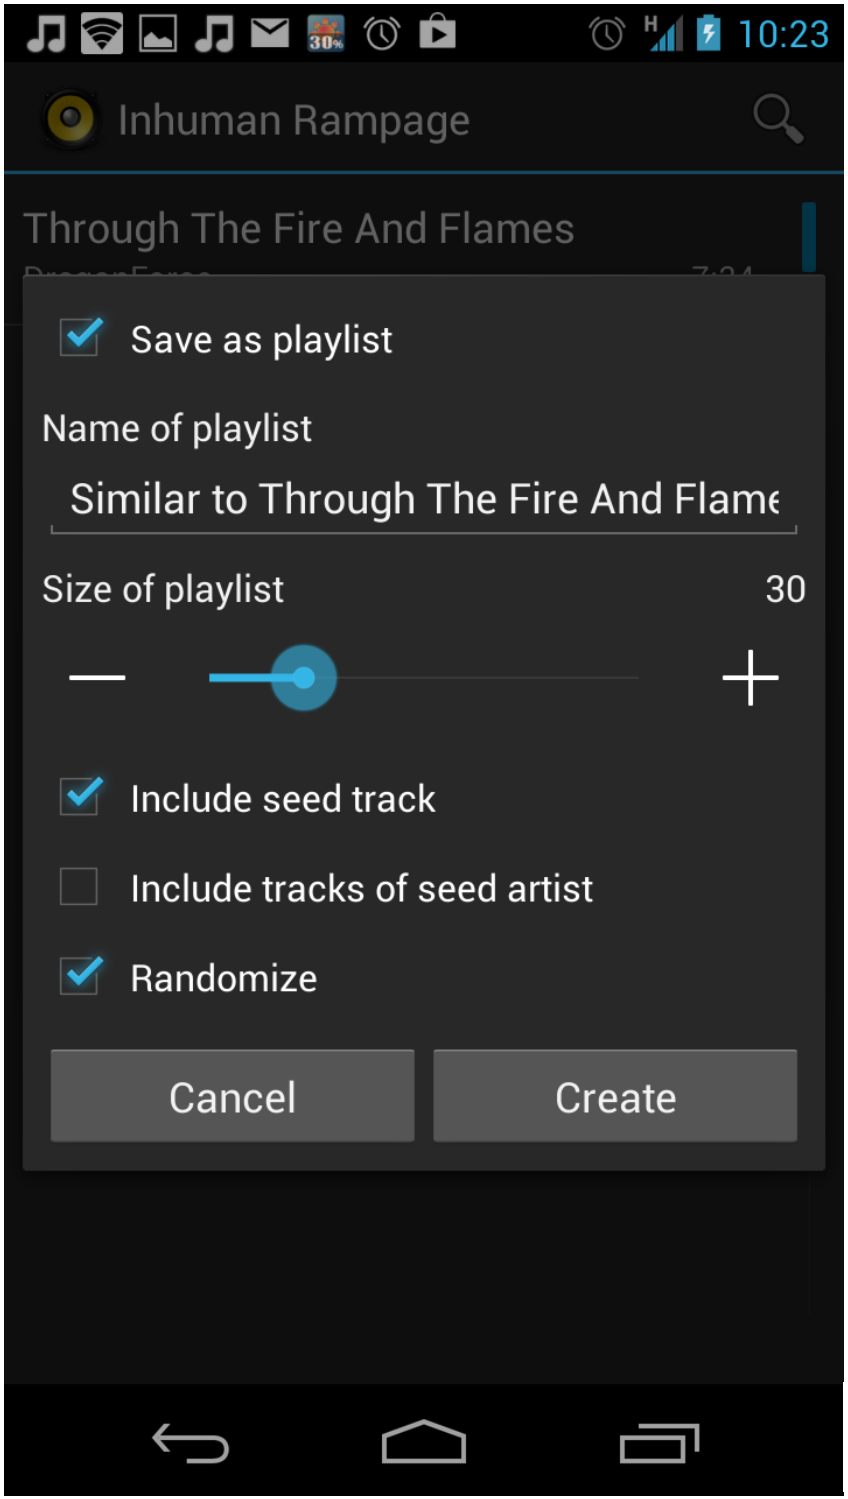

Mobile Music Genius

Automatic playlist generation based on music context (features and similarity computed based on Last.fm tags)

Department of Computational Perception

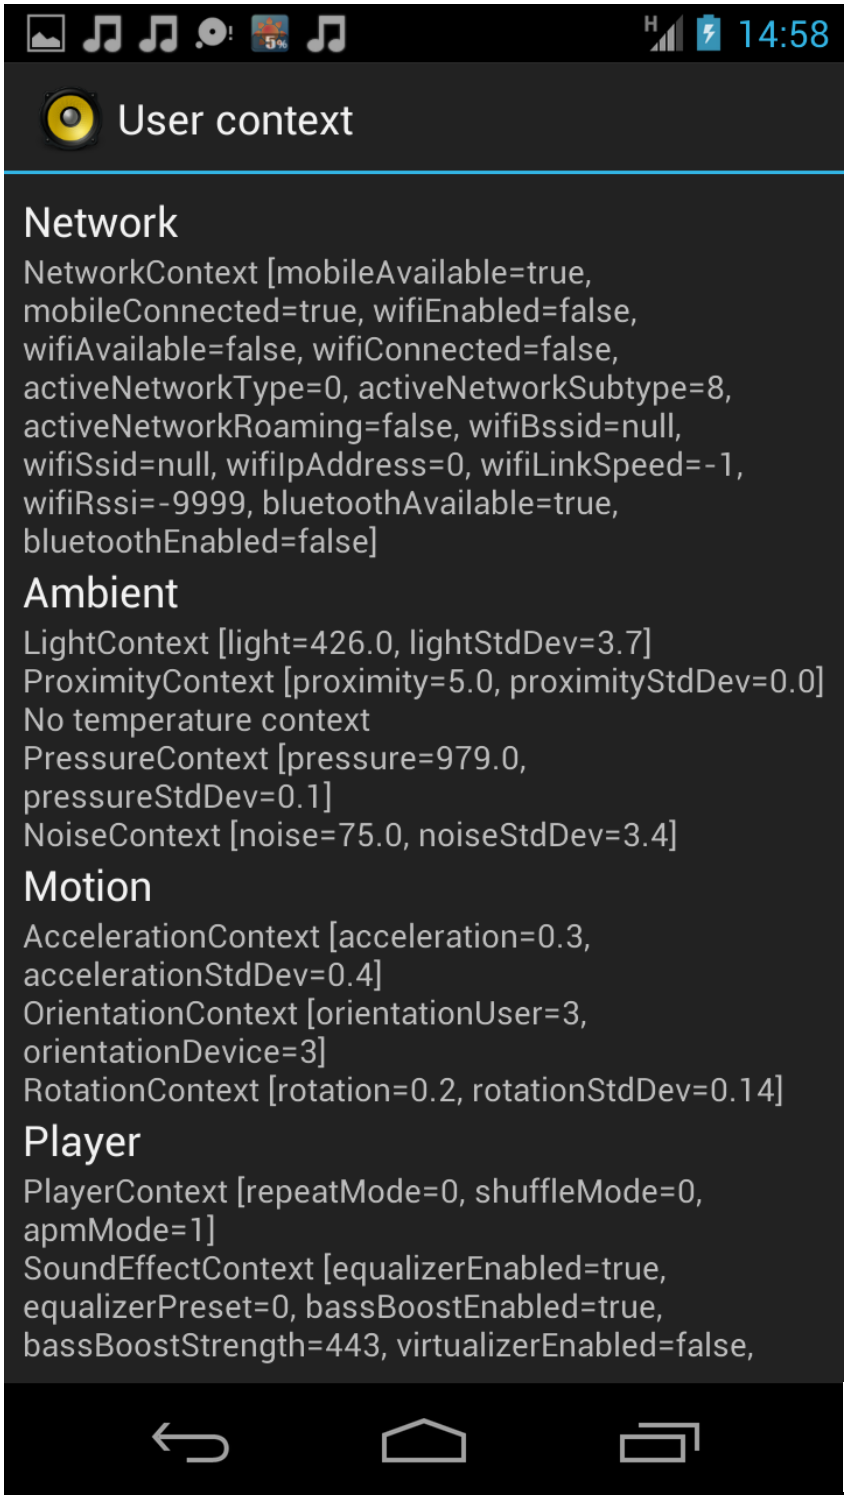

Mobile Music Genius

Some user context features gathered while playing

Department of Computational Perception

User Context Features from Android Phones

Time: timestamp, time zone

Personal: userID/eMail, gender, birthdate

Device: devideID (IMEI), sw version, manufacturer, model, phone state, connectivity, storage, battery, various volume settings (media, music, ringer, system, voice)

Location: longitude/latitude, accuracy, speed, altitude

Place: nearby place name (populated), most relevant city

Weather: wind direction, speed, clouds, temperature, dew point, humidity, air pressure

Ambient: light, proximity, temperature, pressure, noise, digital environment (WiFi and BT network information)

Activity: acceleration, user and device orientation, screen on/off, running apps

Player: artist, album, track name, track id, track length, genre, plackback position, playlist name, playlist type, player state (repeat, shuffle mode),
audio output (headset plugged)

mood and activity (direct user feedback)

Department of
Computational
Perception

Preliminary Evaluation

- collected user context data from 12 participants over a period of 4 weeks
- age: 20-40 years, gender: male
- user context vectors recoded whenever a “sensor” records a change
- 166k data points
- assess different classifiers (Weka) for the task of predicting artist/track/genre/mood given a user context vector: k-nearest neighbor (kNN), decision tree (C4.5), Support Vector Machine (SVM), Bayes Network (BN)
- cross-fold validation (10-CV)

To be analyzed:

(i) Which granularity/abstraction level to choose for representation/learning?

