

# DANICE: Domain adaptation without forgetting in neural image compression

Sudeep Katakol<sup>1\*</sup> Luis Herranz<sup>2</sup> Fei Yang<sup>2</sup> Marta Mrak<sup>3</sup>

<sup>1</sup>Univ. of Michigan, Ann Arbor <sup>2</sup>Computer Vision Center, UAB, Barcelona <sup>3</sup>BBC R&D, London

## Abstract

Neural image compression (NIC) is a new coding paradigm where coding capabilities are captured by deep models learned from data. This data-driven nature enables new potential functionalities. In this paper, we study the adaptability of codecs to custom domains of interest. We show that NIC codecs are transferable and that they can be adapted with relatively few target domain images. However, naive adaptation interferes with the solution optimized for the original source domain, resulting in forgetting the original coding capabilities in that domain, and may even break the compatibility with previously encoded bitstreams. Addressing these problems, we propose Codec Adaptation without Forgetting (CAwF), a framework that can avoid these problems by adding a small amount of custom parameters, where the source codec remains embedded and unchanged during the adaptation process. Experiments demonstrate its effectiveness and provide useful insights on the characteristics of catastrophic interference in NIC.

## 1. Introduction

Lossy image and video coding have been the cornerstone of visual content sharing and communication, achieving high compression rates by allowing a small amount of distortion. With the recent advances in machine learning with neural networks, neural image compression (NIC) [13, 3, 12, 14] has emerged as a new powerful paradigm where codecs optimize their parameters directly using data from the domain of interest, leading to very competitive coding performance. The main limitations for practical deployment are high memory and computation requirements and lack of flexibility. Recent works have addressed some of these practical limitations, including memory and computational complexity [18, 8], and variable rate [13, 7, 19, 6, 12]. But the NIC paradigm can also enable novel functionalities.

Traditional coding formats (such as those from JPEG [15] and MPEG families [17, 11, 5]) follow a care-

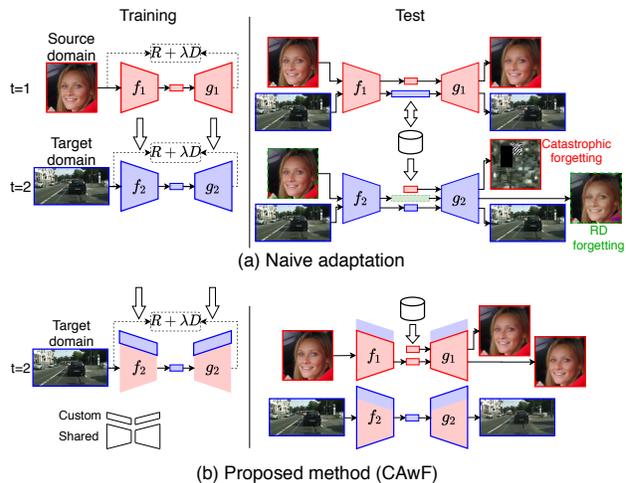


Figure 1. Domain adaptation and forgetting in NIC: (a) naive adaptation results in forgetting, while (b) CAwF prevents forgetting.

fully designed syntax and reference decoder, which are strictly defined by standards to enable interoperability and backward compatibility. In contrast, NIC codecs discover an implicit latent syntax through learnable parametric models. These parameters could be readjusted to improve the performance in a custom domain or functionality. Fig. 1 shows possible scenarios involving such type of customization, and illustrates potential compatibility issues. For instance, an off-the-shelf learned codec optimized for a certain domain of interest (shown in red at  $t = 1$  in Fig. 1a) may not necessarily be optimal for another domain, with performance suffering due to higher bitrate and/or higher distortion. The trainable nature of learned codecs is a convenient mechanism to customize the off-the-shelf codec, resulting in an adapted codec with lower rate and distortion in that domain of interest (shown in blue in Fig. 1a at  $t = 2$ ). There are many potential use cases of this functionality, such as privately adapted codecs for photographers' personal collections, or an on-board camera with a codec adapted to the specific vehicle. In traditional codecs, adapting to specific applications would require designing specific tools or even new coding formats. In NIC, this functionality can be reinterpreted as domain adaptation [16], which in machine learning refers to the reuse and adaptation of the

\* This work was done during an internship at the CVC, Barcelona.

knowledge [1] learned in a source model to improve performance in a target domain where data is scarce.

