

Inductive Bias of Multi-Channel Linear Convolutional Networks with Bounded Weight Norm

Meena Jagadeesan

*UC Berkeley**

MJAGADEESAN@BERKELEY.EDU

Ilya Razenshteyn

CipherMode Labs

ILYA.RAZENSHTEYN@GMAIL.COM

Suriya Gunasekar

Microsoft Research

SURIYAG@MICROSOFT.COM

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

We provide a function space characterization of the inductive bias resulting from minimizing the ℓ_2 norm of the weights in multi-channel convolutional neural networks with linear activations and empirically test our resulting hypothesis on ReLU networks trained using gradient descent. We define an *induced regularizer* in the function space as the minimum ℓ_2 norm of weights of a network required to realize a function. For two layer linear convolutional networks with C output channels and kernel size K , we show the following: (a) If the inputs to the network are single channeled, the induced regularizer for any K is *independent* of the number of output channels C . Furthermore, we derive the regularizer is a norm given by a semidefinite program (SDP). (b) In contrast, for multi-channel inputs, multiple output channels can be necessary to merely realize all matrix-valued linear functions and thus the inductive bias *does* depend on C . However, for sufficiently large C , the induced regularizer is again given by an SDP that is independent of C . In particular, the induced regularizer for $K = 1$ and $K = D$ (input dimension) are given in closed form as the nuclear norm and the $\ell_{2,1}$ group-sparse norm, respectively, of the Fourier coefficients of the linear predictor. We investigate the broader applicability of our theoretical results to implicit regularization from gradient descent on linear and ReLU networks through experiments on MNIST and CIFAR-10 datasets.

1. Introduction

In the study of generalization and model capacity, complexity measures based on magnitude of parameters have long been argued to play an important role in learning overparametrized models (Bartlett, 1996; Bartlett and Mendelson, 2002; Neyshabur et al., 2015; Zhang et al., 2017; Bartlett et al., 2017). In particular, the ℓ_2 norm of weights (or parameters) is a prominent complexity measure of interest in the current practice of deep learning, with connections to explicit regularization (Krogh and Hertz, 1991; Wei et al., 2019) as well as implicit regularization from optimization algorithms (Ji and Telgarsky, 2019; Gunasekar et al., 2018a) (see additional discussion in Section 1.2). Importantly, recent results (Lyu and Li, 2020; Nacson et al., 2019; Ji and Telgarsky, 2020) show that in many (not all) instances of overparametrized classification problems, gradient descent asymptotically leads to solutions that implicitly control the ℓ_2 norm of the parameters (see Section 5 for a formal statement).

We study the question: *What is the nature of functions learned by controlling ℓ_2 norm of parameters?* Consider a model class of functions (or network architecture) $\Phi(\theta; \cdot)$ with parameters

* This work was partly done while M. Jagadeesan and I. Razenshteyn were at Microsoft Research. M. Jagadeesan was supported in part by the Paul and Daisy Soros Fellowship.

(or weights) θ . The function space view of controlling ℓ_2 norm of parameters (denoted as $\|\theta\|$) can be understood in terms of its *representation cost*, i.e., the minimum ℓ_2 norm of weights needed to realize a function using a given network architecture $\Phi(\theta, \cdot)$. This defines an *induced complexity measure* over functions, which we also refer as induced regularizer, given by

$$\mathcal{R}_\Phi(f) := \inf_{\theta} \|\theta\|^2 \text{ s.t., } \forall \mathbf{x}, f(\mathbf{x}) = \Phi(\theta, \mathbf{x}). \quad (1)$$

Note that a learning objective with ℓ_2 norm regularization of the parameters $\min_{\theta} \mathcal{L}(\Phi(\theta, \cdot)) + \lambda \|\theta\|^2$ is equivalent to the corresponding \mathcal{R} -regularization over functions $\min_f \mathcal{L}(f) + \lambda \mathcal{R}_\Phi(f)$.¹

Even for neural networks that realize the same function class, minimizing or bounding the ℓ_2 norm of weights in different architectures can lead to remarkably different effects in function space. For example, consider networks with fully connected and convolution layers and linear activations. These architectures are simply different parameterizations of the same model class of linear functions. Gunasekar et al. (2018b) showed that for fully connected linear networks, the induced regularizer is the ℓ_2 norm of the linear map realized by the network, while for linear convolutional network with *full dimensional kernels*, it is the ℓ_1 norm of Fourier coefficients of the linear map. This function space view reveals that minimizing the ℓ_2 norm of weights in these networks has fundamentally different implications for learned predictors depending on the parametrization of function class.

1.1. Our contribution

In this work, we investigate the induced regularizer in (1) for multi-channel linear convolutional networks. In particular, we study two layer networks that have C output channels, R input channels, and kernel size K . We characterize the role of the number of channels C on the induced regularizer, for networks with arbitrary kernel size K . Our main contribution is that for inputs with a single channel, having multiple output channels in the network surprisingly does not reduce the ℓ_2 -norm representational cost, despite increasing the number of parameters in the network.

Theorem 1 (Informal) *For two layer convolutional networks with any kernel size K , if the inputs are single-channeled, then the induced regularizer is independent of the output channel size C .*

Proof technique. To prove Theorem 1, we construct an *semidefinite program (SDP) relaxation* which corresponds to the induced regularizer when $C = \infty$. We then prove that this SDP relaxation is in fact tight for all $C \geq 1$, which leads to our main result. The SDP further implies a convex structure of the induced regularizer for any C . In our proof of SDP tightness, we use a polynomial representation of convolutions to implicitly argue the existence of a rank-1 optimal solution. A key lemma in our proof (Lemma 7) shows an interesting property about convolutions in \mathbb{R}^D with kernel of size $K < D$. To our knowledge, this property as well as the proof technique involving polynomial representations are new and are of independent interest.

Extension to multi-channel inputs. We further extend our findings to networks with multi-channel inputs. For multi-channel inputs of dimensions $D \times R$, even realizing all linear functions over the inputs can require multiple output channels C (see Lemma 11). Hence, the induced regularizer *does* depend on C , although for large enough C , we show a restricted form of invariance. In particular, we prove the induced regularizer is invariant to the number of output channels when $C \geq R \cdot K$, and

1. \mathcal{R}_Φ can be equivalently defined as any monotonic function of $\|\theta\|$. We use $\|\theta\|^2$ to align with the standard regularizer (see also, (Gunasekar et al., 2018b; Savarese et al., 2019; Ongie et al., 2020; Dai et al., 2021)).

conjecture invariance when $C \geq R$. We then characterize the induced regularizer in the special cases of $K = 1$ and $K = D$ as the nuclear norm and the $\ell_{2,1}$ group sparse norm of the Fourier coefficients, respectively (see Theorems 14-15).

Experiments for gradient descent. Finally, we connect our results to the implicit regularization of gradient descent. When combined with prior work (e.g., [Lyu and Li \(2020\)](#)), our results also extend to asymptotic predictors learned by gradient descent on networks with ReLU or linear activations. We thus formulate and study an empirically testable hypothesis that the ℓ_2 norm complexity of networks learned using gradient descent is invariant to C as long as $C \geq R$. We validate this hypothesis on MNIST and CIFAR-10 datasets on linear convolutional neural networks with circular and zero padding. The behavior also holds on MNIST in networks with ReLU non-linearity.

1.2. Related Work

There is a rich literature of work connecting ℓ_2 norm minimization of weights with explicit regularization ([Krogh and Hertz, 1991](#); [Wei et al., 2019](#)) and implicit regularization from gradient descent ([Neyshabur et al., 2015](#); [Zhang et al., 2017](#); [Bartlett et al., 2017](#); [Gunasekar et al., 2018a,b](#); [Ji and Telgarsky, 2018, 2019](#); [Nacson et al., 2019](#); [Lyu and Li, 2020](#); [Ji and Telgarsky, 2020](#)). While implicit regularization from gradient descent trajectory is not always connected to ℓ_2 norm for regression (see counter examples in [Dauber et al. \(2020\)](#); [Razin and Cohen \(2020\)](#); [Li et al. \(2021\)](#)), the connection is prominent in many settings of interest. Most relevant to our work is the result by [Lyu and Li \(2020\)](#) (stated in Section 5), showing that parameters learned using gradient descent on logistic loss asymptotically converge in the direction of max- ℓ_2 -margin solution. We combine our results with this prior work to demonstrate that the conclusions from our analysis also extend to gradient descent solutions in classification problems.

Motivated by the connections to implicit and explicit regularization and generalization, other prior work also studies induced regularizers corresponding to minimizing ℓ_2 of weights in different architectures. Among recent work, [Savarese et al. \(2019\)](#); [Ongie et al. \(2020\)](#) provided a characterization of induced regularizer for infinite width two layer ReLU neural networks on 1D and higher dimensional inputs respectively. In a work closely related to ours, [Gunasekar et al. \(2018b\)](#) characterized the induced regularizer for fully connected networks and for linear convolutional network with *full dimensional kernels* ($K = D$) and *single-channel networks* ($R = C = 1$). [Yun et al. \(2021\)](#) extended [Gunasekar et al. \(2018b\)](#) and showed a general connection for linear networks between implicit ℓ_1 norm minimization in an orthonormal basis and the existence of data-independent diagonalizations of the linear operator in each layer. [Zhang et al. \(2020\)](#) empirically demonstrated such differences arising from fully connected vs convolutional architectures. However, even within the class of two-layer linear convolutional networks, the conclusions in prior work [Gunasekar et al. \(2018b\)](#); [Yun et al. \(2021\)](#) do not generalize to nontrivial kernel and channel sizes. For example, in the other extreme kernel size of $K = 1$, the induced regularizer is in fact the ℓ_2 norm of the linear function which is fundamentally different from the ℓ_1 norm of the Fourier coefficients for $K = D$, thus emphasizing the importance of our analysis for multi-channel networks with arbitrary kernel sizes. Subsequent to our work appearing as a preprint, [Dai et al. \(2021\)](#) also studied the induced regularizer for convolutional neural networks, though they do not investigate the role of the number of output channels for general kernel sizes.

Lastly, in a complementary approach, a line of work by [Ergen and Pilanci](#) and coauthors ([Pilanci and Ergen, 2020](#); [Ergen and Pilanci, 2021, 2020](#); [Sahiner et al., 2021](#)) study the induced regularizer

of neural networks including convolutional networks by looking at the bi-dual convex relaxation of the ℓ_2 regularized least squares loss. We elaborate on the connections more in Section 3.3.1. Briefly, in the context of linear convolutional networks, our results are significantly stronger, as their analysis shows invariance to number of output channels only in the limit of large C , while we show independence for all $C \geq 1$.

1.3. Notation

We typeset vectors, matrices, and tensors using bold characters, *e.g.*, $\mathbf{v}, \mathbf{x}, \boldsymbol{\theta}, \mathbf{W}, \mathbf{U}$. We will use zero-based indexing with notation $[D] = \{0, 1, \dots, D - 1\}$, and python style slicing notation to specify the sub-entries of an array variable: *e.g.*, given $\mathbf{Z} \in \mathbb{R}^{D_1 \times D_2}$, the d_1^{th} row and d_2^{th} column are denoted as $\mathbf{Z}[d_1, :] \in \mathbb{R}^{D_2}$ and $\mathbf{Z}[:, d_2] \in \mathbb{R}^{D_1}$, respectively. Complex numbers are specified in the polar form as $z = |z| e^{i\phi_z}$ with $\phi_z \in [0, 2\pi)$; or in Cartesian form as $z = \text{Re}(z) + i \text{Im}(z)$ (ref. $i = \sqrt{-1}$ is the imaginary unit). The complex conjugate is denoted as $\bar{z} = |z| e^{-i\phi_z}$. For $\mathbf{a}, \mathbf{b} \in \mathbb{C}^D$, the standard inner product is $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top \bar{\mathbf{b}}$, and analogously extends to matrices.

We use $\|\cdot\|$ to denote the standard Euclidean norm, *i.e.*, ℓ_2 norm of entries. For arrays \mathbf{a}, \mathbf{b} , $\mathbf{a} \odot \mathbf{b}$ denotes entry-wise multiplication and $\mathbf{a} \propto \mathbf{b}$ implies proportionality up to positive scaling. Finally, we define the convolution operator \star as it is used in the neural networks literature.²

Definition 2 (Circular convolution) For $\mathbf{u} \in \mathbb{R}^K$ and $\mathbf{v} \in \mathbb{R}^D$ with $K \leq D$, their D dimensional circular convolution, denoted by $\mathbf{u} \star \mathbf{v}$, is a vector in \mathbb{R}^D given as follows:

$$\forall_{d \in [D]}, (\mathbf{u} \star \mathbf{v})[d] = \frac{1}{\sqrt{D}} \sum_{k=0}^{K-1} \mathbf{u}[k] \mathbf{v}[(d+k) \bmod D].$$

1.4. Multi-channel linear convolutional network

We consider two layer linear convolutional networks with multiple channels in the convolution layer. We first focus on multi-output channel convolutions with single channel inputs described below. We will discuss networks with multi-channel inputs (*e.g.*, RGB color channels) in Section 4.

The inputs to the network are vectors³ of dimension D denoted as $\mathbf{x} \in \mathbb{R}^D$. The first layer is a convolutional layer with kernel size K and number of output channel C whose weights (parameters) are denoted by $\mathbf{U} \in \mathbb{R}^{K \times C}$. The output of the convolution layer, denoted as $h(\mathbf{U}; \mathbf{x}) \in \mathbb{R}^{D \times C}$, is given by $h(\mathbf{U}; \mathbf{x})[:, c] = \mathbf{U}[:, c] \star \mathbf{x}$ for all $c \in [C]$. The second layer is a single output linear layer with weights $\mathbf{V} \in \mathbb{R}^{D \times C}$. Thus, the output of the network, denoted as $\Phi(\mathbf{U}, \mathbf{V}; \mathbf{x})$, is given by:

$$\Phi(\mathbf{U}, \mathbf{V}; \mathbf{x}) = \langle \mathbf{V}, h(\mathbf{U}; \mathbf{x}) \rangle = \sum_{c=0}^{C-1} \langle \mathbf{V}[:, c], \mathbf{U}[:, c] \star \mathbf{x} \rangle. \quad (2)$$

Since, the network described above does not have any non-linearity, the output function $\Phi(\mathbf{U}, \mathbf{V}; \cdot)$ is equivalent to a linear representation $w(\mathbf{U}, \mathbf{V}) \in \mathbb{R}^D$ such that $\forall \mathbf{x}, \Phi(\mathbf{U}, \mathbf{V}; \mathbf{x}) = \langle w(\mathbf{U}, \mathbf{V}), \mathbf{x} \rangle$.

2. In signal processing, \star is known as the cross-correlation operator. To simplify analysis, we use circular padding in the definition (where mod refers to the modulo operator, *i.e.*, $p \bmod D = p - D \lfloor \frac{p}{D} \rfloor$). For convolutions with zero-padding, there will be different edge effects, but we expect qualitatively similar behavior for small padding sizes. We also use a scaling of $1/\sqrt{D}$ —this is merely to simplify notation and does not change the analysis.

3. For simplicity we consider 1D vectors $\mathbf{x} \in \mathbb{R}^D$ as inputs, but all our results can be extended to 2D inputs $\mathbf{x} \in \mathbb{R}^{W \times H}$, such as images, with the corresponding 2D convolutional operator.

Using standard algebraic manipulations on (2), one can derive $w(\mathbf{U}, \mathbf{V})$ as follows:

$$w(\mathbf{U}, \mathbf{V}) = \sum_{c=0}^{C-1} \left(\mathbf{U}[:, c] \star \mathbf{V}[:, c]^\downarrow \right)^\downarrow, \quad (3)$$

where \mathbf{z}^\downarrow denotes the flipped vector of $\mathbf{z} \in \mathbb{R}^D$, given by $\mathbf{z}^\downarrow[d] = \mathbf{z}[D - d - 1]$ for $d = 0, 1, \dots, D-1$.

Remark Even for the smallest network in this class with $K = C = 1$, any linear predictor $\mathbf{w} \in \mathbb{R}^D$ can be realized as $w(\mathbf{U}, \mathbf{V})$ in eq. (3) (e.g., using $\mathbf{U} = \mathbf{1}$, $\mathbf{V} = \mathbf{w}$). In fact, every linear predictor can be represented by multiple networks with different weights \mathbf{U}, \mathbf{V} .

Fourier representation. The convolution operation in Definition 2 permits a simple form in the Fourier domain arising from the *convolution theorem*. Let $\mathbf{F} \in \mathbb{C}^{D \times D}$ denote the unitary discrete Fourier transform (DFT) matrix for \mathbb{R}^D , i.e., $\forall_{k, l \in [D]}$, $\mathbf{F}[k, l] = \frac{1}{\sqrt{D}} e^{-\frac{2\pi i k l}{D}}$ and $\mathbf{F}\mathbf{F}^\top = \mathbf{F}^\top \mathbf{F} = \mathbf{I}$. For any $1 \leq K \leq D$, let $\mathbf{F}_K \in \mathbb{C}^{D \times K}$ denote the submatrix of \mathbf{F} with the first K columns. For a vector $\mathbf{a} \in \mathbb{R}^K$, we denote its D dimensional Fourier representation as $\widehat{\mathbf{a}} = \mathbf{F}_K \mathbf{a} \in \mathbb{C}^D$. From the definition of Fourier transform, we have that $\mathbf{F}\mathbf{a}^\downarrow = \overline{\mathbf{F}\mathbf{a}} = \widehat{\mathbf{a}}$ and we can derive the *convolution theorem* for our operator \star (Definition 2) as $\mathbf{F}(\mathbf{a} \star \mathbf{b}) = \widehat{\mathbf{a}} \odot \widehat{\mathbf{b}}$.

Let $\widehat{w}(\mathbf{U}, \mathbf{V}) := \mathbf{F}w(\mathbf{U}, \mathbf{V})$ denote the Fourier transform of the linear predictor $w(\mathbf{U}, \mathbf{V})$ realized by our network (see eq. 3). We can now express \widehat{w} as follows: Let $\widehat{\mathbf{U}} = \mathbf{F}_K \mathbf{U} \in \mathbb{C}^{D \times C}$ and $\widehat{\mathbf{V}} = \mathbf{F}\mathbf{V} \in \mathbb{C}^{D \times C}$ denote the D dimensional Fourier representation of \mathbf{U}, \mathbf{V} , respectively. We have

$$\widehat{w}(\mathbf{U}, \mathbf{V}) = \sum_{c=0}^{C-1} \widehat{\mathbf{U}}[:, c] \odot \widehat{\mathbf{V}}[:, c] = \text{diag}(\widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top). \quad (4)$$

2. Induced regularizer in the function space

For the network Φ described above, we now turn to the function space view of controlling the ℓ_2 norm of the weights (\mathbf{U}, \mathbf{V}) . Recall that this inductive bias is captured by the *induced regularizer* or the function space representation cost (1). For our linear convolutional network Φ , the function class realized is exactly the set of linear predictors $\mathbf{w} \in \mathbb{R}^D$ and the induced regularizer is given by:

$$\mathcal{R}_{K,C}(\mathbf{w}) := \min_{\mathbf{U} \in \mathbb{R}^{K \times C}, \mathbf{V} \in \mathbb{R}^{D \times C}} \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 \quad \text{s.t.,} \quad w(\mathbf{U}, \mathbf{V}) = \mathbf{w}. \quad (5)$$

Remark 3 It immediately follows from (5) that $\mathcal{R}_{K,C}(\mathbf{w})$ is weakly decreasing in both K and C , i.e., $\forall_C, \mathcal{R}_{1,C}(\mathbf{w}) \geq \mathcal{R}_{2,C}(\mathbf{w}) \geq \dots \mathcal{R}_{D,C}(\mathbf{w})$ and $\forall_K, \mathcal{R}_{K,1}(\mathbf{w}) \geq \mathcal{R}_{K,2}(\mathbf{w}) \geq \dots$

Even within the class of two layer linear convolutional networks, the induced regularizer can exhibit strikingly different properties for different choices of K and C . For example, we recall the following result from Gunasekar et al. (2018b) that for full dimensional kernel $K = D$, the induced regularizer is equal to the ℓ_1 norm of the Fourier transform of the predictor.

Lemma 4 ($K = D$) (Lemma 7 in Gunasekar et al., 2018b) For any $\mathbf{w} \in \mathbb{R}^D$, $\mathcal{R}_{D,1}(\mathbf{w}) = 2\|\widehat{\mathbf{w}}\|_1$.

On the other hand, we have the following characterization for $K = 1$ (full proof is in Appendix B).

Lemma 5 ($K = 1$) For any $\mathbf{w} \in \mathbb{R}^D$, $\mathcal{R}_{1,1}(\mathbf{w}) = 2\sqrt{D}\|\widehat{\mathbf{w}}\|_2 = 2\sqrt{D}\|\mathbf{w}\|_2$.

The induced regularizer thus behaves fundamentally differently for $K = D$ and $K = 1$. In particular, the ℓ_2 regularization of $\mathcal{R}_{1,1}(\mathbf{w})$ is basis agnostic and does not induce sparse solutions, while the ℓ_1 regularization of $\mathcal{R}_{D,1}(\mathbf{w})$ promotes sparsity in the Fourier basis.

Since $K = 1$ and $K = D$ permit closed-form solutions in the Fourier space, one might hope to obtain similarly clean characterizations for other kernel sizes as well. However, neither the proof technique for Lemma 4 nor the proof technique for Lemma 5 extend to the case of general kernel sizes. The proof of Lemma 4 uses the fact that for $K = D$, the weights $\mathbf{U}, \mathbf{V} \in \mathbb{R}^D$ are unconstrained in Fourier space. For networks with smaller kernels, the argument breaks as $\widehat{\mathbf{U}} = \mathbf{F}_K \mathbf{U}$ is constrained to be in a $K < D$ dimensional space spanned by the columns of \mathbf{F}_K . The proof of Lemma 5 again uses the special structure for $K = 1$ that $\widehat{\mathbf{U}} = \mathbf{F}_K \mathbf{U}$ points in the direction of $[1, \dots, 1]$, which does not extend to larger kernel sizes. In fact, we show in Appendix B.2 that even for $K = 2$, the induced regularizer $\mathcal{R}_{2,1}(\mathbf{w})$ takes a much more complex form.

Even though we do not obtain closed form solutions for all kernel sizes K , we derive important properties about the induced regularizer for general kernel sizes in the following sections that also generalize the above results to networks with multiple output channels.

3. Main technical tool: SDP formulation of induced regularizer

To investigate the induced regularizer for general kernel sizes and channel sizes, we construct a semidefinite program (SDP) relaxation, which would be a key tool in our analysis. In this section, we describe and analyze this SDP formulation for multi-output channel networks on inputs with a single channel. (We discuss generalizations to the case of multi-channel inputs in Section 4.)

We first reformulate $\mathcal{R}_{K,C}$ as an SDP with a rank constraint, which immediately motivates an SDP relaxation that provides a lower bound on $\mathcal{R}_{K,C}$. As we will show in Theorem 6, this SDP relaxation is actually tight for all K and C , which enables us to deduce a number of interesting properties of the induced regularizer.

