

Online Hybrid Lightweight Representations Learning: Its Application to Visual Tracking

Ilchae Jung^{1,3}, Minji Kim^{1,2}, Eunhyeok Park³ and Bohyung Han^{1,2}

¹ASRI, Seoul National University

²ECE & IPAI, Seoul National University

³CSE, POSTECH

chey0313@postech.ac.kr, minji@snu.ac.kr, eh.park@postech.ac.kr, bhhan@snu.ac.kr

Abstract

This paper presents a novel hybrid representation learning framework for streaming data, where an image frame in a video is modeled by an ensemble of two distinct deep neural networks; one is a low-bit quantized network and the other is a lightweight full-precision network. The former learns coarse primary information with low cost while the latter conveys residual information for high fidelity to original representations. The proposed parallel architecture is effective to maintain complementary information since fixed-point arithmetic can be utilized in the quantized network and the lightweight model provides precise representations given by a compact channel-pruned network. We incorporate the hybrid representation technique into an online visual tracking task, where deep neural networks need to handle temporal variations of target appearances in real-time. Compared to the state-of-the-art real-time trackers based on conventional deep neural networks, our tracking algorithm demonstrates competitive accuracy on the standard benchmarks with a small fraction of computational cost and memory footprint.

1 Introduction

Recent studies on deep neural representation learning [He *et al.*, 2016; Szegedy *et al.*, 2015] have achieved remarkable progress in various computer vision tasks. However, since their success depends heavily on high-performance GPUs and large memory, it is still challenging to deploy trained models to real-world applications in resource-hungry environments, *e.g.*, visual surveillance, robot navigation, and embedded systems. Such drawbacks are aggravated in the problems handling videos, which critically demand efficient representation learning methodologies for real-time processing.

The mainstream approaches for reducing computational cost in neural networks are categorized into two groups: low-bit quantization and channel pruning. Low-bit quantization methods [Jung *et al.*, 2019; Bulat and Tzimiropoulos, 2019] attempt to reduce the computational burden of deep neural networks by using low-precision weights and activations. Although these approaches are successful under simple offline

learning scenarios for the tasks in the image domain, *e.g.*, image classification [Russakovsky *et al.*, 2015], they still suffer from the challenges in streaming data including online adaptation and fine-grained prediction with limited precision. The other paradigm of model compression, channel pruning focuses on reducing the number of channels in each layer of deep neural networks [He *et al.*, 2017; Gao *et al.*, 2018; Kang and Han, 2020]. Recently, a couple of channel pruning methods [Jung *et al.*, 2020; Che *et al.*, 2018] have shown competitive performance to the baseline models with full parameters on online tasks for streaming data. However, their compression efficiency concerning the size and inference cost remains inferior to the low-bit quantization methods.

By observing the complementary characteristics of the two methods, we propose a novel framework to learn hybrid representations for streaming data based on an ensemble of low-bit quantization and compact full-precision networks. In the fusion of the hybrid networks, the low-bit quantized network delivers primary information of data efficiently while the lightweight channel-pruned network augments missing details in the representation and facilitates online model update to handle temporal variations of data. We evaluate the proposed algorithm on online visual tracking, which requires online model updates and highly accurate bounding box alignments in real-time streaming data. Our algorithm is unique in the sense that the output model is lightweight and suitable for systems with resource constraints. Moreover, we empirically verify that our hybrid representation model is applicable to various types of tracking approaches [Jung *et al.*, 2018; Li *et al.*, 2019; Chen *et al.*, 2021]. Our contributions are three-folded as follows:

- We propose a hybrid representations learning framework based on low-bit quantization and channel pruning by learning an efficient low-bit-width primary network and a precise residual network, respectively.
- We successfully deal with the trade-off between two popular model compression techniques and facilitate online adaptation of the hybrid network on streaming data by incorporating a residual architecture.
- Our learning algorithm achieves competitive performance in terms of accuracy and efficiency and, in particular, illustrates competency in the resource-aware environment.

2 Related Works

2.1 Low-bit Quantization

Recent studies on low-bit quantization [Jung *et al.*, 2019; Park and Yoo, 2020] have made successful progress in image classification. The research related to binary quantization [Hubara *et al.*, 2016; Bulat and Tzimiropoulos, 2019] shows the potential of low-precision optimization, achieving reasonable accuracy only with binarized activations and parameters. Other approaches based on the 4- or 5-bit quantization [Jung *et al.*, 2019; Park and Yoo, 2020] propose more practical solutions, accomplishing lower or negligible accuracy loss while enabling efficient fixed-point computation.

On the other hand, there exists another research trend that attempts to boost the performance of quantized models with assistance from full precision networks [Zhuang *et al.*, 2020]. Specifically, they train the quantization networks simultaneously with parameter-shared full precision networks to allow the relevant gradient flow for better accuracy. However, the full precision network is abandoned for inference, which loses the full potency of the hybrid structure. We judiciously design the framework that utilizes the benefit of both full-precision and quantized networks for visual tracking.

