SAND2020-6381C

Synthetic Training Images for Real-World Object Detection

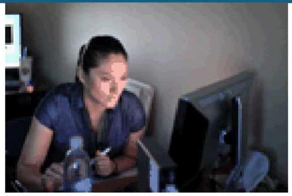

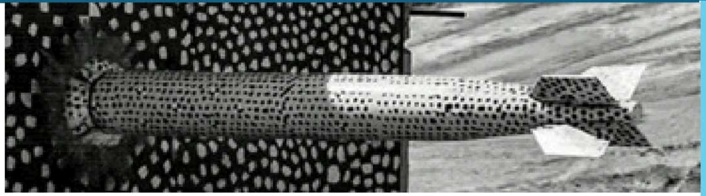

Presented By

Zoe Gastelum

Research Team: Timothy Shead, Michael Higgins

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Potentially Safeguards-Relevant Data

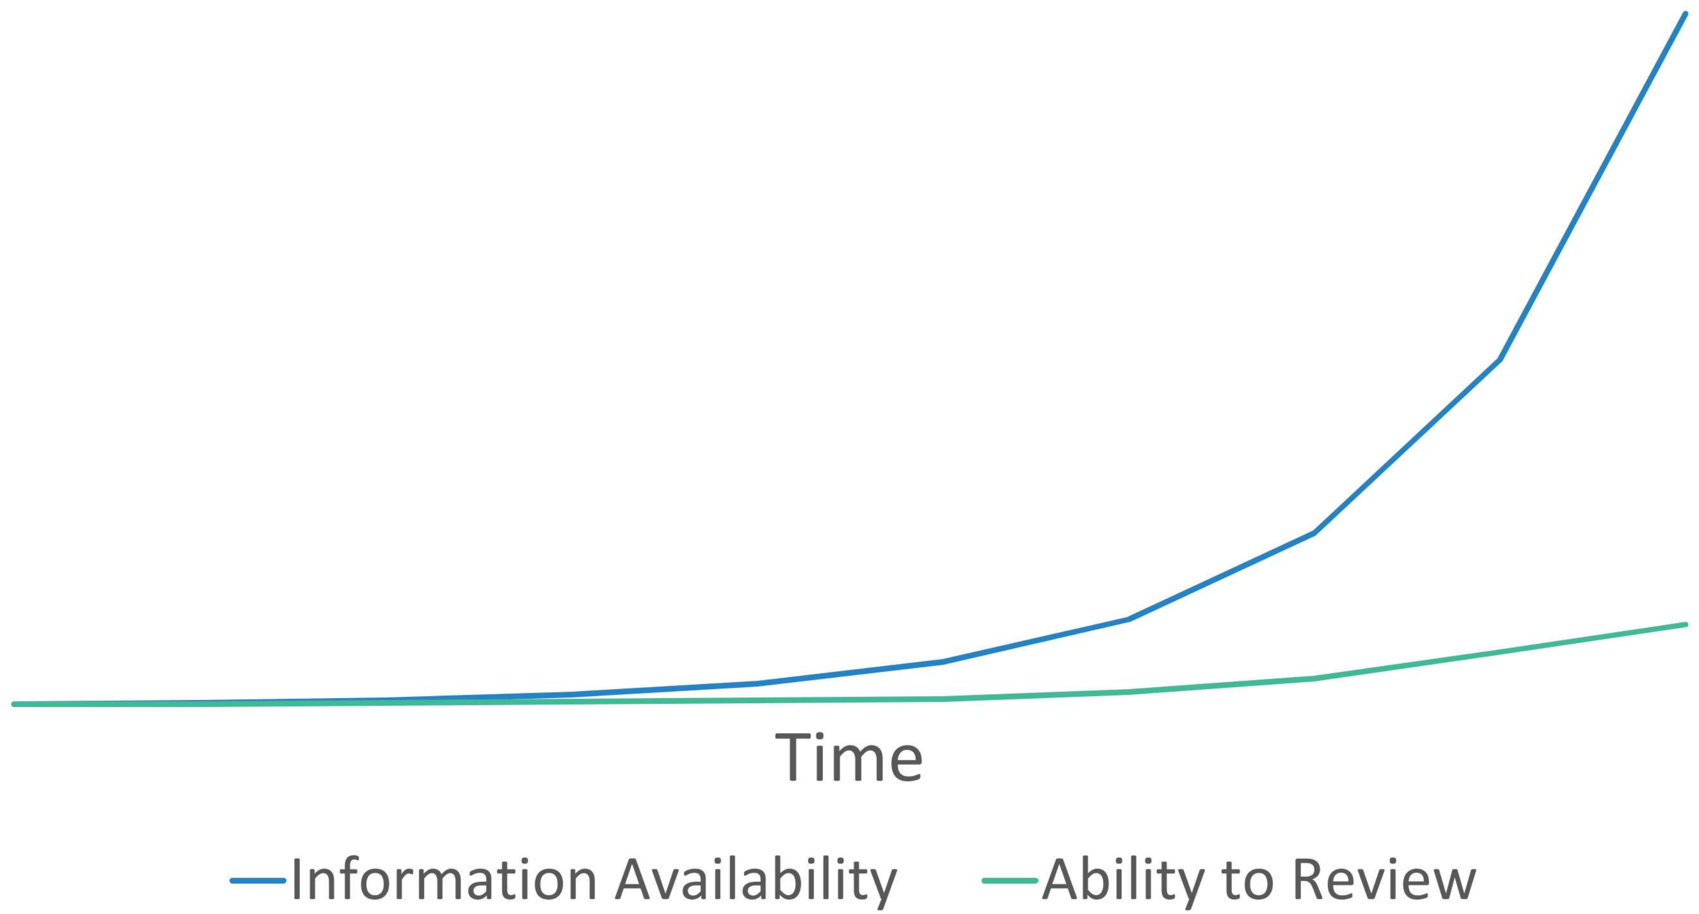

Computer Vision: *What objects are in an image, where they are, how they relate to the background, and the meaning of the image as a whole?*

Classification:

