


# Wrapping Millions of Documents Per Day

and How that's Just the Beginning











# Background: About me



Lecturer in Databases and Co-investigator VADA  
Fellow of Oxford Martin School & OMI



Co-founder and CTO



Senior Research Manager  
Meltwater

Mission





5

The image shows a vertical list of four real estate listings. Each listing includes a photo, a title, location, price, and description. Annotations with arrows point to specific elements:

- Records:** Points to the price (\$372,000) and the monthly payment (\$1,742/mo) in the first listing.
- Data Areas:** Points to the price (\$1,404,995) in the second listing.
- Fields:** Points to the '3,377+ sqft' and 'Preserve at Ardsley is a community of 11 single-family ... More' text in the second listing.
- Descriptions:** Points to the 'LAKELAND SCHOOLS with LOW TAXES...' text in the third listing.

**Listing 1: 2482 Bound Brook Ln**  
Yorktown Heights, NY 10598 | Map | \$372,000 | \$1,742/mo Get Pre-Approved  
3 bd | 1 ba | 1,741 sqft / 0.5 acres | Single-Family Home  
Terrific split level. Fabulous great room addition with a wall of built-ins, vaulted ceiling, beams, skylights, e ...  
Houlihan Lawrence Somers Brokerage | View Details

**Listing 2: Preserve at Ardsley**  
Scarsdale, NY 10583 | Map | From \$1,404,995  
4 Bd | 3,377+ sqft | 3.5 Ba | New Community  
Preserve at Ardsley is a community of 11 single-family ... More  
Toll Brothers

**Listing 3: 3468 Carol Ct**  
Yorktown Heights, NY 10598 | Map | \$379,000 | \$1,774/mo Get Pre-Approved  
3 bd | 2 ba | 1,602 sqft / 0.38 acres | Single-Family Home  
LAKELAND SCHOOLS with LOW TAXES. Don't let this one slip away. Nestled away on a quiet cul-de-sac, backing u ...  
Houlihan Lawrence Somers Brokerage | View Details

**Listing 4: 2906 Hickory St**  
Yorktown Heights, NY 10598 | Map | \$329,000 | \$1,540/mo Get Pre-Approved  
3 bd | 2 ba | 1,290 sqft / 0.26 acres | Single-Family Home  
Incredible updates, since last on the market. Renovated kitchen with luxurious counters and back splash, renovate ...  
Houlihan Lawrence Yorktown Brokerage | View Details





## 25 LOCATIONS FOUND NEAR YOU



### BEACH HAVEN II

600 N. Bay Avenue,  
Beach Haven, NJ 08008  
609-492-6300

Send Address To [E-Mail](#)[Mobile](#)

8.22 miles

[Get Driving Directions](#)

Gift Cards Sold, Gift Cards Accepted



### BORGATA

1 Borgata Way,  
Atlantic City, NJ 08401  
609-317-8206

Send Address To [E-Mail](#)[Mobile](#)

18.04 miles

[Get Driving Directions](#)

[Scoop Website](#)

Catering, Gift Cards Sold, Gift Cards Accepted



### 30TH STREET STATION

30TH STREET STATION, Amtrak - 30th  
Street Station  
Philadelphia, PA 19104  
215-222-2996

Send Address To [E-Mail](#)[Mobile](#)

48.70 miles

[Get Driving Directions](#)

[Scoop Website](#)

Address or Postcode

Bridge Rd, Little Egg Harbor Township,

United States of America

SEARCH

Offer catering  Offer ice cream cakes

© OpenStreetMap contributors



Mission



## 25 LOCATIONS FOUND NEAR YOU



### BEACH HAVEN II

600 N. Bay Avenue,  
Beach Haven, NJ 08008  
609-492-6300

Send Address To [E-Mail](#)[Mobile](#)

8.22 miles

[Get Driving Directions](#)

Gift Cards Sold, Gift Cards Accepted



### BORGATA

1 Borgata Way,  
Atlantic City, NJ 08401  
609-317-8206

Send Address To [E-Mail](#)[Mobile](#)

18.04 miles

[Get Driving Directions](#)

[Scoop Website](#)

Catering, Gift Cards Sold, Gift Cards

Address or Postcode

Bridge Rd, Little Egg Harbor Township,

United States of America

SEARCH

Offer catering  Offer ice cream cakes

© OpenStreetMap contributors



NAME	STREET ADDRESS	LOCALITY	STATE	POSTCODE	PHONE
BEACH HAVEN II	600 N. Bay Avenue	Beach Haven	NJ	08008	609-492-6300
BORGATA	1 Borgata Way	Atlantic city	NJ	08401	609-317-8206
30TH STREET STATION	30TH STREET STATION	Philadelphia	PA	19104	215-222-2996

January 21, 2017 9:36 PM

DATE

TIME



White House press secretary Sean Spicer speaks in the media briefing room in Washington, D.C., on Saturday. (Olivier Douliery/Abaca Press/TNS)

## White House vows to fight media 'tooth and nail' over Trump coverage; says it presented 'alternative facts'

By [Doina Chiacu](#) and [Jason Lange](#)  
Reuters

WASHINGTON — The White House vowed on Sunday to fight the news media “tooth and nail” over what it sees as unfair attacks, with a top adviser saying the Trump administration had presented “alternative facts” to counter low inauguration crowd estimates.

On his first full day as president, Trump said he had a “running war”

# change

**Exclusive:** The five-yearly assessment of what will happen to the UK as the world warms says one of an array of potential threats is the 'significant risk' to supplies of food

by [Ian Johnston](#), [Tom Batchelor](#) Environment Correspondent | 8 hours ago | [230 comments](#)

Click to follow  
The Independent Online



your visit to our site and to bring you advertisements that might interest you. Read our [Privacy](#) and [Cookie](#) Policies to find out more.



In the report, the Government said high-risk issues that needed to be addressed included the damage expected to be caused by flooding  
Getty

The Government has been accused of trying to bury a major report about the potential dangers of global warming to Britain - including the doubling of the deaths during heatwaves, a “significant risk” to supplies of food and the prospect of infrastructure damage from flooding.

**BODY INCLUDE**

The UK Climate Change Risk Assessment Report, which by law has to be produced every five years, was published

**indy100** TRENDING



Kellyanne Conway calls misleading errors 'alternative facts'. Is ripped apart by Twitter

The CIA aren't happy about Donald Trump's first speech at their

Itwater

DEMO GODS



HELP US PLEASE

memegenerator.net



[https://youtu.be/j\\_0IZdNJ-aw](https://youtu.be/j_0IZdNJ-aw)

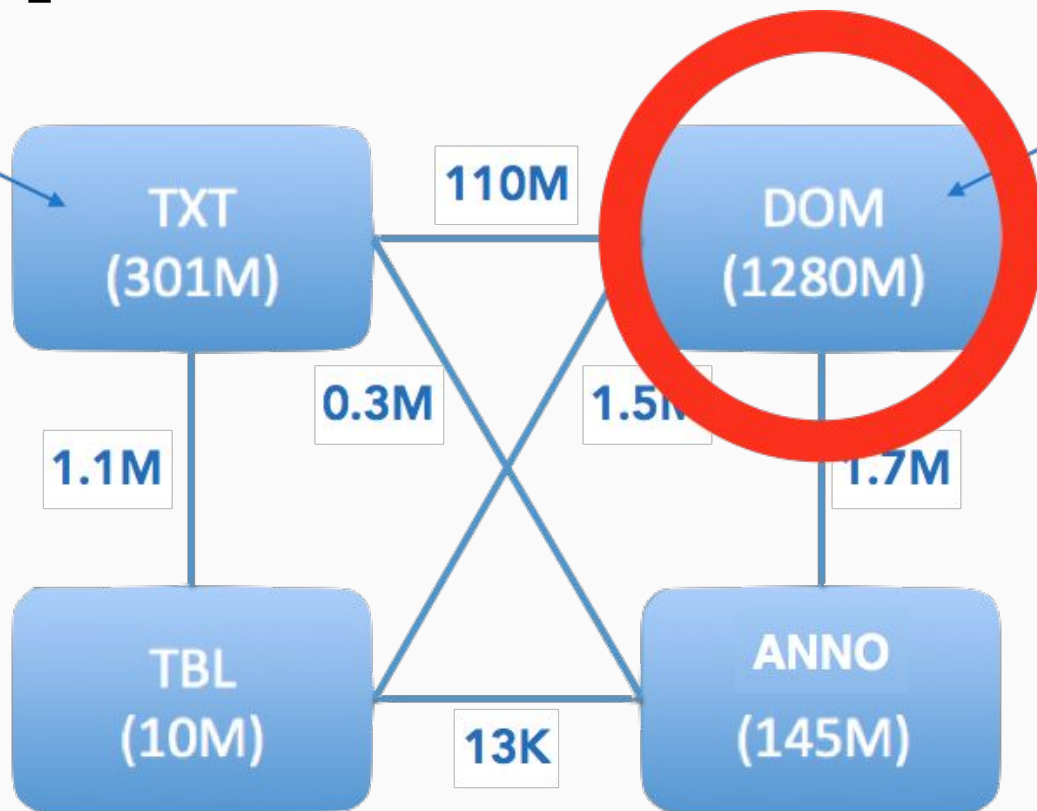
World class **crawling** platform  
to largely automate  
**outside content collection**

