



(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:
PCT/US2017/061159

(22) International Filing Date:
10 November 2017 (10.11.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/420,524 10 November 2016 (10.11.2016) US
15/388,998 22 December 2016 (22.12.2016) US

(71) Applicant: **SYNAPTICS INCORPORATED** [US/US];
1251 McKay Drive, San Jose, California 95131 (US).

(72) Inventors: **KRISHNAPURA, Rohini**; 388 Mayten Way,
Fremont, California 94539 (US). **RUSO, Anthony P.**;
1251 McKay Drive, San Jose, California 95131 (US).

(74) Agent: **BOBEL, Arthur J.**; Leydig, Voit & Mayer, Ltd,
180 N. Stetson Ave STE 4900, Chicago, Illinois 60601
(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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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,

(54) Title: SYSTEMS AND METHODS FOR SPOOF DETECTION BASED ON LOCAL INTEREST POINT LOCATIONS




FIG. 3

(57) Abstract: Disclosed are a system and method for performing spoof detection. The method includes: receiving, by processor from a biometric sensor, an input image of a biometric; obtaining, by the processor, keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image; computing, by the processor, one or more anti-spoof metrics of the input image based on the keypoint locations; and, determining, by the processor, whether the input image is a replica of the biometric based on the one or more anti-spoof metrics.

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:

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

SYSTEMS AND METHODS FOR SPOOF DETECTION BASED ON LOCAL INTEREST POINT LOCATIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 62/420,524, filed on November 10, 2016, which is hereby incorporated by reference in its entirety.

FIELD

[0002] This disclosure generally relates to biometric sensors and, more particularly, to systems and methods for spoof detection based on local interest point locations.

BACKGROUND

[0003] Biometric authentication systems are used for authenticating and/or verifying users of devices incorporating the authentication systems. Biometric sensing technology provides a reliable, non-intrusive way to verify individual identity for recognition purposes.

[0004] Fingerprints, like various other biometric characteristics, are based on distinctive personal characteristics and are, thus, a reliable mechanism for recognizing an individual. There are many potential applications for using fingerprint sensors. For example, fingerprint sensors may be used to provide access control in stationary applications, such as security checkpoints. Electronic fingerprint sensors may also be used to provide access control in mobile devices, such as cell phones, wearable smart devices (e.g., smart watches and activity trackers), tablet computers, personal data assistants (PDAs), navigation devices, and portable gaming devices. Accordingly, some applications, in particular applications related to mobile devices, may require authentication systems that are both small in size and highly reliable.

[0005] Biometric “spoofing” is any attempt to circumvent biometric security using a replica of a user’s sensed biometric. In the context of fingerprint authentication systems, some examples of spoofing materials include a three-dimensional (3D) gelatin mold of a finger, a graphite mold of a finger, a wood glue mold of a finger, and printed two-dimensional (2D) image of a finger, among others. In the context of facial recognition, an example spoofing material could be a photo of person’s face. In the context of voice recognition, an example spoofing material could be a vocal imitation or playback.

[0006] In order to maintain the integrity of biometric authentication systems, there is a need for anti-spoofing systems and methods, also referred to as “liveness detection”, that can detect when an authentication attempt is a spoof and, upon spoof detection, properly deny authentication.

SUMMARY

[0007] One embodiment provides a device comprising a biometric sensor and a processing system. The processing system is configured to: receive, from the biometric sensor, an input image of a biometric; obtain keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image; compute one or more anti-spoof metrics of the input image based on the keypoint locations; and, determine whether the input image is a replica of the biometric based on the one or more anti-spoof metrics.

[0008] Another embodiment provides a method for performing spoof detection. The method includes: receiving, by processor from a biometric sensor, an input image of a biometric; obtaining, by the processor, keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image; computing, by the processor, one or more anti-spoof metrics of the input image based on the keypoint locations; and, determining, by the processor, whether the input image is a replica of the biometric based on the one or more anti-spoof metrics. Some embodiments further include a non-transitory computer-readable storage medium storing instructions that, when executed by a processor, perform the method for spoof detection.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is a block diagram of an example of an electronic device that includes a sensor device and a processing system, according to an embodiment.

[0010] FIG. 2 is a block diagram illustrating a system and method for spoof detection, according to an embodiment.

[0011] FIG. 3 is a block diagram of a spoof detection system, according to an embodiment.

[0012] FIG. 4A is an example of a live fingerprint showing keypoint locations, according to an embodiment.

[0013] FIG. 4B is an example of a spoofed fingerprint showing keypoint locations, according to an embodiment.

[0014] FIG. 5 is a flow diagram illustrating method steps for spoof detection based on local interest point locations, according to an embodiment.

DETAILED DESCRIPTION

[0015] The following detailed description is merely exemplary in nature and is not intended to limit the disclosure or the application and uses of the disclosure. Furthermore, there is no intention to be bound by any expressed or implied theory presented in the preceding technical field, background, summary, brief description of the drawings, or the following detailed description. Turning to the drawings, and as described in greater detail herein, embodiments provide systems and methods for a metric for spoof detection.

[0016] Fingerprint images can have considerable variation from person to person, from sensor to sensor, and even from different images for the same person. This variation reflects the fact that a person's finger is a living thing and changes as the finger responds to external factors, such as sweating, stretching when touched, etc. When a finger is imaged with a fingerprint sensor, these changes are captured to some extent and, hence, the sensed images include some variation, even with different impressions of the same finger. Also, different sensors from the same sensing technology can add small differences.

