
MCUCoder: Adaptive Bitrate Learned Video Compression for IoT Devices

Ali Hojjat, Janek Haberer, Olaf Landsiedel
Kiel University, Germany

{ali.hojjat, janek.haberer, olaf.landsiedel}@cs.uni-kiel.de

Abstract

The rapid growth of camera-based Internet of Things (IoT) devices demands the need for efficient video compression, particularly for edge applications where devices face hardware constraints, often with only 1 or 2 MB of RAM and unstable internet connections. Traditional and deep video compression methods are designed for high-end hardware, exceeding the capabilities of these constrained devices. Consequently, video compression in these scenarios is often limited to Motion-JPEG (M-JPEG) due to its high hardware efficiency and low complexity. This paper introduces MCUCoder, an open-source adaptive bitrate video compression model tailored for resource-limited IoT settings. MCUCoder features an ultra-lightweight encoder with only 10.5K parameters and a minimal 350KB memory footprint, making it well-suited for edge devices and Microcontrollers (MCUs). While MCUCoder uses a similar amount of energy as M-JPEG, it reduces bitrate by 55.65% on the MCL-JCV dataset and 55.59% on the UVG dataset, measured in MS-SSIM. Moreover, MCUCoder supports adaptive bitrate streaming by generating a latent representation that is sorted by importance, allowing transmission based on available bandwidth. This ensures smooth real-time video transmission even under fluctuating network conditions on low-resource devices. Source code available at <https://github.com/ds-kiel/MCUCoder>.

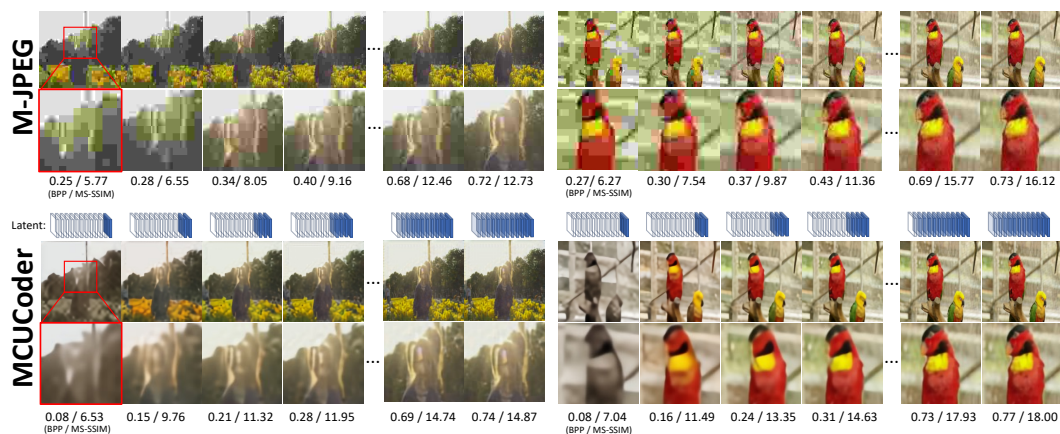


Figure 1: Qualitative comparison of MCUCoder and M-JPEG across various compression rates on two videos from the MCL-JCV [1] and UVG [2] datasets. As we can see, MCUCoder offers a significantly better MS-SSIM/bpp trade-off. For instance, at 0.15 bpp in the left example, with MCUCoder we can see the person’s face whereas with M-JPEG we need at least 0.34 bpp to make out the face. Note that the images in each column do not necessarily have the same bitrate. More examples are reported in Appendix D.

1 Introduction

Motivation: The number of camera-based IoTs devices using always-on MCU is growing rapidly, reaching tens of billions [12]. These devices are widely used in applications such as surveillance cameras [13, 14, 15], wearable cameras [16], robotics [17], wildlife monitoring [18], road monitoring [19], and smart farming [20]. Typically, they capture raw frames through a camera sensor, encode them, and transmit the compressed version to a server via the Internet for further processing, including human observation or AI tasks such as object detection and classification [19]. Therefore, a video encoder is necessary to efficiently compress the captured frames before transmission. However, in IoT environments, there are two primary limitations: constrained hardware resources and limited communication bandwidth.




