

How to Train Neural Field Representations: A Comprehensive Study and Benchmark

Samuele Papa^{1,2}, Riccardo Valperga¹, David Knigge^{1,2}, Miltiadis Kofinas¹, Phillip Lippe¹,
 Jan-Jakob Sonke², Efstratios Gavves¹

¹ University of Amsterdam, ² Netherlands Cancer Institute

s.papa@{uva,nki}.nl, {r.valperga, dm.knigge, m.kofinas, p.lippe}@uva.nl,
 j.sonke@nki.nl, e.gavves@uva.nl

Abstract

Neural fields (NeFs) have recently emerged as a versatile method for modeling signals of various modalities, including images, shapes, and scenes. Subsequently, a number of works have explored the use of NeFs as representations for downstream tasks, e.g. classifying an image based on the parameters of a NeF that has been fit to it. However, the impact of the NeF hyperparameters on their quality as downstream representation is scarcely understood and remains largely unexplored. This is in part caused by the large amount of time required to fit datasets of neural fields.

In this work, we propose a JAX-based library¹ that leverages parallelization to enable fast optimization of large-scale NeF datasets, resulting in a significant speed-up. With this library, we perform a comprehensive study that investigates the effects of different hyperparameters on fitting NeFs for downstream tasks. In particular, we explore the use of a shared initialization, the effects of overtraining, and the expressiveness of the network architectures used. Our study provides valuable insights on how to train NeFs and offers guidance for optimizing their effectiveness in downstream applications. Finally, based on the proposed library and our analysis, we propose Neural Field Arena, a benchmark consisting of neural field variants of popular vision datasets, including MNIST, CIFAR, variants of ImageNet, and ShapeNetv2. Our library and the Neural Field Arena will be open-sourced to introduce standardized benchmarking and promote further research on neural fields.

1. Introduction

Neural Fields (NeFs) [43] have recently emerged as a versatile method for reconstructing discretely sampled continuous signals, with great success on various modalities from images [39–41], videos [32], 3D scenes and shapes [27, 28,

31, 46], to medical imaging [30, 47], audio [14, 39], and physical systems [19, 36, 45]. The continuous nature of coordinate-based models $f_\theta: x \mapsto y, x \in \mathbb{R}^d, y \in \mathbb{R}^c$ (e.g. $d = 2$ for images and $d = 3$ for 3D shapes), allows for modeling signals without the limitations in resolution inherent in discrete representations; their capacity for modeling fine detail is only limited by the expressivity of the underlying NeF architecture. Furthermore, some modalities, such as scenes and shapes, do not have array representations that can be easily adapted to existing deep learning methods. Doubling down on their flexibility, a growing body of literature [2, 10, 12, 20, 29] has further proposed to use NeF parameterizations f_θ directly as *representations of data*. They are used to model images, videos, or shapes not as tensors but as functions on which to run downstream tasks (e.g., classify the function of an image instead of the pixel tensor of an image). In this setting, the parameters θ of a NeF – sometimes along with their computational structure – serve as input for a downstream model.

Despite their promise, NeFs lag behind as image representations, while showing promise for shape datasets [20, 29]. We argue that two challenges contribute to their lack of performance: a practical one, and a theoretical one. The practical challenge is that a NeF is a neural network fit on a single discretely sampled signal, for instance, the pixels of an image. Fitting a new NeF for every sample in a dataset means training a new neural network from scratch, thus scaling up to whole datasets – especially large datasets with high-resolution signals – is prohibitively time-consuming [10, 12, 29]. The theoretical challenge stems from the fact that in this setting we treat the parameters of a neural network as a signal representation, and given our lack of understanding of how neural networks behave as a function of their architecture and optimization, we do not know what makes the parameters of a NeF a good representation for downstream tasks. Moreover, the scaling inefficiencies described above hinder us from forming hypothe-

