

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 June 2002 (13.06.2002)

PCT

(10) International Publication Number
WO 02/47381 A2

- (51) International Patent Classification⁷: **H04N 7/00** **Kaushal**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). **GUTTA, Srinivas**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (21) International Application Number: PCT/EP01/13737
- (22) International Filing Date: 23 November 2001 (23.11.2001) (74) Agents: **CHARPAIL, François** et al.; Internationaal Octrooibureau B.V., Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 09/730,205 5 December 2000 (05.12.2000) US (81) Designated States (*national*): CN, JP, KR.
- (71) Applicant: **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL). (84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).
- Published:
— *without international search report and to be republished upon receipt of that report*

- (72) Inventors: **SCHAFER, James, D.**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). **LEE, Kwok, P.**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). **KURAPATI,** *For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: METHOD AND APPARATUS FOR SELECTIVE UPDATING OF A USER PROFILE

(57) Abstract: A television programming recommender is disclosed that selectively obtains feedback information from a user to update one or more profiles associated with the user. Previously obtained implicit and explicit preferences are utilized to selectively focus the collection of feedback information to further update and refine the implicit and explicit preferences. The present invention obtains feedback from a user in a manner that maximizes the value of the obtained information and improves the performance of the television programming recommender. The present invention automatically requests feedback from the user upon the occurrence of predefined criteria. The nature of the requested feedback, and the manner in which the obtained feedback is used to adjust a profile, can vary.

WO 02/47381 A2

Method and apparatus for selective updating of a user profile

Field of the Invention

The present invention relates to methods and apparatus for making recommendations to a user, such as recommendations of television programming, and more particularly, to techniques for selectively updating the user profiles that are utilized to
5 generate such recommendations.

Background of the Invention

As the number of channels available to television viewers has increased, along with the diversity of the programming content available on such channels, it has become
10 increasingly challenging for television viewers to identify television programs of interest. Historically, television viewers identified television programs of interest by analyzing printed television program guides. Typically, such printed television program guides contained grids listing the available television programs by time and date, channel and title. As the number of television programs has increased, it has become increasingly difficult to effectively
15 identify desirable television programs using such printed guides.

More recently, television program guides have become available in an electronic format, often referred to as electronic program guides (EPGs). Like printed television program guides, EPGs contain grids listing the available television programs by time and date, channel and title. Some EPGs, however, allow television viewers to sort or
20 search the available television programs in accordance with personalized preferences. In addition, EPGs allow for on-screen presentation of the available television programs.

While EPGs allow viewers to identify desirable programs more efficiently than conventional printed guides, they suffer from a number of limitations, which if overcome, could further enhance the ability of viewers to identify desirable programs. For
25 example, many viewers have a particular preference towards, or bias against, certain categories of programming, such as action-based programs or sports programming. Thus, the viewer preferences can be applied to the EPG to obtain a set of recommended programs that may be of interest to a particular viewer.

Thus, a number of tools have been proposed or suggested for recommending television programming. The Tivo™ system, for example, commercially available from Tivo, Inc., of Sunnyvale, California, allows viewers to rate shows using a "Thumbs Up and Thumbs Down" feature and thereby indicate programs that the viewer likes and dislikes, respectively. Thereafter, the TiVo receiver matches the recorded viewer preferences with received program data, such as an EPG, to make recommendations tailored to each viewer.

Implicit television program recommenders generate television program recommendations based on information derived from the viewing history of the viewer, in a non-obtrusive manner. FIG. 1 illustrates the generation of a viewer profile 240 using a conventional implicit television program recommender 160. The implicit viewer profile 140 is derived from a viewing history 125, indicating whether or not a given viewer watched each program. As shown in FIG. 1, the implicit television program recommender 160 processes the viewing history 225, in a known manner, to derive an implicit viewer profile 140 containing a set of inferred rules that characterize the preferences of the viewer. Thus, an implicit television program recommender 160 attempts to derive the viewing habits of the viewer based on the set of programs that the viewer watched or did not watch.

Explicit television program recommenders, on the other hand, explicitly question viewers about their preferences for program features, such as title, genre, actors, channel and date/time, to derive viewer profiles and generate recommendations. FIG. 2 illustrates the generation of a viewer profile 240 using a conventional explicit television program recommender 260. The explicit viewer profile 140 is generated from a viewer survey 225 that provides a rating for each program feature, for example, on a numerical scale that is mapped to various levels of interest between "hates" and "loves," indicating whether or not a given viewer watched each program feature. As shown in FIG. 2, the explicit television program recommender 260 processes the viewer survey 125, in a known manner, to generate an explicit viewer profile 240 containing a set of rules that implement the preferences of the viewer.

While such television program recommenders identify programs that are likely of interest to a given viewer, they suffer from a number of limitations, which if overcome, could further improve the quality of the generated program recommendations. For example, explicit television program recommenders typically do not adapt to the evolving preferences of a viewer. Rather, the generated program recommendations are based on the static survey responses. In addition, to be comprehensive, explicit television program recommenders require each user to respond to a very detailed survey. For example, assuming there are 180

different possible values for the “genre” feature, and the user merely specifies his or her “favorite five genres,” then no information is obtained about the user’s preferences for the other 175 possible genres. Similarly, implicit television program recommenders often make improper assumptions about the viewing habits of a viewer that could have easily been
5 identified explicitly by the viewer.

A need therefore exists for a method and apparatus for updating the user profiles that are utilized to generate the recommendations.

Summary of the Invention

10 Generally, a television programming recommender is disclosed that selectively obtains feedback from a user to update one or more profiles for a given user. Previously obtained implicit and explicit preferences are utilized to selectively focus the collection of feedback information to further update and refine the implicit and explicit preferences. The present invention obtains feedback from a user in a manner that maximizes the value of the
15 obtained information and improves the performance of the television programming recommender. In addition, the present invention reduces the obtrusive nature of the feedback query.

The present invention automatically requests feedback from the user upon the occurrence of predefined criteria. For example, feedback can be requested to update the
20 profile(s) if (i) viewing behavior is inconsistent with information recorded in a profile or with generated program recommendation scores; (ii) a neutral recommendation score (neither a positive or negative recommendation) is generated by an implicit or explicit program recommenders; (iii) conflicting recommendation scores are generated by the implicit and explicit program recommenders; or (iv) any combination of the foregoing. The predefined
25 criteria can be compared in real-time (or offline) to the generated recommendation scores and/or viewing behavior to automatically trigger the request for feedback information.