Figure 2: Number of parameters of MCUCoder and other learned image compression [3, 4, 5, 6, 7, 8, 9] and video compression models [10, 11].

constrained devices. Similarly, while Neural Networks (NNs) and AI-based compression methods outperform traditional models [10, 11], they also often require considerable RAM and GPU resources. For instance, just storing a model with 1M parameters requires around 4 MB of RAM; see Fig. 2. As a result, in such settings, devices are typically limited to using M-JPEG [24], a video compression format where each frame is compressed individually as a JPEG image, which is efficient and hardware-friendly.

2 - Limited Internet: Many IoT devices are located in remote areas where Internet connection is weak and unstable, making it necessary for the encoder to have an **Adaptive Bitrate Encoding** that can generate video streams with varying bitrate. This feature allows the encoder to dynamically adjust its quality according to the available bandwidth, ensuring continuous and smooth playback. This is especially important for real-time applications like live monitoring, where it is crucial to avoid interruptions and maintain a consistent user experience despite fluctuating network conditions. However, implementing an adaptive bitrate encoder adds complexity, as it requires mechanisms to prioritize bit stream information based on its impact on frame quality (e.g., PSNR or MS-SSIM), which is challenging for constrained devices.

Approach: To address these challenges, we introduce MCUCoder, *an adaptive bitrate deep video compression model tailored for resource-limited IoT devices*. Our approach focuses on creating an "asymmetric" compression model that features an ultra-lightweight encoder designed to be both computationally efficient and memory-friendly. Also, MCUCoder produces an "adaptive bitrate" bitstream. Specifically, in MCUCoder, we train the encoder using stochastic dropout such that, instead of explicitly detecting the important parts, it produces latent channels that are sorted based on importance. Afterward, based on the available internet bandwidth, the encoder transmits the first k channels to the decoder; see Fig. 1. This approach is beneficial for low-power MCUs since it shifts the complexity of identifying important data to the training phase rather than the inference phase. Also, by employing stochastic dropout training, the decoder can reconstruct the frame even with partial data availability, which is essential for maintaining smooth and uninterrupted video transmission in real-time applications, where network conditions can vary. Additionally, MCUCoder’s encoder is INT8 quantized, allowing it to utilize Digital Signal Processor (DSP) and CMSIS-NN [25] accelerators for faster processing and reduced power consumption.




Figure 3: Overview of MCUCoder architecture. The encoder compresses the input frame into a sorted latent space. Afterward, channels are independently quantized and transmitted based on available bandwidth. The decoder reconstructs the frame by zeroing out missing channels.

Contributions:

1. MCUCoder has an ultra-lightweight encoder with only 10.5K parameters and a minimal memory footprint of roughly 350KB RAM on nRF5340 and STM32F7 MCUs, making it suitable for such low-resource IoT devices.
2. MCUCoder has an energy-efficient INT8 quantized encoder, which leverages the MCU’s DSP and CMSIS-NN accelerators to achieve JPEG-level energy efficiency. Compared to its main baseline, M-JPEG, it saves 55.65% overall bit rate on the MCL-JCV dataset and 55.59% on the UVG dataset, measured in MS-SSIM.
3. MCUCoder produces a progressive bitstream that enables adaptive bitrate streaming, allowing robust video transmission under varying network conditions.

2 Related Work

Traditional and NN based video compression: Video compression is a field that has been evolving for decades. Beyond traditional codecs like H.264 [21], H.265 [22], and H.266 [23], deep learning-based approaches often replace conventional modules such as motion compensation [10, 26], transform coding [9, 27], and entropy coding [28, 29]. Also, some work has been done regarding the end-to-end optimization of video compression models [30, 31, 32]. Lu et al. [11] introduce DVC, the first end-to-end deep video compression model. Hu et al. [33, 34] extend DVC to operate in both pixel and feature domains. Li et al. [35] and Lie et al. [36] reduce bitrates by modeling probabilities over video frames using conditional coding. Also, in recent years, there has been growing interest in using implicit neural representations for video compression [37, 38]. However, due to their substantial hardware requirements, these models are unsuitable for deployment on low-resource IoT devices.

Video compression for IoT: We can categorize IoT-based video encoders into two parts: hardware-based and software-based. Hardware approaches primarily focus on designing more power-efficient camera sensors [39, 40, 41] and more efficient MCU circuits and processors [42, 43, 44]. Due to its simplicity, scalability, low latency, and very low energy consumption, the most common software-based video encoder on IoT devices is M-JPEG [24]. Nevertheless, there have been few works exploring alternative software-based models: Veluri et al. [16] employ M-JPEG on the encoder to capture black-and-white and colorized frames at two different resolutions and uses super-resolution methods to interpolate and colorize frames on the decoder. However, unlike MCUCoder, it is not adaptive and relies on a JPEG encoder on MCUs. Hu et al. [13] propose a deep image encoder model for MCUs, but it is also non-adaptive. Additionally, they patchify the input, which significantly increases encoding time, making it impractical for real-time video compression. MCUCoder combines the advantages of both worlds: it offers the adaptive bitrate feature of more complex encoders, while maintaining the efficiency necessary for low-resource devices, making it an ideal solution for IoT video compression.