To mitigate the instability of training on network quantization, [Qin *et al.*, 2020] learns binary neural networks with an additional regularizer that maximizes the entropy of the distribution of quantized model parameters. On the other hand, [Yu *et al.*, 2020] proposes a new regularizer to let the distribution of quantized model parameters mimic the uniform distribution. Although those methods achieve competitive performance in offline problems such as image classification, their accuracy is degraded substantially during online updates, making the application to visual tracking not straightforward. This paper introduces the newly designed normalization method considering the quantization interval to maintain the feature quality during online model updates.

2.2 Channel Pruning

Another prominent deep neural network lightening method, channel pruning, aims to reduce the number of channels in model parameters. The pioneering work based on the LASSO regularization [He *et al.*, 2017] enforces the sparsity of the model while maintaining the accuracy in the image classification task. [Gao *et al.*, 2018] has proposed the feature boosting method by dynamically pruning channels at run-time using the channel saliency predictor. Besides, [Jung *et al.*, 2020] adopts the conventional channel pruning method [He *et al.*, 2017] to online object tracking, showing considerable performance gain. While the power of channel pruned models is valid for both online and offline learning, their compression rates remain inferior to low-bit quantization.

Our hybrid representation learning algorithm aims to improve accuracy and efficiency by combining complementary representations from a low-bit quantization network and a lightweight full precision network. While maintaining primary information in the quantization network obtained from rich offline pretraining, the full precision network compressed by channel pruning allows straightforward the online model update as well as accurate real-valued representation learning.

3 Hybrid Representation Learning

3.1 Overview

Our goal is to enhance the quality of representations using an ensemble of two lightweight deep neural networks obtained from quantization and channel pruning. The quantized network is responsible for feature encoding of the primary time-invariant information while the channel-pruned model covers high-precision time-varying residual information.

Given an input image A_0 , the standard convolutional neural network (CNN) learns a feature extractor $\mathbf{f}_{\text{conv}}(\cdot; \cdot)$ and estimates a feature map A_{conv} as follows:

$$A_{\text{conv}} \equiv A_{N_l} = \mathbf{f}_{\text{conv}}(A_0; \mathcal{W}_{\text{conv}}), \quad (1)$$

where $\mathcal{W}_{\text{conv}} = \{W_1, \dots, W_{N_l}\}$ is a collection of convolutional filters throughout the N_l layers. In the i^{th} intermediate layer, we apply a convolution with auxiliary operations and derive a feature map, which is given by

$$A_i = \sigma(A_{i-1} * W_i), \quad (2)$$

where $\sigma(\cdot)$ denotes a non-linear activation function and $*$ indicates the convolution operator. Contrary to the conventional CNNs, our framework utilizes two networks in parallel, a quantized network $\mathbf{f}^q(\cdot; \cdot)$ and a channel-pruned network $\mathbf{f}^p(\cdot; \cdot)$, which generate outputs respectively as

$$A^q = \mathbf{f}^q(A_0; \mathcal{W}^q), \quad \mathcal{W}^q = \{W_1^q, \dots, W_{N_l}^q\}, \quad (3)$$

$$A^p = \mathbf{f}^p(A_0; \mathcal{W}^p), \quad \mathcal{W}^p = \{W_1^p, \dots, W_{N_l}^p\}. \quad (4)$$

Then, the proposed hybrid representation, A^h , is given by the sum of two outputs as follows:

$$A^h = A^q + A^p = \mathbf{f}^q(A_0; \mathcal{W}^q) + \mathbf{f}^p(A_0; \mathcal{W}^p). \quad (5)$$

We can train prediction networks of arbitrary tasks followed by the feature extraction. Note that the hybrid representation is also utilized in the prediction networks; channel pruning is applied to all layers of the prediction networks while low-bit quantization is partially applied to the activation quantization.

3.2 Low-Bit Quantization of Primary Network

To maximize the benefit of the reduced precision network, we quantize both model parameters and activations in all N_l layers. We adopt the quantization method proposed in [Jung *et al.*, 2019], and introduce the additional regularization loss terms to the quantized data for fast and stable training.

