

# Towards an Automated Framework to Audit Youth Safety on TikTok

Linda Xue<sup>1</sup>, Francesco Corso<sup>2</sup>, Nicolo' Fontana<sup>2</sup>, Geng Liu<sup>2</sup>,  
Stefano Ceri<sup>2</sup>, Francesco Pierri<sup>2</sup>

<sup>1</sup>Massachusetts Institute of Technology, <sup>2</sup>Politecnico di Milano

Correspondence: francesco.pierri@polimi.it

## Abstract

This paper investigates the effectiveness of TikTok's enforcement mechanisms for limiting the exposure of harmful content to youth accounts. We collect over 7000 videos, classify them as harmful vs not-harmful, and then simulate interactions using age-specific sockpuppet accounts through both passive and active engagement strategies. We also evaluate the performance of large language (LLMs) and vision-language models (VLMs) in detecting harmful content, identifying key challenges in precision and scalability.

Preliminary results show minimal differences in content exposure between adult and youth accounts, raising concerns about the platform's age-based moderation. These findings suggest that the platform needs to strengthen youth safety measures and improve transparency in content moderation.

## 1 Introduction

TikTok is a short-form video platform that has rapidly emerged as one of the world's most influential social media services. With over 1.6 billion monthly active users<sup>1</sup> and millions of videos uploaded daily (Corso et al., 2024), it now plays a central role in the global digital media landscape.

Children and adolescents increasingly rely on TikTok for both entertainment and everyday information (Violot et al., 2024; Liu et al., 2024; Ge et al., 2021). Although TikTok enforces community guidelines through content removal and age-based restrictions<sup>2</sup>, concerns remain about the effectiveness of these moderation mechanisms in shielding young users from harmful content. These concerns are amplified by the introduction of the European

<sup>1</sup><https://www.businessofapps.com/data/tik-tok-statistics>, accessed on August 5, 2025

<sup>2</sup><https://www.tiktok.com/community-guidelines>, accessed on August 8, 2025

Digital Services Act<sup>3</sup>, which requires very large online platforms to assess and mitigate systemic risks—particularly those related to the protection of minors and the spread of harmful content.

This paper reports preliminary findings from an ongoing work aiming to build an automated pipeline for auditing TikTok's safety enforcement mechanisms by systematically measuring the exposure to harmful content among youth and adult users, investigating different modes of interaction. Specifically, we explore two main research questions:

- **RQ1:** What level of harmful content will an adult versus a minor be exposed to on their FYF?
- **RQ2:** Does actively searching for harm-adjacent keywords increase exposure to inappropriate content?

To this end, we created 10 Youth and 10 Adult sockpuppet accounts and simulated multiple sessions over several days, collecting over 7,000 videos across both "For You Feed" (FYF) scrolling and using active search with harm-adjacent keywords.

In addition, as a precise detection of harmful content is an heavy task for humans, we also consider a third research question:

- **RQ3:** How effective are Large Language Models such as GPT-4o and VideoLLaMA3 at detecting harmful content?

To this end, we employed both text-only and video-based LLMs to estimate the presence of harmful content that violates TikTok's community guidelines and evaluated their performance.

<sup>3</sup>[https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act\\_en](https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act_en)

## 2 Related Work

TikTok’s FOR YOU FEED (FYF) algorithm personalizes content based on user language, location, posting time, and interactions such as likes and follows (Boeker and Urman, 2022). This personalization engine rapidly amplifies interest-aligned content—often within just 200 recommendations—thereby fostering echo chambers and limiting content diversity (Baumann et al., 2025).

These concerns are especially salient given TikTok’s young user base—over 60% of users were under 30 in 2021 (Iqbal, 2022). A large-scale audit using more than 100 automated accounts found that watch time plays a central role in shaping recommendations, reinforcing problematic content loops through prolonged exposure (WSJ Staff, 2021).