¹<https://fit-a-nef.github.io/>

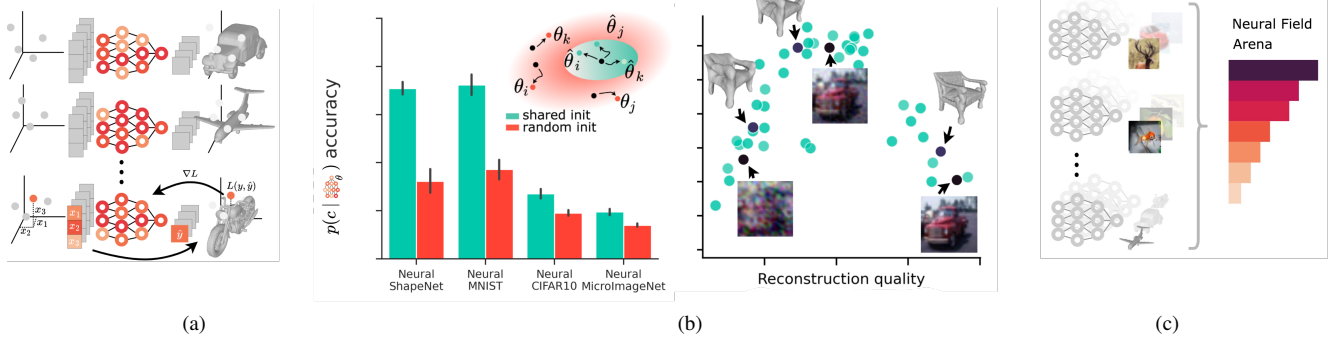


Figure 1. To investigate the use of Neural Fields (NeFs) as representations, we propose `Fit-a-NeF` (1a), a JAX-based library for efficient parallelized fitting of NeFs to datasets of signals, obtaining a $100\text{--}1,300\times$ speed-up. Owing to this efficiency, we are able to uncover the impact of NeF hyperparameter choices on their usability as representations – evaluated in downstream classification (Fig. 1b) – obtaining two important insights. First, 1b-left, it is vital to *group NeFs in parameter-space*, which we propose to enforce by sharing their network initializations. Second, 1b-right, *improved reconstruction quality does not necessarily result in improved representation quality*, implying an optimal combination of NeF expressivity and optimization for learning on NeFs. Incorporating these insights, we create a suite of NeF-based variants of classical CV datasets, Fig. 1c. We bundle these *Neural Datasets* into a benchmark for learning on Neural Fields – which we name *Neural Field Arena* – hoping to enable standardized comparison in order to promote further research into this field.

ses and testing them experimentally, inhibiting systematic research on NeF representations.

In this work, we address both challenges. Firstly, the existing libraries for learning neural networks are not necessarily optimal for fitting NeFs. To this end, we propose a JAX-based library [4, 16] –which we dub `Fit-a-NeF`– that leverages both in-device and across-device parallelization to reduce the time required for NeF fitting drastically.

Secondly, we explore the architectural considerations and optimization choices for making a NeF perform well as a representation. Equipped with an efficient library, we form structured hypotheses on the nature of parameter space –i.e. the space of the parameters of neural networks– and test them empirically.

In this work, we posit and experimentally verify that by constraining distinct NeFs to be close in parameter space, we improve the ability of downstream models to learn on the resulting sets of parameters. Further, we posit and experimentally verify that reconstruction quality with NeFs is not equivalent to the quality of NeF representations. From our analysis, we note that generalization to off-grid sampling should be used as an additional measure to gauge the quality of NeF representations.

Following our previous argument, we fit NeFs on popular datasets, which we call *Neural Datasets*, starting with MNIST, CIFAR, a modified version of ImageNet, and ShapeNet. To foster further research on NeF representations, we set up *Neural Field Arena*, a series of benchmarks on these Neural Datasets analogous to the benchmarks for the original datasets. In short, our contributions are as follows.

- We provide a JAX-based library to quickly and easily

fit millions of neural fields for different modalities, including images and shapes, using accelerators with parallelization on and across several devices.

- We use the library to perform a thorough investigation on how to obtain NeFs for datasets of signals to facilitate better downstream performance, expressed in downstream classification accuracy. We motivate and subsequently show that limiting parameter-space distances between NeFs during optimization is critical to obtaining NeFs that may serve as viable downstream representations. Furthermore, we demonstrate that the design choices for learning strong NeF representations are not necessarily the same as the design choices for good NeF reconstructions.
- We process NeF variants of popular computer vision datasets, creating a benchmark for Neural Field-based learning that we coin *Neural Field Arena* in the hopes of streamlining research in this direction.

We open-source code for the experiments and provide *Neural Field Arena*, a benchmark for learning on Neural Fields.

2. Representations in Parameter Space

Neural Fields are functions $f_{\theta} : x \mapsto y$, parameterized by a neural network whose parameters θ we optimize to reconstruct the continuous signal y on coordinates x . As with regular neural networks, fitting NeFs relies on gradient descent minimization. NeF representations are simply the parameters θ – possibly alongside their corresponding computational graph – that result from this optimization. As a consequence of the above, NeF representations inherit the intricacies and complexities of the topological space of neural networks, namely the extremely non-convex and high-

dimensional optimization landscapes resulting from severe overparameterization. Any downstream model that uses NeFs as signal representation is then required to learn on this highly complex space of representations.

To be more specific, for any two functions f_1 and f_2 (which we would like to fit with NeFs) that are ϵ -close in function space, it is not always theoretically possible to fit neural networks to these respective functions that are also close in their parameter space [34]. In other words, two signals that are similar might not have NeF representations which are themselves similar. Furthermore, the same neural network can be realized with parameters that are far apart in parameter space [1, 5, 15]. Importantly, a recent line of work shows that it is possible to align such neural networks that are close in function space but distant in parameter space, for example by finding optimal permutations in the weight space [1] or finding optimal transport solutions to (soft) neuron alignment [38]. Together, these properties dramatically complicate the task of learning on NeFs.

2.1. Grouping Representations in Parameter Space

Ainsworth et al. [1] posit that –modulo the permutation symmetries naturally present in neural networks– the local minima reached by optimization of different randomly initialized neural network weights form a convex basin. Following these insights, [20, 29, 48] show that downstream models that are invariant to these ambiguities improve performance dramatically over models that do not account for them. Both of these findings support the intuition that enforcing alignment in parameter space simplifies downstream learning. Our hypothesis is that the smaller pairwise distances in parameter space may simplify training because grouped NeF representations span a smaller region in parameter space, preventing, at least partially, very different NeFs from corresponding to semantically similar signals.

