Weather Tunes

An attempt at predicting music listening behavior based on the weather.

Stijn van den Brink
University of Amsterdam
5922127
stijn.vandenbrink@student.uva.nl

Alex Olieman
University of Amsterdam
6077285
alex.olieman@student.uva.nl

Michael Wolbert
University of Amsterdam
10277331
michael.wolbert@student.uva.nl

1. INTRODUCTION

The automatic generation of playlists is a powerful way of exploiting large catalogues of music (Pratchet et al.,1999). It has therefore been widely studied and various different methods have been used to achieve this. One factor that, to the best of our knowledge, has not been researched yet for the automatic creation of playlists is the role that the weather might play in listening behavior. We hypothesize that weather reports can be used to predict which general acoustic features will be more prevalent in the music that people are listening to. In order to test this claim we propose to create a system that assesses listening behavior for a given weather condition of a location and predict listening behavior based on a weather report or forecast.

To do this the Last.fm API[1] is used to collect recorded listening behavior (number of song listeners, aggregated per metropolitan area) for several different locations over a period of two years (one 'chart' per week per location). We then use Last.fm and Echonest[2] to extract relevant music features of the tracks (e.g. artist, tempo and danceability). The next step is to retrieve historic weather data using the Weather Underground API[3], and represent this weather data as a collection of features (e.g. temperature, wind speed, occurrence of rain or snow) and analyze whether any of the music features are temporally correlated with the weather. Finally, we select significantly correlated features, and use these to predict what music features will be most listened to given certain weather conditions. We decided to choose for the prediction of features over the prediction of songs because by predicting the features the method would be separate from an actual music library and the output could thus more easily be used as another variable for existing recommendation and playlist generation systems.

A short literature review revealed that often music recommendations are based on long-term preferences (e.g. for certain artists, genres, etc.), and are not yet able to take into account how listening preferences may fluctuate during shorter time periods. It is possible that since the weather influences our moods, that mood influences our listening behavior. This is why we want to investigate if and how weather influences listening behavior. Music recommendation algorithms could take the current weather at a user's location into account with the creation of a playlist. This could be useful for recommending new music, but also for automatically generating playlists from a user's personal music library or 'favorite tracks'.

Predictions of which music features will be popular based on the (forecasted) weather can also be useful for people who select music for an audience. Radio DJs, for example, could make use of this system to get an idea of what kind of tracks to play during their show (now, later today, or tomorrow). The generated recommendations would for this use case consist directly of the features (e.g. a song with a "danceability" between *x* and *y*) that are most highly correlated with the weather (forecast) for the given day. Here, we view the DJ as a "human recommender", who takes these features into account in their own song selection process.

The remainder of the document will be outlined as follows. A short literature review will be presented in chapter 2. Chapter 3 will describe the data and methods used in this research and in chapter 4 the results will be presented. Chapter 5 and 6 will be used for the discussion and conclusion respectively.

2. RELATED WORK

There is already a fair amount of research available on the effects of weather on mood (Keller et al., 2005), and the same goes for research looking into how music affects mood (Bruner, 1990). Keller et al. found that higher temperature or barometric pressure led to a better mood, but only during spring time as the time spent outside increased. Lu et al. (2006) investigated detecting mood from music. But to the best of our knowledge research looking into the effect of weather on the listening behavior is lacking.

Similarly for the creation of playlists. Previous papers described different ways of generating playlists based on metadata. One way of doing this is by letting the user give one or more hints (seed songs) on which the algorithm can base the content of the playlist (Plat et al., 2002; Pauws & Eggen, 2002). AutoDJ by Plat et. al take multiple seedsongs from the user and uses the metadata from these songs to create a set of vectors of the seed songs. The program then tries to find matching songs from the library. Pauws and Eggen (2002) created PATS (Personalized Automatic Track Selection); this method only takes one seed song from the user to create a vector with metadata and match it to the songs in the available musical library. One other difference with AutoDJ is that PATS tries to both create a coherent and a varied playlist.

A third way of automatically creating playlists based on metadata is using constraints (Autocoutier and Pachet, 2002). Examples of such constraints are: "50% of the songs must be of the genre 'Rock' and the tempo must increase throughout the playlist".

[1] http://www.last.fm/api
[2] http://developer.echonest.com/acoustic-attributes.html
[3] http://www.wunderground.com/weather/api/

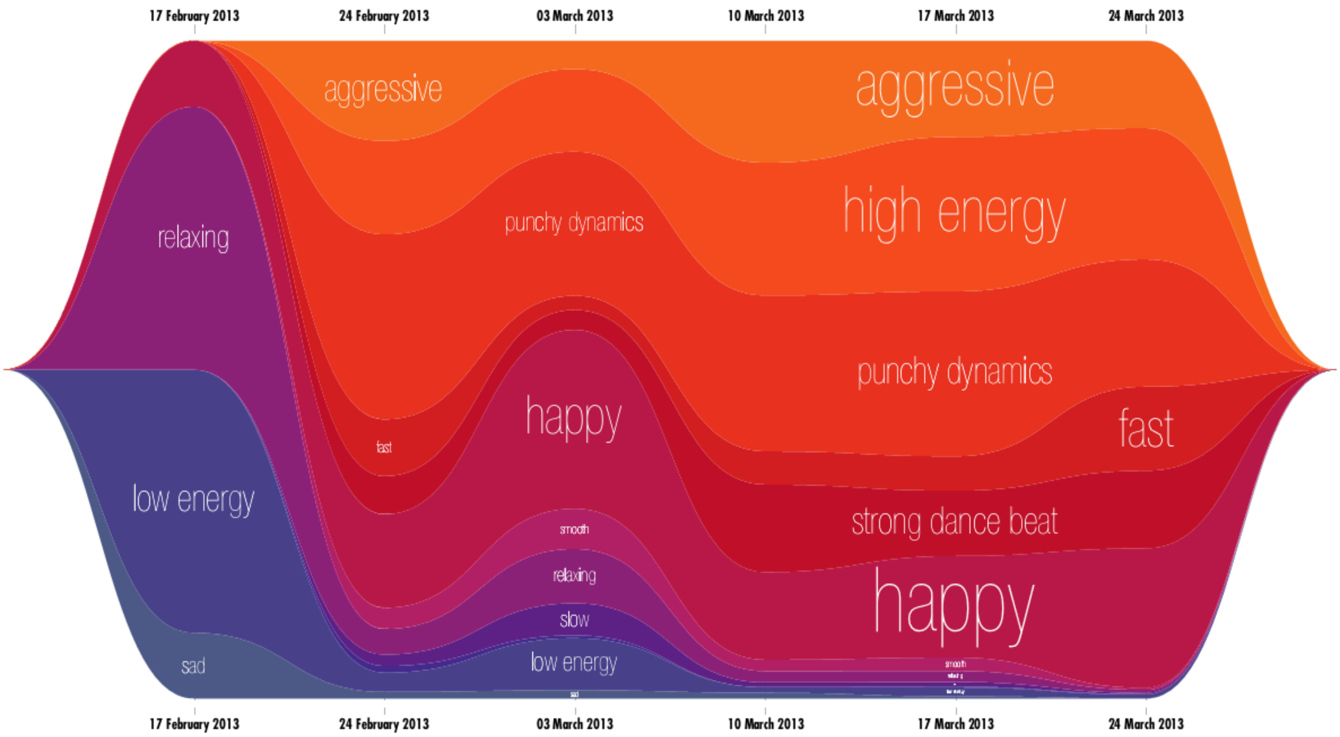

Figure 1. An example of a Last.fm mood report.

But generating a playlist based on metadata is prone to errors in this day and age with the internet as a potential music source (Andric & Haus, 2006). Because of this potential problem Andric and Haus (2006*)* looked for a method that was less sensitive to errors or missing values in the metadata, which is based on the listening behavior of a user. For this they looked at the frequency at which different songs were played in the past and tried to infer whether or not the user will want to listen to it again. Studies have also shown that collaborative filtering methods outperform content based methods (Barrington et al., 2009).

Flexer et al. (2008) were able to create a playlist by letting the user define a start and an end song. In the resulting playlist the songs in the beginning were similar to the start song and the ones at the end were similar to the end song. The similarity was calculated using audio analysis and other metadata was not used in this process. Maillet et al. calculated song similarity by looking at radio playlists and recorded all two and three song sequences (i.e. 'song *n*-grams'). Chen et al. (2012) also used existing playlists, but modeled them as a Markov chain and used that to create a new playlist. Fields et al. (2010) showed that it is possible to create a playlist by combining existing playlists and Last.fm tags.

Shavitt and Weinsberg (2009) calculated song similarity by analyzing peer-to-peer networks. They found that users with a similar taste often share the same files and that this data can be used to create relatively accurate recommendations. Levy and Sandler (2008) argue that current methods of extracting information about a song using audio features perform poorly because these features are not rich enough to describe something as complex as music. They opted for the usage of social tags. These are tags that anyone can assign to a song with the idea that users will gravitate to similar descriptive tags. They build models using these tags to try and improve their ability to search through the music collection and improve techniques for recommending songs to users. Last.fm used automatically computed machine tags to create a 'Mood Report' for individual users (Figure 1). These tags are based on audio analysis as they found that the social tags for songs rarely described the actual sound or mood of a song, and instead were more often used for describing genre, era or nationality of the artist (Levy, 2012).

Another difficulty of automatically creating playlists are social gatherings because you cannot make the assumption that everyone will like the same music. PATS already tried to do this by also making a varied playlist. Bauer et al. (2011) also tried to select music based on mood. They, however, devised a method that tracks ongoing conversations and then tries to extract the current mood. It then picks matching songs based on their tags and puts those songs in the playlist.

3. METHODS

The approach that is taken in this study is structured as follows. First, the necessary data is gathered and pre-processed. Second, the data is analyzed for possible relationships between aspects of the weather and attributes of the music that people listened to. Finally, we use the resulting statistical models in an experiment to predict a popularity ranking of songs based on weather features. The evaluation of the predictions occurs against a ground truth of observed (historical) listening behavior.

This paper focuses on the *how* and *what* of our approach. For an overview of the code that was instrumental in executing this approach, please consult Appendix A.

3.1 Gathering Data

Three main categories of data are needed to uncover possible relationships between aspects of the weather and attributes of the music that people listened to. Records of listening behavior should indicate *when* someone listened to *which song*, and at *what location*. For these locations and times weather records should yield values for (perceivable) parameters of the weather. Because our aim is to predict general attributes of music, the final category of data should consist of values of these attributes for each listened song.

