Self-Supervised Learning Captures Improved Spatial Variation from Satellite Imagery

David Carlson

Assistant Professor, Duke University

March 30, 2022

Consortium for
Enabling Technologies & Innovation

Much of deep learning relies on supervised data

- Remote sensing often has relatively little ground truth data
- Lots of unsupervised data!

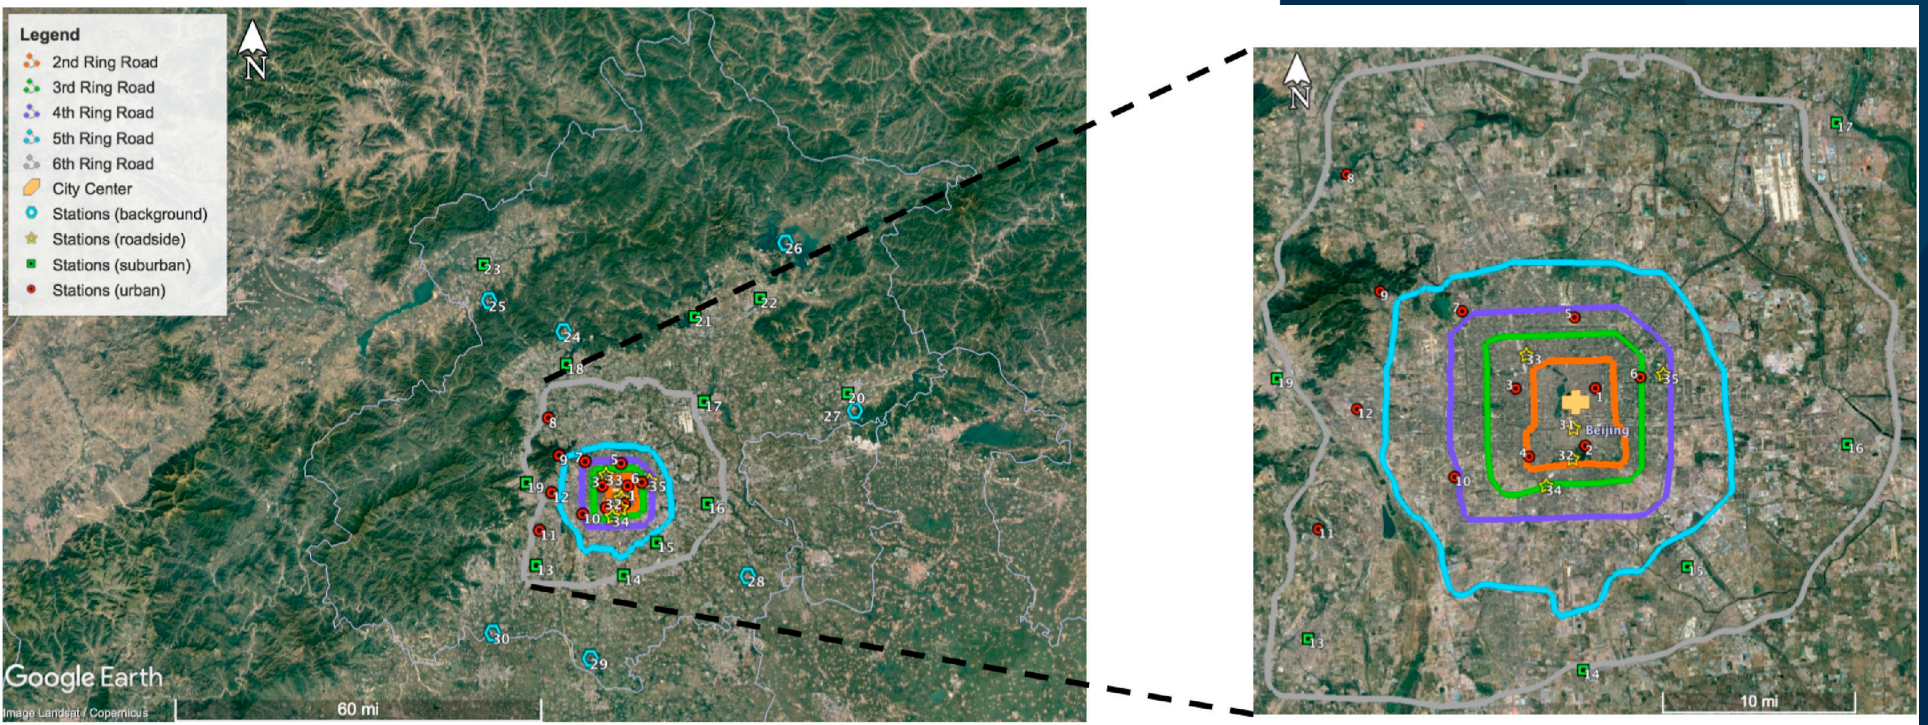

Few Locations with Sensors

