

The Effect of Hiding Dislikes on the Use of YouTube’s Like and Dislike Features

Soon-Gyo Jung
Qatar Computer Research Institute,
Hamad Bin Khalifa University, Doha,
Qatar
sjung@hbku.edu.qa

Joni Salminen
University of Vaasa, Vaasa, Finland
joni.salminen@uwasa.fi

Bernard J. Jansen
Qatar Computer Research Institute,
Hamad Bin Khalifa University, Doha,
Qatar
jjansen@acm.org

ABSTRACT

Using data from a major international news organization, we investigate the effect of hiding the count of dislikes from YouTube viewers on the propensity to use the video like/dislike features. We compare one entire month of videos before ($n = 478$) and after ($n = 394$) YouTube began hiding the dislikes counts. Collectively, these videos had received 450,200 likes and 41,892 dislikes. To account for content variability, we analyze the likes/dislikes by sentiment class (positive, neutral, negative). Results of chi-square testing show that while both likes and dislikes decreased after the hiding, dislikes decreased substantially more. We repeat the analysis with four other YouTube news channels in various languages (Arabic, English, French, Spanish) and one non-news organization, with similar results in all but one case. Findings from these multiple organizations suggest that YouTube hiding the number of dislikes from viewers has altered the user-platform interactions for the like/dislike features. Therefore, comparing the like/dislike metrics before and after the removal would give invalid insights into users’ reactions to content on YouTube.

CCS CONCEPTS

• **Human-centered computing** → Collaborative and social computing.

KEYWORDS

User metrics, User interaction, Web analytics, Social Signaling

ACM Reference Format:

Soon-Gyo Jung, Joni Salminen, and Bernard J. Jansen. 2022. The Effect of Hiding Dislikes on the Use of YouTube’s Like and Dislike Features. In *14th ACM Web Science Conference 2022 (WebSci ’22)*, June 26–29, 2022, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3501247.3531546>

1 INTRODUCTION

On 10 November 2021, in a major change, YouTube retained the video dislike feature but hid the count of video dislikes from viewers,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci ’22, June 26–29, 2022, Barcelona, Spain

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9191-7/22/06...\$15.00

<https://doi.org/10.1145/3501247.3531546>

deviating from a long-standing position of showing this metric in the user interface (UI) [14], as illustrated in Figure 1

The stated rationale for removing the dislike count was that the dislikes were sometimes a reflection of the user opinion of the entire channel rather than the video content itself [10], although no data supporting this was presented [19]. YouTube communications further stated, via internal testing, that removing the dislike count did not result in meaningful viewership differences, but it did reduce dislike attacks [10]. It was also noted that the dislike button remains on the YouTube site, and dislikes will continue to be factored into YouTube’s recommendation algorithm [10], impacting the videos that are suggested to users [4].

The move was met with substantial criticism, including from the co-founder of YouTube, Jawed Karim [11]. Likes, dislikes, and related social signals [2] are employed by users across many online platforms to evaluate the quality of content. Even though several platforms have experimented with options of hiding or removing such signals [20], these signals can assist users in both locating valuable and avoiding unhelpful or harmful content [18] and thus represent valued engagement metrics [1]. However, there are questions concerning the worth and/or harm of using social expressions of approval/disapproval, e.g., in the form of cyberattacks [5, 9, 16], and there appear to be age and gender differences related to the use of these features [8, 15]. Due to potential benefits and risks, the long-term effects of hiding the dislike count remain to be seen and require research in several areas, such as online advertising [17].

As a starting point for this line of research, in this study, we examine the effects of hiding the dislike count from viewers on user interactions with the YouTube like and dislike features using actual data from the channel of a major news organization.

2 RESEARCH QUESTIONS AND HYPOTHESES

Guiding our study, our research question is: How does hiding the number of dislikes from viewers affect user interaction with the YouTube like/dislike features?

