

Unified Coding for Both Human Perception and Generalized Machine Analytics with CLIP Supervision

Kangsheng Yin^{1,3*}, Quan Liu^{1,3*}, Xuelin Shen^{1†}, Yulin He¹, Wenhan Yang^{2†}, Shiqi Wang⁴

¹Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ)

²Peng Cheng Laboratory

³Shenzhen University

⁴City University of Hong Kong

Abstract

The image compression model has long struggled with adaptability and generalization, as the decoded bitstream typically serves only human or machine needs and fails to preserve information for unseen visual tasks. Therefore, this paper innovatively introduces supervision obtained from multimodal pre-training models and incorporates adaptive multi-objective optimization tailored to support both human visual perception and machine vision simultaneously with a single bitstream, denoted as Unified and Generalized Image Coding for Machine (UG-ICM). Specifically, to get rid of the reliance between compression models with downstream task supervision, we introduce Contrastive Language-Image Pre-training (CLIP) models into the training constraint for improved generalization. Global-to-instance-wise CLIP supervision is applied to help obtain hierarchical semantics that make models more generalizable for the tasks relying on the information of different granularity. Furthermore, for supporting both human and machine visions with only a unifying bitstream, we incorporate a conditional decoding strategy that takes as conditions human or machine preferences, enabling the bitstream to be decoded into different versions for corresponding preferences. As such, our proposed UG-ICM is fully trained in a self-supervised manner, *i.e.*, without awareness of any specific downstream models and tasks. The extensive experiments have shown that the proposed UG-ICM is capable of achieving remarkable improvements in various unseen machine analytics tasks, while simultaneously providing perceptually satisfying images.

Code — <https://github.com/YinKangSheng/UG-ICM>

Introduction

Recent years have witnessed the rise of machine-vision-oriented application scenarios, *e.g.* autonomous driving, smart cities, and the Internet of Things. Existing technologies need to shift towards meeting the demands of machine vision. Therefore, on the coding side, the demand for machine-oriented coding also emerges. Some preliminary attempts, *e.g.*, feature coding (Chen et al. 2020, 2019), are implemented to compress the intermediate deep features,

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

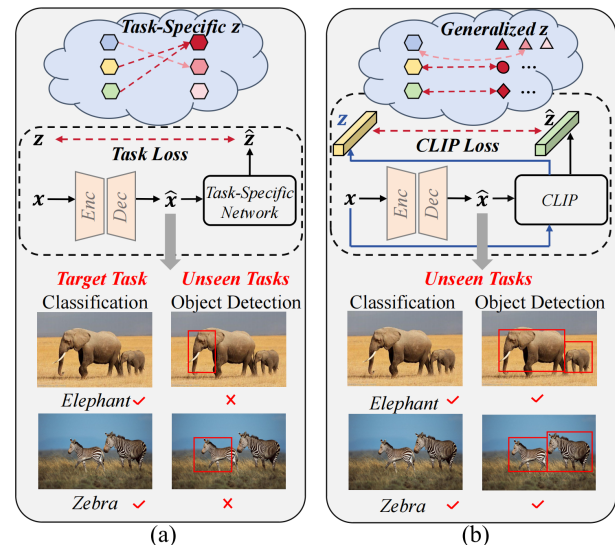


Figure 1: (a) Most existing ICM models rely on task-specific supervision and fall short of generalization capacity in unseen scenarios. (b) z and \hat{z} denote the latent representations satisfying downstream tasks from the original image x and compressed image \hat{x} , respectively.

which satisfies only the machine vision needs by totally discarding low-level information.

To make up this issue for obtaining both human and machine-friendly images, Image Coding for Machine (ICM) is taken as a more promising direction with significant attention (Yang et al. 2024). Preliminary ICM explorations focus on conventional block-based codecs (Sullivan et al. 2012; Bross et al. 2021), making them aware of machine vision preferences by identifying semantic-related side information during preprocessing (Huang et al. 2021; Choi and Han 2020). This information contributes to deriving appropriate block-wise coding parameters for analytics-friendly bit allocation. In recent years, notable progress has been made on Learned Image Compression (LIC), providing the flexibility to integrate entire codec modules specifically tailored for machines. Such codecs are jointly optimized under the supervision of downstream machine analytics models with respect to *rate-perception-analytics* criteria (Chamain

et al. 2021; Li et al. 2022), which offers improved machine analytics performance. However, the inherent gap between the representations of pixels and semantic features makes it challenging to achieve a good trade-off between perceptual quality and machine analytic performance (Liu, Zhang, and Xiong 2019). To pursue enhanced machine analytics without sacrificing the perceptual quality, transfer-based methods (Liu, Sun, and Katto 2022; Chen et al. 2023) have emerged as a new research trend, *i.e.*, transferring human vision-oriented LIC to machine analytic purpose in a plug-and-play manner. However, these methods take separate bitstreams for different tasks, which falls short in terms of flexibility and overall compression efficiency. Moreover, their generalization has long been overlooked and needs to be augmented, since existing methods are coupled with specific downstream machine analytics models for supervision.

Under these circumstances, our work explores to build a Unified and Generalized ICM (UG-ICM) paradigm aimed at satisfying both human perception and various machine analytics tasks (including unseen ones) with a unified bitstream. For *unification*, we introduce a conditional decoding strategy to turn such unified bitstream into different versions for human or machine preference by only adjusting a controllable parameter. Specifically, we develop a novel Preference Conditional Decoding Module (PCDM), which incorporates the controllable parameter into the decoder at different stages. The PCDM produces a set of bias features based on the indexed human or machine preference, which are then incorporated into the image features to guide the decoding process. Furthermore, we incorporate a multi-stage training strategy, for human perception and machine analytics alternately, that stabilizes training and improves the training efficiency.