(ii) Which user context features are the most important to predict music preference?

Preliminary Evaluation: Results

(i) Which granularity/abstraction level to choose for representation/learning?

Predicting class *track*

Results barely above baseline.

Predicting particular tracks is hardly feasible with the amount of data available.

Dataset	0-R	KNN	C4.5	SVM	BN	Max.rel.	Avg.rel.
Time	1.13	1.26	0.90	1.30	1.31	116.04%	105.84%
Location	1.13	1.40	1.57	1.42	1.58	139.76%	132.06%
Location - state	1.13	1.36	1.69	0.96	0.82	150.26%	107.28%
Location - place	1.13	1.31	1.47	1.46	2.23	197.49%	143.46%
Weather	1.13	1.17	0.91	1.19	1.07	105.25%	96.21%
Ambient	1.13	0.79	0.63	1.08	1.12	98.97%	79.99%
Ambient - no n.	1.13	0.64	0.63	0.97	1.10	97.49%	73.97%
Ambient - noise	1.13	0.45	0.67	1.28	1.11	113.38%	77.77%
Motion	1.13	0.54	0.97	1.06	1.32	117.15%	86.25%
Motion - acc.	1.13	0.58	0.58	1.39	1.10	123.50%	80.75%
Motion - orient.	1.13	1.09	1.33	0.94	1.41	124.76%	105.78%
Task	1.13	1.43	1.96	1.57	1.73	173.61%	148.36%
Task - display	1.13	1.75	1.68	1.76	1.76	156.47%	154.21%
Task - tasks	1.13	1.16	1.60	1.13	1.53	141.76%	120.03%
Phone	1.13	1.12	0.97	0.70	0.99	99.41%	83.85%
Network	1.13	1.43	1.34	1.26	1.82	161.79%	129.88%
Network - state	1.13	1.31	1.75	1.58	1.82	161.79%	143.27%
Network - env.	1.05	1.79	1.45	1.44	1.08	170.20%	137.07%
Device	1.13	1.07	1.56	1.12	1.24	138.14%	110.74%
Device - battery	1.13	0.71	1.12	1.23	1.12	109.39%	92.78%
Device - storage	1.13	0.95	1.07	1.44	1.42	127.49%	108.09%
Device - memory	1.13	0.92	0.79	1.24	1.30	115.59%	94.46%
Device - audio	1.13	0.46	0.63	0.96	1.30	114.93%	74.26%
Player	1.13	1.29	1.36	1.35	1.35	120.77%	118.46%
All	1.13	0.90	1.78	1.14	1.14	158.02%	110.05%

Preliminary Evaluation: Results

(i) Which granularity/abstraction level to choose for representation/learning?

Predicting class *artist*

Best results achieved, significantly outperforming baseline.

Relation {context → artist} seems to be predictable.

Dataset	0-R	KNN	C4.5	SVM	BN	Max.rel.	Avg.rel.
Time	28.54	60.83	57.10	59.68	58.70	213.15%	207.01%
Location	28.54	42.69	41.42	37.80	40.04	149.58%	141.86%
Location - state	28.54	41.71	41.83	33.11	37.05	146.55%	134.64%
Location - place	28.54	35.74	36.99	36.07	36.28	129.62%	127.09%
Weather	28.54	63.46	63.25	56.06	61.34	222.35%	213.84%
Ambient	28.54	34.70	36.83	31.17	35.18	129.03%	120.77%
Ambient - no n.	28.54	33.54	34.87	31.43	34.46	122.19%	117.65%
Ambient - noise	28.54	26.12	30.55	28.75	29.81	107.04%	100.94%
Motion	28.54	35.08	36.10	37.14	35.11	130.15%	125.65%
Motion - acc.	28.54	26.54	27.87	28.93	28.62	101.36%	98.07%
Motion - orient.	28.54	36.22	35.63	36.54	35.17	128.02%	125.75%
Task	28.54	60.75	60.65	59.63	56.20	212.86%	207.81%
Task - display	28.54	28.12	28.31	28.62	28.34	100.29%	99.33%
Task - tasks	28.54	61.35	61.28	60.28	55.23	214.97%	208.60%
Phone	28.54	37.30	38.74	31.33	33.74	135.74%	123.61%
Network	28.54	36.38	36.44	37.93	34.87	132.90%	127.56%
Network - state	28.54	34.95	33.14	34.58	34.17	122.45%	119.86%
Network - env.	21.90	25.01	26.42	27.43	22.69	125.26%	115.92%
Device	28.54	70.42	68.68	54.95	65.31	246.76%	227.20%
Device - battery	28.54	39.10	47.15	36.41	46.02	165.23%	147.76%
Device - storage	28.54	61.17	60.37	40.96	57.92	214.33%	193.08%
Device - memory	28.54	39.22	40.56	32.11	36.53	142.10%	130.01%
Device - audio	28.54	47.92	47.71	41.42	42.76	167.90%	157.50%
Player	28.54	38.18	38.36	38.30	38.25	134.41%	134.10%
All	28.54	69.56	69.01	69.87	67.66	244.83%	241.86%

Preliminary Evaluation: Results

(i) Which granularity/abstraction level to choose for representation/learning?

Predicting class *genre*

Prediction on more general level than for artist.

Still genre is an ill-defined concept, hence results inferior to artist prediction.

Dataset	0-R	KNN	C4.5	SVM	BN	Max.rel.	Avg.rel.
Time	29.80	46.75	44.99	46.46	46.27	156.88%	154.76%
Location	29.80	32.92	34.17	34.45	32.05	115.61%	112.08%
Location - state	29.80	32.25	33.41	32.48	30.44	112.12%	107.87%
Location - place	29.80	29.75	32.54	32.38	32.45	109.19%	106.65%
Weather	29.80	49.68	50.61	43.77	46.70	169.83%	160.03%
Ambient	29.80	28.30	34.12	31.38	33.27	114.50%	106.61%
Ambient - no n.	29.80	31.52	33.39	31.42	33.34	112.04%	108.79%
Ambient - noise	29.80	23.38	29.92	29.67	29.77	100.40%	94.57%
Motion	29.80	32.23	34.34	34.56	34.39	115.98%	113.69%
Motion - acc.	29.80	25.67	28.55	30.50	30.41	102.35%	96.59%
Motion - orient.	29.80	34.49	35.22	34.28	34.39	118.21%	116.10%
Task	29.80	43.89	46.47	44.55	41.85	155.95%	148.29%
Task - display	29.80	28.57	29.04	28.78	28.78	97.44%	96.61%
Task - tasks	29.80	44.71	47.62	44.94	42.31	159.81%	150.66%
Phone	29.80	31.17	33.43	31.33	30.13	112.20%	105.77%
Network	29.80	32.31	31.96	33.93	31.73	113.85%	109.00%
Network - state	29.80	31.70	31.14	32.07	31.26	107.63%	105.85%
Network - env.	26.10	26.17	27.02	29.78	27.28	114.08%	105.58%
Device	29.80	49.65	50.03	43.16	48.00	167.88%	160.11%
Device - battery	29.80	31.58	38.03	33.42	35.85	127.61%	116.51%
Device - storage	29.80	47.76	47.55	37.25	46.56	160.29%	150.28%
Device - memory	29.80	30.79	36.87	31.76	36.60	123.73%	114.11%
Device - audio	29.80	40.19	41.12	38.16	37.02	137.99%	131.29%
Player	29.80	35.79	36.34	36.08	35.59	121.96%	120.65%
All	29.80	46.75	49.22	50.41	48.51	169.15%	163.50%

Preliminary Evaluation: Results

(i) Which granularity/abstraction level to choose for representation/learning?

