



(51) International Patent Classification:
H01L 21/66 (2006.01)

(21) International Application Number:
PCT/US2017/052093

(22) International Filing Date:
18 September 2017 (18.09.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/408,402 14 October 2016 (14.10.2016) US
15/694,719 01 September 2017 (01.09.2017) US

(71) Applicant: **KLA-TENCOR CORPORATION** [US/US];
Legal Department, One Technology Drive, Milpitas, California 95035 (US).

(72) Inventors: **ZHANG, Jing**; 510 Mansion Court, Santa Clara, California 95054 (US). **ROULO, Mark**; 115 Chetwood Drive, Mountain View, California 94043 (US). **BHASKAR, Kris**; 1061 Queensbridge Court, San Jose, California 95120 (US). **DONAPATI, Ravi C.**; 1929 Plymouth Road, Apt. 3006, Ann Arbor, Michigan 48105 (US).

(74) Agent: **MCANDREWS, Kevin** et al.; KLA-Tencor Corp., Legal Department, One Technology Drive, Milpitas, California 95035 (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).

(54) Title: DIAGNOSTIC SYSTEMS AND METHODS FOR DEEP LEARNING MODELS CONFIGURED FOR SEMICONDUCTOR APPLICATIONS




Fig. 1

(57) Abstract: Methods and systems for performing diagnostic functions for a deep learning model are provided. One system includes one or more components executed by one or more computer subsystems. The one or more components include a deep learning model configured for determining information from an image generated for a specimen by an imaging tool. The one or more components also include a diagnostic component configured for determining one or more causal portions of the image that resulted in the information being determined and for performing one or more functions based on the determined one or more causal portions of the image.

Published:

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

DIAGNOSTIC SYSTEMS AND METHODS FOR DEEP LEARNING MODELS CONFIGURED FOR SEMICONDUCTOR APPLICATIONS

BACKGROUND OF THE INVENTION

5

1. Field of the Invention

The present invention generally relates to diagnostic methods and systems for deep learning models configured for semiconductor applications.

10

2. Description of the Related Art

The following description and examples are not admitted to be prior art by virtue of their inclusion in this section.

15

Fabricating semiconductor devices such as logic and memory devices typically includes processing a substrate such as a semiconductor wafer using a large number of semiconductor fabrication processes to form various features and multiple levels of the semiconductor devices. For example, lithography is a semiconductor fabrication process that involves transferring a pattern from a reticle to a resist arranged on a semiconductor wafer. Additional examples of semiconductor fabrication processes include, but are not limited to, chemical-mechanical polishing (CMP), etch, deposition, and ion implantation. Multiple semiconductor devices may be fabricated in an arrangement on a single semiconductor wafer and then separated into individual semiconductor devices.

25

Inspection processes are used at various steps during a semiconductor manufacturing process to detect defects on specimens to drive higher yield in the manufacturing process and thus higher profits. Inspection has always been an important part of fabricating semiconductor devices. However, as the dimensions of semiconductor devices decrease, inspection becomes even more important to the successful manufacture of acceptable semiconductor devices because smaller defects can cause the devices to fail.

30

Defect review typically involves re-detecting defects detected as such by an inspection process and generating additional information about the defects at a higher resolution using either a high magnification optical system or a scanning electron
5 microscope (SEM). Defect review is therefore performed at discrete locations on specimens where defects have been detected by inspection. The higher resolution data for the defects generated by defect review is more suitable for determining attributes of the defects such as profile, roughness, more accurate size information, etc.

10 Advances in deep learning have made deep learning an attractive framework for use in defect detection and classification. While such frameworks have proven useful for defect classification and other functions, the frameworks themselves also make it difficult to know whether or not the deep learning framework is operating correctly. For example, in the case of defect classification, currently used approaches for performing quality
15 assurance on a classifier include classical metrics in machine learning such as accuracy, confusion matrix, and sensitivity on an offline test dataset and online/in-field evaluation. In addition, currently used approaches for performing data augmentation include having a domain expert or an algorithm expert guide the valid and critical transformation/augmentation or performing data augmentation based on physic laws
20 and/or common sense.

There are, however, a number of disadvantages for currently used quality assurance and data augmentation methods and systems. For example, the currently used quality assurance approaches described above cannot identify the situation in which a
25 classifier makes a correct prediction based on wrong causal features, especially for deep learning classifiers. In another example, the currently used quality assurance approaches described above treat machine learning and deep learning algorithms as a black box. In an additional example, the currently used approaches for performing data augmentation cannot be used to directly improve/correct a poorly trained classifier.

30

Accordingly, it would be advantageous to develop systems and methods for performing diagnostic functions for a deep learning model that do not have one or more of the disadvantages described above.

5

SUMMARY OF THE INVENTION

The following description of various embodiments is not to be construed in any way as limiting the subject matter of the appended claims.

10

One embodiment relates to a system configured to perform diagnostic functions for a deep learning model. The system includes one or more computer subsystems and one or more components executed by the one or more computer subsystems. The one or more components includes a deep learning model configured for determining information from an image generated for a specimen by an imaging tool. The one or more

15 components also include a diagnostic component configured for determining one or more causal portions of the image that resulted in the information being determined and for performing one or more functions based on the determined one or more causal portions of the image. The system may be further configured as described herein.

20

An additional embodiment relates to another system configured to perform diagnostic functions for a deep learning model. This system is configured as described above. This system also includes an imaging tool configured for generating images of a specimen. The computer subsystem(s) are, in this embodiment, configured for acquiring the images. This embodiment of the system may be further configured as described

25 herein.

30

Another embodiment relates to a computer-implemented method for performing diagnostic functions for a deep learning model. The method includes determining information from an image generated for a specimen by an imaging tool by inputting the image to a deep learning model. The method also includes determining one or more

causal portions of the image that resulted in the information being determined by inputting the information to a diagnostic component. In addition, the method includes performing one or more functions based on the determined one or more causal portions of the image with the diagnostic component. The deep learning model and the diagnostic component are included in one or more components executed by one or more computer systems.

Each of the steps of the method described above may be further performed as described further herein. In addition, the embodiment of the method described above may include any other step(s) of any other method(s) described herein. Furthermore, the method described above may be performed by any of the systems described herein.

Another embodiment relates to a non-transitory computer-readable medium storing program instructions executable on one or more computer systems for performing a computer-implemented method for performing diagnostic functions for a deep learning model. The computer-implemented method includes the steps of the method described above. The computer-readable medium may be further configured as described herein. The steps of the computer-implemented method may be performed as described further herein. In addition, the computer-implemented method for which the program instructions are executable may include any other step(s) of any other method(s) described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

Further advantages of the present invention will become apparent to those skilled in the art with the benefit of the following detailed description of the preferred embodiments and upon reference to the accompanying drawings in which:

Figs. 1 and 2 are schematic diagrams illustrating side views of embodiments of a system configured as described herein;

Figs. 3 and 4 are flow charts illustrating embodiments of steps that may be performed by systems described herein for performing diagnostic functions for a deep learning model; and

5

Fig. 5 is a block diagram illustrating one embodiment of a non-transitory computer-readable medium storing program instructions for causing one or more computer systems to perform a computer-implemented method described herein.

10 While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and are herein described in detail. The drawings may not be to scale. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all
15 modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

20 The terms “design,” “design data,” and “design information” as used interchangeably herein generally refer to the physical design (layout) of an IC and data derived from the physical design through complex simulation or simple geometric and Boolean operations. In addition, an image of a reticle acquired by a reticle inspection system and/or derivatives thereof can be used as a “proxy” or “proxies” for the design.
25 Such a reticle image or a derivative thereof can serve as a substitute for the design layout in any embodiments described herein that use a design. The design may include any other design data or design data proxies described in commonly owned U.S. Patent Nos. 7,570,796 issued on August 4, 2009 to Zafar et al. and 7,676,077 issued on March 9, 2010 to Kulkarni et al., both of which are incorporated by reference as if fully set forth
30 herein. In addition, the design data can be standard cell library data, integrated layout

data, design data for one or more layers, derivatives of the design data, and full or partial chip design data.

5 In addition, the “design,” “design data,” and “design information” described herein refers to information and data that is generated by semiconductor device designers in a design process and is therefore available for use in the embodiments described herein well in advance of printing of the design on any physical specimens such as reticles and wafers.

10 Turning now to the drawings, it is noted that the figures are not drawn to scale. In particular, the scale of some of the elements of the figures is greatly exaggerated to emphasize characteristics of the elements. It is also noted that the figures are not drawn to the same scale. Elements shown in more than one figure that may be similarly configured have been indicated using the same reference numerals. Unless otherwise
15 noted herein, any of the elements described and shown may include any suitable commercially available elements.

One embodiment relates to a system configured to perform diagnostic functions for a deep learning model. Some embodiments described herein are configured as
20 systems with optional visualization capability for causal understanding and guided training of a deep learning model for semiconductor applications such as inspection and metrology. For example, the embodiments described herein provide a system configured to perform quality assurance and causal understanding for a deep learning model. In particular, as described further herein, the embodiments are configured for generating
25 causal information (e.g., causal image/vector) through several possible methods and/or algorithms. In addition, by using the causal information, the embodiments can quantitatively determine the model performance. Furthermore, the systems can use the information gained by quality assurance and/or causal understanding to perform one or more functions such as providing guidance on data augmentation and/or fine-tuning the
30 process to further improve the accuracy of the deep learning model. In other words, by

using causal information (e.g., causal image/vector) in augmentation, the embodiments can improve the deep learning model further. Moreover, the embodiments described herein provide semi-supervised defect or region detection, which can advantageously reduce manual labeling efforts.

5

As described further herein, deep learning models can perform functions such as image composition, feature extraction, and classification by using an end-to-end neural network structure such as those described further herein. Although such a strategy grants the deep learning model excellent accuracy for prediction, it also essentially makes the deep learning model a black box (in that the inner workings of the deep learning model are not necessarily available for human comprehension). The embodiments described herein are aimed to provide a system for diagnosing a trained deep learning model and optionally to use a visualization tool to augment and/or generate regions of interest (ROIs) in the input data for applications such as fine-tuning the deep learning model.

