



- (51) International Patent Classification:  
*G06F 17/30* (2006.01)    *H04W 4/02* (2009.01)
- (21) International Application Number:  
PCT/US2014/011375
- (22) International Filing Date:  
14 January 2014 (14.01.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
13/743,339    17 January 2013 (17.01.2013)    US
- (71) Applicant: MICROSOFT CORPORATION [US/US];  
One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventors: LYMBEROPOULOS, Dimitrios; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). LIU, Jie; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). WANG, He; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (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, LT, 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, 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:**

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

- (54) Title: ACCUMULATION OF REAL-TIME CROWD SOURCED DATA FOR INFERRING METADATA ABOUT ENTITIES




FIG. 1

- (57) Abstract: Various technologies pertaining to crowd sourcing data about an entity, such as a business, are described. Additionally, technologies pertaining to inferring metadata about the entity based upon crowd sourced data are described. A sensor in a mobile computing device is activated responsive to a user of the mobile computing device checking in at an entity. Metadata, such as occupancy at the entity, noise at the entity, and the like is inferred using the data captured by the sensor. A search result for the entity includes the metadata.

WO 2014/113347 A2

- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*
- Published:**
  - *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

## ACCUMULATION OF REAL-TIME CROWD SOURCED DATA FOR INFERRING METADATA ABOUT ENTITIES

### BACKGROUND

5 [0001] Search engines are computer-implemented systems that are configured to provide information to a user that is relevant to the information retrieval intent of such user. Generally, this intent is at least partially represented by terms of a query submitted to the search engine by the user. Responsive to receipt of the query, the search engine provides the user with a ranked list of search results, wherein search results determined to be most  
10 relevant by the search engine are positioned most prominently (e.g., highest) in the ranked list of search results.

[0002] Some search engines have been further adapted to perform local searches, where search results are retrieved based upon a location provided by a user (either explicitly or implicitly based upon a sensed current location of the user). For example, when the user  
15 employs a mobile telephone to perform a local search, a GPS sensor on the telephone can be employed to provide the search engine with the current location of the user. The user then sets forth a query, and search results relevant to both the provided location and the query are provided to the user by the search engine. Oftentimes, when performing a local search, the user is looking for a particular type of business, such as a restaurant, pub, retail  
20 store, or the like.

[0003] Generally, search results retrieved by the search engine for a local search are ranked in accordance with one or more features, such as a distance between the location provided by the user and an entity represented by a respective search result, user rankings for the entity represented by the respective search result, popularity of the entity  
25 represented by the respective search result (e.g., based upon a number of user clicks), or the like. Information about the entity included in the search result, however, may be stale. For instance, a most recent user review about an entity may be several months old, and is therefore not indicative of current or recent activity at the entity. Furthermore, reviews or ratings for a large number of local entities might not be available.

30

### SUMMARY

[0004] The following is a brief summary of subject matter that is described in greater detail herein. This summary is not intended to be limiting as to the scope of the claims.

[0005] Described herein are various technologies pertaining to crowd sourcing data that is indicative of a state of at least one parameter of an entity, and utilizing crowd sourced

data to provide information to users as to the state of the at least one parameter of the entity. In an example, the entity may be a business, such as an eatery, the parameter may be noise level of the eatery, and the state of such parameter may be “low,” “normal,” “high,” or “very high.” Other parameters that crowd sourced data may be indicative of include, but are not limited to, lighting level of an entity, occupancy of the entity, type of music being played at the entity, level of noise of music being played at the entity, level of human chatter (conversation) at the entity, a particular song being played at the entity, whether the entity has an outdoor area, temperature at the entity, humidity level at the entity, barometric pressure at the entity, amongst other parameters.

5  
10 **[0006]** In an exemplary embodiment, sensors of mobile computing devices of users can be leveraged to crowd source data that is indicative of a state of at least one parameter of an entity. Pursuant to an example, a user of a mobile computing device may employ such device to publish her presence at a particular entity. For example, the user may choose to “check-in” to a certain entity (business or other predefined location) through utilization of  
15 a social networking application that is accessible to the user on the mobile computing device. Since it is known that the user is at the entity (due to the user checking into the entity), and it is further known that the mobile computing device is in the hand of the user (rather than a pocket or bag), then it can be presumed that at least one sensor of the mobile computing device is exposed to a current environmental condition at the entity. Thus, the  
20 sensor, at least for a relatively short period of time, can capture data that is indicative of a state of at least one parameter of the entity. Therefore, immediately responsive to the user indicating her presence at the entity, at least one sensor of the mobile computing device can be activated and caused to capture data. In an exemplary embodiment, the at least one sensor can be a microphone that is configured to capture an audio stream. For instance, the  
25 microphone can be activated for some relatively small threshold amount of time, such as 5 seconds or 10 seconds. The audio stream captured by the microphone can be streamed as it is captured at the mobile computing device to another computing device by way of a wireless network connection. In another embodiment, a data packet that comprises an entirety of the audio stream can be transmitted to the another computing device subsequent  
30 to the audio stream being captured at the mobile computing device. The process of capturing data from sensors of mobile computing devices as their respective users check in to certain entities can be undertaken for numerous mobile computing devices at several entities, resulting in crowd sourced data for numerous entities. Furthermore, in an exemplary embodiment, raw audio data can be transmitted to the another computing

device, wherein the another computing device thereafter generates features that will be described below. Alternatively, the mobile computing device can generate the features, and can subsequently send such features (without the raw audio data) to the another computing device.

5 [0007] Data captured by a mobile computing device in the manner described above can be employed to provide a relatively recent and accurate picture of a current state of at least one parameter of an entity. For instance, 10 seconds worth of data captured by a mobile computing device at an entity can be received and processed at a server or at the mobile computing device, wherein the processed data can be used to estimate the state of the at  
10 least one parameter of the entity. Such processing can be undertaken shortly after the data is received. For instance, the received sensor data can be segmented into a plurality of relatively small time segments (e.g., one second segments), and features can be extracted from each segment, thereby creating a respective feature vector for each segment. Such feature vectors, in an example, may then be transmitted to the server. Subsequently, for  
15 each feature vector, a classifier can probabilistically determine the state of the parameter using a respective feature vector as input. Therefore, for example, the classifier can output a classification for the state of the parameter for each segment in the plurality of segments. A majority voting technique may then be employed to determine a final classification of the state of the parameter, and such classification can be stored in a computer-readable  
20 data repository together with information included in a search result for the entity. Thus, for instance, if a search result for the entity is presented to an issuer of a query, the search result can include data indicative of the estimated state of the aforementioned parameter of the entity. In an example, therefore, the entity may be a restaurant, and the search result can indicate that there is currently a high level of occupancy at the restaurant, there is  
25 currently a low noise level at the restaurant, etc.

[0008] In another exemplary embodiment, a search engine index can be updated to include real-time or near real-time metadata about an entity, such that search results can be ranked based upon the real-time or near real-time metadata. Pursuant to an example, data about a business can be crowd sourced in the manner described above. With respect to  
30 a particular business, inferred metadata about the business can indicate that currently, there is a high level of occupancy at the business, and currently rock music is being played at the business. A person searching for a particular place to visit with friends may submit the query “crowded bar playing rock music” to the search engine, and the search engine can highly rank a search result for the business since it maps to the information retrieval intent

of the user (e.g., despite the business not being closest in proximity to the issuer of the query).

[0009] Other aspects will be appreciated upon reading and understanding the attached figures and description.

5

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0010] Fig. 1 illustrates an exemplary mobile computing device that can be employed in connection with crowd sourcing data about an entity.

[0011] Fig. 2 is a functional block diagram of an exemplary mobile computing device.

[0012] Fig. 3 illustrates crowd sourcing of data about numerous entities.

10

[0013] Fig. 4 is a functional block diagram of an exemplary computing device that facilitates processing crowd sourced data about entities to infer states of respective parameters of the entities.

[0014] Fig. 5 illustrates an exemplary search result for an entity.

15

[0015] Fig. 6 is a functional block diagram of an exemplary search engine that can rank search results based at least in part upon crowd sourced data about entities represented by such search results.

[0016] Fig. 7 illustrates an exemplary query and corresponding exemplary search results provided responsive to receipt of such query.

20

[0017] Fig. 8 is a functional block diagram of an exemplary system that facilitates training a classifier to infer metadata about an entity.

[0018] Fig. 9 is a flow diagram that illustrates an exemplary methodology for configuring a mobile computing device to provide data about an entity.

25

[0019] Fig. 10 is a flow diagram that illustrates an exemplary methodology for utilizing a classifier to infer a state of a parameter of an entity based at least in part upon data about the entity received from a mobile computing device.

[0020] Fig. 11 is a flow diagram that illustrates an exemplary methodology for providing a search result responsive to receipt of a query, wherein the search result depicts information based upon real-time or near real-time metadata about an entity corresponding to the search result.

30

[0021] Fig. 12 is a flow diagram that illustrates an exemplary methodology for ranking search results based at least in part upon real-time or near real-time metadata about entities represented by the search results.

[0022] Fig. 13 is an exemplary computing system.

### **DETAILED DESCRIPTION**

[0023] Various technologies pertaining to crowd sourcing data about entities, and utilizing such data to infer states of parameters of the entities, will now be described with reference to the drawings, where like reference numerals represent like elements  
5 throughout. In addition, several functional block diagrams of exemplary systems are illustrated and described herein for purposes of explanation; however, it is to be understood that functionality that is described as being carried out by certain system components may be performed by multiple components. Similarly, for instance, a component may be configured to perform functionality that is described as being carried  
10 out by multiple components. Additionally, as used herein, the term “exemplary” is intended to mean serving as an illustration or example of something, and is not intended to indicate a preference.

[0024] As used herein, the terms “component” and “system” are intended to encompass computer-readable data storage that is configured with computer-executable instructions  
15 that cause certain functionality to be performed when executed by a processor. The computer-executable instructions may include a routine, a function, or the like. It is also to be understood that a component or system may be localized on a single device or distributed across several devices.