In addition, the present invention allows the nature of the requested feedback to vary, as well as how such feedback should be employed to update the profile(s). In one implementation, the user is requested to rate a program (i) being watched (or not watched)
30 that appears inconsistent with information in the profile(s) or an assigned program recommendation score, or (ii) has been assigned a neutral or conflicting recommendation score by the implicit and/or explicit program recommenders.

In one embodiment, the requested feedback is stored in a log file, referred to herein as a “feedback request list,” for subsequent presentation to the user. A feedback

control process coordinates the timing and the number of feedback requests that are presented to the user from the feedback request list during a given feedback request session in order to (i) minimize the obtrusive nature of the requests, (ii) maximize the quality of the obtained feedback information, or (iii) a combination of the foregoing.

5 Based on the indicated feedback, the present invention determines whether to adjust the information contained in the explicit or implicit viewer profile (or both), and by how much. The user-supplied program rating that is received in response to the feedback request can be, for example, a score indicating the strength of the user's like or dislike of the program. The user-supplied program rating can be used to update the implicit profile, as if
10 the user had watched the program. In addition, if the user-supplied program rating satisfies predefined criteria, such as exceeding a minimum threshold, the program itself can be added to the explicit profile. In a further variation, the user can have the option of updating any conflicting information in the explicit profile 500 that triggered the feedback request.

 A more complete understanding of the present invention, as well as further
15 features and advantages of the present invention, will be obtained by reference to the following detailed description and drawings.

Brief Description of the Drawings

 FIG. 1 illustrates the generation of an implicit profile using a conventional
20 implicit television program recommender;

 FIG. 2 illustrates the generation of an explicit profile using a conventional explicit television program recommender;

 FIG. 3 is a schematic block diagram of the television program recommender in accordance with the present invention;

25 FIG. 4 is a schematic diagram illustrating the process flow of a television program recommender in accordance with the present invention;

 FIG. 5 is a table illustrating an exemplary explicit viewer profile of FIG. 3;

 FIG. 6 is a table illustrating an exemplary implicit viewer profile of FIG. 3;

 FIG. 7 is a sample table from the profile influence rules database of FIG. 3;

30 FIG. 8 is a flow chart describing the selective profile update process of FIG. 3 embodying principles of the present invention; and

 FIG. 9 is a flow chart describing the feedback control process of FIG. 3 embodying principles of the present invention.

Detailed Description

FIG. 3 illustrates a television programming recommender 300 in accordance with the present invention. As shown in FIG. 3, the television programming recommender 300 evaluates each of the programs in an electronic programming guide (EPG) 310 to
5 identify programs of interest to a particular viewer. The set of recommended programs can be presented to the viewer, for example, using a set-top terminal/television (not shown) using well known on-screen presentation techniques. While the present invention is illustrated herein in the context of television programming recommendations, the present invention can be applied to any automatically generated recommendations that are based on a behavior
10 history, such as a viewing history or purchase history.

FIG. 4 provides a schematic diagram of the television programming recommender 300 from a process point of view. As shown in FIG. 4, each viewer uses an explicit profile interface 450 to rate their preferences for various program features, including, for example, days and viewing times, channels, actors, and categories (genres) of television
15 programs. The user-supplied explicit preferences are used to generate an explicit profile 500, discussed further below in conjunction with FIG. 5. The explicit profile 500 is, in turn, utilized to generate program recommendation scores by an explicit program recommender 460, in a known manner.

Likewise, an implicit profile 600, discussed further below in conjunction with
20 FIG. 6, is derived by a profiler 440 from a viewing history 430, indicating whether or not a given viewer watched programs with each program feature. The viewing history 430 is obtained from a set-top terminal 425 that monitors the viewing behavior of the user. The implicit profile 600 is, in turn, utilized to generate program recommendation scores by an implicit program recommender 470, in a known manner.

25 According to one feature of the present invention, the television programming recommender 300 selectively obtains feedback from a user to update the implicit or explicit viewer profiles 500, 600 (or both) for a given user. Generally, previously obtained implicit and explicit preferences are utilized to selectively focus the collection of feedback information to update such implicit and explicit preferences. Thus, the television
30 programming recommender 300 can obtain feedback from a user in a manner that maximizes the value of the obtained information and thereby improves the performance of the television programming recommender 300, while minimizing the obtrusive nature of the feedback query.

In one implementation, the present invention employs profile influence rules 700, discussed below in conjunction with FIG. 7, during step 475 that are operable to automatically request feedback from the user upon the occurrence of predefined criteria, such as specified events. As discussed below in conjunction with FIG. 7, the established profile influence rules 700 may determine the timing and nature of the feedback that is requested from a user during step 480, and how such feedback should be employed to update the profile(s) 500, 600 during step 485. Based on the indicated feedback, the television programming recommender 300 can determine whether to adjust the information contained in the explicit or implicit viewer profile 500, 600 (or both), and by how much.

As discussed further below in conjunction with FIGS. 8 and 9, the feedback requested during step 480 can be requested immediately upon the detection of an appropriate feedback trigger condition, or the feedback request can be logged in a feedback request list 350 (FIG. 3) for subsequent processing to minimize the obtrusive nature of the requests or to maximize the quality of the obtained feedback information (or both).

The profile influence rules 700 of the present invention may request feedback to update the profile(s) 500, 600, for example, if (i) viewing behavior is inconsistent with information recorded in a profile or with generated program recommendation scores; (ii) a neutral recommendation score (neither a positive or negative recommendation) is generated by an implicit or explicit program recommenders; (iii) conflicting recommendation scores are generated by the implicit and explicit program recommenders; or (iv) any combination of the foregoing. For example, viewing behavior can be inconsistent with profile information or generated program recommendation scores if, e.g., (i) a program is watched having features that do not match the profile(s) 500, 600; a program is watched that was assigned a low program recommendation score; or (iii) a program receives a high program recommendation score but is not watched in favor of one or more program(s) receiving lower program recommendation scores.

As shown in FIG. 4, once the profile influence rules 700 are established, the profile influence rules may be compared in real-time (or offline) during step 475 to the generated recommendation scores and/or viewing behavior, as well as other factors, in order to automatically determine the applicability of one or more of the profile influence rules 700. Each profile influence rule 700 may comprise the predefined criteria specifying the conditions under which the profile influence rule should be initiated, and, optionally, a feedback request command defining the appropriate information that should be requested in order to influence the profile(s).

In the illustrative embodiment described herein, the feedback request command requests the user to rate a program (i) being watched (or not watched) that appears inconsistent with information in the profile(s) 500, 600 or an assigned program recommendation score, or (ii) has been assigned a neutral or conflicting recommendation score by the implicit and/or explicit program recommenders. The feedback request may optionally indicate the program recommendation score assigned to the program and identify one or more program features that significantly contributed to the program recommendation score (for, example, the top-N contributing program features).

