

Latent-INR: A Flexible Framework for Implicit Representations of Videos with Discriminative Semantics

Anonymous ECCV 2024 Submission

Paper ID #2356

Abstract. Implicit Neural Networks (INRs) have emerged as powerful representations to encode all forms of data, including images, videos, audios, and scenes. With video, many INRs for video have been proposed for the compression task, and recent methods feature significant improvements with respect to encoding time, storage, and reconstruction quality. However, these encoded representations lack semantic meaning, so they cannot be used for any downstream tasks that require such properties, such as retrieval. This can act as a barrier for adoption of video INRs over traditional codecs as they do not offer any significant edge apart from compression. To alleviate this, we propose a flexible framework that decouples the spatial and temporal aspects of the video INR. We accomplish this with a dictionary of per-frame latents that are learned jointly with a set of video specific hypernetworks, such that given a latent, these hypernetworks can predict the INR weights to reconstruct the given frame. This framework not only retains the compression efficiency, but the learned latents can be aligned with features from large vision models, which grants them discriminative properties. We align these latents with CLIP and show good performance for both compression and video retrieval tasks. By aligning with VideoLlama, we are able to perform open-ended chat with our learned latents as the visual inputs. Additionally, the learned latents serve as a proxy for the underlying weights, allowing us perform tasks like video interpolation. These semantic properties and applications, existing simultaneously with ability to perform compression, interpolation, and superresolution properties, are a first in this field of work.

Keywords: Implicit Neural Representation · Video Compression · Video Understanding

1 Introduction

In today's age of content explosion, large quantities of data are created every second, and storing them reliably and efficiently is of utmost importance for many applications. A scalable compression technique enables companies to provide better services at reduced cost and helps the end consumer by improving their access to high-fidelity data in addition to decongesting the network. Since

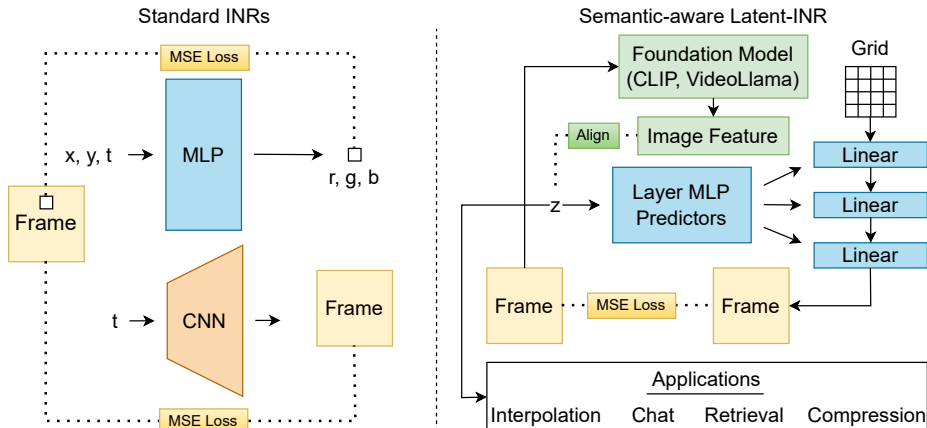


Fig. 1: Existing INRs for video (left) typically take some time-coordinate, or time and positional coordinates and train a single network to reconstruct a video. In contrast to these, we propose an INR system where a dictionary of implicit latent codes is learned for a video, one latent per frame. The latents are aligned to the image features of a large vision model, while simultaneously an INR system is learned which, given these latent codes, generates a positional INR which can reconstruct the frame. With this framework, we successfully develop an INR which performs both reconstructive tasks like compression, and semantic downstream tasks like retrieval and interactive chat.

the early 90s, several compression techniques have been created and widely deployed for this exact purpose. Out of these, JPEG [53] for images, HEVC [47], AV1 [12], and H.264 [54] for videos have emerged as the most popular choices, owing to their simple design and scalable performance.

In the past decade, the rise of deep learning led to a renaissance in computer vision, eventually impacting the visual data compression landscape [4, 15, 30]. Despite their success, these ML-based codecs have not seen widespread adoption like traditional codecs. This is in part due to failure to generalize, since ML codecs trained on large datasets can give sub-optimal compression for data points that differ significantly from their training set [7, 57]. Implicit Neural Representations (INR) attempt to avoid the generalization issue by operating internally. Instead of training large models that learn to identify *general* patterns in training data and apply them to specific out-of-distribution data, implicit techniques involve training a small model to exploit the *specific* patterns for the given data point. That is, for video compression, this approach would train one network per video, and for image compression, it would train one network per image. The resulting model is essentially a function that represents the underlying signal in spatial/temporal space.

Despite these advances, neural video compression remains unsolved. Various methods address issues of compression quality [9, 22], but two crucial questions remain unanswered – (i) how to scale for longer videos given architectural rigidity and (ii) how to reduce long encoding time due to training a network for every

061 video. Although recent works make some progress for these [28], the training time 061
062 is still quite long, and INR behavior for lossy compression is not well-understood, 062
063 limiting potential for practical adoption. 063

064 Furthermore, these approaches for INR tackle only one axis of the problem, 064
065 i.e., how to formulate video INRs with the primary goal of compression. These 065
066 aim to solve problems like long encoding time directly, by reducing it. In contrast 066
067 to these works, we instead aim to justify the compute and time needed 067
068 to train implicit representations. So, as a step towards ML-based codecs with 068
069 compelling real-world potential, we present Latent-INR – a new flexible frame- 069
070 work for formulating video INRs, where in addition to compression, the INR 070
071 enables downstream tasks like retrieval and video question answering, without 071
072 the need to decode the video. Our framework consists of two parts: (i) a diction- 072
073 ary of learnable latents, one for each frame, and (ii) a set of hypernetworks 073
074 learned on the entire video which, given a latent as input, predict frame-specific 074
075 weight modulations on the shared base network. This shared base takes a spatial 075
076 coordinate grid as input and outputs the specific frame 076

