

HIGH EFFICIENCY COMPRESSION FOR OBJECT DETECTION

Hyomin Choi and Ivan V. Bajić

School of Engineering Science, Simon Fraser University, Burnaby, BC, Canada

ABSTRACT

Image and video compression has traditionally been tailored to human vision. However, modern applications such as visual analytics and surveillance rely on computers “seeing” and analyzing the images before (or instead of) humans. For these applications, it is important to adjust compression to computer vision. In this paper we present a bit allocation and rate control strategy that is tailored to object detection. Using the initial convolutional layers of a state-of-the-art object detector, we create an importance map that can guide bit allocation to areas that are important for object detection. The proposed method enables bit rate savings of 7% or more compared to default HEVC, at the equivalent object detection rate.

Index Terms— Bit allocation, rate control, HEVC, object detection, YOLO

1. INTRODUCTION

Human perceptual quality has always been among the main guiding principles of image and video compression. This influence can be seen throughout the history of development of image and video codecs: from perceptually-optimized quantization matrices in JPEG [1] to the perceptual rate control [2, 3] for High Efficiency Video Coding (HEVC) [4]. However, modern multimedia applications do not have humans as the only users. In many cases, for example surveillance and visual analytics, computers must “see” and examine images or video before humans do. Often, the first step of computer vision would be to detect objects, after which higher-level analytics such as activity recognition or anomaly detection can be performed.

Despite its importance for these applications, image and video coding tailored to computer (as opposed to human) vision has been largely unexplored. Among the few studies to tackle this topic is gradient-preserving quantization [5], which attempts to adjust quantization in image compression in order to preserve gradient information. The motivation is that gradients are useful features in a number of computer vision problems, so well-preserved gradients will likely improve the accuracy of the vision pipeline. Another recent work [6] develops a rate control scheme for H.264/AVC video coding that preserves SIFT [7] and SURF [8] features, which have also been found useful in many computer vision problems. These

studies ([5, 6]) have proposed ways to preserve well-known handcrafted features through the compression process, without focusing on any particular problem. However, the recent trend in computer vision has been away from handcrafted features and towards learnt features, especially the features learnt by deep neural networks (DNNs) [9] for specific problems.

In this paper we develop a bit allocation and rate control method that improves object detection of a DNN-based state-of-the-art object detector called YOLO9000 [10]. We utilize the outputs of the initial convolutional layers of this detector to create the importance map, which is used to guide bit allocation towards regions that are important for object detection. The resulting strategy offers significant bit savings of 7% or more compared to the default HEVC at the equivalent object detection rate. For the same bitrate, the proposed strategy offers more accurate object detection and classification compared to the default HEVC.

The paper is organized as follows. Section 2 describes the creation of object importance maps from the outputs of convolutional layers, and presents the related bit allocation and rate control strategies. Section 3 presents the experimental results and Section 4 concludes the paper.

2. PROPOSED METHODS

2.1. Background

In a convolutional neural network, convolutional layers compute cross-correlation between the input and a set of filters [9]. The cross-correlation is usually followed by max-pooling, which selects the local maximum within each small window of the cross-correlation output. Large values therefore tend to propagate through the network towards the final layers, where they contribute to the final output. It is important to appreciate that filter coefficients are computed during the training process to maximize the performance on a given task. Hence, DNN-based object detectors have filters whose coefficients have been tuned to extract the features relevant to detecting the objects that the network was trained on. And because max-pooling suppresses small outputs, it follows that large outputs are the ones that are relevant for detection.

