Concept-Based and Multimodal Methods for Medical Case Retrieval

PhD Defense

Mario Taschwer

May 17, 2017

Outline 1

1. Introduction
   - Medical Case Retrieval (MCR)
   - Problem Statement
   - Contributions
2. Biomedical Concept Mapping
3. Multimodal MCR
4. Further Work
5. Conclusion

Medical Case Retrieval (MCR)

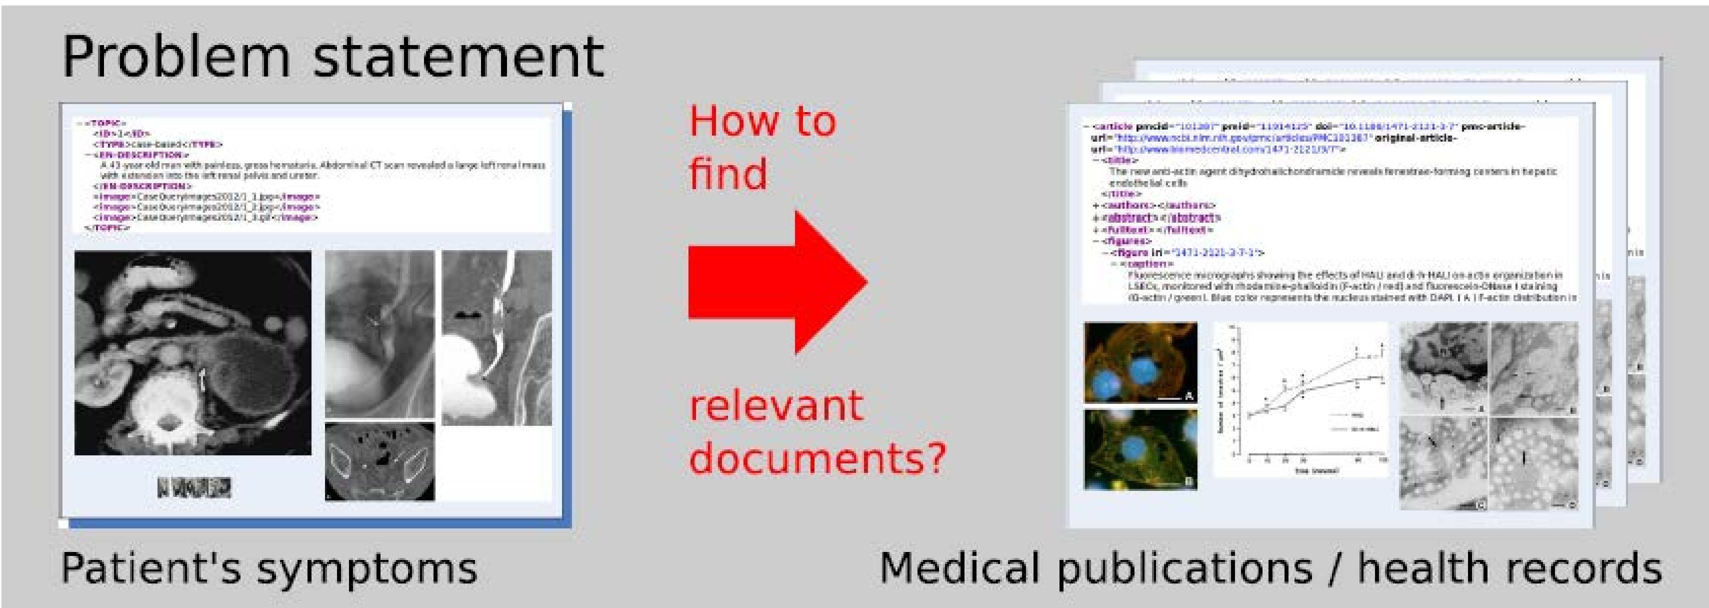

- Major component of medical decision support systems based on case-based reasoning
- Solution may help to generate datasets for medical education and research

Problem Statement

- State of the art for MCR on general datasets:
  - Best systems employ purely textual techniques
- Main research problem:
  - How to improve MCR methods using textual and visual information?
- Hypothesis:
  - Biomedical concepts may help – with techniques:
  - Query or document expansion for text retrieval
  - Concept-based retrieval
  - Fusion of text and concept-based retrieval

Contributions of PhD Thesis

- Novel automatic methods for compound figure classification and separation
- Evaluation of concept mapping techniques:
  - New and existing methods of mapping text or images to biomedical concepts
- Comparison of query and document expansion by biomedical concepts for text-based MCR
- Novel framework combining text and concept-based retrieval, improving over state of the art