077 This design allows us to separate the spatial and temporal aspects of the 077
078 video by modeling them separately. We can view the set of hypernetworks as a 078
079 base model that learns the general structure and style of the video, while each 079
080 learned latent conditions it to output a specific frame. The latent here acts as a 080
081 proxy for the weights of the frame-specific INR. This property is apparent from 081
082 the video interpolation ability of our model - a task that other video INR repre- 082
083 sentations struggle to perform. Like other video INRs, our method is competitive 083
084 for compression, but uniquely retains the properties of original coordinate-based 084
085 INR. That is, our continuous representations of frames allows for spatial inter- 085
086 polation, which can be leveraged for superresolution and a decoding paradigm we 086
087 refer to as “any-resolution inference.” That is, at inference/decoding time, our 087
088 same model, with no changes to latents or architecture, can decode a video at 088
089 any resolution - a key feature missing from traditional codecs. This latent is also 089
090 quite flexible, and according to the procedure shown in Figure 1, we can align 090
091 it with the features from a large vision model, such as CLIP [33] to encode the 091
092 visual semantics of the frame while retaining nice properties such as alignment 092
093 with CLIP text embeddings. This allows for a whole spectrum of applications, 093
094 including frame, concept, and whole video retrieval with text queries. 094

095 In summary, our framework gives that extra edge apart from compression to 095
096 ML-based codecs, paving the way for their widespread adoption. Concretely, 096

- 097 – We propose an auto-decoder latent-based framework with spatio-temporal 097
098 decoupling for implicit video representations. Compared to other video INR 098
099 methods, this is a new way of formulating the problem. 099
- 100 – Our system has good compression performance, competing well with other 100
101 ML-based codecs for PSNR, BPP, and decoding speed while also enabling 101
102 any-resolution inference. 102
- 103 – The learnt latent embeddings from our framework demonstrate internal gen- 103
104 eralization from the encoded dataset, achieving video interpolation, a task 104
105 that other INR based methods struggle to achieve. 105

- We align our latents with large foundational models like CLIP [33], thus making our representations useful for retrieval tasks.
- We align our entire dictionary with video features for VideoLlama [56] to enable chat-style applications, including video question answering and captioning.

2 Related Work

Implicit Neural Representations (INR’s) are a class of neural networks designed with the intention of representing a given data point or dataset perfectly rather than exploiting general patterns and generalizing for unseen data. SIREN [42] pioneered the use of periodic activations to train simple MLP’s that worked well across images, SDF and audio. This was followed by a host of works that improved the training process of INR’s by making them faster [32, 35, 50] work across multiple scales [36] and encode multiple data points [14]. Models that used meta learning [46, 49] started gaining ground as they offered the advantages of compression along with generalization. [38, 48] further made improvements to this line of work by directly learning sparse-INR’s leading to improved compression and improved optimization by dataset selection respectively.

Hypernetworks are a class of networks optimized for predicting parameters of another network, with the aim of generalizing across unseen tasks [16]. Some utilized these for scenes [13, 43, 44]. Trans-INR [11] introduced the paradigm of using a transformer based hypernetwork to convert data directly from image-space to INR’s. [21] improved upon this idea and made the important observation that it is sufficient to modulate only the first hidden layer of an INR to represent a dataset of points. Unfortunately, these hypernetworks act on input data points which require test-time optimizations, making them unsuitable for compression tasks. [39] try to overcome this with an “auto-decoder” framework, where learnable latents represent a dataset of videos, with each latent corresponding to a single video, such that no encoder is needed. Others have investigated this paradigm for a variety of modalities [5, 37, 40]. Still, the lack of decoupling space from time prohibits the method from scaling to real-world videos.

Video INRs have recently gained popularity for compression. [9] was the first implicit representation which modelled a video as a function mapping the temporal coordinates to the corresponding frames. Later works [2, 8, 18, 25] iterated on this method, providing improvements in performance. [22] enhanced this concept by incorporating hash-grid [32] representations to speed up encoding times. NIRVANA [28] represented a video using a series of smaller INR models trained in an autoregressive manner to scale for longer videos.

Video Interpolation has been a fundamental task in computer vision, helping in creating smoother visual experiences. Over the past few years, deep learning based methods have vastly improved the quality of these interpolations [19, 41]. However, current INR-based video encoders lack this feature (see discussion in [8, 10], for example), hindering their widespread usage.

Video Retrieval is an essential process in the digital media landscape, where the objective is to efficiently search and extract specific video content from

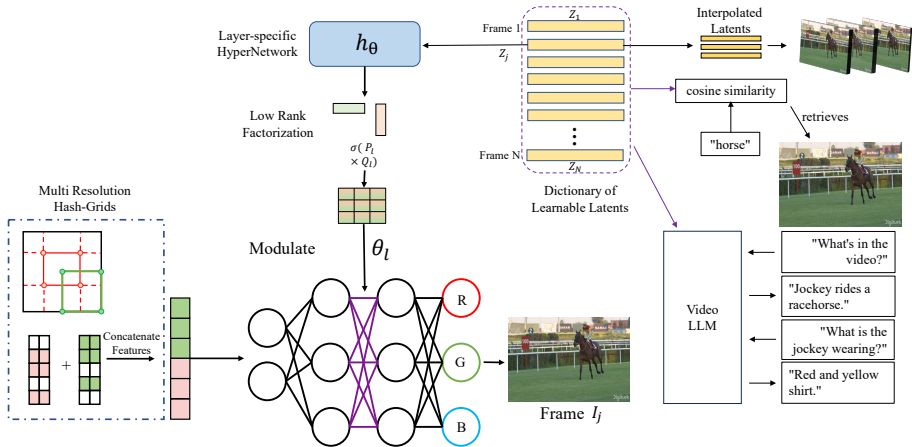


Fig. 2: We propose a new framework for video INR models by decoupling the spatial and temporal aspects of modeling. Our framework consists of auto-decoder based learnable latents that modulate the base network using a hypernetwork, via low-rank modulation. Once encoded, the resulting latents act as a proxy for the underlying weights of the representation. On the right, we show the use of these latents for additional tasks like video interpolation. By aligning these latents to the embedding space of foundational models like CLIP and VideoLlama, we also perform retrieval and chat.

expansive datasets. The complexity of understanding and indexing diverse video content has traditionally posed significant challenges. However, with the advent of machine learning-based methods, there has been a remarkable improvement in both the accuracy and efficiency of video retrieval systems [1, 3, 26]. These advances are limited to systems requiring an additional model, which can act as a burden on the system as they do not compress the data.

