



- (51) International Patent Classification:  
G06F 9/46 (2006.01)
- (21) International Application Number:  
PCT/US2014/058211
- (22) International Filing Date:  
30 September 2014 (30.09.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: NUANCE COMMUNICATIONS, INC.  
[US/US]; One Wayside Road, Burlington, Massachusetts  
01803-4613 (US).
- (72) Inventors: DUCATELLE, Frederick; Haagstraat 4, B-  
9700 Oudenaarde (BE). GRÖBER, Marcus; Am Bullen-  
berg 15, 38518 Gifhorn (DE). MARTENS, Gaetan;  
Torenstraat 6, B-9500 Zarlanninge (BE).
- (74) Agents: DURKEE, Paul, D. et al.; Daly, Crowley, Mof-  
ford & Durkee, LLP, 354A Tumpike Street, Suite 301A,  
Canton, Massachusetts 02021 (US).
- (81) Designated States (unless otherwise indicated, for every  
kind of national protection available): AE, AG, AL, AM,  
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,  
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,  
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,  
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,  
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,  
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,  
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,  
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States (unless otherwise indicated, for every  
kind of regional protection available): ARIPO (BW, GH,  
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,  
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,  
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,  
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,  
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,  
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,  
GW, KM, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

— of inventorship (Rule 4.17(iv))

**Published:**

— with international search report (Art. 21(3))

(54) Title: METHODS AND APPARATUS FOR MODULE ARBITRATION

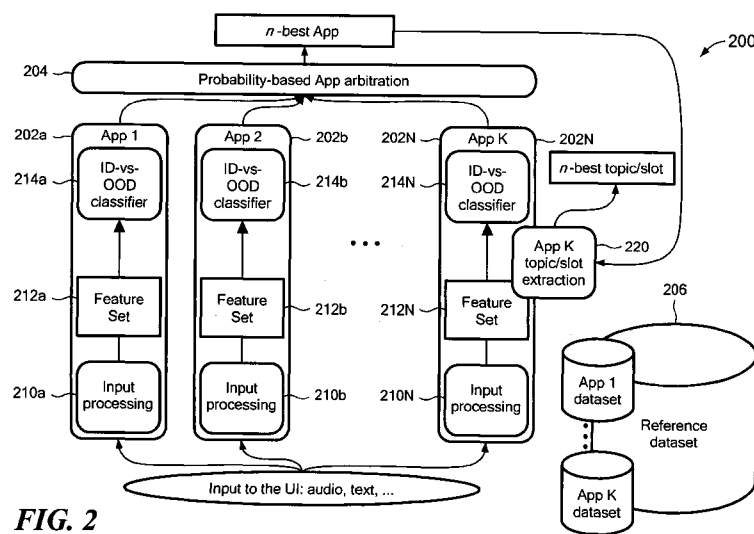


FIG. 2

(57) Abstract: Methods and apparatus to process a user input on independent applications that provide classifier outputs to an arbitration module, which selects one of the application to respond to the user input. The classifier outputs include a probability that the user input is in domain for the application functionality.

WO 2016/053276 A1

## METHODS AND APPARATUS FOR MODULE ARBITRATION

## BACKGROUND

As is known in the art, in some systems a single speech recognition UI system can process  
5 user utterances for a number of applications. An arbitration module can select one of the  
applications to respond to the user utterance. In conventional systems, a multi-class  
classifier performs arbitration for the various applications. However, such a multi-class  
classifier is trained for known applications. Thus, the multi-class classifier is not readily  
extensible for independently developed applications.

10

## SUMMARY

Embodiments of the invention provide methods and apparatus for a single user interface  
(UI) system having multiple modules, where each module implements a subset of the UI  
functionality of the full system. A module typically corresponds to an application  
15 supported by the system. At runtime, arbitration is applied by one arbitration system that  
analyzes the request passed to the UI system by a user, and depending on the requested  
functionality, the UI system passes control to the relevant module(s). Embodiments of the  
invention are directed to UI systems based on speech recognition, where user inputs are  
voice commands, but it could also be used for other UI systems that deal with possibly  
20 ambiguous input.

In general, modules operate independently from each other and can be developed  
independently from each other so that the framework is readily extensible with other  
modules without the need to change existing modules or to change the arbitration system.  
25 Embodiments of the invention provide a UI system from an undefined number of  
independently developed UI modules to which control can be given.