[0025] With reference now to Fig. 1, an exemplary depiction 100 of a mobile computing  
20 device being employed in connection with crowd sourcing data about an entity is illustrated. In the example shown in Fig. 1, a mobile computing device 102 is at a particular entity 104. In an exemplary embodiment, the entity 104 can be one of a plurality of entities that are predefined in a social networking application. For example, the entity  
25 104 may be a business, such as a restaurant, a pub, a retail outlet, a movie theater, an amusement park, a golf course, etc. In another example, the entity 104 may be a monument or other known location, such as an airport, a stadium, a park, a tourist attraction, an arena, a library, or the like.

[0026] The mobile computing device 102, as will be described in greater detail below, may comprise a plurality of sensors, wherein such sensors may include a microphone, a  
30 GPS sensor (or other positioning sensor), a gyroscope, a barometer, a humidity sensor, a thermometer, amongst other sensors. While the mobile computing device 102 is shown as being a mobile telephone, it is to be understood that the mobile computing device 102 may be some other suitable mobile computing device, such as a tablet computing device

(sometimes referred to as a slate computing device), a portable media player, an e-reader, or the like.

[0027] The mobile computing device 102 can be configured to access the above-mentioned social networking application, wherein the user of the mobile computing device  
5 can publish messages via the social networking application. For instance, the mobile computing device 102 may have the social networking application installed thereon. In another exemplary embodiment, the mobile computing device 102 may have a web browser installed thereon, such that the social networking application is accessible by way of the web browser. In still yet another exemplary embodiment, the mobile computing  
10 device 102 may be configured with an operating system that includes functionality that enables the operating system to communicate with the social networking application.

[0028] The mobile computing device 102 is shown as including a touch-sensitive display 106, wherein a button 108 is displayed thereon. When the button 108 is depressed, the user of the mobile computing device 102 is indicating a desire to publish that such user  
15 is at the entity 104. In some social networking applications, such publishing is referred to as the user “checking in” at the entity 104. When the user selects the button 108 by pressing a finger on such button 108, for example, at least one sensor of the mobile computing device 102 can be activated to capture data that is indicative of a current state of at least one parameter of the entity 104. For instance, immediately responsive to the  
20 user of the mobile computing device 102 pressing the button 108 to indicate that the user of the mobile computing device 102 is at the entity 104, a microphone of the mobile computing device 102 can be activated to capture an audio stream, wherein the audio stream is indicative of a level of occupancy at the entity 104, a level of volume of music being played at the entity 104, a level of background human chatter at the entity 104, a  
25 level of noise generally at the entity 104, a type of music being played at the entity, a particular song being played at the entity, etc. Similarly, for instance, responsive to the user selecting the button 108, a luminance sensor or camera can be activated to capture data that is indicative of level of luminance at the entity 104 (e.g., and therefore may be indicative of whether the mobile computing device 102 is outdoors when the user is  
30 checking in at the entity 104). In another example, a gyroscope of the mobile computing device 102 can be activated immediately responsive to the user of the mobile computing device 102 selecting the button 108, wherein data captured by the gyroscope can be indicative of movement of the mobile computing device 102 over a relatively short timeframe. Other exemplary sensors that can capture data that is indicative of a state

(level) of at least one parameter of the entity have been noted above, although such list of sensors is not intended to be exhaustive.

[0029] After a predefined event has occurred, the sensor activated responsive to the user selecting the button 108 can be deactivated. In an example, the predefined event may be the passage of some relatively short threshold period of time (e.g., on the order of 5 seconds, 10 seconds, or 15 seconds). In another example, the predefined event may be depression of a power button on the mobile computing device 102, the mobile computing device 102 entering a low power (e.g., sleep) mode, or some other suitable event. In yet another example, the predefined event may be the obtainment of a suitable reading from the sensor (such as a convergence to a certain value, the obtainment of a value with some confidence of the value being correct, etc.) This may preserve the battery life of the mobile computing device 102, and may cause sensors to capture data about the entity 104 with relatively high accuracy. In the example of the microphone, since the user must select the button 108, it can be inferred that the mobile computing device 102 is not in a pocket or bag of the user of the mobile computing device 102, and that data captured by the microphone is therefore not muffled or otherwise affected.

[0030] The mobile computing device 102 is then configured to transmit data captured by at least one sensor of the mobile computing device 102 to a computing device 110 by way of a suitable wireless network. In an exemplary embodiment, the mobile computing device 102 can be configured to stream data as it is captured by the at least one sensor to the computing device 110. In another exemplary embodiment, the mobile computing device 102 can be configured to generate a data packet that includes the data captured by the sensor when in an active state, and transmit such data packet to the computing device 110. Optionally, the mobile computing device 102 can be configured to compress such data packet prior to transmitting the data packet to the computing device 110. Still further, rather than transmitting raw sensor data to the computing device 110, the mobile computing device 102 can perform processing on the sensor data and transmit the result of such processing to the computing device 110. For instance, as will be described in greater detail below, the mobile computing device 102 can segment the sensor data into numerous segments, and then generate a feature vector for each segment. The mobile computing device 102 may then transmit the feature vectors (rather than the raw sensor data) to the computing device 110.

[0031] In an exemplary embodiment, a check can be performed prior to activating the sensor of the mobile computing device 102 and transmitting the data to the computing

device 110. For instance, when the user checks in at the entity 104, a current location of the mobile computing device 102 (as output by a GPS sensor or other position sensor) can be compared with a known location of the entity 104; if the current location of the mobile computing device 102 does not map to the known location of the entity 104, then the sensor of the mobile computing device can remain in an inactive state. Thus, the at least one sensor may be activated only if the mobile computing device 102 is found to be at the known location of the entity 104.

**[0032]** As will be described in greater detail below, the computing device 110 can be configured to collect crowd sourced data about the entity 104 and other entities being checked into by other users of mobile computing devices to infer respective states of parameters of entities. With respect to the crowd sourced data received from the mobile computing device 102 shown in Fig. 1, the computing device 110 can process such data through utilization of one or more classifiers to infer, for example, the level of occupancy of the entity 104 at the time that the data was captured at the mobile computing device 102.

**[0033]** States of the parameters that can be inferred/determined based upon the data captured by the sensors of the mobile computing device 102 can be employed in a variety of different settings. For example, such states of parameters can be provided as portions of local search results about the entity 104. In an exemplary embodiment, the entity 104 may be a restaurant, and a user may set forth a query for restaurants that are proximate to the current location of the user. Search results retrieved responsive to the query being submitted can include a search result for the entity 104, and the search result can include information typically found in local search results, such as the name of the restaurant, the address of the restaurant, the telephone number of the restaurant, and the like.

Additionally, information pertaining to the states of the aforementioned parameters can be included in the search result. Thus, the issuer of the query can quickly determine (assuming that there is relatively recent data about the entity 104) that the entity 104 has a particular level of occupancy, that the entity 104 has an outdoor area, that a particular type of music is being played at the entity at a certain volume, etc.

**[0034]** In other embodiments, the data generated at the mobile computing device 102 provided to the computing device 110 can be employed for the purpose of trend detection/analysis. That is, over time, patterns for states of parameters of the entity 104 can be recognized based upon crowd sourced data about the entity 104. A search result for the entity can include or make reference to such patterns; thus, for instance, a search result

for a restaurant can include information such as “on Friday nights between 6:00 p.m. and 8:00 p.m. the restaurant is historically very crowded.” In still yet another exemplary embodiment, the data captured at the mobile computing device 102 provided to the computing device 110 can be utilized in connection with ranking search results. Thus, 5 estimated states of parameters of the entity 104 can be included in a search engine index, thereby allowing a search engine to provide contextually relevant search results for queries such as “crowded restaurant in the warehouse district.”

[0035] Now referring to Fig. 2, a functional block diagram of an exemplary mobile computing device 200 is illustrated. As noted above, the mobile computing device 200, in 10 exemplary embodiment, is a mobile telephone, although in other embodiments the mobile computing device 200 may be some other mobile computing device. The mobile computing device 200 comprises a processor 202 and a memory 204, wherein the processor 202 is configured to execute instructions retained in the memory 204. The processor 202 is intended to encompass a conventional processor utilized in mobile 15 computing devices as well as a system on a chip (SoC) or cluster on a chip (CoC) systems. The memory 204 includes an operating system 206, which is configured to manage hardware resources of the mobile computing device 200. The memory 204 additionally includes, for instance, an application 208 that is executed by the processor 202, wherein the application 208 may be a social networking application or a browser that can access 20 the social networking application.

[0036] The mobile computing device 200 further comprises a display 210, wherein in an exemplary embodiment, the display 210 may be a touch-sensitive display. The display 210 can be configured to display data corresponding to the application 208 when the application 208 is executed by the processor 202.

25 [0037] The mobile computing device 200 further comprises a plurality of sensors 212, wherein such sensors 212 may include a microphone, a GPS sensor, a humidity sensor, a thermometer, a camera, a luminance sensor, a barometer, or other suitable sensor. The sensors 212 can be selectively activated and deactivated by the processor 202. The mobile computing device 200 further comprises a battery 214 that is configured to provide energy 30 to the hardware resources of the mobile computing device 200.

[0038] The mobile computing device 200 further comprises a wireless radio 216 that can be employed to receive and transmit data by way of a wireless network. In an exemplary embodiment, the wireless radio 216 can be configured to receive and transmit data by way of a cellular network. In another exemplary embodiment, the wireless radio 216 can be

configured to receive and transmit data by way of a Wi-Fi network. In still yet another exemplary embodiment, the wireless radio 216 may be configured to receive and transmit data by way of a short range communications protocol, such as Bluetooth. It is to be understood, however, that the wireless radio 216 may be configured to receive and  
5 transmit data over any suitable type of wireless network.