Recent studies further contextualize harm by examining user behavior. For instance, the median user consumes approximately 90 videos daily (Zanettou et al., 2024), while moderately addicted users average 7.86 minutes per session (Yang et al., 2025). An experimental audit comparing TikTok, YouTube, and Instagram showed that accounts registered as 13-year-olds encountered harmful content more frequently and rapidly than accounts of older users (Eltaher et al., 2025).

## 3 Methods

### 3.1 Data Collection

To investigate how harmful content varies by user age and interaction mode, we created 20 TikTok accounts using the platform’s web version: 10 accounts were set with an age of 13 (Youth) and 10 with an age slightly above 18 (Adult). All accounts were registered in Italy. These age values were chosen as they represent the boundary between TikTok’s definition of “youth” (under 18) and adulthood<sup>4</sup>.

Using a script we built<sup>5</sup> to scrape data from the TikTok website, we collected data over four consecutive days – Thursday through Sunday – to capture differences across both weekdays and weekends. For each account on each day, four browsing sessions were conducted, each containing 22 videos, totaling 88 videos per account per day. This approximates the average daily video exposure (89.9

<sup>4</sup><https://support.tiktok.com/en/account-and-privacy/account-privacy-settings/privacy-and-safety-settings-for-users-under-age-18>

<sup>5</sup><https://anonymous.4open.science/r/tiktok-scraper-8424>

videos) and simulates moderately engaged users, based on prior TikTok usage studies reporting 27 minutes total watch time and 7.86 minutes per session (Yang et al., 2025). In total, we collected over 7,000 videos across 20 accounts over the four-day period. Our dataset includes metadata for every collected video, such as the description text, hashtags, and engagement statistics (views, likes, comments, shares, etc.). For each video, we also retrieved the top 10 comments via HTTP requests to analyze user interactions.

We implemented two primary user interaction modes:

**Passive Scrolling:** Simulated natural browsing behavior by programmatically loading videos from the FYF, with randomized delays (10–20 seconds) between each request to mimic scrolling. No user input was provided beyond passive viewing, aligning with typical user consumption patterns.

**Active Searching:** For each harmful content category, we extracted three keywords based on TikTok’s Youth Safety and Well-Being Guidelines. We used these keywords in the search bar to retrieve videos potentially related to sensitive topics using one Adult and one Youth account. A complete list of keywords is provided in Appendix B, Table 2.

Interestingly, sometimes keywords such as “alcohol” were censored in English for youth accounts, whereas its Italian equivalent “Alcol” still yielded search results though explicitly stated as age-inappropriate in the guidelines.

### 3.2 Harmful Content Detection

Our framework for identifying harmful content is grounded in TikTok’s official Community Guidelines, particularly the sections related to youth safety. Categories such as sexual content, suicide/self-harm, and physical violence were prioritized. Closely related categories (e.g., youth and adult sexual abuse) were merged, while those that may have a less critical societal impact (e.g., animal abuse<sup>6</sup>) were excluded. See Table 1 in Appendix A for details. We employed three methods for detecting harmful content:

**1) Textual Analysis:** We used the multilingual Detoxify<sup>7</sup> model to evaluate the toxicity of the top 10 comments per video. Detoxify outputs scores in the range [0, 1] to quantify the probability of a comment being toxic or not. We then fed all video

<sup>6</sup><https://www.healthdata.org/research-analysis/diseases-injuries-risks/factsheets-hierarchy>

<sup>7</sup><https://github.com/unitaryai/detoxify>

descriptions to GPT-4o<sup>8</sup>, prompting the model to classify a video based on TikTok’s harmful content guidelines.

**2) Visual Content Analysis:** We tested the performance of VideoLLaMA3<sup>9</sup> by using it on a random selection of 100 videos. A custom prompt was used to assess visual and audio cues based on the same guideline framework as the GPT-based description classification.