One possible solution is to train all NeFs for the same dataset using the same randomly selected but shared initialization instead of randomly initializing every NeF separately, see Fig. 1b. Intuitively, with such a shared initialization, the NeF representations do not stray away too far, thus making subsequent downstream learning easier. Based on that, we form our first research question.

RQ1: Do neural datasets obtained starting from a shared initialization lead to better downstream performance, when measured in terms of accuracy?

2.2. Reconstruction vs. Representation Quality

The primary purpose of NeFs in literature has been reconstructing the target signal, whether for image representation, shape modeling, or 2D-to-3D scene reconstruction with NeRFs [28, 39, 41]. As a consequence, much of the research into NeFs primarily focused on improving reconstruction quality [17, 22, 25], or providing architectures

| Neural Dataset | Reconstruction Quality | | | Ground Truth |
|----------------------|------------------------|--------|------|--------------|
| | Low | Medium | High | |
| Neural MNIST | | | | |
| Neural CIFAR10 | | | | |
| Neural MicroImageNet | | | | |
| Neural ShapeNet | | | | |

Figure 2. From left to right, samples from neural datasets with increasing reconstruction quality. The right-most column shows the ground truth used for fitting.

with favorable optimization properties [13]. However, improvements that lead NeFs to better reconstruction quality do not necessarily yield NeF parameters that serve as good representations. The reason is that reconstruction learning is a transductive problem where the primary goal is fitting a single function as accurately as possible. By contrast, representation learning is an inductive problem where generalization to out-of-sample data is also important.

2.2.1 Overtraining Neural Fields

Our hypothesis is that by overtraining we might obtain NeFs that achieve excellent reconstruction quality whose parameters, however, form poor representations. Neural Fields, like neural networks in general, are susceptible to overfitting. In NeFs, this would be reflected by sets of parameters that only approximate the underlying continuous signal well on the locations for which training data is available, and diverge from the signal elsewhere in the domain. Such NeF representations would be potentially ambiguous, leading to detrimental downstream performance. As such, we investigate the impact of overtraining on reconstruction quality *outside* the training domain by comparing with a surrogate of reconstruction quality in such locations and the number of steps that the network has been trained for. Ultimately, we assess its relation with NeF representation quality.

RQ2: Can overtraining NeFs negatively affect representation quality in terms of downstream accuracy, even if it improves the reconstruction of the given signals?

2.2.2 Expressivity of Neural Field Architectures

Similar to overtraining, our hypothesis is that NeF architectures originally proposed to improve reconstruction quality might not be offering better representational benefits despite their increased ability to reconstruct signals in high fidelity. In fact, we believe overly expressive architectures might have limited use as downstream representations, since their increased expressivity intuitively transfers as increased complexity in any downstream learning. A possible solution is to balance the expressivity of the neural network - i.e. its ability to reconstruct the signals of the dataset well - with its representation quality. As such, we propose to investigate the validity of this hypothesis on the reconstruction-representation trade-off and the viability of balancing reconstruction and representation through hyperparameter choices.

RQ3: Can the expressivity of NeF architecture be detrimental to their representation quality even if it helps with reconstruction quality?

3. Experimental Design

We start by giving a short description of our library that we coin `Fit-a-NeF`. Then, we describe the experimental design used to test our hypotheses and answer the research questions.

3.1. `Fit-a-NeF`: Fitting Neural Fields at Scale

The problem with fitting neural fields at scale is that training millions of small networks sequentially is not well-suited for accelerators, i.e. a GPU or a TPU, that are made to perform large matrix multiplications. As emphasized by Dupont et al. [10], fitting neural fields to large datasets can be computationally prohibitive. Modern Deep Learning frameworks offer `vmap`, a higher-order function used for automatic vectorization that applies a function element-wise to arrays, enabling batch processing and efficient vectorized computations by mapping over array axes without explicit loops. Our library uses JAX’s `vmap` [4] and Flax [16] to parallelize the NeFs training process. Together with the Just-In-Time (JIT) compilation of JAX, the library attains considerable speedups.

The primary goal of our library is to enable fast generation of Neural Datasets without the need of several GPUs. The library is sufficiently flexible to support datasets of different signals types, and new types can be integrated easily. A secondary goal is to offer a framework for loading and processing NeF representations.

Speed comparisons. We show comparisons of `Fit-a-NeF` with a naive approach in Fig. 3. Specifically, we compare the time it takes to fit various neural datasets with our library and with a naive approach consisting of fitting one NeF at a time, on the same GPU. The gain




Figure 3. Speedup obtained using `Fit-a-NeF` over a naive sequential approach using SIRENs using different hidden dimensions and number of layers. The evaluation was performed using 30k samples from MNIST on an A100 GPU. Smaller networks show the biggest speedup, as the parallelization is more effective.

depends on various factors such as the GPU memory and the size of the datapoints. For example, we can fit Neural MNIST in 17 seconds² vs. the 3 hours it would take with the naive approach. In short, with `Fit-a-NeF` it took 62 GPU days on A100s to run all the experiments reported in this work, whilst a naive approach would have required approximately 30 GPU years.

