AdaRadar: Rate Adaptive Spectral Compression for Radar-based Perception

Jinho Park[1], Se Young Chun[2], Mingoo Seok[1]

[1] Columbia University, [2] Seoul National University

Sponsored by:

CVPR DENVER COLORADO JUNE 3-7, 2026

Project page

Radar Complements Camera & LiDAR

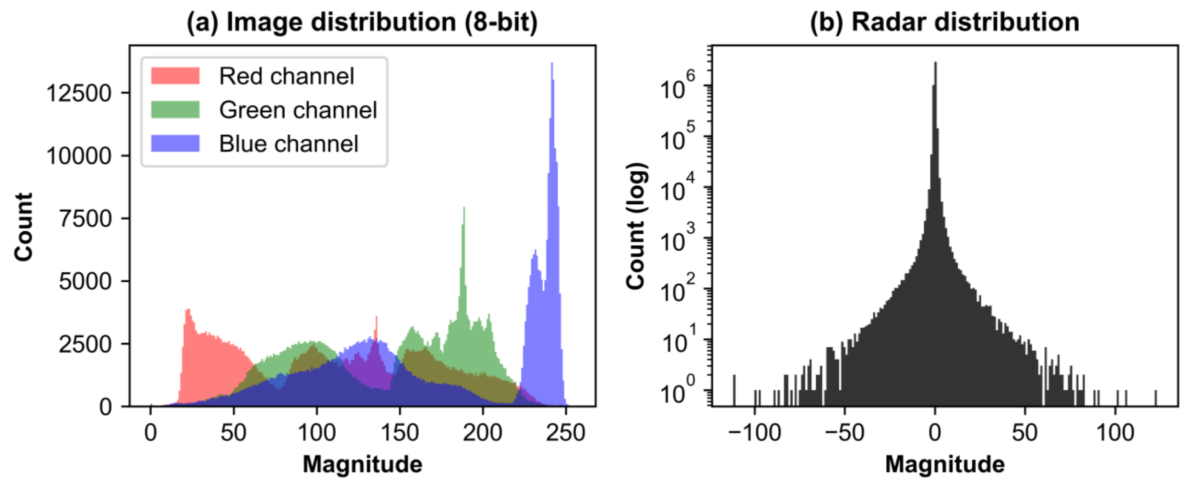

- **Radar perception** is a critical modality in autonomous driving, complementing camera & LiDAR.
- **Natural images** span a broad dynamic range.
- Radar signals exhibit extreme **sparsity** and a heavy-tailed profile (1st/99th percentiles at ±1.74).

