- **General knowledge bases** such as Wikipedia [Ciampaglia et al., 2015]
- **Heterogeneous open data**, e.g., FactMinder: enrichment of online articles with open data [Goasdoué et al., 2013].

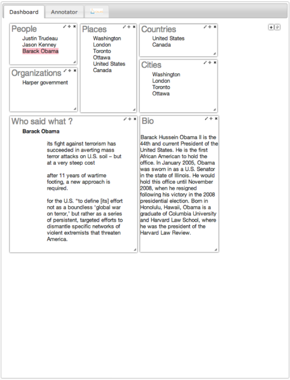

- **Proprietary data**: usually high-quality; data vendors

Building reference data sources: truth discovery

Partially overlapping Web sources require arbitrating between their information.

Example

NY restaurant information [Dong et al., 2009]

Source	Coverage	Exactness	Freshness	#Closed-rest
MenuPages	.66	.98	.86	29
TasteSpace	.44	.97	.3	106
NYMagazine	.43	.98	.54	59
NYTimes	.43	.98	.38	72
ActiveDiner	.41	.95	.86	70
TimeOut	.38	.99	.68	33
SavoryCities	.27	.99	.41	33
VillageVoice	.22	.94	.4	37
FoodBuzz	.18	.92	.3	59
NewYork	.13	.92	.45	28
OpenTable	.12	.92	.45	9
DiningGuide	.1	.9	.09	48
GoogleMaps	-	-	-	212

Extracted from 12 sources + manually checked

Probabilistic approach for determining the true value, based on coverage, exactness and freshness, and on who copied whom.

Building reference data sources: truth discovery

Truth discovery survey [Li et al., 2016]:

Input: a set of values for an object, each from a different source

Output: most likely value and trustworthiness of each source

Principle:

- A source whose value for an object was deemed correct, will be considered more trustworthy
- ... and values coming from a trustworthy source will be considered more likely to be correct

Methods: iterative; optimization-based (error minimization); probabilistic graphical models

Constructing reference data sources: data integration

Valuable information is sometimes found across several data sources

Data integration approaches:

- **Warehouse**: extract and consolidate all data sources into one

  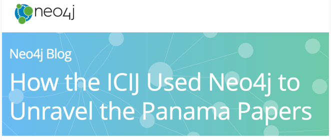

  Text (contracts) and relational (screen company coordinates) data sources fused into one (Neo4J) graph database
  Easy to use; needs to be redone for every new dataset
  Generic system: ConnectionLens [Chanial et al., 2018]
  (⇒ **Demo Group C, Thu 11:30**)
- **Mediator**: structured data sources remain unchanged and are queried together under a unified schema

Constructing reference data sources: data integration (cont'd)

Valuable information is sometimes found across several data sources

Data integration approaches (cont'd):

- **Data space**: structured and unstructured data sources queried through keywords [Franklin et al., 2005, Chanial et al., 2018]
- **Data lake**: large number of structured and unstructured data sources w/o unified schema; subsets of these are exploited together in mediator style[4] e.g. [Bonaque et al., 2016]
- **Dataflow**: data journalism analytical pipelines[5]

[4] https://www.ibm.com/analytics/data-management/data-lake, https://blogs.oracle.com/bigdata/the-new-data-lake-you-need-more-than-hdfs
[5] http://jonathanstray.com/introducing-the-cj-workbench

High-quality reference data, e.g., published by statistic institutes, may be hard to query

1. **Extract** data into RDF Linked Open Data (preserving table and header structure) [Cao et al., 2017]
2. **Search** for exact or closest approximate answer to a keyword query [Cao et al., 2018]

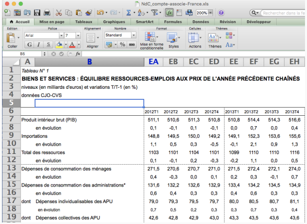

Tableau N° 1

BIENS ET SERVICES : ÉQUILIBRE RESSOURCES-EMPLOIS AUX PRIX DE L'ANNÉE PRÉCÉDENTE CHAÎNÉS

niveaux (en milliards d'euros) et variations T/T-1 (en %)

données CJO-CVS

	2012T1	2012T2	2012T3	2012T4	2013T1	2013T2	2013T3	2013T4
Produit intérieur brut (PIB)	511,1	510,6	511,3	510,8	510,8	514,4	514,3	516,6
en évolution	0,1	-0,1	0,1	-0,1	0,0	0,7	0,0	0,4
Importations	148,8	149,5	150,0	149,2	149,1	152,3	153,6	155,6
en évolution	1,1	0,5	0,3	-0,5	-0,1	2,1	0,9	1,3
Total des ressources	1103	1101	1104	1101	1099	1110	1110	1116
en évolution	0,1	-0,1	0,2	-0,3	-0,2	1,1	0,0	0,5
Dépenses de consommation des ménages	271,5	270,6	270,7	271,0	271,5	272,4	272,1	274,0
en évolution	0,4	-0,3	0,0	0,1	0,2	0,3	-0,1	0,7
Dépenses de consommation des administrations*	131,6	132,2	132,6	132,9	133,4	134,2	134,5	134,9
en évolution	0,6	0,4	0,3	0,3	0,4	0,6	0,3	0,3
dont Dépenses individualisables des APU	79,0	79,3	79,5	79,7	80,0	80,5	80,7	81,1
en évolution	0,7	0,5	0,2	0,3	0,3	0,7	0,3	0,4
dont Dépenses collectives des APU	42,6	42,8	42,9	43,0	43,3	43,5	43,6	43,6

Source: insee.fr/fr/statistiques/3292347?sommaire=3292415

Searching for truth in statistic tables

Query: "youth unemployment France August 2017"

	Seasonally adjusted youth (under 25s) unemployment				
	Number of persons (in thousands)				
	Oct-2016	Jul-2017	Aug-2017	Sep-2017	Oct-2017
Belgium	:	77	77	77	:
Bulgaria	27	22	21	19	19
Czech Republic	34	27	25	23	23
Denmark	62	54	53	49	47
Germany	293	283	283	283	283
France	663	629	625	623	625

Answer: **625**, link to the spreadsheet as result proof (provenance, justification)

- Extraction needs to cope with nested headers
- Off-line source indexing
- Search for (i) relevant datasets and (ii) most relevant cells in each dataset

Claim recognition: claim extraction

- **Topic-driven extraction** from media articles [Levy et al., 2014].
  - Task: Given a topic (*context*), find related claims, e.g.:

Topic	**Selling violent video games to minors should be banned**
Related claim	*Violent video games can increase children's aggression*

  - Approach: fully supervised learning.
- **Locate disputed claims** covered by the reference database [Ennals et al., 2010].
  - Task: Given a text, extract claims disputed by a trusted source, e.g.:

    Many vaccines contain mercury, aluminium and other toxins that should have parents asking questions before immunizing their children.
  - Approach: keyword retrieval against a claim database.

Claim recognition: claim extraction (continued)

- **Entity disambiguation applied on claims**, using a reference knowledge base: DeFacto [Gerber et al., 2015].
  - Task: Given a text, extract 10 types of predefined relations between named entities.[6]
  - Example:
    Input: Albert Einstein was awarded the Nobel Prize in Physics.
    Output:

    dbr:Albert_Einstein dbo:award dbr:Nobel_Prize
  - Approach: rule-based
- Research into extraction from text feeds, audio, video:
  FullFact [Babakar and Moy, 2016]
  (technical details not available at this time).

