



(51) International Patent Classification:

G06T 5/00 (2006.01)

G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2020/062125

(22) International Filing Date:

25 November 2020 (25.11.2020)

(25) Filing Language:

English

(26) Publication Language:

English

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

(72) Inventors: **LIN, Qian**; 1501 Page Mill Rd, Palo Alto, California 94304-1100 (US). **NEVES, Deangeli**; Av. Alan Turing, 275, 13083-898 Campinas (BR). **VIANA, Thar-sis**; Av. Alan Turing, 275, 13083-898 Campinas (BR). **VALENTE, Augusto**; Av. Alan Turing, 275, 13083-898 Campinas (BR).

(74) Agent: **JENNEY, Michael** et al.; HP Inc., Mail Stop 35, 3390 E Harmony Rd, 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, IT, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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 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: IDENTIFYING ANOMALY LOCATION

(57) Abstract: A region of interest is extracted from a captured image of a physical object. An autoencoder model is applied to the extracted region of interest to reconstruct the region of interest. A location of an anomaly of the physical object within the extracted region of interest, if any, is identified based on the extracted and reconstructed regions of interest.



IDENTIFYING ANOMALY LOCATION

BACKGROUND

[0001] Modern-day electronic and other types of devices are often manufactured by assembling multiple components together. A defect or other
5 anomaly within any component, or as a result of attaching components together, can result in failure of the overall device. Therefore, constituent components of a device are often tested for anomalies before inclusion in the device, and after manufacture the device itself is also tested.

BRIEF DESCRIPTION OF THE DRAWINGS

10 [0002] FIG. 1 is a flowchart of an example method for identifying the location of an anomaly within a physical object, if any.

[0003] FIGs. 2A, 2B, 2C, and 2D are diagrams depicting example region of interest (ROI) extraction from a captured image of a physical object, at which the object may have an anomaly.

15 [0004] FIGs. 3A, 3B, 3C, 3D, and 3E are diagrams of example training images for training a segmentation model to extract an ROI from a captured image of a physical object, at which the object may have an anomaly.

[0005] FIG. 4 is a diagram of an example autoencoder model to reconstruct an ROI from an extracted ROI of a captured image of a physical
20 object.

[0006] FIGs. 5A, 5B, 5C, 5D, and 5E are diagrams depicting example anomaly location identification within an extracted ROI of a captured image of a physical object.

[0007] FIG. 6 is a diagram of an example non-transitory computer-readable data storage medium.

[0008] FIG. 7 is a flowchart of an example method.

[0009] FIG. 8 is a diagram of an example computing device.

5

BACKGROUND

[0010] As noted in the background, devices and device components, which are more generally physical objects, are often tested for defects and other anomalies. As one example, an inkjet printhead die may be attached to an inkjet cartridge body during assembly of an inkjet cartridge. As part of this attachment process, a flexible ribbon cable may be secured to the printhead die and encapsulated with a material at the connection point between the cable and the die to form what is known as an encap. The encap protects the electrical connection between the cable and the printhead die from ink during usage of the inkjet cartridge in a printing device to eject ink.

10 [0011] One way by which components may be tested for anomalies like defects is to use machine vision, due to the large number of components that have to be inspected and the small size of potential anomalies. Machine vision involves capturing an image of a component and then performing image processing to detect whether the component has any anomalies. For example, with respect to the encap of an inkjet cartridge, machine vision may be used to detect defects in the encap that may permit ink to reach the electrical connections between the cable and the printhead die, which can result in the failure of the cartridge or even the printing device of which it is a part.

15 20

[0012] Techniques described herein permit anomalies of physical objects to be detected, and their locations on the objects to be identified. A region of interest (ROI) is extracted from a captured image of a physical object, and an autoencoder model is applied to the extracted ROI to
5 reconstruct the ROI. The location of an anomaly of the physical object, if any, is identified within the extracted ROI based on the extracted and reconstructed ROIs.

[0013] FIG. 1 shows an example method 100 for identifying the location of an anomaly within a physical object, if the object has any such anomaly.
10 The method 100 can be implemented as program code stored on a non-transitory computer-readable data storage medium and executed by a processor of a computing device. The computing device may be a computer that is communicatively connected to an image sensor that captures images of physical objects as or after they are manufactured, for instance, to identify
15 any objects having defects or other anomalies.