Traditional music charts (e.g. Top 40, Billboard) are not suitable as records of listening behavior. They are based on record sales or radio plays, and therefore do not directly reflect when people choose to listen to which music. A more direct source of observed listening behavior is produced by automatic track logging (i.e. scrobbling), most notably through the Audioscrobbler[4] software. Audioscrobbler facilitates scrobbling on more than 50 software music players, on several Hi-Fi systems with an internet connection, and on several mobile music players, including iPods, iPhones, Android and Windows phones. The listened tracks that are logged by Audioscrobbler are available through the Last.fm API, which constitutes the largest collection of publicly available listening behavior records (Geleijnse, Schedl, & Knees, 2007).

The Last.fm API offers track logs for individual users in their most granular form (i.e. including a timestamp for each listened song) through the *user.getRecentTracks*[5] API method. However, Last.fm does not publish the location of individual users, which is needed to find the relevant weather records. Instead, we have to rely on *charts* of the 200 most popular songs in a given week (Monday to Sunday), that have been aggregated per metropolitan area[6]. These charts measure the popularity of a song in terms of the number of people who have scrobbled it at least once during this week, from this metropolitan area. The total number of plays for a track is not given, because this metric is easily abused (e.g. by repeatedly playing a single album or track).

Historical weather records are obtained from the Weather Underground API[7]. This service offers detailed weather observations (every 10 minutes) for the majority of populated areas. The use of weekly charts, however, does not allow us to benefit from such detailed weather measurements. Instead, we rely on daily summaries of the weather which give minimum, maximum, and mean values for the measured continuous variables. From all observations that Weather Underground offers, we use the following as our *weather features*:

- Hail (yes or no)
- Snow (yes or no)
- Fog (yes or no)
- Tornado (yes or no)
- Rain (yes or no)
- Thunder (yes or no)
- Mean/min/max pressure
- Mean/min/max dew point
- Mean/min/max visibility
- Mean/min/max wind speed
- Mean/min/max temperature
- Mean/min/max humidity
- Mean/min/max snow depth
- Amount of snowfall
- Amount of precipitation
- Location, date and time(zone).

Because there still is a temporal discrepancy between the (weekly) charts and the (daily) weather summaries, some form of alignment needs to be performed. We consider two possibilities, which both rely on a simplifying assumption. The first method consists of aligning the weekly charts to the daily weather summaries, by assuming that the weekly listener counts for songs can be averaged over each day of the week. This alignment thus produces identical music features for each day of the week. The second method consists of aligning the weather summaries to the weekly charts, by producing a weekly weather summary from the daily summaries. The values of the weather features are re-summarized by computing new mean, minimum, and maximum values where applicable, taking the sum of snowfall and precipitation, and by assigning a “yes” to the binary variables if the value is “yes” for at least one day of the week.

The final category of needed data, music features, is retrieved from two sources. Last.fm offers basic song metadata, such as artist name and song duration, and the top five tags that were assigned to the song by the greatest number of people. With the aim of predicting general attributes of music in mind, we consider the track title and album title to be too specific. As an additional source of music features we use The Echo Nest to retrieve acoustic metadata through their Song API[8]. Several high-level attributes, such as “energy” and “danceability”, have been produced by models which have been trained to predict their value based on lower-level acoustic analysis, and against a ground truth of musicians’ annotations (Whitman, 2013). Two additional song attributes that we gather from The Echo Nest are derived from global listening behavior: artist “familiarity”, and song “hotttnesss” (Jehan, Lamere, & Whitman, 2010). The full initial selection of music features is listed in Table 1.

Table 1. The listing of initially selected music features.

Feature	Description	Type
artist_name	The name of the recording artist, as provided by Last.fm.	Nom
artist_familiarity	“The probability that a music fan will have heard of an artist” (Jehan, Lamere, & Whitman, 2010, pp. 245)	Num
artist_latitude	The latitude of the recording artist’s main location, as provided by The Echo Nest.	Num
artist_longtitude	The latitude of the recording artist’s main location, as provided by The Echo Nest.	Num
duration	The duration of the song in	Int

[4] Audioscrobbler: http://www.audioscrobbler.net/

[5] Individual track logs: http://www.last.fm/api/show/user.getRecentTracks

[6] Metropolitan charts: http://www.last.fm/api/show/geo.getMetroTrackChart

[7] Weather Underground API: http://api.wunderground.com/history/

[8] The Echo Nest Song API: http://developer.echonest.com/docs/v4/song.html

	seconds, as provided by Last.fm.	
top_tags	The five tags that have been assigned to the song by the greatest number of Last.fm users.	Nom
danceability	"The ease with which a person could dance to a song, over the course of the whole song."[9]	Num
liveness	A measure of how much the song sounds like a live recording (e.g. by detection of crowd noise).	Num
hotttnesss	"The daily measure of how [...] listened to [a song] is" (Jehan, Lamere, & Whitman, 2010, pp. 245).	Num
energy	A measure of how energetic the performance sounds (e.g. by using loudness and segment durations).	Num
speechiness	The probability that the track consists of exclusively speech (i.e. a high value denotes mostly speech, while a value near zero indicates no speech at all).	Num
tempo	The number of beats per minute, as provided by The Echo Nest.	Int
time_signature	The time signature of the song, limited to 3/4, 4/4, 5/4, and 7/4.	Nom
mode	Whether a song uses a minor or major scale.	Bool
key	The key of the scale that is used throughout the song.	Nom

The final decision that needs to be made regarding the collection of data, is: for what time range will we collect the Last.fm charts, and for which metropolitan areas? Arguably, we need to retrieve the charts for at least a full year, because the weather varies significantly throughout a yearly seasonal cycle. One year of charts might, however, still give a distorted picture of relations between weather and listening behavior through popularity of a few tracks (e.g. due to an album release) that coincides with a season. To mediate this expected distortion, we retrieve charts and weather records for the period January 2011 - December 2012.

To select metropolitan areas (hereafter 'metros') for which to retrieve charts and weather records, we employ four criteria. Most importantly, we wish to select metros for which charts of the entire 2011-2012 period are available. Secondly, we prefer metros with a large number of Last.fm users. By applying these criteria to the available metros on Last.fm, we narrow the choice down into an intermediate set of 10 metros: Berlin, Chicago, Istanbul, Los Angeles, Madrid, Melbourne, Moscow, New York City, São Paulo, and Sydney. From this intermediate set, we prefer metros with significantly different climates, and with different "listening cultures". Our final selection, which satisfies these preferences, consists of Moscow, New York City, and São Paulo.

A total number of 106 charts and 750 daily weather summaries were retrieved per metro. Figure 2 shows how the total number of listener counts (frequency) varies across the chosen time period. The frequency is computed as the sum of the number of unique listeners for each of the 200 songs in a chart:

$$(1)\ Frequency = \sum_{i=1}^{200} L_i,$$

where L_i indicates the unique listener count of a song. This is not a measure of overall unique listeners, because it is likely that users will have listened to multiple songs in the chart. Figure 2 shows a higher average frequency in 2012 than in 2011, with a remarkable week-to-week fluctuation throughout the period. Most notably, there are several charts for which the frequency in São Paulo is significantly lower than for the surrounding weeks (e.g. the two weeks following 1 July 2011). We suspect that this is caused by a malfunction in Last.fm's logging systems, and that these particular charts do not give an accurate reflection of listener counts.

3.2 Data Analysis

The aim of our data analysis is to create a predictive statistical model for each music feature (see Table 1). The goal is not to blindly end up with one statistical model per music feature, but rather to select the music features that show most potential to be predicted on the basis of the weather features. The first step in this analysis is pre-processing of the data. Several variables need to be removed from the dataset. The collected data is organized as rows of music feature - weather feature tuples (*mf, wf*) per song in a chart. To incorporate the number of people that listened to a given song during a week into our number of observations, the (*mf, wf*) tuples need to be multiplied (i.e. duplicated) by the listener count.

After the data is pre-processed, we inspect the data visually by creating plots of each possible pair of a music (response variable) and weather feature (explanatory variable). We use scatter plots for pairs of numeric variables (interval and ratio), and use box plots for pairs of a numeric and a non-numeric (nominal or Boolean) variable. This approach is not feasible for *artist_name* and *top_tags* due to their large number of categories. Specifically for pairs of numeric variables, we check if there seems to be a linear relationship between the variables. If not, it might be possible to apply logarithmic or power transforms to arrive at a linear relationship. Because of the large number of observations in the dataset, we use random samples of the data to produce readable scatter plots.

The statistical models are not created from the entire dataset, since we want to keep "unseen" data for use in the evaluation. The dataset is split into test- and train data by randomly selecting 20% of the days from the entire period 2011-2012. For each day in this selection, the entire (mf, wf) tuple is placed in the test dataset; the remaining tuples are used as training data. We prefer this over using a time period (e.g. Sep-Dec 2012) as test data, because of the yearly seasonal cycle.

[9] As explained by by J. Sundram: http://runningwithdata.com/post/1321504427/danceability-and-energy

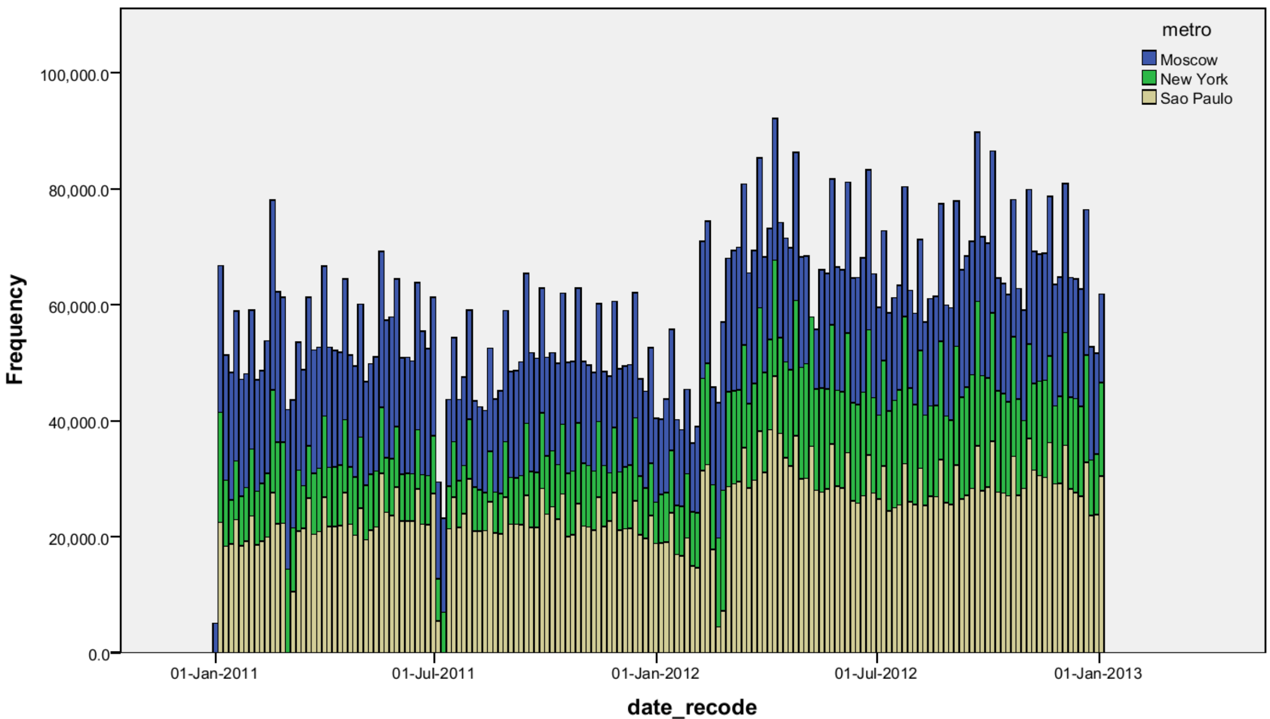

Figure 2. Timeline of the scrobble frequency for 2011-2012 (colored per metro).

