

Video Analytics for Industrial Safety

Technical Field

- 5 The present invention relates, in general terms, to systems and methods for improving industrial safety using video analytics.

Background

- 10 Worksites such as construction, building and industrial sites have several safety hazards that present a risk to operational health and safety. Safety breaches may not only impact the health of workers but also the productivity of the worksite. The hazards may originate from various sources such as dangerous machinery or an inherently dangerous work
15 environment such as a mining operation. One approach to addressing such safety hazards is to place IoT sensors on dangerous machinery and workers and trigger an alarm when a worker is in unsafe proximity to the dangerous machinery. The IoT sensor-based approach requires a large number of IoT sensors on large worksites. It further necessitates the
20 calibration and management of a large number of IoT sensors and the networking infrastructure to support operation of the IoT sensors.

It would be desirable to overcome or alleviate at least one of the above-described problems, or at least to provide a useful alternative.

25

Summary

The present disclosure relate to a system for detecting an adverse event, comprising:

one or more image capture devices, each device being positioned to capture images of a respective zone; and

a server comprising a plurality of neural networks, each neural network receiving, and analysing, at least a portion of said images to
5 identify an object or objects in the respective zone, the server being configured to detect if the respective zone comprises an adverse event based on the identified object or objects; and

an alert server for identifying, for each detected adverse event, one or more stakeholders for the respective zone and transmitting an
10 alert to the one or more stakeholders, the alert comprising identification information for the detected adverse event.

The one or more image capture devices may comprise a plurality of image capture devices. Each image capture device may be configured to send
15 video feed comprising the images of the respective zone.

Each image capture device may be positioned based on:

a relationship between a boundary of the respective zone and a field of view of the respective image capture device;
20 a distance from the respective zone, a resolution of the respective image capture device and a size of a smallest said object sought to be identified in the respective zone.

A minimum distance of the image capture device from the respective
25 zone may be determined based on the boundary being wholly within the field of view, and a maximum distance of the image capture device from the respective zone is determined based on the resolution enabling detection of the smallest said object at a furthest point in the field of view.

Each neural network may be configured to identify a different object. At least one neural network may be configured to detect a human, and wherein the server may be configured to if the respective zone comprises an adverse event if a human is identified by the at least one neural network.

5

The server may be configured to detect if the zone comprises an adverse event by calculating a relationship between a human and another identified object.

- 10 The identified object may comprise equipment for use by a human, the server being configured to calculate the relationship by determining if the equipment is present in the images with the at least one human.

The server may be configured to detect that the respective zone
15 comprises an adverse event when either:

the equipment is absent from the images; or

a human and the safety equipment are present in the images, but not in association with one another.

- 20 The equipment and human may be considered to be in association with one another if the equipment is being used as it is intended to be used.

The identification information may comprise at least a location of the adverse event and one or more said images in which the adverse event
25 is visible.

Each neural network may be configured to label said at least a portion of the images analysed by the respective neural network, based on the respective neural network identifying the object or objects. The
30 identification information may comprise said label.

The alert server may be configured to send the alert to the human, the identification information specifying the equipment that is sought to be identified in association with the human.

5

The system may further comprise an intermediary layer, the intermediary layer specifying particular neural networks of the plurality of neural networks for identifying a predetermined said adverse event.

- 10 The system may further comprise an intermediary layer, the intermediary layer specifying particular neural networks of the plurality of neural networks for use with images from each respective image capture device.

The disclosure also relates to a method for detecting an adverse event,
15 comprising:

capturing one or more images using one or more capture devices positioned to capture images of a respective zone; and

receiving at least a portion of said images at a plurality of neural networks, the neural networks being configured to identify an object or
20 objects in the respective zone;

detecting if the respective zone comprises an adverse event based on the identified object or objects

identifying, for each detected adverse event, one or more stakeholders for the respective zone; and

- 25 transmitting an alert to the one or more stakeholders, the alert comprising identification information for the detected adverse event.

The disclosure also relates to a system for detecting an adverse event, comprising:

one or more image capture devices, each device being positioned to capture images of a respective zone; and

a server comprising a plurality of neural networks, each neural network receiving, and analysing, at least a portion of said images to
5 identify an object or objects in the respective zone, the server being configured to detect if the respective zone comprises an adverse event based on the identified object or objects; and

the server further identifying, for each detected adverse event, one or more stakeholders for the respective zone and transmitting an
10 alert to the one or more stakeholders, the alert comprising identification information for the detected adverse event.

