

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 COMBINE-ICMH: A DUAL-ADAPTER CO-TUNING FRAMEWORK IN IMAGE COMPRESSION FOR MACHINE AND HUMAN VISION

006
007 **Anonymous authors**
008 Paper under double-blind review

011 ABSTRACT

013 To reduce the high training overhead of models for Image Compression for Ma-
014 chine and Human Vision (ICMH), the paradigm of fine-tuning pre-trained mod-
015 els has gained increasing attention. Among these, lightweight adapter-based
016 approaches have emerged as efficient solutions. However, we argue that this
017 paradigm suffers from two critical, yet overlooked flaws. First, existing frequency-
018 domain adapters lack adaptability, often suppressing high-frequency details cru-
019 cial for machine tasks. Second, fine-tuning the transform module alone introduces
020 a “transform-entropy mismatch,” as the frozen entropy model cannot adapt to the
021 altered latent distribution. To address these challenges, we propose Combine-
022 ICMH, a novel framework that enables the synergistic co-optimization of both the
023 transform and entropy models. Specifically, we design a Spatial-Wavelet Modula-
024 tion Adapter (SWMA) to enhance frequency adaptability and introduce a Channel
025 Modulation Adapter (CMA) to directly fine-tune the entropy model, resolving the
026 mismatch. Extensive experiments demonstrate that our method consistently out-
027 performs state-of-the-art approaches on various downstream tasks, including clas-
028 sification, detection, and segmentation, while maintaining comparable parameter
029 efficiency.

031 1 INTRODUCTION

033 The surge of visual data from applications like autonomous driving demands efficient compression
034 for analysis under limited bandwidth. While Image Compression for Machine (ICM) (Chen et al.,
035 2019) optimizes for machine tasks, its severe degradation of visual quality makes the results unsuit-
036 able for human oversight. Therefore, Image Compression for Machine and Human Vision (ICMH)
037 (Wang et al., 2021; Lin et al., 2023) has become a critical research topic.

038 Currently, the dominant paradigm for achieving ICMH has shifted from designing multi-task
039 pipelines (Bai et al., 2022; Codevilla et al., 2021) to a tuning framework that adapts a pre-trained im-
040 age codec optimized for human vision (Liu et al., 2023a; 2022). The former, though capable of par-
041 allel processing for both machine and human vision, suffers from high training overhead, redundant
042 task-specific decoders, and degraded R-D performance, making large-scale deployment impractical
043 (Li et al., 2024). In contrast, this fine-tuning paradigm is built upon a shared codec backbone, inte-
044 grating lightweight, task-specific plug-and-play modules that add only a few trainable parameters.
045 This design achieves dynamic switching between human and machine vision applications through
046 a highly efficient, plug-and-play adaptation, all while crucially preserving the full reconstruction
047 fidelity of the pre-trained model.

048 However, existing fine-tuning methods still exhibit significant limitations. Some works, such as
049 ICMH-Net (Liu et al., 2023b) and TransTIC (Chen et al., 2023), employ masking or prompting
050 (Lester et al., 2021; Li & Liang, 2021) for adaptation, but these techniques are tightly coupled with
051 specific architectures (e.g., Transformers (Liu et al., 2021)) and incur substantial computational over-
052 head. Another fine-tuning paradigm based on lightweight adapters, represented by Adapt-ICMH (Li
053 et al., 2024), better balances complexity and performance by inserting lightweight adapters into the
transform module.

Nevertheless, we identify two overlooked flaws in current adapter-based fine-tuning frameworks (Li et al., 2024; Zhao et al., 2025). First, they lack sufficient adaptability in the frequency domain. The Fourier-based frequency modulation adapter in Adapt-ICMH (Li et al., 2024) not only tends to suppress high-frequency components crucial for detection and segmentation, but also provides only limited enhancement for low-frequency components that carry critical global context for classification. The second, more critical flaw is the “transform–entropy mismatch”. Current methods introduce adapters solely into the transform network (e.g., encoder or decoder), altering the latent representations. Meanwhile, the frozen entropy model is optimized for the original human-vision task and cannot accurately estimate the distribution of these altered latent representations. This misalignment leads to suboptimal bitrate estimation and degraded performance.

To address these critical challenges, we propose two targeted innovations. First, to combat insufficient frequency adaptability, we design the novel Spatial-Wavelet Modulation Adapter (SWMA). This adapter eliminates spatial redundancy through a dedicated spatial branch, while its parallel wavelet branch surpasses the limitations of Fourier-based methods by simultaneously enhancing crucial low-frequency components and preserving high-frequency ones. More critically, to resolve the “transform–entropy mismatch,” we propose the Channel Modulation Adapter (CMA). By directly fine-tuning the entropy model, CMA realigns its probability estimates with the altered latent distribution, thereby addressing a core limitation of existing adapter-based methods.

Building upon these two innovations, we propose Combine-ICMH, an ICMH framework that co-optimizes the feature transformation and its entropy estimation. As shown in Figure 1, we insert SWMA into the base codec’s transform network and integrate CMA into its entropy model. This synergistic fine-tuning strategy effectively addresses the limitations of existing methods. Our main contributions are as follows:

- We propose the Spatial-Wavelet Modulation Adapter (SWMA), which addresses the insufficient frequency adaptability in prior work. It processes spatial and frequency information in parallel, leading to superior fine-tuning performance.
- We design the lightweight Channel Modulation Adapter (CMA), which resolves the “transform–entropy mismatch” in existing frameworks by aligning the entropy model with the shifted latent feature distribution.
- Based on the proposed adapters, we develop Combine-ICMH, an ICMH framework that synergistically fine-tunes both the transform and entropy models. Experimental results demonstrate that our method consistently outperforms existing ICMH frameworks across various machine vision tasks while maintaining a comparable level of complexity.

2 RELATED WORKS

2.1 LEARNED IMAGE COMPRESSION

Learned image compression (LIC) has advanced significantly since Ballé et al. (2017) introduced the first end-to-end image compression framework. Nowadays, state-of-the-art (SOTA) LIC models outperform the best traditional codecs, such as VVC (Bross et al., 2021) and HEVC (Sullivan et al., 2012). Progress in this field has primarily focused on two core components: the transform module and the entropy model.

In the domain of transform module, Cheng et al. (2019) improved compression by stacking residual blocks, while Ghorbel et al. (2023) employed ConvNeXt modules to balance performance and complexity. In addition, attention- and Transformer-based methods have been explored: Zou et al. (2022) introduced window-based local attention, and Zhu et al. (2022) built a pure Swin Transformer architecture.

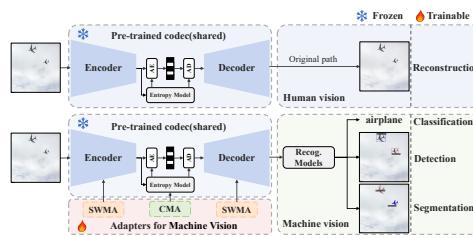


Figure 1: The framework of our Combine-ICMH. Both human and machine vision tasks share a pre-trained codec. By disabling the inserted adapters, the proposed framework could convert to base codec for human vision tasks.

108 In parallel, improving the entropy model is essential for efficient coding. For instance, Ballé et al.
 109 (2018) introduced the hyperprior network, leveraging side information to capture spatial correlations.
 110 Building on this, Minnen et al. (2018) proposed an autoregressive model. To further enhance coding efficiency, He et al. (2021) and Jiang & Wang (2023) fused spatial–channel and multi-
 111 reference contexts, respectively.

113 However, LIC models, optimized for human vision, are suboptimal for machine tasks. Existing
 114 fine-tuning approaches generally adjust only the transform module, overlooking its tight coupling
 115 with the entropy model. This oversight causes a “transform–entropy mismatch,” which results in
 116 inefficient rate estimation and degraded performance.

117

118 2.2 IMAGE COMPRESSION FOR MACHINE AND HUMAN VISION

119

120 To address the conflict between machine vision (Ding et al., 2022; Sun et al., 2019; Xie et al.,
 121 2021a) and human perception in image coding, Image Compression for Machine and Human Vision
 122 (ICMH) has emerged.

123

124 The primary objective of ICMH is to reconcile human perception with machine analysis. For in-
 125 stance, Choi & Bajić (2022) divided the latent representation into a “base layer” and an “enhance-
 126 ment layer,” which are transmitted via a scalable bitstream to decouple machine vision features from
 127 human vision information. Another work (Liu et al., 2022) introduced a framework that performs
 128 machine analysis directly on compressed representations. While such designs improve perceptual
 129 quality over earlier ICM approaches (Chamain et al., 2021; Feng et al., 2022; Le et al., 2021), they
 130 suffer from substantial training and storage overhead, limiting their practical utility (Li et al., 2024).

131

132 Recently, fine-tuning frameworks that adapt pre-trained image codecs (Fischer et al., 2022; Chen
 133 et al., 2023) have gained more attention. Notably, Chen et al. (2023) proposed a prompt-based tun-
 134 ing framework. These approaches leverage pre-trained models to reduce training overhead without
 135 compromising the human perceptual quality of the pretrained model. However, such methods are
 136 often restricted to specific architectures and suffer from low fine-tuning efficiency. Furthermore,
 137 Li et al. (2024) extended their applicability, proposing a fine-tuning paradigm based on lightweight
 138 adapters.

139

140 3 PROPOSED METHOD

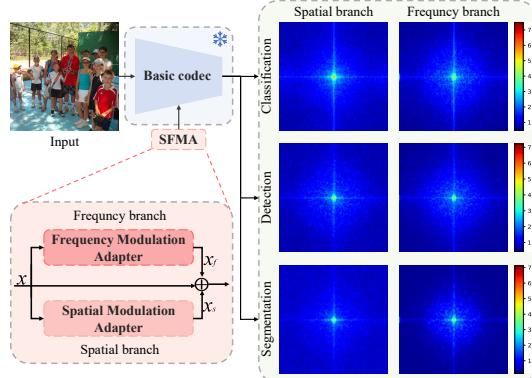
141

142 3.1 EMPIRICAL ANALYSIS AND MOTIVATIONS

143

144 To efficiently adapt human-oriented LIC models to machine vision tasks without compromising R–D
 145 performance of the pretrained models, recent works (Li et al., 2024; Zhao et al., 2025) have adopted
 146 fine-tuning paradigms based on lightweight adapters. The core idea of this paradigm involves fine-
 147 tuning lightweight, plug-and-play adapters while freezing the base codec. Although this paradigm
 148 is promising, we identify two key limitations in current adapter designs.

149 First, we observe a critical limitation in the
 150 frequency adaptability of existing frequency
 151 modulation adapters. Different machine vi-
 152 sion tasks exhibit distinct preferences for fre-
 153 quency components: fine-grained tasks like in-
 154 stance segmentation rely on mid- and high-
 155 frequency components, whereas coarse-grained
 156 tasks like image classification focus more on
 157 low-frequency components. An ideal adapter
 158 should flexibly cater to these diverse de-
 159 mands. However, as shown in Fig. 2, ex-
 160 isting Fourier-based adapters excessively at-
 161 tenuate high-frequency components while fail-
 162 ing to sufficiently amplify low-frequency ones.
 163 This lack of frequency adaptivity exposes the
 164 shortcomings of current frequency-modulation
 165 strategies. To address these issues, we propose



166 Figure 2: Power spectral density (PSD) maps of
 167 intermediate features from the Li et al. (2024) for
 168 three downstream tasks.

162 a Spatial–Wavelet Modulation Adapter (SWMA) that leverages the multi-level wavelet transform to
 163 deliver more precise and effective feature modulation.

164 Beyond frequency adaptivity, we identify a crucial architectural issue in existing adapter-based fine-
 165 tuning frameworks: the “transform-entropy mismatch”. Learned image compression (LIC) com-
 166 prises two key components: nonlinear transform module and entropy model. Previous methods
 167 (Chen et al., 2023; Li et al., 2024) typically fine-tune only one component, usually the transform
 168 module, which introduces a substantial side effect. Specifically, fine-tuning the transform alone
 169 alters the latent distribution. Consequently, the entropy model, which was originally optimized for
 170 human vision, becomes unsuitable for the new latent representation tailored for machine vision. This
 171 mismatch results in inefficient bitrate estimation (e.g., increased bitrate) and degraded performance.

