NLM's Participation in the 2013 Medical Retrieval Tasks of ImageCLEF

Matthew S. Simpson

Lister Hill National Center for Biomedical Communications, U.S. National Library of Medicine, NIH, Bethesda, MD, USA

November 15, 2013

ITI Research Group


Matthew S. Simpson

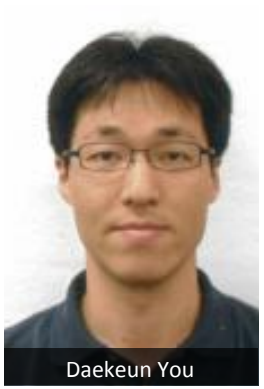
Daekeun You


Md Mahmudur Rahman


Zhiyun Xue


Dina Demner-Fushman


Sameer Antani

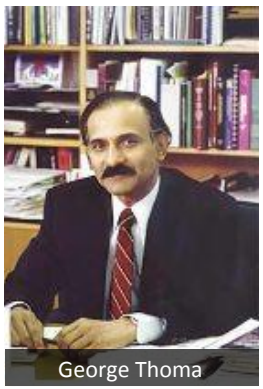
George Thoma

Task Participation

- Image-based retrieval
- Case-based retrieval
- Modality classification
- Compound figure separation

Image-Based Retrieval Task

Example

Topic 1: Osteoporosis x-ray images

Goal

- Retrieve x-ray images from the articles in PubMed Central that depict osteoporosis

