CORTEXICA
AI SOLUTIONS FOR BUSINESS

Real-time PPE Monitoring on the Edge

powered by the latest ultra-low-power high-performance Intel Myriad X VPU

Martin Peniak, PhD — Innovation Department

+44 7757 793 991

martin@cortexica.com

Abstract — Personal Protective Equipment (PPE) is the equipment that protects the user against health or safety risks. Not wearing PPE dramatically increases the chances of injuries and in many situations also of financial losses due to fines for injuries and death of workers as well as contamination caused by not wearing gloves, hairnets, shoe covers, etc. Recent advancements in edge computing hardware coupled with ever more efficient software have enabled novel solutions with the potential to prevent injuries, save lives as well as money and time. This paper describes the world's first edge compute solution for PPE ingress and real-time PPE compliance monitoring.

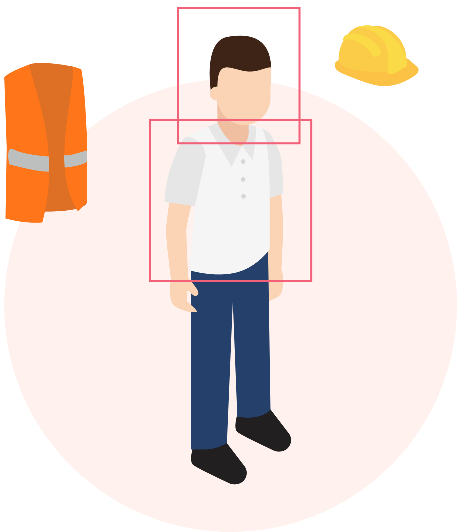

£14.9
billion per year
overall cost of work-related injury
Source: hse.gov.uk

31.2
million days lost
due to non-fatal accidents and ill health

Introduction

Cortexica has been at the forefront of the machine vision revolution. We have reverse-engineered parts of the human visual cortex, which allowed us to develop a powerful image search engine nowadays widely used to solve many real-world problems. Partnering with various world-leading businesses allows us to identify numerous real-world problems that can be addressed by applying machine vision methods. Some of the so-lutions to these problems are often very specific while others are widely applicable and have the potential to save lives. Probably the best example of this is our family of solutions designed for PPE ingress and real-time PPE compliance monitoring (see Figures 1 and 2).

PPE is designed to protect users from serious injuries or illnesses resulting from a physical, mechanical, elec-trical, chemical or radiological contact. The importance of PPE is paramount because it serves as the last line of defence against an injury or death. Unfortunately, studies have demonstrated that 98% of workers said they have seen others not wearing PPE when they should have been and 30% of those said this happens regularly. Head injuries, constituting 9% of all injuries, can be fatal and yet 84% of these were caused by not wearing a helmet [1]. Most of these injuries could have been prevented had there been a system in place that would continuously monitor for PPE compliance.

Recent advancements in edge computing gave raise to novel applications designed to process data right at its source effectively minimising latency and allowing real-time processing. This paper presents a solution for PPE ingress and compliance monitoring that runs in real-time and entirely on the edge.

[1] U.S. Department of Labor, Bureau of Labor Statistics, Accidents Involving Head Injuries, Report 605, (Washington, D.C., Government Printing Office, July 1980) p. 1

Fig. 1
This figure shows a Cortexica PPE ingress solution installed at Axis Experience Centre, Madrid. This solution detects if a person is wearing a helmet and a high visibility jacket. A person is only allowed to enter a workplace once all the PPE checks have been approved.

Click here to watch a video demonstration →

Fig. 2
This figure shows a Cortexica real-time PPE compliance monitoring solution installed at Axis Experience Centre, Madrid. This solution is continuously scanning workplace and raises an alarm if a person is not wearing a helmet and a high visibility jacket. The scanning is done using Axis pan/tilt/zoom (PTZ) camera and Axis radar capable of detecting motion within 50 meters range. The radar detects motion and then automatically moves the camera to the point of interest.

