

TOWARDS CODING FOR HUMAN AND MACHINE VISION: A SCALABLE IMAGE CODING APPROACH

Yueyu Hu, Shuai Yang, Wenhan Yang, Ling-Yu Duan and Jiaying Liu*

Peking University, Beijing, China

ABSTRACT

The past decades have witnessed the rapid development of image and video coding techniques in the era of big data. However, the signal fidelity-driven coding pipeline design limits the capability of the existing image/video coding frameworks to fulfill the needs of both machine and human vision. In this paper, we come up with a novel image coding framework by leveraging both the compressive and the generative models, to support machine vision and human perception tasks jointly. Given an input image, the feature analysis is first applied, and then the generative model is employed to perform image reconstruction with features and additional reference pixels, in which compact edge maps are extracted in this work to connect both kinds of vision in a scalable way. The compact edge map serves as the basic layer for machine vision tasks, and the reference pixels act as a sort of enhanced layer to guarantee signal fidelity for human vision. By introducing advanced generative models, we train a flexible network to reconstruct images from compact feature representations and the reference pixels. Experimental results demonstrate the superiority of our framework in both human visual quality and facial landmark detection, which provide useful evidence on the emerging standardization efforts on MPEG VCM (Video Coding for Machine)¹. Our project website is available at <https://williamyang1991.github.io/projects/VCM-Face/>.

Index Terms— Video Coding for Machine, Image Coding, Scalable Coding, Generative Compression

1. INTRODUCTION

Image compression has been one of the most fundamental techniques in media sharing and storage. The typical goal of image compression is to preserve as much signal fidelity with the bit-rate constraint. The mainstream hybrid coding scheme for images, such as JPEG and JPEG 2000 which typically include transform, quantization, and entropy coding modules, has been developed for decades. It improves the signal fidelity-driven metrics significantly and benefits human vision continuously.

However, in the big data era when massive amounts of data generated everyday needs to be compressed, stored and analyzed, existing compression methods get into troubles to fulfill the needs of both machine and human vision. The sequential compression and analysis paradigm is expensive and even intractable if we expect to maintain the quality of the reconstructed videos. On the other way, when the compression ratio is high [1], the performance of machine vision tasks degrades significantly.

Several works have made efforts in addressing the problem of video analytics on massive data by directly extracting and compressing features used for machine vision tasks into a compact form, rather than compressing the whole high-quality videos. Several typical features are developed, *e.g.* Scale-Invariant Feature Transform (SIFT) [2], Compact descriptors for visual search (CDVS) [3] for image understanding, and skeleton for human action recognition [4]. In this way, the process of feature extraction, compression and transmission becomes light-weighted and less amount of bitstreams are to be handled.

Though these features are compact and highly effective for machine vision tasks, they cannot support machine and human vision tasks jointly in a flexible way, which is expected in the new coding paradigm of video coding for machine (VCM). This is due to the huge gap between feature coding for machine vision and signal encoding for human vision. Existing solutions only pay attention to one of these two aspects. In the big data context, it is still an open problem to support a scalable coding paradigm to satisfy both kinds of vision. Some works show potential ways to address this problem. In [5, 6], generative models are used to reconstruct the images based on the encoded features with very few bits towards conceptual coding. In [7], the bitstreams generated by a Variational Auto-Encoder (VAE) is used for image understanding. However, these attempts are still far away from the ideal targets of VCM: the requirement of machine vision is first satisfied to provide the fast analysis, and more bit-rates are additionally used to further improve the visual quality in the reconstruction.

In our work, we take a further step to bridge the gap between image compression for both machine and human vision. By leveraging both compressive and generative models, a scalable image coding framework is constructed to support machine and human vision tasks jointly. In this framework,

*Corresponding Author. This work was supported by National Natural Science Foundation of China under contract No.61772043, Beijing Natural Science Foundation under contract No.L182002, and the National Key R&D Program of China under Grand No.2018AAA0102700.

