

# **Sad or Glad? Corpus Creation for Odia Poetry with Sentiment Polarity Information**

by

Gaurav Mohanty, Pruthwik Mishra, Radhika Mamidi

in

*19th International Conference on Computational Linguistics and Intelligent Text Processing  
(CICLing-2018)*

Hanoi, Vietnam

Report No: IIIT/TR/2018/-1



Centre for Language Technologies Research Centre  
International Institute of Information Technology  
Hyderabad - 500 032, INDIA  
March 2018

# Sad or Glad? Corpus Creation for Odia Poetry with Sentiment Polarity Information

Gaurav Mohanty, Pruthwik Mishra, and Radhika Mamidi

Language Technologies Research Center (LTRC)  
Kohli Center on Intelligent Systems (KCIS)  
International Institute of Information Technology-Hyderabad, India  
{gaurav.mohanty,pruthwik.mishra}@research.iiit.ac.in,  
radhika.mamidi@iiit.ac.in

**Abstract.** Resource poor languages, like Odia, inherently lack the necessary resources and tools for the task of sentiment analysis to give promising results. With more user-generated raw data readily available today, it is of prime importance to have annotated corpora from various domains. This paper is a first attempt towards building an annotated corpus of Odia poetry with sentiment labels. The annotated corpus is further used for sentiment classification using machine learning techniques in order to establish a baseline. Stylistic variations and structural differences between poetic and non-poetic texts make the task of sentiment classification challenging for the former. Using the annotated corpus of poems, we obtained comparable accuracies across various classification models. Linear-SVM outperformed other classifiers with a macro F1-Score of 0.68. The annotated corpus contains a total of 730 Odia Poems of various genres with a vocabulary of more than 23k words. Fleiss Kappa score of 0.83 was obtained which corresponds to near perfect agreement among the annotators.

## 1 Introduction

Sentiment analysis comprises of extraction and analysis of subjective information present in natural language data. The inception of Web 2.0 served as a gateway to rapid increase in user generated textual content. Opinions are expressed at an ever growing pace in current times on various social media websites. Hence, sentiment analysis systems are widely used for social media monitoring tasks, customer feedback and product review by several commercial organizations <sup>1</sup>. In the area of governance, public feedback via social media and various survey systems is being monitored on a large scale.

With more data available in the native vernacular, the task of sentiment analysis becomes challenging for resource-poor languages. These lack several essential tools and annotated corpora which aid in the task of opinion extraction. Odia is one such language. It is an Indo-Aryan language spoken in various parts

---

<sup>1</sup> [www.sas.com](http://www.sas.com)

of eastern India and has over 45 million native speakers spread across the globe. There is an abundance of Odia data in the form of stories, poems, news articles, blogs, etc over the internet. People have a preference over the genre of textual content they consume depending on their mood. Classification of such content on the basis of the sentiment they evoke, hence is useful. Annotated corpora would therefore be necessary, in order to build automated sentiment classifiers for the same. No such annotated corpus of poems currently exists in literature for any Indian language.

The corpus comprises of an annotated collection of over 700 Odia poems with sentiment information. Poems and songs are unique among other textual content as they do not follow the same syntactic structure and word order for a language. They contain sentiment information at entity, stanza as well as at document level. These poems have been classified as either of positive or negative class. The annotated class is cross-validated by also identifying several emotions that the poems evoke. The baseline experiments have also been conducted in order to compare performance of various classifiers on the annotated corpus.

This is the first corpus of Odia poems, with annotated sentiment information, existing in literature as per our knowledge. It is written in Odia script and hence avoids the pre-processing cost of text normalization. Other than sentiment identification, the corpus is meant to serve as a useful dataset in the task of emotion polarity detection. In terms of application, the emotion tags in the corpus can be used to train models to identify genre of poems. Models can also be trained which help identify user's mood based on the kind of content the user prefers. This can further be used to build recommendation systems.