Brief description of the drawings

15 Embodiments of the present invention will now be described, by way of non-limiting example, with reference to the drawings in which:

Figure 1 illustrates a system for detecting an adverse event;

20 Figure 2 illustrates a method for detecting an adverse event; and

Figure 3 illustrates an exemplary computer system for implementing the server or alert server of the system for detecting an adverse event.

25 Detailed description

The disclosure relates to systems and methods for detecting adverse events. These systems and methods can be implemented in workplaces and other areas that have visually identifiable safety risks, to reduce

reliance on the alertness of humans that would otherwise monitor those workplaces and other areas. The systems and methods rely on video analytics for adverse event detection. The disclosed systems advantageously enable the automatic detection of hazards or adverse events such as absence of barricades, use of unsuitable equipment, non-use of necessary equipment and other such safety hazards at worksites. The disclosed systems also advantageously generate notifications regarding the detected hazards and transmit the notifications to relevant stakeholders or personnel enabling a timely response to the detected adverse event.

One such system 100 is shown in Figure 1. System 100 comprises one or more image capture devices (cameras) 110, each device being positioned to capture images of a respective zone 105. The video analytics server 120 comprises a plurality of neural networks 127. Each neural network 127 analyses some of the images captured by the camera 110 to identify an object or objects in the respective zone 105 – e.g. each neural network may analyse a subset of the images, or a portion of an image or images in which a particular safety hazard is likely to occur. Each neural network 127 may also be configured to label at least a portion of the images analysed by the respective neural network 127, the label indicating a class identifier of the object identified by the neural network 127. Based on the inferences generated by the neural networks 127, the video analytics server identified various adverse events. Adverse events could be detected based on the presence or absence of objects in images captured by camera 120. Adverse events could also be detected based on the relative position of objects in images captured by camera 120.

System 100 also comprises an alert server 130 for identifying, for each detected adverse event, one or more stakeholders for the respective zone and transmitting an alert to the one or more stakeholder computing

device 140, the alert comprising identification information for the detected adverse event. The identification information comprises a label or class information associated with each object detected in images associated with the adverse event.

- 5 The video analytics server 120 comprises at least one processor 122 in communication with memory 124. Memory 124 comprises a neural networks module 125, an image labelling module 128 and a user interface module 129. The alert server 130 comprises at least one processor 132 and a memory 134. Memory 134 comprises a stakeholder database 136.
- 10 The stakeholder database 136 comprises information including identification of stakeholder computing device identifiers and associations between stakeholders and the one or more respective zones 105. The stakeholder database 136 allows the direction of alerts regarding adverse events to a specific stakeholder computing device 140 belonging to a
- 15 stakeholder responsible for the safety of the respective zone. The alert server 130 and the video analytics server 120 could be implemented as a common server, with the common server performing the functionality of both the alert server 130 and the video analytics server 120. The stakeholder computing device 140 may include an end-user computing
- 20 device such as a smartphone, or personal computing device accessible to a stakeholder responsible for one or more respective zones 105.

Each camera 110 is positioned to capture images of a respective zone 105(1)-105(N). Each camera is configured to send a video feed comprising the images of the respective zone 105 to the video analytics

25 server 120.

The cameras 110 are positioned to capture images with sufficient level of detail or at a sufficient resolution to accurately detect adverse events in the respective zones. The cameras are also positioned such that their

field of view covers all objects or areas of interest for detection of adverse events. Each camera 110 is positioned based on a relationship between a boundary of the respective zone 105 and a field of view of the respective camera 110. Also relevant for positioning of cameras 110 is the distance
5 from the respective zone 105, a resolution of the respective camera 110 and a size of a smallest said object sought to be identified in the respective zone 105.

A minimum distance of the camera 110 from the respective zone 105 is determined based on the boundary of an area of interest being wholly
10 within the field of view. A maximum distance of the camera 110 from the respective zone 105 is determined based on the resolution enabling detection of the smallest said object at a furthest point in the field of view of the camera 110. For example, a camera 110 capturing images of a higher resolution is placed farther away from the observed area while a
15 camera 110 capturing images of a lower resolution is placed closed to the observed area.

The camera 110 may include an IP camera and it streams a video feed from the area of interest (respective zone 105). The respective zone 105 may correspond to a zone where high-risk activities may occur. The high-
20 risk activities may relate to high-risk activities in construction sites. The video footage captured by the camera 110 is transmitted through a network 115 to a Network Video Recorder (NVR) 117. The Network Video Recorder 117 may be located on the same site that is being monitored by system 100. A video analytics server 120 in communication with the
25 Network Video Recorder 117 then processes the video feed with the video analytics algorithms to identify safety hazards. Based on the video analytics performed by the video analytics server 120 when a safety violation or adverse event is observed, an alert is sent to a relevant stakeholder computing device 140 by an alert server 130 operating in

communication with the video analytics server 120. The alert could be transmitted using a messaging platform such as the Telegram messaging platform.

