THE IMPLICIT BIAS OF GRADIENT DESCENT ON GENERALIZED GATED LINEAR NETWORKS

Samuel Lippl[1], **L.F. Abbott**[1], and **SueYeon Chung**[1,2]

[1]Department of Neuroscience, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY
[2]Center for Computational Neuroscience, Flatiron Institute, New York, NY
{sl4742,lfabbott}@columbia.edu,schung@flatironinstitute.org

ABSTRACT

Understanding the asymptotic behavior of gradient-descent training of deep neural networks is essential for revealing inductive biases and improving network performance. We derive the infinite-time training limit of a mathematically tractable class of deep nonlinear neural networks, gated linear networks (GLNs), and generalize these results to gated networks described by general homogeneous polynomials. We study the implications of our results, focusing first on two-layer GLNs. We then apply our theoretical predictions to GLNs trained on MNIST and show how architectural constraints and the implicit bias of gradient descent affect performance. Finally, we show that our theory captures a substantial portion of the inductive bias of ReLU networks. By making the inductive bias explicit, our framework is poised to inform the development of more efficient, biologically plausible, and robust learning algorithms.

1 Introduction

Even after a task has been learned perfectly over a training set, learning typically continues to modify network parameters, often indefinitely. Understanding the form of the infinite-time, asymptotic solution has important implications for understanding the inductive bias of the network and improving its learning algorithm. For linear networks, $y = \langle \beta, x \rangle$, Soudry et al. (2018) have shown that gradient descent for most common losses used for classification asymptotically approaches a fixed margin classifier β with minimum L_2 norm (or equivalently a maximum margin classifier with fixed norm, that is, an SVM). Gunasekar et al. (2018b) extended this analysis to deep linear networks, showing that densely connected linear networks converge to the SVM solution regardless of their depth.

The results of Soudry et al. (2018) and Gunasekar et al. (2018b) have been generalized to homogeneous nonlinear predictors by Nacson et al. (2019) and Lyu and Li (2020). They demonstrate that gradient descent converges to a fixed margin classifier that minimizes the L_2 norm over all weights parameterizing the network. It remains unclear, however, what effect this penalty on the weights has on the inductive bias of the network they parameterize.

To answer this question, we extend the approach by Gunasekar et al. (2018b) to a nonlinear class of networks, in which β is a set of different predictors and the linear predictor for a specific input is selected by a context system. We assume that β is constructed from homogeneous polynomials of a global pool of weights w and that the network learns by gradient descent on w. Two interesting examples of such networks are gated linear networks (GLNs; Veness et al., 2017) and what we call frozen-gate ReLU networks (FReLUs). Because we focus on families of gated predictors, our approach covers certain models that escape the analysis by Nacson et al. (2019) and Lyu and Li (2020).

In this paper, we present several key findings.

- We derive and prove that gradient descent in GLNs asymptotically constructs a fixed margin solution with minimum norm on the collection of context-dependent linear predictors β, but with two important modifications: because all predictors are constructed from a shared set of weights, the norm of these vectors is minimized subject to certain equivariance constraints. For the same reason, gradient descent operates on a set of weights that are activated for multiple contexts and, as a consequence, the norm that β minimizes is different from the usual L_2 norm of an SVM.

- Unlike prior work, we use the asymptotic limit of gradient descent to directly train SVMs that share the GLNs' inductive bias. These models allow us to separately study how the equivariance constraints and gradient descent's implicit bias affect generalization.

- We then analyse the implicit bias of these GLNs in detail. Our analysis highlights that gradient descent (without any explicit regularization) incentivizes higher similarity between predictors that share part of their context and that this improves the GLN's performance.

- Finally, we investigate the asymptotic learning behavior of ReLU networks by applying a similar approach to FReLUs. We show that ReLU networks outperform GLNs in part because they can modify their context throughout learning, whereas GLNs cannot.

Taken together, our results have implications for understanding learning in deep neural networks and draw a connection between deep learning and SVMs that may inspire new and improved learning algorithms.

2 Gated Linear Networks

GLNs, like other deep neural networks, process their input through multiple layers of hidden units to compute their output. However, while conventional deep neural networks attain expressivity through nonlinearities, a hidden unit in a GLN computes its output without any nonlinearity. Instead, it attains expressivity by using different weights for different regions of the input space.[1]

Consider an example: a two-dimensional input (Figure 1d) with two units in the first hidden layer (Figure 1a). The first unit may use one set of weights to compute its output for $x_1 > 0$, and a different set of weights to compute its output for $x_1 \leq 0$ (as reflected by the two regions in Figure 1b and the two hyperplanes in Figure 1e). Similarly, the second unit may choose its set of weights differently for $x_2 > 0$ and $x_2 \leq 0$ (Figure 1c,f). As we combine multiple such hidden units, the GLN becomes increasingly expressive. For example, if a single unit now reads out the two hidden units, this readout will have learned a different linear predictor for all four quadrants (as reflected by the four hyperplanes in Figure 1g).

More generally, a GLN learns a different weight vector, β_γ, for each global context γ. Most inputs will not share this global context, but will have the same *local* context for particular hidden units. This means that the GLN uses, and therefore updates, overlapping sets of weights. For example, the bottom-right and top-right quadrant in Figure 1 share the local context of the first hidden unit and therefore use and update the same weights for this unit. (Moreover, all units in this example use the same set of weights for the second layer.) Compared to a shallow GLN using the same partitioning of the input space (which will activate a non-overlapping set of weights for each context), we will see that a deep GLN's linear predictors are equivariant with respect to the local contexts they share, imposing architectural constraints. These constraints change the network's inductive bias by reducing the search space. In addition, we will see that this equivariance changes the implicit bias of gradient descent.

This means that GLNs allow us to ask a fundamental question about gradient descent in deep networks: how do local changes in the weights affect the inductive bias of the global network parameterized by these weights? GLNs are conventionally trained using a local learning rule, where every hidden unit attempts to predict the output. However, we are interested in them specifically because their particular parameterization allows us to exactly characterize their asymptotic behavior under gradient descent.

3 Exact Asymptotic Behavior of Learning in Gated Linear Networks

3.1 Background: Learning in Linear Networks

To characterize the asymptotic behavior of learning in GLNs, we first turn to the asymptotic behavior of learning in linear networks, as characterized by Soudry et al. (2018) and Gunasekar et al. (2018b). Soudry et al. were concerned with gradient descent on a linearly separable dataset $(x^{(n)}, y^{(n)})$ using the *exponential loss* $\exp(-y\langle\beta, x\rangle)$, or a loss that has a similar tail[2] such as the cross-entropy $\ln(1 + \exp(-y\langle\beta, x\rangle))$.

---

[1]Veness et al. (2017) motivate GLNs through opinion pooling. We omit this motivation in favor of a simpler (but equivalent) description of the architecture. The motivation through opinion pooling also meant that they required their input to be scaled to $(0, 1)$, which usually amounted to a squashed version of the input (Veness et al., 2021). Since opinion pooling subsequently expands these probabilities using the inverse sigmoid, this does not tend to be very different from the way we set up our GLNs.

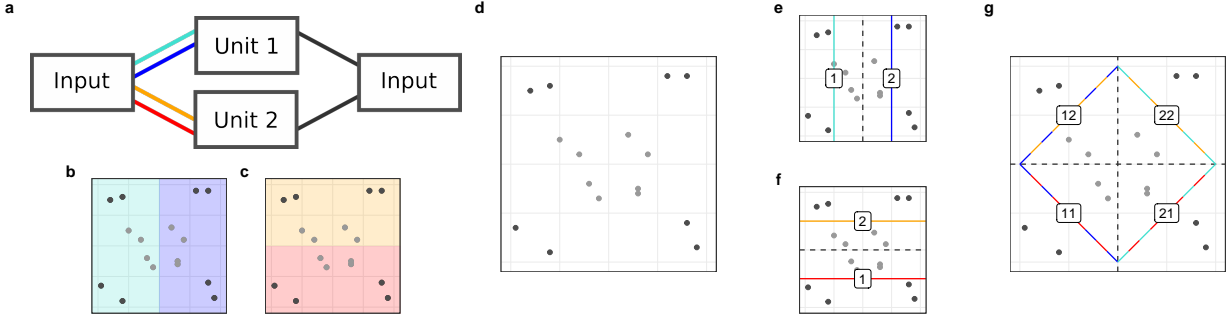[2]We call this class of loss functions *exponential-like*, see Definition B.10.

Figure 1: Example GLN. **a** Sketch of the GLN's structure. Every line represents a weight and units are linear. The colors indicate different weights, which are chosen depending on the region in which the input is located, as illustrated by panels **b** and **c**. For simplicity, we only use context-dependent weights for inputs to (not outputs from) the hidden layer. **d** The input data with labels indicated by shading. **e, f** The two hidden units each learn two weight vectors (as represented by the resulting hyperplanes). The weight is chosen depending on the context. **g** Output function. The output combines both hidden units and therefore has weights in four different regions. Each weight (as represented by the multicolored hyperplanes) is composed from two of the hidden units' linear weights.

If, through gradient descent, the overall loss approaches zero, the linear predictor's norm $\|\beta^{(t)}\|_2$ must diverge. Thus, all of the theorems we discuss (including ours) refer to the asymptotic direction that the weight vector converges to (provided it converges to a fixed direction). In particular, if we define $\hat{\beta}$ as the unit vector pointing in the same direction as the asymptotically diverging vector $\beta^{(\infty)}$, Soudry et al. (2018) show that $\hat{\beta}$ is a maximum margin predictor, or equivalently, that it is proportional to the solution of the optimization problem

$$\min \|\beta\|_2, \quad \text{s.t.} \ y^{(n)}\langle \beta, x^{(n)}\rangle \geq 1. \tag{1}$$

Using the Karush-Kuhn-Tucker (KKT) conditions (Karush, 1939; Kuhn and Tucker, 1951), this minimizing vector can be written as

$$\hat{\beta} = \sum_{n \in S} \lambda_n y^{(n)} x^{(n)}, \quad \lambda_n \geq 0. \tag{2}$$

Here, S is the set of data points for which the margin inequality is tight, i.e. $S = \{n | y^{(n)}\langle \beta, x^{(n)}\rangle = 1\}$. These are the *support vectors*.

Gunasekar et al. (2018b) extend this result to deep linear networks for which the output can be written as $\langle \mathcal{P}(w), x \rangle$, where $\mathcal{P}$ is a polynomial mapping the weights $w \in \mathbb{R}^P$ onto a linear predictor $\beta \in \mathbb{R}^D$. P and D are the number of weights and the input dimension, respectively. For example, a densely connected linear network with two layers has $f_w(x) = w_2^T w_1 x$, so $\mathcal{P}(w) = w_2^T w_1$. They require that $\mathcal{P}$ is homogeneous, that is, $\mathcal{P}(\alpha w) = \alpha^\nu \mathcal{P}(w)$, where ν is the *degree* of $\mathcal{P}$ (this excludes skip connections and bias units). For the example above, $\nu = 2$. They prove that if $w^{(t)}$ converges in direction to $\hat{w}$, $\hat{w}$ is proportional to a solution of the optimization problem

$$\min \|w\|_2, \quad \text{s.t.} \ y^{(n)}\langle \mathcal{P}(w), x^{(n)}\rangle \geq 1. \tag{3}$$

Whereas (1) directly minimizes the norm of the linear predictor, this problem penalizes the overall norm of internal weights that parameterize that predictor. The fixed margin constraint, however, still operates on the linear predictor $\mathcal{P}(w)$.