Adaptive Compression Rate Control

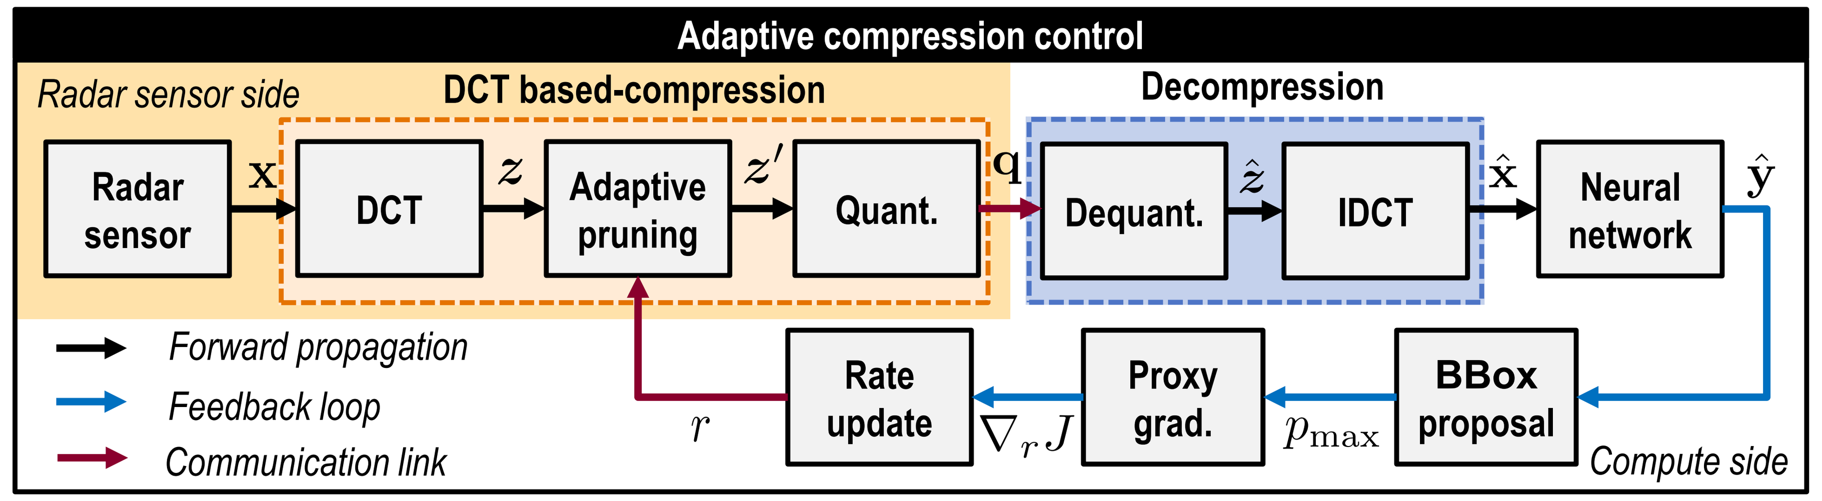

- Our proposed method introduces a **feedback loop** in which the proxy gradient is computed from the detection outputs to update the compression ratio adaptively.
- This avoids the need for **backpropagation** through the communication channel.
- DCT-based compression on the radar sensor side and decompression on the compute side.