[6] Relations are: award, birth, death, foundationPlace, leader, NBAteam, publicationDate, spouse, starringActor, subsidiary

Claim recognition: classifying check-worthiness

1. Verifiability: Verifiable *vs.* Unverifiable
   [Park and Cardie, 2014, Guggilla et al., 2016, Gencheva et al., 2017]
2. Factuality and worthiness:
   Non-factual (e.g., opinions or subjective content) *vs.*
   Factual but not interesting (consensual, general) *vs.*
   Factual and interesting (that is, check-worthy).
   [Hassan et al., 2015, Hassan et al., 2017]
3. Opinion: Facts *vs.* opinions [Yu and Hatzivassiloglou, 2003]
4. Dialogic and argumentative markers:
   - Degrees of agreement with a previous post
   - Cordiality, audience-direction, combativeness, assertiveness, emotionality of argumentation, sarcasm

   [Walker et al., 2012]

All these approaches are based on fully-supervised systems with expert- or crowd-sourced data.

Stance detection

Is a text in favor of a given target, against it, neutral or unrelated?

- Target: legalization of abortion
- Negative stance: *"A foetus has rights too! Make your voice heard"*.

- Target: Donald Trump
- Positive stance: *"@realDonaldTrump is the only honest voice of the @GOP"*.

Sources can be general claims, debates in online forums, student essays, but mostly news or political speeches, debates, tweets.

Approaches are all based on supervised learning.

[Levy et al., 2014, Bar-Haim et al., 2017, Somasundaran and Wiebe, 2009, Murakami and Raymond, 2010, Hasan and Ng, 2013, Faulkner, 2014, Thomas et al., 2006, Rajadesingan and Liu, 2014, Mohammad et al., 2016, Ferreira and Vlachos, 2016], FakeNewsChallenge (2017).

Claim accuracy assessment

- Find evidence potentially proving the claim as Web page text snippets, sufficiently close to the claim [Lehmann et al., 2012, Gerber et al., 2015, Barrón-Cedeño et al., 2018].
- Try to match claim against trusted repository of previously checked claims (e.g. PolitiFact etc.); if this fails, revert to Web search and question answering systems such as Wolfram Alpha [Hassan et al., 2015, Hassan et al., 2017] .

Use as evidence a path found in reference data sources [Ciampaglia et al., 2015, Chanial et al., 2018]; use node degree to asses truth/relevance of a candidate path

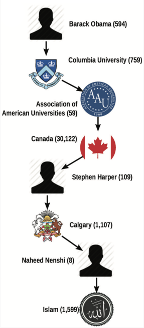

Claim accuracy assessment (continued)

- The Fast and Furious FactCheck Challenge[7] proposed to classify news articles (not claims) among: TRUE, FALSE, SOMEWHAT TRUE and SOMEWHAT FALSE w/ human and/or automated tools;
- Les Décodeurs[8] (Le Monde) developed:
  - A database of manually checked claims w/ analysis and rumor propagators.
  - A web navigator plugin w/ a trust score from the aggregated outputs of previous fact checks, where available.

[7] https://herox.com/factcheck/
[8] http://www.lemonde.fr/les-decodeurs/

Claim accuracy assessment: related tasks

These well-known NLP tasks have never really been applied to fact-checking problems as such:

- Textual entailment compares two texts and decides whether one implies the other [Dagan et al., 2005].
- The SemEval's Semantic Textual Similarity task offers a graded and typed definition of semantic similarity [Agirre et al., 2013].
- Rumor detection classifies a set of posts/tweets as rumor or not rumor, or studies the birth and propagation of rumors [Ma et al., 2016, Zubiaga et al., 2016].

Fake news detection

- Flourishing field
- Growing number of challenges, hackathons and data sets available
  - Around 160 news-related datasets and 70 public kernels on **Kaggle**
  - **BuzzFeedNews**[9]: Sample of news published on Facebook prior to the 2016 U.S. elections
  - **LIAR**[10]: A Politifact archive
  - **BS Detector**[11]: data collected through the BS detector browser extension.
  - **CREDBANK**[12]: A Large-scale Social Media Corpus With Associated Credibility Annotations

[9] https://github.com/BuzzFeedNews/2016-10-facebook-fact-check/tree/master/data
[10] https://www.cs.ucsb.edu/ william/data/liar dataset.zip
[11] https://github.com/bs-detector/bs-detector
[12] https://github.com/compsocial/CREDBANK-data

CLEF 2018 Fact Checking Lab

- "Automatic Identification and Verification of Claims in Political Debates"
- Task 1: Check-worthiness. Predict which claim in a political debate should be prioritized for fact-checking.
- Task 2: Factuality. Checking the factuality of the identified worth-checking claims. 5 participants, winners' mean absolute error (MAE) of 0.7
- `http://alt.qcri.org/clef2018-factcheck/index.php?id=overview`

Putting claims into perspective

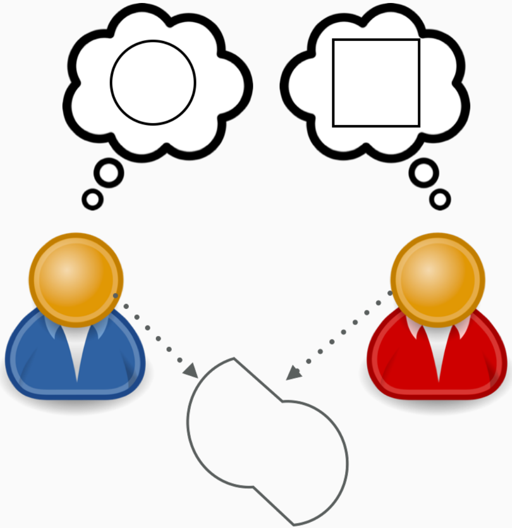

Putting claims into perspective

- Search for interesting additional elements.
  - Query perturbation [Wu et al., 2014, Wu et al., 2017a].
  - Context dependent reasoning [Leblay, 2017].
  - Exceptional Model Mining (Data mining) [Belfodil et al., 2017].
- Build and visualize a general picture of a complex issue.

The query perturbation approach [Wu et al., 2014, Wu et al., 201

Giuliani's claim: "Adoptions went up 65 to 70 percent when [he] was mayor [of New York City]."

```
SELECT after.total / before.total
FROM (SELECT SUM(number) AS total FROM adopt
    WHERE year BETWEEN t-w-d+1 AND t-d) AS before,
  (SELECT SUM(number) AS total FROM adopt
    WHERE year BETWEEN AND t-w+1 AND t) AS after;
```

Query Response Surface (QRS)

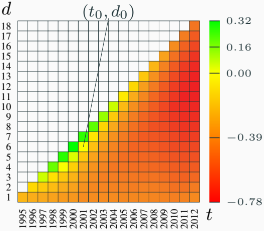

(a) Relative strength of results

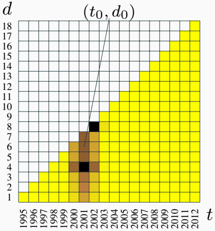

(b) Sensibility of parameter settings

The query perturbation approach (continued)

Relative strength and relative sensitivity are used to

- Find counter-argument (that weakens the original claim), and reverse-engineer vague claims
- Robustness: All perturbations result in stronger or equally strong claims
- Other notions such as fairness, and uniquess

Also introduce ways to check window aggregate comparison claims, and time series similarities claims.

