

ResQ: Residual Quantization for Video Perception

Davide Abati Haitam Ben Yahia Markus Nagel Amirhossein Habibian

Qualcomm AI Research*

{dabati,hyahia,markusn,ahabibia}@qti.qualcomm.com

Abstract

This paper accelerates video perception, such as segmentation and human pose estimation, by leveraging cross-frame redundancies. Unlike the existing approaches, which avoid redundant computations by warping the past features using optical-flow or by performing sparse convolutions on frame differences, we approach the problem from a different perspective: low-bit quantization. We observe that residuals, as the difference in network activations between two neighboring frames, exhibit properties that make them highly quantizable. Based on this observation, we propose a novel quantization scheme for video networks coined as *Residual Quantization*. ResQ extends the standard, frame-by-frame, quantization scheme by incorporating temporal dependencies that lead to better performance in terms of accuracy vs. bit-width. Furthermore, we extend our model to dynamically adjust the bit-width proportionally to the amount of changes in the video. We showcase the superiority of our model, against the standard quantization and existing efficient video perception models, using various architectures on semantic segmentation, video object segmentation and human pose estimation benchmarks.

1. Introduction

Despite the great progress made in optimizing the computational efficiency of deep neural networks, real-time video inference is still an open challenge in many cases, especially when deploying on resource-constrained devices [46, 15, 68, 14, 9, 21, 25]. The most effective strategy to accelerate video inference is to exploit the temporal redundancy, *i.e.* by leveraging what has been processed already in the past. Early works relied on feature warping and adaptation using optical flow [68, 25] or self-attention [21]. More recently, some encouraging results have been obtained by decomposing video snippets into a keyframe followed by residuals, that can be more efficiently processed

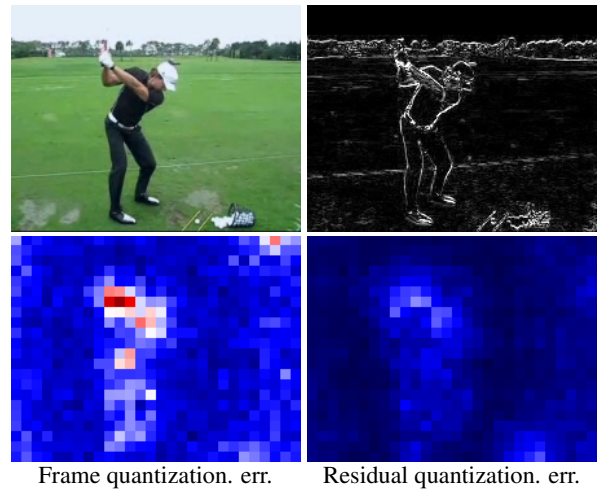


Figure 1: **Visualization of the quantization error** (bottom) when using frame representations (left) and residual representations (right). The residual quantization error is typically lower (blue) than the one on individual frames (red).

with distilled networks [15] or sparse operators [14, 46, 9].

For many applications, neural network quantization has emerged over the years as one of the key techniques for accelerating floating point models [10, 42, 30, 39, 2, 23], and for their deployment in integer precision. However, when dealing with video inputs, quantized models still operate by processing every frame independently, neglecting the opportunity to exploit redundancy among frames.

In this paper, we observe that residuals, as the difference in network activations between adjacent frames, exhibit properties that make them highly quantizable. Specifically, we illustrate that residuals typically exhibit a smaller variance w.r.t. the frame activations, which results in a reduction in quantization error as illustrated in Fig. 1. Following this observation, we propose Residual Quantization, coined as *ResQ*, a novel quantization scheme tailored for video perception. ResQ employs two sets of quantizers: one at a higher precision to quantize keyframes, one at a lower precision to quantize the residuals for subsequent frames. Both quantizers interact during the inference to combine the

*Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

high-precision details from keyframes with complementary information from residuals. Furthermore, motivated by the fact that the range of residuals depends on the scene, we extend ResQ to dynamically adjust the quantization bit-width on the fly. More specifically, we propose a lightweight policy function that assigns a minimally acceptable bit-width when the residuals are small, such as in static scenes.

We extensively evaluate our proposals on three perception tasks, namely semantic segmentation, video object segmentation and human pose estimation. We experiment with various architectures and quantization techniques, *i.e.* post-training quantization and quantization-aware training. ResQ and its dynamic counterpart consistently outperform the standard quantization schemes, and perform favorably compared to state-of-the-art in efficient video processing.

We summarize our contributions as follows: *i)* We formally and empirically show the benefits of using frame residuals in reducing the quantization error (Sec. 3); *ii)* We propose ResQ, a novel quantization scheme for video networks leveraging the residual quantization (Sec. 4); *iii)* We extend our model to dynamically adjust the quantization level based on the residual content (Sec. 4.2); *iv)* We validate our proposals on three tasks, where our models achieve an optimal trade-off of accuracy vs. efficiency.

2. Related Work

Neural network quantization. Progress in building efficient neural networks followed from multiple directions, *e.g.* by pruning parameters [8, 18, 31], decomposing weights into low-rank tensors [18, 65], searching for optimal architectures [5, 38] and distilling knowledge into light models [19, 51]. Besides these strategies, low-precision network quantization [13, 23, 30, 42, 40] has proven to be extremely effective, mostly due to the robustness of deep models to noise in representations and to the wide support in the acceleration of fixed-point inference. In Post-Training Quantization (PTQ), a trained neural network is quantized without fine-tuning its parameters. This can be achieved by optimizing the quantization ranges [30, 2], correcting biases in quantization error [42, 11], performing layer-wise rescaling [42, 37], or using sophisticated rounding operations [39, 32]. Whenever labeled data is available it is possible to perform Quantization-Aware Training (QAT), to fine-tune the model while simulating quantized operations. In this respect, the gradient of rounding operations can be approximated by the straight-through estimator (STE) [4, 13, 23] or its variations [12, 22, 28]. For lower bit-widths, it is beneficial to learn the quantization ranges through gradient decent [10, 6, 26], possibly jointly with bin regularization [16] or oscillation damping [41]. Unlike all these efforts, we leverage temporal redundancies, a novel and complementary aspect to improve the quantization of models when being applied on a sequence of frames.

