

303 Appendix

304 A The Functional Space of ReLU Networks

305 We consider the class of 1D shallow ReLU functions with exactly m neurons:

$$\mathcal{F}_m = \{f_{\theta} : \mathbb{R} \rightarrow \mathbb{R}\}, \quad f_{\theta}(x) = \sum_{i=1}^m c_i [a_i x + b_i]_+. \quad (23)$$

306 It is easy to see that \mathcal{F}_m is a subset of the set CPL_m of *continuous piecewise-linear maps* with at
307 most m knots.

308 **Proposition 5.** *The space \mathcal{F}_m is a strict subset of CPL_m , but it contains CPL_{m-2} .*

309 *Proof.* The fact that $\mathcal{F}_m \subsetneq CPL_m$ can be argued by observing that \mathcal{F}_m has $2m$ degrees of freedom,
310 whereas CPL_m has $2m + 2$ degrees of freedom. More precisely, we can partition the parameter
311 space of f_{θ} based on the signs $\sigma_i = \text{sign}(a_i)$. Within each region, we may write

$$f_{\theta}(x) \sum_{i=1}^m c_i [a_i(x - e_i)]_+ = \sum_{i=1}^m u_i [\sigma_i(x - e_i)]_+, \quad u_i = |a_i|c_i. \quad (24)$$

312 We see that interpolatory constraints on $f_{\theta}(x)$ correspond to linear conditions on u_i . In particular,
313 there are a finite number of functions in \mathcal{F}_m with fixed assigned (generic) values at the knots
314 $e_1 < \dots < e_m$. This contrasts with CPL_m , where the slope in the intervals $[-\infty, e_1]$ and $[e_m, \infty]$
315 may be chosen arbitrarily. Finally, by setting $e_1 = e_2$ and $e_{m-1} = e_m$ we can control the two
316 remaining slopes, so we recover that $CPL_{m-2} \subset \mathcal{F}_m$. \square

317 B Spline Kernels

318 **Proposition 6.** *If either (i) $a(s)$ is identically 1 or (ii) measures of supports $a_+(s)$ and $a_-(s)$ as
319 $m \rightarrow \infty$ on any subinterval of $[k_0, k_1]$ have the same expectation, the kernel $K(x, x')$ defined as*

$$K(x, x') = \int a_+(s)^2 [x - s]_+ [x' - s]_+ ds + \int a_-(s)^2 [s - x]_+ [s - x']_+ ds$$

320 *is a piecewise cubic polynomial in x and x' . In particular, in the overparameterized setting, a function
321 $\hat{f}(x) = \sum_{i=1}^s \alpha_i K_r(x, x_i)$ that interpolates the samples x_1, \dots, x_s will be cubic spline, with knots
322 at these samples.*

323 *Proof.* In case (i), assuming that $x < x'$, we have that

$$\begin{aligned} K_1(x, x') &= \int_{k_0}^{k_1} [x - s]_+ [x' - s]_+ ds = \int_{k_0}^x (x - s)(x' - s) ds \\ &= \left[\frac{1}{3} s^3 - \frac{1}{2} s^2(x + x') + sxx' \right]_{k_0}^x \\ &= -\frac{1}{6} (2k_0 + x - 3x')(k_0 - x)^2 \end{aligned} \quad (25)$$

324 In the second case, each of the finite sums approximating the integrals remains constant, and contains
325 half of the indices, which are distributed densely both for a_+ and a_- . The limit integrals are

$$\begin{aligned} K_2(x, x') &= \frac{1}{2} \int_{k_0}^{k_1} [x - s]_+ [x' - s]_+ ds + \frac{1}{2} \int_{k_0}^{k_1} [-x + s]_+ [-x' + s]_+ ds \\ &= \frac{1}{2} \int_{k_0}^x (x - s)(x' - s) ds + \frac{1}{2} \int_{x'}^{k_1} (s - x)(s - x') ds. \\ &= -\frac{1}{12} (2k_0 + x - 3x')(k_0 - x)^2 + \frac{1}{12} (2k_1 - 3x + x')(k_1 - x')^2. \end{aligned} \quad (26)$$

326 Both $K_1(x, x')$ and $K_2(x, x')$ are piecewise cubic and C^2 in both arguments. This immediately
 327 implies that the solution to the least squares problem $\hat{f}(x) = \sum_{i=1}^s \alpha_i K_t(x, x_i)$ ($t = 1, 2$) is a *cubic*
 328 *spline* interpolating the samples x_1, \dots, x_s . \square

329 Finally the following simple fact shows that the coefficient function $c(x)$ effectively corresponds to
 330 the second derivative (or *linearized curvature*) $f''_{\mathbf{z}}(x)$ of $f_{\mathbf{z}}(x)$.

