

# Speeding up VP9 Intra Encoder with Hierarchical Deep Learning Based Partition Prediction

Somdyuti Paul, Andrey Norkin, and Alan C. Bovik, *Fellow, IEEE*

**Abstract**—In VP9 video codec, the sizes of blocks are decided during encoding by recursively partitioning  $64 \times 64$  superblocks using rate-distortion optimization (RDO). This process is computationally intensive because of the combinatorial search space of possible partitions of a superblock. Here, we propose a deep learning based alternative framework to predict the intra-mode superblock partitions in the form of a four-level partition tree, using a hierarchical fully convolutional network (H-FCN). We created a large database of VP9 superblocks and the corresponding partitions to train an H-FCN model, which was subsequently integrated with the VP9 encoder to reduce the intra-mode encoding time. The experimental results establish that our approach speeds up intra-mode encoding by 69.7% on average, at the expense of a 1.71% increase in the Bjøntegaard-Delta bitrate (BD-rate). While VP9 provides several built-in speed levels which are designed to provide faster encoding at the expense of decreased rate-distortion performance, we find that our model is able to outperform the fastest recommended speed level of the reference VP9 encoder for the *good* quality intra encoding configuration, in terms of both speedup and BD-rate.

**Index Terms**—VP9, video encoding, block partitioning, intra prediction, convolutional neural networks, machine learning.

## I. INTRODUCTION

VP9 has been developed by Google [1] as an alternative to mainstream video codecs such as H.264/AVC [2] and High Efficiency Video Coding (HEVC) [3] standards. VP9 is supported in many web browsers and on Android devices, and is used by online video streaming service providers such as Netflix and YouTube.

As compared to both its predecessor, VP8 [4] and H.264/AVC video codecs, VP9 allows larger prediction blocks, up to size  $64 \times 64$ , which results in a significant improvement in coding efficiency. In VP9, sizes of prediction blocks are decided by a recursive splitting of non-overlapping spatial units of size  $64 \times 64$ , called superblocks. This recursive partition takes place at four hierarchical levels, possibly down to  $4 \times 4$  blocks, through a search over the possible partitions at each level, guided by a rate-distortion optimization (RDO) process. The Coding tree units (CTUs) in HEVC, which are analogous to VP9's superblocks, have the same default

Manuscript received June 21, 2019; revised November 30, 2019 and May 28, 2020; accepted July, 11 2020. This work was supported by Netflix, Inc. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Aline Roumy. (*Corresponding author: Somdyuti Paul*)

S. Paul and A. C. Bovik are with the Department of Electrical and Computer Engineering at The University of Texas at Austin, Austin, TX, 78712 USA (email: somdyuti@utexas.edu, bovik@ece.utexas.edu).

A. Norkin is with Netflix, Inc. Los Gatos, CA, 95032 USA (email: anorkin@netflix.com).




Fig. 1. Recursive partition of VP9 superblock at four levels showing the possible types of partition at each level.

maximum size of  $64 \times 64$  and minimum size of  $8 \times 8$ , which can be further split into smaller partitions ( $4 \times 4$  in the intra-prediction case). However, while HEVC intra-prediction only supports partitioning a block into four square quadrants, VP9 intra-prediction also allows rectangular splits. Thus, there are four partition choices at each of the four levels of the VP9 partition tree for each block at that level: no split, horizontal split, vertical split and four-quadrant split. This results in a combinatorial complexity of the partition search space since the square partitions can be split further. A diagram of the recursive partition structure of VP9 is shown in Fig. 1.

Although the large search space of partitions in VP9 is instrumental to achieve its rate-distortion (RD) performance, it causes the RDO based search to incur more computational overhead as compared to VP8 or H.264/AVC, making the encoding process slower. Newer video codecs, such as AV1 [5] and future Versatile Video Coding (VVC) [6], allow for prediction units of sizes from  $128 \times 128$  to  $4 \times 4$ , giving rise to even deeper partition trees. As ultra high-definition (UHD) videos become more popular, the need for faster encoding algorithms will only escalate. One important way to approach this issue is to reduce the computational complexity of the RDO-based partition search in video coding.

While state of the art performances in visual data processing technologies, such as computer vision, rely heavily on deep learning, mainstream video coding and compression technology remains dominated by traditional block-based hybrid codecs. However, deep learning based image compression

techniques such as [7]–[9] are also being actively explored, and have shown promise. A few of such image based deep learning techniques have also been extended to video compression with promising results [10]–[12]. A second category of work uses deep learning to enhance specific aspects of video coding, such as block prediction [13], [14], motion compensation [15], [16], in-loop filtering [17]–[19], and rate control [20], with the objective of improving the efficiency of specific coding tools. Given these developments, there is a possibility of a future paradigm shift in the domain of video coding towards deep learning based techniques. Our present work is motivated by the success of deep learning-based techniques such as [21], [22], on the current and highly practical task of predicting the HEVC partition quad-trees.

In this paper, we take a step in this direction by developing a method of predicting VP9 intra-mode superblock partitions in a novel bottom-up way, by employing a hierarchical fully convolutional network (H-FCN). Unlike previous methods of HEVC partition prediction, which recursively split blocks starting with the largest prediction units, our method predicts block merges recursively, starting with the smallest possible prediction units, which are  $4 \times 4$  blocks in VP9. The bottom-up approach optimized using an H-FCN model allows us to use a much smaller network than [22]. By integrating the trained model with the reference VP9 encoder, we are able to substantially speed up intra-mode encoding at a reasonably low RD cost, as measured by the Bjøntegaard delta bitrate (BD-rate) [23], which quantifies differences in bitrate at a fixed encoding quality level relative to another reference encode. Our method also surpasses the higher speed levels of VP9 in terms of speedup, while maintaining a lower BD-rate as we show in the experimental results.

The main steps of our work presented in this paper can be summarized as follows:

- 1) We created a large and diverse database of VP9 intra encoded superblocks and their corresponding partition trees using video content from the Netflix library.
- 2) We developed a fast H-FCN model that efficiently predicts VP9 intra-mode superblock partition trees using a bottom-up approach.
- 3) We integrated the trained H-FCN model with the VP9 encoder to demonstrably speed up intra-mode encoding.

These steps are summarized by the block diagram in Fig. 2. The source code of our model implementation, including the modifications made to the reference VP9 decoder and encoder for creating the database of superblock partitions, and using the trained H-FCN model for faster intra encoding, respectively, is available online at [24].

The rest of the paper is organized as follows. In Section II, we briefly review earlier works relevant to the current task. Section III describes the VP9 partition database that we created to drive our deep learning approach. Section IV elaborates the proposed method. Experimental results are presented in Section V. Finally, we draw conclusions and provide directions for future work in Section VI.

## II. RELATED WORK

The earliest machine learning based methods that were designed to infer the block partition structures of coded videos from pixel data relied heavily on feature design. A decision tree based approach was used to predict HEVC partition quadtrees for intra frames from features derived from the first and second order block moments in [25], reportedly achieving a 28% reduction in computational complexity along with 0.6% increase in BD-rate. Using a support vector machine (SVM) classifier on features derived from measurements of the variance, color and gradient distributions of blocks, a 36.8% complexity reduction was gained against a 3% increase in BD-rate in [26], on the screen content coding extension of HEVC in the intra-mode.

With the advent of deep learning techniques in recent years, significant further breakthroughs were achieved [21], [22], [27]. In [21], a parallel convolutional neural network (CNN) architecture was employed to reduce HEVC intra encoding time by 61.1% at the expense of a 2.67% increase in BD-rate. Three separate CNN models were used to learn the three-level intra-mode partition structure of HEVC in [27], obtaining an average savings of 62.2% of encoding time against a BD-rate increase of 2.12%. This approach was extended in [22] to reduce the encoding time of both intra and inter modes using a combination of a CNN with a long short-term memory (LSTM) architecture. This approach reduced the average intra-mode encoding time by 56.9-66.5% against an increase of 2.25% in BD-rate, while in the inter mode, a 43.8%-62.9% average reduction was obtained versus an increase of 1.50% in BD-rate.

However, there has been little work reported on the related problem of reducing the computational complexity of RDO based superblock partition decisions in VP9, and even less work employing machine learning techniques. A multi-level SVM based early termination scheme for VP9 block partitioning was adopted in [28], which reduced encoding time by 20-25% against less than a 0.03% increase in BD-rate in the inter mode. Although superblock partition decisions using RDO consume bulk of the compute expense of intra-mode encoding in VP9, to the best of our knowledge, there has been no prior work on predicting the complete partition trees of VP9 superblocks.

The problem of VP9 superblock partition prediction is a hierarchical decision process, which involves choosing one of four types of partitions for each block, at every level of the partition tree. A hierarchical structure was introduced in [29], using a two-level CNN that was trained to perform coarse-to-fine category classification, achieving improvements in the classification accuracy. Based on similar principles, [30] extended the hierarchical classification approach to more than two levels, using a branched CNN model with multiple output layers yielding coarse-to-fine category predictions. The ability of these architectures to capture the hierarchical relationships inherent in visual data motivates our model design. At the same time, unlike global image tasks such as classification, inferring partition trees from superblock pixels is a spatially dense prediction task. For example, on a  $64 \times 64$  superblock,

