

SCALED : Surrogate-gradient for Codec-Aware Learning of Downsampling in ABR Streaming

Esteban Pesnel^{*†}, Julien Le Tanou^{*}, Michael Ropert^{*}, Thomas Maugey[†], Aline Roumy[†]

^{*}MediaKind, Rennes, France

{esteban.pesnel, julien.letanou, michael.ropert}@mediakind.com

[†]INRIA, Rennes, France

{esteban.pesnel, aline.roumy, thomas.maugey}@inria.fr

Abstract—The rapid growth in video consumption has introduced significant challenges to modern streaming architectures. Over-the-Top (OTT) video delivery now predominantly relies on Adaptive Bitrate (ABR) streaming, which dynamically adjusts bitrate and resolution based on client-side constraints such as display capabilities and network bandwidth. This pipeline typically involves downsampling the original high-resolution content, encoding and transmitting it, followed by decoding and upsampling on the client side. Traditionally, these processing stages have been optimized in isolation, leading to suboptimal end-to-end rate-distortion (R-D) performance. The advent of deep learning has spurred interest in jointly optimizing the ABR pipeline using learned resampling methods. However, training such systems end-to-end remains challenging due to the non-differentiable nature of standard video codecs, which obstructs gradient-based optimization. Recent works have addressed this issue using differentiable proxy models, based either on deep neural networks or hybrid coding schemes with differentiable components such as soft quantization, to approximate the codec behavior. While differentiable proxy codecs have enabled progress in compression-aware learning, they remain approximations that may not fully capture the behavior of standard, non-differentiable codecs. To our knowledge, there is no prior evidence demonstrating the inefficiencies of using standard codecs during training. In this work, we introduce a novel framework that enables end-to-end training with real, non-differentiable codecs by leveraging data-driven surrogate gradients derived from actual compression errors. It facilitates the alignment between training objectives and deployment performance. Experimental results show a 5.19% improvement in BD-BR (PSNR) compared to codec-agnostic training approaches, consistently across the entire rate-distortion convex hull spanning multiple downsampling ratios.

I. INTRODUCTION

Video consumption has grown exponentially during the last decades, representing more than 70% of network traffic [1]. With the emergence of HTTP-based video streaming solutions such as DASH [2] or HLS [3], modern video platforms heavily rely on Adaptive Bitrate streaming (ABR), which dynamically adapts the bit rate and resolution of the delivered content to meet client bandwidth requirements and display. The high-resolution content is first downsampled through various scale ratios $s \in \mathbb{Q}^+$ and separately encoded. Then, the bit stream is decoded by the client and upsampled to the device display resolution.

Traditional downsampling methods such as bicubic, lanczos or bilinear filters [4], [5] rely on linear content-invariant filters, targeting low-pass filtering prior to subsampling. These filters

enjoy low computational complexity, despite visible artifacts (e.g. blurring, ringing, etc.). Recognizing the limitations of content-invariant designs, non-linear methods adapt to local image structure. For example, bilateral filtering [6] combines spatial proximity and intensity similarity to perform edge-aware smoothing, preserving fine structures while reducing noise. More recent adaptive downsampling strategies leverage semantic content and visual saliency to guide resolution reduction [7]–[10]. While classical methods remain integral to video streaming systems, they treat downsampling independently of subsequent processing stages—particularly upsampling and compression. This decoupling often results in sub-optimal reconstructions, especially when aggressive compression introduces quantization and blocking artifacts. An emerging direction addresses this by co-designing downsamplers with knowledge of the upscaler and codec.

With the emergence of deep learning, several studies have focused on integrating deep learned pre- or post-processing around traditional codecs. Among these solutions, several works introduced neural network model for resampling tasks. Joint learning of downsampling and upsampling [11] leads to significant improvements over traditional filtering baselines. However, learned upsamplers increase the client-side computational complexity and may conflict with standardization requirements for technology deployment and mass adoption. To mitigate this problem, other studies [12], [13] propose CNN-based learned downsampling based on bicubic upsampler at training, subject to end-to-end MSE loss function minimization. However, these models are trained without considering the compression process due to its non-differentiable nature, which is suboptimal for video streaming applications. Indeed, the non-differentiability poses a challenge for gradient-based optimizers, in order to minimize a convex objective function.