331 **Lemma 7.** Consider a function of the form

$$f(x) = \int_{k_0}^{k_1} c^+(s)[x-s]_+ ds + \int_{k_0}^{k_1} c^-(s)[s-x]_+ ds, \quad k_0 \leq x \leq k_1. \quad (27)$$

332 Then we have that

$$f'(x) = \int_a^x c^+(\theta) d\theta - \int_x^b c^-(\theta) d\theta, \quad f''(x) = c^+(x) + c^-(x). \quad (28)$$

333 *Proof.* This follows by observing that

$$f(x) = \int_{k_0}^x c^+(s)(x-s) ds + \int_x^{k_1} c^-(s)(s-x) ds, \quad (29)$$

334 and applying the Leibniz integral differentiation rule. \square

335 C Mean Field Computations

336 Making use of notation introduced in Section 3.1 we have that if $\mathbf{w} = (\hat{r}, \theta)$, with $\theta \in \mathcal{A}_k$, then

$$\begin{aligned} \nabla_{\theta} V(\mathbf{w}; \mu_t) &= -\nabla_{\theta} F(\mathbf{w}) + \int_{\mathcal{D}} \nabla_{\theta} K(\mathbf{w}, \mathbf{w}') \mu_t(d\mathbf{w}') \\ &= -\hat{r} \left(\sum_{j \in \mathcal{C}_k} y_j \langle \tilde{x}_j, t(\theta) \rangle - \int_{\mathcal{D}} r' \sum_{j \in \mathcal{C}_k} \langle \tilde{x}_j, t(\theta) \rangle \langle \tilde{x}_j, d(\theta') \rangle_+ \mu_t(d\hat{r}', d\theta') \right) \\ &= -\hat{r} \sum_{j \in \mathcal{C}_k} \langle \tilde{x}_j, t(\theta) \rangle \left(y_j - \int_{\mathbb{R} \times \mathcal{B}_j} \hat{r}' \langle \tilde{x}_j, d(\theta') \rangle \mu_t(d\hat{r}', d\theta') \right) \\ &= \hat{r} \left\langle \sum_{j \in \mathcal{C}_k} \rho_j(t) \tilde{x}_j, t(\theta) \right\rangle, \end{aligned} \quad (30)$$

337 where $\rho_j(t) = f_{\mu_t}(x_j) - y_j = \int_{\mathbb{R} \times \mathcal{B}_j} c(\tilde{x}_j, \theta) \mu_t(d\hat{r}, d\theta) - y_j$ is the residual at point x_j at time t .

338 Similarly, the field in the direction of the charges is given by

$$\nabla_{\hat{r}} V(\mathbf{w}; \mu_t) = \left\langle \sum_{j \in \mathcal{C}_k} \rho_j(t) \tilde{x}_j, \theta \right\rangle. \quad (31)$$

339 We also observe that for each j ,

$$\begin{aligned} \dot{\rho}_j(t) &= \partial_t f_{\mu_t}(x_j) = \partial_t \left(\int_{\mathcal{D}} \varphi(\mathbf{w}; x_j) \mu_t(d\mathbf{w}) \right) \\ &= - \int_{\mathcal{D}} \langle \nabla_{\mathbf{w}} \varphi(\mathbf{w}; x_j), \nabla V(\mathbf{w}; \mu_t) \rangle \mu_t(d\mathbf{w}) \\ &= - \int_{\mathcal{D}} (\nabla_{\theta} \varphi(\mathbf{w}; x_j) \cdot \nabla_{\theta} V(\mathbf{w}; \mu_t) + \nabla_r \varphi(\mathbf{w}; x_j) \cdot \nabla_r V(\mathbf{w}; \mu_t)) \mu_t(d\mathbf{w}) \\ &= - \sum_{k; \mathcal{A}_k \subset \mathcal{B}_j} \int_{\mathbb{R} \times \mathcal{A}_k} \left(r^2 \tilde{x}_j^{\top} (t(\theta) t(\theta)^{\top}) \left(\sum_{j' \in \mathcal{C}_k} \rho_{j'}(t) \tilde{x}_{j'} \right) + \tilde{x}_j^{\top} (\theta \theta^{\top}) \left(\sum_{j' \in \mathcal{C}_k} \rho_{j'}(t) \tilde{x}_{j'} \right) \right) \mu_t(d\mathbf{w}) \\ &= -\tilde{x}_j^{\top} \sum_{k; \mathcal{A}_k \subset \mathcal{B}_j} \Sigma_k(t) \left(\sum_{j' \in \mathcal{C}_k} \rho_{j'}(t) \tilde{x}_{j'} \right), \end{aligned} \quad (32)$$

