ESnet
ENERGY SCIENCES NETWORK

mkiran@es.net

# AI for Networking: an Engineering Perspective

**Mariam Kiran**
Research Scientist
*Energy Sciences Network*
*Lawrence Berkeley National Laboratory*

INDIS Workshop @SC20

12th Nov 2020

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

*Funding Support: DOE ASCR, LBNL*

# Outline

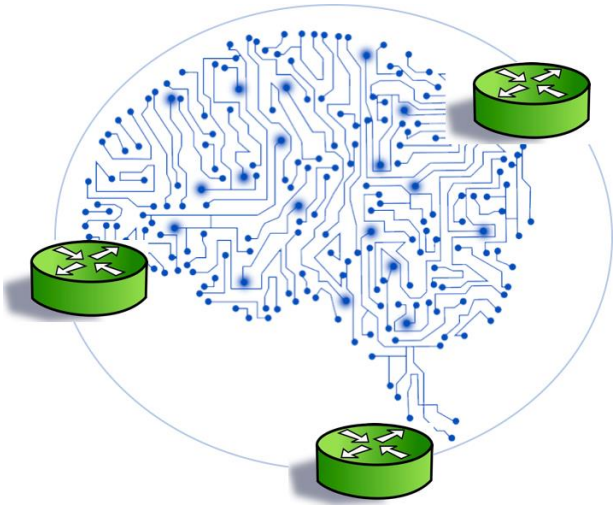

- My Background:
  - Working with Science R&E Networks
  - Focused on deep learning solutions for a ‘self-learning’ network
- Motivation: AI applications for Network Operations
- Quick ML/AI overview
  - Supervised, unsupervised and reinforcement learning
- Mathematical Challenges of AI
- Software: NetPredict (congestion patterns), DeepRoute (traffic engineering)
- Summary

ESnet

# Why use AI?

ESnet

- Networks are growing and becoming more complex
- Multiple vendors and their hardware
- Growing demand for reliable connectivity (science and home networks)
- Help engineers focus on innovation rather than mundane everyday tasks

# Why use AI?

ESnet

# AI Proposals for Networks

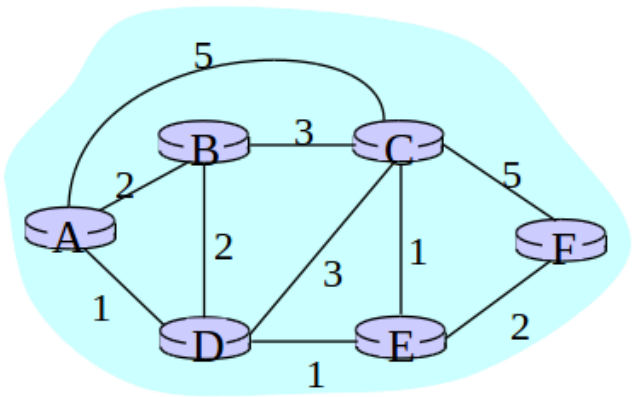


*Improving Traffic Engineering*

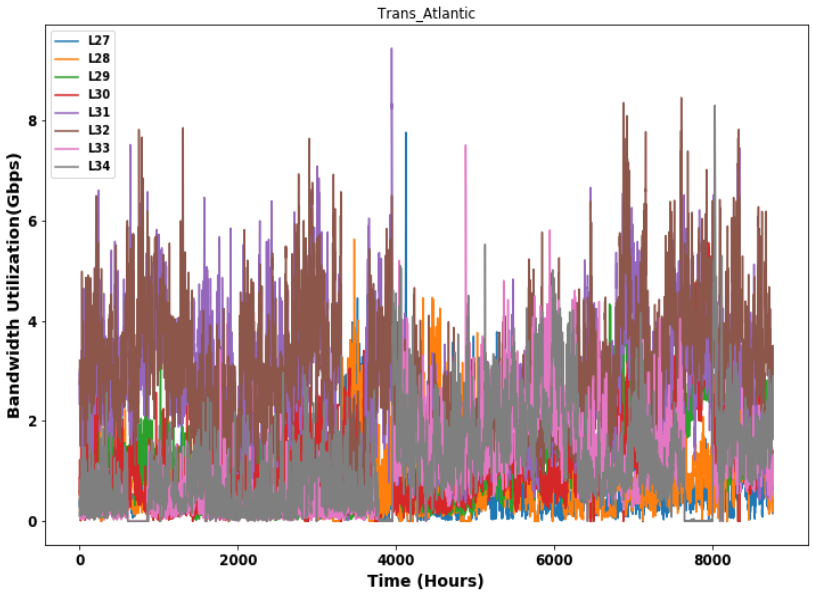


*Predicting Network Surges*

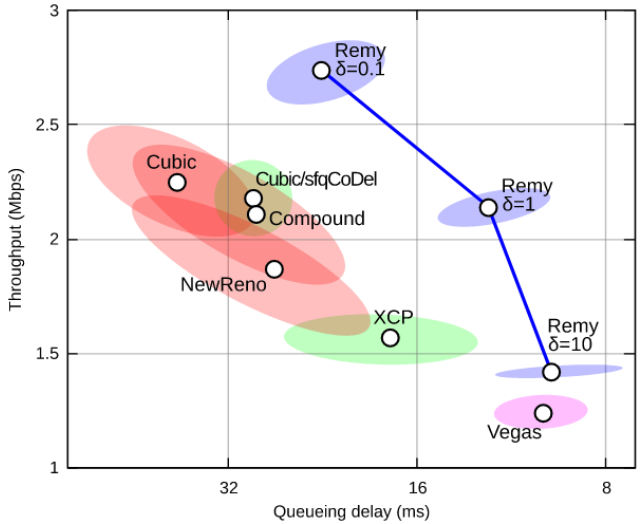


*Improving the TCP Congestion Model*

- *Security Anomalies*
- *Elephant versus mice flows*
- *And more …*

ESnet

# Science WANs Challenge 1: Resources are often underutilized

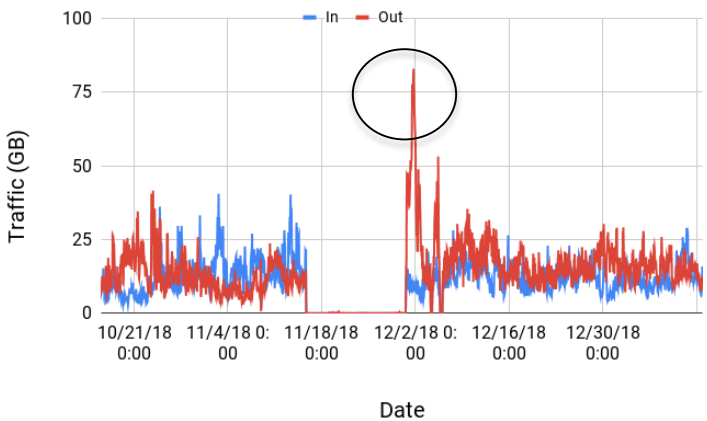


- Traffic is highly variable and 'bursty' (big versus small transfers)
- Flow congestion leads to packet loss
- Networks are designed to be resilient
- Under-utilized resources

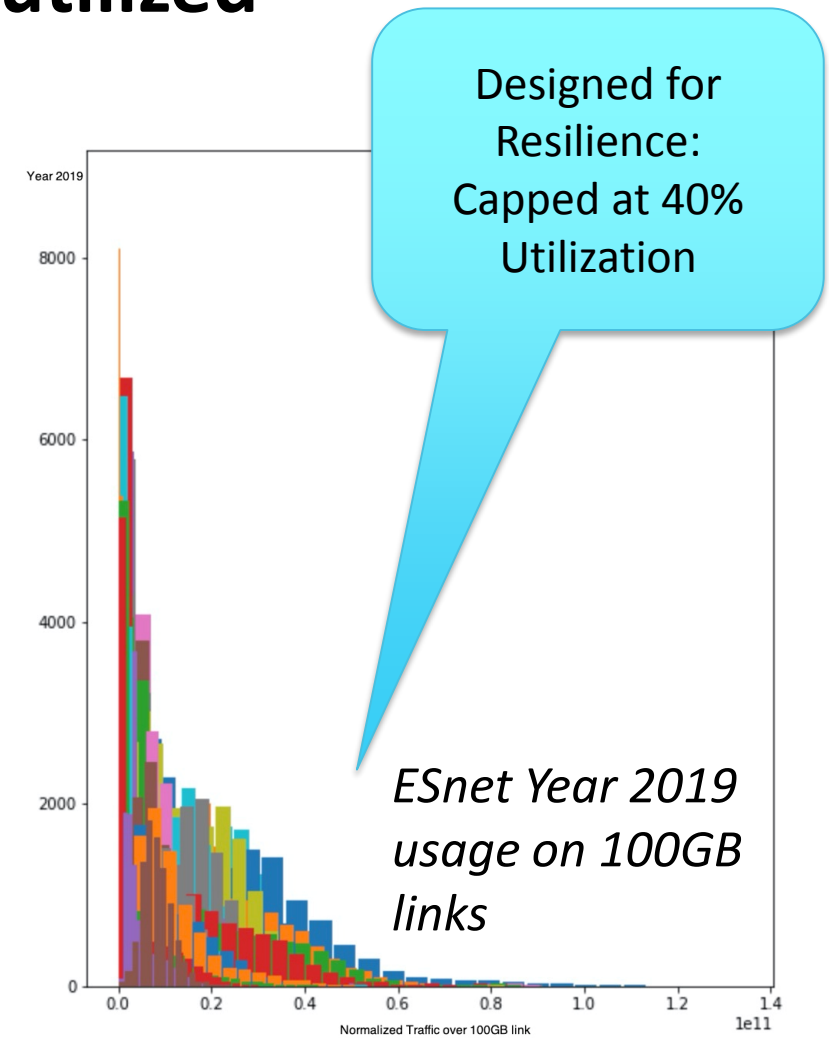


ESnet

# Science WANs Challenge 2: TCP Performance can be Fragile

- Science uses Transmission Control Protocol (TCP) to send data
- 95% of science traffic uses TCP