10

15

One embodiment of a system configured to perform diagnostic functions for a deep learning model is shown in Fig. 1. The system includes one or more computer subsystems (e.g., computer subsystem 36 and computer subsystem(s) 102) and one or more components 100 executed by the one or more computer subsystems. In some embodiments, the system includes imaging tool 10 configured for generating images of a specimen. In the embodiment of Fig. 1, the imaging tool is configured for scanning light over or directing light to a physical version of the specimen while detecting light from the specimen to thereby generate the images for the specimen. The imaging tool may also be configured to perform the scanning (or directing) and the detecting with multiple modes.

20

25

In one embodiment, the specimen is a wafer. The wafer may include any wafer known in the art. In another embodiment, the specimen is a reticle. The reticle may include any reticle known in the art.

In one embodiment, the imaging tool is configured as an optical based imaging tool. In this manner, in some embodiments, the images are generated by an optical based imaging tool. In one such example, in the embodiment of the system shown in Fig. 1, optical based imaging tool 10 includes an illumination subsystem configured to direct light to specimen 14. The illumination subsystem includes at least one light source. For example, as shown in Fig. 1, the illumination subsystem includes light source 16. In one embodiment, the illumination subsystem is configured to direct the light to the specimen at one or more angles of incidence, which may include one or more oblique angles and/or one or more normal angles. For example, as shown in Fig. 1, light from light source 16 is directed through optical element 18 and then lens 20 to specimen 14 at an oblique angle of incidence. The oblique angle of incidence may include any suitable oblique angle of incidence, which may vary depending on, for instance, characteristics of the specimen.

The imaging tool may be configured to direct the light to the specimen at different angles of incidence at different times. For example, the imaging tool may be configured to alter one or more characteristics of one or more elements of the illumination subsystem such that the light can be directed to the specimen at an angle of incidence that is different than that shown in Fig. 1. In one such example, the imaging tool may be configured to move light source 16, optical element 18, and lens 20 such that the light is directed to the specimen at a different oblique angle of incidence or a normal (or near normal) angle of incidence.

In some instances, the imaging tool may be configured to direct light to the specimen at more than one angle of incidence at the same time. For example, the illumination subsystem may include more than one illumination channel, one of the illumination channels may include light source 16, optical element 18, and lens 20 as shown in Fig. 1 and another of the illumination channels (not shown) may include similar elements, which may be configured differently or the same, or may include at least a light source and possibly one or more other components such as those described further herein. If such light is directed to the specimen at the same time as the other light, one or more

characteristics (e.g., wavelength, polarization, etc.) of the light directed to the specimen at different angles of incidence may be different such that light resulting from illumination of the specimen at the different angles of incidence can be discriminated from each other at the detector(s).

5

In another instance, the illumination subsystem may include only one light source (e.g., source 16 shown in Fig. 1) and light from the light source may be separated into different optical paths (e.g., based on wavelength, polarization, etc.) by one or more optical elements (not shown) of the illumination subsystem. Light in each of the different optical paths may then be directed to the specimen. Multiple illumination channels may be configured to direct light to the specimen at the same time or at different times (e.g., when different illumination channels are used to sequentially illuminate the specimen). In another instance, the same illumination channel may be configured to direct light to the specimen with different characteristics at different times. For example, in some instances, optical element 18 may be configured as a spectral filter and the properties of the spectral filter can be changed in a variety of different ways (e.g., by swapping out the spectral filter) such that different wavelengths of light can be directed to the specimen at different times. The illumination subsystem may have any other suitable configuration known in the art for directing the light having different or the same characteristics to the specimen at different or the same angles of incidence sequentially or simultaneously.

10

15

20

In one embodiment, light source 16 may include a broadband plasma (BBP) light source. In this manner, the light generated by the light source and directed to the specimen may include broadband light. However, the light source may include any other suitable light source such as a laser. The laser may include any suitable laser known in the art and may be configured to generate light at any suitable wavelength or wavelengths known in the art. In addition, the laser may be configured to generate light that is monochromatic or nearly-monochromatic. In this manner, the laser may be a narrowband laser. The light source may also include a polychromatic light source that generates light at multiple discrete wavelengths or wavebands.

25

30

Light from optical element 18 may be focused onto specimen 14 by lens 20. Although lens 20 is shown in Fig. 1 as a single refractive optical element, it is to be understood that, in practice, lens 20 may include a number of refractive and/or reflective optical elements that in combination focus the light from the optical element to the specimen. The illumination subsystem shown in Fig. 1 and described herein may include any other suitable optical elements (not shown). Examples of such optical elements include, but are not limited to, polarizing component(s), spectral filter(s), spatial filter(s), reflective optical element(s), apodizer(s), beam splitter(s), aperture(s), and the like, which may include any such suitable optical elements known in the art. In addition, the imaging tool may be configured to alter one or more of the elements of the illumination subsystem based on the type of illumination to be used for imaging.

The imaging tool may also include a scanning subsystem configured to cause the light to be scanned over the specimen. For example, the imaging tool may include stage 22 on which specimen 14 is disposed during inspection. The scanning subsystem may include any suitable mechanical and/or robotic assembly (that includes stage 22) that can be configured to move the specimen such that the light can be scanned over the specimen. In addition, or alternatively, the imaging tool may be configured such that one or more optical elements of the imaging tool perform some scanning of the light over the specimen. The light may be scanned over the specimen in any suitable fashion such as in a serpentine-like path or in a spiral path.

The imaging tool further includes one or more detection channels. At least one of the one or more detection channels includes a detector configured to detect light from the specimen due to illumination of the specimen by the tool and to generate output responsive to the detected light. For example, the imaging tool shown in Fig. 1 includes two detection channels, one formed by collector 24, element 26, and detector 28 and another formed by collector 30, element 32, and detector 34. As shown in Fig. 1, the two detection channels are configured to collect and detect light at different angles of

collection. In some instances, both detection channels are configured to detect scattered light, and the detection channels are configured to detect light that is scattered at different angles from the specimen. However, one or more of the detection channels may be configured to detect another type of light from the specimen (e.g., reflected light).

5

As further shown in Fig. 1, both detection channels are shown positioned in the plane of the paper and the illumination subsystem is also shown positioned in the plane of the paper. Therefore, in this embodiment, both detection channels are positioned in (e.g., centered in) the plane of incidence. However, one or more of the detection channels may be positioned out of the plane of incidence. For example, the detection channel formed by collector 30, element 32, and detector 34 may be configured to collect and detect light that is scattered out of the plane of incidence. Therefore, such a detection channel may be commonly referred to as a “side” channel, and such a side channel may be centered in a plane that is substantially perpendicular to the plane of incidence.

10

15

Although Fig. 1 shows an embodiment of the imaging tool that includes two detection channels, the imaging tool may include a different number of detection channels (e.g., only one detection channel or two or more detection channels). In one such instance, the detection channel formed by collector 30, element 32, and detector 34 may form one side channel as described above, and the imaging tool may include an additional detection channel (not shown) formed as another side channel that is positioned on the opposite side of the plane of incidence. Therefore, the imaging tool may include the detection channel that includes collector 24, element 26, and detector 28 and that is centered in the plane of incidence and configured to collect and detect light at scattering angle(s) that are at or close to normal to the specimen surface. This detection channel may therefore be commonly referred to as a “top” channel, and the imaging tool may also include two or more side channels configured as described above. As such, the imaging tool may include at least three channels (i.e., one top channel and two side channels), and each of the at least three channels has its own collector, each of which is configured to collect light at different scattering angles than each of the other collectors.

20

25

30

As described further above, each of the detection channels included in the imaging tool may be configured to detect scattered light. Therefore, the imaging tool shown in Fig. 1 may be configured for dark field (DF) imaging of specimens. However, the imaging tool may also or alternatively include detection channel(s) that are configured for bright field (BF) imaging of specimens. In other words, the imaging tool may include at least one detection channel that is configured to detect light specularly reflected from the specimen. Therefore, the imaging tools described herein may be configured for only DF, only BF, or both DF and BF imaging. Although each of the collectors are shown in Fig. 1 as single refractive optical elements, it is to be understood that each of the collectors may include one or more refractive optical element(s) and/or one or more reflective optical element(s).

The one or more detection channels may include any suitable detectors known in the art. For example, the detectors may include photo-multiplier tubes (PMTs), charge coupled devices (CCDs), time delay integration (TDI) cameras, and any other suitable detectors known in the art. The detectors may also include non-imaging detectors or imaging detectors. In this manner, if the detectors are non-imaging detectors, each of the detectors may be configured to detect certain characteristics of the scattered light such as intensity but may not be configured to detect such characteristics as a function of position within the imaging plane. As such, the output that is generated by each of the detectors included in each of the detection channels of the imaging tool may be signals or data, but not image signals or image data. In such instances, a computer subsystem such as computer subsystem 36 may be configured to generate images of the specimen from the non-imaging output of the detectors. However, in other instances, the detectors may be configured as imaging detectors that are configured to generate image signals or image data. Therefore, the imaging tool may be configured to generate the images described herein in a number of ways.

It is noted that Fig. 1 is provided herein to generally illustrate a configuration of an imaging tool that may be included in the system embodiments described herein or that may generate images that are used by the system embodiments described herein.

Obviously, the imaging tool configuration described herein may be altered to optimize the performance of the imaging tool as is normally performed when designing a commercial imaging tool. In addition, the systems described herein may be implemented using an existing tool (e.g., by adding functionality described herein to an existing tool) such as the 29xx/39xx and Puma 9xxx series of tools that are commercially available from KLA-Tencor, Milpitas, Calif. For some such systems, the embodiments described herein may be provided as optional functionality of the tool (e.g., in addition to other functionality of the tool). Alternatively, the imaging tool described herein may be designed “from scratch” to provide a completely new imaging tool.

