Control-based classification with neural ordinary differential equations

X Partial differential equations, optimal design and numerics

Antonio Álvarez-López

Joint work with Rafael Orive Illera and Enrique Zuazua

Department of Mathematics,
Universidad Autónoma de Madrid

August 27, 2024

Supervised Learning

Goal

Input space $(\mathcal{X}, \mu^*) \subset \mathbb{R}^d$ $\xrightarrow{\quad F^* \quad}$ Output space $\mathcal{Y} \subset \mathbb{R}^m$

Approximate (*learn*) F^* from a dataset $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N \subset \mathcal{X} \times \mathcal{Y}$:

$$\mathbf{x}_n \sim \mu^*, \qquad \mathbf{y}_n = F^*(\mathbf{x}_n), \qquad n = 1, \ldots N.$$

$x_1 \in \mathbb{R}^{3N_{pixels}}$

$x_2 \in \mathbb{R}^{3N_{pixels}}$

F^*

$y_1 = (1, 0) \equiv$ Cat

$y_2 = (0, 1) \equiv$ Dog

Classification

$\downarrow$

range(F^*) is finite

Motivation

E. Weinan, "A proposal on machine learning via dynamical systems" (2017).

In this note, we go one step further, by exploring the possibility of producing nonlinear functions using continuous dynamical systems, pushing the compositional approach to an infinitesimal limit. In the framework of supervised learning, this gives rise to a new class of control problems. In this view, the deep neural networks can be thought of as being discrete dynamical systems. Compared with deep neural networks, there are several potential advantages with a continuous approach.

1. Mathematically it is easier to think about and deal with continuous dynamical systems. Continuous formulation offers more flexibility (for example adding constraints, adapting the dynamical system to the problem, imposing structure on the dynamical system), and it is easier to analyze continuous dynamical systems than discrete ones.
2. Deep neural network can be thought of as a discretization of the continuous dynamical systems. However, from the viewpoint of discretizing dynamical systems, there are many possibilities that one can explore. For example, one can use adaptive time step size, which corresponds to choosing the layers adaptively. One can use high order or even implicit discretization, and these do not yet have an analog in deep neural networks. One can also use advanced numerical methods for training, such as the multi-grid method or the parallel shooting method (see [4,5]).
3. Most models in physical sciences (physics, chemistry, etc) are represented using dynamical systems in the form of differential equations. The continuous dynamical systems approach makes it easier to combining ideas from machine learning and physical modeling.
4. The vast majority of the applied mathematics community is familiar with differential equations. The continuous dynamical systems approach to machine learning will be of particular interest to them.

Neural ODEs

$$\dot{\mathbf{x}}(t) = \sum_{i=1}^{p} \mathbf{w_i}(t)\sigma\left(\mathbf{a_i}(t) \cdot \mathbf{x} + b_i(t)\right), \qquad t \in (0, T).$$

Neural ODEs

$$\dot{\mathbf{x}}(t) = \sum_{i=1}^{p} \mathbf{w_i}(t)\sigma\left(\mathbf{a_i}(t) \cdot \mathbf{x} + b_i(t)\right), \qquad t \in (0, T).$$

- **Control**: $\quad \theta := (\mathbf{w}_i, \mathbf{a}_i, b_i)_{i=1}^{p}, \quad \theta(t) \in L^{\infty}\left((0, T); \left(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}\right)^p\right).$
- **ReLU activation**: $\quad \sigma(z) = (z)_+ \quad$ Lipschitz, nonlinear.
- **Flow map** in time T generated by (3) is well defined:

$$\begin{array}{rcl} \Phi_T(\cdot; \theta) : \mathbb{R}^d & \to & \mathbb{R}^d \\ \mathbf{x_0} & \mapsto & \mathbf{x}(T; \mathbf{x_0}). \end{array}$$

