



(51) International Patent Classification:

G06N 3/02 (2006.01) G06T 7/30 (2017.01)

(21) International Application Number:

PCT/US2018/049484

(22) International Filing Date:

05 September 2018 (05.09.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 10300 Energy Drive, Spring, Texas 77389 (US).

(72) Inventors: **CHEN, Eli**; Einstein 10, Kiryat Weizmann, 76101 Ness Ziona (IL). **HAIK, Oren**; Einstein 10, Kiryat Weizmann, 76101 Ness Ziona (IL). **PERRY, Oded**; Einstein 10, Kiryat Weizmann, 76101 Ness Ziona (IL).

(74) Agent: **PERRY, Garry A.** et al.; HP Inc., Mail Stop 35, 3390 E. Harmony Road, Fort Collins, Colorado 80528 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,

KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

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

(54) Title: OPTICAL FLOW MAPS




Fig. 1

(57) Abstract: In an example, a method includes receiving a first version and a second version of an image, wherein the orientation of the first version is different from the orientation of the second version. The first and second versions of the image may be provided to a convolutional neural network, which may calculate an optical flow map representing an optical flow between the first and second versions of the image. The method may further comprise smoothing at least one discontinuity in the optical flow map and applying the smoothed optical flow map to the first version of the image to adjust the orientation of the first version of the image to match the orientation of the second version of the image.



OPTICAL FLOW MAPSBACKGROUND

5 **[0001]** In printing, print agents such as inks, toners, coatings and the like (generally, 'print agents') may be applied to a substrate. Substrates may in principle comprise any material, for example paper, card, plastics, fabrics or the like.

10 **[0002]** In some examples, the resulting print may be analysed in order to identify potential or actual defects. In some examples, a printed substrate is scanned, and the captured image is compared to a reference image, for example an image which formed the basis of a print instruction, or a previously printed image which has been determined to meet certain criteria.

BRIEF DESCRIPTION OF DRAWINGS

15 **[0003]** Non-limiting examples will now be described with reference to the accompanying drawings, in which:

[0004] Figure 1 is a flowchart of an example method of adjusting an orientation of an image;

20 **[0005]** Figure 2 is another example of a flowchart of an example method of adjusting an orientation of an image;

[0006] Figure 3 shows example optical flow maps;

[0007] Figure 4 is a flowchart of an example method which comprises part of the method of Figure 2.

25 **[0008]** Figure 5 is an example of a machine readable medium in association with a processor; and

[0009] Figure 6 is a diagram of an example print apparatus.

DETAILED DESCRIPTION

5 [0010] When comparing two images (for example, a scanned image of a printed substrate and a reference image) to detect a defect in one of the images (e.g. the printed image), an alignment or register process may take place so that the images have the same orientation, before defect detection takes place. This can reduce false alarms in relation to defect detection which may occur when an apparent defect is actually the result of a registration error. Such false alarms due to mis-registration may be more common where images include, 10 for example, small text and/or a low amount of data.

[0011] Figure 1 is an example of a method, which may be a method for correcting for relative rotation between two versions of an image. The method may in some examples be a computer implemented method.

15 [0012] Block 102 comprises receiving a first version and a second version of an image, wherein the orientation of the first version is different from the orientation of the second version. In some examples, the first version of the image may be, or may be derived from, a digital reference version of an image and the second version of the image may be, or may be derived from, a scan of a printed version of the same image which has been printed by a print apparatus. In some examples, the digital reference image may be based on an 20 image used to provide print instructions to print the printed version of the image. In such examples, the method may be used as part of a method to compare the first and second versions of the image in order to identify defects in the printed version of the image. The versions of the image may for example be 25 received over a network, and/or retrieved from a local memory. In some examples, the first and second versions of the image are received from different sources.

[0013] Block 104 comprises providing the first and second versions of the image to a convolutional neural network (CNN). In some examples, the CNN 30 is a CNN which is configured, for example, trained, to calculate or estimate optical flow. In some examples, more than one CNN may be provided, as described in greater detail below.