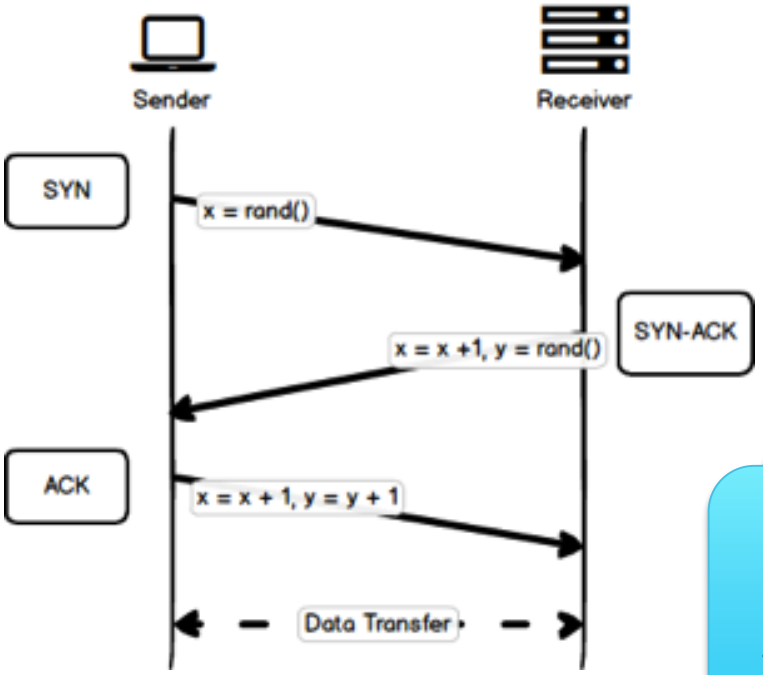


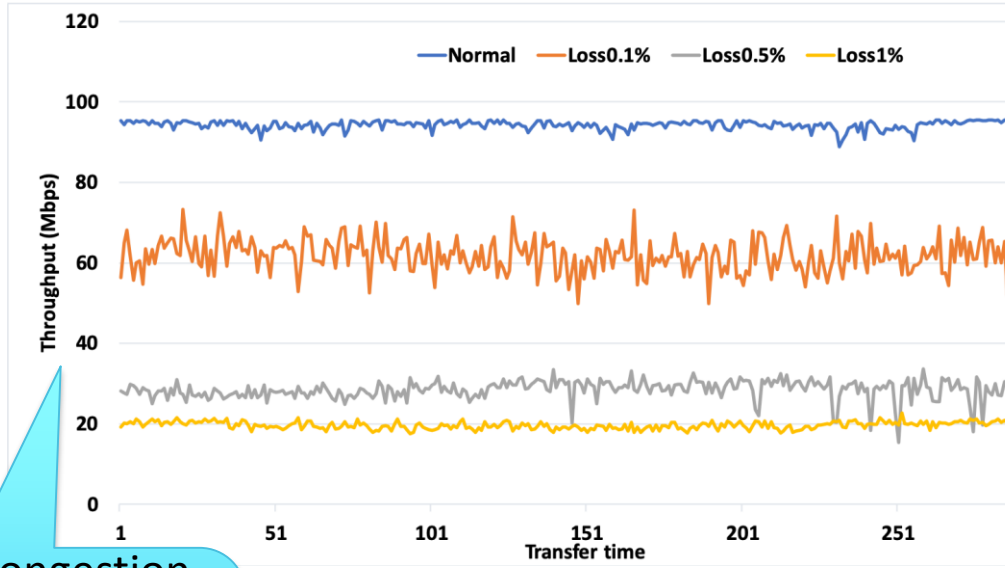


Congestion causes packet loss that affects throughput

ESnet

# Science WANs Challenge 3: TCP Congestion Control for Performance

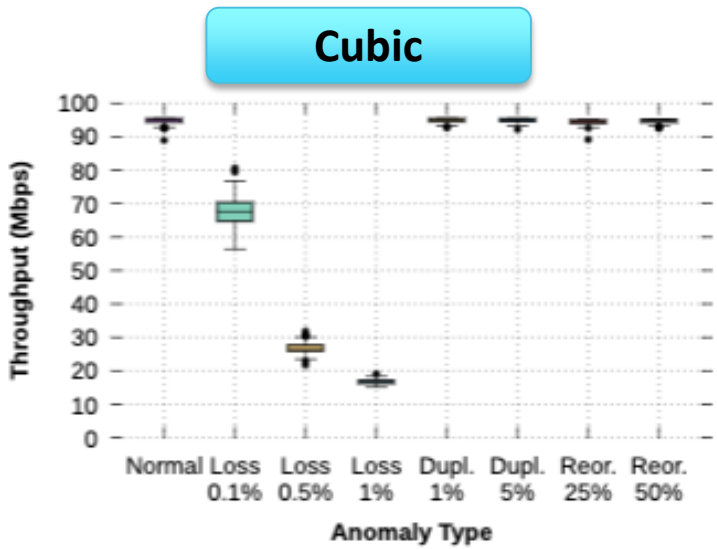


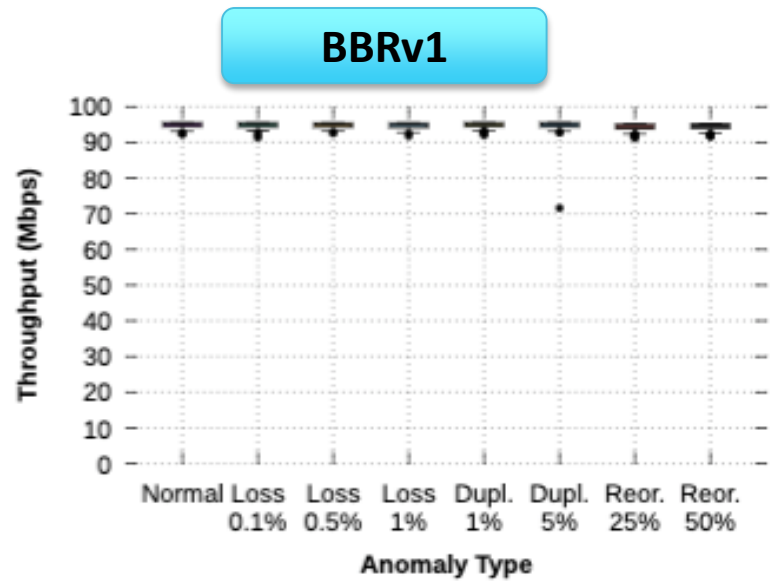


- Each TCP uses complex algorithms to optimize performance (slow down or quicken data sending)*
- Examples include video buffering on Youtube, VOIP (zoom) or big file transfers in few hours (Astronomy/physics from LHC)

*Building Classifiers for TCP congestion algorithms, Machine Learning Journal

# The Fusion Challenge : Network + AI/ML

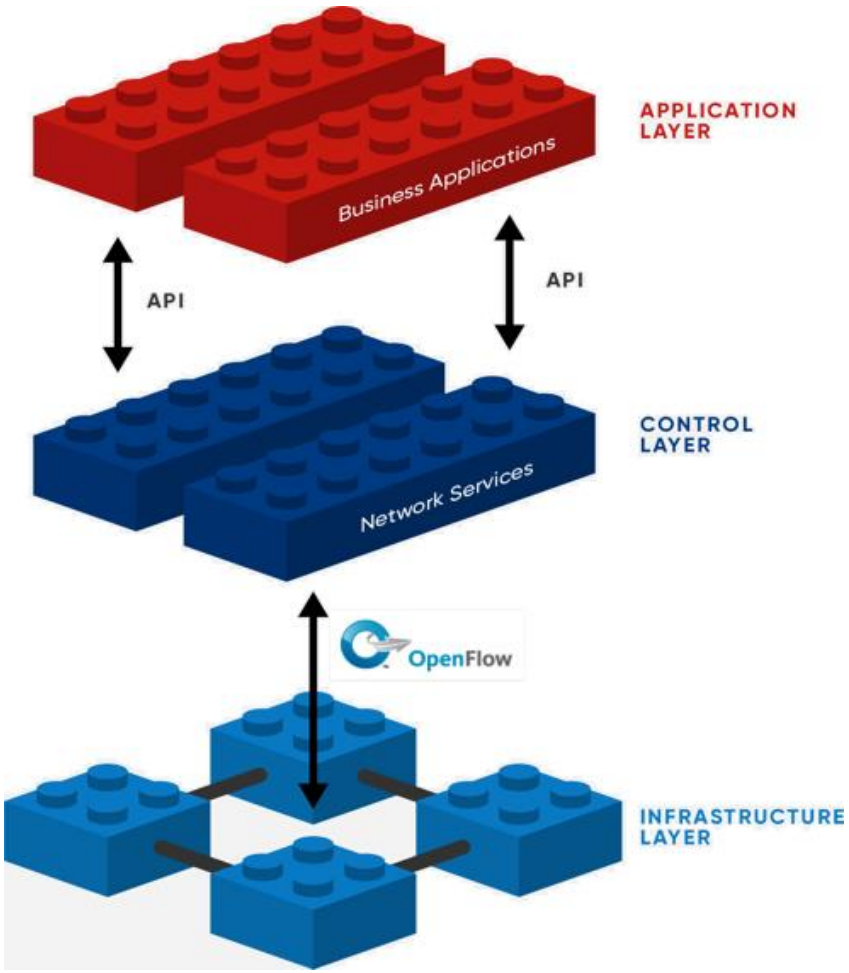


- Deep Learning trend
  - Image recognition e.g. Cats 2011
  - Self-playing game e.g. AlphaGo 2016
- Hardware acceleration
  - GPU advances
  - FPGAs
- Industry and Academic Efforts
  - Smart NICs e.g. Barefoot
  - AI @ Control Plane e.g. juniper, cisco)
  - AI enabled TCP e.g. [1, 2]
  - Traffic patterns e.g. [3]

[1] https://web.mit.edu/remy/
[2] Ghosal et al. Model Predictive Congestion Control for TCP Endpoints
[3] Lakhina et al. Diagnosing network-wide traffic anomalies

# Quick Background of AI

# This talk focuses on Deep Learning (DL) compared to general Machine Learning (ML)

*Trained to recognize cats,*
*"object identification"*

CAT

*Trained to play games,*

*"Best strategies for winning game"*

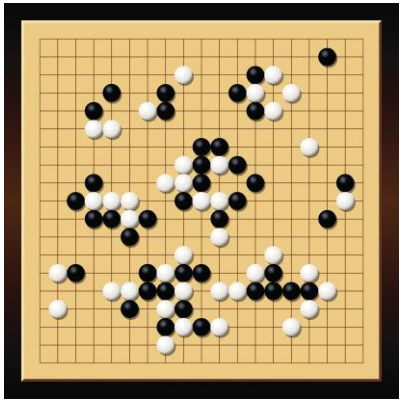

**Deep learning (neural networks - NN) introduces 'data-driven learning' to build bespoke solutions**

ESnet

# Supervised versus Unsupervised ML

## When Labelled Data is Available

Classification of good versus bad TCP flows with labeled data

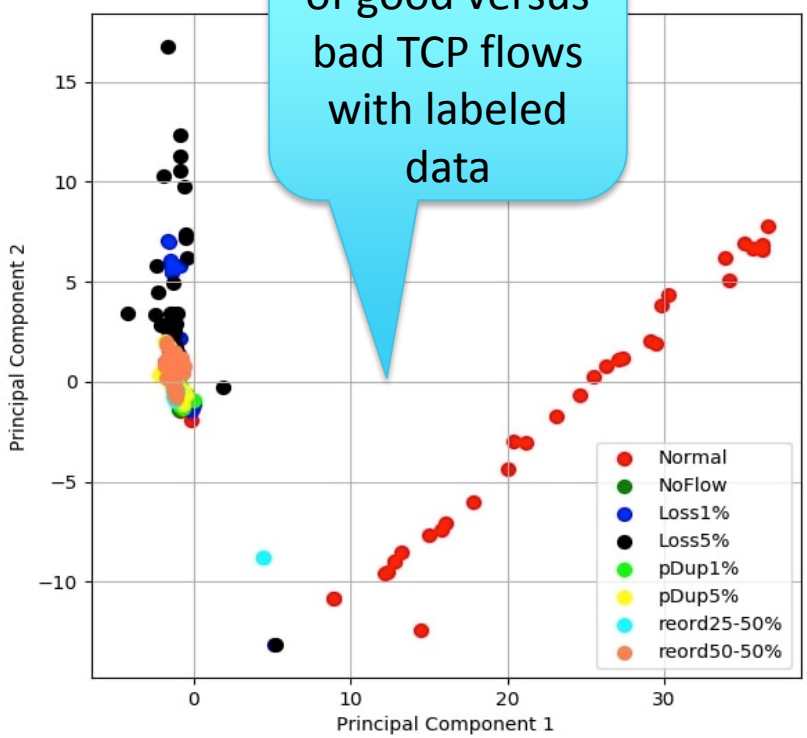


## When NO Labelled data in Available

Learn underlying rules in the data

Classification of abnormal behavior in unlabeled data* (Covid effect on NERSC)

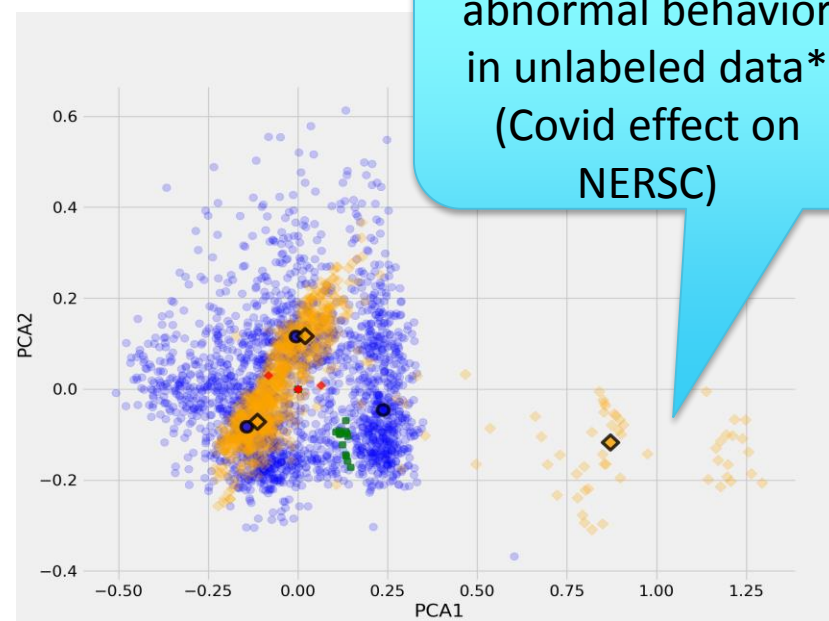


*Unsupervised Traffic Analysis of Data Centers, SMC 2020

# Reinforcement learning (RL)

<u>When NO data is available</u>

Learn via trial and error

- Agent views the state and chooses an action
- Gets a reward
- Over time, learns optimal actions for best rewards