As a result, recent studies have attempted to overcome this issue by modeling the codec with a differentiable proxy. Son *et al.* [14] propose to mimic traditional codec behavior using deep neural networks. More recently, Guleryuz *et al.* [15] introduced a differentiable codec proxy that implements the classical paradigm of hybrid video coding. Specifically, it incorporates a straight-through quantizer, a motion estimation module derived from UFlow [16], and a loop filter proxy based on U-Net [17]. The usage of a differentiable codec during the training tends to improve the learning of pre- and post-

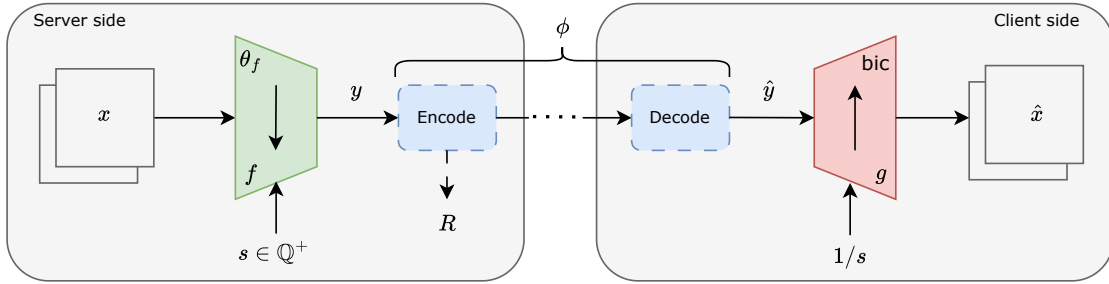


Fig. 1: ABR streaming pipeline. The server downsamples and encodes the high-resolution video at scale ratio s , then the client decodes and upsamples it to match the display. Blue dashed boxes are non-differentiable; the red box is fixed.

processing networks, such as downsampling and upsampling. However, the training efficiency deeply relies on the codec proxy accuracy. Indeed, learned codec proxies are trained to minimize the MSE with real codecs, which may be insufficient to fully capture the specific characteristics of video codecs.

In this context, we propose an alternative learning strategy that integrates the ground-truth standard codec into the downsampling optimization. To address the codec’s non-differentiability, we introduce a surrogate gradient in the backward pass, defined as the derivative of the standard deviation of the ground-truth compression error, modulated by the compression error itself. While during the forward pass, the actual compression error is directly added into the optimization process. We then propose and assess two learning scheme variants based on this strategy against codec-agnostic training schemes. In the first variant, we aim to minimize the reconstruction distortion by leveraging the actual coding artifacts at a given bit rate. In the second variant, we incorporate a rate constraint into the training loss, leveraging on a differentiable rate proxy from the state of the art to enable efficient rate-distortion optimization. Our contribution outperforms codec agnostic baseline for deep learned downsampling by 5.19% in BD-BR(PSNR) over then entire bit rate-quality convex hull of different scaling ratios, demonstrating the effectiveness of using standard, non-differentiable codecs during training.

II. TRAINING SCHEME FOR DOWNSAMPLING OPTIMIZATION

A. Problem formulation

A typical ABR streaming scheme is illustrated in Figure 1. The client first requests a video chunk at a specific bit rate from the server. The high-resolution content x is then downsampled using a scale ratio s and encoded with a standard codec ϕ . The resulting representation is transmitted to the client, which decodes and upsamples the content to match the display resolution, yielding a reconstructed content \hat{x} .

To optimize this video delivery pipeline, we introduce a learned downsampling function with associated weights $f(x; \theta_f)$, and adopt a fixed bicubic upsampler g for reconstruction. The selection of a linear fixed filter for the upsampling task is motivated by the will to avoid non-standard

and computationally intensive technology at the end-client, enabling broad solution adoption. Learned downsampler f aims at minimizing the end-to-end distortion D , for the smallest bit rate budget R , given by a ground-truth rate r .

$$\theta_f^* = \arg \min_{\theta_f} \|x - g(\phi(f(x; \theta_f)))\|_2^2 + \lambda r(\phi(f(x; \theta_f))) \quad (1)$$

However, both ϕ and r are non differentiable, preventing the direct optimization of (1).