For *generalization*, we construct a composite semantic supervision based on the pretrained Contrastive Language-Image Pre-training (CLIP) model (Radford et al. 2021), as illustrated in Fig. 1. Specifically, the feature similarity of compressed images and original ones is measured and regularized in the CLIP feature space with the cosine similarity as the generalized constraint that promotes machine vision analytics. Moreover, such CLIP-based supervision covers hierarchical visual information from global-local-instance granularity, enhancing the capacity for capturing contextual information of different levels, termed Multi-Scale CLIP (MS-CLIP) loss. It is noted that, the entire UG-ICM is trained in a self-supervised manner, without using any specific downstream model and annotations, but offering significant performance gains on a series of visual tasks. We summarize our contributions as follows:

- We make the first exploration of leveraging the CLIP model to construct an intrinsic and generalized semantic constraint in the ICM field, pursuing strong generalization capacity in a self-supervised manner.
- We propose a unified ICM paradigm capable of providing satisfactory perceptual quality and improved analytics performance through a single bitstream.
- We introduce a preference conditional decoding module, enabling the decoding of a single bitstream into differ-

ent versions, each specifically tailored to meet human or machine needs.

- Extensive experimental results have demonstrated that the proposed UG-ICM outperforms state-of-the-art ICM works in both generalized machine analytics and perceptual quality.

Related Works

Learned Image Compression

In the past decade, significant efforts have been devoted to Learned Image Compression (LIC) with considerable advancements. As a result, recent LIC developments have already shown superior performance compared to the state-of-the-art coding standard, VVC (Bross et al. 2021). In special, Ballé *et al.* proposed the first end-to-end image compression network via the VAE architecture (Ballé, Laparra, and Simoncelli 2017), and subsequently introduced a hyperprior (Ballé et al. 2018) and a local context-based (Minnen, Ballé, and Toderici 2018) entropy model to enhance compression performance for *rate-distortion*. Along this vein, numerous efforts have been dedicated to further improving the entropy module. For instance, in (Guo et al. 2021), Guo *et al.* introduced global context information for enhanced entropy coding. Meanwhile, checkerboard context models and channel-wise context were introduced to reduce entropy coding computational complexity in (He et al. 2021) and (Minnen and Singh 2020), respectively. Moreover, by incorporating local spatial, global spatial, and channel contexts, Jiang *et al.* (Jiang et al. 2023) proposed a multi-reference entropy model that achieved state-of-the-art performance. Besides the entropy model, other works have adopted more powerful backbone networks, such as residual networks (Chen, Xu, and Wang 2022), invertible neural networks (Xie, Cheng, and Chen 2021), and Swin-Transformers (Liu, Sun, and Katto 2023).

Image Coding for Machine

To simultaneously satisfy both human perception and machine analytics, preliminary Image Coding for Machine (ICM) works share a similar methodology that cascades the downstream machine analytics network to the output end of LIC. These are jointly trained with respect to the *rate-distortion-analytics* criteria (Chamain et al. 2021; Le et al. 2021; Wang et al. 2021b; Liu et al. 2021; Torfason et al. 2018). However, as described in Liu *et al.* (Liu, Zhang, and Xiong 2019), the inherent pixel-semantic gap inevitably leads this naive methodology to a trade-off scenario, *i.e.*, improved machine analytics inevitably comes at the expense of perceptual quality. Subsequently, some works explored the scalable coding strategy that initially transmits bottleneck features from downstream networks to fulfill machine analytics needs, low-level visual information is subsequently compressed for image reconstruction. To this end, considerable effort has been invested in tailoring the framework for specific tasks (Yang et al. 2021) and developing methods for decoupling semantic and low-level information (Choi and Bajić 2022; Wang et al. 2021a; Liu et al. 2023). Although these methods can fulfill machine analytics needs

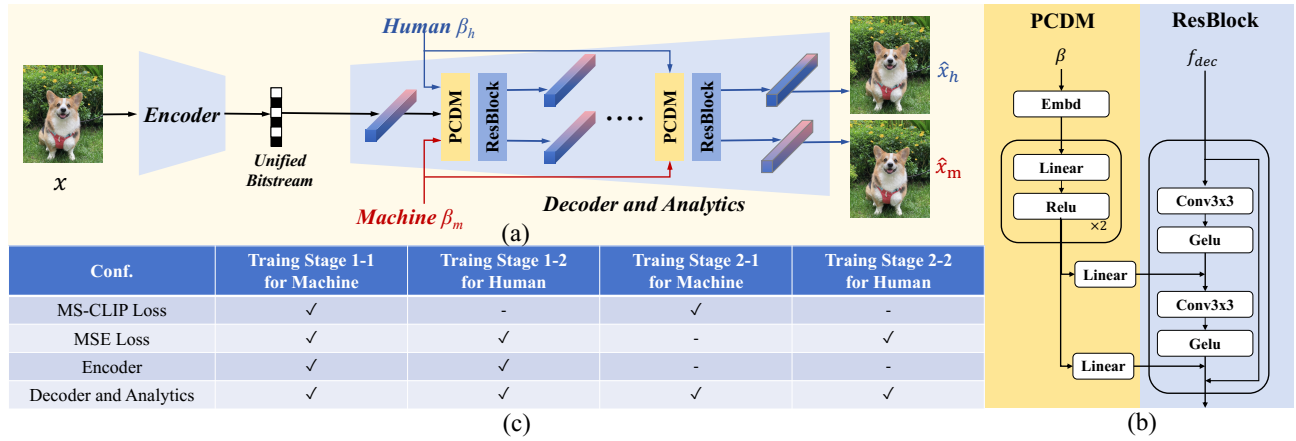


Figure 2: (a) Compressing pipeline of the proposed UG-ICM. (b) Details of the proposed PCDM. (c) Modules and loss terms involved in the multi-stage training process.