$\mathcal{R}_{K,C}$ as an SDP with a rank constraint. Combining the definition of $\mathcal{R}_{K,C}(\mathbf{w})$ in (5) with the Fourier representation of $w(\mathbf{U}, \mathbf{V})$ in (4), we have the following:

$$\mathcal{R}_{K,C}(\mathbf{w}) = \min_{\mathbf{U} \in \mathbb{R}^{K \times C}, \mathbf{V} \in \mathbb{R}^{D \times C}} \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 \quad \text{s.t.,} \quad \text{diag}(\widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top) = \widehat{\mathbf{w}}. \quad (6)$$

The optimization in (6) over $\mathbf{U} \in \mathbb{R}^{K \times C}, \mathbf{V} \in \mathbb{R}^{D \times C}$, can be specified in terms of a rank C positive semi-definite matrix $\mathbf{Z} \in \mathbb{R}^{(D+K) \times (D+K)}$ that we define below:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{U}^\top & \mathbf{V}^\top \end{bmatrix} = \begin{bmatrix} \mathbf{U}\mathbf{U}^\top & \mathbf{U}\mathbf{V}^\top \\ \mathbf{V}\mathbf{U}^\top & \mathbf{V}\mathbf{V}^\top \end{bmatrix} \succcurlyeq 0. \quad (7)$$

We refer to \mathbf{Z} as the lifted space of parameters \mathbf{U}, \mathbf{V} . In the lifted space, the objective and constraints of (6) can now be expressed as linear functions of \mathbf{Z} . The objective is given by $\|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 = \langle \mathbf{U}\mathbf{U}^\top, \mathbf{I} \rangle + \langle \mathbf{V}\mathbf{V}^\top, \mathbf{I} \rangle = \langle \mathbf{Z}, \mathbf{I} \rangle$. The constraints of (6) are given by $\forall_{d \in [D]}, \langle \widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top, \mathbf{e}_d \mathbf{e}_d^\top \rangle = \widehat{\mathbf{w}}[d]$, where $\{\mathbf{e}_d\}_{d \in [D]}$ denotes the standard basis. Alternatively, using $\widehat{\mathbf{U}} = \mathbf{F}_K \mathbf{U}$ and $\widehat{\mathbf{V}} = \mathbf{F}\mathbf{V}$, the constraints are given by $\forall_{d \in [D]}, \langle \mathbf{U}\mathbf{V}^\top, \overline{\mathbf{F}}_K^\top \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}} \rangle = \widehat{\mathbf{w}}[d]$, which in the lifted space is given by

$\langle \mathbf{Z}, \mathbf{A}_d^{\text{real}} \rangle = 2 \cdot \text{Re}(\widehat{\mathbf{w}}[d])$ and $\langle \mathbf{Z}, \mathbf{A}_d^{\text{img}} \rangle = 2 \cdot \text{Im}(\widehat{\mathbf{w}}[d])$, where we define $(\mathbf{A}_d^{\text{real}}, \mathbf{A}_d^{\text{img}})$ as follows:

$$\mathbf{A}_d^{\text{real}} = \begin{bmatrix} 0_K & \overline{\mathbf{F}}_K^\top \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}} \\ \overline{\mathbf{F}}^\top \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}}_K & 0_D \end{bmatrix} \quad \text{and} \quad \mathbf{A}_d^{\text{img}} = \begin{bmatrix} 0_K & i \cdot \overline{\mathbf{F}}_K^\top \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}} \\ -i \overline{\mathbf{F}}^\top \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}}_K & 0_D \end{bmatrix}.$$

Now, we can formulate $\mathcal{R}_{K,C}(\mathbf{w})$ as follows:

$$\begin{aligned} \mathcal{R}_{K,C}(\mathbf{w}) = \min_{\mathbf{Z} \succeq 0} \langle \mathbf{Z}, \mathbf{I} \rangle \quad \text{s.t.}, \quad & \forall_{d \in [D]}, \langle \mathbf{Z}, \mathbf{A}_d^{\text{real}} \rangle = 2 \text{Re}(\widehat{\mathbf{w}}[d]) \\ & \forall_{d \in [D]}, \langle \mathbf{Z}, \mathbf{A}_d^{\text{img}} \rangle = 2 \text{Im}(\widehat{\mathbf{w}}[d]) \\ & \text{rank}(\mathbf{Z}) \leq C. \end{aligned} \quad (8)$$

The formulation in eq. (8) is non-convex due to the rank constraint. We obtain a natural convex relaxation by dropping the rank constraint, leading to the following SDP:

$$\begin{aligned} \mathcal{R}_K^{\text{SDP}}(\mathbf{w}) = \min_{\mathbf{Z} \succeq 0} \langle \mathbf{Z}, \mathbf{I} \rangle \quad \text{s.t.}, \quad & \forall_{d \in [D]}, \langle \mathbf{Z}, \mathbf{A}_d^{\text{real}} \rangle = 2 \text{Re}(\widehat{\mathbf{w}}[d]) \\ & \forall_{d \in [D]}, \langle \mathbf{Z}, \mathbf{A}_d^{\text{img}} \rangle = 2 \text{Im}(\widehat{\mathbf{w}}[d]). \end{aligned} \quad (9)$$

Remark *By construction, the relaxation provides lower bounds on the induced regularizer: for any $K \leq D$, any C , and any $\mathbf{w} \in \mathbb{R}^D$, it holds that $\mathcal{R}_{K,C}(\mathbf{w}) \geq \mathcal{R}_K^{\text{SDP}}(\mathbf{w})$.*

Remark *The symmetry properties of Fourier coefficients of real signals gives us that for any D and any $\mathbf{w} \in \mathbb{R}^D$, $\widehat{\mathbf{w}}[p] = \overline{\widehat{\mathbf{w}}[D-p]}$ for $p \in [D]$. Thus, although the optimization problems in (8) and (9) are specified with $2 \cdot D$ constraints for simplicity, only D of them are unique.*

3.1. Tightness of the SDP Relaxation

Our main technical result is that for any kernel size K , the SDP relaxation is tight (in the case of networks with single-channel inputs). Thus, the induced regularizer $\mathcal{R}_{K,C}$ is equivalent to an SDP that only depends upon on the kernel size K .

Theorem 6 [SDP tightness] *For any K, C , and $\mathbf{w} \in \mathbb{R}^D$, it holds that $\mathcal{R}_{K,C}(\mathbf{w}) = \mathcal{R}_K^{\text{SDP}}(\mathbf{w})$.*

Proof sketch. We can show directly from the KKT conditions that any minimizer \mathbf{Z} of the SDP must have rank at most K . However, to prove Theorem 6, we need to show that there exists a rank 1 solution that has the same objective value as \mathbf{Z} and satisfies the SDP constraints—this does not follow directly from the KKT conditions.

Constructing this rank-1 solution is the main technical hurdle in the proof of Theorem 6. In particular, the following lemma is a key intermediate result about the convolutional operation and is of independent interest beyond this paper.

Lemma 7 *For any $1 \leq K \leq D$, and for any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, there exists a vector $\mathbf{c} \in \mathbb{R}^K$ such that $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b} = \mathbf{c} \star \mathbf{c}$, where convolutions are w.r.t. dimension D .*

For $K = D$, Lemma 7 follows easily from the Fourier space representation (using $\mathbf{z} \star \mathbf{z} = |\widehat{\mathbf{z}}|^2$), since in this case $\widehat{\mathbf{c}}$ is unconstrained and can be explicitly constructed as the square root of $|\widehat{\mathbf{a}}|^2 + |\widehat{\mathbf{b}}|^2$. However, this construction does not generalize to kernel sizes $K < D$. In fact, \mathbf{c} does not appear to have an explicit closed-form for general $K < D$. At the core of our proof, we first provide an argument for the existence of \mathbf{c} in the special case of $D = 2K - 1$ and then show that the case of general $1 \leq K \leq D$ follows from this special case.

Polynomial representation: In our proof, we show existence of \mathbf{c} in Lemma 7 by using the representation of convolutions as polynomial multiplication. This representation relies on the isomorphism between \mathbb{R}^D and polynomials of degree $\leq D - 1$ with real coefficients, *i.e.*, $\mathbf{a} \in \mathbb{R}^D \equiv p_{\mathbf{a}}(x) = \mathbf{a}[0] + \mathbf{a}[1]x + \mathbf{a}[2]x^2 + \dots + \mathbf{a}[D - 1]x^{D-1}$. Lemma 7 can be written in terms of polynomials as follows: for any real-coefficient polynomials $p_{\mathbf{a}}, p_{\mathbf{b}}$ of degree at most $K - 1$, there exists a real-coefficient polynomial $p_{\mathbf{c}}$ of degree at most $K - 1$ such that:

$$x^{K-1}p_{\mathbf{c}}(x)p_{\mathbf{c}}(1/x) = x^{K-1}p_{\mathbf{a}}(x)p_{\mathbf{a}}(1/x) + x^{K-1}p_{\mathbf{b}}(x)p_{\mathbf{b}}(1/x).$$

The remainder of the proof involves implicitly constructing $p_{\mathbf{c}}$ in terms of its roots (and leading coefficient). To do so, we show that the roots of the polynomial $x^{K-1}p_{\mathbf{a}}(x)p_{\mathbf{a}}(1/x) + x^{K-1}p_{\mathbf{b}}(x)p_{\mathbf{b}}(1/x)$ satisfy certain structural properties which allow us to establish the existence of the desired real-coefficient polynomial $p_{\mathbf{c}}$. The full proof is in Appendix A with additional details on the proof technique in Appendix A.3.

3.2. Implications of SDP Tightness in Theorem 6

The first implication of Theorem 6 is that although the optimization in (5) is non-convex, the SDP formulation allows us to efficiently compute $\mathcal{R}_{K,C}(\mathbf{w})$ exactly. In the remainder of the section, we discuss a number of other interesting properties that we can deduce from Theorem 6.

3.2.1. $\mathcal{R}_{K,C}$ IS INDEPENDENT OF NUMBER OF OUTPUT CHANNELS C

Theorem 6 directly implies that $\mathcal{R}_{K,C}$ is independent of C . This means that the linear predictors obtained by fitting training data and minimizing $\mathcal{R}_{K,C}(\mathbf{w})$ will be invariant to C (apart from $\mathcal{R}_{K,C}(\mathbf{w})$ having multiple minimizers). Based on previous work (e.g. Lyu and Li (2020)), this has implications for the asymptotic behavior of gradient descent. In particular, we can hypothesize that for networks with single channel input, the number of output channels does not influence the asymptotic predictor learned from gradient descent.

We provide a detailed empirical evaluation of this hypothesis (along with a generalization of this hypothesis for networks with multi-channel inputs) in Section 5. As a preview, in Figure 1 we show the predictors learned by gradient descent on an MNIST task on a two-layer linear convolutional network with kernel size $K = 3$. Furthermore, our experiments also suggest that our theoretical results might also extend to some cases of networks with ReLU non linearity and bias (see Figure 2). In both cases, we see that the induced regularizer is invariant to the number of output channels. We defer a more extensive empirical evaluation to Section 5.

3.2.2. $\mathcal{R}_{K,C}$ IS A NORM

Another interesting corollary of Theorem 6 is that the induced regularizer is a norm for any K .

Corollary 8 *For any K and any C , $\mathcal{R}_{K,C}(\mathbf{w})$ is a norm.*




Figure 1: Linear predictors learned by two layer linear convolutional network for the task of classifying digits 0 and 1 in MNIST. The sub-figures depict predictors learned by using gradient descent on the exponential loss for over-parameterized networks with kernel size $K = (3, 3)$ and number of output channels $C \in \{1, 2, 4, 8\}$ (left to right).

C	$K : (1, 1)$	$K : (3, 3)$	$K : (8, 8)$
1	10.581	4.948	3.875
2	10.571	4.945	3.910
4	10.578	4.945	3.912
8	10.576	4.946	3.881

Figure 2: $\widehat{\mathcal{R}}_{K,C}(f_{\text{GD}}) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the MNIST 0-vs-1 classifiers learned by gradient descent on two-layer ReLU convolutional networks with bias parameters. We show the median values over 5 trials.

For the end cases of $K = 1$ and $K = D$, this norm can be explicitly specified: $\mathcal{R}_{1,C}(\mathbf{w}) = 2\sqrt{D}\|\widehat{\mathbf{w}}\|$ (Lemma 5) and $\mathcal{R}_{D,C}(\mathbf{w}) = 2\|\widehat{\mathbf{w}}\|_1$ (Lemma 4), respectively. For intermediate kernel sizes, $\mathcal{R}_{K,C}(\mathbf{w})$ interpolates between the ℓ_2 norm and the ℓ_1 norm of the Fourier coefficients of the linear predictor. We further use the SDP in (9) to bound $\mathcal{R}_{K,C}(\mathbf{w})$ in terms of the ℓ_1 and ℓ_2 norms.

Lemma 9 For any K, C , and $\mathbf{w} \in \mathbb{R}^D$:

$$2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_2 \leq \mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{D}\|\widehat{\mathbf{w}}\|_2 \quad \text{and} \quad 2\|\widehat{\mathbf{w}}\|_1 \leq \mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{\lceil \frac{D}{K} \rceil}\|\widehat{\mathbf{w}}\|_1.$$

Remark For the lower bounds, $2\|\widehat{\mathbf{w}}\|_1$ is tight for $\mathbf{w} = [1, 0, \dots, 0]$ and $2\sqrt{\frac{D}{K}}\|\mathbf{w}\|$ is tight for $\mathbf{w} = [1, 1, \dots, 1]$. For the upper bounds, $2\sqrt{\lceil \frac{D}{K} \rceil}\|\widehat{\mathbf{w}}\|_1$ is tight when $K \mid D$ for patterned vectors (see Lemma 10), and $2\sqrt{D}\|\widehat{\mathbf{w}}\|_2$ is tight for $[1, 0, \dots, 0]$.

Lemma 9 demonstrates that when K is a small constant, $\mathcal{R}_{K,C}(\mathbf{w})$ is close to $\mathcal{R}_{1,C}(\mathbf{w}) = 2\sqrt{D}\|\widehat{\mathbf{w}}\|_2$, whereas once K is comparable to D , $\mathcal{R}_{K,C}(\mathbf{w})$ is close to $\mathcal{R}_{D,C}(\mathbf{w}) = 2\|\widehat{\mathbf{w}}\|_1$.

3.3. $\mathcal{R}_{K,C}$ for patterned vectors

Aside from general bounds on $\mathcal{R}_{K,C}$, the SDP formulation can also be used to analyze the behavior of the induced regularizer of special classes of vectors. One interesting case is of *patterned vectors* described as follows: Consider vectors of the form $\mathbf{w}(\mathbf{p}) = [\mathbf{p}, \mathbf{p}, \dots, \mathbf{p}] \in \mathbb{R}^D$ consisting of repetitions of a P dimensional pattern $\mathbf{p} \in \mathbb{R}^P$. A useful property of linear predictors of this form is that they incorporate invariance to periodic translations.

We show a relation between the representation cost $\mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p}))$ of realizing patterned vectors in \mathbb{R}^D and the analogous cost (denoted as $\mathcal{R}_{K,C}^{(P)}(\mathbf{p})$) of realizing \mathbf{p} as a linear predictor in \mathbb{R}^P using a network with the same values of K and C .

Lemma 10 Consider vectors $\mathbf{w}(\mathbf{p}) = [\mathbf{p}, \mathbf{p}, \dots, \mathbf{p}] \in \mathbb{R}^D$ specified by $\mathbf{p} \in \mathbb{R}^P$ s.t., P divides D .

(a) For $K \leq P$, it holds that $\forall C: \mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p})) = \frac{D}{P} \cdot \mathcal{R}_{K,1}^{(P)}(\mathbf{p})$.

(b) For $K \geq P$ if $K = P \cdot T$ for integer T , then $\forall C: \mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p})) = \frac{2D}{\sqrt{TP}}\|\widehat{\mathbf{p}}\|_1 = 2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_1$.

We see that the induced regularizer of repeated patterned vectors is closely related to that of the pattern itself. In particular, for $K \leq P$, we have $\mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p})) \propto \mathcal{R}_{K,1}^{(P)}(\mathbf{p})$.

3.3.1. CONNECTION TO PREVIOUS WORK

Comparison to Pilanci and Ergen (2020); Ergen and Pilanci (2021, 2020); Sahiner et al. (2021). These works study the induced regularizer of neural networks by looking at the bi-dual convex relaxation of the ℓ_2 regularized least squares loss. In comparison to our SDP, this method is a complementary approach to derive lower bounds on the induced regularizer when minimizing convex losses over datasets. In both cases, the relaxations are trivially tight in the limit of *infinitely many* output channels. However, in our work, we use the SDP formulation to show a significantly stronger result than those in prior works.

Phrased in the terminology of our work, the results in Pilanci and Ergen (2020) on linear convolutional networks show that for networks with a single input channel, *if the number of output channels (or width) C is larger than a data-dependent threshold*, then $\mathcal{R}_{K,C}(\mathbf{w}) = \mathcal{R}_K^{\text{SDP}}(\mathbf{w})$. That is, they show that the induced regularizer is independent of C after C is above a certain large finite value that can be large as the dataset size. In contrast, we show the induced regularizer independent of C for any $C \geq 1$. Further, our analysis holds regardless of the dataset and training loss.

Comparison to Gunasekar et al. (2018b). Gunasekar et al. (2018b) characterized the induced regularizer for single-channel networks with full-dimensional kernels. As we discussed in Section 2, the conclusions and closed-form solution for $K = D$ do not extend to the full class of two-layer linear convolutional networks. In contrast, our result (Theorem 6) implies properties of the induced regularizer for networks with arbitrary kernel size.

Moreover, these prior works focus on networks with single-channel inputs, whereas we extend our results to networks with multi-channel inputs in the next section.

4. Networks with multi-channel inputs

While we focused on networks with single-channel inputs in the previous sections, we now expand our results to networks with multiple input channels (*e.g.*, RGB color channels). How do these conclusions change for multiple input channels? How does the induced regularizer, now denoted as $\mathcal{R}_{K,C,R}$, depend on the number input channels R ?

We again consider two layer convolutional networks akin to Section 1.4. We first introduce additional notation: The multi-channel inputs are denoted as $\mathbf{X} \in \mathbb{R}^{D \times R}$, where R denotes the number of input channels. The convolutional first layer now has kernel size K , output channel size C and input channel size R with weights denoted by a set of R matrices $\mathbf{U} = \{\mathbf{U}_r\}_{r \in [R]}$ with $\mathbf{U}_r \in \mathbb{R}^{K \times C}$. The output of this convolution layer $h(\mathbf{U}; \mathbf{X}) \in \mathbb{R}^{D \times C}$ is given as follows:

$$\forall_{c \in C}, h(\mathbf{U}; \mathbf{X})[:, c] = \sum_{r=0}^{R-1} \mathbf{U}_r[:, c] \star \mathbf{X}[:, r]. \quad (10)$$

The second layer is the same as before: a single output linear layer with weights $\mathbf{V} \in \mathbb{R}^{D \times C}$. We denote the equivalent linear predictor for this network by $W(\mathbf{U}, \mathbf{V}) \in \mathbb{R}^{D \times R}$. Following similar calculations as for single input channels, $W(\mathbf{U}, \mathbf{V})$ in signal and Fourier domain (denoted as $\widehat{W}(\mathbf{U}, \mathbf{V}) = \mathbf{F}W(\mathbf{U}, \mathbf{V})$) are given as follows:

$$\forall_{r \in [R]}, W(\mathbf{U}, \mathbf{V})[:, r] = \sum_{c=0}^{C-1} \left(\mathbf{U}_r[:, c] \star \mathbf{V}[:, c]^\downarrow \right)^\downarrow, \text{ and } \widehat{W}(\mathbf{U}, \mathbf{V})[:, r] = \text{diag}(\widehat{\mathbf{U}}_r \widehat{\mathbf{V}}^\top). \quad (11)$$

For multi-channel inputs, the set of all linear predictors is the space of matrices $\mathbf{W} \in \mathbb{R}^{D \times R}$, and we define the induced complexity measure over this matrix space as follows:

$$\mathcal{R}_{K,C,R}(\mathbf{W}) := \inf_{\mathbf{U}, \mathbf{V}} \sum_{r \in [R]} \|\mathbf{U}_r\|^2 + \|\mathbf{V}\|^2 \quad \text{s.t.}, \quad W(\mathbf{U}, \mathbf{V}) = \mathbf{W}. \quad (12)$$

4.1. Role of output channel size C

For multi-channel inputs, we first observe that multiple output channels can be necessary to realize all linear maps. To see this, we show that the sub-network corresponding to each output channel can realize a matrix in $\mathbb{R}^{D \times R}$ of rank at most K , which places an upper bound on the total rank achievable by the full network (see a proof in Appendix D.1). This implies the following lemma:

Lemma 11 *For any K, C and R , in order for the the network represented by $W(\mathbf{U}, \mathbf{V})$ in eq. (11) to realize all linear maps in $\mathbb{R}^{D \times R}$ it is necessary that $K \cdot C \geq \min\{R, D\}$.*

In contrast to single input channels, Lemma 11 demonstrates that, the model class realized by linear convolutional networks over multi-channel inputs, and consequently the induced regularizer, *does* depend on number of output channels C . Nonetheless, similar to single input channel networks, we can again obtain an SDP relaxation $\mathcal{R}_{K,R}^{\text{SDP}}(\mathbf{W})$ for $\mathcal{R}_{K,C,R}(\mathbf{W})$ that is independent of C .

4.2. SDP relaxation for multi-channel input networks

The SDP relaxation for $\mathcal{R}_{K,C,R}$ is derived similarly to for networks with a single input channel (see additional details in Appendix D). For $\mathbf{U} = \{\mathbf{U}_r \in \mathbb{R}^{K \times C}\}_{r \in [R]}$, $\mathbf{V} \in \mathbb{R}^{D \times C}$, we can represent the objective and constraints of (12) as linear functions of a rank C positive semidefinite matrix $\mathbf{Z} \in \mathbb{R}^{(D+K \cdot R) \times (D+K \cdot R)}$, that represents:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{U}_0 \\ \vdots \\ \mathbf{U}_R \\ \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{U}_0^\top & \dots & \mathbf{U}_R^\top & \mathbf{V}^\top \end{bmatrix} \succeq 0.$$

We also define Hermitian matrices $\mathbf{A}_{d,r}^{\text{real}}, \mathbf{A}_{d,r}^{\text{img}} \in \mathbb{R}^{(D+K \cdot R) \times (D+K \cdot R)}$ for $d \in [D], r \in [R]$ as follows: let $\mathbf{Q}_d = \bar{\mathbf{F}}_K^\top \mathbf{e}_d \mathbf{e}_d^\top \bar{\mathbf{F}}$, these matrices are give by

$$\mathbf{A}_{d,0}^{\text{real}} = \begin{bmatrix} \mathbf{0}_{(R \cdot K) \times (R \cdot K)} & \mathbf{Q}_d \\ \hline \bar{\mathbf{Q}}_d & \mathbf{0}_{D \times D} \end{bmatrix}, \quad \mathbf{A}_{d,0}^{\text{img}} = \begin{bmatrix} \mathbf{0}_{(R \cdot K) \times (R \cdot K)} & i \cdot \mathbf{Q}_d \\ \hline -i \cdot \bar{\mathbf{Q}}_d & \mathbf{0}_{D \times D} \end{bmatrix},$$

$$\vdots$$

$$\mathbf{A}_{d,R}^{\text{real}} = \begin{bmatrix} \mathbf{0}_{(R \cdot K) \times (R \cdot K)} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{Q}_d \end{bmatrix}, \quad \mathbf{A}_{d,R}^{\text{img}} = \begin{bmatrix} \mathbf{0}_{(R \cdot K) \times (R \cdot K)} & \mathbf{0} \\ \hline \mathbf{0} & i \cdot \mathbf{Q}_d \end{bmatrix}.$$

We now have the following relaxation of $\mathcal{R}_{K,C,R}(\mathbf{W})$ in terms of Fourier coefficients $\widehat{\mathbf{W}} = \mathbf{F}\mathbf{W}$:

$$\begin{aligned} \mathcal{R}_{K,R}^{\text{SDP}}(\mathbf{W}) &= \min_{\mathbf{Z} \succeq 0} \langle \mathbf{Z}, \mathbf{I} \rangle \\ \text{s.t. } &\forall d \in [D], r \in [R] \quad \langle \mathbf{Z}, \mathbf{A}_{d,r}^{\text{real}} \rangle = 2\text{Re}(\widehat{\mathbf{W}}[d, r]) \\ &\forall d \in [D], r \in [R] \quad \langle \mathbf{Z}, \mathbf{A}_{d,r}^{\text{img}} \rangle = 2\text{Im}(\widehat{\mathbf{W}}[d, r]). \end{aligned} \quad (13)$$

We can check that the SDP in (13) with a rank constraint of $\text{rank}(\mathbf{Z}) \leq C$ is equivalent $\mathcal{R}_{K,C,R}(\mathbf{W})$ and the SDP thus provides a lower bound: *i.e.*, $\forall \mathbf{W}$, $\mathcal{R}_{K,C,R}(\mathbf{W}) \geq \mathcal{R}_{K,R}^{\text{SDP}}(\mathbf{W})$.

4.3. Tightness of SDP Relaxation

Unlike networks with single channel input, the SDP relaxation here is not always tight when $R > 1$, since a sufficiently large C is required to merely realize all matrix-valued linear function over the input space. We can however show a weaker form SDP tightness from the KKT conditions when there are sufficiently many output channels:

Lemma 12 *For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq RK$, it holds that $\mathcal{R}_{K,C,R}(\mathbf{W}) = \mathcal{R}_{K,R}^{\text{SDP}}(\mathbf{W})$.*

Note that the above bound on C for SDP tightness is not sharp, as we showed for $R = 1$ in Theorem 6. Based on our insights from the proof of single-channel SDP tightness in Theorem 6 and additional empirical evidence in Appendix E, we conjecture that SDP tightness holds when $C \geq R$:

Conjecture 13 *For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq R$, it holds that $\mathcal{R}_{K,C,R}(\mathbf{W}) = \mathcal{R}_{K,R}^{\text{SDP}}(\mathbf{W})$.*

In the next subsection, we prove Conjecture 13 in the special cases of $K = 1$ and $K = D$. As a consequence, we show that once C is large enough to realize all linear maps, $\mathcal{R}_{K,C,R}(\mathbf{W})$ can be expressed as interesting closed form norms independent of C in these special cases.