[0014] Examples of CNNs include those which extract features from an image by convolving the image with a plurality of filters to produce a feature map. In

5 some examples, after the feature map is produced, a non-linear function may be applied to the feature map to form a rectified feature map. In some examples, a pooling operation may be applied to the feature map (such as Max. Pooling) which reduces the dimensionality of the feature map whilst retaining the relevant information from the feature map. The convolutional neural network may perform such operations (i.e. the convolution, applying a non-linear function and/or pooling operations) a number of times and may produce a number of rectified feature maps for an image. The convolutional neural network may then classify the image, based on a training dataset and using the rectified feature maps and a fully connected layer.

[0015] An example of a CNN for estimating optical flow (i.e. a relative shift of at least one feature between two images) is FlowNet 2.0.

15 [0016] In some examples, a coarse rotation correction may be applied to either or both of the first and second versions of the image before providing the first and second versions of the image to the CNN, for example by using a classical computer vision approach such as a block matching technique which divides a reference image into a matrix of 'macro blocks', i.e. regions of the image which have dimensions of a few pixels square, which are then each compared with a corresponding block and its adjacent neighbors in a target frame to create a vector that describes the movement of a macro block from one location to another in the target frame. The movement calculated for all the macro blocks that make up the image then constitutes a motion estimated in the reference image, and this motion may be used to determine the coarse rotation correction. In one example, a coarse rotation correction may rotate the first and/or second versions of the image so that the difference in angle between the two images is within around ± 2 degrees. This may largely correct a difference due to a paper jam or the like, which has caused a significant skew.

25 [0017] In some examples, a region of interest may be selected in one or both of the first and second image versions and one or both images may be cropped to the region of interest. In some examples, this may be carried out such that the first and second image versions have the same dimensions. In other examples, at least one (or at least a region of interest within) the first and/or second image versions may be scaled such that the first and second image versions have the same dimensions before being provided to the CNN. In some

30

examples, the size of the first and second images may be scaled, for example reduced, before being provided to the CNN. In some examples, the first and second image versions may be scaled down by a factor of between 1.5 and 3.0. This can increase the efficiency of processing, as the number of computations performed by the CNN may be related to image size. In some examples the first and second image versions may be scaled down by a factor of two before being provided to the CNN.

[0018] Block 106 comprises calculating, by the CNN, an optical flow map representing an optical flow between the first and second versions of the image. In some examples, calculating the optical flow map may firstly comprise a contracting stage in which the first and second versions of the image are input into the CNN and a high level feature representation is extracted. Then an expanding stage may take place in which the optical flow map is generated, using 'up-convolutional layers' and features obtained in the contracting stage. In the expanding stage, the high level feature representation is upsampled in order to produce an optical flow map with the same resolution as the input images. The CNN may have been trained end to end with training data. For example, the CNN may be based on FlowNet 2.0 trained using a suitable dataset such as the 'Flying Chairs' dataset to carry out optical flow estimation, prior to performing method 100. Other examples of datasets used for training CNNs to carry out optical flow estimations include the Sintel and KITTI data sets.

[0019] In this way, the CNN effectively compares the first and second versions of the image and produces an optical flow map having the same dimensions as the image, in which each pixel of the optical flow map is a vector which represents a relative shift in position of a corresponding pixel of the image between the two versions.

[0020] Block 108 comprises smoothing at least one discontinuity in the optical flow map. Because the optical flow map represents a relative rotation between the two images, it is expected that the optical flow map should be substantially continuous (i.e. it is expected that neighboring pixels will have moved by a similar amount so their optical flow will be a similar value). Therefore, at block 108 'post-processing' techniques are applied to smooth out any local discontinuities that do not fit with the expected relative rotation behavior.