In illustrative embodiments, the user input to the UI system is passed to the respective  
modules, possibly after some first processing (e.g., applying automated speech recognition  
30 on an audio utterance). Each module outputs a probability to express a level of relevance  
the input has to the functionality of the respective module. The modules operate  
independently and provide input to an arbitration module, which can process input from  
any practical number of modules. The role of the arbitration module is to gather the output

of the respective modules, and propose one or more of them as the ones that should be allowed to process the command from the user.

Embodiments of the invention are applicable to various applications where an extensible  
5 framework is desirable, such as connecting different devices that contain one or more  
speech-enabled applications to a head unit in a car, for example, and/or integrated systems  
that allow different ASR/NLU technologies to share a single microphone in a system.  
Further applications include an extensible system that provides an ASR/NLU system that  
is extensible with other modules downloaded and executed on different devices. It is  
10 understood that the modules do not need to reside on different devices. Also on a single  
device, with the provided flexibility, modules can be developed at different points in time  
and by independent teams, they can be chosen and combined freely by the user, etc.

In one aspect of the invention, a method comprises: receiving, in response to a user  
15 utterance, a first input from a first classifier of a first application, the first input including a  
probability of at least a portion of the user utterance being in domain for the first  
application, wherein in domain for the first application means that the first application  
contains the functionality required to respond to the utterance; receiving, in response to the  
user utterance, a second input from a second classifier of a second application, the second  
20 input including a probability of at least a portion of the user utterance being in domain for  
the second application; and performing arbitration on the user utterance using the first and  
second inputs to prioritize a first one of the first and second applications for responding to  
the user utterance.

25 The method can further include one or more of the following features: performing  
automated speech recognition (ASR) on the user utterance by an ASR module in the first  
application specific to the first application and generating a first n-best list for the user  
utterance, processing the first n-best list by the first classifier in the first application, the  
arbitration is performed by an arbitration module on a head unit of a vehicle, the first  
30 application is located on the vehicle and the second application is located on a mobile  
phone wirelessly connected to the head unit, the first classifier and the second classifier  
were trained using a reference dataset, the first classifier includes an in domain/out of  
domain classifier trained using input data relevant to the first application for the in domain

data and input data that is not relevant to the first application for out of domain data, where the latter is extracted from the reference dataset, the first and second applications were independently developed, but using the same reference dataset to obtain out of domain data for training, and/or the first classifier receives a feature set as input derived  
5 from the user utterance.

In another aspect of the invention, an article comprises: a non-transitory computer readable storage medium having stored instructions that enable a machine to: receive, in response to a user utterance, a first input from a first classifier of a first application, the  
10 first input including a probability of at least a portion of the user utterance being in domain for the first application, wherein in domain for the first application refers to functionality provided by the first application; receive, in response to the user utterance, a second input from a second classifier of a second application, the second input including a probability of at least a portion of the user utterance being in domain for the second application; and  
15 perform arbitration on the user utterance using the first and second inputs to prioritize a first one of the first and second applications for responding to the user utterance.

The article can further include one or more of the following features: performing automated speech recognition (ASR) on the user utterance by an ASR module in the first  
20 application specific to the first application and generating a first n-best list for the user utterance, processing the first n-best list by the first classifier in the first application, the arbitration is performed by an arbitration module on a head unit of a vehicle, the first application is located on the vehicle and the second application is located on a mobile phone wirelessly connected to the head unit, the first classifier and the second classifier  
25 were trained using a reference dataset, the first classifier includes an in domain/out of domain classifier trained using input data relevant to the first application for the in domain data and input data that is not relevant to the first application for out of domain data, where the latter is extracted from the reference dataset, the first and second applications were independently developed, but using the same reference dataset to obtain out of  
30 domain data, and/or the first classifier receives a feature set as input derived from the user utterance.

In a further aspect of the invention, a system comprises: an interface to communicate with first and second applications; an arbitration means coupled to the interface, the arbitration means to arbitrate on a user utterance from a user based on a first input from a first application in response to the user utterance and a second input from a second application in response to the user utterance, wherein the first input is received from a first classifier of the first application, the first input including a probability of at least a portion of the user utterance being in domain for the first application, wherein in domain for the first application corresponds to functionality provided by the first application, and wherein the second input is received from a second classifier of the second application, the second input including a probability of at least a portion of the user utterance being in domain for the second application, the arbitration means to prioritize a first one of the first and second applications for responding to the user utterance.