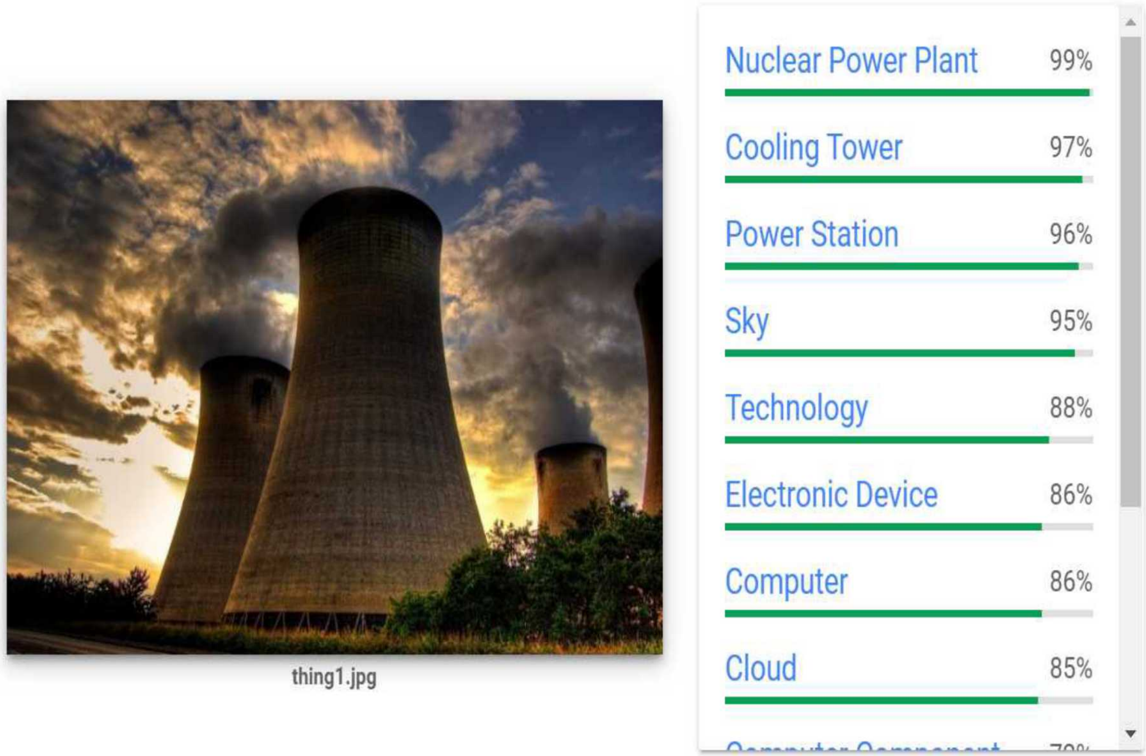

Object Detection:

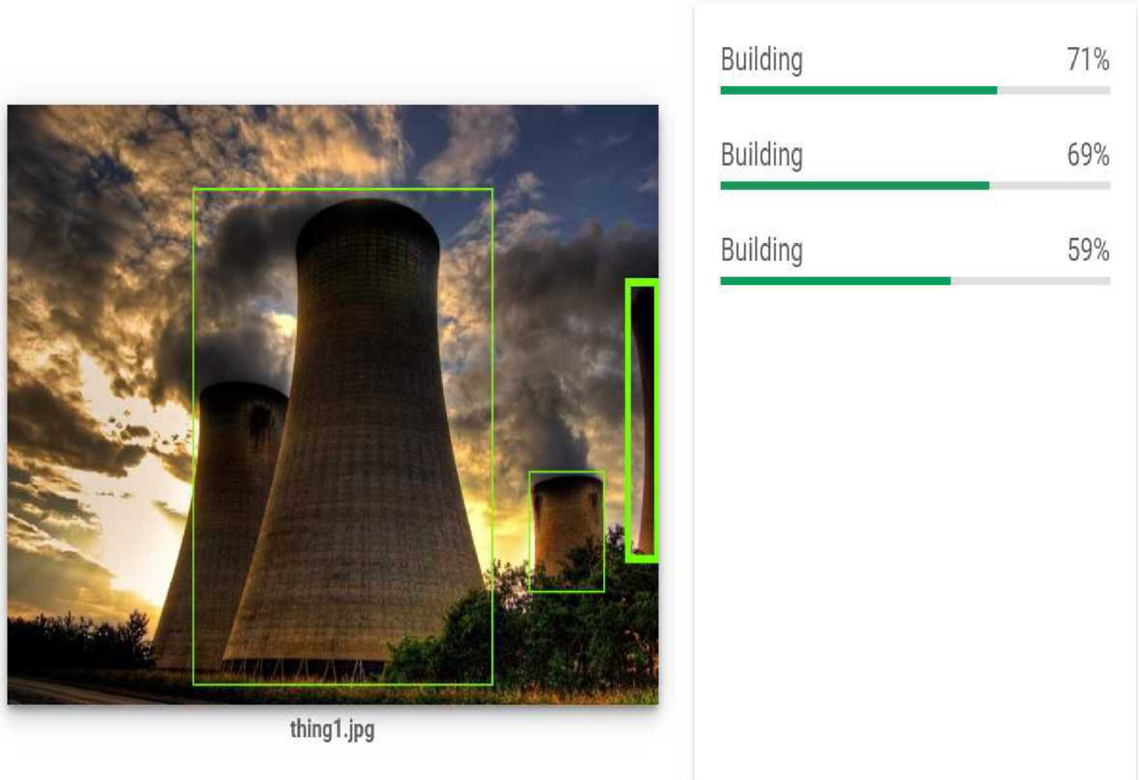

Training Deep Computer Vision Models

MNIST: 70,000 handwritten digits

Training Deep Computer Vision Models

COCO: > 200,000 Labeled images; 1.5 million object instances

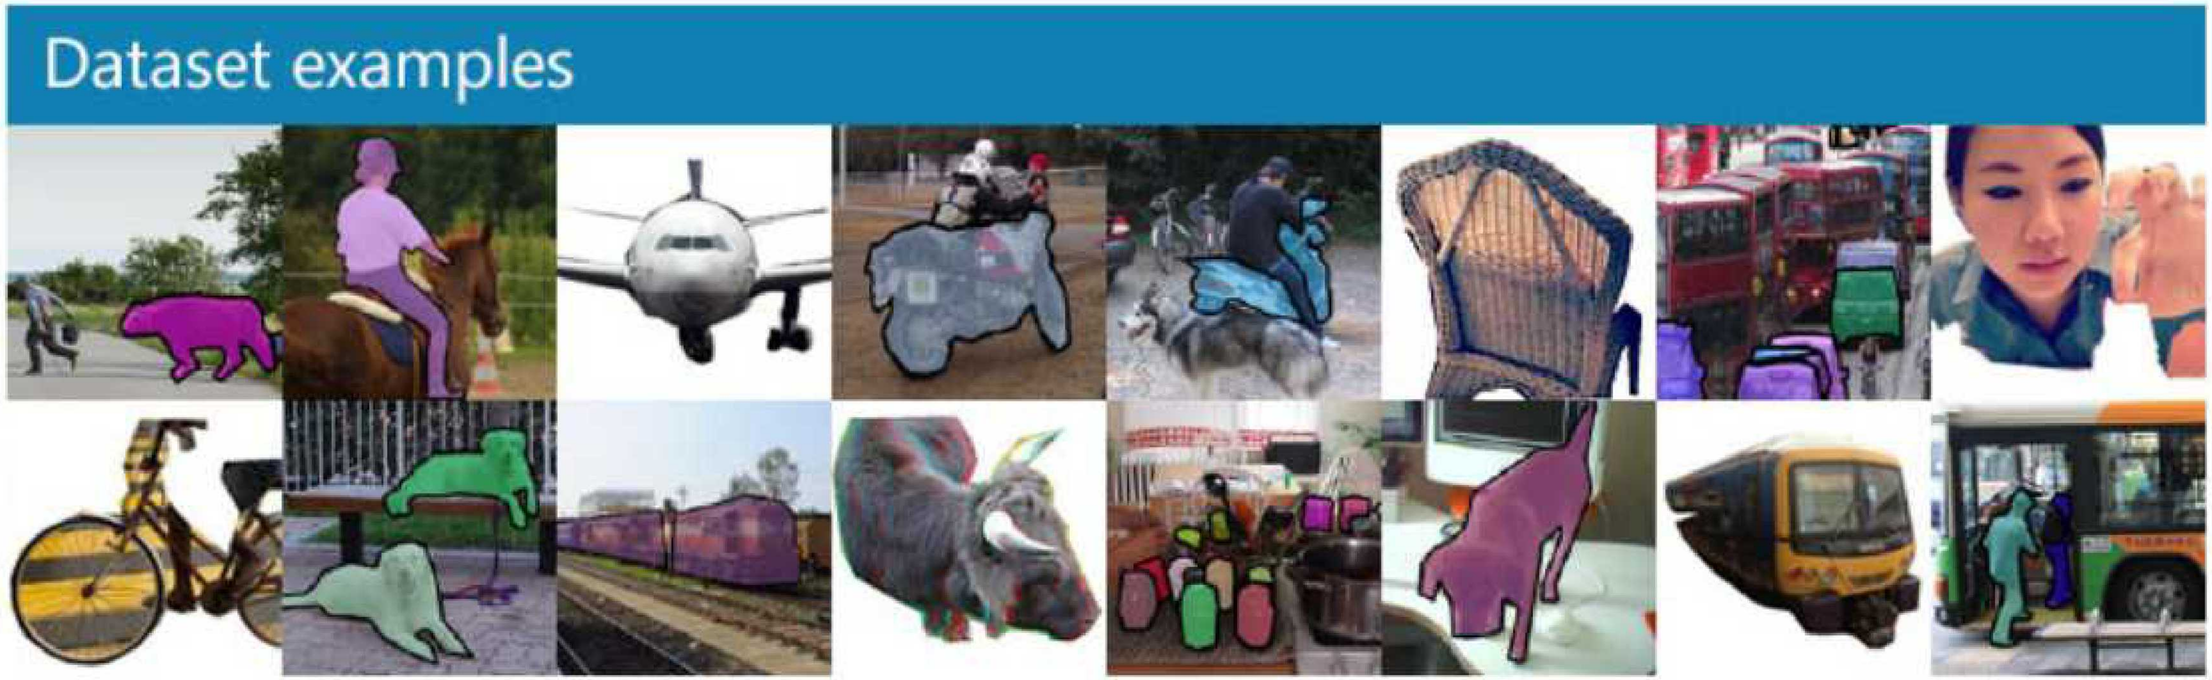

Training Deep Computer Vision Models

ImageNet: > 14 million Images; >1 million bounding boxes

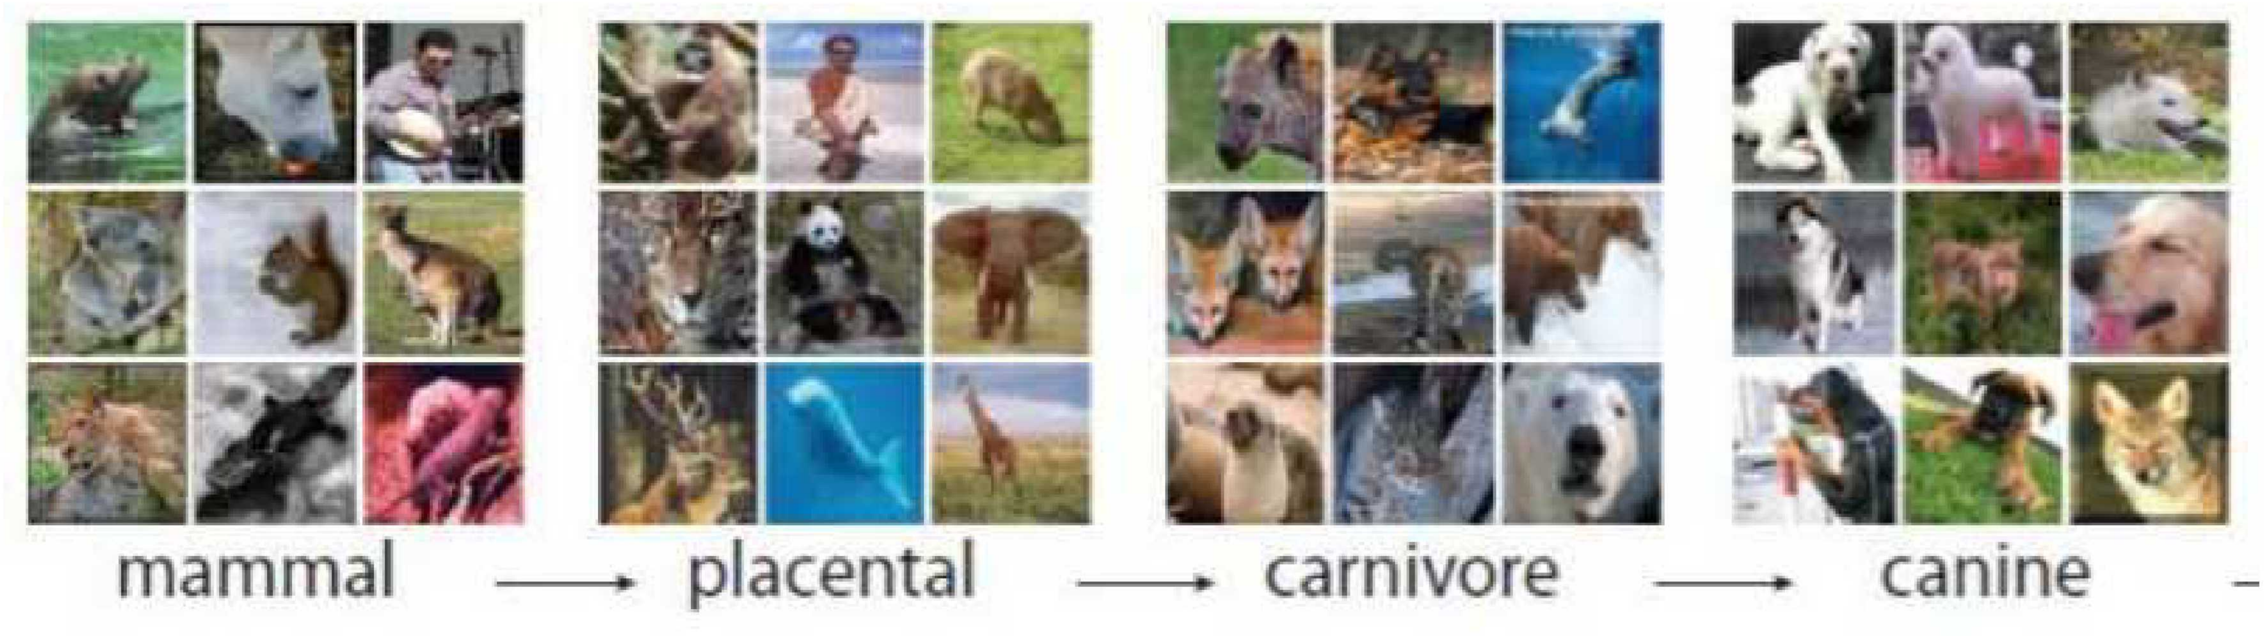

Safeguards training data challenges:

1) Proliferation-relevant images may be rare due to their sensitivity or the limited availability of a technology.