Contributions of PhD Thesis

- Novel automatic methods for compound figure classification and separation
- Evaluation of concept mapping techniques:
  - New and existing methods of mapping text or images to biomedical concepts
- Comparison of query and document expansion by biomedical concepts for text-based MCR
- Novel framework combining text and concept-based retrieval, improving over state of the art

Outline 2

1. Introduction
2. Biomedical Concept Mapping
   - Medical Subject Headings (MeSH)
   - Text-to-Concept Mapping
3. Multimodal MCR
4. Further Work
5. Conclusion

Medical Subject Headings

- Controlled vocabulary of biomedical concepts:
  - ~27k primary terms, ~161k synonyms
  - "More general than" relations between primary terms impose 16 tree structures (maximal depth 11)
- Used to annotate biomedical publications

Primary MeSH Term	*Node Identifier*	*Specialty*
Eye Neoplasms	C04.588.364	2
Neoplasms by Site	C04.588	1
Neoplasms	C04	0
Eye Neoplasms	C11.319	1
Eye Diseases	C11	0

Text-to-Concept Mapping 1

- Existing systems:
  - MetaMap, Open Biomedical Annotator: slow
  - Whatizit MeshUp: kNN classifier, short text input only
- Novel, more efficient string matching approach:
  - based on inverted index of MeSH terms
  - finds (partial) occurrences of MeSH terms in single pass through text document
- Effectiveness evaluated for two objectives:
  - classification: reproducing manual MeSH annotations
  - concept-based retrieval on MCR dataset (~75k docs)

TEXT-TO-CONCEPT MAPPING 2

Concept-based retrieval on MCR dataset
MGT: “ideal” concept mapping using ground-truth MeSH terms

