arXiv:2506.09949v1 [eess.IV] 11 Jun 2025

Sampling Theory for Super-Resolution with Implicit Neural Representations

Mahrokh Najaf, Gregory Ongie*

June 12, 2025

Abstract

Implicit neural representations (INRs) have emerged as a powerful tool for solving inverse problems in computer vision and computational imaging. INRs represent images as continuous domain functions realized by a neural network taking spatial coordinates as inputs. However, unlike traditional pixel representations, little is known about the sample complexity of estimating images using INRs in the context of linear inverse problems. Towards this end, we study the sampling requirements for recovery of a continuous domain image from its low-pass Fourier samples by fitting a single hidden-layer INR with ReLU activation and a Fourier features layer using a generalized form of weight decay regularization. Our key insight is to relate minimizers of this non-convex parameter space optimization problem to minimizers of a convex penalty defined over an infinite-dimensional space of measures. We identify a sufficient number of Fourier samples for which an image realized by an INR is exactly recoverable by solving the INR training problem. To validate our theory, we empirically assess the probability of achieving exact recovery of images realized by low-width single hidden-layer INRs, and illustrate the performance of INRs on super-resolution recovery of continuous domain phantom images.

1 Introduction

Many inverse problems in signal and image processing are naturally posed in continuous domain. For example, in magnetic resonance imaging, measurements are modeled as samples of the Fourier transform of a function defined over continuous spatial coordinates [17]. Often these continuous domain inverse problems are approximated using discrete image or signal representations, along with an appropriate discretization of the forward model, such as the discrete Fourier transform in place of the continuous Fourier transform. However, this can lead to undesirable discretization artifacts, such as Gibb's ringing [1, 51, 34, 35]. Such

*Department of Mathematical and Statistical Sciences, Marquette University, Milwaukee, WI, USA. Corresponding author: Gregory Ongie (gregory.ongie@marquette.edu)

artifacts may be reduced by increasing the resolution of the discretization grid, but at the cost of additional computational complexity and potentially making the problem ill-posed without additional regularization.

Recently, continuous domain image representations using neural networks have gained traction in a variety of inverse problems arising in computer vision and computational imaging (see, e.g., the recent review [54]). These so-called *implicit neural representations* (INRs), also variously known as *coordinate-based neural networks* or *neural radiance fields*, parameterize a signal or image as a neural network defined over continuous spatial coordinates [27, 48, 49]. The main benefits of INRs over discrete representations are that (1) they offer a compact (non-linear) representation of images with potentially fewer parameters than a discrete representation, (2) they allow for more accurate continuous domain forward modeling, which can reduce the impact of discretization artifacts, and (3) standard neural network training can be used to fit an INR to measurements, sidestepping the need for custom iterative solvers.

However, the fundamental limits of image estimation using INRs have not been explored in detail. One basic unanswered question is: How many linear measurements are necessary and sufficient for accurate reconstruction of an image by fitting an INR? Additionally, how does the architecture of the INR influence the recovered image? As a step in the direction of answering these questions, we consider the recovery of an image from its low-pass Fourier samples using INRs. In particular, we focus on a popular INR architecture originally introduced in [49] that combines a fully connected neural network with a Fourier features layer.

Recent work has shown that training a single hidden-layer ReLU network with ℓ^2 regularization on its parameters (also known as "weight decay" regularization) imparts a sparsity effect, such that the network learned has low effective width [44, 36, 37, 38]. Bringing this perspective to INRs, we study the conditions under which an image realizable by a single-hidden layer, low-width INR is exactly recoverable from low-pass Fourier samples by minimizing a generalized weight decay objective. Our main theoretical insight is to relate minimizers of this non-convex parameter space optimization problem to minimizers of a convex optimization problem posed over a space of measures.

Using this framework, we prove two sampling theorems. First, we identify a sufficient number of low-pass Fourier samples such that an image realizable by a width-1 INR is the unique minimizer of the INR training problem regularized with a modified form of weight decay. Additionally, under more restrictive assumptions on the frequencies used to define the Fourier features layer, we extend the result to images realizable by a general width-s INR.

Finally, we validate our theory with a series of experiments. First, we demonstrate that random images generated by low-width INRs are exactly recoverable from low-pass Fourier samples by solving the generalized weight decay regularized INR training problem. Next, we illustrate the performance of INRs on super-resolution recovery of continuous domain phantom images, highlighting the effects of generalized weight regularization in practical INR training. Lastly, we empirically explore the role of depth of INRs for super-resolution

recovery. We show that deeper INRs trained with weight decay regularization perform better at recovering piecewise constant phantom images, suggesting that deep INRs may exhibit a bias towards piecewise constant representations.

Preliminary versions of the results in this paper were presented in [29]. The majority of proofs were omitted from [29], which are included here. Additionally, this paper expands the sampling theory to the general width-s INR case, and contains an expanded experiments section.

1.1 Related Work

Closely related to the Fourier sampling model studied in this work, several papers have explored INRs for MRI reconstruction, including super-resolution imaging [53, 52, 50]. Other applications of INRs to MRI reconstruction include: multi-contrast imaging [26], dynamic imaging [21, 16, 23], and compressed sensing with prior-image constraints [45].

While this work focuses on one type of INR architecture, a wide variety of other INR architectures have been proposed. Early INR architectures focused on fully connected ReLU networks with an initial *positional encoding* layer, inspired by transformer architectures [27]. Later work investigated alternative positional encoding layers given by random Fourier features [49] (the focus of this work), and a multiresolution hash encoding [28]. Other approaches to INR architecture design do not use an initial positional encoding layer, but instead consider fully connected networks with non-standard activation functions, including periodic [48] or variable-periodic [25] functions, Gaussian functions [41], sinc functions [43], and wavelet-inspired activation functions [42, 46]. See the recent review [15] for a detailed comparison of various popular INR architectures.

Several recent works have also sought to mathematically characterize the expressive power and inductive bias of INRs. A neural tangent kernel (NTK) analysis of INRs is given in [49], which shows that a tunable random Fourier features layer can help alleviate the low spatial-frequency spectral bias associated with fully connected ReLU networks. Relatedly, [55] studies the expressive power of a variety of INR architectures (including those proposed in [48, 49]) through a generalized Fourier series analysis. Their analysis shows that only functions given by a linear combination of certain harmonics of the initial feature map, with additional structural constraints on the coefficients, are realizable with a given INR architecture. Follow-up work expands on the Fourier series perspective specifically for INRs with sinusoidal activations [31] and its implications for robust INR initialization [32]. Additionally, recent work investigates incorporating batch normalization in INR layers as a means to reduce low spatial-frequency spectral bias of INR architectures [7].

Most closely related to this work, [46] introduces a means to measure function space regularity of a shallow INR representation using the *variation norm*, which is the function space regularizer induced by weight decay regularization. Similarly, in this work we consider regularizing INRs by quantities analogous to the variation norm. However, the present work goes beyond [46] in proving super-resolution recovery guarantees for images realized by low-width INRs.

Finally, the proof techniques used in this work draw from the "off-the-grid" compressed

sensing literature. This includes works focused on super-resolution of sparse measures [5, 9, 8, 10] and multidimensional generalizations [40, 12, 39, 24]. In particular, the image model investigated in this work is closely related to (though distinct from) the continuous-domain piecewise constant image model investigated in [34, 35, 33]. Algebraic facts regarding the zero-sets of two-variable trigonometric polynomials established in [35] are central to key proofs in this work; higher-dimensional extensions of these results have recently been explored in [56].

2 Problem formulation

We consider the recovery of a d-dimensional real-valued continuous domain signal/image $f : [0,1]^d \to \mathbb{R}$ from given a finite sample of its Fourier coefficients

$$\hat{f}[\boldsymbol{k}] = \int_{[0,1]^d} f(\boldsymbol{x}) e^{-2\pi \mathrm{i} \boldsymbol{k}^\top \boldsymbol{x}} d\boldsymbol{x},$$

where $\boldsymbol{k} \in \mathbb{Z}^d$ is any d-tuple of integer spatial frequencies. In particular, suppose we sample frequencies belonging to a uniformly sampled square region in frequency domain $\Omega = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_\infty \leq K\}$ where K is a positive integer. Let $\mathcal{F}_\Omega f = (\hat{f}[\boldsymbol{k}])_{\boldsymbol{k}\in\Omega}$ denote the vector of Fourier coefficients of f restricted to Ω. We study the inverse problem of recovering the function f given $\boldsymbol{y} = \mathcal{F}_\Omega f$.

Clearly, this is an ill-posed inverse problem if no further assumptions are made on f. One approach to resolve this ill-posedness is to assume f is well-approximated by a function realized by a "simple" neural network $f_\theta : [0,1]^d \to \mathbb{R}$, i.e., an INR, with trainable parameters θ. Under this assumption, one can attempt to recover f from its measurements $\boldsymbol{y}$ by solving the (non-linear) least squares problem

$$\min_\theta \tfrac{1}{2} \|\mathcal{F}_\Omega f_\theta - \boldsymbol{y}\|_2^2. \tag{1}$$

However, this problem may still be ill-posed (i.e., (1) may have multiple solutions) unless the number of trainable parameters used to define the INR f_θ is heavily constrained. Though, even without constraining the number of parameters, useful solutions of (1) may be found by relying on implicit regularization induced by practical gradient-based algorithms [20].

As an alternative to using an overly constrained network architecture or relying on implicit regularization to resolve ill-posedness, we propose incorporating an explicit parameter space regularizer $R(\theta)$ to the least squares objective:

$$\min_\theta \tfrac{1}{2} \|\mathcal{F}_\Omega f_\theta - \boldsymbol{y}\|_2^2 + \lambda R(\theta), \tag{2}$$

where $\lambda > 0$ is a regularization parameter. In particular, we focus on a class of regularizers $R(\theta)$ that generalize *weight decay regularization* [22], the squared ℓ^2-norm of all trainable parameters; this class is described in more detail below in Section 2.2.

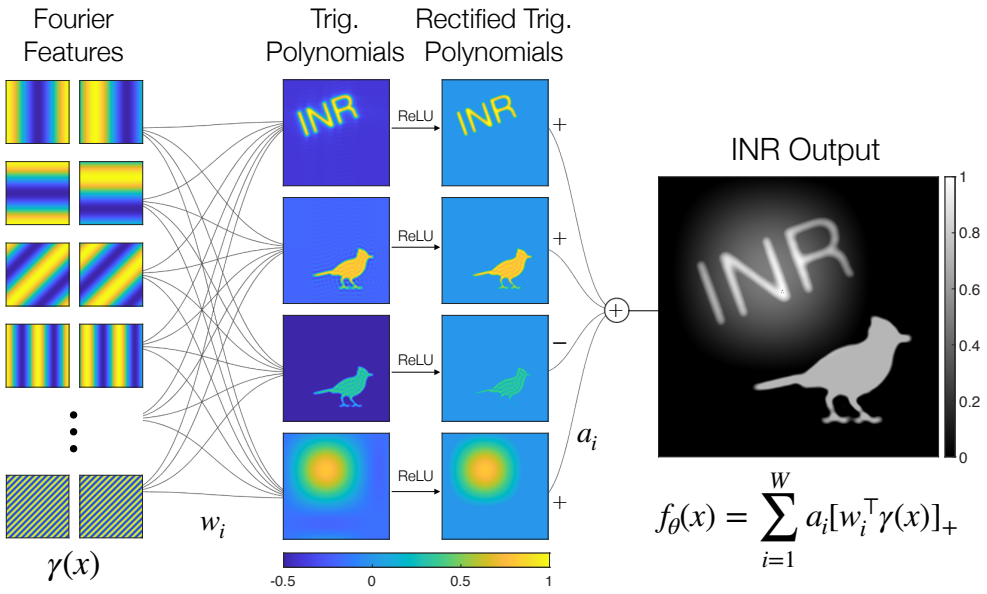

Figure 1: The shallow INR architecture considered in this work represents an image as a linear combination of rectified trigonometric polynomials. We study the sample complexity of estimating an image of this type from its low-pass Fourier samples.

In a noise-free scenario, the regularization parameter λ in (2) needs to be very small to ensure data consistency is tightly enforced. As a model for this situation, we will focus on the equality constrained problem

$$\min_{\theta} R(\theta) \quad s.t. \quad \mathcal{F}_\Omega f_\theta = \boldsymbol{y}, \tag{3}$$

which can be thought of as the limiting case of (2) as $\lambda \to 0$.

Our main goal is to characterize global minimizers of (3). In particular, we are interested in its *function space minimizers*, i.e., functions f_{θ^*} where θ^* is a global minimizer of (3). Given $\boldsymbol{y} = \mathcal{F}_\Omega f$ where f is realizable as an INR, we ask: when is f the unique function space minimizer of (3)? Next, we describe the architectural assumptions we put on the INR f_θ and the associated regularizer $R(\theta)$.

2.1 INR Architecture

In this work, we focus on an INR architecture first proposed in [49] that combines a trainable fully connected network with an initial Fourier features layer. For simplicity, we primarily consider a "shallow" INR architecture, where the fully connected network has a single-hidden layer with ReLU activation.

Specifically, consider a shallow ReLU network $g_\theta : \mathbb{R}^D \to \mathbb{R}$ of width-W and input dimension D given by

$$g_\theta(\boldsymbol{v}) = \sum_{i=1}^{W} a_i[\boldsymbol{w}_i^\top \boldsymbol{v}]_+, \tag{4}$$

where $[t]_+ := \max\{0, t\}$ is the ReLU activation function, and $\theta = ((a_i, \boldsymbol{w}_i))_{i=1}^W$ denotes the vector of all trainable parameters with $a_i \in \mathbb{R}$ and $\boldsymbol{w}_i \in \mathbb{R}^D$ for all $i \in [W]$. We denote the resulting parameter space by $\Theta_W \simeq \mathbb{R}^{W(D+1)}$. By choosing $D = d$, this architecture could be used as an INR directly. However, as shown in [49], it is challenging to recover high-frequency content in images using such an architecture.

Instead, to enhance recovery of high-frequency content, [49] proposed adding an initial Fourier features layer to the architecture: given a collection of non-zero frequency vectors $\Omega_0 = \{\boldsymbol{k}_1, ..., \boldsymbol{k}_p\} \subset \mathbb{R}^d$ define the Fourier features embedding $\boldsymbol{\gamma} : \mathbb{R}^d \to \mathbb{R}^D$ with $D = 2p + 1$ by

$$\boldsymbol{\gamma}(\boldsymbol{x}) = [1, \sqrt{2}\cos(2\pi \boldsymbol{k}_1^\top \boldsymbol{x}), ..., \sqrt{2}\cos(2\pi \boldsymbol{k}_p^\top \boldsymbol{x}), \sqrt{2}\sin(2\pi \boldsymbol{k}_1^\top \boldsymbol{x}), ..., \sqrt{2}\sin(2\pi \boldsymbol{k}_p^\top \boldsymbol{x})]^\top. \quad (5)$$

This gives the shallow INR architecture:

$$f_\theta(\boldsymbol{x}) = g_\theta(\boldsymbol{\gamma}(\boldsymbol{x})) = \sum_{i=1}^W a_i [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+. \quad (6)$$

To better understand the impact of the Fourier features embedding, for any weight vector $\boldsymbol{w} \in \mathbb{R}^D$, let us write $\boldsymbol{w} = [w^{(0)}, \boldsymbol{w}^{(1)}, \boldsymbol{w}^{(2)}]$ where $w^{(0)} \in \mathbb{R}$ and $\boldsymbol{w}^{(1)}, \boldsymbol{w}^{(2)} \in \mathbb{R}^p$. Then defining $\tau(\boldsymbol{x}) := \boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})$, we have

$$\tau(\boldsymbol{x}) = w^{(0)} + \sqrt{2} \sum_{j=1}^p \left(w_j^{(1)} \cos(2\pi \boldsymbol{k}_j^\top \boldsymbol{x}) + w_j^{(2)} \sin(2\pi \boldsymbol{k}_j^\top \boldsymbol{x}) \right). \quad (7)$$

This shows τ is a real-valued *trigonometric polynomial*, i.e., τ is a function with frequency support contained in a finite set. Therefore, a shallow INR with ReLU activation and a Fourier features embedding is a weighted linear combination of *rectified trigonometric polynomials*, i.e., functions of the form $\boldsymbol{x} \mapsto [\tau(\boldsymbol{x})]_+$ where τ is a real-valued trigonometric polynomial; see Figure 1 for an illustration.

Originally, [49] proposed defining the frequency set Ω_0 by sampling a fixed number of normally distributed random vectors $\boldsymbol{k} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I})$, where the variance σ^2 is treated as a tunable parameter. To simplify the mathematical theory, in this work we focus on the case where Ω_0 is a set of uniformly sampled integer-valued frequencies $\boldsymbol{k} \in \mathbb{Z}^d$. This implies that the INR (6) is a periodic function on $[0, 1]^d$, or equivalently, a function defined over the d-dimensional torus $\mathbb{T}^d := (\mathbb{R}/\mathbb{Z})^d$. We note that using uniformly sampled integer frequencies in place of random frequencies for an INR Fourier features layer was also investigated in [3].

2.2 Generalized Weight Decay Regularization

We focus on a family of parameter space regularizers $R(\theta)$ that strictly generalize "weight decay" regularization [22], i.e., the squared ℓ^2-norm of all trainable parameters. Consider any non-negative *weighting function* η defined over inner-layer weight vectors $\boldsymbol{w} \in \mathbb{R}^D$. We say η is an *admissible* weighting function if it satisfies the following properties:

(i) η is continuous,

(ii) η is positive 1-homogeneous: $\eta(\alpha \boldsymbol{w}) = \alpha \eta(\boldsymbol{w})$ for all $\boldsymbol{w} \in \mathbb{R}^d$ and $\alpha \geq 0$,

(iii) $\eta(\boldsymbol{w}) = 0$ implies $[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ = 0$ (or, equivalently, $[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ \neq 0$ implies $\eta(\boldsymbol{w}) > 0$).

Then, given any admissible weighting function η, we define the *generalized weight decay penalty* applied to any parameter vector $\theta = ((a_i, \boldsymbol{w}_i))_{i=1}^W$ by

$$R(\theta) = \frac{1}{2} \sum_{i=1}^{W} \left(|a_i|^2 + \eta(\boldsymbol{w}_i)^2 \right). \tag{8}$$

For example, if $\eta(\boldsymbol{w}) = \|\boldsymbol{w}\|_2$ we recover standard weight decay regularization. For the sampling theory given in Section 3.3, we consider alternative weighting functions $\eta(\boldsymbol{w})$ defined in terms of penalties applied to the corresponding hidden-layer ReLU unit function $[\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$.

3 Theory

Our main theoretical insight is to show that minimizers of the non-convex INR training problem (3) coincide with minimizers of a *convex* optimization problem defined over an infinite-dimensional space of measures. We use this equivalence to prove our main theorems, which identifies a sufficient number of Fourier samples needed to recover a continuous domain image realizable as shallow INR. Proofs of all results in this section are given in the Appendix.

3.1 Generalized Weight Decay Induces Unit Sparsity

Given an admissible weighting function η, and INR training width $W > 0$, the main optimization problem we consider in this work is

$$\min_{\theta \in \Theta_W} \frac{1}{2} \sum_{i=1}^{W} \left(|a_i|^2 + \eta(\boldsymbol{w}_i)^2 \right) \quad \text{s.t.} \quad \mathcal{F}_\Omega f_\theta = \boldsymbol{y}. \tag{$P_{\theta,W}$}$$

First, we show that $(P_{\theta,W})$ is equivalent to optimizing a modified cost function over a restricted set of parameter vectors. In particular, let $\Theta'_W \subset \Theta_W$ denote the subset of parameter vectors $\theta = ((a_i, \boldsymbol{w}_i))_{i=1}^W$ that satisfy the normalization constraint $\|\boldsymbol{w}_i\|_2 = 1$ for all $i \in [W]$. Then we have the following equivalence:

Lemma 1. *Let η be any admissible weighting function. Then function space minimizers of $(P_{\theta,W})$ coincide with function space minimizers of*

$$\min_{\theta \in \Theta'_W} \sum_{i=1}^{W} |a_i| \eta(\boldsymbol{w}_i) \quad \text{s.t.} \quad \mathcal{F}_\Omega f_\theta = \boldsymbol{y}. \tag{$P'_{\theta,W}$}$$

The proof, given in Appendix A.1, follows by an argument similar to the proof of [44, Lemma 1], and relies on the positive 1-homogeneity of the ReLU activation and of the weighting function η. This result shows minimizing a generalized weight decay penalty implicitly minimizes a weighted ℓ^1-norm on the outer-layer weights a_i subject to a normalization constraint on the inner-layer weight vectors $\boldsymbol{w}_i$. Due to the sparsity-promoting property of the ℓ^1-norm, intuitively this result shows that minimizing generalized weight decay penalty should promote functions that are realizable by a "sparse" single-hidden layer INR, i.e., an INR with few active units.

3.2 Infinite-width Convexification

One technical hurdle to characterizing minimizers of $(P'_{\theta,W})$ is that set of parameters $\theta \in \Theta'_W$ satisfying the constraint $\mathcal{F}_\Omega f_\theta = \boldsymbol{y}$ is non-convex. To sidestep this issue, we show how to reformulate $(P'_{\theta,W})$ as a convex optimization problem posed over a space of measures that represents the limit as the hidden-layer width $W \to \infty$. We note that similar convex formulations for single hidden-layer neural networks have been explored in [4, 2, 14].

First, fix any set of normalized INR parameters $\theta = ((a_i, \boldsymbol{w}_i))_{i=1}^W \in \Theta'_W$. By linearity of $\mathcal{F}_\Omega$ we have

$$\mathcal{F}_\Omega f_\theta = \sum_{i=1}^W a_i \mathcal{F}_\Omega [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+. \tag{9}$$

Now we show the sum on the right above is expressible as the integration of a function against a discrete measure, i.e., a linear combination of Dirac deltas. In particular, given any admissible weight function $\eta(\boldsymbol{w})$, define the set

$$S_\eta := \{\boldsymbol{w} \in \mathbb{R}^D : \|\boldsymbol{w}\|_2 = 1, \eta(\boldsymbol{w}) > 0\}, \tag{10}$$

which is an open subset of the $D-1$ dimensional unit sphere, $\mathbb{S}^{D-1} \subset \mathbb{R}^D$. Also, for any signed measure μ defined over S_η, consider the linear operator $\mathcal{K}_\Omega$ defined by

$$\mathcal{K}_\Omega \mu := \int_{S_\eta} \mathcal{F}_\Omega [\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ d\mu(\boldsymbol{w}). \tag{11}$$

Then, for the measure $\mu^* = \sum_{i=1}^W a_i \delta_{\boldsymbol{w}_i}$ where $\delta_{\boldsymbol{w}}$ denotes a Dirac delta measure centered at $\boldsymbol{w} \in S_\eta$, we have $\mathcal{K}_\Omega \mu^* = \mathcal{F}_\Omega f_\theta$. This shows the constraint $\mathcal{F}_\Omega f_\theta = \boldsymbol{y}$ in $(P'_{\theta,W})$ is equivalent to $\mathcal{K}_\Omega \mu^* = \boldsymbol{y}$.

Next, we define the η-weighted total variation norm $\|\mu\|_{TV,\eta}$ of any signed measure μ over S_η by[1]

$$\|\mu\|_{TV,\eta} := \int_{S_\eta} \eta(\boldsymbol{w}) d|\mu|(\boldsymbol{w}),$$

where $|\mu|$ is the total variation measure associated with μ. For example, we have $|\mu^*| = \sum_{i=1}^W |a_i| \delta_{\boldsymbol{w}_i}$, and so $\|\mu^*\|_{TV,\eta} = \sum_{i=1}^W |a_i| \eta(\boldsymbol{w}_i)$, which coincides with the objective in $(P'_{\theta,W})$.

[1]By continuity of η, there exists a constant B such that $\eta(\boldsymbol{w}) \leq B$ for all $\|\boldsymbol{w}\| = 1$, which ensures the integral above is well-defined.

Therefore, letting $\mathcal{M}_W(S_\eta)$ denote the set of signed measures defined over S_η expressible as a weighted linear combination of at most W Diracs, we see that $(P'_{\theta,W})$ is equivalent to the optimization problem

$$p^*_W = \min_{\mu \in \mathcal{M}_W(S_\eta)} \|\mu\|_{TV,\eta} \;\; s.t. \;\; \mathcal{K}_\Omega \mu = \boldsymbol{y}. \tag{$P_{\mu,W}$}$$

In particular, the measure $\mu^* = \sum_{i=1}^W a_i \delta_{\boldsymbol{w}_i}$ is a minimizer of $(P_{\mu,W})$ if and only if the parameter vector $\theta^* = ((a_i, \boldsymbol{w}_i))_{i=1}^W \in \Theta'_W$ is a minimizer of $(P'_{\theta,W})$. Also, by Lemma 1, $(P'_{\theta,W})$ has the same function space minimizers as the original INR training problem $(P_{\theta,W})$. This shows that we can describe all function space minimizers of $(P_{\theta,W})$ by identifying the minimizers of $(P_{\mu,W})$.