In order to achieve our goal of predicting music features on the basis of multiple weather features, we use multiple regression. To choose specific types of models, a distinction needs to be made between nominal and numerical response variables. Multinomial logistic regression is needed to predict nominal variables. This modeling approach most commonly used in maximum entropy classifiers, and often performs better than Naive Bayes because it does not assume independence of the explanatory variables (Jurka, 2012). It is not possible to assume independence for the selected weather features, because many of them are significantly correlated.

For the numerical variables it is possible to use multiple linear regression. First we create a (linear) maximal model which includes all possible explanatory variables. Then, we use stepwise regression to remove less informative explanatory variables, until we arrive at the minimal adequate model. Because we are uncertain about finding any relations whatsoever, we try to create an ordering of variables from strongest to weakest predictors. This procedure needs to be repeated for each music feature. On the basis of the Adjusted R-squared of these models, we will select the music features which can be predicted relatively well to continue with.

The final product of this stage is a selection of music features that show most promise to be predicted. For each selected music feature, two statistical models will be produced: one using a daily alignment method, and one using a weekly alignment method between charts and weather records (see section 3.1).

3.3 Predictions and evaluation of song rankings

With the resulting trained statistical models we intend to predict the selected music features for observed weather features. For the Radio DJ scenario, it seems useful to produce confidence intervals in order to not only predict the best fit, but also consider a range in which each predicted music feature will lie with high confidence. By having these lower and upper prediction values, the DJ has a broader selection criteria for inclusion of songs into his playlist. For evaluation purposes we do not include confidence intervals, but rely on the best fit functions instead.

To evaluate the predictive ability of the trained models, we argue that it would be worthwhile to let the models predict song rankings (charts) from the test dataset based on weather features which have been aligned to these charts. This approximates the task of ranking a user's music library to select the top songs for inclusion in a playlist. Several components are needed to do so, and they are described in evaluation steps below;

1. For each song in a chart (unseen *mf*) in the test dataset, predict a music feature vector; $\mathrm{pmf} = (\mathrm{mf}_0, \ldots, \mathrm{mf}_n)$ given a weather feature (seen *wf*).
2. For each song in the chart let *smf* be the seen music feature vector and let *d* be the distance metric between vectors *smf* and *pmf*.
3. Let $chart_o$ be the original chart ranked by number of listeners, and let $chart_d$ be the chart ranked by distance.
4. For each chart in the test dataset, obtain Spearman's rank correlation coefficient (ρ) between $chart_o$ and $chart_d$ where the selected variable for the test is *d*.
5. To account for insignificant ρ values do; ρ = if p-value >.05: 0, otherwise ρ. Where the p-value is retrieved from the Spearman's rank correlation function.

For the distance calculation of step 2 we will use the cosine similarity. The function is implemented as follows by the scipy[10]

[10] http://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.cosine.html

module; 1 - *cosine similarity*. Spearman's ρ is also available through scipy for step 4.

Since each chart features 200 songs at a maximum, we use 10, 50, 100, 150 and 200 number of songs per chart during the evaluation. This decision is based on evaluating the effect of re-ranking various sizes of charts from step 3.

4. RESULTS

The following section is organized as follows: First, we elaborate on descriptives regarding the dataset and argue how we processed missing values. Second, we present the definitive selection of music features. Third, we elaborate on the results of the 30 (see[11]) trained regression models for the music features *danceability*, *energy*, *tempo*, *loudness* and *speechiness*. Finally, we test the hypothesis against the results of the evaluation.

4.1 Data descriptives & missing data

The complete dataset contained two years of charts with a total of 4662 unique tracks (Moscow, N = 1914; New York, N = 2632; São Paulo, N = 1309)[12]. Since not all acoustic metadata of a given track were available through The Echo Nest, these tracks were removed from the dataset. In summary, missing music features accounted for 9.50% (N = 443) of the total unique tracks with music features (Moscow, N = 1689; New York, N = 2484; São Paulo, N = 1218). An interesting observation from the data is that most Russian and Brazilian artists did not have any retrievable music features from The Echo Nest. Furthermore, the artist name can be composed of a collaboration between artists (i.e. 'feat.') which similarly led to missing music features because these were not available.

For each metro, the dataset contained weather features for two years and each day, accumulating to 2190 daily weather observations. None of the metros reported a *tornado*, and therefore this feature was omitted from the dataset. A similar approach was employed for Moscow's *humidity*, *snowfall* and *snow-depth*, New York's *humidity* and *hail*, and São Paulo's *humidity*, *snow*, *snowfall*, *snow-depth*. Incidentally, *precipitation* and *snowfall* contained the value "T", which was respectively recoded into 5 and none for Moscow, 15 and 2.5 for New York, and 20 and none for São Paulo.

4.2 Music feature selection

Because of the large amount of observations multinomial logistic regression had trouble running, likely due to its iterative training approach. The regressions took several hours to run, and always failed when the model grew larger than the size of the workstation's available memory. Therefore no nominal variables could be used in the models as such.

Correlations between the selected weather- and music features did not yield strong significant relationships. However, as expected, mutually correlated music - and weather features exhibited stronger significant relationship (i.e. temperature and dew, or energy and speechiness; see Appendix B). Since these relationships exhibited elliptical clusters, no non-linear relationship modeling was assumed.

The nominal variables *mode*, *time_signature* & *key* were interpreted as numerical values and using these values multiple linear regressions were run on the numerical values. Even though the *hotttnesss* variable fitted relatively well (mean Adjusted R^2 of 0.007285), it was excluded from the final selection because this variable's value changes daily, and thus contained arbitrary values from the day we retrieved the data, instead of from the date of the respective chart.

The variables *key*, *mode* and *time_signature* were dropped first and foremost because the fit was not very good (mean Adjusted R^2 of 0.000592, 0.000349, 0.000326 respectively), but also because they were interpreted as a numerical values, and this is a doubtful practice.

The model fit of *duration*, *liveness* and *tempo* was also relatively low (mean Adjusted R^2 of 0.000767, 0.000760, 0.000708 respectively). However, because previous literature mentioned *tempo* as an important factor for radio DJs and playlist creation, it was decided to use that variable instead of *duration* and *liveness*.

After this selection process the following features were used for further research: *energy*, *tempo*, *speechiness*, *loudness* and *danceability*.

4.3 Daily and weekly multiple linear regression models

A multiple linear regression model was employed to describe the linear relationship between a music feature (dependent variable) and weather features (independent variables) for daily and weekly datasets using a backward stepwise approach. Equation 2. depicts this linear function where mf is the music feature and c_i is the coefficient for weather feature wf_i.

$$(2)\ \mathrm{mf} = \beta + c_1{*}\mathrm{wf}_1 + c_2{*}\mathrm{wf}_2 + \ldots + c_n{*}\mathrm{wf}_n$$

For each metro and music feature, a model was build on train data (80% of the data). The Adjusted R^2 scores are presented in table 2 for each of these models. These scores reflect the model's ability to account for the variance in *mf*. Moreover, it expresses the fit of the model (i.e. Moscow's weekly loudness model has its selected *wf* terms accounting for 1.41% of the variance in loudness).

The Adjusted R^2 scores for both model types is low, indicating that the selected weather features per model do not describe the relation to *mf* very well. An interesting observation from table 2, is that Moscow's weekly danceability scores are relatively low compared to loudness, while degrees of freedom reports the first to have 7 (83794 - 83787) more *wf* terms in the model. Similar observations can be made for New York's daily energy and speechiness model scores. Overall, the weekly Adjusted R^2 scores suggest a better fit after the aggregation of weather features into a weekly summary, except for New York's loudness - and a tie on São Paulo's danceability model.

For an overview of the included and excluded weather features (independent variables) from the models see appendix C. While all included variables are significant, there are only slight differences in inclusion and exclusion of *wf* terms between the models.

4.3.1 Music feature predictions

From the resulting models the fit, lower - and upper confidence bounds (level = 95%) functions were obtained. For evaluation purposes the fit function was used, while the upper- and lower confidence interval functions were assumed to be useful in playlist generation. Table 3 illustrates the result of all predicted music features for New York's models with the following weather features input;

[11] Model types (2) * metros (3) * music features (5) = 30

[12] Note that there is overlap in unique tracks per metro.

Table 2. Multiple linear regression model fit metrics

		Moscow		*New York*		*Sao Paulo*	
response	**model**	**adjusted R^2***	**df**	**adjusted R^2***	**df**	**adjusted R^2***	**df**
loudness	*weekly*	0.0141	83787	0.0067	61006	0.0172	122270
	daily	0.0073	641428	0.0068	406946	0.0020	807634
energy	*weekly*	0.0152	83788	0.0041	61009	0.0033	122275
	daily	0.0085	641426	0.0008	406946	0.0011	807641
Speechiness	*weekly*	0.0101	83788	0.0054	61006	0.0015	122277
	daily	0.0040	641427	0.0004	406950	0.0009	807640
Danceability	*weekly*	0.0027	83794	0.0086	61008	0.0032	122274
	daily	0.0007	641426	0.0005	406947	0.0032	807642
Tempo	*weekly*	0.0027	83790	0.0020	61009	0.0023	122279
	daily	0.0005	641430	0.0009	406954	0.0007	807640

* all reported adjusted R^2 scores are significant at $p < .05$

Prediction input:

hail=no, snow=yes, thunder=no, tornado=no,rain=yes, fog=yes, presh=1032, presl=1017, press=1025.8, temph=-6, templ=-10, temp=-8, dewh=-8, dewl=-12, dew=-10, snowd=3, vish=10, visl=3, vis=8.6, windh=25, windl=0, wind=15, humh=93, huml=74, hum=86, precip=5, snowf=0

Table 3. Predictions for model fit, lower- and upper confidence bounds (level = 95%). Note that not all weather features are included in each model (see appendix C).

predicted feature	*model*	*fit*	*lower 2.5%*	*upper 97.5%*
loudness	weekly	-9.087	-18.848	0.673
	daily	-10.151	-20.028	-0.273
energy	weekly	0.474	0.117	0.831
	daily	0.465	0.105	0.824
speechiness	weekly	0.078	-0.084	0.239
	daily	0.072	-0.086	0.230
danceability	weekly	0.496	0.200	0.791
	daily	0.542	0.247	0.837
tempo	weekly	106.92	62.730	151.108
	daily	107.70	63.032	152.365

4.4 Evaluation

While the models' Adjusted R^2 scores indicate a low model fit, an evaluation, based on a random sample of 20% of the data (test data), was still carried out. From the descriptives of the dataset, some 2190 data points for each daily weather observation were gathered. Therefore, 435 daily weather features were randomly selected along with the corresponding ranked tracks (charts) for that day. The evaluation, as elaborated in the method section, utilizes Spearman's rank correlation coefficient (ρ) to assess the models' capability of predicting the ranking of the original chart based on a vector of the predicted music features. In table 4 the results of the evaluation are presented for five chart sizes (see '# tracks' column).

