

Responsable: Dario Prandi, dario.prandi@centralesupelec.fr

Chargé de TD: Kevin Le Balc'h, kevin.le-balc-h@inria.fr

Exercise 1: Constrained optimization and Fenchel duality

We consider the optimization problem (P)

$$\min_{(x,y) \in \mathbb{R}^2} x^2 + 2y^2 - 2x - 4y,$$

subject to the constraints

$$x + y \leq 1, \quad x - y \leq 0.$$

1. Show that (P) admits a unique solution.
2. Write the Lagrangian associated with (P) and state the KKT conditions.
3. Verify that Slater's condition holds. Solve (P) using KKT's theorem.
4. Write the dual problem and verify strong duality.

Solution

1. The objective $f(x, y) = x^2 + 2y^2 - 2x - 4y$ is strictly convex (its Hessian is $\text{diag}(2, 4) \succ 0$), hence continuous and coercive. The constraint set is a closed convex polyhedron and is non-empty (e.g., $(0, 0)$ is feasible). Therefore, (P) admits a unique solution.

2. The Lagrangian is

$$L(x, y, \lambda_1, \lambda_2) = x^2 + 2y^2 - 2x - 4y + \lambda_1(x + y - 1) + \lambda_2(x - y), \quad \lambda_1, \lambda_2 \geq 0.$$

The KKT conditions are:

- Primal feasibility: $x + y \leq 1, x - y \leq 0$.
- Dual feasibility: $\lambda_1, \lambda_2 \geq 0$.
- Complementary slackness: $\lambda_1(x + y - 1) = 0, \lambda_2(x - y) = 0$.
- Stationarity:

$$\begin{cases} 2x - 2 + \lambda_1 + \lambda_2 = 0, \\ 4y - 4 + \lambda_1 - \lambda_2 = 0. \end{cases}$$

3. Slater's condition is satisfied: the point $(0, 1/2)$ satisfies $0 + 1/2 < 1$ and $0 - 1/2 < 0$, so both constraints are strictly satisfied.

We solve the KKT system by cases.

Case 1: $\lambda_1 = 0, \lambda_2 = 0$. Stationarity gives $x = 1, y = 1$. Primal feasibility: $1 + 1 = 2 > 1$. **Infeasible.**

Case 2: $\lambda_1 > 0, \lambda_2 = 0$ (first constraint active, second inactive). From $x + y = 1$ we get $x = 1 - y$. Stationarity gives:

$$2(1 - y) - 2 + \lambda_1 = 0 \implies \lambda_1 = 2y, \quad 4y - 4 + \lambda_1 = 0 \implies 4y - 4 + 2y = 0 \implies y = \frac{2}{3}.$$

Hence $x = \frac{1}{3}, \lambda_1 = \frac{4}{3}$. Primal feasibility: $x - y = \frac{1}{3} - \frac{2}{3} = -\frac{1}{3} \leq 0 \checkmark$. Dual feasibility: $\lambda_1 = \frac{4}{3} > 0 \checkmark$.

Case 3: $\lambda_1 = 0, \lambda_2 > 0$ (second constraint active, first inactive). From $x = y$. Stationarity: $2x - 2 + \lambda_2 = 0$ and $4x - 4 - \lambda_2 = 0$. Adding: $6x - 6 = 0 \implies x = 1$, so $y = 1$. Primal feasibility: $x + y = 2 > 1$. **Infeasible.**

Case 4: $\lambda_1 > 0$, $\lambda_2 > 0$ (both constraints active). From $x + y = 1$ and $x = y$: $x = y = \frac{1}{2}$. Stationarity: $\lambda_1 + \lambda_2 = 1$ and $\lambda_1 - \lambda_2 = 2$, giving $\lambda_1 = \frac{3}{2}$, $\lambda_2 = -\frac{1}{2} < 0$. **Dual infeasible.**

