Visual Intelligence Lab
UC SAN DIEGO

UCSDCSE
Computer Science and Engineering

SCAN ME

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

Poster #287 WED (06/21) PM

GeoNet: Benchmarking Unsupervised Adaptation across Geographies

Tarun Kalluri *Wangdong Xu* *Manmohan Chandraker*

GeoNet Overview

1. GeoNet Dataset

205 #classes 600

400,000 #images 300,000

Scene Recognition | *Object Classification*

Restaurant

Freight Car

Shopfront

Headlight

Geographical Image Distribution in GeoNet

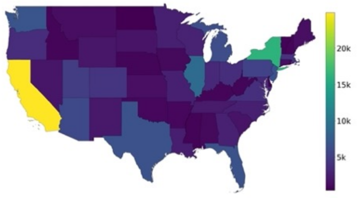

USA

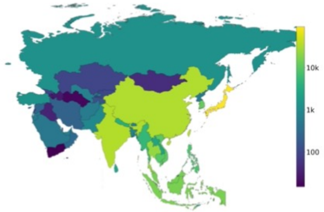

Asia

We study the significant accuracy drop observed on images from under-represented geographies.

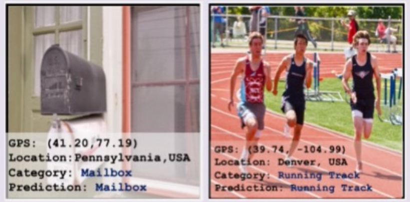

Training Domain USA

Context Shift

Design Shift

Testing Domain Asia

2. GeoNet Analysis

$$P(x, y) = P(b_x|y) \cdot P(f_x|y) \cdot P(y)$$

Context Shift

$P_s(b_x|y) \neq P_t(b_x|y)$

Background of scenes shift across domains

Design Shift

$P_s(f_x|y) \neq P_t(f_x|y)$

Design of objects change across domains

Label Shift

$P_s(y) \neq P_t(y)$

Class-distribution shift across geographies

Label Distribution in GeoNet

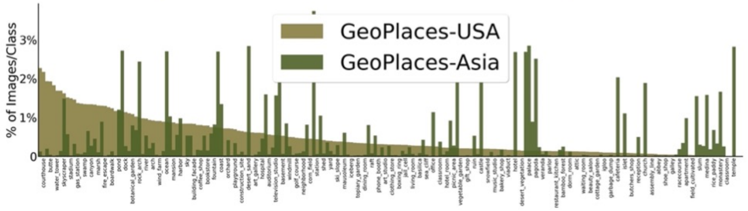

3. GeoNet Benchmarking

Cross-domain Performance

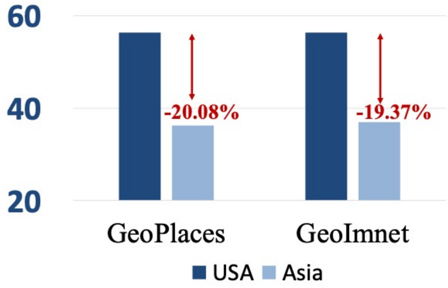

Unsupervised Domain Adaptation

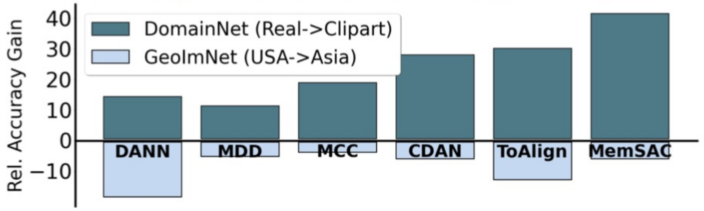

Large-Scale Pretraining

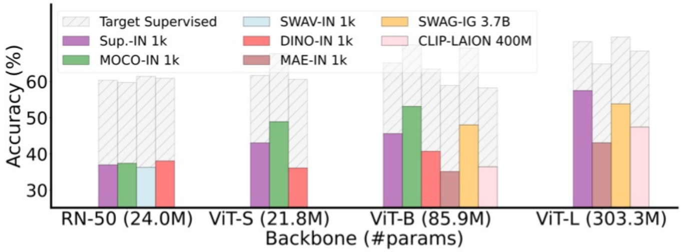

Unsupervised adaptation and large-scale pre-training do not suffice for bridging geographical disparity between domains.

Robustness in Computer Vision

- Models trained on one domain perform poorly on new domains encountered at test-time.
- Dataset-bias prevents *generalization*.

Training Domain

≠

Testing Domain

Acc = 90%

cartoons

Acc = 60%

real world

Prior Works

- **Style**
- **Capture Variations**
- **Photo-realism**
- **Lighting/Brightness**
- **Pose/Shape**

Geographical Bias in Datasets

Training Domain
USA

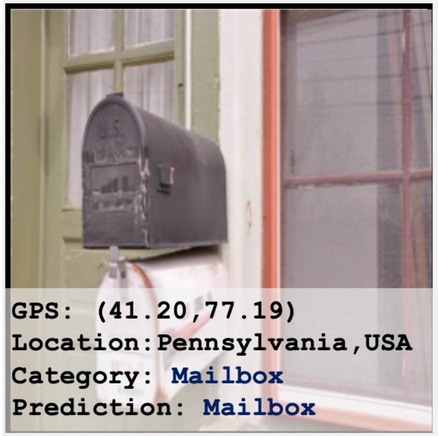

Context Shift

Design Shift

Testing Domain
Asia

- Models trained on these biased datasets generalize poorly to new geographies.
  - Where the model is nevertheless deployed.
- Has deep implications on *fairness* and *inclusivity*.
  - Model deployed on low-resource demography showcases poor performance.
  - Unfair towards targeted sub-populations.

