A Compact Implicit Neural Representation for Efficient Storage of Massive 4D Functional Magnetic Resonance Imaging

Ruoran Li[1], Runzhao Yang[1], Wenxin Xiang[2], Yuxiao Cheng[1], Tingxiong Xiao[1], Lu Yang[3], Jinli Suo[1]*

[1]Department of Automation, Tsinghua University, Beijing 100084, China
[2]School of Biomedical Engineering, Tsinghua University, Beijing 100084, China
[3]School of Mechanical Engineering, Tianjin University of Technology, Tianjin 300384, China
lrr24@mails.tsinghua.edu.cn; jlsuo@tsinghua.edu.cn

Abstract

Functional Magnetic Resonance Imaging (fMRI) data is a widely used kind of four-dimensional biomedical data, which requires effective compression. However, fMRI compressing poses unique challenges due to its intricate temporal dynamics, low signal-to-noise ratio, and complicated underlying redundancies. This paper reports a novel compression paradigm specifically tailored for fMRI data based on Implicit Neural Representation (INR). The proposed approach focuses on removing the various redundancies among the time series by employing several methods, including (i) conducting spatial correlation modeling for intra-region dynamics, (ii) decomposing reusable neuronal activation patterns, and (iii) using proper initialization together with nonlinear fusion to describe the inter-region similarity. This scheme appropriately incorporates the unique features of fMRI data, and experimental results on publicly available datasets demonstrate the effectiveness of the proposed method, surpassing state-of-the-art algorithms in both conventional image quality evaluation metrics and fMRI downstream tasks. This work in this paper paves the way for sharing massive fMRI data at low bandwidth and high fidelity.

Introduction

Functional Magnetic Resonance Imaging (fMRI), as a widely available imaging tool, has been widely used in cognitive neuroscience, clinical psychology, and psychiatry. As large-scale fMRI datasets continue to proliferate, there is an urgent necessity for a high-quality fMRI compression paradigm to achieve efficient storage and transmission.

However, there are several challenges in compressing fMRI data. Firstly, fMRI signals have a lower signal-to-noise ratio(Diedrichsen and Shadmehr 2005), and exhibit special patterns that are completely different from natural images. These issues pose challenges for existing algorithms in identifying and modeling the valid information within fMRI data. Secondly, there is significant variability in neural dynamics(Krohn et al. 2023) across different fMRI files. Current image compression algorithms are predominantly based on pre-trained encoder-decoder architectures, which cannot be fine-tuned to model the specific characteristics of each file restricted by generalization issues and time cost. Thirdly, most dynamic image compressors focus on reducing redundancies by modeling correlations between adjacent frames because of the high continuity of them. However, distinct from natural videos, the correlations in fMRI data predominantly reside within and between temporal signals of various brain regions (McKeown et al. 1998; Van Den Heuvel and Pol 2010), rather than adjacent frames, making it inefficient to apply these algorithms to fMRI.

To address these issues, we attempt to employ Implicit Neural Representation (INR) as the framework for fMRI compression for the first time, which possesses unique advantages. Unlike encoder-decoder architectures, INR trains a distinct network for each compression target data, which bypasses the generalization issue and can efficiently model the unique features in each data. This approach aligns with the specific characteristics of fMRI and the requirement of high modeling accuracy for medical data compression. Furthermore, INR excels in modeling internal data correlations(Sitzmann et al. 2020), which aligns with strong inter-regional correlations typical of fMRI(McKeown et al. 1998; Van Den Heuvel and Pol 2010). Additionally, INR-based compression algorithms are highly flexible in modeling correlations and are not limited to those between adjacent frames, allowing for better reduction of redundancies in fMRI compared to existing algorithms.

To better leverage the advantages of INR, we designed a novel INR architecture specifically for fMRI, which is fundamentally different from existing INR frameworks. Firstly, from an overall perspective, unlike NeRV's frame-index-to-frame mapping(Chen et al. 2021), our network architecture employs spatial-coordinate-to-time-series mapping, which aligns better with the higher temporal signal correlations among adjacent voxels within brain regions in fMRI, reducing intra-region redundancies. Moreover, our network decomposes the data into reusable neuronal activation patterns and their corresponding spatial distributions to effectively remove inter-region redundancies. Subsequently, we further reduce redundancies in the spatial distribution of activation patterns and introduce a feature fusion network to simulate the integration of neuronal activation patterns in real fMRI data, enhancing the encoding precision. Finally, we introduce an initialization scheme for neuronal activation patterns based on Independent Component Analysis (ICA)

*Corresponding author.

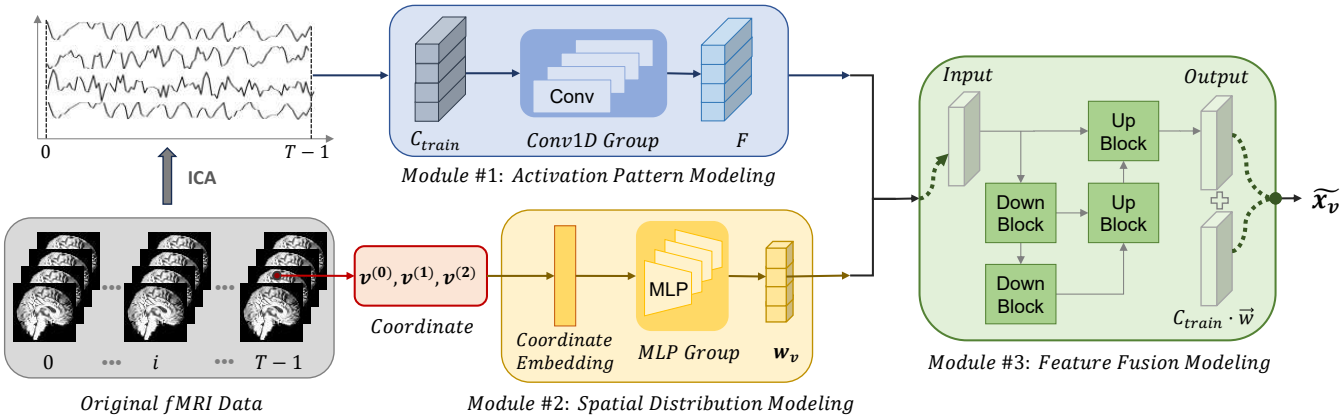

Figure 1: The basic structure of the proposed compression approach. Character T in the figure means T time slices.

to find the primary and unique neuronal activation patterns among each fMRI data, which can better retain the more significant information across brain regions to align with the requirement of fidelity.

