Moving Object Detection with Deep Convolutional Networks and LSTM

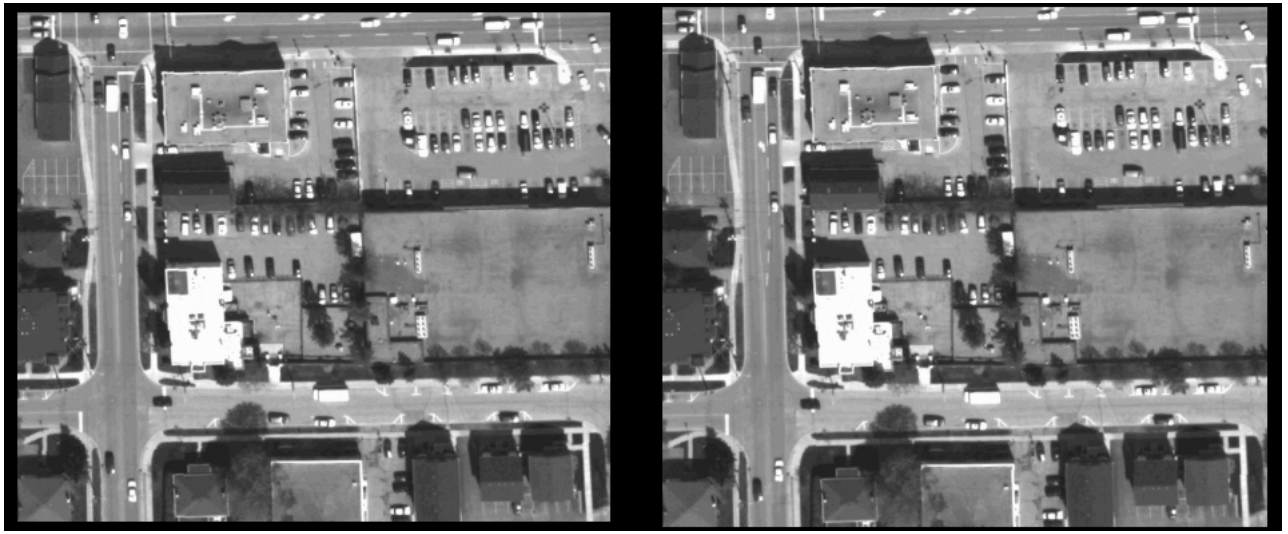
Gregoire Robinson
Advisors: Dong Zhang and Dr. Mubarak Shah

This Week...

- Fixed all bugs, and network is learning
- Keras documentation was unclear
- Currently running two experiments, one with input of 5 frames, the other using background subtraction

