

Technical Disclosure Commons

Defensive Publications Series

April 2020

Filtering Videos Based On Protected Topics

N/A

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

N/A, "Filtering Videos Based On Protected Topics", Technical Disclosure Commons, (April 22, 2020)
https://www.tdcommons.org/dpubs_series/3176



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

FILTERING VIDEOS BASED ON PROTECTED TOPICS

ABSTRACT

Disclosed herein is an improved mechanism for filtering videos based on protected topics. The mechanism can begin by identifying a particular portion of search queries for videos (e.g., the top X percent of search queries for videos). The mechanism can then determine a group of topics to be protected based on the identified search queries. The mechanism can, for each protected topic in the group of topics, identify a subset of channels that provide videos that are associated with the protected topic (e.g., the top N channels that provide videos associated with the protected topic). The mechanism can, for each channel in the identified subset of channels, calculate a channel trust score (e.g., based on channel ranking among channel results for search queries for a topic and based on a number of policy violations associated with the channel). The mechanism can then identify channels associated with a channel trust score greater than a channel trust score threshold. Videos from these channels having a channel trust score greater than a channel trust score threshold can be provided for recommendation (e.g., in response to search queries). For example, a video content sharing service can provide videos associated with a trusted topic to be obtained from one or more trusted channels.

BACKGROUND

Video content sharing services receive uploaded videos from many different users and entities. Some video content sharing services may provide videos for children. These services may occasionally receive uploaded videos that are inappropriate for children, for example, because they include inappropriate content or language. Identifying videos that are inappropriate for children can be difficult. For example, a content uploader may obscure content of an

uploaded video by giving the uploaded video a title that is likely to be associated with a topic searched for by children. As a more particular example, a content uploader may upload a video that includes inappropriate content but that has a title associated with a popular children's movie.

Thus, there is a need for an improved mechanism for filtering videos based on protected topics.

DESCRIPTION

The systems and techniques described in this disclosure relate to filtering videos based on protected topics. The system can be implemented on a server, such as a server associated with a video content sharing service. Note that, in some instances, a video content sharing service can be a video content sharing service that hosts and provides videos for viewing by children. FIG. 1 illustrates an example process for filtering videos based on protected topics.

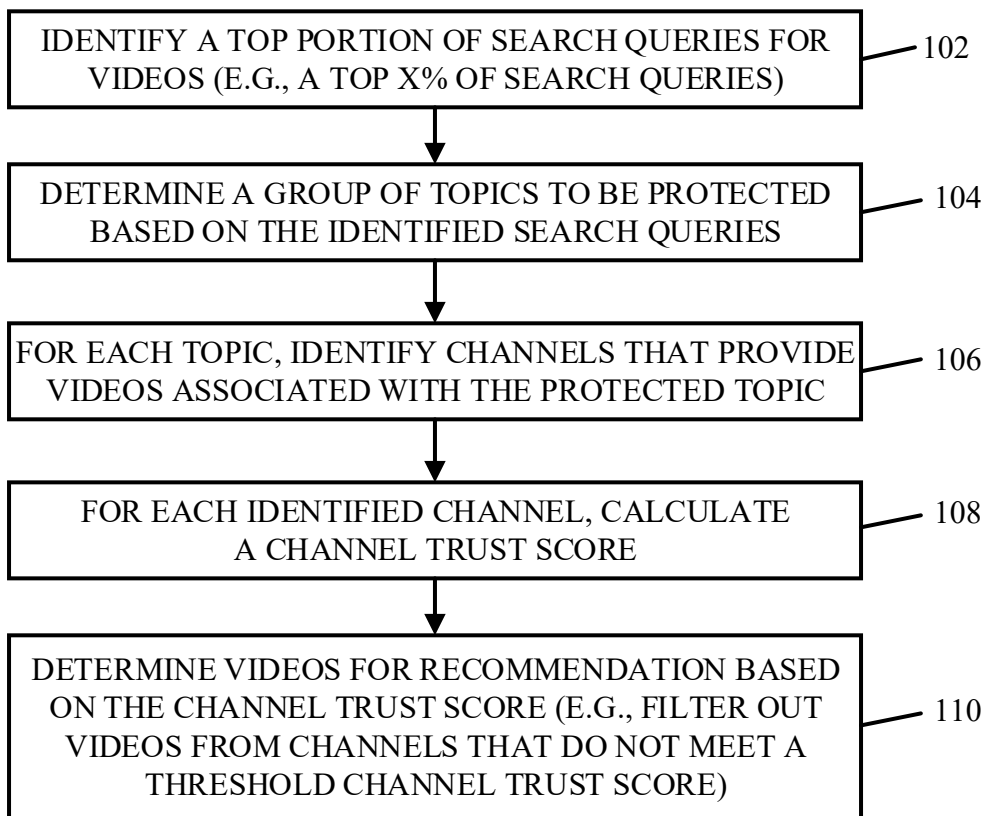


FIG. 1

At 102, the system can begin by identifying a portion of search queries for videos on a video content sharing service. This portion of search queries can include, for example, identifying a top X percent of search queries for videos on a video content sharing service (e.g., 5%, 10%, 20%, and/or any other suitable percentage).

