LyaNet: A Lyapunov Framework for Training Neural ODEs

Ivan D. Jimenez Rodriguez, Aaron D. Ames, Yisong Yue



AeroVironment™

ARGO AI

Caltech

Challenges

- Dynamics for single data point (x, y).
- 2-Dimensional Hidden Logits State
- Binary Classification: Red Class is correct

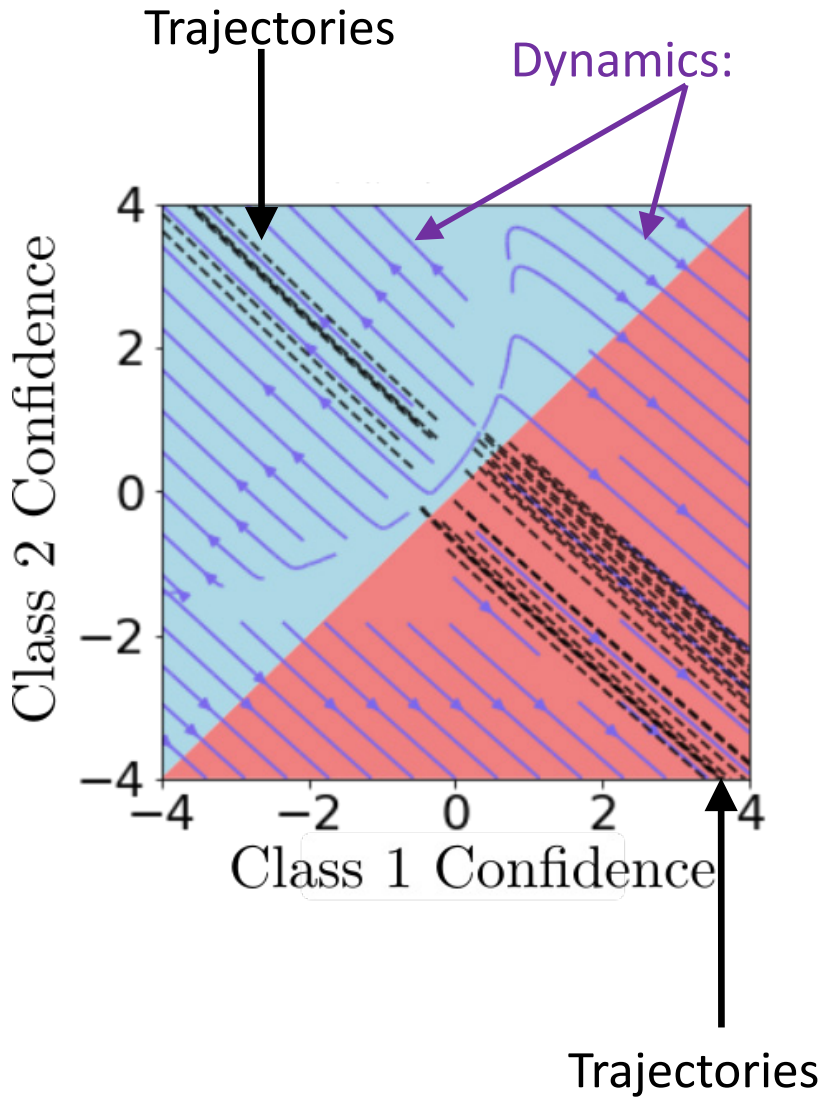

Undesirable Chaotic Dynamics:

- Dynamics often flow in the wrong direction
- Solutions are "fragile"
- Poor Generalization

