



US011194879B2

(12) **United States Patent**  
**Sardar et al.**

(10) **Patent No.:** **US 11,194,879 B2**  
(45) **Date of Patent:** **Dec. 7, 2021**

(54) **CUSTOM COMPILATION VIDEOS**

(71) Applicant: **Valve Corporation**, Bellevue, WA (US)

(72) Inventors: **Adil Sardar**, Renton, WA (US);  
**Anthony John Cox**, Seattle, WA (US);  
**Mark Zbikowski**, Medina, WA (US);  
**Christian Carollo**, Seattle, WA (US);  
**Martin Otten**, Bellevue, WA (US);  
**Taylor Sherman**, Seattle, WA (US);  
**Alden Kroll**, Seattle, WA (US); **Donald**  
**Ichiro Lambe**, Watertown, MA (US)

(73) Assignee: **Valve Corporation**, Bellevue, WA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 219 days.

(21) Appl. No.: **16/657,819**

(22) Filed: **Oct. 18, 2019**

(65) **Prior Publication Data**

US 2021/0011939 A1 Jan. 14, 2021

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 16/505,112, filed on Jul. 8, 2019.

(51) **Int. Cl.**  
**G06F 16/9538** (2019.01)  
**G06F 16/9535** (2019.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 16/9538** (2019.01); **G06F 16/9535** (2019.01)

(58) **Field of Classification Search**  
CPC ..... G06F 16/9535; G06Q 30/0255  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

9,762,629 B1 9/2017 Bhargava et al.  
2005/0089835 A1 4/2005 Soldavini et al.  
2005/0261062 A1\* 11/2005 Lewin ..... A63F 13/358  
463/42  
2006/0282304 A1\* 12/2006 Bedard ..... G06Q 30/0201  
705/7.33

(Continued)

OTHER PUBLICATIONS

Gong, Jin, Yizhou Ye, and Kostas Stefanidis. "A Hybrid Recommender System for Steam Games." Information Search, Integration, and Personalization: 13th International Workshop, ISIP 2019, Heraklion, Greece, May 9-10, 2019, Revised Selected Papers. Springer, 2020. (Year: 2020).\*

(Continued)

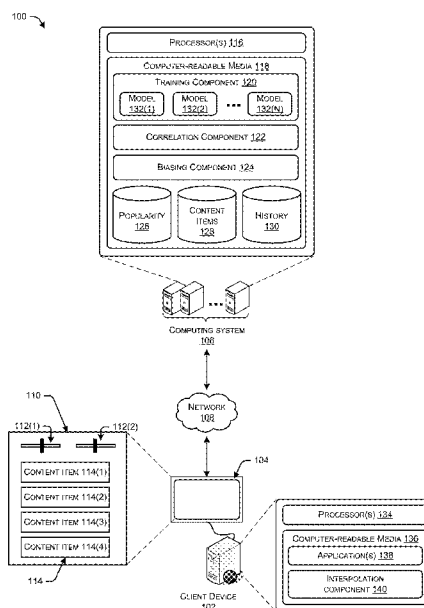
*Primary Examiner* — William Spieler

(74) *Attorney, Agent, or Firm* — Lee & Hayes, P.C.

(57) **ABSTRACT**

Described herein are, among other things, techniques, devices, and systems for generating one or more trained machine-learning models. Also described herein are techniques, devices, and systems for applying a consumption history of a particular user to the trained model(s) to generate score data indicating a correlation between each content-item title and the consumption history. The techniques then determine a ranked list of content items having a highest correlation to the consumption history, which may be used to retrieve videos associated with the most-correlated content items for generating a compilation video composed of these retrieved videos.

**20 Claims, 10 Drawing Sheets**



(56)

References Cited

U.S. PATENT DOCUMENTS

2007/0244742 A1 \* 10/2007 Short ..... G06Q 30/0201  
705/7.28  
2011/0054980 A1 \* 3/2011 Simons ..... G06Q 30/02  
705/7.37  
2011/0302103 A1 12/2011 Carmel et al.  
2012/0023095 A1 \* 1/2012 Wadycki ..... G06F 16/24578  
707/723  
2012/0060077 A1 3/2012 Mate et al.  
2013/0080224 A1 \* 3/2013 O'Brien ..... G06Q 10/10  
705/14.14  
2014/0179439 A1 6/2014 Miura et al.  
2015/0099589 A1 4/2015 Smith  
2015/0269810 A1 \* 9/2015 Wolf ..... G07F 17/3225  
463/25  
2015/0365729 A1 12/2015 Kaya et al.  
2016/0256775 A1 9/2016 Gustafson et al.  
2018/0063276 A1 3/2018 Foged  
2018/0130075 A1 \* 5/2018 Roy ..... H04N 21/6543  
2018/0150894 A1 5/2018 Srivastava et al.  
2021/0011958 A1 1/2021 Cox et al.

OTHER PUBLICATIONS

Gabmeier, Jorg. "Recommender Systems in the Domain of Video Games." 2016. Graz University of Technology, Master's Thesis. (Year: 2016).\*

Anwar, Syed Muhammad, et al. "A game recommender system using collaborative filtering (GAMBIT)." 2017 14th International Bhurban Conference on Applied Sciences and Technology (IBCAST). IEEE, 2017. (Year: 2017).\*

Bertens, Paul, et al. "A machine-learning item recommendation system for video games." 2018 IEEE Conference on Computational Intelligence and Games (CIG). IEEE, 2018. (Year: 2018).\*

Choi, Brian; Ryu, MiSun; Panchal, Dharmesh; Le, Andrew (2019). Building Recommender Systems for Video Games on Steam. In Data Science & Engineering Master of Advanced Study (DSE MAS) Capstone Projects. UC San Diego Library Digital Collections. <https://doi.org/10.6075/J0FQ9TZM> (Year: 2019).\*

Notten, Bram. "Improving Performance for the Steam Recommender System using Achievement Data." 2017. Tilburg University, Master's Thesis (Year: 2017).\*

Kang, Wang-Cheng, and Julian McAuley. "Self-attentive sequential recommendation." 2018 IEEE International Conference on Data Mining (ICDM). IEEE, 2018. (Year: 2018).\*

Politowski, Cristiano, et al. "Learning from the past: A process recommendation system for video game projects using postmortems experiences." Information and Software Technology 100 (2018): 103-118. (Year: 2018).\*

Ben-Porat, Omer, and Moshe Tennenholtz. "A game-theoretic approach to recommendation systems with strategic content providers." arXiv preprint arXiv:1806.00955 (2018). (Year: 2018).\*

Wan, Mengting, and Julian McAuley. "Item recommendation on monotonic behavior chains." Proceedings of the 12th ACM conference on recommender systems. 2018. (Year: 2018).\*

Pathak, Apurva, Kshitiz Gupta, and Julian McAuley. "Generating and personalizing bundle recommendations on steam." Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2017. (Year: 2017).\*

Breese, et. al., "Empirical Analysis of Predictive Algorithms for Collaborative Filtering", arXiv: 1301.7363. Jan. 30, 2013, Retrieved on Aug. 19, 2020 from <<<https://arxiv.org/ftp/arxiv/papers/1301/1301.7363.pdf>>>.

PCT Invitation to Pay Additional Fees, dated Aug. 13, 2020, for PCT Application No. PCT/US2020/039885, 2 pages.

PCT Search Report and Written Opinion dated Sep. 3, 2020 for PCT application No. PCT/US20/39925, 9 pages.

Hruska, "Valve Introduces Machine Learning Algorithm to Recommend New Steam Games", retrieved on Sep. 30, 2020 at <<<https://www.extremetech.com/gaming/294859-valve-introduces-machine-learning-algorithm-to-recommend-new-steam-games>>>, Jul. 12, 2019, 6 pages.

PCT Search Report and Written Opinion dated Oct. 13, 2020 for PCT Application No. PCT/US2020/039885, 13 pages.

Trneny, "Machine Learning for Predicting Success of Video Games", retrieved on Sep. 30, 2020 at <<[https://is.muni.cz/th/396121/fi\\_m/diploma\\_thesis\\_trneny.pdf](https://is.muni.cz/th/396121/fi_m/diploma_thesis_trneny.pdf)>>, Dec. 31, 2017, 68 pages.