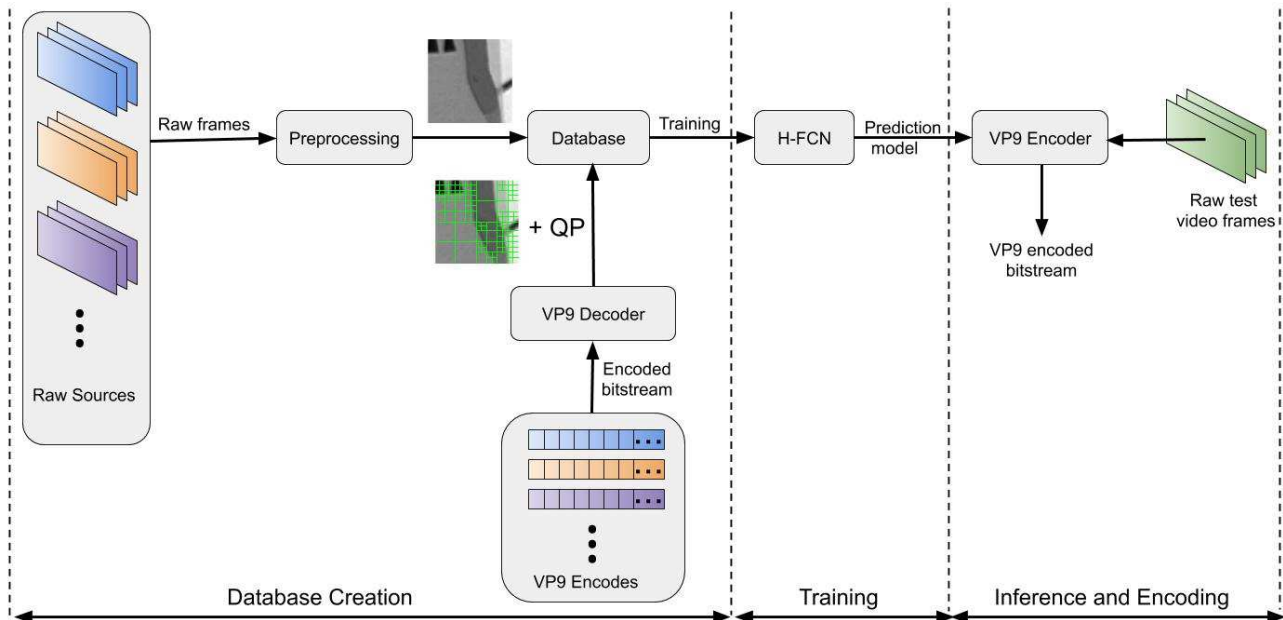


Fig. 2. Flow diagram of our VP9 partition prediction. The inputs to the database are schematically shown for each superblock.

there can be a maximum of  $64 \times 8$  blocks whose partitions are to be inferred. This means that as many as 64 localized predictions must be made on a  $64 \times 64$  superblock at the lowest level of the partition tree.

Fully convolutional networks (FCNs) have been shown to perform remarkably well on a variety of other spatially dense prediction tasks, such as semantic segmentation [31], [32], depth estimation [33], saliency detection [34], object detection [35], visual tracking [36] and dense captioning [37]. Moreover, convolution layers are typically faster than fully connected layers for similar input and output sizes. Likewise, we have found that the H-FCN model that we have developed and explain here imparts the ability to simultaneously handle hierarchical prediction and dense prediction with significant speedup.

### III. VP9 INTRA-MODE SUPERBLOCK PARTITION DATABASE

In order to facilitate a data-driven training of our H-FCN model, we constructed a large database of VP9 intra encoded superblocks, corresponding QP values, and partition trees. In the absence of a publicly available superblock partition database for VP9 similar to the one developed in [22] for HEVC CTU partitions, this was a necessary first step for our work.

#### A. Partition Tree Representation

Since the number of possible partition trees of a superblock is too large to be represented as distinct classes in a multi-class classification problem, we need a simple and concise description of the partition tree to ensure effective learning. The partition tree representation we adopt is similar to [22], but represents block merges instead of splits to facilitate bottom-up prediction. In our approach, the partition tree is represented by

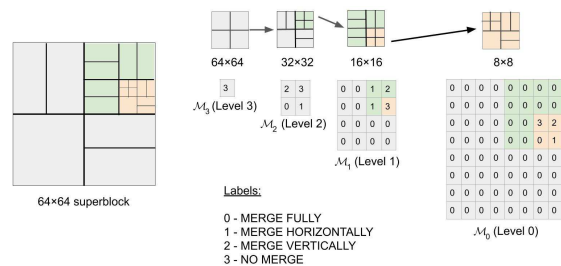


Fig. 3. Matrix representation of the four-level partition tree.

four matrices  $\mathcal{M}_0, \dots, \mathcal{M}_3$  that correspond to the four levels of the VP9 partition tree. An example of a superblock partition tree is illustrated in Fig. 3. The four possible merges of the blocks at each level (including the possibility of no merge) are indicated by the numbers 0 to 3 as shown in Fig. 3. Each element of the matrices indicates the type of merge of the group of four blocks corresponding to that element's location at that level. For example, each element of the  $8 \times 8$  matrix  $\mathcal{M}_0$  indicates how the four  $4 \times 4$  blocks at corresponding locations are merged at level 0. Similarly,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$  and  $\mathcal{M}_3$  represent merges of non-overlapping groups of four  $8 \times 8$ ,  $16 \times 16$  and  $32 \times 32$  blocks, respectively, at the higher levels. We denote the partition tree of a superblock by  $\mathcal{P} = \{\mathcal{M}_0, \dots, \mathcal{M}_3\}$ . This succinct representation allows us to formulate the problem as a multi-level, multi-class classification task, with 4 levels, and 4 classes corresponding to each matrix element.

#### B. Database Creation

Our database was created using content from the Netflix video catalog, encoded using the reference VP9 encoder from the *libvpx* package [38], in VP9 profile 0 (8 bits/sample and 4:2:0 chroma subsampling), using speed level 1 and the quality setting *good*. The contents were drawn from 89

cinematic productions and 17 television episodes, each from a unique television series. The contents selected were drawn from different genres, such as action, drama, animation, etc.

We encoded each content at three resolutions:  $1920 \times 1080$ ,  $1280 \times 720$  and  $960 \times 540$ . The partition pattern selected by the RDO on each superblock depends on both its visual content, as well as the QP value chosen for that superblock. Thus, we modified the VP9 decoder from the *libvpx* package to record the computed partition tree  $\mathcal{P}$  of each superblock, in the form described in Section III-A, along with the corresponding QP value  $Q$  while decoding the intra frames of the VP9 encoded bitstreams.

In order to obtain the raw pixel data corresponding to the superblocks of each VP9 encoded video, the source videos were converted to a YCbCr 4:2:0 8-bit representation, then downsampled to the encode resolution via Lanczos resampling, if the source and the encode were at different resolutions. Following this process, the luma channels of non-overlapping  $64 \times 64$  blocks were extracted from the source frames at the encode resolution. These steps constitute the preprocessing stage shown in Fig. 2. Denote this superblock pixel data by  $\mathcal{S}$ . Thus, each sample of our database may be expressed by a tuple  $(\mathcal{S}, Q, \mathcal{P})$ . If the frame width or height (or both) is not exactly divisible by 64, the VP9 encoder zero pads the frame boundaries to construct superblocks of size  $64 \times 64$  at the boundaries during encoding. However, we excluded the boundary superblocks with partial zero padding from our database.

The HEVC intra-mode partition database [22] is limited to only 4 QP values. By contrast, our database encompasses internal QP values in the range 8-105, where the internal QP value range for VP9 is 0-255 (the corresponding external QP value range is 0-63). The range of QP values used in our database is a practical range used for encoding intra frames in adaptive streaming, where rather than using higher QP values to stream videos at low target bitrates, a lower resolution video is streamed, which is suitably upsampled later at the viewers' end [39]. We refer the reader to Appendix B for a comparison of the encoding performance on our database against on the intra-mode database of [22].

Our new database contains content from only a single episode of each television series. We divided our database into training and validation sets, as summarized in Table I, where the letters M and E indicate cinematic "movies" and television "episodes," respectively. The division was conducted such that there was no overlap in content between the training and validation sets, i.e. entirely different movies and television series were used for each set. Table I also specifies the percentage of computer graphics image (CGI) content in each set. Although the superblocks database of Netflix content cannot be made publicly available (due to the content licenses), we do provide the code for the modified VP9 decoder at [24], which can be used to generate a similar database from a set of VP9 encoded bitstreams and their original source videos. The test set was created from publicly available content, different from the Netflix contents used in the training and validation sets, as described later in Section V-D.

TABLE I  
SUMMARY OF VP9 INTRA-MODE SUPERBLOCK PARTITION DATABASE

Database	Contents	% of CGI content	# of samples
Training	62 (M) + 12 (E)	12.16	11,990,384
Validation	27 (M) + 5 (E)	12.50	4,698,195

#### IV. PROPOSED METHOD

As mentioned earlier, in our approach the partition tree is constructed in a bottom-up manner, where merges of the smallest possible blocks, which are of size  $4 \times 4$  in VP9, are predicted first at the lowest level of the partition tree, followed by merge predictions on larger blocks at the upper levels. This is unlike the coarse-to-fine approaches used in many image analysis applications, but it is well-motivated here. In a multi-layered CNN model, the early layers learn simple low level features depicting local image characteristics, such as edges and corners. These simple features are progressively combined to form textures and more complex features representing higher level global semantics by the deeper layers. The intuition behind our bottom-up block merge strategy follows from this observation, which is supported by feature visualization techniques applied to successive CNN layers [40]. Since the smaller blocks are more spatially localized, low-level local features, such as edges and corners that are learned by the early layers of a CNN, should be adequate to predict the merge types of  $4 \times 4$  blocks i.e. the partition types of  $8 \times 8$  blocks. The most notable benefit of this approach is that early prediction of the merge types of the smaller blocks at the lower levels of the partition tree, saves computation time and network parameters that can instead be invested in predicting the merge types of larger blocks at the upper levels, on which the CNN needs to learn more complex patterns. This allows us to design a deeper network for the higher levels, having many fewer trainable parameters than for example, the ETH-CNN of [22] with three parallel branches and 1,287,189 trainable parameters, which uses just three convolutional layers at each level of prediction, despite the larger number of trainable parameters. A deeper network is able to learn additional levels of hierarchical abstractions of the data, which is relevant in the context of hierarchical partition prediction in VP9. An analysis of the efficacy of this bottom-up prediction strategy is given in Appendix A, by comparing it with an alternative top-down strategy.

