

# Product Spam on YouTube: A Case Study

Janek Bevendorff  
Leipzig University  
Germany

Martin Potthast  
Leipzig University and ScaDS.AI  
Germany

Matti Wiegmann  
Bauhaus-Universität Weimar  
Germany

Benno Stein  
Bauhaus-Universität Weimar  
Germany

## ABSTRACT

YouTube videos are a popular medium for online product reviews. They are not only informative and entertaining, but may also be perceived as quite credible under the viewer’s impression of a personal product demonstration by an expert. As the world’s largest online video platform, YouTube’s content is included prominently in the results of most general-purpose web search engines. Consequently, online marketers are using classic Search Engine Optimization (SEO) techniques also for placing their video content in search engines. Over the years, we have noticed an ever increasing noise floor of low-quality SEO content in product search results and in this study, we show that this trend has spilled over into videos as well. We examine YouTube video reviews for several thousand products retrieved from three commercial search engines and conduct spam detection experiments based directly on the videos’ subtitle transcripts rather than relying on metadata and comments. We find that at least a third of the retrieved videos can be regarded as spam or low-quality productions. We are further able to distinguish these spam product reviews accurately from higher-quality videos with a semi-supervised n-gram classification approach.

## CCS CONCEPTS

• Information systems → Content match advertising; Spam detection.

## KEYWORDS

Web Spam, Content Quality, SEO, YouTube

### ACM Reference Format:

Janek Bevendorff, Matti Wiegmann, Martin Potthast, and Benno Stein. 2024. Product Spam on YouTube: A Case Study. In *Proceedings of the 2024 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR ’24)*, March 10–14, 2024, Sheffield, United Kingdom. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3627508.3638303>

## 1 INTRODUCTION

In an effort to improve their visibility on the web, most commercial websites today use some form of search engine optimization (SEO). A sensible amount of SEO, if applied “with good intentions,” may

actually improve both the on-site user experience and the effectiveness of the search engine itself. A properly optimized website is more accessible and more efficient to parse, making it easier for the search engine to identify relevant information. On the other hand, SEO is often used with malicious intent to gain undeserved visibility—just pretending to be of value to the user by gaming the search engine’s algorithm. This kind of “blackhat” SEO spam has taken over many competitive, low-margin environments, of which a prominent example is product search. In this market, many sellers compete for attention to sell their products online, relying heavily (like most other websites) on SEO to gain new customers. While it has been shown that strong SEO may negatively affect users’ perception of a website’s expertise [18], a high rank is often more important. Users trust their search engines a fair amount [11, 15], thus they may lend extra credibility to well-ranked pages.

In addition to search engine optimization (SEO) and marketing (SEM), larger sellers have started offering affiliate programs, which have since become another massive market for which SEO and SEM play just as much of an important role. In online affiliate marketing, the participant (the affiliate) directs traffic from their own website via special product links to a seller (the affiliate partner), earning a commission for each successful referral. The entry barrier for online affiliate marketing is very low, for all that is needed is a website and an agreement with an affiliate partner or partner network. As a result, affiliate marketing has become a popular source of income particularly for bloggers, social media influencers, and product review portals. The low barrier of entry into affiliate marketing has also lured spammers into the market who try to place a bulk of low-effort or even fake product reviews in search engines to harvest affiliate clicks—a trend that is bound to accelerate with the rise of generative AI. Today, SEO-driven affiliate spam is already too pervasive and ubiquitous to contain even for the large search engines, which drives users to other venues and communities, such as Reddit and YouTube to satisfy their information needs. Review videos on YouTube in particular are often visually appealing and entertaining, more informative (since the product can be seen in use), and they may appear more trustworthy if the review is presented by an actual (expert) person. Well-produced videos are also more costly to create, and creators have a greater interest in growing their channel with high-quality content. Hence, one might come to believe that video reviews (though often sponsored) are the last source of true hands-on reviews on the web simply due to their (still) higher production costs. However, we can show that this is, in fact, no longer true and that video product reviews are already a popular medium for spammers.