where

$$\Sigma_k(t) = \int_{\mathbb{R} \times \mathcal{A}_k} (r^2 t(\theta) t(\theta)^\top + \theta \theta^\top) \mu_t(dr, d\theta)$$

340 tracks the covariance of the measure along each cylindrical region. Equation (32) defines a system of
 341 ODEs for the residuals $\rho(t)$, but its coefficients are time-varying, and behave roughly as quadratic
 342 terms in $\rho(t)$ (since they are second-order moments of the measure whereas the residuals are first-
 343 order moments). It may be possible to obtain asymptotic control of the oscillations $\rho(t)$ by applying
 344 Duhamel's principle.

345 D Changing Metric in the Dynamics

346 **Lemma 8.** *If $\mathbf{z}(t) = (\mathbf{a}(t), \mathbf{b}(t), \mathbf{c}(t))$ is a solution of the gradient flow (5), then the quantities*

$$\boldsymbol{\delta} = (\delta_i = c_i(t)^2 - a_i(t)^2 - b_i(t)^2)_{i=1}^m \quad (33)$$

347 *remain constant for all t . In particular, given a reduced neuron (r_i, θ_i) , we can uniquely recover the*
 348 *original neuron (a_i, b_i, c_i) , since*

$$c_i^2 = \frac{\delta_i + \sqrt{\delta_i^2 + 4r_i^2}}{2}. \quad (34)$$

349 *Proof.* The gradient equations of the loss $L(\mathbf{z})$ can be written as

$$\begin{aligned} \nabla_{a_i} L(\mathbf{z}) &= c_i \sum_{j=1}^s \mathbb{1}[a_i x_j + b_i \geq 0] x_j r_j, \\ \nabla_{b_i} L(\mathbf{z}) &= c_i \sum_{j=1}^s \mathbb{1}[a_i x_j + b_i \geq 0] r_j, \\ \nabla_{c_i} L(\mathbf{z}) &= \sum_{j=1}^s \mathbb{1}[a_i x_j + b_i \geq 0] (a_i x_j + b_i) r_j. \end{aligned} \quad (35)$$

350 From these expressions we see that

$$\begin{aligned} \dot{\delta}_i &= 2c_i \dot{c}_i - 2a_i \dot{a}_i - 2b_i \dot{b}_i \\ &= 2c_i \nabla_{c_i} L(\mathbf{z}) - 2a_i \nabla_{a_i} L(\mathbf{z}) - 2b_i \nabla_{b_i} L(\mathbf{z}) \\ &= 0. \end{aligned}$$

351 Using $r_i^2 = c_i \sqrt{a_i^2 + b_i^2}$, we see that $c_i^2 - \frac{r_i^2}{c_i} = \delta_i$ implies $c_i^4 - \delta_i c_i^2 - r_i^2 = 0$, and thus (34). \square

352 **Theorem 9.** *Let $\mathbf{z}(t)$ be a solution gradient flow (5) of $L(\mathbf{z})$, and let $\boldsymbol{\delta} = (\delta_i) \in \mathbb{R}^m$ be the vector*
 353 *of invariants (15), which depend only on the initialization $\mathbf{z}(0)$. If $\mathbf{w}(t) = (\mathbf{r}(t), \boldsymbol{\theta}(t))$ is curve of*
 354 *reduced parameters corresponding to $\mathbf{z}(t)$, then we have that*

$$\dot{\mathbf{w}}_i(t) = \mathbf{P}_i \cdot \nabla_{\mathbf{w}_i} \tilde{L}(\mathbf{w}), \quad i = 1, \dots, m,$$

355 where

$$\mathbf{P}_{\delta_i}(r_i) = \begin{bmatrix} a_i^2 + b_i^2 + c_i^2 & 0 \\ 0 & \frac{1}{a_i^2 + b_i^2} \end{bmatrix} = \begin{bmatrix} \frac{r_i^2}{c(r_i)^2} + c(r_i)^2 & 0 \\ 0 & \frac{c(r_i)^2}{r_i^2} \end{bmatrix},$$

356 and $c(r_i)^2 = \frac{\delta_i + \sqrt{\delta_i^2 + 4r_i^2}}{2}$.

