

Image Compression for Machine and Human Vision with Spatial-Frequency Adaptation

Han Li¹, Shaohui Li^{2(✉)}, Shuangrui Ding³, Wenrui Dai^{1(✉)},
Maida Cao¹, Chenglin Li¹, Junni Zou¹, and Hongkai Xiong¹

¹ Shanghai Jiao Tong University

² Tsinghua Shenzhen International Graduate School, Tsinghua University

³ The Chinese University of Hong Kong

{qingshi9974,daiwenrui,caomaida,lcl1985,zoujunni,xionghongkai}@sjtu.edu.cn
lishaohui@sz.tsinghua.edu.cn, ds023@ie.cuhk.edu.hk

Abstract. Image compression for machine and human vision (ICMH) has gained increasing attention in recent years. Existing ICMH methods are limited by high training and storage overheads due to heavy design of task-specific networks. To address this issue, in this paper, we develop a novel lightweight adapter-based tuning framework for ICMH, named Adapt-ICMH, that better balances task performance and bitrates with reduced overheads. We propose a spatial-frequency modulation adapter (SFMA) that simultaneously eliminates non-semantic redundancy with a spatial modulation adapter, and enhances task-relevant frequency components and suppresses task-irrelevant frequency components with a frequency modulation adapter. The proposed adapter is plug-and-play and compatible with almost all existing learned image compression models without compromising the performance of pre-trained models. Experiments demonstrate that Adapt-ICMH consistently outperforms existing ICMH frameworks on various machine vision tasks with fewer fine-tuned parameters and reduced computational complexity. Code will be released at <https://github.com/qingshi9974/ECCV2024-AdpatICMH>.

Keywords: Learned Image Compression · Machine Vision · Adapter

1 Introduction

The booming of computer vision [18, 25, 26, 46, 56, 60, 63, 65] raises the demands for excessive images serving machine vision systems and accomplishing tasks such as classification, detection, segmentation. In practice, images from users have to be compressed and transmitted to exploit the off-the-shelf recognition models deployed in the cloud. However, compression directly using existing image codecs, especially learned image compression (LIC) models [3, 4, 35, 42, 50, 51, 68] optimized for human vision, could cause the loss of semantic information contained in images and dramatically reduce the accuracy of downstream tasks. To meet the demands of both machine and human vision systems, unified codecs

Correspondence to Wenrui Dai and Shaohui Li.

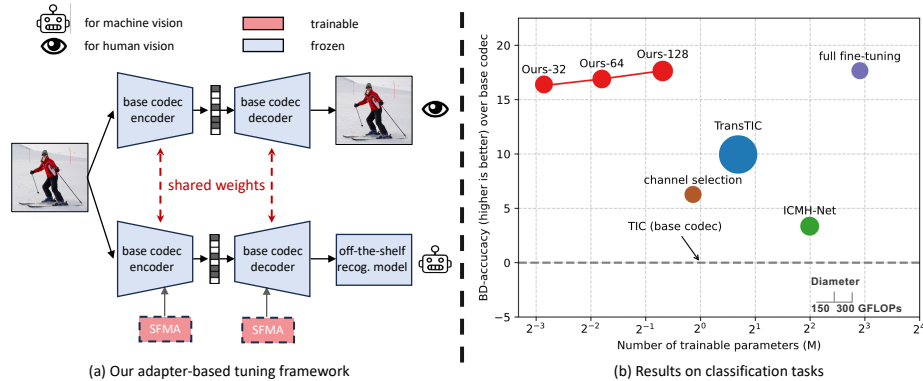


Fig. 1: Left: our adapter-based tuning framework. **Right:** Rate-accuracy performance comparison on classification for ImageNet-*val* [17]. We compare our methods (Ours- n indicates n middle dimensions for SFMA) with full fine-tuning, TransTIC [11], ICMH-Net [44], and channel selection [41]. BD-accuracy is computed by replacing the PSNR in BD-PSNR [6] with top-1 accuracy and setting the base codec of TIC [47] as the anchor. The size of circles indicates GFLOPs for inference during encoding.

to achieve image compression for machine and human vision (ICMH) have been explored [1, 11, 14, 16, 20, 22, 39–41, 43, 44, 64].

A multi-task pipeline has been built for ICMH in [1, 14, 16], where a task-agnostic encoder and multiple task-specific decoders are **jointly trained from scratch** for human vision and multiple machine vision tasks. However, the multi-task pipeline results in degraded rate-distortion (R-D) performance for human vision and is restricted in adapting to newly coming machine vision tasks.

Since LIC models are usually fitted to the distributions of natural images [33, 49], recent works [11, 22, 40, 41, 44] tend to adopt a tuning framework that adapts a **pre-trained** image codec optimized for human vision (called base codec) to diverse machine vision tasks without sacrificing R-D performance. This pipeline enables ICMH by training only a lightweight module specifically for the machine vision task, while freezing the base codec parameters for sharing across human and machine vision. In ICMH-Net [44], a spatial-channel mask generator is trained to determine a subset of quantized latent for a specific machine vision task. Chen *et al.* [11] utilized visual prompt tuning [30] to transfer the transformer-based base codec [47] to machine vision tasks. Unfortunately, these methods suffer from evidently degraded performance compared to codecs fully fine-tuned for machine vision tasks. Besides, they introduce significant overheads of computational and model complexity to the base codec due to their task-specific designs. *It remains a challenge to achieve efficient fine-tuning of pre-trained human vision-oriented image codecs on machine vision tasks.*

In this paper, we consider the fact that almost all spatial and frequency components of an image are exploited by the base codec to improve the visual quality for human vision but could be redundant for machine vision tasks. We are

thereby motivated to modulate the intermediate feature of the base codec in both spatial and frequency domains such that the redundancy in the latent representation can be reduced and the reconstructed image can be efficiently adapted to various machine vision tasks.

To this end, we propose a play-and-plug module of spatial-frequency modulation adapter (SFMA) to better balance the machine vision task performance and transmission cost. Specifically, the proposed SFMA consists of two parallel sub-modules, *i.e.*, a spatial modulation adapter (SMA) that eliminates non-semantic redundancies in the spatial domain and a frequency modulation adapter (FMA) that amplifies the frequency components that contribute more to the machine vision task (*e.g.*, high-frequency components are more crucial for segmentation) and suppresses the redundant frequency components. Furthermore, we develop an adapter-based tuning framework for ICHM, named Adapt-ICMH, by plugging the proposed SFMA into both the encoder and decoder of the base codec, as depicted in Fig. 1. It should be noted that, different from [11] that can only adapt the *transformer-based* base codec, Adapt-ICMH is *architecture-agnostic* such that it is compatible with almost all existing LIC models for adaptation to various machine vision tasks. In summary, our contributions include:

- We propose a spatial-frequency modulation adapter (SFMA) to efficiently update the intermediate feature and better balance the performance and bitrates for diverse machine vision tasks.
- Based on SFMA, we develop a novel tuning framework named Adapt-ICMH to achieve image coding for machine and human vision with a shared base codec, significantly reducing the training and storage overhead.
- Our SFMA is plug-and-play, which could be incorporated with all existing LIC models to achieve ICMH. Experiments show that our method consistently outperforms other ICMH methods in various machine vision tasks with even reduced model and computational complexity.