The system can further include one or more of the following features: automated speech recognition (ASR) performed on the user utterance by an ASR module in the first application specific to the first application and generating a first n-best list for the user utterance, the first n-best list processed by the first classifier in the first application, the arbitration performed by an arbitration module on a head unit of a vehicle, the first application located on the vehicle and the second application located on a mobile phone wirelessly connected to the head unit, the first classifier and the second classifier trained using a reference dataset, the first classifier including an in domain/out of domain classifier trained using input data relevant to the first application for the in domain data and input data that is not relevant to the first application for out of domain data, where the latter is extracted from the reference dataset, the first and second applications were independently developed, but using the same reference dataset to obtain out of domain data, and/or the first classifier receives a feature set as input derived from the user utterance.

30

## BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing features of this invention, as well as the invention itself, may be more fully understood from the following description of the drawings in which:

5 FIG. 1 is a schematic representation of a system having module extensible arbitration;

FIG. 2 is a schematic representation showing further detail of a system having module extensible arbitration;

10 FIG. 3 is a flow diagram showing an illustrative sequence of steps for module extensible arbitration; and

FIG. 4 is a schematic representation of an illustrative computer that can perform at least a portion of the processing described herein.

15

## DETAILED DESCRIPTION

FIG. 1 shows a system 100 including an automobile head unit 102 in communication with a number of devices 104a-N. The head unit 102 includes a number of applications 106a-  
20 M, such as a phone application 106a, a navigation application 106b, and an entertainment application 106N. It is understood that a wide variety of applications 106 can be provided on an automobile head unit 102. The applications 106 are coupled to an interaction module 108, which includes an arbitration module 110.

25 A first device 104a, such as a mobile phone, includes an application 112, such as a social media application, coupled to an interaction module 114. A second device 104N, such as a tablet computer, includes a number of applications 116a-O, such as a navigation application and a SMS application, coupled to an interaction module 118.

30 As described more fully, below, the arbitration module 110 determines which application on the head unit 102 or devices 104 should respond to a user voice command, for example, from a passenger in the vehicle. The arbitration may also determine that multiple

applications can respond to the user voice command. In this case, its role is to determine the relative relevance of each application.

While illustrative embodiments of the invention are shown and described in conjunction  
5 with an automobile head unit, it is understood that embodiments of the invention are applicable to any system having various applications for which arbitration is desirable.

FIG. 2 shows a system 200 having a number of modules 202a-N implementing a subset of the user interface (UI) functionality. In embodiments, each of the modules 202  
10 corresponds to an application supported by the system. Examples could be a module supporting navigation functionality (voice controlled navigation), a module supporting phone functionality (voice aided dialing), and a module supporting control of a music player. The modules 202 can be provided on a vehicle head unit or connected device. At runtime, arbitration is provided by a single arbitration module 204 that analyzes a request  
15 passed to the UI system by a user. Depending upon the requested functionality, the arbitration module 204 passes control to the relevant module(s).

Each of the applications 202 provides different functionality, e.g., mobile phone, navigation, weather, music, from the vehicle head unit or a connected (or unconnected)  
20 device. It is understood that devices can be coupled to the vehicle head unit using any suitable communication protocol, such as BLUETOOTH. In general, head unit and connected applications use microphones and loudspeakers of the vehicle. In embodiments, the arbitration module 204 resides on the vehicle head end but can be located at any practical local or remote location.

25 In embodiments, the modules 202 operate independently from each other. In general, the modules 202 can be developed independently from each other to provide a readily extensible framework without changing existing modules or the arbitration system 204. Embodiments of the invention provide an UI system from an undefined number of  
30 independently developed UI modules and assigns control to the 'correct' module(s), as described more fully below.

Each of the modules 202 takes input of the same commonly defined format. Typically, in a voice UI system this will be an audio command of the user to the system (e.g., “play song X”). It could however also be a different format. For example, one can first place an ASR module, which transforms the audio command to a text command and is common to all  
 5 modules, and then use this as input to the different modules.

Each of the modules 202 includes functionality to provide an output to the arbitration module 204. This output format is a single number, which indicates the probability that received input is relevant to the current module, e.g., that the module possesses the  
 10 required functionality to respond to the received UI command. This probability is comparable over the different modules, as explained more fully below. This comparability enables operation of the arbitration module 204.

In one embodiment, the functionality present in each module 202 includes a respective  
 15 input processing module 210a-N and respective ID-vs-OOD classifier module 241a-N. Input processing 210 analyzes the input and derives from it a feature set 212a-N. In case the input is in the form of audio commands, such input processing would normally take the form of an ASR system. The feature set which forms the output of such a system can include a vector of numbers or Boolean values, for example, indicating:

- 20 - Whether or not certain keywords relevant to the module were recognized in the audio signal (e.g., for a navigation module, the input processor may search for words like “drive”, “navigate”, etc.)
- What is the confidence score that the ASR system attributes to this recognition (how certain is it to have recognized the sentence successfully).
- 25 - What is the language model score of the command recognized by the ASR system (what is the probability of a command containing these words in this order).

