

Improved CNN-based Learning of Interpolation Filters for Low-Complexity Inter Prediction in Video Coding

Luka Murn, Saverio Blasi, Alan F. Smeaton, *Fellow, IEEE*, and Marta Mrak, *Senior Member, IEEE*

The versatility of recent machine learning approaches makes them ideal for improvement of next generation video compression solutions. Unfortunately, these approaches typically bring significant increases in computational complexity and are difficult to interpret into explainable models, affecting their potential for implementation within practical video coding applications. This paper introduces a novel explainable neural network-based inter-prediction scheme, to improve the interpolation of reference samples needed for fractional precision motion compensation. The approach requires a single neural network to be trained from which a full quarter-pixel interpolation filter set is derived, as the network is easily interpretable due to its linear structure. A novel training framework enables each network branch to resemble a specific fractional shift. This practical solution makes it very efficient to use alongside conventional video coding schemes. When implemented in the context of the state-of-the-art Versatile Video Coding (VVC) test model, 0.77%, 1.27% and 2.25% BD-rate savings can be achieved on average for lower resolution sequences under the random access, low-delay B and low-delay P configurations, respectively, while the complexity of the learned interpolation schemes is significantly reduced compared to the interpolation with full CNNs.

Index Terms—video compression, motion compensation, interpolation, machine learning, complexity reduction

I. INTRODUCTION

TO meet the increasing demands for video content at better qualities and higher resolutions, new video compression technology is being researched and developed. The efficiency and versatility of recent Machine Learning (ML) based approaches paved the way for researching more extensive ways to integrate ML solutions into next generation video coding schemes. In this context, the recent MPEG and ITU Versatile Video Coding (VVC) standard [1] already includes some ML-based tools, and it is capable of achieving considerably higher compression efficiency than older video compression technologies.

Deep Convolutional Neural Network (CNN) approaches have been particularly successful in solving tasks such as single image super-resolution [2], image classification, object detection [3], and colourisation [4], some fundamental and challenging computer vision problems. However, such approaches are typically very complex, while video compression applications require a careful design of the coding tools, in order to meet the strict complexity requirements posed by practical encoder and decoder implementations. The inference process required by conventional Deep Neural Network (DNN) approaches is inherently complex, limiting their applicability to video coding. As such, learning-based tools within VVC, for instance the Low-Frequency Non-Separable Transform (LFNST) and Matrix Intra-Prediction (MIP) [5], were carefully designed and simplified throughout the course of the standardisation process. Some other, efficient approaches based on CNNs were proposed to solve video

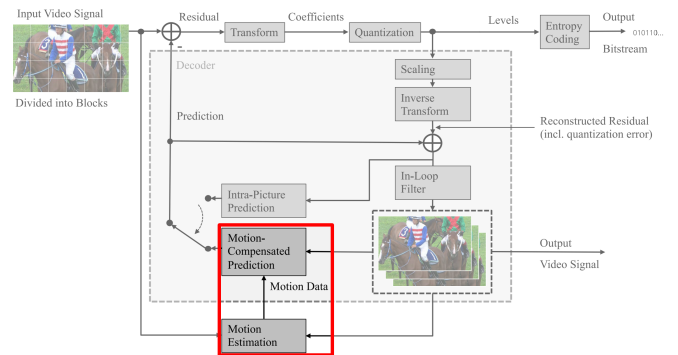


Fig. 1. Block diagram of a modern video encoder (e.g. VVC [7]), with highlighted inter-prediction modules

compression tasks such as chroma intra-prediction [6], with very promising results.

Additionally, methods based on learning may not generalise to the vast amount of different video content available, especially if the training process is not carefully designed. As such, the algorithms learned by these schemes need to be made transparent and explained, to mitigate potential unexpected outcomes.

Most video compression technology relies on a block-based hybrid coding scheme, as illustrated in Fig 1. Each frame is split into a number of blocks, using a flexible partitioning process. A block can then be predicted from already reconstructed pixels in either the current frame (intra-prediction) or by means of Motion Compensation (MC) on other reference frames (inter-prediction), using different prediction schemes. The decoder extracts information from the bitstream necessary to perform the correct prediction. Conversely, at the encoder, the different options are usually tested following a Rate-Distortion (RD) strategy. An RD cost is computed, typically a combination of a distortion metric and an estimate of the bitrate required to compress when using such an option. The option with the minimum RD cost is then selected.

Manuscript submitted May 18, 2021. The work described in this paper has been conducted within the project JOLT funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 765140.

L. Murn, S. Blasi and M. Mrak are with BBC Research & Development, The Lighthouse, White City Place, 201 Wood Lane, London, UK (e-mail: luka.murn@bbc.co.uk, saverio.blasi@bbc.co.uk, marta.mrak@bbc.co.uk).

A. F. Smeaton is with Dublin City University, Glasnevin, Dublin 9, Ireland (e-mail: alan.smeaton@dcu.ie)

Among the various video coding modules, inter-prediction is an essential tool that exploits temporal redundancies among different frames to achieve compression. A Motion Vector (MV) is derived at the encoder side by means of a Motion Estimation (ME) process, and transmitted to the decoder, to identify the spatial displacement between the position of the current block being encoded and the portion of a reference frame that is used for the MC prediction. In order to increase the accuracy of the MC process, sub-pixel interpolation has been widely used in many generations of video coding standards. Interpolation filters are applied to generate sub-pixel locations, which are then used alongside real integer-located pixels to perform MC. Sub-pixel ME simulates frame-to-frame displacements and sub-pixel shifts. Modern standards, such as VVC, may include new, refined interpolation filter sets. For instance, VVC includes new filters for Affine motion and for Reference Picture Resampling, along with an alternate half-pixel filter [7]. These developments indicate that the traditional fixed interpolation filters, based on the Discrete Cosine Transform (DCT), may not describe the displacements well enough or capture the diversity within the video data.

Given the successful application of ML to super-resolution applications, neural networks may be employed within a video coding scheme to produce new, more efficient interpolation filters. In a previous work [8], an approach for CNN-based sub-pixel interpolation filters for fractional motion estimation was presented in the context of the VVC standard. This is the first approach that shows gains in the challenging VVC framework, while also reducing the computational requirements to an acceptable level. This paper expands on such work [8], uncovering new ways to boost the prediction performance of the learned interpolation, without affecting the computational requirements of the prediction process. In particular this paper introduces:

- A multi-layer and multi-branch linear convolutional model architecture with shared layers for unified training of the interpolation filter set. The new architecture leverages the commonalities of the interpolation filters in order to train the full set with a single network;
- A three-stage training framework enabling competition of network's branches to increase learned prediction performance. Each network branch is trained to resemble a specific quarter-pixel fractional shift, whose weights can then be extracted as interpolation filter coefficients; and
- Detailed experimental analysis of training options and evaluation of the proposed solution in the VVC framework. The evaluation showcases how even a simple network without any non-linear functions and biases can achieve significant coding gains for the fractional interpolation task within video coding.

The proposed methodology relies on the powerful learning capabilities of deep CNNs to derive a new quarter-pixel interpolation filter set, even when its non-linear properties are removed to enable a low-complexity inference process required for its implementation within video codecs. The experimental results provide justifications for the proposed designs, and demonstrate that significant compression gains

can be achieved for challenging content used for development of video coding standards.

The paper is organised as follows: Section II provides a brief background of the related work, Section III describes the proposed methodology and new network architecture, Section IV shows experimental results, with conclusion drawn in Section V.

II. BACKGROUND

Video coding standards typically utilise a set of fixed sub-pixel interpolation filters. In practical solutions, most commonly used interpolation filters are designed for half and quarter pixel positions [9]. To boost prediction performance, more advanced video coding solutions such as VVC may employ higher sub-pixel precision in certain circumstances, e.g. Affine motion [7].

In addition to these approaches where filters were designed based on well known signals, new research has recently emerged where ML is used to derive or refine the sub-pixel samples. In these approaches, the target framework for improvement has been the MPEG/ITU-T High Efficiency Video Coding (HEVC/H.265) standard [10]. Several methods based on CNNs have been proposed that either improve the input to the interpolation process, by fusing the reference blocks in bi-prediction in a non-linear way [11], [12], or the output of the process, utilising the neighbouring reconstructed region to refine the prediction [13].