The input size of the YOLO9000 object detector [10] is fixed at 416×416 . If an input image has different resolution, say $W \times H$, the image resolution is first scaled (while keeping

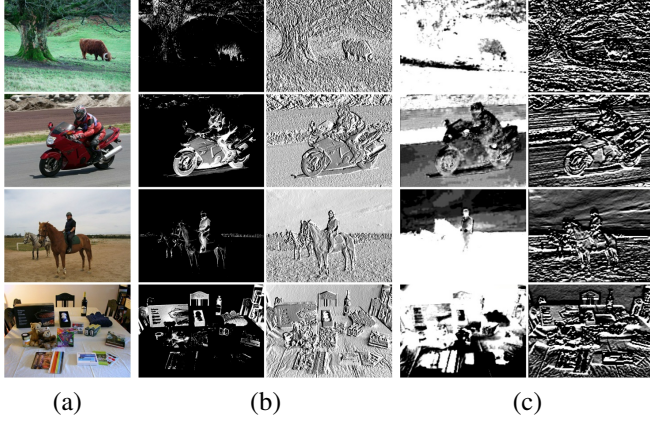


Fig. 1. Examples of (a) test images and outputs of selected filters in the (b) first and (c) third convolutional layers

the aspect ratio) and centered so that it fits the input. The scaling constants for various layers are $S_l = C_l / \max\{W, H\}$, where C_l is the spatial dimension of layer l , so $C_1 = 416$, $C_2 = 208$, etc.

The first convolutional layer employs 32 filters with kernel size 3×3 , and produces 32 outputs. This is followed by max-pooling over 2×2 windows. The subsequent convolutional layers operate on the previous layer's outputs. There are total 32 layers in the YOLO9000 architecture. Fig. 1 shows several input images, and the corresponding outputs of several filters in the first and third convolutional layer. The brighter pixels in the output indicate the higher correlation with the associated filter. As seen in this figure, even the early layers in the convolutional network are able to provide some information about the objects, although precise object location and class is not available until upper layers of the network complete their processing.

Based on this reasoning, we propose an object detection-friendly compression framework shown in Fig. 2. The input image is processed by the initial convolutional layers of the object detector. The filters in each layer can run in parallel, so this process is highly parallelizable. From the resulting filter outputs, we construct an object importance map, which guides bit allocation and rate control in HEVC. The resulting image turns out to be more object detection-friendly, as demonstrated in Section 2.2.

2.2. Object importance map

The object importance map is meant to indicate how important is each pixel to object detection. The YOLO9000 architecture employs leaky activation, which means that layer outputs can be negative. We first clamp the outputs to the range $[0, 1]$ as

$$\hat{\psi}_l^{(n)}(x, y) = \max \left\{ 0, \min \left\{ 1, \psi_l^{(n)}(x, y) \right\} \right\} \quad (1)$$

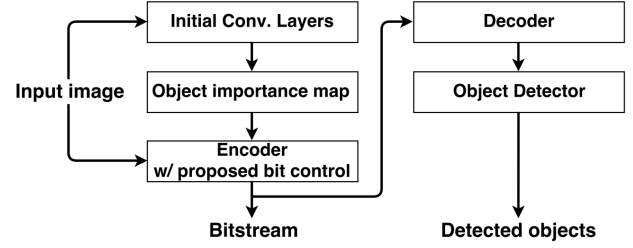


Fig. 2. The proposed object detection-friendly compression framework

where l indicates the layer, x and y are spatial coordinates, and n is the filter index in the given layer. Then, all clamped outputs are stacked in a tensor

$$v_l(x, y) = \left[\alpha_l^{(1)} \hat{\psi}_l^{(1)}(x, y), \alpha_l^{(2)} \hat{\psi}_l^{(2)}(x, y), \dots, \alpha_l^{(N)} \hat{\psi}_l^{(N)}(x, y) \right] \quad (2)$$

where $\alpha_l^{(n)}$ is a weight factor for the n -th filter in layer l .