Computer subsystem 36 of the imaging tool may be coupled to the detectors of the imaging tool in any suitable manner (e.g., via one or more transmission media, which may include “wired” and/or “wireless” transmission media) such that the computer subsystem can receive the output generated by the detectors during scanning of the specimen. Computer subsystem 36 may be configured to perform a number of functions described further herein using the output of the detectors.

The computer subsystems shown in Fig. 1 (as well as other computer subsystems described herein) may also be referred to herein as computer system(s). Each of the computer subsystem(s) or system(s) described herein may take various forms, including a personal computer system, image computer, mainframe computer system, workstation, network appliance, Internet appliance, or other device. In general, the term “computer system” may be broadly defined to encompass any device having one or more processors, which executes instructions from a memory medium. The computer subsystem(s) or system(s) may also include any suitable processor known in the art such as a parallel processor. In addition, the computer subsystem(s) or system(s) may include a computer

platform with high speed processing and software, either as a standalone or a networked tool.

If the system includes more than one computer subsystem, then the different computer subsystems may be coupled to each other such that images, data, information, instructions, etc. can be sent between the computer subsystems as described further herein. For example, computer subsystem 36 may be coupled to computer subsystem(s) 102 as shown by the dashed line in Fig. 1 by any suitable transmission media, which may include any suitable wired and/or wireless transmission media known in the art. Two or more of such computer subsystems may also be effectively coupled by a shared computer-readable storage medium (not shown).

Although the imaging tool is described above as being an optical or light-based imaging tool, in some embodiments, the imaging tool may be configured as an electron beam based imaging tool. In this manner, in some embodiments, the images are generated by an electron beam based imaging tool. In one such embodiment shown in Fig. 2, the imaging tool includes electron column 122 coupled to computer subsystem 124. As also shown in Fig. 2, the electron column includes electron beam source 126 configured to generate electrons that are focused to specimen 128 by one or more elements 130. The electron beam source may include, for example, a cathode source or emitter tip, and one or more elements 130 may include, for example, a gun lens, an anode, a beam limiting aperture, a gate valve, a beam current selection aperture, an objective lens, and a scanning subsystem, all of which may include any such suitable elements known in the art.

Electrons returned from the specimen (e.g., secondary electrons) may be focused by one or more elements 132 to detector 134. One or more elements 132 may include, for example, a scanning subsystem, which may be the same scanning subsystem included in element(s) 130.

The electron column may include any other suitable elements known in the art. In addition, the electron column may be further configured as described in U.S. Patent Nos. 8,664,594 issued April 4, 2014 to Jiang et al., 8,692,204 issued April 8, 2014 to Kojima et al., 8,698,093 issued April 15, 2014 to Gubbens et al., and 8,716,662 issued May 6,
5 2014 to MacDonald et al., which are incorporated by reference as if fully set forth herein.

Although the electron column is shown in Fig. 2 as being configured such that the electrons are directed to the specimen at an oblique angle of incidence and are scattered from the specimen at another oblique angle, it is to be understood that the electron beam
10 may be directed to and scattered from the specimen at any suitable angles. In addition, the electron beam based imaging tool may be configured to use multiple modes to generate images of the specimen as described further herein (e.g., with different illumination angles, collection angles, etc.). The multiple modes of the electron beam based imaging tool may be different in any image generation parameters of the imaging
15 tool.

Computer subsystem 124 may be coupled to detector 134 as described above. The detector may detect electrons returned from the surface of the specimen thereby forming electron beam images of the specimen. The electron beam images may include any
20 suitable electron beam images. Computer subsystem 124 may be configured to perform one or more functions described further herein for the specimen using output generated by detector 134. Computer subsystem 124 may be configured to perform any additional step(s) described herein. A system that includes the imaging tool shown in Fig. 2 may be further configured as described herein.

25 It is noted that Fig. 2 is provided herein to generally illustrate a configuration of an electron beam based imaging tool that may be included in the embodiments described herein. As with the optical based imaging tool described above, the electron beam based imaging tool configuration described herein may be altered to optimize the performance
30 of the imaging tool as is normally performed when designing a commercial imaging tool.

In addition, the systems described herein may be implemented using an existing tool (e.g., by adding functionality described herein to an existing tool) such as the eSxxx and eDR-xxxx series of tools that are commercially available from KLA-Tencor. For some such systems, the embodiments described herein may be provided as optional functionality of the tool (e.g., in addition to other functionality of the tool). Alternatively, the tool described herein may be designed “from scratch” to provide a completely new tool.

Although the imaging tool is described above as being an optical based or electron beam based imaging tool, the imaging tool may be an ion beam based imaging tool. Such an imaging tool may be configured as shown in Fig. 2 except that the electron beam source may be replaced with any suitable ion beam source known in the art. In addition, the imaging tool may be any other suitable ion beam based imaging tool such as those included in commercially available focused ion beam (FIB) systems, helium ion microscopy (HIM) systems, and secondary ion mass spectroscopy (SIMS) systems.

As noted above, the imaging tool is configured for scanning energy (e.g., light or electrons) over a physical version of the specimen thereby generating actual images for the physical version of the specimen. In this manner, the imaging tool may be configured as an “actual” tool, rather than a “virtual” tool. For example, a storage medium (not shown) and computer subsystem(s) 102 shown in Fig. 1 may be configured as a “virtual” tool. In particular, the storage medium and the computer subsystem(s) are not part of imaging tool 10 and do not have any capability for handling the physical version of the specimen. In other words, in tools configured as virtual tools, the output of its one or more “detectors” may be output that was previously generated by one or more detectors of an actual tool and that is stored in the virtual tool, and during the “scanning,” the virtual tool may replay the stored output as though the specimen is being scanned. In this manner, scanning the specimen with a virtual tool may appear to be the same as though a physical specimen is being scanned with an actual tool, while, in reality, the “scanning” involves simply replaying output for the specimen in the same manner as the specimen may be scanned. Systems and methods configured as “virtual” inspection tools are

described in commonly assigned U.S. Patent Nos. 8,126,255 issued on February 28, 2012 to Bhaskar et al. and 9,222,895 issued on December 29, 2015 to Duffy et al., both of which are incorporated by reference as if fully set forth herein. The embodiments described herein may be further configured as described in these patents. For example, the one or more computer subsystems described herein may be further configured as described in these patents. In addition, configuring the one or more virtual systems as a central compute and storage (CCS) system may be performed as described in the above-referenced patent to Duffy. The persistent storage mechanisms described herein can have distributed computing and storage such as the CCS architecture, but the embodiments described herein are not limited to that architecture.

As further noted above, the imaging tool may be configured to generate images of the specimen with multiple modes. In general, a “mode” can be defined by the values of parameters of the imaging tool used for generating images of a specimen or the output used to generate images of the specimen. Therefore, modes that are different may be different in the values for at least one of the imaging parameters of the imaging tool. For example, in one embodiment of an optical based imaging tool, at least one of the multiple modes uses at least one wavelength of light for illumination that is different from at least one wavelength of the light for illumination used for at least one other of the multiple modes. The modes may be different in the illumination wavelength as described further herein (e.g., by using different light sources, different spectral filters, etc.) for different modes. In another embodiment, at least one of the multiple modes uses an illumination channel of the imaging tool that is different from an illumination channel of the imaging tool used for at least one other of the multiple modes. For example, as noted above, the imaging tool may include more than one illumination channel. As such, different illumination channels may be used for different modes.

In one embodiment, the imaging tool is configured as an inspection tool. For example, the optical and electron beam imaging tools described herein may be configured as inspection tools. In this manner, the image inputs to the deep learning model described

herein are generated by an inspection tool in some embodiments. In another embodiment, the imaging tool is configured as a metrology tool. For example, the optical and electron beam imaging tools described herein may be configured as metrology tools. In particular, the embodiments of the imaging tools described herein and shown in Figs. 1 and 2 may
5 be modified in one or more parameters to provide different imaging capability depending on the application for which they will be used. In one such example, the imaging tool shown in Fig. 1 may be configured to have a higher resolution if it is to be used for metrology rather than for inspection. In other words, the embodiments of the imaging tool shown in Figs. 1 and 2 describe some general and various configurations for an
10 imaging tool that can be tailored in a number of manners that will be obvious to one skilled in the art to produce imaging tools having different imaging capabilities that are more or less suitable for different applications.

The one or more computer subsystems may be configured for acquiring the
15 images for the specimen generated by an imaging tool described herein. Acquiring the images may be performed using one of the imaging tools described herein (e.g., by directing light or an electron beam to the specimen and detecting light or an electron beam, respectively, from the specimen). In this manner, acquiring the images may be performed using the physical specimen itself and some sort of imaging hardware.
20 However, acquiring the images does not necessarily include imaging the specimen using imaging hardware. For example, another system and/or method may generate the images and may store the generated images in one or more storage media such as a virtual inspection system as described herein or another storage media described herein. Therefore, acquiring the images may include acquiring the images from the storage media
25 in which they have been stored.

The component(s), e.g., component(s) 100 shown in Fig. 1, executed by the computer subsystem(s), e.g., computer subsystem 36 and/or computer subsystem(s) 102, include deep learning model 104. The deep learning model is configured for determining

information from an image generated for a specimen by an imaging tool. Deep learning model 104 may have one of the configurations described further herein.

5 Generally speaking, “deep learning” (also known as deep structured learning, hierarchical learning or deep machine learning) is a branch of machine learning based on a set of algorithms that attempt to model high level abstractions in data. In a simple case, there may be two sets of neurons: ones that receive an input signal and ones that send an output signal. When the input layer receives an input, it passes on a modified version of the input to the next layer. In a deep network, there are many layers between the input and output (and the layers are not made of neurons but it can help to think of it that way),
10 allowing the algorithm to use multiple processing layers, composed of multiple linear and non-linear transformations.