Weight Quantization

The weight quantization function $Q_W^l(\cdot) : w_l \rightarrow w_l^q$ is composed of two steps. First, the weights are linearly transformed to $[-1, 1]$ by the learnable parameters c_W^l and d_W^l , which define the center and the width of the quantization interval, respectively. Only the weight values within the interval $[c_W^l - d_W^l, c_W^l + d_W^l]$ are quantized while the rests are clamped. The formal definition of the first step is given by

$$\hat{w}_l = \begin{cases} \text{sign}(w_l) & \text{if } |w_l| > c_W^l + d_W^l \\ 0 & \text{if } |w_l| < c_W^l - d_W^l \\ \text{sign}(w_l) \cdot (\alpha_W^l |w_l| + \beta_W^l) & \text{otherwise,} \end{cases} \quad (6)$$

where $\text{sign}(\cdot) \in \{-1, 1\}$ means the sign of an input, $\alpha_W^l = 0.5/d_W^l$, and $\beta_W^l = -0.5c_W^l/d_W^l + 0.5$. Second, the normalized weights are discretized as

$$W_l^q = \text{round}(\hat{W}_l \cdot (2^{N_b-1} - 1)), \quad (7)$$

where $\text{round}(\cdot)$ denotes the element-wise rounding and N_b is the bit-width, leading to $(2^{N_b} - 1)$ quantization levels.

Activation Quantization

Similarly, the activation quantization function $Q_A^l(\cdot) : a_l \rightarrow a_l^q$ has transformation and round steps. First, the activations $a_l \in A_l$ in the l^{th} layer are linearly transformed to the values in $[0, 1]$ as follows:

$$\hat{a}_l = \begin{cases} 1 & \text{if } a_l > c_A^l + d_A^l \\ 0 & \text{if } a_l < c_A^l - d_A^l \\ \alpha_A^l A_l + \beta_A^l & \text{otherwise,} \end{cases} \quad (8)$$

where $\alpha_A^l = 0.5/d_A^l$ and $\beta_A^l = -0.5c_A^l/d_A^l + 0.5$. The quantization interval is defined by the learnable parameters for its center (c_A^l) and width (d_A^l). We assume that the activations are non-negative due to the ReLU function. The final quantized output A_l^q is obtained by applying the rounding function to the individual elements in \hat{A}_l ,

$$A_l^q = \text{round}(\hat{A}_l \cdot (2^{N_b} - 1)). \quad (9)$$

Operation of Quantization Network

The output from each convolutional layer is now given by $A_{l+1} = \sigma(Q_A^l(A_l) * Q_W^{l+1}(W_{l+1}))$, where the standard convolution operation is replaced by a hardware-friendly fixed-point operation for additional acceleration. By applying the quantization to the whole network, the representation of the final convolutional layer is denoted by $A^q \equiv \mathbf{f}^q(A_0; \mathcal{W}^q)$, where $\mathbf{f}^q(\cdot; \cdot)$ is the feature extraction network via the quantization of both weights and activations.

Training Quantization Networks with Normalization

We now present our training scheme for robust and stable quantization. In particular, we discuss a novel regularization loss for quantization, $\mathcal{L}_{\text{quant}}$, which is divided into weight and activation normalization losses as

$$\mathcal{L}_{\text{quant}} = \mathcal{L}_{\text{qW}} + \mathcal{L}_{\text{qA}}. \quad (10)$$

Weight normalization. Clipping model parameters using the learned interval $[c_W^l - d_W^l, c_W^l + d_W^l]$ leads to error in the quantized representations and makes the optimization procedure unstable. It turns out that the final accuracy is sensitive to the initialization of $\{c_W^l, d_W^l\}$, especially with a lower bit-width. To alleviate the limitation, we propose a simple but effective regularizer, \mathcal{L}_{qW} , enforcing a majority of parameters to be located within a certain range, which is given by

$$\mathcal{L}_{\text{qW}} = \sum_{i=1}^{N_l} (\text{avg}(\alpha_W^i |W_i| + \beta_W^i) - \mu_0)^2 + (\text{std}(\alpha_W^i |W_i| + \beta_W^i) - \sigma_0)^2, \quad (11)$$

where $\text{avg}(\cdot)$ and $\text{std}(\cdot)$ denote mean and standard deviation of all entries in their input matrices, respectively, $|\cdot|$ returns

an absolute value matrix of the same size to its input, and (μ_0, σ_0) are hyperparameters for desirable mean and standard deviation values. The proposed regularizer encourages the transformed model parameters to be in $[0, 1]$ and reduces the error resulting from the clamping process. Note that it allows the quantized network to identify more appropriate parameters $\{c_W^l, d_W^l\}$ and undergo more stable optimization.

Activation normalization. Clipping activations outside of the interval $[c_A^l - d_A^l, c_A^l + d_A^l]$ incurs critical information loss in practice. To mitigate this issue, we define the following regularization loss term \mathcal{L}_{qA} :

$$\mathcal{L}_{\text{qA}} = \sum_{i=1}^{N_l} [\text{avg}(A_i^{\text{bn}}) + 2 \cdot \text{std}(A_i^{\text{bn}}) - (c_A^l + d_A^l)]_+ + [\sigma_i - \text{std}(A_i^{\text{bn}})]_+ + [\mu_i - \text{avg}(A_i^{\text{bn}})]_+, \quad (12)$$

where $A^{\text{bn}} = \{A_i^{\text{bn}}\}_{i=1}^{N_l}$ is a set of activations after batch normalization and $[\cdot]_+$ denotes the ReLU function. Note that σ_i and μ_i are determined by combination of current quantization interval parameters (c_A^l, d_A^l). Assuming that the activation A_i^{bn} follows a Gaussian distribution, the regularization loss enforces the activations larger than the mean of the Gaussian and lower than $\text{avg}(A_i^{\text{bn}}) + 2 \cdot \text{std}(A_i^{\text{bn}})$ to be within the active range. Statistically speaking, it implies that almost 96% of positive values are trained to be located in the active range.