5 The alert may consist of a short video of the detected violation or adverse event with details on the type of violation as well as the location at which the violation was detected. Details of the violations may be stored in the video analytics server 120 or the alert server 130 and may be used to generate dashboards and/or graphs to analyse construction site safety.

10 The system 100 is directed to particular applications that are focused on behavioural safety and industrial safety. The system 100 also generates real-time or near real-time alerts which immediately notify stakeholders such as project managers or safety officers to take action that could potentially prevent a hazard from occurring. The video analytics server 120 comprises neural networks that are specifically trained with relevant and specific to construction images and hazards. The video analytics server 120 also provides a user interface through an image labelling module 128 that allows a user to demarcate or label zones of interest in images of a video feed captured by the camera for the algorithms for adverse event detection to work effectively.

20 A Neural network 127 is configured to detect a human, and the video analytics server 120 could determine if the respective zone 105 comprises an adverse event if a human is identified by the neural network 127 in images of the respective zone 105. The video analytics server 120 could detect if zone 105 comprises an adverse event by calculating a relationship between a human and another identified object detected in images of the respective zone 105. Calculating the relationship between a human and another identified object includes calculating a perceived proximity or distance or overlap of the human to the other identified

25

object based on the images of the respective zone 105. The other identified object may comprise equipment for use by a human. The video analytics server 120 calculates the relationship by determining if the equipment is present in the images with at least one human. For
5 example, the identified object may include a safety harness and the video analytics server 120 could determine whether the human is wearing the safety harness to detect the occurrence of an adverse event.

The video analytics server 120 could detect that the respective zone 105
10 comprises an adverse event when either: a necessary equipment is absent from the images, or a human and the safety equipment are present in the images, but not in association with one another. The equipment and human may be considered to be in association with one another if the equipment is being used as it is intended to be used, for
15 example a safety harness or a mask must be worn by a human. The video analytics server 120 operating in concert with the alert server 130 could also transmit an alert to the human detected in the adverse event as not operating safely, for example operating without the requisite safety equipment.

20 The system 100 detects various adverse events in various industrial or work settings. The following section illustrates examples of some adverse events.

Detection of the presence of barricades

Barricades are present in industrial settings to prevent workers from
25 falling out of heights. System 100 could detect the presence of barricades and triggers sending of an alert when a barricade isn't present when it is supposed to be. The alert can be sent to notify workers or supervisors of workers that the worksite is unsafe and a barricade has not been detected.

Detection of the presence of platform ladders

Platform ladders have a larger surface area such that workers have a more stable area to step onto before performing a task. The system 100 could detect the presence of a platform ladder to evaluate whether a safer
5 option of a ladder is chosen and identify the occurrence of an adverse event if a platform ladder is not detected.

Detection of the presence of A-Frame ladders

A-Frame ladders have a small area to work on and may cause workers to fall off while using the ladder, therefore the system 100 could detect if
10 the right type of ladder is used and identify the occurrence of an adverse event if an unsuitable ladder such as an A-Frame ladder is used.

Detection of the presence of excavators

Excavators are dangerous vehicles in the construction sites and system 100 could detect where excavators are located within a worksite to
15 identify if work is conducted safely and determine the occurrence of adverse events. For example, the detection of an excavator in a restricted zone on a worksite may be identified as an adverse event.

Detection of the presence of safety harness

Workers working at heights always need to have safety harness on.
20 Working without a harness may cause them to fall from heights. The system 100 could detect the non-use or improper use of safety harnesses and identify such non-use or improper use as an adverse event.

Once any of the safety hazards based on the described applications or other such applications are detected by the video analytics server 120,
25 notifications to a stakeholder computing device 140 are sent by the alert server 130. The alerts may be in the form of a Telegram message,

WhatsApp message, WeChat message, Email or SMS or any other suitable messaging or communication technology. The messages could be sent in real-time or near real-time with respect to the occurrence or detection of the adverse event.

