Analyzing YouTube Kids' recommendation algorithm for content diversity

OLAF ADAMS, University of Twente, The Netherlands

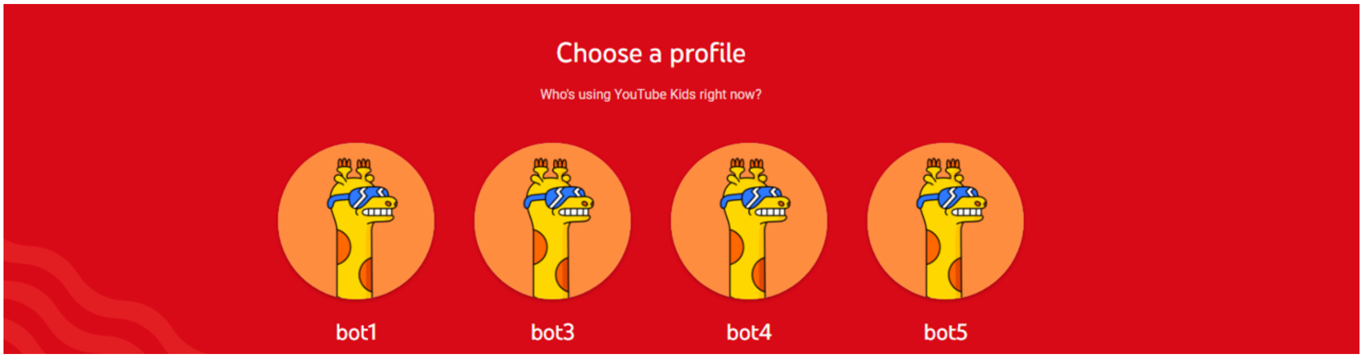

Fig. 1. Profile selection screen from YouTube Kids

In this paper, the behavior of the YouTube Kids recommendation algorithm and its potential influence on the formation of content filter bubbles is investigated. YouTube Kids videos are collected by a Playwright pipeline, which are then included in a recommendation tree. The results show that the entropy remains stable along the average recommendation sequence, going against the claim of the algorithm pushing users into filter bubbles. Transition matrices do showcase potential biases in the algorithm towards recommending certain content, particularly videos in the 'Cartoons' category. The matrices also show high self-transition probabilities, indicating a high likelihood for the algorithm to recommend the same genre of video multiple times in a row. These findings suggest moderate content diversity on the platform, but also call for action to improve the algorithms' behavior to positively affect children's exposure to a variety of content.

Additional Key Words and Phrases: YouTube Kids, recommendation algorithm, content diversity, filter bubbles, zero-shot classification, Shannon entropy, transition matrix, children's media

1 INTRODUCTION

YouTube Kids is an online video-watching platform, designed to be safe for children. These children, including those as young as 8 years old, may spend up to 3 hours a day watching videos online [11, 22]. Often, children use the platform unsupervised, becoming fully dependent on the recommendation algorithm to keep them occupied and invested in the platform. An important topic for a platform such as YouTube Kids is whether the recommendation algorithm promotes filter bubbles.

Filter bubbles [17] are a phenomenon that occurs whenever a recommendation algorithm continuously recommends similar content to a user. This phenomenon can be harmful to children [24], as they are still in a developmental phase and benefit from being exposed to a wide variety of content. Filter bubbles could potentially restrict learning and curiosity in a child. They essentially construct a digital echo chamber [6, 23] in which a child's viewing habits are controlled without them realizing it.

1.1 Background

YouTube Kids is a platform specifically designed for children to watch safe YouTube videos. Content types vary widely from educational to mostly entertainment. Parents have great incentive to favor YouTube Kids over YouTube, as the kids' version is generally perceived as much safer and designed with children in mind. While regular YouTube allows users to interact with videos by leaving a like or a comment, YouTube Kids only allows users to watch a video and proceed to one of the next recommendations. The user interface is also simplified, only displaying video thumbnails and their titles. The platform allows parents to set up custom profiles for their children under a single Google account, and allows the parent to enter the child's age range. Parents are given the option to choose from three age ranges: 4 years and under, 5 to 8 years, and 9 to 12 years old. However, much of the YouTube Kids experience is dictated by the recommendation algorithm, automatically serving users a list of potentially interesting videos to watch next.

These recommendation algorithms shape most social media platforms by predicting user preferences. Although this type of personalization enhances user engagement, it risks the formation of 'filter bubbles', where users are repeatedly shown similar content. This is already a concern among adults [7], but for children, the implications of this may be even more impactful. Research suggests that prolonged exposure to these recommendation algorithms can lead to difficulties in concentration, information retention, and ultimately affect attention span [4]. Content diversity will later be quantified through Shannon entropy and genre distributions.

Author's address: Olaf Adams, o.adams@student.utwente.nl, University of Twente, P.O. Box 217, Enschede, The Netherlands, 7500AE.