172 This inefficiency can be formally analyzed. In the hyperprior entropy model (Ballé et al., 2018),
 173 the rate of latent representation $R(\hat{y})$ is determined by the estimated conditional entropy. When the
 174 learned entropy model mismatches the actual distribution of the representation (i.e., the hyperprior \hat{z}
 175 deviates from the ideal \tilde{z}), the actual rate can be decomposed into the ideal rate and a penalty term,
 176 Δbpp . This decomposition is given by (see Appendix E for detailed derivation):

$$R(\hat{y}) = \mathbb{E} [-\log_2(p_{\hat{y}|\hat{z}}(\hat{y} | \hat{z}))] = \mathbb{E} [-\log_2(p_{\hat{y}|\hat{z}}(\hat{y} | \hat{z}))] + \Delta bpp \quad (1)$$

177 To validate our mismatch hypothesis and demonstrate the performance penalty it incurs, we con-
 178 ducted a staged freezing experiment. This experiment employed the TIC architecture (Lu et al.,
 179 2021), where we selectively fine-tuned layers of the 8-layer analysis/synthesis transforms (g_a, g_s)
 180 and the 4-layer hyperprior encoder/decoder (h_a, h_s). Table 1 specifies the indices of the unfrozen
 181 (i.e., trainable) layers for each configuration. Performance was evaluated using bits per pixel (bpp),
 182 mAP, and the latent correlation ρ as defined in Zhu et al. (2022).

183 Table 1: Analysis of a stage-wise freezing experiment on the TIC (Lu et al., 2021) for object detec-
 184 tion on COCO (Lin et al., 2014). In this experiment, specific network components were frozen to
 185 show the effects of fine-tuning only the transform modules (g_a, g_s) or fine-tuning them along with
 186 the entropy models (h_a, h_s).

g_a	g_s	h_a	h_s	bpp	mAP	ρ	Param(M)
0-7	0-7	0-3	0-3	0.0639	37.63	0.239	7.51
1,3,5	2,4,6	\	\	0.0829	37.21	0.289	3.22
0-7	0-7	\	\	0.0881	37.32	0.314	5.48
1,3,5	2,4,6	0-3	0-3	0.0685	37.42	0.260	5.25

190 As highlighted by rows 3 and 4 of Table 1, forgoing entropy model fine-tuning incurs a critical
 191 penalty. Despite employing more parameters, fine-tuning only the transform module simultaneously
 192 reduces task accuracy and compression efficiency (i.e., increases bitrate). We attribute this counter-
 193 intuitive result to the “transform-entropy mismatch,” where the frozen entropy model fails to adapt
 194 to the distribution of the latent representation in the new domain, leading to significant coding re-
 195 dundancy. The sharp increase in latent correlation ρ of the normalized latents further corroborates
 196 this conclusion.

197 To resolve this mismatch, we introduce the Channel Modulation Adapter (CMA), which modulates
 198 the entropy model via channel-wise adaptation to match shifted latent distributions. Finally, we
 199 combine our two specialized modules, the SWMA and the CMA, into a unified framework. This
 200 framework co-optimizes the transform and entropy modules, mitigating performance bottlenecks
 201 caused by distribution shifts and enabling efficient adaptation.

202 3.2 FRAMEWORK

203 To address the core problem of the mismatch between the transform and entropy models in existing
 204 adapter-based fine-tuning, we propose Combine-ICMH, a new framework that synergistically fine-
 205 tunes both modules through two efficient, plug-and-play adapters.

206 As illustrated in Figure 3, we insert SWMA into the transform modules to adapt them for down-
 207 stream machine vision tasks. Unlike previous Fourier-based adapters, SWMA leverages a multi-
 208 level wavelet transform to enable level-by-level processing of frequency components. SWMA thus

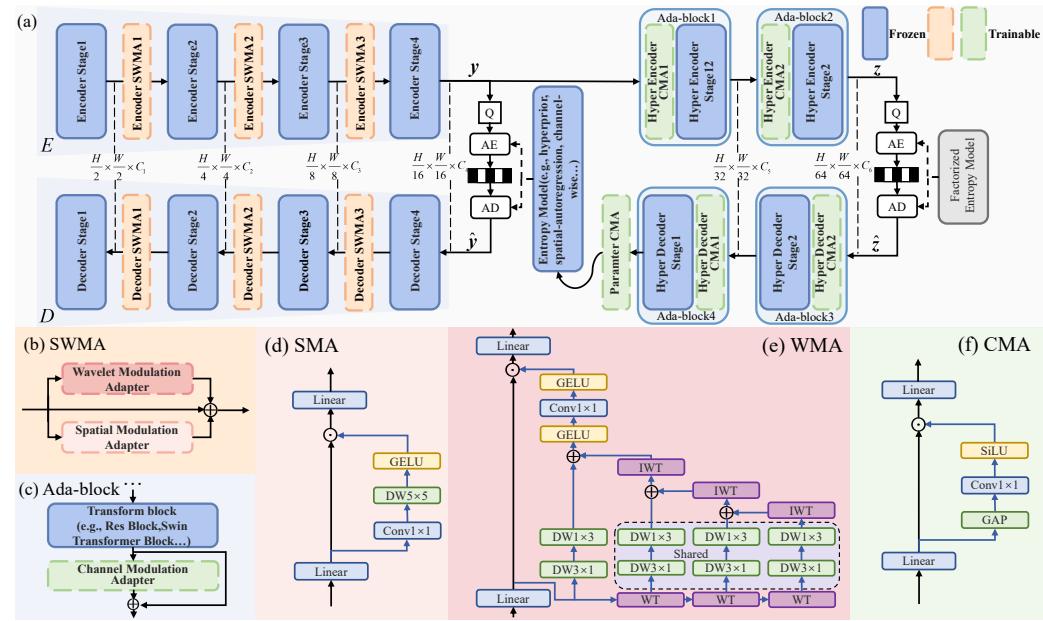


Figure 3: Overall architecture of Combine-ICMH. (a) Main pipeline. SWMA is integrated into the encoder (E) and decoder (D) of the base codec for feature adaptation, while CMA fine-tunes features from the hyperprior network to mitigate transform–entropy mismatch. (b) The proposed SWMA, comprising an SMA (d) and a WMA (e). (c) The structure of the Ada-block. (d) Architecture of the SMA. (e) Architecture of the WMA. (f) Architecture of the CMA.

enables a task-aware modulation of the frequency spectrum, selectively preserving high-frequency components (i.e., edges and boundaries) for precise localization in detection and segmentation, while enhancing low-frequency ones that capture global context for classification.

However, fine-tuning only the transform modules typically leads to the “transform–entropy mismatch”. To tackle this challenge, we further integrate the CMA into the hyperprior network (h_a, h_s). This adapter adjusts the entropy model via channel-wise modulation of the hyperprior, enabling it to precisely align with the latent distribution after fine-tuning.

In line with prior works (Chen et al., 2023; Li et al., 2024; Park et al., 2025), our strategy involves freezing the pre-trained codec and optimizing only the adapters, which significantly reduces training overhead while preserving the codec’s performance for human vision. Crucially, these adapters can be disabled at inference time, allowing our framework to losslessly revert to the base codec and maintain identical R-D performance to the baseline codec.

For a fair comparison with prior works, we adopt the loss function from (Chen et al., 2023):

$$\mathcal{L} = \mathcal{R} + \lambda \cdot \mathcal{D}_{task} \quad (2)$$

where \mathcal{R} denotes the bitrates estimated by the entropy model, \mathcal{D}_{task} is the task-specific perceptual distortion (detailed in Appendix A), and λ is the trade-off hyper-parameter.

3.3 SPATIAL-WAVELET MODULATION ADAPTER(SWMA)

The proposed Spatial-Wavelet Modulation Adapter (SWMA) adopts a parallel dual-branch structure to efficiently refine features. Its Spatial Modulation Adapter (SMA) branch removes spatial redundancy, while the Wavelet Modulation Adapter (WMA) branch modulates frequency components.

Spatial Modulation Adapter (SMA). The SMA, shown in Figure 3 (d), is designed to reduce spatial redundancy. To achieve this, an input feature x is first projected into the lower-dimensional feature x_s via a linear layer W_{down}^s . A spatial map (x_{smap}) is then generated from x_s through a 1×1 convolution and a 5×5 depth-wise convolution for a larger receptive field. Finally, the feature x_s is

270 modulated by the activated spatial map, and the result is up-projected to produce the output x_{sout} .
 271 The process is formulated as:

$$x_s = W_{down}^s(x) \quad (3)$$

$$x_{smap} = \text{DWConv5} \times 5(\text{Conv1} \times 1(x_s)) \quad (4)$$

$$x_{sout} = W_{up}^s(x_s \odot \text{GELU}(x_{smap})) \quad (5)$$

276 where W_{down}^s and W_{up}^s are linear projection layers, GELU denotes the Gaussian Error Linear Unit
 277 activation (Hendrycks & Gimpel, 2016) and \odot represents element-wise multiplication.
 278

279 **Wavelet Modulation Adapter (WMA).** Our WMA, detailed in Fig. 3 (e), is designed to hierar-
 280 chically process frequency-domain features. It operates within a bottleneck, where the intermediate
 281 feature x_w undergoes a multi-level Haar wavelet decomposition. The decomposed sub-bands are
 282 passed through parameter-shared 3×1 and 1×3 depth-wise convolutions that efficiently capture di-
 283 rectional patterns. As depicted in Figure 3 (e), these processed sub-bands are then progressively
 284 fused through a three-level inverse wavelet transform. The reconstructed output is processed by a
 285 1×1 convolution to generate a map x_{wmap} that modulates the intermediate feature x_w . Finally, the
 286 modulated feature is projected back to the higher dimension. The process can be formulated as:

$$x_w = W_{down}^w(x) \quad (6)$$

$$x_{wout} = W_{up}^w(x_w \odot x_{wmap}) \quad (7)$$

289 where W_{down}^w and W_{up}^w are linear projection layers, \odot represents element-wise multiplication, and
 290 x_{wmap} denotes the wavelet processing branch detailed in Figure 3 (e), and the full formulation of
 291 x_{wmap} is provided in the Appendix B.
 292

293 3.4 CHANNEL MODULATION ADAPTER (CMA)

295 Our proposed CMA module resolves the mismatch between the entropy model and the transform
 296 module by modulating channel features within the hyperprior network. Specifically, it first down-
 297 projects the input feature x and then generates a channel attention vector x_{cmap} using global average
 298 pooling, a 1×1 convolution, and a SiLU activation (Elfwing et al., 2018). The final output is then
 299 obtained by up-projecting the down-projected feature x_c after being modulated by x_{cmap} . The
 300 process is formulated as:

$$x_c = W_{down}^c(x) \quad (8)$$

$$x_{cmap} = \text{SiLU}(\text{Conv1x1}(\text{GAP}(x_c))) \quad (9)$$

$$x_{cout} = W_{up}^c(x_{cmap} \odot x_c) \quad (10)$$

305 where W_{down}^c and W_{up}^c are the down/up-projection layers (with a ratio of 1/8) and \odot represents
 306 element-wise multiplication.
 307

308 4 EXPERIMENTS

310 **Implementation Details.** Following Chen et al. (2023), we adopt a simplified version of the TIC
 311 (Lu et al., 2021) as our base codec. The proposed SWMA, which uses an intermediate bottleneck
 312 dimension of 64, is integrated into the first three stages of the transform module, adhering to the setup
 313 in Li et al. (2024). The CMA, with an intermediate dimension of 16, is inserted into the hyperprior
 314 network. As detailed in Figure 3 (c), four of these CMA modules are placed within the hyperprior
 315 network, and the final one is placed after it to ensure alignment with the latent representation.

316 **Training Details and Datasets.** We evaluate our framework on three downstream machine vision
 317 tasks: classification, object detection, and instance segmentation. For training, we use the ImageNet-
 318 train dataset (Deng et al., 2009) and the COCO2017-train dataset (Lin et al., 2014), aligning with
 319 prior works (Chen et al., 2023).