- 5 Memory 124 of the video analytics server 124 comprises a neural network module 125 implementing deep learning algorithms that process the video feed or images captured by the camera 110 to detect the occurrence of adverse events. The neural network module 125 comprises a plurality of neural networks 127. Each neural network 127 could detect
- 10 a specific object, such as a person or a safety harness or a specific kind of equipment. The use of neural networks implementing deep learning algorithms provides the flexibility of detection of adverse events associated with various use-cases mentioned above and a high level of accuracy of the detection.
- 15 While each neural network 127 is configured to detect a specific object or a class of objects, the neural network module 125 comprises additional logic or additional neural networks such as the intermediary layer 126 embodying the logic to determine the occurrence of adverse events. Each object detection neural network 127 detects a specific object in an image
- 20 in a video feed captured by the camera 110. The object detection neural network 127 could also determine a bounding box around the detected object in an image. The multiple neural networks 127 detect multiple objects of interest in a single image and define a bounding box around each detected object. The relative position of the detected objects as
- 25 demarcated by the respective bounding boxes is analysed by the neural network module 125 or other logic implemented by the video analytics server 120 to determine the occurrence of an adverse event.

For example, a first neural network 127 could detect a person and a second neural network 127 could detect safety harnesses in a common image. If the bounding boxes identified around the detected person and the safety harness sufficiently overlap, then this overlap is inferred by
5 the neural network module 125 as a safe operation, i.e. no adverse event is detected. However, if the bounding boxes identified around the detected person and the safety harness do not sufficiently overlap or no safety harness is detected, then this condition or outcome is flagged as an adverse event by the neural network module 125.

10 The neural network module 125 is trained using a training dataset comprising images with labelled objects and bounding boxes or segments defining objects of interest in the training images. In addition, the training dataset also includes a label associated with each training image indicating whether the training image corresponds to an adverse event
15 or not. The various neural networks of the neural network module 125 are trained to both detect objects and to classify or identify adverse events based on the examples in the training dataset.

Commercial Applications

This system 100 addresses several critical industrial safety related
20 problems. Firstly, every construction site needs safety officers for observing the safety of the operations. A limited number of trained safety officers may be available to fully monitor a worksite. The safety officers may spend a significant amount of time ensuring safety regulations are complied with instead of focusing on their core task of detection of
25 adverse events. Adverse event detection by the system 100 reduces the workload or responsibilities a safety officer is typically responsible for.

If adverse events are not proactively detected, then they may lead to a safety hazard that may result in the issuance of a stop-work order. Such

stop-work orders may result in the waste of all the hired manpower, machinery and equipment on site. As a result of the stop-work order, there is also an increased chance of delay in the delivery of a project. With the assistance of adverse event detection by the system 100 in place, safety officers can spend their time more effectively. System 100 provides all around the clock safety monitoring and sends real-time alerts to safety managers to take rectification actions before the adverse events result in the issuance of non-compliance or stop-work orders.

Figure 2 illustrates a method 200 for detecting an adverse event that may be performed by the system 100 of Figure 1.

Step 210: Capture Video

Step 210 comprises capturing one or more images using one or more capture devices 110 positioned to capture images of a respective zone 105. One or more images are captured in the form of a video. The video is captured through a video acquisition module which is in the form of a hardware device or camera 110 that is used to capture video. The hardware module consists of a camera that may be placed at an angle of inclination of between 30 degrees to 80 degrees and overlooks a region of high risk in an industrial setting or a worksite. The camera is placed at a distance far enough to view the entire or a substantial part of the high-risk activity that could potentially occur.

The video streams are captured from multiple hardware modules of multiple image capture devices 110 overlooking various types of high-risk activities in industrial sites. They are continually streamed to the video server 120 which performs the analytics. The video streams from various sources are stored in the network video recorder 117. The video analytics server 120 could query the network video recorder 117 to retrieve video streams associated with a specific camera 110 and

captured over a specific period on a particular date. The video streams stored in the network video recorder 117 could be queried on an ad-hoc basis by the stakeholder computing device 140 and the results of the query are transmitted by the Video analytics server to the stakeholder
5 computing device 140. The video feeds captured by the camera 110 could be streamed in real-time or near real-time to the stakeholder computing device 140.

The video stream captured by the camera 110 may be of a resolution equal to or greater than 480p. The video feed may be captured at any
10 predefined FPS in an H.264 format. Examples of hardware modules that capture these kinds of video streams include IP cameras (Hikvision, Dahua, Pelco, etc.), mobile phones, and wearable cameras, etc.

Step 220: Receiving at least a portion of said images at a plurality of neural networks to detect objects of interest

15 At step 220, the at least a portion of the images captured by the cameras 110 are received at the processor 122 of the video analytics server 120 and processed by the neural networks embodied in the neural network module 125 to identify an object or objects in the respective zone.

The video streams originating from the camera 110 are transmitted over
20 a 3G/4G/5G or a WIFI or a wired network to the video analytics server 120 located physically on-site or in the cloud. The video stream may comply with the Real-Time Streaming Protocol (RTSP), Hypertext Transfer Protocol Live Streaming (HLS) Protocol, Session Description Protocol (SDP), or Audio Video Interleave (AVI) Protocol, etc.

