



- (51) **International Patent Classification:** Not classified
- (21) **International Application Number:** PCT/US2014/015654
- (22) **International Filing Date:** 10 February 2014 (10.02.2014)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:** 61/762,820 8 February 2013 (08.02.2013) US
- (71) **Applicant:** EMOTIENT [US/US]; 4435 Eastgate Mall, Suite 320, San Diego, CA 92121 (US).
- (72) **Inventors:** MOVELLAN, Javier; 8707 Caminito Abrazo, La Jolla, CA 92037 (US). BARTLETT, Marian, Steward; 13056 Candela Place, San Diego, CA 92130 (US). FASEL, Ian; 3066 Ivy Street, San Diego, CA 92104 (US). LITTLEWORT, Gwen, Ford; 145 Cliff Street, Solana Beach, CA 92075 (US). SUSSKIND, Joshua; 8433 Via Mallorca, #97, La Jolla, CA 92037 (US). WHITEHILL, Jacob; 27 Wheel Street, Unit 217, Cambridge, MA 02138 (US).
- (74) **Agent:** CHAMBERS, Daniel, M.; Acuity Law Group, P.C., 12707 High Bluff Drive, Suite #200, San Diego, CA 92130 (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).
- Published:**
— without international search report and to be republished upon receipt of that report (Rule 48.2(g))



(54) **Title:** COLLECTION OF MACHINE LEARNING TRAINING DATA FOR EXPRESSION RECOGNITION

(57) **Abstract:** Apparatus, methods, and articles of manufacture for implementing crowdsourcing pipelines that generate training examples for machine learning expression classifiers. Crowdsourcing providers actively generate images with expressions, according to cues or goals. The cues or goals may be to mimic an expression or appear in a certain way, or to "break" an existing expression recognizer. The images are collected and rated by same or different crowdsourcing providers, and the images that meet a first quality criterion are then vetted by expert(s). The vetted images are then used as positive or negative examples in training machine learning expression classifiers.

COLLECTION OF MACHINE LEARNING TRAINING DATA FOR EXPRESSION RECOGNITION

5

10

15

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims priority from United States provisional patent application Serial Number 61/762,820, entitled SYSTEM FOR COLLECTING MACHINE LEARNING TRAINING DATA FOR FACIAL EXPRESSION RECOGNITION, filed on February 8, 2013, which is hereby incorporated by reference in its entirety as if fully set forth herein, including text, figures, claims, tables, and computer program listing appendices (if present), and all other matter in the United States provisional patent application.

25

FIELD OF THE INVENTION

[0002] This document relates generally to apparatus, methods, and articles of manufacture for collection of training examples for visual machine learning classifiers of expressions of emotions, affective states, action units, and similar psychological states; the use of such training examples in training machine learning classifiers; and machine learning classifiers of expressions of emotions, affective states, action units, and similar psychological states.

30

BACKGROUND

[0003] There are many uses for automated recognition of expressions of emotions, affective states, and similar psychological states. Perhaps the best known use is the smile detection in digital cameras. But there are others, including detection of facial reactions in response to various external stimuli, such as consumer reactions to advertisements, product displays, labeling, packaging, and pricing; and voter facial reactions to talking points and evaluation of debate performance. This list is far from exclusive.

[0004] It is desirable to automate recognition of expressions of emotions, affective states, and similar psychological states. Such automation reduces the costs of recognition, and also provides a measure of objectivity to the result of the recognition process. Automated (or machine) recognition of expressions of emotions, affective states, and similar psychological states is not a trivial design task, because facial expressions, poses, gestures and other face/body movements are not easy to define using standard programming techniques. Machine learning is a collection of techniques that may be used for this purpose. Machine learning allows artificial intelligence systems to learn from examples, in effect performing self-adaptation based on the training data. These techniques typically require large and carefully collected datasets of training examples, for example, a large number of sample images of different people, in different illumination conditions, of various ethnicities and different ages, and with a range of facial artifacts. The examples are needed to allow the machine classifier (recognizer) to discover the features that discriminate between different expression categories. While it may be relatively easy to collect examples of posed smiles from the Web, it is difficult to collect a large number of examples of real or realistic expressions of emotions such as fear, contempt, anger, disgust, and others.

[0005] A need exists for improved methods and apparatus for automatically detecting and classifying psychological states as exhibited by facial expressions, poses, gestures and other face/body movements, whether in reaction to known stimuli or otherwise. A need also exists for efficient methods to generate and collect data that can be used for training automatic classifiers of expressions of emotions, affective states, and similar psychological states.

SUMMARY

[0006] Embodiments described in this document are directed to methods, apparatus, and articles of manufacture that may satisfy one or more of the above described and other needs.

[0007] In an embodiment, a computer-implemented method includes providing cues to a first plurality of untrained providers for mimicking a predetermined expression; receiving from the first plurality of untrained providers images created in response to the step of providing cues; sending requests to rate the images to a second plurality of untrained providers; in response to the step of sending requests, receiving ratings of the second plurality of untrained providers; applying a first quality check to the images rated by the second plurality of untrained providers, the first quality check being based on the ratings of the second plurality of untrained providers, the step of applying the first quality check resulting in a plurality of images that passed the first quality check; sending the plurality of images that passed the first quality check to one or more experts, for rating by the one or more experts; in response to the step of sending the plurality of images, receiving ratings from the one or more experts; applying a second quality check to the images rated by the one or more experts, the second quality check being based on the ratings of the one or more experts, the step of applying the second quality check resulting in one or more images that passed the second quality check; training a classifier using the one or more images that passed the second quality check; and using the classifier after the step of training.

[0008] In an embodiment, a computer-implemented method comprises providing goals to a first plurality of untrained providers for breaking a classifier of a predetermined expression; receiving from the first plurality of untrained providers images created in response to the step of providing goals, resulting in a plurality of received images; checking the received images with the classifier and discarding images that do not meet objective goal of the goals provided, resulting in objectively qualified images; sending requests to rate the objectively qualified images to a second plurality of untrained providers; in response to the step of sending requests, receiving ratings of the second plurality of untrained providers; applying a first quality check to the images rated by the second plurality of untrained providers, the first quality check being based on the ratings of the second plurality of untrained providers, the step of applying the first quality check resulting in a plurality of images that passed the first quality check; sending the plurality of images that passed the first quality check to one or more experts, for rating by the

one or more experts; in response to the step of sending the plurality of images, receiving ratings from the one or more experts; applying a second quality check to the images rated by the one or more experts, the second quality check being based on the ratings of the one or more experts, the step of applying the second quality check resulting in one or more images that passed the second
5 quality check; training a classifier using the one or more images that passed the second quality check; and using the classifier after the step of training.

[0009] These and other features and aspects will be better understood with reference to the following description, drawings, and appended claims.

10 BRIEF DESCRIPTION OF THE FIGURES

[0010] Figure 1 is a simplified block diagram representation of a computer-based system configured in accordance with selected aspects of the present description;

[0011] Figure 2 illustrates selected steps of a process for generating training data for automatic expression classifiers;

15 **[0012]** Figure 3 illustrates selected steps of a process for generating negative examples of training data for automatic expression classifiers;

[0013] Figure 4 illustrates selected steps of a process for generating recognizer “breaking” data for training automatic expression classifiers; and

[0014] Figure 5 illustrates selected steps/block(s) of an iterative process for generating
20 training data for automatic expression classifiers.

DETAILED DESCRIPTION

[0015] In this document, the words “embodiment,” “variant,” “example,” and similar expressions refer to a particular apparatus, process, or article of manufacture, and not necessarily
25 to the same apparatus, process, or article of manufacture. Thus, “one embodiment” (or a similar expression) used in one place or context may refer to a particular apparatus, process, or article of manufacture; the same or a similar expression in a different place or context may refer to a different apparatus, process, or article of manufacture. The expression “alternative embodiment”

and similar expressions and phrases may be used to indicate one of a number of different possible embodiments. The number of possible embodiments/variants/examples is not necessarily limited to two or any other quantity. Characterization of an item as “exemplary” means that the item is used as an example. Such characterization of an embodiment/variant/example does not necessarily mean that the embodiment/variant/example is a preferred one; the embodiment/variant/example may but need not be a currently preferred one. All embodiments/variants/examples are described for illustration purposes and are not necessarily strictly limiting.

[0016] The words “couple,” “connect,” and similar expressions with their inflectional morphemes do not necessarily import an immediate or direct connection, but include within their meaning connections through mediate elements.

[0017] “Affective” information associated with an image or video includes various types of psychological reactions, such as affective, cognitive, physiological, and/or behavioral responses, including both recorded raw signals and their interpretations. Relevant information that represents or describes a particular person’s reaction(s) toward a stimulus in terms of the person’s affective, cognitive, physiological, or behavioral responses is referred to in the present description as affective information. The affective information can be attributable to psychological and physiological reactions such as memories, associations, and the like.

