

SYSTEMS AND METHODS FOR TRAINING GENERATIVE ADVERSARIAL NETWORKS AND USE OF TRAINED GENERATIVE ADVERSARIAL NETWORKS

TECHNICAL FIELD

[01] The present disclosure relates generally to the field of neural networks and the use of such networks for image analysis and object detection. More specifically, and without limitation, this disclosure relates to computer-implemented systems and methods for training generative adversarial networks and using the same. The systems and methods and trained neural networks disclosed herein may be used in various applications and vision systems, such as medical image analysis and systems that benefit from accurate object detection capabilities.

BACKGROUND

[02] In many object detection systems, an object is detected in an image. An object of interest may be a person, place, or thing. In some applications, such as medical image analysis and diagnosis, the location of the object is important as well. However, computer-implemented systems that utilize image classifiers are typically unable to identify or provide the location of a detected object. Accordingly, extant systems that only use image classifiers are not very useful.

[03] Furthermore, training techniques for object detection may rely on manually annotated training sets. Such annotations are time-consuming when the detection network being trained is one that is bounding box-based, such as a You Only Look Once (YOLO) architecture, a Single Shot Detector (SSD) architecture, or the like. Accordingly, large datasets are difficult to annotate for training, often resulting in a neural network that is trained on a smaller dataset, which decreases accuracy.

[04] For computer-implemented systems, extant medical imaging is usually built on a single detector network. Accordingly, once a detection is made, the network simply outputs the detection, e.g., to a physician or other health care professional. However, such detections may be false positives, such as non-polyps in endoscopy or the like. Such systems do not provide a separate network for differentiating false positives from true positives.

[05] Furthermore, object detectors based on neural networks usually feed features identified by a neural network into the detector, which may comprise a second neural network. However, such networks are often inaccurate because feature detection is performed by a generalized network, with only the detector portion being specialized.

[06] Finally, many extant object detectors function on a delay. For example, medical images may be captured and stored before analysis. However, some medical procedures, such as endoscopy, are diagnosed on a real-time basis. Consequently, these systems are usually difficult to apply in the required real-time fashion.

SUMMARY

[07] In view of the foregoing, embodiments of the present disclosure provide computer-implemented systems and methods for training a generative adversarial network and using the same for applications such as medical image analysis. The systems and methods of the present disclosure provide benefits over extant system and techniques, including improved object detection and location information.

[08] In accordance with some embodiments, a computer-implemented system is provided that includes an object detector network that identifies features-

of-interest (i.e., abnormalities or objects of interest), along with locations thereof, and an adversarial network that distinguishes true positives from false positives.

Moreover, embodiments of the present disclosure provide a two-loop technique for training the object detector network. This training process uses annotations based on reviewing a detection, such that manual annotation may occur much faster and therefore with a larger dataset. Moreover, this process may be used to train a generative adversarial network to distinguish false positives from true positives.

[09] In addition, disclosed systems are provided that combine an object detector network with a generative adversary network. By combining such networks, false positives may be distinguished from true positives, thereby providing more accurate outputs. By reducing false positives, a physician or other health care professional may give increased attention to outputs from the network on account of the increased accuracy.

[010] Furthermore, embodiments of the present disclosure include neural networks that do not use generic feature identification by one neural network combined with a specialized detector. Rather, a single, seamless neural network is trained for the object detector portion, which results in greater specialization, as well as increased accuracy and efficiency.

[011] Finally, embodiments of the present disclosure are configured for displaying real-time video (such as endoscopy video or other medical images) along with object detections on a single display. Accordingly, embodiments of the present disclosure provide a video bypass to minimize potential problems from errors and other potential drawbacks with the object detector. Moreover, the object detections may be displayed in particularized ways designed to better draw the attention of the physician or other health care professional.

[012] In one embodiment, a system for training a generative adversarial network using images including representations of a feature-of-interest may comprise at least one memory configured to store instructions and at least one processor configured to execute the instructions to perform operations. The operations may comprise: provide a first plurality of images that include representations of the feature-of-interest and indicators of the locations of the feature-of-interest in images of the first plurality of images, and using the first plurality of images and indicators of the feature-of-interest, train an object detection network to detect the feature-of-interest. The operations may further comprise: provide a second plurality of images that include representations of the feature-of-interest, and apply the trained object detection network to the second plurality of images to produce a first plurality of detections of the feature-of-interest. The second plurality of images may comprise a larger number of images than that included in the first plurality of images. The operations may further comprise: provide manually set verifications of true positives and false positives with respect to the first plurality of detections; using the verifications of the true positives and false positives with respect to the first plurality of detections, train a generative adversarial network; and retrain the generative adversarial network using at least one further set of images and detections of the feature-of-interest, together with further manually set verifications of true positives and false positives with respect to the further detections of the feature-of-interest.

[013] In some embodiments, the at least one processor may be further configured to retrain the generative adversarial network by providing verifications of false negatives for missed detections of the feature-of-interest in two or more images.

[014] In any of the embodiments, the object detection network may be a convolutional neural network.

[015] In any of the embodiments, the number of images in the second plurality of images may be at least 100 times larger than that included in the first plurality of images.

[016] In any of the embodiments, the first plurality of images and the second plurality of images may comprise medical images. For example, the medical images may comprise images of a gastro-intestinal organ.

[017] In any of the embodiments, at least one the first plurality of images and the second plurality of images comprise images from an endoscopy device. Additionally or alternatively, at least one the first plurality of images and the second plurality of images may comprise images from imaging device used during at least one of a gastroscopy, a colonoscopy, an enteroscopy, or an upper endoscopy, such as an esophagus endoscopy.

[018] In any of the embodiments, the feature-of-interest may be an abnormality. For example, the abnormality may comprise a change in human tissue from one type of cell to another type of cell. Additionally or alternatively, the abnormality may comprise an absence of human tissue from a location where the human tissue is expected. Additionally or alternatively, the abnormality may comprise a formation on or of human tissue.

[019] In any of the embodiments, the abnormality may comprise a lesion. For example, the lesion may comprise a polypoid lesion or a non-polypoid lesion.

[020] In one embodiment, a method for training a neural network system to detect abnormalities in images of a human organ may comprise storing, in a database, a plurality of videos including representations of abnormalities; selecting a

first subset of the plurality of videos; and applying a perception branch of an object detection network to frames of the first subset of the plurality of videos to produce a first plurality of detections of abnormalities. The method may further comprise selecting a second subset of the plurality of videos; and using the first plurality of detections and frames from the second subset of the plurality of videos, training a generator network to generate a plurality of artificial representations of abnormalities. The plurality of artificial representations may be generated through residual learning. The method may further comprise training an adversarial branch of the discriminator network to differentiate between the artificial representations of the abnormalities and true representations of abnormalities; applying the adversarial branch of the discriminator network to the plurality of artificial representations to produce difference indicators between the artificial representations of abnormalities and true representations of abnormalities included in frames of the second subset of plurality of videos; applying the perception branch of the discriminator network to the artificial representations to produce a second plurality of detections of the abnormalities; and retraining the perception branch based on the difference indicators and the second plurality of detections. These steps may be performed by at least one processor.

