

392 **A Proof of Theorem 1**

393 For the simplicity of our argument, we work with the following assumptions on the training set.

394 **Assumption 1.** *Label noise: Exactly p fraction of the sample points have their first coordinate*
 395 *flipped.*

396 **Assumption 2.** *Orthogonality: the non-zero coordinates are distinct for all n data points (except*
 397 *for the first coordinate)*

398 Notice that by the fact that $n = o(\sqrt{d})$ and a simple union bound, Assumption 2 holds with high
 399 probability. For each $i \in [d]$, we let $j(i)$ denote the index j that satisfies $\mathbf{x}_j(i) = 1$, if it exists. To
 400 simplify the notation, we assume that all labels y_i are 1; this is without loss of generality, since one
 401 can always replace \mathbf{x}_i with $y_i \mathbf{x}_i$.

402 In order to prove Theorem 1, we will precisely characterize the limiting behavior of SGD in this set-
 403 ting. We remark that as the optimization objective is strongly convex, SGD and GD with appropriate
 404 choices of step size is guaranteed to converge to the global minimum.

405 **Lemma 1** (Convergence of gradient descent). *Assume the setting above. Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be the data*
 406 *matrix. Initialized at point \mathbf{w}_0 , (stochastic) gradient descent converges to*

$$\mathbf{w}' = \mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-1} (\mathbf{1} - \mathbf{X} \mathbf{w}_0) + \mathbf{w}_0, \quad (1)$$

407 *as the number of steps goes to infinity. Moreover, we have*

$$\mathbf{w}'(1) = \left(\frac{n}{n+1} \right) \eta - \frac{\mathbf{s}^T \mathbf{X} \mathbf{w}_0}{n+1} + \mathbf{w}_0(1), \quad (2)$$

408 *where \mathbf{s} is the first column of \mathbf{X} and $\eta = 1 - 2p$, and for $i \neq 1$*

$$\mathbf{w}'(i) = \begin{cases} \mathbf{w}_0(i) & \text{if } \mathbf{x}_j(i) = 0 \text{ for all } \mathbf{x}_j, \\ \beta \left(1 + \mathbf{s}_{j(i)} \left(\frac{-n\eta + \mathbf{s}^T \mathbf{X} \mathbf{w}_0}{n+1} \right) \eta - \mathbf{x}_{j(i)}^T \mathbf{w}_0 \right) + \mathbf{w}_0(i) & \text{if } \mathbf{x}_j(i) = \beta \in \{\pm 1\} \text{ for some } \mathbf{x}_j \end{cases} \quad (3)$$

409 *Proof.* We focus on the gradient descent case; the proof for SGD is analogous. Consider the empirical
 410 loss $\mathcal{L}(\mathbf{w})$. By the definition of gradient descent, the iterations always stay in the affine space
 411 $V = \mathbf{w}_0 + \text{RowSpan}(\mathbf{X})$. Gradient descent solves the linear least squares problem $\min_{\mathbf{w} \in V} \mathcal{L}(\mathbf{w})$.
 412 We claim that \mathbf{w}' is indeed the optimal solution to this program, and thus gradient descent converges
 413 to it.

414 First, one can check $\mathbf{X} \mathbf{X}^T$ is non-singular under Assumption 2, since $\mathbf{X} \mathbf{X}^T = \mathbf{I} + \mathbf{s} \mathbf{s}^T$ and
 415 $\mathbf{s}^T \mathbf{s} \neq -1$. Moreover, $\mathbf{w}' \in V$ since $\mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-1}$ is the orthogonal projector onto the row space
 416 of \mathbf{X} . Finally, we check that \mathbf{w}' achieves zero empirical loss

$$\mathbf{X} \mathbf{w}' = \mathbf{X} \mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-1} (\mathbf{1} - \mathbf{X} \mathbf{w}_0) + \mathbf{X} \mathbf{w}_0 = \mathbf{1}. \quad (4)$$

417 For the first coordinate of \mathbf{w}' , by the Sherman-Morrison formula [12, p. 51],

$$(\mathbf{X} \mathbf{X}^T)^{-1} = (\mathbf{I} + \mathbf{s} \mathbf{s}^T)^{-1} = \mathbf{I} - \frac{\mathbf{s} \mathbf{s}^T}{1 + \mathbf{s}^T \mathbf{s}} = \mathbf{I} - \frac{\mathbf{s} \mathbf{s}^T}{n+1}. \quad (5)$$