We demonstrate the largely superior performance of our method on publicly available datasets in order to validate it. It's worth noting that, in terms of evaluation metrics, in addition to the common metrics like PSNR and SSIM, which measure pixel-wise differences, we also introduced additional evaluation criteria as part of our assessment which are downstream tasks widely used in fMRI data analysis.

In summary, this paper makes the following contributions:

- **INR-based Four-dimensional Data Compressor.** Development of a novel INR-based compression paradigm tailored for challenges of fMRI data, which is also a novel attempt to apply INR to 4D data compression.
- **Coordinate-to-time-series Mapping for Intra-region Correlations.** A novel perspective viewing the dynamic image data as a mapping from coordinates to time series to eliminate redundancies within local brain regions.
- **Decomposition and Fusion of Neuronal Activation Patterns for Inter-region Correlations.** A technique to describe fMRI dynamics with a nonlinear fusion of reusable activation patterns, minimizing the inter-regional redundancies.
- **Initialization with Main Neuronal Activation Patterns by ICA.** An ICA-based method to identify the dominant activation patterns in brain regions for initialization, in order to preserve salient information better.

Related Work

Data Compression. Over the past few decades, there has been rapid development in data compression algorithms, leading to tremendous successes with algorithms like JPEG(Wallace 1992), H.264(Wiegand et al. 2003), and H.265(Sullivan et al. 2012). The techniques which have been adopted by these compressors such as discrete wavelet transform(Heil and Walnut 1989) and block-size motion compensation,have later become extensively applied in the field of data compression. In recent years, with the emergence of deep learning, data compression techniques based on deep learning have gained momentum, resulting in the development of various new compression algorithms(Lu et al. 2019; Agustsson et al. 2020). However, when it comes to high-dimensional data, such as four-dimensional data, proper compressors are still lacking.

INR-based Compression. Recent progress in data compression has introduced INR for more compact data representation, yielding promising results(Chen et al. 2021; Li et al. 2022; Dupont et al. 2021; Damodaran et al. 2023). Unlike traditional data compression algorithms, INR-based data compression algorithms leverage the powerful information capacity of neural networks to implicitly store data information within network parameters. Some recent works have also introduced INR into the compression of medical imaging data(Yang et al. 2023; Yang 2023). Existing algorithms tend to model images as spatiotemporal continuous functions, but fMRI data exhibits poor temporal continuity and thus we take a distinct perspective by decomposing data into temporal profiles and spatial distributions instead. Besides, INR-based compression typically uses either time-to-frame or spatial-coordinate-to-intensity projection, while we use voxel-coordinate-to-time-series, which avoids the high complexity of the former and the slow speed of the latter.

High-Dimensional Medical Data Compression. There is a lack of adequate compression methods for four-dimensional medical data. The technical routes of these algorithms can be roughly divided into two categories. The first category is based on motion compensation, which uses motion vectors to model the difference between frames to reduce redundancies(Nguyen et al. 2011; Sanchez, Nasiopoulos, and Abugharbieh 2009, 2008a,b). The second category is based on transform, such as wavelet transform and so on.(Lalgudi et al. 2005; Liu and Pearlman 2007; Rajeswari and Rajesh 2009) The transform applied to images or signals is modified to higher dimensions and then applied to medical imaging data. This modification is usually a combination of 1D transform and 3D transform, leading to limited ability to sparsely capture more complex, higher-order discontinuities(Bruylants, Munteanu, and Schelkens 2015).

Method

Mathematical Representation of fMRI Data

The fMRI data $\mathbf{X}$ can be represented as a set of time series

$$\mathbf{X} = \{\boldsymbol{x_v} | \boldsymbol{v}^{(0)} \in [0, W), \boldsymbol{v}^{(1)} \in [0, H), \boldsymbol{v}^{(2)} \in [0, D)\}. \tag{1}$$

Here $\boldsymbol{v} = [\boldsymbol{v}^{(0)}, \boldsymbol{v}^{(1)}, \boldsymbol{v}^{(2)}]$ represents the 3D spatial coordinate of voxels, and $\boldsymbol{x_v} \in \mathbb{R}^T$ represents the fMRI signal time series at the corresponding spatial location

$$\boldsymbol{x_v} = [\boldsymbol{x}_{\boldsymbol{v}}^{(0)}, \boldsymbol{x}_{\boldsymbol{v}}^{(1)}, \cdots, \boldsymbol{x}_{\boldsymbol{v}}^{(T-1)}], \tag{2}$$

with W, H, D, T respectively denoting the width, height, depth, and length of the time series in the fMRI data.

The principle of localization in brain function organization suggests that brain functions are carried out in a set of brain regions (McKeown et al. 1998). This implies that the brain space can be divided into several brain regions based on function, and there are similarities in neuronal activation within each brain region. Therefore, in fMRI data, the set of time series for each brain region can be represented as

$$\mathbf{A}_i = \{\boldsymbol{x_v} | \boldsymbol{v} \in \mathbf{V}_i\}, i \in [0, N), \tag{3}$$

in which $\mathbf{V}_i$ represents the set of three-dimensional coordinates for the brain region labeled as i, and N represents the number of brain region labels. Therefore, we have

$$\mathbf{X} = \bigcup_{i=0}^{N-1} \mathbf{A}_i, \tag{4}$$

i.e., $\mathbf{V}$ can be represented as the union of $\mathbf{A}_i$.