Moscow's models mean ρ is mostly negative (except for the daily model containing 200 tracks), meaning that after re-ranking the original chart based on distance metrics of the prediction, its relationship between the original chart is explained oppositely. In figure 3 and 4 the box plots for the evaluation are given for the weekly - and daily model respectively. From these box plots one can observe that for Moscow charts exist which were significantly positive related, but that more charts were significantly negative related, resulting in mean ρ going towards -1. For New York's models, similar observations can be made. For São Paulo, however, the relationship between the original - and re-ranked chart can be explained by 6.91% (daily model containing 10 tracks) of the variance of music features in the original chart. For São Paulo and New York mean ρ decreases when more tracks are added to the charts for evaluation, while for Moscow it increases.

Appendix D contains an example of the evaluation of a chart with 10 tracks for the daily - and weekly models of all metros.

4.5 Hypothesis testing

At the outset of this paper, the hypothesis which was established was; Weather reports can be used to predict which general acoustic features will be more prevalent in the music that people are listening to. To test the hypothesis; the scores presented in the evaluation (see table 4) are tested for mean $\rho > 0$. Since each metro was isolated during modeling and evaluation, the hypothesis was rejected for Moscow and New York, but provisionally accepted for São Paulo.

5. DISCUSSION

The first, and probably the biggest, deciding factor in the listening behavior is the musical preference of users. Even if people do decide on what song to listen to based on the weather, songs from different genres could match the same mood but have totally different musical features. Or even have musical features that match the features of a song from a different genre that matches a different mood.

Table 4. Evaluation metrics for all models.

		Moscow		*New York*		*Sao Paulo*		*All*	
# tracks	**model**	**mean ρ***	**% sig.****	**mean ρ***	**% sig.****	**mean ρ***	**% sig.****	**mean ρ***	**% sig.****
10	*weekly*	-0.0295	4.35	0.0045	6.52	0.0551	13.39	0.0100	13.39
	daily	-0.0354	5.11	0.0197	5.76	0.0691	15.20	0.0178	15.20
50	*weekly*	-0.0084	7.46	-0.0166	6.52	0.0285	9.09	0.0010	9.09
	daily	-0.0062	8.27	-0.0132	5.00	0.0231	7.46	0.0011	7.46
100	*weekly*	-0.0342	17.07	-0.0218	17.60	0.0258	11.63	-0.0102	11.63
	daily	-0.0248	15.20	-0.0268	18.55	0.0228	9.92	-0.0097	9.92
150	*weekly*	-0.0086	8.27	-0.0236	14.84	0.0277	16.13	-0.0017	16.13
	daily	-0.0076	8.27	-0.0250	15.75	0.0279	16.13	-0.0017	16.13
200	*weekly*	-0.0002	9.92	-0.0258	19.51	0.0172	10.77	-0.0031	10.77
	daily	0.0049	17.07	-0.0275	21.49	0.0127	9.09	-0.0035	9.09

* mean significant **ρ** is calculated based on p < .05 (see step 5 in section 3.3)
** based on total of Moscow = 144, New York = 147, Sao Paulo = 144

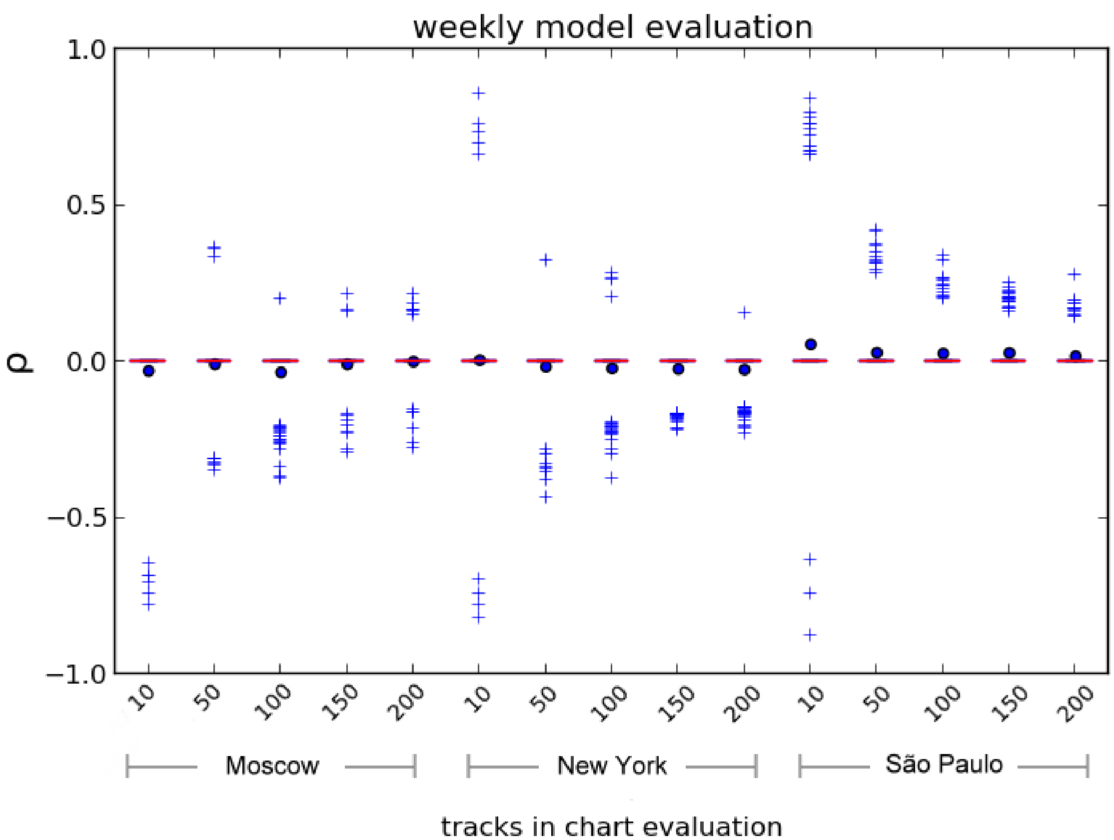

Figure 3. Plot of ρ of weekly models per metro and number of tracks. Note that the 'dot' corresponds to the mean as seen in Table 3.

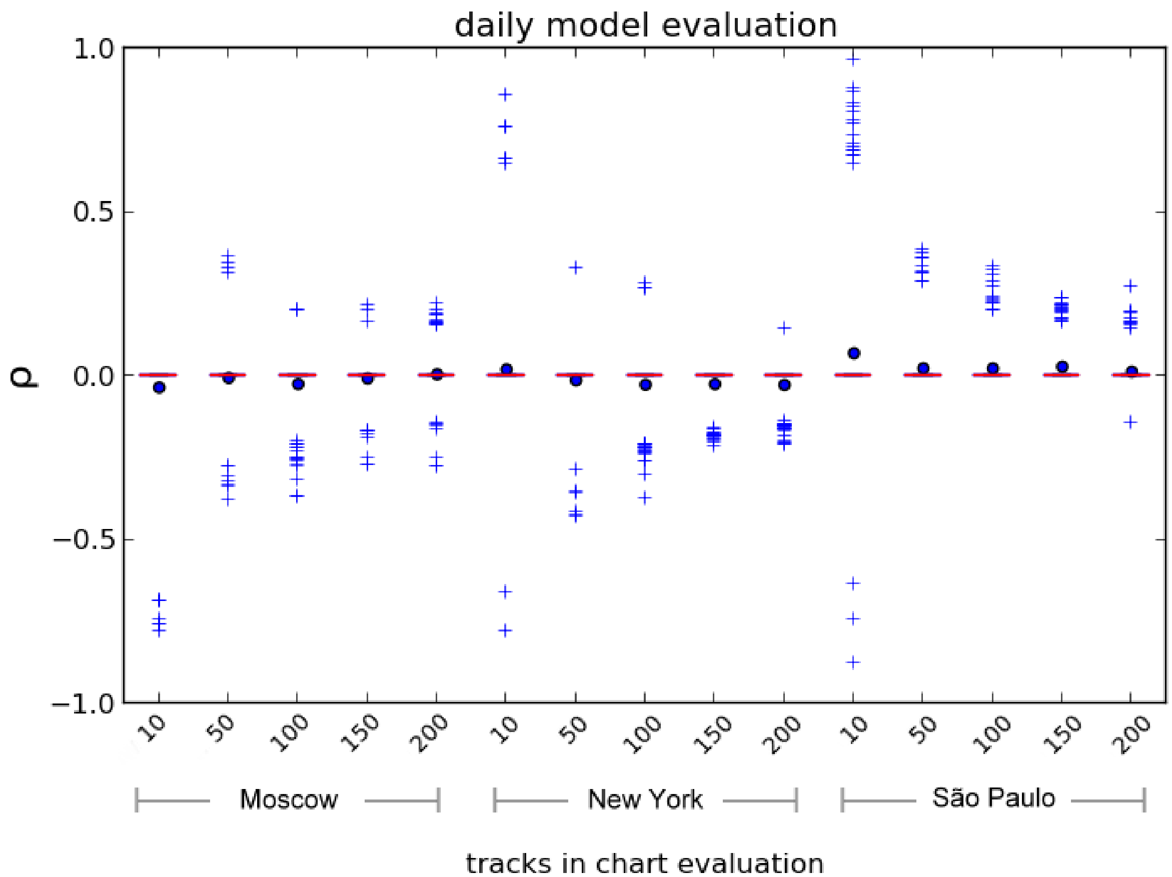

Figure 4. Plot of ρ of daily models per metro and number of tracks.