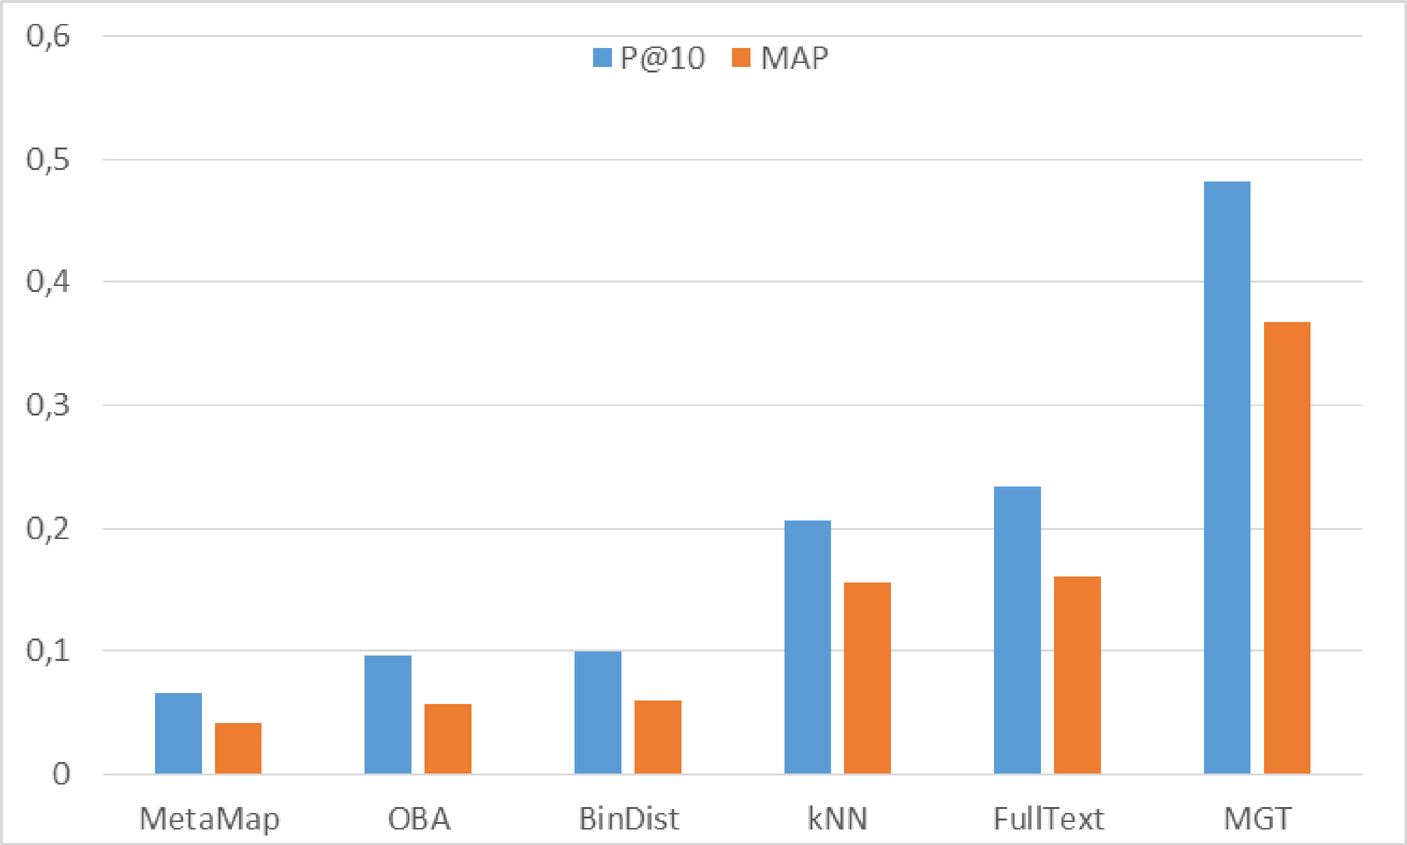

35 queries
75k docs
BinDist index

Algorithms used for query mapping

OUTLINE 3

1. Introduction
2. Biomedical Concept Mapping
3. Multimodal MCR
   - Framework for Text- and Concept-Based Retrieval
   - Fusion Methods
   - Results
4. Further Work
5. Conclusion

