



HAL
open science

Teasing journalistic findings out of heterogeneous sources: a data/AI journey

Ioana Manolescu

► **To cite this version:**

Ioana Manolescu. Teasing journalistic findings out of heterogeneous sources: a data/AI journey. DEBS 2022: The 16th ACM International Conference on Distributed and Event-based Systems, Jun 2022, Copenhagen, Denmark. pp.1-1, 10.1145/3524860.3544406 . hal-03945733

HAL Id: hal-03945733

<https://inria.hal.science/hal-03945733v1>

Submitted on 18 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Inria



Teasing journalistic findings out of heterogeneous sources: a data/AI journey

Ioana Manolescu

Inria Saclay-Île-de-France, Institut Polytechnique de Paris



Outline

1. Motivation: why journalism?
2. Data integration problems raised by data journalism
3. Graph-based data integration in ConnectionLens
4. Related work and perspectives



Motivation: Why journalism?

Bad memories: Romania, 1989



Bad memories: Romania, 1989



Ceaușescu re-elected
at the 14th congress!

Bad memories: Romania, 1989



Ceaușescu re-elected
at the 14th congress!

He had held power
since 1965.

Bad memories: Romania, 1985



1990: things got better



... kind of



1000 dead (approx.)
No one convicted.

Democratic societies crucially need the press

- ❑ To debate and express dissent
- ❑ To analyze, confirm or refute public statements
- ❑ To expose and explain society functioning



Socialist Romania, 1984

Fact-checking

(Data) journalism



Fact-checking


Not everyone agrees, however, that Democrats are not flip-flopping on the issue.

Mark Krikorian, executive director of the Center for Immigration Studies, a think tank that advocates for lower immigration, said that because the public doesn't know exactly what border barriers the Trump administration wants to build, Mulvaney's statement is not an "exact" comparison. But, he said, to dismiss it simply on that basis would be "tendentiously literal."

"The fact is that, other than the 'Mexico will pay for it' stuff, Trump is simply channeling the 2006 Secure Fence Act, and Schumer et al. who voted for it out of political calculation are indeed hypocrites for opposing the attempt to finally bring that law to fruition," Krikorian told us via email.

At the surface level, it is true in a broad sense that Democrats including Schumer, Obama and Clinton have in the past supported border fencing. All three voted for the Secure Fence Act of 2006, and all three supported the 2013 Senate immigration overhaul that passed the Senate, and which called for tougher border security including some additional fencing. But to claim that those measures are the same as what Trump is proposing is a stretch.

Share The Facts





Mick Mulvaney
Director, Office of Management and Budget

"We don't understand why the Democrats are so wholeheartedly against [President Trump's border wall]. They voted for it in 2006."

Fox News Sunday – Sunday, April 23, 2017

SHARE READ MORE


White House Office of Management and Budget Director Mick Mulvaney made an apples-to-oranges comparison when he said he couldn't understand why Democrats opposed supplemental funding for a border wall since many of them were for it back in 2006. Mulvaney is referring to the Secure Fence Act of 2006, which called for construction of 700 miles of fencing and enhanced surveillance technology, such as unmanned drones, ground-based sensors, satellites, radar coverage and cameras. Sen. Chuck Schumer and then-Sens. Barack Obama and Hillary Clinton were among a bipartisan majority that voted in favor of the legislation, and it was signed into law by President George W. Bush. In a very general sense, the Democrats named by Mulvaney supported a bill to build more

ASK FACTCHECK

Q: Did the Supreme Court rule that public schools cannot teach students about Islam?

A: No. That false claim was spread by a network of fake news websites.

TRUE
POLITIFACT TRUTH-O-METER™

MOSTLY TRUE
POLITIFACT TRUTH-O-METER™

HALF TRUE
POLITIFACT TRUTH-O-METER™

MOSTLY FALSE
POLITIFACT TRUTH-O-METER™

FALSE
POLITIFACT TRUTH-O-METER™

PANTS ON FIRE!
POLITIFACT TRUTH-O-METER™

Data journalism

Panama Papers (International Consortium of Investigative Journalism, ICIJ)

The screenshot shows a web browser displaying the ICIJ website. The page title is "The Panama Papers". The main content area features a profile for Jérôme Cahuzac, a French politician. To the right of the profile is a corporate ownership diagram. The diagram shows "MONFORT CAPITAL PARTNERS JLT" as the registered owner of "CERMAN GROUP LIMITED". "CERMAN GROUP LIMITED" is the beneficial owner of "TALWAY INTERNATIONAL CORP." and also has Jérôme Cahuzac and "Mr. Jerome Andre C." as beneficiaries. The registered address for TALWAY INTERNATIONAL CORP. is listed as "85 avenue de Breteuil Paris-France".

