

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

Public Domain Mark 1.0

<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback



Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Neural-Based Compression Scheme for Solar Image Data

ALI ZAFARI , Graduate Student Member, IEEE

ATEFEH KHOSHKAHTINAT 
West Virginia University, Morgantown, WV USA

JEREMY A. GRAJEDA 
New Mexico State University, Las Cruces, NM USA

PIYUSH M. MEHTA 
NASSER M. NASRABADI , Fellow, IEEE
West Virginia University, Morgantown, WV USA

LAURA E. BOUCHERON , Member, IEEE
New Mexico State University, Las Cruces, NM USA

BARBARA J. THOMPSON 
MICHAEL S. F. KIRK

DANIEL E. DA SILVA 
NASA Goddard Space Flight Center, Greenbelt, MD USA

Studying the solar system and especially the Sun relies on the data gathered daily from space missions. These missions are data-intensive and compressing this data to make them efficiently transferable to the ground station is a twofold decision to make. Stronger compression methods, by distorting the data, can increase data throughput at the cost of accuracy, which could affect scientific analysis of the data. On the other hand, preserving subtle details in the compressed data requires a high amount of data to be transferred, reducing the

Manuscript received 4 May 2023; revised 9 September 2023 and 11 October 2023; accepted 31 October 2023. Date of publication 13 November 2023; date of current version 9 February 2024.

DOI: No. 10.1109/TAES.2023.3332056

Refereeing of this contribution was handled by H. Mir.

This work was supported by the National Aeronautics and Space Administration (NASA), under the title of Adaptive and Scalable Data Compression for Deep Space Data Transfer Applications using Deep Learning, under Grant 80NSSC21M0322.

Authors' addresses: Ali Zafari, Atefeh Khoshkhahtinat, Piyush M. Mehta, and Nasser M. Nasrabadi are with the West Virginia University, Morgantown, WV 26506 USA, E-mail: (az00004@mix.wvu.edu; ak00043@mix.wvu.edu; piyush.mehta@mail.wvu.edu; nasser.nasrabadi@mail.wvu.edu); Jeremy A. Grajeda and Laura E. Boucheron are with the New Mexico State University, Las Cruces, NM 88003 USA, E-mail: (jgra@nmsu.edu; lboucher@nmsu.edu); Barbara J. Thompson, Michael S. F. Kirk, and Daniel E. da Silva are with the NASA Goddard Space Flight Center, Greenbelt, MD 20771 USA, E-mail: (barbara.j.thompson@nasa.gov; michael.s.kirk@nasa.gov; daniel.e.dasilva@nasa.gov). (*Corresponding author: Ali Zafari.*)

0018-9251 © 2023 IEEE

desired gains from compression. In this work, we propose a neural network-based lossy compression method to be used in NASA's data-intensive imagery missions. We chose NASA's Solar Dynamics Observatory (SDO) mission, which transmits 1.4 TB of data each day as a proof of concept for the proposed algorithm. In this work, we propose an adversarially trained neural network, equipped with local and nonlocal attention modules to capture both the local and global structure of the image resulting in a better tradeoff in rate-distortion (RD) compared with conventional hand-engineered codecs. The RD variational autoencoder used in this work is jointly trained with a channel-dependent entropy model as a shared prior between the analysis and synthesis transforms to make the entropy coding of the latent code more effective. We also studied how optimizing perceptual losses could help our neural compressor to preserve high-frequency details of the data in the reconstructed compressed image. Our neural image compression algorithm outperforms currently-in-use and state-of-the-art codecs, such as JPEG and JPEG-2000, in terms of the RD performance when compressing extreme-ultraviolet (EUV) data. As a proof of concept for use of this algorithm in SDO data analysis, we have performed coronal hole detection using our compressed images, and generated consistent segmentations, even at a compression rate of ~ 0.1 bits per pixel (compared with 8 bits per pixel on the original data) using EUV data from SDO.

I. INTRODUCTION

Learning-based image compression outperforms [1] almost all traditional codecs including JPEG [2] and JPEG-2000 [3]. With a basis in convolutional neural networks (CNNs), the performance of said learned compression methods has been improved through various investigations: enhanced entropy model [4], [5], [6], learned representation augmentation via attention [7], [8] and incorporating adversarial training for improved perceptual quality of reconstruction [9], [10], [11].

In this work, we demonstrate the potential of neural network compression codecs to serve as the go-to approach for future space missions for onboard data compression. We baseline our novel algorithm on the solar image data collected by the Atmospheric Imaging Assembly instrument onboard the Solar Dynamics Observatory (SDO) spacecraft and compare the performance with traditional methods. Finally, we show that a CH segmentation [12] is minimally affected even at extremely low bit-rate compression provided by our proposed neural compression method. Extremely compressed images using our network preserves the required details for the task of CH segmentation on images downloaded from the SDO mission.

Contributions of This Work: Preliminary results of our proposed algorithm was published in [13], introducing the application of neural image compression on downloaded solar imagery data. The primary contribution of the current work is to critically evaluate the performance of our algorithm on downstream science and/or operational tasks (e.g., CH segmentation). In addition, we provide a detailed description of the joint forward and backward adapted entropy model used to improve the rate-distortion performance of the proposed algorithm, which was not available in [13].

The rest of this article is organized as follows. Section II provides a review of neural compression autoencoders and their potential application for a solar mission with a discussion of the downstream scientific task of segmentation on

the SDO data. Section III is devoted to our proposed method. Experimental results and ablation studies are described in Section IV. Finally, Section V concludes this article.

II. RELATED WORK

A. Neural Image Compression

Transform coding-based image compression algorithms share four main steps to compress an image [14]. First, a transform is applied to the image to transfer the pixels from their spatial domain to a transform domain that reduces correlation between pixels and thereby, results in many small coefficients (i.e., information packing). Second, the transformed image is quantized to discard less significant information from the data in the transform domain. Third, entropy coding is utilized to losslessly encode the quantized samples into a stream of ones and zeros. This bitstream is the compressed image. The fourth and final step occurs at the receiving end (or at the reconstructing step), which is responsible for decoding the quantized values to the original space of the input image. The first and most widely used architecture to mimic this scenario in deep neural networks is the convolutional autoencoder, which has shown its superiority in terms of rate-distortion performance in the literature [15]. Both the encoding and decoding parts of the traditional transform codec can be imitated by an autoencoder [16].

In an end-to-end optimization of an autoencoder, problems arise when we want to quantize the bottleneck to remove redundancies in order to reach high compression ratios. Artificial neural network (ANNs) are optimized using gradient descent algorithms, which update the parameters of the network by backpropagating the gradients of the loss function. Gradients of the quantization process are not useful for optimizing the loss function as its value is either zero or infinity. As a result, we need to approximate hard discrete quantization with an operation yielding informative gradients for updating the parameters of the network. The most widely used approach was proposed by [17], inherited from [18], in which they showed that adding independent and identically distributed unit uniform noise can be interpreted equivalently as doing scalar quantization on the bottleneck. This method is thoroughly discussed in Section III-A3. By applying this noise to the bottleneck, we can optimize the differential entropy of the continuous approximation as a variational upper bound [19] to reduce the entropy of the bottleneck. Low entropy messages are compressed more efficiently into bitstreams [20].

Optimal compression in theory can be achieved by vector quantization [21]. In vector quantization, each data point is represented by a prototype, and a collection of prototypes, called a codebook, is shared between sender and receiver. The application of vector quantization in ANN-based compression has been investigated by Agustsson et al. [22], with the cost of a complicated training procedure. To make the training more accessible, neural compression algorithms

follow the classical approach to avoid the complexity of vector quantization. In classical image compression schemes, e.g., JPEG [2], to get the best out of the quantization process, the first step is to apply an invertible linear transform and translate the image into decorrelated coefficients using a linear transform, e.g., discrete cosine transform in JPEG [2]. By doing so, scalar quantization can reach a reasonable performance close to vector quantization [14]. On the other hand, it has been shown [15], [16] that a joint-optimized learned nonlinear transform, i.e., neural network, followed by scalar quantization is sufficient to approximate a parametric form of vector quantization. As will be shown in Section III, replacing the actual quantization of the latent code/bottleneck of the autoencoder with a uniform noise approximation in the bottleneck of a vanilla autoencoder during training of the network [23] will transform it to a variational autoencoder (VAE) [24]. The difference between the rate-distortion VAE and its vanilla version is the chosen prior for the latent variables. In autoencoder-based image compression, the Gaussian prior of the VAE is replaced with a unit uniform distribution centered on integer numbers to imitate the scalar quantization process.

B. Generative Adversarial Training

Generative adversarial networks (GANs) [25] have emerged as a transformative technology in the fields of computer vision and aerospace systems, offering a wide array of innovative applications. In the realm of computer vision, GANs have played a pivotal role in image generation [26], [27], enhancement [28], and manipulation [29]. Researchers have leveraged GANs to generate high-resolution images from low-resolution inputs [30], thereby enhancing the quality of visual data, a crucial advancement for various computer vision applications. GANs have also demonstrated their ability to generate realistic synthetic images [26], [27], which find application in data augmentation for training deep neural network models, particularly in scenarios where labeled datasets are limited. In addition, GANs have been instrumental in style transfer [31], enabling the creation of unique visual content by merging the style of one image with the content of another [32]. GANs are also used in domain translation tasks to make realistic photos out of sketches or colorizing grayscale images [33], [34].

