



(51) International Patent Classification:

G06F 19/24 (2011.01) *G06K 9/00* (2006.01)
G06F 17/30 (2006.01) *G06K 9/62* (2006.01)

(21) International Application Number:

PCT/US2016/024359

(22) International Filing Date:

25 March 2016 (25.03.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/177,825 25 March 2015 (25.03.2015) US

(71) Applicant: **CARNEGIE MELLON UNIVERSITY** [US/US]; Center For Technology Transfer and Enterprise Creation, 4615 Forbes Avenue, Suite 302, Pittsburgh, PA 15213 (US).(72) Inventors: **LAPUT, Gierad**; 640 Maryland Avenue, 2R, Pittsburgh, PA 15232 (US). **HARRISON, Christopher**; 4927 Penn Avenue, Pittsburgh, PA 15224 (US). **BIGHAM, Jeffrey, P.**; 1401 Beechwood Boulevard, Pittsburgh, PA 15217 (US). **LASECKI, Walter, S.**; 2609 Grant Drive, Ann Arbor, MI 48108 (US). **XIAO, Bo, Robert**; 6347 Marchand Street, Apt. 2, Pittsburgh, PA 15206 (US). **WIESE, Jason**; 45 Lansing Street, Unit 501, San Francisco, CA 94105 (US).(74) Agents: **OBERDICK, David, G.** et al.; Meyer, Unkovic & Scott LLP, 535 Smithfield Street, Suite 1300, Pittsburgh, PA 15222 (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).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR ADAPTIVE, RAPIDLY DEPLOYABLE, HUMAN-INTELLIGENT SENSOR FEEDS




Fig. 1

(57) Abstract: The disclosure describes a sensor system that provides end users with intelligent sensing capabilities, and embodies both crowd sourcing and machine learning together. Further, a sporadic crowd assessment is used to ensure continued sensor accuracy when the system is relying on machine learning analysis. This sensor approach requires minimal and non-permanent sensor installation by utilizing any device with a camera as a sensor host, and provides human-centered and actionable sensor output.

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

TITLE

System and Method for Adaptive, Rapidly Deployable,
Human-Intelligent Sensor Feeds

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit under 35 U.S.C. § 119 of Provisional Serial No. 62/177,825, filed March 25, 2015, which is incorporated herein by reference.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] This invention was made with Government support under IIS1443760 awarded by the National Science Foundation. The Government has certain rights in this invention.

BACKGROUND OF THE INVENTION

[0003] This invention relates to adaptive sensors. More particularly, this invention relates to sensors capable of detecting a wide-range of conditions or events and deployable with minimal development because the sensor output is generated initially with the assistance of crowd workers until machine-learning aspects of the sensor system are able to provide accurate output.

[0004] For decades, “smart” environments have promised to improve our lives by inferring context, activity, and events in diverse environments, ranging from public spaces, offices, and labs, to homes and healthcare facilities. To achieve this vision, smart environments require sensors, and many of them. These systems are expensive because they require specific sensors to be placed throughout the environment. In addition, these systems are special purpose and often invasive (e.g., running power to several sensors).

[0005] An even more challenging problem is that sensor output rarely matches the types of questions humans wish to ask and expected of an intelligent environment. For example, a door opened/closed sensor may not answer the user's true question: "Are my children home from school?" Similarly, a restaurateur may want to know: "How many patrons need their beverages refilled?" and graduate students want to know "Is there free food in the kitchenette?"

[0006] Unfortunately, these sophisticated, multidimensional, and often contextual questions are not easily answered by the simple sensors deployed today. Although advances in sensing, computer vision (CV), and machine learning (ML) have improved sensing abilities, systems that generalize across these broad and dynamic contexts do not yet exist, while less adaptive systems are expensive and complicated. For example, some smart environments monitor instrumented space (scanners, RFID tags, etc.), making the system specific to a particular question to be answered and expensive due to the number of sensors and other equipment. It would therefore be advantageous to develop a sensor system that answers a broad range of user-defined questions based on real-time sensor feeds with reduced complexity and costs.

BRIEF SUMMARY OF THE INVENTION

[0007] According to one embodiment of the present invention is a sensor system that allows minimal and non-permanent sensor installation and is capable of generating human-centered and actionable sensor output, and a method of using the same. In one embodiment of the present invention, answers from crowd workers and machine-learning approaches are fused to provide instant, human-intelligent sensor systems, which end users can set up in minutes. This system can enable a variety of rich end-user

applications and moves the field of human-computer interaction closer to responsive, intelligent environments. In one embodiment, the sensor system uses sporadic crowd assessment and machine learning to ensure continued sensor accuracy.

[0008] This sensor approach requires minimal and non-permanent sensor installation by utilizing any electronic device with sensing capabilities as a sensor host (such as camera on a cellular phone, for example). The sensor system is capable of answering diverse natural language “sensor” questions with reasonably high accuracy. In addition, the sensor system is easy enough to be used by non-experts, requiring zero training and providing live sensor output within seconds.