Deep learning is part of a broader family of machine learning methods based on
15 learning representations of data. An observation (e.g., an image) can be represented in many ways such as a vector of intensity values per pixel, or in a more abstract way as a set of edges, regions of particular shape, etc. Some representations are better than others at simplifying the learning task (e.g., face recognition or facial expression recognition). One of the promises of deep learning is replacing handcrafted features with efficient
20 algorithms for unsupervised or semi-supervised feature learning and hierarchical feature extraction.

Research in this area attempts to make better representations and create models to learn these representations from large-scale unlabeled data. Some of the representations
25 are inspired by advances in neuroscience and are loosely based on interpretation of information processing and communication patterns in a nervous system, such as neural coding which attempts to define a relationship between various stimuli and associated neuronal responses in the brain.

In another embodiment, the deep learning model is a machine learning model. Machine learning can be generally defined as a type of artificial intelligence (AI) that provides computers with the ability to learn without being explicitly programmed. Machine learning focuses on the development of computer programs that can teach
5 themselves to grow and change when exposed to new data. In other words, machine learning can be defined as the subfield of computer science that “gives computers the ability to learn without being explicitly programmed.” Machine learning explores the study and construction of algorithms that can learn from and make predictions on data – such algorithms overcome following strictly static program instructions by making data
10 driven predictions or decisions, through building a model from sample inputs.

The machine learning described herein may be further performed as described in “Introduction to Statistical Machine Learning,” by Sugiyama, Morgan Kaufmann, 2016, 534 pages; “Discriminative, Generative, and Imitative Learning,” Jebara, MIT Thesis,
15 2002, 212 pages; and “Principles of Data Mining (Adaptive Computation and Machine Learning),” Hand et al., MIT Press, 2001, 578 pages; which are incorporated by reference as if fully set forth herein. The embodiments described herein may be further configured as described in these references.

20 In some embodiments, the deep learning model is a generative model. A “generative” model can be generally defined as a model that is probabilistic in nature. In other words, a “generative” model is not one that performs forward simulation or rule-based approaches. Instead, as described further herein, the generative model can be learned (in that its parameters can be learned) based on a suitable training set of data. In
25 one embodiment, the deep learning model is configured as a deep generative model. For example, the model may be configured to have a deep learning architecture in that the model may include multiple layers, which perform a number of algorithms or transformations.

In another embodiment, the deep learning model is configured as a neural network. In a further embodiment, the deep learning model may be a deep neural network with a set of weights that model the world according to the data that it has been fed to train it. Neural networks can be generally defined as a computational approach
5 which is based on a relatively large collection of neural units loosely modeling the way a biological brain solves problems with relatively large clusters of biological neurons connected by axons. Each neural unit is connected with many others, and links can be enforcing or inhibitory in their effect on the activation state of connected neural units. These systems are self-learning and trained rather than explicitly programmed and excel
10 in areas where the solution or feature detection is difficult to express in a traditional computer program.

Neural networks typically consist of multiple layers, and the signal path traverses from front to back. The goal of the neural network is to solve problems in the same way
15 that the human brain would, although several neural networks are much more abstract. Modern neural network projects typically work with a few thousand to a few million neural units and millions of connections. The neural network may have any suitable architecture and/or configuration known in the art.

20 In one embodiment, the information includes a classification for a defect detected on the specimen. In one such embodiment, the deep learning model is configured as an AlexNet. For example, an AlexNet includes a number of convolutional layers (e.g., 5) followed by a number of fully connected layers (e.g., 3) that are, in combination, configured and trained to classify images. Examples of neural networks configured as
25 AlexNets are described in “ImageNet Classification with Deep Convolutional Neural Networks” by Krizhevsky et al., NIPS 2012, which is incorporated by reference as if fully set forth herein. The deep learning models described herein may be further configured as described in this reference.

In another such embodiment, the deep learning model is configured as a GoogleNet. For example, a GoogleNet may include layers such as convolutional, pooling, and fully connected layers such as those described further herein configured and trained to classify images. While the GoogleNet architecture may include a relatively high number of layers (especially compared to some other neural networks described herein), some of the layers may be operating in parallel, and groups of layers that function in parallel with each other are generally referred to as inception modules. Other of the layers may operate sequentially. Therefore, GoogleNets are different from other neural networks described herein in that not all of the layers are arranged in a sequential structure. Examples of neural networks configured as GoogleNets are described in “Going Deeper with Convolutions,” by Szegedy et al., CVPR 2015, which is incorporated by reference as if fully set forth herein. The deep learning models described herein may be further configured as described in this reference.

In a further such embodiment, the deep learning model is configured as a VGG network. For example, VGG networks were created by increasing the number of convolutional layers while fixing other parameters of the architecture. Adding convolutional layers to increase depth is made possible by using substantially small convolutional filters in all of the layers. Like the other neural networks described herein, VGG networks were created and trained to classify images. VGG networks also include convolutional layers followed by fully connected layers. Examples of neural networks configured as VGG are described in “Very Deep Convolutional Networks for Large-Scale Image Recognition,” by Simonyan et al., ICLR 2015, which is incorporated by reference as if fully set forth herein. The deep learning models described herein may be further configured as described in this reference.

In some such embodiments, the deep learning model is configured as a deep residual network. For example, like some other networks described herein, a deep residual network may include convolutional layers followed by fully connected layers, which are, in combination, configured and trained for image classification. In a deep

residual network, the layers are configured to learn residual functions with reference to the layer inputs, instead of learning unreferenced functions. In particular, instead of hoping each few stacked layers directly fit a desired underlying mapping, these layers are explicitly allowed to fit a residual mapping, which is realized by feedforward neural networks with shortcut connections. Shortcut connections are connections that skip one or more layers. A deep residual net may be created by taking a plain neural network structure that includes convolutional layers and inserting shortcut connections which thereby takes the plain neural network and turns it into its residual learning counterpart. Examples of deep residual nets are described in “Deep Residual Learning for Image Recognition” by He et al., NIPS 2015, which is incorporated by reference as if fully set forth herein. The deep learning models described herein may be further configured as described in this reference.

In a further such embodiment, the deep learning model includes one or more fully connected layers configured for classifying defects on the specimen. A “fully connected layer” may be generally defined as a layer in which each of the nodes is connected to each of the nodes in the previous layer. The fully connected layer(s) may perform classification based on the features extracted by convolutional layer(s), which may be configured as described further herein. The fully connected layer(s) are configured for feature selection and classification. In other words, the fully connected layer(s) select features from a feature map and then classify the defects in the image(s) based on the selected features. The selected features may include all of the features in the feature map (if appropriate) or only some of the features in the feature map.

If the deep learning model outputs a classification for a defect detected on the specimen, the deep learning model may output an image classification, which may include a classification result per image with a confidence associated with each classification result. The results of the image classification can also be used as described further herein. The image classification may have any suitable format (such as an image

or defect ID, a defect description such as “pattern,” “bridge,” etc.). The image classification results may be stored and used as described further herein.

5 In some embodiments, the information determined by the deep learning model includes features of the images extracted by the deep learning model. In one such embodiment, the deep learning model includes one or more convolutional layers. The convolutional layer(s) may have any suitable configuration known in the art and are generally configured to determine features for an image as a function of position across the image (i.e., a feature map) by applying a convolution function to the input image
10 using one or more filters. In this manner, the deep learning model (or at least a part of the deep learning model) may be configured as a convolution neural network (CNN). For example, the deep learning model may be configured as a CNN, which is usually stacks of convolution and pooling layers, to extract local features. The embodiments described herein can take advantage of deep learning concepts such as a CNN to solve the normally
15 intractable representation inversion problem. The deep learning model may have any CNN configuration or architecture known in the art. The one or more pooling layers may also have any suitable configuration known in the art (e.g., max pooling layers) and are generally configured for reducing the dimensionality of the feature map generated by the one or more convolutional layers while retaining the most important features.

20

The features determined by the deep learning model may include any suitable features described further herein or known in the art that can be inferred from the input described herein (and possibly used to generate the output described further herein). For example, the features may include a vector of intensity values per pixel. The features
25 may also include any other types of features described herein, e.g., vectors of scalar values, vectors of independent distributions, joint distributions, or any other suitable feature types known in the art.

30 In a further embodiment, the information includes a simulated image generated from the image. The deep learning model may be configured to generate a simulated

image as described in U.S. Patent Application Publication Nos. 2017/0148226 published May 25, 2017 by Zhang et al., 2017/0193680 published July 6, 2017 by Zhang et al., 2017/0194126 published July 6, 2017 by Bhaskar et al., 2017/0200260 published July 13, 2017 by Bhaskar et al., and 2017/0200265 published July 13, 2017 by Bhaskar et al., and
5 U.S. Patent Application Serial No. 15/603,249 filed May 23, 2017 by Zhang et al., which are incorporated by reference as if fully set forth herein. The embodiments described herein may be further configured as described in these publications and applications.

In another embodiment, the information determined by the deep learning model
10 includes one or more segmentation regions generated from the image. In one such embodiment, the deep learning model includes a proposal network configured for identifying the segmentation region(s) (based on features determined for the image) and generating bounding boxes for each of the segmentation regions. The segmentation regions may be detected based on the features (determined for the images by the deep
15 learning model or another method or system) to thereby separate regions in the images based on noise (e.g., to separate noisy regions from quiet regions), to separate regions in the images based on specimen features located therein, to separate regions based on geometric characteristics of the output, etc. The proposal network may use features from a feature map, which may be generated or determined as described further herein, to
20 detect the segmentation region(s) in the image based on the determined features. The proposal network may be configured to generate bounding box detection results. In this manner, the deep learning model may output bounding boxes, which may include a bounding box associated with each segmentation region or more than one segmentation region. The deep learning model may output bounding box locations with each bounding
25 box. The results of the segmentation region generation can also be stored and used as described further herein.

A proposal network or a “region proposal network” can be generally defined as a
fully convolutional network that detects objects in images and proposes regions
30 corresponding to those detected objects. Examples of proposal networks are described in

“Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks” by Ren et al., NIPS, 2015, which is incorporated by reference as if fully set forth herein. The proposal network described herein may be further configured as described in this reference. Proposal networks may be attractive for use in the embodiments described
5 herein because they are relatively fast and are able to detect objects of different size. The proposal network may overlay a sliding window on a feature map generated by a convolution layer(s) thereby generating k anchor boxes. The results of overlaying the sliding window on a feature map may be input to an intermediate layer of the proposal network, which may generate 2k scores via a cls layer and 4k coordinates via a reg layer.

