

Exercises for Multimedia Databases

Sheet 10 (until 04.02.2010)

Please note that you need **50%** of all exercise points to be admitted for the final exams. Exercises have to be turned in until **Thursday** before the next lecture and should be completed in teams of two students each. Write both names and “Matrikelnummer” on each page. If you have multiple pages, staple them together! Please hand in your solutions on **paper** into the mailbox at the IFIS floor or to our secretary (Mühlenpfordtstraße 23, 2nd floor). You may answer in either German or English.

Possible exam questions:

Exercise 1: Audio Retrieval (17P)

- Two important characteristics of sound are the frequency and the amplitude. What influence do these features have on human perception of sound? (1P)
- What is sampling? (1P)
- In which scenarios are audio databases typically used? (1P)
- Enumerate and explain 3 low-level features used in audio retrieval. (2P)
- A central audio feature is the pitch. What is the pitch? What are the theoretical models which explain the transformation from sound waves into the perceived pitch? (4P)
- How can we automatically determine the pitch? Please differentiate between the pitch and the fundamental frequency. (2P)
- What is “query by humming” and what are the problems one needs to solve in implementing such a system? (2P)
- Which distance measure is typically used for comparing note sequences? (2P)
- What is a Hidden Markov Model? Why is it useful for audio retrieval? (2P)

Exercise 2: Video Retrieval (6P)

- Describe the general structure of a video. (1P)
- Why is shot detection useful? Which threshold based methods for shot detection were presented in the lecture and how do they work? (2P)
- How can we establish the similarity between 2 videos? (1P)
- What approaches are there for result presentation in the case of video retrieval? (2P)

Exercise 3: Indexing (9P)

- How does a B-Tree work? Can we use it for indexing multimedia data? Why? (2P)
- What is the structure of an R-Tree? (1P)
- Why do we need a split operation in R-Trees? What is the linear method for splitting? (2P)

- d) Why should we keep the growth volume to a minimum when inserting new objects?
(2P)
- e) How can we index a collection of geometric shapes represented through chain codes? (2P)