[021] In some embodiments, the abnormality may comprise a change in human tissue from one type of cell to another type of cell. Additionally or alternatively, the abnormality may comprise an absence of human tissue from a location where the human tissue is expected. Additionally or alternatively, the abnormality may comprise a formation on or of human tissue.

[022] In any of the embodiments, the abnormality may comprise a lesion. For example, the lesion may comprise a polypoid lesion or a non-polypoid lesion.

[023] In any of the embodiments, each artificial representation may provide a false representation of an abnormality that is highly similar to a true representation of an abnormality.

[024] In any of the embodiments, the generator network may comprise a generative adversarial network.

[025] In any of the embodiments, the discriminator network may comprise a convolutional neural network.

[026] In one embodiment, a system for detecting a feature-of-interest in images of human organ may comprise at least one memory storing instructions and at least one processor configured to execute the instructions to perform operations. The operations may comprise: select frames from a video of a human organ; apply a trained neural network system to the frames to produce at least one detection of the feature-of-interest; generate an indicator of a location of the at least one detection on one of the frames; re-encode the frames into a video; and output the re-encoded video with the indicator. The neural network system may be trained according to any of the embodiments set forth above.

[027] Additional objects and advantages of the present disclosure will be set forth in part in the following detailed description, and in part will be obvious from the description, or may be learned by practice of the present disclosure. The objects and advantages of the present disclosure will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims.

[028] It is to be understood that the foregoing general description and the following detailed description are exemplary and explanatory only, and are not restrictive of the disclosed embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

[029] The accompanying drawings, which comprise a part of this specification, illustrate several embodiments and, together with the description, serve to explain the principles and features of the disclosed embodiments. In the drawings:

[030] FIG. 1 is a schematic representation of an exemplary computer-implemented system for overlaying object detections on a video feed, according to embodiments of the present disclosure.

[031] FIG. 2 is an exemplary two phase training loop for an object detection network, according to embodiments of the present disclosure.

[032] FIG. 3 is a flowchart of an exemplary method for training an object detection network, according to embodiments of the present disclosure.

[033] FIG. 4 is a schematic representation of an exemplary object detector with a discriminator network and a generative network, according to embodiments of the present disclosure.

[034] FIG. 5 is a flowchart of an exemplary method for detecting a feature-of-interest using a discriminator network and a generator network, according to embodiments of the present disclosure.

[035] FIG. 6 is a schematic representation of a computer-implemented system using an object detector network, according to embodiments of the present disclosure.

[036] FIG. 7 is a flowchart of an exemplary method for overlaying object indicators on a video feed using an object detector network, according to embodiments of the present disclosure.

[037] FIG. 8A is an example of a display with an overlay for object detection in a video, according to embodiments of the present disclosure.

[038] FIG. 8B is another example of a display with an overlay for object detection in a video, according to embodiments of the present disclosure.

[039] FIG. 8C is an example of a display with an overlay for object detection in a video, according to embodiments of the present disclosure.

DETAILED DESCRIPTION

[040] The disclosed embodiments relate to computer-implemented systems and methods for training generative adversarial networks and using the same. Advantageously, the exemplary embodiments can provide improved trained networks and fast and efficient object detection. Embodiments of the present disclosure can also provide improved object detection for medical image analysis, with reduced false positives.

[041] Embodiments of the present disclosure may be implemented and used in various applications and vision systems. For example, embodiments of the present disclosure may be implemented for medical image analysis systems and other types of systems that benefit from object detection where the objects may be true or false positives. Although embodiments of the present disclosure are described herein with general reference to medical image analysis and endoscopy, it will be appreciated that the embodiments may be applied to other medical image procedures, such as gastroscopy, colonoscopy, enteroscopy, and upper endoscopy, such as esophagus endoscopy. Further, embodiments of the present disclosure are not limited to other environments and vision systems, such as those for or including LIDAR, surveillance, auto-piloting, and other imaging systems.

[042] According to an aspect of the present disclosure, a computer-implemented system is provided for training a generative adversarial network using images including representations of a feature-of-interest. The system may include at least one memory configured to store instructions and at least one processor configured to execute the instructions (see, e.g., FIGs. 1 and 6). The at least one processor may provide a first plurality of images. For example, the at least one processor may extract the first plurality of images from one or more databases. Additionally or alternatively, the first plurality of images may comprise a plurality of frames extracted from one or more videos.

[043] As used herein, the term “image” refers to any digital representation of a scene or field of view. The digital representation may be encoded in any appropriate format, such as Joint Photographic Experts Group (JPEG) format, Graphics Interchange Format (GIF), bitmap format, Scalable Vector Graphics (SVG) format, Encapsulated PostScript (EPS) format, or the like. Similarly, the term “video” refers to any digital representation of a scene or area of interest comprised of a plurality of images in sequence. The digital representation may be encoded in any appropriate format, such as a Moving Picture Experts Group (MPEG) format, a flash video format, an Audio Video Interleave (AVI) format, or the like. In some embodiments, the sequence of images may be paired with audio.

[044] The first plurality of images may include representations of the feature-of-interest (i.e., an abnormality or object of interest) and indicators of the locations of the feature-of-interest in images of the first plurality of images. For example, the feature-of-interest may comprise an abnormality on or of human tissue. In some embodiments, the feature-of-interest may comprise an object, such as a vehicle, person, or other entity.

[045] In accordance with the present disclosure, an “abnormality” may include a formation on or of human tissue, a change in human tissue from one type of cell to another type of cell, and/or an absence of human tissue from a location where the human tissue is expected. For example, a tumor or other tissue growth may comprise an abnormality because more cells are present than expected. Similarly, a bruise or other change in cell type may comprise an abnormality because blood cells are present in locations outside of expected locations (that is, outside the capillaries). Similarly, a depression in human tissue may comprise an abnormality because cells are not present in an expected location, resulting in the depression.

[046] In some embodiments, an abnormality may comprise a lesion. Lesions may comprise lesions of the gastro-intestinal mucosa. Lesions may be histologically classified (e.g., per the Vienna classification), morphologically classified (e.g., per the Paris classification), and/or structurally classified (e.g., as serrated or not serrated). The Paris classification includes polypoid and non-polypoid lesions. Polypoid lesions may comprise protruded, pedunculated and protruded, or sessile lesions. Non-polypoid lesions may comprise superficial elevated, flat, superficial shallow depressed, or excavated lesions.

[047] In regards to detected abnormalities, serrated lesions may comprise sessile serrated adenomas (SSA); traditional serrated adenomas (TSA); hyperplastic polyps (HP); fibroblastic polyps (FP); or mixed polyps (MP). According to the Vienna classification, an abnormality is divided into five categories, as follows: (Category 1) negative for neoplasia/dysplasia; (Category 2) indefinite for neoplasia/dysplasia; (Category 3) non-invasive low grade neoplasia (low grade adenoma/dysplasia); (Category 4) mucosal high grade neoplasia, such as high grade adenoma/dysplasia, non-invasive carcinoma (carcinoma in-situ), or suspicion of invasive carcinoma; and

(Category 5) invasive neoplasia, intramucosal carcinoma, submucosal carcinoma, or the like.