The user-supplied program rating that is received in response to the feedback request can be, for example, a score indicating the strength of the user's like or dislike of the program. The user-supplied program rating can be used to update the implicit profile 600 in a well-known manner, as if the user had watched the program. In addition, if the user-supplied program rating satisfies predefined criteria, such as exceeding a minimum threshold, the program itself can be added to the explicit profile 500. In other words, an entry can be added to the explicit profile 500 in the form of {if title="program_name" then assigned score= user-supplied program rating}. In a further variation, the user can have the option of updating any conflicting information in the explicit profile 500 that triggered the feedback request.

Thus, as shown in FIG. 3, the television programming recommender 300 includes a feedback request list 350 which may be, for example, a log file containing a list of feedback requests accumulated by the television programming recommender 300. In addition, the television programming recommender 300 includes the explicit viewer profile 500, the implicit viewer profile 600, each discussed further below in conjunction with FIGS. 5 and 6, respectively, and a profile influence rule database 700, discussed further below in conjunction with FIG. 7, containing the profile influence rules.

In addition, a selective profile update process 800 and a feedback control process 900, are discussed further below in conjunction with FIGS. 8 and 9, respectively. Generally, the selective profile update process 800 compares the profile influence rules 700 to, e.g., the generated recommendation scores and/or viewing behavior, and populates the feedback request list 350 with an appropriate feedback request when a given profile influence rule 700 is triggered. The feedback control process 900 coordinates the timing and the extent of the feedback requests that are presented to the user from the feedback request list 350 during a given feedback session to minimize the obtrusive nature of the requests or to maximize the quality of the obtained feedback information (or both).

The television program recommender 300 may be embodied as any computing device, such as a personal computer or workstation, that contains a processor 315, such as a central processing unit (CPU), and memory 320, such as RAM and ROM. In addition, the television programming recommender 300 may be embodied as any available television program recommender, such as the Tivo™ system, commercially available from Tivo, Inc., of Sunnyvale, California, or the television program recommenders described in United States Patent Application Serial No. 09/466,406, filed December 17, 1999, entitled "Method and Apparatus for Recommending Television Programming Using Decision Trees," (Attorney Docket No. 700772), United States Patent Application Serial No. 09/498,271, filed Feb. 4, 2000, entitled "Bayesian TV Show Recommender," (Attorney Docket No. 700690) and United States Patent Application Serial No. 09/627,139, filed July 27, 2000, entitled "Three-Way Media Recommendation Method and System," (Attorney Docket No. 700913), or any combination thereof, as modified herein to carry out the features and functions of the present invention.

FIG. 5 is a table illustrating an exemplary explicit viewer profile 500. As shown in FIG. 5, the explicit viewer profile 500 contains a plurality of records 505-513 each associated with a different program feature. In addition, for each feature set forth in column 540, the viewer profile 500 provides a numerical representation in column 550, indicating the relative level of interest of the viewer in the corresponding feature. As discussed below, in the illustrative explicit viewer profile 500 set forth in FIG. 5, a numerical scale between 1 ("hate") and 7 ("love") is utilized. For example, the explicit viewer profile 500 set forth in FIG. 5 has numerical representations indicating that the user particularly enjoys programming on the Sports channel, as well as late afternoon programming.

In an exemplary embodiment, the numerical representation in the explicit viewer profile 500 includes an intensity scale such as:

Number	Description
1	Hates
2	Dislikes
3	Moderately negative
4	Neutral
5	Moderately positive
6	Likes
7	Loves

FIG. 6 is a table illustrating an exemplary implicit viewer profile 600 corresponding to the same viewer as the explicit viewer profile 600, discussed above. As shown in FIG. 6, the implicit viewer profile 600 contains a plurality of records 605-613 each associated with a different program feature. In addition, for each feature set forth in column 640, the implicit viewer profile 600 provides the corresponding positive and negative counts, in a known manner, in columns 645 and 650, respectively, indicating the number of times the viewer watched and did not watch, respectively, programs having each feature. For each positive and negative program example (i.e., programs watched and not watched), a number of program features are classified in the user profile 600. For example, if a given viewer watched a given sports program ten times on Channel 2 in the late afternoon, then the positive counts associated with these features in the implicit viewer profile 600 would be incremented by 10, and the negative counts would be 0 (zero). Since the implicit viewing profile 500 is based on the user's viewing history, the data contained in the profile 500 is revised over time, as the viewing history grows.

FIG. 7 illustrates an exemplary table of the profile influence rule database 700 that records each of the profile influence rules that dynamically obtain user feedback and adjust the profile(s) 500, 600, if the predefined criteria for initiating the profile influence rule is satisfied. Each profile influence rule 700 may comprise the predefined criteria specifying the conditions under which the profile influence rule should be initiated, and, optionally, a feedback request command defining the appropriate feedback that should be requested in order to influence the profile(s). In illustrative embodiment, the default feedback request queries the user to rate a program (i) being watched (or not watched) that appears inconsistent with information in the profile(s) 500, 600 or an assigned program recommendation score, or (ii) has been assigned a neutral or conflicting recommendation score by the implicit and/or explicit program recommenders. The feedback request may optionally indicate the program recommendation score assigned to the program and identify one or more program features that significantly contributed to the program recommendation score (for, example, the top-N contributing program features).

As shown in FIG. 7, the exemplary profile influence rule database 700 maintains a plurality of records, such as records 705-709, each associated with a different profile influence rule. For each profile influence rule, the profile influence rule database 700 identifies the rule criteria in field 750. In a further variation of the profile influence rule database 700, an additional field (not shown) can be included to record the corresponding feedback request that should be implemented for a given satisfied rule.

FIG. 8 is a flow chart describing the selective profile update process 800 embodying principles of the present invention. As previously indicated, the television programming recommender 300 implements the selective profile update process 800 to monitor viewing behavior and generated recommendation scores, and determine whether the predefined criteria associated with any profile influence rule is satisfied. As previously indicated, each profile influence rule may comprise (i) predefined criteria specifying the conditions under which the profile influence rule should be initiated, and (ii) a profile feedback request command defining the appropriate response that should be implemented in order to influence the profile(s). The feedback request command may be a query to obtain feedback from the user (that in turn can be used to adjust the information in the profile(s) 500, 600), or an appropriate adjustment to information in the profile(s) 500, 600. Thus, once the predefined criteria of a given profile influence rule is satisfied, the selective profile update process 800 will implement the corresponding profile feedback request command to influence the profile(s) in the desired manner.

