

# Never miss a bus again with this one crazy trick

Cooper Sloan



# Background

- Undergrad: MIT, BS in Computer Science
- Masters: MIT, MEng in Artificial Intelligence

# Overview

- Problem
- Risks/Challenges
- Implementation
- Final Design
- Takeaways

# Problem

- Bus schedules are unreliable
- Complicated interactions
- Deep learning may have the answer

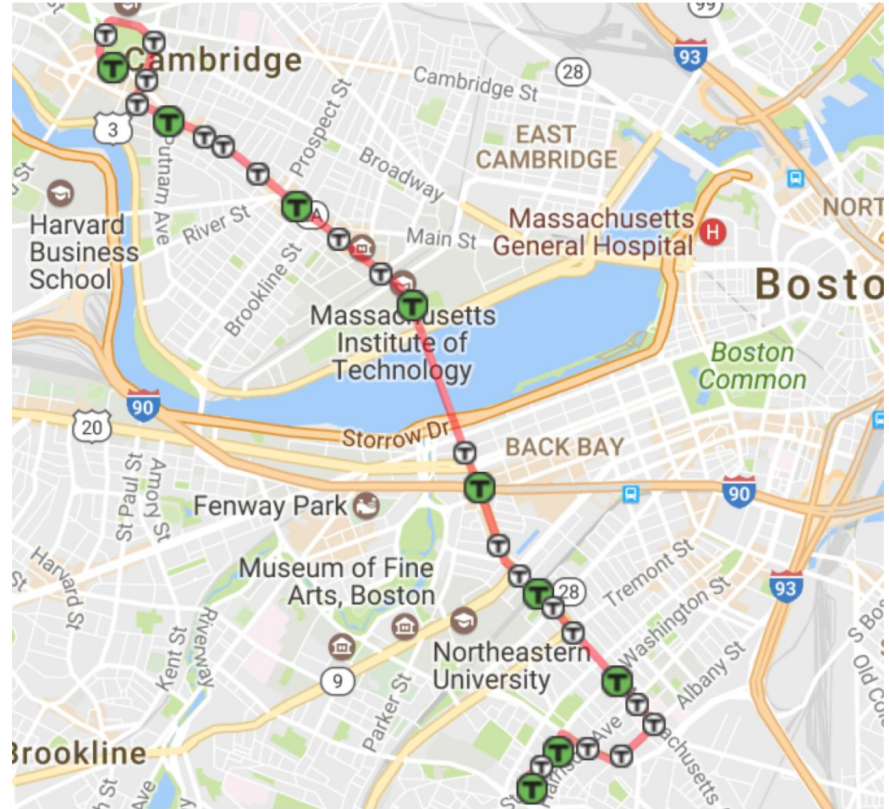


Problem cont.