\* cited by examiner

100




FIG. 1

200 ↘

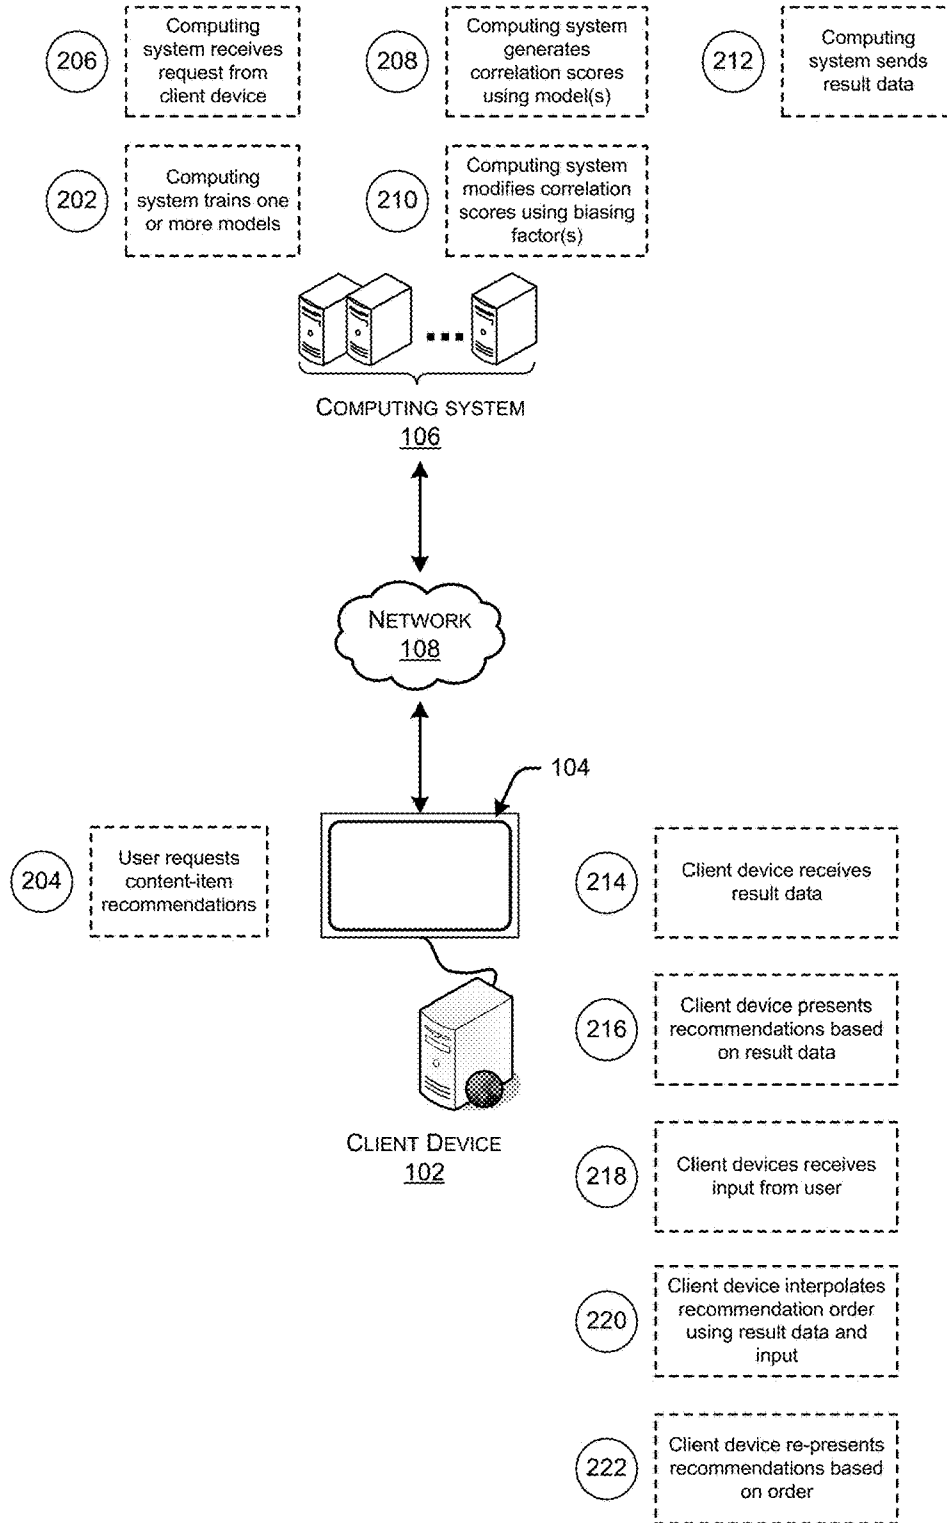


FIG. 2

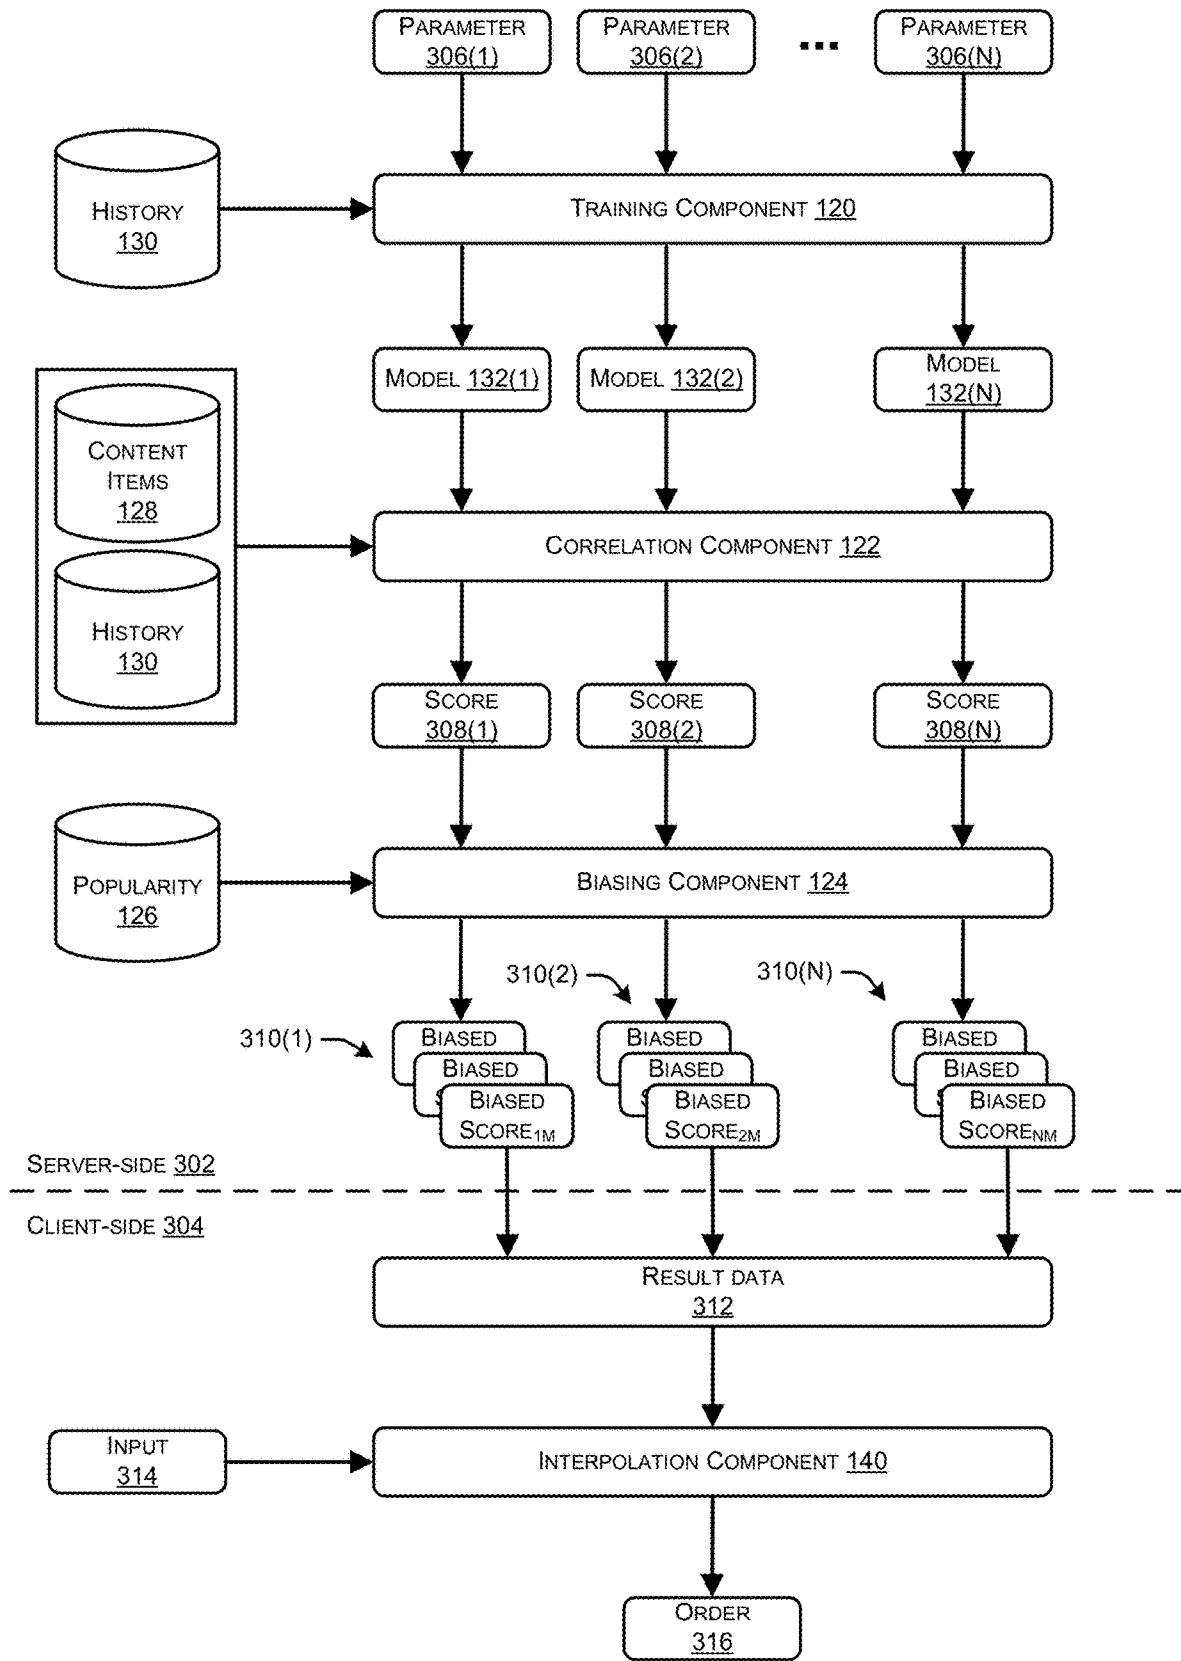


FIG. 3

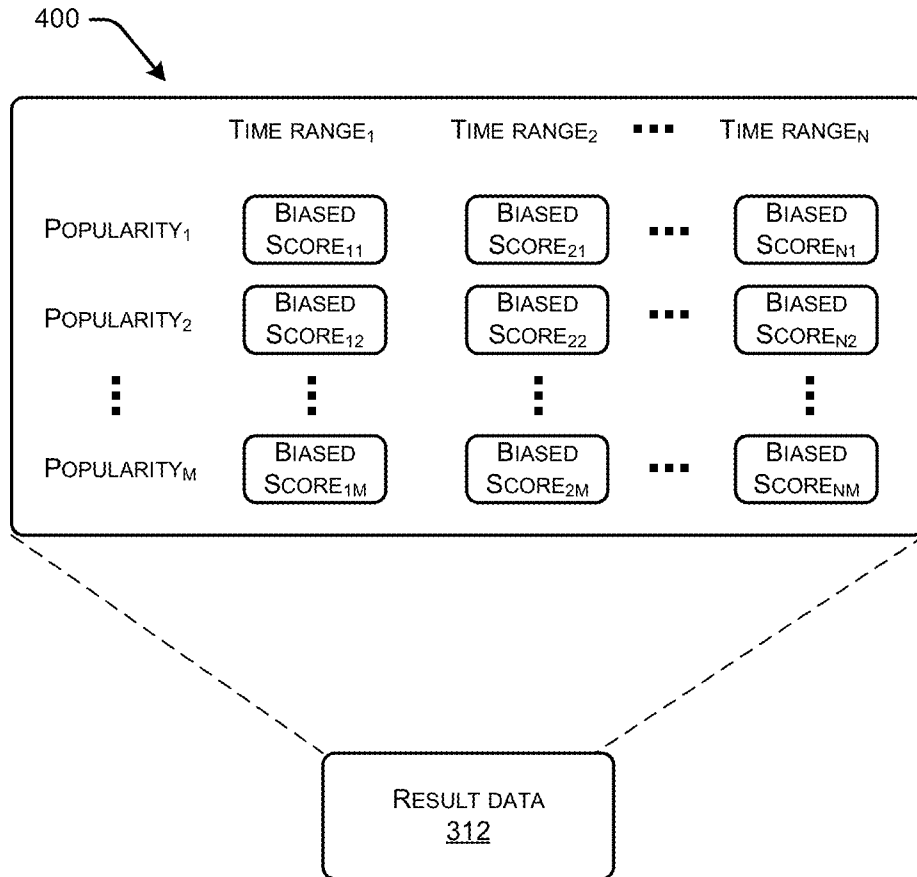


FIG. 4




FIG. 5




FIG. 6



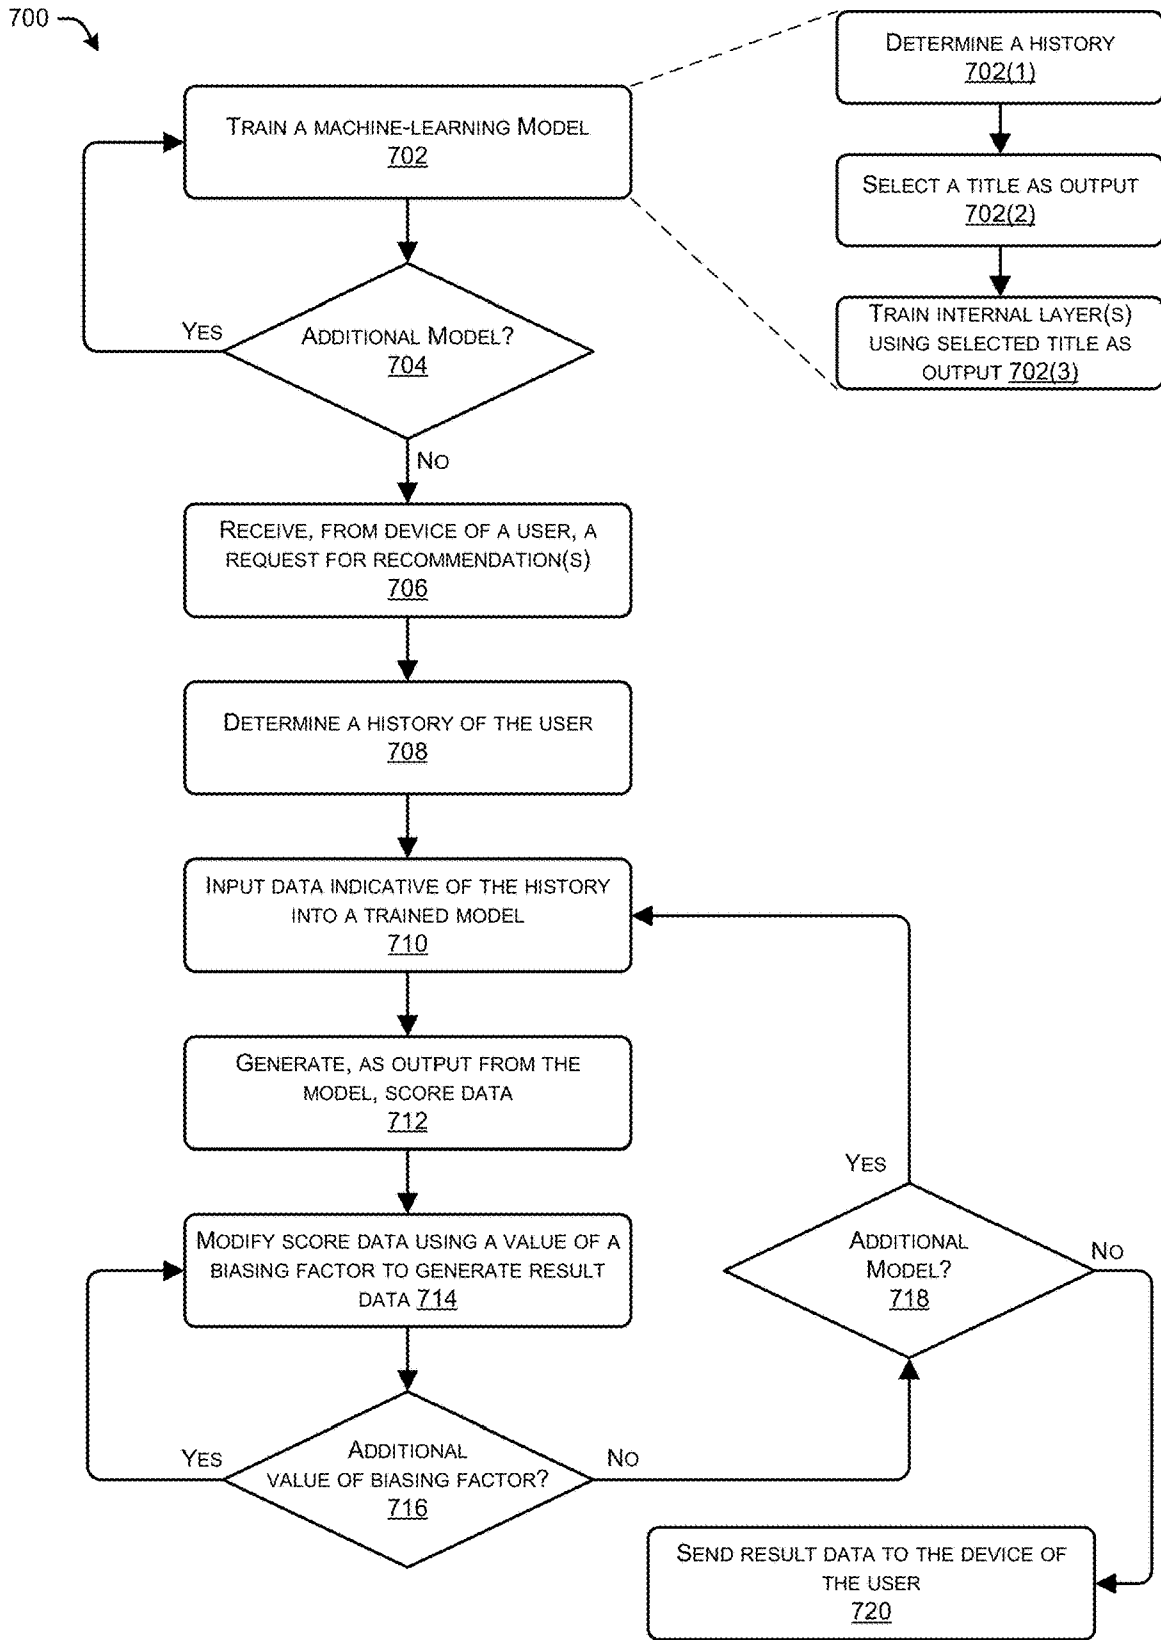


FIG. 7

800 ↘

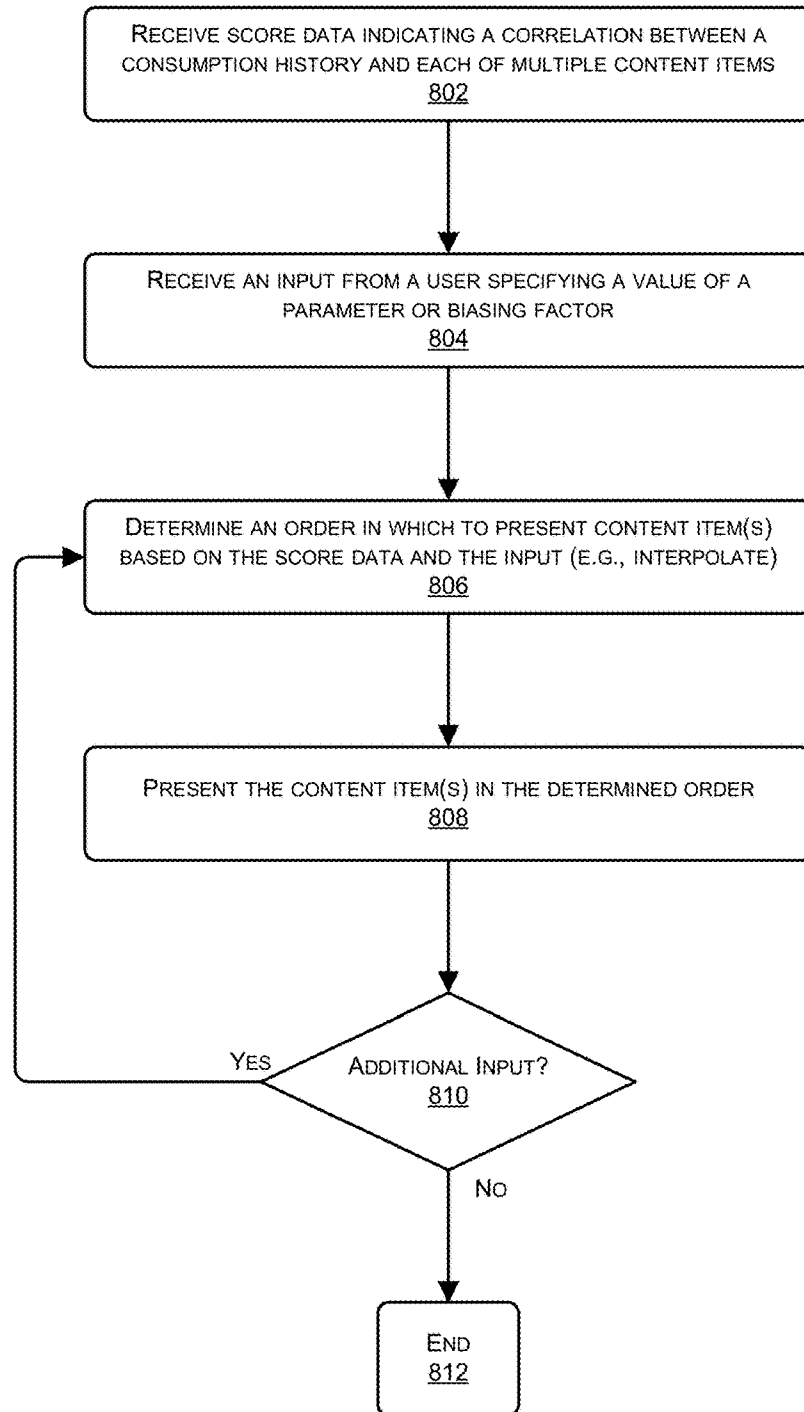


FIG. 8

900

