Computer Vision for International Border Legibility

Trevor Ortega, Thomas Nelson, Skyler Crane, Josh Myers-Dean, Scott Wehrwein

Western Washington University

Task: Border Legibility Estimation

Given an image of a border and the location of the border in the image, predict the legibility of the border. The goal is *not* to predict the location of the border, but rather how visible the border is to the naked eye.

High Legibility — Medium Legibility — Low Legibility

Tile Segment Division

Three regions are considered: each side of the border (A, C), and a margin around the border (B).

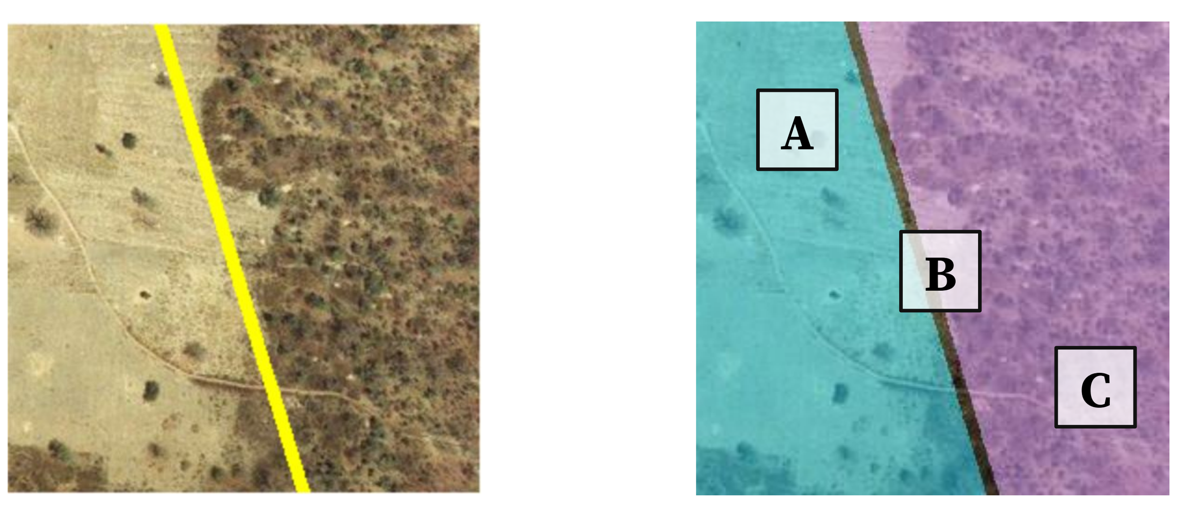

Data Collection

Overhead borders dataset:

- 612,347 aerial image tiles from Bing Maps
- Dense coverage of all global land borders

Validation set:

- 12,000 pairwise legibility judgments among 1000 random tiles

Code and data are available on our project page.

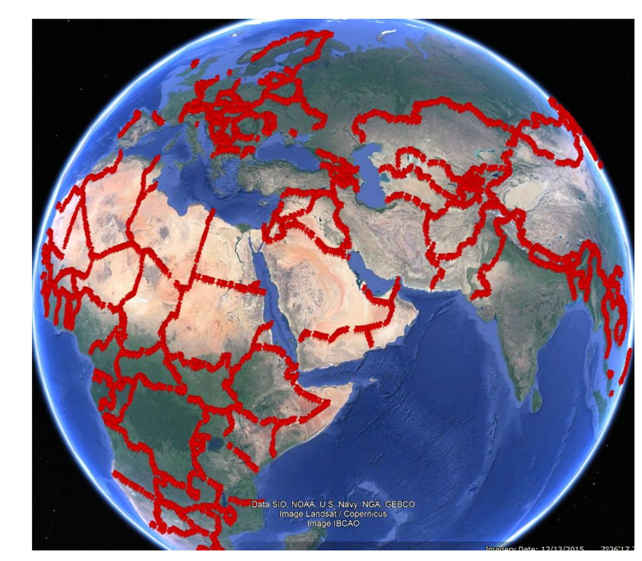

Dataset Sample Locations

BorderCut: A Contrastive Siamese Approach

- Generate synthetic pairs of border tiles, modified to be artificially legible.
- Train a Siamese model to predict which tile is more legible.
- CutMix-style augmentation increases legibility by replacing segments with pixels from another random tile.

