In-N-Out: Using Auxiliary Information for OOD Robustness

Robbie Jones*, Sang Michael Xie*, Ananya Kumar*,
Fereshte Khani, Tengyu Ma, Percy Liang

     

Remote Sensing Applications

- Remote sensing used for tracking deforestation, crop land prediction, land cover classification

- Example: crop type classification

x: Landsat image

→ CNN → **y:** crop type

Obstacle: Lack of labeled data

- Prohibitively expensive to collect labels via surveys esp. in developing countries
- **Goal**: Make good predictions **globally**

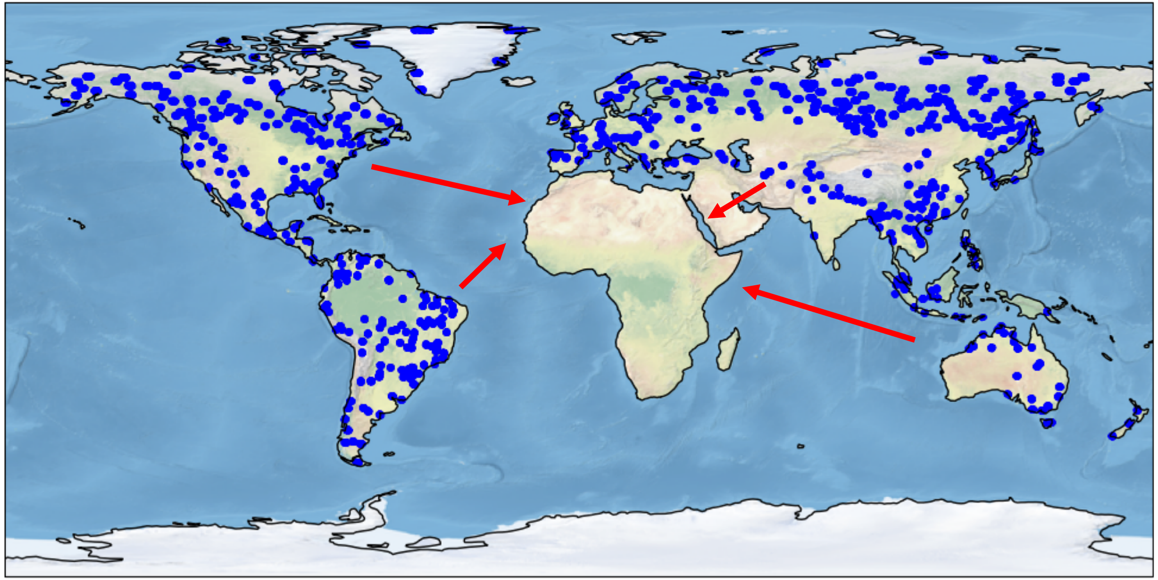

Incorporating auxiliary data as inputs

- Often have **auxiliary data** in addition to satellite input
  - E.g., temperature, precipitation, data from other satellites
- **Standard approach (aux-inputs)**: use auxiliary data as another input