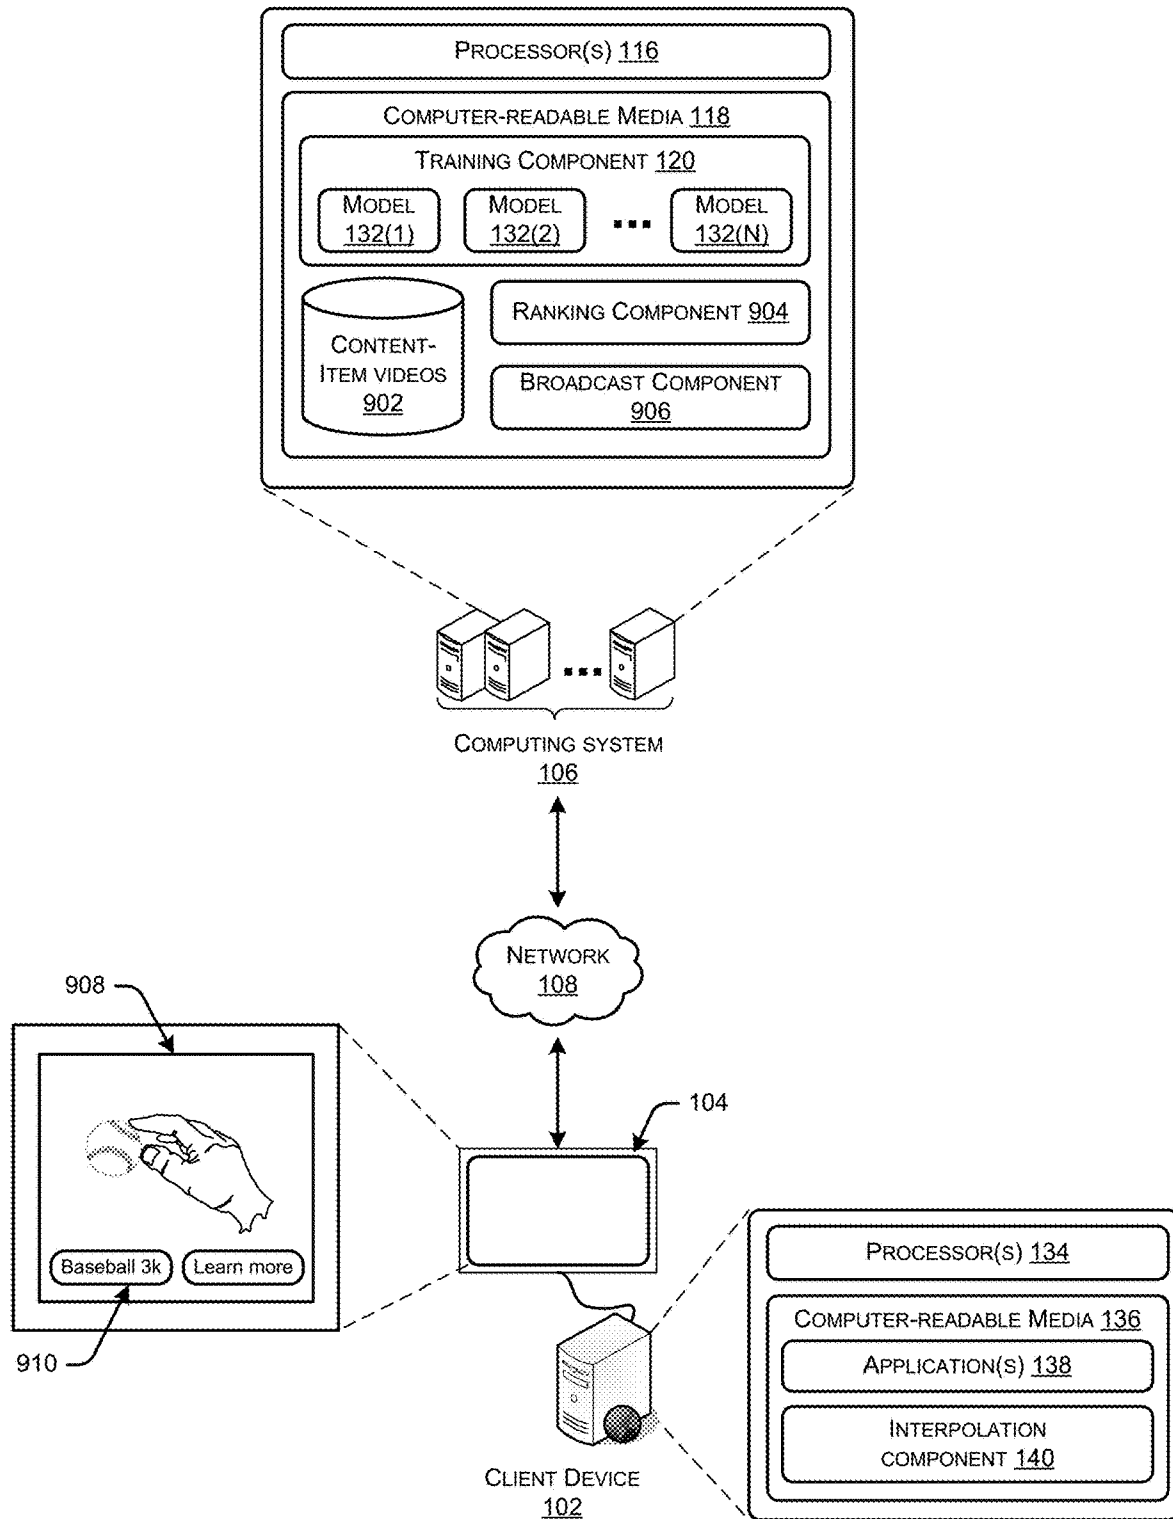


FIG. 9

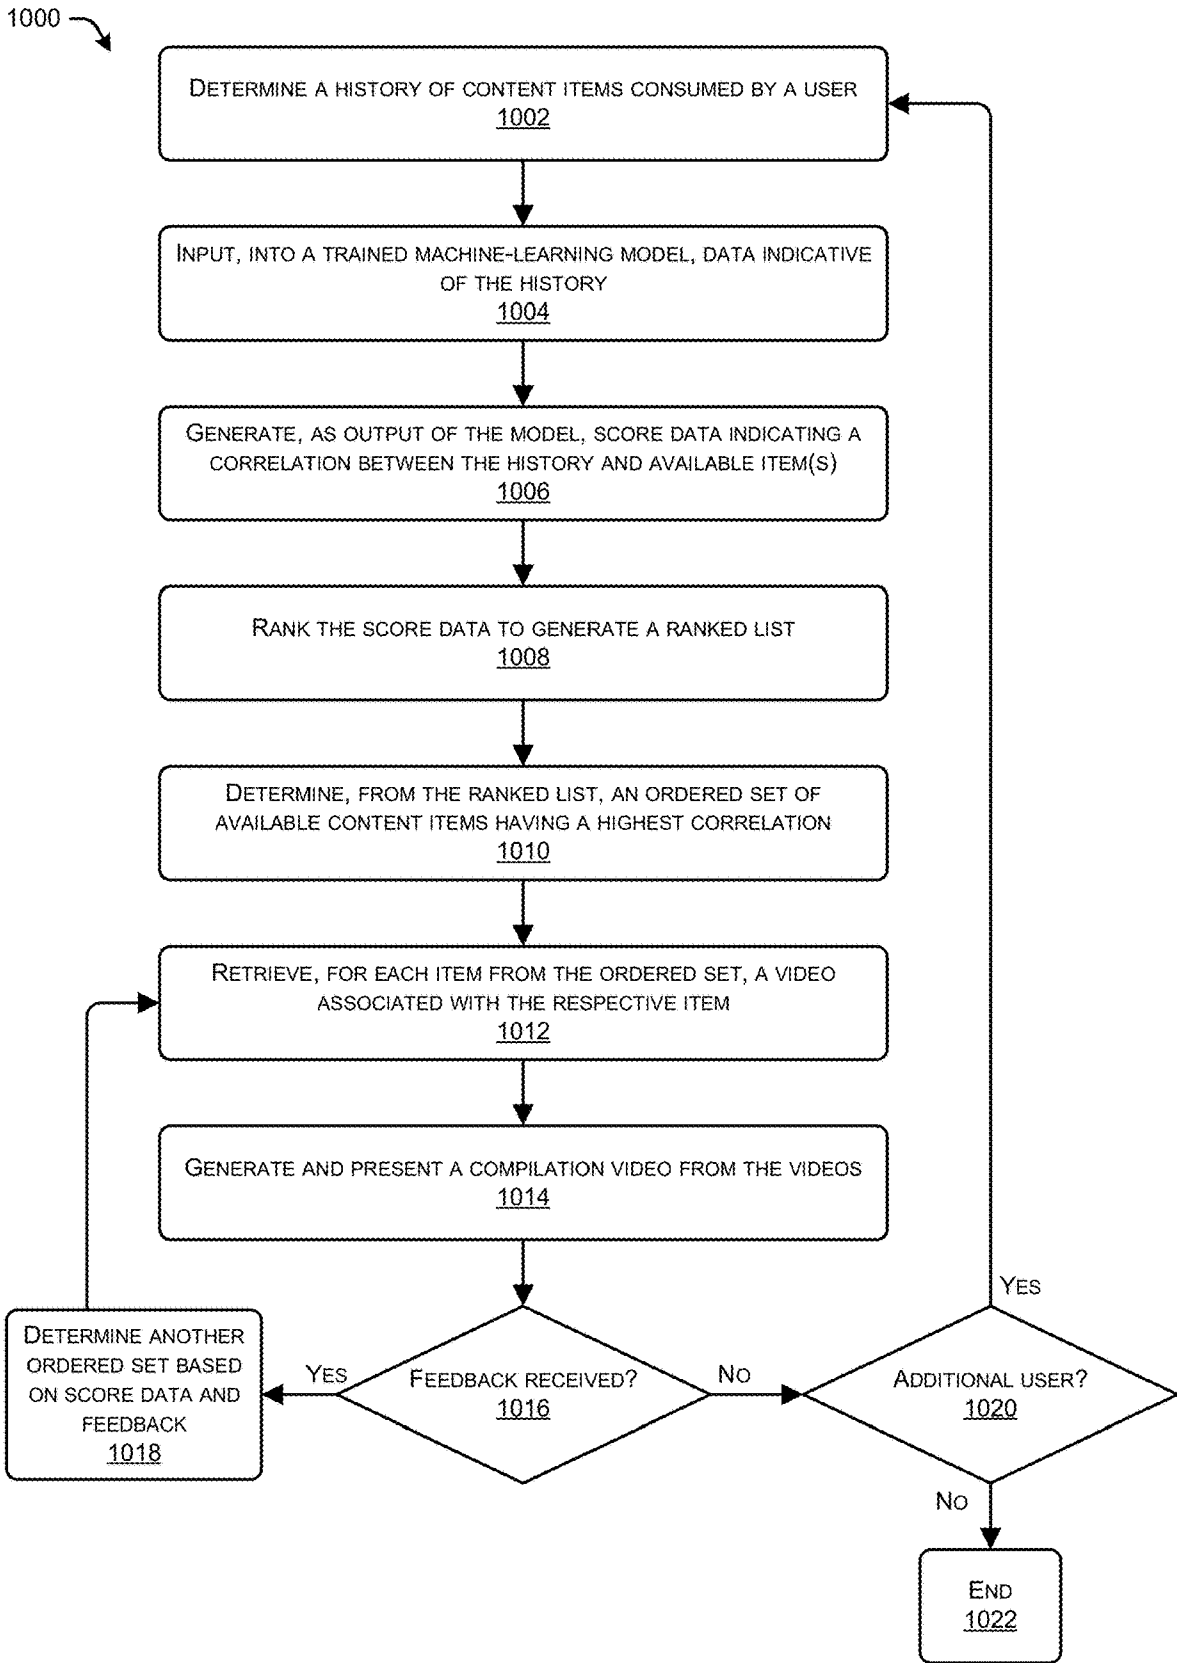


FIG. 10

## CUSTOM COMPILATION VIDEOS

## RELATED APPLICATION

This application is a continuation-in-part of, and claims 5 priority to, U.S. patent application Ser. No. 16/505,112, entitled CONTENT-ITEM RECOMMENDATIONS and filed on Jul. 8, 2019.