[0014] Physical objects detected as having anomalies may be separated from anomaly-free objects so that they are discarded and not subsequently used, or so that they can be inspected at the locations of the detected anomalies for potential repair or manufacturing process analysis.
20 For example, there may be an expected defect rate during the manufacture of physical objects. If the defect rate suddenly increases or exceeds a threshold, the manufacturing process may be analyzed and modified in an automated or manual manner to return the defect rate back to normal.

[0015] Therefore, the method 100 provides for technological
25 improvements in a number of different ways. For instance, the method 100

can ensure that manufactured physical objects are less likely to include defects and other anomalies. The method 100 can similarly ensure that subsequently manufactured devices are less likely to be defective due to the inclusion of defective constituent components. The method 100 can also

5 improve the manufacturing process of physical objects, by providing an automated or manual feedback loop as to the quality of the objects being manufactured. The method 100 has proven to be more accurate and faster than competing machine vision-based defect identification techniques.

[0016] The method 100 includes extracting an ROI from a captured

10 image of a physical object (102). The ROI is a prespecified area within the captured image that is to be inspected for potential anomalies of the physical object. The captured image may include other parts of the physical object that are not of interest, for instance, and may similarly include background regions that are not of interest. Therefore, the ROI is extracted from the

15 captured image to focus the part of the image that is subsequently analyzed.

[0017] The ROI may be extracted from the captured image on one of two ways, by performing parts 104A and 104B, collectively referred to as the parts 104, or by performing parts 106A and 106B, collectively referred to as the parts 106. As to the former, the method 100 can include aligning the

20 captured image against a reference image of another physical object of the same type (104A). For example, a transformation matrix may be calculated that aligns the captured image to the reference image. The method 100 can include then cropping the aligned image based on a bounding box identifying a corresponding ROI within the reference image (104B). For example, an

25 inverse of the calculated transformation matrix may be applied to the

bounding box as defined on the reference image and the captured image then cropped using the inverse-applied bounding box. The cropped aligned image constitutes the ROI.

[0018] As to the second way in which the ROI may be extracted from the captured image, the method 100 can include applying an object segmentation model to the captured image to generate a segmentation mask (106A), and then applying the segmentation mask to the captured image. The object segmentation model may instead output the ROI of the captured image, instead of a segmentation mask that is then applied to the captured image. The object segmentation model may be a regional convolutional neural network (R-CNN) trained using training images of physical objects of the same type as the physical object of the captured image, and on which corresponding ROIs have been preidentified. When an objection segmentation model is used to extract the ROI from the captured image, no reference image has to be provided, as compared to the first way in which the ROI may be extracted.

[0019] FIGs. 2A-2D show example performance of the ROI extraction of part 102. As to the first way in which the ROI may be extracted from a captured image, FIG. 2A shows an example captured image 200 of an inkjet cartridge, including an encap, and FIG. 2B shows an example reference image 210 of an inkjet cartridge of the same type (e.g., of the same model, and which may be manufactured in the same way and/or at the same facility), and also including an encap. A preidentified bounding box 212 encapsulating the encap of the cartridge within the reference image 210 is depicted in FIG.

2B as well. The encaps of the images 200 and 210 differ in that the former includes a defect or anomaly at location 202, whereas the latter does not.

[0020] A transformation matrix can be calculated that aligns the captured image 200 to the reference image 210, and an inverse of the transformation matrix applied to the bounding box 212 as preidentified within the reference image 210. The captured image 200 can then be cropped in correspondence with resulting inverse-applied bounding box 212 to extract the ROI from the captured image 200. FIG. 2C shows an example extracted ROI 220 of the captured image 200, which includes the encap of the inkjet cartridge within the captured image 200, in correspondence with the bounding box 212 preidentified within the reference image 210. The extracted ROI 220 includes the location 222 of the anomaly of the encap, in correspondence with the location 202 of the captured image 200.

[0021] As to the second way in which the ROI may be extracted from a captured image, an object segmentation model may instead be applied to the captured image 200 of FIG. 2A. FIG. 2D shows an example segmentation mask 230 that may result from application of the segmentation model to the captured image 200, where a white region 232 indicates the ROI. The segmentation mask 230 can then be applied to the captured image 200 to extract the ROI by effectively cropping the captured image 200 in correspondence with the mask 230. The resulting extracted ROI is the same as the ROI 220 of FIG. 2C, except that the ROI is not rectangular and instead conforms to the shape of the segmentation mask 230. The reference image 210 of FIG. 2B is not used when a segmentation model is applied to the captured image 200 to extract the ROI.