The terms recognition confidence, language model score, and feature set are well known in the art. The practice of calculating a feature set based on an input signal as input for a classifier is also well known in the art.

30

The ID-vs-OOD classifier 214 takes as input the feature set 212, and produces a probability that the received input signal (which was processed to obtain the feature set) is relevant to the current module. We refer to this probability as  $P(App_a|S)$ , where  $App_a$  is

the current module, and  $S$  is the input signal. Note that a classifier is used in the illustrative embodiment that produces a posterior probability estimate as output. Examples known in the art are logistic regression or neural networks. Other such classifiers can be found in the literature related to this field and are readily understood by one of ordinary skill in the art.

5 The input format accepted by the classifier can be quite free depending upon the needs of a particular application. It should be noted that the output produced by the input processing (the feature set) conforms with the type of input expected by the classifier.

One aspect of embodiments of the invention is that the output (the probability) of the classifier 214 of each module 202 is directly comparable to that of other classifiers. This is obtained through the training of the classifier 214. The training is a process in which the classifier 214 is adapted to the specific use case, by using a training process and a training dataset. The training process depends on the kind of classifier that is chosen; any suitable technique known in the art can be used. The training dataset, on the other hand, makes the

10

15 output of the classifiers 214 comparable.

The training dataset to be used to train the ID-vs-OOD classifier 214a of module 202a should contain examples of input data that are relevant to the module, which we refer to as In-Domain (ID) data, and input that that are not relevant to the modules, which we refer to

20 as Out-Of-Domain (OOD) data. If we assume that module 202a supports music functionality, ID data refer to commands related to the music functionality supported by the module, e.g., “Play song X”, or “Switch off the radio,” while OOD data can be anything, from “Navigate to LA” to “What’s the weather like tomorrow?” In an embodiment, each module 202 should be associated with a set of ID data, a broad and

25 general set of examples of input that are relevant to the module.

To explain how we make the output probability of each classifier comparable, we first consider a conventional integrated system, which is not modular (where it is not needed that modules can be developed individually and combined freely). In this case, one would

30 train a single multi-class classifier, which has  $N$  outputs, one for each module, where each output would correspond to the posterior probability of one module, e.g.,  $P(App_a|S)$ . Classifiers for such multi-class classification and training for them are known in the art.

The training dataset of such classifier would correspond to the conjunction of ID data for all  $N$  modules.

As a first step towards modularity in accordance with embodiments of the invention, we  
5 can change this approach into training  $N$  classifiers, one for each class, with a single  
output  $P(App_a|S)$ . The output of this classifier 214a indicates the posterior probability of  
one module 202a. We refer to such single output classifier 214a as a binary classifier. In  
this case the binary classifier 214a is trained with its own ID data labeled as ID, and the  
data of all other modules 202b-N as OOD.

10

Now we consider the fully modular system targeted by this invention, in which modules  
are developed independently and combined freely. During the development of a module  
 $App_a$ , one cannot easily compose a good training dataset with ID and OOD data, where  
the OOD data are the conjunction of all ID data of all other modules. While the ID of  
15  $App_a$  are known, the ID data of all other modules are normally not known. This is because  
modules are developed independently and will also be put together independently at  
runtime. This means that we do not know the OOD part of the dataset, and we also do not  
know its size. To address this, illustrative embodiments include the following:

20

- We use a reference dataset (206). This dataset contains a wide variety of examples  
of input data. It should be as broad as possible, and generally represent input that  
may be given to the system. Whenever we train a classifier for a new module  
 $App_a$ , we use as OOD data any sentence of the reference dataset that is not  
relevant to modules  $App_a$ . This reference dataset is the common point of reference  
that ties together the training processes of different modules.

25

- We train with a fixed proportion of ID data and OOD data. This proportion can be  
chosen freely, but it should be the same for all modules, e.g., 25% ID data.

30

The arbitration module 204 produces an arbitration decision by comparing the  
probabilities that were calculated by each module 202 including the binary ID-vs-ODD  
and ID probability. New applications can be readily added since the classifier for the new  
module can generate output in a format that is comparable to the other modules.

In general, the arbitration module 204 can receive input from any practical number of application modules 202. In addition, weights can be generated to modulate the behavior of the arbitration module 204, such as favoring certain modules over others. Such weights can reflect user preferences, learning results, training, and the like, taken alone or in  
5 combination.