357 *Proof.* The Jacobian of the mapping $\boldsymbol{\pi}$ from parameters to reduced parameters is given by

$$\text{Jac}(\boldsymbol{\pi})(a_i, b_i, c_i) = \begin{bmatrix} \frac{ca}{\sqrt{a_i^2 + b_i^2}} & \frac{cb}{\sqrt{a_i^2 + b_i^2}} & \sqrt{a_i^2 + b_i^2} \\ -\frac{b}{a_i^2 + b_i^2} & \frac{a}{a_i^2 + b_i^2} & 0 \end{bmatrix}, \quad i = 1, \dots, m.$$

358 This implies that the tangent kernel $\mathbf{P}_{\delta_i}(r_i) = \text{Jac}(\boldsymbol{\pi}) \text{Jac}(\boldsymbol{\pi})^T$ is as in (17). We emphasize that the
 359 fact that this kernel can be written only as a function of \mathbf{w} (and, in fact, only of \mathbf{r}) relies in essential
 360 manner on Lemma 2. \square

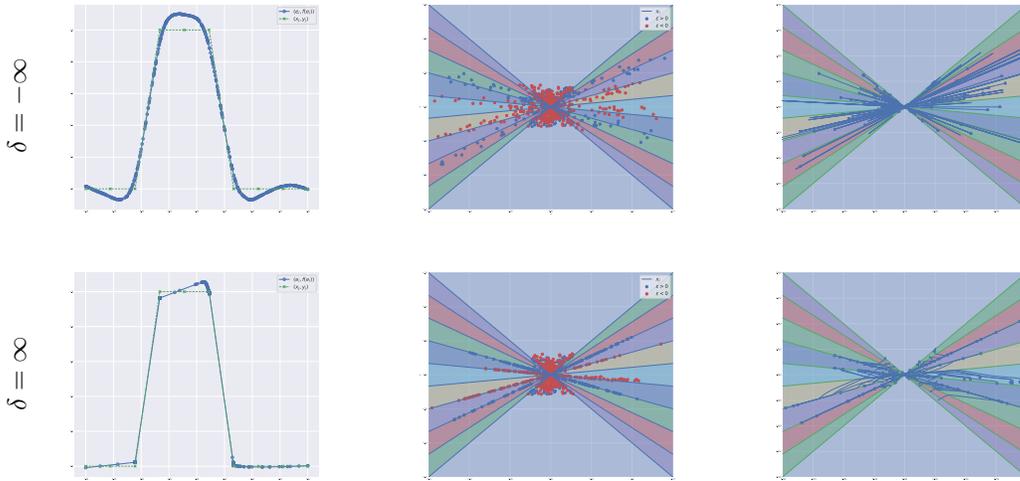


Figure 5: Evolution of 1000 neurons over 10000 epochs for $\delta = \pm\infty$ while fitting 10 points sampled from a square wave. Left: plotted network function after training. Middle: state of the network in uv space after training. Right: training trajectories of 100 random neurons.

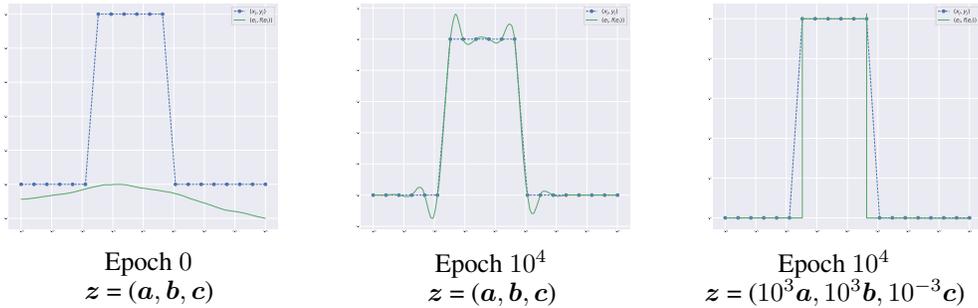


Figure 6: Left: A network (green) with an initial set of parameters initial parameters (a, b, c) is used to approximate a given function (blue). Scaling the initial parameters (right) leads to a very different fit (middle).