¹<https://lists.aau.at/mailman/listinfo/mpeg-vcn>

the source image is represented via a compressive model as edge maps and sparse key reference pixels. The edges are parameterized into vectors as the base layer of the coding bits to obtain a compact feature representation, which only takes a small portion of coding bits. Furthermore, the information in our edge maps is shown to be efficient for machine vision tasks, *e.g.* face landmark detection. To better reconstruct the high-quality frame, reference pixels, sampled in accordance with the edges, can be transmitted as a second layer to the decoder. With the reference pixel values, the decoder is able to faithfully reconstruct the image. We adopt a generative model to reconstruct high-quality images from the sparse edge representations. Experiments on both machine and human vision show significant improvements compared with existing methods, which provide useful evidence on the emerging standardization efforts on MPEG VCM.

In summary, the contributions of this work are threefold:

- We propose an image coding framework that leverages the compressive model to extract highly compact representations of an image and faithfully reconstruct the original image from the bitstreams with the generative model.
- We design the vision-driven compact representations for image compression, where the critical image structure and color information is sparsely encoded. A deep generative network is further proposed to effectively recover images from our compact representations.
- A good balance between human and machine vision is stricken, where we achieve 90% and 73% human vision preferences in terms of fidelity and aesthetics, respectively, and achieve an error drop of 44.75% in the machine vision facial landmark detection task.

The rest of this paper is organized as follows. Section 2 reviews related works. Section 3 presents the proposed scalable image coding method. Experimental results are shown in Section 4 and concluding remarks are given in Section 5.

2. RELATED WORK

Feature based Image Coding. Besides the mainstream transform based codecs [8, 9], there have been other approaches to explore encoding representative image features for reconstruction. In [6], a generative compression framework is proposed to encode an image into low-bit-rate latent code and exploit recurrent generative networks for reconstruction. With compressive variational auto-encoders (VAE) [10], generative networks are also utilized in [5] to reconstruct images from edges and latent features produced by neural networks. Though these frameworks encode compact feature representations of images, they are not shown to both satisfy the need of human and machine vision. In [7], a deep-based encoder is designed to produce latent code that simultaneously serves for machine vision tasks and image reconstruction. However, the encoded feature representation is non-scalable as the full

bit-stream is needed to support the machine vision tasks, neglecting the sparsity features for machine vision. In this work, we explore to encode a base layer of features to facilitate machine vision and an additional layer to improve signal fidelity.

Image Generation. Image generation aims to generate new images. Recent image generation methods focus on the powerful generative adversarial networks (GAN) [11] to learn data distribution using two adversarial networks. By incorporating additional information such as the text, labels, segmentation maps and edges as inputs, users are able to control the output with these conditions. The advanced GAN has shown impressive capability of data distribution learning to recover abundant information that well matches human visions from limited conditions. Such an advantage is also verified by the closely related image inpainting task, where plausible image content is generated from very sparse contextual information [12]. It demonstrates the potential for vision-driven image coding, which forms our research focus in this paper.

3. PROPOSED METHOD

In this section, we describe our vision-driven image compression framework. As shown in Fig. 1, we first extract sparse edges to depict the key structure information of the input image (Section 3.1). We then extract vision-driven compact representations through compressive analysis over the edges and the original images. (Section 3.2). Finally, we train a deep neural network to reconstruct the original image from our compact representation. (Section 3.3).

3.1. Sparse Edge Extraction

Edges are one of the most highly abstract and sparse image representations. Edges depict the key structure information of the image, which is consistent with the human vision. Humans are able to identify the objects from several lines and even infer fine details such as the colors and textures. To this end, we are inspired to build our compact representation using sparse edges. We will show later that images can be plausibly reconstructed purely from its edges based on the robust data distribution learned by GAN.

Specifically, for an input image I , we first use fast edge detection [13] based on structured forests to detect the edge map of I . Then, we follow the post process suggested by pix2pix [14] to binarize the edge maps and discard trivial edges that contain less than 10 pixels.