To investigate this research question, we analyze the number of likes and dislikes of videos from the YouTube channel of a major international news organization. The metrics utilized for this research are the number of likes and dislikes, reserving other metrics for future analysis. We conduct sentiment analysis using a state-of-art sentiment classifier of the video titles to account for possible changes in news content, classifying each as positive, natural, or negative. Therefore, our hypotheses (H) are:

- **H1:** Hiding the number of dislikes for positive sentiment videos changes user interaction with the like/dislike features.

YouTube Like and Dislike Features



Figure 1: Snapshots of YouTube Like and Dislike Feature Before and After the Change of Hiding the Dislike Count.

- **H2:** Hiding the number of dislikes for neutral sentiment videos changes user interaction with the like/dislike features.
- **H3:** Hiding the number of dislikes for negative sentiment videos changes user interaction with the like/dislike features.

We conduct sentiment analysis to control for topical variation of the online content, especially in the news domain, which might change over time. Because there could be a difference in the sentiment of the content in the *pre-* and *post-periods*, users might, therefore, react to content in a fundamentally different way, affecting their likes or dislikes of the content. By grouping the content in sentiment classes, we can examine the effects of the change of not displaying the number of dislikes while controlling for this potential shift in news sentiment. Although there may be other factors at play as well, content sentiment is a valuable control factor given that the use of the like/dislike feature is our research focus.

3 METHODOLOGY

3.1 Data Collection and Preparation

YouTube announced the removal of the number of dislikes on 10 November 2021 [14], so we refer to this date as the *date of the change/removal*. The period before this date is the *pre-period*, and the period after this is the *post-period*. We first collected daily data from 1 September 2021 through 31 January 2022 for all videos posted to the channel using the YouTube Reporting API [21].

The change was implemented mid-month, but we desired an entire month for data collection for both the *pre-* and *post-period*. As November 2021 was a transitional month for implementing the change globally [11]—i.e., changes in the YouTube UI are not necessarily rolled out to all users simultaneously—we determined it best not to include any data from this month, which eliminated videos published in October as the like/dislike period would extend into this transitional period. So, our pre-period sample was all videos posted on the channel in September (*pre-period*) and all videos posted in December 2021 (*post-period*), which resulted in balanced data collection for both the pre- and post-periods.

We then determined the window for collecting the number of likes and dislikes of each video. We analyzed more than 25,000 videos published over five years on this channel. On average, the videos gained 90% of views within 17 days of the publication date, so we use 17 days for the likes/dislikes data collection, which normalizes the counts by the publication date. For example, if a video was published on 1 September, the period for collection of the likes and dislikes is September 1–17. If a video was posted on 31 December, the period for collection of the likes and dislikes is 31 December 2021 – 16 January 2022. This approach ensures that each video has an equitable data collection period.

Therefore, for the *pre-period*, we used all videos posted to this channel during September 2021, with data collection for likes and dislikes for these videos is the period 1 September through 16 October 2021, inclusive. For the *post-period*, we used all videos posted to this channel during December 2021, with data collection for likes and dislikes of these videos being the period of 1 December 2021 through 16 January 2022, inclusive.

3.2 Normalization

As some videos might go viral while others are not popular at all, using the raw number of likes or dislikes would not account for outliers within the dataset. Therefore, we normalized the number of likes and dislikes by the number of views (e.g., normalized likes = # likes / # views; normalized dislikes = # dislikes / # views). For example, assume Video A received 3000 views, 2000 likes, and 1000 dislikes. The normalized likes would be 0.67 (e.g., 2000/3000), and the normalized dislikes would be 0.33 (e.g., 1000/3000). This normalization of likes and dislikes accounts for the unbalanced popularity of videos and represents a typical method for dealing with outliers when analyzing social media data [1]. We thus use the normalized likes and dislikes for our statistical analysis.