Modeling Intra-region Correlations

The high similarities within the neural activation in a local brain region tell that $\boldsymbol{x_v}$ in the same $\mathbf{V}_i$ are highly correlated, with the region $\mathbf{V}_i$ being a continuous brain volume. In other words, if we model the fMRI data as the mapping from spatial coordinates $\boldsymbol{v}$ to the time series $\boldsymbol{x_v}$

$$F(\boldsymbol{v}) \to \boldsymbol{x_v}, \tag{5}$$

the local correlations of the function $F(\cdot)$ match well with the advantageous performance of INR-based methods in modeling continuous local structures. Combining the mathematical representation of fMRI, our parameterization function for the neural network can be written as

$$f(\boldsymbol{v}|\theta) \to \widetilde{\boldsymbol{x_v}}. \tag{6}$$

Here θ denotes the network parameters, and $\widetilde{\boldsymbol{x_v}}$ represents the prediction of the ground truth $\boldsymbol{x_v}$. The detailed network structure is shown in Fig. 1, which leverages INR's superb capability of modeling internal correlations and can effectively eliminate the redundancies of time series within local brain regions.

Modeling Inter-region Correlations

fMRI Time-series Signal Decomposition. From the studies in (Boynton et al. 1996; Boynton, Engel, and Heeger 2012), the generation process of fMRI time-series signals can be theoretically viewed as the output of a Linear Time-Invariant System (LTI System):

$$\boldsymbol{x}_{\boldsymbol{v}}^{(t)} = H(\sum_{i=0}^{K-1} \boldsymbol{w}_{\boldsymbol{v}}^{(i)} \cdot u_i(t)) = \sum_{i=0}^{K-1} \boldsymbol{w}_{\boldsymbol{v}}^{(i)} \cdot H(u_i(t)). \tag{7}$$

Here, H denotes the system generating fMRI signals, $u_i(t)$ represents the input stimulus, $\boldsymbol{w}_{\boldsymbol{v}}^{(i)}$ refers to the stimulus distribution of signal intensity, and $H(u_i(t))$ is the neuronal activation pattern. Therefore, all $\boldsymbol{x_v}$ can be decomposed into the weighted superposition of several time series which characterize basic neuronal activation patterns. For example, as illustrated in Fig. 2, we extracted three neuronal activation patterns from one fMRI data, whose weighted superposition can represent all $\boldsymbol{x_v}$ in this fMRI data. The weights of these neuronal activation patterns can be mapped to the brain map, which shows the spatial distributions of these patterns.

Inspired by this, as shown in Module #1 of Fig. 1, we have established a learnable matrix $C_{\text{train}} \in \mathbb{R}^{K \times T}$. Each row of C_{train} represents a reusable neuronal activation pattern and is mapped into the feature space through Conv1d blocks.

Spatial Distribution of the Activation Patterns. While modeling several reusable neuronal activation patterns, as shown in Module #2 of Fig. 1, we use INR groups to model their spatial distributions to further decrease the redundancies among each distribution

$$g_i(\boldsymbol{v}|\sigma_i) \to \boldsymbol{w}_{\boldsymbol{v}}^{(i)}, i \in [0, K). \tag{8}$$

Here function $g_i(\cdot)$ represents the ith INR network with σ_i being its parameters, which models the distribution of the ith neuronal activation pattern. For each spatial coordinate $\boldsymbol{v}$, we concatenate the outputs of these K INR networks into

$$\boldsymbol{w_v} = [\boldsymbol{w}_{\boldsymbol{v}}^{(0)}, \boldsymbol{w}_{\boldsymbol{v}}^{(1)}, \cdots, \boldsymbol{w}_{\boldsymbol{v}}^{(K-1)}]. \tag{9}$$

The components of $\boldsymbol{w_v} \in \mathbb{R}^K$ represent the spatial distribution values in coordinate $\boldsymbol{v}$ of each activation pattern. In this way, we can further eliminate the redundancies within the distributions of neuronal activities.

Subsequently, we use the Feature Fusion Block to fuse the features of various neuronal activation patterns, as illustrated in Module #3 of Fig. 1. This module is composed of several downsampling and upsampling operations. The internal structures of the Down Block and Up Block are shown in supplementary document. We concatenate the outputs of the Down Block and Up Block, which allows better utilization of the information contained in the outputs of each sampling module(Ronneberger, Fischer, and Brox 2015), ultimately modeling the target signal $\boldsymbol{x_v}$.

Initialization of Activation Patterns by ICA. To improve the convergence of the neural network, it is crucial to provide a suitable initialization for the neuronal activation patterns, i.e. matrix C_{train}. As described in (McKeown et al. 1998), the brain can be divided into local regions, where significant correlations exist among the neuronal dynamics. Drawing inspiration, we use a set of activation patterns to describe the main neuronal dynamics in each region. Specifically, we adopted the ICA algorithm to decompose the original fMRI data into basic time series as the neuronal activation patterns used to initialize the learnable matrix C_{train}, as

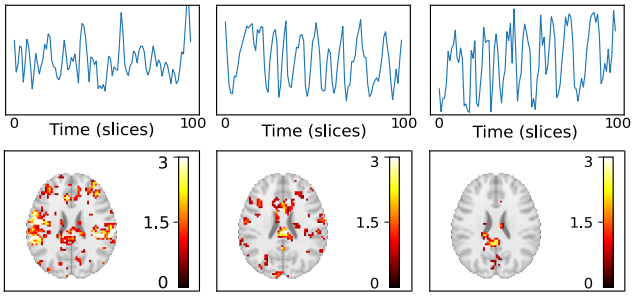

Figure 2: The signal decomposition of an exemplar fMRI time-series. Upper: top 3 neuronal activation patterns. Lower: corresponding spatial distributions.