[048] The indicators of the locations of an abnormality or feature-of-interest may comprise points (e.g., coordinates) or regions (e.g., a rectangle, a square, an oval, or any other regular or irregular shape). The indicators may comprise manual annotations on or of the images. In some embodiments, the first plurality of images may comprise medical images, such as images of a gastro-intestinal organ or other organ or area of human tissue. The images may be generated from a medical imaging device, such as those used during an endoscopy, a gastroscopy, a colonoscopy, an enteroscopy, or an upper endoscopy, such as an esophagus endoscopy, procedure. In such embodiments, if the feature-of-interest is a lesion or other abnormality, a physician or other health care professional may annotate the images to place indicators of the abnormality in the images.

[049] The processor(s) of the system may use the first plurality of images and indicators of the feature-of-interest to train an object detection network to detect the feature-of-interest. For example, the object detection network may comprise a neural network with one or more layers configured to accept an image as input and to output an indicator of a location of a feature-of-interest. In some embodiments, the object detection network may comprise a convolutional network.

[050] Training the object detection network may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, weights of the object detection network may be adjusted to minimize a loss function associated with the network. In some embodiments, the loss function may comprise a square loss function, a hinge loss function, a logistic loss function, a cross entropy loss function, or any other

appropriate loss function or combination of loss functions. In some embodiments, activation (or transfer) functions of the object detection network may be modified to improve the fit between one or more models of the node(s) and the input to the node(s). For example, the processor(s) may increase or decrease the power of a polynomial function associated with the node(s), may change the associated function from one type to another (e.g., from a polynomial to an exponential function, from a logarithmic functions to a polynomial, or the like), or perform any other adjustment to the model(s) of the node(s).

[051] The system processor(s) may further provide a second plurality of images that include representations of the feature-of-interest. For example, the processor(s) may extract the first plurality of images from one or more databases, whether the same database(s) that stored the first plurality of images or one or more different databases. Additionally or alternatively, the second plurality of images may comprise a plurality of frames extracted from one or more videos, whether the same video(s) used to extract the first plurality of images or one or more different videos.

[052] In some embodiments, the second plurality of images may comprise medical images, such as images from an endoscopy device. In such embodiments, the feature-of-interest may comprise a lesion or other abnormality.

[053] In some embodiments, the second plurality of images may comprise a larger number of images than that included in the first plurality of images. For example, the second plurality of images may include at least one hundred times more images than the first plurality of images. In some embodiments, the second plurality of images may include the first plurality, at least in part, or may be different images from the first plurality. In embodiments where the second plurality of images are extracted, at least in part, from one or more videos from which at least part of the

first plurality of images were extracted, the second plurality of images may comprise different frames than the first plurality from the same video(s).

[054] The system processor(s) may apply the trained object detection network to the second plurality of images to produce a first plurality of detections of the feature-of-interest. For example, in embodiments where the trained object detection network comprises a neural network, the at least one processor may input the second plurality of images to the network and receive the detections. The detections may comprise indicators of locations of the feature-of-interest in the second plurality of images. If the second plurality of images does not include the feature-of-interest, the indicator may comprise a null indicator or other indicator of no feature-of-interest.

[055] The system processor(s) may further provide manually set verifications of true positives and false positives with respect to the first plurality of detections. For example, the verifications may be extracted from one or more databases or received as input. In embodiments where the feature-of-interest comprises a lesion or other abnormality, the verifications may be entered by a physician or other health care professional. For example, the processor(s) may output the detections for display to the physician or other health care professional and receive the verifications in response to the displayed detections.

[056] The system processor(s) may use the verifications of the true positives and false positives with respect to the first plurality of detections to train a generative adversarial network. For example, a generative branch of the network may be trained to generate artificial representations of the feature-of-interest. Accordingly, the generative branch may comprise a convolutional neural network.

[057] Similar to the object detection network, training the generative branch may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, as explained above, weights of the generative branch may be adjusted to minimize a loss function associated with the network. Additionally or alternatively, activation (or transfer) functions of the generative branch may be modified to improve the fit between one or more models of the node(s) and the input to the node(s).

[058] Moreover, the adversarial branch of the network may be trained to distinguish false positives from true positives based on the manual verifications. For example, the adversarial branch may comprise a neural network accepting an image and one or more corresponding detection as input and producing a verification as output. In some embodiments, the processor(s) may further retrain the generative network by providing verifications of false negatives for missed detections of the feature-of-interest in two or more images. By providing artificial representations from the generative branch as input to the adversarial branch and recursively using output from the adversarial branch, the adversarial branch and generative branch may perform unsupervised learning.

[059] Similar to the generative branch, training the adversarial branch may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, as explained above, weights of the adversarial branch may be adjusted to minimize a loss function associated with the network. Additionally or alternatively, activation (or transfer) functions of the adversarial branch may be modified to improve the fit between one or more models of the node(s) and the input to the node(s).

[060] Accordingly, in embodiments where the feature-of-interest comprises a lesion or other abnormality, the generative branch may be trained to generate representations of non-abnormalities that look similar to abnormalities, and the adversarial branch may be trained to distinguish artificial non-abnormalities from abnormalities in the second plurality of images.

[061] The system processor(s) may retrain the generative adversarial network using at least one further set of images and detections of the feature-of-interest, together with further manually set verifications of true positives and false positives with respect to the further detections of the feature-of-interest. For example, the processor(s) may extract the further set of images from one or more databases, whether the same database(s) that stored the first plurality of images and/or the second plurality of images or one or more different databases. Additionally or alternatively, the further set of images may comprise a plurality of frames extracted from one or more videos, whether the same video(s) used to extract the first plurality of images and /or the second plurality of images or one or more different videos. Similar to training, retraining the adversarial branch may include further adjustments to the weights of one or more nodes of the network and/or further adjustments to the activation (or transfer) functions of one or more nodes of the network.

[062] According to another aspect of the present disclosure, a computer-implemented method is provided for training a neural network system to detect abnormalities in images of a human organ. The method may be implemented by at least one processor (see, e.g., processor 607 of FIG. 6).

[063] According to the exemplary method, the processor(s) may store, in a database, a plurality of videos including representations of abnormalities. For

example, the videos may comprise endoscopy videos. The videos may be encoded in one or more formats, such as a Moving Picture Experts Group (MPEG) format, a flash video format, an Audio Video Interleave (AVI) format, or the like.

[064] The method may further include selecting, with the processor(s), a first subset of the plurality of videos. For example, the processor(s) may randomly select the first subset. Alternatively, the processor(s) may use one or more indices of the database to select the first subset. For example, the processor(s) may select the first subset as videos indexed as including representations of abnormalities.

[065] The method may further include applying, with the processor(s), a perception branch of an object detection network to frames of the first subset of the plurality of videos to produce a first plurality of detections of abnormalities. For example, the object detection network may comprise a neural network trained to accept images as input and to output the first plurality of detections. The first plurality of detections may comprise indicators of locations of abnormalities in the frames, such as a point or a region of a detected abnormality. A lack of an abnormality may result in a null indicator or other indicator of no abnormality. The perception branch may comprise a neural network (e.g., a convolutional neural network) configured to detect abnormalities and output indicators of locations of any detected abnormalities.

[066] The method may further include selecting, with the processor(s), a second subset of the plurality of videos. In some embodiments, the second subset may include, at least in part, the first subset or may be different videos from the first subset.