Diversity and representativity [Wu et al., 2017b]

Follow-up work: when many counter arguments exist, select a subset maximizing utility, diversity and representativeness.

Problem: Find a **Diverse Set of *k* High-Value Representatives** from numerical data, for counter-argument generation and computational lead finding [Wu et al., 2017b].

Three interesting areas (plus one noisy spike) are hidden in the data. The first three methods fail to find them and/or to ignore the spike.

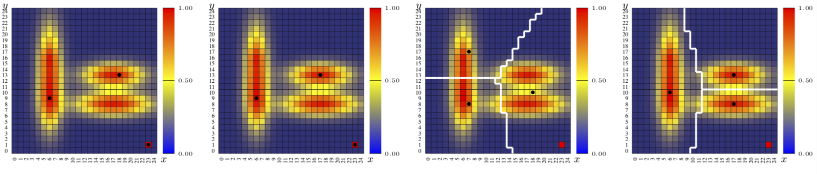

(a) MMR (b) Graph-based diversification (c) Weighted k-means (d) k-DHR

Optimization method to automatically select k-DHR [Wu et al., 2017b].

Putting claims into perspective (continued) - Context dependent reasoning

Context-dependent reasoning can be used to a veracity score to all possible contexts of a claim [Leblay, 2017]

Example

"John Doe is a Eurosceptic."

- Depends on what we mean by "Eurosceptic"
- Not everybody agrees!

Key idea: annotate the data and axioms with contextual details.

Axioms

bachelorFrom(X,Y) $\rightarrow$ CollegeGrad(X)	$:\{([-\infty,+\infty],\{bbc,cnn\})\}$	Hard
party(X,Tories) $\rightarrow$ Eurosceptic(X)	$:\{([2007,2013],\{afp, bbc\})\}$	
party(X,Y) $\wedge$ opposes(Y, Z) $\rightarrow$ opposes(X, Z)	$:\{([-\infty,+\infty],\{afp, bbc\})\}$	
supports(X,Y) $\wedge$ opposes(X,Y) $\rightarrow \perp$	$:\{([-\infty,+\infty],\{afp, bbc, cnn\})\}$	
CollegeGrad(X) $\rightarrow$ opposes(X, Brexit)	: 6.0	Soft
party(X,Tories) $\rightarrow$ support(X, Brexit)	: 3.0	

Putting claims into perspective (continued)

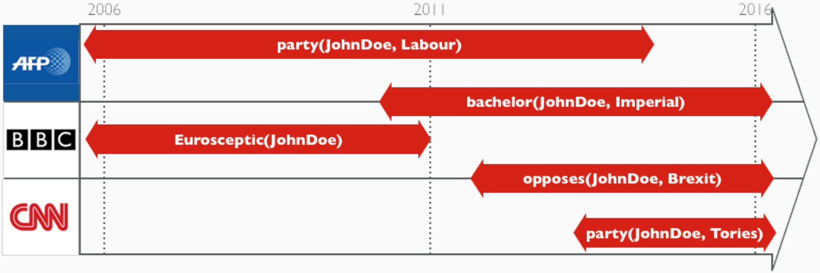

Eurosceptic(JohnDoe)?

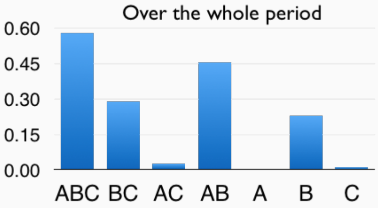

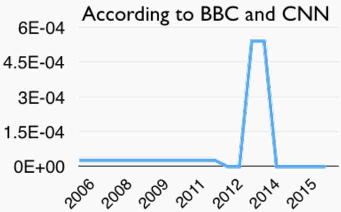

Putting claims into perspective (continued) - data mining

- Data mining techniques used to highlight true, false or misleading claims by providing knowledge from data sets.
- In case of behavioral data sets (e.g. voting, rating, consuming...) highlight by identifying groups of individuals and situations where their agreement significantly differs from their usual [Belfodil et al., 2017].

Putting claims into perspective (continued) - data mining

Example

Claim about European parliament:

- "*Socialists and Democrats deputies (left wing) usually disagree with Conservatives and Reformists (right wing)*".

Factchecking - **enlightening** with respect to ParlTrack [Marsiske, 2018] dataset:

- True in general, considering all votes.
- But they tend to have convergent opinions on ballots concerning the specific theme "*bilateral agreement and relations with countries external to the union*"[13].

[13] E.g. "*implementation of the Free Trade Agreement between the EU and the Republic of Korea*".

Putting claims into perspective (continued) - data mining

Example

Claim about medicine consumption in France:

- "*Women consume more medicines than men*".

Factchecking - **enlightening** with respect to OpenMedic [Maladie, 2018] dataset:

- True in general (1.32 times more).
- A salient point: women consume 5.13 times more medicines for thyroid therapy than men.
- But, another salient point: men consume 3.0 times more drugs against gout sickness than women.

Putting claims into perspective (continued) - data mining

For claims

- comparing groups' behaviors,
- checkable over behavioral data sets (votes, consumption, ratings. . .).

Approach

- Check the claim.
- Analyze the general behavior.
- Search for situations leading to salient behaviors.

Putting claims into perspective (continued) - data mining

Problem: identify when pairwise group behavior goes against their usual likeness or alikeness.

Patterns $\langle g_1, g_2, c \rangle$ where:

- g_1 and g_2 are descriptions of groups: conjunction of conditions over individuals' attributes (e.g. sex, age, nationality, political group...).
- c is the description of a context: conjunction of conditions over entities (e.g. themes of ballots, date of ballots...).

Compare the behavior of g_1 and g_2 in a specific context c to their reference behavior obtained considering the usual context. The difference is noted $\varphi(\langle g_1, g_2, c \rangle)$.

Extract the patterns showing an exceptional difference.

Putting claims into perspective (continued) - data mining

- Avoid redundancy by considering only the most informative pattern among equivalent patterns in terms of their support.

- Prune unpromising sub-search spaces by using tight optimistic estimates on φ.

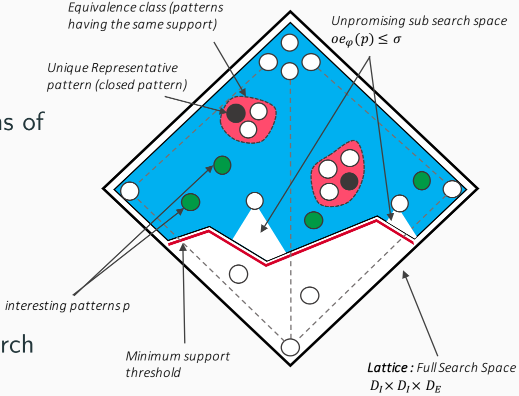

Putting claims into perspective (continued) - build a general picture

- Attempts to build a general and balanced picture of a complex issue [Sato et al., 2015].