The main limitation of such codec adaptation process is that the new parameters are not optimal for the source domain anymore (phenomenon known as forgetting [9] in continual learning [10], see Fig. 1a). This can manifest as lower rate-distortion (RD) performance when encoding images of the source domain (i.e. RD forgetting). This drop in performance is more significant in small models, often required in resource-limited scenarios or when efficiency is a concern. Furthermore, the codec becomes incompatible with bitstreams encoded with the original version of the codec. While the decoder is able to generate an image, the result is disastrous (i.e. catastrophic forgetting). Thus, preventing forgetting is crucial when the codec must remain compatible and with good performance in the source domain.

In this work, we first study the adaptability of learned codecs. For cases where forgetting is a concern, we then propose an incremental architecture (see Fig. 1b) that prevents both RD and catastrophic forgetting by fixing the source codec and using a small number of additional custom parameters to learn target-specific patterns. In this way, the previous version of the codec remains embedded in the codec and can be used independently when necessary.

In summary, the main contributions of this work are:

- We introduce the problem of domain adaptation in neural image compression (DANICE), and the related problem of forgetting the original coding capabilities. In this way we connect concepts of traditional image coding and machine learning, such as backward compatibility, domain adaptation and forgetting.
- We study the adaptability of NIC codecs and propose selective fine tuning for cases with few target images.
- We characterize how forgetting manifests in NIC codecs, proposing a framework that prevents forgetting by design, while ensuring optimal rate-distortion performance and backward compatibility.

## 2. Adapting NIC codecs

### 2.1. NIC framework

The most common framework consists of a learnable feature autoencoder<sup>1</sup>, combined with quantization and learnable entropy coding [3, 12, 4]. The encoder  $\mathbf{b} = f(\mathbf{x}; \theta, \nu)$  maps the input image  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^N$  to the bitstream  $\mathbf{b}$ . The decoder reconstructs the image as  $\hat{\mathbf{x}} = g(\mathbf{b}; \phi, \nu)$ . The full model is determined by its parameters  $\psi = (\theta, \phi, \nu)$  (i.e. encoder, decoder, entropy model).

<sup>1</sup>Our autoencoder combines convolutional and (I)GDN [2] layers.

The parameters  $\psi$  are learned by minimizing<sup>2</sup> a combination of rate and distortion over a training set  $\mathcal{X}^{\text{tr}}$  sampled from the domain of interest  $\mathcal{X}$

$$J(\mathcal{X}^{\text{tr}}, \psi; \lambda) = R(\mathcal{X}^{\text{tr}}, \psi) + \lambda D(\mathcal{X}^{\text{tr}}, \psi), \quad (1)$$

where  $\lambda$  is the (fixed) tradeoff between rate<sup>3</sup>  $R(\mathcal{X}^{\text{tr}}, \psi)$  and distortion<sup>4</sup>  $D(\mathcal{X}^{\text{tr}}, \psi)$ . See Appendix A for more details.

### 2.2. Adapting to new domains

We introduce the problem of *domain adaptation in neural image compression* (DANICE), where a codec trained on a source domain  $\mathcal{X}_1$  is leveraged to improve compression in a target domain  $\mathcal{X}_2$ .

A straightforward approach to DANICE is *naive fine tuning*, where we minimize  $J(\mathcal{X}_2^{\text{tr}}, \psi; \lambda)$  with a target training set  $\mathcal{X}_2^{\text{tr}}$  in order to obtain a target model  $\psi_2$ . Naive fine tuning can easily lead to overfitting when the number of images is small relative to the total number of parameters.

We can mitigate overfitting and improve training stability by reducing the number of tunable parameters. Concretely, we tune the GDN layers [2] and the entropy model, while keeping the convolutional layers fixed. We introduce a small number of channel-specific parameters to adjust scales and biases ( $\tilde{z}_i = \alpha_i z_i + \beta_i$ , for the  $i$ -th channel of a feature  $\mathbf{z}$ ). Through this *selective fine tuning* approach, parameters that are tuned are reduced to only  $\sim 5\%$  of the original amount, improving adaptation under few images.