Note that the 'dot' corresponds to the mean as seen in Table 3.

There are a couple of possible explanations that could explain the results. First of all it is unknown how much of an influence the weather has on the listening behavior of people who spend most of their time inside. Next there are occasions where whole albums were present in the charts. This could be explained by people listening to whole albums, either because these albums were new releases or because they just happen to really like an older album. Another possible explanation of the possible disconnect between the weather and the charts is that users can create a playlist of for example their whole music library and play this list on shuffle. By listening to music this way the users could skip numbers they are not in the mood for, but it is also possible that they do not really mind what song is playing as that they probably already like these songs.

From the literature is seems that there might me a slight correlation between weather and mood (only in certain conditions). It is however possible that there is both a positive and a negative relation between listening behavior and mood. For example someone who is sad might want to elevate his spirits by listening to a more up-tempo song, while someone else could decide to listening to a song that actually matches his mood. Besides possible differences in listening behavior of different users, the same user could also at one point listen to music matching their mood, while at another point in time decide to listen to music that contrasts their mood. This brings us to the next point: how important are the actual lyrics when deciding to pick what song to choose. In this research the music features that were used are based on audio analysis and the actual topic of the song, which could be extracted by analyzing the lyrics, was not taken into account. The actual topic of a song could play an important role in the decision of whether a song matches or is contrasting to the mood of a user. To counter these highly personal preferences in deciding what song to play would be to use user profiles instead of charts for certain regions. This way you will take personal (musical) preference into account and by doing so you get a list of recommendations that can be evaluated by said users.

The Last.fm API does support the extraction of user information, but because of the difficulty of having to select active users with different musical preferences and in different locations, the decision was made to use the much broader charts.

Although understandable it is a pity it was impossible to retrieve track information for all the tracks present in the charts. As was mentioned earlier, finding music features for certain tracks by (local) artists in the charts from Moscow and São Paulo was impossible. Because these are mostly national artists, the songs they produce could have musical features that when used to train models, have these models be better trained for each metro and its culture.

Regarding the actual results of this research, some confidence intervals are not really that useful because they are simply too large. For example the information that songs with a BPM between 80 and 120 will be more popular today is not very useful for a radio DJ as this not specific enough. Although when combined multiple intervals these intervals could narrow the selection down further. Calculating the confidence intervals also required considerable more amounts of memory and computing power compared to predicting music features without these intervals. In case of a radio DJ wanting to predict features for tomorrow's playlist this might not be that big of an issue, but you cannot expect every individual user to have these amounts of memory available in their desktop pc or online recommendation services to have enough memory to serve a decent user base. Therefore the described methods are not efficient enough for users who want a generated playlist right now either through an online service or by using their own pc.

Furthermore, because cross-validation of the data would give us more information to work with, it would exclude pure chance as a possible explanation for the results for São Paulo. Unfortunately due to time limits this step was not performed in this research.

And finally, the question remains whether or not our selected method actually fits our data and goal. Further research should be conducted in order to see if different kinds of models would produce better and more reliable results. Unfortunately due to time constraints only linear models were used in this research.

6. CONCLUSION

This study researched how our listening behavior might be influenced by the daily weather in context of three metropolitan areas. These areas were fixed to Moscow, New York and São Paulo and the data from its citizens listening behavior and weather observations were retrieved from several web APIs for the year 2011 and 2012. A total of 30 multiple regression models were built and evaluated.

Our hypothesis was; weather reports can be used to predict which general acoustic features will be more prevalent in the music that people are listening to. While the results indicate that the hypothesis was only accepted for São Paulo, the predictive power of São Paulo's models were rather weak. Therefore, the chance of an accurate prediction is suggested to be low, and hence is not useful for the recommendation of songs for playlists given several weather features. From the evaluation for São Paulo, we obtained a relatively small set (7% - 21%) of significant ranked charts. This relatively small set indicated that a great number of these evaluated charts were scrambled, and had no predictive power with respect to the original charts.

In conclusion this research showed that an accurate prediction of listening behavior based solely on the weather is near impossible. We cannot, however, conclude that the weather does not influence listening behavior at all. We argue that existing music recommendation systems can, however, still benefit by also taking weather based predictions into account. In order for these existing systems to benefit from our predictions, future research is needed in the direction of further development of current and new methods for incorporating weather features into a music recommendation system.

7. REFERENCES

Andric, A., & Haus, G. (2006). Automatic playlist generation based on tracking user's listening habits. *Multimedia Tools and Applications*, *29*(2), 127-151.

Aucouturier JJ, Pachet F (2002, July 3–5) Scaling up music playlist generation. Poster Session on the 3rd CIRP International Seminar on Intelligent Computation in Manufacturing Engineering, Ischia (Naples), Italy

Bruner, G. C. (1990). Music, mood, and marketing. *The Journal of Marketing*, 94-104.

Barrington, L., Oda, R., & Lanckriet, G. (2009, October). Smarter than genius? human evaluation of music recommender systems. In International Symposium on Music Information Retrieval.

Bauer, J. S., Jansen, A., & Cirimele, J. (2011, October). MoodMusic: a method for cooperative, generative music playlist creation. In Proceedings of the 24th annual ACM symposium adjunct on User interface software and technology (pp. 85-86). ACM.

Chen, S., Moore, J. L., Turnbull, D., & Joachims, T. (2012, August). Playlist prediction via metric embedding. In Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 714-722). ACM.

Fields, B., Rhodes, C., & d'Inverno, M. (2010). Using song social tags and topic models to describe and compare playlists. In 1st Workshop On Music Recommendation And Discovery (WOMRAD), ACM RecSys, 2010, Barcelona, Spain.

Flexer, A., Schnitzer, D., Gasser, M., & Widmer, G. (2008). Playlist generation using start and end songs. In International Symposium on Music Information Retrieval (ISMIR'08) (pp. 173-178).

Geleijnse, G., Schedl, M., & Knees, P. (2007, September). The quest for ground truth in musical artist tagging in the social web era. In *Proceedings of the Eighth International Conference on Music Information Retrieval* (pp. 525-530).

Jehan, T., Lamere, P., & Whitman, B. (2010, March). Music retrieval from everything. In *Proceedings of the international conference on Multimedia information retrieval* (pp. 245-246). ACM.

Jurka, T. P. (2012). maxent: An R Package for Low-memory Multinomial Logistic Regression with Support for Semi-automated Text Classification. *The R Journal, 4(1)*, 56-59.

Keller, M. C., Fredrickson, B. L., Ybarra, O., Côté, S., Johnson, K., Mikels, J., Conway, A. & Wager, T. (2005). A warm heart and a clear head the contingent effects of weather on mood and cognition. *Psychological Science*, *16*(9), 724-731.

Levy, M., & Sandler, M. (2008). Learning latent semantic models for music from social tags. *Journal of New Music Research*, *37*(2), 137-150.

Levy, M. (2012, August 30). How are you feeling today? *Last.fm - the Blog*. From http://blog.last.fm/2012/08/30/how-are-you-feeling-today

Lu, L., Liu, D., & Zhang, H. J. (2006). Automatic mood detection and tracking of music audio signals. *Audio, Speech, and Language Processing, IEEE Transactions on*, *14*(1), 5-18.

Maillet, F., Eck, D., Desjardins, G., & Lamere, P. (2009, October). Steerable playlist generation by learning song similarity from radio station playlists. In Proceedings of the 10th International Conference on Music Information Retrieval.

Pauws S, Eggen B (2002, October 13–17) PATS: Realization and user evaluation of an automatic playlist generator. In: Proc. of the 3rd International Conference on Music Information. Retrieval (ISMIR 2002), IRCAM-Centre Pompidou, Paris, France

Platt, J. C., Burges, C. J., Swenson, S., Weare, C., & Zheng, A. (2002). Learning a Gaussian process prior for automatically generating music playlists.*Dietterich et al.[50]*, 1425-1432.

Pachet, F., Roy, P., & Cazaly, D. (2000). A combinatorial approach to content-based music selection. Multimedia, IEEE, 7(1), 44-51.

Shavitt, Y., & Weinsberg, U. (2009, December). Song clustering using peer-to-peer co-occurrences. In Multimedia, 2009. ISM'09. 11th IEEE International Symposium on (pp. 471-476). IEEE.

Whitman, B. (2013). How music recommendation works -- and doesn’t work. Brian Whitman @ variogr.am. Retrieved on 1 January 2013 From: http://notes.variogr.am/post/37675885491/how-music-recommendation-works-and-doesnt-work

Appendix A: Code

For this research multiple Python and R scripts were used to obtain the results that are described:

- main_scraper.py was used to get the charts and weather information using lastfm_scraper.py and wunderground_scraper.py respectively.
- Next music-features.py was used to collect the features from the songs in the charts from both Echo Nest and Last.fm.
- query_examples.py was used to create the training and test datasets.
- lm.py was used to create the feature and ranking predictions using the models created by the R scripts in statistics.
- General helper scripts can be found in helpers.