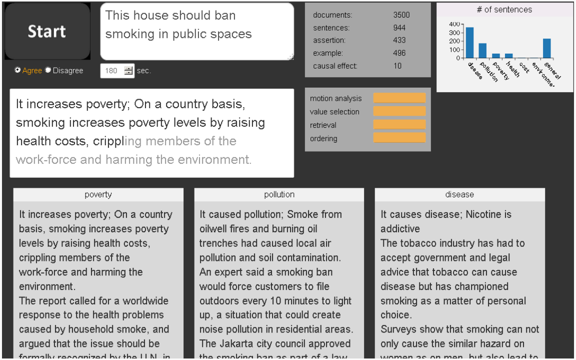

Sharing and publishing fact checking results

- DeFacto shares outputs as RDF graphs with provenance information [Lehmann et al., 2012];
- ClaimBuster provides access to their fact checking outputs [Hassan et al., 2015, Hassan et al., 2017];
- FactCheck.org and PolitFact provide API access, and their output is already used by several other tools

Structured Journalism

- Structured Journalism[14] encourages journalists to publish database items to simplify aggregating, mashing and referencing stories

Source: https://project.wnyc.org/traffic-deaths-2015/

[14] https://reporterslab.org/structured-journalism/

Web Annotations

- W3C's Web Annotation Working Group published recommendation's data model, vocabulary and protocol

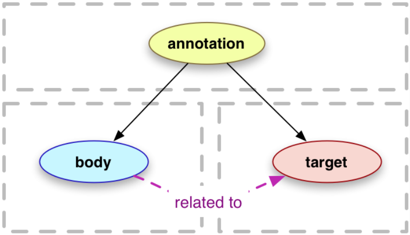

Source: w3.org/TR/annotation-model/

Source: web.hypothes.is/journalism/

- ClaimReview[15] was introduced to Schema.org in 2017.
- Used by search engines to quickly find analysis on past claims.
- Share the facts by Jigsaw, an Alphabet innovation incubator, and Duke Reporters' Lab facilitates sharing fact checking articles.

Source: www.sharethefacts.org/

[15] http://schema.org/ClaimReview

Publishing ClaimReview using MicroFormat

```
<div itemscope="" itemtype="http://schema.org/ClaimReview">
  An example paragraph reviewing a claim expressed in another document.
  <dl>
    <dt>Date published:</dt>
    <dd itemprop="datePublished">2014-07-23</dd>
    <dt>Review url:</dt>
    <dd itemprop="url">http://www.politifact.com/texas/statements/2014/jul/23/rick-perry/
    <dt>Review by:</dt>
    <dd>
     <span itemprop="author" itemscope="" itemtype="http://schema.org/Organization">
        <span itemprop="name"><a itemprop="url" href="http://www.politifact.com/">Politi
        <img itemprop="image" src="http://static.politifact.com/mediapage/jpgs/politifac
        <link itemprop="sameAs" href="http://twitter.com/politifact"/>
     </span>
    </dd>
  </dl>
  <h3>Claim reviewed:</h3>
    <blockquote itemprop="claimReviewed">
    More than 3,000 homicides were committed by 'illegal aliens' over the past six years.
    </blockquote>
    <span itemprop="reviewRating" itemscope="" itemtype="http://schema.org/Rating">
      Rating: <span itemprop="ratingValue">1</span>
      (best score: <span itemprop="bestRating">6</span>),
      "<span itemprop="alternateName">True</span>".
      <img itemprop="image" src="http://static.politifact.com.s3.amazonaws.com/rulings/to
```

ClaimReview as used in Google News

Source: news.google.com

End-to-end systems: ClaimBuster [Hassan et al., 2017]

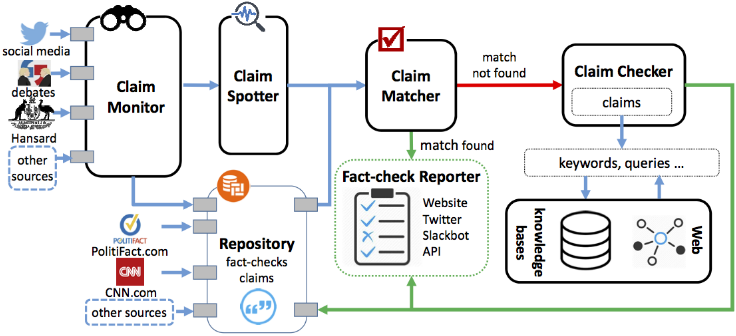

- Assigns a score to each sentence based on how factual it is (low = subjective or opinionated phrase.)
- Above a certain threshold, a claim is matched against a set of fact-checking websites
- Complementary evidence collected from general knowledge bases, otherwise search engines

End-to-end system: ClaimBuster (continued)

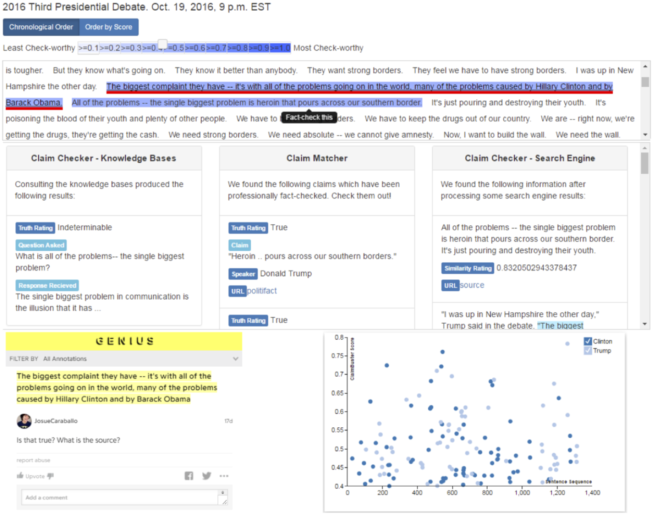

[Hassan et al., 2017]

End-to-end system: FullFact.org [Babakar and Moy, 2016]

End-to-end system: FullFact.org (continued)

[Babakar and Moy, 2016]

End-to-end system: DeFacto [Gerber et al., 2015]

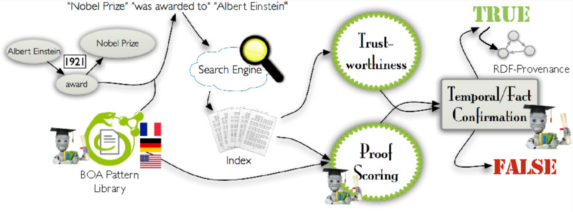

- The system takes as input an RDF Triple, or a sentence that can be translated into one.
- Returns a set of pages, or excerpts thereof, w/ source trustworthiness (relying on PageRank, and page authority on a given topic)
- Confidence score computed based on the number of proofs found and source trustworthiness.
- Try to match the triple against the Linked Open Dataset
- The search for matches is done by verbalizing the input RDF triples and relying on search engines.