In the context of directly deriving the sub-pixel samples, an approach based on super-resolution CNNs to interpolate half-pixel samples was introduced in [14], bringing 0.9% Bjøntegaard delta-rate (BD-rate) [15] bit-rate reductions under low-delay P (LDP) configuration when replacing HEVC luma filters. Furthermore, training separate networks for luma and chroma channels was presented in [16]. The resulting models were tested within the HEVC reference software as a switchable interpolation filter, achieving 2.9% BD-rate coding gains under the LDP configuration. As a follow-up to [14], a sub-pixel MC was formulated as an inter-picture regression problem rather than an interpolation problem [17]. The resulting method uses 15 networks, one for each half and quarter-pixel fractional shift. The input to each network was the decoded reference block for that position, where the ground truth was the original content of the current block. Different NNs were trained for uni-prediction and bi-prediction and for different Quantisation Parameter (QP) ranges, resulting in a total of 120 NN-based interpolation filters. Two NN structures were used to train the interpolation filters, where a simpler structure was based on the Super-Resolution CNN (SRCNN) [2], and a more complex model was based on the Variable-filter-size Residue learning CNN (VRCNN) [18]. When tested on 32 frames under LDP configuration, 2.9% BD-rate gains were reported for VRCNN and 2.2% for SRCNN with respect to HEVC. The VRCNN implementation had increased the decoding time approximately 220 times, while the SRCNN increased it nearly 50 times. In subsequent work, the authors of [14], [17] have proposed a training scheme driven by invertibility, a marked property of fractional interpolation filters, reaching 4.7% BD-rate reductions for the LDB configuration with a VRCNN

architecture, and increasing the decoding time 700 times on average [19]. Lastly, an approach that groups all sub-pixel positions at one level and is trained for different QPs was proposed in [20], achieving 2.2% coding gains under the LDP configuration.

While these approaches provide powerful prediction performance thanks to NN-generated sub-pixel samples, they have very high computational requirements, increasing the encoding complexity by up to 380 times [8], which makes them inapplicable in real world scenarios. Furthermore, these NN-based interpolation filters are still used as “black-boxes”, i.e. without being able to fully understand how each predicted sample is predicted within the network. As such, many of the produced filters may actually behave very similar to the already existing, fixed interpolation filters. Some of these produced filters may not be useful, unnecessarily over-complicating the video coding scheme. Without complete knowledge on the actual filtering process, it is difficult to compare these filters with the conventional filters, possibly introducing redundancy within the video encoder. Additionally, most of these approaches require GPU-based implementations, which limit their applicability in more conventional CPU-based codec environments.

Recent research in the context of NN-based prediction showed that linear network segments can be deployed in place of more complex schemes with advantages in terms of reducing the complexity of the inference process [21]. Application of linear segments has already been presented in the context of intra-prediction for video coding [22]. Furthermore, in previous work [8], a simplified CNN-based filtering process was proposed in the context of inter-prediction in VVC, as opposed to the approach in [14], [16], [20] which used HEVC as its baseline encoder. Due to the fact VVC is considerably more efficient than HEVC, it is a much more challenging baseline from which to achieve compression efficiency gains.

In ScratchCNN [8], biases and activation functions were removed from the network architecture, and padding between network layers is not applied in order to reduce the input size from one layer to the next. This resulted in a new, simpler model, with one network trained for each quarter-pixel position, as visualised in Fig. 2. ScratchCNN had the ability of collapsing its three-layer structure so only a single matrix multiplication would need to be used at inference time, meaning that filter coefficients could be directly extracted from the learned network weights. These new, alternative learned filter coefficients could then be analysed easily to understand what exactly had been learned, rather than blindly implementing a full network with several layers. Moreover, the full, three-layer network with its large set of weights is no longer needed. The prediction process only requires the new, much smaller number of coefficients. Such an approach significantly decreased the storage and complexity requirements of its implementation, and also allowed tests to run on conventional, CPU-based computing architectures.

However, the existing approach still requires a significant number of networks to be trained, one for each fractional pixel location. Additionally, the training process could be improved to better utilise a neural network’s learning capabilities. This paper proposes to solve these issues. The ScratchCNN model

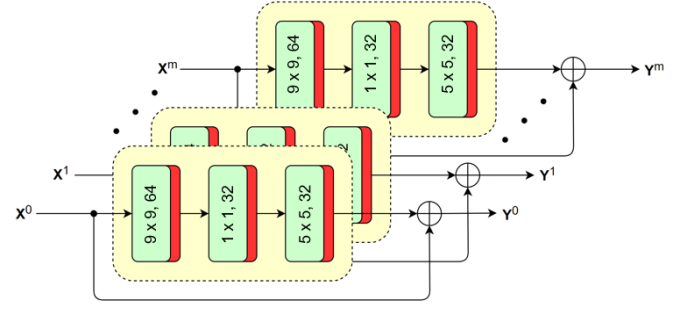


Fig. 2. Quarter-pixel interpolation filtering with *ScratchCNN*. Each convolutional kernel size is indicated as (kernel height \times kernel width, number of kernels)

[8] in Fig. 2 was used as basis for this approach, as it is the only model with acceptable complexity.

III. METHODOLOGY

This section introduces the general network architecture of the proposed approach and the novel methodology employed to produce the learned filters, as described in detail in Section III-A. The framework enables the training of filters with shared layers, referred to as SharedCNN (Section III-B), and in competition with VVC filters, referred to as CompetitionCNN (Section III-C). Additionally, Section III-D explains the multi-stage training scheme which ensures that the learned interpolation filter set of CompetitionCNN is unlike the already present VVC filter set. In Section III-E, details are provided on how the learned predictors can be simplified into interpolation filters similar to those in [8]. The produced filters can be easily integrated into the VVC framework as switchable options. This means that for a given sub-pixel MV, the encoder will test both the conventional filters as well as the proposed, trained filters, and select the best option, which is signalled to the decoder in the bitstream.

A. General network architecture

The learned filters can be used to derive up to quarter-pixel precision predictions. Each pixel therefore requires 15 fractional samples to be produced, to account for all quarter and half pixel locations between a given integer sample, and its neighbouring integer samples on the bottom and on the right. For a given reference block, or input \mathbf{X}^m , the ScratchCNN network [8] as in Fig. 2 predicts residuals \mathbf{R}^m , where $m = 0, \dots, 14$ corresponds to each of 15 half and quarter pixel positions. Each residual is then added to the input reference samples to form the prediction \mathbf{Y}^m representing values at a quarter-pixel position as:

$$\mathbf{Y}^m = \mathbf{R}^m + \mathbf{X}^m. \quad (1)$$

To obtain \mathbf{Y}^m , the central pixel from the 13×13 window of \mathbf{X}^m is used in Eq. 1. The proposed model is a three-layer convolutional network, which contains 64 9×9 convolutional kernels in the first layer, followed by 32 1×1 kernels in the second layer, while the last layer consists of 32 5×5

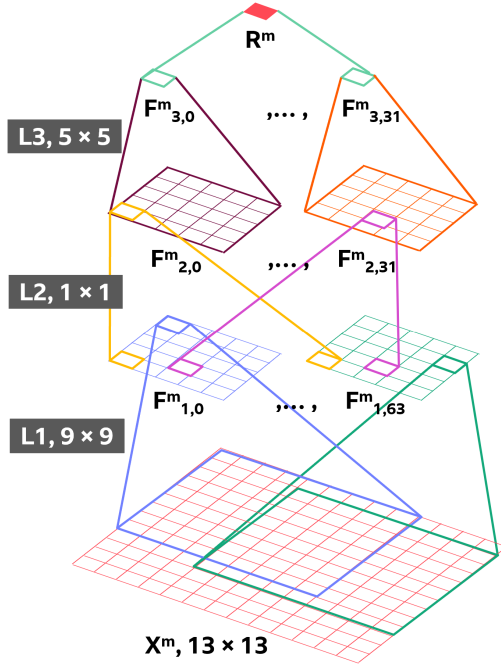


Fig. 3. Fractional pixel derivation process for branch m of an NN interpolation filter set. The proposed approach requires 13×13 input samples.

convolutional kernels. This structure is based on the well-known SRCNN model [2]. However, it is modified by removing activation functions and biases. Additionally, padding between layers is not applied in order to use only the available input samples. Thus, the model can be expressed as a series of linear operations from the input to the output. The process of acquiring a single pixel value from the input area is illustrated in Fig. 3. Due to the convolutional kernel dimensions (9×9 , 1×1 , 5×5), an input area of 13×13 is required to compute an interpolated output sample.