Case-Based Retrieval Task

Example

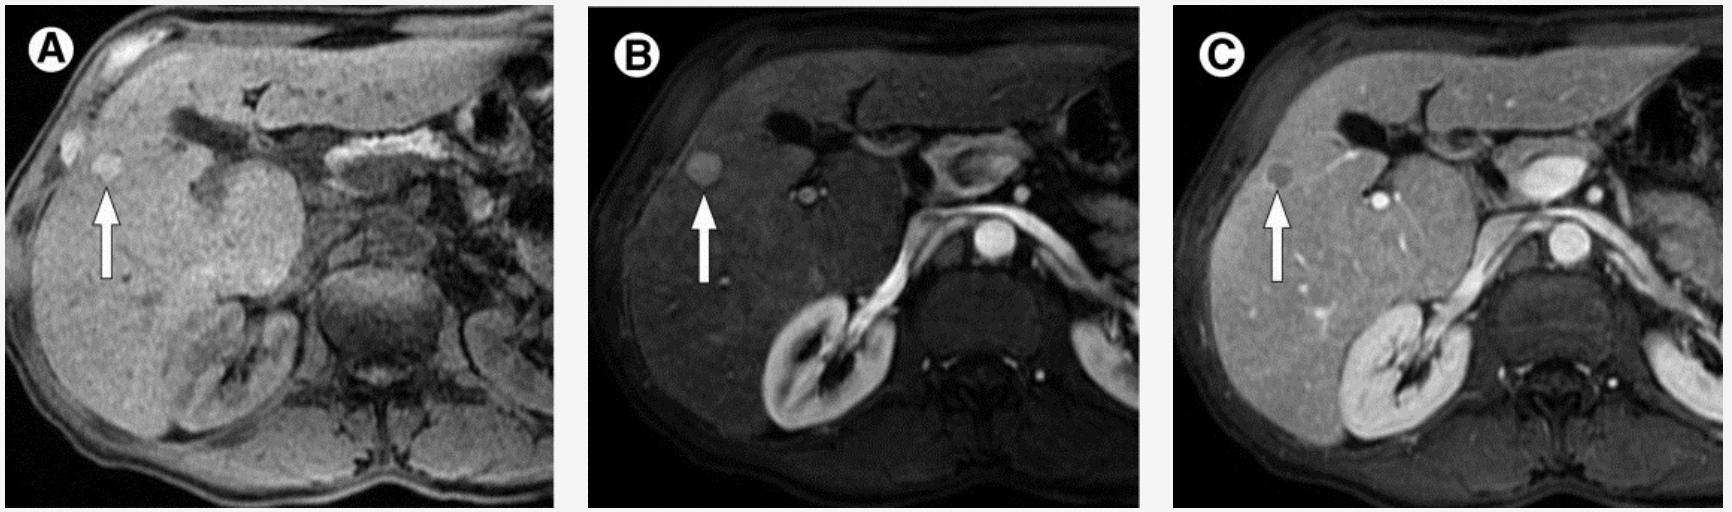

Topic 8: A 56-year-old woman with Hepatitis C, now with abdominal pain and jaundice. Abdominal MRI shows T1 and T2 hyperintense mass in the left lobe of the liver which is enhanced in the arterial phase.

Goal

- Retrieve articles from PubMed Central that describe similar patient cases

Overview

1. Image Representation
2. Textual Retrieval Strategies
3. Visual Retrieval Strategies
4. Mixed Retrieval Strategies
5. Results

Image Representation

Textual Features

- Article title, abstract, and MeSH terms
- Image caption and mention

Visual Features

- Autocorrelation
- Edge frequency
- Fuzzy color and texture (FCTH)
- Gabor moment
- Gray-level co-occurrence matrix (GLCM)
- Local binary pattern (LPB)
- Scale-invariant features (SIFT)
- Shape moment
- Tamura moment
- Edge histogram (EHD)
- Color and edge directivity (CEDD)
- Primitive length
- Color layout (CLD)
- Color moment
- Semantic concept (SCONCEPT)