3 Approach

3.1 Background

Implicit Neural Representations parameterize a function,

$$f_{\theta} : X \rightarrow Y \quad \text{where} \quad X = \{(x_i, y_i) | 0 \leq x_i \leq W, 0 \leq y_i \leq H\}$$

represents the coordinate space, with *height* H and *width* W , and Y represents the underlying signal. In the standard case of an RGB signal, Y can be represented as

$$Y = \{(R, G, B) \mid R, G, B \in [0, 255]\}$$

This parametrization is usually trained with a standard MSE-loss, where we try to minimize the MSE-loss $\|f_{\theta}(X) - Y\|_2$. For a given video $V \in \mathbb{R}^{N \times H \times W \times 3}$ containing N frames, [42] represents them as pixels moving across time, i.e.,

$$f_{\theta}(x, y, t) = Y_t$$

where T_i denotes the boundary of a step. Other formulations exist which learn frame-based [9] or patch-based [28] representation, yet in each of these formulations, the focus is on representing the underlying data, with the added motivation of compressing it. However, none of these systems are designed with the goal of making these representations, f_θ , useful for downstream tasks. Instead, we utilize a learnable latent, z , as a part of an auto-decoder framework, along with a hypernet h to not only compress but to create useful representations.

$$f_\theta((x, y)|\theta_t) = Y_t \quad \theta_t = h(z_t) \quad (1)$$

The resulting latent z can be used for various downstream tasks like interpolation and retrieval, as we show in our work.

3.2 Latent-INR

Directly predicting the weights θ of the base network f , using the hypernet h , is expensive, parameter-heavy, and unsuitable for compression. Hence, we follow [45] [37] and instead predict low-rank matrices, which are then applied to the base network weights. This type of modulation acts as a form of subnetwork selection, analogous to systems proposed in [17] [34]. For a base network f with L layers, our formulation now looks like

$$f_\theta((x, y)|\theta_t^{l_1}, \theta_t^{l_2} \dots \theta_t^{l_L}) = Y_t \quad (2)$$

$$\theta_t^l = \sigma(P^l \times Q^l) \circ \theta^l \quad h_l(z_t) = [P^l, Q^l]$$

where θ^l represents the weights of the l -th layer and θ_t^l denotes the modulated weights for frame t . Here, σ signifies an activation function on the matrix-product of low rank matrices P^l , Q^l , which are of dimensions $R^{K \times r}$ and $R^{r \times K}$ where K is the width of the base network f_θ and rank $r \ll K$. These matrices are responsible for adjusting the weights θ_l as dictated by the corresponding hypernetwork h_l . Note that all hypernetworks use the same latent $z_t \in R^D$ as input. The rank r and the number of modulated layers essentially act as hyperparameters that control the compression-performance trade-off.

3.3 Model architecture

In our experiments, both the base network f_θ and hypernetworks h_l are feedforward MLP's that take in a coordinate input. Following [28], we also propose a variation to the base network with an additional convolutional up-sample block, which accepts coordinates of centroids as input and gives frame patches as output. We use the standard ReLU for base network and tanh for the hypernetwork as the respective non-linearities. The latents Z are initialized to be a standard normal with small variance, as we found empirically that this made the convergence faster. The complete model architecture is presented in Figure 2. For more details, see Appendix.

3.4 Model Compression

We train this entire system end-to-end with MSE-loss as the objective function. Once trained, we apply a standard quantization to all network parameters, further reducing the required storage. Given ϕ , a flattened parameter tensor, we transform it according to the following equations

$$\phi_i = \left\lceil \frac{\phi_i - \phi_{min}}{2^b} \right\rceil \quad \text{scale} = \frac{\phi_{max} - \phi_{min}}{2^b} \quad (3)$$

where the $\lceil \cdot \rceil$ (round) operation converts its argument to the nearest integer as dictated by bit width b of the quantization process. We also store the scale, ϕ_{max} , ϕ_{min} and the parameter shapes. These quantized values for all parameters are concatenated and further compressed using Huffman encoding.

3.5 Interpolation

Given a video of N frames and a scale α , the task of interpolation involves creating $\alpha \cdot N$ coherent frames that essentially increase the FPS of the video. Once we encode a video using our framework, we do linear interpolation on the frame latents $\{z_t\}$ and pass the resulting latent through the hypernetwork. This gives us the weight modulation required in the INR, and the updated base network is used to obtain the interpolated frames.

$$z_{inter} = \beta_i \cdot z_t + (1 - \beta_i) \cdot z_{t-1} \quad Y_{inter} = f_{\theta}(X; h(z_{inter})) \quad (4)$$

where,

$$\beta_i \in \left[\frac{1}{\alpha}, \frac{2}{\alpha}, \dots, \frac{\alpha-1}{\alpha} \right]$$

essentially generating $\alpha - 1$ frames between any two given frames. We train with held out frames and show results for $\alpha \in \{2, 4, 8\}$.

3.6 Downstream Tasks

Retrieval. Video retrieval involves searching and retrieving videos or clips from a large database based on similarity to given user search queries that are usually in the form of text. This can be viewed as a function R mapping query q to a set of corresponding videos V .

$$R : q \rightarrow V \quad (5)$$

The function R can use a similarity measure like cosine, euclidean, or nearest neighbors to retrieve closest matches. We encode a dataset of videos using our Latent-INR framework and use the resulting trained latents as our frame level representation. To ensure that the latents share the same space as the text queries, we add a cosine similarity loss between the latents and the CLIP image

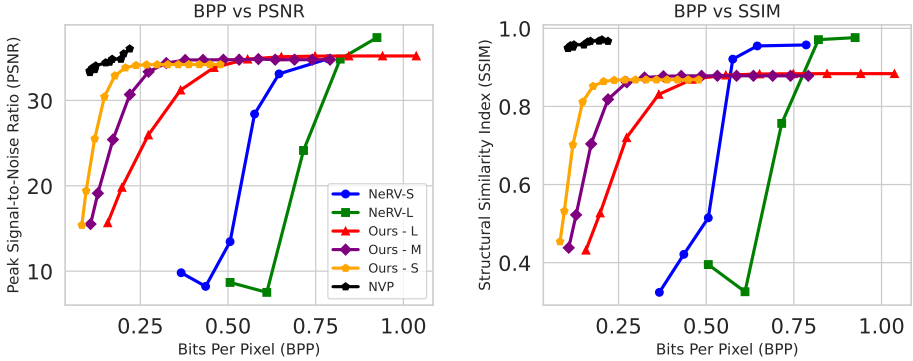


Fig. 3: We plot the rate distortion curves on PSNR and SSIM to compare compression with other methods. We observe that our large model achieves comparable PSNR to current SOTA [22]. Note that, while not plotted here, our decoding FPS is superior. Additional per-video results are available in the Supplementary.

embeddings of the corresponding frames. Our loss function during encoding now becomes:

$$L = L_{MSE} + \lambda \cdot L_{\text{clip}}(Z_t, Z_t^{\text{clip}}) \quad (6)$$

where Z_t^{clip} is the CLIP Image embedding of the input frame and λ controls the strength of this loss. We set the value of λ to be 0.01 in all our experiments.

