University Defence Research Collaboration (UDRC)
Signal Processing in a Networked Battlespace

Reliable Pedestrian (and Object) Detection

E_WP6: Efficient Computation of Signal Processing Algorithms

WP Lead: Prof. John Thompson,
University of Edinburgh
Associated Academics: Dr. Neil Robertson &
Prof. Andy Wallace, *Heriot-Watt University*

Researcher: Calum Blair,
University of Edinburgh, c.blair@ed.ac.uk

Introduction

Existing object detectors perform well at finding objects (humans, cars, etc.) in images. This has many civilian as well as networked battlespace applications — not limited to video.

Giving each detection a probability measure helps ranking of detection confidence and identification of uncertain bounding boxes. Then, slower, more accurate algorithms or human operators can classify uncertain regions.

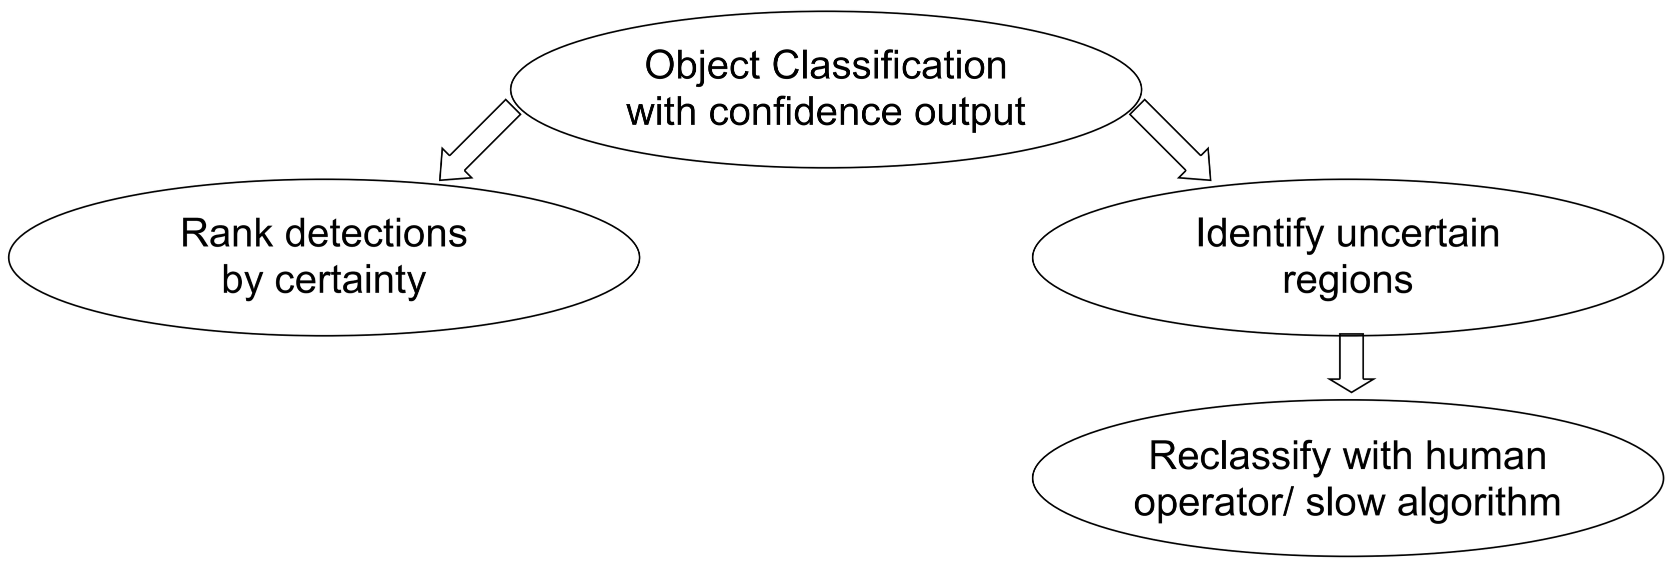

Existing object detectors provide unreliable confidence output: all detections have the same probability(certainty) of 1. The Adaboost pedestrian detector below has state-of-the-art accuracy but is massively overconfident:

Available Detectors

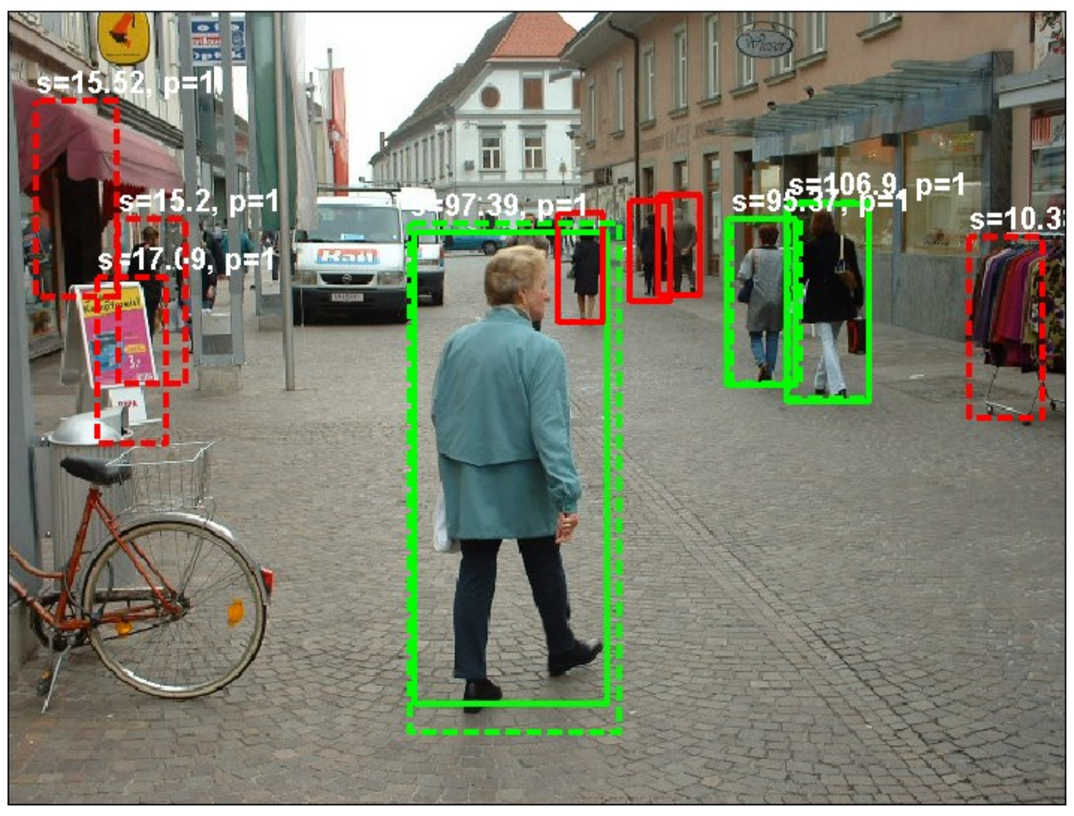

Example output from FPDW detector showing Adaboost scores (s) and probabilities (p).

The current state of the art detector is FPDW[1], a sliding window classifier with 5120–dimensional feature vectors, followed by Adaboost with decision trees as weak learners.

We use the same features for all classifiers and evaluate Adaboost, support vector machines (SVM), relevance vector machines (RVM) and two Gaussian Process classifiers: linear (Linear-GP) and squared exponential (SE-GP) [2].

Adaboost and SVM generate scores so we convert these to probabilities using sigmoid, platt [3] and isotonic regression [4] methods.