The paper is divided into 5 sections. Section 2 briefly discusses related work. Section 3 elaborates on creation and annotation of the corpus. The adopted annotation scheme has been provided in the same. Inter-annotator agreement has also been calculated. Section 4 describes the experimental setup for training a model using various classifiers. This helps in establishing the baseline for sentiment classification of these Odia Poems. Possible future work using the annotated corpus is briefly discussed in Section 5.

## 2 Related Work

So far, sentiment analysis has been majorly focused around classification of non-poetic texts. Moreover, among songs and poems, research on the former is much more than that on the latter. Music classification has been carried out using lyrics [8], audio [13] and even multi-modal features [11] for English. Similar work has been carried out for mood classification of Telugu [1] and Hindi songs [15]. In the case of traditional literary works such as poetry, a lexicon creation methodology has been discussed for analyzing classical Chinese poetry [6]. The authors propose a weakly supervised approach based on Weighted Personalized Page Rank (WPPR) to create the sentiment lexicon. Such sentiment lexicons are useful for extracting aspect and sentence level sentiment information. Recently a linked WordNet based approach has been proposed in literature to create a sen-

timent lexicon for Odia [14]. Even though annotated corpora is available for several major languages in India, no such sizable corpus exists for poetic texts, let alone in Odia. Our work aims at bridging this gap.

### 3 Data Collection and Annotation

User-generated Odia text is readily available over the web. These exist in the form of blog posts, news articles, short stories, poetry, songs lyrics, etc. A good amount of these texts is however present in the form of images of the Odia text, rather than a scrapable text format. Manually converting every single image would be very time consuming. Automatic recognition systems for digitized Odia script documents do exist in literature [2]. Fortunately sufficient Odia content is also available in Odia script in utf-8 character encoding. A few good sources for the same include the Samaja News Website<sup>2</sup> and Odia Wikipedia<sup>3</sup>. For literature, Ame Odia Magazine Website<sup>4</sup> and Odia Gapa<sup>5</sup> serve as useful sources.

#### 3.1 Dataset Source

Odia poems and lyrics are commonly available in transliterated Roman script. For ease of processing, a source where text was available in Odia script was preferred. The **Ame Odia** website hence, was the choice for source data. It contains

**Table 1.** Initial Dataset Statistics

Initial Poem Count	788
Total number of Tokens	98782
Total Number of Unique Words	23532
Average Token Count per Poem	~125

a large and diverse collection of Odia poems along with short stories and blogs. The website has over 800 Odia poems, 400 short stories, 130 essays, and several other texts of various literary forms. It serves as a dynamic database with more and more literary text added to it on a regular basis. At the time of extraction, the website contained 788 Odia poems. These were collected along with meta-data information for each poem. Meta-data includes the title of the poem, name of the poet, and date of publication of the poem. Since poems differ from prose in syntactic structure and have stylistic variations, pre-processing becomes a necessary step. Instead of sentences, poems are composed of several stanzas

<sup>2</sup> <http://www.thesamaja.in>

<sup>3</sup> <http://or.wikipedia.org>


<sup>4</sup> <http://www.ameodia.com>

<sup>5</sup> <http://www.oddiagapa.com>

and these are sometimes numbered. Stanza numbering does not carry any sentiment information and can be treated similar to functional words. Hence these are removed from the poems. The name of the poet and the date of publication also do not serve the purpose of sentiment classification and therefore were not used in baseline experiments. The title of the poem is retained as it may carry sentiment information. Table 1 provides details on the initial statistics of the dataset before annotation.

### 3.2 Annotation Scheme

A well defined annotation scheme is necessary in order to assign proper sentiment labels to all poems. The task of sentiment analysis can be carried out at three different levels [12]. The identification of positive or negative sentiment is carried out at a defined level. Sentiment analysis can be done at an aspect level [7] or sentence level (or stanza) or at the level of the whole document [22]. In the case of poems, it is possible that different parts of the poem elicit different emotions. Since the task is to identify sentiment of the poem as a whole, annotation is carried out only at an overall document level. A polarity identification questionnaire and taxonomy of emotions is provided in order to help the annotators identify sentiment for each poem.