Thus, as shown in FIG. 8, the selective profile update process 800 initially stores the profile influence rules in the profile influence rule database 700 during step 805. As previously indicated, the profile influence rules are operable to automatically request feedback from the user upon the occurrence of predefined criteria.

In addition, the selective profile update process 800 receives the viewing behavior and/or generated recommendation scores during step 810. Thereafter, the selective profile update process 800 compares the received viewing behavior and/or generated recommendation score data to the profile influence rules criteria recorded in the profile influence rule database 700 during step 815. It is noted that the comparison performed during step 815 may be executed periodically, continuously, or at irregular intervals.

A test is performed during step 820 to determine if the predefined criteria for at least one profile influence rule is satisfied. If it is determined during step 820 that the predefined criteria for at least one profile influence rule is not satisfied, then program control returns to step 815 to continue evaluating the received viewing behavior and/or generated recommendation score data in the manner described above.

If, however, it is determined during step 820 that the predefined criteria for at least one profile influence rule is satisfied, then an entry is created in the feedback request list 350 containing the corresponding feedback request during step 825. As discussed further below in conjunction with FIG. 9, the frequency with which feedback requests are presented to the user from the feedback request list 350 and the number of feedback requests that are

presented to the user during a given feedback session can be controlled to minimize the obtrusive nature of the requests or to maximize the quality of the obtained feedback information (or both).

For example, in the illustrative embodiment the default feedback request
5 command queries the user to rate a program (i) being watched (or not watched) that appears inconsistent with information in the profile(s) 500, 600 or an assigned program recommendation score, or (ii) has been assigned a neutral or conflicting recommendation score by the implicit and/or explicit program recommenders. The feedback request may optionally indicate the program recommendation score assigned to the program and identify
10 one or more program features that significantly contributed to the program recommendation score (for, example, the top-N contributing program features).

TIMING AND EXTENT OF FEEDBACK REQUESTS

As previously indicated, the frequency with which feedback requests are
15 presented to the user from the feedback request list 350 and the number of feedback requests that are presented to the user during a given feedback session can be controlled to minimize the obtrusive nature of the requests or to maximize the quality of the obtained feedback information (or both).

FIG. 9 is a flow chart describing an exemplary feedback control process 900
20 that coordinates the timing and the extent of the feedback requests that are presented to the user from the feedback request list 350 during a given feedback session to minimize the obtrusive nature of the requests or to maximize the quality of the obtained feedback information (or both). In addition, the feedback control process 900 can improve its current knowledge by learning from the user reaction to each feedback session. As discussed
25 hereinafter, the feedback control process 900 may employ a number of rules that control the timing and the extent of the feedback requests based on situation-defining parameters.

The rules and associated situation-defining parameters might specify, for example, (i) specific times and days when feedback should or should not be requested; (ii) the number of feedback requests to present during a given feedback request session; (iii) the
30 duration of each feedback request session; and (iv) the minimum time that should separate any two feedback request sessions (i.e., a blackout time period). It is noted that times and days in the feedback gathering rules may be specified in terms of absolute values or relative to a current or future time or event, such as the next time the user powers up the device.

As discussed further below, the feedback gathering rules and/or associated situation-defining parameters can be specified, for example, by the user employing a menu-driven interface, or by an expert in human-machine interactions. Furthermore, the feedback gathering rules and/or associated situation-defining parameters can be predefined or
5 dynamically determined, as discussed below. Generally, the feedback gathering rules and associated situation-defining parameters should be informed by research that make the interactions most tolerable to the human participants and most likely to produce good feedback information over time.

Furthermore, the television programming recommender 300 can be initiated
10 with default values for the situation-defining variables based, for example, on user testing research, that can be modified over time in response to the user's reaction to the feedback gathering process.

As shown in FIG. 9, the feedback control process 900 initially determines if there are currently any feedback requests to be processed in feedback request list 350 during
15 step 910. If it is determined during step 310 that there currently are no feedback requests in the feedback request list 350, then program control terminates. If, however, it is determined during step 310 that there are feedback requests in the feedback request list 350, then the feedback control process 900 computes the time to initiate each feedback request during step 920.

20 For example, the computed time can generally be conditioned on the presence of the user(s) associated with the profiles 500, 600. The presence of a user can be determined, for example, using well-known situation-awareness methods, such as cameras or heat sensors, or an inference that the user is present when the device is turned on.

In addition, the number of requests to include in each feedback request session
25 is determined during step 930. If the number of requests to include in the session exceeds the number of requests in the feedback request list 350, each of which can vary with time, then the feedback requests are prioritized during step 940.

A test is performed during step 950 to determine if it is time to initiate a feedback request session. If it is determined during step 950 that it is not time to initiate a
30 feedback request session, then program control returns to step 950 until the indicated time.

If, however, it is determined during step 950 that it is time to initiate a feedback request session, then the feedback request is initiated during step 960. The requested feedback and other situation defining variables, such as a flag indicating not to

query the user for feedback, e.g., when other people are in the room, or when the user is on the phone, are collected during step 970.

Finally, the feedback management rules are updated during step 980 with the new situation defining variables and the appropriate user profile(s) 500, 600 are updated during step 990 with the obtained feedback. It is noted that the appropriate user profile(s) 500, 600 can be updated, for example, in accordance with the techniques described in United States Patent Application Serial No. 09/627,139, filed July 27, 2000, entitled "Three-Way Media Recommendation Method and System," (Attorney Docket No. 700913), assigned to the assignee of the present invention and incorporated by reference herein.

As previously indicated, the situation-defining variables used by the feedback control process 900 to determine the timing and the extent of the feedback requests can be predefined or dynamically determined. In one implementation, the television programming recommender 300 can be initiated with default values or user-specified values indicating how often feedback request session should be scheduled and how many feedback requests the user is willing to process during each feedback request session.

Thereafter, the television programming recommender 300 can employ a trial-and-error process to refine the initial values. For this approach, the situation-defining parameters can be considered random variables with some probability distribution that needs to be learned, or they may be considered fuzzy functions with confidence weightings.

For example, to determine how many shows to offer during a feedback session, a default position may treat this value as a normally distributed random variable with a mean of 10 and a range of +/- 5. During each feedback request session, a random number will be selected from this distribution and that number of feedback requests will be presented to the user. Thereafter, the feedback request session may be terminated by the user in one of three ways: (i) responding to all requests for feedback and then stopping; (ii) terminating the session before responding to all requests (including a refusal to respond to any requests); or (iii) completing all requests for feedback and then requesting additional feedback requests. Whatever the outcome, an observed value of the random variable that the user determined is accumulated. Over time, these accumulated values are used by the television programming recommender 300 to compute a new probability distribution that more accurately reflects the tolerance of the specific user.