361 E Additional Numerical Experiments

362 In Figure 5, we plot the trajectories of neurons for $\delta = \pm\infty$ over 10000 epochs. We see that, if
 363 $\delta = -\infty$, the neurons move radially away from the origin and thus the knot positions do not change
 364 (top row). In stark contrast, if $\delta = \infty$, the neurons adapt to the input data, and the knots “stick” to
 365 input samples (bottom row).

366 We remark in Figure 6 that the same initial function can yield extremely different results depending
 367 on δ .

368 We now show the effect of varying the number of neurons during training. In this example we
 369 fit 20 samples from a sine wave using 20, 200, and 2000 neurons respectively. In PyTorch, the
 370 default initialization is such that $a, b \sim U(-1, 1)$ and $c \sim U(-1/m, 1/m)$. Thus, as we scale down
 371 the number of neurons, the value of δ grows, making the network function adapt more to the data.
 372 Figure 7 shows the results of this experiment.

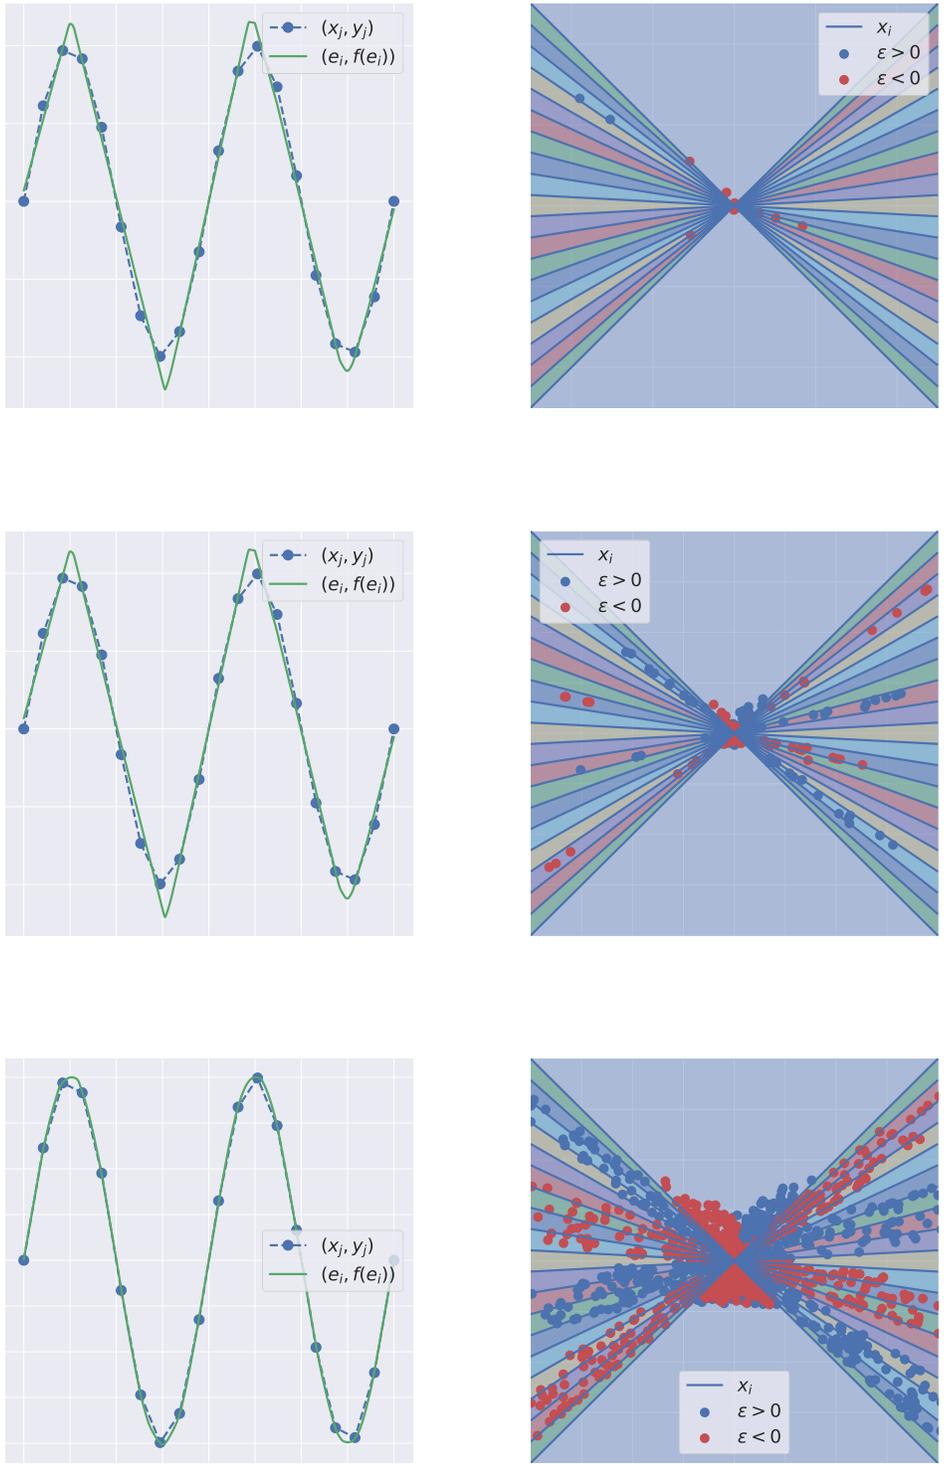


Figure 7: The effect of varying the number of neurons m the top image uses 20 neurons, the middle uses 200 and the bottom uses 2000. Observe that with fewer neurons, the function is adaptive to the data since δ gets larger.

373 E.1 Visualizing Attractor Samples

We can visualize the vector field $(\partial_t \hat{r}, \theta)$ by considering the change of metric from $w = (\hat{r}, \theta)$ to (u, v) with the map

$$\pi_{(u,v)}(\hat{r}, \theta) = (|\hat{r}| \cos \theta, |\hat{r}| \sin \theta) = (u, v).$$

374 Assuming we know the sign of \hat{r} , the vector field

$$\begin{bmatrix} \partial_t u \\ \partial_t v \end{bmatrix} = D\pi_{(u,v)} D\pi_{(u,v)}^T \begin{bmatrix} \partial_t r \\ \partial_t \theta \end{bmatrix} \quad (36)$$

Observing that $D\pi_{(u,v)} D\pi_{(u,v)}^T = I$, we have simply that