## 3. Forgetting and compatibility

Now, consider the case where we are interested in adapting a source codec (optimized for  $\mathcal{X}_1$ ) to  $\mathcal{X}_2$ , whilst retaining good compression performance on images from domain  $\mathcal{X}_1$ . Moreover, we may also want to decode any previously encoded  $\mathcal{X}_1$  images, i.e. have backward compatibility.

### 3.1. Characterizing forgetting

This new setting is more challenging than adaptation, because optimizing for the new domain *interferes* [9] with the solution for the source domain, leading to *forgetting*, i.e. the performance of the model in the source domain drops.

We can distinguish between two cases, leading to two different types of forgetting (described below). For simplicity, we use  $f_t$  to refer to the encoder at time  $t$ , i.e.  $f(\mathbf{x}; \theta_t, \nu_t)$ . Similarly, we use  $g_t$  to refer to  $g(\mathbf{b}; \phi_t, \nu_t)$ .

**Rate-distortion forgetting** Consider that a source image  $\mathbf{x} \in \mathcal{X}_1$  is encoded with the encoder  $f_2$  and decoder  $g_2$  optimized for the target domain  $\mathcal{X}_2$ , i.e.  $\hat{\mathbf{x}} = g_2(f_2(\mathbf{x}))$ . In this case, encoder and decoder belong to  $t = 2$ , so they

<sup>2</sup>During training, quantization is replaced by a differentiable proxy to allow end-to-end training via backpropagation [3].

<sup>3</sup>Approximated by the entropy during training.

<sup>4</sup>Here measured as the average mean square error (MSE).

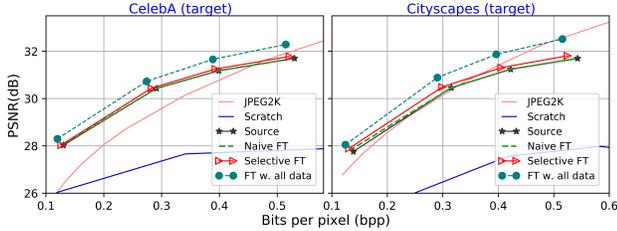


Figure 2. Domain adaptation with 25 images of target domain (CLIC → CelebA/Cityscapes).

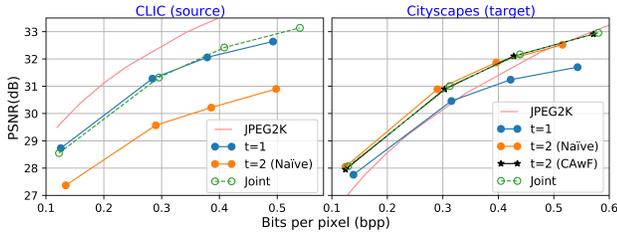


Figure 3. Changes in rate-distortion curves from  $t = 1$  to  $t = 2$  in both domains for the different experiments (left: source, right: target). Note: on the left figure  $t = 2$  (CAwF) corresponds to  $t = 1$ .

remain compatible. However, as the codec is optimized for a different domain, forgetting occurs, resulting in new artifacts and a worse rate-distortion performance compared to using  $f_1$  and  $g_1$  (see Fig. 1a).

**Catastrophic forgetting** In this case, forgetting is closely related to backward compatibility. When an image  $\mathbf{x} \in \mathcal{X}_1$  is encoded with  $f_1$  and saved, and then decoded with an updated version of the decoder  $g_2$ , i.e.  $\hat{\mathbf{x}} = g_2(f_1(\mathbf{x}))$ , the image cannot be recovered at all (see Fig. 1a).

### 3.2. Preventing forgetting and incompatibility

In order to address the aforementioned problems, we propose *Codec Adaptation without Forgetting* (CAwF), an image coding architecture composed of *shared* and *custom* parameters. The source codec, optimized for the source domain, is stored in the shared parameters,  $\psi_1$ . The custom parameters,  $\Delta\psi$  are used together with shared ones to obtain the target codec, i.e.  $\psi_2 = \psi_1 \cup \Delta\psi$ . The adaptation process also minimizes  $J(\mathcal{X}_2^u, \psi; \lambda)$ , but the shared parameters remain fixed. In our case, the custom parameters are distributed across the different layers of the feature auto-encoder, and a separate entropy model is learned. This architecture prevents both RD forgetting and interference between domains by design. The user can choose the version of the encoder, which is always signalled in the bitstream.