shown in Module #1 of Fig. 1. The matrix form of ICA can be represented as $\mathbf{Y} = \mathbf{AS}$, which decomposes the original signal matrix $\mathbf{Y}$ into a mixing matrix A and a source matrix $\mathbf{S} = \{\vec{s}_1, \vec{s}_2, \cdots\}$. Correspondingly, applying ICA to fMRI signals gets

$$\boldsymbol{x_v} = \boldsymbol{s_v} \cdot [\boldsymbol{a}_0, \boldsymbol{a}_1, \cdots, \boldsymbol{a}_{K-1}]^T . \tag{10}$$

In this equation, $\boldsymbol{x_v}$ represents the fMRI time series, $\boldsymbol{a}_i \in \mathbb{R}^{T\times 1}$ denotes the activation patterns, with the vector $\boldsymbol{s_v} = [\boldsymbol{s_v}^{(0)}, \boldsymbol{s_v}^{(1)}, \cdots, \boldsymbol{s_v}^{(K-1)}]$ describing their distributions. ICA maximizes the independence among the distributions, allowing the activation patterns to cover the principal neural dynamics in each region. This initialization promises to preserve the dominant information, fulfilling the high fidelity of medical data.

Overview of Network's Workflow

Built on the above formulation, we summarize the workflow of the proposed compression network, as illustrated in Fig. 1. Here we pursue a representative network for each fMRI file for high data fidelity. Considering the high temporal redundancy of the fMRI data, before compression, we subtract the average frame of the target sequence to focus on the sparse among-frame differences reflecting the neuron activities. For the target fMRI data, the intensity and spatial coordinates undertake respective processing before applying Channel Attention. For the intensity, we first mask the black areas and use their ICA coefficients to initialize C_{train}, and then each row vector of C_{train} passes through K Conv1d blocks and the output is concatenated into a feature map F. Meanwhile, the coordinate $\boldsymbol{v}$ is embedded in the same way as in (Mildenhall et al. 2021) and passes through K MLPs, whose output $\boldsymbol{w}^{(i)}$ concatenated into the weight vector $\boldsymbol{w}$. During Channel Attention, $\boldsymbol{w}$ and F are integrated and fed into the Feature Fusion Block, which contains several Down Blocks and Up Blocks respectively for downsampling and upsampling, with structures shown in supplementary document. Finally, a Conv1D block is applied to generate the output vector $\widetilde{\boldsymbol{x_v}}$.

Regarding model training, we employ a weighted combination of L_2 loss and SSIM loss as a loss function, which is formulated as

$$\begin{aligned}\mathcal{L} = &\left(1 - SSIM\left(\{\boldsymbol{x_v}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}, \{\widetilde{\boldsymbol{x_v}}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}\right)\right) * \sigma \\ &+ L_2\left(\{\boldsymbol{x_v}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}, \{\widetilde{\boldsymbol{x_v}}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}\right).\end{aligned} \tag{11}$$

Here $\{\boldsymbol{x_v}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}$ is a batch of training data and $\{\widetilde{\boldsymbol{x_v}}\}_{\boldsymbol{v=a}}^{\boldsymbol{v=b}}$ is their network prediction, with σ being a parameter balancing the contribution of two terms.

Model Compression

Model pruning, quantization, and entropy coding are widely used techniques in model compression(Han, Mao, and Dally 2016). After fitting the target data with the INR network, we proceed to quantize the network parameters and utilize Huffman entropy coding to compress the model further. Here we do not apply model pruning to ensure high compression efficiency.

Experiments and Analysis

This section presents the core experimental results. Please refer to the supplement for ablation study and more results.

Implementation Details

We conducted experiments on four fMRI datasets collected for different downstream tasks: Three datasets are from OpenfMRI[1], an open-source repository for the free and open sharing of fMRI datasets. The fourth dataset is the widely used Haxby dataset, a pioneering study of brain pattern recognition(Haxby et al. 2001), which has extended time series and suitable for the fMRI classification task. And the compute workers we used were four NVIDIA GeForce RTX 3090 with 24 GB memory.

For the first three datasets, in order to facilitate subsequent analysis, we uniformly preprocessed the fMRI data by aligning it to a standard brain template. This resulted in voxel dimensions of 64×64×48 and a time series length of 100. The Haxby dataset comes with a pre-matched mask. And we did not perform registration, as it is unnecessary for the classification task. Due to its longer time series, we addressed CUDA memory limitations by slicing the data in Haxby.

In the experiments, our MLP network was configured with five layers, and the frequency of the coordinate embedding was set to 10. We used JPEG to compress the mean of the data, and utilized the Adamax optimizer with an initial learning rate of 8e-4. The training epoch was set to 1500. The ICA number K was set to 15∼25 empirically. We also utilized ICA's output network pre-training. The compression ratio of the network can be adjusted by varying parameters such as the number of ICA components, average frame compression quality, count of the MLP parameters, number of convolution channels and network layers. All these parameters can be adjusted using stored YAML files.

[1]These three datasets were obtained from the OpenfMRI database, with accession numbers ds000007, ds000101, ds00102, respectively.