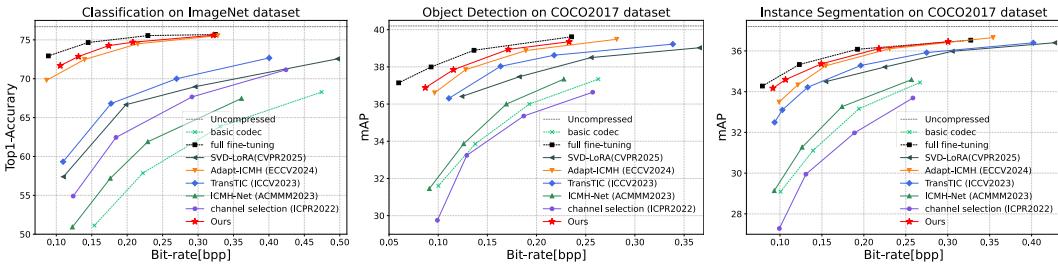
320 We adopt a parameter-efficient fine-tuning strategy, optimizing only the proposed adapters while
 321 keeping the base codec frozen. The training is guided by a task-aware loss (Eq. 2) computed using
 322 pre-trained downstream models: ResNet-50 (He et al., 2016) for classification, Faster R-CNN (Ren
 323 et al., 2015) for object detection, and Mask R-CNN (He et al., 2017) for instance segmentation.
 For data preparation, all images are randomly cropped to 256×256 . The classification adapters

324 are trained for 8 epochs with a batch size of 16, while the detection and segmentation adapters are
 325 trained for 40 epochs with a batch size of 8. All experiments are implemented in PyTorch and run on
 326 a single NVIDIA RTX 4090 with Adam (learning rate is 1e-4). Further details are in the appendix.
 327

328 **Evaluation Metrics.** We evaluate the compression rate in bits per pixel (bpp) alongside task-specific
 329 performance. Specifically, we report top-1 accuracy on the ImageNet-val dataset for classification
 330 (using a pre-trained ResNet-50), and mAP on the COCO2017-val dataset for object detection and in-
 331 stance segmentation (using models from the Detectron2 library (Wu et al., 2019)). Rate-performance
 332 curves are generated by training models at different λ values. To summarize these curves, we report
 333 two BD metrics: BD-Rate and BD-acc/mAP.

334 4.1 MAIN RESULTS

336 To demonstrate the effectiveness of the proposed method, we compare it with five SOTA approaches:
 337 Channel Selection (Liu et al., 2022), ICMH-Net (Liu et al., 2023b), TransTIC (Chen et al., 2023),
 338 Adapt-ICMH (Li et al., 2024), and SVD-LoRA (Park et al., 2025). The evaluation is conducted
 339 on three standard machine vision tasks. For a comprehensive comparison, we adopt two estab-
 340 lished metrics: the Bjontegaard Delta Rate (BD-rate) (Bjontegaard, 2001) to measure bitrate savings
 341 (lower is better), and BD-acc/mAP (Chen et al., 2023) to quantify the average performance gain at
 342 equivalent bitrates (higher is better). Figure 4 visualizes the rate-performance curves, while Table 2
 343 provides a detailed quantitative summary.



353 Figure 4: Rate-Accuracy performance comparison under different machine vision tasks.

356 Figure 4 and Table 2 show that our method consistently outperforms competing approaches across
 357 the three tasks. It maintains a trainable parameter count comparable to the SOTA method (3.8 %
 358 vs. 4.2 %) while reducing computational complexity (kMACs), thus achieving a superior perfor-
 359 mance-efficiency trade-off. To further demonstrate SWMA’s efficiency, we evaluate a lightweight
 360 variant that excludes CMA. As Table 2 indicates, despite having fewer trainable parameters and
 361 lower computational cost, this variant still reaches SOTA performance, underscoring SWMA’s ex-
 362 ceptional effectiveness and frequency-adaptation ability.

364 Table 2: Performance comparison of different methods on three machine tasks, using TIC as the
 365 base codec. We report the number of trainable parameters, computational complexity, and two BD
 366 metrics (Bjontegaard, 2001): BD-rate and BD-acc/mAP.

368 Method	Classification		Detection		Segmentation		kMACs (/pixel)	Trainable Params ↓(M)
	369 BD-rate↓	369 BD-acc↑	369 BD-rate↓	369 BD-mAP↑	369 BD-rate↓	369 BD-mAP↑		
370 full fine-tuning	/	17.67	-75.03%	4.56	-67.95%	3.76	331.0	7.511(100%)
371 channel selection	-37.18%	6.28	6.24%	-0.50	16.51%	-0.95	167.7	0.915(12.2%)
372 ICMH-Net	-18.76%	3.36	-9.82%	0.67	-10.94%	0.66	364.2	3.982(53.0%)
373 TransTIC	-58.53%	9.96	-47.50%	2.83	-46.48%	2.69	534.6	1.620(21.6%)
374 SVD-LoRA	-58.39%	7.83	-34.87%	1.96	-43.43%	1.69	<u>331.0</u>	0.093(1.2%)
375 Adapt-ICMH	-87.06%	16.68	-56.22%	3.60	-52.33%	3.21	360.4	0.287(3.8%)
376 Ours(w/o CMA)	<u>-87.45%</u>	<u>16.85</u>	<u>-60.90%</u>	<u>3.85</u>	<u>-56.60%</u>	<u>3.36</u>	359.6	0.276(3.7%)
377 Ours	-91.38%	17.15	-61.68%	4.10	-62.20%	3.49	359.9	0.319(4.2%)

378 4.2 ABLATION STUDY
379

380 **Effectiveness of the SWMA.** To evaluate the contributions of our SMA and WMA, we conduct
381 a component-wise ablation study with their counterparts from SFMA (Li et al., 2024) (Table 3).
382 The results show the superiority of our proposed modules. First, our SMA consistently improves
383 performance with fewer parameters (e.g., row 2 vs. 4 and row 1 vs. 3). Second, our WMA yields a
384 substantial performance gain over its counterpart from SFMA (Li et al., 2024) (e.g., row 1 vs. 4 and
385 row 2 vs. 3). By integrating these two components, our SWMA achieves SOTA performance with
386 lower trainable parameters.

387 Table 3: Ablations on different variants of SWMA. We replaced the spatial and frequency adapters
388 from Li et al. (2024) with our proposed adapters, achieving a significant performance improvement.
389

FMA	WMA (ours)	SMA (ours)	SMA (ours)	Classification		Detection		Segmentation		Params ↓(M)
				BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
✓	✓			-87.06%	16.68	-56.22%	3.60	-52.33%	3.21	0.287(3.8%)
	✓	✓		-87.45%	16.85	-60.90%	3.85	-56.60%	3.36	0.276(3.7%)
✓		✓		-87.12%	16.81	-56.94%	3.73	-54.45%	3.23	0.263(3.5%)
	✓	✓		-87.38%	16.85	-58.48%	3.77	-56.47%	3.31	0.300(4.0%)

390 **Effectiveness of the CMA.** To verify the efficacy of the proposed CMA, we conduct experiments
391 based on SFMA and our proposed SWMA, with results presented in Table 4. The findings reveal two
392 key points. First, the CMA significantly boosts the performance of both SFMA and SWMA with
393 minimal parameter overhead, highlighting the effectiveness of our channel modulation approach in
394 addressing the “transform–entropy mismatch”. Second, applying the CMA alone to the basic codec
395 yields negligible improvement (or even worse). This is expected because the CMA is designed as a
396 corrective module to mitigate distribution drift from fine-tuning, not as a general feature enhancer.
397

404 Table 4: Ablations on CMA.
405

SFMA	SWMA	CMA	Classification		Detection		Segmentation		Params ↓(M)
			BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
	✓		-87.45%	16.85	-60.90%	3.85	-56.60%	3.36	0.276(3.7%)
✓			-87.06%	16.68	-56.22%	3.60	-52.33%	3.21	0.287(3.8%)
✓	✓		-87.67%	17.04	-57.46%	3.91	-56.97%	3.26	0.331(4.4%)
	✓	✓	-91.38%	17.15	-61.68%	4.10	-62.20%	3.49	0.319(4.2%)
	✓		0.37%	-0.06	-1.88%	0.11	-1.49%	0.08	0.044(0.6%)

413 **Effect on the middle dimension.** Next, we analyze the middle dimension of SWMA (see Table
414 5). The results show a significant performance improvement as the dimension increases from 32 to
415 64. However, further increases beyond 64 yield diminishing gains while the parameter count keeps
416 growing. This indicates that a dimension of 64 offers the best trade-off between performance and
417 efficiency, so we adopt it as our final configuration.
418

419 Table 5: Ablations on the middle dimension of SWMA.
420

Middle Dimension	Classification		Detection		Segmentation		Params ↓(M)
	BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
32	-83.10%	16.48	-57.43%	3.75	-58.11%	3.33	0.170(2.3%)
64	-91.38%	17.15	-61.68%	4.10	-62.20%	3.49	0.319(4.2%)
128	-92.09%	17.80	-64.67%	4.25	-65.49%	3.58	0.692(9.2%)

427 4.3 APPLICATION ON LARGER SOTA METHOD
428

430 To further validate the generalization of our framework, we apply it to another SOTA codec, DCAE
431 (CVPR’25) (Lu et al., 2025). DCAE presents a particularly challenging baseline due to its signif-
432 icantly larger parameter count (119M) and a more complex, dictionary-based entropy model. We

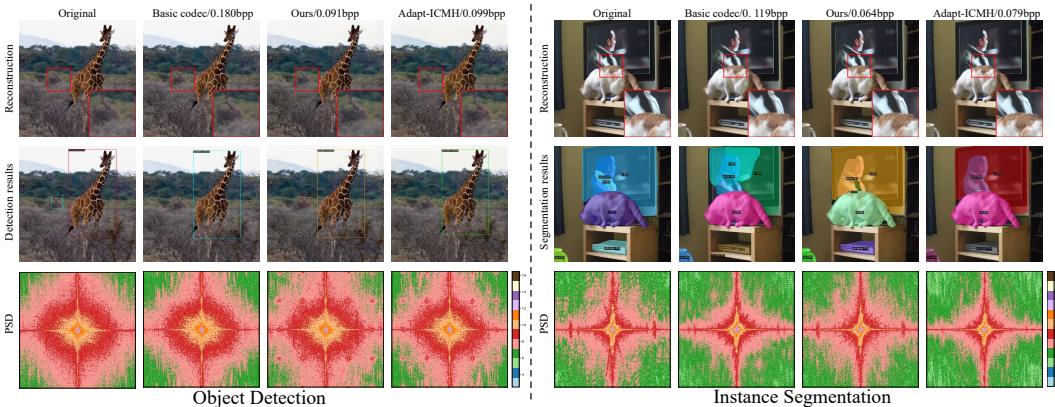
432 evaluate our method with competing ICMH methods Chen et al. (2023); Li et al. (2024) on the object
 433 detection task. The results in Table 6 show that our method outperforms competing methods while
 434 maintaining a comparable parameter count. These experiments confirm the strong generalization
 435 ability of our framework, showcasing its robustness across diverse backbones and its adaptability to
 436 different entropy models.

437
 438 Table 6: Performance comparison for object detection on the COCO2017 dataset using the chal-
 439 lenge DCAE model.

Method	Detection		Trainable Params ↓(M)
	BD-rate↓	BD-mAP↑	
full fine-tuning	-84.68%	4.00	119.760 (100%)
TransTIC	-49.27%	2.34	1.128 (0.94%)
Adapt-ICMH	-64.61%	2.86	0.359 (0.30%)
Ours	-67.87%	3.43	0.434 (0.36%)

448 449 4.4 QUALITATIVE RESULTS

450 Figure 5 provides visual comparisons that underscore the superiority of our method. Even at lower
 451 bitrates, our approach captures fine-grained details, such as the person’s tie, that are missed by
 452 competing methods. A direct comparison with Adapt-ICMH further shows that our method not only
 453 preserves critical high-frequency information (e.g., edges and textures) but also effectively enhances
 454 essential low-frequency components. We attribute these advantages to the frequency-adaptation
 455 capability of our SWMA module.



470 Figure 5: Qualitative comparison of different methods on detection and segmentation. The figure
 471 displays (from top to bottom): the original and decoded images, their corresponding results, and the
 472 log Power Spectral Density (PSD) of these images.

