

# A Taxonomy of Miscompressions: Preparing Image Forensics for Neural Compression

Nora Hofer  
University of Innsbruck, Austria  
nora.hofer@uibk.ac.at

Rainer Böhme  
University of Innsbruck, Austria  
rainer.boehme@uibk.ac.at

**Abstract**—Neural compression has the potential to revolutionize lossy image compression. Based on generative models, recent schemes achieve unprecedented compression rates at high perceptual quality, but they compromise semantic fidelity. Details of decompressed images may appear optically flawless, but semantically different from the originals, making compression errors difficult or impossible to detect. We explore the problem space and propose a provisional taxonomy of miscompressions. It defines three types of “what happens” and has a binary “high impact” flag indicating miscompressions that alter symbols. We discuss how the taxonomy can facilitate risk communication and research into mitigations.

**Index Terms**—neural image compression, miscompression, semantic changes, forensics

## I. INTRODUCTION

A turning point in the investigation of the 2013 Boston Marathon bombing was a bystander’s cellphone photo that allowed police to identify one of the suspects in a crowd [1] [2]. Remarkably, the relevant part of the image comprised just 0.2% of all pixels. In this paper, we ask the question whether digital images will continue to serve as reliable sources in a future where neural compression becomes the default.

Neural image compression employs learning-based elements in the image compression pipeline, achieving high perceptual quality at unprecedented compression rates [3]. State of the art schemes use generative networks to synthesize parts of an image [4], [5]. However, a drawback of this approach is that the synthesized details appear plausible and of high perceptual quality, but may be semantically different from the original.

To illustrate this, Fig. 1 compares a small crop (0.4%) of an uncompressed image (left), to a version from the *HiFiC* neural compression scheme [4] (middle), and the JPEG compressed image at quality factor (QF) 20 (right). This factor was chosen to match the compression rate in bits per pixel between *HiFiC* and JPEG. While the neural compression retains clear readability of the numbers, closer inspection reveals that they differ: the upper row changes from 2264–7668 in the original to 2254–7664 in the *HiFiC* reconstruction. We confirmed that not only human observers, but also Google’s Cloud AI optical character recognition came to this conclusion. Given the high apparent quality of the image, observers unaware of its processing history might be inclined to fully trust the image and its semantic content. By contrast, the visible compression

We thank Benedikt Lorch, Martin Beneš, Judith Senn, and Verena Lachner for valuable discussions and comments. Computational results were achieved using the LEO HPC infrastructure at the University of Innsbruck.



(a) Lossless original (4.1 MB, 11.90 bpp) (b) Neural compression (159 kB, 0.46 bpp, [4]) (c) JPEG QF 20 (158 kB, 0.46 bpp)

Fig. 1. State-of-the-art neural compression schemes can alter the semantic in details of the decompressed images. The high fidelity and the lack of visible compression artifacts make false reconstructions look more authentic than JPEG, which introduces visible distortion. (Crop of image 0831 of DIV2K [6], 0.41% of the original.) All figures are best viewed on screen and magnified.

artifacts in the JPEG image not only make the numbers difficult to read, but also signal low reliability and dissuade users from interpreting the numbers with confidence.

In this paper, we propose the term “miscompression” to describe semantic changes resulting from lossy compression. This new<sup>1</sup> phenomenon arose with neural compression and deserves the attention of researchers and forensic practitioners. To facilitate the conversation, we develop a taxonomy of miscompressions based on the explorative visual inspection of three benchmark datasets, examining five different neural compression schemes at different quality settings. The examples used for illustration in this paper are produced using two schemes of the latest generation. We derive implications for forensics and society at large, and outline objectives for research into mitigations.

We emphasize that all authors of recent neural compression schemes generally acknowledge the risk of hallucinations, warn against using their schemes for critical applications [4, p. 10], and occasionally point to examples of text becoming unreadable. However, miscompressions take many forms and pose a risk to society if left unaddressed. We see our work as a first step towards solutions.