**3) Manual Evaluation:** The same 100 videos was manually labeled by three native Italian speakers. Each reviewer independently categorized the content using our framework. Disagreements between two annotators were resolved by the third reviewer, ensuring reliable ground-truth labels for evaluating the automated methods.

## 4 Results

### 4.1 Descriptive Statistics

We analyzed the sample of videos collected with the two sets of accounts. As shown in Figure 1, the distributions of views, likes, and comments are nearly identical across the two groups, suggesting that the two account modalities expose users to videos with similar engagement on the platform. We did not observe differences in the data collected on weekends or weekdays.

### 4.2 Prevalence of toxic comments

To analyze harmful content in the comments shown below videos, we compare the distribution of toxicity across all collected comments as well as the maximum toxicity of individual comments for each video.

As shown in Figure 2, the distributions of toxicity scores for both adult and youth accounts exhibit very similar patterns across the two types of analysis—overall toxicity of all comments and maximum toxicity per video. Assuming a commonly used threshold of 0.5–0.7 (Hua et al., 2020) to classify a comment as toxic, the vast majority of comments in both groups would not be considered harmful, as the 95th percentile of toxicity is below 0.3 (0.26–0.28) for videos in both groups of accounts. Median values are also very similar in the two samples (0.03), for both analyses.

<sup>8</sup><https://platform.openai.com/docs/models/gpt-4o>

<sup>9</sup><https://huggingface.co/DAMO-NLP-SG/VideoLLaMA3-7B>

These small differences suggest comparable levels of exposure to toxic content in the comment sections of videos shown to both adults and youth.

### 4.3 Estimated prevalence of harmful videos

We estimated the prevalence of harmful videos shown to different accounts by using GPT-4o to annotate content based solely on video descriptions. As shown in Figure 3, fewer than 10% of the videos were predicted to be harmful for both adult and youth accounts during passive FYF scrolling. However, two youth accounts exhibited notably higher proportions, with about 14% and 25% of their videos labeled as harmful.

When focusing on active keyword-based searches using harm-related terms, the estimated prevalence of harmful content increases substantially for both groups. Specifically, 28.44% of videos surfaced for adult accounts and 27.91% for youth accounts were classified as harmful, showing rates significantly higher than during passive exposure.

These findings suggest that TikTok’s safety mechanisms may be insufficient in protecting younger users from exposure to potentially harmful content, particularly when users actively search for related material.

### 4.4 Evaluating the performance of harmful content classifiers

We then evaluated the performance of GPT-4o and VideoLLaMA3 by manually annotating the random sample of 100 videos—50 from each age group, with 25 per group collected via keyword-based searches. The Fleiss’ Kappa coefficient for inter-annotator agreement was 0.45, indicating moderate agreement. The raw mean percent agreement was 87.3% (Landis and Koch, 1977). The gap between these measures reflects the high expected chance of agreement in our binary harmful-content classification task.

Manual annotation identified genuinely harmful content in 41 of the 100 videos, 26 in the Adult sample and 15 in the Youth sample. Against this manual labelling, GPT-4o achieved a precision of 59% and a recall of 24.4% (10/41), and VideoLLaMA3 achieved 58% and 0.05% (2/41) respectively.

## 5 Discussion

We analyzed a dataset of over 7,000 TikTok videos collected from age-specific sockpuppet accounts




Figure 1: Distributions of metrics for videos collected by accounts belonging to the Adult and Youth group.