Figure 3: Empirical analysis on generalization of taking CLIP as supervision constraint.

without sacrificing perceptual quality, the scalable strategies fall short of leveraging the correlation between high-level and low-level information, resulting in suboptimal compression rates. Recently, transfer-based methods have emerged as a new research trend, *i.e.*, transferring human vision-oriented LIC for machine analytic codecs in a plug-and-play manner (Liu, Sun, and Katto 2022; Chen et al. 2023; Shen et al. 2024). However, these transfer-based methods typically involve a recompression process that produces separate bitstreams for human and machine needs, which lacks flexibility and efficiency in practical implementation.

Proposed Method

Overview and Motivations

To address the neglected issues of existing ICM works, our work introduces a new paradigm, UG-ICM, characterized by its capacity to deliver satisfactory perceptual quality and generalized analytic performance with only a single bitstream. The overall framework is illustrated in Fig. 2.

Specifically, the input image x would be first fed to an encoder $G(\cdot)$ to obtain the unified compact latent representation \hat{y} , containing the information required by both human perception and machine analytics,

$$\hat{y} = G(x). \quad (1)$$

An entropy model $E(\cdot)$ can follow to estimate and constrain the entropy via $E(\hat{y})$, which will be elaborated later. Subsequently, the unified compact representation \hat{y} , along with the corresponding preference condition $\beta \in \{\beta_h, \beta_m\}$ indicating human or machine preference, are jointly fed into the decoder $D(\cdot)$,

$$\hat{x}_h = D(\hat{y}, \beta_h), \quad \hat{x}_m = D(\hat{y}, \beta_m), \quad (2)$$

where \hat{x}_h and \hat{x}_m are reconstructed images favored by human perception and generalized machine analytics, respectively.

In our work, we aim for the model to own two key characteristics: first, \hat{y} should be a *general* representation that supports reconstruction for both human and machine vision; second, \hat{y} can be *generalized* to effectively handle a range of machine vision tasks. To fulfill these goals, the preference-guided decoding process is realized by a Preference Conditional Decoding Module (PCDM), which generates preference-tailored features according to β and biases the decoding process towards distinct versions. To ensure the ICM fully recognizes the disparities between human and machine needs, we have specifically designed a multi-stage training strategy, with each stage supervised by human and machine-oriented criteria, respectively. Moreover, to achieve strong generalization across various downstream tasks, we leverage the CLIP model to provide hierarchical machine-oriented supervision at global-local-instance levels, denoted as Multy-Scale (MS)-CLIP loss. The proposed MS-CLIP loss, PCDM, and corresponding training strategy are detailed in the following subsections.

Multi-Scale CLIP Supervision

Empirical Analysis on Generalization of Taking CLIP as Supervision Constraint. It is well known that the latent representation extracted by the CLIP model is capable of satisfying various unseen machine analysis tasks. However, in the context of ICM, the latent representation is required to own both compactness and pixel-level reconstruction capacity. As such, we are motivated to first make a preliminary

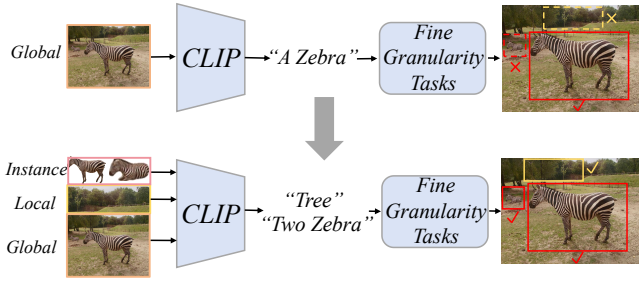


Figure 4: Illustration of the proposed Multi-Scale CLIP loss.

exploration of the generalization capacity of CLIP via leveraging it as the additional supervision for image construction rather than constraining latent representations. To this end, we especially design an offline toy example to examine whether CLIP-based supervision is capable of facilitating generalized machine analytics meanwhile maintaining the perceptual quality.

In detail, denoting \hat{x} as a reconstructed version of x from a certain LIC, we conduct an offline image updating process under the supervision of CLIP,

$$\begin{aligned}\hat{x}^{t+1} &\leftarrow \hat{x}^t - s \cdot \text{sign}[\nabla_{\hat{x}^t}(1 - \cos(\text{CLIP}(\hat{x}^t), \text{CLIP}(x))), \\ \hat{x}^{t+1} &\leftarrow \hat{x} + \text{clamp}(\hat{x}^{t+1} - \hat{x}, -\delta, \delta),\end{aligned}\quad (3)$$

where $\text{CLIP}(\cdot)$ refers to the image modal feature extractor, ∇ signifies the corresponding gradient, t denotes the updating iterations, s represents the updating step size, $\cos(\cdot)$ measures the cosine similarity in CLIP feature space, and the δ is a pixel-wise constraint to ensure the perceptual quality. Herein, we adopt 3,000 images from PASCAL VOC (Everingham et al. 2010) and compress them by MLIC++ (Jiang and Wang 2023), updating them via Equ. (3) with a set of $\delta \in \{1, 2, 8\}$. The updated images are subsequently fed into unseen object detection (Redmon and Farhadi 2018) and semantic segmentation (Chen et al. 2018) models to assess the machine analytics performance.