[0022] FIGs. 3A, 3B, 3C, 3D, and 3E show portions of different example training images 310, 320, 330, 340, and 350, respectively, for training an object segmentation model for extracting an ROI from a captured image. The training images 310, 320, 330, 340, and 350 specifically show portions of encaps of inkjet print cartridges. The training images 310, 320, 330, 340, and 350 have preidentified ROIs 312, 322, 332, 342, and 352, respectively, that are used during model training for the generation of corresponding segmentation masks. An object segmentation model can better identify which parts of an encap should be included within an ROI, as opposed to including the entirety of the ROI as is the case when a reference image is instead used for ROI extraction from a captured image.

[0023] For instance, the preidentified ROI 312 of FIG. 3A excludes a spurious region 314 at the lower left of the encap in the image 310, which is not considered a defect or other anomaly affecting encap quality. Similarly, the preidentified ROI 322 of FIG. 3B excludes a spurious region 324 at the upper right of the encap in the image 320, and the preidentified ROI 332 of FIG. 3C excludes a spurious region 334 at the lower left of the encap in the image 330. The preidentified ROI 342 of FIG. 3D excludes spikes or ridges 344 at the boundary of the encap in the image 340, which are not considered defects or other anomalies affecting encap quality. Similarly, the preidentified ROI 352 of FIG. 3E excludes porosity 354 at the boundary of the encap in the image 350, which is not considered a defect or other anomaly affecting encap quality. Therefore, the object segmentation model can be trained to more accurately extract an ROI from a captured image, such as via outputting a

segmentation mask that upon application to the captured image results in more accurate ROI extraction.

[0024] Referring back to FIG. 1, the method 100 includes applying an autoencoder model to the extracted ROI to reconstruct the ROI (108). The autoencoder model includes an encoder phase that generates an internal representation from the extracted ROI, and a decoder phase that generates the reconstructed ROI from this internal representation. The autoencoder can be a fully convolutional network without a linear layer, so that the input images do not have to be resized. The autoencoder model is trained using training images of anomaly-free (e.g., defect-free) physical objects of the same type as the physical object of the captured image. Therefore, the reconstructed ROI will largely be identical to the extracted ROI except at locations where the extracted ROI includes defects or other anomalies, because the autoencoder model will not have been trained on any such anomalies. More specifically, the autoencoder model is able to encode and decode high-frequency components of the training images with good precision.

[0025] Corresponding high-frequency components of the extracted ROI will thus be well represented within the reconstructed ROI output by the autoencoder model upon application of the model to the extracted ROI.

However, high-frequency components of the extracted ROI that correspond to defects and other anomalies will not be well represented within the reconstructed ROI output by the autoencoder model. Usage of the autoencoder model to detect anomalies may therefore assume that the anomalies have high frequency – i.e., that the anomalies are not represented as relatively large amorphous regions of slowly changing intensity in the

captured image. Usage of the model may further assume that the captured image has a similar background to that of the training images.

[0026] FIG. 4 shows an example autoencoder model 400 that may be used in part 108. The autoencoder model 400 includes an encoder phase 402 and a decoder phase 404. The encoder phase 402 receives as input an extracted ROI 424 and provides as output an internal representation 425 of the extracted ROI 424. The encoder phase 402 includes a normalization layer 406 that normalizes the extracted ROI 424, followed by multiple layer pairs 408, such as four in the example of FIG. 4. Each layer pair 408 includes a convolutional layer 410 that abstracts its input as well as a rectified linear unit (ReLU) layer 412, which is an activation function that may linearize its input if positive and otherwise output zero. The resulting internal representation 425 is an abstracted encoding of the extracted ROI 424.

[0027] The decoder phase 402 in turn receives as input the internal representation 425 of the extracted ROI 424 and provides as output the reconstructed ROI 426. The reconstructed ROI 426 will faithfully correspond to the extracted ROI 424 at locations at which the ROI 424 does include anomalies like defects, since the autoencoder model 400 is trained on anomaly-free training images. However, at locations at which the extracted ROI 424 includes anomalies, the resulting reconstructed ROI 426 will less likely mirror the extracted ROI 424. The decoder phase 402 includes a number of upsampling convolutional layer groups 414, such as three in the example of FIG. 4. Each group 414 includes an upsampling layer 416 that increases the resolution of its input, followed by a convolutional layer 418, which is more specifically a deconvolutional layer that deabstracts its input.