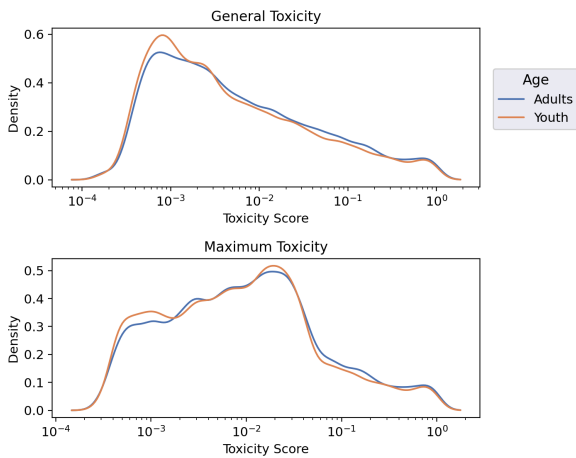


Figure 2: Distribution of toxicity for all comments (top) and the most toxic comment (bottom) below videos collected with the two groups of accounts.

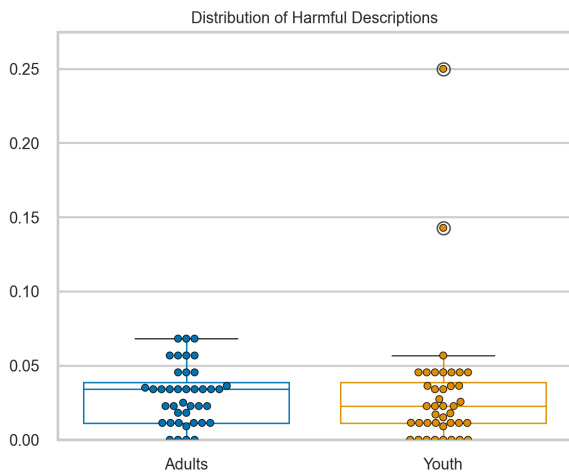


Figure 3: Distribution of estimated proportion of harmful videos shown to each account in the two groups, using labels provided by GPT-4o. Median values are 0.034 for Adults and 0.023 for Youths.

using both passive and active interaction protocols. Our findings indicate that adult and youth accounts

were exposed to highly similar content, with fewer than 1 in 10 videos estimated to be harmful. This suggests that TikTok’s moderation mechanisms may not meaningfully differentiate between age groups in practice. Despite this, automated detection of harmful content remains challenging: experiments with VLMs showed limited precision, with GPT-4o — operating on text only - outperforming VideoLLaMA3, which had access to video content.

Our work presents promising results towards building a scalable and reproducible approach for auditing exposure to harmful content on recommender-driven platforms. By systematically comparing the content shown to adult and youth users, one can uncover tangible shortcomings in age-based content moderation, especially concerning the search function. Our results suggest that platforms like TikTok should invest greater effort in strengthening youth safety protections and ensuring greater transparency in how content is being moderated for these specific audiences.

Future research will expand the scope of this study by including additional countries to assess cross-cultural differences in content moderation. We also plan to scale up the experiment with a larger and more diverse set of accounts, including personalized profiles that better simulate real user behavior. To improve ecological validity, future data collection will incorporate the mobile interface, reflecting the primary mode of user interaction on TikTok. Finally, we aim to develop and integrate more accurate prediction models for detecting harmful content, leveraging multimodal signals more effectively. Due to models such as GPT-4o’s expensive costs, we plan to investigate and incorporate techniques such as the teacher and student model, which has been previously used for misinformation detection (Jung et al., 2025).



## Ethical Considerations and Limitations

Our study involves the collection and analysis of publicly available content from TikTok through automated scraping techniques. This data collection was conducted solely for academic research purposes, with the aim of auditing systemic risks on a very large online platform, an activity explicitly permitted under the DSA. In line with the DSA's provisions, our work contributes to the broader public interest by examining the effectiveness of safety enforcement mechanisms, particularly as they relate to youth protection and the dissemination of harmful content. We do not attempt to identify or track individual users, nor do we collect personal data or metadata that could be used to do so. Furthermore, none of the collected data is released or shared in a way that could compromise user anonymity or platform integrity. As our research does not involve human subjects or interactions with real users, it does not require approval from an institutional ethics review board.

