IBIMA Publishing
Communications of International Proceedings
https://ibimapublishing.com/p-articles/44AI/2024/4424924/
Vol. 2024 (22), Article ID 4424924

Youtubenomics – The Importance of the Recommendation Algorithm for Financial Results and Ad Revenue*

Krzysztof SIUDA

Wrocław University of Economics and Business, Department of Mathematical Economics, Poland

Correspondence should be addressed to: Krzysztof SIUDA, krzysztof.siuda@ue.wroc.pl

* Presented at the 44th IBIMA International Conference, 27-28 November 2024 Granada, Spain

Abstract

The main source of revenue for the YouTube platform is advertising. YouTube acts as an intermediary between advertisers and viewers, which means that the primary goal of the service is to maximize the number of displayed ads by increasing the average daily time spent on the platform by users, but also to precisely target ads to the right audience groups and create a favorable environment for the presentation of the aforementioned ads. The above-mentioned goals can be effectively realized on such a large platform only with the help of an efficient recommendation algorithm, which necessitates the need for its constant improvement. The article proposes the use of comment analysis via natural language processing to select the right content to promote in order to generate higher ad revenues due to their greater effectiveness and thus value. Evaluating the sentiment of comments is proving to be a major challenge due to the difficulty in recognizing context, sarcasm, irony, and mockery. Associating polarity with kurtosis and skewness significantly facilitates the recognition of negative and positive comments.

Keywords: YouTube, NLP, Sentiment Polarity, Recommendation Algorithm

Introduction

YouTube is a global platform for sharing and watching video content. The service was established in 2005, and was acquired by Google (now part of Alphabet Inc.) back in 2006. YouTube in a broader context can be categorized as a social media platform, for instance because of the ability to comment and rate uploaded content, i.e. to interact with other viewers and creators. Regardless of how we define the relevant market in which YouTube competes and the metrics used to measure the success of the platform, YouTube is one of the most important companies in the market. In fact, if YouTube is considered a social network, only Facebook has more monthly active users. Facebook in the April of 2024 had more than 3 billion active users, and YouTube 2.5 billion. Instagram and WhatsApp rounded out the podium with 2 billion users (We Are Social, DataReportal, Meltwater, 2024). In 2023, YouTube generated \$31.5 billion in ad revenue (Alphabet Inc., 2023), which is about 10% of Google's total revenue (Statista, 2024), while Facebook generated \$36.3 billion in revenue in the same year, Instagram \$21.3 billion, and TikTok \$10.1 billion (bizcommunity.com, 2024). Table 1 shows the revenue generated by YouTube between 2017 and 2023, share of these revenues in the total revenues of Google and the number of users in those years. In addition, changes in revenue and the number of users year-over-year were calculated, as well as the change relative to the base year - 2017. It can be noted that in the last included year relative to 2017, the number of users increased by 80%, and revenue by almost 290%. Only in 2022 the growth rate of the number of users was higher than the growth rate of revenues. In the remaining years, revenues grew more than proportionally relative to the number of users. This means that at the current stage of YouTube's development, it is more difficult to find and convince new people to use the platform, but there is great potential to optimize the platform's performance by improving the offer to potential advertisers. A key element is the higher effectiveness of the recommendation algorithm, which will not only keep users on the platform for longer throughout the day, but improve the

Cite this Article as: Krzysztof SIUDA, Vol. 2024 (22) “Youtubenomics – The Importance of the Recommendation Algorithm for Financial Results and Ad Revenue " Communications of International Proceedings, Vol. 2024 (22), Article ID 4424924, https://doi.org/10.5171/2024.4424924

experience of using the platform, which will positively affect the reception of the ads presented. Besides, the reception of ads is influenced by the content being viewed, and well-performing recommendations can direct viewers' attention to content that not only produces engagement, but also positive emotions. This shows how important, from the point of view of corporate finances and increasing business efficiency, is the interaction of financial analytics activities, in the broader business view, and improving the technical layer with new tools - in this case, natural language processing to better identify users' emotions. The share of YouTube ad revenue in Google's total revenue increased from 7.3% in 2017 to 10.25% in 2023. The highest shares were recorded during the pandemic (2020 - 2021), which was directly related to lockdowns and thus increased the amount of time people spent on entertainment without being able to leave their homes.