Chat. We modify the formulation from retrieval slightly, aligning our dictionary of features to VideoLlama [56] instead of CLIP. Since the shapes are not compatible, we treat our latents as tokens and project the dimension to match the VideoLlama space. With this, we are able to integrate with the powerful LLM, substituting our latents for the raw video input tokens. We can then perform any task that VideoLlama can, in particular question answering and captioning. We wish to emphasize that our latents are flexible – we can align well with any large model, such that to improve the VQA or captioning performance, one would only need to align with a more powerful or efficient model.

4 Experiments

4.1 Video Compression

We perform comparative analysis for video compression on the standard Ultra Video Group (UVG) dataset [31].

This dataset comprises seven high-quality videos, each featuring diverse scenes shot at 120fps over a duration of five seconds. While most videos contain 600 frames, the ‘shakendry’ video is an exception with 300 frames, all at a resolution of 1080x1920. To assess the visual quality, we use standard metrics such as Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity index (SSIM). For measuring the storage efficiency of these methods, we use the standard bits




Fig. 4: With the same model, we can perform inference at any resolution, with speeds competitive or beating HEVC. We show sample frames for each resolution.



Fig. 5: We achieve high quality reconstruction and are able to reproduce even the finer details like water dountains and the hair on the horse.

per pixels (BPP) as our metric. As mentioned earlier, we use feedforward MLPs for both the base network f_θ and hypernetworks h_l . The base network consists of 6 layers with layer size of 512 and each hypernetwork that modulates a selected layer has one hidden layer of size 128 with tanh non-linearity, followed by the output layer. In the case where we use patch centroids as inputs, we add a convolutional layer followed by a pixel-shuffle [27] for upsampling.

We use hash-grids [32] for positional encoding due to their high quality reconstruction, although it should be noted we can use other schemes, such as Fourier features [50] to exchange some quality for faster training (see Appendix). We compare our method against NeRV [9] and NVP [22], with each of them encoding a video per model, and the results are presented in Figure 3. We observe that compression from our framework is comparable to baselines at similar bpp ranges, in addition to the other downstream benefits it offers.

Due to our architecture, we are also able to operate in a novel paradigm, “**any-resolution inference.**” Without changing the network architecture at all, we can decode the video at arbitrary smaller resolutions, as well as at higher resolutions (super-resolution) by leveraging the continuous resolution property of our hash grids and MLPs. We show our FPS decoding at various resolutions in Figure 4, although it should be noted that HEVC, the standard codec we compare to, must encode separately for every resolution while we can store all in the same model. We show some sample reconstructions in Figure 5 to showcase our method’s fidelity.

4.2 Video Interpolation

In our framework, learned latents serve as a mapping for the model weights, enabling valid frame outputs upon interpolation. We conduct experiments on two datasets: the “big buck bunny sequence” and a selection of ten videos from the Taichi test set. Frames are held out at a scale stride α during encoding. During testing, we interpolate the resulting latents on the held out frames and evaluate their performance.

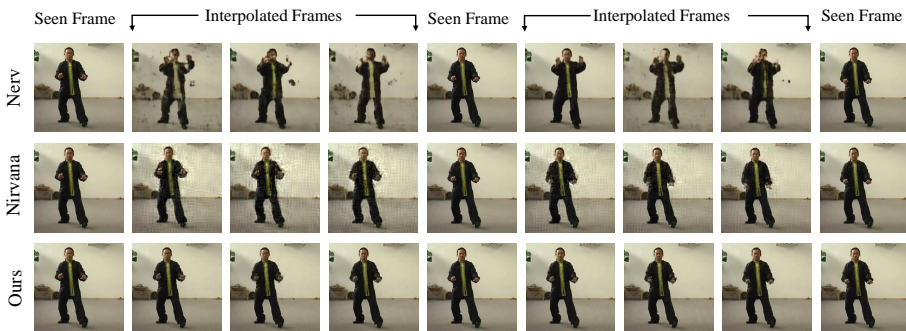
We use the same INR models utilized for compression as our baselines, with a reduction in network layer size and modulating mask rank. While NeRV [9] and NVP [22] interpolate time positions used as input, NIRVANA interpolates

Table 1: Interpolation Performance (PSNR), for different scale strides (α).

Dataset	α	NeRV	NIRVANA	NVP	Ours
Bunny	2	15.92	19.14	20.10	33.17
	4	15.43	18.90	19.11	28.08
	8	13.68	18.67	18.08	25.88
TaiChi	2	16.91	18.19	19.33	35.13
	4	17.14	17.71	18.52	31.84
	8	15.72	16.21	17.7	27.72

Table 2: Reconstruction and retrieval ablations of CLIP on MSR-VTT.

CLIP λ	Reconstruction		Retrieval (T2V)		
	PSNR	R@1	R@5	R@10	
0.0	30.03	0.1	0.3	0.8	
1e-3	29.83	28.4	50.8	60.6	
1e-2	29.46	30.2	52.4	61.0	
1e-1	28.93	29.7	51.5	61.8	
1.0	28.61	30.2	51.4	61.3	

**Fig. 6:** We compare interpolation with Latent-INR to NVP and NIRVANA. We find that our method has less artifacts and smoother motion in the interpolated frames.

the weights. In Table 1, we observe that while other INR methods fail to produce perceptual frames at scale of 2, our model can give reasonable interpolations even at a scale of 8. We confirm this qualitatively also, by inspecting interpolated frames such as those shown in Figure 6. Our outputs have noticeably fewer artifacts, and while imperfect, handle the motion better. Compared to other video INR methods, our approach of using learnt latents facilitates the model to have an internal representation of the video content.

4.3 Downstream Tasks

Retrieval

To showcase the flexibility of our latents, we align them with CLIP and evaluate their performance on standard retrieval tasks. We utilize the validation set of COIN dataset [51] and a subset of Howto100m dataset to evaluate performance. We first encode each video in our split using our Latent-INR framework with a loss that encourages the latents to be closer to the CLIP-Image embeddings of the frames, in addition to the standard reconstruction loss. We consider two distinct problems – retrieval of the correct class across all videos and retrieval

Table 3: Class and segment retrieval. Our method often exceeds CLIP performance.

Dataset	Method	Class Level			Segment Level		
		R@1	R@5	R@10	R@1	R@5	R@10
COIN	CLIP	31.60	44.70	50.70	6.60	13.10	16.50
	Ours	34.40	45.10	50.50	6.40	13.30	17.00
HowTo100m*	CLIP	31.58	36.84	47.37	21.13	37.32	40.85
	Ours	31.58	42.11	47.36	23.24	43.67	48.60

Table 4: Whole video retrieval. Our method matches CLIP performance.