**Fig. 1.** Russell's Circumplex Model [18] classifying 28 Affect words on the basis of positive and negative polarity and arousal.

**Identifying emotions** Poems are the most sophisticated form of language [19]. They evoke several emotions in the mind of the reader. In order to detect these emotions, a proper taxonomy is necessary. Russell’s Circumplex Model of 28 affect words [18] serves as an appropriate reference for emotion identification [21]. The model spots several human emotions on a two dimensional plane of sentiment polarity and arousal as illustrated in Figure 1. For a given poem, the identified emotions are also tagged by the annotators in order to help validate the poem’s annotated sentiment.

**Polarity Identification Questionnaire** Once the emotion tags for a given poem are identified by an annotator, the following questionnaire is used to help determine the polarity label for the poem as a whole. Which of the following best describes the kind of language the poet is using?

- (1) For the whole poem, the poet is using positive language, such as expression of support, motivation, admiration, positive attitude, cheerfulness, forgiving nature, positive emotional state, etc. The emotional states identified are tending to the positive side of Russell’s model, for example, happy, excited, calm, serene, etc. - **Positive**
- (2) For the whole poem, the poet is using negative language, such as expressions of judgement, negative attitude, criticism, failure, sadness, negative emotional state etc. The emotional states identified are tending to the negative side of Russell’s model, for example, alarmed, angry, miserable, tired, etc. - **Negative**
- (3) The poet is majorly using positive language with a minority in the form of negative language. - **Positive**
- (4) The poet is majorly using negative language with a minority in the form of positive language. - **Negative**
- (5) The poet is using a mix of both positive and negative language, where it is difficult to claim majority of one over the other.

Focus has been given to the language used by the poet. The emotions were identified based on the kind of language used and not by making prior assumptions on what the poet’s possible state of mind was when writing the poem. Annotators should not worry about whether they agree or disagree with the poet’s views. As the poems are to be classified into two classes, options 3 and 4 focus on determining the sentiment which is expressed more often throughout the poem. The questionnaire along with the identified emotion tags should help determine the dominant sentiment class for each poem. Having option 5 helps annotators in case they are confused about the dominant sentiment for a given poem.

### 3.3 Evaluation of Dataset

Each Odia poem was annotated as positive or negative by three annotators who are native Odia speakers who speak and write in the language on a daily

basis. Poems satisfying options 1 and 3 from the questionnaire were tagged as positive whereas those satisfying options 2 and 4 were tagged as negative. Poems satisfying option 5 were separated from the dataset for future study.

Each annotator was to independently annotate the poems without any communication with the other annotators. The name of the poet was not provided to the annotators as this might induce preconditioned bias. For example, certain poets always write poems which evoke emotions of sadness or anger. Hence the annotator might have a bias towards annotating such poems as negative, given that the annotator knows the name of the poet.

After the annotation process, 342 poems were tagged as positive whereas 388 poems were tagged as negative. 58 poems were classified as option 5 from the questionnaire. Since the scope of this work entails only positive and negative sentiment classification, these are best kept separated from the final corpus for now. The final classification of a poem was determined by majority rule over all three annotations. In order to capture inter-annotator agreement, Fleiss Kappa<sup>6</sup> score for the annotated sample set was also calculated. Inter-Annotator Agreement is a measure of how well the annotators make the same annotation decision for the same category. Fleiss Kappa score is calculated with three annotators for two categories (positive/negative) as parameters. A balanced sample set of 342 positive and 342 negative poems is used for Fleiss Kappa [10]. A score of  $\kappa = 0.83$  is reported for the annotated corpus which corresponds to "almost perfect agreement". Results of annotation are presented in Table 2.

