Automated Livestock Detection and Assistance Technology

Dylan Bedetti

Mohammed Bennamoun
Computer Science
University of Western Australia

Prof. Farid Boussaid
Electrical, Electronic and Computer Engineering
University of Western Australia

Jacob Pirone & Anton Wagner
CEED Client: BHP

Abstract

The remote nature of BHP mining operations leads to a number of Light Vehicle (LV) interaction events with livestoc. These incidents pose massive risk to both employees and livestock and can be fatal, with two deaths Australia-wide (non-BHP related) in the last recorded year. Incidents commonly occur during dusk, dawn and at night-time, where there is limited vision and cattle can blend in with their surroundings.

This project proposes a solution comprised of a small Field of View (FoV) thermal camera to feed images to a Convolutional Neural Network (CNN) where computation is completed by the Google Coral, a Tensor Processing Unit (TPU). If the model recognises any livestock, it will alert the user through a speaker. The device has been designed to run 'at the edge' and should operate completely without any internet connection. Early results show promising performance in detecting livestock, especially at critical times such as dusk, dawn and night-time. This research is largely a proof of concept for the technology, which will encourage future development in the area of livestock detection and assistance technology.

1. Introduction

The remote nature of BHP mining operations leads to a number of Light Vehicle (LV) interaction events with. LV's have a gross vehicle mass of up to 4,500 kg and for the purpose of this project is the only vehicle type that will be considered, primarily because they provide the least protection to vehicle occupants in collisions. These incidents can be fatal, with two deaths already Australia-wide in the last recorded year (Australian Associated Motor Insurers, 27), (Government of Western Australia Road Safety Commision, 2016). Incidents commonly occur during dusk, dawn and at night-time, where there is limited vision and cattle can blend in with their surroundings.

Drivers are far more likely to have a collision in regional areas compared to urban areas (Government of Western Australia, 2019). Most of BHP's mine sites operate within 'very remote' areas, which is an index of remoteness derived from the measure of road distance between populated localities and service centres, a term defined by the Australian Department of Infrastructure (Department of Infrastructure, Regional Development and Cities, 2018). These areas have greater than eleven times the rate of road deaths compared to that of major cities (Department of Infrastructure, Regional Development and Cities, 2018), (Australian Institute of Health and Welfare, Rural, regional and remote health - A guide to remoteness classifications). This risk of is amplified further when driving at night time, where limited vision and lack of lighting can decrease a driver's perception and response time, especially for objects that blend in with their background (Green, 2000).

The current controls to avoid LV incidents are limited to the two lowest and least effective methods, Personal Protective Equipment (PPE) and Administrative Controls (AC), shown below in Figure 1. While these controls can mitigate risk of an incident, they are considered minimally effective, and will not eliminate the risk (Australia, 2018).

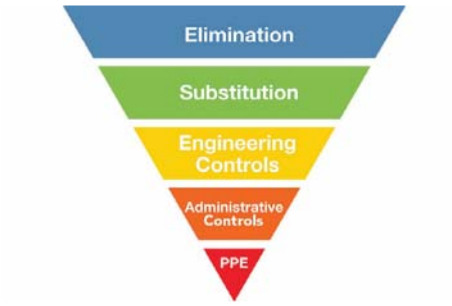

Figure 1 Hierarchy of Controls (Australia, 2018)

The problem of LV livestock collision presents itself as an opportunity to develop a technological risk control through the novel combination of thermal imagery, Deep Learning (DL) and edge computing. To the best of our knowledge, this will be the first research project to date which combines Deep Learning with thermal imagery livestock detection.

2. Process

2.1 Data Collection

Data collection consisted of filming cows on a farm in South-West Western Australia. A normal RGB camera was placed in parallel to the thermal camera to give images a source of ground truth. An output of this is shown in Figure 2 below.

Figure 2 Data collection using RGB and thermal camera

Choosing a thermal camera requires balancing cost and image resolution constraints. The scope of this project limited the cost of the product to a maximum of $3000. Because of this cost restraint, the only thermal cameras available have low image resolution. However, a small FoV camera can provide high pixel density, even at lower resolutions. The thermal camera chosen was an AXIS P1280-E with a horizontal FoV of 35.4°, a resolution of 208 x 156 pixels and a frame rate of 8.3 frames per second (FPS) (AXIS Communications, n.d.).

Figure 3 shows the FoV of the AXIS thermal camera when mounted on a vehicle. The camera can see past the boundaries of the road after 11.5 meters. After 40 meters the camera will show a front-facing cow with 10 pixels, and a sideways facing cow with 30 pixels.