The paper is structured as follows. Section II introduces the principles of neural compression by comparing it to the JPEG pipeline. Section III defines miscompressions, describes our empirical approach, and presents the proposed taxonomy using examples. Section IV discusses implications and outlines a research agenda. Section V concludes our paper.

<sup>1</sup>The closest related work we are aware of is an attempt to introduce copy-evident marks into images which only appear after JPEG compression [9].

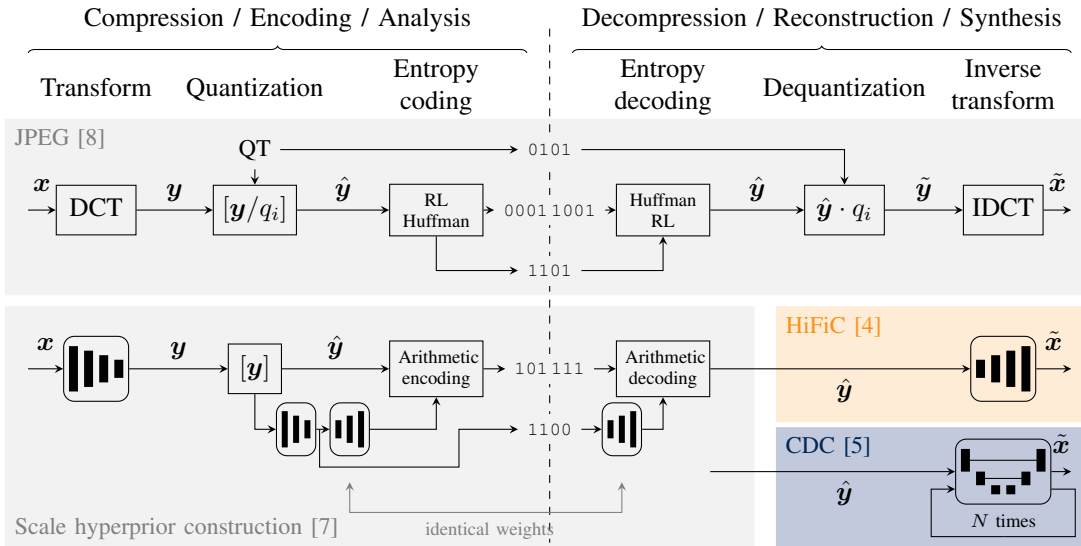


Fig. 2. Comparison of the conventional JPEG compression (top) and the neural compression schemes used in this paper (bottom). Elements with rounded corners are CNNs trained on datasets and used in inference mode during encoding and reconstruction.

## II. PRIMER ON NEURAL COMPRESSION

The top row of Fig. 2 shows the key components of the lossy image compression and decompression pipeline. The first component **transforms** an input image from the spatial domain into a domain where pixels are decorrelated and the variance is concentrated in fewer coefficients. A useful property of common transformations is that the distribution of the coefficients can be approximated parametrically. The second component, **quantization**, is a deliberately lossy process. It maps a range of values to a single discrete value. The quantization steps can be adjusted according to the relevance of each coefficient for the reconstruction of the signal. The final component of the compression pipeline is lossless **entropy coding**, which approximates Shannon’s theorem [10] by encoding a sequence of quantized values into a sequence of bits. The better the input distribution is known, the tighter the approximation and the shorter the resulting bitstream.

Over decades, JPEG [8] has been the most popular lossy image compression scheme. It implements the transform by the linear discrete cosine transform (DCT) [11] on non-overlapping blocks of  $8 \times 8$  pixels. The resulting coefficients are then quantized by dividing them by frequency-specific quantization factors and subsequent rounding to the nearest integer. The quantized coefficients are arranged in zigzag order and entropy-coded using run-length (RL) and Huffman encoding [12]. The resulting image file contains the quantization tables (QT), the quantized DCT coefficients, and the Huffman tables. JPEG decompression reverts this sequence of steps.

Neural compression replaces components of this pipeline with learnable elements, typically deep convolutional neural networks (CNN). This emerging field has its own jargon. Encoding and analysis are used to describe compression. Reconstruction and synthesis denote decompression (cf. Fig. 2).

