Data-Driven Datasets

Deep Active Learning for Autonomous Vehicles, ML and Beyond

Adam Lesnikowski @lesnikow
Senior Software Perception Engineer, NVIDIA
NVIDIA GTC Silicon Valley 2019 Wednesday March 20th, 2019

Data-Driven Datasets

Deep Active Learning for Autonomous Vehicles, ML and Beyond

Adam Lesnikowski @lesnikow
Senior Software Perception Engineer, NVIDIA
NVIDIA GTC Silicon Valley 2019 Wednesday March 20th, 2019

Data-Driven Datasets

Deep Active Learning for Autonomous Vehicles, ML and Beyond

Adam Lesnikowski @lesnikow
Senior Software Perception Engineer, NVIDIA
NVIDIA GTC Silicon Valley 2019 Wednesday March 20th, 2019

Motivations

+

=

Xu et. al

+

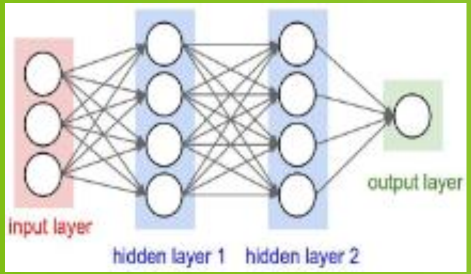

=

So how do we go from:

Xu et. al

+

Not All Data are Created Equally...

vs.

Some Are Much More Informative

Two Motivating Challenges

Two Motivating Challenges

1. How do we improve the cost complexity of our labeling?

Two Motivating Challenges

1. How do we improve the cost complexity of our labeling?

2. How do we find the most informative unlabeled images to get the safest autonomous vehicles possible?

Traditional Machine Learning

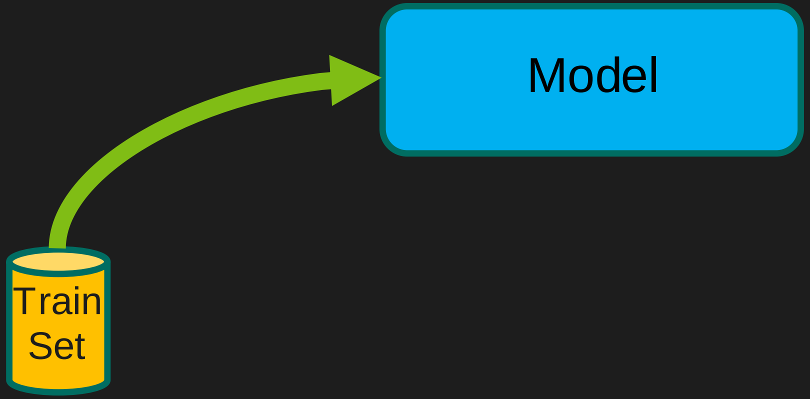

The Active Learning Approach

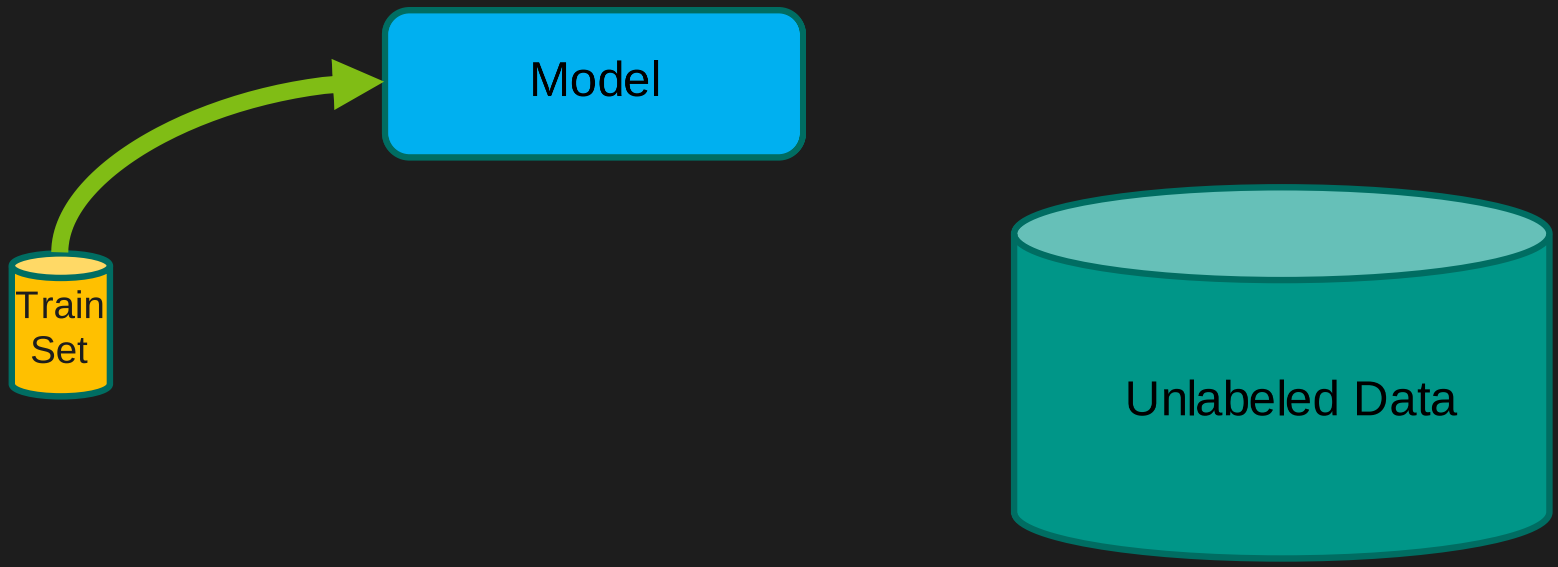

The Active Learning Approach

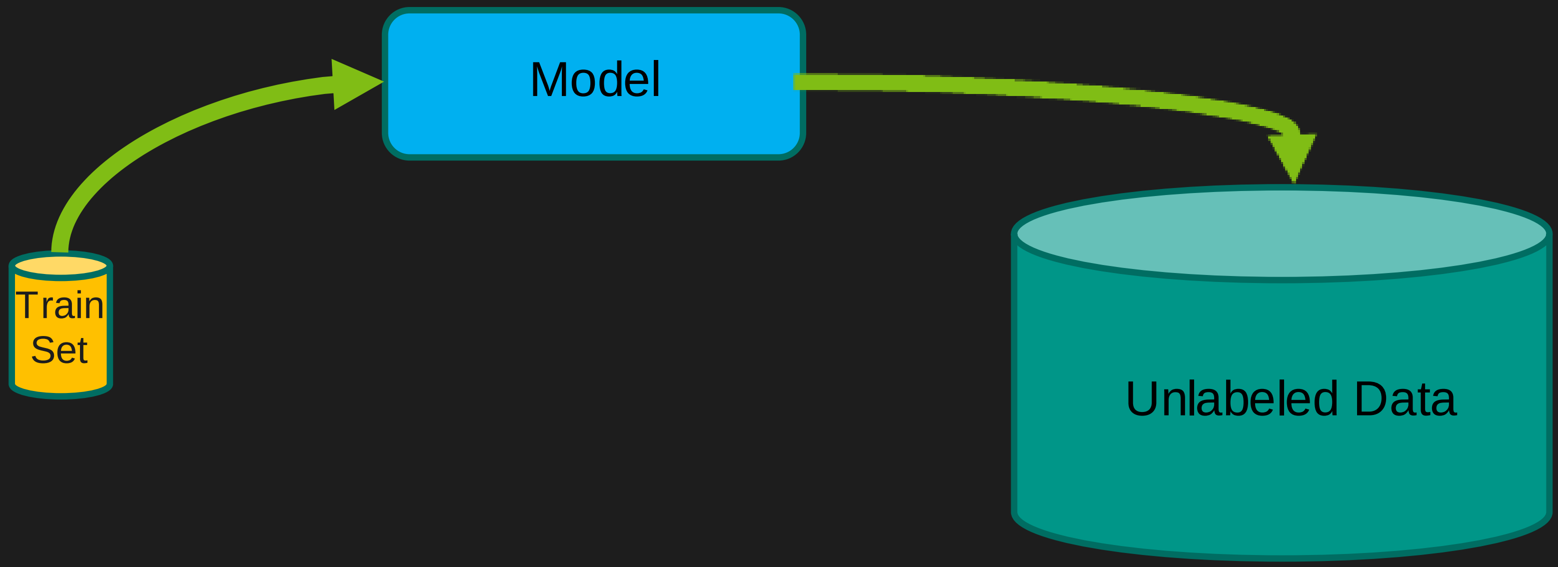

The Active Learning Approach

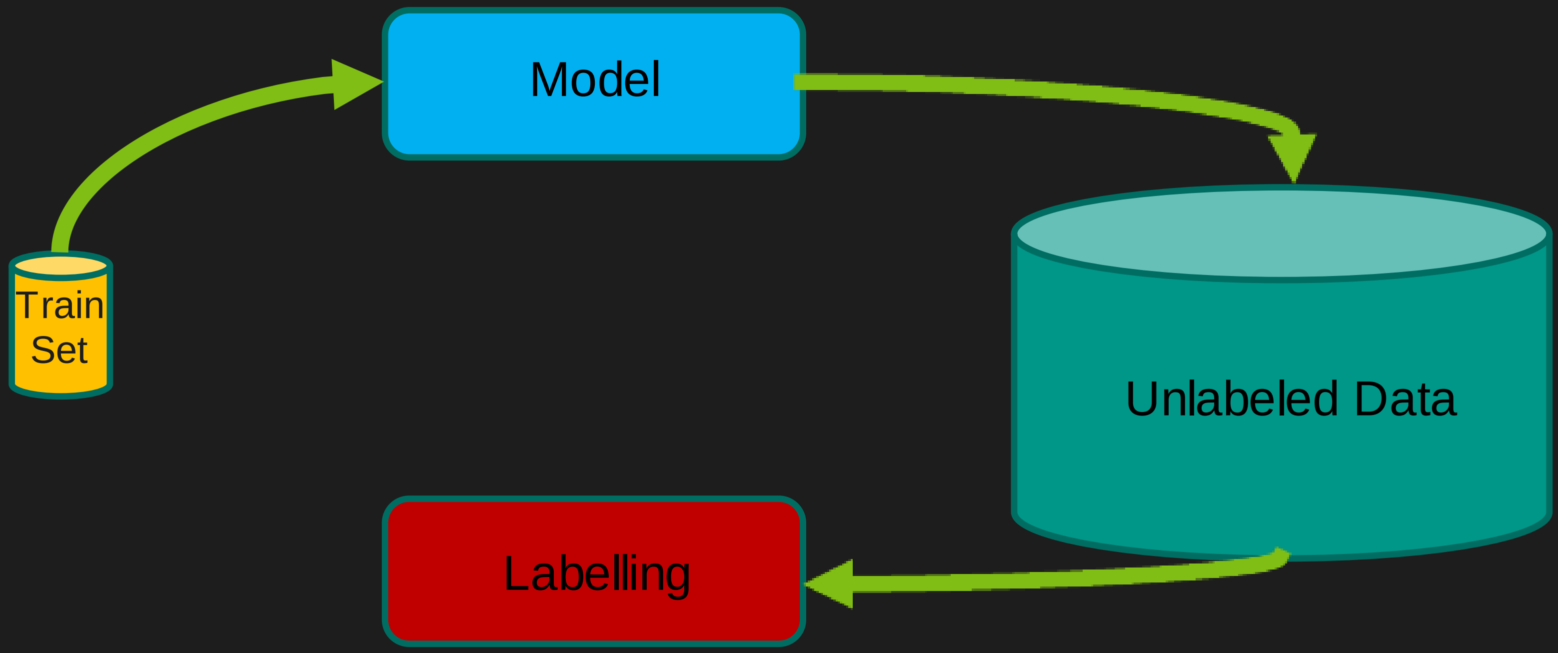

The Active Learning Approach

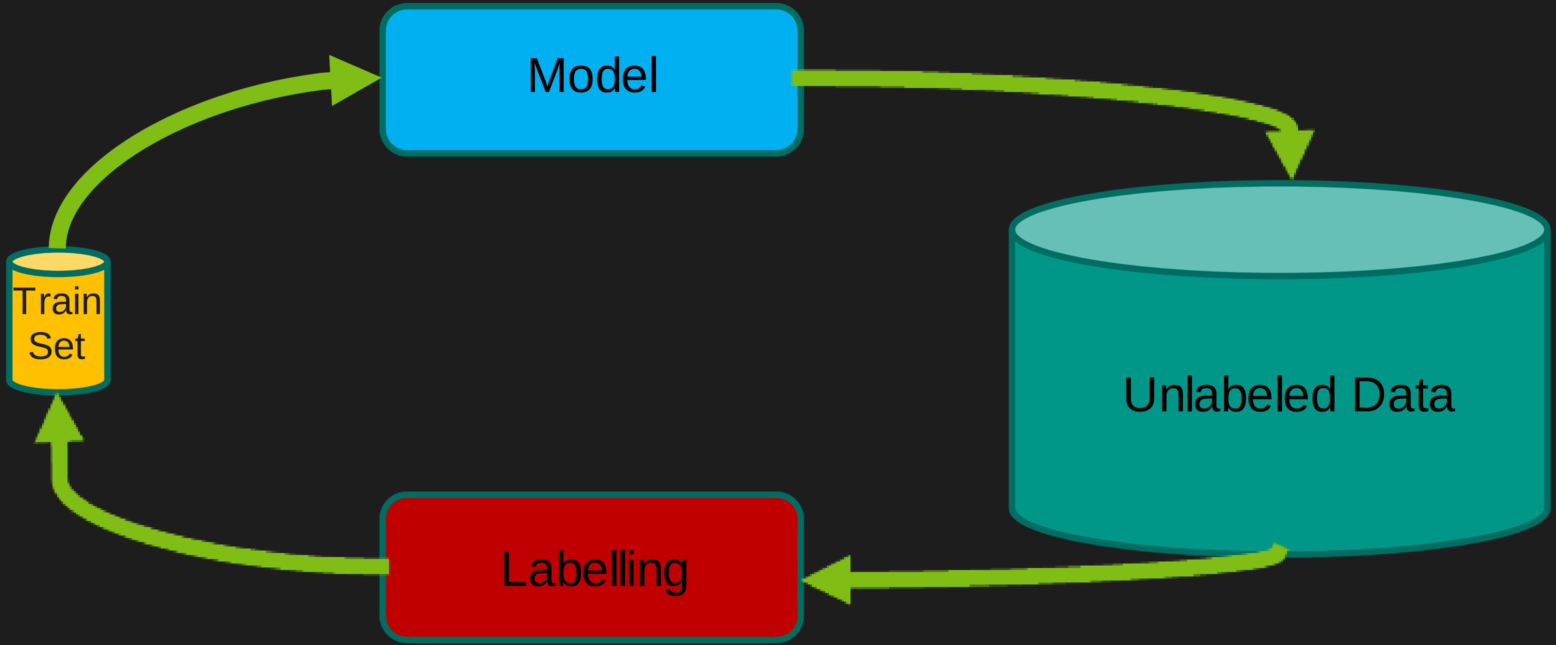

Datasets: Classification

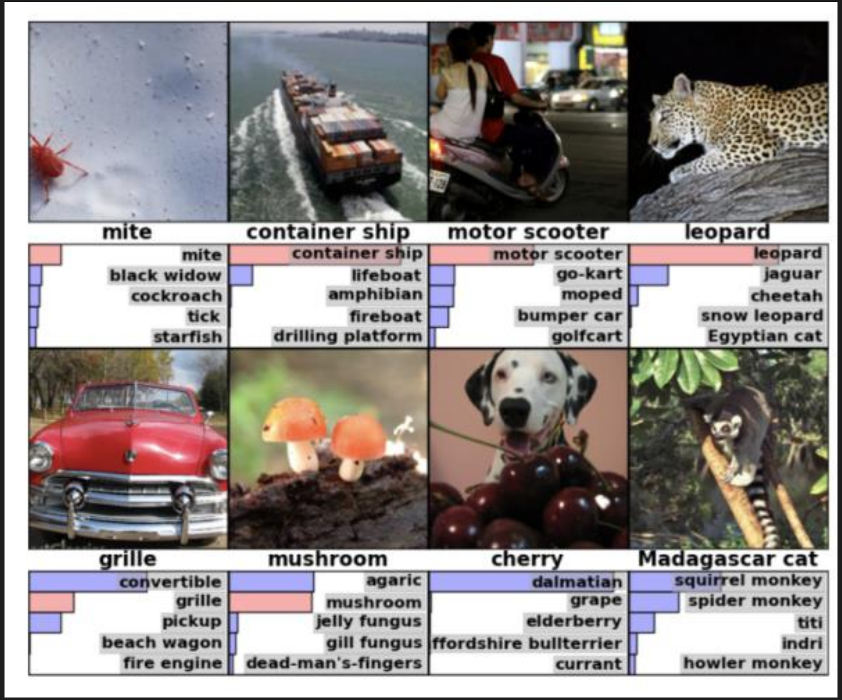

Datasets: MNIST hand-drawn digits, 60K train 10K test split , ImageNet 1K: ~1.3M train images, 50K test images.

Datasets: Lane Segmentation

Dataset: 100K train, 10K val split with pixel segmentations for lanes

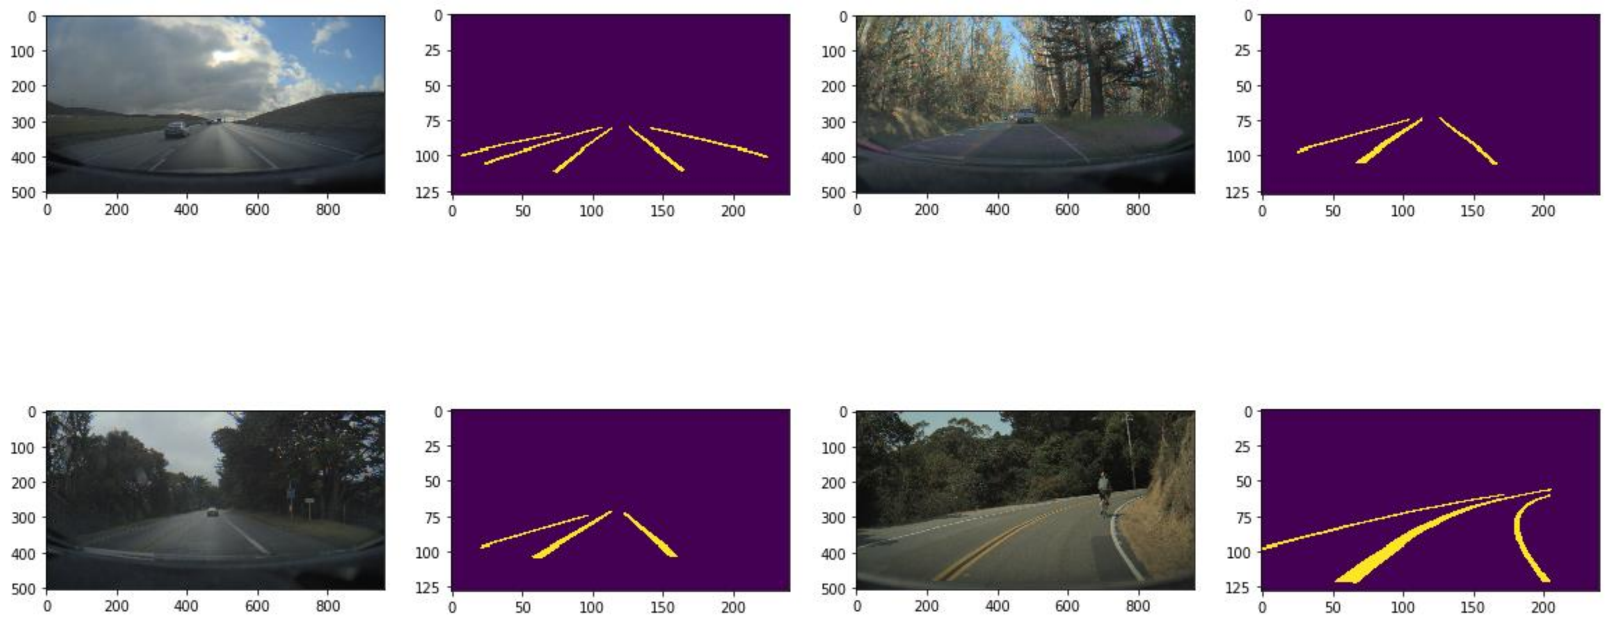

Datasets: Road Scene Object Detection

Dataset: 150K train, 10K val split, with bounding box labels

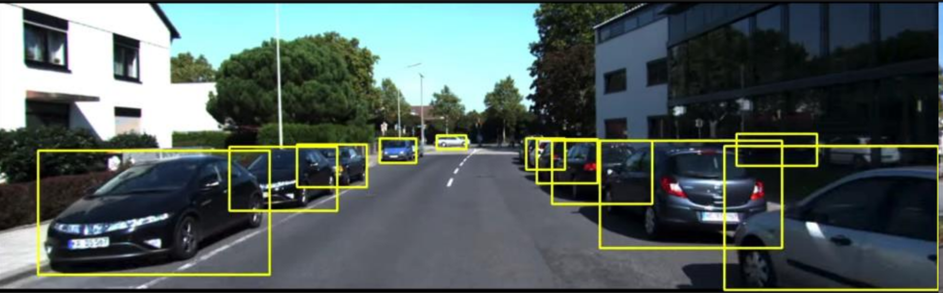
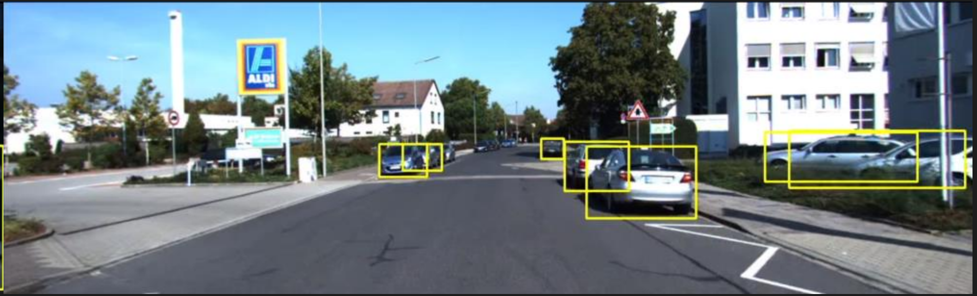