End-to-end systems: DeFacto (continued)

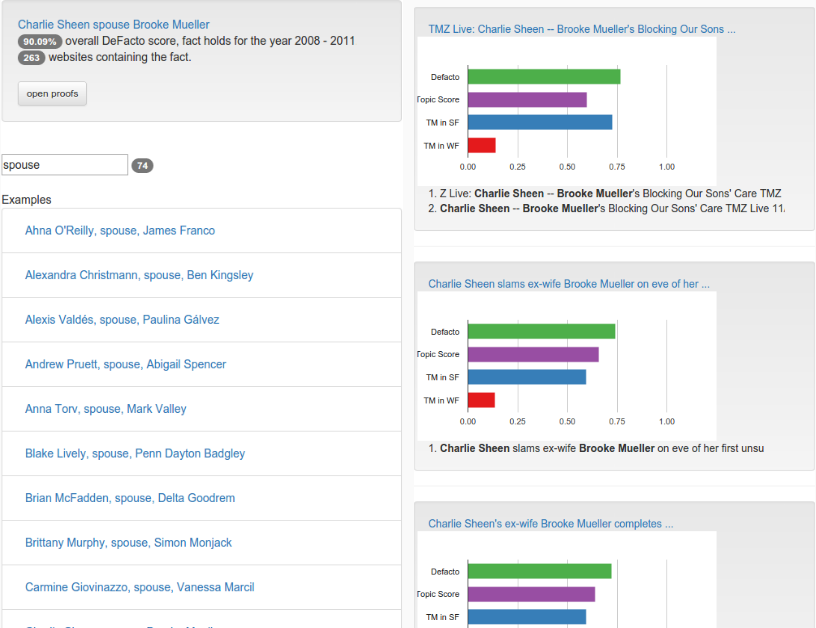

Online fact checking: Truthteller (2013)

Source: truthteller.washingtonpost.com (now discontinued)

- Now defunct
- **Task**: Given a video of a discourse/debate, identify claims and link them to a trusted source of fact-checked claims (FactCheck.org).
- **Approach**: Speech recognition and basic similarity metrics between texts (from videos and from the trusted database).

Overview

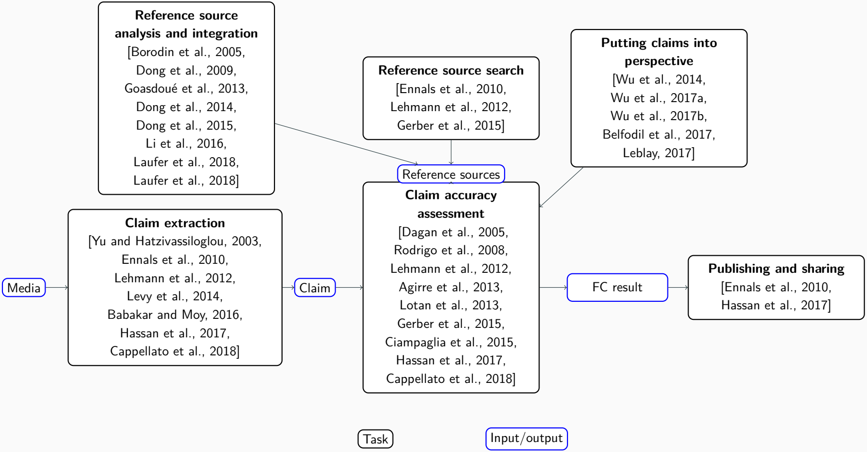
Reference source analysis and integration [Borodin et al., 2005, Dong et al., 2009, Goasdoué et al., 2013, Dong et al., 2014, Dong et al., 2015, Li et al., 2016, Laufer et al., 2018, Laufer et al., 2018]
Reference source search [Ennals et al., 2010, Lehmann et al., 2012, Gerber et al., 2015]
Putting claims into perspective [Wu et al., 2014, Wu et al., 2017a, Wu et al., 2017b, Belfodil et al., 2017, Leblay, 2017]
Reference sources
Claim accuracy assessment [Dagan et al., 2005, Rodrigo et al., 2008, Lehmann et al., 2012, Agirre et al., 2013, Lotan et al., 2013, Gerber et al., 2015, Ciampaglia et al., 2015, Hassan et al., 2017, Cappellato et al., 2018]
Media
Claim extraction [Yu and Hatzivassiloglou, 2003, Ennals et al., 2010, Lehmann et al., 2012, Levy et al., 2014, Babakar and Moy, 2016, Hassan et al., 2017, Cappellato et al., 2018]
Claim
FC result
Publishing and sharing [Ennals et al., 2010, Hassan et al., 2017]
Task
Input/output

Perspectives

Outline

Better foundations

Agreed-upon notions on event, issue, **claim**, **context** and stance would help

- validate new approaches,
- evaluate their coverage and efficiency,
- compare their capabilities.

Quality control

- Facebook discontinued the "Disputed Stories" experiment, following complaint over quality and potential bias[16]

- Finer-grained check-worthiness recognition.
  Current systems rate on a scale from factual to opinionated. It would be useful to rate how context-dependent a factual claim is.
  E.g., "This city's taxes have gone up 20% since the last elections" cannot be checked without context [Babakar and Moy, 2016].

[16] newsroom.fb.com/news/2017/12/news-feed-fyi-updates-in-our-fight-against-misinformation/

Transparency, interpretability, accountability

- Transparency is technically easy, but usually not enough

Example (Fake news detection)

Publishing the machine learning model for a fake news detection system goes in the right direction, but such models are hardly interpretable.

- Interpretability is harder to achieve and typically requires foundations [Ribeiro et al., 2016, Molnar, 2018].

Example (Expert systems)

Expert systems used to have an “explain facility”. We probably need it back!

- Accountability concerns ownership of statements, i.e., who-said-what. vs. who-reported-where. The vast literature on provenance likely has a role to play.

Collaboration

- From exchanging trusted data or previous fact checks, to coordinated work to face difficult investigations
- When fact checkers are of different sensibilities, fact checking becomes less partisan and credibility improves.
- Collaboration empowered by content management tools is a strong trend in journalism, promoted by organizations such as the ICIJ[17]
- CrossCheck[18] is a premier example of this trend.

[17] International Consortium of Investigative Journalists, behind the Panama Papers and other such high-profile international investigations.
[18] https://crosscheck.firstdraftnews.com/

Collaboration (continued)

CJ Workbench: an online data curation and sharing platform for journalists.

Source: cjworkbench.org

Standardization

- Beyond "ClaimReview" more standards are needed to cover fact checking protocols, tools and functions.
- A common and open framework for naming issues and events, and describing their interaction
- A common framework for managing time in Web data

Pluridisciplinarity

- Social and cognitive sciences useful to help devise psychologically effective fact-checking tools
- A recent whitepaper makes recommendations toward making fact checking more convincing, making it reach a larger audience, and avoiding viral misinformation[19]
- Interactions between computer scientists and journalists have been extremely fruitful for both sides [Diakopoulos, 2012]

Example (Computational lead finding)

An analysis of the way Wisconsin voting districts are drawn[20], highlighting the (very) low probability that they may result from an "honest" design.
In the article's words, "it's math versus math, with democracy at stake".

[19] americanpressinstitute.org/publications/reports/white-papers/future-of-fact-checking/

[20] nytimes.com/2017/10/06/opinion/sunday/computers-gerrymandering-wisconsin.html

Focus on issues over claims

- Most newsworthy questions are usually broader than just a claim. E.g., a misleading statement about criminal activity of refugees in the countries receiving them participates to a larger discussion about immigration, and the way different political parties argue it should be handled...
- One of the main points discussed in recent report by the American Press Institute[19].

Education

- Data literacy, envisioned as a set of math and statistic skills, through dedicated education modules at all levels.
- Some news outlets, e.g. France24's The Observer, have dedicated content on critical thinking and news verification.
- The Google News Initiative and CrossCheck FR now organize fact checking classes.
- Computer literacy is gaining ground in school curricula[21] Understanding the way media and communication works gives further tools to discern manipulation, statistic or otherwise.

[21] See e.g. the course "Calling Bullshit: Data Reasoning in a Digital World" created at U. Washington, `http://callingbullshit.org/syllabus.html`

Adapting the delivery

- Timely, sharp and balanced results.
- Avoiding frontal attack on one's convictions and beliefs.
- Choice of the best media for fact checking to reach each audience group.
- Engage and entertain the audience
  - Fact checking success is (also) judged by the audience it can gather and retain.

