

Exploring Concept Representations for Concept Drift Detection

Oliver Becher (CWI), Laura Hollink (CWI) & Desmond
Elliott (U.Edinburgh)

Can we detect changes in concepts over time from an annotated collection of articles?

Types of Concept Drift

Concept Drift as:

- Change in the intention of the concept, defined as the definition or the properties of the concept
- Change in the extension, or the instances, of a concept
- Change in the label of the concept.

Referring to changing Concepts can be Problematic

- Concepts in a KOS are used
 - for annotations (eg. meta data in library)
 - for background knowledge (eg. reasoning)

→ Interoperability between datasets & enable structured access

- Compromised if concepts change (drift)
- Problems after drift:
 - Correspondences between two concepts in different KOS may be incorrect in extension of one changes
 - User keyword query on historic corpus may be interpreted incorrectly if prevalent label to refer to a concept has changed

Aim at first concept change capture & evaluation

- Understand how concepts can be effectively represented to capture change
- Co-occurrence based:
 - Two concepts co-occur if they are annotations of the same document
 - Compare Concept representations at different points in time
- Qualitative evaluation based on visualization
- Quantitative evaluation based on Wikipedia edits

The ION Dataset



The screenshot shows a web browser displaying a New York Times article. The URL is highlighted in red. The article title is "Hillary Clinton Directs Aides to Give Email Server and Thumb Drive to the Justice Department", highlighted in purple. The author is Michael S. Schmidt, and the date is August 11, 2015. The article text is annotated with blue boxes: one for the first paragraph, one for the second paragraph, and one for the third paragraph. A blue box labeled "Category About (topics & entities)" is positioned to the left of the second paragraph. A red box labeled "Image URL" and "Image Features" points to an image of Hillary Clinton speaking at a podium. A red box labeled "Caption" points to the text below the image.

URL

POLITICS

Hillary Clinton Directs Aides to Give Email Server and Thumb Drive to the Justice Department

By MICHAEL S. SCHMIDT AUG. 11, 2015 Publishing Date

Headline

Email Share Tweet Save More

Hillary Rodham Clinton has directed her aides to give the Justice Department an email server that housed the personal account that she used exclusively while secretary of state, along with a thumb drive that contained copies of the emails, her presidential campaign said on Tuesday.

The Justice Department and the F.B.I. have sought the server and the thumb drive as they investigate how classified information was handled in connection with the account. Earlier on Tuesday, the inspector general for the intelligence community told members of Congress that Mrs. Clinton had "top secret" information — the highest classification of government intelligence — in two emails among the 40 from the private account that the State Department has allowed him to review.

Image URL
Image Features

Image URL
Image Features

Caption

What is a Concept?

Our practical definition:

- Each unique wikipedia annotation of a newspaper article
- Representation: Vector of co-occurrences with other concepts

General Approach

Qualitative exploration phase

Exploring
concepts

Find
representation
and
measurement

Comparing
concepts at
different points
in time

Visualize
concepts

Evaluation

Reducing sparsity

Normalize infrequent
reporting

Qualitative
comparison

IDF weighting

Cosine Vector
comparison

Example: Concept Vector

wiki/Hillary_Clinton

7	3	5	8	1	0
---	---	---	---	---	---

wiki/Bill_Clinton
wiki/Democratic
wiki/Chicago
wiki/Election
wiki/Foreign_Politics
wiki/Basketball

<http://en.wikipedia.org/wiki/>

Example: Temporal Concept Vectors with 3 bins

wiki/Hillary_Clinton

Aug - Dec

Jan - April

May - July

	<i>wiki/Bill_Clinton</i>	<i>wiki/Democratic</i>	<i>wiki/Chicago</i>	<i>wiki/Election</i>	<i>wiki/Foreign_Politics</i>	<i>wiki/Basketball</i>
Aug - Dec	7	3	5	8	1	0
Jan - April	5	3	4	9	4	0
May - July	1	9	9	1	7	2