RETRIEVAL FRAMEWORK

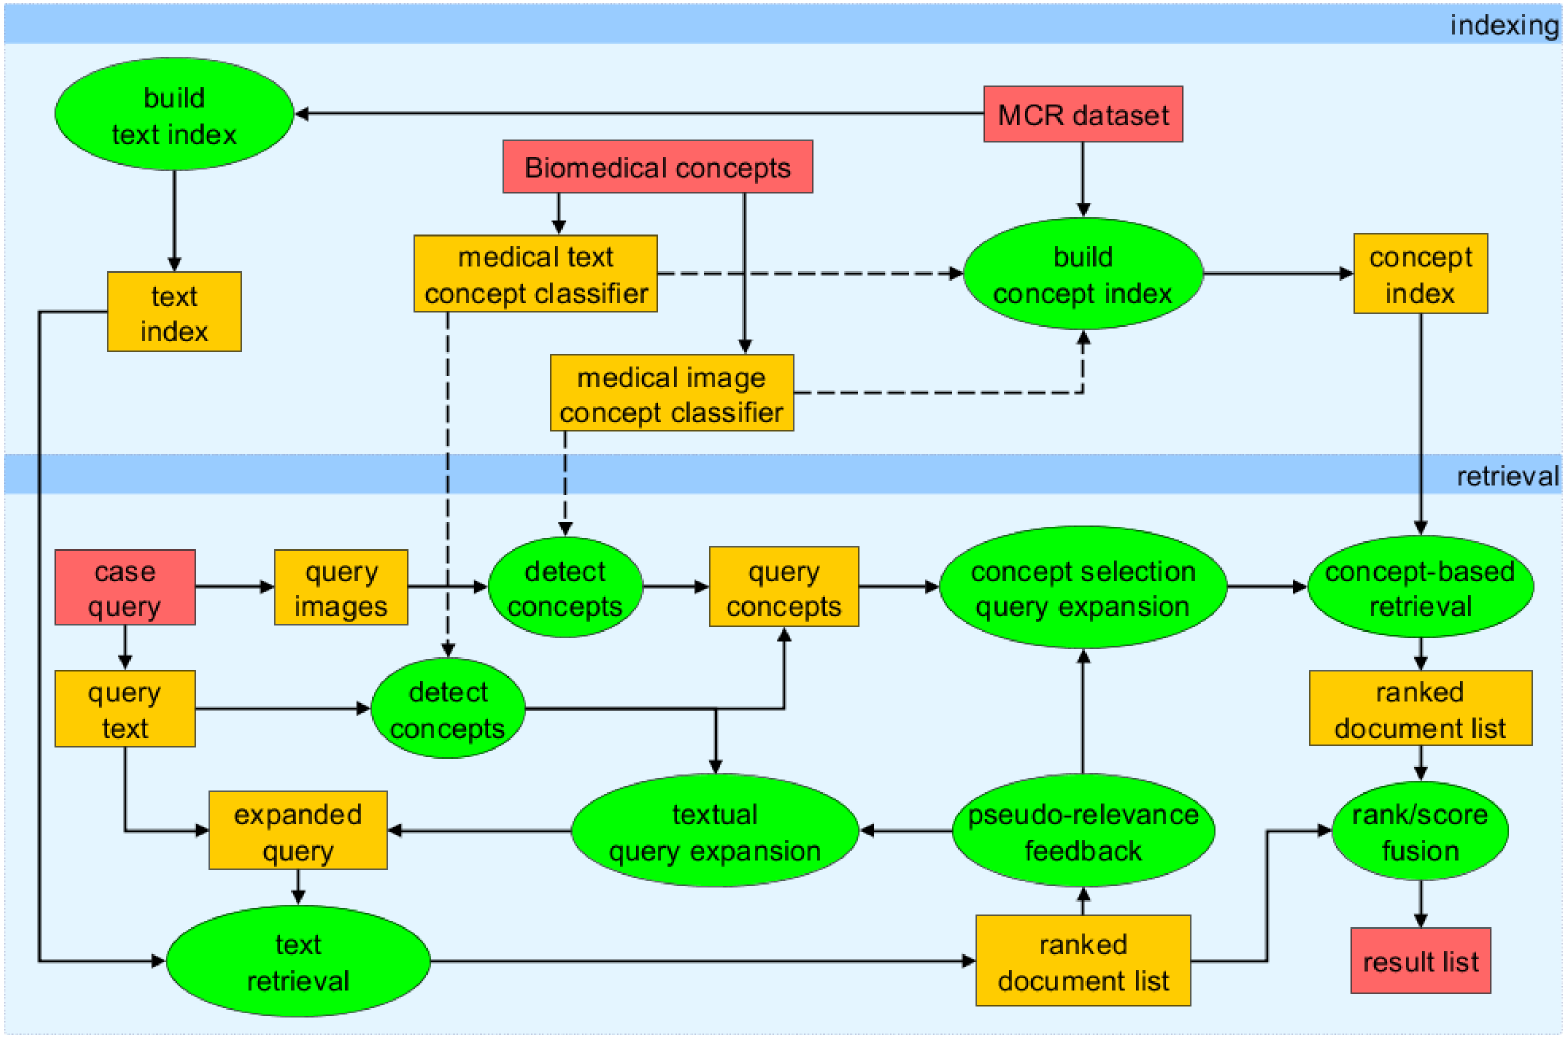

Fusion Methods

- Fuse result lists of retrieval methods A and B
- Linear fusion: $s = \beta * s_A + (1- \beta) * s_B$
  - Combine retrieval scores with fixed weight β
  - s_A, s_B: logistic score normalization from rank positions
- Query-adaptive fusion (QAF):
  - For each query q, choose weight β depending on q
  - E.g. by estimating performance of A and B for q

$$\beta = p_A^2 / (p_A^2 + p_B^2)$$

  - "Ideal" QAF: use an oracle returning true average precision for q, used as p_A and p_B

Fusion Results on MCR Dataset

F: fulltext retrieval (R)
T: text-based R (query expansion)
L: linear fusion

C: practical concept-based R
C+: ideal concept-based R
Q: ideal query-adaptive fusion