Click here to watch a video demonstration →

Development Kit Overview

It takes more than algorithms to deliver an AI driven product that solves a real-life problem. Working in partnership with UP we have developed UP Squared AI Edge-PPE monitoring, a development kit (see Figure 6) for health and safety professionals to create proof of concepts (POC) to full blown AI applications ready for live deployment. The development kit offers the following:

- Real-time video analysis with advanced algorithms and machine learning to ensure employees are wearing the correct PPE for their working envi-ronment
- Parallel detection of PPE, Face, Person, and Body parts leveraging CPU, GPU and VPU processors
- Powered by the latest ultra-low-power high-performance Intel Myriad X VPU (see Figure 7)
- Single image mode
- Real-time mode
- Made for POCs

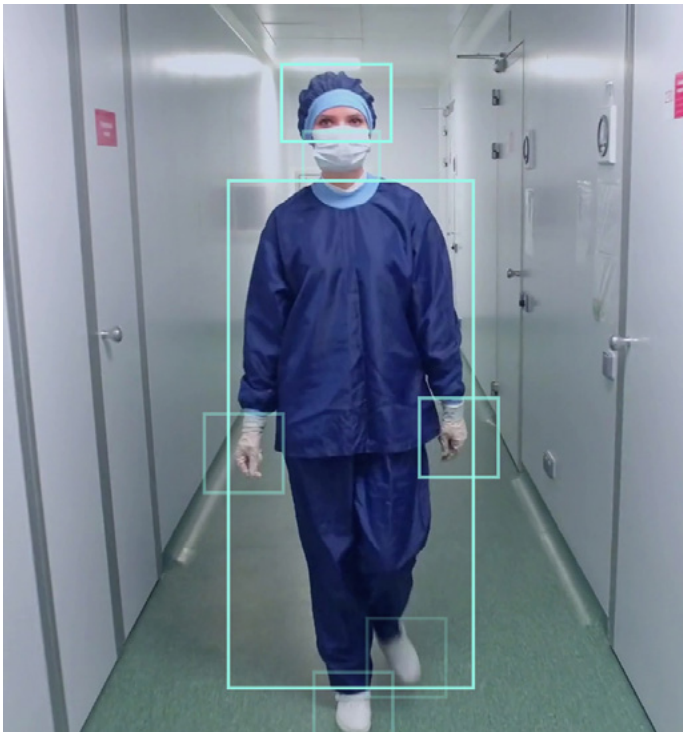

Fig. 3
This figure portraits an example of a Cortexica PPE ingress solution designed for a pharmaceutical company. This solution de-tects if an employee is wearing boots, overall, gloves, hairnet, beardnet, head cover and transparent glasses. An employee is only allowed to enter a workplace once all the PPE checks have been approved.

Fig. 4
Cortexica was the first and only company to present a working demonstration utilising the latest Intel Myriad X VPU. This image shows one of our AI Safety (PPE) demos at the Intel booth. The second demo was installed at AAEON (Asus asoc.) stand.

Fig. 5
The PPE monitoring development kit is available to purchase as an off-the-shelf solution, complete with hardware and software configuration, and detailed step-by-step guides to help you start prototyping your next AI monitoring or video surveillance project. For health and safety professions, this development kit assists the creation of a PPE monitoring system, and for industries at large, an AI driven surveillance application.

Fig. 6
UP Squared featuring Intel Atom® x7-E3950 8GB RAM/64GB eMMC system with WiFi (2T2R) and Bluetooth enabled.

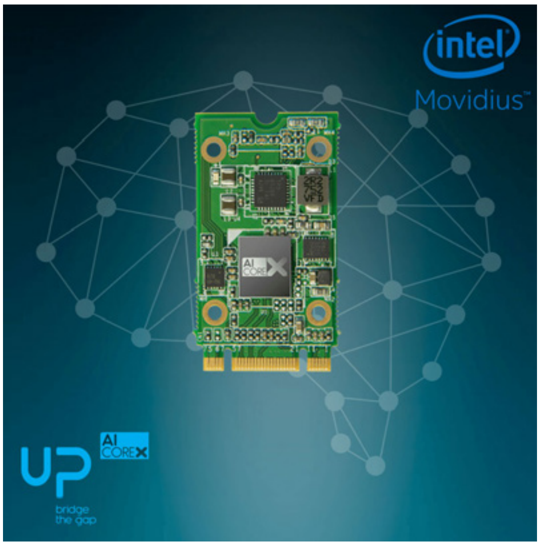

Fig. 7
The AI CORE X is powered by the recently released Intel® Movidius™ Myriad™ X, a third-generation vision processing unit (VPU) that is the first in its class to include a Neural Compute Engine – a dedicated hardware accelerator for deep neural networks, trainable with industry-standard tools.