In contrast to the linear predictor, $\hat{w}$ is not necessarily a global minimum of (3). Instead, stationarity is akin to a local minimum in the context of minimizing a (potentially nonconvex) objective function, but additionally takes into account the margin constraints. More specifically, stationarity requires that

$$\hat{w} = \nabla_w \mathcal{P}(w) \sum_{n \in S} \lambda_n y_n x_n, \quad \lambda_n \geq 0, \tag{4}$$

so the weights are still constructed from a nonnegative sum of support vectors. However, these support vectors must be projected from the input space $\mathbb{R}^D$ to the weight space $\mathbb{R}^P$. This is achieved by the polynomial's Jacobian $\nabla_w \mathcal{P}(w) \in \mathbb{R}^{P \times D}$.

3

3.2 Extension to Generalized GLNs

In contrast to linear predictors and multi-layer linear networks, which are characterized by a single linear predictor, GLNs are characterized by a different weight vector β_γ for each global context γ. Each β_γ is given by a polynomial function of the weights w, and, if we leave out bias units beyond the first layer, this polynomial is homogeneous. Thus, it may appear that we can extend the analysis by Gunasekar et al. (2018b) simply by applying their theorem to each context-specific predictor β_γ. However, predictors for different contexts are parameterized by an overlapping set of weights, so this is not possible. This highlights the critical question we pose in our analysis: *how do the shared weights between different linear predictors affect the inductive bias of gradient descent?*

Motivated by these considerations, we extend the previous analysis by considering a set of (global) contexts $\gamma \in \Gamma$, where each context uses a different homogeneous polynomial $\mathcal{P}_\gamma$ to connect the global pool of weights to the context-specific linear predictor. This means that our model

$$f_w(x; \gamma) := \langle \mathcal{P}_\gamma(w), x \rangle, \tag{5}$$

depends on both the input x and the context γ. We call this class of functions *Generalized Gated Linear Networks* and, in particular, it covers GLNs (without bias units).

We are then able to prove the following (see Appendix B.1 for assumptions and proof):

Theorem 3.1. *Consider a dataset $(x^{(n)}, y^{(n)}, \gamma^{(n)})$, where $x^{(n)} \in \mathbb{R}^D$ is the input, $y^{(n)} = \{-1, 1\}$ is the label, and $\gamma^{(n)}$ is a global context. Then if $w^{(t)}$ converges to a fixed direction and the loss approaches zero, the limiting direction $\hat{w}$ is proportional to a stationary point of*

$$\min \|w\|_2^2, \quad s.t. \ y^{(n)} f_w(x^{(n)}; \gamma^{(n)}) \geq 1. \tag{6}$$

Here and throughout the article, $\|w\|_2^2$ denotes the sum of the squares of all the elements in all of the weight matrices of the network. This stationary point $\hat{w}$ is given by

$$\hat{w} = \sum_{\gamma \in \Gamma} \nabla_w \mathcal{P}_\gamma(w) \sum_{n \in S_\gamma} \lambda_n y_n x_n, \quad \lambda_n \geq 0, \tag{7}$$

where S_γ is the context-specific set of support vectors. Just as in (4), we sum over this set of support vectors and project it into the weight space $\mathbb{R}^P$. The stationary point $\hat{w}$ is then given by the sum of these contextwise projections.

3.3 Proof Sketch

We provide here an outline of the proof of Theorem 3.1.[3] Because this theorem is a relatively straightforward extension of the theorems by Soudry et al. (2018) and Gunasekar et al. (2018b), we begin with an outline of their proofs.

3.3.1 Background: Sketch of Previous Proofs

Soudry et al. (2018) consider the loss function

$$\mathcal{L}(\beta) = \sum_{n=1}^{N} \exp(-y_n \langle \beta, x_n \rangle), \tag{8}$$

and gradient descent updates $-\eta_t \nabla_\beta \mathcal{L}(\beta)$, where $\eta_t > 0$ is the learning rate. They then rely on two facts: that the gradient descent updates converge to some limit direction and that early updates are eventually forgotten. These imply that for large t, the weight direction approaches the limiting direction of the gradient descent updates.

Because

$$-\nabla_\beta \mathcal{L}(\beta) = \sum_{n=1}^{N} \exp(-y_n \langle \beta, x_n \rangle) y_n x_n, \tag{9}$$

the weight $\exp(-\|\beta^{(t)}\|_2 y_n \langle \hat{\beta}, x_n \rangle)$ will converge to zero as $\|\beta^{(t)}\|_2$ increases. However, the weights of the data points with the smallest margins, i.e. the support vectors, will converge to zero exponentially slower than all other data points. Thus, the support vectors dominate the gradient's direction and we can write

$$\hat{\beta} = \sum_{n \in S} \lambda_n y_n x_n, \tag{10}$$

---

[3] A rigorous proof can be found in Appendix B.1.

which is (2). The individual values for λ_n are determined by the rate with which the loss of individual data points, $\exp(-y_n\langle\beta, x_n\rangle)$, approaches zero.

Gunasekar et al. (2018b) extend this result by decomposing the loss gradient for polynomial predictors into $\nabla_w\mathcal{L}(w) = \nabla_w\mathcal{P}(w)(-\nabla_{\mathcal{P}(w)}\mathcal{L}(w))$. The latter part, $-\nabla_{\mathcal{P}(w)}\mathcal{L}(w)$, corresponds to the gradient in the linear predictor. Even though gradient descent is not performed on this linear predictor directly, the result of Soudry et al. (2018) generalizes: if $-\nabla_{\mathcal{P}(w)}\mathcal{L}(w)$ converges to some limit direction, we can[4] again infer that the support vectors will, at large times, dominate this direction. This implies that we can approximate the loss gradient (and thus the weight directional limit) as a sum of support vectors that is projected into the weight space by the Jacobian, as is given by (4).

3.3.2 Proof Sketch of Theorem 3.1

To extend these results to GLNs, we decompose the loss function into a sum over the losses specific to each context,

$$\mathcal{L}_\gamma(w) = \sum_{n:\gamma^{(n)}=\gamma} \exp(-y_n\langle\mathcal{P}_\gamma(w), x_n\rangle). \tag{11}$$

The gradient can be decomposed similarly into

$$\nabla_w\mathcal{L}(w) = \sum_\gamma \nabla_w\mathcal{L}_\gamma(w). \tag{12}$$

Extending the strategy of Gunasekar et al. (2018b) to a sum of loss functions, we asymptotically express $w^{(t)}$ as a weighted sum of the individual limit directions,

$$w^{(t)} \approx \|w^{(t)}\|_2 \sum_\gamma \xi_\gamma \nabla_w\mathcal{P}_\gamma(w)(-\nabla_{\mathcal{P}_\gamma(w)}\mathcal{L}_\gamma(w)). \tag{13}$$

(The linear weights ξ_γ are necessary because the different components of the loss might be scaled differently.) The contextwise gradients are again dominated by the support vectors and we can absorb ξ_γ into λ_n to arrive at (7).

4 Gated Linear Networks with Two Layers

Theorem 3.1 allows us to characterize the implicit bias of gradient descent. However, the minimized norm in this theorem is that of the weights parameterizing the context-dependent linear predictors, not the predictors themselves. To connect our results directly to the linear predictors, we consider a special case: GLNs of depth 2, with one output neuron and one context for this output neuron. This means that we have two hyperparameters for our architecture: the number of hidden units H in the first layer, and the number of contexts per hidden unit C. As a consequence, the global context is given by $\gamma \in \{1, \ldots, C\}^H$. The GLN's weights are given by

$$w = (w^{(1)}, w^{(2)}), \quad w^{(1)} \in \mathbb{R}^{H \times C \times D}, w^{(2)} \in \mathbb{R}^H, \tag{14}$$

and the resulting linear predictors are

$$\beta_\gamma = \sum_{h=1}^H w_h^{(2)} w_{h\gamma_h}^{(1)} \in \mathbb{R}^D. \tag{15}$$

Throughout this exposition, we consider as a simple example $H = C = 2$, as in Figure 1.

As we noted in Section 2 and are now able to analyse in more detail, the parameterization of these networks affects their inductive bias in two ways. First, it imposes *architectural constraints* on the resulting linear predictors β_γ. This manifests in an equivariance condition on neighboring predictors: the difference between two linear predictors is invariant to the contexts they share. In the case of $C = H = 2$, this condition is given by

$$\beta_{21} - \beta_{22} = \beta_{11} - \beta_{12}. \tag{16}$$

As we can see, the two predictors on the left share the first unit's local context and so changing this context does not affect their difference. More generally (see Appendix B.2.1), this means that even though β specifies a set of C^H linear predictors, we can only choose $(C-1)H + 1$ of them freely.

Second, whereas shallow networks minimize the L_2 norm, gradient descent on deep GLNs implicitly minimizes a different norm, which we call the *GLN norm* and denote by $\|\cdot\|_{\mathrm{GLN}}$. Importantly, this norm operates on the linear predictors β_γ instead of the underlying global pool of weights. We characterize $\|\cdot\|_{\mathrm{GLN}}$ in Section 4.2. First, however, we would like to illustrate why it is important to understand the difference between $\|\cdot\|_2$ and $\|\cdot\|_{\mathrm{GLN}}$. To this end, the next section illustrates that the $\|\cdot\|_{\mathrm{GLN}}$ is not only more consistent with a GLN trained with gradient descent, but also leads to better generalization on MNIST.

---

[4]Gunasekar et al. (2018a) prove this for the exponential loss and note that they expect the result to generalize to exponential-like losses. We prove this generalization in Appendix B.3.
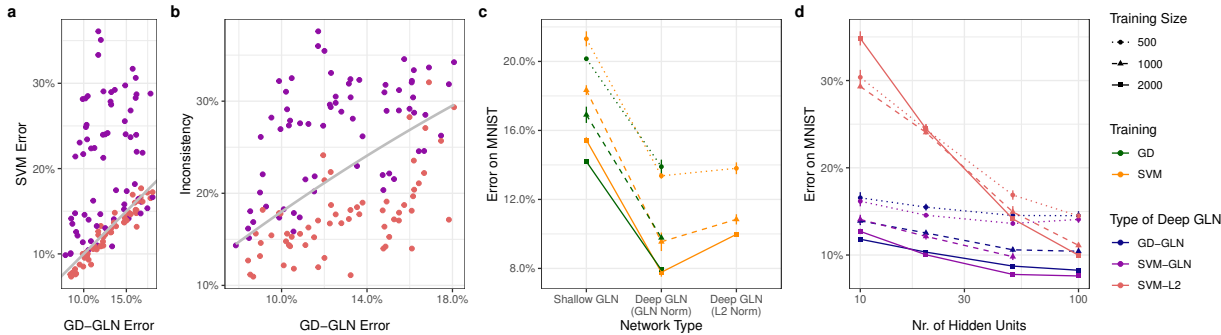
Figure 2: Experiments on GLNs. **a** The error of the GD-GLN plotted against the error of the SVMs. The grey line represents identical performance. **b** Inconsistency between the GD-GLN and the SVM plotted against the GD-GLN's error on MNIST. The grey line represents the inconsistency we would expect from a predictor with matched error rate, but no further correlation with the network (see Footnote 6). **c** Performance of the best shallow and deep GLNs across the number of contexts and (in the case of deep GLNs) hidden units. **d** Error of the SVMs and the GD-GLN plotted against the number of hidden units. This plot depicts the networks with two contexts per hidden unit, see Figure 4 for the networks with four contexts per hidden unit.

4.1 The GLN-Norm Improves Generalization Over the L_2 Norm

Using our theorem, we can examine the impact of the architectural constraints on the deep GLN alone or together with the resulting changes in implicit bias. We fit two support vector machines that respect the deep GLN's architectural constraints, minimizing either the L_2 norm (SVM-L2) or the GLN norm (SVM-GLN) while maintaining a fixed margin. Note that it is the second of these that reflects the full result of our theorem. Both optimization problems are convex, so we can use convex optimization algorithms (Diamond and Boyd, 2016) and are guaranteed to find a global minimum. In addition, we trained a GLN using gradient descent (GD-GLN) for 3200 steps in PyTorch (Paszke et al., 2019).