Assume θ piecewise constant in $(0, T)$, $\quad \underbrace{L \text{ discontinuities}}_{\sim \text{ Transitions between layers}}$

$$\dot{\mathbf{x}}(t) = \sum_{i=1}^{p}\sum_{j=1}^{L} \mathbf{w_{i,j}}\sigma\left(\mathbf{a_{i,j}} \cdot \mathbf{x} + b_{i,j}\right)\mathbf{1}_{(t_{j-1}, t_j)}(t), \qquad t \in (0, T).$$

Problem statement

- **Worst-case scenario**:

$$\mathbf{x_n} \sim \mu = U([0,1]^d) \quad \text{for all } n.$$

-

-

Framework

- **Worst-case scenario**:

$$\mathbf{x_n} \sim \mu = U([0,1]^d) \quad \text{for all } n.$$

- **Binary classification**:

$$\mathcal{D} = \{(\mathbf{x_n}, y_n)\} \subset \mathbb{R}^d \times \{1, 0\}.$$

-

Framework

- **Worst-case scenario**:

$$\mathbf{x_n} \sim \mu = U([0,1]^d) \quad \text{for all } n.$$

- **Binary classification**:

$$\mathcal{D} = \{(\mathbf{x_n}, y_n)\} \subset \mathbb{R}^d \times \{1, 0\}.$$

-

$$1 \longleftrightarrow \Omega_1 := \{x^{(1)} > 0.5\}, \qquad 0 \longleftrightarrow \Omega_0 := \{x^{(1)} < 0.5\}.$$

Problem statement

$$\text{Fix } p = 1 : \qquad \dot{\mathbf{x}}(t) = \mathbf{w}(t)\, \sigma(\mathbf{a}(t) \cdot \mathbf{x}(t) + b(t)). \qquad (1)$$

Problem statement

$$\text{Fix } p = 1 : \qquad \dot{\mathbf{x}}(t) = \mathbf{w}(t)\,\sigma(\mathbf{a}(t) \cdot \mathbf{x}(t) + b(t)). \qquad (1)$$

For any $T > 0$, find a control $\theta = (\mathbf{w}, \mathbf{a}, b)$ for (1) such that

$$\Phi_T(\mathbf{x_n}; \theta) \in \Omega_{y_n} \qquad \text{for all } n$$

with the minimal possible complexity (number of discontinuities L).

Problem statement

$$\text{Fix } p = 1: \qquad \dot{\mathbf{x}}(t) = \mathbf{w}(t)\,\sigma(\mathbf{a}(t) \cdot \mathbf{x}(t) + b(t)). \tag{1}$$

For any $T > 0$, find a control $\theta = (\mathbf{w}, \mathbf{a}, b)$ for (1) such that

$$\Phi_T(\mathbf{x_n}; \theta) \in \Omega_{y_n} \qquad \text{for all } n$$

with the minimal possible complexity (number of discontinuities L).

Human's designed algorithm
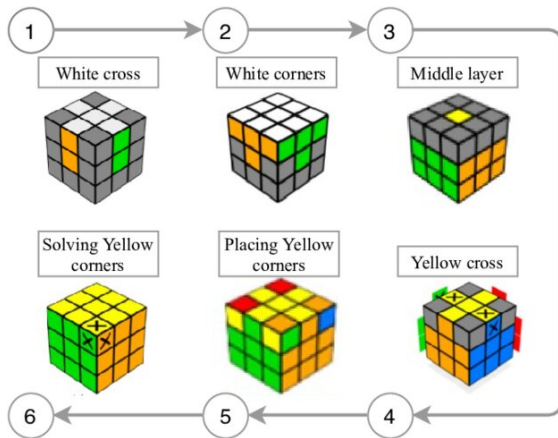
Source: "Two approaches in solving Rubik's cube with Hardware-Software Co-design", E. Barucija et al.

God's algorithm

God's algorithm is a notion originating in discussions of ways to solve the Rubik's Cube puzzle,[1] but which can also be applied to other combinatorial puzzles and mathematical games.[2] It refers to any algorithm which produces a solution having the fewest possible moves. The allusion to the deity is based on the notion that an omniscient being would know an optimal step from any given configuration.

