


timbral and semantic features for music playlists

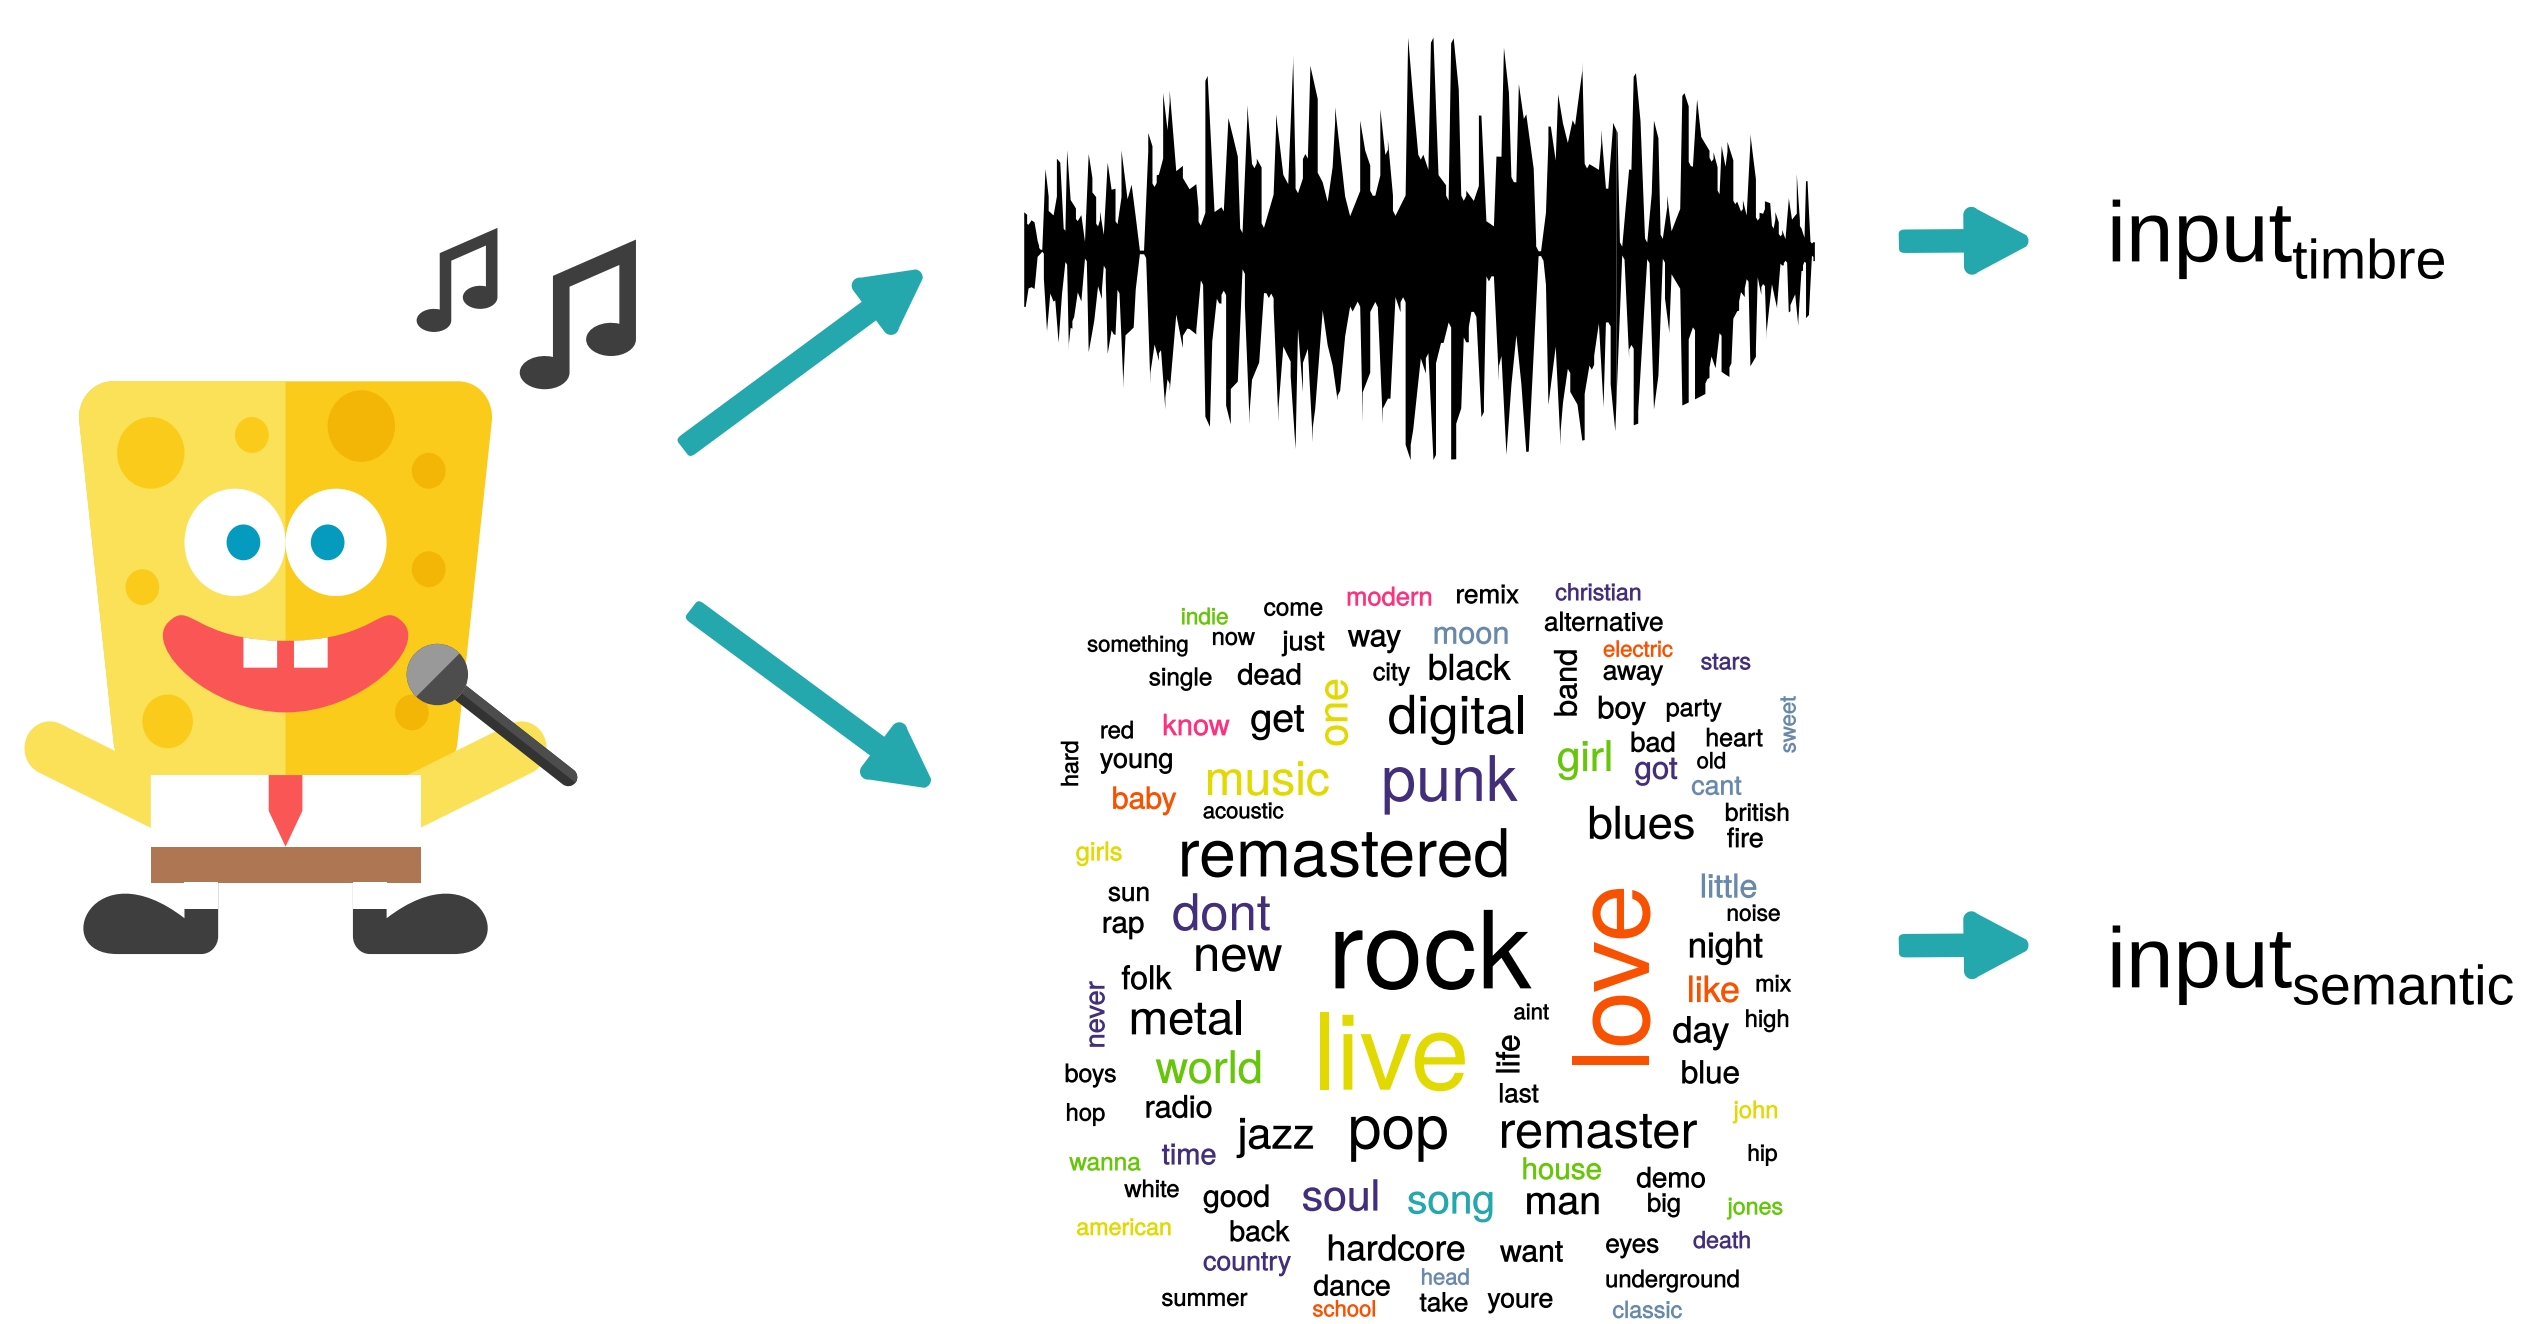
Andreu Vall, Hamid Eghbal-zadeh,
Matthias Dorfer, Markus Schedl
Department of Computational Perception
Johannes Kepler University, Linz, Austria
andreu.vall@jku.at