4.4. Induced regularizer when $K = 1$ and $K = D$

Theorem 14 *For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq \min\{R, D\}$, the induced regularizer for $K = 1$ is given by the scaled nuclear norm $\|\cdot\|_*$:*

$$\mathcal{R}_{1,C,R}(\mathbf{W}) = 2\sqrt{D}\|\mathbf{W}\|_* = 2\sqrt{D}\|\widehat{\mathbf{W}}\|_*.$$

Theorem 15 *For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq 1$, the induced regularizer for $K = D$ is given by*

$$\mathcal{R}_{D,C,R}(\mathbf{W}) = 2\|\widehat{\mathbf{W}}\|_{2,1} := \sum_{d=0}^{D-1} \sqrt{\sum_{r=0}^{R-1} |\widehat{\mathbf{W}}[d, r]|^2}.$$

From Theorems 14-15 it is evident that the number of input channels R fundamentally changes the nature of induced complexity measure in the function space and introduces additional structures along the input channels. Even in the simplest setting of scalar convolution kernels with $K = 1$, the induced regularizer is no longer a Euclidean or RKHS norm, and is instead a richer nuclear norm that encourages low-rank properties. For the case of $K = D$, the induced regularizer is group-sparse

norm on the Fourier coefficients that encourages similar weighting across channels, while promoting sparsity across frequency components. In comparison to the ℓ_1 norm of all Fourier coefficients, this group-sparse norm is a more structured inductive bias for multi-channel inputs. Additionally, like with the single input channel case, we also observe that the induced bias has a more intuitive and interesting interpretation in Fourier domain which is not directly observed in the signal domain.

5. Experiments

We now explicitly connect our findings to the implicit regularization of gradient descent. We formally state the following result paraphrased from [Lyu and Li \(2020\)](#) that relate the asymptotic implicit bias of gradient descent to ℓ_2 norm minimization of the parameters:

Theorem (Paraphrased from [Lyu and Li, 2020](#)) *Assume that Φ is locally Lipschitz and positive homogeneous with order $L > 0$, i.e., $\forall \theta, \alpha > 0, \Phi(\alpha\theta; \cdot) = \alpha^L \Phi(\theta; \cdot)$. Consider minimization of logistic or exponential loss over a separable binary classification dataset $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ using a gradient flow trajectory denoted as $\theta(t)$. Under the assumption that the training data points are correctly classified in finite time (i.e., mean logistic/exponential loss $\mathcal{L}(\theta(t_0)) < 1$ at finite t_0), the limit points of the direction of parameters along the gradient flow correspond to a first order stationary point (KKT point) of the following max- ℓ_2 margin problem in parameter space:*

$$\min_{\theta} \|\theta\|_2^2 \quad \text{s.t.}, \quad \forall_n y_n \Phi(\theta; \mathbf{x}_n) \geq 1. \quad (14)$$

Remark *With additional assumptions, [Lyu and Li \(2020\)](#) also prove an analogous result for gradient descent. This result was further refined in [Ji and Telgarsky \(2020\)](#), which proves the directional convergence of the gradient descent under additional structure on the model class.*

Based on these results, we thus expect the implicit bias from gradient descent to be related to the max- \mathcal{R} -margin problem

$$\min_f \mathcal{R}_\Phi(f) \quad \text{s.t.}, \quad \forall_n y_n f(\mathbf{x}_n) \geq 1.$$

However, theoretically speaking, there is an important caveat: the theorem by [Lyu and Li \(2020\)](#) shows convergence of gradient flow direction to a *stationary point* of the optimization in eq. (14).

Nonetheless, if we overlook these caveats, our findings would then have important implications for predictors learned from gradient descent. In particular, our result regarding the invariance of the induced regularizer with respect to the number of output channels suggests that the asymptotic behavior of gradient descent is similarly invariant to the number of output channels. We formalize this as a testable hypothesis.

Hypothesis 1 *For a separable binary classification task with R input channels, let \mathbf{w}_{GD} be the predictor learned using stochastic gradient descent on a two-layer convolutional network with kernel size K , C output channels, and R input channels (where \mathbf{w}_{GD} is normalized to have unit margin on the training data). Then, as long as $C \geq R$, the induced regularizer $\mathcal{R}_{K,C,R}(\mathbf{w}_{GD})$ is invariant in the number of output channels C .*

We show experimental support for the hypothesis on small linearly separable subsets of MNIST (with 128 images of size 28×28 balanced across 2 classes) and CIFAR-10 (with 512 images of size 32×32 balanced across 2 classes) datasets. Most of our experiments are for multi-channel linear

convolutional networks trained using stochastic gradient descent. We also provide some experiments on ReLU networks, where we see support of our hypothesis well beyond our theoretical study.

Throughout the experiments sections, since we cannot always compute $\mathcal{R}_{K,C,R}(\mathbf{w}_{\text{GD}})$, we approximate it using the weight norms of the trained network $\widehat{\mathcal{R}}_{K,C,R}(\mathbf{w}_{\text{GD}}) = \sum_{\mathbf{U}, \mathbf{V}} \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$, where \mathbf{U}, \mathbf{V} here denote the weights of the trained network.⁴ The experiments are deferred to Appendix E and we summarize the findings below.

1. *Single input channel binary classification on MNIST.* In linear networks, we compare the predictors learned by gradient descent for $K \in \{1, 3, 8, 16, 28\}$ and $C \in \{1, 2, 4, 8\}$ across 10 runs with random initialization. We see that both the values of estimated regularizer $\widehat{\mathcal{R}}$ as well as the visualization of linear predictors in signal and frequency domain are nearly invariant to C (the values overlap within one standard deviation across runs).
2. *3-input channel binary classification on CIFAR-10.* In a similar setup to MNIST, we compare $K \in \{1, 3, 8, 20\}$ and $C \in \{1, 2, 3, 4, 8\}$. As expected from our theory, we see differences in the induced regularizer $\widehat{\mathcal{R}}$ for $C < 3$, but observe invariance to C once $C \geq 3$.
3. *ReLU networks for binary classification on MNIST.* Although our theory is only for linear networks, our hypothesis as stated above can also be tested on networks with non-linearity. We repeat our MNIST experiments on networks with ReLU non-linearity (with and without bias parameters). Interestingly, we observe that the estimated induced regularizer $\widehat{\mathcal{R}}$ is invariant to C suggesting a broader scope for our hypothesis.

These findings support Hypothesis 1 beyond the scope of our theoretical results.

6. Discussion

We showed that when minimizing ℓ_2 norm of weights, the two basic architectural components of convolutional networks—number of output channels (width) and kernel size—have interesting effects even in the simple case of two-layer linear networks. Our results also inspire a broader hypothesis about the impact of number of output channels for networks learned with gradient descent, which we tested and provided support for in our experiments.

Interesting directions for future work include proving tightness of the SDP relaxation for multiple input channels (formalized in Conjecture 13); formally establishing the limiting behavior of gradient descent; and exploring architectural features such as pooling or multiple layers. Furthermore, it would be interesting to conduct an in-depth empirical investigation of the impact of non-linearity.

References

- Peter L. Bartlett. For valid generalization the size of the weights is more important than the size of the network. In *Advances in Neural Information Processing Systems*, 1996.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, pages 463–482, 2002.
- Peter L. Bartlett, Dylan J. Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6241–6250, 2017.

4. In theory $\widehat{\mathcal{R}}$ only provides an upper bound on \mathcal{R} —but in case of predictors learned by SGD, upon checking instances where \mathcal{R} has a closed form solution, we found that the approximation is quite accurate.

Zhen Dai, Mina Karzand, and Nathan Srebro. Representation costs of linear neural networks: Analysis and design. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

Assaf Dauber, Meir Feder, Tomer Koren, and Roi Livni. Can implicit bias explain generalization? stochastic convex optimization as a case study. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

Tolga Ergen and Mert Pilanci. Implicit convex regularizers of cnn architectures: Convex optimization of two- and three-layer networks in polynomial time. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, 2020.

Tolga Ergen and Mert Pilanci. Revealing the structure of deep neural networks via convex duality. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139, pages 3004–3014. PMLR, 2021.

Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1827–1836, 2018a.

Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems*, 2018b.

Ziwei Ji and Matus Telgarsky. Risk and parameter convergence of logistic regression. *CoRR*, abs/1803.07300, 2018.

Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations (ICLR)*, 2019.

Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In *Advances in Neural Information Processing Systems*, 2020.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

Anders Krogh and John A. Hertz. A simple weight decay can improve generalization. In *Advances in Neural Information Processing Systems*, pages 950–957, 1991.

Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.

Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.

Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations (ICLR)*, 2020.

- Mor Shpigel Nacson, Suriya Gunasekar, Jason D. Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 4683–4692, 2019.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations (ICLR), Workshop Track Proceedings*, 2015.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width relu nets: The multivariate case. In *International Conference on Learning Representations (ICLR)*, 2020.
- Mert Pilanci and Tolga Ergen. Neural networks are convex regularizers: Exact polynomial-time convex optimization formulations for two-layer networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119, pages 7695–7705, 2020.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Jason D. M. Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 713–719, 2005.
- Arda Sahiner, Tolga Ergen, John M. Pauly, and Mert Pilanci. Vector-output relu neural network problems are copositive programs: Convex analysis of two layer networks and polynomial-time algorithms. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Proceedings of the Conference on Learning Theory (COLT)*, pages 2667–2690, 2019.
- Colin Wei, Jason D. Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets v.s. their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9709–9721, 2019.
- Chulhee Yun, Shankar Krishnan, and Hossein Mobahi. A unifying view on implicit bias in training linear neural networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Michael C. Mozer, and Yoram Singer. Identity crisis: Memorization and generalization under extreme overparameterization. In *International Conference on Learning Representations (ICLR)*, 2020.

Appendix A. Proof of Theorem 6: SDP tightness

Theorem 6 [SDP tightness] For any K, C , and $\mathbf{w} \in \mathbb{R}^D$, it holds that $\mathcal{R}_{K,C}(\mathbf{w}) = \mathcal{R}_K^{\text{SDP}}(\mathbf{w})$.

The high-level idea of the proof of Theorem 6 is to take an optimal solution \mathbf{Z} of the $\mathcal{R}_K^{\text{SDP}}(\mathbf{w})$ problem in (9), and construct a rank 1 solution that obtains the same objective and satisfies the constraints. We reiterate the SDP formulation for easy reference:

$$\begin{aligned} \mathcal{R}_K^{\text{SDP}}(\mathbf{w}) = \min_{\mathbf{Z} \succcurlyeq 0} \langle \mathbf{Z}, \mathbf{I} \rangle \\ \text{s.t., } \forall d \in [D], \langle \mathbf{Z}, \mathbf{A}_d^{\text{real}} \rangle = 2 \operatorname{Re}(\widehat{\mathbf{w}}[d]) \\ \forall d \in [D], \langle \mathbf{Z}, \mathbf{A}_d^{\text{img}} \rangle = 2 \operatorname{Im}(\widehat{\mathbf{w}}[d]). \end{aligned} \quad (\text{SDP})$$

The outline of the proof is detailed below:

1. In A.1, we look at the KKT conditions of (SDP) and show that Theorem 6 follows from Lemma 7, which states that for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, there exists $\mathbf{c} \in \mathbb{R}^K$ such that $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b} = \mathbf{c} \star \mathbf{c}$, where the convolutions are w.r.t. D .
2. In A.2 we provide the proof of Lemma 7, which is further subdivided as follows:
 - (a) A.2.1 reduces Lemma 7 to the special case where $D = 2K - 1$.
 - (b) A.2.2 introduces a polynomial representation of convolutions which leads to a reformulation of Lemma 7 with $D = 2K - 1$.
 - (c) A.2.3 contains the core argument for existence of the desired $\mathbf{c} \in \mathbb{R}^K$ via reasoning about the roots of polynomial reformulation.

A.1. KKT conditions for (SDP)

The D constraints involving $\operatorname{Re}(\mathbf{w})$ correspond to a dual vector $\boldsymbol{\lambda}^{\text{real}} \in \mathbb{R}^D$; the D constraints involving $\operatorname{Im}(\mathbf{w})$ correspond to a dual vector $\boldsymbol{\lambda}^{\text{img}} \in \mathbb{R}^D$. To simplify these conditions, we take $\boldsymbol{\lambda} \in \mathbb{C}^D$ to be $\boldsymbol{\lambda}^{\text{real}} + i \cdot \boldsymbol{\lambda}^{\text{img}}$ (when $\boldsymbol{\lambda}$ is a dual-optimal solution, $\boldsymbol{\lambda}$ is also a Fourier transform of a real vector). The dual variable for the PSD constraint corresponds to a matrix $\boldsymbol{\Gamma} \succcurlyeq 0$. In this notation, the KKT conditions are primal feasibility, along with the following constraints:

$$\begin{aligned} \boldsymbol{\Gamma} = \mathbf{I} - \begin{bmatrix} \mathbf{0}_K & \mathbf{F}_K^\top \boldsymbol{\Lambda} \overline{\mathbf{F}} \\ \mathbf{F} \overline{\boldsymbol{\Lambda}} \mathbf{F}_K & \mathbf{0}_D \end{bmatrix} \succcurlyeq 0 \\ \mathbf{Z} \overline{\boldsymbol{\Gamma}} = 0. \end{aligned}$$

Now, to simplify these conditions, suppose that \mathbf{Z} is rank L , in which case we can express it as $\mathbf{Z} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{U}^\top & \mathbf{V}^\top \end{bmatrix}$, where $\mathbf{U} \in \mathbb{R}^{K \times L}$ and $\mathbf{V} \in \mathbb{R}^{D \times L}$. Using that $\begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix}$ is full rank along with some algebraic manipulations, we obtain the following formulation of the KKT conditions:

$$\sum_{l=0}^L \widehat{\mathbf{U}}[:, l] \odot \widehat{\mathbf{V}}[:, l] = \widehat{\mathbf{w}} \quad (\text{KKT 1})$$

$$\begin{bmatrix} \mathbf{0}_K & \mathbf{F}_K^\top \boldsymbol{\Lambda} \overline{\mathbf{F}} \\ \mathbf{F} \overline{\boldsymbol{\Lambda}} \mathbf{F}_K & \mathbf{0}_D \end{bmatrix} \preccurlyeq \mathbf{I}_{D+K} \quad (\text{KKT 2})$$

$$\widehat{\mathbf{V}} = \overline{\boldsymbol{\Lambda}} \widehat{\mathbf{U}} \quad (\text{KKT 3})$$

$$\widehat{\mathbf{U}} = \mathbf{F}_K \overline{\mathbf{F}}_K^\top \boldsymbol{\Lambda} \widehat{\mathbf{V}}. \quad (\text{KKT 4})$$

The KKT conditions give useful properties of the solution \mathbf{Z} . For $0 \leq l \leq L-1$, we let $\mathbf{u}_l = \mathbf{U}[:, l]$ and $\mathbf{v}_l = \mathbf{V}[:, l]$. From (KKT 3), we see that

$$\forall_{l \in [L]}, \hat{\mathbf{v}}_l = \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}_l.$$

Combining the above with (KKT 1), we obtain:

$$\hat{\mathbf{w}} = \sum_{l=0}^{L-1} \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}_l \odot \hat{\mathbf{u}}_l.$$

We now make the following assertion.

Claim *We claim that in order to prove Theorem 6, it suffices to find a vector $\mathbf{u} \in \mathbb{R}^K$ such that:*

$$\overline{\hat{\mathbf{u}}} \odot \hat{\mathbf{u}} = \sum_{l=0}^{L-1} \overline{\hat{\mathbf{u}}}_l \odot \hat{\mathbf{u}}_l. \quad (\text{CONV-REDUCTION})$$

The most technical part of the proof is to show (CONV-REDUCTION) which we show follows from Lemma 7. But before that we will first justify our above claim.

Proof [Proof of claim] Assume (CONV-REDUCTION) holds for $\mathbf{u} \in \mathbb{R}^K$. Let $\mathbf{v} = \mathbf{F}^{-1}(\boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}})$ such that we have $\hat{\mathbf{v}} = \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}$. We can take \mathbf{Z}^* to be $\begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \begin{bmatrix} \mathbf{u}^\top & \mathbf{v}^\top \end{bmatrix}$.

We can see that (\mathbf{u}, \mathbf{v}) satisfies the following:

$$\hat{\mathbf{u}} \odot \hat{\mathbf{v}} = \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}} \odot \hat{\mathbf{u}} = \sum_{l=0}^{L-1} \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}_l \odot \hat{\mathbf{u}}_l = \hat{\mathbf{w}}. \quad (15)$$

Thus, \mathbf{Z}^* satisfies the feasibility condition for $\mathcal{R}_K^{\text{SDP}}(\mathbf{w})$. Moreover, we show that the solution also achieves the optimum objective value for $\mathcal{R}_{K,C}(\mathbf{w})$ as follows:

$$\begin{aligned} \langle \mathbf{Z}^*, \mathbf{I} \rangle &= \|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2 = \|\hat{\mathbf{u}}\|_2^2 + \|\hat{\mathbf{v}}\|_2^2 \\ &= \sum_{d=0}^{D-1} |\hat{\mathbf{u}}_d|^2 + \sum_{d=0}^{D-1} |\hat{\mathbf{v}}_d|^2 \\ &\stackrel{(a)}{=} \sum_{d=0}^{D-1} \sum_{l=0}^{L-1} |(\hat{\mathbf{u}}_l)_d|^2 + \sum_{d=0}^{D-1} \sum_{l=0}^{L-1} |\boldsymbol{\lambda}_d|^2 |(\hat{\mathbf{u}}_l)_d|^2 \\ &\stackrel{(b)}{=} \sum_{l=0}^{L-1} \sum_{d=0}^{D-1} |(\hat{\mathbf{u}}_l)_d|^2 + \sum_{l=0}^{L-1} \sum_{d=0}^{D-1} |(\hat{\mathbf{v}}_l)_d|^2 \\ &= \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 \\ &= \langle \mathbf{Z}, \mathbf{I} \rangle = \mathcal{R}_K^{\text{SDP}}(\mathbf{w}) \end{aligned} \quad (16)$$

where (a) follows from (CONV-REDUCTION) and using $\hat{\mathbf{v}} = \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}$, which follows by definition of \mathbf{v} , and (b) follows from (KKT 3) that $\hat{\mathbf{v}}_l = \boldsymbol{\lambda} \odot \overline{\hat{\mathbf{u}}}_l$.

Thus, we have show that (CONV-REDUCTION) implies that there exists a rank 1 solution \mathbf{Z}^* that achieves the same objective value as \mathbf{Z} and satisfies the constraints, and hence is minimizer of the SDP as desired. \blacksquare

Proof of (CONV-REDUCTION) from Lemma 7 It is more convenient to write (CONV-REDUCTION) in signal space. Taking inverse Fourier transforms of (CONV-REDUCTION), we need to show that given $\mathbf{u}_l = \mathbf{F}_K^\top \hat{\mathbf{u}}_l \in \mathbb{R}^K$ there exists $\mathbf{u} \in \mathbb{R}^K$ such that the following holds:

$$\mathbf{u} \star \mathbf{u} = \sum_{l=0}^{L-1} \mathbf{u}_l \star \mathbf{u}_l, \quad (17)$$

where the convolutions are taken in D dimensional space, i.e., $\{\mathbf{u}_l\}_l$ are padded with $D - K$ zeros so that $\mathbf{u}_l \star \mathbf{u}_l \in \mathbb{R}^D$.

We can now see that 17 (and hence (CONV-REDUCTION)) indeed holds by recursively applying Lemma 7, which was stated earlier in the main text and is reiterated below.

Lemma 7 *For any $1 \leq K \leq D$, and for any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, there exists a vector $\mathbf{c} \in \mathbb{R}^K$ such that $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b} = \mathbf{c} \star \mathbf{c}$, where convolutions are w.r.t. dimension D .*

We have thus far shown that Theorem 6 follows from Lemma 7: we first established that it suffices to find $\mathbf{u} \in \mathbb{R}^K$ satisfying (17) (or equivalently (CONV-REDUCTION)), which in turn holds from recursively applying Lemma 7. Now, it suffices to prove Lemma 7; we prove this lemma in the following subsections.

A.2. Proof of Lemma 7: the convolutional property

For $K = D$, Lemma 7 follows easily from the Fourier space representation, since the Fourier space representation of the vector \mathbf{c} can be explicitly constructed as the square root of the Fourier representation of $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b}$. However, this construction does not generalize to kernel sizes $K < D$, and Lemma 7 is thus non-trivial in general. In fact, the vector \mathbf{c} does not even appear to have a clean closed-form characterization for general K . To sidestep this issue, we use a proof technique that enables us to implicitly construct the vector \mathbf{c} . We believe that this proof technique could be of independent interest.

A.2.1. REDUCING LEMMA 7 TO $D = 2K - 1$ CASE

The first step in the proof of Lemma 7 is to show that Lemma 7 for general K, D follows from the special case where $D = 2K - 1$.

Lemma 16 *For any $K \geq 1$, for $D = 2K - 1$, and for any vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, there exists a vector $\mathbf{c} \in \mathbb{R}^K$ such that $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b} = \mathbf{c} \star \mathbf{c}$, where convolutions are taken in dimension D .*

Proof of Lemma 7 from Lemma 16 By Lemma 16, we know that Lemma 7 holds when $D = 2K - 1$. We now show that this implies the statement for a general value of D . Exclusively in this proof, let \star_D denote the convolutional operator w.r.t dimension D . For $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, let $\mathbf{c} \in \mathbb{R}^K$ be the vector from Lemma 16 such that $\mathbf{a} \star_{2K-1} \mathbf{a} + \mathbf{b} \star_{2K-1} \mathbf{b} = \mathbf{c} \star_{2K-1} \mathbf{c}$. We will now show that the same \mathbf{c} also satisfies $\mathbf{a} \star_D \mathbf{a} + \mathbf{b} \star_D \mathbf{b} = \mathbf{c} \star_D \mathbf{c}$ for all D .

Case 1: $D \geq 2K - 1$. First, notice that for K dimensional vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$ and $\mathbf{c} \in \mathbb{R}^K$, their convolutions for $D \geq 2K - 1$ will have at most $2K - 1$ nonzero entries, located at indices $[-K + 1, K - 1] \bmod D$. Thus, for $d' \notin [-K + 1, K - 1]$

$$(\mathbf{a} \star_D \mathbf{a} + \mathbf{b} \star_D \mathbf{b})_{d' \bmod D} = 0 = (\mathbf{c} \star_D \mathbf{c})_{d' \bmod D}. \quad (18)$$

At the same time, for $d \in [-K + 1, K - 1]$, we have that:

$$\begin{aligned} (\mathbf{a} \star_D \mathbf{a} + \mathbf{b} \star_D \mathbf{b})_{d \bmod D} &= (\mathbf{a} \star_{2K-1} \mathbf{a} + \mathbf{b} \star_{2K-1} \mathbf{b})_{d \bmod (2K-1)} \\ &= (\mathbf{c} \star_{2K-1} \mathbf{c})_{d \bmod (2K-1)} \\ &= (\mathbf{c} \star_{2K-1} \mathbf{c})_{d \bmod (D)}. \end{aligned} \quad (19)$$

Case 2: $D < 2K - 1$. Let $K \leq D < 2K - 1$. For any vector $\mathbf{z} \in \mathbb{R}^K$, we can check the following:

$$\forall_{d \in [D]}, (\mathbf{z} \star_D \mathbf{z})_d = \sum_{\substack{0 \leq d' < 2K-1 \\ \text{s.t., } d \equiv d' \bmod D}} (\mathbf{z} \star_{2K-1} \mathbf{z})_{d'}. \quad (20)$$

Thus, for $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$ and $\mathbf{c} \in \mathbb{R}^K$ satisfying Lemma 16, we get the desired result as follows

$$\begin{aligned} \forall_{d \in [D]}, (\mathbf{a} \star_D \mathbf{a} + \mathbf{b} \star_D \mathbf{b})_d &= \sum_{\substack{0 \leq d' < 2K-1 \\ \text{s.t., } d \equiv d' \bmod D}} (\mathbf{a} \star_{2K-1} \mathbf{a} + \mathbf{b} \star_{2K-1} \mathbf{b})_{d'} \\ &= \sum_{\substack{0 \leq d' < 2K-1 \\ \text{s.t., } d \equiv d' \bmod D}} (\mathbf{c} \star_{2K-1} \mathbf{c})_{d'} \\ &= (\mathbf{c} \star_D \mathbf{c})_d. \end{aligned} \quad (21)$$

This concludes the proof of showing Lemma 7 from Lemma 16. \square

The remainder of the section of devoting to proving Lemma 16. This statement trivially holds with \mathbf{c} as the zero vector if $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b} = 0$, so for the remainder of the proof, we assume that $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b}$ is nonzero.