Corresponding results are shown in Fig. 3, where improvements in various machine analytics tasks can be easily observed with also favorable perceptual quality. These results demonstrate that CLIP is also capable of providing strong generalization capacity from the perspective of image reconstruction guidance, which has motivated us to further study and incorporate it into the ICM framework.

Multi-Scale CLIP Loss. Although the strong generalization capacity of CLIP has been demonstrated, there is still an issue that CLIP itself is an image-level modeling tool and trained to match an image as a whole to a text description. It is difficult to provide region- or pixel-level fine-grained information and may overlook the contextual information of less salient regions (Zhong et al. 2022; Gao et al. 2022). To construct a comprehensive semantic constraint, we propose a novel MS-CLIP loss that applies CLIP supervision at global, local, and instance levels, as shown in Fig. 4.

To comprehensively capture the local contextual information from the input image, we constraint the perspective field of the CLIP model to the instance level. In special, an

instance segmentation model (Wang et al. 2022) is leveraged to obtain a set of instance-wise pairs $[i_1, i_2, \dots, i_k]$ and $[\hat{i}_1, \hat{i}_2, \dots, \hat{i}_k]$ from the input image x and the reconstructed image \hat{x} . The cosine similarity between them would be calculated,

$$\mathcal{L}_{ins}(x, \hat{x}) = \sum_{j=1}^k (1 - \cos(\text{CLIP}(i_j), \text{CLIP}(\hat{i}_j))). \quad (4)$$

Moreover, to obtain more fine-grained information for potentially more generalized semantic clues to capture relationships among different models, we incorporate a local-wise supervision method to enhance generalization capabilities across various local perspectives. In special, during the training process, we random crop a local-wise patch pairs $[l, \hat{l}]$ from the input x and reconstructed version \hat{x} , and calculated their cosine similarity at the CLIP feature space,

$$\mathcal{L}_{loc}(x, \hat{x}) = 1 - \cos(\text{CLIP}(l), \text{CLIP}(\hat{l})). \quad (5)$$

Thus, by incorporating global-wise CLIP-based supervision \mathcal{L}_{glo} ,

$$\mathcal{L}_{glo}(x, \hat{x}) = 1 - \cos(\text{CLIP}(x), \text{CLIP}(\hat{x})), \quad (6)$$

the established multi-scale CLIP loss \mathcal{L}_{MC} is expected to provide a comprehensive understanding of the semantic divergence between input image pairs, which can be formulated by,

$$\mathcal{L}_{MC}(x, \hat{x}) = \mathcal{L}_{glo}(x, \hat{x}) + \mathcal{L}_{loc}(x, \hat{x}) + \mathcal{L}_{ins}(x, \hat{x}). \quad (7)$$

Preference Conditional Decoding Module

The decoder of the proposed UG-ICM is responsible for generating different versions of reconstructed images tailored to human or machine needs from a unified bitstream. This capability is enabled by a specially designed Preference Conditional Decoding Module (PCDM), which is integrated into each decoding block. As illustrated in Fig. 2 (b), during the decoding process, the preference condition is indicated by $\beta \in \{\beta_h, \beta_m\}$ for human or machine needs respectively. The preference condition is subsequently fed into a 2-layer MLP, yielding a corresponding preference conditional feature, f_β . This feature, f_β , is then fused with the image feature to bias the decoding process towards its indicated preference,

$$f_{PCDM} = f_\beta W_d + f_{dec}, \quad (8)$$

where f_{PCDM} and f_{dec} denote the output and the input features of the PCDM module respectively, W_d is the feature element-wise weighting map with d denoting the channel depth aligned with the input features. In coding practice, the $[\beta_h, \beta_m]$ are valued at $[0, 1]$.

Training and Supervision

We develop a two-stage training strategy as shown in Fig. 2 (c) to achieve optimal performance in two distinct objectives within the proposed UG-ICM framework. In detail, the first stage optimizes the entire framework with the goal of fully extracting visual information for both human perception and machine analytics. The second stage specifically emphasizes the conditional decoder to further improve the performance with respect to two distinct objectives.




Figure 5: Machine analytics performance comparisons between the proposed UG-ICM and the backbone compression networks.




Figure 6: Perceptual quality comparisons between the proposed UG-ICM and the backbone compression networks.

First Stage. As for the first stage, the encoder G_θ , entropy model E_ϕ , and the conditional decoder D_ψ , parameterized by θ , ϕ , ψ , respectively, are jointly optimized by alternating between *rate-perception* and *rate-analytics* criteria. In particular, we split the training iterations within one epoch into two sessions, and the corresponding training process can be formulated by:

$$S\ 1-1: \bar{\theta}, \bar{\phi}, \bar{\psi} = \arg \min_{(\phi, \psi, \theta)} \sum_{(x \in X)} \lambda \cdot \mathcal{L}_2(x, D_\psi(G_\theta(x), \beta_h)) + E_\phi(\hat{y}),$$

$$S\ 1-2: \tilde{\theta}, \tilde{\phi}, \tilde{\psi} = \arg \min_{(\theta, \phi, \psi)} \sum_{(x \in X)} \lambda \cdot (\mathcal{L}_2(x, D_{\tilde{\psi}}(G_{\tilde{\theta}}(x), \beta_m)) + \mathcal{L}_{MC}(x, D_{\tilde{\psi}}(G_{\tilde{\theta}}(x), \beta_m)) + E_{\tilde{\phi}}(\hat{y})),$$

where λ is a Lagrange parameter to obtain the ICM models at different compression levels.