Learning the transform promises that irrelevance in the input

signal can be isolated better in the so-called latent space than with known structured transformations, such as block-wise DCT. While the networks have shown to derive basis functions similar to those in linear transforms [13], nonlinear transforms offer better adaptation to varying data distributions and can be optimized for specific distortion metrics. The loss function used for training has two terms: distortion and rate. By weighting these terms, different tradeoffs between image quality and file size can be achieved. Finding the right distortion metric for neural compression is an active field of research [14], [15]. Once trained, the weights are stored in the encoder and decoder. The CNNs are used in inference mode for the encoding and reconstruction of images.

Quantization in neural compression typically involves rounding and truncation [16], [17]. Unlike JPEG, it does not use and transmit QTs. The quantization step size is controlled by the scaling in last layer of the transform CNN and thus learned. Therefore, neural compression schemes commonly require a separately trained model for each target quality.

Also, entropy coding requires modifications. A drawback of learning a transform is the lack of a statistical model of the latent space. Here, the answer to machine learning is machine learning. The distribution of the latent space is modelled with a trained auto encoder. The prediction of this model is used to parameterize an arithmetic encoder. As the distribution is data dependent, the latent space of this prediction model must itself be transmitted to the decoder to enable reconstruction. Ballé et al.’s scale hyperprior construction [7] is the basis of the two schemes evaluated in this work.

*HiFiC* [4] and *CDC* [5] improve on previous approaches to neural compression by using generative models for the inverse transform. *HiFiC* uses a **generative adversarial network** (GAN). GANs are trained with a rivaling discriminator network that regularizes the generator network towards

producing outputs of high perceptual quality [18]. To control the content of the reconstruction, the generator is conditioned with the latent representation of the encoded image. *CDC* implements a **diffusion model** [19] for the inverse transform. It uses the latent representation to condition the denoising diffusion probabilistic model [20] (DDPM). Both schemes are trained end-to-end, allowing the variational autoencoder in the transform component to learn how to turn an input image to a condition. The rate estimate of the loss function is taken from the hyperprior model and the distortion estimate is a weighted sum of the perceptual loss and the mean squared error between the input and the reconstructed image. Increasing the weight of the perceptual metric gives the model the flexibility to deviate from the input signal and “make up” details during reconstruction. This enables high perceptual quality at unprecedented compression rates, but compromises semantic fidelity.

The digital image forensics community has only recently started to address the impact of neural compression: Berthet et al. revisit copy-move forgery detection [21] and source social network identification [22], Bergmann et al. use traces in the frequency and spatial domain for detection [23], [24], and Chen et al. show a vulnerability of different neural compression schemes to adversarial perturbations in the input image [25]. Jalilian et al. propose CNNs to compress biometric images [26]. However, to the best of our knowledge, nobody has yet investigated semantic changes and their implications.

### III. MISCOMPRESSIONS

Semantic interpretation is the understanding of a perceived scene by applying domain terminology, *i.e.*, semantic concepts [27]. It is carried out by a human observer and is heavily influenced by their prior semantic and conceptual knowledge of the domain [28]. The *semantic meaning* of a scene or an image is the result of semantic interpretation.

We define **miscompressions** as reconstruction errors that occur when there is a discrepancy between the semantic meaning of an original image (detail) and its reconstructed version after neural compression. As a test, we require that a human observer asked to verbally describe the relevant part of the image would come up with a different description. Note that this definition applies to entire images as well as individual image details. Digital images used as evidence in forensic investigations often capture relevant details unconsciously, as can be seen from the Boston Marathon bombing example. Such photographs often preserve details that serve as objective representations of reality. Consequently, forensic investigators typically focus on analyzing the semantic meaning of specific image details, such as objects or individuals in the background, rather than interpreting the semantics of the entire image.