Motivation for Compression & Adaptation

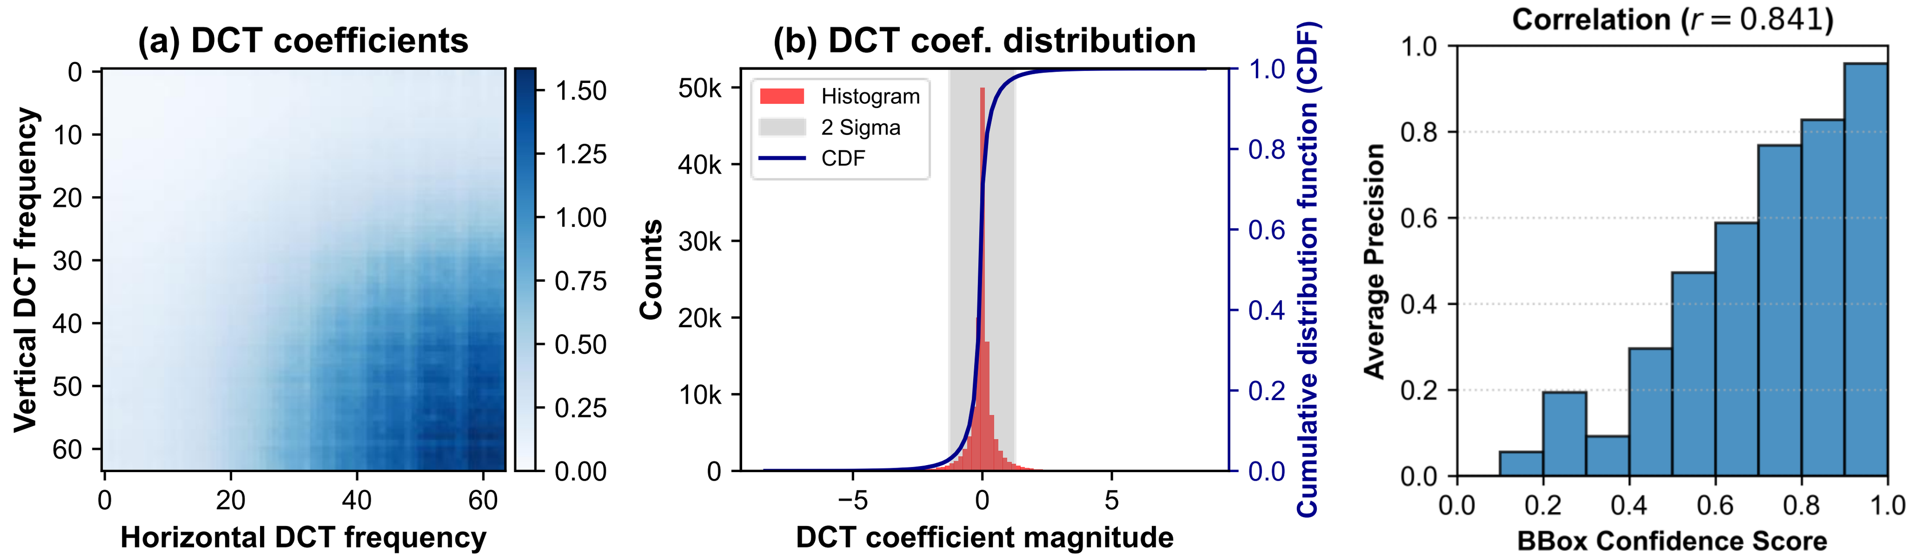

- The DCT coefficient **magnitudes** are clustered in the high-frequency bins.
- Their histogram is sharply peaked—strong **sparsity** and clear opportunities for compression.
- BBox confidence shows great correlation with the ground truth label.

Feedback with Proxy Gradient

Task objective:

$$\max_{\{r_t\}_{t=1}^{T}} \mathbb{E}[J_t(r_t)], \quad J_t(r_t) = h(\mathbf{x}_t, r_t) - \lambda \cdot B(r_t)$$

hyper-parameter to control accuracy-bandwidth trade-off

Instantaneous bit-rate

Gradient w.r.t bit-rate (r):

$$\nabla_r J = \nabla_r h(\mathbf{x}, r) - \lambda \cdot B'(r)$$

Proxy gradient:

$$\hat{\nabla}_r h(\mathbf{x}, r) \approx \frac{\underbrace{h(\mathbf{x}, r)}_{\text{Pass \#1}} - \underbrace{h(\mathbf{x}, r-\epsilon)}_{\text{Pass \#2}}}{\epsilon} = \frac{p - p^-}{\epsilon}$$

Algorithm 1: Adaptive rate control

Input: Radar sequence $(\{\mathbf{x}_t\}_{t=1}^{T})$, weight $(\boldsymbol{\Theta})$, initial pruning ratio (r_1), learning rate (η), perturbation (ϵ)

Output: Detection map $(\{\hat{\mathbf{y}}_t\}_{t=1}^{T})$

for $t := 1$ *to* T **do**
- $\mathbf{q}_t, \mathbf{Q}_t \leftarrow \text{Compression}(\mathbf{x}_t, r_t)$;
  - `// Transfer` $\{\mathbf{q}, \mathbf{Q}\}$ `once`
- $\hat{\mathbf{x}}_t \leftarrow \text{Decompression}(\mathbf{q}_t, \mathbf{Q}_t)$;
- $\hat{\mathbf{y}}_t \leftarrow f_{\boldsymbol{\Theta}}(\hat{\mathbf{x}}_t)$; ← Network inference pass #1
- $\{(b_k, p_k)\}_{k=1}^{K} = \text{Propose}(\hat{\mathbf{y}}_t)$;
- **if** *adaptive* **then**
  - $\mathbf{q}_t^- \leftarrow \text{Pruning}(\mathbf{q}_t, \epsilon)$;
  - $\hat{\mathbf{x}}_t^- \leftarrow \text{Decompression}(\mathbf{q}_t^-, \mathbf{Q}_t)$;
  - $\hat{\mathbf{y}}_t^- \leftarrow f_{\boldsymbol{\Theta}}(\hat{\mathbf{x}}_t^-)$; ← Network inference pass #2
  - **if** $k > 0$ **then**
    - $r_{t+1} \leftarrow r_t - \eta \hat{\nabla}_{r_t} J$;

- Task objective: h encapsulates compression, decompression, forward inference, and BBox proposal.