[0009] An example implementation includes an application programming interface (API) to enable access to sensor data streams. The API further provides a mobile application for easy, end-user authoring of on-demand sensor questions needed to create the intelligent sensors. The API also provides a tool for end-user programming of case-based events that turn sensor output into meaningful actions, such as providing an electronic alert.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0010] Fig. 1 is a block diagram of the sensor system architecture, according to one embodiment.

[0011] Fig. 2 presents an illustrated overview of the sensor system architecture.

[0012] Fig. 3 shows examples of image obfuscation methods employed by the system, including but not limited to (from top to bottom): raw image, light blurring, heavy blurring, and edge masking.

[0013] Fig. 4 presents the web API, according to one embodiment.

[0014] Fig. 5 shows an image and question presented to crowd workers.

[0015] Fig. 6 shows a plot of the accuracy by the system over time, where the sensor system can toggle between the crowd and machine learning to adapt to environment changes.

[0016] Figs. 7A-7D show a sensor image time series. Associated questions ask “Do you see a parked car?” (Fig. 7A), “How many dishes are in the sink?” (Fig. 7B), “Do you see a dishwasher door?” (Fig. 7C), and “How orderly is the line?” (Fig. 7D).

[0017] Figs. 8A-8C illustrate ‘live’ accuracy evaluation of three sensors. This is the accuracy of a sensor over the course of its deployment, as viewed “live” by its operator. X-axes represent training set cutoff time t .

[0018] Fig. 9 illustrates “future” accuracy evaluation of three sensors. This is the accuracy of the sensor assuming ML handoff at time t . X-axes represent ML handoff time t .

DETAILED DESCRIPTION OF THE INVENTION

[0019] In accordance with one embodiment, Fig. 1 shows the architecture of a sensor system 100. In this embodiment, the system 100 comprises a sensor 101 that provides a sensor feed for further processing and analysis by different components of the system 100. In this particular embodiment, the sensor 101 comprises a mobile electronic device with a camera. Mobile electronic devices, such as cellular phones, are especially suited for use as a rapidly deployable sensor 101 because of their built-in computer, battery, and communication abilities. However, many other devices capable of providing sensor output can be used, such as tablets, security cameras, wi-fi cameras, and Kinects. The time-series data from these types of sensors 101 offer rich contextual information about an

activity or environment far more than what basic sensors (e.g., proximity sensors) can provide.

[0020] The feed or data from the sensor 101, such as an image, is sent to the interaction module 102, which provides a software-based interface to allow a user develop an initial sensor configuration. More specifically, the user interaction module 102 provides a platform for the user to select a region of the image captured by the sensor 101 that is relevant to the question the user wants answered. The interaction module 102 further allows the user to define a question and select an obfuscation method, if desired. For example, if the user wants to know if there are available parking spots in front of his apartment, using the interface (such as a mobile application (App.), web portal, or other software-based interface) provided by the interaction module 102, he would select the area of image corresponding to the road immediately adjacent to his apartment. Next, the user would enter “Is there an available parking spot?”

[0021] Once the initial setup is completed, the sensor 101 begins sending a sensor feed (or image stream for a camera-based sensor 101) to an image similarity detector 140 and a dispatcher 110. (While Fig. 1 is discussed in terms of images captured by the sensor 101, alternate embodiments can utilize other types of data, such as audio, for example, depending on the type of sensor 101 used.) The image similarity detector 140, which will be discussed in further detail, determines if the latest sensor feed has changed from the previously analyzed sensor feed. If no changes are present, the prior sensor value is used. If there are changes, indicating that the environment being sensed by the sensor 101 has changed, the sensor feed is passed to a dispatcher 110. The dispatcher 110 determines if the data set, comprising the sensor feed and question, should be sent to a machine learning module 130 or a crowd module 120, or both. The action of the dispatcher 110 depends partially on the state of the system 100.

[0022] For example, when the sensor system 100 is first configured, the dispatcher 110 sends the data set to the crowd module 120. Here, the data set (i.e., the question and a subregion of the sensor image in this example embodiment) are sent to the crowd 121, or a network of distributed human workers that process the request through a web-based interface. Crowdsourcing allows the system 100 to access human intelligence through online marketplaces, such as Mechanical Turk for example. Alternatively, the crowd 121 could comprise the user of the sensor system 100, who would be tasked with answering her own question, possibly for several images, before ML techniques are used to analyze the data set.

[0023] To ensure the accuracy of the crowd 121, the crowd module 120 further comprises an optional quality assurance module 122, which determines a consensus answer generated by the crowd 121. In one embodiment, several answers are collected per image, then fused together using the quality assurance module 122 (*e.g.*, voting) to determine the best response for a given instance. The quality assurance module 122 also can be used to reduce noise and reject malicious workers.