Table 1: Annual Financial and User Growth Overview of Youtube (Alphabet Inc., 2017-2023)

Year	Revenue (billion USD)	Share of Google revenues (%)	Number of Users (billion)	Revenue YoY Change (%)	Users YoY Change (%)	Revenue Change vs 2017 (%)	User Growth vs 2017 (%)
2017	8,1	7,3	1,5	-	-	0%	0%
2018	11,1	8,1	1,9	37%	27%	37%	27%
2019	15,1	9,4	2	36%	5%	86%	33%
2020	19,7	10,9	2,3	30%	15%	143%	53%
2021	28,8	11,2	2,5	46%	9%	256%	67%
2022	29,2	10,33	2,56	1%	2%	260%	71%
2023	31,5	10,25	2,7	8%	5%	289%	80%

Market share can be measured by the time users spend on a given platform in a unit of time. In this case, YouTube loses out to Netflix and TikTok, which are used by adult Americans in 2024 for 62 and 58 minutes each day, respectively. In comparison, YouTube's daily usage time is only 49 minutes (Insider Intelligence, 2023). Indians are the largest ethnic group on YouTube accounting for 476 million users, followed by Americans with 238 million, and Brazilians close out the podium with 147 million (DataReportal, We Are Social, Meltwater, 2024). In addition to advertising revenue, revenue from paid subscriptions to YouTube Music and YouTube Premium services are also important. Since 2020, the number of users paying for the aforementioned services has more than tripled to 100 million (YouTube, 2024). YouTube allows any person to share their digital content without restrictions. This also means that some people will use the service as a private network storage drive, bringing no benefit to the platform, but only generating costs. This means that the platform must balance such attitudes by effectively finding and using publicly available content to generate interest and thus advertising revenue. Most important for that purpose is an efficient recommendation algorithm that functions as a content curator, promoting those videos that are most likely to interest the platform's users based on a diverse and constantly evolving set of parameters identifying such content. YouTube's big challenge is to create a platform where videos of any length (from short to hours long) and live broadcasts are properly valued by the recommendation algorithm. In 2005, 2 million new videos were shared every day, in 2020, 500 hours of new videos were uploaded every minute, this illustrates the remarkable growth of the platform over the years, but most importantly, the enormous challenge of analyzing data and creating a recommendation system (Eksombatchai et al., 2018). Recommendation algorithms are a key part of how such large platforms operate, allowing users to navigate the site and discover content that may be of interest to a particular person. In the early days of YouTube, recommendations were created based on simple metrics such as number of views, likes and comments. This approach focused on the quality of the video while ignoring user preferences (Davidson et al., 2010). In subsequent years this has been recognized, and between 2010 and 2015 watch time, a measure of viewer engagement, became an important metric. Such a measure is in every way better than a simple view count, as YouTube aims to keep viewers on its platform for as long as possible in order to display enough ads. A 10-minute video that is watched for 1 minute on average is obviously less profitable for the platform than a video of the same length, but watched for 50% of its total run time. Algorithms currently use machine learning to predict what content a particular user might be interested in based on information collected on that user, such as age, gender, viewing history, likes, time of day, and device type (Babbar, Anand and Aggrawal, 2024). YouTube has clearly moved away from objectively evaluating content and has taken the direction of providing users with a personalized experience. Challenges include aspects such as the lack of data on videos, how users prefer newly added content, and selecting videos from the top of the recommended list (Zhao et al., 2019). This paper aims to indicate the usefulness of adding another parameter within the recommendation algorithm and the actual feasibility of implementing such a solution, which is based on sentiment analysis in the comments section of Youtube videos. More effectively recommending content to users that not only is interesting, but also evokes positive emotions can positively affect the financial results of the service by increasing the time spent on Youtube, and positive emotions should positively affect the reception of the presented ads. Better reception of advertisements means greater effectiveness of promotional campaigns (Vrtana and Krizanova, 2023), that is, in the context of Youtube, it increases its competitiveness in the advertising market and may lead to favorable changes in advertising prices for Youtube. Emotions that are generated by advertisements are of great importance in terms of product perception (Holbrook and O'Shaughnessy, 1984; Weinberger and Gulas, 1992) and purchasing decisions (Cartwright, McCormick and Warnaby, 2016), but no less important are the emotions of potential customers just before they watch the ads (Batra and Stayman, 1990). Youtube has no influence on the content of the ads, but it can try to provide optimal conditions for advertisers to present their advertising campaigns. Precisely for this reason, YouTube is constantly making changes to its recommendation algorithm, and understanding the content of comments seems to be the very best way to assess what emotions the video being viewed provides, which is the foundation for the commercials displayed during it. Other methods of evaluating content do not allow to detect nuances in the emotions felt by viewers.