Mission



# Crawling Space & Volume

Information  
Extraction



Data  
Extraction

Mission

Source: Xin Luna Dong (Google) - PVLDB '14



# Crawling Coverage

For many kinds of information  
one has to extract from  
**thousands of sites**  
in order to build a  
**comprehensive** database

Mission





Vertical 1 & 2:

# Real Estate & Used Cars, UK



Results



**Records**

**Data Areas**

**Fields**

**Descriptions**

**FEATURED**  
**2482 Bound Brook Ln** **\$372,000**  
Yorktown Heights, NY 10598 | Map \$1,742/mo Get Pre-Approved  
3 bd 1 ba 1,741 sqft / 0.5 acres  
Single-Family Home  
Terrific split level. Fabulous great room addition with a wall of built-ins, vaulted ceiling, beams, skylights, e ...  
Houlihan Lawrence Somers Brokerage  
View Details

**NEW HOMES**  
**Preserve at Ardsley** From **\$1,404,995**  
Scarsdale, NY 10583  
4 Bd 3,377+ sqft  
3.5 Ba New Community  
Preserve at Ardsley is a community of 11 single-family ... More  
Toll Brothers

**FEATURED**  
**3468 Carol Ct** **\$379,000**  
Yorktown Heights, NY 10598 | Map \$1,774/mo Get Pre-Approved  
3 bd 2 ba 1,602 sqft / 0.38 acres  
Single-Family Home  
LAKELAND SCHOOLS with LOW TAXES. Don't let this one slip away. Nestled away on a quiet cul-de-sac, backing u ...  
Houlihan Lawrence Somers Brokerage  
View Details

**FEATURED**  
**2906 Hickory St** **\$329,000**  
Yorktown Heights, NY 10598 | Map \$1,540/mo Get Pre-Approved  
3 bd 2 ba 1,290 sqft / 0.26 acres  
Single-Family Home  
Incredible updates, since last on the market. Renovated kitchen with luxurious counters and back splash, renovate ...  
Houlihan Lawrence Yorktown Brokerage  
View Details





# 10,493

**Sites** from real-estate  
and used-car



# 92%

**Effective wrappers**  
(where we get all the data)



# 96%

**Precision** of extracted  
primary attributes



# 20

**Days** (one expert) to adjust  
system to a new domain



	wrapper		
	effective	wrong or missing data	no data
UK real estate	<b>91%</b>	7%	2%
Oxford real estate	<b>90%</b>	6%	4%
ViNTs <sup>10</sup>	<b>4%</b>	5%	91%
UK used cars	<b>93%</b>	4%	3%
US real estate	<b>90%</b>	5%	5%

```

doc('http://www.wvagency.com/')//label[@for='sale_type_id']/following-sibling::select/{0 /}
  //form/div[@class='formbtn-ctn'] [last()]/button[@class='formbtn']/{click /}
  /.:<data_area>[?./div[@class='pagenumlinks'] [1]//span/text():<number_results=.>]
  /((div[contains(@class, 'proplist_wrap')]/following-sibling::div//a[@class='pagenum'] [last()]{nextclick /})*
  //div[contains(@class, 'proplist_wrap')]:<record>[? .:<origin_url=current-url>]
  [? ./span[@class='prop_price']/text():<price=normalize-space(.)> ]
  [? ./span[.='Type:']/following-sibling::strong/text():<property_type=normalize-space(.)> ]
  [? ./div[@class='prop_statuses']//text():<property_status=normalize-space(.)> ]
  [? ./span[.='Bathrooms:']/following-sibling::strong/text():<bathroom_number=normalize-space(.)> ]
  [? ./span[.='Bedrooms:']/following-sibling::strong/text():<bedroom_number=normalize-space(.)> ]
  [? ./strong[@class='orange']/preceding-sibling::text():<location_raw=string(.)> ]
  [? ./strong[@class='orange']/text():<postcode=normalize-space(.)> ]
  [? ./strong/preceding-sibling::strong/text():<street_address=normalize-space(.)> ]
  [? ./@src:<image=normalize-space(.)> ]
  [? ./div[@class='prop_statuses']/following-sibling::a/@href:<url=normalize-space(.)> ]
  [? ./div[@class='prop_maininfo']:<description=normalize-space(.)> ]
  
```

Vertical 3:

# News Articles



January 21, 2017 9:36 PM

DATE

TIME



White House press secretary Sean Spicer speaks in the media briefing room in Washington, D.C., on Saturday. (Olivier Douliery/Abaca Press/TNS)

## White House vows to fight media 'tooth and nail' over Trump coverage; says it presented 'alternative facts'

By [Doina Chiacu](#) and [Jason Lange](#)  
Reuters

WASHINGTON — The White House vowed on Sunday to fight the news media “tooth and nail” over what it sees as unfair attacks, with a top adviser saying the Trump administration had presented “alternative facts” to counter low inauguration crowd estimates.

On his first full day as president, Trump said he had a “running war”

# change

**Exclusive:** The five-yearly assessment of what will happen to the UK as the world warms says one of an array of potential threats is the 'significant risk' to supplies of food

by [Ian Johnston](#), [Tom Batchelor](#) Environment Correspondent | 8 hours ago | [230 comments](#)

Click to follow  
The Independent Online



your visit to our site and to bring you advertisements that might interest you. Read our [Privacy](#) and [Cookie](#) Policies to find out more.



In the report, the Government said high-risk issues that needed to be addressed included the damage expected to be caused by flooding  
Getty

The Government has been accused of trying to bury a major report about the potential dangers of global warming to Britain - including the doubling of the deaths during heatwaves, a “significant risk” to supplies of food and the prospect of infrastructure damage from flooding.

The UK Climate Change Risk Assessment Report, which by law has to be produced every five years, was published

**indy100** TRENDING



Kellyanne Conway calls misleading errors 'alternative facts'. Is ripped apart by Twitter


The CIA aren't happy about Donald Trump's first speech at their





**~40,000**

**US/UK news sources** from  
Meltwater media intelligence



**85%**

**Effective wrappers**  
(where we get all the data)



**89%**

**Precision of extracted**  
ADICT+ attributes



**90**

**Days to adjust system to this**  
vertical (why? 30min refresh)





```
"site":{  
  "name":"The Guardian (U.S. edition)",  
  "url":"http://www.theguardian.com/us",  
  "socialHandlers":{  
    "twitter":"@guardian"  
  }  
},  
"startUrls":[  
  "http://www.theguardian.com/us"  
],  
"sectionTpls":[ ],  
"articleTpls":{  
  "urlPatterns":[  
    "(?<wordset>([a-zA-Z]{1},[:\\-\\.]{1,}){1,5}[a-zA-Z]{1,})/(?<numberset>([0-9]{1,}){1,1})/(?<wordset1>([a-zA-Z]{1,}){1,1})/(?<numberset1>([0-9]{1,}){1,1})"  
  ],  
  "canonicalUrlXPath":"wrtly:normalize-url(//link[@rel='canonical']/@href)",  
  "titleXPath":"wrtly:normalize-space(//h1[contains(@class,'content__headline')])",  
  "alternativeUrls":[ ],  
  "bylineXPath":"//span[@itemprop='name']",  
  "datePublished":{  
    "datelineXPath":"//meta[@property='article:published_time']/@content"  
  },  
  "ingressXPath":"wrtly:normalize-space(//div[@class='hide-on-mobile']/p[1])",  
  "keywordsXPath":"//head//@content",  
  "contentXPath":{  
    "includeXPath":"wrtly:normalize-space(wrtly:string-join(//div[@class='content__article-body from-content-api js-article__body']/node())[self::div])"  
  },  
  "engagementPatterns":{  
    "valueXPath":"//span[@class='commentcount2__value tone-colour js_commentcount_actualvalue']",  
    "type":"comments"  
  },  
  "commentCount":{  
    "valueXPath":"//span[@class='commentcount2__value tone-colour js_commentcount_actualvalue']"  
  }  
}
```

# Why different wrapper format?

- **XPath**: perfect for **interactive, search engine style** websites
- However: in **media intelligence** – **freshness of data** is critical
  - **30min** maximum between publishing and indexing time
  - (almost) every article has an **indexable, unique URL**
  - large **variety of different article** templates
- **Decompose XPath wrapper into single page segments**
  - memorise set of **section pages** encountered in a run
  - recrawl stored section pages in next run
    - to find new article (& section) pages