##### A. H-FCN Architecture

The architecture of our H-FCN model is shown in Fig. 4. It consists of four output branches and a *trunk* from which the branches emanate. This branched architecture is similar to [30]. However, unlike [30], our model is fully convolutional and uses a bottom-up prediction scheme. The inputs to the model are the superblocks  $\mathcal{S}$  and corresponding QP values  $Q$ . The trunk has a conventional CNN structure with convolutional layers followed by rectified linear unit (ReLU) nonlinearities [41] and using batch normalization [42]. Also  $2 \times 2$  max pooling is applied following every two convolutional layers in the trunk. The outputs of the model are the matrices



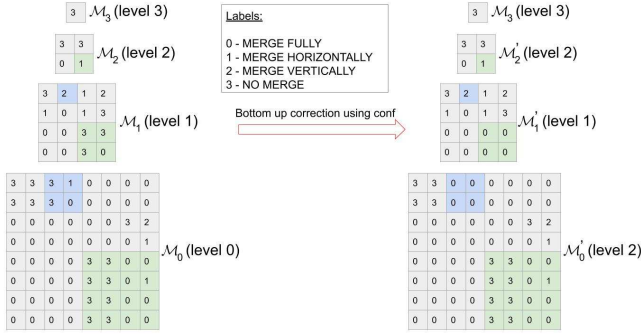


Fig. 5. Top-down inconsistency correction.

superblocks completely contained within the frame boundaries were collectively predicted as a batch, using the integrated model.

Since the different levels of the superblock partition are modeled independently of one another, the predictions of any two adjacent levels might be mutually inconsistent, producing an invalid partition tree. For example, a “full merge” may be predicted for a group of four  $16 \times 16$  blocks at level 2, whereas the four  $8 \times 8$  subblocks within one of these  $16 \times 16$  blocks can be predicted to have no merge in level 1. This is inconsistent, because in this case the merger of four  $16 \times 16$  blocks indicates that a  $32 \times 32$  block is not split, and thus the  $16 \times 16$  subblocks within it cannot have a split either, although a split is required by the “no merge” prediction of the corresponding group of  $8 \times 8$  blocks. Thus, it may be necessary to correct the predictions in  $\mathcal{P}$  to obtain a valid partition tree that can be used while encoding. We devised a top-down correction procedure that is illustrated in Fig. 5, where the colored regions are inconsistent with predictions at other levels. Consistency is enforced between adjacent levels by correcting any of the other three merge predictions to “full merge,” at all such blocks at the lower of the two levels enclosed by a larger block predicted to have “full merge” at the higher level. Beginning with level 2, predictions at each level are successively corrected to be consistent with its corrected next higher level in this manner. In other words, first  $\mathcal{M}_2$  at level 2 is corrected to be consistent with  $\mathcal{M}_3$ ; let the corrected matrix at level 2 be  $\mathcal{M}'_2$ . At the next step,  $\mathcal{M}_1$  is corrected to be consistent with  $\mathcal{M}'_2$  to obtain  $\mathcal{M}'_1$  and so on. At the end of this procedure, we have  $\mathcal{P}' = \{\mathcal{M}_3, \mathcal{M}'_2, \mathcal{M}'_1, \mathcal{M}'_0\}$ , where  $\mathcal{P}'$  is a valid partition tree.

Although the predictions made by the H-FCN are independent at each level, they were found to be remarkably consistent across levels. Our experiments revealed that only about 5.17% of the predicted superblock partition trees from the validation set were inconsistent, and thus needed the aforementioned correction. This suggests that although the consistency requirement was not explicitly enforced, the H-FCN model implicitly learns to make consistent merge predictions most of the time. The motivation behind our choice of the inconsistency correction approach is explained in Section V-D.

Predictions in  $\mathcal{P}'$  are then ordered to form a preorder




Fig. 6. H-FCN loss with training progress.

traversal of the partition tree from left to right and top to bottom, which corresponds to the order in which the blocks are encoded in VP9. The predictions thus ordered are then recursively used to replace the RDO module of the encoder to decide the partitions of all the superblocks except those extending beyond frame boundaries. Since these boundary cases were not included in our database, we simply invoke the RDO module to encode them.

## V. EXPERIMENTAL RESULTS

In order to experimentally evaluate the performance of our H-FCN model relative to RDO, we compared the encoding time and RD performance of the two approaches on a set of test video sequences at three different resolutions. To further validate the efficacy of our approach, we also extended the comparison to include the highest recommended speed level of VP9 for our encoding configuration (*cpu-used* 4 at *good* quality).

### A. System Settings

We developed and trained our H-FCN model using Tensorflow (version 1.12) with the Keras API. The model was trained on a system with an Intel Core i7-6700K CPU @4 GHz, with 8 cores and 32 GB RAM running a 64 bit Ubuntu 16.04 operating system. The training was accelerated with a Nvidia Titan X Pascal GPU with 12 GB of memory.

Since the loss is not calculated during inference, the matrix slicing operation to generate multiple outputs is unnecessary, and removing it reduces the inference time by a significant fraction in our Keras implementation. Thus, we removed the slicing layers during inference as indicated by the dashed box in Fig. 4, directly obtaining the matrices  $\mathcal{M}_0, \dots, \mathcal{M}_3$  as the network outputs. The trained Keras model was then converted to a Tensorflow computation graph prior to deployment in the VP9 encoder. Integration of the trained model with the reference VP9 encoder was done using *libvpx* version 1.6.0, which is same as the one used to encode the VP9 bitstreams to create the partition trees in our database. Since the *libvpx* implementation of VP9 is in C language, we used the Tensorflow C API, which was compiled with support for optimizations that use Intel’s Math Kernel Library for Deep Neural Networks. This enabled us to seamlessly embed our model within the VP9 encoder. The Tensorflow C API was also found to be considerably faster than both Keras and Tensorflow’s native

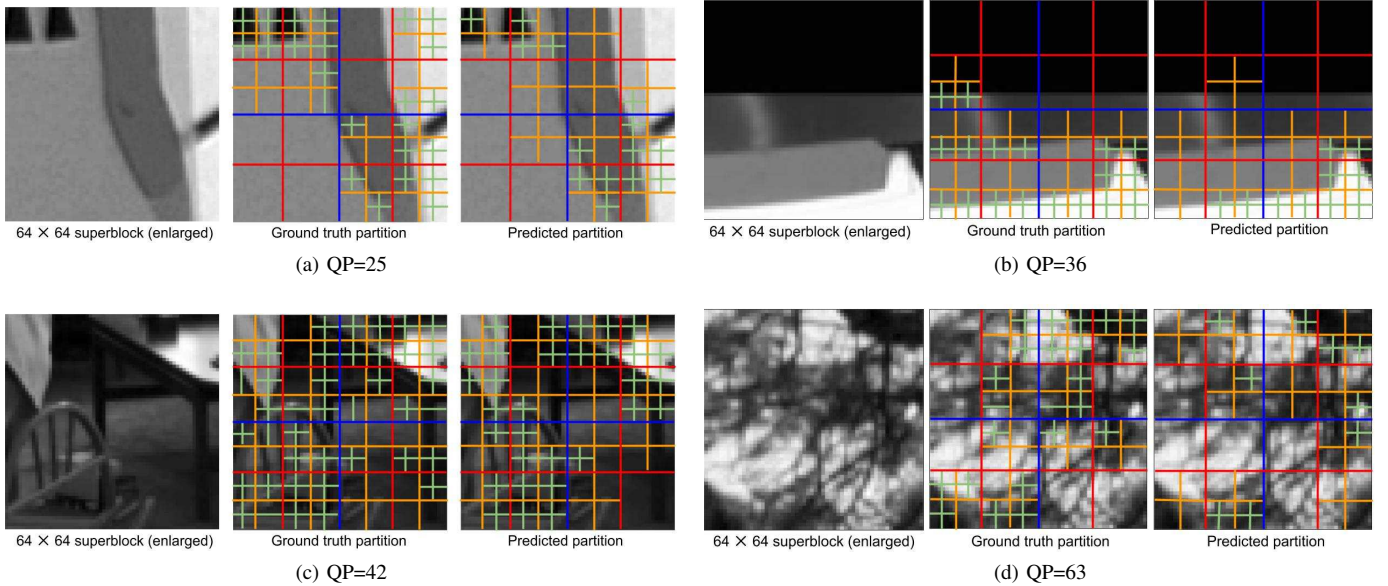


Fig. 7. Superblock partitions predicted by the trained H-FCN model compared with ground truth partitions on four superblocks encoded with different QP values.

TABLE II  
PREDICTION ACCURACY OF H-FCN MODEL

# of parameters	Training (%)				Validation (%)			
	Level 0	Level 1	Level 2	Level 3	Level 0	Level 1	Level 2	Level 3
336,608	91.96	86.63	85.70	90.54	91.43	85.70	84.60	90.80
70,408	91.89	86.46	85.33	90.23	91.34	85.53	84.28	90.81
26,336	91.73	86.07	84.42	89.42	91.18	85.13	83.47	90.27

Python API for inference, which is an added advantage of this choice. While a GPU was used for training, all encoding tests were performed without a GPU on a single core Intel Core i7-7500U CPU @2.70GHz with 8 GB of RAM running 64 bit Ubuntu 16.04.

### B. Training Details and Hyperparameters

The H-FCN model as depicted in Fig. 4 has 26,336 trainable parameters and executes  $\sim 10.8$ M floating point operations (FLOPs) per sample, at a per sample inference time of  $\sim 1$ ms, where each sample in the inference batch is a  $64 \times 64$  superblock. We also experimented with two larger versions of the model, having 336,608 and 70,408 trainable parameters, which were identical to that in Fig. 4, except that we used a larger number of filters in certain layers.<sup>1</sup> The number of training and validation samples are as mentioned in Table I. The H-FCN model was trained with a batch size of 128 using the Adam optimizer [43], with a step size of 0.001 for over  $10^7$  iterations. The weights of each layer were initialized with randomly drawn samples from uniform distributions determined by the size of the inputs to each layer [44]. Fig. 6 illustrates the variation of the training and validation losses of the H-FCN model against the number of training iterations, where the validation loss was evaluated on

$10^5$  samples from the validation set at the end of every 1000 iterations.

### C. Prediction Performance

We evaluated the prediction accuracy of the trained model on the validation set. The average accuracies of the model at the four levels of prediction were evaluated on  $10^5$  randomly drawn samples of the validation set, and are summarized in Table II. The corresponding performance on an equal number of random samples from the training set is also provided for reference.