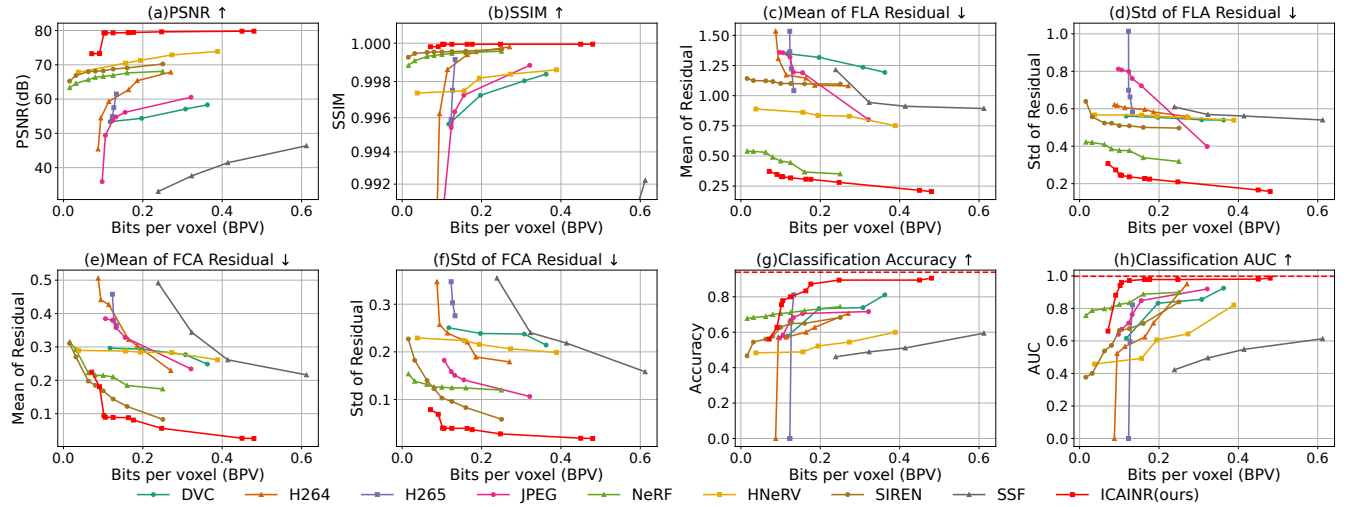

Figure 3: The performance of the proposed approach and the benchmark algorithms.

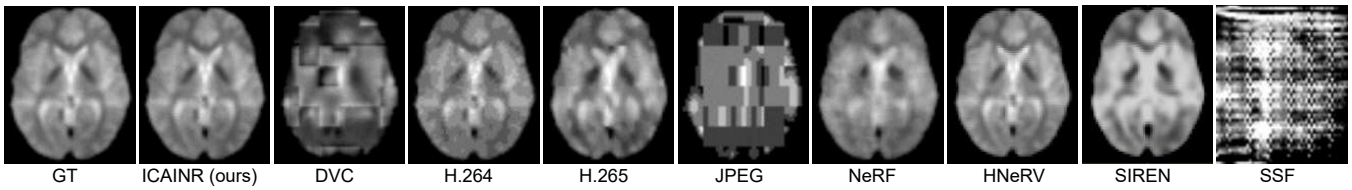

Figure 4: The fMRI slices from the ground truth and decompressed versions of different algorithms. Here the size of the whole fMRI data volume is 64 × 64 lateral pixels, 48 layers, and 100 frames, while the slices we picked are the central slice.

Benchmark Methods and Evaluation Metrics

To demonstrate the advantages of the proposed approach, we comprehensively compare it with eight state-of-the-art (SOTA) algorithms, which can be broadly categorized into three groups: widely used traditional compression algorithms like JPEG(Wallace 1992), H.264(Wiegand et al. 2003), and H.265 (HEVC)(Sullivan et al. 2012), data-driven deep learning compression algorithms such as SSF(Agustsson et al. 2020) and DVC(Lu et al. 2019), and implicit neural representation-based compressors, including HNeRV(Chen et al. 2023), NeRF(Mildenhall et al. 2021), and SIREN(Sitzmann et al. 2020). For the competitors designed for videos, we compressed the 4D fMRI data slice by slice, and for NeRF, we replaced it with ReLU-MLP with Positional Encoding and simplified the input to 4D coordinates since rendering is unnecessary for compression.

We utilized the OpenCV implementation of JPEG and the FFmpeg implementation of H.264 and H.265. We set the compression ratio by calculating the corresponding bit rate. For DVC and SSF, which are data-driven methods, we changed the compression ratio by specifying the quality parameters and fine-tuning the pretrained models provided by their authors. For SIREN and NeRF, we set the MLP layer numbers to 7, and for HNeRV, we specified the network parameters to achieve different compression ratios.

The evaluation metrics employed in our experiments can be divided into two parts. The first part involves traditional image quality evaluation metrics PSNR and SSIM, which are compared across various compression ratios. The second part pertains to downstream tasks based on fMRI data. We selected three downstream tasks, which are all classical and widely used methods in fMRI analysis to comprehensively evaluate the compression quality and fidelity. These downstream tasks encompass the following:

i) General linear model First Level Analysis (FLA): This task involves fitting and hypothesis testing of fMRI waveforms to compute the strength of association between specific stimuli/tasks and various brain regions(Friston et al. 1994), which is the most widely used in fMRI statistical parametric mapping(Smith 2004). We conduct FLA on three datasets and use the mean absolute difference between the statistical maps obtained from the original and compressed data for model evaluation.

ii) Brain regions Functional Connectivity Analysis (FCA): The brain connectome characteristics have offered valuable insights to explain the diversity of pathological conditions and behaviors across different peoples(Mohanty et al. 2020). It is defined as the correlation coefficients between the voxel waveforms of different brain areas(Salvador et al. 2005). In this task, the same dataset was used as in the FLA task, with the MSDL atlas template provided in nilearn(Varoquaux et al. 2011). Again, we use the mean absolute difference between the correlation matrices from the ground truth and decompressed fMRI for performance evaluation.