[0017] The goal of a robust fingerprint matcher is to be agnostic to many of these changes for the same finger, which improves the usability of the system for a genuine user. However, if spoofs are constructed, e.g., with latent prints from the user, the spoof images can be sufficiently similar to the real finger that they can be used to gain entry into the system. Fingerprint anti-spoof technology attempts to distinguish images from spoof fingers from those of live fingers by deriving properties from the images that can be used to differentiate them.

[0018] As described, anti-spoof technology helps to improve the security of a biometric authentication system by preventing a successful authentication using a spoofed fingerprint, for example a spoof fingerprint created using the latent fingerprint of a genuine user of the device. Since the spoof fingerprint is a copy of the fingerprint that is enrolled on the device, without anti-spoof technology, a fingerprint matcher could match the spoofed fingerprint image to an enrolled image and grant access to the device.

[0019] Anti-spoof technology analyzes features in the spoof fingerprint image and attempts to distinguish the spoof fingerprint image from an image of a real finger. Described

embodiments provide a system and method for spoof detection based on computing one or more metrics that are relative to prior images of the same biometric.

[0020] Once the one or more metrics are computed, the metrics are input to a “classifier,” which produces a score that can be compared against a threshold to determine if the finger is a live finger or a spoof. In some implementations, many different metrics are input into the classifier, and the classifier is configured to compute an overall liveness score based on the combination of metrics. In one example, the classifier is a neural network, but any classifier is within the scope of the disclosure. The classifier itself can be developed using machine learning methods, where a training set and a test set are created to train and validate the classifier performance.

[0021] Software-based anti-spoof technology detects spoofs by extracting features in an input image and analyzing those features to distinguish a spoof image from an image of a real finger. In some embodiments, anti-spoof solutions are performed independent of the underlying matcher (i.e., the matcher that is used to determine whether the input image provides a match to an enrollment template (or “enrollment image”) for purposes of authentication, verification, or identification, for example), such that a match/non-match decision and a spoof/non-spoof decision are made separately and independently of one another.

[0022] Determining locations for, and subsequently describing, salient local regions of interest in images is a popular method in computer vision to carry out further actions, such as pattern recognition. These salient local regions of interest can be referred to as “keypoints.” There are several methods for determining locations of keypoints and determining keypoint descriptions, or “descriptors,” for the keypoints. Methods for determining locations of keypoints aim to find locations of the keypoints in an image. Methods for determining keypoint descriptors aim to categorize a given keypoint with distinctive information. One example technique for determining locations of keypoints and determining keypoint descriptors is known as BRISK (Binary Robust Invariant Scalable Keypoints). Other examples include SIFT (Scale-Invariant Feature Transform), SURF (Speeded-Up Robust Features), KAZE, AKAZE, ORB, among others. The various techniques for determining locations of keypoints and determining keypoint descriptors can vary in terms of computation complexity.

[0023] The disclosed embodiments provide a technique for spoof detection based on the locations of keypoints in an image. The disclosed embodiments, however, do not rely on the determined keypoint descriptors and are based solely on the locations for the keypoints.

[0024] In one implementation, keypoints (such as, for example, BRISK keypoints) are detected and locations of these keypoints are used to derive certain metrics for spoof detection, without using the associated keypoint descriptors. The keypoint locations can be computed by an anti-spoof module or may be supplied to the anti-spoof module from an external resource, such as for example from a matcher that determines whether the input image provides a match to an enrollment template.

[0025] Once the keypoint locations are obtained, one or more anti-spoof metrics are computed based on the locations of the keypoints. Example anti-spoof metrics include: a count of the keypoints in the image, a count of the keypoints in an overlap region of the image to an enrollment template, a count of the keypoints in a non-overlap region of the image to an enrollment template, an average number of keypoints in the image, an average number of keypoints in an overlap region of the image to an enrollment template, an average number of keypoints in a non-overlap region of the image to an enrollment template, a count of keypoint clusters included in the image, and a count of the number of keypoints included in a cluster having the largest number of keypoints, average distances between keypoints, among others.

[0026] The metrics computed based on the keypoint locations are passed to a classifier that makes a spoof/non-spoof decision as to whether the image is of a real finger or a spoofed finger.

[0027] Turning to the figures, FIG. 1 is a block diagram of an example of an electronic device 100 that includes a sensor device 102 and a processing system 104, according to an embodiment. By way of example, basic functional components of the electronic device 100 utilized during capturing, storing, and validating a biometric match attempt are illustrated. The processing system 104 includes a processor(s) 106, a memory 108, a template storage 110, an operating system (OS) 112, and a power source(s) 114. Each of the processor(s) 106, the memory 108, the template storage 110, and the operating system 112 are interconnected physically, communicatively, and/or operatively for inter-component communications. The power source 114 is interconnected to the various system components to provide electrical power as necessary.

[0028] As illustrated, processor(s) 106 are configured to implement functionality and/or process instructions for execution within electronic device 100 and the processing system 104. For example, processor 106 executes instructions stored in memory 108 or instructions stored on template storage 110 to identify a biometric object or determine whether a biometric authentication attempt is successful or unsuccessful. Memory 108, which may be a

non-transitory, computer-readable storage medium, is configured to store information within electronic device 100 during operation. In some embodiments, memory 108 includes a temporary memory, an area for information not to be maintained when the electronic device 100 is turned off. Examples of such temporary memory include volatile memories such as random access memories (RAM), dynamic random access memories (DRAM), and static random access memories (SRAM). Memory 108 also maintains program instructions for execution by the processor 106.