Meanwhile, color is another critical important information for human perception. Color constitutes the main characteristics of the spaces circumscribed by edge lines. Moreover, as a basic low-level feature, it can even impact some high-level concepts such as emotions. Thus, in addition to the edges extracted, we are going to extract the compact color representations, which will be detailed in the next subsection.

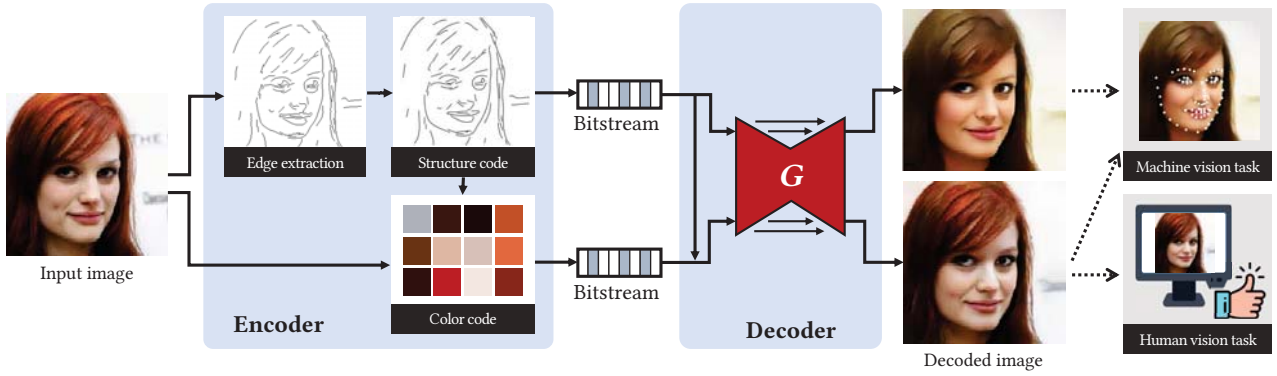


Fig. 1: Overview of the proposed vision-driven image coding framework.

3.2. Compact Representation Extraction

Although edge maps are sparse representations of images, coding such maps into compact bit-streams is still not straight-forward. Existing works in feature-based image compression exploit recurrent generative neural networks [6] or resort to HEVC Screen Content Coding [5, 15]. However, these approaches do not fully exploit the sparsity of the edge maps, as they are mostly based on pixel-level representation or partitioning, and not designed to trace edges. It results in inefficiency in coding binary maps consisting only of edges of uniform width.

To explore a more effective way to encode the edge maps, in our approach, we propose to trace the edges into vector graphics. We adopt the image tracing tool [16] to convert the binary edge image into vectorized representations. The edges are approximated into straight lines and Bézier curves, following the Scalable Vector Graphics (SVG) syntax. To be specific, we use three kinds of operation markers, namely **M**ove, **L**ine and **C**urve. Operation $\mathbf{M}(x, y)$ indicates moving to point (x, y) without drawing a line. $\mathbf{L}(x, y)$ refers to drawing a straight line from the last point (either moved to or ended a line or curve) to the target point (x, y) . $\mathbf{C}(p_a, p_b, p_t)$ denotes the operation to draw a cubic Bézier curve from the current point to the target point p_t , with the intermediate points p_a and p_b . As edge maps of natural images are usually smooth, they can be well approximated by the above-mentioned lines and curves, which only takes a small number of parameters. To further squeeze out redundancy in the parameters, we adopt the Prediction for Partial Matching (PPM) [17] compression scheme to losslessly compress the quantized parameters for the lines and curves into compact bit-streams.