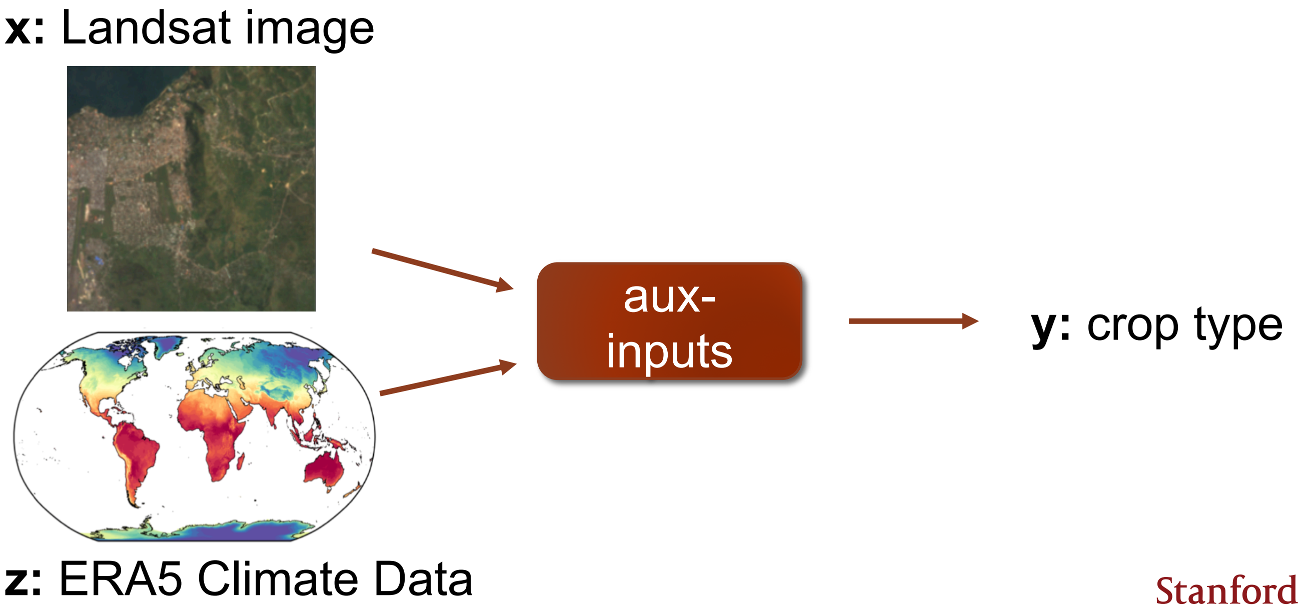

Using auxiliary data as inputs can hurt under-resourced countries

- Aux-inputs improves accuracy in countries with labeled data
- **We find that it worsens performance in under-resourced (OOD) countries**
  - 90% -> 84% on cropland prediction
  - 58% -> 55% on landcover prediction

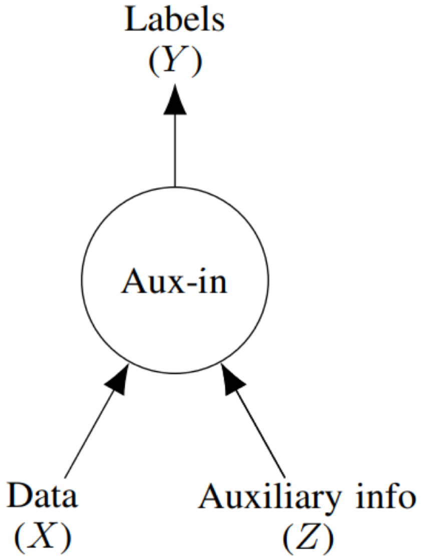

Incorporating auxiliary data as outputs

“**Aux-outputs**”: pre-train by predicting auxiliary information as outputs on unlabeled data