[0029] Template storage 110 comprises one or more non-transitory computer-readable storage media. In the context of a fingerprint sensor, the template storage 110 is generally configured to store enrollment views for fingerprint images for a user's fingerprint or other enrollment information. The enrollment views can include multiple images of the same finger. Further, the enrollment views can include view of multiple different fingers of the user. More generally, the template storage 110 may be used to store information about an object. The template storage 110 may further be configured for long-term storage of information. In some examples, the template storage 110 includes non-volatile storage elements. Non-limiting examples of non-volatile storage elements include magnetic hard discs, solid-state drives (SSD), optical discs, floppy discs, flash memories, or forms of electrically programmable memories (EPROM) or electrically erasable and programmable (EEPROM) memories, among others.

[0030] The processing system 104 also hosts an operating system (OS) 112. The operating system 112 controls operations of the components of the processing system 104. For example, the operating system 112 facilitates the interaction of the processor(s) 106, memory 108 and template storage 110. The processing system 104, although shown as including a processor 106 and memory 108, may further include a microprocessor, microcontroller and/or dedicated circuitry.

[0031] According to various embodiments, the processor(s) 106 implement hardware and/or software to obtain data describing an image of an input object. The processor(s) 106 may also align two images and compare the aligned images to one another to determine whether there is a match. The processor(s) 106 may also operate to reconstruct a larger image from a series of smaller partial images or sub-images, such as fingerprint images when multiple partial fingerprint images are collected during a biometric process, such as an enrollment or matching process for verification or identification.

[0032] The processing system 104 includes one or more power sources 114 to provide power to the electronic device 100. Non-limiting examples of power source 114 include

single-use power sources, rechargeable power sources, and/or power sources developed from nickel-cadmium, lithium-ion, or other suitable material as well power cords and/or adapters which are in turn connected to electrical power.

[0033] Sensor device 102 can be implemented as a physical part of the electronic device 100, or can be physically separate from the electronic device 100. As appropriate, the sensor device 102 may communicate with parts of the electronic device 100 using any one or more of the following: buses, networks, and other wired or wireless interconnections. In some embodiments, sensor device 102 is implemented as a fingerprint sensor to capture a fingerprint image of a user. In accordance with the disclosure, the sensor device 102 uses optical sensing for the purpose of object imaging including imaging biometrics such as fingerprints. The sensor device 102 can be incorporated as part of a display, for example, or may be a discrete sensor. In some embodiments, the sensor device 102 may perform optical imaging. In various other embodiments, the sensor device 102 can be replaced with a capacitive sensor device, ultrasonic sensor device, or another sensor device that uses some other sensing technology for object imaging, as described in greater detail herein.

[0034] The electronic device 100 may utilize any suitable combination of sensor components and sensing technologies to detect user input in the sensing region. Some implementations utilize arrays or other regular or irregular patterns of multiple sensing elements to detect the input. Example sensing techniques that the electronic device 100 may use include capacitive sensing techniques, optical sensing techniques, acoustic (e.g., ultrasonic) sensing techniques, pressure-based (e.g., piezoelectric) sensing techniques, resistive sensing techniques, thermal sensing techniques, inductive sensing techniques, elastive sensing techniques, magnetic sensing techniques, and/or radar sensing techniques.

[0035] For example, the electronic device 100 may use resistive sensing techniques where contact from an input object closes an electrical circuit and can be used to detect input. In one example technique, the sensor device 102 includes a flexible and conductive first layer separated by one or more spacer elements from a conductive second layer. During operation, one or more voltage gradients are created across the layers. Pressing the flexible first layer may deflect it sufficiently to create electrical contact between the layers, resulting in voltage outputs reflective of the point(s) of contact between the layers. These voltage outputs may be used to determine spatial information corresponding to the input object.

[0036] In another example, the electronic device 100 may use inductive sensing techniques where one or more sensing elements pick up loop currents induced by a resonating

coil or pair of coils. Some combination of the magnitude, phase, and frequency of the currents may then be used to determine spatial information corresponding to the input object.

[0037] In another example, the electronic device 100 may use acoustic sensing techniques where one or more acoustic sensing elements detect sound waves from nearby input objects. The sound waves may be in audible frequencies or ultrasonic frequencies. The detected sound waves may include echoes of ambient sound waves and/or echoes of sound waves emitted by the input device that are reflected from surfaces of the input object. Some combination of the amplitude, phase, frequency, and or time delay of the electrical signals may be used to determine spatial information corresponding to the input object.

[0038] One example acoustic sensing technique utilizes active ultrasonic sensing to emit high frequency source waves that propagate to the sensing region. One or more ultrasonic transmitter elements (also “ultrasonic emitters”) may be used to emit high frequency sound waves to the sensing region, and one or more ultrasonic receiving elements (also “ultrasonic receivers”) may detect echoes of the emitted sound waves. Separate elements may be used to transmit and receive, or common elements that both transmit and receive may be used (e.g., ultrasonic transceivers). In some instances, emitted ultrasonic waves are able to penetrate sub-surfaces of the input object, such as dermal layers of a human finger.

[0039] In another example, the electronic device 100 may use optical sensing techniques where one or more sensing elements detect light from the sensing region. The detected light may be reflected from the input object, transmitted through the input object, emitted by input object, or some combination thereof. The detected light may be in the visible or invisible spectrum (such as infrared or ultraviolet light). Example optical sensing elements include photodiodes, CMOS image sensor arrays, CCD arrays, thin-film detectors, and other suitable photosensors sensitive to light in wavelength(s) of interest. Active illumination may be used to provide light to the sensing region, and reflections from the sensing region in the illumination wavelength(s) may be detected to determine input information corresponding to the input object.