Methodology

Natural language processing (NLP) is a field of artificial intelligence that enables computers to interact with human languages by using advanced techniques to analyze and process text or speech (Chang, 2023). The machines are thus able to understand and interpret human language, in effect generating natural language that imitates humans. Sentiment analysis is a part of natural language processing that focuses on relating the analyzed text to its emotional overtones (Kusal et al., 2023). The sentiment of the analyzed text will be determined on a numerical range from -1 to 1, and the range boundaries between negative, neutral and positive tones of speech can be determined by the end user based on his or her own preferences. Sentiment analysis can be performed at three main levels (Sharma, Ali and Kabir, 2024; Khurana et al., 2023). The first of these is the entire document (**document-level analysis**), which is the highest and most general level, while being the most difficult to correctly assess the represented emotion. The second level involves single sentences (**sentence-level analysis**), which makes it possible to detect a variety of emotions within a single, longer statement. For example, a product review may contain both compliments and some criticism. The last level (**aspect-based sentiment analysis**) of sentiment analysis performs a thorough evaluation of individual aspects within a single complex sentence to detect different emotions. The example would be an opinion of a restaurant where the quality of the food was satisfactory, but the decor of the establishment or the service was not good enough (Bian et al., 2019).

There are currently two main approaches to sentiment analysis. The first and much simpler approach involves using dictionaries of words and phrases directly linked to specific (positive, negative, neutral) emotions. Evaluating a text in this approach is very simple, because each word has its own polarity value, unfortunately, the words linked together form a specific context, which can completely change the emotional charge of an expression, which will not be noticed by such a simple algorithm. Statements containing significant amounts of sarcasm and irony in fact have a completely different tone than the literal meaning of the words used. The second approach is machine learning. In this case, the basis is the proper training of classification models to correctly recognize emotions in statements. Popular algorithms include: recurrent neural networks, Support Vector Machines, Random Forest (Salur and Aydin, 2020). NLP provides a number of processes that are useful for text polarity analysis, but are also essential in other aspects of text analysis and processing. Among the most important NLP processes are tokenization, stemming, lemmatization, stop words removal and N-grams analysis. Tokenization involves dividing text into smaller parts, usually single words. For example, the phrase natural language processing would be divided as follows: [natural, language, processing]. Stemming and lemmatization are similar, although stemming aims to simplify words to their root form by cutting off endings, while lemmatization aims to reduce words to their root forms. An example of stemming is replacing the word “processing” with “process,” while lematization would replace the word “smarter” with “smart”. The term stop words is used to describe words that do not have a specific meaning or do not carry relevant information, but are frequently used in a language. For English such words are: “the”. “of”, ‘a’, ‘an’, so they are removed. The last process to be described as part of natural language processing is N-grams. N-grams are sequences of adjacent words, the identification of which allows to understand the underlying context. Depending on the number of words, we distinguish between bigrams (two words), trigrams (three words) and so on. In practice, N-grams consisting of more than three words are rarely encountered due to low utility (Chang, 2023).

The main consideration in enriching the recommendation algorithm with sentiment analysis of comments under YouTube videos is choosing the right approach. In the case of a possible need to analyze sentiment under a huge number of videos on the site, it is important to choose an approach that is simple, efficient, and yet the results can be interpreted, and then decisions can be made to promote or not promote a particular video. In analyzing the sentiment of comments under videos, the aim is not to identify every possible type of emotion, but to limit ourselves to the basic and most prominent ones. In view of the above criteria, a hybrid approach combining lexicon-based polarity analysis enriched with some NLP elements was proposed.