**Table 2.** Results of Annotation

	Positive	Negative	Total
Poem Count	342	388	730
Token Count	40546	52142	92688
<b>Fleiss' Kappa agreement score = 0.83</b>			
Removed poems (opt 5) = 58			

## 4 Baseline for Sentiment Classification

In order to establish baseline results for the annotated corpus, a few experiments were conducted. The task was to classify Odia poems as carrying positive or negative sentiment by training appropriate classification models. Initially three different classifiers were employed for this task and the results of each were compared. Term frequency-inverse document frequency (TF-IDF) [20] was used to create a vector representation for an entire poem. We also explore usage of character level n-grams and word embeddings to evaluate the performance of these classifiers.

<sup>6</sup> [https://en.wikipedia.org/wiki/Fleiss'kappa](https://en.wikipedia.org/wiki/Fleiss%27kappa)

## 4.1 Experimental Setup

The dataset was split into a ratio of 4:1 for the purpose of training and testing. For initial experiments, TF-IDF features for word n-grams and character n-grams were used for classification. Subsequently, word and character embeddings were also used as features for training the classification models. Experiments were conducted using 'scikit-learn' [16], an open source Python library<sup>7</sup>. Precision, Recall and F1-score are the three evaluation metrics which were calculated using 5-fold cross-validation.

## 4.2 Classifiers

The following three classifiers were used for initial baseline experiments:

- **Naive Bayes** : In Naive Bayes, for the task of text classification, the instance is assigned to the class having the highest conditional probability  $P(C|X)$ , where  $C$  is the sentiment class and  $X$  is the group of words for that instance.
- **Logistic Regression** : Logistic Regression is a multi-class logistic model which is used to estimate the probability of a class response based on predictor variables which have one or more independent variables that establish an outcome. The expected values of predictor variables are formed based upon the combination of values taken by the predictors.
- **SVM** : Support Vector Machine [4] is a supervised learning algorithm which makes use of a hyperplane in order to identify the decision boundary. Each data item is plotted as a point in n-dimensional space with the value of each feature being the value of a particular coordinate. The hyperplane divides these data points into separate classes.

Linear-SVM was used for training the classification model when using word and character embeddings as features.

## 4.3 TF-IDF Features

TF-IDF assigns weights to words (or n-grams) and serves as a statistical measure for evaluating how important a word is to a document in a corpus. TF-IDF was calculated for unigrams, bigrams and trigrams. **Table 3** illustrates the results of the same for the aforementioned classifiers.

**Using character n-grams** Even though 730 poems is a sizable corpus for the task at hand, it doesn't show a significant increase in accuracy especially with added bi-gram and tri-gram features. This is because most bi-grams and tri-grams occur sparsely in the entire corpus. In order to tackle the problem of sparsity, we conducted experiments using n-grams at character level. For the baseline, 2-6 and 3-6 character n-grams<sup>8</sup> were used to calculate character level TF-IDF features. The results of the same are illustrated in **Table 4**.

<sup>7</sup> <http://www.scikit-learn.org>

<sup>8</sup> Read as 2 to 6 and 3 to 6 character n-grams respectively

**Table 3.** Sentiment analysis with Word Level TF-IDF Features

Model	Features	Class	Precision	Recall	F1-Score
Linear-SVM	uni	Negative	<b>0.682</b>	<b>0.707</b>	<b>0.691</b>
		Positive	0.651	0.622	0.632
	uni-bi	Negative	0.675	<b>0.762</b>	<b>0.713</b>
		Positive	<b>0.685</b>	0.582	0.625
	uni-bi-tri	Negative	0.664	<b>0.796</b>	<b>0.727</b>
		Positive	<b>0.722</b>	0.575	0.630
Naive Bayes	uni	Negative	<b>0.704</b>	0.621	0.652
		Positive	0.630	<b>0.705</b>	<b>0.658</b>
	uni-bi	Negative	0.656	<b>0.750</b>	<b>0.661</b>
		Positive	<b>0.675</b>	0.504	0.523
	uni-bi-tri	Negative	0.640	<b>0.778</b>	<b>0.629</b>
		Positive	<b>0.684</b>	0.396	0.406
Logistic Regression	uni	Negative	0.669	<b>0.761</b>	<b>0.707</b>
		Positive	<b>0.682</b>	0.573	0.615
	uni-bi	Negative	0.626	<b>0.872</b>	<b>0.722</b>
		Positive	<b>0.759</b>	0.406	0.510
	uni-bi-tri	Negative	0.595	<b>0.924</b>	<b>0.718</b>
		Positive	<b>0.777</b>	0.279	0.381