475 476 5 CONCLUSION

477 In this paper, we address the limitations of existing adapter-based tuning frameworks by proposing
 478 Combine-ICMH, a novel framework for the synergistic co-tuning of both transform and entropy
 479 models. Specifically, the proposed Spatial-Wavelet Modulation Adapter (SWMA) performs effi-
 480 cient feature adaptation in both spatial and frequency domains to meet the demands of downstream
 481 machine tasks. Concurrently, our proposed Channel Modulation Adapter (CMA) directly fine-tunes
 482 the entropy model, resolving the critical mismatch between the adapted features and the pre-trained
 483 entropy model. Experiments show that our method outperforms existing ICMH approaches across
 484 multiple machine vision tasks, while maintaining a comparable or even fewer number of trainable
 485 parameters. Additional experiments on more advanced backbones further demonstrate the robust-
 ness of our approach.

486

6 REPRODUCIBILITY STATEMENT

487
488 We have taken several steps to facilitate the reproduction of our results.
489

490 • **Datasets.** Section 4 details the dataset selection process and the complete pre-processing
491 pipeline.

492 • **Hyper-parameters.** Appendix C enumerates every hyper-parameter used in the main ex-
493 periments.

494 • **Compute.** Appendix C also summarizes the software and hardware specifications.

495 • **Evaluation.** Appendix A describes the task-specific perceptual-distortion losses.

496

497 REFERENCES

498
499 Yuanchao Bai, Xu Yang, Xianming Liu, Junjun Jiang, Yaowei Wang, Xiangyang Ji, and Wen Gao.
500 Towards end-to-end image compression and analysis with transformers. In *Proceedings of the*
501 *AAAI conference on artificial intelligence*, volume 36, pp. 104–112, 2022.

502
503 Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compres-
504 sion. In *5th International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=rJxdQ3jeg>.

505
506 Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational
507 image compression with a scale hyperprior. In *International Conference on Learning Represen-
508 tations*, 2018.

509
510 Gisle Bjontegard. Calculation of average psnr differences between rd-curves. *ITU-T SG16, Doc.
511 VCEG-M33*, 2001.

512
513 Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer
514 Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transac-
515 tions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021.

516
517 Lahiru D Chamain, Fabien Racapé, Jean Bégaint, Akshay Pushparaja, and Simon Feltman. End-to-
518 end optimized image compression for multiple machine tasks. *arXiv preprint arXiv:2103.04178*,
2021.

519
520 Yi-Hsin Chen, Ying-Chieh Weng, Chia-Hao Kao, Cheng Chien, Wei-Chen Chiu, and Wen-Hsiao
521 Peng. Transtic: Transferring transformer-based image compression from human perception to
522 machine perception. In *Proceedings of the IEEE/CVF International Conference on Computer
523 Vision*, pp. 23297–23307, 2023.

524
525 Zhuo Chen, Kui Fan, Shiqi Wang, Lingyu Duan, Weisi Lin, and Alex Chichung Kot. Toward intel-
526 ligent sensing: Intermediate deep feature compression. *IEEE Transactions on Image Processing*,
29:2230–2243, 2019.

527
528 Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Deep residual learning for image
529 compression. In *Cvpr Workshops*, pp. 0, 2019.

530
531 Hyomin Choi and Ivan V Bajić. Scalable image coding for humans and machines. *IEEE Transac-
532 tions on Image Processing*, 31:2739–2754, 2022.

533
534 Felipe Codevilla, Jean Gabriel Simard, Ross Goroshin, and Chris Pal. Learned image compression
535 for machine perception. *arXiv preprint arXiv:2111.02249*, 2021.

536
537 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
538 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
539 pp. 248–255. Ieee, 2009.

540
541 Shuangrui Ding, Maomao Li, Tianyu Yang, Rui Qian, Haohang Xu, Qingyi Chen, Jue Wang,
542 and Hongkai Xiong. Motion-aware contrastive video representation learning via foreground-
543 background merging. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
544 recognition*, pp. 9716–9726, 2022.

540 Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network
 541 function approximation in reinforcement learning. *Neural networks*, 107:3–11, 2018.
 542

543 Donghui Feng, Zhengxue Cheng, Shen Wang, Ronghua Wu, Hongwei Hu, Guo Lu, and Li Song.
 544 Linear attention modeling for learned image compression. In *Proceedings of the Computer Vision
 545 and Pattern Recognition Conference*, pp. 7623–7632, 2025.

546 Ruoyu Feng, Xin Jin, Zongyu Guo, Runsen Feng, Yixin Gao, Tianyu He, Zhizheng Zhang, Simeng
 547 Sun, and Zhibo Chen. Image coding for machines with omnipotent feature learning. In *European
 548 Conference on Computer Vision*, pp. 510–528. Springer, 2022.
 549

550 Kristian Fischer, Fabian Brand, and André Kaup. Boosting neural image compression for machines
 551 using latent space masking. *IEEE Transactions on Circuits and Systems for Video Technology*,
 552 2022.

553 Ahmed Ghorbel, Wassim Hamidouche, and Luce Morin. Convnext-charm: Convnext-based trans-
 554 form for efficient neural image compression. In *2023 11th European Workshop on Visual Infor-
 555 mation Processing (EUVIP)*, pp. 1–6. IEEE, 2023.
 556

557 Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context
 558 model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on
 559 Computer Vision and Pattern Recognition*, pp. 14771–14780, 2021.

560 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
 561 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
 562 770–778, 2016.
 563

564 Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the
 565 IEEE international conference on computer vision*, pp. 2961–2969, 2017.
 566

567 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint
 568 arXiv:1606.08415*, 2016.
 569

570 Wei Jiang and Ronggang Wang. Mlic++: Linear complexity multi-reference entropy model-
 571 ing for learned image compression. In *ICML 2023 Workshop Neural Compression: From In-
 572 formation Theory to Applications*, 2023. URL <https://openreview.net/forum?id=hxIpcSoz2t>.
 573

574 Nam Le, Honglei Zhang, Francesco Cricri, Ramin Ghaznavi-Youvalari, and Esa Rahtu. Image
 575 coding for machines: an end-to-end learned approach. In *ICASSP 2021-2021 IEEE International
 576 Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1590–1594. IEEE, 2021.
 577

578 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt
 tuning. *arXiv preprint arXiv:2104.08691*, 2021.
 579

580 Han Li, Shaohui Li, Shuangrui Ding, Wenrui Dai, Maida Cao, Chenglin Li, Junni Zou, and Hongkai
 581 Xiong. Image compression for machine and human vision with spatial-frequency adaptation. In
 582 *European Conference on Computer Vision*, pp. 382–399. Springer, 2024.
 583

584 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv
 preprint arXiv:2101.00190*, 2021.
 585

586 Hongbin Lin, Bolin Chen, Zhichen Zhang, Jielian Lin, Xu Wang, and Tiesong Zhao. Deepsvc:
 587 Deep scalable video coding for both machine and human vision. In *Proceedings of the 31st ACM
 588 International Conference on Multimedia*, pp. 9205–9214, 2023.
 589

590 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 591 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European
 592 conference on computer vision*, pp. 740–755. Springer, 2014.
 593

594 Jinming Liu, Heming Sun, and Jiro Katto. Improving multiple machine vision tasks in the com-
 595 pressed domain. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pp. 331–
 596 337. IEEE, 2022.

594 Jinming Liu, Xin Jin, Ruoyu Feng, Zhibo Chen, and Wenjun Zeng. Composable image coding
 595 for machine via task-oriented internal adaptor and external prior. In *2023 IEEE International*
 596 *Conference on Visual Communications and Image Processing (VCIP)*, pp. 1–5. IEEE, 2023a.
 597

598 Lei Liu, Zhihao Hu, Zhenghao Chen, and Dong Xu. Icmh-net: Neural image compression towards
 599 both machine vision and human vision. In *Proceedings of the 31st ACM International Conference*
 600 *on Multimedia*, pp. 8047–8056, 2023b.

601 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
 602 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the*
 603 *IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.

604 Jingbo Lu, Leheng Zhang, Xingyu Zhou, Mu Li, Wen Li, and Shuhang Gu. Learned image compres-
 605 sion with dictionary-based entropy model. In *Proceedings of the Computer Vision and Pattern*
 606 *Recognition Conference*, pp. 12850–12859, 2025.

607 Ming Lu, Peiyao Guo, Huiqing Shi, Chunlong Cao, and Zhan Ma. Transformer-based image com-
 608 pression. *arXiv preprint arXiv:2111.06707*, 2021.

609 David Minnen, Johannes Ballé, and George D. Toderici. Joint autoregressive and hierarchical priors
 610 for learned image compression. *Advances in neural information processing systems*, 31, 2018.

611 Unki Park, Seongmoon Jeong, Youngchan Jang, Gyeong-Moon Park, and Jong Hwan Ko. Test-
 612 time fine-tuning of image compression models for multi-task adaptability. In *Proceedings of the*
 613 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4430–4440,
 614 June 2025.

615 Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object
 616 detection with region proposal networks. *Advances in neural information processing systems*, 28,
 617 2015.

618 Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high
 619 efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video*
 620 *technology*, 22(12):1649–1668, 2012.

621 Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning
 622 for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and*
 623 *pattern recognition*, pp. 5693–5703, 2019.

624 Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with
 625 discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF*
 626 *Conference on Computer Vision and Pattern Recognition*, pp. 7939–7948, 2020.

627 Shurun Wang, Shiqi Wang, Wenhan Yang, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen
 628 Gao. Towards analysis-friendly face representation with scalable feature and texture compression.
 629 *IEEE Transactions on Multimedia*, 24:3169–3181, 2021.

630 Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2.
 631 <https://github.com/facebookresearch/detectron2>, 2019.

632 Enze Xie, Wenhui Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Seg-
 633 former: Simple and efficient design for semantic segmentation with transformers. *Advances in*
 634 *neural information processing systems*, 34:12077–12090, 2021a.

635 Yueqi Xie, Ka Leong Cheng, and Qifeng Chen. Enhanced invertible encoding for learned image
 636 compression. In *Proceedings of the 29th ACM international conference on multimedia*, pp. 162–
 637 170, 2021b.

638 Xu Zhang, Peiyao Guo, Ming Lu, and Zhan Ma. All-in-one image coding for joint
 639 human-machine vision with multi-path aggregation. In A. Globerson, L. Mackey, D. Bel-
 640 grave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural In-*
 641 *formation Processing Systems*, volume 37, pp. 71465–71503. Curran Associates, Inc.,
 642 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/8395fdf356059eaa92afd39e3952a677-Paper-Conference.pdf.

648 Jiancheng Zhao, Xiang Ji, Zhuoxiao Li, Zunian Wan, Weihang Ran, Mingze Ma, Muyao Niu, Yifan
649 Zhan, Cheng-Ching Tseng, and Yinqiang Zheng. All-in-one transferring image compression from
650 human perception to multi-machine perception. *arXiv preprint arXiv:2504.12997*, 2025.
651
652 Yinhao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International
653 Conference on Learning Representations*, 2022.
654 Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based
655 attention for image compression. In *Proceedings of the IEEE/CVF conference on computer vision
656 and pattern recognition*, pp. 17492–17501, 2022.
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

Appendix

A TASK-SPECIFIC PERCEPTUAL LOSS

To effectively guide the model’s training, we employ a task-specific perceptual loss (D_{task}) (following Chen et al. (2023)). This perceptual loss is defined by calculating the Mean Squared Error (MSE) between features of the original image x and the reconstructed image \hat{x} (generated by image codecs), where the features are extracted from a pre-trained, task-relevant network.

For different downstream tasks, we utilize specific network feature layers and corresponding loss formulations. The detailed configurations are summarized in Table 7.

To visually illustrate the feature layers used for loss computation, we highlight their positions within the network architecture in Figure 6.