The code snippets presented below illustrate the most important steps leading to the identification of sentiment in the comments. The first step is to import the necessary data, i.e. the content of the comment section. This is possible in several ways, but the easiest is to use the YouTube API (Application Programming Interface). Necessary for this is the google-api-python-client library, specifically the build function from the discovery module. The arguments used define the name of the service whose API we want to connect to, the version of that API and the authorization key.

```
youtube_data = build("youtube", "v3", developerKey=api_key)
```

A list of unique IDs of YouTube videos is then created, from which data such as title and comments will be extracted. The ID list can be of any length, although four YouTube videos were selected for this particular application example.

```
video_ids = []
```

A for loop allows you to iterate on the contents of the video_ids object. In order to obtain the video titles, the get function from the requests library was used based on a dynamic URL for each unique ID. The data obtained in this way was converted to JSON format.

```
response = requests.get(url)

data = json.loads(response.text)
```

Importing comments is possible using the .commentThreads().list().execute() method set on the youtube_data object.

```
Comments = []

results = youtube_data.commentThreads().list(part="snippet", videoId=video_id, textFormat="plainText").execute()

comments.extend(results["items"])
```

Using a while loop, comments from all pages were retrieved.

```
while "nextPageToken" in results:
```

The imported comments inserted in the comments list are ultimately processed into a DataFrame table, which is a pandas library object. Within the DataFrame object, the data is assigned to the appropriate columns and sentiment analysis is performed by creating an instance of the TextBlob class.

```
comments_df['sentiment'] = comments_df['text'].apply(lambda x: TextBlob(x).sentiment.polarity)
```

The new column with polarity values for each comment is created. The polarity takes values ranging from -1 to 1, which indicates the nature of the comment (negative, neutral, positive). TextBlob assigns each word a polarity value based on the lexicon, and then averages the result for the entire expression (comment). This method of text analysis is easy to use, and allows quick application for new data, which is especially important for a site like YouTube due to the amount of new content each day. On the other hand, the results may be less accurate compared to more advanced models, although the amount of content tends to rule out alternatives while remaining cost-effective. All comments are annotated with the time of publication, which are converted to date format and assigned to the nearest week in order to group the data and perform statistical analysis. The following statistics are calculated: average sentiment polarity for each week, weekly total comments, average polarity value for the entire video, standard deviation, kurtosis and skewness.

```
sentiment_stats = describe(comments_df['sentiment'], ddof=0)
```

The calculation of these descriptive statistics enables a more complete interpretation of sentiment distribution. In addition to this, a tokenization of the text is performed, that is, a division into individual words, which allows the removal of stop words from the text, which are words without significant semantic meaning, but often appearing in a given language - for example: “a”, ‘an’, ‘the’, ‘of’.

```
comments_df['text_tokens'] = comments_df['text'].apply(lambda x: word_tokenize(x))

comments_df['text_tokens'] = comments_df['text_tokens'].apply(lambda x: [word for word in x if word.lower()
not in stop_words])
```

If the user wants to enrich the recommendation algorithm, any visualization of the results is unnecessary, but for the purpose of this article, comment polarity histograms were created showing the distribution of negative, positive and neutral posts. In addition, density plots, violin plots, word clouds and comment polarity plots over time and bar charts with the most frequent negative and positive words were generated using matplotlib, seaborn, WordCloud libraries.

Results and discussion

In the presented analysis four videos from YouTube were analyzed. The highest average polarity value was received by the playthrough video of The Last Of Us - with a value of 0.12; the lowest rating was given to the Witcher: Blood Origin trailer - with a value of 0.02. The results may change over time due to the creation of new comments. Polarization values for individual videos are within 0.1, which indicates that the textblob library has noteworthy problems in recognizing the tone of the ongoing discussion. If a commenter directly expresses his or her positive or negative opinion using words with a strong positive or negative tone, textblob has no problem classifying such opinion, but any discussion about the film watched is a major challenge due to the lack of words with a specific emotional tone. The lack of deep understanding of context is a huge challenge for the future when it comes to improving natural language processing. Such observations are also confirmed by other authors, who furthermore point out the problems of classifying not only sarcasm, but also cynical and mocking statements (Koyel Chakraborty et al., 2019). One potential way to deal with the ambiguity of statements is through multimodal analysis, but the combination of different types of data and their availability limits the usefulness of such an approach (Soleymani et al., 2017). There is still little research on emotion detection in short texts (Liu, Chen and Sun, 2021).