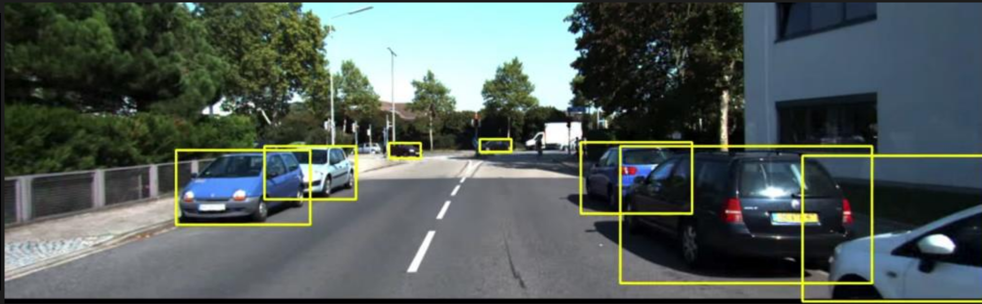

Segmentation Results

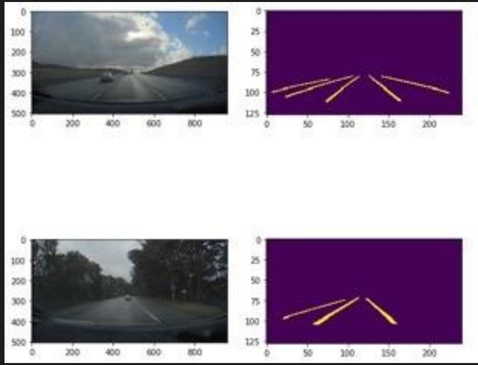

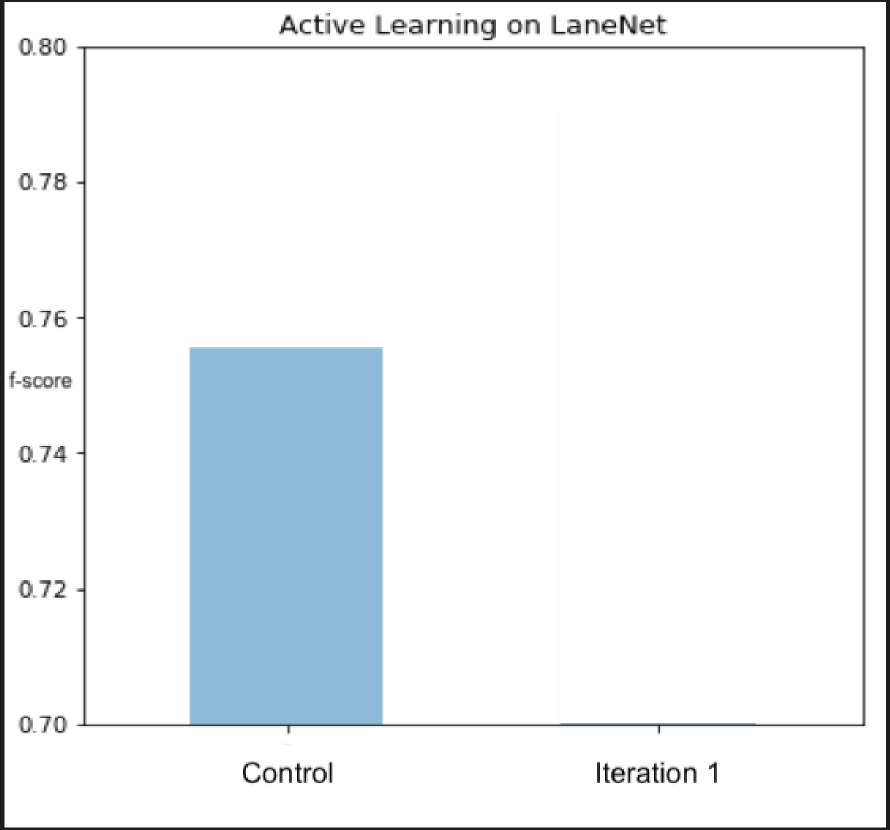

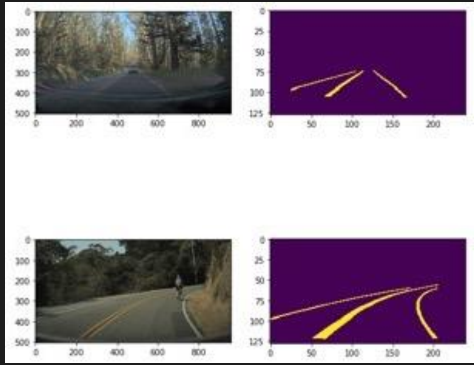

Results of active learning on pixel segmentation for road markings

Segmentation Results

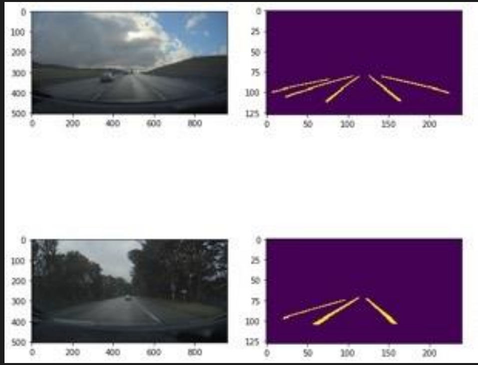

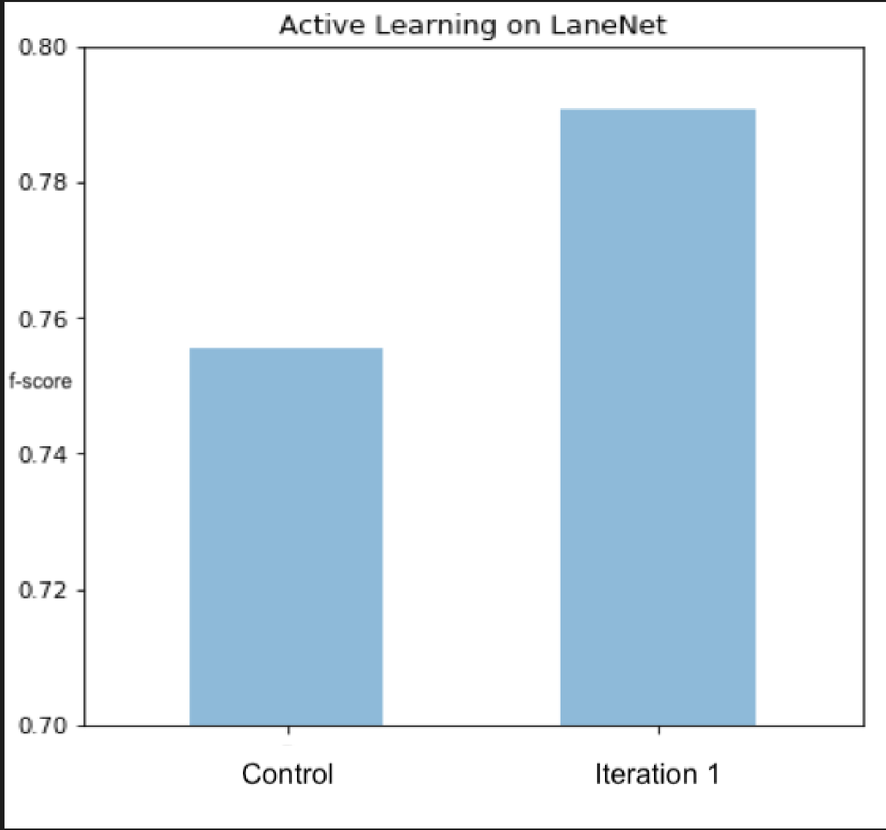

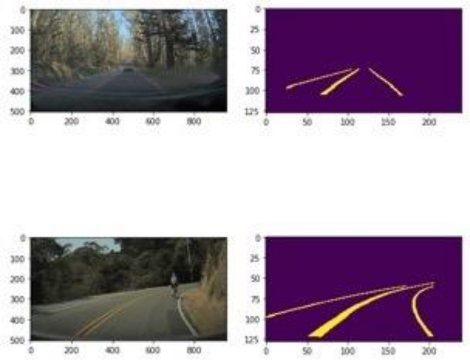

Results of active learning on pixel segmentation for road markings [*Lesnikowski, 2017*]

Segmentation Results

Robustness analysis. Results on segmentation for road markings, across multiple runs of the active learning cycle.

Segmentation Results

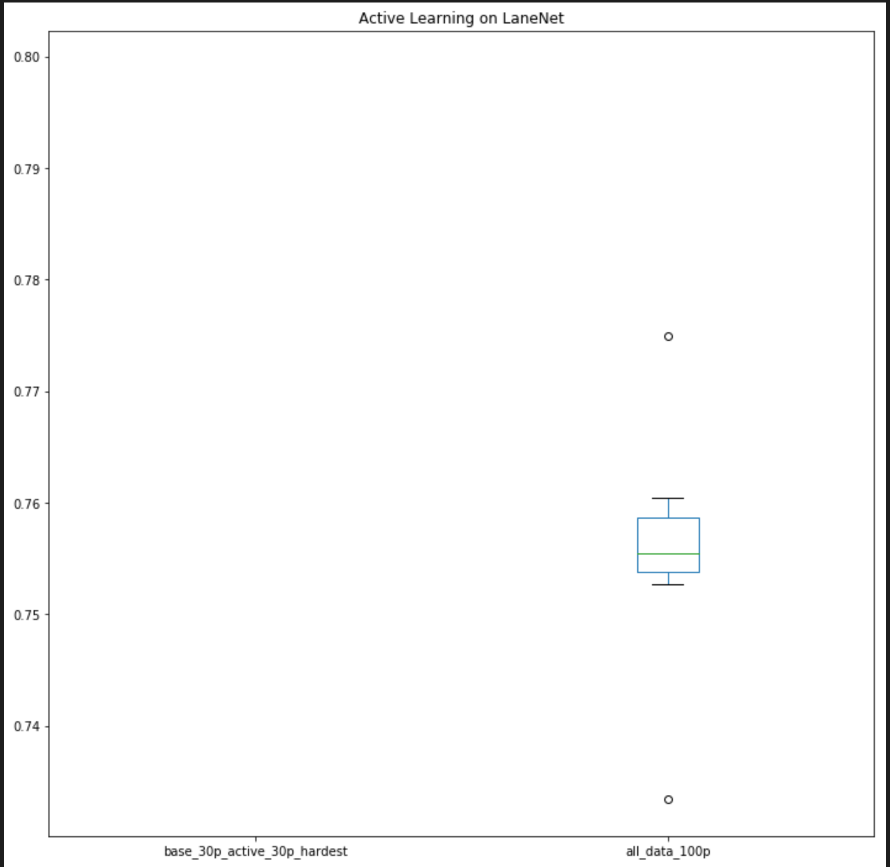

Robustness analysis. Results on segmentation for road markings, across multiple runs of the active learning cycle.

Segmentation Results

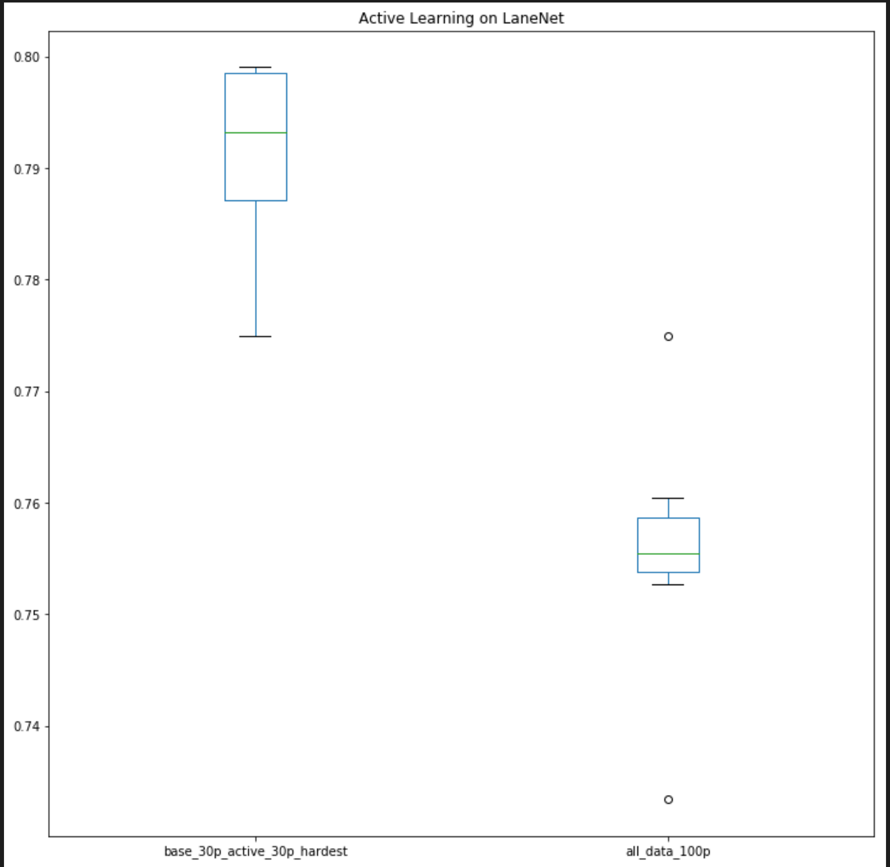

Robustness analysis. Results on segmentation for road markings, across multiple runs of the active learning cycle.

[Lesnikowski, 2017]

