Safety and Resilience of Large-scale Networks via Contraction Theory

Saber Jafarpour

Decision and Control Laboratory
Georgia Institute of Technology

March 3, 2022

Large-scale Nonlinear Networks

Power grid

Transportation network

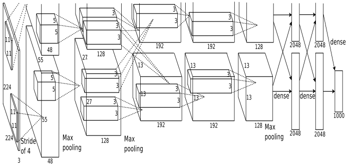

Artificial neural network

Nonlinearity:

- Multiple equilibria
- Transient stability
- Congestion

Large-scale:

- Stochastic
- Distributed

- *"... in Oahu, Hawaii, at least 800,000 micro-inverters interconnect photovoltaic panels to the grid... "* [IEEE Spectrum, 2015]
- *"... Americans lost an average of 97 hours a year due to congestion, costing them nearly 87 billion dollars in 2018 ... "* [INRIX report, 2018]
- *"... autonomous vehicles should fuse a large amount of data from cameras, radar and LiDAR sensors ... "*

Large-scale Nonlinear Networks

Stability and Robustness

Critical task: ensuring safe and reliable operation.

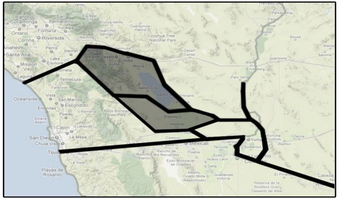

2011 US Southwest blackout

Traffic congestion in Beijing

Tesla self-driving accident

This task is challenging:

- large size of the networks
- nonlinear interactions
- unknown components
- dynamic and stochastic environment

My contribution: Tools and techniques from control theory and dynamical systems

> Rigorous systematic approaches to ensure safety and resilience

Stability and control of large-scale power grids

1. threshold of frequency synchronization
2. multi-stability via partitioning the state-space
3. dynamic stability of low-inertia power grids

SJ, E. Y. Huang, K. D. Smith, and F. Bullo, *Flow and Elastic Networks on the n-torus: Geometry, Analysis, and Computation*, SIAM Review, Research Spotlight, 2021.

SJ and F. Bullo, *Synchronization of Kuramoto Oscillators via Cutset Projections*, IEEE Transactions on Automatic Control, 2019.

Robustness of neural networks

1. elements of a non-Euclidean contraction theory
2. ℓ_∞-norm robustness analysis of implicit neural networks

SJ and A. Davydov and A. Proskurnikov and F. Bullo. *Robust Implicit Networks via Non-Euclidean Contractions*. NeurIPS 2021.

A. Davydov and SJ and F. Bullo. *Non-Euclidean Contraction Theory for Robust Nonlinear Stability*. arXiv: https://arxiv.org/abs/2103.12263, 2021.

Resilience of dynamic flow networks

1. extensions of contraction theory for large-scale real-world networks
2. robustness of transportation systems to large-size attacks or perturbations

SJ and P. Cisneros-Velarde and F. Bullo. *Weak and Semi-Contraction for Network Systems and Diffusively-Coupled Oscillators*. IEEE Transactions on Automatic Control, 2021.

SJ and S. Coogan. *Resilience of Input Metering in Dynamic Flow Networks*. American Control Conference, to appear, 2022.

Presentation outline

- **Robustness of neural networks**
  - implicit neural networks
  - well-posedness using contraction theory
  - robustness via Lipschitz bounds
- Resilience of dynamic flow networks
  - dynamic flow networks
  - dichotomy in asymptotic behavior
  - robustness of transportation networks to failures

Neural Networks in Autonomous Systems

Robustness issues

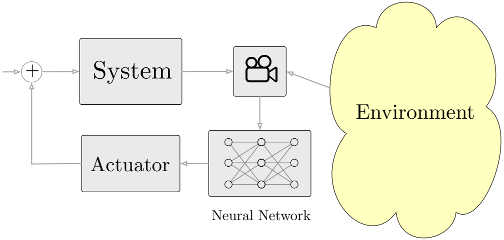

Promising performance

- large amount of data is available
- high-dimensional input
- complicated behaviors
- little knowledge about the system

Robustness challenges

- vulnerable to input perturbations

- Safety- and security-critical applications:

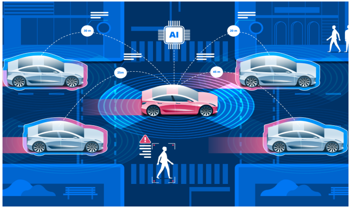

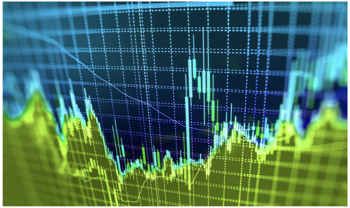

Adversarial Perturbations
Definition and Examples

Small changes in the input lead to *large* changes in the output

C. Szegedy and et. al. Intriguing properties of neural networks. In *ICLR*, 2014

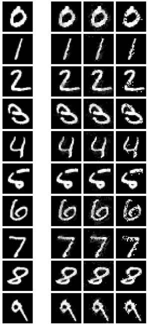

Left column is classified correctly but the three right columns are classified as $i + 1$ (mod 10)

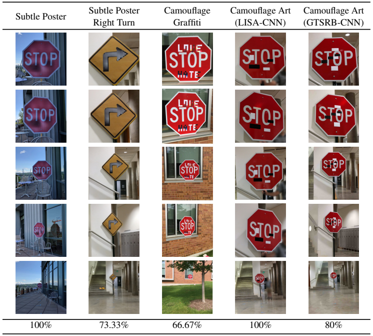

All images are after perturbation and are classified as 45 mph speed limit sign

A rigorous measure for input-output sensitivity of neural networks is **Lipschitz constant**

Input-output Lipschitz constant

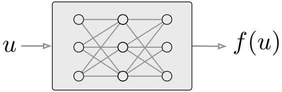

$$\|f(u) - f(v)\| \leq L\|u - v\|, \qquad \text{for all} \quad u, v \in \mathbb{R}^n$$

1. most common norms: ℓ_2 and ℓ_∞
   - ℓ_2-norm Lipschitz constant: change in energy-level
   - ℓ_∞-norm Lipschitz constant: component-wise change
2. computing the input-output Lipschitz constant is NP-hard

   A. Virmaux and K. Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *NeurIPS*, 2018
3. extensive research on estimating Lipschitz constant of neural networks

   M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. J. Pappas. Efficient and accurate estimation of Lipschitz constants for deep neural networks. In *NeurIPS*, 2019