3.3 Sentiment Analysis

The videos between the *pre-* and *post-periods* might have different content, perhaps evoking sentiments that affect how users like or dislike the video. To account for this possible content difference, we classified each video using the video title employing the *IBM Watson Natural Language Understanding* API (version is 2021-08-01) [22]. For our research, the employment of the title for sentiment analysis was a judicious approach, as it has been shown that the title of the video provides key information for viewers to watch/not watch a video [3, 13]. In terms of sentiment analysis, IBM Watson is reliable relative to other services [7], freely available, and extensively employed. As such, it is an excellent choice for this research. IBM Watson determines and returns sentiment labels (e.g., positive, neutral, negative). We use the sentiment label in our research, employing the 872 video titles as input, reserving investigation of other sentiment services and approaches for other research.

3.4 Statistical Analysis Approach

For statistical analysis, we have two periods (i.e., *pre* and *post*) with two measures (i.e., number of likes and dislikes) for each of the three sentiments (i.e., positive, neutral, negative). This results in three (i.e., one for each sentiment) two-by-two contingency tables (i.e., Before and After × Likes and Dislikes) for addressing our hypotheses. We conducted three chi-square tests of independence

Table 1: Comparison of Three Sentiment of Videos Before (B) and After (A) the Change. Percentage Difference (% Diff) shows the Increase (+) or Decrease (-) in the Number of Sentiment Videos by Category in Each Period.

	Before (B)	After (A)	Total	% Diff
Positive	21	17	38	-0.1%
Neutral	324	272	596	1.3%
Negative	133	105	238	-1.2%
Total	478	394	872	0.0%

Table 2: Comparison of Likes and Dislikes Before (B) and After (A) the Change. Percentage Difference (% Diff) shows the Increase (+) or Decrease (-) in the Number of Likes (or Dislikes) by Category in Each Period. % Diff = ((A/#V)/(B/#V))-Total

Likes	Before (B)	After (A)	Total	% Diff
Positive	9,002	10,094	19,096	0.86%
Neutral	191,786	117,299	309,085	-0.05%
Negative	83,280	38,481	121,761	-0.21%
Total	284,068	165,874	449,942	-0.04%
Dislikes	Before (B)	After (A)	Total	% Diff
Positive	885	513	1,398	-0.86%
Neutral	19,215	9,286	28,501	-0.09%
Negative	8,599	3,379	11,978	-0.27%
Total	28,699	13,178	41,877	-0.06%

per YouTube channel, one test for each contingency table. The chi-square test for independence examines an association between two categorical variables within the same population by comparing two variables within a sample set to one another, informing whether given combinations occur more frequently than one would expect by chance. Given that the variables are categorical for the corresponding contingency tables, the chi-square test is appropriate for this research.

4 RESULTS

4.1 Exploratory Analysis

We first compare the frequencies of positive, neutral, and negative sentiment videos published in the pre- and post-periods. As shown in Table 1, there are slight differences between content sentiment in the two periods, with the number of videos for each sentiment relatively the same before and after the change, indicating that the content did not dramatically change. So, our approach of examining by sentiment is perhaps overly cautious but still warranted, as there was some change.

We then compare the likes and dislikes between the pre- and post-periods. As shown in Table 2, there are differences in the ratios of likes and dislikes in the two periods when normalized by the number of videos in that period for that sentiment. % Diff is standardized

in each period by dividing by the number of videos (#V) in that period and then by the total likes/dislikes for that period. There was an increase in likes and a corresponding decrease in dislikes for positive videos. However, there was a decrease in likes and dislikes for neutral and negative videos, with the decrease in dislikes being higher in both cases.

We then generate boxplots of our normalized data, as boxplots provide an easily digestible graph of the spread of values by displaying the information in quartiles. The boxplots also display the median, the interquartile range, minimum, maximum, and outliers. As shown in Figure 2 (next page), the set of boxplots (displayed as *After* and then *Before* the change) indicates that likes and dislikes generally decreased after hiding the number of dislikes from users. The results of this exploratory analysis indicate that further statistical analysis is warranted.