The only KKT point is $(x^*, y^*) = (\frac{1}{3}, \frac{2}{3})$ with multipliers $\lambda_1^* = \frac{4}{3}$, $\lambda_2^* = 0$, and optimal value

$$f\left(\frac{1}{3}, \frac{2}{3}\right) = \frac{1}{9} + \frac{8}{9} - \frac{2}{3} - \frac{8}{3} = 1 - \frac{10}{3} = -\frac{7}{3}.$$

4. The dual function is

$$g(\lambda_1, \lambda_2) = \inf_{(x,y) \in \mathbb{R}^2} L(x, y, \lambda_1, \lambda_2).$$

Since L is strictly convex in (x, y) , the infimum is attained at

$$x = \frac{2 - \lambda_1 - \lambda_2}{2}, \quad y = \frac{4 - \lambda_1 + \lambda_2}{4}.$$

Substituting back:

$$g(\lambda_1, \lambda_2) = -\frac{(2 - \lambda_1 - \lambda_2)^2}{4} - \frac{(4 - \lambda_1 + \lambda_2)^2}{8} - \lambda_1 = -3 + \lambda_1 - \frac{3\lambda_1^2 + 2\lambda_1\lambda_2 + 3\lambda_2^2}{8}.$$

The dual problem is $\max_{\lambda_1, \lambda_2 \geq 0} g(\lambda_1, \lambda_2)$. On the face $\lambda_2 = 0$: $g(\lambda_1, 0) = -3 + \lambda_1 - \frac{3}{8}\lambda_1^2$, whose maximum is at $\lambda_1 = \frac{4}{3}$ (from $1 - \frac{3}{4}\lambda_1 = 0$), giving

$$g\left(\frac{4}{3}, 0\right) = -3 + \frac{4}{3} - \frac{3}{8} \cdot \frac{16}{9} = -3 + \frac{4}{3} - \frac{2}{3} = -\frac{7}{3}.$$

One checks this is the global maximum (the unconstrained critical point has $\lambda_2 < 0$, and g is decreasing in λ_2 for $\lambda_2 \geq 0$).

Strong duality holds since Slater's condition is satisfied: $f(x^*, y^*) = -\frac{7}{3} = g(\lambda_1^*, \lambda_2^*)$.

Exercise 2: Projected gradient descent and convergence

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function of class C^1 with L -Lipschitz gradient, and let $C \subset \mathbb{R}^n$ be a non-empty, closed, convex set. We consider the *projected gradient descent* algorithm defined by

$$x^{(k+1)} = \Pi_C \left(x^{(k)} - \alpha \nabla f(x^{(k)}) \right),$$

where Π_C denotes the orthogonal projection onto C and $\alpha > 0$ is the step size.

1. Recall why Π_C is well-defined on \mathbb{R}^n . Show that, for any $x \in \mathbb{R}^n$ and $y \in C$,

$$\langle x - \Pi_C(x), y - \Pi_C(x) \rangle \leq 0.$$

Hint: for the second part, letting $p = \Pi_C(x)$, justify and use the fact that $\|x - p\|^2 \leq \|x - z\|^2$ for $z = p + t(y - p)$, $t \geq 0$.

2. Show that Π_C is 1-Lipschitz, i.e., for all $x, z \in \mathbb{R}^n$,

$$\|\Pi_C(x) - \Pi_C(z)\| \leq \|x - z\|.$$

3. Show that $x^* \in C$ is a minimizer of f on C if and only if, for any $\alpha > 0$,

$$x^* = \Pi_C(x^* - \alpha \nabla f(x^*)).$$

Hint: reformulate the problem as an unconstrained one, using χ_C , the $0 - \infty$ characteristic function of C . Apply the minimization criterion for the subdifferential and use question 1.