While edge maps provide much of the information about the structure, the information to restore color representation is lost during the parameterization. To support the scalable coding scheme, we propose to embed pixel-level representation as a second layer in accordance with the encoded structural description. We sparsely sample pixels near the lines and curves. As shown in Fig. 2, for a straight line, we sample two




Fig. 2: Illustration of our vectorized structure representation and point samplings for color representation. (a) A vectorized edge map. (b) For straight segments, two points are selected as the reference, according to the slope α . (c) For Bézier curves, one inner point is selected.

points near the midpoint. The slope of the line is calculated to determine whether the two points are chosen horizontally or vertically. If the line is more close to horizontal ($\alpha < 45^\circ$), the two reference points are sampled vertically, and it goes horizontal if $\alpha > 45^\circ$. For a Bézier curve with starting point p_s , intermediate points p_1, p_2 and target point p_t , we first extract the contact point of the curve and the tangent line in parallel with the vector $\overrightarrow{p_s p_t}$. We calculate the slope of the tangent line to determine whether to choose the point vertically or horizontally, just like the straight lines. Additionally, to control the bit-rate and maintain the most informative information from the pixels, we only sample the point at the inner side of the curve, which is expected to have greater gradients and contain more information. The pixel, represented in RGB value, is signaled to the decoder in order as a second layer to provide more fidelity in color. The decoder places the received reference color points following the same rules that the encoder extracts those points, based on the edge maps. Thus, no additional bits are needed to record the positions of the selected pixels.