However, analyzing $(P_{\mu,W})$ directly is still challenging since $\mathcal{M}_W(S_\eta)$ is not a closed vector space. Instead, we pass to the larger Banach space $\mathcal{M}(S_\eta)$ of signed Radon measures defined over S_η with finite total variation norm, and consider the optimization problem

$$p^* = \min_{\mu \in \mathcal{M}(S_\eta)} \|\mu\|_{TV,\eta} \;\; s.t. \;\; \mathcal{K}_\Omega \mu = \boldsymbol{y}. \tag{P_μ}$$

This can be thought of as the "infinite-width" analogue of $(P_{\mu,W})$, or equivalently, $(P'_{\theta,W})$. However, unlike these problems, (P_μ) is a *convex* optimization problem, albeit one posed over an infinite-dimensional space of measures.

Problem (P_μ) also has a function space interpretation: for any $\mu \in \mathcal{M}(S_\eta)$, define

$$f_\mu(\boldsymbol{x}) := \int_{S_\eta} [\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ d\mu(\boldsymbol{w}).$$

Then f_μ can be thought of as "infinite-width" shallow INR, for which $\mathcal{F}_\Omega f_\mu = \mathcal{K}_\Omega \mu$. Similar to before, we call f a function space minimizer of (P_μ), if $f = f_{\mu^*}$ where μ^* is any minimizer of (P_μ).

The existence of sparse minimizers of (P_μ) has direct implications for $(P_{\mu,W})$. First, since $\mathcal{M}_W(S_\eta) \subset \mathcal{M}(S_\eta)$ we always have $p^* \leq p^*_W$. Additionally, if a minimizer μ^* of (P_μ) is a finite linear combination of at most W Diracs, i.e., $\mu^* \in \mathcal{M}_W(S_\eta)$, then μ^* must also be a minimizer of $(P_{\mu,W})$, and so $p^* = p^*_W$. Our next result shows that, for sufficiently large training widths W, such a minimizer always exists assuming the set of frequencies used to define the Fourier features embedding contains a base set of frequencies.

Lemma 2. *Suppose the set of frequencies $\Omega_0 \subset \mathbb{Z}^d$ used to define the Fourier features embedding $\boldsymbol{\gamma}$ given in (5) is such that either $\boldsymbol{e}_i \in \Omega_0$ or $-\boldsymbol{e}_i \in \Omega_0$ for all $i = 1, ..., d$ where $\boldsymbol{e}_i$ is the ith elementary basis vector. Assume $W \geq |\Omega|$. Then, given any $\boldsymbol{y} \in \mathrm{Im}(\mathcal{F}_\Omega)$, (P_μ) has a minimizer $\mu^* \in \mathcal{M}_W(S_\eta)$.*

The proof of Lemma 2, which makes use of the abstract "representer theorem" of [6], is given in Appendix A.2. This result implies that when $W \geq |\Omega|$ the minimums of (P_μ) and $(P_{\mu,W})$ must coincide, i.e., $p^* = p^*_W$, and a sparse measure $\mu^* \in \mathcal{M}_W(S_\eta)$ is a minimizer of (P_μ) if and only if it is a minimizer of $(P_{\mu,W})$. Another immediate implication of Lemma 2 is that $W \geq |\Omega|$ is sufficient to guarantee the existence of feasible solutions for $(P_{\theta,W})$, i.e., given arbitrary Fourier data $\boldsymbol{y} \in \mathrm{Im}(\mathcal{F}_\Omega)$ there is always a shallow INR f_θ of width $W \leq |\Omega|$ such that $\mathcal{F}_\Omega f_\theta = \boldsymbol{y}$.

3.3 Exact Recovery of Shallow INRs

Now we give our main results, which identify a sufficient number of Fourier samples for which an image realizable as a shallow INR is exactly recoverable as the unique function space minimizer of the INR training problem $(P_{\theta,W})$.

Our first result focuses on recovery of a function f realizable as *width-1* shallow INR.

Theorem 1 (Width-1 INR, general Fourier features, $d \geq 2$)**.** *Consider the shallow INR architecture given in* (6) *with input dimension $d \geq 2$, where the set of frequencies $\Omega_0 \subset \mathbb{Z}^d$ used to define the Fourier features embedding $\boldsymbol{\gamma} : \mathbb{R}^d \to \mathbb{R}^D$ given in* (5) *is such that $\|\boldsymbol{k}\|_\infty \leq K_0$ for all $\boldsymbol{k} \in \Omega_0$. Let $f = a_1[\boldsymbol{w}_1^\top \boldsymbol{\gamma}(\cdot)]_+$ for some $a_1 \in \mathbb{R}$ and $\boldsymbol{w}_1 \in \mathbb{R}^D$. Suppose $\Omega \supseteq \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_\infty \leq 3K_0\}$, and let $\boldsymbol{y} = \mathcal{F}_\Omega f$. Then, for all training widths $W \geq 1$, f is the unique function space minimizer of the INR training problem $(P_{\theta,W})$ with weighting function $\eta(\boldsymbol{w}) = \|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2$.*

See Appendix B.1 for the proof. This result shows that $O(K_0^d)$ low-pass Fourier samples are sufficient to uniquely recover a width-1 shallow INR with a Fourier features embedding having a maximum frequency K_0 (in an ℓ^∞ sense). In particular, if the Fourier sampling set is $\Omega = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_\infty \leq K\}$, then $K \geq 3K_0$ is sufficient to ensure unique recovery.

Our second result focuses on the case of functions f realizable as a width-s shallow INR. However, we are unable to directly extend Theorem 1 to this setting in all generality. Instead, we restrict our attention to input dimension $d = 2$ (i.e., 2D images), and we consider an INR architecture defined in terms of a restricted set of Fourier features. In particular, we assume the Fourier features embedding $\boldsymbol{\gamma} : \mathbb{T}^2 \to \mathbb{R}^4$ has the form

$$\boldsymbol{\gamma}(\boldsymbol{x}) = \sqrt{2}[\cos(2\pi x_1), \sin(2\pi x_1), \cos(2\pi x_2), \sin(2\pi x_2)]^\top \text{ for all } \boldsymbol{x} = [x_1, x_2]^\top \in \mathbb{T}^2. \quad (12)$$

Note here that the constant feature 1 is omitted from $\boldsymbol{\gamma}$. Also, we assume that f is a *positive linear combination* of rectified trigonometric polynomials of the form $[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$. Finally, for technical reasons, we will assume that none of the weight vectors $\boldsymbol{w} \in \mathbb{R}^4$ belong to the set

$$V = \{\boldsymbol{w} \in \mathbb{R}^4 : w_1 = w_2 = 0, \text{ or } w_3 = w_4 = 0, \text{ or } w_1^2 + w_2^2 = w_3^2 + w_4^2\}. \quad (13)$$

Observe that V is a measure-zero subset of $\mathbb{R}^4$, i.e., our result still holds for almost all weight vectors $\boldsymbol{w} \in \mathbb{R}^4$.

Theorem 2 (Width-s INR, restricted Fourier features, $d = 2$)**.** *Consider the shallow INR architecture given in* (6) *with input dimension $d = 2$ and restricted Fourier features embedding $\boldsymbol{\gamma}$ given in* (12)*. Let $f = \sum_{i=1}^s a_i[\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ where $a_i > 0$ and $\boldsymbol{w}_i \in \mathbb{R}^4/V$ for all $i \in [s]$. Suppose $\Omega \supseteq \{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_1 \leq 2s\}$, and let $\boldsymbol{y} = \mathcal{F}_\Omega f$. Then, for all training widths $W \geq s$, f is the unique function space minimizer of the INR training problem $(P_{\theta,W})$ with weighting function $\eta(\boldsymbol{w}) = \int_{\mathbb{T}^2}[\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ d\boldsymbol{x}$.*

See Appendix B.2 for the proof. This result shows that $O(s^2)$ low-pass Fourier samples are sufficient to uniquely recover a 2D image realized by a width-s shallow INR having restricted

Fourier features. In particular, if the Fourier sampling set is $\Omega = \{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_\infty \leq K\}$, then $K \geq 2s$ is sufficient to ensure unique recovery.

The main idea of the proofs of Theorem 1 and Theorem 2 is to pass to the convex dual of (P_μ), which reduces to a semi-infinite program, i.e., a finite dimensional linear program with infinitely many constraints. Similar to previous works that have focused on super-resolution of sparse measures [8, 40, 12], we prove unique recovery is equivalent to the construction of a particular dual optimal solution known as a *dual certificate*; see Appendix A.3 for more details.

Note that both these results rely on a generalized form of weight decay regularization that differs from standard weight decay: standard weight decay corresponds to a weighting function $\eta(\boldsymbol{w}) = \|\boldsymbol{w}\|_2$, while in Theorem 1 we assume $\eta(\boldsymbol{w}) = \|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2$ and in Theorem 2 we assume $\eta(\boldsymbol{w}) = \int_{\mathbb{T}^2}[\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ d\boldsymbol{x}$. Ultimately, these generalized weighting functions arise as a consequence of our proof techniques. However, their use as regularizers is still intuitive: they can be seen as (approximately) penalizing the L^p-norm ($p = 1$ or 2) contribution made by each unit $[\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ to the overall image. This is immediately apparent in the case of $\eta(\boldsymbol{w}) = \int_{\mathbb{T}^2}[\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ d\boldsymbol{x} = \|[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_{L^1(\mathbb{T}^d)}$. And for $\eta(\boldsymbol{w}) = \|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2$, by Parseval's Theorem, we have $\eta(\boldsymbol{w}) \approx \|[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_{L^2(\mathbb{T}^d)}$ assuming the low-pass frequency set Ω is sufficiently large.

Despite strong theoretical guarantees and possible benefits in terms of interpretability, these more general forms of weight decay are more complicated and computationally expensive to implement than standard weight decay, which may limit their practical utility. However, we conjecture that recovery guarantees similar to Theorem 1 and Theorem 2 can be extended to the simpler case of standard weight decay regularization (i.e., weighting function $\eta(\boldsymbol{w}) = \|\boldsymbol{w}\|_2$), though doing so may require new proof techniques. This conjecture is verified empirically in the next section.

4 Experiments

Here we present a series of experiments to validate our sampling theory and to illustrate the effects of training INRs with (generalized) weight decay regularization in super-resolution recovery of 2D images. First, to validate the sampling guarantees given in Theorems 1 and 2, we analyze the probability of achieving exact recovery of random continuous-domain images realizable by shallow INRs from their low pass Fourier coefficients. Next, we illustrate the performance of shallow INRs in super-resolution recovery of various continuous-domain phantom images; additional experiments investigate the sparsity-inducing effects of (generalized) weight decay regularization. Finally, going beyond the theory developed in this work, we empirically investigate the role of depth in Fourier features INRs, demonstrating that deeper INRs trained with weight decay regularization achieve improved reconstruction on piecewise constant phantom images relative to shallow INRs, and compare the results with other INR architectures besides ReLU networks with Fourier features layers.

Implementation Details In all our experiments, to approximate the continuous domain Fourier sampling operator $\mathcal{F}_\Omega$, we implement a discretized version $\tilde{\mathcal{F}}_\Omega$ using the discrete Fourier transform. In particular, given a continuous domain image $f : [0,1]^2 \to \mathbb{R}$, we define $\tilde{\mathcal{F}}_\Omega f = \boldsymbol{F}_\Omega \mathcal{E} f$ where $\mathcal{E} f$ is a 2D pixel array obtained by evaluating f on an $N \times N$ uniform grid within $[0,1]^2$, and $\boldsymbol{F}_\Omega$ is the $N^2 \times N^2$ 2D discrete Fourier transform matrix restricted to rows corresponding to the frequency set Ω; we set $N = 1024$. Additionally, the weight functions η given in Theorem 1 and Theorem 2 used to define generalized weight decay penalties are implemented via the approximations $\tilde{\eta}(\boldsymbol{w}) = \|\tilde{\mathcal{F}}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2$ and $\tilde{\eta}(\boldsymbol{w}) = \frac{1}{N^2}\|\mathcal{E}[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_1$, respectively. All experiments were run in the PyTorch numerical computing environment using a single NVIDIA V100 GPU with 32 GB RAM. Additional implementation details are described in the Supplementary Materials E. Code to reproduce all experiments is available at `https://github.com/gregongie/super_inrs`.

4.1 Super-Resolution with Shallow INRs

4.1.1 Exact Recovery Experiments

To validate the exact recovery guarantees Theorem 1 and Theorem 2, we consider a "student-teacher" setup. Specifically, we generate random continuous domain images realizable as a "teacher" shallow INR of varying widths W (Figure 2 shows an example for $W = 3$), and investigate the maximum sampling frequency K (in an ℓ^∞ sense) needed to recover the image by solving the optimization problem (3) using a "student" shallow INR with the same architecture but larger width. As an attempt to solve the equality constrained problem (3) we apply the Augmented Lagrangian (AL) method (see, e.g., [30, Ch 7]); see Appendix C for a detailed algorithmic description.

We consider two settings, matching the settings of Theorem 1 and Theorem 2, which differ in terms of the Fourier features used and the form of generalized weight decay penalty used. In the setting of Theorem 1, we use the general Fourier features embedding (5) with maximum frequency $K_0 = 2$ resulting in $D = 25$ features, while in the setting of Theorem 2, we use the restricted Fourier features embedding (12) consisting of $D = 4$ features only. The generalized weight decay penalties defined in Theorem 1 and Theorem 2 we call *modified WD-I* and *modified WD-II*, respectively. We additionally compare against standard weight decay (i.e., the squared ℓ^2-norm of all trainable parameters), which we call *standard WD*. In both settings, we vary both the maximum sampling frequency cutoff K and the teacher width W, and run 10 random trials for each (K, W) pair to approximate the probability of exact recovery. A random teacher network is defined by drawing inner-layer weight vectors $\boldsymbol{w}_1, ..., \boldsymbol{w}_W$ uniformly at random from the sphere, and drawing outer-layer weights $a_1, ..., a_W$ uniformly at random from the interval $[1, 5]$. The student network width is fixed at 100 in all cases. We say the recovery is "exact" if the pixel-wise mean-squared error (MSE) between the ground truth image and student INR rasterized onto a 1024×1024 pixel grid drops below 10^{-9} at any point during training. More details for these experiments are provided in Supplementary Materials E.

Figure 3 shows the results of these experiments as probability tables. In the probability

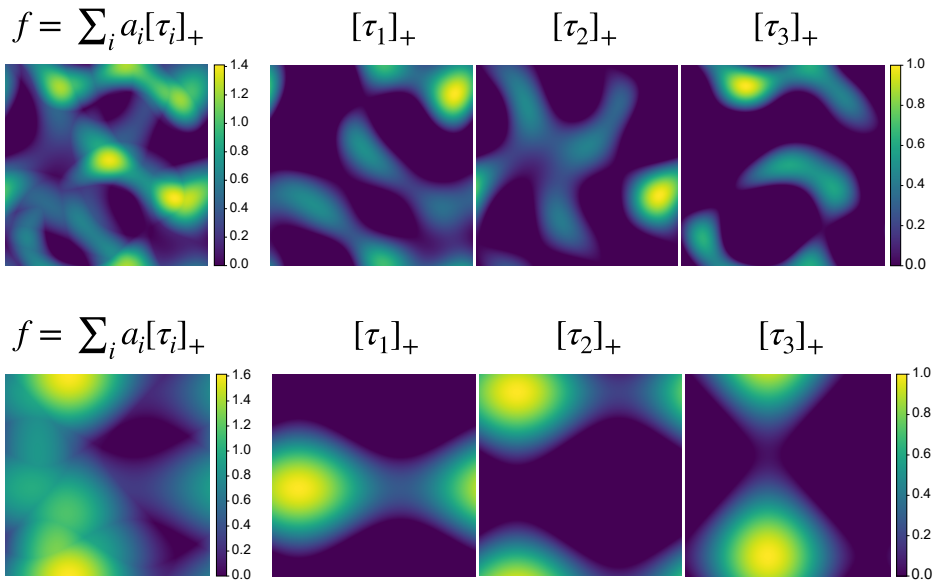

Figure 2: Example of a random width 3 "teacher" INR used in our exact recovery experiments to validate Theorem 1 (top panel) and Theorem 2 (bottom panel). The image f is a weighted sum of the three randomly generated rectified trigonometric polynomials $[\tau_i]_+ = [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ shown on the right. Note that the two settings consider INRs defined with Fourier features embeddings $\boldsymbol{\gamma}(\cdot)$ having different frequencies.

tables corresponding to Theorem 1, for standard WD regularization, we see that a linear trend between the teacher width W and the minimal sampling cutoff K needed to achieve exact recovery with high probability up to width 5, while exact recovery generally fails for teacher widths greater than 5. For modified WD-I regularization, although the linear trend is less pronounced, exact recovery is achieved in all random trials when the teacher width is 1 (consistent with Theorem 1) and remains attainable in many cases for teacher widths greater than one, but generally fails for widths greater than 5. Note that, for both regularization methods, the parameter space optimization problem is non-convex, and so convergence of the AL method to a global minimizer is not guaranteed. Therefore, it is unclear if the failure at achieving exact recovery for large widths in the setting of Theorem 1 is due to a difficulty in reaching a global minimizer with the proposed optimization method, or a more fundamental lack of identifiability.

In the probability tables corresponding to Theorem 2, the linear transition is clear for both standard WD and modified WD-II regularization. This is consistent with Theorem 2, which proves that $K = O(W)$ Fourier samples are sufficient for exact recovery. Also, in this setting, exact recovery is achieved for teacher widths greater than 5. We conjecture that this is because restricted Fourier features used in this setting result in a better conditioned problem.

This set of experiments suggests that standard weight decay regularization performs similar or better at achieving exact recovery of random phantom images relative to the modified forms of weight decay. However, there may be practical benefits to using modified weight

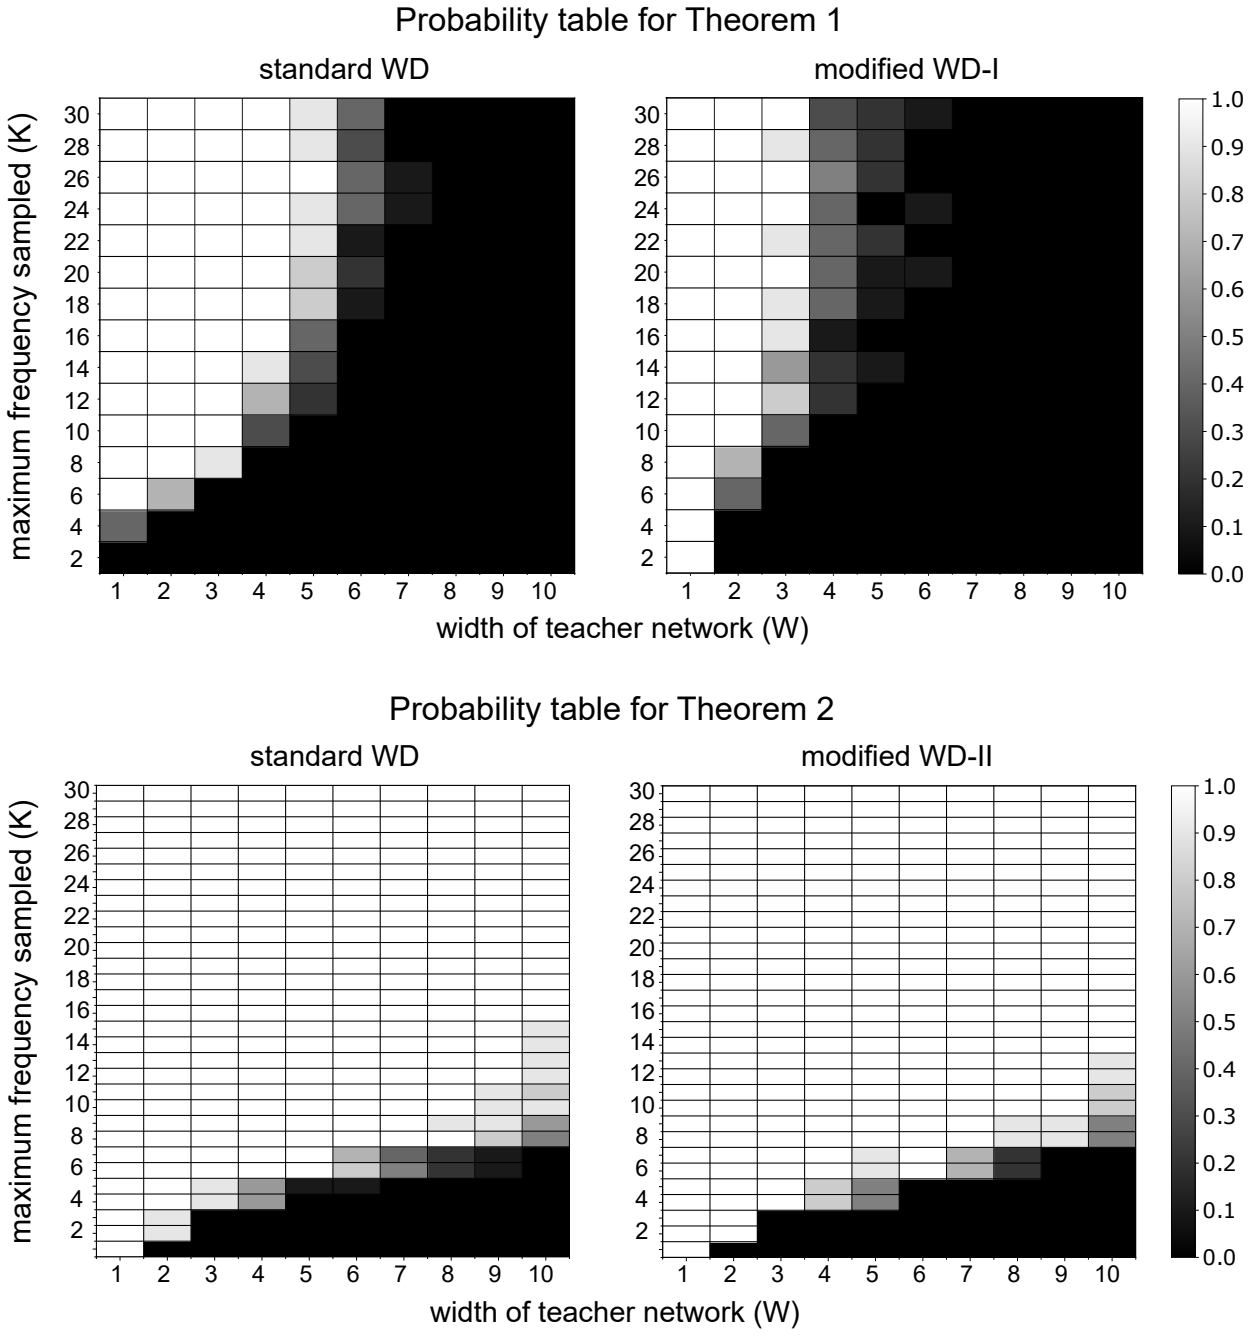

Figure 3: Empirical probability of exactly recovering an random image representable by a single hidden-layer INR of width-W from low-pass Fourier samples with maximum sampling frequency K by solving $(P_{\theta,W})$. The left panel corresponds to the setting of Theorem 1, comparing standard weight decay (standard WD) with modified weight decay regularization proposed for Theorem 1 (modified WD-I) and the right panel corresponds to the setting of Theorem 2, comparing standard weight decay (standard WD) with the modified weight decay proposed for Theorem 2 (modified WD-II). Empirical probabilities are estimated from 10 random trials.

decay regularization over standard weight decay regularization in achieving *approximate* recovery, as we show in the next set of experiments with more realistic phantom images.

4.1.2 Recovery of Continuous Domain Phantom Images

In Figure 4 and Figure 5, we illustrate the use of shallow INRs for the super-resolution recovery of four continuous domain phantoms whose Fourier coefficients are known exactly: (1) DOT, a "dot phantom" image, created for this study, that contains randomly generated disk-shaped features and is exactly realizable by a shallow INR architecture with a width of $W \geq 50$ and maximum Fourier features frequency $K_0 \geq 8$, (2) SL, the Shepp-Logan phantom [47], (3) PWC-BRAIN a piecewise constant MRI "brain phantom" introduced in [19], and (4) PWS-BRAIN a piecewise smooth version of the MRI brain phantom obtained by multiplying PWC-BRAIN by smooth function. Note that the SL, PWC-BRAIN, and PWS-BRAIN phantoms are all discontinuous functions, so they are not exactly realizable with the INR architecture considered in this work, whose output must be a continuous function. We also note all these four phantoms are not bandlimited, such that a simple reconstruction obtained by zero-filling in the Fourier domain and applying the IFFT results in significant ringing artifacts (see the second column in each panel of Figure 4 and Figure 5).