3.2. Datasets and Metrics

In this work, we explore both shapes and images for NeF representations. For the shape experiments, we use a 10-class version of ShapeNetv2 [7] for occupancy prediction, which we call ShapeNetv2-10. Based on the ShapeNetv2 shapes, we create the dataset by sampling 200k points at random, 100k of which are near the surface of the objects, as in [12]. For images, we use MNIST [23] and CIFAR10 [21]. Additionally, we created a 10-class version of ImageNet [9] with 64×64 resolution images, which we call MicroImageNet. The latter presents a more challenging signal to fit with the neural fields. We opt for a 10-class version of ShapeNetv2 and ImageNet for consistency. See Fig. 2 for examples.

Classification accuracy. In order to address our research questions and establish the quality of NeF representations, we train a downstream model for downstream tasks. More specifically, we choose NeF classification as a downstream task and report test accuracy, following prior work on representation learning [3, 8, 24]. For a downstream model,

²this is true if the GPU memory allows you to train the whole MNIST in parallel, e.g., on an NVIDIA A100

| Neural dataset | Initialization | Test accuracy | Gain (%) |
|-----------------------|----------------|---------------------------------|----------|
| Neural MNISTs | Random | 0.37 \pm 0.18 | |
| | Shared | 0.72\pm0.17 | 94.69 |
| Neural CIFAR10s | Random | 0.19 \pm 0.05 | |
| | Shared | 0.27\pm0.08 | 42.22 |
| Neural MicroImageNets | Random | 0.14 \pm 0.03 | |
| | Shared | 0.20\pm0.05 | 39.93 |
| Neural ShapeNets | Random | 0.32 \pm 0.18 | |
| | Shared | 0.71\pm0.10 | 119.63 |

Table 1. We measure the average *test accuracy* (mean \pm standard deviation) and percentage gain across 110 Neural Datasets trained using shared initialization and random initialization with different number of steps and hidden dimensions.

it is important to use a method that respects the neural network permutation symmetries. We choose the work of Kofinas et al. [20] that represents neural fields as computational graphs and processes them with standard graph networks. We create a variant that relies on a simpler message passing neural network (MPNN). We provide details about the classifier implementation in the appendix.

On-grid and off-grid reconstruction quality. We use *peak signal-to-noise ratio* (PSNR) in images and *intersection over union* (IOU) in shapes to measure the quality of reconstruction or generation. Additionally, we propose to extend both to what we call their *off-grid* counterparts; calculated using coordinates that are different from those used to fit the NeF to the signal. For *images*, the off-grids points lie in-between pixels, and the ground-truth image is obtained through a $\beta^{(1)}$ -spline interpolation of the original image at those coordinates. For *shapes*, new points and their respective occupancy is calculated for all shapes in the dataset.

We refer to them as “off-grid-PSNR” and “off-grid-IOU”, and use PSNR and IOU only to refer to their *on-grid* versions, which are computed on the coordinates used to fit the signals. In the experiments, we find the ratio between the two metrics $\frac{\text{off-grid-PSNR}}{\text{on-grid-PSNR}}$, and $\frac{\text{off-grid-IOU}}{\text{on-grid-IOU}}$ particularly informative.

Normalized Mutual Information. Intuitively, high-quality NeFs representations are semantically structured. We propose to use *Normalized Mutual Information* (NMI) to measure the information shared between the assignments of some clustering algorithm, here *k-means*, with *k* equal to the number of classes in the dataset, and the ground truth labels of the represented signals in a Neural Dataset.

3.3. Optimization

To obtain NeF representations of images, we use a mean-squared error loss function of the predicted pixel intensity of a pixel coordinate compared to the ground truth pixel in-

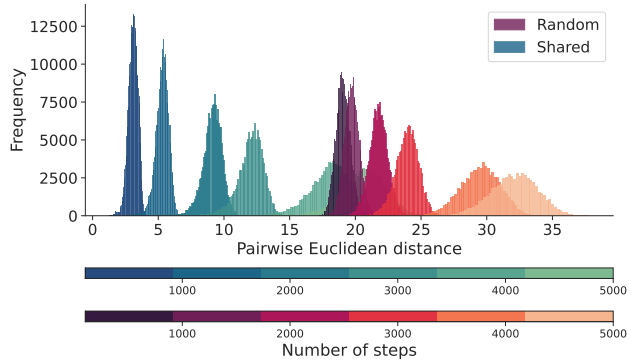


Figure 4. The histograms show the distribution of pairwise distances of the NeF representations. We fit 10 ShapeNet-10 Neural Datasets varying initialization and total number of steps. Shared initialization produces more grouped representations, and pairwise distances increase with the number of steps.

tensity. For shapes, we use binary cross-entropy measuring whether a point is correctly classified to be inside or outside the shape. We optimize the NeFs with Adam [18]. The experiments consist of a vast grid search on various hyperparameters including the NeF architecture, number of steps, and learning rate.