Subsequent to arbitration, a chosen module 202 can also perform additional classification tasks that are specific to the module, for example, “play My Way by Frank Sinatra” may involve a topic/slot extraction 220 to play (topic) a particular song (slot) by a particular  
10 artist (slot). This task is specific to each module, and is independent from arbitration. This means that it can be trained using only ID data, which is available during the independent development of a single module  $App_a$ . Any algorithm for topic/slot extraction known in the art can be used at this point.

15 If topic/slot extraction is done using a classifier that estimates posterior probabilities (like described for ID vs OOD classification), we can refer to its output as  $P(I|App_a, S)$ , where  $I$  refers to a combination of topic and slots. This probability is dependent on the input signal  $S$ , and the module  $App_a$ , because topic/slot extraction is specific to a module. In this case, one can calculate a probability of a topic/slot combination that is comparable over all  
20 modules, as  $P(I|App_a, S)P(App_a|S)$  (where we assume that a given topic/slot combination is only possible within one module).

FIG. 3 shows an illustrative sequence of steps for providing arbitration between independent pluggable modules. In step 300, an input is received, such as a user utterance.  
25 In step 302, first and second modules process the user utterance using ASR generate respective n-best outputs, which form the feature sets that serve as input to their respective classifiers. In step 304, a first classifier for the first module processes the first n-best list and a second classifier in the second module processes the second n-best list. The classifiers output probabilities. In step 306, an arbitration module processes the outputs  
30 from the first and second classifiers. In step 308, the arbitration module assigns one of the first and second modules to respond to the user utterance.

It is understood that the illustrative steps described above can include various types of inputs from which a feature set can be generated and used by an ID-OOD classifier.

It is understood that the term application as used herein refers to a set of functionalities.

5 Normally, such set of functionalities will be somehow related, e.g. because they refer to the same physical device (e.g., a “phone” application), or because they refer to semantically related concepts. Nevertheless, in principle, they can also refer to a more random combination of functionalities (e.g., an application covering both “phone” and “navigation” related functionality). Also, applications do not need to be mutually  
10 exclusive: there can be overlap in the functionality they support. Such flexible definition of the term application ensures a broad applicability of this system.

FIG. 4 shows an exemplary computer 400 that can perform at least part of the processing described herein. The computer 400 includes a processor 402, a volatile memory 404, a  
15 non-volatile memory 406 (e.g., hard disk), an output device 407 and a graphical user interface (GUI) 408 (e.g., a mouse, a keyboard, a display, for example). The non-volatile memory 406 stores computer instructions 412, an operating system 416 and data 418. In one example, the computer instructions 412 are executed by the processor 402 out of volatile memory 404. In one embodiment, an article 420 comprises non-transitory  
20 computer-readable instructions.

Processing may be implemented in hardware, software, or a combination of the two.

Processing may be implemented in computer programs executed on programmable computers/machines that each includes a processor, a storage medium or other article of  
25 manufacture that is readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices.

Program code may be applied to data entered using an input device to perform processing and to generate output information.

30 The system can perform processing, at least in part, via a computer program product, (e.g., in a machine-readable storage device), for execution by, or to control the operation of, data processing apparatus (e.g., a programmable processor, a computer, or multiple computers). Each such program may be implemented in a high level procedural or object-oriented

programming language to communicate with a computer system. However, the programs may be implemented in assembly or machine language. The language may be a compiled or an interpreted language and it may be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a  
5 computing environment. A computer program may be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network. A computer program may be stored on a storage medium or device (e.g., CD-ROM, hard disk, or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the  
10 computer when the storage medium or device is read by the computer. Processing may also be implemented as a machine-readable storage medium, configured with a computer program, where upon execution, instructions in the computer program cause the computer to operate.

15 Processing may be performed by one or more programmable processors executing one or more computer programs to perform the functions of the system. All or part of the system may be implemented as, special purpose logic circuitry (e.g., an FPGA (field programmable gate array) and/or an ASIC (application-specific integrated circuit)).

20 Elements of different embodiments described herein may be combined to form other embodiments not specifically set forth above. Various elements, which are described in the context of a single embodiment, may also be provided separately or in any suitable subcombination. Other embodiments not specifically described herein are also within the scope of the following claims.

25

Having described exemplary embodiments of the invention, it will now become apparent to one of ordinary skill in the art that other embodiments incorporating their concepts may also be used. The embodiments contained herein should not be limited to disclosed  
embodiments but rather should be limited only by the spirit and scope of the appended  
30 claims. All publications and references cited herein are expressly incorporated herein by reference in their entirety.