An algorithm to determine the minimum number of moves to solve Rubik's Cube was published in 1997 by Richard Korf.[10] While it had been known since 1995 that 20 was a lower bound on the number of moves for the solution in the worst case, Tom Rokicki proved in 2010 that no configuration requires more than 20 moves.[11] Thus, 20 is a sharp upper bound on the length of optimal solutions. Mathematician David Singmaster had "rashly conjectured" this number to be 20 in 1980.[4]

Source: Wikipedia

How to move our Rubik's cube

$$\dot{\mathbf{x}}(t) = \mathbf{w}(t)\,\sigma(\mathbf{a}(t) \cdot \mathbf{x}(t) + b(t))$$

- $\mathbf{a}(t), b(t)$ determine the hyperplane in $\mathbb{R}^d$ given by

$$H(\mathbf{x}) = \mathbf{a}(t) \cdot \mathbf{x} + b(t) = 0.$$

- $\sigma(z) = (z)_+$ "activates" $H(\mathbf{x}) > 0$ and "freezes" $H(\mathbf{x}) \leq 0$.

- $\mathbf{w}(t)$ determines the direction of the field in $H(\mathbf{x}) > 0$.

From left to right: Compression, laminar motion, expansion.

Prior work

Theorem (D. Ruiz-Balet, E. Zuazua)

Let $d \geq 2$, and $R := \{\mathbf{x_n}\}$, $B := \{\mathbf{x_m}\} \subset \mathbb{R}^d$ finite datasets of different points. For any $T > 0$, there exists a control
$$\theta \in L^\infty \left((0, T); \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R} \right)$$
such that
$$\Phi_T(\mathbf{x_n}; \theta)^{(1)} > 0.5 \quad and \quad \Phi_T(\mathbf{x_m}; \theta)^{(1)} < 0.5,$$
for all $n = 1, \ldots, |R|$, $m = 1, \ldots, |B|$.
Furthermore, θ is piecewise constant and the number of switches is
$$L = 3(|R| + |B|) + 1.$$

Prior work

Number of discontinuities: $L = O(N)$. Classification algorithm from D. Ruiz-Balet, E. Zuazua, "Neural ode control for classification, approximation and transport" (2021).

But... how does the computer classify?

Probabilistic complexity

Theorem (A. Á-L, R. Orive-Illera, E. Zuazua)

Assume that $|R| = |B| = N$, and $\mathbf{x_n}, \mathbf{x_m} \sim U\left([0,1]^d\right)$, for all $n, m = 1, \ldots, N$.
For any $T > 0$, there exists a control $\theta = (\mathbf{w}, \mathbf{a}, b)$ where
$$\mathbf{a} \in \{\mathbf{e_1}, \ldots, \mathbf{e_d}\} \subset \mathbb{S}^{d-1} \quad const, \qquad (\mathbf{w}, b) \in L^{\infty}((0, T); \mathbb{R}^d \times \mathbb{R}),$$
such that
$$\Phi_T(\mathbf{x_n}; \theta)^{(1)} > 0.5 \quad and \quad \Phi_T(\mathbf{x_m}; \theta)^{(1)} < 0.5, \qquad for \ all \ n, m = 1, \ldots, N.$$
Furthermore, θ is piecewise constant with L discontinuities satisfying:

$$\mathbb{P}(L \leq k) = 1 - \left(\sum_{p=\lceil \frac{k}{2}+1 \rceil}^{N} \binom{N-1}{p-1}^2 + \sum_{p=\lceil \frac{k+1}{2} \rceil}^{N-1} \binom{N-1}{p} \binom{N-1}{p-1} \right)^d \left(\frac{2(N!)^2}{(2N)!} \right)^d,$$

for $k = 0, \ldots, 2N - 2$.