3.3 Channel Pruning of Residual Network

We propose a stochastic channel pruning algorithm using a Gumbel-softmax trick to generate discrete mask samples from the model with a continuous output space. Our algorithm learns a set of channel selection probability vectors, $\{\mathbf{b}_1, \dots, \mathbf{b}_{N_l}\}$, in each layer, where \mathbf{b}_l indicates which channel is likely to be sampled in the l^{th} layer. By employing the Gumbel-softmax trick introduced in [Jang *et al.*, 2017], \mathbf{b}_l generates discrete channel selection mask \mathbf{M}_l as follows:

$$\bar{\mathbf{M}}_l[i] = \frac{\exp\left(\frac{g_i + \log(\mathbf{b}_l[i])}{\tau}\right)}{\exp\left(\frac{g_i + \log(\mathbf{b}_l[i])}{\tau}\right) + \exp\left(\frac{g'_i + \log(1 - \mathbf{b}_l[i])}{\tau}\right)},$$

$$\mathbf{M}_l[i] = \text{round}(\bar{\mathbf{M}}_l[i]), \quad (13)$$

where g_i and g'_i represent random noise samples from the Gumbel distribution, and τ denotes a temperature. We apply the straight-through Gumbel estimator for gradient approximation in backpropagation similarly to [Jang *et al.*, 2017]. In the forward pass of training, we apply the pruning mask \mathbf{M}_l to the output activation A_l^p in the l^{th} layer, which is given by

$$A_l^p = \sigma(A_{l-1}^p * W_l^p) \odot \mathbf{M}_l, \quad l = 1 \dots N_l \quad (14)$$

where \odot denotes an channel-wise multiplication. Then, the representation of the final convolution layer is obtained by applying the channel pruning network $\mathbf{f}^p(\cdot; \cdot)$, as follows: $A^p = \mathbf{f}^p(A_0; \mathcal{W}^p)$. Note that a channel with a selection probability lower than the threshold is removed at testing. For training, the channel selection probabilities $\{\mathbf{b}_l\}_{l=1}^{N_l}$ are learned by minimizing the following loss term

$$\mathcal{L}_{\text{prune}} = \sum_{l=1}^{N_l} \|\bar{\mathbf{M}}_l\|_1, \quad (15)$$

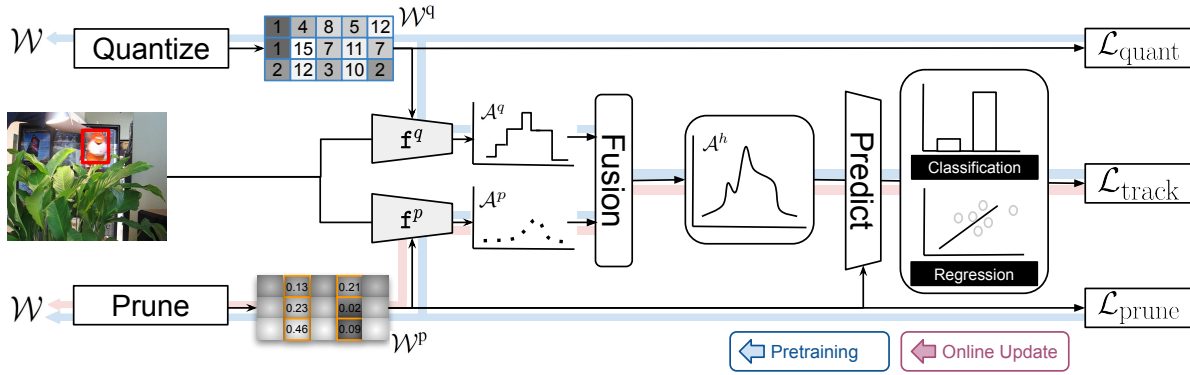


Figure 1: Overall framework of our algorithm. We train hybrid representation for both low-bit quantized network (f^q) and lightweight full precision network (f^p) on pretraining time while only the lightweight network is fine-tuned in test phase. Red arrow denotes the gradient backpropagation of full precision network in online model update while blue arrow denotes the backpropagation in pretraining.

where the number of active channels is penalized.