4.2 Hypothesis Testing

We now test our hypotheses. Table 3 shows three contingency tables for positive, neutral, and negative videos and the use of the like/dislike YouTube feature, with values indicating the summations of the normalized likes and dislikes as discussed.

A chi-square test of independence was performed to examine the relation between pre- and post-period and the use of the like/dislike features for videos of *positive sentiment*. The chi-square test showed a significant association between periods and the use of the like/dislike features, $X^2(1, N = 20,494) = 135.66, p < .001$. Likes increased and dislikes decreased in the post-period. So, **H1 is fully supported**. For the number of likes in the post-period, the percentage increase is substantially smaller (12.1%) compared to the percentage decrease in the number of dislikes (-42.0%).

A chi-square test was performed to examine the relation between pre- and post-period and the use of the like/dislike features for videos of *neutral sentiment*. The chi-square test showed a significant association between periods and the use of like/dislike features, $X^2(1, N = 337,586) = 320.74, p < .001$. So, **H2 is fully supported**. Relative to the pre-period, both likes and dislikes decreased in the post-period. Dislikes decreased more (-51.7%) relative to the likes (-38.8%).

A chi-square test was performed to examine the relation between pre- and post-period and the use of the like/dislike features for videos of *negative sentiment*. The chi-square test showed a significant relation between periods and the use of like/dislike features, $X^2(1, N = 337,586) = 58.25, p < .001$.

So, **H3 is fully supported**. Relative to the pre-period, both likes and dislikes decreased in the post-period. Dislikes decreased more (-60.7%) than likes (-53.8%).

4.3 Repeated Analysis

We carried out the above analysis using the YouTube channel of a large online news organization. To evaluate the robustness of our findings, we repeated the analysis using the same procedure outlined above for the YouTube channels of four other news organizations and one non-news organization (a commercial service company). Each of the four news organizations publishes in a different language (Arabic, English, French, and Spanish), and the non-news organization publishes in English. The rate of videos