For a given hyperparameter choice, we fit a *neural dataset* to the images or shapes in the training dataset and then train a downstream classifier on it for a fixed number of epochs. The fitted Neural Dataset is, then, split into train, validation, and test subsets and the checkpoint classifier model with the highest validation accuracy is tested on the test set to obtain the test accuracy.

3.4. Architecture

In the main text of the paper we using SIREN [39]. In the supplementary material, we additionally experiment with different architectures classically used in NeF literature, RFFNets [41], and MFNs (FourierNet) [13] to assess their impact on representation quality, as well as the effects of the learning rate and weight decay. Given a certain architecture, we also explore its depth and width.

4. Results and Discussion

Here, we describe the experiments to test the aforementioned research questions and analyze the results. In the experiments that follow, each dot in the figures corresponds to fitting a whole dataset using a different set of hyperparameters that relates to a particular experiment. For instance, in Fig. 5, for the ShapeNet plot we have fitted 54 datasets for the 54 points plotted. We ran around 30,000 smaller preliminary experiments to determine the importance of the hyperparameters. Following this preliminary study we ran in-depth studies with the most relevant ones.




Figure 5. Results of the *test accuracy* (\uparrow) vs *NMI* (\uparrow) using different initialization on 220 Neural Datasets created using different hidden dimensions and the number of steps. Different datasets are stylized using different markers. Shared initialization leads to semantically structured NeF representation and, generally to better performance. The NMI of CIFAR10 and MicroImageNet are lower than those of ShapeNet and MNIST, however, are still clearly separated from their random initialization counterpart.

4.1. Shared vs. Random Initialization

Description. We experiment with two different initialization schemes for the neural fields before fitting them to the signals. With *random* initialization we refer to the standard setting of randomly initializing the NeF afresh for every new image or shape. With *shared* initialization, we denote the approach of initially randomly sampling an initialization for the NeF parameters and subsequently using the same initialization to fit all images or shapes.

Analysis. As reported in Table 1, shared initialization consistently yields higher accuracy.

To explain why shared initialization results in better performances we report in Fig. 4 the distribution of pairwise Euclidean distances between the NeF representations obtained with shared initialization and with standard initialization for different numbers of gradient descent steps.

Figure 4 shows that NeF representations trained for a larger number of steps are much more spread in representation space. Given the results of our study, we posit that this spread causes a decrease in the quality of the representations, likely making separability harder. Figure 5 shows how NeF representations obtained using the shared initialization already display a semantic structure. For NeFs trained under the shared initialization scheme, when measuring the mutual information between the assignment obtained with simple k-means and the ground truth labels, we can see how semantically similar neural fields are clustered together. Overall, a higher NMI between these two assignments results in higher downstream classification accuracy.




Figure 6. Results for the *reconstruction quality* (\uparrow) vs *test accuracy* (\uparrow) experiment. By fixing the NeF’s architecture we can more clearly see that there is a trade-off between visual quality and classification accuracy.

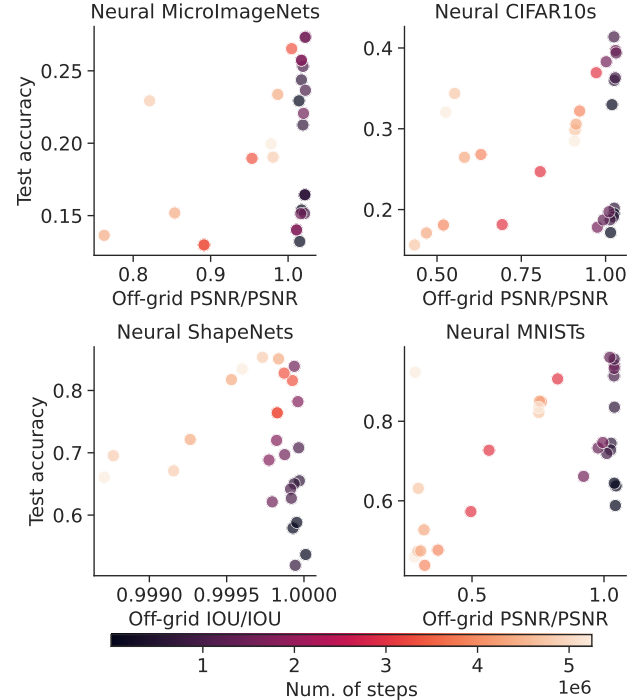


Figure 7. We fit 220 Neural Datasets using different hidden dimensions and the number of steps while keeping the same shared initialization. We find that the ratio of off-grid reconstruction quality and in-grid reconstruction quality can be used to form a heuristic that correlates with high test accuracy.

4.2. Overtraining Neural Fields

Description. We experiment with different hidden dimensions and number of steps with shared initialization to study the relationship between reconstruction quality and representation quality. We observe the relationship between reconstruction quality, off-grid reconstruction quality, and classification accuracy.

