
Neural Weight Step Video Compression

Mikolaj Czerkawski

Javier Cardona

Robert Atkinson

Craig Michie

Ivan Andonovic

Carmine Clemente

Christos Tachatzis

University of Strathclyde
99 George St, G1 1RD
Glasgow, United Kingdom
mikolaj.czerkawski@strath.ac.uk

Abstract

A variety of compression methods based on encoding images as weights of a neural network have been recently proposed. Yet, the potential of similar approaches for video compression remains unexplored. In this work, we suggest a set of experiments for testing the feasibility of compressing video using two architectural paradigms, coordinate-based MLP (CbMLP) and convolutional network. Furthermore, we propose a novel technique of neural weight stepping, where subsequent frames of a video are encoded as low-entropy parameter updates. To assess the feasibility of the considered approaches, we will test the video compression performance on several high-resolution video datasets and compare against existing conventional and neural compression techniques.

1 Introduction

Video compression is a relevant topic of ongoing research in the current stage of the digital age. The amount of shared video data continues to grow due to the increasing popularity of services such as video-on-demand, real-time communication, or distance learning [1]. To accommodate for this rapid expansion with the available bitrate budget, efforts to maximize the amount of recoverable information for every physically transmitted bit play a crucial role. This optimization process relies

35th Pre-registration workshop (NeurIPS 2021), Remote.




Figure 1: The weight step video compression technique employs a neural network to parameterize a single image of the sequence and encodes subsequent frames as low-entropy weight steps applied to the network. Only parameters in dark blue need to be transmitted.

primarily on various compression techniques. These techniques can be divided into conventional, deterministic video coding methods [2]–[5], widely employed in practice, as well as fairly new neural video compression techniques [6]–[13].

Most often, the video compression techniques based on neural networks exhibit close resemblance to the traditional pipelines, that is, they train an encoding module to produce a compressed bitstream that a decoding module can reconstruct the original message from [6]–[13]. We propose an alternative solution, where a neural network is used directly as the compressed representation of the video. Furthermore, we propose a technique, where the shared information in consecutive frames is exploited. This is done by encoding an image sequence I_0, I_1, \dots, I_N as low-entropy network weight updates $\Delta\theta_1, \dots, \Delta\theta_{N-1}$ applied to the network parameters θ_0 representing the anchor frame I_0 , as demonstrated in Figure 1.

If successful, this approach will have the considerable advantage of being data-agnostic, unlike other neural video compression techniques published to date that must be pretrained on an external dataset [6]–[13]. How the proposed method compares to the state-of-the-art video codecs [4], [5] remains to be tested in the experimental stage.

2 Related Work

Several published works provide some foundation components of the solution proposed in this work. The technique of parameterizing images using neural network weights (either coordinate- or convolution-based) has already been explored to some degree [14], [15]. Further, neural networks have already been employed for compression, albeit in fundamentally different schemes. Here, we provide a summary of these past contributions.

2.1 Neural Implicit Representations

The proposed techniques are closely related to the recent work on compression with coordinate-based MLP (CbMLP) networks [14]. There, natural images are encoded using an multilayer perceptron (MLP) network that maps input space coordinates to the appropriate pixel color values. The experiments demonstrate that COIN is **not** capable of beating the state-of-the-art compression schemes (including standardized codecs and neural compression baselines) [14]. However, the potential for applying neural implicit compression for videos has not been explored. Furthermore, the proposed parameter step technique can potentially enhance the performance compared to methods based on direct frame encoding.

Furthermore, a surge in interest pertaining to 3D representation learning was sparked by the introduction of the Neural Radiance Field (NeRF) technique [16]. Similarly to COIN, the approach heavily relies on a CbMLP, albeit with an input space of higher dimensionality; three spatial coordinates (x, y, z) and two viewing direction coordinates (θ, ϕ) . These additional dimensions are learned in the same manner as the temporal dimension of a video representation would be. This approach, in fact, will be one of the configurations tested against the weight step technique proposed in this work.

2.2 Deep Convolutional Priors

A growing body of work addresses the potential of encoding images in parameters of a convolutional neural network. This technique has been demonstrated to be useful for several image inverse problems, such as inpainting, super-resolution, or denoising [15]. Some work on extending these techniques for similar inverse tasks in video domain has been done [17], however, not in the context of image compression. The prior introduced by the use of convolutional kernels appears to have a capability of biasing the solution towards a domain of narrow texture distributions, based on the existing information in the synthesized signal. We hypothesize that this domain narrowing effect can be beneficial for finding compressed representations of the signal.