**Table 4.** Sentiment analysis with Character Level TF-IDF Features

Model	Features	Class	Precision	Recall	F1-Score
Linear-SVM	(2-6)-gram	Negative	0.681	<b>0.796</b>	<b>0.727</b>
		Positive	<b>0.722</b>	0.575	0.630
	(3-6)-gram	Negative	0.681	<b>0.808</b>	<b>0.734</b>
		Positive	<b>0.728</b>	0.569	0.631
Naive Bayes	(2-6)-gram	Negative	<b>0.718</b>	0.643	0.658
		Positive	0.655	<b>0.710</b>	<b>0.663</b>
	(3-6)-gram	Negative	<b>0.716</b>	0.654	<b>0.660</b>
		Positive	0.663	<b>0.700</b>	0.659
Logistic Regression	(2-6)-gram	Negative	0.634	<b>0.855</b>	<b>0.720</b>
		Positive	<b>0.756</b>	0.439	0.534
	(3-6)-gram	Negative	0.633	<b>0.876</b>	<b>0.726</b>
		Positive	<b>0.781</b>	0.419	0.520

**Observation** As observed in Table 3, Linear-SVM performs better with increasing n-grams. It is easy to mistake Logistic Regression to be performing at par with Linear-SVM. However, the former’s performance drops drastically for positive class across the table. Linear-SVM, on the other hand, provides an overall better prediction for both classes, among all classifiers. Through Table 4, it is observed that usage of character level n-grams outperformed that of word level n-grams for all classifiers. This comes from the fact that many different words (or n-grams) can share the same character prefixes. Words with common character prefixes should have similar level of importance to a poem. Even when using

character level n-grams, Linear-SVM outperformed other classifiers in terms of overall prediction for both positive and negative class.

Experimental results show that Precision, Recall and F1-score for poems with negative sentiment are consistently higher than ones with positive sentiment. This could be due to existence of more explicit negative words than positive ones as shown in examples from two different poems in Table 5. Several poems manually classified as positive did not have any explicit positive words, yet expressed overall positive sentiment at stanza level(s) and document level. This is because positive sentiment is sometimes not carried through just affect words, but through the overall meaning of the utterance.

**Table 5.** Examples of explicit positive/negative words in poem stanzas

Utterance(Roman script)	English Gloss	Theme	Words	Class
Au ethara pachaku dekhani ho sanghaatha daga daga kari chaala, agaku, ahuri agaku	Do not look back Oh friend, Keep moving forward and forward	Motivating	-	Positive
Anyaaya anithi badhe nithi nithi dayabahi rakhe nighaa	Injustice and dishonesty increases day by day while you keep watching	Complaining about God	Anyaaya Anithi	Negative

#### 4.4 Word and Character-based Embeddings as Features

**Source** We learned 50-dimensional Glove [17] word embeddings on Odia news articles corpus<sup>9</sup> containing 500k sentences and 127k unique tokens. Embeddings were computed with Glove parameters set to default. It is to be noted that only 10k of these tokens overlapped with the Poem corpus vocabulary(  $\sim 23k$  ).