[0018] “Causing to be displayed” and analogous expressions refer to taking one or more actions that result in displaying. A computer or a mobile device (such as a smart phone or tablet), under control of program code, may cause to be displayed a picture and/or text, for example, to the user of the computer. Additionally, a server computer under control of program code may cause a web page or other information to be displayed by making the web page or other information available for access by a client computer or mobile device, over a network, such as the Internet, which web page the client computer or mobile device may then display to a user of the computer or the mobile device.

[0019] “Causing to be rendered” and analogous expressions refer to taking one or more actions that result in displaying and/or creating and emitting sounds. These expressions include within their meaning the expression “causing to be displayed,” as defined above. Additionally, these expressions include within their meaning causing emission of sound.

[0020] References to “images” used as training data refer to still images, videos, and both still images and videos. A “picture” is a still image. “Video” refers to motion graphics.

[0021] “Facial expressions” as used in this document signifies the facial expressions of primary emotion (such as Anger, Contempt, Disgust, Fear, Happiness, Sadness, Surprise, Neutral); as well as expressions of affective state of interest (such as boredom, interest, engagement); and so-called “action units” (movements of a subset of facial muscles, including movement of individual muscles). The term “facial expressions” includes head poses and gestures.

[0022] “Mental state” as used in this document means emotion, affective state, or similar psychological state; “expression of emotion, affective state, and similar psychological state” means expression of emotion, affective state, or similar psychological state.

[0023] As used in this document, a classifier “breaks” or is “broken” when it makes a mistake in classifying a data sample. Thus, the classifier breaks when it detects a target emotion or affective state in an image that to people does not have the appearance of an expression of the target emotion or affective state. The classifier also breaks when it fails to detect the target emotion or affective state in image that has the appearance of an expression of the target emotion or affective state.

[0024] Other and further explicit and implicit definitions and clarifications of definitions may be found throughout this document.

[0025] Reference will be made in detail to several embodiments that are illustrated in the accompanying drawings. Same reference numerals are used in the drawings and the description to refer to the same apparatus elements and method steps. The drawings are in a simplified form, not to scale, and omit apparatus elements and method steps that can be added to the described systems and methods, while possibly including certain optional elements and steps.

[0026] Advantageously, mimicry can be employed to generate realistic training examples of expressions of mental states, for use as training data, and/or for other purposes. In an example, data is collected through the Web. Users have camera-equipped communication devices, and can send images of their faces to a server. This may be done using a crowdsourcing service, like the Amazon Mechanical Turk.

[0027] The Amazon Mechanical Turk is a crowdsourcing Internet-based resource that enables “requesters” to manage the use of human intelligence for various tasks, such as tasks that computers are not well equipped to perform at this time. The requesters are able to post tasks that “providers” or “Turkers” can perform for compensation set by the requesters. Note that the use of the terms “Turker,” “requester,” or “provider” and related terms does not necessarily signify Amazon’s resource, but applies to any crowdsourcing resource.

[0028] Figure 1 is a simplified block diagram representation of a computer-based system 100, configured in accordance with selected aspects of the present description to interact with crowdsourcing providers at provider machines 180 (such as personal computers or PCs, smartphones, tablets, network-enabled cameras) through a communication network 190. In particular, the system 100 and the provider machines 180 may be configured for use in obtaining training data through the use of mimicry, as is described in this document.

[0029] Figure 1 does not show many hardware and software modules of the system 100 or of the provider machines 180, and omits several physical and logical connections. The system 100 may be implemented as a special purpose data processor, a general-purpose computer, a computer system, or a group of networked computers or computer systems configured to perform the steps of the methods described in this document. In some embodiments, the system 100 is built on a personal computer platform, such as a Wintel PC, a Linux computer, or a Mac computer. The personal computer may be a desktop or a notebook computer. The system 100 may function as a server computer. In some embodiments, the system 100 is implemented as a plurality of computers interconnected by a network, such as the network 190, or another network.

[0030] As shown in Figure 1, the system 100 includes a processor 110, read only memory (ROM) module 120, random access memory (RAM) module 130, network interface 140, a mass storage device 150, and a database 160. These components are coupled together by a bus 115. In the illustrated embodiment, the processor 110 may be a microprocessor, and the mass storage device 150 may be a magnetic disk drive. The mass storage device 150 and each of the memory modules 120 and 130 are connected to the processor 110 to allow the processor 110 to write data into and read data from these storage and memory devices. The network interface 140 couples the processor 110 to the network 190, for example, the Internet. The nature of the

network 190 and of the devices that may be interposed between the system 100 and the network 190 determine the kind of network interface 140 used in the system 100. In some embodiments, for example, the network interface 140 is an Ethernet interface that connects the system 100 to a local area network, which, in turn, connects to the Internet. The network 190 may, therefore, in fact include a collection of networks.

[0031] The database 160 may be used for organizing and storing data that may be needed or desired in performing the method steps described in this document. The database 160 may be a physically separate system coupled to the processor 110. In alternative embodiments, the processor 110 and the mass storage device 150 may be configured to perform the functions of the database 160.

[0032] The processor 110 may read and execute program code instructions stored in the ROM module 120, the RAM module 130, and/or the storage device 150. Under control of the program code, the processor 110 may configure the system 100 to perform the steps of the methods described or mentioned in this document. In addition to the ROM/RAM modules 120/130 and the storage device 150, the program code instructions may be stored in other machine-readable tangible storage media, such as additional hard drives, floppy diskettes, CD-ROMs, DVDs, Flash memories, and similar devices. The program code can also be transmitted over a transmission medium, for example, over electrical wiring or cabling, through optical fiber, wirelessly, or by any other form of physical transmission. The transmission can take place over a dedicated link between telecommunication devices, or through a wide area or a local area network, such as the Internet, an intranet, extranet, or any other kind of public or private network. The program code may also be downloaded into the system 100 through the network interface 140 or another network interface.

[0033] In embodiments, the system 100 may function as a crowdsourcing server that enables the providers at the machines 180 to perform the human intelligence tasks (“HITS”) specified by the system 100. The descriptions of the tasks may be predetermined and stored in one of the storage devices of the system 100, or received by the system 100 from another device or person. Those skilled in the art should be familiar with the Amazon Web Services implementing the Amazon Mechanical Turk that performs crowdsourcing functions.

[0034] The system 100 may store the information received from the providers in a database, which database may be implemented with the components of the system 100, and/or other components. The database storing the information from the providers may be the database 160.

5 [0035] Figure 2 illustrates selected steps of a process 200 for presenting crowdsourcing providers examples of target facial expressions, asking the providers to mimic those expressions, and selecting/using training examples from the images. In embodiments, some or all of the examples may be chosen based on relationships between emotions and their expressions detectable through facial movements described in terms of elements of a system for describing
10 facial movements, such as action units of the Facial Action Coding System (FACS), or other research.

[0036] At flow point 201, the system 100 and at least some provider machines are configured to communicate with each other and to perform the other steps of the process. Note that various providers may communicate with the system 100 at different times, and accordingly
15 their machines may be configured at different times.

[0037] In step 205, one or more providers of a first set of providers are given cues about which features of the facial expression they should mimic/reproduce; the cues may include verbal (sound, written) instructions of what to do, pictures/videos to mimic or reproduce, references to online or other materials describing the expressions, other cues, and combinations
20 of any of these. The providers of the first set need not be specially trained in facial expression recognition. For example, the system 100 and the respective machine of the provider cause the example expression (or portion thereof) to be displayed to the provider, and/or causes to be displayed or rendered to the provider instructions for mimicking the expression or the portion of the expression (“look like this,” “look afraid,” “look bored,” “look happy,” “look excited,” “look
25 scared,” “look engaged,” “make this kind of wrinkle,” “raise your eyebrows as if you are surprised,” “look at and mimic that part of the displayed face”). The example and/or the instructions may be stored in the database 160 or elsewhere, and transmitted to the provider through the network 190, by mail, or otherwise.

[0038] In step 210, the providers of the first set create the expressions and capture them,
30 for example, making facial expressions and using still or video cameras in their respective

machines 180. Again, here and elsewhere, the providers may perform the steps at different times.

[0039] In step 215, the providers of the first set submit the images with the expressions to the system 100, for example, causing the images to be transmitted over the network 190.

5 [0040] In step 220, the system 100 receives the images with the expressions, for example, through the network 190.

