Deep Learning for Transportation

Fast Estimation of Travel Times Using Historical Routes

Dmitry Kudinov

dkudinov@esri.com

Why? The need for performance and flexibility

Large fleet optimizations and hard to model aspects

Traditional methods are slow to compute large volumes efficiently:

1. Large transportation companies route thousands of vehicles daily.
2. Origin-Destination Cost Matrix is the foundational piece, but
3. A 1,000 set of stops to route translates into 1,000,000 travel time estimations.

Hard to model aspects which can be captured with Deep Learning:

- Seasonal traffic fluctuations.
- User preferred routes.
- Multimodal transportation in urban areas.
- Parking delays of individual locations.

Multilayer Perceptron (MLP) predicts Travel Times

MLP architecture:

- 10M trainable parameters x 16 dense hidden layers
- Dropout with keep_prob = 0.9 before the single-neuron output layer

Training Data:

- Preprocessed GPS reports combined into “journeys”, and / or
- Synthetic data produced by traditional routing algorithms.
- ~500M samples covering California and Nevada road network

```
mc [dkudinov@dmitry00]:~/TF/VASYA/data/CA_NV/training_data/mix_300MrandomRo...
/home/dkudinov/TF/VASY~Routes-p3-1000000.csv        1316/7357K        0%
X1,Y1,X2,Y2,START_TIME,NA_COST
-121.27603,38.227002,-121.118766,38.878532,1491561882657,56.024480662457
-121.304219,38.275807,-121.134275,38.862371,1491562097358,50.927425051281
-121.320508,38.309179,-121.170968,38.831802,1491562237716,45.8648692683069
-121.344763,38.360109,-121.191444,38.814621,1491562445663,40.8302296522048
-121.377522,38.392772,-121.221721,38.78919,1491562606646,35.7981504295781
-121.396035,38.425326,-121.258288,38.76284,1491562746631,30.7620486428579
-121.413511,38.463263,-121.281674,38.732524,1491562910897,25.6313266725365
-121.445626,38.495886,-121.311729,38.705218,1491563076316,20.5662843060309
-121.465856,38.521839,-121.35057,38.669656,1491563200034,15.3671588807884
-119.870684,35.495821,-119.773688,36.891619,1491600699785,166.738953396278
-122.27821,38.311581,-121.840912,37.668734,1491325222105,88.8718196026759
-122.277136,38.308901,-121.843662,37.670817,1491325337880,83.7967544798079
-122.275565,38.294956,-121.860476,37.673595,1491325490531,78.7275330741487
-122.274308,38.280665,-121.874319,37.677534,1491325683927,73.4141386984336
-122.271149,38.258982,-121.878589,37.690973,1491325854488,68.2228442326205
-122.263463,38.236635,-121.873202,37.700431,1491326050708,63.158025050229
-122.253395,38.222612,-121.902432,37.701483,1491326201784,58.1262838525141
 1Help  2UnWrap 3Quit  4Hex  5Goto  6  7Search 8Raw  9Format10Quit
```

MLP and Training data

Training

- **Nvidia GP100** (Intel® Xeon(R) CPU E5-1620 v2 @ 3.70GHz × 8)
- **2 x Nvidia GV100 + NVLink** (Intel® Xeon(R) Gold 6136 CPU @ 3.00GHz × 47)
- TensorFlow 1.4 / 1.5
- Training time for up to 4,800 epochs with exponential minibatch size