Each group 414 other than the last group 414 includes a ReLU layer 420, whereas the last group 414 includes a sigmoid layer 422 as a different type of activation layer.

[0028] Referring back to FIG. 1, the method 100 includes identifying the

5 location of an anomaly of the physical object within the extracted ROI of the captured image, if any, based on the extracted and reconstructed ROIs (110).

More generally, there may be zero, one, or multiple anomaly locations, and there may be one or multiple anomalies at each such location. To perform anomaly location identification, the method 100 can include first generating a

10 residual map between the extracted ROI and the reconstructed ROI (112), such as by subtracting the value of each pixel of the reconstructed ROI from the value of the corresponding pixel of the extracted ROI.

[0029] The method 100 can include next removing any pixel of the residual map having a value less than a threshold (114). The threshold may

15 be a static threshold or an adaptive threshold that is based on the residual map itself. As an example of an adaptive threshold, the threshold may be calculated as the mean of the values of the pixels of the residual map, plus a product of a parameter and the standard deviation of the values of the pixels of the residual map. The parameter may be prespecified, and governs the
20 reconstruction error severity that is considered anomalous. For example, a higher parameter value indicates that less severe errors in the reconstructed ROI as compared to the extracted ROI are not considered anomalies, whereas a lower parameter values indicates that such less severe errors are considered anomalies.

[0030] The method 100 can include then applying a morphological operation to the residual map from which pixels having values less than the threshold have been removed (116). The morphological operation that is performed can include an opening morphological operation (118) and/or a closing morphological operation (120). The opening operation, which may also be considered an erosion operation, denoises the residual map of isolated extraneous pixels having values greater than the threshold but which may not in actuality correspond to anomalies. By comparison, the closing operation, which may be considered a dilation operation, connectively groups discontinuous pixels having values greater than the threshold and located near one another, so that they are considered as corresponding to the same anomaly.

[0031] The method 100 can include, if after morphological operation application the residual map includes any pixels having values greater than the threshold (122), determining that the physical object has anomalies like defects (124). The location the anomaly of the anomaly within the extracted ROI corresponds to the locations of the pixels having values greater than the threshold within the residual map. For instance, the location of each group of contiguous pixels having values greater than the threshold within the residual map may be considered as the location of a corresponding anomaly. The method 100 can include, if after morphological operation application the residual map does not include any pixels having values greater than the threshold (124), by comparison determining that the physical object does not have anomalies (126).

[0032] The method 100 can include outputting a segmentation mask that identifies the location of the anomaly of the physical object within the extracted ROI, if any (128). For example, the result of parts 112, 114, and 116 is a post-processed version of the residual map, in which there are white areas and black areas. The white areas correspond to the identified location of the anomaly of the physical object within the extracted ROI of the captured image of the physical object, whereas the black areas do not. This post-processed version of the residual map constitutes the segmentation mask. Therefore, the segmentation mask can be applied to the extracted ROI of the captured image to highlight the location of the anomaly of the physical object, if any.

[0033] FIGs. 5A-5E show example performance of the anomaly location identification of part 110. FIG. 5A shows an example extracted ROI 510 of an encap of an inkjet cartridge having an anomaly at location 512. FIG. 5B shows an example reconstructed ROI 520 corresponding to the extracted ROI 510. The reconstructed ROI 520 is largely identical to the extracted ROI 510, except the reconstructed ROI 520 does not include the anomaly at the location 522 that corresponds to the location 512 in the extracted ROI 510.

[0034] FIG. 5C shows the resulting residual map 530 between the extracted ROI 510 and the reconstructed ROI 520. The residual map 530 is black at locations where the ROIs 510 and 520 are identical. The residual map 530 highlights in white location 532 corresponding to the location 512 of the anomaly within the extracted ROI 510, but is also not black at other locations 534 as well. The locations 534 primarily correspond to the periphery

of the encap at which the autoencoder model did not faithfully reproduce the extracted ROI 510 within the reconstructed ROI 520.