[0041] In step 225, one or more providers of a second set of providers are asked to rate the images with the expressions from the first set of providers, with respect to a set of predetermined criteria or criterion. The providers of the second set also need not be specially
10 trained in facial expression recognition. For example, the system 100 and/or the machines 180 of the respective providers of the second set cause to be displayed to the providers of the second set the criteria and/or the questions. For example, this may be done by displaying an image and asking “does this look like a smile?” The displayed information (criteria/questions) may ask the providers of the second set, how well does the displayed expression matches a specific of
15 emotion, such as anger, contempt, disgust, fear, happiness, sadness, surprise, or another expression, affective state, or action unit; the specific emotion, affective state, or action unit may be the same one that the providers of the first state were asked to mimic. The displayed question may also or instead ask, how well does the displayed expression (previously obtained in the steps 210-220) mimic the sample expression (which was given or defined to the first set of providers
20 in the steps 205). In embodiments, the providers of the second set may be asked to pass or fail each expression. In embodiments, the providers of the second set may be asked to grade the expressions; the grades may be as follows: 1. Looks nothing like it; 2. A poor attempt; 3. OK; 4. A pretty good attempt; 5. Incredibly good mimicking.

[0042] In embodiments, the providers of the second set are all different from the
25 providers of the first set. In embodiments, there may be at least some or even a complete overlap between the providers of the second set and the providers of the first set. In variants that are not necessarily limiting, no provider is asked to rate his or her own submission(s) of the images with the expressions.

[0043] In step 230, which may be performed in response to the step 225, the providers of
30 the second set rate how well each of the rated expressions matches the mimicked expression, or

how well each of the expressions is under the predetermined criteria or criterion. The machines 180 may receive the ratings given by the providers of the second set and transmit them to the system 100, for example, over the network 190. The system 100 receives the ratings, for example, over the network 190. The ratings may be simple yes/no, or more complicated, such as ratings on a scale from 1 to 10, or a letter-grade scale.

[0044] In step 232, the images received from the first set of providers are tested using a first predetermined quality criterion or criteria, and, in step 235, the images that pass the first predetermined quality criterion (or criteria) are sent (*e.g.*, electronically from the system 100, over the network 190) to be vetted by a pool of facial expression experts. The first quality criterion may be, for example, at least N of untrained providers agree that the expression matches that of an emotion X . The number N may be an absolute predetermined threshold number (five, for example), a scaled predetermined threshold number (N out of M , such as 6 out of 10), a predetermined threshold percentage (65%, for example). The percentages and numbers given here are exemplary and not necessarily limiting.

[0045] In step 240, the system 100 receives the decisions of the experts for each of the vetted images, and stores the decisions. For example, the system 100 may receive the decisions of the experts over the network 100.

[0046] In step 245, a second predetermined quality criterion (or criteria) is (or are, as the case may be) used to accept an image as a positive example. The second quality criterion or criteria may be set in a way that is analogous to the way the first predetermined quality criterion was set, for example, at least P experts agree that the expression matches that of an emotion X . The number P may be an absolute number, a scaled number, or a percentage.

[0047] In embodiments, only a single expert is used. The single expert may be asked, for example, to rate the images on a pass/fail basis, or to select and assign a grade, from among two, three or more available grades, to the image with respect to each particular type of expression. In other embodiments, two, three, or more experts may be used.

[0048] In step 250, the system selects some or all of the images that have been vetted by the experts as training data for an expression recognition machine (classifier). For example, all or a predetermined portion of the available training images or a predetermined number of the images may be selected.

[0049] In step 255, the expression recognition machine is trained using the training data. Some of the machine learning techniques that may be employed in this step are described below and elsewhere in this document. Other machine learning and optimization techniques may also be employed.

5 [0050] In step 260, the expression recognition machine trained in the step 255 is used. For example, the machine may be implemented in a camera to allow the camera's user to capture automatically images with the specific expressions. As another example, the expression recognition machine judges people's reactions resulting from exposure to certain stimulus or stimuli, such as advertisements. As still another example, the expression recognition machine
10 may be implemented on a user's computer or over the network (through a browser, perhaps) to allow the user to self-train in creating various expressions; the user may be enabled to capture the user's image (picture/video) and rate the image for presence of a given expression (concern, attention, joy, etc.) using the classifier.

[0051] At flow point 299, the process ends; it may be repeated in whole or in part as
15 needed.

[0052] The process 200 yields what is known a "positive examples." In other words, the examples are of a particular expression which is sought to be matched. Importantly, positive examples from one expression (x) may be considered as negative examples of another expression (y), if those two expressions are known to be incompatible (*e.g.*, "happy" and "sad"). Moreover,
20 the images that rate poorly as expressions of a particular emotion, affective state, or action unit may be used in machine learning as negative examples of that same emotion, affective state, or action unit. Thus, the training in the step 255 may be for positive recognition of the expression of the positive example (x), or for training for non-recognition of the same expression as that of the other (negative) emotion (y).

25 [0053] The steps 225 through 250 of the process 200 may be considered to be a "pipeline" for reducing the cost of generating both positive and negative examples for training expression recognition machines. Thus, all kinds of facial images may be collected by various means (from a web search, for example), and passed through the pipeline to generate either positive or negative examples. For negative examples, the providers of the second set may be

asked whether each image is clearly not the expression of the target emotion. A process 300 of Figure 3 illustrates this approach to the generation of negative examples.

[0054] The step 325 is analogous to the step 225, but geared to the generation of negative rather than positive examples. Thus, one or more providers (of a second set or of a third set, which may be similar to the second set) are asked to rate the images with the expressions from the first set of providers, with respect to a set of predetermined criteria or criterion. These providers also need not be specially trained in facial expression recognition. For example, the system 100 and/or the machines 180 of the respective providers cause to be displayed to the providers the criteria and/or the questions. For example, displaying an image and asking “does this look like a smile?” The displayed information (criteria/questions) may ask the providers of the second set, how well does the displayed expression matches a specific emotion, such as anger, contempt, disgust, fear, happiness, sadness, surprise, or another expression or affective state; the specific emotion, affective state, or action unit may be the same one that the providers of the first state were asked to mimic. The emotion, affective state, or action unit may be different from the emotion, affective state, or action unit mimicked by the providers who generated the images. In embodiments, the providers of the second set may be asked to grade the expressions pass/fail. In embodiments, the providers of the second set may be asked to grade the expressions; the grades may be as follows: 1. Looks nothing like it; 2. A poor attempt; 3. OK; 4. A pretty good attempt; 5. Incredibly good mimicking.

[0055] In embodiments, the providers of the step 325 may all be different from the providers of the first set, or there may be at least some or even a complete overlap between the providers of the second set and the providers of the first set. In variants that are not necessarily limiting, no provider is asked to rate his or her own submission(s) of the images with the expressions.

[0056] In step 330, which is identical or analogous to the step 230, the system 100 receives the responses of the providers, which response were generated in response to the step 325.

[0057] In step 332, which is identical or analogous to the step 232, the images received from the first set of providers are tested using a first predetermined quality criterion or criteria. In step 335, which is analogous to the step 235, the images that pass the first predetermined

quality criterion (or criteria) are sent (*e.g.*, electronically from the system 100, over the network 190) to be vetted by a pool of facial expression experts. The first quality criterion may be, for example, at least N of untrained providers agree that the expression is a bad expression of emotion X . The number N may be an absolute predetermined threshold number (five, for example), a scaled predetermined threshold number (N out of M , such as 6 out of 10), a predetermined threshold percentage (65%, for example). The percentages and numbers given here are exemplary and not necessarily limiting.

[0058] In step 340, the decision(s)/rating(s) of the expert(s) is/are received, for example, by the system 100 through the network 190.

[0059] In step 345, which is identical or analogous to the step 245, a second predetermined quality criterion (or criteria) is (are) used to accept an image as a negative example. The second quality criterion/criteria may be set in a way that is analogous to the way the first predetermined quality criterion was set, for example, at least P experts agree that the expression is a bad expression of the emotion X . The number P may be an absolute number, a scaled number, or a percentage. In embodiments, only a single expert is used. The single expert may be asked, for example, to rate the images on a pass/fail basis, or to select and assign a grade, from among two, three or more available grades, to the image with respect to each particular type of expression. In other embodiments, two, three, or more experts may be used.

[0060] In step 350, which is identical or analogous to the step 250, the system selects some or all of the images that have been vetted by the experts as training data for an expression recognition machine. Here, however, the training data are negative examples. For example, all or a predetermined portion of the available training images or a predetermined number of the images may be selected.

[0061] A goal of the pipelines (both positive and negative) may be to minimize the cost and time per image that is accepted as an example. Setting the parameters for selection (the first and second predetermined criteria) may thus be based on a cost function. The cost function may be based on a weighted combination of time and cost, or simply the actual cost that accounts for both the cost of the untrained providers and of the experts. Mathematical models may be used to minimize the cost function as a function of the number of untrained subjects and of the strength

of the response of the experts. These parameters may be further tuned using reinforcement learning approaches applied to the machine learning.

[0062] Thus, the training data collected using the approaches described above may be used to develop a first generation expression recognition system using machine learning methods (e.g., pattern recognizers, neural networks, support vector machines, adaboost, and possibly other types of classifiers). The techniques can be further extended to collecting the data that “breaks” some version (e.g., the current version) of the expression recognizer. Thus, the system may be configured to ask the crowdsourcing providers to send example images for which the system makes mistakes. This can also be done by embedding applications with the ability to send data to a central server. The new break-me examples may then be vetted using an approach (pipeline) analogous to that described above for collecting positive and/or negative examples. The new data are then used for training, to develop a second generation expression recognition system. The process can be iterated to develop a third generation system, etc.