We trained all our models on MNIST (LeCun et al., 2010), which we turn into a binary classification problem by grouping together the digits 0-4 and 5-9. Since we use full-batch convex optimization, we are restricted in the size of our training data, using subsets of 500, 1000, and 2000 data points. To evaluate generalization, we use a validation dataset with 12000 data points.

We consider GLNs with 10, 20, 50, and 100 hidden units and two or four contexts per hidden unit. We assigned contexts by partitioning the input space using randomly sampled hyperplanes (as is illustrated in Figure 1), similar to Veness et al. (2017). For every architecture, we used three random seeds to sample these hyperplanes. Figure 2 depicts the mean and standard deviation across these three runs.[5]

Figure 2a compares the error of the GD-GLN with the SVMs that use the same contexts, number of hidden units, and number of contexts per hidden units. SVM-GLN matches the GD-GLN in accuracy and even outperforms it slightly. In contrast, SVM-L2 performs much worse and its performance is only weakly correlated with that of the GD-GLN with matching hyperparameters.

Figure 2b depicts the proportion of the validation data for which the SVM predicts different labels than the GD-GLN (inconsistency). If the GD-GLN had truly converged to the SVM-GLN, the inconsistency would be zero. Instead the two predictors make inconsistent predictions on a substantial proportion (more than 10%) of the data. The inconsistency tends to be lower than we would expect from two models with a matching error rate, but no further correlation (grey line).[6] It is also much lower than that between the SVM-L2 and the GD-GLN. Still, this result highlights a substantial difference between the infinite-time predictor we derived and its finite-time counterpart (see Discussion).

Next, we looked at how different choices of depth and width interact with the implicit bias of gradient descent. Figure 2c depicts the best-performing shallow and deep GLN (i.e. one and two layers) across all hyperparameters. This illustrates that the architectural constraints paired with the L_2 norm already allow the deep GLN to find better solutions than

---

[5]More details on the experimental setup can be found in Appendix A. Code to reproduce all experiments can be found at https://github.com/sflippl/implicit-bias-glns.

[6]Suppose this error rate is p. Inconsistent labels mean that one model makes an error and the other does not. The probability of this happening is $2p(1-p)$.

the shallow GLN (which simply learns an SVM for each context). However, we again see that the SVM-GLN further improves in performance over the SVM-L2. Isolating the effects that making the network deeper has on the functions it can express, and on the solutions that gradient descent discovers in practice, would not have been possible without our theory.

Finally, Figure 2d illustrates that the SVM-L2 performs much worse than the SVM-GLN for few hidden units in particular. This panel also makes particularly apparent that the SVM-GLN tends to slightly outperform the GD-GLN with the same hyperparameters. This is exactly what we would expect if the GD-GLN slowly converges to the SVM-GLN and if its generalization performance improves throughout infinite training.

4.2   Understanding the GLN Norm

Having seen that it provides a useful inductive bias, we now turn to understanding the GLN norm. (15) makes apparent that w only affects β through the auxiliary variable

$$\zeta_{h\gamma_h} := w_h^{(2)} w_{h\gamma_h}^{(1)}. \tag{17}$$

For a fixed ζ, which choices of w minimize $\|w\|_2^2$? Intuitively, the L_2 norm incentivizes us to distribute magnitudes across parameters equally. Since $\zeta_h = w_h^{(2)} w_h^{(1)}$, the two parameters should share the magnitude $\|\zeta_h\|_2$ equally, i.e.

$$|w_h^{(2)}| = \|w_h^{(1)}\|_2 = \sqrt{\|\zeta_h\|_2}. \tag{18}$$

This implies[7] that

$$\|w\|_2^2 = \sum_{h=1}^{H} \|w_h^{(1)}\|_2^2 + |w_h^{(2)}|^2 \propto \sum_{h=1}^{H} \|\zeta_h\|_2, \tag{19}$$

and thus

$$\|\beta\|_{\text{GLN}} = \min_{\zeta} \sum_{h=1}^{H} \|\zeta_h\|_2, \quad \text{s.t.} \ \beta_\gamma = \sum_{h=1}^{H} \zeta_{h\gamma_h}. \tag{20}$$

This means that, when expressed in terms of ζ, the GLN norm takes on the form of a group lasso (Yuan and Lin, 2006). Importantly, this norm is different from the L_2 norm on ζ, which would involve summing up $\|\zeta_h\|_2^2$ instead of $\|\zeta_h\|_2$. Because it computes the L_2 norm of ζ_h, the GLN norm encourages the magnitude of this vector to be as small as possible. But because it sums up the norm itself instead of its square, it also incentivizes setting entire components ζ_h to zero, similar to how the L1 norm incentivizes setting single entries of a vector to zero. Put differently, the GLN norm incentivizes sparsity in the components ζ_h. Because each component ζ_h encodes differences in the predictor as a consequence of the different local contexts of the hidden unit h, this norm therefore encourages the set of linear predictors to only learn differences between unit-specific contexts if this is actually useful.

We can further illustrate the difference between $\| \cdot \|_{\text{GLN}}$ and the L_2 norm by considering the special case $H = C = 2$. In this case,

$$\|\beta\|_{\text{GLN}}^2 = \|\beta\|_2^2 + \tfrac{1}{2} \sum_{i,j} \|\beta_{ij} - \beta_{\bar{i}j}\|_2 \|\beta_{ij} - \beta_{i\bar{j}}\|_2, \tag{21}$$

where $\overline{k}$ denotes the local context opposite to k, i.e. $\overline{k} = 2$ if $k = 1$ and $\overline{k} = 1$ if $k = 2$. The GLN norm therefore adds to the L_2 norm a component that encourages neighboring predictors to be more similar. Without any explicit regularization, the equivariant interactions of the GLN cause predictors that share parts of their global context (and thus overlap in their weights) to become more similar to each other.

5   Frozen-Gate ReLU Networks

While we were motivated by understanding GLNs, Theorem 3.1 applies to other architectures as well. We apply the theorem to a particular variation on ReLU networks that makes them generalized gated linear networks: frozen-gate ReLU networks.

A single hidden unit in a ReLU network computes its activation as $z = \max(\langle w, x \rangle, 0)$. That is, it first computes a linear function and then sets any negative values to zero. For fixed weights, we can also implement this with a gated linear predictor. More specifically, there are two contexts associated with the hidden unit, depending on the sign of $\langle w, x \rangle$. If $\langle w, x \rangle > 0$, we use the weight w to compute the hidden unit. If $\langle w, x \rangle \leq 0$, we instead use a zero vector. The

---

[7]Technically, we have to check equivalence of the KKT conditions. We do so in Appendix B.2 and the same intuition applies.
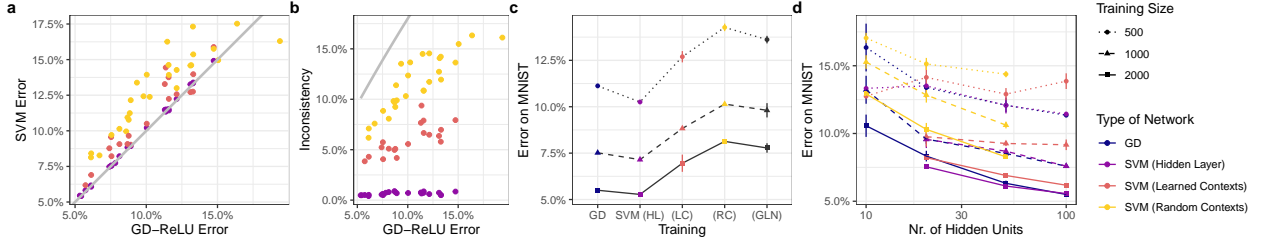
Figure 3: Experiments on ReLU networks. **a** The error of the GD-ReLU plotted against the SVMs. The grey line represents identical performance. **b** Inconsistency between the GD-ReLU and the SVMs plotted against the GD-ReLU's error on MNIST. The grey line represents the inconsistency we would expect from a predictor with matched error rate, but no further correlation with the network (see Footnote 6). **c** Performance of the best ReLU networks as well as the best deep GLN. To make sure the deep GLN does not have more parameters than the FReLU, we only use those GLNs with two contexts and at most 50 hidden units. **d** Error of the SVMs and the GD-ReLU plotted against the number of hidden units.

strategy of separating the gates in this way is similar to Lakshminarayanan and Vikram Singh (2020), who use it to define a neural tangent kernel.

For fixed weights, this gated linear predictor is exactly equivalent to the ReLU network. However, we train it by only changing the linear weights, freezing the gates that determine the context for each hidden unit. We thus call this architecture *frozen-gate ReLU networks* (FReLUs). Throughout gradient descent, as the weights change but the contexts remain fixed, the FReLU diverges from the standard ReLU network. To mitigate this divergence, we also consider networks in which the weights are learned in the usual way for a period of time and then the gates are frozen to apply the asymptotic analysis.

Using FReLUs as an approximation, can our theory shed light on the inductive bias of ReLU networks and how it is different from that of GLNs? To investigate this, we first characterize the implicit bias of FReLUs. We then compare FReLUs with ReLU networks trained with gradient descent.

5.1 The Implicit Bias of FReLUs

FReLUs are structured almost like GLNs, except that one of the two context-gated weight vectors is set to zero. It is therefore not surprising that they also minimize a group Lasso norm. Specifically, for a given context $\gamma \in \{0, 1\}^H$, we can parameterize the resulting linear predictor as $\beta_\gamma = \sum_{h:\gamma_h=1} \zeta_h$, where $\zeta \in \mathbb{R}^{H \times D}$ is an auxiliary variable defined similarly as in Section 4.2. FReLUs then minimize the norm

$$\|\zeta\|_{\mathrm{FReLU}} := \sum_{h=1}^{H} \|\zeta_h\|_2. \tag{22}$$

Since the group Lasso encourages sparsity, this means that unless it would otherwise increase the network's margin, ζ_h will be low or set to zero. Since the ζ_h's induce the kinks in the network's separating hypersurface, this means that gradient descent (again without any explicit regularization) encourages this surface to be as straight as possible.

5.2 Comparing ReLU networks and FReLUs

Can we use this insight to better understand generalization in ReLU networks trained with gradient descent? To investigate this, we trained a ReLU network on the binary MNIST task using gradient descent (GD-ReLU). We trained networks with 10, 20, 50, and 100 hidden units and used three random seeds for initialization. We then compared each network to an SVM trained on $\|\cdot\|_{\mathrm{FReLU}}$ using random contexts and a matching architecture (SVM-RC). This network's performance is already reasonably correlated with that of the matching GD-ReLU (Figure 3a). However, a substantial proportion of its predictions do not match that of the GD-ReLU (Figure 3b), although they are still more consistent than we would expect from an uncorrelated network with matching accuracy.

One reason that the SVM-RC may have a different inductive bias from the GD-ReLU is that the latter adapts its contexts. To see if this can explain part of the disparity, we trained an SVM on $\|\cdot\|_{\mathrm{FReLU}}$ using the context that GD-ReLU has learned at the end of its training (SVM-LC). Indeed, this predictor generalizes better (Figure 3a) and is considerably more consistent with the GD-ReLU (Figure 3b). However, an SVM trained on the hidden layer's activation (SVM-HL) still performs better (it even outperformed the GD-ReLU, Figure 3c) and is more consistent with the GD-ReLU. This is not surprising as the SVM-HL has many fewer free parameters. Nevertheless, this highlights that there is a nontrivial disparity between SVM-LC and GD-ReLU.