3 MCUCoder

In this section, we introduce MCUCoder, an adaptive bitrate asymmetric video compression model, specifically designed for IoT settings. We begin by detailing the asymmetric encoder-decoder architecture of MCUCoder, including the customized quantization processes. Then, we present the stochastic dropout training method, which trains the encoder of MCUCoder to store information in its channels based on importance.




Figure 4: MCUCoder latent channels: Early channels (important ones) capture low-frequency features, while later channels capture high-frequency features, similar to the DCT in JPEG.

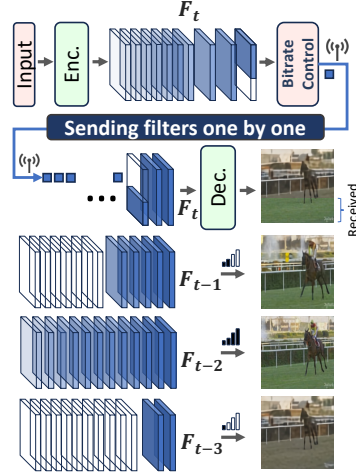


Figure 5: An example of MCUCoder bitrate adaptation under dynamic network bandwidth, where the bitrate control module acts as a gate to determine the number of channels to send.

Asymmetric Compression: MCUs are characterized by highly constrained hardware resources, such as limited RAM, CPU, FLASH, and power availability. Additionally, existing MCU-specific NN frameworks like TFLite Micro support only a limited set of NN layers [13]. To address these constraints, we propose an asymmetric [45] encoder-decoder architecture optimized for constrained devices. Due to hardware constraints, MCUCoder encodes each frame independently, as inter-frame compression is not feasible. The encoder contains only 10.5K parameters, while the decoder utilizes approximately 3M parameters and leverages SOTA image decompression blocks; see Fig. 12. The encoding process begins by passing input frame f_t through three convolutional layers. To maximize the data range for subsequent quantization, no activation function is applied in the final encoder layer, avoiding the negative truncation caused by ReLU. Afterward, each channel of the latent is quantized into INT8 individually, followed by a further reduction to 5-bit precision to enhance compression efficiency. For the decoder, inspired by [46], we integrate a combination of attention blocks [47] and residual bottleneck blocks [48] to reconstruct the frame; see Fig. 12.

Stochastic dropout training: Bitrate adaptation is a feature that typically introduces additional complexity to the encoding process, which can be challenging to implement on MCUs due to resource constraints. In the literature, dropout [49] serves as a powerful tool for enhancing generalization in NNs. Building on this insight, we employ a "biased" version of dropout to train MCUCoder in a way that instead of random dropping, it drops from the tail of the latent [50, 51]. Specifically, on each iteration, after the encoder E gets the input frame f_t , it generates the latent representation z_N , where N is the number of the channels of the latent. Afterward, from a uniform distribution, denoted as $\mathcal{U}_{(0,1)}$, it generates a number, denoted as k , and drops (zero out) the last $[k \times N]$ channels from z_N . As a result, instead of z_N , the decoder D gets $z_{[0:[k \times N]]}$, fills the missing channels with zero, and then reconstructs the output.

$$f_t \rightarrow E(f_t) \rightarrow z_N \xrightarrow{k \sim \mathcal{U}_{(0,1)}} z_{[0:[k \times N]]} \rightarrow D(z_{[0:[k \times N]]}) \rightarrow \hat{f}_t \quad (1)$$

This tailored version of dropout biases the training to prioritize the earlier channels over the later ones. Consequently, the encoder learns to encode more critical information (low frequency) in the initial feature maps and less important (high frequency) details in the subsequent ones; see Fig 6. This prioritization enables flexible bitrate adaptation: upon encoding each frame, the encoder starts transmitting the most significant channels first. Depending on the available bandwidth, the bitrate control module determines how many channels need to be sent to the decoder to ensure uninterrupted streaming; see Fig 5. Importantly, because the latent features are pre-ordered by significance, the