Efficient video perception. Most relevant to our effort are the works that leverage temporal redundancies to accelerate video processing. The seminal work Deep feature flow (DFF) [68] and its extensions [25, 67, 33, 34] avoid expensive feature computation by warping features from key-frames using optical-flow. Despite its effectiveness for expensive backbones, feature warping has become less appealing recently, as the cost of optical-flow estimation might outweigh feature computation. In TDNet [21] a wide backbone is divided into several sub-models, to whom frames are cyclically assigned; current and past representations are then merged via self-attention. Differently, our method saves computation by reducing quantization precision rather than network width, and does not require any additional temporal modeling or other modifications to the original architecture. Recent works focus on representing videos as residuals w.r.t. keyframes [14, 15, 46, 9]. Skip-convolutions [14] showcase how dense features can be approximated by performing few localized operations over residuals, leading to impressive theoretical speed-ups. However, their dependency on sparse convolutions makes them hard to deploy in practice. Despite more recent efforts [46, 9, 15], the development of practical solutions to process residuals is still an open challenge.

Despite the potential impact, no prior work explored how to leverage correlation of video frames to improve neural network quantization. A notable exception is VideoIQ [55], dedicated for action recognition. Specifically, it dynamically quantizes frames at high or low precision based on their importance to classify a video: higher bit-widths for *foreground* frames and lower bit-widths for *background* frames. As we tackle per-frame prediction tasks, *e.g.* semantic segmentation and pose estimation, ResQ aims at a precise representation on *all frames*. This aim is achieved by leveraging frame residuals, which is missing in VideoIQ. Moreover, VideoIQ relies on a complex recurrent policy model, trained jointly with the action classifier using discrete optimization and knowledge distillation. By relying only on the residual representation, ResQ is simply applicable to any off-the-shelf model with minimal modifications.

3. Background and Motivation

A floating-point tensor \mathbf{x} is quantized into a fixed-point tensor $\hat{\mathbf{x}}$ in b -bits using a quantization function $\hat{\mathbf{x}} = q(\mathbf{x}; \Theta)$, where Θ denotes the quantizer parameters. In the case of uniform affine symmetric quantization [40], the quantization function is defined as:

$$q(\mathbf{x}; \Theta) = s \left[\text{clamp} \left(\left\lfloor \frac{\mathbf{x}}{s} \right\rfloor, -2^{b-1}, 2^{b-1} - 1 \right) \right], \quad (1)$$

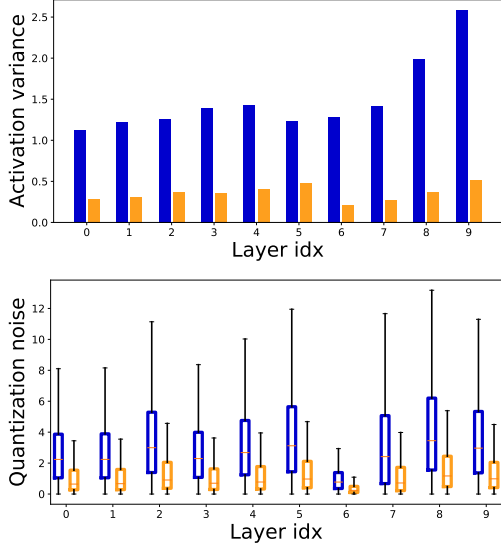


Figure 2: **Residuals vs. Frames** for 10 random layers from an HRNet. Residuals have a smaller variance than frames (top), translating to a lower quantization error (bottom).

where the quantizer parameters include the scaling factor s , and $\lfloor \cdot \rfloor$ is a rounding operation.

A floating-point convolutional function $\mathbf{z} = \mathbf{w} * \mathbf{x}$ is performed in fixed-point $\hat{\mathbf{z}} = \hat{\mathbf{w}} * \hat{\mathbf{x}}$ by quantizing its weight and input tensors, \mathbf{w} and \mathbf{x} , using the quantization function Eq. (1). The rounding and clipping operations introduce some quantization error denoted as $\epsilon = \mathbf{z} - \hat{\mathbf{z}}$. Reasonably, the smaller the quantization error is, the better the fixed-point model performs.

The motivation behind this work is that the inherent redundancies in video frames can be leveraged to reduce the quantization error. Indeed, we observe that residuals exhibit a significantly lower variance than the frame activations, and can be quantized with lower error. Fig. 2 supports this intuition by reporting the variances and quantization errors for frame vs. their residuals in several layers within an HRNet [59] for human pose estimation: as the figure illustrates, the aforementioned hypothesis is supported for several randomly sampled layers in the architecture. Indeed, in the case of the normally distributed activations, it is even possible to derive formally that highly correlated samples (neighboring frames) yield residuals with lower variance: we refer the interested reader to [49] for a full derivation.

Next, we further elaborate on the relation between activation variances and the quantization error, by analyzing the cases of weight and activation quantization separately.

Impact on weight quantization. We consider the output of a linear layer for which only weights are quantized, as

$\hat{\mathbf{z}} = \mathbf{x} * \hat{\mathbf{w}}$. The quantization error $\epsilon_{\mathbf{w}}$ can be expressed as:

$$\epsilon_{\mathbf{w}} = \mathbf{x} * \mathbf{w} - \mathbf{x} * \hat{\mathbf{w}} = \mathbf{x} * (\mathbf{w} - \hat{\mathbf{w}}). \quad (2)$$

As the equation shows, the error $\epsilon_{\mathbf{w}}$ is contributed to by the quantization error on weights ($\mathbf{w} - \hat{\mathbf{w}}$), that is not data dependent, and by the magnitude on the input \mathbf{x} . Considering the lower variance and average magnitude of residuals, convolving the quantized weight on them is likely to incur in a lower quantization error.