The fact that a FReLU with learned contexts outperforms one with random contexts indicates that learnable contexts can be beneficial for a network. In particular, Figure 3c demonstrates that the SVM-LC outperforms a deep GLN of similar size whereas a FReLU using random contexts does not (and is, in fact, slightly worse). Figure 3d shows that the SVM-LC uniformly outperforms the SVM-RC for any number of hidden units.[8] However, its disparity to the GD-ReLU and the SVM-HL increases with increasing latent dimensions.

Our analysis demonstrates that the SVM-LC captures a substantial portion of the ReLU network's inductive bias. FReLUs therefore promise a new perspective on why ReLU networks generalize well: they can adapt their gating function and use the resulting contexts as sparsely as possible. Still, the SVM-LC is also substantially different from the GD-ReLU. This may be because of finite-time effects or because the fact that ReLU networks learn their weights and gates in an entangled manner changes their inductive bias. We leave investigating this question (for example using the results by Lyu and Li (2020)) to future work.

6 Discussion

In this article, we characterized the asymptotic behavior of gradient-descent training of Generalized Gated Linear Networks. We used this theory to exactly characterize the norm minimized by a deep GLN and confirmed that this allows us to train an SVM that captures its performance. This allowed us to tease apart the contributions of architectural constraints and the implicit bias of gradient descent, demonstrating that the implicit bias is essential for good generalization. We also confirmed that this allows us to capture a substantial portion of the inductive bias of ReLU networks, attributing part of their generalization performance to the fact that they (a) learn their contexts and (b) use them as sparsely as possible. This suggests that we might be able to take inspiration from ReLU networks to devise context learning algorithms for GLNs. Conversely, a perspective that decomposes gradient-descent training in ReLU networks into context and weight learning, may shed new light on their inductive bias.

Our experiments indicate our theory's potential to help us understand the inductive bias of deep neural networks. To realize this potential, we must address the fact that the infinite-time deep GLN still makes substantially different predictions from its finite-time counterpart. This may be due to finite-time effects (Arora et al., 2019). Alternatively, gradient descent and convex optimization may have converged to different subsets of the stationary points characterized by our theory (Nacson et al., 2019).

Still, the fact that we can train an SVM that matches (and even outperforms) its corresponding deep GLN indicates that our theory allows us to successfully disentangle the particular optimization procedure used from the inductive bias it implements. This means that we can consider alternative learning algorithms that find the same stationary points, but have other benefits, for example faster convergence, more efficient computations, or higher biological plausibility. To this end, comparing the inductive bias of gradient descent to that of the local learning rule conventionally applied to GLNs (for instance using the results by Ji and Telgarsky (2019)) may help us design new local learning rules that generalize better. Finally, our framework connects networks trained with gradient descent to SVMs, which have formal adversarial protections (Mangasarian, 1999; Gentile, 2003). This perspective may therefore allow us to learn more robust networks, either by imposing the results of infinite-time training or by changing the inductive bias.

Acknowledgements

---

[8]Missing data points indicate that constraints could not be satisfied, or that the optimizer did not converge in the allocated number of iterations (see Appendix A).

References

Agrawal, A., Verschueren, R., Diamond, S., and Boyd, S. (2018). A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60.

Arora, S., Cohen, N., Hu, W., and Luo, Y. (2019). Implicit Regularization in Deep Matrix Factorization. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F. d., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5.

Domahidi, A., Chu, E., and Boyd, S. (2013). ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, pages 3071–3076.

Gentile, C. (2003). The Robustness of the p-Norm Algorithms. *Machine Learning*, 53(3):265–299.

Grant, M., Boyd, S., and Ye, Y. (2006). Disciplined Convex Programming. In Liberti, L. and Maculan, N., editors, *Global Optimization: From Theory to Implementation*, pages 155–210. Springer US, Boston, MA.

Gunasekar, S., Lee, J., Soudry, D., and Srebro, N. (2018a). Characterizing Implicit Bias in Terms of Optimization Geometry. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1832–1841. PMLR.

Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. (2018b). Implicit Bias of Gradient Descent on Linear Convolutional Networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Ji, Z. and Telgarsky, M. (2019). The implicit bias of gradient descent on nonseparable data. In Beygelzimer, A. and Hsu, D., editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1772–1798. PMLR.

Karush, W. (1939). *Minima of functions of several variables with inequalities as side conditions*. PhD Thesis, Thesis (S.M.)–University of Chicago, Department of Mathematics, December 1939.

Kuhn, H. and Tucker, A. (1951). Nonlinear Programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492. University of California Press.

Lakshminarayanan, C. and Vikram Singh, A. (2020). Neural Path Features and Neural Path Kernel : Understanding the role of gates in deep learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5227–5237. Curran Associates, Inc.

LeCun, Y., Cortes, C., and Burges, C. (2010). MNIST handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2.

Lyu, K. and Li, J. (2020). Gradient Descent Maximizes the Margin of Homogeneous Neural Networks. In *International Conference on Learning Representations*.

Mangasarian, O. L. (1999). Arbitrary-norm separating plane. *Operations Research Letters*, 24(1):15–23.

Nacson, M. S., Gunasekar, S., Lee, J., Srebro, N., and Soudry, D. (2019). Lexicographic and Depth-Sensitive Margins in Homogeneous and Non-Homogeneous Deep Models. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4683–4692. PMLR.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F. d., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.

Pedersen, T. L. (2020). *patchwork: The Composer of Plots*.

R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Saxe, A., McClelland, J., and Ganguli, S. (2014). Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. International Conference on Learning Represenatations 2014.

Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. (2018). The Implicit Bias of Gradient Descent on Separable Data. *Journal of Machine Learning Research*, 19(70):1–57.

Stellato, B., Banjac, G., Goulart, P., Bemporad, A., and Boyd, S. (2020). OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 12(4):637–672.

Veness, J., Lattimore, T., Bhoopchand, A., Grabska-Barwinska, A., Mattern, C., and Toth, P. (2017). Online Learning with Gated Linear Networks. *arXiv*. arXiv: 1712.01897.

Veness, J., Lattimore, T., Budden, D., Bhoopchand, A., Mattern, C., Grabska-Barwinska, A., Sezener, E., Wang, J., Toth, P., Schmitt, S., and Hutter, M. (2021). Gated Linear Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(11):10015–10023.

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67. _eprint: https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9868.2005.00532.x.

A   Experimental Setup

A.1   Training the Finite-Time Predictors

We train the models that use gradient descent with PyTorch (Paszke et al., 2019) and PyTorch Lightning (https://github.com/PyTorchLightning/pytorch-lightning) in Python 3.9. We train all models for 1600 steps with a learning rate of 0.04 and for another 1600 steps with a learning rate of 0.01.

For the deep and shallow GLNs, we use orthogonal initialization (Saxe et al., 2014) for the weights and determine the contexts using hyperplanes, following Veness et al. (2017). In particular, we randomly sample hyperplanes $w \in \mathbb{R}^D$ in the input space using a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 36$. We randomly sample a cutoff $b \in \mathbb{R}$ with mean $\mu = 0$ and standard deviation $\sigma = 9$. The corresponding context function is then given by

$$C(x) = \begin{cases} 1 & \text{if } w^T x - b \geq 0, \\ 0 & \text{if } w^T x - b < 0. \end{cases}$$

To generate a context function with more than two possible contexts, we compose multiple hyperplanes, mapping each unique region produced by these multiple hyperplanes to its own context. Unlike Veness et al. (2017), we additionally use a cutoff b that makes C balanced, i.e. maps half the training data to 1 and half the training data to 0. We use this median cutoff in the results presented in the main article. Appendix C.1 discusses results on a random cutoff.

For the ReLU networks, we use Kaiming normal initialization (He et al., 2015).

A.2   Convex Optimization

To solve the convex optimization problems, we used the `cvxpy` library in Python (Diamond and Boyd, 2016; Agrawal et al., 2018), which follows the paradigm of Disciplined Convex Programming (DCP) (Grant et al., 2006). DCP follows a set of conventions on how to formalize convex optimization problems. We trained the shallow GLN as well as the SVM-L2 and SVM-HL using the OSQP algorithm (Stellato et al., 2020) with at most 10,000 iterations. We trained the SVM-GLN, SVM-RC, and SVM-LC using the ECOS algorithm (Domahidi et al., 2013) with at most 200 iterations. We sampled the random contexts for the FReLUs using the same method as for the GLNs with median initialization.

Whereas the other architectures were easily translated into the DCP conventions, the SVM-L2 predictor required a bit more attention. More specifically, it is more natural to express the architectural constraints by specifying ζ instead of β and so we wanted to compute the equivalent of $\|\beta\|_2^2$ for ζ. It turns that this is given by

$$\|A\zeta\|_2^2, \quad (A\zeta)_{hc} := \frac{1}{HC} \sum_{h',c'} \zeta_{h'c'} - \sum_{c'} \zeta_{hc'} + C\zeta_{hc}, \tag{23}$$

which we prove below.

Proof. We have to prove that $\|\beta\|_2^2 = \|A\zeta\|_2^2$. Since β satisfies our architectural constraints, we know that there is a ζ such that

$$\beta_\gamma = \sum_{h=1}^H \zeta_{h\gamma_h}$$

for all γ. Thus,

$$\|\beta\|_2^2 = \sum_{\gamma \in \{1,...,C\}^H} \left\| \sum_{h=1}^H \zeta_{h\gamma_h} \right\|_2^2 = \sum_{\gamma \in \{1,...,C\}^H} \sum_{h,h'=1}^H \langle \zeta_{h\gamma_h}, \zeta_{h'\gamma_{h'}} \rangle = \Delta_1 + \Delta_2,$$

where we define Δ_1, Δ_2 by changing the order of the summation operators and splitting the summation over h and h' into the case $h = h'$ (for Δ_1) and $h \neq h'$ (for Δ_2), i.e.

$$\Delta_1 = \sum_{h=1}^H \sum_{\gamma \in \{1,...,C\}^H} \|\zeta_{h\gamma_h}\|_2^2, \quad \Delta_2 = \sum_{h \neq h'} \sum_{\gamma \in \{1,...,C\}^H} \langle \zeta_{h\gamma_h}, \zeta_{h\gamma_{h'}'} \rangle.$$

We can now simplify these equations by noting that $\zeta_{h\gamma_h}$ is invariant to all but one dimension of our iterator γ:

$$\Delta_1 = C^{H-1} \sum_{h=1}^H \sum_{c=1}^C \|\zeta_{hc}\|_2^2, \quad \Delta_2 = C^{H-2} \sum_{h \neq h'} \sum_{c,c'=1}^C \langle \zeta_{hc}, \zeta_{h'c'} \rangle.$$

Defining

$$M \in \mathbb{R}^{(H \times C \times D)^2}, \quad M_{\substack{hcd \\ h'c'd'}} := \delta_{dd'} \left(C \delta_{hh'} \delta_{cc'} + (1 - \delta_{hh'}) \right) = \delta_{dd'} \left(1 + C(\delta_{hh'} \delta_{cc'}) - \delta_{hh'} \right), \tag{24}$$

we can rewrite this norm as

$$\|\beta\|_2^2 = C^{H-2} \sum_{\substack{hcd \\ h'c'd'}} M_{\substack{hcd \\ h'c'd'}} \zeta_{hcd} \zeta_{h'c'd'}.$$

All that remains to show is that $A^2 = M$ (note that A is symmetric).