For each phantom, the maximum sampling frequency K, maximum Fourier features frequency K_0, and the width of the INR W, are set as shown in Table 1. For this set of

Phantom	K	K_0	W
DOT	32	10	100
SL	48	10	100
PWC-BRAIN	64	20	500
PWS-BRAIN	64	20	500

Table 1: Parameter settings for recovery of all phantoms by fitting a shallow INR

experiments, we fit the INR by minimizing the regularized least squares formulation (2) using the Adam optimizer for $40,000$ iterations with a learning rate of 1×10^{-3} followed by an additional $10,000$ iterations with a learning rate of 1×10^{-4}. The weight decay hyperparameter λ was selected via grid search to minimize the pixel-wise mean-squared error (MSE) between the ground truth image and INR rasterized onto a 1024×1024 pixel grid.

Unlike the exact recovery experiments, where standard weight decay regularization seems to perform better in achieving exact recovery, for all phantoms we find that using a shallow INR trained with modified weight decay regularization results in lower image MSE and fewer artifacts in the reconstructions compared to standard weight decay regularization. In particular, the MSE is a magnitude lower on the DOT phantom with modified weight decay and the reconstruction has far fewer ringing artifacts. A more modest improvement is shown on the other phantoms, where training with modified weight decay yields a slight improvement in MSE, but some ringing artifacts are still visible.

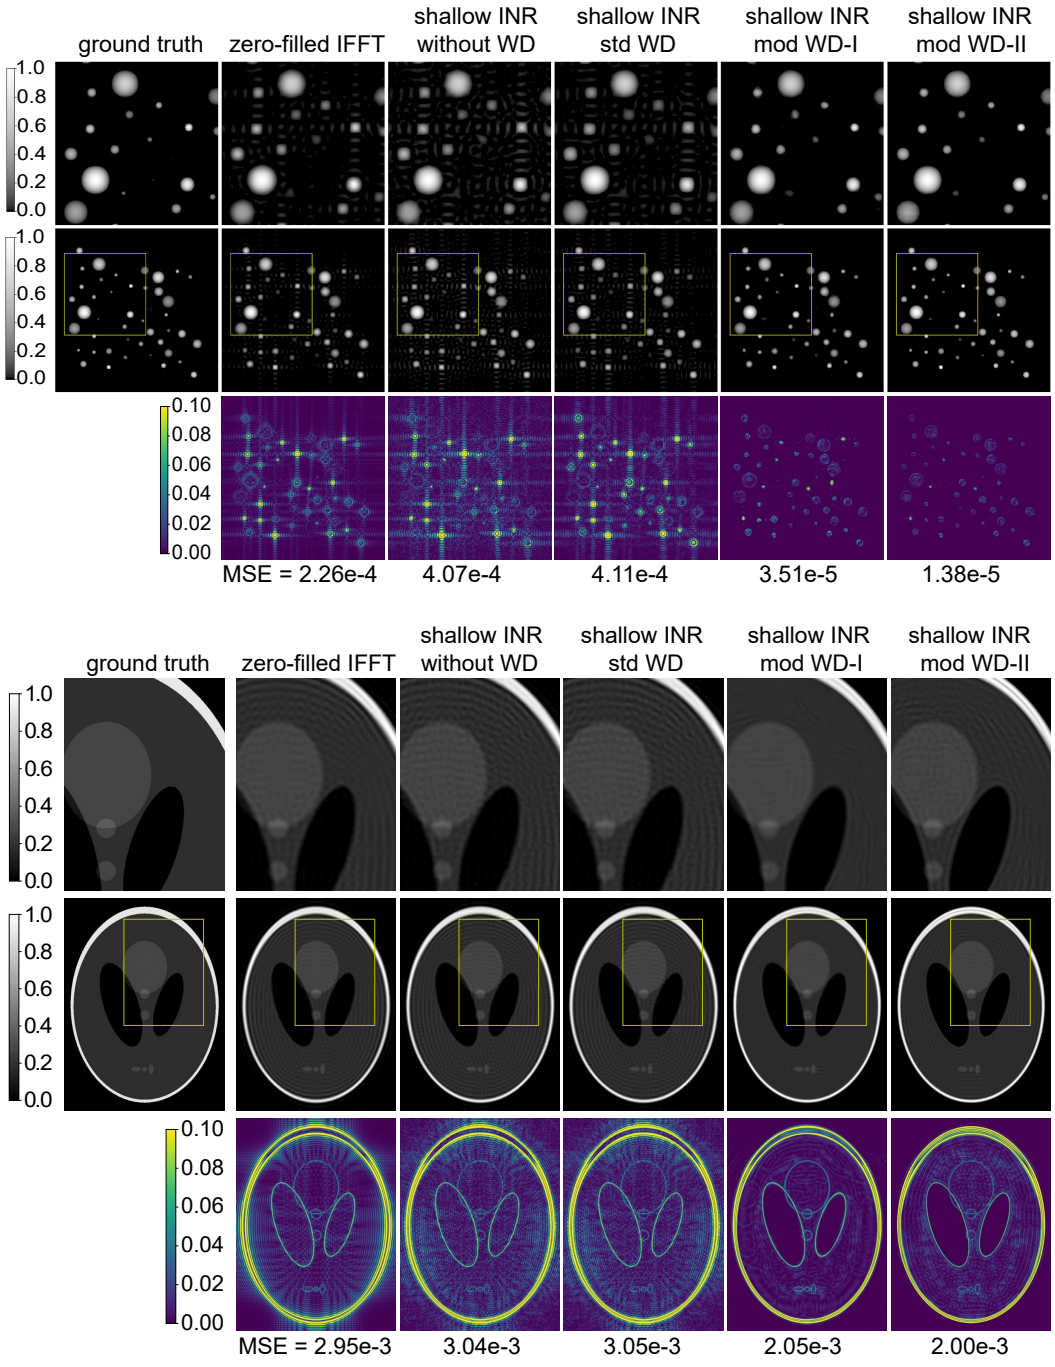

Figure 4: **Recovery of continuous domain DOT and SL phantoms from low-pass Fourier samples.** We compare the zero-filled IFFT reconstruction with a shallow INR trained without regularization (without WD), with standard weight decay regularization (shallow INR, std WD), the modified weight decay regularization proposed for Theorem 1 (shallow, mod WD-I), and the modified weight decay regularization proposed for Theorem 2 (shallow, mod WD-II). The bottom row shows the absolute value of the difference with the ground truth.

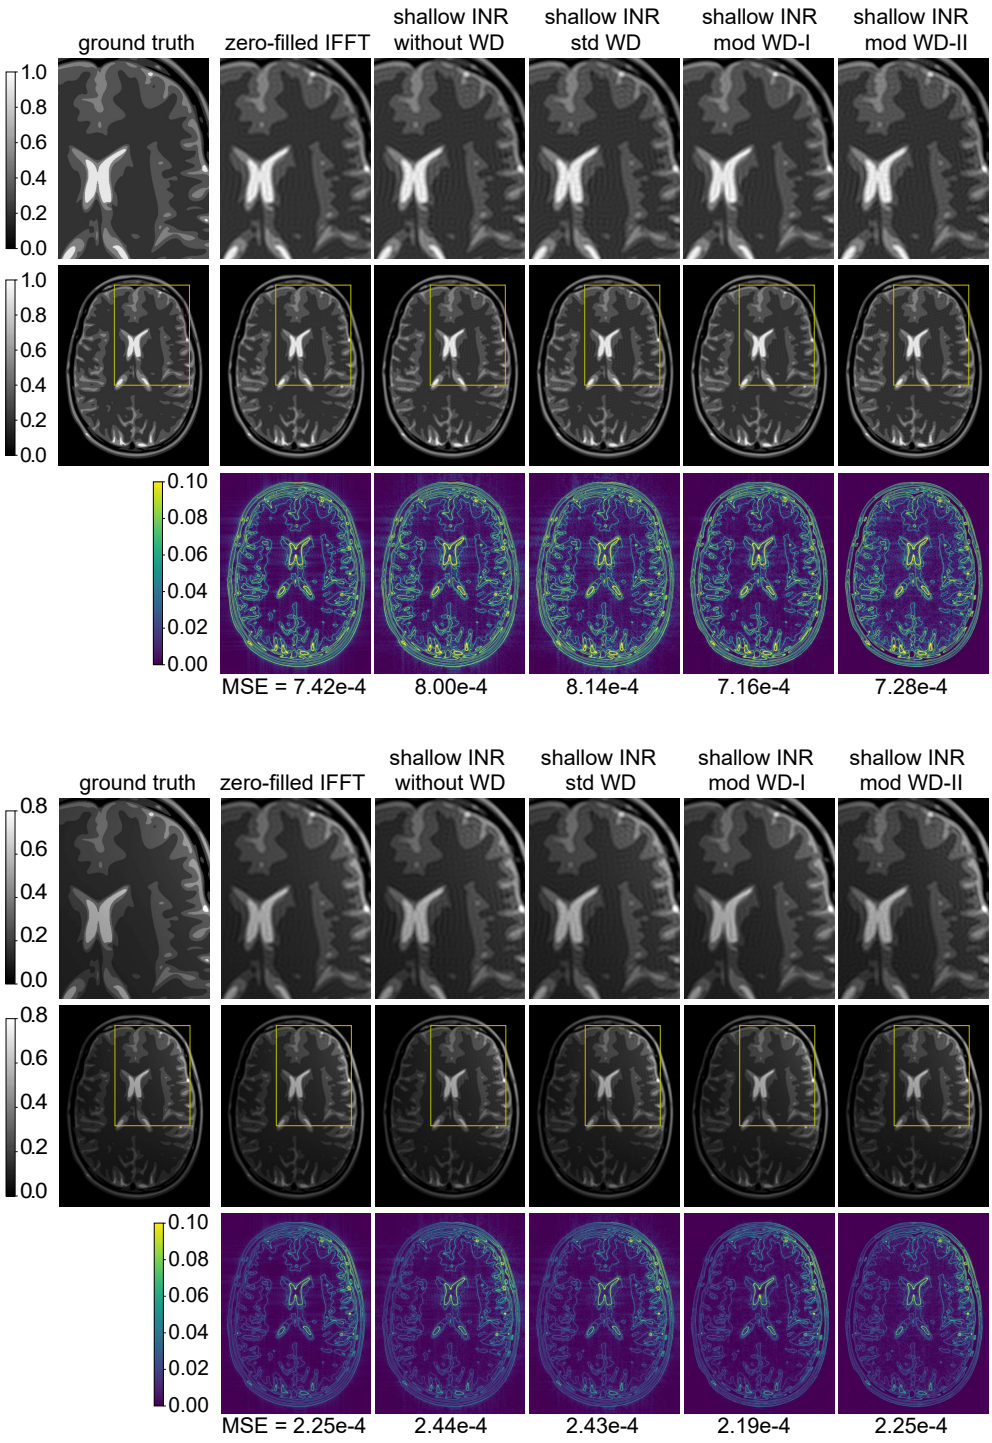

Figure 5: **Recovery of continuous domain PWC-BRAIN and PWS-BRAIN phantoms from low-pass Fourier samples.** We compare the zero-filled IFFT reconstruction with a shallow INR trained without regularization (without WD), with standard weight decay regularization (shallow INR, std WD), the modified weight decay regularization proposed for Theorem 1 (shallow, mod WD-I), and the modified weight decay regularization proposed for Theorem 2 (shallow, mod WD-II). The bottom row shows the absolute value of the difference with the ground truth.

To illustrate the impact of varying the maximum Fourier features frequency, K_0, in the Supplementary Materials D we show the reconstruction of the SL phantom with $K_0 = 5, 10, 15, 20$, trained with standard WD, modified WD-I, and modified WD-II. Setting $K_0 = 10$ resulted in reconstructions with the lowest image MSE for both modified WD-I and modified WD-II.

Additionally, in Figure 6, we demonstrate the impact of weight decay regularization strength on the reconstructed PWC-BRAIN phantom. First, we see that including weight decay regularization (i.e., setting $\lambda > 0$) yields reconstructions with lower MSE and fewer visual ringing artifacts than training without regularization ($\lambda = 0$). However, using too large of a λ value yields an over-smoothed reconstruction. In Figure 6, we also investigate the impact of regularization strength on the sparsity of the trained INR weights. In particular, given a ReLU unit of the form $a[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$, we define its unit size to be $|a|\eta(\boldsymbol{w})$, where η is the weighting function used to define the generalized weight decay regularization term in (8). By Lemma 1, the sum of the unit sizes should implicitly be minimized when using generalized weight decay. From the histograms of the unit sizes shown in Figure 6, we see a large spread in unit sizes when not using regularization ($\lambda = 0$), while as we increase the regularization strength, the number of units with size close to zero increases markedly. The effect is especially pronounced for large λ in the modified weight decay methods, where nearly all unit sizes are clustered close to zero except for a few large outliers, indicating the trained INR has very low effective width.

In Figure 7, we visualize the four largest active units from the unit size histograms presented in Figure 6 for both standard and modified weight decay regularization applied to the PWC-BRAIN phantom. As predicted by our theory, applying stronger weight decay regularization (i.e., larger λ) results in a sparse INR, where only a few units contribute significantly to the output image. Consequently, these active units closely resemble the final reconstructed images. As the effect of regularization decreases (i.e., smaller λ), the number of active units increases. In particular, without regularization (i.e., $\lambda = 0$), the trained INR is no longer sparse, resulting in a larger number of active units, which appear random in structure and do not visually resemble the target phantom image. For the DOT phantom, similar figures representing the effect of regularization on the phantom reconstruction, unit sizes of a trained shallow INR, and the visualization of the four largest active units are provided in the Supplementary Materials D.

4.2 Super-Resolution with Deep INRs

4.2.1 Effect of Depth for Fourier Feature INRs

Though our theory focuses on shallow INRs, deeper INR architectures are often used in practice. Here, we empirically explore the effect of INR depth on super-resolution recovery. Note that the generalized weight decay penalties investigated previously are only defined for shallow INRs, so for these experiments we focus on standard weight decay regularization only.

In Figure 8, we show the results of fitting an INR of width 100 with depths of $2, 5, 10$, 15,

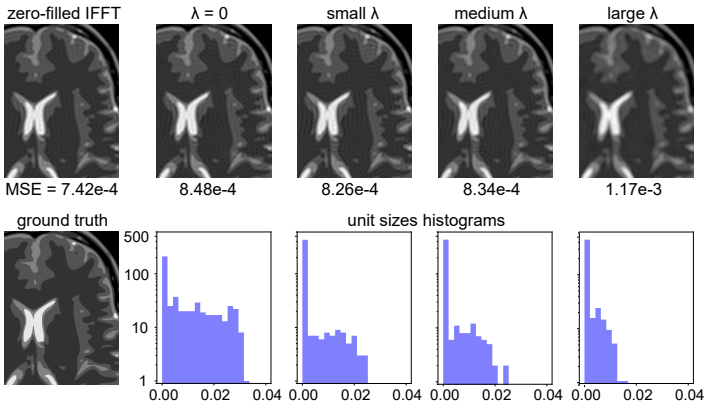

(a) standard weight decay regularization

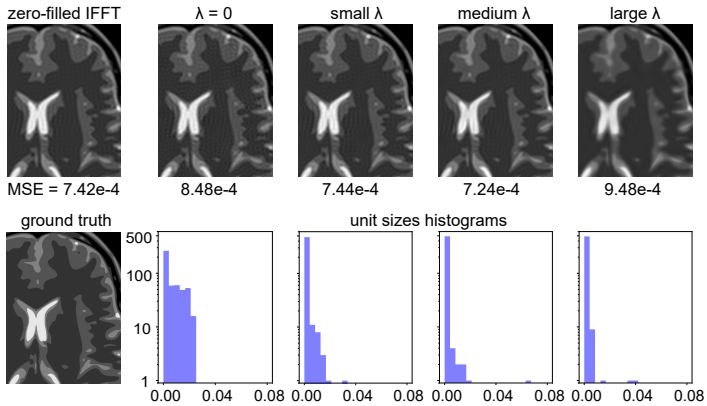

(b) modified weight decay regularization proposed in Theorem 1

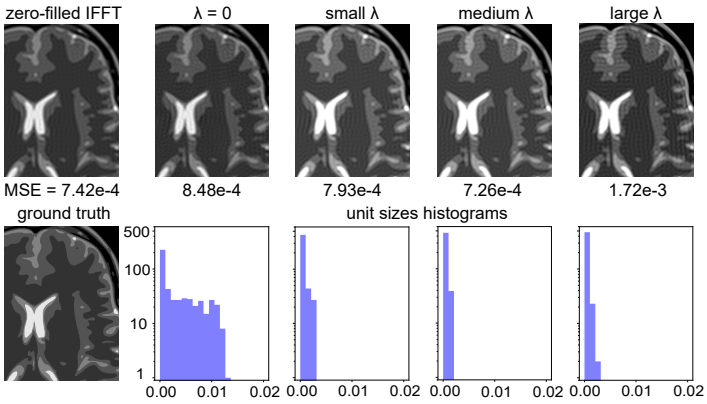

(c) modified weight decay regularization proposed in Theorem 2

Figure 6: **Effect of regularization on the PWC BRAIN phantom reconstruction and unit sizes of a trained depth 2 INR.** In the first row of each panel, we compare the zero-filled IFFT reconstruction with the reconstructions obtained with a shallow INR trained with standard/modified weight decay regularization for different values of regularization strength λ. The histograms in the second row of each panel represent the distribution of unit sizes $\{|a_i|\eta(\boldsymbol{w}_i)\}_{i=1}^{500}$ of the trained INRs across different values of λ. As λ increases, the unit sizes cluster at zero, indicating the trained INR is sparse in the sense that it has a few active units. (Note: histograms are plotted on a log scale.)

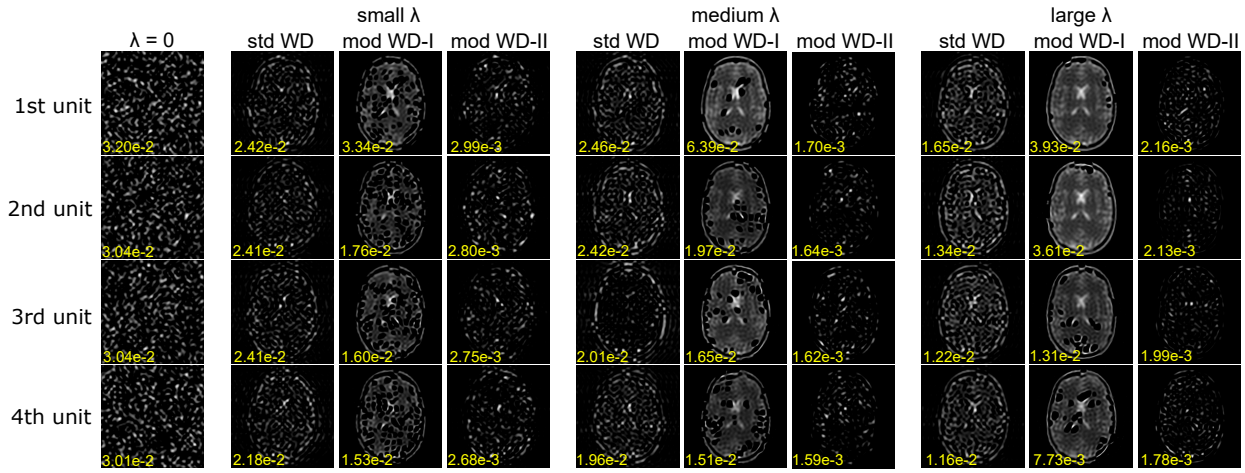

Figure 7: **Visualization of the four largest active units for the PWC-BRAIN phantom, obtained with a shallow INR.** The INR is trained with standard weight decay regularization (std WD), The modified weight decay regularization proposed in Theorem 1 (mod WD-I), and the modified weight decay regularization proposed in Theorem 2 (mod WD-II). All images are normalized to the [0,1] range. The corresponding "unit size" $|a_i|\eta(\boldsymbol{w}_i)$ is displayed on the image in yellow.

and 20, with and without weight decay regularization for the recovery of the SL and PWC BRAIN phantoms. These results show that increasing INR depth results in lower MSE and fewer ringing artifacts, indicating that deeper INRs may be biased towards piecewise constant images. Additionally, for both phantoms, comparing reconstructions with and without regularization shows that, at any fixed depth, applying weight decay regularization leads to lower image MSE and better visual quality. This suggests that weight decay regularization can also help improve recovery in deep INR networks. For more details on the hyperparameter settings used for these results see Supplementary Materials E.2.

4.2.2 Comparison with Different INR Architectures

In this work, we focus on one specific INR architecture inspired by [49] that combines a Fourier features embedding with a fully connected neural network with ReLU activation function. However, there are several alternative INR architectures that use different activation functions and/or feature embeddings. Here, we verify that the INR architecture investigated in this work gives comparable results to other popular INR architectures on the super-resolution recovery problem.

Figure 9 and Figure 10 present the reconstructions of the SL and PWC-BRAIN phantoms using various INR architectures. Specifically, we compare with three established INR architectures: Hash Encoding networks [28], SIREN [48], and ReLU networks with *random* Fourier features [49], in contrast to the *non-random* Fourier features used in previous experiments in this work. In addition, motivated by our earlier findings that deeper INR architectures tend to improve the recovery of piecewise constant phantoms, we evaluate SIREN at two depths: one with a depth of 5, denoted as "SIREN" and one with a depth of 14, denoted as "deep

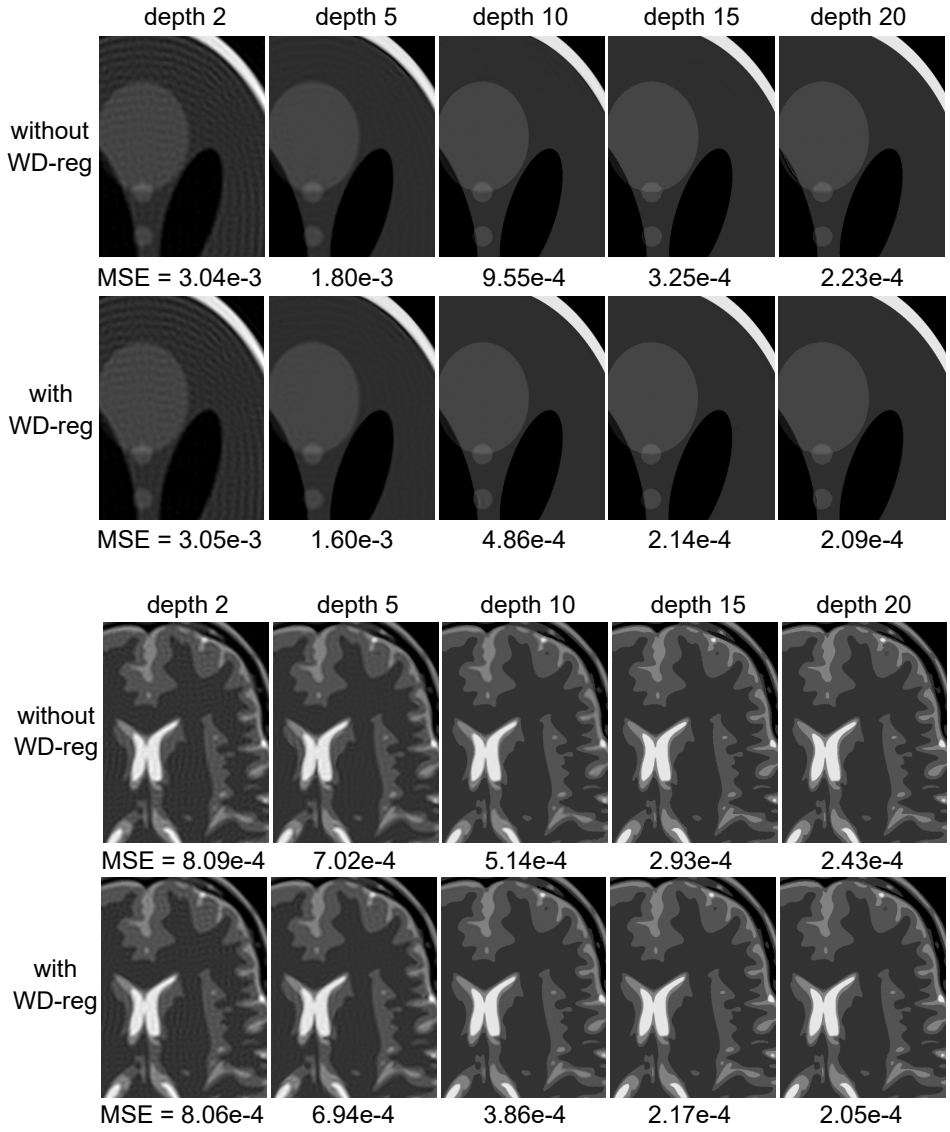

Figure 8: **Effect of INR depth on SL and PWC-BRAIN phantom recovery.** The top row shows comparisons of reconstructions of an INR of width 100 across depths 2, 5, 10, 15, and 20 trained with no regularization (i.e., with the weight decay hyperparameter $\lambda = 0$). The bottom row shows comparisons of reconstructions of an INR with varying depths, trained with the standard weight decay regularization. All images are displayed on a scale of [0,1].