4. Show that, for any x^* minimizer of f over C and for any $k \geq 0$,

$$\|x^{(k+1)} - x^*\|^2 \leq \|x^{(k)} - x^*\|^2 - \alpha \left(\frac{2}{L} - \alpha \right) \|\nabla f(x^{(k)}) - \nabla f(x^*)\|^2 - 2\alpha \left(f(x^{(k)}) - f(x^*) \right).$$

Hint: use the characterization of the projection from question 1, and the quadratic upper bound

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|y - x\|^2, \quad \forall x, y \in \mathbb{R}^n$$

5. Show that for $0 < \alpha < 2/L$, the sequence $\{f(x^{(k)})\}$ is non-increasing.
6. Deduce that for $0 < \alpha < 2/L$, the sequence $\{f(x^{(k)})\}$ converges to $f(x^*)$ and

$$f(x^{(k)}) - f(x^*) \leq \frac{\|x^{(0)} - x^*\|^2}{2\alpha k}.$$

7. Now consider the specific problem of minimizing $f(x) = \frac{1}{2}\|Ax - b\|^2$ over $C = \{x \in \mathbb{R}^n : x \geq 0\}$ (non-negative least squares), where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.
- (a) Write ∇f and determine the Lipschitz constant L of the gradient.
- (b) Write explicitly the projected gradient iteration for this problem.

Solution

1. C is a non-empty, closed, convex set. The function $y \mapsto \|y - x\|^2$ is strongly convex and coercive, hence it admits a unique minimizer on C , which defines $\Pi_C(x)$.

Let $p = \Pi_C(x)$. For any $y \in C$ and $t \in [0, 1]$, $p + t(y - p) \in C$ by convexity. Since p minimizes $\|z - x\|^2$ over C ,

$$\|p - x\|^2 \leq \|p + t(y - p) - x\|^2 = \|p - x\|^2 + 2t\langle p - x, y - p \rangle + t^2\|y - p\|^2.$$

Dividing by $t > 0$ and letting $t \rightarrow 0$: $\langle p - x, y - p \rangle \geq 0$, i.e., $\langle x - p, y - p \rangle \leq 0$.

2. Let $p = \Pi_C(x)$ and $q = \Pi_C(z)$. From question 1:

$$\langle x - p, q - p \rangle \leq 0 \quad \text{and} \quad \langle z - q, p - q \rangle \leq 0.$$

Adding: $\langle x - p - (z - q), q - p \rangle \leq 0$, i.e.,

$$\langle (x - z) - (p - q), q - p \rangle \leq 0 \implies \|p - q\|^2 \leq \langle x - z, p - q \rangle \leq \|x - z\| \cdot \|p - q\|.$$

Dividing by $\|p - q\|$ (if $p \neq q$) gives $\|p - q\| \leq \|x - z\|$.

3. Minimizing f over C is equivalent to minimizing $F := f + \chi_C$ over \mathbb{R}^n . By the theorem on subdifferential characterization of minimizers, x^* minimizes F iff $0 \in \partial F(x^*) = \nabla f(x^*) + \partial \chi_C(x^*)$.

By definition of subdifferential applied to χ_C : since $\chi_C(x^*) = 0$ and $\chi_C(y) = 0$ for all $y \in C$, we have $v \in \partial \chi_C(x^*)$ iff $\langle v, y - x^* \rangle \leq 0$ for all $y \in C$.

Hence x^* is a minimizer iff $-\nabla f(x^*) \in \partial \chi_C(x^*)$, i.e.,

$$\langle \nabla f(x^*), y - x^* \rangle \geq 0, \quad \forall y \in C.$$

Setting $z = x^* - \alpha \nabla f(x^*)$, this is equivalent to $\langle z - x^*, y - x^* \rangle \leq 0$ for all $y \in C$, which by question 1 means exactly $x^* = \Pi_C(z) = \Pi_C(x^* - \alpha \nabla f(x^*))$.

