



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

(21) International Application Number:
PCT/US2014/065831

(22) International Filing Date:
14 November 2014 (14.11.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/905,063 15 November 2013 (15.11.2013) US
14/542,157 14 November 2014 (14.11.2014) US

(71) Applicant: **KOFAX, INC.** [US/US]; 15211 Laguna Canyon Road, Irvine, California 92618-3603 (US).

(72) Inventors: **MA, Jiyong**; 10507 Blue Granite Dr., San Diego, California 92127 (US). **MACCIOLA, Anthony**; 11214 Spectrum, Irvine, California 92618 (US). **KILBY, Steven**; 71 San Sebastian, Rancho Santa Margarita, California 92688 (US). **AMTRUP, Jan W.**; 1201 Burton St.,

Silver Spring, Maryland 20910 (US). **AQUADRO, Justin**; 38 Linseed Rd., West Hatfield, Massachusetts 01088 (US).

(74) Agent: **KOTAB, Dominic M.**; P.O. Box 721120, San Jose, California 95172 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,

[Continued on next page]

(54) Title: SYSTEMS AND METHODS FOR GENERATING COMPOSITE IMAGES OF LONG DOCUMENTS USING MOBILE VIDEO DATA




FIG. 3

(57) Abstract: Systems, methods, and computer program products are disclosed and include: initiating a capture operation using an image capture component of the mobile device, the capture operation comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation. The systems, techniques, and computer program products also include detecting a document depicted in the video data; tracking a position of the detected document throughout the video data; selecting a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generating a composite image based on at least some of the selected plurality of images.

WO 2015/073920 A1

DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, **Published:**
LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, — *with international search report (Art. 21(3))*
SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

SYSTEMS AND METHODS FOR GENERATING COMPOSITE IMAGES OF LONG DOCUMENTS USING MOBILE VIDEO DATA

FIELD OF INVENTION

[0001] The present invention relates to digital video capture and digital video data processing, more particularly to capturing and processing digital video data using a mobile device, and even more particularly to capturing video data, each frame of which depicts at least a portion of a “long” document and processing the captured video data to generate a single composite image depicting the entire “long” document.

BACKGROUND OF THE INVENTION

[0002] Modern mobile devices are well adapted to capturing images of a variety of objects, including documents, persons, automobiles, etc. Improvements to the mobile device image capture component capabilities and/or processing power make applications for capturing and/or processing digital image data using a mobile device increasingly attractive in an increasingly mobile-device-driven economy.

[0003] However, limitations of the mobile device hardware and practical limitations of capturing images using a mobile device present major challenges to efficient and effective digital image processing. For example, in line with the widespread adoption of electronic documents as a preferred medium of information transfer in many business contexts, along with the advent of mobile technology as a competent means of capturing high quality image data, many businesses have turned to a mobile-device oriented document capture and processing approach. In particular contexts, certain documents may be too large to capture in a single image having sufficient resolution to accomplish the necessary processing for downstream purposes.

[0004] Using traditional image capture and processing algorithms, especially those configured for use with a mobile device serving as the capture device, this has prevented implementation of the very useful techniques in contexts that would otherwise benefit greatly from a more modern, automated approach.

[0005] Accordingly, it would be beneficial to provide systems, methods, and/or computer program products capable of capturing and/or processing image data in a manner that overcomes

the challenges presented above and improve users' ability to capture and process data represented on "long" documents, especially using mobile devices.

SUMMARY OF THE INVENTION

[0006] System and methods for long document stitching with images captured by mobile devices are disclosed.

[0007] In one approach, a method includes initiating a capture operation using an image capture component of the mobile device, the capture operation comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detecting a document depicted in the video data; tracking a position of the detected document throughout the video data; selecting a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generating a composite image based on at least some of the selected plurality of images.

[0008] In another embodiment, a system includes a mobile device configured to execute logic, the logic being configured to cause the mobile device, upon execution thereof, to: initiate a capture operation using an image capture component of the mobile device, the capture operation comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detect a document depicted in the video data; track a position of the detected document throughout the video data; select a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generate a composite image based on at least some of the selected plurality of images.

[0009] In yet another embodiment, a computer program product includes a computer readable medium having stored thereon instructions executable by a mobile device, the instructions being configured to cause the mobile device, upon execution thereof, to: initiate a capture operation using an image capture component of the mobile device, the capture operation comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detect a document depicted in the video data; track a position of the detected document throughout the video data; select a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generate a composite image based on at least some of the selected plurality of images.

[0010] Of course, the foregoing summary is to be understood as a mere example of the presently disclosed inventive concepts and embodiments, and should not be considered limiting on the scope of the present descriptions in any manner.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 depicts a simplified schematic of a network computing environment, according to one embodiment.

[0012] FIG. 2 depicts a schematic of a computer workstation in communication with a network, according to one embodiment.

[0013] FIG. 3 depicts an exemplary schematic of a long document, according to one embodiment.

[0014] FIGS. 4A-4C depict portions of the long document depicted in FIG. 3 at various stages in a long document capture and processing algorithm, according to several embodiments.

[0015] FIG. 5 is a flowchart of a method, according to one embodiment.

[0016] FIG. 6 is a flowchart of a method, according to one embodiment.

[0017] FIG. 7 is a flowchart of a method, according to one embodiment.

DETAILED DESCRIPTION

[0018] The following description is made for the purpose of illustrating the general principles of the present invention and is not meant to limit the inventive concepts claimed herein. Further, particular features described herein can be used in combination with other described features in each of the various possible combinations and permutations.

[0019] Unless otherwise specifically defined herein, all terms are to be given their broadest possible interpretation including meanings implied from the specification as well as meanings understood by those skilled in the art and/or as defined in dictionaries, treatises, etc.

[0020] It must also be noted that, as used in the specification and the appended claims, the singular forms "a," "an" and "the" include plural referents unless otherwise specified.

[0021] The present application refers to image processing.

[0022] Images (e.g. pictures, figures, graphical schematics, single frames of movies, videos, films, clips, etc.) are preferably digital images captured by image capture components, especially image capture components of mobile devices. As understood herein, a mobile device is any device capable of receiving data without having power supplied via a physical connection (e.g. wire, cord, cable, etc.) and capable of receiving data without a physical data connection (e.g. wire, cord, cable, etc.). Mobile devices within the scope of the present disclosures include exemplary devices such as a mobile telephone, smartphone, tablet, personal digital assistant, iPod®, iPad®, BLACKBERRY® device, etc.

[0023] However, as it will become apparent from the descriptions of various functionalities, the presently disclosed mobile image processing algorithms can be applied, sometimes with certain modifications, to images coming from scanners and multifunction peripherals (MFPs). Similarly, images processed using the presently disclosed processing algorithms may be further processed using conventional scanner processing algorithms, in some approaches.

[0024] Of course, the various embodiments set forth herein may be implemented utilizing hardware, software, or any desired combination thereof. For that matter, any type of logic may be utilized which is capable of implementing the various functionality set forth herein.

[0025] One benefit of using a mobile device is that with a data plan, image processing and information processing based on captured images can be done in a much more convenient, streamlined and integrated way than previous methods that relied on presence of a scanner. However, the use of mobile devices as document(s) capture and/or processing devices has heretofore been considered unfeasible for a variety of reasons.

[0026] In one approach, an image may be captured by an image capture component of a mobile device. The term "image capture component" should be broadly interpreted to include

any type of device capable of capturing an image of a physical object external to the device, such as a piece of paper. The term “image capture component” does not encompass a peripheral scanner or multifunction device. Any type of image capture component may be used. Preferred embodiments may use image capture components having a higher resolution, e.g. 8 MP or more, ideally 12 MP or more. The image may be captured in color, grayscale, black and white, or with any other known optical effect. The term “image” as referred to herein is meant to encompass any type of data corresponding to the output of the image capture component, including raw data, processed data, etc.

[0027] As recited herein, the term “long document” should be understood to include any type of document incapable of being captured in a single still image with sufficient resolution to accomplish downstream processing of the document and/or document contents, e.g. sufficient resolution to discern the position and identity of individual characters, sufficient resolution to discern the position and identity of document features such as lines, images, reference objects such as barcodes or registration marks (e.g. substantially representing a “+” symbol), and/or sufficient resolution to distinguish the document itself from background textures also depicted in the image data depicting the document. In preferred approaches, “sufficient resolution” is to be understood as a resolution no less than a resolution corresponding to about 200 dots per inch (DPI) or 200 pixels per inch (PPI).

[0028] As discussed in greater detail below, exemplary forms of “long document” may be understood to include receipts, legal documents (e.g. a document size of approximately 8.5 inches wide by 14 inches long), promissory notes, mortgage documents, titles, deeds, posters, banners, prints, forms, envelopes, etc., as would be understood by one having ordinary skill in the art upon reading the present descriptions.

[0029] Put another way, in various embodiments a document may be considered “long” whenever the document exceeds a length of about 11 inches along a longest dimension thereof, and/or whenever the document exhibits an aspect ratio of at least about 2.5:1.

[0030] In embodiments where a document being imaged is “long” it may be particularly advantageous to orient the image capture component and the wide document so that longitudinal axes thereof are perpendicular during the capture operation. This increases the effective resolution of the images captured, as more of the document may be contained within the viewfinder at a given distance from the document than when the longitudinal axes of the document and the camera are aligned in parallel.

[0031] In addition, as referred to herein the term “textual information” should be understood to include any and all types of information that may be contained in, represented by, or derived

from, text. For example, textual information may be understood to include the position of text on a document, the identity of one or more characters (e.g. letters, numbers, symbols, etc.) depicted on the document, an identity of a series of characters (*i.e.* a “string” of text) depicted on the document, a partial or complete shape of one or more characters depicted on the document, a size of one or more characters (absolute or relative, in varied approaches), a color of one or more characters, a font type corresponding to one or more characters, etc. as would be understood by a person having ordinary skill in the art upon reading the present descriptions.

[0032] As to be distinguished from the character identity, in preferred approaches character shape refers to the appearance of markings present on the document, without necessarily including the entire marking or, in the case where the marking corresponds to a character, without necessarily including the identity of the character represented by the marking.

[0033] As recited herein, the term “document features” should be understood to include any and all types of identifying characteristic of a document other than “textual information.” For example, in various approaches document features may include a size or shape of the document itself. Document features may also include presence, absence, size, shape and/or position of any number of markings represented on the document, such as lines, images, logos, signatures, holograms, watermarks, etc. as would be understood by one having ordinary skill in the art upon reading the present descriptions. Document features may further include color information corresponding to part or all of a document, e.g. a color of the document background, a color distribution corresponding to a region of interest within the document (such as a region depicting an image, logo, hologram, signature, etc.), and/or a determination of whether or not a document depicts color information at all.

[0034] In one approach, an image capture component motion tracker is applied to track the image capture component motion relative to a long document being imaged. In order to speed up the tracking process and reduce the power consumption of smartphones, a fast and efficient image capture component tracking algorithm is applied. In the image capture component tracking algorithm, the resolution of an original captured image is reduced, and pixels in the low resolution image are downsampled. A direct image matching of those sampled pixels between a reference frame and a test frame is applied. A best matching is found as the one with minimum matching error.