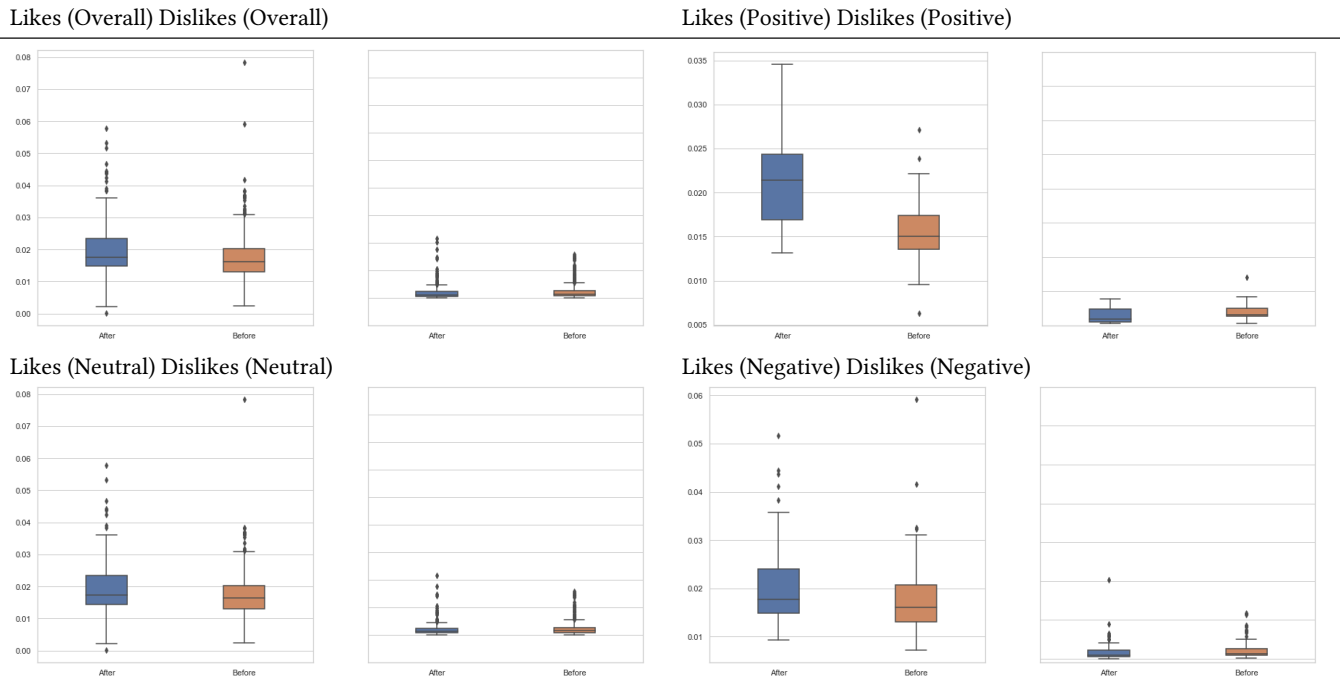


Figure 2: Boxplots of Like and Dislikes (After and Before). Boxplots for All Videos, Positive Videos Only, Neutral Videos Only, and Negative Videos Only. The Boxplots Show That in Three Categories, The Number of Likes and the Dislikes Generally Decreased After the Total Number of Dislikes Was No Longer Displayed on YouTube. There was an Increase in Likes for Positive Videos.

Table 3: Two-by-Two Contingency Tables for Positive, Neutral, and Negative Sentiment YouTube Videos and the Relationship Between Pre- and Post-Period on the Use of the Like/Dislike Features.

	Positive			Neutral			Negative				
	Like	Dislike	Total	Like	Dislike	Total	Like	Dislike	Total		
Pre	0.335	0.044	0.379	Pre	5.682	0.791	6.473	Pre	2.298	0.288	2.586
Post	0.360	0.024	0.379	Post	5.275	0.542	6.473	Post	2.105	0.186	2.586
Total	0.695	0.068	0.758	Total	10.957	1.333	12.946	Total	4.403	0.474	5.172

published varied among the organizations, so there were no videos published in the data sampling periods for a given sentiment for some organizations. In these cases, we were unable to conduct the analysis.

In the interests of space, we report the results in Table 4, which shows that tests are significant for all but one combination (News Org (Arabic) for Neutral Sentiment Videos), and results were highly similar to our primary analysis, including for likes for positive sentiment videos. In the chi-square test in Table 4, all degrees of freedom are one ($df=1$). We replicated all the tests using Fisher’s exact test to account for small sample sizes [12], and the significance of the results did not change.

5 DISCUSSION AND IMPLICATIONS

We evaluate the effect of hiding the number of dislikes for YouTube videos on users’ use of the like/dislike features. We use the like

and dislike counts from the YouTube channel of a major news organization, using videos posted one full month before the change and one month after the changes, allowing each video 17 days of user interaction data. We conduct sentiment analysis on the video titles, classifying all videos as positive, neutral, or negative to control for possible changes in the news content between the two periods, surmising that this might affect users’ use of the like or dislike features. We then conduct chi-square tests to evaluate changes in the like/dislike counts between the two periods for the three sentiment groups of videos.

All hypotheses were fully supported. The dislikes decreased and, in most cases, so did likes, but not as drastically as the dislikes. Removing the dislike count seems to increase the number of likes for positive sentiment videos. The statistical results hold for (nearly) all channels, all languages, and all publishing volumes, i.e., the results were the same if the channel published many videos or

Table 4: Chi-square Test Results Comparing Frequency of Likes and Dislikes. * is Significant at the $p < .01$. The Shaded Row Indicates a Non-Significant Finding.