25 The video streams are then processed by the video analytics server 120. The video streams first reach an intermediary layer 126 which acts as an

intermediary between the video stream data and the various neural networks 127 present in the neural network module 125.

The intermediary layer 126 specifies one or more neural networks 127 of the plurality of neural networks for identifying an adverse event. The
5 video analytics server 120 comprises the following 14 custom trained exemplary neural networks 127:

1) Neural network for Detection of excavators

This neural network is trained on several custom obtained images from various open-source locations and photos from various industrial sites.

10 The neural network detects the presence of an excavator from any angle and at a distance more than 80m away.

2) Neural network for Detection of workers

This neural network is trained to detect the presence of workers at any location in the industrial site video frame.

15 3) Neural network for Detection of Personal Protective Equipment

This neural network is trained on a dataset that allows the detection of the helmet, safety vest, boots, gloves, and goggles present on a worker detected in the industrial site.

4) Neural network for Detection of safety harnesses

20 This neural network enables the detection of the presence of safety harnesses present on workers working at height.

5) Neural network for Detection of barricades

This neural network enables the detection of the presence of barricades of different colours in the zone of interest which is defined by the front-
25 end user interface of the video analytics platform

6) Neural network for Detection of guardrails

This neural network enables the detection of the presence of guardrails of different colours in the zone of interest which is defined by the front-end user interface of the video analytics platform

5 7) Neural network for Detection of A-Frame ladders

This neural network enables the detection of A-frame ladders present in the video stream.

8) Neural network for Detection of platform ladders

10 This neural network enables the detection of platform ladders present in the video stream.

9) Neural network for Detection of industrial traffic congestion

This neural network enables the detection of various types of traffic conditions present in the industrial site, dense, and sparse traffic as well as detects accidents.

15 10) Neural network for Detection of stagnant water

This neural network enables the detection of the presence of stagnant water in the video stream.

11) Neural network for Detection of potholes

20 This neural network enables the detection of the presence of potholes in the video stream.

12) Neural network for Detection of housekeeping hazards

This neural network enables the detection of the presence of housekeeping hazards such as cluttered piping, and over-stacked materials.

13) Neural network for Detection of forklifts

This neural network is trained on several custom obtained images from various open-source locations and photos from various industrial sites. The neural network detects the presence of a forklift from any angle and
5 at a distance more than 80m away.

14) Neural network for Detection of masks

This neural network enables the detection of masks present on or worn by workers.

The above 14 neural networks 127 are utilized in combination with one
10 another to identify safety risks present in the industrial site. A front-end user interface provided by the user interface module 129 allows a user to select images captured by a specific camera 110 and view which of the above-noted safety hazards have been detected.

15 Step 230: Detecting if the respective zone comprises an adverse event based on the identified object or objects

At step 230, based on an output of the neural networks 127, the video analytics server 120 detects if the respective zone comprises an adverse event based on the object or objects identified by the neural networks 127.

20 The intermediary layer 126 between the input layer receiving the images captured by the cameras 110 and the 14 neural networks acts as the channel through which the safety hazard details are broken down and the relevant neural network is activated to identify the presence of a specific safety hazard. The following is a list of examples of safety hazards or
25 adverse events:

1) Worker is closer than 1m to an excavator or a forklift

- 2) An A-frame ladder is used instead of a platform ladder
- 3) Worker is not wearing personal protective equipment when necessary
- 4) Worker is not wearing a mask when necessary
- 5) Worker is not wearing a safety harness when working at heights
- 5 6) Worker is leaning on guardrails or barricades at heights
- 7) Barricades are not present at open edges at heights
- 8) Barricades are not present around dangerous machinery
- 9) Accidents or heavy traffic exists around the industrial site
- 10) Worker crosses a restricted zone in the industrial site
- 10 11) Stagnant water is present around the industrial site
- 12) Potholes and uneven surfaces are present on the site
- 13) Poor housekeeping in the form of cluttered pipes, cables and over-stacked material are present on the site

The neural networks 127 work in cohesion to identify the presence of
15 each of the above exemplary safety hazards or adverse events and other such adverse events. Each detected safety hazard or adverse event contribute to a safety score level which is all added together to provide an overall industrial site safety score.

Step 240: Identifying, for each detected adverse event, one or more
20 stakeholders for the respective zone

Memory 124 of the video analytics server 120 comprises a stakeholder database 132. The stakeholder database 132 could be stored on a computing device other than the video analytics server 120 (such as the alert server 130) and may be accessible to the video analytics server 120.

