Streamable Neural Fields

Junwoo Cho^{1*}, Seungtae Nam^{1*}, Daniel Rho¹, Jong Hwan Ko^{1,2}, and Eunbyung Park^{1,2†}

¹ Department of Artificial Intelligence, Sungkyunkwan University

² Department of Electrical and Computer Engineering, Sungkyunkwan University {jwcho000,stnamjef,daniel231,jhko,epark}@skku.edu

Abstract. Neural fields have emerged as a new data representation paradigm and have shown remarkable success in various signal representations. Since they preserve signals in their network parameters, the data transfer by sending and receiving the entire model parameters prevents this emerging technology from being used in many practical scenarios. We propose streamable neural fields, a single model that consists of executable sub-networks of various widths. The proposed architectural and training techniques enable a single network to be streamable over time and reconstruct different qualities and parts of signals. For example, a smaller sub-network produces smooth and lowfrequency signals, while a larger sub-network can represent fine details. Experimental results have shown the effectiveness of our method in various domains, such as 2D images, videos, and 3D signed distance functions. Finally, we demonstrate that our proposed method improves training stability, by exploiting parameter sharing. Our code is available at https://github.com/jwcho5576/streamable_nf.

1 Introduction

Neural fields [47] have emerged as a powerful representation of real-world signals. It uses a multilayer perceptron (MLP) that takes inputs as the spatial or temporal coordinates and produces signal values in arbitrary resolutions. Thanks to recent advances such as input feature encoding [26, 42, 53] and periodic activation functions [37], it can faithfully reconstruct complex and high-frequency signals. It has achieved great success in various signal representations such as images [4, 20, 27], 3D shapes [2, 5, 7, 11, 15, 25, 29, 36], and novel view synthesis [3, 18, 21, 26, 28, 34, 51].

There are still many challenges that prevent this emerging technique from being used in practical scenarios. In neural fields, the network itself is a data representation [6] (the signals are stored as the parameters of neural networks), and the signal transmissions are done by sending and receiving the entire model parameters. Thus, finding the optimal model size is crucial for lower latency and higher throughput. A naive approach would train different size networks multiple

^{*}Equal contribution, alphabetically ordered.

[†]Corresponding author.

 $\mathbf{2}$



Fig. 1. Representing a varying signal with streamable neural fields. It is a *single* neural network executable at different widths that can reconstruct varying signal domains. *Spectral growing*: larger network reconstructs more high frequency details. *Spatial growing*: larger network reconstructs more pixel locations. *Temporal growing*: larger network reconstructs more pixel locations.

times by increasing depths and widths. However, it would not be an affordable solution since training even a single network takes a long time to converge for deep neural networks. It is tempting to predetermine various network configurations (e.g., widths and depths) for different sizes and types of signals, yet it is not a feasible solution either since the required size of the parameters is determined by the complexity of the signals, not the size or type of the data.

In addition, the raw signals often need to be transmitted in different resolutions or qualities. For example, in media streaming services, users want to receive various quality signals according to their circumstances, e.g., high-resolution videos at home and lower quality on mobile devices. Real-time encoding ondemand is not viable since it requires a long latency gradient descent on deep neural networks. As an alternative approach, we could locally store multiple sizes of the networks representing different qualities of the signals beforehand. However, it is a waste of storage space and not an acceptable solution given the exponential growth of media data.

Unlike most of the standard compression algorithms, such as JPEG [30] and MPEG [10] which are designed to be easily broken down into smaller pieces for potential use cases, including streaming service or partial reconstruction in poor network connections, a neural field cannot be decoupled into meaningful chunks.

All the weight parameters are highly intertwined, and missing a small part of them would result in catastrophic failures for signal reconstruction.

We propose streamable neural fields to overcome the issues mentioned above. We suggest training techniques and architectural designs that enable a single trained network to be separated into executable sub-networks of various widths. With single training procedure, the proposed algorithm can generate parameters of a single network that are streamable over time and capable of reconstructing various qualities of signals (Fig. 1). Each sub-network is in charge of representing some portion of the signals. For example, a small sub-network can only generate the signals in specific quality or a specific temporal (or spatial) range of the signals. A wider network that subsumes narrower sub-networks can represent the additional signals that are not encoded in the narrower sub-networks. By streaming the network parameters (from narrower to wider sub-networks), the signals will be progressively reconstructed in visual quality and temporal (or spatial) orders, which is desirable in many useful scenarios.