2 Related Work

Learned Image Compression (LIC) has attracted wide attention due to its superior rate-distortion (R-D) performance. Research in LIC can be categorized into two main paths. The first focuses on enhancing the nonlinear transforms [2], evolving from convolutional neural networks (CNN)-based [3, 4, 50, 51] to attention-based transforms [13], and more recently, transformer-based transforms [35, 42, 47, 68]. The second path involves designing more powerful entropy models. Prior works [3, 4] proposed factorized and hyperprior models, while recent efforts [32, 50, 51, 55] tend to introduce spatial or channel-wise autoregression in the entropy model. However, these methods focus solely on the human visual quality of the decoded image, while ignoring the application in downstream machine vision tasks.

Image Compression for Machine and Human Vision (ICMH) has become a hot topic in recent years due to the growing number of images that need to be processed for various machine vision tasks, such as classification [17, 26, 29] ,

detection [9, 23, 67], segmentation [25, 58, 59, 66] and pose estimation [8, 36, 37, 60]. One straight way is to fully fine-tune the base codec on machine vision tasks. However, this requires storing and deploying a separate copy of the entire codec network for each individual task, which severely hampers practical applicability.

Earlier works [1, 14, 16, 41] presented a multi-task pipeline to unify the image coding for machine and human vision. However, it still requires multiple task-specific decoders, leading to increased storage overhead. In addition, the entire network must be trained from scratch, resulting in significant training overhead and impeding the adaptation to newly coming machine vision tasks. To address this problem, Liu *et al.* [44] proposed to leverage the pre-trained human vision-oriented base codec and train an extra spatial-channel mask generator to transmit a subset of latent for machine vision task. Unfortunately, it cannot be easily extended to autoregression-based entropy models [50, 51] and fails to achieve adaptation of the nonlinear transforms. Chen *et al.* [11] proposed a prompt-based tuning framework that introduces prompts in the transformer-based nonlinear transforms for machine vision. However, it cannot be compatible with the mainstream CNN-based LIC methods [3, 4, 13] and the performance is still highly inferior to full fine-tuning. In contrast, our framework generalizes to almost all existing LIC models with different nonlinear transforms and entropy models and achieves comparable rate-accuracy performance compared to full fine-tuning. See Appendix B for more discussions about ICMH.

Parameter-Efficient Fine-Tuning (PEFT) has been extensively studied for NLP [24, 28, 34, 45, 53, 54]. It can effectively fine-tune a pre-trained large-scale model on downstream tasks or datasets with the major model frozen. With the advent of vision transformers [9, 19, 46], researchers have applied PEFT tools to machine vision tasks, including prompt tuning [30, 31], adapters [10, 12, 62], and LoRA [27]. In the field of learned image compression, Feng *et al.* [21] first introduced prompts for machine-oriented image coding, while Liu *et al.* [11] further explored the application of prompt-tuning for ICMH. [48, 57, 61] transmitted the parameters of additional adapters for content-adaptive optimization [7, 33] to improve the human visual quality of the decoded image, which is completely different from ours in terms of methodology and application scenarios. In addition, naive adapters [10, 28, 40] update the feature without considering the transmission cost of different underlying frequencies and spatial components of the image, which is vital for ICMH. In our paper, we propose a novel spatial-frequency modulation adapter specifically for ICHM to address this limitation.

3 Methods

3.1 Empirical Findings by Full Fine-tuning

We first conduct a simple experiment to highlight the key insight of this paper: image compression for machine vision tends to prioritize transmitting *distinct spatial and frequency information* compared to human vision. This preference also varies across specific machine vision tasks.

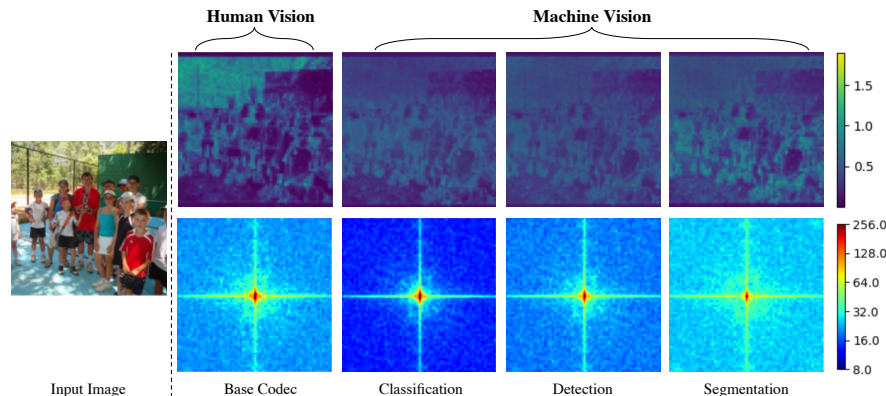


Fig. 2: Visualization of the bit allocation maps (first row) and power spectral density maps (second row) of the latent \hat{y} . The left part shows the raw input image. Each column of the right part corresponds to a codec for each task, including the base codec for human vision and three fine-tuned codecs for machine vision tasks. The bit allocation map is computed by averaging the negative log-likelihood (*i.e.*, $-\log_2 p(\hat{y})$) across channels. The power spectral density map is computed by applying the Fast Fourier Transform (FFT) to \hat{y} with a shift operation to center the zero frequency component, and then averaging its absolute value across channels.

Taking the pre-trained mean-scale hyperprior⁴ model [50] optimized for human vision as the base codec, we fully fine-tune⁵ it on three different machine vision tasks, including classification, object detection, and instance segmentation. Then, we feed the same image to the base codec and three fine-tuned codecs, respectively. The bit allocation map and power spectral density map of latent \hat{y} for each codec are shown in Fig. 2.

From the bit allocation map, we can see that the base codec consumes more bitrates in the complex background region, while the codecs that are fully fine-tuned on machine tasks pay more attention to the semantic object. In addition, we can see from the power spectral density maps that the base codec tends to encode frequency components as many as possible to achieve higher visual quality. However, many of these frequency components are redundant for machine vision. For example, instance segmentation prefers more mid- and high-frequency components for accurate edge and contour reconstruction, while low- and mid-frequency components are sufficient for simpler tasks (*e.g.*, classification).