The approach, as described in [8], requires separate networks for each of the interpolated pixel samples. As illustrated in Fig. 3, the first layer, L1, extracts features from regions of input area \mathbf{X}^m by applying 9×9 convolutional kernels, denoted as $\mathbf{K}_{1,i}^m$, to obtain the first layer output, \mathbf{F}_1^m . It is defined as:

$$\mathbf{F}_{1,i}^m = \mathbf{K}_{1,i}^m * \mathbf{X}^m, \quad (2)$$

where $i = 0, \dots, 63$. The second layer, L2, combines features extracted from the same regions to retrieve its output, \mathbf{F}_2^m , as:

$$\mathbf{F}_{2,j}^m = \sum_{i=0}^{63} K_{2,j}^m \mathbf{F}_{1,i}^m, \quad (3)$$

where $K_{2,j}^m$ corresponds to 1×1 convolutional kernels and $j = 0, \dots, 31$. The network then combines all extracted features to reconstruct the pixel \mathbf{R}^m from \mathbf{F}_3^m feature maps comprising of 5×5 convolutional kernels $\mathbf{K}_{3,j}^m$ as:

$$\mathbf{F}_{3,j}^m = \mathbf{K}_{3,j}^m * \mathbf{F}_{2,j}^m, \quad (4)$$

and their summation for each $j = 0, \dots, 31$ as:

$$\mathbf{R}^m = \sum_{j=0}^{31} \mathbf{F}_{3,j}^m. \quad (5)$$

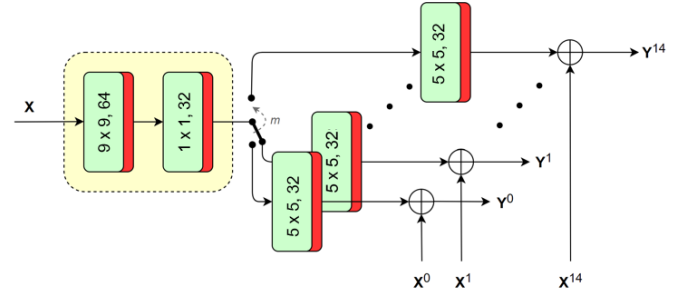


Fig. 4. SharedCNN architecture with a trunk with shared layers

Lastly, the required, non-separable 2D interpolation filter set \mathbf{M}^m can be obtained from the series of convolutional kernel weights $\mathbf{K}_{1,2,3}^m$ which were multiplied with values in the input \mathbf{X}^m . The interpolated sample can be described as:

$$\mathbf{R}^m = \mathbf{M}^m * \mathbf{X}^m = \mathcal{F}(\mathbf{K}_3^m, K_2^m, \mathbf{K}_1^m, \mathbf{X}^m). \quad (6)$$

The coefficients of each filter \mathbf{M}^m represent the contribution of reference samples towards the final prediction in a 13×13 surrounding area. The filter set corresponds to the 15 possible quarter-pixel fractional positions, each trained with 15 ($m = 0, \dots, 14$) separate sets of samples \mathbf{X}^m already classified before training.

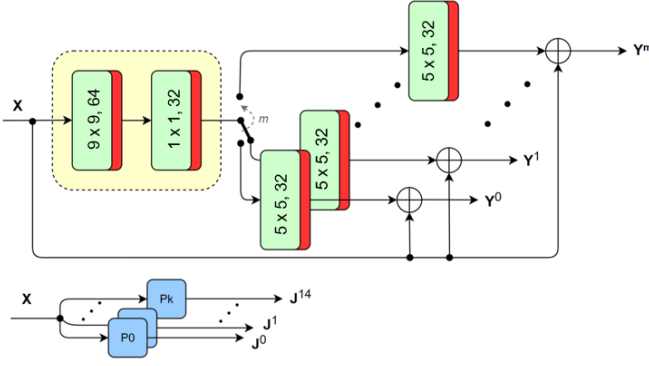
B. SharedCNN

The ScratchCNN architecture requires completely independent networks to be trained for each of the 15 sub-pixel positions. This is an inadequate solution. The networks are not flexible, as they do not share any information amongst themselves, while at the same time trying to approximate a very similar problem. The training process is computationally expensive, requiring to retrain each network whenever there is a need to expand or modify the training set. For that reason, a new architecture called SharedCNN is proposed where a so-called trunk, comprising of the first two layers in the network, is shared among all branches which still correspond to individual sub-pixel positions. This is illustrated in Fig. 4 where the branches are trained for only one sub-pixel position, and the trunk is trained for all positions. Thanks to such an architecture, the networks require considerably fewer coefficients to be trained and updated, which can improve learning performance.

Both the ScratchCNN linear model, as well as the new shared model architecture SharedCNN in Fig. 4, require labelled samples for training so that specific ScratchCNN or specific branch in the network with shared trunk can be updated for each fractional shift m at a time. The equations describing the SharedCNN network are equivalent to Eq. 1 - 6, with specifically defined convolutional layers 1 and 2 which are the same for all sub-pixel positions:

$$\mathbf{K}_1 = \mathbf{K}_1^m, K_2 = K_2^m, \forall m = 0, \dots, 14. \quad (7)$$

From the inference perspective, shared layers are used with each individual branch to produce a non-separable 2D filter for each fractional location, as in ScratchCNN.

Fig. 5. *CompetitionCNN* network architecture

C. *CompetitionCNN*

All frameworks discussed in this section so far rely on training with labelled data which forces interpolation filters to be specialised for specific pre-defined sub-pixel positions. However, such labelling may not be optimal and may limit adaptability of learned filters. Furthermore, learned interpolation filters are intended to be integrated into a video codec which can select between the standard filters and learned filters on a block basis. This means that during training it may be better to prevent back-propagation for some blocks which are better predicted with standard filters. In order to mitigate these problems, a new framework can be defined, on top of the architecture with shared layers in Fig. 4, by evaluating the loss of each prediction from both all neural network filters, as well as standard filters (P_0 to P_{14}). The network is named *CompetitionCNN* and is detailed in Fig. 5. The model contains two shared layers, with the last layer consisting of $b = 0, \dots, 14$ branches which, during training, produce predictions for every input block (i.e. not only for pre-defined sub-pixel positions). For the inference stage, these branches will be used to create 15 different filters \mathbf{M}^m . The process of acquiring an interpolated sample is equivalent to the one in *SharedCNN* and it corresponds to:

$$\mathbf{R}^m = \mathbf{M}^m * \mathbf{X} = \mathcal{F}(\mathbf{K}_3^m, K_2, \mathbf{K}_1, \mathbf{X}). \quad (8)$$

The input-output pairs for training the proposed approaches are formed by VVC. The ground truth is the uncompressed rectangular block in the current frame, while the input to the network is the decoded reference block that was chosen for sub-pixel interpolation. An important aspect of the generated pairs that needs to be preserved is the sub-pixel shift m from the reference input block to the output. The resulting dataset D contains subsets d^m and their relationship is described as:

$$D = \bigcup d^m. \quad (9)$$

Due to its architecture constraint, each *ScratchCNN* model as well as *SharedCNN* had to be trained with a specific subset d^m , consisting of pairs $(\mathbf{X}^m, \mathbf{Y}^m)$. However, the *CompetitionCNN* relaxes this fragmentation, as explained in the next subsection, using the whole dataset D for all branches.

D. *Three-stage training for CompetitionCNN*

In general, the network architecture of *SharedCNN* and *CompetitionCNN* are the same. However, their main difference is showcased in their respective training frameworks. To ensure that each branch of the final layer in the proposed *CompetitionCNN* competes with others when trained on dataset D , a three-stage training framework is proposed. The framework is inspired by a similar approach when training unique intra-prediction modes based on neural networks for its inclusion in the VVC standard [5].