Method	ICAINR(ours)	H.264	H.265	JPEG	NeRF	HNeRV	SIREN	SSF	DVC
Compression Ratio↑	127.83×*	97.49×	125.99×#	102.95×	99.77×	102.71×	99.76×	66.99×	81.13×
PSNR(dB)↑	79.31*	65.37	61.49	56.14	67.63	70.49#	69.09	33.03	57.09
1 - SSIM↓	5.54E-5*	1.96E-3	8.38E-4	2.74E-3	4.83E-4	2.62E-3	4.01E-4#	9.39E-2	4.39E-4
Mean of FLA Residual↓	0.32*	1.15	1.04	1.19	0.37#	0.86	1.10	1.21	1.32
Std of FLA Residual↓	0.24*	0.60	0.58	0.72	0.34#	0.57	0.50	0.61	0.55
Mean of FCA Residual↓	0.09*	0.32	0.37	0.33	0.18	0.29	0.12#	0.50	0.29
Std of FCA Residual↓	0.04*	0.22	0.28	0.14	0.12	0.22	0.08#	0.35	0.24

Table 1: The performance of different compressors at a specific compression ratio, roughly 100×. Methods marked with the * suffix are the best-performing methods, while those marked with the # suffix are the second-best-performing methods.

iii) fMRI-based Classification Task Analysis (CTA): Decoding or pattern recognition techniques are a significant part of fMRI analysis(Haxby et al. 2001). We use a linear classifier based on SVM(De Martino et al. 2008), which is most commonly used in fMRI classification(Naselaris et al. 2011), to do our experiment. In this task, we employed the Haxby dataset with various kinds of images serving as stimuli. We trained the SVM with the voxel waves, to distinguish stimuli images of the house and face. We applied 10-fold cross-validation and used the cross-validation classification accuracy of the decompressed fMRI as the evaluation metric.

Experiment Results

PSNR & SSIM. We first validate the data fidelity after fMRI compression and compare our approach against existing compression algorithms. Here we use PSNR and SSIM as quantitative evaluation metrics, and provide the scores of different algorithms across varying compression ratios, as plotted in Fig. 3a and Fig. 3b. Notably, under similar compression ratios, our method consistently outperforms existing SOTA algorithms. When changing the compression ratio, some algorithms fluctuate, in contrast to the remarkably stable performance of our approach.

Visual Results. The visual comparison of the data decompressed by various algorithms in a fixed compression ratio (∼100×) is shown in Fig. 4. Note that the achieved compression ratio of different algorithms differs slightly, as shown in Tab. 1, because one cannot specify the final compression ratio exactly. The results show that despite yielding relatively high PSNR and SSIM values, many algorithms failed to deliver satisfactory visual quality. For instance, JPEG and DVC suffer from block effect(Lee, Kim, and Park 1998), with noticeable fragmentation between adjacent image blocks. Some INR-based algorithms, like NeRF and SIREN, sacrificed a considerable amount of high-frequency details, thus showing over-smoothness. In contrast, other algorithms such as H.264 and H.265 exhibited noticeable noise, and SSF almost failed to model the fMRI data.

FLA Results. To evaluate the behaviors of our compressor and other competitors in FLA, we calculated the mean and standard deviation of the residue between the compressed fMRI data and the ground-truth FLA results, as plotted in Fig. 3c and Fig. 3d. The curves indicate that our method has the smallest fluctuation in residue and performance in all compression ratios among these algorithms. Furthermore, we display the visualized FLA obtained by various algorithms at a compression ratio of about ∼100×, as illustrated in Fig. 5. Our method exhibits high similarity to the ground truth, while others show significant differences.

FCA Results. Regarding the influence of compression on FCA, similar to FLA, we illustrate the mean and standard deviation of residual after decompression in Fig. 3e and Fig. 3f. The curves show that our method has much less information loss than other algorithms if the compression ratio is lower than 170×, and the average loss approximates zero when the compression ratio is around 30∼40×. From the standard deviation, one can conclude that our approach exhibits the lowest performance fluctuation among all algorithms working at a similar compression ratio and achieves stable performance (standard deviation close to 0) when the compression ratio is lower than $100times$. We also compare the visualized brain connectivity at a compression ratio of about 100× in Fig. 6, from which one can observe that our compressor presents the highest similarity to ground truth.

CTA Results. To test the effectiveness of our compressor and its advantage over previous competitors in the successive classification task (CTA), we calculated the 10-fold cross-validation accuracy and AUC of the decompressed data compared to the original data before compression, as shown in Fig. 3g and Fig. 3h. The results demonstrate that our algorithm outperforms other methods in compression ratios lower than 170×, and can achieve accuracy and AUC close to ground truth, 93. 89% and 0.9975, respectively, when the compression ratio falls below 100×.

In summary, our algorithm presents superior results in terms of both PSNR and SSIM, as well as impressive scores in downstream tasks, while some previous algorithms focus on PSNR and SSIM but suffer from degraded performance in successive analyzes. This robust performance substantiates our algorithm's capacity to effectively retain crucial information in fMRI, providing the necessity to take into account the physical meaning of fMRI data during compression and presenting the potential to perform transmission, processing, and analysis on fMRI data at low bandwidth.

Summary and Discussions

In this paper, we introduce a novel INR-based compression paradigm for fMRI data. By leveraging the strong representation capability of deep neural networks and taking into account of fMRI's unique characteristics, we aim to compactly describe the neural activation patterns and their spatial distributions, which reduces both temporal and spatial redundancies in the raw recordings.

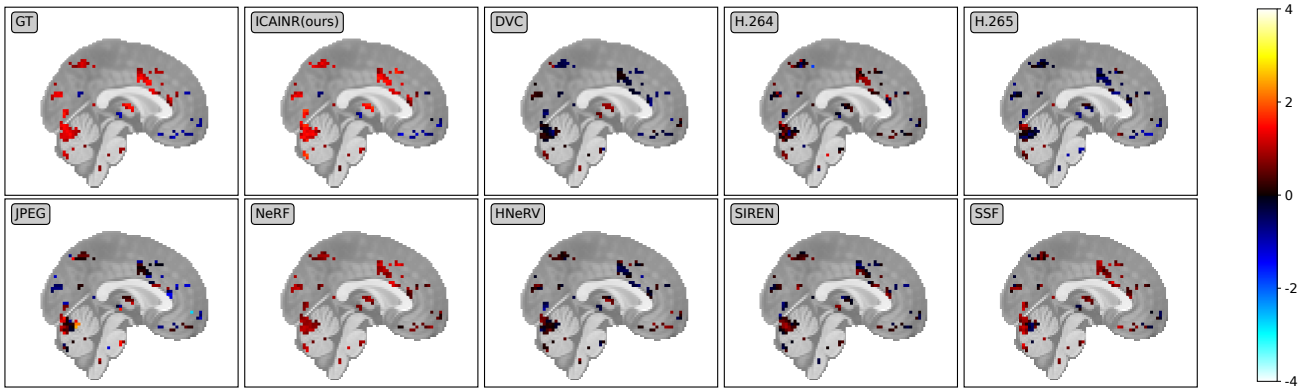

Figure 5: The brain t-score map calculated from the ground truth and decompressed results by different algorithms. We set the confidence level to be 90% to label the voxels which are highly correlated with the stimuli. Here the displayed intersecting plane of the brain map is the center frame along y axis.