B. Existing methods and limitations

Few studies relaxed this constraint by optimizing the downsampler independently of the codec, which was excluded from the training process. Li *et al.* [12] proposed a 10-layer convolutional neural network (CNN) for image downsampling task. Building on a similar architecture, Chen *et al.* [13] extended the approach to support arbitrary fractional down-scaling ratios, with $s \in \mathbb{Q}^+$. The original content x is first downsampled. The low-resolution content \tilde{x} is then bicubic upsampled through a linear function g . The reconstructed signal is then compared to the original one, leading to a codec-agnostic D-only minimization:

$$\theta_f^* = \arg \min_{\theta_f} \|x - g(f(x; \theta_f))\|_2^2 \quad (2)$$

However, training without considering the compression process is sub-optimal for video streaming applications, as both the distortion induced by compression and output bit rate are neglected.

Hence, codec-aware training scheme becomes essential to obtain significant gain in rate distortion. Thus, some recent works [14], [18]–[20] have attempted to model the codec process in a differentiable manner, by designing $\hat{\phi}$ and \hat{r} proxy functions. A common approach among these methods relies on CNN [18] or U-Net [14] to mimic classical codec behavior. Alternatively, other works introduced codec proxy based on hybrid image or video compression scheme. Xing *et al.* [21] introduced a JPEG estimator based on DCT and soft quantization. Guleryuz *et al.* [15] extend this paradigm to video compression by incorporating a motion estimation proxy based on U-Flow [16], a loop filter proxy based on U-Net [17] and a straight-through estimated quantizer [22].

All of these methods implement differentiable entropy estimators or rate proxies, \hat{r} , to enable end-to-end Rate Distortion optimization. Although numerous approaches rely on deep learning models to estimate bit rate [14], [19], [20], some studies exploit the DCT coefficients of the compressed image, modeling the rate through affine mappings [15]. The corresponding optimization function is :

$$\theta_f^* = \arg \min_{\theta_f} \|\mathbf{x} - g(\hat{\phi}(f(\mathbf{x}; \theta_f)))\|_2^2 + \lambda \hat{R} \quad (3)$$

Here, the set of frames \mathbf{x} is first downsampled through f . The codec proxy $\hat{\phi}$ then produces a sequence of frames that mimic those of a compressed video. The rate proxy \hat{r} estimates the associated bit rate \hat{R} . The result is then bicubic upsampled and compared to \mathbf{x} . We then obtain a rate-distortion optimization problem, based on codec and rate proxies.

However, another challenge arises from codec proxy training. Most proxy models are optimized using a mean squared error loss.

$$\theta_{\hat{\phi}}^* = \arg \min_{\theta_{\hat{\phi}}} \|\phi(\mathbf{x}) - \hat{\phi}(\mathbf{x})\|_2^2 \quad (4)$$

However, the MSE metric may be insufficient to capture the true characteristics and artifacts introduced by the codec, such as ringing or blocking effects. As a result, even if the proxy minimizes the MSE, it may fail to reproduce these codec-specific artifacts accurately. This mismatch is sub-optimal for the optimization of the overall objective (4).

More specifically, we tested the codec proxy from sandwiched video compression [15] in the context of ABR streaming. We trained a deep downsampler alongside the codec proxy, with fixed upsampler for each considered bit rate. An example of the obtained rate-distortion curve is shown in Figure 2. We can see that the solution based on codec proxy slightly improves the codec-agnostic solution in high and medium bit rates but obtains lower performance at low bit rate. This weakness may be explained by several limitations in the proxy design. First, quantization error is neglected during backpropagation. Moreover, prediction proxies are estimated on source, and syntax costs are omitted - making the estimators invariant to the target rate.

Instead of relying on a codec proxy, we opted to directly use a real, non-differentiable codec during training. While codec proxies can facilitate gradient-based optimization, their design is highly non-trivial and must be carefully engineered to accurately emulate the behavior of real encoders. In contrast, real codecs already implement the full set of compression tools and can be readily integrated into the pipeline. This choice enables more faithful supervision and simplifies the training setup, while opening the door to further optimization strategies without requiring additional proxy design efforts.