[0035] For embodiments operating on a sequence of frames in a video, the accumulated image capture component motion trajectory is estimated. When the accumulated image capture component motion displacement reaches a pre-defined threshold, a picture is taken. The captured picture is either from in a video recording mode or in a picture mode. The tracking system may

notify users that the image capture component should not be moved during the picture is taken to avoid image blur.

[0036] For example, in several approaches concerning a long document, a sequence of pictures are taken, each of them is a partial image of the long document. The tracked overlap regions between the captured adjacent pictures provide the constraints to reduce the ambiguity in detailed-overlap matching or text block matching afterwards. For the detailed-overlap matching, textual information, including but not limited to: character shape, character position, character identity, character size, character color, character font, etc. are applied to recognize the text in the overlap regions of images.

[0037] The detailed-overlap matching can be based on a text block matching technique. In order to do the text block matching, a robust text line detector is applied to the recognized characters with their associated bounding boxes. The robust text line detector clusters the recognized characters based on their locations and group them in different text lines. After text line detection, a text block matching algorithm is applied to find the best text line match. The text block matching algorithm searches the best matched text line by comparing the correlation between two text blocks with different alignment hypotheses. After the best text line is found, the transform matrix from a successive image to the present image is estimated with the two text line bounding boxes. The successive image is mapped to the present image plane, and an image warping and blending procedure is applied.

[0038] In this way, a single, composite long document image is generated having sufficient resolution to enable downstream processing of the long document in many useful contexts and applications.

[0039] For example, in one general embodiment, a method includes initiating a capture operation using an image capture component of the mobile device, the capture operation comprising: capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detecting a document depicted in the video data; tracking a position of the detected document throughout the video data; selecting a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generating a composite image based on at least some of the selected plurality of images.

In another general embodiment, a system includes a mobile device configured to execute logic, the logic being configured to cause the mobile device, upon execution thereof, to: initiate a capture operation using an image capture component of the mobile device, the capture operation

comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detect a document depicted in the video data; track a position of the detected document throughout the video data; select a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generate a composite image based on at least some of the selected plurality of images.

[0040] In still yet another general embodiment, a computer program product includes a computer readable medium having stored thereon instructions executable by a mobile device, the instructions being configured to cause the mobile device, upon execution thereof, to: initiate a capture operation using an image capture component of the mobile device, the capture operation comprising; capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation; detect a document depicted in the video data; track a position of the detected document throughout the video data; select a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors; and generate a composite image based on at least some of the selected plurality of images.

[0041] The description herein is presented to enable any person skilled in the art to make and use the invention and is provided in the context of particular applications of the invention and their requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0042] In particular, various embodiments of the invention discussed herein are implemented using the Internet as a means of communicating among a plurality of computer systems. One skilled in the art will recognize that the present invention is not limited to the use of the Internet as a communication medium and that alternative methods of the invention may accommodate the use of a private intranet, a Local Area Network (LAN), a Wide Area Network (WAN) or other means of communication. In addition, various combinations of wired, wireless (e.g., radio frequency) and optical communication links may be utilized.

[0043] The program environment in which one embodiment of the invention may be executed illustratively incorporates one or more general-purpose computers or special-purpose devices such as hand-held computers. Details of such devices (e.g., processor, memory, data storage, input and output devices) are well known and are omitted for the sake of clarity.

[0044] It should also be understood that the techniques of the present invention might be implemented using a variety of technologies. For example, the methods described herein may be implemented in software running on a computer system, or implemented in hardware utilizing one or more processors and logic (hardware and/or software) for performing operations of the method, application specific integrated circuits, programmable logic devices such as Field Programmable Gate Arrays (FPGAs), and/or various combinations thereof. In one illustrative approach, methods described herein may be implemented by a series of computer-executable instructions residing on a storage medium such as a physical (e.g., non-transitory) computer-readable medium. In addition, although specific embodiments of the invention may employ object-oriented software programming concepts, the invention is not so limited and is easily adapted to employ other forms of directing the operation of a computer.

[0045] The invention can also be provided in the form of a computer program product comprising a computer readable storage or signal medium having computer code thereon, which may be executed by a computing device (e.g., a processor) and/or system. A computer readable storage medium can include any medium capable of storing computer code thereon for use by a computing device or system, including optical media such as read only and writable CD and DVD, magnetic memory or medium (e.g., hard disk drive, tape), semiconductor memory (e.g., FLASH memory and other portable memory cards, etc.), firmware encoded in a chip, etc.

[0046] A computer readable signal medium is one that does not fit within the aforementioned storage medium class. For example, illustrative computer readable signal media communicate or otherwise transfer transitory signals within a system, between systems e.g., via a physical or virtual network, etc.

[0047] FIG. 1 illustrates an architecture 100, in accordance with one embodiment. As shown in FIG. 1, a plurality of remote networks 102 are provided including a first remote network 104 and a second remote network 106. A gateway 101 may be coupled between the remote networks 102 and a proximate network 108. In the context of the present network architecture 100, the networks 104, 106 may each take any form including, but not limited to a LAN, a WAN such as the Internet, public switched telephone network (PSTN), internal telephone network, etc.

[0048] In use, the gateway 101 serves as an entrance point from the remote networks 102 to the proximate network 108. As such, the gateway 101 may function as a router, which is capable

of directing a given packet of data that arrives at the gateway 101, and a switch, which furnishes the actual path in and out of the gateway 101 for a given packet.

[0049] Further included is at least one data server 114 coupled to the proximate network 108, and which is accessible from the remote networks 102 via the gateway 101. It should be noted that the data server(s) 114 may include any type of computing device/groupware. Coupled to each data server 114 is a plurality of user devices 116. Such user devices 116 may include a desktop computer, laptop computer, hand-held computer, printer or any other type of logic. It should be noted that a user device 111 may also be directly coupled to any of the networks, in one embodiment.

[0050] A peripheral 120 or series of peripherals 120, e.g. facsimile machines, printers, networked storage units, etc., may be coupled to one or more of the networks 104, 106, 108. It should be noted that databases, servers, and/or additional components may be utilized with, or integrated into, any type of network element coupled to the networks 104, 106, 108. In the context of the present description, a network element may refer to any component of a network.

[0051] According to some approaches, methods and systems described herein may be implemented with and/or on virtual systems and/or systems which emulate one or more other systems, such as a UNIX system which emulates a MAC OS environment, a UNIX system which virtually hosts a MICROSOFT WINDOWS environment, a MICROSOFT WINDOWS system which emulates a MAC OS environment, etc. This virtualization and/or emulation may be enhanced through the use of VMWARE software, in some embodiments.

[0052] In more approaches, one or more networks 104, 106, 108, may represent a cluster of systems commonly referred to as a "cloud." In cloud computing, shared resources, such as processing power, peripherals, software, data processing and/or storage, servers, etc., are provided to any system in the cloud, preferably in an on-demand relationship, thereby allowing access and distribution of services across many computing systems. Cloud computing typically involves an Internet or other high speed connection (e.g., 4G LTE, fiber optic, etc.) between the systems operating in the cloud, but other techniques of connecting the systems may also be used.

[0053] FIG. 1 illustrates an architecture 100, in accordance with one embodiment. As shown in FIG. 1, a plurality of remote networks 102 are provided including a first remote network 104 and a second remote network 106. A gateway 101 may be coupled between the remote networks 102 and a proximate network 108. In the context of the present architecture 100, the networks 104, 106 may each take any form including, but not limited to a LAN, a WAN such as the Internet, public switched telephone network (PSTN), internal telephone network, etc.

[0054] In use, the gateway 101 serves as an entrance point from the remote networks 102 to the proximate network 108. As such, the gateway 101 may function as a router, which is capable of directing a given packet of data that arrives at the gateway 101, and a switch, which furnishes the actual path in and out of the gateway 101 for a given packet.

[0055] Further included is at least one data server 114 coupled to the proximate network 108, and which is accessible from the remote networks 102 via the gateway 101. It should be noted that the data server(s) 114 may include any type of computing device/groupware. Coupled to each data server 114 is a plurality of user devices 116. Such user devices 116 may include a desktop computer, lap-top computer, hand-held computer, printer or any other type of logic. It should be noted that a user device 111 may also be directly coupled to any of the networks, in one embodiment.

[0056] A peripheral 120 or series of peripherals 120, e.g., facsimile machines, printers, networked and/or local storage units or systems, etc., may be coupled to one or more of the networks 104, 106, 108. It should be noted that databases and/or additional components may be utilized with, or integrated into, any type of network element coupled to the networks 104, 106, 108. In the context of the present description, a network element may refer to any component of a network.

[0057] According to some approaches, methods and systems described herein may be implemented with and/or on virtual systems and/or systems which emulate one or more other systems, such as a UNIX system which emulates a MAC OS environment, a UNIX system which virtually hosts a MICROSOFT WINDOWS environment, a MICROSOFT WINDOWS system which emulates a MAC OS environment, etc. This virtualization and/or emulation may be enhanced through the use of VMWARE software, in some embodiments.

[0058] In more approaches, one or more networks 104, 106, 108, may represent a cluster of systems commonly referred to as a "cloud." In cloud computing, shared resources, such as processing power, peripherals, software, data processing and/or storage, servers, etc., are provided to any system in the cloud, preferably in an on-demand relationship, thereby allowing access and distribution of services across many computing systems. Cloud computing typically involves an Internet or other high speed connection (e.g., 4G LTE, fiber optic, etc.) between the systems operating in the cloud, but other techniques of connecting the systems may also be used.

[0059] FIG. 2 shows a representative hardware environment associated with a user device 116 and/or server 114 of FIG. 1, in accordance with one embodiment. Such figure illustrates a typical hardware configuration of a workstation having a central processing unit 210, such as a microprocessor, and a number of other units interconnected via a system bus 212.

[0060] The workstation shown in FIG. 2 includes a Random Access Memory (RAM) 214, Read Only Memory (ROM) 216, an I/O adapter 218 for connecting peripheral devices such as disk storage units 220 to the bus 212, a user interface adapter 222 for connecting a keyboard 224, a mouse 226, a speaker 228, a microphone 232, and/or other user interface devices such as a touch screen and a image capture component (not shown) to the bus 212, communication adapter 234 for connecting the workstation to a communication network 235 (e.g., a data processing network) and a display adapter 236 for connecting the bus 212 to a display device 238.

[0061] The workstation may have resident thereon an operating system such as the Microsoft Windows® Operating System (OS), a MAC OS, a UNIX OS, etc. It will be appreciated that a preferred embodiment may also be implemented on platforms and operating systems other than those mentioned. A preferred embodiment may be written using JAVA, XML, C, and/or C++ language, or other programming languages, along with an object oriented programming methodology. Object oriented programming (OOP), which has become increasingly used to develop complex applications, may be used.

[0062] An application may be installed on the mobile device, e.g., stored in a nonvolatile memory of the device. In one approach, the application includes instructions to perform processing of an image on the mobile device. In another approach, the application includes instructions to send the image to a remote server such as a network server. In yet another approach, the application may include instructions to decide whether to perform some or all processing on the mobile device and/or send the image to the remote site.

[0063] In various embodiments, the presently disclosed methods, systems and/or computer program products may utilize and/or include image processing operations such as page detection, rectangularization, detection of uneven illumination, illumination normalization, resolution estimation, blur detection, etc.