[067] The method may further include using the first plurality of detections and frames from the second subset of the plurality of videos to train a generator

network to generate a plurality of artificial representations of abnormalities. For example, the generator network may comprise a neural network configured to generate the artificial representations. In some embodiments, the generator network may comprise a convolutional neural network. The plurality of artificial representations may be generated through residual learning.

[068] As explained above, training the generative network may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, as explained above, weights of the generative network may be adjusted to minimize a loss function associated with the network. Additionally or alternatively, activation (or transfer) functions of the generative network may be modified to improve the fit between one or more models of the node(s) and the input to the node(s).

[069] The method may further include training, with the processor(s), an adversarial branch of the discriminator network to differentiate between the artificial representations of the abnormalities and true representations of abnormalities. For example, the adversarial branch may comprise a neural network that accepts representations as input and outputs indications of whether the input representation is artificial or true. In some embodiments, the neural network may comprise a convolutional neural network.

[070] Similar to the generative branch, training the adversarial branch of the discriminator network may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, as explained above, weights of the adversarial branch of the discriminator network may be adjusted to minimize a loss function associated with the network. Additionally or alternatively, activation (or transfer) functions of the

adversarial branch of the discriminator network may be modified to improve the fit between one or more models of the node(s) and the input to the node(s).

[071] The method may further include applying, with the processor(s), the adversarial branch of the discriminator network to the plurality of artificial representations to produce difference indicators between the artificial representations of abnormalities and true representations of abnormalities included in frames of the second subset of plurality of videos. For example, the artificial representations may comprise representations of non-abnormalities that look similar to abnormalities. Accordingly, each artificial representation may provide a false representation of an abnormality that is highly similar to a true representation of an abnormality. The adversarial branch may learn to identify differences between non-abnormalities (the false representations) and abnormalities (the true representations), particularly non-abnormalities that are similar to abnormalities.

[072] The method may further include applying, with the processor(s), the perception branch of the discriminator network to the artificial representations to produce a second plurality of detections of the abnormalities. Similar to the first plurality of detections, the second plurality of detections may comprise indicators of locations of abnormalities in the artificial representations, such as a point or a region of a detected abnormality. A lack of an abnormality may result in a null indicator or other indicator of no abnormality.

[073] The method may further include retraining the perception branch based on the difference indicators and the second plurality of detections. For example, retraining the perception branch may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, as explained above, weights of the perception

branch may be adjusted to minimize a loss function associated with the network. Additionally or alternatively, activation (or transfer) functions of the perception branch may be modified to improve the fit between one or more models of the node(s) and the difference indicators and the second plurality of detections.

[074] The exemplary method of training described above may produce a trained neural network system. The trained neural network system may form part of a system used for detecting a feature-of-interest in images of a human organ (e.g., a neural network system may be implemented as part of overlay device 105 of FIG. 1). For example, such a system may include at least one memory configured to store instructions and at least one processor configured to execute the instructions. The at least one processor may select frames from a video of a human organ. For example, the video may comprise an endoscopy video.

[075] The system processor(s) may apply a trained neural network system to the frames to produce at least one detection of the feature-of-interest. In some embodiments, the feature-of-interest may comprise an abnormality. The at least one detection may include an indicator of a location of the feature-of-interest. For example, the location may comprise a point of or a region including the detected feature-of-interest. The neural network system may have been trained to detect abnormalities as explained above.

[076] In some embodiments, the system processor(s) may further apply one or more additional classifiers and/or neural networks to the detected feature-of-interest. For example, if the feature-of-interest comprises a lesion, the at least one processor may classify the lesion into one or more types (e.g., cancerous or non-cancerous, or the like). Additionally or alternatively, the neural network system

may further output whether the detected feature-of-interest is a false positive or a true positive.

[077] The system processor(s) may generate an indicator of a location of the at least one detection on one of the frames. For example, the location of the feature-of-interest may be extracted from the indicator and a graphic indicator of the location placed on the frame. In embodiments where the location comprises a point, the graphic indicator may comprise a circle, star, or any other shape placed on a point. In embodiments where the location comprises a region, the graphic indicator may comprise a border around the region. In some embodiments, the shape or border may be animated; accordingly, the shape or border may be generated for a plurality of frames such that it track the location of the feature-of-interest across the frames as well as appearing animated when the frames are shown in sequence. As explained further below, the graphic indicator may be paired with other indicators, such as a sound and/or a vibrational indicator.

[078] Any aspect of the indicator may depend on a classification of the feature-of-interest, e.g., into one or more types or as a false or true positive. Accordingly, a color, shape, pattern, or other aspect of the graphical indicator may depend on the classification. In embodiments also using a sound and/or vibrational indicator, a duration, frequency, and/or amplitude of the sound and/or vibration may depend on the classification.

[079] The system processor(s) may re-encode the frames into a video. Accordingly, after generating the (graphic) indicator and overlaying it on the frame(s), the frames may be re-assembled as a video. The processor(s) of the system may thus output the re-encoded video with the indicator.

[080] According to another aspect of the present disclosure, a computer-implemented system (see, e.g., FIGs. 1 and 6) for processing real-time video is described. The system may comprise an input port for receiving real-time video. For example, the input port may comprise a video graphics array (VGA) port, a high-definition multimedia interface (HDMI) port, a digital visual interface (DVI) port, a Serial Digital Interface (SDI), or the like. The real-time video may comprise a medical video. For example, the system may receive the real-time video from an endoscopy device.

[081] The system may further comprise a first bus for transferring the received real-time video. For example, the first bus may comprise a parallel connection or a serial connection and may be wired in a multidrop topology or a daisy chain topology. The first bus may comprise a PCI Express (Peripheral Component Interconnect Express) bus, a Universal Serial Bus (USB), an IEEE 1394 interface (FireWire), or the like.

[082] The system may comprise at least one processor configured to receive the real-time video from the first bus, perform object detection on frames of the received real-time video, and overlay a border indicating a location of at least one detected object in the frames. The processor(s) may perform object detection using a neural network system trained to produce at least one detection of the object. In some embodiments, the at least one object may comprise a lesion or other abnormality. Accordingly, the neural network system may have been trained to detect abnormalities as explained above.

[083] The processor(s) may overlay the border as explained above. For example, the border may surround a region including the object, the region being received with the at least one detection by the processor(s).

[084] The system may further comprise a second bus for receiving the video with the overlaid border. For example, similar to the first bus, the second bus may comprise a parallel connection or a serial connection and may be wired in a multidrop topology or a daisy chain topology. Accordingly, like the first bus, the second bus may comprise a PCI Express (Peripheral Component Interconnect Express) bus, a Universal Serial Bus (USB), an IEEE 1394 interface (FireWire), or the like. The second bus may comprise the same type of bus as the first bus or may comprise a different type of bus.

[085] The system may further comprise an output port for outputting the video with the overlaid border from the second bus to an external display. The output port may comprise a VGA port, an HDMI port, a DVI port, a SDI port, or the like. Accordingly, the output port may be the same type of port as the input port or may be a different type of port.