Rate-accuracy Trade-off

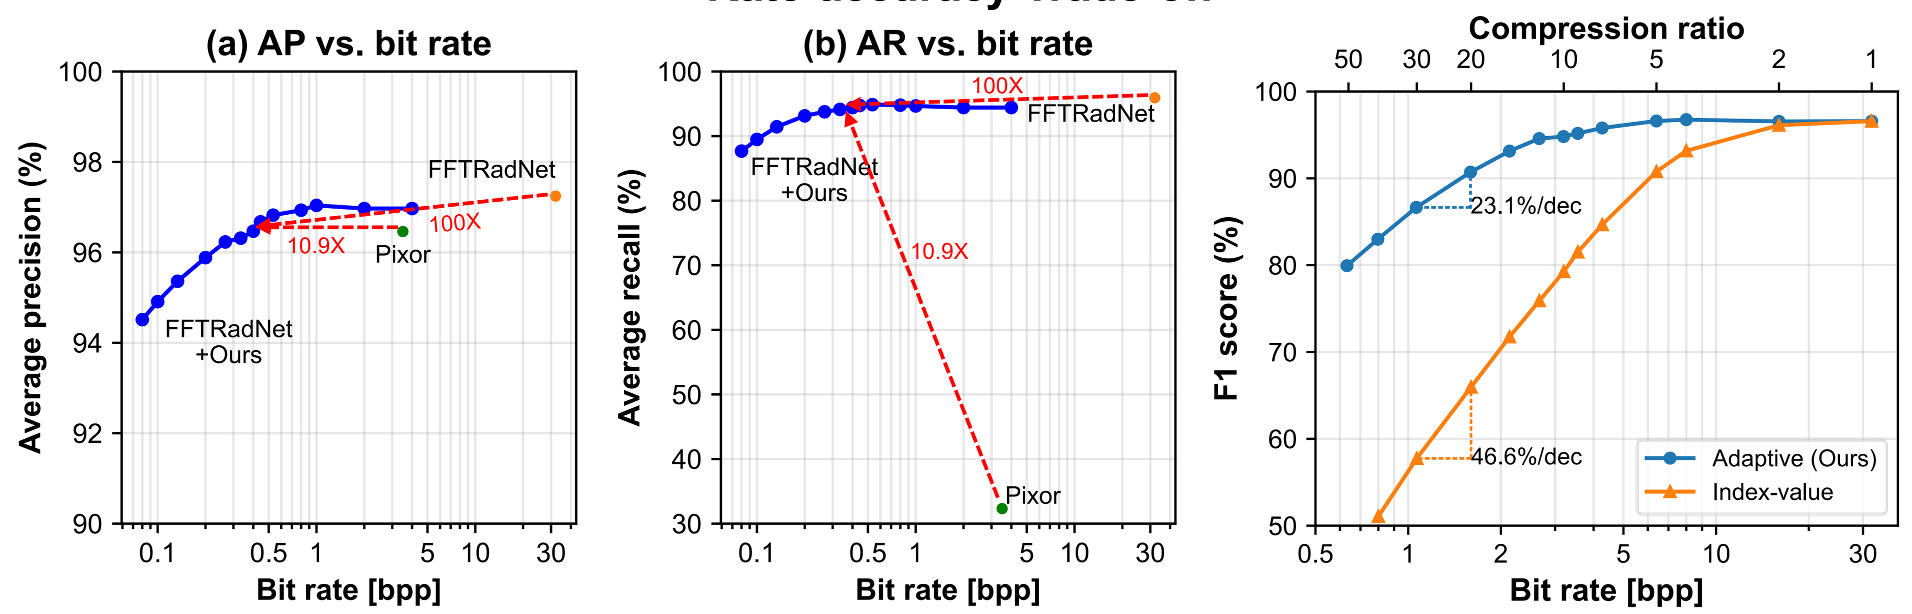

- Our compression scheme achieves **100x** reduction in the radar feature map size while having a **1%p decrease** in the performance compared to the baseline.
- Our method remains **stable** until 5x compression whereas the performance drop-off begins immediately with the index-value pair-based method.
- The roll-off gradient of **23.1%/dec** is much more gentle for the spectral pruning-based compression vs. **46.6%/dec** of the spatial-domain counterpart.