Figure 6: Stochastic dropout training

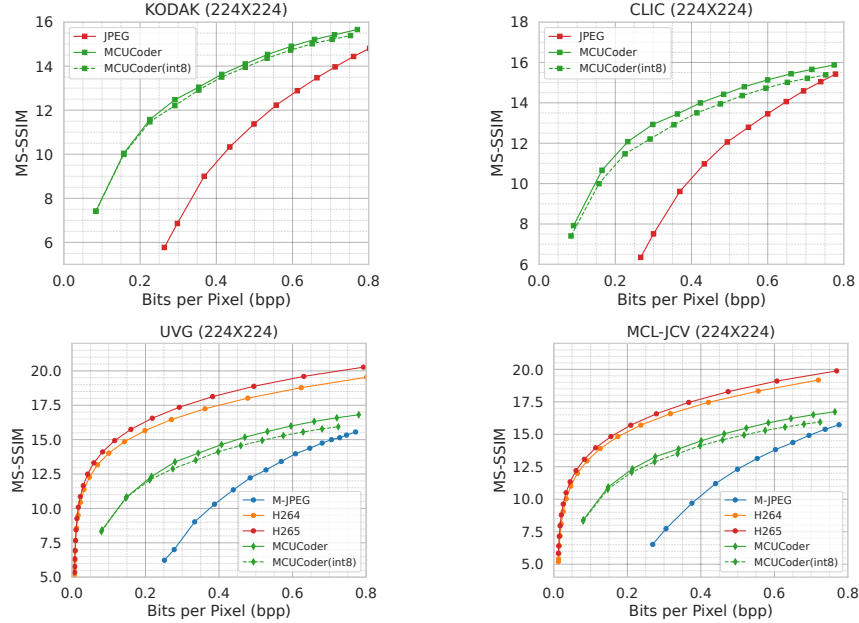


Figure 7: Comparison of MCUCoder (quantized and non-quantized model) and baselines on the image (KODAK [52], CLIC [53]) and video (MCL-JCV [1], UVG [2]) compression datasets. For context, we also compare with H.264 and H.265 on video datasets, despite being impractical for MCUs due to high hardware demands. All datasets are resized to 224×224 .

bitrate control module basically acts like a simple gate and does not add any extra computational complexity to the encoder.

4 Evaluation

We train MCUCoder on the 300K largest ImageNet images [54] and apply noise-downsampling preprocessing [55, 3]. We use Adam with an initial learning rate of 10^{-4} and a batch size of 16, and train for 1M iterations, lowering the learning rate to 10^{-5} in the final 50K iterations [46]. To address quantization effects, we add random noise to the latent. Since MCUCoder is specifically designed for IoT environments, where the structure of the output is more critical than fine details, we use MS-SSIM as the loss function. We also quantize inputs, weights, and activations to INT8 for RAM efficiency and to leverage DSP and CMSIS-NN accelerators [25] in MCUs. We use post-training quantization existing in TFLite-Micro [56] to reduce latency, processing power, and model size with minimal degradation in model accuracy. For all comparisons, we report performance metrics for both the FLOAT32 and INT8 models.

4.1 Quantitative results

Due to the limited hardware resources of MCUs, inter-frame compression is not practically feasible. As a result, in such devices, video compression is limited to M-JPEG where each frame is compressed independently. Therefore, in addition to evaluating MCUCoder and its baselines from the perspective of video compression, we also assess its performance on image compression datasets. Given the lower resolution commonly encountered in IoT scenarios, we resize all the videos and images to 224×224 .

Video compression: We evaluate MCUCoder on the UVG [2] and MCL-JCV [1] datasets, comparing its performance to M-JPEG, see Fig. 7. For additional context, we include comparisons with traditional video codecs such as H.264 [21] and H.265 [22], even though these codecs are impractical for deployment on MCUs due to their significant computational and hardware demands. Also, we report the Bjøntegaard Delta (BD) rate [57] for both datasets in Table 1. The results indicate that MCUCoder achieves a significantly higher MS-SSIM per bit compared to M-JPEG, highlighting its ability to deliver better video quality at lower bitrates. This is especially valuable for IoT applications, where achieving high compression rates with minimal computational overhead is crucial due to

Table 1: MCUCoder (Quantized) BD-rate results. The anchor is M-JPEG.

