

Function Space and Critical Points of Linear Convolutional Networks*

Kathlén Kohn[†], Guido Montúfar[‡], Vahid Shahverdi[†], and Matthew Trager[§]

Abstract. We study the geometry of linear networks with one-dimensional convolutional layers. The function spaces of these networks can be identified with semialgebraic families of polynomials admitting sparse factorizations. We analyze the impact of the network’s architecture on the function space’s dimension, boundary, and singular points. We also describe the critical points of the network’s parameterization map. Furthermore, we study the optimization problem of training a network with the squared error loss. We prove that for architectures where all strides are larger than one and generic data, the nonzero critical points of that optimization problem are smooth interior points of the function space. This property is known to be false for dense linear networks and linear convolutional networks with stride one.

Key words. semialgebraic set, neural network, critical points

MSC codes. 68T07, 14B05, 14E99, 14J99, 14N05, 14P10, 90C23

DOI. 10.1137/23M1565504

1. Introduction. *Linear networks* are artificial neural networks with linear activation functions. Despite only representing linear functions, linear networks have been widely studied as a simplified model for analyzing the behavior of deep neural architectures. Previous work investigated various aspects of linear networks, including the optimization landscape and critical points of the loss function [Bal89, BH95, Kaw16, LK17, ZL18, LvB18, ZSEW20, TKB20, MCTH22, BH23], the dynamics of training [SMG14], and the convergence of gradient flow [BRTW21] and gradient descent [NRT21]. In particular, [TKB20] provided a detailed analysis of “pure” and “spurious” critical points, which are critical points arising from the local geometry of function space (i.e., the set of end-to-end linear functions) and the parameterization. In this context we may also mention [KTB19], which studied the geometry of the set of functions represented by networks with polynomial activation.

Linear convolutional networks (LCNs) are a type of linear network in which each linear map is a convolution. This requirement imposes linear constraints on the entries of the weight matrices—conditions sometimes known as “weight sharing” and “restricted connectivity.”

*Received by the editors April 13, 2023; accepted for publication (in revised form) January 24, 2024; published electronically May 21, 2024.

<https://doi.org/10.1137/23M1565504>

Funding: The first author was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The second author has been supported by NSF CAREER award 2145630, NSF award 2212520, DFG SPP 2298 grant 464109215, ERC Starting Grant 757983, and BMBF in DAAD project 57616814.

[†]Department of Mathematics, KTH Royal Institute of Technology, Stockholm, 10044, Sweden (kathlen@kth.se, vahidsha@kth.se).

[‡]Departments of Mathematics and Statistics, UCLA, Los Angeles, CA 90095 USA, and Max Planck Institute for Mathematics in the Sciences, Leipzig, 04103, Germany (montufar@mis.mpg.de).

[§]Amazon AWS AI Labs, New York, NY, USA. Work done outside of Amazon (matthewtrager@gmail.com).

Although convolutional neural networks (CNNs) are widely used in computer vision applications, LCNs have not received as much attention as their fully connected counterparts. In [GLSS18], LCNs were studied from the perspective of the implicit bias of local parameter optimization. That work, however, considered only nonlocal convolutions with filters of maximal size. More recent works have studied the effect on the function space of regularizing the parameters of the network, obtaining results for stride-one LCNs with arbitrary filter size fixed across layers [DKS21] and for certain two-layer multichannel LCNs [PE20, JRG22]. Most closely related to our present work is [KMMT22], which studied the geometry of the function space represented by LCNs for varying filter size sequences. That work showed that the function space of LCNs can be viewed as a semialgebraic set consisting of polynomials admitting certain factorizations. Several theoretical results on the function space were presented, including a characterization of the boundary and its singularities for LCNs with stride one. For convolutions of higher strides, it was shown that the function space is always contained in a lower-dimensional algebraic set, although many questions remained open in that case.

In this paper, we aim to fill this gap by studying the geometry of the set of functions represented by LCNs with arbitrary strides. We consider networks with an arbitrary number of layers of one-dimensional convolutions having arbitrary filter sizes. Our main results are a characterization of the dimension, boundary, and singularities of the function spaces, as a function of the network's architecture (Theorems 2.4, 2.7, 2.9, and 2.10). We also describe the critical points of the parameterization map (Theorem 2.11). Based on that description, we prove the following for architectures where all strides are larger than one: For generic training data, the nonzero critical points (in parameter space) of the squared error loss correspond to smooth interior points of the function space that are critical points of the loss on that function space (i.e., they are not induced by the network's parameterization map) (Theorem 2.12). Our results show that LCNs with arbitrary strides have a rich structure that is manifested in the geometry of certain families of polynomials with structured roots. To the best of our knowledge, these polynomial families have not been previously explored in detail and could be of independent interest.

We interpret LCNs in terms of reduced LCN architectures (with stride larger than one) composed with stride-one subarchitectures. The reduced architectures can be regarded as defining an initial set of constraints and the stride-one subarchitectures as imposing additional inequality constraints. This can be used in architecture design as it allows us to determine which function spaces are contained in each other and control the inequality constraints by choosing the stride-one subarchitectures.

Our analysis on LCNs provides insights into the complex geometric properties of neural networks, some of which may transfer to networks with nonlinear activations. Unlike densely connected feedforward linear networks, LCNs have a function space that is semialgebraic; i.e., it is a manifold with boundary (and singularities). Moreover, unlike the stride-one setting considered in [KMMT22], for arbitrary strides the function space is generally a low-dimensional subset of its ambient space. We believe that both of these qualitative aspects are important features of general neural architectures. Interestingly, increasing the stride size in LCNs also leads to nonlinear behavior, since it requires associating filters with polynomials of variables in higher degrees. Geometrically, this means that increasing the stride can be seen as

Table 1.1
Table of symbols and notation.

Notation	Description
$\mathbf{k} = (k_1, \dots, k_L)$	sequence of filter sizes for each layer $l = 1, \dots, L$
$\mathbf{s} = (s_1, \dots, s_{L-1}, 1)$	sequence of strides for each layer $l = 1, \dots, L$ with $s_L = 1$
(\mathbf{k}, \mathbf{s})	LCN architecture with filters \mathbf{k} and strides \mathbf{s}
S_l	shorthand for $\prod_{i=1}^{l-1} s_i$
$\mathcal{M}_{\mathbf{k}, \mathbf{s}}$	function space of an LCN architecture (\mathbf{k}, \mathbf{s}) , as a subset of \mathbb{R}^k with $k = \sum_{l=1}^L (k_l - 1)S_l + 1$
π_s	polynomial coefficient map $\mathbb{R}^k \rightarrow \mathbb{R}[x^s, y^s]_{k-1}$; see (2.2)
$\mu_{\mathbf{k}, \mathbf{s}}$	LCN parameterization map $\prod_i^L \mathbb{R}^{k_i} \rightarrow \mathbb{R}^k$; see (2.4)
$(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$	reduced architecture associated with (\mathbf{k}, \mathbf{s}) ; see Definition 2.5
$\overline{\mathcal{M}_{\mathbf{k}, \mathbf{s}}}$	Zariski closure of $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$
$\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$	Euclidean relative boundary of $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$

“twisting” the function space. As we discuss, the geometry of the function space plays a crucial role in optimization, since boundary points and singular points are typically more exposed during training. In LCNs, these special points correspond to functions with “more structure,” since they are functions that can be expressed by more restrictive architectures or that can be obtained as compositions of repeated filters. This sort of stratification is also characteristic of general neural networks. The algebraic nature of LCNs enables however a precise and quantitative description of this structure, revealing, for example, unexpected differences in behavior between “reduced” (where all strides are larger than one) and “nonreduced” architectures. In the future, our analysis could be further extended by considering convolutions of higher dimensions or introducing algebraic activation functions as in [KTB19].

This paper is organized as follows. In section 2 we present our main results about LCNs with arbitrary strides. We provide proofs of these results in subsequent sections. Specifically, in section 3 we prove results on the Zariski closure of the function space; in section 4 we focus on critical points of the parameterization map; in section 5 we analyze the singular points of the function space, and in section 6 we describe its boundary; finally, in section 7 we discuss the optimization of the squared loss. We keep track of the major notational concepts in Table 1.1.

2. Main results. Linear convolutional networks are families of linear maps parameterized as compositions of convolutions. In this work, we focus on single-channel convolutions for one-dimensional signals but we allow arbitrary strides. In this setting, a convolution or convolutional layer is associated with a filter $w \in \mathbb{R}^k$, a stride $s \in \mathbb{N}$, and an output dimension $d' \in \mathbb{N}$. The associated convolution is a linear map $\alpha_{w,s} : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ with input dimension $d := s(d' - 1) + k$ defined by

$$(2.1) \quad \alpha_{w,s}(\mathbf{x})[i] = \sum_{j=0}^{k-1} w[j] \cdot \mathbf{x}[is + j] \quad \text{for } \mathbf{x} \in \mathbb{R}^d \text{ and } i = 0, 1, \dots, d' - 1.$$

Note that this relation ensures that $is + j$ ranges from 0 to $d - 1$. The map (2.1) can also be represented as a generalized Toeplitz matrix $T_{w,s} \in \mathbb{R}^{d' \times d}$. However, for most of our analysis,

we will not be required to specify the input and output dimensions of convolutions, since compositions of layers can be defined purely in terms of filter vectors and strides. This will be clear from Proposition 2.2 below.

Definition 2.1. *The function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ of a linear convolutional network (LCN) architecture with filter sizes $\mathbf{k} = (k_1, \dots, k_L)$ and strides $\mathbf{s} = (s_1, \dots, s_L)$ is the set of all linear maps α that can be expressed as a composition $\alpha = \alpha_L \circ \dots \circ \alpha_1$, where α_l is a convolution of filter size k_l and stride s_l . Here L is the number of layers of the LCN.*

In the following, we assume that $k_l > 1$ for all $l = 1, \dots, L$ (layers with filter size $k_l = 1$ yield only scalar multiplication and can be discarded without loss of generality; see Proposition 2.2 below). Each stride s_l can be an arbitrary positive integer. The linear maps α in an LCN function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ are convolutions of filter size $k := k_1 + \sum_{l=2}^L (k_l - 1) \prod_{i=1}^{l-1} s_i$ and stride $s := s_1 \dots s_L$ [KMMT22, Proposition 2.2]. Since each such convolution is uniquely determined by its filter $w \in \mathbb{R}^k$, we can view $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ as a subset of \mathbb{R}^k .

To study $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$, we use the fact that compositions of convolutions can be described using polynomial multiplication. For any positive integer s and filter $w = (w[0], \dots, w[k-1]) \in \mathbb{R}^k$, we consider the polynomial

$$(2.2) \quad \pi_s(w) := w[0]x^{s(k-1)} + w[1]x^{s(k-2)}y^s + \dots + w[k-2]x^sy^{s(k-2)} + w[k-1]y^{s(k-1)} \in \mathbb{R}[x^s, y^s]_{k-1}.$$

The map π_s is an isomorphism between \mathbb{R}^k and the vector space $\mathbb{R}[x^s, y^s]_{k-1}$ of all homogeneous polynomials of degree $k-1$ in the variables (x^s, y^s) .

Proposition 2.2 ([KMMT22]). *The function space of the LCN architecture (\mathbf{k}, \mathbf{s}) can be identified with the following subset of \mathbb{R}^k :*

$$(2.3) \quad \mathcal{M}_{\mathbf{k}, \mathbf{s}} = \left\{ w \in \mathbb{R}^k : \pi_1(w) = \prod_{l=1}^L \pi_{S_l}(w_l), w_l \in \mathbb{R}^{k_l} \right\}, \quad \text{where } S_l := \prod_{i=1}^{l-1} s_i.$$

Here, π_s is the map from (2.2). Equivalently, $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is the image of the parameterization map

$$(2.4) \quad \mu_{\mathbf{k}, \mathbf{s}} : \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L} \rightarrow \mathbb{R}^k, \quad (w_1, \dots, w_L) \mapsto \pi_1^{-1} \left(\prod_{l=1}^L \pi_{S_l}(w_l) \right).$$

In light of this result, we often view $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ as a family of homogeneous polynomials admitting a sparse factorization as in (2.3). Note that the final stride s_L has no effect on the function space; for this reason, we assume from now on that $s_L = 1$. We say that an LCN architecture is *reduced* if all strides other than s_L are greater than one.

Example 2.3. Consider the architecture $\mathbf{k} = (3, 2)$ and $\mathbf{s} = (2, 1)$. Then we have that $L = 2$, $k = 5$, $S_1 = 1$, $S_2 = 2$. Thus, $\mu_{\mathbf{k}, \mathbf{s}} : \mathbb{R}^3 \times \mathbb{R}^2 \rightarrow \mathbb{R}^5$; $(w_1, w_2) \mapsto \pi_1^{-1}(\pi_2(w_2)\pi_1(w_1))$. Writing $w_1 = (w_1[0], w_1[1], w_1[2])$ and $w_2 = (w_2[0], w_2[1])$, we see that the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ consists of all $w \in \mathbb{R}^5$ with $\pi_1(w) = \pi_2(w_2)\pi_1(w_1) = (w_2[0]x^2 + w_2[1]y^2)(w_1[0]x^2 + w_1[1]x^1y^1 + w_1[2]y^2)$. Multiplying out the latter expression and collecting the coefficients of individual monomials, we see that $w = (w_2[0]w_1[0], w_2[0]w_1[1], w_2[0]w_1[2] + w_2[1]w_1[0], w_2[1]w_1[1], w_2[1]w_1[2])$. These

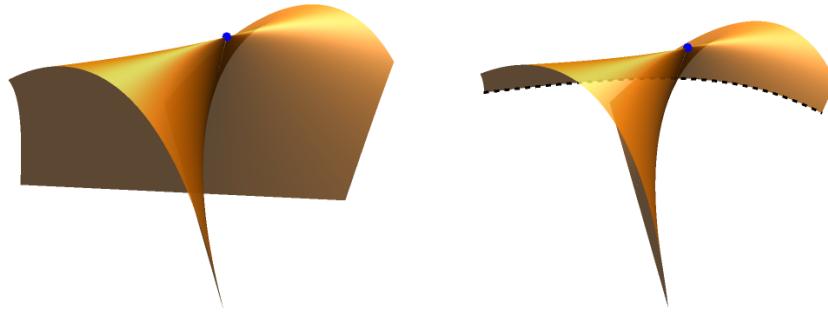


Figure 2.1. Left: Slice of the semialgebraic set $AD^2 + B^2E - BCD = 0$, $C^2 - 4AE \geq 0$, obtained by setting $A = 1$ and $C = -1$. This set corresponds to the function space $\mathcal{M}_{(3,2),(2,1)} \subseteq \mathbb{R}^5$ in Example 2.3. Right: The same set intersected with $B^4 - 4AB(BC - AD) \geq 0$, $D^4 - 4DE(CD - BE) \geq 0$ and $(AE \leq 0 \text{ or } AC \leq 0)$. This intersection corresponds to the function space $\mathcal{M}_{(2,2,2),(1,2,1)}$ discussed in Example 2.6. The reduced boundary points and the stride-one boundary points are depicted as a blue point and a black dashed curve, respectively; see Theorem 2.10.

are precisely the filters of the end-to-end convolutions represented by products $T_{w_2,s_2}T_{w_1,s_1}$ of two generalized Toeplitz matrices T_{w_1,s_1} and T_{w_2,s_2} with filter size and stride pairs $k_1 = 3, s_1 = 2$ and $k_2 = 2, s_2 = 1$, which in the concrete case of end-to-end functions $\mathbb{R}^5 \rightarrow \mathbb{R}^1$ take the form $T_{w_1,s_1} = \begin{bmatrix} w_1[0] & w_1[1] & w_1[2] & 0 & 0 \\ 0 & 0 & w_1[0] & w_1[1] & w_1[2] \end{bmatrix}$, $T_{w_2,s_2} = \begin{bmatrix} w_2[0] & w_2[1] \end{bmatrix}$. According to [KMMT22, Example 4.12], the implicit description of the function space for this architecture is given by

$$\mathcal{M}_{k,s} = \{w = (A, B, C, D, E) : AD^2 + B^2E - BCD = 0 \text{ and } C^2 - 4AE \geq 0\} \subset \mathbb{R}^5.$$

The Zariski closure $\overline{\mathcal{M}_{k,s}}$ of this function space is visualized in Figure 2.1, which displays a 3-dimensional slice.

As the previous example suggests, LCN function spaces are semialgebraic sets; that is, they are subsets of \mathbb{R}^k that are finite unions of solution sets of finitely many polynomial equalities and inequalities.

Theorem 2.4. *The LCN function space $\mathcal{M}_{k,s}$ is a semialgebraic Euclidean-closed subset of \mathbb{R}^k . Its dimension does not depend on the strides and is equal to $k_1 + \dots + k_L - (L - 1)$.*

Our main goal in this work is to investigate how the geometric properties of LCN function spaces are affected by the choice of architecture (the sequences of filter sizes and strides) and to describe how changes in the geometry impact the optimization of a training loss. The case of stride-one architectures was studied in detail in [KMMT22]. In this work, we see that the situation for arbitrary strides is considerably more complex. We focus especially on the following basic qualitative features:

- *Thick versus thin:* We say that a function space is *thick* if $\dim(\mathcal{M}_{k,s}) = k$, that is, if $\mathcal{M}_{k,s}$ is a full-dimensional semialgebraic subset of its ambient space \mathbb{R}^k or, equivalently, if its Zariski closure $\overline{\mathcal{M}_{k,s}}$ equals \mathbb{R}^k . We say that the function space is *thin* if $\dim(\mathcal{M}_{k,s}) < k$, that is, if it is contained in a proper algebraic subset of \mathbb{R}^k .
- *Zariski closed versus nonclosed:* A function space is Zariski closed (equivalently, it is an algebraic set) if it can be described using only polynomial equalities. It is Zariski

nonclosed if its characterization as a semialgebraic set necessarily involves polynomial inequalities.