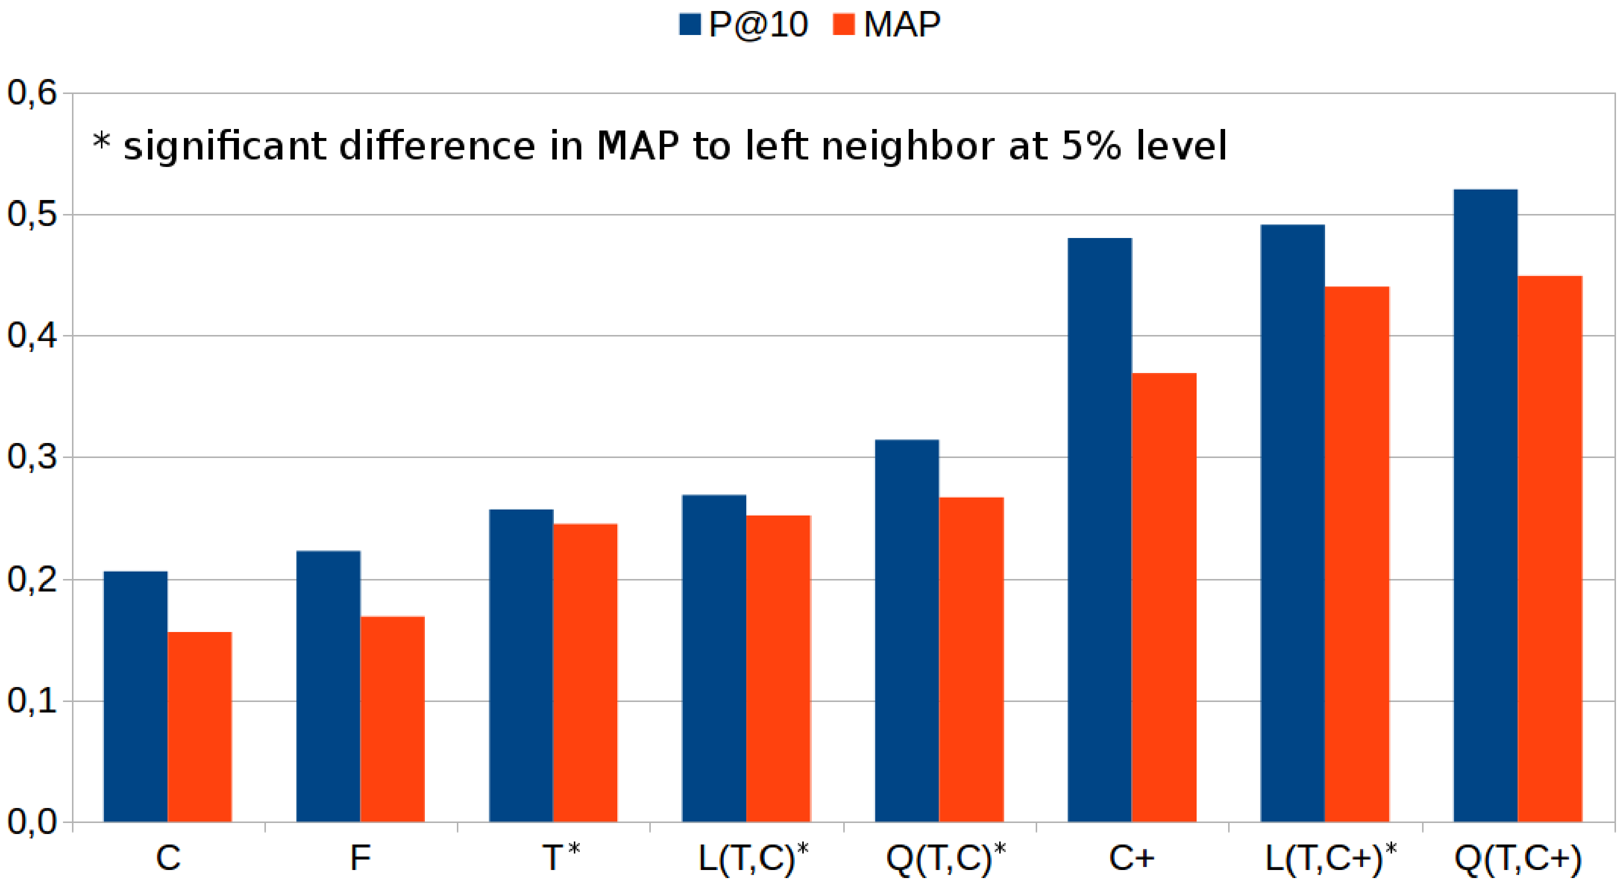

Outline 4

1. Introduction
2. Biomedical Concept Mapping
3. Multimodal MCR
4. Further Work
   - Concept Mapping
   - Retrieval in Multi-View Latent Space
5. Conclusion

Further Work 1

- Concept mapping:
  - Apply multi-view learning
    - Textual and visual modalities can be used as views and non-linearly mapped to a shared latent space
    - Concept mapping is learned by linear projections or kNN techniques in latent space
    - Conceptual and experimental work partly done
  - Apply deep learning to concept mapping
    - Recent advances in image caption generation may provide a starting point

FURTHER WORK 2

- Retrieval in multi-view latent (MVL) space:
  - Assumption: nearby points in latent space represent semantically similar cases