TScIT 43, July 4, 2025, Enschede, The Netherlands

1.2 Problem Statement

Despite widespread use of YouTube Kids, research regarding content diversity on the platform is limited. The recommendation algorithm could limit children's exposure to new learning opportunities and might reduce engagement with a wide range of topics. This is especially relevant in today's digital world.

1.3 Research Objectives

This research will attempt to answer the following question:
To what extent does the YouTube Kids recommendation algorithm contribute to the formation of content filter bubbles?
To help answer this question, several sub-questions are proposed.

(1) **How does content diversity evolve through recommendation sequences on YouTube Kids?**
(2) **What roles do specific video subcategories play in recommending content to YouTube Kids users?**
(3) **Does the YouTube Kids recommendation algorithm steer users into filter bubbles over time?**

2 RELATED WORK

To gather related literature, sources such as Google Scholar, Scopus, and IEEE were used. Search terms such as "Filter Bubbles", "YouTube Kids", "Recommendation Algorithms", and "Topic Classification" were used to gather state-of-the-art papers to form a basis for the research.

2.1 Filter Bubbles

The concept of filter bubbles was first introduced by Pariser [17], who explored the topic and argued that recommendation algorithms can limit exposure to diverse content and viewpoints. Later, studies conducted by Brown et al. [6] and Haroon et al.[9] indicate that conventional YouTube's recommendation algorithm can lead users into rabbit holes and filter bubbles. In this research, the platform repeatedly recommends content supporting the same viewpoints or adhering to similar categories. These concerns may have been established for YouTube, but research on these issues on YouTube Kids is missing.

2.2 Content Diversity

Measuring diversity is an effective method for evaluating a recommendation algorithm. Studies by Adomavicius and Kwon [1] and Antikacioglu and Ravi [3] have proposed ranking-based techniques for recommending content to users based on diversity. Metrics were proposed to quantify content diversity, such as the Gini coefficient [21] and Shannon entropy [15]. Entropy has proven to offer a clear perspective on content diversity and has also been applied to study diversity in music [2].

2.3 YouTube Kids

While YouTube Kids is designed to be a safe alternative to conventional YouTube, studies suggest that not all content is suitable for children. Papadamou et al. [16] and Tahir et al. [25] have emphasized these dangers of disturbing content on the platform.

2.4 YouTube Video Classification

Accurate video classifications are vital for analyzing recommendation sequences for content diversity. Traditional approaches, such as Karpathy et al. [12], Ramesh et al. [19] have relied on analyzing video frames with convolutional neural networks to classify videos. However, supervised approaches such as this one are resource-intensive, and recent work, such as Raza et al. [20] has shifted its focus towards utilizing text-based video metadata, such as description, title, tags, and transcripts. In particular, zero-shot learning enables moderately accurate classification without the need for task-specific training data. Transformer-based models such as BART [13] have shown promise in effectively labeling data in various domains [5]. When combined with prompt engineering, they become a useful tool enabling scalable and adaptable genre categorization.

3 METHODOLOGY

Performing the research involved several steps, including data collection, data parsing, classification, and diversity analysis. Firstly, a YouTube Kids video dataset was constructed that can be analyzed for content diversity and signs of filter bubbles. This dataset contained relevant information to distinguish and analyze videos. The videos were then classified into subcategories based on this metadata. Finally, the data are analyzed for content diversity. All actions were performed using Python programs [18].

3.1 Data Collection

Constructing a comprehensive dataset was vital for the project's success. To achieve this, care was taken to ensure the collected data was as representative as possible of the platform as a whole.
To start, a list of subcategories was constructed by manually analyzing the dataset of videos and the distribution of YouTube categories. Then, an attempt was made to expand on these categories, making them more descriptive and representative of the content on the platform. This list not only provided the main classification categories for the video classification but also served as a starting point for collecting video recommendations. For each subcategory, a video representative of the category was manually selected. This set of videos served as seed videos for the crawling process.
To make effective use of the chosen seed videos, a pipeline was constructed emulating a regular YouTube Kids user. The pipeline made use of Playwright [14], a tool for automating browser interactions. This tool allowed for the construction of recommendation graphs guided by the YouTube Kids recommendation algorithm. Through Playwright, automatic navigation through the platform became trivial.
The pipeline started from one specific seed video and collected the video IDs of the first 3 video recommendations. Only the first 3 recommendations were selected for parsing, due to the way the YouTube Kids interface is structured (See Figure 2). Unless the user scrolls down, only 3 initial recommendations are visible. The pipeline performed a breadth-first search traversal through the YouTube Kids video recommendations, starting from each seed video on a brand new YouTube Kids profile. Using a fresh account for each search is important, as it eliminates potential algorithmic biases based on previous video-watching behavior. It then constructed

a weighted directed graph, representing possible user navigation through the recommendation network. A depth of 6 was chosen to ensure sufficient data coverage while simultaneously adhering to time and hardware constraints. A sample of such a graph can be found in Figure 3. Table 1 showcases a comparison between the custom classifications, YouTube's categorizations, and the number of videos collected for each of them.