Code and data: <https://github.com/webis-de/CHIIR-24>



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

CHIIR ’24, March 10–14, 2024, Sheffield, United Kingdom  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0434-5/24/03.  
<https://doi.org/10.1145/3627508.3638303>

In this paper, we extend our work on analyzing large-scale SEO product review spam [5, 6] in search engines also to video results. Putting ourselves in the shoes of users trying to find trustworthy product reviews via regular web search engines, we analyze the transcripts of all YouTube videos embedded in the web search results for 7,392 product search queries from Startpage (who get their results directly from Google), Bing, and DuckDuckGo. Based on audio transcripts of the videos, we train a simple spam detector using simple part-of-speech (POS) n-gram features. We find that more than a third of the videos are clearly of very low quality or outright spam. Many videos disguise themselves as product reviews, but are really only product listings compiled from the web, presented like commercials using stock footage and stock music. Some even use automatic text-to-speech synthesis in place of a real human narrator. We find that a basic n-gram model—though by design it cannot assess the real value and factual quality of the video contents—is able to separate the scripted, commercial-style, and mostly low-quality videos accurately from videos with a real person reviewing products in front of the camera. We also find that of the three search engines, Startpage (due to Google being the owner of YouTube) returns by far the most YouTube results, whereas Bing and DuckDuckGo prefer textual reviews.

## 2 RELATED WORK

There is surprisingly little research on (spam) classification of YouTube videos based on subtitle transcripts. To the best of our knowledge, current YouTube video spam or scam classification systems use only the more easily accessible metadata, such as link counts, hashtags, likes, views, network features, and comments [8, 20, 21], yet these are only proxy data for the actual video contents. Related research goals such as the detection of clickbait through misleading titles and thumbnails is also metadata-based [9, 23] with few exceptions [22], even though a content analysis would make sense and datasets exist [16]. Other research is focused on user comments rather than the videos themselves in order to detect harmful or spam-like user reactions [1, 14, 19].

A field in which scholars have utilized subtitle transcripts more frequently is creating effective filters for violence or other content that is unsafe for kids [2–4, 13]. Notably, Binh et al. [7] show how incorporating subtitle features improves upon metadata-only classification for unsafe content.

Apart from this specialized direction, transcript-based video classification boils down to a classic text classification or spam detection task. Like the rest of the natural language processing community, spam classification has also shifted to (large) deep neural models and feature representations. However, traditional and much more efficient machine learning algorithms such as Naive Bayes, SVMs, or Logistic Regression trained on bag of words (BoW) text representations are still producing acceptable if not competitive results quite often [10, 17].

## 3 DATA ACQUISITION

We compiled a dataset of 4,755 videos with high-quality transcripts by (1) collecting appropriate product search queries (2) searching for product reviews on three search engines and extracting the unique reviews from the search engine results pages (SERPs), and

**Table 1: Top-20 search engine results for 7,392 product queries from Startpage (Google), Bing, and DuckDuckGo. Website counts are calculated after stripping domain names of their subdomain parts using Mozilla’s Public Suffix List.<sup>3</sup>**

Filtering steps	Startpage	Bing	DDG	All
Total results	147,658	147,592	143,823	439,073
Unique URLs	128,854	122,775	112,702	258,400
Unique websites	41,514	26,853	22,862	60,947
YouTube URLs	5,033	1,127	847	7,007
Unique YouTube URLs	4,588	1,098	810	5,902
Transcripts available	4,242	785	426	4,993
Transcripts (filtered)				4,755
Test data (ground truth obtained via manual labeling)				200
Training data (unsupervised labeling through clustering)				4,555

(3) downloading the video transcripts via the YouTube API and filtering the results. Table 1 summarizes the main dataset statistics.

To find product reviews via web search engines, we compiled a list of product categories from two publicly-available e-commerce product taxonomies: (1) the GS1 Global Product Classification<sup>1</sup> and (2) the Google Product Taxonomy.<sup>2</sup> We combined both taxonomies and removed categories such as fresh produce, live animals, as well as near duplicates. The final list contains 7,392 categories for which we constructed queries of the form “best <product category>.”