Type	Dataset	MS-SSIM	PSNR
Video	MCL-JCV	-55.65%	-47.39%
	UVG	-55.59%	-35.28%
Image	KODAK	-55.75%	-43.01%
	CLIC	-49.54%	-38.02%

Table 2: Resource demands of MCUCoder on nRF5340 and STM32F7 MCUs.

	nRF5340	STM32F7
Exec (ms)	1,969	237
RAM (KB)	344 (33%)	360 (17%)
Flash (KB)	100 (10%)	107 (5%)




Figure 8: MS-SSIM and bpp for the SunBath video from UVG [2] dataset. [0:k] shows the use of the first k channels (out of 12) for decoding.

limited hardware resources. Additionally, MCUCoder has 12 "stacked" channels in its latent space, which provides 12 levels of quality that can be dynamically adjusted based on the available network bandwidth. In Fig. 8, we illustrate the bpp and MS-SSIM for each frame in a video from the UVG dataset for all 12 levels of quality. The results show that using more channels for decoding leads to a higher MS-SSIM, which verifies the effectiveness of the proposed stochastic dropout training. The PSNR results are reported in the Appendix B.

Image compression: To assess the image compression capabilities of MCUCoder, we conduct experiments on the CLIC [53] and KODAK [58] datasets, see Fig. 7. The results in Table 1 show that MCUCoder achieves an impressive average bitrate reduction of 55.75% on the KODAK dataset and 49.54% on the CLIC dataset, compared to JPEG. The PSNR results are reported in the Appendix A.

Latent ordering and DCT-JPEG alignment: Fig. 4 shows the 12 latent channels obtained after training with the stochastic dropout method. The initial channels capture low-frequency information, while subsequent channels focus on high-frequency details. Interestingly, this behavior mirrors the Discrete Cosine Transform (DCT) basis matrix employed in JPEG compression.

Performance on MCUs: We implement MCUCoder using TFLite-Micro [56] and Zephyr RTOS [59] on STM32F7 and nRF5340 MCUs. The STM32F7 has 2 MB Flash, 2 MB RAM, and a Cortex-M7 processor, while the nRF5340 has 1 MB Flash, 512 KB RAM, and a Cortex-M33 processor, with both supporting DSP and CMSIS-NN acceleration [25]. As reported in Table 2, MCUCoder uses 360 KB of RAM on the STM32F7 and 344 KB on the nRF5340, which is remarkably low and suitable for such constrained IoT devices. To compare MCUCoder’s energy consumption against M-JPEG, we measured the energy consumption of MCUCoder and the optimized version of JPEG encoder for the Cortex-M series [60] on the nRF5340, see Fig. 9. The results indicate that MCUCoder matches JPEG’s energy consumption while significantly outperforming it in terms of BD-rate, see Table 7. The nRF5340 shows considerably slower performance than the STM32F7 for both MCUCoder and M-JPEG, suggesting that it is better suited for event-driven applications rather than real-time streaming.




Figure 9: Energy consumption of MCUCoder compared to M-JPEG for compressing one frame on the nRF5340.

5 Conclusion

We introduced MCUCoder, an ultra-lightweight asymmetric video compression model for resource-constrained IoT devices. With just 10.5K parameters and a 350KB memory footprint, compared to M-JPEG, MCUCoder reduces bitrate by over 55% on both the MCL-JCV and UVG datasets while matching the efficiency of M-JPEG. Its adaptive bitrate streaming ensures smooth video transmission under fluctuating network conditions, making it ideal for edge applications.

Acknowledgments and Disclosure of Funding

This project has received funding from the Federal Ministry for Digital and Transport under the CAPTN-Förde 5G project grant no. 45FGU139_H and Federal Ministry for Economic Affairs and Climate Action under the Marispace-X project grant no. 68GX21002E.