These observations confirm that image compression has different spatial and frequency domain preferences for different tasks. However, fully fine-tuning the base codec on different machine vision tasks is costly in terms of training time and memory. Thus, in this paper, we propose a spatial-frequency modulation adapter to conduct efficient fine-tuning, and effectively reduce the redundancies in both spatial and frequency domains for machine vision tasks.

⁴ We use the pre-trained model (named *mbt2018-mean*) offered by CompressAI [5].

⁵ We use loss function (1) to conduct the full fine-tuning.

3.2 Framework Overview

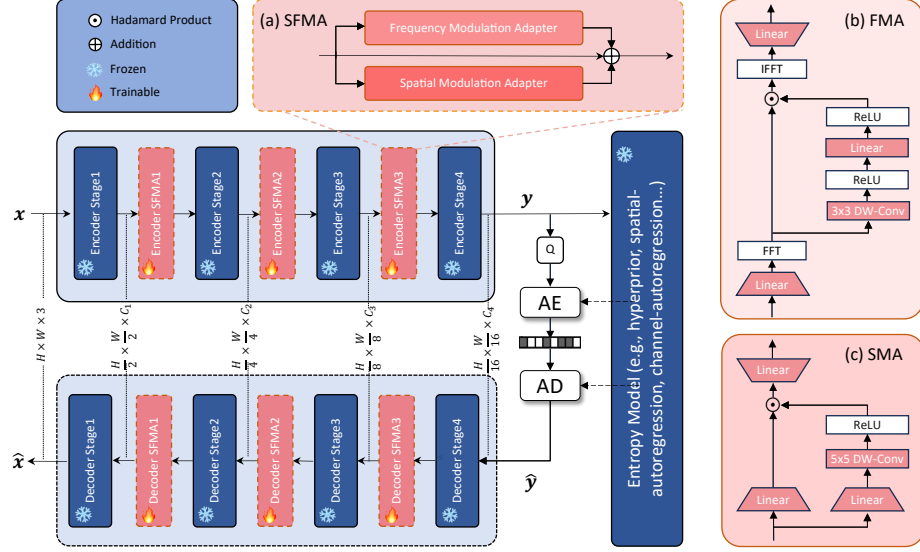


Fig. 3: Overview of our proposed Adapt-ICMH framework. Multiple spatial-frequency modulation adapters (SFMA) are plugged into the encoder g_a and decoder g_s of the base codec for feature adaptation. During the adaptation to the machine vision task, the base codec is frozen and only these adapters are trainable. For brevity, we do not illustrate the specific architecture of the encoder, decoder stage, and entropy model, as it depends on the specific base codec. Please see Appendix E for the detailed architecture.

Our compression framework Adapt-ICMH, shown in Fig. 3, is designed to adapt the existing learned image compression model that is optimized for human vision (*i.e.*, base codec) to downstream machine vision tasks. To this end, we propose the Spatial-Frequency Modulation Adapter (SFMA), a plug-and-play bottleneck module, which is plugged into the encoder g_a and decoder g_s of the base codec for feature adaptation.

During the adaptation stage, the weights of the base codec (blue part) are loaded from the pre-trained checkpoint and kept frozen. Therefore, we only optimize the added adapters (red part) for specific machine vision tasks with the following loss function:

$$\mathcal{L} = \mathcal{R} + \lambda \cdot \mathcal{D}(\mathbf{x}, \hat{\mathbf{x}}; \mathcal{G}), \quad (1)$$

where \mathcal{R} denotes the overall estimated bitrates, \mathcal{D} calculates the task-specific perceptual distortion between raw image \mathbf{x} and reconstructed image $\hat{\mathbf{x}}$ with an off-the-shelf recognition model \mathcal{G} (*e.g.*, Faster RCNN [56] for object detection).

λ is a trade-off term to balance the task performance and bitrates. Please refer to Appendix C for more details about loss function.

After adaptation, we can only modify the weights of lightweight adapters for different machine vision tasks, while keeping the parameters of the base codec unchanged. If an image needs to be compressed for human vision, we can directly remove the adapters to obtain the original base codec, so that the R-D performance for human vision is not affected. In this way, a single base codec model can be used for both human vision and multiple machine vision tasks.

3.3 Spatial-Frequency Modulation Adapter

The proposed Spatial-Frequency Modulation Adapter (SFMA) consists of a Spatial Modulation Adapter (SMA) and a Frequency Modulation Adapter (FMA), which are configured in a parallel way. As shown in Fig. 3, the encoder g_a of base codec generally consists of four stages, denoted by $g_a = g_{a1} \circ g_{a2} \circ g_{a3} \circ g_{a4}$. Given the input image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$, where $H \times W$ is the spatial resolution, we define the output feature of the j -th intermediate stage of g_a as $\mathbf{x}_j \in \mathbb{R}^{C_j \times H_j \times W_j}$, where C_j is the number of channels, $H_j = H/2^j$, $W_j = W/2^j$, and $j \in \{1, 2, 3\}$. Supposing that the SFMA plugged after the j -th encoder stage is denoted by SFMA_j , we can obtain the adapted feature $\tilde{\mathbf{x}}_j$ by aggregating original input feature \mathbf{x}_j with the adapted component. The process can be defined as follows:

$$\begin{aligned} \tilde{\mathbf{x}}_j &= \text{SFMA}_j(\mathbf{x}_j) \\ &= \mathbf{x}_j + \underbrace{\text{FMA}_j(\mathbf{x}_j) + \text{SMA}_j(\mathbf{x}_j)}_{\text{adapted component}}, \end{aligned} \quad (2)$$

where FMA_j and SMA_j denote the j -th frequency and spatial modulation adapter, respectively. The process of the adapter for decoder g_s can be derived similarly.

