January 2022

Displaying Promotional Content Based on Elastic Frequency Caps

Tanmay Khirwadkar

Wenxuan Dai

Sanjay Mangla

DISPLAYING PROMOTIONAL CONTENT BASED ON ELASTIC FREQUENCY CAPS

ABSTRACT

Disclosed herein is an improved mechanism for displaying promotional content based on elastic frequency caps. The mechanism can use a machine learning model to perform an elastic frequency capping decision by predicting whether the marginal benefit of showing promotional content to a user is higher than the marginal cost of showing the promotional content to the user by a particular threshold value. The mechanism can, for example, determine the marginal benefit of showing the promotional content to the user by predicting the clicks or conversions of such promotional content. The mechanism can also, for example, determine the marginal cost of showing the promotional content to the user by predicting the reduction in watch time of content due to showing the promotional content.

BACKGROUND

Users frequently use websites associated with particular services or applications associated with particular services to view media content (e.g., stream or download media content, such as music, television shows, movies, videos, etc.), connect with other users (e.g., via social networking services), and/or perform tasks (e.g., purchase items, edit documents, receive or send messages, etc.). In many cases, a service may want to notify a user of different features available through the service. For example, a service that provides media content may want to notify a user about a premium option through which media content may be downloaded rather than streamed. As another example, a service may want to notify a user about a premium option through which media content may be viewed without advertisements or with fewer advertisements. However, it can be difficult to maximize clicks or conversions of such

promotional content while minimizing the cost (e.g., the adverse impact on the user experience). For example, in many systems, cost is controlled through the use of frequency caps – that is, showing an impression to a user when the particular number of promotions shown within a particular time period is less than a frequency cap.

Accordingly, it is desirable to provide new methods, systems, and media for displaying promotional content based on elastic frequency caps.

DESCRIPTION

The systems and techniques described in this disclosure relate to displaying promotional content based on elastic frequency caps. In particular, the mechanisms described herein can determine whether promotional content associated with a service is to be shown to a user at a current time. In some embodiments, a service can include any suitable type of service used by a user, such as a media content streaming service, a social networking service, a content storage and/or sharing service, and/or any other suitable type of service. In some embodiments, the promotional content can be any suitable promotion associated with the service. For example, in some embodiments, the promotion can include an advertisement of a subscription, a product, and/or a feature provided by the service. As a more particular example, in some embodiments, in an instance in which the service is a media content streaming service, the promotion can include a subscription which, when purchased by the user, allows the user to stream media content with fewer or no advertisements. As another more particular example, in some embodiments, in an instance in which the service is a media content streaming service, the promotion can include an indication of an application provided by the service for playback of a particular type of media content provided by the service, such as music.

In some embodiments, the mechanisms described herein can determine whether promotional content associated with a service is to be shown to a user based on an elastic frequency cap in any suitable manner. For example, in some embodiments, the mechanisms can use a machine learning model to perform an elastic frequency capping decision by predicting whether the marginal benefit of showing promotional content to the user is higher than the marginal cost of showing the promotional content to the user by a particular threshold value. In continuing this example, the machine learning model can determine the marginal benefit of showing the promotional content to the user by predicting the clicks or conversions of such promotional content. Additionally, the machine learning model can determine the marginal cost of showing the promotional content to the user by predicting the reduction in watch time of content (e.g., video content) due to showing the promotional content.

FIG. 1 is an illustrative example of a method for displaying promotional content based on elastic frequency caps. This method can be implemented on a server, such as a server that manages uploaded media content items and provides media content items to user devices.