Organization (Org.)	Sentiment	χ^2	p	N
News Org. 1 (Arabic)	Positive	14.09 *	0.0002	13,687
	Neutral	0.004	0.9450	25,074
News Org. 2 (English)	Positive	10.07 *	0.0015	1,058
	Neutral	337.37 *	0.0000	17,750
	Negative	14.51 *	0.0001	13,198
News Org. 3 (French)	Neutral	160.39 *	0.0000	66,749
News Org. 4 (Spanish)	Neutral	15.10 *	0.0001	838
Non-News Org. (English)	Positive	20.04 *	0.0000	1,089
	Neutral	6.67 *	0.0098	5,732

only published a few videos. There was one channel (News Org 1 Arabic Neutral) where the results were not significant. Why this one channel was not significant requires further investigation, as there is no notable difference in the channel videos or metrics relative to the other organizations from a review of the channel's video content or metrics.

The main implication is that for the like/dislike feature, analysis of user behavior on a given channel cannot be compared before and after the change of hiding the count of dislikes. This implication raises several exciting avenues for further investigation into whether this change has affected other user behaviors, such as subscription rates or video comments, and it also opens research avenues into the implications of other social media platforms' like/dislike feature changes. As there are unintended consequences of providing like/dislike features [5, 6], it appears that there are also unintended consequences of, in this case, hiding the number of dislikes, as it appears to have affected user interaction with the video content. Interesting future research in this area includes user studies exploring *why* removing the dislike counts from viewers has altered users' interactions behaviors.

6 CONCLUSION

Investigating the effect of hiding the dislike count from viewers on their propensity to use the video like/dislike feature, we find statistically significant changes in the use of the like/dislike feature after the removal. Generally, other than likes for positive videos, both likes and dislikes decreased after removal, but dislikes decreased much more than likes decreased, showing that hiding dislikes from the user interface seems to have altered user behavior with the online content on YouTube.