4 Hybrid Representation for Online Tracking

We apply the proposed hybrid representation learning method to online visual tracking, which aims to sequentially estimate the location of a target object in a streamed video given the ground-truth bounding box at the first frame. This task is appropriate for testing the proposed algorithm since visual tracking requires efficient online target appearance modeling and accurate bounding box regression.

The overall framework with online hybrid representation learning for visual tracking is illustrated in Figure 1, which presents two pathways of backpropagation for offline pretraining and online model updates. In the pretraining stage, our framework trains the parameters of both the quantized network and the channel-pruned network. On the other hand, during the online model update in test time, we freeze the parameters of the quantized network to maintain static coarse information while allowing the lightweight full precision network to adapt to target appearance changes. We apply our algorithm on top of three baseline trackers, RT-MDNet [Jung *et al.*, 2018], SiamRPN++ [Li *et al.*, 2019], and TransT [Chen *et al.*, 2021]. We pretrain each tracker with a loss term

$$\mathcal{L}_{\text{hyb}} = \mathcal{L}_{\text{baseline}} + \lambda_1 \mathcal{L}_{\text{quant}} + \lambda_2 \mathcal{L}_{\text{prune}}, \quad (16)$$

where $\mathcal{L}_{\text{baseline}}$ denotes a loss term proposed from baseline tracker, and $\{\lambda_1, \lambda_2\}$ are the balancing hyperparameters.

Hybrid RT-MDNet. RT-MDNet first generates a feature map from a whole image and then predicts candidate windows sampled from the feature map. Hybrid RT-MDNet applies our hybrid representation network to the feature extraction network while the prediction network employs the channel-pruned network only. We perform online updates with the channel-pruned network.

Hybrid SiamRPN++. SiamRPN++ deals with a tracking algorithm via the cross-correlation between the target representation and the candidate feature map using Siamese networks. The predictor network, which is composed of a binary classifier and a bounding box regressor, estimates the confidence score of each adjusted candidate location based on the

cross-correlation results. To formulate Hybrid SiamRPN++, we apply our hybrid representation network to the Siamese backbone and adopt channel-pruned network in the predictor.

Hybrid TransT. TransT is a variant of SiamRPN++ that applies the Transformer encoder-decoder networks [Vaswani *et al.*, 2017] to the feature correlation. Similarly to Hybrid SiamRPN++, we incorporate the hybrid network into the backbone and employ the channel-pruned network in other parts. Additionally, for all self-attention in the Transformer, we perform the activation quantization to reduce the computation burden in the dot-product operations.

5 Experiments

5.1 Preparation for Training

To pretrain the hybrid trackers, we employ the same training datasets as the baseline networks [Jung *et al.*, 2018; Li *et al.*, 2019; Chen *et al.*, 2021]. We adopt the pretrained model of each baseline tracker for initialization and fine-tune them using the proposed loss function. The learned hybrid models are tested on the standard visual tracking benchmarks, OTB2015 [Wu *et al.*, 2015], UAV123 [Mueller *et al.*, 2016], LaSOT [Fan *et al.*, 2019], VOT2016 [Kristan *et al.*, 2016], and VOT2018 [Kristan *et al.*, 2018], for comparisons with recently proposed competitors. We evaluate algorithms using Precision rate (Prec) and Success rate (Succ) in OTB2015, Expected Average of Overlap (EAO) in VOT2016/2018, and Normalized Precision (Norm. Pr) in LaSOT.

5.2 Implementation Details

We employ the same learning rates as the baselines but reduce the rate of the quantized network to 1/10 of the original value in Hybrid SiamRPN++. As in [Jung *et al.*, 2019], we use the full precision in the first convolution layer in Hybrid SiamRPN++ and Hybrid TransT.

For Hybrid RT-MDNet, we train the model with batch size 48 and run 120K iterations while optimizing model parameters and auxiliary parameters using SGD and Adam, respectively. For Hybrid SiamRPN++, training is performed on 8 GPUs for 100K iterations with batch size 22, and the same optimization methods are employed as Hybrid RT-MDNet.

	HybRT-MDNet						HybSiamRPN++					HybTransT				
	conv1	conv2	conv3	fc4	fc5	fc6	ℓ_1	ℓ_2	ℓ_3	ℓ_4	RPN	ℓ_1	ℓ_2	ℓ_3	att	pred
Bit (A/W)	32/4	4/4	4/4	-	-	-	32/5	5/5	5/5	5/5	5/32	32/5	5/5	5/5	6/32	32/32
LayerRatio	0.11	0.25	0.32	0.29	0.03	0.00	0.02	0.04	0.17	0.51	0.26	0.02	0.06	0.24	0.65	0.03
RC _{prune}	0.19	0.04	0.09	0.53	0.26	0.49	0.09	0.08	0.07	0.06	1.00	0.07	0.06	0.06	0.39	1.00
RC _{quant}	0.13	0.03	0.03	-	-	-	0.08	0.04	0.04	0.04	-	0.08	0.04	0.04	0.16	-
RC _{hybrid}	0.32	0.07	0.12	0.53	0.26	0.49	0.17	0.12	0.12	0.10	1.00	0.15	0.10	0.11	0.55	1.00
RC (BOP _{total})	0.26 (= 4.9B / 19.0B)						0.32 (= 5.1M / 15.2M)					0.42 (= 4.4M / 10.5M)				