The utilization of GANs has been particularly noteworthy in the aerospace field of study. GANs have been employed to enhance satellite imagery for Earth observation and environmental monitoring [35]. They facilitate the removal of sensor noise [36], the enhancement of image resolution, and the synthesis of missing data [37], or improving the overall quality of satellite-based remote sensing [38]. GANs have shown great potential in synthesizing RF micro-Doppler signatures to tackle the challenge of limited sample availability and facilitate the training of more complex deep neural networks to improve RF

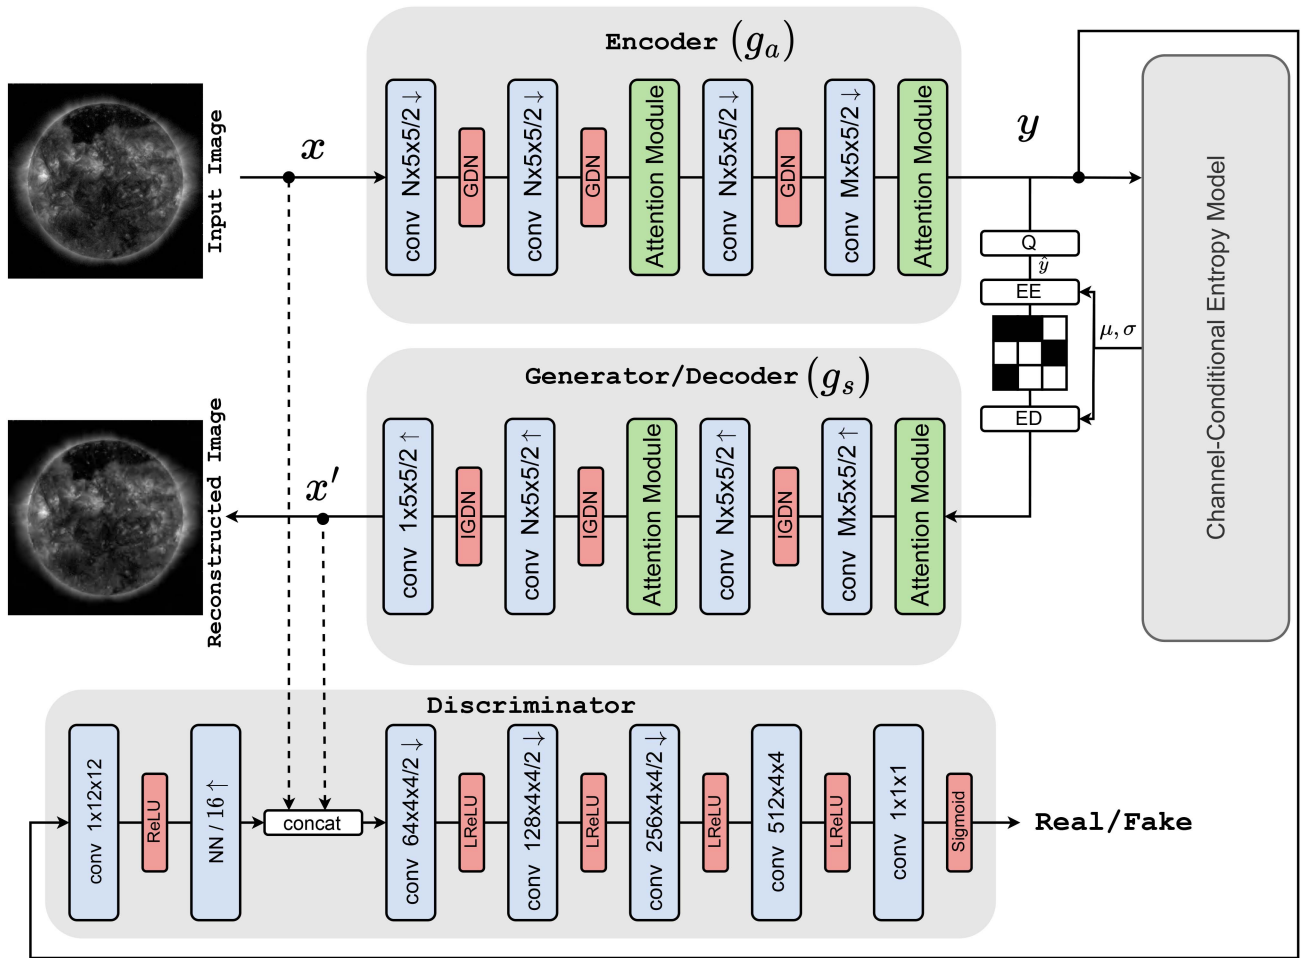


Fig. 1. Network architecture. The input image traverses through a series of convolutional layers yielding 16 times smaller (spatially) feature maps than the original input dimensions. To reconstruct the image, the decoder follows the same dimensions of the encoder network in a reverse order using transpose convolution layers. A conditional discriminator encourages the generator (decoder) toward better perceptual quality. The number of channels in the encoder and decoder are set as $N = 192$ and $M = 320$, respectively. Q performs scalar quantization. EE and ED indicate the entropy encoder and decoder, respectively. The checkerboard box represents the bitstream of the compressed image including only zeros and ones. μ and σ are predicted parameters of the latent code probability distribution, defined by the entropy estimation model, as shown in Fig. 2. LReLU denotes leaky rectified linear unit activation function. GDN and IGDN correspond to GDN nonlinearity and its inverse, discussed in Section IV-B.

signal classification performance [39]. Training neural networks in a GAN framework also improved unsupervised domain adaptation in satellite pose estimation [40]. These advancements have the potential to significantly impact industries, from enhancing image analysis to improving the safety and efficiency of aerospace operations.

C. Attention Mechanism in Neural Networks

When it comes to computer vision, deep CNNs are the de facto standard despite their poor performance in capturing long-range dependencies [41], [42]. The performance degradation of CNNs is attributed to their local receptive field, which is mainly because of the limited kernel size [43] of the filters. Efforts have been made to help CNNs capture a more robust representation of the input image. One naïve solution is to make the network deeper, but other problems will arise in training such networks, which have led to the introduction of deep residual networks [44]. Although

increasing the parameters of a network will generally lead to a richer representation and better performance, it will make training such networks harder since overfitting can easily happen in these overparametrized neural networks. Attention mechanisms have been proposed to address the problem of the local receptive field by keeping the depth of the network unchanged. In [45], a single module was proposed to be included in between sequential convolutional layers, consisting of two branches, the *trunk* to process local features and the *mask* to decide which of the local features in the trunk are more important to be passed to the next convolutional layer, as in Fig. 3(a).

In contrast to local attention, Wang et al. [46] first discussed how nonlocal attention can be viewed as a special case of a nonlocal algorithm, which was traditionally used as a method to denoise images [47]. The idea was to find similar pixels/patches in the image/feature map and replace them with a weighted sum over all the others, with higher




Fig. 2. Joint FA and BA entropy model with the channel conditioning assumption in the probability distribution of the latents. Conv represents a convolutional layer with a kernel size of 3×3 with stride and padding of size 1. Due to limited space, this figure shows channelwise conditioning on only three slices (red, green, and blue), while in the implementation the latent code is divided into ten slices. After decoding, the first slice (\hat{y}_1) is used to decode the second one (\hat{y}_2), and both of them are used for decoding the third slice (\hat{y}_3). ReLu stands for a rectified linear unit. Q, EE, and ED denote quantization, and entropy encoder/decoder, respectively. Checkerboard boxes represent bitstream codes. The estimated probability distribution that is parametrized by μ and σ is shared as the prior over the latent code both on the encoder and decoder side at the time of entropy coding.

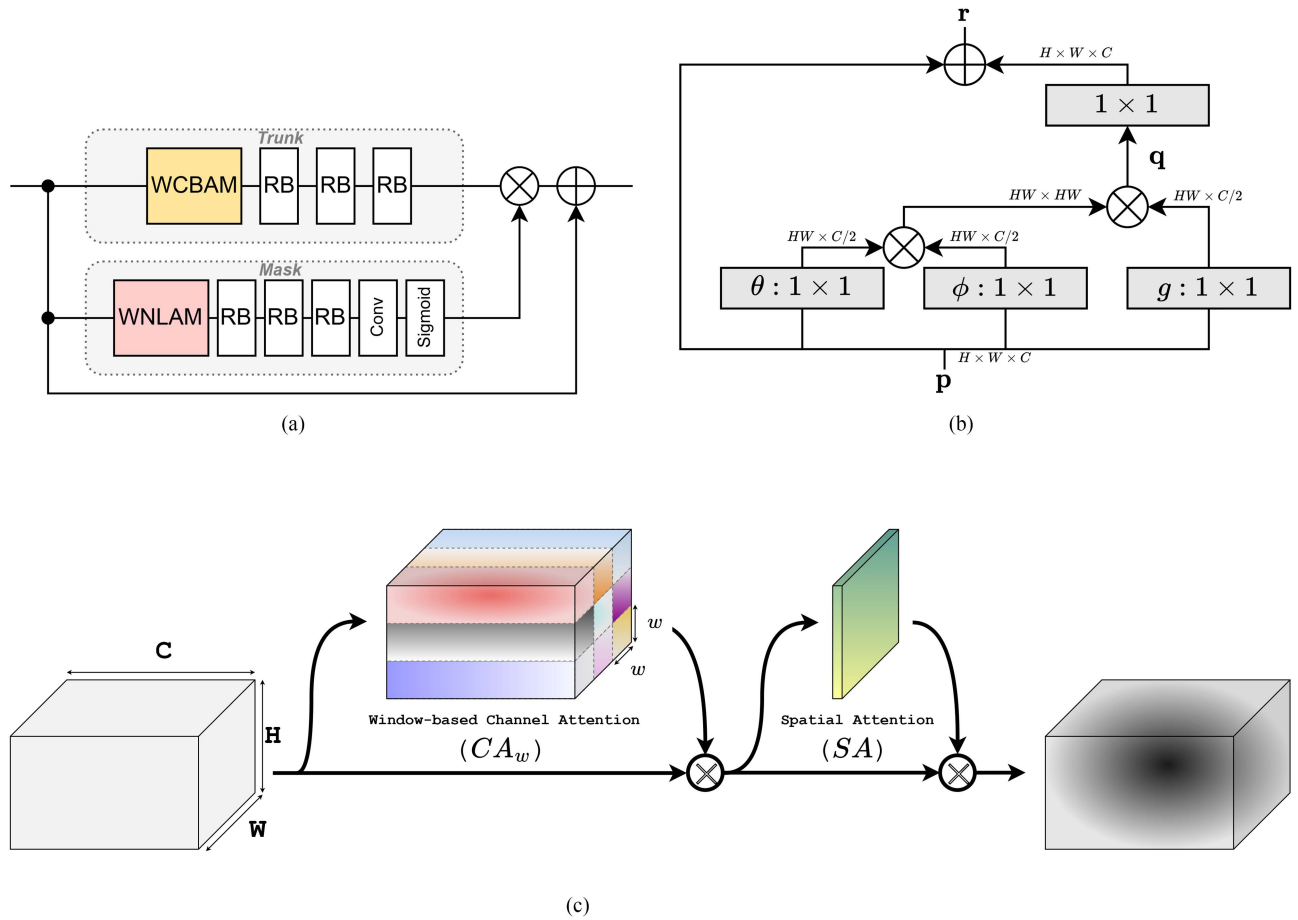


Fig. 3. Attention module architecture. (a) Attention module with skip connection. RB denotes a residual block. WCBAM denotes a window-based CBAM, and WNLAM is a window-based nonlocal attention module. (b) WNLAM. (c) WCBAM. A feature map of C channels with spatial dimensions $H \times W$ is the input to the WCBAM block. w is the window size over which to calculate CA.

weights for more similar ones. It can be inferred from [46] that vision transformers [48] are all special cases of the nonlocal attention mechanism.