[0024] After several initial data sets are analyzed by the crowd module 120, subsequent data sets can be sent to the machine learning module 130 by the dispatcher 110. The machine learning module comprises a trainer 131 and a predictor 132. The trainer utilizes the images labeled by the crowd module 120 to develop a machine learning classifier, which is used by the predictor 132 to label images captured by the sensor 101.

[0025] The labeled sensor feeds 103, created by the crowd module 120 or the machine learning module 130, are ultimately sent to an end user interface 150. However, prior to being sent to the end user interface 150, the labeled sensor feeds 103 from both the crowd module 120 and machine learning module 130 are sent to a comparator 160 to evaluate whether the answer to a questions for a

particular sensor feed generated by the machine learning module 130 is similar to the answer generated by the crowd module 120. If the machine learning module 130 is producing reasonably accurate output (assuming the crowd module 120 is providing superior output based on direct observation and human level intelligence), then the sensor system 100 will continue relying on the machine learning module 130. Conversely, if the accuracy of the answers falls below a threshold, the dispatcher 110 will be instructed to send the data set to the crowd module 120 again. Subsequent answers from the crowd module 120 can be used to provide further training for the machine learning module 120 in this situation.

[0026] The end user interface 150, comprising an App. in one embodiment, allows the end user to visualize the labeled sensor feeds 103, set notifications based on the labeled sensor feeds 103, and otherwise take action based on the data received from the upstream components of the sensor system 100. The interaction module 102 and end user interface 150 can be part of the same App. in some embodiments.

[0027] In an alternative embodiment, the interaction module 102 and end user interface 150 are a web-based interface that allows sensor management for multiple sensors 101. In this embodiment, users link one or more sensors 101 to a web account, where they can create, modify, and synchronize sensors 101 across all of their devices. The web-based user interface also makes it possible to create new sensors 101 remotely. For example, users can “stick” a sensor device 101 at an elevated vantage point (e.g., for viewing an entire parking lot), and then manage and create the sensor 101 without having to physically touch the device 101.

[0028] Additionally, the end user interface 150 allows users to design event-based notifications using data from one or more sensors 101 (e.g., “Send an email when the stove is ON and ZERO people are in the house.”). These directives can be chained together as set of

conjunctions (“and” clauses), or for alternative responses (“or” clauses). Multiple chains can be defined to represent different configurations. These disjunctions of conjunctions comprise a fully expressive set of logical terms. One implementation allows users to select from an email, text message, audio, or chat alert. For each alert type, users are able to define a custom message, which can also display the current value of the sensor by using a specially marked variable.

[0029] Fig. 2 depicts an example implementation of the sensor system 100. The following description outlines how a user interacts with this example implementation. In the example system 100 shown in Fig. 2, a restaurant proprietor uses a smartphone (or tablet) and affixes it to the wall of his restaurant. He installs and launches an App. on the phone to load the interaction module 102. In this manner, a mobile application serves as the primary end-point for creating and modifying the sensor system 100. The phone uses the front facing camera as the sensor 101 to provide a live overview of the restaurant. On the App. provided by the interaction module 102, he presses the “New Sensor” button and circles the bar countertop, thus specifying a region of interest. He then enters a plain text question, for example: *“How many drinks are almost empty?”*

[0030] By pressing “Go”, the sensor 101 is activated and starts providing real-time data. In other words, the device periodically takes snapshots of the region of interest and forwards the data to the dispatcher 110. He can now use the end user interface 150 (or App.) to see a real time visualization of how many drinks need to be refilled, or use an end user programming tool 151 to have the system automatically message a co-worker requesting help if the number exceeds a certain limit, as shown in the lower-left corner of Fig. 2. Within a few minutes, he could similarly set up sensors for: *“Does table four need to be cleaned?”*, *“Are customers wearing their coats*

inside?”, “*Is a check sitting on the table?*” and other questions relevant to the dining experience.

[0031] Unbeknownst to the user, his sensor system 100 is initially powered by crowd workers 121 interpreting his plain text question, providing immediate human-level accuracy, as well as rich, human-centered abstractions. However, because using crowd workers 121 can be costly and difficult to scale, the crowd 121 is only used temporarily.

[0032] Seamlessly to the user, the answers from the crowd 121 are used to train and test image-based machine learning classifiers in the machine learning module 130. The machine learning module 130 tests against the ground truth provided by the crowd module 120, where the ground truth is information provided by direct observation (i.e. crowd workers 121) rather than by inference. If the machine learning module 130 begins to achieve human-like accuracies, the labeled sensor feeds 103 from the machine learning module 130 will be used alongside the feeds 103 provided by the crowd module 120.

[0033] Eventually, if sufficiently robust, the machine learning module 130 is used exclusively. This human-computer handoff is seamless and invisible to end users; as far as users are concerned, they have a sensor system 100 with human-level accuracy from initial set-up. Even if the machine learning module 130 does not achieve the needed level of accuracy, savings measures have been designed into the crowd-based method to conserve costs.