Table 1: The computational cost of the proposed approach. We compare the relative amount of computation in channel-pruned networks, quantized networks, and our hybrid representation models compared to the original network. Bit (A/W) denotes the bit-width of input activations (A) and parameters (W), and LayerRatio means the normalized computational overhead of each layer observed in the baseline network. RC_{prune}, RC_{quant}, and RC_{hybrid} indicate the relative cost of each layer compared to the baseline in the channel-pruned, quantized, and hybrid networks, respectively. Finally, BOP_{total} presents the total BOPs of the proposed hybrid model and the baseline network.

Method	Bit	RC	OTB2015	
			Prec	Succ
RT-MDNet	32	1.00	85.3	61.9
Q	4	0.20	82.0	58.7
Q + N (10)	4	0.20	83.7	61.8
Q + N (10) + P (15)	4	0.26	84.9	63.1
SiamRPN++	32	1.00	90.5/87.6	69.5/66.3
Q	5	0.24	86.1	66.5
Q + N (10)	5	0.24	87.3	67.1
Q + N (10) + P (15)	5	0.32	89.5	68.5

Table 2: Ablative experiment for training objectives. Q, N, and P denote naïve quantization, normalization, and pruning, respectively. In the OTB results of SiamRPN++, we present the scores reported in the paper and the reproduced accuracies.

	Relative cost (RC)							OTB2015	
	conv1	conv2	conv3	fc4	fc5	fc6	total	Prec	Succ
PQ	0.32	0.07	0.12	0.53	0.26	0.49	0.26	84.9	63.1
QQ	0.26	0.06	0.06	0.06	0.03	1.00	0.12	8.3	12.1
PP	0.80	0.32	0.78	0.58	0.35	0.58	0.58	86.6	61.8

Table 3: Analysis of the combinations of the quantized network and the channel-pruned network in HybRT-MDNet. We compare our hybrid representation (PQ) to Quantization-only ensemble representation (QQ) and Pruning-only ensemble representation (PP).

Hybrid TransT is trained for 400 epochs (1K iterations per epoch) using the AdamW optimizer with batch size 40, where the learning rate is decayed by a factor of 10 after 250 epochs.

Two variants of SiamRPN++ are tested, HybSiamRPN++ and OnHybSiamRPN++, which correspond to the models based on offline and online hybrid representation learning, respectively. Specifically, during the online update stage, OnHybSiamRPN++ fine-tunes the last linear layer in the classification branch stemmed from the first cross-correlation block every 10 frames for 15 iterations using the data collected from the latest 20 frames. The sample features are collected after the target prediction and annotated by the same labeling strategy used in the pretraining stage.

Our algorithm is implemented in PyTorch with Intel I7-6850k and NVIDIA Titan Xp Pascal GPU for RT-MDNet and Quadro RTX4000 for SiamRPN++ and TransT.

	Relative cost (RC _{prune})							OTB2015	
	conv1	conv2	conv3	fc4	fc5	fc6	total	Prec	Succ
1	0.19	0.04	0.09	0.53	0.26	0.49	0.22	84.9	63.1
2	0.32	0.10	0.10	0.53	0.26	0.51	0.25	85.2	62.5
3	0.41	0.16	0.17	0.53	0.26	0.51	0.30	85.3	62.5
4	0.51	0.24	0.27	0.53	0.31	0.52	0.37	87.1	62.7
5	0.62	0.35	0.35	0.63	0.36	0.55	0.46	86.7	62.4
6	0.72	0.50	0.51	0.69	0.46	0.62	0.58	86.6	62.7

Table 4: Accuracy-computation trade-off given by varying channel-pruning ratios in HybRT-MDNet.

Bit-width	2bit	3bit	4bit	5bit
Prec	84.7	84.3	84.9	85.5
Succ	59.9	60.7	63.1	63.2

Table 5: Accuracy-computation trade-off given by varying bit-widths in HybRT-MDNet.

5.3 Effects on Acceleration

Following [Baskin *et al.*, 2018], we measure the number of bit-wise operations (BOPs) for unit convolution to estimate the performance benefit of the low-precision computation. For each layer with $\{C^{\text{in}}, C^{\text{out}}\}$ input and output channels based on $K \times K$ kernels, BOPs is given by

$$\text{BOPs} = C^{\text{in}} C^{\text{out}} K^2 (b_a b_w + b_a + b_w + \log_2 C^{\text{in}} K^2), \quad (17)$$

where b_a and b_w are the precisions of input activations and parameters, respectively. Table 1 shows the layer-wise analysis of the effect of our optimization method.