To show this, we make a parameterized guess, following the ansatz

$$A \in \mathbb{R}^{(H \times C \times D)^2}, \quad A_{\substack{hcd \\ h'c'd'}} = \delta_{dd'} \left(\alpha + \beta \delta_{hh'} + \gamma \delta_{cc'} + \kappa \delta_{hh'} \delta_{cc'} \right)$$

We thus require

$$M_{\substack{hcd \\ h'c'd'}} = \sum_{h'',c'',d''} A_{\substack{hcd \\ h''c''d''}} A_{\substack{h''c''d'' \\ h'c'd'}} = \delta_{dd'} \left(HC\alpha + C\beta \delta_{hh'} + H\gamma \delta_{cc'} + \kappa \delta_{hh'} \delta_{cc'} \right),$$

and therefore

$$\alpha = \tfrac{1}{HC}, \beta = -1, \gamma = 0, \kappa = C. \tag{25}$$

$\square$

A.3 Data Analysis

We performed all data analysis in R. (R Core Team, 2021) All figures (except for Figure 1a, for which we used Inkscape) were created using ggplot2 (Wickham, 2016) and patchwork (Pedersen, 2020). In the supplementary material, we provide an R package that contains all data as well as the code reproducing all figures.

B   Proofs

B.1   Proof of Theorem 3.1

We begin by restating the theorem:

Theorem 3.1. *Consider a dataset $(x^{(n)}, y^{(n)}, \gamma^{(n)})$, where $x^{(n)} \in \mathbb{R}^D$ is the input, $y^{(n)} = \{-1, 1\}$ is the label, and $\gamma^{(n)}$ is a global context. Then if $w^{(t)}$ converges to a fixed direction and the loss approaches zero, the limiting direction $\hat{w}$ is proportional to a stationary point of*

$$\min \|w\|_2^2, \quad s.t. \ y^{(n)} f_w(x^{(n)}; \gamma^{(n)}) \geq 1. \tag{6}$$

The theorem is operating under the following fundamental assumptions:

Assumption B.1. ℓ is an exponential-like loss.

Assumption B.2. $\mathcal{L}(w^{(t)}) \to 0$.

Assumption B.3. $w^{(t)}$ converges in direction to some $\hat{w}$.

In addition, we consider two assumptions excluding pathological cases, Assumptions B.4 and B.5, which we introduce at the point where they become necessary.

To prove the theorem, we must demonstrate (1) primal feasibility and (2) stationarity.

Primal feasibility simply involves proving that there is some $w^* = \alpha \hat{w}$ such that

$$y^{(n)} f_{w^*}(x^{(n)}; \gamma^{(n)}) \geq 1.$$

This is fairly straightforward, with a minor complication being created by the fact that different polynomials $\mathcal{P}_\gamma$ may have different degrees ν_γ.

For any given γ, we define the minimal margin of this context's data,

$$m_\gamma := \min_{n:\gamma^{(n)}=\gamma} y^{(n)} f_{\hat{w}}(x^{(n)}; \gamma). \tag{26}$$

13

By assumption, we know that $L(w^{(t)}) \to 0$, and so we are guaranteed that for all n and large enough t, $y^{(n)} f_{w^{(t)}}(x^{(n)}) > 0$. However, we wish to exclude the pathological case, in which the margin of the normalized weight $\hat{w}^{(t)}$ still converges to 0 instead of a positive value. This motivates the following assumption:

Assumption B.4. For all n, $y^{(n)} f_{\hat{w}}(x^{(n)}; \gamma^{(n)}) > 0$.

From this assumption, we know that for all γ, $m_\gamma > 0$. Scaling w by some α results in the changed margin

$$y^{(n)} f_{\alpha \hat{w}}(x^{(n)}; \gamma^{(n)}) = \alpha^{\nu_{\gamma^{(n)}}} y^{(n)} f_{\hat{w}}(x^{(n)}; \gamma^{(n)}).$$

So for each γ, we can set the minimal margin to 1 by scaling w by m_γ^{-1/ν_γ}. Since we want to make sure that the margin for each context is not smaller than 1, we scale w by

$$\alpha := \max_\gamma m_\gamma^{-1/\nu_\gamma}, \tag{27}$$

setting $w^* := \alpha \hat{w}$. Thus, we know that w^* satisfies the margin constraints (i.e. primal feasibility) and are left to check whether it also satisfies stationarity.

To demonstrate **stationarity**, we follow the same strategy as Gunasekar et al. (2018b). To a large extent, we can make arguments that are exactly analogous to theirs and we refer to their proof in these cases. For each γ, we consider the resulting sequence of linear predictors $\beta_\gamma^{(t)} = \mathcal{P}_\gamma(w^{(t)})$. We would like to apply Lemma B.11 (which allows us to consider any exponential-like loss, in contrast to Gunasekar et al. (2018b)). For this purpose, we consider the contextwise loss function

$$\mathcal{L}_\gamma(\beta_\gamma) = \sum_{n : \gamma^{(n)} = \gamma} \ell(y^{(n)} \langle \beta_\gamma, x^{(n)} \rangle). \tag{28}$$

Since the sum of all these loss functions converges to zero and all $\mathcal{L}_\gamma$ or nonnegative, we immediately know that $\mathcal{L}(\beta^{(t)}) \to 0$. Analogous to Gunasekar et al. (2018b), we also know that $\beta_\gamma^{(t)} / \|\beta_\gamma^{(t)}\|_2 \to \hat{\beta}$, where

$$\hat{\beta} := \frac{\mathcal{P}(\hat{w})}{\|\mathcal{P}(\hat{w})\|}. \tag{29}$$

All that is left is to exclude pathological cases where the gradient of the loss in the linear predictor does not converge:

Assumption B.5. For all γ, $\nabla_\beta \mathcal{L}_\gamma(\beta_\gamma^{(t)})$ converges in direction to some $\hat{z}_\gamma$.

This allows us to apply Lemma B.11 and infer that for each γ,

$$\hat{z}_\gamma = \sum_{n \in S_\gamma} \lambda_n y_n x_n, \quad \lambda_n \geq 0. \tag{30}$$

The gradient update $\Delta w^{(t)} = \eta_t \nabla_w \mathcal{L}(w)$ can be decomposed into contextwise updates

$$\Delta w_\gamma^{(t)} = \eta_t \nabla_w \mathcal{L}_\gamma(w). \tag{31}$$

Similarly, we write

$$w_\gamma^{(t)} := w^{(0)}/|\Gamma| + \sum_\gamma \Delta w_\gamma^{(t)}, \tag{32}$$

which implies that

$$w^{(t)} = \sum_\gamma w_\gamma^{(t)} \tag{33}$$

(We could distribute the initial value $w^{(0)} = \sum_{\gamma \in \Gamma} w^{(0)}/|\Gamma|$ in different ways and only write it this way for convenience's sake. If Γ is an infinite set, we leave away all empty contexts without loss of generality.) From Gunasekar et al. (2018b), equation (24), we know that we can write

$$w_\gamma^{(t)} = \left(\nabla_w \mathcal{P}_\gamma(w) \hat{z} + \delta_\gamma^{(t)} \right) \sum_{u < t} \eta_u p_\gamma(u) g(u)^{\nu_\gamma - 1}, \tag{34}$$

where

$$p(u) = \|z_\gamma^{(t)}\|_2, \quad g(u) = \|w^{(t)}\|_2, \tag{35}$$

and $\delta_\gamma^{(t)} \to 0$ is analogous to $\delta_3^{(t)}$ in their article.

For the purpose of a shorter notation, we now define

$$k_\gamma^{(t)} := \sum_{u<t} \eta_u p_\gamma(u) g(u)^{\nu_\gamma - 1}, \tag{36}$$

which means that

$$w^{(t)} = \sum_{\gamma \in \Gamma} k_\gamma^{(t)} \left(\nabla_w \mathcal{P}_\gamma(w^{(\infty)}) \hat{z}^{(\infty)} + \delta_\gamma^{(t)} \right). \tag{37}$$

This $k_\gamma^{(t)}$ serves two purposes: it encodes the fact that $w^{(t)}$ diverges (since $k_\gamma^{(t)}$ diverges), and it specifies the scale of contributions to each contextwise gradient. We thus disentangle these two purposes by defining the overall scale

$$k^{(t)} := \sum_{\gamma \in \Gamma} k_\gamma^{(t)}, \tag{38}$$

and the weights

$$\hat{k}_\gamma^{(t)} := k_\gamma^{(t)} / k^{(t)}. \tag{39}$$

(Since all $k_\gamma^{(t)}$ diverge, we consider a large enough t such that all $k_\gamma^{(t)} > 0$.) Defining

$$\tilde{\delta}^{(t)} = \sum_{\gamma \in \Gamma} \hat{k}_\gamma^{(t)} \delta_\gamma^{(t)} \to 0, \tag{40}$$

we can rewrite $w^{(t)}$ in a way that makes it more obvious how normalization will affect it:

$$w^{(t)} = k^{(t)} \left(\sum_\gamma \hat{k}_\gamma^{(t)} \nabla_w \mathcal{P}_\gamma(w^{(\infty)}) \hat{z}^{(\infty)} + \tilde{\delta}^{(t)} \right). \tag{41}$$

We thus know that the normalized sequence of weights is given by

$$\frac{w^{(t)}}{\|w^{(t)}\|} = \frac{\sum_\gamma \hat{k}_\gamma^{(t)} \nabla_w \mathcal{P}_\gamma(w^{(\infty)}) \hat{z}^{(\infty)} + \tilde{\delta}^{(t)}}{\left\| \sum_\gamma \hat{k}_\gamma^{(t)} \nabla_w \mathcal{P}_\gamma(w^{(\infty)}) \hat{z}^{(\infty)} + \tilde{\delta}^{(t)} \right\|}. \tag{42}$$

We are left with two observation that will allow us to determine the limit of this equation and thus prove the theorem. First, we can infer, analogous to Claim 1 in Gunasekar et al. (2018b),

$$\left\| \sum_\gamma \hat{k}_\gamma^{(t)} \nabla_w \mathcal{P}_\gamma(w^{(\infty)}) \hat{z}^{(\infty)} \right\| > 0, \tag{43}$$

for large enough t that all $\hat{k}_\gamma^{(t)} > 0$. Second, we must consider the limit of $\hat{k}_\gamma^{(t)}$. Here, we face a small complication introduced by the context-gated setup: though we know that $\hat{k}_\gamma^{(t)}$ is upper bounded by 1, we do not know whether it converges. It is possible, for instance, that this sequence oscillates between two different values. However, because the sequence is bounded, we know that it has a convergent subsequence $(t_s)_{s \in \mathbb{N}}$. (For example, this subsequence may take into account only one of the two values between which $\hat{k}_\gamma^{(t)}$ may oscillate.) We choose some bounded subsequence and define the limit

$$\hat{k}_\gamma^{(t_s)} \to \hat{k}_\gamma. \tag{44}$$

These two observations together allow us to infer that

$$\hat{w} = \lim_{t \to \infty} \frac{w^{(t)}}{\|w^{(t)}\|} = \frac{\sum_\gamma \hat{k}_\gamma \nabla_w \mathcal{P}_\gamma(\hat{w}) \hat{z}}{\left\| \sum_\gamma \hat{k}_\gamma \nabla_w \mathcal{P}_\gamma(\hat{w}) \hat{z} \right\|}. \tag{45}$$

This is clearly a linear sum of support vectors and since w^* is a positive scaling of $\hat{w}$, it, too can be written as a linear sum of support vectors, proving the theorem.

B.2 The Inductive Bias of GLNs with Two Layers

B.2.1 Architectural Constraints

We stated in the main article that architectural constraints are given by the fact that the difference between two predictors is invariant to the contexts they share, and that for a given architecture, we can choose exactly $(C-1)H+1$ linear predictors freely. Lemma B.6 formalizes the first statement, Lemma B.7 the second one.