Enriched Citations

Logical Document Representations

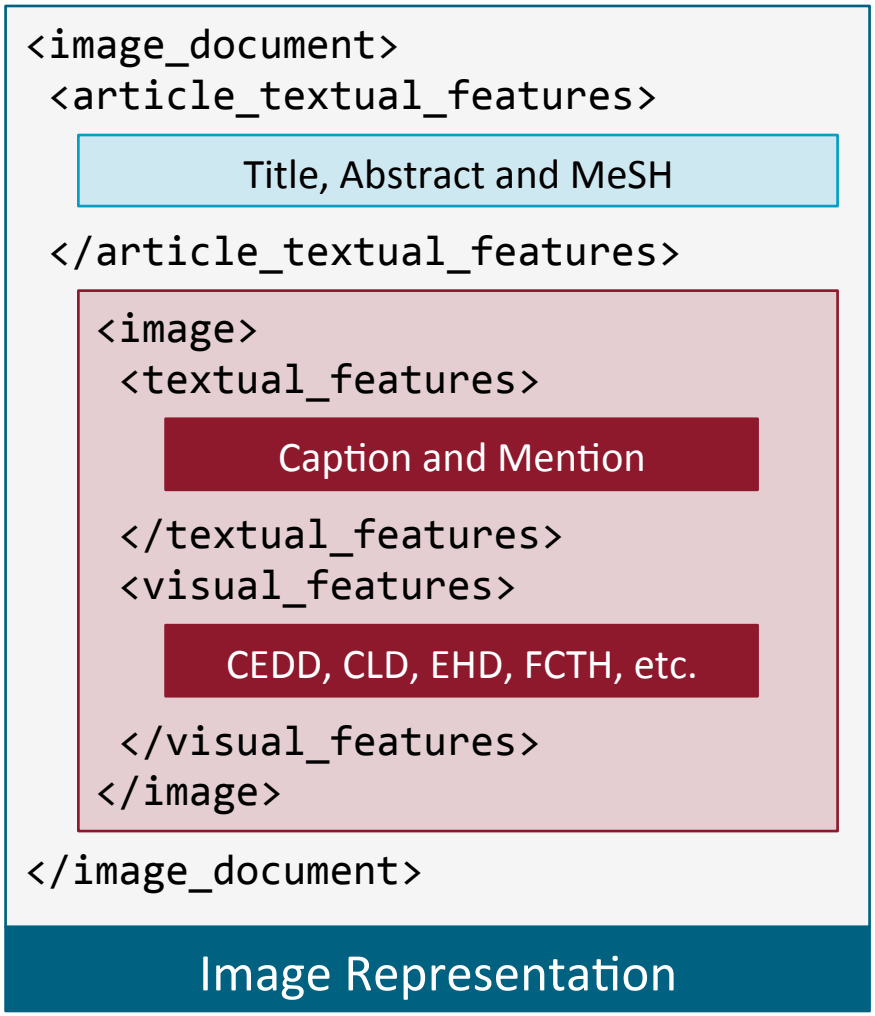

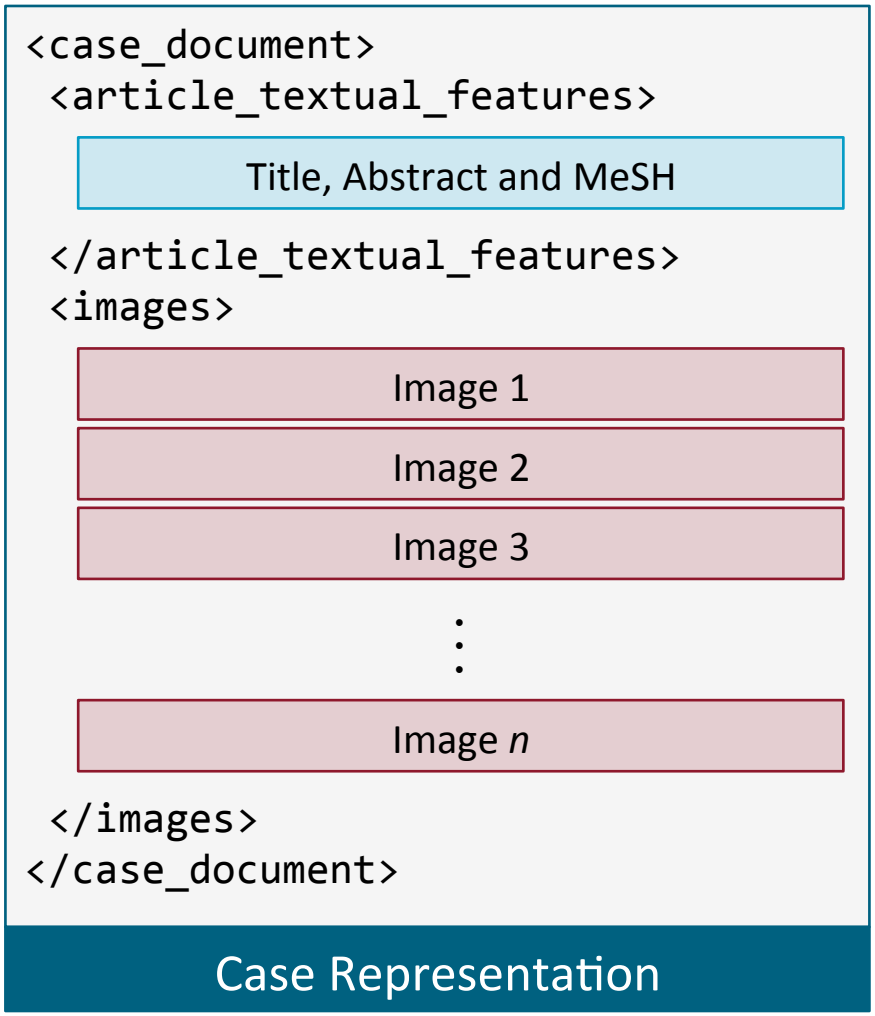

Textual Retrieval Strategy

Basic Approach

Indexing

- Essie biomedical retrieval system
- Expands terms along the synonymy relationships in the UMLS
- Weights term occurrences based on their location
- Implements a variation of the query likelihood model

Querying

- Lossy expansion of verbatim topic

Scoring

- Most probable documents given the query
- Documents ranked by the maximum of each (sentence, field) pair

Textual Retrieval Strategy

Improving Precision with Structured Queries

```
<query>
 <sentence>
  <verbatim>
```

osteoporosis x-ray images

```
  </verbatim>
  <modality>
```

x-ray

```
  </modality>
  <concept>
```

osteoporosis

```
  </concept>
 </sentence>
</query>
```

Structured Query

Querying

1. Disjunction of modalities must occur in caption or mention; concepts can occur anywhere
2. Verbatim topic can occur anywhere