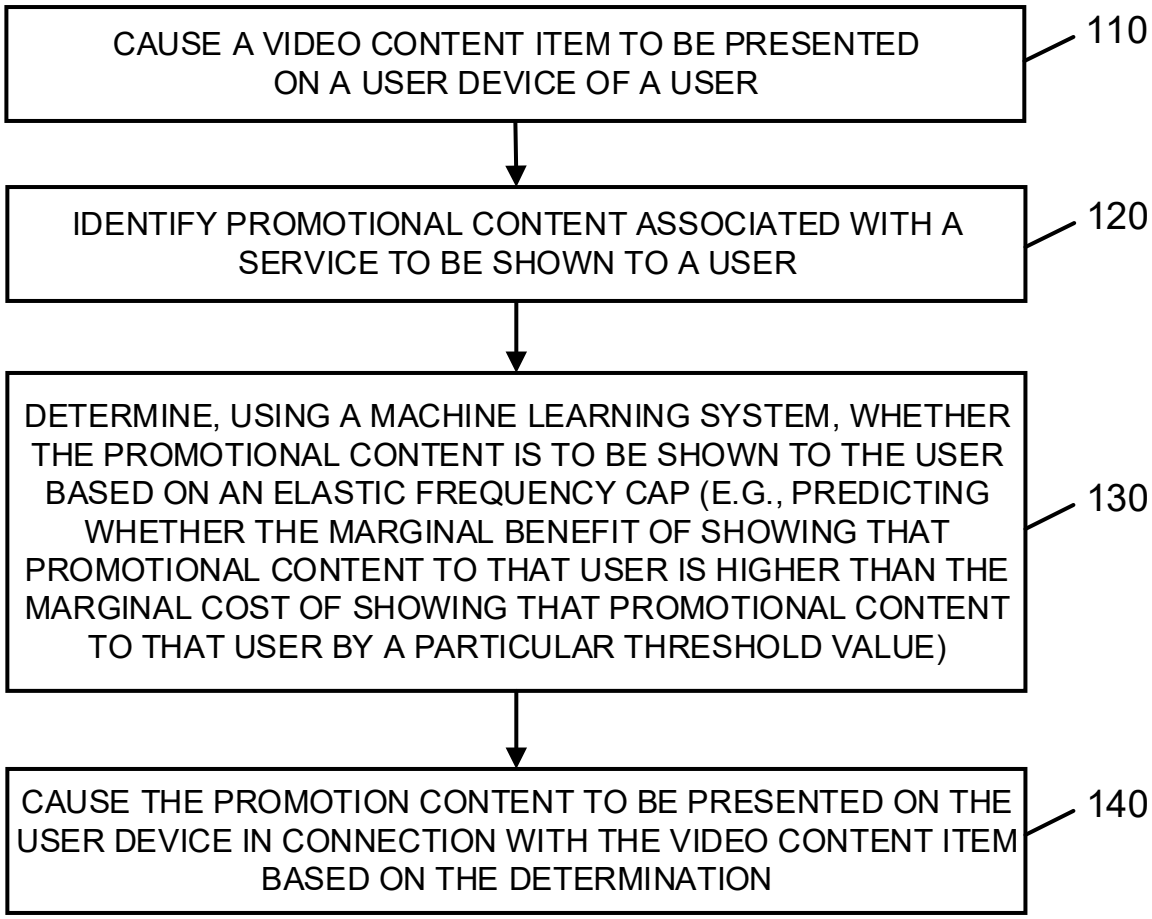

FIG. 1

At 110, the server can cause a video content item to be presented on a user device of a user. In some embodiments, the video content item can be any suitable type of video content item (e.g., a video, a music video, a television show, live-streamed video content, a video content item in a playlist of video content items, and/or any other suitable type of video content item). In some embodiments, the user device can be any suitable type of user device (e.g., a mobile phone, a tablet computer, a virtual assistant device, a television, a smart television, a streaming media device connected to a display device, a wearable computer, and/or any other suitable type of user device). In some embodiments, the server can cause the video content item to be presented on the user device in any suitable manner. For example, in some embodiments, in an instance in which

the video content item is streamed from the server, the server can transmit video data and audio data that corresponds to the video content item. In some embodiments, the server can begin causing the video content item to be presented on the user device in response to receiving a request from the user device.

At 120, the server can identify promotional content to be shown to a user. In some embodiments, the promotional content can be associated with a particular service currently being accessed by the user. For example, in some embodiments, the promotional content can be associated with a media content sharing service or a social networking service that is currently being accessed by the user. Note that, in some embodiments, the service can be accessed by the user on a user device of the user in any suitable manner, such as via a website associated with the service, via an application associated with the service that is executing on the user device, and/or in any other suitable manner.

In some embodiments, the promotional content can include any suitable content. For example, in some embodiments, the promotional content can include an indication of an available feature associated with the service. For example, in some embodiments, the promotional content can indicate that a paid feature associated with the service is available for purchase or subscription (e.g., a feature that allows a user to view media content with no or fewer advertisements, a feature that allows a user to download media content to the user device for offline viewing, and/or any other suitable feature). As another example, in some embodiments, the promotional content can include any suitable selectable user interface controls, such as a push button that, in response to being selected, causes a page to be presented that presents more information related to a feature corresponding to the promotional content, a push button that, in

response to being selected, dismisses the promotional content, and/or any other suitable user interface controls.

In some embodiments, the server can identify the promotional content in any suitable manner. For example, in some embodiments, the server can identify promotional content that indicates a feature available in connection with the service that the user has not yet purchased or activated. For example, in an instance in which the promotional content indicates a feature that allows media content to be downloaded directly to a user device, the server can determine whether the user has purchased or subscribed to the feature. Continuing with this example, in some embodiments, in response to determining that the user has not yet purchased or activated a particular feature, the server can identify promotional content corresponding to the feature. As another example, in some embodiments, the server can identify promotional content relating to a relatively new feature (e.g., released within the previous week, released within the previous month, and/or any other suitable relatively new feature) available in connection with the service.