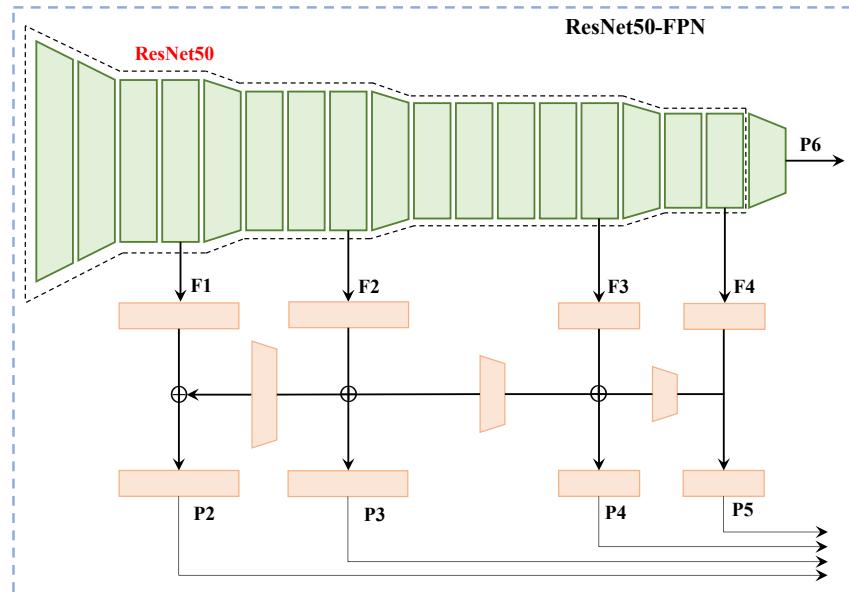


Figure 6: Network architecture of ResNet50-FPN.

B DETAIL INFORMATION FOR WMA

This section provides supplementary details on the calculation formula for the wavelet modulation map w_{map} .

Table 7: Perceptual Loss Configurations for Different Downstream Tasks.

Downstream Task	Network	Feature Layers Used	Loss Formulation
Classification	ResNet-50 (He et al., 2016)	F_1, F_2, F_3, F_4	$D(x, \hat{x}, G) = \frac{1}{4} \sum_{j=1}^4 \text{MSE}(F_j(x), F_j(\hat{x}))$
Object Detection	Faster R-CNN (Ren et al., 2015)	P_2, P_3, P_4, P_5, P_6	$D(x, \hat{x}, G) = \frac{1}{5} \sum_{j=2}^6 \text{MSE}(P_j(x), P_j(\hat{x}))$
Instance Segmentation	Mask R-CNN (He et al., 2017)	P_2, P_3, P_4, P_5, P_6	

756 In the computation, the intermediate feature map w_m , along with the output of each wavelet trans-
 757 form level, is processed by a sequence of 3x1 and 1x3 depthwise convolutions. The process is as
 758 follows:

$$759 \quad \text{DW}(T) = \text{DWConv1} \times 3(\text{DWConv3} \times 1(T)) \quad (11)$$

$$760 \quad T^{(0)} = \text{DW}(X) \quad (12)$$

$$761 \quad T^{(l)} = \text{DW}(\text{WT}(T^{(l-1)})), \quad l = 1, 2, 3 \quad (13)$$

762 where, X denotes the input features, and WT denotes the Haar wavelet transform. Specifically,
 763 when processing feature in wavelet domain, a parameter sharing strategy is adopted for better effi-
 764 ciency.

765 The Haar wavelet transform decomposes a feature into four components: LL, HL, LH, and HH. We
 766 define LL as the low-frequency component T_L and group HL, LH, and HH as the high-frequency
 767 component T_H .

$$768 \quad T_L^{(l)}, T_H^{(l)} = T^{(l)} \quad (14)$$

769 Leveraging the linearity of the wavelet transform, we process the wavelet components independently
 770 before being recomposed for the inverse transform. This process yields the aggregated wavelet
 771 feature $R^{(1)}$.

$$772 \quad R^{(3)} = \text{IWT}(T^{(3)}), \quad (15)$$

$$773 \quad R^{(2)} = \text{IWT}(T_L^{(2)} + R^{(3)}, T_H^{(2)}), \quad (16)$$

$$774 \quad R^{(1)} = \text{IWT}(T_L^{(1)} + R^{(2)}, T_H^{(3)}), \quad (17)$$

775 First, the aggregated wavelet feature $R^{(1)}$ is summed with the $T^{(0)}$. This combined feature, cap-
 776 turing diverse frequency characteristics, is then processed sequentially by a GELU activation and a
 777 1x1 convolution for fusion. The final output is the wavelet modulation map x_{map} .

$$785 \quad x_{map} = \text{GELU}(\text{Conv1x1}(\text{GELU}(R^{(1)} + T^{(0)}))). \quad (18)$$

787 C TRAINING DETAILS AND HYPERPARAMTERS

788 All experiments used PyTorch. The detailed software and hardware configurations are listed below:

789 **Software Environment:**

- 790 • Operating System: Ubuntu 20.04
- 791 • Framework: PyTorch 2.4
- 792 • Libraries: CUDA 12.4, cuDNN 8.9

793 **Hardware Platform:**

- 794 • GPU: 1x NVIDIA GeForce RTX 4090 (24 GB VRAM)
- 795 • CPU: Intel(R) Xeon(R) Silver 4310 @ 2.10 GHz
- 796 • System Memory: 256 GB RAM

802 Table 8 shows all the hyperparameters used in experiment.

804 D DETAILS OF ARCHITECTURE FOR DIFFERENT BASE CODECS

805 In Figures 7, 8, and 9, we provide the detailed network architectures of the proposed Combine-
 806 ICMH framework for different image codecs. The Combine-ICMH framework consistently outper-
 807 forms previous methods varying in architecture, all while preserving the pretrained model’s R-D
 808 performance.

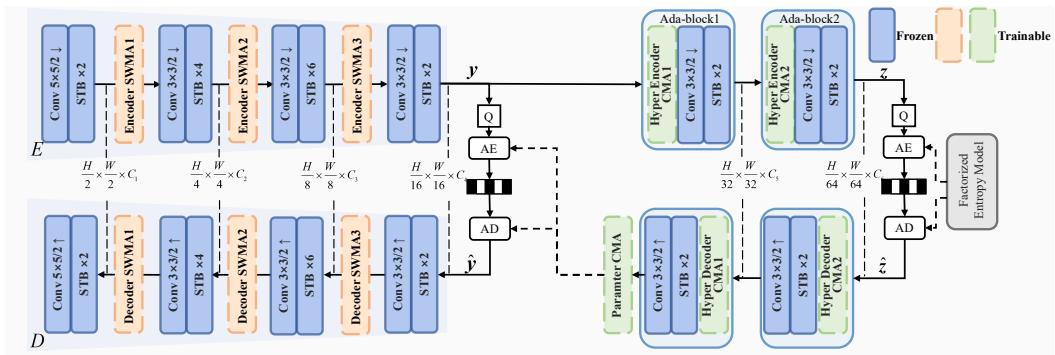


Figure 7: For a fair comparison, we follow the (Chen et al., 2023) and adopt a simplified version of the TIC model (Lu et al., 2021) as our base codec. In this model, STB denotes the Swin-Transformer Block (Liu et al., 2021).

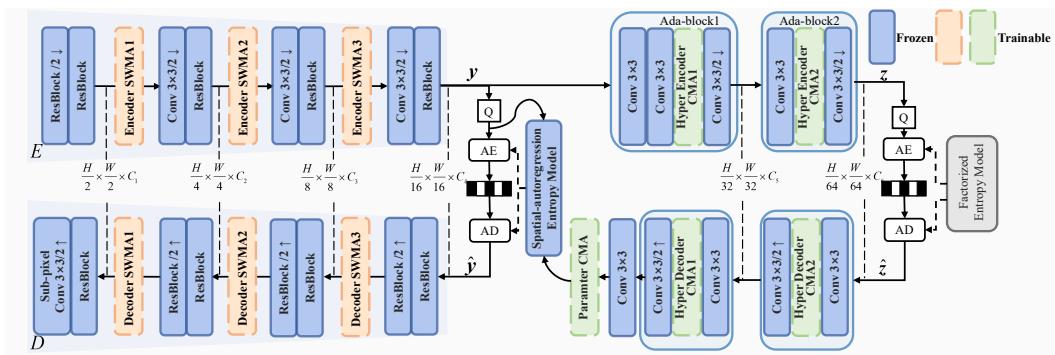


Figure 8: For further comparison, we use Cheng2020-anchor model Cheng et al. (2020) as base codec. In this model, Resblock denotes the residual block.

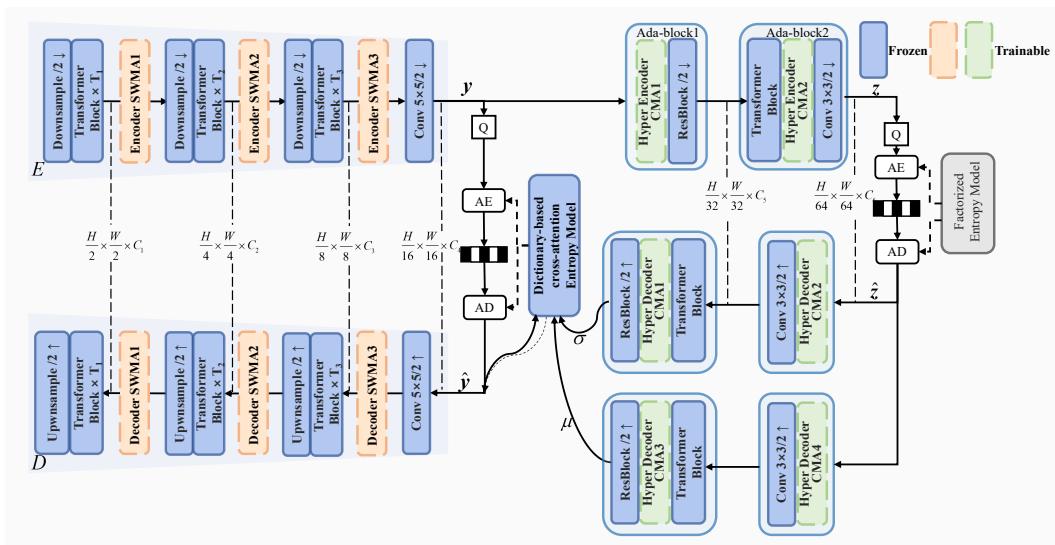


Figure 9: For further comparison, we use DCAE model Lu et al. (2025) as base codec. In this model, *Downsample* denotes three stacked residual blocks and a stride convolution, and *Upsample* denotes three stacked residual blocks and sub-pixel upsampling on the last convolution.

Table 8: Training hyperparamters for experiments.

	Classification	Detection	Segmentation
Optimizer	Adam	Adam	Adam
Batch size	16	8	8
Trade-off term λ	[2.5, 3.5, 5, 6.7, 13]	[0.5, 0.875, 1.75, 3]	[0.35, 0.5, 0.875, 1.75, 3]
Epochs	8	40	40
Learning rate schedule	MultiStepLR	-	-
Milestones	[5,7]	-	-
Learning rate decay	0.5	-	-
Base learning rate	1e-4	1e-4	1e-4

E ANALYSIS OF CHANNEL MODULATION

In a standard Learned Image Compression (LIC) framework, the non-linear transform (g_a, g_s) and the entropy model (h_a, h_s) are jointly optimized. The transform converts an input signal into a compact latent representation \mathbf{y} , while the entropy model estimates the probability distribution of its quantized version, $\hat{\mathbf{y}}$, to minimize coding redundancy.

In hyperprior-based models (Ballé et al., 2018), $\bar{\mathbf{y}} \stackrel{\Delta}{=} (\mathbf{y} - \boldsymbol{\mu})/\boldsymbol{\sigma}$ is modeled as a standard spherical normal vector, Gaussianizing the source distribution to minimize coding overhead (Zhu et al., 2022). The parameters $(\boldsymbol{\mu}, \boldsymbol{\sigma})$ are generated by the hyperprior decoder h_s from the hyper-latents $\hat{\mathbf{z}}$, which are in turn produced by the hyperprior encoder h_a :

$$\mathbf{z} = h_a(\mathbf{y}; \boldsymbol{\theta}_{h_a}) \quad (19)$$

$$\hat{\mathbf{z}} = Q(\mathbf{z}) \quad (20)$$

$$(\boldsymbol{\mu}, \boldsymbol{\sigma}) = h_s(\hat{\mathbf{z}}; \boldsymbol{\theta}_{h_s}) \quad (21)$$

where $Q(\cdot)$ denotes quantization.