## BACKGROUND

As the consumption of content items on electronic devices has continued to proliferate, so has the amount of available content items. For example, the number of songs, movies, television shows, and games available for streaming or 15 download has increased substantially in the recent past. While this increase in available content items is generally beneficial to users, identifying content items of interest for individual users remains difficult.

## BRIEF DESCRIPTION OF THE DRAWINGS

The detailed description is described with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical components or features.

FIG. 1 is a diagram illustrating an example environment 30 that includes a remote computing system configured to train and use machine-learning model(s) for recommending content items to different users. In some instances, the remote computing system trains multiple models and, thereafter, 35 inputs a consumption history of a user into each of the models for generating score data indicating a correlation between each of multiple content-item titles and the consumption history. Further, the computing system may use one or more biasing factors to generate result data, which the computing system may send to a client device of a user. The client device may use this result data, as well as different 40 user inputs, to determine order(s) in which to present content-item recommendations to the user.

FIG. 2 illustrates an example process of operations for 45 generating machine-learning models, applying them to history data of a particular user, sending generated result data to a client device of the user, and using this result data and received user input to determine order(s) in which to present content-item recommendations to the user.

FIG. 3 illustrates example components, and interplay thereof, for generating the models, biasing the resulting scores, applying the models to history data of a particular user, and using this result data and received user input to determine order(s) in which to present content-item recom- 55 mendations to the user.

FIG. 4 illustrates example result data that the remote computing system may generate and send to a client device of a user. As illustrated, the result data may comprise an MxN matrix of scores, which the client device may inter- 60 polate between as needed in response to receiving varying user input.

FIG. 5 illustrates an example user interface (UI) that a client device may present in response to receiving a request from a user for content-item recommendations.

FIG. 6 illustrates an example UI that the client device may present in response to the user altering a value of a first

parameter (in this example, popularity of the content items) and a value of a second parameter (in this example, recency of the content items).

FIG. 7 illustrates a flow diagram of an example process that the remote computing system may employ for generat- 5 ing one or more trained machine-learning models and generating result data for use by a client device in providing content-item recommendations to a user.

FIG. 8 illustrates a flow diagram of an example process 10 that a client device of a user may employ for determining an order in which to present one or more content items to the user.

FIG. 9 is a diagram illustrating an example environment that includes a remote computing system configured to train and use machine-learning model(s) for determining which content-item videos to present to a particular user, and for creating a video comprised of these determined videos. In some instances, the remote computing system trains multiple models and, thereafter, inputs a consumption history of 15 a user into each of the models for generating score data indicating a correlation between each of multiple content-item titles and the consumption history. A broadcast component then generates a video for presentation to the user, with the video comprising trailers (or other videos) associated with the highest-correlated content items.

FIG. 10 illustrates a flow diagram of an example process for determining which content-item videos to present to the user and an order in which to present the videos.

## DETAILED DESCRIPTION

Described herein are, among other things, techniques, devices, and systems for generating one or more trained machine-learning models used for generating content-item recommendations. Also described herein are techniques, devices, and systems for applying a consumption history of a particular user to the trained model(s) to generate score data indicating a correlation between each content-item title and the consumption history, as well as modifying this score data using one or more biasing factors for generating result data. In addition, the techniques, devices, and systems may use this result data, along with received user input, for determining an order in which to present one or more content items to the user. For example, this may include determining which content items to recommend to a user and in which order to do so.

The disclosed techniques may be implemented, at least in part, by a remote computing system that distributes content items (e.g., video games, movies, television shows, songs, etc.) to client devices of a user community (e.g., as part of a video game service, movie service, song service, etc.). These client devices may individually install a client application that is configured to execute the content items received (e.g., downloaded, streamed, etc.) from the remote computing system. In the example of a video game services, the video-game platform enables registered users of the community to play video games as "players." For example, a user can load the client application, login with a registered user account, select a desired video game, and execute the video game on his/her client machine via the client application. Of course, while some examples are described with reference to video-game titles, it is to be appreciated that the described techniques apply to any other type content item, as described above.

In one example, the computing system, which is typically remote from client devices of individual users, generates multiple trained models, each associated with a particular 65

value of a parameter. For example, the computing system may generate a first trained model for use in recommending content items that have been released within a first time window (e.g., the prior six months), a second trained model for use in recommending content items that have been released in a second time window (e.g., the prior year), and so forth. As will be described below, these models may be applying to individual user data for generating tailored content-item recommendations.

While the computing system may train these models in any suitable manner, in one example the computing system begins by accessing respective consumption histories (e.g., gameplay histories, moving histories, etc.) of multiple users associated with the computing-system community. For example, the computing system may access a gameplay history of a first user of a video-game-service community and may use this history for training the first model. For example, if the first model is for use in recommending game titles that have been released within the past six months, the computing system may select, from the gameplay history of the user, one game title that has been released within the past six months and may indicate to the first model that this selected game title is to be the output of the classifier. Thereafter, the computing system may provide the entire gameplay history of the user as input to the first model and the selected game title as output of the first model for training one or more internal layers of the first model. For instance, if the first model comprises an artificial neural network, then selecting the game title as output and providing information regarding the gameplay history of the user as input may be effective to train one or more internal layers of the neural network. It is to be appreciated that the input and output of the neural network, and other types of models that may be used, may comprise features of the corresponding game titles or other content items, such as release data, genre, game-type, and/or the like. In addition, the input to the models may include information associated with the user, such as a geographical location associated with the user, demographic information associated with the user, and/or the like. In addition, while a few example inputs to the models have been described, it is to be appreciated that any other type of input may be provided to the models.

In some instances, the computing system may take into account the amount of time that the user has played each game in her history in selecting the game title to be used as output for the machine-learning model. For example, the computing system may initially determine an amount of time that the user has played each individual game title referenced in the gameplay history of the user. This may be determined either absolutely (e.g., in terms of time played) or relatively (e.g., in terms of time played relative to other users that have also played the game title). This information may be applied as a weight to the individual game title or may be passed into the model as an input. After applying a weight to each game title associated with the gameplay history or inputting this information to the models (after, potentially, filtering out some game titles due to a very small amount of gameplay by the user), the computing system may then select one (or more) game title to be used as output by the first model. By doing so, the computing system increases the chances that the games played the most by the user are selected as the output of the first model, thus training the internal layer(s) of the first model in a way that results in the model being more likely to select games of interest to the user. It is to be appreciated, however, that while the above description illustrates one manner in which to train the

models, other models for outputting content-item recommendations may be trained in any other number of ways.

In addition to training a first model, in some instances the computing device may also train multiple other models, such as a model for recommending games released within the last year, a model for recommending games released within the last three years, and so forth. In each instance, the computing system may access a gameplay history of a user and select, as the output to the particular model, a game title that is both indicated in the gameplay history and that was released within the amount of time associated with the particular model (e.g., within one year, three years, etc.). Furthermore, while the above example describes accessing a single gameplay history for training each of the multiple machine-learning models, it is to be appreciated that the computing system may analyze thousands or more of the gameplay histories for training the models.

After training the machine-learning models, the computing system may then use the models for generating data that may be used by client devices for surfacing content-item (e.g., video game) recommendations. For example, envision that a particular user sends a request, via a client device, for video-game recommendations. Upon receiving the request, the computing system may identify the user (e.g., via a device or user identifier) and may access a gameplay history of the user. The computing system may then input the gameplay history into each of the multiple machine-learning models along with information regarding available games associated with each model. For example, for the first trained model (which has been trained to more heavily weight games released within the last six months), the computing system may input the gameplay history of the user and information regarding games available to be recommended, for the second trained model (which has been trained to more heavily weight games released within the last year), the computing system may input the gameplay history of the user and information regarding those games available to be recommended, and so forth. That is, while the entire catalog of available games may be input to the first model, the second model, and so forth, the first model may be more likely to recommend games having been released within the last six months, the second model may be more likely to recommend games that have been released within the last year, and so forth.

Each trained machine-learning model may output score data that indicates, for each of multiple game titles, a correlation between the gameplay history of the user and the particular game title. As the reader will appreciate, those games that are highly correlated to the gameplay history will have a relatively high score, and vice versa. Further, if the computing system implements “N” number of models, then the computing system may generate “N” number of score data. It is also to be appreciated that the “correlation” might not be a statistical correlation, but rather simply a representation of a likelihood that a user with the same gameplay history of the subject user would choose to play this game.

In addition, the computing system may be configured to bias the score data based on different values of one or more biasing factors. The biasing factors may comprise any number of factors, such as a popularity of each game, a cost of each game, a genre of each game, or the like. In one example, the biasing factor may comprise a popularity of each game, as determined by sales of the game, an amount of gameplay of the game, and/or the like. In some instances, each score data (determined by an individual trained model) may be biased at different levels of popularity such that the resulting biased data will more closely match that particular

level of popularity if so desired by the user. That is, the computing system may bias the score for each game title (and for each model) at a number of “M” target levels of popularity ranging from very popular to very unpopular (or very “niche”). As the reader will appreciate, the computed scores of those game titles that are very popular will be boosted for a target level of popularity of “very high” and will be penalized for a target level of popularity of “very low”. Conversely, the computed scores of more niche game titles will be penalized for a target level of popularity of “very high” and will be boosted for a target level of popularity of “very low”.

Finally, given that the computing system has now computed, for a particular user, score data for each of “N” number of trained models, and has biased each of these scores using “M” number of popularity values, then computing system may now have generated an amount of scores for the user in an amount of “M×N”. For example, if the computing system utilizes six models and five levels of popularity, then the computing system will have generated a total of thirty scores for any particular user (given that each of the six score data generated by the model will be biased five times). Of course, it is to be appreciated that the computed score data may have any other number of values and may be sampled in any other (e.g., nonlinear) manner. Further, while this example describes inputting values of a first parameter prior to training the model (in this case, based on release date) and values of a second after the training (in this case, based on popularity), any other number of parameters may be applied as input to the model and/or as part of the postprocessing described with reference to the popularity data. It is also noted that while the postprocessing parameters are described as biasing factors, these biasing factor(s) represent parameter(s) that may be applied as inputs to the models or as part of postprocessing.

Upon computing this score data, in the form of a matrix or otherwise, the computing system may then send this “result data” to the client device of the user. The client device of the user may then use this result data, as well as input from the user, for determining an order in which to present one or more of the game titles available to the user. That is, the client device may use the result data and input data from the user for determining which video games to recommend to the user, and how to change this recommendation as the user provides different input data.

To continue the example from above, the client device may enable the user to provide, as input, a selection of a desired release recency and a selection of a desired level of popularity. That is, the user may indicate, to the client device, whether the user prefers to see, as game recommendations, games that have been more recently released or games that are older. In addition, the user may specify whether she would like to receive recommendations for games that are in the mainstream (that is, popular), or games that are more niche (that is, less popular). While the examples below illustrate and describe the user as providing this input via one or more sliders, it is to be appreciated that this input may be provided in any other manner.

Upon receiving an input from a user, such as input indicating that the user would like to see recommendations for games that are generally newer and more niche, the client device may use the result data to determine which games to recommend to the user and in which order. In some instances, the client device may identify, from the input, which one or more of the computed score data that the client device has received for the user is closest to the input data of the user. For instance, in the example of the user speci-

fying relatively new games that are less popular, the client device may identify which of the scores from the computed scored data correspond to: (1) the model(s) that have been trained for this time range, and (2) the scores that have been biased to boost less popular titles. After identifying these scores, the client device may use interpolation for determining scores associated with individual games and may then present the game recommendations to the user based on these scores. For example, the client device may utilize bilinear interpolation or any other form of interpolation to compute scores for the game titles and may sort these game titles according to the computed scores, with those game titles having highest scores being placed at the top of the list. Further, if the user provides different input data—such as by moving one or both sliders described above—the client device may again identify the most pertinent score data from the matrix or other representation of scores and may interpolate this score data for re-ranking the game titles.

