

Grundlagen des IR: MIR Exercise

<http://www.ifs.tuwien.ac.at/mir>

Andreas Rauber

**Department of Softwaretechnology and
Interactive Systems**

Vienna University of Technology

<http://www.ifs.tuwien.ac.at/~andi>

Exercise MIR

- Work in groups of 2
- Download the music collection and feature sets from
http://kronos.ifs.tuwien.ac.at/ISMIRgenre_mp3_44khz_64kbit_mono_30sec.zip
http://kronos.ifs.tuwien.ac.at/ISMIRgenre_featureSets.zip
 - Alternatively you can choose to use your own collection of at least comparable size, in which case, however, you will need to extract 30-sec segments as well as the audio features to be submitted with the assignment for evaluation purposes)

- **Sec 1a: Collection Analysis**

Try to identify sub-groups of related songs

- e.g. different genres/sub-genres, ... and describe them in your report (how many? which? Styles/characteristics?)
don't rely solely on the genre labels provided!
- Pick 3 query songs from the collection
- Pick 2 (preferably rather different) feature sets

Exercise MIR

■ Sec 1b: Query Song Characteristics

Take the query songs and for each

- describe its characteristics acoustically, rhythmically, genre, style,...
- describe which songs you would expect to be similar based on your exploration of the music collection so far, as well as which other songs that you know and that are not necessarily included in the given music collection, that you would expect/describe as being similar to each of these

Exercise MIR

Retrieval Experiment:

- Download the SOM Toolbox for performing the similarity retrieval from <http://www.ifs.tuwien.ac.at/dm/somtoolbox/download-releases.php> or use your own implementation for similarity retrieval
- For each of the query songs, retrieve the 100 most similar songs using the 2 different feature sets using Euclidean distance

Exercise MIR

■ Sec 3: Retrieval Experiments

For each query, **analyze** the top 1-5, the 10th, 15th, 20th, 50th, 100th song by listening to the results, **describe** the songs found and interpret the results:

- Include a table in the report and provide a Powerpoint / OpenOffice slide per query song per feature set with
 - Query Song & Feature Set & Distance Metric as table header
 - Rank, Song-ID, Distance, Rating, Reason, Analysis as fields, where
 - Rating: 1 (very closely matching) – 5 (not matching at all)
 - **Reason:** why you provide the specific rating for the song
 - **Analysis:** why you think the song was retrieved using the given feature set
- Following **each table**, provide a **short overall description/summary**:
 - How does the perceived similarity correlate with the computed similarity?
 - How similar are the songs? Does the ranking make sense?
 - How fast do they become “dissimilar” down the ranking list?

Exercise MIR

▪ Sec 4: Cross Feature-Set Comparison

For each pair of queries across two different feature sets provide a summary:

- **Analysis** per query: how do the two result lists compare?
- Which feature set is better suited for the given query?
- **Why** is this so? Which characteristics are captured by the given feature set?
- What can be said in general of the two feature sets

Exercise MIR

- **Sec 5: Summary**

Provide an overall summarizy:

- Which feature set best captures **your** perception of similarity?
- For which genres/styles?
- Which songs are not matching well? **Why?**
- Bonus: which combination of features and metrics seems to work best / worst?

- Describe your overall impression

Exercise MIR

To submit: zip/tgz/rar file with

- Report describing your findings wrt. tasks listed on previous slides (song characteristics, results per query song, summary per feature set / overall)
 - Follow the ACM formatting guidelines, using the templates provided at <https://www.acm.org/publications/proceedings-template>.
(Proceedings Style File: LaTeX2e - Strict Adherence to SIGS style)
LaTeX recommended, but Word/OpenOffice is also ok.
 - Put your group number, names and your student IDs in the report! (as author info)
- An MS Powerpoint or OpenOffice presentation file containing links to the query songs and retrieved results with ratings
(1 slide per query, presentation in same directory as audio files to ensure links are local)