In sum, we present a single neural network that can represent multiple visual qualities and spatial (or temporal) ranges and decode signals in a streamline. The proposed network architecture and training strategy maximize the use of learned partial signals preserved in the small sub-networks. The larger networks explicitly make use of them, resulting in a more stable training procedure, improving reconstruction performance, and increasing parameter efficiency. We show the proposed method's effectiveness in various signals, including images, videos, and 3D shapes.

2 Related Work

Neural fields and spectral bias Neural fields, also known as coordinate-based neural representations or implicit neural representations have shown great success in representing natural signals such as images [4,20,27], videos [17,37], audios [37], 3D shapes [2,5,7,11,15,25,29,36] and view synthesis [3,18,21,26,28,34,51]. They struggled to represent high-frequency details due to low-dimensional inputs and spectral bias in training procedure [31,42]. Fourier feature encodings [26,42] and periodic non-linear activation function [37] enabled networks to represent fine details and have been successful. Although spectral bias is an unpleasant training behavior in many practical tasks, our work exploits this phenomenon to implement a neural field that can decode signals in various qualities with a single neural network.

Learning decomposed signals Several studies on neural fields represent spatially partitioned signals using voxel grids [9, 38, 48], latent codes [4, 23], and a group of neural networks [32]. Voxel-based approaches [9, 38, 48] directly bake a radiance field into the feature grids. The feature grids can be transmitted in a streamline, but their size is very large, which is unfavorable in streaming and compression. Another line of works [4, 23] divides an image into tiles and encodes them as latent vectors. While the latent vectors are much smaller than 4

the feature grids, a decoder is required on the client-side. Our method allows the network parameters to be streamed instead of the latent vectors and does not require an additional decoder to reconstruct the signal. KiloNeRF [32] subdivides a single network into numerous tiny ones. Similar to our method, the parameters are streamable. However, thousands of networks should be trained independently. In addition, we also found that learning a spatially partitioned scene with independent models yields line artifacts (Fig. 10), while our method seamlessly reconstructed the scenes.

On the other branch of work, a signal is partitioned in a frequency domain and learned hierarchically. Takikawa et al. [39] proposed to learn multiresolution codebooks similar to [40], allowing variable bitrate streaming. In a similar spirit to ours, recent works [16, 19, 35] have suggested a single network representing a signal with various bandwidths. Input layers are laterally connected to each intermediate layer, and the intermediate [16] (or additional output [19, 35]) layers reconstruct band-limited signals. To constrain the bandwidth, [19,35] initialize and fix the parameters of each input layer such that they are uniformly distributed over a certain frequency range. Progressive implicit networks (PINs) [16] sort sampled frequencies in ascending order, divide them into subsets, and use each subset as the Fourier encoding [26, 42] frequencies. Although these works [16, 19, 35]share some similarities with our method, there are significant differences. First, we exploit the spectral bias to learn the optimal frequency bandwidths given the limited network capacity while [19, 35] manually constrains the bandwidths, which may result in inefficient use of the network capacity. Second, our method is agnostic to input encoding methods while [16] designed customized algorithms for particular input encoding methods. Finally, we also proposed to grow widths of networks instead of depths [16, 19, 35].

Dynamic neural networks Unlike traditional neural networks with static architecture and size, dynamic neural networks can expand or shrink in size on the fly during training and inference. They can adapt to various computing environments and achieve a trade-off between efficiency and accuracy. One branch of dynamic neural networks related to our work is slimmable neural networks (SNN) [49,50], which dynamically expands the channel width of convolution filters during training. In SNN, without re-training each different network architecture, the model is executable for multiple predefined widths. Each trained sub-network has similar or better performance than individually trained models with the help of knowledge distillation [14] and parameter sharing.

In the field of lifelong learning (or incremental learning) [43], neural networks learn from a sequence of multiple tasks. Progressive neural network [33] dynamically expands while transferring knowledge from prior tasks to new ones to handle more tasks and overcome catastrophic forgetting behavior [12,22]. In this work, we focus on representing signals and propose a neural field that can dynamically grow network size to represent a higher-quality or broader range of signals while preserving the representations from small sub-networks.