Vertical 4:

# Company Extractors



Results





PERSON\_IMAGE  
**John Lyseggen**  
 RECORD\_VALUE AGE  
 Director, Founder  
 PERSON\_ROLE PERSON\_ROLE



PERSON\_IMAGE  
**Martin Hernandez**  
 RECORD\_VALUE AGE  
 Director, Officer  
 PERSON\_ROLE




PERSON\_IMAGE  
**Kaveh Rostampour**  
 RECORD\_VALUE AGE  
 Director, Americas  
 PERSON\_ROLE



PERSON\_IMAGE  
**Paal Larsen**  
 RECORD\_VALUE AGE  
 Director, EMEA  
 PERSON\_ROLE



PERSON\_IMAGE  
**John Box**  
 RECORD\_VALUE AGE  
 Director, APAC  
 PERSON\_ROLE



PERSON\_IMAGE  
**Niklas de Besche**  
 RECORD\_VALUE AGE  
 Director, Products  
 PERSON\_ROLE

Results





### Ann Arbor

2300 Traverwood Dr  
Ann Arbor, MI 48106  
United States

Phone: +1 734-332-6500

[Directions](#)



### Atlanta

1010 10th Street NE  
Atlanta, GA 30309  
United States

Phone: +1 404-487-9000

[Directions](#)



### Austin

9606 North MoPac Expressway  
Austin, TX 78759  
United States

Phone: +1 512-343-5283

[Directions](#)



### Birmingham

325 North Old Woodward  
Birmingham, AL 35203  
United States

[Directions](#)



### Boulder

2600 Pearl Street  
Boulder, CO 80302  
United States

Phone: +1 303-245-0080

[Directions](#)



### Cambridge

355 Main Street  
Cambridge, MA 02149  
United States

Phone: +1 617-575-1300

[Directions](#)



### Chapel Hill

200 West Franklin Street



### Chicago

320 N. Morgan, Suite 600

Results



# Company Extraction: Goals

- Given **only** a company website
  - Extract as much relevant information from **structured sources**
    - executive team, locations, subsidiaries, ...
  - Identify unstructured sources
    - press releases, financial reports, ...
- **Scale** to millions of companies in multiple languages

# Company Extraction: Results

		% of companies			
recall	CEO & execs	description	location	logo	
	<b>85%</b>	<b>74%</b>	<b>82%</b>	<b>89%</b>	
present	<b>58%</b>	<b>79%</b>	<b>69%</b>	<b>74%</b>	



THIS SEEMS LIKE  
**BULLSHIT**



# Restaurant locations

Don't believe us? You aren't the first – major US technology company

- **Need:** US restaurant locations (including chains) for check-ins
- **Problem:** existing location databases incomplete and full of errors
- **Want:** Get that data from the “authoritative” source, i.e.,
  - the restaurant (chain) websites

• They evaluated state-of-the-art – most solutions to crude

- Settled on scrapy, but: 2 months for top 20 US chains
- Very worried about maintenance



## 25 LOCATIONS FOUND NEAR YOU



### BEACH HAVEN II

600 N. Bay Avenue,  
Beach Haven, NJ 08008  
609-492-6300

Send Address To [E-Mail](#)[Mobile](#)

8.22 miles

[Get Driving Directions](#)

Gift Cards Sold, Gift Cards Accepted



### BORGATA

1 Borgata Way,  
Atlantic City, NJ 08401  
609-317-8206

Send Address To [E-Mail](#)[Mobile](#)

18.04 miles

[Get Driving Directions](#)

[Scoop Website](#)

Catering, Gift Cards Sold, Gift Cards

Address or Postcode

Bridge Rd, Little Egg Harbor Township,

United States of America

SEARCH

Offer catering  Offer ice cream cakes

© OpenStreetMap contributors



NAME	STREET ADDRESS	LOCALITY	STATE	POSTCODE	PHONE
BEACH HAVEN II	600 N. Bay Avenue	Beach Haven	NJ	08008	609-492-6300
BORGATA	1 Borgata Way	Atlantic city	NJ	08401	609-317-8206
30TH STREET STATION	30TH STREET STATION	Philadelphia	PA	19104	215-222-2996

# Restaurant location: Results

- After **1 month** applying Wrapidity technology:
  - over 300 US chains, over **100k** websites
  - more than **3M** locations in total
  - fully automated maintenance for those sources

coverage

**85%**

precision

**95%**

- But: *they still didn't believe*
  - hired Accenture to assess quality of the data
  - result: **over 97% precision**

Results



# Restaurant location: Independent Evaluation

	<b>Present</b>	and <b>correct</b> data & extraction	but <b>wrong</b> extraction	but <b>wrong</b> <b>data</b>	but <b>raters</b> <b>disagree</b>
<b>city</b>	100%	<b>99.3%</b>	<b>0.7%</b>	0.0%	0.0%
<b>street</b>	100%	<b>96.4%</b>	<b>1.7%</b>	1.9%	0.0%
<b>postcode</b>	99%	<b>97.1%</b>	<b>0.1%</b>	0.0%	2.8%
<b>latlong</b>	89%	<b>99.7%</b>	<b>0.1%</b>	0.0%	0.1%
<b>hours</b>	47%	<b>98.2%</b>	<b>0.0%</b>	1.3%	0.5%
<b>name</b>	100%	<b>99.5%</b>	<b>0.5%</b>	0.0%	0.0%
<b>phone</b>	86%	<b>98.7%</b>	<b>1.3%</b>	0.0%	0.0%
<b>category</b>	100%	<b>98.9%</b>	<b>0.0%</b>	0.0%	1.1%
	90%	<b>98.5%</b>	<b>0.5%</b>	0.4%	0.6%

This evaluation is done by independent, external evaluators on a sample of 1000 locations.



# Summary



~50k sources, 30 min  
recrawl interval

87%

90%

over 300 US chains, over  
100k websites;  
more than 3M locations



over 30 attributes, where  
present; 1M+ company's site  
crawl ongoing

91%



# Approach



# Web-Scale Wrapper Induction

- We need to **scale to the web**
  - minimize supervision per source
- But: we can afford **prior knowledge**
  - about entities and attributes
  - mostly in form of known **labels & instances** and “**appearance**”
    - expressed as Gazetteers or rules for local, textual information
    - higher-level rules or classifiers for complex structures

