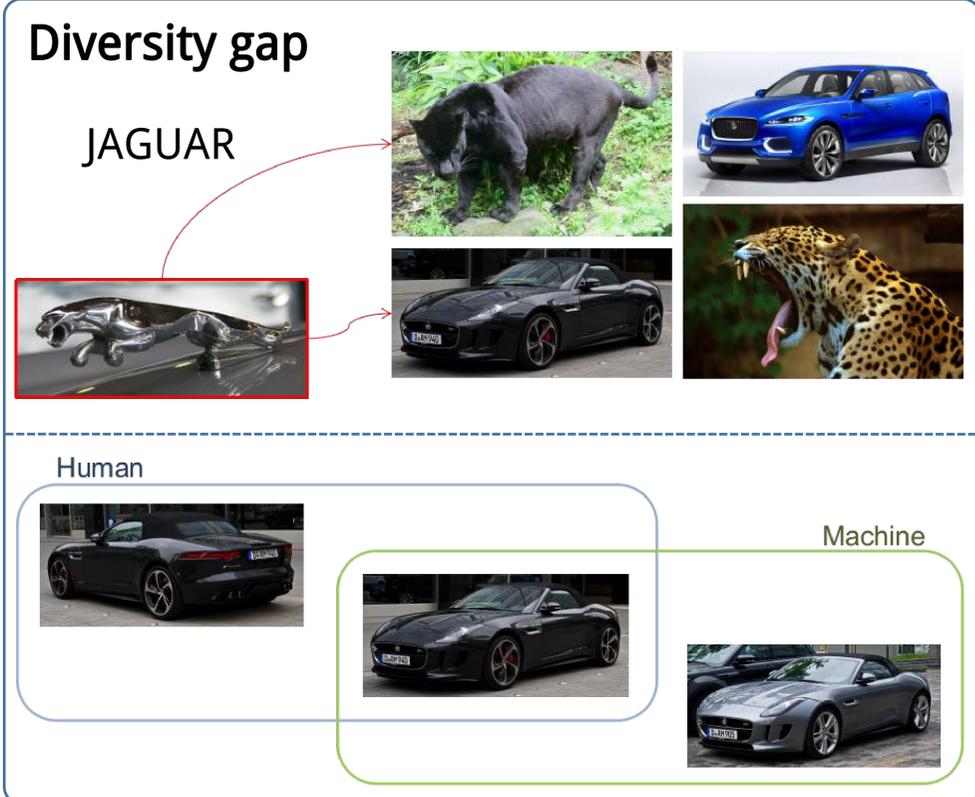
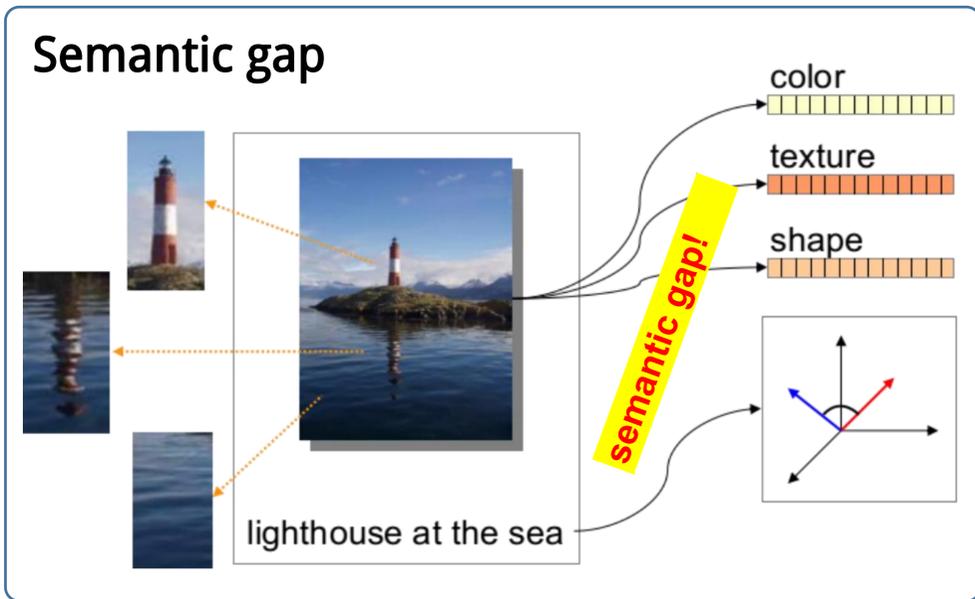