4. Let $g_k = \nabla f(x^{(k)})$ and $g^* = \nabla f(x^*)$. Since x^* is a minimizer on C , by question 3,

$$x^* = \Pi_C(x^* - \alpha g^*).$$

By non-expansiveness of the projection (question 2):

$$\|x^{(k+1)} - x^*\|^2 \leq \|(x^{(k)} - \alpha g_k) - (x^* - \alpha g^*)\|^2 = \|x^{(k)} - x^*\|^2 - 2\alpha \langle g_k - g^*, x^{(k)} - x^* \rangle + \alpha^2 \|g_k - g^*\|^2.$$

For convex L -smooth f , we use

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{1}{2L} \|\nabla f(y) - \nabla f(x)\|^2,$$

with $(x, y) = (x^{(k)}, x^*)$, which gives

$$\langle g_k, x^{(k)} - x^* \rangle \geq f(x^{(k)}) - f(x^*) + \frac{1}{2L} \|g_k - g^*\|^2.$$

Also, optimality of x^* on C implies $\langle g^*, x^{(k)} - x^* \rangle \geq 0$, hence

$$\langle g_k - g^*, x^{(k)} - x^* \rangle \geq f(x^{(k)}) - f(x^*) + \frac{1}{2L} \|g_k - g^*\|^2.$$

Substituting into the previous estimate:

$$\|x^{(k+1)} - x^*\|^2 \leq \|x^{(k)} - x^*\|^2 - 2\alpha(f(x^{(k)}) - f(x^*)) - \alpha\left(\frac{2}{L} - \alpha\right) \|g_k - g^*\|^2.$$

5. For $0 < \alpha < 2/L$, by smoothness,

$$f(x^{(k+1)}) \leq f(x^{(k)}) + \langle g_k, x^{(k+1)} - x^{(k)} \rangle + \frac{L}{2} \|x^{(k+1)} - x^{(k)}\|^2,$$

and by projection optimality with $y = x^{(k)}$,

$$\langle g_k, x^{(k+1)} - x^{(k)} \rangle \leq -\frac{1}{\alpha} \|x^{(k+1)} - x^{(k)}\|^2.$$

Hence

$$f(x^{(k+1)}) \leq f(x^{(k)}) - \left(\frac{1}{\alpha} - \frac{L}{2}\right) \|x^{(k+1)} - x^{(k)}\|^2 \leq f(x^{(k)}).$$

6. For $0 < \alpha < 2/L$, the term $\alpha\left(\frac{2}{L} - \alpha\right) > 0$. Dropping the non-negative gradient term in question 4:

$$2\alpha(f(x^{(k)}) - f(x^*)) \leq \|x^{(k)} - x^*\|^2 - \|x^{(k+1)} - x^*\|^2.$$

Summing from $k = 0$ to $K - 1$ (telescoping):

$$2\alpha \sum_{k=0}^{K-1} (f(x^{(k)}) - f(x^*)) \leq \|x^{(0)} - x^*\|^2.$$

Using question 5, $f(x^{(k)})$ is non-increasing, hence

$$f(x^{(K)}) - f(x^*) \leq \frac{1}{K} \sum_{k=0}^{K-1} (f(x^{(k)}) - f(x^*)) \leq \frac{\|x^{(0)} - x^*\|^2}{2\alpha K}.$$

7. (a) $\nabla f(x) = A^\top(Ax - b)$, which is Lipschitz with constant $L = \|A^\top A\| = \sigma_{\max}(A)^2$, the square of the largest singular value of A .

(b) The projected gradient iteration is

$$x^{(k+1)} = \max(x^{(k)} - \alpha A^\top(Ax^{(k)} - b), 0),$$

where $\max(\cdot, 0)$ is applied component-wise.

Exercise 3: Controllability and LQ optimal control

Consider the linear control system in \mathbb{R}^n :

$$\dot{x} = Ax + Bu, \quad x(0) = x_0,$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $u \in L^\infty([0, T], \mathbb{R}^m)$. Consider the LQ optimal control problem on $[0, T]$: minimize

$$J(u) = \frac{1}{2} \int_0^T (x(t)^\top Q x(t) + u(t)^\top R u(t)) dt + \frac{1}{2} x(T)^\top S x(T), \quad (\text{LQ-cost})$$

where $Q, S \in \mathbb{R}^{n \times n}$ are symmetric positive semi-definite and $R \in \mathbb{R}^{m \times m}$ is symmetric positive definite.