**[0039]** The operating system 206 includes a sensor control component 218 that can selectively activate and deactivate the sensors 212. As noted above, a user of the mobile computing device 200 can access the social networking application by way of the application 208 in the memory 204. When the user is at a particular location (entity), the  
10 user may desire to publish her being at the location by way of the application 208. For instance, as shown in Fig. 1, the user can indicate that they are at a particular restaurant or other business, and “check in” to that restaurant or business by way of the application 208. The sensor control component 218 can receive the indication that the user desires to publish that she is at the location, and can selectively activate at least one sensor in the  
15 sensors 212, thereby causing the at least one sensor to capture data that is indicative of a state of at least one parameter of the location. As noted above, exemplary parameters of the location include, but are not limited to, occupancy level, music volume, background human chatter volume, noise volume, a music type being played at the location, a particular song being played at the location, temperature at the location, etc. Accordingly,  
20 the state of such parameter can be, for instance, low, normal, high, very high, may be a particular number (such as a certain temperature sensed by a thermometer), or the like.

**[0040]** After a predefined event as occurred, such as the mobile computing device 200 entering a low power mode, the passage of a threshold amount of time after the sensor control component 218 has activated the at least one sensor, or the obtainment of a  
25 suitable sensor reading, the sensor control component 218 can deactivate the at least one sensor. The operating system 206 may further include a data transmitter component 220 that is configured to cause the wireless radio 216 to transmit the data captured by the at least one sensor to the computing device 110, where such data can be processed.

**[0041]** Now referring to Fig. 3, a depiction 300 of numerous sensors being used to crowd source data about various entities (e.g., businesses) is illustrated. A plurality of entities can be predefined in a social networking application. One of such entities may be the entity 104 shown in Fig. 1 (referred to as the first entity). Other locations can include a second entity 302 through an  $n$ th entity 304. The mobile computing device 102 may be at the first entity 104, and when the user of the mobile computing device 102 checks in at the

first entity 104, as described above, data that is indicative of a state of at least one parameter of the first entity 104 can be transmitted to the computing device 110. Two sensors 306 and 308 may additionally be at the first entity 104. In an exemplary embodiment, the sensors 306 and 308 may be sensors in mobile computing devices that are activated in a manner similar to the manner that the sensor of the mobile computing device 102 is activated (when users of mobile computing devices that respectively include such sensors 306 and 308 “check in” at the first entity 104). In other embodiments, however, at least one of the sensors 306 or 308 may be another type of sensor that can capture data that is indicative of a state of at least one parameter of the first entity 104. For instance, the sensor 306 may be a microphone that is permanently located at the entity captures audio data at the first entity 104 (e.g., is not a portion of a mobile computing device).

**[0042]** Sensors 310 and 312 are shown as being at the second entity 302, and can output data that is indicative of a state of at least one parameter of the second entity 302 to the computing device 110. Such sensors 310 and 312, in an exemplary embodiment, may be sensors in respective mobile computing devices that are activated responsive to users of the mobile computing devices “checking in” at the second entity 302.

**[0043]** Four sensors 314-320 may be at the *n*th entity 304, and can transmit data about the *n*th entity 304 to the computing device 110. One or more of such sensors may be relatively permanent sensors that are associated with the *n*th entity 304, such as video cameras, microphones, thermometers, etc., while other sensors in the sensors 314-320 may be included in respective mobile computing devices and are selectively activated upon users of the mobile computing devices “checking in” at the *n*th entity 304.

**[0044]** In an exemplary embodiment, since, for example, the *n*th entity 304 has multiple data streams received by the computing device 110 pertaining to the state of the at least one parameter, the computing device 110 can aggregate such data to reduce the possibility of incorrectly estimating the state of the at least one parameter. Moreover, it can be ascertained that the computing device 110 receives real-time data from various sensors across a plurality of different entities. Still further, in addition to receiving real-time data from sensors, it is to be understood that the computing device 110 may retrieve data about the entities 104, 302, and 304 from textual sources. For instance, a user at the first entity 104 may employ her mobile computing device to publish, by way of a public feed, that the entity 104 is crowded. The computing device 110 can be configured to mine such public

feeds and correlate user locations with textual descriptions of parameters of such locations to obtain data about the state of the at least one parameter.

[0045] Now referring to Fig. 4, a functional block diagram of the computing device 110 is illustrated. The computing device 110 comprises a data repository 402 that retains  
5 sensor data 404 received from a sensor of a mobile computing device (such as the mobile computing device 102). The computing device 110 additionally includes a segmenter component 406 that segments the sensor data 404 into a plurality of time segments. For instance, if the sensor data 404 comprises a ten second signal, the segmenter component 406 can segment such signal into ten non-overlapping one second segments. In another  
10 exemplary embodiment, the segmenter component 406 may segment the sensor data 404 into a plurality of segments, at least some of which are partially overlapping. Length of time of the segments and how such segmenting is undertaken over the sensor data 404 can be determined empirically.

[0046] A feature extractor component 408 receives segments of the sensor data 404  
15 from the segmenter component 406, and extracts features from the segments that are indicative of respective states of parameters of the entity from which the sensor data 404 has been received. Such features can include temporal features and/or frequency features of the segments, such that a respective temporal signature and a respective frequency signature can be generated by the feature extractor component 408 for each segment. Such  
20 signatures can have the form of respective feature vectors, and a feature vector including temporal features and a feature vector including frequency features for a segment of the sensor data 404 can be combined into a single feature vector for the segment. Exemplary features that can be extracted from an audio signal by the feature extractor component 408 that are indicative of states of different parameters of an entity will be described below.  
25 While the computing device 110 has been described as including the segmenter component 406, and the feature extractor component 408, it is to be understood that the mobile computing device 102 may comprise such components.

[0047] The computing device 110 also comprises a classifier 410 that receives feature vectors for respective segments of the sensor data 404. The classifier 410 is trained to  
30 output a classification as to a state of a particular parameter of the entity from which the sensor data is received for each segment of the sensor data 404 based upon the feature vectors for a respective segment. In an exemplary embodiment, the classifier 410 can be trained to output a classification as to a level of occupancy of the entity (e.g., business) from which the sensor data is received. Therefore, the classifier 410, continuing with the

example set forth above, can output ten separate classifications as to the occupancy level of the location from which the sensor data is received.

[0048] An output component 412 can receive the classifications output by the classifier 410, and can output a final classification as to the state of the parameter for the sensor data 404 based upon the respective classifications for the segments of the sensor data 404. In an exemplary embodiment, a majority voting technique can be employed by the output component 412 to output a classification as to the state of the parameter for the sensor data 404.

[0049] The computing device 110 may further comprise a database 414. In an exemplary embodiment, the database 414 can comprise information that is to be included in a search result for the entity from which the sensor data has been received, such as the name of the entity, address of the entity (or geographic coordinates), telephone number of the entity, and the like, and the state of the at least one parameter output by the output component 412 can be retained in the database 414 with such information.

[0050] The computing device 110 may further include a search component 416 that receives a query, wherein a search result retrieved responsive to receipt of the query is for the entity from which the sensor data was received. The search result includes up-to-date (recent) information about the state of the at least one parameter of the location that has been output by the output component 412 as described above. Accordingly, for example, the search result for the entity may include the name of the entity, the street address of the entity, contact information for the entity (telephone number), as well as an indication that the entity is currently believed to be highly crowded (e.g., based upon recently received sensor data about the entity). Accordingly, the issuer of the query can obtain information about a current or recent state of a parameter of the entity of interest, which has heretofore been unavailable to searchers.

[0051] The computing device 110 may optionally include a purger component 418 that can remove stale information from the database 414. In an exemplary embodiment, the purger component 418 can remove information from the database 414 that is over 30 minutes old.

[0052] Additionally, the output component 412 can access the database 414 and generate classifications using the sensor data 404 (which has been recently received from a mobile computing device at the entity) and previously computed classifications for the entity existent in the database 414, wherein such previously computed classifications have been computed relatively recently. Accordingly, for example, the output component 412

can consider classifications (e.g., in a sliding time window) in the database 414 when computing new classifications.

[0053] While the computing device 110 is shown as including a single classifier 410, is to be understood that the computing device 110 may include multiple different classifiers that are configured to output classifications of states for multiple different parameters of the entity from which the sensor data was received. For instance, a first classifier may be configured to output a classification as to occupancy level of the entity, a second classifier may be configured to output a classification as to volume of music being played at the entity, a third classifier may be configured to output a classification as to an amount (volume) of background human chatter at the entity, a fourth classifier may be configured to output a classification as to the general noise level at the location, a fifth classifier may be configured to output an indication as to whether there is an outdoor area at the location (e.g., as evidenced by data received from a luminance sensor, GPS, or camera of a mobile computing device), etc. Such classifiers can act in a manner similar to what has been described with respect to the classifier 410, in that the classifiers can output classifications with respect to parameters for multiple segments of received sensor data, and the output component 412 can output a final classification for states of parameters based upon classifications assigned to individual segments.

[0054] Moreover, while the purger component 418 has been described as purging data from the database 414, it is to be understood that trend analysis can be undertaken over data in the database 414, such that if the database 414 does not include recently generated classifications with respect to states of one or more parameters of an entity, trend data indicative of historic states of parameters of the entity can be included in the database 414. For instance, if, over time, it can be ascertained, based upon analysis of crowd sourced data, that a particular entity is typically very crowded and loud music is typically played on Saturdays from 5:00 PM to 10:00 PM, and at a particular Saturday data has not been received for such entity, then the database 414 can be configured to include the trend data that indicates that it can be expected that the entity is crowded and plays loud music between the aforementioned hours on Saturdays. Therefore, if a query is issued on a Saturday between 5:00 PM and 10:00 PM, a search result can be returned to an issuer of the query that indicates that, while sensor data has not recently been received, the issuer of the query can expect the entity to be highly crowded and loud music to be playing.

[0055] Now referring to Fig. 5, an exemplary search result 500 for an entity is illustrated, wherein data about the entity has recently been received. In an exemplary

embodiment, the entity can be a business, and therefore the search result can for the business. An issuer of the query can specify a location (explicitly or implicitly), and can note in the query that she is searching for a business of a type that corresponds to the entity. The search result 500 can be retrieved, based upon the query, in a manner similar to  
5 how conventional local search results are retrieved. For instance, the search result may be ranked in a list of search results based upon proximity of the entity to the location specified by the user. In another exemplary embodiment, the search result 500 can be ranked in a list of search results based upon proximity to the specified location and average rating given by users those that have visited the business.