GeoNet Contributions

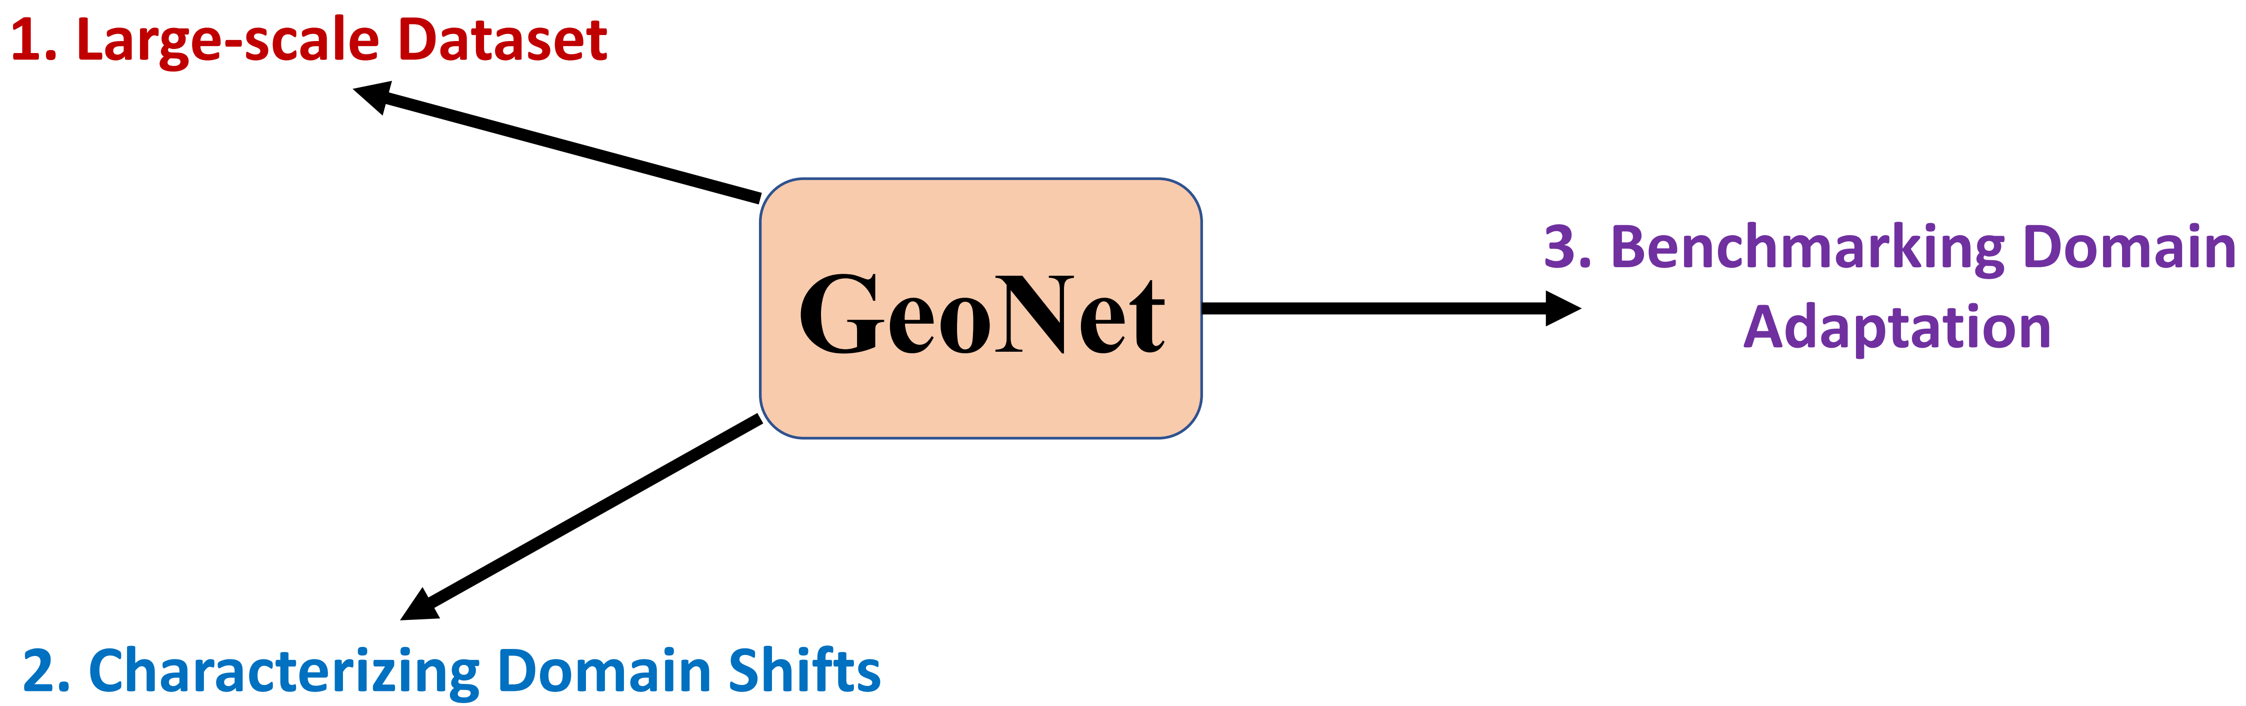

GeoNet Dataset

- GeoNet has data for scene classification and object classification.