[0040] One example optical technique utilizes direct illumination of the input object, which may or may not be in contact with an input surface of the sensing region depending on the configuration. One or more light sources and/or light guiding structures are used to direct light to the sensing region. When an input object is present, this light is reflected directly from surfaces of the input object, which reflections can be detected by the optical sensing elements and used to determine input information about the input object.

[0041] Another example optical technique utilizes indirect illumination based on internal reflection to detect input objects in contact with an input surface of the sensing region. One or more light sources are used to direct light in a transmitting medium at an angle at which it is internally reflected at the input surface of the sensing region, due to different refractive indices at opposing sides of the interface defined by the input surface. Contact of the input surface by the input object causes the refractive index to change across this boundary, which alters the internal reflection characteristics at the input surface. Higher contrast signals can often be achieved if principles of frustrated total internal reflection (FTIR) are used to detect the input object, where the light is directed to the input surface at an angle of incidence at which it is totally internally reflected, except at locations where the input object is in contact and causes the light to scatter and partially transmit across this interface at the region of contact by the input object. An example of this is presence of a finger introduced to an input surface defined by a glass to air interface. The higher refractive index of human skin compared to air causes light incident at the input surface at the critical angle of the interface to air to be partially transmitted across the input interface and scattered by the finger, where it would otherwise be totally internally reflected at the glass to air interface. This optical response can be detected by the system and used to determine spatial information. In some embodiments, this can be used to image small scale surface variations of the input object, such as fingerprint patterns, where the internal reflectivity of the incident light differs depending on whether a ridge or valley of the finger is in contact with that portion of the input surface.

[0042] In another example, the electronic device 100 may use capacitive techniques where voltage or current is applied to create an electric field. Nearby input objects cause changes in the electric field, and produce detectable changes in capacitive coupling that may be detected as changes in voltage, current, or the like. Sensor electrodes may be utilized as capacitive sensing elements. Arrays or other regular or irregular patterns of capacitive sensing elements may be used to create electric fields. Separate sensor electrodes may be ohmically shorted together to form larger sensing elements.

[0043] One example technique utilizes “self capacitance” (or “absolute capacitance”) sensing methods based on changes in the capacitive coupling between sensor electrodes and an input object. An input object near the sensor electrodes alters the electric field near the sensor electrodes, thus changing the measured capacitive coupling. An absolute capacitance sensing method may operate by modulating sensor electrodes with respect to a reference voltage (e.g. system ground), and by detecting the capacitive coupling between the sensor

electrodes and the input object. For example, the sensing element array may be modulated, or a drive ring or other conductive element that is ohmically or capacitively coupled to the input object may be modulated. The reference voltage may be a substantially constant voltage or a varying voltage, or the reference voltage may be system ground.

[0044] Another example technique utilizes “mutual capacitance” (or “transcapacitance”) sensing methods based on changes in the capacitive coupling between sensor electrodes. An input object near the sensor electrodes may alter the electric field between the sensor electrodes, thus changing the measured capacitive coupling. A transcapacitive sensing method may operate by detecting the capacitive coupling between one or more transmitter sensor electrodes (also “transmitter electrodes”) and one or more receiver sensor electrodes (also “receiver electrodes”). Transmitter sensor electrodes may be modulated relative to a reference voltage to transmit transmitter signals. Receiver sensor electrodes may be held substantially constant relative to the reference voltage to facilitate receipt of resulting signals. The reference voltage may be a substantially constant voltage or system ground. The transmitter electrodes are modulated relative to the receiver electrodes to transmit transmitter signals and to facilitate receipt of resulting signals. A resulting signal may comprise effect(s) corresponding to one or more transmitter signals, and/or to one or more sources of environmental interference (e.g. other electromagnetic signals). Sensor electrodes may be dedicated transmitters or receivers, or may be configured to both transmit and receive. Also, sensor electrodes may be dedicated transcapacitance sensing elements or absolute capacitance sensing elements, or may be operated as both transcapacitance and absolute capacitance sensing elements.

[0045] Some non-limiting examples of electronic devices 100 include personal computers of all sizes and shapes, such as desktop computers, laptop computers, netbook computers, tablets, web browsers, e-book readers, and personal digital assistants (PDAs). Additional example electronic devices 100 include composite input devices, such as physical keyboards and separate joysticks or key switches. Further example electronic devices 100 include peripherals such as data input devices (including remote controls and mice) and data output devices (including display screens and printers). Other examples include remote terminals, kiosks, video game machines (e.g., video game consoles, portable gaming devices, and the like), communication devices (including cellular phones, such as smart phones), and media devices (including recorders, editors, and players such as televisions, set-top boxes, music players, digital photo frames, and digital cameras).

[0046] FIG. 2 is a block diagram illustrating a system and method for spoof detection according to an embodiment. At step 202, a sensor captures an image of a fingerprint. The fingerprint can be either from a live finger or a spoofed finger. Although the disclosed embodiments are described in the context of fingerprint spoof detection, is noted that embodiments of the disclosure apply generally to any spoof detection system and are not limited to fingerprints.

[0047] At step 204, a processor computes a spoof detection metric based on local interest point locations, as described in greater detail below. One or more metrics may be computed at step 204. The local interest points location are referred to as keypoints in the image. In some implementations, the metrics computed at step 204 are based on the locations of the keypoints in the image, and do not rely on any keypoints descriptors associated with the keypoints.