[0063] In an embodiment, the system asks the crowdsourcing providers to create examples of expression of a certain emotion (e.g., anger) that would be generally recognized as that expression by people but that would cause the expression recognizer (i.e., detector or classifier) not to identify it as that particular emotion such as anger (and possibly to identify it as another type of emotion, such as happiness); or *vice versa*, the system asks the crowdsourcing providers to create examples of expression that do not appear to people as expressions corresponding to a certain emotion (e.g., anger), but that would cause the expression recognizer (i.e., detector or classifier) to identify it as that particular emotion. The providers would then send the images that they identify as such to the system; the system would receive and store them with appropriate descriptions. The images would then be sent through the pipeline such as that described in relation to Figure 2 (steps 225-260) and Figure 3. Figure 4 illustrates a method for obtaining such data that “breaks” a detector.

[0064] In step 405, one or more providers of a first set of providers are given goals (similar to the cues/instructions described above) for creating images wherein the facial expression appears as a first type of expression but which the classifier does not classify as the first type; alternatively, a goal may be to create a facial expression that does not look like the first type of expression, but that the classifier classifies as the first type of expression; or *vice*

versa. For example, the providers of the first set may be asked to create expressions that look like anger, but are not classified as anger by the current generation classifier; or the providers may be asked to create expressions that do not look like anger, but are classified as anger by the classifier. The providers may be supplied with a means to classify the images they create with the classifier. For example, the classifier may be made available online, or installed on the providers' machines 180. As before, the providers of the first set need not be specially trained in facial expression recognition. The goals and/or the cues/instructions may be stored in the database 160 or elsewhere.

[0065] In step 410, the providers of the first set create the expressions, capture them, and test them against the classifier. For example, the providers capture their images using the cameras in their respective machines 180, and test them also using the machines 180 to access the classifier.

[0066] If a particular provider is satisfied that he or she has achieved the goal(s) with a particular image, the provider submits the image to the system 100, for example, by selecting a menu item that causes the image to be transmitted over the network 190. This is done in step 415.

[0067] In step 420, the system 100 receives the images submitted by the providers.

[0068] In step 422, the system tests the received images using the same classifier as the one that was used by the providers of the first set, and/or one or more other classifiers. If a particular image does not meet the objective goal, the image is discarded. If, for example, the goal was to obtain images that are classified as anger (but does not look like anger), and when the system submits the image to the classifier the output of the classifier is not anger, then the image may be discarded (not used further in the method). Similarly, if the goal was to obtain images that are not classified as anger (but appear as anger to people), then the image may be discarded if the classifier output corresponds to anger.

[0069] In step 425, one or more providers of a second set of providers are asked to rate the images with respect to the subjective goal used by the first set of providers. (The providers of the second set also need not be specially trained in facial expression recognition.) Continuing with the example of the immediately preceding paragraph, the providers of the second set may be asked whether an image is of an angry face.

[0070] In embodiments, the providers of the second set are all different from the providers of the first set. In embodiments, there may be at least some or even a complete overlap between the providers of the second set and the providers of the first set. In embodiments, there may be at least some or even a complete overlap between the providers of the second set and the providers of the first set. In variants that are not necessarily limiting, no provider is asked to rate his or her own submission(s) of the images with the expressions.

[0071] In step 430, the providers of the second set rate the images based on the subjective criteria, the machines 180 receive the ratings given by the providers of the second set and transmit them to the system 100, for example, over the network 190. The system 100 receives the ratings. The ratings may be simple yes/no (looks like anger or does not look like anger), or more complicated, such as a rating on a scale from 1 to 10.

[0072] In step 432, the images are tested using a first predetermined quality criterion or criteria, and, in step 435, the images that pass the first predetermined quality criterion (or criteria) test are sent (*e.g.*, electronically from the system) to be vetted by a pool of facial expression experts. The first quality criterion may be, for example, at least N of untrained providers agree that the expression matches that of an emotion X . The number N may be an absolute predetermined threshold number (five, for example), a scaled predetermined threshold number (N out of M , such as 6 out of 10), a predetermined threshold percentage (65%, for example). The percentages and numbers given here are exemplary and not necessarily limiting.

[0073] In step 440, the system 100 receives the decisions of the experts for each of vetted images, and stores the decisions.

[0074] In step 445, the system applies a second predetermined quality criterion (or criteria) to accept an image as a “break” the classifier example. The second quality criterion or criteria may be set in a way that is analogous to the way the first predetermined quality criterion was set, for example, at least P experts agree that the expression matches (or does not match) that of an emotion X . The number P may be an absolute number, a scaled number, or a percentage.

[0075] In embodiments, only a single expert is used. The single expert may be asked, for example, to rate the images on a pass/fail basis, or to assign a grade of three or more available

grades to the image with respect to each particular type of expression. In other embodiments, two, three, or more experts are used.

[0076] In step 450, the system selects some or all of the images that have been vetted by the experts as training data for the classifier that was “broken” by the data.

5 [0077] In step 455, the classifier is trained by the “breaking” data, and becomes the next generation classifier. The above steps may be iterated one or more times, to refine the classifier further.

[0078] In step 460, the classifier is used in a manner such as that described above in relation to the step 260.

10 [0079] At flow point 499, the process ends, to be repeated in whole or in part as needed.

[0080] Figure 5 shows selected steps/block(s) of an iterative method 500, which is based on the method 200 described above in connection with Figure 2. Here, however, a decision is made in decision block 556 whether to proceed with finding discrepancies between the labels assigned by the human labelers/provider (which can be the providers of the second set) and the
15 corresponding outputs of the trained classifier. The decision may be based on whether a predetermined number of iterations (*e.g.*, two, three, or more) has been reached, or on another criteria. If the decision is not to proceed, the process flow advances to step 560 (similar or identical to the step 260 of the method 200), to use the trained expression recognition machine.

[0081] Otherwise, the process flow continues to a step 557, where images with
20 discrepancies between human ratings and outputs of expression recognition machine are selected; for example, all such images may be selected.

[0082] From the step 557, the process flow continues to step 558, to request human labelers (such as the providers of the second set) for an additional opinion to confirm or disconfirm the discrepancy. When such discrepancy is found, the process continues to ask the
25 experts for a second opinion. This second opinion is included in the second (or subsequent) round of training of the expression recognition machine. The process may be iterated a few times (*e.g.*, asking for a third, fourth opinion, etc.) Basically, this is an iterative mechanism to refine the human-assigned labels. According to this process, second, third, and subsequent opinions are not requested for all the labels, but just for additional opinions when there is a

discrepancy between the human-assigned labels and the system outputs. For example, if the experts indicate that image X shows fear but the expression recognizer indicates that it does not show fear, the image X is sent back to the experts for an additional opinion.

[0083] Figure 5 shows a modification of the process 200. Analogous changes may be made to the processes 300 and 400, to effect iterations in order to refine these processes.

[0084] The system and process features described throughout this document may be present individually, or in any combination or permutation, except where presence or absence of specific feature(s)/element(s)/limitation(s) is inherently required, explicitly indicated, or otherwise made clear from the context.

[0085] Although the process steps and decisions (if decision blocks are present) may be described serially in this document, certain steps and/or decisions may be performed by separate elements in conjunction or in parallel, asynchronously or synchronously, in a pipelined manner, or otherwise. There is no particular requirement that the steps and decisions be performed in the same order in which this description lists them or the Figures show them, except where a specific order is inherently required, explicitly indicated, or is otherwise made clear from the context. Furthermore, not every illustrated step and decision block may be required in every embodiment in accordance with the concepts described in this document, while some steps and decision blocks that have not been specifically illustrated may be desirable or necessary in some embodiments in accordance with the concepts. It should be noted, however, that specific embodiments/variants/examples use the particular order(s) in which the steps and decisions (if applicable) are shown and/or described.

[0086] The instructions (machine executable code) corresponding to the method steps of the embodiments, variants, and examples disclosed in this document may be embodied directly in hardware, in software, in firmware, or in combinations thereof. A software module may be stored in volatile memory, flash memory, Read Only Memory (ROM), Electrically Programmable ROM (EPROM), Electrically Erasable Programmable ROM (EEPROM), hard disk, a CD-ROM, a DVD-ROM, or other forms of non-transitory storage medium known in the art, whether volatile or non-volatile. Exemplary storage medium or media may be coupled to one or more processors so that the one or more processors can read information from, and write

information to, the storage medium or media. In an alternative, the storage medium or media may be integral to one or more processors.