Jérôme Cahuzac
Budget minister at the Ministry of the Economy, Finance and External Trade (2012-2013); Deputy, National Assembly of France (1997-2002, 2007-2012)

Related countries
France

The lies told by Jérôme Cahuzac in 2013 triggered one of the most spectacular downfalls of a public official in the annals of French scandals. As a government minister waging a campaign against tax evasion, Cahuzac was forced to admit he lied to President François Hollande, former colleagues in Parliament and the French people when he repeatedly denied owning foreign bank accounts. He said he stashed over \$750,000 in a Swiss bank account for 20 years, moving the money to Singapore in 2009. His ex-wife disclosed an account opened in Great Britain in 1997. Cahuzac, who made a fortune as a cosmetic surgeon, resigned his ministry post and awaits trial for tax fraud.

Embed

MONFORT CAPITAL PARTNERS JLT registered

CERMAN GROUP LIMITED Beneficial owner

TALWAY INTERNATIONAL CORP. Shareholder

Jérôme Cahuzac Beneficiary

Mr. Jerome Andre C. Beneficiary

registered address

85 avenue de Breteuil Paris-France

Projects and collaborations

Google Award (2015) with U. Paris Sud

ANR ContentCheck (2016-2020) with
Les Décodeurs (Le Monde) <https://contentcheck.inria.fr>

Inria Associated Team WebClaimExplain (2017-2019),
with AIST Japan

Collaboration with H. Galhardas (University of Lisbon),
A. Anadiotis, O. Balalau, E. Pietriga (Inria, Polytechnique)

ANR SourcesSay AI Chair (2020-2023),
with Le Monde and WeDoData <https://sourcessay.inria.fr>

Fact-checking for **RadioFrance** (2022-2023)



Le Monde

LES DÉCODEURS
VENONS-EN AUX FAITS

Inria **AIST**




radiofrance

Inria

Project discussed with Les Décodeurs: fake news detection and propagation on Twitter

Online fact-checks: (semi)structured data sources (JSON, XML) listing

- Link to claim (media, social network etc.), **claim author**
- **Fact-check**, containing: analysis (details), final assessment, fc author, date, institution




Among the first published: <https://www.lemonde.fr/web-service/decodex/updates>

Years later: **ClaimReview** by Google and others (<https://www.claimreviewproject.com/>)

Project discussed with Les Décodeurs: fake news detection and propagation on Twitter


Online fact-checks: (semi)structured data sources (JSON, XML) listing

- Link to claim (media, social network etc.), **claim author**
- **Fact-check**, containing: analysis (details), final assessment, fc author, date




Décodeurs' database of French public figures (Excel)


- First name, last name, Twitter ID, position, political party when known



Question: When does a fake news post first cross into a community, e.g., members of the Parliament?

- Looking for tweets *connected to* a fake news author, *and to a community member*; arbitrary paths (chains of author/likes/retweets/inParty/...)







Integrating heterogeneous journalistic datasets with ConnectionLens

<https://team.inria.fr/cedar/connectionlens/>

Graph-based integration of heterogeneous data sources

- ❑ The sources are **not RDF**. They can be **(semi)structured**, or **unstructured** (text).
- ❑ The sources may be very **dynamic** (projects started and abandoned as per news cycle and data availability).
- ❑ There is no schema. Data producers often uncollaborative.
- ❑ For most journalists, databases do not come naturally, and IT support is limited. They know keyword-based search...