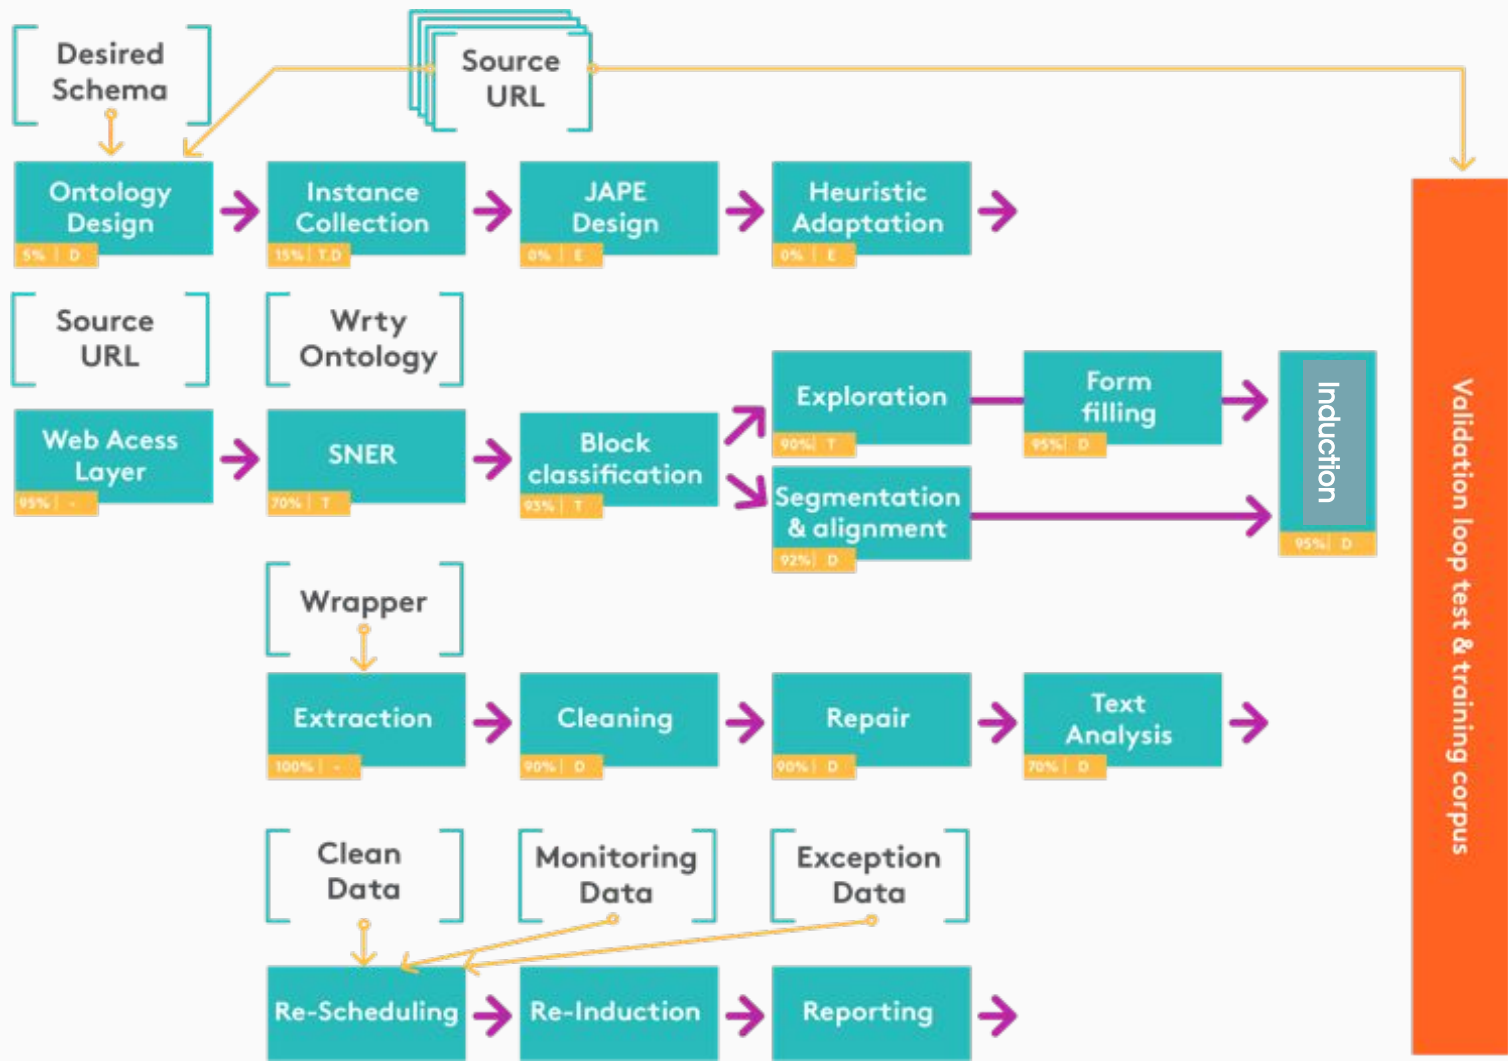
# Web-Scale Wrapper Induction

- **Problem:** application of prior knowledge is costly & noisy
  - wrapper induction to generalise to other pages of site
  - “template” hypothesis
- **Solution:** Generate “wrapper” program from examples
  - then apply to all pages of a site
  - when to apply which extractor
- **Full site extraction** needs to also deal with
  - Interactivity such as pagination & form filling (deep web)
  - Detecting complex structures such as lists, tables, ...



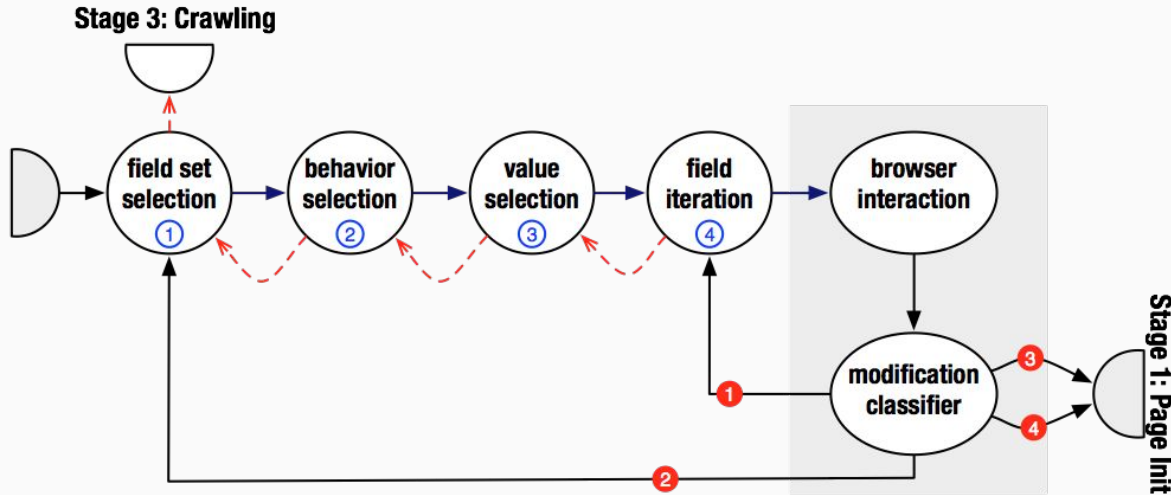
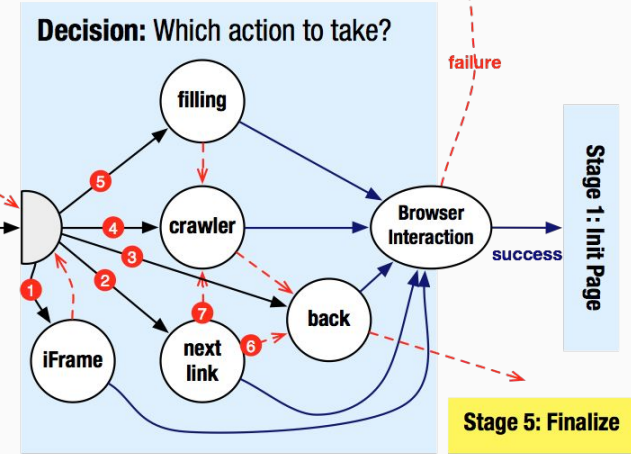


# Approach



# Exploration: Self-Adaptive

- Self-adaptive, dynamic exploration plans
  - planners expressed as guarded FSTs
  - with Datalog rules as guards
  - 1000's of unique exploration plans



Results

# NER for DOMs: Labels, structure, ...



Labels and instances, visible and invisible (HTML structure, Javascript values)

Running costs

Urban mpg	38.7 mpg
Extra Urban mpg	58.9 mpg
Average mpg	49.6 mpg
CO <sub>2</sub> emissions	136 g/km
Annual Tax	£130

Labels: Urban mpg, Extra Urban mpg, Average mpg, CO<sub>2</sub> emissions, Annual Tax, 38.7 mpg, 58.9 mpg, 49.6 mpg, 136 g/km, £130

Instances: 38.7 mpg, 58.9 mpg, 49.6 mpg, 136 g/km, £130

Trulia. Your home for real estate.

Buy Search by Neighborhood, City, or Address Search

Labels: Buy, Search by, Neighborhood, City, or Address, Search

Javascript values: Neighborhood, City, or Address

Crick Road, Norham Manor

Guide Price **£6,000,000**

A substantial detached family home built over four floors located in one of the most desirable address... [Read more](#)

8 Bedrooms, 4 Bathrooms, 5 Living Rooms, 5,683 ft<sup>2</sup>

+ Save property View details

```
<div class="icon first">
  
  <br>8
</div>
<div class="icon">
  
  <br>4
</div>
```

Approach



# Form understanding



- Sub-problems: form **labeling**, form **segmentation**, **classification**
- **Combines** structural, textual, visual, and semantic clues
  - **structural** = structure of the DOM, e.g., distance
  - **visual** = rendering of the form, e.g., for alignment
  - **textual** = detectors for a vertical's types (e.g., "LHR")
  - **semantic** = class, id, ... with semantic labels (e.g. "finput\_dest")
- **Polynomial time** labeling, grouping, and classification algorithm

ICQ dataset	HA [14]	ExQ [41]	StatParser [36]	DIADEM [17]
<b>F<sub>1</sub> for labeling</b>	92%	96%	96%	98%



Approach

# Pick a Path: Wrapper induction

- Pick robust, “semantic” paths
  - less affected by changes
  - over time and within a template