Fig. 2. (a) Network architecture and training procedure of streamable neural fields. When training for each sub-network is done, it can grow the width into arbitrary size. To keep the output of each trained sub-network, previous weights are kept frozen. (b) Parameter initialization for newly added weights. Leaving the lateral connections (red) to default initialization, the remaining weights (green) are initialized to zero for fast convergence. We used SIREN initialization [37] as default.

3 Streamable Neural Fields

This section explains training techniques and architectural designs of *streamable neural fields* that consist of executable sub-networks of various widths. Once training is completed, a single network can represent signals in various qualities without re-training (spectral growing). Narrower sub-networks preserve low-frequency signals and wider sub-networks contain high-frequency details. This functionality can support many practical applications by selecting a width to offer requested quality signals on demand. In general, there is no straightforward solution to manually divide weights to achieve this goal since the network parameters are highly entangled.

In addition, streamable neural fields also support spatial and temporal growing. Each runnable sub-network can only represent a specific part of the entire signal. As depicted in Fig. 1, one would expect to send or receive a video sequentially over time, and one would prefer to receive a small part of a large image. It would be beneficial for transmitting large-size signals, such as high-resolution images or videos.

3.1 Network architecture and progressive training

Our network architecture and training procedure are illustrated in Fig. 2 (a). The model starts training with a small and narrow MLP to predict a target signal. Once converged, it grows its width by arbitrary size. Similar to progressive neural networks architecture [33], we remove the weights that connect from newly added hidden units to the previous units, which prevents the added units from affecting the small network's output. We also freeze weights in the small network and only update the newly added network parameters. This progressive training strategy encourages the large network to use the knowledge learned by the previous small

network and only learn the residual signals that the small network cannot capture. We keep iterating this process until the desired signal quality or spatial/temporal size is fulfilled.

Progressive training vs slimmable training Our general purpose is to create a single network that is executable at various widths. We found that the training technique in slimmable networks [50] can also achieve the goal for image and video fitting tasks. Unlike the proposed progressive training, it iterates over predefined widths, takes a sub-network of the corresponding width, and computes loss using the target signal prediction. The gradients of the sub-networks are accumulated until it visits every width, and weights are updated all at once. Experimental results have shown that progressive training outperforms slimmable training in reconstruction quality and convergence speed. One possible explanation is that the target residual signals for wider networks change over the training course, which results in slowing down convergence speed. The details of the progressive and slimmable training algorithms are described in the supplementary materials.

Training loss In every task, the training objective is to minimize the mean squared error (MSE) between prediction and ground truth target value:

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^{N} \|f_{\theta}(x_i) - y_i\|_2^2, \tag{1}$$

where $\{(x_i, y_i)\}_{i=1}^N$ is coordinate and corresponding signal value pair, e.g., $x_i \in \mathbb{R}^2$ and $y_i \in \mathbb{R}^3$ for image signals, e.g., RGB colors. f_θ is a neural network parameterized by θ . For spectral growing, the target signal is fixed to y_i . On the other hand, for spatial and temporal growing, we divided the input coordinate and the ground truth signal into the desired size for each network to represent. For example, to only represent a specific part of a certain sub-network, our objective becomes:

$$\min_{\theta} \frac{1}{N} \left(\sum_{i \notin S} \|f_{\theta}(x_i)\|_2^2 + \sum_{i \in S} \|f_{\theta}(x_i) - y_i\|_2^2 \right),$$
(2)

where S is the set of indices that belongs to the part. We train the network to predict zero for coordinate locations that do not belong to the part of interest. We empirically found that the first term in Eq. 2 is critical to seamlessly reconstructing the entire signals. The network usually predicts garbage outputs for the locations that are not part of training coordinates. It causes severe artifacts in the boundary areas when we stitch different parts of signals.

Initialization A careful initialization is required when using periodic activation, in order to achieve high performance and convergence speed. Fig. 2 (b) shows the suggested initialization scheme. We used the initialization method in SIREN [37] for lateral connections: $w \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$, where *n* is the number of input neurons in a particular layer. It helps the network maintain the activation distribution throughout layers. The remaining weights that connect between newly added neurons are set to zero. In a consequence, only information from the small sub-networks flows at the beginning of the training process. It would encourage the large network to utilize the small networks' knowledge first and keep newly added parameters from learning redundant signals already learned by the small networks. We empirically found that this technique increases network parameter efficiency and improves the convergence speed.