Frequency Modulation Adapter. Our frequency modulation adapter (FMA) is designed to adapt the feature in the frequency domain. By this modulation operation, we can eliminate frequency redundancies and amplify the important frequency components for machine vision tasks. Specifically, the input $\mathbf{x}_j \in \mathbb{R}^{C_j \times H_j \times W_j}$ firstly be down-projected to a bottleneck middle dimension \hat{C} ($\hat{C} \leq C_j$) by a linear layer $\mathbf{W}_{\text{down}}^f \in \mathbb{R}^{C_j \times \hat{C}}$ and then be transformed into frequency domain by fast Fourier transform (FFT), obtaining the middle feature \mathbf{x}_j^f . Subsequently, we generate a frequency modulation matrix \mathbf{m}_j^f and then use Hadamard product to calculate the modulated result $\bar{\mathbf{x}}_j^f$. Finally, $\bar{\mathbf{x}}_j^f$ is inversely transformed using inverse FFT (IFFT) and up-projected by another linear layer $\mathbf{W}_{\text{up}}^f \in \mathbb{R}^{\hat{C} \times C_j}$. The overall process of FMA can be formulated as follows:

$$\begin{aligned} \mathbf{x}_j^f &= \mathcal{F}(\mathbf{x}_j \cdot \mathbf{W}_{\text{down}}^f), \\ \mathbf{m}_j^f &= \sigma(\text{DW-Conv}_{3 \times 3}(\mathbf{x}_j^f)) \cdot \mathbf{W}_{\text{middle}}^f, \\ \bar{\mathbf{x}}_j^f &= \mathbf{x}_j^f \odot \sigma(\mathbf{m}_j^f), \\ \tilde{\mathbf{x}}_j^f &= \mathcal{F}^{-1}(\bar{\mathbf{x}}_j^f) \cdot \mathbf{W}_{\text{up}}^f \end{aligned} \quad (3)$$

where σ denotes ReLU [52], DW-Conv $_{3 \times 3}$ denotes depth-wise convolution [15] with kernel size 3×3 , $\mathbf{W}_{\text{middle}}^f \in \mathbb{R}^{\hat{C} \times \hat{C}}$ is a linear layer, \odot denotes Hadamard product, and $\mathcal{F}(\cdot)$ and $\mathcal{F}^{-1}(\cdot)$ denote the FFT and IFFT, respectively.

Spatial Modulation Adapter. Our spatial modulation adapter (SMA) is designed to adapt features in the spatial domain, thus guiding the codec to focus more on the semantic region and reduce the spatial redundancies in the background. Similarly, the input $\mathbf{x}_j \in \mathbb{R}^{C_j \times H_j \times W_j}$ is first down-projected to a bottleneck middle dimension \hat{C} by a linear layer $\mathbf{W}_{\text{down1}}^s \in \mathbb{R}^{C_j \times \hat{C}}$, obtaining the middle feature \mathbf{x}_j^s . Then, we produce the spatial modulation matrix \mathbf{m}_j^s by feeding \mathbf{x}_j to another down-projection linear layer $\mathbf{W}_{\text{down2}}^s \in \mathbb{R}^{C_j \times \hat{C}}$ followed by a 5×5 depth-wise convolution layer. We then apply an up-projection linear layer $\mathbf{W}_{\text{up}}^s \in \mathbb{R}^{\hat{C} \times C_j}$ to the modulated feature to get the spatial adapted result. The overall process of SMA can be formulated as follows:

$$\begin{aligned} \mathbf{x}_j^s &= \mathbf{x}_j \cdot \mathbf{W}_{\text{down1}}^s, \\ \mathbf{m}_j^s &= \text{DW-Conv}_{5 \times 5}(\mathbf{x}_j^s \cdot \mathbf{W}_{\text{down2}}^s), \\ \bar{\mathbf{x}}_j^s &= \mathbf{x}_j^s \odot \sigma(\mathbf{m}_j^s), \\ \tilde{\mathbf{x}}_j^s &= \bar{\mathbf{x}}_j^s \cdot \mathbf{W}_{\text{up}}^s, \end{aligned} \tag{4}$$

4 Experiments

4.1 Training Details and Datasets.

To demonstrate the generalization ability of our framework, we adopt three mainstream LIC methods, *i.e.*, stacked convolution-based [50], residual network-based [13], and transformer-based [47] image compression model. We use their pre-trained models from *CompressAI* [5] library as base codes, denoted by *mbt2018-mean*, *cheng2020-anchor* and *Lu2022-TIC*, respectively. For benchmark machine vision tasks, we evaluate our framework on classification, object detection, and instance segmentation, respectively. By freezing these base codecs pre-trained for human vision, we then train our SFMAs using the task-specific loss of Eq. (1) with the bottleneck middle dimension \hat{C} set to 64. Specifically, we train the classification task on ImageNet-*train* dataset [17] for 8 epochs with a batch size of 16, and train the object detection and instance segmentation tasks on COCO2017-*train* dataset [38] both for 40 epochs with a batch size of 8. The images are randomly cropped and resized to 256×256 for training all tasks. For the calculation of task-specific perceptual distortion term in the Eq. (1), we follow [11] to use the pre-trained ResNet50 [26], Faster R-CNN [56] and Mask R-CNN [25] for classification, object detection, and instance segmentation, respectively. More training details are included in Appendix C.

4.2 Evaluation

We follow the setting of [11] to evaluate the performance of our method. Specifically, we use bits per pixel (BPP) to evaluate bitrates for compression. For image

classification, the evaluations are conducted on the ImageNet-*val* dataset [17] and images are resized and center cropped to 256×256 . We adopt the pre-trained ResNet50 from *torchvision* library as the off-the-shelf recognition model and use top-1 accuracy as the quality metric. For objection detection and instance segmentation, the evaluations are both done on the COCO2017-*val* dataset [38] by using the pre-trained Faster R-CNN and Mask R-CNN from *detectron2* library as the off-the-shelf recognition model, respectively. The test images are resized to 800 pixels for the short side and we use mean average precision (mAP) with an Intersection of Union (IoU) threshold of 0.5 as the quality metric.

4.3 Rate-Accuracy Comparison

We compare the rate-accuracy performance of our methods with the state-of-the-art (SOTA) methods⁶, including *TransTIC* [11], *ICMH-Net* [44], and *channel selection* [41]. Note that our method has an equal R-D performance for human vision with these methods, since we all adopt the pre-trained image codec as base codec and freeze it. We also evaluate the rate-accuracy performance of the *base codec* and further *full fine-tuning* it on machine vision tasks end-to-end.

The rate-accuracy curves for competing methods are shown Fig. 4. In addition, we also summarize the average bitrates savings, accuracy improvements, and the trainable parameters count of different methods in Tab. 1. Specifically, we report the BD-rate [6] results to quantify the average bitrates savings with equal task accuracy, we also follow [11] to report BD-acc/BD-mAP results to measure the average task accuracy improvements with equal bitrates. BD-acc/BD-mAP is calculated by replacing the PSNR in BD-PSNR with top-1 accuracy/mAP.