What is claimed is:

1. A method, comprising:

receiving, in response to a user input, a first input from a first classifier of a first  
5 application, the first input including a probability of at least a portion of the user input  
being in domain for the first application, wherein in domain for the first application refers  
to functionality provided by the first application;

receiving, in response to the user input, a second input from a second classifier of a  
10 second application, the second input including a probability of at least a portion of the user  
input being in domain for the second application; and

performing arbitration on the user input using the first and second inputs to  
prioritize a first one of the first and second applications for responding to the user input.

2. The method according to claim 1, further including performing automated speech  
15 recognition (ASR) on the user input by an ASR module in the first application specific to  
the first application and generating a first feature list for the user input.

3. The method according to claim 2, further including processing the first feature list by  
20 the first classifier in the first application.

4. The method according to claim 1, wherein the arbitration is performed by an  
arbitration module on a head unit of a vehicle.

5. The method according to claim 1, wherein the first application is located on the vehicle  
25 and the second application is located on a mobile phone wirelessly connected to the head  
unit.

6. The method according to claim 1, wherein the first classifier and the second classifier  
were trained using a reference dataset.

30

7. The method according to claim 1, wherein the first classifier includes an in domain/out  
of domain classifier trained using input data relevant to the first application for the in

domain data and input data that is not relevant to the first application for out of domain data.

8. The method according to claim 1, wherein the first and second applications were  
5 independently developed.

9. The method according to claim 1, wherein the first classifier receives a feature set as input derived from the user input.

10 10. An article, comprising:

a non-transitory computer readable storage medium having stored instructions that enable a machine to:

receive, in response to a user input, a first input from a first classifier of a first application, the first input including a probability of at least a portion of the user input  
15 being in domain for the first application, wherein in domain for the first application refers to functionality provided by the first application;

receive, in response to the user input, a second input from a second classifier of a second application, the second input including a probability of at least a portion of the user input being in domain for the second application; and

20 perform arbitration on the user input using the first and second inputs to prioritize a first one of the first and second applications for responding to the user input.

11. The article according to claim 10, wherein the arbitration is performed by an arbitration module on a head unit of a vehicle.

25

12. The article according to claim 10, wherein the first application is located on the vehicle and the second application is located on a mobile phone wirelessly connected to the head unit.

30 13. The article according to claim 10, wherein the first classifier and the second classifier were trained using a reference dataset.

14. The article according to claim 10, wherein the first classifier includes an in domain/out of domain classifier trained using input data relevant to the first application for the in domain data and input data that is not relevant to the first application for out of domain data.

5

15. The article according to claim 10, wherein the first and second applications were independently developed.

16. The article according to claim 10, wherein the first classifier receives a feature set as  
10 input derived from the user input.

17. A system, comprising:

an interface to communicate with first and second applications;


an arbitration means coupled to the interface, the arbitration means to arbitrate on a  
15 user input from a user based on a first input from a first application in response to the user input and a second input from a second application in response to the user input, wherein the first input is received from a first classifier of the first application, the first input including a probability of at least a portion of the user input being in domain for the first application, wherein in domain for the first application corresponds to functionality  
20 provided by the first application, and wherein the second input is received from a second classifier of the second application, the second input including a probability of at least a portion of the user input being in domain for the second application, the arbitration means to prioritize a first one of the first and second applications for responding to the user input.

25 18. The system according to claim 17, wherein the first application is located on a vehicle head unit and the second application is located on a device wirelessly connected to the head unit.

19. The system according to claim 17, wherein the first classifier includes an in  
30 domain/out of domain classifier trained using input data relevant to the first application for the in domain data and input data that is not relevant to the first application for out of domain data.

20. The system according to claim 17, wherein the first classifier receives a feature set from an automated speech recognition module as input derived from the user input.