[0034] Because image data is being sent to the crowd 121, users may desire to obfuscate the image captured by the sensor 101 to preserve privacy. Utilizing the entire raw image can easily violate privacy, especially when personally identifying information is present, or when images depict people in sensitive situations. To mitigate this issue, the sensor system 100 asks users to select an arbitrarily shaped subregion relating to the question they wish to ask, where

image data outside the subregion is masked away. The region may be arbitrarily shaped so only the region pertinent to the question is selected. In other words, the shape is not limited to a square or other predefined shape. By selecting a subregion, the system 100 reduces file size, removes unnecessary image elements, and simplifies the sensing effort for both human raters and ML techniques.

[0035] For users wishing to add an additional level of privacy, the system 100 employs several image obfuscation methods. Obfuscation is applied to images before they are sent to the crowd 121 by selecting an obfuscation method in the interaction module 102. Users can choose to leave the image subregion unaltered, or select from four obfuscation methods: light blur, heavy blur, median filter, and edge filter. Examples of obfuscated images are shown in Fig. 3. The level of obfuscation will depend on the level of privacy desired and the need for detail in the image to answer the question. For instance, in the drink example, a higher level of image detail would be required to determine if a drink refill is needed.

[0036] Questions define the “capabilities” of a sensor, and thus, the quality of answers depends on several factors, such as the question’s context and relevance. To further add context and relevance, in one embodiment, the system 100 requires users to define a data type when creating new sensors. Data types curb variance and constrain the range of possible values returned by a sensor system 100 (i.e., the answer to the sensor’s question), and facilitate simple automated processing of the data stream. To this end, the system 100 categorizes questions into four example data types:

[0037] Yes/No – This type of question can be answered by either a ‘Yes’ or ‘No’ response. It is analogous to an ON/OFF sensor mechanism. Examples include: “Is the door open or closed?”, “Is there food in the kitchen?”, or “Can you see a bus in this image?” Unlike ON/OFF sensor mechanisms, however, the Yes/No question may still require context to be fully understood.

[0038] Number – Number data types are intended for questions that require counting. Numbers are continuous and are bound between a minimum and maximum range. Examples include: “How many cars do you see in the parking lot? (min=0, max=30)”, and “What percentage of the water tank is full? (min=0, max=100)”

[0039] Scale – Scale data types are analogous to Likert-scale questions. Answers belong to discrete values specified within an increasing/decreasing scale. For this data type, users are required to supply scale-value pairs. Examples include: “How messy is this table? (1= Clean, 2=Average, 3= Messy)”, or “How happy does the person look? (1=Very Sad, 2=Sad, 3=Neutral, 4=Happy, 5=Very Happy)”

[0040] Multiple Choice – When creating multiple choice questions, users are required to specify the list of choices. Unlike scale data types, choice order is irrelevant. Examples include: “What type of food do you see? (None; Indian; Thai, Asian; Salad; Bagels; Other)” and “What are people doing? (Reading; Using computers; Eating; Other)”.

[0041] As part of the initial set-up when creating sensors, users can specify the frequency at which sensor 101 feeds are generated. Depending on the question, frequency readings can range from near real-time (e.g., every one or two seconds for questions like “Is the refrigerator door open? [Yes/No]”), to extended periods (e.g., once per day for questions like “What product is advertised on the billboard? [MultipleChoice]).” Referring to Fig. 4, which is a screenshot of an interface provided by the interaction module 102 according to one embodiment, the frequency selection is shown in the third input box from the top.

[0042] As previously discussed, the sensor system 100 comprises an image similarity detector 140. Sensor image streams often have periods of little variation (e.g., buildings after closing hours, outdoor scenes at night). Thus, to avoid soliciting redundant and thus costly

responses from the crowd 121 on highly similar images, the image similarity detector 140 collapses runs of similar images using a simple technique, in one embodiment, where the number of changed pixels between two images is analyzed.

[0043] Using this method, the detector 140 counts the number of pixels that have changed from the previous frame using the image's RGB values and compares the count against a predetermined pixel difference threshold. Additionally, if the number of changed pixels in an image exceeds a certain image area percentage threshold, the detector 140 considers the image to be different. While this simple technique is used in this example embodiment, more sophisticated approaches, such as those described by Khan, Acharya, and Vern (Comparison between different illumination independent change detection techniques; in Proc. ICCCS '11) could be used in alternative embodiments.

[0044] By way of further example of one embodiment of the image similarity detector 140, its optimal parameters were determined from a corpus of roughly 6000 time-stamped images taken from multiple sensor streams. Each image was manually labeled as to whether or not it was the same as the previous image, providing a ground truth set. The image similarity detector 140 was then run, seeded with all combinations of the following thresholds: i) 2% to 40% pixel difference threshold, in 2% increments, and ii) 0.1% to 5.0% image area percentage threshold, in 0.1% increments. This produced 130 result sets, which was compared to the ground truth using Jaccard's distance metric. By using a pixel difference and image area threshold of 10% and 1.0% respectively, a Jaccard distance of .64 is achieved. On average, this process removes roughly 40% of images – a significant saving.