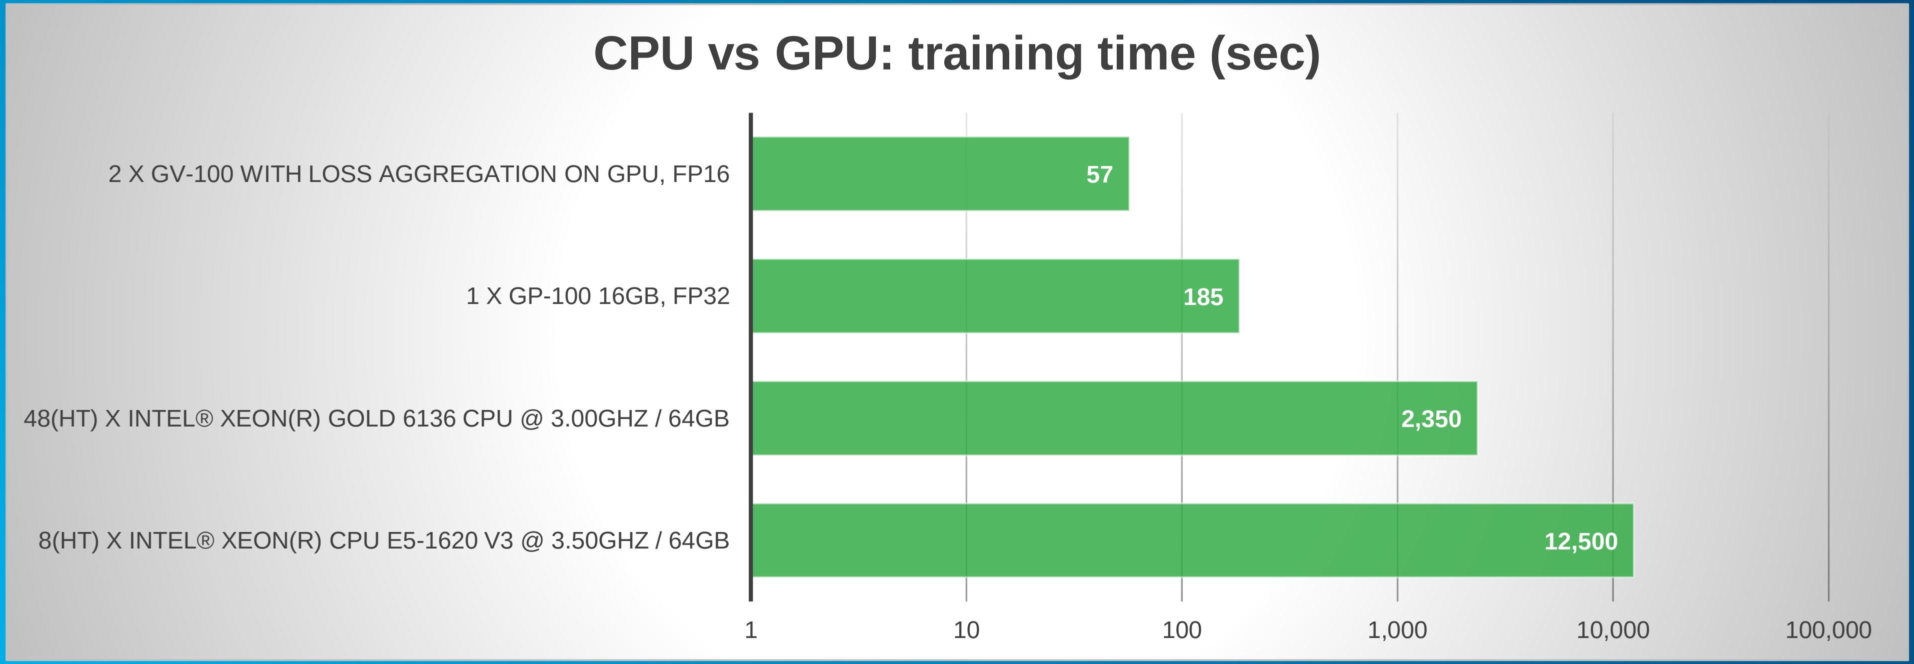

Power of GPU

How to check prediction quality beyond single MSE value?

1. **Standard Deviation** of predicted ETA versus Ground Truth as a function of route length.