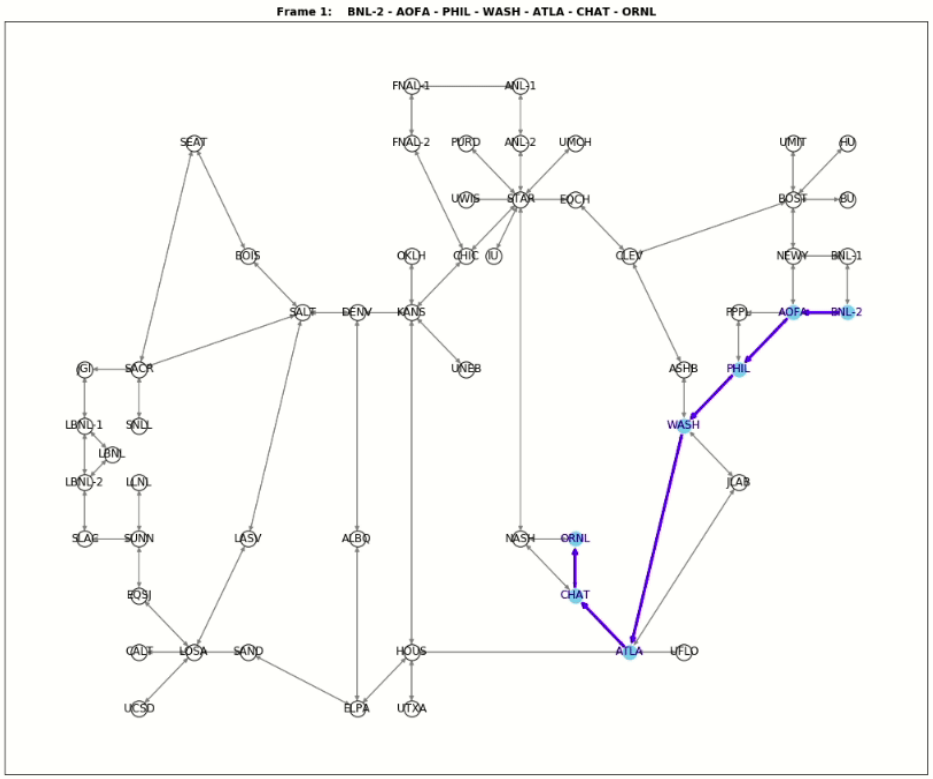


*Self-discovery: Learning flow performance over various paths*

# Mathematics of AI

# AI has Much Randomness

- AI researchers still exploring:
  - Explainable AI
  - Uncertainty quantification
- Choosing appropriate ML algorithms – different accuracies
- Hyperparameter tuning can dramatically improve the accuracy
- AI training needs large training data, often gets stuck in local optima
- Optimization functions can affect how well the NN learns

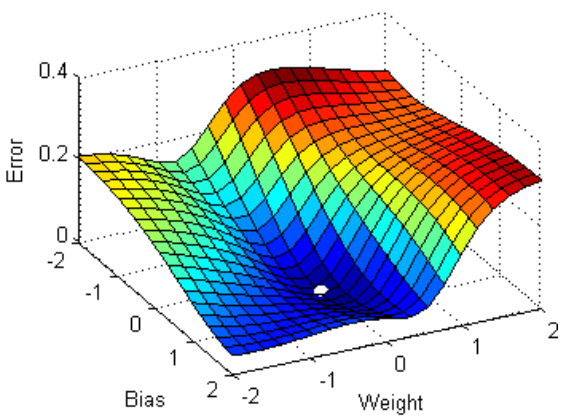


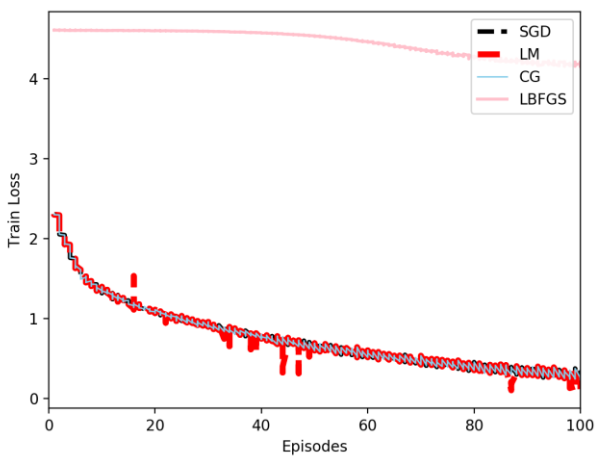


*Levemberg-Marquadt (LM) can lead to better solutions than Stochastic Gradient Descent (SGD)**

*Do optimization methods in deep learning applications matter? arXiv:2002.12642

# AI is Power Hungry

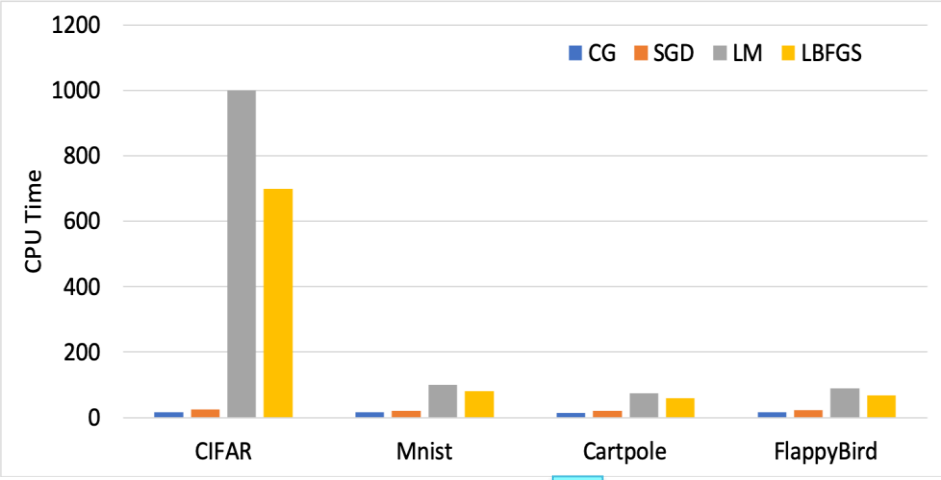


- Access to advanced hardware matters for training complex models
- For (training) Learning-at-the-edge, using GPUs and FPGAs

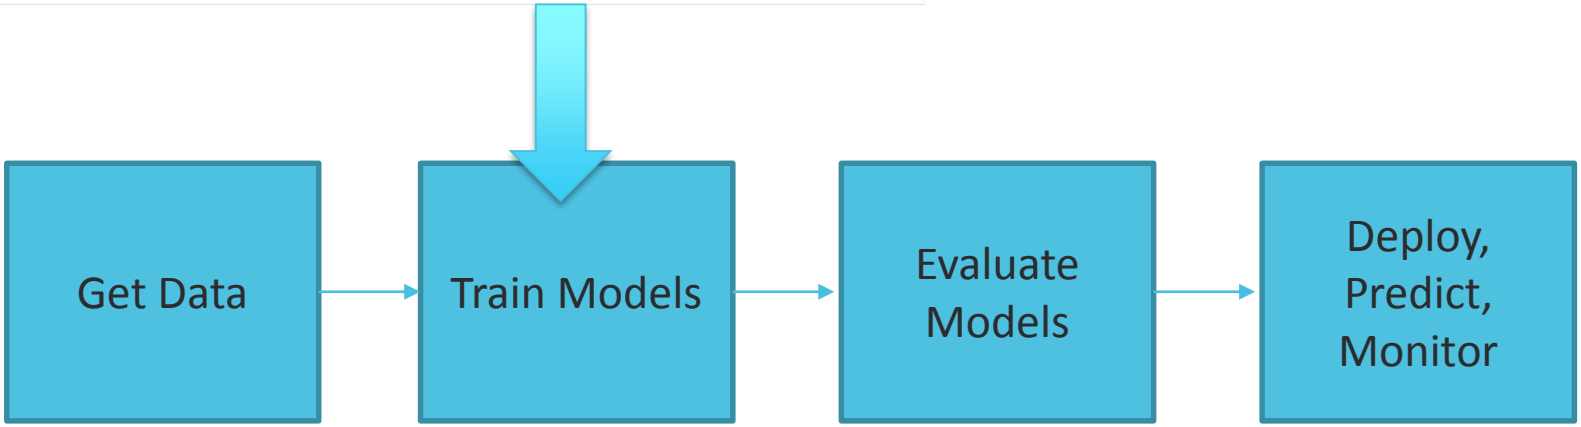


* Deploy Models into Production

ESnet

# Accessing the ‘Right’ Data from Hardware

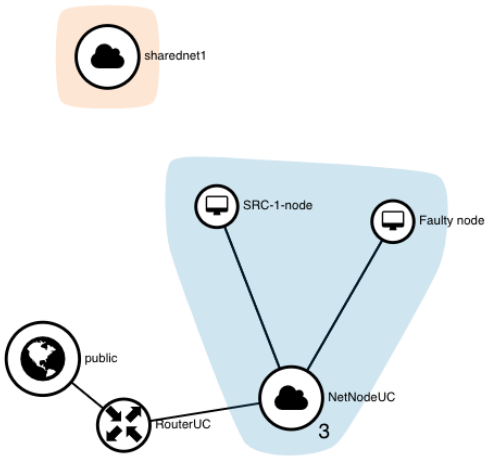


- Testbeds and Hardware capability available (e.g. kind of switches, sflow-enabled)
- Most monitoring can introduce ‘noise’ (e.g. UDP, ICMP packets)*
- Available training data for research

*NetGraf : A collaborative Network Monitoring Stack for Network Experimental Testbeds SC20 poster

*R. Cziva's (ESnet) work

# Handling Streaming AI for Massive Network Data

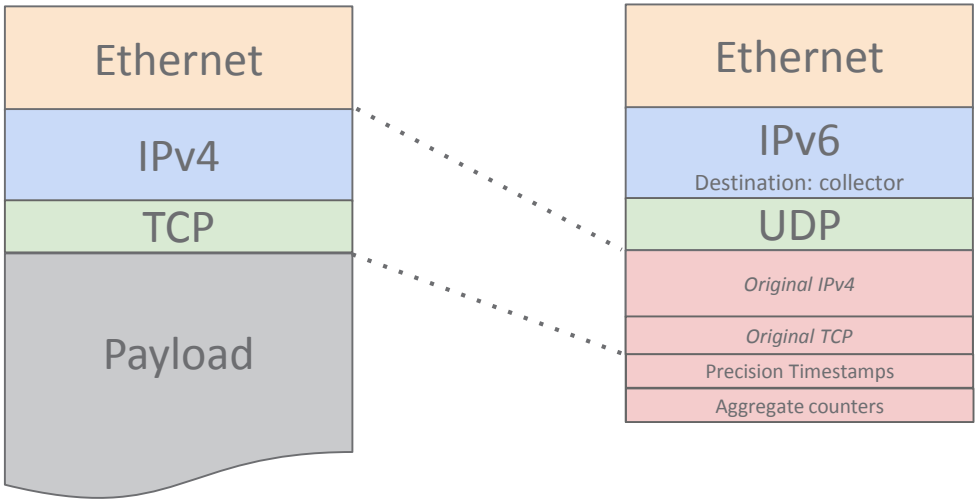


*Original packet of a TCP flow* → *High-Touch Telemetry Packet*

| Packet size | Rate | Telemetry PPS | Telemetry Rate |
|---|---|---|---|
| 1500B | 10Gb/s | 812K | 1,079Mb/s |
| 1500B | 100Gb/s | 8,127K | 10,790Mb/s |
| 9000B | 10Gb/s | 138K | 183Mb/s |
| 9000B | 100Gb/s | 1,383K | 1,833Mb/s |

- Leveraging data tools e.g. splunk
- Building fast streaming solutions
- Security and access
- Build streaming AI solutions becomes a software challenge