[0045] Referring again to the figures, crowd workers 121 process unanswered instances through a web-based interface, as shown in Fig. 5. The interface can vary based on the question/response type.

Each sensor instance is answered (i.e., labeled) by several different crowd workers and voting is used to determine the best response using the quality assurance module 122. For example, in one embodiment, three crowd workers 121 answer each sensor instance and if an answer is provided by two or three of the workers 121, then that answer is used. Greater or fewer workers can be used depending on the complexity of the question. If the question is relatively simple, the quality assurance module 122 can be omitted and only a single worker is used. The answers provided by the crowd workers 121 are stored in a database for further processing by the sensor system 100.

[0046] For the system 100 to work efficiently, the sensor system 100 presents workers 121 with a simple task that they can answer quickly. As such, one image is presented to each worker and a single response is collected. For example, see Fig. 5, which depicts one image and one question presented to the worker 121. If workers 121 are unable to answer an image-based question, they can mark it as an exception. For example, the representative interface shown in Fig. 5 includes an “I Can’t Tell!” button 201, which informs the system 100 that there is something amiss with the sensor 201 itself (e.g., occlusion, insufficient lighting) or that the question was poorly defined. In addition, workers 121 are prompted to provide textual descriptions when exceptions occur. This approach provides actionable user feedback to help remedy the problem.

[0047] To recruit users fast enough to receive answers in real time, one embodiment of the system 100 uses LegionTools, a toolkit for quickly recruiting workers from Mechanical Turk using a web-based interface. It leverages a retainer model, which pre-recruits workers so they are ready to respond to a task within as little as two seconds. When sensors 101 are crowd-powered, this sets the lower-bound on the system latency. As such, answers to sensor questions can be provided in nearly real-time. In alternative embodiments, the

crowd module 120 sources answers in a different manner, yet still relies on the crowd 121.

[0048] Because solutions using crowd-power alone can be costly and difficult to scale (i.e., more sensors 101 requires more people), as previously discussed, the sensor system 100 reduces its reliance on crowd workers 121 over time by using crowd-provided answers to train the machine learning module 130 comprising machine learning classifiers, which are fast and inexpensive. However, even after machine learning has taken over processing sensor feeds, crowd workers 121 are still needed to provide a periodic human-accuracy baseline to ensure high accuracy.

[0049] To initiate the machine learning module 130, the trainer 131 receives all prior input data produced by the crowd module 120, except a portion of the data will be used for evaluation rather than training purposes. In one training technique, histogram equalization is applied to each image to reduce the effect of lighting variances. Then, each input image is processed into a large number of global features. Additionally, each sensor image (which is itself a subregion of the larger original image) is broken into a grid of sub-images. In addition to a 1x1 grid (simply the image unchanged), the trainer 131 also uses 4x4, 7x7 and 10x10 grids, for example. Each of these variously-sized sub-images is then converted to a luminance image, and the mean, standard deviation and mean-squared error across the window are used as numerical features. This produces a total of 332 image features. Feature selection is used as a post-process to extract exemplar features for a given sensor feed. While this embodiment describes one such method, a person having skill in the art will understand that other machine learning techniques can be used.

[0050] The machine learning module 130 uses correlation-based feature selection to select features from the training set, coupled with a backtracking best-first attribute searcher. Both algorithms are implemented in the Weka machine learning toolkit, in one

embodiment. Other machine learning frameworks can be used in alternative embodiments, based on similar techniques. The selection process typically chooses between 10 and 30 features. Notably, the feature set is typically different for different sensors/question environments. That is, there is little overlap in the selected feature set for different sensor implementations.

[0051] After the feature set is selected, the trainer 131 constructs classifiers. In one embodiment, a “pre-classifier” is first trained to distinguish between exceptions and non-exceptions, to ensure that the main classifier is not polluted by incorrect data. For continuous (numeric or scale) sensors, a support vector machine regression classifier is trained using the SMOReg algorithm, for example. For discrete sensors (yes/no, multiple choice), a one-versus-one multiclass quadratic-kernel SVM trained with the SMO algorithm is used. For simple binary sensors, a single SVM is trained. In one embodiment, the SVM is used as the basic classifier because of its ease of training and predictable behavior, though other classification approaches can be used as valid techniques. In an alternative embodiment, the sensor system 100 maintains a library of feature extractors and classification algorithms as part of the machine learning module 130 and selects those exhibiting the best performance for a given sensor 101 implementation.

[0052] As the training corpus grows from crowd labeled instances, the accuracy of the machine learning module 130 typically improves. Once the accuracy of the machine learning exceeds a predefined threshold, such as 95% in one embodiment, the sensor hands off classification tasks to the machine learning module 130. It is also possible to do a soft handoff, where the relative weighting between crowd module 120 and the machine learning module 130 shifts over time.