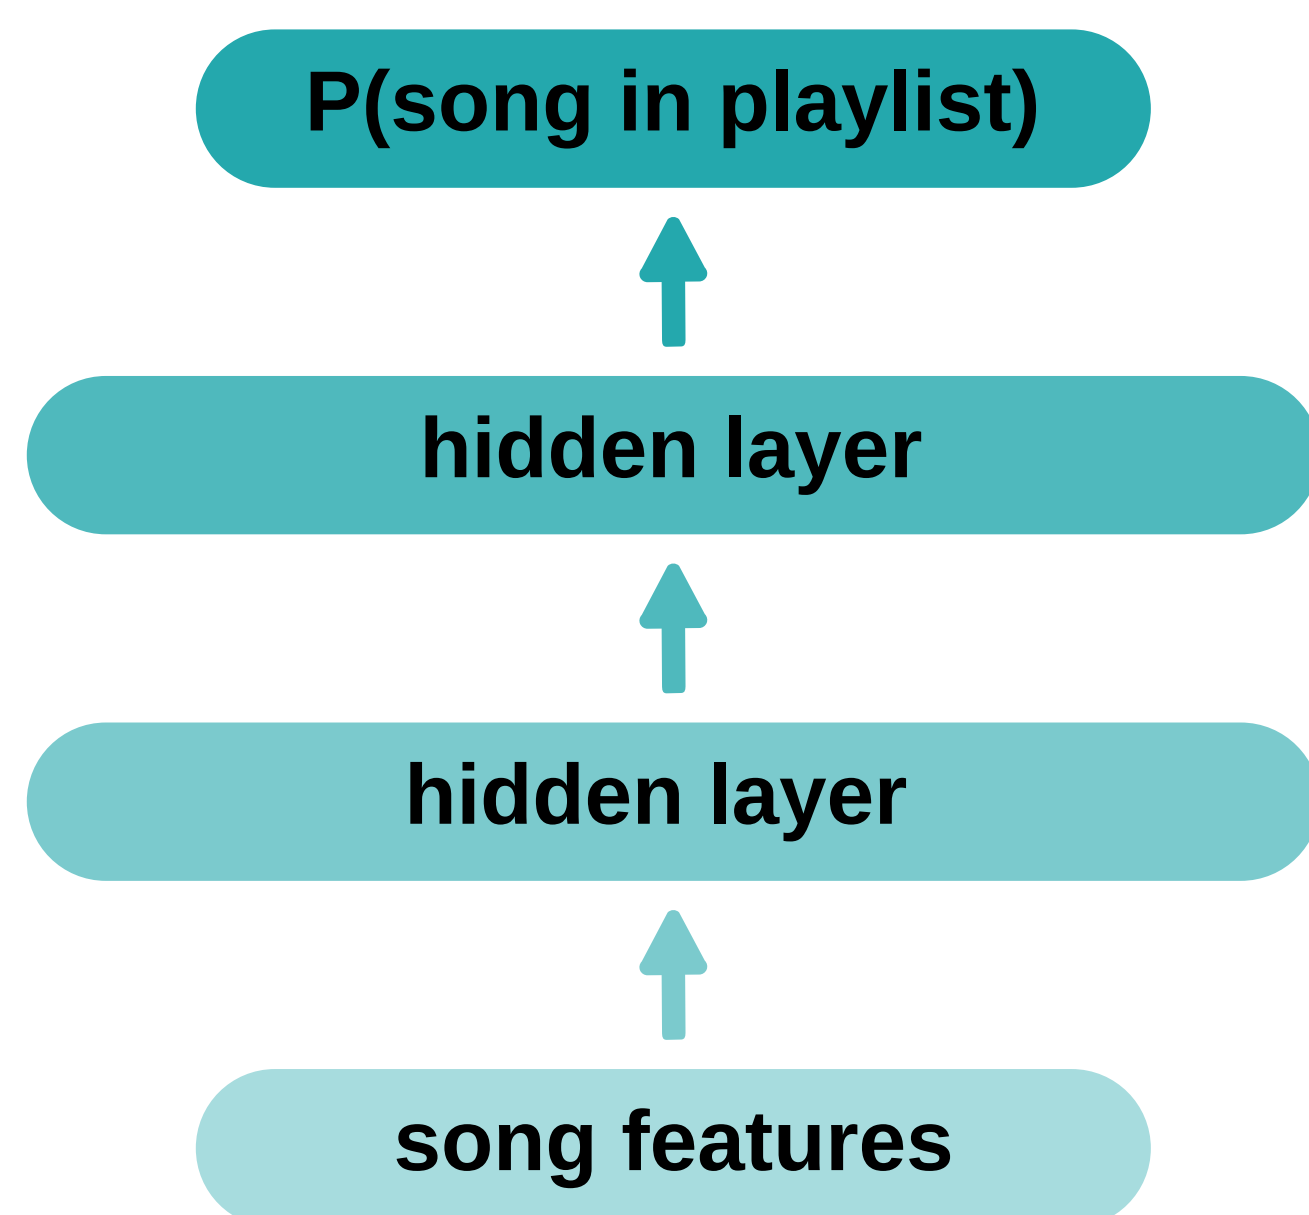
We analyze the predictive power of timbral and semantic features to decide whether a song is suitable or not for a playlist. Empirical results on a dataset of hand-curated playlists indicate that semantic descriptors are better suited for this task. Still, the multiplicity of factors and criteria relevant in the playlist curation process can make the problem ill-defined.