Predicting class *mood*

Poor results as mood in music is quite subjective and hence hard to predict.

Which mood anyway: composers intention? mood expressed by performers? mood evoked in listeners?

Dataset	0-R	KNN	C4.5	SVM	BN	Max.rel.	Avg.rel.
Time	24.00	24.79	27.73	24.56	24.29	115.53%	105.59%
Location	24.00	23.27	23.89	25.05	24.62	104.38%	100.86%
Location - state	24.00	23.44	23.97	25.25	24.79	105.20%	101.51%
Location - place	24.00	21.99	23.99	23.80	23.67	99.94%	97.33%
Weather	24.00	25.13	27.05	27.86	25.39	116.07%	109.82%
Ambient	24.00	17.04	19.41	23.59	24.04	100.17%	87.58%
Ambient - no n.	24.00	21.14	23.18	23.87	24.00	100.00%	96.03%
Ambient - noise	24.00	16.70	21.38	23.79	23.96	99.83%	89.40%
Motion	24.00	19.88	26.54	24.78	24.65	110.56%	99.84%
Motion - acc.	24.00	20.86	22.75	24.32	23.96	101.34%	95.72%
Motion - orient.	24.00	23.99	27.82	24.99	24.65	115.91%	105.68%
Task	24.00	22.94	24.32	24.58	25.00	104.18%	100.87%
Task - display	24.00	24.45	24.58	24.97	24.88	104.06%	103.00%
Task - tasks	24.00	23.56	25.20	24.99	24.13	105.00%	101.95%
Phone	24.00	19.34	24.64	26.75	26.74	111.45%	101.52%
Network	24.00	22.81	24.20	23.92	24.28	101.17%	99.17%
Network - state	24.00	23.48	24.39	24.01	24.28	101.64%	100.17%
Network - env.	27.78	27.68	28.36	29.24	27.78	105.26%	101.74%
Device	24.00	21.45	24.72	25.79	24.86	107.46%	100.86%
Device - battery	24.00	16.09	26.31	23.94	24.06	109.64%	94.17%
Device - storage	24.00	25.57	26.69	25.36	24.48	111.19%	106.36%
Device - memory	24.00	13.92	21.39	23.59	23.81	99.22%	86.16%
Device - audio	24.00	26.33	26.50	25.48	24.43	110.43%	107.03%
Player	24.00	24.81	25.57	25.37	25.45	106.54%	105.41%
All	24.00	22.43	26.16	24.81	26.11	109.00%	103.66%

Preliminary Evaluation: Results

(ii) Which user context features are the most important to predict music preference?

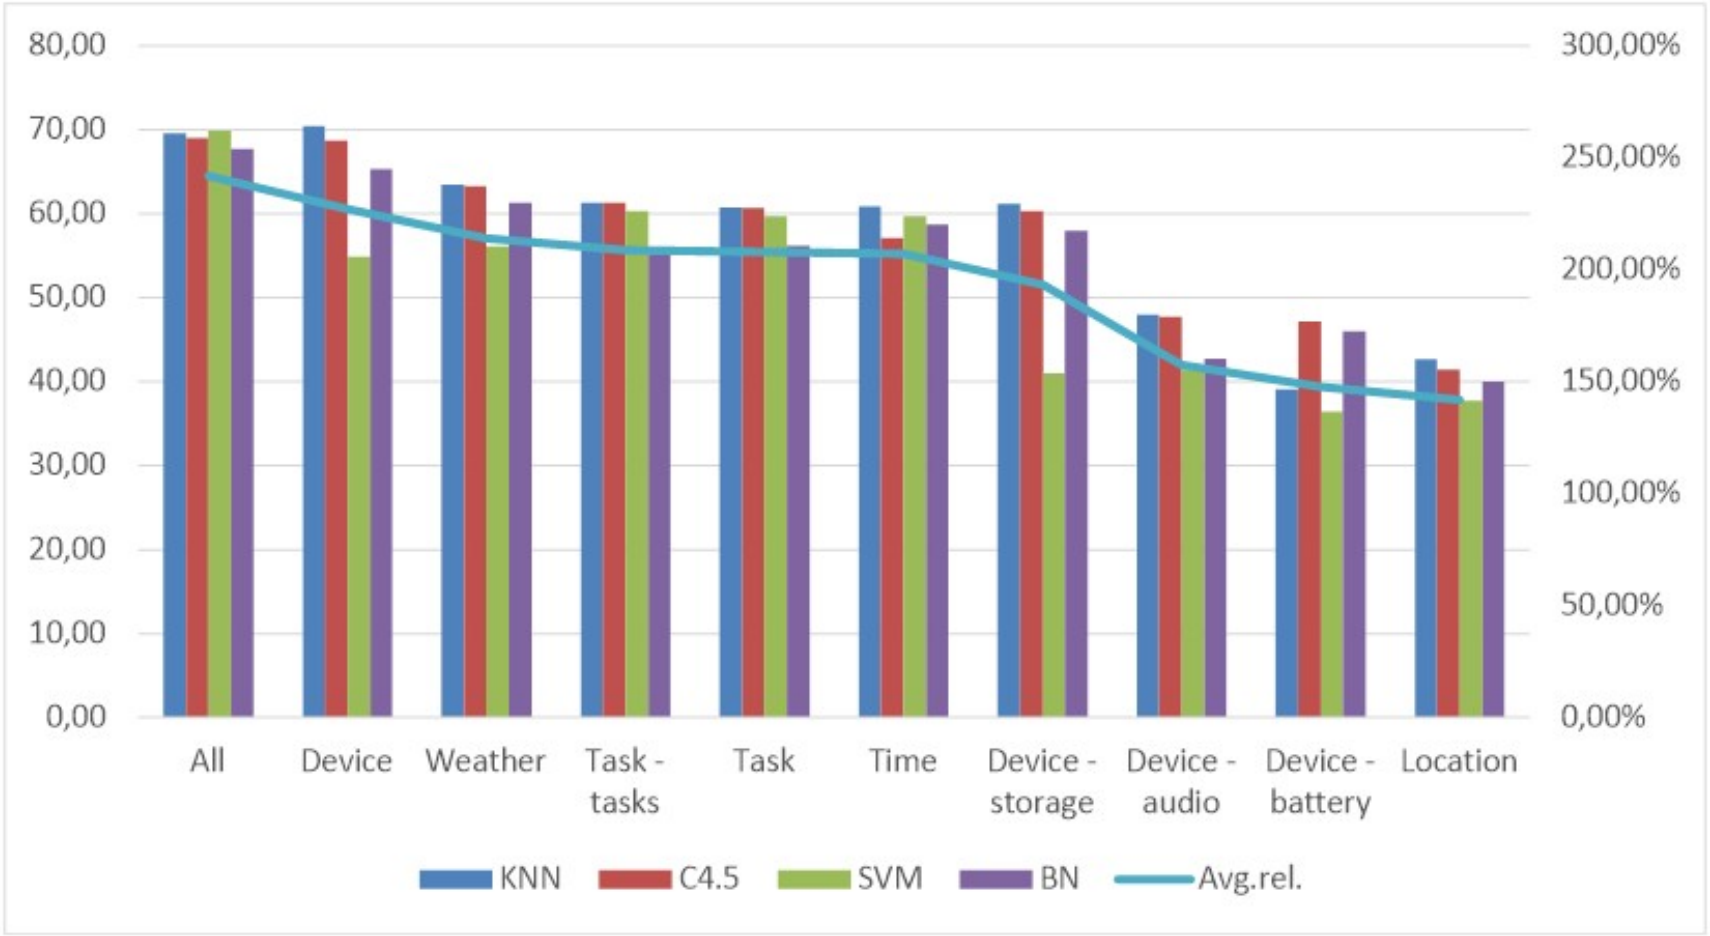

Making use of all features yields best results.