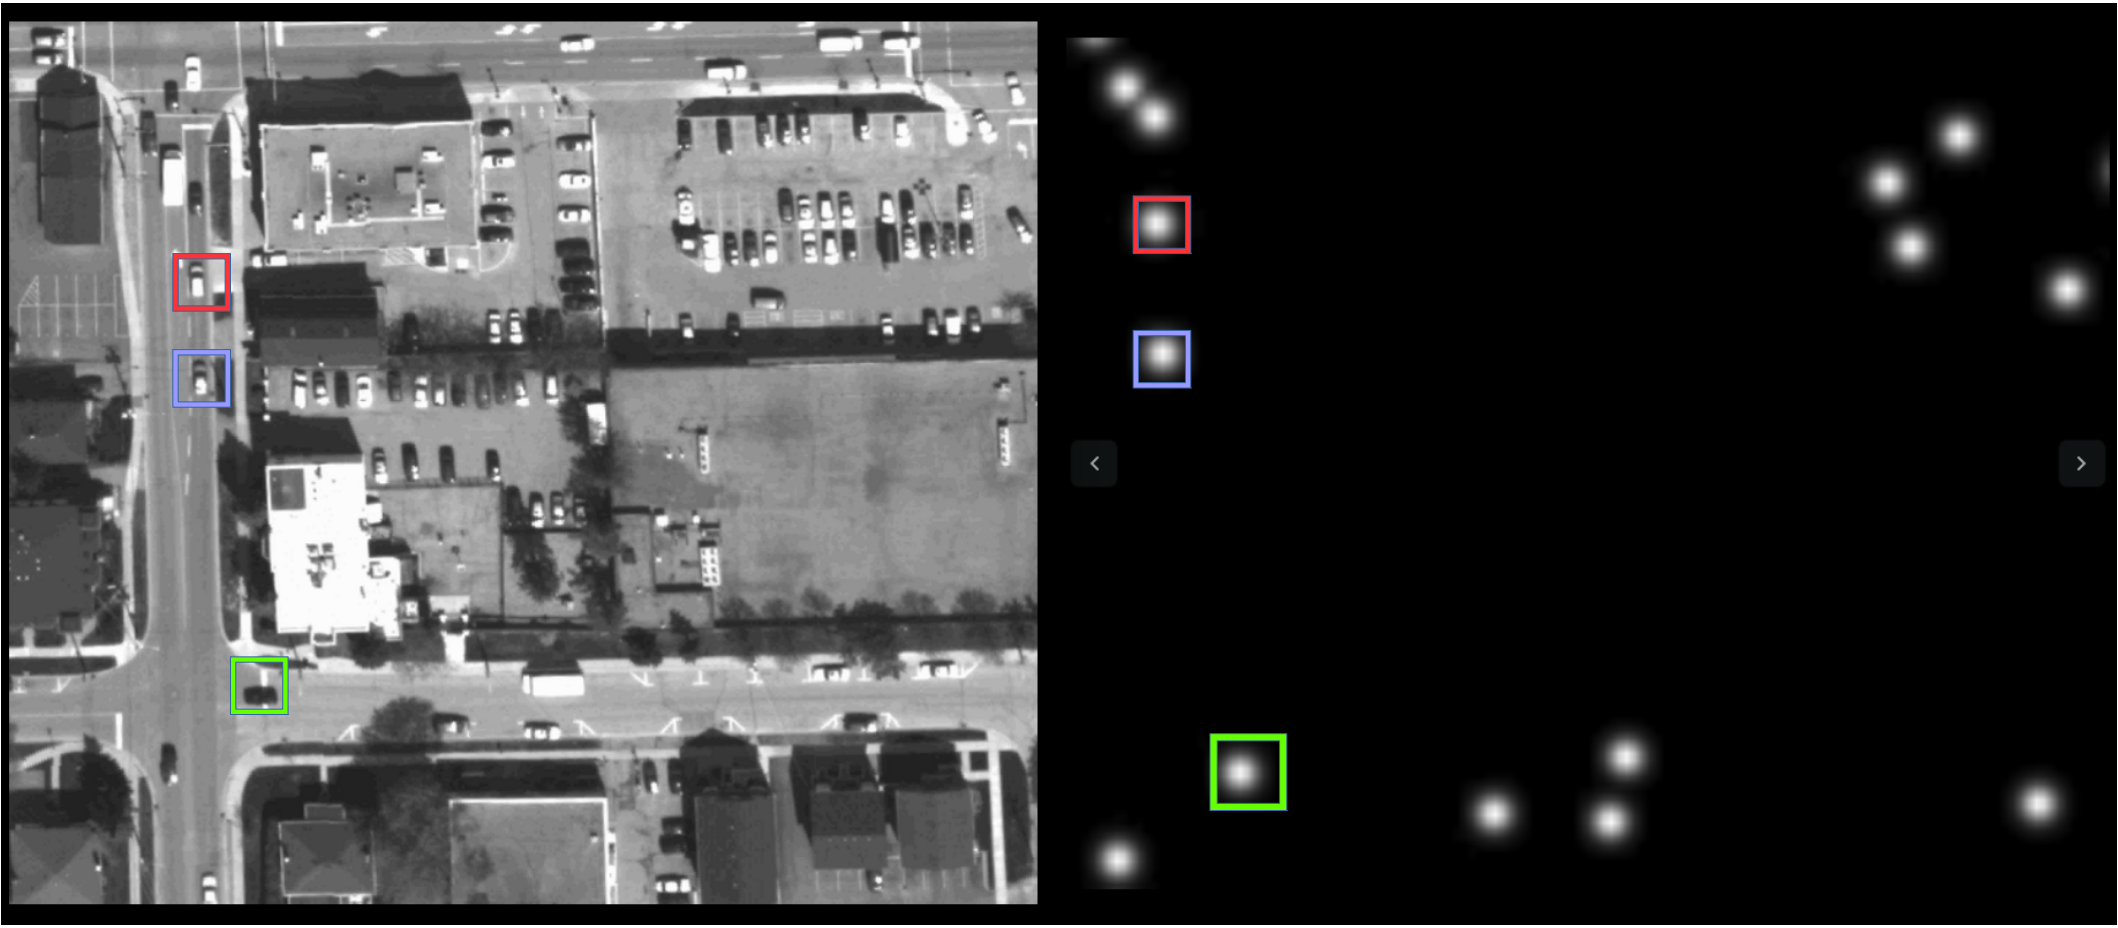
Five Frames





Five Frames

- Train using a single heatmap. Choose the third image in the input sequence