SIREN".

For a fair comparison, we conducted experiments with different training settings, adjusting the INR depth, the learning rate, weight decay regularization strength λ, and other architecture-specific hyperparameters. In Figure 9 and Figure 10, we display the reconstruction that achieved the lowest image MSE across all tuned hyperparameters. See Supplementary Materials E.3 for the hyperparameter settings used in these experiments.

As shown in these figures, all INR architectures result in lower image MSE values com-

pared to zero-filled IFFT reconstruction. Reconstructions using Hash Encoding and SIREN architectures are comparable to reconstructions obtained using a shallow INR with Fourier features layer in terms of image MSE and visual quality. The deep INR architectures (deep SIREN, and deep ReLU networks with Fourier features) give a marked improvement, eliminating most visible ringing artifacts, and achieve a substantially lower image MSE than shallower architectures. Additionally, a figure representing the PWS-BRAIN phantom recovery using different INR architectures is provided in the Supplementary Materials 11. Here a similar improvement is seen with deeper INR architectures, though the relative change MSE and visual quality is smaller than with the piecewise constant phantoms.

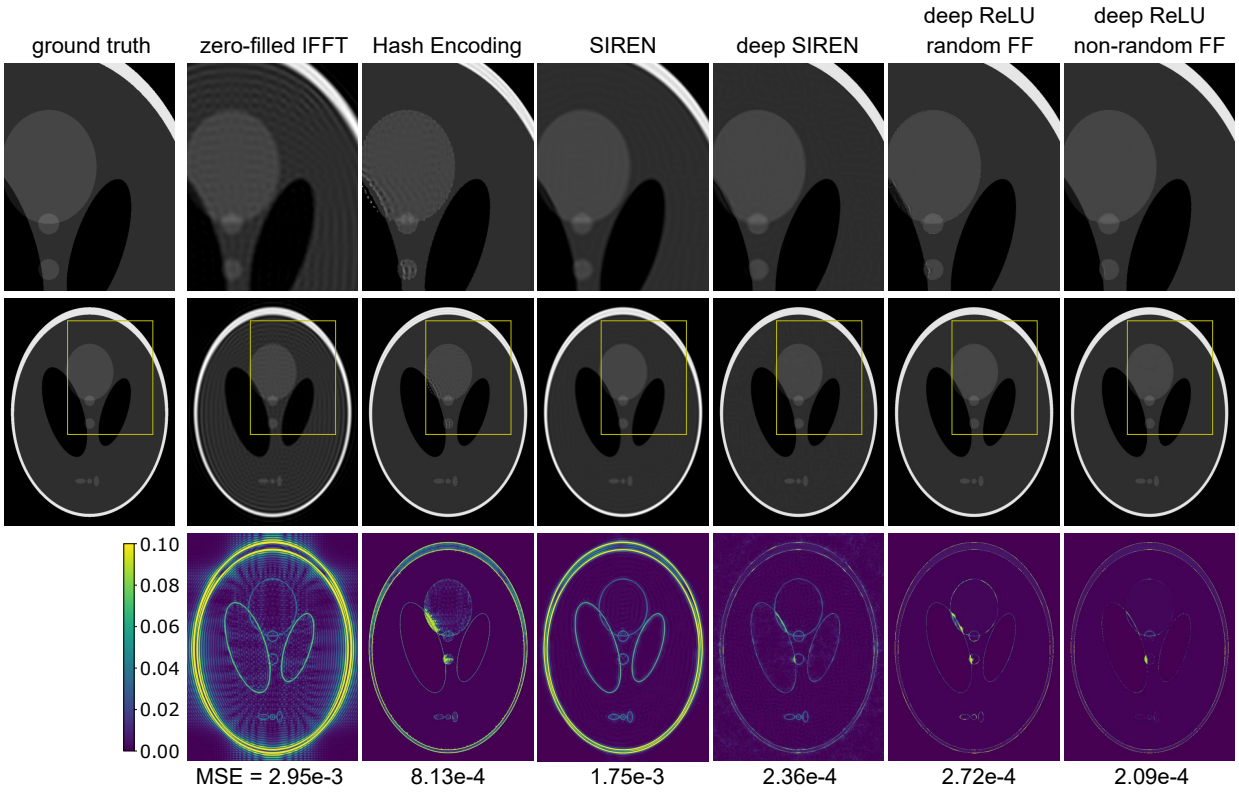

Figure 9: **SL recovery using different INR architectures.** Top row: we compare the zero-filled IFFT reconstruction with the phantom recovery from a Hash Encoding INR, depth 5 SIREN (SIREN), depth 14 SIREN (deep SIREN), and depth 20 ReLU INR (deep ReLU) with *random* Fourier features and *non-random* Fourier features. The bottom row shows the absolute value of the difference with the ground truth. All grayscale images are shown on the scale [0, 1].

5 Conclusions

In this work, we introduce a mathematical framework for studying exact recovery of continuous domain images from low-pass samples with a specific INR architecture. We show

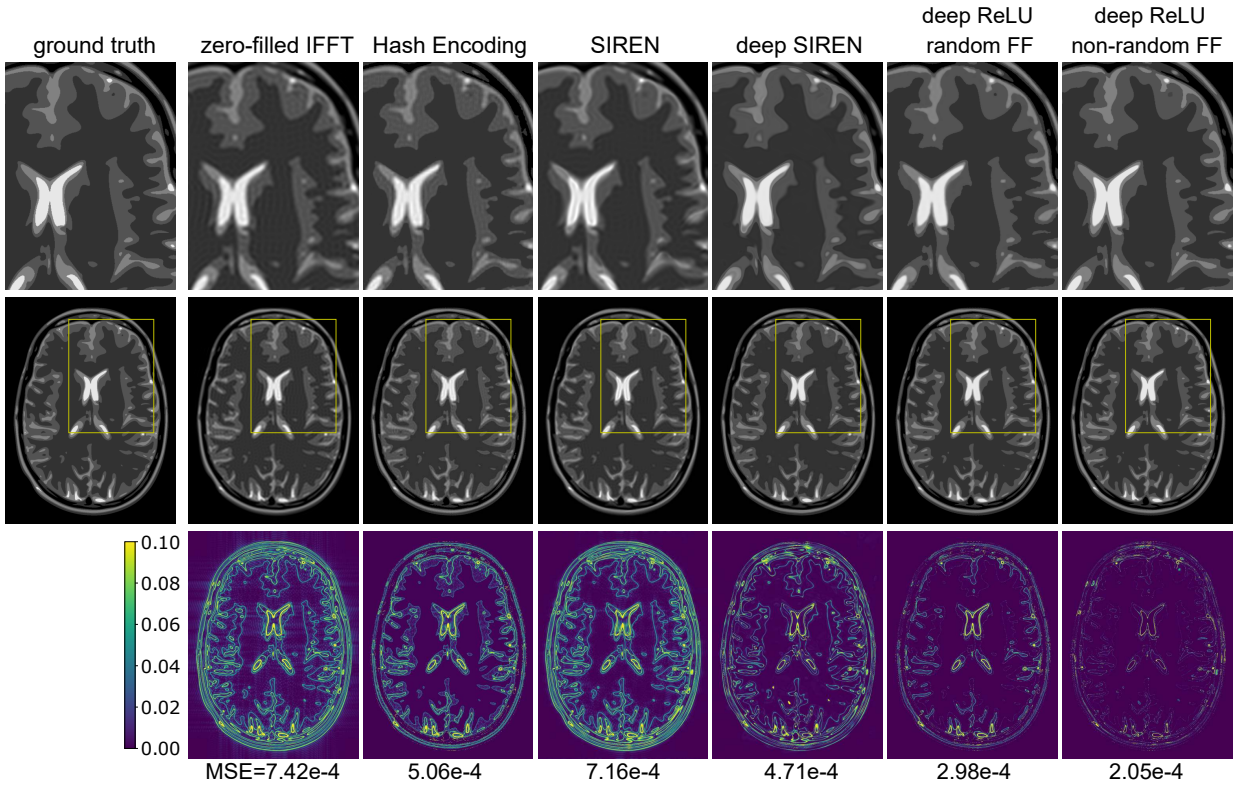

Figure 10: **PWC-BRAIN recovery using different INR architectures.** Top row: we compare the zero-filled IFFT reconstruction with the phantom recovery from a Hash Encoding INR, depth 5 SIREN (SIREN), depth 14 SIREN (deep SIREN), and depth 20 ReLU INR (deep ReLU) with *random* Fourier features and *non-random* Fourier features. The bottom row shows the absolute value of the difference with the ground truth. All grayscale images are shown on the scale [0, 1].

that fitting a shallow INR architecture having ReLU activation and an initial Fourier features layer with a generalized form of weight decay regularization is equivalent to a convex problem optimization defined over a space of measures. Using this perspective, we identify a sufficient number of Fourier samples to uniquely recover an image realizable as INR in two settings: (i) the image is realizable as a width-1 shallow INR with a general Fourier features embedding defined in terms of integer frequencies, and (ii) the image is realizable as a width-s shallow INR with a restricted Fourier features embedding defined in terms of two base frequencies and positive outer-layer weights. Furthermore, we conjecture that similar exact recovery results hold with standard weight decay regularization, which is verified empirically.

While this work focused on super-resolution imaging, the convex optimization framework developed here naturally extends to other linear inverse problems, such as MRI reconstruction from non-uniform Fourier samples or computed tomography reconstruction. An important next step is to investigate whether our proof techniques can yield exact recovery guarantees for INRs over a broader class of inverse problems.

Our analysis characterizes the global optima of an INR training problem, showing that its solutions coincide with those of a convex optimization problem. However, since the original INR training objective is non-convex, convergence to global optima via gradient-based methods is not guaranteed. Despite this, our empirical results demonstrate that standard INR training procedures can converge to global optima under certain conditions. Establishing theoretical convergence guarantees for practical gradient-based INR training algorithms is an interesting open problem.

Finally, the current analysis is specific to the INR architecture proposed in [49]. Extending our results to other popular INR architectures (as discussed in Section 1.1) is a natural direction for further exploration. Moreover, for mathematical tractability, we restricted our focus to shallow networks with a single hidden layer. Generalizing our analysis to deeper INR architectures and examining the role of weight decay regularization in such settings will be a focus of future work.

A Proofs of Lemmas in Section 3

A.1 Proof of Lemma 1

Given any parameter vector $\theta = ((a_i, \boldsymbol{w}_i))_{i=1}^W \in \Theta_W$ we construct a new parameter vector $\theta' = ((a_i', \boldsymbol{w}_i'))_{i=1}^W \in \Theta_W$ that satisfies $f_\theta = f_{\theta'}$ as follows: For all i such that $\eta(\boldsymbol{w}_i) > 0$, define $c_i = \sqrt{|a_i|/\eta(\boldsymbol{w}_i)}$, and set $a_i' = a_i/c_i$ and $\boldsymbol{w}_i' = c_i \boldsymbol{w}_i$. For all i such that $\eta(\boldsymbol{w}_i) = 0$, set $\boldsymbol{w}_i' = \boldsymbol{0}$ and $a_i' = 0$. Then, by positive 1-homogeneity of the ReLU activation and of η, and since $\eta(\boldsymbol{w}_i) = 0$ implies $[\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+ = 0$, we see that $a_i'[(\boldsymbol{w}_i')^\top \boldsymbol{\gamma}(\cdot)]_+ = a_i[\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ for all i, and so $f_\theta = f_{\theta'}$.

Now, by the inequality of arithmetic means and geometric means, we have

$$R(\theta) = \frac{1}{2}\sum_{i=1}^{W}\left(|a_i|^2 + \eta(\boldsymbol{w}_i)^2\right) \geq \sum_{i=1}^{W}|a_i|\eta(\boldsymbol{w}_i) = R(\theta').$$

This shows that every parameter vector $\theta \in \Theta_W$ can be mapped to another parameter vector $\theta' \in \Theta_W$ that realizes the same function and achieves the modified cost $\sum_{i=1}^{W}|a_i|\eta(\boldsymbol{w}_i)$ which is also a lower bound on $R(\theta)$. Therefore, we see that $(P_{\theta,W})$ has the same function space minimizers as

$$\min_{\theta\in\Theta_W}\sum_{i=1}^{W}|a_i|\eta(\boldsymbol{w}_i) \;\; s.t. \;\; \mathcal{F}_\Omega f_\theta = \boldsymbol{y}.$$

Finally, since the quantity $|a_i|\eta(\boldsymbol{w}_i)$ is invariant to the re-scaling $\boldsymbol{w}_i \mapsto \boldsymbol{w}_i/\|\boldsymbol{w}_i\|$ and $a_i \mapsto \|\boldsymbol{w}_i\|a_i$, and this re-scaling of parameters realizes the same function, we may additionally constrain the optimization to all parameter vectors satisfying $\|\boldsymbol{w}_i\| = 1$ for all $i \in [W]$, i.e., to all $\theta \in \Theta'_W$, which gives the equivalence between function space minimizers of $(P_{\theta,W})$ and $(P'_{\theta,W})$.

A.2 Proof of Lemma 2

To prove Lemma 2, we apply the abstract representer theorem of [6, Theorem 3.3]. A simplified version of this result is stated as follows: if Q is a norm on a locally convex space $\mathcal{X}$, and $\mathcal{A} : \mathcal{X} \to H$ is a continuous linear operator mapping surjectively onto a finite dimensional Hilbert space H, then given any $y \in H$ the optimization problem

$$\min_{x\in\mathcal{X}} Q(x) \;\; s.t. \;\; \mathcal{A}x = y,$$

has a minimizer of the form $x^* = \sum_{i=1}^{s} a_i u_i$ where $s \leq \dim H$, $a_i \in \mathbb{R}$, and u_i are extremal points of the unit ball $\{x \in \mathcal{X} : \phi(x) \leq 1\}$.

In the present setting, we consider the locally convex space $\mathcal{X} = \mathcal{M}(S_\eta)$ of finite signed Radon measures over S_η equipped with the weak-$*$ topology. Here, we identify $\mathcal{M}(S_\eta)$ as the continuous dual of $\mathcal{C}_0(S_\eta)$, the Banach space of continuous functions on S_η vanishing on the boundary set $\partial S_\eta \subseteq \{\boldsymbol{w} \in \mathbb{S}^{D-1} : \eta(\boldsymbol{w}) = 0\}$ equipped with the supremum norm. We also consider the norm $Q = \|\cdot\|_{TV,\eta}$ on $\mathcal{M}(S_\eta)$, and the linear operator $\mathcal{A} = \mathcal{K}_\Omega$ mapping into $H = \mathbb{C}_{\mathrm{sym}}(\Omega)$, the space of complex-valued conjugate symmetric arrays indexed by Ω, which is isomorphic to $\mathbb{R}^{|\Omega|}$.

By a similar argument to [6, Theorem 4.2], it is straightforward to show that the extreme points of the unit ball $\{\mu \in \mathcal{M}(S_\eta) : \|\mu\|_{TV,\eta} \leq 1\}$ coincide with the set of scaled Dirac measures $\{\pm\eta(\boldsymbol{w})^{-1}\delta_{\boldsymbol{w}} : \boldsymbol{w} \in S_\eta\}$. Hence, assuming $\mathcal{K}_\Omega$ is weak-$*$ continuous and surjective, we see that (P_μ) has a sparse solution $\mu^* = \sum_{i=1}^{s} a_i\delta_{\boldsymbol{w}_i}$ where $a_i \in \mathbb{R}$ and $\boldsymbol{w}_i \in S_\eta$, and $s \leq \dim_{\mathbb{R}}\mathbb{C}_{\mathrm{sym}}(\Omega) = |\Omega|$. If $W \geq |\Omega|$, this implies μ^* is a minimizer of $(P_{\mu,W})$ and so $p^*_W = p^*$, as claimed.

All that remains so show is that $\mathcal{K}_\Omega : \mathcal{M}(S_\eta) \to \mathbb{C}_{\mathrm{sym}}(\Omega)$ is weak-$*$ continuous and surjective. First, to prove that $\mathcal{K}_\Omega$ is weak-$*$ continuous, it suffices to show $\mathcal{K}_\Omega$ is the adjoint

of a bounded linear operator [11, Prop. 1.3, pg. 167]. In particular, consider the linear operator $\mathcal{L}_\Omega : \mathbb{C}_{\text{sym}}(\Omega) \to \mathcal{C}_0(S_\eta)$ defined for any $\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)$ by

$$[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}) = \langle \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+, \boldsymbol{z} \rangle \quad \forall \boldsymbol{w} \in S_\eta. \tag{14}$$

By the Cauchy-Schwarz inequality, we have

$$|[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w})| \leq \|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2 \|\boldsymbol{z}\|_2,$$

and Parseval's Theorem gives the bound

$$\begin{aligned}
\|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2 &= \|\mathcal{F}_\Omega^* \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_{L^2(\mathbb{T}^d)} \\
&\leq \|[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_{L^2(\mathbb{T}^d)} \\
&\leq \|\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)\|_{L^2(\mathbb{T}^d)} \\
&= \|\boldsymbol{w}\|_2 = 1,
\end{aligned}$$

where in the last equality we used the fact that the components $\{\gamma_i\}_{i=1}^D$ of the Fourier features embedding $\boldsymbol{\gamma}$ defined in (5) are orthonormal as functions in $L^2(\mathbb{T}^d)$. Therefore, we have shown $|[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w})| \leq \|\boldsymbol{z}\|_2$ for all $\boldsymbol{w} \in S_\eta$, and so $\mathcal{L}_\Omega$ is a bounded linear operator. Now we show $\mathcal{K}_\Omega = \mathcal{L}_\Omega^*$: for any $\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)$ and $\mu \in \mathcal{M}(S_\eta)$, we have

$$\begin{aligned}
\langle \mu, \mathcal{L}_\Omega \boldsymbol{z} \rangle &= \int_{S_\eta} [\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}) d\mu(\boldsymbol{w}) \\
&= \int_{S_\eta} \langle \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+, \boldsymbol{z} \rangle d\mu(\boldsymbol{w}) \\
&= \left\langle \int_{S_\eta} \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ d\mu(\boldsymbol{w}), \boldsymbol{z} \right\rangle \\
&= \langle \mathcal{K}_\Omega \mu, \boldsymbol{z} \rangle,
\end{aligned}$$

which proves $\mathcal{K}_\Omega = \mathcal{L}_\Omega^*$, as claimed.

Finally, we show that $\mathcal{K}_\Omega$ is surjective, i.e., $\text{Im}(\mathcal{K}_\Omega) = \mathbb{C}_{\text{sym}}(\Omega)$, or equivalently, $\operatorname{rank} \mathcal{K}_\Omega = |\Omega|$. For this, it suffices to show $\mathcal{L}_\Omega$ has a trivial nullspace, since by the rank-nullity theorem this implies $\operatorname{rank} \mathcal{L}_\Omega = |\Omega|$, and so $\operatorname{rank} \mathcal{L}_\Omega^* = \operatorname{rank} \mathcal{K}_\Omega = |\Omega|$, as well.

Towards this end, suppose $\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)$ is non-zero. We will show $\mathcal{L}_\Omega \boldsymbol{z}$ is not identically zero, i.e., there exists a vector $\boldsymbol{w}_1' \in S_\eta$ such that $[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}_1') \neq 0$. First, observe that for all $\boldsymbol{w} \in S_\eta$ we have

$$\begin{aligned}
[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}) &= \langle \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+, \boldsymbol{z} \rangle \\
&= \int_{\mathbb{T}^d} [\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ [\mathcal{F}_\Omega^* \boldsymbol{z}](\boldsymbol{x}) \, d\boldsymbol{x}.
\end{aligned}$$

Define $\Phi = \mathcal{F}_\Omega^* \boldsymbol{z}$, which is a non-zero real-valued trigonometric polynomial. Without loss of generality, assume the set of values of which Φ is positive $P = \{\boldsymbol{x} \in \mathbb{T}^d : \Phi(\boldsymbol{x}) > 0\}$ is

non-empty (if it is empty, the argument below can be repeated for the set over which Φ is negative, instead). By continuity of Φ, P is an open set, and so it contains an open ball $B(\boldsymbol{x}_0, \epsilon)$ for some radius $\epsilon > 0$. Let

$$\tau_0(\boldsymbol{x}) = \frac{1}{2} + \frac{1}{2d}\sum_{j=1}^{d} \cos(2\pi x_j),$$

which belongs to the range of the Fourier features embedding by assumption, i.e., there exists $\boldsymbol{w}_0 \in \mathbb{R}^D$ such that $\tau_0(\boldsymbol{x}) = \boldsymbol{w}_0^\top \boldsymbol{\gamma}(\boldsymbol{x})$. Note that $0 \leq \tau_0(\boldsymbol{x}) \leq 1$ for all $\boldsymbol{x} \in \mathbb{T}^d$, and it is easy to prove that $\boldsymbol{x} = \boldsymbol{0}$ is the unique global maximizer of τ_0. Now, define $\tau_1(\boldsymbol{x}) := \tau_0(\boldsymbol{x} - \boldsymbol{x}_0) - \alpha$ where we choose $0 < \alpha < 1$ to be such that $[\tau_1]_+$ is non-zero with support contained in $B(\boldsymbol{x}_0, \epsilon)$. Since τ_1 is simply a translation and constant shift of τ_0 it has the same frequency support as τ_0. Therefore, there exists a non-zero vector $\boldsymbol{w}_1 \in \mathbb{R}^D$ such that $\tau_1(\boldsymbol{x}) = \boldsymbol{w}_1^\top \boldsymbol{\gamma}(\boldsymbol{x})$. Defining $\boldsymbol{w}_1' = \boldsymbol{w}_1 / \|\boldsymbol{w}_1\|_2$, we see that

$$[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}_1') = \|\boldsymbol{w}_1\|_2^{-1} \int_{\mathbb{T}^d} [\tau_1(\boldsymbol{x})]_+ \Phi(\boldsymbol{x})\, d\boldsymbol{x} = \|\boldsymbol{w}_1\|_2^{-1} \int_P [\tau_1(\boldsymbol{x})]_+ \Phi(\boldsymbol{x})\, d\boldsymbol{x} > 0.$$

Finally, since $[(\boldsymbol{w}_1')^\top \boldsymbol{\gamma}(\cdot)]_+ = \|\boldsymbol{w}_1\|_2^{-1} [\tau_1(\cdot)]_+ \neq 0$, by the admissibility conditions it must be the case that $\eta(\boldsymbol{w}_1') > 0$, which gives $\boldsymbol{w}_1' \in S_\eta$. This shows there exists $\boldsymbol{w}_1' \in S_\eta$ such that $[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}_1') \neq 0$, as claimed.

A.3 Description of the Dual Problem

The optimization problem (P_μ) is convex. We now describe the associated convex dual problem[2].

Lemma 3. *Let p^* be the minimum of the primal problem (P_μ). Then we have $p^* = d^*$ where*

$$d^* = \sup_{\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)} \langle \boldsymbol{y}, \boldsymbol{z} \rangle \quad s.t. \quad |[\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w})| \leq \eta(\boldsymbol{w}) \quad \forall \boldsymbol{w} \in S_\eta, \tag{D}$$

where $\mathcal{L}_\Omega : \mathbb{C}_{\text{sym}}(\Omega) \to \mathcal{C}_0(S_\eta)$ is the bounded linear operator defined in (14).