III. PROPOSED APPROACH

Learning with non-differentiable function is a fundamental research topic in machine learning. Several works have attempted to circumvent this issue, notably using the Straight-Through Estimator (STE) [22], the injection of additive noise

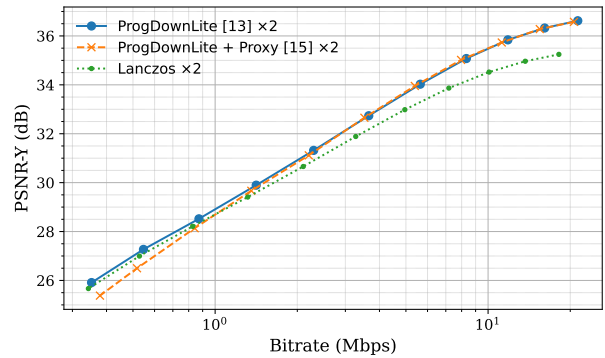


Fig. 2: Rate-distortion curves of WestWindEasy sequence with scale $s = 1/2$, comparing ProgDownLite [13] with and without the codec proxy [15], against a Lanczos baseline.

from uniform distribution [23], or soft-to-hard annealing [24]. They all aimed at addressing the non-differentiability introduced by components such as quantization in codecs. A common way to implement STE, and handle non-differentiable components, is to use the stop-gradient (sg) operator, which explicitly blocks the gradient flow during backpropagation through selected terms of the computation graph. The function $\text{sg}(\mathbf{x})$ acts as the identity function in the forward pass but returns a zero gradient in the backward pass.

$$\text{sg}(\mathbf{x}) = \begin{cases} \mathcal{F} : \mathbf{x} \\ \nabla \mathcal{L} : \frac{\partial \text{sg}(\mathbf{x})}{\partial \mathbf{x}} = 0 \end{cases} \quad (5)$$

Here, \mathcal{F} denotes the forward pass and $\nabla \mathcal{L}$ the backward pass. Thus, \mathbf{x} is considered as a static constant during the forward pass. The application of this operator for STE implementation in the context of codec-aware design is used to detach the compression error ϵ from the optimization path:

$$\hat{\mathbf{y}} = \underbrace{f(\mathbf{x}; \theta_f)}_{\mathbf{y}} + \text{sg}(\underbrace{\phi(f(\mathbf{x}; \theta_f)) - f(\mathbf{x}; \theta_f)}_{\epsilon}) \quad (6)$$

Here, ϕ denotes the true codec function, and ϵ is treated as additive noise during the forward pass, whose unknown gradient is set to zero during the backward pass. We explored the use of the Straight-Through Estimator (STE) in our context. However, this approach led to the same challenges identified by Mack *et al.* [25], which confirms the limitations of regular STE for compression applications. Indeed, since $\frac{\partial \hat{\mathbf{y}}}{\partial \mathbf{y}} = 1$, the optimization disconnects the compression error in the gradient flow which leads to an increase in $|f(\mathbf{x}; \theta_f)|_1$, that prevents from proper convergence during training [25].

To address this issue, we propose a method based on a modification of the STE from Mack *et al.* [25] to enhance the surrogate gradient of the compression error during backpropagation, as shown in (7) and (8).

$$\hat{\mathbf{y}} = \mathbf{y} + \text{sg}(\epsilon) \left[\frac{\sigma(\epsilon)}{\text{sg}(\sigma(\epsilon))} \right] \quad (7)$$

We denote by σ the standard deviation operator, and by $\hat{\mathbf{y}}$ the encoded version of the downsampled content with the

proposed equation. The modified-STE corresponds to a reparameterization trick [26] to account for the compression noise in the backpropagation gradient, resulting in the definition of an enhanced surrogate gradient based on the derivative of the standard deviation of the true compression error, modulated by the compression error itself.

$$\left(\frac{\partial \hat{\mathbf{y}}}{\partial \mathbf{y}}\right)_{\text{surrogate}} = \mathbf{I} + \frac{\boldsymbol{\epsilon}}{\sigma(\boldsymbol{\epsilon})} \frac{\partial \sigma(\boldsymbol{\epsilon})}{\partial \mathbf{y}} = \mathbf{I} - \frac{\boldsymbol{\epsilon}(\boldsymbol{\epsilon} - \bar{\boldsymbol{\epsilon}})^T}{N\sigma^2(\boldsymbol{\epsilon})} \quad (8)$$

where $\bar{\boldsymbol{\epsilon}}$ is a vector containing the average value of $\boldsymbol{\epsilon}$.