[0053] To ensure continued accuracy after the handoff to the machine learning module 130, the system 100 periodically submits

small batches of fresh sensor feeds to the crowd 121 for labeling, which is done to benchmark the accuracy of the machine learning module 130. If accuracy is sufficiently robust, the machine learning module 130 continues to power the sensor system 100. However, if the accuracy has fallen below a threshold, the system 100 can revert to crowd-power. This serves two immediate purposes: 1) the sensor system 100 immediately regains human-intelligence level accuracy, and 2) the system 100 can incorporate new labels in its training for a future handoff to the machine learning module 130.

[0054] Fig. 6 shows the relative accuracy for answers generated by the crowd module 120 and the machine learning module 130. As shown at the top of the graph in Fig. 6, the sensor system 100 initially relies on the crowd module 120 for a period of time until the machine learning module 130 reaches similar accuracy. Periodic validation of the machine learning module 130 is conducted by polling the crowd 121 at regular intervals. In this manner, the sensor system 100 automatically handles infrequent changes (such as the first snow fall in a parking lot) that would prove challenging for most computer vision-driven systems, which are trained before being deployed. Referring again to Fig. 6, this change is shown as a dip in the accuracy of the machine learning module 130, where the system reverts to the crowd module 120 prior to a re-handoff to machine learning module 130 when its accuracy improves. This ability to seamlessly toggle between crowd and automatic approaches, without sensor interruption, makes the sensor system 100 highly adaptive and robust.

[0055] In one example embodiment, the mobile application (or user interface provided by the interaction module 102) was developed on Android, and deployed on a fleet of low cost tablets; however, any smart device or operating system may be used. The user interface 102, backend server running the various modules, and end-user interface 150 were implemented using PHP, MySQL, jQuery, d3.js,

node.js, and socket.io. The machine learning module 130 utilized the Weka toolkit. A web API that allows services and end-user applications to leverage sensor data can also be provided.

[0056] Figs. 7A-7D show several examples of sensor feeds for various sensor systems 100. In this example, 16 sensors 101 were deployed across four diverse environments: a parking lot (Fig. 7A), a home kitchen (Fig. 7B), office kitchenette (Fig. 7C), and a building food court (Fig. 7D). Sensor questions ranged from “Do you see a parked car?” to “How many dishes are in the sink?” Moreover, a range of reporting frequencies (from once per minute to twice per hour) and deployment durations (10 days to 3 weeks) were represented in the example sensor systems 101 shown in Figs. 7A-7D.

[0057] To analyze how well the sensors 101 can quickly provide accurate sensing data, the precision, recall, and latency of the aggregated crowd responses were measured for another set of sensor systems 100. Figs. 8A-8C show the accuracy of crowd workers' ratings, using expert labels generated by the user--who had full understanding of the context of the question--as the ground truth. That is, Figs. 8A-8C show the accuracy of a sensor system 100 over the course of its deployment as viewed ‘live’ by the user. Fig. 8A asked the question “Do you see a dishwasher door?”; Fig. 8B asked “How messy is the counter?”; and Fig. 8C asked “Which parking spots are occupied?”

[0058] Cohen's kappa is calculated to mitigate the effects of skewed class distribution (e.g., a leftover food sensor 101 returned “no” over 75% of the time in another demonstration). Crowd accuracy reaches as high as 96.8% (kappa score 0.859), with mean accuracy 77.4% (median 76.0%). The crowd 121 performed very well on three sensor systems 100 (accessible parking spots occupied, number of cars in parking lot, and dishwasher door), moderately well on one sensor system 100 (leftover food), and poorly on three sensor systems 100 (food type, line length sensor, countertop messy).

[0059] The food type sensor system 100 required users to distinguish between seven types of cuisine (“I do not see any food”, “Pizza”, “Sandwiches”, “Cake or pastries”, “Asian or Indian”, “Salad”, “Bagels or doughnuts”, “Other cuisine or I can't tell”) based on a very low-resolution image, while the line length sensor and countertop sensors both involved subjective judgments (e.g. “Is the line orderly?” and “How messy is the countertop?”). These are subjective or open-ended questions. By contrast, quantitative questions (“Is there food here?”, “Is the door closed?”, and “How many cars are there?”) generally had superior performance.

[0060] In designing questions to be posed to the crowd 121, operators may make assumptions that are not obvious to crowd workers 121, leading to incorrect results. In one example, workers 121 were asked to identify the presence of food on a kitchen countertop. The countertop has a permanent candy jar, which the user assumed would not be classified as food, yet several crowd workers 121 marked the otherwise empty countertop as having food. Based on the observed results, the question was amended to explicitly exclude candy, after which the expected results were obtained.