Dataset	Method	Text to Video			Video to Text		
		R@1	R@5	R@10	R@1	R@5	R@10
MSR-VTT	CLIP	30.10	51.50	61.50	24.70	49.30	61.90
	Ours	30.20	52.40	61.10	25.40	49.90	61.70
ActivityNet*	CLIP	38.4	74.8	86.6	36.2	73.6	84.8
	Ours	38.5	73.9	86.4	36.1	73.5	84.7

**Fig. 7:** Nearest Neighbours for segment-level matching of sample queries from COIN validation set. The green boxes denote the true positives and the red ones are false positives. We show the inner product similarity between the image and the corresponding query inside the green boxes at the bottom of each image

of the correct segment within a video. These two use cases cover both ends of the spectrum, from localizing an event in a given video to searching for similar events across videos. We utilize the standard recall at K , where we have selected $k \in [1, 5, 10]$ to evaluate the efficacy of our method. The results are presented in Table 3. We can see that our method matches CLIP in its retrieval performance and even exceeds it in some cases. The qualitative results are presented in Figure 7, where we visualize the top 5 nearest neighbours of the text query that map to trained latents across all videos. Further results can be found in the supplementary.

We even find that our method can perform whole-video retrieval on MSR-VTT [55] and a custom 1,000 video sample from the ActivityNet Captions [23] ‘val-1’ split. We average-pool both our features and CLIP features (similar to [6]) and use CLIP features computed on video captions. In Table 4 we find that our retrieval is quite competitive to retrieval using the CLIP features themselves, showing that the learnt latents have similarly good averaging and summarizing



Fig. 8: Latent-INR LLM. We show results for aligning our learned latents to a VideoLlama model, which allows for interactive chat. We show a success case (left) and failure case (right) each for a generic (top) and more targeted prompt (bottom).

properties even over longer (180 seconds) videos, as well as alignment even to the paragraph-length captions used in ActivityNet.

Video-based Chat

We evaluate the performance of our trained latents, when aligned to intermediate VideoLlama features. This alignment enables access to the full scope of text chat with video understanding. We show a sample of such results, in the form of text and video prompts with text response, in Figure 8. These results show the LLM is able to understand video inputs when provided in the form of INR latents rather than raw video tokens. While not perfect, we infer the majority of the shortcomings of this system are primarily the fault of the LLM we align to.

Furthermore, on the basis of our success in aligning with CLIP and now VideoLlama, we believe our latents can be aligned to any representation. So, for more powerful chat, one simply needs to align to a more powerful chatbot. We thus provide these results two purposes. First, we show our model’s capability to power efficient open-ended captioning and question answering, while still retaining reconstruction capabilities. Second, we point to the immense potential of our model (or a similar paradigm) to continue to be leveraged with such models as they expand in their size and performance.

4.4 Visualizing Trained Latents

The trained latents, representing the modulated frames, offer intriguing insights when visualized in a reduced dimensional space. Utilizing Uniform Manifold Approximation and Projection (UMAP) [29] we project the embeddings Z_t into a 2D space, allowing for an intuitive interpretation of their relationships. In Figure 9, we plot the UMAP for three distinct videos from the UVG dataset: ‘Bosphore,’ ‘Honeybee,’ and ‘Jockey,’ each offering unique characteristics for examination. ‘Bosphore’, characterized by its slow-moving object and relatively




Fig. 9: We visualize the trained latents Z_t projected to 2D using UMAP. We show that the trained latents from our framework capture meaningful semantics of the underlying data. Left to right: Latents for Bosphore, Honeybee and Jockey from UVG dataset. Dark to Light color indicates frame numbers ranging from 0 to 600.




Fig. 10: Ablations to study the effect of layer modulations in the hypernetwork and the effect of patch size on reconstruction quality (PSNR).

static foreground, exhibits a smooth latent trajectory in the 2D space. This smoothness reflects the minimal variance in frame content, suggesting that our method effectively captures the subtle dynamics of the scene. In contrast, the ‘Honeybee’ video, with its repetitive frames, results in latents that cluster tightly together, signifying our model’s ability to recognize and encode repetitive patterns efficiently. The most dynamic of the three, ‘Jockey’, presents a more complex scenario with rapid changes in both the foreground and background. Here, the latents form clusters around similar scenes, yet maintain a discernible trajectory through the 2D space. These visualizations illustrate the semantic richness embedded within the latents obtained from our framework even when trained only for compression.

5 Ablation Studies

CLIP λ . We investigate the impact of the large model alignment weighting term on both reconstruction and retrieval for MSR-VTT. In Table 2, we find

that PSNR decreases slightly as λ increases. However, the retrieval performance seems to saturate at $\lambda = 0.01$. So, we suggest not tuning the λ too high for any application, given the diminishing returns.

Layer Modulations. In our approach, we have separate hypernetworks that modulated the selected layers. To evaluate the importance of each, we design an experiment where they are modulated in isolation. We use the same setup as the compression experiments with the modulating mask rank fixed at 20 for all models. In Figure 10, we can clearly see that the first few layers have a significant impact on the encoding performance. This matches the observations from [21] about the out sized impact of first few layers while modulating INRs.

Patch Size. Scaling to higher-resolution videos can be memory-intensive. This is particularly true when employing memory-demanding positional encoding schemes such as hash-grids [32]. To investigate this aspect further, we experiment with models that process centroids of fixed-size patches, directly predicting the corresponding frame patches, to save memory. From Figure 10 we observe that smaller patches have similar performance, with a steep drop observed for higher patch sizes.

6 Conclusion

Limitations. Our latents are somewhat restricted by the quality of the embeddings they are aligned to. Additionally, more work is still required to match standard codecs in terms of storage and encoding time, in spite of impressive gains in terms of quality and decoding speed. Future work could both improve the compression, and leverage more powerful vision models.

Broader Impacts. Our method for simultaneously compressing and learning useful features for recognition could reduce the need to decode videos for these tasks and thus save computational resources, cutting costs and helping the environment. However, work that advances performance for compression and recognition also has applications in surveillance and warfare.

In this work, we propose a new framework, Latent-INR, where we decouple the temporal aspect from the spatial into a dictionary of learnable latents. These auto-decoder based learnable latents modulate the layers of the base INR network via low-rank modulation using hypernetworks. Latent-INR is not only well-suited to video compression, but the resulting latents learn an internal representation of the data they encode that lends itself to SOTA interpolation for video INRs. Additionally, we also augment these latents by training them to be aligned with CLIP and VideoLlama, which allows us to bring the power of foundational models to compressed representations, and perform retrieval and chat-based applications like captioning and question answering. Our work thus opens up new possibilities of research in the implicit neural space where downstream tasks can be performed by these model without the need for decoding.