## 4. Experiments

### 4.1. Settings

We evaluate DANICE with few target images (section 4.2) and when forgetting is a concern (section 4.3). We focus our attention on scenarios involving lightweight mod-

Table 1. Average difference in bitrate due to adaptation to CelebA/Cityscapes with different number of images from target domain (measured as BD-rate (%) with respect to model obtained by fine tuning source model using all target data).

| Source model            | CLIC → CelebA     |                       | CLIC → Cityscapes |                       |
|-------------------------|-------------------|-----------------------|-------------------|-----------------------|
|                         | 19.24             |                       | 23.93             |                       |
| Number of target images | Naive fine tuning | Selective fine tuning | Naive fine tuning | Selective fine tuning |
| 10                      | 19.24             | <b>16.46</b>          | 22.96             | <b>17.54</b>          |
| 25                      | 18.76             | <b>14.93</b>          | 18.44             | <b>15.79</b>          |
| 50                      | 15.59             | <b>13.73</b>          | 16.29             | <b>15.33</b>          |
| JPEG2K                  | 46.96             |                       | 24.55             |                       |

Table 2. Average differences in bitrate of various DANICE approaches when forgetting is a concern (measured as BD-rate (%) with respect to the model trained jointly using source and target data). Lower the value, better the performance.

| Time      | CelebA → Cityscapes |              |              | Cityscapes → CelebA |              |              |
|-----------|---------------------|--------------|--------------|---------------------|--------------|--------------|
|           | CelebA              | Cityscapes   | Avg.         | Cityscapes          | CelebA       | Avg.         |
| 1         | <b>-1.14</b>        | 32.77        | 15.82        | <b>-4.75</b>        | 84.32        | 39.79        |
| 2 (naive) | 96.25               | <b>-4.79</b> | 45.73        | 26.25               | <b>-2.45</b> | 11.90        |
| 2 (CAwF)  | <b>-1.14</b>        | -0.40        | <b>-0.77</b> | <b>-4.75</b>        | 4.01         | <b>-0.37</b> |
| JPEG2K    | 38.26               | 16.34        | 27.3         | 16.34               | 38.26        | 27.3         |

| Time      | CLIC → CelebA |              |             | CLIC → Cityscapes |              |              |
|-----------|---------------|--------------|-------------|-------------------|--------------|--------------|
|           | CLIC          | CelebA       | Avg.        | CLIC              | Cityscapes   | Avg.         |
| 1         | <b>5.97</b>   | 15.59        | 10.78       | <b>-2.13</b>      | 18.90        | 8.39         |
| 2 (naive) | 10.59         | <b>-2.05</b> | 4.27        | 74.28             | <b>-3.18</b> | 35.55        |
| 2 (CAwF)  | <b>5.97</b>   | 1.14         | <b>3.56</b> | <b>-2.13</b>      | 0.27         | -0.93        |
| JPEG2K    | -23.57        | 40.66        | 8.54        | -28.37            | 18.38        | <b>-4.99</b> |

els and narrow target domains.

**Domains** We consider three datasets CLIC, Cityscapes (resized to  $512 \times 256$ ) and CelebA-HQ (resized to  $256 \times 256$ ) for our experiments. CLIC is a wide domain (photos of a variety of topics), while the domains of CelebA and Cityscapes are narrower (faces and street scenes taken from a car, respectively). We experimented with the following pairs of source-target domains: CelebA-Cityscapes, Cityscapes-CelebA, CLIC-CelebA and CLIC-Cityscapes.

**Implementation** Our framework is based on the *factorized-prior* model from [4]<sup>5</sup>, where the encoder has four convolutional layers separated by GDN layers, and the decoder mirrors them. The model is trained for 500k-750k iterations at four RD tradeoffs,  $\Lambda = \{\lambda_1 = 0.002, \lambda_2 = 0.008, \lambda_3 = 0.016, \lambda_4 = 0.032\}$ . Our framework has 64 shared filters and 16 custom filters per conv. layer. We include JPEG2K<sup>6</sup> RD curves as reference.

**Metrics** In addition to the usual average rate (in bits-per-pixel, bpp) and average distortion (in average RGB PSNR) curves, we also use the BD-rate metric for comparisons.