Table 1. Comparison of Custom Classifier Labels and YouTube Categories

Custom Classification	Count	YouTube Category	Count
Cartoons	1948	Film & Animation	1903
Minecraft gameplay	607	Entertainment	1326
Real-Life Challenges	571	Education	1162
Language learning	418	Gaming	837
Gaming	414	People & Blogs	190
Children's Songs	381	Sports	90
Animated short films	233	Music	68
Sports	220	Howto & Style	61
Science education	182	Science & Technology	47
Drawing Videos	171	Pets & Animals	42
Dance videos	154	Travel & Events	13
Math education	110	Comedy	8
Comedy and Skits	90	Autos & Vehicles	7
DIY Crafts	60		
Cooking for kids	59		
Roblox Gameplay	59		
Toy Reviews	47		
LEGO Building	30		

Fig. 2. YouTube Kids Video Interface

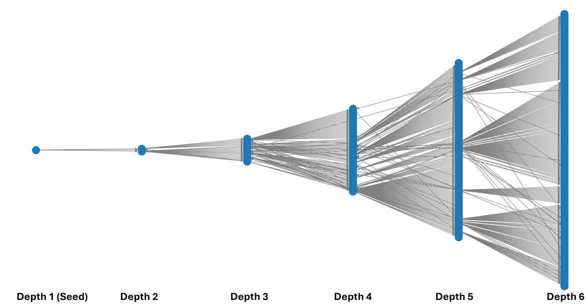

Fig. 3. Recommendation Graph Sample

3.2 Data Parsing

In this context, a subcategory refers to a manually defined content class and includes classes such as 'Gaming', 'Cartoons', or 'Science'. These content classes were used to collect seed videos for the crawling process, as well as to label the resulting data. The categories were determined through manual inspection of a broad sample of YouTube Kids videos, with emphasis on capturing distinct content themes on the platform.

Once a recommendation tree had been constructed for all 18 distinct subcategories, the trees were merged into one recommendation tree. This resulted in a weighted directed graph, with the edges pointing to the recommended videos for each video in the dataset. The weights of the edges were calculated based on how many times the video was recommended from that video and how high the video was ranked on average in the recommendation list. This ranking refers to the video being the first, second, or third recommendation. The edges were weighted in this way to ensure consistently higher-ranked recommendations have more influence in the resulting graph.

Let

$$R_v = \{(r_i, \mathrm{rank}_i)\}$$

be the list of video recommendations from video v and

$$s_i = 3 - rank_i$$

the score of a video i in the recommendation list. The normalized edge weight from video v to r_i then becomes

$$w(v, r_i) = \frac{3 - \mathrm{rank}_i}{S_v}, \quad \text{where} \quad S_v = \sum_{(r_j, \mathrm{rank}_j) \in R_v} (3 - \mathrm{rank}_j)$$

The final tree contains 5754 unique videos, collected through the YouTube Kids recommendation algorithm.

Thanks to the constructed recommendation tree, it became trivial to collect recommendation sequences from the platform. As the data collection was already done, 1.000.000 random walks through the recommendation tree were collected. The number of walks was chosen to be sufficiently large to ensure good representation of the overall dataset. The walks were collected by starting from a seed video initially used for the data collection process. Then, the edge weights were used as probabilities for the simulated user to navigate to a specific recommendation. Whenever a video recommendation was encountered that had already been seen in that specific walk, it was discarded, as the chance of a user re-watching a video is very small.

3.3 Classification

Effectively analyzing content diversity on a platform such as YouTube Kids requires the collection of as much useful metadata as possible. Unfortunately, YouTube Kids does not have an official API; however, one discovery that was made is that all YouTube Kids videos are also present on YouTube. Because of this, extracting the video ID from