Visual Retrieval Strategy

Basic Approach

Indexing

- Extract low-level visual features from all images

Querying

- Extract low-level visual features form query images
- Compute Euclidean distances in brute-force fashion for each feature

Scoring

- Merge distances as a weighted linear combination of features (weights either predetermined or based on modality classification accuracy)
- Rank images in increasing order of normalized distance

Visual Retrieval Strategy

Improving Precision with Modality Filter

Querying

1. Filter image collection based on modality of query images
2. Perform basic visual approach on remaining image collection

Combining Textual and Visual Approaches

Early Fusion

- Represent visual features as text

Late Fusion

- Merge textual and visual retrieval results by score
- Merge textual and visual retrieval results by rank

Early Fusion

Representing Visual Features as Text

Process

1. Extract visual features form all images and group them into clusters
2. Map each cluster to a unique alphanumeric word and assign it to all images whose features are members of the cluster
3. Add the visual words assigned to an image to its enriched citation and index with Essie
4. Search the index with multimodal queries by first assigning visual words to a topic's example images and then by combining them with the topic's description

Enriched Citations

Logical Documents with Visual Features Represented as Words

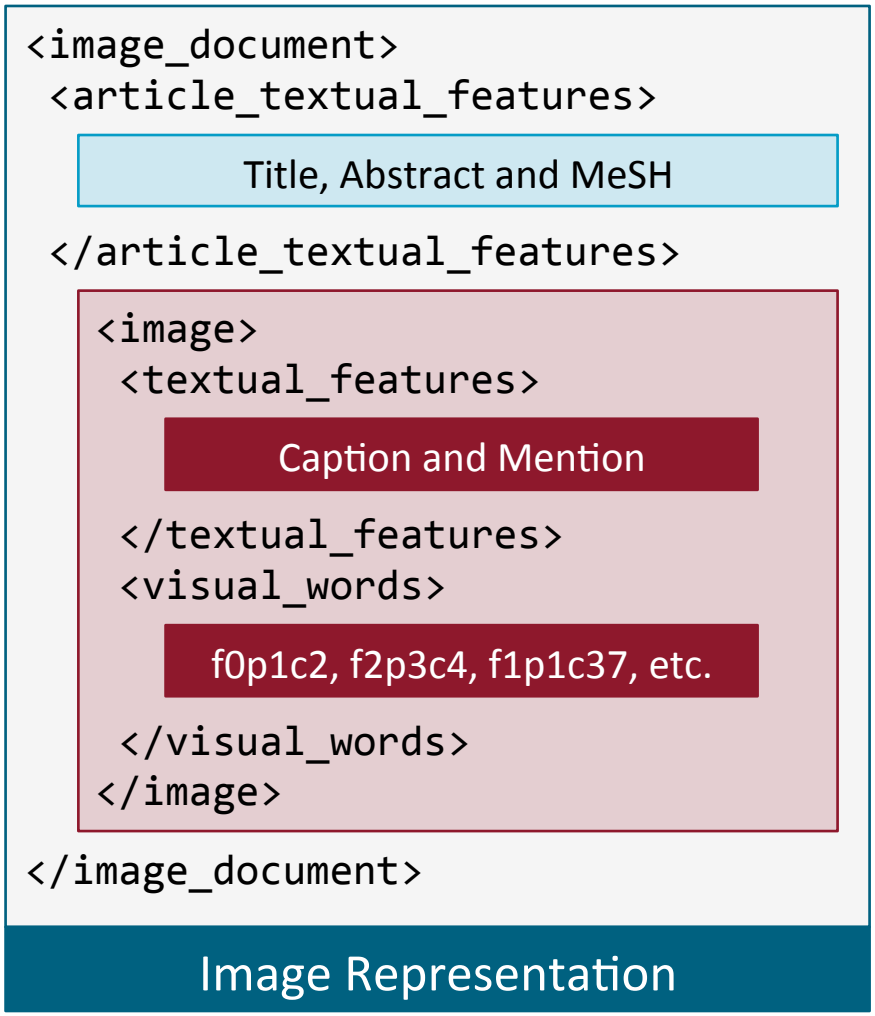

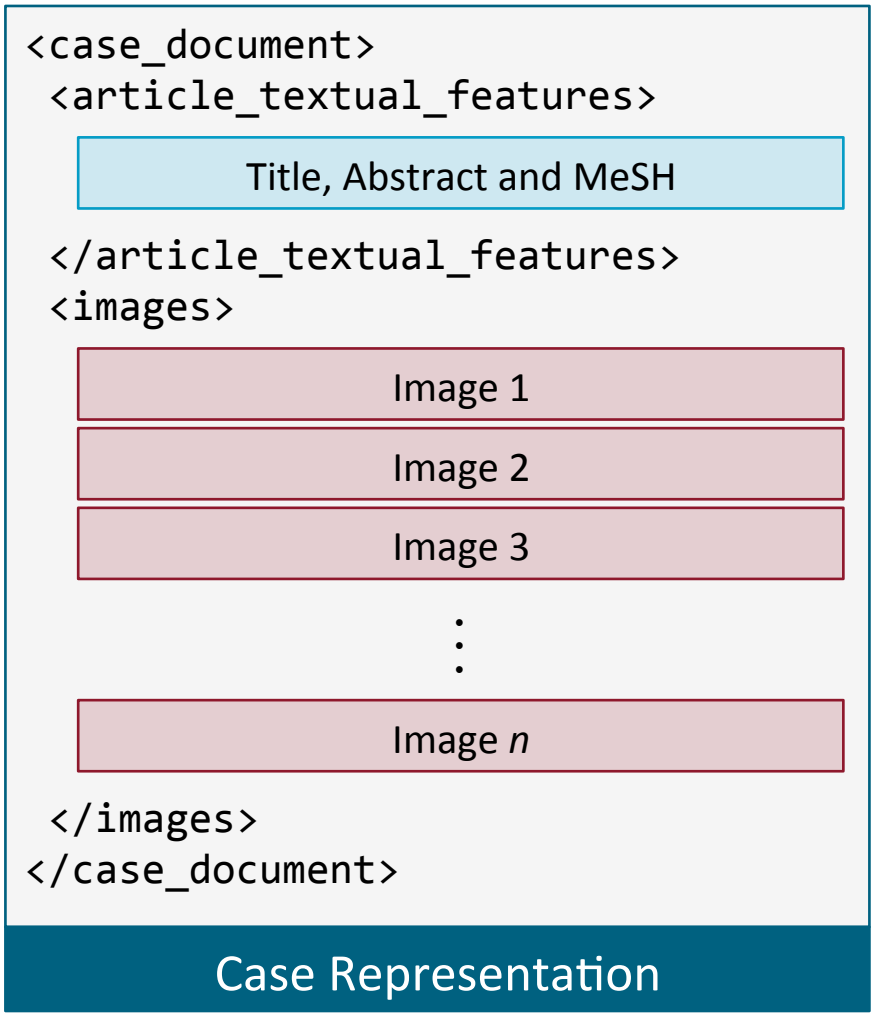

Early Fusion

Retrieval Strategy

```
<query>
 <sentence>
  <verbatim>
```

osteoporosis x-ray images

```
  </verbatim>
  <modality>
```

x-ray

```
  </modality>
  <concept>
```

osteoporosis

```
  </concept>
 </sentence>
 <visual_words>
```

f0p1c2, f2p3c4, f1p1c37, etc.

```
 </visual_words>
</query>
```

Structured Query

Querying

1. Disjunction of modalities must occur in caption or mention; concepts can occur anywhere
2. Verbatim topic can occur anywhere
3. Disjunction of visual words must occur in visual features field