[086] The system may comprise a third bus for directly transmitting the received real-time video to the output port. The third bus may carry the real-time video from the input port to the output port passively, to be effective even when the overall system is turned off. In some embodiments, the third bus may be the default bus that is active when the overall system is off. In such embodiments, the first and the second bus may be activated when the overall system is activated, and the third bus may be deactivated accordingly. The third bus may be re-activated when the overall system is turned off, or upon receipt of an error signal from the processor(s). For example, if the object detection implemented by the processor malfunctions, the processor(s) may activate the third bus, thereby allowing continued output of the real-time video stream without interruption due to the malfunction.

[087] In some embodiments, the overlaid border may be modified across frames. For example, the overlaid border may comprise a two-dimensional shape that is displayed around a region of the image including the at least one detected object, the border being a first color. After an elapsed period of time, the processor(s) may modify the border to a second color if the at least one detected object is a true positive and to a third color if the at least one detected object is a false positive. Additionally or alternatively, the processor(s) may modify the border based on a classification of the detected object. For example, if the object comprises a lesion or other abnormality, the modification may be based on whether the lesion or formation is cancerous or otherwise abnormal.

[088] In any of the embodiments described above, the overlaid indicator may be paired with one or more additional indicators. For example, the processor(s) may transmit a command to one or more speakers to produce a sound when the at least one object is detected. In embodiments where the border is modified, the processor(s) may transmit the command when the border is modified. In such embodiments, at least one of duration, tone, frequency, and amplitude of the sound may depend on whether the at least one detected object is a true positive or a false positive. Additionally or alternatively, at least one of duration, tone, frequency, and amplitude of the sound may depend on a classification of the detected object.

[089] Additionally or alternatively, the processor(s) may transmit a command to at least one wearable apparatus to vibrate when the at least one object is detected. In embodiments where the border is modified, the processor(s) may transmit the command when the border is modified. In such embodiments, at least one of duration, frequency, and amplitude of the vibration may depend on whether the at least one detected object is a true positive or a false positive. Additionally or

alternatively, at least one of duration, tone, frequency, and amplitude of the vibration may depend on a classification of the detected object.

[090] According to another aspect of the present disclosure, a system for processing real-time video is described. Similar to the processing system described above, the system may comprise an input port for receiving real-time video; at least one processor configured to receive the real-time video from the input port, perform object detection by applying a trained neural network on frames of the received real-time video, and overlay a border indicating a location of at least one detected object in the frames; and an output port for outputting the video with the overlaid border from the processor to an external display.

[091] The system may further comprise an input device for receiving a sensitivity setting from a user. For example, the input device may comprise a knob, one or more buttons, or any other device suitable for receiving one command to increase the setting and another command to decrease the setting.

[092] The system processor(s) may adjust at least one parameter of the trained neural network in response to the sensitivity setting. For example, the processor(s) may adjust one or more weights of one or more nodes of the network to either increase or decrease the number of detections produced by the network, based on the sensitivity setting. Additionally or alternatively, one or more thresholds of the output layer of the network and/or applied to the detections received from the output layer of the network may be increased or decreased in response to the sensitivity setting. Accordingly, if the sensitivity setting is increased, the processor(s) may decrease the threshold(s) such that the number of detections produced by the network is increased. Similarly, if the sensitivity setting is decreased, the

processor(s) may increase the threshold(s) such that the number of detections produced by the network is decreased.

[093] FIG. 1 is a schematic representation of an exemplary system 100 including a pipeline for overlaying object detections on a video feed, consistent with embodiments of the present disclosure. As shown in the example of FIG. 1, system 100 includes an operator 101 who controls image device 103. In embodiments where the video feed comprises a medical video, operator 101 may comprise a physician or other health care professional. Image device 103 may comprise a medical imaging device, such as an X-ray machine, a computed tomography (CT) machine, a magnetic resonance imaging (MRI) machine, an endoscopy machine, or other medical imaging device that produces videos or one or more images of a human body or a portion thereof. Operator 101 may control image device 103 by controlling a capture rate of device 103 and/or a movement of device 103, e.g., through or relative to the human body. In some embodiments, image device 103 may comprise a Pill-Cam™ device or other form of capsule endoscopy device in lieu of an external imaging device, such as an X-ray machine, or an imaging device inserted through a cavity of the human body, such as an endoscopy device.

[094] As further depicted in FIG. 1, image device 103 may transmit the captured video or images to an overlay device 105. Overlay device 105 may comprise one or more processors to process the video, as described above. Also, in some embodiments, operator 101 may control overlay 105 in addition to image device 103, for example, by controlling the sensitivity of an object detector (not shown) of overlay 105.

[095] As depicted in FIG. 1, overlay device 105 may augment the video received from images device 103 and then transmit the augmented video to a

display 107. In some embodiments, the augmentation may comprise the overlaying described above. As further depicted in FIG. 1, overlay device 105 may also be configured to relay the video from image device 103 directly to display 107. For example, overlay device 105 may perform a direct relay under predetermined conditions, such as if an object detector (not shown) included in overlay device 105 malfunctions. Additionally or alternatively, overlay device 105 may perform a direct relay if operator 101 inputs a command to overlay 105 to do so. The command may be received via one or more buttons included on overlay device 105 and/or through an input device such as a keyboard or the like.

[096] FIG. 2 is a schematic representation of a two phase training loop 200 for an object detection network, consistent with embodiments of the present disclosure. Loop 200 may be implemented by one or more processors. As shown in FIG. 2, Phase I of loop 200 may use a database 201 of images including a feature-of-interest. In embodiments where the images comprise medical images, the feature-of-interest may include an abnormality, such as a lesion.

[097] As explained above, database 201 may store individual images and/or one or more videos, each video including a plurality of frames. During Phase I of loop 200, one or more processors may extract a subset 203 of images and/or frames from database 201. The one or more processors may select subset 203 randomly or, at least in part, using one or more patterns. For example, if database 201 stores videos, the one or more processors may select no more than one, two, or the like number of frames from each video included in subset 203.

[098] As further depicted in FIG. 2, feature indicators 205 may comprise annotations to subset 203. For example, the annotations may include a point of or a region including the feature-of-interest. In some embodiments, an operator may

view the video or images and manually input the annotations via an input device (e.g., any combination of a keyboard, mouse, touch screen, and display) to the processor(s). Annotations may be stored as in a data structure separate from the image, in formats such as JSON, XML, text, or the like. For example, in embodiments where the images are medical images, the operator may be a physician or other health care professional. Although depicted as added to subset 203 after extraction, subset 203 may have been annotated before storage in database 201 or at another earlier time. In such embodiments, the one or more processors may select subset 203 by selecting the images in database 201 having feature indicators 205.

[099] Subset 203 together with feature indicators 205 comprise training set 207. The one or more processors may train a discriminator network 209 using training set 207. For example, discriminator network 209 may comprise an object detector network, as described above. As explained further above, training the discriminator network may include adjusting weights of one or more nodes of the network and/or adjusting activation (or transfer) functions of one or more nodes of the network. For example, weights of the object detection network may be adjusted to minimize a loss function associated with the network. In another example, activation (or transfer) functions of the object detection network may be modified to improve the fit between one or more models of the node(s) and the input to the node(s).

[0100] As shown in FIG. 2, during Phase II of loop 200, the one or more processors may extract a subset 211 of images (and/or frames) from database 201. Subset 211 may comprise, at least in part, some or all of the images from subset 203 or may comprise a different subset. In embodiments where subset 203 comprises a

plurality of frames from one or more videos, subset 211 may include adjacent or other frames from one or more of the same videos. Subset 211 may comprise a larger number of images than subset 203, e.g., at least 100 times more images.