Segmentation Results

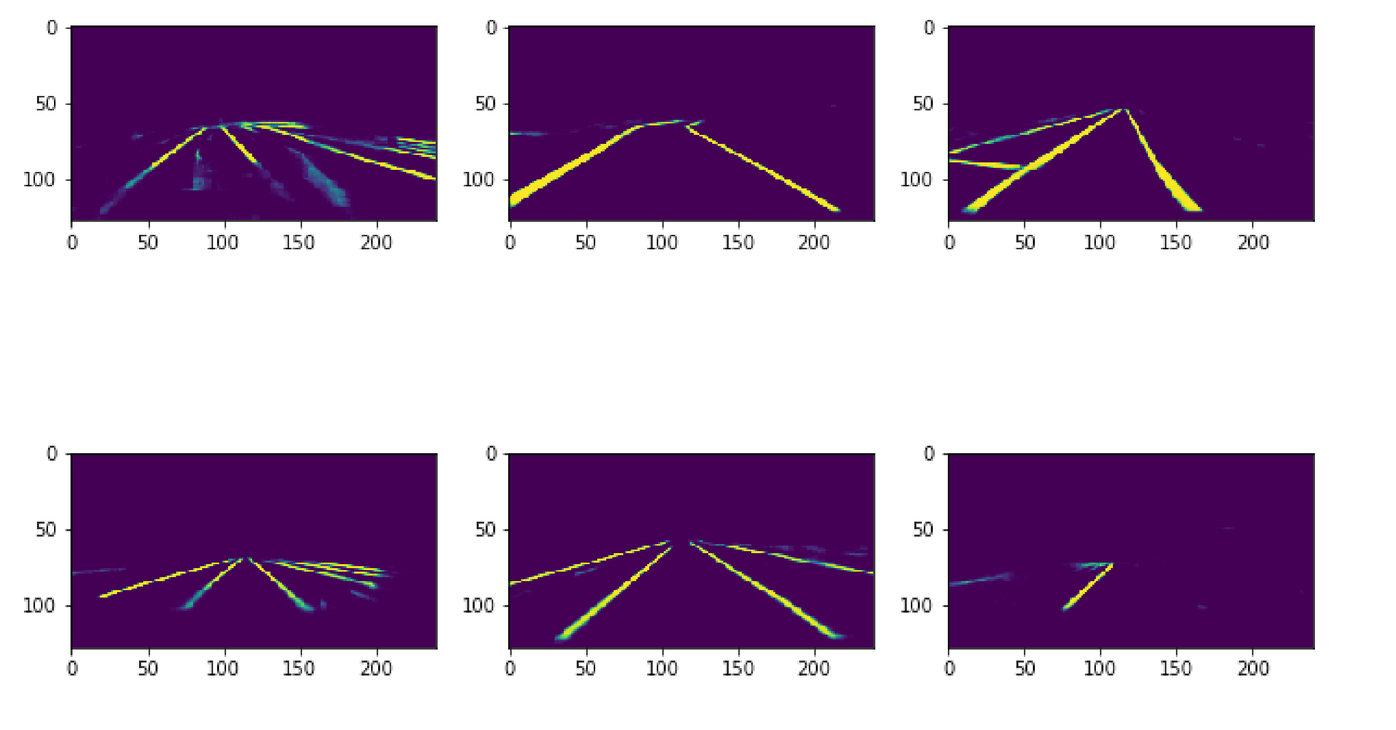

Sample predictions from our network

Segmentation Results

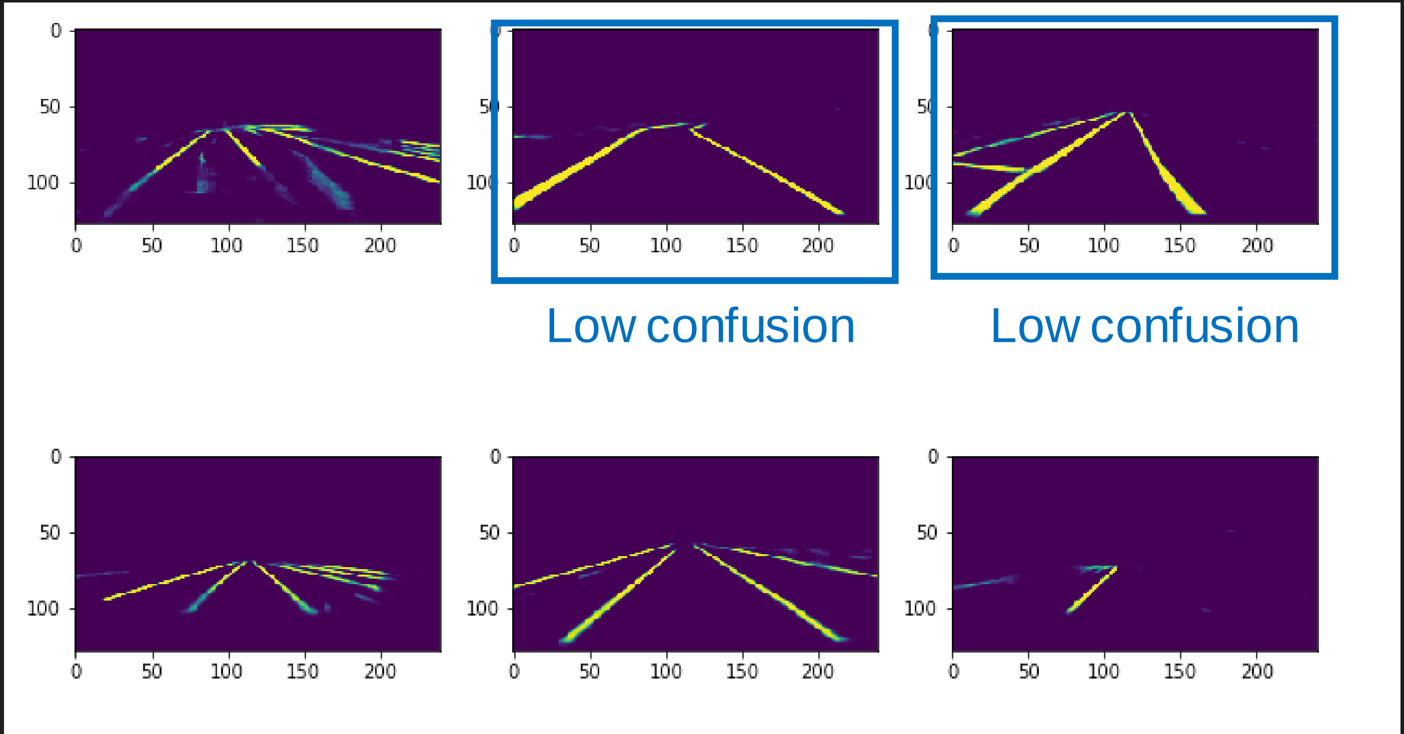

Sample predictions from our network

Segmentation Results

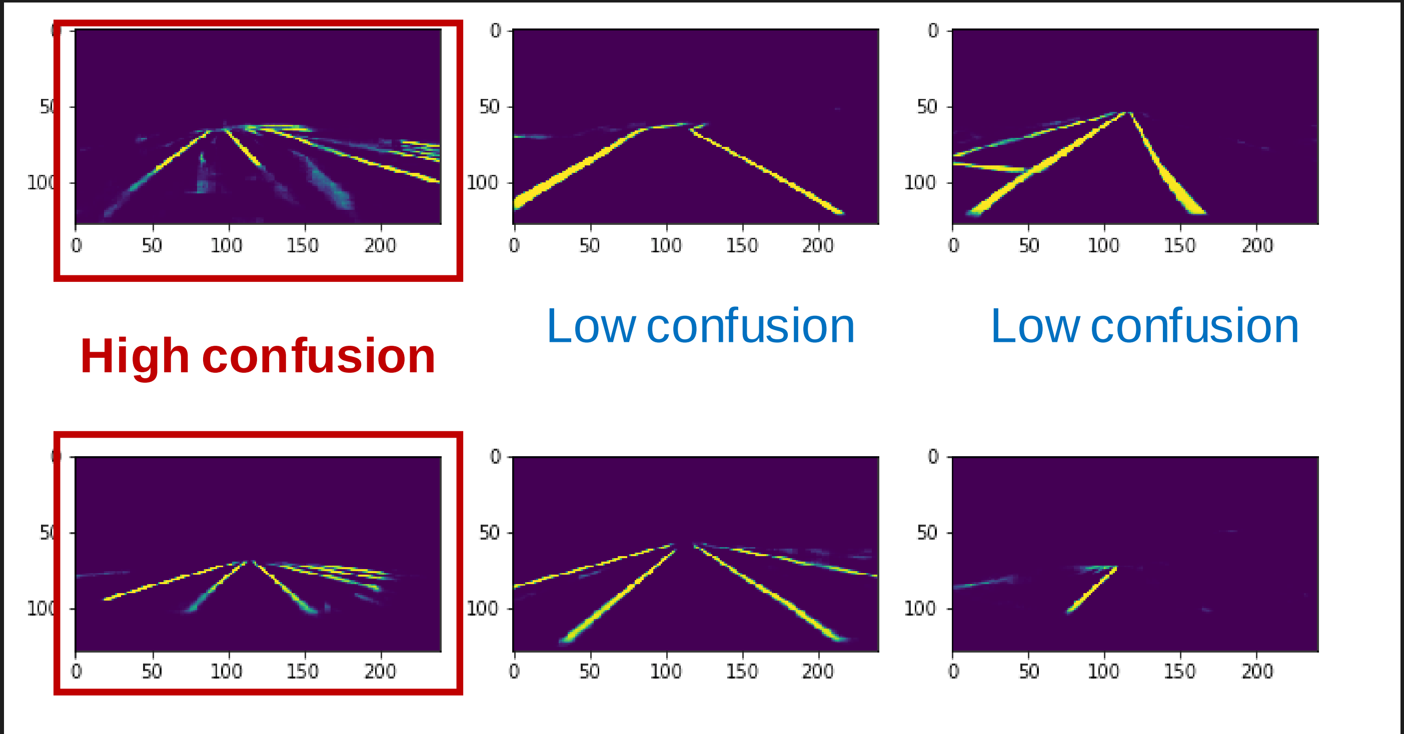

Sample predictions from our network

Detection Results

Detection Experiment Setup

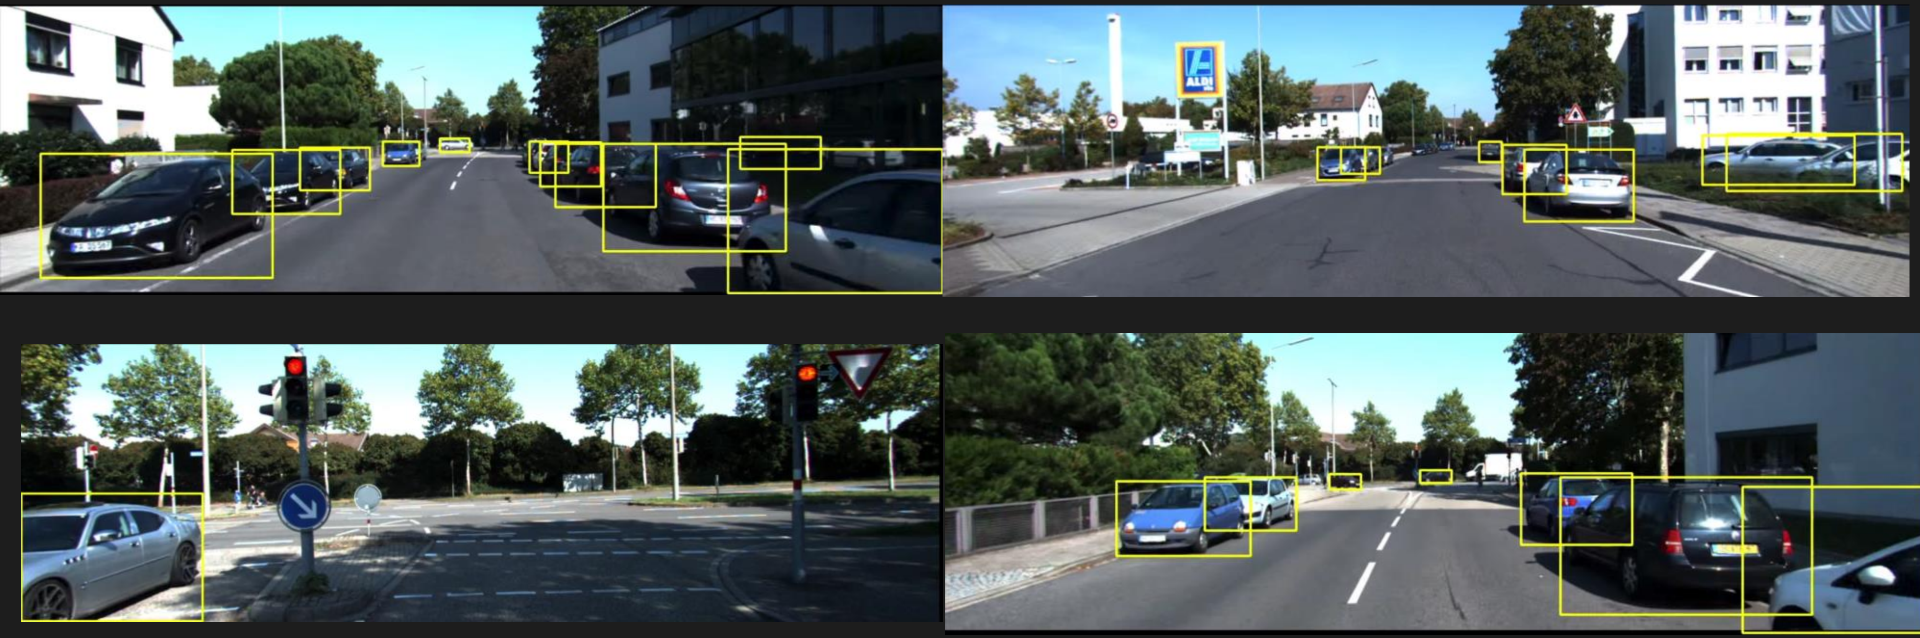

Detection Results

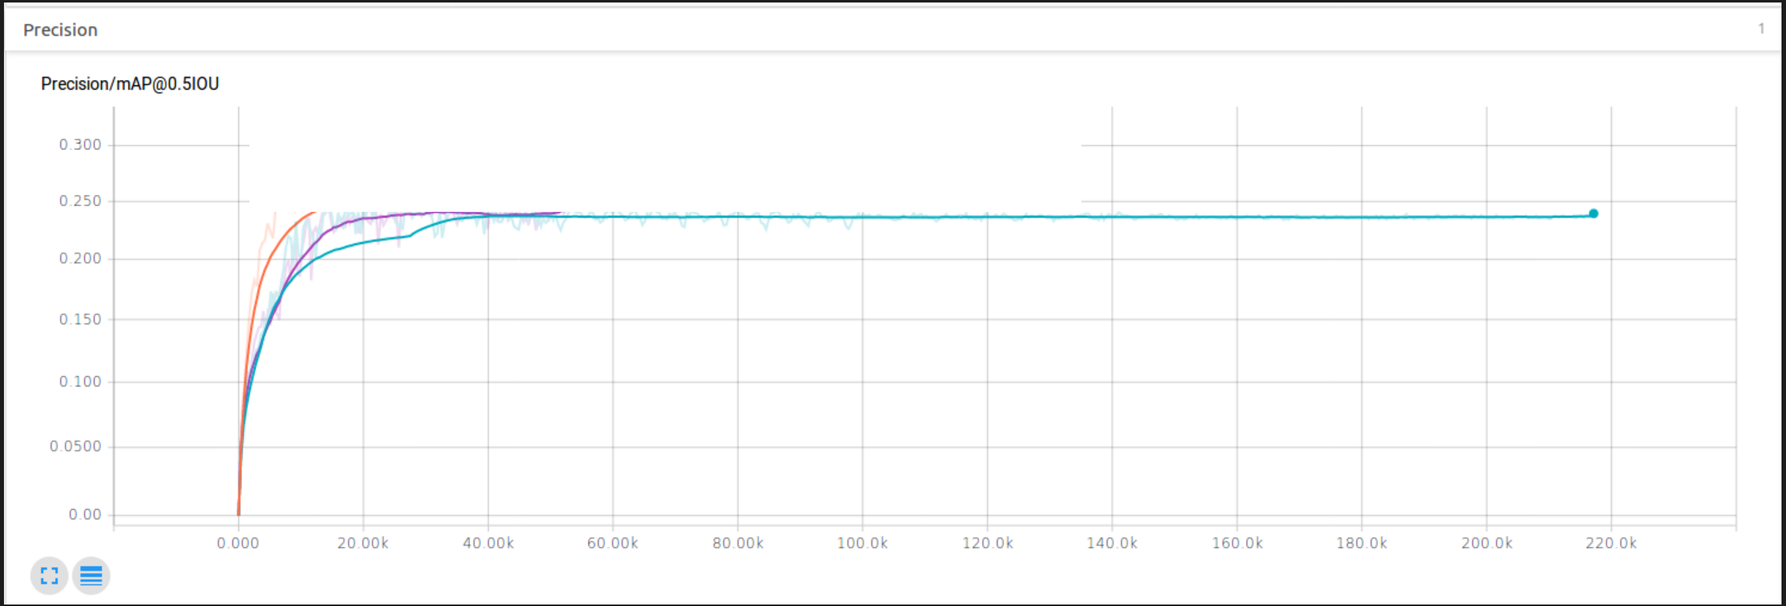

Object detection performance. mAP as as function of epochs, for base model (blue), random strategy (purple) and active strategy (orange).

Detection Results

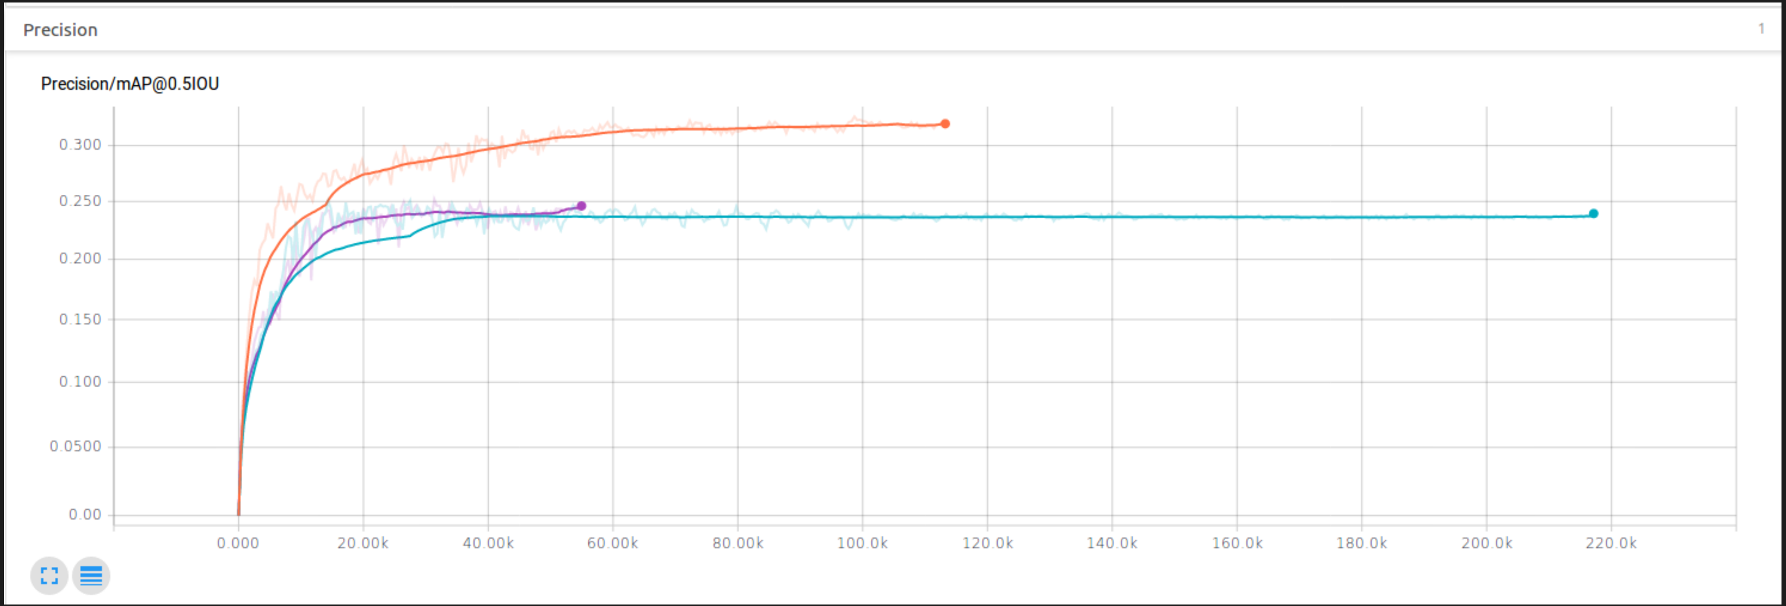

Object detection performance. mAP as as function of epochs, for base model (blue), random strategy (purple) and active strategy (orange).

[Lesnikowski, Plump, 2017]

Detection Results

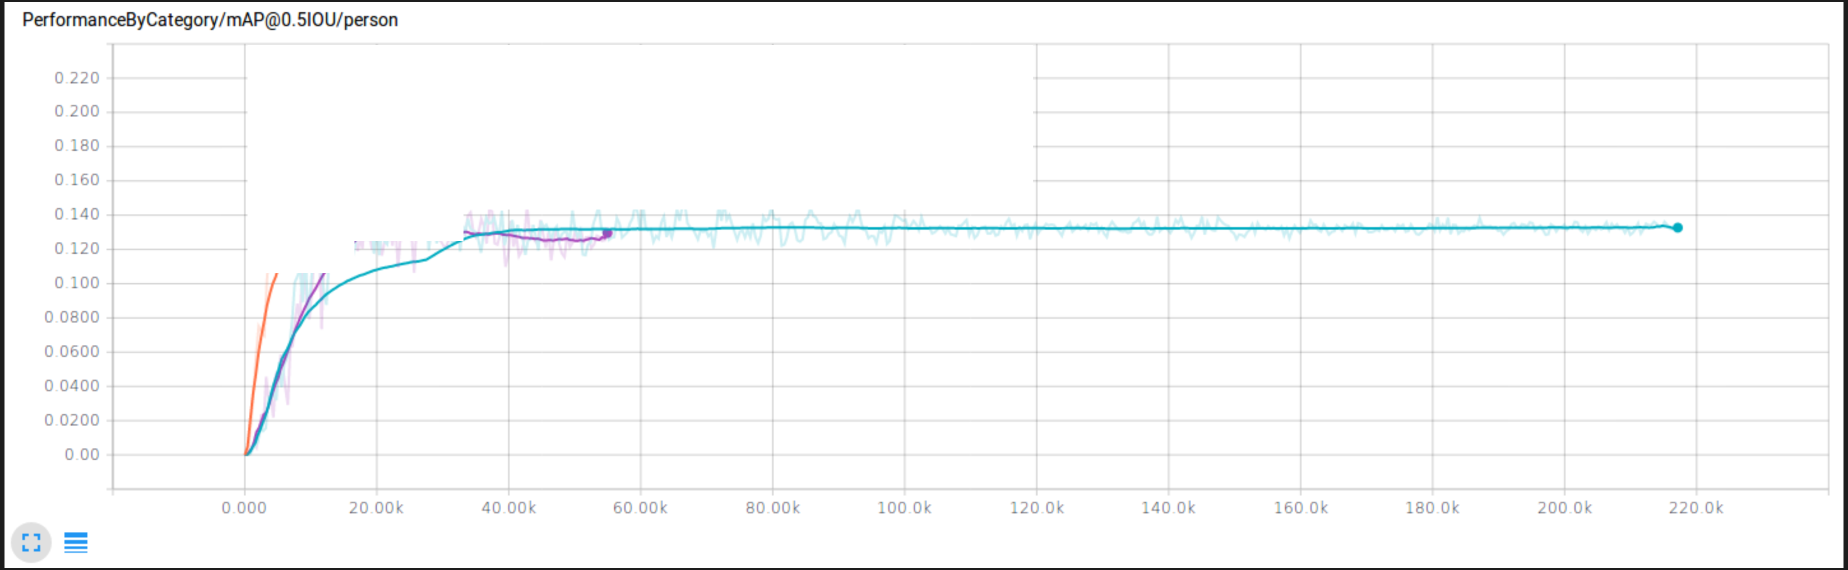

Object detection performance. mAP for persons as as function of epochs, for base model (blue), random strategy (purple) and active strategy (orange).

[Lesnikowski, Plump, 2017]

Detection Results

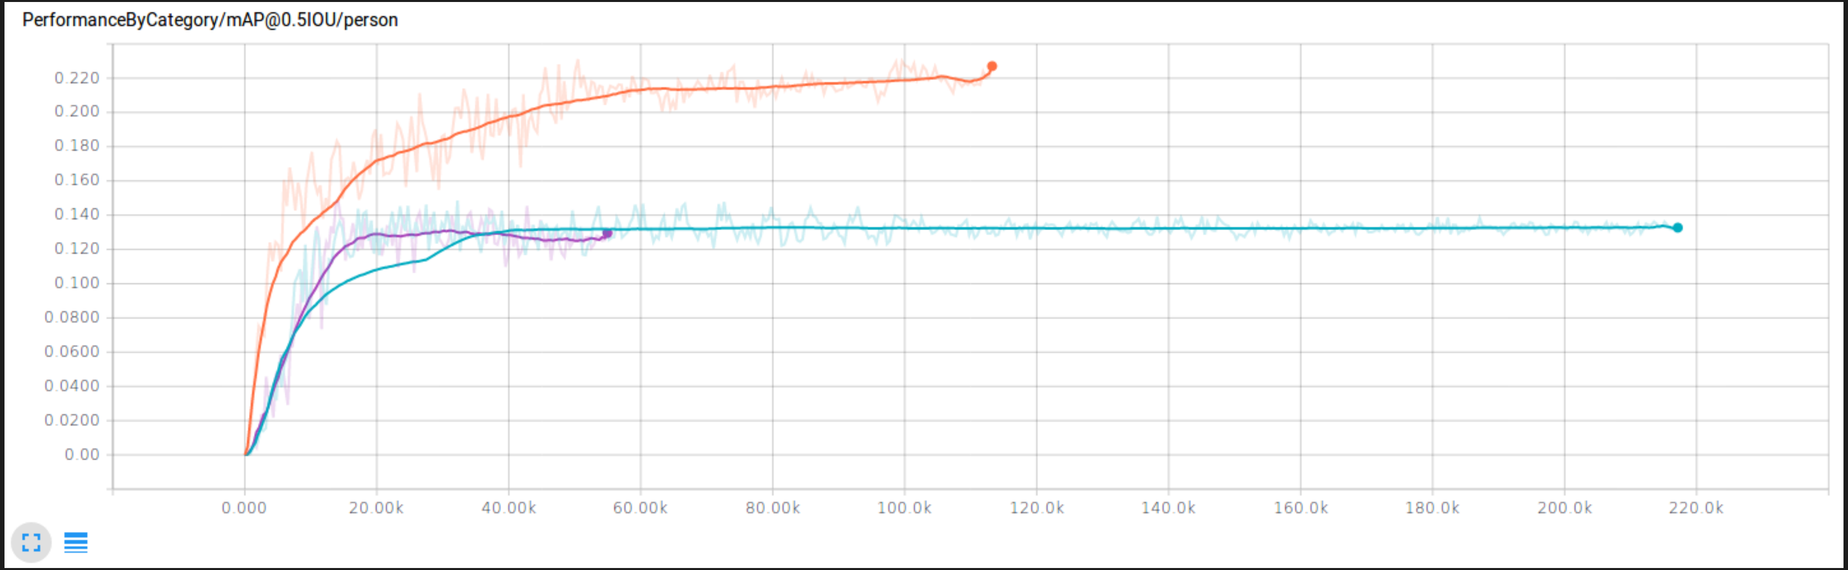

Object detection performance. mAP for persons as as function of epochs, for base model (blue), random strategy (purple) and active strategy (orange).