418 Substituting this into (1) and simplifying yields claims (2) and (3). \square

419 We now recall Theorem 1:

420 **Theorem 1.** *Consider training a linear classifier via minimizing the empirical square loss using*
 421 *SGD. Let $\varepsilon > 0$ be a small constant and let the initial vector \mathbf{w}_0 satisfy $\mathbf{w}_0(1) \geq -n^{0.99}$, and*
 422 *$|\mathbf{w}_0(i)| \leq 1 - 2p - \varepsilon$ for all $i > 1$. Then, with high probability, sample accuracy approaches 1 and*
 423 *population accuracy approaches $1 - p$ as the number of gradient steps goes to infinity.*

424 *Proof.* Let $k(j)$ denote the non-zero coordinate of \mathbf{x}_j (besides the first coordinate). First we note
 425 that

$$\mathbf{s}^T \mathbf{X} \mathbf{w}_0 = n \mathbf{w}_0(1) + \sum_{j=1}^n \mathbf{s}_j \mathbf{x}_j(k(j)) \mathbf{w}_0(k(j))$$

426 Because each $x_j(k(j))$ is a Bernoulli(p) random variable and every $w_0(k(j)) \leq 1$, we can apply
 427 Chebyshev’s inequality to further deduce that with high probability, $s^T \mathbf{X} \mathbf{w}_0 = n w_0(1) + O(\sqrt{n})$.
 428 Substituting this into (2) of Lemma 1, letting \mathbf{w}' be the weight vector SGD converges to, we obtain

$$\mathbf{w}'(1) = \left(\frac{n}{n+1} \right) \eta - \frac{n w_0(1) + O(\sqrt{n})}{n+1} + w_0(1) = \eta - o(1)$$

429 By (3) of Lemma 1, for every coordinate i such that $x_j(i) = 0$ for all j , we have $w'(i) = w_0(i)$.
 430 Consider a point \mathbf{x} drawn from the population, and let k be the index of the non-zero coordinate of
 431 \mathbf{x} . With high probability, $k \neq k(j)$ for all $j \in [n]$. With probability $1 - p$, $\mathbf{x}(1) = 1$, and in this case
 432 we obtain

$$\langle \mathbf{w}', \mathbf{x} \rangle = w'(1)\mathbf{x}(1) + w_0(k)\mathbf{x}(k) \geq \eta - o(1) - (\eta - \varepsilon) = \varepsilon - o(1)$$

433 For sufficiently large n (corresponding to sufficiently large d) this quantity is always positive, so
 434 with probability approaching $1 - p$ the model correctly classifies \mathbf{x} . \square

435 B Experimental Setup and Results for Sections 3 and 4

436 **Dataset description.** We used the following four datasets in our experiments.

- 437 (i) Binary MNIST: predict whether the image represents a number from 0 to 4 or from 5 to 9.
 438 It admits a linear classifier with accuracy $\approx 87\%$,
- 439 (ii) CIFAR-10 Animals vs Objects: predict whether the image represents an animal or an
 440 object. In order not to enforce bias towards any of the classes we included all the 4 object
 441 classes (airplane, automobile, ship, truck) and only 4 out of 6 of the animal ones (bird, cat,
 442 dog, horse). Hence the number of positive and negative samples are the same. CIFAR-10
 443 Animals vs Objects admits a linear classifier with accuracy $\approx 75\%$,
- 444 (iii) CIFAR-10 First 5 vs Last 5: predict whether an image belongs to any of the first 5 classes
 445 of CIFAR10 (airplane, automobile, bird, cat, deer) or the last 5 classes (dog, frog, horse,
 446 ship, truck). CIFAR-10 First 5 vs Last 5 does not admit a linear classifier with satisfying
 447 accuracy. The best linear classifier achieves accuracy of $\approx 58\%$,
- 448 (iv) High-dimensional Sinusoid: predict $y := \text{sign}(\langle w, x \rangle + \sin\langle w', x \rangle)$ for standard Gaussian
 449 $x \in \mathbb{R}^d$ where $d = 100$. The vector w is also chosen uniformly from S^{d-1} and w' is an
 450 orthogonal vector to the hyperplane. High-dimensional sinusoid admits a linear classifier
 451 with accuracy $\approx 80\%$.