- Suitable as “foundation” for

- **template discovery**

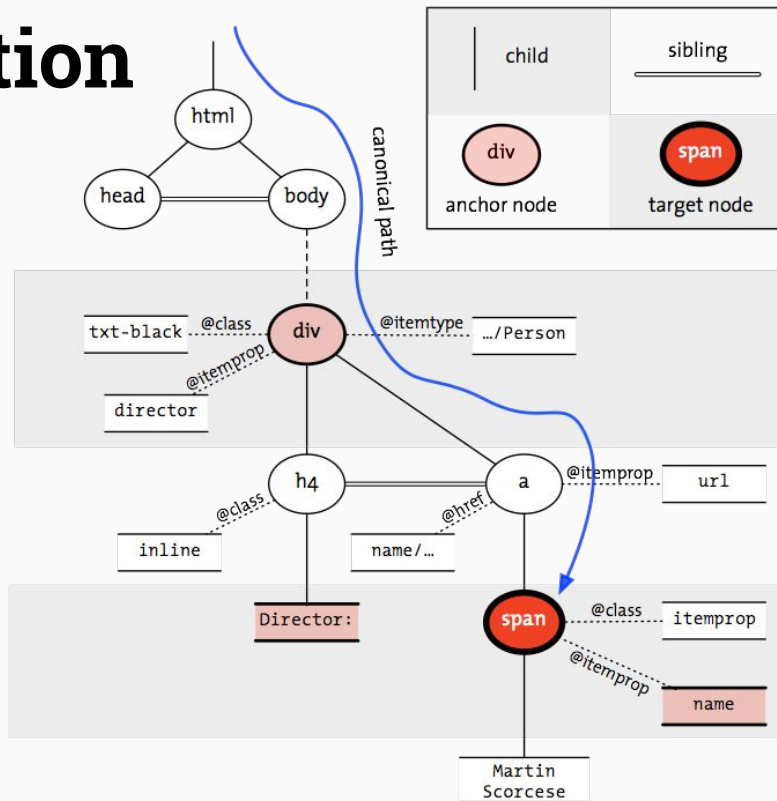
- E.g.: Select the director

- Firebug (“canonical” XPath)

```
/html[1]/body[1]/ ... /div[4]/a[1]/span[1]
```

- Ours:

```
//div[starts-with(., "Director:")]//span[(@class="itemprop" )
```



# Wrapper Repair

ICDE'16


Postcode	Phone	Locality	State	Street Address
San Diego, CA 92101	619-234-1802	San Diego, CA 92101	<NULL>	471 Horton Plaza, near Westland park
Boise, ID 83702	208-342-1992	Boise, ID 83702	<NULL>	103 North 10th Street
Portland, OR 97209	503-796-3033	Portland, OR 97209	<NULL>	301 NW 10th Avenue, near the Fish Market

Postcode	Phone	Locality	State	Street Address
92101	619-234-1802	San Diego	CA	471 Horton Plaza
83702	208-342-1992	Boise	ID	103 North 10th Street
97209	503-796-3033	Portland	OR	301 NW 10th Avenue

San Diego, CA 92101

Boise, ID 83702

Portland, OR 97209



■ Locality

■ State

■ Postcode

■ Phone



# Wrapper Repair

- **Joint repair** of wrapper and output data (relation)
- Problem related to **table segmentation problem**
  - generally NP-complete
  - but we show that it's **polynomial under atomic misplacement**
- **Atomic misplacement:** attribute value is either
  - entirely misplaced, or
  - its fragments are in adjacent fields

① **ROSeAnn** (VLDB'14)  
Entity extraction from DOMs

② **OPAL** (WWW'12, VLDBJ'13)  
Form understanding & filling

③ **AMBER** (ICWE'11)  
Record identification for lists

④ **OXPath** (VLDB'11, VLDBJ'13)  
Extraction language

⑤ **Robust XPath** (SIGMOD'16)  
Change-resistant wrappers

⑥ **Oxtractor** (Coling'16)  
Attribute extraction

⑦ **WaDaR** (ICDE'16)  
Joint wrapper & relation repair

⑧ **VADA** (EDBT'16)  
Wrangling of extracted data (in progress)



⑥ **DIADEM** (VLDB'14)  
World-first accurate, automatic **full-site extraction system**





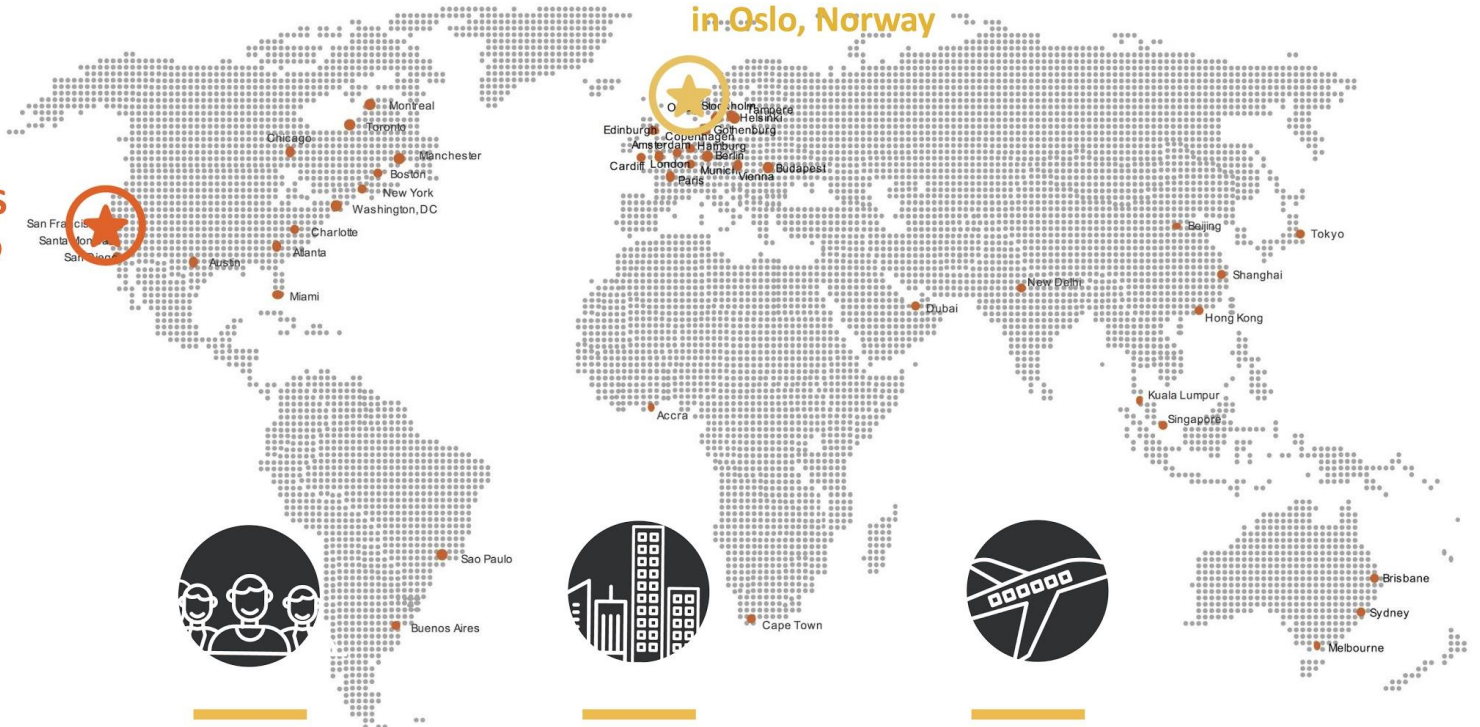




# Meltwater: Who are we?

HEADQUARTERS  
in San Francisco

FOUNDED 2001  
in Oslo, Norway



1500 employees  
worldwide



26,000 Business customers  
in 108 countries



60 offices across  
27 Countries



# Meltwater: Media Intelligence

Sources: Editorial, Social, Broadcasts



media exposure



trends



influencers



sentiment analysis

More than 300k different types of user queries

The dashboard displays a 'Content Stream' with several news items. The first item is from 'iberkisses.com' dated May 26, 5:59 pm, titled 'Fourteen Williams Seniors Win Fellowships for Post-Grad Study at Cambridge'. The second item is from 'BBC World Service Radio' dated May 6, 4:06 am, titled 'BBC World Service Radio at May 6th 2017 11:03 PM'. The third item is from 'BBC Radio 4' dated May 6, 4:05 am, titled 'BBC Radio 4 at May 6th 2017 4:06 AM'. The fourth item is from 'BBC World Service Radio' dated May 6, 4:05 am, titled 'BBC World Service Radio at May 6th 2017 12:05 AM'. The fifth item is from 'BBC World Service Radio' dated Apr 9, 9:15 am, titled 'BBC World Service Radio at April 9th 2017 9:15 AM'. The sixth item is from 'ABC West Victoria AM' dated Apr 5, 3:15 pm, titled 'ABC West Victoria AM at April 5th 2017 4:06 AM'. A 'Media Exposure' chart shows a bar graph with a peak in May 2017. A 'Trending Themes' chart shows a world map with highlighted regions.