A.2.2. INTRODUCING THE POLYNOMIAL REPRESENTATION

The key idea for the proof of Lemma 16 is to use the polynomial formulation of convolutions.⁵ We then use factorization of the polynomials to implicitly construct \mathbf{c} . We use the following notation. Let $\mathcal{P}_k \subseteq \mathbb{R}[x]$ denote the set of degree $\leq k - 1$ polynomials with *real coefficients*. For a vector $\mathbf{z} \in \mathbb{R}^K$, we define the *polynomial representation* $p_{\mathbf{z}}(x) \in \mathcal{P}_{K-1}$ to be the polynomial $\mathbf{z}[0] + \mathbf{z}[1]x + \dots + \mathbf{z}[K-1]x^{K-1}$. Using the polynomial representations, convolutions can be expressed as polynomial multiplication:

Fact 1 *Let $\mathbf{a} \in \mathbb{R}^K$. The $D = 2K - 1$ dimensional convolution $\mathbf{a} \star \mathbf{a}$ has polynomial representation $p_{\mathbf{a} \star \mathbf{a}}(x)$ that is equivalent to the polynomial $x^{K-1} p_{\mathbf{a}}(x) \cdot p_{\mathbf{a}}(1/x) \in \mathcal{P}_{2K-2}$ up to permuting the coefficients appropriately.*

5. E.g., see <https://en.wikipedia.org/wiki/Convolution>

The polynomial representation enables us to construct a vector \mathbf{c} in terms of the roots of the relevant polynomials. We now reformulate Lemma 16 using the polynomial representation. Recall that we wish to show that there exists a vector $\mathbf{c} \in \mathbb{R}^K$ such that $\mathbf{c} \star \mathbf{c} = \mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b}$.

In the polynomial representation, we equivalently want $p_{\mathbf{c} \star \mathbf{c}}(x) = p_{\mathbf{a} \star \mathbf{a}}(x) + p_{\mathbf{b} \star \mathbf{b}}(x)$. Now, applying Fact 1, we see that the polynomial formulation of the Lemma 16 is the following:

$$\exists \mathbf{c} \in \mathbb{R}^K \text{ s.t., } x^{K-1}p_{\mathbf{c}}(x) \cdot p_{\mathbf{c}}(1/x) = x^{K-1}p_{\mathbf{a}}(x)p_{\mathbf{a}}(1/x) + x^{K-1}p_{\mathbf{b}}(x)p_{\mathbf{b}}(1/x).$$

To simplify notation, we denote the right-hand-side of the previous equation by $Q(x)$:

$$Q(x) := x^{K-1}p_{\mathbf{a}}(x)p_{\mathbf{a}}(1/x) + x^{K-1}p_{\mathbf{b}}(x)p_{\mathbf{b}}(1/x) \in \mathcal{P}_{2K-2}. \quad (22)$$

Since there is a 1-to-1 correspondence between polynomials in \mathcal{P}_{K-1} and vectors in \mathbb{R}^K . In this notation, our goal is to show that there exists a polynomial $p \in \mathcal{P}_{K-1}$ such that:

$$\text{(to show)} \quad Q(x) = x^{K-1}p(x)p(1/x). \quad (23)$$

The coefficients of such a p would then give us the desired $\mathbf{c} \in \mathbb{R}^K$ in Lemma 16. The remainder of the proof boils down to constructing $p \in \mathcal{P}_{K-1}$ such that (23) is satisfied.

A.2.3. PROVING THE POLYNOMIAL REPRESENTATION VERSION OF THE LEMMA STATEMENT

The first property of $Q(x)$ that we leverage is that $Q(x)$ is a *palindromic polynomial* of order $2K - 2$ (i.e., its coefficients zero padded to degree $2K - 2$ forms a palindrome) with real coefficients. To see that our $Q(x)$ is palindromic, note any polynomial $Q(x) \in \mathcal{P}_{2K-2}$ is a palindromic polynomial if and only if $Q(x) = x^{2K-2}Q(1/x)$, which is satisfied by our definition in (22).

At first glance, it would appear that (23) follows immediately from standard properties of palindromic polynomials with real coefficients whose complex roots are known to come in reciprocal pairs as $(\alpha, 1/\alpha)$.⁶ However, we cannot obtain eq. (23) from the palindromic property alone. To see this, consider the following example:

Example 1 Consider the palindromic polynomial with real coefficients $x^2 + 1 = ix(x - i)(1/x - i)$. This polynomial is not expressible as $xp'(x)p'(1/x)$ for any *real* polynomial p' .

The proof of (23) thus must leverage further structure of $Q(x)$, which we will ultimately extract through examining the roots of $Q(x)$. Using that \mathbb{C} is algebraically closed, we can factor $Q(x)$ into a polynomial $c_Q \prod_i (x - \alpha_i)$ with exactly $2K - 2$ roots where the α_i need not be distinct. To show (23), it suffices to show that $x^{K-1}p(x) \cdot p(1/x)$ has roots (with multiplicities) given by the multi-set $S_Q = \{\alpha_i\}$ and has the same leading coefficient c_Q . Drawing upon this formulation, we will construct p implicitly by the multi-set S_p of its roots (with multiplicities) and its leading coefficient $c_p \neq 0$: that is, so that $p = c_p \prod_{\alpha \in S_p} (x - \alpha)$.

The remainder of the proof is structured as follows:

- Step 1.** We establish key properties of the multi-set S_Q of $2K - 1$ complex roots of $Q(x)$.
- Step 2.** We then construct the roots S_p of p from exactly half of elements in S_Q .
- Step 3.** We show that S_p is the “right” set of roots in the sense that with any nonzero choice of leading coefficient c_p , the resulting polynomial p is of degree at most $K - 1$, and the multi-set of roots of $x^{K-1}p(x)p(1/x)$ match S_Q .
- Step 4.** Finally, we choose a real leading coefficient c_p of $p(x)$ to get the desired property in (23).

6. This is a standard fact: e.g., see https://en.wikipedia.org/wiki/Reciprocal_polynomial.

Step 1. Properties of the complex roots of $Q(x)$

Before we construct S_p and c_p , it is helpful to establish properties of the multi-set of roots S_Q .

1. (P1) For every root α with multiplicity m , $\bar{\alpha}$ is a root and has multiplicity m .
2. (P2) If $\alpha \neq 0$ is a root with multiplicity m , then $1/\alpha$ is a distinct root with multiplicity m .
3. (P3) If α such that $|\alpha| = 1$ is a root, then α has even multiplicity.

The first two properties (P1) and (P2) follow from the fact that $Q(x)$ is a palindromic polynomial with real coefficients. In particular, we see that (P1) follows from the fact that $Q(x)$ has real coefficients so that the roots come in conjugate pairs, and (P2) follows from standard properties of palindromic polynomials.⁷

The last property, (P3), uses deeper aspects of the structure of $Q(x)$. In particular, it uses that $Q(x)$ is the sum of polynomials of the form $x^{K-1}p'(x)p'(1/x)$ where $p' \in \mathcal{P}_{K-1}$ has real coefficients, rather than just an arbitrary palindromic polynomial. To see this, let's return to Example 1 and observe that $x^2 + 1$ does not satisfy (P3) since its roots are i and $-i$, each with odd-multiplicity of 1. Hence, we use further structure of $Q(x)$ and we show:

Lemma 17 *Consider vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, and let $p_{\mathbf{a}}, p_{\mathbf{b}} \in \mathcal{P}_{K-1}$ be their polynomial representation. If $\alpha \in \mathbb{C}$ such that $|\alpha| = 1$ is a root of $p_{\mathbf{a}}(x)(x^{K-1}p_{\mathbf{a}}(1/x)) + p_{\mathbf{b}}(x)(x^{K-1}p_{\mathbf{b}}(1/x))$, then α has even multiplicity.*

Proof Suppose that α is a root of $Q(x)$ and $|\alpha| = 1$. We see that

$$\begin{aligned}
 0 &= Q(\alpha) \\
 &= p_{\mathbf{a}}(\alpha)(\alpha^{K-1}p_{\mathbf{a}}(1/\alpha)) + p_{\mathbf{b}}(\alpha)(\alpha^{K-1}p_{\mathbf{b}}(1/\alpha)) \\
 &= p_{\mathbf{a}}(\alpha)(\alpha^{K-1}p_{\mathbf{a}}(\bar{\alpha})) + p_{\mathbf{b}}(\alpha)(\alpha^{K-1}p_{\mathbf{b}}(\bar{\alpha})) \\
 &= \alpha^{K-1} \left(p_{\mathbf{a}}(\alpha)\overline{p_{\mathbf{a}}(\alpha)} + p_{\mathbf{b}}(\alpha)\overline{p_{\mathbf{b}}(\alpha)} \right) \\
 &= \alpha^{K-1} (|p_{\mathbf{a}}(\alpha)|^2 + |p_{\mathbf{b}}(\alpha)|^2).
 \end{aligned}$$

Since $\alpha \neq 0$, this means that $|p_{\mathbf{a}}(\alpha)|^2 + |p_{\mathbf{b}}(\alpha)|^2 = 0$. Thus, $p_{\mathbf{a}}(\alpha) = 0$ and $p_{\mathbf{b}}(\alpha) = 0$. Now, it suffices to show that α is a root with even multiplicity in $p_{\mathbf{a}}(x)(x^{K-1}p_{\mathbf{a}}(1/x))$ and in $p_{\mathbf{b}}(x)(x^{K-1}p_{\mathbf{b}}(1/x))$.

We show that α has even multiplicity in $p_{\mathbf{a}}(x)(x^{K-1}p_{\mathbf{a}}(1/x))$ (an analogous argument shows this for $p_{\mathbf{b}}(x)(x^{K-1}p_{\mathbf{b}}(1/x))$). Suppose that α has multiplicity m in $p_{\mathbf{a}}(x)$. Since $p_{\mathbf{a}}(x)$ has real coefficients, we know that $\bar{\alpha}$ is a root of $p_{\mathbf{a}}(x)$ with multiplicity m . We also know that $1/\alpha = \bar{\alpha}$ is a root with multiplicity m of $x^{K-1}p_{\mathbf{a}}(1/x)$. Since $x^{K-1}p_{\mathbf{a}}(1/x)$ has real coefficients, we know that $\bar{\bar{\alpha}} = \alpha$ is a root with multiplicity m of $x^{K-1}p_{\mathbf{a}}(1/x)$. This means that α has multiplicity $2m$ in $p_{\mathbf{a}}(x)(x^{K-1}p_{\mathbf{a}}(1/x))$ as desired. \blacksquare

We note that (P3) follows immediately from Lemma 17.

⁷ This is a standard fact: e.g., see https://en.wikipedia.org/wiki/Reciprocal_polynomial.

Step 2: Constructing the roots of S_p

We now construct the multi-set of roots S_p . To do this, we begin by constructing the *nonzero* roots in S_p , and then we add in the zero roots with the appropriate multiplicities at the end.

Constructing the nonzero roots. The high-level intuition for the construction is that we ultimately need the nonzero roots of $x^{K-1}p(x)p(1/x)$ for a degree $K - 1$ polynomial $p \in \mathcal{P}_{K-1}$ to exactly match the nonzero roots in S_Q . Since the nonzero roots of $x^{K-1}p(1/x)$ are the reciprocals of the roots of $p(x)$ (with multiplicity preserved), and since polynomials with real coefficients have roots in conjugate pairs, we wish to divide the real nonzero roots of $Q(x)$ into disjoint pairs $(\alpha, 1/\alpha)$ (so that $p(x)$ has root α and $x^{K-1}p(1/x)$ has root $1/\alpha$) and the complex roots into disjoint quadruples $(\alpha, \bar{\alpha}, 1/\alpha, 1/\bar{\alpha})$ (so that $p(x)$ has roots α and $\bar{\alpha}$ and $x^{K-1}p(1/x)$ has roots $1/\alpha$ and $1/\bar{\alpha}$).

Let's formalize this high level argument by constructing an undirected graph G where the vertices are the nonzero roots in S_Q (a root with multiplicity m corresponds to m separate vertices). We will now add edges such that G forms a bipartite graph with following properties : (a) each edge is of form $(\alpha, 1/\alpha)$ without any self-loop; and (b) each vertex is connected to one and exactly one other vertex (*i.e.*, all edges are disjoint).

The edges are defined as follows. When $\alpha \neq \pm 1$ is a nonzero root of $Q(x)$, we have $1/\alpha \neq \alpha$ and it follows from (P2) that α and $1/\alpha$ will have same multiplicity, say m . Thus, we can create m disjoint edges by connecting each vertex with value α to exactly one vertex with value $1/\alpha$. This ensures that the subgraph of G with values α or $1/\alpha$ and respects our desired construction. Using (P3), we can also handle roots $\alpha = \pm 1$ in a similar construction. In this case, we will have even number of vertices with value α and hence can form non-self-loop edges of the form (α, α) such that all the edges are disjoint.

Having constructed G as above, let G^{real} be the subgraph of G consisting of vertices corresponding to real roots, and let G^{complex} be the subgraph of G consisting of vertices corresponding to roots with nonzero imaginary part. It is easy to see that $G = G^{\text{real}} \cup G^{\text{complex}}$ and these graphs are disjoint.

Let us now use these graphs to construct the set of nonzero roots in S_p , which will contain half of the vertices in G . For edges of the form $(\alpha, 1/\alpha)$ in G^{real} , we add one vertex from each edge in G^{real} to S_p . For G^{complex} , we first pair up the edges in the subgraph G^{complex} as follows. When $|\alpha| \neq 1$, we can pair up the edges $(\alpha, 1/\alpha)$ and $(\bar{\alpha}, 1/\bar{\alpha})$ so that the pairs are disjoint (follows from (P1)) and further no edge is paired with itself (since $1/\alpha \neq \bar{\alpha}$ when $|\alpha| \neq 1$). When $|\alpha| = 1$, we have $1/\alpha = \bar{\alpha}$. Here we use the fact that α and $\bar{\alpha}$ have same even multiplicity (see (P3)), say $2m$. Thus, we can again pair up m edges $(\alpha, \bar{\alpha}) \equiv (\alpha, 1/\alpha)$ with m distinct edges $(\alpha, \bar{\alpha}) \equiv (\bar{\alpha}, 1/\bar{\alpha})$ in such a way that all the pairs continue to be disjoint and no edge is paired with itself. We have thus paired all edges in G^{complex} as $(\alpha, 1/\alpha)$ and $(\bar{\alpha}, 1/\bar{\alpha})$ such that no edge is paired with itself and all pairings are disjoint. Now, from each distinct pair of edges $(\alpha, 1/\alpha)$ and $(\bar{\alpha}, 1/\bar{\alpha})$, we add α from first edge and $\bar{\alpha}$ from second edge to S_p .

Adding the zero roots. To construct S_p , all that remains is to determine the multiplicity of the zero roots. Let m be the multiplicity of 0 in S_Q . We then simply add m copies of 0 to S_p .

In summary, all the real nonzero roots of S_Q were partitioned into disjoint pairs $(\alpha, 1/\alpha)$ (with α from each pair included in S_p) and the complex nonzero roots of S_Q were partitioned into disjoint quadruples $(\alpha, \bar{\alpha}, 1/\alpha, 1/\bar{\alpha})$ (with $\alpha, \bar{\alpha}$ from each quadruple included in S_p). We then included all zero entries of S_Q in S_p .

Step 3. Proof that S_p is the correct multi-set of roots

We will derive the leading coefficient c_p of our desired polynomial shortly, but first show that S_p is the right multi-set of roots for p . We first prove that S_p corresponds to the roots of a degree $K - 1$ polynomial with real coefficients; we then show that with any nonzero leading coefficient c_p , the resulting polynomial $p(x)$ is such that the multiset of roots of $x^{K-1}p(x)p(1/x)$ is equal to S_Q .

Proof that S_p corresponds to the roots of a polynomial in \mathcal{P}_{K-1} . For S_p to be a valid multi-set of roots for a polynomial $p \in \mathcal{P}_{K-1}$, we need to ensure that S_p consists of at most $K - 1$ elements and that the complex roots come in conjugate pairs.

The fact that the complex roots (with nonzero imaginary components) come in conjugate pairs follows from the construction of the graphs above. Recall the complex roots of S_Q were partitioned into disjoint quadruples $(\alpha, \bar{\alpha}, 1/\alpha, 1/\bar{\alpha})$ so that S_p has roots α and $\bar{\alpha}$ from each quadruple. This ensures that the roots comes up in conjugate pairs as desired.

We now show that S_p consists of exactly $K - 1$ elements using the following root counting argument. Recall that we defined $Q(x) = x^{K-1}p_a(x)p_a(1/x) + x^{K-1}p_b(x)p_b(1/x)$ to be a palindromic polynomial of order $2K - 2$, i.e., for $d = 0, 1, \dots, 2K - 2$, coefficient of x^d is the same as coefficient of x^{2K-2-d} . Thus, if $Q(x)$ has a zero root with multiplicity m , i.e., the coefficients of $x^0, x^1, x^2, \dots, x^{m-1}$ are zero, then the palindromic property would ensure that the coefficients of $x^{2K-2}, x^{2K-3}, \dots, x^{2K-1-m}$ are also zero. Which in turn implies that the max degree of $Q(x)$ is of degree $2K - 2 - m$. Thus, we see that $Q(x)$ has $2K - 2 - m$ roots including multiplicity, of which m are zero roots and $2K - 2 - 2m$ nonzero roots. By the construction of S_p , we included exactly half of nonzero roots of S_Q and exactly m zero roots in S_p . Thus S_p has $K - 1 - m$ nonzero entries and exactly m zero entries leading to a total of $K - 1$ roots.

Proof that the multi-set of roots of $x^{K-1}p(x)p(1/x)$ equals the multi-set S_Q . We denote the monic polynomial given by the roots of S_p as $p_{\text{monic}}(x) = \prod_{\alpha \in S_p} (x - \alpha)$ and consider

$$Q_1(x) := p_{\text{monic}}(x)x^{K-1}p_{\text{monic}}(1/x).$$

It suffices to show that the multi-set of roots of $Q_1(x)$ with multiplicities is equal to the multi-set S_Q .

For the nonzero roots, we use the construction of S_p . Recall that the real roots of S_Q were partitioned into disjoint pairs $(\alpha, 1/\alpha)$ (so that $p_{\text{monic}}(x)$ has root α) and the complex roots of S_Q were partitioned into disjoint quadruples $(\alpha, \bar{\alpha}, 1/\alpha, 1/\bar{\alpha})$ (so that $p_{\text{monic}}(x)$ has roots α and $\bar{\alpha}$). This, coupled with the fact that the nonzero roots of $x^{K-1}p_{\text{monic}}(1/x)$ are the inverses of the nonzero roots of $p_{\text{monic}}(x)$, means that the multi-set of nonzero roots of $Q_1(x)$ which is the union of the multi-set of nonzero roots in $p_{\text{monic}}(x)$ with the multi-set of nonzero roots of $x^{K-1}p_{\text{monic}}(1/x)$ is exactly equal to the multi-set of nonzero roots of $Q(x)$.

For the zero roots, we simply need to show that $x^{K-1}p_{\text{monic}}(1/x)$ has no roots that are 0. Using the fact that the constant term of $x^{K-1}p_{\text{monic}}(1/x)$ is equal to the coefficient of x^{K-1} in $p_{\text{monic}}(x)$, it suffices to show that the coefficient of x^{K-1} in $p_{\text{monic}}(x)$ is nonzero. The latter condition follows from the argument in the previous paragraph where we showed S_p to have exactly $K - 1$ roots (with multiplicity) and hence the degree of $p_{\text{monic}}(x)$ must be $K - 1$ with a nonzero coefficient of x^{K-1} . We are now done as the zero roots of $Q(x)$ matches those of $p_{\text{monic}}(x)$ with $x^{K-1}p_{\text{monic}}(1/x)$ adding no additional zero roots.

This concludes the argument that the roots of $Q_1(x)$ with multiplicities matches those of $Q(x)$.

Step 4. Constructing the leading coefficient of p

Now, we need to just construct the leading coefficient of p . As above, let $p_{\text{monic}}(x) = \prod_{\alpha \in S_p} (x - \alpha)$ be the monic polynomial given by the roots of S_p , and consider $Q_1(x) = p_{\text{monic}}(x)x^{K-1}p_{\text{monic}}(1/x)$. Since $Q_1(x)$ and $Q(x)$ have the same set of roots with multiplicities, we know that $Q(x) = \gamma \cdot Q_1(x)$ for some $\gamma \neq 0$. Let's take $c_p = \sqrt{\gamma}$. With this choice of c_p , we define $p(x) = c_p \prod_{\alpha \in S_p} (x - \alpha)$ and have established by construction that $Q(x) = x^{K-1}p(x)p(1/x)$. The final component is to show that c_p is real so that $p(x) = c_p \prod_{\alpha \in S_p} (x - \alpha)$ is a polynomial with real coefficients.

To show that c_p is real, it suffices to show that γ is positive. We show this as follows. Let $\mathbf{c}_{\text{monic}} \in \mathbb{R}^K$ denote the vector corresponding to the polynomial representation $p_{\text{monic}}(x)$ such that $Q_1(x) = p_{\text{monic}}(x)x^{K-1}p_{\text{monic}}(1/x)$ is equivalent to the polynomial representation of $\mathbf{c}_{\text{monic}} \star \mathbf{c}_{\text{monic}}$ (using Fact 1). Recall that $Q(x)$ is analogously equivalent to the polynomial representation of $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b}$. Further, note that the component at index 0 of self-convolution operations satisfy, $(\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b})_0 = \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2$ and $(\mathbf{c}_{\text{monic}} \star \mathbf{c}_{\text{monic}})_0 = \|\mathbf{c}_{\text{monic}}\|^2$. One can check that, these 0-index components in turn appear as the coefficients of x^{K-1} in $Q(x)$ and $Q_1(x)$, respectively. This along with $Q(x) = \gamma Q_1(x)$ implies that $\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 = \gamma \|\mathbf{c}_{\text{monic}}\|^2$. Since the LHS is strictly positive (as $Q(x) \neq 0$ without loss of generality), we must have $\gamma > 0$ as desired.

Concluding Lemma 16. We have shown that $p(x) = c_p \prod_{\alpha \in S_p} (x - \alpha)$ is a polynomial of degree $K - 1$ with real coefficients. We have shown that $Q(x)$ and $x^{K-1}p(x) \cdot p(1/x)$ have the same multi-set of roots and the same leading coefficients, thus proving the polynomial formulation of Lemma 16 in (23). This concludes the proof of Lemma 16.

A.3. Discussion of the proof technique

We conclude with a discussion of the analysis and highlight the main parts of the proof. At the beginning of the section, we used the KKT conditions to show that it suffices to prove **CONV-REDUCTION**, an additive property about convolutions of for kernel size. We then showed that it suffices to prove a version of this statement for the sum of two such convolutions, i.e. Lemma 7. We believe that this property could be of independent interest.

The bulk of the proof boils down to proving Lemma 7 in the special case of $D = 2K - 1$, i.e., Lemma 16. Proving Lemma 16 was the core technical contribution in this section. Since \mathbf{c} does not necessarily always a clean closed-form solution as a function of \mathbf{a} and \mathbf{b} , we needed to construct \mathbf{c} *implicitly*. The polynomial representation of convolutions enabled us to implicitly construct \mathbf{c} via its roots. To construct \mathbf{c} and ensure that the corresponding polynomial representation p had real coefficients, we needed to leverage the structure of the polynomial representation $Q(x)$ of $\mathbf{a} \star \mathbf{a} + \mathbf{b} \star \mathbf{b}$ beyond its palindromic structure. (This additional property was proven in Lemma 17) With this structure, we can factor $Q(x)$ and partition its roots in order to construct the roots of p .

Appendix B. Proofs for Section 2: Induced regularizer in special cases

For a single channel convolutional network (i.e., $C = 1$), we denote the weights in the first and second layer as $\mathbf{U} \in \mathbb{R}^K$ and $\mathbf{V} \in \mathbb{R}^D$, respectively. Recall that the D dimensional discrete Fourier transform of the weights \mathbf{U}, \mathbf{V} and the linear predictor $\mathbf{w} \in \mathbb{R}^D$ are denoted as $\widehat{\mathbf{U}} = \mathbf{F}_K \mathbf{U}$, $\widehat{\mathbf{V}} = \mathbf{FV}$, $\widehat{\mathbf{w}} = \mathbf{Fw}$, respectively. Moreover, the Fourier transform is normalized to be unitary such that for any $\mathbf{z} \in \mathbb{R}^K$, $\|\mathbf{z}\| = \|\widehat{\mathbf{z}}\|$ and $\mathbf{F}\mathbf{F}^\top = \overline{\mathbf{F}}\mathbf{F}^\top = \mathbf{I}$. Thus, for all the quantities of interest, we use

the ℓ_2 norm in signal domain interchangeably with ℓ_2 norm in the Fourier domain, e.g., $\|\mathbf{U}\| = \|\widehat{\mathbf{U}}\|$, $\|\mathbf{V}\| = \|\widehat{\mathbf{V}}\|$, and $\|\mathbf{w}\| = \|\widehat{\mathbf{w}}\|$.

In the following proofs, we use the formulation of $\mathcal{R}_{K,C}(\mathbf{w})$ in eq. (6) for the case of $C = 1$ as:

$$\mathcal{R}_{K,1}(\mathbf{w}) = \min_{\mathbf{U} \in \mathbb{R}^K, \mathbf{V} \in \mathbb{R}^D} \|\widehat{\mathbf{U}}\|^2 + \|\widehat{\mathbf{V}}\|^2 \quad \text{s.t.,} \quad \widehat{\mathbf{U}} \odot \widehat{\mathbf{V}} = \widehat{\mathbf{w}}.$$

B.1. Proof of Lemma 5

Lemma 5 ($K = 1$) For any $\mathbf{w} \in \mathbb{R}^D$, $\mathcal{R}_{1,1}(\mathbf{w}) = 2\sqrt{D}\|\widehat{\mathbf{w}}\|_2 = 2\sqrt{D}\|\mathbf{w}\|_2$.