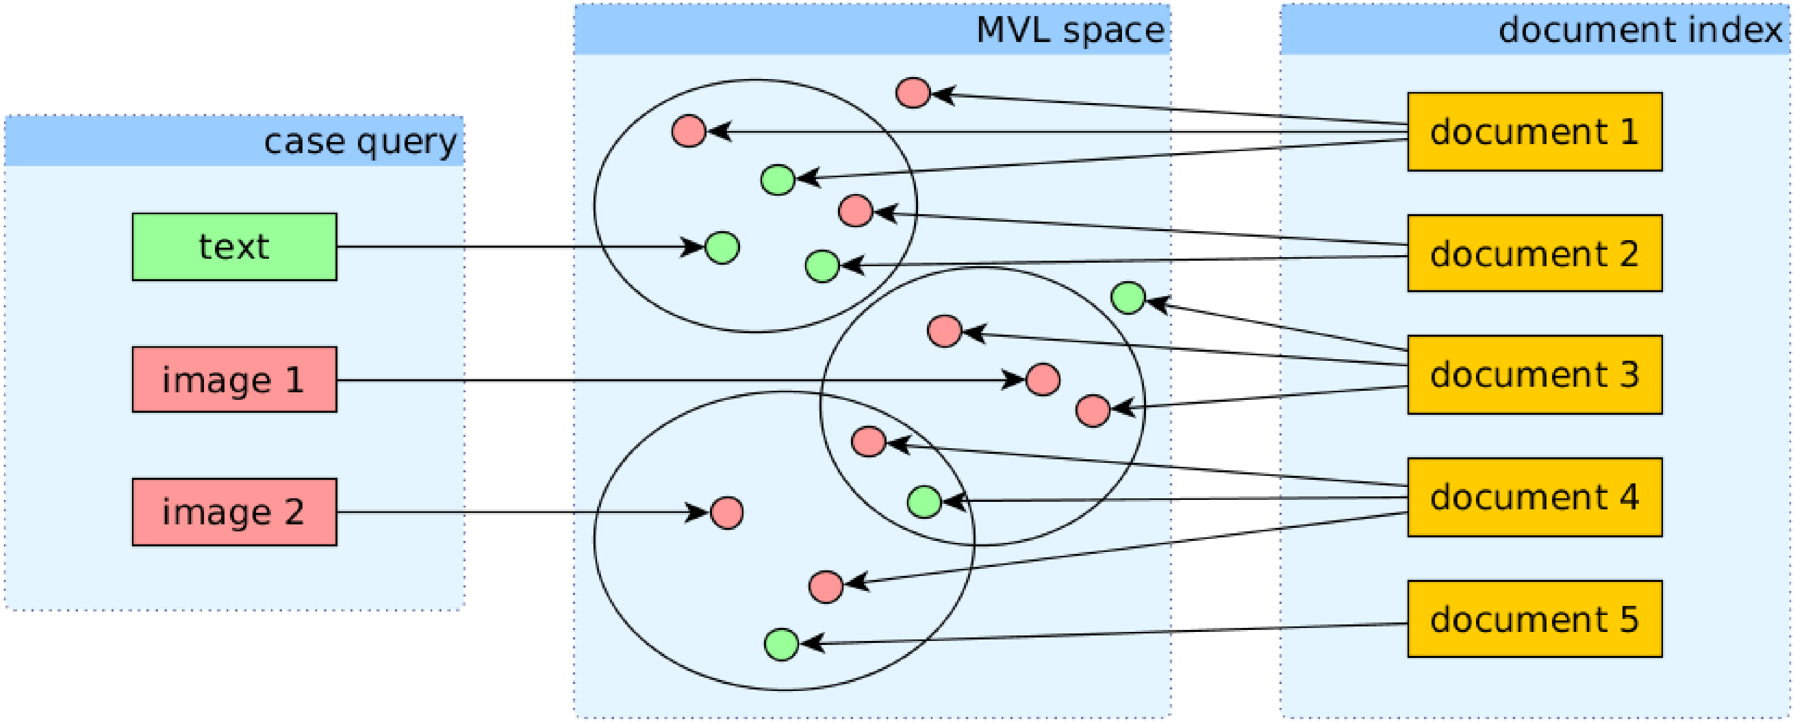

Conclusion

- Biomedical concepts can help to improve MCR over fulltext retrieval
  - Text-based query expansion increased MAP by 45%
  - Multimodal fusion with practical concept-based retrieval added another 13%
- There is room for future improvements of concept-based and multimodal techniques
  - Ideal concept-based retrieval and fusion improved MAP by 161% w.r.t. fulltext retrieval