The nonlocal attention block [see Fig. 3(b)] helps the *mask* branch to efficiently learn the most informative parts of features (in the *trunk*) for the task at hand [49]. Zhang et al. [49] also added a skip connection to help the output feature maps be richer, by letting the module have access to both attended and raw features. This skip connection prevents vanishing gradients as well.

Another recently proposed mechanism to incorporate attention in CNNs has been introduced in [50]. It is an enhanced version of the squeeze-and-excitation network [51], to apply attention on both spatial and channel feature maps separately. This way of applying attention is simpler and computationally more efficient than previous computation-heavy attention mechanisms based on pretrained networks or complex calculations [52].

III. METHODS

A. Generative Image Compression

Autoencoder-based learned image compression networks, such as the one we have proposed in Fig. 1, generally consist of two major parts. The first part includes the encoder/decoder network and the second part is the bottleneck entropy estimation network. The former is discussed in this section, and Section III-A2 is devoted to describing the functionality of the latter. As illustrated in Figs. 1 and 2, the network input (\mathbf{x}) and output (\mathbf{x}') relations can be summarized as follows:

$$\begin{aligned} \mathbf{x}' &= g_s(\text{ED}(\text{EE}(\hat{\mathbf{y}})); \theta_g) \\ \hat{\mathbf{y}} &= \lfloor g_a(\mathbf{x}; \phi_g) \rfloor \\ \hat{\mathbf{z}} &= \lfloor h_a(\mathbf{y}; \phi_h) \rfloor \end{aligned} \quad (1)$$

in which $\lfloor \cdot \rfloor$ denotes the quantization of a real value to the nearest integer number, $\hat{\mathbf{y}}$ is the quantized latent variable, and $\hat{\mathbf{z}}$ is its quantized hyperprior, which is discussed in Section III-A2. ED and EE denote the entropy decoder and entropy encoder, respectively. The encoder and decoder nonlinear transforms are represented by g_a and g_s with their learned parameters, ϕ_g and θ_g , respectively. The subscripts a and s refer to *analysis* and *synthesis* as they are common designations for the compression and decompression processes in the area of transform coding-based compression. h_a is the analysis transform to get the hyperpriors of the entropy estimation model, defined by its parameters ϕ_h . Throughout the following section, we use the terms quantized bottleneck and quantized latent code interchangeably.

1) *Learning Objective*: Any learned image compression network tries to find an optimal tradeoff in the rate-distortion plane by trading off distortion for the expected bit-rate (or vice versa), governed by a Lagrangian coefficient λ . The rate-distortion tradeoff can be described as

$$R + \lambda D_r \quad (2)$$

where R and D_r correspond to the estimated entropy of the latent code and reconstruction distortion, respectively. The

estimated entropy of the quantized bottleneck, R , represents the rate term. Optimizing the parameters of the neural network will enforce this objective to be minimized. The probability distribution of the latent code is variationally approximated by hyperprior \mathbf{z} . Then, the quantized $\hat{\mathbf{z}}$ is transmitted alongside the compressed image as side information to build the shared prior on the decoder side. Therefore, the entropy of both the quantized bottleneck and its hyperprior should be optimized

$$R = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [-\log_2 P_{\hat{\mathbf{y}}|\hat{\mathbf{z}}}(\hat{\mathbf{y}}|\hat{\mathbf{z}}; \theta_h) - \log_2 P_{\hat{\mathbf{z}}}(\hat{\mathbf{z}}; \psi)] \quad (3)$$

where θ_h and ψ are parameters of the learned entropy model on the latent code ($\hat{\mathbf{y}}$) and hyperprior ($\hat{\mathbf{z}}$), respectively. $P_{\hat{\mathbf{y}}|\hat{\mathbf{z}}}$ is the probability mass function of the discrete bottleneck and $P_{\hat{\mathbf{z}}}$ denotes the probability mass function for the discretized hyperprior.

In (2), D_r accounts for the distortion between the input and output image of the network, which can be measured by any desired metric. The prevalently used criterion to measure distortion between input and output is the mean squared error (mse), which is heavily criticized in the context of computer vision [53] because it often results in the reconstruction of blurry images. Efforts have been made to propose metrics that can adhere perceptually to the human visual system, e.g., the multiscale structural similarity index (MS-SSIM) [54]. Even these metrics have shown weaknesses when intensely scrutinized [55]. Recently, perceptual-aware metrics based on features generated by pretrained neural networks have been proposed. Learned perceptual image patch similarity (LPIPS) introduced by [56] uses trained AlexNet/VGGNet features to compare patches of an image with a corresponding reference. In training our neural compressor, we will enhance its reconstruction loss by exploiting this perceptual metric.

To make the reconstruction closer to the input image, we also consider adversarial training for our decoder network. GANs [9], [10], [25], [57] consisting of a generator and a discriminator subnetwork, are able to better match the distribution of data at reconstruction. In our network, the decoder plays the role of the generator. In the GAN framework, the discriminator forces the decoder output to preserve the distribution of the input image at the time of reconstruction. The proposed objective to be optimized for adversarial training of the generator is a combination of distortion and perception as follows:

$$\begin{aligned} D_r &= \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [\lambda_{\text{recon}} \text{MSE}(\mathbf{x}, \mathbf{x}') \\ &\quad + \lambda_{\text{perc}} \text{LPIPS}(\mathbf{x}, \mathbf{x}') \\ &\quad - \lambda_{\text{adv}} \log D(\mathbf{x}', \mathbf{y})] \end{aligned} \quad (4)$$

where $D(\mathbf{x}', \mathbf{y})$ denotes the classification decision of the discriminator, which is based on two inputs fed into its network, reconstructed image \mathbf{x}' and bottleneck \mathbf{y} , which serves as a condition. In this setting, λ_{adv} regularizes how much the discriminator should be optimized to classify the generator reconstructed image (*fake* image in GAN terminology) as an original image. As a result, the generator parameters are

optimized with a combination of (4) as a distortion and (3) as a rate penalty.

To make the adversarial training feasible we need the discriminator to judge whether its input sample came from the true distribution of data or is a fake generated image, which is the reconstructed image in our network, i.e., \mathbf{x}' . The discriminator will need to be optimized by a separate auxiliary loss, given as

$$L_{\text{disc.}} = \mathbb{E}_{\mathbf{x} \sim p_X} [-\log(D(\mathbf{x}, \mathbf{y}))] + \mathbb{E}_{\mathbf{x}' \sim p_X} [-\log(1 - D(\mathbf{x}', \mathbf{y}))] \quad (5)$$

which is the cross entropy loss between the label provided by the discriminator network (D) and the true labels, assuming label 1 for the original image and 0 for the reconstructed image by the generator.

It has been shown analytically that increasing the perceptual quality of a generator can result in degradation in terms of distortion measures [58]. GANs are a solution to encourage better perceptual quality in the reconstructed image by tolerating an acceptable amount of distortion. More detailed experiments have been adopted in [11] to prove in practice the idea that GANs improve perceptual quality at the cost of a small increase in distortion. Therefore, it would be an expected behavior to have a lower peak signal-to-noise ratio (PSNR) value on a decoder trained adversarially in contrast to a decoder trained merely on distortion metrics. These adversarially trained networks are expected, however, to perform better when measured with perceptually motivated metrics.

2) *Entropy Modeling*: The performance of any learned image compression scheme depends heavily on how well it can estimate the true entropy of the bottleneck. Thus, the objective will be to minimize the cross entropy between the probability model and the latent code's true probability distribution. To make entropy estimation possible, several probability estimation methods have been proposed in the literature, including empirical histogram density estimation [22], [59], piecewise linear models [15], conditioning on a latent variable (hyperprior) [23], and context modeling based on autoregressive models [4].

From a high-level overview, entropy estimation models can be divided into two main categories: forward adaptation (FA) and backward adaptation (BA) models. The former suffers from a low capacity to capture all dependencies in the probability distribution of the latent code and the latter's disadvantage is that the decoding process cannot be parallelized. Learned FA models [16], [23] will only use the information provided during the encoding of the image, while BA methods, which are based on autoregressive models [4], need information from the decoded message as well. In the following section, we discuss the functionality of each of them. We emphasize that using a combined entropy model of both FA and BA is the approach we take in this work.

Forward Adaptation: To model the probability distribution of latent code $\hat{\mathbf{y}}$, some assumptions must be made to make the learning feasible and efficient. The simplest way

to model any multivariate random variable is to assume independence between all its dimensions, which is called the fully factorized model [60], i.e.

$$P_{\hat{\mathbf{y}}}(\hat{\mathbf{y}}) = P_{\hat{y}_1}(\hat{y}_1) P_{\hat{y}_2}(\hat{y}_2) \dots P_{\hat{y}_m}(\hat{y}_m) = \prod_i P_{\hat{y}_i}(\hat{y}_i) \quad (6)$$

where m is the dimension of latent code $\hat{\mathbf{y}}$.

On the other hand, the most flexible and expressive model is to use an autoregressive probability model to capture all dependencies between dimensions of the latent code

$$P_{\hat{\mathbf{y}}}(\hat{\mathbf{y}}) = P_{\hat{y}_1}(\hat{y}_1) P_{\hat{y}_2}(\hat{y}_2 | \hat{y}_1) P_{\hat{y}_3}(\hat{y}_3 | \hat{y}_1, \hat{y}_2) \dots P_{\hat{y}_{m-1}}(\hat{y}_{m-1} | \hat{y}_{<m-1}) P_{\hat{y}_m}(\hat{y}_m | \hat{y}_{<m}) = \prod_i P_{\hat{y}_i}(\hat{y}_i | \text{pa}(\hat{y}_i)) \quad (7)$$

where $\text{pa}(\hat{y}_i) = \hat{y}_{<i}$ denotes the parents of \hat{y}_i , i.e., $\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{i-1}$, whose probability density is conditioned on them. However, letting all the dependencies be visible in the model make it infeasible for practical applications. The curse of dimensionality arises when you model all the conditional dependencies [60], preventing the entropy model from being realized. Even if there was enough computational power to learn this fully visible model, the required time to train the model can easily approach infinity as the dimension of the latent code increases. Therefore, an essential decision is how to modify the modeling to capture only essential dependencies while ignoring irrelevant ones. By doing so the modeling accuracy can be compromised at a reasonable rate but much more efficiently [4].