The proposed hybrid representation learning approach reduces the computation overhead by maintaining a small fraction of parameters that are critical to accuracy. Specifically, we observe more channels pruned in mid-layers (conv2-3 and ℓ_2 -4) than the parameters in head layers (conv1 and ℓ_1) and the prediction layers (fc, RPN, att, and pred), which tend to be intact for the stability of training. Consequently, the proposed model with hybrid representations only requires a fraction of BOPs (0.26 to 0.42 as presented in Table 1) compared to the original network.

5.4 Effect of Hybrid Representation

We analyze the impact of each component in the proposed hybrid representation learning. Note that we use HybRT-

Method	Bit	RC	OTB		UAV123		LaSOT		VOT2016	VOT2018
			Succ	Prec	Succ	Prec	Succ	Norm. Pr	EAO	EAO
ATOM [Danelljan <i>et al.</i> , 2019]	32	-	67.1	87.9	64.4	51.5	51.5	57.6	-	0.401
DaSiamRPN [Zhu <i>et al.</i> , 2018]	32	-	65.8	88.0	58.6	79.6	41.5	49.6	0.411	0.326
MDNet [Nam and Han, 2016]	32	-	67.8	90.9	52.8	74.7	39.7	46.0	0.257	-
RT-MDNet [Jung <i>et al.</i> , 2018]	32	1.00	65.0	88.5	52.8	77.2	30.8	36.0	0.325	0.176
HybRT-MDNet	5	0.29	64.5	88.0	52.6	79.5	29.5	35.1	0.329	0.194
SiamRPN++ [Li <i>et al.</i> , 2019]	32	1.00	69.5	90.5	60.6	79.8	49.3	57.7	0.408	0.340
SiamRPN++ (Rep)	32	1.00	66.3	87.6	57.0	76.9	45.6	54.5	0.401	0.305
HybSiamRPN++	5	0.32	68.5	89.5	59.7	78.8	45.9	53.6	0.383	0.299
OnHybSiamRPN++	5	0.34	69.4	89.8	61.2	81.8	46.1	54.1	0.412	0.344
TransT [Chen <i>et al.</i> , 2021]	32	1.00	68.2	88.3	66.0	85.2	64.2	73.5	0.387	0.298
HybTransT	5	0.42	68.2	88.4	67.1	86.5	64.5	73.9	0.418	0.285

Table 6: Experiments on several benchmarks with a comparison of other competitive trackers. Our hybrid models (HybRT-MDNet, HybSiamRPN++, OnHybSiamRPN++, HybTransT) present competitive performance compared to other trackers while their computational costs are significantly reduced.

MDNet pretrained in VOT-OTB dataset for the ablative study.

Benefit of each component. Table 2 illustrates the benefit of each component in our hybrid representation. The quantization with normalization (Q+N) (10) achieves better accuracy than the model by a simple quantization (Q), which implies the effectiveness of our normalization strategies for stable quantization. In addition, the hybrid representation model (Q+N+P) trained with both (10) and (15) outperforms the Q+N model, demonstrating that the proposed hybrid representation learning improves performance effectively. Finally, our final models present competitive accuracy to the baseline models while reducing the computation cost significantly.

Combinations of quantization and pruning. Table 3 illustrates the performance resulting from several combinations of quantized and pruned networks, which include the proposed hybrid network (PQ), the ensemble of pruned networks (PP), and the ensemble of quantized networks (QQ). We perform this experiment on the RT-MDNet baseline. The numbers in the table denote the relative computational cost in terms of BOPs between the model corresponding to each row and the RT-MDNet baseline. According to the results, the proposed hybrid representation presents a competitive performance in terms of both accuracy and efficiency, whereas the ensemble of quantized networks and pruned networks suffer from performance degradation in accuracy and computational cost, respectively.

Impact of pruning rates. Table 4 analyzes the computational cost by varying channel pruning ratios for our hybrid model. Note that our model based on the hybrid representation maintains competitive accuracy in both measures, precision and success rates, even after pruning almost 80% of channels.

Impact of quantization bit-width. Table 5 presents the experimental results by varying the bit-width for quantization in our hybrid model. The results show that the success rate drops sharply with less than 4-bit quantization while the precision rate is hardly affected by the bit-width variation.

5.5 Comparison on Multiple Benchmarks

Table 6 compares the performance of our hybrid networks with conventional models in visual tracking on the standard benchmarks, where SiamRPN++ and SiamRPN++ (Rep) indicate the official model and its reproduced version employed in our experiment, respectively. The proposed hybrid networks such as HybRT-MDNet, HybSiamRPN++, OnHybSiamRPN++, and HybTransT achieve competitive performance to the baseline models while significantly reducing the computational cost (0.29 to 0.42 in RC). Besides, the modified Hybrid SiamRPN++ with online update (OnHybSiamRPN++) improves tracking accuracy compared to the baseline SiamRPN++, which implies the effectiveness of our online hybrid representation learning for streaming data. Our model with the hybrid representation based on Transformer, HybTransT, also accomplishes outstanding accuracy with the significantly reduced computational burden.