10 **[0056]** The search result 500 can include information that is included in conventional local search results, such as the identity of the business, the address of the business, the telephone number of the business, an average rating given by reviewers of the business, and the like. The search result 500 can also include information about a current or recent  
15 state of a parameter of the business. As shown, a viewer of the search result 500 can ascertain that the business is currently relatively busy (the occupancy is high), that the level of human chatter is normal, that the level of music volume is low, that an outdoor section exists at the business, and that the overall noise level is low.

**[0057]** Now referring to Fig. 6, an exemplary search engine 600 that can rank search results as a function of metadata about entities (e.g., computed/inferred based upon data  
20 received from mobile computing devices known to be at such entities) is illustrated. The search engine 600 includes a search engine index 602. The search engine index 602 includes an index of search results that can be retrieved responsive to the search engine receiving queries. The search engine 600 additionally comprises an updater component 604 that updates the search engine index 602 based at least in part upon crowd sourced  
25 data. With more particularity, the updater component 604 can receive classifications output by the output component 412 and can update the search engine index 602 based upon such classifications.

**[0058]** The search engine 600 further comprises the search component 416 that receives a query and executes a search over the search engine index 602. A ranker component 606  
30 can rank search results retrieved by the search component 406 based at least in part upon the crowd sourced data. The updating of the search engine index 602 allows the search engine 600 to rank search results based upon current or recent states of parameters of entities represented by the search results. For instance, a restaurant determined to be relatively quiet and playing classical music (through analysis of crowd sourced data) can

be returned as a highly ranked search result for a query “quiet four-star restaurant playing classical music”.

[0059] With reference now to Fig. 7, an exemplary query 700 that may be provided to a search engine and exemplary search results 702 that can be retrieved and ranked

5 responsive to receipt of the query are illustrated. The query 700 is requesting information about a “crowded restaurant in the warehouse district”. Responsive to receiving the query, the search engine 600 can set forth the list of search results 702, which can include search results 704-708. The search results 702 can be ordered based at least in part upon crowd  
10 sourced data received from sensors at the entities represented by the search results 704-708. For example, the search result 704 displayed highest in the ranked list of search results may be found to be relatively busy based upon data recently received from at least one sensor that is at the entity represented by the search result 704. The search engine 600 can consider other features when ranking search results, such as location of the entity represented by the search result, average user rating for the entity represented by the  
15 search result, when data about the entity represented by the search result was received, etc.

[0060] Turning now to Fig. 8, an exemplary system 800 that facilitates training the classifier 410 to output a classification with respect to a state of a parameter of an entity from which sensor data is received is illustrated. The system 800 includes a data store 802 that comprises labeled training data 804. For the purposes of explanation, the labeled  
20 training data 804 will be described herein as being audio data received from microphones of respective mobile computing devices at different businesses. It is to be understood, however, that the labeled training data 804 is not so limited to audio data and that entities are not to be limited to businesses. Pursuant to an example, human labelers can be instructed to assign labels similar to those shown below in Table 1, wherein the parameters  
25 of interest are occupancy, background chatter, music level, and general noise level, and wherein respective states of such parameters are “very high”, “high”, “normal”, and “low”, with corresponding definitions for such states. Additionally, labels for the parameters are for an entirety of an audio stream, which may be on the order of ten seconds, fifteen seconds, or the like (e.g., labels need not be assigned to one second  
30 segments).

**Table 1**

Metadata	Class (State)			
	Very High	High	Normal	Low
<b>Occupancy</b>	>80%	60%-80%	30%-60%	<30%
<b>Background Chatter</b>	Need to yell to be heard	Need to talk loudly to be heard	Normal talking, clearly hear other people	Barely hear other people
<b>Music Level</b>	Need to yell to be heard	Need to talk loud to be heard	Normal talking, clearly listen to music	Barely hear music or no music
<b>Noise Level</b>	Loud noise	Loud enough to distract you	Typical indoor environmental noise	Barely hear any noise or no noise

[0061] The segmenter component 406 can receive an audio stream from the label training data 804 and segment such audio stream into a plurality of segments. Each segment can be assigned the label that has been assigned to the entirety of the audio stream, thereby creating numerous segments of audio data from the audio stream. The segmenter component 406 can perform such segmentation for each labeled audio stream in the label training data 804.

[0062] The feature extractor component 408 may then extract features from each segment provided thereto by the segmenter component 406. Each segment output by the segmenter component 406 can be treated independently, in the sense that feature extraction and classifier training takes place for each segment separately. Segmentation of audio streams can be employed to ensure high classification accuracies; since the overall recording time of each audio stream can be on the orders of tens of seconds, the characteristics of the different sound sources can easily change multiple times during a recording. For instance, in a beginning portion of a recording, there may be people speaking directly in front of the mobile computing device while music is playing. Later, the music can be stopped for a short amount of time as the next song is beginning to be played, and subsequently the people in front of the mobile computing device can stop speaking, such that only background chatter is recorded, etc. Generally, sound sources can significantly change within seconds, resulting in major variations in recorded audio

streams. Therefore, using lengthy audio segments that span multiple seconds can pollute the feature extraction process, which in turn can lead to erroneous classification.

Segmentation of audio streams can remove such variations, thereby allowing more robust inferences to be made over multiple shorter time windows during a recording.

5 [0063] The feature extractor component 408 is tasked with encoding unique characteristics of different types of sounds, such as background chatter, music, and noise, in each of the segments provided thereto by the segmenter component 406. This approach enables inference models to accurately recognize different sound levels in an audio stream, even though all of such sounds are simultaneously recorded using a single microphone. To  
10 achieve this, smoothness and amplitude of audio signals can be examined by the feature extractor component 408 in both the temporal and frequency domains. First, the feature extractor component 408 can generate short-term features over sub-second audio windows, wherein the duration of such windows can vary, depending upon the exact feature that is being extracted. The feature extractor component 408 may then generate  
15 second long-term features by examining statistics of the short term features over all sub-second windows within a segment. Accordingly, for every feature, and for each segment, its mean, minimum, and maximum values can be recorded over all different sub-second windows, as well as its overall variation across these windows. This set of long-term features can form the actual feature vector that describes each segment and can be what is  
20 output by the feature extractor component 408.

[0064] Exemplary features that can be extracted by the feature extractor component 408 and for training the classifier 410 and for use by the classifier 410 to classify a state of at least one parameter are now set forth. Such features are exemplary and are not intended to be interpreted as an exhaustive list. Furthermore, for the purposes of an example, the  
25 features relate to features that can be extracted from audio streams captured by a microphone of a mobile computing device. Again, this is set forth for purposes of explanation, and is not intended to be limiting.

[0065] In the temporal domain, a recorded audio stream describes the amplitude of the audio stream over time. Absolute amplitude information can be utilized for estimating the  
30 loudness (volume) of the audio stream, either as music, human chatter, or people talking close to a microphone of a mobile computing device. To capture differences in the absolute amplitude, depending upon the type of parameter (music, human chatter, or near phone talking), the energy  $E$  of a recorded audio signal can be calculated as the root mean square of the audio samples  $s_i$  as follows:

$$E = \sqrt{\frac{\sum_{i=1}^N \text{sign}(s_i)^2}{N}} \quad (1)$$

where  $N$  is the total number of samples. Energy can be calculated over sliding windows of some threshold time duration (e.g., 50 ms). Using a 50 ms duration, if a microphone sampling rate is 16 KHz, and segments are one second in length, an energy value can be computed a total of 15,201 times within each segment.

[0066] Relative amplitude information, and in particular the smoothness of amplitude over time, can additionally provide insight about the recorded audio stream. In general, human talking includes continuous repetitions of consonants and vowels, resulting in audio signals with high energy and zero cross rate (ZCR) variations in short time windows. It has been found through experiment that a recording of background music tends to be a far “smoother” audio signal than a recording of human talking. For example, the energy in ZCR variations in the signal are almost negligible compared to the ones found in audio signals when people are actually talking. This difference holds even when comparing a person talking with the person singing the exact same word or sentence. The energy variance during normal human speech is significantly smoothed out during singing because of the different pronunciation of consonants and vowels (e.g., prolonging vowels, etc.). To capture this fundamental difference between human talking and music in an audio signal, ZCR, as well as the variation of both ZCR and energy can be computed as follows:

$$ZCR = \frac{\sum_{i=2}^N |\text{sign}(s_i) - \text{sign}(s_{i-1})|}{2} \quad (2)$$

$$ZCR_{var} = \sqrt{\frac{\sum_{i=1}^N |ZCR_i - ZCR_{mean}|^2}{N-1}} \quad (3)$$

$$E_{var} = \sqrt{\frac{\sum_{i=1}^N |E_i - E_{mean}|^2}{N-1}}. \quad (4)$$

25

[0067] Similar to energy, ZCR can be computed over sliding time windows of some threshold duration (e.g., 50 ms). Using a 50 ms duration, if a microphone sampling rate is 16 KHz, and segments are one second in length, 15,201 ZCR values can be computed for each segment. In an exemplary embodiment,  $ZCR_{var}$  and  $E_{var}$  can be computed using longer sliding windows that overlap with a predefined step size. If the segment duration is one segment, duration of the sliding window is 500 ms, and the step size is 100 ms, five

30

$ZCR_{var}$  and five  $E_{var}$  values can be computed for each segment. Using the segment duration, microphone sampling rate, subwindow durations, and step size set forth above, 30,412 features can be computed for each segment:

$$\begin{aligned}
 5 \quad F_{s.t.}^{temp.} = & \{E_{s.t.} = E^1, \dots, E^{15201}\}, ZCR_{s.t.} [ZCR^1, \dots, ZCR^{15201}], E_{s.t.}^{var} = \\
 & [E_{VAR}^1, \dots, E_{VAR}^5], ZCR_{s.t.}^{var} = [ZCR_{VAR}^1, \dots, ZCR_{VAR}^5] \quad (5)
 \end{aligned}$$

[0068]  $F_{s.t.}^{temp.}$  represents short-term features generated from the sub-second audio segment processing. Such features are not directly used as input to the classifier training stage, but instead statistics for each of the four different types of short-term temporal features are computed. Specifically, the minimum, maximum, mean, and variation values of  $E$ ,  $ZCR$ ,  $ZCR_{var}$ , and  $E_{var}$  are computed over all values in Eq. (5) as follows:

$$\begin{aligned}
 15 \quad F_{l.t.}^{temp.} = & \{\{min, max, mean, var\}(E_{s.t.}), \{min, max, mean, var\}(ZCR_{s.t.}), \\
 & \{min, max, mean, var\}(E_{s.t.}^{var}), \{min, max, mean, var\}(ZCR_{s.t.}^{var})\}. \quad (6)
 \end{aligned}$$

$F_{l.t.}^{temp.}$  therefore 16 long-term features for each segment. This set of temporal features represents a temporal signature of each segment and can be used as input during training of the classifier 410, and can likewise be used as input for classification undertaken by the classifier 410.