[0087] This document describes in considerable detail the inventive apparatus, methods, and articles of manufacture for obtaining and using training data for machine learning systems. This was done for illustration purposes. The specific embodiments or their features do not necessarily limit the general principles underlying the invention. The specific features described herein may be used in some embodiments, but not in others, without departure from the spirit and scope of the invention as set forth herein. Various physical arrangements of components and various step sequences also fall within the intended scope of the invention. Many additional modifications are intended in the foregoing disclosure, and it will be appreciated by those of ordinary skill in the pertinent art that in some instances some features will be employed in the absence of a corresponding use of other features. The illustrative examples therefore do not necessarily define the metes and bounds of the invention and the legal protection afforded the invention, which function is carried out by the claims and their equivalents.

15 **[0088]** What is claimed is:

CLAIMS

1. A computer-implemented method comprising steps of:

5 providing one or more cues to a first plurality of untrained providers for mimicking a predetermined expression;

receiving from the first plurality of untrained providers images created in response to the step of providing one or more cues;

sending requests to rate the images to a second plurality of untrained providers;

10 receiving ratings of the images from the second plurality of untrained providers, in response to the step of sending requests;

applying a first quality check to the images rated by the second plurality of untrained providers, the first quality check being based on the ratings of the second plurality of untrained providers, the step of applying the first quality check resulting in one or more images that passed the first quality check;

15 sending the one or more images that passed the first quality check to one or more experts, for rating by the one or more experts;

in response to the step of sending the one or more images that passed the first quality check, receiving one or more ratings from the one or more experts; and

20 applying a second quality check to the images rated by the one or more experts, the second quality check being based on the one or more ratings of the one or more experts, the step of applying the second quality check resulting in one or more images that passed the second quality check.

2. A computer-implemented method as in claim 1, further comprising:

training a machine-learning classifier of an emotion, affective state, or action unit associated with the predetermined expression using the one or more images that passed the second quality
5 check; and

using the classifier, after the step of training, to classify new images that have not been used in training of the classifier;

wherein the images that fail the first quality check are not sent to the one or more experts.

10 3. A computer-implemented method as in claim 2, wherein the predetermined expression is an expression of a predetermined emotion, affective state, or action unit, the classifier is a classifier of the predetermined emotion, affective state, or action unit, the step of providing one or more cues comprises providing one or more cues for mimicking the predetermined emotion, affective state, or action unit, and the step of training the classifier comprises training the classifier of the
15 predetermined emotion, affective state, or action unit.

4. A computer-implemented method as in claim 3, wherein the steps of providing one or more cues, receiving from the first plurality of untrained providers images created in response to the step of providing one or more cues, sending requests to rate the images to a second plurality of
20 untrained providers, receiving ratings of the second plurality of untrained providers, and receiving ratings of the second plurality of untrained providers are performed by a computer-

based system over a network coupling the computer-based system to provider machines of the providers of the first and second pluralities of untrained providers.

5 5. A computer-implemented method as in claim 4, wherein the step of providing one or more cues comprises at least one of: sending instructions over the network from the computer-based system to the provider machines of the providers of the first plurality of untrained providers, and sending one or more sample images over the network from the computer-based system to the provider machines of the providers of the first plurality of untrained providers.

10 6. A computer-implemented method as in claim 4, wherein the step of using the classifier comprises programming the classifier into a camera to enable the camera to capture pictures with the predetermined expression.

15 7. A computer-implemented method as in claim 4, wherein the step of using the classifier comprises programming the classifier into a device configured to judge reactions of people resulting from exposure to one or more stimuli.

20 8. A computer-implemented method as in claim 4, wherein the step of using the classifier comprises using the classifier to allow a user to self-train in creating the predetermined expression, wherein the classifier is employed to rate a picture of the user for appearance of the predetermined expression.

9. A computer-implemented method as in claim 2, further comprising:

selecting from the one or more images that passed the second quality check discrepant images that cause a discrepancy between labels assigned by the second plurality of untrained workers and labels assigned by the machine-learning classifier;

5 sending the discrepant images to the second plurality of untrained providers;

receiving ratings of the discrepant images from the second plurality of untrained providers;

applying the first quality check to the discrepant images rated by the second plurality of untrained providers, the step of applying the first quality check to the discrepant images resulting in one or more discrepant images that passed the first quality check;

10 sending the one or more discrepant images that passed the first quality check to the one or more experts, for rating by the one or more experts;

in response to the step of sending the one or more discrepant images that passed the first quality check, receiving one or more ratings of discrepant from the one or more experts; and

15 applying the second quality check to the one or more discrepant images rated by the one or more experts;

training the classifier with the one or more discrepant images that passed the second quality check; and

using the classifier, after the step of training the classifier with the one or more discrepant images, to classify additional new images that have not been used in training of the classifier.

20

10. A computer-implemented method comprising steps of:

providing one or more goals to a first plurality of untrained providers for breaking a classifier, the classifier being a machine-learning classifier of a predetermined expression of a predetermined emotion, affective state, or action unit associated with the predetermined expression;

5 receiving from the first plurality of untrained providers images created in response to the step of providing one or more goals, resulting in a plurality of received images;

checking the received images with the classifier and discarding images that do not meet a predetermined standard applied to output of the classifier, resulting in objectively qualified images that pass the predetermined standard;

10 sending to a second plurality of untrained providers requests to rate the objectively qualified images with respect to appearance of the predetermined emotion, affective state, or action unit in the objectively qualified images;

receiving ratings of the second plurality of untrained providers, in response to the requests;

15 applying a first quality check to the received images rated by the second plurality of untrained providers, the first quality check being based on the ratings of the second plurality of untrained providers, the step of applying the first quality check resulting in a plurality of images that passed the first quality check;

20 sending the plurality of images that passed the first quality check to one or more experts, for rating the images of the plurality of images that passed the first quality check by the one or more experts with respect to appearance of the predetermined emotion, affective state, or action unit in the images of the plurality of images that passed the first quality check;

receiving ratings from the one or more experts, in response to the step of sending the plurality of images that passed the first quality check; and

applying a second quality check to the images rated by the one or more experts, the second quality check being based on the ratings of the one or more experts, the step of applying the
5 second quality check resulting in one or more images that passed the second quality check.

11. A computer-implemented method as in claim 10, wherein:

the one or more goals comprise creation of positive example images that cause the classifier to indicate non-appearance of the predetermined emotion, affective state, or action unit, and the
10 example images appear as expression of the predetermined emotion, affective state, or action unit; and

the objectively qualified images that do not pass the first quality check are not sent to the one or more experts;

the method further comprising:

15 training the classifier using the one or more images that passed the second quality check as positive examples of appearance of the predetermined emotion, affective state, or action unit; and

using the classifier after the step of training.

20 12. A computer-implemented method as in claim 11, wherein the steps of providing one or more goals, receiving from the first plurality of untrained providers images created in response to the

step of providing one or more goals, sending to a second plurality of untrained providers requests to rate, and receiving ratings of the second plurality of untrained providers are performed by a computer-based system over a network coupling the computer-based system to provider machines of the providers of the first and second pluralities of untrained providers.

5

13. A computer-implemented method as in claim 12, wherein the step of providing one or more goals comprises sending instructions over the network from the computer-based system to the provider machines of the providers of the first plurality of untrained providers.

10 14. A computer-implemented method as in claim 12, wherein the step of providing one or more goals comprises sending one or more images over the network from the computer-based system to the provider machines of the providers of the first plurality of untrained providers.

15 15. A computer-implemented method as in claim 12, wherein the step of using the classifier comprises programming the classifier into a camera to enable the camera to capture pictures with the appearance of the predetermined emotion, affective state, or action unit.

20 16. A computer-implemented method as in claim 12, wherein the step of using the classifier comprises programming the classifier into a device configured to judge reactions of people resulting from exposure to one or more stimuli.

17. A computer-implemented method as in claim 12, wherein the step of using the classifier comprises using the classifier to allow a user to self-train in creating an appearance of the predetermined emotion, affective state, or action unit, wherein the classifier is employed to rate a pictured of the user for appearance of the predetermined emotion, affective state, or action unit.

5

18. A computer-implemented method as in claim 10, wherein:

the one or more goals comprise creation of negative example images that cause the classifier to indicate appearance of the predetermined emotion, affective state, or action unit, and the example images do not appear as expression of the predetermined emotion, affective state, or
10 action unit; and

the objectively qualified images that do not pass the first quality check are not sent to the one or more experts;

the method further comprising:

training the classifier using the one or more images that passed the second quality check as
15 negative examples of appearance of the predetermined emotion, affective state, or action unit;
and

using the classifier after the step of training.