[0035] FIG. 5D shows the resulting residual map 540 after removing pixels having values less than the threshold from the residual map 530. The residual map 540 still highlights in white location 542 corresponding to the location 512 of the anomaly within the extracted ROI 510. However, most of the other non-black locations 534 of the residual map 530 have been removed in the residual map 540, although a few such locations 544 remain. FIG. 5E then shows the resulting residual map 550 after performing a morphological operation on the residual map 540, which removes all or nearly all of the remaining non-black locations other than location 552 that corresponds to the location 512 of the anomaly within the extracted ROI 510. As such, an anomaly is considered as having been detected at the location 552.

[0036] FIG. 6 shows an example non-transitory computer-readable data storage medium 600 storing program code 602. The program code 602 is executable by a processor of a computing device to perform processing. The processing includes extracting a region of interest from a captured image of a physical object (604), and applying an autoencoder model to the extracted region of interest to reconstruct the region of interest (606). The processing includes identifying a location of an anomaly of the physical object within the extracted region of interest, if any, based on the extracted and reconstructed regions of interest (608).

[0037] FIG. 7 is a flowchart of an example method 700. The method 700 may be performed by a processor of a computing device, and may be implemented as program code stored on a non-transitory computer-readable

data storage medium and executed by the processor. The method 700 includes training an autoencoder model using training images of anomaly-free physical objects of a same type (702). The method 700 includes using the autoencoder model to identify a location of an anomaly of a physical object of the same type as the anomaly-free physical objects within an extracted ROI of a captured image of the physical object (704).

[0038] FIG. 8 shows an example computing device 800. The computing device 800 includes a processor 802 and memory 804 storing instructions 806. The instructions 806 are executable by the processor 802 to preprocess an image of a physical object to crop the image (808). The cropped image may be considered the extracted ROI of the image. The instructions 806 are executable by the processor 802 to apply an autoencoder model to the preprocessed image to generate a reconstructed image (810). The instructions 806 are executable by the processor 802 to identify a location of an anomaly of the physical object within the image, if any, based on a residual map between the preprocessed image and the reconstructed image (812).

[0039] Techniques have been described for detecting anomalies of physical objects within captured images and identifying the locations of the anomalies within the captured images. The techniques have been described in relation to encaps of inkjet cartridges as one type of physical object. However, the techniques are applicable to other types of physical objects as well, including other types of manufactured device components. The techniques provide for an automated manner of anomaly detection using an autoencoder model, in which a captured image is preprocessed to extract an

ROI thereof and is postprocessed to better distinguish between noise and actual anomalies like defects.

We claim:

1. A non-transitory computer-readable data storage medium storing program code executable by a processor to perform processing comprising:
extracting a region of interest from a captured image of a physical
5 object;

applying an autoencoder model to the extracted region of interest to reconstruct the region of interest; and

identifying a location of an anomaly of the physical object within the extracted region of interest, if any, based on the extracted and reconstructed
10 regions of interest.

2. The non-transitory computer-readable data storage medium of claim 1, wherein extracting the region of interest comprises:

aligning the captured image of the physical object against a reference image of another physical object of a same type as the physical object; and

15 cropping the aligned captured image based on a bounding box

identifying a corresponding region of interest within the reference image,

wherein the cropped aligned captured image constitutes the region of interest.

3. The non-transitory computer-readable data storage medium of claim 2,

20 wherein aligning the captured image against the reference image comprises calculating a transformation matrix that aligns the captured image to the reference image,

and wherein cropping the aligned captured image based on the

bounding box comprises applying an inverse of the transformation matrix to the bounding box and cropping the captured image using the inverse-applied bounding box.

4. The non-transitory computer-readable data storage medium of claim 1,
5 wherein extracting the region of interest comprises:

applying an object segmentation model to the captured image.

5. The non-transitory computer-readable data storage medium of claim 4,
wherein the object segmentation model is a regional convolutional neural
network (R-CNN) trained using a plurality of training images of other physical
10 objects of a same type as the physical object and on which corresponding
regions of interest have been preidentified.

6. The non-transitory computer-readable data storage medium of claim 1,
wherein the autoencoder model is trained on training images of anomaly-free
physical objects of a same type as the physical object.