YouTube Kids allows the use of the YouTube Data API v3 [8]. This API provides useful metadata, in the form of video title, description, tags, channel name, and user-picked video category.
These metadata form a good understanding of how the video presents itself; however, it is also important to take the contents of the video into account. To achieve this, the video transcripts can be of use. The transcripts for all videos in the dataset were collected through the YouTube-transcript-api [10]. Unfortunately, transcripts were not available for all videos in the dataset, but for roughly 90% of videos, transcripts were available.
Although videos have already been partitioned into categories by the creator of the video, it is very important to construct a distinct set of possible classifications. This is because a creator of a video may select a video category that does not match the actual content. Furthermore, it should also be noted that YouTube video categories were originally created with regular YouTube in mind, not YouTube Kids, resulting in several categories not having substantial representation on the YouTube Kids platform. Examples of these types of categories include 'Autos & Vehicles', 'Nonprofits & Activism', and various others.
In order to accurately classify each video into one of the 18 manually defined content categories, a zero-shot classifier was used, utilizing a natural language processing model (NLP). For the classification of the videos, a pre-trained zero-shot classifier was used from Huggingface's Transformers Library, specifically the bart-large-mnli model [13]. This model is based on the BART or RoBERTa architecture, which is popular for natural language inference (NLI). Zero-shot classification is particularly useful for this use case, as there is a wide range of topics used for classification. A different approach may have required manually labeling a large dataset, ensuring sufficient coverage for training a custom model. This time-consuming task can now be automated thanks to the zero-shot classifier. The NLI model can be used for classification purposes by cleverly structuring the data fed into it. The model is generally used for evaluating a given hypothesis, and for this use case, the data can be reformulated as such. For classifying a specific video, the hypothesis can be formulated as "This video is about <topic>", where the model evaluates the hypothesis based on a given premise, which in this case is the video metadata. In this context, the model returned a confidence score for each of the provided classifications, clearly indicating which class fit the videos best.
As the model is based on natural language inference, it is vital for the premise to be formatted in such a way that it conforms to the general expectations of the model. If the input was noisy or inconsistently formatted, the model may become confused, leading to lower confidence scores or even incorrect classifications. For this reason, each field of metadata was clearly labeled. Specifically, each premise began with the video's title, followed by the top 10 associated tags, a truncated portion of the video description (up to 300 characters), as well as a snippet of the video transcript (up to 500 characters). All components of the premise were cleaned by removing line breaks and excess whitespace to maintain consistency within input data. Fields were also concatenated using clear separators; in this case, the " | " character was used. This formatting ensured the classifier was able to effectively interpret the available data, even when certain fields were sparse or fully missing.

4 RESULTS

In this section, the results of the study and data analysis will be discussed. First, the metrics used to evaluate content diversity on YouTube Kids will be covered, followed by an interpretation of the results.

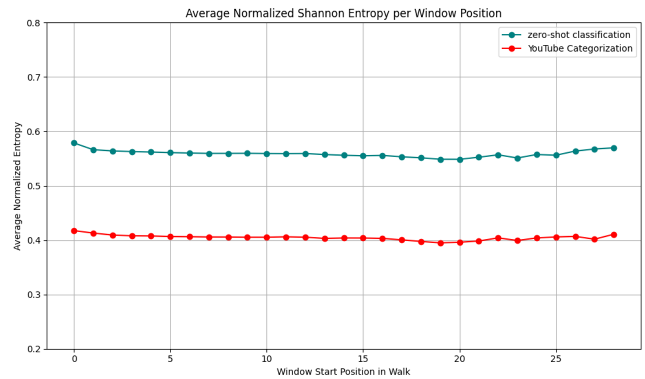

Fig. 4. Average normalized entropy for sliding window through recommendation walks

4.1 Diversity over Time

Getting a good grasp on how content diversity evolves as a user navigates through the platform can produce valuable insights into how the recommendation algorithm functions as a whole. One such metric that can help understand this better is Shannon entropy, which quantifies the unpredictability or diversity of a given distribution. In the context of YouTube Kids, the metric can be applied effectively to the distribution of video categories and classifications within recommendation sequences. A high entropy value would indicate a diverse mix of content categories, supporting the idea of an algorithm promoting content diversity. Conversely, a relatively low entropy suggests low diversity.
Additionally, it proves useful to measure Shannon entropy throughout walks. This can be achieved by implementing a sliding window approach and measuring Entropy as the algorithm carries the user further into a specific walk.

Let a set of videos have empirical probabilities $p_1, p_2, \ldots, p_k$, and let T be the total number of known possible categories. Then the normalized Shannon entropy is defined as:

$$H_{\text{norm}} = -\frac{1}{\log_2(\min(T, n))} \sum_{i=1}^{k} p_i \log_2 p_i$$

Where:

- p_i is the proportion of occurrences of item i,
- k is the number of unique observed items,
- n is the number of items in the sequence,
- T is the total number of possible topics.

Figure 4 showcases the average normalized Shannon entropy for all random walks through the recommendation network. Entropy was calculated with YouTube's default content categorizations as

Fig. 5. Transition matrix (YouTube Categories)

well as the zero-shot classifications. As is visible from the relatively straight line, there is very little change in entropy as the average user traverses a random walk. Shannon entropy stays relatively constant at a value of 0.4, which, although on the lower end, suggests a persistent medium level of content diversity. The algorithm may not fully collapse into a fully homogeneous set of recommendations (yielding an entropy near 0), but it also does not maintain a high level of diversity (approaching a level of 1). This finding implies that the recommendation algorithm tends to limit user exposure to a limited subset of available content classes and categories.
Aside from YouTube's categorizations, entropy can also be analyzed through the lens of our zero-shot classifier, also visible in Figure 4. As is visible, the average normalized Shannon entropy is significantly higher when it is based on zero-shot classification than on YouTube's classification. This increase in average entropy can be explained by the fact that the classifier provides a more fine-grained and diverse view of the content. While YouTube's categorizations tend to group content in broad labels, the classifier can distinguish between subtle variations.
More importantly, the average entropy for YouTube's classifications and the zero-shot classifier stay relatively constant, indicating that content diversity does not change significantly as a user watches more videos during a singular viewing session.