Proof. First, we may re-cast (D) as the convex minimization problem:

$$-d^* = \inf_{\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)} F(\boldsymbol{z}) + G(\mathcal{L}_\Omega \boldsymbol{z}),$$

where $F(\boldsymbol{z}) = -\langle \boldsymbol{y}, \boldsymbol{z} \rangle$ and G is the indicator function of the convex set

$$C = \{\phi \in \mathcal{C}_0(S_\eta) : |\phi(\boldsymbol{w})| \leq \eta(\boldsymbol{w}), \forall \boldsymbol{w} \in S_\eta\}.$$

[2]Technically speaking, this is the pre-dual problem, since the Banach space $\mathcal{M}(S_\eta)$ is not reflexive. However, formally it has the same properties as a convex dual, so we call it the dual problem for simplicity.

Since $\mathcal{L}_\Omega$ is a bounded linear operator whose adjoint is $\mathcal{K}_\Omega : \mathcal{M}(S_\eta) \to \mathbb{C}_{\text{sym}}(\Omega)$, by Fenchel-Rockafellar duality [13, Theorem II.4.1] we have

$$-d^* \geq \sup_{\mu \in \mathcal{M}(S_\eta)} -G^*(\mu) - F^*(\mathcal{K}_\Omega \mu), \tag{15}$$

where G^* and F^* are the convex conjugates of G and F, respectively. By the characterization of the total variation $\|\cdot\|_{TV}$ as a dual norm, we have:

$$\begin{aligned}\|\mu\|_{TV,\eta} = \|\eta \cdot \mu\|_{TV} &= \sup_{\substack{\psi \in \mathcal{C}_0(\mathbb{S}^{D-1}) \\ |\psi(\boldsymbol{w})| \leq 1, \forall \boldsymbol{w} \in \mathbb{S}^{D-1}}} \int_{\mathbb{S}^{D-1}} \psi(\boldsymbol{w}) \eta(\boldsymbol{w}) d\mu(\boldsymbol{w}) \\ &= \sup_{\substack{\phi \in \mathcal{C}_0(S_\eta) \\ |\phi(\boldsymbol{w})| \leq \eta(\boldsymbol{w}), \forall \boldsymbol{w} \in S_\eta}} \int_{S_\eta} \phi(\boldsymbol{w}) d\mu(\boldsymbol{w}) = \sup_{\phi \in C} \langle \mu, \phi \rangle = G^*(\mu).\end{aligned}$$

Additionally, it is straightforward to show F^* is the indicator function of the singleton $\{\boldsymbol{y}\}$. Therefore, eq. (15) is equivalent to

$$d^* \leq p^* := \inf_{\mu \in \mathcal{M}(S_\eta)} \|\mu\|_{TV,\eta} \;\text{ s.t. }\; \mathcal{K}_\Omega \mu = \boldsymbol{y}.$$

Finally, since it is clear that F and G are proper convex lower semi-continuous functions, $F(\boldsymbol{z})$ is finite at $\boldsymbol{z} = \boldsymbol{0}$, and $G(\phi)$ is finite and continuous at $\phi = \mathcal{L}_\Omega \boldsymbol{0} = 0$, by [13, Theorem II.4.1] strong duality holds, i.e., $p^* = d^*$, as claimed. Additionally, [13, Theorem II.4.1] ensures the primal problem has a solution, i.e., there exists a $\mu^* \in \mathcal{M}(S_\eta)$ attaining the infimum, so that the use of min in place of inf in (P_μ) is justified. □

Now we show how the existence of a particular feasible variable for the dual problem (D) ensures optimality of a sparse measure for the primal problem (P_μ). Suppose $\boldsymbol{y} = \mathcal{F}_\Omega f$ where $f = \sum_{i=1}^s a_i [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ such that $a_i \neq 0$, $\boldsymbol{w}_i \in S_\eta$ for all $i \in [s]$, and $\boldsymbol{w}_i \neq \boldsymbol{w}_j$ for all $i \neq j$. Then for $\mu^* = \sum_{i=1}^s a_i \delta_{\boldsymbol{w}_i}$ we have $\|\mu^*\|_{TV,\eta} = \sum_{i=1}^s |a_i| \eta(\boldsymbol{w}_i)$, and by linearity of $\mathcal{F}_\Omega$, we see that

$$\langle \boldsymbol{y}, \boldsymbol{z} \rangle = \langle \mathcal{F}_\Omega f, \boldsymbol{z} \rangle = \sum_{i=1}^s a_i \langle \mathcal{F}_\Omega [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+, \boldsymbol{z} \rangle = \sum_{i=1}^s a_i [\mathcal{L}_\Omega \boldsymbol{z}](\boldsymbol{w}_i).$$

Therefore, if there exists a dual feasible $\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)$ such that $\Phi = \mathcal{L}_\Omega \boldsymbol{z}$ satisfies $\Phi(\boldsymbol{w}_i) = \text{sign}(a_i) \eta(\boldsymbol{w}_i)$ for all $i \in [s]$, we see that

$$\langle \boldsymbol{y}, \boldsymbol{z} \rangle = \sum_{i=1}^s |a_i| \eta(\boldsymbol{w}_i) = \|\mu^*\|_{TV,\eta},$$

i.e., the dual objective matches the primal objective, which implies μ^* is a global minimizer of (P_μ), and so $f = f_{\mu^*}$ is a function space minimizer.

However, this argument does not determine whether f_{μ^*} is the unique function space minimizer of (P_μ). A sufficient condition that guarantees uniqueness is as follows:

Lemma 4. *Suppose $\mu^* = \sum_{i=1}^s a_i \delta_{\boldsymbol{w}_i}$ is feasible for (P_μ) where the $\boldsymbol{w}_i \in S_\eta$ are distinct and $a_i \neq 0$ for all $i \in [s]$. Also, suppose there exists a finite set $\mathcal{W} \subset S_\eta$ such that $\{\boldsymbol{w}_1, ..., \boldsymbol{w}_s\} \subseteq \mathcal{W}$, and for any measure $\nu_0 \in \mathcal{M}(S_\eta)$ with $\operatorname{supp}(\nu_0) \subseteq \mathcal{W}$ we have $\mathcal{K}_\Omega \nu_0 = 0$ implies $f_{\nu_0} = 0$. Furthermore, if there exists $\boldsymbol{z} \in \mathbb{C}_{\text{sym}}(\Omega)$ such that $\Phi = \mathcal{L}_\Omega \boldsymbol{z}$ satisfies*

(i) $|\Phi(\boldsymbol{w})| \leq \eta(\boldsymbol{w})$ *for all* $\boldsymbol{w} \in S_\eta$,

(ii) $\Phi(\boldsymbol{w}_i) = \operatorname{sign}(a_i)\eta(\boldsymbol{w}_i)$ *for all* $i \in [s]$, *and*

(iii) $|\Phi(\boldsymbol{w})| < \eta(\boldsymbol{w})$ *for all* $\boldsymbol{w} \in S_\eta / \mathcal{W}$,

then $f_{\mu^*} = \sum_{i=1}^s a_i [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ is the unique function space minimizer of (P_μ).

Proof. Suppose $\mu \in \mathcal{M}(S_\eta)$ is another measure feasible for (P_μ) (i.e., $\mathcal{K}_\Omega \mu = \boldsymbol{y}$) such that $f_\mu \neq f_{\mu^*}$. Define $\nu = \mu - \mu^*$. By Lebesgue's Decomposition Theorem, there exist measures $\nu_0, \nu_1 \in \mathcal{M}(S_\eta)$ such that $\nu = \nu_0 + \nu_1$ where ν_0 is absolutely continuous with respect to the measure $\delta_{\mathcal{W}} := \sum_{\boldsymbol{w} \in \mathcal{W}} \delta_{\boldsymbol{w}}$ and ν_1 is mutually singular to $\delta_{\mathcal{W}}$. In particular, we have $\operatorname{supp}(\nu_0) \subseteq \mathcal{W}$, and so ν_1 is also mutually singular to μ^*.

First, we show $\nu_1 \neq 0$ by way of contradiction: if $\nu_1 = 0$ then

$$\mathcal{K}_\Omega \nu_0 = \mathcal{K}_\Omega \nu = \mathcal{K}_\Omega \mu - \mathcal{K}_\Omega \mu^* = 0,$$

and so $f_{\nu_0} = 0$ by assumption. This implies $f_\nu = 0$, or equivalently, $f_\mu = f_{\mu^*}$, a contradiction. Therefore, $\nu_1 \neq 0$, and in particular $\|\nu_1\|_{TV,\eta} > 0$ since $\|\cdot\|_{TV,\eta}$ is a norm.

Next, using the fact that ν_1 is mutually singular to both μ^* and ν_0, we have

$$\begin{aligned}
\|\mu\|_{TV,\eta} &= \|\mu^* + \nu\|_{TV,\eta} \\
&= \|\mu^* + \nu_0\|_{TV,\eta} + \|\nu_1\|_{TV,\eta} \\
&> \langle \mu^* + \nu_0, \Phi \rangle + \langle \nu_1, \Phi \rangle \\
&= \langle \mu^*, \Phi \rangle + \langle \nu, \Phi \rangle \\
&= \|\mu^*\|_{TV,\eta} + \langle \nu, \Phi \rangle \\
&= \|\mu^*\|_{TV,\eta} + \langle \nu, \mathcal{L}_\Omega \boldsymbol{z} \rangle \\
&= \|\mu^*\|_{TV,\eta} + \langle \mathcal{K}_\Omega \nu, \boldsymbol{z} \rangle \\
&= \|\mu^*\|_{TV,\eta},
\end{aligned}$$

where the strict inequality comes from the fact that $\|\nu_1\|_{TV,\eta} > \langle \nu_1, \Phi \rangle$ since $\operatorname{supp}(\nu_1) \subset S_\eta / \mathcal{W}$ and the assumption that $|\Phi(\boldsymbol{w})| < \eta(\boldsymbol{w})$ for all $\boldsymbol{w} \in S_\eta / \mathcal{W}$. Therefore, $\|\mu\|_{TV,\eta} > \|\mu^*\|_{TV,\eta}$ for all feasible μ with $f_\mu \neq f_{\mu^*}$, which shows $f_{\mu^*} = \sum_{i=1}^s a_i [\boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)]_+$ is the unique function space minimizer of (P_μ). □

We call the function Φ appearing in Lemma 4 a *dual certificate* for the primal feasible measure μ^* associated with the support set $\mathcal{W}$. The proofs of our exact recovery theorems, given below, rely on explicitly constructing dual certificates.

B Proof of Sampling Theorems

For our proofs in this section, it will be convenient to consider a complex Fourier basis for trigonometric polynomials. For any $\Gamma \subset \mathbb{Z}^d$ that is finite and symmetric (i.e., $\Gamma = -\Gamma$), we let $\mathsf{TP}(\Gamma) \subset L^2(\mathbb{T}^d)$ denote the space of real-valued trigonometric polynomials $\tau : \mathbb{T}^d \to \mathbb{R}$ with Fourier coefficients supported in Γ. Equivalently, we have

$$\mathsf{TP}(\Gamma) = \left\{ \sum_{\boldsymbol{k} \in \Gamma} c[\boldsymbol{k}] e^{i2\pi \boldsymbol{k}^\top \boldsymbol{x}} : \boldsymbol{c} \in \mathbb{C}_{\mathrm{sym}}(\Gamma) \right\},$$

where recall that we use $\mathbb{C}_{\mathrm{sym}}(\Gamma)$ to denote all complex-valued conjugate symmetric arrays indexed by Γ, i.e., $\boldsymbol{c} = (c[\boldsymbol{k}] : \boldsymbol{k} \in \Gamma) \in \mathbb{C}_{\mathrm{sym}}(\Gamma)$ if and only if $c[\boldsymbol{k}] = \overline{c[-\boldsymbol{k}]}$ for all $\boldsymbol{k} \in \Gamma$. Note that $\mathsf{TP}(\Gamma) \simeq \mathbb{C}_{\mathrm{sym}}(\Gamma)$ by identifying any $\tau \in \mathsf{TP}(\Gamma)$ with its array of Fourier coefficients $\mathcal{F}_\Gamma \tau = (\hat{\tau}[\boldsymbol{k}] : \boldsymbol{k} \in \Gamma) \in \mathbb{C}_{\mathrm{sym}}(\Gamma)$.

In particular, suppose the Fourier features embedding $\boldsymbol{\gamma}$ given in (5) is defined in terms of frequencies $\Omega_0 = \{\boldsymbol{k}_1, ..., \boldsymbol{k}_p\} \subset \mathbb{Z}^d$ such that $\boldsymbol{k}_i \neq \pm \boldsymbol{k}_j$ for all $i \neq j$ and $\boldsymbol{k}_i \neq \boldsymbol{0}$ for all $i \in [p]$. Then we have

$$\left\{ \boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot) : \boldsymbol{w} \in \mathbb{R}^{2p+1} \right\} = \mathsf{TP}(\Omega_0^*)$$

where $\Omega_0^* = \Omega_0 \cup -\Omega_0 \cup \{\boldsymbol{0}\}$. Furthermore, in this case the components of $\boldsymbol{\gamma}$ are orthonormal as functions in $L^2(\mathbb{T}^d)$, and so if $\tau = \boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)$ we have $\|\tau\|_{L^2(\mathbb{T}^d)} = \|\boldsymbol{w}\|_2$.

Finally, we will frequently use the fact that if $\tau, \sigma \in \mathsf{TP}(\Gamma)$ where $\Gamma \subset \mathbb{Z}^d$ is finite and symmetric, then by the convolution theorem, their product $\tau\sigma \in \mathsf{TP}(2\Gamma)$ where $2\Gamma := \{\boldsymbol{k} + \boldsymbol{\ell} : \boldsymbol{k}, \boldsymbol{\ell} \in \Gamma\}$, i.e., the Minkowski sum of Γ with itself. For example, if $\Gamma = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_\infty \leq K_0\}$ for some positive integer K_0, then $2\Gamma = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_\infty \leq 2K_0\}$. Likewise, if $\Gamma = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_1 \leq K_0\}$ then $2\Gamma = \{\boldsymbol{k} \in \mathbb{Z}^d : \|\boldsymbol{k}\|_1 \leq 2K_0\}$.

B.1 Proof of Theorem 1

Before giving the proof of Theorem 1, we prove the following key lemma, which shows that a rectified trigonometric polynomial is uniquely identifiable from sufficiently many of its low-pass Fourier samples:

Lemma 5. *Under the same assumptions on Ω_0 and Ω in Theorem 1, the map $\boldsymbol{w} \mapsto \mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$ is injective over the set $\{\boldsymbol{w} \in \mathbb{R}^D : [\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ \neq 0\}$.*

The key idea behind the proof of Lemma 5 is to show that given $\boldsymbol{c} = \mathcal{F}_\Omega[\tau]_+$ where $\tau = \boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)$ we can construct a linear operator $\mathcal{Q}_{\boldsymbol{c}}$ acting on trigonometric polynomials whose nullspace is one-dimensional and spanned by τ. Below we first show how to construct $\mathcal{Q}_{\boldsymbol{c}}$, and then formally prove this claim, and finally give the proof Lemma 5.

First, fix any $\tau \in \mathsf{TP}(\Omega_0^*)$ such that $[\tau]_+ \neq 0$. Let $H : \mathbb{R} \to \mathbb{R}$ denote the Heaviside step function, i.e., $H(t) = 1$ if $t \geq 0$ and $H(t) = 0$ otherwise. Observe that $[\tau]_+ = \tau H(\tau)$, and furthermore

$$\nabla [\tau]_+ = H(\tau) \nabla \tau \quad a.e.$$

Therefore,

$$\tau\nabla[\tau]_+ = \tau H(\tau)\nabla\tau = [\tau]_+\nabla\tau \quad a.e.$$

or, equivalently,

$$\tau\nabla[\tau]_+ - [\tau]_+\nabla\tau = 0 \quad a.e. \tag{16}$$

Passing to Fourier domain, by the convolution theorem, we have

$$\left(\widehat{\nabla[\tau]_+} * \widehat{\tau}\right)[\boldsymbol{k}] - \left(\widehat{[\tau]_+} * \widehat{\nabla\tau}\right)[\boldsymbol{k}] = \mathbf{0}, \quad \forall \boldsymbol{k} \in \mathbb{Z}^d, \tag{17}$$

where $*$ denotes the linear convolution of d-dimensional arrays. Using the fact that $\tau \in \mathsf{TP}(\Omega_0^*)$, and for any weakly differentiable functions g we have $\widehat{\nabla g}[\boldsymbol{k}] = 2\pi i \boldsymbol{k}\widehat{g}[\boldsymbol{k}]$ for all $\boldsymbol{k} \in \mathbb{Z}^d$, we see that (17) is equivalent to

$$\sum_{\boldsymbol{\ell}\in\Omega_0^*} (\boldsymbol{\ell}-\boldsymbol{k})\widehat{[\tau]_+}[\boldsymbol{\ell}-\boldsymbol{k}]\widehat{\tau}[\boldsymbol{\ell}] - \sum_{\boldsymbol{\ell}\in\Omega_0^*} \boldsymbol{\ell}\,\widehat{[\tau]_+}[\boldsymbol{\ell}-\boldsymbol{k}]\,\widehat{\tau}[\boldsymbol{\ell}] = \mathbf{0}, \quad \forall \boldsymbol{k} \in \mathbb{Z}^d. \tag{18}$$

Note that for any fixed frequency $\boldsymbol{k} \in \mathbb{Z}^d$ the equation (18) is linear in terms of the Fourier coefficients $(\widehat{\tau}[\boldsymbol{\ell}] : \boldsymbol{\ell} \in \Omega_0^*)$. Also, given Fourier samples $\boldsymbol{c} = \mathcal{F}_\Omega[\tau]_+ = (\widehat{[\tau]_+}[\boldsymbol{k}] : \boldsymbol{k} \in \Omega)$, the equations in (18) are realizable for all frequencies $\boldsymbol{k} \in \mathbb{Z}^d$ belonging to the set $\Omega|\Omega_0^* := \{\boldsymbol{k} \in \mathbb{Z}^d : \boldsymbol{\ell} - \boldsymbol{k} \in \Omega,\ \forall \boldsymbol{\ell} \in \Omega_0^*\}$.

This motivates the following definition: for any $\boldsymbol{c} \in \mathbb{C}_{\mathrm{sym}}(\Omega)$, define the linear operator $\mathcal{Q}_{\boldsymbol{c}} : \mathsf{TP}(\Omega_0^*) \to \mathbb{C}_{\mathrm{sym}}(\Omega)^d$ that acts on any $\varphi \in \mathsf{TP}(\Omega_0^*)$ by

$$(\mathcal{Q}_{\boldsymbol{c}}\varphi)[\boldsymbol{k}] = \sum_{\boldsymbol{\ell}\in\Omega_0^*} (\boldsymbol{\ell}-\boldsymbol{k})c[\boldsymbol{\ell}-\boldsymbol{k}]\widehat{\varphi}[\boldsymbol{\ell}] - \sum_{\boldsymbol{\ell}\in\Omega_0^*} \boldsymbol{\ell}\,c[\boldsymbol{\ell}-\boldsymbol{k}]\,\widehat{\varphi}[\boldsymbol{\ell}], \ \forall \boldsymbol{k} \in \Omega|\Omega_0^*.$$

In particular, when $\boldsymbol{c} = \mathcal{F}_\Omega[\tau]_+$ working backwards from (18), we see that

$$\mathcal{Q}_{\boldsymbol{c}}\varphi = \mathcal{F}_{\Omega|\Omega_0^*}(\varphi\nabla[\tau]_+ - [\tau]_+\nabla\varphi),$$

and by the identity given in (16), we see that $\mathcal{Q}_{\boldsymbol{c}}\tau = 0$, i.e., τ belongs to the nullspace of $\mathcal{Q}_{\boldsymbol{c}}$. The next result shows that when $\boldsymbol{c} = \mathcal{F}_\Omega[\tau]_+$, and Ω is sufficiently large, the nullspace of $\mathcal{Q}_{\boldsymbol{c}}$ is one-dimensional and spanned by τ.

Lemma 6. *Suppose* $\Omega|\Omega_0^* \supseteq 2\Omega_0^*$. *Let* $\boldsymbol{c} = \mathcal{F}_\Omega[\tau]_+$ *where* $\tau \in \mathsf{TP}(\Omega_0^*)$ *with* $[\tau]_+ \neq 0$. *Then* $\ker \mathcal{Q}_{\boldsymbol{c}} = \mathrm{span}\{\tau\}$.

Proof. The condition $\varphi \in \ker \mathcal{Q}_{\boldsymbol{c}}$ is equivalent to

$$\mathcal{F}_{\Omega|\Omega_0^*}(\varphi\nabla[\tau]_+ - [\tau]_+\nabla\varphi) = \mathbf{0}.$$

This implies for all vector-valued functions $\boldsymbol{\rho} \in \mathsf{TP}(\Omega|\Omega_0^*)^d$, i.e., $\boldsymbol{\rho} = (\rho_1, ..., \rho_d)$ where $\rho_1, ..., \rho_d \in \mathsf{TP}(\Omega|\Omega_0^*)$, we have

$$\int_{\mathbb{T}^d} (\varphi\nabla[\tau]_+ - [\tau]_+\nabla\varphi)^\top \boldsymbol{\rho}\, d\boldsymbol{x} = 0. \tag{19}$$

Recalling that $[\tau]_+ = H(\tau)\tau$ and $\nabla[\tau]_+ = H(\tau)\nabla\tau$ a.e., we see that equation (19) is equivalent to

$$\int_{\mathbb{T}^d} H(\tau)(\varphi\nabla\tau - \tau\nabla\varphi)^\top \boldsymbol{\rho}\, d\boldsymbol{x} = 0.$$

Consider the specific choice $\boldsymbol{\rho} = \varphi\nabla\tau - \tau\nabla\varphi$. Then the components of $\boldsymbol{\rho}$ belong to $\mathsf{TP}(2\Omega_0^*)$, and so $\boldsymbol{\rho} \in \mathsf{TP}(\Omega|\Omega_0^*)^d$ by the assumption that $\Omega|\Omega_0^* \supseteq 2\Omega_0^*$. Therefore, this implies

$$\int_{\mathbb{T}^d} H(\tau)\|\varphi\nabla\tau - \tau\nabla\varphi\|_2^2\, d\boldsymbol{x} = 0, \tag{20}$$

which shows $\varphi(\boldsymbol{x})\nabla\tau(\boldsymbol{x}) - \tau(\boldsymbol{x})\nabla\varphi(\boldsymbol{x}) = \mathbf{0}$ for all $\boldsymbol{x}$ in the open set $P := \{\boldsymbol{x} \in \mathbb{T}^d : \tau(\boldsymbol{x}) > 0\}$. Finally, we will show this implies $\varphi(\boldsymbol{x}) = c \cdot \tau(\boldsymbol{x})$ for some constant $c \in \mathbb{R}$.

Let B be any closed ball contained in P. By continuity, τ must be bounded away from zero over B, hence $\frac{\varphi}{\tau}$ is smooth over B. Therefore, by the quotient rule

$$\nabla\left(\frac{\varphi}{\tau}\right)(\boldsymbol{x}) = \frac{\tau(\boldsymbol{x})\nabla\varphi(\boldsymbol{x}) - \varphi(\boldsymbol{x})\nabla\tau(\boldsymbol{x})}{\tau(\boldsymbol{x})^2} = \mathbf{0} \quad \forall \boldsymbol{x} \in B.$$

This implies $\frac{\varphi}{\tau}$ is constant over B, and so there exists a constant $c \in \mathbb{R}$ such that $\varphi(\boldsymbol{x}) = c \cdot \tau(\boldsymbol{x})$ for all $\boldsymbol{x} \in B$, or equivalently, $\varphi(\boldsymbol{x}) - c \cdot \tau(\boldsymbol{x}) = 0$ for all $\boldsymbol{x} \in B$. However, since a trigonometric polynomial is a real analytic function, it cannot vanish on a ball contained $\mathbb{T}^d$ unless it is identically zero. Therefore, it must be the case that $\varphi(\boldsymbol{x}) = c \cdot \tau(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathbb{T}^d$, as claimed. □

Now we prove Lemma 5.