REFERENCES

- [1] Kholoud Khalil Aldous, Jisun An, and Bernard J. Jansen. 2019. View, Like, Comment, Post: Analyzing User Engagement by Topic at 4 Levels across 5 Social Media Platforms for 53 News Organizations. In *Proceedings of the International AAAI Conference on Web and Social Media*, 47–57. Retrieved February 3, 2021 from <https://ojs.aaai.org/index.php/ICWSM/article/view/3208>
- [2] Kholoud Khalil Aldous, Jisun An, and Bernard J. Jansen. 2022. Measuring 9 Emotions of News Posts from 8 News Organizations across 4 Social Media Platforms for 8 Months. *ACM Transactions on Social Computing* 4, 4 (2022), 1–31. DOI:<https://doi.org/10.1145/3516491>
- [3] Balakumaran. 2019. YouTube Video Title that Gets Views: 6 Best Ideas in 2022. *Stranger Show*. Retrieved January 28, 2022 from [youtube-video-title/](https://strangershow.com/youtube-video-title/)
- [4] Nathan Bartley, Andres Abeliuk, Emilio Ferrara, and Kristina Lerman. 2021. Auditing Algorithmic Bias on Twitter. In *13th ACM Web Science Conference 2021 (WebSci '21)*, Association for Computing Machinery, New York, NY, USA, 65–73. DOI:<https://doi.org/10.1145/3447535.3462491>
- [5] Emiliano De Cristofaro, Arik Friedman, Guillaume Jourjon, Mohamed Ali Kaafar, and M. Zubair Shafiq. 2014. Paying for Likes? Understanding Facebook Like Fraud Using Honey Pots. In *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14)*, Association for Computing Machinery, New York, NY, USA, 129–136. DOI:<https://doi.org/10.1145/2663716.2663729>
- [6] Hridoy Sankar Dutta, Mayank Jobanputra, Himani Negi, and Tanmoy Chakraborty. 2021. Detecting and Analyzing Collusive Entities on YouTube. *ACM Transactions on Intelligent Systems and Technology* 12, 5 (November 2021), 64:1–64:28.
- [7] Soon-gyo Jung, Joni Salminen, and Bernard J. Jansen. 2022. Engineers, Beware! Commercial Tools Disagree on Social Media Sentiment. In *Proceedings of the 14th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS'22)*, Sophia Antipolis, France.
- [8] M. Laeeq Khan. 2017. Social media engagement: What motivates user participation and consumption on YouTube? *Computers in Human Behavior* 66, (January 2017), 236–247. DOI:<https://doi.org/10.1016/j.chb.2016.09.024>
- [9] Sarah Lutz and Frank M. Schneider. 2021. Is receiving Dislikes in social media still better than being ignored? The effects of ostracism and rejection on need threat and coping responses online. *Media Psychology* 24, 6 (2021), 741–765. DOI:<https://doi.org/10.1080/15213269.2020.1799409>
- [10] 172 Shares and 935 Reads. 2022. YouTube CEO Defends Removal of Dislike Counts. *Search Engine Journal*. Retrieved January 28, 2022 from <https://www.searchenginejournal.com/youtube-ceo-defends-removal-of-dislike-counts/435092/>
- [11] 214 Shares and 6 9k Reads. 2021. YouTube Co-Founder Is Against Removing Dislikes. *Search Engine Journal*. Retrieved January 28, 2022 from <https://www.searchenginejournal.com/youtube-co-founder-is-against-removing-dislikes/427590/>
- [12] Peter Sprent. 2011. Fisher Exact Test. In *International Encyclopedia of Statistical Science*, Miodrag Lovric (ed.). Springer, Berlin, Heidelberg, 524–525. DOI:https://doi.org/10.1007/978-3-642-04898-2_253
- [13] Wondwesen Tafesse. 2020. YouTube marketing: how marketers' video optimization practices influence video views. *Internet Research* 30, 6 (January 2020), 1689–1707. DOI:<https://doi.org/10.1108/INTR-10-2019-0406>
- [14] The YouTube Team. An update to dislikes on YouTube. *blog.youtube*. Retrieved January 28, 2022 from <https://blog.youtube/news-and-events/update-to-youtube/>
- [15] Marika Tiggemann, Susannah Hayden, Zoe Brown, and Jolanda Veldhuis. 2018. The effect of Instagram "likes" on women's social comparison and body dissatisfaction. *Body Image* 26, (September 2018), 90–97. DOI:<https://doi.org/10.1016/j.bodyim.2018.07.002>
- [16] Elaine Wallace and Isabel Buil. 2021. Hiding Instagram Likes: Effects on negative affect and loneliness. *Personality and Individual Differences* 170, (February 2021), 110509. DOI:<https://doi.org/10.1016/j.paid.2020.110509>
- [17] Yanwu Yang and Panyu Zhai. 2022. Click-through rate prediction in online advertising: A literature review. *Information Processing & Management* 59, 2 (March 2022), 102853. DOI:<https://doi.org/10.1016/j.ipm.2021.102853>
- [18] 2020. Do Dislikes Matter on YouTube? [More Than You Might Think!]. *Alan Spicer - YouTube Certified Expert*. Retrieved January 29, 2022 from <https://alanspicer.com/do-dislikes-matter-on-youtube/>

- [19] YouTube is removing the dislike count on all videos across its platform. *TechCrunch*. Retrieved January 29, 2022 from <https://social.techcrunch.com/2021/11/10/youtube-is-removing-the-dislike-count-on-all-videos-across-its-platform/>
- [20] YouTube tests hiding dislike counts on videos. *TechCrunch*. Retrieved January 29, 2022 from <https://social.techcrunch.com/2021/03/30/youtube-tests-hiding-dislike-counts-on-videos/>
- [21] Channel Reports | YouTube Analytics and Reporting APIs. *Google Developers*. Retrieved January 28, 2022 from https://developers.google.com/youtube/reporting/v1/reports/channel_reports
- [22] IBM Watson Natural Language Understanding | IBM. Retrieved January 28, 2022 from <https://www.ibm.com/cloud/watson-natural-language-understanding>