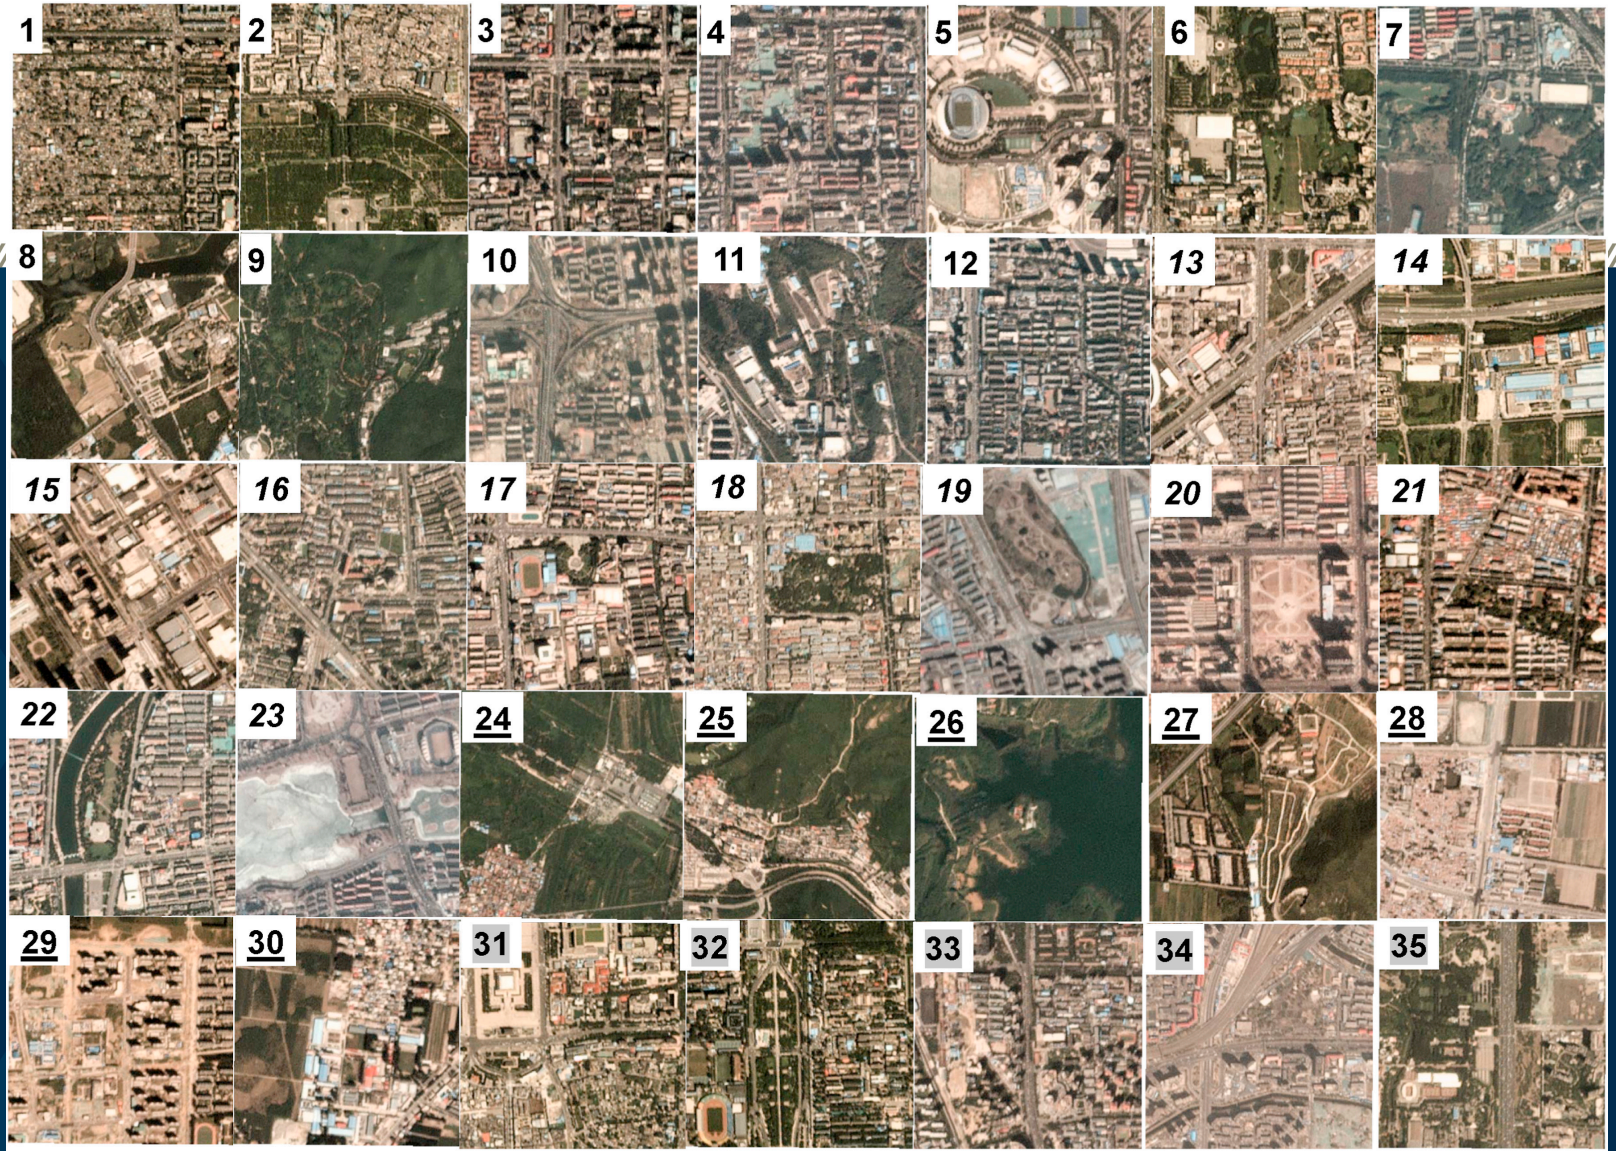

Standard Practice is to use a Convolutional Neural Network (CNN) with Transfer Learning

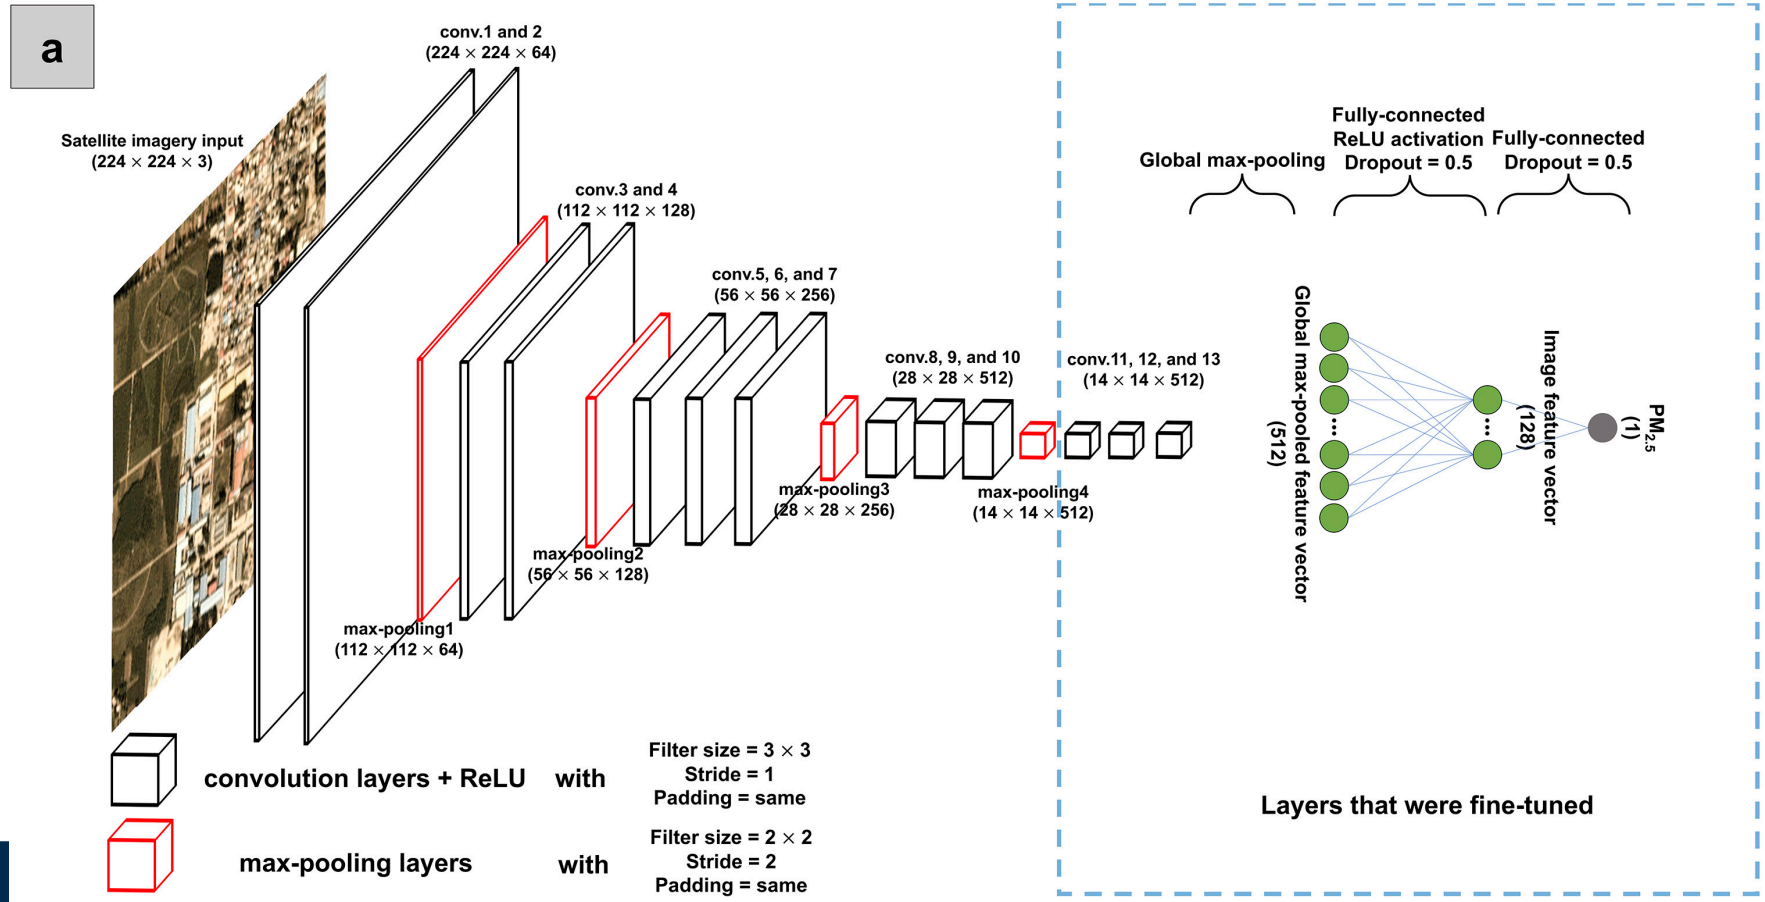

CNN parameters pretrained on ImageNet (Image Classification)

Zheng et al, Atmospheric Environment 2020

CNN is pre-trained on natural images

Child in a bee costume

An actual bee

Does pretraining on natural imagery capture the most relevant information?