By generating this score data at the remote computing system and performing the interpolation locally at the client device, the described techniques greatly improve the fluidity of the process for ranking and re-ranking content-item recommendations at the client device. That is, because the techniques involve sending multiple score data to the client device rather than a single result, and because the client device is configured to perform real-time interpolation of these score data in response to receiving different user input, the techniques enable the ranking and re-ranking of recommendations without requiring additional roundtrip requests to the remote computing system. Stated otherwise, the techniques enable unlimited reranking of content-item recommendations at the cost of a single request to and response from the remote computing system, as opposed to a requirement of making a call to the system and awaiting a response from the system each time the user provides different input.

In addition, the techniques described herein provide for a very fine-grain level of content recommendation. By enabling a user to specify a desired value of a first parameter (e.g., cost, release date, etc.) and a desired value of a second parameter (e.g., popularity, cost, etc.), and by performing client-side interpolation of these inputs, the techniques provide highly-tailored content-item recommendations.

It is to be appreciated that while some of the techniques are described herein with reference to two parameters (release data and popularity), the techniques may be used for any other number of parameters (e.g., one, five, etc.) and for any other parameter values (e.g., cost, genre, etc.). It is also to be appreciated that while the recommendations are described as being presented on one or more UIs, the recommendations may be output in any other manner (e.g., audibly, visually, etc.).

In addition, described herein are, among other things, techniques, devices, and systems for generating one or more trained machine-learning models used for generating which content-item videos (e.g., trailers, etc.) to present to a user, as well as an order in which to present these videos. Further, the techniques, devices, and systems may, similar to the techniques above, apply a consumption history of a particular user to the trained model(s) to generate score data indicating a correlation between each content-item video and the consumption history, as well as modifying this score data using one or more biasing factors for generating result data. In addition, the techniques, devices, and systems may use this result data, along with received user input, for determining an order in which to present the one or more

content-item videos to the user. For example, this may include determining which videos to recommend to a user and in which order to do so.

In one example, the techniques train one or more models using the techniques described above with reference to content-item recommendations. For example, the techniques may begin by accessing respective consumption histories (e.g., gameplay histories, moving histories, etc.) of multiple users associated with the computing-system community. For example, the computing system may access a gameplay history of a first user of a video-game-service community and may use this history for training the first model. For example, the computing system may select, from the gameplay history of the user, one game title and may indicate to the first model that this selected game title is to be the output of the classifier. Thereafter, the computing system may provide the entire gameplay history of the user as input to the first model and the selected game title as output of the first model for training one or more internal layers of the first model. For instance, if the first model comprises an artificial neural network, then selecting the game title as output and providing information regarding the gameplay history of the user as input may be effective to train one or more internal layers of the neural network. It is to be appreciated that the input and output of the neural network, and other types of models that may be used, may comprise features of the corresponding game titles or other content items, such as release data, genre, game-type, and/or the like. In addition, the input to the models may include information associated with the user, such as a geographical location associated with the user, demographic information associated with the user, and/or the like. In addition, while a few example inputs to the models have been described, it is to be appreciated that any other type of input may be provided to the models. Further, it is to be appreciated that multiple models may be trained for the first user (e.g., corresponding to a first time range, a second time range, or the like) as described above, or a single model may be trained.

In some instances, the computing system may take into account the amount of time that the user has played each game in her history in selecting the game title to be used as output for the machine-learning model. For example, the computing system may initially determine an amount of time that the user has played each individual game title referenced in the gameplay history of the user. This may be determined either absolutely (e.g., in terms of time played) or relatively (e.g., in terms of time played relative to other users that have also played the game title). This information may be applied as a weight to the individual game title or may be passed into the model as an input. After applying a weight to each game title associated with the gameplay history or inputting this information to the models (after, potentially, filtering out some game titles due to a very small amount of gameplay by the user), the computing system may then select one (or more) game title to be used as output by the first model. By doing so, the computing system increases the chances that the games played the most by the user are selected as the output of the first model, thus training the internal layer(s) of the first model in a way that results in the model being more likely to select games of interest to the user. It is to be appreciated, however, that while the above description illustrates one manner in which to train the models, other models for outputting content-item recommendations may be trained in any other number of ways. Furthermore, while the above example describes accessing a single gameplay history for training one or more machine-learning models, it is to be appreciated that the computing

system may analyze thousands or more of the gameplay histories for training the models.

After training the machine-learning models, the computing system may then use the models for generating data that may be used by client devices for surfacing content-item videos to recommend to a user. In some instances, these recommendations are used to determine the videos to surface to a particular user and an order of the videos. With this information, the computing system may generate what appears to the user as a single video, but in reality comprises multiple individual content-item videos played to the user in a serial (or concurrent) manner.

For example, envision that a particular user sends a request, via a client device, to view a video showing individual videos of content items that are recommended for the user. That is, envision that the user requests to view a video showing a series of video-game trailers for video games that are recommended for this particular user. Upon receiving the request, the computing system may identify the user (e.g., via a device or user identifier) and may access a gameplay history of the user. The computing system may then input the gameplay history into a machine-learned model. In response, the machine-learned model (or models in examples of multiple models being utilized) may output score data that indicates, for each of multiple game titles, a correlation between the gameplay history of the user and the particular game title. As the reader will appreciate, those games that are highly correlated to the gameplay history will have a relatively high score, and vice versa. It is also to be appreciated that the “correlation” might not be a statistical correlation, but rather simply a representation of a likelihood that a user with the same gameplay history of the subject user would choose to play this game. In addition, the computing system may be configured to bias the score data based on different values of one or more biasing factors, as described above. The biasing factors may comprise any number of factors, such as a popularity of each game, a cost of each game, a genre of each game, or the like.

Regardless of whether the score data has been biased, the computing system may then determine, from the score data, the top “X” number of video games that are to be recommended for the particular user. For example, the score data may be used to identify the top ten video-game titles recommended for the user. The computing system may then send an instruction to a broadcast component, described below, to generate a video for the user corresponding to trailers for these particular games. That is, the computing system may instruct the broadcast component to present, to the user, a video comprising a compilation of trailers corresponding to the video games recommended for the user. In some instances, the broadcast component may present the trailers serially according to a prescribed order, while in other instances the broadcast layer may present the trailers concurrently. For example, the broadcast component may present a first trailer at a first location on a display of the user, a second trailer at a second location on the display of the user, and so forth. In some instances, a top-ranked trailer (i.e., a trailer corresponding to the top-ranked game for the particular user) may be deemed the primary trailer, and music or other audio for that primary trailer may be played for the entirety of the video of the top “X” trailers.

In some instances, the broadcast component may also present additional information atop the presented trailers. For example, the broadcast component may retrieve and present a link for purchasing a video game corresponding to a currently playing trailer, details of the game associated with the trailer (e.g., name, release date, etc.), and so forth.



The broadcast component may also present video controls to enable a user to control playback of the video (e.g., pause, rewind, etc.).

In some instances, the broadcast component may create the video of trailers by obtaining relatively small portions of each individual trailer and presenting these individual portions serially and in an order defined by the score data described above. For example, the broadcast component may determine an order in which to present the ten (or other number of) trailers and may acquire two-second, for example, increments of these trailers in the defined order. Further, the broadcast component may present these ten trailers in the order defined by the model in a way that appears, to the end user, as a single video rather than a compilation of multiple individual videos.

While a single video (of multiple trailers) may be determined and presented for a particular user in some instances, in other instances the techniques may further enable a user to specify one or more value of one or more different parameters for causing the system to present a different video of multiple trailers. For example, the system may present, to the user, interactive elements that enable a user to specify a desired release recency, a desired level of popularity, and/or the like. Upon receiving this selection, the system may use one or more of the trained models to determine score data for identifying the trailers to present to the user and an order in which to present them. The broadcast component may then receive this information and generate the new video for presentation to the user. In some instances, the system may enable a user to provide the input via tags. For example, the user may request to exclude, from the video, trailers for games associated with a particular tag and/or may request to include only trailers for games associated with a particular tag.

In some instances, the techniques for generating tailored videos for users may limit the universe of games (or other content items) from which the trailers are selected. For example, the techniques may apply a model to a particular subset of games, may receive score data from the model to identify the top “X” recommended games (and an order thereof), and may generate a video comprising trailers of these recommended games in the defined order. For example, this closed universe may correspond to games already owned (e.g., previously purchased) by a user, games on sale, games from a particular publisher, or the like.

As noted above, upon identifying the top “X” recommended games and an order in which to present trailers corresponding to these games, the broadcast component may generate the file for presenting these “X” number of trailers in the determined order. In addition, the broadcast component may include one or more “intro” or “outro” videos at the beginning, end, and/or mixed in with the “X” number of trailers. For example, if the broadcast component determines to present ten trailers to a particular user, the broadcast component may present, prior to the ten trailers, an intro video introducing the ten games, with this video including information associated with these games, such as the genre they are associated with, links to each game, and the like. Similarly, the broadcast component may present, at the end of the ten videos, an outro video, indicating information associated with these trailers, with this information being the same or different as that presented in the intro video.

In still another example, these intro and/or outro videos may be interspersed within the ten (or other example number of) trailers. For example, if the first five trailers are associated with a first tag or genre (e.g., adventure) and the second five trailers are associated with a second tag or genre (e.g.,

puzzle games), then the broadcast component may generate and provide a first intro video prior to presenting the first five trailers, a first outro video after presenting the first five trailers, a second intro video after presenting the first outro video, and a second outro video after presenting the last five trailers. Again, the intro and/or outro videos may include information associated with the corresponding trailers. For instance, the first intro and/or outro videos may indicate the five adventure-related trailers, with information for obtaining these games or the like. Similarly, the second intro and/or outro videos may indicate the five puzzle-related trailers, with information for obtaining these games and/or additional information. The type of content of these intro and/or outro videos may comprise static content (e.g., static images, links, etc.) and/or dynamic content (e.g., voiceover data introducing the games, animations, prerecorded video of a curator introducing the content, or the like).

In yet another example, the output of the trained models may be used to make recommendations to users other than content-item (e.g., game) recommendations. For example, the output of these model(s) may be used to recommend content-item curators to the users in some instances. For example, continuing the example from above where a model determines, using a consumption history of a particular user, a set of the top “X” video games (or other content items) to recommend to a user, this output may be compared to information regarding one or more curators to determine whether to recommend one or more curators to the user. For example, the techniques may compute a similarity between the content items output by the model and the content items recommended by each of multiple human curators.

That is, the techniques may compare these model-recommended games to each set of games recommended by respective human curators. The techniques may then use this similarity data to suggest, to the user, that the user connect with (e.g., begin “following”) one or more human curators. For example, the techniques may recommend that top “P” number of curators according to the computed similarity scores, those curators associated with a similarity score that is greater than a threshold similarity score, and/or the like.

FIG. 1 is a diagram illustrating an example environment that includes a client device **102**, associated with a display **104**, and a remote computing system **106** configured to train and use machine-learning model(s) for recommending content items to different users. In some instances, and as described in further detail below, the remote computing system **106** trains multiple models and, thereafter, inputs a consumption history of a user into each of the models for generating score data indicating a correlation between each of multiple content-item titles and the consumption history. Further, the computing system **106** may use one or more biasing factors to generate result data, which the computing system may send to the client device **102** of a user. The client device **102** may use this result data, as well as different user inputs, to determine order(s) in which to present content-item recommendations to the user, on the display **104** or otherwise.

The client device **102** may be implemented as any suitable type of computing device configured to execute content items, such as video games, movies, songs, and/or the like. The client device **102** may comprise, without limitation, a personal computer (PC), a desktop computer, a laptop computer, a mobile phone (e.g., a smart phone), a tablet computer, a portable digital assistant (PDA), a wearable computer (e.g., virtual reality (VR) headset, augmented reality (AR) headset, smart glasses, etc.), an in-vehicle (e.g., in-car) computer, a television (smart television), a set-top-box

11

(STB), a game console, a music player, a voice-controlled assistant, and/or any similar computing device. The client device **102** may communicate with the remote computing system **106** (sometimes shortened herein to “computing system **106**”) over a computer network **108**. The computer network **108** may represent and/or include, without limitation, the Internet, other types of data and/or voice networks, a wired infrastructure (e.g., coaxial cable, fiber optic cable, etc.), a wireless infrastructure (e.g., radio frequencies (RF), cellular, satellite, etc.), and/or other connection technologies. The computing system **106** may, in some instances be part of a network-accessible computing platform that is maintained and accessible via the computer network **108**. Network-accessible computing platforms such as this may be referred to using terms such as “on-demand computing”, “software as a service (SaaS)”, “platform computing”, “network-accessible platform”, “cloud services”, “data centers”, and so forth.

In some embodiments, the computing system **106** acts as, or has access to, a video game platform that implements a video game service to distribute (e.g., download, stream, etc.) video games (or any other type of content item) to the client device **102**. In an example, the client device **102** may each install a client application thereon. The installed client application may be a video-game client (e.g., gaming software to play video games). A client device **102** with an installed client application may be configured to download, stream, or otherwise receive programs (e.g., video games, and content thereof) from the computing system **106** over the computer network **108**. Any type of content-distribution model can be utilized for this purpose, such as a direct purchase model where programs (e.g., video games) are individually purchasable for download and execution on a client device **102**, a subscription-based model, a content-distribution model where programs are rented or leased for a period of time, streamed, or otherwise made available to the client devices **102**. Accordingly, an individual client device **102** may include one or more installed video games that are executable by loading the client application.

The client device **102** may be used to register with, and thereafter login to, a video game service. A user may create a user account for this purpose and specify/set credentials (e.g., passwords, PINs, biometric IDs, etc.) tied to the registered user account. As a plurality of users interact with the video game platform (e.g., by accessing their user/player profiles with a registered user account, playing video games on their respective client devices **102**, etc.), the client devices **102** send data to the remote computing system **106**. The data sent to the remote computing system **106**, for a given client machine **104**, may include, without limitation, user input data, video game data (e.g., game performance statistics uploaded to the remote system), social networking messages and related activity, identifiers (IDs) of the video games played on the client device **102**, and so on. This data can be streamed in real-time (or substantially real-time), sent the remote system **106** at defined intervals, and/or uploaded in response to events (e.g., exiting a video game). As described in further detail below, this data may be used to determine a gameplay history of the user of the client device **102**, which may be used for determining scores for recommending game titles to the user.