Lemma B.6. *Consider two pairs of context $\gamma^{(1)}, \gamma^{(2)}$ and $\tilde{\gamma}^{(1)}, \tilde{\gamma}^{(2)}$. We call this pair* different only in contexts they share *if the following two conditions are true:*

1. *If $\gamma_h^{(1)} = \gamma_h^{(2)}$, then $\tilde{\gamma}_h^{(1)} = \tilde{\gamma}_h^{(2)}$,*

2. *If $\gamma_h^{(1)} \neq \gamma_h^{(2)}$, then $\gamma^{(i)} = \tilde{\gamma}^{(i)}$ for $i = 1, 2$.*

For any such pair of pairs,

$$\beta_{\gamma^{(2)}} - \beta_{\gamma^{(1)}} = \beta_{\tilde{\gamma}^{(2)}} - \beta_{\tilde{\gamma}^{(1)}}. \tag{46}$$

Proof. We know that

$$\beta_{\gamma^{(2)}} - \beta_{\gamma^{(1)}} = \sum_{h=1}^{H} w_h^{(2)} \left(w_{h\gamma_h^{(2)}}^{(1)} - w_{h\gamma_h^{(1)}}^{(1)} \right). \tag{47}$$

If we prove that

$$w_{h\gamma_h^{(2)}}^{(1)} - w_{h\gamma_h^{(1)}}^{(1)} = w_{h\tilde{\gamma}_h^{(2)}}^{(1)} - w_{h\tilde{\gamma}_h^{(1)}}^{(1)}, \tag{48}$$

for all h, we have proven the lemma.

To prove the equation, we consider two cases. If $\gamma_h^{(1)} = \gamma_h^{(2)}$, both sides of the equation are zero. If $\gamma_h^{(1)} \neq \gamma_h^{(2)}$, then $w_{h\gamma_h^{(i)}}^{(1)} = w_{h\tilde{\gamma}_h^{(i)}}^{(1)}$ for $i = 1, 2$ and again, the equation holds true. $\square$

Lemma B.7. *Consider the set of contexts $\overline{\Gamma}$, where at most one hidden unit's local context is different from one, $\gamma_h \neq 1$. We can pick w to parameterize an arbitrary set of linear predictors β_γ for all $\gamma \in \overline{\Gamma}$ and this set, in turn, uniquely determines all other linear predictors. This means we can pick exactly $|\overline{\Gamma}| = (C-1)H+1$ linear predictors freely.*

Proof. Consider an arbitrary set of predictors β_γ, $\gamma \in \overline{\Gamma}$. Let us denote by $\gamma_1 = (1)_{h=1,\ldots,H}$ the vector where all contexts are 1 and by $\gamma_{ch} = (1 + (c-1)\delta_{h'h})_{h'=1,\ldots,H}$ the vector where all entries are 1 except for the h-th context, which is c. Any $\gamma \in \overline{\Gamma}$ can be expressed as either γ_1 or γ_{hc}.

We then define the weights $w^{(1)} \in \mathbb{R}^{C \times H \times D}$ as

$$\begin{aligned} w_{11}^{(1)} &:= \beta_{\gamma_1}, \quad \forall_{c>1} w_{c1}^{(1)} := \beta_{\gamma_{c1}}, \\ \forall_{h>1} w_{1h}^{(1)} &:= 0, \qquad \forall_{c>1} w_{ch}^{(1)} := \beta_{\gamma_{ch}} - \beta_{\gamma_1}, \end{aligned} \tag{49}$$

and $w^{(2)} = (1)_{h=1,\ldots,H}$.

Using this definition, we can see that the predictor $\tilde{\beta}$ parameterized by w is identical to β for all $\gamma \in \overline{\Gamma}$:

$$\tilde{\beta}_{\gamma_1} = w_1^{(2)} w_{11}^{(1)} = \beta_{\gamma_1}, \forall_{c>1} \tilde{\beta}_{\gamma_{c1}} = w_1^{(2)} w_{c1}^{(1)} = \beta_{\gamma_{c1}},$$

$$\forall_{h>1} \forall_{c>1} \tilde{\beta}_{\gamma_{ch}} = w_{11}^{(1)} + w_{ch}^{(1)} = \beta_{\gamma_{ch}}.$$

We now show that, given this set of predictors, we can use (46) to define all other predictors. We use finite induction: consider some $h = 1, \ldots, H$ and suppose that we have already uniquely defined all β_γ if $\gamma_k = 1$ for all $k \geq h$ (for $h = 1$, this is trivially true as $\gamma_1 \in \overline{\Gamma}$). For the induction, we must uniquely define any context γ where $\gamma_h = c \neq 1$ and $\gamma_k = 1$ for all $k \geq h+1$.

For any such context we consider the pair of contexts γ, γ_{ch} and $\tilde{\gamma}, \gamma_1$, where $\tilde{\gamma}_k = \gamma_k$ for all $k \neq h$ and $\tilde{\gamma}_h = 1$. Since $\gamma_{ch}, \gamma_1 \in \overline{\Gamma}$, $\beta_{\gamma_{ch}}, \beta_{\gamma_1}$ have already been defined. Since $\tilde{\gamma}_k = 1$ for all $k \geq h$, $\beta_{\tilde{\gamma}}$ has been defined by the induction's assumption. Since these pairs are different only in contexts they share (namely only in dimension h), this immediately implies that

$$\beta_\gamma = \beta_{\gamma_{ch}} + \beta_{\tilde{\gamma}} - \beta_{\gamma_1}, \tag{50}$$

is uniquely defined as well. $\square$

B.2.2 The GLN Norm

To prove (20), we prove the following statement:

Proposition B.8. *For a deep GLN with two layers, under Assumptions B.1 to B.5, $\hat{\beta}$ is proportional to a minimum of*

$$\|\beta\|_{\mathrm{GLN}}, \quad \text{s.t. } y^{(n)}\langle\beta_{\gamma^{(n)}}, x^{(n)}\rangle \geq 1, \tag{51}$$

where

$$\|\beta\|_{\mathrm{GLN}} = \min_{\zeta}\sum_{h=1}^{H}\|\zeta_h\|_2, \quad \text{s.t. } \beta_\gamma = \sum_{h=1}^{H}\zeta_{h\gamma_h}. \tag{20 restated}$$

Proof. Since w converges in direction, ζ also converges in direction. We aim to prove that its limit direction, $\hat{\zeta}$, is proportional to a stationary point of

$$\sum_{h=1}^{H}\|\zeta_h\|_2, \quad \text{s.t. } y^{(n)}\left\langle\sum_{h=1}^{H}\zeta_{h\gamma_h}, x^{(n)}\right\rangle \geq 1 \tag{52}$$

Since (52) is convex (because its objective function is convex and its constraints are linear), this will be sufficient to prove the proposition.

From Theorem 3.1, we know that $\hat{w}$ is proportional to a stationary point of (6) and can therefore be characterized as

$$\hat{w}_{ch}^{(1)} = \hat{w}_h^{(2)}\sum_{\gamma:\gamma_h=c}\phi_\gamma, \quad \hat{w}_h^{(2)} = \sum_{c=1}^{C}\left\langle\hat{w}_{ch}^{(1)}, \sum_{\gamma:\gamma_h=c}\phi_\gamma\right\rangle. \tag{53}$$

,

where ϕ_γ is the sum of support vectors for a given context:

$$\phi_\gamma = \sum_{n\in S_\gamma}\lambda_n x^{(n)}y^{(n)}. \tag{54}$$

$\hat{w}$ parameterizes $\hat{\zeta}$ as

$$\hat{\zeta}_h = \hat{w}_h^{(2)}\hat{w}_h^{(1)} \in \mathbb{R}^{C\times D} \tag{55}$$

A stationary point of (52) is characterized by the equation

$$\zeta_{hc} = \|\zeta_h\|_2\sum_{\gamma:\gamma_h=c}\phi_\gamma, \tag{56}$$

so since the margin constraints are equivalent, we must only show that $\hat{\zeta}$ satisfies (56).

If $\hat{w}_h^{(2)} = 0$, $\hat{\zeta}_h = 0$ and (56) holds true. We therefore assume that $\hat{w}_h^{(2)} \neq 0$. From (53), we can infer that

$$\hat{w}_h^{(2)} = \hat{w}_h^{(2)}\sum_{c=1}^{C}\left\|\sum_{\gamma:\gamma_h=c}\phi_\gamma\right\|_2^2,$$

which, since $\hat{w}_h^{(2)} \neq 0$, implies that

$$\sum_{c=1}^{C}\left\|\sum_{\gamma:\gamma_h=c}\phi_\gamma\right\|_2^2 = 1. \tag{57}$$

That in turn implies that

$$\|\hat{w}_h^{(1)}\|_2^2 = \sum_{c=1}^{C}\|\hat{w}_{hc}^{(1)}\|_2^2 = \left(\hat{w}_h^{(2)}\right)^2\sum_{c=1}^{C}\left\|\sum_{\gamma:\gamma_h=c}\phi_\gamma\right\|_2^2 = \left(\hat{w}_h^{(2)}\right)^2, \tag{58}$$

and therefore

$$\hat{\zeta}_{hc} = \hat{w}_h^{(2)} \hat{w}_{hc}^{(1)} = \left(\hat{w}_h^{(2)}\right)^2 \sum_{\gamma:\gamma_h=c} \phi_\gamma = |\hat{w}_h^{(2)}| \|\hat{w}_h^{(1)}\|_2 \sum_{\gamma:\gamma_h=c} \phi_\gamma = \|\hat{\zeta}_h\|_2 \sum_{\gamma:\gamma_h=c} \phi_\gamma,$$

which proves (56).

To prove the reverse direction, we must show that if $\hat{\zeta}$ satisfies (56), we can find a $\hat{w}$ satisfying (53) such that $\hat{\zeta}_h = \hat{w}_h^{(2)} \hat{w}_h^{(1)}$. If $\hat{\zeta}_h = 0$, we define $\hat{w}_h^{(2)} = 0$ and $\hat{w}_h^{(1)} = 0$. We therefore assume $\hat{\zeta}_h \neq 0$ and define

$$\hat{w}_h^{(1)} := \hat{\zeta}_h / \sqrt{\|\hat{\zeta}_h\|_2}, \quad \hat{w}_h^{(2)} := \sqrt{\|\hat{\zeta}_h\|_2}. \tag{59}$$

Clearly this parameterizes ζ and so all that is left to show is that it satisfies (53). (56) implies

$$\hat{w}_h^{(1)} = \hat{\zeta}_h / \sqrt{\|\hat{\zeta}_h\|_2} = \sqrt{\|\hat{\zeta}_h\|_2} \sum_{\gamma:\gamma_h=c} \phi_\gamma = w_h^{(2)} \sum_{\gamma:\gamma_h=c} \phi_\gamma.$$

Moreover, (56) immediately implies (57) and thus

$$\sum_{c=1}^C \left\langle \hat{w}_{hc}^{(1)}, \sum_{\gamma:\gamma_h=c} \phi_\gamma \right\rangle = \frac{1}{\sqrt{\|\hat{\zeta}_h\|_2}} \sum_{c=1}^C \left\langle \hat{\zeta}_{hc}, \sum_{\gamma:\gamma_h=c} \phi_\gamma \right\rangle =$$

$$\frac{1}{\sqrt{\|\hat{\zeta}_h\|_2}} \sum_{c=1}^C \left\langle \|\hat{\zeta}_h\|_2 \sum_{\gamma:\gamma_h=c} \phi_\gamma, \sum_{\gamma:\gamma_h=c} \phi_\gamma \right\rangle = \sqrt{\|\hat{\zeta}_h\|_2} \sum_{c=1}^C \left\| \sum_{\gamma:\gamma_h=c} \phi_\gamma \right\|_2^2 = \sqrt{\|\hat{\zeta}_h\|_2} = w_h^{(2)},$$

which proves (53) and thus the proposition. $\qquad\square$

Note that we can prove the inductive bias of FreLUs (22) in an analogous manner.