The loss function used for training all neural networks in this paper is the Sum of Absolute Differences (SAD) loss. SAD is also used in *CompetitionCNN* to decide which branch will be used for back-propagation. For each branch b of the *CompetitionCNN* network, the loss l^b is calculated as:

$$l^b = \text{SAD}(\mathbf{Y}^b, \mathbf{GT}) = E[|\mathbf{Y}^b - \mathbf{GT}|], \quad (10)$$

where \mathbf{GT} denotes the ground truth, i.e. the original uncompressed block which needs to be predicted. The three stages of the proposed training are performed according to the following description.

Stage 1: To set the weights of the shared trunk and to also avoid the issue of branch starvation, where only a few branches in the final layer of the network are back-propagated while other \mathbf{K}_3^m kernels are mostly ignored due to their initial weight setting, the training loss L_1 of the first training epoch jointly updates all layers of the network with the loss of the whole dataset D in all branches, as follows:

$$L_1 = \sum_b l_D^b. \quad (11)$$

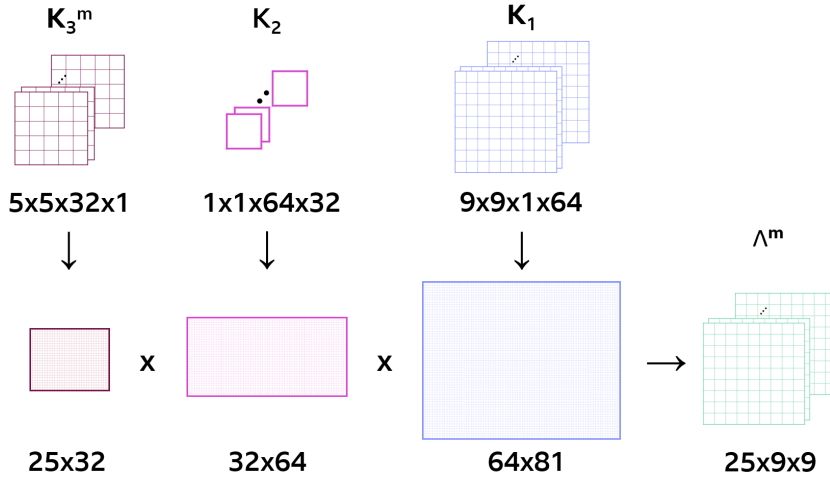
Stage 2: The critical requirement of the proposed architecture is the emergence of 15 unique interpolation filters, that mostly correspond to particular quarter-pixel fractional shifts. During the second epoch, each branch b of the last layer of the neural network is pre-trained with a separate subset d^m of the dataset, such that $m = b$, equivalent to each quarter-pixel position. The total loss L_2 is defined as:

$$L_2 = \sum_{b,m=b} l_{d^m}^b, \quad (12)$$

where only $l_{d^m}^b$, with $m = b$, is used to update each branch. By applying this step, a higher correlation between a branch output and a distinct quarter-pixel filter is established to further prevent starvation. Earlier experiments without this step resulted in several filters being a combination of similar filters.

Stage 3: Lastly, after the two pre-training steps are finished, training is commenced with a loss L_3 that updates the two shared layers and only the branch of the network which minimizes the loss for particular blocks, if the prediction it produces has lower SAD than any of the standard predictions \mathbf{J}^n , $n = 0, \dots, 14$, obtained by corresponding standard filters which are denoted as P_n , Fig. 5. Index μ of the network's branch that produces minimal loss for given block from D is found as:

CNN Weights Combination



Interpolation Filter Calculation

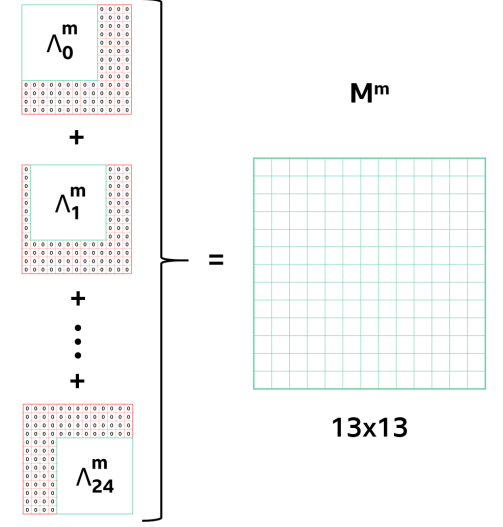


Fig. 6. A visualisation of the interpolation filter extraction process. Weights from all layers are combined into Λ^m . Each 9×9 coefficient matrix is placed within an 13×13 zero array before being summed up to obtain M^m .

$$\mu = \underset{b}{\operatorname{argmin}} l_D^b. \quad (13)$$

If a network filter fares better than all J^n , i.e. if

$$l^\mu < \min_n SAD(J^n, GT), \quad (14)$$

then back-propagation in μ^{th} branch is commenced for a given block.

This stage ensures that the network competes with the traditional VVC filters during the network training process already, leading to a closer representation of its final implementation form – a switchable interpolation filter set.

E. Learned interpolation filter coefficients

After training a network with shared layers, whether it is SharedCNN or CompetitionCNN, each 2D interpolation filter of size 13×13 can be constructed from the network's trunk combined with each branch using Eq. 8, by following the principle visualised in Fig. 3.

In Fig. 6, this process is explained in more detail. Since the architecture of a network with shared layers does not contain any non-linear functions or biases, a master kernel matrix Λ^m can be formulated as a combination of network weights K_3^m , K_2 , K_1 for each branch $m = 0, \dots, 14$ of the network. As illustrated in Fig. 3, the Λ^m matrix is applied onto 9×9 patches of an input X^m . By placing each of the $25 \times 9 \times 9$ matrices into the correct position within a 13×13 array and summing them all up, the interpolation filter M^m is created from a trained CNN. Note that the networks described in this paper learn residuals, so a value of 1 needs to be added to the central position of the learned interpolation filter M^m before using it in a video codec environment.

The filter coefficients corresponding to each sub-pixel shift, for CompetitionCNN, are visualized in Fig. 7. It can be

observed how the learned coefficients are correlated to the VVC filter coefficients. Only a portion of these coefficients actually contribute to the final pixel prediction, implying that a handful of them can be discarded with a negligible decline in filter efficiency. For example, a list of significant coefficients can be extracted for each filter M^m . Additionally, extracted filter coefficients can also be quantised to integer values to resemble filter sets already present in the VVC standard, which are whole numbers \mathbb{Z} in the range $[-2^6, 2^6]$. In this paper, filter coefficients are integrated within VVC in floating point precision, in the $[-1, 1]$ range. Any simplification approaches would need to be highly optimised to further reduce coding complexity without a significant decline in coding performance. This is left for future work as such optimisations would have to be harmonised with the design of other ML tools in a video codec.

The learned interpolation filter set is implemented into VVC's test model as switchable interpolation filters. For every potential fractionally interpolated block, both the existing filters and new, learned filters are considered before deciding the preferred choice. The selection is handled at a Coding Unit (CU) level, indicated as a flag in the compressed bitstream. In merge mode, blocks to be merged inherit the flag and consequently the choice of the interpolation filter.