Outline

Data modeling and storage

Journalism is one of the last industries to adopt digital tools.
Many industries have successfully carried this transition, however for journalists, it is complicated by:

- Historical focus on **text** (not structured records)
- Strong focus on **creativity** and **speed** over **procedure**, **discipline**, **long-run**
- Lack of a **single application domain** (across the newroom); doable for specialized journalists or teams
- Limited financial means, with some notable exceptions (Ouest France)
- In many newrooms, there is no long-term persistent content management plan beyond archiving own articles
- A start: **reference databases**, e.g., of sports teams, precincts, public figures, companies...

Data modeling and storage

What kinds of data journalists need to use?

- Whatever they can get their hands on
- Popular formats: PDF, JSON, CSV, XLS etc.
- Need automatic data types detection

All these data types need back-up mechanisms, e.g. (cloud-storage), CMS functionalities...

Data matching, linking and integration

Data comes in heterogeneous data models, schemas

Data sources partially overlap (or have similar topics) but have been produced in isolation.

This is the perfect setting in need of:

1. Entity recognition: identifying in text, mentions of a known structured entiy
   - Link incoming text article to the entities it features, as they are described in the reference database
2. Entity linking: recognizing when two structured objects are the same
   - Well-known problem in the Web context

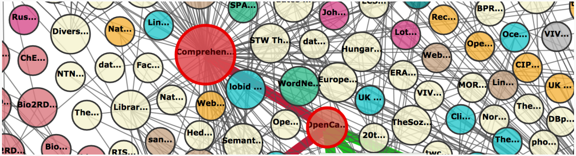

   - Particular twist: strong bonus for trusted data

Natural Language Processing

- Basic NLP functionalities are used in some newsrooms but hardly exploited in workflows
  - Named entity recognition (person, location, organization names)
  - Smart search
  - Voice recognition

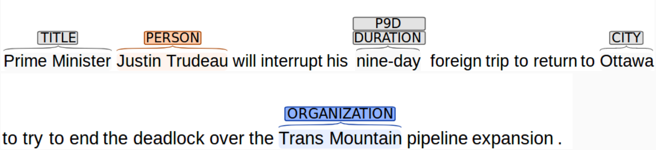

Source: Stanford CoreNLP output

Natural Language Processing (continued)

- Advanced approaches are not ready for production use by non-specialists, and require important human annotation effort:
  - Domain-specific entity extraction

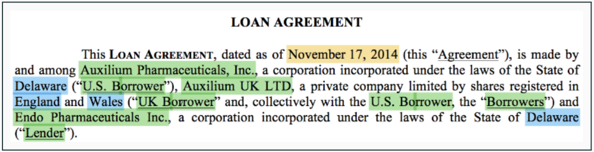

LOAN AGREEMENT

This LOAN AGREEMENT, dated as of November 17, 2014 (this "Agreement"), is made by and among Auxilium Pharmaceuticals, Inc., a corporation incorporated under the laws of the State of Delaware ("U.S. Borrower"), Auxilium UK LTD, a private company limited by shares registered in England and Wales ("UK Borrower" and, collectively with the U.S. Borrower, the "Borrowers") and Endo Pharmaceuticals Inc., a corporation incorporated under the laws of the State of Delaware ("Lender").

Source: [Alvarado et al., 2015]

  - Relation extraction

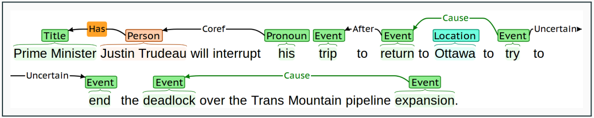

  - Knowledge discovery

Natural Language Processing (continued)

- Bottlenecks currently tackled by NLP researchers:
  - Data quality: how to perform a good extraction from noisy or sparse data
  - Data heterogeneity: how to deal with knowledge distributed over structured, semi-structured and unstructured datasets.
  - Supervision: current effective approaches requires an important amount of human-annotated data. Reducing the need for human supervision is critical (distant supervision, active learning, domain adaptatation, transfer learning, etc.)
  - Reasoning and inference is still limited.
  - Interpretability is a key challenge.
  - Industrial grade systems are still not the rule.
- Efficient NLP systems for fact checking will have to be crosslingual!

Time management

Almost **everything** is time-dependent

- Facts, beliefs and data evolves in time and have a limited period of validity.
- Events have start and end points.
- Fact check results become outdated, also!

The time dimension can be the news!

The time when someone does, say or learns something can make the difference between

- A willful lie or ignorance
- Lawful or criminal behavior, e.g., insider trading, lying to investigators

Example (Comey vs. Trump)

"you have to understand the chronology. The underlying question is whether Trump's firing of Comey constituted obstruction of justice, which has a great deal to do with Flynn"[22]

Follows a chronology on **11 dates** and the conclusion:

"If we accept Comey's account [...], then Trump asked Comey to drop the investigation of Flynn **after** members of his staff knew he had lied to the VP about it, and might even have had reason to believe he had lied to the FBI as well".

[22] https://www.washingtonpost.com/blogs/plum-line/wp/2018/04/20/heres-another-telling-revelation-in-the-comey-memos

Time management (continued)

Tracing data and its evolution for accuracy, transparency, reproducibility.

A FCMS should record and permanently store time information such as

- Data creation time stamp
- Acquisition times
- Statement date
- Version management
- In text, temporal expressions and their relations with events

Yet, despite a W3C recommendation (OWL-Time), there is no widely used standard for representing time in Web data.

Data quality management

Applying data life cycle management tools to reference sources and fact checks.

[Auer et al., 2012]

Enlisting experts

- Prodiving tools for experts to validate complex claims

Enlisting experts (continued)

Source: climatefeedback.org/

Support for reproducibility

Enabling to "replay" fact checking effort and get the same results.

- Fact checking can be seen as a scientific or forensic work
  $\rightarrow$ Reproducibility is needed
- This means:
  - Defining and structuring the fact-checking process, inputs and outputs
  - Keeping trace of manual fact-checking processes
  - Building multilingual benchmarks more complex than binary fake-news benchmarks
- Sharing reproducible results can help:
  - Strengthening a scientific community and accelerating the research.
  - Preserving (or regaining) citizens' trust.
  - Important for fact-checking in general, not only for automated solutions.

Thank you for your attention.

Questions?

References i

Agirre, E., Cer, D., Diab, M., and Gonzalez-Agirre, A. (2013).
***Sem 2013 Shared Task: Semantic Textual Similarity.**
In Proceedings of the Second Joint Conference on Lexical and Computational Semantics and the Shared Task: Semantic Textual Similarity, pages 32–43.

Alvarado, J. C. S., Verspoor, K., and Baldwin, T. (2015).
Domain adaption of named entity recognition to support credit risk assessment.
In Proceedings of the Australasian Language Technology Association Workshop 2015, pages 84–90, Parramatta, Australia.

Auer, S., Bühmann, L., Dirschl, C., Erling, O., Hausenblas, M., Isele, R., Lehmann, J., Martin, M., Mendes, P. N., Van Nuffelen, B., et al. (2012).
Managing the life-cycle of linked data with the lod2 stack.
In International semantic Web conference, pages 1–16. Springer.