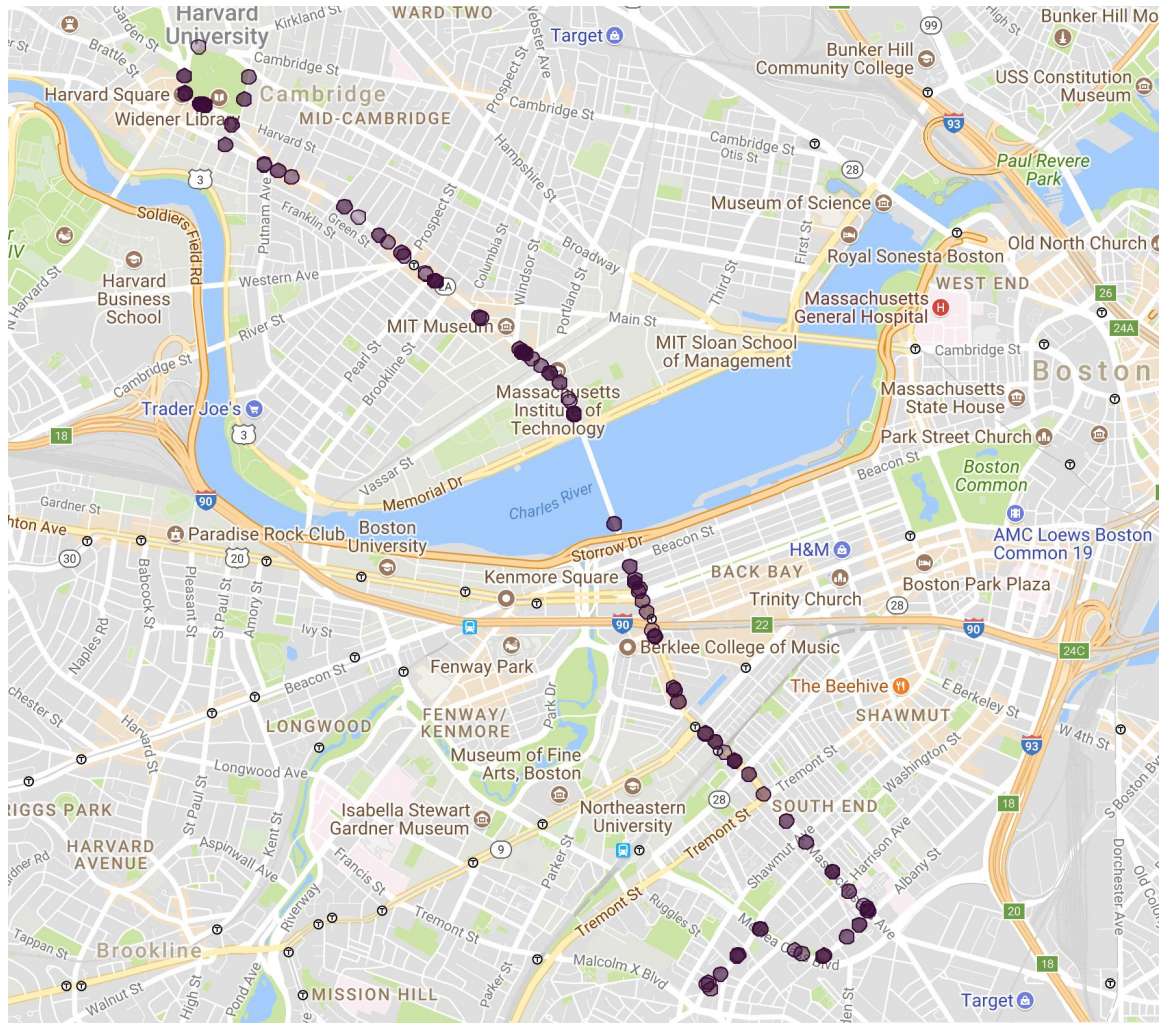
# Challenge 1: Feature Selection

- Clumping
- Dwell Time
- Travel Time
- Schedule Adherence
- Temporal Features



## Challenge 2: Noisy/Incomplete Data

- Urban Valleys
- 1: Naive Approach
- 2: Interpolation





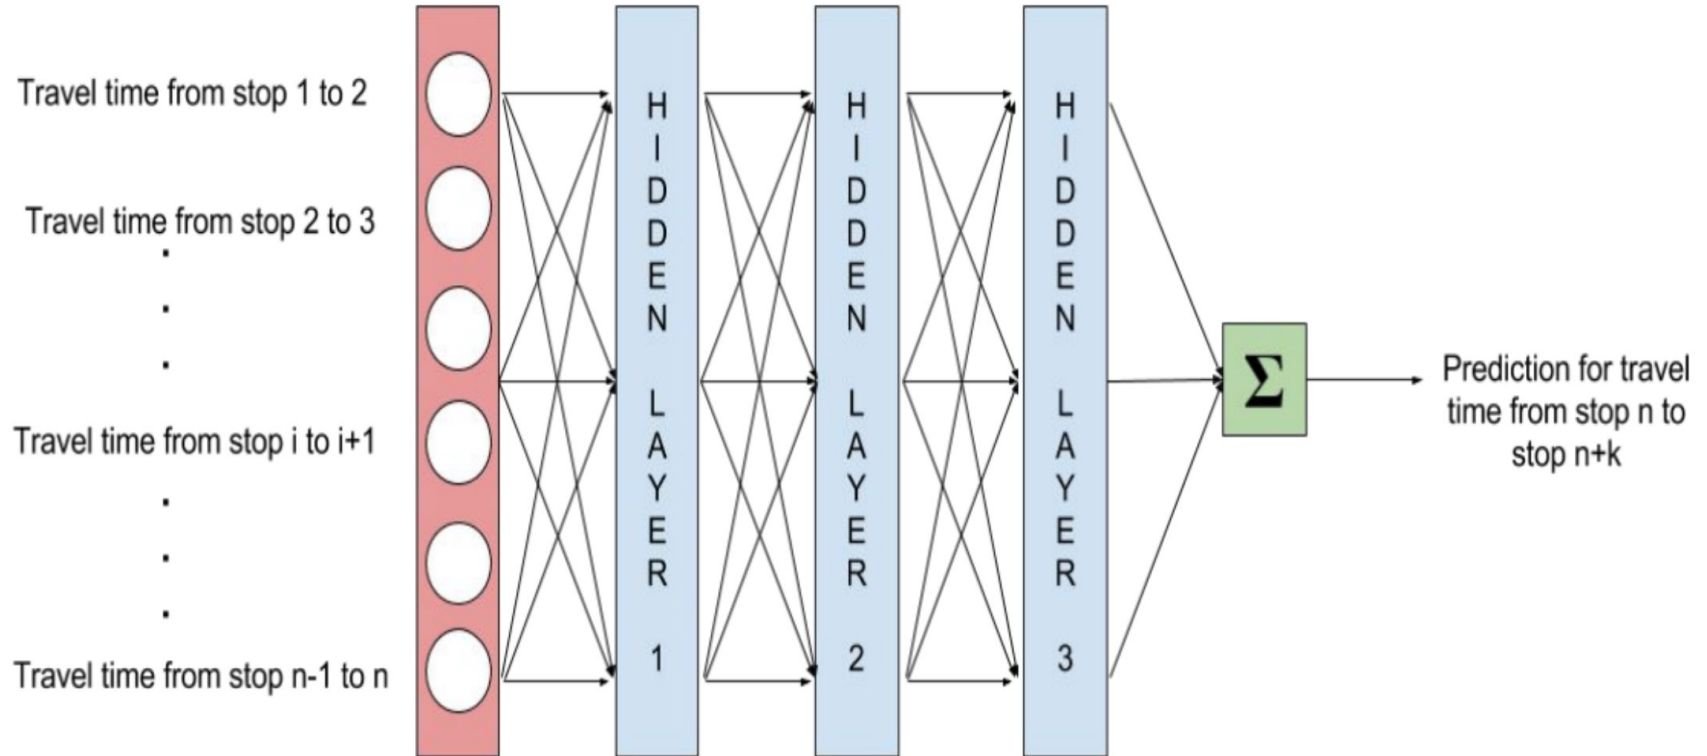
## Challenge 3: Model Design

- Train/test split
- Data partitioning
- Routes
- Architecture
- Evaluation

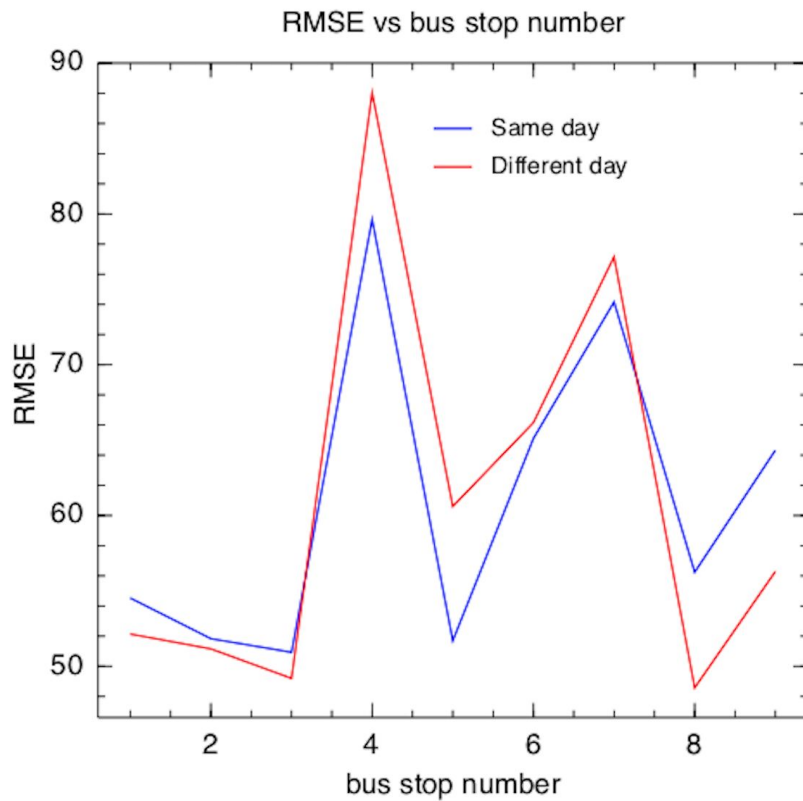
# Iteration

- Overfitting
- Signal to noise
- Temporal data

# Architecture



# Results



# Comparison

- Knet
  - Pytorch
  - Imperative
  - Easy Debugging
  - More expressiveness
- Mocha
  - Tensorflow/Keras
  - Declarative
  - Easy to write (usually...)
  - Hard Debugging
  - Good for simple models
  - Less expressive

# Summary

- Predict bus arrival times using neural networks
- GPS data from MBTA
- Use travel time between stops as features
- 3 hidden layers + RNN