2. Isochron–based visualization (how far can you get from a single point on a road within next 30 minutes?)

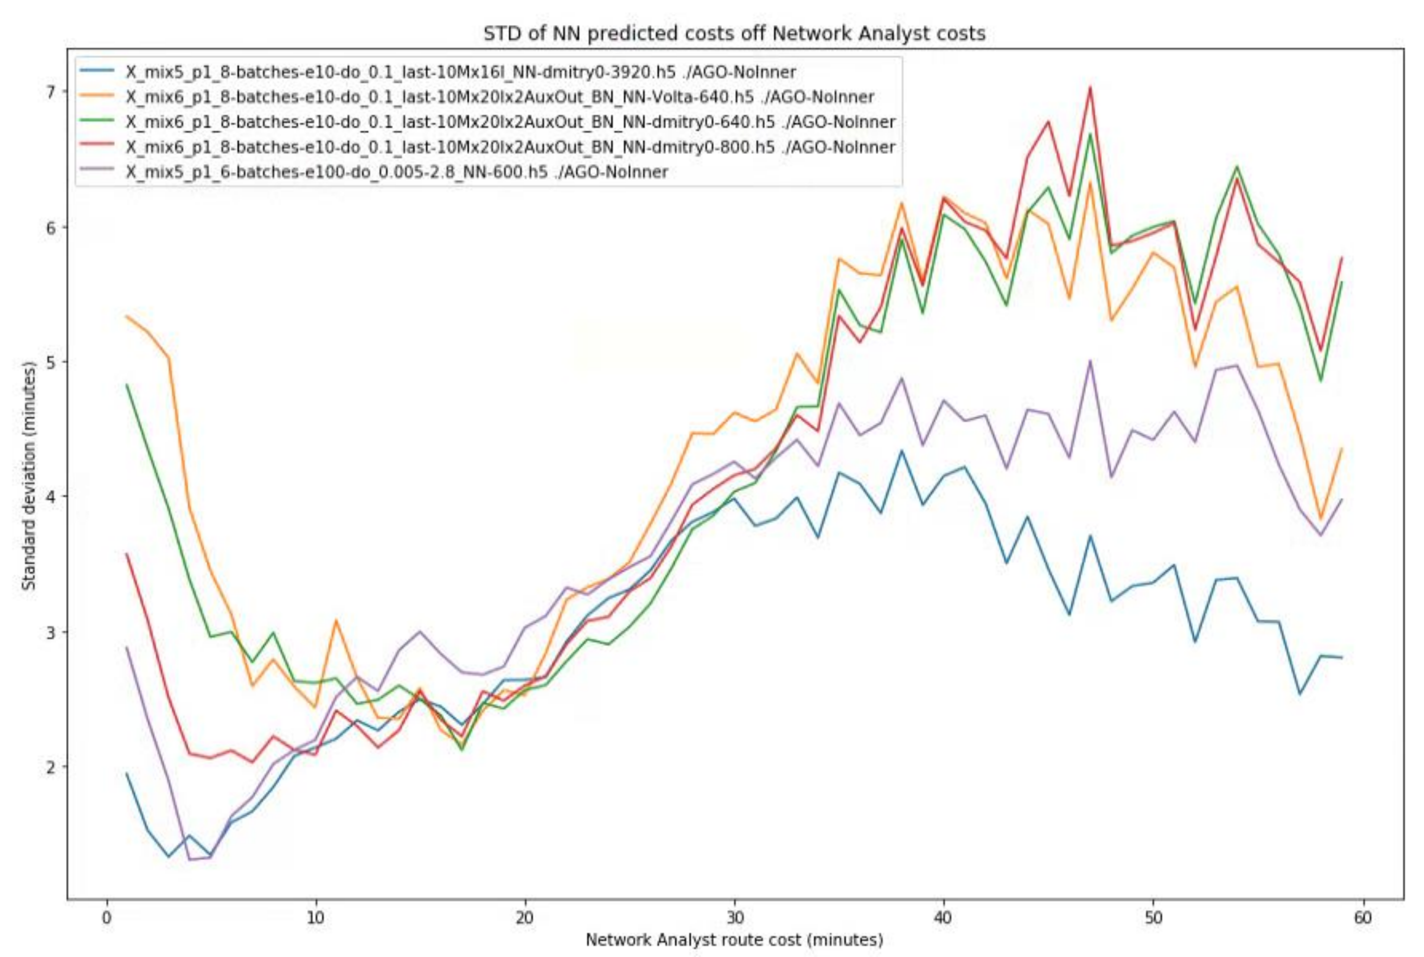

Quality Control

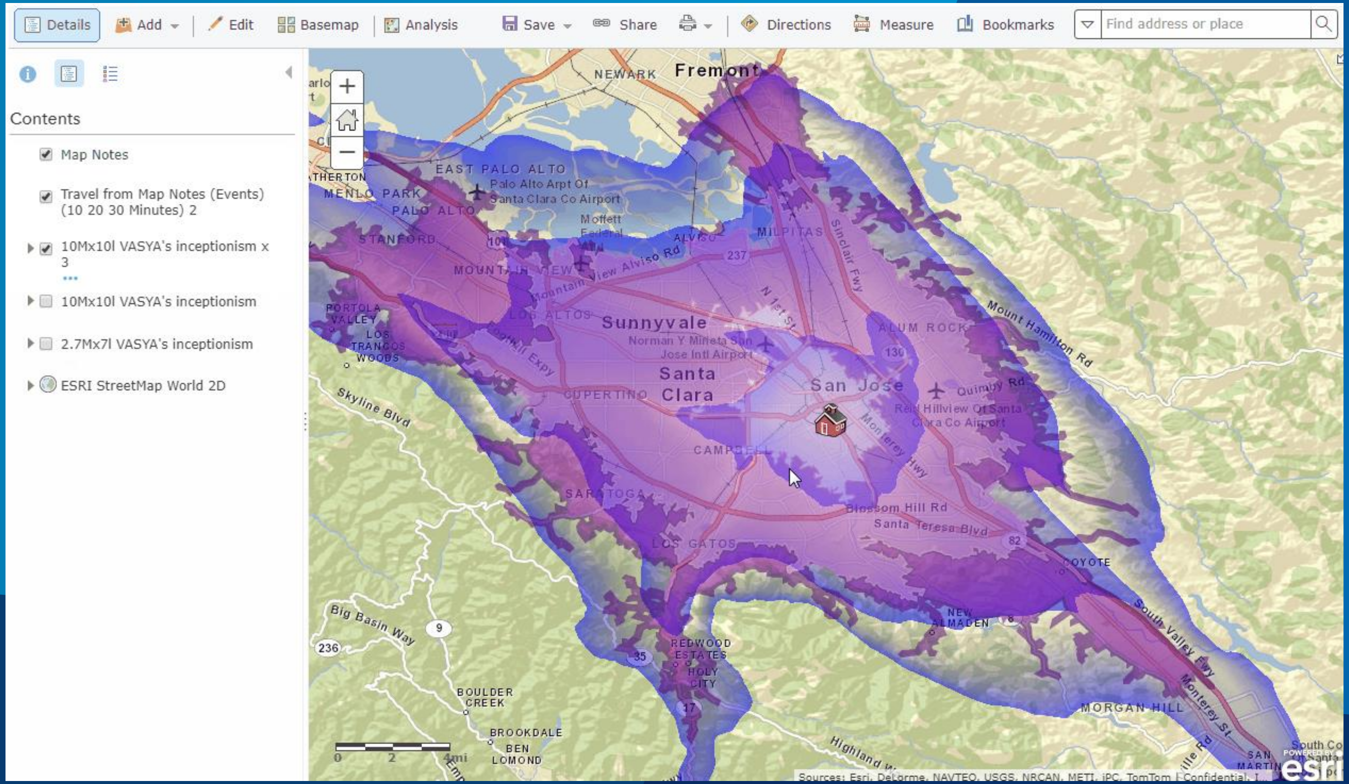

DEMO

MLP produced isochrons on the map

Curious case of MLP-based travel predictions

Road congestion patterns which are captured from simple GPS reports during training show, that the suggested Deep Learning architecture is capable of storing and accurately representing hard-to-model transportation aspects.