GeoPlaces

Scene Recognition

Restaurant

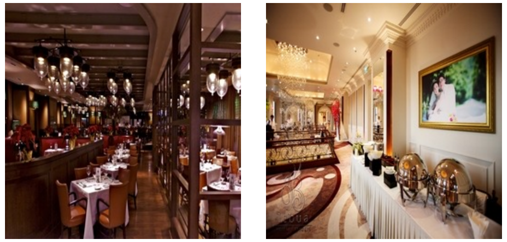

Shopfront

GeoImnet

Object Classification

Freight Car

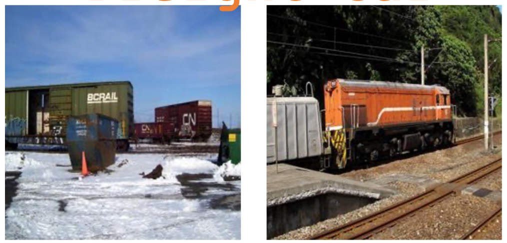

Headlight

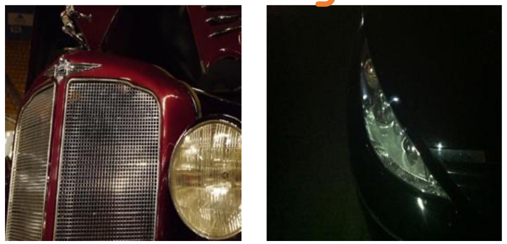

GeoPlaces		GeoImnet
205	#classes	600
400,000	#images	300,000

GeoNet Dataset

- GeoNet has data for scene classification and object classification.
- Largest d

Dataset publicly available, scan the following QR Code!

SCAN ME

Scene Recognition

205

400,000

#Images

600

300,000

Object Classification

Characterizing Distribution Shifts

- With some reasonable assumptions, we can split the joint image-label probability into context, design and label shifts.

- b_X: **Background.**
- f_X: **Foreground**
- **P(y): label distribution**

$$P(x,y) = \underbrace{P(b_x|y)}_{\text{context}} \cdot \underbrace{P(f_x|y)}_{\text{design}} \cdot \underbrace{P(y)}_{\text{prior}}$$

$$P_s(b_x|y) \neq P_t(b_x|y)$$

Background of scenes shift across domains

Characterizing Distribution Shifts

- With some reasonable assumptions, we can split the joint image-label probability into context, design and label shifts.

- **b_X: Background.**
- **f_X: Foreground**
- **P(y): label distribution**

$$P(x, y) = \underbrace{P(b_x|y)}_{\text{context}} \cdot \underbrace{P(f_x|y)}_{\text{design}} \cdot \underbrace{P(y)}_{\text{prior}}$$

$$P_s(f_x|y) \neq P_t(f_x|y)$$

Design of objects changes across domains

Characterizing Distribution Shifts

- With some reasonable assumptions, we can split the joint image-label probability into context, design and label shifts.

- **b_X: Background.**
- **f_X: Foreground**
- **P(y): label distribution**

$$P(x, y) = \underbrace{P(b_x|y)}_{\text{context}} \cdot \underbrace{P(f_x|y)}_{\text{design}} \cdot \underbrace{P(y)}_{\text{prior}}$$