[0101] The one or more processors may apply discriminator network 209' (which represents the discriminator network 209 after the training of Phase I is completed) to subset 211 produce a plurality of feature indicators 213. For example, feature indicators 213 may comprise a point of or a region including a feature-of-interest detected by discriminator network 209'.

[0102] As further depicted in FIG. 2, verifications 215 may comprise annotations to feature indicators 213. For example, the annotations may include an indicator of whether each feature indicator is a true positive or a false positive. An image that had no feature-of-interest detected but includes the feature-of-interest may be annotated as a false negative.

[0103] Subset 211 together with feature indicators 213 and verifications 215 comprise training set 217. The one or more processors may train a generative adversarial network 219 using training set 217. For example, generative adversarial network 219 may comprise a generative network and an adversarial network, as described above. Training the generative adversarial network may include training the generative network to produce artificial representations of the feature-of-interest or of a false feature-of-interest that looks similar to a true feature-of-interest and training the adversarial network to distinguish the artificial representations from real representations, e.g., those included in subset 211.

[0104] Although not depicted in FIG. 2, verifications 213 may further be used to retrain discriminator network 209'. For example, weights and/or activation (or transfer) functions of discriminator network 209' may be adjusted to eliminate

detections in images annotated as false positives and/or adjusted to produce detections in images annotated as false negatives.

[0105] FIG. 3 is a flowchart of an exemplary method 300 for training an object detection network. Method 300 may be performed by one or more processors. At step 301 in FIG. 3, at least one processor may provide a first plurality of images that include representations of the feature-of-interest and indicators of the locations of the feature-of-interest in images of the first plurality of images. The indicators may comprise manually set indicators. The manually set indicators may be extracted from a database or received as input from an operator.

[0106] At step 303, the at least one processor may, using the first plurality of images and indicators of the feature-of-interest, train an object detection network to detect the feature-of-interest. For example, the object detection network may be trained as explained above.

[0107] At step 305, the at least one processor may provide a second plurality of images that include representations of the feature-of-interest, the second plurality of images comprising a larger number of images than that included in the first plurality of images. In some embodiments, the second plurality of images may overlap, at least in part, with the first plurality of images. Alternatively, the second plurality of images may consist of different images than those in the first plurality.

[0108] At step 307, the at least one processor may apply the trained object detection network to the second plurality of images to produce a first plurality of detections of the feature-of-interest. In some embodiments, as explained above, the detections may include indicators of locations of detected features-of-interest. For example, the object detection network may comprise a convolutional neural network outputting one or more matrices, each matrix defining coordinates and/or regions of

any detected features-of-interest, optionally with one or more associated confidence scores for each detection.

[0109] At step 309, the at least one processor may provide manually set verifications of true positives and false positives with respect to the first plurality of detections. For example, the at least one processor may extract the manually set verifications from a database or receive them as input from an operator.

[0110] At step 311, the at least one processor may, using the verifications of the true positives and false positives with respect to the first plurality of detections, train a generative adversarial network. For example, the generative adversarial network may be trained as explained above.

[0111] At step 313, the at least one processor may retrain the generative adversarial network using at least one further set of images and detections of the feature-of-interest, together with further manually set verifications of true positives and false positives with respect to the further detections of the feature-of-interest. In some embodiments, the further set of images may overlap, at least in part, with the first plurality of images and/or the second plurality of images. Alternatively, the further set of images may consist of different images than those in the first plurality and those in the second plurality. Step 313 may thus comprise applying the trained object detection network to the further set of images to produce further detections of the feature-of-interest, providing manually set verifications of true positives and false positives with respect to the further detections, and retraining the generative adversarial network using the verifications with respect to the further detections.

[0112] Consistent with the present disclosure, the example method 300 may include additional steps. For example, in some embodiments, method 300 may include retraining the generative adversarial network by providing verifications of

false negatives for missed detections of the feature-of-interest in two or more images. Accordingly, the manually set verifications extracted from a database or received as input, may include verifications of false negatives as well as verifications of true positives and false positives. The false negatives may be used to retrain the generative adversarial network. Additionally or alternatively, the false negatives may be used to retrain the object detection network.

[0113] FIG. 4 is a schematic representation of an object detector 400. Object detector 400 may be implemented by one or more processors. As shown in FIG. 4, object detector 400 may use a database 401 of videos including a feature-of-interest. In embodiments where the images comprise medical image, the feature-of-interest may include an abnormality, such as a lesion. In the example of FIG. 4, database 401 comprises a database of endoscopy videos.

[0114] As further depicted in FIG. 4, detector 400 may extract a subset 403 of videos from database 401. As explained above with respect to FIG. 2, subset 403 may be selected randomly and/or using one or more patterns. Detector 400 may apply a perception branch 407 of a discriminator network 405 to frames of subset 403. Perception branch 407 may comprise an object detection network, as described above. Perception branch 407 may have been trained to detect the feature-of-interest and identify a location (e.g., a point or a region) associated with a detected feature-of-interest. For example, perception branch 407 may detect abnormalities and output bounding boxes including the detected abnormalities.

[0115] As shown in FIG. 4, perception branch 407 may output detections 413. As explained above, detections 413 may include points or regions identifying locations of detected features-of-interest in subset 403. As further depicted in FIG. 4, detector 400 may extract a subset 411 of videos from database 401. For

example, subset 411 may overlay, at least in part, with subset 403 or consist of different videos. Subset 411 may have a larger number of videos than subset 403, e.g., at least 100 times more videos. Detector 400 may use subset 411 and detections 413 to train a generative network 415. Generative network 415 may be trained to produce artificial representations 417 of the feature-of-interest, e.g., abnormalities. Artificial representations 417 may comprise false representations of the feature-of-interest that look similar to true representations of the feature-of-interest. Accordingly, generative network 415 may be trained to fool perception branch 407 into making detections that are false positives.

[0116] As further depicted in FIG. 4, generative network 415, once trained, may produce artificial representations 417. Detector 400 may use artificial representations 417 to train an adversarial branch 409 of discriminator network 405. As described above, adversarial branch 409 may be trained to distinguish artificial representations 417 from subset 411. Accordingly, adversarial branch 409 may determine difference indicators 419. Difference indicators 419 may represent any feature vectors or other aspects of an image that are present in artificial representations 417 but not in subset 411, present in subset 411 but not in artificial representations 417, or subtractive vectors or other aspects representing differences between feature vectors or other aspects of artificial representations 417 and those of subset 411.

[0117] As depicted in FIG. 4, detector 400 may retrain perception branch 407 using difference indicators 419. For example, in embodiments where artificial representations 417 comprise false representations of the feature-of-interest, detector 400 may retrain perception branch 407 such that the false representations do not result in detections true representation in subset 411.