ESnet

# Learning Features: Correlation Patterns

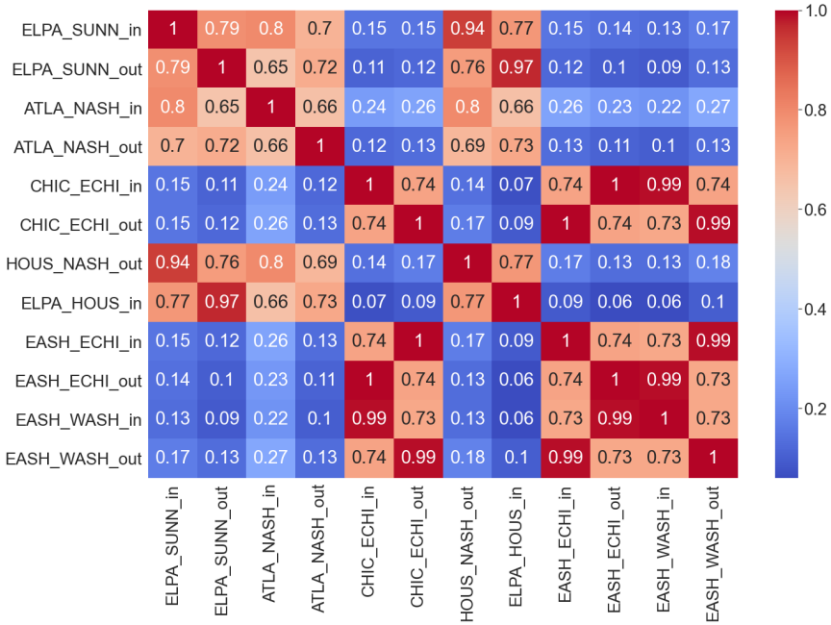


$$\text{Correlation} = \frac{Cov(x,y)}{\sigma x * \sigma y}$$

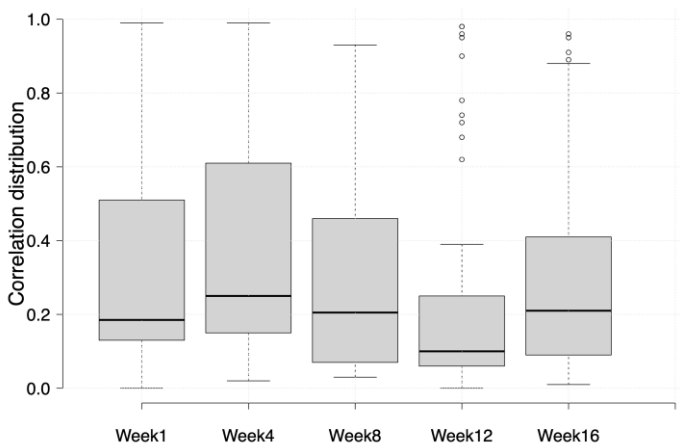


*First 16 weeks of year 2019 and shows network is extremely dynamic*

# Learning Communities: Forming and Dispersing

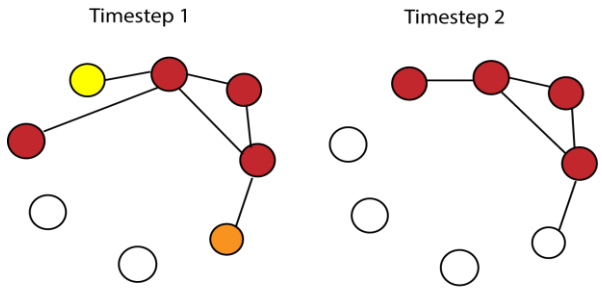


$$C_i(t_k, tk_{+1}) \propto \left|\left|S_i \epsilon (t_i, Ei, Ni)\right|\right|$$

Clustering Coefficient with temporal perturbation recording edge-change

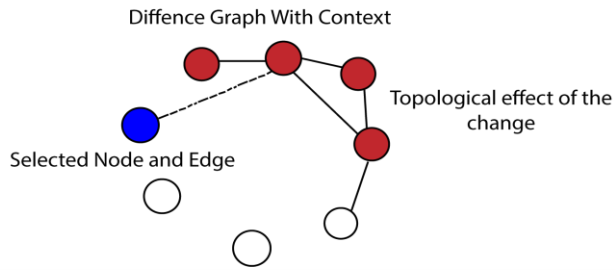


Learn 'Day and Night' Traffic patterns

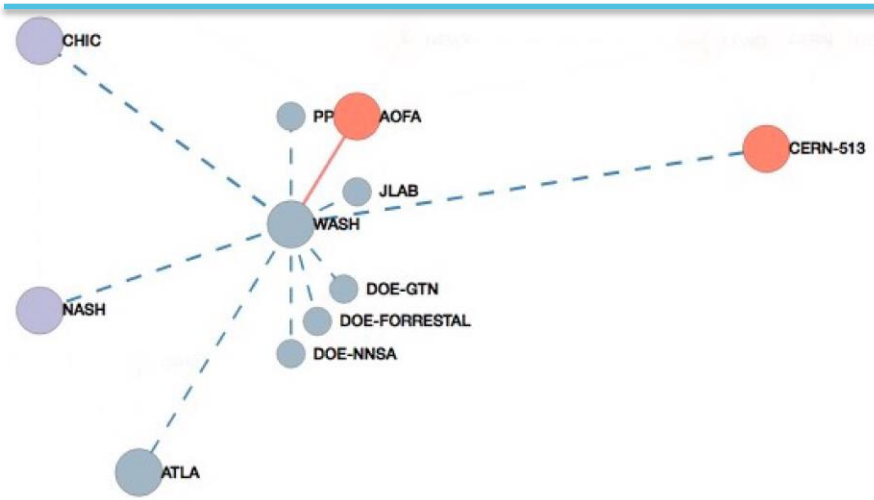


TimeStep 21- 22

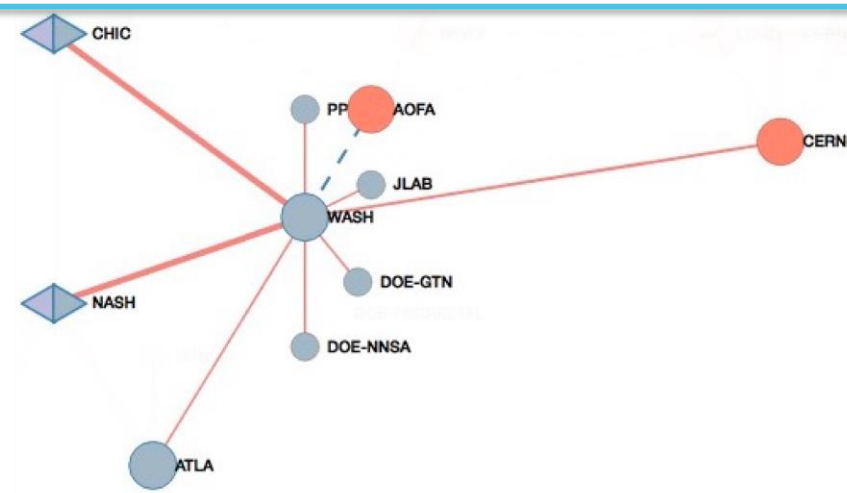


TimeStep 22- 23

# Graph Neural Networks for Traffic Prediction

*+ features*

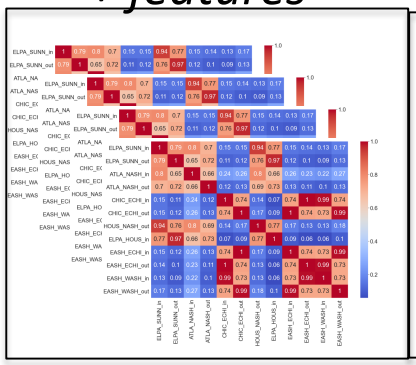

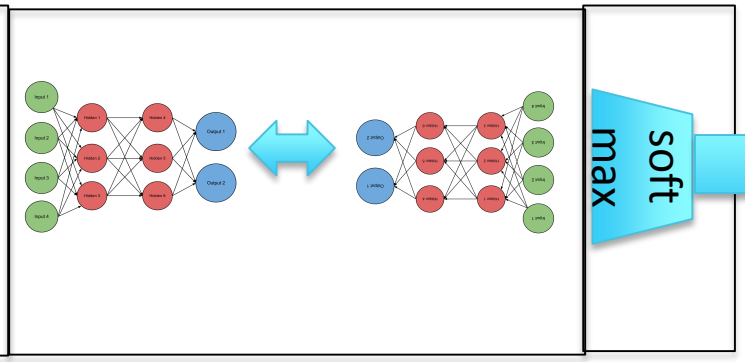

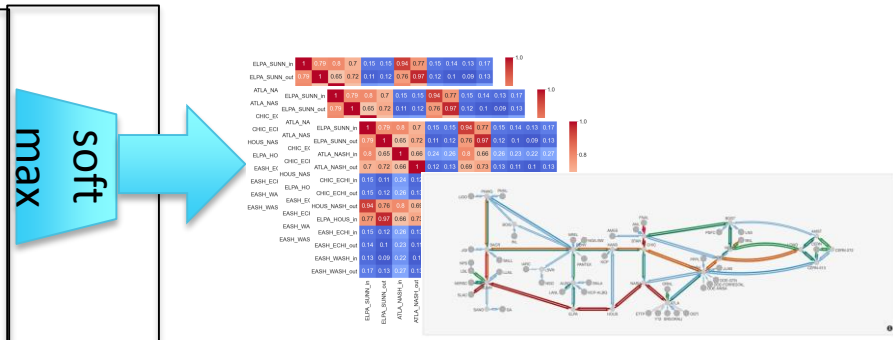


*Input*
*W=36 hrs*

*Multiple layers of bi-directional recurrent NNs*

*Output*
*W=24 hrs*

- **Benchmarking** Prediction algorithms
- No need to learn long term traffic patterns (short term is enough)
- Deployed for internal testing with real-time data

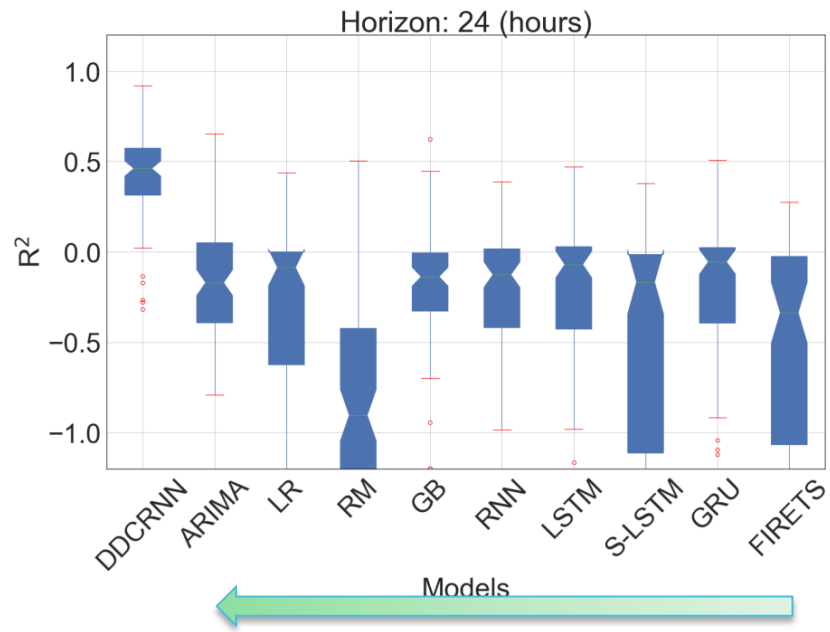


*Data-driven Learning to Predict WAN Network Traffic, SNTA'20
*Dynamic Graph Neural Network for Traffic Forecasting in WAN, IEEE BigData20