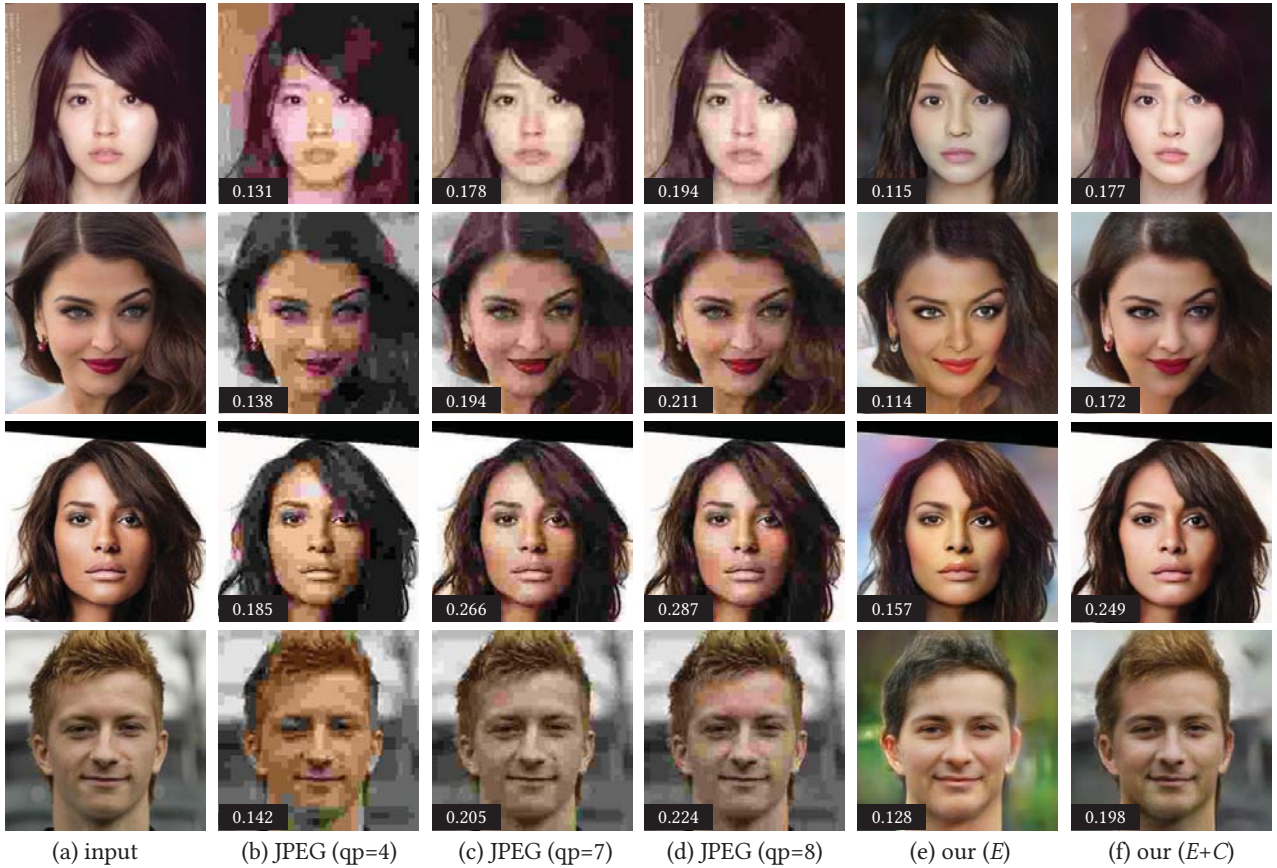


Fig. 3: Visual comparison with JPEG compression. (a) Input image. (b)-(d) Images compressed by JPEG using quality parameter of 4, 7 and 8, respectively. (e) Our decoded images using the encoded edge representations. (f) Our decoded images using both the encoded edge representation and color representation. For each reconstructed image, its bit-rate (bit per pixel, bpp) is shown in the lower left black box.

3.3. Adversarial-based Image Reconstruction

Given the proposed compact representation of edges and colors, we are going to recover an image as close as possible to the original image. The main idea is to leverage GAN to learn robust data distribution, which maps our sparse representation back to the original image spaces and benefits both human visual quality and machine visual tasks.

Specifically, we first convert our compact representation back to the image domain by rendering the vector graphic as a normal bitmap E . The sparsely sampled points are rendered as a one-channel image mask M where 1 means the corresponding pixel is sampled and 0 vice versa. And finally, another three-channel image C is provided with the color values of the sampled pixels at the corresponding locations. The remaining unknown pixels are set to 0. Through the conversion, we transform our decoding task as a standard machine vision task of image inpainting augmented with extra edge information. C can be regarded to be obtained by the original image I with missing regions indicated by M .

Taking use of the advancement of image inpainting re-

search, we design our decoding network as pix2pix [14]. It contains fully convolutional encoders and decoders, where the low-level information is conveyed to the outputs via skips connections to enforce the structure and color constraints from the inputs. Let our generator and discriminator denoted as G and D , respectively. Then G is used to map the input of E , C and M to a reconstructed image $I_G = G(E, C, M)$ to approach I in both color and structure senses through a reconstruction loss:

$$\mathcal{L}_r = \mathbb{E} [\lambda_1 \|I_G - I\| + \lambda_2 \text{SSIM}(I_G, I)], \quad (1)$$

where L_1 measures the color discrepancy between the reconstructed image and I and SSIM [18] emphasizes the structural similarity, weighted by λ_1 and λ_2 , respectively. In addition to these human-perceptual criteria, we incorporate perceptual loss [19] to enhance the machine-perceptual quality of I_G ,

$$\mathcal{L}_p = \mathbb{E} [\lambda_3 \text{PERC}(I_G, I)]. \quad (2)$$

Finally, we use hinge loss [20] as our adversarial objective function to learn the data distribution:

$$\mathcal{L}_G = -\mathbb{E}[D(I_G, E, M)], \quad (3)$$




Fig. 4: Illustration of the averaged normalized point-to-point error (NPE) on facial landmark detection and bit-rate of JPEG compression and the proposed method.

$$\mathcal{L}_D = \mathbb{E}[\text{ReLU}(\tau + D(I_G, E, M))] + \mathbb{E}[\text{ReLU}(\tau - D(I, E, M))], \quad (4)$$

where $\tau = 10$ is a margin parameter. Here we use channel-wise concatenation to feed multiple inputs into G and D .