[Lesnikowski, Plump, 2017]

Detection Acquisition Examples

Sample predictions from our object detection model.

Detection Acquisition Examples

Sample predictions from our object detection model.

Detection Acquisition Examples

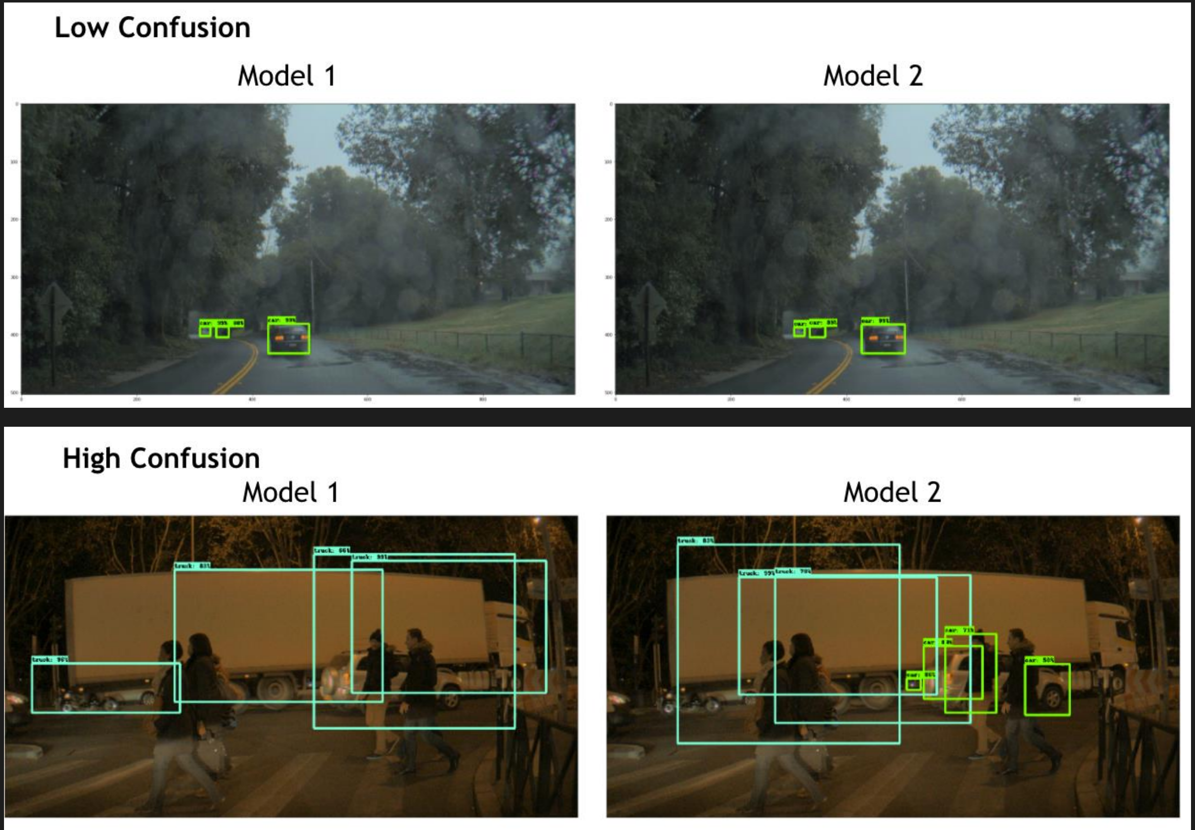

Sample predictions from our object detection model.

Bayesian Deep Active Learning for Autonomous Vehicles

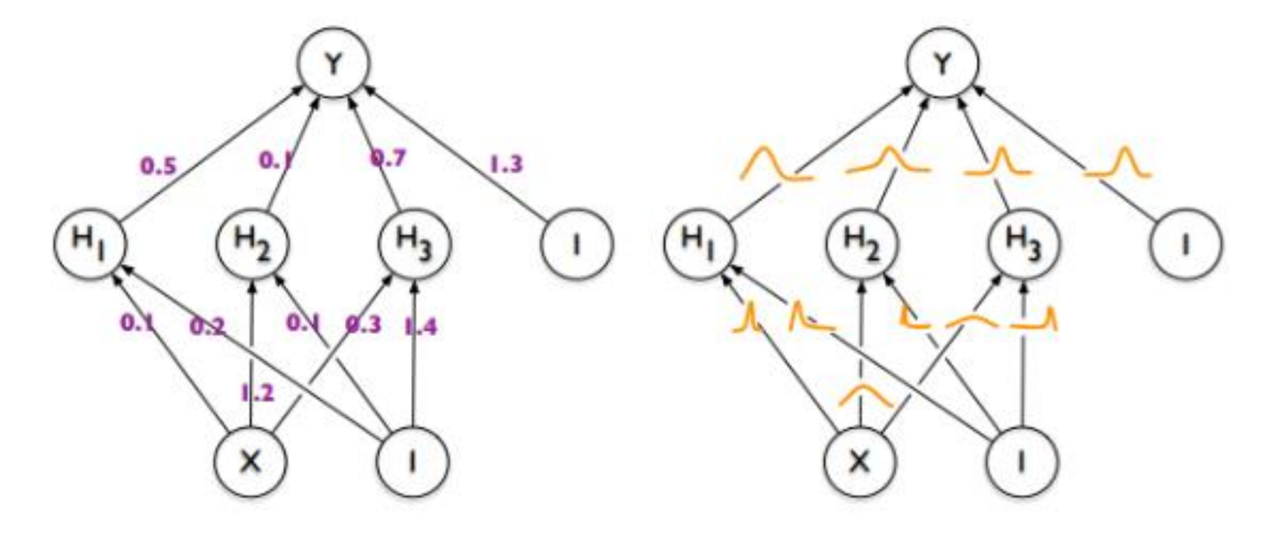

[Zeng, Lesnikowski, Alvarez, 2018]

TensorFlow Probability

MNIST

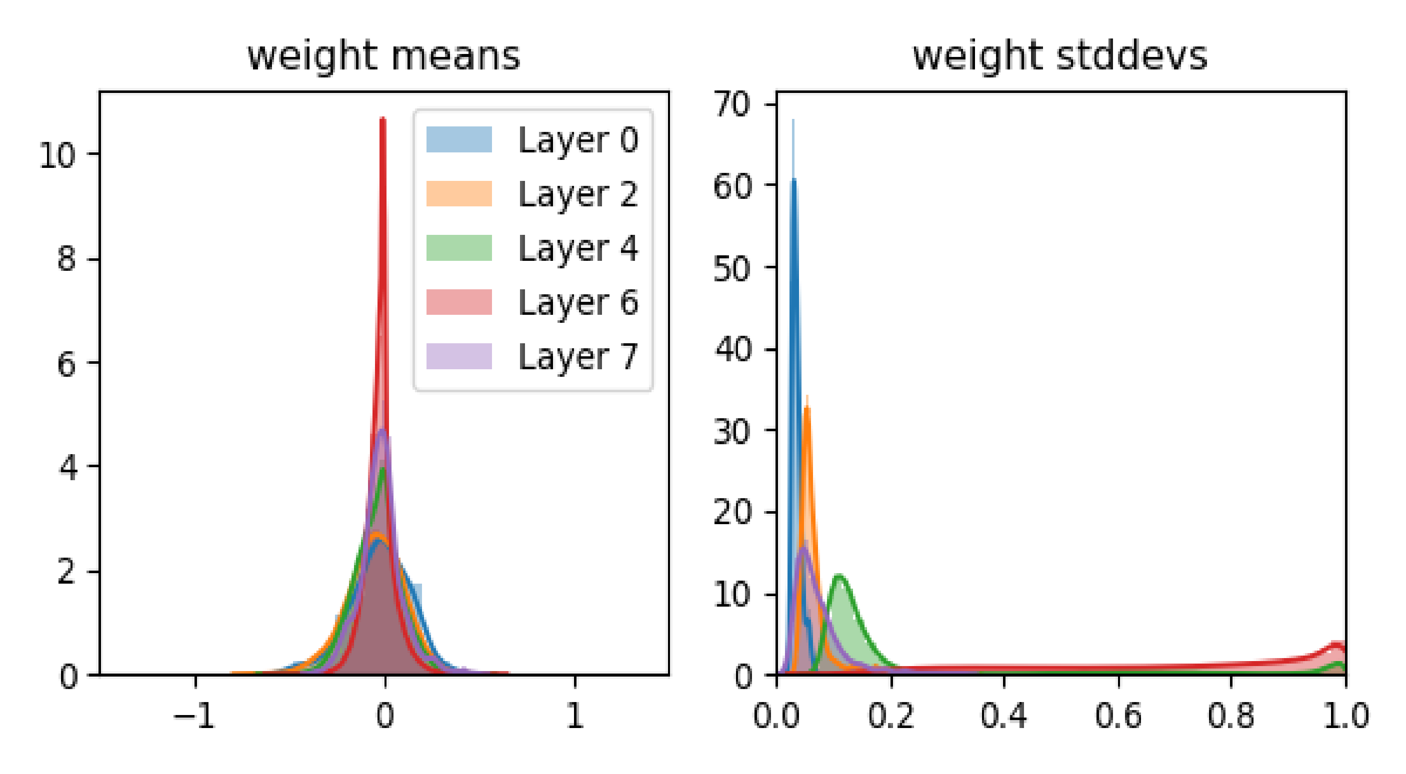

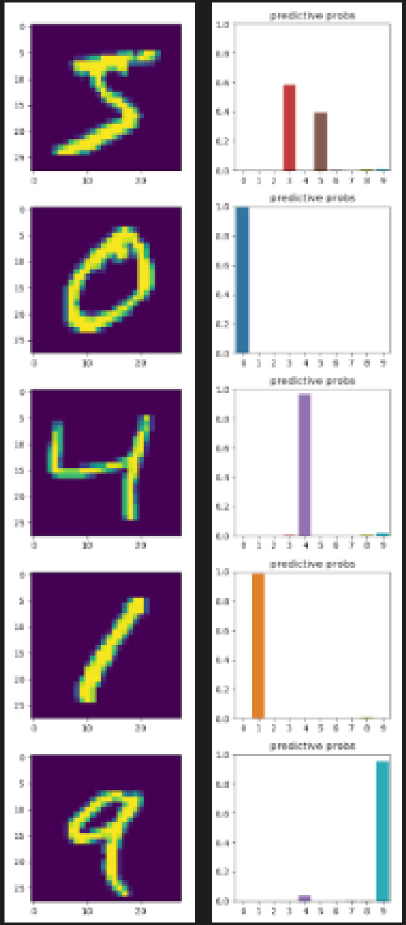

Network: ResNet18 **Training:** 1000 epochs **Testing Accuracy:** 81%

Example incorporated into the official TFP repository.

TensorFlow Probability

CIFAR-10

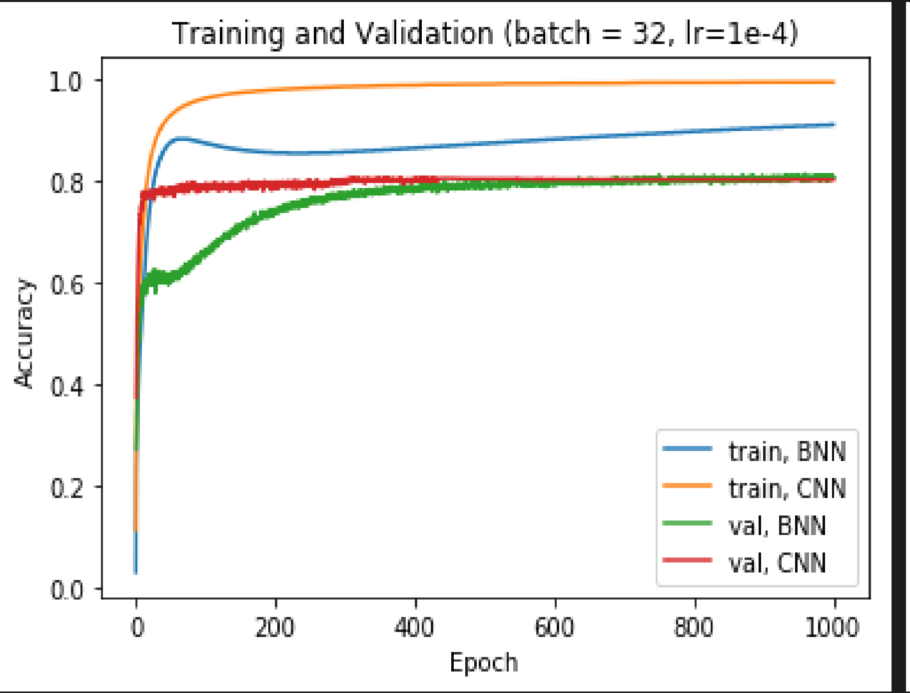

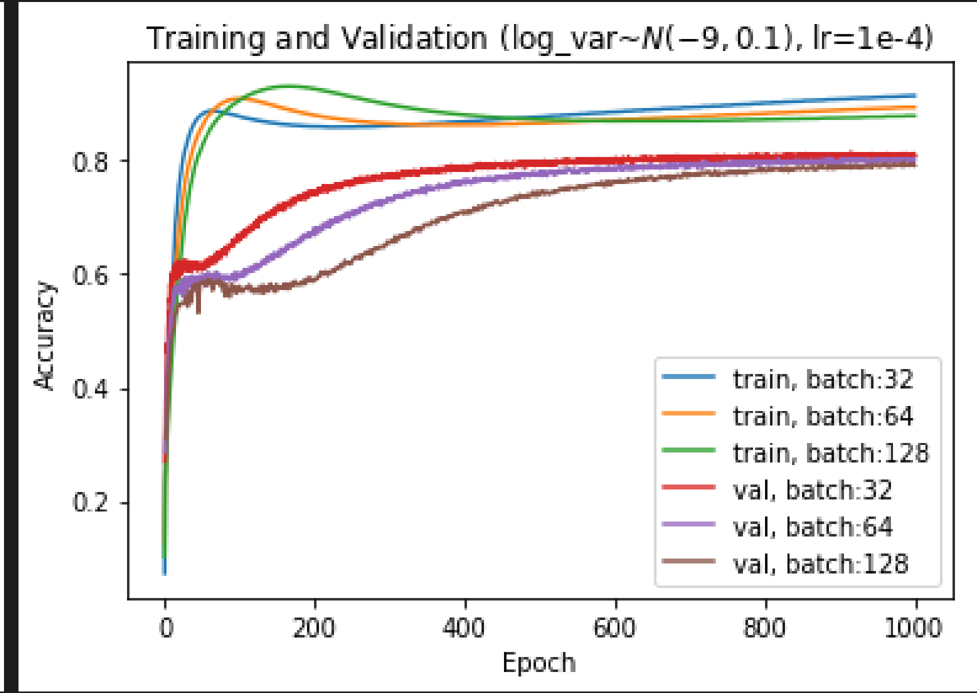

Network: ResNet18 **Training:** 1000 epochs **Testing Accuracy:** 81%

[*Zeng, Lesnikowski, Alvarez, NeurIPS BDL workshop 2018*]

Example incorporated into the official TFP repository.

Bayesian Regularized Ensembles for Segmentation

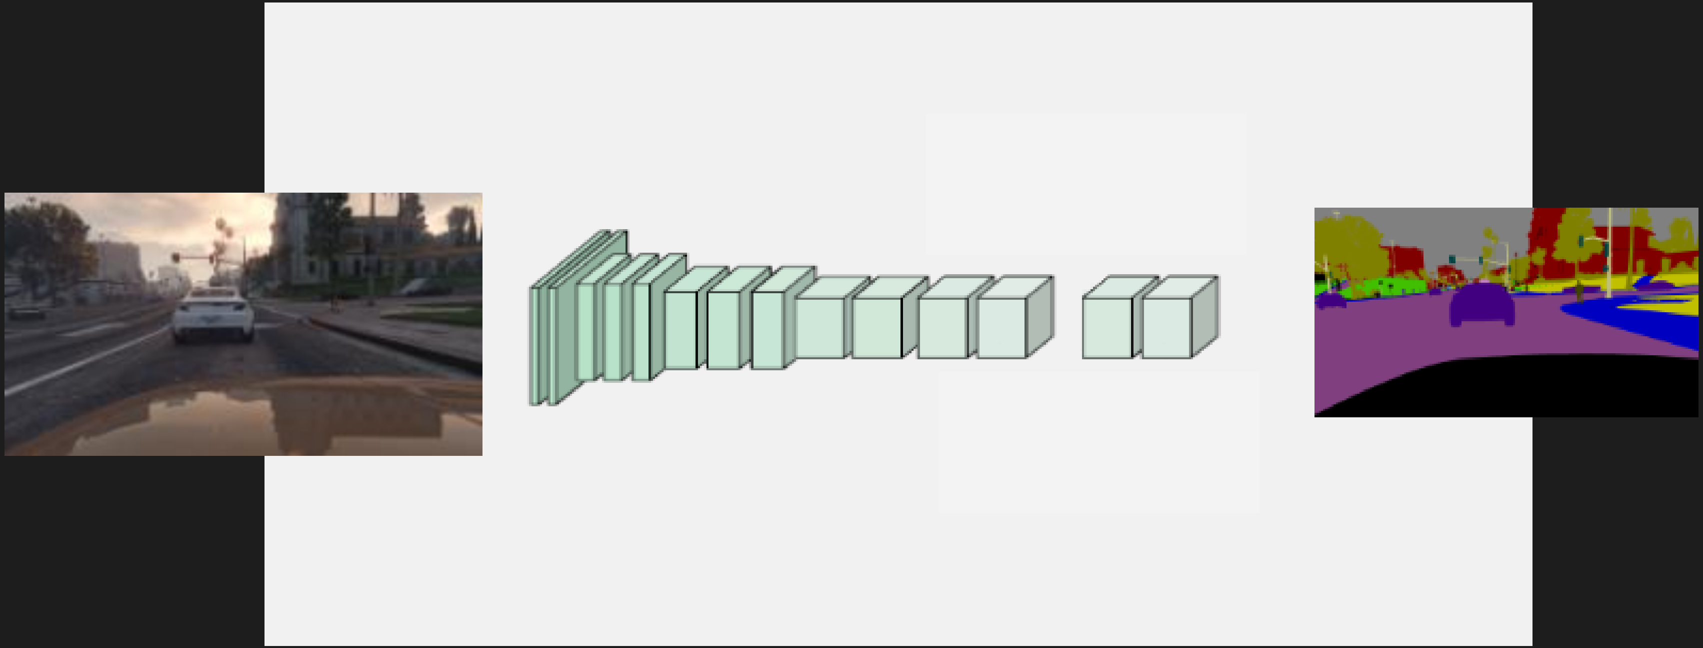

[*Chitta, Alvarez, Lesnikowski NeurIPS BDL workshop, 2018*]

Bayesian Regularized Ensembles for Segmentation

[Chitta, Alvarez, Lesnikowski NeurIPS BDL workshop, 2018]