# Model Development Cycle

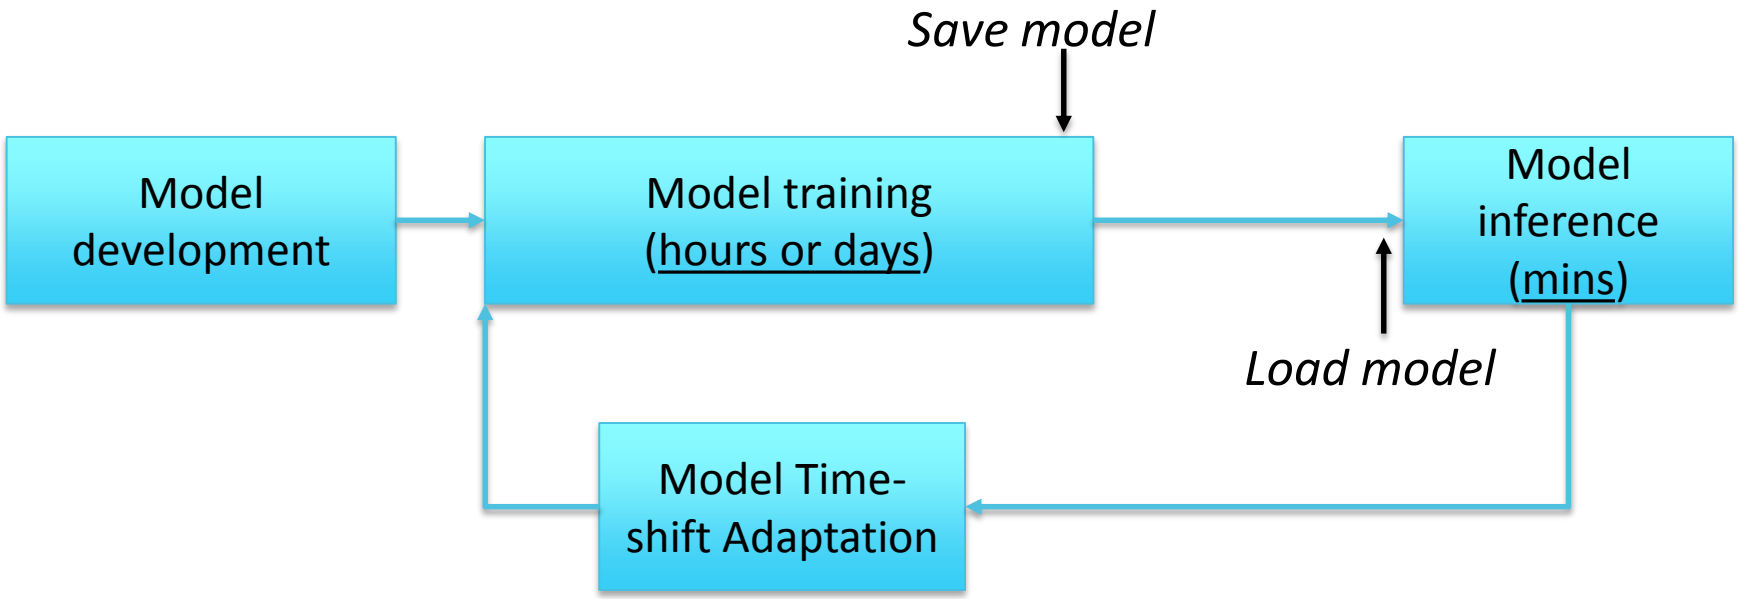


- Use HPC to optimize model training and "save model"
  - Costly and time consuming
- Model shift: Testing to update where the model is not performing well and needs a rethink

ESnet

# NetPredict: Predict traffic 1 week ahead

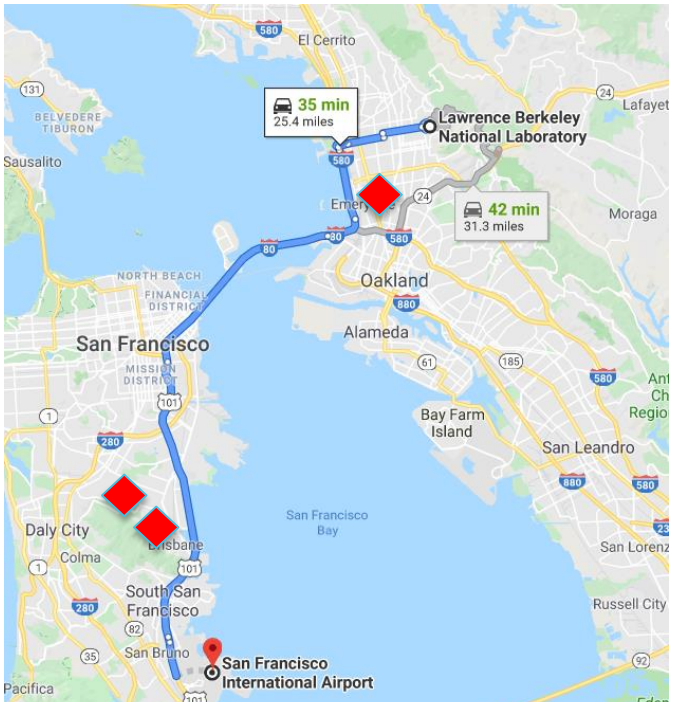


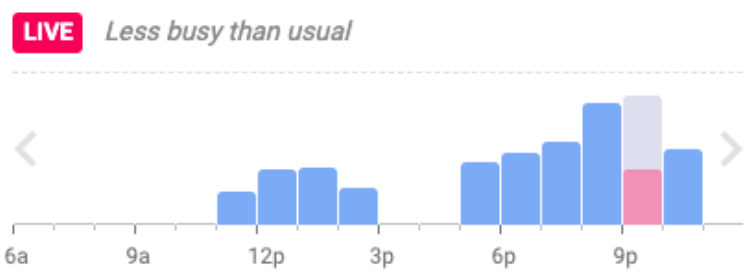


DDCRNN model

An easy to use 'Google-map' style to offer Bandwidth Usage Prediction for Planning and Scheduling.