Even greater problems can be seen in the case of “Velma”, the reception of the teaser could be considered extremely negative, but the polarity rating is positive (0.07), and even higher than the also badly received Witcher: Blood Origin trailer. The reason for this was the commenting trend of extreme sarcastic comments, i.e., criticism began with the words “I love”, “You gotta love” and so on. Such a style of commenting completely confounded textblob and probably any other natural language processing library, as a broad context and understanding of sarcasm and irony is required to evaluate such remarks. It is worth noting that most comments are written in the first two or three weeks. This means that the polarity values over time for subsequent weeks fluctuate considerably and may not adequately represent the average polarity value, as several weeks after the release, individual comments may have been added that are inconsistent with the emerging consensus in the initial post-publication period. Careful analysis of frequency histograms can reveal viewers' true feelings about a given material. For positively received films, the histogram bars slowly increase/decrease in the direction to/from the mean, and even the bars for the lowest polarity values may not exist. For disliked films, the first bar (indicating the existence of negative comments) is clearly taller than the next few. This can be clearly seen by comparing Figure 1 with Figure 2.

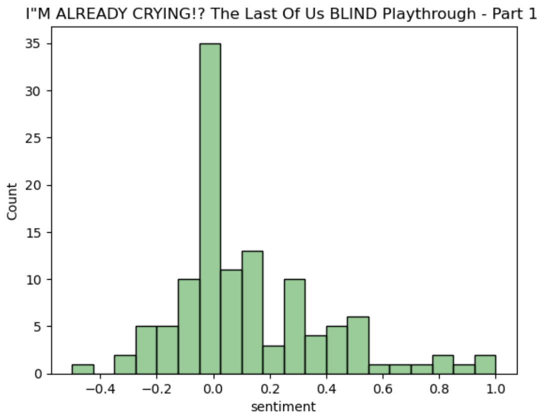

Figure 1: Sentiment Distribution for Video with the Highest Polarity Score

In Figure 1 there is a frequency histogram for the video with the highest polarity, and in Figure 2 there are two histograms for the worst rated videos.

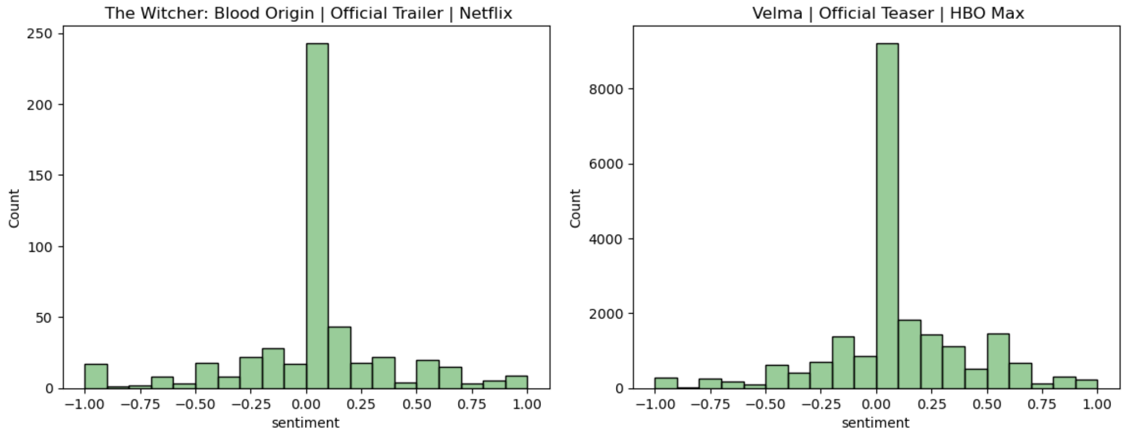

Figure 2: Sentiment Distribution for Videos with the Lowest Polarity Scores

In the case of the well-rated video, the histogram starts at -0.4, while the poorly rated ones use the scale in its entirety. It can also be seen that for well rated content there is a concentration of ratings around the mean, and for poorly rated there are more extreme opinions.