4.2 Transition matrices

More valuable insights can be retrieved by analyzing topic transitions within the network. Figure 5 displays this, specifically for YouTube's categorizations. It should be noted that video categories with less than 50 entries have been excluded from the results, as they may disproportionally skew the statistics. This transition matrix was constructed by analysis of the random walks through the recommendation network. Probabilities are calculated by counting the number of transitions from one topic to another, then normalizing the values with the total number of transitions. The figure showcases the empirically observed probabilities of transitioning from one topic to another. The diagonal of the transition matrix represents the transitions of certain categories into themselves. The data suggests that when watching a video categorized as 'Gaming' or 'Education', there is a more than 70% chance that the next video a user watches will be in the same category.
To further investigate potential algorithmic biases, the concept of dominant genres can be introduced. A genre can be considered dominant if the mean of its incoming transition probabilities significantly exceeds the global mean. Specifically, a genre is dominant if its mean incoming transition probability exceeds the global mean by more than one standard deviation. This threshold allows for concretely measuring genres that receive a disproportionately high number of incoming transitions. It is an objective measure for identifying content categories that may disproportionally benefit from the recommendation algorithm.
When this definition is applied to the YouTube category matrix, it is found that:

- The global mean transition probability: 0.125
- Standard deviation: 0.175
- Dominance threshold: 0.3

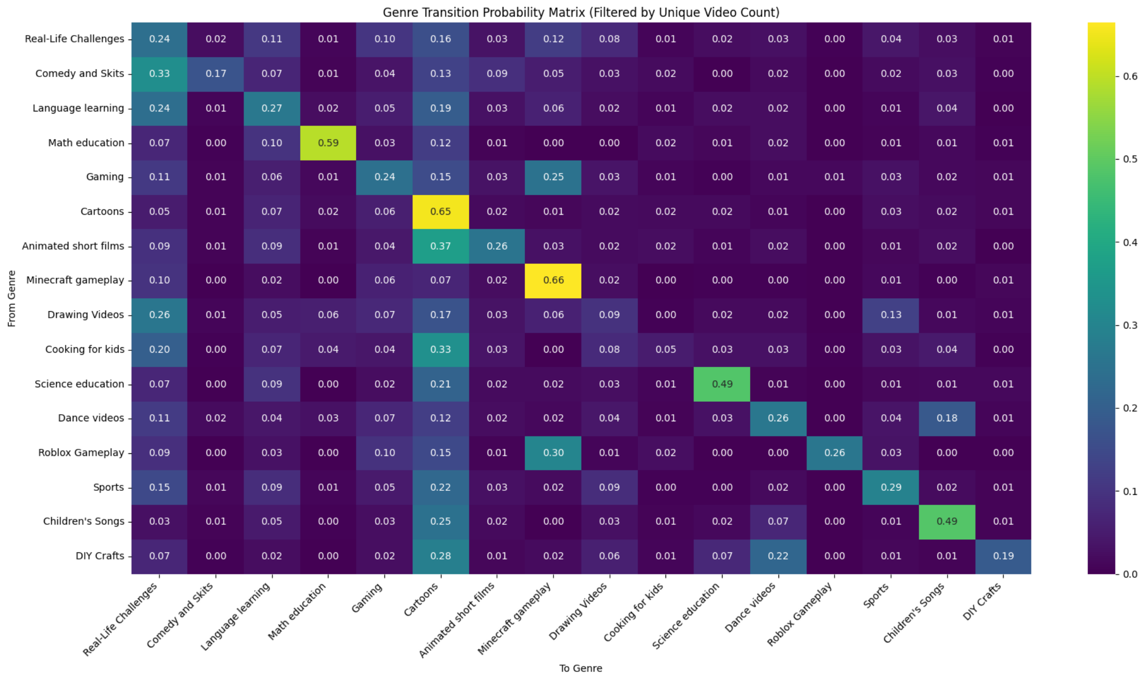

Fig. 6. Transition matrix (Zero-shot Classifications)

Going off of these values, no singular category exceeds the dominance threshold in the data.
Similarly, the zero-shot classification matrix in Figure 6 can be analyzed for dominant genres as well.

- The global mean transition probability: 0.077
- Standard deviation: 0.149
- Dominance threshold: 0.226

In this case, there is a genre exceeding the threshold, namely 'Cartoons'.
While these findings suggest that certain genres attract a disproportionate share of transitions, this result is more descriptive than causal. Though the 'Cartoons' genre appears more frequently as recommendations, it cannot be directly inferred that this is due to algorithmic bias.