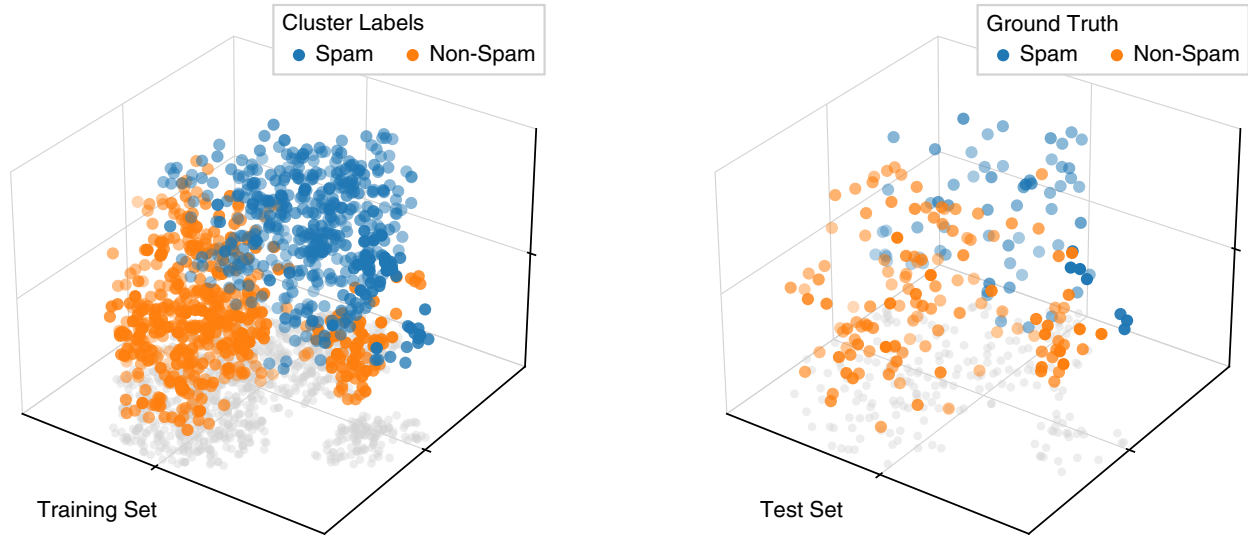
The prepared queries were sent to Startpage (as a proxy for Google), Bing, and DuckDuckGo between May 24–25th, 2023, requesting the top 20 results each. From Startpage, we retrieved 4,588 unique YouTube URLs (5,033 including duplicates) which make for 3.4 % of all search results. We used the normal web search and not the search engines’ dedicated video search, so the majority of results were non-video hits. Given 20 hits per query, 68 % of all SERPs contained at least one YouTube URL on average. For Bing and DuckDuckGo, the number was much lower with only 1,127 (0.7 %) and 847 (0.05 %) URLs, respectively. In total, we retrieved 7,007 URLs, of which 5,902 were unique. Half the videos (3,620) were ranked among the top-10 results. The median result rank across all videos and search engines was 9 (zero-indexed). The median result rank of videos on the first page was 6. Google ranked videos significantly higher (*Median* = 8) than Bing (*Median* = 12) or DuckDuckGo (*Median* = 13). The maximum number of videos per SERP was 9, but the majority of SERPs contained at most one.

For 4,993 of the video URLs, the English-language transcripts were downloadable via the YouTube API. Transcripts were provided either by the video author (less common) or were auto-generated using YouTube’s own speech recognition system. The transcripts contain time codes but neither punctuation to mark sentence structures, nor speaker diarization. Periods during which only music is played are indicated by a special “[Music]” token. We removed all transcripts that contained fewer than 100 spoken words (excluding music), which resulted in a final set of 4,755 video transcripts.

<sup>1</sup><https://www.gs1.org/standards/gpc>, Nov. 2021

<sup>2</sup><https://www.google.com/basepages/producttype/taxonomy.en-US.txt>, Sep. 2021

<sup>3</sup><https://publicsuffix.org/>



**Figure 1:** Feature vectors from the training data (left) and test data (right) embedded into a 3D space for visualization via  $t$ -SNE. The training data are labeled with their automatically assigned cluster labels, the test data with human-annotated ground-truth labels. To reduce visual clutter, only one quarter of the training data points are shown (randomly sampled).