Second Stage. In the second stage, the parameters of the encoder and entropy model are frozen, thereby focusing optimization efforts solely on the decoder. Moreover, the iterations are half split into two sessions for each epoch, the training processes are formulated by,

$$S\ 2-1: \bar{\psi} = \arg \min_{(\psi)} \sum_{(x \in X)} \mathcal{L}_2(x, D_{\bar{\psi}}(G_\theta(x), \beta_h)),$$

$$S\ 2-2: \tilde{\psi} = \arg \min_{(\tilde{\psi})} \sum_{(x \in X)} \mathcal{L}_{MC}(x, D_{\tilde{\psi}}(G_\theta(x), \beta_m)).$$

Experiments

We first implement the proposed UG-ICM on different backbone LIC networks, examining its effectiveness by compar-

ing the LIC networks with and without our implementations. Subsequently, we compare the UG-ICM with cutting-edge ICM methods in terms of perceptual quality and machine analytics performance improvement across a variety of downstream tasks. Moreover, comprehensive ablation studies have been conducted to demonstrate the effectiveness of the proposed PCDM and Multi-Scale CLIP Loss.

Experimental Settings

Benchmark. To comprehensively examine the generalized machine analytics performance of the proposed UG-ICM, four commonly employed downstream tasks are employed, including classification, object detection, instance segmentation, and semantic segmentation, all of which are unseen by our method. Additionally, perceptual quality is thoroughly assessed using various indices.

- **Classification:** 25,000 images from the ILSVRC 2012 dataset (Deng et al. 2009) (ImageNet-1k), spanning 1,000 categories, are employed for testing. ResNet101 (He et al. 2016) well-trained on ImageNet is employed for assessment. The performance is indexed by Top-1 accuracy.
- **Object Detection:** 2,150 images from the testing set of PASCAL VOC (Everingham et al. 2010) are employed for testing. The YOLOv3 (Redmon and Farhadi 2018) well-trained by VOC is used for assessment. The performance is indexed by mAP for IoU = 0.5.
- **Instance Segmentation:** The subset *val2017* of COCO dataset (Lin et al. 2014) are employed for testing. The Mask-RCNN (He et al. 2017) well-trained by COCO is employed for assessment. Corresponding performance is indexed by mAP for IoU $\in [0.5 : 0.05 : 0.95]$.
- **Semantic Segmentation:** 1,450 images from PASCAL VOC are employed for testing. DeepLabv3+ (Chen et al.

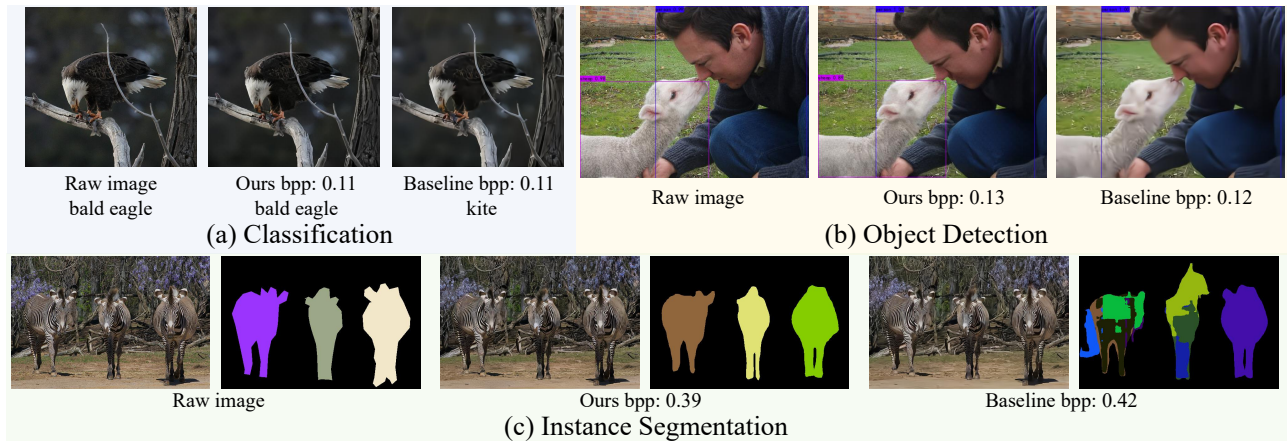


Figure 7: Visualization of the machine analytics performance regarding classification (a), object detection (b), and instance segmentation (c), where the ‘Baseline’ denotes the backbone compression network MLIC++.

2018) well-trained by VOC is employed for assessment. The performance is indexed by mIoU.

- **Human Perception:** The Kodak24 dataset is employed to assess the perceptual quality using three full-reference image quality assessment measures, *i.e.*, PSNR, SSIM, and LPIPS (Zhang et al. 2018).

Anchors. We implement the proposed UG-ICM on two LIC networks, including the milestone *Bmshj2018* (Ballé et al. 2018) and the cutting-edge MLIC++ (Jiang and Wang 2023). To ensure fairness, the LIC with and without the implementation of UG-ICM are both trained from scratch, with the same training set.