Proof This statement is trivially true for $\mathbf{w} = 0$, so it suffices to show this for $\mathbf{w} \neq 0$. When the $K = 1$, the first layer weight $\mathbf{U} \in \mathbb{R}^{1 \times 1}$ is a scalar. Let this scalar be $\mathbf{U} = u \neq 0$. We then have $\widehat{\mathbf{U}} = \frac{1}{\sqrt{D}}[u, u, \dots, u]$. Since $\widehat{\mathbf{U}} \odot \widehat{\mathbf{V}} = \widehat{\mathbf{w}}$, we have $\widehat{\mathbf{V}} = \frac{\sqrt{D}}{u}\widehat{\mathbf{w}}$. This means that $\mathcal{R}_{1,1}(\mathbf{w}) = \min_u \|\widehat{\mathbf{U}}\|^2 + \|\widehat{\mathbf{V}}\|^2 = \min_u u^2 + \frac{D}{u^2}\|\widehat{\mathbf{w}}\|_2^2$. By using the AM-GM inequality ($a^2 + b^2 \geq 2ab$), this is at most $2\sqrt{D}\|\widehat{\mathbf{w}}\|_2$. Moreover, we can pick $u^2 = \sqrt{D}\|\widehat{\mathbf{w}}\|_2$ to achieve equality. \blacksquare

B.2. Induced Regularizer for $K = 2$

In the following lemma we show that for $K = 2$ the characterization of $\mathcal{R}_{2,1}(\mathbf{w})$ in Fourier space involves a maximization over a high-degree rational function, and is thus unlikely to admit clean closed-form solutions.

Lemma 18 For any $\mathbf{w} \in \mathbb{R}^D$, it holds that:

$$\mathcal{R}_{2,1}(\mathbf{w}) = 2\sqrt{D} \sqrt{\inf_{\alpha \in (-1,1)} \sum_{d=0}^{D-1} \frac{|\widehat{\mathbf{w}}[d]|^2}{1 + \alpha \cos(2\pi d/D)}}.$$

Although Lemma 18 does not yield closed form solutions for $\mathcal{R}_{2,1}(\mathbf{w})$, we observe that it hints at some form of band-pass frequency structure: for any α from the inner optimization, the resulting regularizer is a weighted sum of Fourier coefficients such that the nearby frequency components of $\widehat{\mathbf{w}}$ are weighted with nearby values. This band-pass nature was also observed in a complementary result by Yun et al. (2021) in the context of implicit bias from gradient descent on a single data point: for any \mathbf{x} , it was shown that $\min_{\mathbf{w}} \mathcal{R}_{2,1}(\mathbf{w})$ s.t., $\mathbf{w}^\top \mathbf{x} > 1$ corresponds to a low-pass or high pass filter depending on the sign of $\mathbf{x}^\top \mathbf{x}^\downarrow$.

Proof [Proof of Lemma 18] We first note that for any $\mathbf{U} \in \mathbb{R}^K$, $\mathbf{V} \in \mathbb{R}^D$ we can re-scale the norms so that $\|\mathbf{U}\| = \|\mathbf{V}\|$ while satisfying the constraints of $\widehat{\mathbf{U}} \odot \widehat{\mathbf{V}} = \widehat{\mathbf{w}}$ in the definition of $\mathcal{R}_{K,1}(\mathbf{w})$. Further, such a scaling would be optimal for minimizing the ℓ_2 norm of weights based on AM-GM inequality that $\|\widehat{\mathbf{U}}\|^2 + \|\widehat{\mathbf{V}}\|^2 \geq 2\|\widehat{\mathbf{U}}\| \cdot \|\widehat{\mathbf{V}}\|$. Thus, in the rest of the proof, we consider the following equivalent formulation of $\mathcal{R}_{K,1}(\mathbf{w})$ as:

$$\mathcal{R}_{K,1}(\mathbf{w}) = \min_{\mathbf{U} \in \mathbb{R}^K, \mathbf{V} \in \mathbb{R}^D} 2\|\mathbf{U}\| \cdot \|\mathbf{V}\| \quad \text{s.t.,} \quad \widehat{\mathbf{U}} \odot \widehat{\mathbf{V}} = \widehat{\mathbf{w}} \quad (24)$$

We see that for any \mathbf{U}, \mathbf{V} satisfying the constraint in the above equation, we have: $\forall_{d \in \text{supp}(\widehat{\mathbf{w}})}$, it holds that $\widehat{\mathbf{V}}[d] = \frac{\widehat{\mathbf{w}}[d]}{\widehat{\mathbf{U}}[d]}$ (where $\text{supp}(\widehat{\mathbf{w}}) = \{d \in [D] : |\widehat{\mathbf{w}}[d]| \neq 0\}$). Moreover, at an optimal solution, it is easy to see that $\widehat{\mathbf{V}}[d] = 0 \iff \widehat{\mathbf{w}}[d] = 0$.

Let $\mathbf{U} = [c_0, c_1]$. This means that the objective can be written as

$$2\|\mathbf{U}\|\|\widehat{\mathbf{V}}\| = 2\|\mathbf{U}\|\sqrt{\sum_{d=0}^{D-1} |\widehat{\mathbf{V}}[d]|^2} = 2\|\mathbf{U}\|\sqrt{\sum_{d \in \text{supp}(\widehat{\mathbf{w}})} |\widehat{\mathbf{V}}[d]|^2} = 2\sqrt{c_0^2 + c_1^2} \sqrt{\sum_{d \in \text{supp}(\widehat{\mathbf{w}})} \frac{|\widehat{\mathbf{w}}[d]|^2}{|\widehat{\mathbf{U}}[d]|^2}}. \quad (25)$$

We write the second term in terms of the signal domain representation of c_0 and c_1 . We see that the Fourier transform of \mathbf{U} is given by $\widehat{\mathbf{U}}[d] = \frac{1}{\sqrt{D}} (c_0 + c_1 e^{-2\pi i d/D}) = (c_0 + c_1 \cos(-2\pi i d/D)) + i(c_1 \sin(-2\pi i d/D))$. We thus have that:

$$\begin{aligned} |\widehat{\mathbf{U}}[d]|^2 &= \frac{1}{D} \left[(c_0 + c_1 \cos(-2\pi i d/D))^2 + (c_1 \sin(-2\pi i d/D))^2 \right] \\ &= \frac{1}{D} (c_0^2 + c_1^2 + 2c_0 c_1 \cos(-2\pi i d/D)). \\ &= \frac{1}{D} (c_0^2 + c_1^2) \left(1 + \frac{2c_0 c_1}{c_0^2 + c_1^2} \cos(-2\pi i d/D) \right). \end{aligned} \quad (26)$$

Let $\alpha = \frac{2c_0 c_1}{c_0^2 + c_1^2}$. Plugging eq. (26) back into the objective in eq. (25) and using $\cos(z) = \cos(-z)$, we get that for any \mathbf{U}, \mathbf{V} satisfying the constraints in the computation of $\mathcal{R}_{2,1}(\mathbf{w})$, the objective is in the desired formulation:

$$2\|\mathbf{U}\|\|\widehat{\mathbf{V}}\| = 2\sqrt{D} \sqrt{\sum_{d \in \text{supp}(\widehat{\mathbf{w}})} \frac{|\widehat{\mathbf{w}}[d]|^2}{1 + \alpha \cos(2\pi i d/D)}}. \quad (27)$$

Let us now consider the domain of α , which is the only unknown in the above equation. Observe that for any c_0, c_1 , $\frac{2c_0 c_1}{c_0^2 + c_1^2} \in [-1, 1]$. Moreover, any $\alpha \in [-1, 1]$ be realized by some values of c_0 and c_1 . Thus all $\alpha \in [-1, 1]$ are valid. Here we further remark that the denominator in eq. (27) is zero if and only if $\widehat{\mathbf{U}}[d] = 0$ for any $d \in D$. However, this can only happen if $\widehat{\mathbf{w}}[d] = 0$ as otherwise the constraints $\widehat{\mathbf{U}} \odot \widehat{\mathbf{V}} = \mathbf{w}$ is not satisfied for any \mathbf{V} . We can thus, minimize the RHS of eq. (27) over $\alpha \in [-1, 1]$ to obtain $\mathcal{R}_{2,1}(\mathbf{w})$.

If we include the terms corresponding to $d \notin \text{supp}(\widehat{\mathbf{w}})$ in the summation in eq. (27), there is a technical condition than can lead to $0/0$ terms in end cases of $\alpha = 1$ (when $\widehat{\mathbf{w}}[0] = 0$) and $\alpha = -1$ (when $\widehat{\mathbf{w}}[D/2] = 0$). To avoid this technicality, we consider the infimum over $\alpha \in (-1, 1)$ rather than minimum over $\alpha \in [-1, 1]$. This is equivalent because the expression is continuous on the set of α on which it is well-defined. This completes the proof. \blacksquare

Appendix C. Remaining proofs of results in Section 3

C.1. Proof of Corollary 8

Corollary 8 For any K and any C , $\mathcal{R}_{K,C}(\mathbf{w})$ is a norm.

Corollary 8 follows from Theorem 6.

Proof [Proof of Corollary 8] It suffices to establish the scalar multiplication property, the triangle inequality, and point separation.

Scalar multiplication. Let $\gamma \in \mathbb{R}$. By definition, we see that

$$\mathcal{R}_{K,C}(\gamma \mathbf{w}) = \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 \quad \text{s.t.}, \quad \text{diag}(\widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top) = \gamma \widehat{\mathbf{w}}.$$

Let's do a change of variables $\mathbf{U} \leftarrow \frac{1}{\sqrt{|\gamma|}} \mathbf{U}$, $\mathbf{V} \leftarrow \frac{1}{\text{sign}(\gamma)\sqrt{|\gamma|}} \mathbf{V}$ to see that

$$\mathcal{R}_{K,C}(\gamma \mathbf{w}) = |\gamma| \left[\min_{\mathbf{U}, \mathbf{V}} \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 \quad \text{s.t.}, \quad \text{diag}(\widehat{\mathbf{U}}\widehat{\mathbf{V}}^\top) = \widehat{\mathbf{w}} \right] = |\gamma| \mathcal{R}_{K,C}(\mathbf{w})$$

as desired.

Triangle inequality. It suffices to show that if $\mathbf{w} = \mathbf{w}_1 + \mathbf{w}_2$, then $\mathcal{R}_{K,C}(\mathbf{w}) \leq \mathcal{R}_{K,C}(\mathbf{w}_1) + \mathcal{R}_{K,C}(\mathbf{w}_2)$. By Theorem 6, it suffices to show that $\mathcal{R}_K^{\text{SDP}}(\mathbf{w}) \leq \mathcal{R}_{K,C}(\mathbf{w}_1) + \mathcal{R}_{K,C}(\mathbf{w}_2)$. Suppose that \mathbf{U}_1 and \mathbf{V}_1 are optimal solutions to $\mathcal{R}_{K,C}(\mathbf{w}_1)$, and suppose that \mathbf{U}_2 and \mathbf{V}_2 are optimal solutions to $\mathcal{R}_{K,C}(\mathbf{w}_2)$. If we define:

$$\begin{aligned} \mathbf{Z}_1 &= \begin{bmatrix} \mathbf{U}_1 \mathbf{U}_1^\top & \mathbf{U}_1 \mathbf{V}_1^\top \\ \mathbf{V}_1 \mathbf{U}_1^\top & \mathbf{V}_1 \mathbf{V}_1^\top \end{bmatrix} \\ \mathbf{Z}_2 &= \begin{bmatrix} \mathbf{U}_2 \mathbf{U}_2^\top & \mathbf{U}_2 \mathbf{V}_2^\top \\ \mathbf{V}_2 \mathbf{U}_2^\top & \mathbf{V}_2 \mathbf{V}_2^\top \end{bmatrix} \\ \mathbf{Z} &= \mathbf{Z}_1 + \mathbf{Z}_2, \end{aligned}$$

then we see that the SDP objective $\langle \mathbf{Z}, \mathbf{I} \rangle = \langle \mathbf{Z}_1, \mathbf{I} \rangle + \langle \mathbf{Z}_2, \mathbf{I} \rangle = \mathcal{R}_{K,C}(\mathbf{w}_1) + \mathcal{R}_{K,C}(\mathbf{w}_2)$. Moreover, \mathbf{Z} is a feasible solution to (9) for \mathbf{w} . This shows $\mathcal{R}_K^{\text{SDP}}(\mathbf{w}) \leq \mathcal{R}_{K,C}(\mathbf{w}_1) + \mathcal{R}_{K,C}(\mathbf{w}_2)$ as desired.

Point separation. Notice that $\mathcal{R}_{K,C}(\mathbf{w}) = 0$, then there exist \mathbf{U} and \mathbf{V} such that $\|\mathbf{U}\|^2 + \|\mathbf{V}\|^2 = 0$. This means that $\mathbf{U} = 0$ and $\mathbf{V} = 0$, which means that $\mathbf{w} = 0$ as desired. Moreover, if $\mathbf{w} = 0$, then it's clear that $\mathcal{R}_{K,C}(\mathbf{w}) = 0$. \blacksquare

C.2. The dual formulation of SDP

In order to analyze the SDP formulation, we consider the dual. We use the formulation of the dual variable in A as $\boldsymbol{\lambda} \in \mathbb{C}^D$. In this form, the dual can be expressed as:

$$\begin{aligned} \max_{\boldsymbol{\lambda} \in \mathbb{C}^D} \quad & \text{Re}(\langle \boldsymbol{\lambda}, \widehat{\mathbf{w}} \rangle) \\ \text{s.t.} \quad & \begin{bmatrix} \mathbf{0}_K & \mathbf{F}_K^\top \overline{\boldsymbol{\lambda}} \mathbf{F} \\ \overline{\mathbf{F}} \boldsymbol{\lambda} \mathbf{F}_K & \mathbf{0}_D \end{bmatrix} \preceq \mathbf{I}. \end{aligned}$$

To simplify the objective $\text{Re}(\langle \boldsymbol{\lambda}, \widehat{\mathbf{w}} \rangle)$, notice that the phases of $\boldsymbol{\lambda}$ can be set to align with $\widehat{\mathbf{w}}$ without affecting the constraint. Thus, we can set the objective to be $|\langle \boldsymbol{\lambda}, \widehat{\mathbf{w}} \rangle|$. For convenience, we also expand out the conic constraint in vector form, and this reformulation incurs a factor of 2 on the objective. We thus obtain the following equivalent formulation of the dual:

$$\begin{aligned} \max_{\boldsymbol{\lambda} \in \mathbb{C}^D} \quad & 2 \sum_{d=0}^{D-1} |\boldsymbol{\lambda}[d]| \cdot |\mathbf{w}[d]| \\ \text{s.t.} \quad & \forall \mathbf{x} \in \mathbb{C}^K, \quad \sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 \leq 1. \end{aligned} \tag{28}$$

We now show that strong duality holds for this SDP.

Proposition 19 *The SDP in (9) satisfies strong duality.*

Proof To show strong duality, it suffices to show Slater's condition. We just need to find a solution $\boldsymbol{\lambda} \in \mathbb{C}^d$ where the inequality constraint is not tight. That is, we need to find $\boldsymbol{\lambda}$ such that $\forall \mathbf{x} \in \mathbb{C}^K, \sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 < 1$. Let's take $\boldsymbol{\lambda} = [1/2, 0, 0, \dots, 0]$. Notice that $\sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 = 0.5|\widehat{\mathbf{x}}[0]|^2 \leq 0.5 < 1$, as desired. \blacksquare

With the dual, along the fact that $\mathcal{R}_K^{\text{SDP}}(\mathbf{w})$ is a norm, we are equipped to prove general upper and lower bounds on the induced regularizer as well as sharper bounds for patterned vectors.

C.3. Proof of Lemma 9

Lemma 9 *For any K, C , and $\mathbf{w} \in \mathbb{R}^D$:*

$$2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_2 \leq \mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{D}\|\widehat{\mathbf{w}}\|_2 \quad \text{and} \quad 2\|\widehat{\mathbf{w}}\|_1 \leq \mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{\lceil \frac{D}{K} \rceil}\|\widehat{\mathbf{w}}\|_1.$$

Proof The bounds of $2\|\widehat{\mathbf{w}}\|_1$ and $2\sqrt{D}\|\widehat{\mathbf{w}}\|_2$ follow in a straightforward way from Lemma 4 and Lemma 5 coupled with Theorem 6.

The lower bound of $2\|\widehat{\mathbf{w}}\|_1$ follows as: $2\|\widehat{\mathbf{w}}\|_1 \stackrel{(a)}{=} \mathcal{R}_{D,1}(\mathbf{w}) \stackrel{(b)}{=} \mathcal{R}_{D,C}(\mathbf{w}) \stackrel{(c)}{\leq} \mathcal{R}_{K,C}(\mathbf{w})$, where (a) follows from Lemma 5, (b) follows from Theorem 6, and (c) follows from Remark 3. Similarly, the upper bound of $2\sqrt{D}\|\widehat{\mathbf{w}}\|_2$ follows as: $2\sqrt{D}\|\widehat{\mathbf{w}}\|_2 \stackrel{(a)}{=} \mathcal{R}_{1,1}(\mathbf{w}) \stackrel{(b)}{=} \mathcal{R}_{1,C}(\mathbf{w}) \geq \mathcal{R}_{K,C}(\mathbf{w})$, where (a) follows from Lemma 5 and (b) follows from Theorem 6, and (c) follows from Remark 3.

The bulk of the proof lies in showing the lower bound of $2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_2$, and an upper bound of $2\sqrt{\lceil \frac{D}{K} \rceil}\|\widehat{\mathbf{w}}\|_1$. We first prove the lower bound, and then we prove the upper bound.

Proof of the lower bound $2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_2$. We prove that $\mathcal{R}_{K,C}(\mathbf{w}) \geq 2\sqrt{\frac{D}{K}}\|\widehat{\mathbf{w}}\|_2$. It suffices to consider a dual feasible vector to eq. (28) that achieves an objective $2\sqrt{\frac{D}{K}}\|\mathbf{w}\|_2$. We consider

$$\boldsymbol{\lambda} = \frac{\mathbf{w}}{\|\mathbf{w}\|} \sqrt{\frac{D}{K}}.$$

We see that the objective is equal to

$$2 \sum_{d=0}^{D-1} |\boldsymbol{\lambda}[d]| |\mathbf{w}[d]| = 2\sqrt{\frac{D}{K}}\|\mathbf{w}\|_2,$$

as desired. It thus suffices to show that $\boldsymbol{\lambda}$ satisfies $\sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 \leq 1$ for all $\mathbf{x} \in \mathbb{C}^d$ such that $\|\mathbf{x}\|_2 \leq 1$. Using Holder's inequality, we can bound:

$$\sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 \leq \left(\max_{0 \leq d \leq D-1} |\widehat{\mathbf{x}}[d]|^2 \right) \left(\sum_{d=0}^{D-1} |\boldsymbol{\lambda}[d]|^2 \right) = \left(\max_{0 \leq d \leq D-1} |\widehat{\mathbf{x}}[d]|^2 \right) \|\boldsymbol{\lambda}\|_2^2.$$

We can bound the first term by:

$$|\widehat{\mathbf{x}}[d]| = \frac{1}{\sqrt{D}} \left| \sum_{k=0}^{K-1} \mathbf{x}[k] e^{-2\pi i k d / D} \right| \leq \frac{1}{\sqrt{D}} \sum_{k=0}^{K-1} |\mathbf{x}[k] e^{-2\pi i k d / D}| = \frac{\|\mathbf{x}\|_1}{\sqrt{D}} \leq \frac{\sqrt{K}}{\sqrt{D}}.$$

Moreover, we see that $\|\boldsymbol{\lambda}\| = \sqrt{\frac{D}{K}}$. This means that $(\max_{0 \leq d \leq D-1} |\widehat{\mathbf{x}}[d]|^2) \|\boldsymbol{\lambda}\|_2^2 \leq 1$, as desired.

Proof of the upper bound $2\sqrt{\lceil \frac{D}{K} \rceil} \|\widehat{\mathbf{w}}\|_1$. We prove that $\mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{\lceil \frac{D}{K} \rceil} \|\widehat{\mathbf{w}}\|_1$. Our main ingredient is Corollary 8 which tells us that $\mathcal{R}_{K,C}(\mathbf{w})$ is a norm. We define $T = \lceil D/K \rceil$ vectors $\mathbf{w}_0, \dots, \mathbf{w}_{T-1} \in \mathbb{R}^D$ where $\mathbf{w} = \sum_{t=0}^{T-1} \mathbf{w}_t$, and apply Corollary 8 to obtain that:

$$\mathcal{R}_{K,C}(\mathbf{w}) \leq \sum_{t=0}^{T-1} \mathcal{R}_{K,C}(\mathbf{w}_t).$$

These vectors are chosen that each $\mathcal{R}_{K,C}(\mathbf{w}_t)$ takes on a simple closed-form solution.

In order to construct the vectors \mathbf{w}_t , we consider $\widehat{\mathbf{q}} = \sqrt{\widehat{\mathbf{w}}}$, defined so that $\mathbf{w} = \mathbf{q}^\downarrow \star \mathbf{q}$ and $\|\widehat{\mathbf{q}}\|^2 = \|\mathbf{q}\|^2 = \|\widehat{\mathbf{w}}\|_1$. We define vectors $\mathbf{r}_0, \dots, \mathbf{r}_{T-1} \in \mathbb{R}^D$ such that $\sum_{t=0}^{T-1} \mathbf{r}_t = \mathbf{q}$ as follows. Roughly speaking these vectors consist of the disjoint subsets of the coordinates of \mathbf{q} corresponding to the t^{th} block of size K . More formally, for $0 \leq t \leq T-1$, let \mathbf{r}_t be defined so that $\mathbf{r}_t[l] = \mathbf{q}[l]$ for $l \in [t \cdot K, \min((t+1) \cdot K - 1, D-1)]$, and $\mathbf{r}_t[l] = 0$ otherwise. Let $\mathbf{w}_t = \mathbf{r}_t^\downarrow \star \mathbf{q}$ for $0 \leq t \leq T-1$. It is evident that $\sum_t \mathbf{r}_t = \mathbf{q}$ and hence $\sum_t \mathbf{w}_t = \mathbf{w}$.

We now show that $\mathcal{R}_{K,C}(\mathbf{w}_t) \leq 2\|\mathbf{r}_t\| \|\mathbf{q}\|$. We show this by explicitly constructing solutions to (5), taking advantage of the fact that \mathbf{r}_t is effectively a vector in \mathbb{R}^K that is zero-padded appropriately. This K dimensional vector $\mathbf{r}'_t \in \mathbb{R}^K$ is given by $\mathbf{r}'_t[k] = \mathbf{r}_t[(t \cdot K + k) \bmod D]$ for $0 \leq k \leq K-1$. Now, we wish to write \mathbf{w}_t as a convolution $(\mathbf{r}'_t)^\downarrow \star \mathbf{q}_t$, for some suitably chosen vector \mathbf{q}_t . Since $\mathbf{w}_t = \mathbf{r}_t^\downarrow \star \mathbf{q}$ and \mathbf{r}'_t is merely a circular shifted version of \mathbf{r}_t , we can take $\mathbf{q}_t \in \mathbb{R}^D$ to be \mathbf{q} with the coordinates shifted appropriately. Now, we can rescale \mathbf{r}'_t and \mathbf{q}_t so that they have equal ℓ_2 norms, and obtain the following vectors: $\left(\sqrt{\frac{\|\mathbf{q}\|_t}{\|\mathbf{r}'_t\|}}\right) \mathbf{r}'_t$ and $\left(\sqrt{\frac{\|\mathbf{r}'_t\|}{\|\mathbf{q}\|_t}}\right) \mathbf{q}_t$. These vectors are a feasible solution to eq. (5) for $\mathcal{R}_{K,C}(\mathbf{w}_t)$ and achieve an objective of $2\|\mathbf{r}'_t\| \|\mathbf{q}_t\| = 2\|\mathbf{r}_t\| \|\mathbf{q}\|$, as desired.