Thus, we now approximate, during backpropagation, how changes in input \mathbf{y} influence the compression error, using its magnitude and statistical spread. The term $\frac{\boldsymbol{\epsilon}}{\sigma(\boldsymbol{\epsilon})}$ acts as a normalized measure of the quantization noise, that stabilizes gradient flow, while the derivative of sigma w.r.t. \mathbf{y} reflects how the noise distribution changes with \mathbf{y} . This gives a dynamically scaled surrogate gradient for ϵ sensitive to compression error variability. Then, leveraging on (7), we will consider the two following proposals with true compression errors:

- an MSE-based training scheme with non-differentiable codec, based on (2) and (7).
- a Rate-Distortion training scheme with non-differentiable codec, based on (3) and (7).

A. Distortion based model with non-differentiable codec

In this variant, (7) is substituted in (2), resulting in the new optimization formulation presented in (10):

$$\boldsymbol{\theta}_f^* = \arg \min_{\boldsymbol{\theta}_f} \left\| \mathbf{x} - g(f(\mathbf{x}; \boldsymbol{\theta}_f) + \text{sg}(\boldsymbol{\epsilon}) \cdot \frac{\sigma(\boldsymbol{\epsilon})}{\text{sg}(\sigma(\boldsymbol{\epsilon}))}) \right\|_2^2 \quad (9)$$

Where $\boldsymbol{\epsilon} = \phi(f(\mathbf{x}; \boldsymbol{\theta}_f)) - f(\mathbf{x}; \boldsymbol{\theta}_f)$, ϕ denotes the true compression function. The new network trained using formulation (10) will be referred to as SCALED_D in the following sections, minimized subject to a certain bit rate budget given to the true encoder.

B. Rate-distortion based model with non-differentiable codec

To account for the trade-off between distortion and rate, we now formulate an optimization objective based on the $D + \lambda \hat{R}$ trade-off. We then leverage on (7), (3) and (1), resulting in the new optimization formulation presented in (11)

$$\boldsymbol{\theta}_f^* = \arg \min_{\boldsymbol{\theta}_f} \left\| \mathbf{x} - g(f(\mathbf{x}; \boldsymbol{\theta}_f) + \text{sg}(\boldsymbol{\epsilon}) \cdot \frac{\sigma(\boldsymbol{\epsilon})}{\text{sg}(\sigma(\boldsymbol{\epsilon}))}) \right\|_2^2 + \lambda \hat{R} \quad (10)$$

Here, \hat{R} denotes a state-of-the-art rate estimation module from Google Sandwiched Compression [15], based on differentiable approximation of the ℓ_0 norm of DCT coefficients. The objective of this variant is to demonstrate the benefits of rate-distortion optimization. The new network trained using formulation (11) will be referred to as SCALED_{RD} in the following sections.

IV. EXPERIMENTS

A. Implementation details

To ensure the practical relevance of our learning framework, we base our experiments on a codec that reflects real-world deployment scenarios. We therefore adopt the H.264 standard via its open-source implementation x264 [28]. We use default medium preset, both for training and experiments.

We adopt the ProgDownLite network [13] as a learned downsampling function f . This model is based on a five-layer convolutional network for image pre-processing, followed by bilinear downsampling. Subsequently, five additional convolutional layers aim to restore the image details that were suppressed during filtering. Since the resampling is performed by a bilinear downsampling, it supports non-integer scale ratio $s \in \mathbb{Q}^+$, which is mandatory in numerous applications. We fixed the upsampler to a bicubic upsampling.

We use the video dataset from Google, which contains 800 video sequences, in YUV444 colorspace. Each one has a resolution of 256×256 pixels and a length of 10 frames. The chroma channels are first bilinearly downsampled and then bicubically upsampled to simulate the chroma degradation of YUV 4:2:0 content used in the experimental dataset.

B. Performance assessment with compression

We consider a multi-scale ratio compression scheme, where $s \in \{2/3, 1/2, 2/5, 1/3, 1/4, 1/5\}$. The downsampled contents are then encoded with x264, with the quantization parameter (QP) varying from 17 to 47 in increments of 3. The downsampled, encoded content is then decoded and upsampled back to the original content resolution. We then computed the convex hull of best operating rate-distortion points across the scale ratios set. We trained one SCALED model for each QP and scale ratios.

To assess the downsampling filters performance in rate-distortion context, we select 24 videos from Xiph dataset, and 7 videos from UVG dataset [27], all in 1920×1080 spatial resolution. For models trained with RGB dataset, each YUV420 video was prior converted to RGB using FFmpeg.