IV. EXPERIMENTAL RESULTS

In this section, a description of the dataset, training and testing configuration is presented, followed by experiments which validate a linear network approach for sub-pixel interpolation filtering¹. Next, prior and proposed linearised network architectures are compared, before a final test is performed on Common Test Condition (CTC) video sequences, as defined

¹An open-source software package with the full implementation of this paper is available at <https://github.com/bbc/cnn-fractional-motion-compensation>

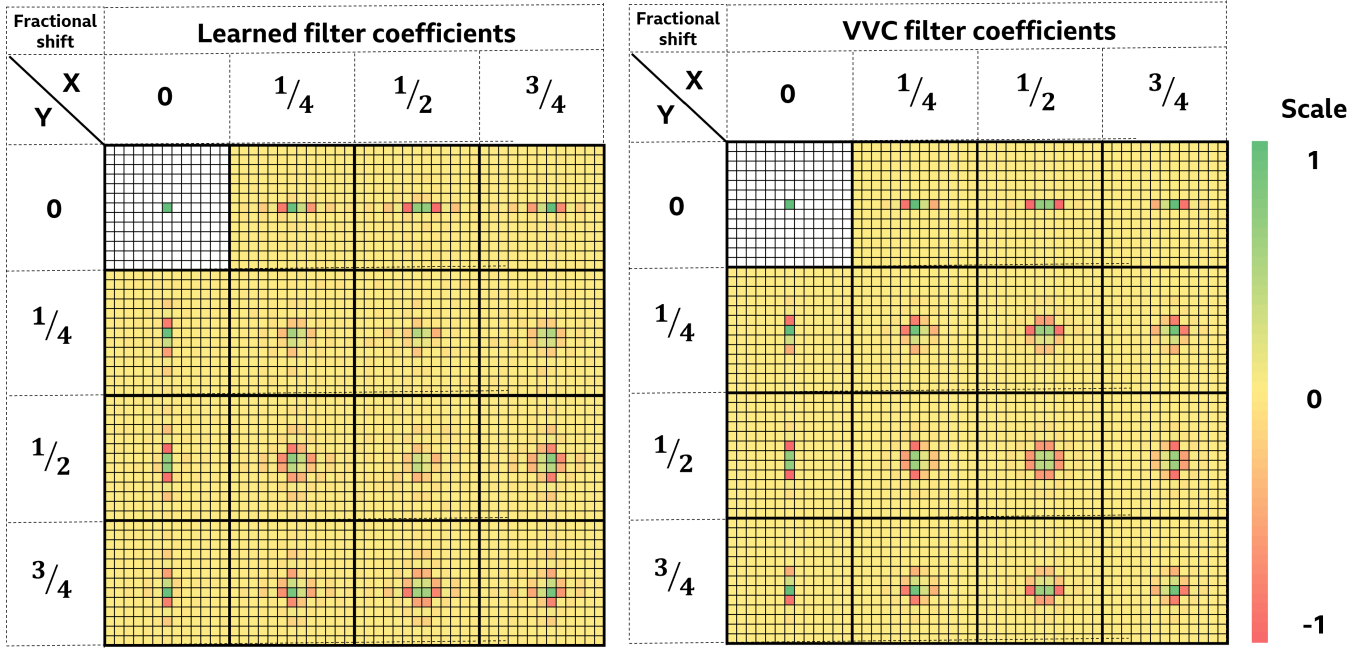


Fig. 7. A comparison of 15 learned 13×13 interpolation filters, one for each quarter-pixel position, and ones already present in VVC. For the $(0, 0)$ position, only the central pixel is considered

by JVET [23], with several encoder configurations. CTC video sequence classes are defined as sequences of the same spatial resolution.

A. Dataset generation

The training dataset for the proposed approach of fractional interpolation using CNNs is created by encoding the video sequence *BlowingBubbles* within VVC Test Model (VTM) 6.0 [24], under the Random Access (RA) configuration.

In order to generate accurate training data, a set of restrictions has been applied to the CTC conditions. The flags are: Triangle=0, Affine=0, DMVR=0, BIO=0, WeightedPredP=0, WeightedPredB=0. Affine, DMVR and Triangle modes of fractional interpolation have been disabled to generate a larger quantity of data from a selected sequence. Weighted prediction is turned off as it may contribute to a less accurate selection of sub-pixel positions m when blocks with lower weights are considered. To avoid training different networks for different Quantisation Parameters (QP), it is set to 27. Previous experiments have shown how filters generated from such a dataset generalize well when tested on sequences coded with various QPs. There is no restriction on block sizes, the network is trained on all rectangular block examples. As NN padding is avoided by the introduction of larger receptive fields, each convolutional operation reduces the block size of an example. Thus, the input reference block needs to be larger than the output prediction, comprising a 13×13 surrounding area around the required predicted pixels. If the input block is on the boundary of the frame, repetitive padding is applied to add the required pixels, a typical approach used in many MPEG standards [25].

After extracting the luma component data from VVC, the dataset is balanced before network training in terms of having

equal number of examples for each extracted fractional shift and block size. As the *BlowingBubbles* sequence is of 416×240 spatial resolution, the dataset contains block sizes up to 32×32 , as larger block partitioning is very scarce for the selected video input.

B. Training and testing configuration

The neural network was trained for 1000 epochs with a batch size of 32, using an Adam optimizer and a learning rate of 0.0001. Early stopping is enabled if the loss does not reduce for more than 50 epochs, i.e. when it converges. To avoid possible large gaps in loss optimization between the pre-training and training phases, gradient clipping by norm is implemented within the network, with the norm value set to 5.0.

The switchable filter implementation was tested in a modified VTM 6.0 encoding environment. The inter-prediction encoding conditions were simplified to resemble those of HEVC, with a number of restrictions imposed to coding tools and algorithms. The same encoding configuration flags were used as for generating the dataset, along with these additional restrictions: MHItra=0, SBT=0, MMVD=0, SMVD=0, IMV=0, SubPuMvp=0, TMVPMODE=0.

C. Using linear, multi-layer networks

To confirm that adopting a linear network approach without biases for creating fractional interpolation filters benefits the prediction performance, a comparison of different network architectures was undertaken. A test was performed on Class D CTC video sequences, comparing SAD losses on fractionally interpolated blocks with ScratchCNN, SharedCNN and CompetitionCNN after training. All input blocks were interpolated

TABLE I
LOSS COMPARISON OF LINEAR AND NON-LINEAR NETWORK STRUCTURES, ON CLASS D SEQUENCES

Model	Convolutional kernels	Number of networks	Activation functions	Biases	SAD
ScratchCNN	L1: 9x9,64 L2: 1x1,32 L3: 5x5,32	15	Yes	Yes	21.74
			Yes	No	21.75
			No	Yes	21.72
			No	No	21.68
	L1: 13x13,1		No	No	21.75

by the described networks and then compared to the ground truth block, as defined in Eq. 10. This loss was then compared with the loss of a block interpolated using VVC filters, with the minimum value of those two losses used as the final error. Using the minimum SAD loss between the network filters and VVC filters mimics the behavior of a switchable filter implementation in an actual video codec.

From Table I, it can be observed that the difference in performance between all three-layer ScratchCNN network architectures is negligible. Furthermore, adding biases and/or activation functions after the first and second layer of the network does not contribute to its performance. Consequently, removing these parameters creates a linear network structure that enables 2D interpolation filters to be extracted from the trained architecture, explained in Eq. 6, which leads to a much more streamlined and low-complexity implementation within a video codec.

Additionally, another test was performed to demonstrate how training a multi-layer linear network leads to a closer representation of the underlying solution than training a simple one-layer network. After training, a multi-layer network collapses into a single layer, behaving the same as a one-layer network, but the learned coefficients are different. In Table I, it can be observed how the network trained with only one 13×13 convolutional kernel has higher SAD loss than other ScratchCNN structures.

To demonstrate how this loss difference is reflected within VVC, interpolation filters extracted from both the three-layer and one-layer ScratchCNN were integrated as switchable filters in the codec. In Table II, the two linear structures are compared, with the multi-layer architecture displaying notably better performance across the RA, low-delay B (LDB) and low-delay P (LDP) encoding configurations. In the case of RA, the coding gains for the luma component, as expressed in the BD-rate, more than double, going from -0.10% to -0.26% . The BD-rate metric expresses the objective video quality when encoding content at different bit rates. Negative values indicate better compression performance for the same objective quality, compared to the baseline.

D. Comparison between three main architectures

The ScratchCNN architecture, as defined in [8], required different networks to be trained for each fractional shift. With the introduction of SharedCNN and CompetitionCNN, a single network can be trained to approximate all possible fractional shifts. This improvement allows the NN itself to

TABLE II
BD-RATE LUMA COMPARISON OF ONE-LAYER AND MULTI-LAYER NETWORK STRUCTURES, ON 32 FRAMES FOR CLASS D SEQUENCES

Model	Convolutional kernels	Encoder configuration		
		RA	LDB	LDP
ScratchCNN	L1: 9x9,64 L2: 1x1,32 L3: 5x5,32	-0.26%	-0.58%	-1.11%
	L1: 13x13,1	-0.10%	-0.37%	-1.00%
SharedCNN	L1: 9x9,64 L2: 1x1,32	-0.34%	-0.53%	-1.21%
CompetitionCNN	L3: 5x5,32	-0.62%	-1.06%	-2.21%

share knowledge through the first two layers to decide how best to learn and define specific interpolation filters. With additional improvements in the training framework, these new architectures can lead to lower loss values when compared to ScratchCNN, and higher compression efficiency when integrated within the codec.