10 In another embodiment, the information includes a multi-dimensional output generated from the image. For example, as described further herein, the information generated by the deep learning model may include one or more properties such as a one-dimensional vector, e.g., defect types, or a multiple dimensional output such as a two-
15 dimensional image-like output, e.g., a simulated image, segmentation region(s), etc., or an n-dimensional output.

In general, the deep learning model described herein is a trained deep learning model. For example, the deep learning model for which one or more diagnostic functions
20 are performed may be previously trained by one or more other systems and/or methods. In addition, the deep learning model may be trained by one or more of the embodiments described herein before the one or more diagnostic functions are performed for the deep learning model. In this manner, the diagnostic functions described herein are unlike deep
25 learning model characterizations that are performed during training of a deep learning model in that in the embodiments described herein, the deep learning model is already generated and trained and then the functionality of the model is determined as described herein, which can then be used to perform one or more additional functions for the deep learning model.

The one or more components include a diagnostic component, e.g., diagnostic component 106 shown in Fig. 1, configured for determining one or more causal portions of the image that resulted in the information being determined. In this manner, the embodiments described herein provide a causal understanding of how the deep learning model generated the information for a particular input or image. In addition, the
5 embodiments described herein provide a causal understanding and assurance of a deep learning model. The following provides a formal definition of causal understanding and causal portions as those terms are used herein.

10 An input to a deep learning model can include a combination of: a) images defined by $x(h, w, c, t, \dots)$, which is an N-dimensional tensor of images with height = h and width = w across other dimensions, e.g., channel c, time t, etc. (In semiconductor applications, x can be an optical image, an electron beam image, a design data image (e.g., CAD image), etc. under different tool conditions.); and b) feature vector $v(m)$,
15 which is a 1-dimensional vector (The dimension can be generalized to be more than 1.).

A deep learning model can be defined as a trainable function of $F(y|x, v; w)$, i.e., given inputs of x and v and model parameters of w, F predicts the property of y. The predicted output y can be one or more properties as a 1-dimensional vector, e.g., defect
20 types. The predicted output y can also be 2-dimensional, image-like output, e.g., a simulated image, segmentation region, etc. In addition, the predicted output y can be an n-dimensional output. Given a trained model, i.e., parameters of w are learned as W, the model response with respect to the input change from (x, v) to (x', v') is defined as $R = F(x', v'; w=W) - F(x, v; w=W)$.

25 In general, the inputs of x and v have relatively high dimensions. For example, given an input image having a width of 128 pixels and a height of 128 pixels, the dimensions are 128 x 128. Given an observed response, not every dimension (i.e., pixel) is equally contributing to the response (e.g., in defect detection applications, the defective

pixels tend to be more important than background pixels, and the neighboring pixels may or may not be important depending on circuit design and process condition).

Generating a causal understanding of a deep learning model as described herein refers to qualitatively and quantitatively identifying the importance of each input pixel and each input feature subject to an observed response. Formally, the importance score of x is called a “causal image,” and the importance score of v as a “causal vector.” In the causal image and causal vector, a positive value means this pixel/feature positively contributes to the response (i.e., the higher the value, the higher the response); a near-zero value means the pixel/feature is irrelevant to the response; and a negative value means the pixel/feature has a negative contribution to the response. Therefore, the regions with positive, near-zero, and negative values are called positive, irrelevant, and negative regions, respectively. The term “causal portions” as used herein refers to any and all of the causal understanding information described above. For example, the causal portions can be defined by both the causal image and the causal vector. In this manner, causal portions may include portion(s) of an image that positively contribute to the response (or information determined by the deep learning model), portion(s) of the image that are largely irrelevant to the response, portion(s) of the image that have a negative contribution to the response, or some combination thereof.

Given a trained deep learning model $F(y|x, v; w=W)$ and a particular input or data (x, v) , the system runs the deep learning model in the forward direction to generate the prediction y and/or related outputs. The related outputs are application-specific, e.g., for a classification task, the related outputs can be the top-N classification results. The response of the model from predefined (x_0, v_0) to (x, v) can be constructed and computed. The “causal image” and/or causal vector is/are computed using one or more methods and/or algorithms described further herein. The choice of (x_0, v_0) is based on application and algorithm.

In a further embodiment, the diagnostic component is configured for determining the one or more causal portions by computing a local sensitivity. For example, as described by Simonyan et al., “Deep inside convolutional networks: Visualizing image classification models and saliency maps,” ICLR Workshop, 2014, which is incorporated
 5 by reference as if fully set forth herein, the causal image can be computed as the local sensitivity,

$$\left\| \frac{\partial}{\partial x} F(x, v; w) \right\|_{\infty} \text{ or } \left\| \frac{\partial}{\partial v} F(x, v; w) \right\|_{\infty}.$$

10

The embodiments described herein may be further configured as described in the above-reference.

In another embodiment, the diagnostic component is configured for determining
 15 the one or more causal portions by causal back propagation. For example, the diagnostic component may perform a top-down backpropagation of the deep learning model. In the backpropagation process at each layer, there are three inputs that are available: (1) activation value from previous layer; (2) trained parameters for this layer; and (3) causal values backpropagated from the next layer. In general, a causal backpropagation function
 20 can be designed for each layer type including, but not limited to, convolutional layer, fully-connected layer, various activation layer, various pooling layers, etc. to assign pixel-wise causal scores in the input space (i.e., dimension) based on the layer computation and the three inputs mentioned above. The causal backpropagation can be started at the highest layer (i.e., the output layer of a network) or any intermediate layer given the
 25 causal values can be initialized in a way. When the causal backpropagation is backpropagated to the lowest layer (i.e., the input layer to a network), the causal weights form the causal image in the image space. The region highlighted in the causal image is

the region that causes the deep learning model to make its decision (i.e., determine the information that it did).

In some embodiments, the diagnostic component is configured for determining the one or more causal portions by causal back propagation performed using a deconvolution heatmap algorithm. A deconvolution heatmap can be viewed as a specific implementation of causal backpropagation. For example, as described by Zeiler et al., “Visualizing and understanding convolutional networks,” ECCV, 2014, pp. 818-833, which is incorporated by reference as if fully set forth herein, the causal image can be computed via mapping activation from the deep learning model’s output back to the pixel/feature (i.e., x and v) space through a backpropagation rule. The embodiments described herein may be further configured as described in this reference.

In still another embodiment, the diagnostic component is configured for determining the one or more causal portions by causal back propagation performed using a layer-wise relevance propagation. A layer-wise relevance propagation can also be viewed as a specific implementation of causal backpropagation. For example, as described by Samek et al., in “Evaluating the visualization of what a Deep Neural Network has learned,” arXiv: 1509.06321, which is incorporated by reference as if fully set forth herein, the causal image/vector can be computed via back-propagating the relevance under the conservation rule. The embodiments described herein may be further configured as described in this reference.

In an additional embodiment, the diagnostic component is configured for determining the one or more causal portions by causal back propagation performed using a deep lift algorithm. A deep lift algorithm can also be considered a specific implementation of causal backpropagation. For example, as described by Shrikumar et al., in “Not Just A Black Box: Learning Important Features through propagating activation differences,” arXiv: 1605.01713, which is incorporated by reference as if fully set forth herein, the causal image/vector can be computed via back-propagating the

activation difference of a deep learning model. The embodiments described herein may be further configured as described in this reference.

In a further embodiment, the diagnostic component is configured for determining the one or more causal portions by global average pooling. As described by Lin et al. in “Network In Network,” arXiv: 1312.4400, which is incorporated by reference as if fully set forth herein, the global average pooling (GAP) is introduced and defined. GAP provides crude pixel-level causal region information, which can be approximately interpreted as causal image/vector. The embodiments described herein may be further configured as described in the above reference.

In some additional embodiments, the diagnostic component is configured for determining the one or more causal portions by computing a path integral on gradients. Similar to Sundararajan et al. in “Axiomatic Attribution for Deep Networks,” arXiv: 1703.01365, which is incorporated by reference as if fully set forth herein, the causal image can be computed as the path integral on gradients,

$$\int_t \frac{\partial F(x, v; w)}{\partial x} \cdot \frac{\partial x}{t} dt$$

where t is a path from (x, v) to (x', v') as discussed further herein. The embodiments described herein may be further configured as described in this reference.

In another embodiment, the diagnostic component is configured for determining the one or more causal portions by computing a partial dependence plot. Partial dependence plot (PDP) is introduced by Friedman in “Greedy Function Approximation: A Gradient Boosting Machine,” The Annals of Statistics, 29(5): 1189-1232, which is incorporated by reference as if fully set forth herein. Causal image can be computed as

$$\int F(x, v; w) dP(x).$$

The embodiments described herein may be further configured as described in this reference.

5

In an additional embodiment, the diagnostic component is configured for determining the one or more causal portions by computing a partial dependence plot with path integral. As the $P(x)$ (or $P(v)$) is defined on a higher dimension, the integral above can be intractable in practice, we can combine the PDP with Path Integral (PDPPI) by defining the $P(x)$ as the probability distribution of x along a path from x to x' . And the formula for PDPPI is:

10

$$\int_t F(x, v; w) dP(x(t)).$$

15

Similarly, PDPPI can also be defined on gradients:

$$\int_t \frac{\partial F(x, v; w)}{\partial x} dP(x(t)).$$

20

The diagnostic component is also configured for performing one or more functions based on the determined one or more causal portions of the image. The one or more functions may include any of the functions described further herein.