Bayesian Regularized Ensembles for Segmentation

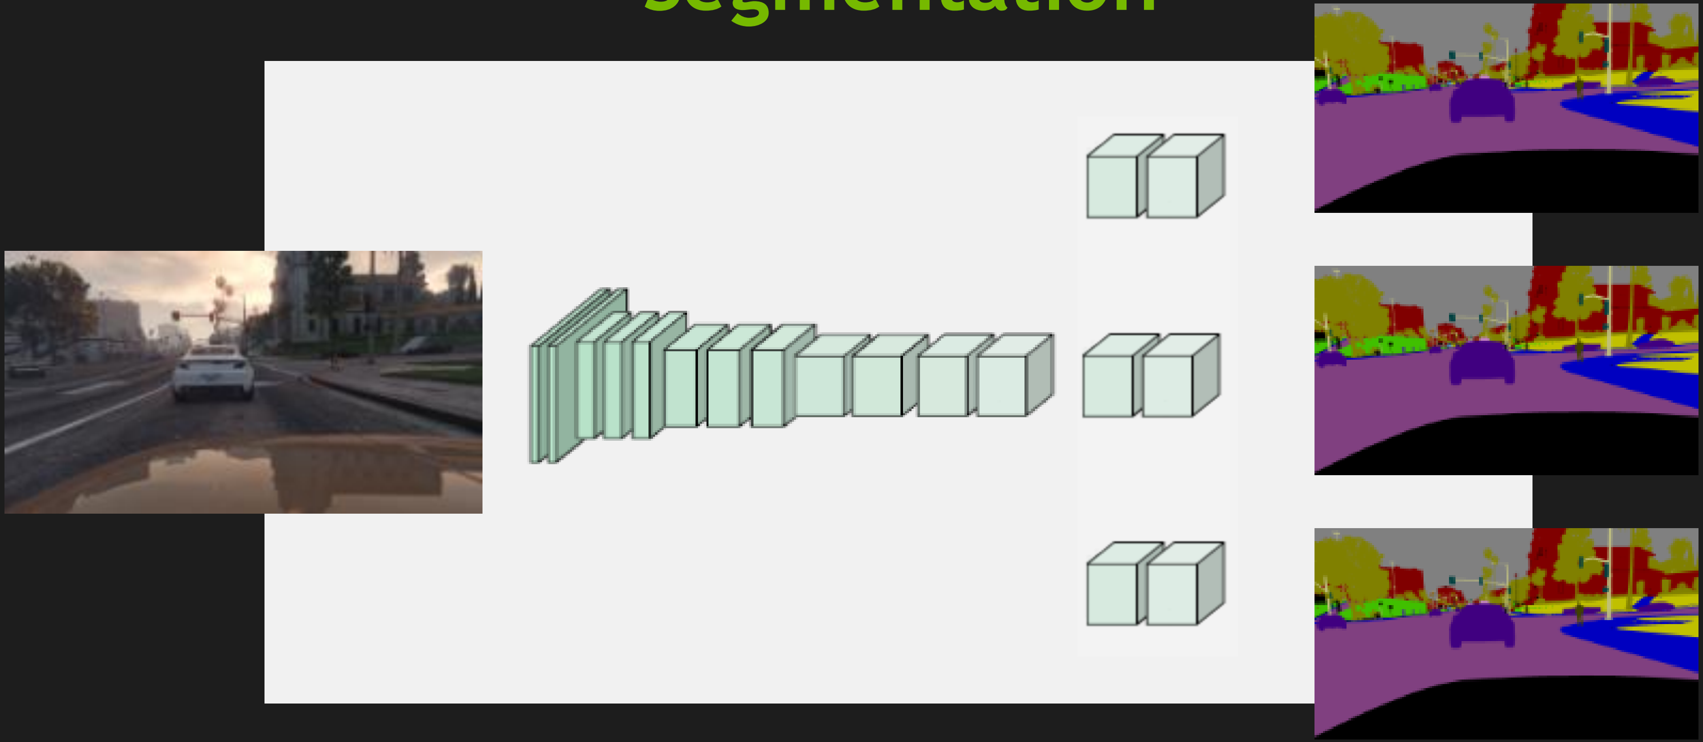

[*Chitta, Alvarez, Lesnikowski NeurIPS BDL workshop, 2018*]

Segmentation: Experimental Setup

Pretrain on CityScapes, evaluate active learning on BDD100k

Segmentation: Experimental Setup

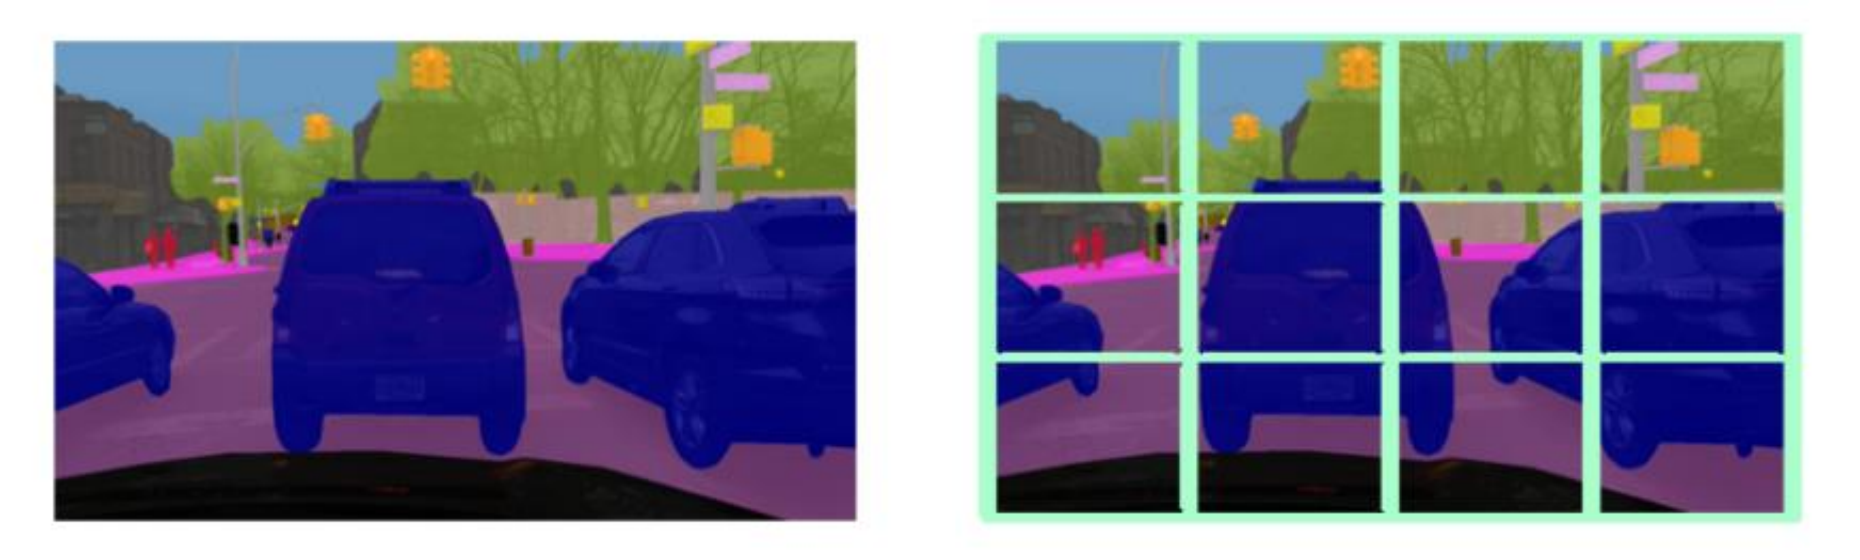

Partition image into a 4x3 grid, giving 12 crops

Results: Class-wise Improvements

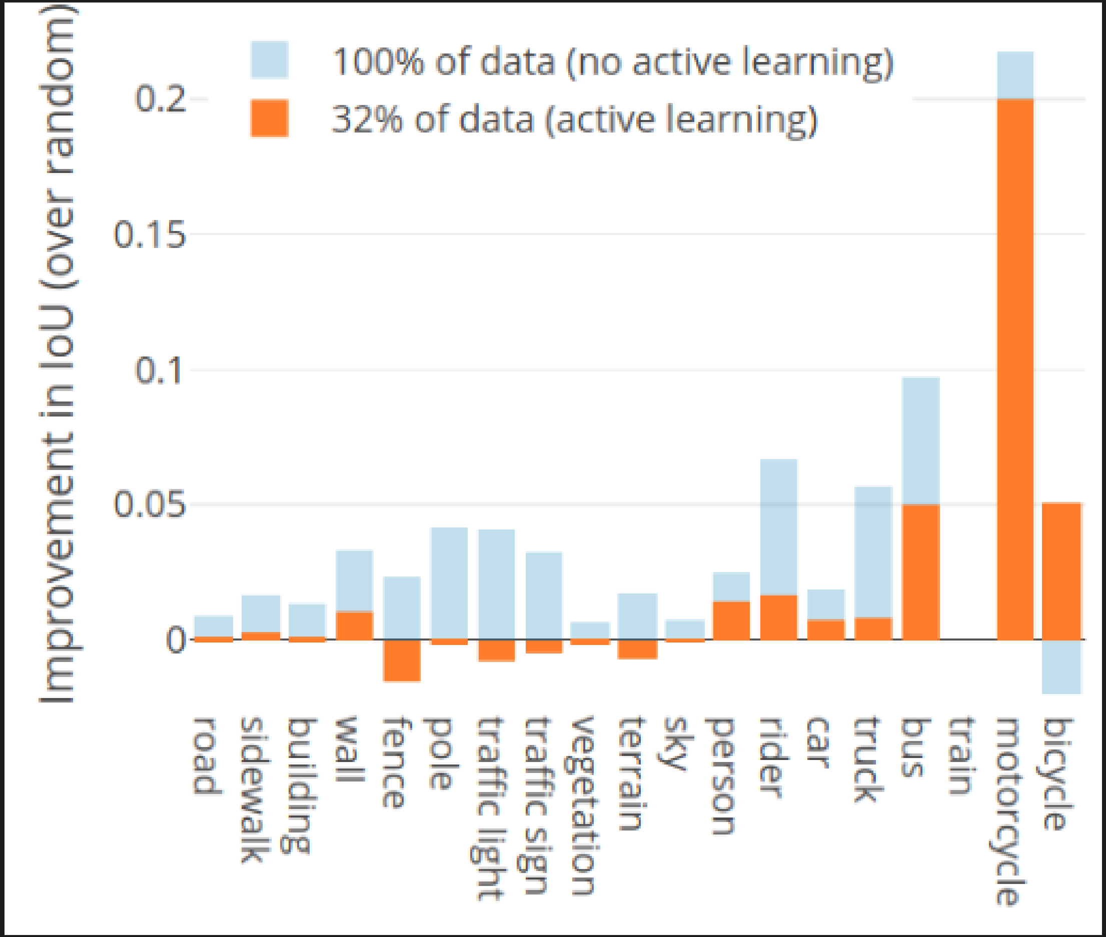

Results: Targeting Classes

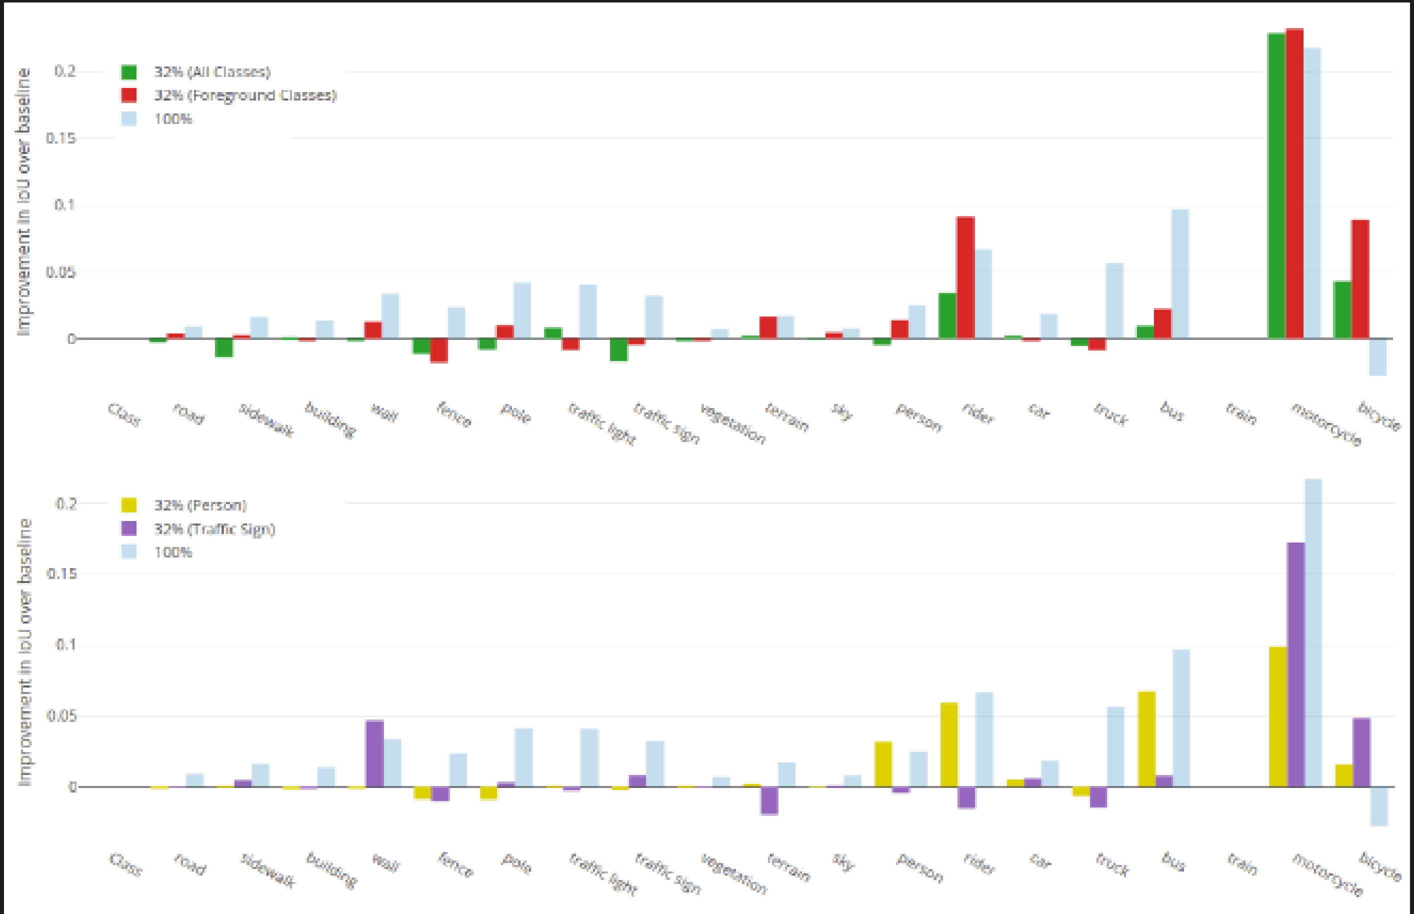

[Chitta, Alvarez, Lesnikowski forthcoming]

Classification Results

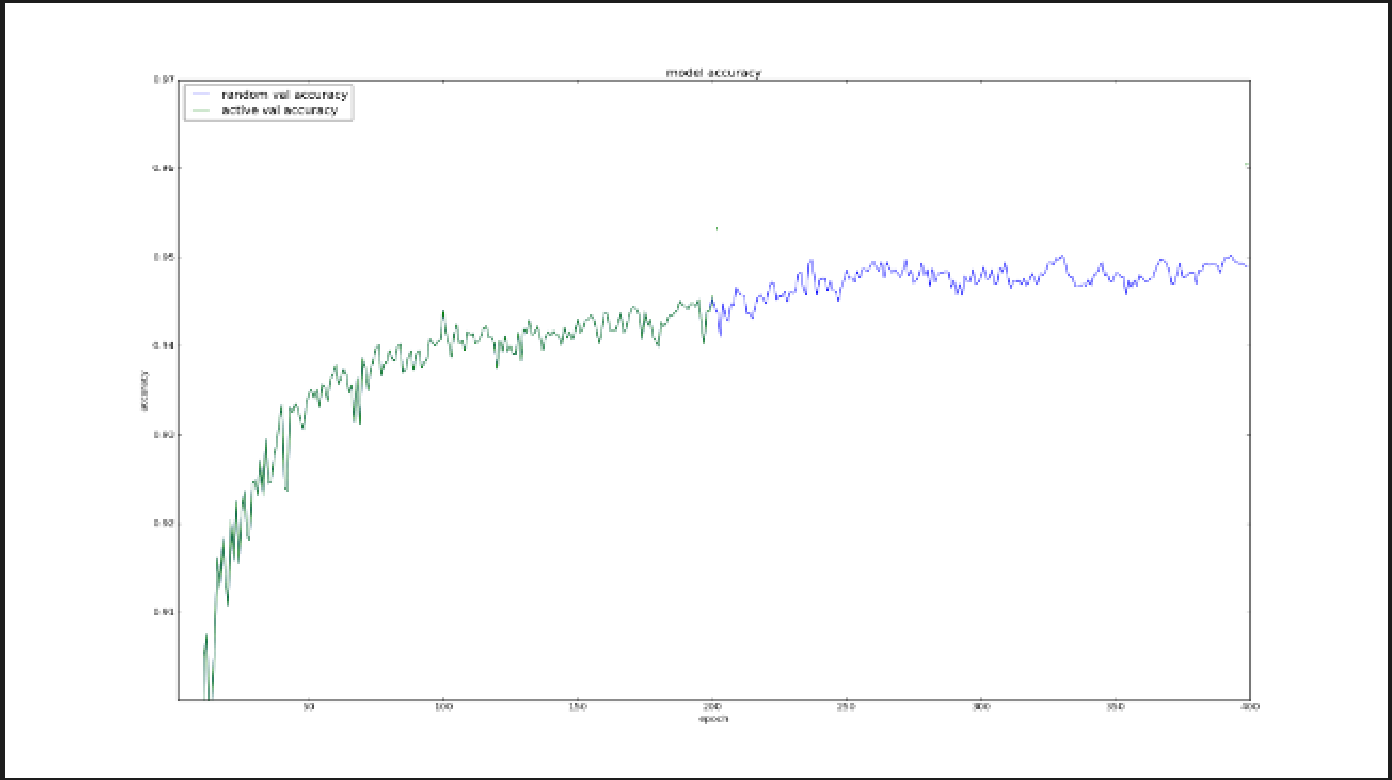

Classification results after one active learning cycle with an entropy policy.
Random baseline (blue) versus active strategy (green).

Classification Results

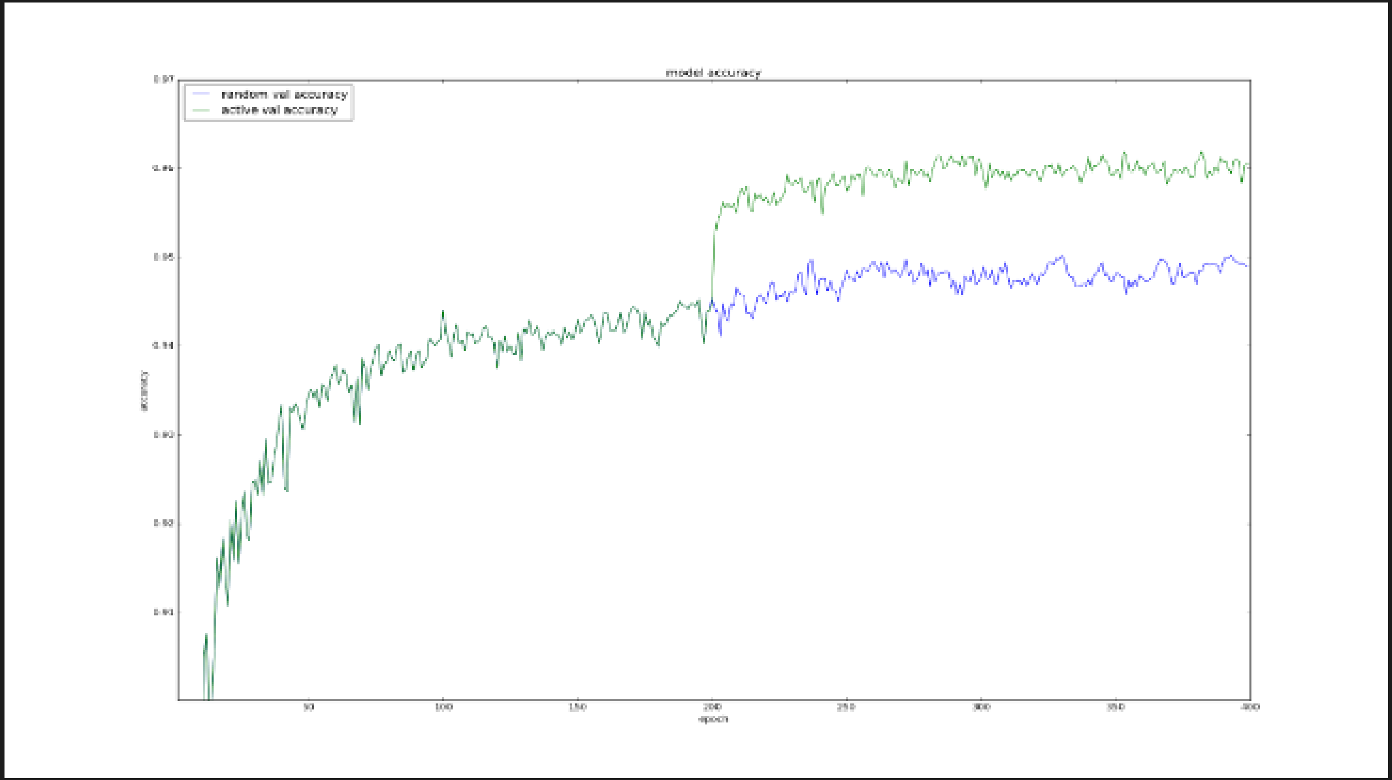

Classification results after one active learning cycle with an entropy policy. Random baseline (blue) versus active strategy (green).