## 4 MANUAL TEST DATA ANNOTATION

We constructed a test set of 200 manually annotated videos to evaluate the effectiveness of our approach. The examples were selected from the (filtered) dataset by uniform random sampling. Each video was annotated via single human annotation as either “Spam” or “Non-Spam”. We used a heuristic annotation approach and constructed two sets of indicators for “Spam” and “Non-Spam” based on our observations of the material. Annotation was done by skipping through the video using the scrub bar until the annotator was convinced that one set of indicators was the more prevalent. The amount of time spent per annotation was around 15–30 seconds for most videos, with some more difficult cases requiring slightly more time.

### 4.1 Indicators for “Spam”:

- Video uses lots of stock footage or is a slide show of commercial product photos.
- No actual hands-on product usage is shown on camera.
- Commentary is fully scripted and narration is performed by a text-to-speech system or a hired voice actor.
- Video or narration give the impression of a commercial rather than an unbiased review.
- Video is mostly a listing of product features and specifications from the web or the manufacturer’s website.
- Product ratings are based on or reference user reviews from online shopping sites.
- The product selection appears random.
- Products are featured in rapid succession without much additional context.

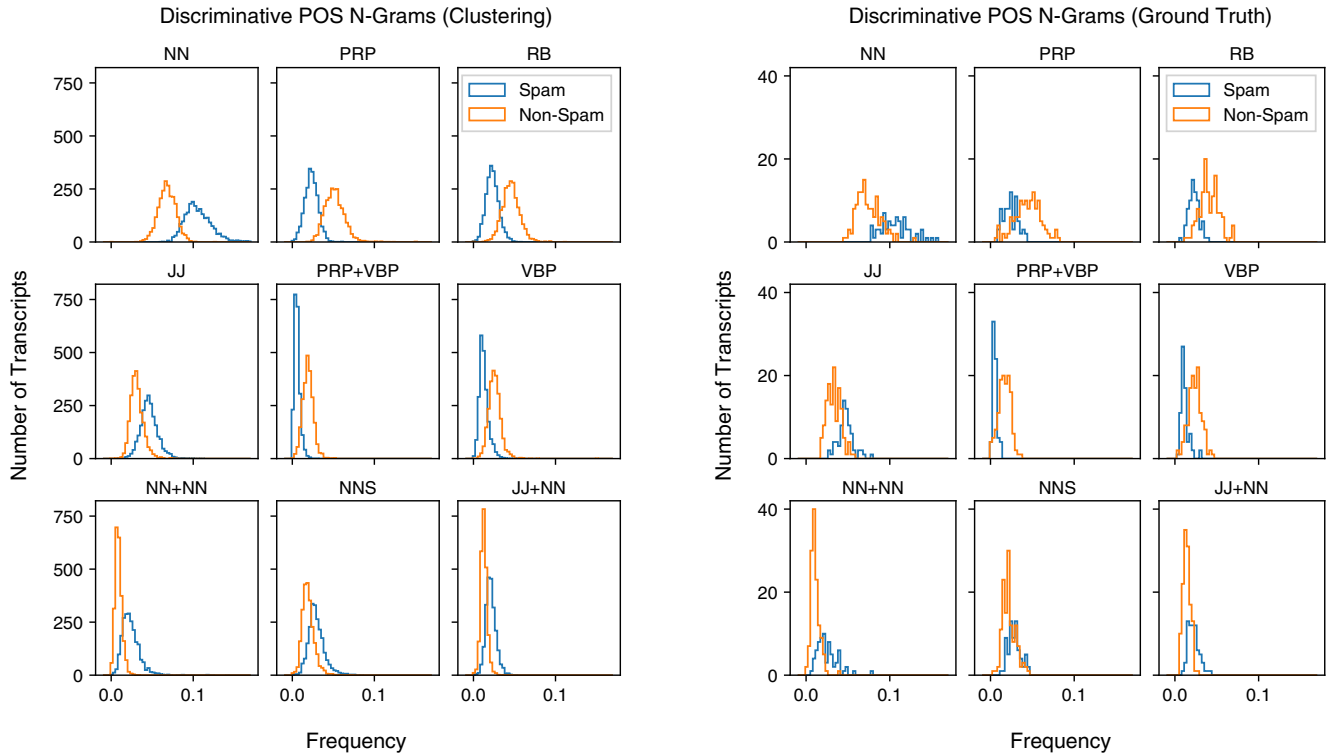
### 4.2 Indicators for “Non-Spam”:

- Video uses original footage.
- Products are shown live and in action, actual testing is performed in front of the camera.
- Video shows (at least parts of) a human protagonist on camera handling the products.
- Protagonist shares expert knowledge, providing additional, non-obvious information about the products and their usage.
- If instead, video uses off-camera commentary, sufficient expert knowledge is provided to make content believable and stand out from the “Spam” class.