2) Creating relevant images through real-world staging is costly and introduces biases into the resulting model.

3) Expert-labeling is expensive, time consuming, and prone to error and dissent.

Hypothesis:

Synthetic images generated from three-dimensional CAD models of an object of interest can be used during training to overcome these problems.

Would you trust the autonomous vehicle that can't identify both signs?

Progress To-Date

1) Synthetic image generation

2) Data and model validation

3) Image background experiments

4) Real:Synthetic ratio experiments

Synthetic image generation in Houdini

- Parameterized model for unlimited poses and configurations
- Randomized camera angle with realistic lighting
- 33 panoramic HDR backgrounds from industrial environments

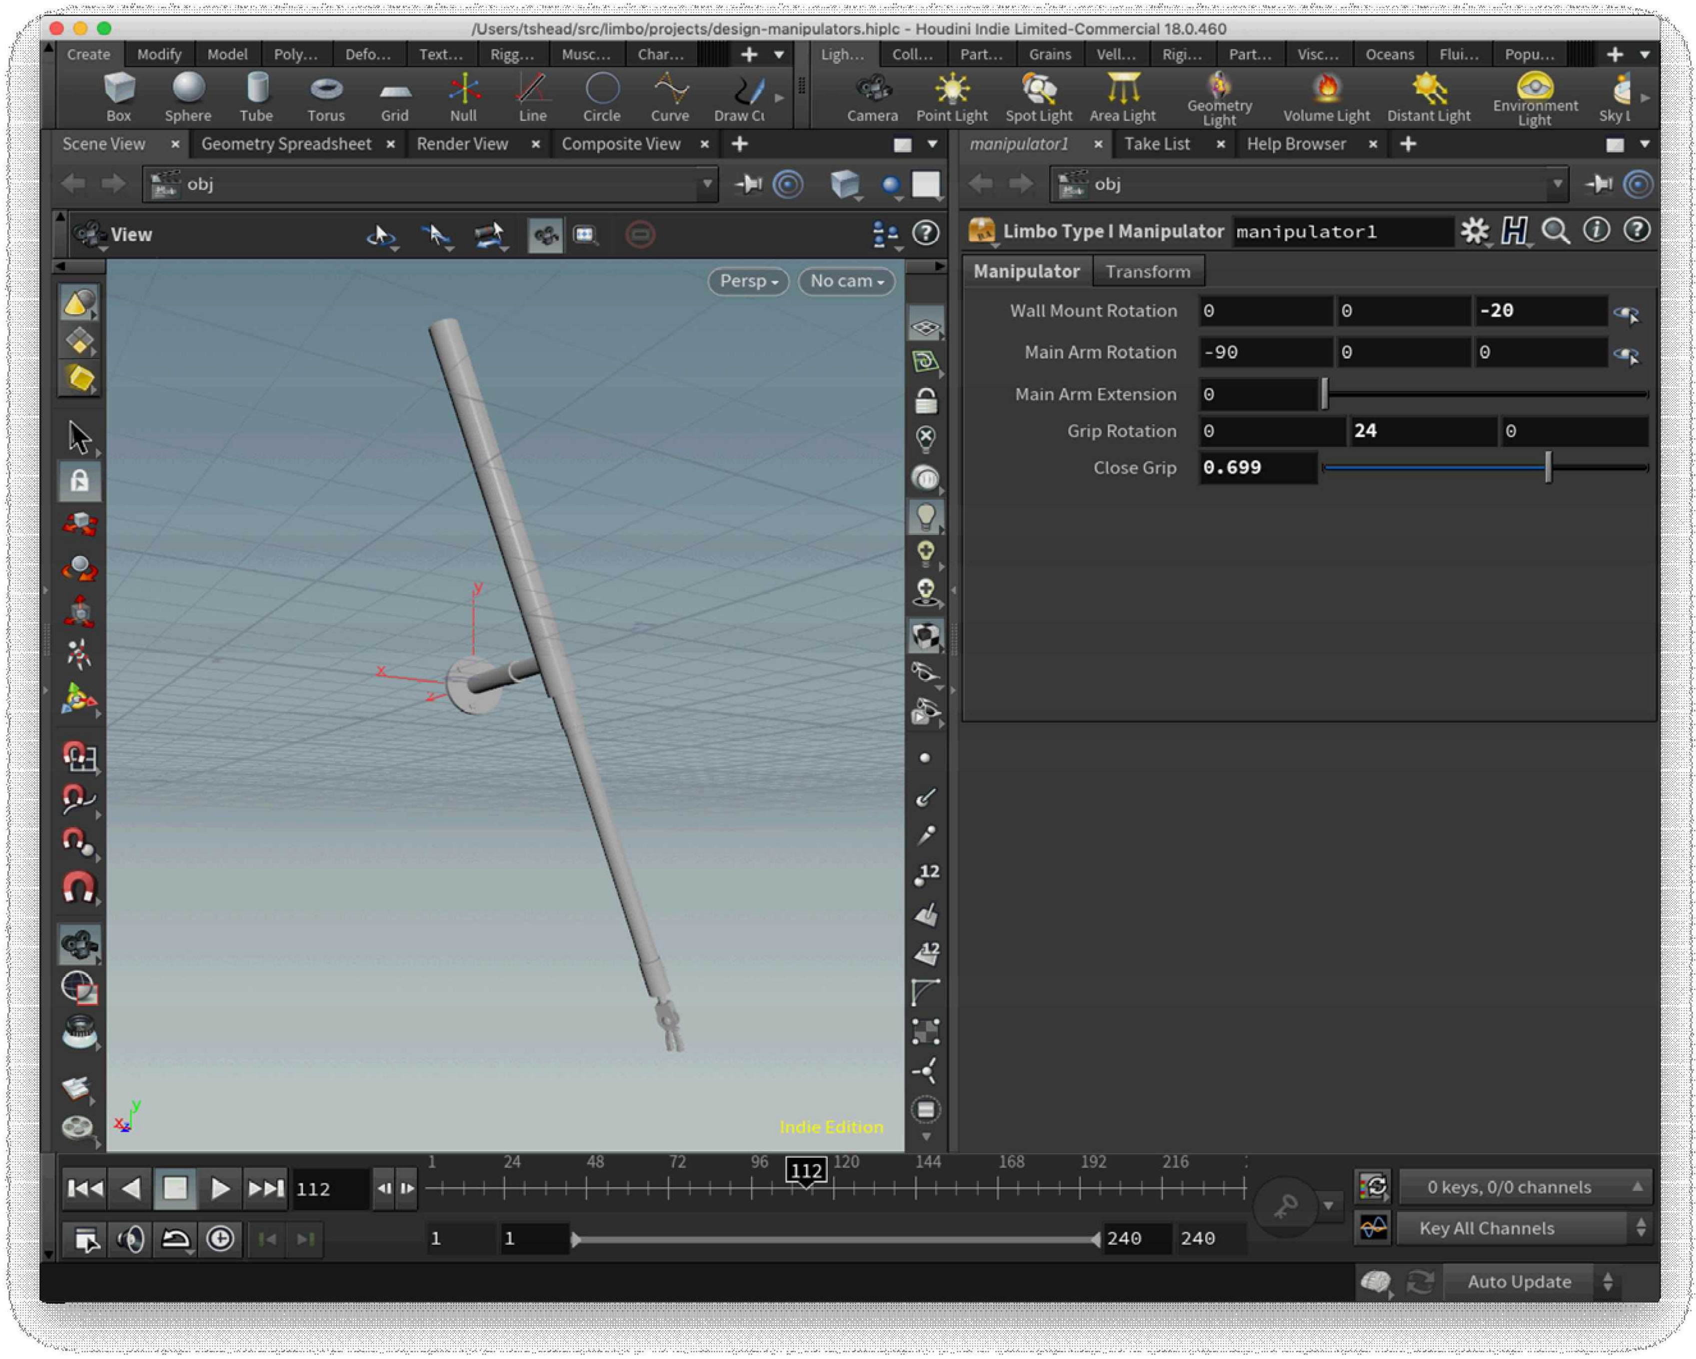

Example images

Data & Model Validation

Do we have sufficient data to fine-tune the model?

Res-Net 50 image classification model pre-trained on ImageNet

Replaced final 1000-ooutput fully connected layer to detect a single class: “manipulators”

1) Fine-tuned using 411 real-world images

2) Fine-tuned using 2000 synthetically generated images

Real-World Image Validation

Average loss and accuracy training and testing using only real-world images with 5x2 cross validation.

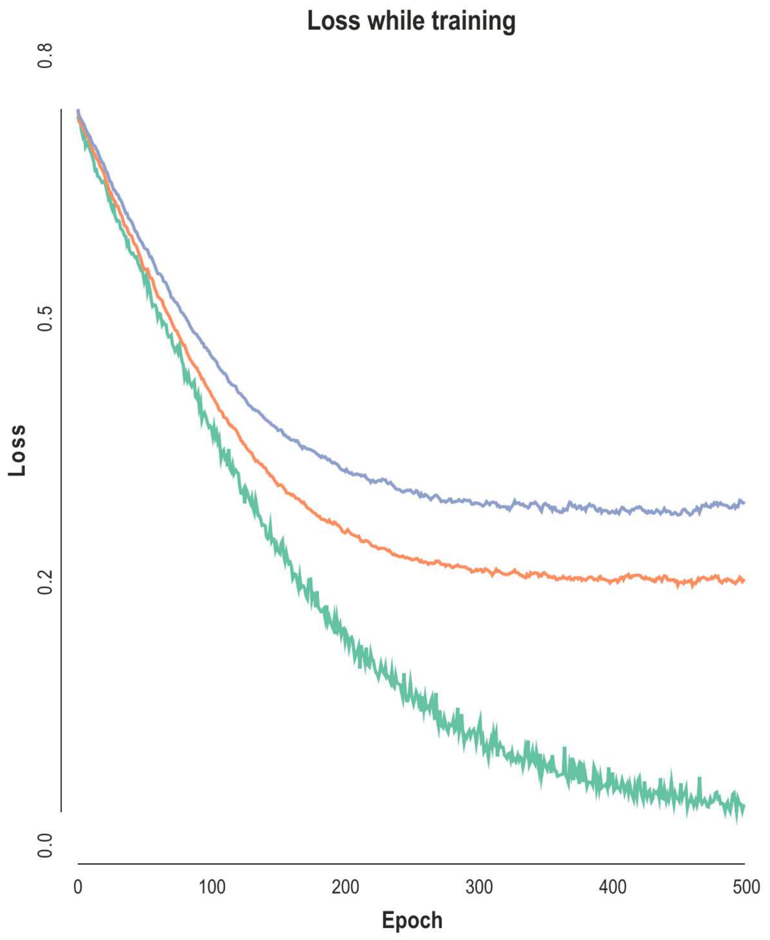

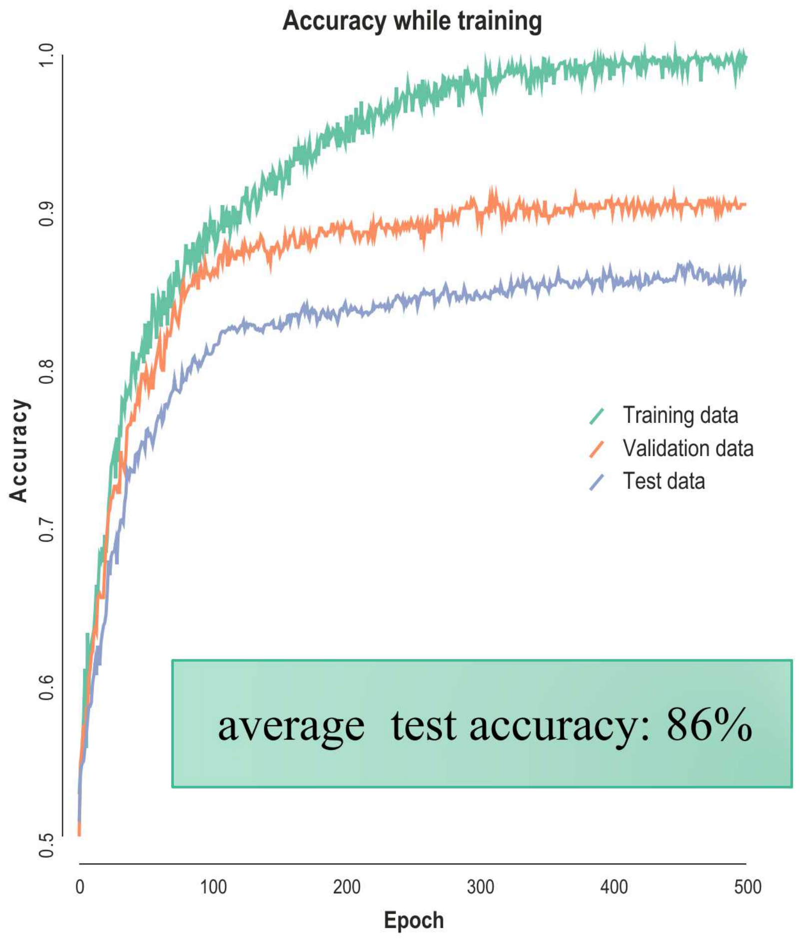

Synthetic Image Validation

Average loss and accuracy training and testing using only synthetic images with 5x2 cross validation.