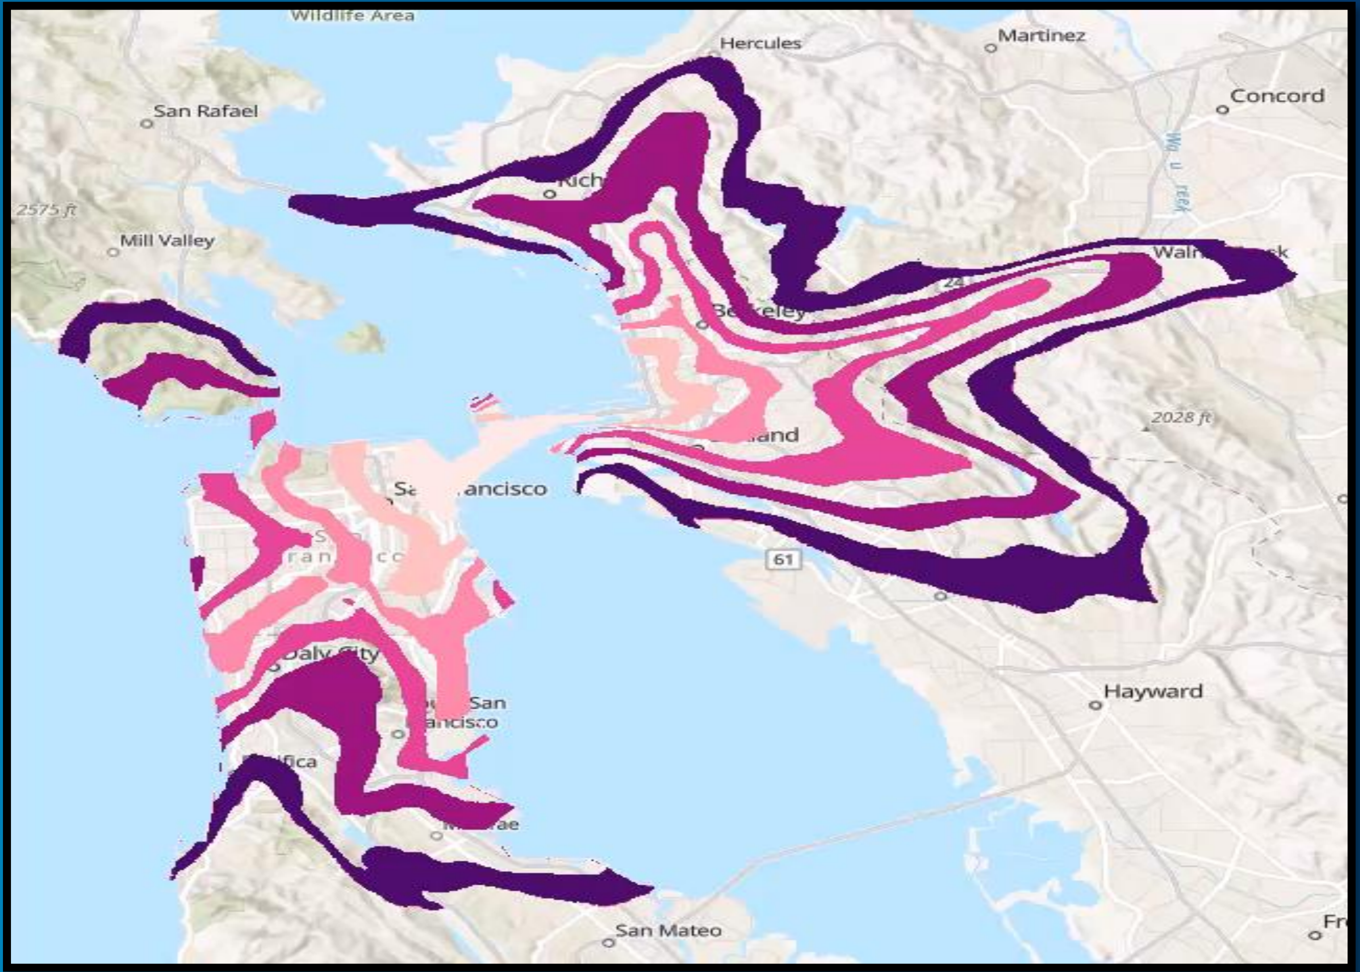

Capturing road congestion from GPS reports

What does it mean: “MLP knows about road network?”