From the results of Fig. 4 and Tab. 1, we can observe that **(1)** our method consistently achieves excellent performance across all three machine vision tasks and outperforms [11, 41, 44], the superiority is more significant on the classification task; **(2)** Our method demonstrates strong generalization ability as it can effectively generalize to three different base codecs, despite their varying architectures of nonlinear transforms (i.e., g_a and g_s) and entropy models. **(3)** our method achieve comparable or even superior performance to *full fine-tuning*. However, *full fine-tuning* requires storing and deploying a separate copy of the entire codec for each task, significantly increasing training and storage overhead and severely hindering practical applicability. **(4)** our method requires fewer trainable parameters (i.e., about 1 ~ 5% of the number of parameters of the base codec) compared to other tuning-based methods (i.e., about 1 ~ 5% of the number of parameters of the base codec), which demonstrates the parameter efficiency of our method. For the application on larger transformer-based image codecs [35, 42, 68], see Appendix H. For the comparison of computational complexity, see Appendix G.

⁶ Since ICMH-Net [44] and channel selection [41] are not open source and use different recognition models and test datasets than ours, we reproduce them with the same training settings as ours. See Appendix D for more details.

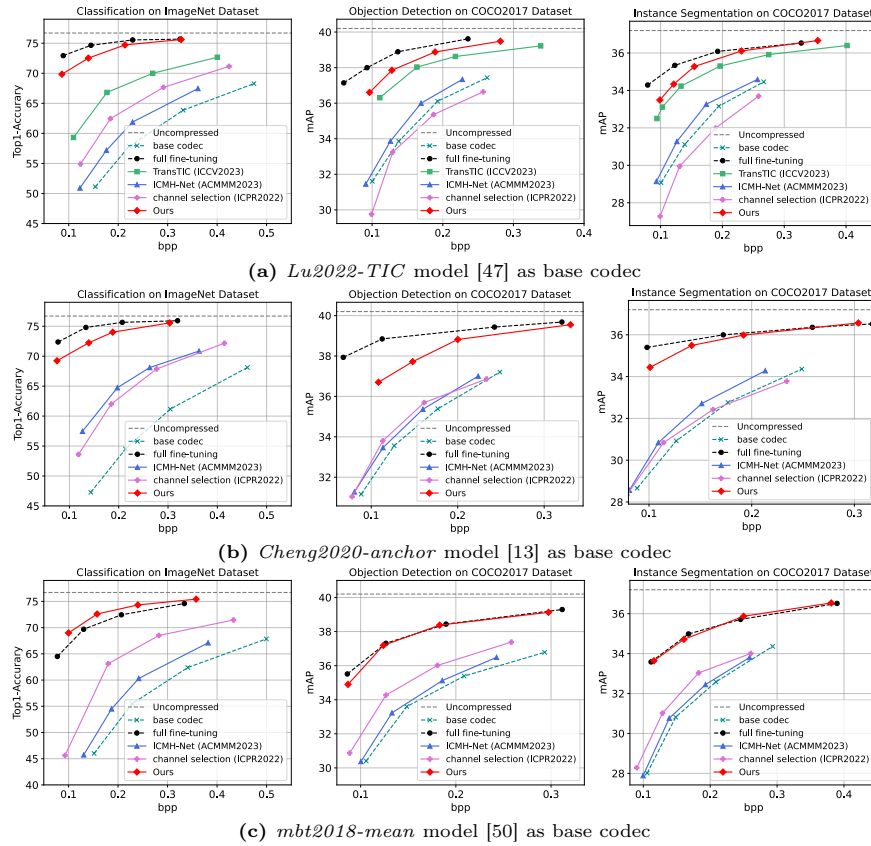


Fig. 4: Rate-Accuracy performance comparison under different machine vision tasks and different base codecs.

4.4 Ablation study

We conduct ablation studies to investigate how each design in our SFMA influences the overall performance. Without loss of generality, we use the pre-trained *mbt2018-mean* as the base codec for convenience.

Effect on the architecture of SFMA. We first explore the impact of the architecture of SFMA on performance by conducting experiments with various variants of SFMA, as summarized in Fig. 5. **(a) Proposed SFMA:** The default configuration of SFMA. **(b) SMA-only:** The FMA branch is removed, retaining only the SMA branch. **(c) FMA-only:** The SMA branch is removed, retaining only the FMA branch. **(d) FMA-SMA-sequential:** FMA and SMA are organized sequentially, where SMA is followed by FMA. **(e) SMA-FMA-sequential:** FMA and SMA are organized sequentially, where FMA is followed by SMA. **(f) SMA-SMA-parallel:** two SMAs are configured in parallel.

Table 1: Comparison of rate-accuracy performance and the number of trainable parameters under different machine vision tasks and different base codecs. BD-rate, BD-acc, BD-mAP are used to quantitatively measure the rate-accuracy performance with the base codec as anchor, respectively. Arrows indicate whether lower is better (\downarrow) or higher is better (\uparrow). (-) indicates the BD-rate can not be computed.

Method	Classification		Detection		Segmentation		Trainable Params \downarrow (M)
	BD-rate \downarrow	BD-acc \uparrow	BD-rate \downarrow	BD-mAP \uparrow	BD-rate \downarrow	BD-mAP \uparrow	
<i>Lu2022-TIC</i>							
full fine-tuning	-	17.68	-73.94%	4.51	-67.98%	3.76	7.51(100%)
TransTIC [11]	-58.32%	9.96	-46.30%	2.77	-46.41%	2.70	1.61(21.4%)
ICMH-Net [44]	-18.75%	3.36	-9.07%	0.63	-10.77%	0.65	3.98(53.0%)
channel selection [41]	-37.17%	6.28	6.84%	-0.55	16.51%	-0.94	0.91(12.1%)
Ours	-88.57%	16.90	-55.14%	3.55	-52.40%	3.21	0.28(3.7%)
<i>Cheng2020-anchor</i>							
full fine-tuning	-	21.04	-59.20%	4.69	-76.03%	3.86	26.60(100%)
ICMH-Net [44]	-47.46%	10.4	-8.81%	0.53	-12.18%	0.74	4.43(16.6%)
channel selection [41]	-41.58%	8.77	-11.66%	0.73	5.22%	0.22	1.34(4.8%)
Ours	-87.56%	20.27	-49.34%	3.13	-59.90%	3.48	0.41(1.5%)
<i>mbt2018-mean</i>							
full fine-tuning	-79.92%	17.97	-60.87%	3.87	-56.64%	3.28	7.03(100%)
ICMH-Net [44]	-15.99%	3.55	-6.02%	0.39	-4.97%	0.31	3.98(56.6%)
channel selection [41]	-41.30%	9.90	-23.07%	1.52	-15.09%	1.05	0.91(12.9%)
Ours	-82.00%	18.71	-56.17%	3.84	-52.65%	3.17	0.28(4.1%)




Fig. 5: Different variants of SFMA: (a) Proposed SFMA. (b) SMA-only.(c) FMA-only. (d) FMA-SMA-sequential. (e) SMA-FMA-sequential. (f) SMA-SMA-parallel.