Reconstruct without RGB. Note that for some high-level machine vision tasks such as image segmentation and image detection that do not rely on color information much, our framework is scalable to reconstruct images purely from E without M and C , which further saves bit-rate. To be specific, we only need to revise the input channel number of the first layer of G and D such that G receives one-channel input E and D receives four-channel input (I, E) or (I_G, E) , respectively. Beyond that, other settings are the same as our aforementioned reconstruction process with color information.

4. EXPERIMENTAL RESULTS

In this section, we present the experimental results of the proposed method for both the task of human vision and machine vision. We first evaluate our method with respect to human visual quality both qualitatively and quantitatively in Section 4.1. Then we test our method on the high-level facial landmark detection task in Section 4.2. We choose the VGGFace2 [21] dataset for evaluation considering the pervasiveness and importance of facial images in our daily life and industry. We filter the images in VGGFace2 that have small resolution and low quality, and finally use 39,122 images from the training set to train our reconstruction network and 20,665 images from the testing set for performance evaluation. To train our network, we set $\lambda_1 = 100$, $\lambda_3 = 1$ and $\tau = 10$. λ_2 is set to 0 and 1000 for the human vision evaluation and machine vision evaluation, respectively.

4.1. Human Vision: Visual Quality Evaluation

Qualitative evaluations. In Fig. 3, we present a visual comparison of the proposed method with JPEG compression un-




Fig. 5: Cumulative error distribution of JPEG compression and the proposed method on facial landmark detection.

Table 1: The preference ratio on fidelity and aesthetics of different methods at different bit-rates.

Method	Bit-Rate (bpp)	Fidelity	Aesthetics
JPEG ($qp = 4$)	0.152	0.00	0.00
our (E)	0.134	0.04	0.24
JPEG ($qp = 7$)	0.214	0.02	0.01
JPEG ($qp = 8$)	0.234	0.04	0.02
our ($E + C$)	0.209	0.90	0.73

der different quality factor parameter (qp), which are selected to match the bit-rate of our method for fair comparison. Specifically, for our reconstructed images decoded without color cues, we show the JPEG compression results with $qp = 4$; while for our reconstructed images decoded with full color and structure cues, we use $qp = 7$ and $qp = 8$. It can be observed that JPEG compression yields distinct block artifacts, which greatly decrease visual quality. By comparison, our method produces more natural results.

Quantitative evaluations. We perform user studies for quantitative evaluations. Besides the four cases shown in Fig. 3, we randomly select six cases to add up to 10 cases from the testing data shown to the participants. Each subject is asked to select one from the five results that best matches the original image (**Fidelity**) and has the best visual quality (**Aesthetics**). A total of 10 subjects participate in this study and a total of 200 selections are tallied. The preference ratio is used as the evaluation metric. It is calculated as the ratio of a method selected in all comparisons with this method. As shown in Table 1, the proposed structure-color-hybrid method obtains the best average preference ratio of 0.90 and 0.73 for both the fidelity and aesthetics, respectively, outperforming JPEG compression under the similar bit-rate. The user study quantitatively verifies the superiority of our method.

4.2. Machine Vision: Landmark Detection

The machine vision performance of our method is verified on the high-level facial landmark detection task. We perform facial landmark detection [22] on the original VG-Face2 [21] dataset and the reconstructed dataset by JPEG and our method. Detection results on the original data are served as ground truth. We then calculate the normalized point-to-point error (NME) [23] between the detection results on the compressed data and the ground truth. Fig. 4 illustrates the averaged NME and the bit-rate of JPEG compression and our method. It can be clearly seen that our method achieves much fewer errors at the similar bit-rate compared to JPEG. Specifically, NME of our method without color cues is only 4.03%, which is 44.75% lower than JPEG under $qp = 4$. Meanwhile, with color cues, our method achieves merely 3.33% NME, 1.15% lower than JPEG under $qp = 8$. Fig. 5 further shows the cumulative error distribution, where more than 90% of the images reconstructed by the proposed method have tiny errors less than 5%, showing great robustness.