Abstract

The diversity promotion in retrieval results has been applied to tackle queries, which refer to multiple information needs, e.g., due to ambiguity. For instance, diverse results are achieved by selecting representative objects from clustered results [1]. Despite the effectiveness of diversity-aware methods, the image wealth of large collections and the subjectivity of human perceptions still bring the semantic gap problem. In turn, the difference between automatic clustering for diversification and human generated partitions define the diversity gap problem. This work proposes attenuating the diversity gap using multimodal approaches based on feature selection, fusion, and consensus.

Retrieval Challenges



Proposed Analysis

How to reduce the clustering errors and allow higher diversity?

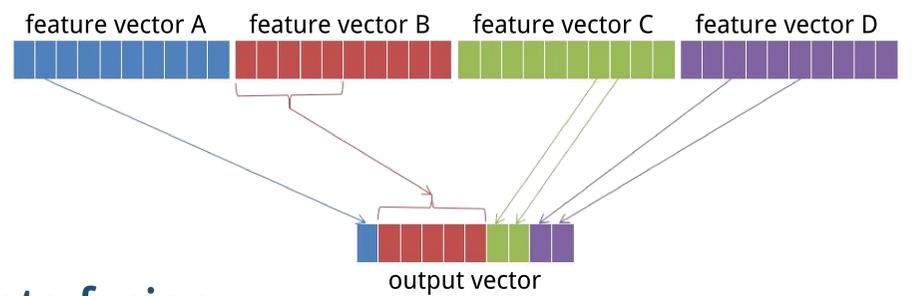
- Assess how automatic grouping relates to partitioning by human
- Reduce the diversity gap through data fusion
 - Early-fusion: Feature fusion and selection
 - Late-fusion: Clustering consensus

References

- [1] Calumby, R. T., Araujo, I. B. A. d. C., Santana, V. P., Munoz, J. A., Penatti, O. A., Li, L. T., Almeida, J., Chiachia, G., Goncalves, M. A., and Torres, R. d. S. (2015). Recod @ mediaeval 2015: Diverse social images retrieval. Working Notes of MediaEval.
- [2] Strehl, A. and Ghosh, J. (2002). Cluster ensembles—a knowledge reuse framework for combining multiple partitions. JMLR, 3(Dec):583–617.
- [3] Ionescu, B., Popescu, A., Lupu, M., Gînsca, A.-L., and Müller, Henning, B. B. (2015). Retrieving diverse social images at mediaeval 2015: Challenge, dataset and evaluation. In MediaEval.

Early-fusion

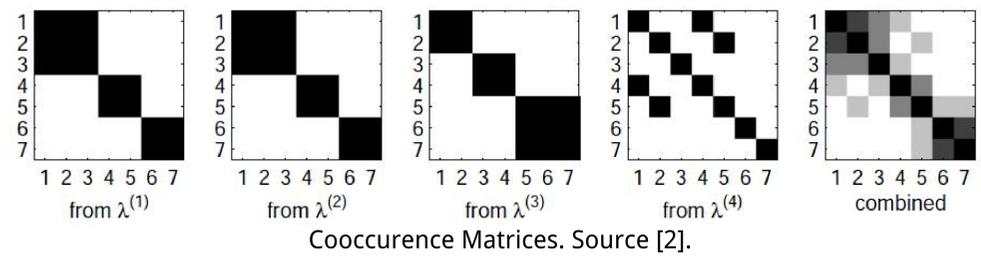
Feature fusion and selection using principal component analysis (PCA)



Late-fusion

Ensemble clustering (Consensus)

Cluster-based Similarity Partitioning Algorithm (CSPA)



Experimental Setup

Dataset: Retrieving Diverse Social Images Task [3]