Miscompressions are a new phenomenon, requiring a precise terminology to describe, mitigate, or ultimately avoid them. This paper takes a first step in this direction and proposes a provisional taxonomy, systematically derived with three key objectives in mind. First, the taxonomy should facilitate research into the risks posed by miscompressions. Distin-

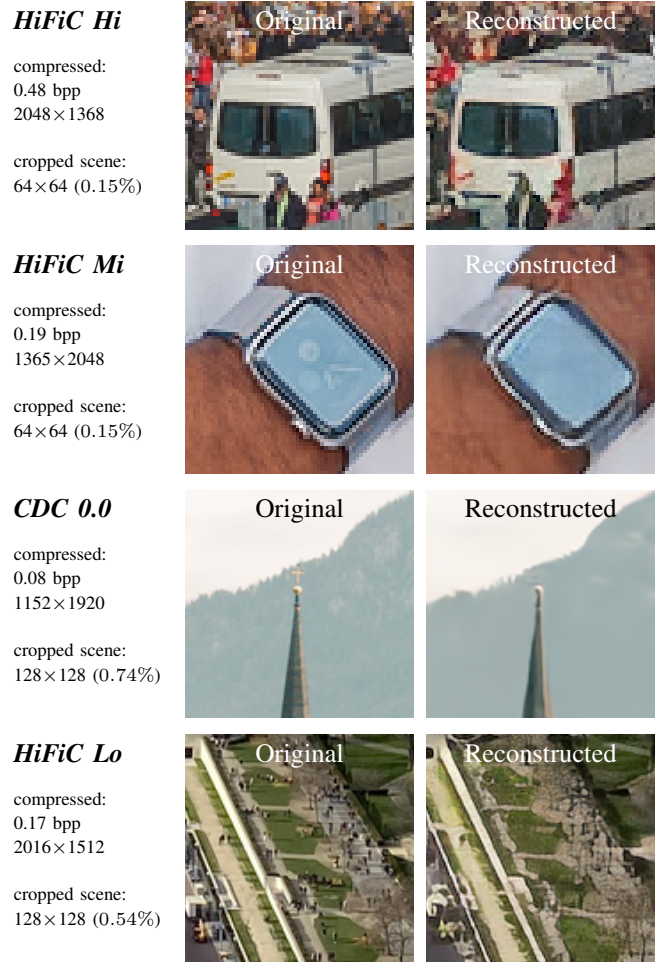


Fig. 3. Category AMPLITUDE: Reconstructions differ in the amplitude of spatial frequencies in the signal, affecting attributes such as brightness, color saturation, and the intensity of high frequency components.

guishing between different forms of reconstruction errors that lead to miscompressions allows us to measure their prevalence, and compare compression schemes and the optimization metrics used. Second, it should facilitate research towards making neural compression safer. While ideal compression schemes would be completely immune to miscompressions, it is uncertain whether this is achievable at competitive compression rates. However, confidence in neural compression would be greatly improved if it could be ensured that certain types of miscompressions are extremely unlikely. The third intended application of our taxonomy is to deal with the remaining risk in practice. It should allow forensics experts to explain miscompressions using references to scientific evidence in order to convince a judge or jury.

#### A. Method

Our approach is exploratory. We focus on five relevant neural compression schemes [4], [5], [7], [17], [29], compress the test images of three widely-used benchmark datasets, *CLIC2020* [30], *DIV2K* [6], and *Kodak* [31] (full dataset), and manually inspect the reconstructions of a total of 552

images to identify miscompressions, using difference images for guidance, where necessary.

We have observed miscompressions in all tested schemes, but decided to shift our focus to *HiFiC* [4] and *CDC* [5]. These schemes employ generative networks and stand out, as their reconstructions are of such high fidelity that they appear deceptively authentic. We use the pre-trained *HiFiC* model with 180 million parameters for the three available compression intensities *high*, *mid*, and *low*,<sup>2</sup> and the pre-trained  $\mathcal{X}$ -parameterization model of *CDC* with 54 million parameters for the widely-adopted LPIPS loss at weights 0.0 and 0.9.<sup>3</sup> We varied the noise seeds for *CDC* and found that miscompressions prevailed.