Integrate heterogeneous sources within a graph, query w/ keywords


ConnectionLens: graph-based integration of heterogeneous data sources

<https://team.inria.fr/cedar/connectionlens/>

Joint work with: J. Leblay (AIST Japan),
H. Galhardas and C. Conceição (U. Portugal),
A. Anadiotis, O. Balalau, N. Barret, T. Bouganim,
F. Chimienti, M.-Y. Haddad, T. Merabti,
P. Upadhyay (CEDAR) + interns

S.Horel (Le Monde, European Press Prize
“Investigative Reporting Award 2018”)

Ongoing work in ANR/DGA AI Chair SourcesSay
(<https://sourcessay.inria.fr>), DIM RFSI



ConnectionLens principles [VLDBdemo2018]

Integrate **any kind of data** into a **graph**

Extract entities from any text node (regardless of the model of the data source where the text comes from)

- ❑ Same entity in two different text nodes =
link among the text nodes (*densification* of the graph)

The graph is **heterogeneous** and **irregular** →

Query it through **keywords**: find trees that connect 1 node matching each kwd

- ❑ Closely related to the Group Steiner Tree Problem (GSTP)

ConnectionLens principles [VLDBdemo2018]

Integrate **any kind of data** into a **graph**

Extract entities from any text node (regardless of the model of the data source where the text comes from)

- ❑ Same entity in two different text nodes = link among the text nodes (*densification* of the graph)

The graph is **heterogeneous** and **irregular** →

Query it through **keywords**: find trees that connect 1 node matching each kwd


- ❑ Closely related to the Group Steiner Tree Problem (GSTP)

Rest of the talk based on current state of the project:

[CIKMdemo2021, Elsevier InfSys 2022, IEEE DataEngBull 2021, ongoing works]

ConnectionLens graph construction

The Balkany and their African connections



ENQUETE


Troisième villa saisie pour les Balkany

Par Emmanuel Fansten — 13 août 2015 à 14:58



A screenshot of a Le Point article. The URL is https://www.lepoint.fr/politique/villas-a-marrakech-fonds-occultes-les-epoux-balkany-juges-lundi-12-05-2019-2312058_20.php#. The article title is 'Villas à Marrakech, fonds « occultes »... : les époux Balkany jugés lundi'. The text below the title reads: 'Soupçonnés d'avoir dissimulé 13 millions d'euros d'avoirs au fisc, les édiles de Levallois-Perret comparaissent pour fraude fiscale et blanchiment. Source AFP'. The article is dated 'Publié le 12/05/2019 à 11:19 | Le Point.fr'. There is a social media share bar at the bottom of the article. To the right of the article is a sidebar with a 'Le Point Politique' header, a 'MENU' button, and a section titled 'EN CONTINU' with a list of other news items. At the bottom right of the screenshot is a purple advertisement for 'qonto' with the text 'Ouvrez un compte pro en ligne, contrôlez vos finances simplement' and a 'Découvrir les offres' button.


The Balkany and their African connections



ENQUETE

Troisième villa saisie pour les Balkany

Par Emmanuel Fansten — 13 août 2015 à 14:58



Actualité > Politique

Villas à Marrakech, fonds « occultes »... : les époux Balkany jugés lundi

Soupçonnés d'avoir dissimulé 13 millions d'euros d'avoirs au fisc, les édiles de Levallois-Perret comparaissent pour fraude fiscale et blanchiment.
Source AFP

Publié le 12/05/2019 à 11:19 | Le Point.fr



PROFITEZ DE VOTRE ABONNEMENT À 1€ LE 1ER MOIS !

De somptueuses villas à Marrakech et dans les Caraïbes, des fonds « occultes » transitant par le Panama ou Singapour... Soupçonnés d'avoir dissimulé plus de 13 millions d'euros d'avoirs au fisc, les édiles de Levallois-Perret Patrick et Isabelle Balkany sont jugés à partir de...



3 RÉSULTATS CORRESPONDANT À VOTRE RECHERCHE


RESPONSABLES PUBLICS

REPRESENTANTS D'INTÉRÊTS

ACTIVITÉS DE REPRÉSENTATION D'INTÉRÊTS

Perte de l'autorité morale : demande d'Anticor au Président de la République pour que soient révoqués un maire et son adjointe qui ont avoué avoir fraudé l'administration fiscale

ANTICOR



The Balkanys and their African connections

Public officials transparency high authority (CSV)

Name	Owner	Location	Type
Dar Gyucy	P. Balkany	Marrakech	Real Estate
Moulin Cossy	I. Balkany	Giverny	Real Estate

The Balkanys and their African connections

Public officials transparency high authority (CSV)