- explicit hidden layers are replaced by a single implicit layer

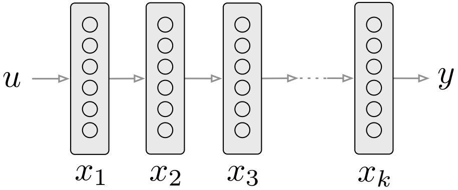

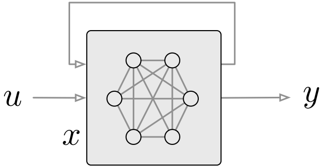

- explicit hidden layers are replaced by a single implicit layer

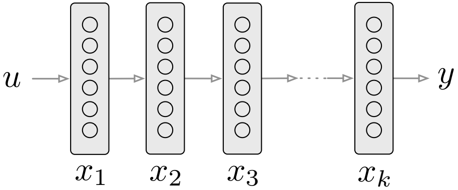

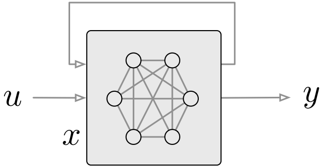

- traditional neural networks:

$$x^{i+1} = \phi_i(A_i x^i + b_i), \quad x^0 = u$$
$$y = A_k x^k + c$$

- activation functions are slope-restricted in $[0, 1]$, i.e., $0 \leq \frac{\phi_i(x)-\phi_i(y)}{x-y} \leq 1$ for all $x, y \in \mathbb{R}$

- explicit hidden layers are replaced by a single implicit layer

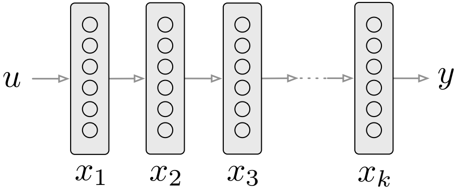

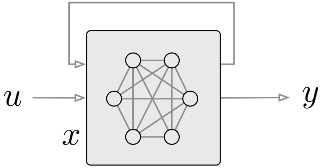

- traditional neural networks:

$$x^{i+1} = \phi_i(A_i x^i + b_i), \quad x^0 = u$$
$$y = A_k x^k + c$$

- implicit neural networks:

$$x = \Phi(Ax + Bu + b)$$
$$y = Cx + c$$

- activation functions are slope-restricted in $[0, 1]$, i.e., $0 \leq \frac{\phi_i(x)-\phi_i(y)}{x-y} \leq 1$ for all $x, y \in \mathbb{R}$
- $\Phi((y_1, \ldots, y_n)) = (\phi_1(y_1), \ldots, \phi_n(y_n))^\top$ is a diagonal activation function.

- Origins:

S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *NeurIPS*, 2019

L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Y. Tsai. Implicit deep learning. *SIMODS*, 2019

A. Kag, Z. Zhang, and V. Saligrama. RNNs incrementally evolving on an equilibrium manifold: A panacea for vanishing and exploding gradients? In *ICLR*, 2020

- Generalizing feedforward neural networks to fully-connected synaptic matrices

Intuition: $x^{i+1} = \phi_i(A_i x^i + b_i) \quad \Longleftrightarrow \quad x = \Phi(Ax + Bu + b)$, where A has upper diagonal structure.

$A_{\text{upper-diagonal}} =$ 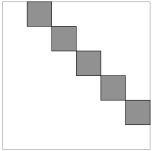

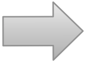

$A_{\text{complete}} =$ 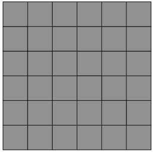

Implicit Neural Networks (INNs)

Origin and Motivations

- comparable accuracy to traditional neural networks with significant memory reduction

  S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *NeurIPS*, 2019

Intuition: implicit neural network = weight-tied infinite-layer network

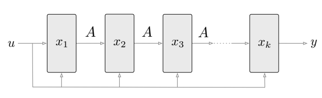

$x^{i+1} = \phi_i(Ax^i + B_i u + b_i) \implies \lim_{i\to\infty} x^i = x^*$ solution to the INN

- suitable for learning constrained optimization problems

  A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *NeurIPS*, 2019

Intuition: casting KKT condition as an implicit layer

Implicit Neural Networks (INNs)

Origin and Motivations

- vanishing and exploding gradient

  A. Kag, Z. Zhang, and V. Saligrama. RNNs incrementally evolving on an equilibrium manifold: A panacea for vanishing and exploding gradients? In *ICLR*, 2020

Intuition: the notion of "autapse" (time-delayed self-feedback) from neuroscience

$A_{\text{upper-diagonal}} =$ 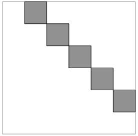

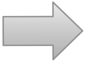

$A_{\text{Autapse}} =$ 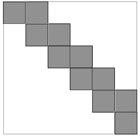

- suitable for learning stiff problems or problems with discontinuity

  S. Pfrommer, M. Halm, and M. Posa. ContactNets: Learning discontinuous contact dynamics with smooth, implicit representations. *arXiv preprint*, 2020

- **Challenge 1**: well-posedness, i.e., existence and uniqueness of

$$x = \Phi(Ax + Bu + b)$$

- **Challenge 2**: convergence stability, i.e., algorithms for computing the solution of

$$x = \Phi(Ax + Bu + b)$$

- **Challenge 3**: computing robustness margin, i.e., estimate Lipschitz bound for INNs

- **Challenge 4**: implementing robustness in the training

> **Goal**: develop a rigorous framework to study these challenges

$\dot{x} = \mathsf{G}(t, x)$ is contractive if its flow is a contraction map

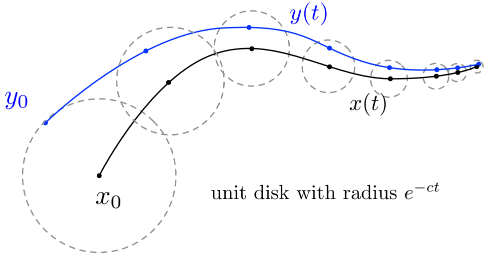

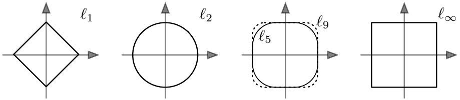

- **Origins**

D. C. Lewis. Metric properties of differential equations. *American Journal of Mathematics*, 71(2):294–312, 1949

B. P. Demidovich. Dissipativity of a nonlinear system of differential equations. *Uspekhi Matematicheskikh Nauk*, 16(3(99)):216, 1961

C. A. Desoer and H. Haneda. The measure of a matrix as a tool to analyze computer algorithms for circuit analysis. *IEEE Transactions on Circuit Theory*, 19(5):480–486, 1972.

- **Origins**

D. C. Lewis. Metric properties of differential equations. *American Journal of Mathematics*, 71(2):294–312, 1949

B. P. Demidovich. Dissipativity of a nonlinear system of differential equations. *Uspekhi Matematicheskikh Nauk*, 16(3(99)):216, 1961

C. A. Desoer and H. Haneda. The measure of a matrix as a tool to analyze computer algorithms for circuit analysis. *IEEE Transactions on Circuit Theory*, 19(5):480–486, 1972.