Class-distribution shift across geographies

$$P_s(y) \neq P_t(y)$$

Cross-Domain Performance Gaps

Cross-Domain Performance Gaps

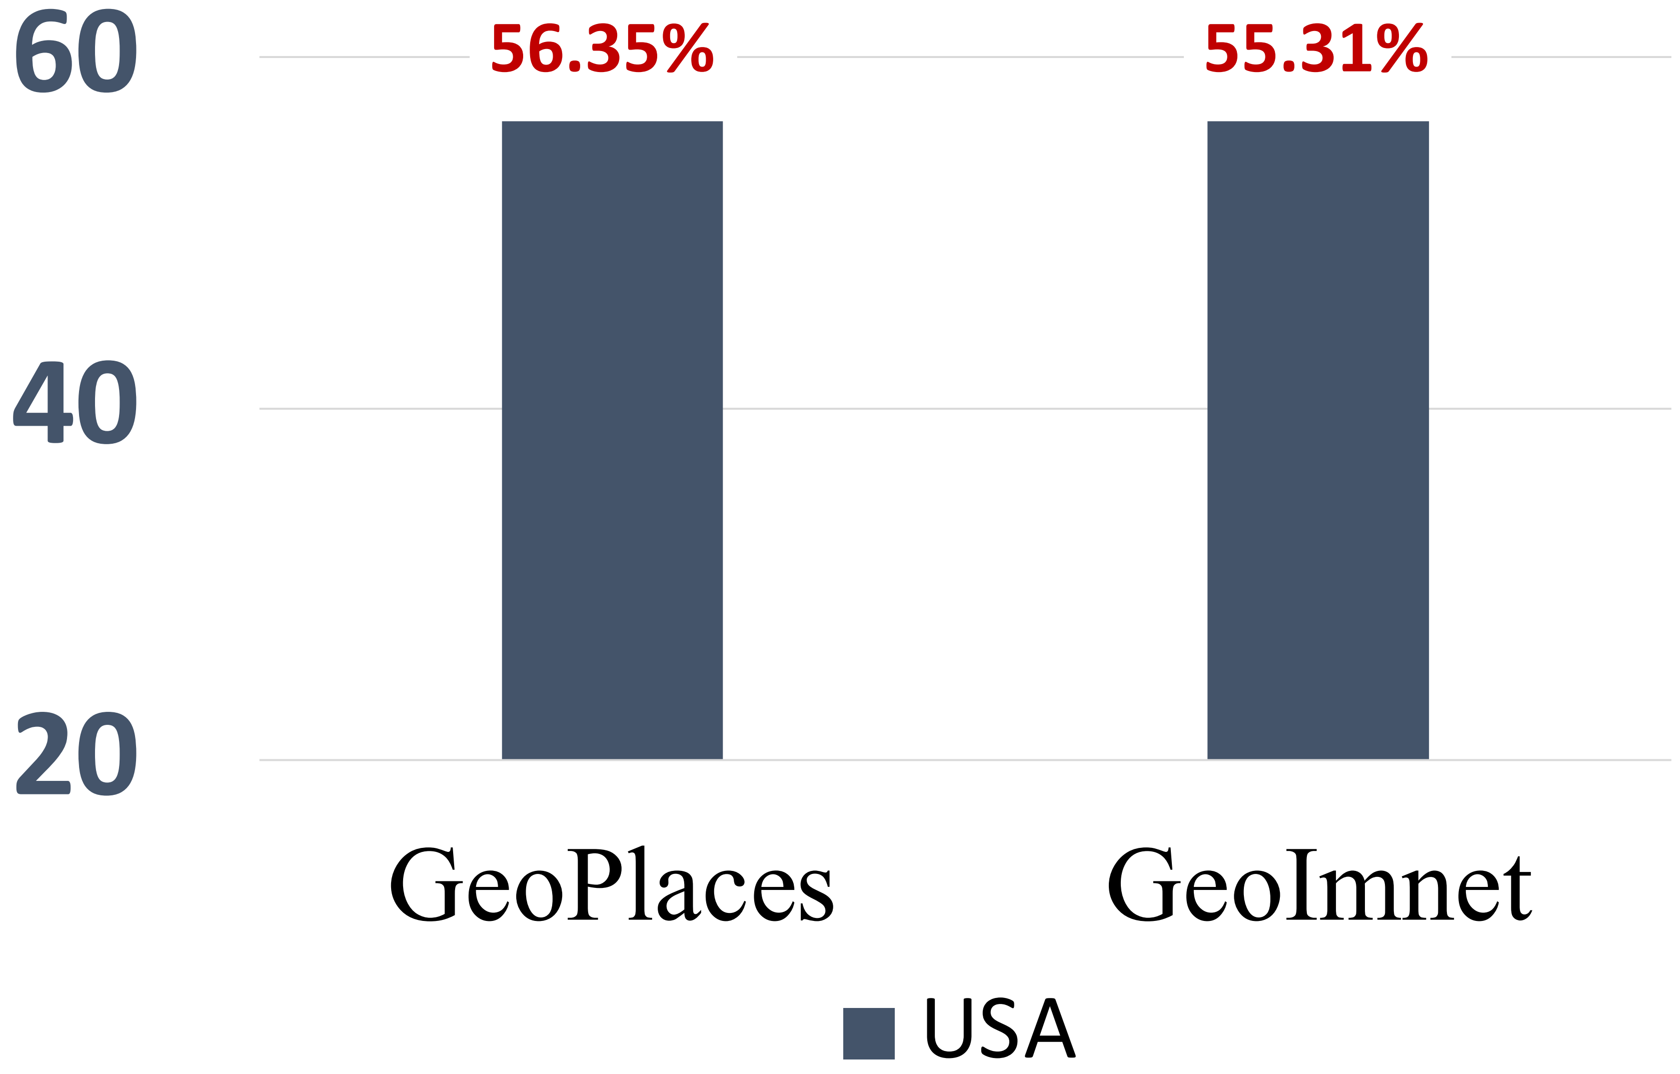