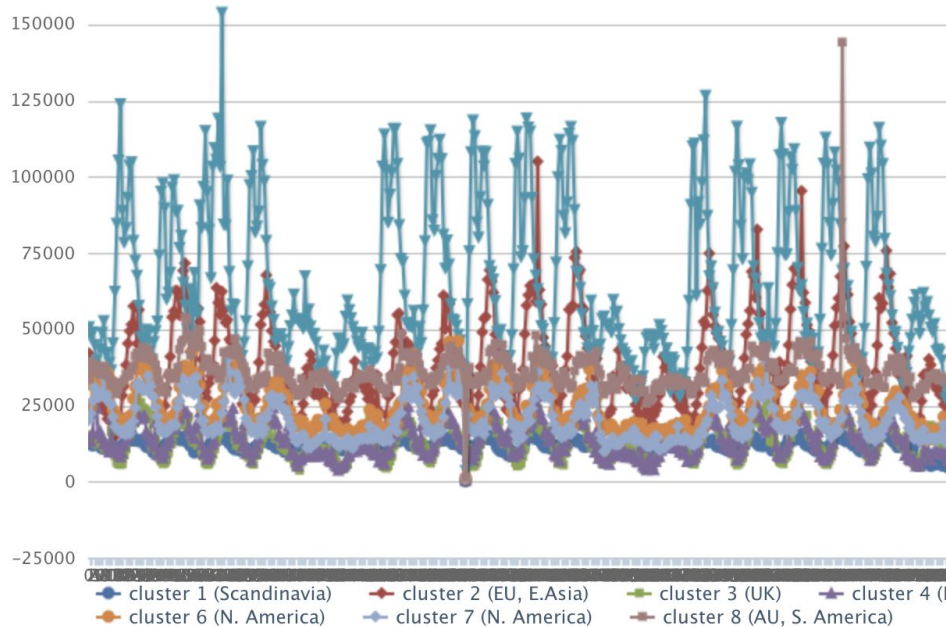
The interface shows a profile for 'Eli Blumenthal Reporter' with contact information: Personal (elblumenthal@usatoday.com), Work (703) 854-3400, 535 Madison Avenue, New York, New York 10022, United States. The 'Boots' section lists 'Digital technology, Mobile communications, Technology (general)'. The 'Recent Results' section shows 'Financial Review (Licensed by Copyright Agency)' dated May 31, 3:01 am, with the content 'Here's what you need to land America's best jobs'. The 'Notes' section is currently set to 'Email'. The 'Home' section features a 'Media Exposure' chart, a 'Sentiment' donut chart showing Positive, Negative, and Neutral sentiment, and a 'Heat Map' of the world. The 'MY INFLUENCERS' section includes a search bar and a list of influencers.

# Meltwater: In Numbers

~200B indexed documents

- Crawlers fetch ~3.3M articles/day from 190k editorial sources
  - re-crawled every 30 minutes
- With the social fire hoses we go up to 30M docs/ day.

Name	Country	Language	Documents
<a href="#">Notiradar</a>	Mexico	Spanish	40700
<a href="#">福建东南新闻网</a>	China	Chinese (simpl.)	23182
<a href="#">中工网</a>	China	Chinese (simpl.)	20953
<a href="#">매일경제</a>	Korea, Republic Of	Korean	20191
<a href="#">جستجوگر اخبار تی نیوز</a>	Iran (islamic Republic Of)	Persian	18055
<a href="#">Match 生活網</a>	Taiwan	Chinese (trad.)	17512
<a href="#">47NEWS</a>	Japan	Japanese	16966
<a href="#">Nambia Press Agency</a>	Namibia	English	9521
<a href="#">中金在线 - 外汇网</a>	China	Chinese (simpl.)	7957
<a href="#">Onet.pl</a>	Poland	English	7127






# Meltwater: Existing Technology Stack

## Ingestion:

- Social media hoses (partnerships)
- Editorial/News (partnerships + web crawling)
- Broadcasts (views on the above)

## Storage and search:

- Elastic search 
- Rabbit MQ (distributed queues) 
- AWS 

## Enrichments (15 languages):

- Text categorization (topic, language)
- NERD (person, location, organization, ...)
- NED ( [https://en.wikipedia.org/wiki/Tim\\_Cook](https://en.wikipedia.org/wiki/Tim_Cook) )
- Sentiment Analysis


## Media Intelligence applications (Custom)

- Boolean queries (keywords / entities)
- Counters
- Aggregates
- Drill downs / pivoting

### Sources




### Hoses + crawlers




### Enrichments



### Elastic search



### Intelligence apps



# Vision: **Insight** Building on **Outside** Data

Build a world class **AI** platform for a new software category

## Outside Insight



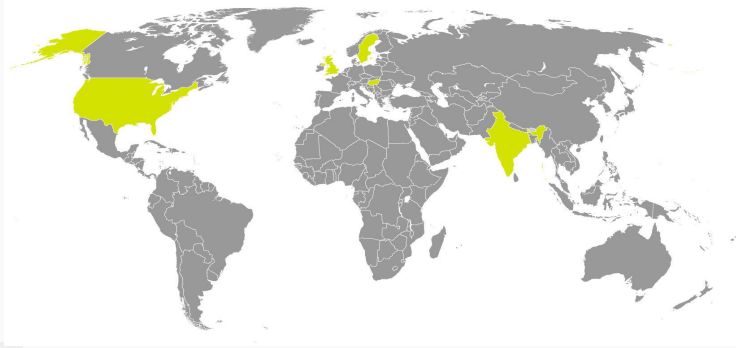
**FAIRHAIR**

Fairhair



# Fairhair: People & Community

## 5 Data Science Research offices



## 6 Data Science Hubs (co-working spaces)

- London
- San Francisco
- Singapore
- Sydney
- Berlin
- New York



## University collaborations



**Stanford  
University**

**Carnegie  
Mellon  
University**



## Meltwater Entrepreneurial School of Technology

- campuses in Ghana and Nigeria
- it's a school for African entrepreneurs
- it's an incubator (33 startups)
- it's a networking hub



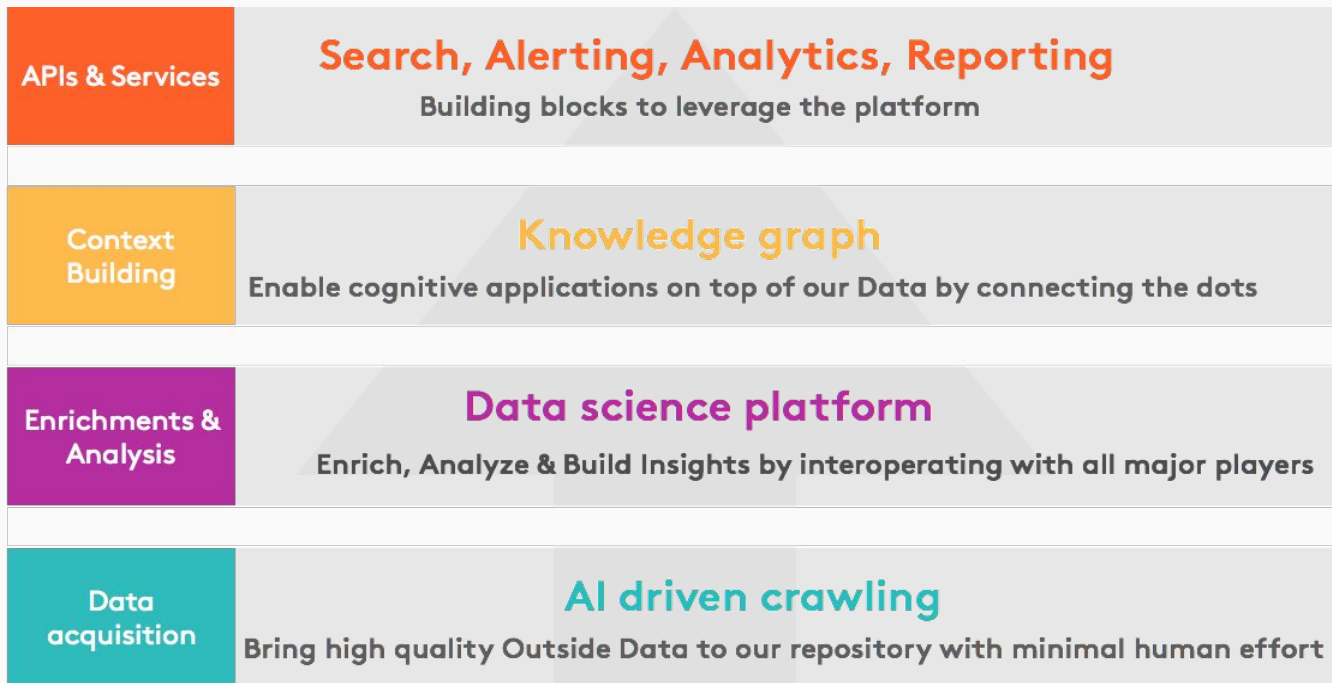


# Fairhair's AI First Approach

Step 1: **Outside Data** acquisition & making it available in a form that's crunchable.