2.3 Neural Video Compression

The approach of encoding the signal into a compressed transmittable representation and then reconstructing back to the original representation [18], [6]–[8] is known as end-to-end neural video coding [1] (other variants identified in [1] include up-sampling approaches and modularized schemes

that are combined with traditional coding frameworks). We focus on the end-to-end coding variant since it is most closely related to the idea proposed in this work. Recently explored neural image compression techniques tend to rely on the training of a convolutional autoencoder network that reconstructs the output from the compressed bitstream [19]–[21]. Analogous methods have been tested in the video domain [9]–[11]. Further approaches to neural video coding take advantage of the correlations in the temporal domain [12], [13]. However, all cited neural video compression techniques require a network to be pretrained on an external dataset before use, which carries a significant computational cost and a risk of overfitting. Our approach is data-agnostic in the sense that it requires no pretraining on any external datasets and adapts to any new video.

3 Compressing Video as Parameter Steps

We wish to learn whether video compression can be achieved by relying on weight changes applied to a neural network that encodes a two-dimensional image representation of an anchor frame (for example, the first frame in the sequence). This approach could have several benefits. First, directly encoding a single anchor frame in the weights of a neural network should be less computationally demanding than encoding a three-dimensional video representation; for architectures based on convolutions the model size tends to grow exponentially with every new dimension of the convolution. If representational neural networks implicitly learn a set of visual abstractions about the image, then making small changes to the network weights may allow to reuse these abstractions and therefore, be an efficient way of encoding motion pictures.

There are at least two fundamentally different techniques for representing images as a neural network. An image can be reconstructed as i) the response of a **convolutional network** to a fixed random input vector [15], or ii) the output distribution of an **CbMLP** network to a set of coordinate inputs (one for each pixel) [22], [23].

3.1 Convolutional Network

Methods of encoding images in the weights of convolutional neural networks have been proposed in seminal works such as [15]. The authors claim that the sole use of convolutional architectures is a source of useful prior for natural images. The setting involves a convolutional network with parameters θ that maps a random fixed noise input z to a given output approximation \hat{I} with a function f . During training, the error between the approximation \hat{I} and ground truth image I is minimized.

$$f(\theta, z) = \hat{I} \tag{1}$$

3.2 Coordinate-based MLP Network

As in COIN [14], we can encode a signal in any number of output dimensions, whether it is two spatial dimensions for a single anchor frame, or a video sequence with one more dimension in the output distribution. This setting is largely based on the recent developments [22], [23] related to encoding an image as the output distribution f of a MLP that maps from coordinates (x, y) to pixel color values $\hat{c}(x, y)$ at the queried coordinate.

$$f(\theta, x, y) = \hat{c}(x, y) \tag{2}$$

For the weight step technique, we only need to encode a single two-dimensional image I containing the anchor frame. The approximated image \hat{I} is obtained by aggregating predicted color $\hat{c}(x, y)$ for all coordinates of ground-truth image I .

Furthermore, the capability of CbMLP networks to produce high-frequency detail in the output distribution has only been achieved recently by utilising techniques such as random Fourier feature encoding [23], or sinusoidal network activations [22]. In this work, we will test both approaches in the first stage of experiments.

The optimization process for both CbMLP and convolutional variants is similar, where the reconstruction error between the ground truth image I and the approximation \hat{I} resulting from $f_{(\theta)}$ is minimized.

For clarity, we describe the high-level image approximation function of each variant as $f_{(\theta)}$ from now on, abstracting away the differences such as the need for a random input vector z , or aggregation of all input coordinates (x, y) .

3.3 Weight Step Encoding

For both representation paradigms, the setting is similar; we want to first encode the parameters θ_0 that result in the generation of the first image I_0 of the video sequence. Any given neural network topology determines a function $f_{(\theta)}$ that yields a specific 2D output.

$$f_{(\theta_0)} = I_0 \tag{3}$$

Then, at each temporal step t , we wish to find a low entropy parameter shift $\Delta\theta_t$ that results in the image at I_{t+1} .