The results are shown in Tab. 2. We can observe that **(1)** removing either SMA or FMA leads to performance degradation, and FMA plays a more important role than SMA in the adaptation process, showing the superiority of frequency domain adaptation. **(2)** parallel configuration outperforms the sequential configuration. **(3)** barely replicating SMA has more trainable parameters compared to our SFMA, but leads to a performance degradation. This shows that the improvement of our SFMA is not only due to an overall increase in model capacity but rather to its specific structure.

Effect on the position of SFMA. We then investigate the impact of the position to plug the SFMA. The results are shown in the Table 3. From shallow to deep, we first plug SFMA after different stages of encoder g_a and decoder g_s . We observe that the SFMA at the shallow layer can benefit more for the rate-accuracy performance, *i.e.*, plugging SFMA after stage1 of g_a and g_s can achieve 69.54% bitrates savings for classification compared to the base codec, while plugging SFMA after stage3 of g_a and g_s can only achieve 64.76% bitrates savings. In addition, employing more SFMA can bring further benefits.

Table 2: Ablations on different variants of SFMA

Method	Classification		Detection		Segmentation		Params (M)
	BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
(a)	-82.00%	18.71	-56.17%	3.84	-52.65%	3.17	0.28
(b)	-73.82%	16.30	-51.94%	3.61	-48.16%	2.92	0.16
(c)	-77.40%	17.29	-52.86%	3.32	-49.58%	2.90	0.12
(d)	-80.49%	18.67	-56.87%	3.81	-51.79%	3.14	0.28
(e)	-83.31%	18.50	-53.77%	3.79	-51.28%	3.02	0.28
(f)	-79.63%	17.59	-54.91%	3.62	-50.96%	3.04	0.32

Table 3: Ablations on the position of SFMA.

Stage			g_a	g_s	Classification		Detection		Segmentation		Params (M)
1	2	3			BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
✓			✓	✓	-69.54%	15.95	-50.72%	3.39	-48.27%	2.92	0.09
	✓		✓	✓	-69.00%	15.69	-49.96%	3.24	-46.40%	2.82	0.09
		✓	✓	✓	-64.76%	13.29	-50.08%	3.17	-39.98%	2.43	0.09
✓	✓	✓	✓	✓	-82.00%	18.71	-56.17%	3.84	-52.65%	3.17	0.28
✓	✓	✓	✓		-64.75%	13.28	-38.84%	2.88	-38.67%	2.28	0.14
✓	✓	✓	✓	✓	-61.87%	12.40	-44.05%	3.01	-40.43%	2.66	0.14

Then, we observe that SFMAs are indispensable at both the encoder and decoder side, such that removing either would lead to a significant degradation in performance. SFMA in the encoder tends to reduce redundancies in the feature for machine vision tasks. SFMA in the decoder tends to adapt the feature so that the reconstructed image is more suitable for downstream tasks.

Effect on the middle dimension \hat{C} . Tab. 4 reports how middle dimension \hat{C} impacts the rate-accuracy performance and complexity of SFMA. The results show that enlarging the middle dimension from 1 to 128 can boost performance while leading to a higher training complexity. Thus, we set \hat{C} to 64 to make a better trade-off between performance and complexity.

4.5 Qualitative Results

Fig. 6 shows the qualitative results of our Adapt-ICMH and other methods for instance segmentation. Our approach achieves superior segmentation results

Table 4: Ablations on the middle dimension \hat{C} of SFMA.

\hat{C}	Classification		Detection		Segmentation		Params (M)
	BD-rate↓	BD-acc↑	BD-rate↓	BD-mAP↑	BD-rate↓	BD-mAP↑	
1	-60.41%	12.58	-48.58%	2.93	-45.24%	2.40	0.005
32	-78.74%	17.52	-53.61%	3.53	-48.50%	2.92	0.14
64	-82.00%	18.71	-56.17%	3.84	-52.65%	3.17	0.28
128	-83.61%	19.50	-63.28%	4.11	-60.50%	3.76	0.62

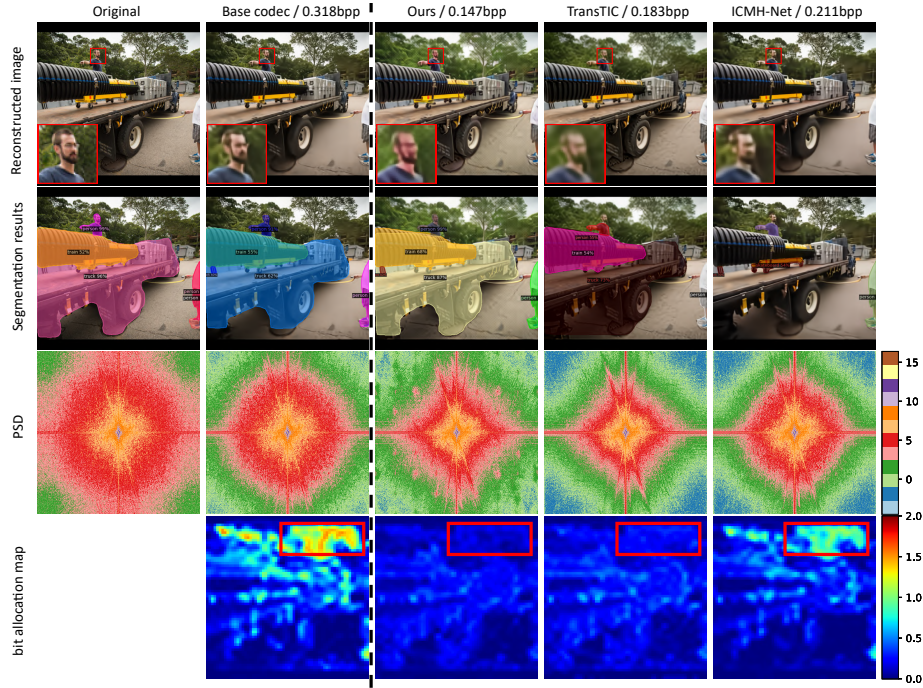


Fig. 6: Qualitative comparison of our Adapt-ICMH with other ICMH methods. **First row:** The original image and decoded image of each method. We show the decoded images for machine vision of three ICMH methods (left). **Second row:** The instance segmentation results of each image. **Third row:** The log power spectral density maps of each image. **Bottom row:** The bit allocation maps for \hat{y} of each method.

with a lower bitrate (0.147bpp). In comparison to other methods, Adapt-ICMH effectively minimizes the bitrates of the background while enhancing the high-frequency edge details of the semantic object. We also observe that our method’s decoded image for machine vision exhibits some distortion in terms of overall chromaticity and brightness. However, these low-frequency distortions do not affect the segmentation result, which supports the goal of our SFMA.