FIG. 1, for example, illustrates that the client device **102** may present, on the display **104**, a UI **110** for recommending one or more content items (e.g., video games) to the user of the client device **102**. In this example, the UI **110** includes one or more UI controls **112(1)** and **112(2)** for providing two different user inputs. For example, the first UI control **112(1)**

12

may enable a user of the client device **102** to provide a desired value of a first parameter, while the second UI control **112(2)** may enable the user to provide a desired value of a second parameter. The client device may then use these values for generating a custom list **114** of recommended content items **114(1)**, **114(2)**, **114(3)**, and **114(4)**. As described and illustrated in further detail below, the first UI control **112(1)** may enable the user to specify a desired level of popularity, while the second UI control **112(2)** may enable the user to specify a desired level of release-date recency of the content items. Further, and also as discussed below, the client device **102** may update the list **114** of recommended content items in response to the user provides input via the UI control(s) and without interacting with the computing system **106**.

FIG. 1 further illustrates that the computing system **106** may include one or more processors **116** (e.g., central processing unit(s) (CPU(s)) and computer-readable media **118**. The computer-readable media **118** may include volatile and nonvolatile memory, removable and non-removable media implemented in any method or technology for storage of information, such as computer-readable instructions, data structures, program modules, or other data. Such memory includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, RAID storage systems, or any other medium which can be used to store the desired information and which can be accessed by a computing device. The computer-readable media **118** may be implemented as computer-readable storage media (“CRSM”), which may be any available physical media accessible by the processor(s) **116** to execute instructions stored on the computer-readable media **118**. In one basic implementation, CRSM may include random access memory (“RAM”) and Flash memory. In other implementations, CRSM may include, but is not limited to, read-only memory (“ROM”), electrically erasable programmable read-only memory (“EEPROM”), or any other tangible medium which can be used to store the desired information and which can be accessed by the processor(s) **116**.

As illustrated, the computer-readable media **118** may store or otherwise have access to a training component **120**, a correlation component **122**, and a biasing component **124**. In addition, the media **118** may store popularity data **126** indicating popularity (e.g., sales, gameplay time, etc.) of respective content items, content items **128** available for acquisition at the client device, and history data **130** indicating respective consumption histories of different users.

The training component **120** may use the history data **130** to train one or more machine-learning models **132(1)**, **132(2)**, . . . , **132(N)**. For example, the training component **120** may train each of multiple machine learning models using a portion of the history data **130** that is associated with a sampled set of user accounts as training data to obtain trained machine-learning models **132(1)-(N)**.

The trained machine-learning model(s) **132** may represent a single model or an ensemble of base-level machine-learning models, and may be implemented as any type of machine-learning model **132**. For example, suitable machine-learning models **132** for use with the techniques and systems described herein include, without limitation, neural networks, tree-based models, support vector machines (SVMs), kernel methods, random forests, splines (e.g., multivariate adaptive regression splines), hidden Markov model (HMMs), Kalman filters (or enhanced Kal-

man filters), Bayesian networks (or Bayesian belief networks), expectation maximization, genetic algorithms, linear regression algorithms, nonlinear regression algorithms, logistic regression-based classification models, or an ensemble thereof. An “ensemble” can comprise a collection of machine-learning models **132** whose outputs (predictions) are combined, such as by using weighted averaging or voting. The individual machine-learning models of an ensemble can differ in their expertise, and the ensemble can operate as a committee of individual machine-learning models that is collectively “smarter” than any individual machine-learning model of the ensemble.

The training data that is used to train each machine-learning model **132** may include various types of data. In general, training data for machine learning can include two components: features and labels. However, the training data used to train the machine-learning model(s) **132** may be unlabeled, in some embodiments. Accordingly, the machine-learning model(s) **216** may be trainable using any suitable learning technique, such as supervised learning, unsupervised learning, semi-supervised learning, reinforcement learning, and so on. The features included in the training data can be represented by a set of features, such as in the form of an n-dimensional feature vector of quantifiable information about an attribute of the training data. The following is a list of example features that can be included in the training data for training the machine learning model (s) **132** described herein. However, it is to be appreciated that the following list of features is non-exhaustive, and features used in training may include additional features not described herein, and, in some cases, some, but not all, of the features listed herein. Example features included in the training data may include, without limitation, a release date a content item, a genre of a content item, a title of a content item, a length of a content item, a rating of a content item, users that have consumed the content item, and/or the like. Further, as part of the training process, the training component **120** may set weights for machine learning. These weights may apply to a set of features included in the training data, as derived from the historical data **130**. In some embodiments, the weights that are set during the training process may apply to parameters that are internal to the machine-learning model(s) (e.g., weights for neurons in a hidden-layer of a neural network). These internal parameters of the machine-learning model(s) may or may not map one-to-one with individual input features of the set of features. The weights can indicate the influence that any given feature or parameter has on the score that is output by the trained machine learning model **132**.

In some instances, the training component **120** may train individual models by accessing individual consumption histories of users, selecting a particular content item indicated by the respective history as the output of the model, inputting the consumption history as input to the model, and training one or more internal layers of the model. For example, and as described above, the training component may access a consumption history of a first user and may use this history for training the first model **132(1)**. For example, if the first model **132(1)** is for use in recommending content-item titles that have been released within the past six months, the training component **120** may select, from the history of the user, one game title that has been released within the past six months and may indicate to the first model **132(1)** that this selected title is to be the output of the model. Thereafter, the training component **120** may provide the entire consumption history of the user as input to the first model **132(1)** and the selected content-item title as output of

the first model **132(1)** for training one or more internal layers of the first model **132(1)**. For instance, if the first model **132(1)** comprises an artificial neural network, then selecting the content-item title as output and providing information regarding the consumption history of the user as input may be effective to train one or more internal layers of the neural network. It is to be appreciated that the input and output of the neural network, and other types of models that may be used, may comprise features of the corresponding game titles or other content items, such as release data, genre, game-type, and/or the like.

In some instances, the training component **120** may take into account the amount of time that the user has consumed each content item in her history in selecting the title to be used as output for the machine-learning model **132(1)**. For example, the training component **120** may initially determine an amount of time that the user has consumed each individual title referenced in the consumption history of the user. This may be determined either absolutely (e.g., in terms of time) or relatively (e.g., in terms of time relative to other users that have also consumed the title). This information may be applied as a weight to the individual content-item title. After applying a weight to each title associated with the history or inputting this information to the models (after, potentially, filtering out some titles due to a very small amount of consumption by the user), the training component **120** may then select one (or more) content-item title to be used as output by the first model **132(1)**. By doing so, the training component **120** increases the chances that the content items consumed the most by the user are selected as the output of the first model **132(1)**, thus training the internal layer(s) of the first model in a way that results in the model being more likely to select content items of interest to the user.

In addition to training a first model **132(1)**, in some instances the training component **120** may also train multiple other models, such as a model for recommending games released within the last year **132(2)**, a model for recommending games released within the last three years **132(3)**, and so forth. In each instance, the training component **120** may access a consumption history of a user and select, as the output to the particular model, a title that is both indicated in the history and that was released within the amount of time associated with the particular model (e.g., within one year, three years, etc.). Furthermore, while the above example describes accessing a single history for training each of the multiple machine-learning models, it is to be appreciated that the training component **120** may analyze thousands or more of the histories for training the models **132**.

The correlation component **122**, meanwhile, may function to generate correlation scores for individual users in response to receiving recommendation requests from these individual users. For example, as illustrated the client device **102** may include one or more processors **134** and computer-readable media **136**, which may store or otherwise having access to one or more client applications **138** and an interpolation component **140**. In some instances, a user of the client device **102** may utilize one of the client applications **138** (e.g., a browser, a dedicated application, etc.) to interact with the computing system **106**, such as to request data for rendering the UI **110**. In response to receiving this request, the correlation component may determine an identifier associated with a user of the client device **102** for generating score data that may be used to recommend one or more content items to the user.

In one example, the user correlation component **122** access, from the history data **130**, a consumption history of the user of the client device **102** after identifying the user. The correlation component may then input this consumption history into each of the “N” trained models **132**. The models **132** may each output score data indicating, for each of multiple content items associated with a time range corresponding to the individual model **132**, a correlation between the content item and the consumption history of the user. As described above, a higher score may indicate a greater correlation. In some instances, each time range associated with a particular model is exclusive, such as in the example where a first model is associated with content items released within the last six months, a second model is associated with content items released between six and twelve months prior, and so forth. In other instances, the time ranges are not exclusive, as in the example where a first model is associated with content released within the last six months, a second model is associated with content items released within the last year (including the prior six months), and so forth.

The biasing component **124**, meanwhile, may function to bias each of the “N” score data generated by the correlation component **122** according to “M” number of values of a particular parameter. This parameter may comprise a level of a popularity, a cost, a release date, or the like. In some instances, the parameter comprises a level of popularity, as indicated by the popularity data **126**.

For example, the biasing component may determine, for each content item associated with each score data output by each model, whether to boost or penalize the individual score for “M” different values of the popularity. In one example, the biasing component **124** boosts or penalizes the score for each content item according to five different values of popularity, ranging from very popular to very unpopular (or, rather, very “niche”). After doing so, the biasing component may have generated a matrix or other representation of scores, one implementation of which is illustrated and described with reference to FIG. 3.

After the matrix or other representation of scores has been generated, the computing system **106** may send this “result data” back to the client device **102** that initially requested the recommendations. The interpolation component **140** may then determine an order in which to present recommended content items based at least in part on the received result data, as well as the values of the UI controls **112(1)** and **112(2)**. For example, for any given set of values of these two controls, the interpolation component **140** may identify the nearest one or more score data from the result data and may interpolate scores for content item(s) based on the identified score data. The interpolation component **140** may employ bilinear interpolation and/or any other method of interpolation. Furthermore, as the user provides input via the UI control **112(1)** and/or **112(2)**, the interpolation component **140** may perform a new interpolation and update the UI **110** to indicate the new recommendation(s).

FIG. 2 illustrates an example process **200** of operations for generating machine-learning models, applying them to history data of a particular user, sending generated result data to a client device of the user, and using this result data and received user input to determine order(s) in which to present content-item recommendations to the user. While this process **200** is described with reference to the environment **100** of FIG. 1, it may apply to other environments.

At an operation **202**, the computing system **106** may train one or more machine-learning models using any of the techniques described above. In some instances, the operation **202** includes training “N” number of models, each associ-

ated with a particular value of a parameter, such as a particular time window during which a content item was released.

In one example, the computing system **106** may begin the training operation by accessing respective consumption histories (e.g., gameplay histories, moving histories, etc.) of multiple users associated with the computing-system community. For example, the computing system may access a gameplay history of a first user of a video-game-service community and may use this history for training the first model. For example, if the first model is for use in recommending game titles that have been released within the past six months, the computing system may select, from the gameplay history of the user, one game title that has been released within the past six months and may indicate to the first model that this selected game title is to be the output of the classifier. That is, the computing system **106** may identify a game title or other content item having the value of the parameter for which the model is being trained.

Thereafter, the computing system **106** may provide the entire gameplay history of the user as input to the first model and the selected game title as output of the first model for training one or more internal layers of the first model. For instance, if the first model comprises an artificial neural network, then selecting the game title as output and providing information regarding the gameplay history of the user as input may be effective to train one or more internal layers of the neural network. It is to be appreciated that the input and output of the neural network, and other types of models that may be used, may comprise features of the corresponding game titles or other content items, such as release data, genre, game-type, and/or the like.

In some instances, the computing system **106** may take into account the amount of time that the user has played each game in her history in selecting the game title to be used as output for the machine-learning model. For example, the computing system may initially determine an amount of time that the user has played each individual game title referenced in the gameplay history of the user. This may be determined either absolutely (e.g., in terms of time played) or relatively (e.g., in terms of time played relative to other users that have also played the game title). This information may be applied as a weight to the individual game title or inputted as information for the models. After applying a weight to each game title associated with the gameplay history or inputting this information to the models (after, potentially, filtering out some game titles due to a very small amount of gameplay by the user), the computing system may then select one (or more) game title to be used as output by the first model. By doing so, the computing system increases the chances that the games played the most by the user are selected as the output of the first model, thus training the internal layer(s) of the first model in a way that results in the model being more likely to select games of interest to the user.

In addition to training a first model, in some instances the computing device may also train multiple other models each associated with a different value of the parameter, such as a model for recommending games released within the last year, a model for recommending games released within the last three years, and so forth. In each instance, the computing system may access a gameplay or other consumption history of a user and select, as the output to the particular model, a game title that is both indicated in the gameplay history and that was released within the amount of time associated with the particular model (e.g., within one year, three years, etc.). Furthermore, while the above example describes accessing

a single gameplay history for training each of the multiple machine-learning models, it is to be appreciated that the computing system may analyze thousands or more of the gameplay histories for training the models.

At an operation 204, the client device 102 receives a request from a user for one or more content-item recommendations. This may include a user using the client device 102 to navigate to a site provided by the computing system 106, interacting with a dedicated application generated by the computing system 106, or the like.

At an operation 206, the computing system 106 receives the request and, at an operation 208, generates a set of “N” correlation scores using the “N” trained models. As described above, this may include identifying a consumption history of the requesting user and inputting the consumption history into each of the “N” trained models to receive “N” sets of correlation scores.

At an operation 210, the computing system 106 modifies each score within the “N” sets of correlation scores using one or more biasing factors. For example, for each biasing factor, the computing system 106 may use “M” number of values of the biasing factor to generate M×N number of sets of now-biased correlation scores. Of course, if multiple biasing factors are used, then a greater number of sets may be generated. For example, if two biasing factors are used, then the computing system 106 may generate M×M×N number of biased scores. As described above, in some instances, the biasing factor may correspond to a level of popularity of each content item. Further, it is also to be appreciated that the computing system 106 may generate the correlation scores and the biased correlation scores nonlinearly. Further, any number of values may be computed.