Name	Owner	Location	Type
Dar Gyucy	P. Balkany	Marrakech	Real Estate
Moulin Cossy	I. Balkany	Giverny	Real Estate

National Directory of Elected Officials (JSON)

```
[{  
  name: "Levallois-Perret",  
  mayor: "P. Balkany",  
  city-council: [  
    {name: "I. Balkany"},  
    ...  
  ]  
}, ...]
```

The Balkanys and their African connections

Public officials transparency high authority (CSV)

Name	Owner	Location	Type
Dar Gyucy	P. Balkany	Marrakech	Real Estate
Moulin Cossy	I. Balkany	Giverny	Real Estate

dbpedia.org (RDF)

```
{
  dbr:Marrakech
    dbr:name      "Marrakech"
    rdf:type      dbo:City ;
    dbo:country   dbr:Morocco .
  dbr:Morocco
    dbr:name      "Morocco"
    rdf:type      dbo:Country
    dbo:locatedIn dbr:Africa .
  dbr:CentralAfricanRepublic
    dbr:name      "Central African Republic"
    dbo:locatedIn dbr:Africa .
}
```

National Directory of Elected Officials (JSON)

```
[{
  name: "Levallois-Perret",
  mayor: "P. Balkany",
  city-council: [
    {name: "I. Balkany"},
    ...
  ]
}, ...]
```

The Balkany and their African connections

Public officials transparency high authority (CSV)

Name	Owner	Location	Type
Dar Gyucy	P. Balkany	Marrakech	Real Estate
Moulin Cossy	I. Balkany	Giverny	Real Estate

dbpedia.org (RDF)

```
{
  dbr:Marrakech
    dbr:name      "Marrakech"
    rdf:type      dbo:City ;
    dbo:country   dbr:Morocco .
  dbr:Morocco
    dbr:name      "Morocco"
    rdf:type      dbo:Country
    dbo:locatedIn dbr:Africa .
  dbr:CentralAfricanRepublic
    dbr:name      "Central African Republic"
    dbo:locatedIn dbr:Africa .
}
```

National Directory of Elected Officials (JSON)

```
[{
  name: "Levallois-Perret",
  mayor: "P. Balkany",
  city-council: [
    {name: "I. Balkany"},
    ...
  ]
}, ...]
```

Libération – Nov. 13, 2014 (Text)

Balkany mineur de fonds

L'élu de **Levallois-Perret** est soupçonné d'avoir touché 5 millions de dollars de commission en 2009 grâce à son rôle d'intermédiaire entre **Areva** et la **Centrafrique** dans le dossier **Uramin**. [...]

How is Levallois-Perret connected to Africa and "real estate"?

Public officials transparency high authority (CSV)

Name	Owner	Location	Type
Dar Gyucy	P. Balkany	Marrakech	Real Estate
Moulin Cossy	I. Balkany	Giverny	Real Estate

dbpedia.org (RDF)

```
{
  dbr:Marrakech
    dbr:name      "Marrakech"
    rdf:type      dbo:City ;
    dbo:country   dbr:Morocco .
  dbr:Morocco
    dbr:name      "Morocco"
    rdf:type      dbo:Country
    dbo:locatedIn dbr:Africa .
  dbr:CentralAfricanRepublic
    dbr:name      "Central African Republic"
    dbo:locatedIn dbr:Africa .
}
```

National Directory of Elected Officials (JSON)


```
[{
  name: "Levallois-Perret",
  mayor: "P. Balkany",
  city-council: [
    {name: "I. Balkany"},
    ...
  ]
}, ...]
```

Libération – Nov. 13, 2014 (Text)

Balkany mineur de fonds

L'élu de **Levallois-Perret** est soupçonné d'avoir touché 5 millions de dollars de commission en 2009 grâce à son rôle d'intermédiaire entre **Areva** et la **Centrafrique** dans le dossier **Uramin**. [...]