- **Application in control theory**:

W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998

Contraction theory

Historical notes

- **Origins**

D. C. Lewis. Metric properties of differential equations. *American Journal of Mathematics*, 71(2):294–312, 1949

B. P. Demidovich. Dissipativity of a nonlinear system of differential equations. *Uspekhi Matematicheskikh Nauk*, 16(3(99)):216, 1961

C. A. Desoer and H. Haneda. The measure of a matrix as a tool to analyze computer algorithms for circuit analysis. *IEEE Transactions on Circuit Theory*, 19(5):480–486, 1972.

- **Application in control theory**:

W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998

- **Reviews**:

Z. Aminzare and E. D. Sontag. Contraction methods for nonlinear systems: A brief introduction and some open problems. In *Proc CDC*, pages 3835–3847, Dec. 2014

M. Di Bernardo, D. Fiore, G. Russo, and F. Scafuti. Convergence, consensus and synchronization of complex networks via contraction theory. In *Complex Systems and Networks: Dynamics, Controls and Applications*, pages 313–339. Springer, 2016

H. Tsukamotoa, S.-J. Chung, and J.-J. E. Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview, 2021. URL https://arxiv.org/abs/2110.00675

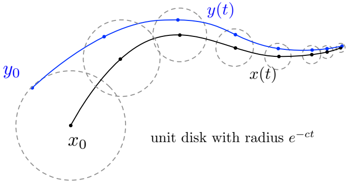

Highly ordered **transient** and **asymptotic** behavior:

1. time-invariant G: unique globally exponential stable equilibrium
   two natural Lyapunov functions

2. periodic G: contracting system entrain to periodic inputs

3. strong robustness properties: contractivity rate is natural measure of robust stability
   input-to-state stability in presence of un-modeled dynamics

4. accurate numerical integration and efficient methods for their equilibrium computation

The **matrix measure** of $A \in \mathbb{R}^{n\times n}$ wrt to $\|\cdot\|$:

$$\mu_{\|\cdot\|}(A) := \lim_{h\to 0^+} \frac{\|I_n + hA\| - 1}{h}.$$

- Directional derivative of norm $\|\cdot\|$ in direction of A,

The **matrix measure** of $A \in \mathbb{R}^{n\times n}$ wrt to $\|\cdot\|$:

$$\mu_{\|\cdot\|}(A) := \lim_{h\to 0^+} \frac{\|I_n + hA\| - 1}{h}.$$

- Directional derivative of norm $\|\cdot\|$ in direction of A,

$$\mu_2(A) = \tfrac{1}{2}\lambda_{\max}(A + A^\top)$$

$$\mu_1(A) = \max_j \left(a_{jj} + \sum\nolimits_{i\neq j} |a_{ij}|\right) \qquad \mu_\infty(A) = \max_i \left(a_{ii} + \sum\nolimits_{j\neq i} |a_{ij}|\right)$$

The **matrix measure** of $A \in \mathbb{R}^{n\times n}$ wrt to $\|\cdot\|$:

$$\mu_{\|\cdot\|}(A) := \lim_{h\to 0^+} \frac{\|I_n + hA\| - 1}{h}.$$

- Directional derivative of norm $\|\cdot\|$ in direction of A,

$$\mu_2(A) = \tfrac{1}{2}\lambda_{\mathsf{max}}(A + A^\top)$$

$$\mu_1(A) = \max_j \left(a_{jj} + \sum\nolimits_{i\neq j} |a_{ij}|\right) \qquad \mu_\infty(A) = \max_i \left(a_{ii} + \sum\nolimits_{j\neq i} |a_{ij}|\right)$$

- One-sided Lipschitz constant

  E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I. Nonstiff Problems*. 1993

- Logarithmic norm

  T. Ström. On logarithmic norms. *SIAM Journal on Numerical Analysis*, 1975

$\dot{x} = \mathsf{G}(t,x)$ is contractive if its flow is a contraction map

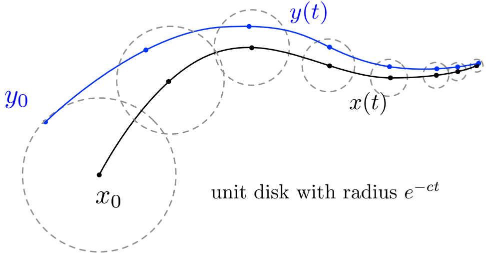

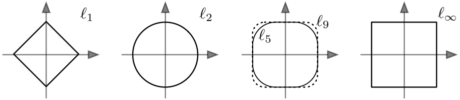

Dynamical system $\dot{x} = \mathsf{G}(t,x)$ is contracting with respect to the norm $\|\cdot\|$ iff

$$\mu(D\mathsf{G}(t,x)) \leq -c, \qquad \text{for all } x, t$$

$$\ell_2 - \textbf{contraction} \qquad\qquad \textbf{LMI}$$

$$\mu_2(D\mathsf{G}(t,x)) \leq -c \quad \Longleftrightarrow \quad D\mathsf{G}(t,x) + D\mathsf{G}(t,x)^\top \preceq -cI$$

- Monotone Operator Theory

  E. K. Ryu and S. Boyd. Primer on monotone operator methods. *Applied Computational Mathematics*, 2016

ℓ_2 – **contraction** | **LMI**

$$\mu_2(D\mathsf{G}(t,x)) \leq -c \quad \Longleftrightarrow \quad D\mathsf{G}(t,x) + D\mathsf{G}(t,x)^\top \preceq -cI$$

- Monotone Operator Theory

  E. K. Ryu and S. Boyd. Primer on monotone operator methods. *Applied Computational Mathematics*, 2016

ℓ_1/ℓ_∞ – **contraction** | **Diagonal Dominance**

$$\mu_\infty(D\mathsf{G}(t,x)) \leq -c, \quad \Longleftrightarrow \quad (D\mathsf{G}(t,x))_{ii} + \sum_{j\neq i} |(D\mathsf{G}(t,x))_{ij}| \leq -c, \quad \forall i$$

$$\mu_1(D\mathsf{G}(t,x)) \leq -c, \quad \Longleftrightarrow \quad (D\mathsf{G}(t,x))_{ii} + \sum_{j\neq i} |(D\mathsf{G}(t,x))_{ji}| \leq -c, \quad \forall i$$

Challenge 1: well-posedness and **Challenge 2:** convergence stability

Problem statement

For a fixed-point equation

$$x = \mathsf{F}(x,u) \qquad \text{(for implicit neural networks } \mathsf{F}(x,u) = \Phi(Ax + Bu + b))$$

1. when do we have a unique solution?
2. how to efficiently compute it?

Challenge 1: well-posedness and **Challenge 2:** convergence stability

Problem statement

For a fixed-point equation

$$x = \mathsf{F}(x,u) \qquad \text{(for implicit neural networks } \mathsf{F}(x,u) = \Phi(Ax + Bu + b)\text{)}$$

1. when do we have a unique solution?
2. how to efficiently compute it?