The stakeholder database 132 comprises a mapping of the respective zones 105 with one or more stakeholder or stakeholder identities. The stakeholder database 132 also comprises information to enable the specific direction of alerts or messages to a particular stakeholder computing device 140. At step 240, based on the adverse event detected at 230, the video analytics server 120 identifies one or more stakeholders responsible for the zone in which the adverse event was detected.

Step 250: Transmitting an alert to the one or more stakeholders

At step 250 an alert is transmitted to the identified stakeholders. The alert comprises identification information associated with the adverse event. The identification information comprises at least a location of the adverse event and one or more said images in which the adverse event is visible. The identification information could also comprise the label or class information associated with objects detected/not detected in relation to the identified adverse event.

Once any safety hazard or violation is detected by the video analytics server 120, the alert is sent to the stakeholder computing device 140. The alert could be in the form of a telegram channel message or WhatsApp group message or WeChat message or Email or text message, for example. The alert comprises information regarding the adverse event enabling decision making in response to the detected adverse event by the stakeholder. The notification or alert comprises a zone identifying information in which the violation or adverse event is detected, the timing of the violation, a short video of the violation and the risk level of the violation.

The alert could be transmitted by the alert server 130. The transmission of the alert by alert server 130 may be initiated by the video analytics server 120.

Data regarding the violations or adverse events detected by the video analytics server 120 is stored in memory 124 or a storage device accessible to the video analytics server 120. Data regarding the adverse events is used to populate a dashboard of adverse events. The dashboard
5 is served or made available to the stakeholder computing device 140 through the user interface module 129 of the video analytics server. The dashboard contains details of the violation or adverse event, location, timing, risk level and a video of the violation. The dashboard comprises graphs containing historical information of the site safety conditions,
10 broken down on a daily, weekly, or monthly basis. A safety/unsafety score may be allocated to each detected adverse event. The safety scores may be added together to produce a safety score for an industrial site or worksite.

Figure 3 illustrates an example computer system 300. The video
15 analytics server 120, alert server 130 and the stakeholder computing device 140 may comprise one or more components described with reference to the computer system 300 to provide the requisite functionality. In particular embodiments, one or more computer systems 300 perform one or more steps of one or more methods described or
20 illustrated herein. In particular embodiments, one or more computer systems 300 provide functionality described or illustrated herein. In particular embodiments, software running on computer system/systems 300 performs steps the methods described or illustrated herein or provides functionality described or illustrated herein. Particular
25 embodiments include one or more portions of one or more computer systems 300. Herein, reference to a computer system may encompass a computing device, and vice versa, where appropriate. Moreover, reference to a computer system may encompass one or more computer systems, where appropriate.

This disclosure contemplates any suitable number of computer systems 300. This disclosure contemplates computer system 300 taking any suitable physical form. As an example and not by way of limitation, computer system 1600 may be an embedded computer system, a system-on-chip (SOC), a single-board computer system (SBC), a desktop computer system, a laptop or notebook computer system, a mesh of computer systems, a mobile telephone, a server, a tablet computer system, or a combination of two or more of these. Where appropriate, computer system 300 may include one or more computer systems 300; be unitary or distributed; span multiple locations; span multiple machines; span multiple data centres; or reside in a cloud, which may include one or more cloud components in one or more networks. Where appropriate, one or more computer systems 300 may perform without substantial spatial or temporal limitation one or more steps of one or more methods described or illustrated herein. As an example and not by way of limitation, one or more computer systems 300 may perform in real-time or in batch mode one or more steps of one or more methods described or illustrated herein. One or more computer systems 300 may perform at different times or at different locations one or more steps of one or more methods described or illustrated herein, where appropriate.

Computer system 300 includes a processor 302, memory 304, storage 306, an input/output (I/O) interface 308, a communication interface 310, and a bus 312. Although this disclosure describes and illustrates a particular computer system having a particular number of particular components in a particular arrangement, this disclosure contemplates any suitable computer system having any suitable number of any suitable components in any suitable arrangement.