Fig. 8
UP UB3.0 HD camera with 3.6mm lens.

Software

This section provides a detailed description of the software architecture, its components and their inte-gration through the REST API.

PPE Service

The PPE service is a C++ application running a HTTP server with a REST API. This service is able to run inference on multiple deep-learning models in par-allel on all the available CPU, GPU and VPU processors to obtain the results in the shortest time possible. The service starts automatically running after the operating system is initialised and then continues running on the port 8081. The REST API has the following endpoints that can be used to interact with the service:

- /ppe – single image mode allowing maximum of 10 requests followed by 180 seconds timeout
- /continuous-ppe – continuous mode allowing un-limited number of images within 20 seconds fol-lowed by 180 seconds timeout

Both endpoints accept multipart/form-data POST re-quests containing an image and a JSON specifying the query options (see Figure 11).

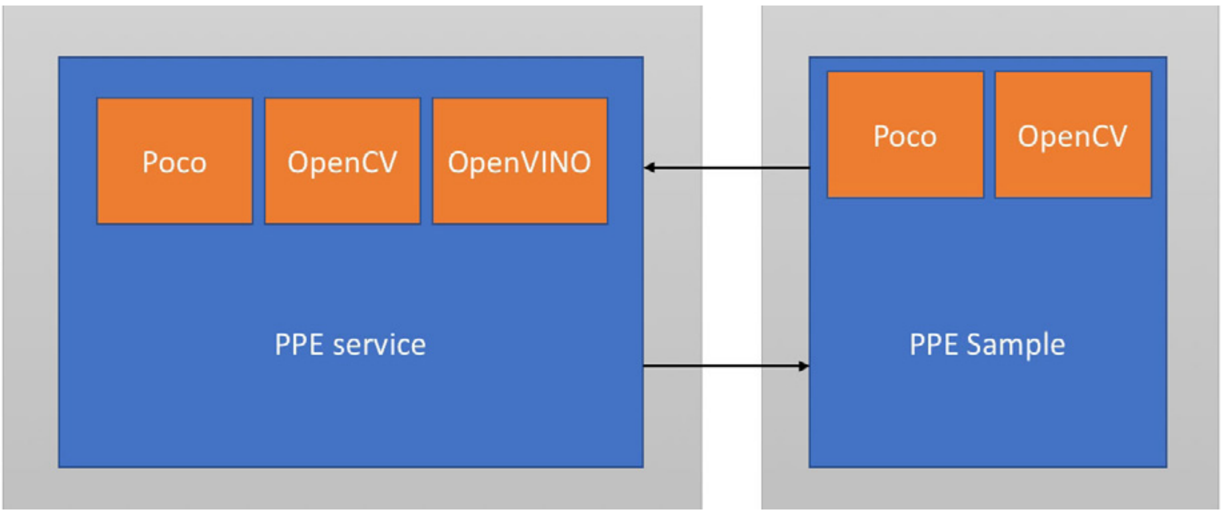

Fig. 9
Software architecture diagram showing the PPE service and the demo sample. The PPE service receives an image and a JSON specifying simple options and returns a JSON with detection results containing bounding boxes, corresponding object classes and the time in milliseconds that it took to run a single inference.

```
Received results: {
  "ms" : 49,
  "result" : [
    {
      "categoryClass" : "vest",
      "confidence" : 0.92055,
      "xEnd" : 0.82968,
      "xStart" : 0.47890,
      "yEnd" : 0.99027,
      "yStart" : 0.36382
    },
    {
      "categoryClass" : "helmet",
      "confidence" : 0.95361,
      "xEnd" : 0.74687,
      "xStart" : 0.58749,
      "yEnd" : 0.21527,
      "yStart" : 0.03449
    },
    {
      "categoryClass" : "face",
      "confidence" : 1,
      "xEnd" : 0.92653,
      "xStart" : 0.61093,
      "yEnd" : 0.41527,
      "yStart" : 0.13194
    }
  ],
  "status" : 0
}
```