In practice, the hyperprior-based LIC framework employs the Gaussian parameters $(\boldsymbol{\mu}, \boldsymbol{\sigma})$ to model the probability distribution of the quantized latent representation $\hat{\mathbf{y}}$ and calculate the bitrate $R(\hat{\mathbf{y}})$ (Feng et al., 2025):

$$R(\hat{\mathbf{y}}) = \mathbb{E} [-\log_2(p_{\hat{\mathbf{y}}|\hat{\mathbf{z}}}(\hat{\mathbf{y}}|\hat{\mathbf{z}}))] \quad (22)$$

$$p_{\hat{\mathbf{y}}|\hat{\mathbf{z}}}(\hat{\mathbf{y}}|\hat{\mathbf{z}}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2) \quad (23)$$

However, during fine-tuning in previous frameworks (Chen et al., 2023; Li et al., 2024), only the transform modules (g_a, g_s) are trained, while the entropy model (h_a, h_s) is kept frozen. This introduces a ‘‘transform-entropy mismatch.’’ Specifically, fine-tuning shifts the original latent \mathbf{y} to a new task-specific latent \mathbf{y}' , altering its statistical distribution.

Consequently, the Gaussian parameters $(\boldsymbol{\mu}', \boldsymbol{\sigma}')$ generated by the frozen entropy model for the new latent \mathbf{y}' are suboptimal. They are derived via the frozen hyperprior network:

$$\mathbf{z}' = h_a(\mathbf{y}'; \boldsymbol{\theta}_{h_a}) \quad (24)$$

$$\hat{\mathbf{z}}' = Q(\mathbf{z}') \quad (25)$$

$$(\boldsymbol{\mu}', \boldsymbol{\sigma}') = h_s(\hat{\mathbf{z}}'; \boldsymbol{\theta}_{h_s}) \quad (26)$$

This suboptimality manifests as a deviation $(\Delta\boldsymbol{\mu}, \Delta\boldsymbol{\sigma})$ from the ideal parameters $(\boldsymbol{\mu}_{ideal}, \boldsymbol{\sigma}_{ideal})$ that model \mathbf{y}' into a standard spherical normal vector. The deviation is defined as:

$$\Delta\boldsymbol{\mu} = \boldsymbol{\mu}' - \boldsymbol{\mu}_{ideal} \quad (27)$$

$$\Delta\boldsymbol{\sigma} = \boldsymbol{\sigma}' - \boldsymbol{\sigma}_{ideal} \quad (28)$$

As a result, the normalized latent $\bar{\mathbf{y}}'$ deviates from the standard Gaussian distribution. This mismatch results in encoding redundancy. We can express the actual bitrate $R(\hat{\mathbf{y}}')$ by decomposing it into an ideal bitrate and a penalty term, Δbpp , which is defined as:

$$R(\hat{\mathbf{y}}') = \mathbb{E} [-\log_2(p_{\hat{\mathbf{y}}'|\hat{\mathbf{z}}'}(\hat{\mathbf{y}}'|\hat{\mathbf{z}}'))] = \mathbb{E} [-\log_2(p_{\hat{\mathbf{y}}'|\hat{\mathbf{z}}}(\hat{\mathbf{y}}'|\hat{\mathbf{z}}))] + \Delta bpp \quad (29)$$

$$p_{\hat{\mathbf{y}}'|\hat{\mathbf{z}}}(\hat{\mathbf{y}}'|\hat{\mathbf{z}}) \sim \mathcal{N}(\boldsymbol{\mu}_{ideal}, \boldsymbol{\sigma}_{ideal}^2) \quad (30)$$

918 where $\tilde{\mathbf{z}}$ is the assumed hyper-latent that enables the latent representation \mathbf{y}' to be modeled as a standard Gaussian distribution, and Δbpp represents the increased bpp compared to the ideal scenario where the entropy model and transform module are perfectly matched.

919
920
921 To mitigate the mismatch, we propose the Channel-Modulation Adapter (CMA) module, which is
922 specifically designed to resolve the “transform-entropy mismatch” by adapting the hyperprior to
923 meet the fine-tuned transform, thereby minimizing coding redundancy.

924
925 The proposed CMA mitigates the “transform-entropy mismatch” by adjusting the entropy model
926 (h_a, h_s). This adaptation is designed to minimize $\Delta\mu$ and $\Delta\sigma$. Ideally, if $\Delta\mu$ and $\Delta\sigma$ are driven
927 to zero, the coding redundancy introduced by the fine-tuning transform module is completely eliminated.

928
929 In the ideal case:

930
931
$$R(\hat{\mathbf{y}}') = \mathbb{E} [-\log_2(p_{\hat{\mathbf{y}}'|\tilde{\mathbf{z}}'}(\hat{\mathbf{y}}'|\tilde{\mathbf{z}}'))] = \mathbb{E} [-\log_2(p_{\hat{\mathbf{y}}'|\tilde{\mathbf{z}}}(\hat{\mathbf{y}}'|\tilde{\mathbf{z}}))] \quad (31)$$

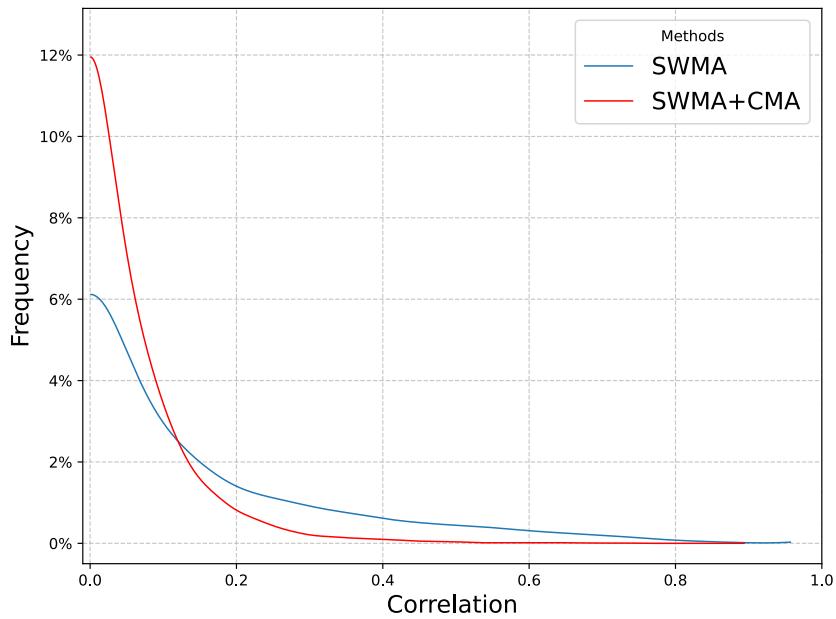
932
$$\Delta bpp = 0 \quad (32)$$

933
934 To further quantitatively analyze the effectiveness of the proposed CMA in reducing coding redundancy,
935 we measure various correlations within the latent representation y , following the Zhu et al.
936 (2022).

937
938 We compare two fine-tuning strategies at the lowest-bitrate setting, where redundancy is most pro-
939 nounced: (i) a baseline that inserts only the SWMA adapter and (ii) the full approach that jointly
940 optimizes SWMA and CMA. Both variants fine-tune the TIC codec (Lu et al., 2021) for object
941 detection on the COCO2017-val dataset.

942 First, we estimate the inter-channel correlation of the latent representation y by calculating the pair-
943 wise channel similarity, as depicted in Figure 10.

944 As shown in Figure 10, the baseline model (SWMA-only) exhibits high inter-channel correlation,
945 implying significant redundancy among channels. In contrast, the model equipped with our CMA
946 (SWMA+CMA) demonstrates a marked reduction in correlation, evidenced by a much steeper decay
947 in the similarity curve. These findings strongly suggest that the CMA promotes the learning of more
948 distinct, decorrelated features, thereby effectively reducing coding redundancy.



950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971 Figure 10: Pairwise channel similarities of the latent representation y for different fine-tuning strate-
gies. Incorporating CMA markedly reduce the inter-channel correlation.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Next, we extend our analysis to the spatial dimension. Following the Zhu et al. (2022), we measure the correlation between nearby spatial positions, which is averaged across all channels. As visualized in Figure 11, the results show a marked improvement. While the SWMA-only baseline exhibits a high average spatial correlation ($\rho=0.3283$), our SWMA+CMA model substantially reduces this to $\rho=0.2439$. This reduction in spatial redundancy directly corresponds to the improved downstream task performance, validating that our framework yields a more compact and effective latent representation.

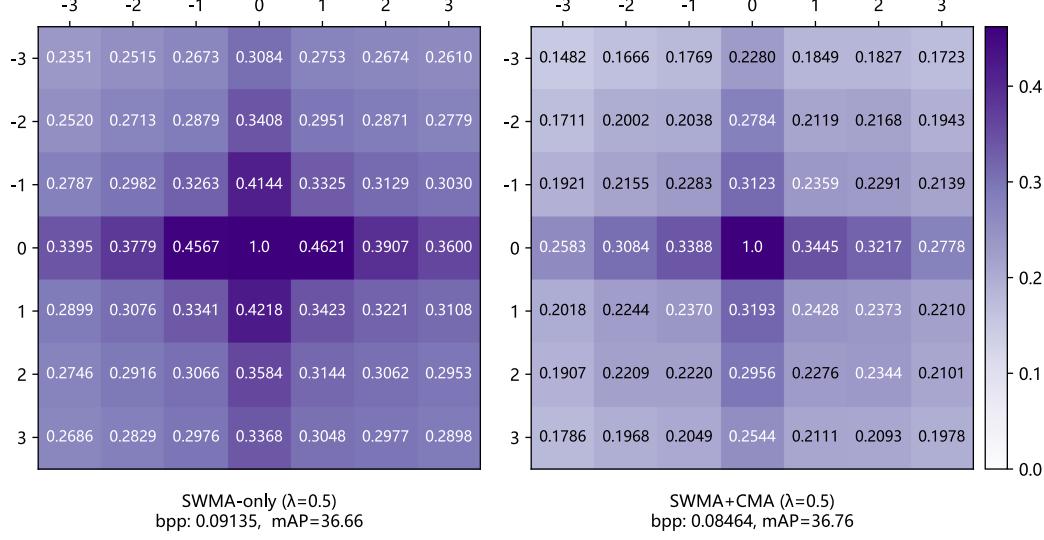


Figure 11: Spatial correlation of $(y - \mu)/\sigma$ with models trained at $\lambda = 0.5$. SWMA+CMA (right) achieves average smaller correlation than SWMA-only (left).

In summary, our comprehensive correlation analysis across both channel and spatial dimensions provides strong evidence for the effectiveness of our proposed fine-tuning paradigm. Prior methods Li et al. (2024); Chen et al. (2023), which only fine-tune the transform module, inevitably cause a “transform-entropy mismatch.” Our approach directly addresses this issue by introducing the CMA module, which facilitates a collaborative optimization between the entropy model and the transform. The significant reduction in inter-channel and spatial correlation directly reflects an improvement in encoding efficiency, underscoring the benefit of this collaborative optimization.

Furthermore, we analyze the CMA’s specific impact on the hyperprior. As illustrated in Figure 3, our approach differs from conventional channel attention mechanisms that multiplicatively scale each channel. Instead, the CMA computes an additive correction based on the input features x to modify their distribution. Specifically, in a hyperprior-based entropy model, the CMA module directly operates on the output of the hyperprior network (h_a, h_s). This allows it to adapt the distribution estimation by modifying the Gaussian parameters (μ, σ) .