[0064] In various embodiments, the presently disclosed methods, systems and/or computer program products may utilize and/or include any classification and/or data extraction operations, including for instance classifying objects depicted in a digital image according to type based at least in part on characteristics of the object, performing custom-tailored image processing using information about the object characteristics and/or object class, building and/or using feature vectors to perform classification, building and/or using feature vectors to develop a data extraction model for the object and/or object class(es), using data extraction models to extract data from digital images, etc.

[0065] Long Document Capture

[0066] FIG. 3 depicts a schematic of an exemplary “long document” image 300 according to one embodiment. The long document image 300 substantially represents a receipt, but one having ordinary skill in the art will appreciate that the long document may include any number or type of “long documents” as defined herein and further as would be understood upon reading the present descriptions. The image 300 as shown in FIG. 3 conspicuously includes an image background 304 and an image foreground 302. The image foreground 302 preferably corresponds to the long document.

[0067] As further shown in FIG. 3, the long document includes a plurality of features such as textual information 306, 306a, a plurality of borders or separating lines 308, a reference object such as a barcode 310, and an image or logo 312. The features may be arranged in any manner throughout the document, and may even exhibit partial or complete overlap, e.g. as demonstrated by overlapping textual information 306 and 306a, in some embodiments.

[0068] FIGS. 4A-4C depict several embodiments of a long document capture process at various stages of completion, as disclosed herein. Each of FIGS. 4A-4C correspond to a selectively captured image 400, 410, 420 (respectively) that will be utilized to generate a composite single image depicting the entire document (e.g. as shown in FIG. 3).

[0069] In one exemplary embodiment, “automatic long document stitching” refers to an automatic process that can stitch partially overlapped document images captured from a camera in a video or in separate pictures. For instance, a commonly used camera in mobile devices (e.g. a camera having a resolution of about eight megapixels) may not have enough high resolution to capture a long receipt in a single picture. In order to get a sufficiently high resolution image of the long receipt with the same device, several partially overlapped images of the long receipt may be captured and stitched together. In FIGS. 4A-4C, three images with overlaps are captured, which may be stitched together as one image, substantially representing the long document as shown in FIG. 3.

[0070] In more approaches, the automatic long document stitching problem is similar to panoramic image stitching. The main difference between these two problems is that for long document stitching, the camera may be close to the document, as a result, a little movement of the camera can cause image blur. Therefore, long document stitching is more challenging. Some techniques developed for panoramic image stitching may be applied to long document stitching. However, there exists artifact at the seams of the stitching document. Because of the limited processing power of mobile devices, applying these techniques directly to long document stitching requires use of additional processing resources such as one or more GPU accelerators and/or multi-core CPU support. Considering the hardware limitations of mobile devices, and

provide an efficient approach to long document stitching. The approach involves document tracking, text block matching, and image composition, as discussed further below.

[0071] With specific respect to document capture, in a preferred embodiment a user puts a long document on a desk with a flat surface, and initiates a capture operation, e.g. within a mobile application. The user continuously moves the mobile device in a preferably straight along the longitudinal axis of the document as shown in FIG. 3. The motion should be as straight as possible to avoid situations where the document is out of camera view.

[0072] The user also preferably keeps the vertical distance between the camera and document substantially constant to avoid changes in apparent document size between the captured images.

[0073] Moreover, the speed of camera movement is preferably kept as substantially constant to make document tracking possible (*i.e.*, within the limitations of the mobile device hardware). In some approaches, constancy of motion may be monitored and the capture operation may be terminated or paused if motion deviates from desired parameters, e.g. as may be accomplished using a motion displacement threshold, described above. As will be understood by skilled artisans upon reading the present descriptions, the amount of motion displacement tolerable in a particular context may be partially dependent on factors such as camera resolution, shutter speed, etc.

[0074] Similarly, vertical camera movement may be detected and restricted based on information obtained from additional mobile device components, such as an accelerometer. Since the size of the long document is not necessarily known *a priori*, it is not desirable to utilize image data to track vertical motion. Instead, it is advantageous to query a device accelerometer and in response to determining the device has moved a predetermined amount in a predefined (e.g. vertical) direction over a predefined span of time (e.g. one centimeter over a span of one second), the capture operation may be terminated or paused.

[0075] The above references to terminating or pausing the capture operation should be understood as quality control measures, rather than as criteria causing abandonment of the presently disclosed long document stitching techniques. Subsequent to termination or pausing of the capture operation, in preferred embodiments a new (or the same) capture operation is initiated (resumed) to complete the requisite capture operation for subsequent processing and stitching.

[0076] With respect to tracking and document detection, in preferred approaches the primary aim of camera motion tracking is to track the motion of camera relative to the document in a

video. Using camera motion information, it is advantageous to estimate how the overlap between two adjacent captured images.

[0077] Various motion tracking and image registration techniques can be applied. For example, the tracking approaches can be pixel-based or feature-based. In a preferred approach a direct pixel-image approach is applied to camera motion tracking. At high capturing rates, e.g. greater than 24 frames/second in one embodiment, greater than 30 frames/sec in another embodiment, and greater than 59 frames/sec in yet another embodiment. It is preferable to ignore the camera movement in space rotations, and vertical motions (i.e. camera motions that will cause the change of the distance between camera and the document). Instead, it is preferred to only track two dimensional displacement motions, i.e., left and right motions and up/down motions along the document's longitudinal axis (as shown in FIG. 3 and FIGS. 4A-4C).

[0078] The camera motion tracking module is preferably used to determine when a picture of the document should be taken, and whether the picture should be captured automatically or manually. For instance, in one approach the first frame of the document is captured when a document detection module once detects there exists a document in the picture. For the following frames of the document image, when an image of the document should be taken is preferably determined by the specified overlap length between two adjacent frames of documents as shown in FIGS. 4A-4C. The specified overlaps between two adjacent frames of images (represented in FIGS. 4A-4C as Ω), can be converted to a number of pixels. If the accumulated camera motion/displacement is close to the specified value (also referred to herein as an "overlap threshold"), the system preferably captures an image of the document.

[0079] In additional and/or alternative approaches, the presently disclosed document tracking techniques include: downsampling captured image data to reduce the original image resolution; sampling image pixels in the downsampled image; and estimating motion vectors.

[0080] For example, in one embodiment estimating motion vectors may include a scenario where two adjacent frames of images are captured, and the first frame is defined as a reference frame, while the second frame is designated as a test frame. The residual errors between intensity of pixels in the test frame and the reference frame are computed for different hypotheses of the actual motion vectors. The best motion vector hypothesis is chosen as the one with minimum of residual errors. Preferably, the residual errors are the accumulated intensity errors of all pixels between reference frame and test frame. In one exemplary approach, assume there are ten hypotheses of motion vectors, (x_1, y_1) , (x_2, y_2) , ... (x_{10}, y_{10}) . The document tracking techniques may compare the image intensities of those ten pixels (e.g. in the test frame) with that a reference

pixel (e.g. in the reference frame. The pixel with minimum matching error would be the best matching.

[0081] Document tracking also include generating edge masks; pixels near the four edges of the reference frame may be out of camera view in the test frame; and a mask may be generated for those pixels so that they are excluded in image matching. As shown in FIG. 4C, the edge mask(s) may be generated so as to have a width Δ , where Δ is preferably a value in a range from about 5% to about 10% of a total document length as detected at the beginning of the tracking process.

[0082] The motion vector estimation and edge masking may be repeated iteratively until the entire document is captured and processed.

[0083] When the accumulated camera motion displacement has reached the pre-defined threshold, the tracking system will automatically capture an image of the partial document, and/or notify the user that a picture of the partial document will be taken. For instance, in FIG. 4B, the image is taken when the camera motion tracking system has detected that the overlap between the first picture shown in the first row and the second picture to be taken approximately equals a pre-defined overlap threshold value (e.g. 40%).

[0084] For FIG. 4C, and as a general concern for the "last" partial image in any sequence of images captured in the course of the presently disclosed techniques, the overlap between the second image and itself may not necessarily equal to the pre-defined value, accordingly the image is taken once the right part of the document is in the camera view.

[0085] In a preferred approach, the three images are taken as shown in FIGS. 4A-4C. The first one (FIG. 4A) is taken once the system has detected there is a document and its top side is in the camera view (shown at left, in FIG. 4A). The second image (FIG. 4B) is taken when the camera tracking system has detected the camera displacement has just reached the pre-defined threshold value. The third image (FIG. 4C) is taken when the bottom part of the document is in the camera view (shown at right).

[0086] Referring now to text block matching, preferred embodiments of the capture and tracking techniques described above generate a sequence of partially overlapped images of the long document. These images have overlaps between two adjacent frames. The lengths of the overlaps are estimated by the camera motion tracking system. These overlap lengths are not necessarily equal to, but are representative of, the exact overlap length between two adjacent frames of images. In other words, the overlap length provides constraints to find the detailed match between the two adjacent frames of images afterwards. These constraints reduce the ambiguity in finding the detailed match.

[0087] Most preferably, however, textual information such as the characters in the document can be used to find the detailed match.

[0088] In other embodiments, optical character recognition (OCR) can be applied. For example, an OCR module will recognize the position and identity of characters depicted in textual information throughout the various images. Bar codes, reference objects, logos, pictures, etc. may be in these images, but are preferably ignored.

[0089] In one embodiment, an OCR module is utilized to process the input image. The output image may be different from the input image because a de-skew process may be applied to the input image, to generate a de-skewed image as output. The OCR module also recognizes the input image and outputs the textual information of the recognized characters and their associated bounding boxes.

[0090] In additional embodiments, in order to find text “lines” from the recognized characters and their associated bounding boxes, it is advantageous to apply a robust text line finding algorithm for them. The robust text line algorithm may employ clustering techniques using the character bounding boxes as input. This algorithm will group characters within one line as a text line, e.g. by locating adjacent pairs of characters, then locating adjacent pairs of character pairs to form character triplets, then locating adjacent character triplets to form adjacent character quadruplets, etc. etc. as would be understood by a skilled artisan upon reading the present descriptions. Subsequently, text lines in the pre-defined region of an image are preferably organized as a text block, which may be used as the basic unit of comparison for text block matching, as described herein

[0091] The text block matching approach, in one embodiment, is as follows: for two text blocks in the overlap regions of two adjacent images compute a correlation between at least two text blocks; find the best matching alignment hypothesis based on the correlation; generate, for the particular alignment hypothesis, a text block matching score based on a number of characters in the two text blocks that match (e.g. exhibit substantially same character identity and character position); and sum the text block matching scores to generate a text line matching score.

[0092] Referring now to image composition and stitching, also known as “generating a composite image,” in various embodiments the presently disclosed inventive concepts include some or all of the following features.

[0093] In one embodiment, and based on the results of the text block matching described above, bounding boxes of the text lines in the best match are used to estimate an affine or homograph transform matrix, also referred to herein as a “first transformation matrix.” The first transform matrix is applied to every pixel in the second image (test frame) to transform the

second image to coordinate system in the first image (reference frame). In this way, the second image is adjusted to the first image plane, and a composite image including information depicted in both the two images is derived.

[0094] In more embodiments, e.g. for a third image which has overlap with the second image, the same procedure mentioned for the first two images is applied to get the second transform matrix to map the third image to the second image plane. The first transform matrix multiplied by the second transform matrix is the accumulated transform matrix which maps the third image to the first image plane. In this way, for any number of images to be composed, the accumulated transform matrices can be derived, and applied to the images.