Preliminary Evaluation: Results

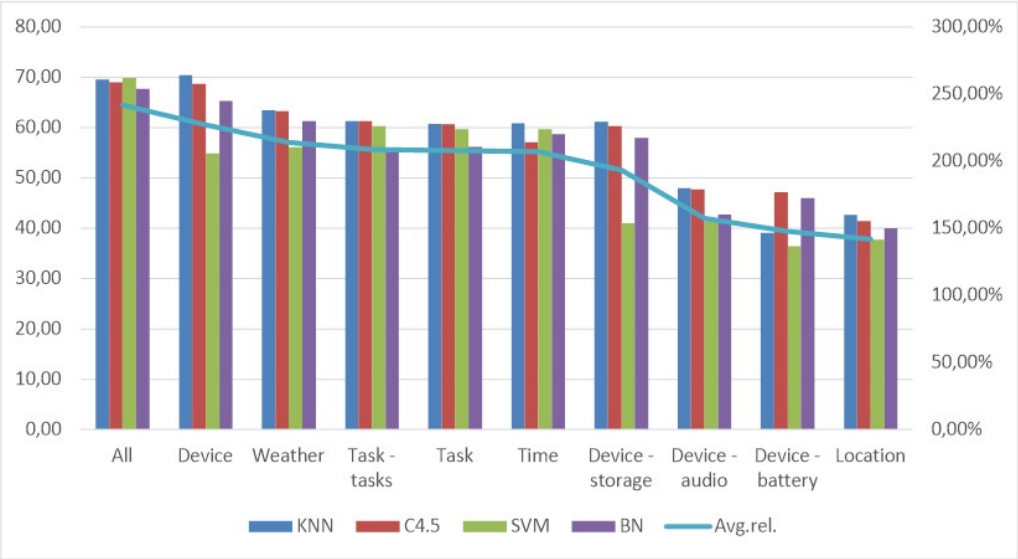

(ii) Which user context features are the most important to predict music preference?

Weka-feature selection confirms most important attributes:

time: weekday, hour of day

location: nearest populated place (better than longitude, and latitude)

weather: temperature, humidity, air pressure, wind speed/direction, and dew point

device: music and ringer volume, battery level, available storage and memory

task: running tasks/apps

Department of Computational Perception

Preliminary Evaluation: Results

- Problems:

– too little data to make significant predictions on the quality of the approach

→ need more data from more participants over a longer period of time

→ large-scale study

– dataset does not incorporate features potentially highly relevant to music listening inclination (user activity and mood)

Large-scale Evaluation

- collected user context data from JKU students over a period of 2 months
- about 8,000 listening data items and corresponding user context gathered

To be analyzed:

(i) How well does our approach perform to predict the preferred artist based on a given user context vector?

Results for predicting class “artist”:

ZeroR (baseline) classifier	15% accuracy
k-nearest neighbors	42% accuracy
JRip rule learner	51% accuracy
J48 decision tree	55% accuracy

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

recommend music that is suited to a place of interest (POI) of the user (context-aware)

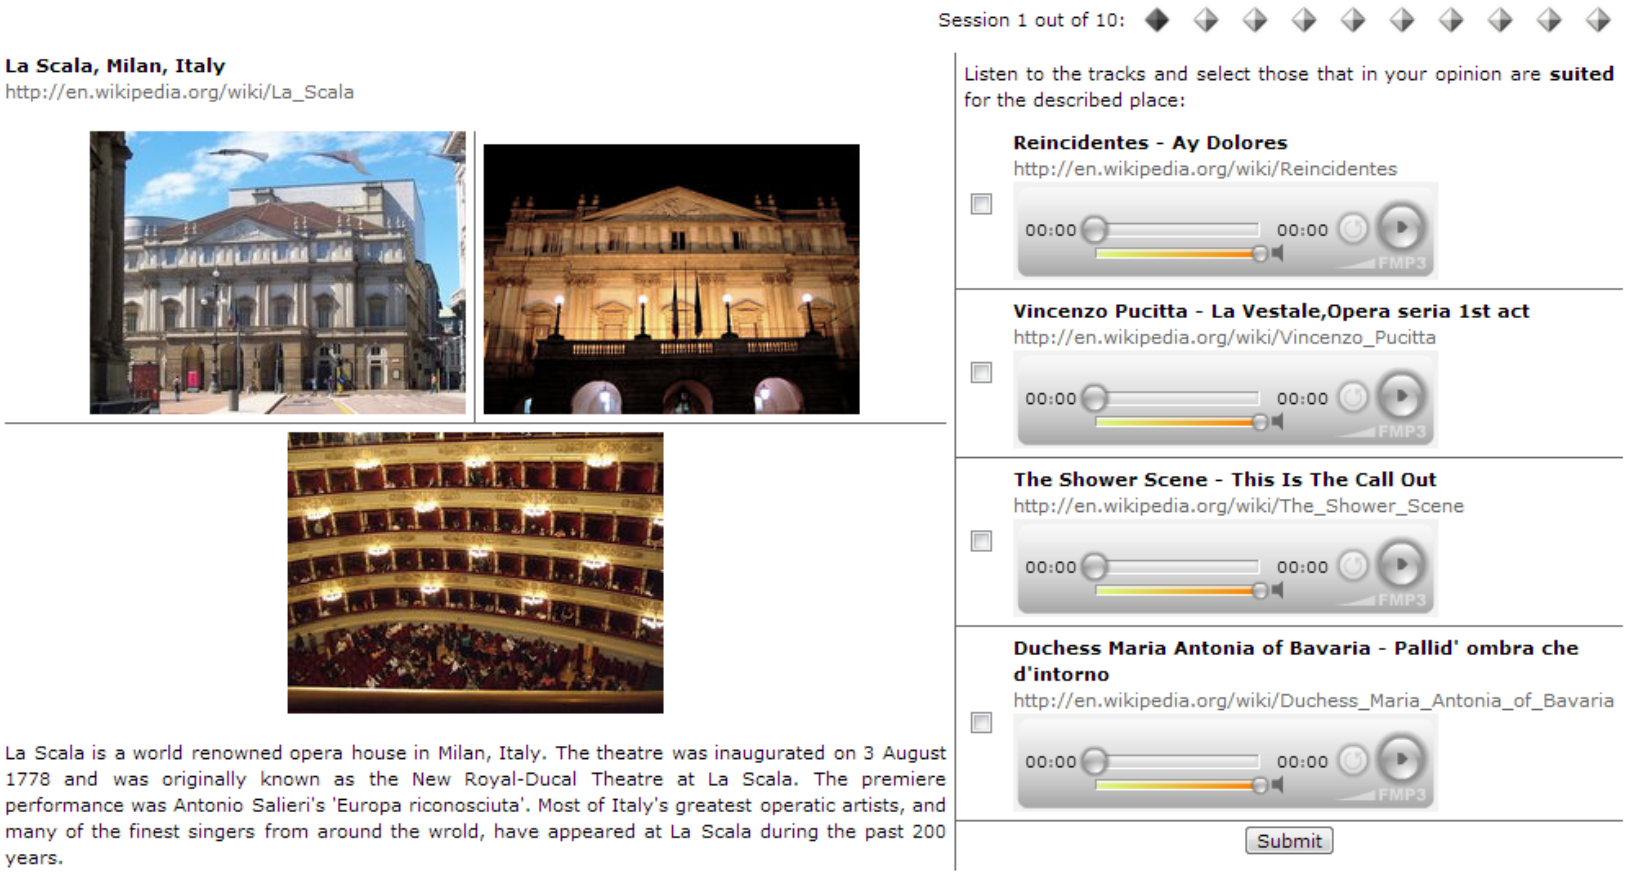

Department of Computational Perception

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *genre-based*: only play music belonging to the user's preferred genres (baseline)

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *knowledge-based*: use the DBpedia knowledge base (relations between POIs and musicians)

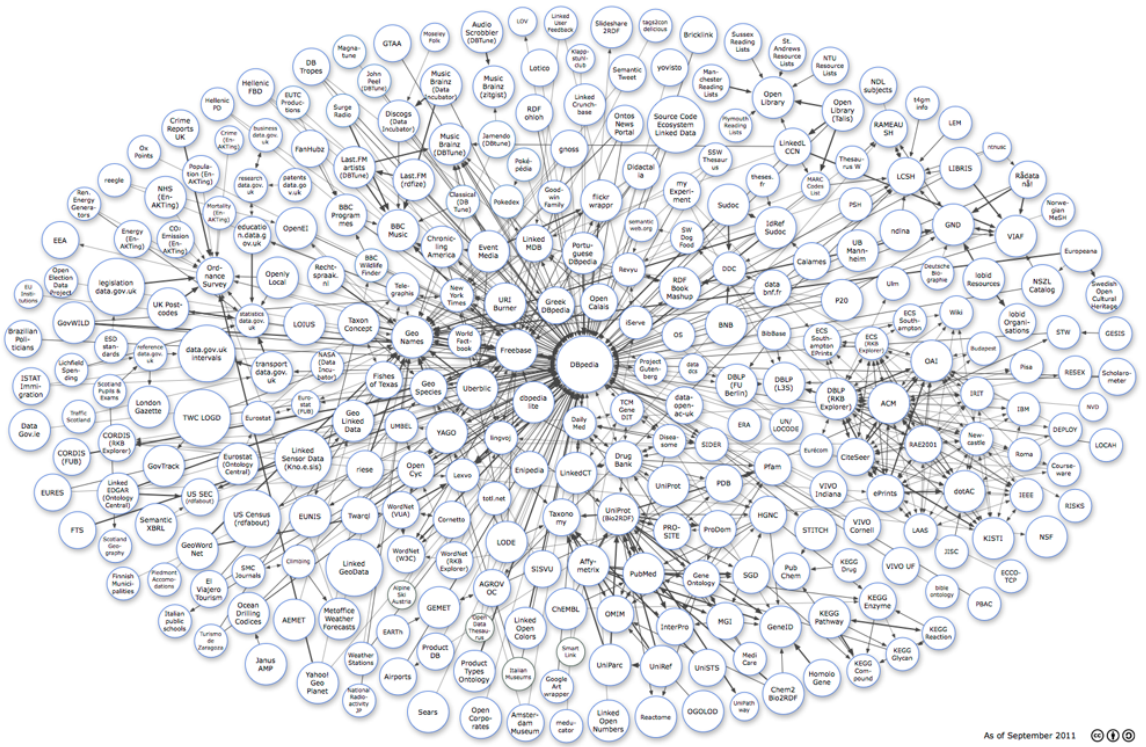

Department of Computational Perception

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *tag-based*: user-assigned emotion tags describing images of POIs and music, Jaccard similarity between music-tag-vectors and POI-tag-vectors