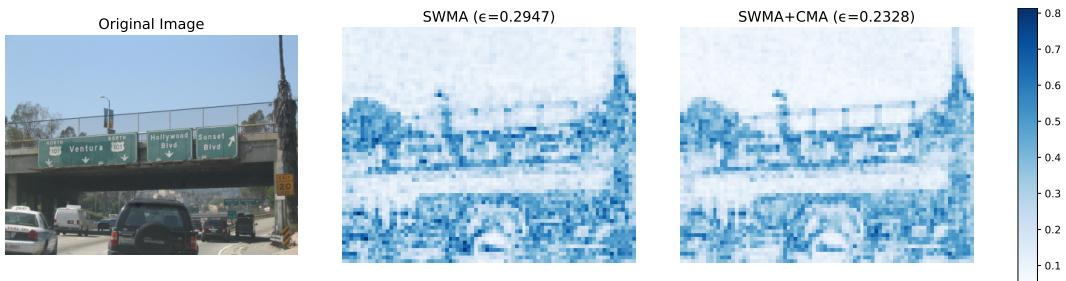
To evaluate this adaptation, we analyze the quantization loss of the latent representation y , following the methodology of (Xie et al., 2021b). The quantization loss during compression is quantified using a “scaling deviation” metric (Xie et al., 2021b), which is defined as $\epsilon = \text{abs}(\hat{y} - y)/\Sigma y$. In the context of the Gaussian entropy model, the quantization process of the latent representation y can be formulated as $\hat{y} = Q(y - \mu) + \mu$, so there is the following equation:

$$\epsilon = \text{abs}(Q(y - \mu) - (y - \mu))/\Sigma y \quad (33)$$

Equation 33 reveals a direct link between the Gaussian parameter μ and the quantization loss. A smaller distance between μ and y results in lower quantization loss. Therefore, we use this metric as a proxy to evaluate our paradigm’s effectiveness in adapting the hyperprior statistics.

Figure 12 presents the scaled deviation map on the COCO2017-val dataset. The results clearly demonstrate that our CMA module significantly reduces the quantization loss, particularly in background regions. This indicates that our paradigm, by adapting the Gaussian parameter μ , more

1026 accurately estimates the statistical distribution of the shifted latent representation \mathbf{y}' , thereby mitigating the “transform-entropy mismatch”.
 1027
 1028



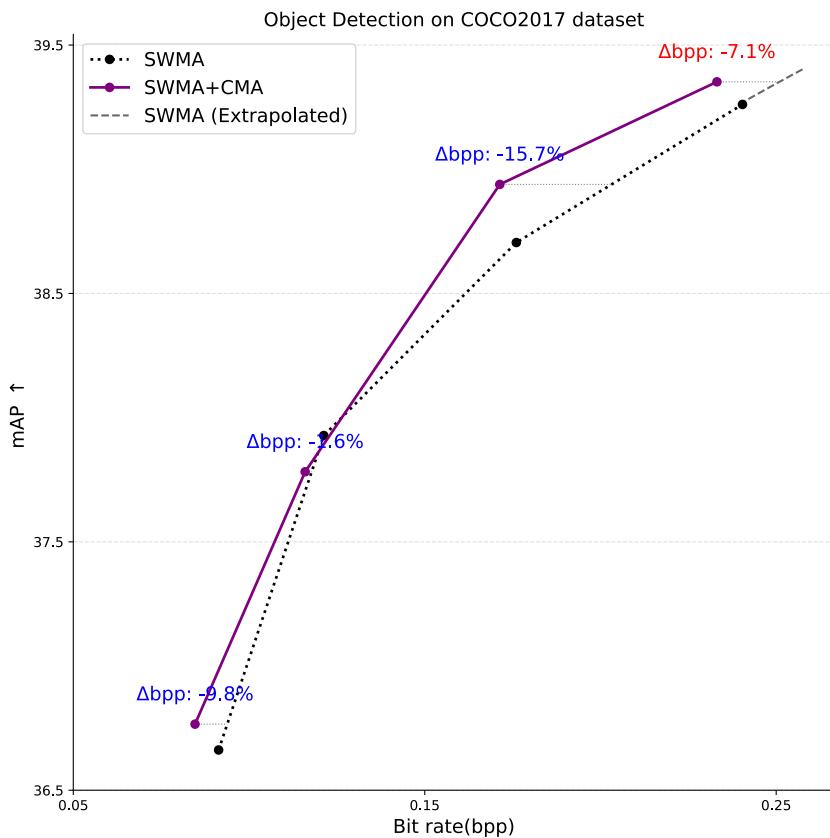
1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079
 1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187
 1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619
 1620
 1621
 1622
 1623
 1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639
 1640
 1641
 1642
 1643
 1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673
 1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727
 1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781
 1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835
 1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889
 1890
 1891
 1892
 1893
 1894
 1895
 1896
 1897
 1898
 1899
 1900
 1901
 1902
 1903
 1904
 1905
 1906
 1907
 1908
 1909
 1910
 1911
 1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943
 1944
 1945
 1946
 1947
 1948
 1949
 1950
 1951
 1952
 1953
 1954
 1955
 1956
 1957
 1958
 1959
 1960
 1961
 1962
 1963
 1964
 1965
 1966
 1967
 1968
 1969
 1970
 1971
 1972
 1973
 1974
 1975
 1976
 1977
 1978
 1979
 1980
 1981
 1982
 1983
 1984
 1985
 1986
 1987
 1988
 1989
 1990
 1991
 1992
 1993
 1994
 1995
 1996
 1997
 1998
 1999
 2000
 2001
 2002
 2003
 2004
 2005
 2006
 2007
 2008
 2009
 2010
 2011
 2012
 2013
 2014
 2015
 2016
 2017
 2018
 2019
 2020
 2021
 2022
 2023
 2024
 2025
 2026
 2027
 2028
 2029
 2030
 2031
 2032
 2033
 2034
 2035
 2036
 2037
 2038
 2039
 2040
 2041
 2042
 2043
 2044
 2045
 2046
 2047
 2048
 2049
 2050
 2051
 2052
 2053
 2054
 2055
 2056
 2057
 2058
 2059
 2060
 2061
 2062
 2063
 2064
 2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075
 2076
 2077
 2078
 2079
 2080
 2081
 2082
 2083
 2084
 2085
 2086
 2087
 2088
 2089
 2090
 2091
 2092
 2093
 2094
 2095
 2096
 2097
 2098
 2099
 2100
 2101
 2102
 2103
 2104
 2105
 2106
 2107
 2108
 2109
 2110
 2111
 2112
 2113
 2114
 2115
 2116
 2117
 2118
 2119
 2120
 2121
 2122
 2123
 2124
 2125
 2126
 2127
 2128
 2129
 2130
 2131
 2132
 2133
 2134
 2135
 2136
 2137
 2138
 2139
 2140
 2141
 2142
 2143
 2144
 2145
 2146
 2147
 2148
 2149
 2150
 2151
 2152
 2153
 2154
 2155
 2156
 2157
 2158
 2159
 2160
 2161
 2162
 2163
 2164
 2165
 2166
 2167
 2168
 2169
 2170
 2171
 2172
 2173
 2174
 2175
 2176
 2177
 2178
 2179
 2180
 2181
 2182
 2183
 2184
 2185
 2186
 2187
 2188
 2189
 2190
 2191
 2192
 2193
 2194
 2195
 2196
 2197
 2198
 2199
 2200
 2201
 2202
 2203
 2204
 2205
 2206
 2207
 2208
 2209
 2210
 2211
 2212
 2213
 2214
 2215
 2216
 2217
 2218
 2219
 2220
 2221
 2222
 2223
 2224
 2225
 2226
 2227
 2228
 2229
 2230
 2231
 2232
 2233
 2234
 2235
 2236
 2237
 2238
 2239
 2240
 2241
 2242
 2243
 2244
 2245
 2246
 2247
 2248
 2249
 2250
 2251
 2252
 2253
 2254
 2255
 2256
 2257
 2258
 2259
 2260
 2261
 2262
 2263
 2264
 2265
 2266
 2267
 2268
 2269
 2270
 2271
 2272
 2273
 2274
 2275
 2276
 2277
 2278
 2279
 2280
 2281
 2282
 2283
 2284
 2285
 2286
 2287
 2288
 2289
 2290
 2291
 2292
 2293
 2294
 2295
 2296
 2297
 2298
 2299
 2300
 2301
 2302
 2303
 2304
 2305
 2306
 2307
 2308
 2309
 2310
 2311
 2312
 2313
 2314
 2315
 2316
 2317
 2318
 2319
 2320
 2321
 2322
 2323
 2324
 2325
 2326
 2327
 2328
 2329
 2330
 2331
 2332
 2333
 2334
 2335
 2336
 2337
 2338
 2339
 2340
 2341
 2342
 2343
 2344
 2345
 2346
 2347
 2348
 2349
 2350
 2351
 2352
 2353
 2354
 2355
 2356
 2357
 2358
 2359
 2360
 2361
 2362
 2363
 2364
 2365
 2366
 2367
 2368
 2369
 2370
 2371
 2372
 2373
 2374
 2375
 2376
 2377
 2378
 2379
 2380
 2381
 2382
 2383
 2384
 2385
 2386
 2387
 2388
 2389
 2390
 2391
 2392
 2393
 2394
 2395
 2396
 2397
 2398
 2399
 2400
 2401
 2402
 2403
 2404
 2405
 2406
 2407
 2408
 2409
 2410
 2411
 2412
 2413
 2414
 2415
 2416
 2417
 2418
 2419
 2420
 2421
 2422
 2423
 2424
 242

1080 and decoder for reducing spatial and frequency redundancy. Another approach (Park et al., 2025)
 1081 has utilized SVD-LoRA to fine-tune the codec, combined with Test-Time Fine-tuning (TTFT) for
 1082 instance-specific optimization. A common thread in these methods is their exclusive focus on adapt-
 1083 ing the transform modules, neglecting the entropy model.

1084 In contrast, our work stems from a key observation: neglecting the entropy model during fine-tuning
 1085 leads to a “transform-entropy mismatch.” This mismatch increases coding redundancy, which de-
 1086 grades both coding efficiency and overall task performance. To address this, we propose the Channel
 1087 Modulation Adapter (CMA), a lightweight adapter designed specifically to adapt the entropy model
 1088 and resolve this mismatch.

1089 To validate our claim, we compared four architectural configurations: SWMA-only, SWMA +
 1090 CMA, SFMA-only, and SFMA + CMA. As illustrated in Figures 14 and 15, integrating our CMA
 1091 yields a significant reduction in bit-per-pixel (bpp) without compromising task performance. This
 1092 result confirms that CMA effectively mitigates coding redundancy in the ICMH fine-tuning process.

1093 The quantitative results in Table 9 further corroborate this finding. Specifically, the CMA provides
 1094 an average bitrate saving of 5.05% when integrated with the SWMA baseline and 3.89% with the
 1095 SFMA baseline. Its effectiveness across different models further highlights the importance of fine-
 1096 tuning the entropy model in the ICMH.



1125 Figure 14: Object-detection performance on the COCO2017-val dataset with TIC (Lu et al., 2021) as
 1126 the basic codec. Incorporating CMA markedly reduces the bitrate without compromising detection
 1127 accuracy.