Since we aimed at only a rough estimate, we deemed 200 instances annotated by three individual annotators (the two main authors of this paper and a third volunteer) sufficient. Most videos fell clearly into one of the two classes, even after watching only short segments. For the few ambiguous cases, the annotators were asked to give a personal value judgment about whether they felt informed given the information conveyed by the video in comparison to reading a chart of product specifications. Following these instructions, the annotators agreed with their label assignments to a very high degree (Krippendorff’s  $\alpha = 0.87$ ). The few cases of disagreement were resolved by majority vote. Of the 200 annotated samples, 70 (35%) fell into the “Spam” class, the other samples were deemed sufficiently believable to be considered “Non-Spam.”

## 5 TRAINING DATA CLUSTERING

For lack of large-scale gold-standard labels, we automatically labeled the training set using an unsupervised clustering approach over a stylometric feature space. Our approach assumes that videos



**Figure 2: Frequencies of the nine most discriminative POS n-grams according to the SVM hyperplane coefficients by cluster labels (left) and by ground-truth labels (right). Spam videos tend to use more singular nouns (NN) and adjectives (JJ), whereas non-spam videos tend to use more personal pronouns (PRP) and adverbs (RB). Noun and personal pronoun unigram frequencies alone explain the majority of the group differences.**

in either class use a distinctly different language that is independent of the actual product topics. This assumption is loosely backed by our impression from annotating the 200 test set instances. To obtain a topic-independent representation of linguistic markers from the transcripts, we removed all “[Music]” markers and transformed the remaining examples to part-of-speech (POS) tags using the Penn Treebank tag set [12] as assigned by the SpaCy library.<sup>4</sup> From the transformed texts, we built a feature matrix with the relative frequencies of the 150 most common POS 1–4-grams over all instances in our training data. The first 20 of these n-grams in the dataset are in descending order of frequency: NN, IN, DT, PRP, RB, JJ, VB, DT+NN, NNS, VBZ, VBP, CC, IN+DT, NN+IN, JJ+NN, PRP+VBP, NN+NN, DT+JJ, VBG, TO.

We tested several clustering approaches from the scikit-learn library<sup>5</sup> to separate the data points into two stable sets. We found that a basic K-Means or spectral clustering produced the most stable clusters. DBSCAN as a density-based clustering was too unstable and sensitive to hyperparameter settings to be practical, which hints at the data being rather uniform in density and thus not easily separable into disjunct clusters with clear boundaries. In the end, we settled with the spectral clustering, as it makes no assumptions about the shapes of clusters (unlike K-Means), but still allows to control the number of clusters.

<sup>4</sup><https://spacy.io/>, Version 3.6.0

<sup>5</sup><https://scikit-learn.org/stable/>, Version 1.3.0

Of the 4,555 clustered samples, 2,202 (48.3%) fell into Cluster 1 and 2,353 (51.7%) into Cluster 2. Figure 1 (left) shows a 3D  $t$ -SNE projection of the produced clusters. Following the class distribution in the test set, we assigned the smaller of the two clusters the class “Spam” and the larger one the class “Non-Spam.” We can see already that the portion of “Spam” instances is 13 percentage points larger in the clustered training set than in our manually labeled test set (35%). This is not surprising considering that either the data itself or the extracted features are not expressive enough and seem to be lacking some amount of structure that would allow for clear separation. In such a case, the clustering approach would try to separate the data into approximately equally-sized parts, which is very likely what we are seeing here. Although not optimal, this is not a huge problem as long as the produced cluster centroids are stable and reproducible. Luckily, a projection of the labeled test data into the same 3D space (Figure 1, right) reveals very similar cluster locations and shapes, differing mostly only in cluster size.