Probabilistic complexity

Theorem (A. Á-L, R. Orive-Illera, E. Zuazua)

Assume that $|R| = |B| = N$, and $\mathbf{x_n}, \mathbf{x_m} \sim U\left([0,1]^d\right)$, for all $n, m = 1, \ldots, N$.
For any $T > 0$, there exists a control $\theta = (\mathbf{w}, \mathbf{a}, b)$ where
$$\mathbf{a} \in \{\mathbf{e_1}, \ldots, \mathbf{e_d}\} \subset \mathbb{S}^{d-1} \quad const, \qquad (\mathbf{w}, b) \in L^\infty((0, T); \mathbb{R}^d \times \mathbb{R}),$$
such that
$$\Phi_T(\mathbf{x_n}; \theta)^{(1)} > 0.5 \quad and \quad \Phi_T(\mathbf{x_m}; \theta)^{(1)} < 0.5, \qquad for\ all\ n, m = 1, \ldots, N.$$
Furthermore, θ is piecewise constant with L discontinuities satisfying:

$$\mathbb{P}(L \leq k) = 1 - \left(\sum_{p=\lceil \frac{k}{2}+1 \rceil}^{N} \binom{N-1}{p-1}^2 + \sum_{p=\lceil \frac{k+1}{2} \rceil}^{N-1} \binom{N-1}{p} \binom{N-1}{p-1} \right)^d \left(\frac{2(N!)^2}{(2N)!} \right)^d,$$

for $k = 0, \ldots, 2N - 2$.

- **Linear separability** ($k = 0$): $\mathbb{P}(L = 0) \geq 1 - \left(\frac{2(N!)^2}{(2N)!} \right)^d \sim 1 - \exp\left\{ -\frac{\sqrt{\pi N}}{2^{2N-1}} d \right\}$.
- **Asymptotics**: $\mathbb{P}(L = 0) \to 1 \iff d \gtrsim 2^N/\sqrt{N}$.
- The **optimal number** of discontinuities $L^*(R, B)$ will satisfy
$$\mathbb{P}(L \leq k) \leq \mathbb{P}(L^* \leq k) \leq 1, \qquad for\ all\ k.$$

Probabilistic complexity

For $k = 0, \ldots, 2N - 2$:

$$\mathbb{P}(L \leq k) = 1 - \left(\sum_{p = \lceil \frac{k}{2} + 1 \rceil}^{N} \binom{N-1}{p-1}^2 + \sum_{p = \lceil \frac{k+1}{2} \rceil}^{N-1} \binom{N-1}{p} \binom{N-1}{p-1} \right)^d \left(\frac{2(N!)^2}{(2N)!} \right)^d,$$

Step 1: Separability in dimension $d = 1$

Let

$$Z(R, B) := \text{Minimum number of points needed to separate } (R, B)$$

- $Z = 2N - 1$ if

- $Z = 1$ if

$$P(Z = 1) = P(Z = 2N - 1) = \frac{2(N!)^2}{(2N)!}.$$

Step 1: Separability in dimension $d = 1$

Let

$Z(R, B) :=$ Minimum number of points needed to separate (R, B)

$$P(Z = k) = \begin{cases} \binom{N-1}{p-1}^2 \frac{2(N!)^2}{(2N)!}, & \text{if} \qquad k = 2p - 1, \\ \\ \binom{N-1}{p} \binom{N-1}{p-1} \frac{2(N!)^2}{(2N)!}, & \text{if} \qquad k = 2p, \end{cases}$$

for $k = 1, \ldots, 2N - 1$.