[0069] Frequency domain features can also be extracted by the feature extractor component 408. The feature extractor component 408 can utilize similar processing of the audio stream to analyze the magnitude of the audio stream across frequencies and its smoothness over time. Such features can capture parts of the underlying structure of the audio stream that temporal features might not be able to accurately capture. When frequency and temporal domain features are combined, a more descriptive feature set is generated, and thus, a more robust basis for accurate classifier training and (and therefore accurate classification) is formed.

[0070] In the frequency domain, the feature extractor component 408 can calculate the spectrogram of the recorded audio stream. This can be undertaken by dividing the audio stream into relatively small non-overlapping windows and computing the Fast Fourier Transform (FFT) for each window. The spectrogram can then be computed by the feature extractor component 408 by concatenating all the different FFTs. In essence, the

spectrogram describes the magnitude of the audio stream at different frequencies over time, and forms the basis for feature extraction.

[0071] Directly encoding a spectrogram as a feature is not a scalable approach, as a large number of features would be generated, posing stringent restrictions on data collection and model training. Instead, building components of the spectrogram can be leveraged to extract a relatively small feature set. For instance, for each segment provided by the segmenter component 406, a 512-point FFT of the audio signal can be calculated (32 ms time window given a 16KHz microphone sampling rate) in 31 non-overlapping windows. For each of the 31 FFTs, the DC component can be deleted and the remaining frequency bins can be normalized, such that the sum of squares is equal to one.  $p_t(i)$  is used to denote the magnitude of the  $i$ th frequency bin of the normalized FFT at time  $t$ . The spectrogram can be summarized by computing spectral centroid (SC), bandwidth (BW), and spectral flux (SF), as follows:

$$SC = \frac{\sum_{i=1}^N i * p(i)^2}{\sum_{i=1}^N p(i)^2}, N = 256, \quad (7)$$

$$BW = \frac{\sum_{i=1}^N (i - SC)^2 * p(i)^2}{\sum_{i=1}^N p(i)^2}, N = 256, \quad (8)$$

$$SF_t = \sum_{i=1}^N (p_t(i) - p_{t-1}(i))^2, N = 256 \quad (9)$$

[0072] Both spectral centroid and bandwidth can be computed for each one of the 31 FFTs over a single segment, resulting in 31 SC and 31 BW features. Additionally, spectral flux can be computed for every consecutive pair of FFTs, resulting in 30 SF features for each segment.

[0073] Intuitively, spectral centroid represents the center frequency of the computed FFT and is calculated as the weighted mean of the frequencies present in the normalized FFT, with the magnitude of these frequencies as the weights. Bandwidth is a measure of the width/range of frequencies in the computed FFT. Finally, spectral flux represents the spectrum difference in adjacent FFTs and is an indication of the variation of spectral density over time.

[0074] In addition to spectral centroid bandwidth and spectral flux, Mel-Frequency Cepstral Coefficients (MFCCs) can be computed by the feature extractor component 408. MFCCs are coefficients that collectively make up an MFC, which is a representation of the short-term power spectrum of a sound. MFC coefficients have been widely used in

speech recognition and speaker identification, and are considered high-quality descriptors of human speech.

[0075] To compute MFC coefficients, numerous sliding windows with particular step sizes can be employed. In an exemplary embodiment, 256 sample sliding windows with a step size of 128 samples (given a 16 KHz microphone sampling rate) can be employed. This results in 124 windows for each segment. For each window, the first twelve MFCC coefficients can be leveraged, and the  $i$ th MFCC coefficient at window  $t$  can be denoted as  $MFCC^t(i)$ .

[0076] As a result, given the exemplary numbers set forth above, the set of short-term frequency domain features extracted over each segment can include 1580 features:

$$F_{s.t.}^{freq} = \left\{ \left\{ \begin{array}{l} SC_{s.t.} = [SC^1, \dots, SC^{31}], BW_{s.t.} = \\ [BW^1, \dots, BW^{31}], SF_{s.t.} = [SF^1, \dots, SF^{30}] \end{array} \right\}, MFCC_{s.t.}(i) = \right. \\ \left. [MFCC^1(i), \dots, MFCC^{124}(i)], i = 1, \dots, 12 \right\} \quad (10)$$

[0077] The long-term features that are eventually used during classifier training are computed directly from the short-term frequency features. Similarly to the temporal domain feature extraction, the minimum, maximum, mean, and variation values of SC, BW and SF, as well as the mean values for each of the 12 MFCC coefficients are computed over all the short-term feature values in Eq. 10, resulting in the following long-term frequency features:

$$F_{l.t.}^{temp.} = \left\{ \begin{array}{l} \{min, max, mean, var\}(SC_{s.t.}), \\ \{min, max, mean, var\}(BW_{s.t.}), \\ \{min, max, mean, var\}(SF_{s.t.}), \\ \{mean\}(MFCC_{s.t.}(i)), i = 1, \dots, 12 \end{array} \right\}. \quad (11)$$

[0078]  $F_{l.t.}^{temp.}$  includes, therefore, 24 features for each segment. This set of long-term features represents the frequency signature of each segment of the audio stream, and is used as input during the classifier training phase, and such features of audio streams are also input to the classifier 410 when performing classification.

[0079] As described above, audio streams can be mapped to feature vectors, thereby allowing the labeled training data 804 to be leveraged to train classifiers to output

classifications to states of parameters that are desirably inferred. The labeled training data 804 can be utilized to bootstrap the training process, and a relatively small number of audio streams (approximately 100) may be sufficient to train accurate classifiers. In an exemplary embodiment, each audio stream in the labeled training data 804 can be associated to one of four different labels, as shown in Table 1 above, for each type of parameter that is desirably inferred. For instance, labels for occupancy, human chatter, music, and noise levels are provided for each audio recording.

[0080] The system 800 includes a trainer component 806 that receives the features extracted by the feature extractor component 408 and the labels assigned to the label training data 804 and learns feature weights for the parameter type for which states are desirably inferred. The trainer component 806 can utilize a variety of machine-learning approaches to multiclass classification to learn a mapping between feature values and labels representing the state of a specific parameter type. In an exemplary embodiment, the classifier 410 can be a decision tree.

[0081] By leveraging real audio streams captured at entities (businesses), wherein such audio streams are properly labeled with information of when people were talking near the microphone (near-phone talking) or not, an additional classifier can be built for inferring when an audio stream includes near-phone talking. Audio stream segmentation, feature extraction, and training for a near phone talking classifier can be undertaken as described above with respect to the classifier 410. The only difference are the binary labels (near phone talking or not) assigned to the feature vectors during the training process. The near phone talking model can be used in two ways. First, audio streams for which near phone talking has been detected can be filtered from being employed as training data, since near-phone talking can dominate the audio stream and hide the background sounds that are desirably; ignoring such streams helps to remove noise. Additionally, instead of completely removing audio streams, the output of the near phone talking classifier can be leveraged at run-time as a binary input feature for other classifiers. In such a way, enough information can be provided in the training phase to enable classifiers to adjust to near-phone talking audio streams and maintain high recognition rates.

[0082] Returning to Fig. 4, at runtime, having trained classifiers for each type of parameter that is desirably the subject of classification, the level or state of various parameters of interest can be inferred. This can occur as described above, where sensor data can be split into segments, and temporal and frequency domain features can be computed for each of such segments. The trained classifier 410 can take computed feature

vectors for each segment as input, and probabilistically map the feature vector to one of several pre-specified states of the parameter. The label with the highest probability can be assigned to each segment, and then majority voting across all segments can be applied to infer the state of the parameter. This process can create a real-time stream of crowd  
5 sourced data about locations (businesses). At any given time, when a user searches for nearby local businesses, search engines can leverage recently inferred metadata to inform the user about the state of the parameter of the business at the time of the query.

[0083] With reference now to Figs. 9-12, various exemplary methodologies are illustrated and described. While the methodologies are described as being a series of acts  
10 that are performed in a sequence, it is to be understood that the methodologies are not limited by the order of the sequence. For instance, some acts may occur in a different order than what is described herein. In addition, an act may occur concurrently with another act. Furthermore, in some instances, not all acts may be required to implement a methodology described herein.

[0084] Moreover, the acts described herein may be computer-executable instructions  
15 that can be implemented by one or more processors and/or stored on a computer-readable storage medium or media. The computer-executable instructions may include a routine, a sub-routine, programs, a thread of execution, and/or the like. Still further, results of acts of the methodologies may be stored in a computer-readable storage medium, displayed on a  
20 display device, and/or the like.

[0085] With reference now to Fig. 9, an exemplary methodology 900 that facilitates the use of a mobile computing device to generate crowd sourced data about an entity is illustrated. The methodology 900 starts at 902, and at 904 an indication from a user is received that the user is checking in at an entity that is predefined in a social networking  
25 application. Such entity may be a business, such as an eatery. At 906, responsive to receiving indication, at least one sensor of the mobile computing device is activated. In other words, the at least one sensor is transitioned from an inactive state to an active state. Pursuant to an example, and as noted above, the at least one sensor may be a microphone that is configured to generate an audio stream/signal.