Cross-Domain Performance Gaps

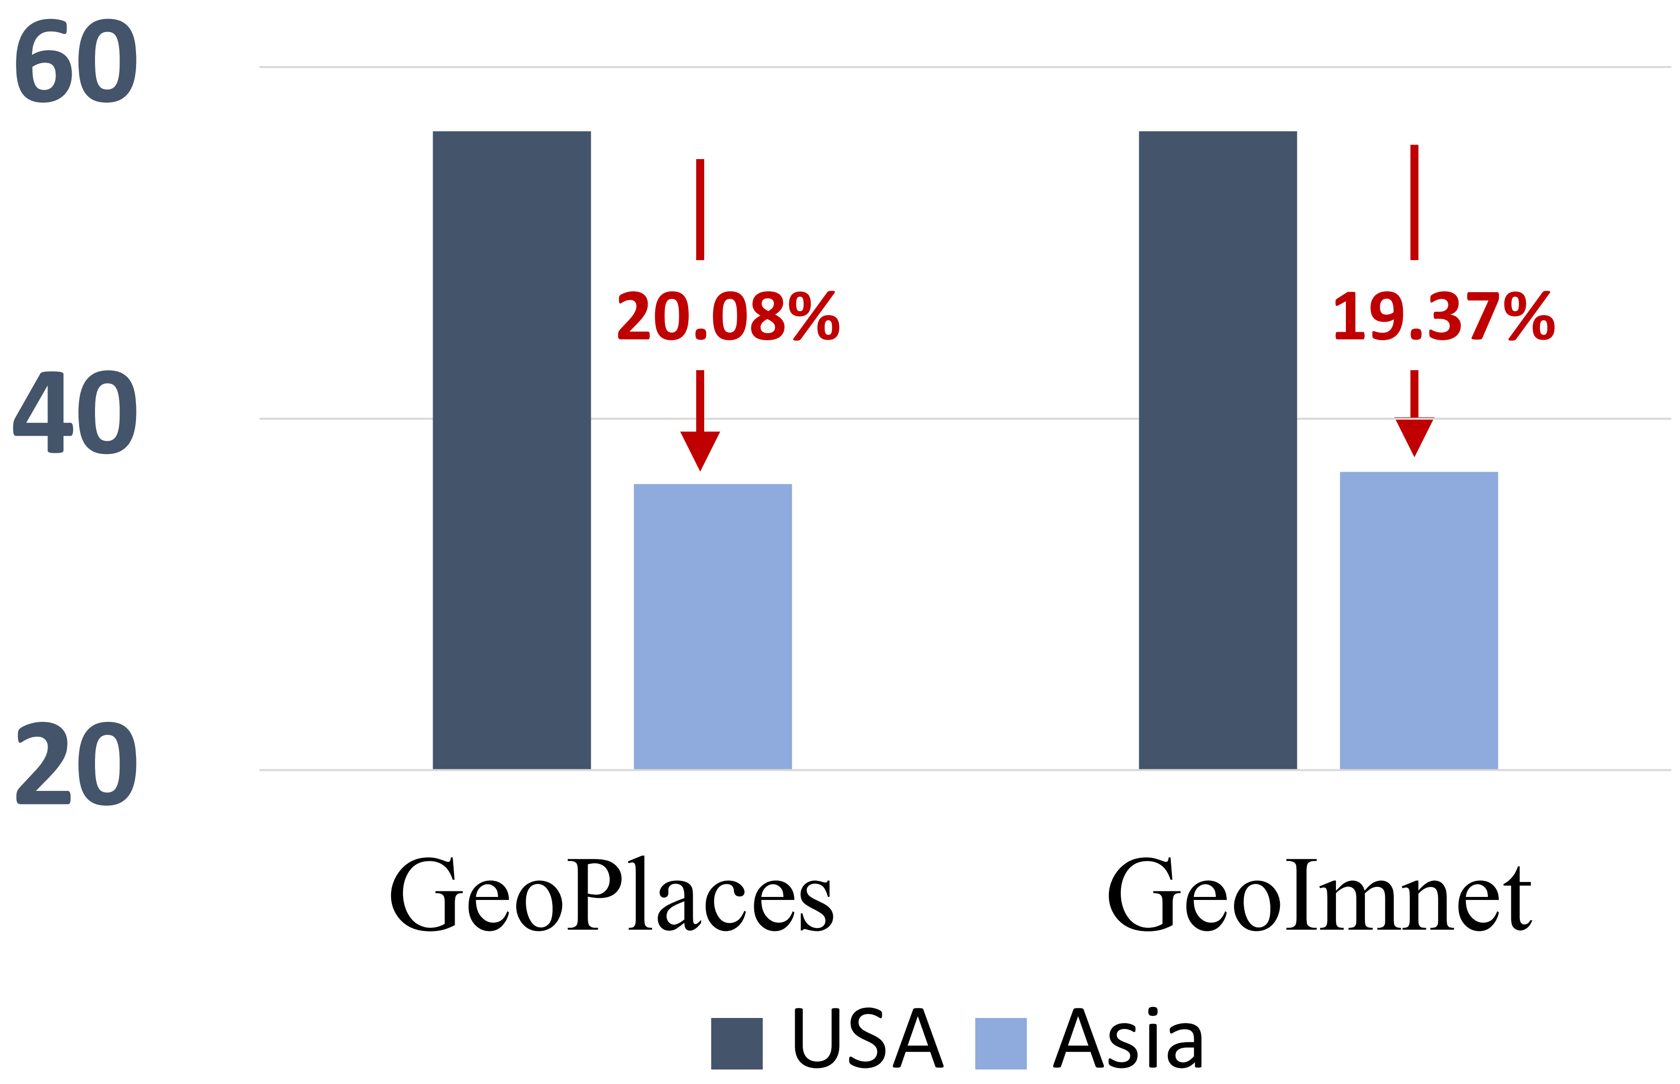

Benchmarking Unsupervised Adaptation on GeoNet

- UDA methods generally designed for covariate shifts.