Latent-INR: A Flexible Framework for Implicit Representations of Videos with Discriminative Semantics

Supplementary Material

A Network Architecture

Base Network: We use an MLP with 10 layers, width of 512 and ReLU non-linearity as our base network f_θ .

Hypernetwork: All hypernetworks h^l used to modulate a layer l of the base network have 3 layers with a hidden dimension of 512 and tanh as non-linearity. Unless specified, we only modulate the first hidden layer of the base network.

Latents: Each latent Z_t corresponding to a frame has a dimension of 512 and is initialized to be standard Gaussian before training. We set our learning rate as $5e-4$ and used the standard Adam optimizer without any weight decay.

B Compression

B.1 Fourier Features

We use the multiresolution hash grid for positional encoding in all our models. In table 5 we show results for full coordinate resolution using fourier features for positional encoding. Due to lack of a hash grid, the resulting models train upto 30% faster, but at the cost of inferior reconstruction.

B.2 Quantization

Instead of quantizing all components equally, we notice that retaining the latents and the base network at full precision provides better reconstruction at negligible additional storage.

B.3 Effect of latent dimension

To study the effect of latent dimension on compression, we train models by varying it and encode the “bosphore” video from UVG dataset. The results are presented in Figure 11. We notice that there is positive gains till dimension 512 and diminishing returns thereafter. Hence we choose that as our default latent size in all our experiments.




Fig. 11: Effect of varying latent dimension across different bitrates.

C Video Retrieval

We perform two retrieval tasks on the COIN dataset [51] - *class-level*, and *segment-level*. In both settings, we use the standard val set as the database. For *class-level*, we use the distinct video-level task names in COIN as our query set. For *segment-level*, we use the set of distinct clip-level captions in COIN as our query set. We get the CLIP ViT-B/32 text embeddings of each of these captions, and these become our query vectors. For database vectors, we use the per-frame learned latents for each video in the database. For comparison with CLIP, we replace these database vectors with the CLIP ViT-B/32 image embeddings for each frame. For *class-level* retrieval, we consider a result frame a positive match if it belongs to a video with the same class label as the queried caption. On the other hand, for *segment-level* retrieval, we consider a result frame a positive match only if it belongs to a segment with the same caption as the query. Further, this search is done over all videos. We use FAISS [20] as our retrieval implementation and use *cosine similarity* as the distance metric.

We perform whole-level video retrieval as described in the main paper. For text, we use CLIP to compute a feature for the paragraph caption. For the video, we compute a per-frame feature for CLIP, or use the learnt latents from Latent-INR. For a single video feature, we then average these per-frame features. We normalize all features, and perform retrieval by finding the closest embeddings using dot product similarity. Both text-to-video and video-to-text are performed in the same manner, the only difference being which features are used as query and key.

Fig.12, shows the retrieval results on the COIN data in the *segment-level* setting. It can be seen that a majority of failure cases could be attributed to visual similarity across different tasks when seen at an individual frame level.

D Video Chat

We interface our latents with learned features from Video-Llama [56] to enable interactive chat with the compressed videos. In [56], the N video frames are



Fig. 12: Nearest Neighbours for segment-level matching of sample queries from COIN validation set. The green boxes denote the true positives and the red ones are false positives. We show the inner product similarity between the image and the corresponding query inside the green boxes at the bottom of each image

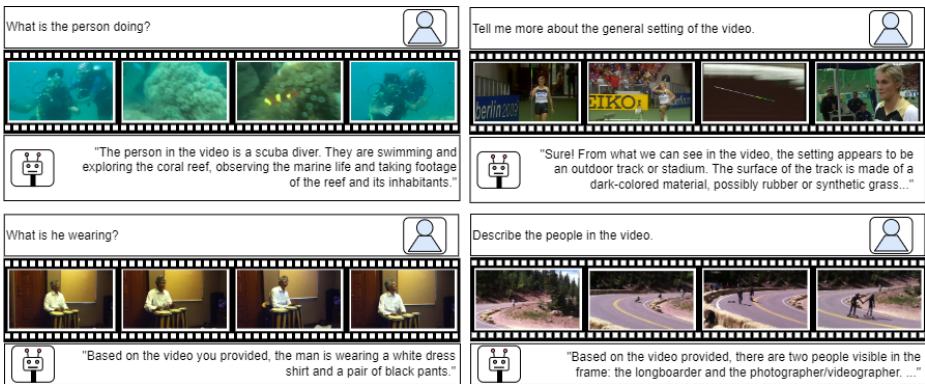


Fig. 13: Additional results for Latent-INR interface with Video-LLM.

passed through a ViT based visual encoder to extract features of size $k \times d$ per frame. These are then passed through a Query Former [24] to obtain a unified video representation of size $k_v \times d_v$. This tensor is then passed to a trainable MLP layer before aligning with an LLM of our choice (LLama-2 [52] in our models).

We align our latents Z with these per-video features of size $k_v \times d_v$ using a linear projection layer which is trained end to end. The loss function is slightly modified to incorporate a cosine similarity loss between the terms.



$$L = L_{MSE} + \lambda \cdot L_{cos}(F_t, F_t^{V-LLM}) \quad (7)$$

where F_t is the predicted feature and F_t^{V-LLM} is the corresponding Video-LLama extracted features. We show additional results of the interactive chat in Figure 13.

Table 5: Fourier Feature Models

Method	PSNR	BPP
Ours- Fourier - S	31.99	0.31
Ours- Fourier - M	33.69	0.62
Ours- Fourier - L	33.19	0.84

E Video-wise results

**Fig. 14:** BPP vs. PSNR, SSIM for beauty.**Fig. 15:** BPP vs. PSNR, SSIM for bosphore.

We plot the results for each video from UVG dataset [31] in Figures 14, 15, 16, 17, 18, 19, and 20. We show three versions of our model based on the dimension of the *low-rank* modulating matrix. The *Ours-s*, *Ours-m*, and *Ours-l* correspond respectively to *size* = 50, 100, 200. The *Ours-m* model achieves reasonable performance when compared to other methods, and at the same time can do the downstream tasks of interpolation and retrieval which none of the compared methods can.




Fig. 16: BPP vs. PSNR, SSIM for honeybee.