19. A computer system comprising:

20 at least one processor;

at least memory device storing one or more cues for mimicking expression of a predetermined emotion, affective state, or action unit, the at least one memory device being coupled to the at least one processor; and

at least one network interface coupled to the at least one processor and to at least one
5 network;

wherein the at least one processor is configured to send through the at least one network interface the one or more cues to a first plurality of untrained providers; to receive from the first plurality of untrained providers through the at least one network interface images created in response to the one or more cues; to send to a second plurality of untrained providers through the
10 at least one network interface requests to rate the images; to receive through the at least one network interface ratings of the second plurality of untrained providers in response to the requests; to apply a first quality check to the images rated by the second plurality of untrained providers, the first quality check being based on the ratings of the second plurality of untrained providers, thereby obtaining one or more images that passed the first quality check; to send the
15 one or more images that passed the first quality check to one or more experts, for rating by the one or more experts; to receive one or more ratings from the one or more experts generated in response to the one or more images that passed the first quality check being sent to the one or more experts; to apply a second quality check to the images rated by the one or more experts, the second quality check being based on the one or more ratings of the one or more experts, thereby
20 obtaining one or more images that passed the second quality check; and to train a classifier using the one or more images that passed the second quality check; and

wherein the at least one processor is configured so that images that fail the first quality check are not sent to the one or more experts.

20. A computer system comprising:

at least one processor;

at least one memory device storing one or more goals for breaking a classifier of expression
5 of a predetermined emotion, affective state, or action unit, the at least one memory device being
coupled to the at least one processor; and

at least one network interface coupled to the at least one processor and to at least one
network;

wherein the at least one processor is configured to send through the at least one network
10 interface the one or more goals to a first plurality of untrained providers; to receive from the first
plurality of untrained providers through the at least one network interface images created in
response to the goals; to check the received images with the classifier and to obtain objectively
qualified images that meet a predetermined standard; to send to a second plurality of untrained
providers through the at least one network interface requests to rate the objectively qualified
15 images with respect to appearance of expression of the predetermined emotion, affective state, or
action unit in the objectively qualified images; to receive through the at least one network
interface ratings of the second plurality of untrained providers in response to the requests to rate
the objectively qualified images; to apply a first quality check to the objectively qualified images
rated by the second plurality of untrained providers, the first quality check being based on the
20 ratings of the second plurality of untrained providers, thereby obtaining one or more images that
passed the first quality check; to send the one or more images that passed the first quality check
to one or more experts, for rating by the one or more experts with respect to appearance of

expression of the predetermined emotion, affective state, or action unit in the one or more images that passed the first quality check; to receive one or more ratings from the one or more experts generated in response to the one or more images that passed the first quality check being sent to the one or more experts; to apply a second quality check to the images rated by the one or more
5 experts, the second quality check being based on the one or more ratings of the one or more experts, thereby obtaining one or more images that passed the second quality check; and to train a classifier using the one or more images that passed the second quality check as examples of appearance or non-appearance of expression of the predetermined emotion, affective state, or action unit; and
10 wherein the at least one processor is configured so that the objectively qualified images that fail the first quality check are not sent to the one or more experts.