There is another approach to avoid density modeling based on the probability chain rule (7). We can use a latent variable model (LVM) to model the dependencies between visible variables, in our case $\hat{\mathbf{y}}$, based on hierarchical invisible latent variables [61], in our case $\hat{\mathbf{z}}$. By introducing a set of hidden variables, i.e., $\hat{\mathbf{z}}$, and the target random variable, i.e., $\hat{\mathbf{y}}$, probabilities will be conditionally independent by definition [61]. This is a crucial improvement toward simplifying the modeling complexity and not sacrificing modeling accuracy at the same time. Since, in learned image compression, we are interested in modeling a shared prior to be used both on the encoder and decoder side, this model is called a hyperprior in the literature [6], [16], [23], [62], [63]. Thus, if we denote the hyperprior by $\hat{\mathbf{z}}$, the shared prior distribution between the encoder and decoder can be written as

$$P_{\hat{\mathbf{y}}}(\hat{\mathbf{y}}) = \prod_i P_{\hat{y}_i}(\hat{y}_i | \hat{\mathbf{z}}). \quad (8)$$

Equation (8) explicitly models the multivariate random variable by the conditional independence on the hyperprior $\hat{\mathbf{z}}$.

Backward Adaptation: Although ideally LVMs are able to capture all dependencies in the dimensions of a random variable, the practical issues of training them and the variational and amortization gaps hinder them from

performing on par with their autoregressive counterparts. The variational gap is the mismatch between the assumed variational density and the true distribution of the latent code. The amortization gap refers to the assumption that the posterior is calculated with only a single input to the encoder network. To address this issue, context should be introduced to the LVM. In addition, to prevent the infeasibility of a global context model (autoregressive model) the amount of context will be enforced on neighboring elements close to the dimension whose probability is being modeled. In the area of computer vision, masked convolutions are the de facto choice to model the local context in a causal manner [4], [64].

Joint FA and BA: To take advantage of both LVM, which is an implementation of FA, and autoregressive entropy models, which implements the BA modeling [4], [5], [23], [65], we define the conditional probability of the latent code as

$$P_{\hat{y}|\hat{z}}(\hat{y}|\hat{z}) = \prod_i P(\hat{y}_i | \hat{y}_{j < i}, \hat{z}; \theta_h). \quad (9)$$

Conditioning on the quantized hyperprior, i.e., \hat{z} , as side information is an example of FA and conditioning on all previously decoded elements of the latent code, i.e., $\hat{y}_{j < i}$, is an example of BA. Spatially autoregressive models have slow decoding time [4], [65] since the decoding time complexity increases quadratically with the spatial dimensions of the latent code. In contrast to spatial autoregressive modeling, Minnen and Singh [5] only considered the conditioning of the probabilities on the channels. With this channelwise autoregressive modeling, the decoding time is only a function of the number of slices assumed in the entropy model thus the spatial dimension of the input image will not affect the decoding latency of our neural image compression network. We have used the same approach as in [5], as shown in Fig. 2, to estimate the entropy and minimize it during training.

3) Relaxed Quantization

The gradients of uniform scalar quantization are either zero or infinity. Thus, this operation is required to be replaced with an approximation that provides informative gradients. These useful gradients will provide the ability to update the parameters of analysis and synthesis transforms using backpropagation [17], [22], [66], [67].

This approximation can be shown by assuming the simplest form of scalar quantization, i.e., rounding to the nearest integer. If an element of a latent code and its quantized version are denoted by y and \hat{y} , respectively, then we have the following:

$$y \sim p_y(y) \xrightarrow{\hat{y}=\lfloor y \rfloor} \hat{y} \sim P_{\hat{y}}(\hat{y}). \quad (10)$$

In this setting, y is a continuous random variable. Although the quantized latent code \hat{y} is a discrete random variable, its generalized probability density function can be written as a train of Dirac delta functions with weights $P_{\hat{y}}(\hat{y} = n)$ at

integer- values of n ($n \in \mathbb{Z}$)

$$p_{\hat{y}}(\hat{y}) = \sum_{n=-\infty}^{+\infty} P_{\hat{y}}(n) \delta(\hat{y} - n) \quad (11)$$

where the weights $P_{\hat{y}}(n)$ define the probability mass function for the discrete random variable \hat{y} .

For every integer-valued \hat{y} , its probability after being quantized with uniform scalar quantization will be

$$\begin{aligned} P_{\hat{y}}(\hat{y} = n) &= P_y\left(n - \frac{1}{2} < y < n + \frac{1}{2}\right) \\ &= \int_{n-\frac{1}{2}}^{n+\frac{1}{2}} p_y(\alpha) d\alpha \\ &= \int_{-\infty}^{+\infty} p_y(\alpha) \text{rect}(n - \alpha) d\alpha \\ &= (p_y * \text{rect})(n) \end{aligned} \quad (12)$$

where $*$ denotes the convolution operation. Convolution of probability density functions for two independent random variables implies the summation of those random variables.

From (12), we can see that if a unit uniform noise with mean zero, i.e., $\mathcal{U}(-\frac{1}{2}, +\frac{1}{2})$, is added to the unquantized latent representation, it will have the same density value at integer points, which are the actual quantized values. Therefore, adding independent zero mean unit uniform noise will act as continuous approximation to the hard uniform scalar quantization.

As a result, we have the relaxed quantized latent code \tilde{y} as

$$\tilde{y} = y + w \quad (13)$$

where $w \sim \mathcal{U}(-\frac{1}{2}, +\frac{1}{2})$.

We emphasize again that only at integer-valued points (quantized values) the value of the probability density of relaxed latent \tilde{y} will be equal to the probability mass of the actual quantized value \hat{y}

$$P_{\hat{y}}(\hat{y} = n) = p_{\tilde{y}}(\tilde{y} = n). \quad (14)$$

B. ATTENTION ASSISTED IMAGE COMPRESSION

The attention mechanism in neural networks (discussed in Section II-C) has been also employed in deep neural compression networks. Zhou et al. [68] applied residual attention, then Chen et al. [69] improved their work by adding a nonlocal attention mechanism to the mask of the residual attention. As a further improvement, Zou et al. [70] applied nonlocal attention limited to small windows of the feature maps. This window-based attention attained better results in compression. Here, we propose to use two kinds of attention mechanisms in a window-based manner.

Solar images have a great amount of spatial redundancy compared with natural scene images. In this view, discarding the redundancy and only keeping the low-frequency content, which is desired in natural image compression, could lead to high distortions unless paying the cost of transmitting high-frequency contents as well. Another important issue when it comes to the compression of solar images

is that the minute details and high-frequency components are important for the analysis of data [for example solar flare detection and coronal hole (CH) segmentation], while in general image compression, these high-frequency details could be discarded without intolerable cost. To address these differences in the compression domain, we propose two separate attention mechanisms. First, to apply a window-based nonlocal attention and second to refine features over a local window in the channel dimension enriching the latent code of the image.

1) Window-Based Nonlocal Attention Module (WNLAM)

To clarify the procedure of the window-based nonlocal attention mechanism, described in Section II-C, a concise review of how this method enriches the representations learned by the CNNs is included in this section. A nonlocal attention block, as shown in Fig. 3(b), is composed of a weighted average, denoted by \mathbf{q} , over a linear transformed version of the block input \mathbf{p} , i.e., $g(\mathbf{p})$

$$\mathbf{q}_i = \frac{1}{\sum_j e^{\theta(\mathbf{p}_i)^T \phi(\mathbf{p}_j)}} \sum_k e^{\theta(\mathbf{p}_i)^T \phi(\mathbf{p}_k)} g(\mathbf{p}_k) \quad (15)$$

where $g(\cdot)$ is a linear transformation, with learnable parameters W_g implemented by a 1×1 convolution layer defined as $g(\mathbf{p}_k) = W_g \mathbf{p}_k$. The weights of the sum in (15) are calculated by the measure of similarity in the embedding space of the input, i.e., $\theta(\mathbf{p}_i) = W_\theta \mathbf{p}_i$ and $\phi(\mathbf{p}_k) = W_\phi \mathbf{p}_k$, where W_θ and W_ϕ are learnable parameters.

As the final operation in nonlocal attention, \mathbf{r}_i is calculated by a linear transformation (W_r) added to the original \mathbf{p}_i as follows:

$$\mathbf{r}_i = W_r \mathbf{q}_i + \mathbf{p}_i. \quad (16)$$

Applying a nonlocal attention mechanism locally through nonoverlapping windows has shown to be more effective in the task of image compression [70] than its global counterpart [7]. In high-bit-rate image compression, restoring edges and high-frequency content are as important as representing the global features in the latent representation [2], [7], [23]. Consequently, a naïve nonlocal attention mechanism can perform worse than local attentions, which are able to capture local redundancies and preserve details on the reconstructed image [70].

2) Window-Based Convolutional Block Attention Module (WCBAM)

A simple to implement attention mechanism in CNNs is the convolutional block attention module (CBAM), which has shown great benefit in classification tasks [50]. It includes two attention mechanisms. First, the channel attention (CA) guides the network to only consider channels with higher importance for the desired task. Second, the spatial attention (SA) dictates where the network should pay more attention. Here, we propose to utilize this attention module in a window-based manner. Instead of globally considering

the whole spatial extent of each channel, we focus only on a cropped window size of w , as shown in Fig. 3(c).

Applying the WCBAM mechanism on the input features X can be summarized as

$$\begin{aligned} X_{CA} &= CA_w \odot X \\ X_{CA,SA} &= SA \odot X_{CA} \end{aligned} \quad (17)$$

where \odot is the Hadamard product. CA_w reweighs the channels over each window. After refining the channels, the SA enforces each of the refined channels (X_{CA}) to highlight their important spatial content for the task of image compression by the Hadamard product with SA.

The window-based CA is calculated by passing the average and max pool through a shared fully connected network (F), as in the following:

$$CA_w = \text{sigmoid}(F(\text{Avg}(X_w)) + F(\text{Max}(X_w))) \quad (18)$$

where X_w is a chosen window over the input feature map X . Next, the SA weights (SA) will be derived by concatenating the average and max pool passed through a convolutional layer as

$$SA = \text{sigmoid}(\text{Conv}([\text{Avg}(X_{CA}), \text{Max}(X_{CA})])) \quad (19)$$

WCBAM helps the network to capture global dependencies by looking over all channels of each chosen window simultaneously and highlighting the spatially important features with a global average/max pooled feature. These global features are needed in transforming the image from pixel space to feature space.

3) Transformers as Attention Modules

The superiority of models based on transformers, which are a special kind of nonlocal attention mechanism, compared with CNNs has been recently proven [48], [73]. Although transformers have shown great benefit in image classification and object detection tasks, their naïve application in image compression networks has failed [70]. The goal of transformers is to capture long-range dependencies in an image as opposed to convolutional-based neural networks, which inherently have a local inductive bias due to the use of a local kernel. On the other hand, the ultimate goal of image compression is to capture both local and global dependencies in order to summarize them efficiently in the latent code. If the latent code includes global information, we can expect a more compact representation. By naïvely applying the transformer blocks in the neural compression networks, it was shown [8], [70] that optimizing the rate-distortion loss leads to a local receptive field, hindering the self-attention from global dependency modeling. Therefore, enriching the attention mechanism in the CNNs, as we show in this work, could lead to better performance in terms of rate-distortion.