$$\begin{bmatrix} \partial_t u \\ \partial_t v \end{bmatrix} = \begin{bmatrix} \partial_t r \\ \partial_t \theta \end{bmatrix}$$

375 Figure 8 shows a plot of this vector field by a single particle in uv in the case where $\delta = \infty$. In this
 376 case, the partial derivative $\partial_t r$ remains unchanged. Furthermore, we remark that at the boundaries
 377 of samples, the vector field can change directions, causing these samples to become ‘‘attractors’’ or
 378 ‘‘repulsors’’ (see Lemma 11 in the main document).

379 F Implicit Regularization of the Kernel Regime

380 We can write the network function (11) in the standard parameterization as a matrix product:

$$f_z(x) = M\mathbf{c}, \quad M \in \mathbb{R}^{s \times m}, \quad M_{ij} = [a_j x_i + b_j]_+ \quad (37)$$

381 Now consider the least-squares problem:

$$\text{minimize } \frac{1}{2} \|M\mathbf{c} - \mathbf{y}\|^2 \quad (38)$$

382 If we are in the kernel regime (see Section 3.3) then we minimize (38) by only changing the
 383 parameter \mathbf{c} . Following gradient flow, we have that \mathbf{c} follows the ODE

$$\partial_t \mathbf{c}(t) = -\nabla \frac{1}{2} \|M\mathbf{c}(t) - \mathbf{y}\|^2 \quad (39)$$

$$= -(M^T M \mathbf{c} - M^T \mathbf{y}). \quad (40)$$

In the finite width case, the kernel K can be written as

$$K(x_i, x_j) = K = (MM^T)_{ij}$$

384 And, thus, $f(t) = M\mathbf{c}(t)$ follows the separable ODE:

$$\partial_t f(t) = M \partial_t \mathbf{c}(t) \quad (41)$$

$$= MM^T \mathbf{y} - MM^T M \mathbf{c} \quad (42)$$

$$= K\mathbf{y} - KM\mathbf{c} \quad (43)$$

$$= K\mathbf{y} - Kf(t) \quad (44)$$

385 whose solutions are of the form:

$$f(t) \propto \exp(-K)t + \mathbf{y} \quad (45)$$

386 We can decompose K into its eigenbasis where \mathbf{e}_i , an eigenvector of K and λ_i is the corresponding
 387 eigenvalue for $i = 1, \dots, s$.