[0095] FIG. 5 depicts an exemplary flowchart of a method 500 for accomplishing long document capture, according to one embodiment of the present disclosures. As would be understood by one having ordinary skill in the art reading these descriptions, the method 500 may be performed in any environment, including those depicted in FIGS. 1-4C, in various embodiments.

[0096] As shown in FIG. 5, method 500 includes operation 502, where a capture operation is initiated using a capture component of a mobile device. The capture operation preferably includes capturing video data, and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation.

[0097] Method 500 also includes operation 504, where a document depicted in the video data is captured. Preferably, the document is a “long document” as defined herein.

[0098] Method 500 further includes operation 506, where a position of the detected document is tracked throughout the video data. As recited above, “throughout” should be understood to include both temporal and data-based measures. In other words, tracking a document “throughout” video data may include tracking the document in each portion of the entirety of the video data (even if performed over a course of several discontinuous spans of time) and/or tracking the document during an entire duration of the time during which video data are captured.

[0099] Method 500 still further includes operation 508, where a plurality of images, each depicting a portion of the document, are selected using the image capture component. The selection is based in whole or in part on the tracked document position and the estimated motion vectors.

[00100] Method 500 also includes operation 510, where a composite image is generated based on at least some of the selected images.

[00101] In various embodiments, the method 500 may be performed exclusively using a mobile device, or parts of the method may be performed using the mobile device and other parts may be performed using other resources such as a workstation or network server. Preferably, where the method is performed across multiple devices, at least the capturing, the detecting, the tracking,

[00102] In various approaches, the presently disclosed long document capture and processing techniques may be embodied as a computer program product, which may have any or all of the features described herein.

[00103] For example, in one embodiment, a computer program product may include a computer readable medium having stored thereon computer readable instructions effective to cause a computing device, upon execution thereof, to perform a method, e.g. method 500 as represented in FIG. 5 and discussed above. Preferably, the computing device is a mobile device, but in alternative approaches the computing device may include any combination of devices such as a mobile device, a computer workstation, a network server, etc. as would be understood by one having ordinary skill in the art upon reading the present descriptions. However, it is also to be understood that the inventive embodiments disclosed herein are specially configured to enable operation of mobile devices in the context of long document capture techniques, which are otherwise not possible using conventional mobile devices and image processing approaches.

[00104] The computer program product may further include instructions configured to cause the mobile device to store at least some of the selected images to a memory of the mobile device in response to selecting the images.

[00105] With respect to tracking, in preferred embodiments the tracking functionality is based exclusively on the estimated plurality of motion vectors. In additional approaches, the tracking functionality may be based on textual information and/or document features instead of, or in addition to, the plurality of motion vectors.

[00106] The computer program product may additionally and/or alternatively include instructions configured to cause the mobile device to: determine at least one motion displacement based on some or all of the estimated plurality of motion vectors, each motion displacement corresponding to the image capture component during the capture operation; and terminate the capture operation in response to determining one of the motion displacement(s) is characterized by a value exceeding a predefined motion displacement threshold.

[00107] The predefined motion displacement threshold may have a value in a range from about 25 microns to about 50 microns, from about 30 microns to about 45 microns, from about 35 microns to about 40 microns, or a value of about 37.5 microns, in various approaches. In other

approaches, the motion displacement threshold may have a value measured in pixels, and be in a range from about 5 pixels to about 25 pixels, about 10 pixels to about 20 pixels, about 5 pixels to about 10 pixels, 5 pixels, or any value in these ranges.

[00108] In more approaches, the instructions configured to cause the mobile device to detect the document may additionally and/or alternatively include instructions configured to cause the mobile device to identify at least one edge of the document depicted in the captured video data.

[00109] Preferably, each of the selected plurality of images depicts a portion of the document, and the composite image depicts an entirety of the document. However, in additional embodiments the composite image may depict only portions of the document, e.g. portions that are relevant to a downstream processing operation or particular transaction to which the document relates. Similarly, in even more approaches a portion of the document that satisfies quality control criteria or other prerequisite criteria (e.g. image format, image resolution, image size, etc.) may be represented in the composite image, while other portions not satisfying one or more of the quality control (or other prerequisite) criteria may be omitted from the composite image.

[00110] The composite image may also be characterized by at least one of: an image resolution greater than an image resolution of any of the selected plurality of images; and an image size greater than an image size of any of the selected plurality of images. For example, the composite image may have a length approximately equal to a sum of lengths of the plurality of images from which the composite image was generated.

[00111] In other embodiments, the composite image may have a length approximately equal to a sum of lengths of the plurality of images from which the composite image was generated, but discounting an amount of overlap between the plurality of images from which the composite image was generated. For instance, if an overlap of approximately one half (50%) is utilized as a threshold overlap, then the length of the composite image may be approximately equal to two-thirds the sum of the lengths of the plurality of images from which the composite image was generated. Similarly, if the overlap threshold is approximately one third (33%), then the length of the composite image may be approximately equal to four-fifths the sum of the lengths of the plurality of images from which the composite image was generated.

[00112] Most preferably, each selected image depicts a portion of the document, and the composite image depicts only portion(s) of the document that correspond to a business event (e.g. financial transaction, contract formation) memorialized by the document.

[00113] In additional embodiments, the computer program product may further include instructions configured to cause the mobile device to: identify, based on the composite image,

one or more portions of the document depicting textual information; classify each identified portion of the document based on the textual information depicted therein; determine whether each classified portion is relevant to the financial transaction or irrelevant to the financial transaction, the determining being based on the portion classification; and remove each portion determined to be irrelevant to the financial transaction from the composite image.

[00114] The computer program product may even further comprise instructions configured to cause the mobile device to: align the portions determined to be relevant to the financial transaction; and generate a second composite image, wherein the second composite image is characterized by: approximately a same image size as an image size of the composite image; approximately a same image resolution as an image resolution of the composite image; excluding textual information irrelevant to the financial transaction; and including textual information relevant to the financial transaction.

[00115] Preferably, a plurality of characters comprising the textual information relevant to the financial transaction are aligned with one another, so that all textual information depicted in the composite image is substantially aligned along a single orientation or angle, as is the case with a single image of a document (assuming all textual information is similarly aligned within the physical document itself).

[00116] In several approaches, the instructions configured to cause the mobile device to select the plurality of images may include instructions configured to cause the mobile device to define a plurality of frame pairs. Each frame pair may consist of a reference frame and a test frame, while each reference frame and each test frame is selected from the video data.

[00117] The instructions configured to cause the mobile device to select the plurality of images may additionally and/or alternatively include instructions configured to cause the mobile device to: determine an amount of overlap between the reference frame and the test frame of each frame pair; and select an image corresponding to each frame pair for which the amount of overlap between the reference frame and the test frame is greater than a predetermined overlap threshold.

[00118] Preferably, the amount of overlap corresponds to the document, as opposed to background textures depicted in the test frame and/or the reference frame. As will be appreciated by those having ordinary skill in the art, alignment based on background textures (or even including background textures, but not necessarily based thereon) have an undesirable tendency to generate false-positive alignment results with respect to the document (or indeed, any other object depicted in the image).

[00119] Even more preferably, in various embodiments the predetermined overlap threshold corresponds to a distance of at least 50%, at least 40%, at least 33%, or at least 25% of a length of the reference frame. In other approaches, the overlap threshold may be defined with respect to the length of the document, as opposed to the length of the portion(s) of the document depicted in a particular reference frame or reference frames.

[00120] The instructions configured to cause the mobile device to generate the composite image, in multiple embodiments, further comprise instructions configured to cause the mobile device to: detect textual information in each of the reference frame and the test frame of at least one frame pair. The textual information is depicted in the document, as opposed to textual information that may appear in the image background.

[00121] In additional approaches, the instructions configured to cause the mobile device to detect textual information in the reference frame(s) and the test frame(s) include instructions configured to cause the mobile device to: define, in the reference frame, at least one rectangular portion of the document depicting some or all of the textual information; define, in the test frame, at least one corresponding rectangular portion of the document depicting some or all of the textual information; and align the document depicted in the test frame with the document depicted in the reference frame. Put another way, alignment operates such that the test frame is aligned with the reference frame, using the document (as opposed, for example, to frame edges or background textures) as the point of reference for the alignment.

[00122] In various embodiments, the alignment may be based on one or more of the following: textual information, document features, document edges, etc. as would be understood by one having ordinary skill in the art upon reading the present descriptions.

[00123] In preferred approaches, the textual information comprises at least one of: an identity of one or more characters represented in the rectangular portion; an identity of one or more characters represented in the corresponding rectangular portion; a sequence of characters represented in the rectangular portion; a sequence of characters represented in the corresponding rectangular portion; a position of one or more characters represented in the rectangular portion; a position of one or more characters represented in the corresponding rectangular portion; an absolute size of one or more characters represented in the rectangular portion; an absolute size of one or more characters represented in the corresponding rectangular portion a size of one or more characters represented in the rectangular portion relative to a size of one or more characters represented in the corresponding rectangular portion; a size of one or more characters represented in the corresponding rectangular portion relative to a size of one or more characters represented in the rectangular portion; a color of one or more characters represented in the

rectangular portion; a color of one or more characters represented in the corresponding rectangular portion; a shape of one or more characters represented in the rectangular portion; and a shape of one or more characters represented in the corresponding rectangular portion.

[00124] In even more approaches, the instructions configured to cause the mobile device to align the document depicted in the test frame with the document depicted in the reference frame include instructions configured to cause the mobile device to perform optical character recognition (OCR) on at least the rectangular portion and the corresponding rectangular portion. In these approaches, alignment may be preferably performed utilizing character location and character identity as primary points of reference.

[00125] Furthermore, the instructions configured to cause the mobile device to generate the composite image may further comprise instructions configured to cause the mobile device to: detect a skew angle (e.g. θ as depicted in FIGS. 4A-4C) in one or more of the reference frame and the test frame of at least one of the frame pairs, the skew angle corresponding to the document and having a magnitude of > 0.0 degrees (as depicted in FIG. 4B); and correct the skew angle in at least one of the reference frame and the test frame. The document depicted in the composite image is characterized by a skew angle of approximately 0.0 degrees (e.g. as depicted in FIG. 3).

[00126] The computer program product may further include instructions configured to cause the mobile device to downsample the video data, e.g. by a factor of 5, and the instructions configured to cause the mobile device to detect the document, track the position of the document, and select the plurality of images is configured to perform the detecting, the tracking, and the selecting using the downsampled video data.

[00127] The preceding descriptions elaborate upon several features involved in the presently disclosed long document capture and processing techniques. The additional details should be understood as capable of being combined in any suitable manner that would be appreciated by one having ordinary skill in the art upon reading the present descriptions. Any synthesis, combination, permutation, etc. of the features disclosed herein may be included or excluded from various embodiments without departing from the scope of the present disclosure.

[00128] Classification and Extraction Implementations

[00129] In one approach, for example, document classification may be performed in a manner substantially similar to the flow diagram 600 shown in FIG. 6. Of course, the flow diagram is presented merely by way of example to facilitate understanding of the inventive concepts disclosed herein, and is not intended to be limiting on the scope of the present application.

[00130] In one approach, document classification may proceed as follows. In operation 602, a rectified image is received, preferably at a mobile device.

[00131] In operation 604, an image processing engine, e.g. a processor of a mobile device or server, synchronizes with a classification knowledgebase. The classification knowledgebase may preferably include a plurality of predefined document classes, defined according to unique features thereof, e.g. via a feature vector and/or plurality of reference feature matrices.

