Land Use Prediction using Electro-Optical to SAR Few-Shot Transfer Learning

Marcel Hussing

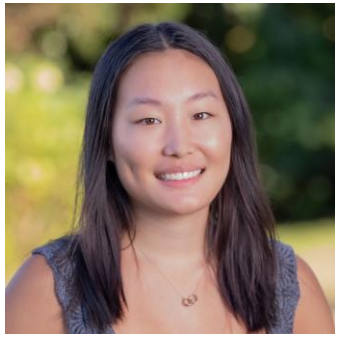

Karen Li

Eric Eaton

University of Pennsylvania

Contributions

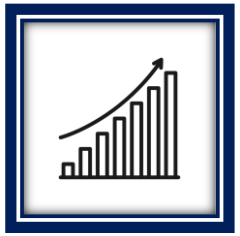

1. Sliced Wasserstein distance (SWD) embedding alignment can be scaled to multi-class setting

2. Instance normalization leads to more stable training and better performance

3. Contrastive learning improves transfer performance of SWD approach

Motivation

"Approximately 20% of the SDG indicators can be interpreted and measured either through the direct use of geospatial data itself or through integration with statistical data." – United Nations

- Land use and land cover mapping can measure the health of populations, urban areas, and ecosystems over time

- SAR data is unaffected by weather conditions and day-and-night cycle and can effectively collect information continuously, but labeled SAR datasets are limited and costly