Submission guidelines

- **Upload ONE [zip/tgz/rar] file to TUWEL that contains all your files** (all Report, Presentation, audio snippets in case you used your own collection) with the **report as a PDF file inside that zip file (no Word files, no TEX sources)**. You must follow this naming convention:
 - GIR2019_Ex2_group_<groupno>_<Matnr.1>_<Matnr.2>.zip
 - Example: A submission of group 99 with 2 students (ids: 00059999, 00039999) looks like this:
GIR2019_Ex2_group_99_00059999_00039999.[zip/tgz/rar]
 - Apply the same naming convention to the report (but obviously with pdf extension)

Similarity Retrieval – How To

Similarity Retrieval

- Documentation at

<http://www.ifs.tuwien.ac.at/dm/somtoolbox/reference.html#SimilarityRetrieval>

- Open a shell, go to the SOMToolbox directory

- cd C:\Program Files\SOMToolbox

- For GUI-Version, run command

```
somtoolbox[.bat|.sh] SimilarityRetrieval -gui
```

Similarity Retrieval

- Similarity retrieval parameters
 - v Feature input vector file (RP, RH, or SSD)
 - n Number of neighbours
 - l name of vector to find neighbours of
- **Example:** `somtoolbox.bat SimilarityRetrieval -v ISMIRgenre.rh -n 5 -v -l rock_pop_artist_97_album_1_track_3.mp3`

Nearest neighbours to 'rock_pop/rock_pop_artist_97_album_1_track_3.mp3', using distance metric: L2

Rank	InputName	Distance
1	jazz_blues/jazz_blues_artist_73_album_3_track_1.mp3	745.0025497659058
2	metal_punk/metal_punk_artist_78_album_1_track_1.mp3	836.2325986964626
3	jazz_blues/jazz_blues_1-like_white_on_rice.mp3	843.7657375996287
4	world/world_artist_122_album_1_track_2.mp3	851.2882748058057
5	rock_pop/rock_pop_artist_87_album_1_track_2.mp3	854.8449095309211



Similarity Retrieval

- For similarity retrieval via GUI run:
somtoolbox.bat SimilarityRetrievalGUI

The screenshot shows the 'Similarity Retrieval GUI' interface. Key elements include:

- Feature files:** A path is entered in the 'Feature files' field, and the 'Load' button is circled in red.
- Options:** The '# to retrieve' is set to 5. The 'Query vector' is a long alphanumeric string. The 'Distances' section has 'relative' selected. The 'Distance metric' is set to 'L2Metric'. The 'Start' button is circled in red.
- Results:** A table with 5 rows showing Rank, Vector Label, and Distance.
- Database Details:** A list of filenames on the right side of the window.

Rank	Vector Label	Distance
1	C:\Users\Andi\ehre\ir_ethiopia\software\mini_newsgroups_SOM\mini_newsgroups\Calt.atheism\C51126	72.6%
2	C:\Users\Andi\ehre\ir_ethiopia\software\mini_newsgroups_SOM\mini_newsgroups\Calt.atheism\C51127	75.58%
3	C:\Users\Andi\ehre\ir_ethiopia\software\mini_newsgroups_SOM\mini_newsgroups\Calt.atheism\C51305	78.609%
4	C:\Users\Andi\ehre\ir_ethiopia\software\mini_newsgroups_SOM\mini_newsgroups\Calt.atheism\C53235	84.325%
5	C:\Users\Andi\ehre\ir_ethiopia\software\mini_newsgroups_SOM\mini_newsgroups\Calt.atheism\C51131	86.931%



Similarity Retrieval

- You can also normalize the feature vectors:
somtoolbox.bat SOMLibVectorNormalization [--gui]
- -m <method> Normalization method
 - UNIT_LEN normalises the vectors in the input file to unit length
 - MIN_MAX normalises each attributes between 0 and 1
 - STANDARD_SCORE normalises each attribute to a mean of 0,
and a max value of the standard deviation
- default: UNIT_LEN
- <input> Name of input vector file to be read
- <output> Name of new vector file to be created
- -G, --gui Show the graphical interface for additional options
- --help Print this help and exit
- --version Print the version and exit