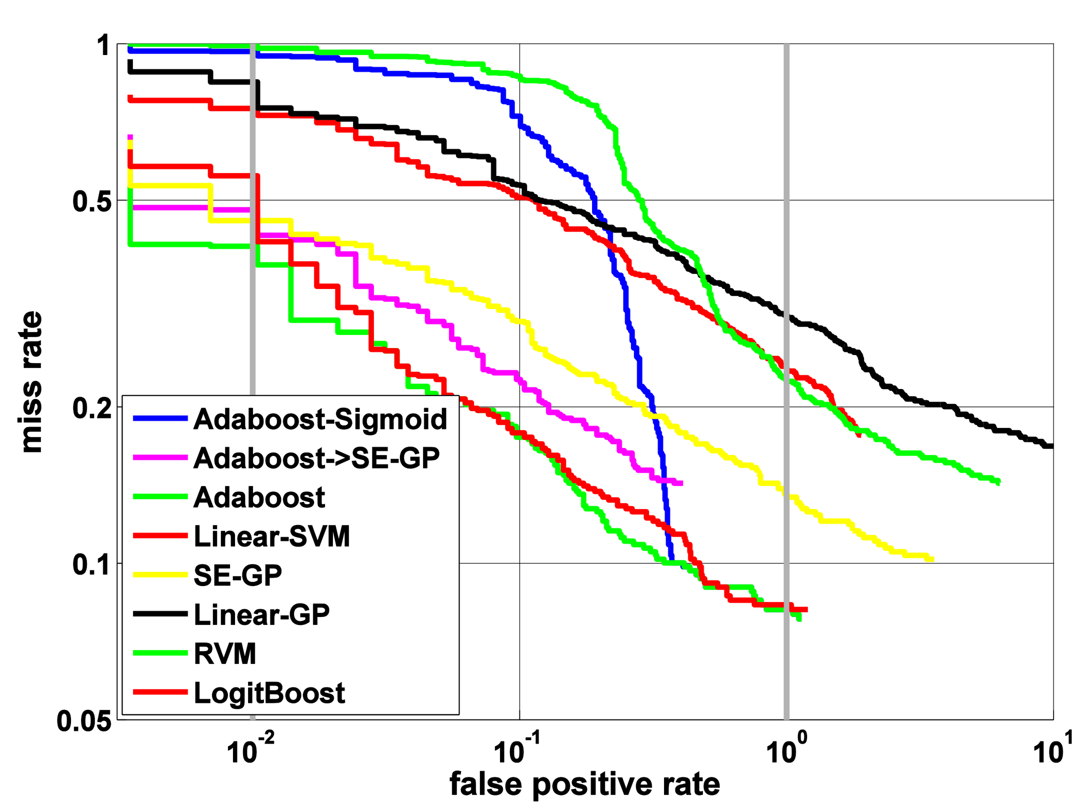

Detection Error Tradeoff curve for pedestrian detectors on the INRIA dataset. State of the Art (Adaboost/FPDW[1]) performs best on this curve.

Measuring Reliable Detections

Error rate measurements using true and false positives (TP, TN, FP, FN) measure accuracy, not reliability. FP and FN are penalised equally and confidence errors are not considered. FN with probability 0.49 and 0.01 are counted the same (right).

Mean squared error addresses this:

$$MSE = \frac{2}{N}\sum_{k=1}^{N}(C_k - p(1|\mathbf{x}_i))^2$$

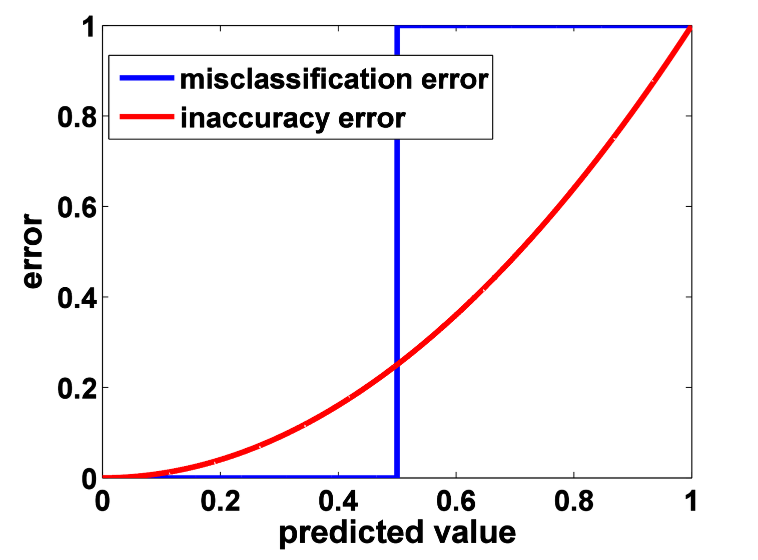

Error rate measurements (blue) only penalize misclassification. Mean-squared error (red) penalizes correct, uncertain detections less.

Classifier	Mapping Method	MSE
Adaboost	sigmoid	0.803
	Platt scaling	0.819
	isotonic regression	1.792
Gaussian Process SE	N/A	0.597
Gaussian Process Linear	N/A	0.978
Linear SVM	Sigmoid	2.667
RVM	N/A	1.209
Adaboost-> SE-GP	N/A	**0.411**

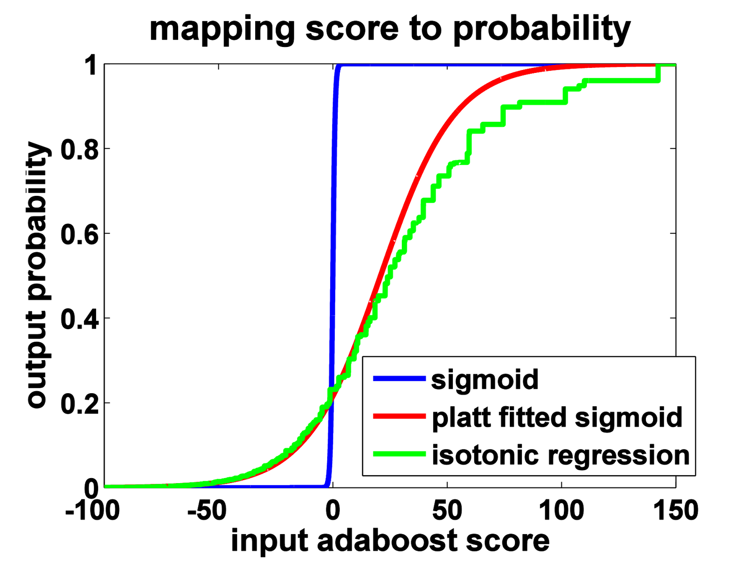

Methods of mapping an input Adaboost or SVM score to output probability: Platt and Isotonic regression are fitted to a validation set.