In Table II, we also include the performance results of the two larger models with 336,608 and 70,408 parameters respectively, to show that our model design allows for smaller architectures with no significant impairment of prediction accuracy. From Table II, we observe that the maximum decline in accuracy on the validation set between the largest and the smallest models was 1.13%, which occurred at level 2. The impact of this decline in accuracy on the BD-rate is examined in Section V-D.

The partitions predicted by the H-FCN model are visualized in Fig. 7, against the corresponding ground truth partitions, on four randomly selected superblocks encoded at different QP values from the validation set, with the  $64 \times 64$  superblocks shown enlarged for visual clarity. These visualizations serve to demonstrate that the global structure and density of the

<sup>1</sup>The architectures of the larger models are provided at <https://drive.google.com/file/d/1nDP42TjoeZhoWAlfhGIHICMi04-CmKUE/view?usp=sharing>

TABLE III  
ENCODING PERFORMANCE OF H-FCN MODEL

Sequence	Video Source	Resolution	# of frames	$\Delta T$ (%)			BD-rate (%)		
				336,608	70,408	26,336	336,608	70,408	26,336
<i>Sintel trailer</i>	Blender	1920×1080	300	-0.1	27.8	<b>56.0</b>	<b>2.78</b>	2.78	3.39
<i>Pedestrian area</i>	TUM		375	36.6	41.6	<b>61.8</b>	<b>2.08</b>	2.12	2.45
<i>Sunflower</i>	TUM		500	18.8	34.4	<b>55.0</b>	<b>2.61</b>	2.76	3.07
<i>Crowd run</i>	VQEG		500	45.2	49.7	<b>65.3</b>	<b>0.99</b>	1.13	1.16
<i>Ducks takeoff</i>	VQEG		500	43.6	48.6	<b>67.9</b>	<b>1.42</b>	1.55	1.72
<i>Narrator</i>	Netflix El fuente		300	36.6	55.8	<b>73.3</b>	0.32	<b>0.25</b>	0.58
<i>Food market</i>	Netflix El fuente		300	37.0	46.4	<b>68.5</b>	1.51	<b>1.50</b>	1.63
<i>Toddler fountain</i>	Netflix Chimera		420	47.8	48.9	<b>70.2</b>	1.37	1.35	<b>1.33</b>
<i>Rainroses</i>	EBU		500	57.9	64.2	<b>73.6</b>	<b>0.83</b>	0.86	0.92
<i>Kidsoccer</i>	EBU		500	74.6	76.5	<b>83.5</b>	0.65	<b>0.59</b>	0.74
Average	-	-	-	39.8	49.4	<b>67.5</b>	<b>1.46</b>	1.49	1.70
<i>Big buck bunny</i>	Blender	1280×720	300	37.6	55.2	<b>69.3</b>	2.26	<b>2.22</b>	3.21
<i>Park run</i>	TUM		504	24.0	50.4	<b>65.2</b>	<b>1.29</b>	2.06	1.74
<i>Shields</i>	TUM		504	54.2	62.9	<b>69.3</b>	<b>1.61</b>	1.88	1.73
<i>Into tree</i>	VQEG		500	52.3	69.4	<b>77.0</b>	<b>1.10</b>	1.23	1.27
<i>Old town cross</i>	VQEG		500	59.4	67.8	<b>77.5</b>	<b>0.38</b>	0.71	0.82
<i>Crosswalk</i>	Netflix El fuente		300	42.4	63.3	<b>74.4</b>	<b>0.91</b>	1.07	1.32
<i>Tango</i>	Netflix El fuente		294	35.8	49.6	<b>75.4</b>	<b>1.52</b>	1.61	1.73
<i>Driving POV</i>	Netflix Chimera		500	35.6	55.9	<b>70.1</b>	1.25	<b>1.24</b>	1.46
<i>Dancers</i>	Netflix Chimera		500	25.5	49.9	<b>71.9</b>	3.08	<b>1.68</b>	2.47
<i>Vegicandle</i>	EBU		500	47.1	54.0	<b>71.6</b>	<b>1.47</b>	1.61	1.74
Average	-	-	-	41.4	57.8	<b>72.2</b>	<b>1.49</b>	1.53	1.75
<i>Elephant's dream</i>	Blender	960×540	545	37.6	54.7	<b>59.9</b>	<b>1.68</b>	1.72	1.82
<i>Euro Truck Simulator 2</i>	Twitch		300	48.0	68.0	<b>73.1</b>	1.65	<b>1.63</b>	1.69
<i>Station</i>	TUM		313	64.0	65.4	<b>80.7</b>	<b>2.01</b>	2.08	2.27
<i>Rush hour</i>	TUM		500	27.3	50.3	<b>66.4</b>	<b>1.71</b>	1.92	1.98
<i>Touchdown pass</i>	VQEG		570	45.2	63.7	<b>72.8</b>	<b>1.45</b>	1.49	1.62
<i>Snow mnt</i>	VQEG		570	35.1	50.8	<b>64.6</b>	1.09	1.05	<b>1.04</b>
<i>Roller coaster</i>	Netflix El fuente		500	46.1	53.6	<b>73.9</b>	<b>1.40</b>	1.46	1.68
<i>Dinner scene</i>	Netflix Chimera		500	23.9	54.0	<b>65.7</b>	1.48	<b>1.38</b>	1.70
<i>Boxing practice</i>	Netflix Chimera		254	43.1	58.4	<b>68.6</b>	1.43	<b>1.42</b>	1.50
<i>Meridian</i>	Netflix Meridian		500	42.4	57.9	<b>69.6</b>	<b>1.36</b>	1.46	1.47
Average	-	-	-	41.3	57.7	<b>69.5</b>	<b>1.53</b>	1.56	1.68
Overall average				40.8	55.0	<b>69.7</b>	<b>1.49</b>	1.53	1.71

superblock partitions predicted by our H-FCN model are in good agreement with ground truth, although some local blocks may be partitioned differently between the two.

#### D. Encoding Performance

Although the encoding time can be reduced by predicting the block partitions using a trained model instead of conducting RDO-based exhaustive search, the RD performance of learned models may suffer due to incorrect partition predictions. Thus, to evaluate the performance of our trained model, it is also necessary to assess its RD performance with respect to that of the reference RDO. Accordingly, we encoded several test video sequences with the original VP9 encoder using the RDO based partition search, and also with the modified VP9 encoder using the integrated H-FCN model to conduct partition prediction. The test video sequences were obtained as raw videos from multiple publicly available sources commonly used for evaluating video codecs.<sup>2</sup> All of the test sequences were converted to the YCbCr 4:2:0 8-bit format prior to encoding. The lengths of the sequences

varied between 254 to 570 frames. Sequences which were originally longer were clipped to less than 550 frames (the number of frames used for each test sequence are reported in Table III). As mentioned earlier in Section III-B, our model was jointly trained on three different spatial resolutions. For each resolution, the test sequences chosen were at the same or higher resolution. The sequences that were originally at higher resolutions were downsampled to the required resolution using Lanczos resampling.

A set of five internal QP values given by {15, 31, 47, 70, 99} (corresponding to external QP values of {20, 30, 35, 40, 45} respectively) was selected to represent a practical quality range of intra-frames used in adaptive streaming. Each sequence was then encoded at these five QP values in one pass, using constant quality intra-mode, one tile per frame, speed level 1 and the *good* quality setting. These settings were chosen to be compatible with the encoding configuration used for our database, as mentioned in Section III-B, although our model is equally amenable to be trained for other configurations and can be parallelized over multiple tiles. The RDO based VP9 encoder at these settings thus formed the baseline for our approach. We used the same encoding settings as the baseline

<sup>2</sup>Our test video sequences were sourced from <https://media.xiph.org/video/derf/> and [https://tech.ebu.ch/hdvtv/hdvtv\\_test-sequences](https://tech.ebu.ch/hdvtv/hdvtv_test-sequences).






Fig. 8. Variation of  $\Delta t$ ,  $\Delta$ bitrate and  $\Delta$ PSNR across the range of QP values at different resolutions.

for the modified encoder with the integrated H-FCN model.

The percentage speedup of our method with respect to the RDO baseline for a test sequence is calculated as:

$$\Delta T = \frac{T_{\text{RDO}} - T_{\text{H-FCN}}}{T_{\text{RDO}}} \times 100 \quad (2)$$

where  $T_{\text{RDO}}$  and  $T_{\text{H-FCN}}$  are the total times taken to encode the sequence at all five QP values with the RDO baseline encoder and the H-FCN integrated encoder, respectively. Since our model is integrated with the encoder,  $T_{\text{H-FCN}}$  includes the total inference time of the H-FCN model on the superblock batches from all of the frames of the video sequence to be encoded, as well as the encoding time without the RDO procedures for superblock partition decision. Thus, a positive value of  $\Delta T$  represents a speedup with respect to the RDO. To measure the RD performance, we also computed the BD-rate relative to the RDO baseline, using peak signal-to-noise ratio (PSNR) as the distortion metric. Table III reports the  $\Delta T$  and BD-rate values of the test sequences when using the H-FCN models with 336, 608, 70, 408 and 26, 336 parameters. The test videos in Table III comprise 30 distinct contents, equally divided between the three resolutions. In Table III, the first sequences of the  $1920 \times 1080$  and  $1280 \times 720$  categories as well as the first two sequences of the  $540 \times 960$  category are CGI content. As would be expected, the smallest model, with just 26, 336 parameters, achieved the highest speedup over all the test videos. Although the model with 336, 608 parameters achieved the highest accuracy on the validation set, we observe that for some test sequences, the smaller models achieved lower BD-rate values. The overall BD-rates achieved by the three models were 1.49%, 1.53% and 1.71%. The smallest model was thus found to maintain a reasonable BD-rate, while achieving the most speedup. All of the analysis that we provide was conducted using the smallest H-FCN model having 26, 336 parameters, unless otherwise mentioned.