Table 2 shows data on polarity, kurtosis, skewness and the correlation between polarity and the aforementioned measures characterizing the distribution. As previously indicated, the worst-rated videos received the lowest polarization values as a result of sentiment analysis, but the differences are sometimes small, and it is difficult to assess the emotions generated by a given video on this basis without a human being overseeing such a process. On the other hand, an in-depth statistical analysis of the distribution reveals measures that, when combined with sentiment analysis, provide an unambiguous numerical answer, i.e. a recommendation algorithm could operate fully autonomously based on such data. As can be seen, the skewness for poorly rated videos takes on negative values, significantly different from positively and neutrally rated content. For example, the first and last videos differ in the polarity rating by only 0.01, but in the case of the skewness, this difference is much greater, at 0.59.

The (positive) skewness on the right indicates the presence of a significant number of observations to the right of the mean, i.e. extremely positive comments. The inverse relationship is found for leftward (negative) asymmetry. The correlation between polarity and skewness is very strong at 0.9. Higher kurtosis for poorly rated films confirms the presence of more extreme opinions. Kurtosis is strongly negatively correlated with polarity (-0.7). A right (positive) skew indicates the presence of a significant number of observations to the right of the mean, i.e., extremely positive comments. The opposite relationship occurs in the case of left (negative) skewness. The correlation between polarity and skewness is very strong at 0.9. Higher kurtosis for poorly rated films confirms the presence of more extreme opinions. Kurtosis is strongly negatively correlated with polarity (-0.7).

Table 2: Correlation Analysis of Polarity, Kurtosis, Skewness for Selected YouTube Videos

No.	Title	ID	Polarity	Kurtosis	Skewness
1	GTA 5 MIRROR WORLD Chaos Mod Speedrun! -Viewers Randomly Mod The Game In A Reversed Los Santos! #1	OififS2GqNRg	0,08	1,34	0,42
2	I"M ALREADY CRYING!? The Last Of Us BLIND Playthrough - Part 1	7N0MGmvYXUw	0,12	1,32	1,05
3	The Witcher: Blood Origin | Official Trailer | Netflix	py3i-BGUDGw	0,02	2,08	-0,27
4	Velma | Official Teaser | HBO Max	lHdtsWn7sgE	0,07	1,67	-0,17
	Correlation between polarization and kurtosis	Correlation between polarization and kurtosis			
	-0,692265024	0,898792281			

Figures 3 and 4 show the changes in polarity over time and the corresponding changes in the number of new comments over the same period for the gameplay snippet from The Last of Us and the Velma series teaser. The graphs show that the majority of comments are published in the first two weeks, and then activity in the comment section almost disappears. This means that the polarity cannot be calculated from week to week anew, as it is subject to large fluctuations.

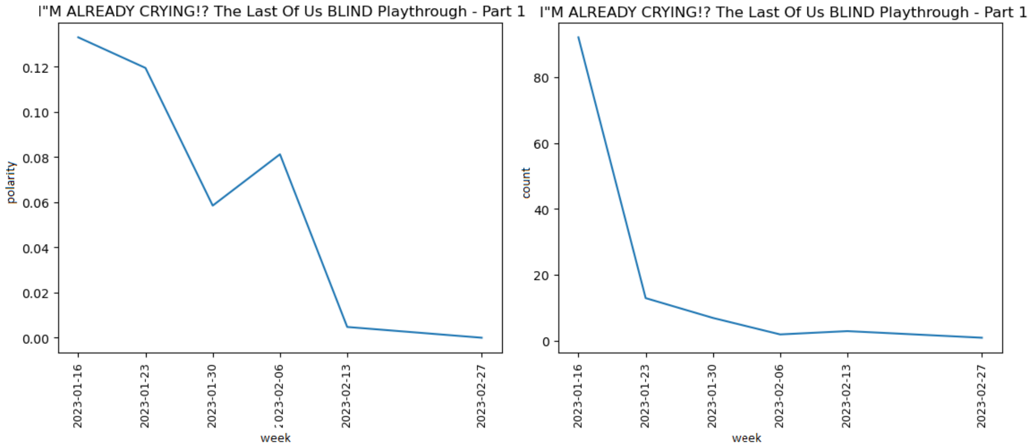

Figure 3: Polarity Over Time and New Comments for Video with Polarity Score = 0.12

The example of the Velma series teaser shows how much trouble textblob has in recognizing sarcasm and irony. In the comments there was a trend of ironic praise of the disliked production, which elevated the polarization to 0.15, after this trend stopped, the new comments are directly critical and the polarization drops below -0.1.