Using that $\mathcal{R}_{K,C}(\mathbf{w}_t) \leq 2\|\mathbf{r}_t\| \|\mathbf{q}\|$ for $0 \leq t \leq T-1$, we obtain the following bound on $\mathcal{R}_{K,C}(\mathbf{w})$:

$$\mathcal{R}_{K,C}(\mathbf{w}) \leq 2\|\mathbf{q}\| \sum_{t=0}^{T-1} \|\mathbf{r}_t\|.$$

Now, notice that since the supports of \mathbf{r}_t for $0 \leq t \leq T-1$ are disjoint, $\sum_{t=0}^{T-1} \|\mathbf{r}_t\|^2 = \|\mathbf{q}\|^2$. Applying AM-GM, this means that

$$\left(\sum_{t=0}^{T-1} \|\mathbf{r}_t\|\right)^2 \leq T \left(\sum_{t=0}^{T-1} \|\mathbf{r}_t\|^2\right) = T \|\mathbf{q}\|^2.$$

Thus, we have that

$$\mathcal{R}_{K,C}(\mathbf{w}) \leq 2\sqrt{T} \|\mathbf{q}\|^2 = 2\sqrt{T} \|\widehat{\mathbf{w}}\|_1 = 2\sqrt{\lceil \frac{D}{K} \rceil} \|\widehat{\mathbf{w}}\|_1. \quad \blacksquare$$

C.4. Proof of Lemma 10

Lemma 10 Consider vectors $\mathbf{w}(\mathbf{p}) = [\mathbf{p}, \mathbf{p}, \dots, \mathbf{p}] \in \mathbb{R}^D$ specified by $\mathbf{p} \in \mathbb{R}^P$ s.t., P divides D .

(a) For $K \leq P$, it holds that $\forall C: \mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p})) = \frac{D}{P} \cdot \mathcal{R}_{K,1}^{(P)}(\mathbf{p})$.

(b) For $K \geq P$ if $K = P \cdot T$ for integer T , then $\forall C: \mathcal{R}_{K,C}(\mathbf{w}(\mathbf{p})) = \frac{2D}{\sqrt{TP}} \|\hat{\mathbf{p}}\|_1 = 2\sqrt{\frac{D}{K}} \|\hat{\mathbf{w}}\|_1$.

Proof [Proof of Lemma 10]

We first prove an upper bound on $\mathcal{R}_{K,C}(\mathbf{w})$, and then we prove a matching lower bound. In these proofs, let $\mathbf{w} = \mathbf{w}(\mathbf{p})$. Both of these proofs use the standard fact that $\hat{\mathbf{w}}[(D/P) \cdot p] = \sqrt{\frac{D}{P}} \hat{\mathbf{p}}[p]$ for $0 \leq p \leq P - 1$, and $\hat{\mathbf{w}}[d] = 0$ if $(D/P) \nmid d$.

Upper bound. For the upper bound, it suffices to consider the case of a single output channel. We explicitly construct a pair (\mathbf{u}, \mathbf{v}) where $\|\mathbf{u}\|^2 + \|\mathbf{v}\|^2$ achieves the desired objective.

First, we consider the case where $K \leq P$. We construct (\mathbf{u}, \mathbf{v}) using an optimal solution $\mathbf{u}_{\mathbf{p}}$ and $\mathbf{v}_{\mathbf{p}}$ to eq. (5) for $\mathcal{R}_{K,1}^{(P)}(\mathbf{p})$. We let \mathbf{u} be defined so that $\mathbf{u}[k] = \sqrt{\frac{D}{P}} \mathbf{u}_{\mathbf{p}}[k]$ for $0 \leq k \leq K - 1$. We let \mathbf{v} be defined to be $\mathbf{v} = [\mathbf{v}_{\mathbf{p}}, \dots, \mathbf{v}_{\mathbf{p}}]$. Notice that:

$$\|\mathbf{u}\|^2 + \|\mathbf{v}\|^2 = \frac{D}{P} \|\mathbf{u}_{\mathbf{p}}\|^2 + \frac{D}{P} \|\mathbf{v}_{\mathbf{p}}\|^2 = \frac{D}{P} \mathcal{R}_{K,1}^{(P)}(\mathbf{p}),$$

as desired. To see that $\mathbf{u} \star \mathbf{v} = \mathbf{w}$, it suffices to show $\hat{\mathbf{u}} \odot \hat{\mathbf{v}} = \hat{\mathbf{w}}$. Notice that $\hat{\mathbf{v}}[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{v}}_{\mathbf{p}}[p]$ for $0 \leq p \leq P - 1$, and $\hat{\mathbf{v}}[d] = 0$ if $(D/P) \nmid d$. This means that $(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[d] = 0 = \hat{\mathbf{w}}[d]$ if $(D/P) \nmid d$ as desired. Thus it suffices to handle $(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[(D/P) \cdot p]$. Notice that $\hat{\mathbf{v}}[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{v}}_{\mathbf{p}}[p]$, and $\hat{\mathbf{u}}[(D/P) \cdot p] = \frac{\sqrt{P}}{\sqrt{D}} \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{u}}_{\mathbf{p}}[p] = \hat{\mathbf{u}}_{\mathbf{p}}[p]$. This means that:

$$(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{u}}_{\mathbf{p}}[p] \hat{\mathbf{v}}_{\mathbf{p}}[p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{p}}[p] = \hat{\mathbf{w}}[(D/P) \cdot p],$$

as desired.

Next, we consider the case where $K = T \cdot P$. We construct (\mathbf{u}, \mathbf{v}) using an optimal solution $\mathbf{u}_{\mathbf{p}}$ and $\mathbf{v}_{\mathbf{p}}$ to eq. (5) for $\mathcal{R}_{P,1}^{(P)}(\mathbf{p})$. We let $\mathbf{u} = \frac{\sqrt{D}}{T^{3/4}\sqrt{P}} [\mathbf{u}_{\mathbf{p}}, \dots, \mathbf{u}_{\mathbf{p}}]$ be a scaled version of T repeated copies of $\mathbf{u}_{\mathbf{p}}$. We let $\mathbf{v} = \frac{1}{T^{1/4}} [\mathbf{v}_{\mathbf{p}}, \dots, \mathbf{v}_{\mathbf{p}}]$ be a scaled version of $\frac{D}{P}$ repeated copies of $\mathbf{v}_{\mathbf{p}}$. Then,

$$\|\mathbf{u}\|^2 + \|\mathbf{v}\|^2 = \frac{D}{\sqrt{TP}} \|\mathbf{u}_{\mathbf{p}}\|^2 + \frac{D}{\sqrt{TP}} \|\mathbf{v}_{\mathbf{p}}\|^2 = \frac{D}{\sqrt{TP}} \mathcal{R}_{P,1}(\mathbf{p}) = 2\frac{D}{\sqrt{TP}} \|\hat{\mathbf{p}}\|_1,$$

as desired. To see that $\mathbf{u} \star \mathbf{v} = \mathbf{w}$, it suffices to show $\hat{\mathbf{u}} \odot \hat{\mathbf{v}} = \hat{\mathbf{w}}$. Notice that $\hat{\mathbf{v}}[(D/P) \cdot p] = \frac{\sqrt{D}}{T^{1/4}\sqrt{P}} \hat{\mathbf{v}}_{\mathbf{p}}[p]$ for $0 \leq p \leq P - 1$, and $\hat{\mathbf{v}}[d] = 0$ if $(D/P) \nmid d$. This means that $(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[d] = 0 = \hat{\mathbf{w}}[d]$ if $(D/P) \nmid d$ as desired. Thus it suffices to handle $(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[(D/P) \cdot p]$. Notice that $\hat{\mathbf{v}}[(D/P) \cdot p] = \frac{\sqrt{D}}{T^{1/4}\sqrt{P}} \hat{\mathbf{v}}_{\mathbf{p}}[p]$, and $\hat{\mathbf{u}}[(D/P) \cdot p] = \frac{T\sqrt{P}}{\sqrt{D}} \frac{\sqrt{D}}{\sqrt{PT^{3/4}}} \hat{\mathbf{u}}_{\mathbf{p}}[p] = T^{1/4} \hat{\mathbf{u}}_{\mathbf{p}}[p]$. This gives

$$(\hat{\mathbf{u}} \odot \hat{\mathbf{v}})[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{u}}_{\mathbf{p}}[p] \hat{\mathbf{v}}_{\mathbf{p}}[p] = \frac{\sqrt{D}}{\sqrt{P}} \hat{\mathbf{p}}[p] = \hat{\mathbf{w}}[(D/P) \cdot p],$$

as desired.

Lower bound. For the lower bound, we use the dual formulation in eq. (28). Our approach is to construct a dual vector $\boldsymbol{\lambda} \in \mathbb{C}^D$ so that eq. (28) achieves the desired objective.

Again, first we consider the case where $K \leq P$. Let $\boldsymbol{\lambda}_{\mathbf{p}} \in \mathbb{C}^P$ be the dual optimal solution for $\mathcal{R}_K^{\text{SDP}}(\mathbf{p})$. Now, let $\boldsymbol{\lambda}[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{P}} \boldsymbol{\lambda}_{\mathbf{p}}[p]$ for $0 \leq p \leq P-1$, and $\boldsymbol{\lambda}[d] = 0$ if $(D/P) \nmid d$. Notice that the objective becomes:

$$\begin{aligned} 2 \sum_{d=0}^D |\boldsymbol{\lambda}[d]| \cdot |\widehat{\mathbf{w}}[d]| &= 2 \sum_{p=0}^{P-1} |\boldsymbol{\lambda}[(D/P) \cdot p]| \cdot |\widehat{\mathbf{w}}[(D/P) \cdot p]| \\ &= 2 \frac{D}{P} \sum_{p=0}^{P-1} |\boldsymbol{\lambda}_{\mathbf{p}}[p]| \cdot |\widehat{\mathbf{p}}[p]| \\ &= 2 \frac{D}{P} \sum_{p=0}^{P-1} |\boldsymbol{\lambda}_{\mathbf{p}}[p]| \cdot |\widehat{\mathbf{p}}[p]| \\ &= \frac{D}{P} \mathcal{R}_K^{\text{SDP}}(\mathbf{p}) \\ &= \frac{D}{P} \mathcal{R}_{K,C}(\mathbf{p}), \end{aligned}$$

where the last equality follows from tightness of the SDP (Theorem 6). It thus suffices to show that $\boldsymbol{\lambda}$ is dual feasible. For $\mathbf{x} \in \mathbb{C}^K$ such that $\|\mathbf{x}\| \leq 1$, consider:

$$\begin{aligned} \sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 &= \sum_{p=0}^{P-1} |\widehat{\mathbf{x}}[(D/P) \cdot p]|^2 \cdot |\boldsymbol{\lambda}[(D/P) \cdot p]|^2 \\ &= \frac{D}{P} \sum_{p=0}^{P-1} |\widehat{\mathbf{x}}[(D/P) \cdot p]|^2 \cdot |\boldsymbol{\lambda}_{\mathbf{p}}[p]|^2. \end{aligned}$$

Now, let $\widehat{x}^{(P)} \in \mathbb{C}^P$ be the Fourier representation of x when the base dimension is P . Observe that $\widehat{\mathbf{x}}[(D/P) \cdot p]$ is equal to $\sqrt{\frac{P}{D}} \widehat{\mathbf{x}}^{(P)}[p]$. Thus the above expression is equal to:

$$\frac{D}{P} \frac{P}{D} \sum_{p=0}^{P-1} |\widehat{\mathbf{x}}^{(P)}[p]|^2 \cdot |\boldsymbol{\lambda}_{\mathbf{p}}[p]|^2 = \sum_{p=0}^{P-1} |\widehat{\mathbf{x}}^{(P)}[p]|^2 \cdot |\boldsymbol{\lambda}_{\mathbf{p}}[p]|^2.$$

Since $\boldsymbol{\lambda}_{\mathbf{p}}[p]$ is dual feasible for the P dimensional problem, we see that this is at most 1, as desired.

Next, we consider the case where $K = T \cdot P$. We consider $\boldsymbol{\lambda}[(D/P) \cdot p] = \frac{\sqrt{D}}{\sqrt{K}}$ for $0 \leq p \leq P-1$, and $\boldsymbol{\lambda}[d] = 0$ if $(D/P) \nmid d$. Notice that the objective in eq. (28) is equal to:

$$2 \sum_{d=0}^D |\boldsymbol{\lambda}[d]| \cdot |\widehat{\mathbf{w}}[d]| = 2 \frac{\sqrt{D}}{\sqrt{K}} \sum_{p=0}^{P-1} |\widehat{\mathbf{w}}[(D/P) \cdot p]| = 2 \frac{\sqrt{D}}{\sqrt{K}} \|\widehat{\mathbf{w}}\|_1,$$

as desired. It thus suffices to show that $\boldsymbol{\lambda}$ is dual feasible. For $\mathbf{x} \in \mathbb{C}^K$ such that $\|\mathbf{x}\| \leq 1$, we consider

$$\sum_{d=0}^{D-1} |\widehat{\mathbf{x}}[d]|^2 \cdot |\boldsymbol{\lambda}[d]|^2 = \frac{D}{K} \sum_{p=0}^{P-1} |\widehat{\mathbf{x}}[(D/P) \cdot p]|^2.$$

Let $\mathbf{x}_0, \dots, \mathbf{x}_{T-1} \in \mathbb{C}^P$ be defined so that $\mathbf{x}_t[p] = \mathbf{x}[t \cdot K + p]$ for $0 \leq p \leq P-1$ and $0 \leq t \leq T-1$. Now, let $\widehat{\mathbf{x}}_t^{(P)}$ denote the Fourier representation of \mathbf{x} when the base dimension is P , and observe that $\widehat{\mathbf{x}}[(D/P) \cdot p] = \sqrt{\frac{P}{D}} \sum_{t=0}^{T-1} \widehat{\mathbf{x}}_t^{(P)}[p]$. Thus we can rewrite the above expression as:

$$\begin{aligned} \frac{D}{K} \frac{P}{D} \sum_{p=0}^{P-1} \left| \sum_{t=0}^{T-1} \widehat{\mathbf{x}}_t^{(P)}[p] \right|^2 &= \frac{1}{T} \sum_{p=0}^{P-1} \sum_{t=0}^{T-1} T \left| \widehat{\mathbf{x}}_t^{(P)}[p] \right|^2 \\ &= \sum_{t=0}^{T-1} \left\| \widehat{\mathbf{x}}_t^{(P)} \right\|^2 \\ &= \sum_{t=0}^{T-1} \left\| \mathbf{x}_t \right\|^2 \\ &= \left\| \mathbf{x} \right\|^2 \leq 1, \end{aligned}$$

as desired. This completes the proof. \blacksquare

Appendix D. Appendix for Section 4: Networks with multi-channel inputs

For all the results in this appendix, we recall that the weights of the first and second layer are denoted as $\mathbf{U} = \{\mathbf{U}_r\}_{r \in [R]}$ with $\mathbf{U}_r \in \mathbb{R}^{K \times C} \forall_{r \in [R]}$ and $\mathbf{V} \in \mathbb{R}^{D \times C}$, respectively.

First, we provide a more formal descriptions of the matrices in the SDP. We provide a block-wise description of only the upper diagonal blocks, with lower diagonal blocks filled to satisfy the Hermitian matrix property. Additionally, for matrices $\{\mathbf{A}_{d,r}^{\text{real}}, \mathbf{A}_{d,r}^{\text{img}}\}_{d \in [D], r \in [R]}$ any unspecified block is by default treated as zero matrix $\mathbf{0}$ of appropriate dimension:

- For $r_1, r_2 \in [R]$ with $r_2 \geq r_1$, the $K \times K$ block with indices $(r_1 : (r_1 + 1)K)$ along rows and $(r_2 : (r_2 + 1)K)$ along columns is given as

$$\mathbf{Z}[r_1 : (r_1 + 1)K, r_2 : (r_2 + 1)K] = \mathbf{U}_{r_1} \mathbf{U}_{r_2}^\top;$$

- For $r \in [R]$, the $K \times D$ blocks with indices $(r : (r + 1)K)$ along rows and $(RK : (D + RK))$ along column, are given as follows:

$$\mathbf{Z}[r : (r + 1)K, RK : (D + RK)] = \mathbf{U}_r \mathbf{V}^\top.$$

Further,

$$\begin{aligned} \forall_{d \in [D]}, \mathbf{A}_{d,r}^{\text{real}}[r : (r + 1)K, RK : (D + RK)] &= \mathbf{Q}_d, \text{ and} \\ \mathbf{A}_{d,r}^{\text{img}}[r : (r + 1)K, RK : (D + RK)] &= i \cdot \mathbf{Q}_d. \end{aligned}$$

Note that for $r' \neq r$, the corresponding blocks in $\mathbf{A}_{d,r'}^{\text{real}}, \mathbf{A}_{d,r'}^{\text{img}}$ remain the default zero.

- Finally, the lower-right $D \times D$ block is given as

$$\mathbf{Z}[RK : (D + RK), RK : (D + RK)] = \mathbf{V} \mathbf{V}^\top.$$

D.1. Realizability of linear functions

We show that multiple output channels can be needed to merely realize all linear maps for multi-input-channel networks.

Lemma 11 *For any K, C and R , in order for the the network represented by $W(\mathbf{U}, \mathbf{V})$ in eq. (11) to realize all linear maps in $\mathbb{R}^{D \times R}$ it is necessary that $K \cdot C \geq \min\{R, D\}$.*

We prove this lemma by showing that the sub-network corresponding to each output channel can realize a matrix in $\mathbb{R}^{D \times R}$ of rank at most K .

Proof We first reiterate the expressions for the linear predictor $W(\mathbf{U}, \mathbf{V})$ in terms of \mathbf{U}, \mathbf{V} :

$$\forall_{r \in [R]}, W(\mathbf{U}, \mathbf{V})[:, r] = \sum_{c=0}^{C-1} \left(\mathbf{U}_r[:, c] \star \mathbf{V}[:, c]^\downarrow \right)^\downarrow. \quad (29)$$

We will now express the above formulation as matrix multiplication using the following new notation: $\forall_{c \in [C]}$ let $\underline{\mathbf{U}}_c \in \mathbb{R}^{K \times R}$ denote the representation of first layer weights corresponding to each output channel such that

$$\forall_{r \in [R]} \forall_{c \in [C]}, \underline{\mathbf{U}}_c[:, r] = \mathbf{U}_r[:, c].$$

For $c \in [C]$, consider the following matrix which consists of first K columns of the circulant matrix formed by $\mathbf{V}[:, c]$:

$$\forall_{c \in [C]}, \tilde{\mathbf{V}}_c = \frac{1}{\sqrt{D}} \begin{bmatrix} \mathbf{V}[0, c] & \mathbf{V}[D-1, c] & \cdots & \mathbf{V}[D-K+1, c] \\ \mathbf{V}[1, c] & \mathbf{V}[0, c] & \cdots & \mathbf{V}[D-K+2, c] \\ \vdots & \vdots & & \vdots \\ \mathbf{V}[D-1, c] & \mathbf{V}[D-2, c] & \cdots & \mathbf{V}[D-K, c] \end{bmatrix} \in \mathbb{R}^{D \times K}$$

Based on this notation, we can check by following the definitions that for all c ,

$$\left(\mathbf{U}_r[:, c] \star \mathbf{V}[:, c]^\downarrow \right)^\downarrow = \tilde{\mathbf{V}}_c \mathbf{U}_r[:, c] = \tilde{\mathbf{V}}_c \underline{\mathbf{U}}_c[:, r].$$

We can thus write $W(\mathbf{U}, \mathbf{V})$ as follows:

$$W(\mathbf{U}, \mathbf{V}) = \sum_{c=0}^{C-1} \tilde{\mathbf{V}}_c \underline{\mathbf{U}}_c. \quad (30)$$

We now observe that each term in the summation $\tilde{\mathbf{V}}_c \underline{\mathbf{U}}_c$ is of rank utmost K as $\tilde{\mathbf{V}}_c \in \mathbb{R}^{D \times K}$ and $\underline{\mathbf{U}}_c \in \mathbb{R}^{D \times K}$. Thus, for any \mathbf{U}, \mathbf{V} , $\text{rank}(W(\mathbf{U}, \mathbf{V})) \leq K \cdot C$. From this we conclude that in order to realize all linear maps in the multi-channel input space of $\mathbb{R}^{D \times R}$, we necessarily need $K \cdot C \geq \min\{R, D\}$.

Additionally, in eq. (30) we see that since $\underline{\mathbf{U}}_c$ are unconstrained, each term in the sum can realize any rank 1 matrix. This implies that $C \geq \min\{R, D\}$ is a sufficient condition for $W(\mathbf{U}, \mathbf{V})$ to realize any $\mathbf{W} \in \mathbb{R}^{D \times R}$. However, from Theorem 15 we know that this condition is not necessary. It is an open question to derive the tightest necessary and sufficient conditions. \blacksquare

We note that a similar proof as above can be shown using the Fourier representation in eq. (11). In the special cases of $K = 1$ and $K = D$, in Theorems 14-15 we show that the SDP is tight once C is large enough to realize all linear functions, which in these cases is $C \geq R/K$. In these end cases, we further derive interesting closed form expressions of $\mathcal{R}_{K,C,R}(\mathbf{W})$.

D.2. Proof of Lemma 12

Lemma 12 For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq RK$, it holds that $\mathcal{R}_{K,C,R}(\mathbf{W}) = \mathcal{R}_{K,R}^{SDP}(\mathbf{W})$.

Proof We show a stronger statement: any optimal solution to the SDP has rank at most RK . This implies the desired result, because $\mathcal{R}_{K,C,R}(\mathbf{W})$ is equivalent to the SDP with a rank constraint of C .

For the remainder of the proof, we let \mathbf{Z} be an optimal solution to the SDP and we prove that $\text{rank}(\mathbf{Z}) \leq RK$. Let $\text{rank}(\mathbf{Z}) = L$. We can write:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{U}_0 \\ \mathbf{U}_1 \\ \dots \\ \mathbf{U}_{R-1} \\ \mathbf{V} \end{bmatrix} [\mathbf{U}_0^\top \quad \mathbf{U}_1^\top \quad \dots \quad \mathbf{U}_{R-1}^\top \quad \mathbf{V}^\top]$$

where the matrices $\mathbf{U}_r \in \mathbb{R}^{K \times L}$ for $0 \leq r \leq R-1$ correspond to the weights in the convolution layer, the matrix $\mathbf{V} \in \mathbb{R}^{D \times L}$ corresponds to the weights in the linear layer.

It suffices to show that there exists a spanning set of the column space of $[\mathbf{U}_0^\top, \mathbf{U}_1^\top, \dots, \mathbf{U}_{R-1}^\top, \mathbf{V}^\top]^\top$ with at most RK elements. The key ingredient of the proof is the KKT conditions. Using similar logic to Appendix A, we can write the KKT conditions in the following form:

$$\text{for all } 0 \leq r \leq R-1 : \widehat{\mathbf{W}}[:, r] = \widehat{\mathbf{U}}_r \odot \widehat{\mathbf{V}} \quad (\text{KKT 1})$$

$$\begin{bmatrix} \mathbf{0}_{(R \cdot K) \times (R \cdot K)} & \begin{bmatrix} \overline{\mathbf{F}}_K^\top \mathbf{\Lambda}_0 \overline{\mathbf{F}} \\ \overline{\mathbf{F}}_K^\top \mathbf{\Lambda}_1 \overline{\mathbf{F}} \\ \vdots \\ \overline{\mathbf{F}}_K^\top \mathbf{\Lambda}_{R-1} \overline{\mathbf{F}} \end{bmatrix} \\ \hline \overline{\mathbf{F}} \mathbf{\Lambda}_0 \mathbf{F}_K \quad \overline{\mathbf{F}} \mathbf{\Lambda}_1 \mathbf{F}_K \quad \dots \quad \overline{\mathbf{F}} \mathbf{\Lambda}_{R-1} \mathbf{F}_K & \mathbf{0}_{D \times D} \end{bmatrix} \preceq \mathbf{I}_{D+KR} \quad (\text{KKT 2})$$

$$\widehat{\mathbf{V}} = \sum_{r=0}^{R-1} \overline{\mathbf{\Lambda}}_r \widehat{\mathbf{U}}_r \quad (\text{KKT 3})$$

$$\text{for all } 0 \leq r \leq R-1 : \widehat{\mathbf{U}}_r = \mathbf{F}_K \overline{\mathbf{F}}_K^\top \overline{\mathbf{\Lambda}}_r \widehat{\mathbf{V}}. \quad (\text{KKT 4})$$

where the matrices $\mathbf{\Lambda}_r \in \mathbb{C}^{D \times D}$ for $0 \leq r \leq R-1$ are diagonal and correspond to the relevant dual variables. Let's use (KKT 4) to construct a spanning set of the column space of this matrix. Let $S_r = \{\mathbf{u}_r \in \mathbb{R}^K\}$ be a basis for the column space of \mathbf{U}_r for each $0 \leq r \leq R$; since \mathbf{U}_r is $K \times L$ dimensional, we see that S_r has at most K elements. We see by (KKT 4) that the set of vectors given by the concatenation of $[\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{R-1}, \overline{\mathbf{F}} (\sum_{r=0}^{R-1} \overline{\mathbf{\Lambda}}_r \widehat{\mathbf{u}}_r)]$ for $\mathbf{u}_r \in S_r$ for each $0 \leq r \leq R$ spans the column space of the desired matrix. By construction, this spanning set has at most RK elements, as desired. ■

D.3. Proofs of Theorems 14-15: induced regularizer for $K = 1$ and $K = D$

Theorem 14 For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq \min\{R, D\}$, the induced regularizer for $K = 1$ is given by the scaled nuclear norm $\|\cdot\|_*$:

$$\mathcal{R}_{1,C,R}(\mathbf{W}) = 2\sqrt{D}\|\mathbf{W}\|_* = 2\sqrt{D}\|\widehat{\mathbf{W}}\|_*$$

Proof For $K = 1$, we have that $\forall_{r \in [R]}$, $\mathbf{U}_r \in \mathbb{R}^{1 \times C}$. For this proof, we stack the vectors \mathbf{U}_r to obtain $\tilde{\mathbf{U}} \in \mathbb{R}^{R \times C}$ such that $\forall_{r \in [R]}$, $\tilde{\mathbf{U}}[r, :] = \mathbf{U}_r$.