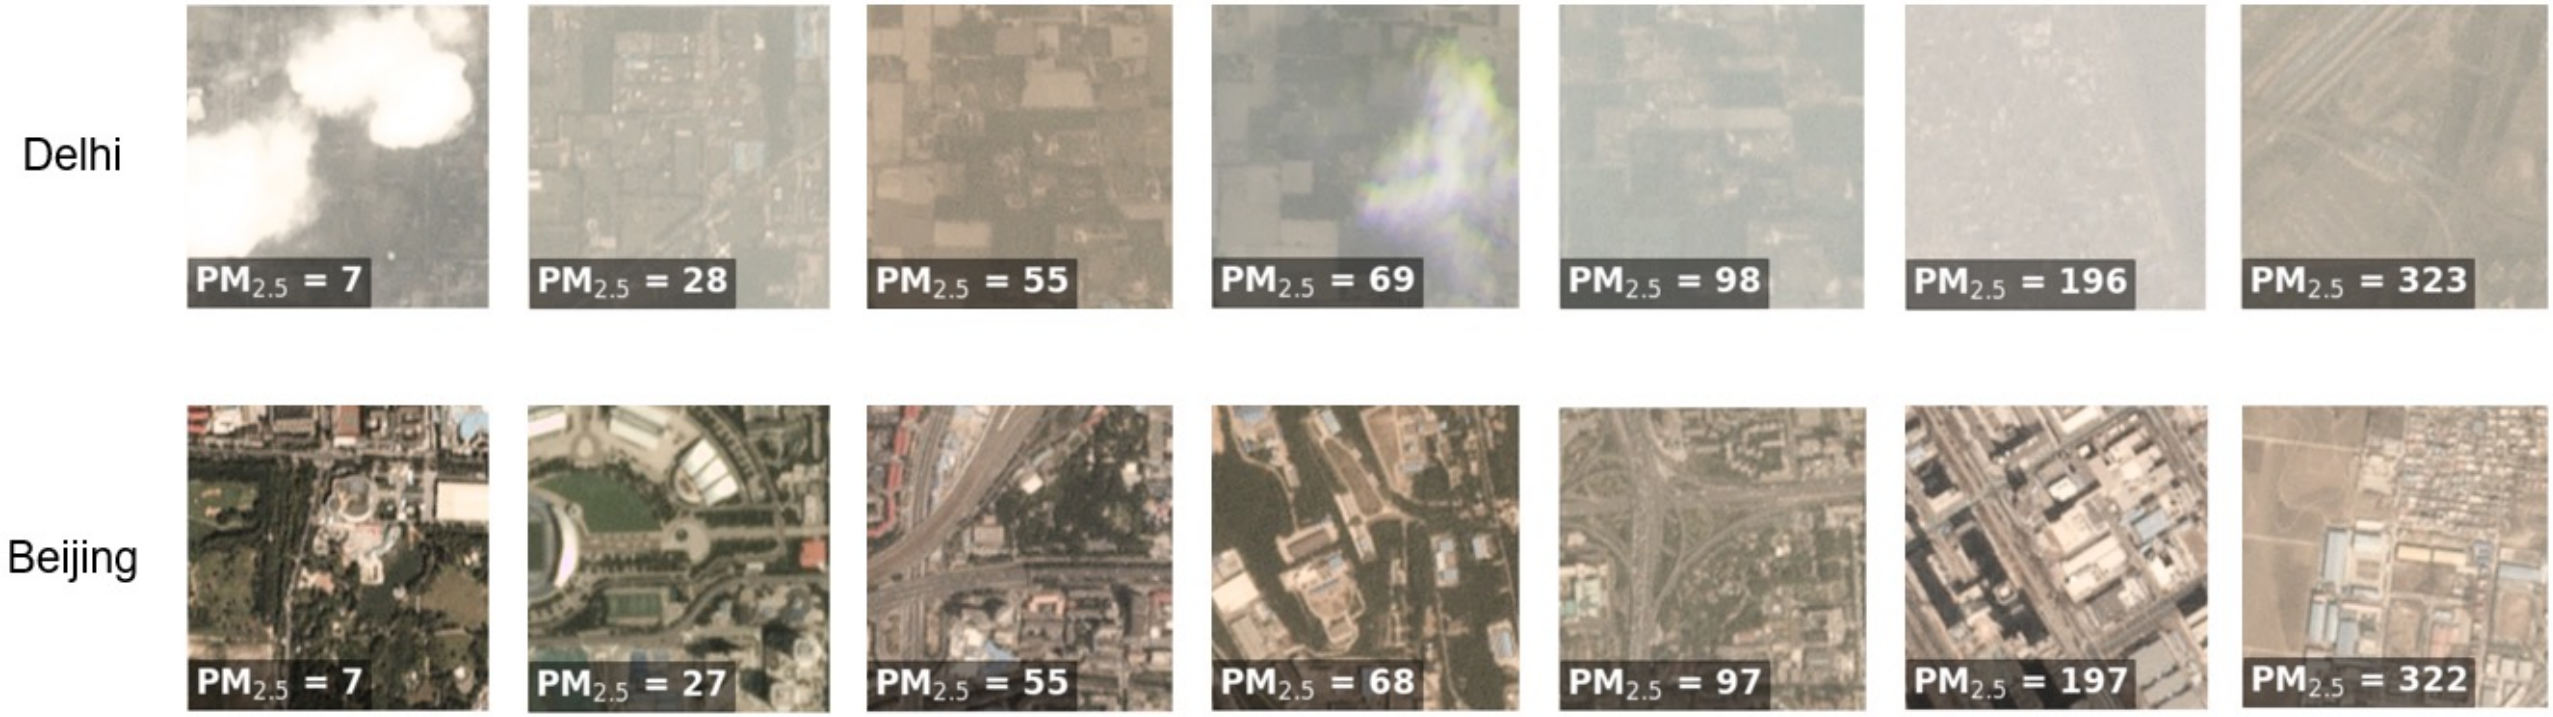

Self-Supervision Pre-Training

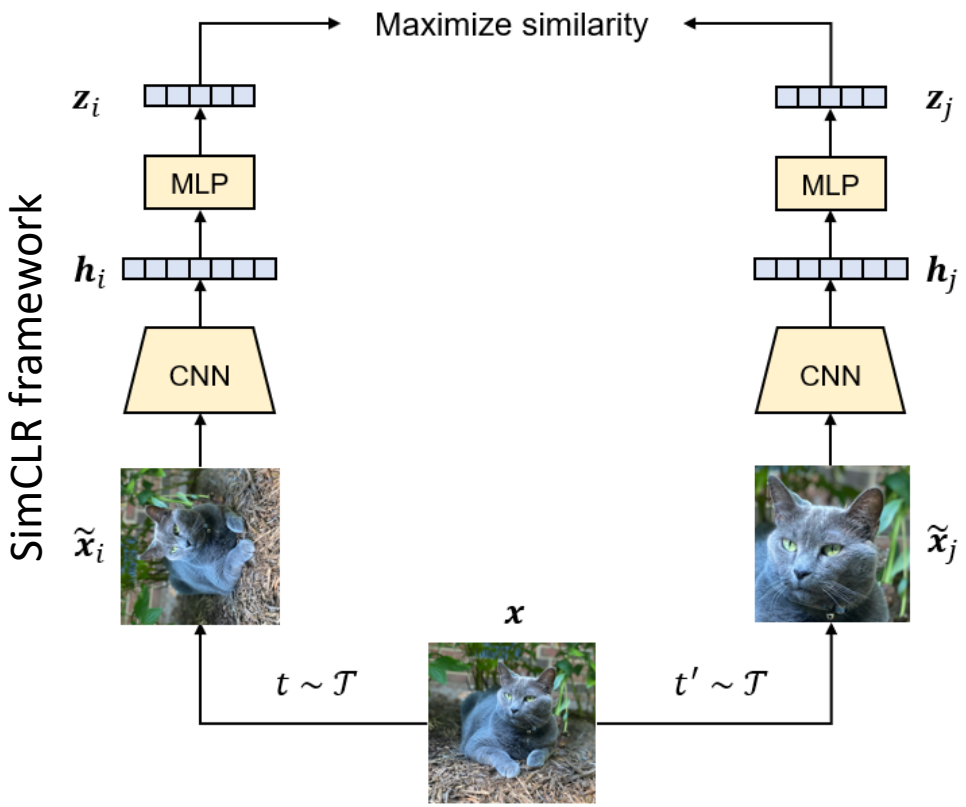

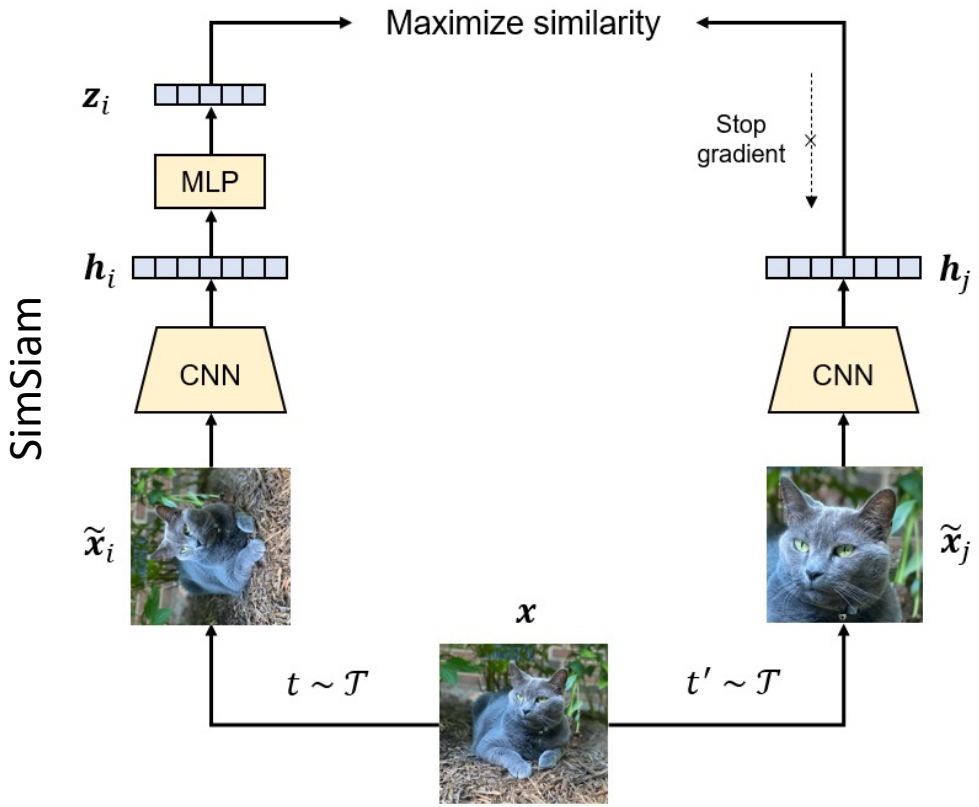

Massively effective for semi-supervised problems in natural imagery!