[00132] In operation 606, classification is performed on the rectified image using the classification knowledgebase.

[00133] In operation 608, a result of the classification operation, e.g., success or failure, is determined.

[00134] In operation 610, in response to determining classification was successful, a document type is automatically assigned to the rectified image. The automatically assigned document type is based on the successful classification result.

[00135] In operation 612, on the other hand, and in response to determining the classification operation result was failure, a document type is manually assigned to the rectified image. Moreover, the classification knowledgebase is preferably updated with the manually assigned document type so that in future situations where similar documents are presented in the rectified image, it will be possible to automatically assign the corresponding document type based on the expanded classification knowledge base, e.g., similar to as described above with reference to operation 610.

[00136] In operation 614, either the automatically assigned document type or the manually assigned document type is reported, preferably to a user or via being displayed on a display of the mobile device.

[00137] In additional approaches, it is advantageous to utilize image classification techniques, which present unique challenges in the context of long document capture and processing, since many long documents (e.g. receipts) are not standard forms of a given size, and the features shown on a particular long document can change position with respect to other types of long document. In such scenarios, it is useful to divide the receipt into several portions (e.g. similar as described above with respect to test and reference frames, and particularly with respect to text block matching), and classify each portion. Based on the classification determined for each portion, an overall document classification may be determined, e.g. in preferred approaches each portion is determined to belong to a same particular classification, which is thus determined to be the classification of the entire document. Of course, in practical applications 100% agreement is

rare, so and the most frequently determined class among the various portion classifications may be utilized as the document classification.

[00138] Similarly, extraction may be performed in a manner substantially similar to the flow diagram 700 shown in FIG. 7. Again, the flow diagram is not to be considered limiting in any way, but merely an illustrative example of one embodiment of the presently described inventive concepts.

[00139] In operation 702, an image depicting a document, and having associated therewith a document type corresponding to the document, is received (preferably at a mobile device).

[00140] In operation 704, an extraction taxonomy is determined based on the document type.

[00141] In operation 706, it is determined whether the extraction taxonomy corresponds to an extraction knowledgebase.

[00142] In operation 708a, if the extraction taxonomy does not correspond to the extraction knowledgebase, content is recognized within the document (e.g. using OCR). The recognized content includes metadata.

[00143] In operation 710a, a new extraction model is trained based on the recognized content. If the extraction taxonomy does not correspond to the extraction knowledgebase, the method 700 proceeds to operation 714.

[00144] On the other hand, in operation 708b, if the extraction taxonomy corresponds to an extraction knowledgebase, the extraction knowledgebase is associated with the document type.

[00145] In operation 710b, content is recognized within the document, the recognition being based at least in part on the extraction knowledgebase, and the recognized content including metadata.

[00146] In operation 712b, the metadata are selectively extracted based on the extraction knowledgebase.

[00147] In operation 714, the metadata are validated based on one or more of associative validation information in an associative validation database, and predefined business rules.

[00148] In operation 716, an intelligent document (preferably a PDF) is generated based on the validated metadata and one or more of the extraction knowledgebase, the predefined business rules, and the document type.

[00149] In more approaches, the presently disclosed methods, systems, and/or computer program products may be utilized with, implemented in, and/or include one or more user interfaces configured to facilitate performing any functionality disclosed herein and/or in the aforementioned related Patent Application, such as an image processing mobile application, a case management application, and/or a classification application, in multiple embodiments.

[00150] In still more approaches, the presently disclosed systems, methods and/or computer program products may be advantageously applied to one or more of the use methodologies and/or scenarios disclosed in the aforementioned related Patent Application, among others that would be appreciated by one having ordinary skill in the art upon reading these descriptions.

[00151] It will further be appreciated that embodiments presented herein may be provided in the form of a service deployed on behalf of a customer to offer service on demand.

[00152] *Exemplary Use Cases*

[00153] Bill Pay

[00154] One major challenge in Bill-Pay scenarios (and for other applications involving mobile capture) is taking a high quality picture of the source document, e.g. invoice, packing slip, remittance slip, etc. Preferably, the document should fill the viewfinder to a large degree, with no clipped corners or edges. The document should also preferably be adequately lighted, in focus, and taken at an angle with relatively small deviations from normal (e.g. the imaging device being oriented in a plane substantially parallel the document) to minimize distortions. It should also have good background separation, and a uniform background with respect to texture, color, and/or illumination, etc. The use cases and development areas described below outline useful applications for technology that can assist a user in achieving these goals.

[00155] Automatic Capture Quality Assistance

[00156] When a user submits a document, such as a bill for payment, the automatic capture should preferably only take a picture when a document is truly positioned in the viewfinder, a situation which may be verified by the imaging device using various techniques.

[00157] For example, in one embodiment the imaging device (e.g. mobile phone) may preprocess the video feed to detect a single document in the video frame. In one exemplary approach, preprocessing involves finding features of a document page (e.g. edges or areas of similar color) and some reasoning about what set of features constitutes a document. In various implementations of this technology, it is possible to use a function such as an “opencv function” to find regions within an image that have been preprocessed using filters such as a Laplacian filter or other similar filter as would be understood by one having ordinary skill in the art upon reading the present descriptions.

[00158] Using such filters, it is possible to provide an indication to the application software when the document is completely within the viewfinder. Also, in preferred embodiments it is possible to measure the stability of the device (e.g. using integrated device components such as the accelerometer) and the angle of the device, e.g. relative to a horizontal surface upon which the document is placed (e.g. using integrated device components such as the gyroscope).

Accordingly, preferred embodiments are capable of initiating the capture process for picture-taking only when one or more, and preferably all three, quality measures are met.

[00159] Stability in detection

[00160] In some approaches, the indicator of the detected document may be unstable, e.g. move around too much to capture a desirably clear image or verify the located document in the video preview. As a result, it is desirable to have more stable document detection in the video preview.

[00161] In one approach, stability may be enhanced by utilizing a procedure where, instead of detecting a single document in a single video frame, a multi-frame approach would be to average the movement of detected edges over a window of time, thereby avoiding rapid movement of the document hypothesis.

[00162] Complex backgrounds

[00163] In various instances, it is advantageous for a document having a predominantly light-toned background (e.g. a white background with black text printed thereon) to be imaged in front of a relatively dark background. Unfortunately, in many cases a user may not have convenient access to a dark, uniformly colored surface. Accordingly, it is favorable for the user to be enabled to take the picture of the document on surfaces with more heterogeneous texture.

[00164] Evaluate lighting conditions

[00165] Ideally, capture device image capture components should automatically evaluate ambient light conditions and optimize capture settings to ensure adequate exposure. However, the light sensor is not directly accessible in some devices. In some approaches, it is possible to overcome this limitation by using the back and front image capture components as indirect light sensors.

[00166] In one embodiment it is possible to evaluate lighting using only the back image capture component (i.e. the one that is eventually taking the picture). The device can evaluate the brightness distribution of a video frame and take a picture only if that distribution matches situations previously found or otherwise known to lead to good exposure.

[00167] In another embodiment, it is possible to evaluate lighting conditions by taking two pictures automatically, one with flash, and one without, and analyzing which one gives the better exposure. It is also possible to mix the two images (flash and no-flash), particularly when captured in short succession (e.g. within less than one second of one another, which may be accomplished utilizing features such as “burst mode” capture and HDR imagery).

[00168] In more embodiments, it is possible to use a mobile device light, e.g. an LED such as a video lamp (torch, flashlight, etc.) to find the best possible capture conditions. The lighting

level of the light can be adjusted, so the device could ramp up the light, take frames along the way, analyze which one gives the best exposure, and take a high quality exposure with that setting. Good exposure may be indicated by any of the exemplary quality measures described above.

[00169] In still more embodiments, the device may take pictures from both image capture components (back and front) and analyze the brightness distribution for both. This approach preferably reveals situations where the main light source is behind the image capture component, e.g. where an image capture component casts a shadow on the document, etc., and the user may be directed to move to a new location and re-evaluate the brightness distribution for better capture conditions.

[00170] For example, the brightness setting of the screen of the device might be accessible through the device's resident operating system (OS) application programming interface (API). In such cases, the brightness setting should preferably be correlated to the amount of light hitting the device (e.g. as may be measured according to an amount of light entering one or more image capture components of the mobile device), although not necessarily the amount of light hitting the document surface. This correlation allows the device and/or software application to get a sense of the ambient light present.

[00171] Avoid glare

[00172] As will be appreciated by skilled artisans reading the instant disclosures, it may be advantageous apply the presently techniques to situations where glare presents an additional challenge, e.g. where a user wants to take a picture of a glossy document, such as a driver's license, a gift card, a credit card, etc. However, when the user attempts to take the picture, there is glare due to the reflections, and/or any embedded hologram may show through the surface of the document. Preferable embodiments of the presently disclosed techniques and products convey abilities to avoid glare in glossy documents.

[00173] For example, in one approach it is preferable if pictures of documents with reflective surfaces or embedded holograms are taken at a slight angle. The presently disclosed techniques may therefore utilize sample frames, e.g. to detect potential glare. It also is advantageous in some approaches to use the detected document within the frame to estimate the current angle of the image capture component to the document (i.e., independent of information that may be provided by other components of the mobile device to determine mobile device orientation, such as an accelerometer, compass, gyroscope, etc. as would be understood by one having ordinary skill in the art upon reading the present descriptions). In this embodiment, the user is provided directions to then guide the user to take a picture at a slight angle to the document, e.g. an angle of about 5

degrees 10 degrees, or 15 degrees deviation from normal, with respect to the predominant planar orientation of the document.

[00174] Show the document in detail

[00175] In several approaches, a user may desire to review a document being captured in detail, while aligning the image capture component of the mobile device with the document. To accomplish this result, the present techniques may utilize a combined deskew and cropping approach, wherein (optionally in response to detecting presence of a complete document depicted within the field of view of the image capture component), a frame is captured, and the frame is cropped and straightened, and the resulting document is shown in full size within the viewfinder.

[00176] Indicate the document category

[00177] When a user frames an image of a document, in some instances it may be advantageous to indicate what category the document is, thereby making more certain that the capture and/or downstream processing will be successful.

[00178] When a document has been found, the presently disclosed techniques may include cropping and straightening the document as described above, followed by performing a classification operation. The success of classification operation may be visually indicated, e.g. with a green overlay over the document, and potentially the category is output to the device display, to memory, to a downstream processing application, library function, call, etc.

[00179] Capture multiple documents, one at a time

[00180] In some instances, it is advantageous to be able to capture two (or more) documents laid out on a surface such as a desk. A user may want to move the image capture component slowly over the documents, have the image capture component automatically detect each, and capture images of each, all without necessarily taking a picture of the same document twice.

[00181] In addition to finding a document in a frame, this use case benefits significantly from the ability to track documents. Tracking allows the system to take an image of the first document, keep tracking it, but not taking another picture of it. As soon as the second document comes into view, the system switches to that document and starts tracking it. When it is in full view, a picture will be taken. Document tracking may be performed substantially as described above with respect to long document capture, in preferred embodiments.