Appendix B. Box- and Scatterplots to illustrate visual inspection of the dataset

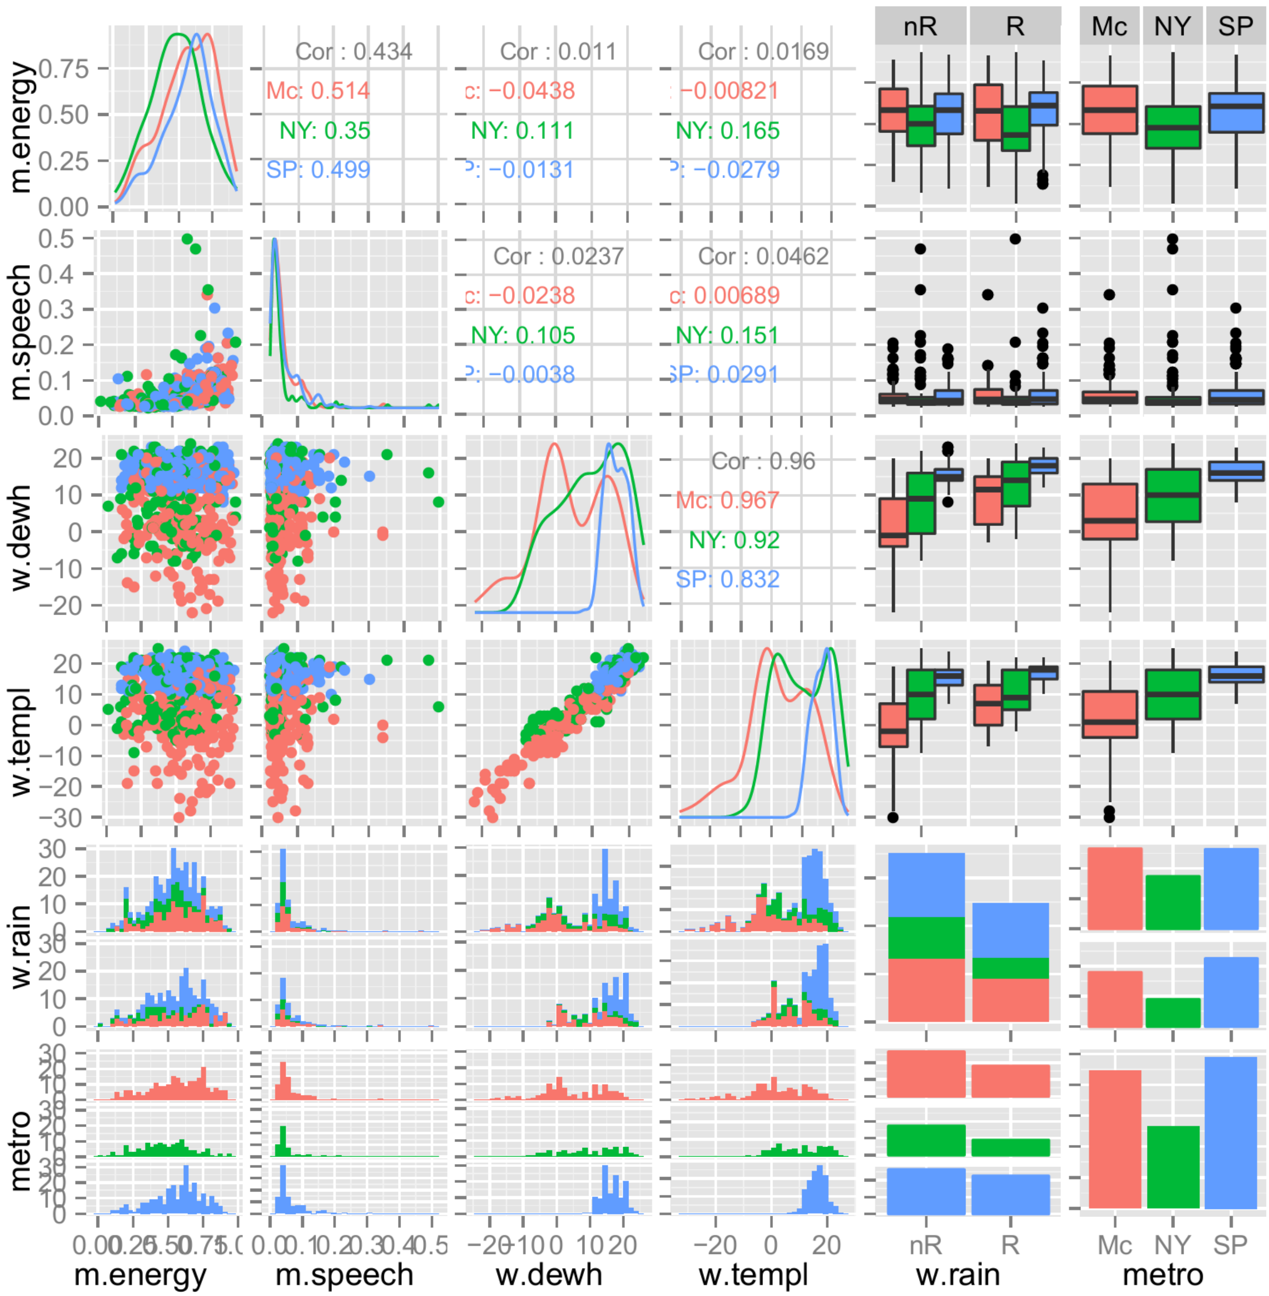

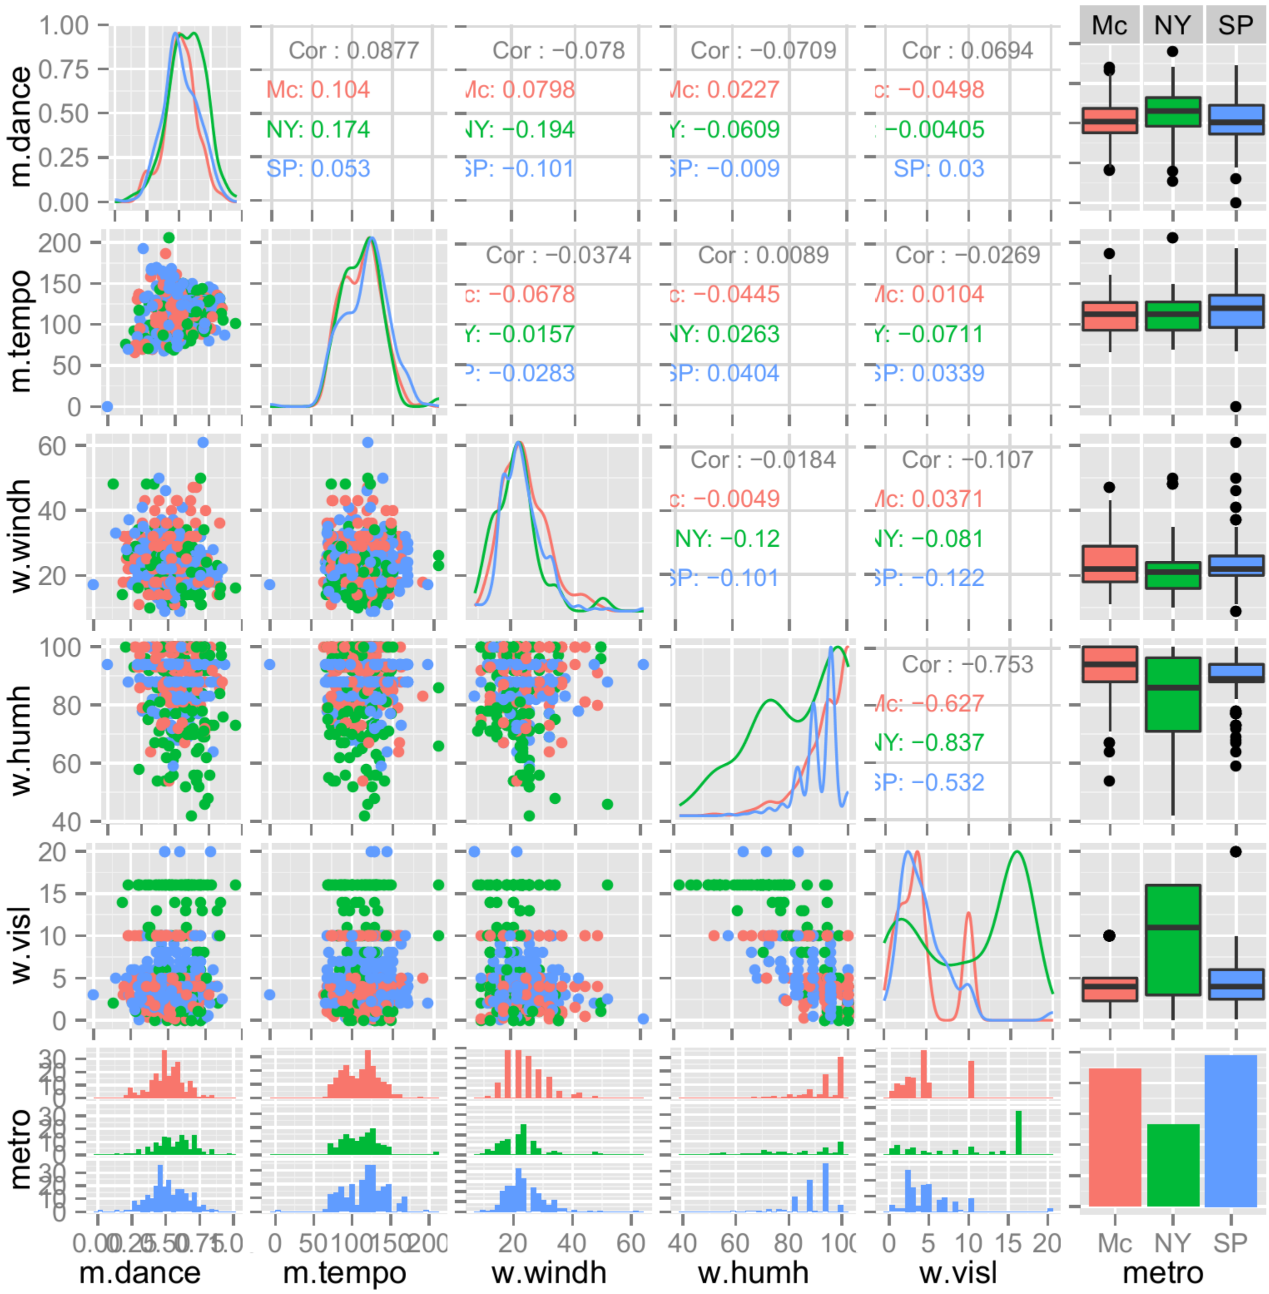
Mc
NY
SP
Cor : 0.0877
Mc: 0.104
NY: 0.174
SP: 0.053
Cor : −0.078
Cor : −0.0709
Cor : 0.0694
Cor : −0.0374
Cor : 0.0089
Cor : −0.0269
Cor : −0.0184
NY: 0.0263
Cor : −0.107
Mc: 0.0104
Cor : −0.753
m.dance
m.tempo
w.windh
w.humh
w.visl
metro
Mc
NY
SP