- 153 queries with ~300 images/query
- Metadata: title, description and tags (fused)

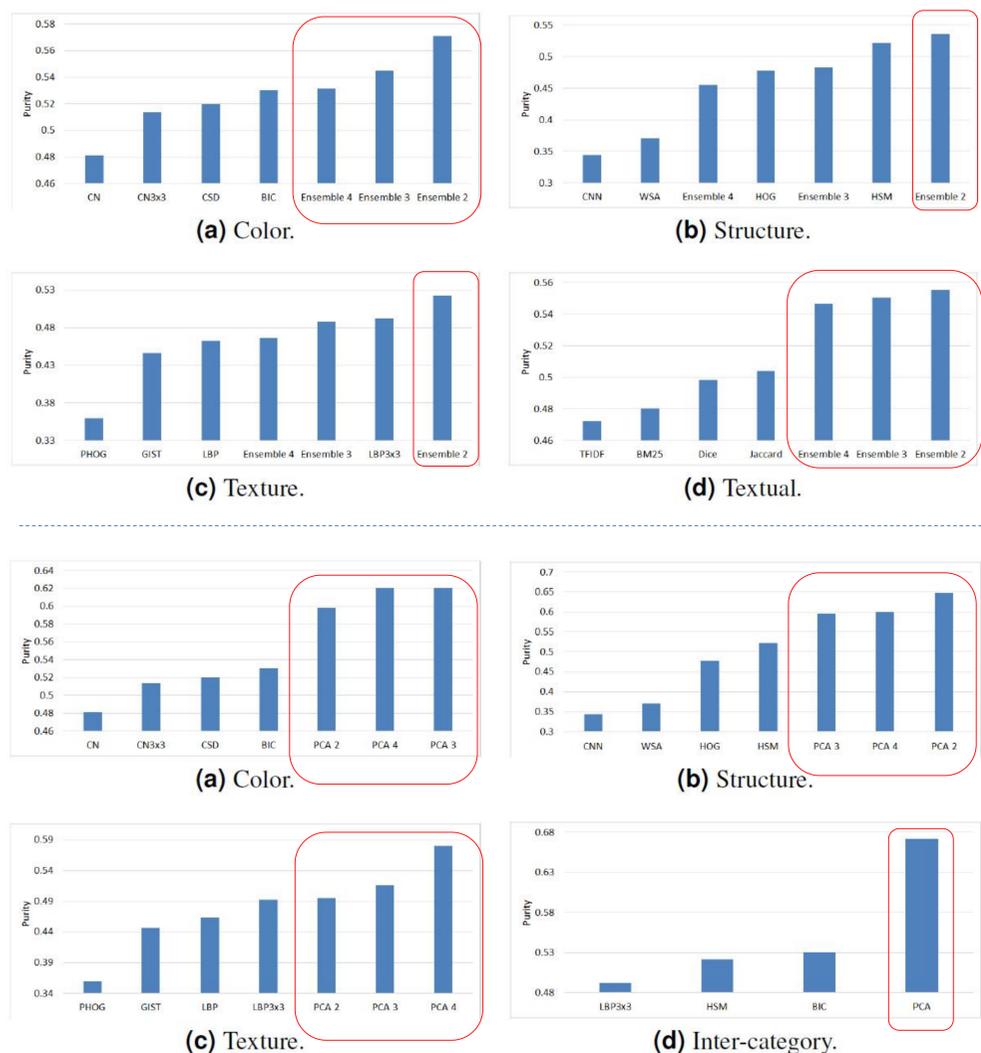
Similarity measures

- Visual: Color (13), Structure (5), and Texture (23)
- Textual: Cosine, BM25, Dice, Jaccard, and TF-IDF

Clustering Algorithms: k-Medoids, Single-link, Complete-link, Average-link, BIRCH, and Chameleon.

Evaluation: Purity, Maximum matching, and F-measure.

Results



Conclusions

Both early-fusion and late-fusion approaches consistently allowed the construction of clusterings more similar to the human-based partitioning. The experiments demonstrate the potential of these approaches to boost diversity-aware engines.