The weights are meant to indicate how informative is a particular filter's output for a given input image. Ideally, the filter's output would be high near the objects of interest and low elsewhere. As seen in Fig. 1, filters' outputs are not equally informative about the objects in the image. Moreover, a certain filter may be very informative on one image, and not very informative on another image, which means that weights should be adapted from image to image. We experimented with entropy of the filter output as a guide to set weights (lower entropy inducing higher weight), but eventually settled for a simpler approach that gave slightly better results. In particular, we set the weight as 1 minus the average clamped output:

$$\alpha_l^{(n)} = 1 - \frac{1}{W_l \cdot H_l} \sum_{y=0}^{H_l-1} \sum_{x=0}^{W_l-1} \hat{\psi}_l^{(n)}(x, y) \quad (3)$$

where $W_l = W/S_l$ and $H_l = H/S_l$ are the width and height of the filter's output on the particular image at level l . When the filter produces high responses across the entire image (i.e., it is not very informative), the average is high, so its weight becomes low. If the filter's output is low on average, its weight becomes high. Therefore, $v_l(x, y)$ in Eq. (2) will be high only when the filter is informative (high weight) and has a high response at the particular (x, y) . Finally, we take the ℓ^2 norm of $v_l(x, y)$, $O_l(x, y) = \|v_l(x, y)\|_2$, and then normalize $O_l(x, y)$ by linearly mapping it to the range $[0, 1]$ to produce the final importance map $\tilde{O}_l(x, y)$. Figure 3 shows several importance maps generated from the first and third layer on different images.

2.3. Bit allocation and rate control

The proposed bit allocation makes use of the object importance map $\tilde{O}_l(x, y)$ to decide how to spend bits.

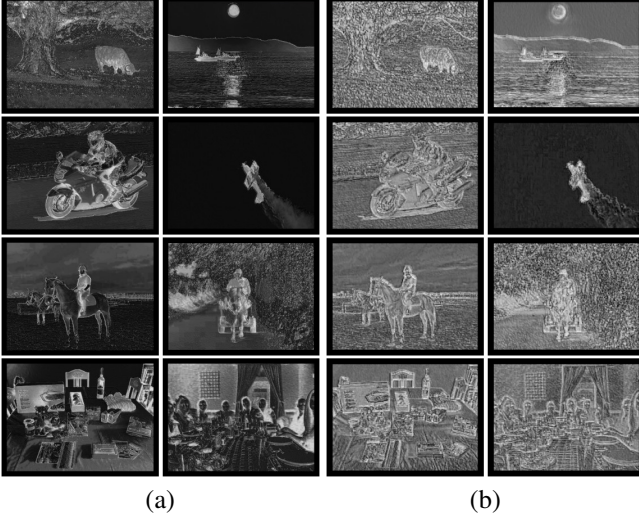


Fig. 3. The object importance maps combined using the outputs of the (a) first and (b) third layer

First, the pixel-wise importance $\tilde{O}_l(x, y)$ is converted into block-wise importance $I(i, j)$. The size of the block is the size of the corresponding coding unit scaled by S_l . Then $I(i, j)$ is computed by summing $\tilde{O}_l(x, y)$ within the corresponding block and dividing by the total sum over the importance map. From here on, we use (i, j) as the coordinates of the top-left corner of the block.

We then calculate the initial coarse estimates of the bits per pixel (bpp_{coarse}) for each block as

$$bpp_{coarse}(i, j) = \frac{1}{N_{pixels}(i, j)} \cdot I(i, j) \cdot T_{bits} \quad (4)$$

where $N_{pixels}(i, j)$ is the number of pixels in the block whose top-left corner is at (i, j) and T_{bits} represents the number of target bits for the image to be coded. In Eq. (4), the calculated $bpp_{coarse}(i, j)$ could possibly be zero, which turns out to be harmful in subsequent encoding. In order to refine this coarse estimate, we run the R- λ model [11, 12], which is default rate control model in the HEVC reference software [13]. Specifically, by inputting T_{bits} and $bpp_{coarse}(i, j)$ into the R- λ model, we compute the slice/picture-level QP_s and the block-level preliminary quantization parameter $QP_p(i, j)$. Note that $QP_p(i, j)$ is bounded by $QP_s \pm 2$ by default, however we extend the range to $QP_s \pm 3$. Then, the preliminary bits per pixel $bpp_p(i, j)$ is computed by inverting the R- λ model with $QP_p(i, j)$ as the input. Finally, $QP_p(i, j)$ is incremented by 1 if $I(i, j) = 0$.