Banach Fixed-point Theorem: if $\|D_x\mathsf{F}(x,u)\| < 1$, then $x = \mathsf{F}(x,u)$ has a unique solution by the Picard iterations

$$x^{k+1} = \mathsf{F}(x^k, u).$$

Solvability of fixed-point equations

A contraction-based framework

Key insight

Fixed-point of	$\iff$	Equilibrium point of
$x = \mathsf{F}(x,u)$		$\dot{x} = -x + \mathsf{F}(x,u)$

- **Contraction theory**: existence and uniqueness of equilibrium point

$$\mu(D_x \mathsf{F}(x,u)) < 1.$$

Solvability of fixed-point equations

A contraction-based framework

Key insight

Fixed-point of $x = \mathsf{F}(x, u)$ $\iff$ Equilibrium point of $\dot{x} = -x + \mathsf{F}(x, u)$

- **Contraction theory**: existence and uniqueness of equilibrium point

$$\mu(D_x \mathsf{F}(x, u)) < 1.$$

Theorem: Fixed-point via matrix measures

If $\mu(D_x \mathsf{F}(x, u)) < 1$ then

1. F has a unique fixed-point x_u^*.
2. $x^{k+1} = (1-\alpha)x^k + \alpha \mathsf{F}(x^k, u)$ converges to x_u^*, for $0 < \alpha \leq \alpha^*$.

Theorem: Fixed-point via norm

If $\|D_x \mathsf{F}(x, u))\| < 1$ then

1. F has a unique fixed-point x_u^*.
2. $x^{k+1} = \mathsf{F}(x^k, u)$ converges to x_u^*.

S. Jafarpour, A. Davydov, A. V. Proskurnikov, and F. Bullo. Robust implicit networks via non-Euclidean contractions. In *NeurIPS*, Dec. 2021b

$\mu(D_x\mathsf{F}(x,u)) < 1$ is less conservative than $\|D_x\mathsf{F}(x,u)\| < 1$.

- $\mathsf{F}(x,u) = Ax + Bu$ with $A = \begin{bmatrix} a & b \\ b & a \end{bmatrix}$

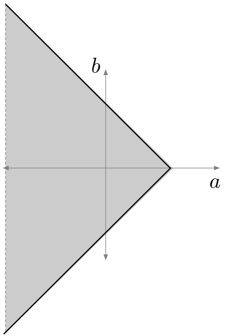

$\mu_\infty(A) \le 1$
unbounded

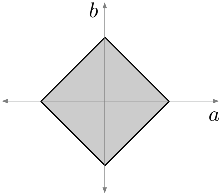

$\|A\|_\infty \le 1$
bounded

Well-posedness of INNs

Computing fixed-points

$$x = \Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\mu_\infty(D\Phi(x)A) < 1$, for every x, then

1. there exists a unique fixed-point,
2. for $\alpha \in]0, (1 - \min_i(a_{ii})_-))^{-1}]$, the average map is a contraction map:

$$\mathsf{N}_\alpha(x) := (1-\alpha)x + \alpha\Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\|D\Phi(x)A\|_\infty < 1$, for every x, then

1. there exists a unique fixed-point,
2. the Picard iterations

$$x^{k+1} := \Phi(Ax^k + Bu + b)$$

is a contraction map.

$$x = \Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\mu_\infty(A) < 1$, then

1. there exists a unique fixed-point,
2. for $\alpha \in]0, (1 - \min_i (a_{ii})_-))^{-1}]$, the average map is a contraction map:

$$\mathsf{N}_\alpha(x) := (1-\alpha)x + \alpha\Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\|A\|_\infty < 1$, then

1. there exists a unique fixed-point,
2. the Picard iterations

$$x^{k+1} := \Phi(Ax^k + Bu + b)$$

is a contraction map.

$$x = \Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\mu_\infty(A) < 1$, then

1. there exists a unique fixed-point,
2. for $\alpha \in \,]0, (1 - \min_i(a_{ii})_-))^{-1}]$, the average map is a contraction map:

$$\mathsf{N}_\alpha(x) := (1-\alpha)x + \alpha\Phi(Ax + Bu + b)$$

Theorem: Fixed-points of INNs

If $\|A\|_\infty < 1$, then

1. there exists a unique fixed-point,
2. the Picard iterations

$$x^{k+1} := \Phi(Ax^k + Bu + b)$$

is a contraction map.

The iteration $x^{k+1} = \mathsf{N}_\alpha(x^k)$ is Euler discretization of

$$\dot{x} = -x + \Phi(Ax + Bu + b)$$

S. Jafarpour, A. Davydov, A. V. Proskurnikov, and F. Bullo. Robust implicit networks via non-Euclidean contractions. In *NeurIPS*, Dec. 2021b

Challenge 3: Robustness margins

Problem statement

How does the fixed-point of $x = \mathsf{F}(x, u)$ change with u?

Challenge 3: Robustness margins

Problem statement

How does the fixed-point of $x = \mathsf{F}(x, u)$ change with u?

Theorem: Input-to-state Lipschitz bounds

x_u^* is a fixed-point of $x = \mathsf{F}(x, u)$ and $\mu(D_x\mathsf{F}) < 1$,

$$\|x_u^* - x_v^*\| \leq \frac{\|D_u\mathsf{F}\|}{1 - \mu(D_x\mathsf{F})}\|u - v\|$$

Theorem: Input-to-state Lipschitz bounds

x_u^* is a fixed-point of $x = \mathsf{F}(x, u)$ and $\|D_x\mathsf{F}\| < 1$,

$$\|x_u^* - x_v^*\| \leq \frac{\|D_u\mathsf{F}\|}{1 - \|D_x\mathsf{F}\|}\|u - v\|$$

S. Jafarpour, A. Davydov, A. V. Proskurnikov, and F. Bullo. Robust implicit networks via non-Euclidean contractions. In *NeurIPS*, Dec. 2021b

$$x = \Phi(Ax + Bu + b),$$
$$y = Cx + c$$

- How to compute Lipschitz bounds in INNs?

$$u \underbrace{\mapsto}_{\mathrm{Lip}_{u \to x^*}} x^* \underbrace{\mapsto}_{\mathrm{Lip}_{x^* \to y}} y \quad \Longrightarrow \quad \mathrm{Lip}_{u \to y} = \mathrm{Lip}_{u \to x^*} \mathrm{Lip}_{x^* \to y}$$

$$x = \Phi(Ax + Bu + b),$$
$$y = Cx + c$$

- How to compute Lipschitz bounds in INNs?

$$u \underbrace{\mapsto}_{\mathrm{Lip}_{u\to x^*}} x^* \underbrace{\mapsto}_{\mathrm{Lip}_{x^*\to y}} y \quad \implies \quad \mathrm{Lip}_{u\to y} = \mathrm{Lip}_{u\to x^*}\mathrm{Lip}_{x^*\to y}$$