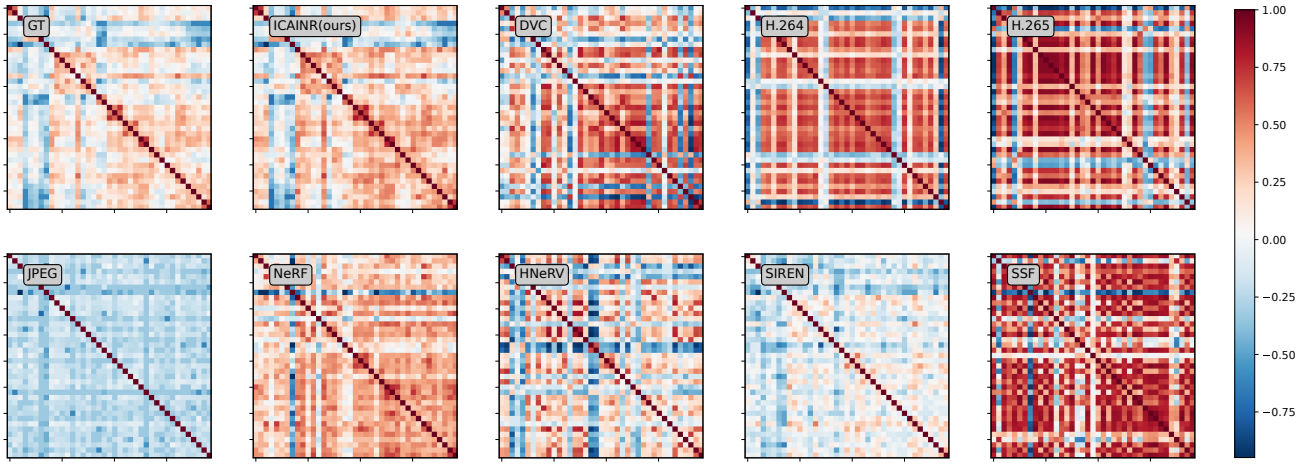

Figure 6: The brain region connectivity matrix involving 39 brain regions calculated from ground truth and decompressed results of all algorithms. Each matrix entry is calculated as the correlation coefficients between two corresponding regions.

Advantages. Based on the comprehensive experiments conducted, our algorithm outperforms existing SOTAs in terms of data fidelity such as PSNR and SSIM, as well as evaluation metrics for downstream fMRI tasks, which meets the fidelity requirements for medical image compression. Notably, our study pioneers the application of INR-based compression methods to four-dimensional biomedical data. Additionally, we present a novel INR-based compression framework, offering valuable insights for future research in INR-based biomedical image compression.

Limitations and Future Work. As a preliminary research on INR-based fMRI compression, our approach can be extended in several aspects before becoming a mature compression tool. Firstly, as a deep learning-based compression algorithm, our approach takes longer time than conventional compressors. Therefore, we are exploring novel techniques, such as meta-learning, for fast model learning. Secondly, we use all the initial components obtained from ICA, but some of them contain little information about the neural activities. We are also working on strategies for extracting informative signals to avoid compressing unnecessary components and thus build a more compact compression. Thirdly, we adopt a widely used method for network compression, and there is room for designing more tailored model compression techniques for our network in the future.

Acknowledgments

This work is jointly funded by Ministry of Science and Technology of China (Grant No. 2024YFF0505703) and National Natural Science Foundation of China (Grant Nos. 61931012 and 62088102).

References

Agustsson, E.; Minnen, D.; Johnston, N.; Balle, J.; Hwang, S. J.; and Toderici, G. 2020. Scale-Space Flow for End-to-End Optimized Video Compression. In *Proceedings of*

the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8503–8512.

Boynton, G. M.; Engel, S. A.; Glover, G. H.; and Heeger, D. J. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, 4207–4221.

Boynton, G. M.; Engel, S. A.; and Heeger, D. J. 2012. Linear systems analysis of the fMRI signal. *NeuroImage*, 975–984.

Bruylants, T.; Munteanu, A.; and Schelkens, P. 2015. Wavelet based volumetric medical image compression. *Signal Processing: Image Communication*, 31: 112–133.

Chen, H.; Gwilliam, M.; Lim, S.-N.; and Shrivastava, A. 2023. Hnerv: A hybrid neural representation for videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10270–10279.

Chen, H.; He, B.; Wang, H.; Ren, Y.; Lim, S. N.; and Shrivastava, A. 2021. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34: 21557–21568.

Damodaran, B. B.; Balcilar, M.; Galpin, F.; and Hellier, P. 2023. RQAT-INR: Improved Implicit Neural Image Compression. In *2023 Data Compression Conference*, 208–217.

De Martino, F.; Valente, G.; Staeren, N.; Ashburner, J.; Goebel, R.; and Formisano, E. 2008. Combining multivariate voxel selection and support vector machines for mapping and classification of fMRI spatial patterns. *NeuroImage*, 43(1): 44–58.

Diedrichsen, J.; and Shadmehr, R. 2005. Detecting and adjusting for artifacts in fMRI time series data. *NeuroImage*, 27(3): 624–634.