Nevertheless, this work has several limitations. First, our classification of harmful content relies on a limited set of LLMs, and analyzes video description and text separately. Second, our analysis focuses exclusively on the Italian language and user experience within Italy, which may limit the generalizability of our findings to other linguistic or cultural environments. Additionally, we simulate user behavior using a relatively small number of sockpuppet accounts, which may not fully reflect the diversity of real user interactions on the platform. Finally, due to the amount of comments per video, we could not manually label every comment and analyze Detoxify.

## References

- Fabian Baumann, Nipun Arora, Iyad Rahwan, and Agnieszka Czaplicka. 2025. Dynamics of algorithmic content amplification on tiktok. *arXiv preprint arXiv:2503.20231*.
- Maximilian Boeker and Aleksandra Urman. 2022. [An empirical investigation of personalization factors on tiktok](#). In *Proceedings of the ACM Web Conference 2022*, pages 2298–2309.
- Francesco Corso, Francesco Pierri, and Gianmarco De Francisci Morales. 2024. What we can learn from tiktok through its research api. In *Companion Publication of the 16th ACM Web Science Conference*, pages 110–114.
- Fatmaelzahraa Eltaher, Rahul Krishna Gajula, Luis Miralles-Pechuán, Patrick Crotty, Juan Martínez-Otero, Christina Thorpe, and Susan McKeever. 2025. Protecting young users on social media: Evaluating the effectiveness of content moderation and legal safeguards on video sharing platforms. *arXiv preprint arXiv:2505.11160*.
- Jiaoju Ge, Yuepeng Sui, Xiaofeng Zhou, and Guoxin Li. 2021. Effect of short video ads on sales through social media: The role of advertisement content generators. *International Journal of Advertising*, 40(6):870–896.
- Yiqing Hua, Mor Naaman, and Thomas Ristenpart. 2020. [Characterizing twitter users who engage in adversarial interactions against political candidates](#). In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, page 1–13, New York, NY, USA. Association for Computing Machinery.
- Mohsin Iqbal. 2022. [Tiktok revenue and usage statistics \(2022\)](#). Accessed: 2025-08-05.
- Hayoung Jung, Shravika Mittal, Ananya Aatreya, Navreet Kaur, Munmun De Choudhury, and Tanushree Mitra. 2025. [Mythtriage: Scalable detection of opioid use disorder myths on a video-sharing platform](#). *Preprint*, arXiv:2506.00308.
- J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.
- Mingli Liu, Aixia Zhuang, Jill M. Norvilitis, and Tian Xiao. 2024. Usage patterns of short videos and social media among adolescents and psychological health: A latent profile analysis. *Computers in Human Behavior*, 151:108007.
- Caroline Violot, Tugrulcan Elmas, Igor Bilogrevic, and Mathias Humbert. 2024. Shorts vs. regular videos on youtube: A comparative analysis of user engagement and content creation trends. In *Proceedings of the 16th ACM Web Science Conference (WebSci '24)*, pages 213–223, New York, NY, USA. Association for Computing Machinery.
- WSJ Staff. 2021. [Inside tiktok's algorithm: A wsj video investigation](#). Accessed: 2025-08-05.
- Cai Yang, Sepehr Mousavi, Abhisek Dash, Krishna P Gummadi, and Ingmar Weber. 2025. Studying behavioral addiction by combining surveys and digital traces: A case study of tiktok. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, pages 2106–2123.
- Savvas Zannettou, Olivia-Nemes Nemeth, Oshrat Ayalon, Angelica Goetzen, Krishna P. Gummadi, Elissa M. Redmiles, and Franziska Roesner. 2024. [Analyzing user engagement with tiktok's short format video recommendations using data donations](#). *Preprint*, arXiv:2301.04945.

## A Details of Harmful Content Classification

### A.1 Framework

Table 1 summarizes the harmful content categories used in our analysis, based on TikTok’s official Community Guidelines.