Classification Results

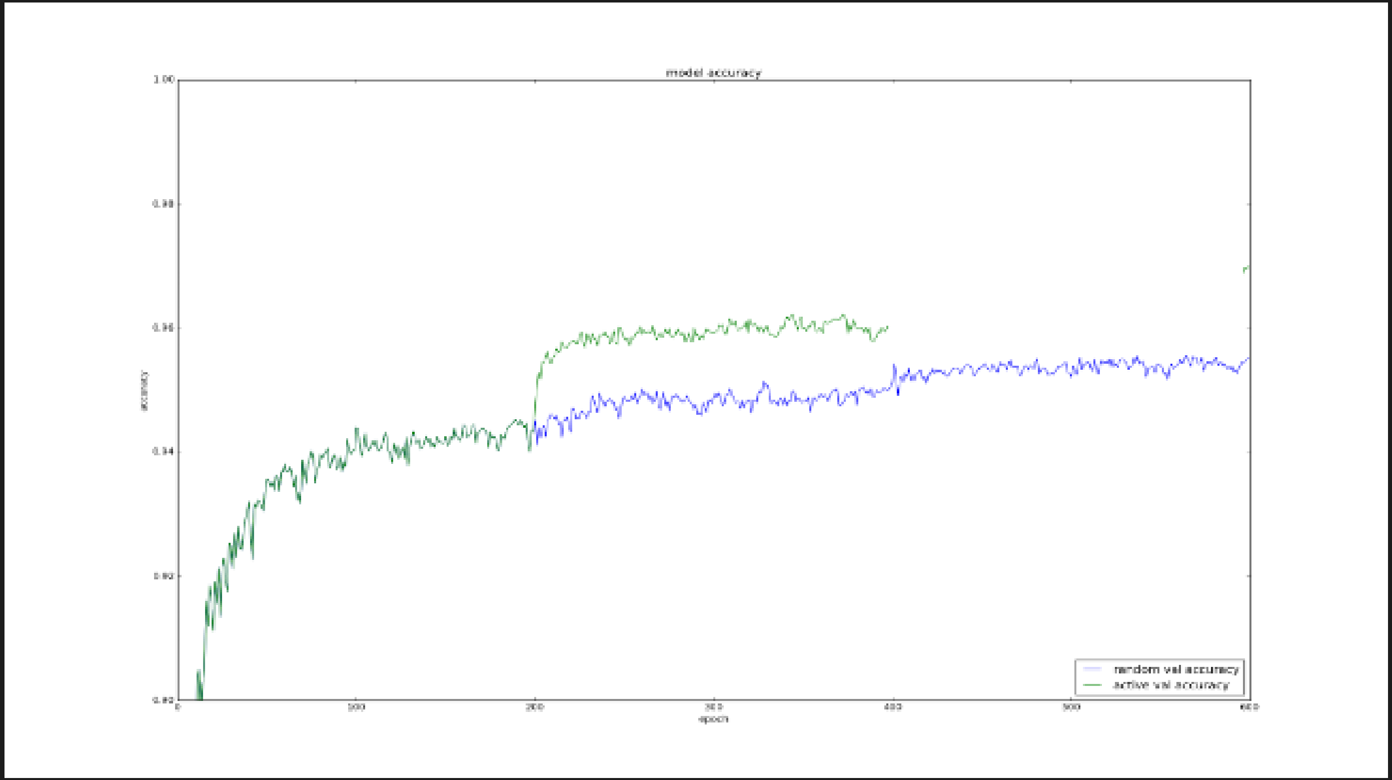

Classification results after two active learning cycles with an entropy policy.
Random baseline (blue) versus active strategy (green).

Classification Results

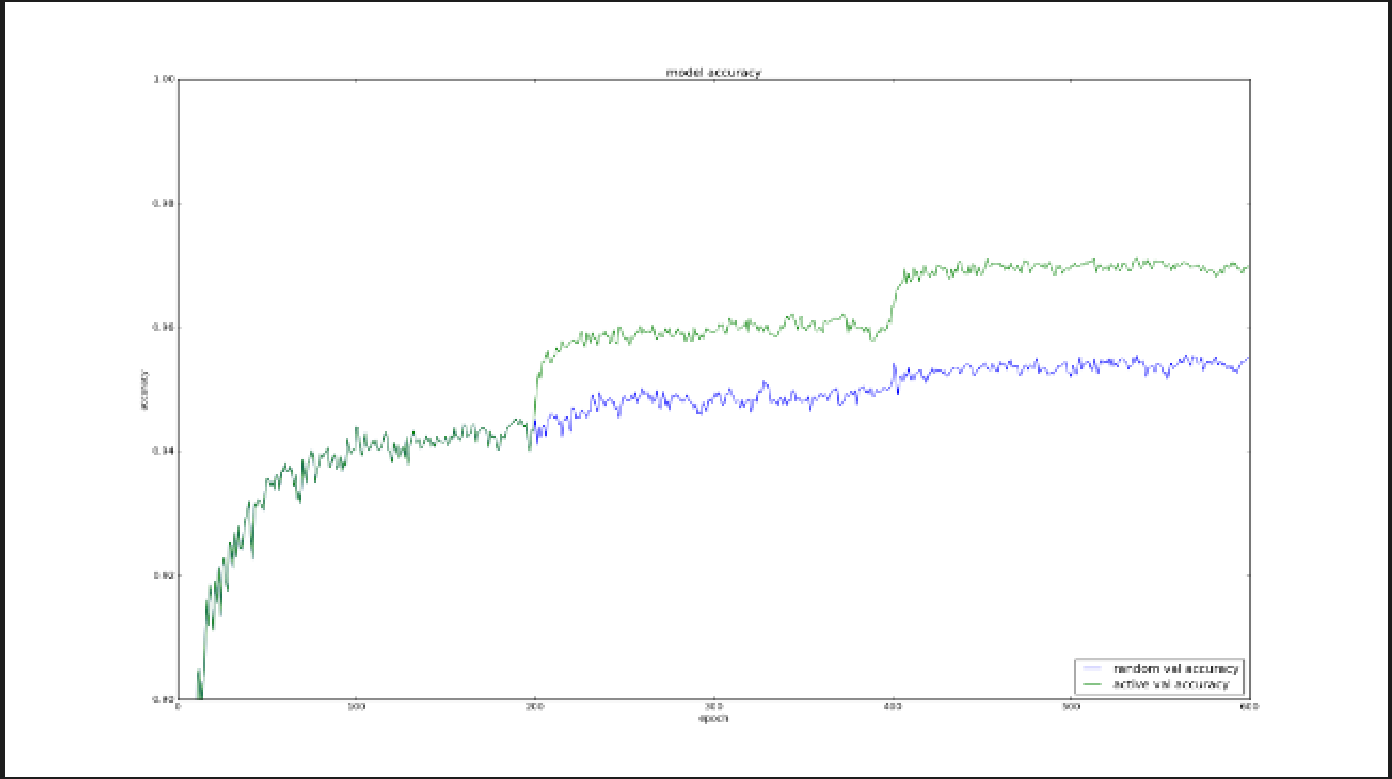

Classification results after two active learning cycles with an entropy policy.
Random baseline (blue) versus active strategy (green).

Classification Results

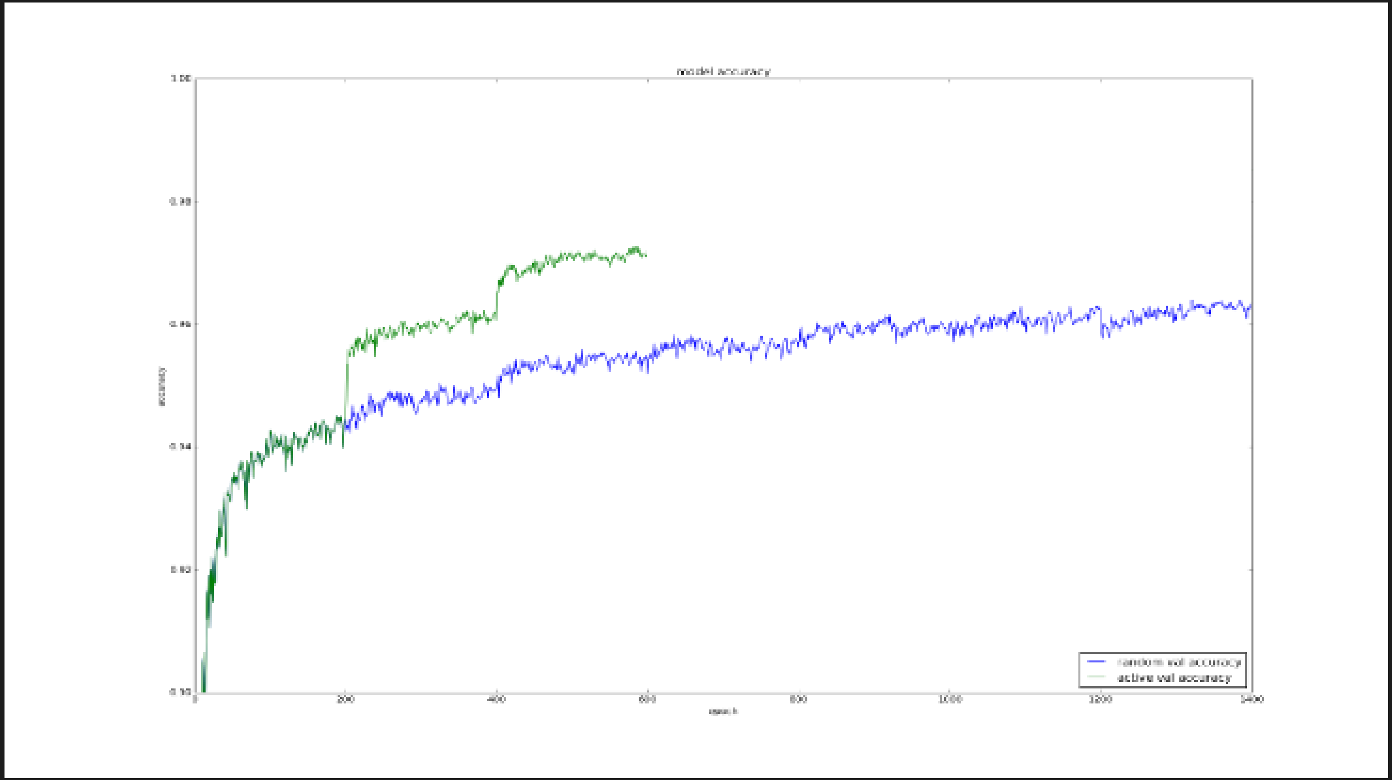

Classification results after six active learning cycles with an entropy policy.
Random baseline (blue) versus active strategy (green).

Classification Results

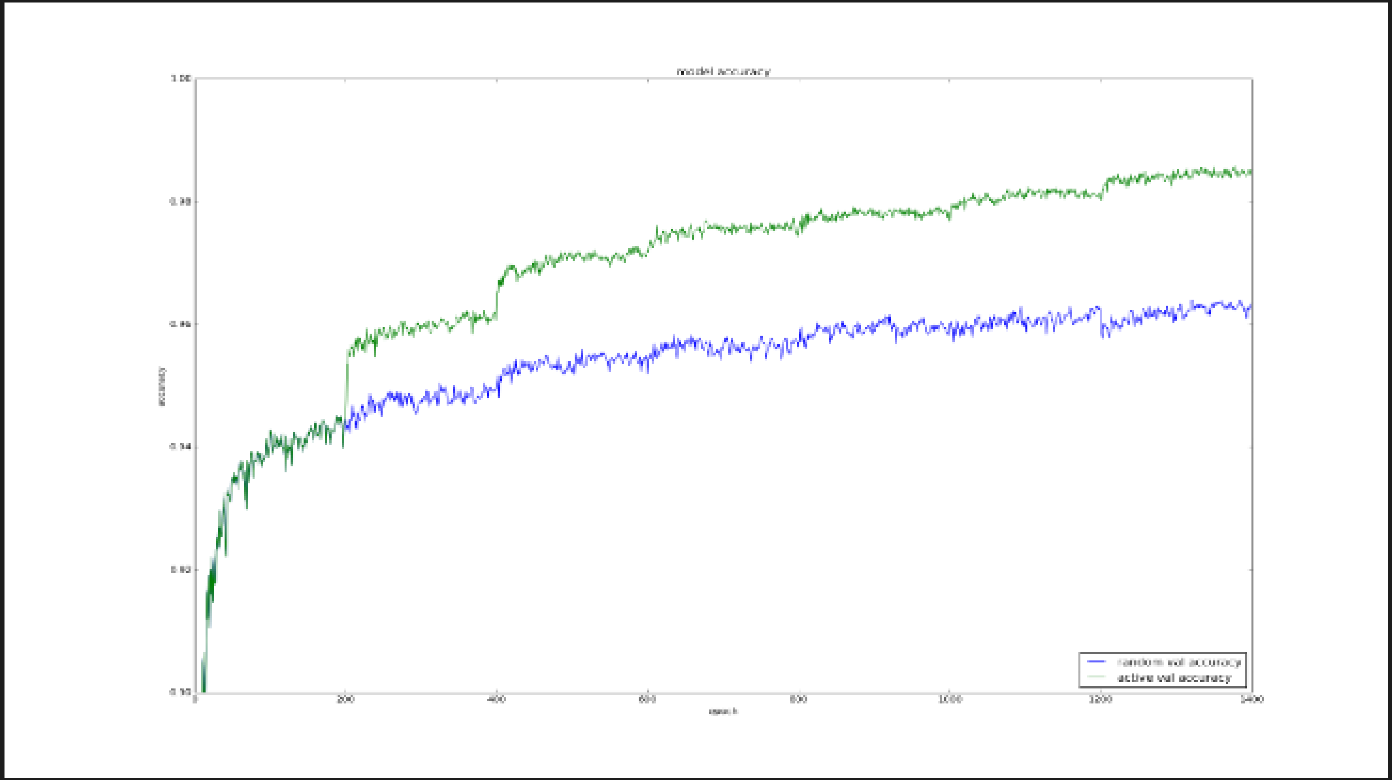

Classification results after six active learning cycles with an entropy policy.
Random baseline (blue) versus active strategy (green).

[Lesnikowski, Farabet, 2017]

Classification Results

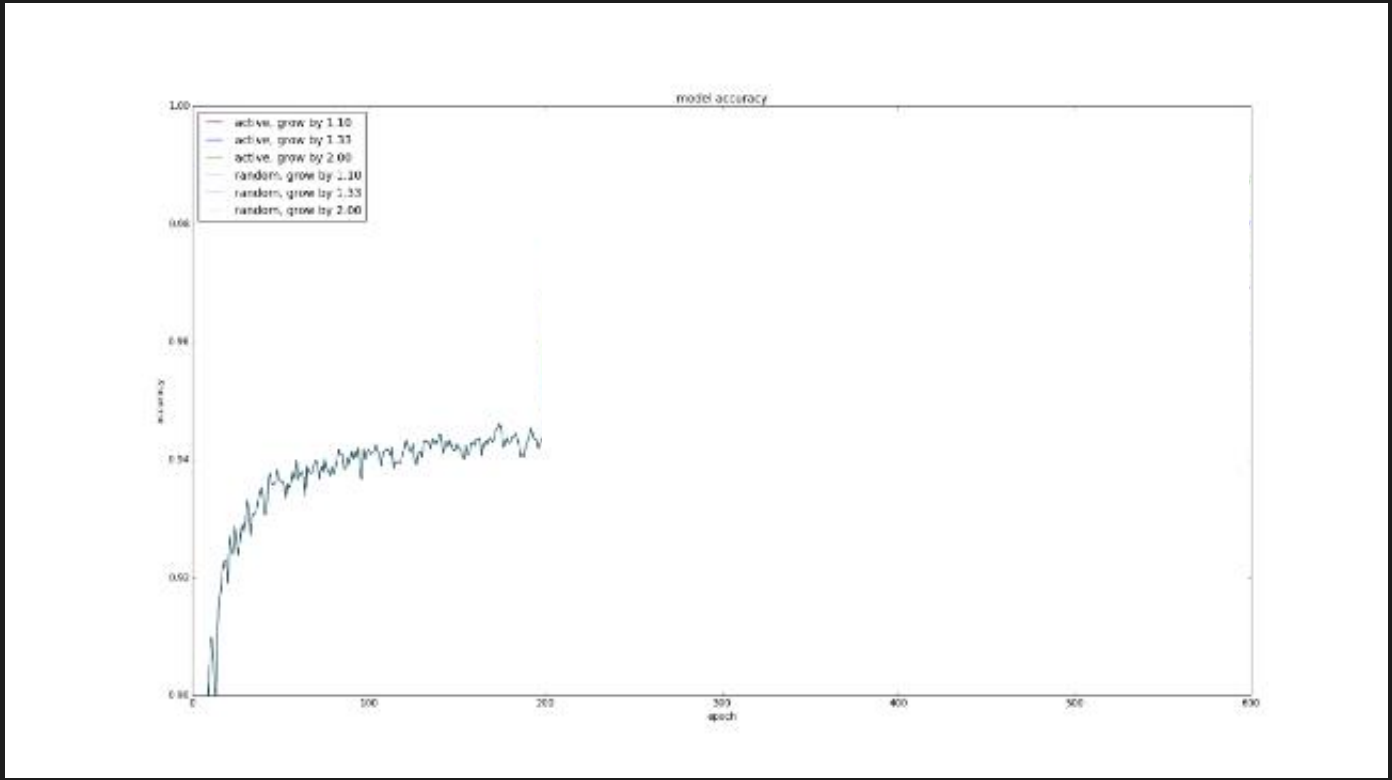

The effect of varying the growth parameter in active learning strategies.

Classification Results

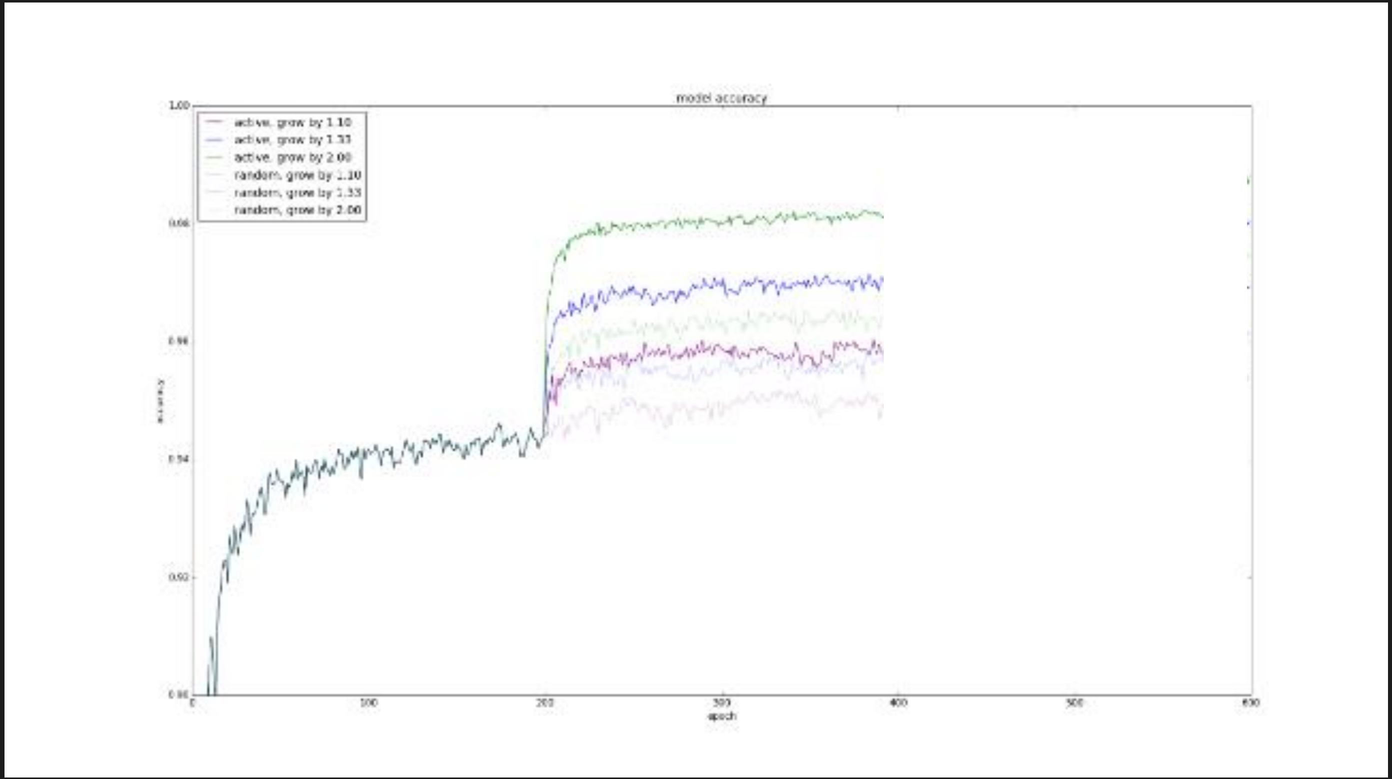

The effect of varying the growth parameter in active learning strategies.
Adding 10% (purple), 33% (blue) and 100% (green).