These collected observations may be further enhanced by collecting additional variables that characterize the situation when the observed value was collected. For instance, the day and time might be noted. This would permit modeling the number of shows as a

multivariate distribution. A system using this method might learn, for example, that the user is willing to respond to more feedback requests on certain days of the week and/or during certain time periods. Similarly, if the genre of the tuned-in show is also noted, the system may learn that the user is willing to respond to more feedback requests when sitcoms are on
5 than when the News is on. The system may also note the program title, so that it may learn that this user is usually willing to supply feedback at the end of a given program, but not at the beginning and generally not with other programs.

Considering parameters that govern the timing of a feedback request, the same methods can be applied. If technology for segmenting the broadcast is used, such as those
10 techniques described, for example, in United States Patent Application Serial Number 09/532,845, filed March 21, 2000, entitled "System and Method for Automatic Content Enhancement of Multimedia Output Device," assigned to the assignee of the present invention and incorporated by reference herein, then the system might learn that feedback requests are more likely to be accepted if offered during the show broadcast itself.

15 It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention.

CLAIMS:

1. A method for updating a user profile (500, 600), comprising the steps of:
obtaining said user profile (500, 600) indicating preferences of a user;
comparing said preferences to an item selection made by said user;
requesting feedback information from said user if said selection is inconsistent
5 with said preferences; and
updating said user profile (500, 600) with said feedback information.
2. The method of claim 1, wherein said step of requesting feedback further
comprises the steps of placing said feedback in a log file and presenting said feedback
10 request to said user at a later time.
3. The method of claim 2, wherein said feedback requests are presented in a
given session during said step of presenting said feedback request based on a user-specified
parameter.
15
4. The method of claim 1, wherein said feedback requests said user to rate a
selected content item that appears inconsistent with information in said user profile (500,
600).
- 20 5. The method of claim 1, wherein said feedback requests said user to rate a
content item that is not selected that appears inconsistent with information in said user profile
(500, 600).
6. The method of claim 1, wherein said feedback requests said user to rate a
25 content item that has been assigned an inconsistent recommendation score by two
recommenders.
7. A system (300) for updating a user profile (500, 600), comprising:

a memory (320) for storing computer readable code and said user profile (500, 600); and

a processor (315) operatively coupled to said memory (320), said processor (315) configured to:

- 5 obtain said user profile (500, 600) indicating preferences of a user;
 compare said preferences to an item selection made by said user;
 request feedback information from said user if said selection is inconsistent
with said preferences; and
 update said user profile (500, 600) with said feedback information.

10

8. A computer program product for updating a user profile (500, 600),
comprising computer readable code which comprise:

- a step to obtain said user profile (500, 600) indicating preferences of a user;
 a step to compare said preferences to an item selection made by said user;
15 a step to request feedback information from said user if said selection is
inconsistent with said preferences; and
 a step to update said user profile (500, 600) with said feedback information.