[0086] At 908, a determination is made regarding whether a predefined event has been  
30 detected. For instance, the predefined event may be receiving an indication that a user has depressed a power button (to put the mobile computing device in a low power mode). In another example, the predefined event may be passage of a threshold amount of time since the sensor has been activated. In yet another example, the predefined event may be the

obtainment of a suitable reading from the sensor. Generally, such threshold amount of time may be relatively short, such as on the order of 5 seconds, 10 seconds, or 15 seconds. In an exemplary embodiment, if the predefined event has not been detected since the at least one sensor was transitioned to the active state, then the methodology 900 proceeds to 5 910, where data captured by the at least one sensor while in the active state can be streamed to a computing device by way of a wireless network. If the predefined event has been detected, then at 912, the at least one sensor is deactivated (transition from the active state to the inactive state).

[0087] In another exemplary embodiment, rather than streaming the data to the 10 computing device, the mobile computing device can be configured to generate a data packet that includes the entirety of an audio stream (which may optionally be compressed and transmit the data packet to the computing device after the threshold amount of time is passed). In still another example, the mobile computing device can process raw sensor data to generate feature vectors that are respectively indicative of at least one state of a 15 parameter of an entity for numerous time segments, and the mobile computing device can transmit the feature vectors to the computing device. The methodology 900 completed 914.

[0088] With reference now to Fig. 10, an exemplary methodology 1000 for assigning a classification as to a state of at least one parameter of an entity (business) is illustrated. 20 The methodology 1000 starts at 1002, and at 1004 data from a sensor of a mobile computing device that is indicative of a state of a parameter of the entity is received. As noted above, in an example, such data may be an audio stream.

[0089] The 1006, the data is segmented into a plurality of segments. As noted above, such segments can be approximately one second non-overlapping segments. In other 25 examples, the segments may be partially overlapping, and duration of such segments may be shorter or longer than one second.

[0090] At 1008, a segment is selected from the plurality of segments, and at 1010 a classification for the state of the parameter of the entity is output for the selected segment. As noted above, a feature vector can be generated that can be viewed as a signature for the 30 segment, and a classifier can map the feature vector to one of several classes based upon feature weights learned during training of the classifier and values of the feature vector. At 1012, a determination is made if the segment over which the classifier has been executed is the last segment in the plurality of segments. If there are more segments in the plurality of segments that have not been analyzed, then the methodology 1000 returns to 1008. If it

is determined at 1012 that the segment is the last segment, then at 1014, a classification for the state of the parameter is computed based at least in part upon classifications for the states of the parameter for all segments.

5 [0091] At 1016, a determination is made regarding whether previous classifications within some threshold time window had been undertaken for the entity/parameter. If there are previous classifications within the time window, then at 1018, a final classification as to the state of the parameter of the entity is computed based upon the classification output at 1014 and the previous classifications. The methodology 1000 completed 1020.

10 [0092] With reference now to Fig. 11, an exemplary methodology 1100 that facilitates provision of real-time information about a parameter of an entity to an issuer of a query is illustrated. The methodology 1100 starts 1102, and at 1104 a query from a client computing device is received. At 1106, a search result is retrieved for a business based upon the received query. At 1108, a classification as to a state of a parameter of the business is retrieved, wherein the classification is based upon recently generated crowd  
15 sourced data. At 1110, the search result and the classification as to the state of the parameter are transmitted for display on the computing device of the user. The methodology completes at 1112.

[0093] Now referring to Fig. 12, an exemplary methodology 1200 for ranking search results based at least in part upon crowd sourced data is illustrated. The methodology 1200  
20 starts at 1202, and at 1204 crowd sourced sensor data pertaining to an entity (business) is received. At 1206, a classification as to a state is assigned to a state of the parameter for the entity based upon the crowd sourced sensor data. At 1208, a search engine index is updated based at least in part upon the classification.

[0094] At 1210, a query is received from a client computing device, wherein the query is  
25 configured to retrieve information about the entity. Further, the query can include a terms that relate to a near-time or real-time context of the entity. At 1212, a plurality of search results is retrieved responsive to receipt of the query, wherein the plurality of search results includes a search result for the entity. At 1214, the search result for the entity is positioned in a ranked list of search results based at least in part upon the classification  
30 assigned at 1206. The methodology 1200 completed 1216.

[0095] Now referring to Fig. 13, a high-level illustration of an exemplary computing device 1300 that can be used in accordance with the systems and methodologies disclosed herein is illustrated. For instance, the computing device 1300 may be used in a system that supports capturing data about an entity. In another example, at least a portion of the

computing device 1300 may be used in a system that supports inferring a state of a parameter of an entity based at least in part upon crowd sourced sensor data about the entity. The computing device 1300 includes at least one processor 1302 that executes instructions that are stored in a memory 1304. The memory 1304 may be or include RAM, ROM, EEPROM, Flash memory, or other suitable memory. The instructions may be, for instance, instructions for implementing functionality described as being carried out by one or more components discussed above or instructions for implementing one or more of the methods described above. The processor 1302 may access the memory 1304 by way of a system bus 1306. In addition to storing executable instructions, the memory 1304 may also store sensor signals, trained classifiers, etc.

[0096] The computing device 1300 additionally includes a data store 1308 that is accessible by the processor 1302 by way of the system bus 1306. The data store may be or include any suitable computer-readable storage, including a hard disk, memory, etc. The data store 1308 may include executable instructions, sensor signals, etc. The computing device 1300 also includes an input interface 1310 that allows external devices to communicate with the computing device 1300. For instance, the input interface 1310 may be used to receive instructions from an external computer device, from a user, etc. The computing device 1300 also includes an output interface 1312 that interfaces the computing device 1300 with one or more external devices. For example, the computing device 1300 may display text, images, etc. by way of the output interface 1312.

[0097] Additionally, while illustrated as a single system, it is to be understood that the computing device 1300 may be a distributed system. Thus, for instance, several devices may be in communication by way of a network connection and may collectively perform tasks described as being performed by the computing device 1300.

[0098] Various functions described herein can be implemented in hardware, software, or any combination thereof. If implemented in software, the functions can be stored on or transmitted over as one or more instructions or code on a computer-readable medium. Computer-readable media includes computer-readable storage media. A computer-readable storage media can be any available storage media that can be accessed by a computer. By way of example, and not limitation, such computer-readable storage media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to carry or store desired program code in the form of instructions or data structures and that can be accessed by a computer. Disk and disc, as used herein, include compact disc (CD),

laser disc, optical disc, digital versatile disc (DVD), floppy disk, and Blu-ray disc (BD), where disks usually reproduce data magnetically and discs usually reproduce data optically with lasers. Further, a propagated signal is not included within the scope of computer-readable storage media. Computer-readable media also includes communication  
5 media including any medium that facilitates transfer of a computer program from one place to another. A connection, for instance, can be a communication medium. For example, if the software is transmitted from a website, server, or other remote source using a coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared, radio, and microwave, then the coaxial cable, fiber optic  
10 cable, twisted pair, DSL, or wireless technologies such as infrared, radio and microwave are included in the definition of communication medium. Combinations of the above should also be included within the scope of computer-readable media.

**[0099]** It is noted that several examples have been provided for purposes of explanation. These examples are not to be construed as limiting the hereto-appended claims.  
15 Additionally, it may be recognized that the examples provided herein may be permuted while still falling under the scope of the claims.

### CLAIMS

1. A method executed on a mobile computing device, the method comprising:
  - receiving an indication that a user of the mobile computing device desires to publish, by way of a social networking application, that the user is at a particular entity;
  - responsive to receiving the indication, activating at least one sensor in the mobile computing device;
  - subsequent to the activating of the at least one sensor in the mobile computing device, transmitting data to another computing device by way of a wireless network, wherein the data is based upon data captured by the at least one sensor;
  - detecting that a predefined event has occurred subsequent to the at least one sensor in the mobile computing device being activated; and
  - deactivating the at least one sensor responsive to the detecting the predefined event.
2. The method of claim 1, wherein the at least one sensor comprises a microphone.
3. The method of claim 1, wherein the activating of the at least one sensor, the transmitting of the data to the another computing device by way of the wireless network, the detecting that the predefined event has occurred, and the deactivating the at least one sensor are undertaken by an operating system of the mobile computing device.
4. The method of claim 1, further comprising:
  - segmenting the data captured by the at least one sensor into a plurality of time segments;
  - for each time segment, generating a respective feature vector, the respective feature vector being indicative of a state of the entity; and
  - transmitting a plurality of feature vectors that correspond to the plurality of time segments to the another computing device.
5. The method of claim 1, further comprising:
  - receiving location information from a location sensor of the mobile computing device;
  - comparing the location information from the location sensor of the mobile computing device with a known location of the particular entity; and

activating the at least one sensor in the mobile computing device only if the location from the location sensor of the mobile computing device is within a predefined threshold distance from the known location of the particular entity.

6. A mobile computing device comprising:
  - a sensor; and
  - a memory, the memory comprising an operating system for the mobile computing device, the operating system being configured to perform acts comprising:
    - receiving input from a user that the user is at a business identified by the user;
    - responsive to receiving the input from the user, causing the sensor to transition from an inactive state to an active state;
    - detecting a predefined event;
    - responsive to detecting the predefined event, causing the sensor to transition from the active state to the inactive state; and
    - transmitting data that is indicative of a state of a parameter of the business to another computing device by way of a wireless network connection, wherein the data that is indicative of the state of the parameter of the business is based upon data captured by the sensor when the sensor was in the active state.
7. The mobile computing device of claim 6, wherein the parameter is one of occupancy of the business when the sensor was active, noise level of the business when the sensor was active, whether the user was outdoors at the business when the sensor was active, or type of music being played at the business when the sensor was active.
8. The mobile computing device of claim 6 being a mobile telephone.
9. The mobile computing device of claim 6, wherein the memory comprises a social networking application, and wherein the input from the user is received by way of the social networking application.
10. The mobile computing device of claim 6, wherein the predefined event is a passage of a threshold amount of time from when the sensor was transitioned from the inactive state to the active state.