452 In the cases of datasets (i), (ii), and (iii) we created the train and tests sets by relabeling the train and
 453 test sets of MNIST and CIFAR10 with $\{0, 1\}$ labels according to the specific dataset (and excluded
 454 the images that are not relevant). All experiments are repeated 10 times with random initialization;
 455 standard deviations are reported in the figures (shaded area).

456 **Model details.** Our results were consistent across various architectures and hyperparameter choices.
 457 For the Sinusoid distribution we train a 2-layer MLP, with ReLU activations. Each layer is 256
 458 neurons. For the MNIST and CIFAR tasks in section 3 we train a CNN with 4 2D-Conv layers; each
 459 layer has 32 filters of size 3×3 . After the first and second layers we have a 2×2 Max-Pooling layer,
 460 and at the end of the 4 convolutional layers we have two Dense layers of 2000 units. The activations
 461 on all intermediate layers are ReLUs. Across all architectures the last layer is a sigmoid neuron.

462 **Training procedure.** We initialize the neural networks with Uniform Xavier [10]. We note that
 463 in all the experiments we use *vanilla SGD* without regularization (e.g. dropout) since we want to
 464 isolate and investigate purely the effect of the optimization algorithm. We use batch size of 32 for
 465 MLP’s and 64 for CNNs. For the MLP’s in section 3, the learning rate is 0.01. For the CNN’s in
 466 Section 3, the learning rate is 0.001. For Section 4, we train the resnet and all smaller CNN’s using
 467 SGD with momentum 0.9, batch size 128 and learning rate 0.01.

468 **Finding the Conditional Models.** For Section 3, in order to find the linear classifier ℓ that best
 469 explains the initial phase of learning, we do the following. For tasks (i), (ii), and (iv), where there

	Linear Model $\mu(F_{T_0}; L)/I(F_{T_0}; Y)$	Null Model $\mu(F_{T_0}; \tilde{L})/I(F_{T_0}; Y)$
MNIST	0.79	0.52
CIFAR (ii)	0.80	0.31
CIFAR (iii)	0.74	0.02
Sinusoid	0.74	0.27

Table 2: Performance Correlation of Linear Model vs. Null Model

	Simple Model $\mu(F_{T_0}; G_i)/I(F_{T_0}; Y)$	Null Model $\mu(F_{T_0}; \tilde{G}_i)/I(F_{T_0}; Y)$
2-layer CNN	0.68	0.20
4-layer CNN	0.72	0.31
6-layer CNN	0.72	0.40

Table 3: Performance Correlation of simple CNN Models vs. Null Model on CIFAR First vs. Last 5

470 exists a unique optimal linear classifier, we set ℓ by training a linear classifier via SGD with logistic
471 loss on the training set of the distribution.

472 For task (iii), we need to break the symmetry among many linear classifier that are roughly of equal
473 performance. Here, we set ℓ by training a linear classifier via SGD to reproduce the outputs of f_{T_0}
474 on the training set. That is, we label the train set using f_{T_0} (outputting labels in $\{0, 1\}$), and train a
475 linear classifier on these labels. Note that we could also have trained on the sigmoid output of f_{T_0} ,
476 not rounding its output to $\{0, 1\}$; this does not affect results in our experiments.

477 In Section 4, we perform a similar procedure: After training f_∞ , we train simple models g_i on the
478 predictions of f_∞ on the train set, via SGD.

479 **Estimating the Mutual Information.** Let f, g be two classifiers, and y be the true labels. In order
480 to estimate our mutual-information metrics, we use the empirical distribution of (F, G, Y) on the
481 test set. For example, to estimate $I(F; Y|G)$ we use the definition

$$I(F; Y|G) = \sum_{(f, y, g) \in \{0, 1\}^3} p(f, y, g) \log \left(\frac{p(f, y|g)}{p(f|g)p(y|g)} \right) \quad (6)$$

482 where $p(f, y, g)$ is the joint probability density function of (F, Y, G) and $p(f|g)$, etc. are the con-
483 ditional density functions. To estimate this quantity, we first compute the empirical distribution
484 of $(f, y, g) \in \{0, 1\}^3$ over the test set. Let $\hat{p}(f, y, g)$ be this empirical density function. Then we
485 estimate $I(F; Y|G)$ by evaluating Equation 6 using \hat{p} in place of p .