Impact on activation quantization. Similarly, for a linear layer with activation quantization, $\hat{\mathbf{z}} = \hat{\mathbf{x}} * \mathbf{w}$, the error can be expressed as $\epsilon_{\mathbf{x}} = (\mathbf{x} - \hat{\mathbf{x}}) * \mathbf{w} = \Delta \mathbf{x} * \mathbf{w}$. Assuming no clipping, the quantization error becomes

$$\Delta \mathbf{x} = \mathbf{x} - \mathbf{q}(\mathbf{x}; \Theta) = x - s \left\lfloor \frac{\mathbf{x}}{s} \right\rfloor = s \left(\frac{\mathbf{x}}{s} - \left\lfloor \frac{\mathbf{x}}{s} \right\rfloor \right), \quad (3)$$

and is therefore bound by $-\frac{s}{2} \leq \Delta \mathbf{x} \leq \frac{s}{2}$. As the quantization scale s is proportional to the magnitude and variance of the input \mathbf{x} , quantizing residuals with a smaller variance will lead to smaller quantization errors $\Delta \mathbf{x}$ and $\epsilon_{\mathbf{x}}$. An intuitive visualization of the relation between scale factor and quantization error is provided in Fig. 3.

4. Residual Quantization

Motivated by the intrinsic benefits of quantizing the residuals compared to the frames, we propose ResQ, a residual-based quantization scheme for video networks. We follow a sigma-delta formulation [44, 14, 15, 46, 9], a common approach to leverage the cross-frame redundancies, and elaborate how fixed-bit quantization can be integrated to accelerate video processing. Using the distributive property of linear functions, the output of a convolution on a activation \mathbf{x}^t can be formulated as:

$$\begin{aligned} \mathbf{z}^t &= (\mathbf{x}^t - \mathbf{x}^k) * \mathbf{w} + \mathbf{x}^k * \mathbf{w}, \\ &= \delta^t * \mathbf{w} + \mathbf{z}^k, \end{aligned} \quad (4)$$




Figure 3: **Impact of residual distribution on the quantization error:** The lower variance of residuals allow for smaller scale factors in their quantization (equivalent to the discrete bin-width in the plots). In turn, the expected error is lower.




Figure 4: **ResQ at inference.** The key-frame activation and weights, in blue, are quantized at higher precision, whereas the residual activations and weights highlighted in orange are quantized at lower precision. The activations at residual timesteps are obtained by summation with the key-frame activations. This scheme is applied for all layers throughout the model.

where δ^t is a residual w.r.t. a keyframe \mathbf{x}^k . Throughout this work, we set keyframes simply at fixed frame intervals, with a given period T . Eq. (4) is performed in fixed-point using the quantization function Eq. (1):

$$\begin{aligned} \hat{\mathbf{z}}^t &= \mathbf{q}(\delta^t; \Theta_a) * \mathbf{q}(\mathbf{w}; \Theta_w) + \hat{\mathbf{z}}^k, \\ \hat{\mathbf{z}}^k &= \mathbf{q}(\mathbf{x}^k; \Phi_a) * \mathbf{q}(\mathbf{w}; \Phi_w), \end{aligned} \quad (5)$$

where Φ_w and Φ_a denote the weight and activation quantization parameters for keyframes, respectively. The same set of quantization parameters is shared for all the keyframes. Similarly, Θ_w and Θ_a denote the weight and activation quantization parameters for residual frames. These can be shared for all residual frames or can be dynamically adjusted based on the residual content as discussed in Sec. 4.2. Separating the quantization parameters between keyframes and residuals allows amortizing high-precision computation on the keyframe with low-precision computations on residuals, for an overall lower quantization error. Inference with our model is represented in Fig. 4.

4.1. Estimating Quantizer Parameters

For quantizing both keyframes and residuals, we rely on a uniform affine symmetric scheme in Eq. (1). The scale factor s depends on the activation range of the quantizer, identified by r_{min}, r_{max} , as follows:

$$s = s(r_{max}, r_{min}) = \frac{2 \max(r_{max}, -r_{min})}{2^b - 1}. \quad (6)$$

At any convolution layer, we employ two different range setters for quantizing the input \mathbf{x} and the weights \mathbf{w} . Specifically, for weights, we employ a simple min-max strategy,

$$r_{max}^w = \max(\mathbf{w}), \quad r_{min}^w = \min(\mathbf{w}), \quad (7)$$

and directly compute the quantized weights $\hat{\mathbf{w}}$.

For estimating the range of activations r_{min}^x, r_{max}^x , we first collect exemplar inputs to the layer $\{\mathbf{x}_i\}_1^c$, by feeding c

calibration samples to the model. We then concatenate them into a batch \mathbf{X} , and build a line-search space \mathcal{S} between its minimum and maximum, with r candidate points, as $\mathcal{S} = \text{linspace}(\min(\mathbf{X}), \max(\mathbf{X}), r)$. We then search the range in \mathcal{S} , by minimizing the following objective, where the scale factor s is derived from the range using Eq. (6):

$$\arg \min_{r_{max}^x, r_{min}^x \in \mathcal{S}} \|\mathbf{X} * \mathbf{w} - \mathbf{q}(\mathbf{X}; \Theta = s(r_{max}^x, r_{min}^x)) * \hat{\mathbf{w}}\|_F. \quad (8)$$

4.2. Dynamic Residual Quantization

As discussed in Sec. 3, residuals have a lower variance compared to frames, thus are typically easier to quantize. However, not all residuals have similar characteristics, as they depend on the video content. For static scenes, residuals usually have a low magnitude, whereas in the presence of motion they can exhibit a higher dynamic range, making their quantization more challenging. This motivates us to design a *dynamic residual quantization* scheme that adaptively adjusts the precision to the amount of changes occurring in a video. To this aim, instead of using one set of parameters for quantizing all timesteps, we adjust the quantizer parameters based on the residual content. We consider this adjustment for the activation quantizer parameters only (Θ_a in Eq. (5)) as weights are static and input independent.