Late Fusion

Merging Textual and Visual Results by Score and Rank

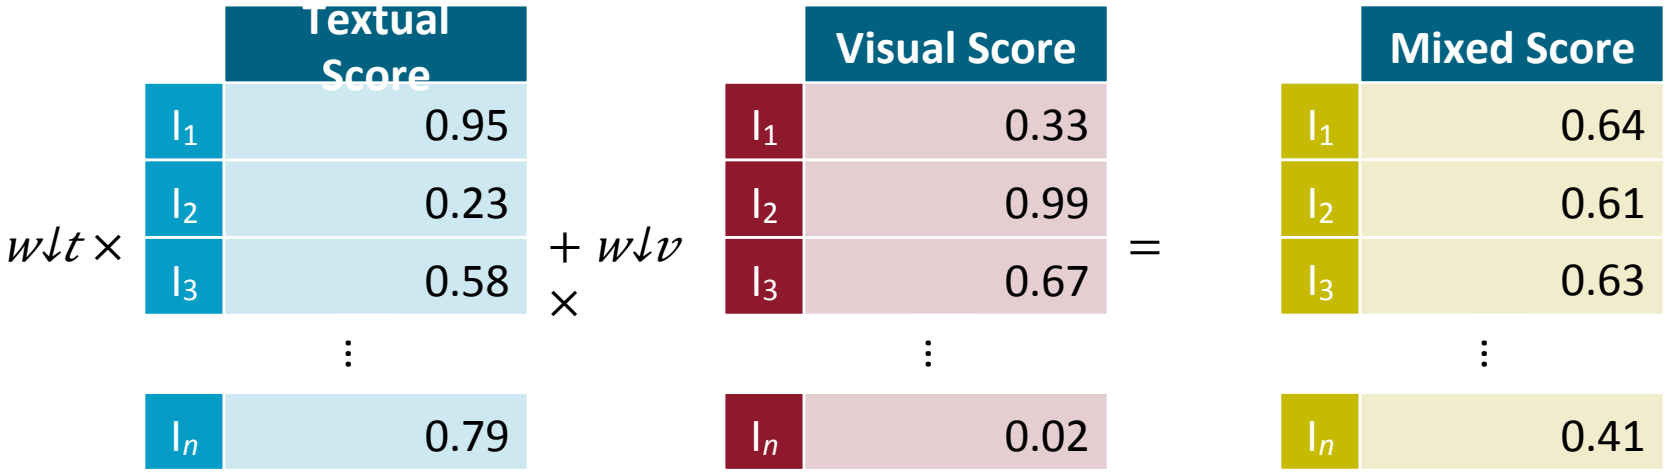

For merging by rank, score as reciprocal rank and then merge by score

Image-Based Retrieval Results

Method	Mode	MAP	bpref	P@10
Structured query with visual words	Mixed	0.3196	0.2983	0.3886
Structured query	Textual	0.3196	0.2982	0.3886
Structured query, basic visual approach with modality weights, score-merged	Mixed	0.3124	0.3014	0.3886
Structured query, basic visual approach with modality weights, rank-merged	Mixed	0.3120	0.2950	0.3771
Structured query, basic visual approach with predefined weights, scoremerged	Mixed	0.3086	0.2938	0.3857
Structured query, basic visual approach with predefined weights, rank-merged	Mixed	0.3032	0.2872	0.3943
Basic visual approach, predefined weights	Visual	0.0110	0.0207	0.0257
Basic visual approach, modality filter	Visual	0.0101	0.0244	0.0343
Basic visual approach, modality weights	Visual	0.0092	0.0179	0.0314
Visual words	Visual	0.0002	0.0021	0.0029

Case-Based Retrieval Results

Method	Mode	MAP	bpref	P@10
Structured query with visual words	Mixed	0.0886	0.0926	0.1457
Structured query	Textual	0.0885	0.0926	0.1457
Visual words	Visual	0.0008	0.0044	0.0057

Conclusions

1. Textual approaches work well for literature-based retrieval tasks
2. Visual approaches add value but also risk hurting performance
3. Structured query representations improve early precision
4. A retrieval system tailored for biomedical language is beneficial
5. Image-based and case-based tasks require different methods

TREC 2014

Clinical Decision Support Track

http://openi.nlm.nih.gov