1. Consider the following system:

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

- Determine whether the system is controllable, and describe the reachable set from a given initial state x_0 .
- State the matrix Riccati equation satisfied by $E(t)$ and give the expression for the optimal control $u^*(t)$ in feedback form.
- Does the optimal control exist and is it unique, even if the system (A, B) is not controllable? Justify.

2. We now consider the *double integrator* ($n = 2, m = 1$):

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad Q = I_2, \quad R = 1.$$

- Verify that (A, B) is controllable.
- Let $E_T(t)$ denote the solution of the finite-horizon Riccati equation for the cost functional J of Equation (LQ-cost) with $S = 0$. Assume that, for fixed t , $E_T(t)$ converges to a limit matrix $E_\infty \succeq 0$ as $T \rightarrow \infty$. Show that E_∞ satisfies the algebraic Riccati equation:

$$Q + A^\top E_\infty + E_\infty A - E_\infty B R^{-1} B^\top E_\infty = 0.$$

- Writing $E_\infty = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, expand the algebraic Riccati equation into a system of scalar equations for a, b, c, d and solve it explicitly.
- Give the optimal feedback control $u^* = K^* x$ corresponding to E_∞ and write the closed-loop matrix $A_{cl} = A - BK^*$. Verify that A_{cl} is Hurwitz (i.e., is such that the closed-loop system $\dot{x} = A_{cl}x$ is asymptotically stable).

Solution

- (a) The Kalman rank condition states that (A, B) is controllable if and only if

$$\text{rank } \mathcal{K}, \quad \mathcal{K} = [B, AB, A^2B, \dots, A^{n-1}B] = n.$$

Here $n = 3$ and we compute:

$$B = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad AB = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \quad A^2B = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix},$$

so $\text{rank } \mathcal{K} = 1 < 3$. The system is **not controllable**.

The dynamics read $\dot{x}_1 = x_2 + u$, $\dot{x}_2 = x_3$, $\dot{x}_3 = 0$. Thus $x_3(t) = x_3(0)$ is entirely unaffected by u : the reachable set from x_0 is $\{y \in \mathbb{R}^3 : y_3 = x_{0,3}\}$, a plane parallel to the (x_1, x_2) -plane.

- The symmetric positive semi-definite matrix $P(t)$ satisfies the backward Riccati ODE

$$-\dot{E}(t) = Q + A^\top E(t) + E(t)A - E(t)BR^{-1}B^\top E(t), \quad E(T) = S.$$

The optimal feedback control is

$$u^*(t) = -R^{-1}B^\top E(t)x(t).$$

- Yes: since $R \succ 0$ the integrand is strictly convex and coercive in u , so J has a unique minimizer for every x_0 . Controllability governs whether arbitrary terminal states can be reached; here we are minimizing a cost, not targeting a specific terminal state.

2. (a) The Kalman matrix is

$$\mathcal{K} = [B, AB] = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

which has rank 2. Hence (A, B) is controllable.

(b) Introduce the time-to-go variable $\tau = T - t$ and define $p(\tau) := P_T(T - \tau)$. Then p solves the autonomous ODE

$$\frac{dp}{d\tau} = Q + A^\top p + pA - pBR^{-1}B^\top p, \quad p(0) = 0.$$

By assumption, $p(\tau) \rightarrow P_\infty$ as $\tau \rightarrow +\infty$. For an autonomous ODE, any limit point must be an equilibrium, hence P_∞ satisfies the algebraic Riccati equation.