IV. EXPERIMENTS

A. Dataset

The dataset of SDO images described in [74] includes images of the Sun at wavelengths of 94, 131, 171, 193, 211, 304, 335, 1600, and 1700 Å at a cadence of 6 min.




Fig. 4. Rate-distortion curves averaged over the test set described in Section IV-A. On the left-hand side, PSNR is calculated from mse using $10 \log_{10} \frac{255^2}{\text{MSE}}$. On the right-hand side, MS-SSIM is reported in logarithmic scale by $-10 \log(1 - m)$ to show the differences better, in which m is the MS-SSIM in the range of zero to one.

We temporally downsampled the images to a cadence of 1 h to decrease dependencies between training samples. In addition, to prevent biases of the images with respect to solar variations at different stages of the solar cycle, we followed the same approach proposed by Salvatelli et al. [75] to divide the dataset based on the month they are taken. Images from January to August of years 2015–2018 are chosen for training and September–December of the same years are reserved for testing. The total number of training images is 21 416 and the test set includes 8257 samples. The results reported in this section are all based on the test portion of the dataset.

B. Implementation Details

As the nonlinearity in our neural network, we have utilized a computationally efficient [76] version of generalized divisive normalization (GDN) [77]. As a result of GDN’s local normalization, statistical dependencies are reduced in the feature maps. By exploiting GDN instead of more conventional nonlinearities, such as rectified linear unit, the statistical dependencies in the feature maps will be reduced significantly [77].

During the evaluation phase, entropy coding of the latent integer values was realized by range asymmetric numeral systems coding [78]. It is worth mentioning that the entropy coding is lossless and doing it during the training phase has no impact on the measured performance or functionality of the algorithm. It is only during the evaluation phase that entropy coding is needed since the performance of algorithms is compared with standard codecs, such as JPEG [2] and JPEG-2000 [3]. Seven models have been trained with empirically chosen hyperparameter $\lambda \in \{0.0015, 0.0035, 0.0070, 0.0125, 0.0250, 0.0410, 0.0550\}$ governing the rate-distortion tradeoff as in (2) for 100 epochs. We have used the Adam [24] optimizer on batches of size 16 consisting of randomly cropped 256×256 patches from the original 512×512 images. The initial value of the learning rate is set to 10^{-4} and annealed during the training to 1.2×10^{-6} , which took about 48 h on a

machine with a single NVIDIA RTX 6000 A graphic card, for each model. In addition, to compare our proposed neural codec with state-of-the-art neural-based codecs we picked ELIC [72] neural compression and trained with exactly the same hyperparameters and training policies as described for our own networks.

Training our autoencoder based on mse and LPIPS will result in outperforming even the state-of-the-art hand-engineered codec, i.e., BPG [71], as shown in Fig. 4. As can be seen in this figure, augmenting the WNLAM attention with WCBAM is capable of improving the rate-distortion performance in terms of both distortion measures PSNR and MS-SSIM. Although the adversarially trained network with the GAN has a distortion performance almost the same as JPEG-2000 [3], the general lower performance of our GAN network is a common issue addressed in [58]. The PSNR or MS-SSIM are unable to capture the perceptual quality of the generated image in a GAN. The perceptual quality of GAN network reconstructions is discussed in Section IV-C, where the quality is measured by the perceptual metric LPIPS.

To evaluate our model’s computational complexity compared with other compression algorithms, we have conducted experiments to measure the encoding and decoding latencies, as reported in Table I. All hand-crafted codecs, including JPEG, JPEG2000, and BPG, are evaluated on a system powered by Intel Core i7-5930K Broadwell-E CPU. Our proposed method and other state-of-the-art neural compression algorithm (i.e., ELIC) are tested on a single NVIDIA RTX 6000A graphic card.

C. Ablation Study

To investigate how much the attention modules contribute to the performance of our neural compressor, we have trained three separate networks including a network with only the WNLAM module, then augmented with the WCBAM module, and finally augmented with the WCBAM module and trained adversarially in a GAN framework. Performance for each of the seven targeted bit-rates is discussed in Section IV-B. The first architecture has only

TABLE I
Encoding/decoding Latency of the Proposed Neural-Based Codec Compared with Other Codecs in Three Different Bit-rate Regimes (Input Image Size of 4096×4096 Pixels)

Codec	$\sim 0.1 \text{ bpp}$		$\sim 0.35 \text{ bpp}$		$\sim 0.7 \text{ bpp}$	
	ENC (ms)	DEC (ms)	ENC (ms)	DEC (ms)	ENC (ms)	DEC (ms)
JPEG [2]	42	48	47	58	50	64
JPEG2000 [3]	312	187	367	221	416	249
BPG [71]	2340	2048	3032	2354	4080	2936
ELIC [72]	3671	3012	3754	3099	3818	3176
ours	3527	3321	3654	3423	3698	3455

ENC and DEC refer to encoding and decoding times, respectively, and are reported in ms.




Fig. 5. Visual comparison of proposed compression schemes (Attention only and GAN+Attention) to other standard codecs. Reported performance is in terms of bit-rate/distortion [bpp/PSNR \uparrow]. GAN outputs are visually closer to the original input despite their inferior performance in terms of PSNR. *Best viewed on screen.*

the WNLAM module (see Fig. 4) whose performance in terms of PSNR and MS-SSIM has been improved by adding the WCBAM attention module.

As emphasized in Fig. 5, the adversarially trained decoder results in better visual quality of the reconstructed image than the autoencoder only trained with the window-based nonlocal and convolutional block attention mechanisms. Conventional metrics, such as PSNR and MS-SSIM are unable to capture the higher perceptual quality of the GAN-reconstructed images. It is empirically shown [56] that LPIPS can show the merit of an adversarially trained network. LPIPS is known as a measure of similarity between

an image and its reconstruction, and it is shown that it is consistent with the human judgment of the quality of reconstructed images [56]. As shown in Fig. 6, the adversarially trained network performs better than the others, which are not trained using GANs. This figure shows that if the human judgment has priority over the PSNR/MS-SSIM, training adversarially is the best option.

D. Example of the Impact on Downstream Use Applications

In order to provide an example of the effects of the proposed compression scheme on data science applications,




Fig. 6. Rate-distortion curve. Distortion is measured by the LPIPS metric (lower is better) as described in Section III-A1. As shown, GAN performance in generating high-quality images can be quantified by this metric.

the CH detection and segmentation scheme outlined in [12], which is an extension of the active contours without edges (ACWE) algorithm of [79], was applied to four 193 Å Solar extreme-ultraviolet (EUV) images. The effects of compression were determined by comparing the similarity of the resulting CH regions at various bit-rates to the CH regions identified on the original Level 1 193 Å EUV images.

CHs are regions of low-density, low-temperature plasma within the Sun's corona. These regions are associated with open magnetic field lines, and are sources of high-speed solar wind [80], [81], [82], [83]. For this reason, accurate delineation of CH boundaries and extents is important for accurate space weather modeling and prediction [82], [84], [85]. The segmentation scheme outlined in [12] combines threshold-based detection, which is a common method of detecting CHs [86], with a region refinement algorithm that further refines the initial segmentation based on the homogeneity of the region. By using the algorithm of [12], we can verify that region intensity and underlying structure necessary for accurate delineation of CHs are both preserved in the compressed images.

ACWE is an iterative process wherein an initial contour or "seed" is manipulated on a pixel-by-pixel basis across multiple iterations in order to minimize an energy functional. The ACWE energy functional is minimized by balancing three forces, which seek to the following:

- 1) minimize the length of the contour,
- 2) maximize the homogeneity of the foreground (or CH region), and
- 3) maximize the homogeneity of the background or non-CH region [79].

Each force is subject to a user-defined weight which defines the relative importance of achieving each goal [79]. Unlike the process outlined in [12], ACWE was performed on the images at the original resolution of 4096×4096 pixels. The rest of the process follows the method outlined in [12], wherein the images are corrected for limb brightening following the method outlined in [87]. An

initial seed is then generated by selecting all on-disk pixels with an intensity $\leq \alpha \times QS$ where QS is the mean intensity of the quiet Sun, and the seeding parameter (α) is defined as $\alpha = 0.3$. From there, ACWE was performed on the on-disk region using the length constraint $\mu = 0$, and the ratio of foreground to background homogeneity parameters $\lambda_i/\lambda_o = 50$. The compression process was evaluated using four 193 Å Solar EUV images with record times (as expressed by the T_REC keyword) 2010-07-25 T12:00:02Z, 2010-08-21 T22:00:02Z, 2010-09-18 T00:00:02Z, and 2013-02-21 T23:00:01Z. In order to prepare the images for the compression process, the original Level 1 EUV image (at their native spatial resolution of 4096×4096 pixels) were clipped to the intensity range of $[20, 2500]$. Once clipped, a \log_{10} transform was applied to the intensity levels within images. The resulting intensities were then mapped to 255 discrete levels before performing the compression using the proposed network (with both WNLAM and WCBAM attention). For this process, eight models were developed using the hyperparameters $\lambda = \{0.0005, 0.0035, 0.007, 0.009, 0.011, 0.0125, 0.019, 0.025\}$. Once a compressed image for each EUV input was generated from each model, the images were then restored by reversing the intensity mapping and reversing the \log_{10} transform prior to performing ACWE.

An example of the effects of the proposed compression scheme on the final CH segmentation is presented in Fig. 7. In this figure the output of compressed images by five different compression algorithms are compared, by showing the original image and its correct segmentation on the first row for four different images of the Sun. The subsequent rows are the highest compression rate achievable by the mentioned algorithm and as it can be seen the proposed neural compression method can still deliver in the sub-0.01 bpp regime. Each image in Fig. 7 is a comparison between the two segmentations (original and compressed) wherein purple regions were only identified as CHs in the original segmentation, yellow regions were identified as belonging to a CH only in the compressed image segmentation, green regions were identified as CHs in both segmentations, and blue regions were not identified as CHs in either segmentation. Within this image set, and the remaining cases tested, the discrepancies between segmentations are generally limited to small-scale structures, usually along the boundary of CH regions. This suggests that the proposed compression scheme is able to preserve the overall structure (and intensity) of CH regions within the solar EUV image.