Step 2: Make **Data Science** (data, algorithms, infrastructure, tools) power everything


Step 3: We can't foresee all uses of data and insights → **Developer APIs & integrations**



# Fairhair's AI Crawlers


Traditional scraping requires a **huge human effort**:

- **Code** wrappers for each source, e.g., in **Scrapy** or MW's source configurations
- **Visually** testing and support tool (ala **Connotate**, **Mozenda**, ...)
- **Automatic** scraping for small number of fixed data types (ala **Diffbot**), e.g., Microdata
- Meltwater (old): **~50 "source engineers"** maintaining manual wrappers
  - sources failing at a rate of **100's** per week, **1-2h** to fix each source effectively



80-90% **lower**  
**human effort**

without loss in quality  
compared with  
state-of-the-art



3-10x more  
**attributes**

and domains than existing  
automated solutions and  
affordable supervised one



10-100x  
**more sources**


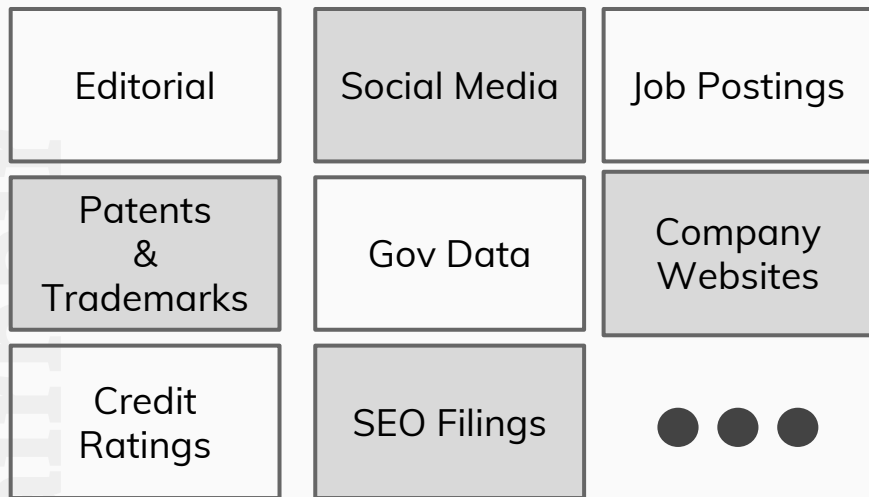
e.g., 100k+ restaurant websites,  
300k+ news sources, 1M+ of  
company websites



# Data and Content Lake

Factual information: **wherefrom?**

Need to **restrict the domain**: focus on the **corporate domain**, i.e., companies, people, products, ...

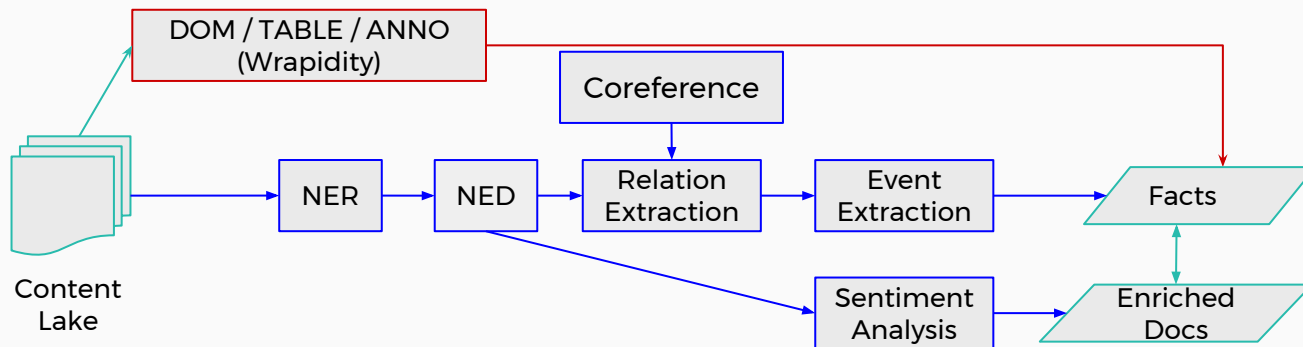


EC2 cluster with 2.8k vCPU, 21TB RAM, 630TB SSD



# Data and Content Lake

Linguistic enrichments to support **semantic retrieval** and **fact extraction**



Summary:

- Scalable and distributed dynamic enrichment workflows
- **CRFs** for NER, **PageRank** (variant) for NED, **CNNs** and **LSTMs** for Relation/Event extraction, sentiment analysis
- **TensorFlow** and **GPUs** for training infrastructure

# Enrichments

We can't foresee all uses of our data: **Developer APIs** to **Integrate** and **orchestrate** third party tools. **Personalization** is key in Data Science: A **flexible data wrangling** infrastructure is required.

The screenshot displays the 'ENRICHMENTS' marketplace interface. The top navigation bar includes 'INSIGHTS', 'ENRICHMENTS', 'DATA SETS', and a 'GO' button. Below the navigation, there are tabs for 'Enrichments', 'Enrichments Marketplace' (which is selected), and 'Unpublished Enrichments'. The main content area is divided into two sections: 'Recent Enrichments' and 'Popular Enrichments'. Each section contains a grid of enrichment service cards. Each card features an icon, a title, and a red plus sign in the bottom right corner, indicating that the service can be added to a pipeline.

**Recent Enrichments:**

- CoreNLP service - Sentences extraction (Icon: Red 'S' with a dollar sign)
- FHAI Basic annotators (Icon: Fairhair logo)
- IBM Keywords (Icon: Blue globe with lines)
- CoreNLP service - named entities extraction (Icon: Red 'S' with a dollar sign)
- IBM Entity Sentiment (Icon: Blue globe with lines)

**Popular Enrichments:**

- IBM Entities (Icon: Blue globe with lines)
- FHAI Concepts (Icon: Fairhair logo)
- FHAI Keyphrases (Icon: Fairhair logo)
- Microsoft Parse Tree (Icon: White bar)
- FHAI Sentiment (Icon: Fairhair logo)
















- Interoperate with state-of-the-art external enrichments
- Chain multiple external enrichments
- Train your own models!

# Connectors to Internal Systems

Goal is to join **Outside Insights** with **Internal Data** and workflows

**Data Ingestion & Insights Delivery** by setting up simple schema mappers

Destinations Library

<p>Google Big Query </p> <p>Get your data into BigQuery</p>	<p>CouchDB </p> <p>Get your data into a CouchDB database</p>	<p>Amazon DynamoD </p> <p>Get your data into an Amazon DynamoDB database</p>	<p>ElasticSearch </p> <p>Get your data into Elastic Search</p>
<p>FTP </p> <p>Get your data as a series of files</p>	<p>HTTP </p> <p>Get your data via Webhooks</p>	<p>MongoDB </p> <p>Get your data into a MongoDB database</p>	<p>MySQL </p> <p>Get your data into MySQL</p>
<p>PostgreSQL </p> <p>Get your data into PostgreSQL</p>	<p>Redis </p> <p>Get your data into a Redis list</p>	<p>Amazon S3 </p> <p>Get your data into an Amazon S3 bucket</p>	<p>SFTP </p> <p>Get your data as a series of files</p>
<p>Splunk Enterprise </p> <p>Analyze Social Data in Splunk</p>	<p>Splunk Storm </p>	<p>Zoomdata </p>	

# Knowledge Graph

Wait... did you say PageRank, triples? So do you have a (Knowledge) Graph?



## Content:

- Companies
- Brands
- Products
- Key people
- Influencers




## Goals:

- Relate facts
- Data mining
- Cognitive applications (higher-order reasoning)



## Challenges:

- Data Cleaning
- Data deduplication / integration
- Truth Finding



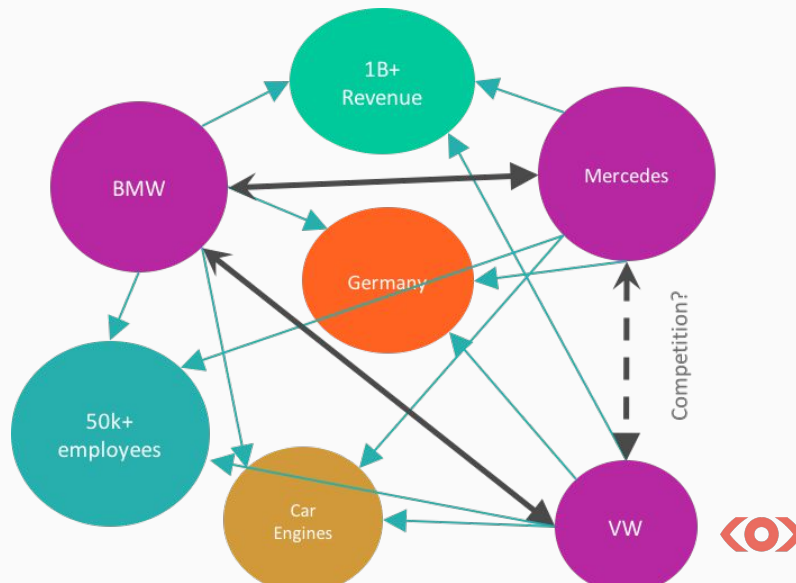
# Cognitive Applications

Infer **high-level insights** from a set of **extracted events/facts**.