### B. Taxonomy

At a high level, our taxonomy separates the signal processing perspective (“What happens?”) from the semantic impact (“How bad is the misinterpretation?”). Based on the apparent transformation of the signal, we define three categories. To illustrate each category we provide examples, cropped from compressed images of three datasets, and specify the compression model used, the bpp of the compressed image, the pixel dimensions of the original image, and the crop, as well as the percentage of the original represented by the displayed crop. The selected examples illustrate the characteristics of each type of miscompression. In practice, many miscompressions exhibit combined effects of multiple types.

a) **AMPLITUDE**: refers to reconstructions that differ in the amplitude of spatial frequencies in the image signal, such as changes in brightness, color saturation, or intensity of high frequency components. Attenuation seems to be more common, although we cannot rule out amplification. Unlike global signal processing operations, these effects tend to be local and content-dependent. Objects that we have found to be particularly susceptible to this type of miscompression include lights, colors of eye, hair, and skin, as well as birthmarks, and tattoos. Attenuation can result in altered colors or “disappearing” objects. Semantic changes occur when the amplitude carries meaning, as illustrated in the examples in Fig. 3. In the top row, the fact that the car was braking, as indicated by the brake lights, is lost in the reconstruction. In the second row, a reconstructed watch appears to be turned off but is actually on and displaying the time. In the third row, the reconstructed image of a church tower does no longer include the Christian cross. In the bottom row, the reconstructed image of a park does not include the people present in the original image.

b) **GEOMETRY**: refers to reconstructions with geometric transformations such as translations, rotation, scaling, and shearing. This includes locally shifted shapes, dissolved contours, and imperfect representations of 3D scenes in 2D pixel matrices. Semantic changes occur when the geometry of an object carries semantic meaning, as illustrated in Fig. 4. The top row shows a reconstruction of a flag that could be mistaken

<sup>2</sup>*HiFiC*: <https://github.com/tensorflow/compression/tree/master/models/hific>

<sup>3</sup>*CDC*: [https://github.com/buggyyang/CDC\\_compression](https://github.com/buggyyang/CDC_compression)



Fig. 4. Category **GEOMETRY**: Reconstructions contain geometric transformations, such as translation, rotation, scaling, and shearing, including shifted shapes and dissolved contours. (The top image illustrates the level of detail at which miscompressions occur. A grid was added to the middle crop.)

for graffiti on the wall in the background. The direction of the nose and chin of the person’s profile in the second row is altered and differs from the original. The bottom row shows a leaf in the foreground that looks like a crack in the floor after reconstruction. Other susceptible objects, not illustrated here, include shadows and reflections. Semantic changes occur if the direction of a shadow changes and suggests a different positioning of the object casting the shadow, or if an object that is placed in front of a reflecting body of water or glass appears to be part of the reflection. This has implications for forensic methods that exploit inconsistencies in lighting direction and shadows for the detection of image manipulations [32, p. 14].

c) **SHAPE**: refers to reconstructions that differ in shape, potentially caused by biases in the retrieval augmentation process. Semantic changes occur when the shape of an object conveys a semantic meaning, as illustrated in Fig. 5. The top row shows a cropped image of a camera lens. The change in shape causes the number 8 in the original image to be mistaken as the number 6 in the reconstruction. In the bottom row, the shape of direction-specific traffic lights changes in the



Fig. 5. Category SHAPE: Reconstructions contain changed contours.

reconstruction from arrows to round lights. This results in a change of semantic meaning, as a round green light typically indicates that drivers can proceed in any direction.

Although we have observed several instances of altered textures, we do not consider them as a distinct type for now because the alterations did not align with our definition of miscompressions, *i.e.*, changes in semantic meaning. However, texture changes have been reported in the super-resolution literature [33, p. 25789], and we retain it as a potential extension to our taxonomy.