Dynamic quantization can be performed at various granularities, *e.g.* at frame-, region-, or pixel-level. We opt for the pixel-level scheme, which shares the same quantizer along the channel dimension for every pixel. This gives us full flexibility in adapting to any spatial differences in residual frames while still being computationally efficient. In dynamic residual quantization, every layer has a single quantizer for weights Θ_w , and a *pool* of quantizers, $\{\Theta_a^1 \dots \Theta_a^n\}$, ordered from low to high bit-width, available for input residuals. All quantizers are fit separately in the calibration stage, as explained in Sec. 4.1, identically to the case of static residual quantization.

During inference, given a residual $\delta \in \mathbf{R}^{c \times h \times w}$, we produce an index map $\Pi \in \{1 \dots n\}^{h \times w}$ as $\Pi = \pi(\delta)$, where π is a policy function selecting the optimal quantizer per pixel. We rely on the quantization error as the criteria to design the policy function. Specifically, for each quantizer Θ_a^i from the pool, we compute an error map $\epsilon_i \in \mathbf{R}^{h \times w}$ as:

$$\epsilon_i = \|\delta * \hat{\mathbf{w}} - \mathbf{q}(\delta; \Theta_a^i) * \hat{\mathbf{w}}\| = \|(\delta - \mathbf{q}(\delta; \Theta_a^i)) * \hat{\mathbf{w}}\|, \quad (9)$$

where the norm is to be intended pixel-wise. However, the estimation in Eq. (9) is troublesome, as it requires projecting the quantization error on residuals, $\delta - \mathbf{q}(\delta; \Theta_a^i)$ through $\hat{\mathbf{w}}$. Using Young's convolution inequality [62], we approximate the error map by multiplying the norms of $\delta - \mathbf{q}(\delta; \Theta_a^i)$ and $\hat{\mathbf{w}}$, where the former has to be intended pixel-wise:

$$\epsilon_i \leq \|\delta - \mathbf{q}(\delta; \Theta_a^i)\| \|\hat{\mathbf{w}}\|. \quad (10)$$

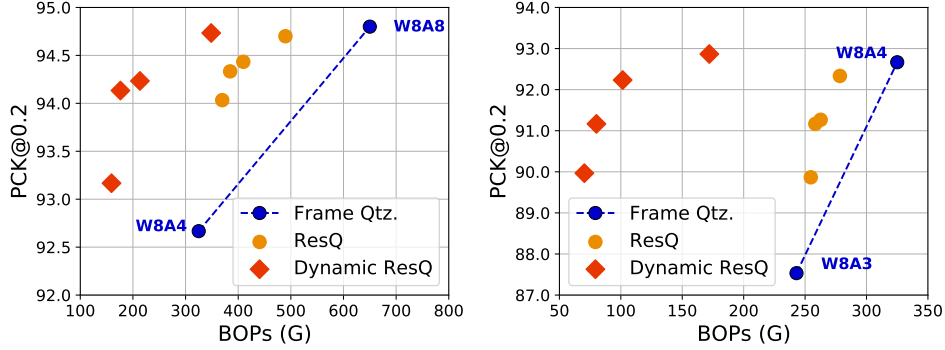


Figure 5: Comparison of quantization schemes on JHMDB. Curves generated by sweeping $T = 2, 4, 6, 8$.

The approximate error in Eq. (10) serves as an upper bound of the objective in Eq. (9) [62]. Once an error map ϵ_i is approximated for every quantizer in the pool $\Theta_a^1 \dots \Theta_a^n$, the policy function starts from the lowest precision and iteratively considers the reduction in quantization error given by selecting the next (higher bits) quantizer:

$$\Pi = \min_{i=1, \dots, n} i \quad \text{s.t.} \quad \epsilon_i - \epsilon_{i+1} < \tau. \quad (11)$$

Intuitively, for every residual pixel we consider the reduction in quantization error granted by the next precision Θ_a^{i+1} and we stop as soon as the gap is smaller than a predetermined threshold τ . For a visualization of the decisions taken by the policy function, we refer the reader to Fig. 6.

5. Experiments

Efficiency metric and evaluation. As current deep learning frameworks (e.g. PyTorch) currently support 32, 16, and 8-bit kernels only, we follow the common practice in quantization [3, 57] and rely on Bit Operations (BOPs) to measure the computational costs: BOPs are computed by re-weighting the multiply-accumulate (MAC) count of a model accounting for bit-widths of weights and activations. Since residual quantization utilizes different bit-width for the keyframe and the residual frames, we report precision in a unified notation: for instance, W8A8→W8A4 indicates bit-width of 8-8, for weights-activation of keyframes, and bit-width of 8-4 for weight-activation of residual frames.

Following [21, 46, 14], we split each video into sequences of length T , whose first frame is treated as keyframes. BOPs and accuracy metrics are averaged over all frames in a sequence. For the Dynamic-ResQ experiments, the policy function is considered in the BOP count.