We work with the definition of the linear predictor realized by network as: $\forall_{r \in [R]}$, $W(\mathbf{u}, \mathbf{V})[:, r] = \sum_{c=0}^{C-1} (\mathbf{U}_r[:, c] \star \mathbf{V}[:, c]^\downarrow)^\downarrow$ (from eq. (11)). For kernel size of 1, we notice that from the definition of convolution in Definition 2, we have the following:

$$\left(\mathbf{U}_r[:, c] \star \mathbf{V}[:, c]^\downarrow\right)^\downarrow = \frac{\tilde{\mathbf{U}}[r, c]}{\sqrt{D}} \mathbf{V}[:, c] \propto \mathbf{V}[:, c]. \quad (31)$$

Plugging this back into the expression of $W(\mathbf{u}, \mathbf{V})$, we have the following:

$$W(\mathbf{u}, \mathbf{V})[:, r] = \frac{1}{\sqrt{D}} \mathbf{V} \tilde{\mathbf{U}}[r, :] \implies W(\mathbf{u}, \mathbf{V}) = \frac{1}{\sqrt{D}} \mathbf{V} \tilde{\mathbf{U}}^\top. \quad (32)$$

In the above formulation $\mathbf{V} \in \mathbb{R}^{D \times C}$, $\tilde{\mathbf{U}} \in \mathbb{R}^{R \times C}$ are of rank C , but otherwise completely unconstrained. Thus, they can realize any rank K matrix. So as long as $C \geq \min\{R, D\}$, the network can realize any linear predictor $\mathbf{W} \in \mathbb{R}^{D \times R}$.

The rest of the proof follows from connecting the above expression into the variational characterization of the nuclear norm. The induced regularizer $\mathcal{R}_{1,C,R}(\mathbf{W})$ from eq. (12) for $C \geq \min\{R, D\}$ can now be expressed as follows:

$$\begin{aligned} \mathcal{R}_{1,C,R}(\mathbf{W}) = \min_{\tilde{\mathbf{U}} \in \mathbb{R}^{R \times C}, \mathbf{V} \in \mathbb{R}^{D \times C}} & \|\tilde{\mathbf{U}}\|^2 + \|\mathbf{V}\|^2 \\ \text{s.t.,} & \sqrt{D}\mathbf{W} = \mathbf{V}\tilde{\mathbf{U}}^\top. \end{aligned} \quad (33)$$

For $C \geq \min\{R, D\}$ eq. (33) is exactly the variational definition of nuclear norm (see Rennie and Srebro (Lemma 1 2005)) and thus $\mathcal{R}_{1,C,R}(\mathbf{W}) = 2\sqrt{D}\|\mathbf{W}\|_*$ for even unbounded C . The fact that $C = \min\{R, D\}$ is sufficient can be seen by obtaining the optimum nuclear norm as upper bound from using $\mathbf{V} = \mathbf{L}\sqrt{\Sigma}$ and $\tilde{\mathbf{U}} = \mathbf{R}\sqrt{\Sigma}$, where $\sqrt{D}\mathbf{W} = \mathbf{L}\Sigma\mathbf{R}^\top$ is the singular value decomposition of $\sqrt{D}\mathbf{W}$. Finally, we note that based on our normalization of Fourier transform, we have $\|\mathbf{W}\|_* = \|\widehat{\mathbf{W}}\|_*$. This completes our proof.

Note: For $K = 1$ we provided the proof of $\mathcal{R}_{1,C,R}(\mathbf{W})$ in the signal space of \mathbf{W} , but the Theorem can also be proved in the Fourier domain (similar to the proof of Theorem 15 given below) by first showing that the SDP relaxation evaluates to the nuclear norm and combining this with the matching upper bound for $\mathcal{R}_{1,C,R}(\mathbf{W})$ shown above. \blacksquare

Theorem 15 For any $\mathbf{W} \in \mathbb{R}^{D \times R}$, and any $C \geq 1$, the induced regularizer for $K = D$ is given by

$$\mathcal{R}_{D,C,R}(\mathbf{W}) = 2\|\widehat{\mathbf{W}}\|_{2,1} := \sum_{d=0}^{D-1} \sqrt{\sum_{r=0}^{R-1} |\widehat{\mathbf{W}}[d, r]|^2}.$$

Proof We begin by expressing induced regularizer for full dimensional kernels $\mathcal{R}_{D,C,R}(\mathbf{W})$ from eq. (12) in terms of the Fourier representation of the linear predictor realized by the network $W(\mathbf{U}, \mathbf{V})$:

$$\begin{aligned} \mathcal{R}_{D,C,R}(\mathbf{W}) &= \inf_{\mathbf{U}, \mathbf{V}} \|\widehat{\mathbf{U}}_r\|^2 + \|\widehat{\mathbf{V}}\|^2 \\ \text{s.t., } \forall_{r \in [R]} \mathbf{W}[:, r] &= \text{diag}(\widehat{\mathbf{U}}_r \widehat{\mathbf{V}}^\top), \end{aligned} \quad (34)$$

where \mathbf{U}, \mathbf{V} are of dimensions $\mathbf{U} = \{\mathbf{U}_r \in \mathbb{R}^{D \times C}\}_{r \in [R]}$, $\mathbf{V} \in \mathbb{R}^{D \times C}$.

Our proof for networks with multi-channel inputs with $K = D$ follows the following structure:

Step 1. We first show an upper bound on the induced regularizer for single output channel $C = 1$ with full dimensional kernel $K = D$ as $\mathcal{R}_{D,1,R}(\mathbf{W}) \leq 2\|\widehat{\mathbf{W}}\|_{2,1}$ by providing a construction of \mathbf{U}, \mathbf{V} . It immediately follows from the monotonicity of $\mathcal{R}_{D,C,R}$ that for all C , $\mathcal{R}_{D,C,R}(\mathbf{W}) \leq \mathcal{R}_{D,1,R}(\mathbf{W}) \leq 2\|\widehat{\mathbf{W}}\|_{2,1}$. This step also consequently shows that when $K = D$, every linear predictor over the multi-channel input space of $\mathbb{R}^{D \times R}$ is realizable by a network with even a single output channel.

Step 2. The bulk of our proof lies in matching the upper bound with a lower bound on the dual problem of the SDP in eq. (13) as $\mathcal{R}_{D,R}^{\text{SDP}}(\mathbf{W}) \geq 2\|\widehat{\mathbf{W}}\|_{2,1}$. This gives us that for all $C \geq 1$, $\mathcal{R}_{D,C,R}(\mathbf{W}) \geq \mathcal{R}_{D,R}^{\text{SDP}}(\mathbf{W}) \geq 2\|\widehat{\mathbf{W}}\|_{2,1}$.

Step 1. Upper bound on the induced regularizer $\mathcal{R}_{D,C,R}$: We first show that $\mathcal{R}_{D,1,R}(\mathbf{W}) \leq 2\|\widehat{\mathbf{W}}\|_{2,1} = 2\sum_{d \in [D]} \|\widehat{\mathbf{W}}[d, :]\|$. Since $\mathcal{R}_{D,R,C}(\mathbf{W})$ is decreasing in C , it suffices to show this for $C = 1$. For $C = 1$ and $K = D$, we have $\mathbf{V} \in \mathbb{R}^D$ and $\forall_{r \in [R]} \mathbf{U}_r \in \mathbb{R}^D$. For full dimensional kernels, the Fourier domain representations $\widehat{\mathbf{U}}_0, \widehat{\mathbf{U}}_1, \dots, \widehat{\mathbf{U}}_R, \widehat{\mathbf{V}} \in \mathbb{C}^D$ are all unconstrained beyond the symmetry properties of Fourier transform of real matrices. Thus consider the following \mathbf{U}, \mathbf{V} :

$$\forall_{d \in [D]} \forall_{r \in [R]}, \widehat{\mathbf{U}}_r[d] = \frac{\widehat{\mathbf{W}}[d, r]}{\sqrt{\|\widehat{\mathbf{W}}[d, :]\|}}, \quad \text{and} \quad \widehat{\mathbf{V}}[d] = \sqrt{\|\widehat{\mathbf{W}}[d, :]\|}. \quad (35)$$

It is easy to see that the above \mathbf{U}, \mathbf{V} satisfy the constraints of $\mathcal{R}_{D,1,R}(\mathbf{W})$ in eq. (34) that is $\forall_{r \in [R]} \mathbf{W}[:, r] = \text{diag}(\widehat{\mathbf{U}}_r \widehat{\mathbf{V}}^\top)$. Further $\widehat{\mathbf{U}}_0, \widehat{\mathbf{U}}_1, \dots, \widehat{\mathbf{U}}_R, \widehat{\mathbf{V}} \in \mathbb{C}^D$ satisfy the required symmetry properties since \mathbf{W} is real.

Now computing the objective, we immediately have that $\|\widehat{\mathbf{V}}\|^2 = \sum_{d \in [D]} \|\widehat{\mathbf{W}}[d, :]\|$, and further,

$$\sum_{r \in [R]} \|\widehat{\mathbf{U}}_r\|^2 = \sum_{d \in [D]} \left[\sum_{r \in [R]} \frac{|\widehat{\mathbf{W}}[d, r]|^2}{\|\widehat{\mathbf{W}}[d, :]\|} \right] = \sum_{d \in [D]} \|\widehat{\mathbf{W}}[d, :]\|.$$

This construction thus gives us the desired upper bound

$$\mathcal{R}_{D,C,R}(\mathbf{W}) \leq \mathcal{R}_{D,1,R}(\mathbf{W}) \leq 2\|\widehat{\mathbf{W}}\|_{2,1}. \quad (36)$$

Step 2: Lower bound on the induced regularizer $\mathcal{R}_{D,C,R}$ We show the lower bound by lower bounding the dual problem of the SDP in eq. (13). For $r \in [R]$, let $\boldsymbol{\lambda}_r^{\text{real}} \in \mathbb{R}^D$ and $\boldsymbol{\lambda}_r^{\text{img}} \in \mathbb{R}^D$ denote the dual variables corresponding to the constraints in the SDP in eq. (13) for the real and imaginary parts, respectively, of $\widehat{\mathbf{W}}[:, r]$. Similar to single input channel proofs, we define $\boldsymbol{\lambda}_r = \boldsymbol{\lambda}_r^{\text{real}} + i \cdot \boldsymbol{\lambda}_r^{\text{img}}$ and $\boldsymbol{\Lambda}_r = \text{diag}(\boldsymbol{\lambda}_r)$. Additionally, we introduce the notation for the matrix obtained by taking $\{\boldsymbol{\lambda}_r\}_r$ as columns: $\Xi = [\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_R] \in \mathbb{C}^{D \times R}$ such that $\forall_r, \Xi[:, r] = \boldsymbol{\lambda}_r$.

Based on weak duality for the SDP in eq. (13), we have the following:

$$\begin{aligned} \mathcal{R}_{D,R}^{\text{SDP}}(\mathbf{W}) &\geq \max_{\Xi = [\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_R]} 2 \cdot \text{Re}(\langle \widehat{\mathbf{W}}, \Xi \rangle) \\ \text{s.t.}, &\quad \sum_{d,r} (\boldsymbol{\lambda}_r^{\text{real}}[d] \cdot \mathbf{A}_{d,r}^{\text{real}} + \boldsymbol{\lambda}_r^{\text{img}}[d] \cdot \mathbf{A}_{d,r}^{\text{img}}) \preccurlyeq \mathbf{I}. \end{aligned} \quad (37)$$

Our rest of the proof obtains the lower bound by constructing an appropriate Ξ satisfying the constraints: From the definitions of $\{\mathbf{A}_{d,r}^{\text{real}}, \mathbf{A}_{d,r}^{\text{img}}\}_{d \in [D], r \in [R]}$ in Appendix 4.2 and using $\mathbf{Q}_d = \overline{\mathbf{F}} \mathbf{e}_d \mathbf{e}_d^\top \overline{\mathbf{F}}$ (note that for $K = D$ and $\mathbf{F}_K = \mathbf{F} = \mathbf{F}^\top$), we have the following:

$$\sum_{d,r} (\boldsymbol{\lambda}_r^{\text{real}}[d] \cdot \mathbf{A}_{d,r}^{\text{real}} + \boldsymbol{\lambda}_r^{\text{img}}[d] \cdot \mathbf{A}_{d,r}^{\text{img}}) = \begin{bmatrix} & & & & & & & \overline{\mathbf{F}} \boldsymbol{\Lambda}_0 \overline{\mathbf{F}} \\ & & & & & & & \overline{\mathbf{F}} \boldsymbol{\Lambda}_1 \overline{\mathbf{F}} \\ & & & & & & & \vdots \\ & & & & & & & \overline{\mathbf{F}} \boldsymbol{\Lambda}_R \overline{\mathbf{F}} \\ \hline \mathbf{F} \boldsymbol{\Lambda}_0 \mathbf{F}_K & \mathbf{F} \boldsymbol{\Lambda}_1 \mathbf{F}_K & \dots & & \mathbf{F} \boldsymbol{\Lambda}_R \mathbf{F}_K & & & \mathbf{0}_{D \times D} \end{bmatrix} \quad (38)$$

We state and prove the following claim:

Claim Ξ satisfies the constraints of the dual problem in the RHS of eq. (37) if $\max_{d \in [D]} \|\Xi[d, :]\| \leq 1$

Proof [Proof of claim.] The relevant constraint in eq. (37) is $\sum_{d,r} (\boldsymbol{\lambda}_r^{\text{real}}[d] \cdot \mathbf{A}_{d,r}^{\text{real}} + \boldsymbol{\lambda}_r^{\text{img}}[d] \cdot \mathbf{A}_{d,r}^{\text{img}}) \preccurlyeq \mathbf{I}$.

It thus suffices to show that for all $\mathbf{y} \in \mathbb{R}^D$ such that $\|\mathbf{y}\| = 1$, it holds that $\sum_{r=0}^R \|\overline{\mathbf{F}} \boldsymbol{\Lambda}_r \overline{\mathbf{F}} \mathbf{y}\|^2 \leq 1$. We have the following set of inequalities that prove the claim $\forall_{\mathbf{y}: \|\mathbf{y}\|=1}$:

$$\begin{aligned} \sum_{r=0}^R \|\overline{\mathbf{F}} \boldsymbol{\Lambda}_r \overline{\mathbf{F}} \mathbf{y}\|^2 &\stackrel{(a)}{=} \sum_{r=0}^R \|\boldsymbol{\Lambda}_r(\overline{\mathbf{F}} \mathbf{y})\|^2 \\ &= \sum_{d=0}^D \sum_{r=0}^R |\boldsymbol{\lambda}_r[d]|^2 \cdot |(\overline{\mathbf{F}} \mathbf{y})[d]|^2 \stackrel{(b)}{=} \sum_{d=0}^D \|\Xi[d, :]\|^2 \cdot |(\overline{\mathbf{F}} \mathbf{y})[d]|^2 \\ &\leq \max_{d \in [D]} \|\Xi[d, :]\|^2 \|\overline{\mathbf{F}} \mathbf{y}\|^2 \\ &\stackrel{(c)}{=} \max_{d \in [D]} \|\Xi[d, :]\|^2, \end{aligned}$$

where (a) follows from $\overline{\mathbf{F}}$ being unitary, (b) from definition of Ξ , and (c) from $\|\overline{\mathbf{F}} \mathbf{y}\| = \|\mathbf{y}\| = 1$. ■

Consider Ξ defined as follows:

$$\forall_{d \in [D]}, \quad \Xi[d, :] = \frac{\widehat{\mathbf{W}}[d, :]}{\|\widehat{\mathbf{W}}[d, :]\|}$$

It is easy to check that $\max_{d \in [D]} \|\Xi[d, :]\| = 1$ and thus based on the claim we proved above Ξ satisfies the constraints of the dual optimization problem in the RHS of eq. (37). Additionally, the objective evaluates to the desired bound of $\text{Re}(\langle \widehat{\mathbf{W}}, \Xi \rangle) = \sum_d \|\widehat{\mathbf{W}}[d, :]\| = \|\widehat{\mathbf{W}}\|_{2,1}$. We thus have the following lower bound:

$$\mathcal{R}_{D,C,R}(\mathbf{W}) \geq \mathcal{R}_{D,R}^{\text{SDP}}(\mathbf{W}) \geq 2 \text{Re}(\langle \widehat{\mathbf{W}}, \Xi \rangle) = 2\|\widehat{\mathbf{W}}\|_{2,1}. \quad (39)$$

Conclusion of the proof. The proof of the Theorem follows from combining the matching upper and lower bounds in eq. (36) and eq. (39), respectively, to obtain $\mathcal{R}_{D,C,R}(\mathbf{W}) = 2\|\widehat{\mathbf{W}}\|_{2,1}$. ■

Appendix E. Experiments for gradient descent

Theorem (Paraphrased from [Lyu and Li, 2020](#)) Assume that Φ is locally Lipschitz and positive homogeneous with order $L > 0$, i.e., $\forall_{\theta, \alpha > 0}, \Phi(\alpha\theta; \cdot) = \alpha^L \Phi(\theta; \cdot)$. Consider minimization of logistic or exponential loss over a separable binary classification dataset $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ using a gradient flow trajectory denoted as $\theta(t)$. Under the assumption that the training data points are correctly classified in finite time (i.e., mean logistic/exponential loss $\mathcal{L}(\theta(t_0)) < 1$ at finite t_0), the limit points of the direction of parameters along the gradient flow correspond to a first order stationary point (KKT point) of the following max- ℓ_2 margin problem in parameter space:

$$\min_{\theta} \|\theta\|_2^2 \quad \text{s.t.}, \quad \forall_n y_n \Phi(\theta; \mathbf{x}_n) \geq 1. \quad (40)$$

Remark With additional assumptions, [Lyu and Li \(2020\)](#) also prove an analogous result for gradient descent. This result was further refined in [Ji and Telgarsky \(2020\)](#), which proves the directional convergence of the gradient descent under additional structure on the model class.

We run our experiments on two layer linear convolutional networks on a subset of MNIST dataset ([LeCun and Cortes, 2010](#)) as well a subset of the CIFAR-10 dataset ([Krizhevsky, 2009](#)). The input images in MNIST are of size 28×28 and have a single input channel. The input images in MNIST are of size 32×32 and have a 3 input channels. We apply 2D convolutions with kernel sizes $K = (K_1, K_2)$ and circular padding for image inputs. We consider binary classification task for both datasets. For MNIST, we predict digits 0 and 1 in MNIST using a balanced sub-sampling of 128 samples as training data, which ensures linear separability. For CIFAR-10, we predict classes “automobile” and “dog” using a balanced sub-sampling of 256 samples as training data, which ensures linear separability. The initialization scale was taken to be 0.001.

We train our network using gradient descent on exponential loss and run gradient descent until the training loss is 10^{-6} . The initialization scale is taken to be 0.001 in order to reduce the variance arising from the randomness of initialization. In order to compare the predictors across different architectures, we normalize the weights learned by gradient descent \mathbf{U}, \mathbf{V} such that the linear predictor $w(\mathbf{U}, \mathbf{V})$ realized by the trained networks has unit margin on the training dataset (i.e., $y \langle w(\mathbf{U}, \mathbf{V}), \mathbf{x} \rangle \geq 1$ for all training samples (\mathbf{x}, y)). Note that for homogeneous models, such positive scaling of weights *does not* change the classification boundary of the learned model.⁸

8. The code is available at <https://github.com/mjagadeesan/inductive-bias-multi-channel-CNN>.




Figure 3: Linear predictors learned by two layer linear convolutional network for the task of classifying digits 0 and 1 in MNIST. The sub-figures depict predictors learned by using gradient descent on the exponential loss for overparameterized networks with $C = 1, 2, 4$ and kernel sizes $K \in \{(1, 1), (3, 3), (8, 8), (16, 16), (28, 28)\}$ (left to right). The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\widehat{w}(\mathbf{U}, \mathbf{V})$.

E.1. Impact of the number of channels on MNIST

For networks with a single-input channel, Hypothesis 1 would imply that $\mathcal{R}_{K,C,R}(w(\mathbf{U}, \mathbf{V}))$ is invariant in the number of output channels C regardless of C . To demonstrate this, we repeat the experimental setup on 28×28 MNIST images on networks with multiple output channels $C \in \{1, 2, 4, 8\}$ and across different kernel sizes. As described earlier, we scale the weights learned by gradient descent \mathbf{U}, \mathbf{V} such that the linear predictors $w(\mathbf{U}, \mathbf{V}) = \mathbf{w}_{\text{GD}}$ have unit margin on training data. Since it is difficult to directly compute $\mathcal{R}_{K,C,R}(\mathbf{w}_{\text{GD}})$, we turn to an approximation. In particular, we compute $\widehat{\mathcal{R}}_{K,C}(w(\mathbf{U}, \mathbf{V})) := \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$. Strictly speaking, this is only an *upper bound* on the induced regularizer $\mathcal{R}_{K,C,R}(w(\mathbf{U}, \mathbf{V}))$.⁹

Table 1 shows $\widehat{\mathcal{R}}_{K,C}(w(\mathbf{U}, \mathbf{V}))$ across different values of K and C . We see that for each kernel size, the differences in $\widehat{\mathcal{R}}_{K,C}(w(\mathbf{U}, \mathbf{V}))$ across different settings of C are minimal and are usually smaller than the standard deviation for fixed settings of C . This suggests that the induced regularizer is indeed invariant to the number of output channels, thus providing evidence for Hypothesis 1 in the case of a single input channel.

Invariance of learned predictors to C . While Hypothesis 1 primarily pertains to the behavior of the induced regularizer, it also suggests that the predictor \mathbf{w}_{GD} will also be independent of the number of channels so long as K is strictly less than D . If we overlook the caveats, we expect gradient

9. For $K = 1$ and $K = D$, we verified that the estimate is close to tight by computing the ℓ_2 and ℓ_1 norms of Fourier transform of the predictor, respectively.




Figure 4: Explicit $\mathcal{R}_{K,C}$ margin predictor on sampled MNIST dataset for kernel sizes $K \in \{(1, 1), (28, 28)\}$ (left to right). The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$.

C	$K = (1, 1)$	$K = (3, 3)$	$K = (8, 8)$	$K = (16, 16)$	$K = (28, 28)$
1	$10.28 \pm 2.34 \times 10^{-5}$	$4.50 \pm 1.51 \times 10^{-3}$	$3.32 \pm 5.35 \times 10^{-2}$	$3.15 \pm 5.81 \times 10^{-2}$	$2.84 \pm 1.16 \times 10^{-1}$
2	$10.28 \pm 2.00 \times 10^{-5}$	$4.50 \pm 1.06 \times 10^{-3}$	$3.30 \pm 3.13 \times 10^{-2}$	$3.10 \pm 2.63 \times 10^{-2}$	$2.79 \pm 1.27 \times 10^{-1}$
4	$10.28 \pm 1.00 \times 10^{-5}$	$4.50 \pm 7.48 \times 10^{-4}$	$3.30 \pm 2.25 \times 10^{-2}$	$3.10 \pm 3.33 \times 10^{-2}$	$2.77 \pm 8.20 \times 10^{-2}$
8	$10.28 \pm 7.83 \times 10^{-5}$	$4.50 \pm 6.25 \times 10^{-4}$	$3.29 \pm 1.68 \times 10^{-2}$	$3.11 \pm 3.27 \times 10^{-2}$	$2.72 \pm 7.31 \times 10^{-2}$

Table 1: $\hat{\mathcal{R}}_{K,C}(w(\mathbf{U}, \mathbf{V})) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the predictor learned by gradient descent on linear convolutional networks with different number of output channels C and kernel sizes K on the MNIST task. We show the mean over 10 trials as well as the standard deviations are also shown.

descent to implicitly learn a max- $\mathcal{R}_{K,C}$ margin predictor: $\min_{\mathbf{w}} \mathcal{R}_{K,C}(\mathbf{w})$ s.t., $\forall_n y_n \langle \mathbf{w}, \mathbf{x}_n \rangle \geq 1$. For $K < D$, our theoretical findings suggest that the induced regularizer is a norm interpolating between the ℓ_2 and ℓ_1 norms. This would mean that there is a *unique* global minimizer, and thus we would expect that \mathbf{w}_{GD} to be invariant to C .