## 6 EVALUATION AND DISCUSSION

As discussed in the previous section, projecting the clustered and the human-labeled examples into the same 3D space (Figure 1) reveals a high agreement between the clustering and the ground truth, which strongly suggests that our approach may be effective. However, the clustering itself has the inherent problem that it

**Table 2: Classification results on 200 test samples (Non-Spam: 130, Spam: 70) after training a linear SVM and a logistic regression model on the cluster labels.**

Model	Class	Prec.	Recall	F1	AUROC
SVM	Non-Spam	1.00	0.78	0.88	0.95
	Spam	0.71	1.00	0.83	
Log. Regression	Non-Spam	0.99	0.82	0.89	0.95
	Spam	0.74	0.99	0.85	

cannot easily be reproduced on new, unseen data to classify spam in a production environment. So instead of calculating an accuracy score directly on the clusters, we used the cluster labels as silver-standard training targets to produce a reusable supervised model, which could then be shipped to production. We trained both a linear SVM and a logistic regression model in this way and evaluated their effectiveness on the human-labeled test set. Table 2 summarizes the classification results on the ground-truth examples. Both classifiers achieve an F1 score of around 0.85–0.89 and an AUROC of 0.95. The “Non-Spam” precision and the “Spam” recall are virtually perfect with 1.0 and 0.99, though it should be taken into account that with 130 examples (65%), “Non-Spam” is also the majority class. Despite the not optimally-chosen cluster boundaries, this confirms that our system does capture the target concept well, hence being a working (although tunable) classification approach for affiliate product review spam on YouTube.

Ranked by the SVM’s highest absolute hyperplane coefficients, we identify the unigram frequencies of NN, PRP, RB, and JJ as the most discriminative features (see Figure 2). These are followed by PRP+VBP, VBP, and NN+NN. Regarding the differences between the classes, “Spam” examples tend to have higher NN and JJ frequencies, whereas the “Non-Spam” examples tend to have higher PRP and RB frequencies. Most of the differences between the two groups are explained by the NN and PRP frequencies alone. Consulting the ground-truth examples, we observe about the same frequency distribution for these n-grams. This means that as their primary distinctive language features, “Spam” transcripts make more frequent use of nouns and adjectives, whereas “Non-Spam” transcripts have higher usage of personal pronouns and adverbs. The finding makes sense considering our annotation guidelines. Videos that are by and large a summary or a listing of product features fulfill key criteria for the “Spam” class. Videos featuring a real human protagonist talking about the product and their personal experience with it, on the other hand, are far more likely to be “Non-Spam.”

## 7 LIMITATIONS

The study shows promising results, but has a few limitations. First of all, the number of ground-truth, but also training samples is quite low. Second, the high annotator agreement and high classification precision indicate that most cases are rather easy to decide to begin with. Third, the classifier is able to predict the “Non-Spam” majority class with higher precision than the “Spam” class, resulting in quite a few false positives. A correction of the decision boundary is quite feasible, but would require actual human-labeled training examples, additional features that are more indicative of the intended

target concept, or at least an externally tuned model bias. Moreover, the (non-)spam classifier is able to identify obvious spam content quite reliably, but it cannot check the factual accuracy of the review contents. Finally, the findings are limited to English-language search results for synthetic product queries from three commercial blackbox web search engines.

## 8 CONCLUSION

We have developed an effective yet simple unsupervised spam classifier for improving product-centric video results of web search engines. We examined the YouTube video results of three major commercial search engines for thousands of product review queries. For all collected video results, we retrieved the automatic subtitle transcripts via the YouTube API and automatically labeled them using a spectral clustering based on POS n-gram frequency representations. We verified the clustering accuracy by training a supervised linear model on the unsupervised labels and compared the results to a small set of human-annotated examples. It turns out that the resulting clusters approximately capture the classes of “Spam” and “Non-Spam” videos and most of the differences are explained by the use of nouns, personal pronouns, adverbs, and adjectives. However, the decision boundary is not learned perfectly ( $F_1 = 0.88$ , AUROC = 0.95), which leaves potential for further optimization.