[0048] The one or more metrics computed at step 204 are passed to a classifier. Optionally, at step 206, the processor may compute other spoof detection metrics and also pass them to the classifier. At step 208, the processor executes the classifier to determine whether the image of the fingerprint captured at step 202 is from a live finger or a spoofed finger.

[0049] FIG. 3 is a block diagram of a spoof detection system, according to an embodiment. The system includes a sensor device 102 and a processor 106. The processor 106 is configured to execute one or more software functional blocks, including an image acquisition module 302, a matcher 304, and a spoof detection module 306. The image acquisition module 302, the matcher 304, and the spoof detection module 306 are implemented as software instructions stored in a memory and executed by one or more processors 106. It is understood that each of the functional blocks may be also implemented by dedicated circuitry instead of or in combination with software.

[0050] Although shown as separate function blocks in FIG. 3, two or more of the image acquisition module 302, the matcher 304, and the spoof detection module 306 may be executed together as a single software module, application, or operating system. Alternatively, the image acquisition module 302, the matcher 304, and the spoof detection module 306 may be executed separately and/or provided by different software vendors. Also, in some embodiments, multiple processors 106 may be used to execute one or more of the image acquisition module 302, the matcher 304, and the spoof detection module 306.

[0051] In some embodiments, an input image, such as a fingerprint image, is captured by the sensor device 102. The input image is passed to the image acquisition module 302, which

determines whether the image is an enrollment image or a verification image. If the input image is an enrollment image, a template associated with the input image is stored in a matcher template storage 308 and/or the matcher template storage 308 is updated based on the new input image.

[0052] If the input image is a verification image, the image acquisition module 302 also passes the input image to the matcher 304, which is configured to determine whether the input image matches any of the enrollment images stored in the matcher template storage 308. In one implementation, the matcher 304 may compare the input image to the enrollment image to determine a difference between the images. In some embodiments, if the difference is below a threshold, a match is found; otherwise, there is no match. In other embodiments, various techniques other than a comparison to a threshold can be used to determine whether the input image is a match to any of the enrollment images. Many different techniques can be used to execute the matcher 304, including point-based techniques, ridge-based techniques, or a combination of point-based and ridge-based techniques.

[0053] In one implementation, before the matcher 304 can compare the input image to the stored enrollment images (or “templates”), the matcher 304 performs alignment. An alignment that most closely aligns the input image to one of the enrollment images is determined, and transformation corresponding to the alignment is applied to the input image. The transformation T can be represented by $T = (T_x, T_y, \theta)$, where T_x is a translation in the horizontal direction, T_y is a translation in the vertical direction, and θ is a rotation. This process is known in the art as image alignment. Various techniques may be used by the matcher 304 to compute the image alignment.

[0054] In one embodiment, after the matcher 304 performs image alignment, the matcher 304 makes a match/non-match decision. In other embodiments, the matcher generates a match score and returns the match score to another entity of the system that called the matcher 304 (e.g., the image acquisition module 302), where the other entity makes the match/non-match decision based on the match score. The match/non-match decision may be based on comparing overlapping regions of the input image and the enrollment image. In one implementation, the matcher 304 may compare the overlapping regions of the aligned input image to the enrollment image to determine a difference between the images. In some embodiments, if the difference is below a threshold, a match is found; otherwise, there is no match. It should be understood that many different techniques can be used for matching and are also within the scope of the disclosure.

[0055] In some embodiments, for enrollment images, the image acquisition module 302 also passes the input image to the spoof detection module 306, which may extract anti-spoof metrics from the input image. Example anti-spoof metrics include: an average gray level of ridges, an average gray level of valleys, one or more values as to whether the input image includes blurred areas, one or more values as to whether the input image includes relative lighter areas, one or more values as to whether the input image includes relative darker areas, texture information (for example, by computing LBP (linear binary patterns) on portions of the input image, among others. In some implementations, the anti-spoof metrics may not be discerning enough to provide adequate fingerprint matching results, i.e., since many spoofed images could satisfy a matcher that relied solely on anti-spoof metrics for matching.

[0056] The anti-spoof metrics extracted from the input image by the spoof detection module 306 are stored in an anti-spoof template in the anti-spoof template storage 310. In some embodiments, the metrics extracted from the input image can be combined with the anti-spoof metrics in the anti-spoof template, for example by averaging the metrics extracted from the input image and the anti-spoof metrics in the anti-spoof template, to generate an updated anti-spoof template. Some embodiments do not store an anti-spoof template, and the spoof/non-spoof decision is based solely on the input image.

[0057] In one implementation, the matcher template storage 308 and the anti-spoof template storage 310 comprise one storage device. In another implementation, the matcher template storage 308 and the anti-spoof template storage 310 comprise separate storage devices.

[0058] In addition, in one implementation, when a user is enrolling enrollment images, the same images are used for updating the matcher template storage 308 and the anti-spoof template storage 310. In other implementations, separate enrollment processes are used to update the matcher template storage 308 and the anti-spoof template storage 310. As such, a given enrollment image could be used to update just one or both of the matcher template storage 308 and the anti-spoof template storage 310. However, as described, other embodiments do not store any anti-spoof templates, and the spoof/non-spoof decision is based solely on the input image.

[0059] In some embodiments, if the matcher 304 does not find a match in the matcher template storage 308, then the matcher 304 takes an appropriate action, such as, for example, denying entry to a mobile device. If the matcher 304 finds a match, then the spoof detection module 306 is configured to determine whether the input image is a spoof of a live finger,

i.e., whether image is that of a real live finger or a other non-derma-based material, such as gelatin or wood glue.