In order to show that the H-FCN model is able to generalize irrespective of the QP values chosen for evaluation within the range of the database, we also report the performance measures for an additional set of four internal QP values  $\{10, 40, 60, 88\}$  in Table IV. The speedup and BD-rate values reported in Table IV are close to the corresponding values from Table III, indicating that the performance is consistent across the range of QP values used in the database. For each of the nine QP values in the combined set  $\{10, 15, 31, 40, 47, 60, 70, 88, 99\}$ , the average changes in encoding time, bitrate and PSNR with respect to the RDO baseline, denoted by  $\Delta t$ ,  $\Delta$ bitrate and  $\Delta$ PSNR respectively, are also plotted in Fig. 8 for the three

resolutions. The curves in Fig. 8, illustrate that a substantial speedup is obtained at every QP value, for a relatively small bitrate increase and PSNR reduction with respect to the RDO baseline. Although the encoding time was predominantly reduced with increases in the QP value using the RDO baseline as well as the proposed method, on the 540p sequences, the average encoding time of the RDO baseline was found to increase at the higher end of the range of QP values, while the proposed method was still faster at these QPs. Consequently, the average  $\Delta t$  values for the 540p sequences increased at the higher end of the QP range, unlike the trend observed for the other two resolutions. Further, the  $\Delta t$ ,  $\Delta$ bitrate and  $\Delta$ PSNR values obtained at the adjacent QP values tested are generally quite close. Subsequent results are hence reported using the initial set of five QP values.

TABLE IV  
ENCODING PERFORMANCE ON ADDITIONAL QP VALUES

Resolution	$\Delta T$ (%)	BD-rate (%)
1080p	69.0	1.80
720p	67.8	1.37
540p	71.9	1.41
Overall	70.7	1.61

TABLE V  
PERFORMANCE TRADE-OFF OF THE INCONSISTENCY CORRECTION APPROACH WITH RESPECT TO USING RDO TO RESOLVE INCONSISTENCIES

Resolution	$\Delta T$ (%)		BD-rate (%)	
	Top-down	RDO	Top-down	RDO
1080p	67.5	63.7	1.70	1.40
720p	72.2	62.3	1.75	1.13
540p	69.5	64.1	1.68	1.40
Overall	69.7	63.4	1.71	1.31

We also evaluated the performance trade-off achieved by the top-down inconsistency correction scheme described in Section IV-C, by exploring an alternative approach of resolving inconsistencies, whereby the RDO was invoked to decide the partitioning of a superblock if its predicted partition tree was determined to be inconsistent. Among all approaches that could be devised to handle inconsistencies, assigning superblocks having inconsistent partition predictions to the RDO in this manner would ensure the best performance in terms of BD-rate, albeit at the possible expense of a decline in the speedup achieved. This trade-off is summarized in Table V, in terms of average BD-rate and speedup of our

TABLE VI  
COMPARISON OF SPEEDUP VERSUS BD-RATE TRADEOFF OF OUR APPROACH WITH VP9 SPEED LEVEL 4

Sequence	Res.	$\Delta T$ (%)		BD-rate (%)		BD-PSNR (dB)	
		speed 4	H-FCN	speed 4	H-FCN	speed 4	H-FCN
<i>Sintel trailer</i>	1080p	46.3	<b>56.0</b>	4.46	<b>3.39</b>	-0.21	<b>-0.17</b>
<i>Pedestrian area</i>		<b>62.8</b>	61.8	2.94	<b>2.45</b>	-0.12	<b>-0.10</b>
<i>Sunflower</i>		50.2	<b>55.0</b>	6.74	<b>3.07</b>	-0.28	<b>-0.13</b>
<i>Crowd run</i>		62.8	<b>65.3</b>	2.30	<b>1.16</b>	-0.21	<b>-0.11</b>
<i>Ducks takeoff</i>		63.2	<b>67.9</b>	1.88	<b>1.72</b>	-0.18	<b>-0.17</b>
<i>Narrator</i>		59.5	<b>73.3</b>	3.72	<b>0.58</b>	-0.13	<b>-0.02</b>
<i>Food market</i>		55.9	<b>68.5</b>	2.39	<b>1.63</b>	-0.18	<b>-0.12</b>
<i>Toddler fountain</i>		61.7	<b>70.2</b>	1.73	<b>1.33</b>	-0.14	<b>-0.11</b>
<i>Rainroses</i>		72.9	<b>73.5</b>	2.50	<b>0.92</b>	-0.14	<b>-0.04</b>
<i>Kidssoccer</i>		<b>85.1</b>	83.5	0.86	<b>0.74</b>	-0.10	<b>-0.08</b>
Average		62.0	<b>67.5</b>	2.95	<b>1.70</b>	-0.17	<b>-0.10</b>
<i>Big buck bunny</i>		720p	58.4	<b>69.3</b>	7.92	<b>3.21</b>	-0.51
<i>Parkrun</i>	61.1		<b>65.1</b>	3.63	<b>1.74</b>	-0.46	<b>-0.22</b>
<i>Shields</i>	<b>70.4</b>		69.3	2.84	<b>1.73</b>	-0.24	<b>-0.14</b>
<i>Into tree</i>	75.6		<b>77.0</b>	1.57	<b>1.27</b>	-0.13	<b>-0.11</b>
<i>Old town cross</i>	75.5		<b>77.5</b>	2.00	<b>0.82</b>	-0.15	<b>-0.06</b>
<i>Crosswalk</i>	71.8		<b>74.4</b>	3.75	<b>1.32</b>	-0.15	<b>-0.05</b>
<i>Tango</i>	62.4		<b>75.3</b>	2.39	<b>1.73</b>	-0.11	<b>-0.08</b>
<i>Driving POV</i>	62.6		<b>70.1</b>	2.00	<b>1.46</b>	-0.16	<b>-0.12</b>
<i>Dancers</i>	70.4		<b>71.9</b>	12.26	<b>2.47</b>	-0.09	<b>-0.00</b>
<i>Vegicandle</i>	<b>74.1</b>		71.6	2.85	<b>1.74</b>	-0.14	<b>-0.08</b>
Average	68.2	<b>72.2</b>	4.12	<b>1.75</b>	-0.21	<b>-0.11</b>	
<i>Elephants dream</i>	540p	59.2	<b>59.9</b>	2.13	<b>1.82</b>	-0.18	<b>-0.16</b>
<i>Euro Truck Simulator 2</i>		72.0	<b>73.1</b>	1.89	<b>1.69</b>	-0.23	<b>-0.20</b>
<i>Station</i>		76.5	<b>80.7</b>	2.06	2.27	<b>-0.13</b>	-0.15
<i>Rush hour</i>		59.9	<b>66.4</b>	2.68	<b>1.98</b>	-0.11	<b>-0.08</b>
<i>Touchdown pass</i>		70.1	<b>72.8</b>	2.56	<b>1.62</b>	-0.15	<b>-0.09</b>
<i>Snow mnt</i>		<b>65.2</b>	64.6	2.31	<b>1.04</b>	-0.30	<b>-0.14</b>
<i>Roller coaster</i>		71.1	<b>73.9</b>	2.21	<b>1.68</b>	-0.15	<b>-0.11</b>
<i>Dinner Scene</i>		60.1	<b>65.7</b>	3.91	<b>1.70</b>	-0.09	<b>-0.03</b>
<i>Boxing practice</i>		65.8	<b>68.6</b>	2.03	<b>1.50</b>	-0.14	<b>-0.10</b>
<i>Meridian</i>		59.5	<b>69.6</b>	1.98	<b>1.47</b>	-0.11	<b>-0.08</b>
Average	65.9	<b>69.5</b>	2.38	<b>1.68</b>	-0.16	<b>-0.12</b>	
Overall Average		65.4	<b>69.7</b>	3.15	<b>1.71</b>	-0.18	<b>-0.11</b>

inconsistency correction approach against that obtained when the RDO is employed to resolve inconsistencies, for each of the three resolutions considered in our work. Thus, by using the RDO to handle inconsistencies, an overall improvement of 0.4% in BD-rate was achieved at the expense of a 6.3% reduction in speedup on the sequences tested. Naturally, the improvement in RD performance that can be achieved by any other correction scheme is upper bounded by 0.4%, and more sophisticated correction approaches such as those based on prediction probabilities indeed achieved negligible gains. Thus, due to its simplicity, we found that the top-down approach described in Section IV-C to be a good choice for inconsistency correction.

#### E. Comparison with VP9 Speed Levels

The VP9 encoder provides speed control settings which allow skipping certain RDO search options and using early termination for faster encoding at the expense of RD performance. Thus, it essentially works towards the same goal as our H-FCN based partition prediction design. VP9 has 9 speed levels designated by the numerals 0-8. To use a particular speed level, the *cpu-used* parameter is set to the corresponding number while encoding. The recommended speed levels for *best* and *good* quality settings, are 0-4, whereas levels 5-8 are reserved for use with the *realtime* quality setting. Speed level 0 is the slowest and is rarely used in practice, whereas speed

level 1 provides a good quality versus speed trade-off.<sup>3</sup> As already pointed out earlier in this section, speed level 1 with the *good* quality setting is the baseline for our method, and the fastest encoding for this configuration occurs by setting the speed level to 4.

The experimental results in Table III reveal that our method is 69.7% faster than the baseline, with an increase of 1.71% in BD-rate in intra-mode. However, any approach that claims to accelerate VP9 encoding is only useful in practice if it performs better than the available speed control setting. Thus, we compared the performance of our model with that of speed level 4, which is the fastest speed level of VP9 intended for this configuration (i.e. with the *good* quality setting). Table VI summarizes the speedups, BD-rates, and BD-PSNRs of our method and those of speed level 4, where all quantities were computed with respect to the baseline. The results in Table VI reveal that on the sequences tested, our H-FCN based approach achieved a 4.3% higher speedup on average than did the RDO at speed level 4, while incurring 1.44% smaller BD-rate penalty. Speed level 4 suffered an average BD-PSNR decrease of -0.18 dB, as compared to -0.11 dB by our method. Our learned model delivers greater speedup, while also yielding better RD performance than the

<sup>3</sup>See <http://wiki.webmproject.org/ffmpeg/vp9-encoding-guide> for recommended settings.

available speed control setting of the VP9 codec operating in intra-mode. The computational efficiency of the H-FCN model, combined with the effectiveness of the partitions that it predicts while maintaining good RD performance, implies that it better optimizes this important trade-off than does the VP9 speed control mechanism.