- *Smooth versus singular.* We call a filter w in the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ *singular* if it is a singular point of the algebraic variety $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$. Otherwise, we say that w is *smooth*. We say the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is *smooth* if every filter $w \in \mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is smooth. Note that the Euclidean relative interior of $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is a manifold. The singular points of that manifold are contained in the singular locus of $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$. Hence, by describing all singular points in the algebraic sense (in Theorem 2.9), we find a superset of all singular points in the manifold sense.

These distinctions are relevant for the study of the optimization of training losses for LCNs. The simplest cases are those where the function space is both thick and Zariski closed. In this case, the LCN function space is equal to its ambient space \mathbb{R}^k . Thus, minimizing a smooth convex loss function on the function space is simply a convex optimization problem.

If the function space is not Zariski closed, it has a nontrivial Euclidean (relative) boundary. That boundary might be more exposed during the optimization, in the sense that many critical points of the optimization problem lie on the boundary. This happens, for instance, for stride-one LCN architectures [KMMT22].

Similarly, if the function space is singular, its singular points might be more exposed as well. This happens for dense linear networks [TKB20, NRT21].

As long as the function space is thick, its relevant boundary is the standard boundary in the Euclidean topology on \mathbb{R}^k . If the function space is thin, we need to consider its *relative boundary*, i.e., the set of points in the function space that are limit points of sequences in $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}}$.

All described geometric qualities can appear in LCN function spaces. We provide minimal example architectures in Table 2.1.

To study these geometric properties, it is useful to decompose an LCN architecture into a reduced architecture and several stride-one subarchitectures. Intuitively, we reduce an architecture by merging all neighboring layers l and $l+1$ where $s_l = 1$ (i.e., $S_{l+1} = S_l$) by multiplying their polynomial factors in (2.3). Formally, we state the following definition.

Definition 2.5. *Given an LCN architecture (\mathbf{k}, \mathbf{s}) , we define its reduced architecture as $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$, where $\tilde{\mathbf{s}} := (s_{l_1}, \dots, s_{l_{M-1}}, 1)$ is the subsequence of strides in \mathbf{s} that are greater than one (with a final stride equal to one) and $\tilde{\mathbf{k}} := (\tilde{k}_1, \dots, \tilde{k}_M)$ with $\tilde{k}_{j+1} := \sum_{i=l_j+1}^{l_{j+1}} (k_i - 1) + 1$ (and $l_0 := 0$, $l_M := L$). We define M associated stride-one architectures $(\tilde{\mathbf{k}}^j, \mathbf{1})$, with $\tilde{\mathbf{k}}^{j+1} = (k_{l_j+1}, \dots, k_{l_{j+1}})$ for $j = 0, \dots, M-1$.*

Table 2.1

Minimal examples of nontrivial LCN architectures (with at least two layers) of the different geometric types classified in Theorems 2.7 and 2.9.

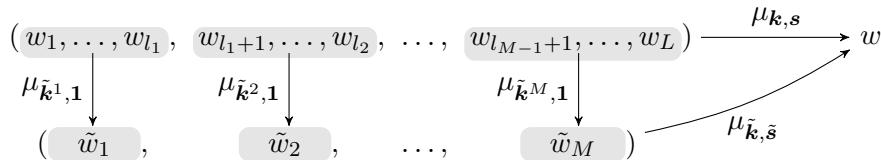
	Zariski closed	Zariski non-closed
Thick	$\mathbf{k} = (3, 2), \mathbf{s} = (1, 1)$	$\mathbf{k} = (2, 2), \mathbf{s} = (1, 1)$
Thin	$\mathbf{k} = (3, 3), \mathbf{s} = (2, 1) \mid \mathbf{k} = (2, 2), \mathbf{s} = (2, 1)$	$\mathbf{k} = (2, 2, 2), \mathbf{s} = (1, 2, 1) \mid \mathbf{k} = (3, 2), \mathbf{s} = (2, 1)$

0 is the only singular point

Note that if (\mathbf{k}, \mathbf{s}) was already reduced, then $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}}) = (\mathbf{k}, \mathbf{s})$ and $\tilde{\mathbf{k}}^j = (k_j)$ for all $1 \leq j \leq L$. The parameterization map $\mu_{\mathbf{k}, \mathbf{s}}$ of the original architecture factorizes as

$$(2.5) \quad \mu_{\mathbf{k}, \mathbf{s}} = \mu_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \circ (\mu_{\tilde{\mathbf{k}}^1, \mathbf{1}}, \dots, \mu_{\tilde{\mathbf{k}}^M, \mathbf{1}}).$$

Here the different arguments w_1, \dots, w_L of the parameterization map are assigned to their respective stride-one subarchitectures, $\mu_{\tilde{\mathbf{k}}^{j+1}, \mathbf{1}}(w_{l_j+1}, \dots, w_{l_{j+1}})$ for $j = 0, \dots, M-1$, as illustrated below.



Note that the function spaces $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ and $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ of the initial LCN architecture (\mathbf{k}, \mathbf{s}) and of its reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$ live in the same ambient space \mathbb{R}^k . We show that reducing an architecture can enlarge the function space but does not affect its Zariski closure:

$$(2.6) \quad \mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \text{ and } \overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} = \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}.$$

See Lemma 3.6. The intuition for this fact is that the function spaces of the associated stride-one architectures, which are fed to the reduced architecture, may satisfy constraints in their natural ambient spaces but are always full-dimensional. Conversely, we can interpret (2.6) as follows: The function space of a reduced architecture with filter sizes $(\tilde{k}_1, \dots, \tilde{k}_M)$ and strides $(\tilde{s}_1, \dots, \tilde{s}_{M-1}, \tilde{s}_M = 1)$ contains the function space of any architecture with filter sizes $(k_1^1, \dots, k_{m_1}^1, \dots, k_1^M, \dots, k_{m_M}^M)$ and strides $(s_1^1, \dots, s_{m_1}^1, \dots, s_1^M, \dots, s_{m_M}^M)$ satisfying $\sum_{i=1}^{m_j} (k_i^j - 1) + 1 = \tilde{k}_j$ and $s_{m_j}^j = \tilde{s}_j$ for $j = 1, \dots, M$ and $s_i^j = 1$ otherwise.

Example 2.6. Consider the architecture $\mathbf{k} = (2, 2, 2)$ and $\mathbf{s} = (1, 2, 1)$. Then $M = 2$ and $l_1 = 2$, $l_2 = 3$. The associated reduced architecture is $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}}) = ((3, 2), (2, 1))$, which we already encountered in Example 2.3. Both architectures appear in Table 2.1. The associated stride-one architectures are $(\mathbf{k}^1, \mathbf{1}) = ((2, 2), (1, 1))$ and $(\mathbf{k}^2, \mathbf{1}) = ((2), (1))$, which have output filters of sizes 3 and 2, respectively, fitting to the domain of the parameterization map of the reduced architecture. The function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is described by the equality $AD^2 + B^2E - BCD = 0$ and inequalities $(AE \leq 0 \text{ or } AC \leq 0)$, $B^4 - 4AB(BC - AD) \geq 0$, $D^4 - 4DE(CD - BE) \geq 0$, and $C^2 - 4AE \geq 0$. Nonetheless, both architectures have the same Zariski closure. This is because one function space is a subset of the other, and they are irreducible and have the same dimension by Theorem 2.4.

We are now ready to describe the qualitative features of the LCN function space.

Theorem 2.7. *Let (\mathbf{k}, \mathbf{s}) be an LCN architecture.*

- (a) *The function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is thick if and only if $\mathbf{s} = \mathbf{1}$.*
- (b) *The function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is smooth if and only if it is thick.*
- (c) *To determine whether $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed:*

- (c1) If (\mathbf{k}, \mathbf{s}) is reduced, the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed if and only if, for all $l = 1, \dots, L$, k_l is odd or $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$.
- (c2) If $\mathbf{s} = \mathbf{1}$, then $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed if and only if at most one of its filter sizes is even.
- (c3) In general, the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed if and only if the function spaces of its associated reduced architecture and of its associated stride-one architectures are all Zariski closed.

In particular, an LCN is a universal approximator of functions in the natural ambient space \mathbb{R}^k if and only if $\mathbf{s} = \mathbf{1}$ and at most one of its filter sizes is even.

Remark 2.8. The condition $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$ from the previous statement will reappear in several other results throughout the paper, e.g., in Theorem 2.9 just below. This relation has a simple interpretation: it means that the composition of the first $l - 1$ layers has stride size (equal to S_l) at least as large as its filter size (equal to $\sum_{i=1}^{l-1} (k_i - 1) S_i + 1$). This in turn implies that the “receptive fields” of the output coordinates of this convolution do not overlap; that is, the sets of input coordinates that influence each output coordinate are disjoint. As we will see, this condition also influences the geometry of the projectivized parameterization map of the function space (Remark 3.1).

If the function space is not smooth (respectively, not Zariski closed), we aim to understand its singular points (respectively, its Euclidean relative boundary). Because of (2.6), to describe the singular locus of a function space’s Zariski closure, it is sufficient to consider reduced architectures.

Theorem 2.9. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture with at least two layers. Then the singular locus of the Zariski closure of the function space is comprised of the zero filter and the union of all LCN function spaces with the same sequence of strides whose Zariski closures are proper subsets of $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$:*

$$\text{Sing}(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}) = \{0\} \cup \bigcup_{\mathbf{k}' \in K} \overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}} = \{0\} \cup \bigcup_{\mathbf{k}' \in K} \mathcal{M}_{\mathbf{k}', \mathbf{s}}, \text{ where}$$

$$K := K_{\mathbf{k}, \mathbf{s}} := \left\{ \mathbf{k}' \in \mathbb{Z}_{>0}^L : \begin{array}{l} \mathbf{k}' \neq \mathbf{k}, \sum_{i=1}^L (k'_i - 1) S_i = \sum_{i=1}^L (k_i - 1) S_i, \\ \sum_{i=l}^L (k'_i - 1) S_i \geq \sum_{i=l}^L (k_i - 1) S_i \text{ for all } l = 1, \dots, L \end{array} \right\}.$$

The set K is empty (i.e., 0 is the only singular point) if and only if $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$ for every layer l .

The conditions for \mathbf{k}' signify that the truncation of the architecture $(\mathbf{k}', \mathbf{s})$ to the first $l - 1$ layers has at most the same end-to-end filter size as the corresponding truncation of (\mathbf{k}, \mathbf{s}) and the same filter size when all layers are considered. The fact that these are precisely the architectures which satisfy $\overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}} \subsetneq \overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$ is shown in Corollary 5.3.

We denote the Euclidean relative boundary of the function space by $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$. Since the function space is Euclidean closed (see Theorem 2.4), we have $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \mathcal{M}_{\mathbf{k}, \mathbf{s}}$. Recall that $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$ consists of all points in $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ that are limits of sequences of points in $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}}$. We distinguish between two types of boundary points, using the fact that the reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$ of (\mathbf{k}, \mathbf{s}) satisfies $\mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \subseteq \overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} = \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ as in (2.6):

- *Reduced boundary points* $\partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^R$: limits in $\mathcal{M}_{\mathbf{k},\mathbf{s}}$ of sequences of points in $\overline{\mathcal{M}}_{\mathbf{k},\mathbf{s}} \setminus \mathcal{M}_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}}$.
- *Stride-one boundary points* $\partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^S$: limits in $\mathcal{M}_{\mathbf{k},\mathbf{s}}$ of sequences of points in $\mathcal{M}_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}} \setminus \mathcal{M}_{\mathbf{k},\mathbf{s}}$.

The boundary of LCN function spaces $\mathcal{M}_{\mathbf{k},\mathbf{1}}$ with stride-one architectures has been fully characterized in terms of the real-root structure of the polynomials in $\pi_1(\mathcal{M}_{\mathbf{k},\mathbf{1}})$ [KMMT22, Proposition 4.4]. The relative boundary in the case of strides larger than one is significantly more complicated. In particular, in contrast to the stride-one case, the reduced boundary points form a semialgebraic set of unexpectedly low dimension (i.e., $\dim \partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^R < \dim \mathcal{M}_{\mathbf{k},\mathbf{s}} - 1$). This can be seen in Figure 2.1, for the architecture with $\mathbf{k} = (2, 2, 2)$ and $\mathbf{s} = (1, 2, 1)$, where the reduced boundary has codimension two while the stride-one boundary has codimension one.

Theorem 2.10. *Let (\mathbf{k}, \mathbf{s}) be an LCN architecture with reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$.*

- Reduced boundary points are on the relative boundary of the reduced architecture; in fact $\partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^R = \partial\mathcal{M}_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}} \cap \mathcal{M}_{\mathbf{k},\mathbf{s}}$.*
- Reduced boundary points are contained in lower-dimensional LCN function spaces; more precisely, we have $\partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^R \subseteq \bigcup_{\tilde{\mathbf{k}}' \in \tilde{K}} \mathcal{M}_{\tilde{\mathbf{k}}',\tilde{\mathbf{s}}} \subseteq \text{Sing}(\overline{\mathcal{M}}_{\mathbf{k},\mathbf{s}})$, where $\tilde{K} := K_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}}$ is as defined in Theorem 2.9.*
- The dimension of $\partial\mathcal{M}_{\mathbf{k},\mathbf{s}}^R$ is at most $\dim \mathcal{M}_{\mathbf{k},\mathbf{s}} - \min\{s_i : s_i > 1\}$; in particular, its relative co-dimension is strictly larger than 1.*

Our discussion until this moment has focused on the implicit geometry of the function space. In practice, we are also interested in the parameterization of this space by the network's parameterization map. Critical points of a loss function can in fact arise from degenerate points of the parameterization (called “spurious critical points” in [TKB20]). We characterize these points in the following result.

Theorem 2.11. *Let $L > 1$. A filter tuple $\theta = (w_1, \dots, w_L)$ is a critical point of the parameterization map $\mu_{\mathbf{k},\mathbf{s}}$ if and only if there exists a layer $l \in \{1, \dots, L\}$ such that $w_l = 0$ or the polynomials $\pi_{S_l}(w_l)$ and $\pi_{S_{l-1}}(w_{l-1}) \cdots \pi_1(w_1)$ have a nontrivial common factor of the form $Q(x^{S_l}, y^{S_l})$. In particular, critical values w correspond to polynomials $\pi_1(w)$ in the discriminant hypersurface, that is, the set of polynomials with a double root.*

Finally, we investigate the minimization of the squared error on LCNs. Given some training data $\mathcal{D} = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) \in \mathbb{R}^{d_0} \times \mathbb{R}^{d_L} : i = 1, \dots, N\}$, the squared error loss on the function space is

$$\ell_{\mathcal{D}}(w) := \sum_{i=1}^N \|\mathbf{y}^{(i)} - T_{w,s} \mathbf{x}^{(i)}\|^2,$$

where $T_{w,s} \in \mathbb{R}^{d_L \times d_0}$ is the generalized Toeplitz matrix associated with a filter $w \in \mathbb{R}^k$ and with stride s (note that $d_0 = s(d_L - 1) + k$). When training an LCN with data \mathcal{D} , we minimize the squared error loss $\mathcal{L}_{\mathbf{k},\mathbf{s}}^{\mathcal{D}} := \ell_{\mathcal{D}} \circ \mu_{\mathbf{k},\mathbf{s}}$ on the parameter space. Commonly we use gradient descent to minimize this objective function, and thus we are interested in its critical points.

Theorem 2.12. *Let $N \geq k$. For almost all¹ data $\mathcal{D} \in (\mathbb{R}^{d_0} \times \mathbb{R}^{d_L})^N$, every critical point θ of $\mathcal{L}_{\mathbf{k},\mathbf{s}}^{\mathcal{D}}$ satisfies one of the following:*

¹For all points except those contained in some proper algebraic subset of $(\mathbb{R}^{d_0} \times \mathbb{R}^{d_L})^N$.

1. θ is a regular point of $\mu_{\mathbf{k},\mathbf{s}}$ and $\mu_{\mathbf{k},\mathbf{s}}(\theta)$ is a smooth, relative interior point of $\mathcal{M}_{\mathbf{k},\mathbf{s}}$ (i.e., $\mu_{\mathbf{k},\mathbf{s}}(\theta) \notin \text{Sing}(\overline{\mathcal{M}}_{\mathbf{k},\mathbf{s}})$ and $\mu_{\mathbf{k},\mathbf{s}}(\theta) \notin \partial\mathcal{M}_{\mathbf{k},\mathbf{s}}$), or
2. $\mu_{\mathbf{k},\mathbf{s}}(\theta) = 0$, or
3. θ is a critical point of $(\mu_{\tilde{\mathbf{k}}^1,1}, \dots, \mu_{\tilde{\mathbf{k}}^M,1})$.

Note that the last condition is not possible for reduced architectures. Hence, for reduced architectures, every critical point maps either to zero in the function space $\mathcal{M}_{\mathbf{k},\mathbf{s}}$ or to a smooth interior point of $\mathcal{M}_{\mathbf{k},\mathbf{s}}$ that is a critical point of $\ell_{\mathcal{D}}|_{\mathcal{M}_{\mathbf{k},\mathbf{s}}}$. In the language of [TKB20], these critical points are “pure,” in the sense that they are critical points in function space, rather than being degenerate points of the parameterization map. We remark that in the case of stride-one LCN architectures critical points frequently correspond to functions located on the boundary of the function space or even to functions situated in the interior of the function space as spurious (i.e., nonpure) critical points that are induced by the parameterization map; see [KMMT22, Example 5.10]. This illustrates the surprising qualitative differences between reduced and stride-one architectures.