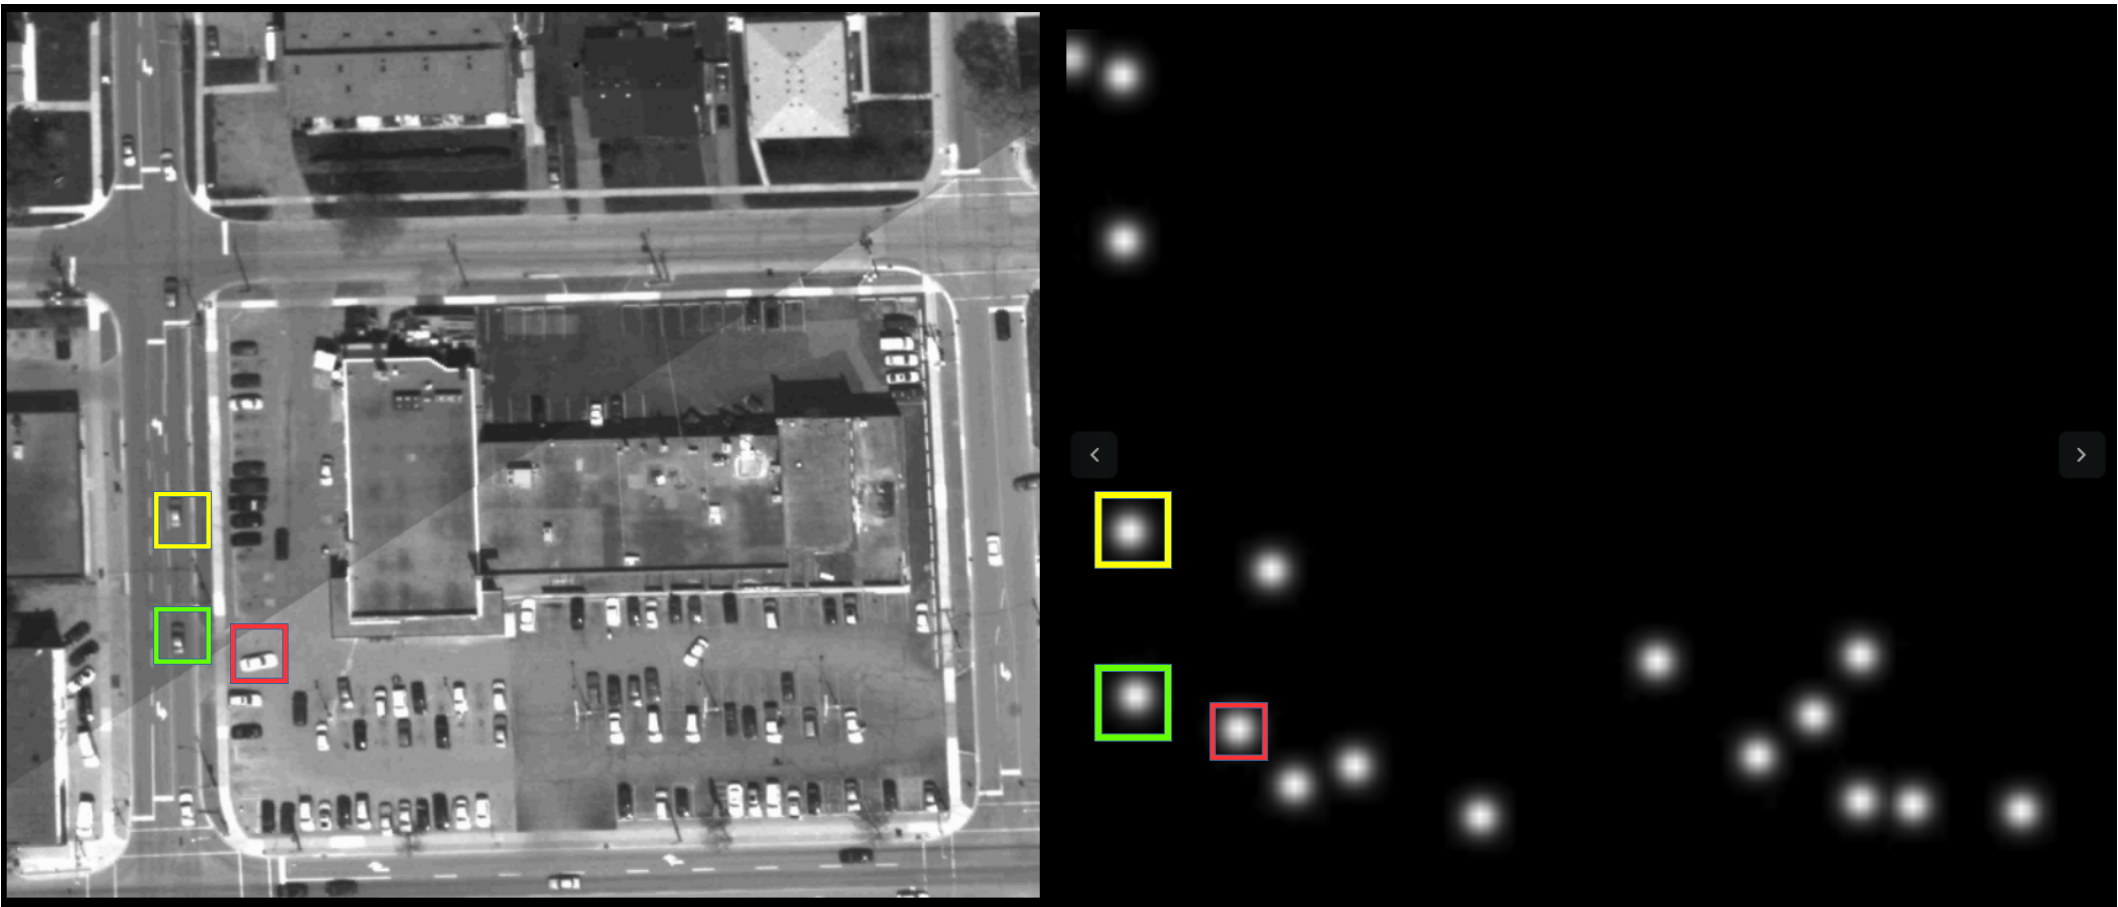
Background Subtraction

- Network input is a single sub-image and its corresponding background-subtracted image



Background Subtraction

- Train using a single heatmap. Annotations correspond to moving vehicles.