[0060] In some embodiments, the spoof detection module 306 is executed as part of the matcher 304. In other embodiments, the spoof detection module 306 is executed separately from the matcher 304.

[0061] In some embodiments, the spoof detection module 306 is executed after the matcher 304 finds a positive match. In other embodiments, the spoof detection module 306 is executed before the matcher 304 makes a match/non-match decision. In still further embodiments, the spoof detection module 306 and the matcher 304 are executed in parallel.

[0062] As described in greater detail herein, the spoof detection module 306 is configured to obtain locations of keypoints in the input image and compute one or more spoof detection metrics based on the locations of keypoints. The one or more metrics are passed to a classifier that makes the anti-spoof decision.

[0063] Also, in some embodiments, the match/non-match decision of the matcher is made by a classifier associated with the matcher, which is the same classifier that makes the spoof/non-spoof decision. In other embodiments, the match/non-match decision is made by a different classifier than the classifier that makes the spoof/non-spoof decision.

[0064] FIG. 4A is an example of a live fingerprint showing keypoint locations 402, according to an embodiment. FIG. 4B is an example of a spoofed fingerprint showing keypoint locations 402, according to an embodiment. The keypoint locations 402 shown in FIGs. 4A-4B are computed using the BRISK keypoint detection system. In other embodiments, any keypoint detection system may be used to determine the keypoint locations 402.

[0065] As can be seen in the example images, the keypoint locations 402 of the live fingerprint in FIG. 4A are more evenly distributed across the image than the keypoint locations 402 of the spoofed fingerprint in FIG. 4B. Also, there is a larger number of keypoint locations 402 is found in the image of the live fingerprint in FIG. 4A as compared to the image of the spoofed fingerprint in FIG. 4B. Still further, the image of the spoofed fingerprint in FIG. 4B shows clusters 404 of keypoints that are not apparent in the image of the live fingerprint in FIG. 4A. Described embodiments exploit these characteristics of keypoints in live versus spoofed fingerprints to aid in making a spoof decision.

[0066] FIG. 5 is a flow diagram illustrating method steps for spoof detection based on local interest point locations, according to an embodiment. In one embodiment, the method

in FIG. 5 is implemented by the spoof detection module 306 in FIG. 3. According to various embodiments, the local interest point locations are referred to as keypoints.

[0067] As shown, at step 502, a processor receives an input image. The input image may be an image of a fingerprint or other biometric, such as a face, an iris, etc. In one implementation, the input image is a verification image used to authenticate a user of a system.

[0068] At step 504, the processor obtains keypoint locations in the input image. As described, many different techniques can be used to determine the keypoint locations, including BRISK, SIFT, SURF, KAZE, AKAZE, ORB, among others.

[0069] In one embodiment, the keypoint locations are obtained by the spoof detection module 306 from an external resource, such as, for example, from a matcher (e.g., matcher 304 in FIG. 3) that determines whether the input image provides a match to an enrollment template. In other embodiments, the spoof detection module 306 directly computes the keypoint locations.

[0070] At step 506, the processor computes one or more anti-spoof metrics based on the keypoint locations. As described herein, the one or more anti-spoof metrics are based on the keypoint locations, and do not rely on any keypoint descriptors or other identifying information about the keypoints themselves. According to certain embodiments, the one or more anti-spoof metrics computed at step 506 may be rotation-invariant.

[0071] One example anti-spoof metric computed at step 506 includes a count of a number of keypoints in the input image. Another example anti-spoof metric includes an average number of keypoints in the input image, for example based on image size.

[0072] Some embodiments may further compute the anti-spoof metric at step 506 based on one or more enrollment images. As described, a matcher is configured to make a match/non-match decision based on comparing the input image to one or more stored enrollment templates. A common part of the matching process is find the best alignment of the input image to a stored enrollment template. The alignment provides an indication of an overlap area and a non-overlap area of the input image to the enrollment template. In some embodiments, the anti-spoof metric computed at step 506 may include: a count of a number of keypoints in an overlapping area of the input image and an enrollment template, a count of a number of keypoints in a non-overlapping area of the input image and an enrollment template, an average number of keypoints in an overlapping area of the input image and an enrollment template (e.g., based on image size), or an average number of keypoints in a non-overlapping area of the input image and an enrollment template.

[0073] Other example anti-spoof metrics computed at step 506 are based on identifying clusters of keypoints. In one implementation, once the keypoint locations are determined for the input image, the processor scales the input image to a smaller size to generate a scaled image (e.g., one-quarter size). The processor then determines a count of a number of keypoints in input image that correspond to the same pixel in the scaled image. Such keypoints may be regarded as forming a cluster of keypoints. In one embodiment, the anti-spoof metric computed at step 506 may comprise a number of clusters of keypoints in the input image. In another embodiment, the anti-spoof metric computed at step 506 may comprise a count of the number of keypoints in the cluster having the greatest number of keypoints.

[0074] Still further examples of anti-spoof metrics computed at step 506 are based on comparing distances between keypoints. In one implementation, for each keypoint, a distance is computed between the keypoint and its n nearest neighbors. For example, n may be about 3 to 5. The average distance between the keypoint and its n nearest neighbors can then be computed to generate a distance score for the keypoint. In one embodiment, the anti-spoof metric computed at step 506 may comprise an average of the distance scores for each keypoint. In another embodiment, the distance scores can be binned into a histogram and the anti-spoof metric computed at step 506 may comprise counts for each bin of the histogram.