Idea: Enforce Always Making Progress

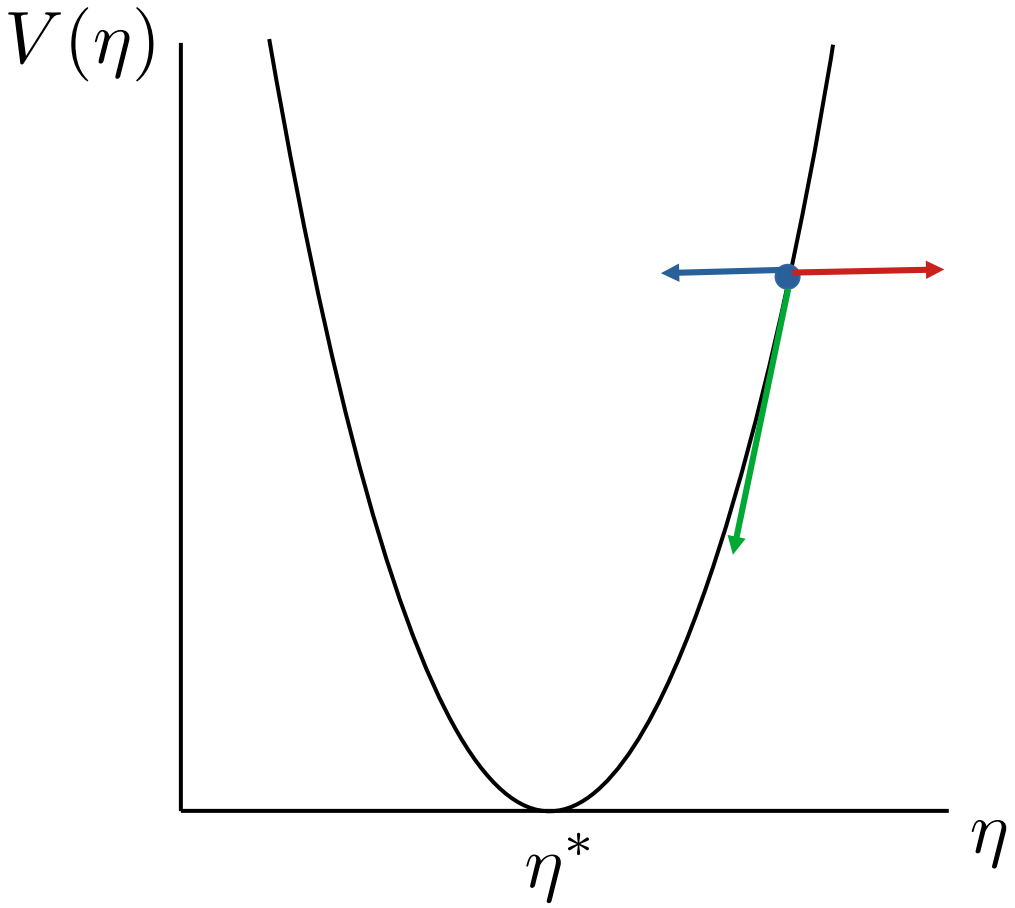

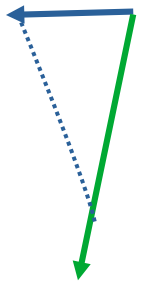

Progress

$$\frac{\partial V}{\partial \eta}^\top f(\eta) \leq -\kappa V(\eta)$$

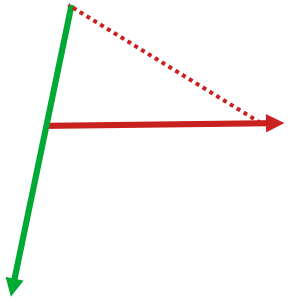

No Progress

$$\frac{\partial V}{\partial \eta}^\top f(\eta) > -\kappa V(\eta)$$

Measuring Progress on Supervised Loss

For a sample, define a potential function:

$$V_{\eta^*}(\eta) = \mathcal{L}(\eta, \eta^*)$$
$$= \|\eta - \eta^*\|_2^2$$

Similar potential function can be selected for Cross-Entropy.

Exponential Convergence (CIFAR-10)