3. The function space and its Zariski closure. From now on, to simplify notation, we omit the index of the function space $\mathcal{M} = \mathcal{M}_{\mathbf{k},\mathbf{s}}$ and its parameterization map $\mu = \mu_{\mathbf{k},\mathbf{s}}$ when the LCN architecture (\mathbf{k}, \mathbf{s}) is clear from the context. As the function space \mathcal{M} is always closed under multiplication by scalars (in other words, it forms a cone in \mathbb{R}^k), it is natural to consider its projectivization, denoted by $\mathbb{P}(\mathcal{M})$. This means that end-to-end filters that differ only by scalar multiplication are treated as equivalent. Similarly, we can projectivize the ambient space of the filters in each layer and replace the parameterization map μ with a morphism ν that composes filters up to scaling:

$$\nu = \nu_{\mathbf{k},\mathbf{s}} : \mathbb{P}_{\mathbb{R}}^{k_1-1} \times \dots \times \mathbb{P}_{\mathbb{R}}^{k_L-1} \rightarrow \mathbb{P}(\mathcal{M}_{\mathbf{k},\mathbf{s}}) \subseteq \mathbb{P}_{\mathbb{R}}^{k-1}.$$

Remark 3.1. The map ν is the composition of a Segre embedding of $\mathbb{P}_{\mathbb{R}}^{k_1-1} \times \dots \times \mathbb{P}_{\mathbb{R}}^{k_L-1}$ followed by a linear projection. Moreover, ν itself is a Segre embedding of $\mathbb{P}_{\mathbb{R}}^{k_1-1} \times \dots \times \mathbb{P}_{\mathbb{R}}^{k_L-1}$ (i.e., the linear projection is not required) if and only if $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$ for every layer l . Note that the latter condition appears in the last statement of Theorem 2.9. That statement can now be reinterpreted as follows: The projectivization of the function space is smooth if and only if it is a Segre variety (and not a proper linear projection from a Segre variety).

The projective setting has several technical advantages, including the fact that the image of a projective morphism is Zariski closed over the complex numbers [Har13b, II, section 4, Theorem 4.9]. Moreover, the morphism ν has finite fibers (see Remark 3.3), which enables us to use techniques from birational geometry.

In order to leverage these advantages, we introduce the *complex function space* $\mathcal{M}^{\mathbb{C}} = \mathcal{M}_{\mathbf{k},\mathbf{s}}^{\mathbb{C}}$, consisting of complex filters that can be factorized according to the network architecture with complex filters in each layer. We also define the complex version of the projective morphism ν :

$$\nu^{\mathbb{C}} = \nu_{\mathbf{k},\mathbf{s}}^{\mathbb{C}} : \mathbb{P}_{\mathbb{C}}^{k_1-1} \times \dots \times \mathbb{P}_{\mathbb{C}}^{k_L-1} \rightarrow \mathbb{P}(\mathcal{M}_{\mathbf{k},\mathbf{s}}^{\mathbb{C}}) \subseteq \mathbb{P}_{\mathbb{C}}^{k-1}.$$

It now follows from basic real algebraic geometry that the real Zariski closure $\overline{\mathcal{M}}$ of the real function space \mathcal{M} is the set of real points in $\mathcal{M}^{\mathbb{C}}$. We provide a formal proof of this fact and summarize these observations using the identification of filters with polynomials.

Proposition 3.2. *If π_1 is the map defined in (2.2) and $\pi_1^{\mathbb{C}}$ is its complex counterpart, then*

$$\begin{aligned}\pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}) &= \{P \in \mathbb{R}[x, y]_{k-1} : P = P_L \cdots P_1, P_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}\}, \\ \pi_1(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}) &= \{P \in \mathbb{R}[x, y]_{k-1} : P = P_L \cdots P_1, P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}\}, \\ \pi_1^{\mathbb{C}}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}) &= \pi_1^{\mathbb{C}}(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}) = \{P \in \mathbb{C}[x, y]_{k-1} : P = P_L \cdots P_1, P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}\}.\end{aligned}$$

Proof. It was already observed in [KMMT22, Remark 2.8] that an end-to-end filter $w = \mu(w_1, \dots, w_L)$ of an LCN with strides $\mathbf{s} = (s_1, \dots, s_L)$ is the coefficient vector of a polynomial with a sparse factorization as follows:

$$\pi_1(w) = \pi_{S_L}(w_L) \cdot \pi_{S_{L-1}}(w_{L-1}) \cdots \pi_{S_2}(w_2) \cdot \pi_{S_1}(w_1),$$

where $S_i = s_{i-1} \cdots s_1$ for $i > 1$ and $S_1 = 1$. That observation immediately implies the first equality of this proposition, which was already stated as Proposition 2.2 in section 2.

The equalities in the last row of the statement are the analogue over the complex numbers of the first row, using that the image $\mathbb{P}(\mathcal{M}^{\mathbb{C}})$ of the map $\nu^{\mathbb{C}}$ is Zariski closed. Finally, the last equality implies the claim in the middle row. This follows from the fact that for any subset $X \subseteq \mathbb{R}^n$, its Zariski closures $\overline{X}^{\mathbb{R}}$ inside \mathbb{R}^n and $\overline{X}^{\mathbb{C}}$ inside \mathbb{C}^n satisfy $\overline{X}^{\mathbb{R}} = \overline{X}^{\mathbb{C}} \cap \mathbb{R}^n$.

That fact can be seen as follows. The inclusion of vanishing ideals $I_{\mathbb{R}}(X) \subseteq I_{\mathbb{C}}(X)$ implies the reverse inclusion of zero loci $\overline{X}^{\mathbb{R}} \supseteq \overline{X}^{\mathbb{C}} \cap \mathbb{R}^n$. Moreover, for any polynomial $g \in I_{\mathbb{C}}(X)$ that vanishes on X , the real polynomials $\text{Real}(g) = \frac{g+\bar{g}}{2}$ and $\text{Imag}(g) = \frac{g-\bar{g}}{2i}$ also vanish on X , meaning that they also vanish on $\overline{X}^{\mathbb{R}}$ and so does $g = \text{Real}(g) + i\text{Imag}(g)$. In other words, we have $I_{\mathbb{C}}(X) \subseteq I_{\mathbb{C}}(\overline{X}^{\mathbb{R}})$, which yields $\overline{X}^{\mathbb{R}} \subseteq \overline{X}^{\mathbb{C}}$.

Now, if $X := \mathcal{M}$, then $\overline{X}^{\mathbb{R}} = \overline{\mathcal{M}}$ and $\overline{X}^{\mathbb{C}} = \mathcal{M}^{\mathbb{C}}$. To see the latter equality, note that $\mathcal{M}^{\mathbb{C}}$ contains X and since it is Zariski closed, we have $\overline{X}^{\mathbb{C}} \subseteq \mathcal{M}^{\mathbb{C}}$. To show the other inclusion, we consider a polynomial $f \in I_{\mathbb{C}}(X)$ in the complex vanishing ideal of X . Then, $f \circ \mu^{\mathbb{C}}$ vanishes on all real inputs; thus $f \circ \mu^{\mathbb{C}} = 0$. This implies $f \in I_{\mathbb{C}}(\mathcal{M}^{\mathbb{C}})$ and, therefore, $I(\overline{X}^{\mathbb{C}}) \subseteq I(\mathcal{M}^{\mathbb{C}})$. The inclusion of ideals yields the reverse inclusion of varieties, $\mathcal{M}^{\mathbb{C}} \subseteq \overline{X}^{\mathbb{C}}$, which finishes the proof. \blacksquare

Remark 3.3. One direct consequence of Proposition 3.2 is that each fiber $\nu^{-1}(w)$ or $(\nu^{\mathbb{C}})^{-1}(w)$ is finite, based on the different arrangements of the roots of the polynomial $P = \pi_1^{\mathbb{C}}(w)$ into factors $P = P_L \cdots P_1$ according to the LCN architecture.

Hence, using the projective morphism ν , we can easily compute the dimension of LCN function spaces and show that they are closed in the Euclidean topology.

Proof of Theorem 2.4. Since the function space \mathcal{M} has a polynomial parameterization, it is semialgebraic by Tarski–Seidenberg. To see that \mathcal{M} is Euclidean closed, we consider the projective spaces appearing in the map ν endowed with the quotient topology of the Euclidean topology on their underlying real vector spaces. Then these spaces are Hausdorff (unlike in the Zariski topology) and compact. Hence, since ν is continuous, its image is closed. Since the function space \mathcal{M} is the affine cone over the image of ν , it is closed in the Euclidean topology. To find the dimension of the function space, we use again the projective morphism ν . By Remark 3.3, every fiber is zero-dimensional, and hence its domain and image have the same dimension by [Har13b, II, section 3, Exercise 3.22]. We conclude that

$$\dim \mathcal{M}_{\mathbf{k}, \mathbf{s}} - 1 = \dim(\mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}})) = \dim(\mathbb{P}^{k_1-1} \times \cdots \times \mathbb{P}^{k_L-1}) = k_1 + \cdots + k_L - L. \quad \blacksquare$$

Our next goal is to prove Theorem 2.7. For that, we study the root structure of the polynomial factors in Proposition 3.2.

Definition 3.4. *Given a positive integer s , an s -hyper-root is any binomial of the form $ax^s + by^s$ with $a, b \in \mathbb{C}$, $(a, b) \neq (0, 0)$. We say that the s -hyper-root is real if $a, b \in \mathbb{R}$. An s -hyper-root R is nonreal if αR is not real for any $\alpha \in \mathbb{C} \setminus \{0\}$.*

Lemma 3.5. *If a nonreal s -hyper-root R divides a real polynomial $P \in \mathbb{R}[x, y]$, then its complex conjugate \bar{R} divides $\frac{P}{R}$.*

Proof. The statement is clear for $s = 1$. For $s > 1$, we write $R = L_1 \cdots L_s$, where $L_i \in \mathbb{C}[x, y]_1$. Since the linear factors L_i correspond to the roots of R , which was assumed to be nonreal, each L_i is nonreal and its complex conjugate \bar{L}_i does not divide R . However, \bar{L}_i must divide the real polynomial P . Hence, $\bar{R} = \bar{L}_1 \cdots \bar{L}_s$ divides $\frac{P}{R}$. ■

With the notion of hyper-roots, we now prove (2.6).

Lemma 3.6. *Let (\mathbf{k}, \mathbf{s}) be an LCN architecture and let $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$ be its associated reduced architecture. Then, $\mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ and $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} = \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$.*

Proof. We start by showing the inclusion $\mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$. By Proposition 3.2, every filter in the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ corresponds to a polynomial that can be factorized as $P = P_L \cdots P_1$, where $P_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}$. Using the notation in Definition 2.5, setting $\tilde{P}_{j+1} := P_{l_{j+1}} \cdots P_{l_j+1}$ yields a factorization $P = \tilde{P}_M \cdots \tilde{P}_1$ according to the reduced architecture, i.e., $\pi_1^{-1}(P) \in \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$. To prove the equality of Zariski closures, it suffices to show that $\overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ is a subset of $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$. By Proposition 3.2, every filter in $\overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ corresponds to a real polynomial P with a factorization $P = \tilde{P}_M \cdots \tilde{P}_1$, where $\tilde{P}_i \in \mathbb{C}[x^{\tilde{S}_i}, y^{\tilde{S}_i}]_{\tilde{k}_i+1}$. Since every complex factor \tilde{P}_i can be written as a product of \tilde{S}_i -hyper-roots, we can find a complex factorization of \tilde{P}_i according to the stride-one architecture $(\tilde{\mathbf{k}}^i, \mathbf{1})$. This yields a complex factorization of the real polynomial P according to the original architecture (\mathbf{k}, \mathbf{s}) , so $\pi_1^{-1}(P) \in \mathcal{M}_{\mathbf{k}, \mathbf{s}}$ by Proposition 3.2. ■

The final ingredient for our proof of Theorem 2.7 is to show that the zero filter is a singular point for all non-stride-one architectures.

Lemma 3.7. *If not all strides are equal to one, then the algebraic degree of $\overline{\mathcal{M}}$ is larger than one. In particular, the zero filter is a singular point of the affine cone $\overline{\mathcal{M}}$.*

Proof. It is sufficient to show that there are two filters $w, w' \in \overline{\mathcal{M}}$ such that their sum is not contained in $\overline{\mathcal{M}}$. We do this by constructing their corresponding polynomials $P = \pi_1(w)$ and $P' = \pi_1(w')$. Let l be the minimal layer such that $s_l > 1$. Note that this implies that $S_{l+1} = s_l$ and $S_j = 1$ for all $j \leq l$. We pick polynomials $P_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}$ for $i > l+1$ arbitrarily. We choose an arbitrary $Q_{l+1} \in \mathbb{R}[x^{s_l}, y^{s_l}]_{k_{l+1}-2}$ and set $P_{l+1} := x^{s_l} Q_{l+1}$. Finally, we pick $Q_l \in \mathbb{R}[x, y]_{k_l-2}$, set $P_l := y Q_l$, and choose $P_j \in \mathbb{R}[x, y]_{k_j-1}$ for $i < l$ such that the product $P_l \cdots P_1$ is not divisible by any s_l -hyper-root. Then, $P := P_L \cdots P_1 \in \pi_1(\mathcal{M})$. The second polynomial $P' := P'_L \cdots P'_1 \in \pi_1(\mathcal{M})$ is constructed by setting $P'_i := P_i$ for $i > l+1$, $P'_{l+1} := y^{s_l} Q_{l+1}$, $P'_l := x Q_l$, and $P'_j := P_j$ for $j < l$. Now we have that $P + P' = P_L \cdots P_{l+2} Q_{l+1} Q_l P_{l-1} \cdots P_1 (x^{s_l} y + xy^{s_l})$. Since $x^{s_l} y + xy^{s_l}$ is not

divisible by any s_l -hyper-root, the sum $P + P'$ does not contain enough s_l -hyper-roots to admit a factorization according to the architecture (\mathbf{k}, \mathbf{s}) , i.e., $P + P' \notin \pi_1(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}})$ by Proposition 3.2. \blacksquare

Proof of Theorem 2.7. By Theorem 2.4, the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is thick if and only if

$$(3.1) \quad k_1 + \sum_{l=2}^L (k_l - 1) = \dim(\mathcal{M}_{\mathbf{k}, \mathbf{s}}) = k = k_1 + \sum_{l=2}^L (k_l - 1)S_l.$$

Since $S_l \geq 1$ and we always assume that $k_l > 1$, the equality in (3.1) holds if and only if $S_2 = \dots = S_L = 1$. The latter is equivalent to $s_1 = \dots = s_{l-1} = 1$. Since we always assume the last stride s_L to be one, we have proven assertion (a).

For assertion (b), we observe that for thick LCN function spaces, their Zariski closures are vector spaces and thus smooth. Thin LCN function spaces have at least one stride larger than one by assertion (a) and are thus singular by Lemma 3.7.

We now prove assertion (c). Part (c2) was shown in [KMMT22, Theorem 4.1]. For parts (c1) and (c3), we observe that the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed if and only if every real polynomial P with a complex factorization $P = P_L \cdots P_1$, $P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$ admits such a factorization with real factors (this follows from Proposition 3.2).

We start by proving one direction of part (c1) and assume that every layer $l \in \{2, \dots, L\}$ satisfies that the filter size k_l is odd or $S_l > \sum_{i=1}^{l-1} (k_i - 1)S_i$. We show by induction on L that every real polynomial $P = P_L \cdots P_1$ with $P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$ admits the analogous real factorization. The base case of the induction with $L = 1$ is trivial. For $L > 1$, we consider the S_L -hyper-roots dividing P . We observe that all roots of the factor P_L correspond to S_L -hyper-roots of P , but that P might also have other S_L -hyper-roots. By Lemma 3.5, all its nonreal S_L -hyper-roots appear in complex conjugated pairs. We now distinguish two cases. First, if k_L is odd, i.e., the degree of P_L is even, then we can rearrange the S_L -hyper-roots of P into a new factorization $P = \tilde{P}_L \cdots \tilde{P}_1$ of the same format such that the new factor \tilde{P}_L becomes real. Second, the inequality $S_L > \sum_{i=1}^{L-1} (k_i - 1)S_i$ means that no S_L -hyper-root fits into the polynomial $P_{L-1} \cdots P_1$ since $\deg(P_{L-1} \cdots P_1) = \sum_{i=1}^{L-1} (k_i - 1)S_i$ is too small. Hence, P_L is the product of all S_L -hyper-roots of P and thus must be real. In either case, we have found a factorization of $P = \tilde{P}_L \cdots \tilde{P}_1$ according to the LCN architecture with \tilde{P}_L real. So $\tilde{P}_{L-1} \cdots \tilde{P}_1$ is real and we can apply the induction hypothesis to the $(L-1)$ -layer LCN that omits the last layer.

For the converse direction of (c1), we assume that the architecture (\mathbf{k}, \mathbf{s}) is reduced and that there exists a layer $l \in \{2, \dots, L\}$ with even filter size k_l and stride relation $S_l \leq \sum_{i=1}^l (k_i - 1)S_i$. We fix $P_l := (x^{S_l} + 2iy^{S_l})(x^{S_l} - 2iy^{S_l}) \cdot (x^{S_l} + 4iy^{S_l})(x^{S_l} - 4iy^{S_l}) \cdots (x^{S_l} + k_l iy^{S_l})$ of odd degree $k_l - 1$ and $R := x^{S_l} - k_l iy^{S_l}$. Note that $P_l \cdot R$ has no S -hyper-root for any $S > S_l$. As $S_l \leq \sum_{i=1}^l (k_i - 1)S_i$, we can choose $P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$ for $1 \leq i < l$ such that R divides $P_{l-1} \cdots P_1$. Moreover, we may choose the P_i such that the quotient $\frac{P_{l-1} \cdots P_1}{R}$ is real and has no S_l -hyper-root (since $S_l > S_{l-1}$). Now, we choose the factors $P_j \in \mathbb{R}[x^{S_j}, y^{S_j}]_{k_j-1}$ for $l < j \leq L$ arbitrarily. The resulting polynomial $P := P_L \cdots P_1$ is real and hence $\pi_1^{-1}(P) \in \overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$. If $P = Q_L \cdots Q_1$ is any other complex factorization of P according to the *reduced* architecture,

then $P_L \cdots P_{l+1} = Q_L \cdots Q_{l+1}$ (up to scaling). Thus, Q_l is the product of S_l -hyper-roots of $\frac{P}{P_L \cdots P_{l+1}}$, i.e., Q_l divides $P_l \cdot R$. In particular, the odd-degree factor Q_l cannot be real and so $\pi_1^{-1}(P) \notin \mathcal{M}_{\mathbf{k}, \mathbf{s}}$.