The final block importance is computed as

$$I_F(i, j) = \frac{bpp_p(i, j)}{\sum \sum bpp_p(i, j)} \quad (5)$$

where the double summation is over all valid (i, j) . Using this, we compute the weight for each block as the normalized

importance of that block

$$w(i, j) = \frac{I_F(i, j)}{\sum \sum I_F(i, j)} \quad (6)$$

where again the double summation is over all valid (i, j) . Then the total target bits for the block can be computed as

$$T_{bits}^{Blk}(i, j) = \left\{ L_{bits}(i, j) + \frac{L_{blk}(i, j)(L_{bits}(i, j) - L_{bits}^{Est}(i, j))}{SW} \right\} w(i, j) \quad (7)$$

where $L_{bits}(i, j)$ are the remaining bits in the bit budget before coding block (i, j) , $L_{blk}(i, j)$ is the number of remaining blocks to be coded, including (i, j) , and SW is the sliding window size used to smooth out the bit variation (we use the default $SW = 4$ from the HM16.12 [13]). $L_{bits}^{Est}(i, j)$ is an estimate of the bits that will actually be used for coding the block (i, j) and all subsequent blocks, computed as

$$L_{bits}^{Est}(i, j) = \sum_{u \geq i} \sum_{v \geq j} bpp_p(u, v) \cdot N_{pixels}(u, v) \quad (8)$$

where the summations are starting at (i, j) and going over all subsequent valid block indices (u, v) .

Finally, we estimate the actual QP values QP_a by inputting $T_{bits}^{Blk}(i, j)/N_{pixels}(i, j)$ to the R- λ model. The resulting QP_a is bounded by $QP_s \pm 2$ by default. However, we shift the bound upward as $[QP_s, QP_s + 4]$ if $QP_p \geq QP_s + 3$. This has the potential to save bits in less important regions.

3. EXPERIMENTS

In this section, we assess the performance of the proposed bit allocation and rate control scheme in terms of its effect on object detection. The proposed methods were implemented in HEVC reference software HM16.12 [13]. The YOLO9000 model in the Darknet framework [14] is used for the object detection performance evaluation.

Bjontegaard Delta (BD) [15] is a standard measurement method for evaluating compression performance. It compares the average bit rates of two coding methods at the equivalent quality metric. Usually, the quality metric is Peak Signal-to-Noise Ratio (PSNR) and we refer to this measurement as BD bitrate for PSNR (BD-BR-PSNR). However, it is also possible to use quality metrics other than PSNR in the BD analysis. Specifically, since our goal is to compare object detection performance between methods, instead of PSNR we use the standard object detection accuracy metric called mean Average Precision (mAP) [16]. The mAP is in the range $[0, 1]$. By computing BD over rate vs. mAP curves, we can obtain the average bit rate saving (or increment) that one compression method would have over another at the equivalent mAP. We call this metric BD bitrate for mAP (BD-BR-mAP).

For testing, we employ the widely used PASCAL VOC 2007 dataset [16], which has 9963 images out of which 4952