song to playlist classification

feature extraction



neural network classifier



Cost function

$$\mathcal{L}(\theta) = \sum_{i,j} -y_{ij} \log f_j(x_i) - (1 - y_{ij}) \log (1 - f_j(x_i))$$

playlist variability

Each point in the figure represents a song. The coordinates correspond to the first two principal components of the timbral and semantic features. The color indicates the playlist they belong to.

- Punk, dammit: *This was made when I first started getting into punk ...*
- Ultimate 80's: *These tracks ... could be called the quintessential 80's mix ...*
- Nothing's a waste: *This mix is really random ... he'll listen to anything ...*

numerical study

Based on subsets of the Art of the Mix 2011 dataset with 19+ known songs

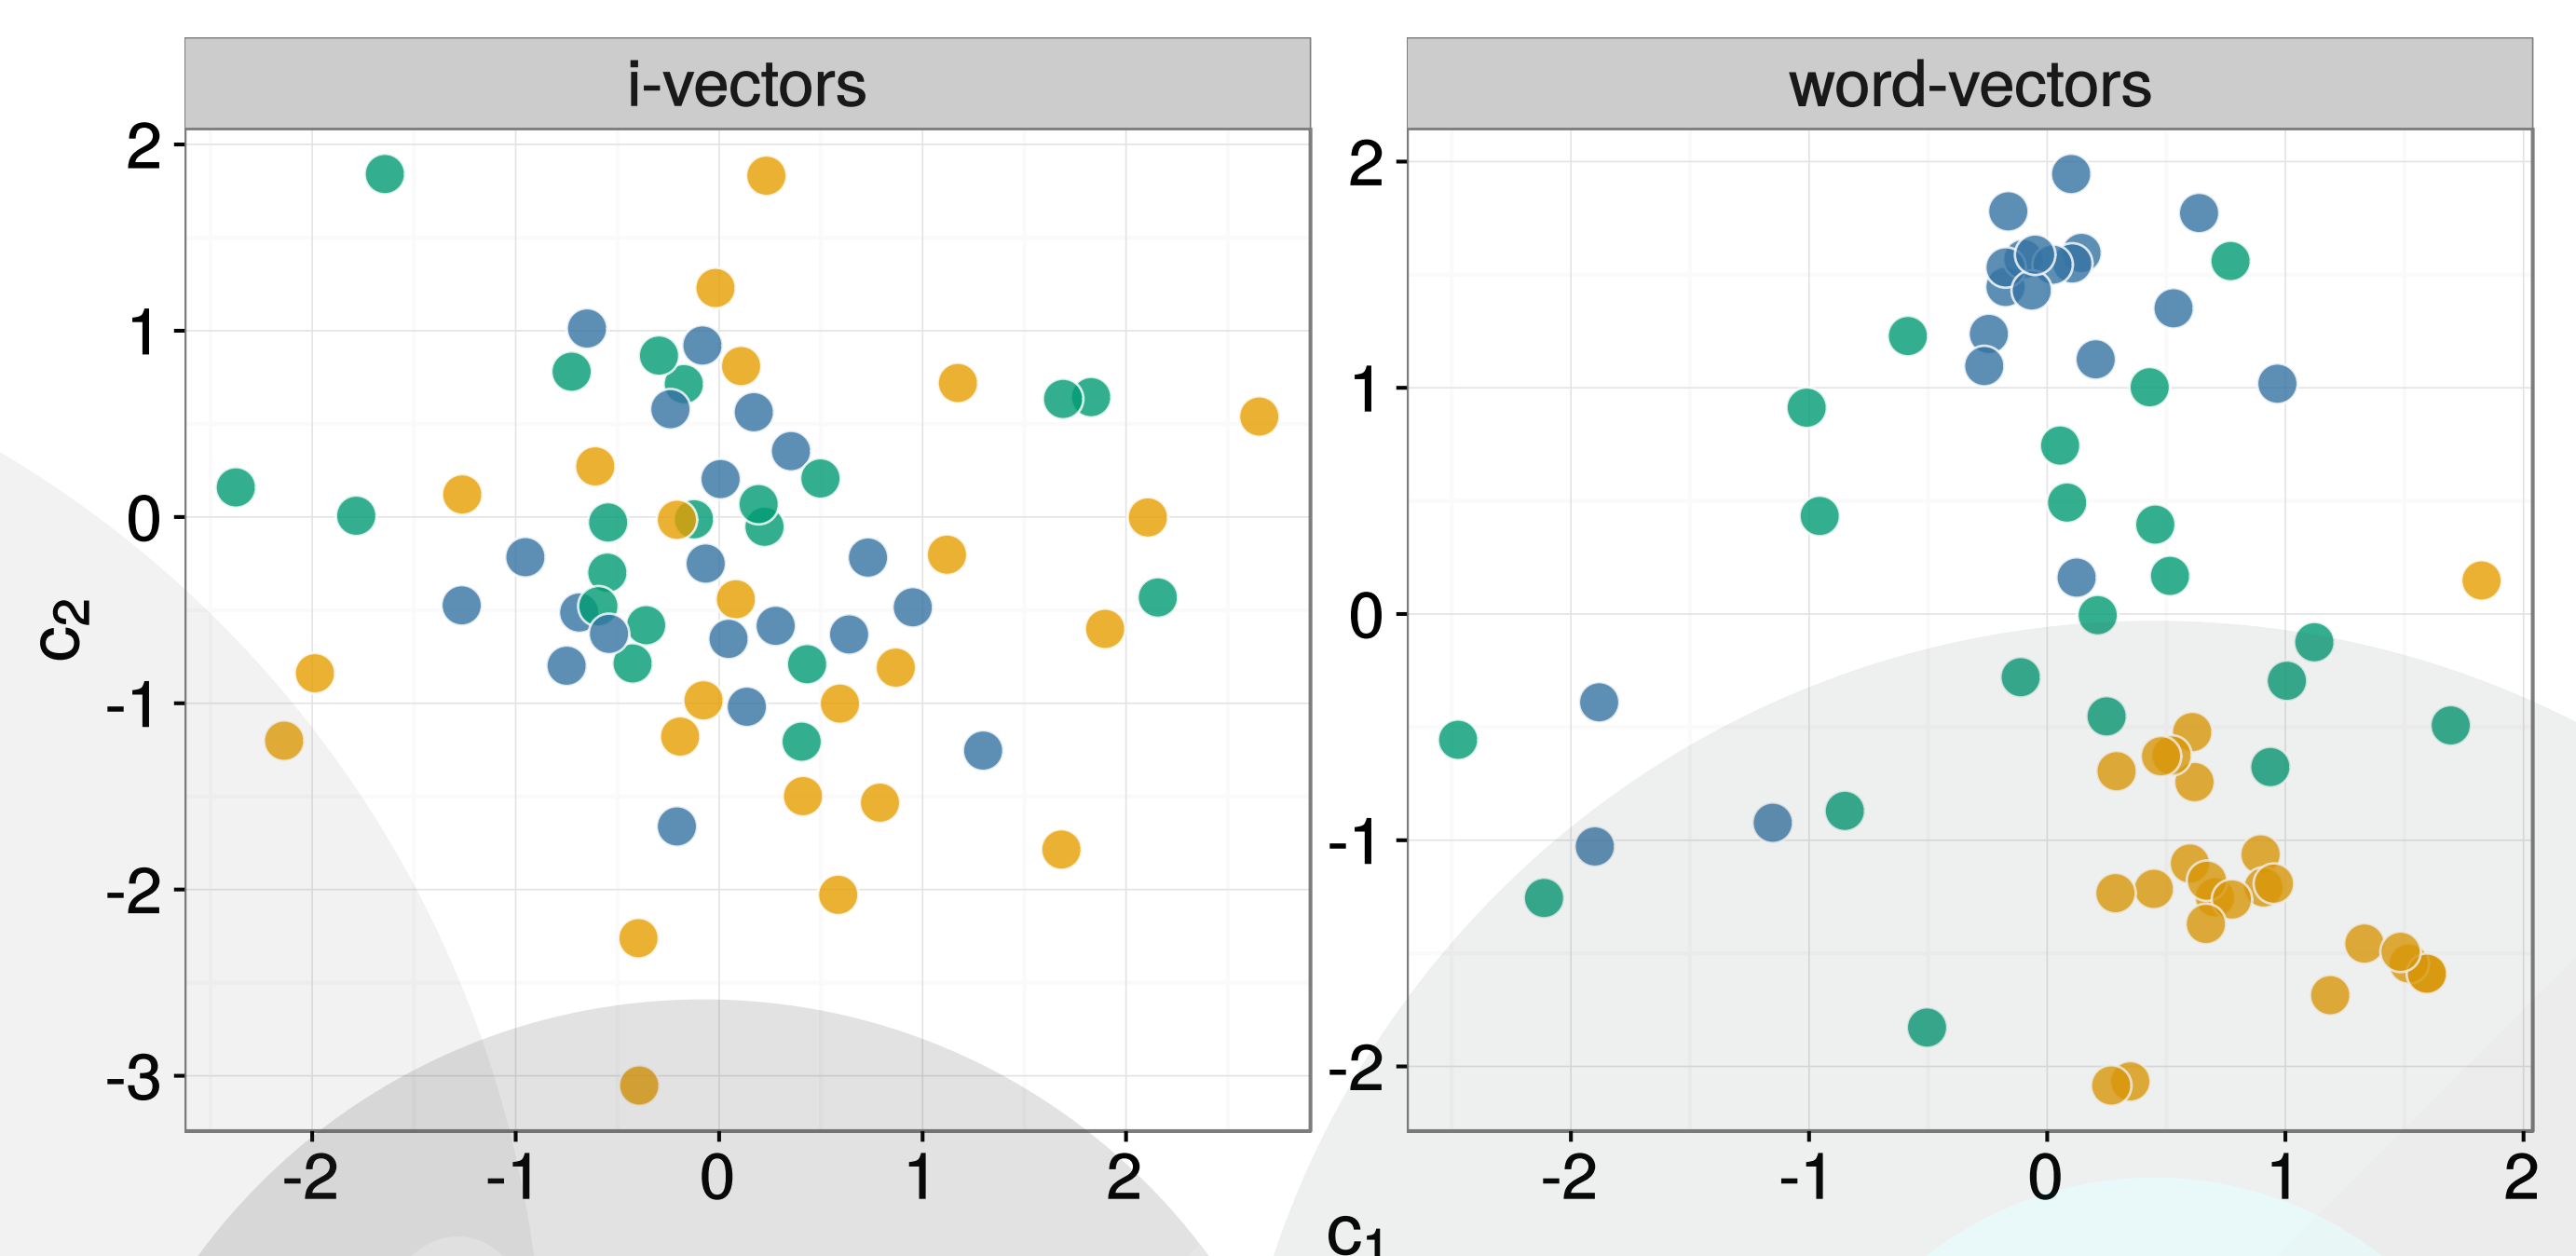
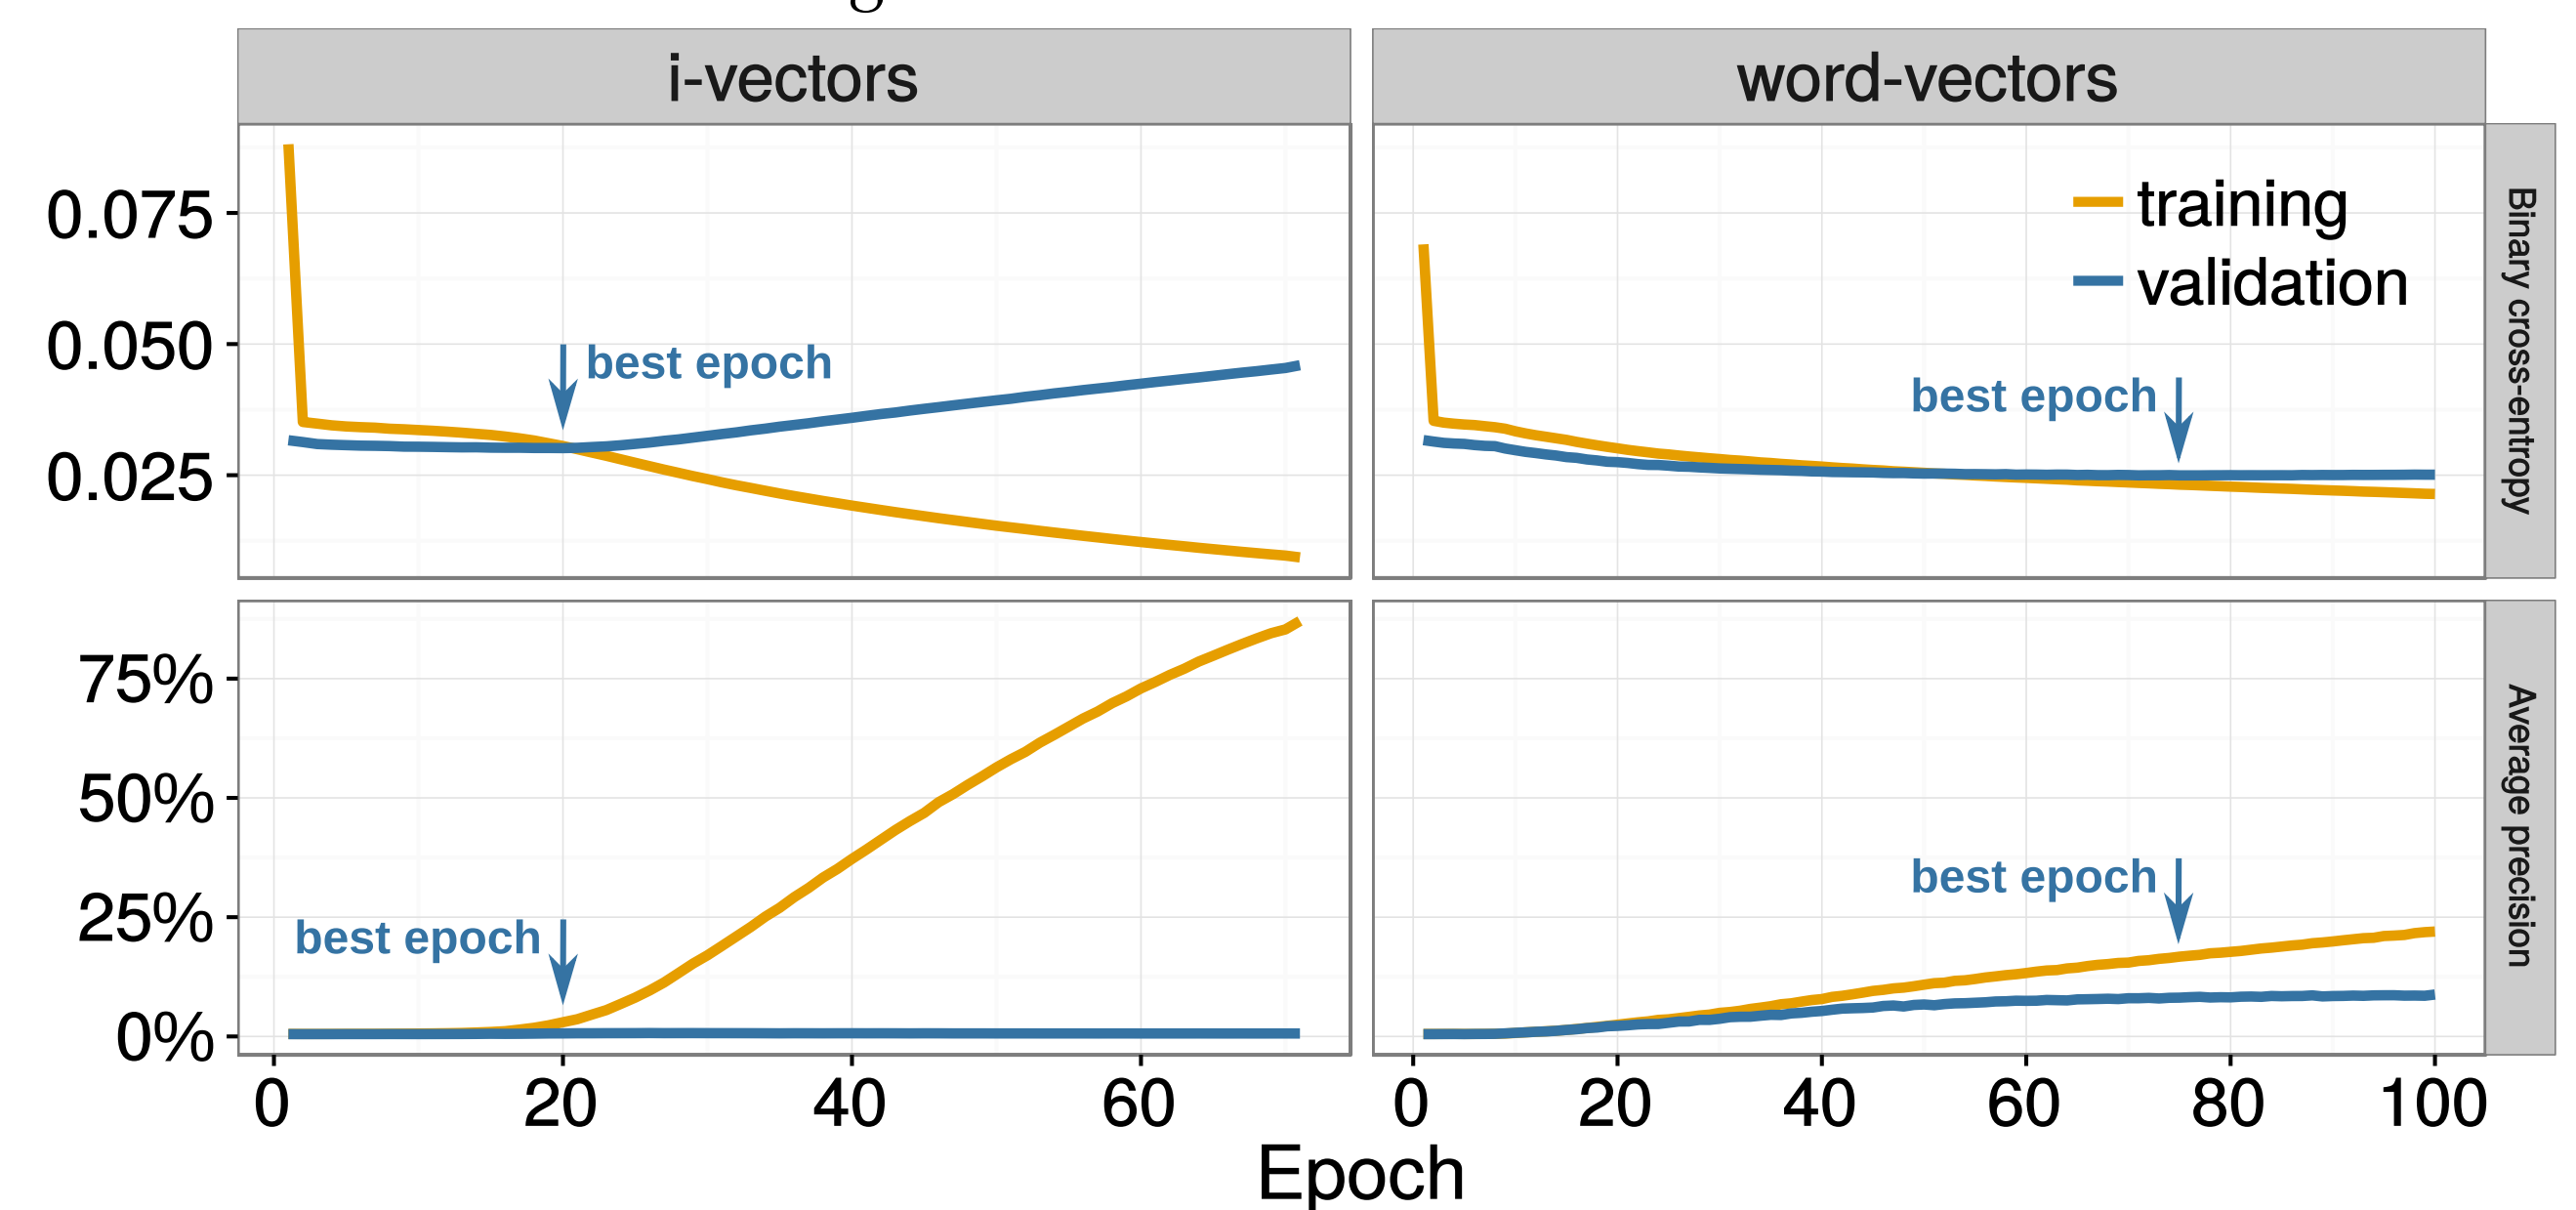
Classifier's performance