388 We can then write (45) as:

$$f(t) = \mathbf{y} + \sum_{i=1}^n \exp(-t\lambda_i) \mathbf{e}_i \quad (46)$$

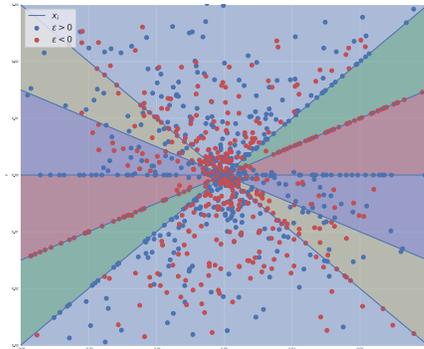
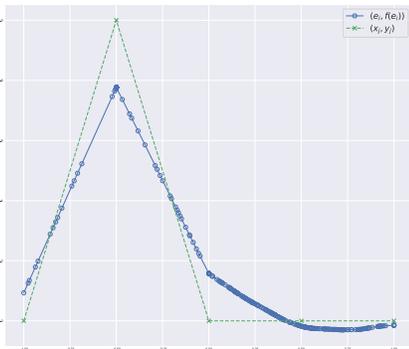
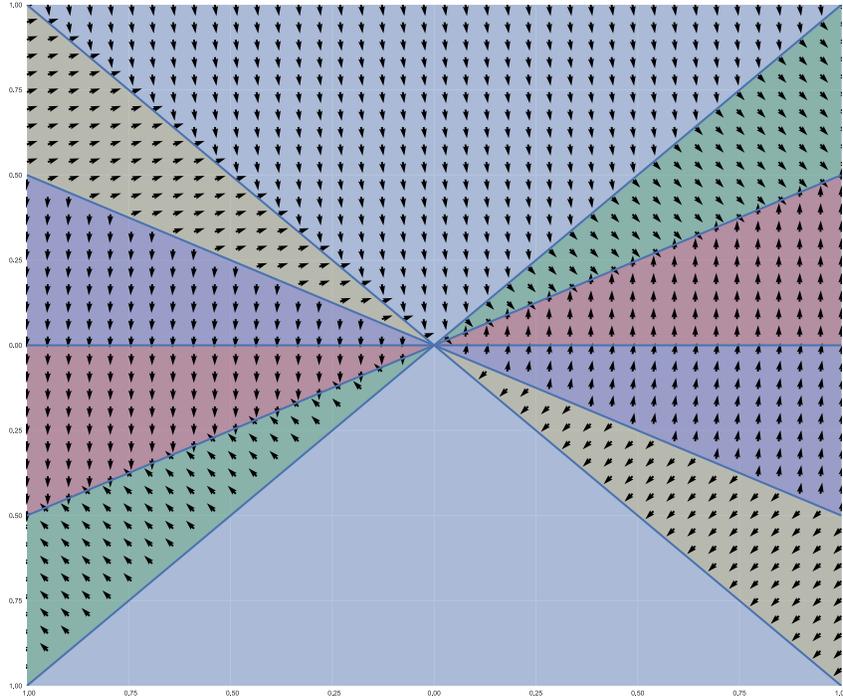


Figure 8: *Top*: The gradient field (18) felt by a particle. Note how the vectors change directions at certain samples. These samples are “attractors” or “repulsors” where particles get stuck or get pushed away from. *Bottom Left*: A plot of the network function for the gradient field in the top image. Observe how there are clusters of neurons (blue circles) aligned with certain samples. *Bottom Right*: A plot of the neurons in uv space. Observe how the red neurons cluster at “attractor” points in the top image

389 which implies the dynamics of the residual, $f(t) - y$ are:

$$f(t) - y = \sum_{i=1}^n \exp(-t\lambda_i \mathbf{e}_i) \quad (47)$$

390 Thus, as $t \rightarrow \infty$, the terms $\exp(-t\lambda_i \mathbf{e}_i)$ decay, at a rate exponential in λ_i and $f(t) \rightarrow y$, which, as
391 mentioned in [16], suggests that early stopping acts as a regularizer by decaying the residual primarily
392 along the principle components of the kernel. These principle components usually correspond to
393 smoother functions. Furthermore, (47) means that if the kernel is not full rank, then the residual
394 cannot go to zero since some λ_i will be zero.