To classify the potential semantic impact of miscompressions across all categories, we define the SYMBOL modifier. The consequence of miscompressions is elevated when the affected objects portray *symbols*, *i.e.*, signs that carry specific meaning to human observers within a given social and cultural context [34]. Examples of symbols include letters, numerals, and signs, as well gestures, body adornments (*e.g.*, religious jewelry or clothing, wedding rings, etc.), traffic signs and lights, watch hands, logos, tattoos, graffiti, etc. This list is not exhaustive, and identifying a SYMBOL is subjective and can be challenging, especially without knowledge of the cultural and societal context of the captured scene. When symbols are involved, small changes in amplitude, geometry, or shape can completely alter the semantic meaning. For instance, the miscompression of a plant in the bottom row of Fig. 4 is likely harmless, whereas missing jewelry, as shown in Fig. 6 could lead to discord.

While screening our data, we have made a number of noteworthy qualitative observations. First, we found that miscompressions commonly occur in small image details (see top row in Fig. 4), and a single image can contain multiple instances of miscompressions. Notably, not every image contains miscompressions. Images that depict single large objects against flat backgrounds are less susceptible. Moreover, we find that CDC is more likely to visibly destroy text, which reduces the risk of misinterpretation of incorrectly reconstructed text. In general, miscompressions occur seemingly unpredictably, and are difficult to distinguish from authentic image details.



Fig. 6. Miscompression of *symbols* such as body adornments or religious jewelry increase the risk of semantic misinterpretation.

#### IV. DISCUSSION

In this section, we briefly reflect on the decisions that have shaped our taxonomy, then outline how it can be applied, before discussing the wider implications of miscompressions and closing with selected avenues for future research.

The definition of miscompressions based on textual descriptions is naturally subjective. It depends on the observer with their experience and on the language, which defines concepts based on culture. Intersubjectivity can be improved by asking multiple observers [35]. The language dependency aligns the definition with relevance in the given cultural context.

The proposed taxonomy of miscompressions is a first qualitative step towards mitigations. The next step is to apply the taxonomy for the annotation of a larger dataset. This will pave the way towards quantifying the prevalence of miscompressions and identifying influencing factors with statistical methods. Importantly, such a dataset could be used to train models that detect and classify miscompressions, removing the human from the loop and allowing even greater scale.

Also image-to-text models can be useful tools to this end. To explore this option, we asked ChatGPT 4.0 to describe the original as well as the reconstructed image of the church tower in Fig. 3. We presented it with a  $256 \times 256$  crop (bigger than in Fig. 3) that included the lower part of the tower and roofs of houses. ChatGPT’s description changed from “*The image depicts a church steeple with a cross at the top, situated in a mountainous area. [...]*” for the original image to “*The image appears to be of a church steeple or a spire set against a mountainous backdrop. [...]*” But this approach did not work for all of our examples. For instance, there was no difference in the description of the watch in Fig. 3. Using targeted object recognition methods to identify specific cases of miscompressions appears more promising, *e.g.*, the use of optical character recognition to identify miscompressed text.

Automatic detectors can be implemented as a safety net at encoding time to catch potential miscompressions. This would allow an increase in the number of bits allocated to areas of the image where miscompressions loom. The next step would be to use annotations of miscompressions as part of the training loss metric in order to harden future neural compression schemes against miscompressions. Our taxonomy can be used to tailor these metrics to the types of miscompressions that should be avoided in a particular application, for example humans in the surveillance of public places and license plates in traffic surveillance. Conversely, neural compression could

be tuned to *deliberately cause* miscompressions of, say, human faces as an integrated privacy-enhancing technology [36]–[38].