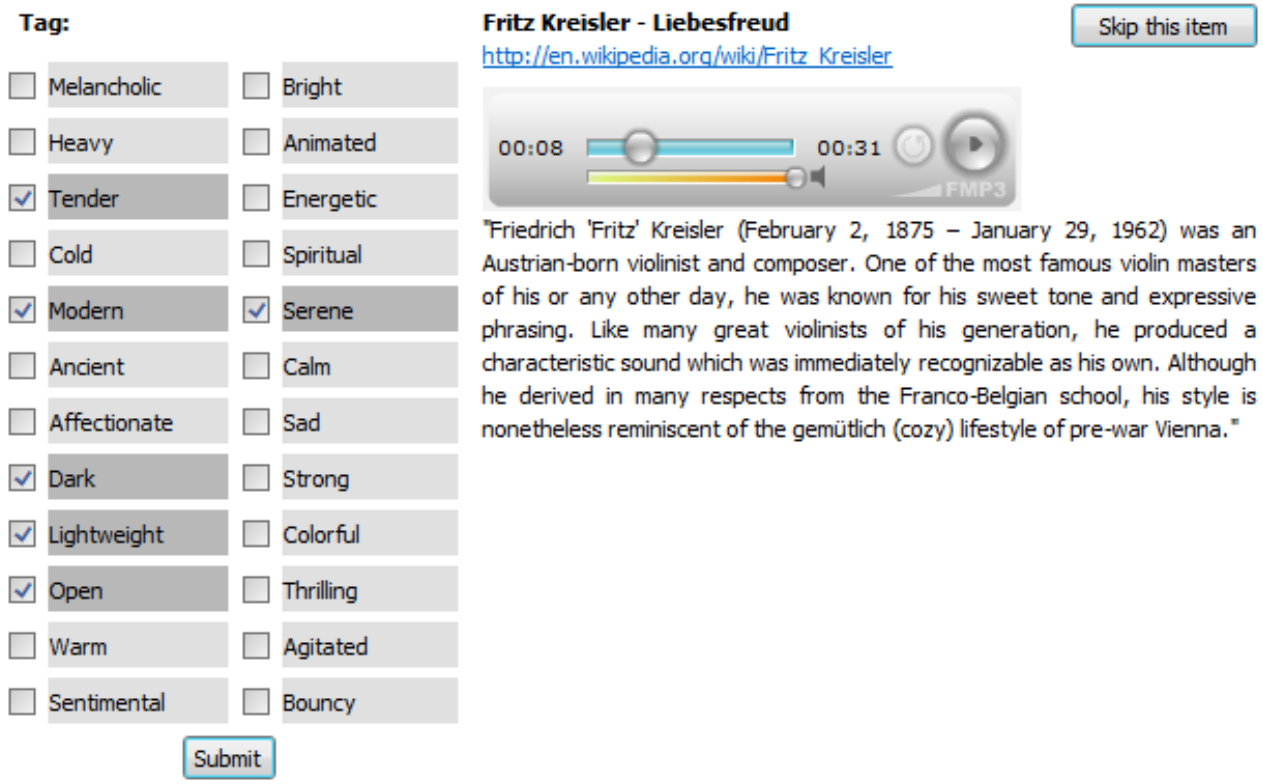

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *auto-tag-based*: use state-of-the-art music auto-tagger based on the Block-level Feature framework to automatically label music pieces; then again compute Jaccard similarity between music-tag-vectors and POI-tag-vectors

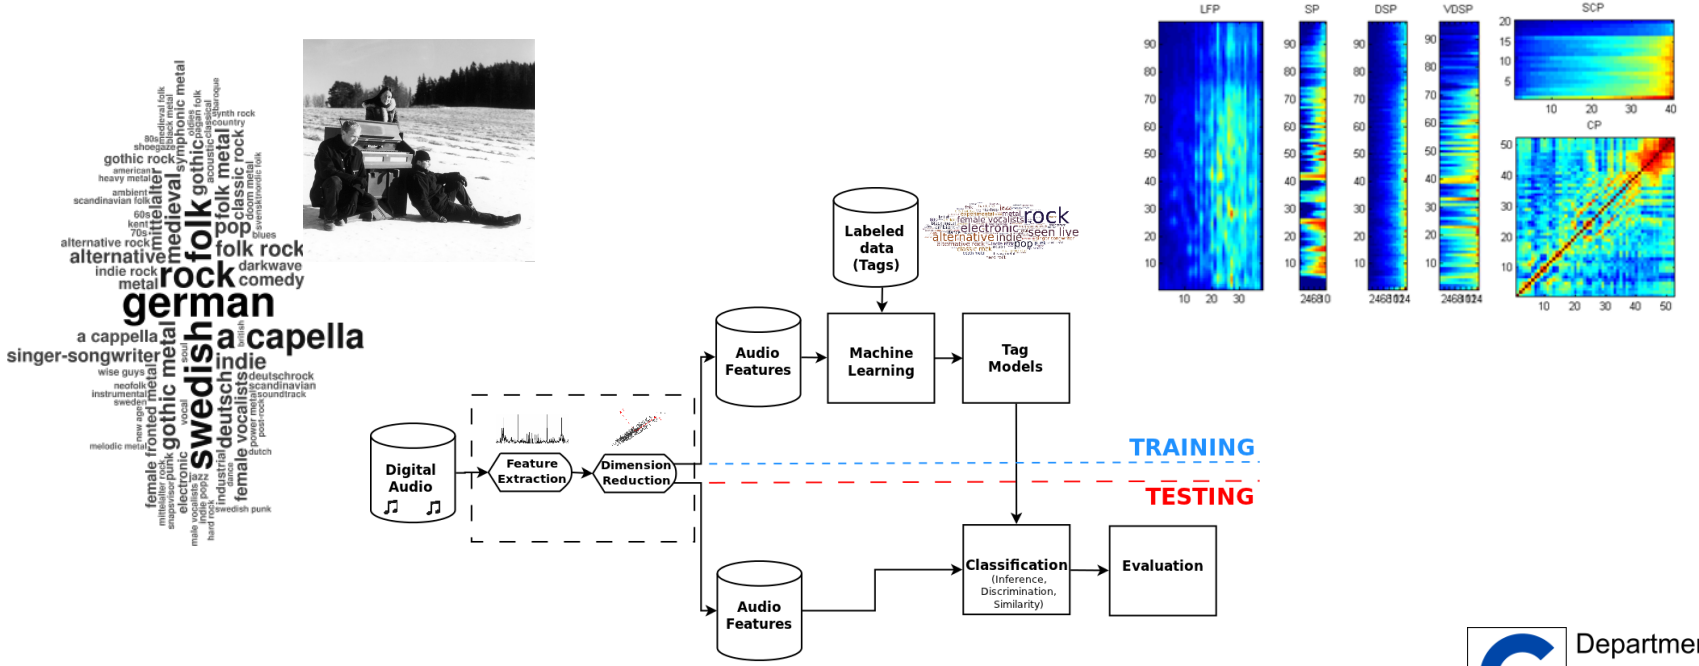

Department of Computational Perception

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *combined*: aggregate music recommendations w.r.t. ranks given by knowledge-based and auto-tag-based approaches