As for the compared ICM works, we employ three works with different methodologies and pipelines for a comprehensive examination. These include two supervised approaches: *Chamain’s* (Chamain et al. 2021) and TransTIC (Chen et al. 2023), along with a self-supervised approach, SA-ICM (Shindo et al. 2024). In detail,

- *Chamain’s*: represents the conventional ICM paradigm, which straightforwardly adopts the machine analytics performance as an extra loss term for network optimization. We implement it on the MLIC++ backbone and re-train it with corresponding downstream models to ensure fair comparisons.
- TransTIC: is the cutting-edge transfer-based ICM, adapting the LIC to machine analytics needs in a plug-and-play manner. It should be noted that, with this method, images need to be recompressed into another bitstream when adapting to a new task.
- SA-ICM: is a self-supervised ICM approach that adopts the segmentation map from SAM (Kirillov et al. 2023) to identify regions of interest, which are then allocated more coding resources during the compression process.

Implementation Details. 80,000 images from the subset *train2017* of COCO are employed for training. The whole network is implemented in Pytorch with CUDA support and is trained on a single NVIDIA A800-80G GPU. The CLIP

(ViT-B/32) is adopted for our MS-CLIP implementation. In the first stage of the training process, we use the Adam optimizer with a learning rate of $1e-4$ and train for 200 epochs. For the second stage, we reduce the learning rate to $1e-5$ and train for an additional 10 epochs. The UG-ICM and corresponding LIC backbone networks are each trained four times with different Lagrange parameters, achieving four distinct compression rates regarding bits-per-pixel (bpp).

Experimental Results

Comparison on Different Backbones. As mentioned above, the proposed UG-ICM is implemented on two LIC backbones, *Bmshj2018* and MLIC++. In this subsection, we compare the UG-ICM with the vanilla backbone in terms of perceptual quality and various unseen downstream models, demonstrating the effectiveness of our implementation. The comparison results for machine analytics are presented in Fig. 5. Encouraging observations have been made, showing that the proposed method brings overall improvement across multiple downstream tasks and bitrates.

Specifically, with the *Bmshj* backbone, the proposed method achieves an average analytics gain of 1.7% in top-1 accuracy, 1.7% in mAP, 3.2% in mIOU, and 1.2% in mAP under the same bits-per-pixel (bpp) level for classification, object detection, semantic segmentation, and instance segmentation, respectively. As for the MLIC++ backbone, the improvement becomes even larger, with an average bpp-analytics gain of 2.8%, 3.9%, 6.0%, and 3.2% regarding the four downstream tasks. Considering that both the downstream networks and corresponding annotations are unseen by our UG-ICM, the observed improvements are significant, demonstrating that the rich image modal knowledge from CLIP is comprehensively leveraged by our scheme. To gain intuition into the analytics performance improvement, a group of examples is provided in Fig. 7.

Comparison results in terms of PSNR, SSIM, and LPIPS are provided in Fig. 6 (a), (b), and (c) respectively. As shown, the human-preference version of our UG-ICM maintains comparable perceptual quality to the corresponding

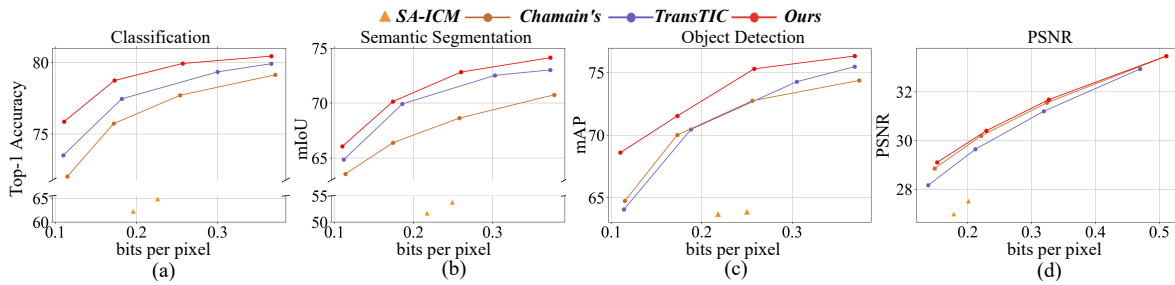


Figure 8: Performance comparison between the proposed UG-ICM and employed anchors regarding classification (a), semantic segmentation (b), objection detection (c), and perceptual quality indexed by PSNR (d).

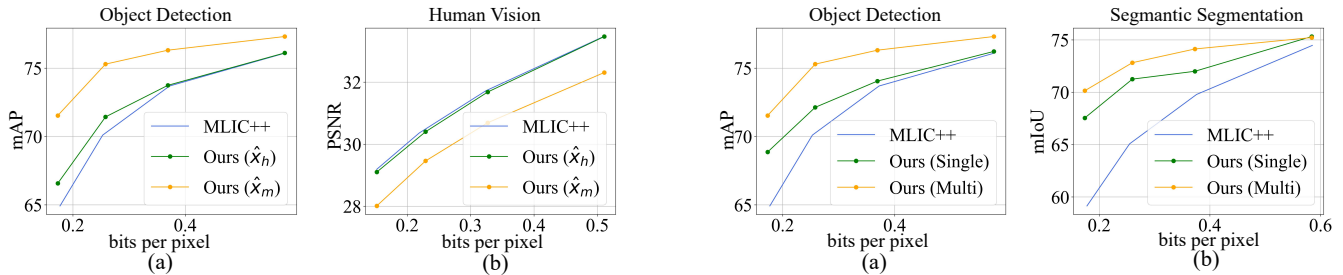


Figure 9: Performance comparisons between decoded images tailored for human and machine preferences.