[0118] Although not depicted in FIG. 4, detector 400 may further use recursive training to improve generative network 415, perception branch 407, and/or adversarial branch 409. For example, detector 400 may retrain generator network 415 using difference indicators 419. Accordingly, the output of adversarial branch 409 may be used to retrain generator network 415 such that the artificial representations look even more similar to true representations. Additionally, retrained generator network 415 may produce a new set of artificial representations used to retrain adversarial branch 409. Accordingly, adversarial branch 409 and generator network 415 may engage in unsupervised learning, the output of each being used to retrain the other in a recursive manner. This recursive training may be repeated until a threshold number of cycles has been reached and/or until a loss function associated with generator network 415 and/or a loss function associated with adversarial branch 409 reaches a threshold. Moreover, during this recursive training, perception branch 407 may also be retrained using each new output of difference indicators, such that a new subset with new detections may be used to further retrain generator network 415.

[0119] FIG. 5 is a flowchart of an exemplary method 500 for detecting a feature-of-interest using a discriminator network and a generator network. Method 500 may be performed by one or more processors.

[0120] At step 501 in FIG. 5, at least one processor may store, in a database, a plurality of videos including representations of a feature-of-interest, such as abnormalities. For example, the videos may have been captured during endoscopy procedures. As part of step 501, the at least one processor may further select a first subset of the plurality of videos. As explained above, the at least one processor may select randomly and/or using one or more patterns.

[0121] At step 503, the at least one processor may apply a perception branch of an object detection network to frames of the first subset of the plurality of videos to produce a first plurality of detections of abnormalities. In some embodiments, as explained above, the detections may include indicators of locations of detected abnormalities. Also, in some embodiments the perception branch may comprise a convolutional neural network, as explained above.

[0122] At step 505, the at least one processor may select a second subset of the plurality of videos. As explained above, the at least one processor may select randomly and/or using one or more patterns. Using the first plurality of detections and frames from the second subset of the plurality of videos, the at least one processor may further train a generator network to generate a plurality of artificial representations of abnormalities, the plurality of artificial representations being generated through residual learning. As explained above, each artificial representation provides a false representation of an abnormality that is highly similar to a true representation of an abnormality.

[0123] At step 507, the at least one processor may train an adversarial branch of the discriminator network to differentiate between the artificial representations of the abnormalities and true representations of abnormalities. For example, as explained above, the adversarial branch may be trained to identify differences between the artificial representations and the true representations in the frames. In some embodiments, the adversarial branch may comprise a convolutional neural network, as explained above.

[0124] At step 509, the at least one processor may apply the adversarial branch of the discriminator network to the plurality of artificial representations to produce difference indicators between the artificial representations of abnormalities

and true representations of abnormalities included in frames of the second subset of plurality of videos. For example, as explained above, the difference indicators may represent any feature vectors or other aspects of an image that are present in the artificial representations but not in the frames, are present in the frames but not in the artificial representations, or are subtractive vectors or other aspects representing differences between feature vectors or other aspects of the artificial representations and those of the frames

[0125] At step 511, the at least one processor may apply the perception branch of the discriminator network to the artificial representations to produce a second plurality of detections of the abnormalities. Similar to the first plurality of detections, the detections may include indicators of locations of detected abnormalities in the artificial representations.

[0126] At step 513, the at least one processor may retrain the perception branch based on the difference indicators and the second plurality of detections. For example, in embodiments where each artificial representation provides a false representation of an abnormality that is highly similar to a true representation of an abnormality, the at least one processor may retrain the perception branch to decrease the number of detections returned from the artificial representations and, accordingly, to increase the number of null indicators or other indicators of no abnormality returned from the artificial representations.

[0127] Consistent with the present disclosure, the example method 500 may include additional steps. For example, in some embodiments, method 500 may include retraining the generative network based on the difference indicators. In such embodiments, method 500 may further include applying the generative network to generate a further plurality of artificial representations of abnormalities and retraining

the adversarial branch based on the further plurality of artificial representations of abnormalities. Such retraining steps may be recursive. For example, method 500 may include applying the retrained adversarial branch to the further plurality of artificial representations to produce further difference indicators between the further artificial representations of abnormalities and true representations of abnormalities included in frames of the second subset of plurality of videos and retraining the generative network based on the further difference indicators. As explained above, this recursive retraining may be repeated until a threshold number of cycles has been reached and/or until a loss function associated with the generative network and/or a loss function associated with the adversarial branch reaches a threshold.

[0128] FIG. 6 is a schematic representation of a system 600 comprising a hardware configuration for a video feed, consistent with embodiments of the present disclosure. As shown in FIG. 6, system 600 may be communicably coupled to an image device 601, such as a camera or other device outputting a video feed. For example, image device 601 may comprise a medical imaging device, such as CT scanner, an MRI machine, an endoscopy device, or the like. System 600 may further be communicably coupled to a display 615 or other device for displaying or storing video. For example, display 615 may comprise a monitor, screen, or other device for displaying images to a user. In some embodiments, display 615 may be replaced with or supplemented by a storage device (not shown) or a network interface controller (NIC) communicably connected to a cloud-based storage system (also not shown).

[0129] As further depicted in FIG. 6, system 600 may include an input port 603 for receiving the video feed from camera 601, as well as an output port 611 for

outputting video to display 615. As explained above, input port 603 and output port 611 may comprise VGA ports, HDMI ports, DVI ports, or the like.

[0130] System 600 further includes a first bus 605 and a second bus 613. As shown in FIG. 6, first bus 605 may transmit video received through input port 603 through at least one processor 607. For example, processor(s) 607 may implement any of the object detector networks and/or discriminator networks described above. Accordingly, processor(s) 607 may overlay one or more indicators, e.g., the exemplary graphical indicator of FIG. 8, on the video received via first bus 602, e.g., by using the exemplary method 700 of FIG. 7. Processor 607 may then transmit the overlaid video via a third bus 609 to output port 611.

[0131] In certain circumstances, the object detector implemented by processor(s) 607 may malfunction. For example, the software implementing the object detector may crash or otherwise stop functioning properly. Additionally or alternatively, processor(s) 607 may receive a command to halt overlaying the video (e.g., from an operator of system 600). In response to the malfunction and/or the command, processor(s) 607 may activate second bus 613. For example, processor(s) 607 may send a command or other signal, as depicted in FIG. 6, to activate second bus 613.

[0132] As depicted in FIG. 6, second bus 613 may transmit received video directly from input port 603 to output port 611, thereby allowing system 600 to function as a pass-through for image device 601. Second bus 613 may allow for seamless presentation of video from image device 601 even if software implemented by processor 607 malfunctions or if an operator of hardware overlay 600 decides to halt the overlaying in the middle of the video feed.

[0133] FIG. 7 is a flowchart of an exemplary method 700 for overlaying object indicators on a video feed using an object detector network, consistent with embodiments of the present disclosure. Method 700 may be performed by one or more processors. At step 701 in FIG. 7, at least one processor may provide at least one image. For example, the at least one image may be extracted from a database or received from an imaging device. In some embodiments, the at least one image may comprise a frame within the video feed.

[0134] At step 703, the at least one processor may overlay a border comprising a two-dimensional shape around a region of the image detected as including the feature-of-interest, the border being rendered in a first color. At step 705, after an elapsed period of time, the at least one processor may modify the border to appear in a second color if the feature-of-interest is a true positive, and to appear in a third color if the feature-of-interest is a false positive. The elapsed period of time may represent a preset period (e.g., a threshold number of frames and/or seconds) and/or may represent an elapsed time between detection of the feature-of-interest and classification thereof as a true or false positive.