4



**FIG. 1**

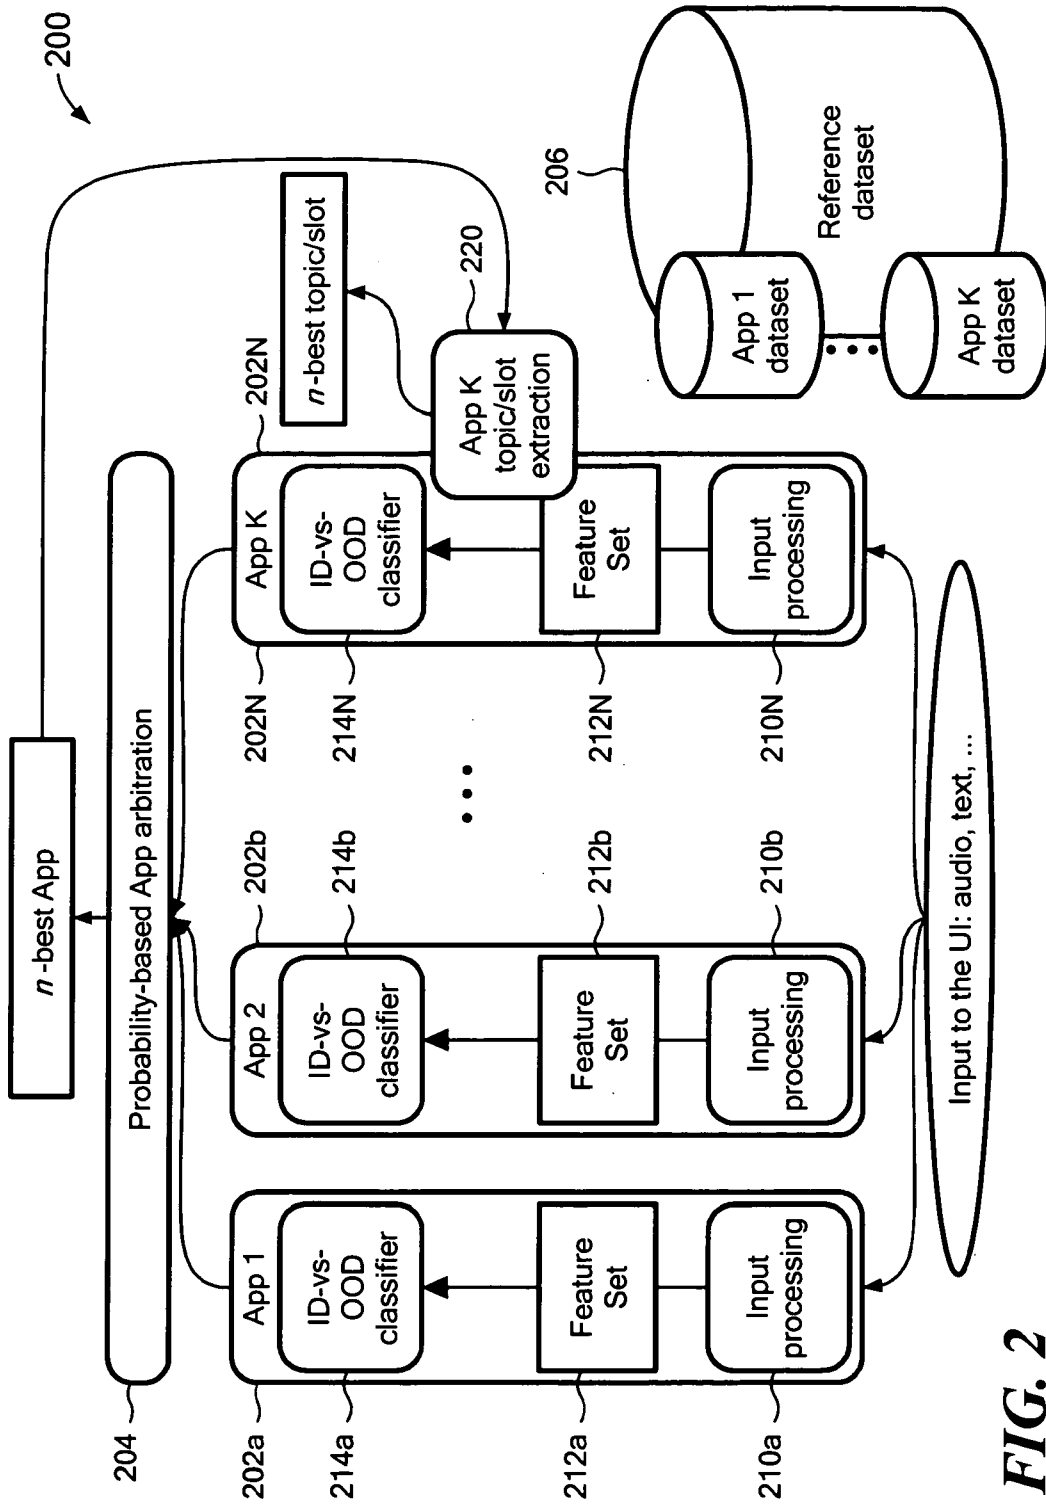
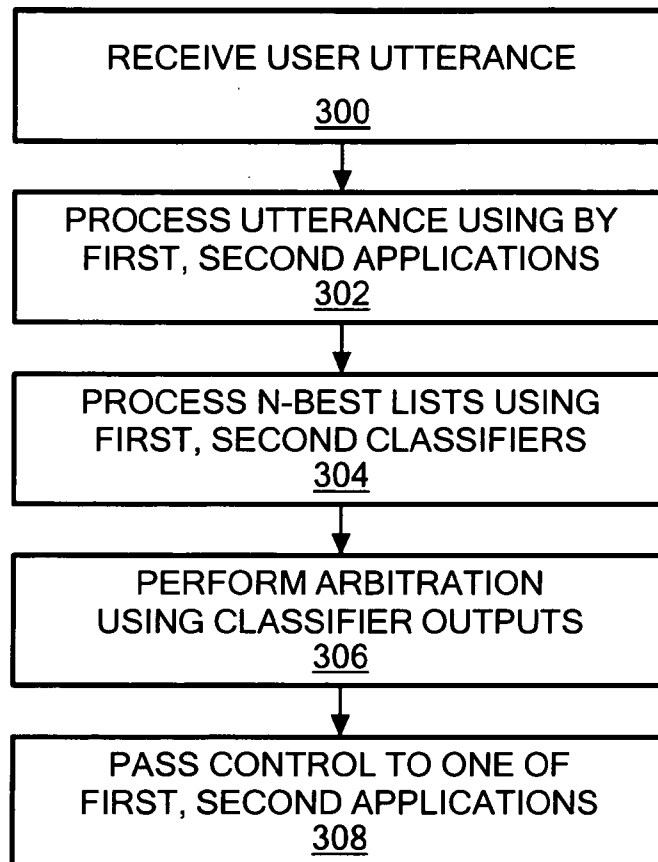



