

Sentiment Analysis of tweets

Topic: lemonade

A Spark Streaming application has been created using Scala that continuously reads data from Twitter about a topic.

A stream of status updates is obtained using TwitterUtils.

The sentiments associated with the statuses are obtained using Stanford CoreNLP library.

The sentiment data is then sent to ElasticSearch for analysis.

To exchange data between the Spark application and ElasticSearch, Kafka is used as the broker.

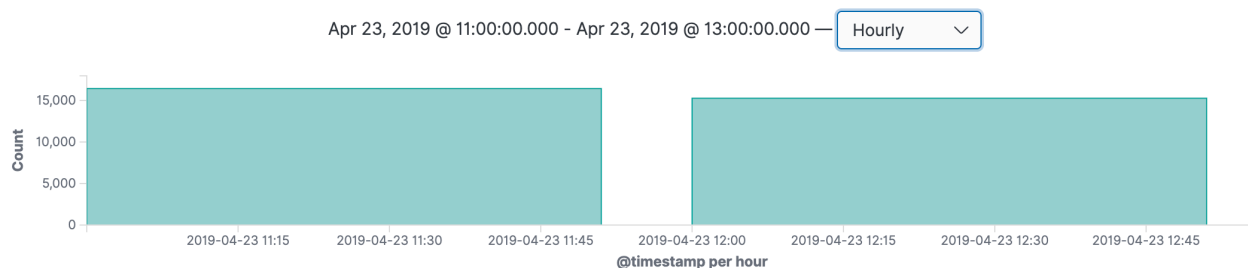
To ingest data from the Kafka server and map it to the required ElasticSearch server, LogStash has been used.

Finally, the sentiments are visualized in Kibana.

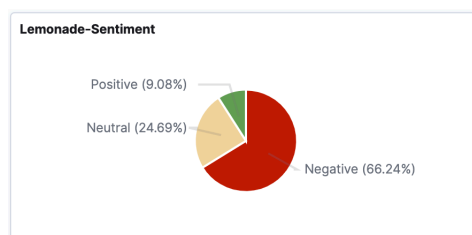
Here, Twitter statuses have been obtained using keyword 'lemonade' as it is one of the trending topics during the arrival of spring. The sentiments of the same have been obtained and visualized to find what percentage of the tweets have 'positive', 'negative', and 'neutral' sentiments.

Pie-chart snapshots in Kibana showing how the sentiment for 'lemonade' varied over time have been saved.