$$f_{(\theta_t + \Delta\theta_t)} = I_{t+1} \tag{4}$$

This, however, should be subject to constraints that enforce efficient encoding. Rather than computing any parameter shift resulting in the subsequent image, we want to find one such that its entropy is low. This can be done by enforcing sparsity in the weight step representation, similarly to existing techniques for network sparsity [24]–[27]. A promising approach would be to add a regularization term $\mathcal{R}(\Delta\theta_t)$ to the objective function. As analyzed in [24], regularizing the weight parameters with ℓ_0 norm is not practical, since the differentiability condition is not satisfied. To that end, we utilize the technique proposed by the same authors in [24] based on the introduction of non-binary stochastic gates and a surrogate sparsity objective. Further, we want to explore two techniques of sparsity regularization. The first technique will involve maximizing sparsity of the direct parameter update array. This will reward solutions where only a small number of network weights is updated at each step:

$$\mathcal{R}(\Delta\theta_t) = \|\Delta\theta\|_0 \tag{5}$$

Alternatively, using the second technique, sparsity can be achieved in a domain orthogonal to the spatial domain. In this case, we propose to test a technique where sparsity is enforced in the spatial frequency domain instead. This can be done by applying the same sparsification technique to the discrete cosine transform (DCT) of each layer block (for convolutional network, DCT is applied to individual sequences of aligned kernel elements), similarly to [28]. The technique allows the parameter update solution to involve changes to a large number of network weights while preserving low information content.

$$\mathcal{R}(\Delta\theta_t) = \|\text{DCT}(\Delta\theta)\|_0 \tag{6}$$

Finally, the optimization process itself can be done in at least two different ways. The network can be trained in two stages, where the first stage optimizes only the starting point θ_0 while the second stage optimizes an $(N - 1)$ -long sequence of $\Delta\theta_t$ originating from that point. Another way would be to train this sequence in a single-stage process, where the complete sequence $\left\{(\theta_0), (\theta_0 + \Delta\theta_0), \dots, (\theta_0 + \sum_{t=0}^{N-1} \Delta\theta_t)\right\}$ is to be found.

4 Experimental Protocol

We plan to test on several video datasets. Our primary focus will be on high-resolution datasets such as UVG (1080p and 4K) [29] and MCL-JVC (1080p) [30]. Furthermore, we will run additional experiments on a lower-fidelity dataset of Video Trace Library (288p) [31].

4.1 Base Architectures

Our two basic architecture types are CbMLP and a convolutional neural network. During the experiments, we will want to compare several alterations of each model corresponding to a different parameter count.

Furthermore, we will want to test how the weight step compression techniques compare to equivalent models that directly encode the entire video, with no weight stepping applied. For the CbMLP, the full video can be captured by learning one additional input dimension t for time ($f(\theta, x, y) = \hat{c}(x, y)$ becomes $f(\theta, x, y, t) = \hat{c}(x, y, t)$). For the convolutional module, one solution to approach this is to define a fixed sequence of input tensors $\{z_0, z_1, \dots, z_N\}$, one z_i for each of N video frames. Similarly, to a single z , this input activation is predefined before use and remains the same for any encoded video.

4.2 Architectural Ablations

Before evaluating the core results, we will want to run several design-related experiments for our method.

Both the SIREN approach [22] and random Fourier feature encoding [23] appear to be a valid design choice for coordinate-based MLP variant. However, it is not clear how the two approaches compare in the context of end-to-end compression. To better understand, which approach can encode signals in a more parameter-efficient manner, we compress each image in the three datasets using the two approaches with the same set of underlying layer topology.

For the convolution-based approach, we would like to test whether the application of long-distance skip connections (like in some of the configurations proposed in the Deep Image Prior work [15]) impacts the reconstruction quality. We hypothesize that the unimpeded flow of information through the skip connections can benefit the network’s capability to reuse encoded abstractions.

4.3 Model Capacity Tests

Several variables are expected to greatly influence the resulting distortion and bitrate. Here, we focus on the model size, but for longer videos, it may be beneficial to split the complete sequence into slices with fewer frames and repeat the process independently for each. This can not only affect the quality of reconstruction but also encoding time (the time it takes to encode a complete video). For that reason, we will test several model sizes and their capability to encode video sequences of various length, as well as the time it takes to encode a complete video. Then, the configurations that offer the best distortion-to-bitrate trade-off will be selected for the subsequent core evaluation stage.

4.4 Core Evaluation

The baselines will include common video compression codecs such as H.264 [4] and H.265 [5] as well as two recent methods of neural video compression of DVC [11] and VCII [12]. For the latter, the models pretrained by the authors that are publicly available will be employed. We will compare these baselines against the best performing variants of our method resulting from the preceding experiments described above. This leads to a total of 10 tested compression techniques, as listed in Table 1.