Fig. 10
JSON results example showing where ms is inference time in milliseconds, categoryClass is the class of the detected object, confidence is the detection confidence level (e.g. 0.95 = 95% confidence) and xStart, xEnd, yStart, yEnd define the bounding boxes of the detected object. When multiplied with the image dimensions, pixel coordinates can be obtained. For example, a bounding box in OpenCV can be created using: cv::Rect rect(cv::Point(xStart * width, yStart * height), cv::Point(xEnd * width, yEnd * height)) where width and height are the image dimensions.


```
{
  "detectPPE": true,
  "detectFaces": true,
  "detectPersons": true,
  "detectBodyParts": true,
  "ppeThreshold": 0.6,
  "faceThreshold": 0.6,
  "personThreshold": 0.6,
  "bodyPartsThreshold": 0.6
}
```


Fig. 11
JSON options example showing all the four models enabled with thresholds set to 60% confidence. The threshold parameters are optional and if no threshold is set then the default value of 50% confidence will be applied.

```
curl --request POST
  --url http : / / localhost :8081/ppe
  --header ' content-type : multipart /
  form-data ;
  boundary=----WebKitFormBoundary'
  --form ' options={
  "detectPPE " : true ,
  "detectFaces " : true ,
  "detectPersons " : true ,
  "detectBodyParts " : true ,
  "ppeThreshold " : 0.5 ,
  "faceThreshold " : 0.5 ,
  "personThreshold " : 0.5 ,
  "bodyPartsThreshold " : 0.5} '
--form image=@example.png
```

Fig. 12
Example curl request.

PPE Demo Sample

This sample application was written in C++ to demonstrate how to send queries to and receive results from the PPE service via the REST API. The application source code is located under /home/upsquared/cortex-ica/sample/ and the compiled binary under the bin subdirectory. The application depends only on OpenCV and Poco libraries. To compile the sample run the following commands:

```
$ cd /home/upsquared/cortexica/sample/
$ mkdir build
$ cd build
$ source /opt/intel/computer_vision_sdk/bin/setupvars.sh
$ cmake ..
$ make
```