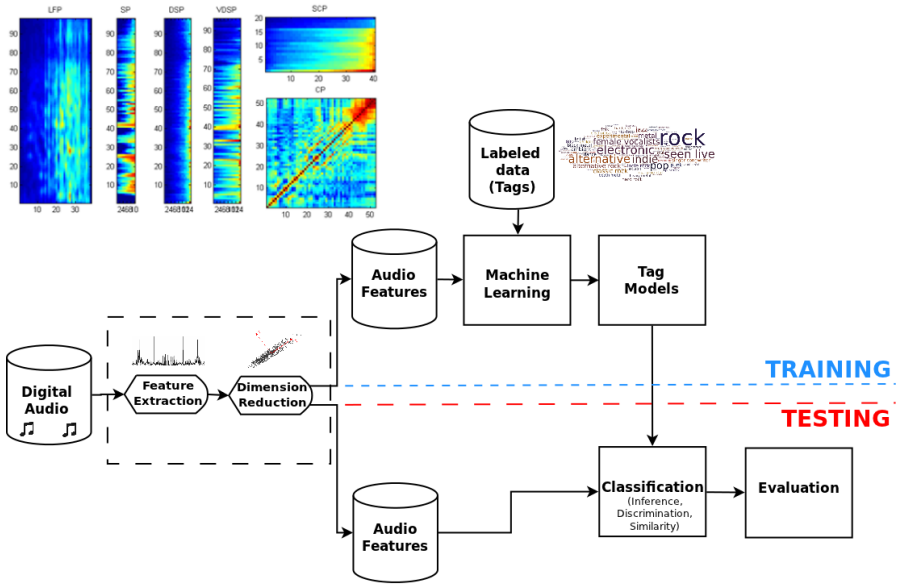

Department of Computational Perception

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Approaches:

- *genre-based*: only play music belonging to the user's preferred genres (baseline)
- *knowledge-based*: using the DBpedia knowledge base (relations between POIs and musicians)
- *tag-based*: user-assigned emotion tags describing images of POIs and music, Jaccard similarity between music-tag-vectors and POI-tag-vectors
- *auto-tag-based*: using state-of-the-art music auto-tagger based on the Block-level Feature Framework to automatically label music pieces; then again use Jaccard similarity between music-tag-vectors and POI-tag-vectors
- *combined*: aggregate music recommendations w.r.t. ranks given by knowledge-based and auto-tag-based approaches

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Evaluation:

- user study via web interface (58 users, 564 sessions)

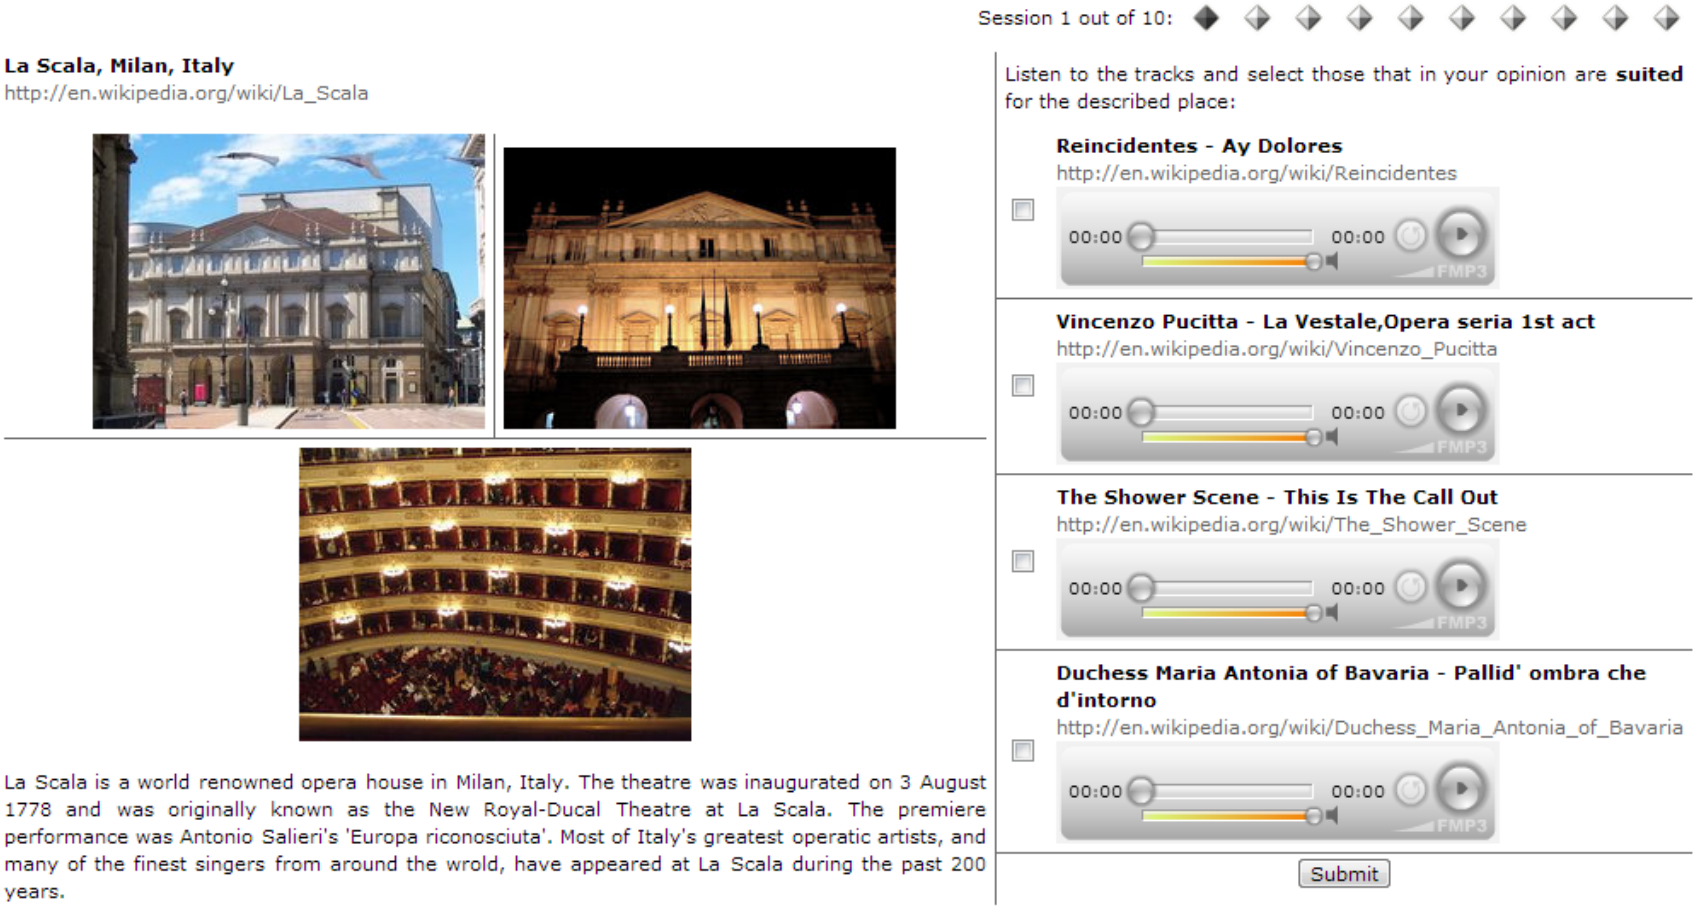

Matching Places of Interest and Music

(Kaminskas et al.; RecSys 2013)

Evaluation:

- Performance measure: number of times a track produced by each approach was considered as well-suited in relation to total number of evaluation sessions, i.e. probability that a track marked as well-suited by a user was recommended by each approach

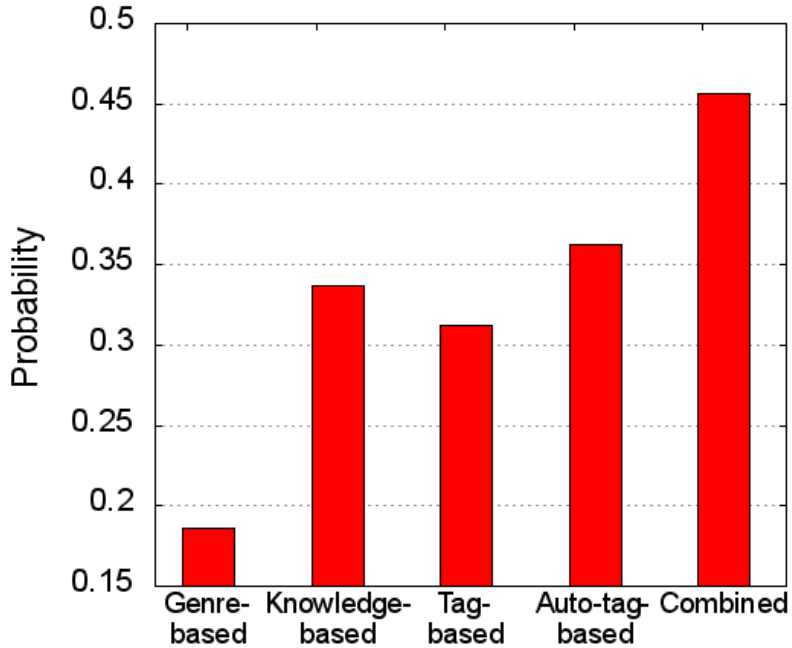

Department of Computational Perception

SUMMARY

Music Information Retrieval is a great field

Various approaches to extract information from the audio signal

Various sources and approaches to extract contextual data and similarity information from the Web

Multi-modal modeling and retrieval is important and allows for exciting applications

Next big challenges:

- modeling user properties and context
- improve personalization and context-awareness
- situation-based retrieval
- new and better suited evaluation strategies

Department of
Computational
Perception

Department of Computational Perception

JKU
JOHANNES KEPLER UNIVERSITY LINZ

References

Markus Schedl
Peter Knees

{markus.schedl, peter.knees}@jku.at

Department of Computational Perception
Johannes Kepler University (JKU)
Linz, Austria

(Arzt, Widmer; 2010) **Towards Effective 'Any-Time' Music Tracking**, Proceedings of the Starting AI Researchers' Symposium (STAIRS 2010), Lisbon, Portugal, 2010.