FIG. 1



**FIG. 2**





FIG. 3



**FIG. 4**

500

SEARCH RESULT 1  
ADDRESS:  
TELEPHONE NUMBER  
RATING

OCCUPANCY: HIGH  
CHATTER: NORMAL  
MUSIC: LOW  
OUTDOOR SECTION: YES  
NOISE: LOW

FIG. 5

600

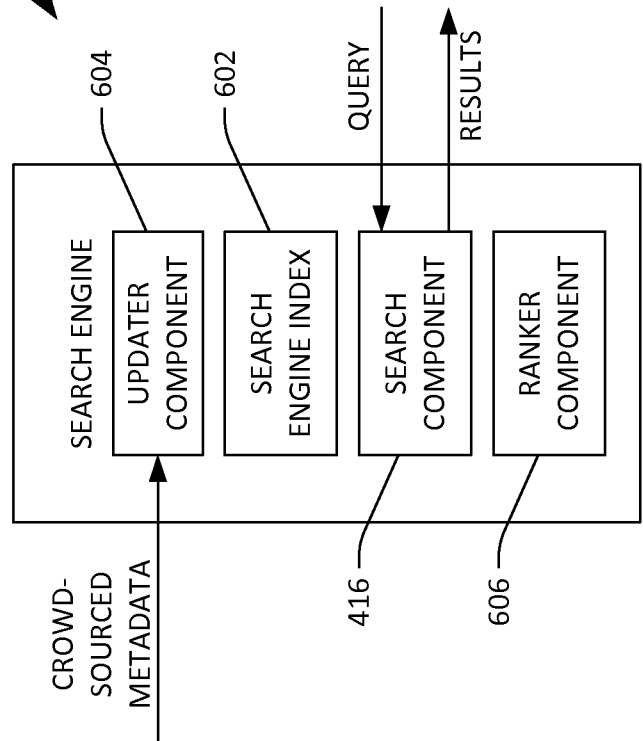


FIG. 6