Table I presents the rate-distortion evaluations. We first reproduce the training conditions of ProgDownLite [13] using DIV2K dataset, referred to as ProgDownLite_{RGB}. We also re-train ProgDownLite on sandwiched video dataset [15], denoted ProgDownLite_{YUV}. Finally, we include the results of our two main contributions, SCALED_D and SCALED_{RD}, both trained in the context of non-differentiable codecs and video dataset, and compared to traditional STE. All the models are trained on 100 epochs, with $lr = 1e^{-4}$ and $(\beta_1, \beta_2) = (0.9, 0.999)$.

It is worth noting that codec agnostic solutions show marginal gains over Lanczos baseline. Indeed, due to the lack of codec consideration at the training, they tend to maintain high frequencies costly to code, resulting in a substantial rate increase. On the other hand, ProgDownLite_{YUV} slightly improves the MSE of ProgDownLite_{RGB} due to the training dataset nature - videos instead of pictures.

Then, our SCALED contributions outperform codec-agnostic solutions, thanks to the consideration of the actual

Filter	Training dataset	Training strategy	Codec aware	XIPH				UVG [27]			
				PSNR	SSIM	VMAF	VMAF _{NEG}	PSNR	SSIM	VMAF	VMAF _{NEG}
ProgDownLite _{RGB}	DIV2K	D-only (2)	No	-0.66	-1.24	-5.67	-3.58	1.22	-1.78	-7.72	-5.24
ProgDownLite _{YUV}	Google [15]	D-only (2)	No	-0.85	-1.19	-4.71	-3.60	1.91	-2.31	-6.11	-5.02
ProgDownLite _{YUV}	Google [15]	STE (6)	Yes	68.57	48.34	11.22	20.69	49.39	45.33	4.46	5.06
ProgDownLite _{YUV}	Google [15]	Proxy [15]	Yes	-1.47	-1.87	-2.58	-1.67	-0.08	-2.99	-5.88	-4.79
ProgDownLite _{YUV}	Google [15]	SCALED _D (10)	Yes	-4.67	-5.21	-9.80	-7.78	-2.84	-5.79	-10.95	-9.07
ProgDownLite _{YUV}	Google [15]	SCALED _{RD} (11)	Yes	-6.11	-4.29	-8.80	-7.59	-5.07	-5.21	-8.88	-8.02

TABLE I: Comparison of filtering methods on the XIPH and UVG datasets. Metrics are computed as BD-BR values with respect to the convex hull of the Lanczos filter, used as the reference rate-distortion baseline.

codec. Applying a rate constraint with SCALED_{RD} increase the performance in MSE, though slightly degrading other metrics.

V. CONCLUSION

In this paper, we introduce SCALED, a coding-aware downsampling learning framework that leverages data-driven surrogate gradients derived from actual compression errors to optimize the downsampling process end-to-end. The proposed approach demonstrates that surrogate gradients can effectively enable the integration of non-differentiable, real codecs into end-to-end training pipelines. Furthermore, we proposed a hybrid approach in which distortion is estimated using the actual codec, while the rate is approximated via a state-of-the-art differentiable rate proxy. This combination yields improved performance over purely distortion-based methods. Our findings suggest promising directions for developing deep downsamplers, or any learned pre-processing, that are explicitly optimized with respect to real-world codec behavior.