At an operation 212, the computing system sends this “result data” to the client device 102, which receives the result data at an operation 214. At an operation 216, the client device 102 presents one or more content-item recommendations to the user based on the received result data. In some instances, the initial recommendations may be based on the current setting of the values of any parameters having values that may be set by the user.

At an operation 218, the client device 102 receives input from the user. This may include altering a value of one or more parameters, such as a desired release data, a desired popularity level, or the like. At an operation 220, the client device 102 using the input data and the result data to interpolate scores for the content items. That is, the client device 102 generates a recommendation order using the result data and the received input. Finally, at an operation 222, the client device 102 may re-present the recommendations, albeit in the newly computed order.

FIG. 3 illustrates example components, and interplay thereof, for generating the models, biasing the resulting scores, applying the models to history data of a particular user, and using this result data and received user input to determine order(s) in which to present content-item recommendations to the user. It is to be appreciated that while FIG. 3 illustrates one example set of components for performing the described techniques, other components may be used in other instances.

As illustrated, in this example some of the components reside on a server-side 302 and other component(s) reside on a client-side 104. For example, the components on the server-side 302 may reside on the computing system 106, while the component(s) on the client-side may reside on the client device 102.

To begin, history data 130 and one or parameters 306(1), 306(2), . . . , 306(N) may be input into the training compo-

nent 120 for generating the respective models 132(1)-(N). That is, the training component 120 may receive an indication of particular values of the parameter for which the training component 120 is to train a model. For example, the parameter 306(1) may indicate a value of “release date within last six months”, the parameter 306(2) may indicate a value of “release date within last year”, and so forth. Of course, in other instances the parameters may correspond to a price of a content item, a popularity of a content item, and so forth. Furthermore, the history data 130 may comprise consumption history of a population of users as described above.

Also as described above, the training component 120 may input individual consumption histories into a respective machine-learning model for training the model. For example, for the model 132(1) having the parameter value of “release date within last six months”, the training component may input individual consumption histories as input to the model and may designate, from each respective history, a content-item title having been released within the previous six months as output for the model 132(1). The training component 120 may then train one or more internal layers of this model 132(1) using this input and output. The training component 120 may perform this operation for an array of different consumption histories from the history data 130 and for each of the models 132(1)-(N). In addition, the input to the models 32 may include information associated with the user, such as a geographical location associated with the user, demographic information associated with the user, and/or the like. In addition, while a few example inputs to the models have been described, it is to be appreciated that any other type of input may be provided to the models.

After the models 132(1)-(N) have been trained, these models may be used for generating content-item recommendations for users. Upon receiving a request for recommendations from a particular user, a consumption history (stored in the history data 130) associated with the user may input into the correlation component 122, along with a library of available content items 128. The correlation component may input this data to each of the models 132(1)-(N), which may output respective correlation scores 308(1), 308(2), . . . , 308(N). For example, the model 132(1) may output recommendations for content-items that have been released within the last six month and that are tailored to the user based on the consumption history of the user, and so forth.

In some instances, these scores 308(1)-(N) may input into the biasing component 124, which may generate a set of biased scores 310(1), 310(2), . . . , 310(N) for each of the correlation scores 308(1)-(N). That is, the biasing component may select one or more values of a particular biasing factor and may boost or penalize the score of each content item within each of the scores 308(1)-(N) based on how well the content item corresponds to the respective value of the biasing factor. For example, the biasing factor may comprise a popularity of a content item, as indicated by the popularity data 126 which may be input into the biasing component. As described above, this popularity data 126 may comprise sales data, consumption time, or the like.

In one example, the biasing component select “M” number of values of the biasing factor, such as M values of popularity ranging from very popular (e.g., having a very high sales number) to very unpopular or niche (e.g., having a very low sales number). Thus, the biasing component 124 may boost the scores of very popular titles for the former value of the biasing factor and penalize very niche titles, and may do the opposite for the latter value of the biasing. In this example, the biasing component 124 may output a number

of biased scores of “M×N” or any other number, given that the sampling may occur non-linearly. In some examples, the computing system 106 sends these biased scores 310(1)-(N) to the client device 102 associated with the user that initiated the request for the recommendations.

As illustrated, the sets of biased scores 310(1)-(N) may be provided as result data 312 to the client device 102. This result data 312 may be input into the interpolation component 140, potentially along with input data 314 generated in response to user input received at the client device 102. For example, the input data 314 may indicate a specific value of the parameter on which the models 132(1)-(N) and/or a specific value of the biasing factor on which the scores were biased. In response to receiving this input data 314, the interpolation component 140 may determine an order 316 in which to present the recommendations to the user. For example, upon receiving the input data 314, the interpolation component 140 may identify one or more nearest biased scores from the result data and may interpolate between these values if need be. For example, the interpolation component 140 may perform bilinear interpolation or any other type of interpolation for determining the order 316 in which to present the recommendations to the user. Furthermore, as different input data 314 is received (e.g., a different value of the parameter or biasing factor), the interpolation component 140 may re-interpolate the order 316 using the input data 314 and the result data 312 and may present the recommendations in the new order 316 to the user. Furthermore, and as described above, the interpolation component 140 may re-compute different orders using the result data 312 and the input data 314 without interacting with the computing system 106, thus enabling real-time resorting on the recommendations on the display of the client device 102.

FIG. 4 illustrates example result data 312 that the remote computing system 102 may generate and send to a client device 102 of a user. As illustrated, in this example the result data 312 may comprise an M×N matrix of scores, which the client device may interpolate between as needed in response to receiving varying user input. It is to be appreciated, however, that the result data may be sampled in any pattern in parameter space of any number of dimensions. For example, the matrix or other representation of scores may be sampled at different points other than the M×N points illustrated. For instance, the sampling may occur non-linearly. Further, the number of sampled points may vary (e.g., may be more less than M×N number of points). In addition, more than two parameters and/or biasing factors may be used, resulting in higher-dimensionality parameter spaces.

In this example, the result data 312 corresponds to a parameter of a time range in which a content item was released and a biasing factor corresponding to a popularity of a content item. Further, the training component 120 has trained “N” number of values of this parameter, as indicated by the result data having “N” number of values of the parameter (time ranges 1-N). In addition, the biasing component 124 may biased each score output by the correlation component 122 using “M” number of values of the biasing factor (popularity 1-M). Thus, in this example, the result data 312 comprises an M×N number of scores.

In response to a user providing an input changing a value of the parameter (time range) or a value of the biasing factor (popularity), the interpolation component may identify the nearest scores and may perform interpolation to computing interpolated scores, which may be used to determine an order in which to recommend content items to the user. For example, envision that a user operating the client device 102

requests to view recommendations for content items that have been released at a time between “time range 1” and “time range 2” and having a popularity that falls between “popularity 1” and “popularity 2”. In response to receiving this input data, the interpolation component 140 may identify the four scores that are nearest this input, such as “biased score 11”, “biased score 12”, “biased score 21”, and “biased score 22”. The interpolation component may then use these scores to perform bilinear interpolation or any other type of suitable interpolation algorithm to generate an interpolated score. This score may then be used to determine an order in which to present the content-item recommendations to the user. For example, a highest-scoring content item may be presented first, a second-highest-scoring content item may be presented second, and so forth.

FIG. 5 illustrates an example user interface (UI) 500 that a client device 102 may present in response to receiving a request from a user for content-item recommendations. As illustrated, the UI 500 may present the custom list 114 of recommended content items 114(1), 114(2), 114(3), and 114(4), with this list 114 being determined based on a consumption history of the user for which the UI 500 was generated and respective values of each of the UI controls 112(1) and 112(2). For example, the UI control 112(1) may comprise a slider or other type of UI control to enable the user to select a desired level of popularity, while the UI control 112(2) may comprise a slider or other type of UI control to enable the user to select a desired recency of the recommended content items. In some instances, the UI 500 presented to the user may comprise a default value for each of these UI controls 112(1), such as in the middle of each range, or the like. Thus, the scores for the content items and, hence, the list 114 may be generated based on these default values, which may be changed as discussed below with reference to FIG. 6.

In addition, the UI 500 may include a UI control 502 in which a user may request to filter out games having a user-designated tag, as well as a UI control 504 in which a user may request to filter out games not having a user-designated tag. In this example, for instance, the user may type in or select via a drop-down menu a tag in either of the controls 502 or 504. In response to receiving a tag in the UI control 502, the client device 102 may change the list 114 by removing any content items associated with the designated tag. In response to receiving a tag in the UI control 504, meanwhile, the client device 102 may change the list 114 by removing any content items that are not associated with the designated tag.

FIG. 6 illustrates an example UI 600 that the client device 102 may present in response to the user altering a value of a first parameter (in this example, popularity of the content items) and a value of a second parameter (in this example, recency of the content items). For instance, in this example the user of the client device has moved the slider on the left side of the UI further to the right (e.g., to request content items that are less popular) while moving the slider on the right side of the UI to the left (e.g., to request content items that are relatively older). Thus, in response, the interpolation component 114 may computed a new list 602 of content items corresponding to these user inputs. For example, the list 602 may include content items that were released within a relatively large time range and that are not overly popular.

FIG. 7 illustrates a flow diagram of an example process 700 that the remote computing system may employ for generating one or more trained machine-learning models and generating result data for use by a client device in providing content-item recommendations to a user. This

process, and each process described herein, is illustrated as a collection of blocks in a logical flow graph, which represent a sequence of operations that can be implemented in hardware, software, or a combination thereof. In the context of software, the blocks represent computer-executable instructions that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular abstract data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described blocks can be combined in any order and/or in parallel to implement the processes.

In some instances, the computing system 106 may be configured to perform some or all of the operations, although in other instances other devices may additionally or alternatively perform some or all of the operations.

At an operation 702, the computing system 106 may train a machine-learning model. As illustrated, this operation may comprise a series of sub-operations. For example, training the model may include, at a sub-operation 702(1), determining a consumption history of one or more users. As described above, this could include a gameplay history, a consumption history of movies or music, and/or the like. A sub-operation 702(2), meanwhile, may comprise selecting a title from each of the histories as output for the model that is to be trained. In some instances, this selection may be made: (1) from those content items in the history that meet the value of the parameter associated with the model currently being trained, and (2) with reference to how much each content item was consumed by the user. For example, the computing system 106 may apply a larger weight to those content items consumed by the respective user more greatly, thus increasing the likelihood that these content items will be selected as output. Finally, at a sub-operation 702(3), the computing system 106 may train one or more internal layers of the model using each consumption history as input and each selected title as output. These sub-operations may be performed serially for each of multiple consumption histories.

At an operation 704, the process 700 determines whether there is an additional model to be trained. For example, if the process 700 is train “N” number of models, the operation 704 represents determining whether any of the N models are still to be trained. If so, then the process 700 returns to the operation 702. If not, then the process 700 proceeds to an operation 706. At the operation 706, the computing system 106 may receive, from a client device of a user, a request for content-item recommendations. At an operation 708, the computing system 106 determines a consumption history of the user and, at an operation 710 inputs this history into a trained model.

At an operation 712, the computing system 106 generates, as output from the trained model, score data indicating a correlation between the consumption history of the user and each respective content item. In some instances, a relatively higher score corresponds to a relatively higher level of correlation. At an operation 714, the computing system 106 modifies each score of the score data based on a particular value of a basing factor, as described above. An operation 716 represents determining whether an additional value of a basing factor is to be used to biased the scores. For example, if the computing system 106 is to bias the score data using “M” number of values of the basing factor, then the operation 716 represents determining whether any of the

“M” values are still to be used to bias the scores. If so, then the process 700 returns to the operation 714 for a new value of the basing factor.

If not, then the process 700 proceeds to an operation 718, which represents determining whether an additional model is to be used to generate score data for the user. For example, if the computing system 106 is configured to use “N” number of models for generating score data that is tailored to the user, then the operation 718 represents determining whether any of the “N” models are still to be used to generate the score data for the user. If so, then the process 700 returns inputting the history data of the user into the next model at the operation 710. If not, then the process 700 proceeds to send the result data (e.g., the matrix or other representation of biased score data) to the client device for enabling the presentation of recommendations to the user.

FIG. 8 illustrates a flow diagram of an example process 800 that a client device of a user may employ for determining an order in which to present one or more content items to the user. While this process 800 is described as being performed on the client device 102, it is to be appreciated that the process 800 may be additionally or alternatively be performed on other devices.

At an operation 802, the client device 102 receives score data indicating a correlation between consumption history of the user of the client device each of multiple content items. In some instances, the score data may comprise the result data 312 discussed above. At an operation 804, the client device 102 receives an input from the user specifying at least one of a value of a parameter on which the model was trained or a value of the basing factor on which the score data was biased. For example, the input may specify a desired recency of the content-items to be recommended (e.g., a time range associated with their respective release dates), a desire cost range of the content items, a desired popularity level of the content items, and/or the like. At an operation 806, the client device 102 determines an order in which to present the recommended content items based at least in part on the score data and the received input. For example, the client device 102 may perform interpolation on the score data based on the value of the parameter and/or basing factor specified by the user.

At an operation 808, the client device 102 presents the content-item recommendations in the determined order and, at an operation 810, determines whether an additional input is received from the user. If so, then the process 800 proceeds to again determine an order in which to (re-) present the content-item recommendations to the user based on the input. If not, then the process 800 ends at an operation 812.

FIG. 9 is a diagram illustrating an example environment 900 that includes a remote computing system configured to train and use machine-learning model(s) for determining which content-item videos to present to a particular user, and for creating a video comprised of these determined videos. In some instances, the remote computing system trains multiple models and, thereafter, inputs a consumption history of a user into each of the models for generating score data indicating a correlation between each of multiple content-item titles and the consumption history. A broadcast component then generates a video for presentation to the user, with the video comprising trailers (or other videos) associated with the highest-correlated content items.

In some instances, the environment includes several of the components described above with reference to the environment 100, such as the computing system 106. In this example, however, the computing system 106 may store, in

memory, one or more content-item videos, such as “trailers” associated with these content items. In some instances, publishers of these content items provide these trailers, while in other instances the computing system 106 may generate these trailers based on popular portions of the content items, randomly selected portions, or the like.