(Aucouturier, Pachet; 2004) **Improving Timbre Similarity: How High's the Sky?**, Journal of Negative Results in Speech and Audio Sciences, 1(1), 2004.

(Baccigalupo, Plaza, Donaldson; 2008) **Uncovering Affinity of Artists to Multiple Genres from Social Behaviour Data**, Proceedings of the 9th International Conference on Music Information Retrieval (ISMIR 2008), Philadelphia, September 2008.

(Baumann, Hummel; 2003) **Using Cultural Metadata for Artist Recommendation**, Proceedings of the 3rd International Conference on Web Delivering of Music (WEDELMUSIC 2003), Leeds, UK, September 2003

(Bertin-Mahieux, Eck, Maillet, Lamere; 2008) **Autotagger: A Model for Predicting Social Tags From Acoustic Features on Large Music Databases**, Journal of New Music Research, 37(2):115-135, 2008.

(Breitschopf; 2013) **Personalized, Context-Aware Music Playlist Generation on Mobile Devices**. Master's thesis, Johannes Kepler University Linz, Austria, 2013.

(Cano, Koppenberger; 2004) **The Emergence of Complex Network Patterns in Music Artist Networks**, Proceedings of the 5th International Symposium on Music Information Retrieval (ISMIR 2004), Barcelona, Spain, October 2004.

(Casey, Veltkamp, Goto, Leman, Rhodes, Slaney; 2008) **Content-based music information retrieval: Current directions and future challenges**, Proceedings of the IEEE 96 (4), 668-696, 2008.

(Celino, Turati, Della Valle, Cerizza; 2006) **Squiggle: a Semantic Search Engine at Work**, Proceedings of the 1st International Conference on Semantics And digital Media Technology (SAMT 2006), Athens, Greece, December 2006

(Celma; 2010) **Music Recommendation and Discovery in the Long Tail – The Long Tail, Long Fail, and Long Play in the Digital Music Space**, Springer, 2010.

(Celma, Cano, Herrera; 2006) **SearchSounds: An Audio Crawler Focused on Weblogs**, Proceedings of the 7th International Conference on Music Information Retrieval (ISMIR 2006), Victoria, Canada, October 2006.

(Cooley, Tukey; 1965) **An Algorithm for the Machine Calculation of Complex Fourier Series**, Mathematics of Computation, 19(90): 297-301, 1965.

(Dror, Koenigstein, Koren, Weimer; 2012) **The Yahoo! Music Dataset and KDD-Cup'11**, Journal of Machine Learning Research, 18:3-18, 2012.

(Elliott, Tomlinson; 2006) **Personalsoundtrack: context-aware playlists that adapt to user pace**. CHI '06: Extended Abstracts on Human Factors in Computing Systems (CHI EA), New York, USA, 2006.

(Ellis, Whitman, Berenzweig, Lawrence; 2002) **The Quest For Ground Truth in Musical Artist Similarity**, Proceedings of the 3rd International Conference on Music Information Retrieval (ISMIR 2002), Paris, France, October 2002.

(Flexer, Schnitzer, Gasser, Pohle; 2010) **Combining Features Reduces Hubness in Audio Similarity**, Proceedings of the 11th International Society for Music Information Retrieval Conference (ISMIR), Utrecht, the Netherlands, 2010.

(Fujishima; 1999) **Realtime chord recognition of musical sound: a system using Common Lisp Music**, Proceedings of the International Computer Music Conference (ICMC), Beijing, China, 1999.

(Geleijnse, Korst; 2006) **Web-based Artist Categorization**, Proceedings of the 7th International Conference on Music Information Retrieval (ISMIR 2006), Victoria, Canada, October 2006.

(Geleijnse, Schedl, Knees; 2007) **The Quest for Ground Truth in Musical Artist Tagging in the Social Web Era**, Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007), Vienna, Austria, September 2007.

(Hauger, Schedl; 2012) **Exploring Geospatial Music Listening Patterns in Microblog Data**, Proceedings of the 10th International Workshop on Adaptive Multimedia Retrieval (AMR 2012), Copenhagen, Denmark, October 2012.

(Hirjee, Brown; 2009) **Automatic Detection of Internal and Imperfect Rhymes in Rap Lyrics**, Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR 2009), Kobe, Japan, October 2009.

(Hu, Downie, Ehmann; 2009) **Lyric Text Mining in Music Mood Classification**, Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR 2009), Kobe, Japan, October 2009

(Hu, Downie, West, Ehmann; 2005) **Mining Music Reviews: Promising Preliminary Results**, Proceedings of the 6th International Conference on Music Information Retrieval (ISMIR 2005), London, UK, September 2005

(Kaminskas, Ricci, Schedl; 2013) **Location-aware Music Recommendation Using Auto-Tagging and Hybrid Matching**, Proceedings of the 7th ACM Conference on Recommender Systems (RecSys), Hong Kong, China, 2013.

(Knees, Pampalk, Widmer; 2004) **Artist Classification with Web-based Data**, Proceedings of the 5th International Symposium on Music Information Retrieval (ISMIR 2004), Barcelona, Spain, October 2004.

(Knees, Pohle, Schedl, Schnitzer, Seyerlehner; 2008) **A Document-centered Approach to a Natural Language Music Search Engine**, Proceedings of the 30th European Conference on Information Retrieval (ECIR), Glasgow, Scotland, UK, 2008.

(Knees, Pohle, Schedl, Widmer; 2006) **Combining Audio-based Similarity with Web-based Data to Accelerate Automatic Music Playlist Generation,** Proceedings of the 8th ACM SIGMM International Workshop on Multimedia Information Retrieval (MIR 2006), Santa Barbara, USA, October 2006.

(Knees, Pohle, Schedl, Widmer; 2007) **A Music Search Engine Built upon Audio-based and Web-based Similarity Measures**, Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2007), Amsterdam, the Netherlands, July 2007.

(Knees, Schedl; 2011) **Towards Semantic Music Information Extraction from the Web Using Rule Patterns and Supervised Learning**, Proceedings of the 2nd Workshop on Music Recommendation and Discovery (WOMRAD), Chicago, USA, October 2011.

(Knees, Schedl; 2014) **Music Similarity and Retrieval**, expected to be published by Springer in 2014.

(Knees, Schedl, Pohle, Widmer; 2006) **An Innovative Three-Dimensional User Interface for Exploring Music Collections Enriched with Meta-Information from the Web**, Proceedings of the 14th ACM International Conference on Multimedia (MM'06), Santa Barbara, California, USA, 2006.

(Krenmair; 2010) **Musikspezifische Informationsextraktion aus Webdokumenten**, Master‘s thesis, Johannes Kepler University Linz, Austria, 2010.

(Law, von Ahn; 2009) **Input-Agreement: A New Mechanism for Collecting Data Using Human Computation Games**, Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI 2009), Boston, MA, USA, April 2009.

(Levy, Sandler; 2007) **A Semantic Space for Music Derived from Social Tags**, Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007), Vienna, Austria, September 2007

(Levy, Sandler; 2008) **Latent Semantic Models for Music**, Journal of New Music Research, 37(2):137-150, 2008.

(Lippens, Martens, Leman, Baets, Meyer, Tzanetakis; 2004) **A Comparison of Human and Automatic Musical Genre Classification**, Proceedings of the IEEE International Conference on Audio, Speech and Signal Processing (ICASSP), Montreal, Canada, 2004

(Liu, Rautenberg; 2009) **Music playlist recommendation based on user heartbeat and music preference**, **Improving Timbre Similarity: How High's the Sky?**, Proceedings of the 2009 International Conference on Computer Technology and Development (ICCTD), Washington DC, USA, 2009.