6 Conclusion

We presented a novel framework of online hybrid representation learning applicable to streaming data by an ensemble of an efficient low-bit quantized network and a precise lightweight full-precision network. Adopting a residual architecture, the online hybrid representation successfully complements the two model compression approaches. We further improved the performance of the networks based on the hybrid representation through effective normalization of weights and activations for low-bit quantization. The proposed algorithm was incorporated into online visual tracking tasks and verified that the proposed hybrid representation makes competitive accuracy to baseline models despite low-bit quantization and channel pruning.

Acknowledgements

This work was supported in part by Samsung Electronics Co., Ltd. (IO210917-08957-01), and by the IITP grants [2021-0-01343, Interdisciplinary Program in AI (SNU); 2021-0-02068, AI Innovation Hub] and the NRF grant [2021M3A9E4080782, 2022R1A2C3012210] funded by the Korean government (MSIT).

References

- [Baskin *et al.*, 2018] Chaim Baskin, Eli Schwartz, Evgenii Zheltonozhskii, Natan Liss, Raja Giryes, Alex M Bronstein, and Avi Mendelson. UniQ: Uniform noise injection for non-uniform quantization of neural networks. *arXiv preprint arXiv:1804.10969*, 2018.
- [Bulat and Tzimiropoulos, 2019] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *arXiv preprint arXiv:1909.13863*, 2019.
- [Che *et al.*, 2018] Manqiang Che, Runling Wang, Yan Lu, Yan Li, Hui Zhi, and Changzhen Xiong. Channel pruning for visual tracking. In *ECCVW*, 2018.
- [Chen *et al.*, 2021] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021.
- [Danelljan *et al.*, 2019] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, 2019.
- [Fan *et al.*, 2019] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, 2019.
- [Gao *et al.*, 2018] Xitong Gao, Yiren Zhao, Łukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331*, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [He *et al.*, 2017] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.
- [Hubara *et al.*, 2016] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. In *NeurIPS*, 2016.
- [Jang *et al.*, 2017] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.
- [Jung *et al.*, 2018] Ilchae Jung, Jeany Son, Mooyeol Baek, and Bohyung Han. Real-Time MDNet. In *ECCV*, 2018.
- [Jung *et al.*, 2019] Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *CVPR*, 2019.
- [Jung *et al.*, 2020] Ilchae Jung, Kihyun You, Hyeonwoo Noh, Minsu Cho, and Bohyung Han. Real-time object tracking via meta-learning: Efficient model adaptation and one-shot channel pruning. *AAAI*, 2020.
- [Kang and Han, 2020] Minsoo Kang and Bohyung Han. Operation-aware soft channel pruning using differentiable masks. In *ICML*, 2020.
- [Kristan *et al.*, 2016] Matej Kristan, Aleš Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Luka Čehovin Zajc, Tomas Vojir, Gustav Häger, Alan Lukežič, and Gustavo Fernandez. The visual object tracking VOT2016 challenge results. In *ECCVW*, 2016.
- [Kristan *et al.*, 2018] Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Luka Čehovin Zajc, Tomas Vojir, Goutam Bhat, Alan Lukežic, Abdelrahman Eldesokey, Gustavo Fernandez, and et al. The sixth visual object tracking vot2018 challenge results. In *ECCVW*, 2018.
- [Li *et al.*, 2019] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *CVPR*, 2019.
- [Mueller *et al.*, 2016] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for UAV tracking. In *ECCV*, 2016.
- [Nam and Han, 2016] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016.
- [Park and Yoo, 2020] Eunhyeok Park and Sungjoo Yoo. Profit: A novel training method for sub-4-bit mobilenet models. In *ECCV*, 2020.
- [Qin *et al.*, 2020] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. Forward and backward information retention for accurate binary neural networks. In *CVPR*, 2020.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015.
- [Szegedy *et al.*, 2015] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [Wu *et al.*, 2015] Y. Wu, J. Lim, and M. Yang. Object tracking benchmark. *TPAMI*, 2015.
- [Yu *et al.*, 2020] Haibao Yu, Tuopu Wen, Guangliang Cheng, Jiankai Sun, Qi Han, and Jianping Shi. Low-bit quantization needs good distribution. In *CVPRW*, 2020.
- [Zhu *et al.*, 2018] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *ECCV*, 2018.
- [Zhuang *et al.*, 2020] Bohan Zhuang, Lingqiao Liu, Mingkui Tan, Chunhua Shen, and Ian Reid. Training quantized neural networks with a full-precision auxiliary module. In *CVPR*, 2020.