Our approach consistently surpasses the RDO at speed level 4 in terms of speedup achieved over the common baseline, across the practical range of QP values for adaptive streaming considered in our work. To emphasize this point further, Fig. 9 plots the net speedup achieved on the test videos against the QP values, for each of the three spatial resolutions considered. In every instance, the net speedup obtained by using the speed level 4 setting of the VP9 encoder fell below that achieved by our system, over all QP values, although at 540p resolution and QP value of 99, it was close to that of our method. This strongly supports the utility of our method for practical encoding scenarios that use a range of QP values, rather than a few specific ones.




Fig. 9. Plot of speedup achieved by H-FCN and RDO at speed 4 relative to baseline against QP value, for three spatial resolutions.

TABLE VII  
PERFORMANCE COMPARISON OF ETH-CNN [22] AGAINST OUR  
APPROACH ON OUR TEST SET

Sequence	Res.	$\Delta T$ (%)		BD-rate (%)	
		ETH-CNN	H-FCN	ETH-CNN	H-FCN
<i>Sintel trailer</i>	1080p	<b>83.1</b>	74.6	6.40	<b>5.00</b>
<i>Pedestrian area</i>		<b>76.6</b>	71.1	5.12	<b>4.60</b>
<i>Sunflower</i>		<b>79.4</b>	73.9	2.99	<b>2.65</b>
<i>Crowd run</i>		54.6	<b>56.1</b>	1.26	<b>1.11</b>
<i>Ducks takeoff</i>		58.7	<b>66.7</b>	2.90	<b>1.49</b>
<i>Narrator</i>		<b>81.0</b>	73.2	3.87	<b>3.52</b>
<i>Food market</i>		<b>60.0</b>	59.9	2.02	<b>1.97</b>
<i>Toddler fountain</i>		53.4	<b>57.6</b>	2.15	<b>1.61</b>
<i>Rainroses</i>		69.8	<b>70.1</b>	4.05	<b>3.23</b>
<i>Kidssoccer</i>		<b>43.7</b>	41.9	1.14	<b>1.08</b>
Average		<b>66.0</b>	64.5	3.19	<b>2.63</b>
<i>Big buck bunny</i>		720p	<b>76.8</b>	71.2	3.21
<i>Parkrun</i>	47.2		<b>47.5</b>	1.64	<b>1.43</b>
<i>Shields</i>	<b>51.7</b>		51.2	2.04	<b>1.67</b>
<i>Into tree</i>	<b>68.1</b>		65.9	1.89	<b>1.35</b>
<i>Old town cross</i>	<b>57.2</b>		53.9	1.57	<b>1.44</b>
<i>Crosswalk</i>	<b>75.8</b>		70.1	3.17	<b>2.85</b>
<i>Tango</i>	<b>69.7</b>		64.8	3.69	<b>3.52</b>
<i>Driving POV</i>	<b>53.8</b>		52.8	1.71	<b>1.58</b>
<i>Dancers</i>	<b>84.4</b>		74.2	4.96	<b>4.84</b>
<i>Vegicandle</i>	<b>65.6</b>		61.7	2.69	<b>2.39</b>
Average	<b>65.6</b>		61.7	2.69	<b>2.39</b>
Overall average	<b>65.8</b>		63.1	2.94	<b>2.51</b>

## F. HEVC Encoding Performance

In this section, we assess the performance of the H-FCN model in reducing the intra-mode encoding complexity of HEVC. The H-FCN model used for this purpose was identical to the model shown in Fig. 4 except that the first branch with 64 outputs was removed to account for the 3 level partition structure of CTUs into coding units (CUs) in the default configuration of the HM reference encoder [45] for HEVC in the intra mode. Also, size of the last convolutional layer of each of remaining three branches was changed from  $1 \times 1 \times 4$  to  $1 \times 1 \times 2$  since the split labels are binary for HEVC intra mode CUs. Thus, minimal changes were made to adapt the H-FCN model for HEVC partition prediction. The resulting model had 21,418 parameters, and was trained using the publicly available CPH-Intra database for HEVC introduced in [22]. The CPH-Intra database was created using four QP values  $\{22, 27, 32, 37\}$  and has 2,446,725 and 143,925 sample CTUs for training and validation respectively [22]. The hyperparameters used for training were identical to the ones mentioned in Section V-B for VP9. Training the modified H-FCN model in the above manner enabled us to compare the performance of our model with that of the ETH-CNN model of [22] trained on the same database.<sup>4</sup> for predicting the intra mode partition of CTUs into CUs. Although it might be possible to improve performance by also predicting the splits of the smallest  $8 \times 8$  CUs into  $4 \times 4$  prediction units (PUs), the CPH-Intra database does not include the corresponding split labels for  $8 \times 8$  CUs, and thus our model does not make this prediction either, which allows for a fair comparison with [22].

The accuracy values obtained by the H-FCN model on the CPH-Intra validation set were 86.90%, 86.87% and 91.39% at levels 3, 2 and 1 respectively while the corresponding accuracy values for the ETH-CNN model with 1,287,189 parameters were 90.98%, 86.42% and 80.42%. The average accuracy of the H-FCN model across all the three levels was 88.38% which was thus better than that of the much larger ETH-CNN model [22], which averaged 85.94% across all levels. The set of test video sequences from Table III were then encoded at QP values  $\{22, 27, 32, 37\}$  with the reference HEVC encoder from the HM 16.5 software [45] using the main profile in intra mode. Each test sequence was encoded thrice, using the RDO based CTU partition decisions of the HM encoder, the partitions predicted by the ETH-CNN model of [22] and our H-FCN model respectively. The 540p test sequences listed in Table III were excluded in this case, as the HM reference encoder requires the frame height to be a multiple of the minimum CU depth; since the minimum CU depth is 8 in default configuration, this condition is not satisfied for 540p sequences.

The comparison of the encoding performance of our approach with that of [22] provided in Table VII indicates that our method attains a slightly smaller speedup than ETH-CNN on average, but is able to consistently achieve better BD-rates on all the test sequences. In this case, we performed the inference purely in Python to allow for a fair comparison

<sup>4</sup>the implementation of ETH-CNN as well as pretrained models are provided at <https://github.com/tianyi2017/HEVC-Complexity-Reduction>

TABLE VIII  
PERFORMANCE COMPARISON OF ETH-CNN [22] AGAINST OUR APPROACH ON THE JCT-VC TEST SET

Class	Resolution	Sequence	$\Delta T$ (%)		BD-rate (%)	
			ETH-CNN	H-FCN	ETH-CNN	H-FCN
A	2560×1600	<i>PeopleOnStreet</i>	56.0	<b>59.3</b>	<b>2.37</b>	2.53
		<i>Traffic</i>	62.8	<b>64.4</b>	<b>2.55</b>	2.63
B	1920×1080	<i>BasketBallDrive</i>	63.7	<b>65.0</b>	4.26	<b>3.60</b>
		<i>BQTerrace</i>	<b>54.1</b>	54.0	1.84	<b>1.78</b>
		<i>Cactus</i>	59.9	<b>60.9</b>	<b>2.26</b>	2.32
		<i>Kimono</i>	<b>79.5</b>	74.9	2.59	<b>1.90</b>
		<i>Parkscene</i>	62.5	<b>64.6</b>	1.96	<b>1.80</b>
E	1280×720	<i>FourPeople</i>	60.6	<b>61.2</b>	<b>3.11</b>	3.13
		<i>Johnny</i>	<b>73.2</b>	67.6	3.82	<b>3.64</b>
		<i>KristenAndSara</i>	<b>70.9</b>	66.1	3.46	<b>3.26</b>
Average			<b>64.3</b>	63.8	2.82	<b>2.66</b>

with [22] which does not integrate the ETH-CNN model with the HM encoder, and performs an external inference step using Python. Integrating the trained model with the encoder using the optimized C API as we have done for VP9 should further improve the speedup achieved by our approach. The encodes used to compute the RD performance of the approach of [22] were obtained using four different trained ETH-CNN models, each separately trained for one QP value, whereas our method was able to obtain better RD performance with a single H-FCN model jointly trained for all four QP values. The performance tradeoff achieved by our approach seems to be particularly promising for the 1080p sequences where 0.56% lower BD-rate was achieved for 1.5% increase in encoding complexity by our approach. The higher speedup attained by the approach of [22] can be explained by the early termination scheme of the ETH-CNN model. On an average our approach reduced encoding time by 63.1% for 2.51% increase in BD-rate, while the the method of [22] resulted in 65.8% reduction in encoding time for a 2.94% increase in BD-rate. Using the same experimental setup that was just described, we also compared the encoding performance of our approach against that of [22] on the 8 bit sequences having resolutions of at least 720p from the JCT-VC test set [46] as reported in Table VIII. Such deep learning based complexity reduction schemes are more beneficial at higher resolutions, as indicated by the lower speedup values reported on resolutions smaller than 720p in [22]. Thereby, we only report the results on the highest three resolutions from the JCT-VC test set in Table VIII. On average, our approach achieved a lower BD-rate on this set of test sequences as well, while being marginally slower than the method of [22].

## VI. CONCLUSION

We have developed and explained a deep learning based partition prediction method for VP9 superblocks, that is implemented using a hierarchical fully convolutional network. We constructed a large database of VP9 superblocks and corresponding partitions on real streaming video content from the Netflix library, which we used to train the H-FCN model. The trained model was found to produce consistent partition trees yielding good prediction accuracy on VP9 superblocks. By integrating the trained H-FCN model into the VP9 encoder,

we were able to show that VP9 intra-mode encoding time can be reduced by 69.7% on average, at the cost of an increase of 1.71% in BD-rate, by using the partitions predicted by the model, instead of invoking an RDO-based partition search during encoding. Further, by selectively employing the RDO to decide the partitions of superblocks that are predicted with inconsistent partitions, the BD-rate was reduced to 1.31% with a corresponding speedup of 63.4%. The experiments we conducted comparing our model against the fastest recommended speed level of VP9 for the *good* quality setting, further corroborated its effectiveness relative to the faster speed settings available in the reference encoder. This strongly suggests that the framework developed here offers an attractive alternative approach to accelerate VP9 intra-prediction partition search. We also believe that our approach is applicable to other video codecs, such as HEVC and AV1, that employ hierarchical block partitioning at multiple levels.

