WMG, IMC Building
University of Warwick
Coventry
CV4 7AL
United Kingdom

Alan Chalmers, Professor of Visualisation
17 January 2022

Report on PhD Thesis: Petr Musil

The PhD thesis of Petr Musil addresses the topic of object detection in video. In particular, the thesis sets out to make three major contributions to the field:

1. Optimising image detecting algorithms by implementing them on a FPGA
2. Developing a new approach for multiscale object detection with the key advantage that there is no need for external memory
3. Clearly showing that the approach investigated and implemented outperforms current state-of-the-art methods by enabling higher levels of object detection for less power requirements.

The novel work was validated on two challenging real-world applications: face recognition and identifying number plates on cars from a moving vehicle.

The thesis is divided into 8 chapters.

Chapter 1 introduces the topic of object detection and justifies the need for efficient computational solutions due to the huge amount of data that is generated from a multitude of cameras. The work is motivated by emphasising the high computational costs of current techniques. FPGA based approaches offer novel low-cost, low-power solutions for real-time video processing.

Chapter 2 presents a detailed review of previous literature relating to the use of embedded devices for object detection. There is a focus on previous research that investigated object detection methods using boosted classifiers and hardware detection methods. Figures 2.3 and Figure 2.4 in the chapter provide a useful comparison of the considered detectors for faces and pedestrians.

An overview of goals and contribution of the research presented in thesis is given in **Chapter 3.** The choices made are justified and the three high quality publications from this research highlighted. A list of other publications in which the candidate was involved is also given. Although not directly related to the research in this thesis, they do give a clear indication of the candidate's ability to undertake high quality research in a range of topics.

The core of the research presented in this thesis is given in Chapters 4,5 and 6 as re-formatted versions of the papers that were published.

Chapter 4 presents the work published in 2013:

- High performance architecture for object detection in streamed videos," in Field Programmable Logic and Applications (FPL), 2013

This paper presented a novel architecture for object detection in images and video using a scanning window and classification of its contents by a WaldBoost classifier. It is nine years since the paper was published. I would have liked to see some indication in the Conclusion section of this paper, what the impact of this work has been, demonstrated, for example, by the number of citations the paper had received. Also, it would have been helpful to know if the method was still applicable, or whether it has since been superseded by other approaches.

Chapter 5 presents the work published in:

+44(0)2476 522012 alan.chalmers@warwick.ac.uk

- "Cascaded stripe memory engines for multi-scale object detection in fpga," IEEE Transactions on Circuits and Systems for Video Technology, vol. 30, no. 1, pp. 267–280, 2019.

This recent paper gives a comprehensive description of a novel architecture for object detection on an FPGA. A number of configurations of detector cascades were suggested and compared.

In **Chapter 6** the very latest research, which is still under review (at the time of the thesis submission) is given:

- "Unconstrained license plate detection in hardware", submitted to the International Conference on Vehicle Technology and Intelligent TransportSystems (VEHITS), 2021.

The novelty of the developed hardware detector is demonstrated by its use within a camera system that is used to detect license plates. This is a very challenging applications as the detection is performed from a moving car and the license plates need to be detected from a wide range of positions and angles. The results clearly show that the developed detector is sufficiently accurate to be used successfully for this challenging application and, furthermore, has the necessary low power consumption, which is critical for the system to be deployed in the "real world".

Chapter 7 describes some applications for the detector and discusses possible future work. As the developed detector is capable of achieving object detection with good accuracy in full HD video at more than 15 frames per second, it has the potential to be used in a number of important application. The main one discussed in this chapter is its use for road and traffic monitoring. This is in collaboration with the industrial partner, Camea. Future work possibilities discussed include enabling the system to cope with HDR video and using multi-channel features.

Chapter 8 is a one-page conclusion of the work. This includes a summary of what was achieved.

Limitations

I found simply inserting previously published papers in the thesis a little limiting. In particular, I would like to have seen more up-to-date information on the paper which had been published in 2013. In addition, I would like to have seen some of the key results restated in the Conclusion and their importance fully re-emphasised.

Strengths

This thesis has presented a significant amount of research. The thesis clearly demonstrates that the candidate fully understands the field. This research has made an original contribution and has the potential for significant commercial impact in its use for license plate recognition and other similar application. There is certainly enough novel work here to qualify for PhD. The thesis is well written, and the candidate's contributions already published in leading journals. and conferences

Questions for the PhD defence

There are a number of questions that arise from this thesis which would be helpful to have answered when the candidate defends his thesis. These include:

1. In license plate recognition scenario, would it make a difference if the car being looked at was also moving?
2. Your system can localise and detect license plates. How much effort would be required to be actually able to read the number plate e.g. to determine in real-time the owner of the car?
3. Will it be possible one day to deploy these methods within a mobile phone, which also has a low power requirement?

Summary

In summary, this thesis is a significant body of original work. The candidate has clearly shown a very good knowledge of the field and used this knowledge to develop a novel, low powered

object detection system for use on FPGAs. The candidate has clearly demonstrated the efficacy of his method on challenging real world applications.

The candidate's thesis is indeed fully worthy of the award of a PhD, and the candidate should be particularly commended on the amount and quality of work undertaken.

Yours sincerely

Professor Alan Chalmers

+44(0)2476 522012 alan.chalmers@warwick.ac.uk