25

In some embodiments, the one or more functions include altering one or more parameters of the deep learning model based on the determined one or more causal portions. For example, the diagnostic component may determine if the one or more causal portions are the correct causal portion(s) of the image, which may be performed as described further herein. If the one or more causal portions are incorrect, the diagnostic

component may be configured to fine-tune or re-train the deep learning model to thereby alter one or more parameters of the deep learning model, which may include any of the parameters described herein. The fine-tuning or re-training of the deep learning model may include inputting additional training images to the deep learning model, comparing
5 the output generated for the training images to known output for the training images (e.g., defect classification(s), segmentation region(s), etc.), and altering one or more parameters of the deep learning model until the output generated for the additional training images by the deep learning model substantially matches the known output for the additional training images. In addition, the diagnostic component may be configured to perform any
10 other method and/or algorithm to alter one or more parameters of the deep learning model based on the determined one or more causal portions.

In one embodiment, the one or more functions include determining if the one or more causal portions that resulted in the information being determined are the correct one
15 or more causal portions of the image. For example, as described further herein, the correct one or more causal portions of the image may be identified by a user or users, and then the correct one or more causal portions can be compared to the determined one or more causal portions to determine if the correct and determined causal portions match and/or to determine differences between the correct and determined causal portions. The
20 correct causal portion(s) of the image may also be acquired from another method or system rather than from a user. For example, by inputting the image to a different deep learning model having known acceptable functionality, the diagnostic component described herein can be configured or used to determine the causal portion(s) of the image that resulted in the information being determined by the different deep learning model.
25 Given the known acceptable functionality of the different deep learning model, those one or more causal portions can be assumed to be the correct one or more causal portions, which can then be compared to the one or more causal portions determined by the diagnostic component for the deep learning model to determine if the deep learning model identified the correct one or more causal portions of the image.

30

In another embodiment, the one or more components include a visualization component configured for displaying at least the image, the determined information, and the determined one or more causal portions to a user. For example, the displayed information/data provided to a user by the visualization component may include, but is not limited to, the input data (x, v), prediction y and related outputs, causal image and/or causal vector, i.e., the causal portion(s), ground truth (if any), and other related data and/or meta data. The visualization component may be deployed with a trained deep learning model on a system such as those described further herein. The visualization component may also be deployed with any suitable graphical user interface (GUI) known in the art. The visualization component may be disabled in normal runtime mode and enabled in the diagnostic mode. When the visualization component is enabled, the causal images are generated and displayed for selected input data. The causal images may be checked by one or more users such as domain experts or crowd sourcing or comparison algorithms and/or methods against region labels (if available). In this manner, the embodiments described herein provide a way for visualizing the causal of the prediction on a particular input by a deep learning model.

Fig. 3 illustrates one or more steps that can be performed by the embodiments described herein. Fig. 3 also illustrates steps that may be performed to develop or create a deep learning model, which may be performed by the embodiments described herein or by another system or method, in addition to one or more steps described further herein. The deep learning model development workflow may include data collection 302 from imaging tool 300, which may include any of the imaging tools described herein. The data that is collected may include output described further herein (e.g., images, image data, etc.) of any of the imaging tools described herein. The data that is collected may be for one or more specimens that may or may not be of the same type (e.g., multiple specimens of the same type or multiple specimens at least one of which is of a type different from the others). Data collection 302 may therefore produce raw data 304.

The deep learning model development workflow may also include data labeling 306 and optional data augmentation 312. The data labeling may be performed based on user input 308 such as input from a crowd or one or more experts. The data labeling may be performed in any suitable manner and the labels may be any suitable labels known in the art. Data labeling 308 may produce labels 310. Data augmentation may be performed as described further herein and may generate processed data 314. The labels and the processed data may then be integrated in integration 316, which may be performed in any suitable manner, thereby generating data and labels 318.

The deep learning model development workflow may also include data partition 320, in which data and labels 318 is separated into training data 322, validation data 324, and test data 326. Data partitioning may be performed in any suitable manner. For example, the data and labels may be simply divided into three different equal portions such that some data (and its corresponding labels) is available and used for training, some other data (and its corresponding labels) is used for validation, and still other data (and its corresponding labels) is used for testing.

Training data 322 may be input to model training 328, which may be performed in any suitable manner. For example, the model training may include inputting the training data to the deep learning model and modifying one or more parameters of the model until the output of the model is the same as (or substantially the same as) the labels assigned to the data. Model training 328 may generate one or more trained models, which may then be sent to model selection 330, which is performed using validation data 324. The results that are produced by each one or more trained models for the validation data that is input to the one or more trained models may be compared to the labels assigned to the validation data to determine which of the models is the best model, e.g., best model 332. For example, the model that produces results that most closely match the validation data labels may be selected as the best model. Test data 326 may then be used for model evaluation 334 of best model 332. Model evaluation may be performed in any suitable manner. For example, the test data may be input to the best model and the results

produced by the best model for the test data may be compared to the labels for the test data to determine how closely the results produced by the best model match the labels. Best model 332 may also be sent to model deployment 338 in which the best model may be sent to imaging tool 300 for use in a production or runtime mode (post-training mode).
5 The best model may then be applied to additional images, data, output, etc. generated by the imaging tool.

The model evaluation results may also be sent to visualization 336, which may be performed by one or more of the systems described herein and may include determining
10 one or more causal portions of the image that resulted in the information being determined and for performing one or more functions based on the determined one or more causal portions of the image. In other words, the “visualization” shown in Fig. 3 may essentially be any of the functionality described herein of the diagnostic component. In some instances, the information determined from an image generated for a specimen
15 by a deep learning model may be provided to the diagnostic component or visualization 336 via model deployment 338. However, the information determined from an image generated for a specimen by a deep learning model may be provided to the diagnostic component or visualization 336 via model evaluation 334.

In one such embodiment, the one or more components include a user interface component configured to receive input from the user after the displaying, and the one or more functions performed by the diagnostic component are determined based on the input from the user. For example, the embodiments described herein allow a user to reinforce the correct causal relationship to be learned by a deep learning model. In one particular
25 example, as described further herein, one or more users may provide input to the embodiments described herein regarding the correct causal portion(s) of the image input. The correct causal portion(s) received from the user may then be used to perform one or more other functions described herein. The user interface component may have any suitable configuration known in the art that can be used to receive the input from the user.

30

In a further embodiment, the one or more functions include determining one or more characteristics of the one or more causal portions and determining, based on the one or more characteristics of the one or more causal portions, if additional images for the specimen should be collected from the imaging tool and used for additional training of the deep learning model. For example, the diagnostic component or visualization 336 may be added after model evaluation as shown in Fig. 3 and it may fall back to data collection 302, if causal assurance failed on a) considerable samples of one type or class; and/or b) considerable samples of several types or classes. If this path is selected, extra data for the error types or classes are collected from imaging tool 300 for further training. For example, as shown in Fig. 3, visualization 336 may send output such as instructions for additional data collection to data collection 302 step, which may be performed using imaging tool 300. The additional data collection may be performed using the same specimens that were used for initial data collection and/or different specimens not previously used for data collection.

The additional data collection may or may not be performed with the same imaging parameters as the original data collection. For example, it may be advantageous to try to collect different or better data with different imaging parameters that are more suitable for generating input data for the deep learning model. In this manner, the embodiments described herein may also provide an iterative feedback loop in which models are trained and evaluated, and if a suitable model cannot be generated and trained, then different imaging data may be generated with different imaging parameters until suitable images or data have been generated that resulted in a suitably trained model being developed and trained. The fall back path involving additional data collection also allows a domain expert and/or other user(s) to reinforce and to fine tune a deep learning model to learn the correct causal relationship at a finer granularity.

In another embodiment, the one or more functions include determining one or more characteristics of the one or more causal portions and altering the image based on the one or more characteristics of the one or more causal portions to thereby generate an

augmented image for input to the deep learning model. For example, the embodiments may augment the input image based on guidance generated by the system from the one or more characteristics. Altering the image based on the one or more characteristics of the causal portion(s) to thereby generate an augmented image may be performed as described further herein.

In an additional embodiment, the one or more functions include determining one or more characteristics of the one or more causal portions and generating a data augmentation method for application to additional images input to the deep learning model. For example, the embodiments may generate one or more parameters of a new data augmentation method and/or alter one or more parameters of an existing data augmentation method based on guidance generated by the system from the one or more characteristics to thereby generate a data augmentation method. The data augmentation method can be applied to the images that were input to the deep learning model for which the causal portion(s) were determined and/or to images that have not yet but will be input to the deep learning model.

In the embodiments described above, the diagnostic component or visualization 336 may be added after model evaluation as shown in Fig. 3 and it may fall back to data augmentation 312, if causal assurance failed on a) a few samples of one type or class; b) ambiguously-labeled samples; c) samples belonging to multiple types or classes; or some combination thereof. If this path is selected, the causal image (from visualization 336 or the diagnostic component) and the labeled causal region or portion (e.g., from domain expert or crowd sourcing) are compared to generate a data augmentation method. The data augmentation method may include applying a mask, function, filter, algorithm, etc. to the data or images generated by the imaging tool. The mask, function, filter, algorithm, etc. may have any suitable type or format such as random, fuzzy, or noisy. The data augmentation method may advantageously enhance the causal portion(s) of the images or data (or the positively contributing portions of the input images or data) and/or de-

emphasize the non-causal portion(s) of the images or data (or the largely irrelevant and/or negatively contributing portion(s) of the images or data).

5 In one such example, the causal information may be generated for an input image and if the relevant region in the causal information matches the defect to be detected (in the case of defect detection or classification), the diagnostic component may determine that no augmentation needs to be performed as the model predicted correctly. However, if the relevant region only matches part of a defect or does not match a defect at all (in the case of defect detection or classification), the diagnostic component may determine that
10 an augmentation method may be advantageous and may request input from a user for a possible augmentation method. The user may then, for example, specify one or more attention portions and/or one or more ignore regions in the input image via bounding boxes, locations, etc. The information for these user-specified portions can be sent to the augmentation step to alter the input image, for example, by randomly perturbing the
15 ignore portion(s) by zeroing or adding noise and/or randomly transforming the attention portion(s).