Quantization methods. We conduct our experiments with both PTQ and QAT, as explained in Sec. 2. For PTQ, we use $c = 64$ calibration samples and $r = 20$ line search points for selecting the activation quantization range. For

QAT, we initialize the scale factor s of the quantizer via PTQ, and propagate gradients through the rounding operation in Eq. (1) by means of a straight-through estimator [4]. We either apply the same scaling factor s for the whole tensor, i.e. tensor quantization, or have separate s for each channel of the weight tensor, i.e. channel quantization. Unless specified otherwise, we utilize tensor quantization. For the dynamic policy, we fix $\tau = 0.0003$ for all experiments.

5.1. Human Pose Estimation

Dataset and accuracy metric. We use JHMDB [27], consisting of a collection of 316 clips, comprising more than 11,000 frames. Every frame is labeled with the location of 15 body joints, for a single person in the scene. The dataset comes with three different train/test splits, that we use accordingly with the standard protocol: every reported operating point is an average over all splits. As a performance metric, we rely on the standard PCK metric, with threshold $\alpha = 0.2$, as utilized in [53, 36, 43, 14].

Implementation details. As a starting floating point model, we rely on HRNet-w32 [54], pretrained on MPII Human Pose [1], that we fine-tune on every split for 100 epochs using the Adam optimizer [29] with a batch size of 16 images. The learning rate is initialized to $1e - 3$ and decayed by a factor 10 at epochs 40 and 80. We follow standard training augmentations such as scaling, rotations, flipping and cropping, following the implementation of [60].

	GBOPs	Hea.	Sho.	Elb.	Wri.	Hip	Kne.	Ank.	Avg
DKDet <i>et al.</i> [43]	553	98.3	96.6	90.4	87.1	99.1	96.0	92.9	94.0
S-SVD [24]	322	97.9	96.9	90.6	87.3	98.7	95.3	91.1	94.3
W-SVD [64]	325	97.9	96.3	87.2	82.8	98.1	93.2	88.8	92.4
L0 [35]	208	97.1	95.5	86.5	81.7	98.5	92.9	88.6	92.1
Skip-Conv[14]	539	98.7	97.7	92.0	88.1	99.3	96.6	91.0	95.1
ResQ	384	98.3	97.2	90.9	87.0	99.0	95.1	91.0	94.3
Dynamic-ResQ	176	98.2	97.0	91.0	87.1	98.7	94.6	90.5	94.1

Table 1: Comparison to state-of-the-art on JHMDB. Dynamic-ResQ achieves the best accuracy vs. efficiency trade-off.




Figure 6: **Dynamic-ResQ policy.** We represent policy decisions for `conv2` in HRNet. Dark blue, light blue and yellow pixels color code pixels that are quantized to 0, 4, and 8 bits respectively. Few regions require high precision, and many stationary regions are processed in low precision. Selected precision generally increases with the distance to the keyframe (Δt).

ResQ vs. frame quantization. We evaluate the effectiveness of ResQ by comparing it with Frame Quantization, the standard quantization strategy processing frames independently. We quantize all models using PTQ with the same setting, in terms of range setters and calibration samples, to foster a fair comparison. Fig. 5 reports the results for $W8A8 \rightarrow W8A4$ (left) and $W8A4 \rightarrow W8A3$ (right) bit-width configurations. ResQ achieves a better trade-off between accuracy vs. efficiency with respect to Frame Quantization, supporting our core hypothesis that quantizing the residuals leads to a smaller quantization error.

Dynamic vs. static ResQ. We evaluate the effectiveness of the dynamic residual quantization, proposed in Sec. 4.2, by comparing Dynamic-ResQ vs. ResQ in configurations of $W8A8 \rightarrow W8A\{0,4,8\}$ as reported in Fig. 5 (left). More specifically, we use 8-bit weights while dynamically adjusting the activation bit-width among 0, 4, and 8 bits. Similarly, Fig. 5 (right) reports the performance for $W8A8 \rightarrow W8A\{0,3,4\}$. As reported, Dynamic-ResQ outperforms both Frame Quantization and ResQ, advocating for the benefits of dynamically selecting the bit-width given the residual content. To gain more insights about the policy function, we represent in color coding its decisions in Fig. 6. As the figure shows, moving objects are quantized to a higher bit-width, whereas background and static areas tend to be assigned low-precision, contributing to a reduction in BOP cost without sacrificing accuracy.

Comparison with the state-of-the-art. Tab. 1 compares our method with the current state-of-the-art for efficient human pose estimation in video: weight compression [24, 64] and pruning [35] methods, as well as dynamic kernel distillation (DKD) [43] and Skip-Convolution [14] that leverage temporal redundancies to accelerate video inference. Except for DKD, all other methods use the same HRNet-w32 [54] backbone. The BOPs of competing methods have

been computed from the MACs reported in the corresponding papers and assuming lossless quantization to W8A8.

The results confirm that Dynamic-ResQ achieves the lowest computational cost, 176 GBOPs, whilst maintaining a high accuracy of 94.1. Such results are only 1% less accurate than the highest PCK reported by Skip-Conv while being $\sim 3\times$ more efficient. Moreover, both ResQ and Dynamic-ResQ perform favorably compared to the alternatives in terms of accuracy vs. efficiency trade-off. Finally, it is worth noting that in this experiment we rely only on PTQ and do not use any form of fine-tuning, contrary to all competing methods that need a supervised training procedure.