References

- [1] Haiqiang Wang, Weihao Gan, Sudeng Hu, Joe Yuchieh Lin, Lina Jin, Longguang Song, Ping Wang, Ioannis Katsavounidis, Anne Aaron, and C.-C. Jay Kuo. Mcl-jcv: A jnd-based h.264/avc video quality assessment dataset. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 1509–1513, 2016.
- [2] Alexandre Mercat, Marko Viitanen, and Jarno Vanne. Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 297–302, 2020.
- [3] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.
- [4] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5306–5314, 2017.
- [5] Jae-Han Lee, Seungmin Jeon, Kwang Pyo Choi, Youngo Park, and Chang-Su Kim. Dpict: Deep progressive image compression using trit-planes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16113–16122, 2022.
- [6] Seungmin Jeon, Kwang Pyo Choi, Youngo Park, and Chang-Su Kim. Context-based trit-plane coding for progressive image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14348–14357, 2023.
- [7] Jinming Liu, Heming Sun, and Jiro Katto. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14388–14397, June 2023.
- [8] Yueqi Xie, Ka Leong Cheng, and Qifeng Chen. Enhanced invertible encoding for learned image compression. In *Proceedings of the 29th ACM international conference on multimedia*, pages 162–170, 2021.
- [9] Yin hao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International Conference on Learning Representations*, 2022.
- [10] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2020.
- [11] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11006–11015, 2019.
- [12] Ji Lin, Wei-Ming Chen, Yujun Lin, Chuang Gan, Song Han, et al. Mccnet: Tiny deep learning on iot devices. *Advances in neural information processing systems*, 33:11711–11722, 2020.
- [13] Pan Hu, Junha Im, Zain Asgar, and Sachin Katti. Starfish: Resilient image compression for aiot cameras. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 395–408, 2020.
- [14] Colleen Josephson, Lei Yang, Pengyu Zhang, and Sachin Katti. Wireless computer vision using commodity radios. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, pages 229–240, 2019.

- [15] Saman Naderiparizi, Mehrdad Hesar, Vamsi Talla, Shyamnath Gollakota, and Joshua R Smith. Towards {Battery-Free}{HD} video streaming. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 233–247, 2018.
- [16] Bandhav Veluri, Collin Pernu, Ali Saffari, Joshua Smith, Michael Taylor, and Shyamnath Gollakota. Neuricam: Key-frame video super-resolution and colorization for iot cameras. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–17, 2023.
- [17] Manabu Nakanoya, Sai Shankar Narasimhan, Sharachchandra Bhat, Alexandros Anemogiannis, Akul Datta, Sachin Katti, Sandeep Chinchali, and Marco Pavone. Co-design of communication and machine inference for cloud robotics. *Autonomous Robots*, 47(5):579–594, 2023.
- [18] Vikram Iyer, Ali Najafi, Johannes James, Sawyer Fuller, and Shyamnath Gollakota. Wireless steerable vision for live insects and insect-scale robots. *Science robotics*, 5(44):eabb0839, 2020.
- [19] Ali Hojjat, Janek Haberer, Tayyaba Zainab, and Olaf Landsiedel. Limitnet: Progressive, content-aware image offloading for extremely weak devices & networks. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services, MOBISYS '24*, page 519–533, New York, NY, USA, 2024. Association for Computing Machinery.
- [20] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International conference on machine learning*, pages 5637–5664. PMLR, 2021.
- [21] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h.264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7):560–576, 2003.
- [22] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012.
- [23] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021.
- [24] William B Pennebaker and Joan L Mitchell. *JPEG: Still image data compression standard*. Springer Science & Business Media, 1992.
- [25] ARM-software. Cmsis-nn: Efficient neural network kernels for arm cortex-m cpus, 2024. Accessed: 2024-09-22.
- [26] Ren Yang, Fabian Mentzer, Luc Van Gool, and Radu Timofte. Learning for video compression with hierarchical quality and recurrent enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6628–6637, 2020.
- [27] Ge Gao, Pei You, Rong Pan, Shunyuan Han, Yuanyuan Zhang, Yuchao Dai, and Hojae Lee. Neural image compression via attentional multi-scale back projection and frequency decomposition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14677–14686, 2021.
- [28] Jinxi Xiang, Kuan Tian, and Jun Zhang. Mimt: Masked image modeling transformer for video compression. In *The Eleventh International Conference on Learning Representations*, 2023.
- [29] Fabian Mentzer, George Toderici, David Minnen, Sung-Jin Hwang, Sergi Caelles, Mario Lucic, and Eirikur Agustsson. Vct: A video compression transformer. *arXiv preprint arXiv:2206.07307*, 2022.
- [30] Gang He, Chang Wu, Lei Li, Jinjia Zhou, Xianglin Wang, Yunfei Zheng, Bing Yu, and Weiyang Xie. A video compression framework using an overfitted restoration neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 148–149, 2020.