1/8

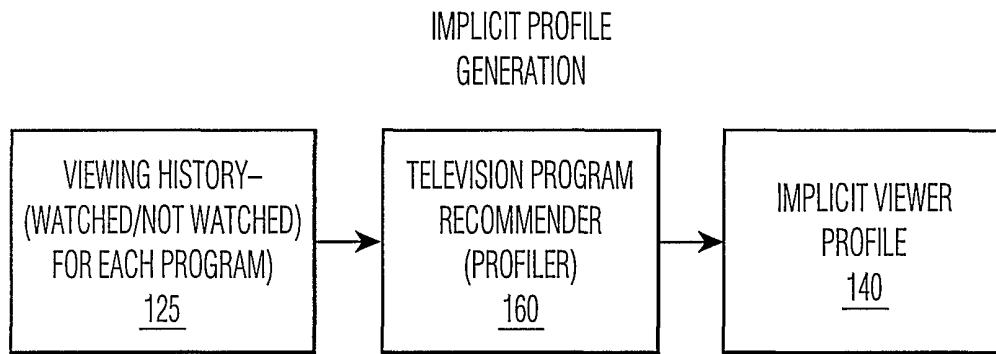


FIG. 1




FIG. 2

2/8

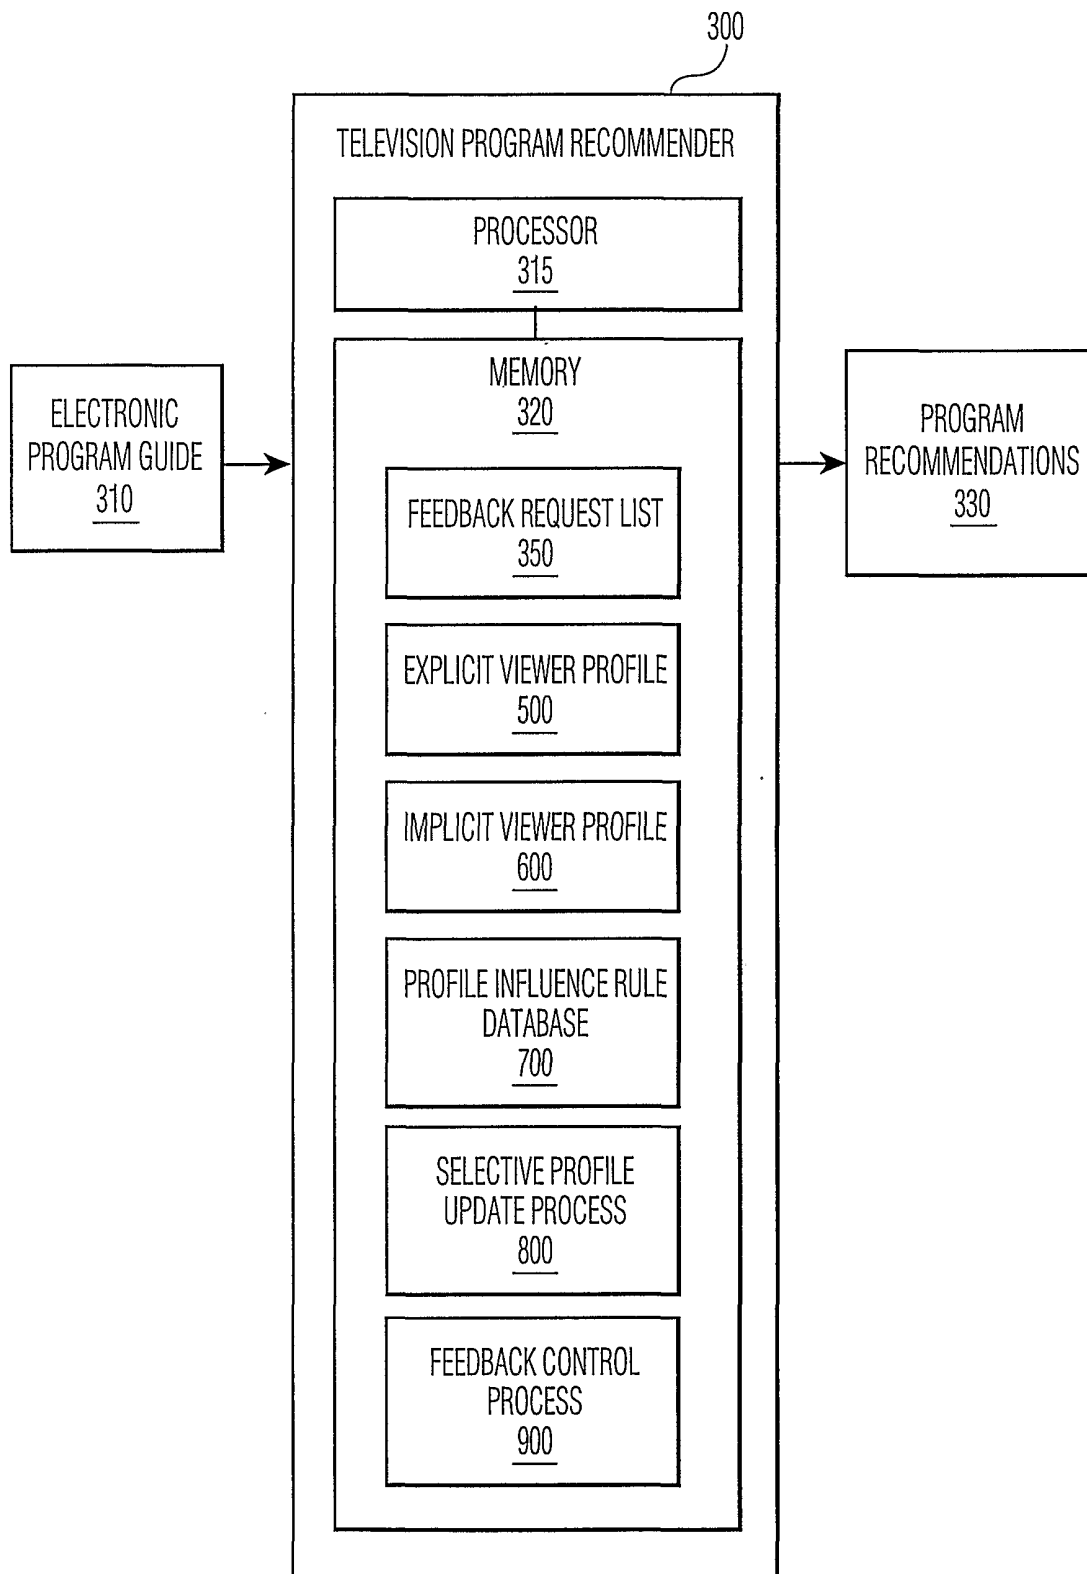


FIG. 3

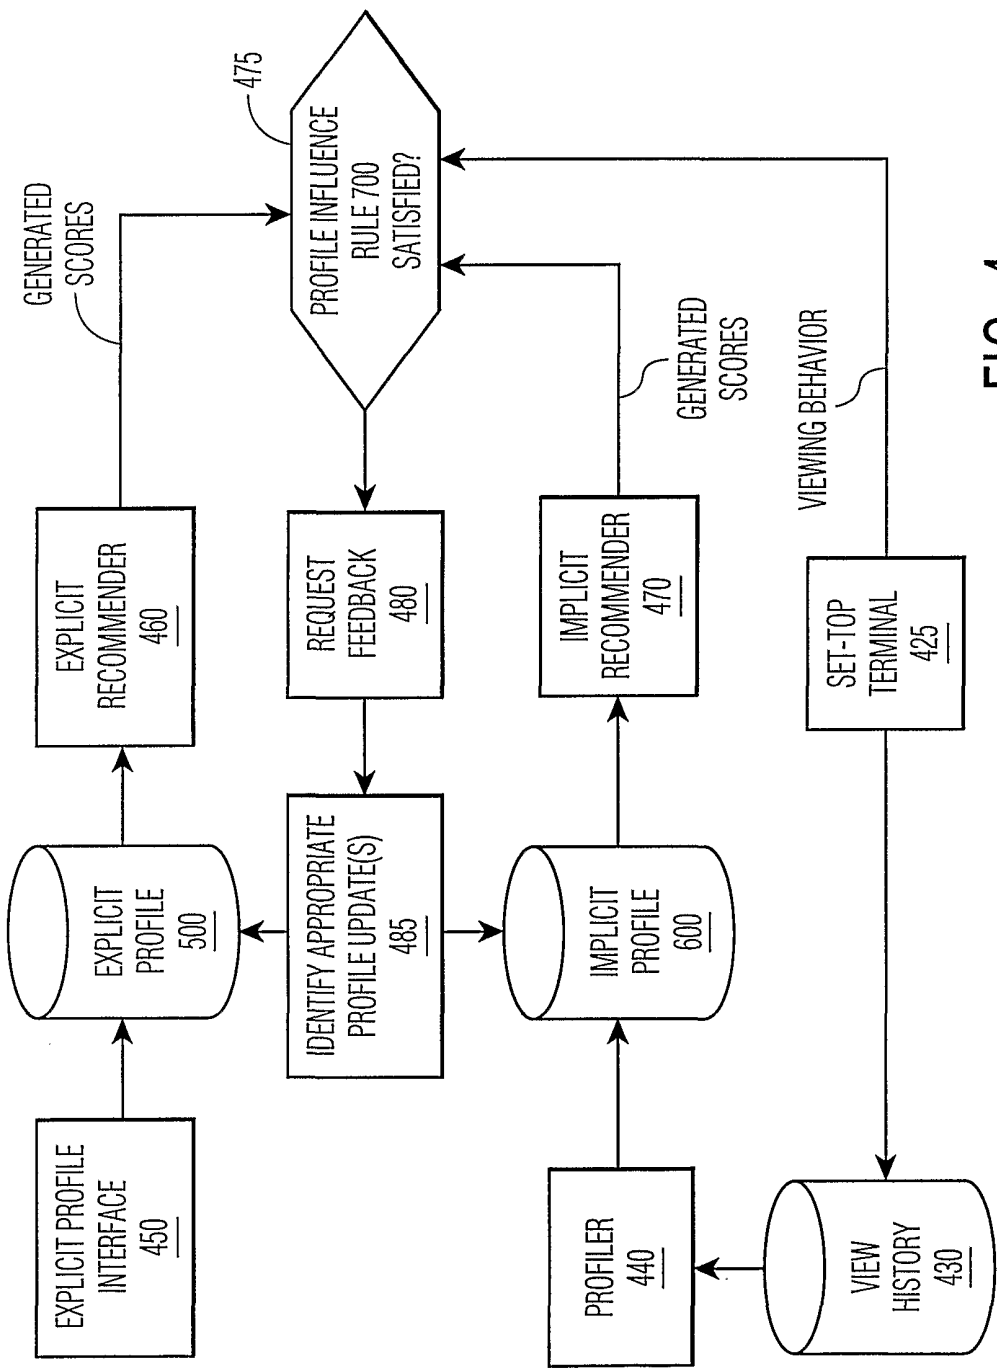


FIG. 4

EXPLICIT VIEWER PROFILE 500

	FEATURE <u>540</u>	NUMERICAL (OR SYMBOLIC) REPRESENTATION <u>550</u>
505	CHANNEL 2	3
506	CHANNEL 4	4
507	CHANNEL 7	3
508	
509	SPORTS CHANNEL	7
....	MUSIC CHANNEL	2
	
510	MORNING PROGRAMS	1
511	EARLY AFTERNOON PROGRAMS	3
512	LATE AFTERNOON PROGRAMS	7
513	EVENING PROGRAMS	5
	...	

FIG. 5

5/8

IMPLICIT VIEWER PROFILE 600

	FEATURE <u>640</u>	POSITIVE COUNTS <u>645</u>	NEGATIVE COUNTS <u>650</u>
605	TOTAL PROGRAMS	45	45
606	CHANNEL 2	10	0
607	CHANNEL 4	3	1
608	CHANNEL 7	4	4
		
609	SPORTS CHANNEL	10	0
....	MUSIC CHANNEL	1	0
		
610	MORNING PROGRAMS	2	2
611	EARLY AFTERNOON PROGRAMS	1	2
612	LATE AFTERNOON PROGRAMS	10	0
613	EVENING PROGRAMS	6	4
	...		

FIG. 6

PROFILE INFLUENCE RULE DATABASE 700

	RULE CRITERIA <u>750</u>
705	PROGRAM IS WATCHED HAVING FEATURES THAT DO NOT MATCH PROFILE(S)
706	PROGRAM IS WATCHED HAVING A LOW PROGRAM RECOMMENDATION SCORE
707	PROGRAM RECEIVES HIGH SCORE BUT IS NOT WATCHED IN FAVOR OF PROGRAM(S) RECEIVING LOWER PROGRAM RECOMMENDATION SCORE(S)
708	NEUTRAL RECOMMENDATION SCORE IS GENERATED BY PROGRAM RECOMMENDER
...	
709	CONFLICTING RECOMMENDATION SCORES ARE GENERATED BY IMPLICIT AND EXPLICIT PROGRAM RECOMMENDERS