Classification Results

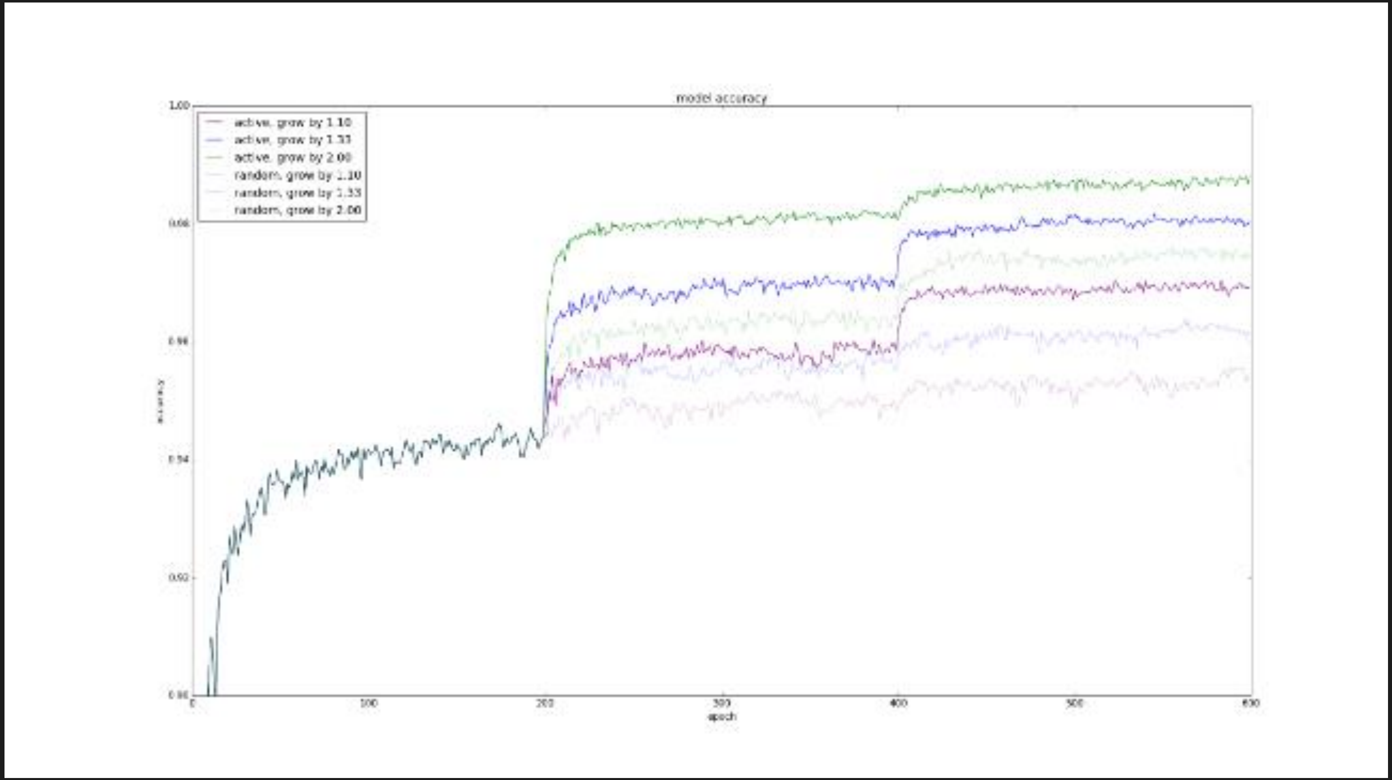

The effect of varying the growth parameter in active learning strategies. Adding 10% (purple), 33% (blue) and 100% (green).

[Lesnikowski, Farabet, 2017]

Classification Results

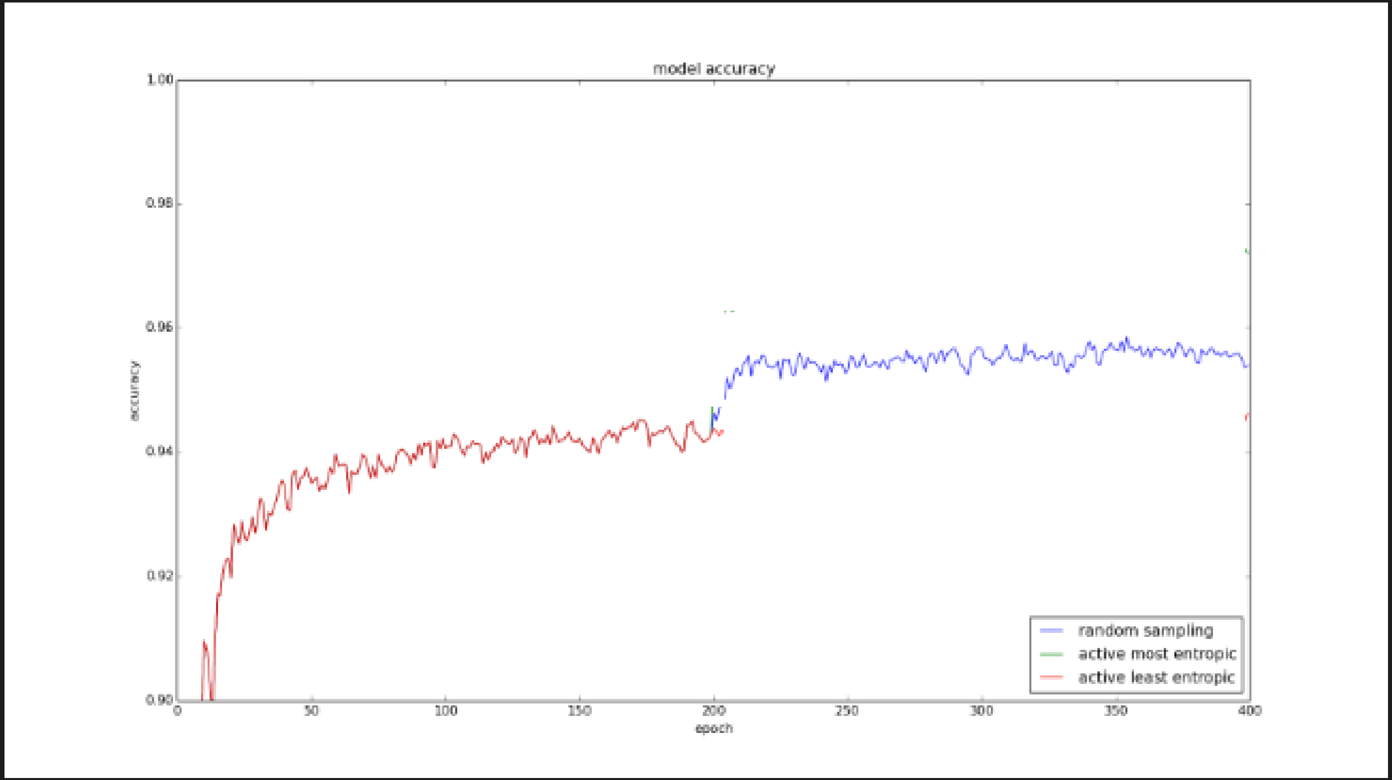

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue.

Classification Results

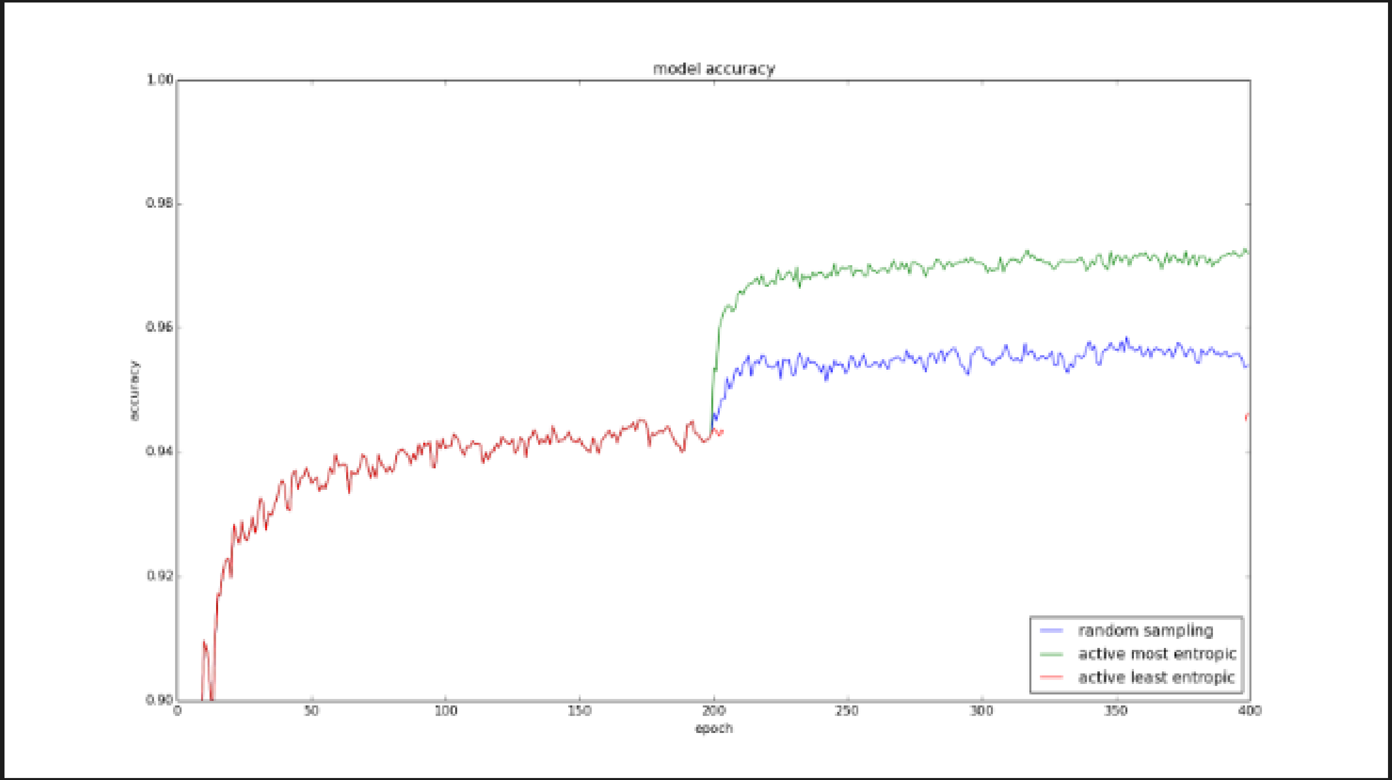

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue.

Classification Results

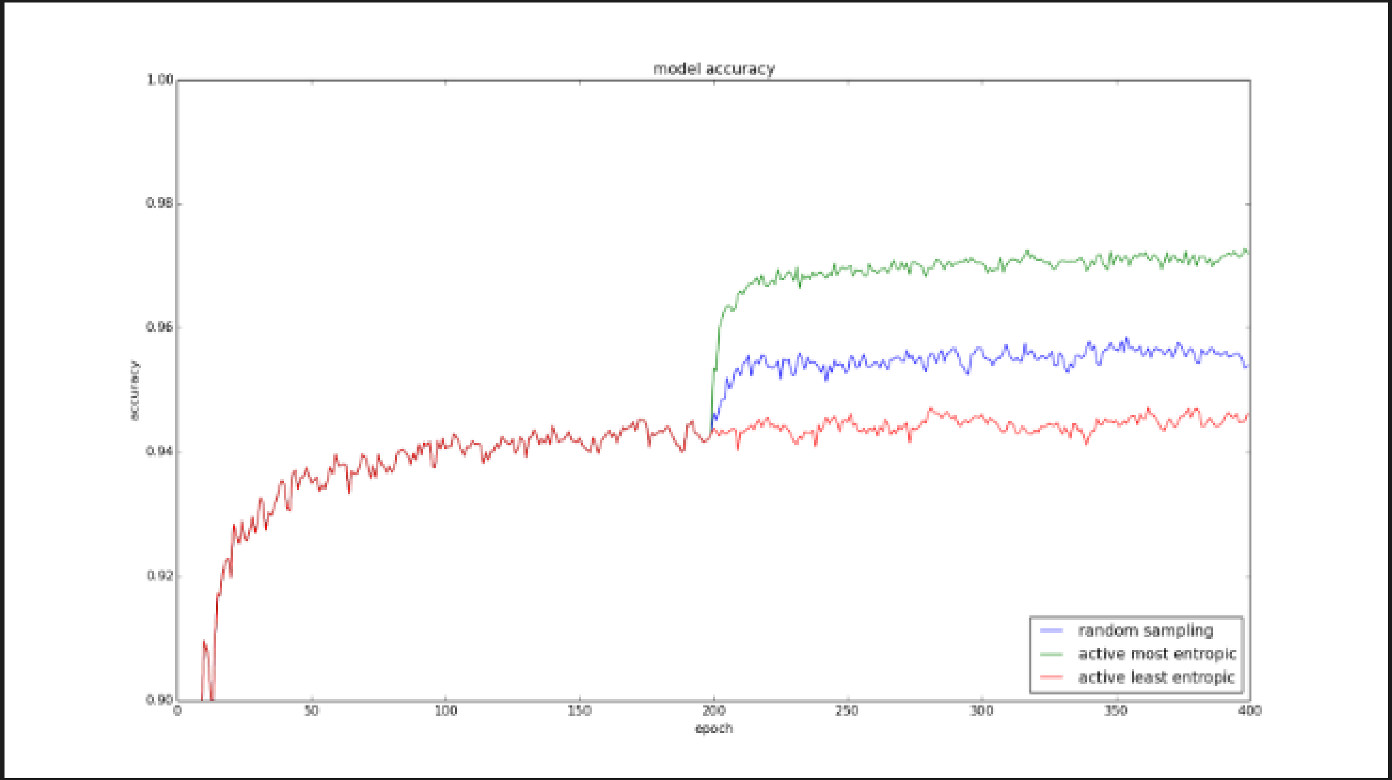

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue.

Classification Results

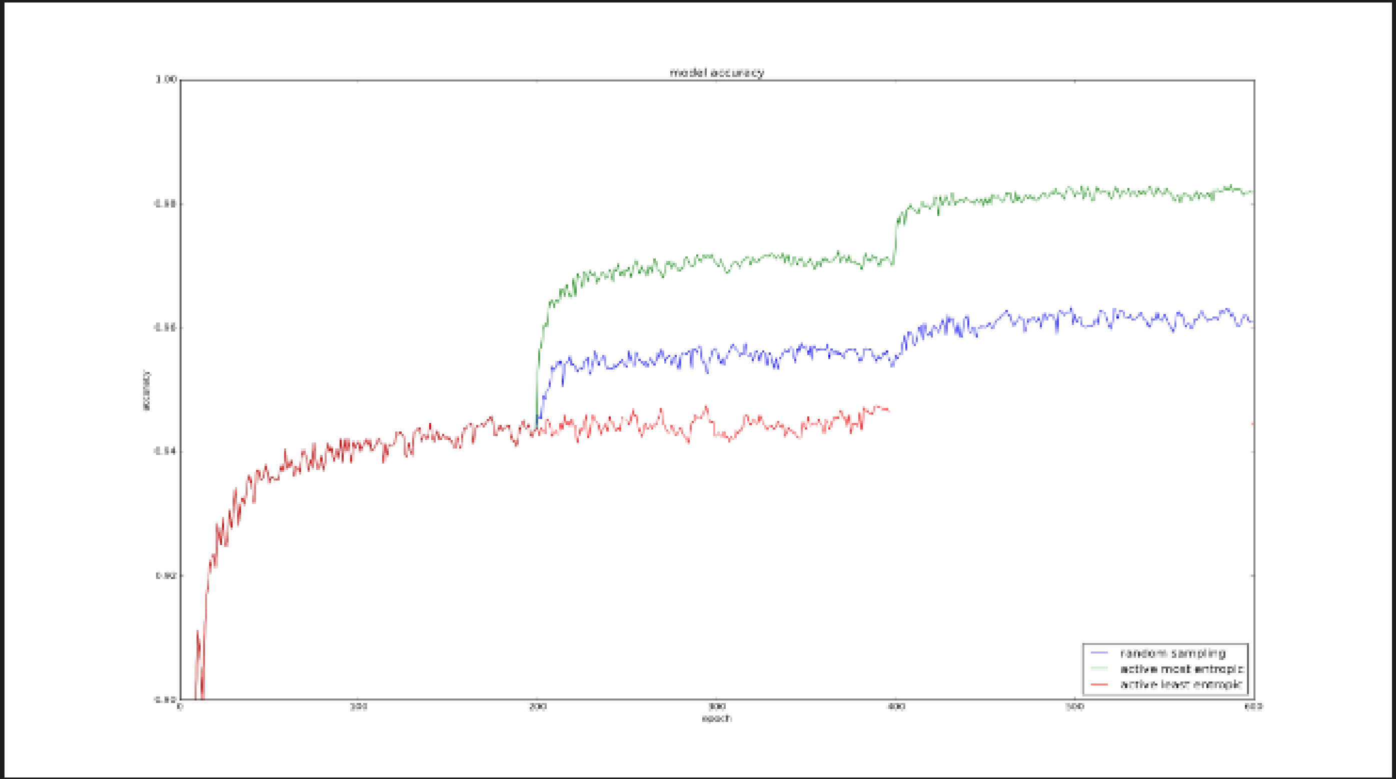

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue. **Two** active learning loops.

Classification Results

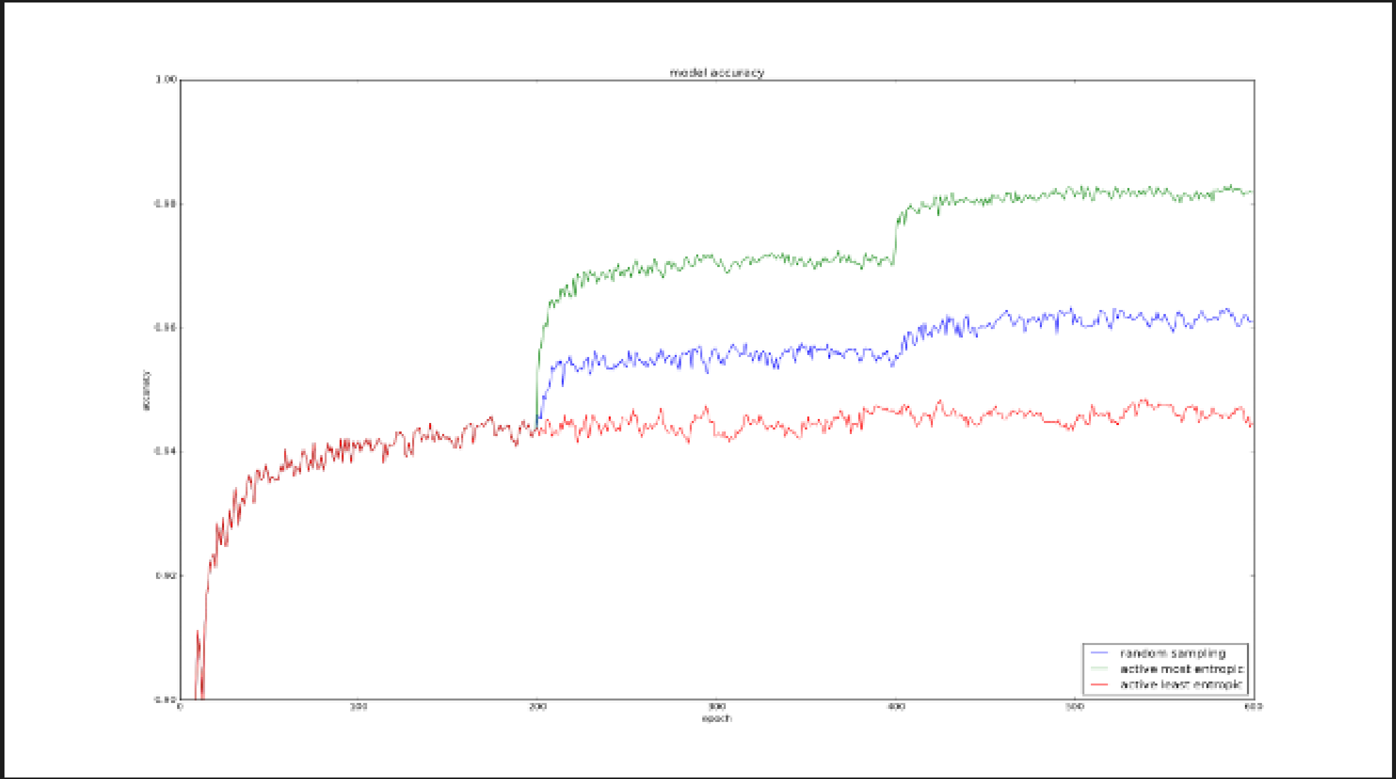

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue. **Two** active learning loops.