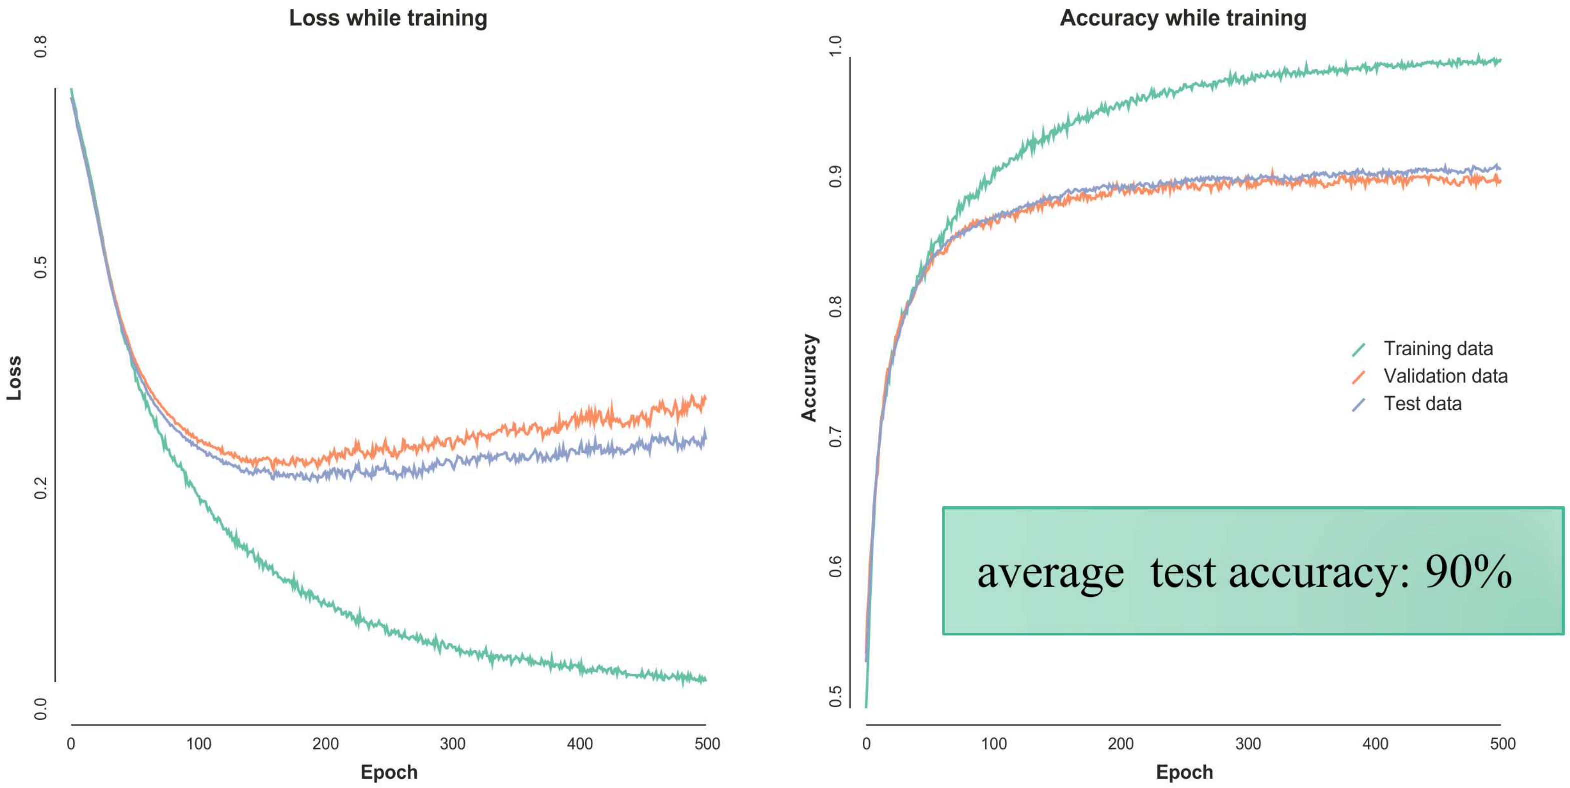

Average performance of synthetic training data with real test data, for ten models.