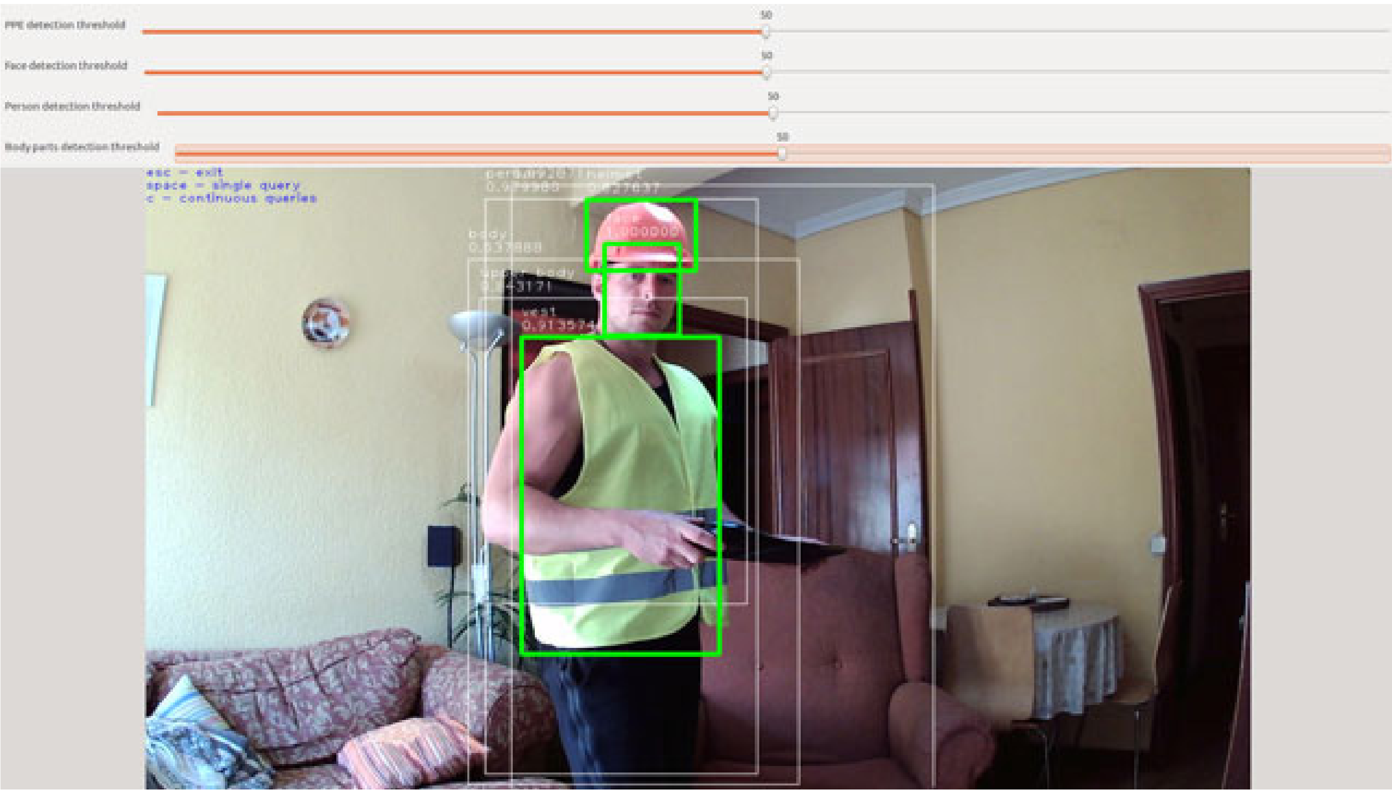

Fig. 13
Sample application front-end showing the detection results plotted as bounding boxes. The four sliders at the top are used to adjust the sensitivity of the models or to turn them off completely.

Performance

The PPE service provides four different models that are mapped to CPU, GPU and VPU processors. There are four models but only three processors and there-fore the GPU can in certain configurations run two models. The Myriad X is able to run face detection at around 25 milliseconds per image (40 FPS). The GPU can run the PPE or person detection at approximately 50 milliseconds (20 FPS). If both PPE and person de-tection models are enabled then this latency increases to around 85 milliseconds (12 FPS) as the models need to share the same resources. The body parts detection runs on the CPU at around 50 milliseconds (20 FPS). The final latency of the PPE service will depend on what models have been enabled. For example, if only face detection is used then the frame-rate will be around 40 FPS. However, if the PPE model is also enabled then the frame-rate drops to around 20 FPS as the final latency can only be as low as the latency of the slowest model, which in this case would be the PPE moel running on the GPU at around 50 milliseconds.

License and Limitations

The development kit has been primarily designed to enable the evaluation of our technology and to facilitate seamless integration with other applications with the goal of creating POCs. If a client is satisfied with the POC and wants to use it for commercial purposes then the client should contact us to obtain a license and a production-grade software with models optimised for that specific use case. As a consequence, using this development kit for commercial purposes is prohibited. The following measures have been implemented to prevent unauthorised use of this development kit:

- Single image mode is restricted to a maximum of 10 consecutive requests. There is a 180 seconds time out afterwards during which no requests can be made. Additional requests can be made after this time out has finished.
- Real-time mode allows unlimited number of frames within a 20 seconds period. There is a 180 seconds time out after these 20 seconds during which no requests can be made. Additional requests can be made after this time out has finished.
- The PPE service is only available as a binary, which will expire after one year from the first time the development kit was powered on.
- Models have been encrypted and therefore their use outside the scope of this RRK is not possible.

For more information contact us

+44 (0) 203 868 8880
info@cortexica.com