- [31] Ties Van Rozendaal, Johann Brehmer, Yunfan Zhang, Reza Pourreza, Auke Wiggers, and Taco S Cohen. Instance-adaptive video compression: Improving neural codecs by training on the test set. *arXiv preprint arXiv:2111.10302*, 2021.
- [32] Mehrdad Khani, Vibhaalakshmi Sivaraman, and Mohammad Alizadeh. Efficient video compression via content-adaptive super-resolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4521–4530, 2021.
- [33] Zhihao Hu, Guo Lu, Jinyang Guo, Shan Liu, Wei Jiang, and Dong Xu. Coarse-to-fine deep video coding with hyperprior-guided mode prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5921–5930, 2022.
- [34] Zhihao Hu, Guo Lu, and Dong Xu. Fvc: A new framework towards deep video compression in feature space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1502–1511, 2021.
- [35] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34:18114–18125, 2021.
- [36] Jerry Liu, Shenlong Wang, Wei-Chiu Ma, Meet Shah, Rui Hu, Pranaab Dhawan, and Raquel Urtasun. Conditional entropy coding for efficient video compression. In *European Conference on Computer Vision*, pages 453–468. Springer, 2020.
- [37] Ho Man Kwan, Ge Gao, Fan Zhang, Andrew Gower, and David Bull. Hinerv: Video compression with hierarchical encoding-based neural representation. *Advances in Neural Information Processing Systems*, 36, 2024.
- [38] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34:21557–21568, 2021.
- [39] Fukashi Morishita, Norihito Kato, Satoshi Okubo, Takao Toi, Mitsuru Hiraki, Sugako Otani, Hideaki Abe, Yuji Shinohara, and Hiroyuki Kondo. A cmos image sensor and an ai accelerator for realizing edge-computing-based surveillance camera systems. In *2021 Symposium on VLSI Circuits*, pages 1–2. IEEE, 2021.
- [40] Suyao Ji, Jing Pu, Byong Chan Lim, and Mark Horowitz. A 220pj/pixel/frame cmos image sensor with partial settling readout architecture. In *2016 IEEE Symposium on VLSI Circuits (VLSI-Circuits)*, pages 1–2. IEEE, 2016.
- [41] Andrea Bejarano-Carbo, Hyochan An, Kyojin Choo, Shiyu Liu, Qirui Zhang, Dennis Sylvester, David Blaauw, and Hun-Seok Kim. Millimeter-scale ultra-low-power imaging system for intelligent edge monitoring. *arXiv preprint arXiv:2203.04496*, 2022.
- [42] Martin Lefebvre, Ludovic Moreau, Rémi Dekimpe, and David Bol. 7.7 a 0.2-to-3.6 tops/w programmable convolutional imager soc with in-sensor current-domain ternary-weighted mac operations for feature extraction and region-of-interest detection. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 64, pages 118–120. IEEE, 2021.
- [43] Davide Rossi, Francesco Conti, Manuel Eggiman, Stefan Mach, Alfio Di Mauro, Marco Guermandi, Giuseppe Tagliavini, Antonio Pullini, Igor Loi, Jie Chen, et al. 4.4 a 1.3 tops/w@ 32gops fully integrated 10-core soc for iot end-nodes with 1.7 μ w cognitive wake-up from mram-based state-retentive sleep mode. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 64, pages 60–62. IEEE, 2021.
- [44] Han Xu, Ziru Li, Ningchao Lin, Qi Wei, Fei Qiao, Xunzhao Yin, and Huazhong Yang. Macsen: A processing-in-sensor architecture integrating mac operations into image sensor for ultra-low-power bnn-based intelligent visual perception. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 68(2):627–631, 2020.
- [45] Shuochao Yao, Jinyang Li, Dongxin Liu, Tianshi Wang, Shengzhong Liu, Huajie Shao, and Tarek Abdelzaher. Deep compressive offloading: Speeding up neural network inference by trading edge computation for network latency. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 476–488, 2020.

- [46] Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5718–5727, 2022.
- [47] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7948, 2020.
- [48] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [49] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [50] Ali Hojjat, Janek Haberer, and Olaf Landsiedel. Progtdt: Progressive learned image compression with double-tail-drop training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1130–1139, June 2023.
- [51] Janek Haberer, Ali Hojjat, and Olaf Landsiedel. Hydravit: Stacking heads for a scalable vit. *arXiv preprint arXiv:2409.17978*, 2024.
- [52] Eastman kodak (1993). kodak lossless true color image suite (photocd pcd0992). <https://r0k.us/graphics/kodak>.
- [53] Workshop and challenge on learned image compression (clic). <http://www.compression.cc>, 2020. 6.
- [54] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [55] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14771–14780, 2021.
- [56] TensorFlow Lite Micro. <https://www.tensorflow.org/lite/micro>. Accessed: [2023].
- [57] G Bjøtegaard. Calculation of average psnr differences between rd-curves (vceg-m33). In *VCEG Meeting (ITU-T SG16 Q. 6), Austin, Texas, USA., Tech. Rep. M*, volume 16090, 2001.
- [58] Eastman Kodak. Kodak lossless true color image suite (photocd pcd0992). <http://r0k.us/graphics/kodak>, 1993. 6.
- [59] Zephyr Project RTOS. <https://github.com/zephyrproject-rtos/zephyr>. Accessed: [2023].
- [60] JPEGEncoder4Cortex. <https://github.com/noritsuna/JPEGEncoder4Cortex-M/tree/master>.