To empirically validate this, we show the learned linear predictors for $C = \{1, 2, 4\}$ in Figure 3. We observe that the linear predictors indeed visually appears to be invariant across different settings of C for all for kernel sizes $K < D$. For $K = (28, 28)$, there appear to be differences in the predictors—this likely arises from the fact that there are multiple linear predictors that minimize the ℓ_1 norm on the dataset. We nonetheless emphasize that the *induced regularizer* still appears to be invariant in this case, although the predictors are not.

E.1.1. NON-LINEAR NETWORKS WITH RELU ACTIVATION

Although our theoretical results are restricted to networks with linear activations, it is nevertheless interesting to evaluate if our conclusions lead to useful heuristics for networks with non-linearity. As a simple demonstration, we repeat our experiment on MNIST on two-layer convolutional networks with ReLU non-linearity with and without bias parameters (*i.e.*, networks with a convolution layer, followed by ReLU layer, followed by linear layer).¹⁰ As before, we first scale the weights learned by gradient descent such that the resulting predictor f_{GD} has unit margin on training data. We

10. The initialization scale was taken to be 0.005 for networks without bias parameters and 0.01 for networks with bias parameters.

then consider the representation cost $\mathcal{R}_{\Phi_{K,C}}(f_{\text{GD}})$, as per equation (1), given by the minimum ℓ_2 norm of the weights needed to realize f . We consider the approximation of $\mathcal{R}_{\Phi_{K,C}}(f_{\text{GD}})$ given by $\widehat{\mathcal{R}}_{\Phi_{K,C}}(f_{\text{GD}}) := \|\mathbf{U}\|_2^2 + \|\mathbf{V}\|_2^2$ where \mathbf{U} and \mathbf{V} are the weights learned by gradient descent. (As before, strictly speaking, this is only an *upper bound* on the representation cost $\mathcal{R}_{\Phi_{K,C}}(f_{\text{GD}})$.)

Table 2 and Table 3 show $\widehat{\mathcal{R}}_{\Phi_{K,C}}(f_{\text{GD}}) := \|\mathbf{U}\|_2^2 + \|\mathbf{V}\|_2^2$ across different settings of C and K , for networks with no bias as well as networks with bias parameters on both the convolution layer and the fully connected layer.¹¹ Like in the case of linear convolutional neural networks, $\widehat{\mathcal{R}}_{\Phi_{K,C}}(f_{\text{GD}})$ is consistent across different settings of C . This suggests that the implicit bias from gradient might result in predictors that are independent of the number of output channels, even when there is a ReLU layer, and Hypothesis 1 might hold in much more generality than the scope of theoretical findings.

C	$K : (1, 1)$	$K : (3, 3)$	$K : (8, 8)$	$K : (16, 16)$	$K : (28, 28)$
1	11.412	5.160	3.998	3.785	3.520
2	11.413	5.155	3.964	3.721	3.539
4	11.414	5.153	3.966	3.719	3.448
8	11.415	5.156	3.971	3.738	3.498

Table 2: $\widehat{\mathcal{R}}_{K,C}(f_{\text{GD}}) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the predictor learned by gradient descent on ReLU convolutional networks without bias parameters, with different number of output channels C and kernel sizes K , on the MNIST task. The values shown are the medians taken over 5 trials.

C	$K : (1, 1)$	$K : (3, 3)$	$K : (8, 8)$	$K : (16, 16)$	$K : (28, 28)$
1	10.581	4.948	3.875	3.714	3.519
2	10.571	4.945	3.910	3.698	3.413
4	10.578	4.945	3.912	3.712	3.399
8	10.576	4.946	3.881	3.697	3.437

Table 3: $\widehat{\mathcal{R}}_{K,C}(f_{\text{GD}}) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the predictor learned by gradient descent on ReLU convolutional networks with bias on both layers, with different number of output channels C and kernel sizes K , on the MNIST task. The values shown are the medians taken over 5 trials.

We note that we observed that gradient descent sometimes leads to outliers where $\mathcal{R}_{\Phi_{K,C}}(f_{\text{GD}})$ is very large. For example, when $K = 1$, the values of $\mathcal{R}_{\Phi_{K,C}}(f)$ are [11.412, 11.412, 109.471, 11.412, 11.412], where 109.471 appears to be an outlier. We anticipate that this outlier arises because gradient descent converges to a stationary point, rather than a local minima, of the max- ℓ_2 margin problem in parameter space (see the discussion in Section 1.2). Since our goal is to investigate the behavior of gradient descent when it does lead to global minima of the $\max - \mathcal{R}_{\Phi}$ margin problem, we compute the median so that these data points do not affect our estimate.

E.2. Impact of the number of channels on CIFAR-10

We carry out a similar investigation of Hypothesis 1 on the CIFAR-10 dataset for networks with 3-channel inputs. As discussed in Section 4, we expect that the induced regularizer is *not* independent of the number of output channels for $C < R$, but begins to exhibit invariance once $C \geq R$. On the CIFAR-10 dataset, where there are 3 input channels, we would expect to see invariance once $C \geq 3$.

11. We note that the representation cost includes the magnitude of the *weights* but not the magnitude of the *biases*.

To demonstrate this, we repeat the experimental setup on 32×32 CIFAR-10 images on networks with multiple output channels $C \in \{1, 2, 4, 8\}$ and across different kernel sizes. As described earlier, we scale the weights learned by gradient descent \mathbf{U}, \mathbf{V} such that the linear predictor $w(\mathbf{U}, \mathbf{V}) = \mathbf{w}_{\text{GD}}$ has unit margin on training data. Since it is difficult to directly compute $\mathcal{R}_{K,C,R}(\mathbf{w}_{\text{GD}})$, we turn to an approximation, as we did in the case of single-input channels. We compute $\widehat{\mathcal{R}}_{K,C,R}(w(\mathbf{U}, \mathbf{V})) := \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ which strictly speaking, this is only an *upper bound* on the induced regularizer $\mathcal{R}_{K,C,R}(w(\mathbf{U}, \mathbf{V}))$.

Table 4 shows $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V}))$ across different values of K and C . We see that for each kernel size, the differences in $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V}))$ across different settings of C are minimal, as long as $C \geq 3$ (and often, even when $C \geq 2$). This suggests that the induced regularizer is indeed invariant to the number of output channels when $C \geq 3$, thus providing evidence for Hypothesis 1 in the case of multiple input channels. Moreover, there are non-trivial differences in $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V}))$ for $C = 1$ and larger C —this aligns with our theoretical findings in Section 4 that the induced regularizer does depend on C when it is below R .

While Hypothesis 1 primarily pertains to the behavior of the induced regularizer, we would also expect that there is a unique global minimizer in most cases, for reasons similar to for the single input channel case. Thus, we would expect \mathbf{w}_{GD} to be invariant to C as long as $C \geq R = 3$. To empirically validate this, we show the learned linear predictors for $C = \{1, 2, 3, 4, 8\}$ in Figures 5-7. We observe that the linear predictors indeed visually appears to be invariant across different settings of C .

C	$K = (1, 1)$	$K = (3, 3)$	$K = (8, 8)$	$K = (20, 20)$
1	246.04	215.21	202.27	131.16
2	246.26	182.77	168.40	124.66
3	245.98	182.80	165.32	123.56
4	246.29	182.83	164.50	123.37
8	245.58	182.82	164.86	123.59

Table 4: $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V})) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the predictor learned by gradient descent on linear convolutional networks with different number of output channels C and kernel sizes K on the CIFAR-10 task.

E.2.1. CONVOLUTIONS WITH ZERO PADDING

We additionally rerun the same setup as in the previous section with zero padding rather than circular padding. Table 5 shows $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V}))$ across different values of K and C . Although convolutions with zero padding goes beyond the scope of our theoretical results, we show that Hypothesis 1 nonetheless still holds.

E.3. Varying kernel sizes

While the number of output channels has little influence on the induced regularizer of the learned predictors, we show that the kernel size can have significant impact, which aligns with our theoretical findings.

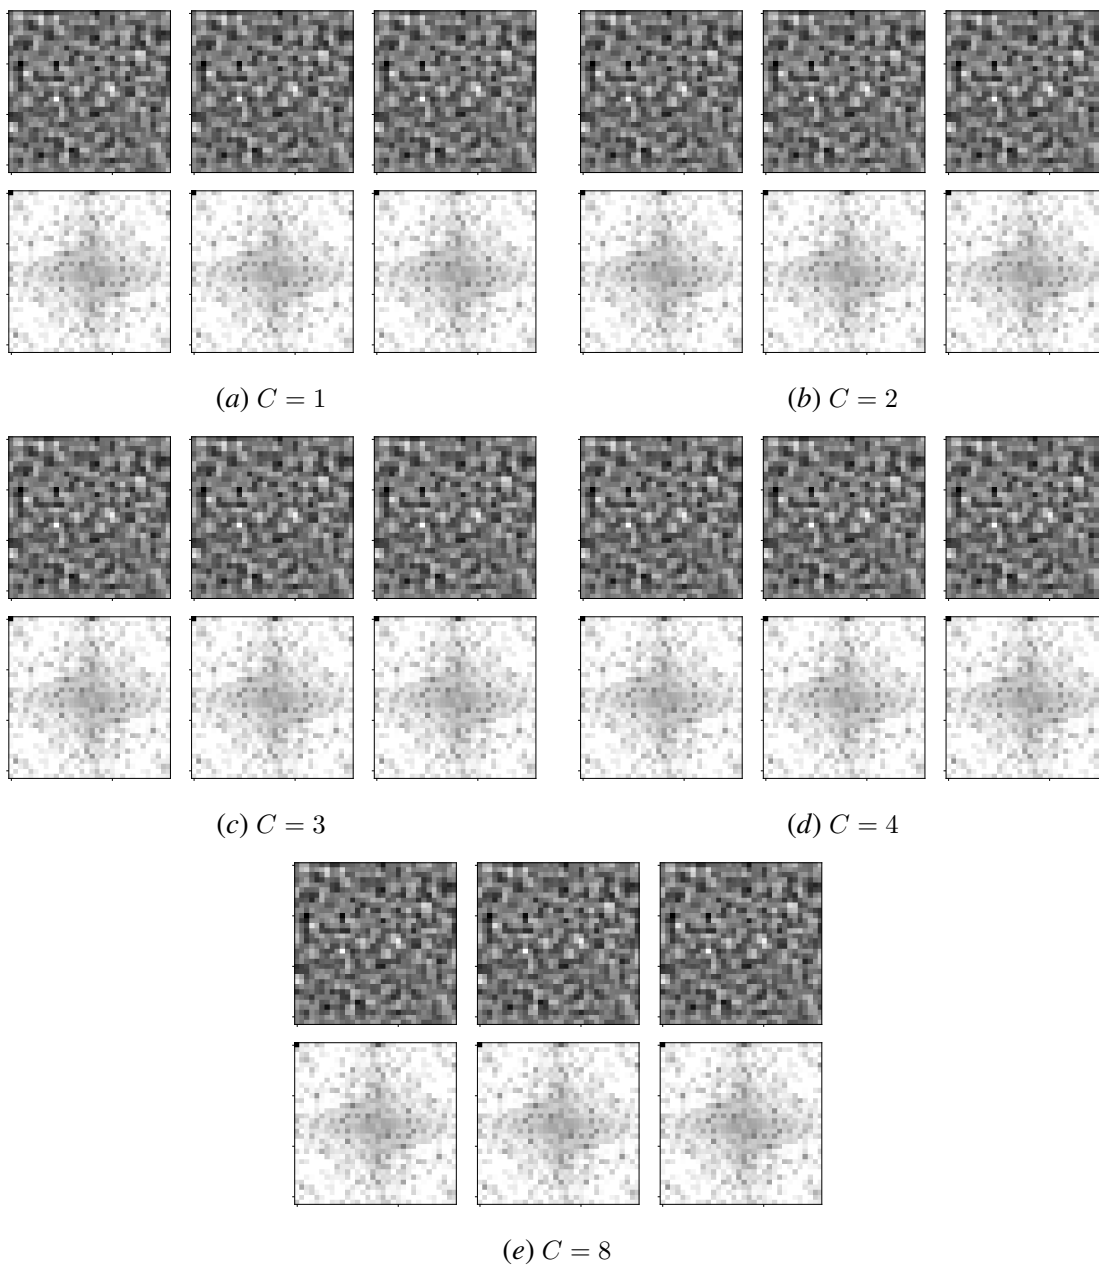


Figure 5: Linear predictors learned by two layer linear convolutional network on CIFAR-10 task. The sub-figures depict predictors learned by using gradient descent on the exponential loss for overparameterized networks with $C \in \{1, 2, 3, 4, 8\}$ and kernel size $K = (1, 1)$. The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$.




Figure 6: Linear predictors learned by two layer linear convolutional network on CIFAR-10 task. The sub-figures depict predictors learned by using gradient descent on the exponential loss for overparameterized networks with $C \in \{1, 2, 3, 4, 8\}$ and kernel size $K = (3, 3)$. The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$.

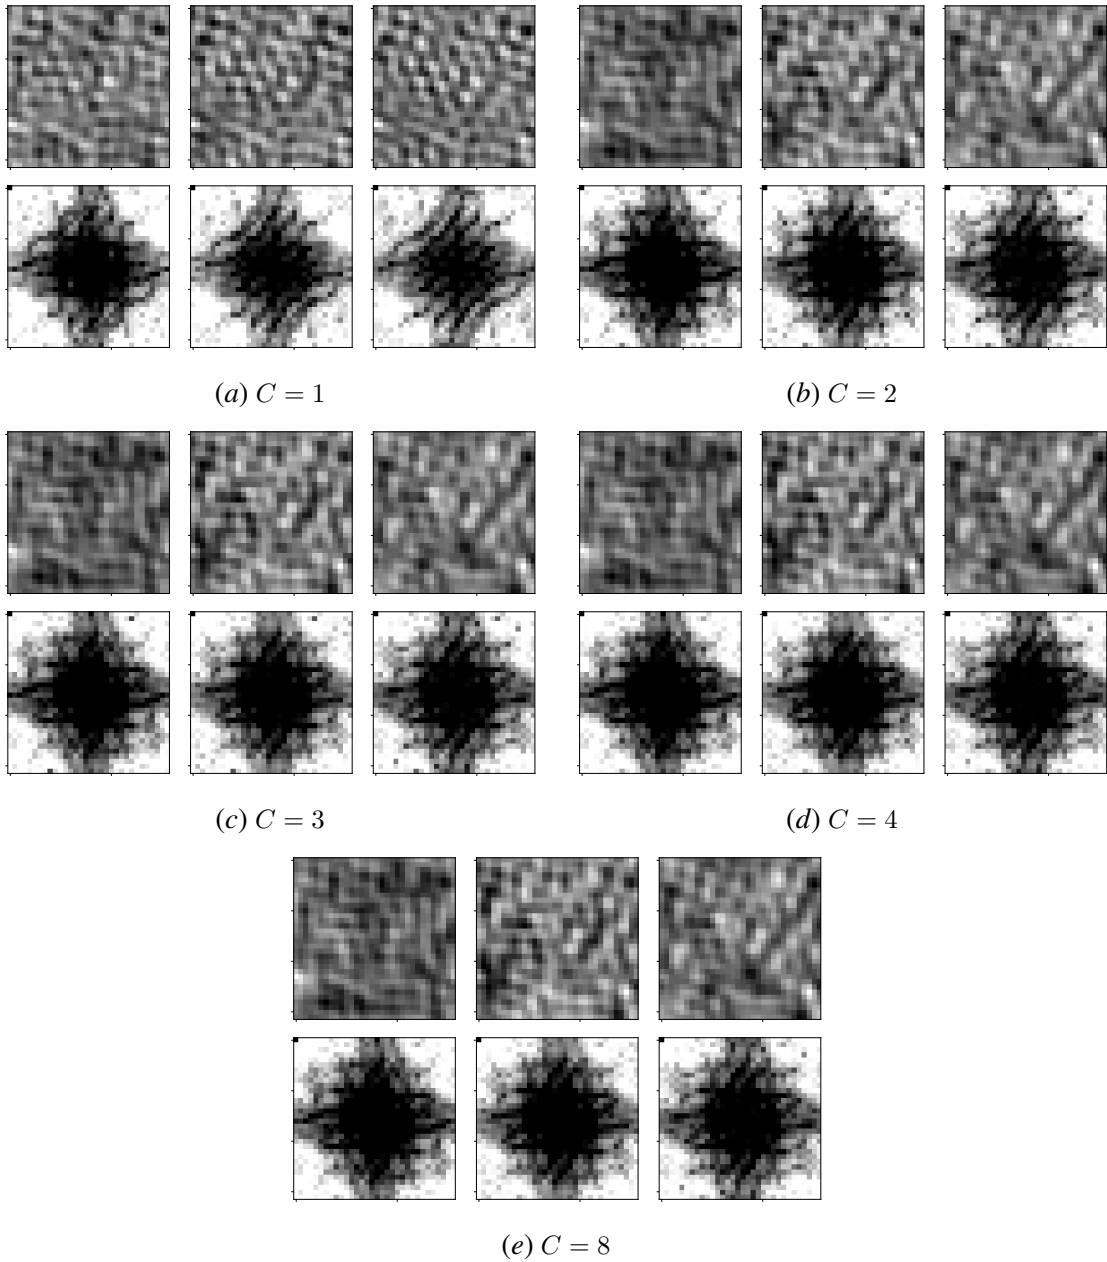


Figure 7: Linear predictors learned by two layer linear convolutional network on CIFAR-10 task. The sub-figures depict predictors learned by using gradient descent on the exponential loss for overparameterized networks with $C \in \{1, 2, 3, 4, 8\}$ and kernel size $K = (8, 8)$. The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$.

C	$K = (1, 1)$	$K = (3, 3)$	$K = (8, 8)$	$K = (20, 20)$
1	246.04	225.64	217.14	177.07
2	245.98	191.28	186.32	152.58
3	245.94	191.34	182.18	154.84
4	245.98	191.20	182.42	152.48
8	245.34	191.38	180.28	150.88

Table 5: $\widehat{\mathcal{R}}_{K,C,3}(w(\mathbf{U}, \mathbf{V})) = \|\mathbf{U}\|^2 + \|\mathbf{V}\|^2$ of the predictor learned by gradient descent on linear convolutional networks with zero padding with different number of output channels C and kernel sizes K on the CIFAR-10 task.

E.3.1. EFFECT OF KERNEL SIZE ON MNIST

Our theoretical findings in Section 3.1 suggest that the induced regularizer interpolates between ℓ_2 and ℓ_1 norms in the Fourier domain. The ℓ_2 regularization of $\mathcal{R}_{1,1}(\mathbf{w})$ does not induce sparse solutions, while the ℓ_1 regularization of $\mathcal{R}_{D,1}(\mathbf{w})$ promotes sparsity in the Fourier basis. This would suggest the following: *larger kernel sizes induce sparsity in the frequency domain.*

Explicit optimal solutions for $K = (1, 1)$ and $K = (D, D)$. First, to illustrate in the extreme cases of $K = 1$ and $K = D$, we explicitly compute the $\mathcal{R}_{K,C}$ margin predictor:

$$\min_{\mathbf{w}} \mathcal{R}_{K,C}(\mathbf{w}) \text{ s.t., } \forall_n y_n \langle \mathbf{w}, \mathbf{x}_n \rangle \geq 1$$

on the dataset using the closed-form solutions for the induced regularizer in these special cases. In Figure 4, we show resulting optimal solutions for $K = D$ (a minimum ℓ_1 solution) and $K = 1$ (a minimum ℓ_2 solution). While the solution for $K = (1, 1)$ exhibits no sparsity in the frequency domain, the solution for $K = (D, D)$ exhibits significant sparsity.

The corresponding values of $\mathcal{R}_{K,C}$ are 9.32 for $K = D$ and 2.10 for $K = 1$. We note that the induced regularizer do not exactly match those computed on the $w(\mathbf{U}, \mathbf{V})$ from gradient descent—this is because the convergence of these values can be quite slow. Moreover, for $K = D$, the difference in the predictors likely stems from the minimum ℓ_1 -norm solution being non-unique. We nonetheless show that the qualitative findings apply to gradient descent, despite the fact that the limiting values have not been reached.

Extension to gradient descent. Consider networks with one output channel $C = 1$ and compute $w(\mathbf{U}, \mathbf{V})$ learned by gradient descent for networks with different kernel sizes in Figure 3-(a). Notice in the frequency domain plots that the predictor learned with kernel size $K = (1, 1)$ is not sparse, the predictor learned with $K = (3, 3)$ already starts to exhibit some sparsity, and the linear predictor learned with $K = (28, 28)$ is highly sparse in the frequency domain.

Since sparsity in the frequency domain promotes a patterned structure in the signal domain, we explore the qualitative behavior of large kernel sizes in the signal domain in more depth. To do this, we construct an augmented version of the dataset with 112×112 dimensional images where the top-left 28×28 region is the original image, while the remaining space is all 0s. Figure 8 shows the linear predictors learned by running gradient descent on single output channel networks with different kernel sizes. As K increases, the nonzero region of the predictor becomes larger, eventually encompassing the full 112×112 dimensional space. For large kernel sizes, we can visually see that the predictors are composed of repetitions of a pattern. This is suggestive of a restricted form of periodic translation invariance, where the shift size aligns with the size of the patterns.




Figure 8: Linear predictors learned by gradient descent on single output channel networks over an augmented input space with kernel sizes $K \in \{(1, 1), (3, 3), (27, 27), (45, 45), (65, 65), (84, 84)\}$ (left to right). The input images from the MNIST dataset are augmented by padding with zeros to obtain an image of size 112×112 with the signal present only in the top-left 28×28 block.

E.3.2. EFFECT OF KERNEL SIZE ON CIFAR-10

We now examine the role of kernel size for multi-channel networks on CIFAR-10. Our theoretical findings in Section 4 suggest that the induced regularizer interpolates between the nuclear norm for $K = 1$ and the $\ell_{2,1}$ norm for $K = D$. This would again suggest the following: *larger kernel sizes induce sparsity in the frequency domain*. The behavior across input channels, however, is more nuanced. Both of these norms favor similarities across different input channels (with the effect intuitively stronger for $K = 1$ since the nuclear norm is closely related to rank). We explore both of these effects in the following experiments.

Explicit optimal solutions for $K = (D, D)$. First, to illustrate in the extreme case of $K = D$, we explicitly compute the $\mathcal{R}_{K,C}$ margin predictor: $\min_{\mathbf{w}} \mathcal{R}_{K,C}(\mathbf{w})$ s.t., $\forall_n y_n \langle \mathbf{w}, \mathbf{x}_n \rangle \geq 1$ on the dataset using the closed-form solutions for the induced regularizer in these special cases. In Figure 9, we show resulting optimal solutions for $K = D$ (a minimum $\ell_{2,1}$ solution) along with the optimal $\ell_{1,1}$ solution. We visually see that the $\ell_{2,1}$ solution favors similarity across input channels at the expense of greater sparsity in the frequency domain.

The corresponding values of $\mathcal{R}_{D,C}$ for the minimum $\ell_{2,1}$ norm solution is 82.85 for $K = (D, D)$. As for the single-input channel case, we note that the induced regularizer do not exactly match the value of 114.85 of the induced regularizer computed on the $w(\mathbf{U}, \mathbf{V})$ from gradient descent—this is because the convergence of the induced regularizer is known to be slow. We nonetheless show that the qualitative findings apply to gradient descent, despite the fact that the limiting value of induced regularizer has not been reached.

Extension to gradient descent. Consider networks with one output channel $C = 3$ and compute $w(\mathbf{U}, \mathbf{V})$ learned by gradient descent for networks with different kernel sizes in Figure 10. First, the higher kernel does indeed favor sparsity in the frequency domain, as in the single-input channel case. Nontrivial sparse structures can be observed even for $K = 3$. Next, we discuss the predictor across different input channels. Let’s first focus on the extreme case of $K = 1$. For sake of comparison, we show the explicit minimum ℓ_2 (Frobenius) predictor in Figure 9 (note that this is *not* the $\mathcal{R}_{1,C}$ margin predictor because the induced regularizer is related to the nuclear norm, not the ℓ_2 norm). As expected, we see that the learned predictor has a greater degree of similarity across channels than the ℓ_2 predictor. For other kernel sizes, Figure 10 also shows some degree of similarity across input channels, although the differences appear to grow as K becomes larger.




Figure 9: Explicit $\ell_{2,1}$, ℓ_1 and ℓ_2 margin predictors on sampled CIFAR-10 dataset. The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$




Figure 10: Linear predictors learned by two layer linear convolutional network on CIFAR-10 task. The sub-figures depict predictors learned by using gradient descent on the exponential loss for overparameterized networks with $C = 3$ and kernel sizes $K \in \{(1, 1), (3, 3), (8, 8), (20, 20), (32, 32)\}$. The top row in each sub-figure is the signal domain representation $w(\mathbf{U}, \mathbf{V})$, and the bottom row is the Fourier domain representation $\hat{w}(\mathbf{U}, \mathbf{V})$.