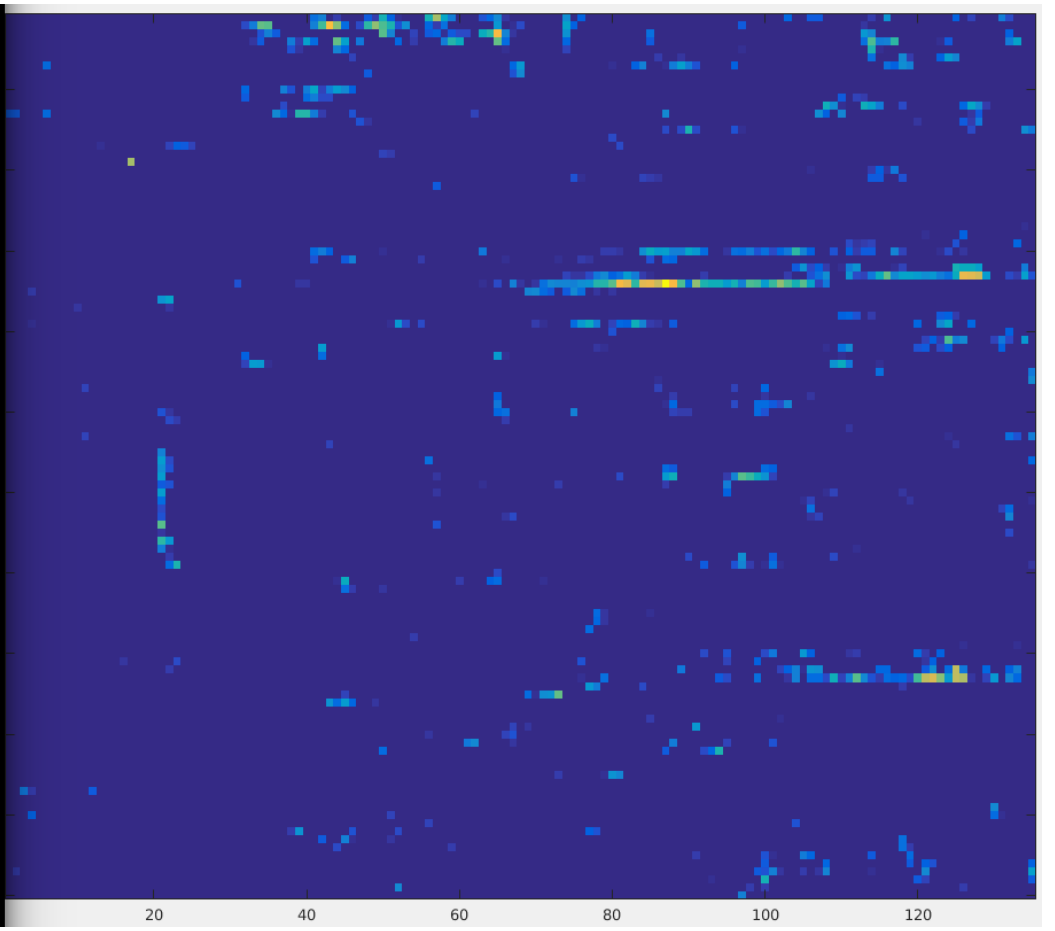


Five Frames Ground Truth

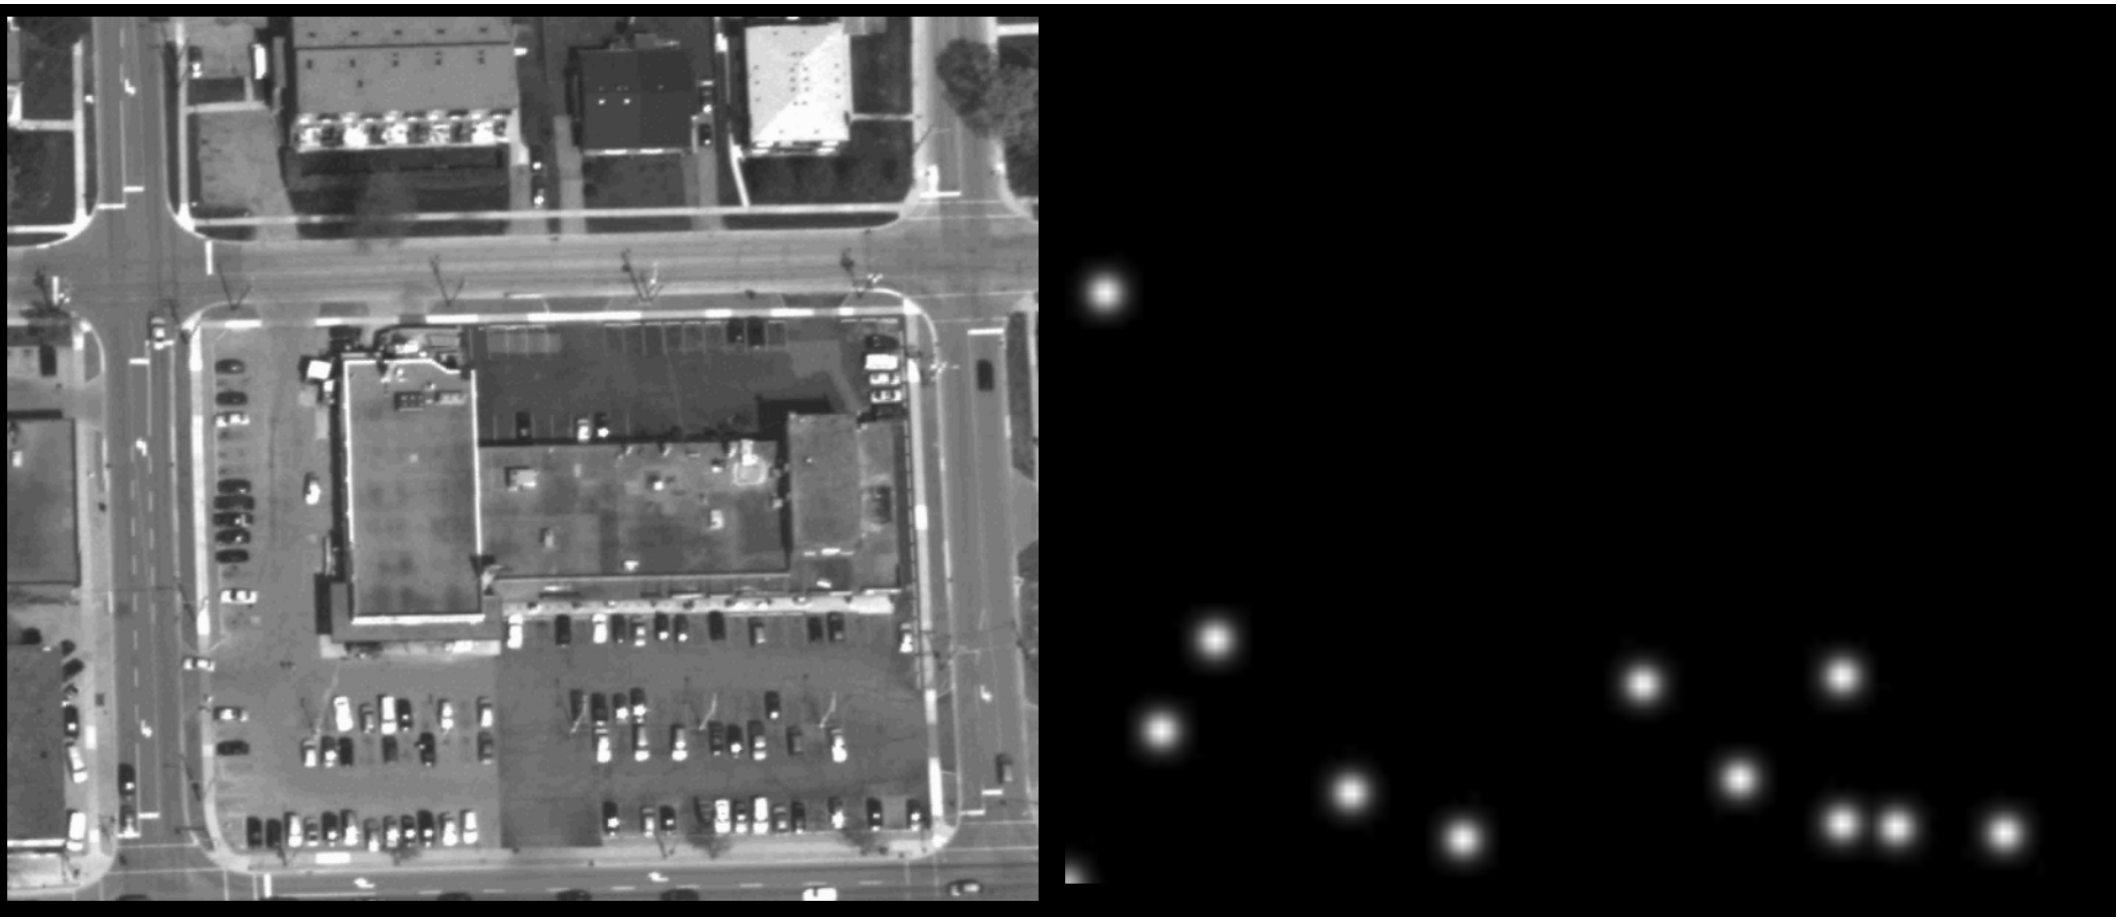


Five Frames Results