As presented in Table II, when tested as switchable filter sets within VVC on 32 frames, both SharedCNN and CompetitionCNN exhibit higher coding gains compared to the previous ScratchCNN approach because of the inclusion of a two-layer trunk shared across all network branches. Another significant improvement introduced with the CompetitionCNN network is the competition of the branches amongst themselves when trained on a dataset, rather than being trained on subsets. As CompetitionCNN has the best coding performance among all approaches, with SharedCNN considered as an intermediate step, the learned filters extracted from it are subsequently tested on full non-CTC and CTC sequences, as presented in the next subsection.

E. Compression performance evaluation

A comparison between ScratchCNN and CompetitionCNN on Class C and D sequences for RA, LDB and LDP configuration is summarised in Table III.

On average, coding gains have improved across all configurations and classes, from -2.04% to -2.58% for Class D, LDP, for example. Significant coding gains are featured for the *BQSquare* sequence, a -5.54% improvement in compression efficiency compared to the baseline. The three-stage training framework has stabilised the performance of CompetitionCNN and allowed it to generalise well across varied video content, as CompetitionCNN displays coding

TABLE III
CODING PERFORMANCE OF THE PROPOSED APPROACH FOR RA, LDB AND LDP CONFIGURATIONS;
BD-RATE FOR LUMA; TESTED ON CTC AND NON-CTC SEQUENCES.

Sequence Set / Class	ScratchCNN [8]			CompetitionCNN (proposed)		
	RA	LDB	LDP	RA	LDB	LDP
<i>BasketballDrill</i> (C)	-0.15%	0.11%	-0.28%	-0.33%	-0.79%	-1.38%
<i>BQMall</i> (C)	-0.32%	-0.69%	-1.25%	-0.66%	-0.87%	-1.79%
<i>PartyScene</i> (C)	-0.82%	-1.92%	-3.22%	-1.32%	-2.18%	-3.78%
<i>RaceHorses</i> (C)	0.14%	0.19%	0.19%	-0.62%	-0.43%	-0.73%
CTC ClassC Overall	-0.29%	-0.58%	-1.14%	-0.73%	-1.07%	-1.92%
<i>BasketballPass</i> (D)	-0.14%	-0.33%	-0.52%	-0.43%	-0.61%	-1.32%
<i>BQSquare</i> (D)	-1.35%	-3.02%	-4.54%	-1.42%	-3.06%	-5.54%
<i>BlowingBubbles</i> (D)	-0.90%	-2.18%	-3.14%	-1.16%	-2.02%	-2.96%
<i>RaceHorses</i> (D)	0.04%	0.21%	0.02%	-0.17%	-0.23%	-0.49%
CTC ClassD Overall	-0.59%	-1.33%	-2.04%	-0.80%	-1.48%	-2.58%
CTC ClassC&D Overall	-0.44%	-0.95%	-1.59%	-0.77%	-1.27%	-2.25%
<i>BangkokMarketVidevo</i> (C)	0.08%	-0.12%	-0.49%	-0.31%	-0.58%	-1.18%
<i>BuntingHangingAcrossHongKongVidevo</i> (C)	0.05%	-0.49%	-1.02%	-0.11%	-0.38%	-1.53%
<i>BusyHongKongStreetVidevo</i> (C)	0.08%	0.07%	-0.10%	-0.04%	-0.31%	-0.79%
<i>CalmingWaterBVITexture</i> (C)	-0.07%	0.00%	-0.07%	-0.34%	-0.36%	-0.81%
non-CTC ClassC Overall	0.03%	-0.14%	-0.42%	-0.20%	-0.41%	-1.08%
<i>SkyscraperBangkokVidevo</i> (D)	-0.23%	-0.89%	-1.00%	-0.70%	-1.53%	-2.93%
<i>VegetableMarketS2LIVENetFlix</i> (D)	-0.24%	-0.59%	-0.73%	-0.54%	-1.24%	-2.28%
<i>VeniceSceneIRIS</i> (D)	-0.49%	-1.32%	-1.61%	-0.87%	-1.80%	-2.64%
<i>WalkingDownKhaoStreetVidevo</i> (D)	0.09%	-0.49%	-1.11%	-0.24%	-0.82%	-1.54%
non-CTC ClassD Overall	-0.22%	-0.82%	-1.11%	-0.59%	-1.35%	-2.35%
non-CTC ClassC&D Overall	-0.09%	-0.48%	-0.77%	-0.39%	-0.88%	-1.71%

TABLE IV
COMPETITIONCNN CODING COMPLEXITY ON CTC AND NON-CTC SEQUENCES
FOR RA, LDB AND LDP CONFIGURATIONS, ENTIRE SEQUENCE.

Class	RA		LDB		LDP	
	EncTime	DecTime	EncTime	DecTime	EncTime	DecTime
CTC ClassC	767%	290%	982%	231%	691%	221%
CTC ClassD	690%	264%	845%	220%	599%	204%
non-CTC ClassC	730%	220%	855%	180%	660%	190%
non-CTC ClassD	878%	281%	1171%	228%	783%	223%

gains for every sequence in every configuration. On the contrary, ScratchCNN performs inconsistently for some sequences like *RaceHorses*, across both C and D classes. Additionally, CompetitionCNN demonstrates only minor coding efficiency differences for different classes in RA, unlike ScratchCNN filters, whose performance notably diminishes when tested on higher resolution sequences.

As stated in Section IV-A, the networks were trained on a Class D CTC sequence, *BlowingBubbles*. In order to verify that both networks are not overfitting, tests on additional C and D sequences were performed, taken from the Bristol database for deep video compression [26]. The database contains a total of 800 video sequences, created from 200 4K sequences, which were carefully selected from public video databases, such as Videvo Free Stock Video Footage [27], IRIS32 Free

4K Footage [28], BVI-Texture [29], and LIVE-Netflix [30] among others. The 4K sequences were spatially downsampled to resolutions of 1920×1080 , 960×540 (Class C) and 480×270 (Class D) and truncated to 64 frames without scene cuts.

For the experiments presented in this paper, 8 video sequences were selected according to their low-level video features, spatial information (SI) and temporal information (TI). The features of these sequences accurately represent the general SI-TI plots of their respective class. CompetitionCNN retains its good predictive performance on these non-CTC sequences, with the compression efficiency reaching -2.93% for a Class D sequence in the LDP configuration. On the contrary, ScratchCNN achieves -1.11% coding gains on average for Class D with the LDP configuration, with this performance