We conclude that since the portion of “Spam” videos seems quite high and we were able to detect them so easily, search engines should apply more careful filtering of their video results. Our case study can be thought of as a first step towards better product spam recognition in online product review videos.

## ACKNOWLEDGMENTS

This publication has received funding from the European Commission under grant agreement № 101070014 (OpenWebSearch.eu).

## REFERENCES

- [1] Túlio C. Alberto, Johannes V. Lochter, and Tiago A. Almeida. 2015. TubeSpam: Comment Spam Filtering on YouTube. In *14th IEEE International Conference on Machine Learning and Applications, ICMLA 2015, Miami, FL, USA, December 9–11, 2015*, Tao Li, Lukasz A. Kurgan, Vasile Palade, Randy Goebel, Andreas Holzinger, Karin Verspoor, and M. Arif Wani (Eds.). IEEE, 138–143. <https://doi.org/10.1109/ICMLA.2015.37>
- [2] Sharifa Alghowinem. 2019. A safer youtube kids: An extra layer of content filtering using automated multimodal analysis. In *Intelligent Systems and Applications: Proceedings of the 2018 Intelligent Systems Conference (IntelliSys) Volume 1*. Springer, 294–308.
- [3] Saeed Ibrahim Alqahtani, Wael MS Yafooz, Abdullah Alsaeedi, Liyakathunisa Syed, and Reyadh Alluhaibi. 2023. Children’s Safety on YouTube: A Systematic Review. *Applied Sciences* 13, 6 (2023), 4044.
- [4] Vishal Anand, Ravi Shukla, Ashwani Gupta, and Abhishek Kumar. 2019. Customized video filtering on YouTube. *CoRR* abs/1911.04013 (2019). arXiv:1911.04013 <http://arxiv.org/abs/1911.04013>
- [5] Janek Bevendorff, Matti Wiegmann, Martin Potthast, and Benno Stein. 2022. The Impact of Online Affiliate Marketing on Web Search. In *4th International Symposium on Open Search Technology (OSSYM 2022)*, Andreas Wagner, Christian Guehl, Michael Granitzer, and Stefan Voigt (Eds.). International Open Search Symposium.
- [6] Janek Bevendorff, Matti Wiegmann, Martin Potthast, and Benno Stein. 2024. Is Google Getting Worse? A Longitudinal Investigation of SEO Spam in Search Engines. In *Advances in Information Retrieval. 46th European Conference on IR Research (ECIR 2024) (Lecture Notes in Computer Science)*. Springer.
- [7] Le Binh, Rajat Tandon, Chingis Oinar, Jeffrey Liu, Uma Durairaj, Jiani Guo, Spencer Zahabizadeh, Sanjana Ilango, Jeremy Tang, Fred Morstatter, Simon S. Woo, and Jelena Mirkovic. 2022. Samba: Identifying Inappropriate Videos for Young Children on YouTube. In *Proceedings of the 31st ACM International*

- Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17–21, 2022*. Mohammad Al Hasan and Li Xiong (Eds.). ACM, 88–97. <https://doi.org/10.1145/3511808.3557442>
- [8] Elijah Bouma-Sims and Bradley Reaves. 2021. A First Look at Scams on YouTube. *CoRR* abs/2104.06515 (2021). arXiv:2104.06515 <https://arxiv.org/abs/2104.06515>
- [9] Ruchira Gothankar, Fabio Di Troia, and Mark Stamp. 2021. Clickbait Detection in YouTube Videos. *CoRR* abs/2107.12791 (2021). arXiv:2107.12791 <https://arxiv.org/abs/2107.12791>
- [10] Thiago S. Guzella and Waldir M. Caminhas. 2009. A review of machine learning approaches to Spam filtering. *Expert Syst. Appl.* 36, 7 (2009), 10206–10222. <https://doi.org/10.1016/j.eswa.2009.02.037>
- [11] Dirk Lewandowski, Sebastian Sünkler, and Nurce Yagci. 2021. The influence of search engine optimization on Google's results: A multi-dimensional approach for detecting SEO. In *Proceedings of the 13th ACM Web Science Conference 2021* (Virtual Event, United Kingdom) (*WebSci '21*). Association for Computing Machinery, New York, NY, USA, 12–20.
- [12] Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a Large Annotated Corpus of English: The Penn Treebank. *Computational Linguistics* 19, 2 (1993), 313–330. <https://aclanthology.org/J93-2004> [https://www.ling.upenn.edu/courses/Fall\\_2003/ling001/penn\\_treebank\\_pos.html](https://www.ling.upenn.edu/courses/Fall_2003/ling001/penn_treebank_pos.html)
- [13] Enrico Mariconti, Guillermo Suarez-Tangil, Jeremy Blackburn, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Jordi Luque Serrano, and Gianluca Stringhini. 2019. "You Know What to Do": Proactive Detection of YouTube Videos Targeted by Coordinated Hate Attacks. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.
- [14] P Nagaraj, K Muthamil Sudar, P Thrived, P Girish Kumar Reddy, Sk Baji Babu, and P Siva Rama Krishna. 2023. Youtube Comment Spam Detection. In *2023 International Conference on Computer Communication and Informatics (ICCCI)*. IEEE, 1–6.
- [15] Kristen Purcell, Joanna Brenner, and Lee Rainie. 2012. *Search Engine Use 2012*. Technical Report. Pew Research Center's Internet & American Life Project.
- [16] Jiani Qu, Amny Marleen Hißbach, Tim Gollub, and Martin Potthast. 2018. Towards Crowdsourcing Clickbait Labels for YouTube Videos. In *6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2018)*, Yiling Chen and Gabrielle Kazai (Eds.).
- [17] Sanjeev Rao, Anil Kumar Verma, and Tarunpreet Bhatia. 2021. A review on social spam detection: Challenges, open issues, and future directions. *Expert Syst. Appl.* 186 (2021), 115742. <https://doi.org/10.1016/j.eswa.2021.115742>
- [18] Sebastian Schultheiß, Helena Häußler, and Dirk Lewandowski. 2022. Does Search Engine Optimization come along with high-quality content? A comparison between optimized and non-optimized health-related web pages. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) (*CHIIR '22*), David Elswiler (Ed.). Association for Computing Machinery, New York, NY, USA, 123–134.
- [19] Sadia Sharmin and Zakia Zaman. 2017. Spam Detection in Social Media Employing Machine Learning Tool for Text Mining. In *13th International Conference on Signal-Image Technology & Internet-Based Systems, SITIS 2017, Jaipur, India, December 4–7, 2017*, Kokou Yétongnon, Albert Dipanda, Richard Chbeir, Luigi Gallo, and Neeta Nain (Eds.). IEEE Computer Society, 137–142. <https://doi.org/10.1109/SITIS.2017.32>
- [20] Ashutosh Tripathi, Kusum Kumari Bharti, and Mohona Ghosh. 2019. A Study on Characterizing the Ecosystem of Monetizing Video Spams on YouTube Platform. In *Proceedings of the 21st International Conference on Information Integration and Web-based Applications & Services, iiWAS 2019, Munich, Germany, December 2–4, 2019*. ACM, 222–231. <https://doi.org/10.1145/3366030.3366078>
- [21] Ashutosh Tripathi, Mohona Ghosh, and Kusum Kumari Bharti. 2022. Analyzing the uncharted territory of monetizing scam Videos on YouTube. *Soc. Netw. Anal. Min.* 12, 1 (2022), 119. <https://doi.org/10.1007/s13278-022-00945-1>
- [22] Deepika Varshney and Dinesh Kumar Vishwakarma. 2021. A unified approach for detection of Clickbait videos on YouTube using cognitive evidences. *Appl. Intell.* 51, 7 (2021), 4214–4235. <https://doi.org/10.1007/s10489-020-02057-9>
- [23] Savvas Zannettou, Sotirios Chatzis, Kostantinos Papadamou, and Michael Sirivianos. 2018. The Good, the Bad and the Bait: Detecting and Characterizing Clickbait on YouTube. In *2018 IEEE Security and Privacy Workshops, SP Workshops 2018, San Francisco, CA, USA, May 24, 2018*. IEEE Computer Society, 63–69. <https://doi.org/10.1109/SPW.2018.00018>