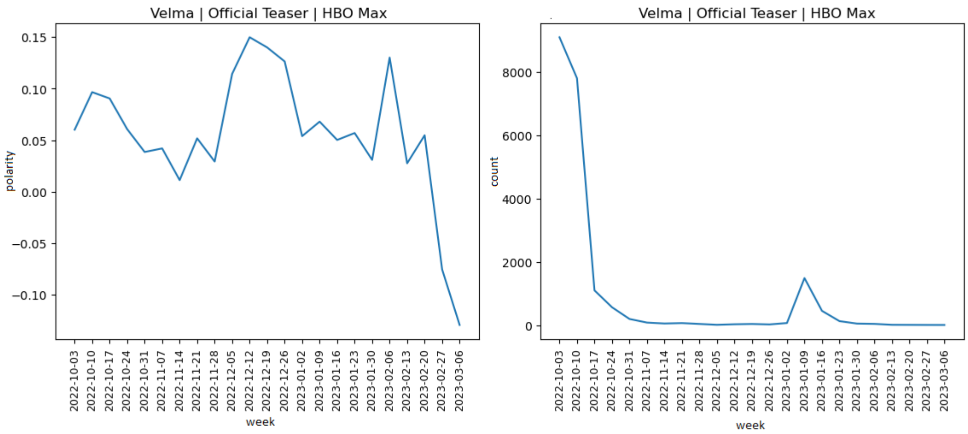

Figure 4: Polarity Over Time and New Comments for Video with Polarity Score = 0.07

This article explores the possibility of enhancing YouTube's recommendation algorithm with sentiment analysis of comments based on simple tools that do not require significant computing power. With the combination of sentiment analysis and statistical distribution analysis, it seems possible to improve the current algorithm with a cost-effective solution that will help generate more ad revenue despite the platform's limited ability to grow based on acquiring new users.

References

- Alphabet Inc. (2017-2023) *Annual Reports on Form 10-K for the Fiscal Years 2017-2023: https://abc.xyz/investor (Accessed: 20 September 2024)* [Online].
- Alphabet Inc. (2023) *Form 10-K for the Fiscal Year Ended December 31, 2023. Washington, D.C.: United States Securities and Exchange Commission. https://abc.xyz/assets/43/44/675b83d7455885c4615d848d52a4/goog-10-k-2023.pdf (Accessed: 20 September 2024).* [Online].