[0075] At step 508, the processor optionally compares the anti-spoof metrics of the input image computed at step 506 to corresponding anti-spoof metrics of an anti-spoof template 310. As described, some embodiments involve the anti-spoof module computing an anti-spoof template that includes anti-spoof metrics derived from enrollment images for the given biometric. In these embodiments, the processor may compute a difference value (or “differential” value) between a value of a given anti-spoof metric in the input image (computed at step 506) and a corresponding value of the given anti-spoof metric in the anti-spoof template 310. These difference values are then passed to the classifier. In embodiments that do not involve an anti-spoof template, the one or more anti-spoof metrics computed at step 506 are passed to the classifier, and step 508 is omitted.

[0076] At step 514, the processor executes the classifier to make a spoof decision. The classifier is configured to make a spoof decision as to whether the input image is of a real finger or of a spoofed finger, i.e., determine whether the input image is a replica of a biometric, such as finger. The classifier may be implemented as a neural network, but any classifier is within the scope of the disclosure. The classifier itself can be developed using

machine learning methods, where a training set and a test set are created to train and validate the classifier performance.

[0077] As described, in some embodiments, other anti-spoof metrics 510 unrelated to metrics based on keypoint locations can also be input into the classifier. In some embodiments, the processor may also optionally adjust, or weigh, these other anti-spoof metrics 510 based on the keypoint locations. In some implementations, the other anti-spoof metrics 510 are passed to the classifier, which weighs the other anti-spoof metrics 510.

[0078] Although this invention describes optical object imaging in the context of fingerprint image sensing, the method and system may be used to image any object.

[0079] The use of the terms “a” and “an” and “the” and “at least one” and similar referents in the context of describing the invention (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The use of the term “at least one” followed by a list of one or more items (for example, “at least one of A and B”) is to be construed to mean one item selected from the listed items (A or B) or any combination of two or more of the listed items (A and B), unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

[0080] Preferred embodiments of this invention are described herein, including the best mode known to the inventors for carrying out the invention. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate, and the inventors intend for the invention to be practiced otherwise than as specifically described herein. Accordingly, this invention includes all modifications and

equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the invention unless otherwise indicated herein or otherwise clearly contradicted by context.

WHAT IS CLAIMED IS:

1. A device, comprising:
a biometric sensor; and
a processing system configured to:
receive, from the biometric sensor, an input image of a biometric;
obtain keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image;
compute one or more anti-spoof metrics of the input image based on the keypoint locations; and
determine whether the input image is a replica of the biometric based on the one or more anti-spoof metrics.
2. The device of claim 1, wherein the keypoint locations are obtained from a matcher configured to determine whether the input image is a match to one or more enrollment images separate from the processing system determining whether the input image is a replica of the biometric.
3. The device of claim 2, wherein the matcher is configured to determine an overlap region of the input image and a first enrollment image, and the one or more anti-spoof metrics comprise a metric corresponding to a number of keypoints in the overlap region of the input image and the first enrollment image.
4. The device of claim 2, wherein the matcher is configured to determine an overlap region of the input image and a first enrollment image and a non-overlap region of the input image and the first enrollment image, and the one or more anti-spoof metrics comprise a metric corresponding to a number of keypoints in the non-overlap region of the input image and the first enrollment image.
5. The device of claim 1, wherein the one or more anti-spoof metrics comprise a metric corresponding to a number clusters of keypoints in the input image.
6. The device of claim 5, wherein the processing system is configured to:
generate a scaled image by scaling the input image to a smaller size; and

determine that two or more keypoints in the input image comprise a cluster of keypoints in response to determining that keypoint locations of the two or more keypoints in the input image correspond to a same pixel in the scaled image.

7. The device of claim 1, wherein the one or more anti-spoof metrics comprise a metric corresponding to an average distance score of the keypoints in the input image, wherein a distance score for a given keypoint is computed by calculating an average distance between the keypoint and a number of its closest neighboring keypoints.

8. The device of claim 1, wherein the keypoint locations are obtained based on the processing system executing a BRISK (Binary Robust Invariant Scalable Keypoints) algorithm on the input image.

9. The device of claim 1, wherein the processing system is further configured to: receive an anti-spoof template corresponding to the biometric; and compute the one or more anti-spoof metrics based on comparing anti-spoof metrics included in the anti-spoof template with the one or more anti-spoof metrics of the input image.

10. The device of claim 1, wherein the one or more anti-spoof metrics comprise a metric corresponding to a number of keypoints in the input image.

11. The device of claim 1, wherein the one or more anti-spoof metrics comprise a metric corresponding to an average number of keypoints in the input image based on a size of the input image.

12. A method for performing spoof detection, comprising:
receiving, by processor from a biometric sensor, an input image of a biometric;
obtaining, by the processor, keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image;
computing, by the processor, one or more anti-spoof metrics of the input image based on the keypoint locations; and
determining, by the processor, whether the input image is a replica of the biometric based on the one or more anti-spoof metrics.

13. The method of claim 12, wherein the keypoint locations are obtained from a matcher configured to determine whether the input image is a match to one or more enrollment images separate from the processor determining whether the input image is a replica of the biometric, wherein the matcher is configured to determine an overlap region of the input image and a first enrollment image, and the one or more anti-spoof metrics comprise a metric corresponding to a number of keypoints in the overlap region of the input image and the first enrollment image.