Self-Supervision with Spatiotemporal Similarity

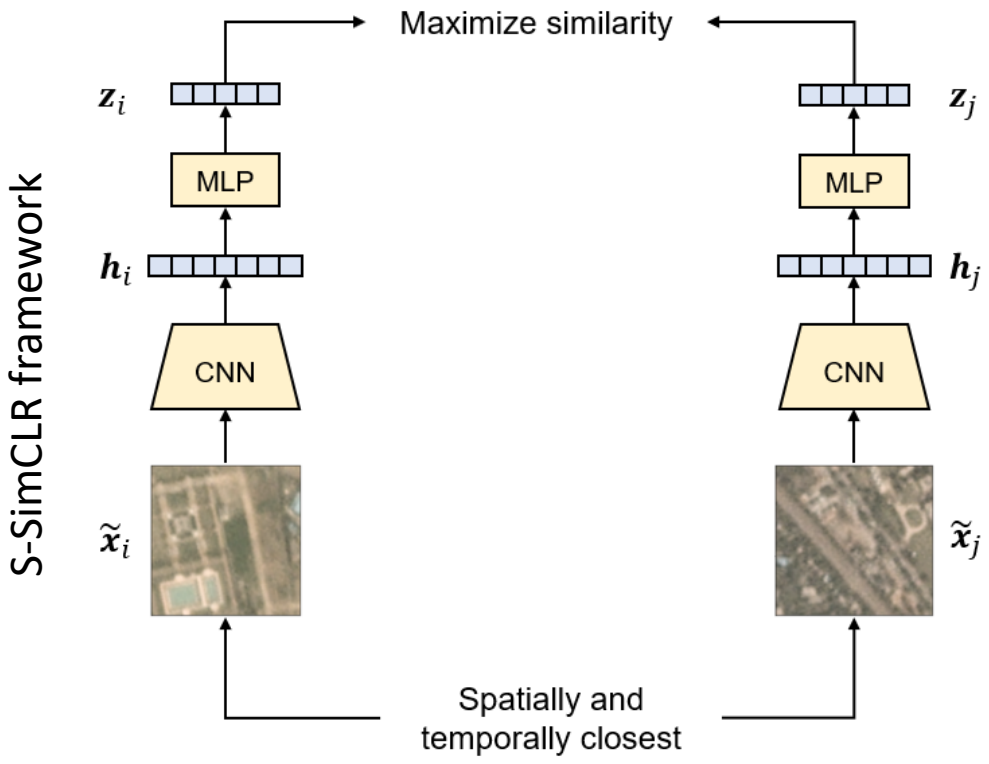

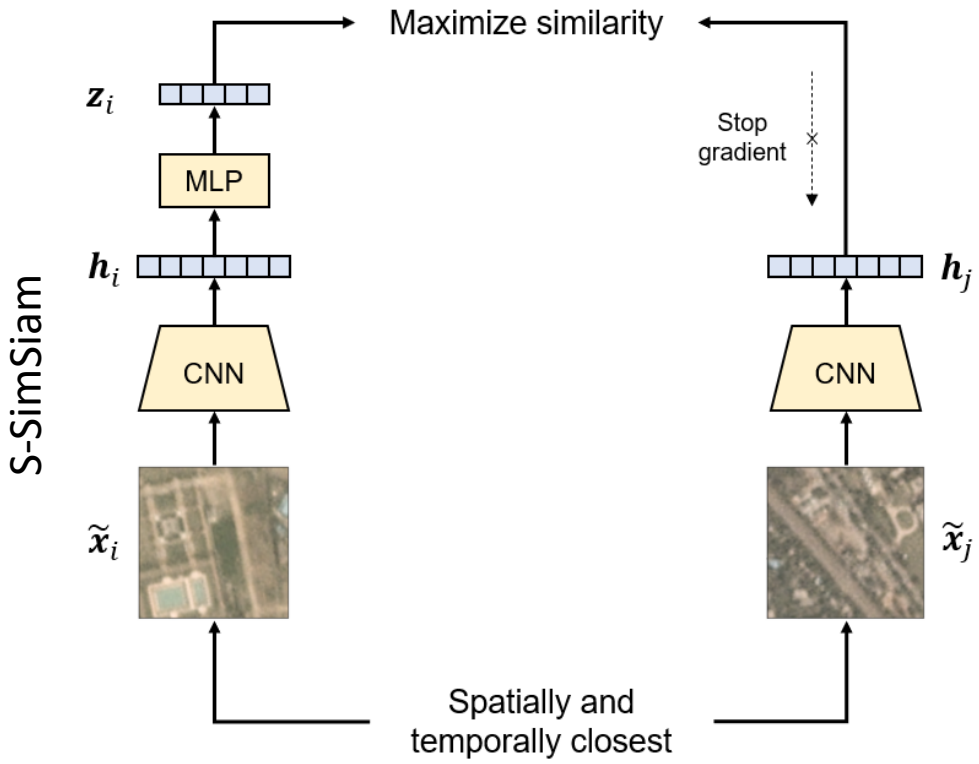

Fine-Tuning

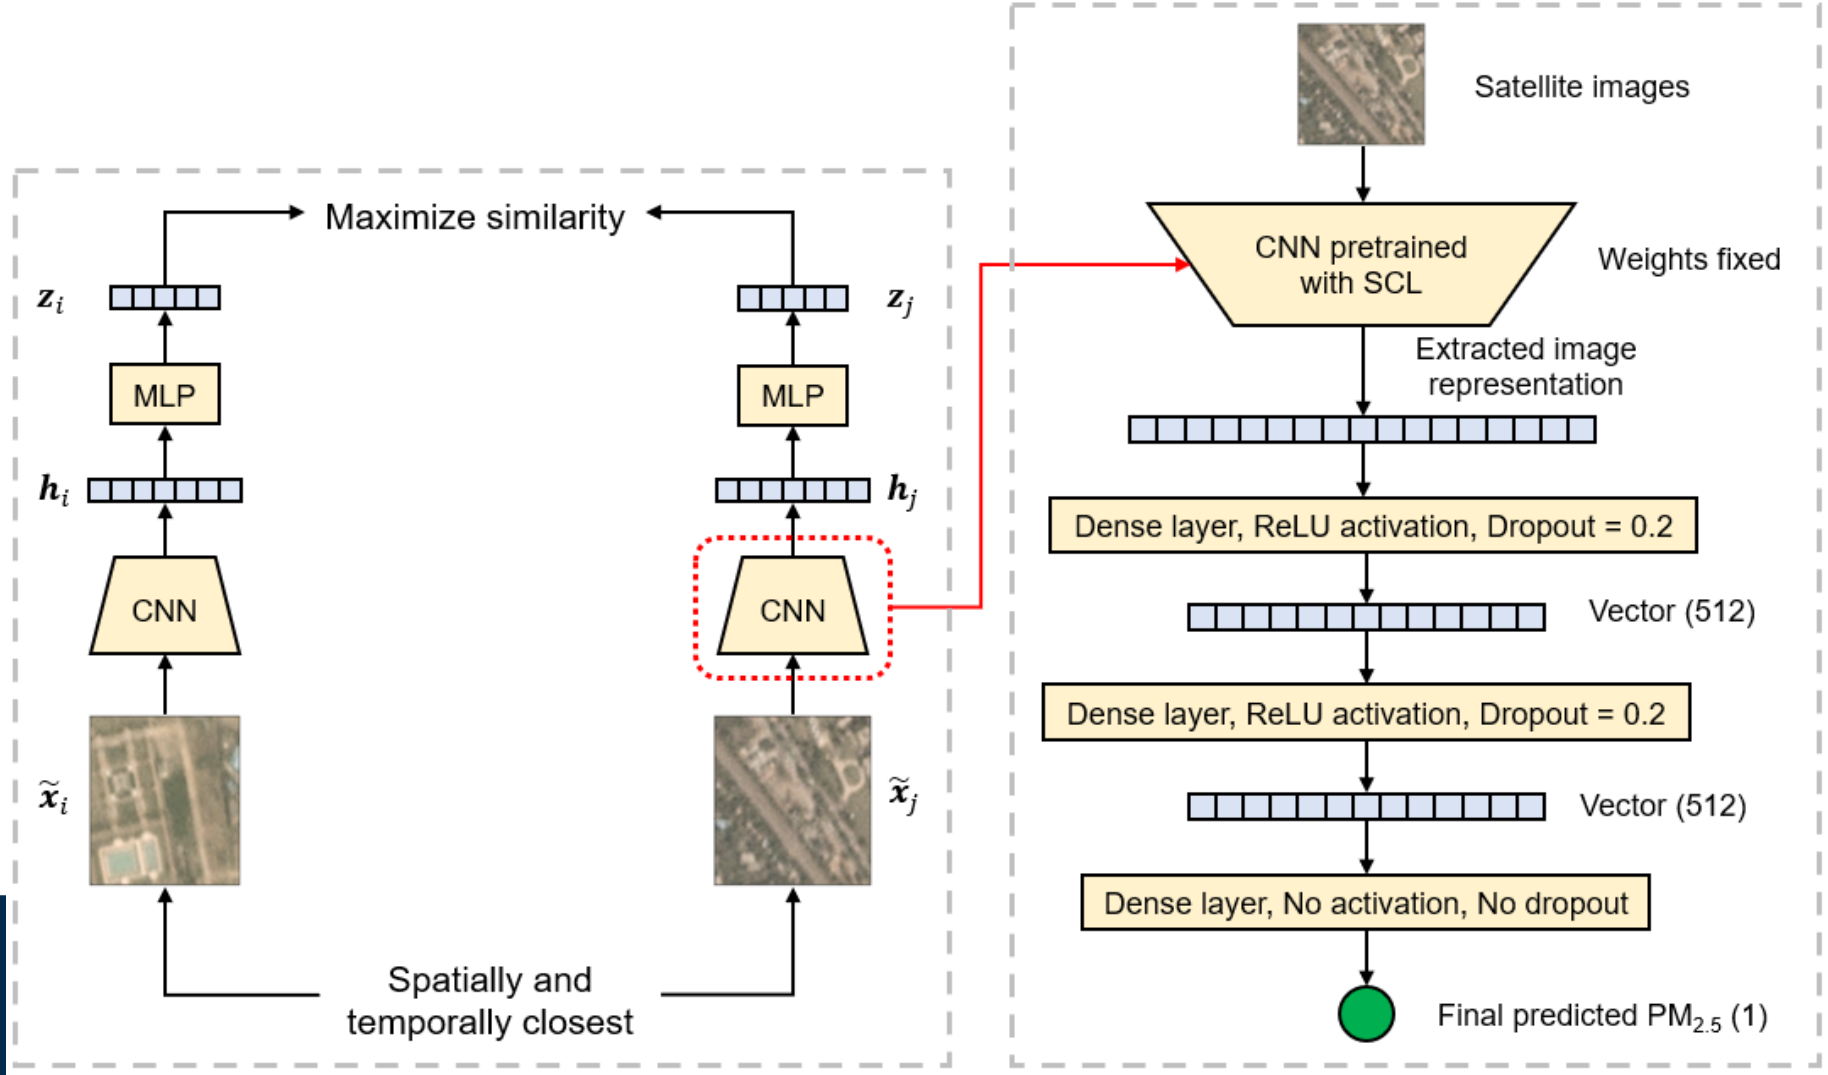
Maximize similarity
z_i
z_j
MLP
MLP
h_i
h_j
CNN
CNN
$\tilde{x}_i$
$\tilde{x}_j$
Spatially and temporally closest
Pretext task
Satellite images
CNN pretrained with SCL
Weights fixed
Extracted image representation
Dense layer, ReLU activation, Dropout = 0.2
Vector (512)
Dense layer, ReLU activation, Dropout = 0.2
Vector (512)
Dense layer, No activation, No dropout
Final predicted $PM_{2.5}$ (1)
Downstream task