(c) Observe that P_∞ is symmetric (it is a limit of symmetric matrices), so $c = b$. Write directly

$$P_\infty = \begin{bmatrix} a & b \\ b & d \end{bmatrix}.$$

Then

$$A^\top P_\infty = \begin{bmatrix} 0 & 0 \\ a & b \end{bmatrix}, \quad P_\infty A = \begin{bmatrix} 0 & a \\ 0 & b \end{bmatrix}, \quad P_\infty BR^{-1}B^\top P_\infty = \begin{bmatrix} b^2 & bd \\ bd & d^2 \end{bmatrix}.$$

Hence the ARE gives

$$(1, 1) : 1 - b^2 = 0, \quad (1, 2) : a - bd = 0, \quad (2, 2) : 1 + 2b - d^2 = 0.$$

From (1, 1), $b = \pm 1$; for the stabilizing positive-semidefinite solution, take $b = 1$. Then (1, 2) gives $a = d$, and (2, 2) gives $d^2 = 3$, so $d = \sqrt{3}$ (positive branch), hence $a = \sqrt{3}$. Therefore

$$P_\infty = \begin{bmatrix} \sqrt{3} & 1 \\ 1 & \sqrt{3} \end{bmatrix}.$$

This is indeed positive definite: $\det P_\infty = 3 - 1 = 2 > 0$ and $\text{tr } P_\infty = 2\sqrt{3} > 0$.

(d) The optimal feedback law has the form

$$u^*(x) = -R^{-1}B^\top P_\infty x = -[1 \quad \sqrt{3}] x.$$

If we denote the (row) gain by

$$K^* := R^{-1}B^\top P_\infty = [1 \quad \sqrt{3}],$$

then $u^*(x) = -K^*x$ and, with the statement convention $A_{\text{cl}} = A - BK^*$,

$$A_{\text{cl}} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \end{bmatrix} [1 \quad \sqrt{3}] = \begin{bmatrix} 0 & 1 \\ -1 & -\sqrt{3} \end{bmatrix}.$$

Its characteristic polynomial is

$$\det(\lambda I - A_{\text{cl}}) = \lambda^2 + \sqrt{3}\lambda + 1,$$

with roots $\lambda = \frac{-\sqrt{3} \pm i}{2}$. Since $\text{Re}(\lambda) = -\frac{\sqrt{3}}{2} < 0$, the matrix A_{cl} is Hurwitz and the closed-loop system is asymptotically stable.

Exercise 4: Neural ODEs and ResNets

Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a nonlinear activation function applied component-wise. Let $\theta : [0, 1] \rightarrow \mathbb{R}^p$ be a time-varying parameter and consider the *neural ODE*

$$\dot{x} = f_\theta(t, x) = W(t) \sigma(V(t)x + c(t)), \quad x(0) = x_0 \in \mathbb{R}^n, \quad (1)$$

where $W(t) \in \mathbb{R}^{n \times q}$, $V(t) \in \mathbb{R}^{q \times n}$, $c(t) \in \mathbb{R}^q$, and $\theta(t) = (W(t), V(t), c(t))$ are learnable parameters.

- Write down the ResNet with L layers corresponding to the Euler discretization of (1) on $[0, 1]$ with step $h = 1/L$.
- Let $\Phi_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denote the flow map of (1), i.e., $x(t) = \Phi_t(x_0)$. Show that

$$\det(D\Phi_t(x_0)) > 0, \quad \forall t \in [0, 1], \forall x_0 \in \mathbb{R}^n.$$

What does this imply about the expressiveness of neural ODEs?

Hint: use the Liouville–Jacobi formula $\frac{d}{dt} \det(D\Phi_t) = \det(D\Phi_t) \operatorname{tr}(D_x f_\theta(t, \Phi_t))$.

- To overcome this limitation, one considers the *augmented neural ODE*:

$$\frac{d}{dt} \begin{pmatrix} x \\ z \end{pmatrix} = g_\theta(t, x, z), \quad \begin{pmatrix} x(0) \\ z(0) \end{pmatrix} = \begin{pmatrix} x_0 \\ 0 \end{pmatrix},$$

Here, $(x, z) \in \mathbb{R}^{n+n'}$ is the augmented state. The function g_θ has the form (1) but with x replaced by (x, z) and learnable parameters of appropriate dimensions. The output map is then $\Psi : x_0 \mapsto x(1)$ (projecting back onto the first n coordinates).