Babakar, M. and Moy, W. (2016).
The state of automated factchecking.
https://fullfact.org/media/uploads/full_fact-the_state_of_automated_factchecking_aug_2016.pdf.

Bar-Haim, R., Bhattacharya, I., Dinuzzo, F., Saha, A., and Slonim, N. (2017).
Stance Classification of Context-Dependent Claims.
In Conf. of the European Chapter of ACL: Volume 1, Long Papers, pages 251–261, Valencia, Spain. Association for Computational Linguistics.

Barrón-Cedeño, A., Elsayed, T., Suwaileh, R., Màrquez, L., Atanasova, P., Zaghouani, W., Kyuchukov, S., Martino, G. D. S., and Nakov, P. (2018).
Overview of the CLEF-2018 checkthat! lab on automatic identification and verification of political claims. task 2: Factuality.
In CLEF (Working Notes), volume 2125 of CEUR Workshop Proceedings. CEUR-WS.org.

Belfodil, A., Cazalens, S., Lamarre, P., and Plantevit, M. (2017).
Flash points: Discovering exceptional pairwise behaviors in vote or rating data.
In PKDD, pages 442–458.

Bode, L. and Vraga, E. K. (2015).
In related news, that was wrong: The correction of misinformation through related stories functionality in social media.
Journal of Communication, 65(4):619–638.

References v

Bonaque, R., Cao, T. D., Cautis, B., Goasdoué, F., Letelier, J., Manolescu, I., Mendoza, O., Ribeiro, S., Tannier, X., and Thomazo, M. (2016).
Mixed-instance querying: a lightweight integration architecture for data journalism.
PVLDB, 9(13):1513–1516.

Borodin, A., Roberts, G. O., Rosenthal, J. S., and Tsaparas, P. (2005).
Link analysis ranking: Algorithms, theory, and experiments.
ACM Trans. Internet Technol., 5(1):231–297.

Cao, T. D., Manolescu, I., and Tannier, X. (2017).
Extracting Linked Data from statistic spreadsheets.
In Proceedings of the SIGMOD workshop on Semantic Big Data (SBD 2017), Chicago, USA.

Cao, T. D., Manolescu, I., and Tannier, X. (2018).
Searching for Truth in a Database of Statistics.
In Proceedings of the 21st International Workshop on the Web and Databases (WebDB 2018), Houston, USA.

Cappellato, L., Ferro, N., Nie, J., and Soulier, L., editors (2018).
Working Notes of CLEF 2018 - Conference and Labs of the Evaluation Forum, Avignon, France, September 10-14, 2018, volume 2125 of CEUR Workshop Proceedings. CEUR-WS.org.

Cazalens, S., Lamarre, P., Leblay, J., Manolescu, I., and Tannier, X. (2018).
A Content Management Perspective on Fact-Checking.
In
The Web Conference 2018 - alternate paper tracks "Journalism, Misinforma
pages 1–10, Lyon, France.

Chanial, C., Dziri, R., Galhardas, H., Leblay, J., Nguyen, M. L., and Manolescu, I. (2018).
Connectionlens: Finding connections across heterogeneous data sources.
PVLDB, 11(12):2030–2033.

Ciampaglia, G. L., Shiralkar, P., Rocha, L. M., Bollen, J., Menczer, F., and Flammini, A. (2015).
Computational fact checking from knowledge networks.
PloS one, 10(6):e0128193.

Cohen, S., T., H. J., and Turner, F. (2011).
Computational journalism.
Commun. ACM, 54(10):66–71.

Dagan, I., Glickman, O., and Magnini, B. (2005).
The PASCAL Recognising Textual Entailment Challenge.
In PASCAL Challenges Workshop for Recognizing Textual Entailment.

Diakopoulos, N. (2012).
Cultivating the landscape of innovation in computational journalism.
http://www.nickdiakopoulos.com/wp-content/uploads/2012/05/diakopoulos_whitepaper_systematicinnovation.p

Dong, X., Gabrilovich, E., Heitz, G., Horn, W., Lao, N., Murphy, K., Strohmann, T., Sun, S., and Zhang, W. (2014).
Knowledge Vault: A Web-scale Approach to Probabilistic Knowledge Fusion.
In Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, pages 601–610, New York, NY, USA. ACM.

Dong, X. L., Berti-Equille, L., and Srivastava, D. (2009).
Truth discovery and copying detection in a dynamic world.
Proceedings of the VLDB Endowment, 2(1):562–573.

References x

Dong, X. L., Gabrilovich, E., Murphy, K., Dang, V., Horn, W., Lugaresi, C., Sun, S., and Zhang, W. (2015).
Knowledge-Based Trust: Estimating the Trustworthiness of Web Sources.
In Proceedings of the VLDB Endowment.

Ennals, R., Trushkowsky, B., and Agosta, J. M. (2010).
Highlighting disputed claims on the web.
In Proceedings of the 19th international conference on World wide web, pages 341–350. ACM.

Faulkner, A. (2014).
Automated Classification of Stance in Student Essays: An Approach Using Stance Target Information and the Wikipedia Link-Based Measure.
In Int.'l Florida Artificial Intelligence Research Society Conference, Pensacola Beach, USA. Association for the Advancement of Artificial Intelligence.

Ferreira, W. and Vlachos, A. (2016).
Emergent: a novel data-set for stance classification.
In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1163–1168, San Diego, California. Association for Computational Linguistics.

Franklin, M., Halevy, A., and Maier, D. (2005).
From databases to dataspaces: A new abstraction for information management.
SIGMOD Record, 34(4):27–33.

Garrett, R. K. (2009).
Echo chambers online?: Politically motivated selective exposure among internet news users1.
Journal of Computer-Mediated Communication, 14(2):265–285.

Garrett, R. K. (2016).
Facebook's problem is more complicated than fake news.
http://theconversation.com/facebooks-problem-is-more-complicated-than-fake-news-68886.

Garrett, R. K. and Weeks, B. E. (2013).
The promise and peril of real-time corrections to political misperceptions.
In Proceedings of the 2013 conference on Computer supported cooperative work, pages 1047–1058. ACM.

Gencheva, P., Nakov, P., Màrquez, L., Barrón-Cedeño, A., and Koychev, I. (2017).
A context-aware approach for detecting worth-checking claims in political debates.
In Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017, pages 267–276. INCOMA Ltd.

Gerber, D., Esteves, D., Lehmann, J., Bühmann, L., Usbeck, R., Ngomo, A.-C. N., and Speck, R. (2015).
DeFacto—temporal and multilingual deep fact validation.
Web Semantics: Science, Services and Agents on the World Wide Web, 35:85–101.

Goasdoué, F., Karanasos, K., Katsis, Y., Leblay, J., Manolescu, I., and Zampetakis, S. (2013).
Fact checking and analyzing the web.
In SIGMOD, pages 997–1000.

Guggilla, C., Miller, T., and Gurevych, I. (2016).
CNN- and LSTM-based claim classification in online user comments.
In Proceedings of the 26th International Conference on Computational Linguistics: Technical Papers (COLING 2016), pages 2740–2751.

Hasan, K. S. and Ng, V. (2013).
Stance Classification of Ideological Debates: Data, Models, Features, and Constraints.
In Proceedings of the Sixth International Joint Conference on Natural Language Processing, pages 1348–1356, Nagoya, Japan. Asian Federation of Natural Language Processing.