ESnet

# NetPredict: Google Map for ESnet

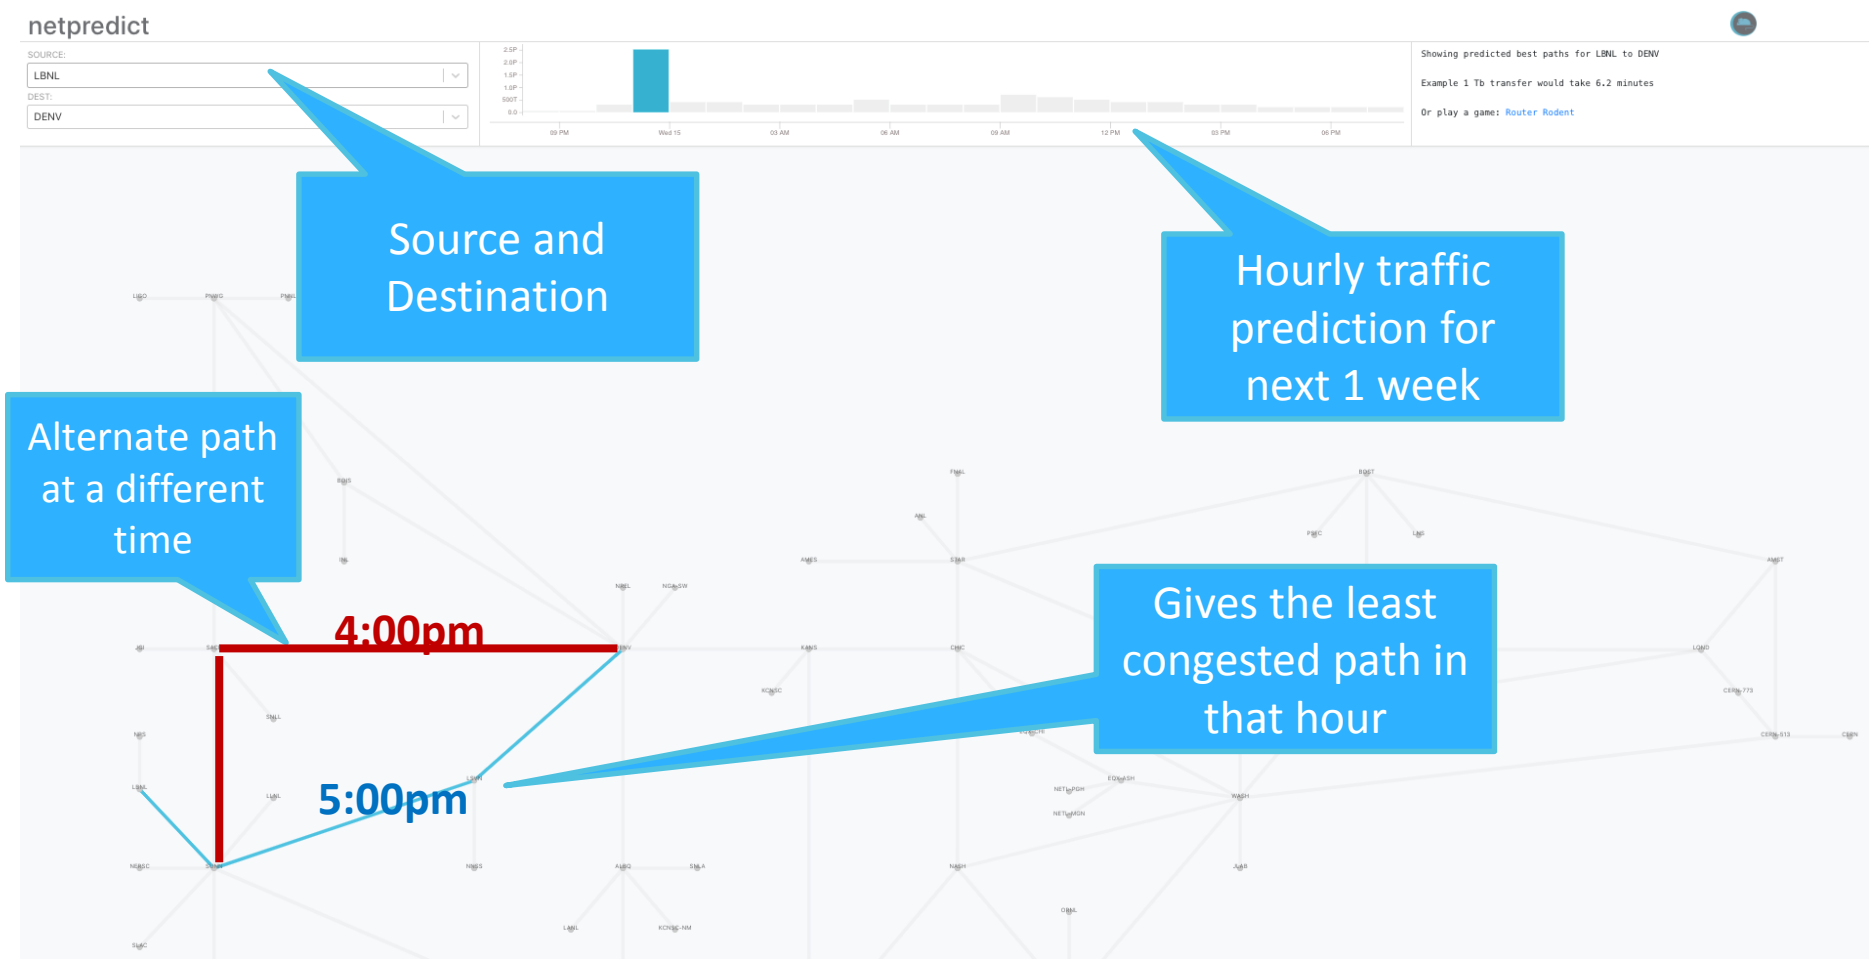


*SC19 Demo Network Research Exhibitions

# System Architecture (hosted on Google Cloud)

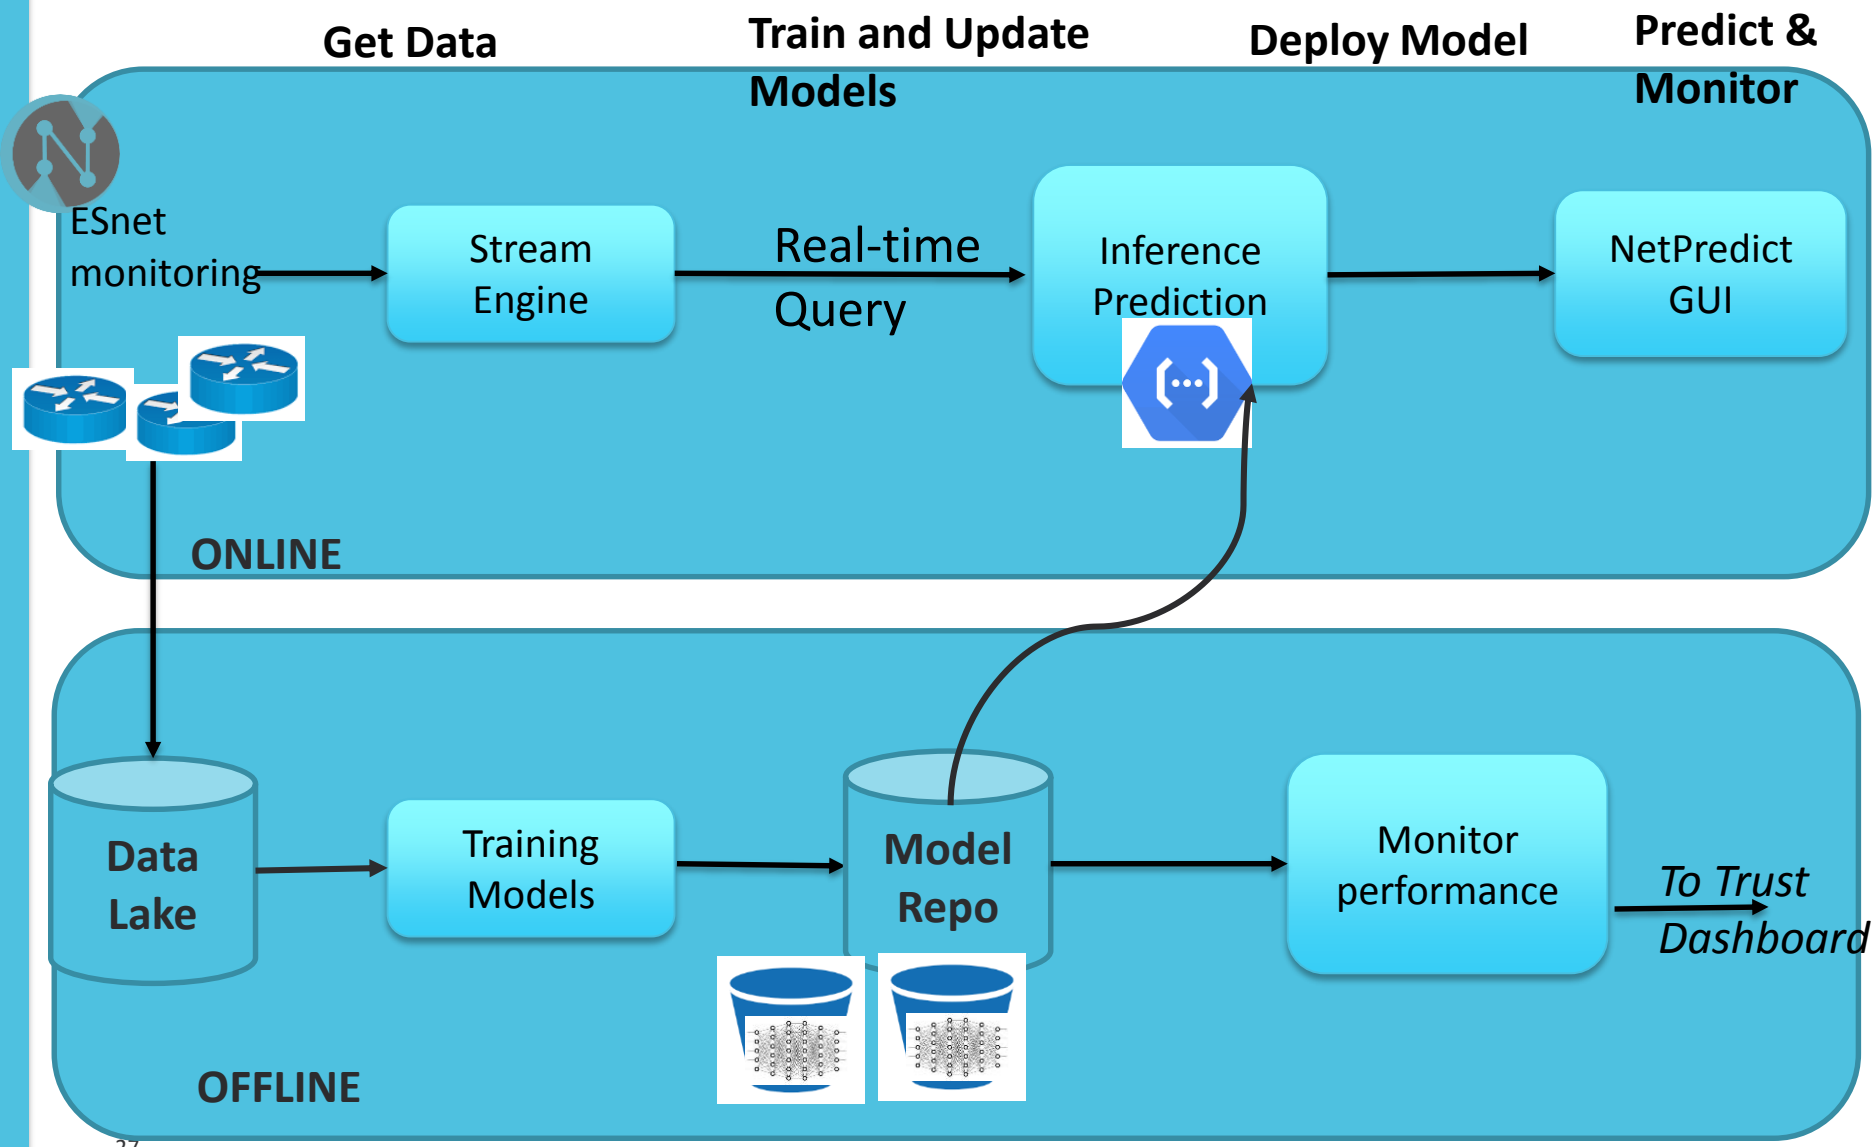

# DeepRoute: Reinforcement Learning for Traffic Engineering

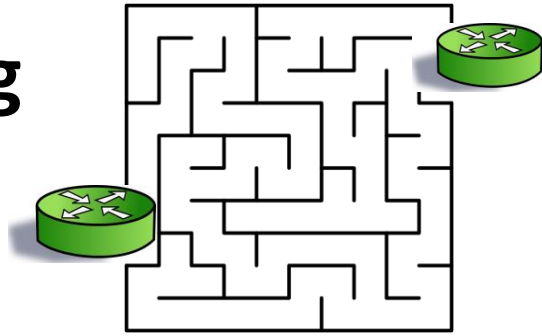

- Multiple ways between source and destination
- Learning the Q function (value function or Bellman equation):

$$Q(s,a) = R(s,a) + \gamma \max_{a'} Q(s',a'))$$

- Learned States: Current utilization, loss and flow details (u,l,f)

| state | action | Reward (flow end time) |
|---|---|---|
| State0 *(u,l,f)* | Move path1 | 10 |
| state1 | Move path2 | 2 |
| … | … | .. |
| … | .. | .. |

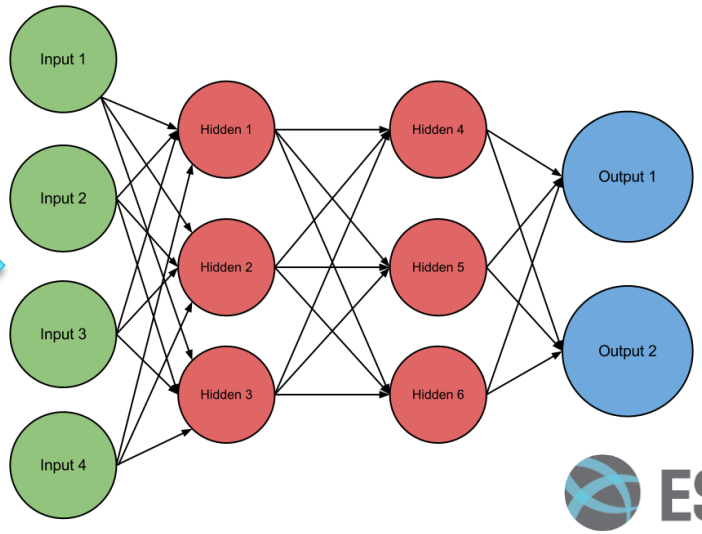


ESnet

# Improving Flow Completion Time

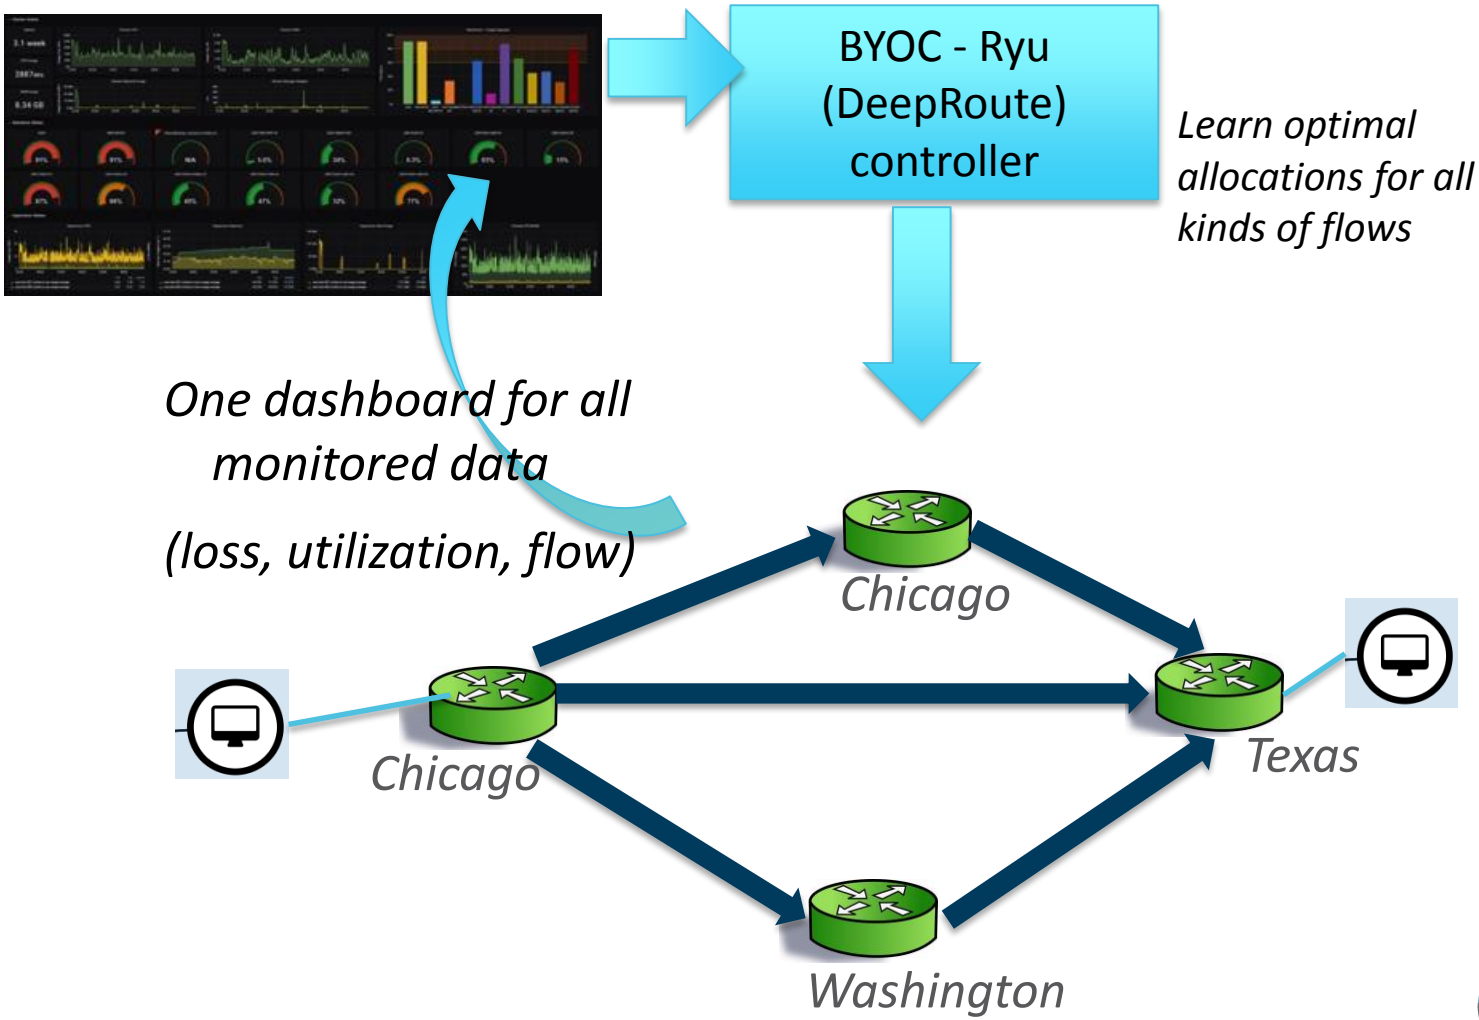


*DeepRoute: Herding Elephant and Mice Flows with Reinforcement Learning, MLN 2019