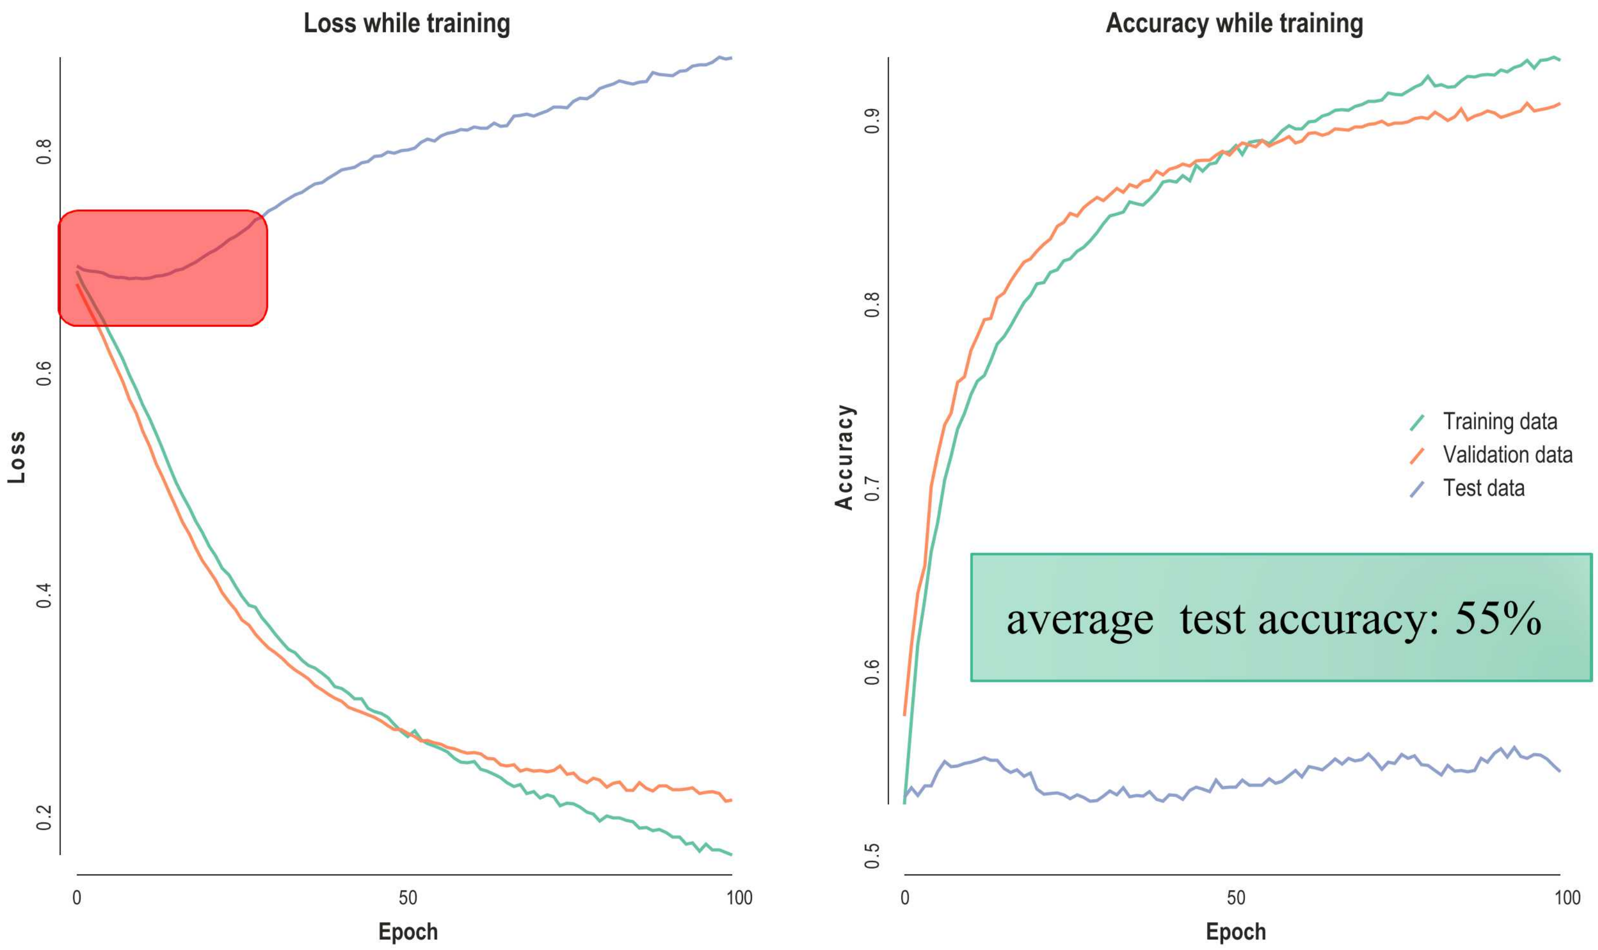

First Run

Trained on synthetic data

Tested on real-world data

Random 20% of synthetic data for validation, each experiment repeated 10 times

Poor performance could be due to low variance in training data

Image Background Experiments

Re-rendered 2000 synthetic images with randomly-selected backgrounds from 191 real-world distractor images

Test loss never decreases – model is not learning, hybrid dataset was of zero value

Average performance of synthetic-foreground + real-background training data, ten models