Finally, we prove (c3). We use the notation in Definition 2.5, and start by assuming that the function spaces $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ of the associated reduced architecture and $\mathcal{M}_{\tilde{\mathbf{k}}^j, \mathbf{1}}$ (for $1 \leq j \leq M$) of the associated stride-one architectures are Zariski closed. Let $P = P_L \cdots P_1$ be a real polynomial with $P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$. Then $\tilde{P}_{j+1} := P_{l_{j+1}} \cdots P_{l_j+1}$ yields a complex factorization $P = \tilde{P}_M \cdots \tilde{P}_1$ according to the reduced architecture. Since $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ is Zariski closed, we can find a real factorization $P = \tilde{Q}_M \cdots \tilde{Q}_1$ of the same format. Moreover, since $\mathcal{M}_{\tilde{\mathbf{k}}^{j+1}, \mathbf{1}}$ is Zariski closed, every real polynomial of degree $\tilde{k}_{j+1} - 1$ can be factorized into real factors according to the stride-one architecture $(\tilde{\mathbf{k}}^{j+1}, \mathbf{1})$. In particular, we can factorize $\tilde{Q}_{j+1} = Q_{l_{j+1}} \cdots Q_{l_j+1}$ such that $Q_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}$.

For the converse direction of (c3), we assume that the function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is Zariski closed. We see directly from Lemma 3.6 that the function space of the associated reduced architecture must be Zariski closed as well. We assume for contradiction that the function space $\mathcal{M}_{\tilde{\mathbf{k}}^j, \mathbf{1}}$ of one of the associated stride-one architectures is not Zariski closed. By (c2), this means that at least two of its filter sizes are even. In particular, any polynomial $\tilde{P}_j \in \mathbb{R}[x, y]_{\tilde{k}_j-1}$ with one or zero real roots (depending on the parity of its degree) cannot be factorized according to the stride-one architecture $(\tilde{\mathbf{k}}^j, \mathbf{1})$. We now fix such a polynomial \tilde{P}_j such that it has no s -hyper-root for any $s > 1$. We further pick $\tilde{P}_i \in \mathbb{R}[x, y]_{\tilde{k}_i-1}$ for $j < i \leq M$ arbitrarily and for $1 \leq i < j$ such that \tilde{P}_i has no s -hyper-root for any $s > 1$. Then $P := \tilde{P}_M(x^{\tilde{S}_M}, y^{\tilde{S}_M}) \cdots \tilde{P}_1(x, y)$ is a factorization of the real polynomial P according to the reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$. Since $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ and $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ are Zariski closed, there is a factorization $P = \tilde{Q}_M \cdots \tilde{Q}_1$ into real factors \tilde{Q}_i according to the reduced architecture such that each \tilde{Q}_i factorizes according to its associated stride-one architecture. However, by our construction, \tilde{Q}_j must be equal to \tilde{P}_j (up to scaling), which contradicts that \tilde{P}_j cannot be factorized according to $(\tilde{\mathbf{k}}^j, \mathbf{1})$. ■

4. Critical points of the parameterization. When determining the critical points of the LCN parameterization map μ , it is once again easier to work with the projective morphism ν . This is because the kernel of the differential of ν at a regular point is zero (due to the finiteness of the fibers of ν). Writing \bar{x} for the equivalence class of $x \in \mathbb{R}^{n+1} \setminus \{0\}$ in $\mathbb{P}_{\mathbb{R}}^n$, the tangent space $T_{\bar{x}}\mathbb{P}_{\mathbb{R}}^n$ is canonically isomorphic to $\text{Hom}_{\mathbb{R}}(\langle x \rangle, \mathbb{R}^{n+1}/\langle x \rangle)$ [Sha13, Example 6.24]. In our calculations below, we view that tangent space as $\mathbb{R}^{n+1}/\langle x \rangle$ ($\cong \text{Hom}_{\mathbb{R}}(\langle x \rangle, \mathbb{R}^{n+1}/\langle x \rangle)$) for simpler notation. This still captures the relevant geometry by modding out the trivial kernel vectors of the differentials of the affine parameterization map μ , as we see below in Lemma 4.1 and its proof. The underlying reason is that for the map $f: \mathbb{R}^{n+1} \setminus \{0\} \rightarrow \mathbb{P}_{\mathbb{R}}^n$ that sends x to \bar{x} , the differential $d_x f: \mathbb{R}^{n+1} \rightarrow T_{\bar{x}}\mathbb{P}_{\mathbb{R}}^n$ is surjective and its kernel is the line $\langle x \rangle$ spanned by x .

Lemma 4.1. *A filter tuple $\theta = (w_1, \dots, w_L) \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L}$ with $w_i \neq 0$ (for all i) is a critical point of μ if and only if its corresponding projective point in $\mathbb{P}^{k_1-1} \times \dots \times \mathbb{P}^{k_L-1}$ is a critical point of ν .*

Proof. The maps μ and ν and their differentials form the following commutative diagrams:

$$\begin{array}{ccc} \prod_{i=1}^L (\mathbb{R}^{k_i} \setminus \{0\}) & \xrightarrow{\mu} & \mathbb{R}^k \setminus \{0\} \\ \downarrow & & \downarrow \\ \prod_{i=1}^L \mathbb{P}^{k_i-1} & \xrightarrow{\nu} & \mathbb{P}^{k-1} \end{array} \quad \begin{array}{ccc} \prod_{i=1}^L \mathbb{R}^{k_i} & \xrightarrow{d_\theta \mu} & \mathbb{R}^k \\ \downarrow & & \downarrow \\ \prod_{i=1}^L (\mathbb{R}^{k_i} / \langle w_i \rangle) & \xrightarrow{d_{\bar{\theta}} \nu} & \mathbb{R}^k / \langle \mu(\theta) \rangle \end{array}$$

In particular, we have that $\text{im } d_{\bar{\theta}} \nu \cong (\text{im } d_\theta \mu) / \langle \mu(\theta) \rangle$, which implies the assertion. \blacksquare

Lemma 4.2. *Let $L > 1$. Every filter tuple $\theta = (w_1, \dots, w_L) \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L}$ where one of the filters w_i equals zero is a critical point of μ .*

Proof. Using our identification π_1 of filters with polynomials, the map μ becomes

$$(4.1) \quad \mu : \mathbb{R}[x^{S_1}, y^{S_1}]_{k_1-1} \times \dots \times \mathbb{R}[x^{S_L}, y^{S_L}]_{k_L-1} \rightarrow \mathbb{R}[x, y]_{k-1}, \quad (P_1, \dots, P_L) \mapsto P_L \cdots P_1,$$

and we can write the differential at $\theta = (P_1, \dots, P_L)$ as

$$\begin{aligned} d_\theta \mu : \mathbb{R}[x^{S_1}, y^{S_1}]_{k_1-1} \times \dots \times \mathbb{R}[x^{S_L}, y^{S_L}]_{k_L-1} &\rightarrow \mathbb{R}[x, y]_{k-1}, \\ (\dot{P}_1, \dots, \dot{P}_L) &\mapsto \dot{P}_L \cdot P_{L-1} \cdots P_1 + \dots + P_L \cdots P_2 \cdot \dot{P}_1. \end{aligned}$$

If one of the polynomials, say P_1 , is zero, the differential simplifies to $d_\theta \mu : (\dot{P}_1, \dots, \dot{P}_L) \mapsto P_L \cdots P_2 \cdot \dot{P}_1$, and so we obtain $\text{rank}(d_\theta \mu) \leq k_1$. Since we assume all filter sizes to be larger than one and $L \geq 2$, we conclude $\text{rank}(d_\theta \mu) \leq k_1 < k_1 + (k_2 - 1) + \dots + (k_L - 1) = \dim \mathcal{M}$. \blacksquare

Remark 4.3. Lemma 4.2 does not apply to single-layer LCN architectures. If $L = 1$, the parameterization map μ is the identity map and thus smooth.

We now compute all critical points of μ for two-layer architectures.

Proposition 4.4. *Consider a two-layer LCN architecture. A filter pair (w_1, w_2) is a critical point of μ if and only if one of the filters is zero or the polynomials $\pi_{s_1}(w_2)$ and $\pi_1(w_1)$ have a common s_1 -hyper-root.*

Proof. By Lemma 4.2, we can assume that neither w_1 nor w_2 are zero. Then, by Lemma 4.1, we can use the projective map ν to determine whether (w_1, w_2) is a critical point. Using our identification with polynomials $P_1 = \pi_1(w_1)$ and $P_2 = \pi_{s_1}(w_2)$, we can write $\nu : \mathbb{P}(\mathbb{R}[x, y]_{k_1-1}) \times \mathbb{P}(\mathbb{R}[x^{s_1}, y^{s_1}]_{k_2-1}) \rightarrow \mathbb{P}(\mathbb{R}[x, y]_{k-1})$, $(P_1, P_2) \mapsto P_2 \cdot P_1$ and

$$\begin{aligned} d_{(P_1, P_2)} \nu : \mathbb{R}[x, y]_{k_1-1} / \langle P_1 \rangle \times \mathbb{R}[x^{s_1}, y^{s_1}]_{k_2-1} / \langle P_2 \rangle &\rightarrow \mathbb{R}[x, y]_{k-1} / \langle P_1 P_2 \rangle, \\ ([\dot{P}_1], [\dot{P}_2]) &\mapsto [\dot{P}_2 P_1 + P_2 \dot{P}_1]. \end{aligned}$$

Since ν has finite fibers (see Remark 3.3), the kernel of its differential at a regular point is trivial (i.e., zero-dimensional). At a critical point (P_1, P_2) , there is a nontrivial kernel element. We start by showing that each critical point (P_1, P_2) satisfies that the polynomials P_1 and P_2 have a common s_1 -hyper-root. A nontrivial kernel element of the differential means that there is $([\dot{P}_1], [\dot{P}_2]) \neq ([P_1], [P_2])$ with $[\dot{P}_2 P_1 + P_2 \dot{P}_1] = [P_1 P_2]$. In particular, P_2 divides $\dot{P}_2 P_1$. Any factor of P_2 that is linear in x, y is part of an s_1 -hyper-root of P_2 . Moreover, if such a linear

factor divides \dot{P}_2 , its whole s_1 -hyper-root must divide \dot{P}_2 . Hence, every s_1 -hyper-root of P_2 must divide either \dot{P}_2 or P_1 . If such an s_1 -hyper-root divides P_1 , we are done. Otherwise, P_2 and \dot{P}_2 are equal up to scaling, and $[P_1 P_2] = [\dot{P}_2 P_1 + P_2 \dot{P}_1] = [P_2 \dot{P}_1]$ implies that P_1 and \dot{P}_1 are equal up to scaling as well, which contradicts that we started from a nontrivial kernel element. For the other direction, if P_1 and P_2 have a common s_1 -hyper-root $r \in \mathbb{R}[x^{s_1}, y^{s_1}]_1$, then we can write $P_2 = rQ_2$ and $P_1 = rQ_1$. We now pick $f \in \mathbb{R}[x^{s_1}, y^{s_1}]_1$ such that $\gcd(P_1 P_2, f) = 1$. Then, $([fQ_1], [-fQ_2])$ is a nontrivial kernel element of the differential $d_{(P_1, P_2)}\nu$, which shows that (P_1, P_2) is a critical point of ν . ■

To determine the critical points of the parameterization map μ for arbitrarily many layers, we start by proving a technical lemma, which provides a partial understanding of the image of the differential of μ .

Lemma 4.5. *Let $\theta = (w_1, \dots, w_L) \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L}$ be such that $w_i \neq 0$ for all i , the polynomials $\pi_{S_1}(w_1), \dots, \pi_{S_L}(w_L)$ are pairwise coprime, and their product $\pi_1(\mu(\theta)) \in \mathbb{R}[x, y]_{k-1}$ is a polynomial in x^s, y^s (i.e., $\pi_1(\mu(\theta)) \in \mathbb{R}[x^s, y^s]_{k'-1}$, where $s = s_1 \dots s_L$ is the product of the strides and $(k'-1)s = k-1$). Then we have that $\mathbb{R}[x^s, y^s]_{k'-1} \subseteq \pi_1(\text{im}(d_\theta \mu))$.*

Proof. We omit writing π_1 . Instead, we directly view μ as the polynomial multiplication map in (4.1) and work with the polynomials $P_i := \pi_{S_i}(w_i)$. We prove the assertion by induction on L . For a single layer (i.e., $L = 1$), the map μ is the identity and the assertion is trivial. For the induction step, we consider the $(L-1)$ -layer LCN architecture that omits the L th layer and denote its parameterization map by $\tilde{\mu}$. Now the tuple $\tilde{\theta} = (P_1, \dots, P_{L-1})$ satisfies the assumptions in Lemma 4.5 since $\tilde{\mu}(\tilde{\theta}) = \frac{\mu(\theta)}{P_L}$ is a polynomial in x^{S_L}, y^{S_L} . Writing $\tilde{k}-1$ for the degree of $\tilde{\mu}(\tilde{\theta})$ and $(\tilde{k}'-1)S_L = \tilde{k}-1$, we have $\tilde{\mu}(\tilde{\theta}) \in \mathbb{R}[x^{S_L}, y^{S_L}]_{\tilde{k}'-1}$. Applying the induction hypothesis yields that

$$(4.2) \quad \mathbb{R}[x^{S_L}, y^{S_L}]_{\tilde{k}'-1} \subseteq \text{im}(d_{\tilde{\theta}} \tilde{\mu}).$$

We now perform a change of variables $\tilde{x} = x^{S_L}, \tilde{y} = y^{S_L}$ and consider the map $\varphi : \mathbb{R}[\tilde{x}, \tilde{y}]_{k_L-1} \times \mathbb{R}[\tilde{x}, \tilde{y}]_{\tilde{k}'-1} \rightarrow \mathbb{R}[\tilde{x}, \tilde{y}]_{k_L+\tilde{k}'-2}$ that multiplies two polynomials. The assumptions of this lemma in particular yield that P_L and $\tilde{\mu}(\tilde{\theta})$ are coprime. Hence, by Proposition 4.4, the pair $(P_L, \tilde{\mu}(\tilde{\theta}))$ is a regular point of the map φ . This means that the differential $d_{(P_L, \tilde{\mu}(\tilde{\theta}))}\varphi$ is surjective. Thus, for any $P \in \mathbb{R}[x^s, y^s]_{k'-1} = \mathbb{R}[\tilde{x}, \tilde{y}]_{(k'-1)s_L} = \mathbb{R}[\tilde{x}, \tilde{y}]_{k_L+\tilde{k}'-2}$, there are polynomials $\dot{P}_L \in \mathbb{R}[\tilde{x}, \tilde{y}]_{k_L-1}$ and $\dot{Q} \in \mathbb{R}[\tilde{x}, \tilde{y}]_{\tilde{k}'-1}$ such that $P = \dot{P}_L \cdot \tilde{\mu}(\tilde{\theta}) + P_L \cdot \dot{Q}$. Moreover, due to (4.2), there are $\dot{P}_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}$ (for $1 \leq i \leq L-1$) such that $\dot{Q} = d_{\tilde{\theta}} \tilde{\mu}(\dot{P}_1, \dots, \dot{P}_{L-1})$. Therefore, $P = \dot{P}_L \cdot \tilde{\mu}(\tilde{\theta}) + P_L \cdot d_{\tilde{\theta}} \tilde{\mu}(\dot{P}_1, \dots, \dot{P}_{L-1}) = d_\theta \mu(\dot{P}_1, \dots, \dot{P}_L) \in \text{im}(d_\theta \mu)$. ■

This lemma enables us to show that the critical points of the parameterization map μ can be understood from the critical points of subnetworks with fewer layers. For a projective LCN parameterization map $\nu : \mathbb{P}^{k_1-1} \times \dots \times \mathbb{P}^{k_L-1} \rightarrow \mathbb{P}^{k-1}$, we denote by $\tilde{\nu} : \mathbb{P}^{k_1-1} \times \dots \times \mathbb{P}^{k_{L-1}-1} \rightarrow \mathbb{P}^{\tilde{k}-1}$ the parameterization map that is obtained by omitting the last layer. Moreover, we write $\varphi : \mathbb{P}^{k-1} \times \mathbb{P}^{k_L-1} \rightarrow \mathbb{P}^{k-1}$ for the two-layer LCN parameterization map that recovers ν from $\tilde{\nu}$, i.e., $\nu = \varphi \circ (\tilde{\nu} \times \text{id}_{\mathbb{P}^{k_L-1}})$.

Proposition 4.6. *Let $L > 1$ and let $\theta = (w_1, \dots, w_L) \in \mathbb{P}^{k_1-1} \times \dots \times \mathbb{P}^{k_L-1}$. Then θ is a critical point of ν if and only if one of the following holds:*

- (a) $\tilde{\theta} = (w_1, \dots, w_{L-1})$ is a critical point of $\tilde{\nu}$ or
 (b) $(\tilde{\nu}(\tilde{\theta}), w_L)$ is a critical point of φ .