Theorem: Input-to-output Lipschitz

if $\mu_\infty(A) < 1$ then

$$\mathsf{Lip}_{u\to y} = \frac{\|B\|_\infty \|C\|_\infty}{1 - \mu_\infty(A)_+}.$$

Theorem: Input-to-output Lipschitz

if $\|A\|_\infty < 1$ then

$$\mathsf{Lip}_{u\to y} = \frac{\|B\|_\infty \|C\|_\infty}{1 - \|A\|_\infty}.$$

S. Jafarpour, A. Davydov, A. V. Proskurnikov, and F. Bullo. Robust implicit networks via non-Euclidean contractions. In *NeurIPS*, Dec. 2021b

Intuition #1: Weight-tied infinite-depth NN $\rightarrow$ fixed-point of INN

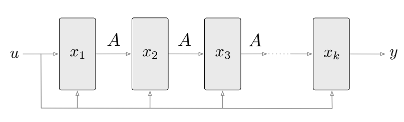

contraction of $x^{i+1} = \Phi(Ax^i + B_i u + b_i) \implies \lim_{i\to\infty} x^i = x^*$ solution to the INN

Intuition #1: Weight-tied infinite-depth NN $\rightarrow$ fixed-point of INN

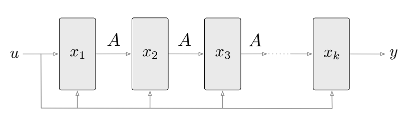

contraction of $x^{i+1} = \Phi(Ax^i + B_i u + b_i) \implies \lim_{i\to\infty} x^i = x^*$ solution to the INN

Intuition #2: Neural ODE model (infinite time) $\rightarrow$ fixed-point of INN

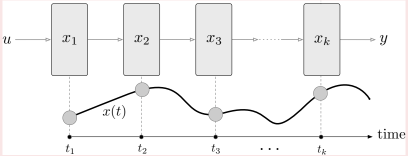

contraction of $\dot{x} = -x + \Phi(Ax + Bu + b) \implies \lim_{t\to\infty} x(t) = x^*$ solution to INN

R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. In *NeurIPS*, 2018

Challenge 4: training of robust and well-posed INNs

1. loss function $\mathcal{L}$
2. training data $(\widehat{u}_i, \widehat{y}_i)_{i=1}^N$

$$\min_{A,B,C,b,c} \quad \sum_{i=1}^{N} \mathcal{L}(\widehat{y}_i, Cx_i + c) \quad +\lambda \quad \mathsf{Lip}_{u \to y}$$

$$x_i = \Phi(Ax_i + B\widehat{u}_i + b)$$

$$\mu_\infty(A) \leq \gamma,$$

- $\gamma < 1$ is a hyperparameter and $\lambda \geq 0$ is a regularization parameter
- training optimization problem is solved via SGD
- at each step of SGD, $x_i = \Phi(Ax_i + B\widehat{u}_i + b)$ is solved using the average-iterations

Challenge 4: training of robust and well-posed INNs

1. loss function $\mathcal{L}$
2. training data $(\widehat{u}_i, \widehat{y}_i)_{i=1}^N$

$$\min_{A,B,C,b,c} \quad \sum_{i=1}^{N} \mathcal{L}(\widehat{y}_i, Cx_i + c) \quad +\lambda \quad \text{Lip}_{u\to y}$$

$$x_i = \Phi(Ax_i + B\widehat{u}_i + b)$$

$$\mu_\infty(A) \leq \gamma,$$

- $\gamma < 1$ is a hyperparameter and $\lambda \geq 0$ is a regularization parameter
- training optimization problem is solved via SGD
- at each step of SGD, $x_i = \Phi(Ax_i + B\widehat{u}_i + b)$ is solved using the average-iterations

$$\mu_\infty(A) \leq \gamma \quad \iff \quad \exists T \text{ s.t. } A = T - \mathrm{diag}(|T|\mathbb{1}_n) + \gamma I_n.$$

State-of-the-art architectures:

Implicit Deep Learning (IDL)

- ℓ_∞-norm well-posedness and robustness analysis
- results in the green boxes

L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Y. Tsai. Implicit deep learning. *SIMODS*, 2019

Monotone operator equilibrium networks (MON)

- ℓ_2-norm well-posedness and robustness analysis
- $\mathrm{Lip}_\infty \leq \sqrt{r}\,\mathrm{Lip}_2$ with r size of the input

E. Winston and J. Z. Kolter. Monotone operator equilibrium networks. In *NeurIPS*, 2020

C. Pabbaraju, E. Winston, and J. Z. Kolter. Estimating Lipschitz constants of monotone deep equilibrium models. In *ICLR*, 2021

- MNIST dataset: 28×28 pixel handwritten digits between $0-9$, $60,000$ training images and $10,000$ test images.
- implicit neural network order: $n = 100$ and $\gamma = 0.95$
- loss function: cross entropy

Test error vs Lipschitz constant on MNIST handwritten digits

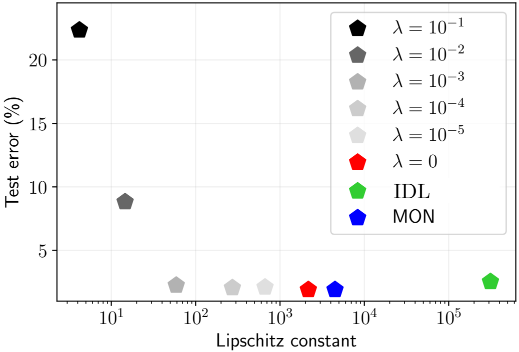

Improvements:

- $(\lambda = 0)$: two orders of magnitude wrt. IDL and wrt. MON
- $(\lambda = 10^{-3})$: three orders of magnitude wrt. IDL and one order of magnitude wrt. MON
- $(\lambda = 10^{-2})$: four orders of magnitude wrt. IDL and two orders of magnitude wrt. MON

- Pareto-optimal curve

- perturbation: inversion attack $u_{\mathrm{adv}} = u + \epsilon\, \mathrm{sign}(\frac{1}{2}\mathbb{1}_{784} - u)$