The data augmentation method may then be applied to corresponding input data to thereby generate new augmented data. For example, as shown in Fig. 3, visualization 336
20 may send the new data augmentation method to data augmentation 312 step, which may be applied to raw data 304. The fall back path involving data augmentation allows a domain expert and/or other user(s) to reinforce and to fine tune a deep learning model to learn the correct causal relationship at a finer granularity.

25 Data augmentation 312 may, therefore, be performed differently depending on when it is performed and which information it is performed based on. For example, in a first run (or in runs performed prior to having the information for the causal portion(s)), the augmentation is optional and may only involve performing physical transformation(s) such as translation, rotation, mirroring, etc. However, in later run(s), when the causal
30 information (e.g., causal map/causal vector) and possibly user input is/are available, the

augmentation may be performed using the causal information determined by the diagnostic component and (optional) user input to generate augmented data.

The visualization step that is performed after model evaluation may also fall back to model deployment if causal assurance is passed. In other words, if visualization 336 does not determine that additional data should be collected and/or that a new or altered data augmentation method should be generated and/or tested, then the model may be deployed as a trained model that is used at runtime or in production.

Fig. 4 illustrates one or more steps that can be performed by the embodiments described herein. Fig. 4 also illustrates steps that may be performed to develop or create a deep learning model, which may be performed by the embodiments described herein or by another system or method, in addition to one or more steps described further herein. The deep learning model development workflow may include data collection 402 from imaging tool 400, which may include any of the imaging tools described herein. The data that is collected may include output described further herein (e.g., images, image data, etc.) of any of the imaging tools described herein. The data that is collected may be for one or more specimens that may or may not be of the same type (e.g., multiple specimens of the same type or multiple specimens at least one of which is of a type different from the others). Data collection 402 may therefore produce raw data 404.

The deep learning model development workflow may also include data labeling 406 and optional data augmentation 412. The data labeling may be performed based on user input 408 such as input from a crowd and/or one or more experts. The data labeling may be performed in any suitable manner and the labels may be any suitable labels known in the art. Data labeling 406 may produce labels 410. Data augmentation may be performed as described further herein and may generate processed data 414. The labels and the processed data may then be integrated in integration 416, which may be performed in any suitable manner, thereby generating data and labels 418.

The deep learning model development workflow may also include data partition 420, in which data and labels 418 is separated into training data 422 and any other data (not shown in Fig. 4) described herein. Data partitioning may be performed in any suitable manner. For example, the data and labels may be simply divided into different equal portions such that some data (and its corresponding labels) is available and used for training and other data (and its corresponding labels) is available and used for other purposes.

Training data 422 may be input to model training 1 424, which may be performed in any suitable manner. For example, the model training may include inputting the data to the deep learning model and modifying one or more parameters of the model until the output of the model is the same as (or substantially the same as) the labels assigned to the data. Model training 1 424 may generate one or more trained models, which may then be sent to model selection 426, which may be performed in any suitable manner. For example, the results that are produced by each one or more trained models may be compared to the labels assigned to validation data to determine which of the models is the best model, e.g., best model 1 428. For example, the model that produces results that most closely match validation data labels may be selected as the best model. Test data may then be used for model evaluation 430 of best model 1 428. Model evaluation may be performed in any suitable manner. For example, test data may be input to best model 1 and the results produced by best model 1 for the test data may be compared to the labels for the test data to determine how closely the results produced by best model 1 match the labels. Best model 1 428 may also be sent to model deployment 444 in which best model 1 may be sent to imaging tool 400 for use in a production or runtime mode (post-training mode). Best model 1 may then be applied to additional images, data, output, etc. generated by the imaging tool.

The model evaluation results may also be sent to detection 432, which may be performed by one or more of the systems described herein and may include determining one or more causal portions of the images that resulted in the information being

determined. In a further embodiment, the one or more functions include identifying the one or more causal portions as one or more regions of interest (ROIs) in the image and tuning the deep learning model based on the one or more ROIs. For example, causal back propagation or another of the causal portion determination methods described herein may be used as semi-supervised ROI detection to better tune the original deep learning model. The reason this is called “semi-supervised” is that the labeling process for best model 1 does not require labeling exactly the bounding box for each object. In this manner, the one or more causal portions may be identified and used as semi-supervised region detection information. In other words, the diagnostic component described herein may be used as a detection module after the model evaluation to automatically generate ROI(s). In other words, the causal portion(s) specify the important pixels or higher order features responsible for the output generation (e.g., classification or prediction). Therefore, the causal portion information provides pixel-level location information.

For example, as shown in Fig. 4, detection 432 may generate ROI 434, which may include information for any one or more causal portions identified as ROIs, which may be used for cropping 436 of the original image to the candidate patch along with the output (e.g., class prediction) from best model 1. In particular, the original image may be cropped to eliminate portion(s) of the original image that do not correspond to the ROI(s). In one such example, cropped image(s) 438 generated by cropping 436 may be output to data partition 420, which may then use the cropped images to generate additional training data 422, which may replace the original training data. The new training data may then be used to tune best model 1. For example, the new training data may be input to model training 1 424, which may be used to tune or fine tune parameters of best model 1, which may output results to model selection 426. Model selection may produce best model 1 428, which would be a modified version of the best model 1 originally produced. The new best model 1 may then be evaluated as described above and used for detection of ROI(s), which can be used to generate still further training data, which can be used to re-tune the best model 1 again. In this manner, the embodiments described herein provide a

system for iteratively tuning a deep learning model based on ROI(s) determined by previous versions of the deep learning model.

In some embodiments, the one or more functions include identifying the one or more causal portions as one or more ROIs in the image, which may be performed as described herein, and training an additional deep learning model based on the one or more ROIs. For example, causal back propagation or another of the causal portion determination methods described herein may be used as semi-supervised ROI detection to train a second “more accurate” deep learning model based on cropped images. The reason this is called “semi-supervised” is that the labeling process for best model 1 does not require labeling exactly the bounding box for each object. As shown in Fig. 4, for example, cropped image 438 may be provided to model training 2 440. Model training 2 may be performed as described herein, but using a different deep learning model than that trained in model training 1 424. Results of model training 2 may produce best model 2 442, which may then be provided to model deployment 444, which may be performed as described further herein.

As can be seen from the description of the embodiments provided herein, therefore, the embodiments described herein provide a number of new and advantageous features and/or functionality compared to previously used methods for evaluating a deep learning model. For example, the embodiments described herein enable a way to feedback user knowledge to a deep learning model development workflow in a systematic way. In addition, the embodiments described herein enable a way to understand why deep learning models make a prediction via various algorithms. Furthermore, the embodiments described herein provide diagnostics for a deep learning model. The embodiments described herein further provide causal assurance on a trained deep learning model. Moreover, the embodiments described herein provide user-guided data augmentation and/or semi-supervised region detection and fine-tuning on cropped ROI images.

The computer subsystem(s) described herein may be further configured for single image detection as described in U.S. Patent Application Publication No. 2017/0140524 published May 18, 2017 by Karsenti et al., which is incorporated by reference as if fully set forth herein. The embodiments described herein may be further configured as
5 described in commonly owned U.S. Patent Application Publication Nos. 2017/0140524 published May 18, 2017 by Karsenti et al., 2017/0148226 published May 25, 2017 by Zhang et al., 2017/0193400 published July 6, 2017 by Bhaskar et al., 2017/0193680 published July 6, 2017 by Zhang et al., 2017/0194126 published July 6, 2017 by Bhaskar et al., 2017/0200260 published July 13, 2017 by Bhaskar et al., 2017/0200264 published
10 July 13, 2017 by Park et al., and 2017/0200265 published July 13, 2017 by Bhaskar et al., and U.S. Patent Application Serial No. 15/603,249 filed May 23, 2017 by Zhang et al., which are incorporated by reference as if fully set forth herein. The embodiments described herein may also be configured to perform any steps described in these patent application publications and applications.

15
The one or more computer subsystems may be configured for generating results of one or more steps described herein. The results may include any of the results described herein such as the determined information, the determined causal portion(s), results of the one or more functions, etc. The results may be generated by the computer subsystem(s)
20 in any suitable manner. The results may have any suitable form or format such as a standard file type. The computer subsystem(s) may generate the results and store the results such that the results can be used by the computer subsystem(s) and/or another system or method to perform one or more functions for the deep learning model, the imaging tool, the specimen or another specimen of the same type. Such functions
25 include, but are not limited to, altering a process such as a fabrication process or step that was performed on the specimen in a feedback manner, altering a process such as a fabrication process or step that will be performed on the specimen in a feedforward manner, etc.

Another embodiment of a system is also configured to perform diagnostic functions for a deep learning model. The system includes an imaging tool configured for generating images of a specimen. The imaging tool may have any configuration described herein. The system also includes one or more computer subsystems, e.g., computer subsystem(s) 102 shown in Fig. 1, configured for acquiring the images. The computer subsystem(s) may be configured as described further herein. The system also includes one or more components, e.g., component(s) 100, executed by the one or more computer subsystems, which may include any of the component(s) described herein. The component(s) include deep learning model 104, which may be configured as described herein. For example, the deep learning model is configured for determining information from an image generated for the specimen by the imaging tool. The component(s) also include diagnostic component 106 configured for determining one or more causal portions of the image that resulted in the information being determined and for performing one or more functions based on the determined one or more causal portions of the image. This system embodiment may be further configured as described herein.

Each of the embodiments of each of the systems described above may be combined together into one single embodiment.

Another embodiment relates to a computer-implemented method for performing diagnostic functions for a deep learning model. The method includes determining information from an image generated for a specimen by an imaging tool by inputting the image to a deep learning model. The method also includes determining one or more causal portions of the image that resulted in the information being determined by inputting the information to a diagnostic component. In addition, the method includes performing one or more functions based on the determined one or more causal portions of the image with the diagnostic component. The deep learning model and the diagnostic component are included in one or more components executed by one or more computer systems.