3.2 Spectral decomposition of streamable neural fields

In this section, we analyze a trained streamable neural field through the lens of spectral bias [31,42]. In neural fields using a simple MLP, the final output is a weighted sum of the previous layer's outputs (assuming no activation functions on the output layer). Any signals can be represented as a sum of different frequency components, and it is theoretically possible to assign values to the weight matrix in the final layer in a way that the final output can progressively represent higher frequency components as increasing the width.

The proposed progressive training scheme does not modify previously learned sub-network outputs, implying that the newly added hidden units only represent the residual signals (Fig. 3) and spectral bias [31,42] claims that the network prioritizes learning low-frequency parts of signals. Therefore, a narrow subnetwork trained by the progressive training would represent low-frequency signal, and wide network will preserve high-frequency details. The output value of a fullyconnected neural network is a linear combination of the final hidden activation. Our model progressively accumulates the final hidden activation learned by each sub-network. More formally, we can formulate the output layer of an MLP as follows:

$$y = f_{\theta}(x) = \sum_{j=1}^{d} w_j \phi_j(x) = \underbrace{\sum_{j=1}^{s} w_j \phi_j(x)}_{\text{low frequency}} + \underbrace{\sum_{k=s+1}^{d} w_k \phi_k(x)}_{\text{high frequency}}, \tag{3}$$

where $x \in \mathbb{R}^n$ is an *n*-dimensional coordinate, $y \in \mathbb{R}$ is a single-channel signal, $w \in \mathbb{R}^d$ is a weight vector in the final layer of the network, $\phi_j(x)$ is the *j*-th hidden unit in the final layer. The right-hand side in Eq. 3 is a decomposition into two partial sums, splitting the summation by index *s* (width of a sub-network). We assume there is no bias term in the final layer.

We can interpret ϕ_j as basis functions and w_j as coefficients. Unlike wellknown basis functions, such as Fourier basis or Chebyshev polynomials, we learn basis functions through progressive training and do not have any constraints e.g., orthogonality and periodicity on a Fourier basis. We solely rely on the inductive bias of MLP architecture and spectral bias of training neural networks to obtain spectral growing neural fields. 8



Fig. 3. Spectral growing and residual representation of streamable neural fields. Due to spectral bias, a neural network trained with a standard MSE loss would naturally build up signals in increasing frequency orders. Narrower sub-networks will first learn low-frequency components and wider sub-networks will increasingly represent higher frequency signals as long as the capacity of the network permit. Pruning the partial weights of the final hidden layer outputs residual signals and linearly combining them (red boxes) gives the full reconstruction.

Quality control The central premise of modern image compression algorithms is that the human eye cannot detect high-frequency components. JPEG [30] uses discrete cosine transform (DCT) and quantization matrix to remove the high-frequency parts in an image that are not recognizable to human eyes. It uses a quality factor that determines the compression rate and the reconstruction quality depending on the user's demand. Our streamable neural fields gradually eliminate the high-frequency component as the width decreases. Thus, choosing the width of the network to reconstruct a signal of desired quality is analogous to choosing a quality factor in image compression algorithms.

4 Experiments

We tested our model on various signal reconstruction tasks: 1D sinusoidal functions, 2D images, videos, and 3D signed distance functions (SDF). Our experimental setting is described in Fig. 4. For spectral growing, since there is no ground-truth residual signal, every individually trained model is trained to reconstruct the original target. On the other hand, for spatial/temporal growing, we divided the ground truth signal into multiple patches/frames and make individual models that have identical network sizes. Each model is then trained to represent specific image patches or video frames. Note that streaming spatial/temporal growing individual models can also achieve sequential signal transmission, while spectral growing individual models don't.

4.1 1D sinusoidal function reconstruction

In this section, we show the spectral growing of a signal by simple 1D scalar function fitting followed by [31]. This experiment intuitively shows the residual

9



Fig. 4. Description of our experimental setting. Every shown image above is the desired output signal and the number of parameters for each model is matched for a fair comparison. (a) Spectral growing experiment. The ground truth signal is fixed. (b) Spatial growing experiment. Individually trained models have constant network capacity.