Proof of Lemma 5. Given $\boldsymbol{w}_0, \boldsymbol{w}_1 \in S_\eta$, define trigonometric polynomials $\tau_0 = \boldsymbol{w}_0^\top\boldsymbol{\gamma}(\cdot)$ and $\tau_1 = \boldsymbol{w}_1^\top\boldsymbol{\gamma}(\cdot)$, and suppose for $\boldsymbol{c}_0 = \mathcal{F}_\Omega[\tau_0]_+$ and $\boldsymbol{c}_1 = \mathcal{F}_\Omega[\tau_1]_+$ we have $\boldsymbol{c}_0 = \boldsymbol{c}_1$. Then $\mathcal{Q}_{\boldsymbol{c}_0} = \mathcal{Q}_{\boldsymbol{c}_1}$ and so $\ker \mathcal{Q}_{\boldsymbol{c}_0} = \ker \mathcal{Q}_{\boldsymbol{c}_1}$. By Lemma 6 this implies $\tau_0 = c \cdot \tau_1$ for some $c \in \mathbb{R}$. Since $\|\tau_0\|_{L^2(\mathbb{T}^d)} = \|\boldsymbol{w}_0\|_2 = 1$ and $\|\tau_1\|_{L^2(\mathbb{T}^d)} = \|\boldsymbol{w}_1\|_2 = 1$, we must have $c = \pm 1$. Suppose it were the case that $c = -1$, i.e., $\tau_1 = -\tau_0$, and so $\mathcal{F}_\Omega[\tau_0]_+ = \mathcal{F}_\Omega[-\tau_0]_+$. By the identity $[t]_+ - [-t]_+ = t$ for all $t \in \mathbb{R}$, we see this implies $\mathcal{F}_\Omega([\tau_0]_+ - [-\tau_0]_+) = \mathcal{F}_\Omega\tau_0 = \mathbf{0}$. Since $\Omega_0 \subset \Omega$, this can only be the case if $\tau_0 = 0$, which is ruled out by our assumption that $[\tau_0]_+ \neq 0$. Therefore, $c = 1$, and so $\tau_0 = \tau_1$ or equivalently, $\boldsymbol{w}_0 = \boldsymbol{w}_1$, which proves the claim. □

Finally, we prove Theorem 1 by constructing a dual certificate and invoking Lemma 4.

Proof of Theorem 1. Let $\eta(\boldsymbol{w}) = \|\mathcal{F}_\Omega[\boldsymbol{w}^\top\boldsymbol{\gamma}(\cdot)]_+\|_2$ and $f = a_1[\boldsymbol{w}_1^\top\boldsymbol{\gamma}(\cdot)]_+$. Define $\boldsymbol{w}_0 = \boldsymbol{w}_1/\|\boldsymbol{w}_1\|$ and $a_0 = \|\boldsymbol{w}_1\|a_1$. Then we have $f = a_0[\boldsymbol{w}_0^\top\boldsymbol{\gamma}(\cdot)]_+$ with $\boldsymbol{w}_0 \in S_\eta$. We will show that $\boldsymbol{z}_0 := \operatorname{sign}(a_0)\eta(\boldsymbol{w}_0)^{-1}\mathcal{F}_\Omega[\boldsymbol{w}_0^\top\boldsymbol{\gamma}(\cdot)]_+$ yields a valid dual certificate $\Phi = \mathcal{L}_\Omega\boldsymbol{z}_0$ for the measure $\mu = a_0\delta_{\boldsymbol{w}_0}$ with respect to the support set $\mathcal{W} = \{\boldsymbol{w}_0\}$.

First, observe that for all $\boldsymbol{w} \in S_\eta$ we have

$$\Phi(\boldsymbol{w}) = \operatorname{sign}(a_0)\eta(\boldsymbol{w}_0)^{-1}\langle\mathcal{F}_\Omega[\boldsymbol{w}^\top\boldsymbol{\gamma}(\cdot)]_+, \mathcal{F}_\Omega[\boldsymbol{w}_0^\top\boldsymbol{\gamma}(\cdot)]_+\rangle.$$

Plugging in $\boldsymbol{w} = \boldsymbol{w}_0$, and since $\eta(\boldsymbol{w}_0) = \|\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+\|_2$, we see that

$$\Phi(\boldsymbol{w}_0) = \text{sign}(a_0)\eta(\boldsymbol{w}_0)^{-1}\|\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+\|_2^2 = \text{sign}(a_0)\eta(\boldsymbol{w}_0).$$

Also, by the Cauchy-Schwarz inequality, for all $\boldsymbol{w} \in S_\eta$ we have

$$|\Phi(\boldsymbol{w})| \leq \eta(\boldsymbol{w}_0)^{-1}\|\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2 \, \|\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+\|_2 = \eta(\boldsymbol{w}),$$

with strict inequality when $\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+$ is not co-linear with $\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$. Therefore, to prove Φ is a valid dual certificate, it suffices to show $\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+$ is not co-linear with $\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$ when $\boldsymbol{w} \neq \boldsymbol{w}_0$.

Towards this end, suppose $\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ = c \cdot \mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+$ for some $c \in \mathbb{R}$. First, note that for all $\boldsymbol{v} \in S_\eta$ we have $\widehat{[\boldsymbol{v}^\top \boldsymbol{\gamma}(\cdot)]_+}[\boldsymbol{0}] = \int_{\mathbb{T}^d} [\boldsymbol{v}^\top \boldsymbol{\gamma}(\cdot)]_+ d\boldsymbol{x} > 0$, since otherwise this would imply $[\boldsymbol{v}^\top \boldsymbol{\gamma}(\cdot)]_+$ is identically zero, and so $\eta(\boldsymbol{v}) = \|\mathcal{F}_\Omega[\boldsymbol{v}^\top \boldsymbol{\gamma}(\cdot)]_+\|_2 = 0$, contrary to assumption that $\boldsymbol{v} \in S_\eta$. This implies $c = \widehat{[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+}[\boldsymbol{0}] / \widehat{[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+}[\boldsymbol{0}] > 0$. Therefore, by linearity of $\mathcal{F}_\Omega$ and positive 1-homogeneity of the ReLU activation, we have $\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ = \mathcal{F}_\Omega[(c\boldsymbol{w}_0)^\top \boldsymbol{\gamma}(\cdot)]_+$. However, by Lemma 5, we see this implies $\boldsymbol{w} = c\boldsymbol{w}_0$, and since we assume $\|\boldsymbol{w}\|_2 = \|\boldsymbol{w}_0\|_2 = 1$, this can be the case only if $c = 1$, and so $\boldsymbol{w} = \boldsymbol{w}_0$. Therefore, if $\boldsymbol{w} \in S_\eta$ is such that $\boldsymbol{w} \neq \boldsymbol{w}_0$ then $\mathcal{F}_\Omega[\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+$ is not co-linear with $\mathcal{F}_\Omega[\boldsymbol{w}_0^\top \boldsymbol{\gamma}(\cdot)]_+$, and so $|\Phi(\boldsymbol{w})| < \eta(\boldsymbol{w})$. This proves Φ is a dual certificate for the measure $\mu = a\delta_{\boldsymbol{w}_0}$, and so by Lemma 4 (with $\mathcal{W} = \{\boldsymbol{w}_0\}$), $f = f_\mu$ is the unique function space minimizer of (P_μ). Finally, since μ is 1-sparse, f is also the unique function space minimizer of $(P_{\mu,W})$ for all $W \geq 1$, and hence also of $(P_{\theta,W})$, since $(P_{\theta,W})$ has the same function space minimizers as $(P_{\mu,W})$. $\square$

B.2 Proof of Theorem 2

The proof of Theorem 2 relies on some basic algebraic facts regarding the zero-sets of trigonometric polynomials as originally developed in [35], which we summarize here.

First, we say $\tau : \mathbb{T}^2 \to \mathbb{C}$ is a *2D trigonometric polynomial* if it has the form

$$\tau(x_1, x_2) = \sum_{(k_1,k_2)\in\Gamma} c_{(k_1,k_2)} e^{i2\pi k_1 x_1} e^{i2\pi k_2 x_2}, \quad (x_1, x_2) \in \mathbb{T}^2, \tag{21}$$

for some constants $c_{(k_1,k_2)} \in \mathbb{C}$ indexed by a finite subset $\Gamma \subset \mathbb{Z}^2$, and use $Z(\tau)$ denote its zero set

$$Z(\tau) = \{(x_1, x_2) \in \mathbb{T}^2 : \tau(x_1, x_2) = 0\}.$$

Following [35], we say $C \subset \mathbb{T}^2$ is a *trigonometric curve* if $C = Z(\tau)$ for some non-zero 2D trigonometric polynomial τ, and C is an infinite set that has no isolated points. In [35] it is shown that every trigonometric curve is a union of finitely many piecewise smooth curves that are closed and connected as subsets of $\mathbb{T}^2$.

Given any 2D trigonometric polynomial τ of the form (21) we may associate a unique complex polynomial in two variables $\mathcal{P}[\tau] \in \mathbb{C}[z_1, z_2]$ by making the substitutions $e^{i2\pi x_1} \mapsto z_1$ and $e^{i2\pi x_2} \mapsto z_2$, and multiplying by the smallest powers of z_1 and z_2 so that the resulting

function has no negative powers of z_1 or z_2. In particular, if τ has the form (21) then $p = \mathcal{P}[\tau]$ has the form

$$p(z_1, z_2) = z_1^{\ell_1} z_2^{\ell_2} \left(\sum_{(k_1,k_2)\in\Gamma} c_{(k_1,k_2)} z_1^{k_1} z_2^{k_2} \right), \quad (z_1, z_2) \in \mathbb{C}^2,$$

for some non-negative integers ℓ_1 and ℓ_2. Note $\tau(x_1, x_2) = 0$ if and only if $p(e^{i2\pi x_1}, e^{i2\pi x_2}) = 0$.

Following [35], we say the 2D trigonometric polynomial τ is *irreducible* if $p = \mathcal{P}[\tau]$ is irreducible as a polynomial in $\mathbb{C}[z_1, z_2]$, i.e., p cannot be written as a product of two or more non-constant polynomials in $\mathbb{C}[z_1, z_2]$. Also, we say the 2D trigonometric polynomial τ_0 is a *minimal polynomial* for the trigonometric curve C if $C = Z(\tau_0)$ and $p_0 = \mathcal{P}[\tau_0]$ is a product of non-constant irreducible polynomials $p_1, ..., p_n \in \mathbb{C}[z_1, z_2]$ that are not scalar multiples of each other and the zero set of each p_i has infinite intersection with the complex unit torus $\mathbb{CT}^2 = \{(z_1, z_2) \in \mathbb{C}^2 : |z_1| = |z_2| = 1\}$. In [35] it is shown there is a unique (up to scaling) real-valued minimal polynomial associated with every trigonometric curve. In particular, if $C = Z(\tau)$ is an infinite set and τ is real-valued and irreducible, then τ is a minimal polynomial for C.

Additionally, we will need the following facts regarding irreducible and minimal polynomials.

Lemma 7 (Adapted from [35])**.**

(a) Suppose τ, σ are irreducible trigonometric polynomials and $\tau \neq c \cdot \sigma$ for any $c \in \mathbb{R}$. Then $Z(\tau) \cap Z(\sigma)$ is a finite set.

(b) Suppose τ is a minimal polynomial for the trigonometric curve C. Then $\nabla\tau(\boldsymbol{x}) = \mathbf{0}$ for at most finitely many $\boldsymbol{x} \in C$.

(c) Suppose τ is a minimal polynomial for the trigonometric curve C, and σ is a trigonometric polynomial that vanishes on C. Then τ divides σ, i.e., there exists a trigonometric polynomial β such that $\sigma = \tau \cdot \beta$.

Part (a) is special case of [35, Corollary A.2], part (b) is [35, Proposition A.4], and part (c) is [35, Proposition A.3].

Let $\Gamma_1 = \{[\pm 1, 0]^\top, [0, \pm 1]^\top\}$. Then $\mathsf{TP}(\Gamma_1)$ is the space of 2D trigonometric polynomials realized by the restricted Fourier features embedding $\boldsymbol{\gamma}$ defined in (12). In particular, $\tau \in \mathsf{TP}(\Gamma_1)$ if and only if

$$\tau(x_1, x_2) = w_1 \cos(2\pi x_1) + w_2 \sin(2\pi x_1) + w_3 \cos(2\pi x_2) + w_4 \sin(2\pi x_2) \tag{22}$$

for some $w_1, w_2, w_3, w_4 \in \mathbb{R}$.

Our first result shows that every $\tau \in \mathsf{TP}(\Gamma_1)$ is such that $C = Z(\tau)$ is a trigonometric curve for which τ is a minimal polynomial, and characterizes when τ is irreducible. The proof of this result, which relies on elementary calculus and algebra, is provided in the Supplementary Materials F.

Lemma 8. *Suppose $\tau \in \mathsf{TP}(\Gamma_1)$ is non-zero. Then $C = Z(\tau)$ is a trigonometric curve, and τ is a minimal polynomial for C. Furthermore, τ is irreducible except in the following cases:*

1. *$\tau(x_1, x_2) = \alpha \sin(2\pi(x_1 - t))$ for some $\alpha \in \mathbb{R}$, $t \in \mathbb{T}$.*
2. *$\tau(x_1, x_2) = \alpha \sin(2\pi(x_2 - t))$ for some $\alpha \in \mathbb{R}$, $t \in \mathbb{T}$.*
3. *$\tau(x_1, x_2) = \alpha(\sin(2\pi(x_1 - t_1)) + \sin(2\pi(x_2 - t_2)))$ for some $\alpha \in \mathbb{R}$, $t_1, t_2 \in \mathbb{T}$.*

Together, these cases correspond to all τ of the form (22) where the weight vector $\boldsymbol{w} = [w_1, w_2, w_3, w_4]^\top$ belongs to the set V defined in (13).

Next, we prove a key lemma, which shows that given any $\tau \in \mathsf{TP}(\Gamma_1)$, a second-order differential operator applied to $[\tau]_+$ yields a distribution (i.e., generalized function [18]) supported on $Z(\tau)$. Later, we use this property to construct a dual certificate.

Below, for a distribution D and $\mathcal{C}^\infty$-smooth test function $\varphi : \mathbb{T}^2 \to \mathbb{R}$, we let $\langle D, \varphi \rangle$ denote the evaluation of D at φ. In particular, when D can be identified with a function $g \in L^1(\mathbb{T}^2)$, we have $\langle D, \varphi \rangle = \int_{\mathbb{T}^2} g\varphi \, d\boldsymbol{x}$.

Lemma 9. *Let $\tau \in \mathsf{TP}(\Gamma_1)$, and let $\Delta = \frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2}$ denote the Laplacian. Then $(\Delta + 4\pi^2)[\tau]_+$ is a distribution acting on $\mathcal{C}^\infty$-smooth test functions $\varphi : \mathbb{T}^2 \to \mathbb{R}$ by*

$$\langle (\Delta + 4\pi^2)[\tau]_+, \varphi \rangle = \oint_{Z(\tau)} \varphi \|\nabla \tau\| ds,$$

where $\oint_C (\cdot) ds$ denotes the line integral of a function over the curve C.

Proof. First, observe that $\Delta\tau = -4\pi^2\tau$ and so $(\Delta + 4\pi^2)\tau = 0$. Also, we have $\nabla[\tau]_+ = H(\tau)\nabla\tau$ almost everywhere. Now, for any C^∞-smooth test function $\varphi : \mathbb{T}^2 \to \mathbb{R}$ we have

$$\begin{aligned}
\langle \Delta[\tau]_+, \varphi \rangle &= -\langle \nabla[\tau]_+, \nabla\varphi \rangle \\
&= -\langle H(\tau)\nabla\tau, \nabla\varphi \rangle \\
&= -\langle H(\tau), \nabla\tau \cdot \nabla\varphi \rangle \\
&= -\langle H(\tau), \nabla \cdot (\varphi\nabla\tau) - \varphi\Delta\tau \rangle \qquad (23) \\
&= -\langle H(\tau), \nabla \cdot (\varphi\nabla\tau) \rangle + \langle H(\tau), \varphi\Delta\tau \rangle \\
&= -\langle H(\tau), \nabla \cdot (\varphi\nabla\tau) \rangle - 4\pi^2 \langle H(\tau), \varphi\tau \rangle \\
&= -\langle H(\tau), \nabla \cdot (\varphi\nabla\tau) \rangle - 4\pi^2 \langle [\tau]_+, \varphi \rangle.
\end{aligned}$$

where in (23) we used the identity $\nabla \cdot (\varphi\nabla\tau) = \nabla\tau \cdot \nabla\varphi + \varphi\Delta\tau$. Rearranging terms above gives

$$\begin{aligned}
\langle (\Delta + 4\pi^2)\tau, \varphi \rangle &= -\langle H(\tau), \nabla \cdot (\varphi\nabla\tau) \rangle \\
&= -\int_{\tau > 0} \nabla \cdot (\varphi\nabla\tau) \, d\boldsymbol{x} \\
&= -\oint_{Z(\tau)} \phi(\nabla\tau \cdot \boldsymbol{n}) \, ds
\end{aligned}$$

where $\boldsymbol{n}$ is an outward unit normal to the curve $Z(\tau)$, and this last step follows from the divergence theorem. Since τ is a minimal polynomial for the trigonometric curve $Z(\tau)$, $\nabla\tau(\boldsymbol{x}) = \boldsymbol{0}$ for at most finitely many $\boldsymbol{x} \in Z(\tau)$. Therefore, $\boldsymbol{n}(\boldsymbol{x}) = -\frac{\nabla\tau(\boldsymbol{x})}{\|\nabla\tau(\boldsymbol{x})\|}$ for all $\boldsymbol{x} \in Z(\tau)$ except possibly finite many points, and so we have

$$\langle(\Delta + 4\pi^2)\tau, \varphi\rangle = \oint_{Z(\tau)} \varphi\|\nabla\tau\|\, ds,$$

as claimed. □

Now we proceed with proof of Theorem 2, which we split into two lemmas. First, we prove that, for sufficient large sampling set Ω, the injectivity condition of Lemma 4 is satisfied assuming the parameter vectors $\boldsymbol{w}_i$ generate irreducible 2D trigonometric polynomials.

Lemma 10. *Let $\boldsymbol{w}_1, ..., \boldsymbol{w}_s \in \mathbb{S}^3$ be such that $\boldsymbol{w}_i \neq \pm\boldsymbol{w}_j$ for all $i \neq j$, and $\tau_i = \boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot) \in \mathsf{TP}(\Gamma_1)$ is irreducible. Suppose $\Omega \supseteq \{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_1 \leq 2(s-1)\}$. Then for any measure $\nu \in \mathcal{M}(\mathbb{S}^3)$ supported on the set $\mathcal{W} = \{\pm\boldsymbol{w}_1, ..., \pm\boldsymbol{w}_s\}$, we have $\mathcal{K}_\Omega\nu = 0$ implies $f_\nu = 0$.*

Proof. Assume ν has the form

$$\nu = \sum_{i=1}^{s} \left(a_i\delta_{\boldsymbol{w}_i} + b_i\delta_{-\boldsymbol{w}_i}\right),$$

for some scalars $a_i, b_i \in \mathbb{R}$, $i \in [s]$. Define $\tau_i = \boldsymbol{w}_i^\top\boldsymbol{\gamma}(\cdot)$ for all $i \in [s]$. If $\mathcal{K}_\Omega\nu = 0$ then

$$\sum_{i=1}^{s} \left(a_i\mathcal{F}_\Omega[\tau_i]_+ + b_i\mathcal{F}_\Omega[-\tau_i]_+\right) = \mathcal{F}_\Omega\left(\sum_{i=1}^{s} a_i[\tau_i]_+ + b_i[-\tau_i]_+\right) = 0. \tag{24}$$

This implies for all $\varphi \in \mathsf{TP}(\Omega)$ we have

$$\left\langle \sum_{i=1}^{s} a_i[\tau_i]_+ + b_i[-\tau_i]_+, \varphi \right\rangle = \sum_{i=1}^{s} a_i\langle[\tau_i]_+, \varphi\rangle + b_i\langle[-\tau_i]_+, \varphi\rangle = 0. \tag{25}$$

Now, fix any $j \in [s]$, and suppose we set $\varphi = (\Delta + 4\pi^2)\psi_j$ where $\psi_j = \prod_{i\neq j}\tau_i^2$, which has frequency support contained in $\{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_1 \leq 2(s-1)\}$, hence belongs to $\mathsf{TP}(\Omega)$ by assumption. Then, starting from equation (25) we have

$$\begin{aligned}
0 &= \sum_{i=1}^{s} a_i\langle[\tau_i]_+, (\Delta + 4\pi^2)\psi_j\rangle + b_i\langle[-\tau_i]_+, (\Delta + 4\pi^2)\psi_j\rangle \\
&= \sum_{i=1}^{s} a_i\langle(\Delta + 4\pi^2)[\tau_i]_+, \psi_j\rangle + b_i\langle(\Delta + 4\pi^2)[-\tau_i]_+, \psi_j\rangle \\
&= \sum_{i=1}^{s}(a_i + b_i) \oint_{Z(\tau_i)} \psi_j\|\nabla\tau_i\| ds \qquad (26) \\
&= (a_j + b_j) \oint_{Z(\tau_j)} \psi_j\|\nabla\tau_j\| ds, \qquad (27)
\end{aligned}$$

where in (26) we applied Lemma 9, and in (27) all but the jth term in the sum vanishes because ψ_j contains a factor of τ_i for all $i \neq j$. Finally, the integrand in (27) is non-negative and not identically zero since $Z(\psi_j) = \cup_{i\neq j} Z(\tau_j)$, and by Lemma 7(a) each of the τ_i can have only finitely many zeros on $Z(\tau_j)$, and by Lemma 7(b), $\nabla\tau_j$ has at most finitely many zeros over $Z(\tau_j)$. Therefore, the integral in (27) non-zero, and so $a_j + b_j = 0$, or equivalently, $b_j = -a_j$ for all $j \in [s]$. However, by the identity $\tau_i = [\tau_i]_+ - [-\tau_i]_+$, from equation (24), this implies

$$\mathcal{F}_\Omega \left(\sum_{i=1}^s a_i \tau_i \right) = 0.$$

But since $\Gamma_1 \subset \Omega$, this gives $\sum_{i=1}^s a_i \tau_i = 0$, and so

$$f_\nu = \int_{S_\eta} [\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+ d\nu(\boldsymbol{w}) = \sum_{i=1}^s a_i \tau_i = 0,$$

as claimed. □

Finally, to finish the proof of Theorem 2, we construct a dual certificate assuming that f is a positive weighted linear combination of rectified *irreducible* trigonometric polynomials belonging to $\mathsf{TP}(\Gamma_1)$.

Lemma 11. *Let $\boldsymbol{w}_1, ..., \boldsymbol{w}_s \in \mathbb{S}^3$ be such that $\boldsymbol{w}_i \neq \pm \boldsymbol{w}_j$ for all $i \neq j$, and $\tau_i = \boldsymbol{w}_i^\top \boldsymbol{\gamma}(\cdot)$ is irreducible for all $i \in [s]$. Suppose $f = \sum_{i=1}^s a_i [\tau_i]_+$ for some scalars $a_i > 0$, and $\boldsymbol{y} = \mathcal{F}_\Omega f$ where $\Omega \supseteq \{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_1 \leq 2s\}$. Let the weighting function η be given by $\eta(\boldsymbol{w}) = \int_{\mathbb{T}^2} [\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ d\boldsymbol{x}$. Then there exists a dual certificate for the measure $\mu = \sum_{i=1}^s a_i \delta_{\boldsymbol{w}_i}$ with respect to the support set $\mathcal{W} = \{\pm \boldsymbol{w}_1, ..., \pm \boldsymbol{w}_s\}$.*

Proof. Let $\rho = \prod_{i=1}^s \tau_i^2$. Then $\hat{\rho}[\boldsymbol{k}]$ has support contained within $\{\boldsymbol{k} \in \mathbb{Z}^2 : \|\boldsymbol{k}\|_1 \leq 2s\}$, and so $\rho \in \mathsf{TP}(\Omega)$ by assumption. Also, let $\varphi = 1 - \alpha(\Delta + 4\pi^2)\rho$, where $\alpha > 0$ is a scale factor to be determined later. Since the Laplacian Δ is a Fourier multiplier, we also have $\varphi \in \mathsf{TP}(\Omega)$. Finally, let $\boldsymbol{z} = \mathcal{F}_\Omega \varphi$, and let $\Phi(\boldsymbol{w}) = \mathcal{L}_\Omega \boldsymbol{z} = \langle \mathcal{F}_\Omega [\boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)]_+, \boldsymbol{z} \rangle$ be the corresponding dual function.

Now fix any $\boldsymbol{w} \in \mathbb{S}^3$, and let $\tau = \boldsymbol{w}^\top \boldsymbol{\gamma}(\cdot)$. Then, since $\mathcal{F}_\Omega^* \boldsymbol{z} = \mathcal{F}_\Omega^* \mathcal{F}_\Omega \varphi = \varphi$, we have

$$\Phi(\boldsymbol{w}) = \langle \mathcal{F}_\Omega [\tau]_+, \boldsymbol{z} \rangle = \langle [\tau]_+, \varphi \rangle = \eta(\boldsymbol{w}) - \alpha \langle (\Delta + 4\pi^2)[\tau]_+, \rho \rangle, \tag{28}$$

where in the last equality we used the fact that $\langle [\tau]_+, 1 \rangle = \int_{\mathbb{T}^2} [\boldsymbol{w}^\top \boldsymbol{\gamma}(\boldsymbol{x})]_+ \, d\boldsymbol{x} = \eta(\boldsymbol{w})$.