Proof. We start by assuming that neither condition (a) nor (b) are satisfied, and we show that θ is a regular point of ν . The converse of condition (a) means that the vector space $V := \text{im } d_\theta(\tilde{\nu} \times \text{id}_{\mathbb{P}^{k_{L-1}}})$ has the expected dimension $\sum_{i=1}^L (k_i - 1) = \dim \mathbb{P}(\mathcal{M})$. Since φ has finite fibers (see Remark 3.3), the converse of condition (b) is that the differential $d_{(\tilde{\nu}(\tilde{\theta}), w_L)} \varphi$ is injective. Together with the chain rule, the converses of (a) and (b) imply that $\dim \text{im } d_\theta \nu = \dim \text{im } ((d_{(\tilde{\nu}(\tilde{\theta}), w_L)} \varphi)|_V) = \dim V = \dim \mathbb{P}(\mathcal{M})$, which shows that θ is a regular point of ν . Next, we assume that condition (a) holds, i.e., $\dim V < \dim \mathbb{P}(\mathcal{M})$. Then the chain rule yields $\dim \text{im } d_\theta \nu = \dim \text{im } ((d_{(\tilde{\nu}(\tilde{\theta}), w_L)} \varphi)|_V) \leq \dim V < \dim \mathbb{P}(\mathcal{M})$, so θ is a critical point of μ . Finally, we show that condition (b) implies that θ is a critical point of ν . This is the technical part of the proof. As in Proposition 4.4, we work directly with the polynomials $P_i := \pi_{S_i}(w_i)$ and view the maps ν , $\tilde{\nu}$, and φ as multiplying polynomials. By Proposition 4.4, condition (b) means that the polynomials $\tilde{P} := \tilde{\nu}(\tilde{\theta}) = P_1 \dots P_{L-1}$ and P_L have a common S_L -hyper-root $r \in \mathbb{P}(\mathbb{R}[x^{S_L}, y^{S_L}]_1)$. Now we distinguish two cases. First, if $r \notin \{x^{S_L}, y^{S_L}\}$, then the hyper-root factorizes as $r = r_1 \dots r_{L-1}$ such that each factor $r_i \in \mathbb{P}(\mathbb{R}[x^{S_i}, y^{S_i}])$ divides P_i and the factors r_i are pairwise coprime (as polynomials in $\mathbb{P}(\mathbb{R}[x, y])$). Hence, we can apply Lemma 4.5 to the map that multiplies the polynomials r_1, \dots, r_{L-1} : This yields that for every $f \in \mathbb{R}[x^{S_L}, y^{S_L}]_1$ there are f_1, \dots, f_{L-1} with $f_i \in \mathbb{R}[x^{S_i}, y^{S_i}]$ and $\deg(f_i) = \deg(r_i)$ such that $f_1 r_2 \dots r_{L-1} + \dots + r_1 \dots r_{L-2} f_{L-1} = f$. Thus, writing $P_L = r Q_L$ and $P_i = r_i Q_i$ for all $i < L$, we see that $([f_1 Q_1], \dots, [f_{L-1} Q_{L-1}], [-f Q_L])$ is in the kernel of the differential $d_{(P_1, \dots, P_L)} \nu$. Choosing $f \in \mathbb{R}[x^{S_L}, y^{S_L}]_1$ such that $\gcd(\tilde{P} P_L, f) = 1$ ensures that kernel element is nontrivial. Since ν has finite fibers, the existence of a nontrivial kernel element shows that $\theta = (P_1, \dots, P_L)$ is a critical point of ν . Second, if $r = x^{S_L}$ or $r = y^{S_L}$, we may assume the former without loss of generality. Since each S_i divides S_{i+1} , there are two layers i and j with $1 \leq i < j \leq L$ such that P_i and P_j have x^{S_j} as a common factor. Writing $P_i = x^{S_j} Q_i$ and $P_j = x^{S_j} Q_j$, we construct a nontrivial kernel element $([\dot{P}_1], \dots, [\dot{P}_L])$ of $d_\theta \nu$ by choosing $\dot{P}_i := y^{S_j} Q_i$, $\dot{P}_j := -y^{S_j} Q_j$, and $\dot{P}_m := 0$ for all other layers. ■

Proof of Theorem 2.11. By Lemma 4.2, we can assume that no filter w_i is zero. Then, by Lemma 4.1, we can use the projective map ν to determine whether the given filter tuple is a critical point. Now Theorem 2.11 follows by induction on L from Propositions 4.4 and 4.6. ■

Remark 4.7. The description of critical points in Theorem 2.11 is the same when working over the complex numbers instead of the reals.

5. Singular points. The goal of this section is to prove Theorem 2.9. Hence, throughout this section, we consider a reduced LCN architecture (\mathbf{k}, \mathbf{s}) . The main idea of the proof is (1) observe that the projective parameterization map $\nu^{\mathbb{C}}$ is birational for reduced architectures; (2) compute the singular points of its image (which is the projectivized complex function space) using the following method.

Fact 5.1 ([KNT17, Lemma 3.2]). *Let $f : X \rightarrow Y$ be a birational finite surjective morphism between irreducible complex projective varieties and let $y \in Y$. The variety Y is smooth at the point y if and only if the fiber $f^{-1}(y)$ contains exactly one point $x \in X$, the variety X is smooth at the point x , and the differential $d_x f : T_x X \rightarrow T_y Y$ is an injection.*

We begin by investigating the fibers of $\nu^{\mathbb{C}}$. For that, given the reduced sequence of strides \mathbf{s} , we define an \mathbf{s} -factorization of a homogeneous polynomial $P \in \mathbb{C}[x, y]$ to be a factorization of the form

$$(5.1) \quad P(x, y) = Q_L(x^{S_L}, y^{S_L}) \cdots Q_2(x^{S_2}, y^{S_2})Q_1(x^{S_1}, y^{S_1})$$

such that, for all $l \in \{L, \dots, 2\}$, the factor $Q_{l-1}(x^{S_{l-1}}, y^{S_{l-1}}) \cdots Q_1(x^{S_1}, y^{S_1})$ is not divisible by any S_l -hyper-root. Note that every homogeneous polynomial $P \in \mathbb{C}[x, y]$ has an \mathbf{s} -factorization and that it is unique (up to scaling of the factors Q_l). In particular, the degrees d_l of the factors Q_l in (5.1) are uniquely determined by P and \mathbf{s} . We refer to the sequence $\mathbf{d} = (d_L, \dots, d_1)$ as the \mathbf{s} -factor degrees of P . Those degrees give us a new perspective on LCN function spaces (see Lemma 5.2) and the smaller function spaces they contain (see Corollary 5.3), which provides us with information on the fibers of $\nu^{\mathbb{C}}$ (see Corollary 5.4).

Lemma 5.2. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture. Then*

$$\pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}) = \left\{ P \in \mathbb{C}[x, y]_{k-1} : \begin{array}{l} \text{the } \mathbf{s} \text{-factor degrees } \mathbf{d} \text{ of } P \text{ satisfy} \\ \sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i \text{ for all } l = L, \dots, 1 \end{array} \right\}.$$

Proof. By Proposition 3.2, the set $\pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ consists of polynomials $P \in \mathbb{C}[x, y]_{k-1}$ of the form $P = P_L \cdots P_1$ with $P_i \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$. If $P = Q_L \cdots Q_1$ is the \mathbf{s} -factorization of any such polynomial, then for any $l = L, \dots, 1$ we have that $P_L \cdots P_l$ divides $Q_L \cdots Q_l$. This implies $\sum_{i=l}^L d_i S_i = \deg(Q_L \cdots Q_l) \geq \deg(P_L \cdots P_l) = \sum_{i=l}^L (k_i - 1) S_i$. Conversely, if the \mathbf{s} -factor degrees of a polynomial P satisfy these inequalities, then there exists at least one factorization of P of the desired form. ■

Corollary 5.3. *Let \mathbf{s} be a reduced sequence of strides, and let $\mathbf{k}, \mathbf{k}' \in \mathbb{Z}_{>0}^L$ be such that $\sum_{i=1}^L (k_i - 1) S_i = \sum_{i=1}^L (k'_i - 1) S_i$. Then $\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} \subsetneq \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ if and only if*

$$(5.2) \quad \forall l = L, \dots, 2 : \sum_{i=l}^L (k'_i - 1) S_i \geq \sum_{i=l}^L (k_i - 1) S_i, \text{ and at least one inequality is strict.}$$

Proof. We start by assuming (5.2) and show the strict inclusion of function spaces. Let $P \in \pi_1(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$. By Lemma 5.2, the \mathbf{s} -factor degrees \mathbf{d} of P satisfy $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k'_i - 1) S_i$ for every layer l . Due to (5.2), this implies $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ for every l , which means $P \in \pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ (again by Lemma 5.2). Hence, we have shown that $\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} \subseteq \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$. We can see that the inclusion is strict by considering a polynomial P whose \mathbf{s} -factor degrees are exactly $(k_L - 1, \dots, k_1 - 1)$. Then, $P \in \pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$. However, since one of the inequalities in (5.2) is assumed to be strict, P cannot be contained in $\pi_1(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$ due to Lemma 5.2.

Now we assume $\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} \subsetneq \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$. This time, we consider a polynomial P whose \mathbf{s} -factor degrees are exactly $(k'_L - 1, \dots, k'_1 - 1)$. Then $P \in \pi_1(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subseteq \pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$, and so Lemma 5.2 implies that the inequalities in (5.2) hold. Moreover, one of the inequalities has to be strict, because otherwise \mathbf{k}' would be equal to \mathbf{k} , which would contradict that the inclusion $\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} \subsetneq \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ is strict. ■

Corollary 5.4. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture. If the fiber of a filter $w \in \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ under $\nu_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ has cardinality larger than one, then w is contained in a strictly smaller function space $\mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subsetneq \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$. In particular, the map $\nu_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ is birational.*

Proof. Let $w \in \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ be a filter that is *not* contained in any $\mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$ with $\mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subsetneq \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$. By Lemma 5.2 and Corollary 5.3, the \mathbf{s} -factor degrees of the polynomial $P = \pi_1(w)$ are $(k_L - 1, \dots, k_1 - 1)$. Hence, its \mathbf{s} -factorization in (5.1) is the unique (up to scaling) factorization of P according to the reduced architecture (\mathbf{k}, \mathbf{s}) . Therefore, the fiber of the corresponding filter w under $\nu^{\mathbb{C}}$ is a singleton. This proves the first assertion.

Since complex LCN function spaces are irreducible varieties (as they are parameterized), the strict inclusion $\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} \subsetneq \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ is equivalent to that $\dim \mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}} < \dim \mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$. Thus, the first assertion of this corollary implies that the generic fiber of $\nu^{\mathbb{C}}$ is a singleton. Therefore, $\nu^{\mathbb{C}}$ is birational (see [Har13a, Exercise 7.8]). ■

Proposition 5.5. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture with at least two layers. Then, we have that $\text{Sing}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}) = \{0\} \cup \bigcup_{\mathbf{k}' \in K} \mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$, where K is the index set from Theorem 2.9.*

Proof. By Remark 3.3, every fiber of $\nu^{\mathbb{C}}$ has a finite number of elements. Since $\nu^{\mathbb{C}}$ is a morphism between projective varieties, this means that it is a finite morphism [Har13b, III, section 11, Exercise 11.2]. Moreover, $\nu^{\mathbb{C}}$ is surjective (since $\overline{\mathcal{M}}^{\mathbb{C}} = \mathcal{M}^{\mathbb{C}}$) and birational (by Corollary 5.4). Hence, we can apply Fact 5.1 to determine the singular points of $\mathbb{P}(\mathcal{M}^{\mathbb{C}})$. Since the domain of $\nu^{\mathbb{C}}$ is smooth, we conclude that a filter w is singular in $\mathbb{P}(\mathcal{M}^{\mathbb{C}})$ if and only if w is a critical value of $\nu^{\mathbb{C}}$ or the fiber $(\nu^{\mathbb{C}})^{-1}(w)$ contains more than one element. Each singular point of the latter type is contained in some $\mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subsetneq \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ by Corollary 5.4. For every critical value w of $\nu^{\mathbb{C}}$, Theorem 2.11 (and Remark 4.7) state that the \mathbf{s} -factor degrees \mathbf{d} of the polynomial $\pi_1(w)$ satisfy at least one of the inequalities $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ (for $l = 2, \dots, L$) strictly, which shows that w is also contained in some $\mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subsetneq \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ by Corollary 5.3. Hence, so far we have shown that $\text{Sing}(\mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})) \subseteq \bigcup_{\mathbf{k}' \in K} \mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$, where K is the index set from Theorem 2.9.

For the reverse inclusion, let us consider a filter $w \in \mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}}) \subsetneq \mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$. Then, the \mathbf{s} -factor degrees \mathbf{d} of the polynomial $P = \pi_1(w)$ satisfy at least one of the inequalities $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ (for $l = 2, \dots, L$) strictly. We fix the maximal such l and consider the \mathbf{s} -factorization $P = Q_L \cdots Q_1$ in (5.1). Any factorization $P = P_L \cdots P_1$ according to the reduced architecture (\mathbf{k}, \mathbf{s}) satisfies that $P_i = Q_i$ (up to scaling) for all $i > l$, that P_l divides Q_l , and that $\frac{Q_l}{P_l}$ divides $P_{l-1} \cdots P_1$. If the factor Q_l is a power of an S_l -hyper-root, then P_l and $P_{l-1} \cdots P_1$ have such a hyper-root in common, and so w is a critical value of $\nu^{\mathbb{C}}$ by Theorem 2.11 and thus a singular point of $\mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ by Fact 5.1. Otherwise, if Q_l has at least two distinct S_l -hyper-roots, then there are at least two distinct factorizations of P according to the architecture (\mathbf{k}, \mathbf{s}) (depending on which hyper-root is dividing P_l and which $P_{l-1} \cdots P_1$). The latter means that the fiber of w under $\nu_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ has cardinality larger than one, and so w is a singular point of $\mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})$ by Fact 5.1.

Now we have shown that $\text{Sing}(\mathbb{P}(\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}})) = \bigcup_{\mathbf{k}' \in K} \mathbb{P}(\mathcal{M}_{\mathbf{k}', \mathbf{s}}^{\mathbb{C}})$. Since the architecture (\mathbf{k}, \mathbf{s}) is reduced and has at least two layers, the zero filter is a singular point of the affine cone $\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$ by Lemma 3.7. This proves the assertion. ■

To prove Theorem 2.9, it remains to transfer Proposition 5.5 to the real numbers and to understand when the index set K is empty. For the first, we investigate the \mathbf{s} -factorizations of real polynomials (see Lemma 5.6); for the latter, we make use of a technical statement on integers (see Lemma 5.8).

Lemma 5.6. *All factors in the \mathbf{s} -factorization of a real homogeneous polynomial $P \in \mathbb{R}[x, y]$ are real.*

Proof. The factor Q_L in (5.1) is the product of all S_L -hyper-roots of P . Those hyper-roots either are real or appear in complex conjugated pairs by Lemma 3.5. Hence, their product Q_L is real. Since $Q_{L-1} \cdots Q_1$ is the \mathbf{s} -factorization of $\frac{P}{Q_L}$, we see inductively that all factors Q_i are real. \blacksquare

Corollary 5.7. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture. Then*

$$\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} = \bigcup_{\mathbf{k}' \in K \cup \{\mathbf{k}\}} \mathcal{M}_{\mathbf{k}', \mathbf{s}},$$

where K is the index set from Theorem 2.9. In particular, $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}} \subseteq \bigcup_{\mathbf{k}' \in K} \mathcal{M}_{\mathbf{k}', \mathbf{s}}$.