Reliability plots (below) show how well-calibrated a probabilistic classifier is. Reliable detectors are those closest to the black line.

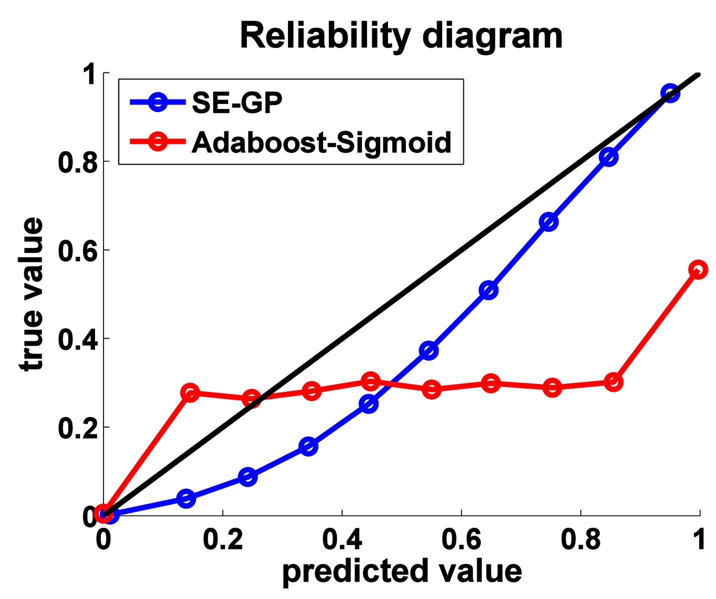

On raw detections (prior to thresholding), the SE-GP classifier is most well-calibrated. The Adaboost classifier, although more *accurate*, is less *reliable*.

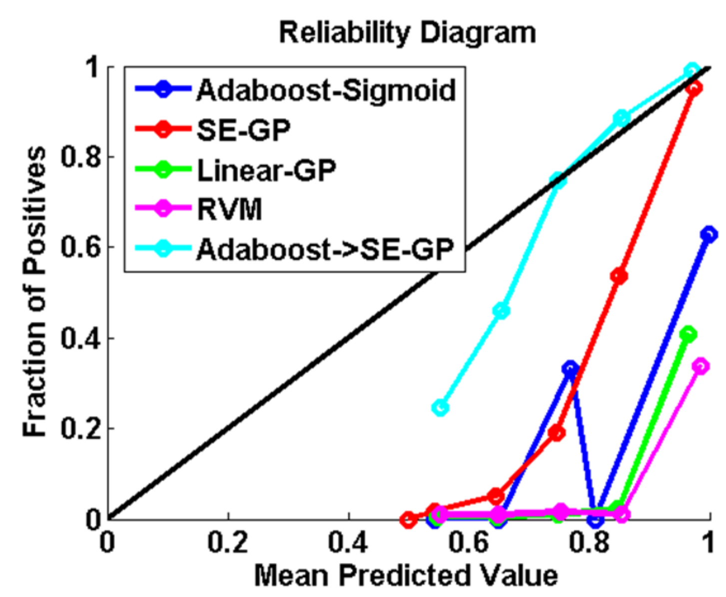

After non-maximal suppression and removal of negative detections, the detection distribution changes and Adaboost followed by SE-GP is closer to the well-calibrated line.

Conclusion

Overall, using detections generated by Adaboost then running SE-GP on these to produce a probability is the most reliable detection scoring method. This is still computationally expensive as SE-GP is $O(n^2)$.

Future Work

We will extend this analysis to other modalities (SAR, sonar imagery) and are producing a GPU-accelerated version of the Gaussian Process classifier to generate near-realtime accurate detections.

We are also investigating ways to identify signals which differ significantly from the training data where any generated detections will be unreliable.

References

[1] Dollar, P., Belongie, S. & Perona, P., 2010. The Fastest Pedestrian Detector in the West. In *Proc. BMVC 2010*, pp. 68.1–68.11.
[2] Rasmussen, C.E. & Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*, University Press Group Limited.
[3] Platt, J., 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In *Advances in large margin classifiers*.
[4]Niculescu-Mizil, A. & Caruana, R., 2005. Predicting good probabilities with supervised learning. In *International Conference on Machine Learning*.

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) and Dstl