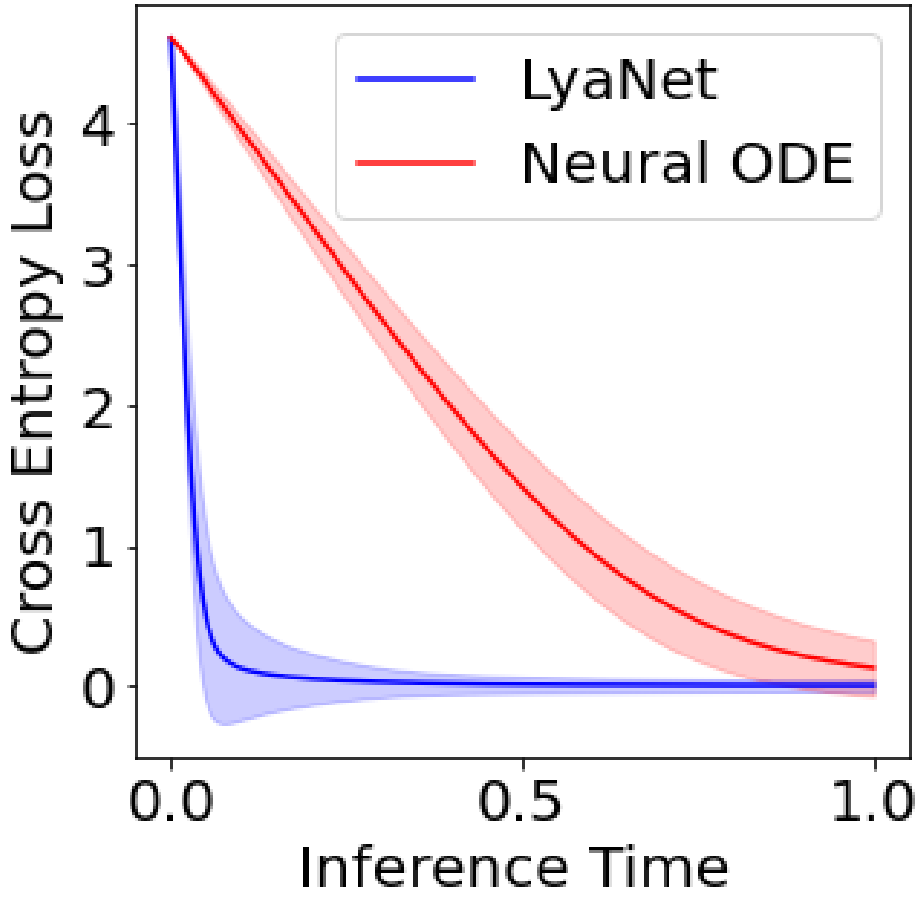

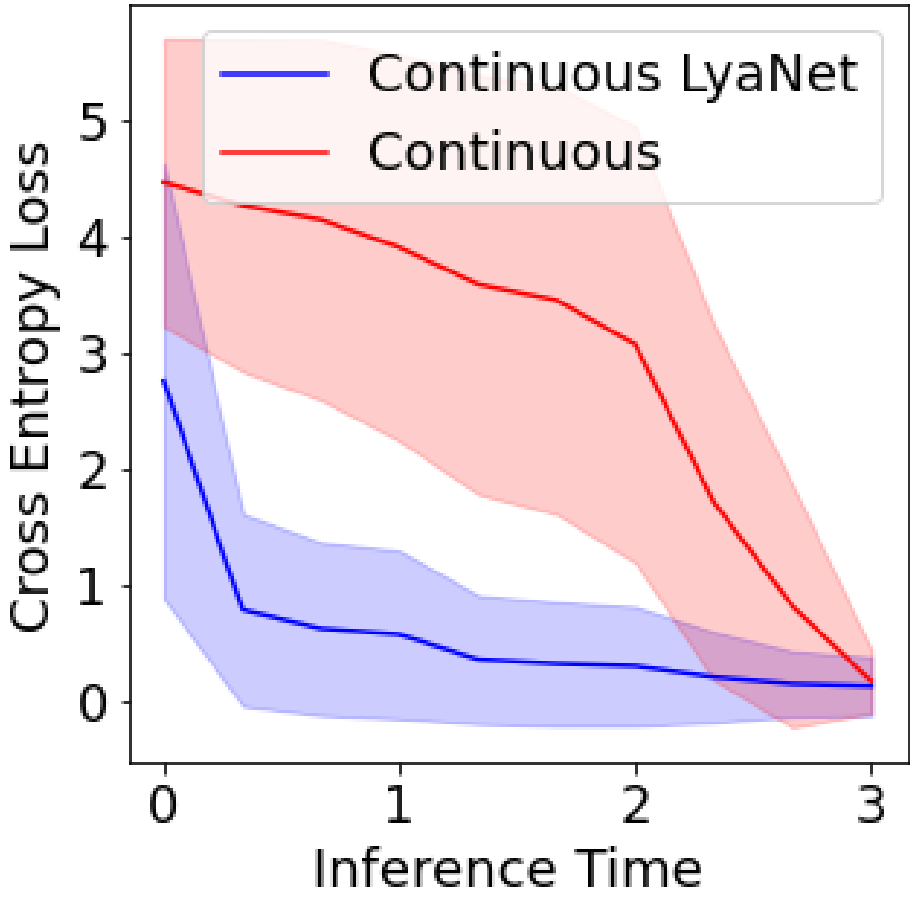

Our Method: Robustness Guarantees (Illustrated)