4.3 Genre Drift Across Recommendation Depths

Analyzing the individual graphs before they are merged into the full recommendation tree can provide additional insight into how certain genres are represented as a user watches more videos. Figure 7 shows one such graph, starting from a video related to basic math education. The graph showcases the distribution of video genres for each recommendation depth. Figure 8 contains graphs for each different genre. As is visible in the graphs, the proportion of the starting genre gradually stabilizes. The graphs mostly display a clear downward trend when focusing on the starting genre, with few outliers. More genres get represented in the distributions as depth increases. Generally, there is an even distribution of genres by depth 6. These results support the idea of the algorithm providing sufficient diversity in its recommendations to children.

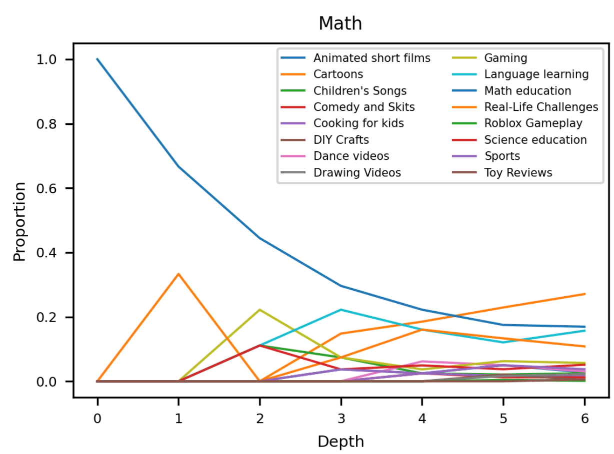

Fig. 7. Content Class distribution with Math-related seed video

5 DISCUSSION

This study has helped shed light on the behavior of the YouTube Kids recommendation algorithm, as well as its influence on the perception of content diversity and filter bubbles on the platform. It has accomplished this by analyzing the Shannon entropy of recommendation sequences, which maintained a stable level, indicating consistent levels of diversity. While this does suggest that users do not fall into narrowing filter bubbles, the Entropy values are not as high as may be expected, were there a lot of content diversity on the platform.

Utilizing the zero-shot classifier has helped reveal more fine-grained patterns compared to YouTube's classifications. However, content diversity remains mostly stable, even with more nuanced classifications, as is evident from the Shannon entropy.

These findings are further supported by the transition matrices through the display of a strong preference for videos to be followed by other videos within the same genre. Particularly, the 'Cartoons' genre from the zero-shot classifier emitted signs of dominance in the transition matrix, suggesting a potential bias in the recommendation algorithm towards recommending certain types of content. Biases such as this may be influenced by content popularity or availability, but still present the risk of children becoming overexposed to specific types of content.

5.1 Limitations

Further research can be done with larger datasets. It could also incorporate more realistic user behavior regarding the weights of the edges in the recommendation network. In this research, only the ranking of the recommendation was taken into account. However, a realistic user would also be influenced by the video title and quality of the thumbnail.

Additionally, video classification could be further explored by incorporating more metadata. Full transcripts could be used, as well as data extracted from video frames and thumbnails for more accurate video classifications.

It is also important to note that, although carefully defined, the manually constructed list of categories has limitations. Future studies could improve upon this by taking larger samples of videos and finding a dynamic methodology for generating content categories.

6 CONCLUSION

This study has combined random walk simulations, entropy measures, and topic transition matrices to analyze the YouTube Kids recommendation algorithm. Videos were classified with a zero-shot classifier, allowing for the assessment of diversity and potential biases on the platform's algorithmic recommendation patterns. Results suggest that while entropy remained stable, many transitions stayed within the same video genre.

Specifically, the 'Cartoons' genre played a significant role in this phenomenon. It exceeded the dominance threshold in the zero-shot classification transition matrix, highlighting a potential algorithmic bias towards this type of content. While not conclusive, it indicates potential filter bubbles that may negatively influence children's viewing experiences. Individual recommendation trees showcase that starting a viewing session from a certain genre barely affects the genres encountered in the rest of the session.

All in all, no strong indications of filter bubble formations or lack of diversity in recommendation chains were found.

7 ACKNOWLEDGEMENTS

I would like to thank M. Liberato and A. Affinito for their continued support as my supervisors throughout this project.

7.1 AI statement

During the writing of this paper, ChatGPT was used to modify several sentences to sound more academic. After using the tool, the author has reviewed and edited the content as needed and takes full responsibility for the content of the work.

REFERENCES

[1] Gediminas Adomavicius and YoungOk Kwon. 2011. Improving aggregate recommendation diversity using ranking-based techniques. *IEEE Transactions on Knowledge and Data Engineering* 24, 5 (2011), 896–911. https://ieeexplore.ieee.org/document/5680904.