- Hassan, N., Li, C., and Tremayne, M. (2015).
  Detecting check-worthy factual claims in presidential debates.
  In Proceedings of International on Conference on Information and Knowledge Management, pages 1835–1838. ACM.
- Hassan, N., Zhang, G., Arslan, F., Caraballo, J., Jimenez, D., Gawsane, S., Hasan, S., Joseph, M., Kulkarni, A., Nayak, A. K., et al. (2017).
  Claimbuster: The first-ever end-to-end fact-checking system.
  Proceedings of the VLDB Endowment, 10(7).
- Laufer, C., Schwabe, D., and Busson, A. (2018).
  Ontologies for representing relations among political agents.
  CoRR, abs/1804.06015.

Leblay, J. (2017).
A declarative approach to data-driven fact checking.
In AAAI, pages 147–153.

Lehmann, J., Gerber, D., Morsey, M., and Ngomo, A.-C. N. (2012).
DeFacto-deep fact validation.
In International Semantic Web Conference, pages 312–327. Springer.

Levy, R., Bilu, Y., Hershcovich, D., Aharoni, E., and Slonim, N. (2014).
Context Dependent Claim Detection.
In COLING, pages 1489–1500.

Li, Y., Gao, J., Meng, C., Li, Q., Su, L., Zhao, B., Fan, W., and Han, J. (2016).
A survey on truth discovery.
Acm Sigkdd Explorations Newsletter, 17(2):1–16.

Lotan, A., Stern, A., and Dagan, I. (2013).
TruthTeller: Annotating Predicate Truth.
In Proceedings of NAACL-HLT 2013, page 752–757, Atlanta, USA.

Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B. J., Wong, K.-F., and Cha, M. (2016).
Detecting rumors from microblogs with recurrent neural networks.
In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16, pages 3818–3824. AAAI Press.

Maladie, A. (2018).
Open medic : base complète sur les dépenses de médicaments interrégimes.
https://www.data.gouv.fr/fr/datasets/open-medic-base-complete-sur-les-depenses-de-medicaments-interregimes/.

Marsiske, S. (2018).
Parltrack european initiative.
http://parltrack.euwiki.org.

Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., and Cherry, C. (2016).
SemEval-2016 Task 6: Detecting Stance in Tweets.
In Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016), pages 31–41, San Diego, California. Association for Computational Linguistics.

Molnar, C. (2018).
Interpretable Machine Learning. A Guide for Making Black Box Models Explainable.
https://christophm.github.io/interpretable-ml-book/.
github: https://github.com/christophM/iml.

Murakami, A. and Raymond, R. (2010).
Support or Oppose? Classifying Positions in Online Debates from Reply Activities and Opinion Expressions.
In COLING, pages 869–875.

Nyhan, B. and Reifler, J. (2010).
When corrections fail: The persistence of political misperceptions.
Political Behavior, 32(2):303–330.

Park, J. and Cardie, C. (2014).
Identifying appropriate support for propositions in online user comments.
In Proceedings of the First Workshop on Argumentation Mining, pages 29–38, Baltimore, Maryland. Association for Computational Linguistics.

Rajadesingan, A. and Liu, H. (2014).
Identifying Users with Opposing Opinions in Twitter Debates.
In Kennedy, W. G., Agarwal, N., and Yang, S. J., editors, Social Computing, Behavioral-Cultural Modeling and Prediction: 7th International Conference, SBP 2014, Washington, DC, USA, April 1-4, 2014. Proceedings, pages 153–160, Cham. Springer International Publishing.

Ribeiro, M. T., Singh, S., and Guestrin, C. (2016).
Why Should I Trust You?": Explaining the Predictions of Any Classifier.
In SIGKDD, pages 1135–1144, San Francisco, USA.
github: https://github.com/marcotcr/lime.

Rodrigo, A., Peñas, A., and Verdejo, F. (2008).
Overview of the Answer Validation Exercise 2008.
In Proceedings of the 9th Workshop of the Cross-Language Evaluation Forum, CLEF 2008, pages 296–313.

Rubin, V. L., Chen, Y., and Conroy, N. J. (2015).
Deception detection for news: three types of fakes.
Proceedings of the Association for Information Science and Technology, 52(1):1–4.

Sato, M., Yanai, K., Miyoshi, T., Yanase, T., Iwayama, M., Sun, Q., and Niwa, Y. (2015).
End-to-end argument generation system in debating.
In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL), System Demonstrations, pages 109–114.

Shin, J., Jian, L., Driscoll, K., and Bar, F. (2017).
Political rumoring on twitter during the 2012 us presidential election: Rumor diffusion and correction.
New Media & Society, 19(8):1214–1235.

Shu, K., Sliva, A., Wang, S., Tang, J., and Liu, H. (2017).
Fake news detection on social media: A data mining perspective.
ACM SIGKDD Explorations Newsletter, 19(1):22–36.

Somasundaran, S. and Wiebe, J. (2009).
Recognizing stances in online debates.
In Joint ACL/AFNLP Conference, pages 226–234, Suntec, Singapore. Association for Computational Linguistics.

Thomas, M., Pang, B., and Lee, L. (2006).
Get out the vote: Determining support or opposition from congressional floor-debate transcripts.
In Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, pages 327–335, Sydney, Australia. Association for Computational Linguistics.

Vosoughi, S., Roy, D., and Aral, S. (2018).
The spread of true and false news online.
Science, 359:1146–1151.

Walker, M., Tree, J. F., Anand, P., Abbott, R., and King, J. (2012).
A corpus for research on deliberation and debate.
In Chair), N. C. C., Choukri, K., Declerck, T., Doğan, M. U., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., and Piperidis, S., editors, Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12), Istanbul, Turkey. European Language Resources Association (ELRA).

Wood, M. J., Douglas, K. M., and Sutton, R. M. (2012).
Dead and alive: Beliefs in contradictory conspiracy theories.
Social Psychological and Personality Science, 3(6):767–773.

Wu, Y., Agarwal, P. K., Li, C., Yang, J., and Yu, C. (2014).
Toward computational fact-checking.
Proceedings of the VLDB Endowment, 7(7):589–600.

- Wu, Y., Agarwal, P. K., Li, C., Yang, J., and Yu, C. (2017a).
  Computational fact checking through query perturbations.
  ACM Transactions on Database Systems (TODS), 42(1):4.
- Wu, Y., Gao, J., Agarwal, P. K., and Yang, J. (2017b).
  Finding diverse, high-value representatives on a surface of answers.
  Proceedings of the VLDB Endowment, 10(7):793–804.
- Yu, H. and Hatzivassiloglou, V. (2003).
  Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences.
  In Proceedings of 2003 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 129–136.

Zubiaga, A., Liakata, M., Procter, R., Wong Sak Hoi, G., and Tolmie, P. (2016).
Analysing How People Orient to and Spread Rumours in Social Media by Looking at Conversational Threads.
PLoS ONE, 11(3).

Computational Fact Checking

A Content Management Perspective

Sylvie Cazalens[1] Philippe Lamarre[1] Julien Leblay[2] Ioana Manolescu[3] Xavier Tannier[4]

PVLDB Conference - August 2018

[1]Université de Lyon, INSA-Lyon – CNRS, LIRIS, France

[2]Artificial Intelligence Research Center – AIST, Japan

[3]Inria, LIX – CNRS and École Polytechnique, France

[4]Sorbonne Université – Inserm – LIMICS, France