Proof. Since $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$ is the real part of $\mathcal{M}_{\mathbf{k}, \mathbf{s}}^{\mathbb{C}}$, it corresponds to the set of real homogeneous polynomials whose \mathbf{s} -factor degrees \mathbf{d} satisfy $\sum_{i=1}^L d_i S_i \geq \sum_{i=1}^L (k_i - 1) S_i$ for all layers l . Thus, for all $\mathbf{k}' \in K$, we conclude that $\overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}} \subseteq \overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}$. This shows the inclusion “ \supseteq ” in the assertion. For the other inclusion “ \subseteq ”, let us consider a polynomial P in the complement $\pi_1(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}})$. That means that P does not have a real factorization according to the architecture (\mathbf{k}, \mathbf{s}) , but since its \mathbf{s} -factorization is real due to Lemma 5.6, its \mathbf{s} -factor degrees \mathbf{d} need to satisfy one of the inequalities $\sum_{i=1}^L d_i S_i \geq \sum_{i=1}^L (k_i - 1) S_i$ strictly. Hence, $\mathbf{k}' := (d_1 + 1, \dots, d_L + 1) \in K$. Moreover, the \mathbf{s} -factorization of P is a real factorization according to the architecture $(\mathbf{k}', \mathbf{s})$, i.e., and $P \in \pi_1(\mathcal{M}_{\mathbf{k}', \mathbf{s}})$. \blacksquare

Lemma 5.8. *Let k_1, \dots, k_{l-1} and S_1, \dots, S_l be positive integers such that S_i divides S_{i+1} for all $i \in \{1, \dots, l-1\}$. If $S_l \leq \sum_{i=1}^{l-1} (k_i - 1) S_i$, then there are integers $0 \leq e_i < k_i$ for $i \in \{1, \dots, l-1\}$ such that $S_l = \sum_{i=1}^{l-1} e_i S_i$.*

Proof. We prove the statement by induction on l . For $l = 2$, we simply put $e_1 := \frac{S_2}{S_1}$. For $l > 2$, we distinguish two cases. If $S_l \leq (k_{l-1} - 1) S_{l-1}$, then we similarly set $e_{l-1} := \frac{S_l}{S_{l-1}}$ and $e_i := 0$ for all $i < l-1$. Otherwise, we have $0 < S_l - (k_{l-1} - 1) S_{l-1} \leq \sum_{i=1}^{l-2} (k_i - 1) S_i$ and we can apply the induction hypothesis to find integers $0 \leq e_i < k_i$ for $i < l-1$ such that $S_l - (k_{l-1} - 1) S_{l-1} = \sum_{i=1}^{l-2} e_i S_i$. Setting $e_{l-1} := k_{l-1} - 1$ concludes the proof. \blacksquare

Proof of Theorem 2.9. Since $\overline{\mathcal{M}}$ is the real part of $\mathcal{M}^{\mathbb{C}}$, transferring Proposition 5.5 to the real numbers yields $\text{Sing}(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}}) = \{0\} \cup \bigcup_{\mathbf{k}' \in K} \overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}}$. Clearly, we have that $\bigcup_{\mathbf{k}' \in K} \mathcal{M}_{\mathbf{k}', \mathbf{s}} \subseteq \bigcup_{\mathbf{k}' \in K} \overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}}$. The reverse inclusion follows from Corollary 5.7. It is left to show that the index set K is empty if and only if every layer l satisfies $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$. We start by assuming that $S_l > \sum_{i=1}^{l-1} (k_i - 1) S_i$ holds for every l . We assume for contradiction that there is some $\mathbf{k}' \in K$. That means that one of the inequalities in (5.2) is strict. Let l be the maximal layer such that $\sum_{i=l}^L (k'_i - 1) S_i > \sum_{i=l}^L (k_i - 1) S_i$. Then we have that $k'_i = k_i$ for all $i > l$ and that $k'_l > k_l$. Since $\sum_{i=1}^L (k'_i - 1) S_i = \sum_{i=1}^L (k_i - 1) S_i$, we obtain that

$\sum_{i=1}^{l-1} (k_i - 1)S_i - \sum_{i=1}^{l-1} (k'_i - 1)S_i = S_l(k'_l - k_l) > S_l > \sum_{i=1}^{l-1} (k_i - 1)S_i$, which would imply that $0 > \sum_{i=1}^{l-1} (k'_i - 1)S_i$, a contradiction. Hence, the set K is empty. Finally, we assume that some layer l satisfies that $S_l \leq \sum_{i=1}^{l-1} (k_i - 1)S_i$. We now set $k'_i := k_i$ for all $i > l$, $k'_l := k_l + 1$, and $k'_j := k_j - e_j$ for all $j < l$, where the e_j are the integers found in Lemma 5.8. Then $\mathbf{k}' \in \mathbb{Z}_{>0}^L$ satisfies for every $j < l$ that $\sum_{i=j}^L (k'_i - 1)S_i - \sum_{i=j}^L (k_i - 1)S_i = S_l - \sum_{i=j}^{l-1} e_i S_i \geq$, with equality if $j = 1$. This shows that $\mathbf{k}' \in K$, so the set K is not empty. \blacksquare

6. The boundary of the function space. The boundary points of the stride-one LCN function spaces have been described in terms of the multiplicities of the real roots of polynomials in [KMMT22]. More specifically, if $P \in \mathbb{R}[x, y]_{k-1}$ is a homogenous polynomial with n distinct real roots with multiplicities $\alpha_1, \dots, \alpha_n$, then Lemma 4.2 and Proposition 4.4 in [KMMT22] state that

$$\begin{aligned} P \in \pi_1(\mathcal{M}_{\mathbf{k}, \mathbf{1}}) &\Leftrightarrow \sum_{i=1}^n \alpha_i \geq e; \\ P \in \pi_1(\partial \mathcal{M}_{\mathbf{k}, \mathbf{1}}) &\Leftrightarrow \sum_{i=1}^n \alpha_i \geq e \text{ and } |\{\alpha_i : \alpha_i \text{ is odd}\}| \leq e - 2, \end{aligned}$$

where $e := |\{k_i : k_i \text{ is even}\}|$. Here, we provide a first extension of that result to LCN architectures with larger strides. For that, we observe that, for every positive integer s , every nonzero homogeneous polynomial $P \in \mathbb{R}[x, y]$ can be uniquely factorized, up to scaling, into real homogeneous polynomials $P(x, y) = Q(x^s, y^s) \cdot R(x, y)$ such that $R \notin \mathbb{R}[x^s, y^s]$. Now the factorization of Q into real irreducible factors yields

$$P = \ell_1(x^s, y^s)^{\rho_1} \cdots \ell_r(x^s, y^s)^{\rho_r} \cdot q_1(x^s, y^s)^{\gamma_1} \cdots q_c(x^s, y^s)^{\gamma_c} \cdot R,$$

where the linear ℓ_i and the quadratic q_j correspond to the real and complex roots of Q . We call the exponents $(\rho_1, \dots, \rho_r, \gamma_1, \dots, \gamma_c)$ the *real s-hyper-root multiplicities* of P .

Theorem 6.1. *Let (\mathbf{k}, \mathbf{s}) be a two-layer reduced LCN architecture whose function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is not Zariski closed. A nonzero filter $w \in \mathbb{R}^k$ is in the relative boundary $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$ if and only if the real s_1 -hyper-root multiplicities of the polynomial $\pi_1(w)$ satisfy $\sum_{i=1}^c \gamma_i + \sum_{j=1}^r \lfloor \frac{\rho_j}{2} \rfloor \geq \frac{k_2}{2}$ and $\sum_{j=1}^r \rho_j \geq 1$.*

Proof. By Theorem 2.7, the function space being non-Zariski closed is equivalent to the condition that k_2 is even and $s_1 \leq k_1 - 1$. Consider a sequence $w^{(j)}$ of filters from the set $\overline{\mathcal{M}} \setminus \mathcal{M}$. By Proposition 3.2, each $w^{(j)}$ is associated with a polynomial $P^{(j)} \in \mathbb{R}[x, y]_{k-1}$ that can be factored as $P_2^{(j)} P_1^{(j)}$, where $P_1^{(j)} \in \mathbb{C}[x, y]_{k_1-1}$ and $P_2^{(j)} \in \mathbb{C}[x^{s_1}, y^{s_1}]_{k_2-1}$. Note that this sequence is taken from outside of the function space, so no permutation of roots can make $P_1^{(j)}$ and $P_2^{(j)}$ real polynomials. This condition is equivalent to $P^{(j)}$ containing at least k_2 s_1 -hyper-roots that all are non-real (this is because $\deg(P_2^{(j)}) = k_2 - 1$ is odd and the nonreal hyper-roots appear in complex conjugated pairs by Lemma 3.5).

Now, suppose that the sequence $P^{(j)}$ converges to a polynomial P that corresponds to a filter w in the function space \mathcal{M} . Then P should have a factorization $P_2 P_1$ with $P_1 \in \mathbb{R}[x, y]_{k_1-1}$ and $P_2 \in \mathbb{R}[x^{s_1}, y^{s_1}]_{k_2-1}$. Observe that for a polynomial in $\overline{\mathcal{M}}$, the existence of such a real factorization is equivalent to the existence of a real s_1 -hyper-root. In the limit, the

number of s_1 -hyper-roots (counted with multiplicity) cannot decrease, but real s_1 -hyper-roots can appear in two ways:

- (1) A factor in $P_1^{(j)}$ that is not an s_1 -hyper-root converges to a real s_1 -hyper-root, which means $\rho_j \geq 1$ for at least one j . Then P satisfies the relations $\sum_{j=1}^r \rho_j \geq 1$ and $\sum_{i=1}^c \gamma_i \geq \frac{k_2}{2}$.
- (2) A complex pair of s_1 -hyper-roots becomes real; i.e., $\rho_j \geq 2$ for at least one j . In this case, $\sum_{j=1}^r \rho_j \geq 2$ and $\sum_{i=1}^c \gamma_i + \sum_{j=1}^r \lfloor \frac{\rho_j}{2} \rfloor \geq \frac{k_2}{2}$.

This shows that relative boundary points correspond to polynomials satisfying $\sum_{j=1}^r \rho_j \geq 1$ and $\sum_{i=1}^c \gamma_i + \sum_{j=1}^r \lfloor \frac{\rho_j}{2} \rfloor \geq \frac{k_2}{2}$. Conversely, if a polynomial P satisfies these inequalities, we can construct a sequence in $\pi_1(\overline{\mathcal{M}} \setminus \mathcal{M})$ that converges to P by replacing double real s_1 -hyper-roots with pairs of complex conjugated ones and the remaining real s_1 -hyper-roots R with $R + \epsilon^{(j)} xy$ such that $\epsilon^{(j)} \rightarrow 0$. \blacksquare

We can use the description of the boundary points for two layers in Theorem 6.1 to prove Theorem 2.10(c) for arbitrarily many layers. First, however, we show parts (a) and (b) of that theorem, and observe that the zero filter is on the relative boundary whenever the latter is nonempty.

Lemma 6.2. *Let (\mathbf{k}, \mathbf{s}) be an LCN architecture whose function space $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ is not Zariski closed. Then $0 \in \mathbb{R}^k$ is contained in $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$.*

Proof. If a polynomial P belongs to the complement of the function space in its Zariski closure, we can generate a converging sequence $\frac{1}{n}P$, which also belongs to the complement and approaches zero in the limit. \blacksquare

Proof of Theorem 2.10(a)–(b). We start with assertion (a) that compares the reduced boundary points $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}^R \subseteq \partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$ with the relative boundary $\partial \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ of the reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$. Both types of boundary points are limits of sequences in $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}}$, with the difference that the limits land in $\mathcal{M}_{\mathbf{k}, \mathbf{s}}$ or $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$, respectively. This implies that $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}^R = \partial \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \cap \mathcal{M}_{\mathbf{k}, \mathbf{s}}$.

For assertion (b), we recall from Corollary 5.7 that $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \subseteq \bigcup_{\tilde{\mathbf{k}}' \in \tilde{K}} \mathcal{M}_{\tilde{\mathbf{k}}', \tilde{\mathbf{s}}}$, where $\tilde{K} := K_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ is the index set from Theorem 2.9. Since LCN function spaces are Euclidean closed, limits of sequences in $\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ are also contained in the finite union of the $\mathcal{M}_{\tilde{\mathbf{k}}', \tilde{\mathbf{s}}}$. In particular, this shows $\partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}^R \subseteq \bigcup_{\tilde{\mathbf{k}}' \in \tilde{K}} \mathcal{M}_{\tilde{\mathbf{k}}', \tilde{\mathbf{s}}}$. The latter union (which is empty in case the reduced architecture $(\tilde{\mathbf{k}}, \tilde{\mathbf{s}})$ has just one layer) is contained in $\text{Sing}(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}})$ by Theorem 2.9. \blacksquare

To estimate the dimension of the reduced boundary of an LCN function space, we use the two-layer description in Theorem 6.1 to deduce the necessary algebraic conditions the reduced boundary points have to satisfy.

Proposition 6.3. *Let (\mathbf{k}, \mathbf{s}) be a reduced LCN architecture. For every nonzero relative boundary point $w \in \partial \mathcal{M}_{\mathbf{k}, \mathbf{s}}$, the associated polynomial $P = \pi_1(w)$ satisfies one of the following:*

1. *P has a real double S_l -hyper-root for some $l \in \{2, \dots, L\}$,*
2. *the \mathbf{s} -factor degrees \mathbf{d} of P satisfy at least two of the inequalities $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ for $l \in \{2, \dots, L\}$ strictly, or*
3. *$\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i + 2 S_i$ for some $l \in \{2, \dots, L\}$.*

Proof. By Theorem 2.10(b), the \mathbf{s} -factor degrees \mathbf{d} of P satisfy at least one of the inequalities $\sum_{i=l}^L d_i S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ strictly. We now assume for contradiction the contrary of the assertion. That means that exactly one of these inequalities is strict, say for the layer $l \in \{2, \dots, L\}$, and that $\sum_{i=l}^L d_i S_i = \sum_{i=l}^L (k_i - 1) S_i + S_l$, as well as that P_l has no real double S_i -hyper-root for any $i \in \{2, \dots, L\}$.

Since P is taken from the relative boundary of the function space, there is a convergent sequence of polynomials $P^{(j)} \in \pi_1(\overline{\mathcal{M}}_{\mathbf{k}, \mathbf{s}} \setminus \mathcal{M}_{\mathbf{k}, \mathbf{s}})$ with limit P . Each real polynomial $P^{(j)}$ can be factored as $P_L^{(j)} \cdots P_1^{(j)}$, where $P_i^{(j)} \in \mathbb{C}[x^{S_i}, y^{S_i}]_{k_i-1}$, and it does not admit an analogous real factorization due to Proposition 3.2. However, the limit polynomial P has a real factorization $P_L \cdots P_1$ with $P_i \in \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1}$.

For any layer i , the number of S_i -hyper-roots cannot decrease in the limit. Hence, the \mathbf{s} -factor degrees $\mathbf{d}^{(j)}$ of $P^{(j)}$ satisfy $\sum_{i=m}^L d_i S_i \geq \sum_{i=m}^L d_i^{(j)} S_i \geq \sum_{i=m}^L (k_i - 1) S_i$ for every layer m . Since we assumed that $\sum_{i=m}^L d_i S_i = \sum_{i=m}^L (k_i - 1) S_i$ holds for all $m \neq l$, we have

$$(6.1) \quad d_m = d_m^{(j)} = k_m - 1 \text{ for } m \in \{L, \dots, l+1\},$$

$$(6.2) \quad d_l = k_l \geq d_l^{(j)} \geq k_l - 1,$$

$$(6.3) \quad d_{l-1} = k_{l-1} - 1 - s_{l-1} \text{ and } d_{l-1}^{(j)} + d_l^{(j)} s_{l-1} = d_{l-1} + d_l s_{l-1},$$

$$(6.4) \quad d_i = d_i^{(j)} = k_i - 1 \text{ for } i \in \{l-2, \dots, 1\}.$$

In particular, if $L > l$, we see from $d_L = d_L^{(j)}$ that every S_L -hyper-root of P is the limit of a sequence of S_L -hyper-roots in $P^{(j)}$. Moreover, $d_L = k_L - 1$ implies that P_L and $P_L^{(j)}$ are the products of the S_L -hyper-roots in P and $P^{(j)}$, respectively. Hence, $P_L^{(j)}$ converges to P_L and $P_L^{(j)}$ is real by Lemma 5.6. Applying the same argument successively for $m = L, \dots, l+1$ shows that $P_m^{(j)}$ is real and converges to P_m for all $m > l$. Analogously, we obtain from $d_{l-1}^{(j)} + d_l^{(j)} s_{l-1} = d_{l-1} + d_l s_{l-1} = (k_{l-1} - 1) + (k_l - 1) s_{l-1}$ that $P_l^{(j)} P_{l-1}^{(j)}$ is real and converges to $P_l P_{l-1}$. Finally, we conclude from (6.4) that $P_i^{(j)}$ is real and converges to P_i for all $i < l-1$.

Therefore, the factors $P_l^{(j)}$ and $P_{l-1}^{(j)}$ are nonreal and their product $P_l^{(j)} P_{l-1}^{(j)}$ does not admit a real factorization of the same format. This means that the sequence $P_l^{(j)} P_{l-1}^{(j)}$ comes from the complement $\overline{\mathcal{M}}_{\mathbf{k}', \mathbf{s}'} \setminus \mathcal{M}_{\mathbf{k}', \mathbf{s}'}$ of the two-layer architecture $(\mathbf{k}', \mathbf{s}') = ((k_{l-1}, k_l), (s_{l-1}, 1))$ (after a change of variables $(x', y') := (x^{S_{l-1}}, y^{S_{l-1}})$). Moreover, it converges to $P_l P_{l-1}$, which corresponds to a filter inside the function space $\mathcal{M}_{\mathbf{k}', \mathbf{s}'}$. By Theorem 6.1, the real s_{l-1} -hyper-root multiplicities of $P_l P_{l-1}$ satisfy $\sum_{i=1}^c \gamma_i + \sum_{j=1}^r \lfloor \frac{\rho_j}{2} \rfloor \geq \frac{k_l}{2}$ and $\sum_{j=1}^r \rho_j \geq 1$. Since we assumed that P does not have any real double S_l -hyper-root, we obtain that $P_l P_{l-1}$ has at least $\frac{k_l}{2}$ pairs of complex conjugated s_{l-1} -hyper-roots plus a real one, meaning that it has at least $k_l + 1$ many s_{l-1} -hyper-roots in total. In terms of the \mathbf{s} -factor degrees of P , this means that $d_l \geq k_l + 1$. This is a contradiction to (6.2). \blacksquare

The final ingredient for the proof of Theorem 2.10 is the following technical statement on integers that we later use to estimate the dimension difference of LCN function spaces.

Lemma 6.4. *Let $\mathbf{s} \in \mathbb{Z}_{\geq 0}^{L-1}$, $S_l := \prod_{i=1}^{l-1} s_i$ for all $l = 1, \dots, L$, and let $\mathbf{k}, \mathbf{k}' \in \mathbb{Z}^L$ be distinct integer tuples such that $\sum_{i=1}^L (k'_i - 1) S_i = \sum_{i=1}^L (k_i - 1) S_i$ and $\sum_{i=l}^L (k'_i - 1) S_i \geq \sum_{i=l}^L (k_i - 1) S_i$ holds for all $l = 2, \dots, L$. Then,*

$$(6.5) \quad \sum_{i=1}^L (k_i - k'_i) \geq \min\{s_1, \dots, s_{L-1}\} - 1.$$

Moreover, in the case that $s_i > 1$ for all i , the inequality (6.5) is strict if one of the following conditions holds:

1. at least two of the inequalities $\sum_{i=l}^L (k'_i - 1)S_i \geq \sum_{i=l}^L (k_i - 1)S_i$ for $l \in \{2, \dots, L\}$ are strict, or
2. exactly one of the inequalities is strict and $k'_l > k_l + 1$ for some $l \in \{2, \dots, L\}$.