While research should strive to avoid miscompressions entirely, in the meantime it is crucial to deal with the existing risk in practice. It is imperative to acknowledge the existence of miscompressions and explain the associated risks of misunderstandings and false accusations to end users of the technology. Neural compression is not ready for use in safety and security critical applications, such as public surveillance or autonomous driving. The benefit of bandwidth savings is disproportionate to the risk of wrongful convictions and potentially fatal accidents. Worryingly, surveillance and autonomous driving are mentioned prominently in the motivation for the upcoming JPEG AI neural compression standard [39, p. 104]. In less critical applications, the use of neural compression should be documented. Suitable annotation could be stored in image metadata, where professionals, such as photo journalists and forensic investigators, can find and interpret them. A quantitative perceptual metric of miscompressions, similar to the metric for photo retouching [40], could be used in image captions, visible watermarks, or icons, and inform consumers about the potential presence of miscompressions. Reliable methods to detect neural compression are needed to enforce such policies [23], [24]. Interesting open research problems remain: Is it possible to detect instances of miscompressions in reconstructed images without access to the original? Can we develop forensic methods to distinguish uncontrolled miscompressions from malicious manipulations?

## V. CONCLUSION

To our knowledge, this is the first study that compares multiple neural compression schemes for their susceptibility to produce semantically different reconstructions. We raise awareness of this novel problem in forensics, propose a provisional taxonomy of what we call *miscompressions*, and support it with existential evidence. We hope that as this taxonomy develops, it will enable quantitative studies of automated detection, prevalence, influencing factors, and mitigations.