Fig. 5. The results of the 1D sinusoidal function experiment. From (a) to (d): Output of each sub-network. As width grows, the model represents high-frequency details. Yellow dashed lines are the target ground truth signal. From (e) to (g): Summation of each sub-network and residual output. It is identical to the network outputs from (b) to (d). From (h) to (j): Residual output of each sub-network. The newly added weights reconstruct high-frequency residuals.

representation of our model. The target function is a mapping $f : [0, 1] \to \mathbb{R}$, constructed by summation of sinusoids with different frequencies and phase angles. Starting with a width of 10 (the channel size of hidden layers), we gradually increased the size up to 40. Each sub-network was trained for 150 epochs. Fig. 5 shows the output signals and residuals learned by each sub-network. As expected, the summation of low-frequency signals from small sub-networks and residual outputs give the same signal learned by larger sub-network.

4.2 Spectral growing in images and 3D shapes

Images We trained on 24 images in the Kodak dataset for spectral growing. The network size grows three times, using the same ground-truth images for four



Fig. 6. Quantitative result of the spectral growing experiment. (a): Averaged PSNR \uparrow , SSIM \uparrow , and LPIPS \downarrow of 24 Kodak images. (b): Chamfer distance (multiplied by 10³) of three 3D shapes.

different executable sub-networks. We compared to baseline method, denoted as *individual*, which trains MLPs with different sizes. The number of hidden units in *individual* models are adjusted to match the total number of parameters in sub-networks of the *streamable* network. *streamable (progressive)* denotes a streamable model trained with the proposed progressive training, and *streamable* (*slimmable*) for the slimmable training.

Fig. 6 (a) shows the averaged PSNR, SSIM [44] and LPIPS [52]. As expected, enlarging the network capacity gives a higher quality image. We found that the progressive training strategy outperforms the slimmable training. We believe the fact that during the slimmable training, large sub-networks affect the outputs of the small sub-networks, and vice versa would hurt the final reconstruction performance and convergence speed. In terms of the final reconstruction quality compared to *individual*, our method performs comparably given the same number of parameters. Especially, LPIPS gives a good score to *streamable (progressive)* model, even though when PSNR and SSIM do not. Since raw Kodak images contain undetectable high frequency components, our model gets rid of these due to spectral bias [31]. Although our model does not *exactly* reconstruct the target signal in terms of PSNR, low LPIPS implies that it sufficiently represents important features in terms of human's visual perception.

Moreover, the *individual* model in spectral growing does not have a streamable functionality, which means that the *individual* requires much more parameters to represent various qualities, e.g., four *individual* models for four different qualities. Table 1 shows the memory comparison between the *streamable (progressive)* model and the *individual* model. We trained on Kodak image 23 and divided it into 15 sub-networks that reconstruct 15 spectral growing images. The memory requirement of our *streamable (progressive)* model increases much more slowly compared to the *individual*. The difference between the two models becomes larger as more sub-networks are created.

In Fig. 7, we show a qualitative result on a large size image. The original image presents many stars of various sizes and brightness. After training each sub-network to the same target, we obtained residual signals by the method

PSNR	25.2	$25.2 \sim 27.8$	$25.2 \sim 29.8$	 $25.2 \sim 38.5$	$25.2 \sim 39.2$	$25.2 \sim 39.8$
ours	1.8K	$5.2 \mathrm{K}$	10.2K	 148K	$171 \mathrm{K}$	195K
individual	1.8K	$6.5 K_{(\times 1.25)}$	$15.0K_{(\times 1.47)}$	 $418K_{(\times 2.82)}$	$499K_{(\times 2.92)}$	590K(×3.03)

Table 1. The total number of parameters required to cover the indicated range (PSNRs in the first row) of image quality. The second row shows the number of parameters of each *sub-network* in our model to reconstruct the PSNR range above, and the rightmost value (195K) subsumes all the numbers left. The values in the third row are obtained by accumulating the number of parameters of *individual* models.