- **Visualizing the training process with geographically bound Isochrons**

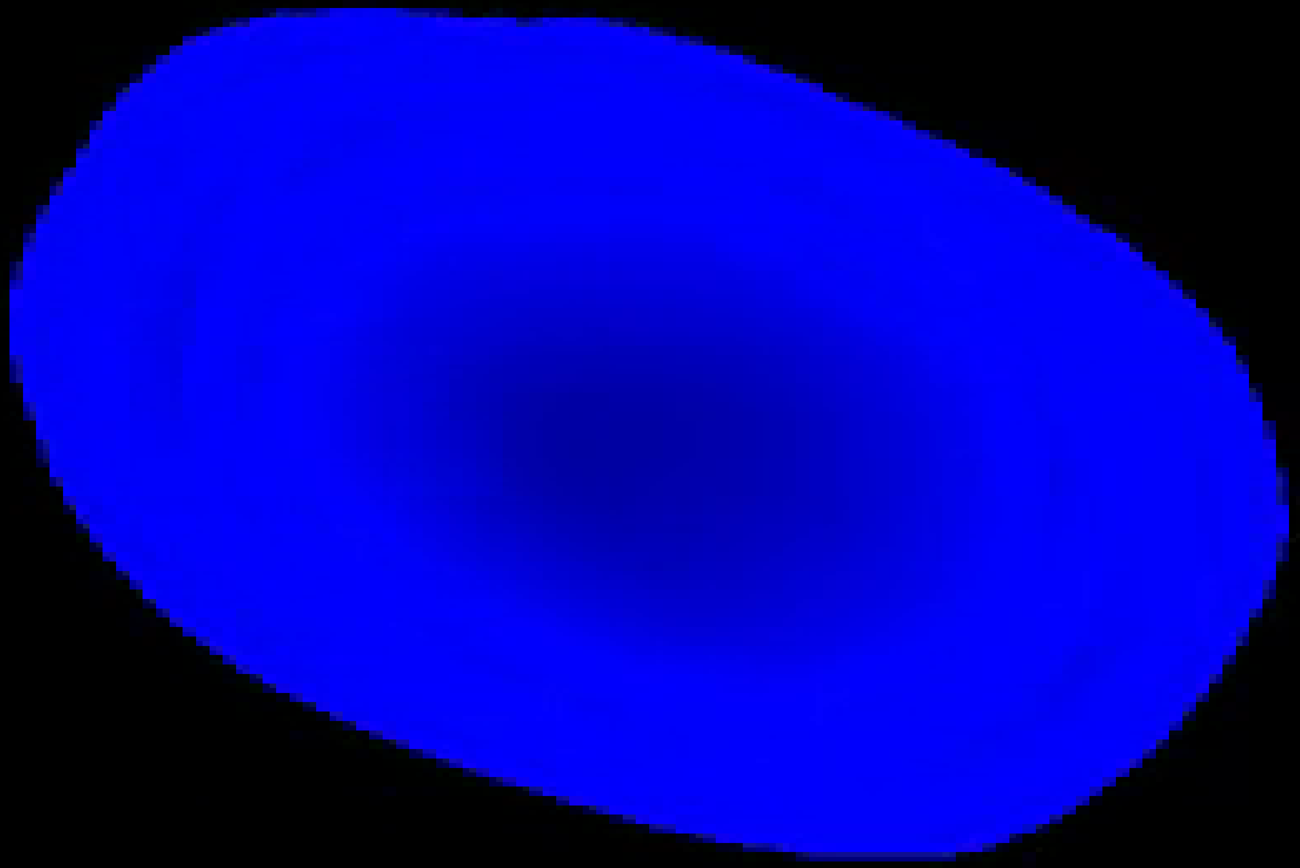

So, how does training look like?

esri

THE SCIENCE OF WHERE