[0061] The demonstration sought to estimate the accuracy of a sensor system 100, over the course of its deployment, as viewed ‘live’ by the user. For each sensor, ten time periods each covering one-tenth of the data were defined, numbered $t=0.1$ through $t=1.0$. To estimate live accuracy at time t , all data up to time t was used for training, and then tested on all data from time t to time $t+0.3$ (i.e. it was tested on a sliding window of 30% of the data). The results for three representative sensor systems 100, as shown in Figs. 8A-8C, are compared against the crowd accuracies. In many cases, a relatively small portion of the data is needed to reach crowd-level accuracies.

[0062] Alternatively, it is equally important to assess what the accuracy of a sensor system 100 would be going forward, assuming a ML handoff occurs at time t . To assess this, a complete ML handoff is

simulated at each time increment. All data up to that point is used for training, while all future data is used for testing. This analysis is stopped when less than 30% of the data is available for testing, to avoid noisy results from insufficient test data. These results are summarized in Figs. 9A-9C, compared against the overall crowd accuracies. The accuracies follow similar curves to the curves of Figs. 8A-8C, suggesting that live classification accuracy may be able to predict future post-handoff classification accuracy.

[0063] For underperforming sensor systems 100, poor accuracy is often the result of questions that are subjective or required additional context. For example, one sensor asked, “How orderly is the line?” (Fig. 7D) with three possible answers: “No people visible in image”, “People present, but no obvious organization”, and “People standing in a line.” Because the answer to whether a line is organized is subjective (e.g., relative to local cultural norms), crowd workers 121 will provide widely varying answers. Another sensor system 100 was tasked with sensing whether a dishwasher was opened or closed (see Fig. 7C). When the question was initially defined as “Is the dishwasher door open?”, crowd workers 121 were confused, reducing sensor accuracy. In the image shown, no dishwasher was visible in the closed state due to the overhead view and obstruction by the countertop. When presented with the question of “Is the dishwasher door open?”, the crowd 121 presumably wondered “What dishwasher?” Rephrasing the question to be more context-free--“Do you see a dishwasher door?”--significantly boosts accuracy.

[0064] There are a number of ways to alleviate “bad” sensor questions. In one embodiment, the system suggests example questions or provides structured question templates (e.g., “Do you see a _____ [in/on] the _____?”), helping end-users formulate questions with less ambiguity. These sample questions can be based on the accuracy of prior sensor system 100 implementations. Additionally, the “I can’t tell” button 201 in the crowd interface could

allow the system 100 to flag sensors 101 causing confusion and suggest the question or image subregion be modified. Another approach is for the crowd labeling interface to provide exemplar images, one for each possible answer (e.g., show examples of both dishwasher door states). In this approach, the user could take images of the door in both the open and closed states to provide the example images. Finally, the system 100 can show crowd workers 121 a random set of previously collected sensor images that hopefully better capture the full range of possible states (e.g., orderly line to chaotic), so they can make better relative judgments.

[0065] While the disclosure has been described in detail and with reference to specific embodiments thereof, it will be apparent to one skilled in the art that various changes and modification can be made therein without departing from the spirit and scope of the embodiments. Thus, it is intended that the present disclosure cover the modifications and variations of this disclosure provided they come within the scope of the appended claims and their equivalents.

CLAIMS

What is claimed is:

1. An adaptive sensor system comprising:

a sensor, wherein the sensor provides a sensor feed of an environment;

a user interface,

wherein a user can enter a question to be answered by the sensor system using the user interface,

wherein the question relates to a condition of the environment,

an image similarity detector that receives the sensor feed, wherein the image similarity detector maintains a current sensor value if there is no change in the condition of the environment over a period of time;

a machine learning module capable of creating a labeled sensor feed;

a crowd module capable of creating an additional labeled sensor feed;

a dispatcher, wherein the dispatcher directs the sensor feed and question to at least one of the crowd module and the machine learning module,

a periodic comparator that evaluates a relative accuracy of the labeled sensor feed generated by the machine learning module as compared to the labeled sensor feed generated by the crowd module.

2. The adaptive sensor system of claim 1, wherein the dispatcher directs the sensor feed to the crowd module upon initiation of the

sensor system and prior to the machine learning module reaching a threshold level of accuracy.

3. The adaptive sensor system of claim 2, wherein the labeled sensor feeds generated by the crowd module are sent to the machine learning module,

wherein the machine learning module uses the labeled sensor feeds to train a classifier.

4. The adaptive sensor system of claim 3, wherein the dispatcher directs the sensor feed to the machine learning module when the machine learning module is operating above a threshold level of accuracy.

5. The adaptive sensor system of claim 1, wherein the crowd module sends the sensor feed and question to crowd workers to generate the labeled sensor feed.

6. The adaptive sensor system of claim 1, wherein the user interface provides a visualization of the labeled sensor feed.

7. The adaptive sensor system of claim 1, wherein the sensor comprises an electronic device having a camera.

8. The adaptive sensor system of claim 1, wherein the sensor feed comprises an image.

9. The adaptive sensor system of claim 1, wherein the question has a data type selected from the group consisting of: yes/no, number, scale, and multiple choice.