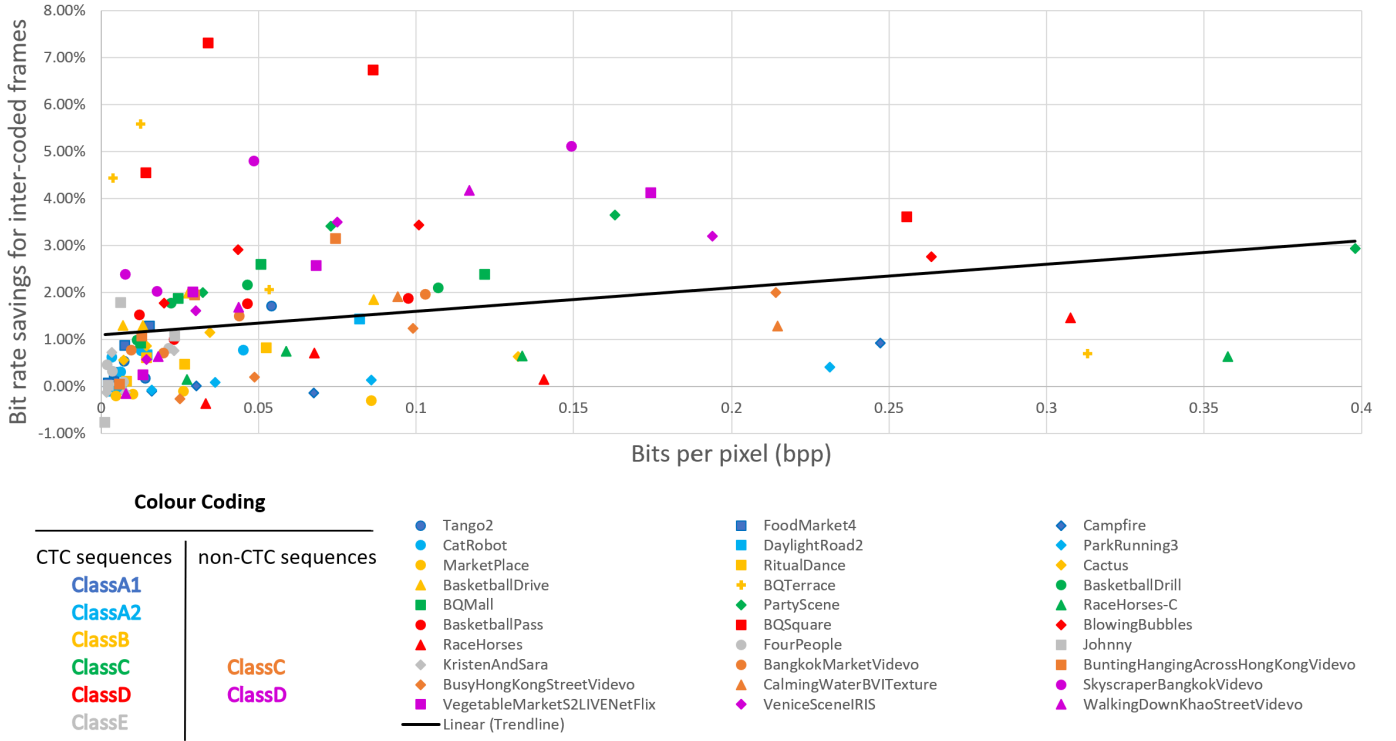


Fig. 8. Bit rate savings of the CompetitionCNN switchable filter implementation in VVC for various video sequences at 4 reference bit rates - for QPs 22, 27, 32, 37. Higher QP indicates lower bit rate. The savings are outlined against the bits per pixel needed to encode the first 31 inter-predicted frames of each sequence with the LDP configuration. A linear trendline is computed that takes into account every plot in the graph.

diminishing significantly when tested on Class C sequences.

Lastly, CompetitionCNN's coding complexity for CTC and non-CTC sequences is outlined in Table IV. While prior approaches that had been proposed for HEVC required GPU coding environments [16], [17], [20], the proposed approach can be applied in conventional CPU-based codec environments, with tolerable complexity increases in both encoding and decoding time. In [8], a full three-layer network was implemented within VVC and demonstrated increases in encoding complexity by up to 380 times. The CompetitionCNN implementation that only requires 13×13 filters displays a 6 to 11 encoding time increase and a 2 decoding time increase, when compared to the baseline. To allow fair comparison, Single Instruction/Multiple Data (SIMD) instructions were disabled for the VVC filters.

As explained in Section III-E, there are possibilities for further complexity optimisations, most prominently in terms of integerisation of the learned coefficients or reducing the size of convolutional kernels, resulting in a smaller filter size.

F. Evaluation on higher resolution video sequences

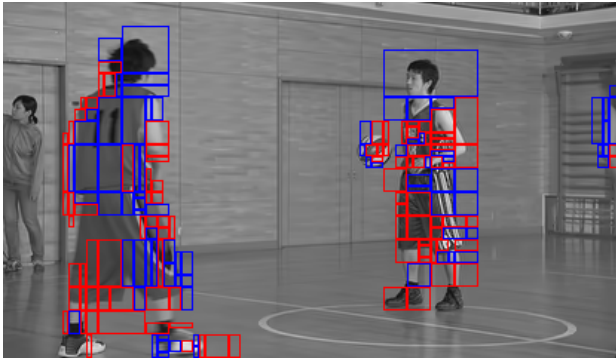
A pattern emerges when examining the results presented in Table III. For both CTC and non-CTC sequences, higher coding gains are achieved for the lower resolution class. This occurrence can partly be explained with the training dataset, a Class D sequence. Therefore, additional experiments have been carried out.

In Fig. 8, the bit rate savings of the proposed switchable filter approach are plotted against the bits per pixel (bpp)

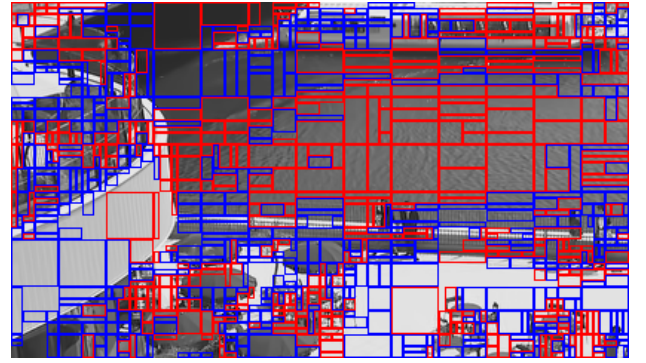
needed to encode the first 31 inter-predicted frames of CTC and non-CTC video sequences, for 4 QPs (22, 27, 32, 37) under the LDP configuration. As opposed to Tables II and III, the savings are calculated only for inter-coded frames and positive values indicate coding gains. The described video quality evaluation methodology was first presented in [31], with the script which computes this verification metric available in [32]. In this paper, the methodology is extended by computing the bit rate savings at specific points defined by reference QP bit rates and against the bits per pixel.

A linear trendline suggests that higher bit rate savings are achieved for sequences which require more bits per pixel. The bit rate of inter-coded frames increases with more inter-predicted blocks and more motion vectors to be transmitted, meaning there are actually savings that can be achieved. It is more important to have methods that enhance compression of such sequences for which low bit rates are difficult to achieve. Video sequences that are easier to encode and for which the CompetitionCNN switchable filter set does not provide significant bit rate savings are concentrated at the bottom-left of the graph. These sequences generally belong to higher resolution classes - Class A1 (4K), Class A2 (4K), Class B (1080p) and Class E (1080p). Both CTC and non-CTC sequences of classes C and D display higher bit rate savings and usually require more bits to encode a pixel. 30 out of the 40 plot points with a bits per pixel value higher than 0.05 belong to these two classes.

As the CompetitionCNN switchable filter implementation works better for sequences with higher bits per pixel values,



(a) BasketballPass, 2nd frame, QP 27, LDP configuration



(b) BQSquare, 465th frame, QP 27, LDP configuration

Fig. 9. Sub-pixel motion compensated blocks filter selection visualisation. Red CUs indicate the choice of CompetitionCNN filter set, blue CUs indicate the choice of the traditional filter set

TABLE V
SELECTION RATIO OF COMPETITIONCNN FILTERS AMONG
FRACTIONALLY INTERPOLATED BLOCKS, AT DIFFERENT QUANTISATION
LEVELS, LDP CONFIGURATION

Class	QP 22	QP 27	QP 32	QP 37
Class C (CTC)	47.20%	49.80%	47.55%	38.03%
Class D (CTC)	46.23%	50.25%	46.92%	38.87%

and those usually correspond to lower resolution classes, we have provided full coding performance results for the selected classes for CTC and non-CTC sequences, shown in Table III.

G. Learned interpolation filter selection analysis

To explore why adding an interpolation filter set extracted from the learned CompetitionCNN brings coding gains, a filter selection ratio test was performed. Each CU within a frame-to-be-encoded can choose between the learned interpolation filter set and the traditional filter set during the Rate-distortion Optimisation (RDO) process. The ratio of CompetitionCNN-based filters being selected among all sub-pixel motion compensated CUs is presented in Table V. The results are presented for Class C and Class D CTC sequences, encoded with the LDP configuration. The hit ratio hovers around 50%, suggesting that the CompetitionCNN training framework, presented in Section III-D, has ensured the development of true switchable filters, and not filters that could fully replace existing ones.

To visualise this selection process on inter-predicted frames, two examples are shown in Fig. 9. In Fig. 9a, the learned interpolation filter set is used on the outlines of the moving basketball players which are partitioned into smaller blocks. In Fig. 9b, the learned filters are predominantly used on CUs positioned over the moving water, a notably challenging pattern to encode in visual data compression.