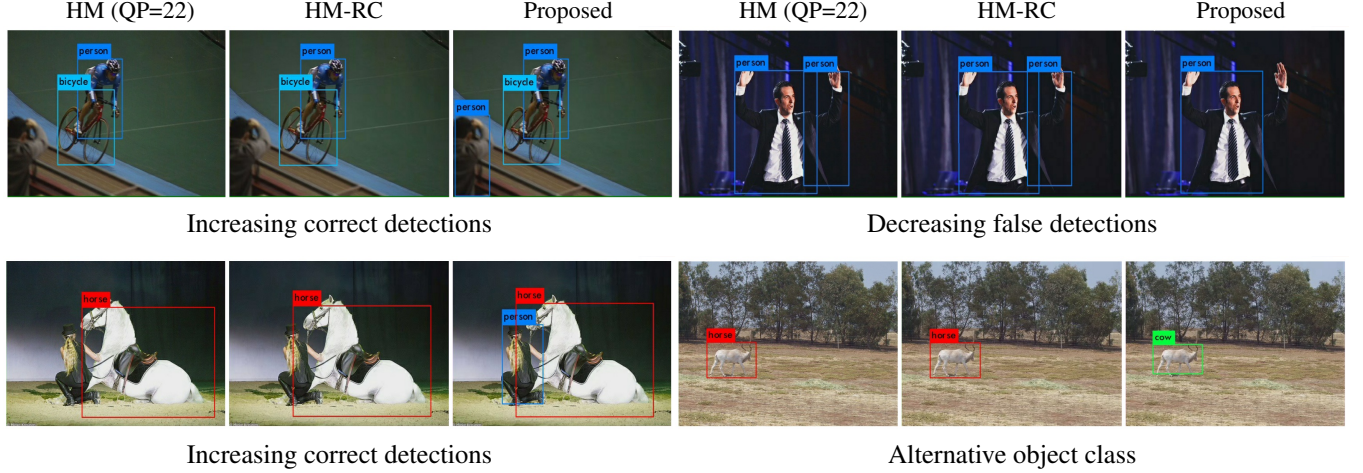


Fig. 4. Examples illustrating different object detection/classification produced by YOLO9000 on images encoded by HM with QP=22, HM-RC at the same bitrate, and the proposed method at the same bitrate, with importance maps computed from the third convolutional layer.

are test images. The images are annotated with 20 different object classes, such as aeroplane, bicycle, bird, and so on. For encoding, 16×16 CTU is adopted and RDOQ tool is off, but other coding parameters follow the common HEVC test conditions [17] of the Main Still Picture Profile [18]. We first encode each test image using the default HM with $QP \in \{22, 27, 32, 37\}$. The resulting bits are used as target bits for the default HM rate control (HM-RC) and our proposed method. For the proposed method we construct the importance maps from the outputs of the first, third, and seventh layer, in order to examine the behaviour of the system with different importance maps.

All encoded images are then decoded and fed to the YOLO9000 object detector. mAP is computed by comparing detector's output with the ground truth. Table 1 shows various comparisons among the three tested codecs: HM, HM-RC and proposed. For the rate control accuracy, Δ_{bpp} is the mean absolute difference (MAD) in bits per pixel (bpp) between the output bits of HM and the two rate control methods (HM-RC and proposed) across all images. HM-RC shows averaged $\Delta_{bpp} = 0.0483$, while our rate control gave smaller deviation in each of the three cases, with importance map computed from the seventh layer being the most accurate. In terms of BD-BR-PSNR, both HM-RC and our proposed method have lower performance (positive BD-BR-PSNR) compared to the default HM, since they both deviate from the optimal rate-distortion allocation in order to achieve different objectives. However, our method achieves significant advantage in BD-BR-mAP over both HM and HM-RC, which was the main design objective. In particular, with the importance map computed from the output of the third convolutional layer, 7.32% bit reduction is achieved over HM, and 8.23% reduction over HM-RC, at an equivalent mAP. This shows that importance

maps can successfully guide bit allocation towards regions that are most relevant for object detection. Fig. 4 shows a few examples where YOLO9000 produces different detections on the images encoded by HM, HM-RC, and proposed methods. Although the images look very similar visually, detection on images encoded by the proposed method is the most accurate. This again illustrates that importance maps successfully guide bit allocation towards regions that are most relevant for object detection.

4. CONCLUSION

We proposed a novel bit allocation and rate control strategy whose goal was to improve object detection after decoding. Using the outputs of the initial convolutional layers of a state-of-the-art object detector, the proposed algorithm successfully achieved efficient bit control and improved object detection performance over the default HEVC implementations. The proposed strategy can be used in many applications where computers “see” and analyze the data before (or instead of) humans.