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a)   (b)

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a)

(b)

# Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a)   (b)

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a) $Z^1(R, B) = 3$ (b)

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a) $Z^1(R, B) = 3$

(b)

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a) $Z^1(R, B) = 3$ (b)

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

(a) $Z^1(R, B) = 3$

(b) $Z^2(R, B) = 4$

Step 2: Separability in dimension $d > 1$

Let

$$\pi_i := i\text{-th canonical projection in } \mathbb{R}^d, \qquad Z_i(R, B) := Z(\pi_i(R), \pi_i(B))$$

Z_i independent and identically distributed to Z for all i, so we can compute:

$$\mathbb{P}\left(\min_{i=1,\dots,d} \{Z_i\} \geq k \right) = \left(\mathbb{P}\left(Z \geq k \right) \right)^d.$$

Step 3: Control

The statement is concluded via $\quad \mathbb{P}(L \leq k) = \mathbb{P}\left(\min_{i=1,\dots,d} \{Z_i\} \leq k+1\right).$

Characterization of the worst-case scenario

Theorem (A. Á-L, R. Orive-Illera, E. Zuazua)

For any $R, B \subset \mathbb{R}^d$ with $|R| = |B| = N$ and $R \cap B = \emptyset$, the maximum number of hyperplanes required to separate their points is $2N - 1$.
Furthermore, the maximum is attained if and only if the $2N$ points of R and B are interspersed along a straight line in $\mathbb{R}^d$.

Characterization of the worst-case scenario

Theorem (A. Á-L, R. Orive-Illera, E. Zuazua)

For any $R, B \subset \mathbb{R}^d$ with $|R| = |B| = N$ and $R \cap B = \emptyset$, the maximum number of hyperplanes required to separate their points is $2N - 1$.

Furthermore, the maximum is attained if and only if the $2N$ points of R and B are interspersed along a straight line in $\mathbb{R}^d$.

For general $\mathbf{a} \in \mathbb{R}^d$, these configurations are very pathological!

Generic minimax bound

Theorem (A. Á-L, R. Orive-Illera, E. Zuazua)

Assume that $R \cup B \subset \mathbb{R}^d$ is in general position[a]. For any $T > 0$, there exists
$$\theta \in L^\infty \left((0, T); \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R} \right) \quad \text{such that}$$
$$\Phi_T(\mathbf{x_n}; \theta)^{(1)} > 0.5 \quad \text{and} \quad \Phi_T(\mathbf{x_m}; \theta)^{(1)} < 0.5, \qquad \text{for all } n = 1, \dots, N.$$

Furthermore, θ is piecewise constant with $L = 4 \left\lceil \frac{\min\{|R|, |B|\}}{d} \right\rceil - 1$ switches.

---

[a]no $d + 1$ points can lie on the same hyperplane

General position configurations in $\mathbb{R}^2$ and $\mathbb{R}^3$. Observe that they are generic when points are sampled from a non-singular measure.

Idea of the proof

Idea of the proof

Idea of the proof

Idea of the proof

Idea of the proof

# Idea of the proof

Idea of the proof

# Idea of the proof

Idea of the proof

Idea of the proof: Algorithm

Numerical experiments: Goal

Evaluate the performance of our algorithm.
How?
We estimate the minimum number of discontinuities L which allows to classify any dataset of N red and N blue points in $\mathbb{R}^d$ and compare it with the complexity required by our algorithms, given by $L_0 := 4\lceil N/d \rceil - 1$.

Numerical experiments: Setting

- **Data**. 10 datasets of $N = 30$ red points and $N = 30$ blue points sampled from $U([0, 1]^d)$ for various dimensions d ranging from 2 to $2N$.
- **Model**. Piecewise constant controls over $L + 1$ intervals of length $\Delta t = \frac{T}{L+1}$, where $T = 60$. Explicit Euler scheme with step size of $0.25\Delta t$.
- **Complexity**. For each d, we verify that $L = L_0$ discontinuities suffice to classify all datasets. Then we reduce L until any dataset fails to be classified.
- **Error**.

$$\frac{1}{|R|} \sum_{\mathbf{x}_n \in R} \text{dist} \left(\Phi_T(\mathbf{x}_n, \theta), \{x^{(1)} < 1\} \right)^2 + \frac{1}{|B|} \sum_{\mathbf{x}_n \in B} \text{dist} \left(\Phi_T(\mathbf{x}_n, \theta), \{x^{(1)} > 0\} \right)^2.$$

- **Training.** Adam optimizer with a learning rate of 0.01.
- **Stopping criteria**.
  1. The maximum number of 70000 epochs is reached.
  2. Slow convergence, if $\mathcal{L} \geq 0.15$ at 20000 epochs or $\mathcal{L} \geq 0.1$ at 40000 epochs, or if the minimum error does not decrease over 5000 consecutive epochs.
  3. Local minima detection, if the maximum relative error on every 50 consecutive epochs exceeds a threshold of 10^{-20}.