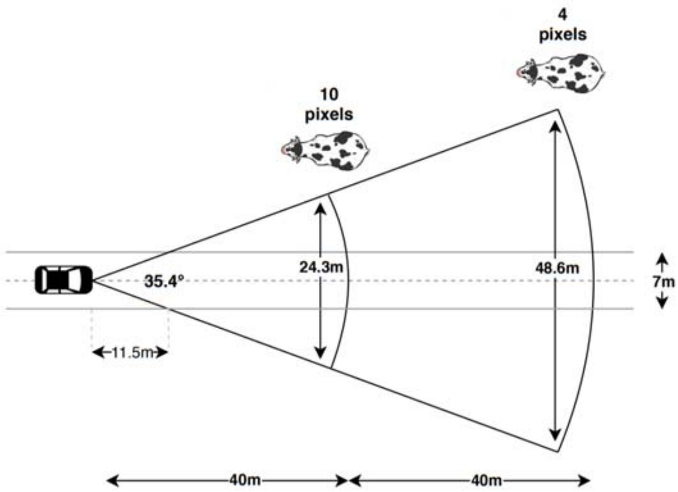

Figure 3 Thermal Camera FoV

2.2 Model Training

A core requirement of the solution is model inference speed, as the faster the inference of the model, the faster a driver can be alerted of cattle ahead. To achieve this speed, the Google Coral (Google, n.d.), an edge compute Tensor Processing Unit (TPU) was used. Currently, the Google Coral only supports MobileNets for detection-based models. The Coral is at the forefront of edge Deep Learning and is still being further developed.

MobileNet Single Shot Detector (SSD) V1 which had been trained on the Common Objects in Context (COCO) dataset was used for transfer learning of the last few layers (COCO, n.d.).

The MobileNet model has input dimensions of 300 x 300 pixels which required resizing of the thermal camera image from 208 x 156 pixels. The tensor required for training of the MobileNet model has 3 input channels, however the thermal image only returns a single channel of thermal intensity. It was therefore decided to duplicate three copies of the thermal image to get the desired number of channels.

A thousand images were chosen from the data collection and were labelled. These were used for training with an 80/20 split for training and testing data. The model was only trained with a single class of ‘cow’. Most of the convolutional layers were frozen, with only the last few pointwise and depthwise convolutions remaining unfrozen, and malleable during training. The training was completed with 500 training steps and 100 evaluation steps. Further training did not significantly increase model performance.

Google also offers an Automatic Machine Learning (AutoML) service, which only requires a training dataset and object labels. The AutoML service will test multiple different model architectures and attempt hyperparameter optimisation on each with the goal of returning the best model for the given input data.

The same dataset of 1000 images and labels was fed into the Google AutoML service. The resulting model achieved a precision of 93.77% and a recall of 59.32% at an IoU and confidence threshold of 50%.

3. Results and Discussion

The MobileNet Edge TPU enabled model was chosen for the MVP over the AutoML model because of the smaller number of model operations, 64 compared to 294, and the greater percentage of Edge TPU enabled operations, 98% compared to 45%.

For the application of this product, speed is more important than accuracy. This is because the possible impact of a false positive has less risk than a false negative. Also, the risk of a collision grows as the delay between capturing an image and alerting the driver increases. Therefore, for this product, a model that achieves fast inference with less accuracy is preferred over one that has slow inference but higher accuracy. Figure 4 below illustrates the model’s ability to detect cows within the test dataset.

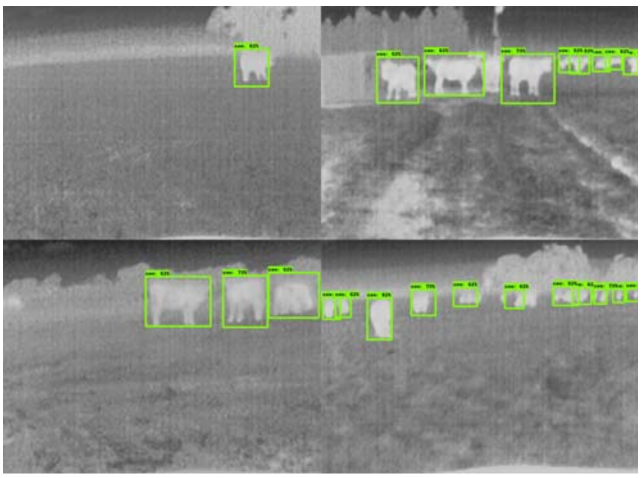

Figure 4 Model detections on testing data

The associated weighted sigmoid focal loss, plotted in Figure 5, indicates the presence of a learning process. However, after testing the inference with live field data, it became evident that the model is not accurate in distinguishing different objects with similar thermal signatures, such as being able to differentiate a cow from either a human or car at distance.