Fig. 17: BPP vs. PSNR, SSIM for jockey.




Fig. 18: BPP vs. PSNR, SSIM for readysteadygo.




Fig. 19: BPP vs. PSNR, SSIM for shakendry.




Fig. 20: BPP vs. PSNR, SSIM for yachtride.

References

1. A scheme for shot detection and video retrieval using spatio temporal features. *International Journal of Recent Technology and Engineering* (2019), <https://api.semanticscholar.org/CorpusID:241499090> 5
2. Bai, Y., Dong, C., Wang, C., Yuan, C.: Ps-nerv: Patch-wise stylized neural representations for videos. In: 2023 IEEE International Conference on Image Processing (ICIP). pp. 41–45. IEEE (2023) 4
3. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 1708–1718 (2021), <https://api.semanticscholar.org/CorpusID:232478955> 5
4. Ballé, J., Minnen, D.C., Singh, S., Hwang, S.J., Johnston, N.: Variational image compression with a scale hyperprior. *ArXiv abs/1802.01436* (2018), <https://api.semanticscholar.org/CorpusID:3611540> 2
5. Bauer, M., Dupont, E., Brock, A., Rosenbaum, D., Schwarz, J., Kim, H.: Spatial functa: Scaling functa to imagenet classification and generation. *arXiv preprint arXiv:2302.03130* (2023) 4
6. Buch, S., Eyzaguirre, C., Gaidon, A., Wu, J., Fei-Fei, L., Niebles, J.C.: Revisiting the "video" in video-language understanding (2022) 11
7. Cao, L., Jiang, A., Li, W., Wu, H., Ye, N.: Oodhdr-codec: Out-of-distribution generalization for HDR image compression. In: Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022. pp. 158–166. AAAI Press (2022). <https://doi.org/10.1609/AAAI.V36I1.19890>, <https://doi.org/10.1609/aaai.v36i1.19890> 2
8. Chen, H., Gwilliam, M., Lim, S.N., Shrivastava, A.: Hnerv: A hybrid neural representation for videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10270–10279 (2023) 4
9. Chen, H., He, B., Wang, H., Ren, Y., Lim, S.N., Shrivastava, A.: Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems* **34**, 21557–21568 (2021) 2, 4, 6, 9
10. Chen, H., Matthew, G., He, B., Lim, S.N., Shrivastava, A.: Cnerv: Content-adaptive neural representation for visual data. In: *BMVC* (2022) 4
11. Chen, Y., Wang, X.: Transformers as meta-learners for implicit neural representations. In: *European Conference on Computer Vision* (2022) 4
12. Chen, Y., Murherjee, D., Han, J., Grange, A., Xu, Y., Liu, Z., Parker, S., Chen, C., Su, H., Joshi, U., et al.: An overview of core coding tools in the av1 video codec. In: 2018 picture coding symposium (PCS). pp. 41–45. IEEE (2018) 2
13. Chiang, P.Z., Tsai, M.S., Tseng, H.Y., Lai, W.S., Chiu, W.C.: Stylizing 3d scene via implicit representation and hypernetwork. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 1475–1484 (2022) 4
14. Dupont, E., Kim, H., Eslami, S.M.A., Rezende, D.J., Rosenbaum, D.: From data to functa: Your data point is a function and you can treat it like one. In: *International Conference on Machine Learning* (2022), <https://api.semanticscholar.org/CorpusID:249395684> 4
15. Ehrlich, M., Davis, L.S.: Deep residual learning in the jpeg transform domain. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019) 2

- 534 16. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation 534
535 of deep networks. In: International Conference on Machine Learning (2017), <https://api.semanticscholar.org/CorpusID:6719686> 4 536
- 537 17. Frankle, J., Carbin, M.: The lottery ticket hypothesis: Finding sparse, trainable 537
538 neural networks. arXiv preprint arXiv:1803.03635 (2018) 6 538
- 539 18. He, B., Yang, X., Wang, H., Wu, Z., Chen, H., Huang, S., Ren, Y., Lim, S.N., Shri- 539
540 vastava, A.: Towards scalable neural representation for diverse videos. In: Proceed- 540
541 ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 541
542 pp. 6132–6142 (2023) 4 542
- 543 19. Jiang, H., Sun, D., Jampani, V., Yang, M.H., Learned-Miller, E.G., Kautz, J.: Super 543
544 slomo: High quality estimation of multiple intermediate frames for video inter- 544
545 polation. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition 545
546 pp. 9000–9008 (2017), [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:10817557) 546
547 [10817557](https://api.semanticscholar.org/CorpusID:10817557) 4 547
- 548 20. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with GPUs. IEEE 548
549 Transactions on Big Data **7**(3), 535–547 (2019) 2 549
- 550 21. Kim, C., Lee, D., Kim, S., Cho, M., Han, W.S.: Generalizable implicit neural 550
551 representations via instance pattern composers. In: Proceedings of the IEEE/CVF 551
552 Conference on Computer Vision and Pattern Recognition. pp. 11808–11817 (2023) 552
553 [4, 14](https://api.semanticscholar.org/CorpusID:1180811817) 553
- 554 22. Kim, S., Yu, S., Lee, J., Shin, J.: Scalable neural video representations with learn- 554
555 able positional features. Advances in Neural Information Processing Systems **35**, 555
556 12718–12731 (2022) [2, 4, 8, 9](https://api.semanticscholar.org/CorpusID:1271812731) 556
- 557 23. Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Niebles, J.C.: Dense-captioning events 557
558 in videos (2017) [11](https://api.semanticscholar.org/CorpusID:1180811817) 558
- 559 24. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: bootstrapping language-image pre- 559
560 training with frozen image encoders and large language models. In: Proceedings 560
561 of the 40th International Conference on Machine Learning. ICML'23, JMLR.org 561
562 (2023) [3](https://api.semanticscholar.org/CorpusID:1180811817) 562
- 563 25. Li, Z., Wang, M., Pi, H., Xu, K., Mei, J., Liu, Y.: E-nerv: Expedite neural video 563
564 representation with disentangled spatial-temporal context. In: European Confer- 564
565 ence on Computer Vision. pp. 267–284. Springer (2022) [4](https://api.semanticscholar.org/CorpusID:1180811817) 565
- 566 26. Luo, H., Ji, L., Zhong, M., Chen, Y., Lei, W., Duan, N., Li, T.: Clip4clip: An 566
567 empirical study of clip for end to end video clip retrieval. Neurocomputing **508**, 567
568 293–304 (2021), <https://api.semanticscholar.org/CorpusID:233296206> [5](https://api.semanticscholar.org/CorpusID:233296206) 568
- 569 27. Luo, H., Chen, Y., Zhou, Y.: An extremely effective spatial pyramid and pixel shuf- 569
570 fle upsampling decoder for multiscale monocular depth estimation. Computational 570
571 Intelligence and Neuroscience **2022** (2022), [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:251272212) 571
572 [CorpusID:251272212](https://api.semanticscholar.org/CorpusID:251272212) [9](https://api.semanticscholar.org/CorpusID:251272212) 572
- 573 28. Maiya, S.R., Girish, S., Ehrlich, M., Wang, H., Lee, K.S., Poirson, P., Wu, P., 573
574 Wang, C., Shrivastava, A.: Nirvana: Neural implicit representations of videos with 574
575 adaptive networks and autoregressive patch-wise modeling. In: Proceedings of the 575
576 IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14378– 576
577 14387 (2023) [3, 4, 6](https://api.semanticscholar.org/CorpusID:1437814387) 577
- 578 29. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and 578
579 projection for dimension reduction (2020) [12](https://api.semanticscholar.org/CorpusID:1180811817) 579
- 580 30. Mentzer, F., Toderici, G., Tschannen, M., Agustsson, E.: High-fidelity gener- 580
581 ative image compression. ArXiv [abs/2006.09965](https://arxiv.org/abs/2006.09965) (2020), [https://api.](https://api.semanticscholar.org/CorpusID:219721015) 581
582 [semanticscholar.org/CorpusID:219721015](https://api.semanticscholar.org/CorpusID:219721015) [2](https://api.semanticscholar.org/CorpusID:219721015) 582