Compression Results – RADIal Dataset

Method	Bit	Prune ratio	Bit rate [bpp] ↓ (Comp. ratio ↑)	Detection P (%) ↑	Detection R (%) ↑	Detection F_1 (%) ↑	Segment. mIoU (%) ↑
Baseline [2]	32	-	32 (1×)	97.24	95.93	96.58	75.97
+Index-value [11]	32	12×	2.67 (12×)	**97.55**	62.12	75.91	49.86
+Adaptive (Ours)	4	12.57×	**0.32 (101×)**	96.25	**94.04**	**95.13**	**79.34**

- An average pruning ratio of **12.57x** combined with **8x** gain from quantization, yields a bit rate of **0.32 bits per pixel** (bpp), corresponding to a **101x** compression ratio.
- On the segmentation task, ours improves the performance by **4%**.
- Our method markedly **surpasses** index-value based method.

Qualitative Results

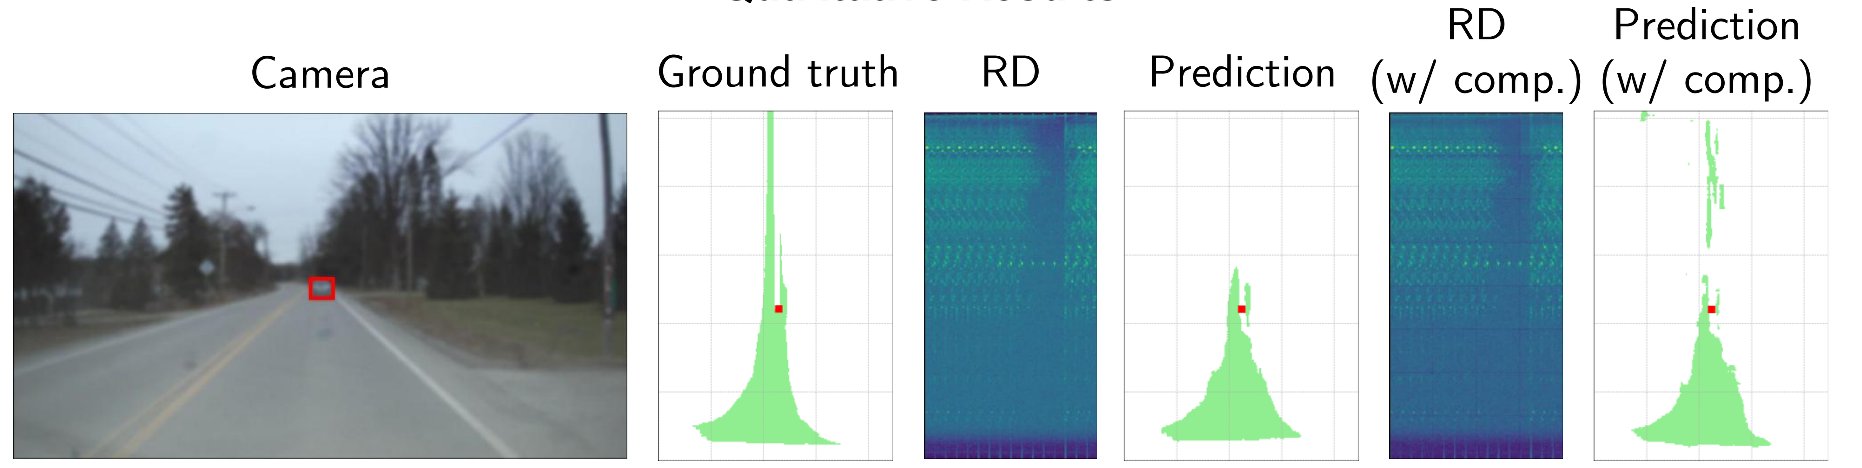

- **Camera images** provide contextual reference for the scene (not fed into network).
- The network predicts both **detection** and **segmentation** outputs.