- 15 7. The non-transitory computer-readable data storage medium of claim 1,
wherein identifying the location of the anomaly of the physical object within
the extracted region of interest, if any, comprises:

generating a residual map between the extracted region of interest and
the reconstructed region of interest; and

- 20 removing any pixel of the residual map having a value less than a
threshold.

8. The non-transitory computer-readable data storage medium of claim 7, wherein the threshold is a static threshold.

9. The non-transitory computer-readable data storage medium of claim 7, wherein the threshold is an adaptive threshold calculated as a mean of values
5 of pixels of the residual map, plus a product of a parameter and a standard deviation of the values of the pixels of the residual map, the parameter governing reconstruction error severity that is considered anomalous.

10. The non-transitory computer-readable data storage medium of claim 7, wherein identifying the location of the anomaly of the physical object within
10 the extracted region of interest, if any, further comprises:

after removing any pixel having a value less the threshold, applying a morphological operation to the residual map.

11. The non-transitory computer-readable data storage medium of claim 10, wherein applying the morphological operation to the residual map
15 comprises:

applying an opening morphological operation to denoise the residual map of extraneous pixels having values greater than the threshold; and

applying a closing morphological operation to connectively group discontiguous pixels having values greater than the threshold.

20 12. The non-transitory computer-readable data storage medium of claim 7, wherein identifying the location of the anomaly of the physical object within the extracted region of interest, if any, further comprises:

determining that the physical object includes the anomaly if the residual map includes any pixels having values greater than the threshold,

wherein the location of the anomaly of the physical object within the extracted region corresponds to locations of pixels having values greater than
5 the threshold within the residual map.

13. The non-transitory computer-readable data storage medium of claim 12, wherein identifying the location of the anomaly of the physical object within the extracted region of interest, if any, further comprises:

determining that the physical object does not include any anomaly if
10 the residual map does not include any pixels having values greater than the threshold.

14. A method comprising:

training an autoencoder model using a plurality of training images of anomaly-free physical objects of a same type; and

15 using the autoencoder model to identify a location of an anomaly of a physical object of the same type as the anomaly-free physical objects within an extracted region of interest of a captured image of the physical object.

15. A computing device comprising:

a processor; and

20 a memory storing instructions executable by the processor to:

preprocess an image of a physical object to crop the image;

apply an autoencoder model to the preprocessed image to

generate a reconstructed image; and

identify a location of an anomaly of the physical object within the image, if any, based on a residual map between the preprocessed image and the reconstructed image.

1/6




FIG 2A




FIG 2B




FIG 2C




FIG 2D




FIG 3A




FIG 3B




FIG 3C




FIG 3D





FIG 3E






FIG 5A




FIG 5B




FIG 5C




FIG 5D







FIG 5E



6/6

**FIG 6****FIG 7****FIG 8**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2020/062125

A. CLASSIFICATION OF SUBJECT MATTER		
G06T 5/00 (2006.01) G06K 9/00 (2006.01) 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)		
G06T 5/00, 7/00, 7/70, G06N 3/00-3/08, 5/00-5/04, 20/00-20/20, G06K 9/00, 9/36, 9/46, 9/62		
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.
X Y A	US 2020/0364842 A1 (FUJITSU LIMITED) 19.11.2020, paragraphs [0062], [0066], [0068]-[0069], [0074], [0102], [0105], [0107]	1, 4, 6, 14 2-3, 5, 15 7-13
Y	WO 2020/081343 A1 (VENTANA MEDICAL SYSTEMS, INC) 23.04.2020, paragraphs [0090]-[0091], [0094]	2-3, 5, 15
Y	US 2020/0134804 A1 (NEC LABORATORIES AMERICA, INC) 30.04.2020, paragraphs [0005], [0048]	15
Y	CN 107463945 A (UNIV ZHEJIANG) 12.12.2017, paragraphs [0057], [0071]-[0076]	3
A	US 2020/0160505 A1 (BNSF RAILWAY COMPANY) 21.05.2020	1-15
☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.		
* Special categories of cited documents:	"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	
"A" document defining the general state of the art which is not considered to be of particular relevance		
"D" document cited by the applicant in the international application		
"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		
Date of the actual completion of the international search	Date of mailing of the international search report	
26 July 2021 (26.07.2021)	19 August 2021 (19.08.2021)	
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 T. Mukhina Telephone No. 8(495)531-65-15	