31. Mercat, A., Viitanen, M., Vanne, J.: Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In: Proceedings of the 11th ACM Multimedia Systems Conference. p. 297–302. MMSys '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3339825.3394937>, <https://doi.org/10.1145/3339825.3394937> 8, 4
32. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**(4), 102:1–102:15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://doi.org/10.1145/3528223.3530127> 4, 9, 14
33. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (2021), <https://api.semanticscholar.org/CorpusID:231591445> 3, 4
34. Ramanujan, V., Wortsman, M., Kembhavi, A., Farhadi, A., Rastegari, M.: What's hidden in a randomly weighted neural network? In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11893–11902 (2020) 6
35. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.: Wire: Wavelet implicit neural representations. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 18507–18516 (2023), <https://api.semanticscholar.org/CorpusID:255749557> 4
36. Saragadam, V., Tan, J., Balakrishnan, G., Baraniuk, R.G., Veeraraghavan, A.: MINER: multiscale implicit neural representations. *CoRR abs/2202.03532* (2022), <https://arxiv.org/abs/2202.03532> 4
37. Schwarz, J.R., Tack, J., Teh, Y.W., Lee, J., Shin, J.: Modality-agnostic variational compression of implicit neural representations. *arXiv preprint arXiv:2301.09479* (2023) 4, 6
38. Schwarz, J.R., Teh, Y.W.: Meta-learning sparse compression networks (2022) 4
39. Sen, B., Agarwal, A., Namboodiri, V.P., Jawahar, C.: Inr-v: A continuous representation space for video-based generative tasks. *arXiv preprint arXiv:2210.16579* (2022) 4
40. Sen, B., Singh, G., Agarwal, A., Agaram, R., Krishna, K.M., Sridhar, S.: Hyp-nerf: Learning improved nerf priors using a hypernetwork. *arXiv preprint arXiv:2306.06093* (2023) 4
41. Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 1874–1883 (2016), <https://api.semanticscholar.org/CorpusID:7037846> 4
42. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions (2020) 4, 5
43. Sitzmann, V., Rezkikov, S., Freeman, B., Tenenbaum, J., Durand, F.: Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems* **34**, 19313–19325 (2021) 4
44. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems* **32** (2019) 4
45. Skorokhodov, I., Ignatyev, S., Elhoseiny, M.: Adversarial generation of continuous images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10753–10764 (2021) 6

- 633 46. Strümpfer, Y., Postels, J., Yang, R., Gool, L.V., Tombari, F.: Implicit neural representations for image compression. In: European Conference on Computer Vision 634 (2021), <https://api.semanticscholar.org/CorpusID:244954443> 4 635
- 636 47. Sullivan, G.J., Ohm, J.R., Han, W.J., Wiegand, T.: Overview of the high efficiency 637 video coding (hevc) standard. IEEE Transactions on circuits and systems for video 638 technology **22**(12), 1649–1668 (2012) 2 639
- 639 48. Tack, J., Kim, S., Yu, S., Lee, J., Shin, J., Schwarz, J.R.: Learning large-scale 640 neural fields via context pruned meta-learning (2023) 4 641
- 641 49. Tancik, M., Mildenhall, B., Wang, T., Schmidt, D., Srinivasan, P.P., Barron, J.T., 642 Ng, R.: Learned initializations for optimizing coordinate-based neural representations. In: CVPR (2021) 4 643
- 644 50. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn 645 high frequency functions in low dimensional domains. NeurIPS (2020) 4, 9 646
- 647 51. Tang, Y., Ding, D., Rao, Y., Zheng, Y., Zhang, D., Zhao, L., Lu, J., Zhou, J.: Coin: 648 A large-scale dataset for comprehensive instructional video analysis (2019) 10, 2 649
- 649 52. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., 650 Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023) 3 651
- 652 53. Wallace, G.K.: The jpeg still picture compression standard. Communications of 653 the ACM **34**(4), 30–44 (1991) 2 654
- 654 54. Wiegand, T., Sullivan, G., Bjontegaard, G., Luthra, A.: Overview of the h.264/avc 655 video coding standard. IEEE Transactions on Circuits and Systems for Video Technology (2003) 2 656
- 657 55. Xu, J., Mei, T., Yao, T., Rui, Y.: Msr-vtt: A large video description dataset for 658 bridging video and language. In: Proceedings of the IEEE Conference on Computer 659 Vision and Pattern Recognition (CVPR) (June 2016) 11 660
- 660 56. Zhang, H., Li, X., Bing, L.: Video-llama: An instruction-tuned audio-visual language 661 model for video understanding. arXiv preprint arXiv:2306.02858 (2023), 662 <https://arxiv.org/abs/2306.02858> 4, 8, 2 663
- 663 57. Zhang, M., Zhang, A., McDonagh, S.G.: On the out-of-distribution generalization 664 of probabilistic image modelling. In: Neural Information Processing Systems 665 (2021), <https://api.semanticscholar.org/CorpusID:237431305> 2 666