Next, by Lemma 9, observe that

$$\langle (\Delta + 4\pi^2)[\tau]_+, \rho \rangle = \oint_{Z(\tau)} \rho \|\nabla \tau\| ds \geq 0,$$

since $\rho \geq 0$. Furthermore, we have the bound

$$\begin{aligned}
\oint_{Z(\tau)} \rho \|\nabla \tau\| ds &\leq \|\rho\|_{L^\infty(\mathbb{T}^2)} \oint_{Z(\tau)} \|\nabla \tau\| ds \\
&= \|\rho\|_{L^\infty(\mathbb{T}^2)} \langle (\Delta + 4\pi^2)[\tau]_+, 1 \rangle \qquad (29) \\
&= \|\rho\|_{L^\infty(\mathbb{T}^2)} \langle [\tau]_+, (\Delta + 4\pi^2)1 \rangle \qquad (30) \\
&= 4\pi^2 \|\rho\|_{L^\infty(\mathbb{T}^2)} \langle [\tau]_+, 1 \rangle \\
&= 4\pi^2 \|\rho\|_{L^\infty(\mathbb{T}^2)} \eta(\boldsymbol{w}),
\end{aligned}$$

where the equality in (29) follows from Lemma 9, and the equality in (30) follows from the definition of distributional derivatives. Therefore, we have shown

$$0 \leq \langle (\Delta + 4\pi^2)[\tau]_+, \rho \rangle \leq 4\pi^2 \|\rho\|_{L^\infty(\mathbb{T}^2)} \eta(\boldsymbol{w}).$$

Choosing α to be any constant satisfying $0 < \alpha < (4\pi^2 \|\rho\|_{L^\infty(\mathbb{T}^2)})^{-1}$, then (28) combined with the inequality above shows that $0 \leq \Phi(\boldsymbol{w}) \leq \eta(\boldsymbol{w})$ for all $\boldsymbol{w} \in \mathbb{S}^3$.

Finally, we prove $\Phi(\boldsymbol{w}) = \eta(\boldsymbol{w})$ if and only $\boldsymbol{w} = \pm \boldsymbol{w}_i$ for some $i \in [s]$. From (28), this is equivalent to showing

$$\langle (\Delta + 4\pi^2)[\tau]_+, \rho \rangle = \oint_{Z(\tau)} \left(\textstyle\prod_{i=1}^s \tau_i^2 \right) \|\nabla \tau\| ds = 0 \qquad (31)$$

if and only if $\tau = \pm \tau_i$ for some $i \in [s]$. First, if $\tau = \pm \tau_i$ then clearly (31) holds. Conversely, if (31) holds, then the integrand $(\prod_{i=1}^s \tau_i^2) \|\nabla \tau\|$ must vanish over $Z(\tau)$. However, by Lemma 7(b), $\nabla \tau$ has at most finitely zeros over $Z(\tau)$, and so $\prod_{i=1}^s \tau_i^2$ must vanish everywhere on $Z(\tau)$. Since τ is a minimal polynomial for $Z(\tau)$, by Lemma 7(c) this implies τ divides $\prod_{i=1}^s \tau_i^2$. Furthermore, since the τ_i are irreducible, and τ must be square-free (by virtue of being a minimal polynomial), this is possible if and only if $\tau = c \prod_{i \in I} \tau_i$ for some constant $c \in \mathbb{R}$ and index set $I \subset [s]$. Finally, since the product of two or more τ_i has Fourier support extending outside Γ_1, the only possibility is $\tau = c \cdot \tau_i$ for some $i \in [s]$, which implies $\boldsymbol{w} = c\boldsymbol{w}_i$ for some i. Finally, since $\|\boldsymbol{w}\|_2 = 1$, we have $c = \pm 1$, which proves the claim.

Therefore, we have shown $\Phi(\boldsymbol{w}) = \eta(\boldsymbol{w})$ for all $\boldsymbol{w} = \pm \boldsymbol{w}_i$, $i \in [s]$, while $0 \leq \Phi(\boldsymbol{w}) < \eta(\boldsymbol{w})$ for all $\boldsymbol{w}$ such that $\boldsymbol{w} \neq \pm \boldsymbol{w}_i$. This shows Φ satisfies the conditions of Lemma 4, i.e., Φ is a dual certificate for the measure $\mu^* = \sum_{i=1}^s a_i \delta_{\boldsymbol{w}_i}$ with respect to the support set $\mathcal{W} = \{\pm \boldsymbol{w}_1, ..., \pm \boldsymbol{w}_s\}$. □

C INR fitting with Augmented Lagrangian Method

Here we describe an Augmented Lagrangian (AL) approach to solving the equality constrained INR fitting problem (3) (using the approximation $\tilde{\mathcal{F}}_\Omega$ in place of $\mathcal{F}_\Omega$). For INR parameters $\theta \in \Theta_W$, we introduce a Lagrange multiplier vector $\boldsymbol{q} \in \mathrm{Im}(\mathcal{F}_\Omega)$ and penalty parameter $\sigma > 0$, and define the AL loss function $AL(\theta, \boldsymbol{q}, \sigma)$ associated with (3) by

$$AL(\theta, \boldsymbol{q}, \sigma) := R(\theta) + \langle \boldsymbol{q}, \tilde{\mathcal{F}}_\Omega f_\theta - \boldsymbol{y} \rangle + \frac{\sigma}{2} \|\tilde{\mathcal{F}}_\Omega f_\theta - \boldsymbol{y}\|_2^2.$$

Then starting from initializations θ_0, $\boldsymbol{q}_0$, and σ_0, for $n = 0, 1, 2, ..., N-1$, we perform the updates

$$\begin{aligned} \theta_{n+1} &= \arg\min_{\theta} AL(\theta, \boldsymbol{q}_n, \sigma_n) \\ \boldsymbol{q}_{n+1} &= \boldsymbol{q}_n + \sigma_n(\tilde{\mathcal{F}}_\Omega f_{\theta_{n+1}} - \boldsymbol{y}) \\ \sigma_{n+1} &= \gamma\sigma_n, \end{aligned} \tag{32}$$

where $\gamma > 1$ is a fixed constant. To approximate the solution of the subproblem (32) we use a standard gradient-based INR training algorithm, warm-starting from the parameters θ_n obtained in the previous outer-loop iteration. See the Supplementary Materials E.1 for more details on the hyperparameter settings used in the experiments in Section 4.1.1.

Acknowledgments

This work was supported by NSF CRII award CCF-2153371.

References

[1] R. Archibald and A. Gelb, *A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity*, IEEE Transactions on Medical Imaging, 21 (2002), pp. 305–319.

[2] F. Bach, *Breaking the curse of dimensionality with convex neural networks*, Journal of Machine Learning Research, 18 (2017), pp. 1–53.

[3] N. Benbarka, T. Höfer, A. Zell, et al., *Seeing implicit neural representations as Fourier series*, in Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2022, pp. 2041–2050.

[4] Y. Bengio, N. Roux, P. Vincent, O. Delalleau, and P. Marcotte, *Convex neural networks*, Advances in neural information processing systems, 18 (2005).

[5] B. N. Bhaskar, G. Tang, and B. Recht, *Atomic norm denoising with applications to line spectral estimation*, IEEE Transactions on Signal Processing, 61 (2013), pp. 5987–5999.

[6] K. Bredies and M. Carioni, *Sparsity of solutions for variational inverse problems with finite-dimensional data*, Calculus of Variations and Partial Differential Equations, 59 (2020), p. 14.

[7] Z. Cai, H. Zhu, Q. Shen, X. Wang, and X. Cao, *Batch normalization alleviates the spectral bias in coordinate networks*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 25160–25171.

[8] E. J. Candès and C. Fernandez-Granda, *Towards a mathematical theory of super-resolution*, Communications on pure and applied Mathematics, 67 (2014), pp. 906–956.

[9] Y. Chi and Y. Chen, *Compressive two-dimensional harmonic retrieval via atomic norm minimization*, IEEE Transactions on Signal Processing, 63 (2014), pp. 1030–1042.

[10] Y. Chi and M. F. Da Costa, *Harnessing sparsity over the continuum: Atomic norm minimization for superresolution*, IEEE Signal Processing Magazine, 37 (2020), pp. 39–57.

[11] J. B. Conway, *A course in functional analysis*, vol. 96, Springer, 2nd ed., 1990.

[12] A. Eftekhari, T. Bendory, and G. Tang, *Stable super-resolution of images: theoretical study*, Information and Inference: A Journal of the IMA, 10 (2021), pp. 161–193.

[13] I. Ekeland and R. Temam, *Convex analysis and variational problems*, SIAM, 1999.

[14] T. Ergen and M. Pilanci, *Convex geometry and duality of over-parameterized neural networks*, Journal of machine learning research, 22 (2021), pp. 1–63.

[15] A. Essakine, Y. Cheng, C.-W. Cheng, L. Zhang, Z. Deng, L. Zhu, C.-B. Schönlieb, and A. I. Aviles-Rivero, *Where do we stand with implicit neural representations? a technical and performance survey*, arXiv preprint arXiv:2411.03688, (2024).

[16] J. Feng, R. Feng, Q. Wu, Z. Zhang, Y. Zhang, and H. Wei, *Spatiotemporal implicit neural representation for unsupervised dynamic MRI reconstruction*, arXiv preprint arXiv:2301.00127, (2022).

[17] J. A. Fessler, *Model-based image reconstruction for MRI*, IEEE signal processing magazine, 27 (2010), pp. 81–89.

[18] F. G. Friedlander, *Introduction to the Theory of Distributions*, Cambridge University Press, 1998.

[19] M. Guerquin-Kern, L. Lejeune, K. P. Pruessmann, and M. Unser, *Realistic analytical phantoms for parallel magnetic resonance imaging*, IEEE Transactions on Medical Imaging, 31 (2011), pp. 626–636.

[20] S. Gunasekar, J. Lee, D. Soudry, and N. Srebro, *Characterizing implicit bias in terms of optimization geometry*, in International Conference on Machine Learning, PMLR, 2018, pp. 1832–1841.

[21] W. Huang, H. B. Li, J. Pan, G. Cruz, D. Rueckert, and K. Hammernik, *Neural implicit k-space for binning-free non-cartesian cardiac MR imaging*, in International Conference on Information Processing in Medical Imaging, Springer, 2023, pp. 548–560.

[22] A. Krogh and J. Hertz, *A simple weight decay can improve generalization*, Advances in neural information processing systems, 4 (1991).

[23] J. F. Kunz, S. Ruschke, and R. Heckel, *Implicit neural networks with fourier-feature inputs for free-breathing cardiac MRI reconstruction*, IEEE Transactions on Computational Imaging, (2024).

[24] B. Kurmanbek and E. Robeva, *Multivariate super-resolution without separation*, Information and Inference: A Journal of the IMA, 12 (2023), pp. 1938–1963.

[25] Z. Liu, H. Zhu, Q. Zhang, J. Fu, W. Deng, Z. Ma, Y. Guo, and X. Cao, *FINER: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 2713–2722.

[26] J. McGinnis, S. Shit, H. B. Li, V. Sideri-Lampretsa, R. Graf, M. Dannecker, J. Pan, N. Stolt-Ansó, M. Mühlau, J. S. Kirschke, et al., *Single-subject multi-contrast MRI super-resolution via implicit neural representations*, in Medical Image Computing and Computer-Assisted Intervention, Springer, 2023, pp. 173–183.

[27] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, *Nerf: Representing scenes as neural radiance fields for view synthesis*, Communications of the ACM, 65 (2021), pp. 99–106.

[28] T. Müller, A. Evans, C. Schied, and A. Keller, *Instant neural graphics primitives with a multiresolution hash encoding*, ACM transactions on graphics (TOG), 41 (2022), pp. 1–15.

[29] M. Najaf and G. Ongie, *Towards a sampling theory for implicit neural representations*, in Asilomar Conference on Signals, Systems, and Computers, 2024, pp. 105–110.

[30] J. Nocedal and S. J. Wright, *Numerical optimization*, Springer, 2 ed., 2006.

[31] T. Novello, *Understanding sinusoidal neural networks*, arXiv preprint arXiv:2212.01833, (2022).

[32] T. Novello, D. Aldana, A. Araujo, and L. Velho, *Tuning the frequencies: Robust training for sinusoidal neural networks*, 2025, `https://arxiv.org/abs/2407.21121`, `https://arxiv.org/abs/2407.21121`.

[33] G. Ongie, S. Biswas, and M. Jacob, *Convex recovery of continuous domain piecewise constant images from nonuniform Fourier samples*, IEEE Transactions on Signal Processing, 66 (2017), pp. 236–250.

[34] G. Ongie and M. Jacob, *Recovery of piecewise smooth images from few Fourier samples*, in 2015 International Conference on Sampling Theory and Applications (SampTA), 2015, pp. 543–547.

[35] G. Ongie and M. Jacob, *Off-the-grid recovery of piecewise constant images from few Fourier samples*, SIAM Journal on Imaging Sciences, 9 (2016), pp. 1004–1041.

[36] G. Ongie, R. Willett, D. Soudry, and N. Srebro, *A function space view of bounded norm infinite width ReLU nets: The multivariate case*, in International Conference on Learning Representations, 2020.

[37] R. Parhi and R. D. Nowak, *Banach space representer theorems for neural networks and ridge splines*, Journal of Machine Learning Research, 22 (2021), pp. 1–40.

[38] R. Parhi and R. D. Nowak, *Deep learning meets sparse regularization: A signal processing perspective*, IEEE Signal Processing Magazine, 40 (2023), pp. 63–74.

[39] C. Poon, N. Keriven, and G. Peyré, *The geometry of off-the-grid compressed sensing*, Foundations of Computational Mathematics, 23 (2023), pp. 241–327.

[40] C. Poon and G. Peyré, *Multidimensional sparse super-resolution*, SIAM Journal on Mathematical Analysis, 51 (2019), pp. 1–44.

[41] S. Ramasinghe and S. Lucey, *Beyond periodicity: Towards a unifying framework for activations in coordinate-MLPs*, in European Conference on Computer Vision, Springer, 2022, pp. 142–158.

[42] V. Saragadam, D. LeJeune, J. Tan, G. Balakrishnan, A. Veeraraghavan, and R. G. Baraniuk, *WIRE: Wavelet implicit neural representations*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 18507–18516.

[43] H. Saratchandran, S. Ramasinghe, V. Shevchenko, A. Long, and S. Lucey, *A sampling theory perspective on activations for implicit neural representations*, in Proceedings of the 41st International Conference on Machine Learning, 2024, pp. 43422–43444.

[44] P. Savarese, I. Evron, D. Soudry, and N. Srebro, *How do infinite width bounded norm networks look in function space?*, in Conference on Learning Theory, PMLR, 2019, pp. 2667–2690.

[45] L. Shen, J. Pauly, and L. Xing, *NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction*, IEEE Transactions on Neural Networks and Learning Systems, 35 (2022), pp. 770–782.

[46] J. Shenouda, Y. Zhou, and R. D. Nowak, *ReLUs are sufficient for learning implicit neural representations*, International Conference on Machine Learning (ICML), (2024).

[47] L. A. Shepp and B. F. Logan, *The Fourier reconstruction of a head section*, IEEE Transactions on nuclear science, 21 (1974), pp. 21–43.

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

[49] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, *Fourier features let networks learn high frequency functions in low dimensional domains*, Advances in Neural Information Processing Systems, 33 (2020), pp. 7537–7547.

[50] D. Van Veen, R. Van der Sluijs, B. Ozturkler, A. D. Desai, C. Bluethgen, R. D. Boutin, M. H. Willis, G. Wetzstein, D. B. Lindell, S. Vasanawala, et al., *Scale-agnostic super-resolution in MRI using feature-based coordinate networks*, in Medical Imaging with Deep Learning, 2022.

[51] J. Veraart, E. Fieremans, I. O. Jelescu, F. Knoll, and D. S. Novikov, *Gibbs ringing in diffusion MRI*, Magnetic resonance in medicine, 76 (2016), pp. 301–314.

[52] Q. Wu, Y. Li, Y. Sun, Y. Zhou, H. Wei, J. Yu, and Y. Zhang, *An arbitrary scale super-resolution approach for 3D MR images via implicit neural representation*, IEEE Journal of Biomedical and Health Informatics, 27 (2022), pp. 1004–1015.

[53] Q. Wu, Y. Li, L. Xu, R. Feng, H. Wei, Q. Yang, B. Yu, X. Liu, J. Yu, and Y. Zhang, *IREM: High-resolution magnetic resonance image reconstruction via implicit neural representation*, in Medical Image Computing and Computer Assisted Intervention, Springer, 2021, pp. 65–74.

[54] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, *Neural fields in visual computing and beyond*, Computer Graphics Forum, 41 (2022), pp. 641–676.

[55] G. Yüce, G. Ortiz-Jiménez, B. Besbinar, and P. Frossard, *A structured dictionary perspective on implicit neural representations*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 19228–19238.

[56] Q. Zou and M. Jacob, *Recovery of surfaces and functions in high dimensions: sampling theory and links to neural networks*, SIAM Journal on Imaging Sciences, 14 (2021), pp. 580–619.

Supplementary Material

D Additional Experimental Results and Figures

In this section, we provide additional reconstruction figures.

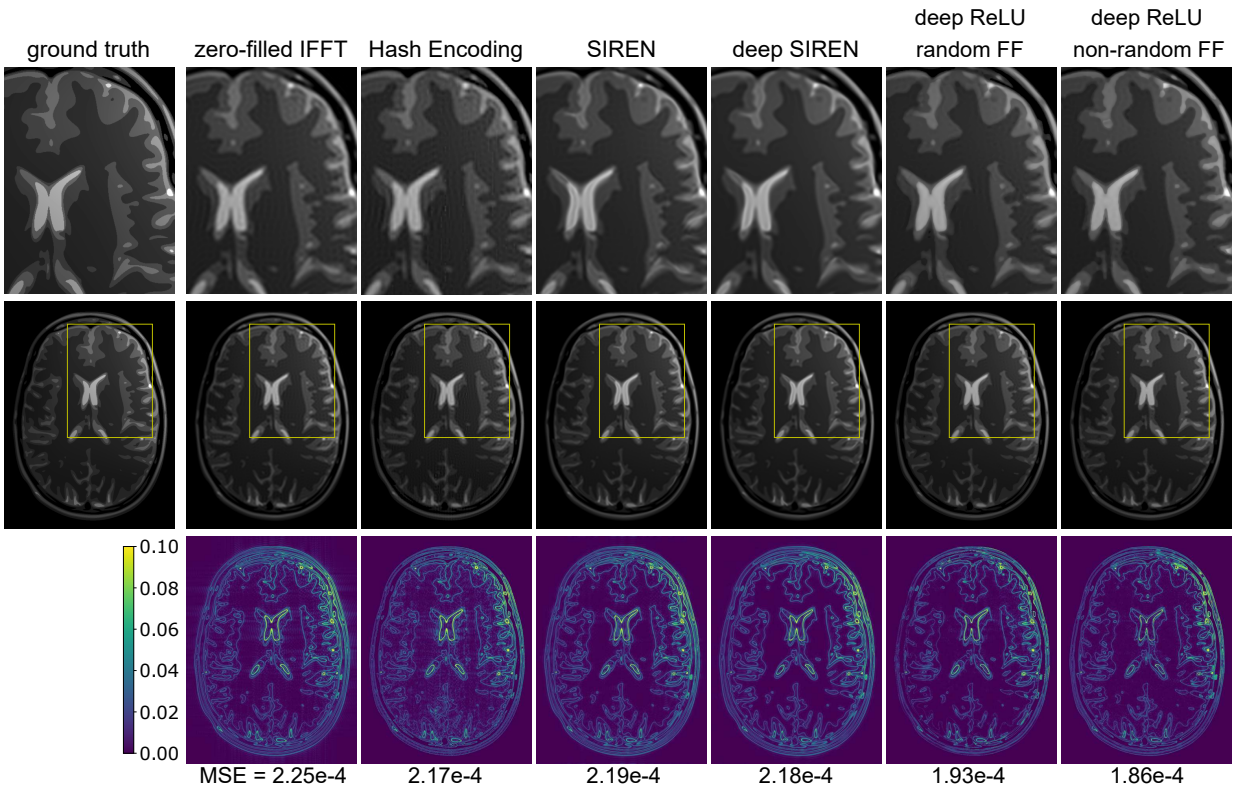

Figure 11: **PWS-BRAIN recovery using different INR architectures.** Top row: we compare the zero-filled IFFT reconstruction with the phantom recovery from Hash Encoding, depth 5 SIREN (SIREN), depth 14 SIREN (deep SIREN), and depth 15 INR (deep INR) with *random* Fourier features and *non-random* Fourier features using ReLU. The bottom row shows the absolute value of the difference with the ground truth. All grayscale images are shown on the scale [0, 1].

E Experimental Details

In this section, we provide additional details for some experiments included in the main text.

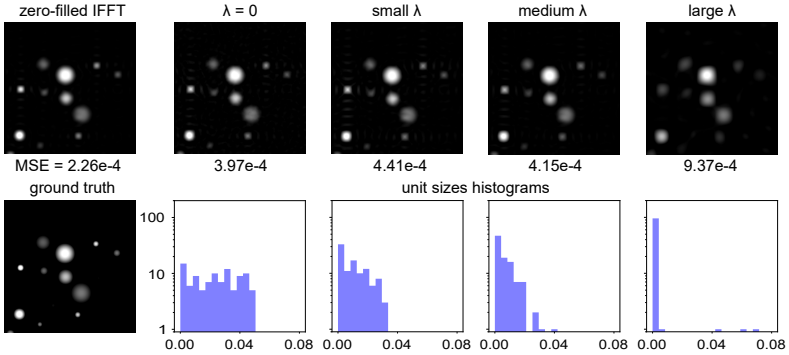

(a) standard weight decay regularization

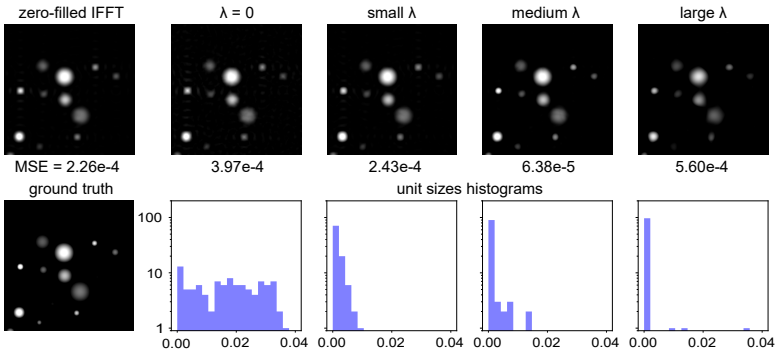

(b) modified weight decay regularization proposed in Theorem 1

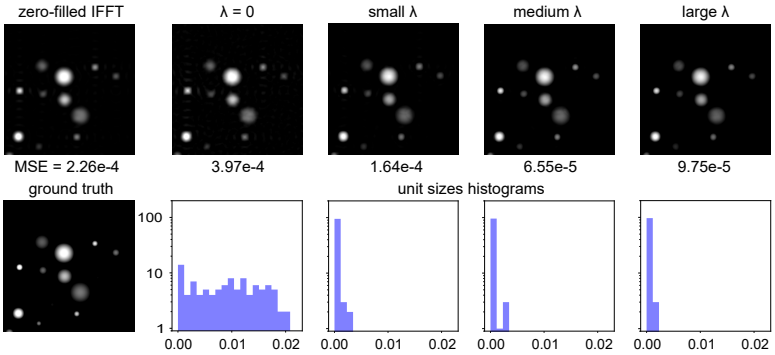

(b) modified weight decay regularization proposed in Theorem 2

Figure 12: **Impact of regularization on the DOT phantom reconstruction and unit sizes of a trained shallow INR.** In the first row of each panel, we compare the zero-filled IFFT reconstruction with the reconstructions obtained with a shallow INR trained with standard/modified weight decay regularization for different values of regularization strength λ. The histograms in the second row of each panel represent the distribution of unit sizes $\{|a_i|\eta(\boldsymbol{w}_i)\}_{i=1}^{100}$ of the trained INRs across different values of λ. As λ increases, the unit sizes cluster at zero, indicating the trained INR is sparse in the sense that it has few active units. (Note: histograms are plotted on a log scale.)