Time Series Analysis

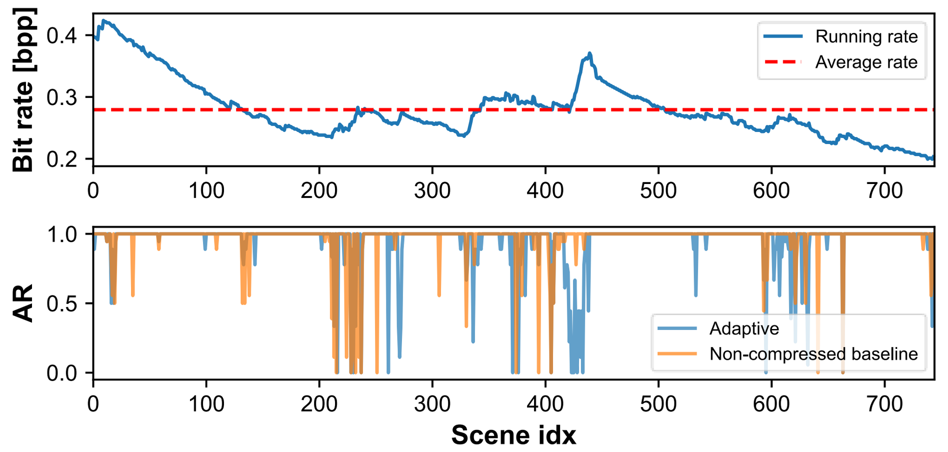

- We **visualize** the online compression with bit rate and detection metric.
- We expect that the controller decreases the pruning ratio to compensate for the AR drop.
- It achieves an average bit rate of **0.279 bpp** with 8-bit quantization, yielding a **115x** compression with AR of **93.91%**.

Effect of Quantization Bit Width & DCT Block Length

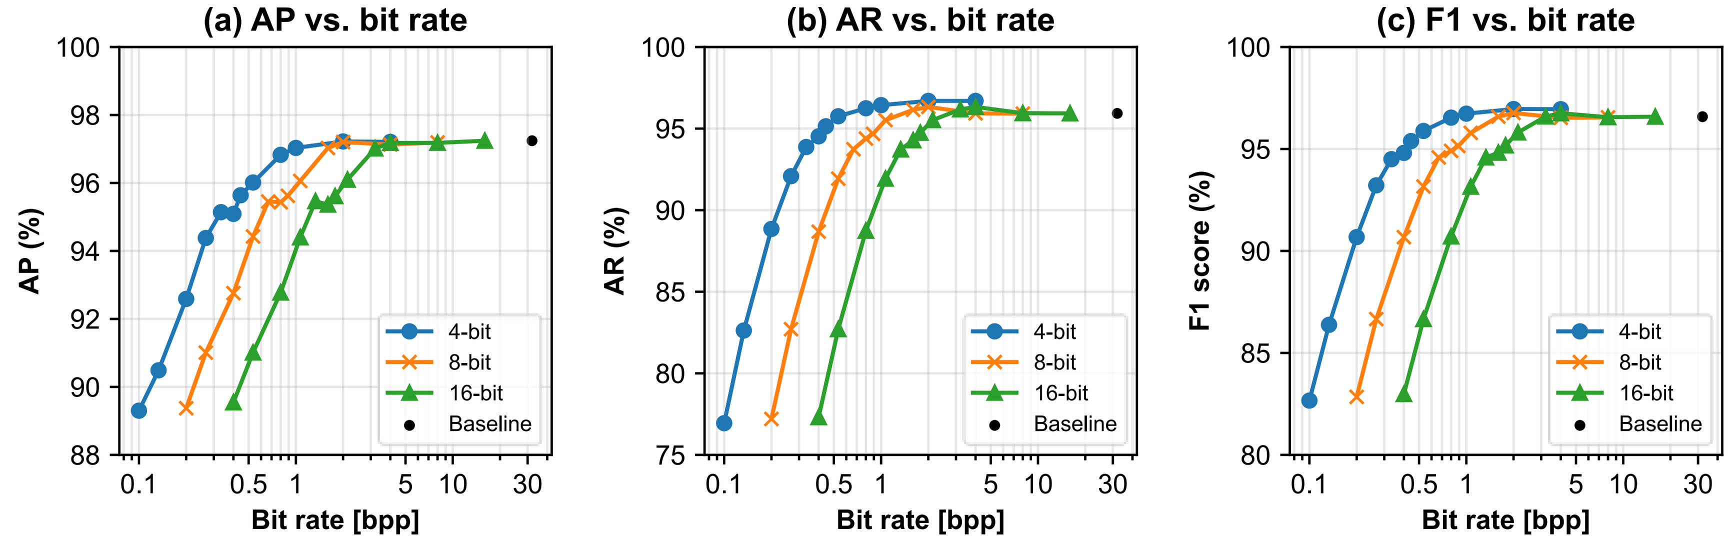

- Quantization up to **4-bit** does not affect the performance compared to that of 8-bit and 16-bit. (M = 64)