In addition, the computing system 106 may store a ranking component 904 and a broadcast component 906. The ranking component 904 may function to determine which of the content-item videos to present to a particular user, such as a user associated with the client device 102, and an order in which to present these videos. For example, the ranking component 904 may input, into one or more of the trained models 132(1)-(N), information associated with a consumption history of the user, as described above. The ranking component 904 may receive, as output of the model(s), a correlation between each content item, of the available content items 128, and the consumption history of the user. The ranking component 904 may then rank these correlations to determine the highest-ranking content items. For example, the ranking component 904 may determine the ten highest-ranking content items for the user (that is, the ten content items having the greatest degree of correlation to the consumption history of the user.

The ranking component 904 may then provide an indication of these content items and, potentially, the ranking of these items, to a broadcast component 906. The broadcast component 906 may then generate a video for presentation to the user based on a compilation of the content-item videos corresponding to the highest-ranked content items. That is, the broadcast component 906 may identify, from the content-item videos 902, the trailers (or other videos) corresponding to the ten (or other number) highest-ranked content items for the particular user. The broadcast component 906 may then present the video to the user according to the determined order.

FIG. 9, for instances, illustrates that the broadcast component 906 may send a video 908 to the client device for presentation on the display 104. The video 908 may present the identified content-item videos according to the order determined by the ranking component 904. For example, a trailer for a highest-ranked content item may be presented first, a trailer for a second-highest-ranked content item may be presented second, and so forth. In some instances, the broadcast component 906 presents the video by acquiring relatively small increments of the each content-item video (e.g., two seconds) and send this increment to the client device 102 for presentation on the display 104. By doing so, the broadcast component 906 is able to present, to the user, what appears to be a single, static video of multiple content-item videos but is, in fact, a video being created on-the-fly by the broadcast component 906.

In some instances, the broadcast component 906 may also present, atop the video 908, one or more interactive controls 910 with which the user of the client device 102 may interact. These controls 910 may include, for example, selectable icons, such as icons to cause the client device 102 to present a webpage showing additional information regarding the content item, a webpage for the purchasing the content item, and/or the like. In addition, these controls 910 may enable a user to provide feedback regarding the video selected for the user, the content item associated with the currently presented video, or the like. Further, the broadcast component 906 may present video controls that enable the user to perform operations on the video, such as fast-forward, rewind, pause, and the like.

Furthermore, in some instances the controls 910 enable a user to specify desired values of certain parameters. For example, and as described above, the user may select a desired level of popularity of content items, a desired release date (e.g., a desired time range), or the like. Upon receiving this information, the ranking component 904 may utilize one or more different models for re-ranking the content items, with the newly identified content items being provided to the broadcast component 906 for generating a new video taking into account the specified parameters.

Furthermore, in some instances the broadcast component 906 may present one or more “intro” or “outro” videos at the beginning, end, and/or mixed in with the ranked content-item videos. For example, if the broadcast component 906 determines to present ten trailers to a particular user, the broadcast component may present, prior to the ten trailers, an intro video introducing the ten games, with this video including information associated with these games, such as the genre they are associated with, links to each game, and the like. Similarly, the broadcast component 906 may present, at the end of the ten videos, an outro video, indicating information associated with these trailers, with this information being the same or different as that presented in the intro video.

In still another example, these intro and/or outro videos may be interspersed within the ten (or other example number of) trailers. For example, if the first five trailers are associated with a first tag or genre (e.g., adventure) and the second five trailers are associated with a second tag or genre (e.g., puzzle games), then the broadcast component 906 may generate and provide a first intro video prior to presenting the first five trailers, a first outro video after presenting the first five trailers, a second intro video after presenting the first outro video, and a second outro video after presenting the last five trailers. Again, the intro and/or outro videos may include information associated with the corresponding trailers. For instance, the first intro and/or outro videos may indicate the five adventure-related trailers, with information for obtaining these games or the like. Similarly, the second intro and/or outro videos may indicate the five puzzle-related trailers, with information for obtaining these games and/or additional information. The type of content of these intro and/or outro videos may comprise static content (e.g., static images, links, etc.) and/or dynamic content (e.g., voiceover data introducing the games, animations, pre-recorded video of a curator introducing the content, or the like).

FIG. 10 illustrates a flow diagram of an example process 1000 for determining which content-item videos to present to the user and an order in which to present the videos. In some instances, computing system 106 may perform some or all of the operations of the process 1000.

At an operation 1002, the computing system 106 may determine a history of one or more content items consumed by a user, as described above. The content items may represent video games, movies, songs, electronic books, and/or any other type of content item(s). At an operation 1004, the computing system 106 may input, into a trained machine-learning model, data indicative of this history. As noted above, the data generated from the history may comprise any sort of data associated with this history, such as an amount of time that the user consumed the item (in general or relative to other users or a total gameplay time of the user, etc.), a release date of the respective content item, sales data associated with the respective content item, and/or the like.



## 25

At an operation **1006**, the computing system **102** may generate, as output from the first trained machine-learning model, score data indicating, for each of multiple available content items, a score representing a correlation between the history and the respective available content item. At an operation **1008**, the computing system may rank the score data to generate a ranked list of the items, from which the computing system **102** may determine, at an operation **1010**, an ordered set of available content items having a highest correlation with the history.

At an operation **1012**, the computing system **102** may retrieve, for each available content item from the ordered set, a video associated with the respective available content item from the ordered set. For example, the computing system **102** may retrieve trailers associated with a video game, a movie, or the like. At an operation **1014**, the computing system **102** may generate a compilation video using the respective videos. As described above, this may include generating the compilation video according to the order of the ordered set of content items, such that the trailer of a highest ranked content item appears first in the compilation video, a trailer of a second-highest ranked content item appears second, and so forth. The process **1000** may further include presenting this compilation video on a client device, which may comprise streaming the video the client device, downloading the compilation video the client device, and/or the like.

An operation **1016**, meanwhile, represents determining whether feedback regarding the compilation video has been received. As described above, for example, in some instances the computing device may present a control selectable by a user to provide feedback regarding the compilation video to the computing device **102**. If feedback is received, then at an operation **1018** the computing system **102** may determine another ordered set of available content items based on the available score data and the feedback. The process **1000** may then return to retrieving the videos associated with the content items of this second ordered set, and so forth. Therefore, in some instances, the process **1000** can create a loop in which a user may receive and consume a compilation video, interact within or otherwise provide feedback on one or more portions of the compilation video, receive a new compilation video based on the feedback, and so forth. Thus, the process **1000** may enable an “infinite loop” of compilation videos based on recommended content items.

If, however, feedback is not received, the process **100** may determine, at an operation **1020**, whether the process **1000** should be applied for another user. If so, then the process **1000** returns to the operation **1002**, where the process **1000** determines a consumption history for the new user and proceeds to perform the process **1000** using this consumption history. If, however, the process **1000** is not to be performed for an additional user, then the process **1000** ends at an operation **1022**.

Although the subject matter has been described in language specific to structural features, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features described. Rather, the specific features are disclosed as illustrative forms of implementing the claims.

What is claimed is:

1. A method comprising:

determining, by a computing system, a history of one or more game titles played by a user;

inputting, by the computing system and into a trained machine-learning model, data indicative of the history;

## 26

generating, as output from the trained machine-learning model, score data indicating, for each of multiple available game titles, a score representing a correlation between the history and the respective available game title;

ranking the score data to generate a ranked list of the multiple available game titles;

determining, from the ranked list, an ordered set of available game titles having a highest correlation with the history;

retrieving, for each available game title from the ordered set, a trailer associated with the respective available game title from the ordered set; and

generating a compilation video using the respective trailers.

2. The method as recited in claim **1**, wherein the generating comprises generating the compilation video in an order defined by the ordered set such that a trailer of an available game title having a highest correlation with the history appears first in the compilation video.

3. The method as recited in claim **1**, further comprising identifying, from the ordered set, an available game title having a highest correlation with the history, and wherein the generating comprises generating the compilation video using audio associated with the available game title having the highest correlation with the history.

4. The method as recited in claim **1**, wherein the ordered set comprises a first ordered set and the compilation video comprises a first compilation video, and further comprising: receiving feedback data indicating feedback from the user regarding the compilation video;

determining a second ordered set of available game titles based at least part on the score data and the feedback data;

retrieving, for each available game title from the second ordered set, a trailer associated with the respective available game title from the second ordered set; and

generating a second compilation video using the respective trailers associated with the respective available game titles from the second ordered set.

5. A method comprising:

determining, by a computing system, a history of one or more content items consumed by a user;

inputting, by the computing system and into a trained machine-learning model, data indicative of the history;

generating, as output from the trained machine-learning model, score data indicating, for each of multiple available content items, a score representing a correlation between the history and the respective available content item;

ranking the score data to generate a ranked list of the multiple available content items;

determining, from the ranked list, an ordered set of available content items having a highest correlation with the history;

retrieving, for each available content item from the ordered set, a video associated with the respective available content item from the ordered set; and

generating a compilation video using the respective videos.

6. The method as recited in claim **5**, wherein the generating comprises generating the compilation video in an order defined by the ordered set such that a video of an available content item having a highest correlation with the history appears first in the compilation video.

7. The method as recited in claim **5**, further comprising identifying, from the ordered set, an available content item

27

having a highest correlation with the history, and wherein the generating comprises generating the compilation video using audio associated with the available content item having the highest correlation with the history.

8. The method as recited in claim 5, wherein the ordered set comprises a first ordered set and the compilation video comprises a first compilation video, and further comprising: receiving feedback data indicating feedback from the user regarding the compilation video; determining a second ordered set of available content items based at least part on the score data and the feedback data; retrieving, for each available content item from the second ordered set, a video associated with the respective available content item from the second ordered set; and generating a second compilation video using the respective videos associated with the respective available content items from the second ordered set.

9. The method as recited in claim 5, wherein the user comprises a first user, the score data comprises first score data, the ranked list comprises a first ranked list, the ordered set comprises a first ordered set, and further comprising: determining, by the computing system, a history of one or more content items consumed by a second user; inputting, by the computing system and into the trained machine-learning model, data indicative of the history of the one or more content items consumed by the second user; generating, as output from the first trained machine-learning model, second score data indicating, for each of the multiple available content items, a score representing a correlation between the history of the one or more content items consumed by the second user and the respective available content item; ranking the second score data to generate a second ranked list of the multiple available content items; determining, from the second ranked list, a second ordered set of available content items having a highest correlation with the history of the one or more content items consumed by the second user; retrieving, for each available content item from the second ordered set, a video associated with the respective available content item from the second ordered set; and generating a compilation video using the respective videos associated with the respective available content item from the second ordered set.

10. The method as recited in claim 5, further comprising presenting, with the compilation video, a control selectable by the user to acquire one or more content items of the ordered set.

11. The method as recited in claim 5, wherein the one or more content items comprise one or more video games, one or more movies, one or more songs, or one or more electronic books.

12. The method as recited in claim 5, wherein the retrieving comprises retrieving, for each available content item, a trailer associated with the respective available content item from the ordered set, and the generating comprises generating the compilation video using the respective trailers.

13. A computing system comprising:

one or more processors; and  
one or more computer-readable media storing computer-executable instructions that, when executed, cause the one or more processors to perform acts comprising:  
determining a history of one or more content items consumed by a user;

28

inputting, into a trained machine-learning model, data indicative of the history;  
generating, as output from the trained machine-learning model, score data indicating, for each of multiple available content items, a score representing a correlation between the history and the respective available content item;  
ranking the score data to generate a ranked list of the multiple available content items;  
determining, from the ranked list, an ordered set of available content items having a highest correlation with the history;  
retrieving, for each available content item from the ordered set, a video associated with the respective available content item from the ordered set; and  
generating a compilation video using the respective videos.

14. The computing system as recited in claim 13, wherein the generating comprises generating the compilation video in an order defined by the ordered set such that a video of an available content item having a highest correlation with the history appears first in the compilation video.

15. The computing system as recited in claim 13, further comprising identifying, from the ordered set, an available content item having a highest correlation with the history, and wherein the generating comprises generating the compilation video using audio associated with the available content item having the highest correlation with the history.

16. The computing system as recited in claim 13, wherein the ordered set comprises a first ordered set, the compilation video comprises a first compilation video, and the acts further comprise:

receiving feedback data indicating feedback from the user regarding the compilation video;  
determining a second ordered set of available content items based at least part on the score data and the feedback data;  
retrieving, for each available content item from the second ordered set, a video associated with the respective available content item from the second ordered set; and  
generating a second compilation video using the respective videos associated with the respective available content items from the second ordered set.

17. The computing system as recited in claim 13, wherein the user comprises a first user, the score data comprises first score data, the ranked list comprises a first ranked list, the ordered set comprises a first ordered set, and the acts further comprise:

determining a history of one or more content items consumed by a second user;  
inputting, into the trained machine-learning model, data indicative of the history of the one or more content items consumed by the second user;  
generating, as output from the first trained machine-learning model, second score data indicating, for each of the multiple available content items, a score representing a correlation between the history of the one or more content items consumed by the second user and the respective available content item;  
ranking the second score data to generate a second ranked list of the multiple available content items;  
determining, from the second ranked list, a second ordered set of available content items having a highest correlation with the history of the one or more content items consumed by the second user;

retrieving, for each available content item from the second ordered set, a video associated with the respective available content item from the second ordered set; and generating a compilation video using the respective videos associated with the respective available content item from the second ordered set. 5

**18.** The computing system as recited in claim **13**, the acts further comprising presenting, with the compilation video, a control selectable by the user to acquire one or more content items of the ordered set. 10

**19.** The computing system as recited in claim **13**, wherein the one or more content items comprise one or more video games, one or more movies, one or more songs, or one or more electronic books.

**20.** The computing system as recited in claim **13**, wherein the retrieving comprises retrieving, for each available content item, a trailer associated with the respective available content item from the ordered set, and the generating comprises generating the compilation video using the respective trailers. 15 20

\* \* \* \* \*