### 4.2. Adaptation with limited data

We evaluate the benefits of selective fine tuning for DANICE, by comparing it against training from scratch (randomly initialized parameters), using the source model as is, and naive fine tuning. Fig. 2 shows that a source model

<sup>5</sup><https://github.com/tensorflow/compression>

<sup>6</sup><https://www.openjpeg.org>

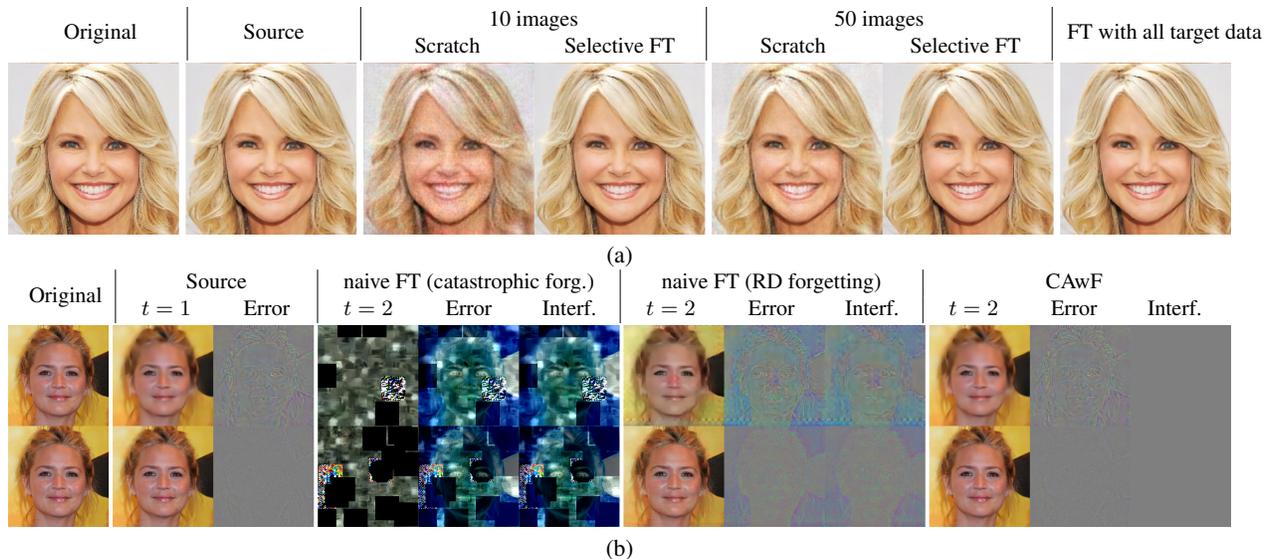


Figure 4. Examples showing that the source model (CLIC) is transferable to CelebA (a), and the benefit of selective fine tuning. Adaptation from CelebA to Cityscapes (b) shows learning in Cityscapes and forgetting in CelebA. First and second rows in each block row show results for low and high rate codecs, respectively. Error shows the difference image with respect to Original, while Interf. and Change show the difference with respect to image at  $t = 1$ . Best viewed in electronic version with zoom.

trained on CLIC achieves competitive RD performance, and performs better than training from scratch using few images. This demonstrates that NIC codecs are transferable and improve RD performance through adaptation with few images (especially using selective fine tuning), which is not possible with a traditional codec such as JPEG2K. Table 1 shows this gain reporting the BD rate of various approaches using the codec fine tuned using all target data ( $\geq 20000$  images) as reference. Fig. 4 (a) compares reconstructed images using different models, showing that not using a source model (i.e. from scratch) leads to poor quality and artifacts such as noise when the training data is limited.

### 4.3. Forgetting and compatibility

**Rate-distortion performance** Fig. 3 shows the RD curves on source and target domains. We can observe that after naive fine tuning (i.e.  $t = 2$ ), the RD performance on the target domain improves (mostly reducing distortion, and slightly reducing rate) while in the source domain often deteriorates significantly, which showcases forgetting. With CAwF, we are able to avoid forgetting in the source domain by design, while being able to attain good performance in the target domain due to the addition of custom parameters. The amount of improvement and decline after adaptation depends on the specific pair of domains, being small for CLIC-CelebA, and large for Cityscapes-CelebA. Table 2 reports the BD-rate savings of various approaches with respect to a codec with 80 (64+16) filters trained jointly with both domains, since it often serves as a (soft) upper bound for sequential training. With CAwF, the average of BD-

rate of source and target domains is close to zero and much lower than naive fine tuning. However, we note that the addition of 16 filters in CAwF can increase the number of parameters up to 55% (see Table 3 in the appendix).