**Methodology** In order to get vector representation of a poem, the word embeddings for individual words in the poem were used. The mean of all word embeddings for words in a poem was calculated and used as the vector representation for the poem. Linear-SVM was adopted to train the classification model. Since every word in the poem corpus does not have a word embedding, we adopted two different methods to tackle this problem. These are outlined as follows:

- **Mean of Available Word Embeddings:** Only use available word embeddings for words in a given poem to calculate its mean.
- **Use Character-based Embeddings:** In order to obtain word embedding for an out-of-vocabulary (OOV) word, we take the mean of the character

<sup>9</sup> Corpus scrapped from [http://www.thesamaja.com/news\\_archive.php](http://www.thesamaja.com/news_archive.php)

embeddings which make up the word. A character’s embedding was calculated from the Glove vectors by taking the mean of embeddings of all the words in news corpus vocabulary in which that character occurred.

**Table 6.** Sentiment analysis with word and character-based embeddings

Model	Class	Precision	Recall	F1-Score
Mean (Word Emb)	Negative	0.663	0.763	0.706
	Positive	0.682	<b>0.562</b>	0.610
Mean (Word+Char Emb)	Negative	<b>0.671</b>	<b>0.805</b>	<b>0.728</b>
	Positive	<b>0.720</b>	0.555	<b>0.619</b>

**Observation** As illustrated in Table 6, using character embeddings to compute word embeddings for OOV words does show some improvement in performance by the classifier. Comparing this with Table 4 shows that Precision, Recall and F1-Score are comparable to that of using TF-IDF character n-gram features to train the SVM classifiers.

**Neural Network - BiLSTM** Bidirectional LSTMs (BiLSTM) [5] often outperform other neural network architecture while representing a sequence because it takes into account the left as well as right context present in the sequence. Hence, we used the BiLSTM model present in the Keras [3] deep learning framework. The BiLSTM network gave a 50-dimensional representation of a poem which was then used to predict its sentiment. The output layer had 2 nodes with softmax activation function to assign probabilities to possible outcomes. For training the model, we used sparse categorical cross entropy loss function with Adam [9] optimizer and 25 epochs. As can be seen in Table 7, BiLSTM does not perform as well as other classification models. We attribute this to the relatively small number of samples used to train the BiLSTM model.

**Table 7.** Summary of results based on different features

Model	Features	Precision	Recall	Macro-F1
Linear-SVM	(3-6 char n-gram)	<b>0.70</b>	<b>0.69</b>	<b>0.68</b>
Linear-SVM	(Word+Char Emb):Mean	<b>0.70</b>	0.68	0.67
BiLSTM	(Word+Char Emb)	0.64	0.63	0.63

## 5 Conclusion and Future Work

In this paper, we present the first corpus of Odia poems of diverse themes annotated with sentiment information. These poems are manually labelled as either

having positive or negative sentiment. In this work, we have described an annotation scheme in which annotators make use of a polarity identification questionnaire along with taxonomy of emotions, in order to assign proper labels to these poems. The emotion tags identified for the poems are also included in the annotation as meta-data. Furthermore, we have also compared the performance of three classifiers on this annotated corpus. Among all the three, Linear-SVM provides an overall better prediction for both classes. Classification models have been built using both word level and character level TF-IDF features. The latter outperforms word level features across all the three classifiers in terms of performance. Using word embeddings as features, it is observed that Linear-SVM gives results comparable to that of TF-IDF character n-gram features. There is sufficient room for addition of domain specific features which in turn can provide better performance. Hence, the current results may serve as a good baseline for the task of sentiment classification in Odia poems using the annotated corpus.

Usage of sentiment lexicon for sentiment extraction at aspect and stanza level can further improve performance of these sentiment classifiers. The baseline experiments primarily focused on short sequence of words. The stanzas in poetry are usually much longer than sentences in prose. As mentioned previously, BiLSTMs require a much larger training data to work with; hence in the future we intend to focus on annotation of poems at stanza level and aspect level.

Through this paper, we hope the annotated corpus of Odia poems would serve as a good resource in order to help evaluate research in sentiment analysis in Odia, especially for emotion classification in poetic texts.