Fig. 7. Visualizing residual representation of elephant's trunk nebula image (2560×1984) . As the network enlarges, darker stars gradually appear. The third row shows the deformed residual images for better visualization. Note that all images shown are the output of a single network.

described in Fig. 3. The figure shows that the residual output images only contain small stars not captured by smaller sub-network.

3D shapes We also tested our method on 3D shapes represented by SDF. We trained on *Armadillo*, *Dragon*, and *Happy Buddha* from the Stanford 3D Scanning Repository. Following [29], we used bidirectional chamfer distance (CD) against the true shape to quantitatively compare different approaches. The results are shown in Fig. 6 (b) and the qualitative result of the *Dragon* shape is illustrated in Fig. 8.

Again, our single network outputs the same shape in various qualities. Despite many hyperparameter searches, the slimmable training approach did not converge on 3D shapes, so we only compared against the *individual*. Even with the streamable functionality in our method, *streamable (progressive)* shows better



Fig. 8. Spectral growing of *Dragon* in Stanford 3D scanning repository. Indicated numbers are chamfer distances against the reference shape.



Fig. 9. (a) Averaged performance of spatial growing of 8 Kodak images. (b) Averaged performance of temporal growing of 7 UVG videos.

reconstruction quality given the same number of parameters. Again, *individual* needs multiple independent models to support various quality outputs, which requires enormous storage space for 3D shapes.

In 3D shapes, training vanilla SIREN [37] was very unstable, so some models fail on reconstruction (see Fig. 6 (b) second column blue line). However, our streamable model shows a steady performance improvement when the model capacity increases. We provide an analysis of this phenomenon in terms of stable training behavior of the proposed progressive training (section 4.4).

4.3 Spatial and temporal growing

Images We trained on 8 images in the Kodak dataset for spatial growing. The network size grows three times and a total of four sub-networks represent spatially growing (horizontal direction) images. We compared to baselines described in Fig. 4 (b). For a fair comparison, model prediction is compared against the signal domain that contains the desired signal in the evaluation stage. The quantitative result in Fig. 9 (a) shows that *streamable (progressive)* models have higher representation power compared to the *individual* models. The performance drop along with the width increase implies that high-frequency details are concentrated

in the center of the images. Though the *streamable (slimmable)* model falls behind, due to the sandwich rule [49], the representation power keeps increasing as the width grows.

We also found that training independent models to spatially separate scenes yields line artifacts between each patch (Fig. 10 (a)) because the entire scene is not taken into account during the training of each individual network. Since these line artifacts are easily recognized, we believe our method has advantages over using individual models when learning spatially decom-



(a) individual

(b) streamable (progressive)

Fig. 10. Individually trained model shows line artifacts while our model doesn't in spatial growing of images.

posed signals (Fig. 10 (b)). Seamlessly combining individually trained scenes is an active research area, and the recent Block-NeRF [41] had to use additional alignment techniques, e.g., inverse distance weighting interpolation.

Videos We trained on 7 videos in UVG [24] dataset, resized to 480×270 for temporal growing. The network size grows two times and a total of three sub-networks represent temporally growing videos. Each sub-network reconstructs 8 frames, giving the largest network a total of 24 frames to reconstruct. We compared to baselines described in Fig. 4 (b). As shown in Fig. 9, for the same network capacity, *streamable (progressive)* models outperform the *individual* models.

4.4 Our Model Stabilizes the Training

As shown in Fig. 11, the PSNR curves of *individual* models drop significantly and then rebound during training, whereas those of *streamable (progressive)* models show no abrupt changes. We reconstructed the predicted RGB values of the largest *individual* model at the drop points and found severe artifacts on the images (Fig. 11 right). This phenomenon is not limited to the example provided. When we trained the *individual* models on 24 Kodak images, we found 19 similar cases. This suggests that the *streamable (progressive)* model can maintain the training process stable.

Since high-frequency components are sensitive to the perturbations in network parameters [31], fine-tuning the weights is required for a neural network to represent the high-frequency parts. Updating the entire parameters will significantly change the output signal and the loss value. This might be one reason why the individual model's PSNR curve fluctuates a lot. On the other hand, our progressive training method freezes pre-trained sub-network parameters and only updates the newly added weights. Updating partial weights results in small changes in the network output and encourages stable fine-tuning.