Dupont, E.; Goliński, A.; Alizadeh, M.; Teh, Y. W.; and Doucet, A. 2021. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*.

Friston, K. J.; Holmes, A. P.; Worsley, K. J.; Poline, J.-P.; Frith, C. D.; and Frackowiak, R. S. 1994. Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping*, 2(4): 189–210.

Han, S.; Mao, H.; and Dally, W. J. 2016. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. arxiv:1510.00149.

Haxby, J. V.; Gobbini, M. I.; Furey, M. L.; Ishai, A.; Schouten, J. L.; and Pietrini, P. 2001. Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539): 2425–2430.

Heil, C. E.; and Walnut, D. F. 1989. Continuous and discrete wavelet transforms. *SIAM Review*, 31(4): 628–666.

Krohn, S.; von Schwanenflug, N.; Waschke, L.; Romanello, A.; Gell, M.; Garrett, D. D.; and Finke, C. 2023. A spatiotemporal complexity architecture of human brain activity. *Science Advances*, 9(5): eabq3851.

Lalgudi, H. G.; Bilgin, A.; Marcellin, M. W.; Tabesh, A.; Nadar, M. S.; and Trouard, T. P. 2005. Four-Dimensional Compression of fMRI Using JPEG2000. In *Medical Imaging 2005: Image Processing*, volume 5747, 1028–1037.

Lee, Y.; Kim, H.; and Park, H. 1998. Blocking effect reduction of JPEG images by signal adaptive filtering. *IEEE Transactions on Image Processing*, 7(2): 229–234.

Li, Z.; Wang, M.; Pi, H.; Xu, K.; Mei, J.; and Liu, Y. 2022. E-NeRV: Expedite Neural Video Representation with Disentangled Spatial-Temporal Context. arxiv:2207.08132.

Liu, Y.; and Pearlman, W. A. 2007. Four-dimensional wavelet compression of 4-D medical images using scalable 4-D SBHP. In *2007 Data Compression Conference*, 233–242. Ieee.

Lu, G.; Ouyang, W.; Xu, D.; Zhang, X.; Cai, C.; and Gao, Z. 2019. Dvc: An End-to-End Deep Video Compression Framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11006–11015.

McKeown, M. J.; Makeig, S.; Brown, G. G.; Jung, T.-P.; Kindermann, S. S.; Bell, A. J.; and Sejnowski, T. J. 1998. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, 6(3): 160–188.

Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing Scenes as Neural Radiance Fields for View Synthesis. *Communications of the ACM*, 65(1): 99–106.

Mohanty, R.; Sethares, W. A.; Nair, V. A.; and Prabhakaran, V. 2020. Rethinking measures of functional connectivity via feature extraction. *Scientific Reports*, 10(1): 1298.

Naselaris, T.; Kay, K. N.; Nishimoto, S.; and Gallant, J. L. 2011. Encoding and decoding in fMRI. *NeuroImage*, 56(2): 400–410.

Nguyen, B. P.; Chui, C.-K.; Ong, S.-H.; and Chang, S. 2011. An Efficient Compression Scheme for 4-D Medical Images Using Hierarchical Vector Quantization and Motion Compensation. *Computers in Biology and Medicine*, 41(9): 843–856.

Rajeswari, R.; and Rajesh, R. 2009. Efficient compression of 4D fMRI images using bandelet transform and fuzzy thresholding. In *2009 World Congress on Nature & Biologically Inspired Computing*, 543–547. Ieee.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, 234–241.

Salvador, R.; Suckling, J.; Coleman, M. R.; Pickard, J. D.; Menon, D.; and Bullmore, E. 2005. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*, 15(9): 1332–1342.

Sanchez, V.; Nasiopoulos, P.; and Abugharbieh, R. 2008a. Efficient 4D Motion Compensated Lossless Compression of Dynamic Volumetric Medical Image Data. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, 549–552.

Sanchez, V.; Nasiopoulos, P.; and Abugharbieh, R. 2008b. Efficient Lossless Compression of 4-D Medical Images Based on the Advanced Video Coding Scheme. *IEEE Transactions on Information Technology in Biomedicine*, 12(4): 442–446.

Sanchez, V.; Nasiopoulos, P.; and Abugharbieh, R. 2009. Novel Lossless fMRI Image Compression Based on Motion Compensation and Customized Entropy Coding. *IEEE Transactions on Information Technology in Biomedicine*, 13(4): 645–655.

Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33: 7462–7473.

Smith, S. M. 2004. Overview of fMRI analysis. *The British Journal of Radiology*, 77(suppl_2): S167–s175.

Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the High Efficiency Video Coding (HEVC) Standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12): 1649–1668.

Van Den Heuvel, M. P.; and Pol, H. E. H. 2010. Exploring the brain network: a review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8): 519–534.

Varoquaux, G.; Gramfort, A.; Pedregosa, F.; Michel, V.; and Thirion, B. 2011. Multi-Subject Dictionary Learning to Segment an Atlas of Brain Spontaneous Activity. In *Information Processing in Medical Imaging: 22nd International Conference 2011, Kloster Irsee, Germany, July 3-8, 2011. Proceedings 22*, 562–573.

Wallace, G. K. 1992. The JPEG Still Picture Compression Standard. *IEEE Transactions on Consumer Electronics*, 38(1): xviii–xxxiv.

Wiegand, T.; Sullivan, G. J.; Bjontegaard, G.; and Luthra, A. 2003. Overview of the H. 264/AVC Video Coding Standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7): 560–576.

Yang, R. 2023. TINC: Tree-structured Implicit Neural Compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18517–18526.

Yang, R.; Xiao, T.; Cheng, Y.; Cao, Q.; Qu, J.; Suo, J.; and Dai, Q. 2023. Sci: A Spectrum Concentrated Implicit Neural Compression for Biomedical Data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4774–4782.