Captures Significantly Improved Spatial Metrics

The spatial self-supervisions methods improved over transfer learning with:

- Significantly reduced Spatial RMSE at multiple locations
- Significantly improved Spatial Pearson r
- Spatial self-supervision improves over image augmentation self-supervision approaches (not shown)

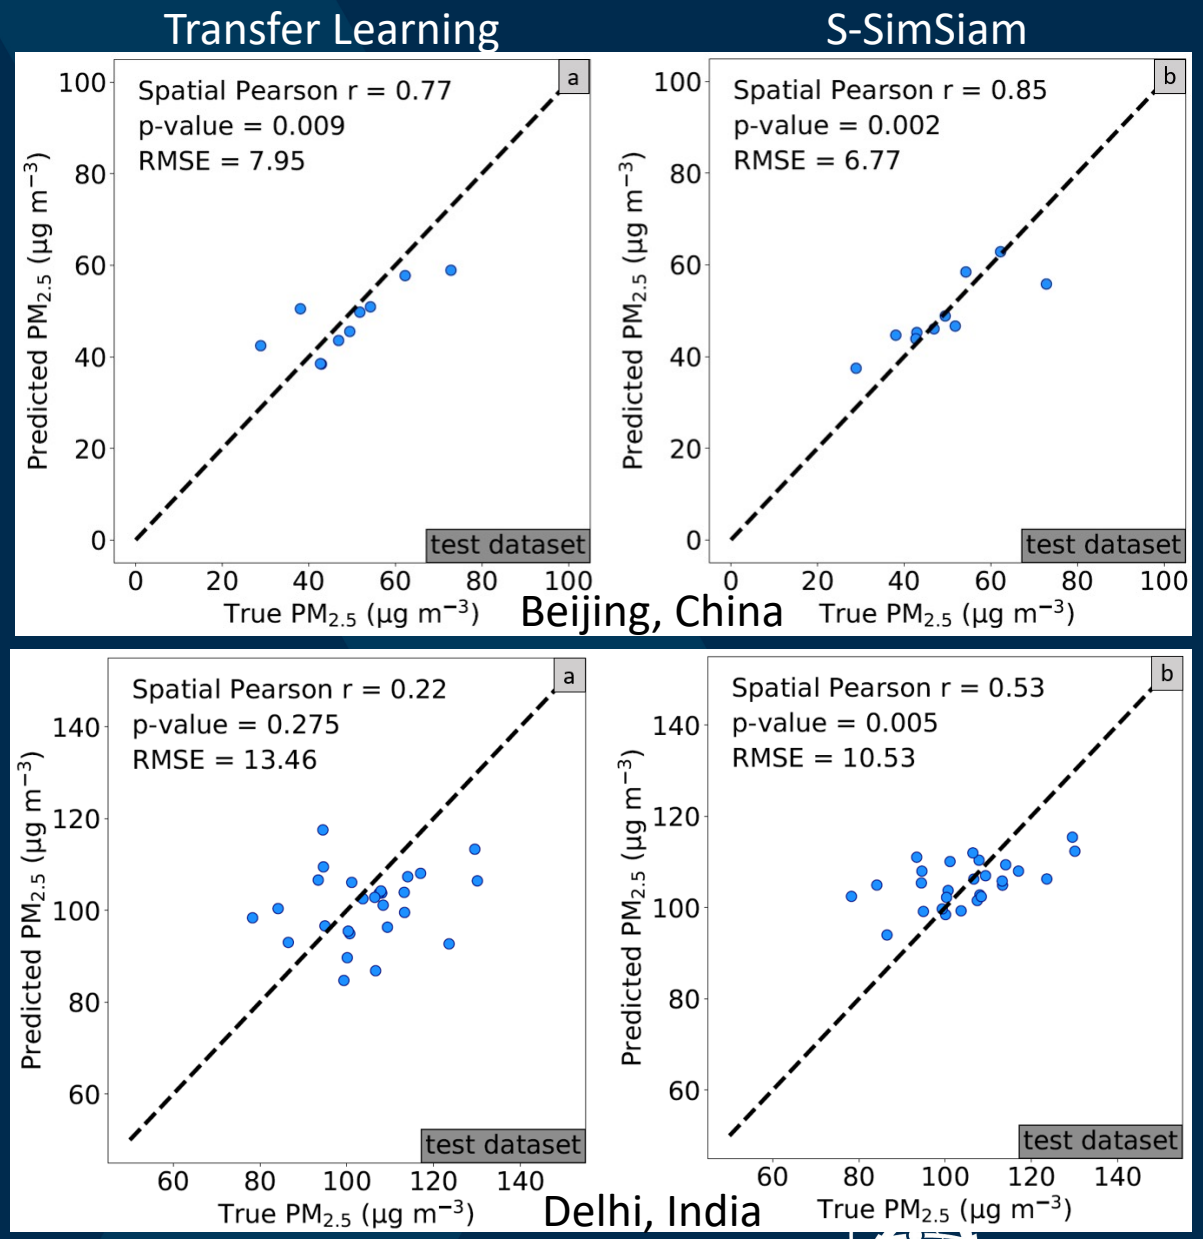

Capturing Uncertainty from Deep Networks

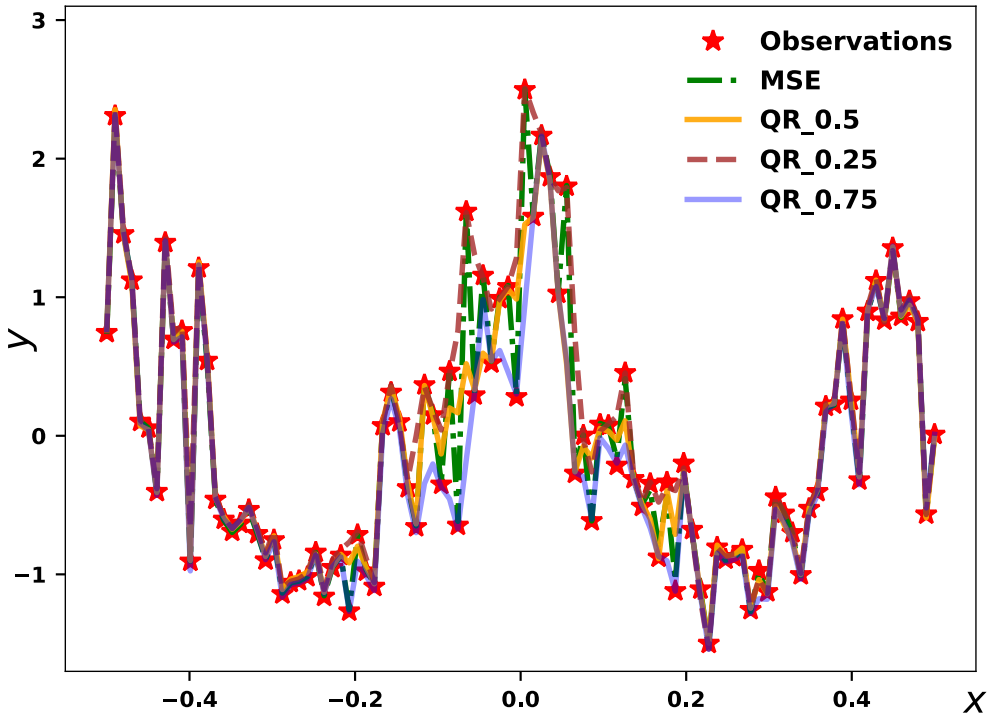

Quantile Regression with Deep Networks

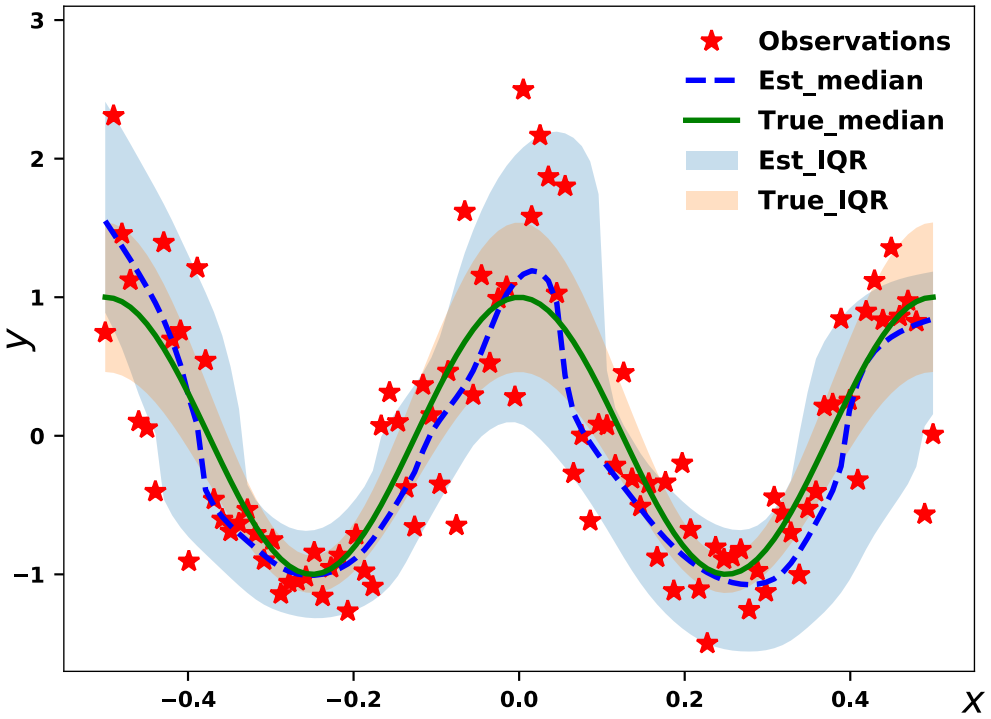

Collaborating Networks