[0135] Additionally or alternatively, the at least one processor may modify the border to the second color if the feature-of-interest is classified in a first category, and modify the border to the third color if the feature-of-interest is classified in a second category. For example, if the feature-of-interest is a lesion, the first category may comprise cancerous lesions and the second category may comprise non-cancerous lesions.

[0136] Consistent with the present disclosure, the example method 700 may include additional steps. For example, in some embodiments, method 700 may include transmitting a command to one or more speakers to produce a sound when

the border is modified and/or transmitting a command to at least one wearable apparatus to vibrate when the border is modified. In such embodiments, at least one of duration, tone, frequency, and amplitude of the sound and/or the vibration may depend on whether the at least one detected object is a true positive or a false positive.

[0137] FIG. 8A illustrates an example overlay 801 for object detection in a video, consistent with embodiments of the present disclosure. In the example of FIG. 8A, as well as FIGs. 8B and 8C, the illustrated video samples 800a and 800b are from a colonoscopy procedure. It will be appreciated from the present disclosure, video from other procedures and imaging devices may be utilized when implementing embodiments of the present disclosure. Thus, the video samples 800a and 800b are non-limiting examples of the present disclosure. In addition, by way of example, the video display of FIGs. 8A-8C may be presented on a display device, such as display 107 of FIG. 1 or display 615 of FIG. 6.

[0138] Overlay 801 represents one example of a graphical border used as an indicator for a detected abnormality or feature-of-interest in a video. As shown in FIG. 8A, images 800a and 800b comprise frames of a video including a detected feature-of-interest. Image 800b includes graphical overlay 801 and corresponds to a frame that is further in sequence, or later in time, than image 800a.

[0139] As shown in FIG. 8A, images 800a and 800b comprise video frames from a colonoscopy, and the feature-of-interest comprises a lesion or polyp. In other embodiments, as described above, images from other medical procedures, such as gastroscopy, enteroscopy, upper endoscopy, such as esophagus endoscopy, or the like, may be utilized and overlaid with a graphical indicator, such as overlay 801. In some embodiments, indicator 801 may be overlaid after detection of the abnormality

and an elapse of time (e.g., a particular number of frames and/or seconds between image 800a and image 800b). In the example of FIG. 8A, overlay 801 comprises an indicator in the form of rectangular border with a predetermined pattern (i.e., solid corner angles). In other embodiments, overlay 801 may be a different shape (whether regular or irregular). In addition, overlay 801 may be displayed in a predetermined color, or transition from a first color to another color.

[0140] In the example of FIG. 8A, overlay 801 comprises an indicator with solid corner angles surrounding the detected location of the feature-of-interest in the video frame. Overlay 801 appears in video frame 800b, which may follow in sequence from video frame 800a.

[0141] FIG. 8B illustrates another example of a display with an overlay for object detection in a video, according to embodiments of the present disclosure. FIG. 8B depicts an image 810a (similar to image 800a) and a later image 810b (similar to image 800b) that is overlaid with an indicator 811. In the example of FIG. 8B, overlay 811 comprises a rectangular border with solid lines on all sides. In other embodiments, overlay 811 may be a first color and/or a different shape (whether regular or irregular). In addition, overlay 811 may be displayed in a predetermined color, or transition from a first color to another color. As shown in FIG. 8B, overlay 811 is placed over the detected abnormality or feature-of-interest in the video. Overlay 811 appears in video frame 810b, which may follow in sequence from video frame 810a.

[0142] FIG. 8C illustrates another example of a display with an overlay for object detection in a video, according to embodiments of the present disclosure. FIG. 8C depicts an image 820a (similar to image 800a) and a later image 820b (similar to image 800b) that is overlaid with an indicator 821. In the example of FIG.

8C, overlay 821 comprises a rectangular border with dashed lines on all sides. In other embodiments, indicator 821 may be a different shape (whether regular or irregular). In addition, overlay 821 may be displayed in a predetermined color, or transition from a first color to another color. As shown in FIG. 8C, overlay 821 is placed over the detected abnormality or feature-of-interest in the video. Overlay 821 appears in video frame 820b, which may follow in sequence from video frame 820a.

[0143] In some embodiments, the graphical indicator (i.e., overlay 801, 811, or 821) may change pattern and/or color. For example, the pattern and/or color of the border of the pattern may be modified in response to an elapse of time (e.g., a particular number of frames and/or seconds between image 800a and image 800b, image 810a and image 810b, or image 820a and image 820b). Additionally, or alternatively, the pattern and/or color of the indicator may be modified in response to a particular classification of the feature-of-interest (e.g., if the feature-of-interest is a polyp, a classification of the polyp as cancerous or non-cancerous, etc.). Moreover, the pattern and/or color of the indicator may depend on the classification of the feature-of-interest. Accordingly, the indicator may have a first pattern or color if the feature-of-interest is classified in a first category, a second pattern or color if the feature-of-interest is classified in a second category, etc. Alternatively, the pattern and/or color of the indicator may depend on whether the feature-of-interest is identified as a true positive or a false positive. For example, the feature-of-interest may be detected by an object detector network (or a perception branch of a discriminator network), as described above, resulting in the indicator, but then determined to be a false positive by an adversarial branch or network, as described above, resulting in the indicator being a first pattern or color. The indicator may be

displayed in a second pattern or color if, instead, the feature-of-interest is determined to be a true positive by the adversarial branch or network.

[0144] The foregoing description has been presented for purposes of illustration. It is not exhaustive and is not limited to precise forms or embodiments disclosed. Modifications and adaptations of the embodiments will be apparent from consideration of the specification and practice of the disclosed embodiments. For example, the described implementations include hardware, but systems and methods consistent with the present disclosure can be implemented with hardware and software. In addition, while certain components have been described as being coupled to one another, such components may be integrated with one another or distributed in any suitable fashion.

[0145] Moreover, while illustrative embodiments have been described herein, the scope includes any and all embodiments having equivalent elements, modifications, omissions, combinations (e.g., of aspects across various embodiments), adaptations and/or alterations based on the present disclosure. The elements in the claims are to be interpreted broadly based on the language employed in the claims and not limited to examples described in the present specification or during the prosecution of the application, which examples are to be construed as nonexclusive. Further, the steps of the disclosed methods can be modified in any manner, including reordering steps and/or inserting or deleting steps.

[0146] The features and advantages of the disclosure are apparent from the detailed specification, and thus, it is intended that the appended claims cover all systems and methods falling within the true spirit and scope of the disclosure. As used herein, the indefinite articles “a” and “an” mean “one or more.” Similarly, the use of a plural term does not necessarily denote a plurality unless it is unambiguous

in the given context. Words such as “and” or “or” mean “and/or” unless specifically directed otherwise. Further, since numerous modifications and variations will readily occur from studying the present disclosure, it is not desired to limit the disclosure to the exact construction and operation illustrated and described, and accordingly, all suitable modifications and equivalents may be resorted to, falling within the scope of the disclosure.

[0147] Other embodiments will be apparent from consideration of the specification and practice of the embodiments disclosed herein. It is intended that the specification and examples be considered as example only, with a true scope and spirit of the disclosed embodiments being indicated by the following claims.