486 **Further Quantitative Results.** In Tables 2 and 3 we provide further quantitative results for our
487 experiments in Sections 3 and 4 respectively.

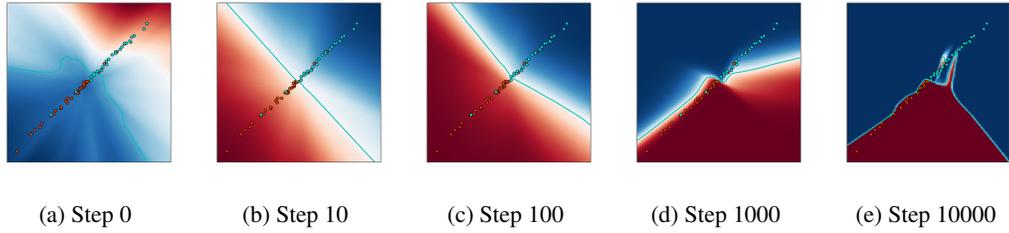


Figure 6: SGD training on a 3-layer, width-100 dense neural network. The data distribution is an (essentially) 1-dimensional Gaussian labeled by a linear classifier with 10% label noise. After a few hundreds of SGD steps the decision boundary is very non-linear outside of the training set, indicating that the neural network does not actually learn a linear classifier but rather a classifier that highly agrees with a linear one. In other words, there exist a linear classifier that explains the predictions of the neural network well, despite the fact that the network itself is highly non-linear.

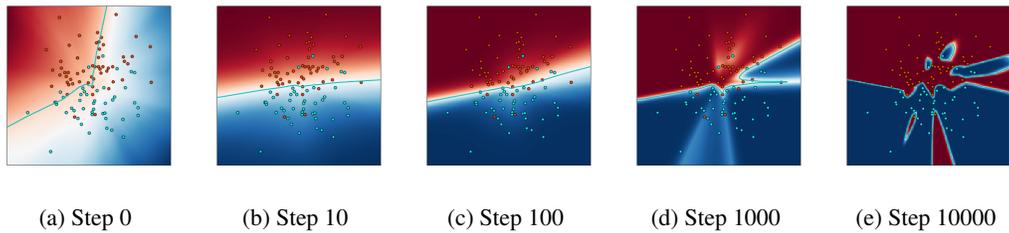


Figure 7: A simplified version of Figure 1. SGD training on a 3-layer, width-100 dense neural network. The data distribution is an isotropic Gaussian in 2-dimensions, labeled by a linear classifier with 10% label noise. The blue line corresponds to the decision boundary of the neural network which becomes more “linear” in the initial stages before starting to overfit to the label noise.

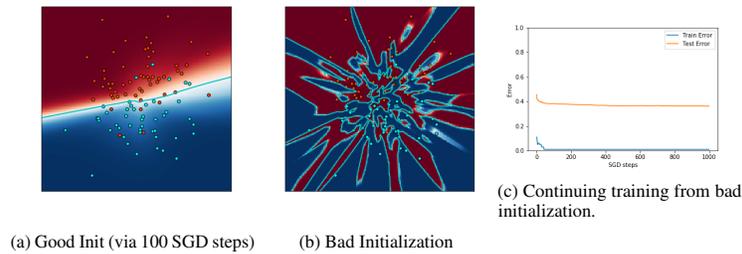


Figure 8: Two different initializations, each with 88% accuracy on the training set. The data distribution is an isotropic Gaussian in 2-dimensions, labeled by a linear classifier with 10% label noise. The good initialization is found after running 100 SGD steps from a random initialization. The bad initialization is found by generating randomly labeled points, and fitting the function to them together with the original training set. One can see that continuing training from the bad initialization allows us to overfit to the training set but does not improve the population accuracy at all. The model here is again a 3-layer, width-100 dense neural network.