→ Cosine Sim 0.893

→ Cosine Sim 0.672

Example: Flexible Time Bins with 3 bins

wiki/Hillary_Clinton

Aug - Sept

Sept - May

June - July

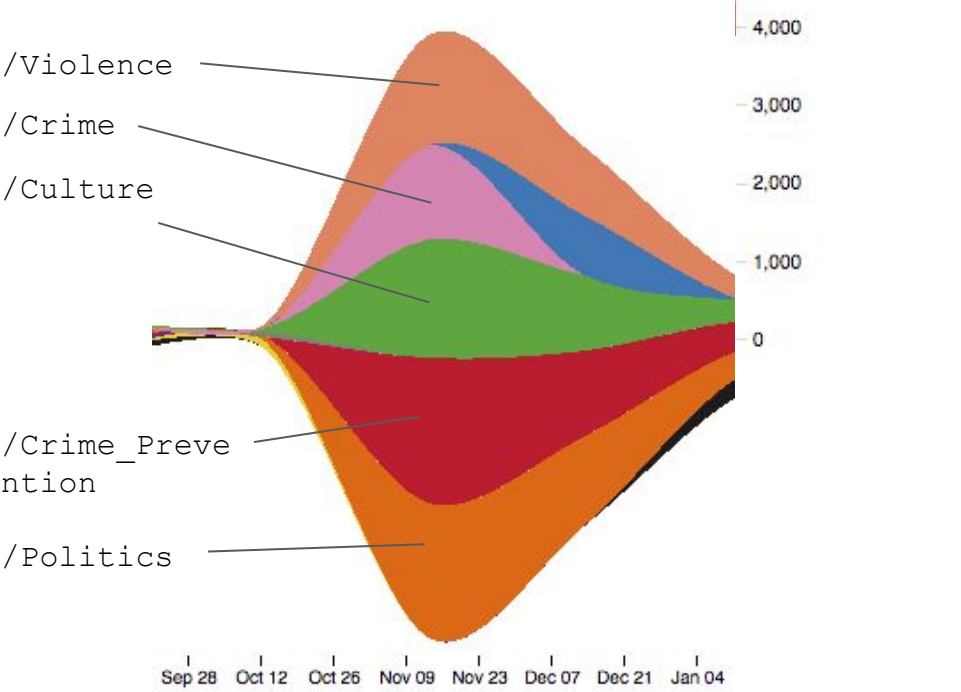
	<i>wiki/Bill_Clinton</i>	<i>wiki/Democratic</i>	<i>wiki/Chicago</i>	<i>wiki/Election</i>	<i>wiki/Foreign_Politics</i>	<i>wiki/Basketball</i>
Aug - Sept	7	3	5	8	1	0
Sept - May	7	3	4	9	2	0
June - July	1	9	9	1	8	2

→ Cosine Sim 0.953

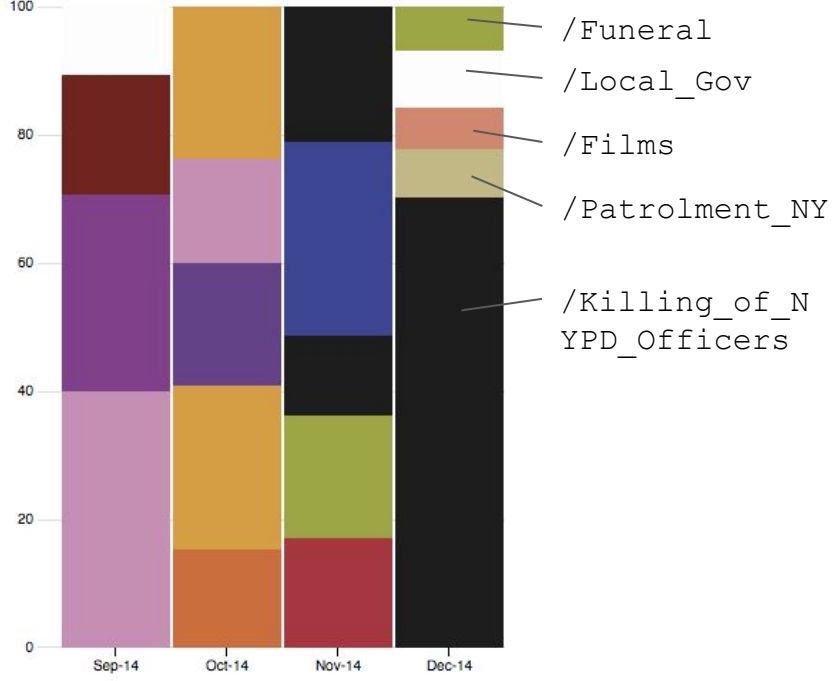
→ Cosine Sim 0.472

Concept: Death_of_Eric_Garner

Core of the concept:



New context of the concept:



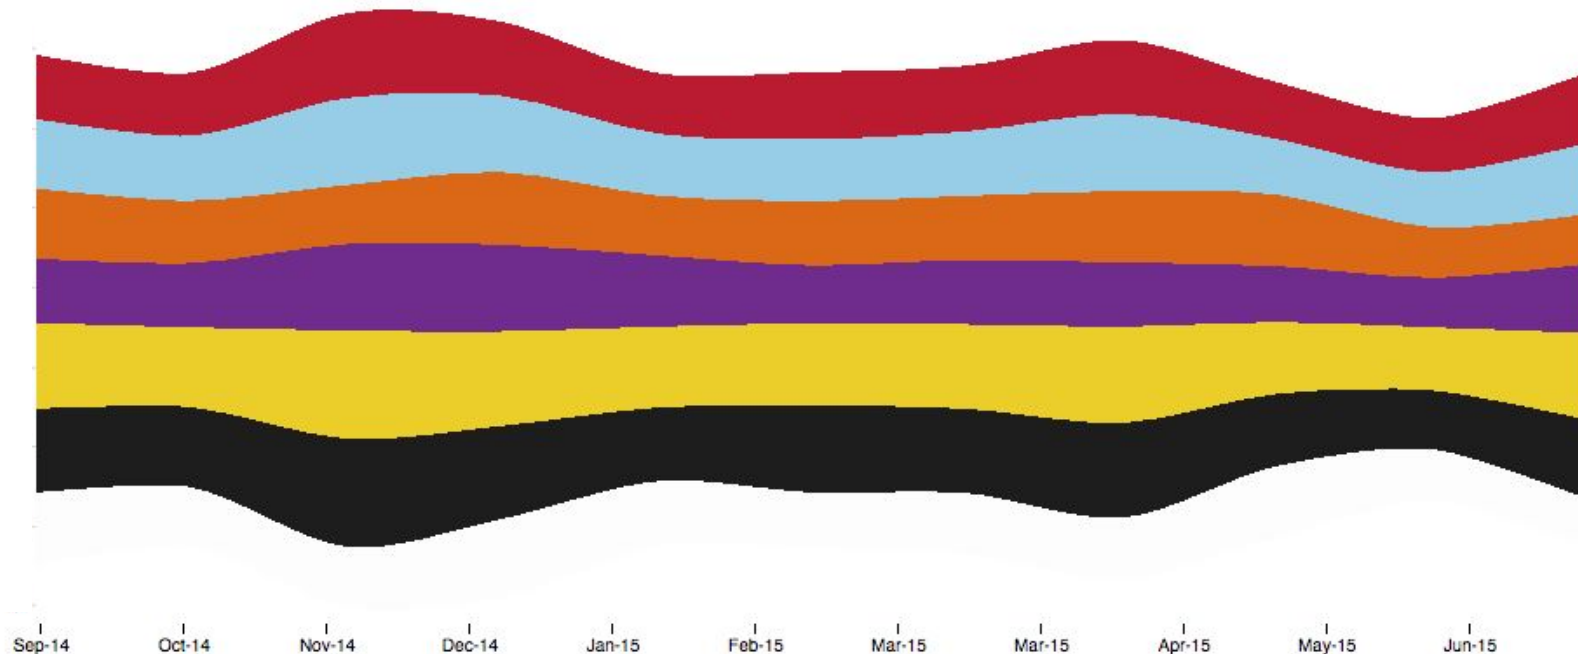
Qualitative evaluation based on cosine similarities

- Visualization based
- Based on streamgraph


Selecting concepts to evaluation by:

- Average cosine similarities for each concept
 - Each concept has one averaged change score
 - Pick concepts with highest averaged score and lowest average score

Highest Averaged Cosine Similarity (WP:Police)



Lowest Averaged Cosines Similarities (WP:New_York_University)



Quantitative Evaluation

How to Evaluate with no Truth

- No datasets available for concept change
- What's a good reference point?

→ Assumption: When a concept is changing, someone will change the wikipedia page

Towards a Quantitative Evaluation

- Retrieval of Wikipedia Edits for each concept
- Create Counts of edits, matching to time frames of temporal concept vectors
- Compare vectors of edit counts to vectors of similarity scores
- Use spearman correlation

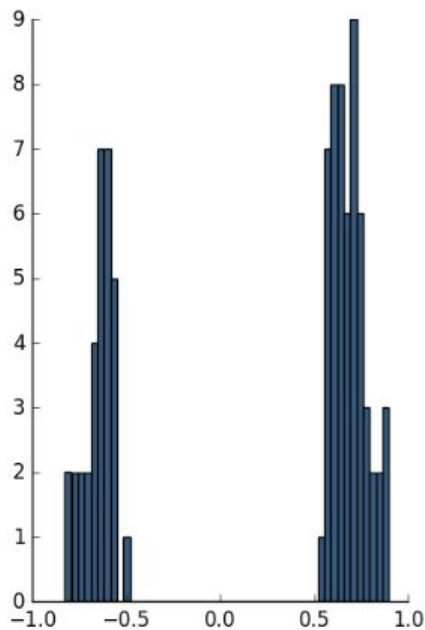
Towards a Quantitative Evaluation

- Different set ups to obtain similarity scores
 - Cosine Similarity / KL Divergence
 - Flexible Time bins / fixed time bins
 - TF-IDF / no TF-IDF
 - 5 different amounts of bins (100, 52, 24, 12, 6)


Results

- Evaluated 964 concepts (stratified sample)
- Flexible time bins much better than fixed time bins
- KL Divergence better than Cosine similarity
- TF-IDF no effect
- Amount of bin needs to be sufficient (at least 24, 52 is better)

Histogram of Spearman correlations with different bin sizes



(a) $n=12$



(b) $n=100$

Types of Concept Drift

Concept Drift as:

- **Change in the intention of the concept, defined as the definition or the properties of the concept**
- **Change in the extension, or the instances, of a concept**
- Change in the label of the concept.

Outlook

- Concept drift may take longer than 1 year
- Nonetheless, measureable differences
- Temporal representation valuable
- No perfect evaluation tool / gold truth available
- More research on which types of concepts correlate to WP edits to understand WP as evaluation tool
- Distinguish actual change and simple growth of WP article

Questions