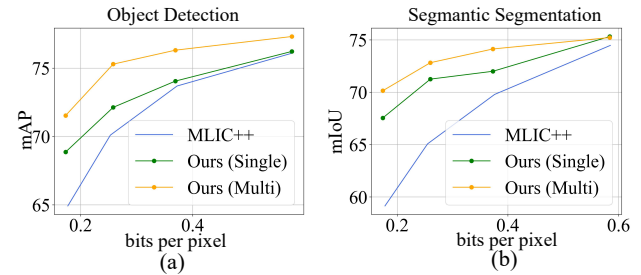


Figure 10: Ablation studies of Multi-Scale CLIP Loss.

backbones, owing to the proposed PCDM module and the multi-stage training strategy that sensitively aligns the compression module with both human and machine needs.

Comparison to State-of-the-art ICM Methods. This subsection compares our UG-ICM (MLIC++) with other ICM methods, covering both machine vision and human vision aspects. Comparison results are provided in Fig. 8, demonstrating the superiority of our proposed UG-ICM against others regarding both machine analytics and perceptual quality. In particular, compared with the conventional ICM framework *Chamain's*, which adopts the same compression backbone as ours, the naive method of incorporating machine analytics preferences and the empirically decided bits allocation strategy fall short in achieving a good balance between perceptual quality and multiple downstream tasks. Moreover, compared with TransTIC, the proposed UG-ICM still possesses significant advantages in machine analytics. This is because TransTIC employs a *plug-and-play* strategy, which avoids changing the parameters of the compression backbone, ultimately leading to limited analytics improvement.

Ablation Studies

Effectiveness Examination of the PCDM. The proposed UG-ICM is capable of decoding a single bitstream into different versions, facilitated by the incorporated PCDM module. Herein, we assess its effectiveness by comparing the decoded versions tailored for human preference, \hat{x}_h , machine preference, \hat{x}_m , as well as the decoded images from the compression backbone. As shown in Fig. 9, \hat{x}_m and \hat{x}_h

exhibit a clear bias towards their corresponding preferences. Specifically, \hat{x}_m achieves an average gain of 3% in bpp-mAP compared to \hat{x}_h regarding object detection, but falls short by 1dB bpp-PSNR concerning human perception.

Ablation Studies of The Multi-Scale CLIP Loss. We ablate the multi-scale strategy from the CLIP-based supervision by retaining the whole framework with global-wise CLIP loss. Comparison results between our full version ‘Ours (Multi)’ and the ablated version ‘Ours (Single)’ regarding object detection and semantic segmentation are provided in Fig. 10. As shown, the multi-scale supervision strategy significantly improves machine analytics by enhancing the ability to capture the semantic information of small objects, which are often overlooked. In contrast, vanilla CLIP models tend to prioritize capturing global context.

Conclusion

This paper proposes a novel Unified and Generalized (UG-ICM). It is capable of providing satisfactory perceptual quality and improved machine analytics performance across various tasks with a single bitstream. Specifically, we first incorporate a conditional decoding module that introduces human or machine needs as extra conditions, enabling the bitstream to be decoded into different versions for corresponding preferences. Furthermore, we explore the generalization capacity of the CLIP model and propose global-to-instance-wise CLIP supervision, accommodating preferences for different scale semantic information across various machine analytics tasks. Experimental results and ablation studies have demonstrated the effectiveness of our design.

Acknowledgments

This work was in part by the Basic and Frontier Research Project of PCL, the Major Key Project of PCL, in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515010454, in part by the Open Research Fund from Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ) under Grant No. GML-KF-24-27, in part by the Natural Science Foundation of Guangdong Province under Grant 2023A1515011667, in part by the Science and Technology Major Project of Shenzhen under Grant 202302D074, in part by the Key Basic Research Foundation of Shenzhen under Grant JCYJ20220818100205012, and in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2023B1515120020.