FIG. 7



**FIG. 8**

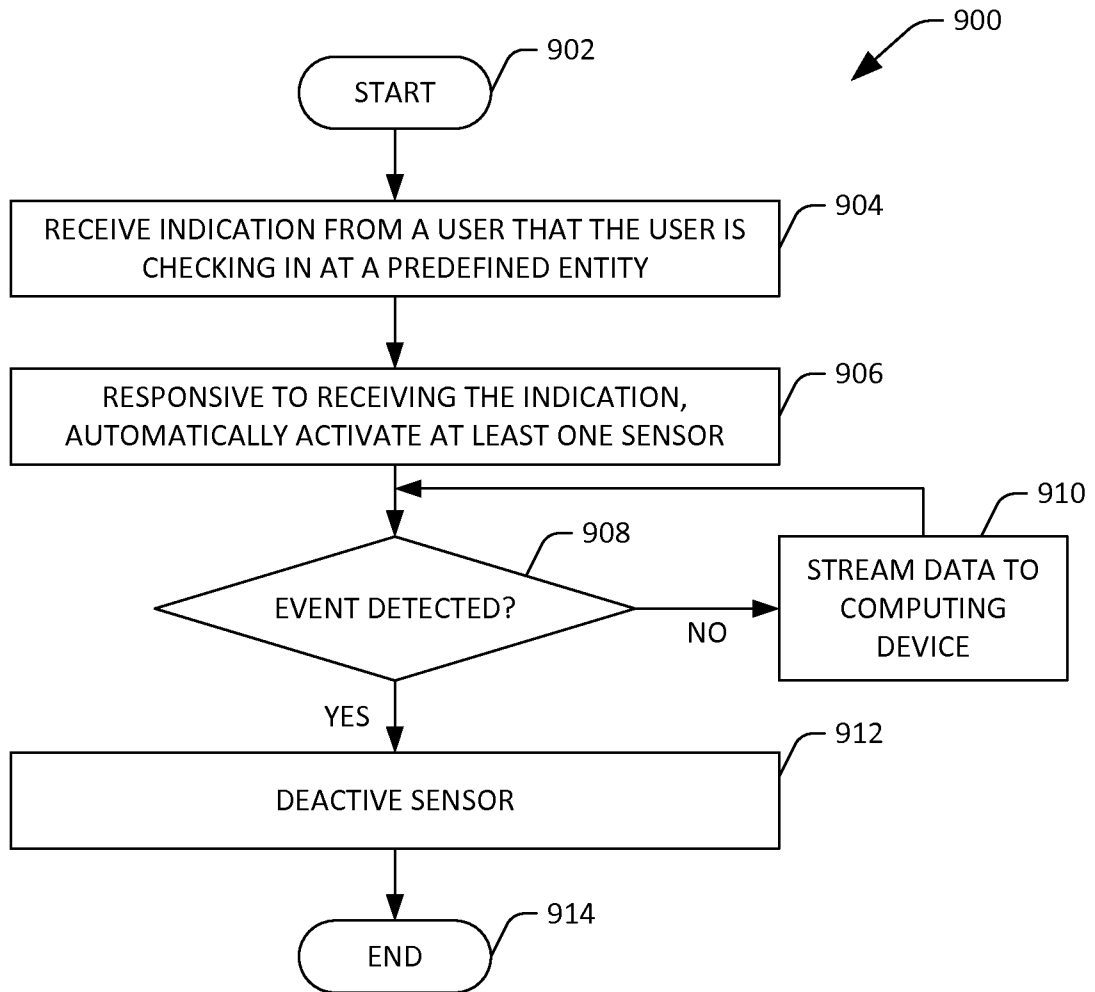


FIG. 9

9/12

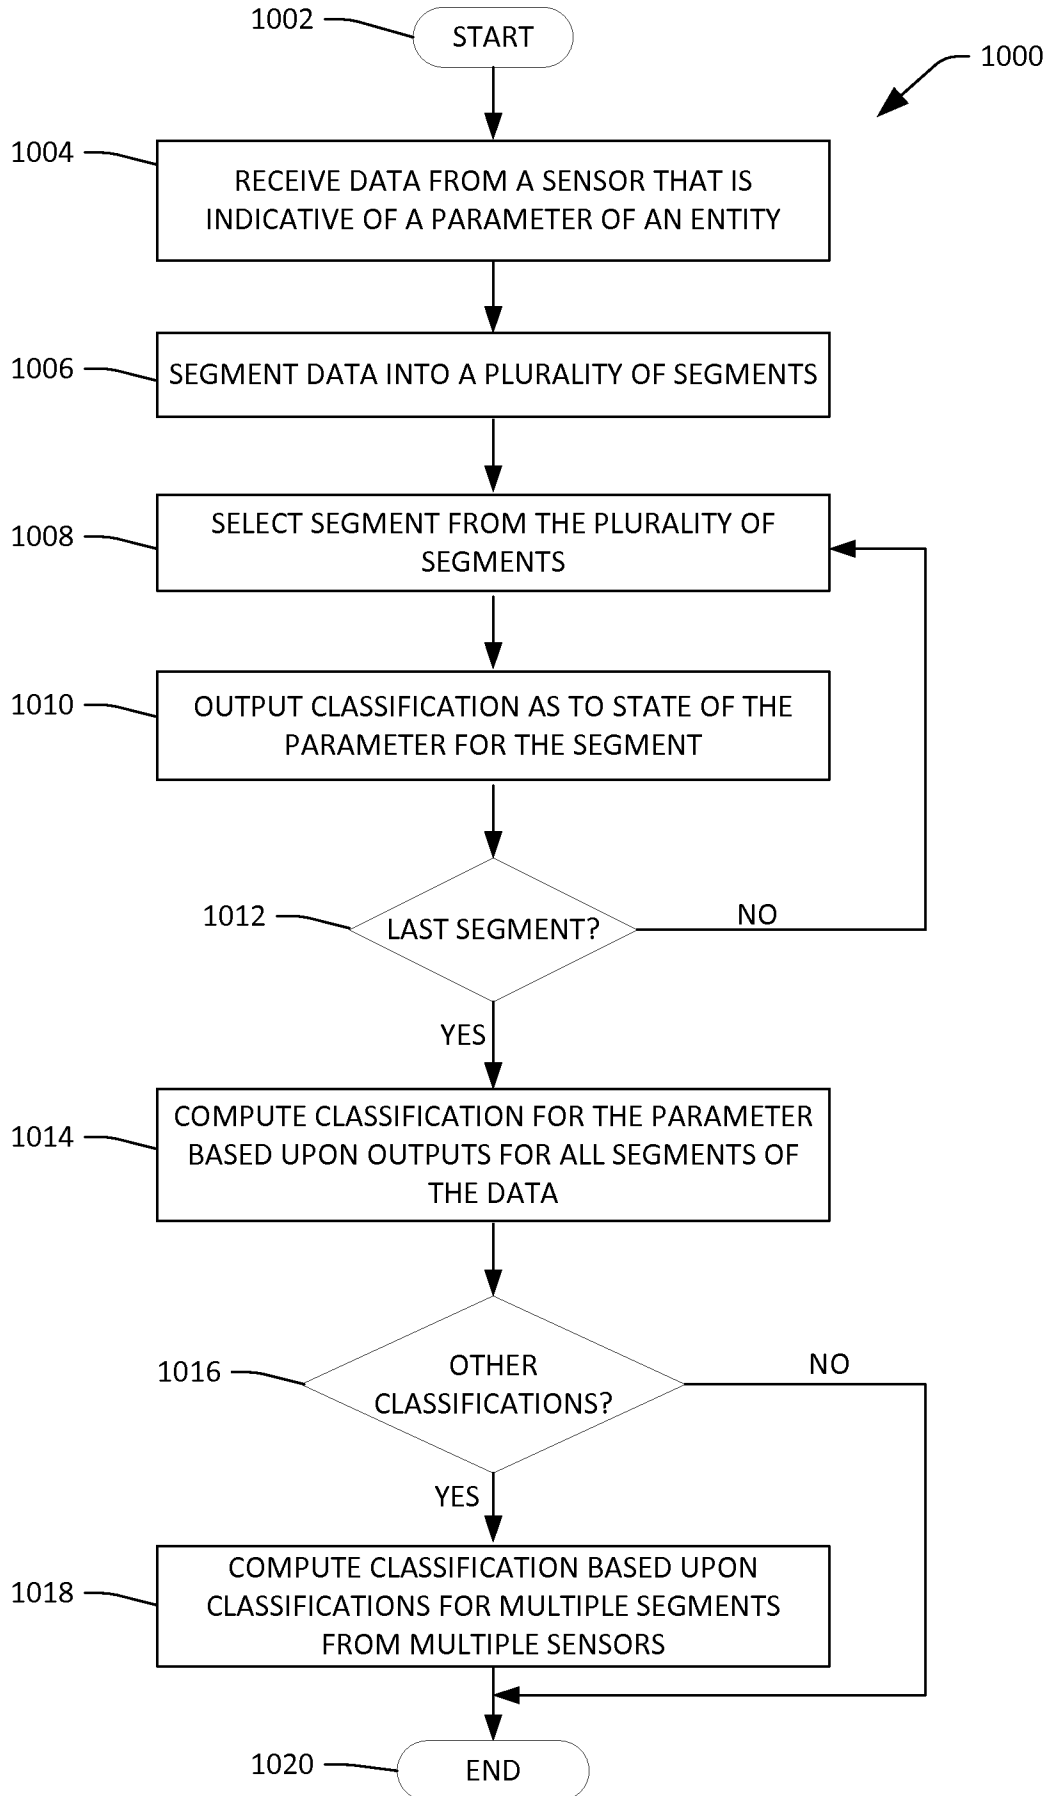


FIG. 10

10/12

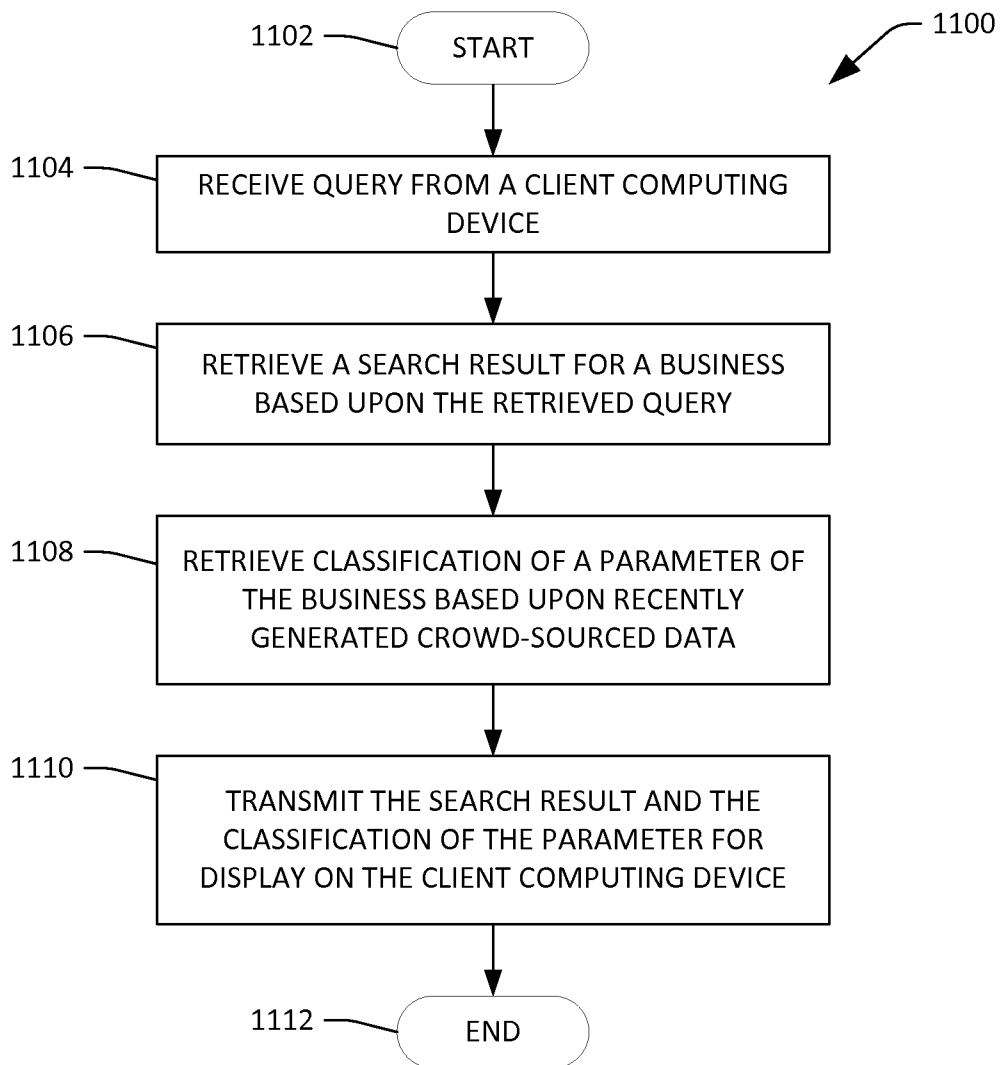


FIG. 11

11/12

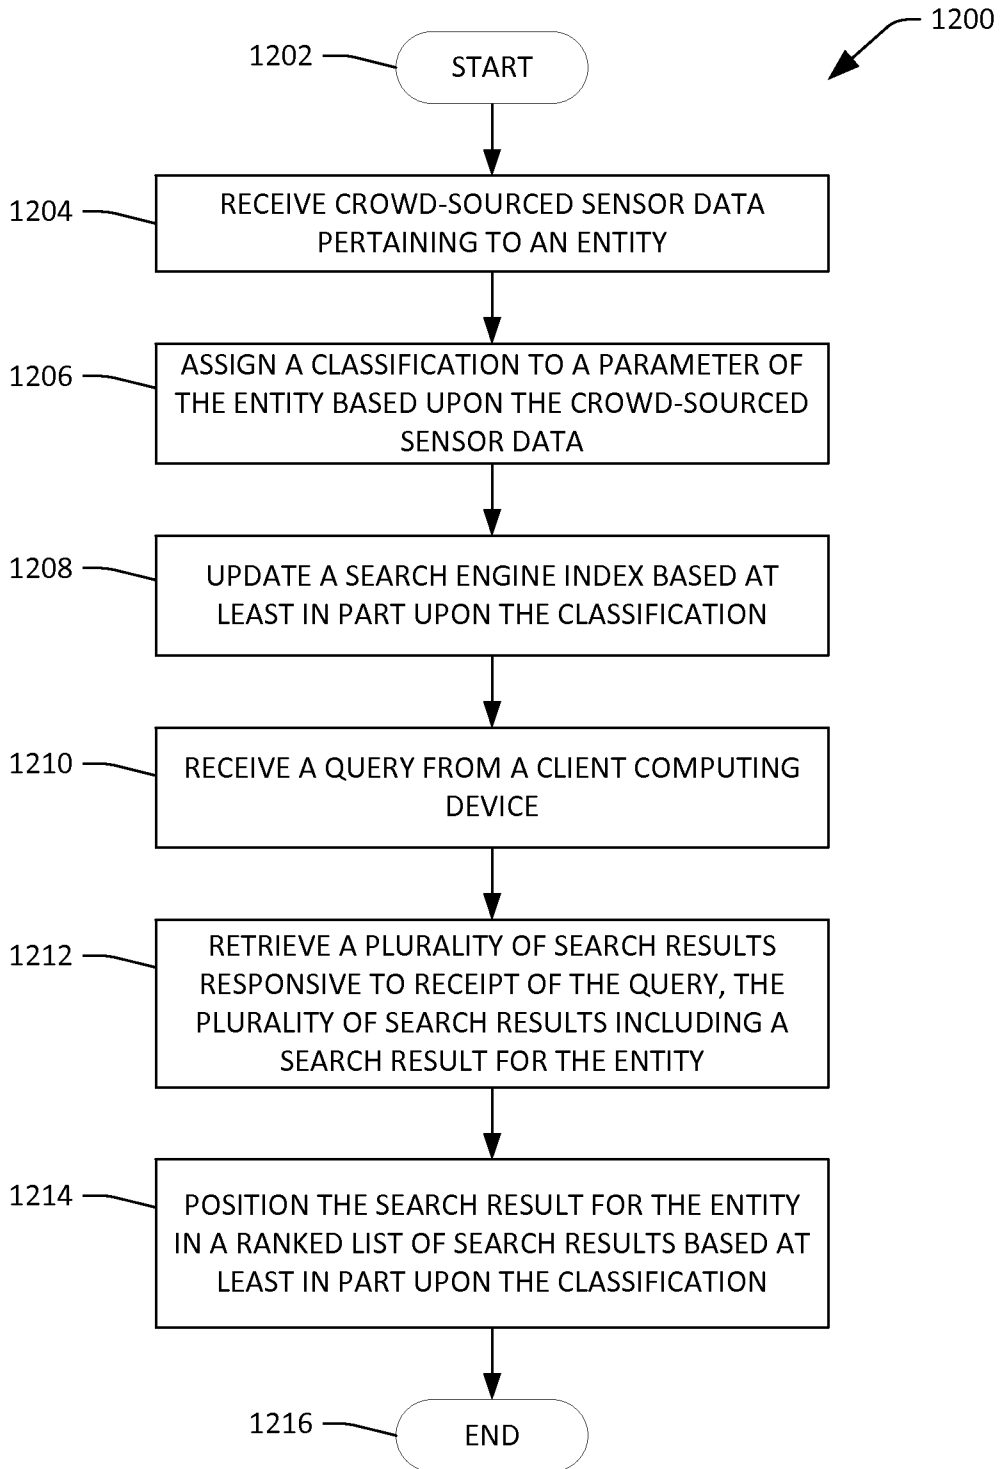


FIG. 12




FIG. 13