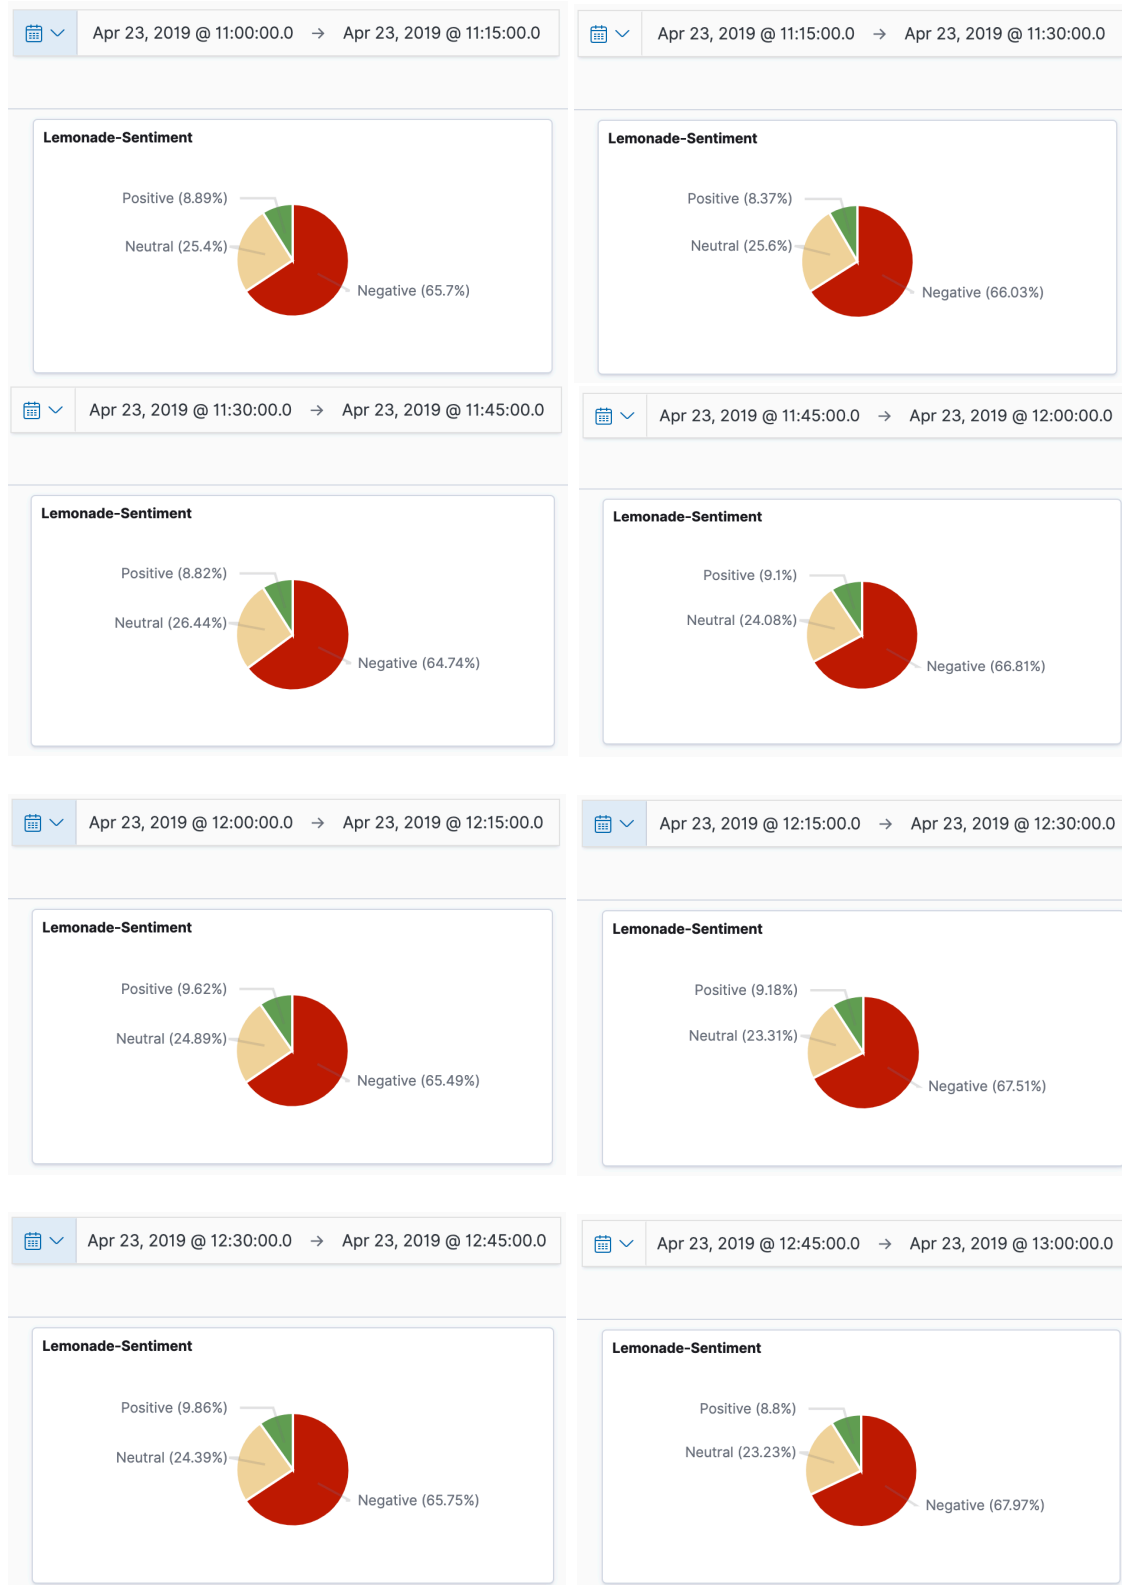
The hourly tweet count involving 'lemonade' on Apr 23, 2019 from 1100 to 1300 hrs. is as below:



Overall sentiment classification during the time is as below:



The variation of the sentiments by 15 min windows is as below:



We can conclude that the sentiments do not vary much and the majority of tweets containing 'lemonade' are associated with negative sentiments.