[00182] Capture multiple documents at the same time

[00183] In more complex situations, a user may need to capture two (or more) documents laid out on a surface such as a desk, and sometimes the documents may be positioned in close proximity to each other, presenting an additional challenge to distinguishing between the two documents. In order to address this challenge, the presently disclosed techniques may direct the

user to move the image capture component slowly over the documents and so that the image capture component automatically detects each document and captures an image or images of the documents, without taking a picture of the same document twice. For example, the mobile device may provide auditory instructions to the user indicating a preferred direction of motion.

[00184] Preferably, the techniques track multiple documents in a single frame. However, it is also within the scope of the present disclosures to track the movement of plural documents over multiple frames. In some implementations, it is preferable to take multiple images and keep only a selected portion of the multiple captured images, e.g. for tracking. From the selected portion, the multiple-document capture algorithms isolate the individual documents that are shown (which might be all of the documents or only some of the documents), and determine which image of each document is the best one. Document tracking may be performed substantially as described above with respect to long document capture, in preferred embodiments.

[00185] Capture and track multiple documents

[00186] Additional applications include capturing multiple documents and tracking those documents in real or near-real time. In such approaches, while the user is hovering the capture device over the documents, the user is preferably provided an indication of what each document is (e.g. document classification) and further indicates which of the documents have been captured in an image as described above.

[00187] For instance, within the viewfinder, one could mark each document that has been processed with an overlay (yellow for detected and desired category, green for captured document, and gray for undesired category).

[00188] In another instance, imagine a scene where a user starts with a high-level overview of their desk. The system detects five documents and is able to classify them. Four of them are relevant to the user's particular use for the documents (e.g. are relevant to a particular transaction, or are suitable for downstream processing, etc.), and are colored yellow. The irrelevant document is colored gray.

[00189] Further, suppose the resolution of images captured (or images that would be captured from the current position where the image capture component is located) and depicting the four relevant documents is not high enough to allow good quality images of all four desired documents. For example, the user (and by extension the image capture component of the mobile device) may be too far away from three of the four documents, such that only one document is close enough to allow for a sufficiently high resolution. The system takes a picture, isolates that document, and performs one or more quality assurance checks (e.g. for image clarity, brightness, etc.), and marks the document with a green overlay in response to determining the quality

assurance checks are passed. The user then moves the image capture component to have another document appear bigger in the viewfinder. The system tracks all documents, and snaps another picture of the document that is now in better view. The image of that document is isolated and checked, and marked with a green overlay. Similarly, the other documents are captured.

[00190] During this procedure, some of the documents already detected and processed might disappear from the view (the user zooms in on one particular document, and the others are not in view). When the user goes back to the bigger picture, the system has to recognize those documents as already having been processed. Therefore, in at least some instances it would be advantageous to perform a “scene analysis” wherein the position of documents relative to each other is recorded and analyzed.

[00191] A computer program product comprising a computer readable medium having stored thereon instructions executable by a mobile device, the instructions being configured to cause the mobile device, upon execution thereof, to: initiate a capture operation using an image capture component of the mobile device. The capture operation includes capturing video data; and estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation. The instructions are also configured to cause the mobile device to detect a document depicted in the video data; track a position of the detected document throughout the video data; select a plurality of images using the image capture component of the mobile device, and generate a composite image based on at least some of the selected plurality of images. The selection is based at least in part on: the tracked position of the detected document; and the estimated motion vectors. The document is a long document. The tracking is based exclusively on the estimated plurality of motion vectors. The instructions are further configured to cause the mobile device to: determine at least one motion displacement based on some or all of the estimated plurality of motion vectors, each motion displacement corresponding to the image capture component during the capture operation; either terminate or pause the capture operation in response to determining one of the motion displacement(s) is characterized by a value exceeding a predefined motion displacement threshold; and either initiate a new capture operation in response to terminating the capture operation; or resume the capture operation in response to pausing the capture operation. The predefined motion displacement threshold has a value in a range from about 5 pixels to about 25 pixels. The instructions configured to cause the mobile device to detect the document are configured to cause the mobile device to identify at least one edge of the document depicted in the captured video data. Each of the selected plurality of images depicts a portion of the document, while the composite image depicts an entirety of the document. Moreover, the composite image is characterized by at least one of: an image

resolution greater than an image resolution of any of the selected plurality of images; and an image size greater than an image size of any of the selected plurality of images. Each selected image depicts a portion of the document, and the composite image depicts only portion(s) of the document that correspond to a financial transaction memorialized by the document. The instructions are also configured to cause the mobile device to: identify, based on the composite image, one or more portions of the document depicting textual information; classify each identified portion of the document based on the textual information depicted therein; determine whether each classified portion is relevant to the financial transaction or irrelevant to the financial transaction, the determining being based on the portion classification; and remove each portion determined to be irrelevant to the financial transaction from the composite image. The computer program product also includes instructions configured to cause the mobile device to: align the portions determined to be relevant to the financial transaction; and generate a second composite image, wherein the second composite image is characterized by: approximately a same image size as an image size of the composite image; approximately a same image resolution as an image resolution of the composite image; excluding textual information irrelevant to the financial transaction; and including textual information relevant to the financial transaction, wherein a plurality of characters comprising the textual information relevant to the financial transaction are aligned. The instructions configured to cause the mobile device to select the plurality of images include instructions configured to cause the mobile device to define a plurality of frame pairs. Each frame pair consists of a reference frame and a test frame, while reference frame and each test frame is selected from the video data. The instructions configured to cause the mobile device to select the plurality of images further comprising instructions configured to cause the mobile device to: determine an amount of overlap between the reference frame and the test frame of each frame pair; and select an image corresponding to each frame pair for which the amount of overlap between the reference frame and the test frame is greater than a predetermined overlap threshold. The amount of overlap corresponds to the document, not to a background depicted in the reference frame, not to a background depicted in the test frame. The predetermined overlap threshold corresponds to a distance of at least 40% of a length of the reference frame. The instructions configured to cause the mobile device to generate the composite image further includes instructions configured to cause the mobile device to: detect textual information in each of the reference frame and the test frame of at least one frame pair, the textual information being depicted in the document. The instructions configured to cause the mobile device to detect textual information in each of the reference frame and the test frame include instructions configured to cause the mobile device to: define, in the reference frame, at

least one rectangular portion of the document depicting some or all of the textual information; define, in the test frame, at least one corresponding rectangular portion of the document depicting some or all of the textual information; and align the document depicted in the test frame with the document depicted in the reference frame. The alignment is based on: the textual information depicted in at least one of the rectangular portion(s); and the textual information depicted in at least one of the corresponding rectangular portion(s). The textual information comprises at least one of: an identity of one or more characters represented in the rectangular portion; an identity of one or more characters represented in the corresponding rectangular portion; a sequence of characters represented in the rectangular portion; a sequence of characters represented in the corresponding rectangular portion; a position of one or more characters represented in the rectangular portion; a position of one or more characters represented in the corresponding rectangular portion; an absolute size of one or more characters represented in the rectangular portion; an absolute size of one or more characters represented in the corresponding rectangular portion; a size of one or more characters represented in the rectangular portion relative to a size of one or more characters represented in the corresponding rectangular portion; a size of one or more characters represented in the corresponding rectangular portion relative to a size of one or more characters represented in the rectangular portion; a color of one or more characters represented in the rectangular portion; a color of one or more characters represented in the corresponding rectangular portion; a shape of one or more characters represented in the rectangular portion; and a shape of one or more characters represented in the corresponding rectangular portion. The instructions configured to cause the mobile device to align the document depicted in the test frame with the document depicted in the reference frame comprise instructions configured to cause the mobile device to perform optical character recognition (OCR) on at least the rectangular portion and the corresponding rectangular portion. The instructions configured to cause the mobile device to generate the composite image further include instructions configured to cause the mobile device to: detect a skew angle in one or more of the reference frame and the test frame of at least one of the frame pairs, the skew angle corresponding to the document and having a magnitude of > 0.0 degrees; and correct the skew angle in at least one of the reference frame and the test frame. The document depicted in the composite image is characterized by a skew angle of approximately 0.0 degrees. The computer program product also includes instructions configured to cause the mobile device to downsample the video data, and wherein the instructions configured to cause the mobile device to detect the document, track the position of the document, and select the plurality of images is configured to perform the detecting, the tracking, and the selecting using the downsampled video data.

[00192] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of an embodiment of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

CLAIMS

What is claimed is:

1. A computer program product comprising a computer readable medium having stored thereon instructions executable by a mobile device, the instructions being configured to cause the mobile device, upon execution thereof, to:
 - initiate a capture operation using an image capture component of the mobile device, the capture operation comprising:
 - capturing video data; and
 - estimating a plurality of motion vectors corresponding to motion of the image capture component during the capture operation;
 - detect a document depicted in the video data;
 - track a position of the detected document throughout the video data;
 - select a plurality of images using the image capture component of the mobile device, wherein the selection is based at least in part on:
 - the tracked position of the detected document; and
 - the estimated motion vectors; and
 - generate a composite image based on at least some of the selected plurality of images.
2. The computer program product as recited in claim 1, wherein the document is a long document.
3. The computer program product as recited in claim 1, wherein the tracking is based exclusively on the estimated plurality of motion vectors.
4. The computer program product as recited in claim 1, further comprising instructions configured to cause the mobile device to:
 - determine at least one motion displacement based on some or all of the estimated plurality of motion vectors, each motion displacement corresponding to the image capture component during the capture operation; and
 - either terminate or pause the capture operation in response to determining one of the motion displacement(s) is characterized by a value exceeding a predefined motion displacement threshold; and
 - either initiate a new capture operation in response to terminating the capture operation; or

resume the capture operation in response to pausing the capture operation.

5. The computer program product as recited in claim 4, wherein the predefined motion displacement threshold has a value in a range from about 5 pixels to about 25 pixels.
6. The computer program product as recited in claim 1, the instructions configured to cause the mobile device to detect the document further comprising instructions configured to cause the mobile device to identify at least one edge of the document depicted in the captured video data.
7. The computer program product as recited in claim 1, wherein each of the selected plurality of images depicts a portion of the document, and
wherein the composite image depicts an entirety of the document.
8. The computer program product as recited in claim 1, wherein the composite image is characterized by at least one of:
an image resolution greater than an image resolution of any of the selected plurality of images; and
an image size greater than an image size of any of the selected plurality of images.
9. The computer program product as recited in claim 1, wherein each selected image depicts a portion of the document, and
wherein the composite image depicts only portion(s) of the document that correspond to a financial transaction memorialized by the document.
10. The computer program product as recited in claim 9, further comprising instructions configured to cause the mobile device to:
identify, based on the composite image, one or more portions of the document depicting textual information;
classify each identified portion of the document based on the textual information depicted therein; determine whether each classified portion is relevant to the financial transaction or irrelevant to the financial transaction, the determining being based on the portion classification; and
remove each portion determined to be irrelevant to the financial transaction from the composite image.