Zhou et al., Journal of Machine Learning Research 2021

Uncertainty by Pairing Deep Networks

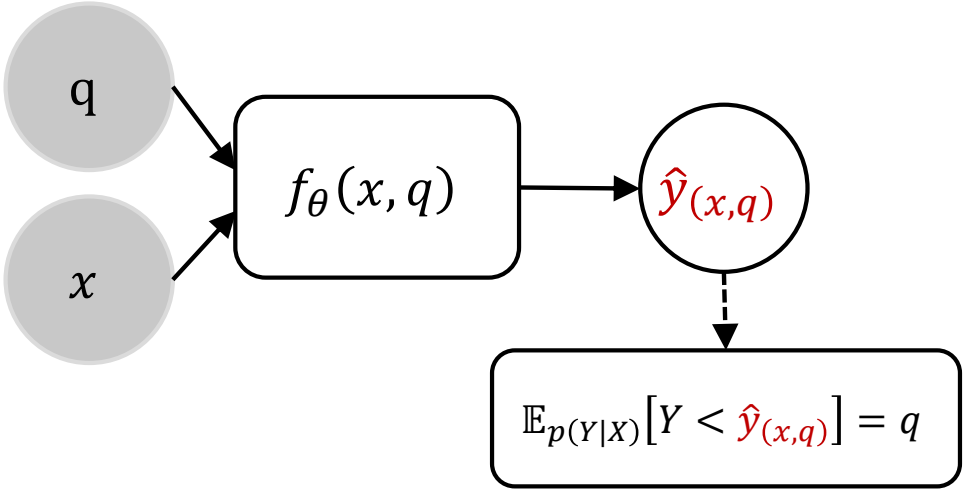

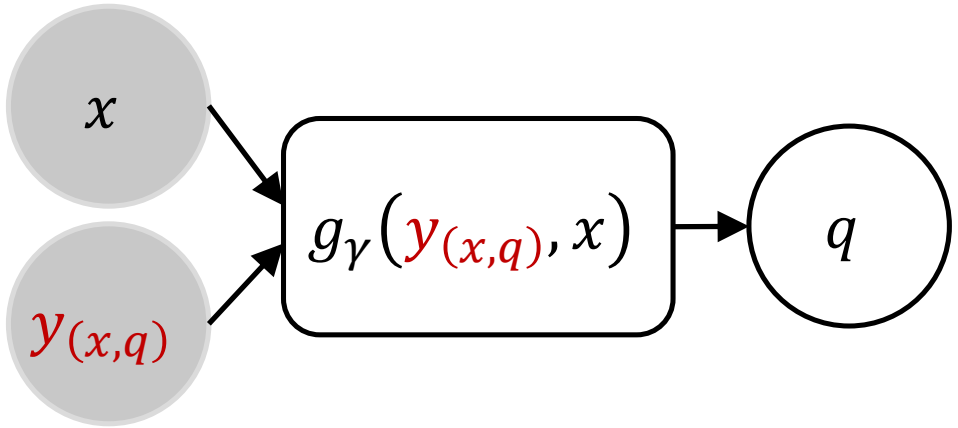

$$\text{f-loss}_\theta : \quad \mathbb{E}_{q\sim p(q),\boldsymbol{x}\sim p(X)}\left[(q - g_\gamma(f_\theta(q,\boldsymbol{x}),\boldsymbol{x}))^2\right]$$

$$\text{g-loss}_\gamma : \quad \mathbb{E}_{q\sim p(q),\boldsymbol{x},y\sim p(X,Y)}\left[\ell(1_{(y<f_\theta(q,\boldsymbol{x}))}, g_\gamma(f_\theta(q,\boldsymbol{x}),\boldsymbol{x}))\right]$$

Can define networks as convolutional neural networks!