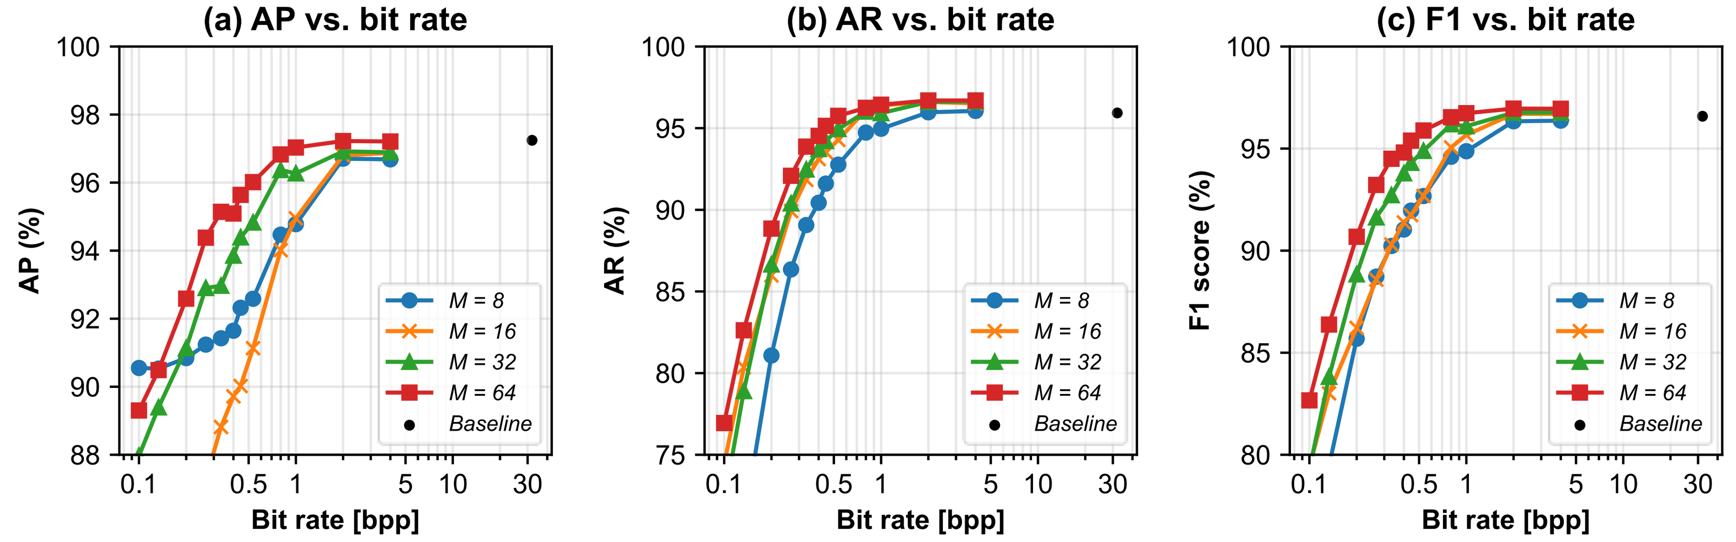

- For the block-size sweep, M **= 64** yields the best trade-off between performance and overhead from the scaling factor. (bit width = 4-bit)

Acknowledgements: This work was supported in part by COGNISENSE, one of seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The work of SY Chun was supported by IITP grants funded by the Korea government(MSIT) [No.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University) / No.RS-2025-02314125, Effective Human-Machine Teaming With Multimodal Hazy Oracle Models].