5. CONCLUSION AND DISCUSSION

In this paper, we present a new image coding framework to facilitate both human vision and machine vision. The input image is first analyzed and compressed as the compact structure and color representations. Leveraging the advanced generative model in machine vision, we train a network to faithfully reconstruct images from the compact representations. Experimental results demonstrate the superiority of the proposed method in both human visual quality and facial landmark detection. This paper presents the first attempt towards VCM with respect to image coding via scalable feature-based compression. As a future direction, we would like to explore temporal feature modeling for video coding to more pervasively benefit human vision and machine vision.

6. REFERENCES

- [1] L. Ding, Y. Tian, H. Fan, Y. Wang, and T. Huang. Rate-performance-loss optimization for inter-frame deep feature coding from videos. 26(12):5743–5757, Dec 2017.
- [2] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [3] L. Duan, V. Chandrasekhar, J. Chen, J. Lin, Z. Wang, T. Huang, B. Girod, and W. Gao. Overview of the mpeg-cdvs standard. pages 179–194, Jan 2016.
- [4] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *Proc. of AAAI Conf. on Artificial Intelligence*, 2017.
- [5] J. Chang, Q. Mao, Z. Zhao, S. Wang, S. Wang, H. Zhu, and S. Ma. Layered conceptual image compression via deep semantic synthesis. In *Proc. IEEE Int'l Conf. Image Processing*, 2019.
- [6] K. Gregor, F. Besse, D. J. Rezende, I. Danihelka, and D. Wierstra. Towards conceptual compression. In *Proc. of Advances In Neural Information Processing Systems*, 2016.
- [7] R. Torfason, F. Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. Van Gool. Towards image understanding from deep compression without decoding. In *Proc. of Int'l Conf. on Learning Representation*, 2018.
- [8] G. K. Wallace. The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1):xviii–xxxiv, 1992.
- [9] C. Christopoulos, A. Skodras, and T. Ebrahimi. The jpeg2000 still image coding system: an overview. *IEEE Transactions on Consumer Electronics*, 46(4):1103–1127, 2000.
- [10] L. Theis, W. Shi, A. Cunningham, and F. Huszár. Lossy image compression with compressive autoencoders. In *Proc. of Int'l Conf. on Learning Representations*, 2017.
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. WardeFarley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Proc. of Advances in Neural Information Processing Systems*, 2014.
- [12] T. Dekel, C. Gan, D. Krishnan, C. Liu, and W. T. Freeman. Sparse, smart contours to represent and edit images. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2018.
- [13] P. Dollár and C. L. Zitnick. Structured forests for fast edge detection. In *Proc. Int'l Conf. Computer Vision*, 2013.
- [14] P. Isola, J. Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2017.
- [15] J. Xu, R. Joshi, and R. A. Cohen. Overview of the emerging hevce screen content coding extension. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(1):50–62, 2015.
- [16] M. Weber. AutoTrace: a program for converting bitmap to vector graphic. 1998. <http://autotrace.sourceforge.net/>.
- [17] J. Cleary and I. Witten. Data compression using adaptive coding and partial string matching. *IEEE Transactions on Communications*, 32(4):396–402, 1984.
- [18] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [19] J. Johnson, A. Alahi, and F. F. Li. Perceptual losses for real-time style transfer and super-resolution. In *Proc. European Conf. Computer Vision*, 2016.
- [20] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Free-form image inpainting with gated convolution. In *Proc. Int'l Conf. Computer Vision*, 2019.
- [21] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *IEEE Int'l Conf. on Automatic Face & Gesture Recognition*. IEEE, 2018.
- [22] A. Bulat and G. Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *Proc. Int'l Conf. Computer Vision*, pages 1021–1030, 2017.
- [23] D. Cristinacce and T. F. Cootes. Feature detection and tracking with constrained local models. In *Proc. British Machine Vision Conference*, 2006.