Numerical experiments: Goal

Evaluate the performance of our algorithm.
How?
We estimate the minimum number of discontinuities L which allows to classify any dataset of N red and N blue points in $\mathbb{R}^d$ and compare it with the complexity required by our algorithms, given by $L_0 := 4\lceil N/d \rceil - 1$.

Close. We conclude that the trained model does not have the capacity to classify with L discontinuities if a stopping criterion is met for 20 randomized initializations of the parameters in any of the 10 datasets.


Available codes in GitHub

Numerical experiments: Results (*L* vs *d*)

(a) Green bars → Our algorithm. Purple bars → Numerical training.

(b) After rescaling to zero mean and unit variance.

Conclusions

- We have found a new minimal value of L that ensures that any given dataset in $\mathbb{R}^d$ can be almost surely classified.
- In any dataset we can generically find clusters of d points that can be isolated from the rest using two parallel hyperplanes.
- For any $N \geq 1$, if $d \gtrsim 2^N/\sqrt{N}$ then two sets of N points can be almost surely classified with $L = 0$ discontinuities (autonomous neural ODE).

Open problems

- Further reduce the maximum number of switches.
- Is $L = 4\lceil N/d \rceil - 1$ sharp? In such case, what configurations attain this maximality?
- Role played by the chosen boundary decision? Which is the optimal one?
- Potential application for initialization of parameters in training.

📄 Antonio Álvarez-López, Rafael Orive-Illera, and Enrique Zuazua.
Optimized classification with neural ODEs via separability.
arXiv preprint arXiv:2312.13807, 2023.

📄 Antonio Álvarez-López, Arselane H. Slimane, and Enrique Zuazua.
Interplay between depth and width for interpolation in neural ODEs.
Neural Networks, 180:106640, 2024.

📄 Domènec Ruiz-Balet and Enrique Zuazua.
Neural ODE Control for Classification, Approximation, and Transport.
SIAM Review, 65(3):735--773, 2023.

Thank you for your attention!

Scan me!

Analogy

Further improvements, and finding God's Number   [edit]

In 2006, Silviu Radu further improved his methods to prove that every position can be solved in at most 27 face turns or 35 quarter turns.[20] Daniel Kunkle and Gene Cooperman in 2007 used a supercomputer to show that all unsolved cubes can be solved in no more than 26 moves (in face-turn metric). Instead of attempting to solve each of the billions of variations explicitly, the computer was programmed to bring the cube to one of 15,752 states, each of which could be solved within a few extra moves. All were proved solvable in 29 moves, with most solvable in 26. Those that could not initially be solved in 26 moves were then solved explicitly, and shown that they too could be solved in 26 moves.[21][22]

Tomas Rokicki reported in a 2008 computational proof that all unsolved cubes could be solved in 25 moves or fewer.[23] This was later reduced to 23 moves.[24] In August 2008, Rokicki announced that he had a proof for 22 moves.[25]

Finally, in 2010, Tomas Rokicki, Herbert Kociemba, Morley Davidson, and John Dethridge gave the final computer-assisted proof that all cube positions could be solved with a maximum of 20 face turns.[2] In 2009, Tomas Rokicki proved that 29 moves in the quarter-turn metric is enough to solve any scrambled cube.[26] And in 2014, Tomas Rokicki and Morley Davidson proved that the maximum number of quarter-turns needed to solve the cube is 26.[3]