Modeling Representation Uncertainty in Concept-Based Multimedia Retrieval

Robin Aly
University Twente
r.aly@ewi.utwente.nl

18.10.2010

Abstract

Representing multimedia documents by means of concepts labels attached to parts of these documents has great potential for improving retrieval performance. The reason is that concepts are independent from how users refer to them and from the modality in which they occur. For example, a Flower and une Fleur refers to the same concept and a singing bird can appear in an image or an audio recording. The question whether a concept occurs in a multimedia document is answered by a concept detector. However, as building concept detectors is difficult the current detection performance is low which causes the retrieval engine to be uncertain about the actual document representation.

This thesis proposes the Uncertain Document Representation Ranking (URR) Framework which deals with this uncertainty by transferring the principles of the Portfolio Selection Theory in finance where the future win of a share is uncertain to the concept-based retrieval problem. Similarly to the distribution of future wins, the retrieval framework considers *multiple* possible concept-based document representations for each document resulting in *multiple* possible scores, which is the main scientific contribution of this thesis. Given an existing retrieval function for a certain representation, documents are ranked by the expected score plus an expression of the score's variance.

From the general URR framework, we derive ranking models for shot and video segment retrieval. The shot retrieval and the video segment model re-use the probability of relevance and a language modeling ranking function respectively, basing themselves on decades of text retrieval research. We show in experiments that the models significantly improve performance over several strong baselines in five TRECVid collections.

Furthermore, current performance of concept-based multimedia retrieval is low. A major reason for this is the performance of the concept detectors on which the simulation is based. Therefore, we predict the influence of improved concept detectors on general concept-based retrieval performance using Monte Carlo simulations. We find that more effort is needed to improve concept detectors, but it is realistic for concept-based retrieval to reach performance suitable for large-scale, real-life applications in the future.

Available at `http://doc.utwente.nl/72019/` (or `http://robin.aly.de`, with errata)