Idea: integrate **all** data sources into a **heterogeneous graph**



Graph construction stages

1. Primary node and edge construction

- ❑ Direct for XML, JSON, RDF, HTML
- ❑ 1 relational tuple=1 node;
primary keys-foreign keys as links
- ❑ Convert information from PDF into:
 - ❑ JSON for text content
 - ❑ RDF describing tables




Graph construction stages

1. Primary node and edge construction

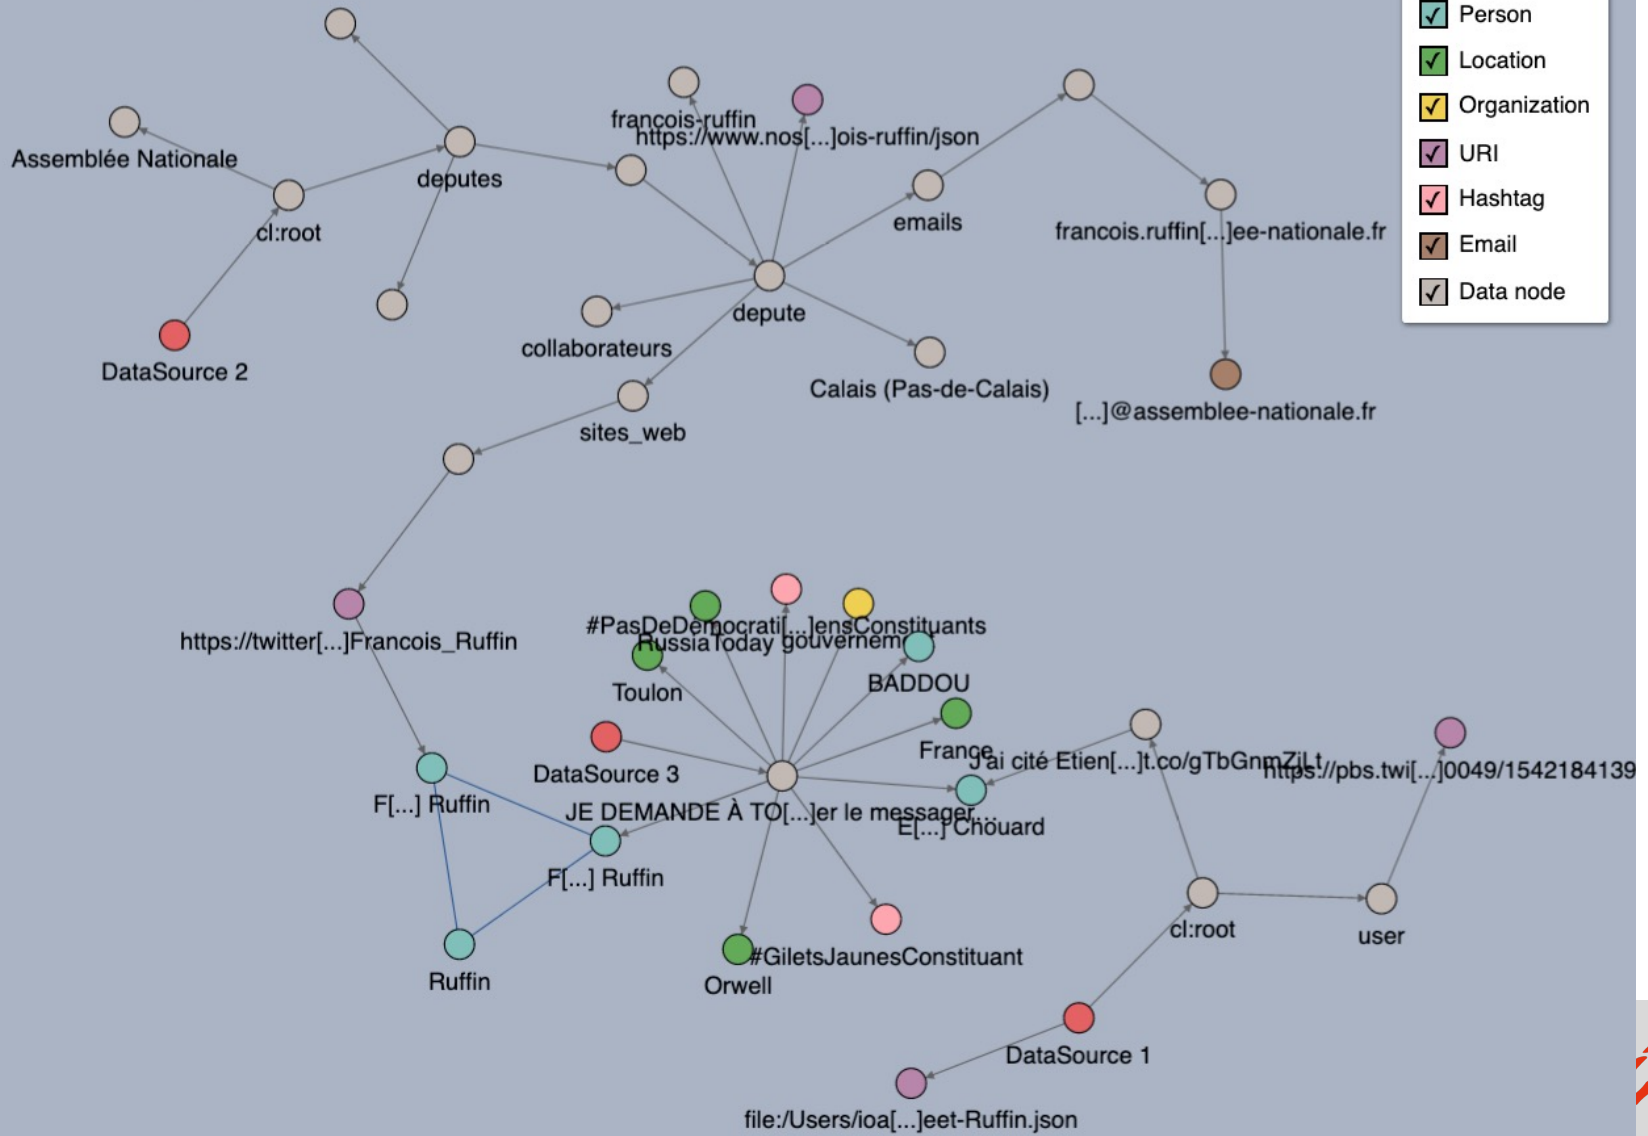
- ❑ Direct for XML, JSON, RDF, HTML
- ❑ 1 relational tuple=1 node; PK-FKs as links
- ❑ [Optional] segment text documents
- ❑ Extract information from PDF into: (a) JSON, and (b) RDF describing tables