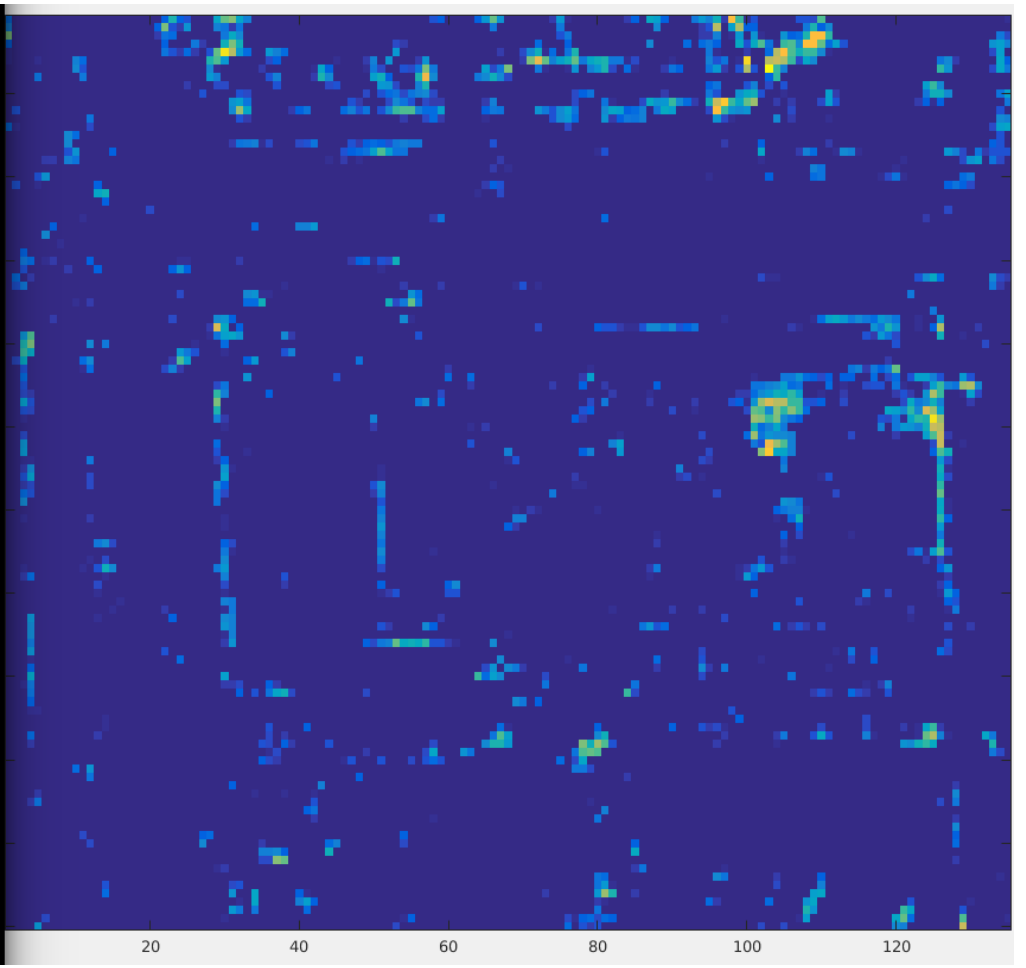


Background Subtraction Ground Truth



Background Subtraction Results





Conclusions

- Difference between Caffe and Theano learning
- Training time on cluster is at least 18 times slower than on a local machine