| score | #artists | i-vectors | | word-vectors | |
|----------------------|----------|-----------|--------|--------------|--------|
| | | train | valid | train | valid |
| Binary cross-entropy | 2 | 0.0279 | 0.0267 | 0.0180 | 0.0204 |
| | 5 | 0.0306 | 0.0302 | 0.0232 | 0.0250 |
| | 19 | 0.0388 | 0.0380 | 0.0306 | 0.0326 |
| Average precision | 2 | 2.36% | 0.66% | 29.97% | 17.32% |
| | 5 | 2.98% | 0.64% | 16.72% | 8.13% |
| | 19 | 2.72% | 0.71% | 11.44% | 5.06% |

Datasets overview

| #artists | #playlists | #songs |
|----------|------------|--------|
| 2+ | 258 | 4,674 |
| 5+ | 226 | 4,123 |
| 19+ | 170 | 3,244 |

Training evolution for the 5+ artists subset



[1] Vall, A., Eghbal-zadeh, H., Dorfer, M. and Schedl, M. Timbral and Semantic Features for Music Playlists. Accepted presentation at the Machine Learning for Music Discovery Workshop within the International Conference on Machine Learning (ICML 2016). New York City, NY, USA.

[2] McFee, B. and Lanckriet, G. Hypergraph Models of Playlist Dialects. In Proc. ISMIR 2012.