[0021] In some examples, block 108 may comprise detecting contours in the optical flow map with high contrast in relation to their environment (for example contours indicating different predetermined pixel counts which are spaced by less than a threshold spacing – i.e. ‘close contours’). In some examples, block 108 may comprise replacing the value of each pixel of the discontinuity with a value based on a value of a neighboring pixel, for example using a “flood fill” method. In some examples, block 108 may comprise applying a median filter to the optical flow map. For example, a median filter with a kernel size of 5x5 pixels may be used. In some examples, block 108 may comprise smoothing the optical flow map using a Gaussian blur. In some examples, block 108 may comprise cropping a margin of the optical flow map and replicating pixel values of pixels bordering the margin to rebuild the margin so that the size of the optical flow map remains consistent but any border discontinuities are removed. In other examples, smoothing may be carried out in some other way and/or using combinations of the techniques described above.

[0022] Smoothing discontinuities in the optical flow map enables outlying pixels which are due to errors to be removed without removing features in the optical flow map that genuinely represent an optical flow caused by a relative rotation between the two images. Smoothing the optical flow map may also reduce differences in noise and/or brightness change between the first and second versions of the image, which can enable more accurate comparison of two image versions.

[0023] In examples where the first and second versions of the image have been reduced by a scaling factor before being provided to the CNN, the optical flow map may be scaled up such that the smoothed optical flow map is the same size as the original first version of the image.

[0024] Block 110 comprises applying the smoothed optical flow map to the first version of the image to adjust the orientation of the first version of the image to match the orientation of the second version of the image. The optical flow map comprises a 2D array of vectors, wherein each vector corresponds to a pixel of the first version of the image and represents the optical flow of that pixel in the x and y directions. Therefore, moving each pixel in the first version of the image by the amount represented by the corresponding vector of the optical flow map corrects for any relative rotation or skew that may be present

between the first and second versions of the image, thereby registering the first version of the image with the second version of the image. In some examples, block 110 comprises applying the smoothed optical flow map to a digital reference image which causes the digital reference image to be rotated so that it is at the same angle as a scan of the printed image (or a scan of the printed image to which a coarse rotation has been applied in some examples).

[0025] The method of Figure 1 enables two versions of an image that are rotated relative to each other to be aligned with a high degree of accuracy (in some examples, following a coarse rotation). This can enable an accurate comparison between the two versions of the image (for example a reference image and a scan of a printed image) for accurate detection of defects in one of the images (for example the printed image).

[0026] Figure 2 is a further example of a method 200 for correcting for relative rotation between two versions of an image.

[0027] Block 202 comprises receiving a first version and a second version of an image, wherein the orientation of the first version is different from the orientation of the second version. In some examples, the first version of the image may be, for example, a digital reference image and the second version of the image may be a scan of a printed version of the image.

[0028] Block 204 comprises detecting the corners of the second version of the image, followed by cropping the second version of the image and rotating it to reduce the relative angle between the second and the first image, for example such that it is within ± 2 degrees. In other words, block 204 may comprises a 'coarse' rotational correction.

[0029] Block 206 comprises selecting a region of interest in the first image version that has an equal size to the second image version and cropping the first image version to the region of interest.

[0030] Block 208 comprises reducing the size of the first and second image versions by scaling the images down by a factor of between 1.5 and 3. In some examples, block 208 comprises scaling the first and second images down by a factor of 2. Scaling the images down before providing them to the CNN may improve the runtime of the rotation correction and/or reduce processing

resources utilized by reducing the number of computations to be performed when compared to the full sized image.

5 **[0031]** Block 210 comprises providing versions of the first and second image to a convolutional neural network. In some examples, the convolutional neural network may be configured to calculate or estimate optical flow.

[0032] Block 212 comprises calculating, by the convolutional neural network, an optical flow map representing a skew or relative rotation between the first and second versions of the image.