14. The method of claim 12, wherein the keypoint locations are obtained from a matcher configured to determine whether the input image is a match to one or more enrollment images separate from the processor determining whether the input image is a replica of the biometric, wherein the matcher is configured to determine an overlap region of the input image and a first enrollment image and a non-overlap region of the input image and the first enrollment image, and the one or more anti-spoof metrics comprise a metric corresponding to a number of keypoints in the non-overlap region of the input image and the first enrollment image.

15. The method of claim 12, wherein the one or more anti-spoof metrics comprise one or more of:

- a metric corresponding to a number clusters of keypoints in the input image;
- a metric corresponding to a number of keypoints in the input image;
- a metric corresponding to an average number of keypoints in the input image based on a size of the input image; and
- a metric corresponding to an average distance score of the keypoints in the input image, wherein a distance score for a given keypoint is computed by calculating an average distance between the keypoint and a number of its closest neighboring keypoints.

16. The method of claim 16, further comprising:
generating a scaled image by scaling the input image to a smaller size; and
determining that two or more keypoints in the input image comprise a cluster of keypoints in response to determining that keypoint locations of the two or more keypoints in the input image correspond to a same pixel in the scaled image.

17. The method of claim 12, wherein the keypoint locations are obtained based on the processor executing a BRISK (Binary Robust Invariant Scalable Keypoints) algorithm on the input image.

18. The method of claim 12, further comprising:
receiving an anti-spoof template corresponding to the biometric; and
computing the one or more anti-spoof metrics based on comparing anti-spoof metrics included in the anti-spoof template with the one or more anti-spoof metrics of the input image.

19. A non-transitory computer-readable storage medium storing instructions that, when executed by a processor, causes a computing device to perform spoof detection, by performing steps comprising:

receiving, from a biometric sensor, an input image of a biometric;
obtaining keypoint locations of keypoints in the input image, wherein keypoints in the input image comprise local regions of interest in the input image;
computing one or more anti-spoof metrics of the input image based on the keypoint locations; and
determining whether the input image is a replica of the biometric based on the one or more anti-spoof metrics.

20. The computer-readable storage medium of claim 19, wherein the one or more anti-spoof metrics comprise one or more of:

a metric corresponding to a number clusters of keypoints in the input image;
a metric corresponding to a number of keypoints in the input image;
a metric corresponding to an average number of keypoints in the input image based on a size of the input image; and
a metric corresponding to an average distance score of the keypoints in the input image, wherein a distance score for a given keypoint is computed by calculating an average distance between the keypoint and a number of its closest neighboring keypoints.

1/5




FIG. 1

2/5




FIG. 2

3/5




FIG. 3




FIG. 4B




FIG. 4A

5/5




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/061159**A. CLASSIFICATION OF SUBJECT MATTER****G06K 9/00(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06K 9/00; G06F 21/32

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & keywords: keypoint, location, image, fingerprint, biometric

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014-0331059 A1 (MITSUBISHI ELECTRIC RESEARCH LABORATORIES, INC.) 06 November 2014 See paragraphs [0025]-[0051], claims 9-10 and figures 1-2B.	1-20
Y	US 6282304 B1 (SERGEY OLEGOVICH NOVIKOV et al.) 28 August 2001 See columns 11-15, claims 1, 3 and figures 6-12.	1-20
Y	US 2010-0061602 A1 (SHIGEFUMI YAMADA) 11 March 2010 See paragraph [0047] and figures 1A-2.	4, 13-14
A	US 2016-0247012 A1 (PERSISTENT SYSTEMS LIMITED et al.) 25 August 2016 See claims 1-6 and figures 3-8.	1-20
A	US 2012-0070043 A1 (TERUYUKI HIGUCHI) 22 March 2012 See claims 17-23 and figures 1-6.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

02 February 2018 (02.02.2018)

Date of mailing of the international search report

05 February 2018 (05.02.2018)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

Kim, Yeonkyung

Telephone No. +82-42-481-3325



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/061159

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0331059 A1	06/11/2014	JP 2014-219967 A JP 6226802 B2 US 9036876 B2	20/11/2014 08/11/2017 19/05/2015
US 6282304 B1	28/08/2001	AU 2000-47140 A1 AU 2000-47141 A1 AU 2000-48510 A1 AU 2000-51350 A1 JP 2003-536121 A US 2002-0031244 A1 US 2002-0031245 A1 US 6487662 B1 WO 00-70542 A1 WO 00-70543 A1 WO 00-70544 A1 WO 00-70545 A1	05/12/2000 05/12/2000 05/12/2000 05/12/2000 02/12/2003 14/03/2002 14/03/2002 26/11/2002 23/11/2000 23/11/2000 23/11/2000 23/11/2000
US 2010-0061602 A1	11/03/2010	EP 2161676 A1 EP 2161676 B1 JP 2010-061576 A JP 5277804 B2 US 8509500 B2	10/03/2010 27/06/2012 18/03/2010 28/08/2013 13/08/2013
US 2016-0247012 A1	25/08/2016	US 9734384 B2	15/08/2017
US 2012-0070043 A1	22/03/2012	CN 102804229 A CN 102804229 B CN 106446808 A EP 2442274 A1 EP 2442274 A4 EP 2442274 B1 JP 2010-282519 A JP 5056798 B2 US 2014-0254895 A1 US 8768015 B2 US 9165178 B2 WO 2010-143671 A1	28/11/2012 12/10/2016 22/02/2017 18/04/2012 19/12/2012 08/03/2017 16/12/2010 24/10/2012 11/09/2014 01/07/2014 20/10/2015 16/12/2010