Table 1: Set of Tested Compression Techniques

Type	Variant
CNN	Full Sequence (as described in 4.1)
	Weight Stepped - Two Stage
	Weight Stepped - Single Stage
CbMLP	Full Sequence (as described in 4.1)
	Weight Stepped - Two Stage
	Weight Stepped - Single Stage
Video Codecs	H.264 [4]
	H.265 [5]
Neural Video Compression	DVC [11]
	VCII [12]

Our results will contain rate-distortion plots for each tested method and dataset. Furthermore, we will provide further analysis of total encoding and decoding time depending on the video length for each method since these can be expected to vary significantly.

5 Conclusion

Compression using neural implicit representations has only been tested for 2D images, where a degree of compressive capability has been demonstrated. In this work, we would like to explore the potential of similar approaches to video compression. Furthermore, a previously unexplored technique of low-entropy weight steps is described, which aims to learn network state sequences with high shared information content.

References

- [1] D. Ding, Z. Ma, D. Chen, *et al.*, “Advances in Video Compression System Using Deep Neural Network: A Review and Case Studies,” *Proceedings of the IEEE*, vol. 109, no. 9, pp. 1494–1520, 2021, ISSN: 15582256. DOI: 10.1109/JPROC.2021.3059994. arXiv: 2101.06341.
- [2] . Tudor, “MPEG-2 Video Compression,” *Electronics & communication engineering journal*, vol. 7, no. 6, pp. 257–264, 1995. DOI: 10.1007/springerreference_72962.
- [3] T. Sikora, “The MPEG-4 video standard verification model,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 7, no. 1, pp. 19–31, 1997, ISSN: 10518215. DOI: 10.1109/76.554415.
- [4] T. Wiegand, G. Sullivan, G. Bjontegaard, and A. Luthra, “Overview of the h.264/avc video coding standard,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 560–576, 2003. DOI: 10.1109/TCSVT.2003.815165.
- [5] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, “Overview of the high efficiency video coding (hevc) standard,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, no. 12, pp. 1649–1668, 2012. DOI: 10.1109/TCSVT.2012.2221191.
- [6] T. Chen, H. Liu, Q. Shen, *et al.*, “DeepCoder: A deep neural network based video compression,” *2017 IEEE Visual Communications and Image Processing, VCIP 2017*, vol. 2018-January, pp. 1–4, 2018. DOI: 10.1109/VCIP.2017.8305033.
- [7] L. I. Dong, L. I. Yue, L. I. Jianping, *et al.*, “Deep learning-based video coding: A review and a case study,” *ACM Computing Surveys*, vol. 53, no. 1, 2020, ISSN: 15577341. DOI: 10.1145/3368405. arXiv: 1904.12462.
- [8] S. Ma, X. Zhang, C. Jia, *et al.*, “Image and Video Compression with Neural Networks: A Review,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 6, pp. 1683–1698, 2020, ISSN: 15582205. DOI: 10.1109/TCSVT.2019.2910119. arXiv: 1904.03567.
- [9] O. Rippel, S. Nair, C. Lew, *et al.*, “Learned video compression,” *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2019-October, pp. 3453–3462, 2019, ISSN: 15505499. DOI: 10.1109/ICCV.2019.00355. arXiv: 1811.06981.
- [10] H. Liu, H. Shen, L. Huang, *et al.*, “Learned video compression via joint spatial-temporal correlation exploration,” *AAAI 2020 - 34th AAAI Conference on Artificial Intelligence*, vol. 1, pp. 11 580–11 587, 2020, ISSN: 2159-5399. DOI: 10.1609/aaai.v34i07.6825. arXiv: 1912.06348.
- [11] G. Lu, W. Ouyang, D. Xu, *et al.*, “Dvc: An end-to-end deep video compression framework,” *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, pp. 10 998–11 007, 2019, ISSN: 10636919. DOI: 10.1109/CVPR.2019.01126. arXiv: 1812.00101.
- [12] C. Y. Wu, N. Singhal, and P. Krähenbühl, “Video compression through image interpolation,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11212 LNCS, pp. 425–440, 2018, ISSN: 16113349. DOI: 10.1007/978-3-030-01237-3_26. arXiv: 1804.06919.