Analysis. The PSNR and IOU measure the reconstruction quality of neural fields, whereas the off-grid PSNR and off-grid IOU measure overfitting. We plot the classification accuracy achieved with Neural Datasets of different re-

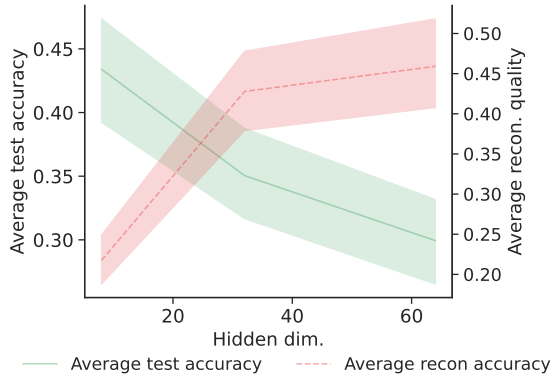


Figure 8. We measure the average test accuracy across all Neural Datasets and different numbers of steps. Additionally, we normalize the PSNR for the image datasets and the IOU for the shape dataset to measure the average reconstruction quality.

construction quality in Fig. 6. We observe that the classification accuracy increases with reconstruction quality but then reaches a maximum and starts decreasing. The general trend is consistent across different datasets, architectures, and optimization choices. In Fig. 7 we plot the classification accuracy against the ratio $\frac{\text{off-grid-PSNR}}{\text{on-grid-PSNR}}$ for images or $\frac{\text{off-grid-IOU}}{\text{on-grid-IOU}}$ for shapes. Interestingly, a clear pattern emerges that is consistent across datasets: when the number of steps is the only hyperparameter that is changed, stopping as soon as the NeF starts showing signs of overtraining leads to the best representation quality. We conclude that overfitting negatively affects downstream performances.

4.3. Expressivity of Neural Field Architectures

Description. We experiment with the same setting of Sec. 4.2. This time we focus on the effects of the architecture choice on test accuracy and reconstruction quality.

Analysis. Figure 8 aggregates results from 110 experiments across the different datasets and shows how increasing the hidden dimension tends to lead to worse representation quality while the higher expressivity of the NeF leads to higher PSNR. When exploring different activation functions and model layouts (see supplementary material) we observe a consistent behavior.

5. Benchmarks

Part of the fast progress of Deep Learning is due to publicly available datasets for benchmarking that allow for fair comparisons. Neural representations are a relatively new approach to representation learning and because of that, we lack good benchmark datasets to compare architectures that perform different downstream tasks on them. Since we observed great variability in downstream performance as we

| Neural Dataset | Test accuracy (%) | | |
|---------------------|----------------------------------|----------------------------------|----------------------------------|
| | GNN baseline | DWS [29] | Transformer [20] |
| Neural MNIST [29] | — | 85.70 \pm 0.60 | 92.40 \pm 0.30 |
| Neural MNIST (Ours) | 96.40\pm0.11 | 93.51\pm0.12 | 95.08\pm0.95 |

Table 2. We compare the proposed neural MNIST to previous work [29]. The insights gained in the study have resulted in a significant performance increase when the same models are used. We note that no hyper-parameter tuning was performed.

| Neural Dataset | Test accuracy (%) | | |
|-----------------------------|-------------------|------------------|------------------|
| | GNN baseline | DWS [29] | Transformer [20] |
| Neural CIFAR10 (Ours) | 39.83 \pm 1.70 | 44.01 \pm 0.48 | 44.11 \pm 0.20 |
| Neural MicroImageNet (Ours) | 25.08 \pm 0.01 | 28.61 \pm 0.48 | 28.50 \pm 0.09 |
| Neural ShapeNet (Ours) | 82.96 \pm 0.02 | 91.06 \pm 0.25 | 90.31 \pm 0.15 |

Table 3. We benchmark the new three proposed neural datasets using state-of-the-art methods from literature [20, 29].

vary the NeF architecture and optimization strategy, proper benchmarking is key to systematic progress hereupon.

After having explored the best optimization and architecture for several datasets, we now establish a new benchmark with existing downstream models that had previously achieved the state-of-the-art classification using neural representations.

Table 2 shows how state-of-the-art results can be achieved with our simple GNN baseline just by choosing the correct neural representation parameters.

We create Neural Datasets variants for CIFAR10, MNIST, MicroImageNet and ShapeNets-10, and bundle these as first entries in the Neural Field Arena. With these Neural Datasets, we evaluate the performance of DWS [29], a Relation Transformer [20], and our simple GNN baseline.

6. Related Work

Neural Fields. Recent years have seen a surging interest in the development of methods to facilitate NeF-based signal reconstructions. A number of works explore changes in neural network architectures to address the *spectral bias* of conventional ReLU MLPs towards low-frequency signals [35], which impedes the learning of high-fidelity signal reconstructions in fine detail. Tancik et al. [41] use a random Fourier feature mapping applied to the input of a NeF to prevent over-smoothing, whereas Sitzmann et al. [39] propose to use periodic activations to artificially inflate the frequency spectrum of the network, achieving better modeling of fine details. Multiplicative filters networks [13] are a different class of functions that avoids traditional compositional network depth and multiplies together simple functions, which has the added benefit of direct control over the network’s frequency spectrum [37]. Importantly, all of these works only investigate NeF architectures for improved

signal reconstruction quality. We instead focus on researching NeF architectures for their viability as *downstream representations*, i.e. learning on the parameter space of NeFs.