Dataset and Base-Method

<u>So2Sat LCZ42 Dataset</u>:
Contains Sentinel 1 & 2 satellite imagery with

- Electro optical (EO) and
- Synthetic aperture (SAR) radar

images for land-use prediction with a total of 17 classes

- 4 low-rise
- 10 urban
- 7 natural

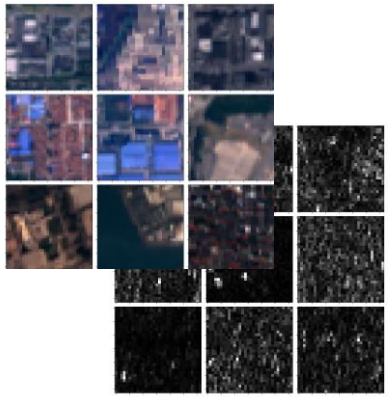

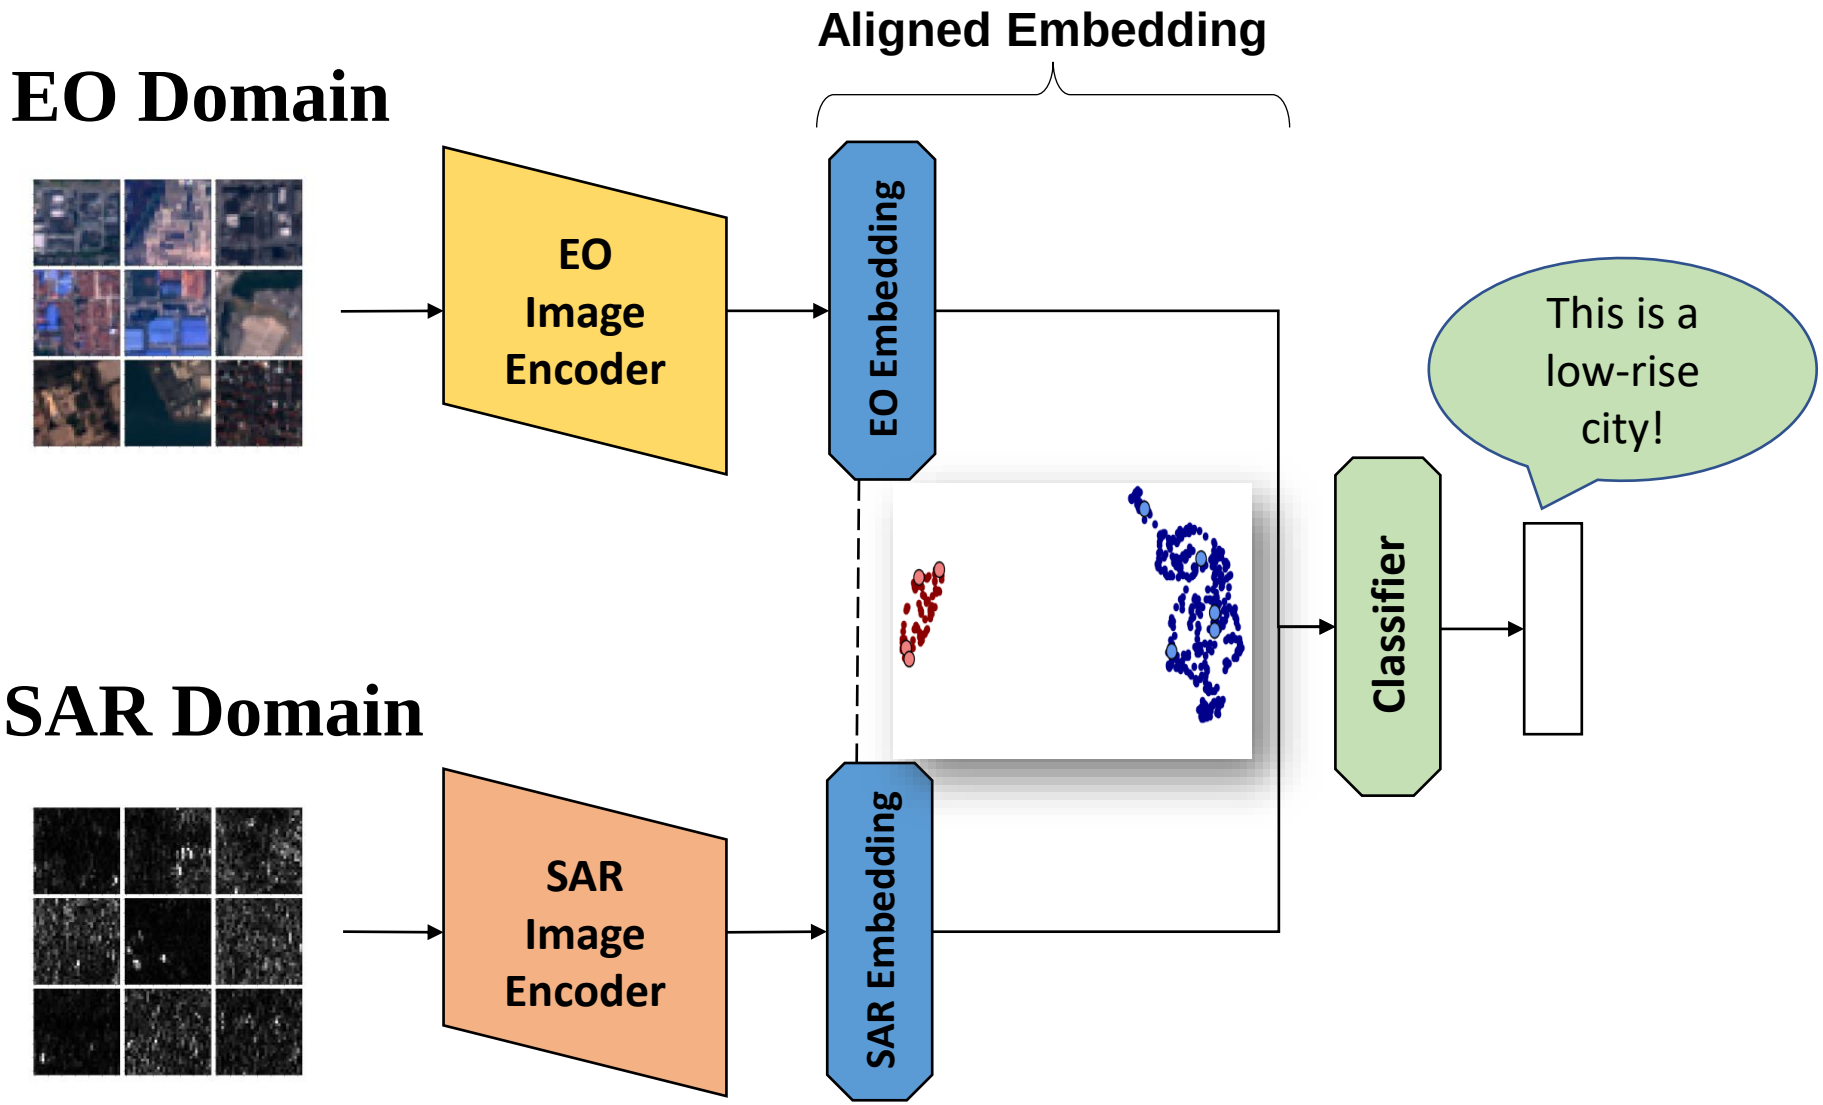

Method Extensions

Scaling to multi-class problems

- Consider more realistic and harder settings than previously analyzed.
- Consider 3 different difficulties of multi-task settings.