[2] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of the web conference 2020*. 2155–2165. https://dl.acm.org/doi/abs/10.1145/3366423.3380281.

[3] Arda Antikacioglu and R Ravi. 2017. Post processing recommender systems for diversity. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 707–716. https://dl.acm.org/doi/10.1145/3097983.3098173.

[4] Mohd Asif and Saniya Kazi. 2024. Examining the influence of short videos on attention span and its relationship with academic performance. *International Journal of Science and Research* 13, 4 (2024), 1877–1883. https://www.researchgate.net/profile/Mohd-Asif-22/publication/380348721_Examining_the_Influence_of_Short_Videos_on_Attention_Span_and_its_Relationship_with_Academic_Performance/links/6637630608aa54017adba6df/Examining-the-Influence-of-Short-Videos-on-Attention-Span-and-its-Relationship-with-Academic-Performance.pdf.

[5] Ken Barker, Parul Awasthy, Jian Ni, and Radu Florian. 2021. IBM MNLP IE at CASE 2021 task 2: NLI reranking for zero-shot text classification. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*. 193–202. https://aclanthology.org/2021.case-1.24/.

[6] Megan A Brown, James Bisbee, Angela Lai, Richard Bonneau, Jonathan Nagler, and Joshua A Tucker. 2022. Echo chambers, rabbit holes, and algorithmic bias: How YouTube recommends content to real users. *Available at SSRN 4114905* (2022). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4114905.

[7] Lauren Valentino Bryant. 2020. The YouTube algorithm and the alt-right filter bubble. *Open Information Science* 4, 1 (2020), 85–90. https://www.degruyterbrill.com/document/doi/10.1515/opis-2020-0007/html.

[8] Google Developers. 2024. YouTube Data API v3. https://developers.google.com/youtube/v3 Accessed 2025-06-18.

[9] Muhammad Haroon, Anshuman Chhabra, Xin Liu, Prasant Mohapatra, Zubair Shafiq, and Magdalena Wojcieszak. 2022. Youtube, the great radicalizer? auditing and mitigating ideological biases in youtube recommendations. *arXiv preprint arXiv:2203.10666* (2022). https://arxiv.org/abs/2203.10666.

[10] Johannes Hofmeister. 2024. youtube-transcript-api: A Python interface for retrieving YouTube transcripts. https://pypi.org/project/youtube-transcript-api/. Accessed: 2025-06-21.

[11] Ikhfi Imaniah, Nurul Fitria Kumala Dewi, and Akhmad Zakky. 2020. YouTube kids channels in developing young children's communication skills in English: Parents' beliefs, attitudes, and behaviors. *Ijlecr-International Journal of Language Education and Culture Review* 6, 1 (2020), 20–30. https://www.academia.edu/download/63811913/IJLECR-Ikhfi__Nurul__Zakky20200702-80141-r8apg1.pdf.

[12] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. 2014. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 1725–1732. https://www.cv-foundation.org/openaccess/content_cvpr_2014/html/Karpathy_Large-scale_Video_Classification_2014_CVPR_paper.html.

[13] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. https://huggingface.co/facebook/bart-large-mnli.

[14] Microsoft. 2024. Playwright. https://playwright.dev/ Version 1.x.

[15] Dimitar Nikolov, Diego FM Oliveira, Alessandro Flammini, and Filippo Menczer. 2015. Measuring Online Filter Bubbles. *arXiv preprint arXiv:1502.07162* (2015). https://www.academia.edu/download/118646623/1502.pdf.

[16] Kostantinos Papadamou, Antonis Papasavva, Savvas Zannettou, Jeremy Blackburn, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Michael Sirivianos. 2020. Disturbed YouTube for kids: Characterizing and detecting inappropriate videos targeting young children. In *Proceedings of the international AAAI conference on web and social media*, Vol. 14. 522–533. https://ojs.aaai.org/index.php/ICWSM/article/view/7320.

[17] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. penguin UK. https://books.google.com/books?hl=en&lr=&id=-FWO0puw3nYC&oi=fnd&pg=PT24&dq=Eli+Pariser&ots=g6KsCstQUW&sig=GHn8MJzmKAeSwfJxs-d0wLwX7Hs.

[18] Python Software Foundation. 2023. Python Language Reference, version 3.11. https://www.python.org/. Accessed: 2025-06-21.

[19] Ishwarya Ramesh, Ishwarya Sivakumar, Kiruthya Ramesh, Vishnu Priya Prasanna Venkatesh, and V Vetriselvi. 2020. Categorization of YouTube videos by video sampling and keyword processing. In *2020 International conference on communication and signal processing (ICCSP)*. IEEE, 56–60. https://ieeexplore.ieee.org/abstract/document/9182158.

[20] Ali Raza, Faizan Younas, Hafeez Ur Rehman Siddiqui, Furqan Rustam, Monica Gracia Villar, Eduardo Silva Alvarado, and Imran Ashraf. 2024. An improved deep convolutional neural network-based YouTube video classification using textual features. *Heliyon* 10, 16 (2024). https://www.cell.com/heliyon/fulltext/S2405-8440(24)11843-9.