We now prove the following statement:

Proposition B.9. *For $C = H = 2$, the norm is given by*

$$\|\beta\|_{\mathrm{GLN}}^2 = (\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)^2 + \|\beta_{12} + \beta_{21}\|_2^2. \tag{60}$$

Equivalently,

$$\|\beta\|_{\mathrm{GLN}}^2 = \|\beta\|_2^2 + \frac{1}{2} \sum_{i,j} \|\beta_{ij} - \beta_{\bar{i}j}\|_2 \|\beta_{ij} - \beta_{i\bar{j}}\|_2, \tag{21 restated}$$

where $\bar{k} = 2$ if $k = 1$ and $\bar{k} = 1$ if $k = 2$.

Proof. We must minimize

$$\min_{\zeta \in \mathbb{R}^{2 \times 2 \times D}} \|\zeta_1\|_2 + \|\zeta_2\|_2, \quad \text{s.t. } \beta_{ij} = \zeta_{1i} + \zeta_{2j}, \tag{61}$$

for $i, j = 1, 2$. For any ζ_{11}, β will uniquely determine ζ such that the constraint holds. More specifically,

$$\zeta_{21} = \beta_{11} - \zeta_{11}, \quad \zeta_{22} = \beta_{12} - \zeta_{11}, \quad \zeta_{12} = \beta_{21} - \zeta_{21} = \beta_{21} - \beta_{11} + \zeta_{11}.$$

Thus we must minimize the function

$$g(\zeta_{11}) := \|\zeta_1\|_2 + \|\zeta_2\|_2 = \sqrt{g_1(\zeta_{11})} + \sqrt{g_2(\zeta_{11})},$$
$$g_1(\zeta_{11}) := \|\zeta_{11}\|_2^2 + \|\beta_{21} - \beta_{11} + \zeta_{11}\|_2^2, \tag{62}$$
$$g_2(\zeta_{11}) := \|\beta_{11} - \zeta_{11}\|_2^2 + \|\beta_{12} - \zeta_{11}\|_2^2,$$

where we leave the dependence of g on β implicit to simplify the notation. g is convex and thus minimized by ζ_{11} if and only if

$$0 = \frac{\partial g}{\partial \zeta_{11}}(\zeta_{11}),$$

18

which is equivalent to

$$0 = \sqrt{g_1(\zeta_{11})}\frac{\partial g_2}{\partial \zeta_{11}}(\zeta_{11}) + \sqrt{g_2(\zeta_{11})}\frac{\partial g_1}{\partial \zeta_{11}}(\zeta_{11}) =$$

$$\sqrt{g_1(\zeta_{11})}(2\zeta_{11} - \beta_{11} - \beta_{12}) + \sqrt{g_2(\zeta_{11})}(2\zeta_{11} + \beta_{21} - \beta_{11}).$$

This means that

$$2g(\zeta_{11})\zeta_{11} = \sqrt{g_1(\zeta_{11})}(\beta_{11} + \beta_{12}) + \sqrt{g_2(\zeta_{11})}(\beta_{11} - \beta_{21}).$$

Defining

$$\alpha = \frac{\sqrt{g_1(\zeta_{11})}}{g(\zeta_{11})} \in [0, 1], \tag{63}$$

we can express

$$\zeta_{11} = \tfrac{1}{2}\left(\alpha(\beta_{12} + \beta_{11}) + (1-\alpha)(\beta_{11} - \beta_{21})\right) = \tfrac{1}{2}\left(\beta_{11} + \alpha\beta_{12} - (1-\alpha)\beta_{21}\right).$$

We have therefore reduced our search space to α, and reexpress the norm in terms of α:

$$
\begin{aligned}
\tilde{g}(\alpha) :=& \sqrt{\tilde{g}_1(\alpha)} + \sqrt{\tilde{g}_2(\alpha)}, \\
\tilde{g}_1(\alpha) :=& g_1(\zeta_{11}) = \|\zeta_{11}\|_2^2 + \|\beta_{21} - \beta_{11} + \zeta_{11}\|_2^2 = \\
& \tfrac{1}{4}(\|\beta_{11} - \beta_{21} + \alpha(\beta_{12} + \beta_{21})\|_2^2 + \|\beta_{21} - \beta_{11} + \alpha(\beta_{12} + \beta_{21})\|_2^2) = \\
& \tfrac{1}{2}(\|\beta_{11} - \beta_{21}\|_2^2 + \alpha^2\|\beta_{12} + \beta_{21}\|_2^2), \\
\tilde{g}_2(\alpha) :=& g_2(\zeta_{11}) = \|\beta_{11} - \zeta_{11}\|_2^2 + \|\beta_{12} - \zeta_{11}\|_2^2 = \\
& \tfrac{1}{4}(\|\beta_{11} + \beta_{21} - \alpha(\beta_{12} + \beta_{21})\|_2^2 + \|(2-\alpha)(\beta_{12} + \beta_{21}) - (\beta_{11} + \beta_{21})\|_2^2) = \\
& \tfrac{1}{2}(\|\beta_{11} + \beta_{21}\|_2^2 + (1 + (1-\alpha)^2)\|\beta_{12} + \beta_{21}\|_2^2 - 2\langle\beta_{12} + \beta_{21}, \beta_{11} + \beta_{21}\rangle) = \\
& \tfrac{1}{2}(\|\beta_{11} - \beta_{12}\|_2^2 + (1-\alpha)^2\|\beta_{12} + \beta_{21}\|_2^2).
\end{aligned} \tag{64}
$$

We now wish to show that

$$\min_\alpha \tilde{g}(\alpha) = \sqrt{\tfrac{1}{2}}\|\beta\|_{\text{GLN}}, \tag{65}$$

according to the first definition, (60). If $\beta_{12} + \beta_{21} = 0$, this follows immediately. We therefore assume that $\beta_{12} + \beta_{21} \neq 0$. Any minimal α must satisfy

$$0 = \frac{\partial \tilde{g}(\alpha)}{\partial \alpha} = \frac{\alpha\|\beta_{12} + \beta_{21}\|_2^2}{\sqrt{\tilde{g}_1(\alpha)}} - \frac{(1-\alpha)\|\beta_{12} + \beta_{21}\|_2^2}{\sqrt{\tilde{g}_2(\alpha)}}. \tag{66}$$

Since $\alpha \in [0, 1]$, both α and $1 - \alpha$ are nonnegative, and (66) is equivalent to

$$\tilde{g}_1(\alpha)(1-\alpha)^2 = \tilde{g}_2(\alpha)\alpha^2, \tag{67}$$

which, in turn, reduces to

$$\alpha^2\|\beta_{11} - \beta_{12}\|_2^2 = (1-\alpha)^2\|\beta_{11} - \beta_{21}\|_2^2. \tag{68}$$

Taking the square root and rearranging results in

$$\alpha = \frac{\|\beta_{11} - \beta_{21}\|_2}{\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2}. \tag{69}$$

This implies

$$\tilde{g}_1(\alpha) = \tfrac{1}{2}\|\beta_{11} - \beta_{21}\|_2^2\left(1 + \frac{\|\beta_{12} + \beta_{21}\|_2^2}{(\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)^2}\right).$$

Since

$$1 - \alpha = \frac{\|\beta_{11} - \beta_{12}\|_2}{\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2},$$

we can analogously infer that

$$\tilde{g}_2(\alpha) = \tfrac{1}{2}\|\beta_{11} - \beta_{12}\|_2^2\left(1 + \frac{\|\beta_{12} + \beta_{21}\|_2^2}{(\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)^2}\right).$$

Finally, these identities imply that

$$g_3(\alpha^*) = \sqrt{\tfrac{1}{2}}(\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)\sqrt{1 + \frac{\|\beta_{12} + \beta_{21}\|_2^2}{(\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)^2}} =$$
$$\sqrt{\tfrac{1}{2}}\sqrt{(\|\beta_{11} - \beta_{12}\|_2 + \|\beta_{11} - \beta_{21}\|_2)^2 + \|\beta_{12} + \beta_{21}\|_2^2},$$

which proves (65) and thus (60).

All that is left to show is that (60) is equivalent to (21). To do so, we note that

$$\beta_{22} - \beta_{21} = \beta_{12} - \beta_{11}. \tag{16 recalled}$$

Therefore

$$\|\beta\|_2^2 = \|\beta_{11}\|_2^2 + \|\beta_{12}\|_2^2 + \|\beta_{21}\|_2^2 + \|\beta_{12} + \beta_{21} - \beta_{11}\|_2^2 =$$
$$2(\|\beta_{11}\|_2^2 + \|\beta_{12}\|_2^2 + \|\beta_{21}\|_2^2 + \langle\beta_{12}, \beta_{21}\rangle - \langle\beta_{11}, \beta_{12}\rangle - \langle\beta_{11}, \beta_{21}\rangle), \tag{70}$$

and

$$\|\beta_{11} - \beta_{12}\|_2^2 + \|\beta_{11} - \beta_{21}\|_2^2 + \|\beta_{12} + \beta_{21}\|_2^2 =$$
$$2(\|\beta_{11}\|_2^2 + \|\beta_{12}\|_2^2 + \|\beta_{21}\|_2^2 + \langle\beta_{12}, \beta_{21}\rangle - \langle\beta_{11}, \beta_{12}\rangle - \langle\beta_{11}, \beta_{21}\rangle) = \|\beta\|_2^2. \tag{71}$$

This, in turn implies that

$$\|\beta\|^2 = \|\beta\|_2^2 + 2\|\beta_{11} - \beta_{12}\|_2\|\beta_{11} - \beta_{21}\|_2 = \|\beta\|_2^2 + \tfrac{1}{2}\sum_{i,j}\|\beta_{ij} - \beta_{\bar{i}j}\|_2\|\beta_{ij} - \beta_{i\bar{j}}\|_2, \tag{72}$$

where the latter equality follows from the fact that for all i, j

$$\|\beta_{ij} - \beta_{\bar{i}j}\|_2\|\beta_{ij} - \beta_{i\bar{j}}\|_2 = \|\beta_{11} - \beta_{12}\|_2\|\beta_{11} - \beta_{21}\|_2, \tag{73}$$

as a result of (16). □

B.3 Extending Gunasekar et al. (2018a) to exponential-like losses

We here consider the same set of loss functions as Soudry et al. (2018) and Lyu and Li (2020). We call this class *exponential-like losses*.

Definition B.10. We call $\ell : \mathbb{R} \to \mathbb{R}$ *exponential-like* if and only if the function satisfies the following assumptions:

1. ℓ is monotonically decreasing, differentiable, and $\lim_{u\to\infty}\ell(u) \to 0$.

2. ℓ' is Lipschitz continuous.

3. $-\ell'$ has a tight exponential tail, i.e. there exist positive constants c, a, μ_+, μ_-, u_+, u_- such that
$$\forall_{u>u_+} -\ell'(u) \le c(1 + \exp(-\mu_+ u))\exp(-au), \tag{74}$$
$$\forall_{u>u_-} -\ell'(u) \ge c(1 - \exp(-\mu_- u))\exp(-au). \tag{75}$$

Gunasekar et al. (2018a,b) only consider the exponential loss, but note that they expect their results to generalize towards exponential-like losses. More specifically, they prove Lemma 8 in Gunasekar et al. (2018a) only for the exponential loss. We here close the small gap their results leave by extending this lemma to exponential-like losses.