REFERENCES

- [1] Ericsson, "Ericsson mobility report," Ericsson, Tech. Rep., 2024, <https://www.ericsson.com/en/reports-and-papers/mobility-report>.
- [2] ISO/IEC, "Information technology — Dynamic adaptive streaming over HTTP (DASH) — Part 1: Media presentation description and segment formats," <https://www.iso.org/standard/77975.html>, 2022, iSO/IEC 23009-1:2022.
- [3] R. Pantos and W. May, "HTTP Live Streaming," <https://datatracker.ietf.org/doc/html/rfc8216>, 2017, rFC 8216.
- [4] R. Keys, "Cubic convolution interpolation for digital image processing," *IEEE transactions on acoustics, speech, and signal processing*, vol. 29, no. 6, pp. 1153–1160, 2003.
- [5] C. E. Duchon, "Lanczos filtering in one and two dimensions," *Journal of Applied Meteorology (1962-1982)*, pp. 1016–1022, 1979.
- [6] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*. IEEE, 1998, pp. 839–846.
- [7] J. Kopf, A. Shamir, and P. Peers, "Content-adaptive image downscaling," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, pp. 1–8, 2013.
- [8] A. C. Oeztireli and M. Gross, "Perceptually based downscaling of images," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–10, 2015.
- [9] N. Weber, M. Waechter, S. C. Amend, S. Guthe, and M. Goesele, "Rapid, detail-preserving image downscaling," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 205–1, 2016.
- [10] J. Liu, S. He, and R. W. Lau, " $L_{\{0\}}$ -regularized image downscaling," *IEEE Transactions on Image Processing*, vol. 27, no. 3, pp. 1076–1085, 2017.
- [11] W. Sun and Z. Chen, "Learned image downscaling for upscaling using content adaptive resampler," *IEEE Transactions on Image Processing*, vol. 29, pp. 4027–4040, 2020.
- [12] Y. Li, D. Liu, H. Li, L. Li, Z. Li, and F. Wu, "Learning a convolutional neural network for image compact-resolution," *IEEE Transactions on Image Processing*, vol. 28, no. 3, pp. 1092–1107, 2018.
- [13] L.-H. Chen, C. G. Bampis, Z. Li, J. Sole, C. Chen, and A. C. Bovik, "Learned fractional downsampling network for adaptive video streaming," *Signal Processing: Image Communication*, vol. 128, p. 117172, 2024.
- [14] H. Son, T. Kim, H. Lee, and S. Lee, "Enhanced standard compatible image compression framework based on auxiliary codec networks," *IEEE Transactions on Image Processing*, vol. 31, pp. 664–677, 2021.
- [15] O. G. Guleryuz, P. A. Chou, B. Isik, H. Hoppe, D. Tang, R. Du, J. Taylor, P. Davidson, and S. Fanello, "Sandwiched compression: Repurposing standard codecs with neural network wrappers," *arXiv preprint arXiv:2402.05887*, 2024.
- [16] S. K. Lodha, A. Pang, R. E. Sheehan, and C. M. Wittenbrink, "Uflow: Visualizing uncertainty in fluid flow," in *Proceedings of Seventh Annual IEEE Visualization '96*. IEEE, 1996, pp. 249–254.
- [17] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [18] P. Eusébio, J. Ascenso, and F. Pereira, "Optimizing an image coding framework with deep learning-based pre-and post-processing," in *2020 28th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 506–510.
- [19] M. Guo, S. Zhao, H. Jiang, J. Li, and L. Zhang, "Video compression with arbitrary rescaling network," *arXiv preprint arXiv:2306.04202*, 2023.
- [20] Z. Chen and L. Chen, "An enhanced video compression framework based on rescaling networks," in *2023 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*. IEEE, 2023, pp. 1–6.
- [21] J. Xing, W. Hu, M. Xia, and T.-T. Wong, "Scale-arbitrary invertible image downscaling," *IEEE Transactions on Image Processing*, vol. 32, pp. 4259–4274, 2023.
- [22] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *CoRR*, vol. abs/1308.3432, 2013.
- [23] J. Ballé, V. Laparra, and E. P. Simoncelli, "End-to-end optimized image compression," *arXiv preprint arXiv:1611.01704*, 2016.
- [24] E. Agustsson, F. Mentzer, M. Tschannen, L. Cavigelli, R. Timofte, L. Benini, and L. V. Gool, "Soft-to-hard vector quantization for end-to-end learning compressible representations," *Advances in neural information processing systems*, vol. 30, 2017.
- [25] W. Mack, A. Mustafa, R. Łaganowski, and S. Hijazy, "Efficient evaluation of quantization-effects in neural codecs," *arXiv preprint arXiv:2502.04770*, 2025.
- [26] D. P. Kingma, M. Welling *et al.*, "Auto-encoding variational bayes," 2013.
- [27] A. Mercat, M. Viitanen, and J. Vanne, "Uvg dataset: 50/120fps 4k sequences for video codec analysis and development," in *Proceedings of the 2020 ACM Multimedia Systems Conference (MMSys '20)*, Istanbul, Turkey, 2020, pp. 297–302.
- [28] J.-B. Kempf and VideoLAN, "x264: a free software library for encoding video streams into the h.264/mpeg-4 avc format," <https://www.videolan.org/developers/x264.html>, 2024, accessed: 2025-07-17.