Note that, in some embodiments, the user can be performing any suitable tasks or actions in connection with the service. For example, in an instance in which the service is a media content sharing service or a social networking service, the user can be browsing pages associated with different content creators, viewing media content items (e.g., streaming media content items, selecting media content items for download to the user device, and/or viewing media content items in any other suitable manner), and/or performing any other suitable actions. In some embodiments, the user can be performing any suitable tasks or actions via an application associated with the service that is executing on the user device. Additionally or alternatively, in some embodiments, the user can be performing any suitable tasks or actions via a website

associated with the service that is presented via a browser application executing on the user device.

At 130, the server can determine, using a machine learning model, whether the identified promotional content is to be shown to a user based on an elastic frequency cap in any suitable manner. For example, the machine learning model can determine the marginal benefit of showing the identified promotional content to the user by predicting the clicks or conversions of such promotional content and the machine learning model can determine the marginal cost of showing the identified promotional content to the user by predicting the reduction in watch time of content (e.g., video content) due to showing the promotional content.

In a more particular example, in some embodiments, the machine learning model can determine occupancy in a certain time period after a decision point or a point in which the server makes a decision to show the identified promotional content or to not show the identified promotional content. The cost can, for example, be calculated by taking the difference between the predicted occupancy of a decision to show the identified promotional content and the predicted occupancy of a decision to not show the identified promotion content. This can be represented as follows:

Label: occupancy in x hours after the decision point.

$$\text{Occupancy} = \text{WT} / \text{GAP}$$

$$\text{pCost} = 1 - \text{pOccupancy(features, show)} / \text{pOccupancy(features, no\_show)}$$

WT: watch time within the next X hours.

GAP: time window length, X.

pOccupancy(features, show): predicted occupancy if we show an impression.

pOccupancy(features, no_show): predicted occupancy if we don't show an impression.

In another more particular example, in some embodiments, the machine learning model can determine occupancy in a time gap between two consecutive decision points. This can, for example, reduce the impact of other promotional content or other impressions during the measured time gap.

In continuing these examples, the machine learning model can then perform an elastic frequency capping decision by predicting whether the marginal benefit of showing the identified promotional content to the user is higher than the marginal cost of showing the identified promotional content to the user by a particular threshold value.

For example, the machine learning model can use a thresholding algorithm for elastic frequency capping. The threshold, z, can be defined as the ratio of recent promotional content loading to visit frequency, which can be viewed as one approach for determining user experience cost. That is, if a user has seen more impressions or promotional content or if a user visits a service less frequently, the cost of showing the user one more impression or one more piece of promotional content can be higher. As such, the threshold, z, can be increased. This can be represented as follows:

$$\frac{pCPI - \mu(pCPI)}{\sigma(pCPI)} > z * \alpha + \beta$$

where

pCPI = predicted conversions per impression

μ, σ = average and standard deviation of pCPI computed over all impressions for the allocation group of the promo

α, β = scaling factor of z score

z = days_impression_show_last_14_days / days_active_last_14_days

By incorporating the thresholding algorithm for elastic frequency capping with the cost model based on occupancy described above, the machine learning model can determine whether:

$$\frac{pCPI - \mu(pCPI)}{\sigma(pCPI)} > pCost * \alpha + \beta$$

$$pCost = pOccupancy(features, no\ show) / pOccupancy(features, show) - 1$$

Based on the output from the machine learning model that predicts whether the marginal benefit of showing the identified promotional content to the user is higher than the marginal cost of showing the identified promotional content to the user by a particular threshold value, an elastic frequency capping decision on whether to show the identified promotional content can be made.

If, at 130, the server determines that the identified promotional content is to be presented to the user based on the elastic frequency capping decision ("yes" at 130), the server can, at 140, cause the identified promotional content to be presented on the user device at the current time. In some embodiments, the identified promotional content can be presented in connection with the video content item in any suitable manner. For example, in some embodiments, the identified promotional content can be presented as an overlay on the video content item. As another example, in some embodiments, the identified promotional content can be presented within a pop-up window that is presented in response to determining that presentation of the video content item on the user device has finished. As yet another example, in some embodiments, the identified promotional content can be presented as an advertisement inserted into the presentation of the video content item.

Additionally or alternatively, the server can also use the machine learning model to determine whether the identified promotional content should be presented to the user at a later time. For example, the machine learning model can incorporate a timing model with the models

described above. In a more particular example, if, at 130, the server determines that the identified promotional content is not to be presented to the user based on the elastic frequency capping decision (e.g., if an impression or other suitable promotional content has a high cost to be shown to the user at the current time), the server can use the machine learning model to determine whether to defer the promotional content to be shown at a later time.

Accordingly, a mechanism for displaying promotional content based on elastic frequency caps is provided.