**Qualitative analysis** Fig. 4 (b) shows several decoded images from the CelebA–Cityscapes experiments. While in all cases images are decodable, we observe that the CelebA bitstreams from  $t = 1$  turn into a collage of random patches, and encoding at  $t = 2$  results in a significant loss of RD performance, with noticeable artifacts caused by the interference. In contrast, CAwF is able to prevent interference and keep the original quality.

## 5. Conclusion

In this paper, we propose a framework to optimize a NIC codec to a target domain of interest, which can be performed directly by the user with a custom dataset. In order to avoid interference that could break the compatibility or harm the coding capabilities in the source domain, target domain knowledge is stored as separate custom parameters. Our work draws from insights from several machine learning areas, chiefly transfer learning, domain adaptation and continual learning, which could provide inspiration for new functionalities in future image and video coding standards.

## Acknowledgments

We acknowledge the support from Huawei Kirin Solution and the Spanish Government funding for projects RTI2018-102285-A-I00 and RYC2019-027020-I.

## References

- [1] H. Azizpour, A. S. Razavian, J. Sullivan, A. Maki, and S. Carlsson. Factors of transferability for a generic convnet representation. *PAMI*, 38, 2015.
- [2] J. Ballé, V. Laparra, and E.P. Simoncelli. Density modeling of images using a generalized normalization transformation. *ICLR*, 2016.
- [3] J. Ballé, V. Laparra, and E. P. Simoncelli. End-to-end optimized image compression. *ICLR*, 2016.
- [4] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston. Variational image compression with a scale hyperprior. *ICLR*, 2018.
- [5] B. Bross, J. Chen, S. Liu, and Y.-K. Wang. Versatile video coding (draft 10). Technical report, Joint Video Experts Team (JVET), 2020.
- [6] Chunlei Cai, Li Chen, Xiaoyun Zhang, and Zhiyong Gao. Efficient variable rate image compression with multi-scale decomposition network. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(12):3687–3700, 2018.
- [7] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3146–3154, 2019.
- [8] Nick Johnston, Elad Eban, Ariel Gordon, and Johannes Ballé. Computationally efficient neural image compression. *arXiv preprint arXiv:1912.08771*, 2019.
- [9] M. McCloskey and N. J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [10] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [11] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand. Overview of the high efficiency video coding (hevc) standard. *TCSVT*, 22(12):1649–1668, 2012.
- [12] L. Theis, W. Shi, A. Cunningham, and F. Huszár. Lossy image compression with compressive autoencoders. *arXiv*, 2017.
- [13] G. Toderici, S. M. O’Malley, S. J. Hwang, D. Vincent, D. Minnen, S. Baluja, M. Covell, and R. Sukthankar. Variable rate image compression with recurrent neural networks. *ICLR*, 2016.
- [14] G. Toderici, D. Vincent, N. Johnston, S. J. Hwang, D. Minnen, J. Shor, and M. Covell. Full resolution image compression with recurrent neural networks. In *CVPR*, pages 5306–5314, 2017.
- [15] G.K. Wallace. The JPEG still picture compression standard. *TCE*, 38(1):18–34, 1992.
- [16] M. Wang and W. Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- [17] T. Wiegand, G. J. Sullivan, G. Bjontegaard, and A. Luthra. Overview of the h. 264/avc video coding standard. *TCSVT*, 13(7):560–576, 2003.
- [18] F. Yang, L. Herranz, Y. Cheng, and M.G. Mozerov. Slimmable compressive autoencoders for practical neural image compression. In *CVPR*, 2021.
- [19] F. Yang, L. Herranz, J. van de Weijer, J.A. Iglesias-Guitián, A.M. López, and M.G. Mozerov. Variable rate deep image compression with modulated autoencoder. *IEEE Signal Processing Letters*, 27:331–335, 2020.