[21] Xiaolong Ren, Linyuan Lü, Runran Liu, and Jianlin Zhang. 2014. Avoiding congestion in recommender systems. *New Journal of Physics* 16, 6 (2014), 063057. https://iopscience.iop.org/article/10.1088/1367-2630/16/6/063057.

[22] Vicky Rideout. 2021. The Common Sense Census: Media Use by Kids Age Zero to Eight in America, A Common Sense Media Research Study,[United States], 2013, 2017. (2021). https://www.icpsr.umich.edu/web/ICPSR/studies/37491/variables.

[23] Camille Roth, Antoine Mazières, and Telmo Menezes. 2020. Tubes and bubbles topological confinement of YouTube recommendations. *PloS one* 15, 4 (2020), e0231703. https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0231703.

[24] Muhammad Azeem Sarwar, Dawood Ahmad, and Shahzad Ahmad. 2023. Exploring the influence of YouTube Kids app on children's cognitive skills. *Journal of Journalism, Media Science & Creative Arts* 3, 1 (2023), 117–136. https://www.academia.edu/download/108576076/Influence_of_Youtube_kids_app.pdf.

[25] Rashid Tahir, Faizan Ahmed, Hammas Saeed, Shiza Ali, Fareed Zaffar, and Christo Wilson. 2019. Bringing the kid back into youtube kids: Detecting inappropriate content on video streaming platforms. In *Proceedings of the 2019 IEEE/ACM international conference on advances in social networks analysis and mining*. 464–469. https://dl.acm.org/doi/abs/10.1145/3341161.3342913.

A FULL GENRE DISTRIBUTIONS

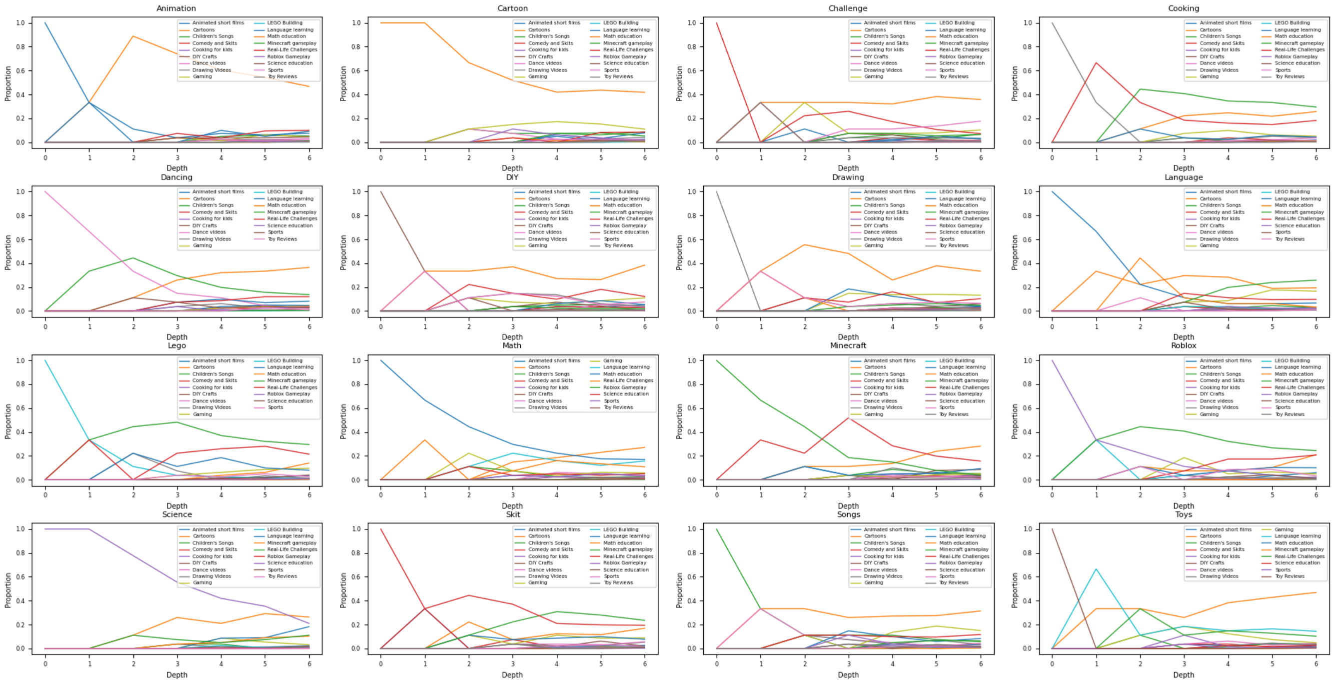

Fig. 8. All genre distribution graphs for different seed videos