10 **[0033]** Block 214 comprises detecting discontinuities in the optical flow map by finding close contours in the optical flow map and/or any feature which has high contrast in relation to its environment. In some examples, contours may be defined as lines within the optical flow map of constant pixel offset. In some examples a discontinuity may be defined as a pixel or area of pixels which has a contrast in comparison to a neighboring area of pixels above a predetermined threshold. In this way, at block 214 the method detects artefacts or outlier pixels in the optical flow map which do not fit with an expected representation of a rotation. In other examples, such artefacts or outlier pixels may be detected in some other way.

20 **[0034]** Block 216 comprises applying post-processing techniques on the optical flow map in order to remove the artefacts detected at block 214 by smoothing the optical flow map. In the method shown in Figure 2, smoothing the optical flow map comprises filling detected outlier pixels and/or areas of outlier pixels with the values of neighboring pixels using a flood fill method. In some examples, a median filter may then be applied to the optical flow map, which in some examples may have a kernel size of between 2x2 pixels and 10x10 pixels. In some examples the kernel size may be 5x5 pixels. Block 216 may also comprise, after applying the median filter, applying a Gaussian blur to the optical flow map to further smooth out any remaining areas of high contrast. In this example, block 218 comprises cropping at least one of the horizontal and vertical margins of the optical flow map and replicating pixel values of pixels neighboring the cropped border to replace or rebuild the border so that the overall size of the optical flow map remains unchanged.

30 **[0035]** Block 220 comprises resizing the optical flow map by scaling up the optical flow map such that the optical flow map matches the size of the first

version of the image prior to block 208, i.e. reversing the scaling factor that was applied to the first and second versions of the image at block 208.

[0036] Block 222 comprises applying the smoothed optical flow map to the first version of the image.

5 **[0037]** The method 200 results in a significant reduction in mis-registration errors when comparing two images. In some examples, with appropriate processing resources, the method may be performed at a speed of 10 frames per second (fps) for image sizes in the order of 1250 x 500 pixels.

10 **[0038]** Figure 3 shows a visual representation of an x-axis optical flow 302 and a visual representation of a y-axis optical flow 304 from an optical flow map representing a relative rotation between two images. As discussed above, an optical flow map that accurately represents a relative rotation between two images will have a smooth gradient. Therefore it is possible to improve the accuracy of an optical flow map where the optical flow map is expected to represent a relative rotation in orientation between two images by removing at least one feature in the optical flow map which is inconsistent with a representation of a relative rotation.

[0039] Figure 4 shows an example method of carrying out block 212 of the method of Figure 2.

20 **[0040]** This example combines two approaches that can be used to determine optical flow using a CNN. A first approach (a complex CNN) comprises, in the contracting stage, providing two separate but identical processing streams to extract features from each image version separately and then correlate the features of the two resultant high level feature representations at different locations in the image versions before calculating the optical flow. A second approach (a simple CNN) begins, in the contracting stage, by stacking both input image versions together and then processing the stacked images jointly by feeding the stacked images through a generic CNN, allowing the CNN to decide how to process the image pair to extract the optical flow or 'motion' information. In some examples, one of these two different approaches may be used. However, in other examples, one of which is described with reference to Figure 4, these two approaches (simple and complex) may be stacked in different combinations to provide the CNN architecture.

25

30

- [0041] In the example shown in Figure 4, the method comprises, at block 402, inputting first and second versions of an image to a first CNN, which may be a complex CNN. The first CNN outputs an estimated optical flow map. The second version of the image is then warped with the estimated optical flow map (block 404) and the resultant warped image is then input to a second CNN, which may be a simple CNN, along with the first version of the image (block 406). Next, the second CNN outputs a second estimated optical flow map and the second version of the image is warped with the second estimated optical flow map (i.e. the method loops back to block 404). The resultant second warped image is then input into a third CNN, which may be a simple CNN, along with the first version of the image. The resultant optical flow which is output by the third CNN in block 408 may then be passed to block 214 of method 200 for post-processing.
- [0042] Using the particular stacked arrangement shown in Figure 4 has been found to result in the production of a high quality optical flow map.
- [0043] Figure 5 is an example of a tangible machine readable medium 500 in association with a processor 502. The machine readable medium comprises non-transitory instructions 504 which, when executed by the processor 502 cause the processor 502 to carry out processes on receipt of a first version and a second version of an image, wherein the orientation of the first version is different from the orientation of the second version. The instructions 504 include instructions 506 to cause the processor 502 to calculate, using a convolutional neural network an optical flow map providing an initial spatial mapping between the first and second versions of the image. The instructions 504 further comprise instructions 508 to cause the processor 502 to, modify the optical flow map to remove at least one feature in the optical flow map which is inconsistent with a representation of a relative rotation in orientation between the first and second versions of the image. The instructions 504 further comprise instructions 510 to cause the processor 502 to apply the modified optical flow map to the first version of the image to adjust the orientation of the first version of the image to match the orientation of the second version of the image. In some examples, the instructions 504 may carry out at least one block of the method of Figure 1, 2 or 3.