# Data-driven Learning for Better Paths

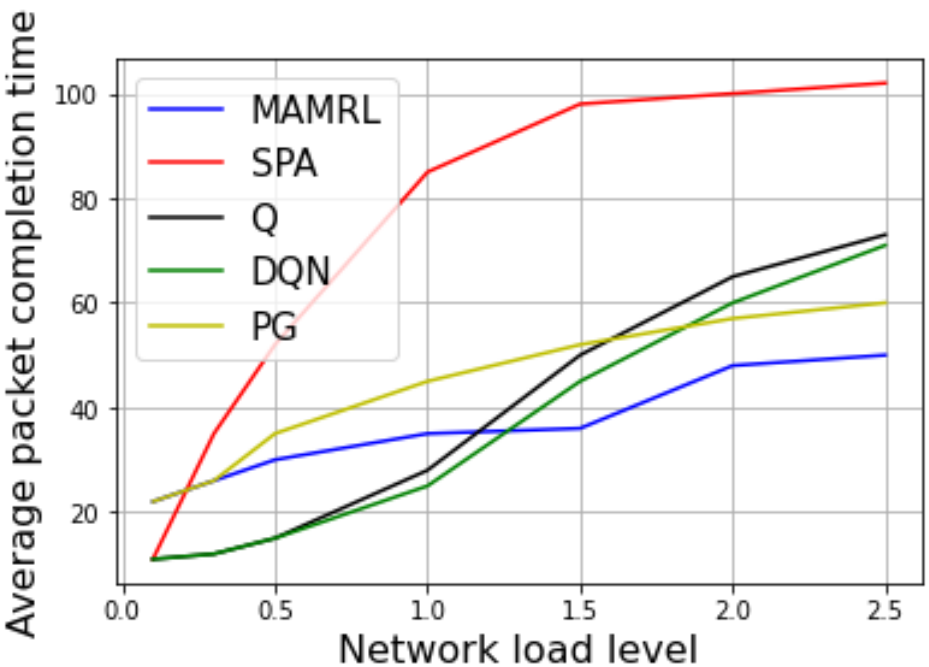


*Assumption: one flow per iteration*

- Can perform multi-objective optimization

Multi-objective optimization =
- Min (flow time)
- Min (packet loss)

- MAMRL – Multi-agent meta reinforcement learning: Higher network load better copes with our approach

ESnet

*Multi-agent meta reinforcement learning for packet routing in dynamic network environments, SC20

# Summary

- Self-driving (autonomous) networks are within reach, with clear objectives
- Share solutions among groups, data and software tools
  - Research papers often omit key aspects of deploying the 'AI'
  - Having Off-the-shelf solutions such as data and tools
- PROPOSAL: Extending NetPredict to other R&E Networks and Groups
- Hardware is expensive, but open testbeds are helping and are necessary (e.g. FABRIC, Chameleon)

ESnet

# Thankyou for listening!

- More information:
  - Open for Collaboration!
  - Deep Learning and Artificial Intelligence for High-Performance Networks (DAPHNE Lab)

(https://sites.google.com/lbl.gov/daphne)

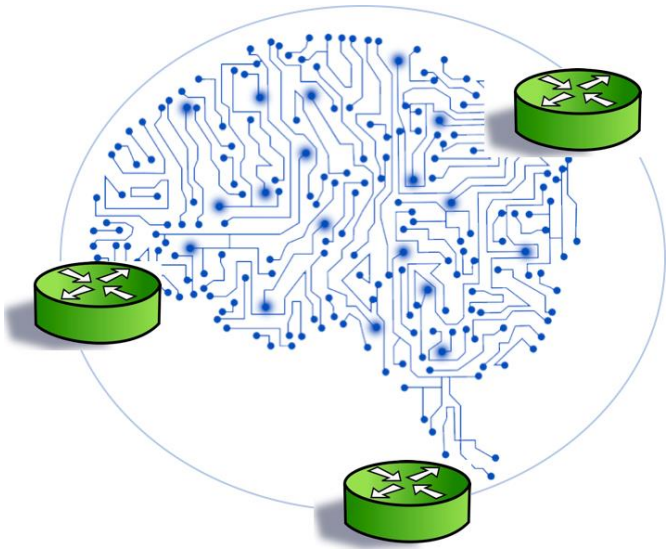

ESnet

- Extra slides

# How to action

- Networks engineers should work with software engineers
- Software engineers are trained to work with data processing challenges, streaming data, tools, testing, etc
- We should be forming a community:
  - Think of what data you can make available
  - Think of what software library also
- How to start: Find a problem with concrete goals of what the ML application should do
  - Collaborate with ML enthusiastic (e.g. students)
  - Deploy, test and test
    - "Search for Bounds at which the AI will fail"

ESnet

# More Problems to Explore

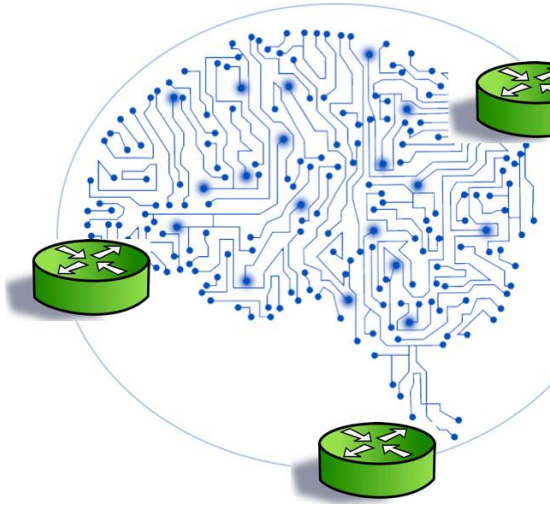

- Integrating AI with Control plane but also AI at the Edge
  - Filter traffic using classification techniques
- Optimizing network configurations
  - Checks, configs
  - Debugging
- Optimize engineer's times
  - Timely alerts if something wrong, bad behaving routers/switches
  - Vendors offering this but it is quite easy to do yourself
- Unexpected versus expected network behaviors
  - Topology or flow changes, day and night patterns
- Intent-driven networks

ESnet

# Testing Multiple ML models

Multiple models tested in real time

MSE is growing, model not relevant any more

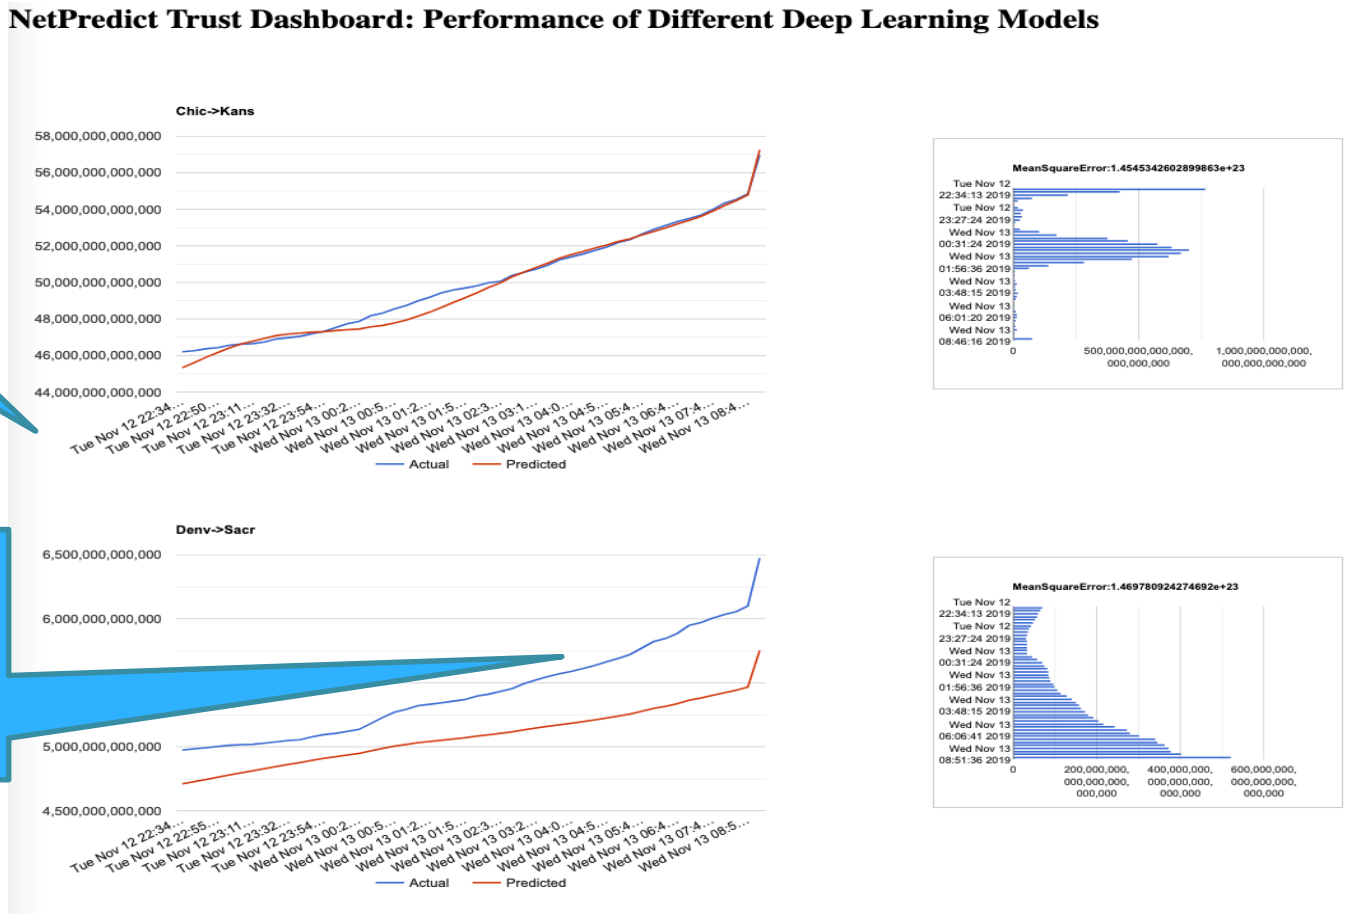

# Supervised Learning

<u>When Labelled Data is Available</u>

- Learning to recognize the labels
- Classification, object detection, anomaly detection
- Works very well if we can identify clear class boundaries

Good versus bad flows

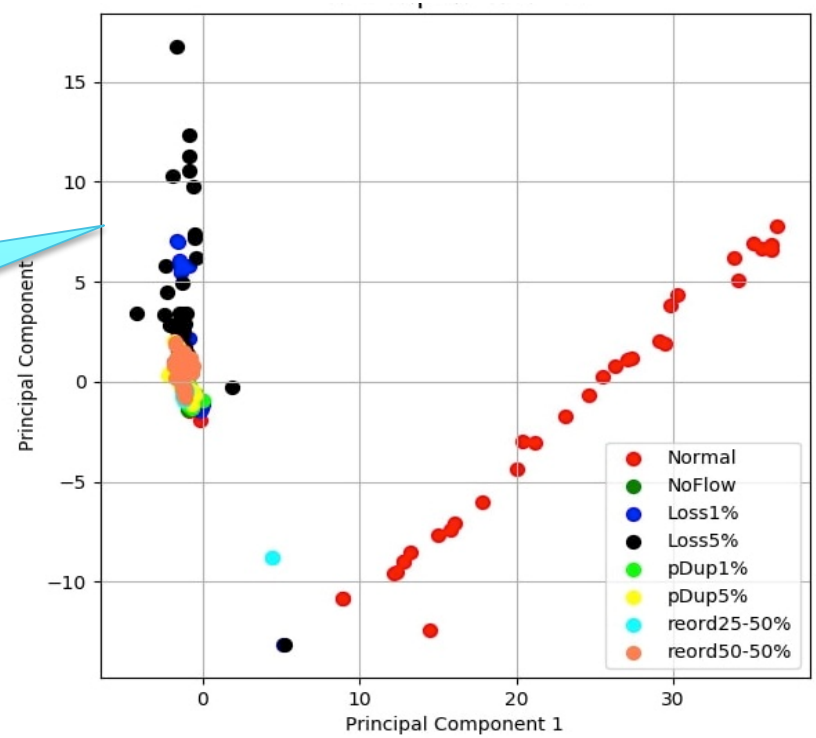


*Plot of anomalous TCP traffic*

*S. Campbell, F. Bannatwala, SMC 2020

# Unsupervised Learning

<u>When NO Labelled data in Available</u>

- Learn underlying rules in the data
- Clustering, feature identification, recognize anomalies in test data