4.6 Towards scalable coding for machine and human vision.

Scalable coding for machine and human vision is of practical significance to adapt the transmission among diverse demands. The framework of Fig 3, followed by [11, 41], is designed to produce a single bitstream tailored for each individual task. However, our proposed SFMA can also benefit existing scalable coding methods to achieve better scalable coding performance for machine and human vision. To demonstrate it, we integrate our SFMA and the spatial-channel mask generator proposed by *ICMH-Net* [44] with the base codec. Specifically, we only

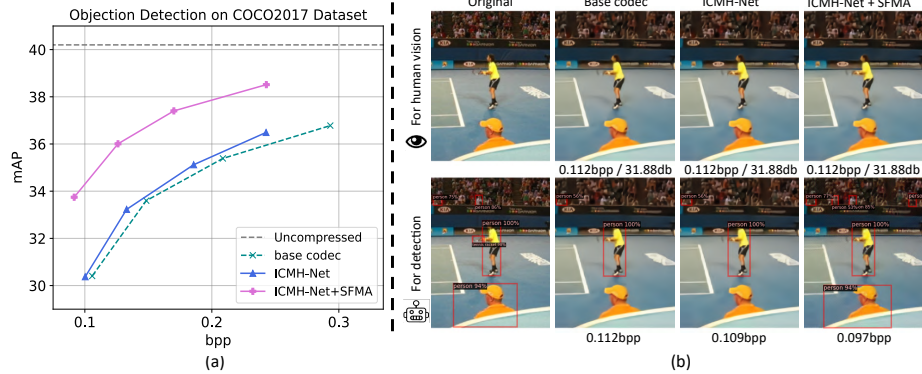


Fig. 7: Quantitative and qualitative comparison of each method for the object detection task. Pre-trained *mbt2018-mean* [50] is used as the base codec. We present the bpp/PSNR results for the decoded image for human vision and the bpp results for the decoded image for the detection task.

plug our SFMA into the decoder g_s and utilize spatial-channel mask generator to produce two bitstreams, the truncated one is decoded by the g_s with SFMAs to reconstruct image for machine vision task, while the full bitstream is decoded by the g_s without SFMAs to reconstruct image for human vision. See Appendix F for detailed network architecture for scalable coding.

Fig. 7 show the comparison of rate-accuracy performance and qualitative results. It demonstrates the superior performance of our SFMA, which significantly improves the rate-accuracy performance of *ICMH-Net*. Note that we do not alter the base codec, so the visual quality of the image reconstructed from the full bitstream is the same as the base codec.

5 Conclusion

In this paper, we develop a novel adapter-based tuning framework for ICMH, named Adapt-ICMH. Specifically, we fine-tune the pre-trained image codec optimized for human vision on multiple machine vision tasks, which is achieved by the proposed spatial-frequency modulation adapter (SFMA). Our SFMA can efficiently reduce the latent redundancies in both spatial and frequency domains, and effectively adapt the decoded image for downstream machine vision tasks. Our method is plug-and-play and compatible with all existing learned image compression models. Experiments show that our method consistently outperforms other ICMH methods in various machine vision tasks, even with a reduced number of trainable parameters and computational complexity. In the future work, we will extend our work to video coding for machines and humans.

6 Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 62125109, Grant 61931023, Grant 61932022, Grant 62371288, Grant 62320106003, Grant 62301299, Grant T2122024, Grant 62120106007, China Postdoctoral Innovation Talents Support Program (BX20230184).

References

1. Bai, Y., Yang, X., Liu, X., Jiang, J., Wang, Y., Ji, X., Gao, W.: Towards end-to-end image compression and analysis with transformers. In: AAAI. vol. 36, pp. 104–112 (2022)
2. Ballé, J., Chou, P.A., Minnen, D., Singh, S., Johnston, N., Agustsson, E., Hwang, S.J., Toderici, G.: Nonlinear transform coding. *IEEE JSTSP* **15**(2), 339–353 (2020)
3. Ballé, J., Laparra, V., Simoncelli, E.P.: End-to-end optimized image compression. In: ICLR (2016)
4. Ballé, J., Minnen, D., Singh, S., Hwang, S.J., Johnston, N.: Variational image compression with a scale hyperprior. In: ICLR (2018)
5. Bégaint, J., Racapé, F., Feltman, S., Pushparaja, A.: Compressai: a pytorch library and evaluation platform for end-to-end compression research. arXiv preprint arXiv:2011.03029 (2020)
6. Bjontegaard, G.: Calculation of average psnr differences between rd-curves. In: VCEG-M33 (2001)
7. Campos, J., Meierhans, S., Djelouah, A., Schroers, C.: Content adaptive optimization for neural image compression. In: CVPRW (2019)
8. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7291–7299 (2017)
9. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. pp. 213–229. Springer (2020)
10. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. In: NeurIPS. vol. 35, pp. 16664–16678 (2022)
11. Chen, Y.H., Weng, Y.C., Kao, C.H., Chien, C., Chiu, W.C., Peng, W.H.: Transtic: Transferring transformer-based image compression from human perception to machine perception. In: ICCV. pp. 23297–23307 (2023)
12. Chen, Z., Duan, Y., Wang, W., He, J., Lu, T., Dai, J., Qiao, Y.: Vision transformer adapter for dense predictions. In: ICLR (2023)
13. Cheng, Z., Sun, H., Takeuchi, M., Katto, J.: Learned image compression with discretized gaussian mixture likelihoods and attention modules. In: CVPR. pp. 7939–7948 (2020)
14. Choi, H., Bajić, I.V.: Scalable image coding for humans and machines. *IEEE TIP* **31**, 2739–2754 (2022)
15. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: CVPR. pp. 1251–1258 (2017)
16. Codevilla, F., Simard, J.G., Goroshin, R., Pal, C.: Learned image compression for machine perception. arXiv preprint arXiv:2111.02249 (2021)
17. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255 (2009)