Processor 302 may include hardware for executing instructions, such as those making up a computer program. As an example and not by way of

limitation, to execute instructions, processor 302 may retrieve (or fetch) the instructions from an internal register, an internal cache, memory 304, or storage 306; decode and execute them; and then write one or more results to an internal register, an internal cache, memory 304, or storage
5 306. Processor 302 may include one or more internal caches for data, instructions, or addresses. Instructions in the instruction caches of processor 302 may be copies of instructions in memory 304 or storage 306, and the instruction caches may speed up retrieval of those instructions by processor 302. Data in the data caches may be copies of
10 data in memory 304 or storage 306 for instructions executing at processor 302 to operate on; the results of previous instructions executed at processor 302 for access by subsequent instructions executing at processor 302 or for writing to memory 304 or storage 306; or other suitable data. Processor 302 may include one or more arithmetic logic
15 units (ALUs); be a multi-core processor; or include one or more processors 302. Although this disclosure describes and illustrates a particular processor, this disclosure contemplates any suitable processor.

Memory 304 includes main memory for storing instructions for processor 302 to execute or data for processor 302 to operate on. As an example
20 and not by way of limitation, computer system 300 may load instructions from storage 306 or another source (for example, another computer system 300) to memory 304. Processor 302 may then load the instructions from memory 304 to an internal register or internal cache. To execute the instructions, processor 302 may retrieve the instructions
25 from the internal register or internal cache and decode them. During or after execution of the instructions, processor 302 may write one or more results (which may be intermediate or final results) to the internal register or internal cache. Processor 302 may then write one or more of those results to memory 304. One or more memory buses (which may

each include an address bus and a data bus) may couple processor 302 to memory 304. Bus 312 may include one or more memory buses. In particular embodiments, one or more memory management units (MMUs) reside between processor 302 and memory 304 and facilitate
5 access to memory 304 requested by processor 302.

Storage 306 may include mass storage for data or instructions. As an example and not by way of limitation, storage 306 may include a hard disk drive (HDD), a floppy disk drive, flash memory, an optical disc, a magneto-optical disc, magnetic tape, or a Universal Serial Bus (USB)
10 drive or a combination of two or more of these. Storage 306 may include removable or non-removable (or fixed) media, where appropriate. Storage 306 may be internal or external to computer system 300, where appropriate. This disclosure contemplates mass storage 306 taking any suitable physical form. Storage 306 may include one or more storage
15 control units facilitating communication between processor 302 and storage 306, where appropriate. Where appropriate, storage 306 may include one or more storage 306. Although this disclosure describes and illustrates particular storage, this disclosure contemplates any suitable storage.

20 In particular embodiments, I/O interface 308 includes hardware, software, or both, providing one or more interfaces for communication between computer system 300 and one or more I/O devices. Computer system 300 may include one or more of these I/O devices, where appropriate. One or more of these I/O devices may enable
25 communication between a person and computer system 300. This disclosure contemplates any suitable I/O devices and any suitable I/O interfaces 308 for them. Where appropriate, I/O interface 308 may include one or more device or software drivers enabling processor 302 to

drive one or more of these I/O devices. I/O interface 308 may include one or more I/O interfaces 308, where appropriate.

Communication interface 310 may include hardware, software, or both providing one or more interfaces for communication (such as, for
5 example, packet-based communication) between computer system 300 and one or more other computer systems 300 or one or more networks. As an example and not by way of limitation, communication interface 310 may include a network interface controller (NIC) or network adapter for communicating with an Ethernet or other wire-based network or a
10 wireless NIC (WNIC) or wireless adapter for communicating with a wireless network, such as a WI-FI network. This disclosure contemplates any suitable network and any suitable communication interface 310 for it. Bus 312 may include hardware, software, or both coupling components of computer system 300 to each other.

15 The term "adverse event" refers to an event giving rise to a potential for injury, equipment damage, idle worker(s) and other undesirable events.

The intermediary layer governs the use of particular neural networks on images from particular image capture devices – i.e. cameras. This ensures neural networks used for detecting adverse events are limited to
20 those relevant to one or both of (i) adverse events (e.g. fall hazards, human-equipment interactions and others) anticipated to occur in the respective zone for the particular image capture device, and (ii) adverse events expected to occur in the workplace in which the image capture device is located. For example, for a camera positioned to detect proper
25 use of harnesses on an elevated platform, there is no need to use a neural network trained to recognise heavy equipment such as excavators. Similarly, in a factory setting, the neural networks may be limited to those identifying personal protective equipment (PPE), forklifts and other

potential adverse events occurring in that factory setting. A user may also specify particular hazards they wish to identify in the zone associated with a particular image capture device.

It will be appreciated that many further modifications and permutations
5 of various aspects of the described embodiments are possible. Accordingly, the described aspects are intended to embrace all such alterations, modifications, and variations that fall within the spirit and scope of the appended claims.