1128
 1129
 1130
 1131
 1132
 1133

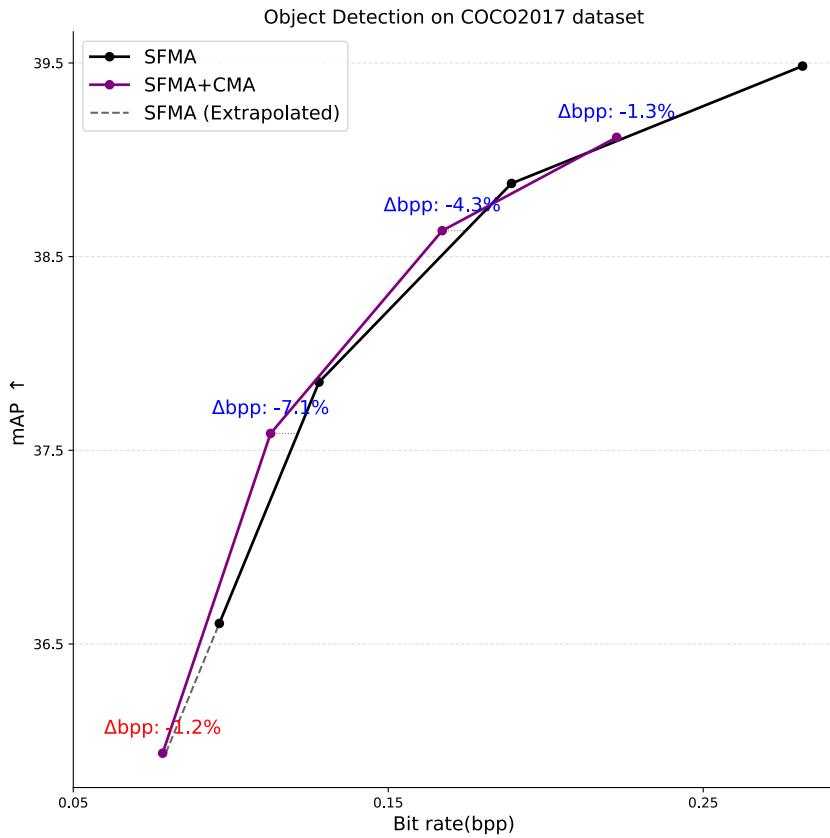


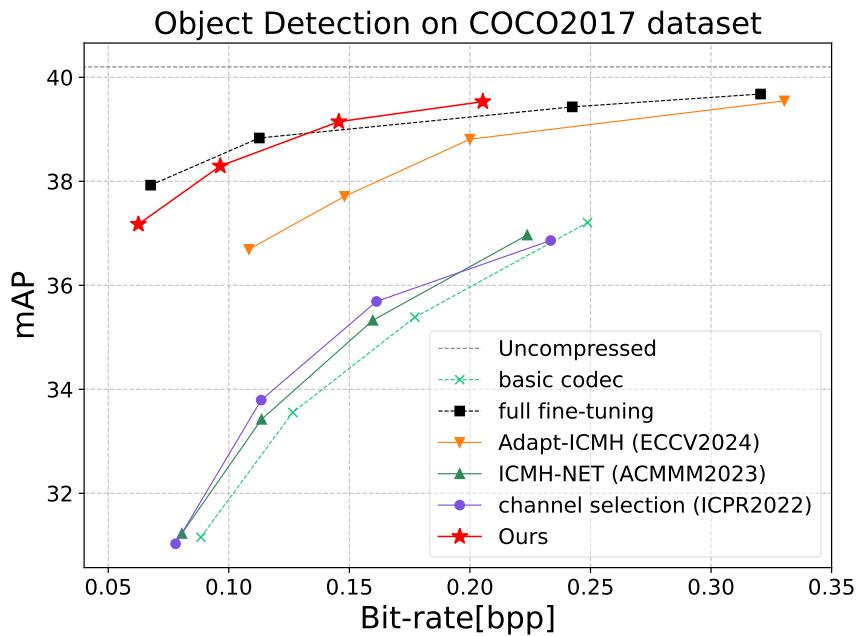
Figure 15: Object-detection performance on the COCO-2017 dataset with TIC (Lu et al., 2021) as the basic codec. Incorporating CMA markedly reduces the bitrate without compromising detection accuracy.

Table 9: Performance comparison of different methods on Object Detection, using TIC as the base codec. We report the number of trainable parameters and two BD metrics (Bjontegaard, 2001): BD-rate and BD-acc/mAP.

Method	Detection		Trainable
	BD-rate \downarrow	BD-mAP \uparrow	Params \downarrow (M)
SWMA-only	0.00%	0.00	0.276 (3.7%)
SWMA+CMA	-5.05%	0.17	0.319 (4.2%)
SFMA-only	0.00%	0.00	0.287 (3.8%)
SFMA+CMA	-3.89%	0.10	0.331 (4.4%)

1188 G MORE RESULTS
11891190 G.1 IMPLEMENTATION ON CNN-BASED IMAGE CODECS
1191

1192 To further validate the generalizability of our method, we conduct supplementary experiments on
1193 cheng20-anchor (Cheng et al., 2020), a CNN-based model featuring an auto-regressive entropy
1194 model. Its CNN-based architecture and moderate complexity (26.6M) provide a sharp contrast to the
1195 low-complexity (TIC (Lu et al., 2021)) and high-complexity (DCAE Lu et al. (2025)) Transformer-
1196 based models from our main experiments, thus offering a rigorous test for our paradigm’s adaptabil-
1197 ity. As shown in Fig. 16, Fig. 17 and Table 10, on the object detection task and instance segmen-
1198 tation task, our combine-ICMH again achieves superior performance with a comparable parameter
1199 count to the baseline ICMH framework. This result provides strong evidence for the robustness and
1200 adaptability of our proposed method across different architectures and entropy models.



1221 Figure 16: Performance comparison for object detection on the COCO2017 dataset using the
1222 cheng2020-anchor.
1223

1224 Table 10: Performance comparison for object detection on the COCO2017 dataset using the
1225 cheng2020-anchor.
1226

Method	Detection		Segmentation		Trainable Params ↓(M)
	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
full fine-tuning	-61.65%	4.69	-73.45%	3.86	26.60 (100%)
ICMH-NET	-8.31%	0.50	-11.67%	0.71	4.43 (16.6%)
channel selection	-11.66%	0.73	-5.16%	0.23	1.34 (4.8%)
Adapt-ICMH	-49.21%	3.12	-60.71%	3.40	0.41 (1.5%)
Ours	-61.45%	4.40	-65.03%	3.55	0.43 (1.6%)

1236 G.2 ABLATION STUDY ON THE POSITION OF THE CMA
1237

1238 We conducted an ablation study on the COCO2017-val dataset to determine the optimal placement
1239 of the CMA module for the object detection task. Our baseline is the Adapt-ICMH framework (at its
1240 first quality level) with the CMA module disabled. We integrated the module at three distinct loca-
1241 tions within the hyperprior network—within the hyper-encoder (h_a), within the hyper-decoder (h_s),
1242 and after the entire network—and evaluated the resulting mAP gain under matched bpp conditions.

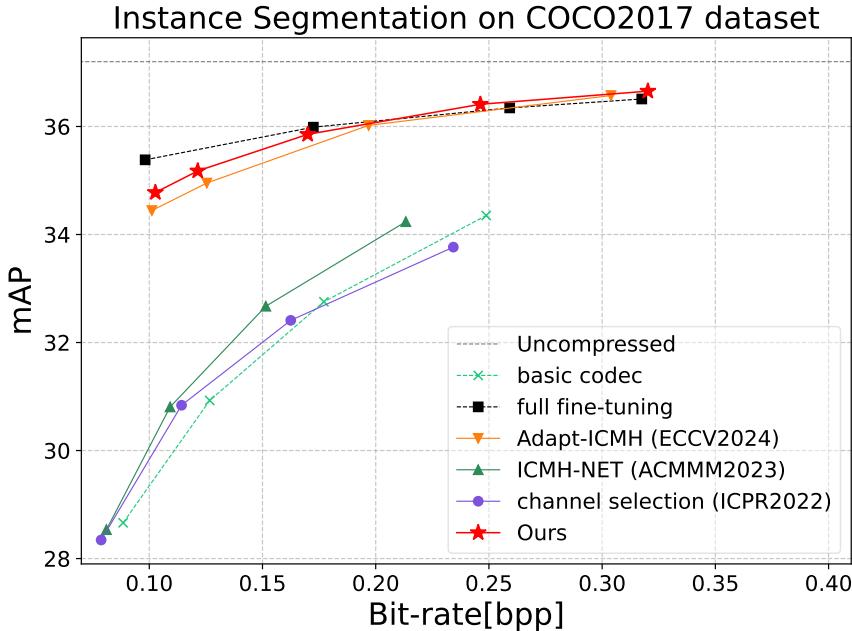


Figure 17: Performance comparison for instance segmentation on the COCO2017 dataset using the cheng2020-anchor.

The results, presented in Table 11. First, they validate the overall effectiveness of the CMA module, as all configurations significantly outperform the baseline. Further comparison indicates that the optimal strategy for maximum performance gain is to place the module after the hyper-prior’s feature extraction module (corresponding to configurations b and e). Interestingly, placing the CMA after the entire hyperprior network also yields substantial improvements. We attribute this enhancement to the module’s ability to directly make adjustments to the final output features of the hyperprior entropy model.

Table 11: Ablations on middle dimension of CMA. Performance comparison for object detection on the COCO2017 dataset.

	h_a	h_s	after h_s	gain	Param(M)
(a)	\	\	\	0.00	0.287(3.8%)
(b)	1,3	0,2	\	0.423	0.305(4.1%)
(c)	1,2,3	0,1,2	\	0.340	0.319(4.2%)
(d)	1,3	1,3	\checkmark	0.545	0.358(4.8%)
(e)	1,3	0,2	\checkmark	0.713	0.343(4.6%)

G.3 ABLATION STUDY ON THE MIDDLE DIMENSION OF THE CMA

To investigate the optimal middle dimension of CMA, we conduct an ablation study on the downstream task of object detection using the COCO2017-val dataset. Our baseline is the Adapt-ICMH framework (at its first quality level) with the CMA module disabled. We then integrate CMA variants with different middle dimensions at the h_a , h_s and evaluate the mAP gain for each configuration under matched bpp conditions.

Table 12 reveals a clear trade-off in the CMA’s middle dimension. Increasing the dimension to 16 substantially boosts performance, but further increases yield negligible gains while the parameter count grows steadily. Thus, we select 16 as the optimal dimension, balancing performance with model complexity.

1296 Notably, a sharp degradation in performance is observed when the middle dimension is increased to
 1297 64. We hypothesize that this is due to over-parameterization, where an excessively complex adapter
 1298 may hinder the fine-tuning process or introduce training instability.
 1299

1300 Table 12: Ablations on middle dimension of CMA. Performance comparison for object detection on
 1301 the COCO2017 dataset.

	middle dimension	gain	Param(M)
(a)	1	0.046	0.289(3.8%)
(b)	16	0.423	0.305(4.1%)
(c)	32	0.354	0.324(4.3%)
(d)	64	-0.047	0.358(4.8%)

G.4 ABLATION STUDY ON THE VARIANT OF THE CMA

1311 We conduct a comparative study of our $GAP \rightarrow Conv1 \times 1 \rightarrow SiLU$ operation within the CMA
 1312 module against two attention mechanisms: ECA (Wang et al., 2020), and MCA (Jiang et al., 2024).
 1313 The experiment was performed on the COCO2017-val dataset. For a fair comparison, all variants
 1314 used the same optimal placement for the CMA (configuration (b) from Table 11), and their perfor-
 1315 mance was measured as the mAP gain over the Adapt-ICMH baseline. As shown in Table 13, the
 1316 results confirm that our proposed operation decisively outperforms both alternatives, highlighting its
 1317 superior design for this task.

1318 Table 13: Ablations on middle dimension of CMA. Performance comparison for object detection on
 1319 the COCO2017 dataset.

Method	gain	Param(M)
Ours	0.423	0.305(4.1%)
ECA	0.073	0.304(4.1%)
MCA	0.269	0.305(4.1%)

REFERENCES

1321
 1322 Wang, Q.; Wu, B.; Zhu, P.; Li, P.; Zuo, W.; and Hu, Q. 2020. ECA-Net: Efficient Channel Attention
 1323 for Deep Convolutional Neural Networks. In *Proc. CVPR*, 11534–11542.
 1324
 1325 Jiang, Y.; Jiang, Z.; Han, L.; Huang, Z.; and Zheng, N. 2024. MCA: Moment Channel Attention
 1326 Networks. In *Proc. AAAI*, 38(3): 2579–2588.
 1327
 1328

H COMPARISON TO OTHER RELATED WORK

1331 While our analysis focuses on single-task baselines, we notice current work in multi-task adaptation,
 1332 such as the multi-path aggregation scheme from Zhang et al. (2024) and the multi task transfer
 1333 techniques from Zhao et al. (2025). Our study, however, follows the single-task, adapter-based
 1334 paradigm (Chen et al., 2023; Li et al., 2024; Park et al., 2025). A direct quantitative comparison is
 1335 not feasible due to the specialized, multi-task loss used in the other works. Nonetheless, adapting
 1336 our analysis for multi-task scenarios remains a valuable avenue for future research.

I THE USE OF LARGE LANGUAGE MODELS

1343 A publicly available large-language model was used exclusively for minor grammar and wording
 1344 corrections after the full manuscript—including all technical content, proofs, and equations—had
 1345 been drafted by the authors. No text, figures, tables, or results were generated by the model. The
 1346 authors manually verified every suggestion and assume full responsibility for the final content. The
 1347 LLM is not listed as an author.