It should be noted that, in some instances, the top X percent of search queries can be identified based on any suitable information or criteria. For example, in some instances, the top X percent of search queries can correspond to the most common search queries that are received by the video content sharing service. As another example, in some instances, the top X percent of search queries can correspond to search queries that resulted in the most selected videos hosted by the video content sharing service. Note that, in some instances, the system can identify a top X percent of search queries for a particular geographic location (e.g., a particular country, a particular city, and/or any other suitable location) and/or received in a particular language.

At 104, the system can identify a group of topics to be protected based on the identified search queries. In some instances, the system can identify topics in the group of topics based on any suitable technique or combination of techniques. For example, in some instances, the system can identify one or more topics corresponding to each identified search query based on keywords included in the search query. As a more particular example, in some instances, the system can use any suitable database or knowledge graph that maps keywords to topics. As a specific example, the system can determine that a first query that includes a first character of a movie and a second query that includes a second character of the movie are both related to the topic of the movie name.

At 106, the system can, for each topic in the group of topics, identify the top N channels that provides one or more videos associated with the protected topic (e.g., the top five channels, the top ten channels, the top twenty channels, and/or any other suitable number of channels). In some instances, the system can identify the top N channels in any suitable manner. For example, in some instances, the system can identify a group of channels that are associated with the topic using any suitable database or knowledge graph. The system can then identify the top N channels of the group of channels based on any suitable criteria. For example, the top N channels can be the channels that are associated with the most popular or most watched videos associated with the topic.

At 108, the system can, for each channel, calculate a channel trust score. In some instances, the system can calculate the channel trust score based on any suitable information and using any suitable technique(s). For example, in some instances, a channel trust score can be based on a rank of the channel that indicates a popularity of the channel. As another example, in some instances, a channel trust score can be based on a number of times content items provided by the channel have violated terms of service associated with the video content sharing service. A more particular example for calculating a channel trust score can be defined as:

$$\text{channel trust score} = 1/(R * PV),$$

where R is the rank of the channel (e.g., where 1 is the top channel associated with a particular topic), and where PV is a number of policy violations corresponding to the channel. As a specific example, a channel that is the top-ranked channel and with 0 or 1 policy violations can have a channel trust score of 1. As another specific example, a channel that is the tenth ranked channel and with ten policy violations can have a channel trust score of 0.01.

At 110, the system can determine videos for recommendation based on the calculated channel trust scores associated with each of the identified channels.

For example, for each video in a corpus of videos associated with a protected topic, the system can identify videos provided by channels associated with a channel trust score greater than a channel trust score threshold. In continuing this example, the system can identify a corpus of videos where each video in the corpus is associated with a protected topic from the group of topics identified at 104. In some instances, the videos in the corpus of videos can be identified based on the database or knowledge graph that maps keywords to topics. For example, the system can identify videos with titles or other metadata related to particular topics. Such videos from these channels having a channel trust score greater than a channel trust score threshold can be provided for recommendation. In a more particular example, in response to a search query on a trusted topic, recommended videos on the trusted topic can be obtained from one or more trusted channels (e.g., one or more channels having a channel trust score greater than a channel trust score threshold).

This can, for example, filter out videos belonging to one or more channels having a channel trust score that does not meet a channel trust score threshold from being included as a recommended video. This can also, for example, protect the topics that users are most likely to view by relying on channel trust or channel reputation (e.g., as opposed to performing a detection approach on each and every video and/or media content item associated with a content sharing service).

In another example, for each video in a corpus of videos associated with a protected topic, the system can identify videos belonging to a channel with a channel trust score below a predetermined threshold. In continuing this example, the system can identify a corpus of videos

where each video in the corpus is associated with a protected topic from the group of topics identified at 104. In some instances, the videos in the corpus of videos can be identified based on the database or knowledge graph that maps keywords to topics. For example, the system can identify videos with titles or other metadata related to particular topics. The system can then determine whether a channel that provides a video in the corpus of videos is associated with a channel trust score below any suitable predetermined threshold (e.g., below a channel trust score threshold of 0.5, below a channel trust score threshold of 0.75, and/or below any other suitable threshold).

In some instances, the system can take any suitable action after identifying videos in the corpus of videos associated with channels with a channel trust score below the predetermined threshold. For example, in some instances, the system can forward an indication of the videos to a human evaluator with a message that the videos are to be manually examined a potentially including inappropriate content. As another example, in some instances, the system can block the videos from being returned in response to any received search query. As yet another example, in some instances, the system can block the video from being presented to a user if the user is below a predetermined age (e.g., less than 18, and/or any other suitable age).

Accordingly, an improved mechanism for filtering videos based on protected topics is provided.