Across the four images, the effects of compression were evaluated by computing the DICE coefficient, defined as

$$DSC = \frac{2|S_1 \cap S_2|}{|S_1| + |S_2|} \quad (20)$$

between the CH segmentation generated by ACWE from the original EUV image (S_1), and the segmentation generated from the compressed image (S_2), where $|\cdot|$ denotes cardinality and \cap denotes intersection. These results, which are shown as a function of bit-rate of the compressed image in Fig. 8, corroborate the observations seen in Fig. 7 by

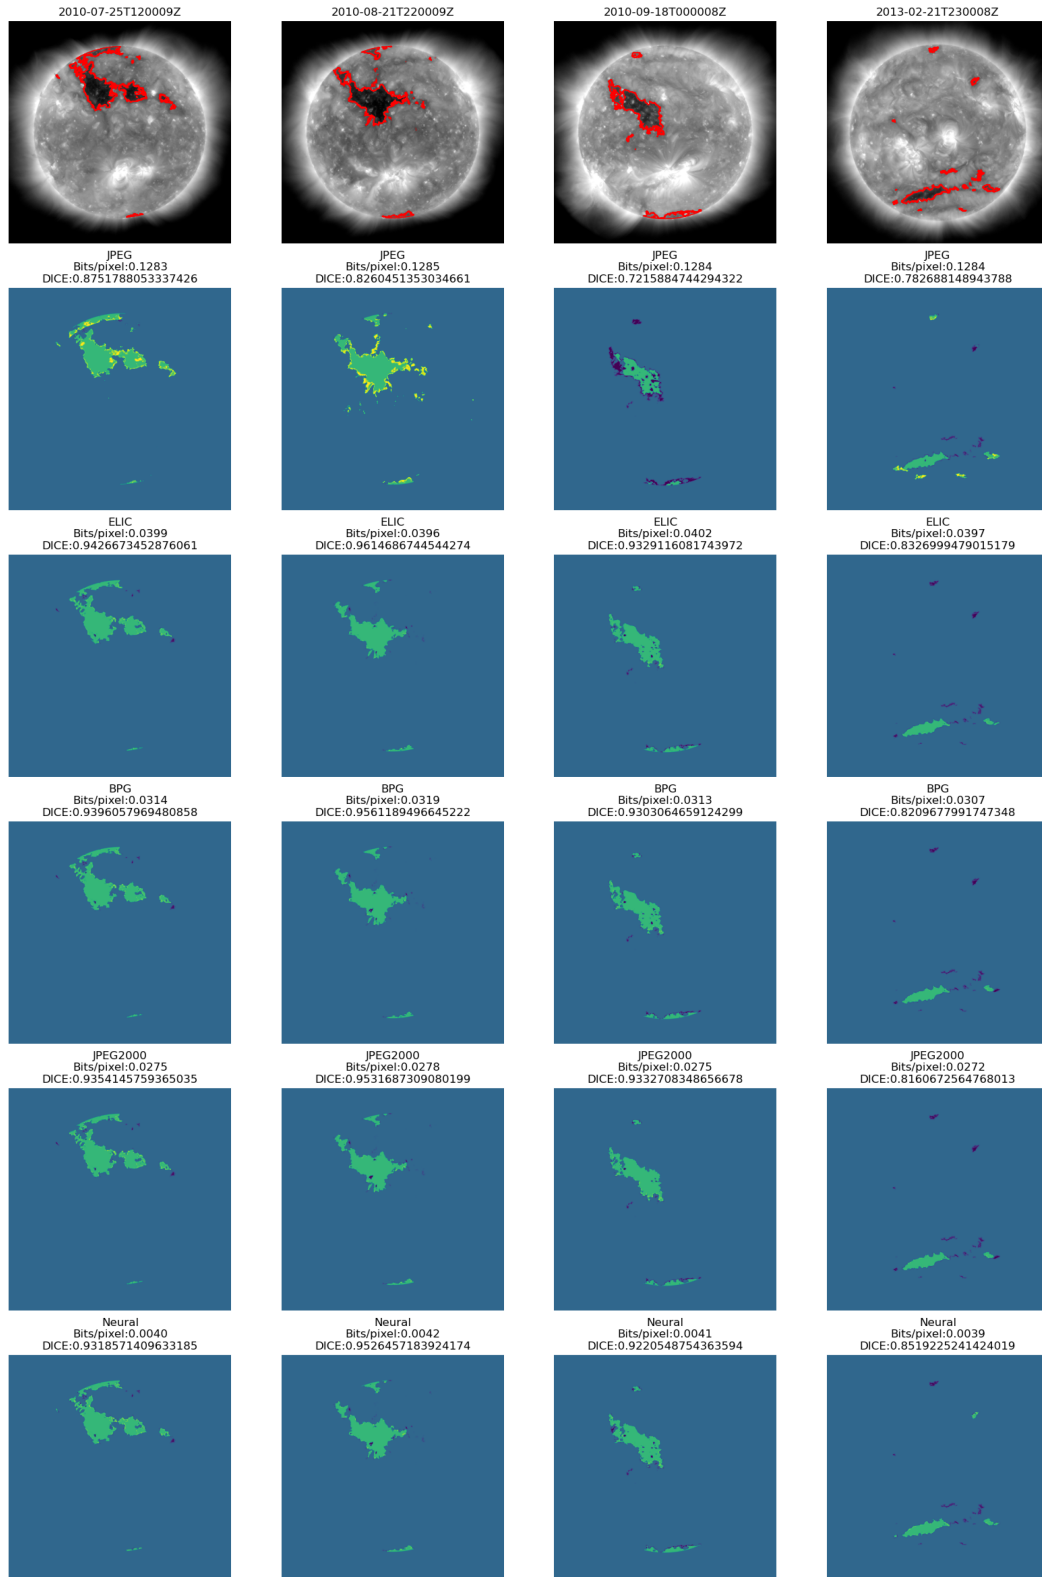


Fig. 7. Effects of the proposed compression scheme on CH detection via ACWE. The top row images are the original solar EUV images. At this wavelength CHs appear as dark regions. The CH segmentation on the original images is shown in the same row. The remaining images are comparisons of the original ACWE segmentation with the segmentation generated from the compressed images with different compression schemes. Last row contains the results of segmentation on reconstructed images through our proposed neural codec at sub ~ 0.01 bit-rates, a bit-rate regime not achievable by other codecs preserving the same DICE coefficient. In each image, purple regions were only identified as CHs in the original segmentation, yellow regions were identified as belonging to a CH only in the compressed image segmentation, green regions were identified as CHs in both segmentations, and blue regions were not identified as CHs in either segmentation.

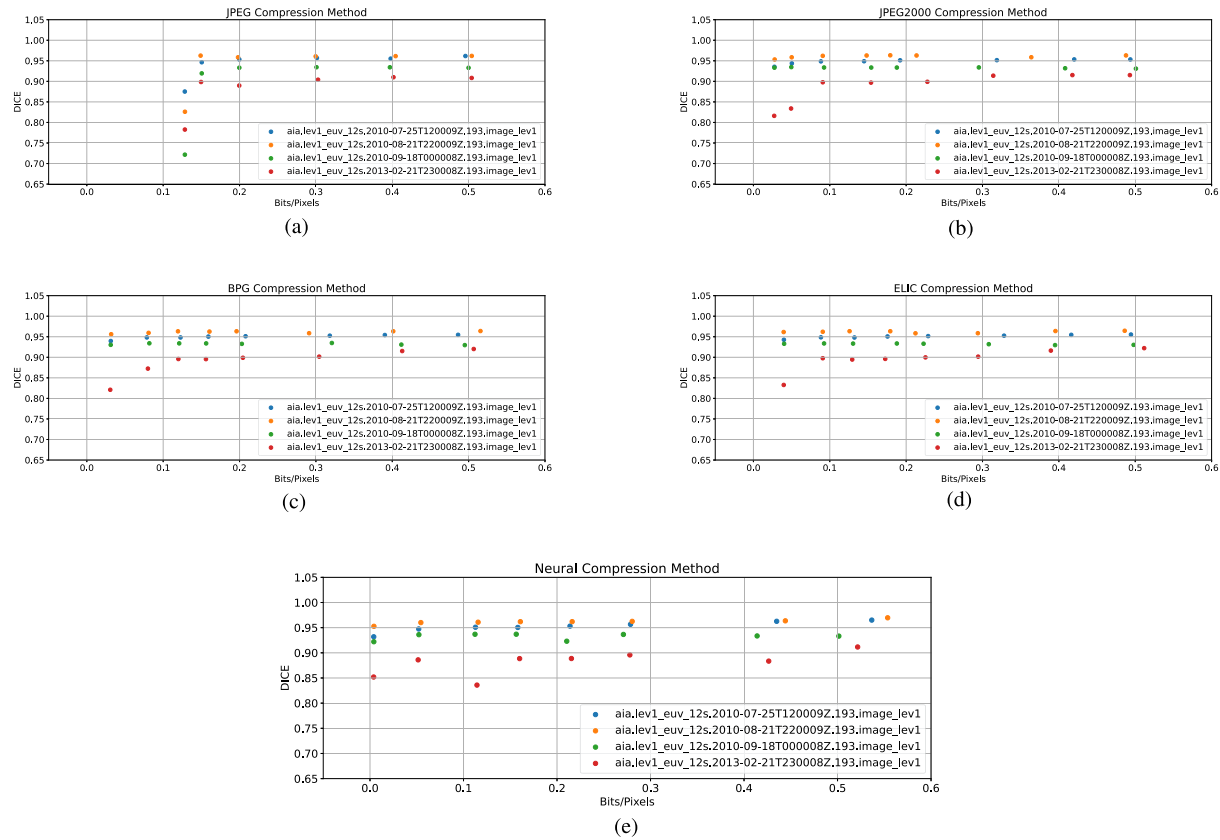


Fig. 8. Effects of the different compression schemes on CH detection via ACWE, expressed as DICE coefficient as a function of bit-rate. (a) JPEG. (b) JPEG-2000. (c) BPG. (d) ELIC. (e) Our neural codec.

showing a high similarity between segmentations across all bit-rates and more desirably in the extremely compressed regimes.

It should be noted that applying a log transformation to the intensities within an image increases the number of output intensity levels used to represent low intensities while decreasing the number of output levels used to represent high intensities. For this reason, more of the 255 discrete intensity levels were allocated to preserving fine detail at the low-intensity range of the image, which, in turn, improved the quality of the ACWE segmentation [88].