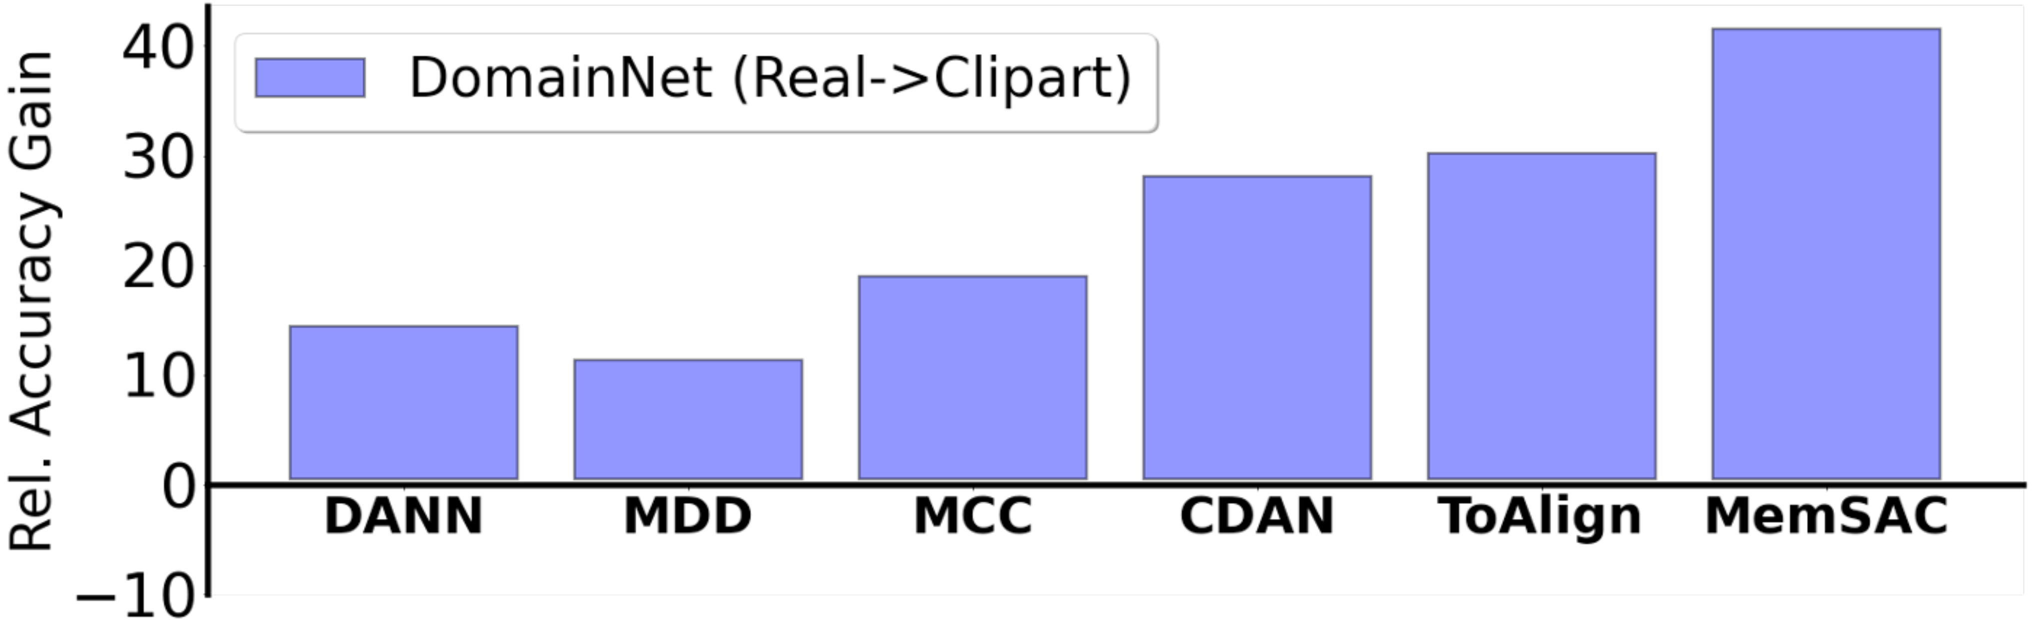

Benchmarking Unsupervised Adaptation on GeoNet

- UDA methods generally designed for covariate shifts, but they do not address geographical shifts.