- 5 **[0044]** Figure 6 is an example of a print apparatus 600 comprising a scanning apparatus 602 to scan a printed image, processing circuitry 604 and a CNN 606. In use of the apparatus 600, the processing circuitry 604 receives a reference version of the image and provides an image based on the scanned version of the image and an image based on the reference version of the image to the convolutional neural network 606. In some examples, the images provided to the convolutional neural network 606 may be (i) a scanned version of the image and (ii) the reference version of the image respectively. In other examples, at least one of the images may be at least one of rotated (for example, a 'coarse' rotation may be applied as described above), cropped, scaled or the like before being provided to the convolutional neural network 606. In some examples, a plurality of convolutional neural network 606 may be provided.
- 10 **[0045]** The convolutional neural network 606 calculates an optical flow map representing an optical flow between the scanned and reference versions of the printed image. The processing circuitry 604 then removes features in the optical flow map which are inconsistent with a representation of relative rotation in orientation between the scanned version of the image and the reference version of the image by applying smoothing to discontinuities in the optical flow map, and applies the optical flow map to the reference version of the image to adjust the orientation of the reference version of the image to match the orientation of the scanned version of the image.
- 15 **[0046]** In some examples, the print apparatus 600 may be an integrated apparatus, i.e. the scanning apparatus 602 may be provided at an output of a printer, and be integral thereto (for example being mechanically fastened to and/or aligned therewith). However the scanning apparatus 602, processing circuitry 604 and CNN 606 could be remote from one another. The CNN 606 may be part of the processing circuitry 604 in some examples.
- 20 **[0047]** In some examples, the print apparatus 600 is a Liquid Electro Photographic (LEP) printing apparatus which may be used to print a print agent such as an electrostatic ink composition (or more generally, an electronic ink). In other examples, the print apparatus may comprise an ink-jet printer, a laser (or LED) printer, a solid ink printer, an offset printer, or some other print apparatus type.
- 25
- 30

[0048] Aspects of some examples in the present disclosure can be provided as methods, systems or machine readable instructions, such as any combination of software, hardware, firmware or the like. Such machine readable instructions may be included on a computer readable storage medium (including but not limited to disc storage, CD-ROM, optical storage, etc.) having computer readable program codes therein or thereon.

[0049] The present disclosure is described with reference to flow charts and block diagrams of the method, devices and systems according to examples of the present disclosure. Although the flow diagrams described above show a specific order of execution, the order of execution may differ from that which is depicted. Blocks described in relation to one flow chart may be combined with those of another flow chart. It shall be understood that at least one block in the flow charts, as well as combinations of the blocks in the flow charts can be realized by machine readable instructions.

[0050] The machine readable instructions may, for example, be executed by a general purpose computer, a special purpose computer, an embedded processor or processors of other programmable data processing devices to realize the functions described in the description and diagrams, and which may, for example, comprise at least part of the processing circuitry 604. In particular, a processor or processing apparatus may execute the machine readable instructions. Thus functional modules of the apparatus and devices may be implemented by a processor executing machine readable instructions stored in a memory, or a processor operating in accordance with instructions embedded in logic circuitry. The term 'processor' is to be interpreted broadly to include a CPU, processing unit, ASIC, logic unit, or programmable gate array etc. The methods and functional modules may all be performed by a single processor or divided amongst several processors.