- Explain why Ψ is no longer constrained to have positive Jacobian determinant.
- Provide an explicit example with $n = n' = 1$ and linear activation function σ , where Ψ realizes the map $x_0 \mapsto -x_0$.

Hint: consider a simple linear system that rotates the state in the (x, z) -plane.

Solution

- The Euler discretization of (1) with step $h = 1/L$ gives the ResNet

$$x_{k+1} = x_k + \frac{1}{L} W_k \sigma(V_k x_k + c_k), \quad k = 0, \dots, L-1,$$

where $W_k = W(k/L)$, $V_k = V(k/L)$, and $c_k = c(k/L)$ are the layer weights. This is a residual network: each layer adds a “correction” $\frac{1}{L} W_k \sigma(\cdot)$ to the identity.

- By the Liouville–Jacobi formula,

$$\frac{d}{dt} \det(D\Phi_t(x_0)) = \det(D\Phi_t(x_0)) \operatorname{tr}(D_x f_\theta(t, \Phi_t(x_0))).$$

Since $\det(D\Phi_0) = \det(\operatorname{Id}) = 1 > 0$, integrating gives

$$\det(D\Phi_t(x_0)) = \exp\left(\int_0^t \operatorname{tr}(D_x f_\theta(s, \Phi_s(x_0))) ds\right) > 0$$

for all t and x_0 , since the exponential is always positive.

Consequence: the flow map Φ_1 always *preserves orientation*. Therefore, a neural ODE cannot learn any map whose Jacobian has negative determinant anywhere — in particular, it cannot realize reflections, orientation-reversing permutations, or any map of degree -1 . This is a fundamental expressiveness limitation of the architecture.

- (a) The augmented flow $\tilde{\Phi}_t : (x_0, 0) \mapsto (x(t), z(t))$ in $\mathbb{R}^{n+n'}$ still satisfies $\det(D\tilde{\Phi}_t) > 0$ by the same argument. However, the map of interest,

$$\Psi = \pi \circ \tilde{\Phi}_1 \circ \iota,$$

where $\iota : x_0 \mapsto (x_0, 0)$ (embedding) and $\pi : (x, z) \mapsto x$ (projection), goes between spaces of the *same* dimension n but passes through a higher-dimensional space. Since ι and π are not square maps, the chain rule gives $D\Psi = D\pi \cdot D\tilde{\Phi}_1 \cdot D\iota$, a product of non-square matrices, whose determinant is unconstrained in sign. In particular, the positivity of $\det(D\tilde{\Phi}_1)$ does not force $\det(D\Psi) > 0$.

(b) Consider the linear system

$$\frac{d}{dt} \begin{pmatrix} x \\ z \end{pmatrix} = \begin{bmatrix} 0 & -\omega \\ \omega & 0 \end{bmatrix} \begin{pmatrix} x \\ z \end{pmatrix}, \quad \begin{pmatrix} x(0) \\ z(0) \end{pmatrix} = \begin{pmatrix} x_0 \\ 0 \end{pmatrix},$$

where $\omega > 0$ is a constant. This can be realised as an augmented neural ODE with linear activation $\sigma(s) = s$ and parameters

$$W = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad V = \begin{bmatrix} 0 & -\omega \\ \omega & 0 \end{bmatrix}, \quad c = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

The flow of this system is a rotation in the (x, z) -plane:

$$\tilde{\Phi}_t(x_0, 0) = \begin{bmatrix} \cos(\omega t) & -\sin(\omega t) \\ \sin(\omega t) & \cos(\omega t) \end{bmatrix} \begin{pmatrix} x_0 \\ 0 \end{pmatrix} = \begin{pmatrix} x_0 \cos(\omega t) \\ x_0 \sin(\omega t) \end{pmatrix}.$$

Choosing $\omega = \pi$ gives $\Psi(x_0) = x(1) = -x_0$, realizing the desired map.