Our approach can be improved to avoid the inconsistency correction step, by deploying a simple mechanism, such as early termination to enforce consistent partition predictions at all levels. Another immediate next step is to extend our approach to the prediction of inter-mode superblock partitions, a task which encounters the additional challenge of spatiotemporal inference. Further, computing the partition tree, as well as other decisions made by the RDO, may be perceptually suboptimal, since they are generally driven by the goal of optimizing the mean squared error (MSE) with respect to the reference video. Since the MSE is generally a poor indicator of the perceptual quality of images and videos [47], other more perceptually relevant criteria could be considered. By instead optimizing the partition decision process in terms of a suitable perceptual quality model like SSIM [48], MS-SSIM [49], VIF [50], ST-RRED [51], or VMAF [52], [53], RD performance could potentially be further improved, which is a direction that we intend to explore as part of our future work. Finally, it is also interesting to extend this approach to optimize the AV1 codec, which, due to its even deeper partition tree, and more computationally intensive RDO search process, could benefit even more from the speedup offered by our system model.

TABLE IX  
PERFORMANCE OF THE TOP-DOWN AND BOTTOM-UP MODEL DESIGNS.

Resolution	$\Delta T$ (%)			BD-rate (%)		
	Bottom-up (70,408)	Bottom-up (26,336)	Top-down (31,456)	Bottom-up (70,408)	Bottom-up (26,336)	Top-down (31,456)
1080p	49.4	<b>67.5</b>	47.4	<b>1.49</b>	1.70	1.60
720p	57.8	<b>72.2</b>	53.3	<b>1.53</b>	1.75	1.63
540p	57.7	<b>69.5</b>	40.2	<b>1.56</b>	1.68	1.63
Overall	55.0	<b>69.7</b>	46.8	<b>1.53</b>	1.71	1.62

#### ACKNOWLEDGMENT

We thank Anush K. Moorthy and Christos G. Bampis of Netflix for their help in creating the database.

#### APPENDIX A

##### COMPARISON WITH TOP-DOWN PARTITION PREDICTION

In order to substantiate the efficacy of our bottom-up hierarchical block merge prediction scheme, we designed a top-down version of the H-FCN model, to predict block splits instead of merges, starting with the split type of the  $64 \times 64$  superblocks. While performing top-down prediction with the hierarchical model, spatial isolation between features from the adjacent smallest blocks need to be maintained until the end of the trunk, which limits the number of max pooling stages that can be used in the trunk. We thus modified the H-FCN architecture of Fig. 4 such that some of the max pooling layers were removed from the trunk and instead added to the branches, so that the output of each branch is a matrix of the desired size. However, the number of convolutional filters used in a branch relative to the number of outputs predicted by that branch were the same as in Fig. 4. A diagram illustrating the architecture of the top-down H-FCN model is available at [54]. The top-down model has 31,456 parameters, which is comparable to the 26,336 parameters in the bottom-up model from Fig. 4 used in our work. The hyperparameters used when training the top-down model were the same as described in Section V-B.

The accuracy values of the trained top-down H-FCN model on  $10^5$  samples from the validation set were 89.35%, 82.80%, 85.12%, and 91.47%, from the level 3 to level 0 of the partition tree. As compared to the corresponding values of 90.27%, 83.47%, 85.13% and 91.18% obtained using the bottom-up model with 26,336 parameters, the accuracy improved at the lowest level, but declined at the other three levels. In fact, the reduction in accuracy was largest at the highest level of the partition tree. This is consistent with our observation that global semantics learned by the deeper convolutional layers aid the partition decisions on larger blocks.

The encoding performance of the top-down and bottom-up partition prediction approaches are compared in Table IX, where the numbers in parentheses denote the number of parameters of each model. The values in Table IX were calculated by encoding the set of test video sequences listed in Table III. While the top-down model has a better BD-rate performance as compared to the bottom-up model with 26,336 parameters, the speedup it can achieve is considerably less. However, a higher average speedup and lower BD-rate than obtained by the top-down model on average can be achieved

by using a bottom-up model with 70,408 parameters, as shown in the second and fifth columns of Table IX. If the number of parameters of the top-down model is reduced, its speedup will approach the value obtained by the bottom-up model with 70,408 parameters, but the BD-rate will increase further. This suggests that the hierarchical model is more efficient when predictions are made in a bottom-up fashion.

The key idea behind the bottom-up partition decision is to reduce the number of computations by introducing max pooling at each stage, which progressively reduces the spatial dimensions of the feature maps that propagate along the trunk of the network. However, as the top-down approach predicts the partitions of the smallest blocks at the last stage, the number of max pooling stages that can be introduced in the trunk is limited by the need to prevent the feature maps from the adjacent smallest blocks from combining. Hence, the prediction of the largest  $8 \times 8$  matrix representing the split types of  $8 \times 8$  blocks at the last branch requires a larger feature map to be carried through the trunk as compared to the bottom-up approach, which limits its speed. The results from Table IX support the intuition behind adopting a bottom-up merge prediction scheme. While other frameworks can be designed that are amenable to top-down split prediction, our hierarchical model is thus found to be more suitable for bottom-up merge prediction.

#### APPENDIX B

##### EVALUATION ON CPH-INTRA DATABASE

In order to evaluate the effect of the contents chosen to train the H-FCN model, we also created a VP9 partition database, from the images provided as part of the CPH-Intra database introduced in [22] for HEVC intra-mode partition prediction. The CPH-Intra database is comprised of 2000 images, with 1700 images provided for training and the rest for validation and testing. As in our original database described in Section III, the contents were encoded at three different resolutions (1080p, 720p and 540p). The preprocessing stage was also the same as described in Section III, where at each encode resolution, all the images of the same or higher resolution were used as sources after Lanczos resampling (for contents that were originally at a higher resolution), and conversion to YCbCr 4:2:0 8-bit representation. However, in this case the content was a set of images, each with a distinct visual content, unlike a set of videos characterized by substantial temporal correlation between frames, as in our original database. Thus, we encoded each image individually as an intra-frame using the *libvpx* reference encoder for VP9 to ensure maximal utilization of the available content. The encoding configuration,

and the process of extracting superblocs from the sources and partition trees from the encodes were exactly the same as in Section III-B. The final database, which we refer to as the CPH-VP9 database, consists of 841,500 training samples. The H-FCN model trained with this CPH-VP9 database was evaluated using the set of test sequences from Table III. At each resolution, we compared the average BD-rate obtained using the CPH-VP9 database and our larger database from Section III, as reported in Table X. The average speedup is essentially the same as in Table III since the same H-FCN model was trained on both databases.

TABLE X  
COMPARISON OF BD-RATE ON CPH-INTRA DATABASE AND OUR DATABASE

Resolution	CPH-VP9 (%)	Ours (%)
1080p	1.65	1.70
720p	1.58	1.75
540p	2.36	1.68
Overall	1.86	1.71

As compared to the H-FCN model trained on our database, the model trained on the CPH-VP9 database achieved a lower average BD-rate on the 1080p and 720p resolutions, while its average BD-rate at 540p was higher. The overall BD-rate performance of the model trained with our database was a little better. Nevertheless, the H-FCN model was able to generalize well on the smaller CPH-VP9 database.