Lemma B.11. *For almost all linearly separable datasets $\mathcal{D} = \left((x^{(n)}, y^{(n)})\right)_{n=1,\ldots,N}$, consider the loss function*

$$\mathcal{L}(\beta) := \sum_{n=1}^{N}\ell(y^{(n)}\langle\beta, x^{(n)}\rangle), \tag{76}$$

where ℓ is an exponential-like loss.

Any sequence $\beta^{(t)}$ such that

$$\mathcal{L}(\beta^{(t)}) \to 0, \tag{77}$$
$$\beta^{(t)}/\|\beta^{(t)}\|_2 \to \hat{\beta}, \tag{78}$$
$$-\nabla_\beta\mathcal{L}(\beta^{(t)})/\|\nabla_\beta\mathcal{L}(\beta^{(t)})\|_2 \to \hat{z}, \tag{79}$$

for some $\hat{\beta}$, $\hat{z}$. Let

$$S_{\mathcal{D}} := \left\{ n \mid y^{(n)} \langle \hat{\beta}, x^{(n)} \rangle = \min_n y^{(n)} \langle \hat{\beta}, x^{(n)} \rangle \right\} \tag{80}$$

be the support. *Then there exists a sequence of nonnegative numbers* $(\alpha_n)_{n \in S_{\mathcal{D}}}$, $\alpha_n \geq 0$ *such that*

$$\hat{z} = \sum_{n \in S_{\mathcal{D}}} \alpha_n y_n x_n. \tag{81}$$

Proof. To prove the lemma, we reduce the general, exponential-like case to the special case where ℓ is the exponential loss.

The gradient is given by

$$-\nabla_\beta \mathcal{L}(\beta) = \sum_{n=1}^{N} -\ell'(y^{(n)} \langle \beta, x^{(n)} \rangle) y^{(n)} x^{(n)}. \tag{82}$$

We now decompose this gradient into the special case where ℓ is the exponential loss, and the deviation of ℓ from the exponential loss. We write

$$-\nabla_\beta \tilde{\mathcal{L}}(\beta) = \sum_{n=1}^{N} -c \exp(-a y^{(n)} \langle \beta, x^{(n)} \rangle) y^{(n)} x^{(n)}, \tag{83}$$

where c, a are the constants from the definition of ℓ. (Put differently, $\tilde{\mathcal{L}}$ defines the dataset's loss if we were using the exponential loss.)

We know that

$$m_n^{(t)} := y^{(n)} \langle \beta^{(t)}, x^{(n)} \rangle, \tag{84}$$

diverges. Thus we can assume a large enough t such that $y^{(n)} \langle \beta^{(t)}, x^{(n)} \rangle > u_+, u_-$ for all n. We can now write

$$-\ell'(m_n^{(t)}) = c \exp(-a m_n^{(t)}) + c \exp(-a m_n^{(t)}) \delta_n^{(t)}, \tag{85}$$

where, due to the fact that $-\ell'$ has a tight exponential tail, we are guaranteed that

$$-\exp(-\mu_- m_n^{(t)}) \leq \delta_n^{(t)} \leq \exp(-\mu_+ m_n^{(t)}), \tag{86}$$

for large enough t. We have thus sucessfully decomposed the gradient:

$$-\nabla_\beta \mathcal{L}(\beta^{(t)}) = -\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)}) + c \sum_{n=1}^{N} \exp(-a m_n^{(t)}) \delta_n^{(t)} y^{(n)} x^{(n)}. \tag{87}$$

We now want to use this decomposition to prove that

$$-\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)}) / \|\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})\|_2 \to \hat{z}, \tag{88}$$

as well. We know that

$$\delta^{(t)} := \frac{\sum_{n=1}^{N} \exp(-a m_n^{(t)}) \delta_n^{(t)} y^{(n)} x^{(n)}}{\exp(-a m_n^{(t)}) y^{(n)} x^{(n)}} \tag{89}$$

is bounded by $\tilde{\delta}^{(t)} := \max_n \delta_n^{(t)}$, and therefore $\delta^{(t)} \to 0$. This, in turn, implies that

$$(1 + \delta^{(t)}) / \|1 + \delta^{(t)}\|_2 \to 1. \tag{90}$$

We can rewrite

$$-\nabla_\beta \mathcal{L}(\beta^{(t)}) = -\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})(1 + \delta^{(t)}), \tag{91}$$

and therefore infer

$$\hat{z} = \lim_{t \to \infty} \frac{-\nabla_\beta \mathcal{L}(\beta^{(t)})}{\|\nabla_\beta \mathcal{L}(\beta^{(t)})\|} = \lim_{t \to \infty} \frac{-\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})}{\|\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})\|} \frac{1 + \delta^{(t)}}{\|1 + \delta^{(t)}\|} = \lim_{t \to \infty} \frac{-\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})}{\|\nabla_\beta \tilde{\mathcal{L}}(\beta^{(t)})\|} \tag{92}$$

Clearly, $\tilde{\mathcal{L}}(\beta^{(t)}) \to 0$, which allows us to apply Lemma 8 from Gunasekar et al. (2018a) and proves the more general lemma. $\square$
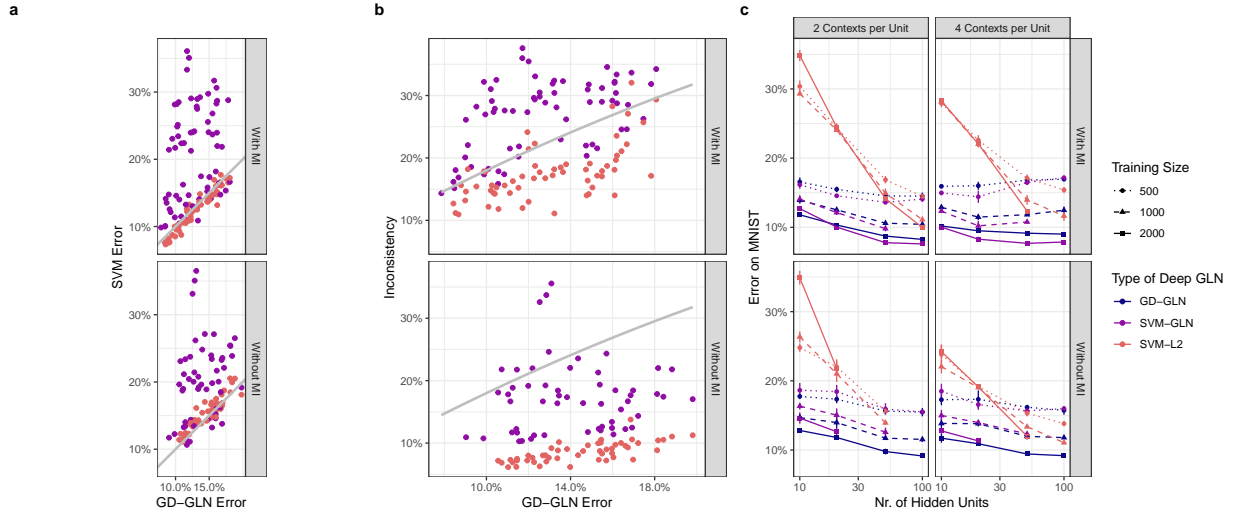
Figure 4: Full data from our experiments on GLNs. **a** The error of the GD-GLN plotted against the error of the SVMs. The grey line represents identical performance. The two panels correspond to the network that used median initialization ("With MI") and those that did not ("Without MI"). **b** Inconsistency between the GD-GLN and the SVM plotted against the GD-GLN's error on MNIST. The grey line represents the inconsistency we would expect from a predictor with matched error rate, but no further correlation with the network (see Footnote 6). **c** Error of the SVMs and the GD-GLN plotted against the number of hidden units.

C Extended Data

For clarity's sake, we only showed a subset of the data in Figures 2 and 3. Figures 4 and 5 depict the full data from our experiments on GLNs and ReLU networks, respectively.

C.1 Gated Linear Networks

Since Veness et al. (2017) did not use median initialization (MI), we trained the GLNs without MI as well. This resulted in slightly worse performance, but the SVM-GLN was still more consistent with the GD-GLN than the SVM-L2 (Figure 4a,b).

In addition, Figure 4c, unlike Figure 2d, also depicts the networks with four contexts per unit. Remarkably, for 100 hidden units, the SVM-L2 is beginning to outperform the SVM-GLN. Other than that, the interpretation of the data remained unaffected. Without MI, the data was qualitatively similar, as well, except that the GD-GLN tended to slightly outperform the SVM-GLN.

C.2 Frozen-Gate ReLU Networks

In Figure 3, we depicted the networks trained with a momentum of 0.9. Since our theorem technically only holds for gradient descent without momentum (though Soudry et al. (2018) saw qualitatively similar behavior with momentum, as well), we additionally trained networks without momentum. Since they do not rely on the ReLU networks at all, this did not change the SVM-RC. It changed the SVM-LC and SVM-HL only insofar as the model from which they used the contexts and hidden layer, respectively, changed. The interpretation of the data remained unaffected. Most notably, both performance and consistency with the SVMs were a bit worse. This is consistent with the interpretation that momentum speeds up training to a limit that is, in part, characterized by the SVMs, and that moving closer to this limit improves performance.
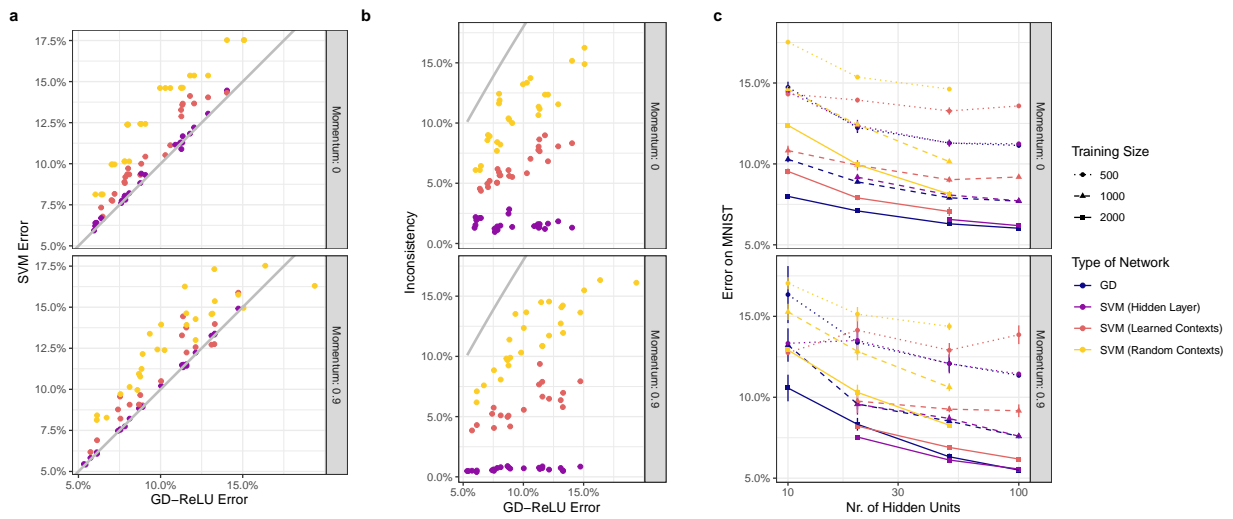
Figure 5: Full data from our experiments on ReLU networks. **a** The error of the GD-ReLU plotted against the SVMs. The grey line represents identical performance. The two panels correspond to training with and without momentum. **b** Inconsistency between the GD-ReLU and the SVMs plotted against the GD-ReLU's error on MNIST. The grey line represents the inconsistency we would expect from a predictor with matched error rate, but no further correlation with the network (see Footnote 6). **c** Error of the SVMs and the GD-ReLU plotted against the number of hidden units.