2. Entity extraction

- ❑ From all text nodes of all the sources: **entity node** child of text node
- ❑ [VLDB2018]: based on Stanford NER
- ❑ [InfSys2021] Developed and trained new entity extractor FR/EN, based on Flair framework
- ❑ Date extraction based on HeidelTime




Search Q



Graph construction stages

2. Entity extraction

- ❑ From all text nodes of all the sources: **entity node** child of text node
- ❑ [VLDB2018]: based on Stanford NER
- ❑ [BDA2020] Developed and trained new entity extractor from French, based on Flair framework



3. Entity disambiguation

- ❑ For each recognized entity, e.g., "Hollande" the place or the person?
- ❑ Built novel disambiguation pipeline for French, based on Ambiverse framework
 - ❑ Based on knowledge bases (WikiData, YAGO) and Wikipedia
 - ❑ Helpful on well-known entities




Graph construction stages

4. Node matching

- ❑ To create **sameAs edges**:
 - ❑ Strong sameAs edges: equivalent nodes **1.**
 - ❑ Weak sameAs edges: similar nodes **.85**
- ❑ Label normalization, distance functions
- ❑ Remains quadratic at the core ☹️, so...

Node factorization (heuristic): create only one node per label per document (or per graph)



More complex entity matching/cleaning pipelines for: Person, Organization, Location

ConnectionLens graph querying

Querying problem statement

- ❑ Given the graph $G = (N, E)$ built out of the datasets D and a query $Q = \{w_1, \dots, w_m\}$, return the k **highest-score minimal answer trees**.
- ❑ An answer tree is a set of edges which (i) form a tree (ii) contain at least one node whose label matches each keyword w_i .
- ❑ We are interested in **minimal answer trees**, that is:
 - ❑ Removing an edge from the tree should make it lack some keyword(s).
 - ❑ If a keyword matches more than one nodes in the answer tree, then all these matching nodes must be equivalent.