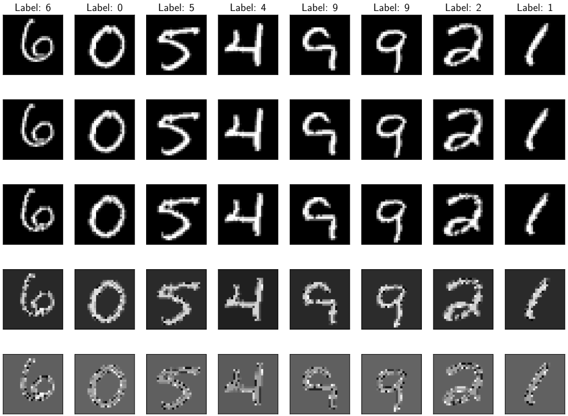

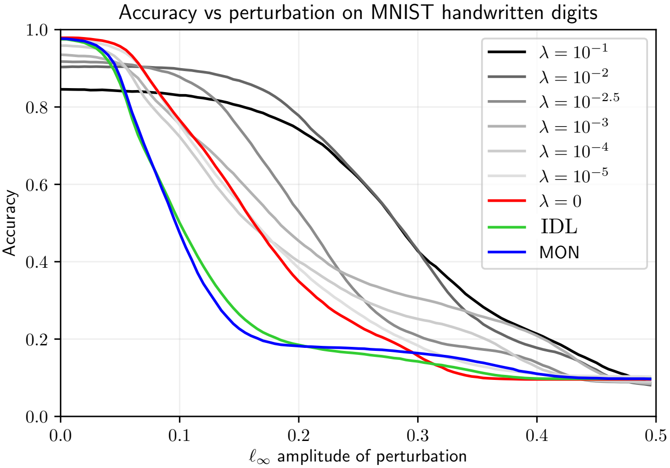

- $(\lambda = 0)$: improved robustness than IDL and MON
- $(\lambda > 0)$: improved robustness at sizable perturbations but losing some percentage accuracy in clean performance

Tradeoff between **clean performance** and **robustness**

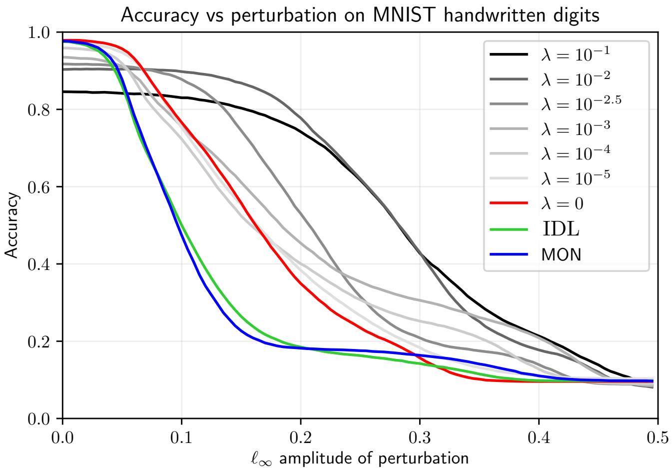

- $(\lambda = 0)$: improved robustness than IDL and MON
- $(\lambda > 0)$: improved robustness at sizable perturbations but losing some percentage accuracy in clean performance

Tradeoff between **clean performance** and **robustness**

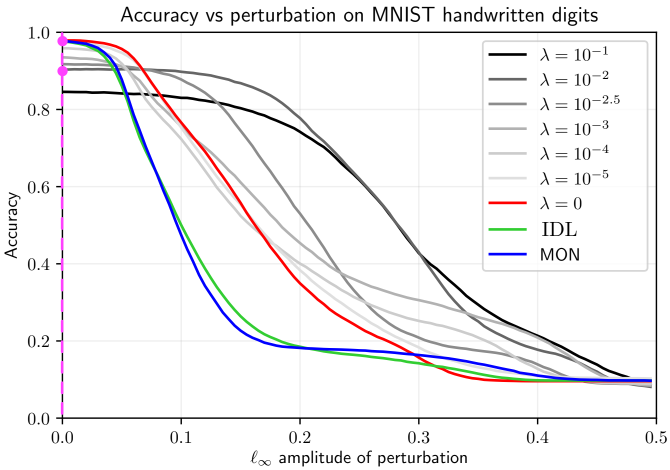

- $(\lambda = 0)$: improved robustness than IDL and MON
- $(\lambda > 0)$: improved robustness at sizable perturbations but losing some percentage accuracy in clean performance

Tradeoff between **clean performance** and **robustness**

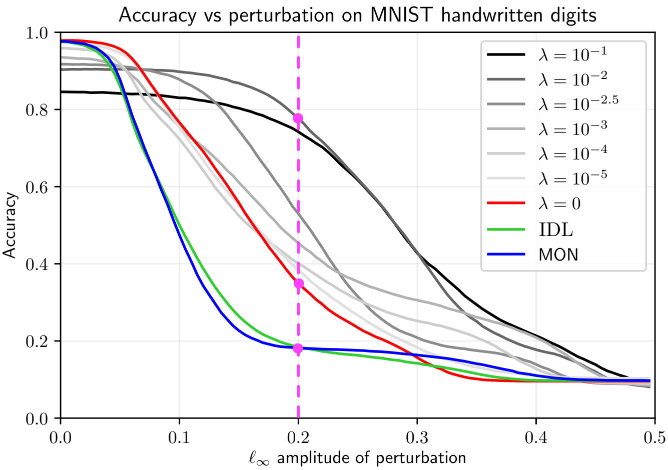

- $(\lambda = 0)$: improved robustness than IDL and MON
- $(\lambda > 0)$: improved robustness at sizable perturbations but losing some percentage accuracy in clean performance

Tradeoff between **clean performance** and **robustness**

Presentation outline

- Robustness of neural networks
  - implicit neural networks
  - well-posedness using contraction theory
  - robustness via Lipschitz bounds
- **Resilience of dynamic flow networks**
  - dynamic flow networks
  - dichotomy in asymptotic behavior
  - robustness of transportation networks to failures

A network of interconnected compartments $\mathcal{L}$
density of the ith compartment is x_i:

$$\dot{x}_i = F_i^{\text{in}}(x) - F_i^{\text{out}}(x) := f_i(x),$$

- $F_i^{\text{in}}(x)$ inflow to compartment i
- $F_i^{\text{out}}(x)$ outflow from compartment i

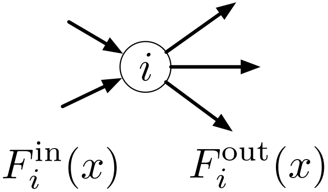

Power networks

Transportation networks

Water networks

Dynamic flow problem

Problem statement

Real-world dynamic flow networks:

- have wide-spread disturbances
- have complex interconnections
- interact with decision-makers
- interact with stochastic environments

Robustness of flow networks with respect to transient failures (adversarial attacks or random disturbances)


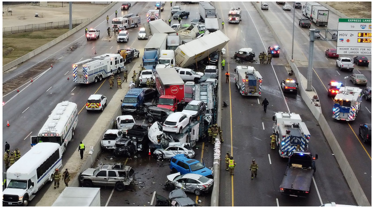


Transient stability of dynamic flow = region of attraction

For a dynamic flow network,

$$\textbf{conservation law} \qquad \mathbb{1}_n^\top f(x) = \mathbb{1}_n^\top f(y) \quad \forall x,y \qquad \Longleftrightarrow \qquad \mathbb{1}_n^\top Df(x) = 0 \quad \forall x$$