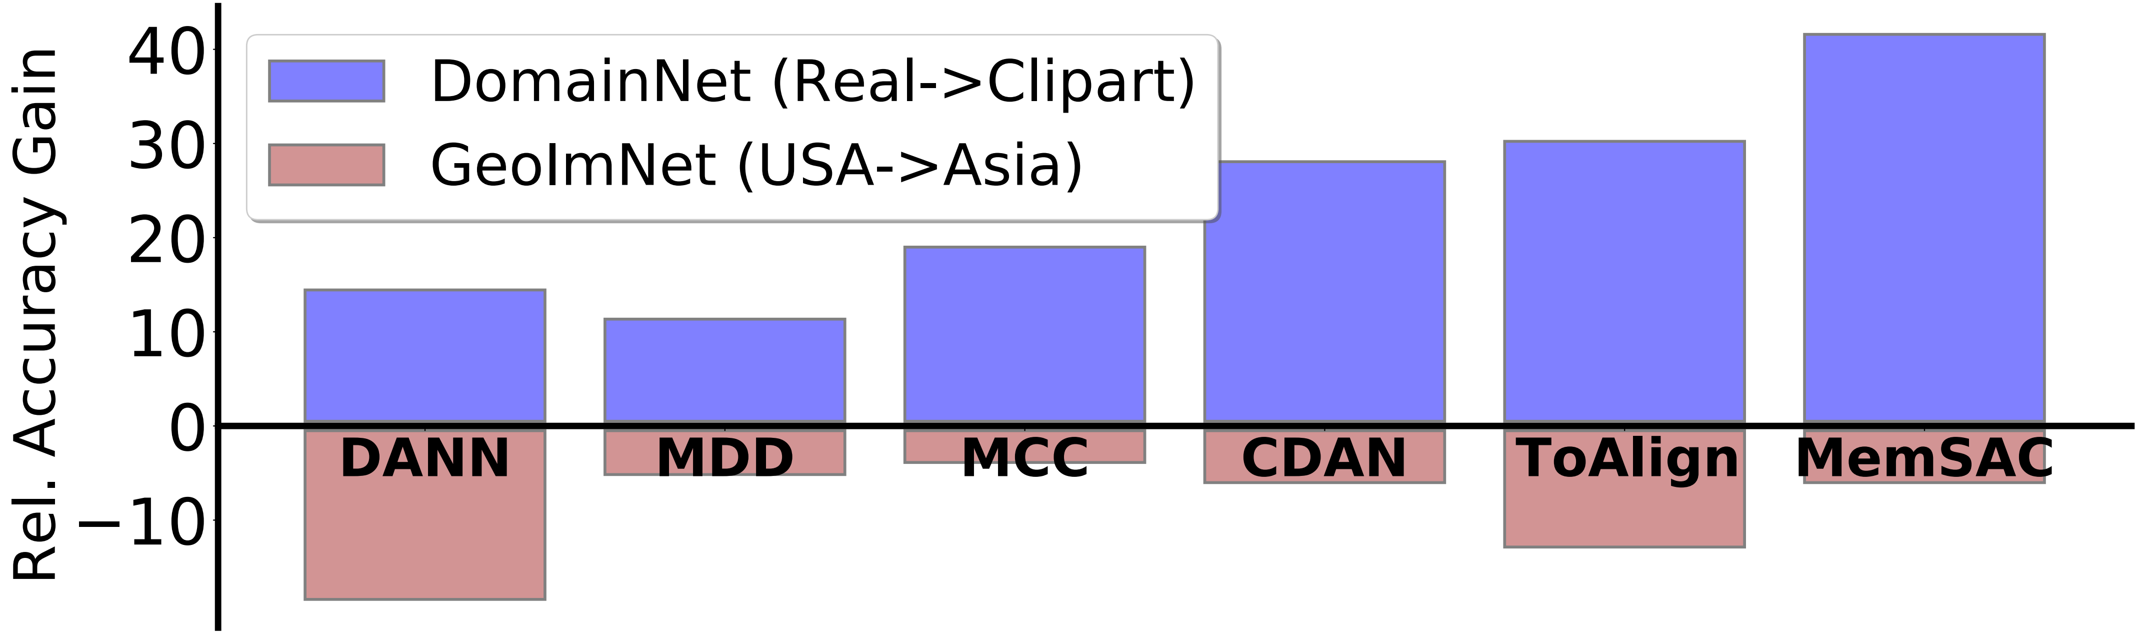

Large-Scale Pre-training Does Not Suffice for GeoDA

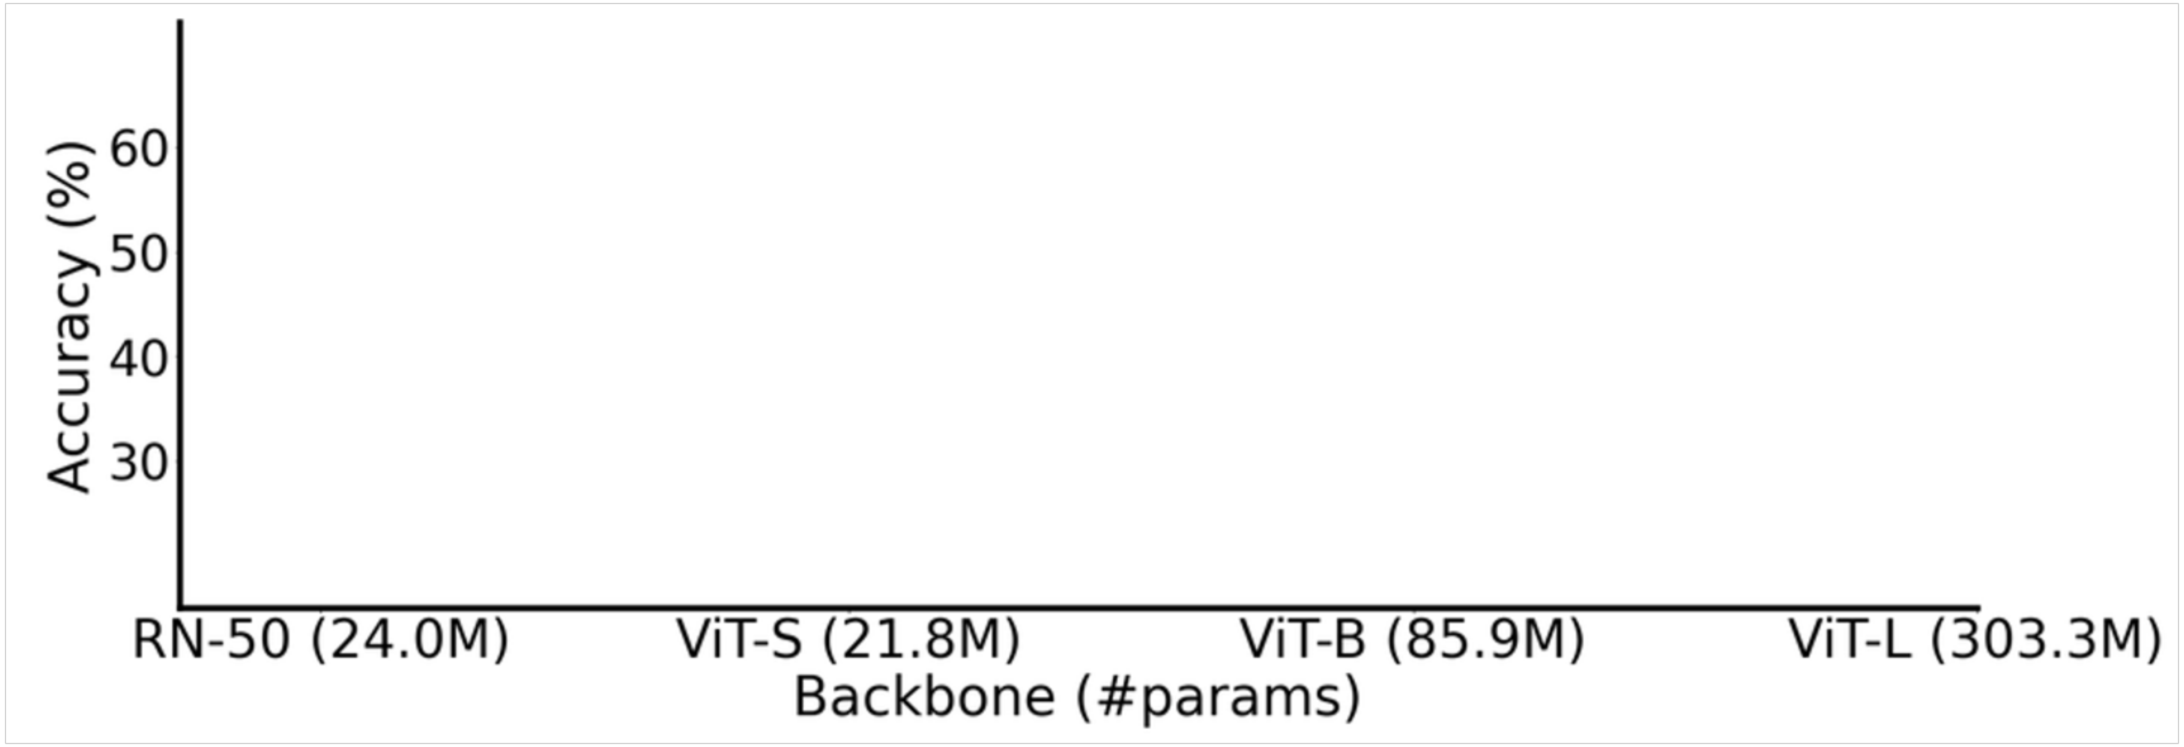

Large-Scale Pre-training Does Not Suffice for GeoDA

- Large-scale pre-training using large data does not suffice to bridge geographical disparities.