(Logan, Kositsky, Moreno; 2004) **Semantic Analysis of Song Lyrics**, Proceedings of the IEEE International Conference on Multimedia and Expo (ICME 2004), Taipei, Taiwan, June 2004.

(Mahedero, Martínez, Cano, Koppenberger, Gouyon; 2005) **Natural Language Processing of Lyrics**, Proceedings of the 13th ACM International Conference on Multimedia, Singapore, November 2005.

(Mayer, Neumayer, Rauber; 2008) **Combination of Audio and Lyrics Features for Genre Classification in Digital Audio Collections**, Proceedings of the ACM Multimedia, Vancouver, Canada; 2008.

(McFee, Bertin-Mahieux, Ellis, Lanckriet; 2012) **The Million Song Dataset Challenge**, Proceedings of 4th International Workshop on Advances in Music Information Research (AdMIRe 2012), Lyon, France, April 2012.

(Moens, van Noorden, Leman; 2010) **D-Jogger: Syncing Music with Walking**, In Proceedings of the 7th Sound and Music Computing Conference (SMC), Barcelona, Spain, July 2010.

(Pachet, Westerman, Laigre; 2001) **Musical Data Mining for Electronic Music Distribution**, Proceedings of the 1st International Conference on Web Delivering of Music (WEDELMUSIC 2001), Florence, Italy, November 2001.

(Pampalk; 2001) **Islands of Music - Analysis, Organization, and Visualization of Music Archives**, Master's Thesis, TU Vienna, 2001.

(Pampalk, Goto; 2007) **MusicSun: A New Approach to Artist Recommendation**, Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007), Vienna, Austria, September 2007

(Pohle, Knees, Schedl, Widmer; 2007) **Meaningfully Browsing Music Services**, Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007), Vienna, Austria, September 2007.

(Raphael; 2003**) Orchestral musical accompaniment from synthesized audio**. Proceedings of the International Computer Music Conference (ICMC). 2003.

(Resnick, Iacovou, Suchak, Bergstrom, Riedl; 1994) **GroupLens: An Open Architecture for Collaborative Filtering of Netnews**. Proceedings of Conference on Computer Supported Coorperative Work (CSCW), Chapel Hill, 1994.

(Schedl; 2011a) **Analyzing the Potential of Microblogs for Spatio-Temporal Popularity Estimation of Music Artists**, IJCAI 2011: International Workshop on Social Web Mining, Barcelona, Spain, July 2011.

(Schedl; 2011b) **Web- and Community-based Music Information Extraction**, In *Music Data Mining*, CRC Press/ Chapman Hall, July 2011.

(Schedl; 2012a) **#nowplaying Madonna: A Large-Scale Evaluation on Estimating Similarities Between Music Artists and Between Movies from Microblogs**, Information Retrieval, 15(3-4), 2012.

(Schedl; 2012b) **Exploiting Social Media for Music Information Retrieval**, In Naeem Ramzan, Roelof van Zwol, Jong-Seok Lee, Kai Clüver, Xian-Sheng Hua (eds.), *Social Media Retrieval*, Springer, December 2012

(Schedl; 2013) **Leveraging Microblogs for Spatiotemporal Music Information Retrieval**, Proceedings of the 35th European Conference on Information Retrieval (ECIR 2013), Moscow, Russia, March 2013.

(Schedl, Flexer; 2012) **Putting the User in the Center of Music Information Retrieval**, Proceedings of the 13th International Society for Music Information Retrieval Conference (ISMIR 2012), Porto, Portugal, Oct 2012

(Schedl, Flexer, Urbano; 2013) **The Neglected User in Music Information Retrieval Research**, Journal of Intelligent Information Systems, 2013.

(Schedl, Hauger; 2012) **Mining Microblogs to Infer Music Artist Similarity and Cultural Listening Patterns**, Proceedings of the 21st International World Wide Web Conference (WWW 2012): 4th International Workshop on Advances in Music Information Research (AdMIRe 2012), Lyon, France, April 2012.

(Schedl, Hauger, Schnitzer; 2012) **A Model for Serendipitous Music Retrieval**, IUI 2012: 2nd International Workshop on Context-awareness in Retrieval and Recommendation (CaRR 2012), Lisbon, Portugal, 2012.

(Schedl, Knees; 2011) **Personalization in Multimodal Music Retrieval**, Proceedings of the 9th Workshop on Adaptive Multimedia Retrieval (AMR 2011), Barcelona, Spain, July 2011.

(Schedl, Knees, Pohle, Widmer; 2006) **Towards Automatic Retrieval of Album Covers**, Proceedings of the 28th European Conference on Information Retrieval (ECIR 2006), London, UK, April 2006.

(Schedl, Knees, Widmer; 2005) **A Web-Based Approach to Assessing Artist Similarity using Co-Occurrences**, Proceedings of the 4th International Workshop on Content-Based Multimedia Indexing, Riga, Latvia, 2005.

(Schedl, Pohle, Knees, Widmer; 2011) **Exploring the Music Similarity Space on the Web**, ACM Transactions on Information Systems, 29(3):14, July 2011.

(Schedl, Pohle, Koenigstein, Knees; 2010) **What's Hot? Estimating Country-Specific Artist Popularity**, 11th International Society for Music Information Retrieval Conference (ISMIR), Utrecht, the Netherlands, 2010.

(Schedl, Schnitzer; 2013) **Hybrid Retrieval Approaches to Geospatial Music Recommendation**, Proceedings of the 35th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2013), Dublin, Ireland, July-August 2013.

(Schedl, Sordo, Koenigstein, Weinsberg; 2013) **Mining User-generated Data for Music Information Retrieval,** In Marie-Francine Moens, Juanzi Li, Tat-Seng Chua (eds.), *Mining of User Generated Content and Its Applications,* CRC Press, 2013.

(Schedl, Widmer; 2007) **Automatically Detecting Members and Instrumentation of Music Bands via Web Content Mining**, Proceedings of the 5th Workshop on Adaptive Multimedia Retrieval (AMR 2007), Paris, France, 2007.

(Schnitzer, Flexer, Schedl, Widmer; 2012) **Local and Global Scaling Reduce Hubs in Space**, Journal of Machine Learning Research, 13:2871-2902, October 2012.

(Schnitzer, Pohle, Knees, Widmer; 2007) **One-Touch Access to Music on Mobile Devices**, Proceedings of the 6th International Conference on Mobile and Ubiquitous Multimedia (MUM 2007), Oulu, Finland, December 2007.

(Shavitt, Weinsberg; 2009) **Songs Clustering Using Peer-to-Peer Co-occurrences**, Proceedings of the IEEE International Symposium on Multimedia (ISM2009): International Workshop on Advances in Music Information Research (AdMIRe 2009), San Diego, USA, December 2009.

(Seyerlehner; 2010) **Content-Based Music Recommender Systems: Beyond Simple Frame-Level Audio Similarity**, PhD Thesis, Johannes Kepler University Linz, December, 2010.

(Seyerlehner, Widmer, Knees; 2010) **A Comparison of Human, Automatic and Collaborative Music Genre Classification and User Centric Evaluation of Genre Classification Systems**, Proceedings of the 8th International Workshop on Adaptive Multimedia Retreival (AMR), Linz, Austria, 2010.

(Seyerlehner, Widmer, Schedl, Knees; 2010) **Automatic Music Tag Classification based on Block-Level Features**, Proceedings of the 7th Sound and Music Computing Conference (SMC), Barcelona, Spain, July 2010.

(Sordo; 2012) **Semantic Annotation of Music Collections: A Computational Approach**, PhD Thesis, Universitat Pompeu Fabra, Barcelona, Spain, February 2012.

(Turnbull, Barrington, Torres, Lanckriet; 2007) **Towards Musical Query-by-Semantic-Description using the CAL500 Data Set**, Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2007), Amsterdam, the Netherlands, July 2007.

(Turnbull, Barrington, Yazdani, Lanckriet; 2009) **Combining Audio Content and Social Context for Semantic Music Discovery**, Proceedings of the 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), Boston, USA, July 2009.

(Urbano, Schedl; 2012) **Minimal Test Collections for Low-Cost Evaluation of Audio Music Similarity and Retrieval Systems**, International Journal of Multimedia Information Retrieval, 2012.

(Whitman, Ellis; 2004) **Automatic Record Reviews**, Proceedings of the 5th International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain, 2004.

(Whitman, Lawrence; 2002) **Inferring Descriptions and Similarity for Music from Community Metadata**, Proceedings of the 2002 International Computer Music Conference (ICMC), Göteborg, Sweden, 2002.

(West, Cox, Lamere; 2006) **Incorporating Machine-Learning Into Music Similarity Estimation**, Proceedings of the 1st ACM Workshop on Audio and Music Computing Multimedia, Santa Barbara, USA, October 2006.