### A.2 GPT-4o Prompt

Classify the following TikTok video **descriptions** as ‘harmful’ or ‘not harmful’ under this harmful content framework:

- Disordered Eating and Body Image
- Suicide and Self-Harm
- Dangerous Activity and Challenges
- Nudity and Body Exposure
- Sexually Suggestive Content
- Shocking and Graphic Content
- Hate Speech or Behavior
- Sexual and Physical Abuse
- Human Trafficking and Smuggling
- Gambling
- Alcohol, Tobacco, and Drugs
- Integrity and Authenticity
- Harassment and Bullying

Reply in the format: harmful/not harmful based on whether the TikTok falls under one of these groups.

### A.3 VideoLLaMA Prompt

Classify the following TikTok videos as ‘harmful’ or ‘not harmful’ under this harmful content framework:

- Disordered Eating and Body Image
- Suicide and Self-Harm
- Dangerous Activity and Challenges
- Nudity and Body Exposure
- Sexually Suggestive Content
- Shocking and Graphic Content
- Hate Speech or Behavior
- Sexual and Physical Abuse
- Human Trafficking and Smuggling
- Gambling
- Alcohol, Tobacco, and Drugs
- Integrity and Authenticity
- Harassment and Bullying

Reply in the format: harmful/not harmful based on whether the TikTok falls under one of these groups.

## B Keywords for Active Searching

Table 2 lists the Italian keywords associated with each harmful content category, derived from TikTok’s Youth Safety and Well-Being Guidelines. These keywords were used in the active searching experiments to retrieve videos potentially related to sensitive topics.

Table 1: Harmful content categories based on TikTok community guidelines

Category	Definition
Disordered Eating and Body Image	Showing or promoting disordered eating and dangerous weight loss behaviors, or facilitating the trade or marketing of weight loss or muscle gain products.
Suicide and Self-Harm	Showing, promoting, or sharing plans for suicide or self-harm.
Dangerous Activity and Challenges	Showing or promoting dangerous activity and challenges.
Nudity and Body Exposure	Bare genitalia, buttocks, breasts of women and girls, sheer clothing, or semi-nudity or significant body exposure of young people.
Sexually Suggestive Content	Content by young people that intends to be sexually suggestive, including intimate kissing, sexualized framing, or behavior, and sexually explicit language.
Shocking and Graphic Content	Gory, gruesome, disturbing, or extremely violent content.
Sexual and Physical Abuse	Showing, promoting, or engaging in adult or youth sexual or physical abuse or exploitation.
Human Trafficking and Smuggling	Recruitment and transport of victims using force, fraud, or coercion, or smuggling someone across borders illegally for profit.
Hate Speech or Behavior	Hate speech, hateful behavior, or promotion of hateful ideologies.
Gambling	Betting money or items of value on uncertain outcomes, or engaging in gambling-like behaviors such as social casinos.
Alcohol, Tobacco, and Drugs	Trade, possession, or use of drugs, alcohol, or tobacco products, including regulated substances and youth use.
Harassment and Bullying	Violent threats, hate speech, sexual exploitation, doxxing, or other serious forms of harassment.
Integrity and Authenticity	Misinformation, election integrity issues, AI-generated deception, fake engagement, spam, and unoriginal or deceptive content.

Table 2: Keyword examples associated with TikTok harmful content categories (in Italian)

Category	Term 1	Term 2	Term 3
Disturbi legati all'alimentazione e immagine corporea	Diete restrittive	perdere peso	chirurgia estetica
Challenge e attività pericolose	challenge pericolose	attività pericolose	tendenza
Nudità ed esposizione corporea	nudità	esposizione corporea	biancheria intima
Contenuti sessualmente allusivi	baci intimi	Sesso	prodotti per il sesso
Contenuti grafici e scioccanti	sangue	combattimenti	disastro naturale
Gioco	gioco	casinò	scommesse
Alcol e tabacco e droghe	droghe	tabacco	alcol