Instance normalization to stabilize training

- Normalization of inputs is very crucial, so is the normalization in embedding space; we find that the type of normalization matters!
- Instead of using batch normalization, we use instance normalization to better handle the wide range of spectrum in the SAR data.

Contrastive Learning for discriminative embedding

- Method assumes that learned EO embedding is sufficiently discriminative.
- We employ contrastive learning to ensure that this is in fact the case.

Experimental Results

Scaling to multi-class problems

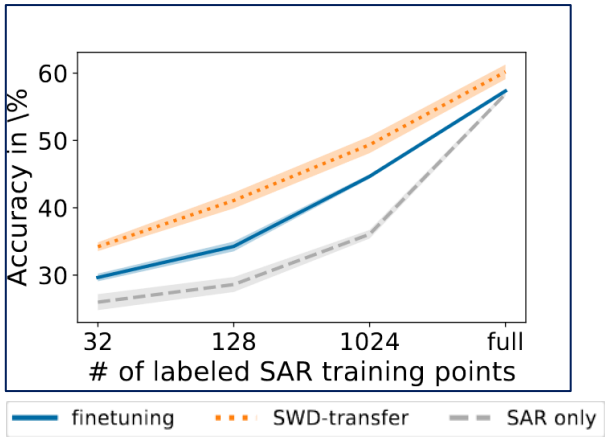

Method outperforms baselines on multi-class problems but leaves room for improvement in low data regime.

Instance normalization to stabilize training

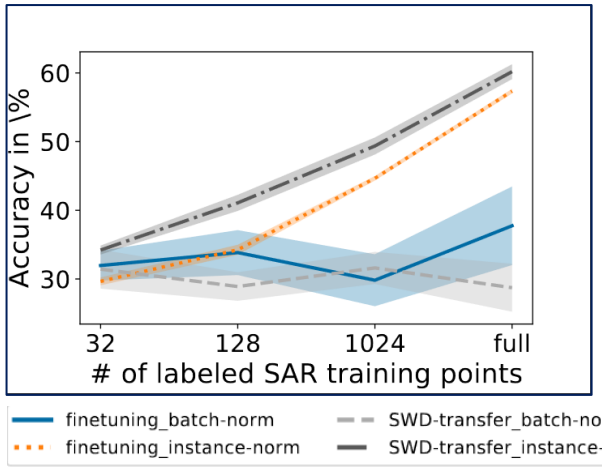

Instance normalization significantly stabilizes training and is *required* to achieve good results.

Contrastive Learning for discriminative embedding

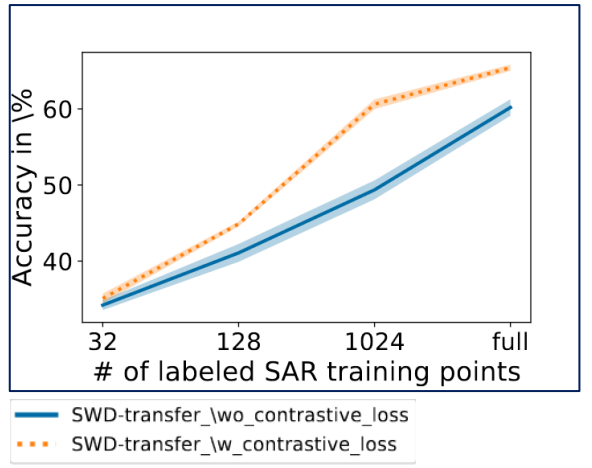

Contrastive learning improves overall transfer performance for most data regimes.

Thank you!

Marcel Hussing

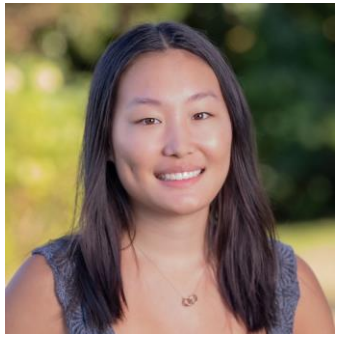

Karen Li

Eric Eaton

Contact: mhussing@seas.upenn.edu karentli@seas.upenn.edu

University of Pennsylvania