Table 1. Performance comparison among HM, HM-RC, and proposed method

Test cases	Δ_{bpp}	σ_{bpp}	BD-BR-PSNR	BD-BR-mAP
HM vs. HM-RC	0.0483	0.1187	3.08%	1.67%
HM vs. Ours w/ 1st L.	0.0385	0.1113	7.10%	-3.90%
HM vs. Ours w/ 3rd L.	0.0372	0.1094	7.15%	-7.32%
HM vs. Ours w/ 7th L.	0.0232	0.1086	6.96%	-6.33%
HM-RC vs. Ours w/ 1st L.	-	-	3.82%	-5.31%
HM-RC vs. Ours w/ 3rd L.	-	-	3.87%	-8.23%
HM-RC vs. Ours w/ 7th L.	-	-	3.68%	-7.10%

5. REFERENCES

- [1] ITU-T Rec. T.81 and ISO/IEC 10918-1, "Information technology - digital compression and coding of continuous-tone still images," 1992.
- [2] H. Zeng, K. N. Ngan, and M. Wang, "Perceptual adaptive lagrangian multiplier for high efficiency video coding," in *Picture Coding Symposium*, Dec. 2013.
- [3] A. Yang, H. Zeng, L. Ma, J. Chen, C. Cai, and K.-K. Ma, "A perceptual-based rate control for HEVC," in *6th Int. Conf. Image Processing Theory Tools and Applications*, Dec. 2016.
- [4] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 22, pp. 1649–1668, Dec. 2012.
- [5] M. Makar, H. Lakshman, V. Chandrasekhar, and B. Girod, "Gradient preserving quantization," in *Proc. IEEE ICIP'12*, Sep. 2012, pp. 2505–2508.
- [6] J. Chao, R. Huitl, E. Steinbach, and D. Schroeder, "A novel rate control framework for SIFT/SURF feature preservation in H.264/AVC video compression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 25, no. 6, pp. 958–972, Jun. 2015.
- [7] D. Lowe, "Distinctive image feature from scale-invariant keypoints," *Int. J. Computer Vision*, vol. 60, no. 2, pp. 91–110, Nov. 2004.
- [8] H. Bay, T. Tuytelaars, and L. V. Gool, "Surf: Speeded up robust features," in *Proc. European Conf. Computer Vision*, May. 2006, pp. 404–417.
- [9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, 2016.
- [10] J. Redmon and A. Farhadi, "YOLO9000: better, faster, stronger," in *IEEE Conf. Computer Vision and Pattern Recognition*, Jul. 2017.
- [11] B. Li, H. Li, L. Li, and J. Zhang, " λ domain rate control algorithm for high efficiency video coding," *IEEE Trans. Image Processing*, vol. 23, no. 9, pp. 3841–3854, Sep. 2014.
- [12] B. Li, J. Xu, D. Zhang, and H. Li, "QP refinement according to Lagrange multiplier for high efficiency video coding," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2013, pp. 477–480.
- [13] "HEVC reference software (HM 16.12)," <https://hevc.hhi.fraunhofer.de/trac/hevc/browser/tags/HM-16.12>, Accessed: 2017-05-27.
- [14] J. Redmon, "Darknet: Open source neural networks in C.," <http://pjreddie.com/darknet/>, 2013-2017, Accessed: 2017-10-19.
- [15] G. Bjøntegaard, "VCEG-M33: Calculation of average PSNR differences between RD curves," in *Video Coding Experts Group (VCEG)*, Apr. 2001.
- [16] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [17] F. Bossen, "Common HM test conditions and software reference configurations," in *ISO/IEC JTC1/SC29 WG11 m28412, JCTVC-L1100*, Jan. 2013.
- [18] T. Nguyen and D. Marpe, "Objective performance evaluation of the HEVC Main Still Picture profile," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 25, no. 5, May 2015.