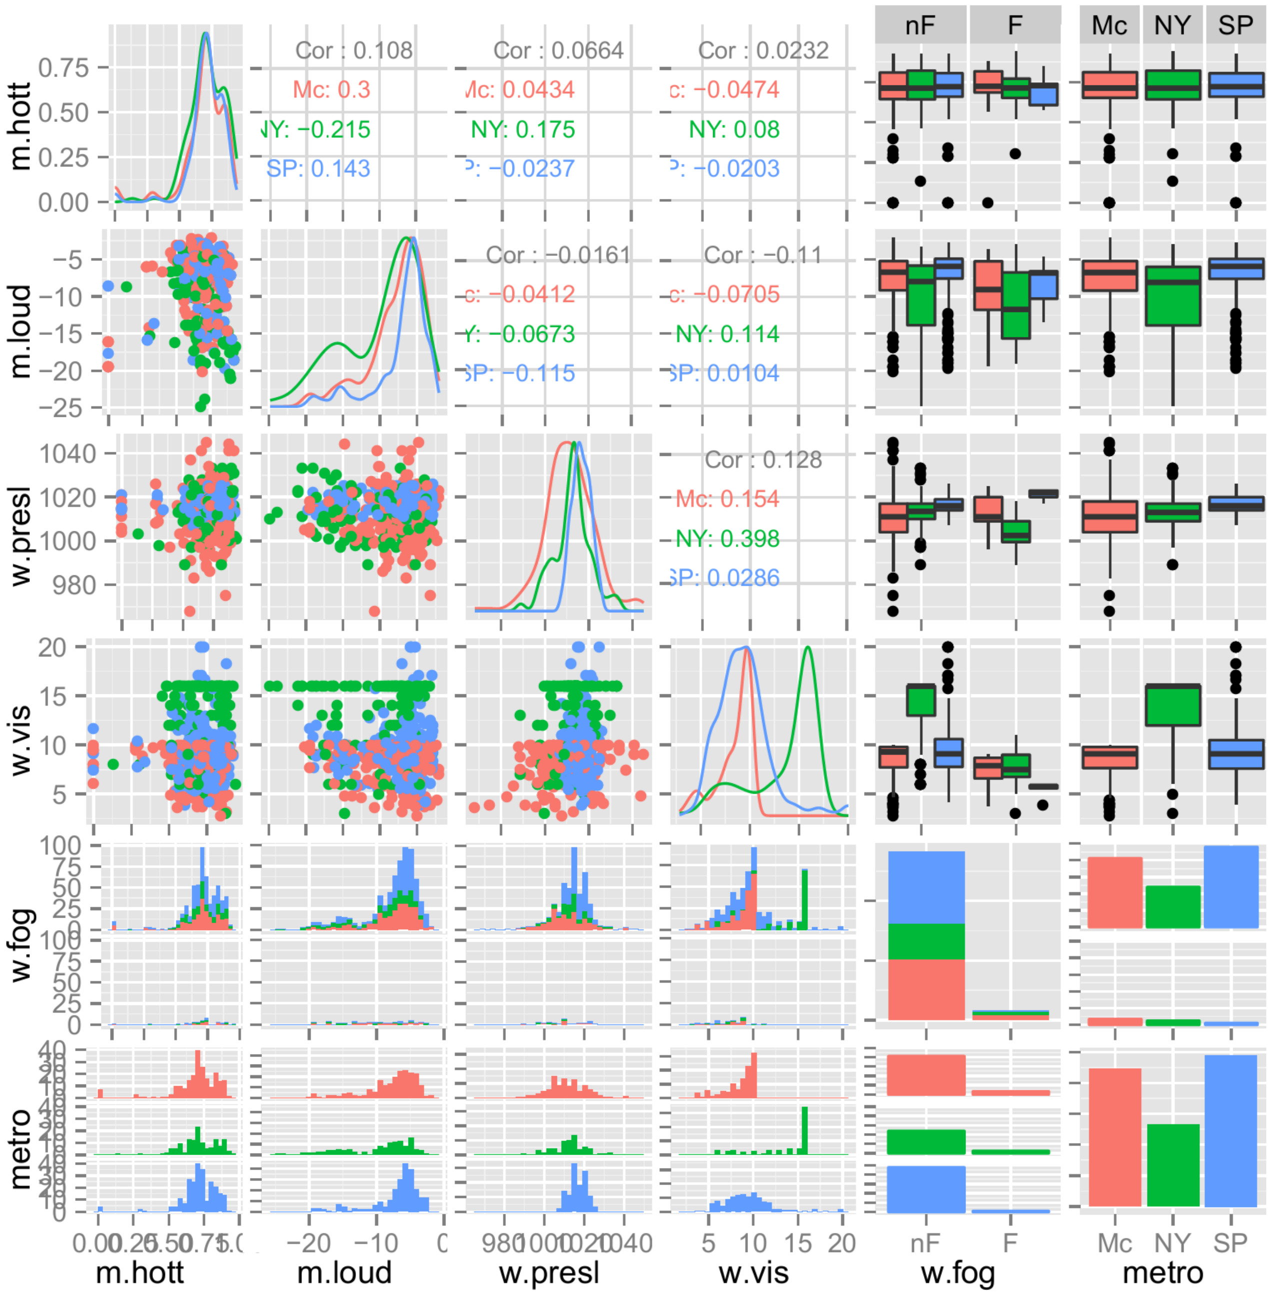

nF
F
Mc
NY
SP
m.hott
m.loud
w.presl
w.vis
w.fog
metro
Cor : 0.108
Mc: 0.3
NY: −0.215
SP: 0.143
Cor : 0.0664
Mc: 0.0434
NY: 0.175
SP: −0.0237
Cor : 0.0232
Mc: −0.0474
NY: 0.08
SP: −0.0203
Cor : −0.0161
Mc: −0.0412
NY: −0.0673
SP: −0.115
Cor : −0.11
Mc: −0.0705
NY: 0.114
SP: 0.0104
Cor : 0.128
Mc: 0.154
NY: 0.398
SP: 0.0286

Appendix C. Table of weather features in models

model	*response*	*rain*	*hail*	*fog*	*snow*	*thunder*	*tornado*	*temp*	*templ*	*temph*	*dew*	*dewl*	*dewh*	*press*	*pressl*	*pressh*	*wind*	*windl*	*windh*	*vis*	*visl*	*vish*	*hum*	*huml*	*humh*	*precip*	*snowd*	*snowf*
Moscow weekly	*loudness*	y***	n	y**	y***	n	n	y***	n	y.	y***	y***	n	y***	y*	y***	y***	y*	y***	y***	y***	n	n	y***	y.	y	n	n
	energy	y***	n	n	y***	y***	n	y***	y*	y**	y***	y**	y*	y***	n	n	y***	n	y*	y***	y***	n	n	y.	y**	y*	n	n
	speechiness	y***	n	n	y***	y***	n	y***	y***	y***	n	n	n	n	y.	y	y***	y***	y***	y***	y***	n	n	y***	y***	y*	n	n
	danceability	n	n	n	y*	y*	n	n	n	n	n	n	n	y***	y***	n	y**	y*	y*	y***	n	n	n	y***	y***	y*	n	n
	tempo	y*	n	n	y***	y***	n	n	n	y***	y***	y**	n	y***	y***	y	y***	y***	n	y***	n	n	n	y***	y***	y**	n	n
Moscow daily	*loudness*	n	n	n	y***	n	n	n	y***	y*	n	y	y**	y***	y***	y***	n	y***	y***	y***	y***	y***	n	y***	y***	n	n	n
	energy	y*	n	n	y***	y***	n	n	y***	y***	n	y*	y.	y***	y***	y***	n	y**	y***	y***	y***	y***	n	y***	y***	n	n	n
	speechiness	y***	n	n	y***	y***	n	y***	n	n	y**	y***	y***	n	y***	y***	y***	y**	y***	n	y***	y***	n	n	y***	y*	n	n
	danceability	n	n	n	y*	y***	n	y***	y**	y***	n	y*	y*	y***	n	y***	y*	y	y	y***	y***	n	n	y.	y***	y***	n	n
	tempo	n	y	n	n	y.	n	n	y***	n	y***	y**	y***	n	n	y***	y***	y***	n	y*	y***	n	n	y**	y***	n	n	n
New York weekly	*loudness*	y***	n	y***	y*	n	n	n	n	y***	n	y***	y***	y***	y***	n	y***	n	y***	y*	y***	n	n	y***	y***	y***	n	y
	energy	n	n	y***	n	n	n	n	y***	y***	n	y***	y***	y***	n	y*	n	n	y***	y*	y***	n	n	y	y***	y*	n	n
	speechiness	n	n	y**	y**	n	n	y***	y***	y***	n	y***	n	y**	y**	y**	y*	n	y***	y***	n	n	n	y***	n	y**	y	y*
	danceability	n	n	n	n	y**	n	y***	y***	y***	n	y***	n	y***	y	n	y*	n	y	y*	n	n	n	n	y***	y*	y***	y**
	tempo	y***	n	y.	n	n	n	n	y**	y***	y	n	y**	n	n	y***	y***	n	y***	y	y*	n	n	y*	n	y.	n	n
New York daily	*loudness*	y*	n	y***	y*	n	n	y***	y***	y***	y*	y***	y***	y***	y***	y***	y***	y***	y**	y*	y***	y***	n	y**	n	y**	y***	y***
	energy	y*	n	y***	n	n	n	y	y***	y***	n	y***	y***	n	y*	y*	y***	y***	n	n	y***	n	n	y***	y*	y*	y.	y***
	speechiness	y***	n	y***	y***	y.	n	y**	y**	y***	y.	y***	y	y**	y***	n	n	n	y***	n	y***	y*	n	n	n	y***	y***	y*
	danceability	y***	n	y*	n	n	n	y.	y***	y***	y***	y*	y***	y***	y***	y**	y*	y***	y*	n	y***	n	n	y***	y***	y***	y***	y
	tempo	n	n	y*	n	n	n	y***	y***	y***	y**	n	y.	n	y***	y***	y***	n	y***	y*	y***	n	n	n	n	y**	y***	n
Sao Paulo weekly	*loudness*	y***	y*	y***	n	y***	n	y***	y***	y***	y***	y***	y	y*	y***	y***	y***	n	y***	y***	n	y***	n	n	y***	y***	n	n
	energy	y***	n	n	n	y***	n	n	y***	n	y.	y***	n	y***	y**	n	y***	y***	y**	y*	y***	y**	n	y.	n	n	n	n
	speechiness	y***	n	n	n	y**	n	n	y***	n	y***	y**	n	y***	n	y***	y*	n	n	n	y*	n	n	n	y***	y	n	n
	danceability	y*	n	n	n	y***	n	y.	y*	n	y**	y***	y***	y***	n	y**	y***	y***	y***	n	y*	y***	n	y.	n	n	n	n
	tempo	y.	n	n	n	n	n	n	y*	y**	n	y***	y*	y***	n	y**	y***	n	n	n	n	y***	n	n	n	y**	n	n
Sao Paulo daily	*loudness*	y***	y.	y***	n	y***	n	y***	y**	y.	y***	y***	y***	n	y***	y***	y***	y***	y***	y*	y***	y***	n	y***	y***	y***	n	n
	energy	y***	n	n	n	y***	n	n	y**	y***	y***	y***	y***	y***	n	y***	y***	n	n	y***	y**	y***	n	y***	n	n	n	n
	speechiness	n	n	y*	n	y**	n	y**	y**	y***	y***	y.	y***	n	n	y***	y***	y*	n	y*	y*	y***	n	y***	n	n	n	n
	danceability	n	n	n	n	y***	n	n	n	n	n	y***	n	y*	y***	y***	y***	y.	y.	y***	y***	y	n	y***	y**	n	n	n
	tempo	y***	n	n	n	y**	n	n	y	y***	y***	y***	y***	y***	n	n	y***	y	y**	y***	y**	y	n	y***	y**	y***	n	n

Appendix D. Original chart and re-ranked chart containing 50 tracks (two examples)