18. Ding, S., Li, M., Yang, T., Qian, R., Xu, H., Chen, Q., Wang, J., Xiong, H.: Motion-aware contrastive video representation learning via foreground-background merging. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9716–9726 (2022)
19. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2020)
20. Feng, R., Gao, Y., Jin, X., Feng, R., Chen, Z.: Semantically structured image compression via irregular group-based decoupling. In: ICCV (2023)
21. Feng, R., Liu, J., Jin, X., Pan, X., Sun, H., Chen, Z.: Prompt-icm: A unified framework towards image coding for machines with task-driven prompts. arXiv preprint arXiv:2305.02578 (2023)
22. Fischer, K., Brand, F., Kaup, A.: Boosting neural image compression for machines using latent space masking. IEEE TCSVT (2022)
23. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
24. He, J., Zhou, C., Ma, X., Berg-Kirkpatrick, T., Neubig, G.: Towards a unified view of parameter-efficient transfer learning. In: ICLR (2022)
25. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
26. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
27. He, X., Li, C., Zhang, P., Yang, J., Wang, X.E.: Parameter-efficient model adaptation for vision transformers. In: AAAI. vol. 37, pp. 817–825 (2023)
28. Houshy, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: ICML. pp. 2790–2799. PMLR (2019)
29. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
30. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: ECCV. pp. 709–727. Springer (2022)
31. Khattak, M.U., Rasheed, H., Maaz, M., Khan, S., Khan, F.S.: Maple: Multi-modal prompt learning. In: CVPR. pp. 19113–19122 (2023)
32. Koyuncu, A.B., Gao, H., Boev, A., Gaikov, G., Alshina, E., Steinbach, E.: Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. In: ECCV (2022)
33. Lee, J., Cho, S., Beack, S.K.: Context-adaptive entropy model for end-to-end optimized image compression. In: ICLR (2019)
34. Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. In: EMNLP (2021)
35. Li, H., Li, S., Dai, W., Li, C., Zou, J., Xiong, H.: Frequency-aware transformer for learned image compression. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=HKGQDDTuvZ>
36. Li, H., Shi, B., Dai, W., Chen, Y., Wang, B., Sun, Y., Guo, M., Li, C., Zou, J., Xiong, H.: Hierarchical graph networks for 3d human pose estimation. arXiv preprint arXiv:2111.11927 (2021)
37. Li, H., Shi, B., Dai, W., Zheng, H., Wang, B., Sun, Y., Guo, M., Li, C., Zou, J., Xiong, H.: Pose-oriented transformer with uncertainty-guided refinement for 2d-to-3d human pose estimation. In: AAAI. vol. 37, pp. 1296–1304 (2023)

38. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
39. Liu, J., Feng, R., Qi, Y., Chen, Q., Chen, Z., Zeng, W., Jin, X.: Rate-distortion-cognition controllable versatile neural image compression. In: ECCV. Springer (2024)
40. Liu, J., Jin, X., Feng, R., Chen, Z., Zeng, W.: Composable image coding for machine via task-oriented internal adaptor and external prior. In: VCIP. pp. 1–5 (2023)
41. Liu, J., Sun, H., Katto, J.: Improving multiple machine vision tasks in the compressed domain. In: ICPR. pp. 331–337. IEEE (2022)
42. Liu, J., Sun, H., Katto, J.: Learned image compression with mixed transformer-cnn architectures. In: CVPR. pp. 14388–14397 (2023)
43. Liu, K., Liu, D., Li, L., Yan, N., Li, H.: Semantics-to-signal scalable image compression with learned reversible representations. *IJCV* **129**(9), 2605–2621 (2021)
44. Liu, L., Hu, Z., Chen, Z., Xu, D.: Icmh-net: Neural image compression towards both machine vision and human vision. In: ACM MM. pp. 8047–8056 (2023)
45. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G.: Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys* **55**(9), 1–35 (2023)
46. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV. pp. 10012–10022 (2021)
47. Lu, M., Guo, P., Shi, H., Cao, C., Ma, Z.: Transformer-based image compression. In: DCC. pp. 469–469 (2022)
48. Lv, Y., Xiang, J., Zhang, J., Yang, W., Han, X., Yang, W.: Dynamic low-rank instance adaptation for universal neural image compression. In: ACM MM. pp. 632–642 (2023)
49. Mentzer, F., Agustsson, E., Tschannen, M., Timofte, R., Gool, L.V.: Practical full resolution learned lossless image compression. In: CVPR. pp. 10629–10638 (2019)
50. Minnen, D., Ballé, J., Toderici, G.D.: Joint autoregressive and hierarchical priors for learned image compression. In: NeurIPS. vol. 31 (2018)
51. Minnen, D., Singh, S.: Channel-wise autoregressive entropy models for learned image compression. In: ICIP. pp. 3339–3343 (2020)
52. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: ICML. pp. 807–814 (2010)
53. Pfeiffer, J., Kamath, A., Rücklé, A., Cho, K., Gurevych, I.: Adapterfusion: Non-destructive task composition for transfer learning. arXiv preprint arXiv:2005.00247 (2020)
54. Pfeiffer, J., Rücklé, A., Poth, C., Kamath, A., Vulić, I., Ruder, S., Cho, K., Gurevych, I.: Adapterhub: A framework for adapting transformers. In: EMNLP. pp. 46–54 (2020)
55. Qian, Y., Sun, X., Lin, M., Tan, Z., Jin, R.: Entroformer: A transformer-based entropy model for learned image compression. In: ICLR (2022)
56. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS. vol. 28 (2015)
57. Shen, S., Yue, H., Yang, J.: Dec-adapter: Exploring efficient decoder-side adapter for bridging screen content and natural image compression. In: CVPR. pp. 12887–12896 (2023)
58. Shi, B., Jiang, D., Zhang, X., Li, H., Dai, W., Zou, J., Xiong, H., Tian, Q.: A transformer-based decoder for semantic segmentation with multi-level context mining. In: European Conference on Computer Vision. pp. 624–639. Springer (2022)

59. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7262–7272 (2021)
60. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: CVPR. pp. 5693–5703 (2019)
61. Tsubota, K., Akutsu, H., Aizawa, K.: Universal deep image compression via content-adaptive optimization with adapters. In: WACV. pp. 2529–2538 (2023)
62. Wang, Y., Shi, B., Zhang, X., Li, J., Liu, Y., Dai, W., Li, C., Xiong, H., Tian, Q.: Adapting shortcut with normalizing flow: An efficient tuning framework for visual recognition. In: CVPR. pp. 15965–15974 (2023)
63. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. In: NeurIPS. vol. 34, pp. 12077–12090 (2021)
64. Yang, S., Hu, Y., Yang, W., Duan, L.Y., Liu, J.: Towards coding for human and machine vision: Scalable face image coding. *IEEE TMM* **23**, 2957–2971 (2021)
65. Zheng, H., Zhou, L., Li, H., Su, J., Wei, X., Xu, X.: Bem: Balanced and entropy-based mix for long-tailed semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 22893–22903 (June 2024)
66. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6881–6890 (2021)
67. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)
68. Zou, R., Song, C., Zhang, Z.: The devil is in the details: Window-based attention for image compression. In: CVPR. pp. 17492–17501 (2022)