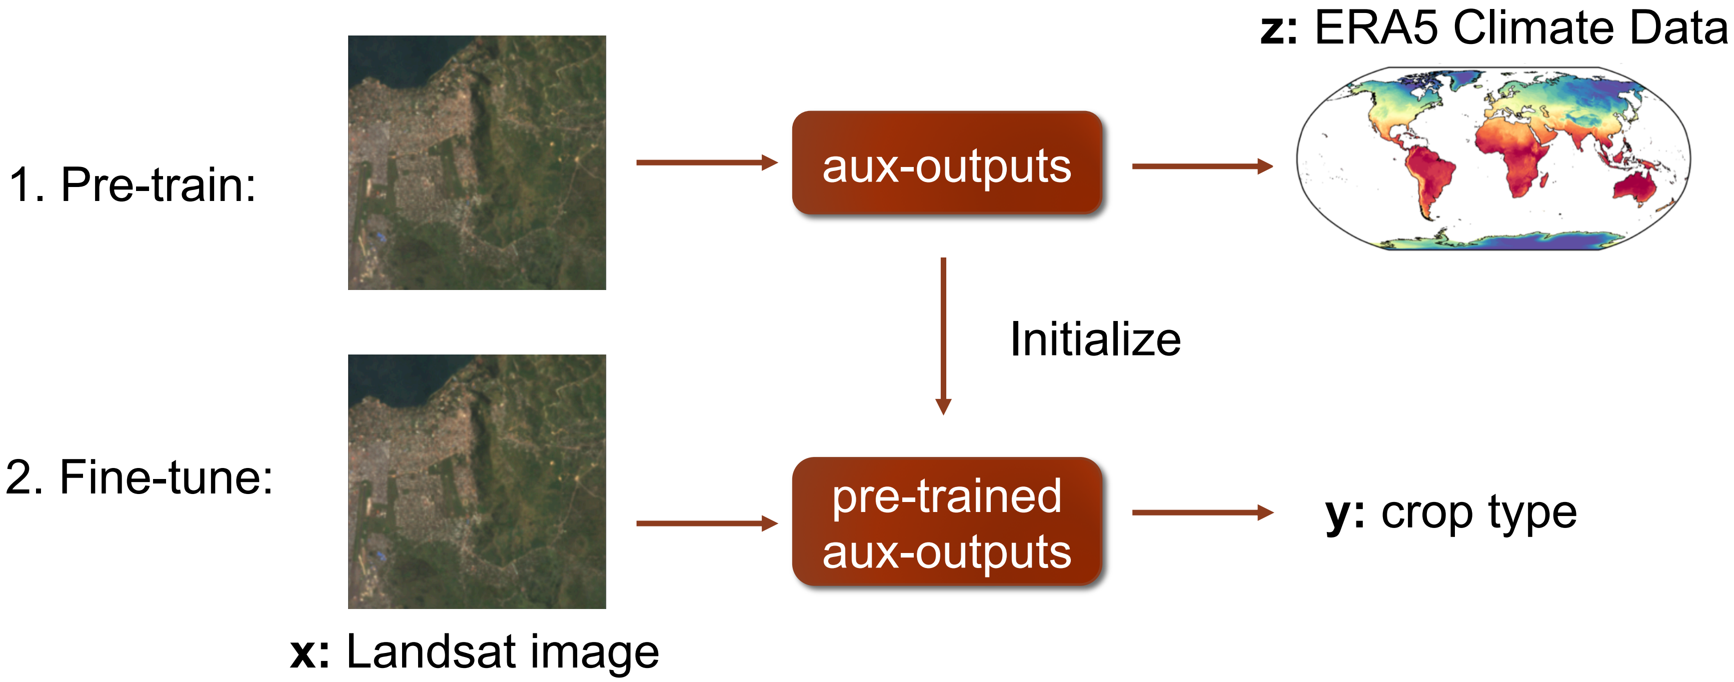

Predicting auxiliary data as outputs helps under-resourced countries

- Intuition: learns good initial features for all countries
  - We prove this improves OOD error theoretically for linear regression
- **Aux-outputs improves accuracy in OOD countries**
  - 90% -> 92% on cropland prediction
  - 58% -> 61% on landcover prediction

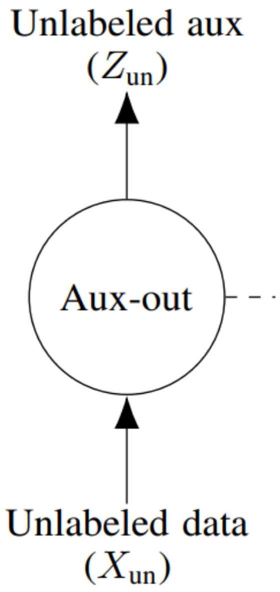

In-N-Out combines benefits from aux-inputs and aux-outputs

1. Train aux-inputs model (better in seen countries)
2. Pre-train aux-outputs (better for unseen countries)
3. Pseudolabel unlabeled data from seen countries with trained aux-inputs

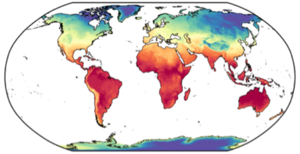

Trained aux-inputs

ŷ: crop type pseudolabel

4. Finetune aux-outputs model on labeled + pseudolabeled data

pre-trained aux-outputs

y: crop type

In-N-Out: Intuition

- Combine benefits of aux-inputs and aux-outputs using self-training
- Use aux-inputs to produce better pseudolabels for unlabeled data in seen countries
- More data improves fine-tuning of a pretrained aux-outputs model

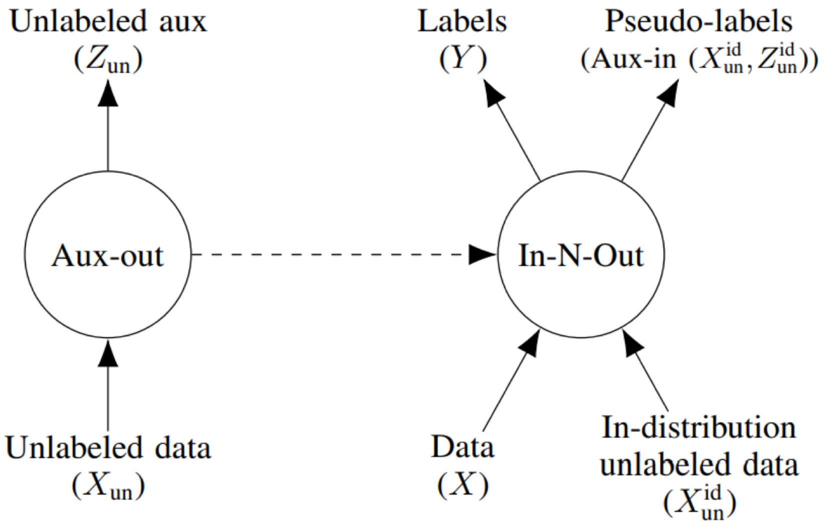

Datasets

Dataset Name	CelebA	Cropland	Landcover
Data Type	RGB Image	Landsat Image	MODIS Time Series
Visualization (x)			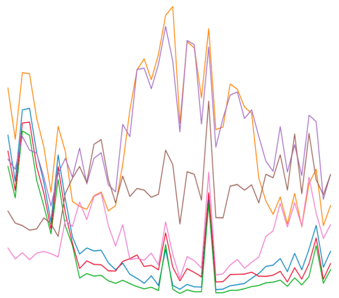
Aux Info (z)	7 binary attributes	Vegetation, Lat/Lon	Meteorological Data
Target (y)	Male/female?	Cropland/not cropland?	Land cover class
ID-Split	People without hats	IA, MN, IL	Outside Africa
OOD-Split	People with hats	IN, KY	Africa

In-N-Out does better in theory and practice

- We prove in a linear regression setting that In-N-Out improves *both* OOD and in-distribution accuracy
  - OOD Accuracy

	Celeb-A	Cropland	Landcover
Standard Training	72.6%	90.3%	58.3%
In-N-Out	**80.4%**	**92.2%**	**62.6%**

  - In-distribution Accuracy

	Celeb-A	Cropland	Landcover
Standard Training	90.5%	94.5%	75.9%
In-N-Out	**93.8%**	**95.5%**	**77.1%**