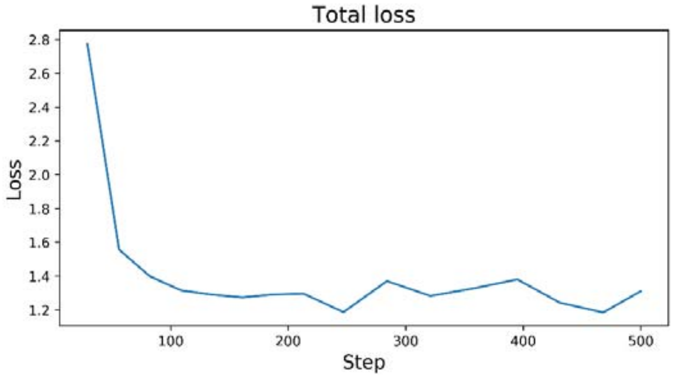

Figure 5 Weighted sigmoid focal loss

This limitation is likely caused by the thermal camera resolution, as at a distance of 40 meters, an object the same size of a forward facing cow will be visible in only 10 pixels of the image (AXIS Communications, n.d.). All visual Deep Learning methods are restricted on using the information available in the image presented, and the information that can be conveyed by 6 pixels in a grayscale image is limited.

4. Conclusions and Future Work

An MVP has been successfully developed, tested and evaluated in and around a BHP mine-site environment for detecting and alerting the presence of cattle. This project was sponsored by BHP, in-partnership with UWA and was supported by many individuals throughout both organisations.

The product delivered is a cattle detection retrofit kit for LV's which uses a combination of GPS position and a Deep Learning model for inference and alerting drivers. The project offers a scalable software platform and recommended technical direction for further development.

Such a project is made possible with the very recent advances in hardware, software and Deep Learning. It is truly exciting to imagine what impact these technologies will have in the coming years as they become faster, smaller and more integrated. It is the hope of the author that this project can be used as both a steppingstone for further development and as inspiration for others to pursue improvement and change through technology.

5. Acknowledgements

I would like to express my sincere gratitude to everyone who supported, challenged and educated me during my final year thesis project. To the University of Western Australia who provided me with the opportunity to pursue a Masters degree in Professional Engineering, and for giving me the tools to succeed. To the Co-operative Education for Enterprise Development program, Jeremy Leggoe & Amanda Bolt for providing me with the opportunity to undertake an industry focused Thesis and for the ongoing support and guidance. To BHP, Jacob Pirone & Anton Wagner for enthusiastically accepting my project proposal, continuing my professional development and for guiding me through the challenging process of building, innovating and developing.To Winthrop Prof. Mohammed Bennamoun & Prof. Farid Boussaid for providing me with technical insight and inspiration, and for furthering my education and passion of machine learning.

6. References

Australia, S. W. (2018). *How to manage work health and safety risks.* Canberra: Safe Work Australia.

Australian Associated Motor Insurers. (27, 5 27). *AAMI animal car accidents data 2019.* (AAMI) Retrieved 8 28, 2019, from https://www.aami.com.au/aami-answers/insurancey/aami-reveals-peak-periods-for-animal-collisions.html

Australian Institute of Health and Welfare. (Rural, regional and remote health - A guide to remoteness classifications). *Australian Institute of Health and Welfare.* Canberra.

AXIS Communications. (n.d.). *AXIS P1280-E Thermal Network Camera.* (AXIS) Retrieved May 3, 2020, from https://www.axis.com/en-au/products/axis-p1280e

BHP. (2019). BHP Event Reporting Portal (ERP). Perth: BHP.

BHP. (2019). Light Vehicle Collision with Livestock at Blackwater Mine. Brisbane: BHP.

COCO. (n.d.). *Common Objects in Context.* Retrieved May 2, 2020, from http://cocodataset.org/#home

Department of Infrastructure, Regional Development and Cities. (2018). *International road safety comparisons 2016.* Canberra.

Google. (n.d.). *Coral.* Retrieved May 2, 2020, from https://coral.ai/

Government of Western Australia. (2019). *2018 Preliminary summary of fatalities of Western Australian roads.* Perth: Road Safety Commision.

Government of Western Australia Road Safety Commision. (2016). *Reported road crashes in Western Australia 2015.* Perth: Road Safety Commision.

Green, M. (2000). "How Long Does It Take to Stop?" Methodological Analysis of Driver Perception-Brake Times. *Transportation Human Factors, 2*, 195-216.

Loughnan, L. (2011). FENCING ROAD RESERVE IN PASTORAL REGIONS MAIN ROADS/PGA & STAKEHOLDERS . Perth: Main Roads Western Australia .

Mainroads Western Australia. (2019, May 8). *Design of Fencing/Walls.* (The Government of Western Australia) Retrieved September 7, 2019, from https://www.mainroads.wa.gov.au/BuildingRoads/StandardsTechnical/RoadandTrafficEngineering/RoadsideItems/Pages/Design_of_Fencing_Walls.aspx

Rebecca Butcher, S. C. (2019, Feburary 13). *The Western Australian beef industry.* (Department of Pimary Industries and Regional Development) Retrieved September 7, 2019, from https://www.agric.wa.gov.au/industry-development/western-australian-beef-industry?nopaging=1

WALGA. (2017). *Reference information for Elected Members on a Regional Road Group.* Perth: Pilbara Regional Road Group.