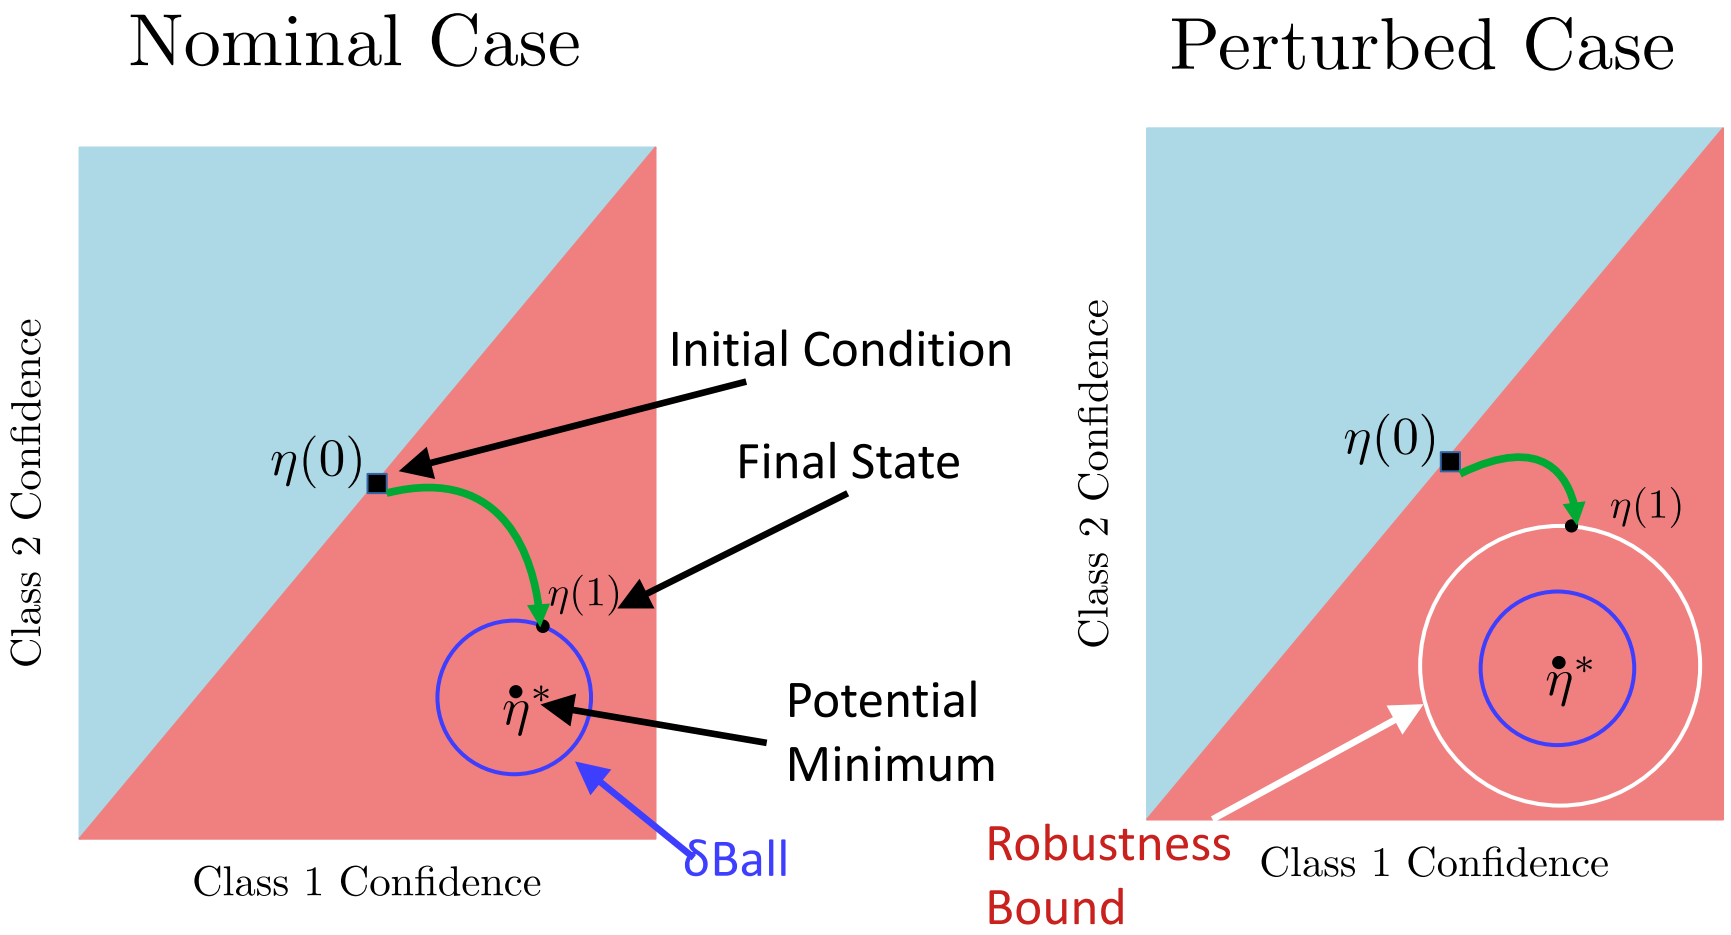

Conclusions

- **LyaNet:** A novel approach for training Neural ODEs using Lyapunov Stability
- **Main benefits:**
  - Fast convergence of inference dynamics
  - Empirical and provable robustness
  - Allows adaptation of control theoretic tools to the learning context
- **Other Applications:**
  - Enforce constraints on the learned dynamics

Paper Extension Website