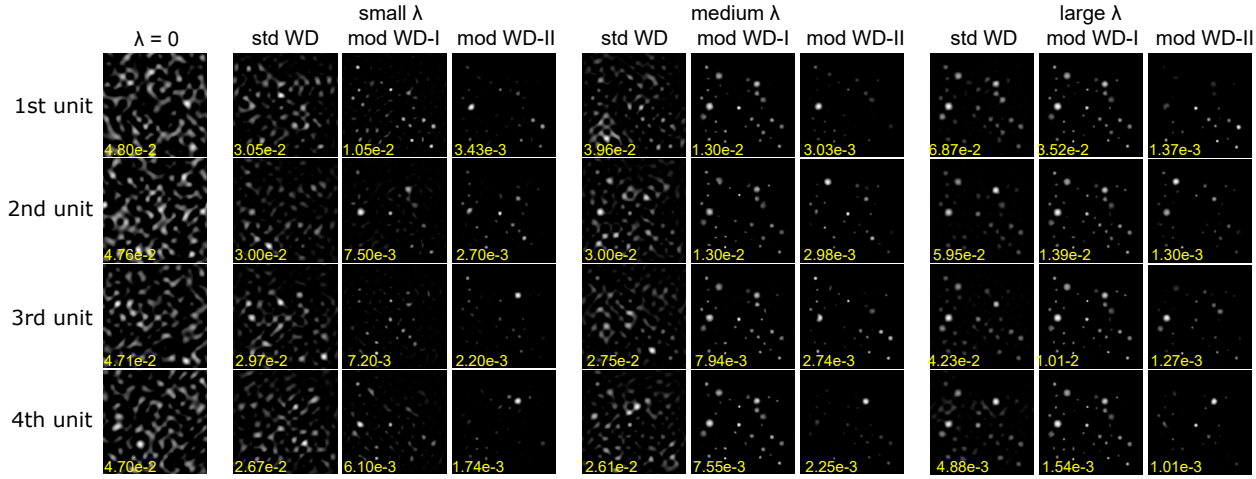

Figure 13: **Visualization of the 4 largest active units for the DOT phantom, obtained with a shallow INR.** The INR is trained with standard weight decay regularization (std WD), modified weight decay regularization proposed in Theorem 1 (mod WD-I), and the modified weight decay regularization proposed in Theorem 2 (mod WD-II). All images are normalized to the [0,1] range. The corresponding "unit size" $|a_i|\eta(\boldsymbol{w}_i)$ is displayed on the image in yellow.

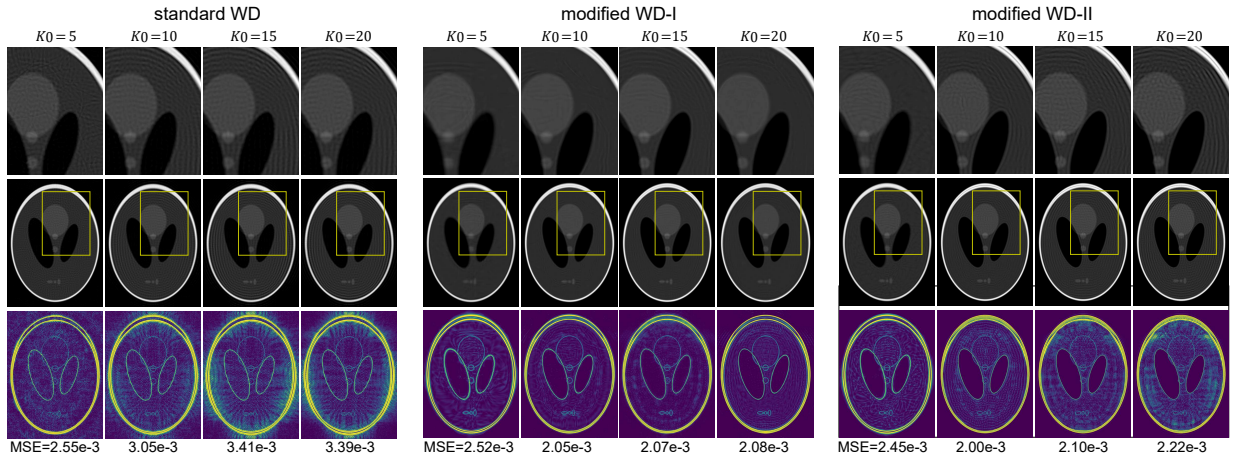

Figure 14: The effect of varying the maximum frequency for Fourier features K_0, on the SL phantom recovery trained using a shallow INR with standard WD, modified WD-I, and modified WD-II regularization. Observe that the lowest MSE is achieved for $K_0 = 10$ for modified WD-I and WD-II.

E.1 Additional Details of Experiments in Section 4.1.1

In the setting of Theorem 1, we use AL parameters $\sigma_0 = 10$, with $\gamma = 1.5$ and perform $N = 60$ AL outer loop iterations. In the setting of Theorem 2, we use AL parameters $\sigma_0 = 10$, with $\gamma = 1.1$, and perform $N = 100$ AL outer loop iterations. In both cases, we run $5,000$ iterations of the Adam optimizer with a learning rate of 1×10^{-3} to approximately minimize the AL inner loop subproblem (32).

E.2 Additional Details of Experiments in Section 4.2.1

This section presents the details of the experiments for reconstructing the SL and PWC-BRAIN phantoms using INR with different depths, with and without weight decay regularization, as shown in Figure 8.

For the SL phantom recovery, we set the maximum sampling frequency to $K = 48$, and set the maximum sampling frequency of the Fourier features map to $K_0 = 10$. For the reconstructions with weight decay regularization, the regularization strength parameter λ is set as follows: 1×10^{-6} for depth 2, 5×10^{-8} for depth 5, 5×10^{-7} for depth 10, 5×10^{-8} for depth 15 and 2×10^{-14} for depth 20.

For the PWC BRAIN phantom recovery, we set $K = 64$ and $K_0 = 10$. Additionally, λ is set as follows: 1×10^{-4} for depth 2, 2×10^{-8} for depths 5 and 10, 1×10^{-7} for depth 15 and 1×10^{-8} for depth 20. In all experiments, both phantoms are recovered by training an INR of width 100 using the Adam optimizer for 40,000 iterations with a learning rate of 1×10^{-3} followed by 10,000 iterations with a learning of 1×10^{-4}.

E.3 Additional Details of Experiments in Section 4.2.2

This section outlines the experimental details for reconstructing the SL, PWC-BRAIN, and PWS-BRAIN phantoms using different INR architectures, as shown in Figures 9, 10 and 11.

Hash Encoding For the hash encoding INR architecture, we use a fully-connected ReLU network with a depth of 3 and a width of 64 and the following parameters for the hash encoding layer: the number of levels $L = 16$, number of feature dimensions per entry $F = 2$, coarsest resolution $N_{min} = 16$, finest resolution $N_{max} = 1024$. Additionally, details on training the INR using hash encoding for each phantom reconstruction are provided below.

SL: We set the hash table size $T = 2^{21}$, and train the INR using the Adam optimizer with a learning rate of 1×10^{-2} for 100,000 iterations, and set the weight decay regularization strength parameter to $\lambda = 1 \times 10^{-4}$.

PWC-BRAIN: We set the hash table size $T = 2^{19}$, and train the INR using the Adam optimizer with a learning rate of 1×10^{-2} for 100,000 iterations, and set $\lambda = 1 \times 10^{-6}$.

PWS-BRAIN: We set the hash table size $T = 2^{19}$, and train the INR using the Adam optimizer with a learning rate of 1×10^{-2} for 100,000 iterations, and set $\lambda = 1 \times 10^{-4}$.

SIREN For SIREN, we set the network's depth to 5 and for the deep SIREN, we set it to 14. For both SIREN and deep SIREN the hidden-layer widths is set to 256 and we set the bandlimit parameter $\omega_0 = 30$. Our experiments showed that increasing the depth from 5 to 14 can lead to lower image MSE. Additionally, in experiments with depth 5 SIREN networks, we did not apply weight decay regularization since it did not have a significant effect on the output. However, our experimental results indicate that with deep SIREN, using weight decay regularization improves the reconstructions and leads to lower image MSE values. Details on training the INR using SIREN for each phantom reconstruction are provided below.

SL: For SIREN, we train the INR using the Adam optimizer with a learning rate of 1×10^{-5} for 50,000 iterations without weight decay regularization. For deep SIREN, we train the INR using the Adam optimizer with a learning rate of 1×10^{-5} for 50,000 iterations and set $\lambda = 5 \times 10^{-3}$.

PWC-BRAIN: For SIREN, we train the network using the Adam optimizer with a learning rate of 1×10^{-6} for 50,000 iterations without weight decay regularization. For deep SIREN, we train the INR using the Adam optimizer for 10,000 iterations with a learning rate of 1×10^{-4} followed by 40,000 iterations with a learning rate of 1×10^{-5}, and set $\lambda = 5 \times 10^{-2}$.

PWS-BRAIN: For SIREN, we train the network using the Adam optimizer with a learning rate of 1×10^{-6} for 50,000 iterations without weight decay regularization. For deep SIREN, we train the INR using the Adam optimizer for 50,000 iterations with a learning rate of 1×10^{-6}, and set λ to 5×10^{-3}.

Deep INR with random Fourier features The INR width is set to 100, and the initial Fourier features layer is defined by sampling 256 random frequency vectors whose entries are i.i.d. Gaussian with mean-zero and standard deviation 10, resulting in a total of $D = 513$ features. Additionally, the network is trained using the Adam optimizer for 40,000 iterations with a learning rate of 1×10^{-3}, followed by 10,000 iterations with a learning rate of 1×10^{-4}. For each phantom, the INR's depth and weight decay hyperparameter are as below:

SL: The network's depth is set to 20 and $\lambda = 1 \times 10^{-11}$.

PWC-BRAIN: The network's depth is set to 20 and $\lambda = 4 \times 10^{-8}$.

PWS-BRAIN: The network's depth is set to 15 and $\lambda = 5 \times 10^{-7}$.

Deep INR with non-random Fourier features In these experiments, an INR of width 100 is trained using the Adam optimizer for 40,000 iterations with a learning rate of 1×10^{-3}, followed by 10,000 iterations with a learning rate of 1×10^{-4}. For each phantom, we define K, K_0, the INR's depth and weight decay hyperparameter λ as below:

SL: The INR's depth is set to 20, $K = 48$, $K_0 = 10$ and $\lambda = 2 \times 10^{-14}$.

PWC-BRAIN: The INR's depth is set to 20, $K = 64$, $K_0 = 10$, and $\lambda = 1 \times 10^{-8}$.

PWS-BRAIN: The INR's depth is set to 20, $K = 64$, $K_0 = 10$, and $\lambda = 5 \times 10^{-8}$.

E.4 Experimental Details of the Reconstructions Shown in Figures 6, 12, 7, and 13

This section describes the experimental setup for the effect of regularization on the PWC BRAIN and DOT phantoms reconstruction and unit sizes of a trained shallow INR as shown in Figures 6 and 12. For the PWC-BRAIN phantom, we set $K = 64$, $K_0 = 20$, width 500. For the DOT phantom, we set $K = 32$, $K_0 = 10$, INR width to 100. For both phantoms, we train the INR networks using the Adam optimizer for 10,000 iterations with a learning rate of 1×10^{-3}. The weight decay regularization hyperparameter λ is selected by a grid search to optimize the image domain MSE in each experiment. Additionally, we use the same experimental settings and parameters to visualize the four largest active units shown in Figures 7 and 13.

E.5 Experimental Details of the Reconstructions Shown in Figure 14

This section outlines the experimental details for reconstructing the SL phantom using a depth 2 INR with different K_0 as shown in Figure 14. In these experiments, we set $K = 48$ and train a shallow INR of width 100 using the Adam optimizer for 40,000 iterations with a learning rate of 1×10^{-3}, followed by 10,000 iterations with a learning rate of 1×10^{-4}. We trained the INR using standard WD, modified WD-I, and modified WD-II regularization and tested varying values of K_0 for each. In each experiment, we selected the regularization parameter λ using a grid search to optimize the image MSE. For the main results included in the paper, we set $K_0 = 10$ for the SL phantom recovery. When $K_0 = 10$, we have $n = 220$ frequency vectors, plus one constant feature, which gives a total of $D = 2n + 1 = 441$ features.

F Proof of Lemma 8

For convenience, we restate Lemma 8 before giving its proof. Recall that $\mathsf{TP}(\Gamma_1)$ is the space of 2D trigonometric polynomials $\tau : \mathbb{T}^2 \to \mathbb{R}$ of the form

$$\tau(x_1, x_2) = a_1 \cos(2\pi x_1) + b_1 \sin(2\pi x_1) + a_2 \cos(2\pi x_2) + b_2 \sin(2\pi x_2)$$

for some $a_1, a_2, b_1, b_2 \in \mathbb{R}$.

Lemma (Restatement of Lemma 8)**.** *Suppose $\tau \in \mathsf{TP}(\Gamma_1)$ is non-zero. Then $C = Z(\tau)$ is a trigonometric curve, and τ is a minimal polynomial for C. Furthermore, τ is irreducible except in the following cases:*

1. *$\tau(x_1, x_2) = \alpha \sin(2\pi(x_1 - t))$ for some $\alpha \in \mathbb{R}$, $t \in \mathbb{T}$.*
2. *$\tau(x_1, x_2) = \alpha \sin(2\pi(x_2 - t))$ for some $\alpha \in \mathbb{R}$, $t \in \mathbb{T}$.*

3. $\tau(x_1, x_2) = \alpha(\sin(2\pi(x_1 - t_1)) + \sin(2\pi(x_2 - t_2)))$ *for some* $\alpha \in \mathbb{R}$, $t_1, t_2 \in \mathbb{T}$.

Proof. Let $\tau \in \mathsf{TP}(\Gamma_1)$ be non-zero, and $C = Z(\tau)$. First, we note that C is nonempty: we have $\hat{\tau}[0,0] = \int\int_{\mathbb{T}^2} \tau(x_1, x_2) dx_1 dx_2 = 0$, which implies τ must change sign over $\mathbb{T}^2$, and so by the intermediate value theorem τ must have at least one zero.

Next, we prove that C has no isolated points (i.e., τ has no isolated zeros). Suppose, by way of contradiction, that τ has an isolated zero at a point (x_1^*, x_2^*). Then, by continuity of τ, it must have a strict local maximum or minimum at (x_1^*, x_2^*). In particular, the Hessian matrix $\nabla^2\tau(x_1^*, x_2^*)$ must be positive semi-definite or negative semi-definite. And since

$$\tau(x_1, x_2) = a_1 \cos(2\pi x_1) + b_1 \sin(2\pi x_1) + a_2 \cos(2\pi x_2) + b_2 \sin(2\pi x_2)$$

for some $a_1, a_2, b_1, b_2 \in \mathbb{R}$, we have

$$\nabla^2\tau(x_1^*, x_2^*) = -4\pi^2 \begin{bmatrix} a_1 \cos(2\pi x_1^*) + b_1 \sin(2\pi x_1^*) & 0 \\ 0 & a_2 \cos(2\pi x_2^*) + b_2 \sin(2\pi x_2^*) \end{bmatrix}.$$

Also, $\tau(x_1^*, x_2^*) = 0$ implies $a_2 \cos(2\pi x_2^*) + b_2 \sin(2\pi x_2^*) = -(a_1 \cos(2\pi x_1^*) + b_1 \sin(2\pi x_1^*))$. Therefore, if $a_1 \cos(2\pi x_1^*) + b_1 \sin(2\pi x_1^*)$ were non-zero, this would imply the Hessian is indefinite, so the only possibility is

$$a_1 \cos(2\pi x_1^*) + b_1 \sin(2\pi x_1^*) = a_2 \cos(2\pi x_2^*) + b_2 \sin(2\pi x_2^*) = 0.$$

This implies

$$a_1 = -\alpha \sin(2\pi x_1^*), b_1 = \alpha \cos(2\pi x_1^*), a_2 = -\beta \sin(2\pi x_2^*), b_2 = \beta \cos(2\pi x_2^*)$$

for some scalars $\alpha, \beta \in \mathbb{R}$. Simple trigonometric identities show that this implies

$$\tau(x_1, x_2) = \alpha \sin(2\pi(x_1 - x_1^*)) + \beta \sin(2\pi(x_2 - x_2^*)).$$

Finally, since (x_1^*, x_2^*) is a critical point of τ, we have

$$\nabla\tau(x_1^*, x_2^*) = 2\pi(\alpha, \beta) = (0, 0)$$

and so $\alpha = \beta = 0$, i.e., τ is identically zero, a contradiction. Therefore, τ cannot have any isolated zeros, and so $C = Z(\tau)$ is a trigonometric curve.

Now we prove that τ is a minimal polynomial for C. To this end, we show that the conversion of τ into a complex polynomial in two variables is either irreducible or factors into a product of distinct irreducible factors. First, expanding τ as a sum of complex exponentials, we have

$$\tau(x_1, x_2) = c_1 e^{i2\pi x_1} + \overline{c_1} e^{-i2\pi x_1} + c_2 e^{i2\pi x_2} + \overline{c_2} e^{-i2\pi x_2}$$

where $c_1, c_2 \in \mathbb{C}$ are such that either $c_1 \neq 0$ or $c_2 \neq 0$. Let $p = \mathcal{P}[\tau]$ be the conversion of τ to a complex polynomial in two variables. Then p has the form

$$p(z_1, z_2) = z_1^{q_1} z_2^{q_2} \left(c_1 z_1 + \overline{c_1} z_1^{-1} + c_2 z_2 + \overline{c_2} z_2^{-1}\right)$$

where $q_1, q_2 \in \{0, 1\}$. We consider the following cases:

Case 1: $c_1 = 0$ or $c_2 = 0$. Then p has the form

$$p(z_1, z_2) = z_i(c_i z_i^{-1} + \overline{c_i} z_i) = c_i + \overline{c_i} z_i^2$$

for some $i = 1, 2$, which factors as

$$p(z_1, z_2) = \overline{c_i}(z_i - r)(z_i + r)$$

where $r = \sqrt{-c_i/\overline{c_i}}$ satisfies $|r| = 1$. Note that the linear factors $z_i + r$ and $z_i - r$ are irreducible. Also, their zero-sets $Z(z_i - r) = \{(z_1, z_2) : z_i = r\}$ and $Z(z_i + r) = \{(z_1, z_2) : z_i = -r\}$ are distinct, and have infinite intersection with the complex unit torus. Therefore, τ is a minimal polynomial for $Z(\tau)$. Converting p back into a trigonometric polynomial gives $\tau(x_1, x_2) = \alpha \sin(2\pi(x_i - t))$ for some $i \in \{1, 2\}$, $\alpha \in \mathbb{R}$ with $\alpha \neq 0$, and $t \in \mathbb{T}$, which corresponds to the first two cases in the statement of the lemma.

Case 2: $c_1 \neq 0$ and $c_2 \neq 0$. Then p has the form

$$p(z_1, z_2) = z_1 z_2 \left(c_1 z_1 + \overline{c_1} z_1^{-1} + c_2 z_2 + \overline{c_2} z_2^{-1}\right) = c_1 z_1^2 z_2 + \overline{c_1} z_2 + c_2 z_2^2 z_1 + \overline{c_2} z_1. \tag{33}$$

If p is irreducible, there is nothing to show since then τ is a minimal polynomial by definition. Instead, assume p is reducible, i.e., p factors as $p = p_1 p_2$ for some non-constant polynomials $p_1, p_2 \in \mathbb{C}[z_1, z_2]$. Since p has total degree 3, the only possibility is that one of the factors has total degree 1 and the other factor has total degree 2, i.e., p has a factorization of the form

$$\begin{aligned} p(z_1, z_2) &= (a_0 + a_1 z_1 + a_2 z_2)(b_0 + b_1 z_1 + b_2 z_2 + b_3 z_1 z_2 + b_4 z_1^2 + b_5 z_2^2) \\ &= a_1 b_4 z_1^3 + a_2 b_5 z_2^3 + (a_1 b_3 + a_2 b_4) z_1^2 z_2 + (a_1 b_5 + a_2 b_3) z_2^2 z_1 \\ &\quad + (a_0 b_4 + a_1 b_1) z_1^2 + (a_0 b_5 + a_2 b_2)\, z_2^2 + (a_1 b_2 + a_2 b_1 + a_0 b_3) z_1 z_2 \\ &\quad + (a_2 b_0 + a_0 b_2) z_2 + (a_0 b_1 + a_1 b_0) z_1 + a_0 b_0 \end{aligned} \tag{34}$$

Comparing coefficient between (33) and (34), we see that:

$$\begin{gathered} a_0 b_0 = a_1 b_4 = a_2 b_5 = a_0 b_4 + a_1 b_1 = a_0 b_5 + a_2 b_2 = a_1 b_2 + a_2 b_1 + a_0 b_3 = 0, \\ c_1 = a_1 b_2 + a_2 b_4 = \overline{a_2 b_0 + a_0 b_2} \neq 0, \\ c_2 = a_1 b_5 + a_2 b_3 = \overline{a_0 b_1 + a_1 b_0} \neq 0. \end{gathered}$$

From the equation $a_0 b_0 = 0$, either $a_0 = 0$ or $b_0 = 0$ (or both). First, suppose $b_0 = 0$. From the equations above, this implies $a_0 b_2 \neq 0$ and $a_0 b_1 \neq 0$, and so a_0, b_1, b_2 are all nonzero. If $a_1 \neq 0$ then the equation $a_1 b_4 = 0$ implies $b_4 = 0$. Together with the equation $a_0 b_4 + a_1 b_1 = 0$ this implies $b_1 = 0$, a contradiction. Therefore, $a_1 = 0$. By the same reasoning, if $a_2 \neq 0$ then this would imply $b_2 = 0$, a contradiction, and so $a_2 = 0$. But if $a_1 = a_2 = 0$ then the first factor $a_0 + a_1 z_1 + a_2 z_2$ is a constant polynomial, and so p is irreducible, contradicting our assumption that p was reducible.

Therefore, if p is reducible, then $b_0 \neq 0$ and $a_0 = 0$. From the equations above this implies $a_2 b_0 \neq 0$ and $a_1 b_0 \neq 0$, i.e., b_0, a_1, a_2 are all non-zero. This in turn implies $b_1 = b_2 = b_4 = b_5 = 0$. Therefore, the only possible factorization of p is

$$p(z_1, z_2) = (a_1 z_1 + a_2 z_2)(b_0 + b_3 z_1 z_2)$$

where $a_1, a_2, b_0, b_3 \in \mathbb{C}$ are all non-zero, and satisfy $a_1 b_3 = \overline{a_2 b_0}$ and $a_2 b_3 = \overline{a_1 b_0}$. In particular, from these relations imply $|b_0| = |b_3|$ and $|a_1| = |a_2|$. Note that $p_1(z_1, z_2) = a_1 z_1 + a_2 z_2$ is irreducible since it is degree 1. Also, the zero-set of p_1 has an infinite intersection with the complex unit torus: substituting $z_i = e^{i2\pi x_i}$, $i = 1, 2$, the condition $a_1 z_1 + a_2 z_2 = 0$ is equivalent to $e^{j2\pi(x_1 - x_2)} = a$ where $a = -a_2/a_1$ is such that $|a| = 1$. This has infinitely many solutions $\{(x_1, x_2) \in \mathbb{T}^2 : x_1 - x_2 = t\}$ where $t \in \mathbb{T}$ is such that $e^{i2\pi t} = a$.

Also, the degree 2 polynomial $p_2(z_1, z_2) = b_0 + b_3 z_1 z_2$ is irreducible since it is easy to show it cannot be written as a product of two degree 1 polynomials. Furthermore, the zero-set of p_2 has infinite intersection with the complex unit torus: substituting $z_i = e^{i2\pi x_i}$, $i = 1, 2$, the condition $b_0 + b_3 z_1 z_2 = 0$ is equivalent to $e^{i2\pi(x_1 + x_2)} = b$ where $b = -b_3/b_0$ is such that $|b| = 1$. This has infinitely many solutions $\{(x_1, x_2) \in \mathbb{T}^2 : x_1 + x_2 = t\}$ where $t \in \mathbb{T}$ is such that $e^{i2\pi t} = b$.

Therefore, we have shown that if p is reducible, it is a product of distinct irreducible factors, each of which vanish on an infinite subset of the complex unit torus, which proves that τ is a minimal polynomial. Finally, converting p back into a trigonometric polynomial, it is straightforward to show that in this case

$$\tau(x_1, x_2) = \alpha(\sin(2\pi(x_1 - t_1)) + \sin(2\pi(x_2 - t_2)))$$

for some $\alpha \in \mathbb{R}$ with $\alpha \neq 0$ and some $t_1, t_2 \in \mathbb{T}$, which corresponds to the third case in the statement of the lemma. □