Classification Results

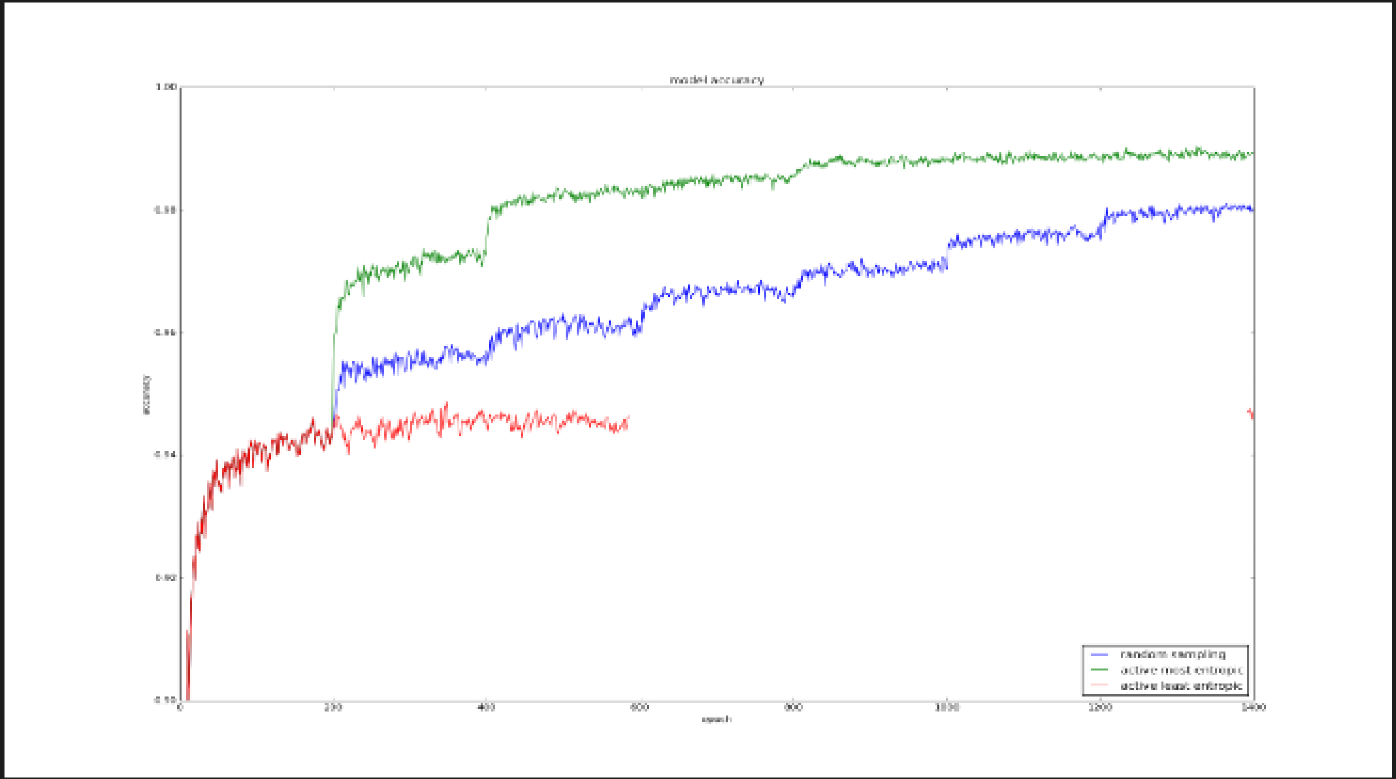

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue. **Six** active learning loops.

Classification Results

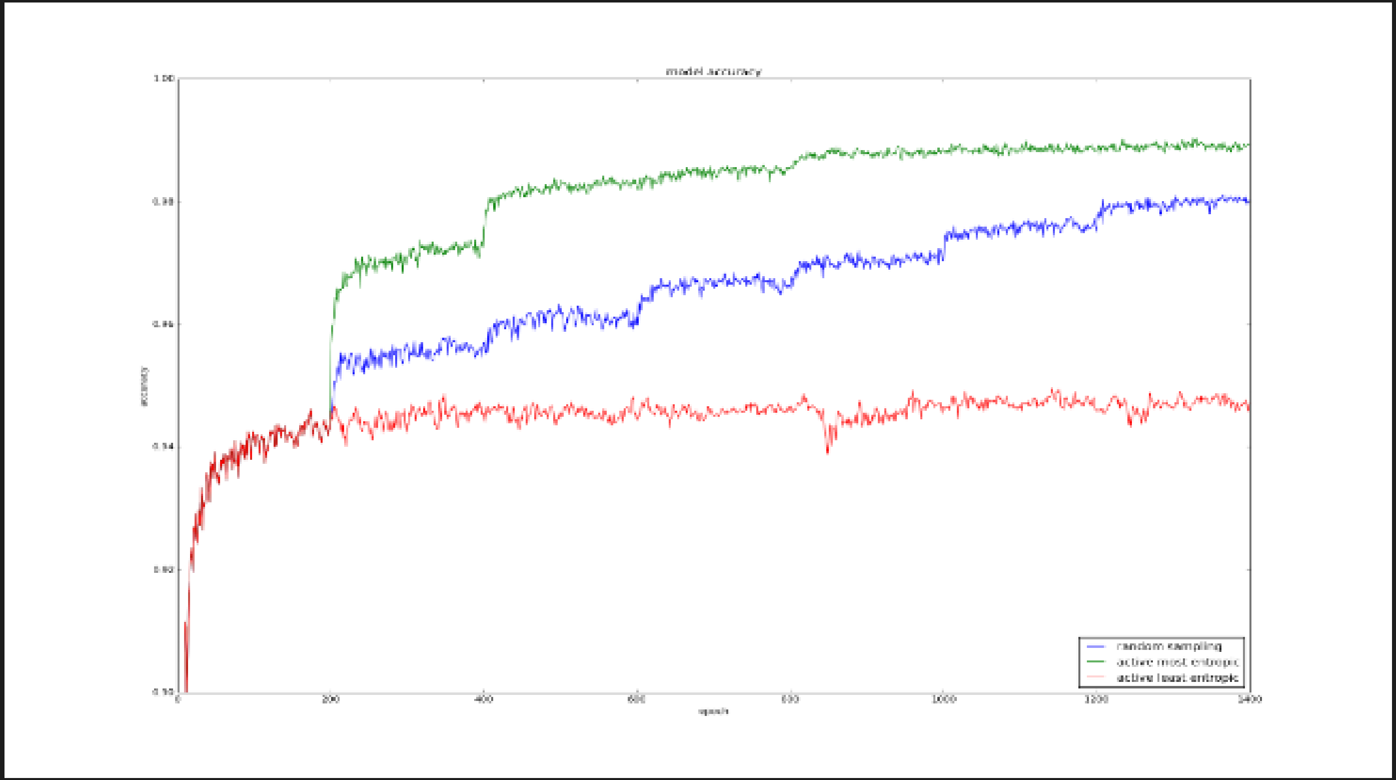

Sampling the least uncertain images (red) versus the most uncertain (green)
The random baseline is in blue. **Six** active learning loops.

[Lesnikowski, Farabet, 2017]

Scalability

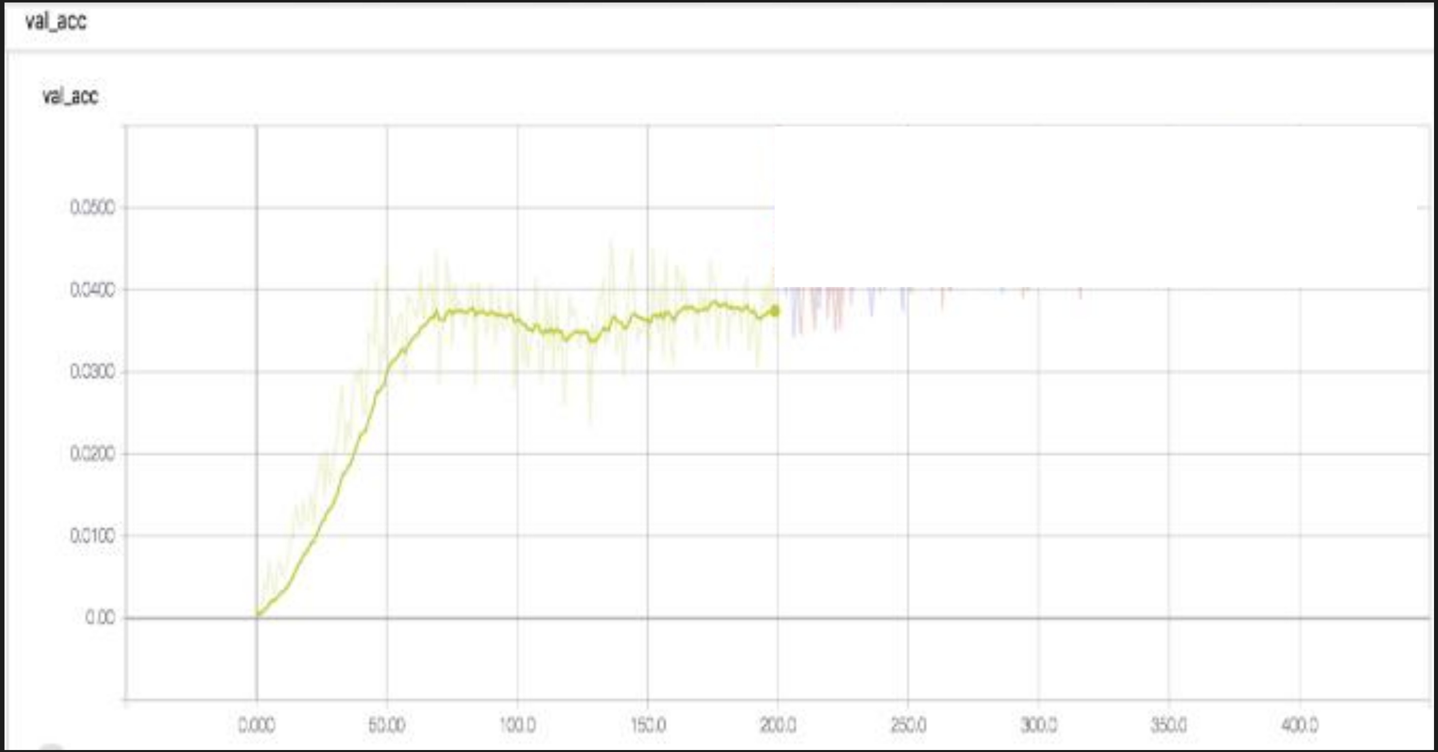

Classification on ImageNet after one active learning cycle with a margin policy.
Random baseline (red) and two runs of the active strategy (blue and purple).

Scalability

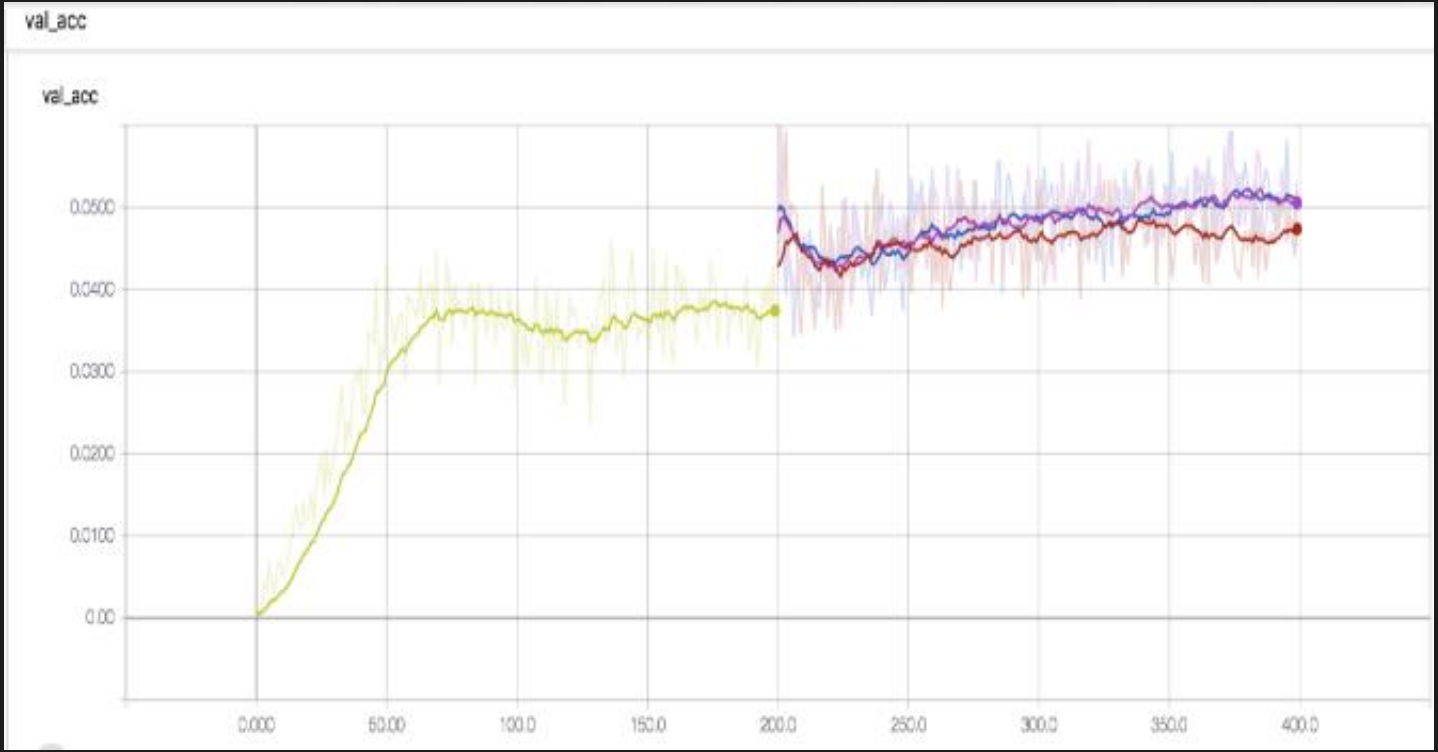

Classification on ImageNet after one active learning cycle with a margin policy. Random baseline (red) and two runs of the active strategy (blue and purple).

[*Lesnikowski, Schoffier, Farabet, 2017*]

Scalability

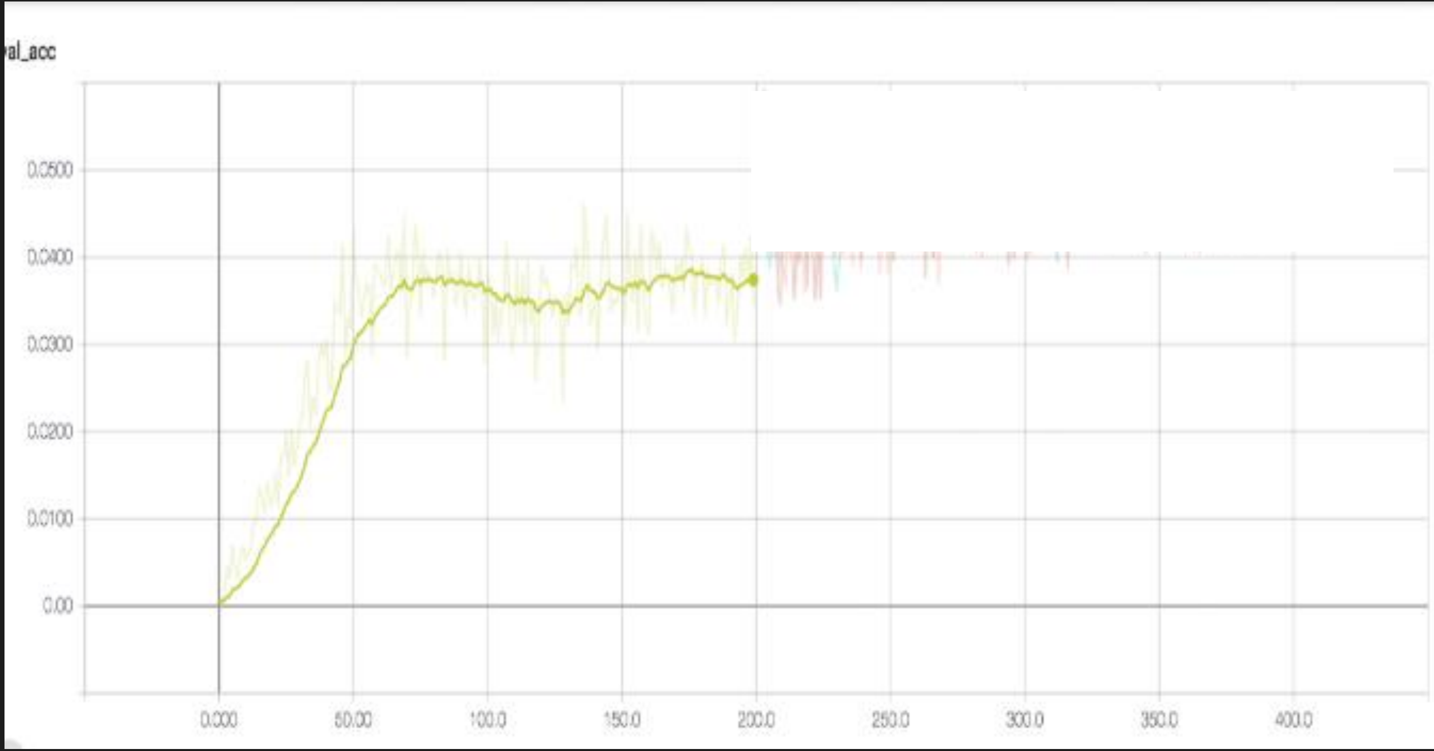

Classification on ImageNet after one active learning cycle with an entropy policy. Random baseline (red) and two runs of the active strategy (blue and orange).

Scalability

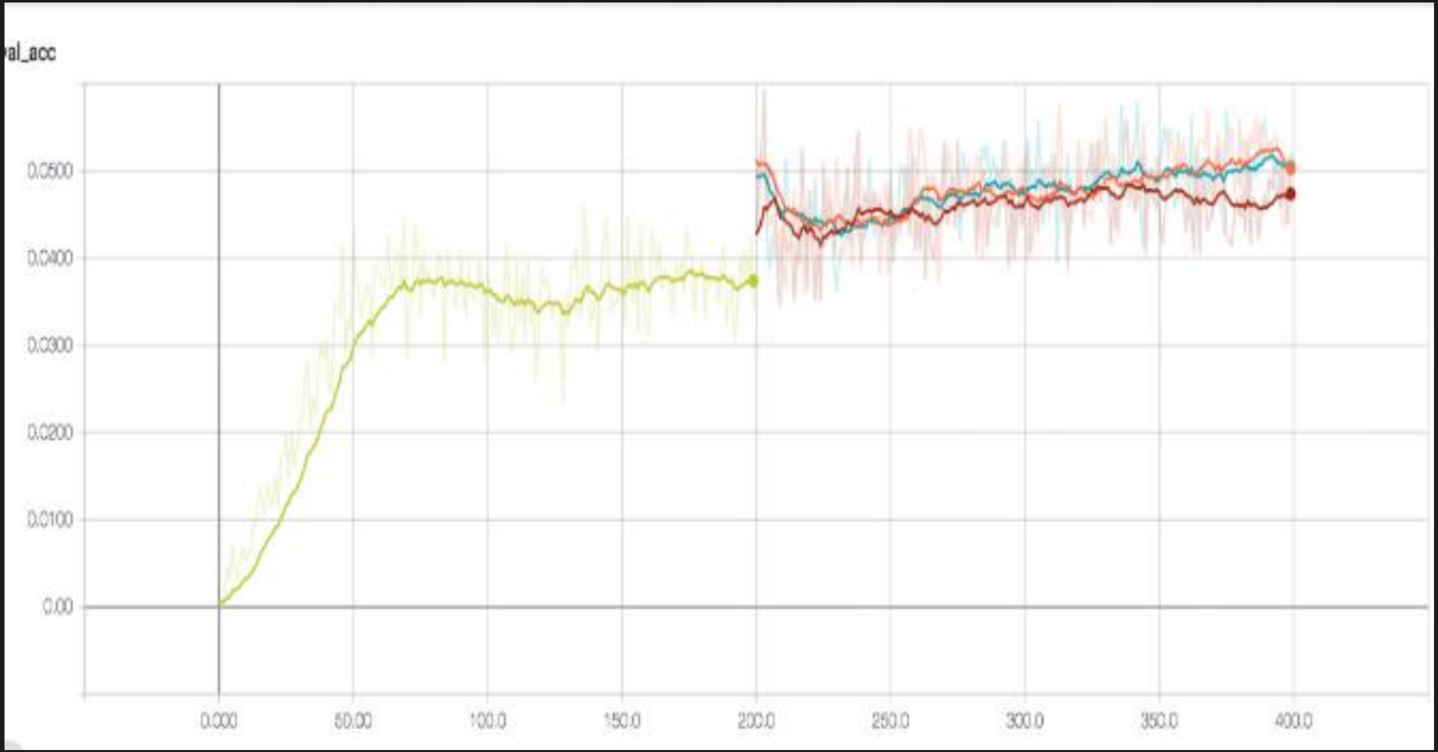

Classification on ImageNet after one active learning cycle with an entropy policy.
Random baseline (red) and two runs of the active strategy (blue and orange).

[Lesnikowski, Farabet, 2017]

Conclusions

Conclusions

* Deep active learning works

* Gives significant efficiency gains in time and money

* Exciting work in progress and more to come!

Thanks!

Jiaming Zeng, Stanford

Yousef Hindy, Stanford

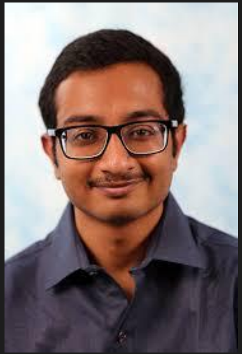

Kashyap Chitta, CMU

Clement Farabet, NVIDIA

Aysegul Dundar, NVIDIA

Tim Plump, MIT

Evan Cater, Brown

Jose Alvarez, NVIDIA

Thanks!

Shalini De Mello, NVIDIA

Stephen Tyree, NVIDIA

Anima Anandkumar, Caltech & NVIDIA

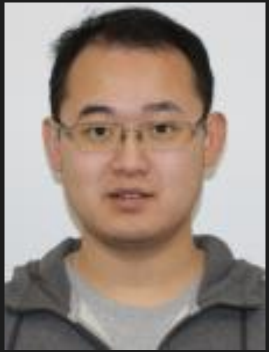

Shiyu Liang, NVIDIA

Qs, As and Discussion

Qs, As and Discussion

Thank you!

Qs, As and Discussion

Thank you!