References

- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2017. End-to-end optimized image compression. In *Proceedings of the International Conference on Learning Representations*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. In *Proceedings of the International Conference on Learning Representations*.
- Bross, B.; Wang, Y.; Ye, Y.; Liu, S.; Chen, J.; Sullivan, G.; and Ohm, J. 2021. Overview of the Versatile Video Coding (VVC) Standard and its Applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10): 3736–3764.
- Chamain, L. D.; Racapé, F.; Bégaint, J.; Pushparaja, A.; and Feltman, S. 2021. End-to-end optimized image compression for machines, a study. In *Proceedings of the Data Compression Conference*, 163–172.
- Chen, F.; Xu, Y.; and Wang, L. 2022. Two-stage octave residual network for end-to-end image compression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 3922–3929.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision*, 801–818.
- Chen, Y.-H.; Weng, Y.-C.; Kao, C.-H.; Chien, C.; Chiu, W.-C.; and Peng, W.-H. 2023. Transtic: Transferring transformer-based image compression from human perception to machine perception. In *Proceedings of the IEEE International Conference on Computer Vision*, 23297–23307.
- Chen, Z.; Fan, K.; Wang, S.; Duan, L.; Lin, W.; and Kot, A. 2020. Toward Intelligent Sensing: Intermediate Deep Feature Compression. *IEEE Transactions on Image Processing*, 29: 2230–2243.
- Chen, Z.; Fan, K.; Wang, S.; Duan, Y.; Lin, W.; and Alex, K. 2019. Lossy intermediate deep learning feature compression and evaluation. In *Proceedings of the ACM International Conference on Multimedia*, 2414–2422.
- Choi, H.; and Bajić, I. V. 2022. Scalable image coding for humans and machines. *IEEE Transactions on Image Processing*, 31: 2739–2754.
- Choi, J.; and Han, B. 2020. Task-aware quantization network for jpeg image compression. In *Proceedings of European Conference on Computer Vision*, 309–324.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88: 303–338.
- Gao, Y.; Liu, J.; Xu, Z.; Zhang, J.; Li, K.; Ji, R.; and Shen, C. 2022. Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *Neural Information Processing Systems*, 35: 35959–35970.
- Guo, Z.; Zhang, Z.; Feng, R.; and Chen, Z. 2021. Causal contextual prediction for learned image compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(4): 2329–2341.
- He, D.; Zheng, Y.; Sun, B.; Wang, Y.; and Qin, H. 2021. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 14771–14780.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Huang, Z.; Jia, C.; Wang, S.; and Ma, S. 2021. Visual analysis motivated rate-distortion model for image coding. In *Proceedings of IEEE International Conference on Multimedia and Expo*, 1–6.
- Jiang, W.; and Wang, R. 2023. MLIC++: Linear Complexity Multi-Reference Entropy Modeling for Learned Image Compression. In *ICML Workshop Neural Compression: From Information Theory to Applications*.
- Jiang, W.; Yang, J.; Zhai, Y.; Ning, P.; Gao, F.; and Wang, R. 2023. Mlic: Multi-reference entropy model for learned image compression. In *Proceedings of the ACM International Conference on Multimedia*, 7618–7627.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Le, N.; Zhang, H.; Cricri, F.; Ghaznavi-Youvalari, R.; and Rahtu, E. 2021. Image coding for machines: an end-to-end learned approach. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 1590–1594.
- Li, B.; Ye, L.; Liang, J.; Wang, Y.; and Han, J. 2022. Region-of-interest and channel attention-based joint optimization of image compression and computer vision. *Neurocomputing*, 500: 13–25.

- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, 740–755.
- Liu, D.; Zhang, H.; and Xiong, Z. 2019. On the classification-distortion-perception tradeoff. In *Proceedings of the International Conference on Neural Information Processing Systems*, 1206–1215.
- Liu, J.; Sun, H.; and Katto, J. 2022. Improving multiple machine vision tasks in the compressed domain. In *Proceedings of the IEEE International Conference on Pattern Recognition*, 331–337.
- Liu, J.; Sun, H.; and Katto, J. 2023. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 14388–14397.
- Liu, K.; Liu, D.; Li, L.; Yan, N.; and Li, H. 2021. Semantics-to-signal scalable image compression with learned revertible representations. *International Journal of Computer Vision*, 129(9): 2605–2621.
- Liu, L.; Hu, Z.; Chen, Z.; and Xu, D. 2023. Icmh-net: Neural image compression towards both machine vision and human vision. In *Proceedings of the ACM International Conference on Multimedia*, 8047–8056.
- Minnen, D.; Ballé, J.; and Toderici, G. D. 2018. Joint autoregressive and hierarchical priors for learned image compression. *Advances in Neural Information Processing Systems*, 31.
- Minnen, D.; and Singh, S. 2020. Channel-wise autoregressive entropy models for learned image compression. In *Proceedings of the IEEE International Conference on Image Processing*, 3339–3343.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763.
- Redmon, J.; and Farhadi, A. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Shen, X.; Yin, K.; Wang, X.; He, Y.; Wang, S.; and Yang, W. 2024. Image Coding for Analytics via Adversarially Augmented Adaptation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 3605–3609.
- Shindo, T.; Yamada, K.; Watanabe, T.; and Watanabe, H. 2024. Image Coding for Machines with Edge Information Learning Using Segment Anything. *arXiv preprint arXiv:2403.04173*.
- Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on Circuits and Systems for Video technology*, 22(12): 1649–1668.
- Torfason, R.; Mentzer, F.; Agustsson, E.; Tschannen, M.; Timofte, R.; and Van Gool, L. 2018. Towards image understanding from deep compression without decoding. *arXiv preprint arXiv:1803.06131*.
- Wang, S.; Wang, S.; Yang, W.; Zhang, X.; Wang, S.; Ma, S.; and Gao, W. 2021a. Towards analysis-friendly face representation with scalable feature and texture compression. *IEEE Transactions on Multimedia*, 24: 3169–3181.
- Wang, S.; Wang, Z.; Wang, S.; and Ye, Y. 2021b. End-to-end compression towards machine vision: Network architecture design and optimization. *IEEE Open Journal of Circuits and Systems*, 2: 675–685.
- Wang, X.; Yu, Z.; De Mello, S.; Kautz, J.; Anandkumar, A.; Shen, C.; and Alvarez, J. M. 2022. Freesolo: Learning to segment objects without annotations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 14176–14186.
- Xie, Y.; Cheng, K.; and Chen, Q. 2021. Enhanced invertible encoding for learned image compression. In *Proceedings of the ACM International Conference on Multimedia*, 162–170.
- Yang, S.; Hu, Y.; Yang, W.; Duan, L.-Y.; and Liu, J. 2021. Towards coding for human and machine vision: Scalable face image coding. *IEEE Transactions on Multimedia*, 23: 2957–2971.
- Yang, W.; Huang, H.; Hu, Y.; Duan, L.; and Liu, J. 2024. Video Coding for Machines: Compact Visual Representation Compression for Intelligent Collaborative Analytics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(07): 5174–5191.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 586–595.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 16793–16803.