If f satisfies the conservation law then

1. $\mu_1(Df(x)) \geq 0$
2. if additionally f is cooperative, then $\mu_1(Df(x)) = 0$

Cooperative vector field

f is a cooperative vector field if $\frac{\partial f_i(x)}{\partial x_j} \geq 0$, for every x and every $i \neq j$

Dynamic flow networks with the conservation law are **not** contracting

Weakly-contracting systems

Definition and properties

$\dot{x} = \mathsf{G}(t, x)$ is weakly-contracting wrt $\|\cdot\|$:

$$\mu(D\mathsf{G}(t,x)) \leq 0, \qquad \text{for all } x, t$$

Dichotomy for weakly-contracting systems

For a weakly-contracting system $\dot{x} = \mathsf{G}(x)$, either

1. G has no equilibrium and every trajectory is unbounded, or
2. G has at least one equilibrium x^* and every trajectory is bounded,
   - if the norm $\|\cdot\|$ is a p-norm, $p \in \{1, \infty\}$ and f is piecewise real analytic, then every trajectory converges to the set of equilibria,

S. Jafarpour, P. Cisneros-Velarde, and F. Bullo. Weak and semi-contraction for network systems and diffusively-coupled oscillators. *IEEE Transactions on Automatic Control*, Feb. 2021a

$\dot{x} = \mathsf{G}(t, x)$ is weakly-contracting wrt $\|\cdot\|$:

$$\mu(D\mathsf{G}(t,x)) \leq 0, \qquad \text{for all } x, t$$

Contracting systems

For a contracting system $\dot{x} = \mathsf{G}(x)$, then

- G has a unique equilibrium x^* and every trajectory converges to it.

S. Jafarpour, P. Cisneros-Velarde, and F. Bullo. Weak and semi-contraction for network systems and diffusively-coupled oscillators. *IEEE Transactions on Automatic Control*, Feb. 2021a

$\dot{x} = \mathsf{G}(t, x)$ is weakly-contracting wrt $\|\cdot\|$:

$$\mu(D\mathsf{G}(t, x)) \leq 0, \qquad \text{for all } x, t$$

Dichotomy for weakly-contracting systems

For a weakly-contracting system $\dot{x} = \mathsf{G}(x)$, either

1. G has no equilibrium and every trajectory is unbounded, or
2. G has at least one equilibrium x^* and every trajectory is bounded,
   - if the norm $\|\cdot\|$ is a p-norm, $p \in \{1, \infty\}$ and f is piecewise real analytic, then every trajectory converges to the set of equilibria,

S. Jafarpour, P. Cisneros-Velarde, and F. Bullo. Weak and semi-contraction for network systems and diffusively-coupled oscillators. *IEEE Transactions on Automatic Control*, Feb. 2021a

Macroscopic models of traffic network

- segments of the roads are modeled as network of compartments $\mathcal{L}$
- vehicles flow from compartment to compartment
- density of link i is $x_i \in [0, \overline{x}_i]$ with jam density $\overline{x}_i$.

$$\dot{x}_i = F_i^{\mathsf{in}}(x) - F_i^{\mathsf{out}}(x) := f_i(x),$$

Demand and supply of link i

(demand)	$x_i \mapsto d_i(x_i)$	increasing
(supply)	$x_i \mapsto s_i(x_i)$	decreasing

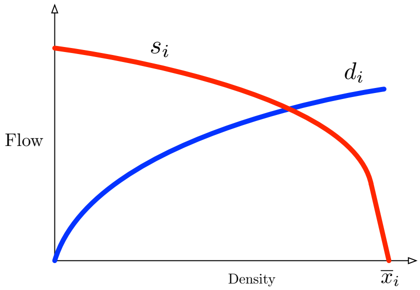

$$\mathcal{L} = \overbrace{\mathcal{R}}^{\text{input links}} \cup \overbrace{\mathcal{O}}^{\text{other links}}$$

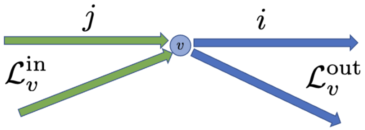

Input metering

$$F_i^{\text{in}}(x) = \min\{u_i, s_i(x_i)\}$$

Fixed routing ratios

$$F_i^{\text{in}}(x) = R_i^v \sum_{j\in\mathcal{L}_v^{\text{in}}} F_j^{\text{out}}(x_j)$$

Conservation of vehicles

$$\sum_{i\in\mathcal{L}_v^{\text{out}}} R_i^v = 1$$

First-In-First-Out (FIFO rule)

$$F_j^{\text{out}}(x) = \alpha^v(x) d_j(x_j),$$

$$\alpha^v(x) = \min_{l\in\mathcal{L}_v^{\text{out}}} \left\{1, \frac{s_l(x_l)}{R_l^v \sum_{k\in\mathcal{L}_v^{\text{in}}} d_k(x_k)}\right\}$$

- Link i is in **congestion** if $\alpha^v(x) < 1$
- Link i is in **free-flow** if $\alpha^v(x) = 1$

The transportation network $\dot{x} = f(x)$ satisfies the conservation law

$$\mathbb{1}_n^\top D f(x) = 0$$

Cooperative behavior and diverging junctions

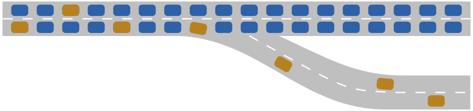

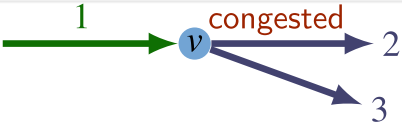

$$F_1^{\text{out}}(x) = \frac{1}{R_2^v} s_2(x_2),$$
$$F_2^{\text{in}}(x) = R_2^v F_1^{\text{out}}(x) = s_2(x_2),$$
$$F_3^{\text{in}}(x) = R_3^v F_1^{\text{out}}(x) = \frac{R_3^v}{R_2^v} s_2(x_2)$$

$$\frac{\partial f_3(x)}{\partial x_2} = \frac{R_3^v}{R_2^v} \frac{\partial s_2(x_2)}{\partial x_2} \leq 0$$

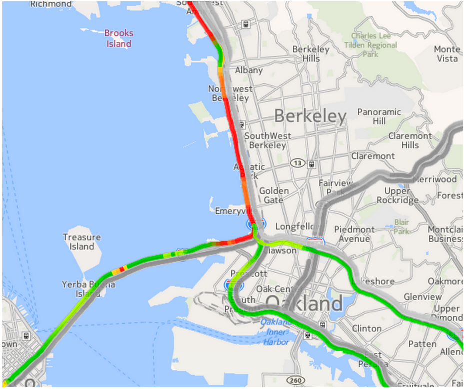

Diverging junctions: source of non-cooperative behavior in traffic flow

Cooperative domain

$$\mathcal{M} = \{x \in [\mathbb{0}_{|\mathcal{L}|}, \overline{x}] \mid f_i^{\text{out}}(x) = d_i(x_i), \text{ for } i \in \mathcal{L}_v^{\text{out}} \text{ with } v \text{ div. junction } \}.$$