1/8

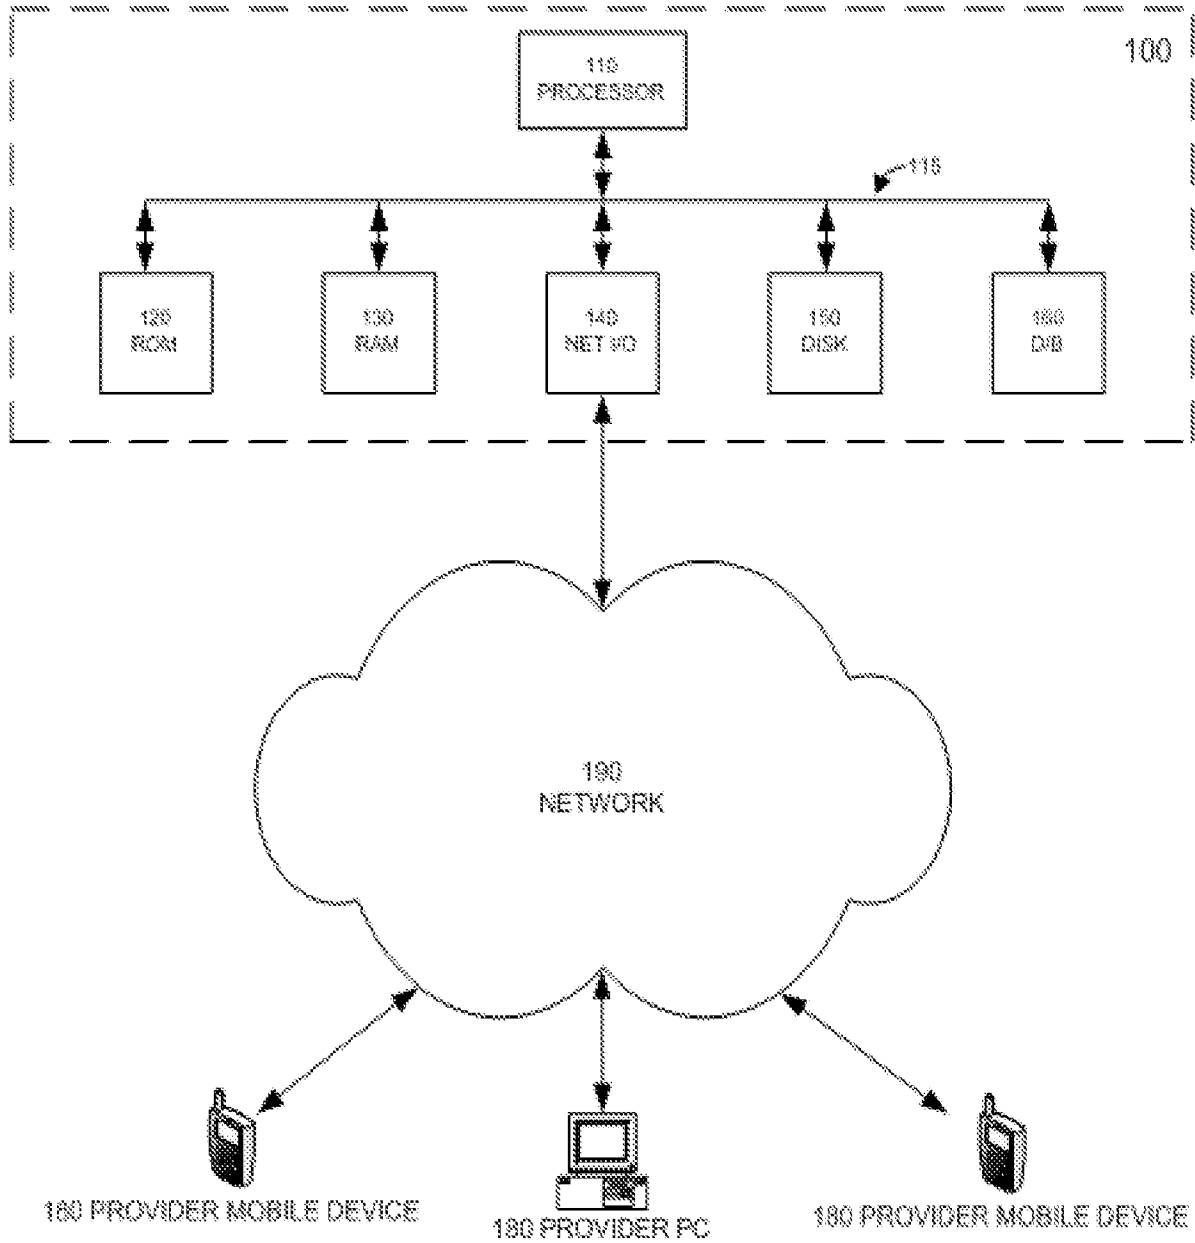


FIG. 1

2/8

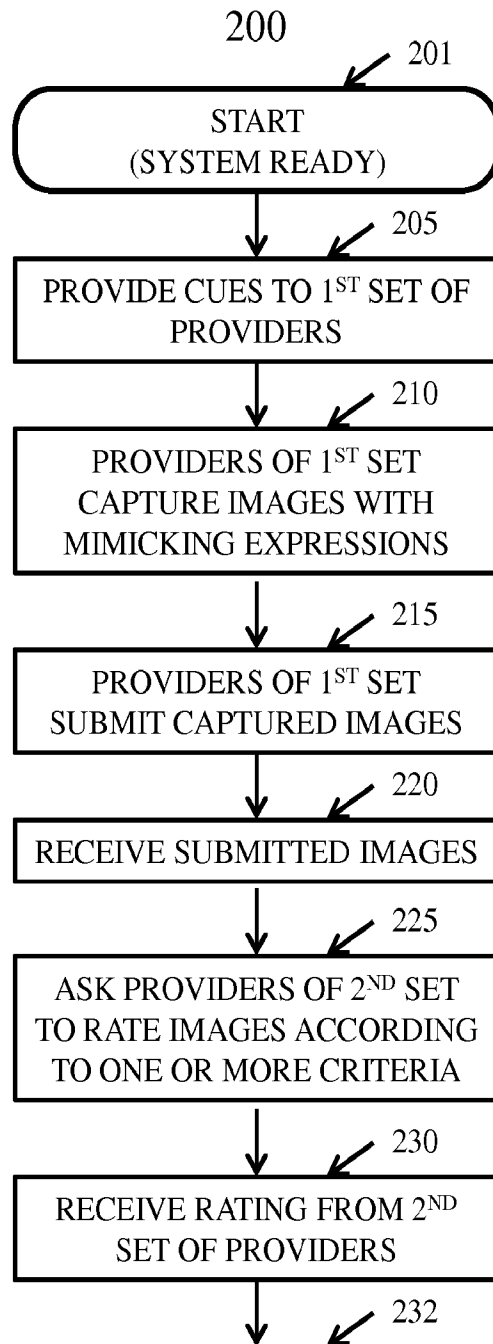


FIG. 2

3/8

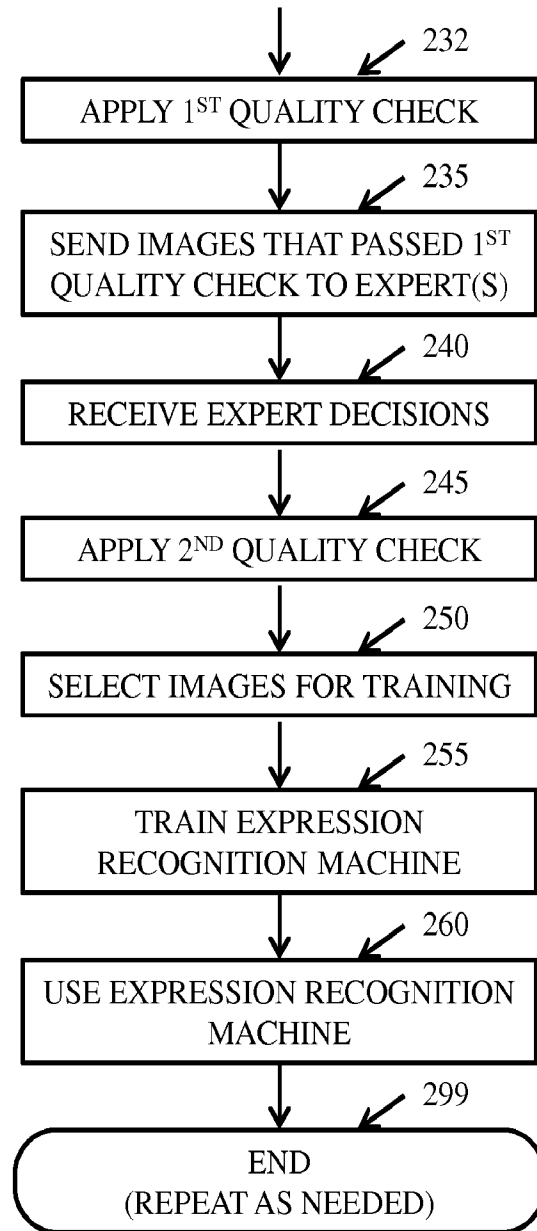


FIG. 2 (CONT.)

4/8

300

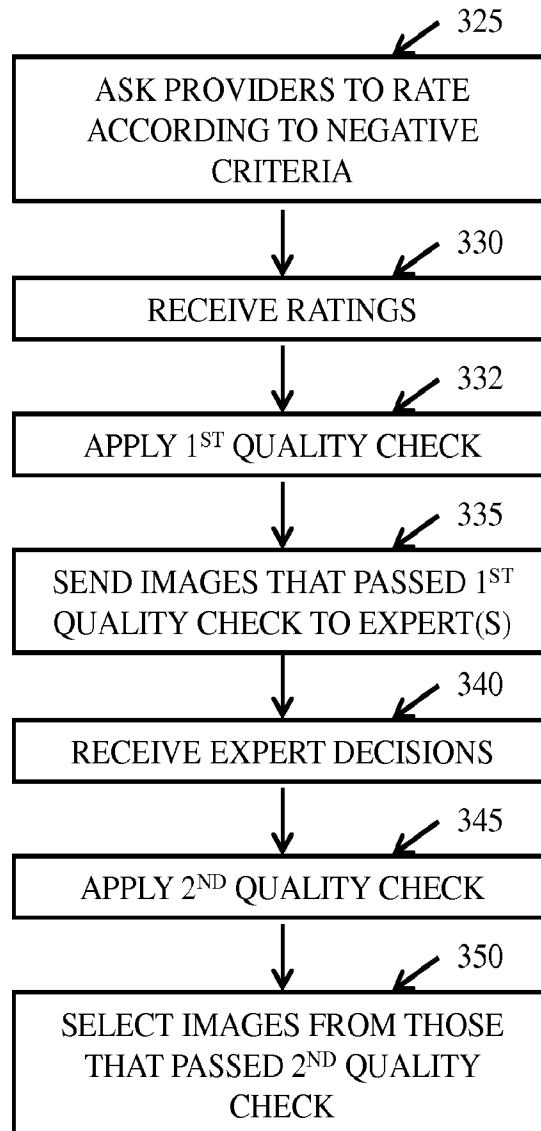


FIG. 3

5/8

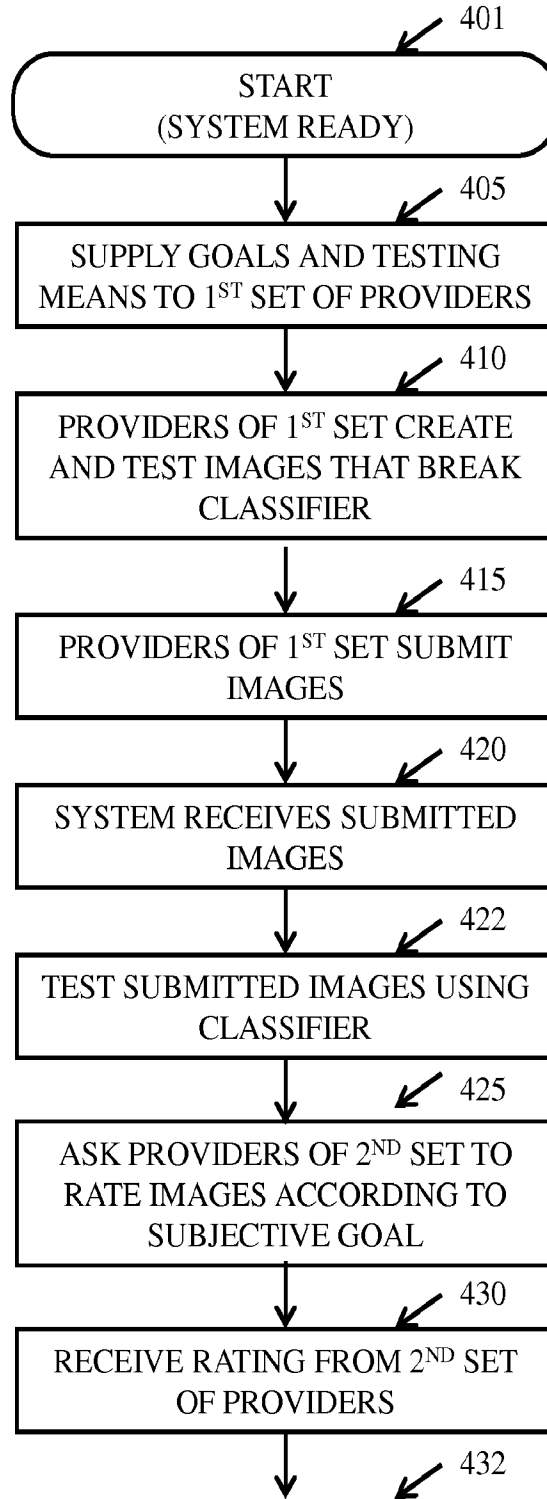


FIG. 4

6/8

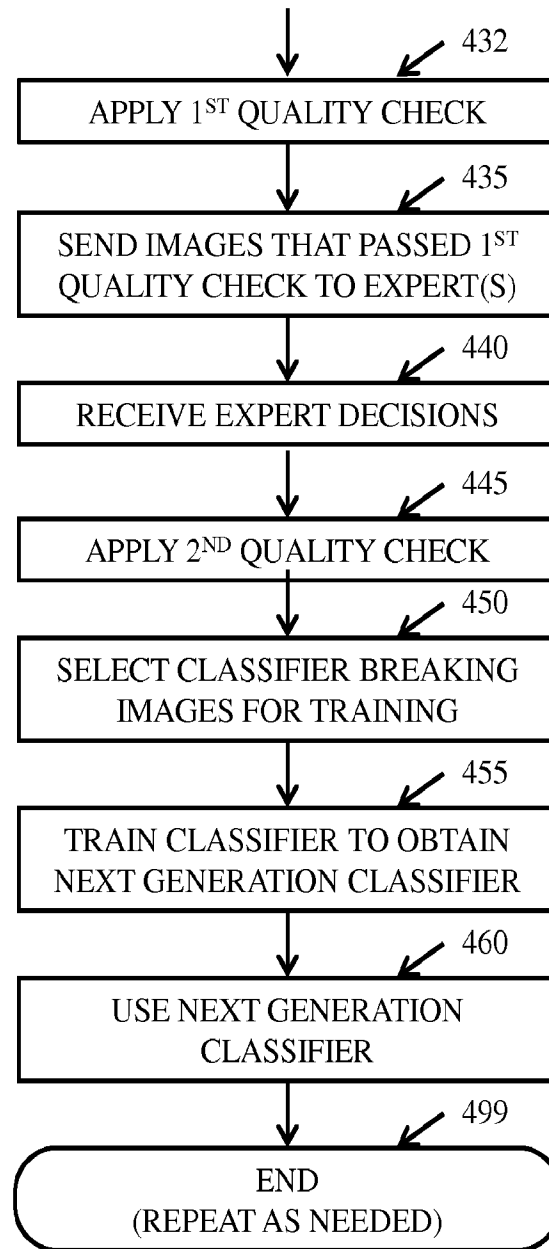


FIG. 4 (CONT.)