V. CONCLUSIONS

A novel approach for developing sub-pixel interpolation filters using simplified, shared and interpreted CNNs was presented in this paper. A novel architecture containing two shared layers and a last layer with 15 branches was introduced,

alongside a new training framework, in order to maximise the coding performance of the newly formed interpolation filters. The linear structure of the network allows layers to collapse into a 2D matrix for application during inference time. In this way, it is possible to directly extract and visualise the filter coefficients, i.e. to interpret what was learned and to further derive the relevant coefficients for the final pixel prediction within the filter set.

The approach was tested in the context of the challenging VVC codec as switchable interpolation filters. As shown in the experimental results, the new NN-based filters, extracted from a simple linear network, display considerable compression performance for lower resolution sequences while offering a much more sensible trade-off in increased coding complexity compared to prior approaches in this area, paving the way for possible practical implementations of NN-based interpolation filters within video encoders.

Further improvements and simplifications to the network training could be done, by adding an adaptive weight quantisation metric in the loss function, potentially resulting in optimised integerised representations of the learned filter coefficients. Finally, tackling the challenging problem of integrating the introduced method with all VVC tools would be beneficial, as it would allow for a verification of coding performance results under VTM CTC, along with a comparison of all prior approaches from related NN work that was proposed for HEVC.

REFERENCES

- [1] B. Bross, J. Chen, and S. Liu, "Versatile Video Coding (VVC) Draft 10," *Document JVET-S2001*, Teleconference meeting, July 2020.
- [2] C. Dong, C. C. Loy, K. He, and X. Tang, "Image Super-Resolution Using Deep Convolutional Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 295–307, 2016.
- [3] Y. Guo, Y. Liu, A. Oerlemans, S.-Y. Lao, S. Wu, and M. Lew, "Deep Learning for Visual Understanding: A Review," *Neurocomputing*, vol. 187, 11 2015.
- [4] M. G. Blanch, M. Mrak, A. F. Smeaton, and N. E. O'Connor, "End-to-End Conditional GAN-based Architectures for Image Colourisation," in *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSP)*, 2019, pp. 1–6.
- [5] P. Helle, J. Pfaff, M. Schäfer, R. Rischke, H. Schwarz, D. Marpe, and T. Wiegand, "Intra Picture Prediction for Video Coding with Neural Networks," in *2019 Data Compression Conference (DCC)*, March 2019, pp. 448–457.

- [6] M. G. Blanch, S. Blasi, A. F. Smeaton, N. E. O'Connor, and M. Mrak, "Attention-based Neural Networks for Chroma Intra Prediction in Video Coding," *IEEE Journal on Selected Topics on Signal Processing*, 2021.
- [7] B. Bross, J. Chen, J. R. Ohm, G. J. Sullivan, and Y. K. Wang, "Developments in International Video Coding Standardization After AVC, With an Overview of Versatile Video Coding (VVC)," *Proceedings of the IEEE*, pp. 1–31, 2021.
- [8] L. Murn, S. G. Blasi, A. F. Smeaton, N. E. O'Connor, and M. Mrak, "Interpreting CNN for Low Complexity Learned Sub-Pixel Motion Compensation in Video Coding," in *IEEE International Conference on Image Processing, ICIP 2020*. IEEE, 2020, pp. 798–802. [Online]. Available: <https://doi.org/10.1109/ICIP40778.2020.9191193>
- [9] K. Ugur, A. Alshin, E. Alshina, F. Bossen, W. Han, J. Park, and J. Lainema, "Interpolation filter design in HEVC and its coding efficiency - complexity analysis," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 1704–1708.
- [10] G. Sullivan, J.-R. Ohm, and T. Wiegand, "Overview of the High Efficiency Video Coding (HEVC) standard," *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 22, 12 2012.
- [11] 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 Transactions on Circuits and Systems for Video Technology*, vol. 29, no. 11, pp. 3291–3301, 2019.
- [12] J. Mao and L. Yu, "Convolutional Neural Network Based Bi-Prediction Utilizing Spatial and Temporal Information in Video Coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 7, pp. 1856–1870, 2020.
- [13] S. Huo, D. Liu, F. Wu, and H. Li, "Convolutional Neural Network-based Motion Compensation Refinement for Video Coding," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018, pp. 1–4.
- [14] N. Yan, D. Liu, H. Li, and F. Wu, "A Convolutional Neural Network Approach for Half-pel Interpolation in Video Coding," *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2017. [Online]. Available: <http://dx.doi.org/10.1109/ISCAS.2017.8050458>
- [15] G. Bjontegaard, "Calculation of average PSNR differences between RD-curves," *VCEG-M33*, 2001.
- [16] C. D. Pham and J. Zhou, "Deep Learning-Based Luma and Chroma Fractional Interpolation in Video Coding," *IEEE Access*, vol. 7, pp. 112 535–112 543, 2019.
- [17] N. Yan, D. Liu, H. Li, B. Li, L. Li, and F. Wu, "Convolutional Neural Network-Based Fractional-Pixel Motion Compensation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 29, no. 3, pp. 840–853, March 2019.
- [18] Y. Dai, D. Liu, and F. Wu, "A Convolutional Neural Network Approach for Post-Processing in HEVC Intra Coding," *Lecture Notes in Computer Science*, p. 28–39, Dec 2016. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-51811-4_3
- [19] N. Yan, D. Liu, H. Li, B. Li, L. Li, and F. Wu, "Invertibility-driven Interpolation Filter for Video Coding," *IEEE Transactions on Image Processing*, vol. 28, no. 10, pp. 4912–4925, 2019.
- [20] J. Liu, S. Xia, W. Yang, M. Li, and D. Liu, "One-for-All: Grouped Variation Network-Based Fractional Interpolation in Video Coding," *IEEE Transactions on Image Processing*, vol. 28, no. 5, pp. 2140–2151, 2019.
- [21] L. Jing, J. Zbontar, and Y. LeCun, "Implicit Rank-Minimizing Autoencoder," *ArXiv*, vol. abs/2010.00679, 2020.
- [22] M. Santamaria, S. Blasi, E. Izquierdo, and M. Mrak, "Analytic Simplification of Neural Network Based Intra-Prediction Modes for Video Compression," *2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, Jul 2020. [Online]. Available: <http://dx.doi.org/10.1109/ICMEW46912.2020.9106027>
- [23] J. Boyce, K. Suehring, X. Li, and V. Seregin, "JVET Common Test Conditions and Software Reference Configurations," *Document JVET-J1010*, Ljubljana, Slovenia, July 2018.
- [24] S. K. J. Chen, Y. Ye, "Algorithm Description for Versatile Video Coding and Test Model 6 (VTM)," *Document JVET-O2002*, Gothenburg, July 2019.
- [25] P. Fleury, S. Bhattacharjee, L. Piron, T. Ebrahimi, and M. Kunt, "MPEG-4 Video Verification Model: A Solution for Interactive Multimedia Applications," *Journal of Electronic Imaging*, vol. 7, 05 2003.
- [26] D. Ma, F. Zhang, and D. R. Bull, "BVI-DVC: A Training Database for Deep Video Compression," 2020.
- [27] Video Free Stock Video Footage. [Online]. Available: <https://www.videvo.net/>
- [28] IRIS32 Free 4K UHD Free Footage. [Online]. Available: <https://www.youtube.com/channel/UCjJee-JAzoRRH5T0wqe7tw/>
- [29] M. A. Papadopoulos, F. Zhang, D. Agrafiotis, and D. Bull, "A video texture database for perceptual compression and quality assessment," in *2015 IEEE International Conference on Image Processing (ICIP)*, 2015, pp. 2781–2785.
- [30] C. G. Bampis, Z. Li, I. Katsavounidis, T.-Y. Huang, C. Ekanadham, and A. C. Bovik, "Towards Perceptually Optimized End-to-end Adaptive Video Streaming," 2018.
- [31] T. K. Tan, R. Weerakkody, M. Mrak, N. Ramzan, V. Baroncini, J.-R. Ohm, and G. J. Sullivan, "Video Quality Evaluation Methodology and Verification Testing of HEVC Compression Performance," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 26, no. 1, pp. 76–90, 2016.
- [32] A. Seixas Dias and M. Mrak, "Visualisation of BD-rate interpolation curves for MOS and PSNR results," *Document JCTVC-U0183*, warsaw, June 2015.