FIG. 2

3/4

**FIG. 3**

4/4



**FIG. 4**

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2014/058211****A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/46(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**Minimum documentation searched (classification system followed by classification symbols)  
G06F 9/46; G10L 15/00; G10L 15/22; G10L 21/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
Korean utility models and applications for utility models  
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
eKOMPASS(KIPO internal) & Keywords: application, arbitration, speech recognition, ASR**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 06098043 A (FOREST; SERGE et al.) 01 August 2000 See abstract;column 3, line 10 - column 6, line 20; claims 1,2,4,6,7,8 and figures 3-5.	1-20
A	EP 2733697 A1 (QNX SOFTWARE SYSTEMS LIMITED) 21 May 2014 See abstract;paragraphs [0011]-[0026];claims 1-10 and figures 1-3.	1-20
A	US 2003-0171929 A1 (STEVE, RUSSEL FALCON et al.) 11 September 2003 See abstract;paragraphs [0040]-[0064];claims 1,14,15,23 and figure 1.	1-20
A	US 2003-0005174 A1 (DANIEL M. COFFMAN et al.) 02 January 2003 See abstract;paragraphs [0010]-[0023];claims 1,21,34 and figures 2,4,6.	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;" document member of the same patent family

Date of the actual completion of the international search

26 June 2015 (26.06.2015)

Date of mailing of the international search report

**26 June 2015 (26.06.2015)**

Name and mailing address of the ISA/KR

International Application Division  
Korean Intellectual Property Office  
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,  
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KU, Dae Sung

Telephone No. +82-42-481-8192



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2014/058211**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 06098043 A	01/08/2000	None	
EP 2733697 A1	21/05/2014	EP 2733697 A9	11/02/2015
US 2003-0171929 A1	11/09/2003	US 2006-0053016 A1 US 2006-0161429 A1 US 2010-0191529 A1 US 7167831 B2 US 7363229 B2 US 7720678 B2 US 8447616 B2	09/03/2006 20/07/2006 29/07/2010 23/01/2007 22/04/2008 18/05/2010 21/05/2013
US 2003-0005174 A1	02/01/2003	AT483200 T CN1224894 C CN1491382 A DE60237825 D1 EP 1410171 A1 EP 1410171 A4 EP 1410171 B1 IL159565A IL159565D0 JP 2005-503609 T JP 2005-503609A JP 3943543B2 US 6839896 B2 WO 03-003193 A1	15/10/2010 26/10/2005 21/04/2004 11/11/2010 21/04/2004 02/04/2008 29/09/2010 16/06/2010 01/06/2004 03/02/2005 03/02/2005 11/07/2007 04/01/2005 09/01/2003