Fig. 11. PSNR curve and intermediate reconstructions of *individual4* model during training. The *individual* model's PSNR curve (blue) fluctuates during the entire training process while *streamable (progressive)* (red) doesn't.

5 Limitation and Discussion

Visiting each width during the training stage is unavoidable if a neural network is to be runnable at varied architectures, and it requires a demanding training time. There are several suggested parameter initialization methods on dynamically growing networks for fast convergence [8,45,46]. Further study on optimization and initialization techniques to improve the training dynamics can make our model more applicable for real-world tasks.

Conventional MLPs used in neural fields are black box models so that we cannot analyze how exactly the model reconstructs the desired signal. The predicted signal of our model can be decoupled into partial reconstructions and the additive sense of predicting the output can give interpretability to the general machine learning and deep learning models [1, 13]. We believe our work can also bring insight into the interpretability of neural fields and encourage further research on neural field architectures in the community.

6 Conclusion

We showed the possibility of decoding neural signal representation in a streamline. Our model is applicable to various signals such as images, videos, and 3D shapes which can vary in quality or grow spatially/temporally. Compared with individually trained models, the streamable neural field has similar or higher representation power, along with efficient memory cost by utilizing parameter sharing. Without re-training multiple individual models, our model can dynamically grow during the training procedure to search for optimal network capacity. We have also shown that partially training the model weights stabilizes training.

Acknowledgements This research was supported by the Ministry of Science and ICT (MSIT) of Korea, under the National Research Foundation (NRF) grant (2021R1F1A1061259), Institute of Information and Communication Technology Planning Evaluation (IITP) grants for the AI Graduate School program (IITP-2019-0-00421) and Artificial Intelligence Innovation Hub program (2021-0-02068).

References

- Agarwal, R., Melnick, L., Frosst, N., Zhang, X., Lengerich, B., Caruana, R., Hinton, G.E.: Neural additive models: Interpretable machine learning with neural nets. In: Advances in Neural Information Processing Systems. vol. 34, pp. 4699–4711. Curran Associates, Inc. (2021) 14
- Chabra, R., Lenssen, J.E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., Newcombe, R.: Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In: European Conference on Computer Vision. pp. 608–625. Springer (2020) 1, 3
- Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: Pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5799–5809 (June 2021) 1, 3
- Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8628–8638 (June 2021) 1, 3
- Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5939–5948 (June 2019) 1, 3
- 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 should treat it like one. arXiv preprint arXiv:2201.12204 (2022) 1
- Erler, P., Guerrero, P., Ohrhallinger, S., Mitra, N.J., Wimmer, M.: Points2surf learning implicit surfaces from point clouds. In: European Conference on Computer Vision. pp. 108–124. Springer (2020) 1, 3
- Evci, U., van Merrienboer, B., Unterthiner, T., Pedregosa, F., Vladymyrov, M.: Gradmax: Growing neural networks using gradient information. In: International Conference on Learning Representations (2022) 14
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5501–5510 (June 2022) 3
- Gall, D.L.: MPEG: A video compression standard for multimedia applications. Commun. ACM 34(4), 46–58 (November 1991) 2
- Genova, K., Cole, F., Vlasic, D., Sarna, A., Freeman, W.T., Funkhouser, T.: Learning shape templates with structured implicit functions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7154–7164 (October 2019) 1, 3
- Goodfellow, I.J., Mirza, M., Xiao, D., Courville, A., Bengio, Y.: An empirical investigation of catastrophic forgetting in gradient-based neural networks. arXiv preprint arXiv:1312.6211 (2013) 4
- Hastie, T.J.: Generalized additive models. In: Statistical models in S, pp. 249–307. Routledge (2017) 14
- Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) 4
- Jiang, C.M., Sud, A., Makadia, A., Huang, J., Niessner, M., Funkhouser, T.: Local implicit grid representations for 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6001–6010 (June 2020) 1, 3