## REFERENCES

- [1] “Boston Bombing Day 2,” Apr 2016, <https://abcnews.go.com/US/boston-bombing-day-improbable-story-authorities-found-bombers/story?id=38375726>, last accessed: July 2024.
- [2] “Dzhokhar A. Tsarnaev,” Mar 2024, [https://www.justice.gov/usao-ma/tsarnaev-exhibits-day-2-exh\\_29.pdf](https://www.justice.gov/usao-ma/tsarnaev-exhibits-day-2-exh_29.pdf), last accessed: July 2024.
- [3] Y. Yang, S. Mandt, and L. Theis, “An introduction to neural data compression,” *Found. Trends Comput. Graph. Vis.*, pp. 113–200, 2023.
- [4] F. Mentzer, G. Toderici, M. Tschannen, and E. Agustsson, “High-fidelity generative image compression,” *NeurIPS*, 2020.
- [5] R. Yang and S. Mandt, “Lossy image compression with conditional diffusion models,” *NeurIPS*, 2024.
- [6] E. Agustsson and R. Timofte, “NTIRE 2017 challenge on single image super-resolution: Dataset and study,” in *CVPR*, 2017.
- [7] J. Ballé, D. Minnen, S. Singh, S. Hwang, and N. Johnston, “Variational image compression with a scale hyperprior,” in *ICLR*, 2018.
- [8] G. Wallace, “The JPEG still picture compression standard,” *Communications of the ACM*, pp. 30–44, 1991.
- [9] A. Lewis and M. Kuhn, “Towards copy-evident JPEG images,” *GI Jahrestagung*, pp. 1582–1591, 2009.
- [10] C. Shannon, “A mathematical theory of communication,” *The Bell System Technical Journal*, pp. 379–423, 1948.
- [11] N. Ahmed, T. Natarajan, and K.R. Rao, “Discrete cosine transform,” *IEEE Transactions on Computers*, pp. 90–93, 1974.
- [12] D. Huffman, “A method for the construction of minimum-redundancy codes,” *IRE*, pp. 1098–1101, 1952.
- [13] J. Duan, M. Lu, Z. Ma, and F. Zhu, “Opening the black box of learned image coders,” in *PCS. IEEE*, 2022, pp. 73–77.
- [14] K. Ding, K. Ma, S. Wang, and E. Simoncelli, “Image quality assessment: Unifying structure and texture similarity,” *TPAMI*, pp. 2567–2581, 2020.
- [15] S. Bhardwaj, I. Fischer, J. Ballé, and T. Chinen, “An unsupervised information-theoretic perceptual quality metric,” in *NeurIPS*, 2020, pp. 13–24.
- [16] L. Theis, W. Shi, Q. Cunningham, and F. Huszár, “Lossy image compression with compressive autoencoders,” *ICLR*, 2017.
- [17] J. Ballé, V. Laparra, and E. Simoncelli, “End-to-end optimized image compression,” *arXiv preprint arXiv:1611.01704*, 2016.
- [18] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *NeurIPS*, 2014.
- [19] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *ICML. PMLR*, 2015, pp. 2256–2265.
- [20] J. Ho, A. Jain, and P. Abbeel, “Denosing diffusion probabilistic models,” *NeurIPS*, pp. 6840–6851, 2020.
- [21] A. Berthet and J. Dugelay, “AI-based compression: A new unintended counter attack on JPEG-related image forensic detectors?,” in *ICIP. IEEE*, 2022, pp. 3426–3430.
- [22] Alexandre Berthet, Chiara Galdi, and Jean-Luc Dugelay, “On the impact of AI-based compression on deep learning-based source social network identification,” in *MMSp. IEEE*, 2023, pp. 1–6.
- [23] S. Bergmann, D. Moussa, F. Brand, A. Kaup, and C. Riess, “Frequency-domain analysis of traces for the detection of AI-based compression,” in *IWBF. IEEE*, 2023, pp. 1–6.
- [24] S. Bergmann, D. Moussa, F. Brand, A. Kaup, and C. Riess, “Forensic analysis of AI-compression traces in spatial and frequency domain,” *Pattern Recognit. Lett.*, pp. 41–47, 2024.
- [25] T. Chen and Z. Ma, “Towards robust neural image compression: Adversarial attack and model finetuning,” *TCSVT*, 2023.
- [26] E. Jalilian, H. Hofbauer, and A. Uhl, “Iris image compression using deep convolutional neural networks,” *Sensors*, vol. 22, no. 7, 2022.
- [27] D. Marr, *Vision: A computational investigation into the human representation and processing of visual information*, MIT press, 2010.
- [28] C. Hudelot, N. Maillot, and M. Thonnat, “Symbol grounding for semantic image interpretation: From image data to semantics,” in *ICCV. IEEE*, 2005, pp. 1875–1875.
- [29] D. Minnen and S. Singh, “Channel-wise autoregressive entropy models for learned image compression,” in *ICIP. IEEE*, 2020, pp. 3339–3343.
- [30] G. Toderici, W. Shi, R. Timofte, L. Theis, J. Balle, E. Agustsson, N. Johnston, and F. Mentzer, “Workshop and challenge on learned image compression (clic2020),” in *CVPR*, 2020.
- [31] R. Franzen, “Kodak photocd dataset,” Nov 1999.
- [32] A. Piva, “An overview on image forensics,” *International Scholarly Research Notices*, p. 496701, 2013.
- [33] B. Li, X. Li, H. Zhu, Y. Jin, R. Feng, Z. Zhang, and Z. Chen, “SeD: Semantic-aware discriminator for image super-resolution,” in *IEEE/CVF CVPR*, 2024, pp. 25784–25795.
- [34] A. Frutiger, *Signs and Symbols*, Weiss Verlag, 1989.
- [35] A. Smaling, “Varieties of methodological intersubjectivity - the relations with qualitative and quantitative research, and with objectivity,” *Quality and Quantity*, pp. 169–180, 1992.
- [36] K. Lander, V. Bruce, and H. Hill, “Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces,” *JARMAC*, pp. 101–116, 2001.
- [37] P. Ilia, I. Polakis, E. Athanasopoulos, F. Maggi, and S. Ioannidis, “Face/off: Preventing privacy leakage from photos in social networks,” in *CCS*, 2015, pp. 781–792, ACM.
- [38] Y. Li, N. Vishwamitra, B. Knijnenburg, H. Hu, and K. Caine, “Effectiveness and users’ experience of obfuscation as a privacy-enhancing technology for sharing photos,” *PACM HCI*, pp. 1–24, 2017.
- [39] J. Ascenso, E. Alshina, and T. Ebrahimi, “The JPEG AI standard: providing efficient human and machine visual data consumption,” *IEEE MultiMedia*, pp. 100–111, 2023.
- [40] E. Kee and H. Farid, “A perceptual metric for photo retouching,” *PNAS*, pp. 19907–19912, 2011.