- Need domain scientists after clusters are recognized

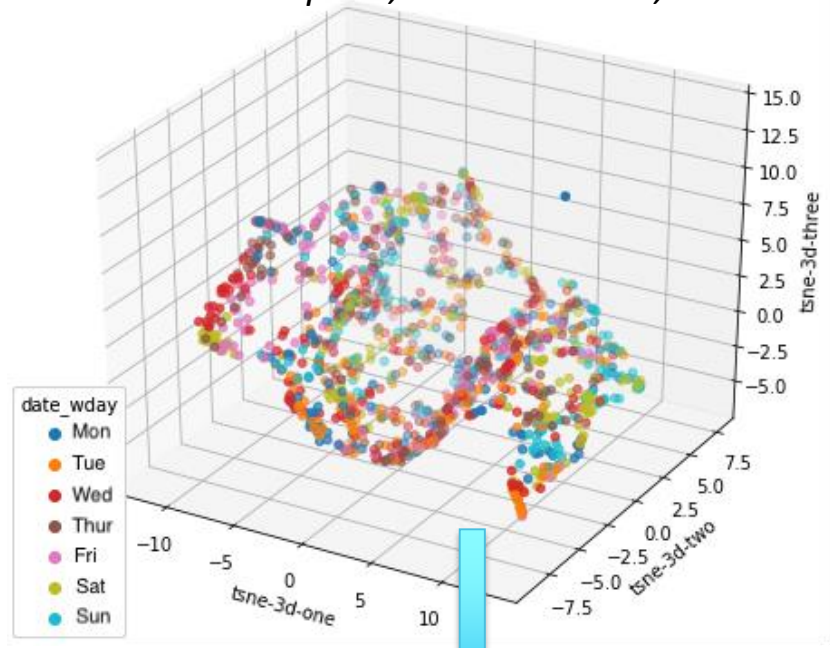


Visualizing NERSC flows into clusters

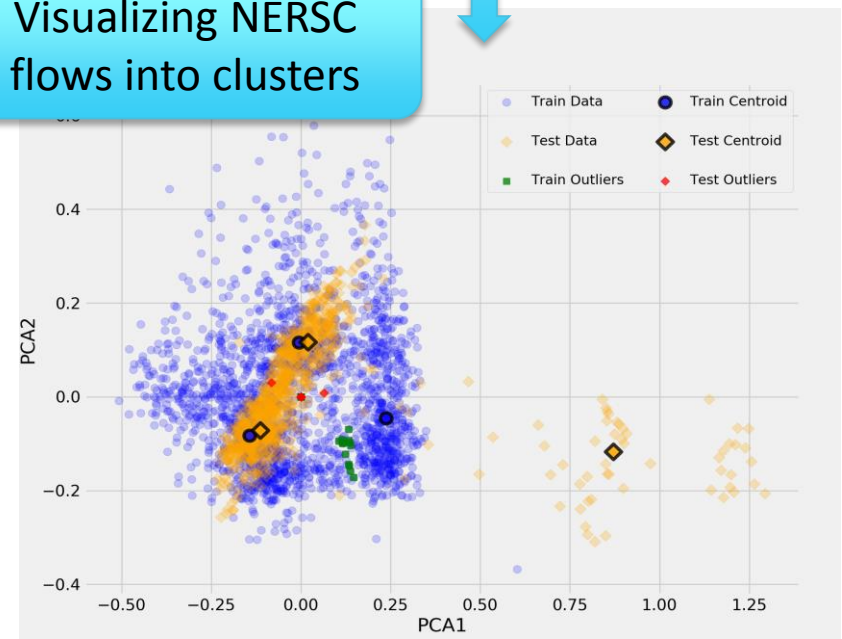


*Plot of NERSC netflow transfers*

# Keep Trying different AI models

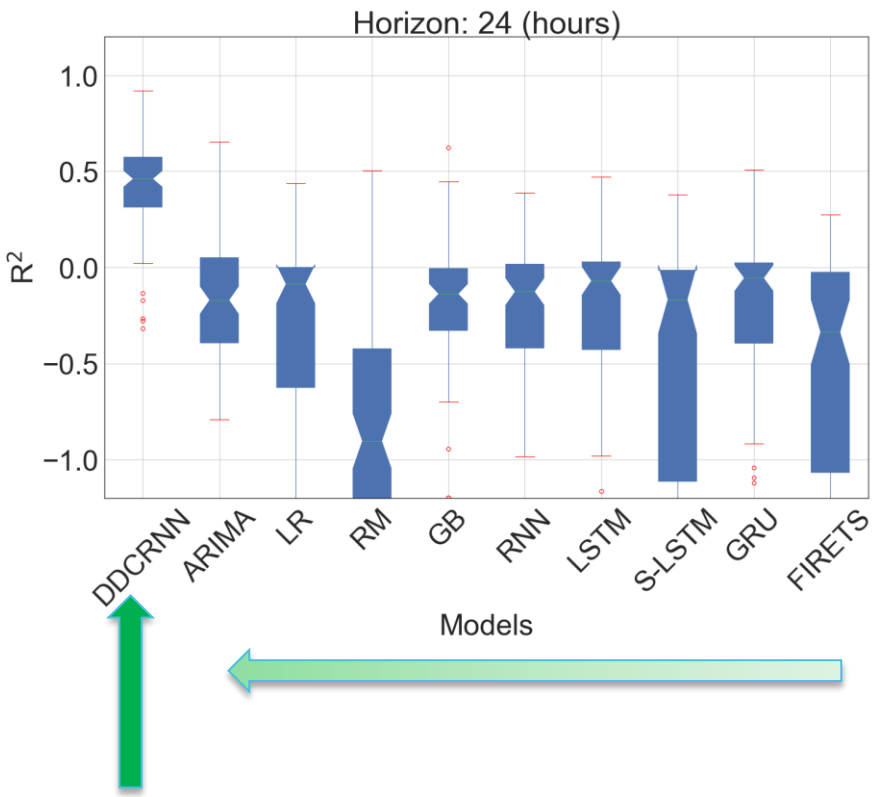


- No need to learn long term traffic patterns (short term is enough) [4]
- Developed our own Graph neural networks to improve prediction results
- Can still be active research
  - Improving time predicted
  - Improving AI model itself



[4] Mallick et al. Dynamic Graph Neural Network for Traffic Forecasting in WAN

***SC19 Demo Network Research Exhibitions**

- NN models pretrained on HPC facilities (NERSC)
  - 2 years of SNMP data
  - Save all AI models
- Deployed on Google Cloud Platform
  - Buckets save tensorflow models
  - VMs to do inference
  - App to visualize results

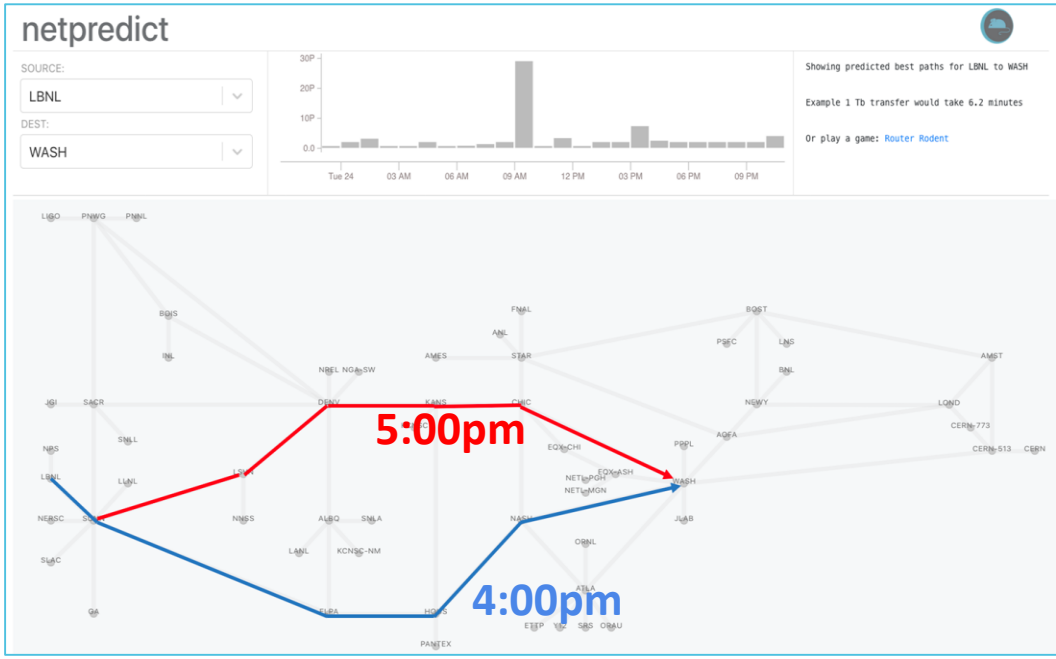


ESnet

# LSTM to model time as Sequences

- Time as discrete sequences

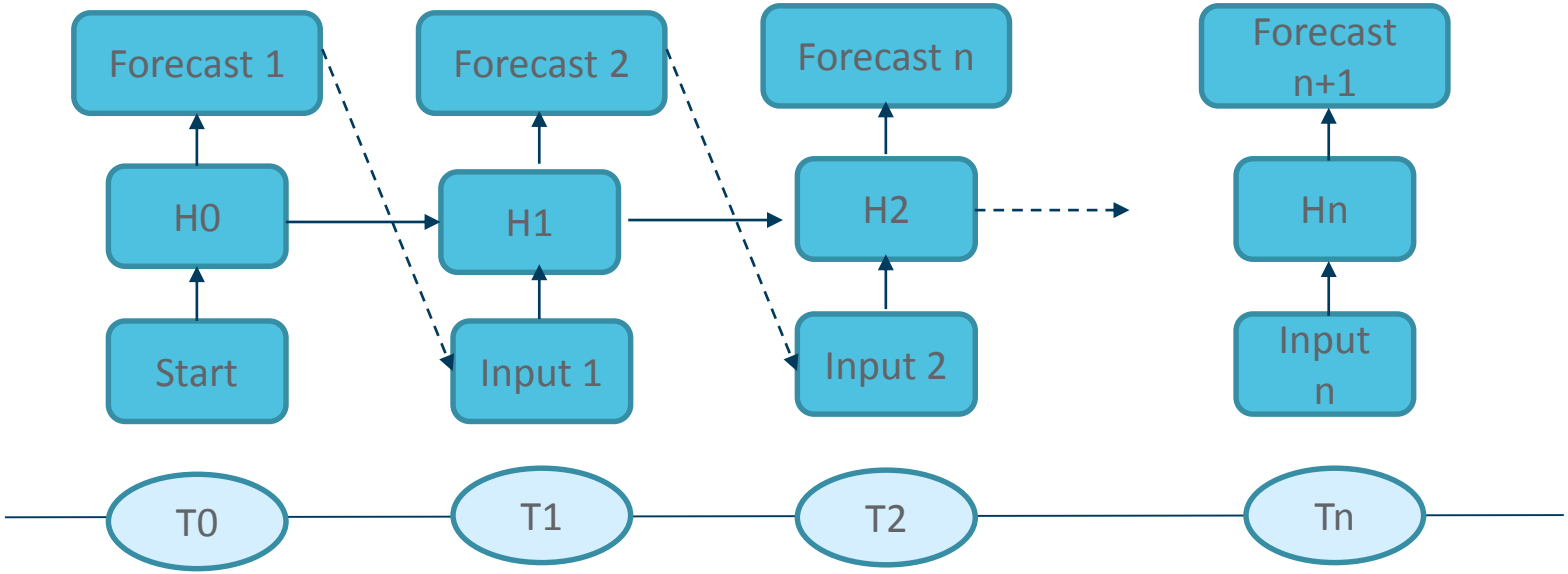

# Add features to LSTM to predict future

- Further signals which might affect data, e.g. time of day, weekday, weekend