Int J Comput Vis (2012) 100:203–215
DOI 10.1007/s11263-011-0461-z

Making a Shallow Network Deep: Conversion of a Boosting Classifier into a Decision Tree by Boolean Optimisation

Tae-Kyun Kim · Ignas Budvytis · Roberto Cipolla

Received: 17 December 2010 / Accepted: 11 May 2011 / Published online: 7 June 2011

Abstract This paper presents a novel way to speed up the evaluation time of a boosting classifier. We make a shallow (flat) network deep (hierarchical) by growing a tree from decision regions of a given boosting classifier. The tree provides many short paths for speeding up while preserving the reasonably smooth decision regions of the boosting classifier for good generalisation. For converting a boosting classifier into a decision tree, we formulate a Boolean optimisation problem, which has been previously studied for circuit design but limited to a small number of binary variables. In this work, a novel optimisation method is proposed for, firstly, several tens of variables i.e. weak-learners of a boosting classifier, and then any larger number of weak-learners by using a two-stage cascade. Experiments on the synthetic and face image data sets show that the obtained tree achieves a significant speed up both over a standard boosting classifier and the Fast-exit—a previously described method for speeding-up boosting classification, at the same accuracy. The proposed method as a general meta-algorithm is also useful for a boosting cascade, where it speeds up individual stage classifiers by different gains. The proposed method is further demonstrated for fast-moving object tracking and segmentation problems.

T.-K. Kim (✉)
Department of Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK
e-mail: tk.kim@imperial.ac.uk

I. Budvytis · R. Cipolla
Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

Keywords Boosting · Decision tree · Decision regions · Boolean optimisation · Boosting cascade · Face detection · Tracking · Segmentation

1 Introduction

Boosting is a popular method in object detection (Viola and Jones 2001, 2004), tracking (Grabner and Bischof 2006) and segmentation (Avidan 2006) problems, where a vast number of image sub-windows, across pixels and scales, need to be classified. Doing the tasks in a reasonable time demands very fast evaluation of a classifier per window. A boosting classifier makes a fast decision by aggregating simple weak-learners such as Haar-like features, whose computations are accelerated by an integral image. Despite its efficiency, it is often required to further reduce the classification time of a boosting classifier. A cascade of boosting classifiers, which could be seen as a degenerate tree (see Fig. 1(a)), effectively improves the classification speed by filtering out majority of negative class samples in its early stages (Viola and Jones 2001, 2004; Xiao et al. 2003). Designing a cascade, however, involves manual efforts for setting a number of parameters: the number of classifier stages, the number of weak-learners and the threshold per stage. See Sect. 7.1 for details.

In this work, we propose a novel way to accelerate the classification (or evaluation) time of a boosting classifier up to an order of magnitude without sacrificing its accuracy, not relying on a conventional cascade. The chance for improvement comes from the fact that a standard boosting classifier can be seen as a very shallow network, see Fig. 1(b), where each weak-learner is a decision-stump and all weak-learners are used to make a decision. The flat structure affords smooth decision regions which help avoid overfitting