7/8

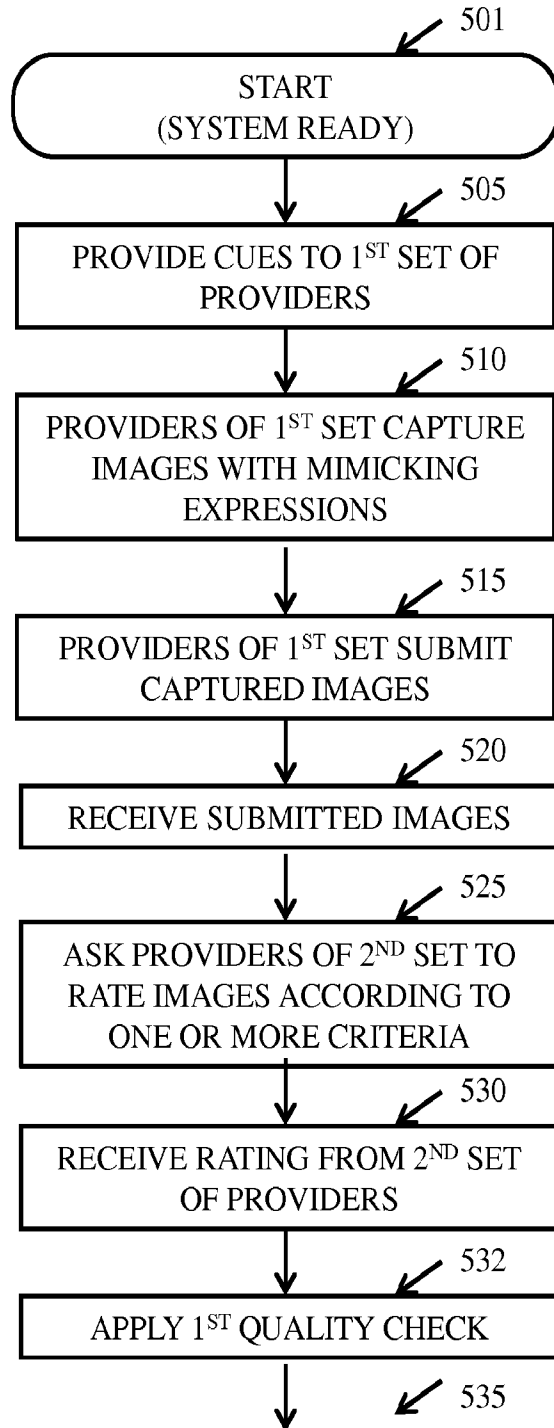


FIG. 5

8/8

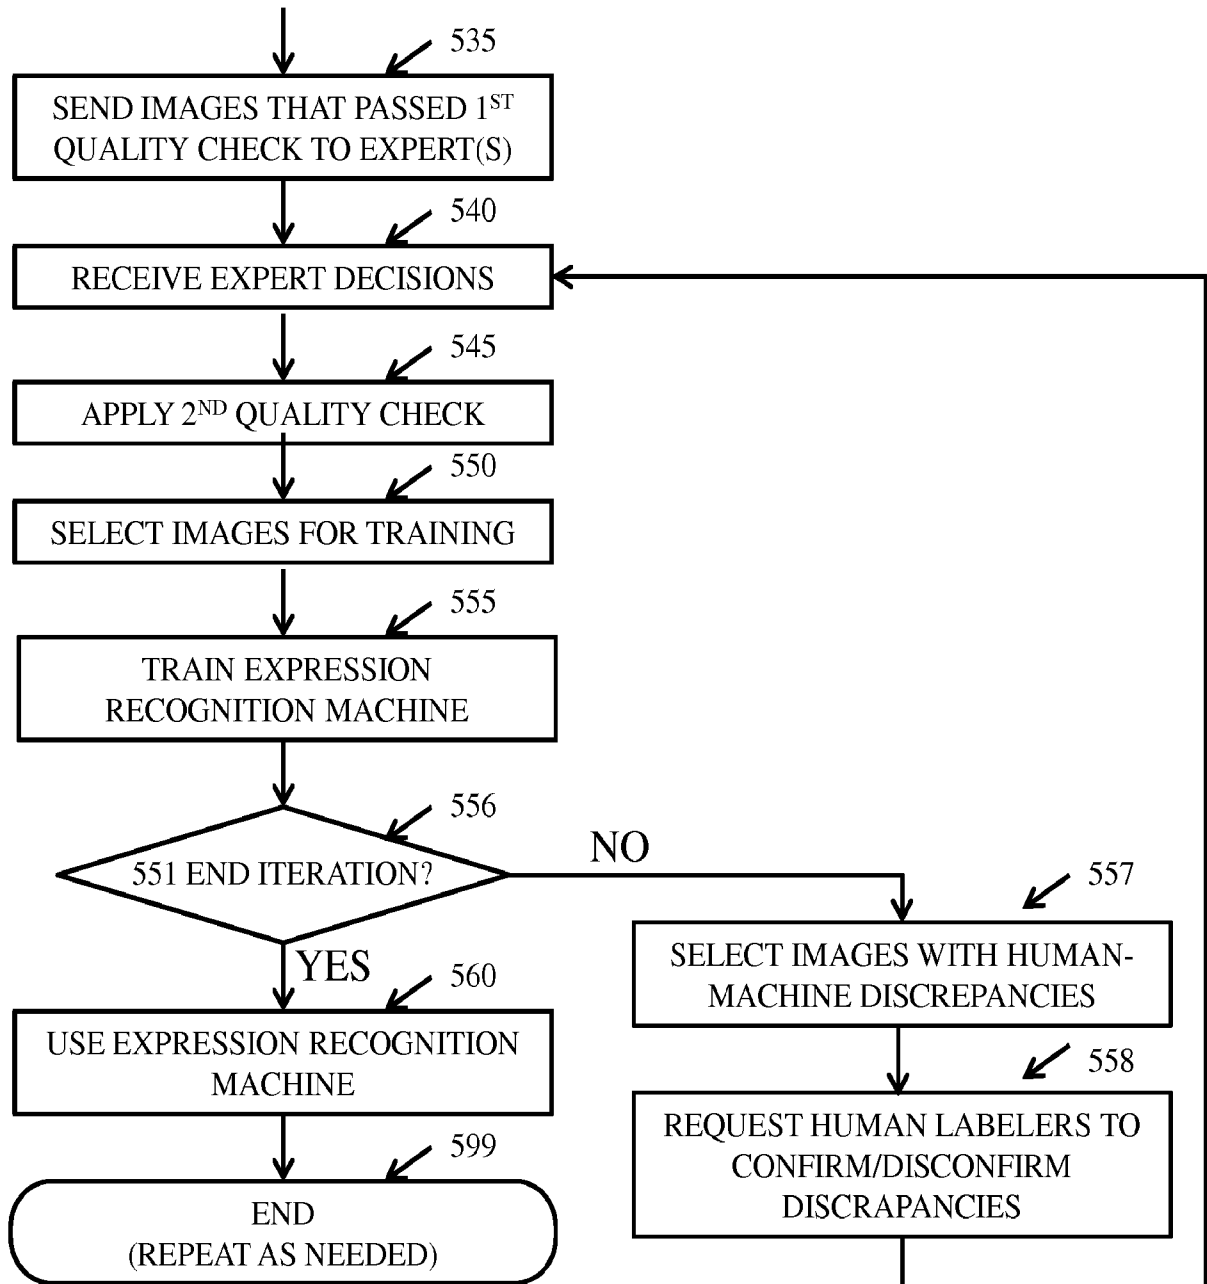


FIG. 5 (CONT.)