11. The computer program product as recited in claim 10, further comprising instructions configured to cause the mobile device to:
 - align the portions determined to be relevant to the financial transaction; and
 - generate a second composite image, wherein the second composite image is characterized by:
 - approximately a same image size as an image size of the composite image;
 - approximately a same image resolution as an image resolution of the composite image;
 - excluding textual information irrelevant to the financial transaction; and
 - including textual information relevant to the financial transaction, wherein a plurality of characters comprising the textual information relevant to the financial transaction are aligned.
12. The computer program product as recited in claim 1, the instructions configured to cause the mobile device to select the plurality of images further comprising instructions configured to cause the mobile device to define a plurality of frame pairs, wherein each frame pair consists of a reference frame and a test frame, and wherein each reference frame and each test frame is selected from the video data.
13. The computer program product as recited in claim 12, the instructions configured to cause the mobile device to select the plurality of images further comprising instructions configured to cause the mobile device to:
 - determine an amount of overlap between the reference frame and the test frame of each frame pair; and
 - select an image corresponding to each frame pair for which the amount of overlap between the reference frame and the test frame is greater than a predetermined overlap threshold.
14. The computer program product as recited in claim 13, wherein the amount of overlap corresponds to the document,
 - wherein the amount of overlap does not correspond to a background depicted in the reference frame, and
 - wherein the amount of overlap does not correspond to a background depicted in the test frame.

15. The computer program product as recited in claim 13, wherein the predetermined overlap threshold corresponds to a distance of at least 40% of a length of the reference frame.
16. The computer program product as recited in claim 12, the instructions configured to cause the mobile device to generate the composite image further comprising instructions configured to cause the mobile device to:
 - detect textual information in each of the reference frame and the test frame of at least one frame pair, the textual information being depicted in the document.
17. The computer program product as recited in claim 16, the instructions configured to cause the mobile device to detect textual information in each of the reference frame and the test frame of at least one frame pair further comprising instructions configured to cause the mobile device to:
 - define, in the reference frame, at least one rectangular portion of the document depicting some or all of the textual information;
 - define, in the test frame, at least one corresponding rectangular portion of the document depicting some or all of the textual information; and
 - align the document depicted in the test frame with the document depicted in the reference frame, the alignment being based on:
 - the textual information depicted in at least one of the rectangular portion(s); and
 - the textual information depicted in at least one of the corresponding rectangular portion(s).
18. The computer program product as recited in claim 17, wherein the textual information comprises at least one of:
 - an identity of one or more characters represented in the rectangular portion;
 - an identity of one or more characters represented in the corresponding rectangular portion;
 - a sequence of characters represented in the rectangular portion;
 - a sequence of characters represented in the corresponding rectangular portion;
 - a position of one or more characters represented in the rectangular portion;
 - a position of one or more characters represented in the corresponding rectangular portion;
 - an absolute size of one or more characters represented in the rectangular portion;
 - an absolute size of one or more characters represented in the corresponding rectangular

portion

a size of one or more characters represented in the rectangular portion relative to a size of one or more characters represented in the corresponding rectangular portion;

a size of one or more characters represented in the corresponding rectangular portion relative to a size of one or more characters represented in the rectangular portion;

a color of one or more characters represented in the rectangular portion;

a color of one or more characters represented in the corresponding rectangular portion;

a shape of one or more characters represented in the rectangular portion; and

a shape of one or more characters represented in the corresponding rectangular portion.

19. The computer program product as recited in claim 17, the instructions configured to cause the mobile device to align the document depicted in the test frame with the document depicted in the reference frame further comprising instructions configured to cause the mobile device to perform optical character recognition (OCR) on at least the rectangular portion and the corresponding rectangular portion.

20. The computer program product as recited in claim 12, the instructions configured to cause the mobile device to generate the composite image further comprising instructions configured to cause the mobile device to:

detect a skew angle in one or more of the reference frame and the test frame of at least one of the frame pairs, the skew angle corresponding to the document and having a magnitude of > 0.0 degrees; and

correct the skew angle in at least one of the reference frame and the test frame,

wherein the document depicted in the composite image is characterized by a skew angle of approximately 0.0 degrees.

21. The computer program product as recited in claim 1, further comprising instructions configured to cause the mobile device to downsample the video data, and

wherein the instructions configured to cause the mobile device to detect the document, track the position of the document, and select the plurality of images is configured to perform the detecting, the tracking, and the selecting using the downsampled video data.