5.2. Semantic Segmentation

Dataset and metric. We use the Cityscapes dataset [7] and rely on the standard training and validation splits with 2,975 and 500 clips, respectively. Cityscapes provides pixel-level annotations into 19 classes for one frame per snippet. Following [15], we extract the per-frame pseudo-

Model	Backbone	Bit-width	BOPs (T) ↓	mIoU ↑
TDNet [21]	PSPNet [66]	W8A8*	34.5	79.9
TDNet [21]	BiseNet [63]	W8A8*	6.4	76.4
DFF [68]	ResNet-101 [17]	W8A8*	7	69.2
Delta Dist. [15]	HRNet-w18s [59]	W8A8*	2.2	75.7
ResQ	HRNet-w18s	W4A4→W3A3	0.9	76.9
Delta Dist. [15]	DDRNet-39 [20]	W8A8*	9.0	79.9
ResQ	DDRNet-39	W8A8→W4A4	9.0	80.8
ResQ	DDRNet-39	W4A4→W3A3	3.2	80.0
Delta Dist. [15]	DDRNet-23 [20]	W8A8*	4.6	78.9
ResQ	DDRNet-23	W8A8→W4A4	4.6	79.3
ResQ	DDRNet-23	W4A4→W3A3	1.6	78.8
Skip-Conv [14]	DDRNet-23s [20]	W8A8*	1.7	75.5
Delta Dist. [15]	DDRNet-23s	W8A8*	1.1	76.2
ResQ	DDRNet-23s	W8A8→W4A4	1.1	77.3
ResQ	DDRNet-23s	W4A4→W3A3	0.4	76.2

Table 2: **Comparison with state-of-the-art on Cityscapes.** ResQ outperforms alternatives with similar or less BOPs. *Models are assumed lossless in performance when quantized to 8 bits.




Figure 7: **Frame vs. Residual Quantization on Cityscapes.** The base model is DDRNet-23 slim. Points are generated for $T = 2, 3, 4, 5$.




Figure 8: **Qualitative comparison** between ResQ and Frame Quantization. ResQ takes advantage of information from the previous keyframe to improve the segmentation of the scene.

annotations required to train video models by applying an off-the-shelf segmentation network [56] on unannotated frames in the training set. We evaluate the accuracy using mean Intersection-over-Union (mIoU) and perform inference by selecting the keyframe at intervals of $T = 3$ frames. To deal with sparse labels in the validation set, we follow the protocol in [21, 15] and average evaluations for all clips in which each labeled frame falls in position $t \leq T$.

Implementation details. We conduct our experiments on DDRNet [20] and HRNet-w18-small [59] backbones, and perform experiments with both PTQ and QAT quantizers. When using QAT, we fine-tune the model quantized with PTQ using the original training protocol [20, 59], except for the learning rate, which is lowered to $5e - 4$. As such, we train with 4 GPUs, a batchsize of 12 using SyncBN for 484 epochs, and we employ online hard example mining [52].

Comparison of quantization schemes Fig. 7 compares Frame Quantization and ResQ on DDRNet-23 slim, for two different bit-width configurations ($W8A8 \rightarrow W4A8$ and $W8A8 \rightarrow W4A4$). Consistent with the previous results on human pose estimation, we observe that ResQ outperforms Frame Quantization. Furthermore Dynamic-ResQ ($W8A8 \rightarrow W4A\{0,4,8\}$) outperforms ResQ and largely recovers the catastrophic degradation hampering

Frame Quantization in low bit-width regimes (W4A4).

Fig. 8 qualitatively compares Frame Quantization and ResQ at 4, 3, and 2 bits. As seen on the highlighted regions, our proposal takes advantage of the high precision representation for the keyframe to better segment the scene.

Comparison with the state-of-the-art We compare ResQ with prior solutions that are devised to accelerate video inference: TDNet [21], DFF [68], Skip Convolutions [14] and Delta Distillation [15]. Again, we assume lossless quantization to 8 bits for competing methods and compare against their upper bound accuracies. ResQ relies on QAT and channel quantization. As shown in Tab. 2, our proposal achieves a better accuracy vs. efficiency trade-off against existing models. More specifically, ResQ improves over Delta Distillation by 0.9, 0.4, and 1.1 points in mIoU for DDRNet-39, DDRNet-23, DDRNet-23s, respectively. Interestingly, we notice our model improves in mIoU over the floating point model we are optimizing. We ascribe this behavior to the fact that, by propagating keyframe representations to future timesteps, with ResQ the model successfully learns to leverage the temporal context.

Temporal stability. Temporal degradation can be a concern in keyframe based models. To assess this aspect, we consider ResQ on DDRNet23s on Cityscapes, with QAT

Model	Bit-width	Single Object VOS (DAVIS-2016)			Multi Object VOS (DAVIS-2017)		
		\mathcal{J} -Mean \uparrow	\mathcal{F} -Mean \uparrow	BOPs (T) \downarrow	\mathcal{J} -Mean \uparrow	\mathcal{F} -Mean \uparrow	BOPs (T) \downarrow
STM [45]	W8A8	83.8	85.9	10.07	75.7	81.1	19.09
	W4A8	00.1	00.5	5.04	01.8	03.5	9.54
	W8A4	00.9	06.0	5.04	03.0	11.6	9.54
+ResQ	W8A8 \rightarrow W4A8	73.4	76.0	6.74	67.3	72.8	12.78
	W8A8 \rightarrow W8A4	70.2	71.9	6.74	65.3	70.5	12.78
+Dynamic-ResQ	W8A8 \rightarrow W8A[0,4,8]	82.8	83.8	6.51	75.1	79.3	11.33

Table 3: **VOS results** on the DAVIS 2016 and 2017 benchmarks. The different BOP counts for the two splits can be ascribed to the multi-object nature of DAVIS-2017, requiring multiple joint encodings of frames and masks [45]. Low-bit quantization disrupts the baseline, while ResQ and Dynamic-ResQ retain segmentation capabilities at lower precision.

