

Continuous Models

Numerical Methods for Deep Learning

Lars Ruthotto

Departments of Mathematics and Computer Science, Emory University

Course Overview

Course Overview

- ▶ Part 1: Linear Models
 1. Introduction and Applications
 2. Linear Models: Least-Squares and Logistic Regression
- ▶ Part 2: Neural Networks
 1. Introduction to Nonlinear Models
 2. Parametric Models, Convolutions
 3. Single Layer Neural Networks
 4. Training Algorithms for Single Layer Neural Networks
 5. Neural Networks and Residual Neural Networks (ResNets)
- ▶ Part 3: Neural Networks as Differential Equations
 1. ResNets as ODEs
 2. Residual CNNs and their relation to PDEs

Minicourse Summary: Numerical Methods for DL

Learning objective: a look under the hood.

- ▶ Part 1 (linear models): foundations for deep networks
- ▶ Part 2 (nonlinear models): feature extraction, challenges in optimization
- ▶ Part 3 (deep continuous models): hard to optimize/analyze. new insights from ODEs/PDEs

We have only (or not even) scratched the surface. Find more challenges in:

- ▶ recurrent neural networks (time sequence data \rightsquigarrow ODEs)
- ▶ generative models and (variational) auto-encoder
- ▶ reinforcement learning