## REFERENCES

- [1] D. Mukherjee *et al.*, “The latest open-source video codec VP9 - An overview and preliminary results,” in *Proc. IEEE Picture Coding Symp.*, 2013, pp. 390–393.
- [2] T. Wiegand, G. J. Sullivan, G. Bjontegaard, and A. Luthra, “Overview of the H. 264/AVC video coding standard,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 7, pp. 560–576, 2003.
- [3] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, “Overview of the high efficiency video coding (HEVC) standard,” *IEEE Trans. on Circuits Syst. Video Technol.*, vol. 22, no. 12, pp. 1649–1668, 2012.
- [4] J. Bankoski, P. Wilkins, and Y. Xu, “Technical overview of VP8, an open source video codec for the web,” in *Proc. IEEE Int. Conf. Multimedia Expo*, 2011, pp. 1–6.
- [5] Y. Chen *et al.*, “An overview of core coding tools in the AV1 video codec,” in *Proc. IEEE Picture Coding Symp.*, 2018, pp. 41–45.
- [6] B. Bross, J. Chen, and S. Liu, “Versatile video coding (draft 5),” JVT Document JVET-N1001, Geneva, CH, Mar. 2019.
- [7] G. Toderici *et al.*, “Full resolution image compression with recurrent neural networks,” in *Proc. IEEE Conf. Comput. Vision Pattern Recogn.*, 2017, pp. 5306–5314.
- [8] E. Agustsson, M. Tschanen, F. Mentzer, R. Timofte, and L. Van Gool, “Generative adversarial networks for extreme learned image compression,” *arXiv preprint arXiv:1804.02958*, 2018.
- [9] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, “Variational image compression with a scale hyperprior,” *arXiv preprint arXiv:1802.01436*, 2018.
- [10] T. Chen, H. Liu, Q. Shen, T. Yue, X. Cao, and Z. Ma, “Deepcoder: a deep neural network based video compression,” in *Proc. IEEE Visual Commun. Image Process.*, 2017, pp. 1–4.
- [11] S. Kim *et al.*, “Adversarial video compression guided by soft edge detection,” *arXiv preprint arXiv:1811.10673*, 2018.
- [12] C.-Y. Wu, N. Singhal, and P. Krahenbuhl, “Video compression through image interpolation,” in *Proc. Euro. Conf. Comput. Vision*, 2018, pp. 416–431.
- [13] Z. Zhao, S. Wang, S. Wang, X. Zhang, S. Ma, and J. Yang, “Enhanced bi-prediction with convolutional neural network for high efficiency video coding,” *IEEE Trans. Circ. Syst. Video Technol.*, Oct. 2018 (Early Access).
- [14] Y. Wang, X. Fan, C. Jia, D. Zhao, and W. Gao, “Neural network based inter prediction for HEVC,” in *Proc. IEEE Int. Conf. Multimedia Expo*, 2018, pp. 1–6.
- [15] J. Liu, S. Xia, W. Yang, M. Li, and D. Liu, “One-for-all: Grouped variation network-based fractional interpolation in video coding,” *IEEE Trans. Image Process.*, vol. 28, no. 5, pp. 2140–2151, May 2019.
- [16] Z. Zhao, S. Wang, S. Wang, X. Zhang, S. Ma, and J. Yang, “CNN-based bi-directional motion compensation for high efficiency video coding,” in *Proc. IEEE Int. Symp. Circuits Syst.*, 2018, pp. 1–4.
- [17] C. Jia *et al.*, “Content-aware convolutional neural network for in-loop filtering in high efficiency video coding,” *IEEE Trans. Image Process.*, vol. 28, no. 7, pp. 3343 – 3356, Jul. 2019.
- [18] Y. Zhang, T. Shen, X. Ji, Y. Zhang, R. Xiong, and Q. Dai, “Residual highway convolutional neural networks for in-loop filtering in HEVC,” *IEEE Trans. Image Process.*, vol. 27, no. 8, pp. 3827–3841, Mar. 2018.
- [19] J. Kang, S. Kim, and K. M. Lee, “Multi-modal/multi-scale convolutional neural network based in-loop filter design for next generation video codec,” in *Proc. IEEE Int. Conf. Image Process.*, 2017, pp. 26–30.
- [20] Y. Li, B. Li, D. Liu, and Z. Chen, “A convolutional neural network-based approach to rate control in HEVC intra coding,” in *Proc. IEEE Vis. Commun. Image Process.*, 2017, pp. 1–4.
- [21] Z. Liu, X. Yu, Y. Gao, S. Chen, X. Ji, and D. Wang, “CU partition mode decision for HEVC hardwired intra encoder using convolution neural network,” *IEEE Trans. Image Process.*, vol. 25, no. 11, pp. 5088–5103, Nov. 2016.
- [22] M. Xu, T. Li, Z. Wang, X. Deng, R. Yang, and Z. Guan, “Reducing complexity of HEVC: A deep learning approach,” *IEEE Trans. Image Process.*, vol. 27, no. 10, pp. 5044–5059, Oct. 2018.
- [23] G. Bjontegaard, “Calculation of average PSNR differences between RD-curves,” *Proc. of the ITU-T Video Coding Experts Group (VCEG) Thirteenth Meeting*, Apr. 2001.
- [24] H-FCN Based VP9 Intra Encoding. Accessed May 27, 2019. [Online]. Available: <https://github.com/Somdyuti2/H-FCN>
- [25] D. Ruiz-Coll, V. Adzic, G. Fernandez-Escribano, H. Kalva, J. L. Martinez, and P. Cuenca, “Fast partitioning algorithm for HEVC intra frame coding using machine learning,” in *Proc. IEEE Int. Conf. Image Process.*, 2014, pp. 4112–4116.
- [26] F. Duanmu, Z. Ma, and Y. Wang, “Fast CU partition decision using machine learning for screen content compression,” in *Proc. IEEE Int. Conf. Image Process.*, 2015, pp. 4972–4976.
- [27] T. Li, M. Xu, and X. Deng, “A deep convolutional neural network approach for complexity reduction on intra-mode HEVC,” in *Proc. IEEE Int. Conf. Multimedia Expo*, 2017, pp. 1255–1260.
- [28] Y. Xian, Y. Wang, Y. Tian, Y. Xu, and J. Bankoski, “Multi-level machine learning-based early termination in VP9 partition search,” *Electron. Imaging*, vol. 2018, no. 2, pp. 1–5, 2018.
- [29] Z. Yan *et al.*, “HD-CNN: hierarchical deep convolutional neural networks for large scale visual recognition,” in *Proc. IEEE Int. Conference Comput. Vision*, 2015, pp. 2740–2748.
- [30] X. Zhu and M. Bain, “B-CNN: Branch convolutional neural network for hierarchical classification,” *arXiv preprint arXiv:1709.09890*, 2017.
- [31] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proc. IEEE Conf. Comput. Vision Pattern Recogn.*, 2015, pp. 3431–3440.
- [32] V. Badrinarayanan, A. Kendall, and R. Cipolla, “SegNet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE Tran. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, Dec. 2017.
- [33] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, “Deeper depth prediction with fully convolutional residual networks,” in *Proc. 4th Int. Conf. 3D Vision*, 2016, pp. 239–248.
- [34] L. Wang, L. Wang, H. Lu, P. Zhang, and X. Ruan, “Saliency detection with recurrent fully convolutional networks,” in *Proc. Euro. Conf. Comput. Vision*, 2016, pp. 825–841.
- [35] J. Dai, Y. Li, K. He, and J. Sun, “R-FCN: Object detection via region-based fully convolutional networks,” in *Proc. Adv. Neural Info. Process. Syst.*, 2016, pp. 379–387.
- [36] L. Wang, W. Ouyang, X. Wang, and H. Lu, “Visual tracking with fully convolutional networks,” in *Proc. IEEE Int. Conf. Comput. Vision*, 2015, pp. 3119–3127.
- [37] J. Johnson, A. Karpathy, and L. Fei-Fei, “DenseCap: Fully convolutional localization networks for dense captioning,” in *Proc. IEEE Conf. Comput. Vision Pattern Recogn.*, 2016, pp. 4565–4574.
- [38] libvpx: VP8/VP9 Codec SDK (version 1.6.0). Accessed May 27, 2019. [Online]. Available: <https://chromium.googlesource.com/webm/libvpx/>

- [39] A. Aaron, Z. Li, M. Manohara, J. De Cock, and D. Ronca. Per-title encode optimization. Accessed: May 27, 2019. [Online]. Available: <https://medium.com/netflix-techblog/per-title-encode-optimization-7e99442b62a2>
- [40] M. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European Conf. Comput. Vision*, 2014, pp. 818–833.
- [41] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 315–323.
- [42] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [43] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proc. IEEE Int. Conf. Comput. Vision*, 2015, pp. 1026–1034.
- [45] JCT-VC. HM Software. Accessed Nov 26, 2019. [Online]. Available: [https://hevc.hhi.fraunhofer.de/svn/svn\\_HEVCSoftware/tags/HM-16.5/](https://hevc.hhi.fraunhofer.de/svn/svn_HEVCSoftware/tags/HM-16.5/)
- [46] F. Bossen. (2013) Common test conditions and software reference configurations. Accessed: May 18, 2020. [Online]. Available: [https://www.itu.int/wftp3/av-arch/jctvc-site/2010\\_07\\_B\\_Geneva/JCTVC-B300.doc](https://www.itu.int/wftp3/av-arch/jctvc-site/2010_07_B_Geneva/JCTVC-B300.doc)
- [47] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? a new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [48] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [49] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *Proc. 37th Asilomar Conf. Signals, Syst. Comput.*, vol. 2, 2003, pp. 1398–1402.
- [50] H. R. Sheikh and A. C. Bovik, "Image information and visual quality," *IEEE Trans. Image Process.*, vol. 15, no. 2, pp. 430–444, Feb. 2006.
- [51] R. Soundararajan and A. C. Bovik, "Video quality assessment by reduced reference spatio-temporal entropic differencing," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 23, no. 4, pp. 684–694, 2012.
- [52] Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy, and M. Manohara. Toward A Practical Perceptual Video Quality Metric. Accessed: May 27, 2019. [Online]. Available: <https://medium.com/netflix-techblog/toward-a-practical-perceptual-video-quality-metric-653f208b9652>
- [53] Z. Li *et al.* VMAF: The Journey Continues. Accessed: May 27, 2019. [Online]. Available: <https://medium.com/netflix-techblog/vmaf-the-journey-continues-44b51ee9ed12>
- [54] Top-down H-FCN architecture. Accessed Nov 26, 2019. [Online]. Available: <https://drive.google.com/file/d/1k4PhhKaUXLTmmiTo38BsnCBz09VgYIT7/view?usp=sharing>




**Somdyuti Paul** received the B.Tech. degree in electronics and communication engineering from West Bengal University of Technology in 2012 and the M.Tech. degree in electrical engineering from Indian Institute of Technology Kanpur in 2015. She is currently pursuing the Ph.D. degree in electrical and computer engineering at The University of Texas at Austin, where she is with the Laboratory for Image and Video Engineering. Her research interests include image and video processing, computer vision, and machine learning.



**Andrey Norkin** received the M.Sc. degree in computer engineering from Ural State Technical University, Yekaterinburg, Russia, in 2001 and the Doctor of Science degree in signal processing from Tampere University of Technology, Tampere, Finland, in 2007.

From 2008 to 2015, he was with Ericsson, Sweden, conducting research on video compression and 3D video. In 2014, he worked on video encoding techniques for broadcast products at Ericsson TV, Southampton, UK. Since 2015, Dr. Norkin has been with Netflix, USA as a Senior Research Scientist working on encoding techniques for OTT video streaming, High Dynamic Range (HDR) video, and new video compression algorithms.

He has participated in ITU-T and MPEG efforts on developing video compression standards, including HEVC, its extensions, and VVC, having multiple technical contributions and coordinating work in certain areas, such as deblocking and loop filters. He has also been involved in the Alliance for Open Media work and actively contributed to the development of AV1 video codec. Dr. Norkin's current research interests include video compression, OTT streaming, HDR video, and machine learning techniques.



**Alan C. Bovik** (F'95) is the Cockrell Family Regents Endowed Chair Professor at The University of Texas at Austin. His research interests include image processing, digital photography, digital television, digital streaming video, and visual perception. For his work in these areas he has been the recipient of the 2019 Progress Medal from The Royal Photographic Society, the 2019 IEEE Fourier Award, the 2017 Edwin H. Land Medal from The Optical Society, a 2015 Primetime Emmy Award for Outstanding Achievement in Engineering Development

from the Television Academy, and the Norbert Wiener Society Award and the Karl Friedrich Gauss Education Award from the IEEE Signal Processing Society. He has also received about 10 best journal paper awards, including the 2016 IEEE Signal Processing Society Sustained Impact Award. His books include *The Essential Guides to Image and Video Processing*. He co-founded and was longest-serving Editor-in-Chief of the *IEEE Transactions on Image Processing*, and also created/chaired the *IEEE International Conference on Image Processing* which was first held in Austin, Texas, 1994.