Search space and complexity

- ❑ Problem related to the **(Group) Steiner Tree Problem**
 - ❑ Given graph G , and nodes n_1, \dots, n_m , the Steiner Tree Problem (STP) requires the smallest tree in G that connects all the nodes. Known NP-hard problem in $|G|$
 - ❑ Group STP: start with m groups of nodes
- ❑ **Differences with our problem:**
 - ❑ Each edge can be taken in **both directions**: exponential increase in search space size
 - ❑ We need the **k smallest-cost trees**, not just one.
 - ❑ Support arbitrary score functions
- ❑ **Large literature (30+) on kwd search in data graphs.**
 - ❑ Differ in search space and/or make limitative assumptions on score
- ❑ **Our approach: enumerate solutions until time-out or max number of solutions reached.**
 - ❑ Return best k solutions found, for given score function

GAM (Grow and Aggressive Merge) Algorithm

- Builds trees exploring around the keyword matches
 - GROW** adds an edge to the root of a tree, **MERGE** merges trees with the same root
- Sample search (tree roots underlined):



GAM (Grow and Aggressive Merge) Algorithm

- ❑ Builds trees exploring around the keyword matches
- ❑ **GROW** adds an edge to the root of a tree, **MERGE** merges trees with the same root

Every solution is a **connection (edge set)**: the root is not meaningful to users

Ideally, we should develop each edge set only once... But:

- ❑ Grow and Merge depend on the tree root → Many edge sets with the same root
- ❑ Also, there are many Grow-Merge combinations that produce the same rooted tree!

Optimizations and careful pruning to keep exploration highly efficient (ongoing)

Sample query answers

The screenshot shows the ConnectionLens web interface. The search bar contains the query "briand tonolli ruffin". Below the search bar, a list of results is displayed:

- Rank: #1 Size: 10 Score: 0.28 2 sources
- Rank: #2 Size: 9 Score: 0.28 1 source
- Rank: #3 Size: 9 Score: 0.28 1 source
- Rank: #4 Size: 9 Score: 0.28 1 source
- Rank: #5 Size: 10 Score: 0.27 1 source
- Rank: #6 Size: 10 Score: 0.26 1 source
- Rank: #7 Size: 11 Score: 0.25 1 source
- Rank: #8 Size: 10 Score: 0.23 1 source


To the right of the results is a network graph showing connections between nodes. The nodes include:

- Julie[...] Briand
- Mme Julie Briand
- collaborateurs
- M. Angelo Tonolli
- Angelo[...] Tonolli
- Francois[...] ruffin

ConnectionLens architecture and performance

Implementation

- ❑ Java (230 classes/46K LOC), Python (24 classes/2800 LOC), JS + CSS
- ❑ Available online: <https://gitlab.inria.fr/cedar/connectionlens>



Implementation [InfSys2022]

- ❑ Graph creation time mostly **linear in the size of the data**
- ❑ Costliest operations involve ML (disambiguation, extraction)
 - ❑ **Batch extraction:** 20x speed-up on GPU, 2x speed-up on regular server
 - ❑ **Extraction policies** replace or avoid extraction in some parts of the data

Graph creation performance: storage, extraction, disambiguation [InfSys2022]




Figure 6: Graph construction time (seconds).

Graph creation performance: batch extraction [InfSys2022]




Figure 7: YAGO loading time (minutes) using Flair.

Application: conflicts of interest in the biomedical domain [IEEE DataEngBull 2021, CIKM2021]

Collaboration with Stéphane Horel (Le Monde)

Data: XML, PDF→JSON, HTML

$ N $	$ E $	$ N $	$ N_P $	$ N_O $	$ N_L $
XML	32,028,429	19,851,904	1,483,631	584,734	126,629
JSON	1,025,307	432,303	75,297	7,320	4,139
HTML	246,636	185,479	3,726	7,227	320
Total	33,300,372	20,469,686	1,562,654	665,167	131,088

Table 3: Statistics on Conflict of Interest application graph.

Application: conflicts of interest in the biomedical domain [IEEE DataEngBull 2021, CIKM2021]

Collaboration with Stéphane Horel (Le Monde)

The screenshot shows the CL-LinkingCOIS web application. The browser address bar displays the URL: <https://cl-linkingcois.saclay.inria.fr/?keyword=org%3ANovartis>. The navigation menu includes Home, Dashboard, Settings, Help, and About. The user is logged in as ioana.manolescu@inria.fr. The main heading is "CL-LinkingCOIS". Below it is a search bar containing "org:Novartis" and a "Search" button. A green badge indicates "315 results".