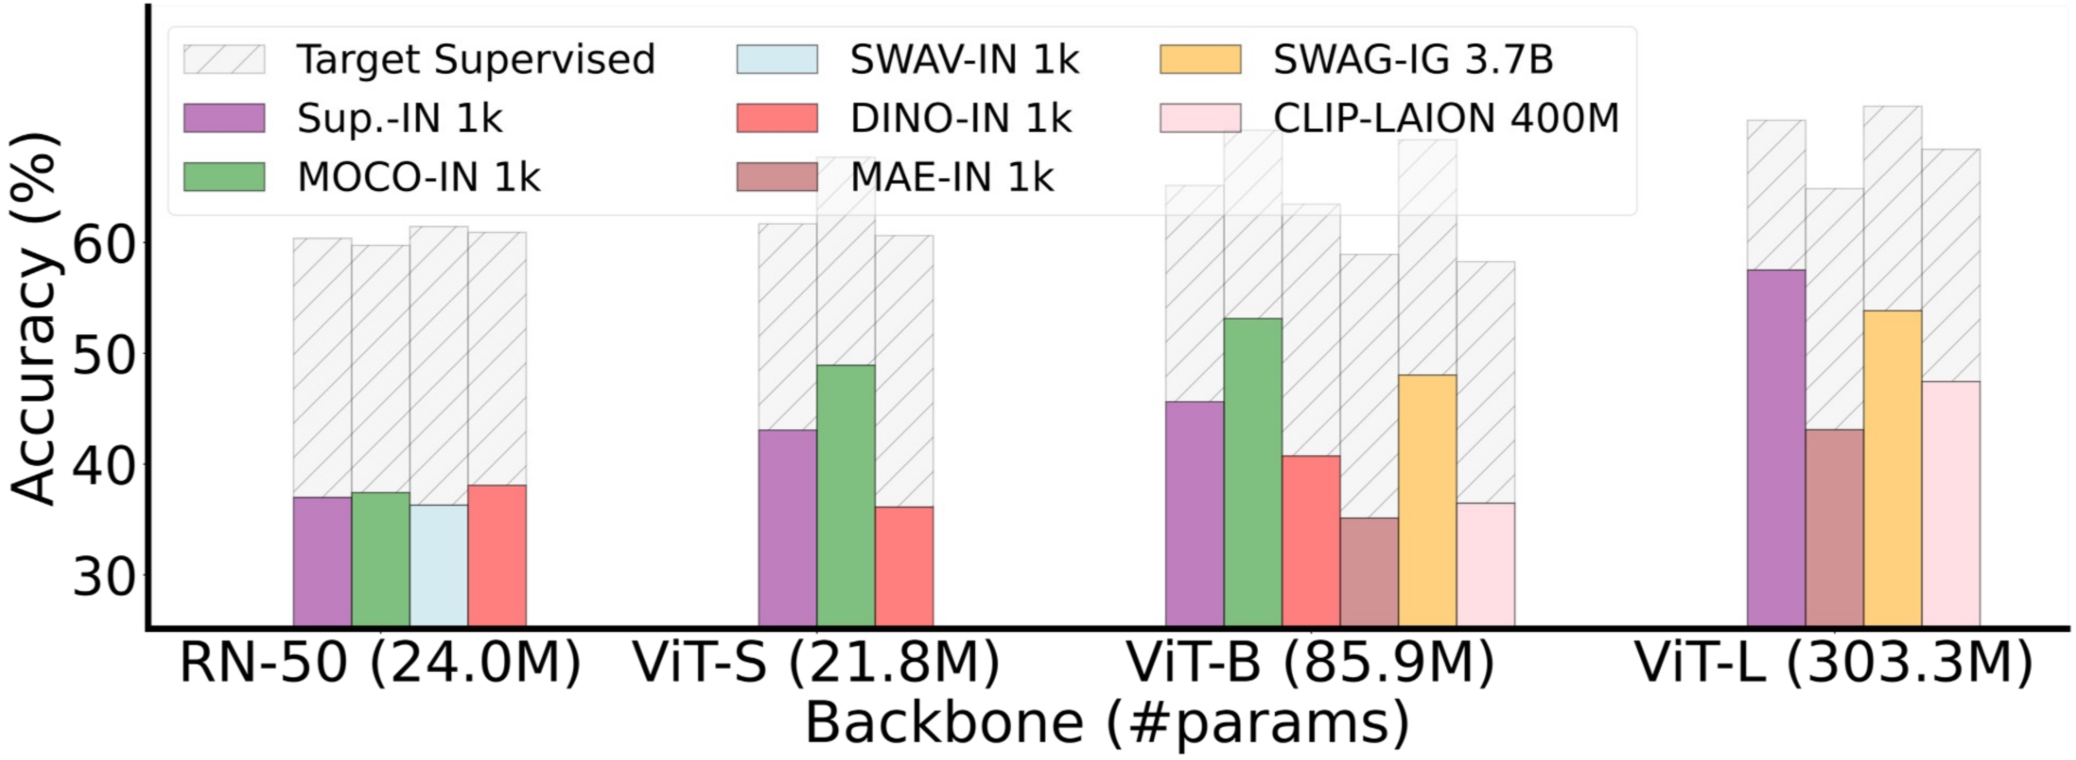

Summary

- GeoNet is a large-scale benchmark useful to study unsupervised adaptation across geographies.
- Existing domain adaptation methods are necessary, but not sufficient to bridge the novel shifts due to geography.
- Novel algorithmic solutions are needed to address the issues and deploy geographically robust models.

Poster #287

WED (06/21) PM

SCAN ME