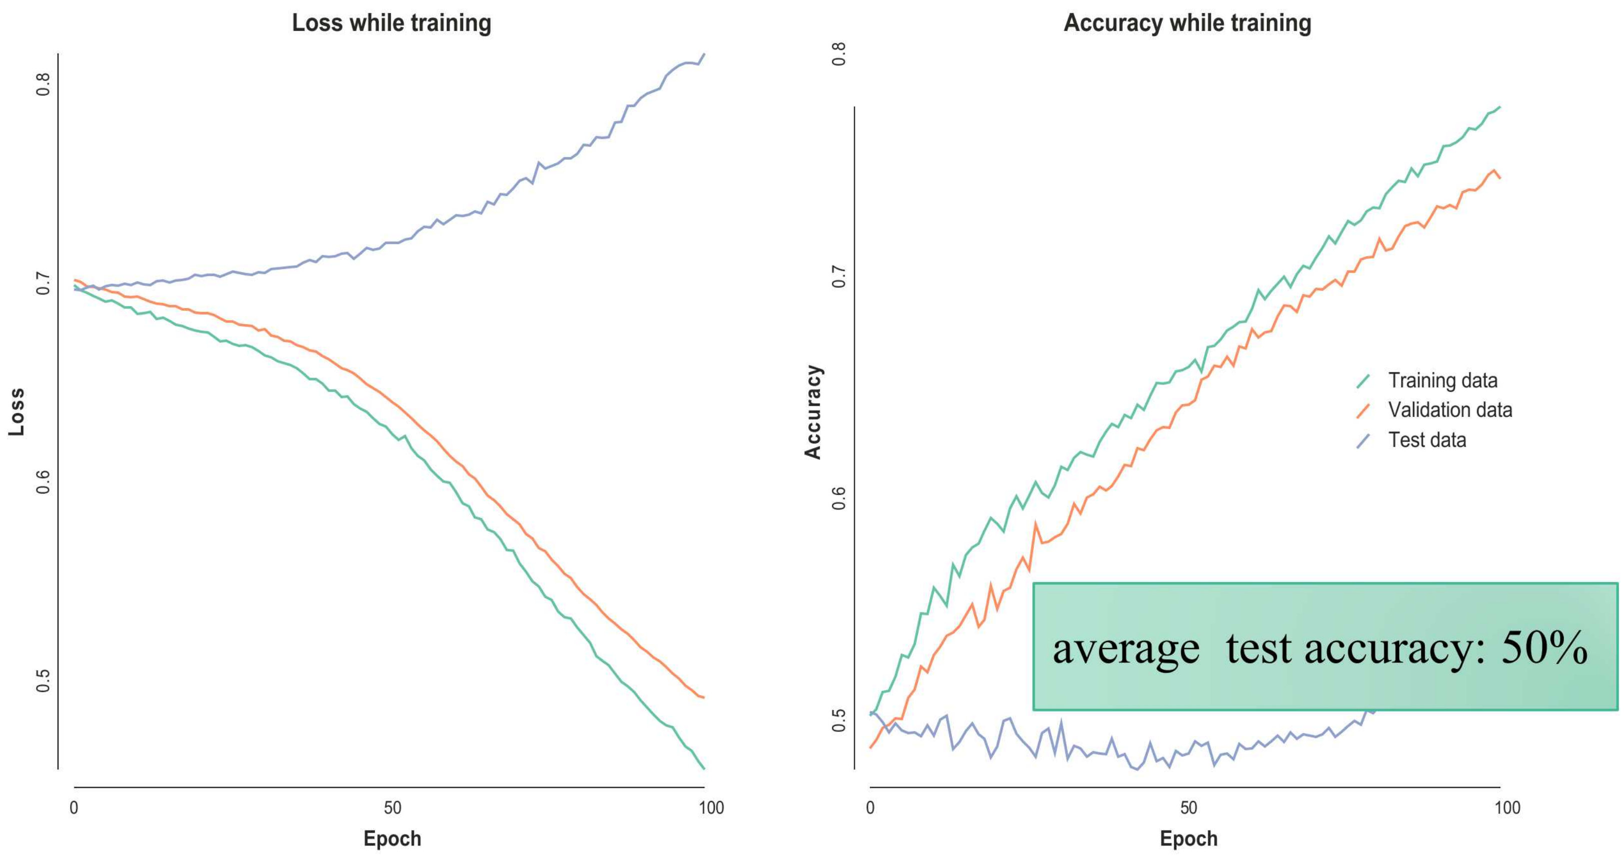

Image Ratio Experiments

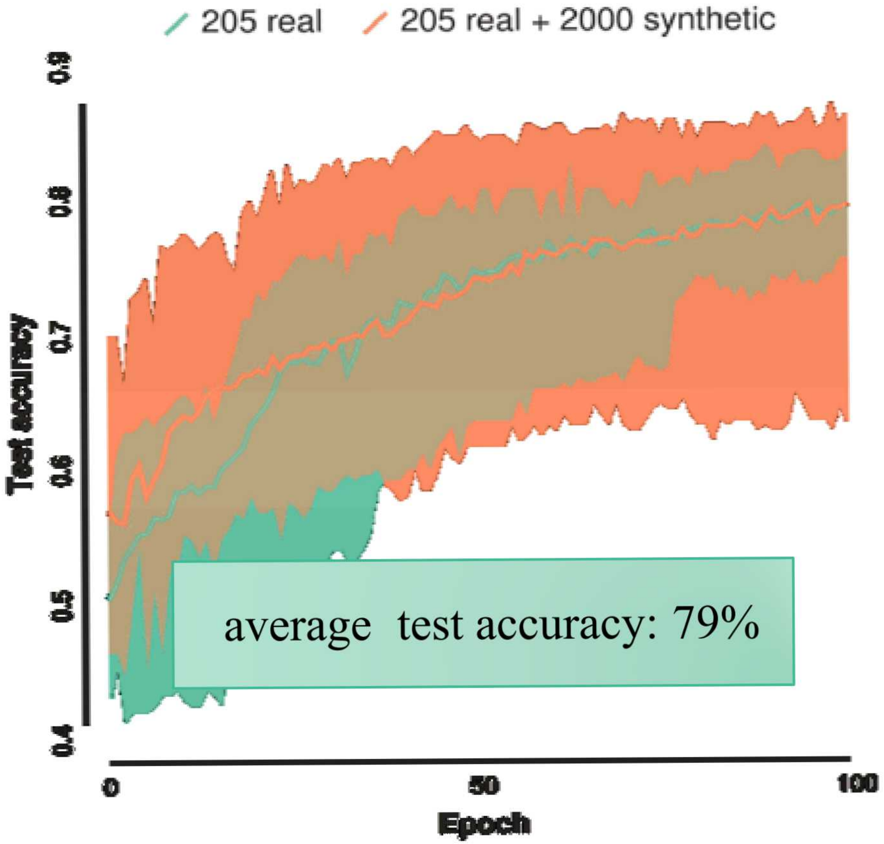

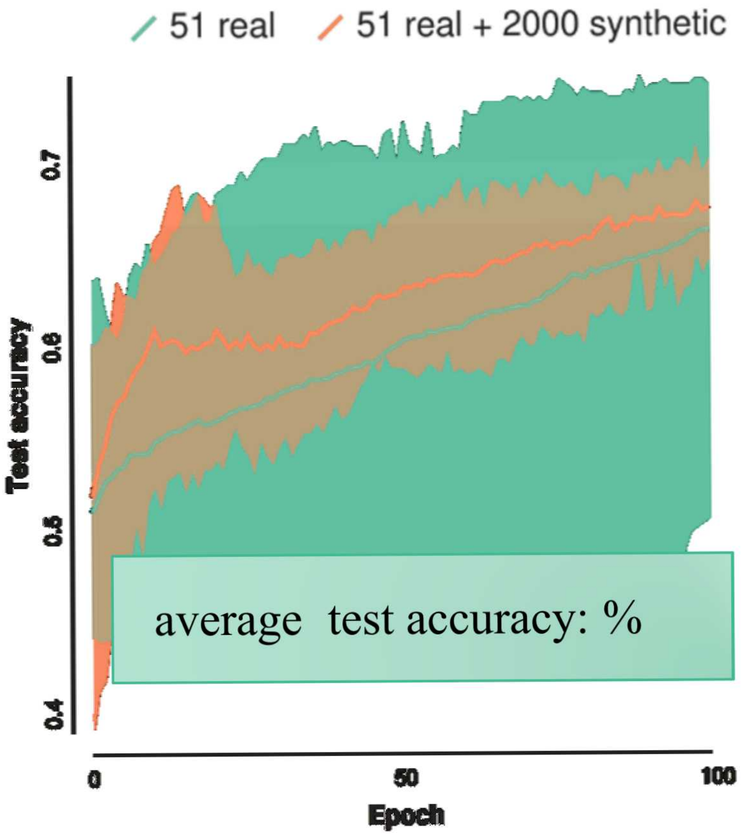

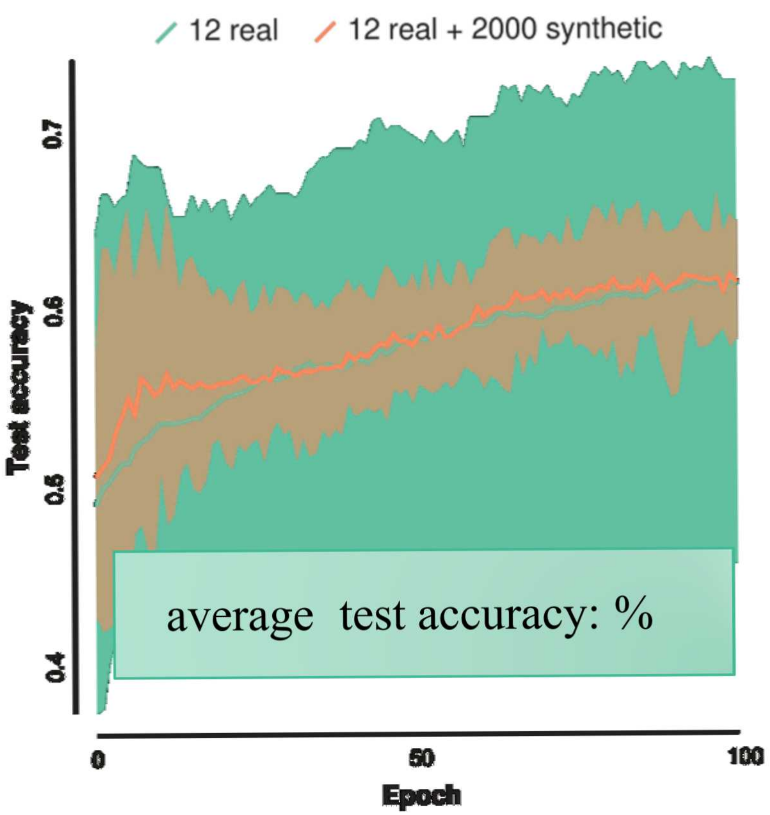

Conclusions

From our results comparing panoramic to flat backgrounds, we speculate that the flat backgrounds led to lower performance because they rarely aligned with the randomly chosen synthetic camera angle. This suggests that realism may be more important than we originally imagined.

Our synthetic data led to useful loss reduction for the real-world prediction task in early epochs, before diverging. This suggests that the distribution of our synthetic data needs to match the distribution of our real-world data more closely, again implying that more realism is better.

Synthetic data in combination with real world data was useful when real world data was limited, and decreased variance in test accuracy even when there was no improvement in mean accuracy. This suggests that synthetic data could be useful for decreasing uncertainty even when sufficient real-world data is available (perhaps by training models on unlikely scenarios).

Acknowlegements

This work was funded by Sandia National Laboratories’ Laboratory-Directed Research & Development program, under the Global Security Investment Area.