Throughout this specification and the claims which follow, unless the
10 context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference in this specification to any prior publication (or information
15 derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

Claims

1. A system for detecting an adverse event, comprising:
 - one or more image capture devices, each device being
 - 5 positioned to capture images of a respective zone; and
 - a server comprising a plurality of neural networks, each neural network receiving, and analysing, at least a portion of said images to identify an object or objects in the respective zone, the server being configured to detect if the respective zone comprises an
 - 10 adverse event based on the identified object or objects; and
 - an alert server for identifying, for each detected adverse event, one or more stakeholders for the respective zone and transmitting an alert to the one or more stakeholders, the alert comprising identification information for the detected adverse event,
 - 15 wherein each image capture device is positioned based on:
 - a relationship between a boundary of the respective zone and
 - a field of view of the respective image capture device;
 - a distance of the image capture device from the respective zone, a resolution of the respective image capture device and a size
 - 20 of a smallest said object sought to be identified in the respective zone.
2. The system of claim 1, wherein the one or more image capture devices comprise a plurality of image capture devices.
- 25 3. The system of claim 1 or claim 2, wherein each image capture device is configured to send video feed comprising the images of the respective zone.

..

4. The system of claim 1, wherein a minimum distance of the image capture device from the respective zone is determined based on the boundary being wholly within the field of view, and a maximum distance of the image capture device from the respective zone is determined based on the resolution enabling detection of the smallest said object at a furthest point in the field of view.
5. The system of any one of claims 1 to 4, wherein each neural network is configured to identify a different object.
6. The system of claim 5, wherein at least one neural network is configured to detect a human, and wherein the server is configured to detect if the respective zone comprises an adverse event if a human is identified by the at least one neural network.
7. The system of claim 6, wherein the server is configured to detect if the zone comprises an adverse event by calculating a relationship between a human and another identified object.
8. The system of claim 7, wherein said another identified object comprises equipment for use by a human, the server being configured to calculate the relationship by determining if the equipment is present in the images with the at least one human.
9. The system of claim 8, wherein the server is configured to detect that the respective zone comprises an adverse event when either:
 - the equipment is absent from the images; or
 - a human and the safety equipment are present in the images, but not in association with one another.

10. The system of claim 9, wherein the equipment and human are in association with one another if the equipment is being used as it is intended to be used.
- 5 11. The system of any one of claims 1 to 10, wherein the identification information comprises at least a location of the adverse event and one or more said images in which the adverse event is visible.
- 10 12. The system of claim 11, wherein each neural network is configured to label said at least a portion of the images analysed by the respective neural network, based on the respective neural network identifying the object or objects.
- 15 13. The system of claim 12, wherein the identification information comprises said label.
14. The system of claim 9, wherein the alert server is configured to send the alert to the human, the identification information specifying the equipment that is sought to be identified in association with the human.
- 20 15. The system of any one of claims 1 to 14, further comprising an intermediary layer, the intermediary layer specifying particular neural networks of the plurality of neural networks for identifying a predetermined said adverse event.
- 25 16. The system of any one of claims 1 to 14, further comprising an intermediary layer, the intermediary layer specifying particular neural networks of the plurality of neural networks for use with images from each respective image capture device.
- 30

17. A method for detecting an adverse event, comprising:

capturing one or more images using one or more capture devices
positioned to capture images of a respective zone, wherein each
5 image capture device is positioned based on:

a relationship between a boundary of the respective zone and
a field of view of the respective image capture device;

a distance of the image capture device from the respective
zone, a resolution of the respective image capture device and a size
10 of a smallest said object sought to be identified in the respective zone;
and

receiving at least a portion of said images at a plurality of
neural networks, the neural networks being configured to identify an
object or objects in the respective zone;

15 detecting if the respective zone comprises an adverse event
based on the identified object or objects

identifying, for each detected adverse event, one or more
stakeholders for the respective zone; and

transmitting an alert to the one or more stakeholders, the
20 alert comprising identification information for the detected adverse
event.

18. A system for detecting an adverse event, comprising:

one or more image capture devices, each device being positioned
25 to capture images of a respective zone, wherein each image capture
device is positioned based on:

a relationship between a boundary of the respective zone and
a field of view of the respective image capture device;

a distance of the image capture device from the respective
30 zone, a resolution of the respective image capture device and a size

of a smallest said object sought to be identified in the respective zone;
and

5 a server comprising a plurality of neural networks, each
neural network receiving, and analysing, at least a portion of said
images to identify an object or objects in the respective zone, the
server being configured to detect if the respective zone comprises an
adverse event based on the identified object or objects; and

10 the server further identifying, for each detected adverse
event, one or more stakeholders for the respective zone and
transmitting an alert to the one or more stakeholders, the alert
comprising identification information for the detected adverse event.