Each of the steps of the method may be performed as described further herein. The method may also include any other step(s) that can be performed by the system, computer subsystem(s), and/or imaging tools described herein. The one or more computer systems and the one or more components may be configured according to any
5 of the embodiments described herein, e.g., computer subsystem(s) 102, component(s) 100, deep learning model 104, and diagnostic component 106. In addition, the method described above may be performed by any of the system embodiments described herein.

An additional embodiment relates to a non-transitory computer-readable medium
10 storing program instructions executable on one or more computer systems for performing a computer-implemented method for performing diagnostic functions for a deep learning model. One such embodiment is shown in Fig. 5. In particular, as shown in Fig. 5, non-transitory computer-readable medium 500 includes program instructions 502 executable on computer system(s) 504. The computer-implemented method may include any step(s)
15 of any method(s) described herein.

Program instructions 502 implementing methods such as those described herein may be stored on computer-readable medium 500. The computer-readable medium may be a storage medium such as a magnetic or optical disk, a magnetic tape, or any other
20 suitable non-transitory computer-readable medium known in the art.

The program instructions may be implemented in any of various ways, including procedure-based techniques, component-based techniques, and/or object-oriented techniques, among others. For example, the program instructions may be implemented
25 using ActiveX controls, C++ objects, JavaBeans, Microsoft Foundation Classes (“MFC”), SSE (Streaming SIMD Extension) or other technologies or methodologies, as desired.

Computer system(s) 504 may be configured according to any of the embodiments described herein.

30

Further modifications and alternative embodiments of various aspects of the invention will be apparent to those skilled in the art in view of this description. For example, methods and systems for performing diagnostic functions for a deep learning model are provided. Accordingly, this description is to be construed as illustrative only
5 and is for the purpose of teaching those skilled in the art the general manner of carrying out the invention. It is to be understood that the forms of the invention shown and described herein are to be taken as the presently preferred embodiments. Elements and materials may be substituted for those illustrated and described herein, parts and processes may be reversed, and certain features of the invention may be utilized
10 independently, all as would be apparent to one skilled in the art after having the benefit of this description of the invention. Changes may be made in the elements described herein without departing from the spirit and scope of the invention as described in the following claims.

WHAT IS CLAIMED IS:

1. A system configured to perform diagnostic functions for a deep learning model, comprising:

5

one or more computer subsystems; and

one or more components executed by the one or more computer subsystems,
wherein the one or more components comprise:

10

a deep learning model configured for determining information from an
image generated for a specimen by an imaging tool; and

15

a diagnostic component configured for determining one or more causal
portions of the image that resulted in the information being
determined and for performing one or more functions based on the
determined one or more causal portions of the image.

20

2. The system of claim 1, wherein the information comprises a classification for a
defect detected on the specimen.

3. The system of claim 1, wherein the information comprises features of the image
extracted by the deep learning model.

25

4. The system of claim 1, wherein the information comprises a simulated image
generated from the image.

5. The system of claim 1, wherein the information comprises one or more
segmentation regions generated from the image.

30

6. The system of claim 1, wherein the information comprises a multi-dimensional output generated from the image.

7. The system of claim 1, wherein the deep learning model is a trained deep learning model.

8. The system of claim 1, wherein the deep learning model is further configured as a neural network.

9. The system of claim 1, wherein the one or more functions comprise altering one or more parameters of the deep learning model based on the determined one or more causal portions.

10. The system of claim 1, wherein the one or more functions comprise determining if the one or more causal portions that resulted in the information being determined are the correct one or more causal portions of the image.

11. The system of claim 1, wherein the one or more components further comprise a visualization component configured for displaying at least the image, the determined information, and the determined one or more causal portions to a user.

12. The system of claim 11, wherein the one or more components further comprise a user interface component configured to receive input from the user after said displaying, and wherein the one or more functions performed by the diagnostic component are determined based on the input from the user.

13. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by computing a local sensitivity.

14. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by causal back propagation.

5 15. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by causal back propagation performed using a deconvolution heatmap algorithm.

10 16. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by causal back propagation performed using a layer-wise relevance propagation.

15 17. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by causal back propagation performed using a deep lift algorithm.

18. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by global average pooling.

20 19. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by computing a path integral on gradients.

20. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by computing a partial dependence plot.

25 21. The system of claim 1, wherein the diagnostic component is further configured for determining the one or more causal portions by computing a partial dependence plot with path integral.

30 22. The system of claim 1, wherein the one or more functions comprise determining one or more characteristics of the one or more causal portions and determining, based on

the one or more characteristics of the one or more causal portions, if additional images for the specimen should be collected from the imaging tool and used for additional training of the deep learning model.

5 23. The system of claim 1, wherein the one or more functions comprise determining one or more characteristics of the one or more causal portions and altering the image based on the one or more characteristics of the one or more causal portions to thereby generate an augmented image for input to the deep learning model.

10 24. The system of claim 1, wherein the one or more functions comprise determining one or more characteristics of the one or more causal portions and generating a data augmentation method for application to additional images input to the deep learning model.

15 25. The system of claim 1, wherein the one or more functions comprise identifying the one or more causal portions as one or more regions of interest in the image and tuning the deep learning model based on the one or more regions of interest.

20 26. The system of claim 1, wherein the one or more functions comprise identifying the one or more causal portions as one or more regions of interest in the image and training an additional deep learning model based on the one or more regions of interest.

25 27. The system of claim 1, wherein the imaging tool is configured as an inspection tool.

28. The system of claim 1, wherein the imaging tool is configured as a metrology tool.

29. The system of claim 1, wherein the imaging tool is configured as an electron beam based imaging tool.

30

30. The system of claim 1, wherein the imaging tool is configured as an optical based imaging tool.

31. The system of claim 1, wherein the specimen is a wafer.

5

32. The system of claim 1, wherein the specimen is a reticle.

33. A system configured to perform diagnostic functions for a deep learning model, comprising:

10

an imaging tool configured for generating images of a specimen;

one or more computer subsystems configured for acquiring the images; and

15

one or more components executed by the one or more computer subsystems, wherein the one or more components comprise:

a deep learning model configured for determining information from an image generated for the specimen by the imaging tool; and

20

a diagnostic component configured for determining one or more causal portions of the image that resulted in the information being determined and for performing one or more functions based on the determined one or more causal portions of the image.

25

34. A non-transitory computer-readable medium, storing program instructions executable on one or more computer systems for performing a computer-implemented method for performing diagnostic functions for a deep learning model, wherein the computer-implemented method comprises:

30

determining information from an image generated for a specimen by an imaging tool by inputting the image to a deep learning model;

5 determining one or more causal portions of the image that resulted in the information being determined by inputting the information to a diagnostic component; and

10 performing one or more functions based on the determined one or more causal portions of the image with the diagnostic component, wherein the deep learning model and the diagnostic component are included in one or more components executed by the one or more computer systems.

35. A computer-implemented method for performing diagnostic functions for a deep learning model, comprising:

15 determining information from an image generated for a specimen by an imaging tool by inputting the image to a deep learning model;

20 determining one or more causal portions of the image that resulted in the information being determined by inputting the information to a diagnostic component; and

25 performing one or more functions based on the determined one or more causal portions of the image with the diagnostic component, wherein the deep learning model and the diagnostic component are included in one or more components executed by one or more computer systems.

1/5




Fig. 1




Fig. 2

3/5




Fig. 3

4/5




Fig. 4

5/5




Fig. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/052093**A. CLASSIFICATION OF SUBJECT MATTER****H01L 21/66(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H01L 21/66; G06F 15/00; G06N 3/08; G06N 3/04; G06K 9/00; G01N 21/88; H01J 37/26; G06T 7/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: deep learning, diagnostic, defect, causal portion

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 9324022 B2 (SIGNALSENSE, INC.) 26 April 2016 See column 8, lines 13-22, column 16, lines 25-43, column 31, line 40 - column 38, line 28 and claims 1, 2, 9.	1-35
Y	US 9171364 B2 (KLA-TENCOR CORPORATION) 27 October 2015 See column 2, lines 14-44, column 14, lines 35-63 and claims 1, 17, 18.	1-35
A	US 2016-0025648 A1 (KLA-TENCOR CORPORATION) 28 January 2016 See paragraphs [0029], [0047], [0053], [0065], [0105] and claims 1, 29.	1-35
A	US 2009-0037134 A1 (ASHOK KULKARNI et al.) 05 February 2009 See claims 1-3.	1-35
A	WO 2015-100099 A1 (KLA-TENCOR CORPORATION) 02 July 2015 See pages 12, 13.	1-35

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

19 December 2017 (19.12.2017)

Date of mailing of the international search report

20 December 2017 (20.12.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

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



Facsimile No. +82-42-481-8578

Authorized officer

CHOI, Sang Won

Telephone No. +82-42-481-8291



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/052093

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 9324022 B2	26/04/2016	US 2015-0254555 A1 WO 2015-134665 A1	10/09/2015 11/09/2015
US 9171364 B2	27/10/2015	CN 105453245 A JP 2016-524334 A KR 10-2016-0022377 A TW 201510516 A US 2014-0376802 A1 WO 2014-205385 A1	30/03/2016 12/08/2016 29/02/2016 16/03/2015 25/12/2014 24/12/2014
US 2016-0025648 A1	28/01/2016	CN 106537125 A DE 112015003362 T5 KR 10-2017-0033875 A SG 11201610822 A TW 201611146 A WO 2016-014714 A1	22/03/2017 13/04/2017 27/03/2017 27/02/2017 16/03/2016 28/01/2016
US 2009-0037134 A1	05/02/2009	IL 203250 A IL 203250 B JP 2010-535430 A JP 5662146 B2 KR 10-1381300 B1 KR 10-2010-0046240 A US 8611639 B2 WO 2009-018337 A1	31/10/2010 31/03/2015 18/11/2010 28/01/2015 04/04/2014 06/05/2010 17/12/2013 05/02/2009
WO 2015-100099 A1	02/07/2015	TW 201531698 A US 2015-0179400 A1 US 9293298 B2	16/08/2015 25/06/2015 22/03/2016