[0051] Such machine readable instructions may also be stored in a computer readable storage that can guide the computer or other programmable data processing devices to operate in a specific mode.

[0052] Such machine readable instructions may also be loaded onto a computer or other programmable data processing devices, so that the computer or other programmable data processing devices perform a series of operations to produce computer-implemented processing, thus the instructions

executed on the computer or other programmable devices realize functions specified by block(s) in the flow charts and/or in the block diagrams.

[0053] Further, the teachings herein may be implemented in the form of a computer software product, the computer software product being stored in a storage medium and comprising a plurality of instructions for making a computer device implement the methods recited in the examples of the present disclosure.

[0054] While the method, apparatus and related aspects have been described with reference to certain examples, various modifications, changes, omissions, and substitutions can be made without departing from the spirit of the present disclosure. It is intended, therefore, that the method, apparatus and related aspects be limited only by the scope of the following claims and their equivalents. It should be noted that the above-mentioned examples illustrate rather than limit what is described herein, and that those skilled in the art will be able to design many alternative implementations without departing from the scope of the appended claims. Features described in relation to one example may be combined with features of another example.

[0055] The word “comprising” does not exclude the presence of elements other than those listed in a claim, “a” or “an” does not exclude a plurality, and a single processor or other unit may fulfil the functions of several units recited in the claims.

[0056] The features of any dependent claim may be combined with the features of any of the independent claims and/or with any other dependent claim(s).

CLAIMS

1. A method comprising:
 - receiving a first version and a second version of an image, wherein an orientation
 - 5 of the first version is different from an orientation of the second version;
 - providing the first and second versions of the image to a convolutional neural network;
 - calculating, by the convolutional neural network, an optical flow map representing an optical flow between the first and second versions of the image;
 - 10 smoothing a discontinuity in the optical flow map; and
 - applying the smoothed optical flow map to the first version of the image to adjust the orientation of the first version of the image to match the orientation of the second version of the image.
- 15 2. A method according to claim 1 wherein the first version of the image is, or is derived from, a digital reference version of an image and the second version of the image is, or is derived from, a scan of a printed substrate.
3. A method according to claim 1, further comprising: prior to providing the first and
- 20 second versions of the image to the convolutional neural network, detecting at least one corner of the second version of the image and performing a coarse rotation to rotate the second version of the image to reduce a relative rotation between the first and second versions of the image.
- 25 4. A method according to claim 1 further comprising: prior to providing the first and second versions of the image to the convolutional neural network, selecting a region of interest in the first version of the image which has an equal size to the second version of the image and cropping the first version of the image to the region of interest.
- 30 5. A method according to claim 1, further comprising: prior to providing the first and second versions of the image to the convolutional neural network, reducing a size of each of the first and second versions of the image by a scale factor and subsequently increasing a size of the calculated optical flow map by the same factor.

6. A method according to claim 1, wherein smoothing a discontinuity in the optical flow map comprises replacing a value of each pixel of the discontinuity with a value based on a value of a neighboring pixel.
- 5 7. A method according to claim 1, wherein smoothing a discontinuity in the optical flow map comprises applying a median filter to the optical flow map.
8. A method according to claim 1, wherein smoothing a discontinuity in the optical flow map comprises smoothing the optical flow map using a Gaussian blur.
- 10 9. A method according to claim 1, wherein smoothing a discontinuity in the optical flow map comprises cropping a margin of the optical flow map and replicating pixel values of pixels bordering the margin to rebuild the margin without discontinuities.
- 15 10. A method according to claim 9, wherein calculating the optical flow map comprises:
passing the first version and the second version of the image through a first convolutional neural network to produce a first optical flow map;
warping the second version of the image with the first optical flow map; and
20 passing the warped second version of the image and the first version of the image through a second convolutional neural network to produce a second optical flow map.
- 25 11. A method according to claim 10, wherein calculating the optical flow map further comprises:
warping the second version of the image with the second optical flow map; and
passing the second version of the image which has been warped with the second optical flow map through a third convolutional neural network to produce a third optical flow map.
- 30 12. A machine-readable medium storing instructions which when executed by a processor cause the processor to, on receipt of a first version and a second version of an image, wherein an orientation of the first version is different from an orientation of the second version:

calculate, using a convolutional neural network, an optical flow map providing an initial spatial mapping between the first and second versions of the image;

modify the optical flow map to remove at least one feature in the optical flow map which is inconsistent with a representation of a relative rotation in orientation between
5 the first and second versions of the image; and

apply the modified optical flow map to the first version of the image to adjust the orientation of the first version of the image to match the orientation of the second version of the image.

10 13. A machine readable medium according to claim 12 wherein modifying the optical flow map to remove at least one feature in the optical flow map which is inconsistent with a representation of a relative rotation in orientation between the first and second versions of the image comprises detecting close contours in the optical flow map.

15 14. A machine readable medium according to claim 13 wherein modifying the optical flow map to remove at least one feature in the optical flow map which is inconsistent with a representation of a relative rotation in orientation between the first and second versions of the image further comprises:

replacing a value of each pixel of the feature with a value based on a value of a
20 neighboring pixel; and/or
applying a median filter to the optical flow map.

15. A print apparatus comprising:
scanning apparatus to scan an image printed by the print apparatus and to
25 provide a scanned version of the printed image;
a convolutional neural network; and
a processor, the processor to:
receive a reference version of the image;
provide images based on the scanned version of the image and the reference
30 version of the image to the convolutional neural network;

wherein the convolutional neural network is to calculate an optical flow map representing an optical flow between images based on the scanned and reference versions of the printed image; and

the processor is to:

remove features in the optical flow map which are inconsistent with a representation of relative rotation in orientation between the scanned version of the image and the reference version of the image by applying smoothing to discontinuities in the optical flow map; and

- 5 apply the optical flow map to the reference version of the image to adjust the orientation of the reference version of the image to match the orientation of the scanned version of the image.

1/6




Fig. 1

2/6




Fig. 2

3/6




Fig. 3

4/6




Fig. 4

5/6




Fig. 5

6/6




Fig. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2018/049484

A. CLASSIFICATION OF SUBJECT MATTER		
G06N 3/02 (2006.01) G06T 7/30 (2017.01) According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols)		
G06N 3/00, 3/02, 3/04, 3/08, G06T 5/00, 7/00, 7/30, 11/00, G06K 9/00, 9/03, 9/62 - 9/66, 9/68, A61B 6/00		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
PatSearch (RUPTO Internal), USPTO, PAJ, Espacenet, Information Retrieval System of FIPS		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 8337405 B2 (RAYTHEON COMPANY) 25.12.2012	1-15
A	CN 104504362 A (NANJING AIKELESI NETWORK TECHNOLOGY CO., LTD.) 08.04.2015	1-15
A	US 2018/0121767 A1 (ADOBE SYSTEMS INCORPORATED) 03.05.2018	1-15
A	WO 2014/133597 A1 (SPINELLA IP HOLDINGS, INC.) 04.09.2014	1-15
☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.		
* Special categories of cited documents: “A” document defining the general state of the art which is not considered to be of particular relevance “E” earlier document 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		Date of mailing of the international search report
17 May 2019 (17.05.2019)		30 May 2019 (30.05.2019)
Name and mailing address of the ISA/RU: Federal Institute of Industrial Property, Berezhkovskaya nab., 30-1, Moscow, G-59, GSP-3, Russia, 125993 Facsimile No: (8-495) 531-63-18, (8-499) 243-33-37		Authorized officer D. Starshinov Telephone No. 8(495) 531-64-81