Proof. We prove the assertion by induction on L . Since \mathbf{k} and \mathbf{k}' are assumed to be distinct, L cannot be equal to one. Thus, the base case of the induction is for $L = 2$. In that situation, we have $(k'_2 - 1)s_1 + (k'_1 - 1) = (k_2 - 1)s_1 + (k_1 - 1)$ and $k'_2 > k_2$. Hence, $k_2 - k'_2 + k_1 - k'_1 = (k'_2 - k_2)(s_1 - 1) \geq s_1 - 1$. Moreover, in the case that $s_1 > 1$, the latter inequality is strict if and only if $k'_2 > k_2 + 1$. Now we consider the case $L > 2$. If the inequalities $\sum_{i=l}^L (k'_i - 1)S_i \geq \sum_{i=l}^L (k_i - 1)S_i$ are equalities for all $i \geq 3$, then $k'_i = k_i$ for $i \geq 3$ and $k'_2 > k_2$, which means that we can argue exactly as in the induction beginning. If one of those inequalities for $i \geq 3$ is strict, we can apply the induction hypothesis to $\bar{\mathbf{s}} := (s_2, \dots, s_{L-1})$, $\bar{\mathbf{k}} := (k_2, \dots, k_L)$, and $\bar{\mathbf{k}}' := (k'_2 - \alpha, k'_3, \dots, k'_L)$, where we define $\alpha := \sum_{i=2}^L (k'_i - 1) \frac{S_i}{s_1} - \sum_{i=2}^L (k_i - 1) \frac{S_i}{s_1} \geq 0$. This yields

$$(6.6) \quad \sum_{i=2}^L (k_i - k'_i) + \alpha \geq \min\{s_2, \dots, s_{L-1}\} - 1 \geq \min\{s_1, \dots, s_{L-1}\} - 1.$$

Note that $\alpha s_1 = k_1 - k'_1$ due to $\sum_{i=1}^L (k'_i - 1)S_i = \sum_{i=1}^L (k_i - 1)S_i$. Therefore, we have that

$$(6.7) \quad \sum_{i=1}^L (k_i - k'_i) = \sum_{i=2}^L (k_i - k'_i) + \alpha + \alpha(s_1 - 1) \geq \sum_{i=2}^L (k_i - k'_i) + \alpha,$$

and (6.5) follows from (6.6). Finally, we assume that $s_i > 1$ for all i . If one of the two conditions in Lemma 6.4 holds, then either $\alpha = 0$ and we see from applying the induction hypothesis that the first inequality in (6.6) is strict, or $\alpha > 0$ and the inequality in (6.7) is strict. In either case, (6.5) is strict. \blacksquare

Proof of Theorem 2.10(c). Since $\partial\mathcal{M}_{\mathbf{k}, \mathbf{s}}^R \subseteq \partial\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ by Theorem 2.10(a), it is enough to show that

$$(6.8) \quad \dim \partial\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \leq \dim \mathcal{M}_{\mathbf{k}, \mathbf{s}} - \min\{s_i : s_i > 1\} = \dim \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} - \min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\}.$$

We first consider the case that $\min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\} > \dim \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$. This implies for every layer $m > 1$ in the reduced architecture that $\tilde{S}_m = S_{m-1}s_{m-1} > S_{m-1} \sum_{i=1}^{m-1} (\tilde{k}_i - 1) \geq \sum_{i=1}^{m-1} S_i (\tilde{k}_i - 1)$. Thus, the relative boundary of $\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$ is empty by Theorem 2.7(c1).

Hence, we assume in the following that $\min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\} \leq \dim \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}}$. We know that $\partial\mathcal{M}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} \subseteq \bigcup_{\tilde{\mathbf{k}}' \in \tilde{K}} \overline{\mathcal{M}}_{\tilde{\mathbf{k}}', \tilde{\mathbf{s}}}$ by Theorem 2.10(b). The dimension difference

$$(6.9) \quad \dim \overline{\mathcal{M}}_{\tilde{\mathbf{k}}, \tilde{\mathbf{s}}} - \dim \overline{\mathcal{M}}_{\tilde{\mathbf{k}}', \tilde{\mathbf{s}}} = \sum_{i=1}^M (\tilde{k}_i - \tilde{k}'_i) \geq \min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\} - 1$$

is estimated for all $\tilde{\mathbf{k}}' \in \tilde{K}$ in Lemma 6.4.

We now distinguish between the three different types of boundary points described in Proposition 6.3. The points in $\partial\mathcal{M}_{\tilde{k},\tilde{s}}$ of the second type are contained in the union of $\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$ whose dimension difference in (6.9) is strict (due to the first condition in the second part of Lemma 6.4). The points in $\partial\mathcal{M}_{\tilde{k},\tilde{s}}$ that are not of the second type but of the third type in Proposition 6.3 are also contained in the union of $\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$ whose dimension difference in (6.9) is strict (due to the second condition in Lemma 6.4). Thus, denoting the set of points in $\partial\mathcal{M}_{\tilde{k},\tilde{s}}$ that are of the second or third type by $\partial\mathcal{M}_{\tilde{k},\tilde{s}}^S$, we conclude that $\dim(\partial\mathcal{M}_{\tilde{k},\tilde{s}}^S) \leq \dim\overline{\mathcal{M}}_{\tilde{k},\tilde{s}} - \min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\}$.

Finally, we consider the set of points in $\partial\mathcal{M}_{\tilde{k},\tilde{s}}$ whose associated polynomials have a real double S_l -hyper-root for some $S_l > 1$. We write $\Delta_{\tilde{k},\tilde{s}}$ for the Zariski closure of that set in \mathbb{R}^k . Since $\Delta_{\tilde{k},\tilde{s}} \subseteq \bigcup_{\tilde{k}' \in \tilde{K}} \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$, we have that $\Delta_{\tilde{k},\tilde{s}} = \bigcup_{\tilde{k}' \in \tilde{K}} (\Delta_{\tilde{k},\tilde{s}} \cap \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}})$. By definition, $\Delta_{\tilde{k},\tilde{s}}$ is contained in the discriminant hypersurface that describes polynomials with double roots. However, not every filter in $\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$ corresponds to a polynomial with a double root, and so $\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$ is not contained in the discriminant hypersurface. Hence, $\Delta_{\tilde{k},\tilde{s}} \cap \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}} \subsetneq \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}$. Since the latter is an irreducible variety, this shows that $\dim(\Delta_{\tilde{k},\tilde{s}} \cap \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}) < \dim(\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}})$. Therefore, we conclude that $\dim\Delta_{\tilde{k},\tilde{s}} = \max_{\tilde{k}' \in \tilde{K}} \dim(\Delta_{\tilde{k},\tilde{s}} \cap \overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}) \leq \dim(\overline{\mathcal{M}}_{\tilde{k}',\tilde{s}}) - 1 \leq \dim\overline{\mathcal{M}}_{\tilde{k},\tilde{s}} - \min\{\tilde{s}_1, \dots, \tilde{s}_{M-1}\}$, where the latter inequality comes from (6.9).

Since $\partial\mathcal{M}_{\tilde{k},\tilde{s}} \subseteq \{0\} \cup \partial\mathcal{M}_{\tilde{k},\tilde{s}}^S \cup \Delta_{\tilde{k},\tilde{s}}$, we have proven (6.8). ■

7. Optimization. In this section, we prove Theorem 2.12. We start by expressing the squared error loss $\ell_{\mathcal{D}}(w) = \sum_{i=1}^N \|\mathbf{y}^{(i)} - T_{w,s}\mathbf{x}^{(i)}\|^2$ directly in terms of the filters w instead of first passing to Toeplitz matrices $T_{w,s}$. For that, we collect the training data $\mathcal{D} = \{(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathbf{x}^{(N)}, \mathbf{y}^{(N)})\} \subseteq \mathbb{R}^{d_0} \times \mathbb{R}^{d_L}$ into two matrices $X \in \mathbb{R}^{d_0 \times N}$ and $Y \in \mathbb{R}^{d_L \times N}$ whose columns are $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$ and $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}$, respectively, and write $\ell_{\mathcal{D}}(w) = \|Y - T_{w,s}X\|_F^2$, where $\|\cdot\|_F$ is the Frobenius norm. We next consider the linear map

$$\chi_{X,s} : \mathbb{R}^k \rightarrow \mathbb{R}^{d_L \times N}, \quad w \mapsto T_{w,s}X.$$

We let \bar{Y} denote the orthogonal projection of Y onto the image of $\chi_{X,s}$ and choose a filter $u_Y \in \mathbb{R}^k$ such that $\chi_{X,s}(u_Y) = \bar{Y}$. With this, we can write the squared error loss as

$$(7.1) \quad \ell_{\mathcal{D}}(w) = \|Y - \bar{Y}\|_F^2 + \|\bar{Y} - T_{w,s}X\|_F^2 = \|Y - \bar{Y}\|_F^2 + \|\chi_{X,s}(u_Y - w)\|_F^2.$$

In this expression, \bar{Y} and thus $\|Y - \bar{Y}\|_F^2$ only depend on the data \mathcal{D} and (k, s) , but not on the filter w . Hence, minimizing $\ell_{\mathcal{D}}(w)$ is equivalent to minimizing $\|\chi_{X,s}(u_Y - w)\|_F^2$. We observe that $\|\cdot\|_{X,s} := \|\chi_{X,s}(\cdot)\|_F$ is a seminorm on \mathbb{R}^k that is induced by the (possibly degenerate) inner product $\langle w_1, w_2 \rangle_{X,s} := \langle \chi_{X,s}(w_1), \chi_{X,s}(w_2) \rangle_F = \text{tr}((T_{w_1,s}X)^\top T_{w_2,s}X)$.

Lemma 7.1. *The inner product $\langle \cdot, \cdot \rangle_{X,s}$ is nondegenerate (i.e., $\|\cdot\|_{X,s}$ is a norm on \mathbb{R}^k) if and only if the linear map $\chi_{X,s}$ is injective. Moreover, if $N \geq k$, then $\chi_{X,s}$ is injective for almost all $X \in \mathbb{R}^{d_0 \times N}$.*

Proof. By definition of the inner product $\langle \cdot, \cdot \rangle_{X,s}$, we see that $\langle w, w \rangle_{X,s} = 0$ if and only if $\chi_{X,s}(w) = 0$, which shows the first part of the assertion. For the second part, we write $\bar{x}^{(i)} \in \mathbb{R}^k$ for the first k entries of the i th column of X . The condition $\chi_{X,s}(w) = 0$ implies in particular that $w \in \mathbb{R}^k$ is orthogonal (with respect to the standard Euclidean inner product)

to each of the vectors $\bar{x}^{(1)}, \dots, \bar{x}^{(N)}$. Hence, if $N \geq k$, almost all choices of $\bar{x}^{(1)}, \dots, \bar{x}^{(N)}$ force w to be zero. \blacksquare

Corollary 7.2. *Let $N \geq k$. For almost all $X \in \mathbb{R}^{d_0 \times N}$, minimizing the squared loss $\ell_D(w)$ is equivalent to minimizing the squared inner product norm $\|u_Y - w\|_{X,s}^2$, where u_Y is the unique filter such that $\chi_{X,s}(u_Y)$ is the orthogonal projection (with respect to the Frobenius norm) of Y onto the image of $\chi_{X,s}$.*

Proof. This follows immediately from (7.1) and Lemma 7.1. \blacksquare

In the following, we fix an LCN architecture (\mathbf{k}, \mathbf{s}) and write $\mu = \mu_{\mathbf{k}, \mathbf{s}}$ and $\langle \cdot, \cdot \rangle_X := \langle \cdot, \cdot \rangle_{X,s}$. In light of Corollary 7.2, we assume from now on that $N \geq k$ and that $X \in \mathbb{R}^{d_0 \times N}$ is such that minimizing ℓ_D is equivalent to minimizing the squared norm $\|u_Y - \cdot\|_X^2$. To prove Theorem 2.12, it is now sufficient to show that, for a fixed X and for almost every data filter $u \in \mathbb{R}^k$, every critical point of $\mathcal{L}_{u,X}(\theta) := \|u - \mu(\theta)\|_X^2$ satisfies one of the three conditions in Theorem 2.12.

A filter tuple θ is a critical point of $\mathcal{L}_{u,X}$ if and only if the data filter u is contained in the *normal space*

$$(7.2) \quad N_X(\theta) := \text{im}(d_\theta \mu)^{\perp_X} + \mu(\theta) \subseteq \mathbb{R}^k,$$

where \perp_X denotes the orthogonal complement with respect to the inner product $\langle \cdot, \cdot \rangle_X$. Hence, to prove that a fixed set Θ of filter tuples does not contain any critical point of $\mathcal{L}_{u,X}$ for almost all data filters u , our proof strategy is to show that the union of the normal spaces $N_X(\theta)$ over all $\theta \in \Theta$ is contained in a proper algebraic subset of \mathbb{R}^k .

Definition 7.3. *We say that a semialgebraic subset $\Theta \subseteq \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L}$ is exposed with respect to $\langle \cdot, \cdot \rangle_X$ if*

$$\dim \left(\bigcup_{\theta \in \Theta} N_X(\theta) \right) = k.$$

Example 7.4. If $L \geq 2$, then $\mu^{-1}(0)$ is exposed. Indeed, for any filter tuple θ with at least two zero filters, $d_\theta \mu = 0$ and thus $\dim(N_X(\theta)) = k$.

Lemma 7.5. *Let $\Theta \subseteq (\mathbb{R}^{k_1} \setminus \{0\}) \times \dots \times (\mathbb{R}^{k_L} \setminus \{0\})$ be a semialgebraic subset that is an affine cone; i.e., for every $\theta \in \Theta$ and $\lambda \in (\mathbb{R} \setminus \{0\})^L$ we have that $(\lambda_1 \theta_1, \dots, \lambda_L \theta_L) \in \Theta$. Then $\dim(\bigcup_{\theta \in \Theta} N_X(\theta)) \leq \dim \Theta - L + 1 + k - \min_{\theta \in \Theta} \text{rank}(d_\theta \mu)$.*

Proof. Let us start by fixing a filter tuple $\theta \in \Theta$. For $\lambda \in (\mathbb{R} \setminus \{0\})^L$, we write $\lambda \cdot \theta := (\lambda_1 \theta_1, \dots, \lambda_L \theta_L)$. Since for every such λ we have $\mu(\lambda \cdot \theta) = \lambda_1 \dots \lambda_L \mu(\theta)$ and $\text{im}(d_{\lambda \cdot \theta} \mu) = \text{im}(d_\theta \mu)$, we see that $\dim(\bigcup_{\lambda \in (\mathbb{R} \setminus \{0\})^L} N_X(\lambda \cdot \theta)) \leq \dim N_X(\theta) + 1$. Hence,

$$\begin{aligned} \dim \left(\bigcup_{\theta \in \Theta} N_X(\theta) \right) &= \dim \left(\bigcup_{\bar{\theta} \in \mathbb{P}(\Theta)} \bigcup_{\lambda \in (\mathbb{R} \setminus \{0\})^L} N_X(\lambda \cdot \bar{\theta}) \right) \\ &\leq \dim \mathbb{P}(\Theta) + \max_{\bar{\theta} \in \mathbb{P}(\Theta)} \dim \left(\bigcup_{\lambda \in (\mathbb{R} \setminus \{0\})^L} N_X(\lambda \cdot \bar{\theta}) \right) \leq \dim \Theta - L + \max_{\theta \in \Theta} (\dim N_X(\theta) + 1). \end{aligned}$$

Since the dimension of the normal space $N_X(\theta)$ is $k - \text{rank}(d_\theta \mu)$, the assertion follows. \blacksquare

To estimate the rank of the differential of μ , we investigate the S_l -hyper-roots that the filters in a tuple θ have in common.

Definition 7.6. Let $\theta = (w_1, \dots, w_L) \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L}$ with $w_l \neq 0$ and $P_l := \pi_{S_l}(w_l)$ for all l . Set $G_l(\theta) := \gcd(P_l, P_{l-1} \cdots P_1) \in \mathbb{R}[x^{S_l}, y^{S_l}]$ for $l \in \{2, \dots, L\}$, and $G_1(\theta) := 1$. The common hyper-root degree of θ is $\text{chd}(\theta) := \sum_{l=2}^L \deg(G_l(\theta))$.

Recall that Theorem 2.11 states that θ with $\mu(\theta) \neq 0$ is a critical point of μ if and only if $\text{chd}(\theta) \geq 1$.

Proposition 7.7. For $\theta \in (\mathbb{R}^{k_1} \setminus \{0\}) \times \dots \times (\mathbb{R}^{k_L} \setminus \{0\})$, we have that $\text{rank}(d_\theta \mu) \geq \dim(\mathcal{M}) - \text{chd}(\theta)$.

Proof. Using our identification π_1 of filters with polynomials, we consider $\theta = (P_1, \dots, P_L)$ as a tuple of polynomials as in (4.1). For every layer l , we define $P'_l := \frac{P_l}{G_l(\theta)}$. Then P'_l and the product $P'_{l-1} \cdots P'_1$ are coprime. Hence, $\theta' := (P'_1, \dots, P'_L)$ is a regular point of $\mu' := \mu_{k', s}$ where $k'_l := \deg P'_l + 1$ (this follows from Theorem 2.11 after omitting all layers l with $k'_l = 1$). We can see $d_{\theta'} \mu'$ as a restriction of $d_\theta \mu$ via the following commutative diagram:

$$\begin{array}{ccc} \prod_{i=1}^L \mathbb{R}[x^{S_i}, y^{S_i}]_{k_i-1} & \xrightarrow{d_\theta \mu} & \mathbb{R}[x, y]_{k-1} \\ \varphi \uparrow & & \uparrow \psi \\ \prod_{i=1}^L \mathbb{R}[x^{S_i}, y^{S_i}]_{k'_i-1} & \xrightarrow{d_{\theta'} \mu'} & \mathbb{R}[x, y]_{k'-1} \end{array}$$

where $\varphi : (\dot{P}'_1, \dots, \dot{P}'_L) \mapsto (\dot{P}'_1 G_1(\theta), \dots, \dot{P}'_L G_L(\theta))$ and $\psi : \dot{P}' \mapsto \dot{P}' G_1(\theta) \cdots G_L(\theta)$. Therefore, we conclude that $\text{rank}(d_\theta \mu) \geq \text{rank}(d_{\theta'} \mu') = \dim(\mathcal{M}_{k', s}) = \sum_{i=1}^L (k'_i - 1) + 1 = \sum_{i=1}^L (k_i - 1 - \deg G_i(\theta)) + 1 = \dim(\mathcal{M}_{k, s}) - \text{chd}(\theta)$. \blacksquare

We now aim to show that the critical points of μ (except those in $\mu^{-1}(0)$) are not exposed. For that, we stratify that set of critical points as the disjoint union (over all $\delta \in \mathbb{Z}_{>0}$) of $C_\delta := \{\theta \in (\mathbb{R}^{k_1} \setminus \{0\}) \times \dots \times (\mathbb{R}^{k_L} \setminus \{0\}) \mid \text{chd}(\theta) = \delta\}$.