- Babbar, G., Anand, A. and Aggrawal, N. (2024) 'Modelling & Analyzing View Growth Pattern of YouTube Videos inculcating the impact of Subscribers, Word of Mouth and Recommendation Systems', *International Journal of Mathematical, Engineering and Management Sciences*, vol. 9, no. 3, pp. 435–450.
- Batra, R. and Stayman, D. M. (1990) 'The Role of Mood in Advertising Effectiveness', *Journal of Consumer Research*, vol. 17, no. 2, pp. 203–214 [Online]. Available at http://www.jstor.org/stable/2626812 (Accessed 20 September 2024).
- Bian, W., Wang, C., Ye, Z. and Yan, L. (2019) 'Emotional Text Analysis Based on Ensemble Learning of Three Different Classification Algorithms', *2019 10th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)*, pp. 938–941.
- bizcommunity.com (Various sources) (2024) *Leading social media platforms in the United States in 2023, by ad revenue (in billion U.S. dollars). Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/1103339/social-media-ad-revenue-platform/* [Online].
- Cartwright, J., McCormick, H. and Warnaby, G. (2016) 'Consumers' emotional responses to the Christmas TV advertising of four retail brands', *Journal of Retailing and Consumer Services*, vol. 29, No. 4, pp. 82–91.
- Chang, K.-H. (2023) 'Natural Language Processing: Recent Development and Applications', *Applied Sciences*, vol. 13, no. 20, p. 11395.
- DataReportal, We Are Social, Meltwater (2024) *Leading countries based on YouTube audience size as of July 2024 (in millions). Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/280685/number-of-monthly-unique-youtube-users/* [Online].
- Davidson, J., Liebald, B., Liu, J., Nandy, P., van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B. and Sampath, D. (2010) 'The YouTube video recommendation system', *Proceedings of the fourth ACM conference on Recommender systems.* Barcelona Spain, 26 09 2010 30 09 2010. New York, NY, USA, ACM, pp. 293–296.
- Eksombatchai, C., Jindal, P., Liu, J. Z., Liu, Y., Sharma, R., Sugnet, C., Ulrich, M. and Leskovec, J. (2018) 'Pixie: A System for Recommending 3+ Billion Items to 200+ Million Users in Real-Time', in *Proceedings of the 2018 World Wide Web Conference,* Lyon, France, International World Wide Web Conferences Steering Committee, pp. 1775–1784.
- Holbrook, M. B. and O'Shaughnessy, J. (1984) 'The role of emotion in advertising', *Psychology & Marketing*, vol. 1, no. 2, pp. 45–64.
- Insider Intelligence (2023) *Average time spent per day on Netflix, TikTok, and YouTube by adults in the United States from 2020 to 2024 (in minutes). Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/1359403/us-time-spent-per-day-netflix-tiktok-youtube/* [Online].
- Khurana, D., Koli, A., Khatter, K. and Singh, S. (2023) 'Natural language processing: State of the art, current trends and challenges', *Multimedia tools and applications*, vol. 82, no. 3, pp. 3713–3744.
- Koyel Chakraborty, Siddhartha Bhattacharyya, Rajib Bag and Aboul Alla Hassanien (2019) '7 - Sentiment Analysis on a Set of Movie Reviews Using Deep Learning Techniques', in Nilanjan Dey, Samarjeet Borah, Rosalina Babo and Amira S. Ashour (eds) *Social Network Analytics,* Academic Press, pp. 127–147.
- Kusal, S., Patil, S., Choudrie, J., Kotecha, K., Vora, D. and Pappas, I. (2023) 'A systematic review of applications of natural language processing and future challenges with special emphasis in text-based emotion detection', *Artificial Intelligence Review*, vol. 56, no. 12, pp. 15129–15215.
- Liu, L., Chen, H. and Sun, Y. (2021) 'A Multi-Classification Sentiment Analysis Model of Chinese Short Text Based on Gated Linear Units and Attention Mechanism', *ACM Transactions on Asian and Low-Resource Language Information Processing*, vol. 20, no. 6, pp. 1–13.
- Salur, M. U. and Aydin, I. (2020) 'A Novel Hybrid Deep Learning Model for Sentiment Classification', *IEEE Access*, vol. 8, pp. 58080–58093.
- Sharma, N. A., Ali, A. B. M. S. and Kabir, M. A. (2024) 'A review of sentiment analysis: Tasks, applications, and deep learning techniques', *International Journal of Data Science and Analytics*, vol. 8, no. 3, p. 58080.
- Soleymani, M., Garcia, D., Jou, B., Schuller, B., Chang, S.-F. and Pantic, M. (2017) 'A survey of multimodal sentiment analysis', *Image and Vision Computing*, vol. 65, no. 11, pp. 3–14.

- Statista (2024) *YouTube's advertising revenues as percentage of Google's global revenues from 2017 to 2023. Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/289659/youtube-share-of-google-total-ad-revenues/* [Online].
- Vrtana, D. and Krizanova, A. (2023) 'The Power of Emotional Advertising Appeals: Examining Their Influence on Consumer Purchasing Behavior and Brand–Customer Relationship', *Sustainability*, vol. 15, no. 18, p. 13337.
- We Are Social, DataReportal, Meltwater (2024) *Most popular social networks worldwide as of April 2024, by number of monthly active users (in millions). Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/* [Online].
- Weinberger, M. G. and Gulas, C. S. (1992) 'The Impact of Humor in Advertising: A Review', *Journal of Advertising*, vol. 21, no. 4, pp. 35–59.
- YouTube (2024) *Number of paying YouTube Music and YouTube Premium subscribers worldwide from 2020 to 2024 (in millions). Statista. Statista Inc. Accessed: September 20, 2024. https://www.statista.com/statistics/1344265/youtube-paying-subscribers/* [Online].
- Zhao, Z., Hong, L., Wei, L., Chen, J., Nath, A., Andrews, S., Kumthekar, A., Sathiamoorthy, M., Yi, X. and Chi, E. (2019) 'Recommending what video to watch next', *Proceedings of the 13th ACM Conference on Recommender Systems.* Copenhagen Denmark, New York, NY, USA, ACM, pp. 43–51.