FIG. 7

7/8

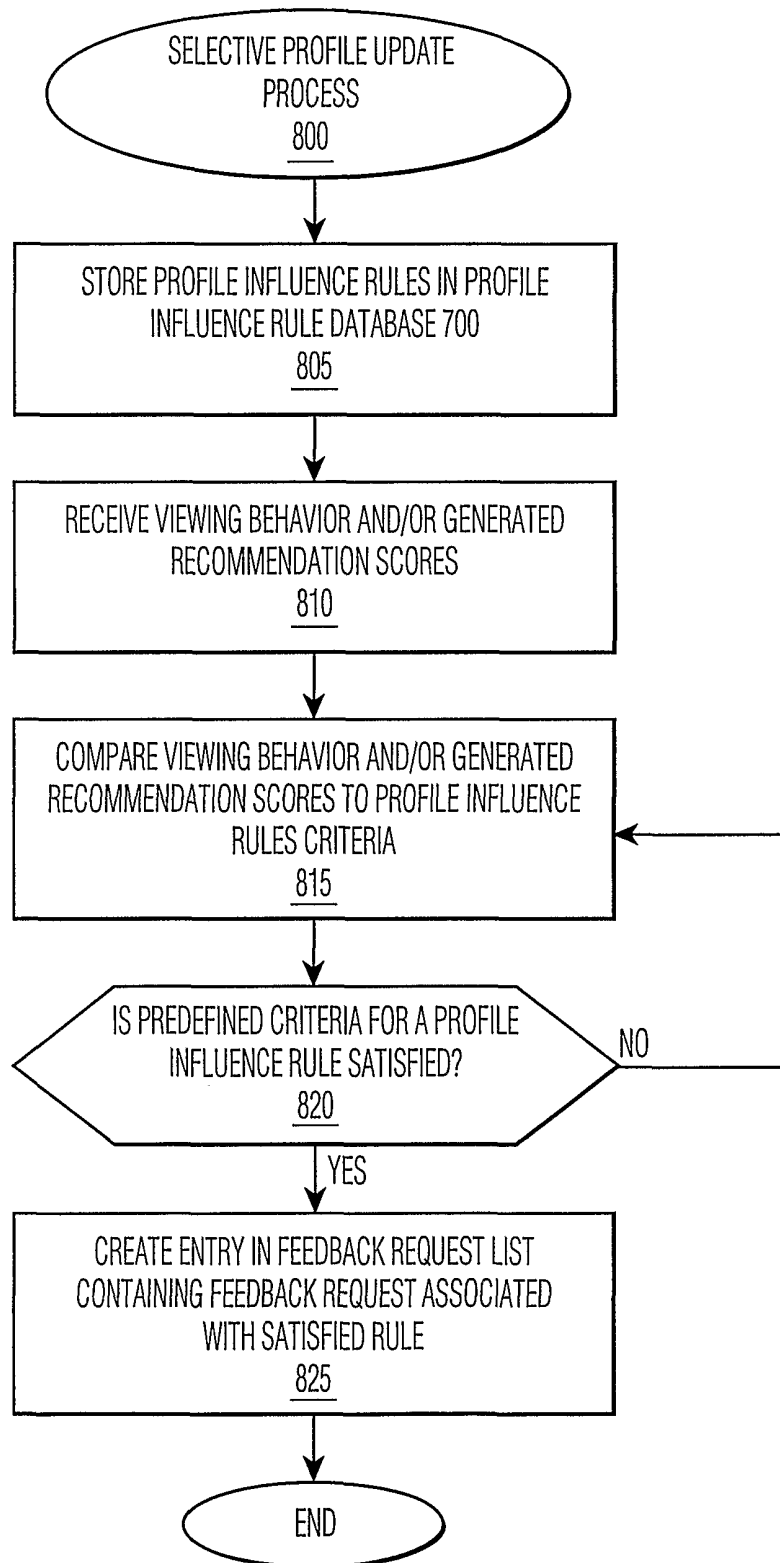


FIG. 8

8/8

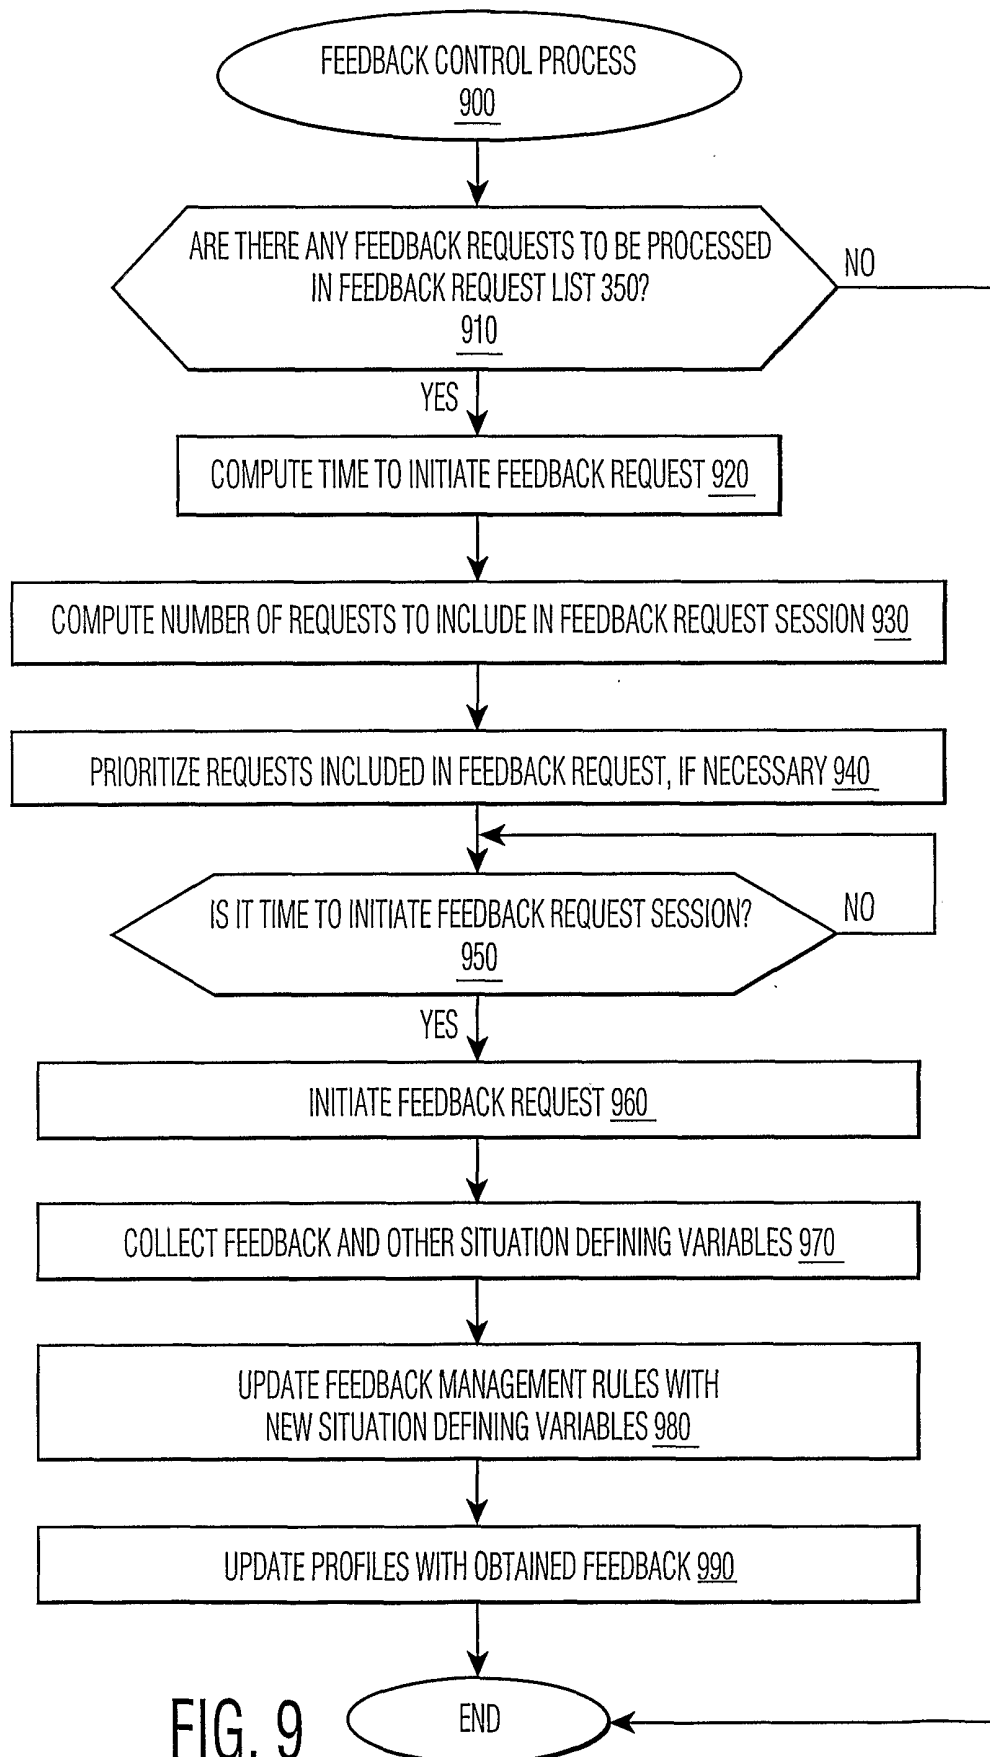


FIG. 9