- [13] A. Djelouah, J. Campos, S. Schaub-Meyer, and C. Schroers, “Neural inter-frame compression for video coding,” *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2019-October, pp. 6420–6428, 2019, ISSN: 15505499. DOI: 10.1109/ICCV.2019.00652.
- [14] E. Dupont, A. Goliński, M. Alizadeh, *et al.*, “COIN: COmpression with Implicit Neural representations,” in *ICLR 2021 Workshop Neural Compression*, 2021, pp. 1–12. arXiv: 2103.03123. [Online]. Available: <http://arxiv.org/abs/2103.03123>.
- [15] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep Image Prior,” *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1867–1888, 2020, ISSN: 15731405. DOI: 10.1007/s11263-020-01303-4. arXiv: 1711.10925.
- [16] B. Mildenhall, P. P. Srinivasan, M. Tancik, *et al.*, “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,” in *ECCV 2020*, vol. 12346 LNCS, 2020, pp. 405–421, ISBN: 9783030584511. DOI: 10.1007/978-3-030-58452-8_24. arXiv: 2003.08934. [Online]. Available: <http://arxiv.org/abs/2003.08934>.
- [17] C. Lei, Y. Xing, and Q. Chen, “Blind video temporal consistency via deep video prior,” *Advances in Neural Information Processing Systems*, vol. 2020-December, no. NeurIPS, pp. 1–11, 2020, ISSN: 10495258. arXiv: 2010.11838.
- [18] J. Ballé, V. Laparra, and E. P. Simoncelli, “End-to-end optimized image compression,” *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017. arXiv: 1611.01704.
- [19] J. Ballé, D. Minnen, S. Singh, *et al.*, “Variational image compression with a scale hyperprior,” *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, 2018. arXiv: 1802.01436.
- [20] D. Minnen, J. Ballé, and G. Toderici, “Joint autoregressive and hierarchical priors for learned image compression,” *Advances in Neural Information Processing Systems*, vol. 2018-December, no. Nips, pp. 10771–10780, 2018, ISSN: 10495258. arXiv: 1809.02736.
- [21] J. Lee, S. Cho, and S. K. Beack, “Context-adaptive entropy model for end-to-end optimized image compression,” *7th International Conference on Learning Representations, ICLR 2019*, no. 2017, 2019. arXiv: 1809.10452.
- [22] V. Sitzmann, J. N. Martel, A. W. Bergman, *et al.*, “Implicit Neural Representations with Periodic Activation Functions,” in *Advances in Neural Information Processing Systems*, 2020, pp. 1–12. arXiv: 2006.09661.
- [23] M. Tancik, P. P. Srinivasan, B. Mildenhall, *et al.*, “Fourier features let networks learn high frequency functions in low dimensional domains,” *Advances in Neural Information Processing Systems*, vol. 2020-December, pp. 1–24, 2020, ISSN: 10495258. arXiv: 2006.10739.
- [24] C. Louizos, M. Welling, and D. P. Kingma, “Learning sparse neural networks through L0 regularization,” *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, pp. 1–13, 2018. arXiv: 1712.01312.
- [25] C. Qiao, Y. Shi, Y. X. Diao, *et al.*, “Log-sum enhanced sparse deep neural network,” *Neurocomputing*, vol. 407, pp. 206–220, 2020, ISSN: 18728286. DOI: 10.1016/j.neucom.2020.04.118. [Online]. Available: <https://doi.org/10.1016/j.neucom.2020.04.118>.
- [26] Y. LeCun, J. S. Denker, and S. A. Solla, “Optimal Brain Damage (Pruning),” *Advances in neural information processing systems*, pp. 598–605, 1990.
- [27] P. Molchanov, S. Tyree, T. Karras, *et al.*, “Pruning convolutional neural networks for resource efficient inference,” *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, no. 2015, pp. 1–17, 2017. arXiv: 1611.06440.
- [28] K. Checinski and P. Wawrzynski, “Dct-conv: Coding filters in convolutional networks with discrete cosine transform,” *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–6, 2020.
- [29] A. Mercat, M. Viitanen, and J. Vanne, “Uvg dataset: 50/120fps 4k sequences for video codec analysis and development,” in *Proceedings of the 11th ACM Multimedia Systems Conference*. New York, NY, USA: Association for Computing Machinery, 2020, pp. 297–302, ISBN: 9781450368452. [Online]. Available: <https://doi.org/10.1145/3339825.3394937>.
- [30] H. Wang, W. Gan, S. Hu, *et al.*, “Mcl-jev: A jnd-based h.264/avc video quality assessment dataset,” in *2016 IEEE International Conference on Image Processing (ICIP)*, 2016, pp. 1509–1513. DOI: 10.1109/ICIP.2016.7532610.

- [31] *Yuv video sequences*. [Online]. Available: <http://trace.eas.asu.edu/yuv/> (visited on 09/15/2021).