To help disentangle the effects of this advantage, this same preprocessing was applied to the same EUV images before compressing the images using other compression, then reversing the process in the same manner. The DICE coefficient (again, compared with the segmentation of the original EUV image) as a function of the bit-rate of the compressed image for the other compression schemes is presented in Fig. 8. It should be noted that the other compression schemes begin to show signs of degradation in image features that negatively impact CH segmentation at compression rates of ~ 0.12 bits per pixel, and cannot produce images with a compression rate < 0.1 bits per pixel for these 4K by 4K images, unless sacrificing DICE measure. One the other hand, the neural compression scheme proposed here continues to perform well beyond this threshold, suggesting that the neural compression scheme is able

to better preserve the features that are relevant to this application.

V. CONCLUSION

In this work, we have shown how an effective image compression scheme based on trainable neural networks could be utilized for ad hoc applications, such as images from NASA's SDO mission. We explored the effectiveness of window-based spatial and cross-CA mechanisms in an adversarially trained neural network to improve the performance of compression in terms of rate-distortion-perception tradeoff. It was shown that neural compression algorithms may be able to benefit data-intensive space missions with minimal degradation in downstream scientific tasks, such as CH segmentation as described in this work.

REFERENCES

- [1] Y. Yang, S. Mandt, and L. Theis, "An introduction to neural data compression," *Found. Trends Comput. Graph. Vis.*, vol. 15, no. 2, pp. 113–200, 2023.
- [2] G. K. Wallace, "The JPEG still picture compression standard," *Commun. ACM*, vol. 34, pp. 30–44, 1991.
- [3] D. S. Taubman and M. W. Marcellin, *Marcellin JPEG2000 - Image Compression Fundamentals, Standards and Practice*. Norwell, MA, USA: Kluwer, 2002.
- [4] D. Minnen, J. Ballé, and G. Toderici, "Joint autoregressive and hierarchical priors for learned image compression," in *Proc. Adv. Neural Inf. Process.*, 2018, pp. 10794–10803.

- [5] D. Minnen and S. Singh, "Channel-wise autoregressive entropy models for learned image compression," in *Proc. IEEE Int. Conf. Image Process.*, 2020, pp. 3339–3343.
- [6] Y. Qian, X. Sun, M. Lin, Z. Tan, and R. Jin, "Entroformer: A transformer-based entropy model for learned image compression," in *Proc. Int. Conf. Learn. Representations*, May 2022.
- [7] Z. Cheng, H. Sun, M. Takeuchi, and J. Katto, "Learned image compression with discretized Gaussian mixture likelihoods and attention modules," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7936–7945.
- [8] Y. Zhu, Y. Yang, and T. Cohen, "Transformer-based transform coding," in *Proc. Int. Conf. Learn. Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=IDwN6xjHnK8>
- [9] E. Agustsson, M. Tschannen, F. Mentzer, R. Timofte, and L. V. Gool, "Generative adversarial networks for extreme learned image compression," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 221–231.
- [10] Y. Blau and T. Michaeli, "Rethinking lossy compression: The rate-distortion-perception tradeoff," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2019, pp. 675–685.
- [11] F. Mentzer, G. D. Toderici, M. Tschannen, and E. Agustsson, "High-fidelity generative image compression," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, Art. no. 999.
- [12] L. Boucheron, M. Valluri, and R. McAteer, "Segmentation of coronal holes using active contours without edges," *Sol. Phys.*, vol. 291, no. 8, pp. 2353–2372, 2016.
- [13] A. Zafari et al., "Attention-based generative neural image compression on solar dynamics observatory," in *Proc. 21st IEEE Int. Conf. Mach. Learn. Appl.*, 2022, pp. 198–205.
- [14] V. Goyal, "Theoretical foundations of transform coding," *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 9–21, Sep. 2001.
- [15] J. Ballé, V. Laparra, and E. P. Simoncelli, "End-to-end optimized image compression," in *Proc. Int. Conf. Learn. Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=rJxdQ3jeg>
- [16] J. Ballé et al., "Nonlinear transform coding," *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 2, pp. 339–353, Feb. 2021.
- [17] J. Ballé, V. Laparra, and E. P. Simoncelli, "End-to-end optimization of nonlinear transform codes for perceptual quality," in *Proc. Picture Coding Symp.*, 2016, pp. 1–5.
- [18] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.
- [19] L. Theis, A. van den Oord, and M. Bethge, "A note on the evaluation of generative models," in *Proc. Int. Conf. Learn. Representations*, 2016, *arXiv:1511.01844*.
- [20] T. M. Cover, *Cover Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1999.
- [21] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, vol. 159. Berlin, Germany: Springer, 2012.
- [22] E. Agustsson et al., "Soft-to-hard vector quantization for end-to-end learning compressible representations," in *Proc. Adv. Neural Inf. Process.*, 2017, pp. 1141–1151.
- [23] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, "Variational image compression with a scale hyperprior," in *Proc. 6th Int. Conf. Learn. Representations*, 2018, *arXiv:1802.01436*.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015, *arXiv:1412.6980*.
- [25] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [26] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *Proc. 7th Int. Conf. Learn. Representations*, 2019, *arXiv:1809.11096*. [Online]. Available: <https://openreview.net/forum?id=B1xsqj09Fm>
- [27] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 12873–12883.
- [28] Y. Poirier-Ginter and J.-F. Lalonde, "Robust unsupervised StyleGAN image restoration," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 22292–22301.
- [29] H. Ling, K. Kreis, D. Li, S. W. Kim, A. Torralba, and S. Fidler, "EditGAN: High-precision semantic image editing," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 16331–16345.
- [30] C. Ledig et al., "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4681–4690.
- [31] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4401–4410.
- [32] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8110–8119.
- [33] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1125–1134.
- [34] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2223–2232.
- [35] S. Jozdani, D. Chen, D. Pouliot, and B. Alan Johnson, "A review and meta-analysis of generative adversarial networks and their applications in remote sensing," *Int. J. Appl. Earth Obs. Geoinf.*, vol. 108, 2022, Art. no. 102734. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0303243422000605>
- [36] N. Toner and J. Fletcher, "Self-attending task generative adversarial network for realistic satellite image creation," in *Proc. IEEE Aerosp. Conf.*, 2022, pp. 1–9.
- [37] L. Hu, W. Wang, Y. Xiang, and J. Zhang, "Flow field reconstructions with GANs based on radial basis functions," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 58, no. 4, pp. 3460–3476, Aug. 2022.
- [38] P. Liu, J. Li, L. Wang, and G. He, "Remote sensing data fusion with generative adversarial networks: State-of-the-art methods and future research directions," *IEEE Geosci. Remote Sens. Mag.*, vol. 10, no. 2, pp. 295–328, Jun. 2022.
- [39] M. M. Rahman, S. Z. Gurbuz, and M. G. Amin, "Physics-aware generative adversarial networks for radar-based human activity recognition," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 3, pp. 2994–3008, Jun. 2023.
- [40] Z. Wang, M. Chen, Y. Guo, Z. Li, and Q. Yu, "Bridging the domain gap in satellite pose estimation: A self-training approach based on geometrical constraints," *IEEE Trans. Aerosp. Electron. Syst.*, early access, Mar. 1, 2023, doi: [10.1109/TAES.2023.3250385](https://doi.org/10.1109/TAES.2023.3250385).
- [41] D. Li et al., "Involution: Inverting the inheritance of convolution for visual recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 12321–12330.
- [42] J. Zhou, V. Jampani, Z. Pi, Q. Liu, and M.-H. Yang, "Decoupled dynamic filter networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 6647–6656.
- [43] P. Ramachandran, N. Parmarand, A. Vaswani, I. Bello, A. Levskaya, and J. Shlens, "Stand-alone self-attention in vision models," in *Proc. Adv. Neural Inf. Process.*, 2019, Art. no. 7.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [45] F. Wang et al., "Residual attention network for image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3156–3164.
- [46] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7794–7803.
- [47] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2005, pp. 60–65.
- [48] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. 9th Int. Conf. Learn. Representations*, 2021.
- [49] Y. Zhang, K. Li, K. Li, B. Zhong, and Y. Fu, "Residual non-local attention networks for image restoration," in *Proc. Int. Conf. Learn. Representations*, 2019.