- Competitor
- Customer
- Investment
- Lawsuit/Litigation
- Partnership
- Supplier
- Acquisition
- Out/under performance
- Expanding Operations
- Compliance
- Funding Developments
- Leadership Changes
- New Offerings
- Bankruptcy
- Restructuring, Cost Cutting

Insight **discovery**:

- Rule/Graph mining (data cleaning)
  - GPAR (VLDB '15)
  - RUDI-K (internal, paper submitted)
- Link prediction (data enrichment, fact checking)
  - Path Ranking Algorithms (PRA)
  - Probabilistic Soft Logic (PSL)
  - Snorkel (Stanford Collaboration)



Fairhair


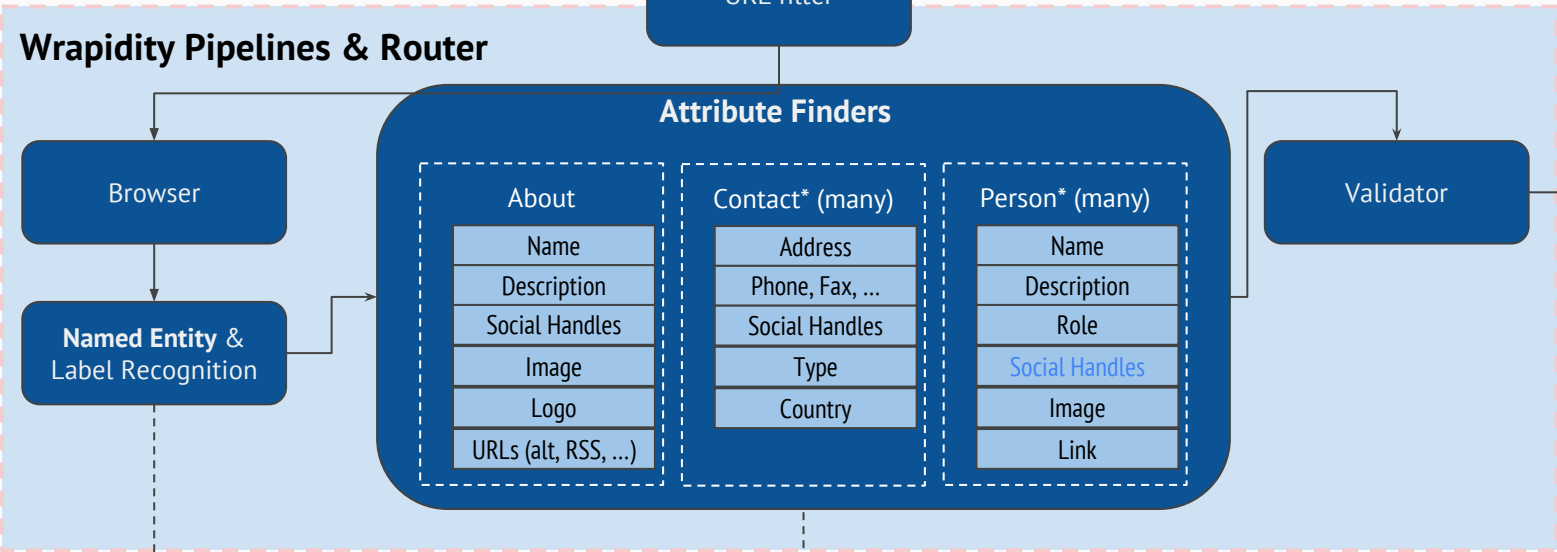

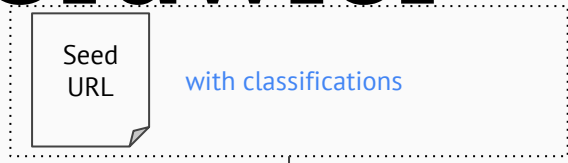




# More?



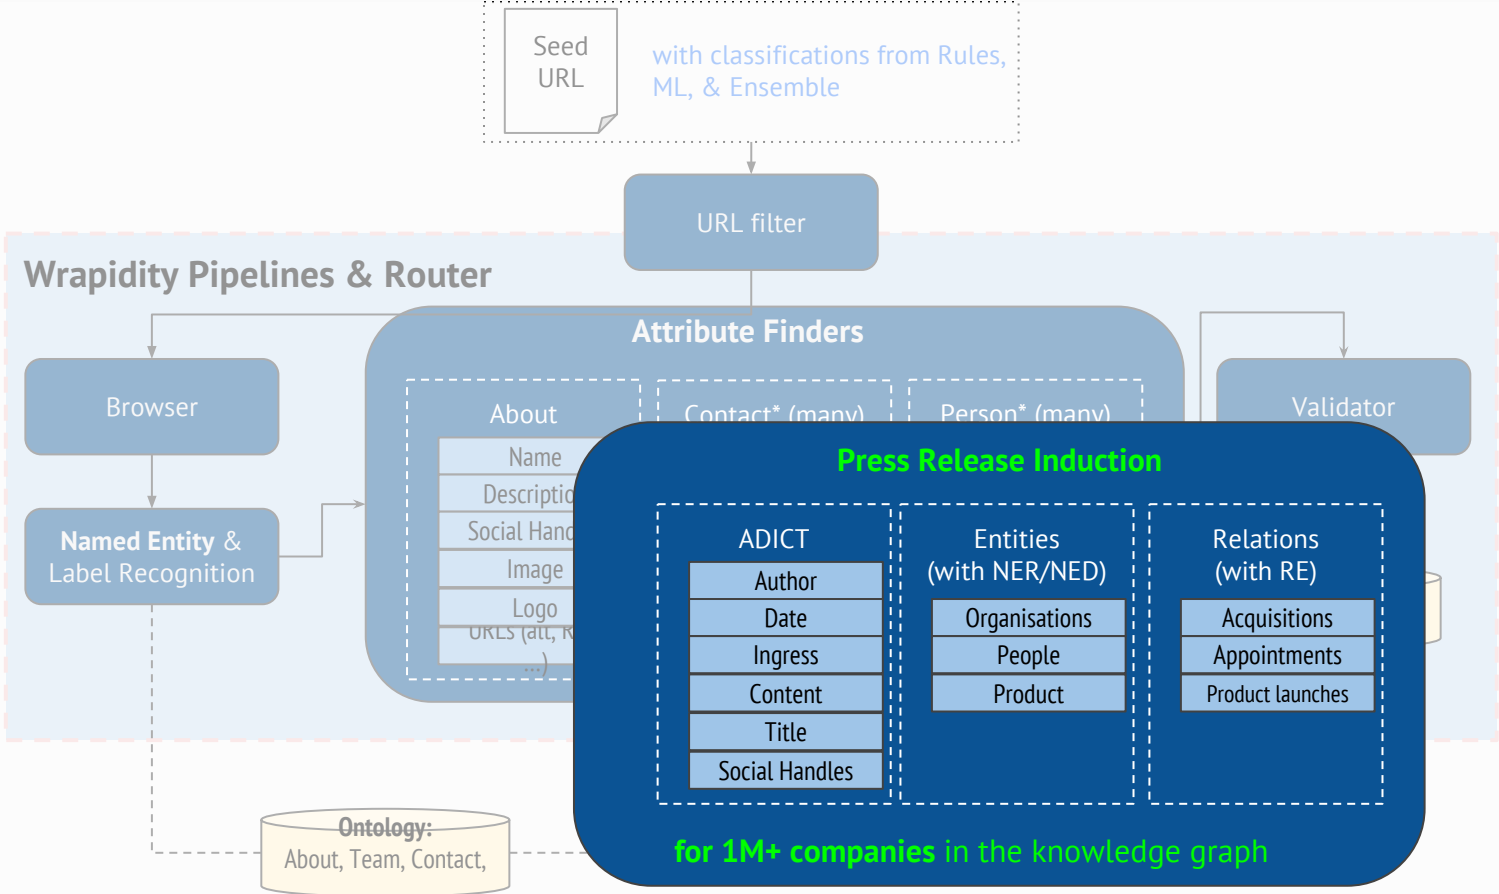
# Company Crawler



Questions



# Company Crawler



Questions