A different direction of research attempts to alleviate the computational burden of fitting individual NeFs to dataset samples, and to reduce the complexity of a NeF as representation. Building on the concept of neural network conditioning through modulation [33], Chan et al. [6], Mehta et al. [26] propose to use a single NeF to represent a diversity of signals. The activations of this NeF, which shares parameters across the entire dataset, are modulated with datapoint-specific low-dimensional latents to yield datapoint-specific outputs. The latents may be obtained using an encoder [26], or learned by using a combination of meta-learning and auto-decoding [31]. Dupont et al. [10] provided the insight that these learned latents may be used as downstream representation instead of the parameters of a sample-specific NeF to simplify the task of learning on a parameter space. Although reducing the complexity of learning on parameter space, this approach limits performance – both in terms of signal reconstruction quality and downstream performance [2]. Bauer et al. [2] show that a spatial grid of latent vectors is needed to reach competitive downstream performance, entailing all limitations that come with discrete data representations. In this work –to obtain the most general and flexible signal parameterizations– we instead focused on NeFs without shared parameters or sample-specific conditioning and explored what factors facilitate learning on their parameter space directly in downstream tasks.

Applications of neural representations. We briefly mention a few applications of neural representations. Several works propose their use for generating shapes using diffusion models [12]. In the context of dynamical systems and PDEs, [45] show that neural representations can be used for dynamics forecasting to the flow of partial differential equations as spatially and time-continuous functions.

Deep Weight Space methods. The exploration of deep learning architectures for processing neural network parameters is relatively recent. Here, we offer an overview of the key studies. The works by [11] and [42] revolve around predicting characteristics of trained neural networks (NNs) by scrutinizing their weight configurations. [11] is primarily focused on estimating the hyperparameters used during network training, while [42] is dedicated to evaluating a network’s generalization capabilities. Both investigations involve applying conventional neural networks to either the flattened weight data or their statistical attributes. On the other hand, [44] introduced a novel concept where NNs are processed using another neural network that operates on a combination of their high-order spatial derivatives. This technique is particularly suitable for NeF representations

where derivative information plays a crucial role. Nevertheless, the adaptability of these networks to broader tasks remains uncertain, and the demand for high-order derivatives can introduce significant computational overhead. Lastly, the works of [20], proposed a transformer-like architecture for compatibility with multiplicative filters networks, and the studies by [29] and [48], which initially emphasized the importance of augmentations and permutation-invariant architectures for processing neural field weights.

7. Conclusion

The goal of this work is to gain a deeper understanding of NeFs as data representations. To this end, we developed a library for fast generation and manipulation of Neural Datasets for different modalities, such as images and shapes. The speed gain in fitting Neural Datasets is substantial compared to naive approaches, practically enabling the use of neural fields for representing datasets at scale.

The study we conducted resulted in significant insights into what constituted better-quality NeF representations. The first key takeaway is that obtaining good NeF representations is a trade-off between good reconstruction and how spread NeFs are in parameter space. Generally, to achieve good reconstruction more gradient descent iterations are needed, which would spread the NeF representations in parameter space, negatively affecting their quality. Because of that, we propose to fit Neural Datasets from the same *shared initialization*, and demonstrate that this simple practice results in significantly higher representation quality across different modalities. Second, *overfitting hurts representation quality*. In other words, good reconstruction helps downstream performances up until the NeFs are overfitting the signals. To address this, we propose to stop fitting the representations early, before overfitting happens. Finally, we show how an increase in the number of parameters used to fit a signal leads to a decrease in downstream performance.

Last, we tackle the lack of benchmarks to compare architectures that perform different downstream tasks on NeF representations by providing the *Neural Field Arena*, covering various modalities.

Limitations. We limited our study to images and shapes leaving out neural radiance fields (NeRFs) and videos despite our library being suited to those types of NeFs. Additionally, for practical reasons, we choose to measure representation quality of NeFs using the classification accuracy of a GNN classifier. Although classification is often used in the representation learning literature to evaluate representation quality [3, 8, 24], we acknowledge that good classification accuracy might not result in good downstream performance for general tasks. We plan to explore other downstream tasks and include them in the Neural Field Arena in future work.