A PSNR on KODAK and CLIC datasets

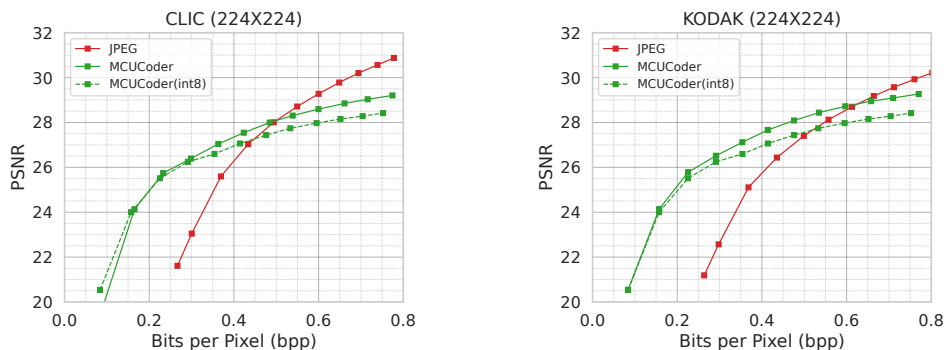


Figure 10: Comparison of MCUCoder (quantized and non-quantized model) and baselines on the KODAK [52] and CLIC [53] datasets. All datasets are resized to 224×224 . Since MCUCoder is specifically designed for IoT environments—where structural integrity is prioritized over fine details—it has been optimized for MS-SSIM. Consequently, M-JPEG achieves better PSNR performance at higher bpp.

B PSNR on the MCL-JCV and UVG datasets

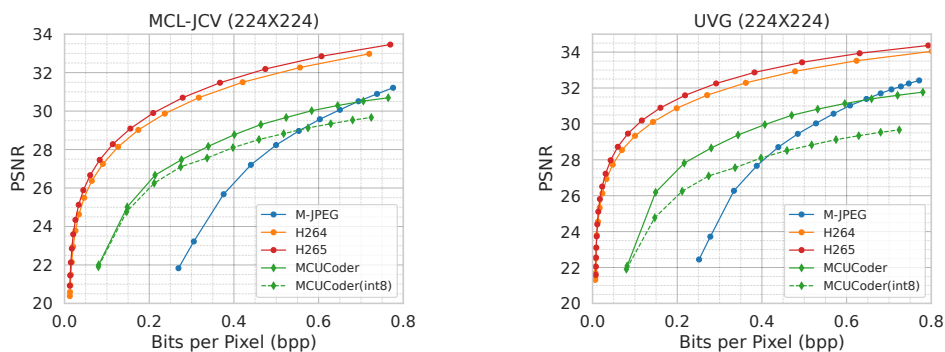


Figure 11: Comparison of MCUCoder and baselines on MCL-JCV [1] and UVG [2]. H.264 and H.265 are included for reference, though impractical for MCUs due to hardware demands. MCUCoder, optimized for IoT with a focus on MS-SSIM, prioritizes structural integrity, while M-JPEG shows better PSNR at higher bpp.

C Training logs of MCUCoder




Figure 12: MS-SSIM values on the KODAK dataset during training. The notation $[0 : k]$ represents the MS-SSIM of the reconstructed image using the first k latent channels out of a total of 12. As shown, with stochastic dropout training, all the sub-latents can be trained simultaneously without overfitting to any particular sub-latent.

D Examples of MCUCoder



Figure 13: Some samples from the MCL-JCV [1] dataset. The columns represent different frames, while the rows display progressively improving levels of quality from top to bottom, produced by MCUCoder.




Figure 14: Some samples from the MCL-JCV [1] dataset. The columns represent different frames, while the rows display progressively improving levels of quality from top to bottom, produced by MCUCoder.