- **Intuition:** in domain $\mathcal{M}$, the downstreams of diverging junctions are in free-flow

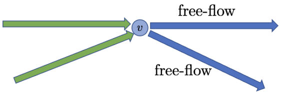

Theorem: traffic in cooperative domain

The transportation system $\dot{x} = f(x)$

1. is cooperative on the domain $\mathcal{M}$
2. satisfies $\mu_1(Df(x)) = 0$, for every $x \in \mathcal{M}$

S. Jafarpour and S. Coogan. Resilience of input metering in dynamic flow networks. In *American Control Conference*, 2022

The network routing matrix $[R_{\mathcal{O}}]_{kl} = \begin{cases} R^v_k, & l,k \in \mathcal{O} \\ 0 & \text{otherwise.} \end{cases}$ and the input routing matrix $[R_{\mathcal{R}}]_{kl} = \begin{cases} R^v_k, & k \in \mathcal{O}, l \in \mathcal{R} \\ 0 & \text{otherwise.} \end{cases}$.

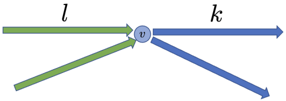

$$P = (I - R_{\mathcal{O}})^{-1} R_{\mathcal{R}}$$

Theorem: free-flow equilibrium point

Let the input metering u be strictly feasible and define $x^e_i(u) = d^{-1}_i(f^e_i(u))$ where

$$f^e_i(u) = \begin{cases} u_i, & i \in \mathcal{R}, \\ [Pu]_i, & i \in \mathcal{O}. \end{cases}$$

1. $x^e(u)$ is the unique equilibrium point in $\mathcal{M}$
2. at the equilibrium point $x^e(u)$, every link is in free-flow.

Theorem: region of attraction

Let $t \mapsto y(t)$ be the solution to

$$\dot{y} = h(y),$$
$$y(0) = \overline{x}$$

and let $t^* = \min\{t \in \mathbb{R}_{\geq 0} \mid y(t) \in \mathcal{M}\}$

1. $\lim_{t\to\infty} y(t) = x^e(u)$
2. every trajectory of f starting from $[\mathbb{0}_{|\mathcal{L}|}, y(t^*)]$ converges to $x^e(u)$.

Cooperative extension h:

$$\alpha^v(x) \mapsto \beta^v(x) = \begin{cases} \alpha^v(x) & x \in \mathcal{M}, \\ 1 & x \notin \mathcal{M}. \end{cases}$$

$$f(x) \mapsto h(x)$$

- h is cooperative on $[\mathbb{0}_{\mathcal{L}}, \overline{x}]$
- $f(x) = h(x)$ for every $x \in \mathcal{M}$

S. Jafarpour and S. Coogan. Resilience of input metering in dynamic flow networks. In *American Control Conference*, 2022

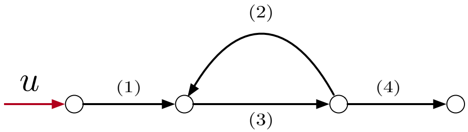

For $i \in \{1, \ldots, 4\}$

$$x_i \in [0, 30]$$
$$d_i(x_i) = \min\{15, x_i\}$$
$$s_i(x_i) = \min\{15, 30 - x_i\}$$

routing ratios

$$R_i^v = \begin{cases} \frac{1}{2} & i = 2, 4, \\ 1 & i = 1, 3 \end{cases}$$

For input metering $u = 5$

- free-flow equilibrium $x^e(u) = (5,\ 10,\ 5,\ 5)^\top$
- $t \mapsto y(t)$ trajectory of cooperative extension
- at $t^* = 24.29$ we have

$$y(t^*) = (20.97,\ \ 15,\ \ 22.5,\ \ 7.5)^\top \in \mathcal{M}$$

Thus $[\mathbb{0}_4, (20.97,\ \ 15,\ \ 22.5,\ \ 7.5)^\top]$ is region of attraction of $x^e(u)$.

Acknowledgment

Francesco Bullo
UCSB

Samuel Coogan
Georgia Tech

Alexander Davydov
UCSB

Pedro Cisneros-Velarde
UIUC

Anton Proskurnikov
Politecnico di Torino, Italy.

Thank you for your attention!

Backup slides

Feature of adversarial perturbations:

- exist for a large class of learning algorithms
- transfer across models (not always!)
- *not* caused by overfitting (empirical evidence)

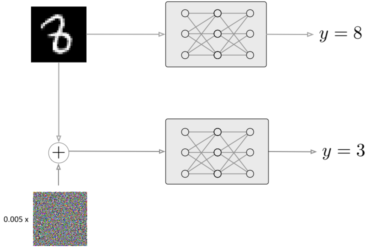

How to mitigate the effect of adversarial perturbations?

Adversarial training

- improve training using an attack
- easy to implement
- no provable guarantee

Robust optimization

- use over-approximation of the output
- hard to implement in training
- provide guarantees

- A large and flexible class of neural networks: includes feedforward neural networks

$$\begin{bmatrix} x^k \\ x^{k-1} \\ x^{\ell-2} \\ \vdots \\ x^2 \\ x^1 \end{bmatrix} = \sigma\left(\begin{bmatrix} 0 & A_{k-1} & 0 & \dots & 0 \\ 0 & 0 & A_{k-2} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & A_1 \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} x^k \\ x^{k-1} \\ x^{k-2} \\ \vdots \\ x^2 \\ x^1 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ A_0 \end{bmatrix} u \right),$$

$$y = \begin{bmatrix} A_k & 0 & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} x^k \\ x^{k-1} \\ x^{\ell-2} \\ \vdots \\ x^2 \\ x^1 \end{bmatrix}$$

1. **Origins**

S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *NeurIPS*, 2019

L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Y. Tsai. Implicit deep learning. *SIMODS*, 2019

A. Kag, Z. Zhang, and V. Saligrama. RNNs incrementally evolving on an equilibrium manifold: A panacea for vanishing and exploding gradients? In *ICLR*, 2020

2. **Monotone operator theory**

E. Winston and J. Z. Kolter. Monotone operator equilibrium networks. In *NeurIPS*, 2020 M. Revay,

R. Wang, and I. R. Manchester. Lipschitz bounded equilibrium networks. 2020. URL https://arxiv.org/abs/2010.01732

3. **Convergence**

K. Kawaguchi. On the theory of implicit deep learning: Global convergence with implicit layers. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=p-NZIuwqhI4 S. W. Fung, H. Heaton, Q. Li, D. McKenzie,

S. Osher, and W. Yin. Fixed point networks: Implicit depth models with Jacobian-free backprop, 2021. URL https://arxiv.org/abs/2103.12803. ArXiv e-print