References

- [1] Samuel K Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*, 2022. [3](#)
- [2] Matthias Bauer, Emilien Dupont, Andy Brock, Dan Rosenbaum, Jonathan Schwarz, and Hyunjik Kim. Spatial functa: Scaling functa to imagenet classification and generation. *arXiv preprint arXiv:2302.03130*, 2023. [1](#), [8](#)
- [3] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013. [4](#), [8](#)
- [4] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. [2](#), [4](#)
- [5] Johann Brea, Berfin Simsek, Bernd Illing, and Wulfram Gerstner. Weight-space symmetry in deep networks gives rise to permutation saddles, connected by equal-loss valleys across the loss landscape. *arXiv preprint arXiv:1907.02911*, 2019. [3](#)
- [6] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 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*, pages 5799–5809, 2021. [8](#)
- [7] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. [4](#)
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [4](#), [8](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [4](#)
- [10] Emilien Dupont, Hyunjik Kim, SM Eslami, Danilo Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. *arXiv preprint arXiv:2201.12204*, 2022. [1](#), [4](#), [8](#)
- [11] Gabriel Eilertsen, Daniel Jönsson, Timo Ropinski, Jonas Unger, and Anders Ynnerman. Classifying the classifier: dissecting the weight space of neural networks. *arXiv preprint arXiv:2002.05688*, 2020. [8](#)
- [12] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. *arXiv preprint arXiv:2303.17015*, 2023. [1](#), [4](#), [8](#)
- [13] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. [3](#), [5](#), [7](#)
- [14] Ruohan Gao, Yen-Yu Chang, Shivani Mall, Li Fei-Fei, and Jiajun Wu. Objectfolder: A dataset of objects with implicit visual, auditory, and tactile representations. *arXiv preprint arXiv:2109.07991*, 2021. [1](#)
- [15] Robert Hecht-Nielsen. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pages 129–135. Elsevier, 1990. [3](#)
- [16] Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas Steiner, and Marc van Zee. Flax: A neural network library and ecosystem for JAX, 2023. [2](#), [4](#)
- [17] Amir Hertz, Or Perel, Raja Giryes, Olga Sorkine-Hornung, and Daniel Cohen-Or. Sape: Spatially-adaptive progressive encoding for neural optimization. *Advances in Neural Information Processing Systems*, 34:8820–8832, 2021. [3](#)
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [5](#)
- [19] Miltiadis Kofinas, Erik J Bekkers, Naveen Shankar Nagaraja, and Efstratios Gavves. Latent Field Discovery in Interacting Dynamical Systems with Neural Fields. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, 2023. [1](#)
- [20] Miltiadis Kofinas, Boris Knyazev, Yan Zhang, Yunlu Chen, Gertjan J Burghouts, Efstratios Gavves, Cees G. M. Snoek, and David W. Zhang. Graph Neural Networks for Learning Equivariant Representations of Neural Networks. In *12th International Conference on Learning Representations (ICLR)*, 2024. [1](#), [3](#), [5](#), [7](#), [8](#)
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [4](#)
- [22] Zoe Landgraf, Alexander Sorkine Hornung, and Ricardo Silveira Cabral. Pins: progressive implicit networks for multi-scale neural representations. *arXiv preprint arXiv:2202.04713*, 2022. [3](#)
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [4](#)
- [24] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR, 2019. [4](#), [8](#)
- [25] Julien NP Martel, David B Lindell, Connor Z Lin, Eric R Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *arXiv preprint arXiv:2105.02788*, 2021. [3](#)
- [26] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14214–14223, 2021. [8](#)
- [27] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks:

- Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 1
- [28] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 3
- [29] Aviv Navon, Aviv Shamsian, Idan Achituve, Ethan Fetaya, Gal Chechik, and Haggai Maron. Equivariant architectures for learning in deep weight spaces. *arXiv preprint arXiv:2301.12780*, 2023. 1, 3, 7, 8
- [30] Samuele Papa, David M Knigge, Riccardo Valperga, Nikita Moriakov, Miltos Kofinas, Jan-Jakob Sonke, and Efstratios Gavves. Neural modulation fields for conditional cone beam neural tomography. *arXiv preprint arXiv:2307.08351*, 2023. 1
- [31] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 1, 8
- [32] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 1
- [33] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 8
- [34] P Petersen, M Raslan, and F Voigtlaender. Topological properties of the set of functions generated by neural networks of fixed size (2018). *arXiv preprint arXiv:1806.08459*. 3
- [35] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019. 7
- [36] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. 1
- [37] David W Romero, Robert-Jan Bruintjes, Jakub M Tomczak, Erik J Bekkers, Mark Hoogendoorn, and Jan C van Gemert. Flexconv: Continuous kernel convolutions with differentiable kernel sizes. *arXiv preprint arXiv:2110.08059*, 2021. 7
- [38] Sidak Pal Singh and Martin Jaggi. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33:22045–22055, 2020. 3
- [39] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020. 1, 3, 5, 7
- [40] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8:131–162, 2007.
- [41] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 1, 3, 5, 7
- [42] Thomas Unterthiner, Daniel Keysers, Sylvain Gelly, Olivier Bousquet, and Ilya Tolstikhin. Predicting neural network accuracy from weights. *arXiv preprint arXiv:2002.11448*, 2020. 8
- [43] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. In *Computer Graphics Forum*, pages 641–676. Wiley Online Library, 2022. 1
- [44] DeJia Xu, Peihao Wang, Yifan Jiang, Zhiwen Fan, and Zhangyang Wang. Signal processing for implicit neural representations. *Advances in Neural Information Processing Systems*, 35:13404–13418, 2022. 8
- [45] Yuan Yin, Matthieu Kirchmeyer, Jean-Yves Franceschi, Alain Rakotomamonjy, and Patrick Gallinari. Continuous pde dynamics forecasting with implicit neural representations. 2022. 1, 8
- [46] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1
- [47] Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: neural attenuation fields for sparse-view cbct reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 442–452. Springer, 2022. 1
- [48] Allan Zhou, Kaien Yang, Kaylee Burns, Adriano Cardace, Yiding Jiang, Samuel Sokota, J Zico Kolter, and Chelsea Finn. Permutation equivariant neural functionals. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 3, 8