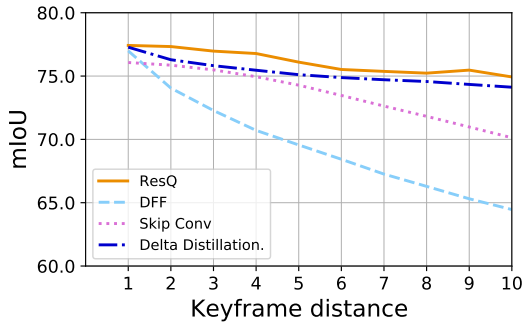


Figure 9: **Temporal stability.** We measure the impact of using residuals for up to 10 frames and compare against similar video efficiency models based on keyframe and residual processing.

precision W8A8 \rightarrow W4A4. Fig. 9 plots its segmentation mIoU against the distance to the last keyframe, along with results from other video efficiency models [68, 14, 15]. As shown, our model proves more robust than competing methods, with improvements ranging from 0.5 to 1 point.

5.3. Video Object Segmentation

Dataset, base model and metric. We conduct additional experiments with Video Object Segmentation (VOS) on the DAVIS dataset. More specifically, we experiment with semi-supervised VOS, where objects are segmented on the first frame of a clip and their masks are propagated to successive frames. As a base model, we choose Space-Time Memory Networks (STM) [45]. We use the released weights, obtained by the authors by training on Youtube-VOS [61], and we quantize the encoder and decoder modules with PTQ, whilst leaving the memory operations at full precision. We test our model on the validation sets of DAVIS-2016 [47], composed of 20 clips each with a single target object, as well as DAVIS-2017 [48], where multiple objects are targeted in 30 videos. We rely on the official evaluation SDK and report performance in terms of region (\mathcal{J}) and boundary (\mathcal{F}) accuracy (as defined in [47]).

Results. Performance of both Frame Quantization and ResQ ($T = 3$) on the 2016 and 2017 splits are reported in Tab. 3. We found the model hard to quantize, possibly due to quantization noise interfering with its memory operations. Indeed, quantization to 8 bits is already lossy w.r.t. the floating point model (4.9 and 3.5 drop in \mathcal{J} -Mean for 16 and 17 benchmarks respectively). As the Table shows, when quantizing to 4 bits, Frame Quantization demans the model, whereas ResQ and Dynamic-ResQ retain performance. Remarkably, Dynamic-ResQ suffers minor drops in \mathcal{J} -Mean of 1 point and 0.6 points for the single and multi-object benchmark respectively, allowing to save approximately 35 to 40% BOPs.

6. Conclusion

We motivated - and supported with theoretical insights - that whenever running quantized models on video streams a lower quantization error can be obtained by using residuals of frame representations. We have therefore proposed ResQ, a technique for amortizing expensive high precision computations on a keyframe over several residual frames, by processing the latter with low precision. We further introduced a dynamic extension, where precision for each residual pixel is adjusted on the fly, based on content of the residual itself. Through several experiments, we have empirically validated that ResQ can obtain favorable results in terms of accuracy and efficiency, outperforming existing methods in this trade-off in three different perception tasks.

Limitations. We highlight some potential limitations of our model to be tackled by future work. One downside to ResQ is that it requires the propagation of representations to future timesteps, leading to a memory overhead potentially impacting latency in memory-bound applications. Moreover, implementing location-specific quantized operations is not trivial and requires specialized hardware or gather-scatter implementations of convolutions, similar to what previously used for sparse processing in [50, 58]. Finally, we note how ResQ is able to reduce the amortized cost of video processing, yet the peak BOPs is not reduced.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2014. 5
- [2] Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. In *Neural Information Processing Systems*, 2019. 1, 2
- [3] Chaim Baskin, Natan Liss, Eli Schwartz, Evgenii Zheltonozhskii, Raja Giryes, Alex M Bronstein, and Avi Mendelson. UniQ: Uniform noise injection for non-uniform quantization of neural networks. *TCOS*, 2021. 5
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2, 5
- [5] Wuyang Chen, Xinyu Gong, Xianming Liu, Qian Zhang, Yuan Li, and Zhangyang Wang. Fasterseg: Searching for faster real-time semantic segmentation. *International Conference on Learning Representations*, 2020. 2
- [6] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. PACT: parameterized clipping activation for quantized neural networks. *arXiv preprint arxiv:805.06085*, 2018. 2
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016. 6
- [8] Misha Denil, Babak Shakibi, Laurent Dinh, Marc’Aurelio Ranzato, and Nando de Freitas. Predicting parameters in deep learning. In *Neural Information Processing Systems*, 2013. 2
- [9] Matthew Dutson and Mohit Gupta. Event neural networks. *Proceedings of the European Conference on Computer Vision*, 2022. 1, 2, 3
- [10] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *International Conference on Learning Representations*, 2020. 1, 2
- [11] Alexander Finkelstein, Uri Almog, and Mark Grobman. Fighting quantization bias with bias. *arXiv preprint arxiv:1906.03193*, 2019. 2
- [12] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. *IEEE International Conference on Computer Vision*, 2019. 2
- [13] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, 2015. 2
- [14] Amirhossein Habibian, Davide Abati, Taco S Cohen, and Babak Ehteshami Bejnordi. Skip-convolutions for efficient video processing. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2021. 1, 2, 3, 5, 6, 7, 8
- [15] Amirhossein Habibian, Haitam Ben Yahia, Davide Abati, Efstratios Gavves, and Fatih Porikli. Delta distillation for efficient video processing. *Proceedings of the European Conference on Computer Vision*, 2022. 1, 2, 3, 6, 7, 8
- [16] Tiantian Han, Dong Li, Ji Liu, Lu Tian, and Yi Shan. Improving low-precision network quantization via bin regularization. In *IEEE International Conference on Computer Vision*, 2021. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016. 6
- [18] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *IEEE International Conference on Computer Vision*, 2017. 2
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [20] Yuanduo Hong, Huihui Pan, Weichao Sun, Yisong Jia, et al. Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes. *arXiv preprint arXiv:2101.06085*, 2021. 6, 7
- [21] Ping Hu, Fabian Caba, Oliver Wang, Zhe Lin, Stan Sclaroff, and Federico Perazzi. Temporally distributed networks for fast video semantic segmentation. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2020. 1, 2, 5, 6, 7
- [22] B. Ham J. Lee, D. Kim. Network quantization with element-wise gradient scaling. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2021. 2
- [23] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [24] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. *British Machine Vision Conference*, 2014. 5, 6
- [25] Samvit Jain, Xin Wang, and Joseph E Gonzalez. Accel: A corrective fusion network for efficient semantic segmentation on video. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019. 1, 2
- [26] Sambhav R. Jain, Albert Gural, Michael Wu, and Chris Dick. Trained uniform quantization for accurate and efficient neural network inference on fixed-point hardware. *arxiv preprint arxiv:1903.08066*, 2019. 2
- [27] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *IEEE International Conference on Computer Vision*, 2013. 5
- [28] Jangho Kim, KiYoon Yoo, and Nojun Kwak. Position-based scaled gradient for model quantization and pruning. In *Neural Information Processing Systems*, 2020. 2

- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014. 5
- [30] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. 1, 2
- [31] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *International Conference on Learning Representations*, 2017. 2
- [32] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brcq: Pushing the limit of post-training quantization by block reconstruction. In *International Conference on Learning Representations*, 2021. 2
- [33] Yule Li, Jianping Shi, and Dahua Lin. Low-latency video semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 2
- [34] Mason Liu and Menglong Zhu. Mobile video object detection with temporally-aware feature maps. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 2
- [35] Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through L0 regularization. In *International Conference on Learning Representations*, 2018. 5, 6
- [36] Yue Luo, Jimmy Ren, Zhouxia Wang, Wenxiu Sun, Jinshan Pan, Jianbo Liu, Jiahao Pang, and Liang Lin. Lstm pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 5
- [37] Eldad Meller, Alexander Finkelstein, Uri Almog, and Mark Grobman. Same, same but different: Recovering neural network quantization error through weight factorization. In *International Conference on Machine Learning*, 2019. 2
- [38] Bert Moons, Parham Noorzad, Andrii Skliar, Giovanni Mariani, Dushyant Mehta, Chris Lott, and Tijmen Blankevoort. Distilling optimal neural networks: Rapid search in diverse spaces. In *IEEE International Conference on Computer Vision*, 2021. 2
- [39] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? Adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, 2020. 1, 2
- [40] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021. 2
- [41] Markus Nagel, Marios Fournarakis, Yelysei Bondarenko, and Tijmen Blankevoort. Overcoming oscillations in quantization-aware training. In *International Conference on Machine Learning*, 2022. 2
- [42] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization through weight equalization and bias correction. *IEEE International Conference on Computer Vision*, 2019. 1, 2
- [43] Xuecheng Nie, Yuncheng Li, Linjie Luo, Ning Zhang, and Jiashi Feng. Dynamic kernel distillation for efficient pose estimation in videos. In *IEEE International Conference on Computer Vision*, 2019. 5, 6
- [44] Peter O'Connor and Max Welling. Sigma delta quantized networks. *International Conference on Learning Representations*, 2017. 3
- [45] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *IEEE International Conference on Computer Vision*, 2019. 8
- [46] Mathias Parger, Chengcheng Tang, Christopher D. Twigg, Cem Keskin, Robert Wang, and Markus Steinberger. Deltacnn: End-to-end cnn inference of sparse frame differences in videos. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 3, 5
- [47] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016. 8
- [48] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 8
- [49] Sahand Rabbani. Proof that the difference of two jointly distributed normal random variables is normal. <https://srabbani.com/bivariate.pdf>, 2007. 3
- [50] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Urtasun. Sbnnet: Sparse blocks network for fast inference. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 8
- [51] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *International Conference on Learning Representations*, 2015. 2
- [52] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016. 7
- [53] Jie Song, Limin Wang, Luc Van Gool, and Otmar Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017. 5
- [54] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019. 5, 6
- [55] Ximeng Sun, Rameswar Panda, Chun-Fu Richard Chen, Aude Oliva, Rogerio Feris, and Kate Saenko. Dynamic network quantization for efficient video inference. In *IEEE International Conference on Computer Vision*, 2021. 2
- [56] Andrew Tao, Karan Sapra, and Bryan Catanzaro. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*, 2020. 7
- [57] Mart Van Baalen, Christos Louizos, Markus Nagel, Rana Ali Amjad, Ying Wang, Tijmen Blankevoort, and Max Welling. Bayesian bits: Unifying quantization and pruning. *Neural Information Processing Systems*, 2020. 5

- [58] Thomas Verelst and Tinne Tuytelaars. Dynamic convolutions: Exploiting spatial sparsity for faster inference. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2020. 8
- [59] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 3, 6, 7
- [60] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *European Conference on Computer Vision (ECCV)*, 2018. 5
- [61] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European Conference on Computer Vision*, 2018. 8
- [62] William Henry Young. On the multiplication of successions of fourier constants. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 87(596):331–339, 1912. 4, 5
- [63] Changqian Yu, Changxin Gao, Jingbo Wang, Gang Yu, Chunhua Shen, and Nong Sang. Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation. *IJCV*, 2021. 6
- [64] Xiangyu Zhang, Jianhua Zou, Kaiming He, and Jian Sun. Accelerating very deep convolutional networks for classification and detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015. 5, 6
- [65] X. Zhang, Jianhua Zou, Kaiming He, and Jian Sun. Accelerating very deep convolutional networks for classification and detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016. 2
- [66] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017. 6
- [67] Xizhou Zhu, Jifeng Dai, Xingchi Zhu, Yichen Wei, and Lu Yuan. Towards high performance video object detection for mobiles. *arXiv preprint arXiv:1804.05830*, 2018. 2
- [68] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017. 1, 2, 6, 7, 8