- [50] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 3–19.
- [51] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7132–7141.
- [52] T. Xiao, Y. Xu, K. Yang, J. Zhang, Y. Peng, and Z. Zhang, "The application of two-level attention models in deep convolutional neural network for fine-grained image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 842–850.
- [53] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 47–57, Mar. 2017.
- [54] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [55] J. Nilsson and T. Akenine-Möller, "Understanding SSIM," 2020, *arXiv:2006.13846*.
- [56] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 586–595.
- [57] F. Mentzer, G. D. Toderici, M. Tschannen, and E. Agustsson, "High-fidelity generative image compression," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 11913–11924.
- [58] Y. Blau and T. Michaeli, "The perception-distortion tradeoff," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 6228–6237.
- [59] L. Theis, W. Shi, A. Cunningham, and F. Huszár, "Lossy image compression with compressive autoencoders," in *Proc. Int. Conf. Learn. Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=rJiNwv9gg>
- [60] C. M. Bishop, *Pattern Recognition and Machine Learning*, vol. 4. Berlin, Germany: Springer, 2006.
- [61] C. M. Bishop, "Latent variable models," in *Learning in graphical models*, vol. 371, 1999, pp. 371–403.
- [62] Y. Qian et al., "Learning accurate entropy model with global reference for image compression," in *Proc. Int. Conf. Learn. Representations*, 2021, *arXiv:2010.08321*.
- [63] J.-H. Kim, B. Heo, and J.-S. Lee, "Joint global and local hierarchical priors for learned image compression," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 5982–5991.
- [64] A. Van den Oord et al., "Conditional image generation with pixelCNN decoders," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 4797–4805.
- [65] J. Lee, S. Cho, and S.-K. Beack, "Context-adaptive entropy model for end-to-end optimized image compression," in *Proc. Int. Conf. Learn. Representations*, 2019, .
- [66] Y. Yang, R. Bamler, and S. Mandt, "Variational Bayesian quantization," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2020, pp. 10670–10680.
- [67] Z. Guo, Z. Zhang, R. Feng, and Z. Chen, "Soft then hard: Rethinking the quantization in neural image compression," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2021, pp. 3920–3929.
- [68] L. Zhou, Z. Sun, X. Wu, and J. Wu, "End-to-end optimized image compression with attention mechanism," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2019.
- [69] T. Chen, H. Liu, Z. Ma, Q. Shen, X. Cao, and Y. Wang, "End-to-end learnt image compression via non-local attention optimization and improved context modeling," *IEEE Trans. Image Process.*, vol. 30, no. 1, pp. 3179–3191, Feb. 2021.
- [70] R. Zou, C. Song, and Z. Zhang, "The devil is in the details: Window-based attention for image compression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 17471–17480.
- [71] F. Bellard, "BPG image format," 2018. [Online]. Available: <https://bellard.org/bpg/>
- [72] D. He, Z. Yang, W. Peng, R. Ma, H. Qin, and Y. Wang, "ELIC: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 5718–5727.
- [73] Z. Liu et al., "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9992–10002.
- [74] R. Galvez et al., "A machine-learning data set prepared from the NASA solar dynamics observatory mission," *Astrophysical J. Suppl. Ser.*, vol. 242, no. 1, p. 7, 2019.
- [75] V. Salvatelli et al., "Using U-Nets to create high-fidelity virtual observations of the solar corona," 2019, *arXiv:1911.04006*.
- [76] N. Johnston, E. Eban, A. Gordon, and J. Ballé, "Computationally efficient neural image compression," 2019, *arXiv:1912.08771*.
- [77] J. Ballé, V. Laparra, and E. P. Simoncelli, "Density modeling of images using a generalized normalization transformation," in *Proc. 4th Int. Conf. Learn. Representations*, Jan. 1, 2016.
- [78] J. Duda, "Asymmetric numeral systems: Entropy coding combining speed of Huffman coding with compression rate of arithmetic coding," 2013, *arXiv:1311.2540*.
- [79] T. Chan and L. Vese, "Active contours without edges," *IEEE Trans. Image Process.*, vol. 10, no. 2, pp. 266–277, Feb. 2001.
- [80] M. D. Altschuler, D. E. Trotter, and F. Q. Orrall, "Coronal holes," *Sol. Phys.*, vol. 26, no. 2, pp. 354–365, Oct. 1972.
- [81] R. H. Munro and G. L. Withbroe, "Properties of a coronal 'hole' derived from extreme-ultraviolet observations," *Astrophysical J.*, vol. 176, Sep. 1972, Art. no. 511.
- [82] Y. M. Wang and N. R. Sheeley, "Solar wind speed and coronal flux-tube expansion," *Astrophysical J.*, vol. 355, pp. 726–732, Jun. 1990.
- [83] Y.-M. Wang, S. H. Hawley, J. Sheeley, and R. Neil, "The magnetic nature of coronal holes," *Science*, vol. 271, no. 5248, pp. 464–469, Jan. 1996.
- [84] C. N. Arge, D. Odstrcil, V. J. Pizzo, and L. R. Mayer, "Improved method for specifying solar wind speed near the Sun," in *Proc. Amer. Inst. Phys. Conf. Ser.*, 2003, pp. 190–193.
- [85] C. N. Arge, J. G. Luhmann, D. Odstrcil, C. J. Schrijver, and Y. Li, "Stream structure and coronal sources of the solar wind during the May 12th, 1997 CME," *J. Atmospheric Sol.- Terr. Phys.*, vol. 66, no. 15/16, pp. 1295–1309, Oct. 2004.
- [86] M. A. Reiss et al., "The observational uncertainty of coronal hole boundaries in automated detection schemes," *Astrophysical J.*, vol. 913, no. 1, May 2021, Art. no. 28. [Online]. Available: <http://dx.doi.org/10.3847/1538-4357/abf2c8>
- [87] C. Verbeeck, V. Delouille, B. Mampaey, and R. De Visscher, "The SPoCA-suite: Software for extraction, characterization, and tracking of active regions and coronal holes on EUV images," *Astron. Astrophys.*, vol. 561, 2014, Art. no. A29. [Online]. Available: <https://doi.org/10.1051/0004-6361/201321243>
- [88] J. A. Grajeda, L. E. Boucheron, M. S. Kirk, A. Leisner, and C. N. Arge, "Quantifying the consistency and characterizing the confidence of coronal holes detected by active contours without edges (ACWE)," *Sol. Phys.*, vol. 298, no. 11, 2023, Art. no. 133.



Ali Zafari (Graduate Student Member, IEEE) received the B.Sc. degree in electrical engineering from the University of Tehran, Tehran, Iran, in 2016, and the M.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, in 2020.

He is currently a Graduate Research Assistant with the West Virginia University's Lane Department of Computer Science and Electrical Engineering, Morgantown, WV, USA, where he is engaged in research related to information theory and computer vision. His research focuses on neural image/video compression, which is involved with data compression using deep probabilistic models.



Atefeh Khoshkhahtinat received the B.Sc. degree in electrical engineering from the K. N. Toosi University of Technology, Tehran, Iran, in 2016, and the M.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in 2020.

She is currently a Graduate Research Assistant with the West Virginia University's Lane Department of Computer Science and Electrical Engineering, Morgantown, WV, USA. Her research interests include deep learning, machine

learning, and computer vision.



Jeremy A. Grajeda received the B.S. and M.S. degrees in electrical engineering from New Mexico State University, Las Cruces, NM, USA, in 2019 and 2022, respectively. They are currently working toward the doctorate degree in developing image processing techniques for coronal hole segmentation with the Klipsch School of Electrical and Computer Engineering, New Mexico State University, Las Cruces, NM, USA.

In 2021, they interned with the Solar Physics Laboratory, NASA's Goddard Space Flight Center, Greenbelt, MD, USA. Their research interests include digital image

processing, computer vision, machine learning, space weather forecasting, improving solar flare predictions via machine learning, and improving coronal hole identification and segmentation using classical techniques.



Piyush M. Mehta received the B.S. and Ph.D. degrees in aerospace engineering from the University of Kansas, Lawrence, KS, USA, in 2009 and 2013, respectively.

He is currently an Assistant Professor of space systems and a J. Wayne and Kathy Richards Faculty Fellow with the Department of Mechanical and Aerospace Engineering, West Virginia University, Morgantown, WV, USA. His research interests include the intersection of astrodynamics, space science and technology, and data science,

investigating the impact of space weather on technological assets in space and on the ground, and broader application of machine learning techniques to challenging problems across domains including unsteady aerodynamics and data compression.

Dr. Mehta is a Member of the NASA Space Weather Council.



Nasser M. Nasrabadi (Fellow, IEEE) received the B.Sc. (Eng.) and Ph.D. degrees in electrical engineering from the Imperial College of Science and Technology, University of London, London, U.K., in 1980 and 1984, respectively.

In 1984, he was with IBM, U.K., as a Senior Programmer. From 1985 to 1986, he was with the Philips Research Laboratory, New York, NY, USA, as a Member of the Technical Staff. From 1986 to 1991, he was an Assistant Professor with the Department of Electrical Engineering,

Worcester Polytechnic Institute, Worcester, MA, USA. From 1991 to 1996, he was an Associate Professor with the Department of Electrical and Computer Engineering, State University of New York at Buffalo, Buffalo, NY, USA. From 1996 to 2015, he was a Senior Research Scientist with the U.S. Army Research Laboratory. Since 2015, he has been a Professor with the Lane Department of Computer Science and Electrical Engineering. His current research interests include deep learning, computer vision, biometrics, statistical machine learning theory, and image processing.

Dr. Nasrabadi is a Fellow of the International Society for Optics and photonics (SPIE) and ARL. He was an Associate Editor for IEEE TRANSACTIONS ON IMAGE PROCESSING, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS.



Laura E. Boucheron (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from New Mexico State University, Las Cruces, NM, USA, in 2001 and 2003, respectively, and the Ph.D. degree in electrical engineering from the University of California Santa Barbara, Santa Barbara, CA, USA, in 2008.

She was a Postdoctoral Fellow (2008–2009), Research Assistant Professor (2010–2011), Assistant Professor (2011–2017), and Associate Professor (2017–2022) with the Klipsch School

of Electrical and Computer Engineering, New Mexico State University, where since 2022, she has been a Professor. She has intern experience with Sandia National Laboratories (1998–2000), Agilent Technologies (2001), and Los Alamos National Laboratories (2003 and 2005–2007). Her current research interests include image processing and machine and deep learning applied to image processing, and in interdisciplinary applications, including in the domains of biomedical and astronomy data.

Dr. Boucheron is currently an Associate Editor for IEEE TRANSACTIONS ON IMAGE PROCESSING for which she was the recipient of the Outstanding Editorial Board Member Award in 2022.



Barbara J. Thompson received the B.A. degree in physics and mathematics from the University of Pennsylvania, Philadelphia, PA, USA, in 1991, and the Ph.D. degree in physics from the University of Minnesota, Minneapolis, MN, USA, in 1996.

She has been a Civil Servant of NASA Goddard Space Flight Center since 1998, and is currently a Research Scientist with the Heliophysics Science Division. She was on numerous space mission teams, with a research emphasis

on dynamic phenomena on the Sun and the inner Heliosphere. She is a part of the leadership team for the Center for Helio Analytics, a Community of Practice to envision solutions using machine learning, knowledge capture, and data analytics to expand the discovery potential for key heliophysics research topics and missions.



Michael S. F. Kirk was born in Salem, OR. He received the B.A. degree in physics and astronomy from Whitman College, Walla Walla, WA, USA, in 2006, and the M.S. and Ph.D. degrees in astronomy from New Mexico State University, Las Cruces, NM, USA, in 2011 and 2013, respectively.

From 2013 to 2016, he was a NASA Postdoctoral Fellow with NASA's Goddard Space Flight Center, the Heliophysics Science Division, Greenbelt, MD, USA. From 2016 to 2021,

he was a Research Scientist with The Catholic University of America, Washington, DC, USA. From 2021 to 2022, he was a Principal Scientist with Atmospheric and Space Technology Research Associates, LLC, Louisville, CO, USA. In 2022, he joined NASA's Goddard Space Flight Center, as a Research Astrophysicist. In this position, he is currently the Principal Investigator of NASA's Heliophysics Education Activation Team (NASA HEAT) and a coinvestigator on the SunCET CubeSat mission, launching in 2024. In addition, he is helping to lead NASA Goddard's Center for HelioAnalytics, which seeks to integrate data science into heliophysics to better the physics of the Sun, the causes of solar variability, and its impacts on Earth.

Dr. Kirk is a Member of the American Astronomical Society, Solar Physics Division, where he is the Press Officer. He is also a Member of the American Geophysical Union, Space Physics, and Aeronomy sections.



Daniel E. da Silva received the B.Sc. degree in mathematics from the University of Maryland, College Park, MD, USA, in 2014, and the M.Sc. degree in space systems engineering from Johns Hopkins University, Baltimore, MD, in 2021. He is currently working toward the doctorate degree in designing and implementation of a new calculation tool for dynamic physically-based magneto-hydro-dynamic (MHD) magnetic field models with the Laboratory for Atmospheric and Space Physics, University of Colorado, Boulder,

BO, USA.

His research interests include applied information theory and data compression, radiation belt modeling, and probabilistic space weather forecasting.