10. A method of creating an adaptive sensor system, comprising:

- providing a sensor that provides a sensor feed;
- receiving from a user interface a question, wherein the question is capable of being answer by observation of the sensor feed;
- sending the sensor feed to a crowd to generate a labeled sensor feed;
- using the crowd labeled sensor feed to train a machine learning module;
- sending a subsequent sensor feed to the machine learning module to generate the labeled sensor feed when the machine learning module exceeds a threshold level of accuracy; and
- sending the labeled sensor feed to a user via the user interface.

11. The method of claim 10, further comprising:

- periodically testing the labeled sensor feed generated by the machine learning module against the labeled sensor feed generated by the crowd to determine the relative accuracy of the machine learning module.

12. The method of claim 11, wherein the sensor feed is sent to the crowd if the machine learning module falls below the threshold level of accuracy.

13. The method of claim 10, further comprising:

choosing a region of interest in the sensor feed prior to sending the sensor feed to the crowd.

14. The method of claim 10, further comprising:

obfuscating the sensor feed.

15. The method of claim 10, further comprising:


rejecting a second sensor feed received subsequent to a first sensor feed if there is no difference between the first sensor feed and the second sensor feed.

16. The method of claim 15, wherein the second sensor feed is rejected if a pixel difference threshold does not exceed 10%.

17. The method of claim 15, wherein the second sensor feed is rejected if an image area percentage threshold does not exceed 1%.

18. The method of claim 10, wherein sending the sensor feed to a crowd to generate a labeled sensor feed further comprises:

sending an example sensor feed to the crowd.






Fig. 2



Fig. 3

X ZENSORS

Use this form to modify the settings of your sensor. All fields required.

Sensor Name
Food Type Sensor

Sensor Question
What type of food do you see?

Reading Frequency
Every 10 Minutes

Sensor Data Type
Multiple Choice

Choices

I do not see any food	X
Pizza	X
Sandwiches	X
Asian or Indian	X
I can't tell	X

Add New Choice

Obfuscation Method
Light Blur

Activate the Sensor?
Yes

**Fig. 4**

**Fig. 5**

201




Fig. 6

Fig. 7A




Fig. 7B




Fig. 7C








Fig. 7D

**Fig. 8A****Fig. 8B**

**Fig. 8C****Fig. 9A**




Fig. 9B




Fig. 9C

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/024359

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 19/24; G06F 17/30; G06K 9/00; G06K 9/62 (2016.01)

CPC - G06F 17/30256; G06F 17/30244; G06F 17/30864; G06K 9/0063; G06K 9/00771 (2016.01)

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)

IPC(8) - G06F 17/30; G06F 19/24; G06K 9/00; G06K 9/62 (2016.01)

CPC - G06F 17/30; G06F 17/30244; G06F 17/30256; G06F 17/30864; G06F 19/24; G06K 9/00; G06K 9/0063; G06K 9/00771; G06K 9/62; G06K 9/6253 (2016.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

USPC - 382/155.000; 382/159.000; 382/170.000; 382/181.000; 382/190.000; 382/218.000; 382/219.000; 707/759.000; 707E17023; 707E17020; 707E17015; 707E17014 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, Google Scholar, Google, Directory of Open Access Journals

Search terms used: machine learning, sensor, feed, question, query, inquiry, picture, image, same, similarity, identical, mirror, input, data, smart sensor, crowd, people, person, someone, dispatch, dispense, direct, redirect, intelligence, opinion, consensus, threshold, limit, classifier, yes/no, data type, multiple choice, pixel, image area, reject

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 20120197852 A1 (DUTTA et al) 02 August 2012 (02.08.2012) entire document	1-18
Y	US 7,961,906 B2 (RUEDIN) 14 June 2011 (14.06.2011) entire document	1-18
Y	EP 2347809 A1 (BLUE BANANA BVBA) 27 July 2011 (27.07.2011) entire document	1-18
Y	US 8,204,316 B2 (PANAHPOUR TEHRANI et al) 19 June 2012 (19.06.2012) entire document	16
Y	US 4,046,071 A (MIZUNO et al) 06 September 1977 (06.09.1977) entire document	17
A	US 8,756,173 B2 (HUNZINGER et al) 17 June 2014 (17.06.2014) entire document	1-18
A	US 20130077820 A1 (MARAIS et al) 28 March 2013 (28.03.2013) entire document	1-18
A	US 20140101083 A1 (NICHOLSON et al) 10 April 2014 (10.04.2014) entire document	1-18
A	US 20130150117 A1 (RODRIGUEZ et al) 13 June 2013 (13.06.2013) entire document	1-18
A	US 8,971,612 B2 (SHOTTON et al) 03 March 2015 (03.03.2015) entire document	1-18
A	US 20130218816 A1 (YU et al) 22 August 2013 (22.08.2013) entire document	1-18



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

"&" document member of the same patent family

Date of the actual completion of the international search

11 May 2016

Date of mailing of the international search report

21 JUN 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774