Zhou et al., Journal of Machine Learning Research 2021

Integrating CNNs to Estimate Uncertainty

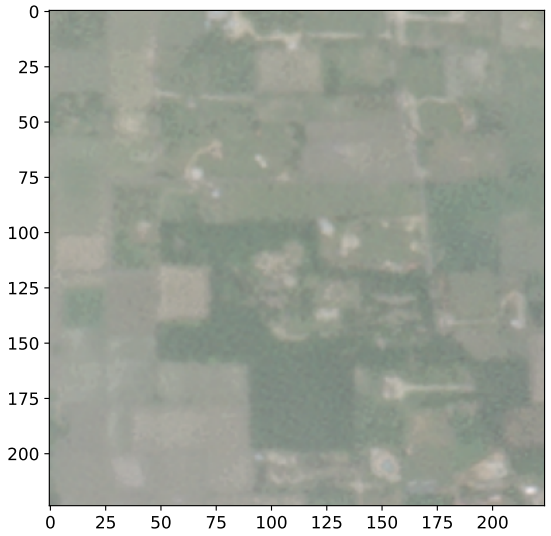

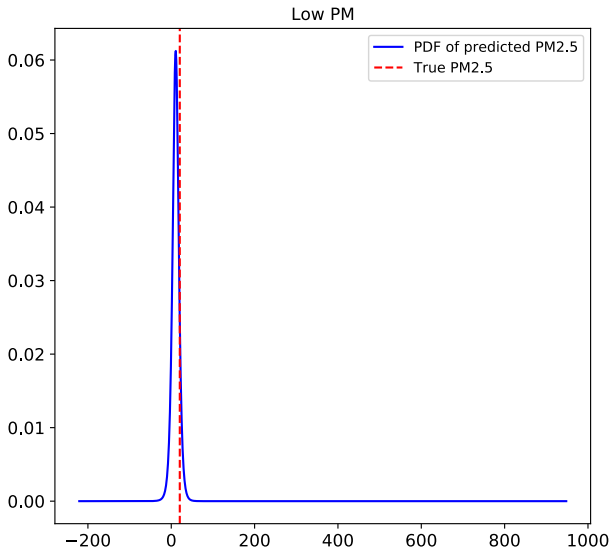

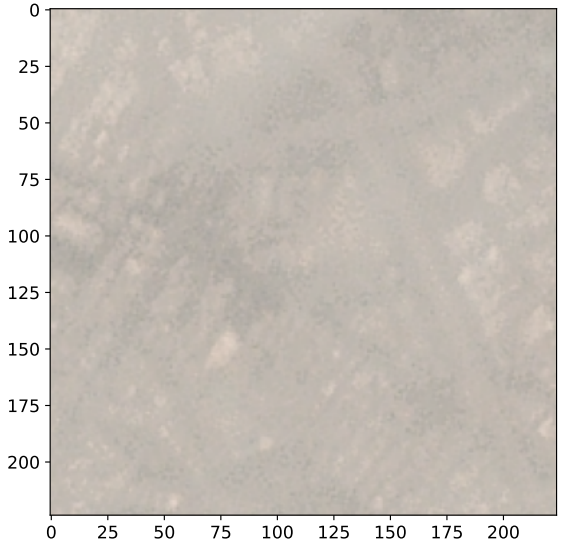

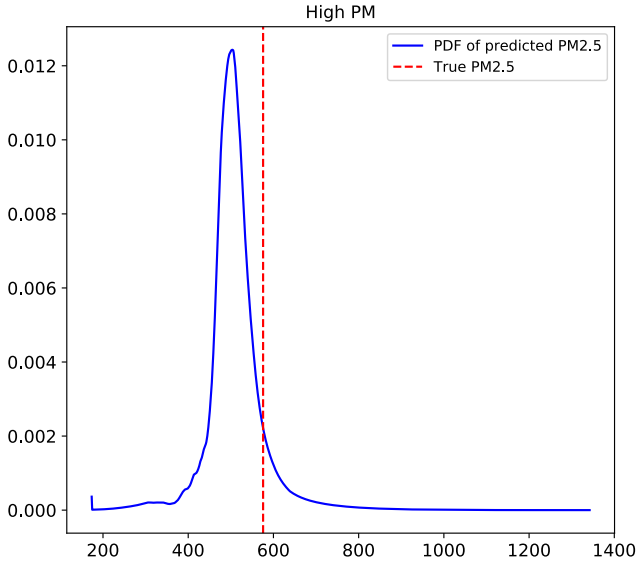

Captures Uncertainty Ranges Well

Calibration is good for relatively small data set

Covers ~80% of samples for a predicted 90% coverage range

Can do posthoc corrections

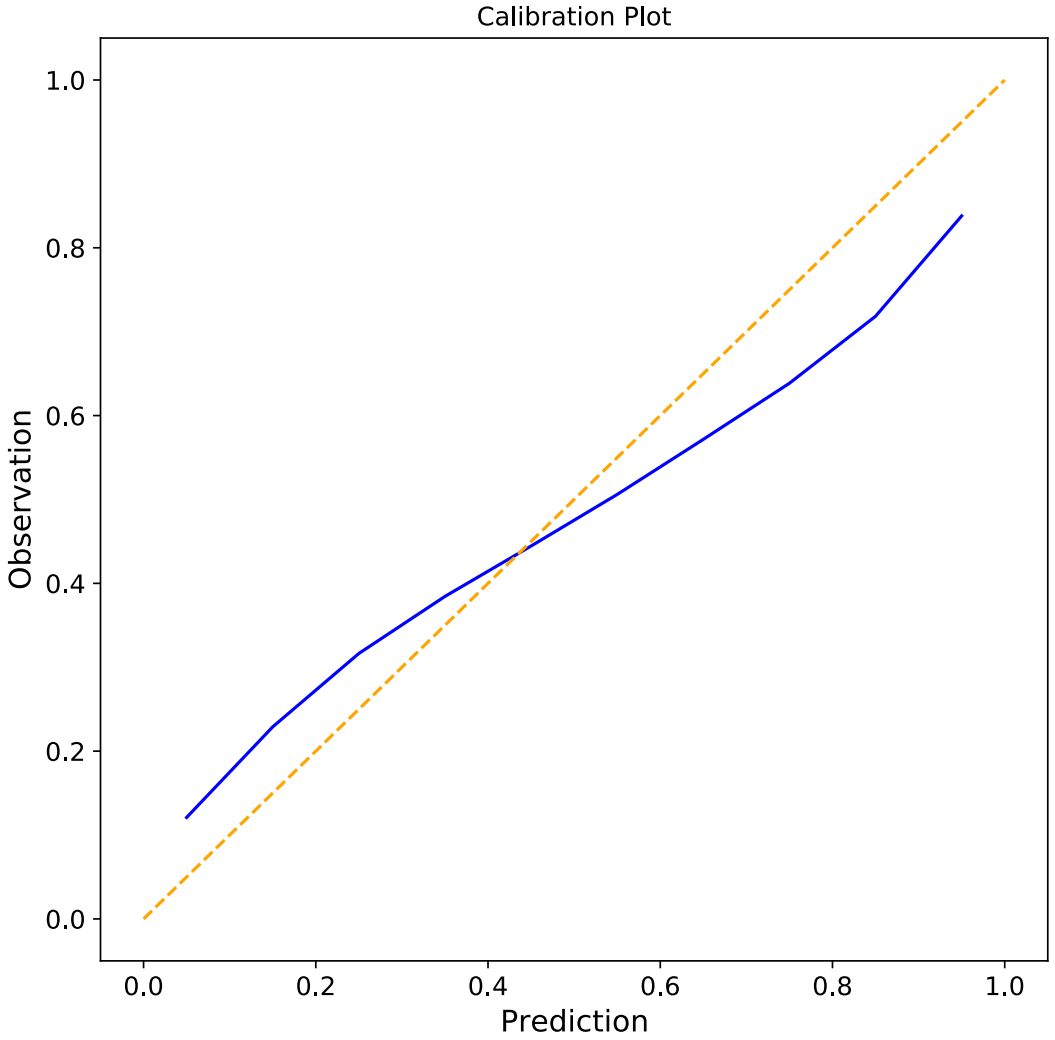

Summary

- Can choose self-supervision tasks to match our prediction goal
- Many tasks are not naturally matched with CNNs for natural images
- Can integrate modern techniques to get reliable uncertainty ranges from deep networks