## References

1. Abburi, H., Akkireddy, E.S.A., Gangashetti, S., Mamidi, R.: Multimodal sentiment analysis of telugu songs. In: Proceedings of the 4th Workshop on Sentiment Analysis where AI meets Psychology (SAAIP 2016) co-located with 25th International Joint Conference on Artificial Intelligence (IJCAI 2016), New York City, USA, July 10, 2016. pp. 48–52 (2016), <http://ceur-ws.org/Vol-1619/paper8.pdf>
2. Chaudhuri, B.B., Pal, U., Mitra, M.: Automatic recognition of printed oriya script. In: 6th International Conference on Document Analysis and Recognition (ICDAR 2001), 10-13 September 2001, Seattle, WA, USA. pp. 795–799 (2001), <https://doi.org/10.1109/ICDAR.2001.953897>
3. Chollet, F., et al.: Keras (2015)
4. Cortes, C., Vapnik, V.: Support vector machine. *Machine learning* 20(3), 273–297 (1995)
5. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks* 18(5), 602–610 (2005)
6. Hou, Y., Frank, A.: Analyzing sentiment in classical chinese poetry. In: Proceedings of the 9th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities, LaTeCH@ACL 2015, July 30, 2015, Beijing, China. pp. 15–24 (2015), <http://aclweb.org/anthology/W/W15/W15-3703.pdf>
7. Hu, M., Liu, B.: Mining opinion features in customer reviews. In: Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference

- on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA. pp. 755–760 (2004), <http://www.aaai.org/Library/AAAI/2004/aaai04-119.php>
8. Hu, Y., Chen, X., Yang, D.: Lyric-based song emotion detection with affective lexicon and fuzzy clustering method. In: Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR 2009, Kobe International Conference Center, Kobe, Japan, October 26-30, 2009. pp. 123–128 (2009), <http://ismir2009.ismir.net/proceedings/PS1-12.pdf>
  9. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
  10. Landis, J.R., Koch, G.G.: The measurement of observer agreement for categorical data. *Biometrics* pp. 159–174 (1977)
  11. Laurier, C., Sordo, M., Serrà, J., Herrera, P.: Music mood representations from social tags. In: Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR 2009, Kobe International Conference Center, Kobe, Japan, October 26-30, 2009. pp. 381–386 (2009), <http://ismir2009.ismir.net/proceedings/OS5-4.pdf>
  12. Liu, B.: *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies, Morgan & Claypool Publishers (2012), <https://doi.org/10.2200/S00416ED1V01Y201204HLT016>
  13. Lu, L., Liu, D., Zhang, H.: Automatic mood detection and tracking of music audio signals. *IEEE Trans. Audio, Speech & Language Processing* 14(1), 5–18 (2006), <https://doi.org/10.1109/TSA.2005.860344>
  14. Mohanty, G., Kannan, A., Mamidi, R.: Building a sentiwordnet for odia. In: Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, WASSA@EMNLP 2017, Copenhagen, Denmark, September 8, 2017. pp. 143–148 (2017), <http://aclanthology.info/papers/W17-5219/w17-5219>
  15. Patra, B.G., Das, D., Bandyopadhyay, S.: Multimodal mood classification framework for hindi songs. *Computación y Sistemas* 20(3), 515–526 (2016), <http://www.cys.cic.ipn.mx/ojs/index.php/CyS/article/view/2461>
  16. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., VanderPlas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in python. *Journal of Machine Learning Research* 12, 2825–2830 (2011), <http://dl.acm.org/citation.cfm?id=2078195>
  17. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
  18. Russell, J.A., Pratt, G.: A description of the affective quality attributed to environments. *Journal of personality and social psychology* 38(2), 311 (1980)
  19. Slinn, E.W.: *Victorian Poetry as Cultural Critique: The Politics of Performative Language*. University of Virginia Press (2003)
  20. Sparck Jones, K.: A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation* 28(1), 11–21 (1972)
  21. Thayer, R.E.: *The biopsychology of mood and arousal*: Oxford university press. New York (1989)
  22. Turney, P.D.: Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA. pp. 417–424 (2002), <http://www.aclweb.org/anthology/P02-1053.pdf>