CoiStatement	PubmedLink
Competing interests : Richard L. Baretto reports grants , personal fees and honoraria for lectures from ThermoFisher , Novartis and ALK Abello outside the submitted work. Mamidipudi Thirumala Krishna received honoraria for lectures from Thermo Fisher and ALK Abello , outside the submitted work.	view the pubmed paper
Conflict of interest : Benjamin Waschki has nothing to disclose. Conflict of interest : Christian Herr has nothing to disclose. Conflict of interest : Christina Magnusser has nothing to disclose. Conflict of interest : Christoph Sinning has nothing to disclose. Conflict of interest : Claus F. Vogelmeier reports grants and personal fees from AstraZeneca , Boehringer Ingelheim , GlaxoSmithKline , Grifols and Novartis , personal fees from CSL Behring , Chiesi , Menarini , Mundipharma , Teva and Cipla , grants from Bayer-Schering , MSD and Pfizer , outside the submitted work. Conflict of interest : Henrik Watz reports personal fees from AstraZeneca , Boehringer Ingelheim , GlaxoSmithKline , BerlinChemie , Chiesi , Novartis and Roche , outside the submitted work. Conflict of interest : Johannes T. Neumann reports personal fees from Abbott Diagnostics and Siemens , outside the submitted work. Conflict of	view the pubmed paper

Graph quality experiments [InfSys2022]

- ❑ PDF extraction accuracy: 63%
- ❑ F1 score for entity extraction from French:
 - ❑ Flair stacked forward and backward embeddings with French fastText embeddings: 73%
 - ❑ Spacy: 63%
 - ❑ StanfordNER: 45%
- ❑ F1 score of disambiguation: 86%

ConnectionLens in the scientific landscape

ConnectionLens in the scientific landscape (1)

Data integration for structured, semistructured and unstructured data

- ❑ “Ad-hoc” (combinations of sources to be unioned, joined, or chained)
- ❑ No schema, ontologies, queries known in advance
- ❑ Reachability queries instead of join: price to pay for powerful integration
- ❑ Comparable work by R. S. Roy, G. Weikum: GST querying of RDF graph enriched on-the-fly with content extracted from Web sources
 - ❑ Web info may be missing or not trusted

Data cleaning: Similarity links require value or entity matching

- ❑ Use reference data sources, heuristics, interactive cleaning (under integrated)

ConnectionLens in the scientific landscape (2)

Graph construction

- ❑ Users of **entity extraction** modules, trained a model for French

Keyword search on structured data

- ❑ Intensely studied for relational, graph, or XML databases: DISCOVER, Banks & Banks2, DPBF, ObjectRank, BLINKS, Spark, ... QGPT (WWW 2021)
 - ❑ Some assume regularity in the graph and translate to SQL or SPARQL
 - ❑ Consider more limited search spaces (edge direction; one tree per root; ...)
 - ❑ Exploit favorable properties of fixed score function; establish approximations...
- ❑ We focus on efficient keyword search for arbitrary graphs (up to time-out)
w/o assumptions on score, w/o sub-optimal structure prop., w/ bidirectional search
- ❑ In-memory graph store and parallel query processor (200x speed-up)

Ongoing work

- ❑ Learning to predict extraction success ~ detecting disguised nulls (BDA 2021, invited TLDKS 2022)
- ❑ Integrating keyword search into structured graph querying (with A. Anadiotis, M. Mohanty)
- ❑ Relationship extraction based on OpenIE (with O. Balalau, P. Upadhyay)
- ❑ Natural language graph querying (with O. Balalau, K. Zhang)
- ❑ Improving the quality of graph linking (with T. Bouganim, H. Galhardas)
- ❑ Abstracting CL graphs (with N. Barret, P. Upadhyay) <https://team.inria.fr/cedar/projects/abstra/>
- ❑ Applications:
 - ❑ Conflicts of Interests in the biomedical domain (w/ S. Horel and G. Fooks, Aston U., UK)
 - ❑ France's business x political sphere (w/ Radio France Investigation)

Why data journalism?

Because I grew up in a dictatorship, and I value free press

Because journalists are threatened and killed still today in Europe



Daphne Galizia, 1964-2017



Jan Kuciak, 1990-2018

Because the press' economic model is threatened by IT giants

Because this industry is currently underserved by IT – and we could really make an impact!

Thank you

Questions?

ConnectionLens: <https://team.inria.fr/cedar/connectionlens/>