#### chart Moscow_14_9_2011 ranked by listeners	#### chart Moscow_14_9_2011 ranked by cosine
#1 Kasabian - Days Are Forgotten (452)	#1 Kasabian - Let's Roll Just Like We Used To (418)
#2 Kasabian - Let's Roll Just Like We Used To (418)	#2 Kasabian - Man of Simple Pleasures (347)
#3 Kasabian - Goodbye Kiss (407)	#3 Kasabian - Switchblade Smiles (362)
#4 Kasabian - Velociraptor! (403)	#4 blink-182 - Snake Charmer (323)
#5 Kasabian - La Fee Verte (384)	#5 Adele - Rolling in the Deep (307)
#6 Kasabian - Re-Wired (378)	#6 Kasabian - I Hear Voices (367)
#7 Nirvana - Smells Like Teen Spirit (371)	#7 blink-182 - Kaleidoscope (301)
#8 Kasabian - I Hear Voices (367)	#8 Radiohead - Creep (240)
#9 blink-182 - Ghost On The Dance Floor (366)	#9 Kasabian - Acid Turkish Bath (Shelter from the Storm) (361)
#10 Kasabian - Switchblade Smiles (362)	#10 blink-182 - Love is Dangerous (266)
#11 Kasabian - Acid Turkish Bath (Shelter from the Storm) (361)	#11 Kasabian - Re-Wired (378)
#12 blink-182 - Natives (358)	#12 Muse - Hysteria (227)
#13 Kasabian - Man of Simple Pleasures (347)	#13 Nirvana - Lithium (197)
#14 blink-182 - After Midnight (344)	#14 blink-182 - Fighting The Gravity (254)
#15 blink-182 - Up All Night (339)	#15 Kasabian - La Fee Verte (384)
#16 blink-182 - Snake Charmer (323)	#16 blink-182 - Wishing Well (309)
#17 Kasabian - Neon Noon (322)	#17 Kasabian - Days Are Forgotten (452)
#18 blink-182 - Wishing Well (309)	#18 blink-182 - MH 4.18.2011 (280)
#19 Adele - Rolling in the Deep (307)	#19 Nirvana - Come as You Are (248)
#20 blink-182 - Heart's All Gone (306)	#20 blink-182 - After Midnight (344)
#21 blink-182 - Heart's All Gone Interlude (306)	#21 Kasabian - Underdog (230)
#22 blink-182 - Kaleidoscope (301)	#22 Red Hot Chili Peppers - The Adventures of Rain Dance Maggie (294)
#23 blink-182 - This Is Home (299)	#23 Nirvana - Smells Like Teen Spirit (371)
#24 Red Hot Chili Peppers - The Adventures of Rain Dance Maggie (294)	#24 Kasabian - Goodbye Kiss (407)
#25 blink-182 - MH 4.18.2011 (280)	#25 blink-182 - Ghost On The Dance Floor (366)
#26 Radiohead - Karma Police (271)	#26 Red Hot Chili Peppers - Ethiopia (244)
#27 blink-182 - Love is Dangerous (266)	#27 Radiohead - Paranoid Android (203)
#28 Red Hot Chili Peppers - Factory of Faith (264)	#28 Red Hot Chili Peppers - Brendan's Death Song (236)
#29 Red Hot Chili Peppers - Monarchy of Roses (260)	#29 blink-182 - Up All Night (339)
#30 Muse - Supermassive Black Hole (256)	#30 Muse - Starlight (228)
#31 blink-182 - Fighting The Gravity (254)	#31 Red Hot Chili Peppers - Californication (248)
#32 blink-182 - Even If She Falls (252)	#32 Muse - Uprising (203)
#33 Red Hot Chili Peppers - Californication (248)	#33 Muse - Supermassive Black Hole (256)
#34 Nirvana - Come as You Are (248)	#34 blink-182 - This Is Home (299)
#35 Red Hot Chili Peppers - Ethiopia (244)	#35 Red Hot Chili Peppers - Did I Let You Know (231)
#36 Radiohead - Creep (240)	#36 Radiohead - Karma Police (271)
#37 Red Hot Chili Peppers - Look Around (239)	#37 Red Hot Chili Peppers - Annie Wants a Baby (229)
#38 Red Hot Chili Peppers - Brendan's Death Song (236)	#38 The Killers - Somebody Told Me (223)
#39 Red Hot Chili Peppers - Did I Let You Know (231)	#39 Kasabian - Velociraptor! (403)
#40 Kasabian - Underdog (230)	#40 blink-182 - Even If She Falls (252)
#41 Red Hot Chili Peppers - Annie Wants a Baby (229)	#41 Red Hot Chili Peppers - Look Around (239)
#42 Muse - Starlight (228)	#42 Red Hot Chili Peppers - Goodbye Hooray (214)

#43 Muse - Hysteria (227)	#43 Red Hot Chili Peppers - Factory of Faith (264)
#44 The Killers - Somebody Told Me (223)	#44 blink-182 - Heart's All Gone (306)
#45 Radiohead - No Surprises (220)	#45 Red Hot Chili Peppers - Happiness Loves Company (197)
#46 Red Hot Chili Peppers - Goodbye Hooray (214)	#46 blink-182 - Natives (358)
#47 Muse - Uprising (203)	#47 Red Hot Chili Peppers - Monarchy of Roses (260)
#48 Radiohead - Paranoid Android (203)	#48 Kasabian - Neon Noon (322)
#49 Nirvana - Lithium (197)	#49 Radiohead - No Surprises (220)
#50 Red Hot Chili Peppers - Happiness Loves Company (197)	#50 blink-182 - Heart's All Gone Interlude (306)
Spearman's rho=0.360864345738, sig.=0.0100370623797	

#### chart Sao Paulo_22_8_2012 ranked by listeners	#### chart Sao Paulo_22_8_2012 ranked by cosine
#1 Lana Del Rey - Blue Jeans (527)	#1 Lana Del Rey - Video Games (480)
#2 Lana Del Rey - Born to Die (525)	#2 Lana Del Rey - Without You (276)
#3 The xx - Angels (497)	#3 Lana Del Rey - Off to the Races (353)
#4 Lana Del Rey - Video Games (480)	#4 Lana Del Rey - National Anthem (464)
#5 Lana Del Rey - National Anthem (464)	#5 Lana Del Rey - Dark Paradise (375)
#6 Carly Rae Jepsen - Call Me Maybe (457)	#6 Lana Del Rey - Born to Die (525)
#7 Foster the People - Pumped Up Kicks (452)	#7 Adele - Set Fire to the Rain (367)
#8 The xx - Chained (452)	#8 Katy Perry - Wide Awake (379)
#9 The xx - Fiction (439)	#9 Adele - Rolling in the Deep (324)
#10 The xx - Try (414)	#10 Lana Del Rey - Carmen (289)
#11 Lana Del Rey - Summertime Sadness (407)	#11 Lana Del Rey - Summertime Sadness (407)
#12 The xx - Reunion (400)	#12 One Direction - What Makes You Beautiful (290)
#13 The xx - Missing (388)	#13 The xx - Crystalised (276)
#14 Katy Perry - Wide Awake (379)	#14 Arctic Monkeys - Fluorescent Adolescent (363)
#15 The xx - Tides (377)	#15 Lana Del Rey - This Is What Makes Us Girls (297)
#16 Lana Del Rey - Dark Paradise (375)	#16 Oasis - Wonderwall (332)
#17 Rihanna - Where Have You Been (370)	#17 Arctic Monkeys - Mardy Bum (280)
#18 Adele - Set Fire to the Rain (367)	#18 Lady Gaga - Marry the Night (293)
#19 The xx - Unfold (363)	#19 Rihanna - Where Have You Been (370)
#20 Arctic Monkeys - Fluorescent Adolescent (363)	#20 Coldplay - Paradise (288)
#21 Lana Del Rey - Off to the Races (353)	#21 Katy Perry - Part of Me (281)
#22 The xx - Swept Away (347)	#22 Arctic Monkeys - 505 (298)
#23 The Strokes - You Only Live Once (343)	#23 Lana Del Rey - Blue Jeans (527)
#24 The xx - Our Song (343)	#24 Lady Gaga - Born This Way (278)
#25 Maroon 5 - One More Night (332)	#25 Nicki Minaj - Starships (293)
#26 Oasis - Wonderwall (332)	#26 The xx - Chained (452)
#27 Adele - Rolling in the Deep (324)	#27 The xx - Fiction (439)
#28 Arctic Monkeys - Teddy Picker (321)	#28 Arctic Monkeys - I Bet You Look Good on the Dancefloor (307)
#29 Foster the People - Houdini (310)	#29 Foster the People - Call It What You Want (304)
#30 Arctic Monkeys - I Bet You Look Good on the Dancefloor (307)	#30 Foster the People - Helena Beat (288)
#31 Foster the People - Call It What You Want (304)	#31 The Wanted - Glad You Came (281)
#32 Lana Del Rey - Radio (304)	#32 Demi Lovato - Give Your Heart a Break (303)
#33 Demi Lovato - Give Your Heart a Break (303)	#33 Lana Del Rey - Radio (304)

#34 Arctic Monkeys - 505 (298)	#34 Foster the People - Houdini (310)
#35 Lana Del Rey - This Is What Makes Us Girls (297)	#35 Arctic Monkeys - Teddy Picker (321)
#36 Lady Gaga - Marry the Night (293)	#36 The Strokes - You Only Live Once (343)
#37 Nicki Minaj - Starships (293)	#37 The xx - Tides (377)
#38 Gotye - Somebody That I Used to Know (292)	#38 Gotye - Somebody That I Used to Know (292)
#39 One Direction - What Makes You Beautiful (290)	#39 The Killers - Somebody Told Me (288)
#40 Lana Del Rey - Carmen (289)	#40 The xx - Swept Away (347)
#41 The Killers - Somebody Told Me (288)	#41 The xx - Reunion (400)
#42 Coldplay - Paradise (288)	#42 Carly Rae Jepsen - Call Me Maybe (457)
#43 Foster the People - Helena Beat (288)	#43 Arctic Monkeys - Crying Lightning (282)
#44 Arctic Monkeys - Crying Lightning (282)	#44 Foster the People - Pumped Up Kicks (452)
#45 The Wanted - Glad You Came (281)	#45 Maroon 5 - One More Night (332)
#46 Katy Perry - Part of Me (281)	#46 The xx - Missing (388)
#47 Arctic Monkeys - Mardy Bum (280)	#47 The xx - Try (414)
#48 Lady Gaga - Born This Way (278)	#48 The xx - Angels (497)
#49 Lana Del Rey - Without You (276)	#49 The xx - Unfold (363)
#50 The xx - Crystalised (276)	#50 The xx - Our Song (343)
Spearman's rho=-0.0767827130852, sig.=0.596121047798	