Proposition 7.8. Let the architecture (k, s) be reduced. If $\delta > 0$ and $C_\delta \neq \emptyset$, then we have that $\text{codim}(C_\delta) > \delta$.

Proof. We prove the assertion by induction on the number L of layers and, as above, identify filters with polynomials. For single-layer architectures, $\text{chd}(\theta) = 0$ for every nonzero θ , and thus there is nothing to show. Hence, the induction beginning is $L = 2$. In that case,

$$C_\delta = \left\{ (RQ_1, RQ_2) \mid \begin{array}{l} R \in \mathbb{R}[x^{s_1}, y^{s_1}]_\delta, Q_2 \in \mathbb{R}[x^{s_1}, y^{s_1}]_{k_2-1-\delta}, Q_1 \in \mathbb{R}[x, y]_{k_1-1-\delta s_1} \\ R \neq 0, Q_2 \neq 0, Q_1 \neq 0, \gcd(Q_2, Q_1) = 1 \end{array} \right\}$$

has dimension $k_2 + (k_1 - \delta s_1)$. Therefore, the codimension of C_δ is δs_1 , which is larger than δ due to $s_1 > 1$ and $\delta > 0$.

For the induction step, we assume $L > 2$. For any partition $\bar{\delta} = (\delta_2, \dots, \delta_L) \in \mathbb{Z}_{\geq 0}^{L-1}$ of δ , we define $C_{\bar{\delta}} := \{\theta \in C_\delta \mid \forall l = 2, \dots, L : \deg G_l(\theta) = \delta_l\}$. Then $C_\delta = \bigcup_{\bar{\delta}} C_{\bar{\delta}}$, where the union runs over all nonnegative partitions of δ into $L - 1$ parts. Hence, it is enough to show for every such partition $\bar{\delta}$ that either $\text{codim}(C_{\bar{\delta}}) > \delta$ or $C_{\bar{\delta}} = \emptyset$. We fix a partition $\bar{\delta}$ with $C_{\bar{\delta}} \neq \emptyset$

and distinguish two cases. First, if $\delta_L = 0$, then for every $\theta' \in C_{(\delta_2, \dots, \delta_{L-1})}$ we have that almost every $P_L \in \mathbb{R}[x^{S_L}, y^{S_L}]_{k_L-1}$ gives a point $(\theta', P_L) \in C_{\bar{\delta}}$. Thus, the induction hypothesis yields $\dim C_{\bar{\delta}} = \dim C_{(\delta_2, \dots, \delta_{L-1})} + k_L < (k_1 + \dots + k_{L-1} - \delta) + k_L$.

Second, if $\delta_L > 0$, we start by observing that

$$(7.3) \quad \dim C_{\bar{\delta}} < \dim C_{(\delta_2, \dots, \delta_{L-1})}.$$

Indeed, for every $(P_1, \dots, P_{L-1}, RQ_L) \in C_{(\delta_2, \dots, \delta_{L-1})}$ with $R \in \mathbb{R}[x^{S_L}, y^{S_L}]_{\delta_L-1}$ dividing the product $P_1 \cdots P_{L-1}$, almost every $Q'_L \in \mathbb{R}[x^{S_L}, y^{S_L}]_{k_L-\delta_L}$ yields a new point $(P_1, \dots, P_{L-1}, RQ'_L) \in C_{(\delta_2, \dots, \delta_{L-1})}$. Hence, every irreducible component of the Zariski closure $\overline{C}_{(\delta_2, \dots, \delta_{L-1})}$ is of the form $\Sigma \times \mathbb{R}^{k_L-\delta_L}$. However, no such component is contained in the Zariski closure $\overline{C}_{\bar{\delta}}$, because the latter imposes an algebraic condition on Q'_L . Since $\overline{C}_{\bar{\delta}} \subseteq \overline{C}_{(\delta_2, \dots, \delta_{L-1})}$ and no irreducible component of $\overline{C}_{(\delta_2, \dots, \delta_{L-1})}$ is equal to $\overline{C}_{\bar{\delta}}$, we have shown (7.3). Applying that inequality δ_L times, we obtain $\dim C_{\bar{\delta}} \leq \dim C_{(\delta_2, \dots, \delta_{L-1}, 0)} - \delta_L$. Now, invoking the first case (where δ_L was assumed to be zero), we get $\dim C_{(\delta_2, \dots, \delta_{L-1}, 0)} < \sum_{i=1}^L k_i - \sum_{j=1}^{L-1} \delta_j$. Putting the last two inequalities together, we conclude that $\dim C_{\bar{\delta}} < \sum_{i=1}^L k_i - \sum_{j=1}^{L-1} \delta_j$. ■

Theorem 7.9. *If the architecture (\mathbf{k}, \mathbf{s}) is reduced, $\text{Crit}^\circ(\mu) := \{\theta \in \text{Crit}(\mu) \mid \mu(\theta) \neq 0\}$ is not exposed.*

Proof. Since $\text{Crit}^\circ(\mu)$ is the disjoint union over all C_δ with $\delta \geq 1$, we have that

$$(7.4) \quad \dim \left(\bigcup_{\theta \in \text{Crit}^\circ(\mu)} N_X(\theta) \right) = \dim \left(\bigcup_{\delta \geq 1} \bigcup_{\theta \in C_\delta} N_X(\theta) \right) = \max_{\delta \geq 1} \dim \left(\bigcup_{\theta \in C_\delta} N_X(\theta) \right).$$

Applying Lemma 7.5 to each nonempty C_δ in the union, we obtain $\dim \left(\bigcup_{\theta \in C_\delta} N_X(\theta) \right) \leq \dim C_\delta - L + 1 + k - \min_{\theta \in C_\delta} \text{rank}(d_\theta \mu)$. Therefore, Propositions 7.7 and 7.8 yield that $\dim \left(\bigcup_{\theta \in C_\delta} N_X(\theta) \right) < (\sum_{i=1}^L k_i - \delta) - L + 1 + k - (\dim \mathcal{M} - \delta) = k$. Since C_δ is nonempty only for finitely many choices of δ , the latter inequality shows that (7.4) is less than k . ■

Corollary 7.10. *Let $\mathcal{Z} \subseteq \mathcal{M}$ be a semialgebraic subset that is an affine cone with $\dim(\mathcal{Z}) < \dim(\mathcal{M})$. If the architecture (\mathbf{k}, \mathbf{s}) is reduced, then $\mu^{-1}(\mathcal{Z} \setminus \{0\})$ is not exposed. In particular, this statement holds for $\mathcal{Z} = \partial \mathcal{M}$ or $\mathcal{Z} = \mathcal{M} \cap \text{Sing}(\overline{\mathcal{M}})$.*

Proof. $\Theta := \mu^{-1}(\mathcal{Z} \setminus \{0\})$ is the disjoint union of $\Theta_C := \Theta \cap \text{Crit}(\mu) = \Theta \cap \text{Crit}^\circ(\mu)$ and $\Theta_R := \{\theta \in \Theta \mid \theta \notin \text{Crit}(\mu)\}$. By Theorem 7.9, Θ_C is not exposed. Hence, it is left to show that Θ_R is not exposed either. Since every $\theta \in \Theta_R$ is a regular point of μ , the rank of the differential $d_\theta \mu$ is equal to $\dim \mathcal{M}$. Thus, applying Lemma 7.5 to Θ_R yields $\dim \left(\bigcup_{\theta \in \Theta_R} N_X(\theta) \right) \leq \dim \Theta - L + 1 + k - \dim \mathcal{M} = \dim \mathcal{Z} + k - \dim \mathcal{M} < k$. ■

Theorem 7.9 and Corollary 7.10 imply Theorem 2.12 for reduced LCN architectures. The general version follows from the following observation.

Lemma 7.11. *Let θ be a regular point of $(\mu_{\tilde{\mathbf{k}}^1, 1}, \dots, \mu_{\tilde{\mathbf{k}}^M, 1})$. Then $N_X(\theta) = N_X(\tilde{\theta})$, where $\tilde{\theta} := (\mu_{\tilde{\mathbf{k}}^1, 1}, \dots, \mu_{\tilde{\mathbf{k}}^M, 1})(\theta)$.*

Proof. Since each $\mu_{\tilde{\mathbf{k}}^i, 1}$ is the parameterization map of a stride-one LCN, its function space is thick. Hence, at a regular point θ of $(\mu_{\tilde{\mathbf{k}}^1, 1}, \dots, \mu_{\tilde{\mathbf{k}}^M, 1})$, the image of the differential

of $(\mu_{\tilde{\mathbf{k}}^1,1}, \dots, \mu_{\tilde{\mathbf{k}}^M,1})$ is equal to the domain of the differential of $\mu_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}}$ at $\tilde{\theta}$. Therefore, (2.5) implies $\text{im}(d_{\theta}\mu_{\mathbf{k},\mathbf{s}}) = \text{im}(d_{\tilde{\theta}}\mu_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}})$, and the assertion follows. \blacksquare

Proof of Theorem 2.12. We consider the set of filter tuples θ that do not satisfy any of the three conditions in Theorem 2.12. More concretely, writing

$$\Theta := \{\theta \in \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_L} : \mu_{\mathbf{k},\mathbf{s}}(\theta) \neq 0, \theta \notin \text{Crit}((\mu_{\tilde{\mathbf{k}}^1,1}, \dots, \mu_{\tilde{\mathbf{k}}^M,1}))\},$$

that set is the union of

$$\Theta_C := \Theta \cap \text{Crit}(\mu_{\mathbf{k},\mathbf{s}}), \Theta_S := \Theta \cap \mu_{\mathbf{k},\mathbf{s}}^{-1}(\text{Sing}(\overline{\mathcal{M}}_{\mathbf{k},\mathbf{s}})), \text{ and } \Theta_B := \Theta \cap \mu_{\mathbf{k},\mathbf{s}}^{-1}(\partial \mathcal{M}_{\mathbf{k},\mathbf{s}}).$$

It is sufficient to show that $\Theta_C \cup \Theta_S \cup \Theta_B$ is not exposed with respect to any $\langle \cdot, \cdot \rangle_X$ that is an inner product. For $\diamond \in \{C, S, B\}$, we set $\tilde{\Theta}_{\diamond} := (\mu_{\tilde{\mathbf{k}}^1,1}, \dots, \mu_{\tilde{\mathbf{k}}^M,1})(\Theta_{\diamond})$. Lemma 7.11 shows that $\bigcup_{\theta \in \Theta_{\diamond}} N_X(\theta) = \bigcup_{\tilde{\theta} \in \tilde{\Theta}_{\diamond}} N_X(\tilde{\theta})$. Hence, it is enough to show that none of the $\tilde{\Theta}_{\diamond}$ is exposed with respect to $\langle \cdot, \cdot \rangle_X$. Since $\tilde{\Theta}_C \subseteq \text{Crit}^{\circ}(\mu_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}})$, we have that $\tilde{\Theta}_C$ is not exposed by Theorem 7.9. Moreover, $\tilde{\Theta}_S \cup \tilde{\Theta}_B \subseteq \mu_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}}^{-1}(\mathcal{Z} \setminus \{0\})$, where $\mathcal{Z} := (\mathcal{M}_{\mathbf{k},\mathbf{s}} \cap \text{Sing}(\overline{\mathcal{M}}_{\mathbf{k},\mathbf{s}})) \cup \partial \mathcal{M}_{\mathbf{k},\mathbf{s}}$. Thus, applying Corollary 7.10 to $\mathcal{Z} \subseteq \mathcal{M}_{\tilde{\mathbf{k}},\tilde{\mathbf{s}}}$, we conclude that $\tilde{\Theta}_S \cup \tilde{\Theta}_B$ is not exposed. \blacksquare

REFERENCES

- [Bal89] P. BALDI, *Linear learning: Landscapes and algorithms*, in Advances in Neural Information Processing Systems 1, D. S. Touretzky, ed., Morgan-Kaufmann, 1989, pp. 65–72.
- [BH95] P. BALDI AND K. HORNIK, *Learning in linear neural networks: A survey*, Trans. Neural Netw., 6 (1995), pp. 837–858.
- [BH23] A. BHARADWAJ AND S. HOŞTEN, *Complex Critical Points of Deep Linear Neural Networks*, [arXiv:2301.12651](https://arxiv.org/abs/2301.12651), 2023.
- [BRTW21] B. BAH, H. RAUHUT, U. TERSTIEGE, AND M. WESTDICKENBERG, *Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers*, Inf. Inference, 11 (2021), pp. 307–353.
- [DKS21] Z. DAI, M. KARZAND, AND N. SREBRO, *Representation costs of linear neural networks: Analysis and design*, in Advances in Neural Information Processing Systems 34, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, eds., 2021.
- [GLSS18] S. GUNASEKAR, J. D. LEE, D. SOUDRY, AND N. SREBRO, *Implicit bias of gradient descent on linear convolutional networks*, in Advances in Neural Information Processing Systems 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds., Curran Associates, 2018, pp. 9461–9471.
- [Har13a] J. HARRIS, *Algebraic Geometry: A First Course*, Grad. Texts in Math. 133, Springer, 2013.
- [Har13b] R. HARTSHORNE, *Algebraic Geometry*, Grad. Texts in Math. 52, Springer, 2013.
- [JRG22] M. JAGADEESAN, I. RAZENSHTEYN, AND S. GUNASEKAR, *Inductive bias of multi-channel linear convolutional networks with bounded weight norm*, in Proceedings of Thirty Fifth Conference on Learning Theory, Proc. Mach. Learn. Res. 178, P.-L. Loh and M. Raginsky, eds., PMLR, 2022, pp. 2276–2325.
- [Kaw16] K. KAWAGUCHI, *Deep learning without poor local minima*, in Advances in Neural Information Processing Systems 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, eds., Curran Associates, 2016, pp. 586–594.
- [KMMT22] K. KOHN, T. MERKH, G. MONTÚFAR, AND M. TRAGER, *Geometry of linear convolutional networks*, SIAM J. Appl. Algebra Geom., 6 (2022), pp. 368–406, <https://doi.org/10.1137/21M1441183>.
- [KNT17] K. KOHN, B. I. U. NØDLAND, AND P. TRIPOLI, *Secants, bitangents, and their congruences*, in Combinatorial Algebraic Geometry, Springer, 2017, pp. 87–112.

- [KTB19] J. KILEEL, M. TRAGER, AND J. BRUNA, *On the expressive power of deep polynomial neural networks*, in Advances in Neural Information Processing Systems 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, eds., Curran Associates, 2019.
- [LK17] H. LU AND K. KAWAGUCHI, *Depth Creates No Bad Local Minima*, [abs/1702.08580](https://arxiv.org/abs/1702.08580), 2017.
- [LvB18] T. LAURENT AND J. VON BRECHT, *Deep linear networks with arbitrary loss: All local minima are global*, in Proceedings of the 35th International Conference on Machine Learning, Proc. Mach. Learn. Res. 80, J. Dy and A. Krause, eds., PMLR, Stockholm, Sweden, 2018, pp. 2902–2907.
- [MCTH22] D. MEHTA, T. CHEN, T. TANG, AND J. D. HAUENSTEIN, *The loss surface of deep linear networks viewed through the algebraic geometry lens*, IEEE Trans. Pattern Anal. Mach. Intell., 44 (2022), pp. 5664–5680.
- [NRT21] G. M. NGUEGNANG, H. RAUHUT, AND U. TERSTIEGE, *Convergence of Gradient Descent for Learning Linear Neural Networks*, [arXiv:2108.02040](https://arxiv.org/abs/2108.02040), 2021.
- [PE20] M. PILANCI AND T. 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, Proc. Mach. Learn. Res. 119, H. Daumé III and A. Singh, eds., PMLR, 2020, pp. 7695–7705.
- [Sha13] I. R. SHAFAREVICH, *Basic Algebraic Geometry 2: Schemes and Complex Manifolds*, Springer, Berlin, Heidelberg, 2013.
- [SMG14] A. M. SAXE, J. L. MCCLELLAND, AND S. GANGULI, *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*, in 2nd International Conference on Learning Representations, Conference Track Proceedings, Y. Bengio and Y. LeCun, eds., 2014.
- [TKB20] M. TRAGER, K. KOHN, AND J. BRUNA, *Pure and spurious critical points: A geometric study of linear networks*, in International Conference on Learning Representations, 2020.
- [ZL18] Y. ZHOU AND Y. LIANG, *Critical points of linear neural networks: Analytical forms and landscape properties*, in International Conference on Learning Representations, 2018.
- [ZSEW20] Z. ZHU, D. SOUDRY, Y. C. ELDAR, AND M. B. WAKIN, *The global optimization geometry of shallow linear neural networks*, J. Math. Imaging Vis., 62 (2020), pp. 279–292.