- 16 J. Cho, S. Nam, D. Rho, JH. Ko and E. Park
- Landgraf, Z., Hornung, A.S., Cabral, R.S.: Pins: Progressive implicit networks for multi-scale neural representations. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 11969–11984. PMLR (July 2022) 4
- Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6498–6508 (June 2021) 3
- Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5741–5751 (October 2021) 1, 3
- Lindell, D.B., Van Veen, D., Park, J.J., Wetzstein, G.: Bacon: Band-limited coordinate networks for multiscale scene representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16252–16262 (June 2022) 4
- Martel, J.N.P., Lindell, D.B., Lin, C.Z., Chan, E.R., Monteiro, M., Wetzstein, G.: Acorn: adaptive coordinate networks for neural scene representation. ACM Trans. Graph. 40(4), 58:1–58:13 (2021) 1, 3
- Martin-Brualla, R., Radwan, N., Sajjadi, M.S.M., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7210–7219 (June 2021) 1, 3
- McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: Psychology of learning and motivation, vol. 24, pp. 109–165. Elsevier (1989) 4
- Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14214–14223 (October 2021) 3
- 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, MMSys. pp. 297–302. ACM (June 2020) 13
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4460–4470 (June 2019) 1, 3
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision. pp. 405–421. Springer (2020) 1, 3, 4
- 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 (2022) 1, 3
- Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11453–11464 (June 2021) 1, 3
- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 165–174 (June 2019) 1, 3, 11
- Pennebaker, W.B., Mitchell, J.L.: JPEG: Still image data compression standard. Springer Science & Business Media (1992) 2, 8

- Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 5301–5310. PMLR (June 2019) 3, 7, 8, 10, 13
- 32. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14335–14345 (October 2021) 3, 4
- Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016) 4, 5
- Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3d-aware image synthesis. In: Advances in Neural Information Processing Systems. vol. 33, pp. 20154–20166. Curran Associates, Inc. (2020) 1, 3
- Shekarforoush, S., Lindell, D.B., Fleet, D.J., Brubaker, M.A.: Residual multiplicative filter networks for multiscale reconstruction. arXiv preprint arXiv:2206.00746 (2022)
 4
- Sitzmann, V., Chan, E., Tucker, R., Snavely, N., Wetzstein, G.: Metasdf: Metalearning signed distance functions. In: Advances in Neural Information Processing Systems. vol. 33, pp. 10136–10147. Curran Associates, Inc. (2020) 1, 3
- 37. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: Advances in Neural Information Processing Systems. vol. 33, pp. 7462–7473. Curran Associates, Inc. (2020) 1, 3, 5, 6, 12
- Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5459–5469 (June 2022) 3
- Takikawa, T., Evans, A., Tremblay, J., Müller, T., McGuire, M., Jacobson, A., Fidler, S.: Variable bitrate neural fields. arXiv preprint arXiv:2206.07707 (2022) 4
- Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11358–11367 (June 2021) 4
- Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8248–8258 (June 2022) 13
- Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: Advances in Neural Information Processing Systems. vol. 33, pp. 7537–7547. Curran Associates, Inc. (2020) 1, 3, 4, 7
- Thrun, S.: A lifelong learning perspective for mobile robot control. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 23–30. IEEE (October 1994) 4
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. 13(4), 600–612 (July 2004) 10
- Wu, L., Liu, B., Stone, P., Liu, Q.: Firefly neural architecture descent: a general approach for growing neural networks. In: Advances in Neural Information Processing Systems. vol. 33, pp. 22373–22383. Curran Associates, Inc. (2020) 14

- 18 J. Cho, S. Nam, D. Rho, JH. Ko and E. Park
- Wu, L., Wang, D., Liu, Q.: Splitting steepest descent for growing neural architectures. In: Advances in Neural Information Processing Systems. vol. 32, pp. 10655–10665. Curran Associates, Inc. (2019) 14
- 47. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. Comput. Graph. Forum 41(2), 641–676 (July 2022), https://neuralfields.cs.brown.edu/1
- Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenoctrees for real-time rendering of neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5752–5761 (October 2021) 3
- Yu, J., Huang, T.S.: Universally slimmable networks and improved training techniques. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1803–1811 (October 2019) 4, 13
- Yu, J., Yang, L., Xu, N., Yang, J., Huang, T.: Slimmable neural networks. In: International Conference on Learning Representations (2019) 4, 6
- 51. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020) 1, 3
- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (June 2018) 10
- Zhong, E.D., Bepler, T., Davis, J.H., Berger, B.: Reconstructing continuous distributions of 3d protein structure from cryo-em images. In: International Conference on Learning Representations (2020) 1