FIG. 1




FIG. 2




FIG. 3

4/7




FIG. 4B



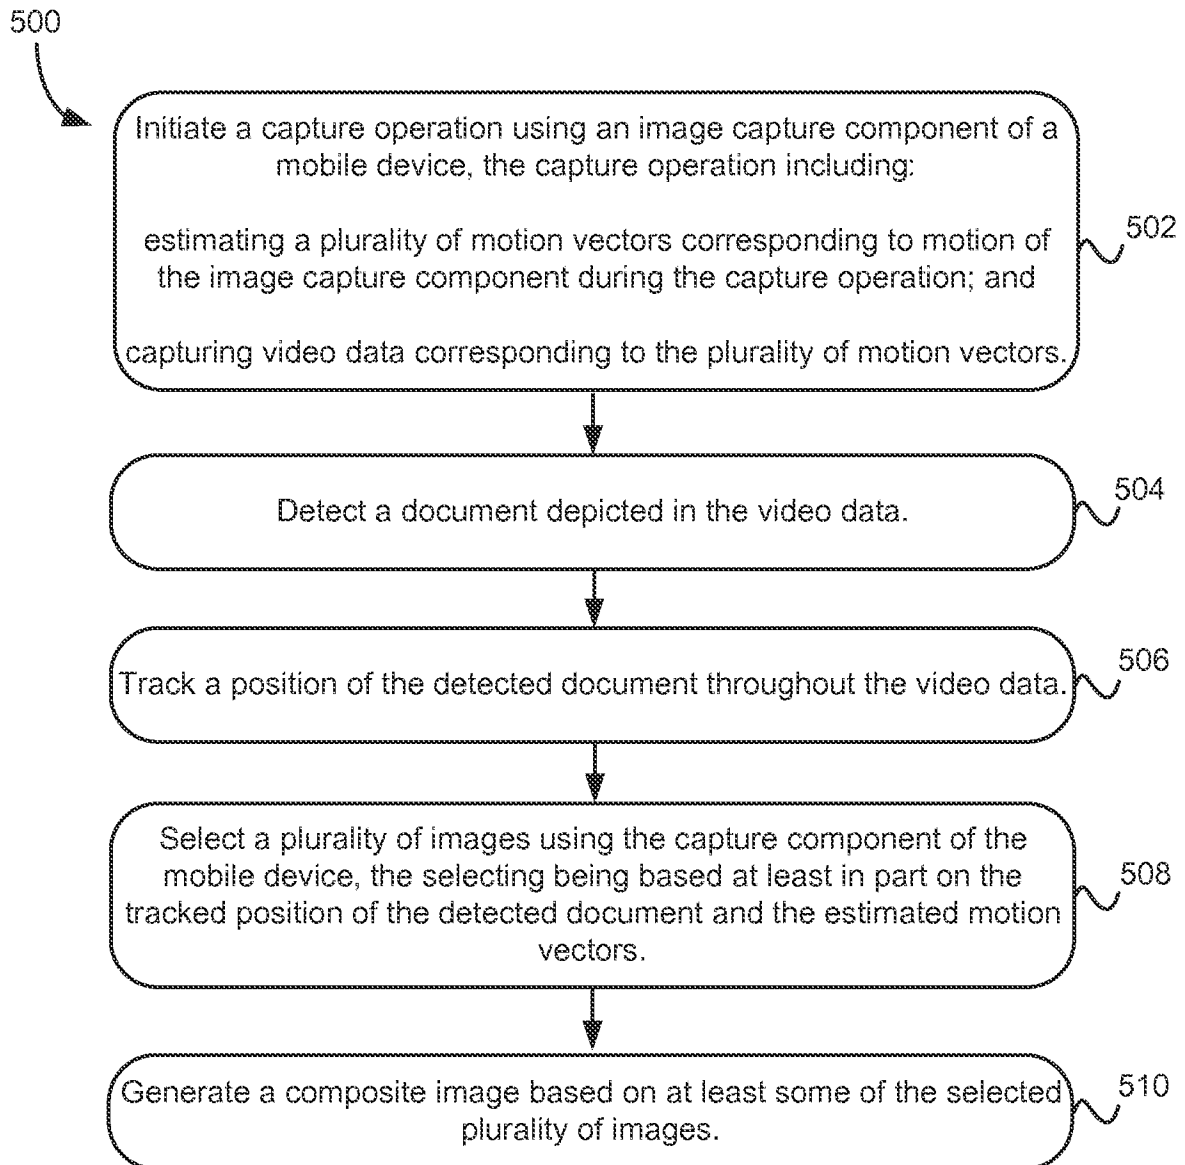


FIG. 4C



5/7

**FIG. 5**

6/7

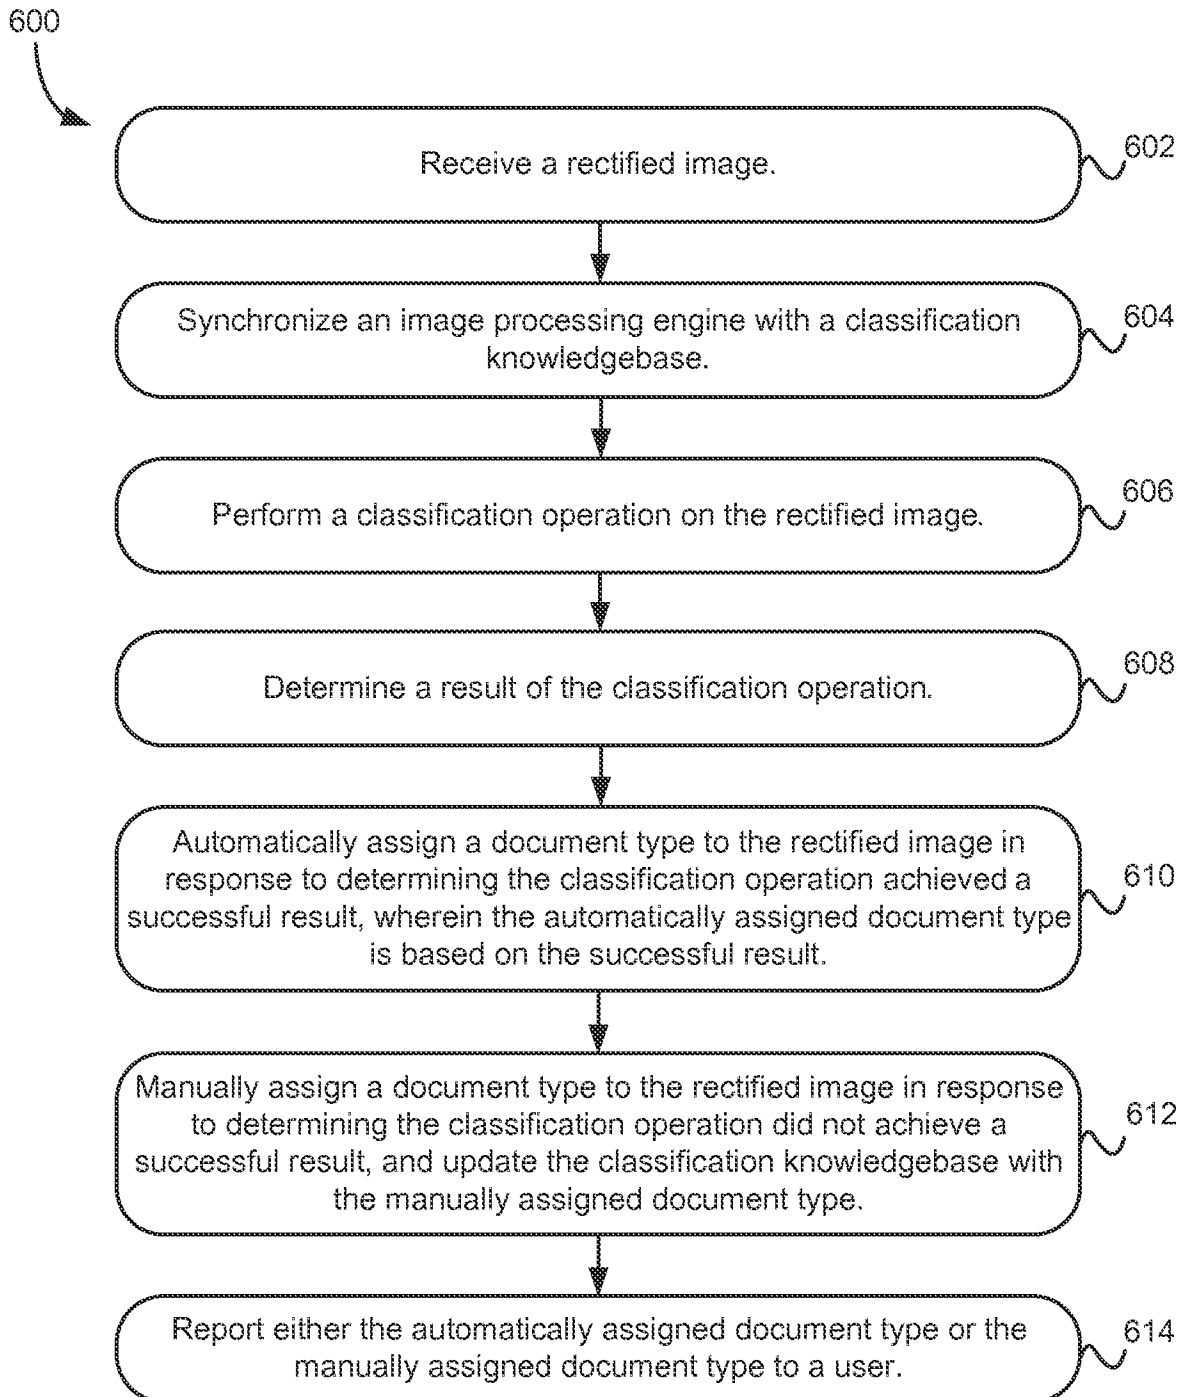


FIG. 6

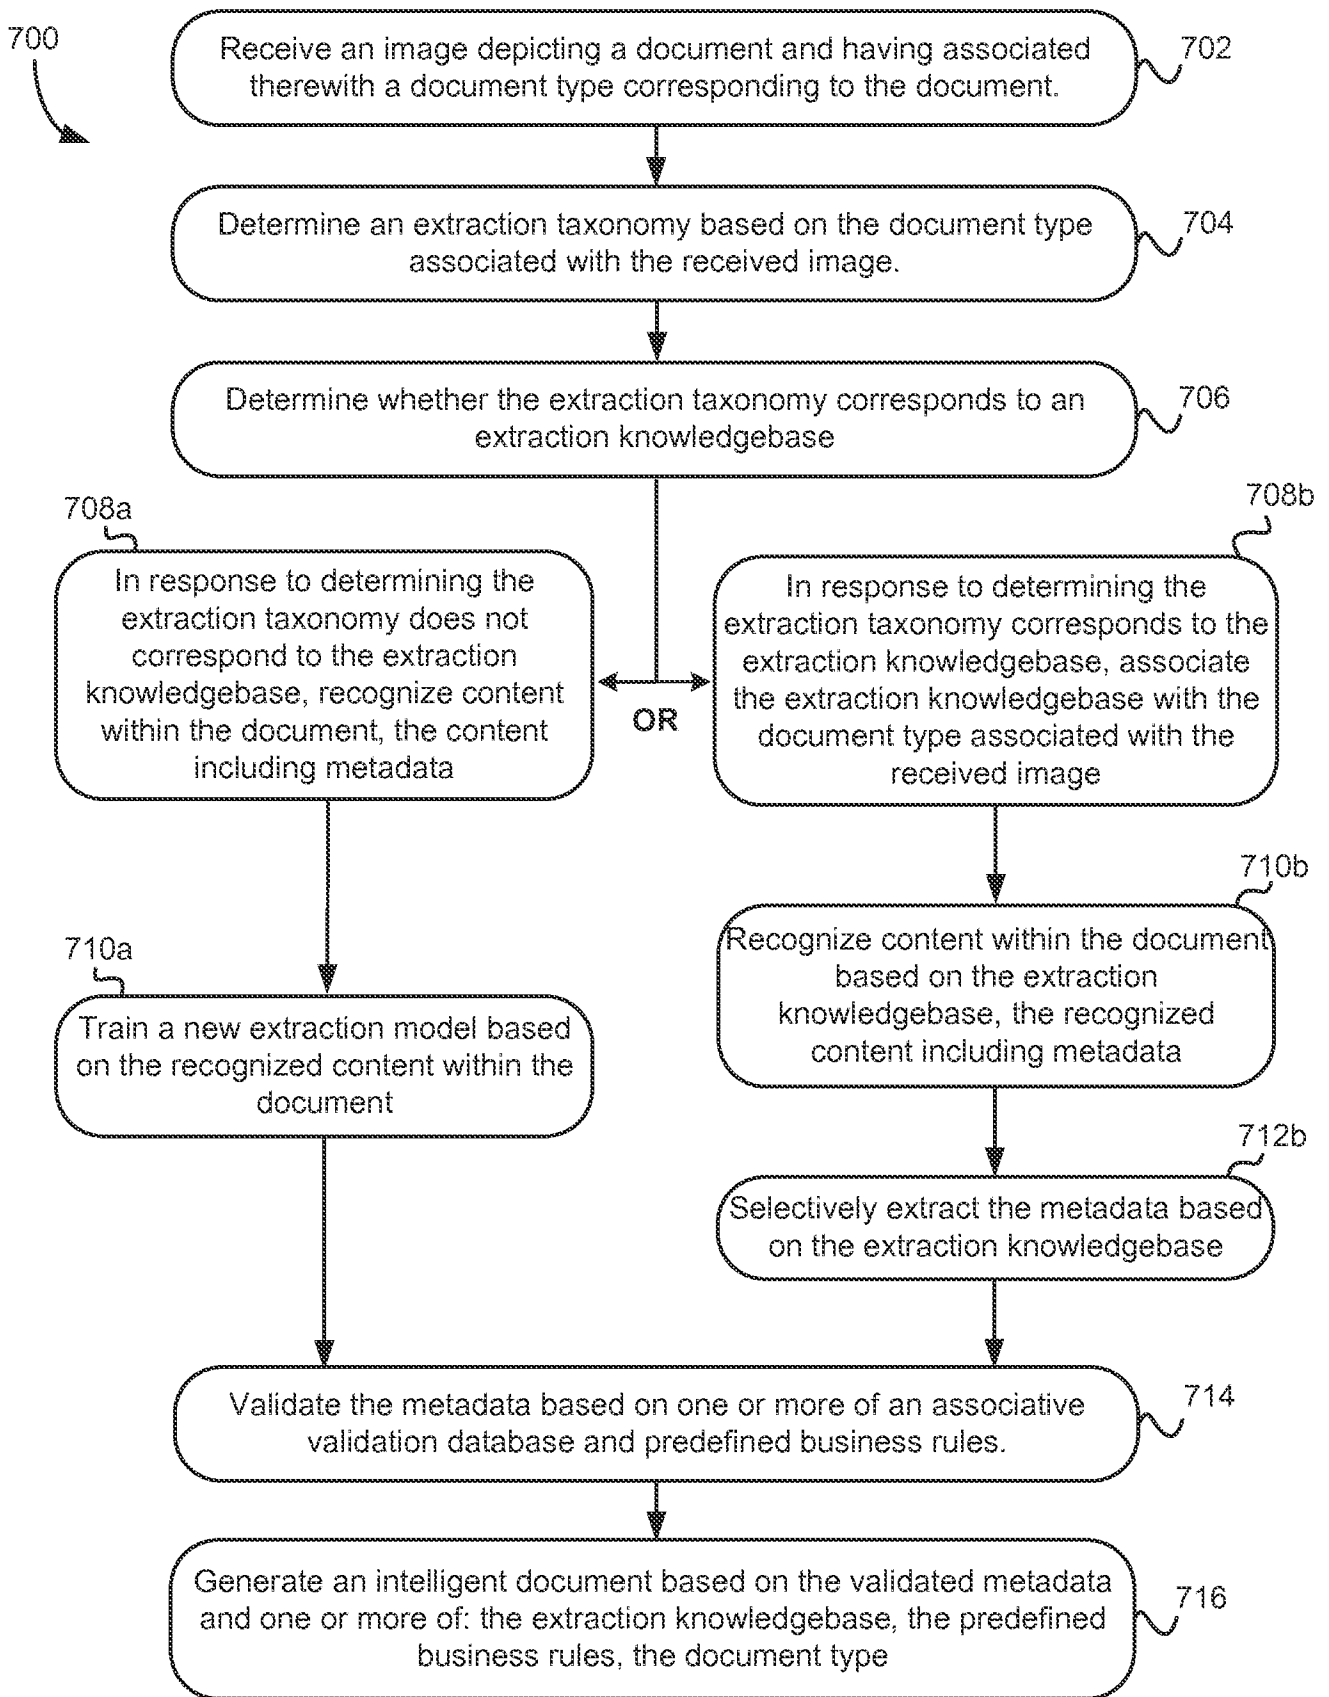


FIG. 7

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06K 9/00 (2014.01)

CPC - G06F 3/033 (2014.11)

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8) - G06F 15/00; G06K 9/00, 15/00, 15/02; H04N 1/00, 1/40, 1/46, 1/387,1/393 (2014.01)

USPC - 358/100, 103, 107, 112, 117,118, 284, 312; 382/1.2, 1.5, 1.9, 2.1, 3.11, 3.12, 302, 400, 401, 406, 448, 451, 488, 504

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
CPC - (2014.11) G06F 3/033; G06K 9/36, 9/46; H04N 5/228 (keyword delimited)

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

Orbit, Google Patents, Google Scholar, Patentlens.net.

Search terms used: Mobile device, video data, motion vectors, detected document, composite images, frame, digital video data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	Showering (US 2011/0285873 A1) 24 Novembre 2011 (24.11.2011) entire document	1-21
A	US 2008/0004073 A1 (John et al.) 03 January 2008 (03.01.2008) entire document	1-21
A	US 2011/0025842 A1 (King et al.) 03 February 2011 (03.02.2011) entire document	1-21
A	US 2008/0166025 A1 (Thorne) 10 July 2008 (10.07.2008) entire document	1-21
A	US 2010/0007751 A1 (Icho et al.) 14 January 2010 (14.01.2010) entire document	1-21
A	US 2010/0214584 A1 (Takahashi) 26 August 2010 (26.08.2010) entire document	1-21
A	US 2006/0164682 A1 (Lev) 27 July 2006 (27.07.2006) entire document	1-21
A	US 2011/0200107 A1 (Ryu) 18 August 2011 (18.08.2011) entire document	1-21

 Further documents are listed in the continuation of Box C.

* 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

26 January 2015

Date of mailing of the international search report

26 FEB 2015

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

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