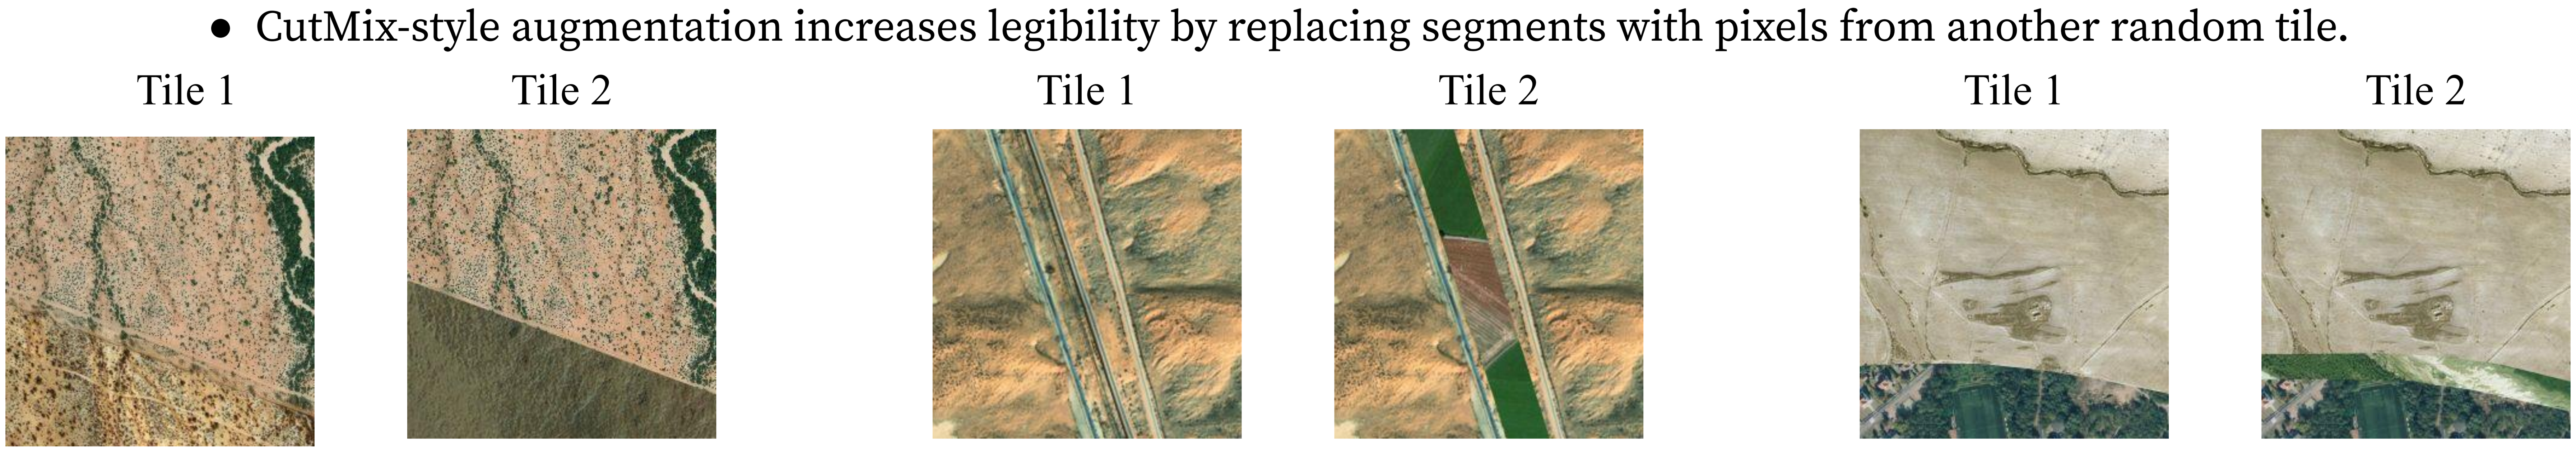

Original tile — Region C replaced — Original tile — Region B replaced — Region C replaced — Regions B and C replaced

Feature Comparison Baselines

- Pairwise differences between segment features.
- Coherence of feature cluster assignments among segments.

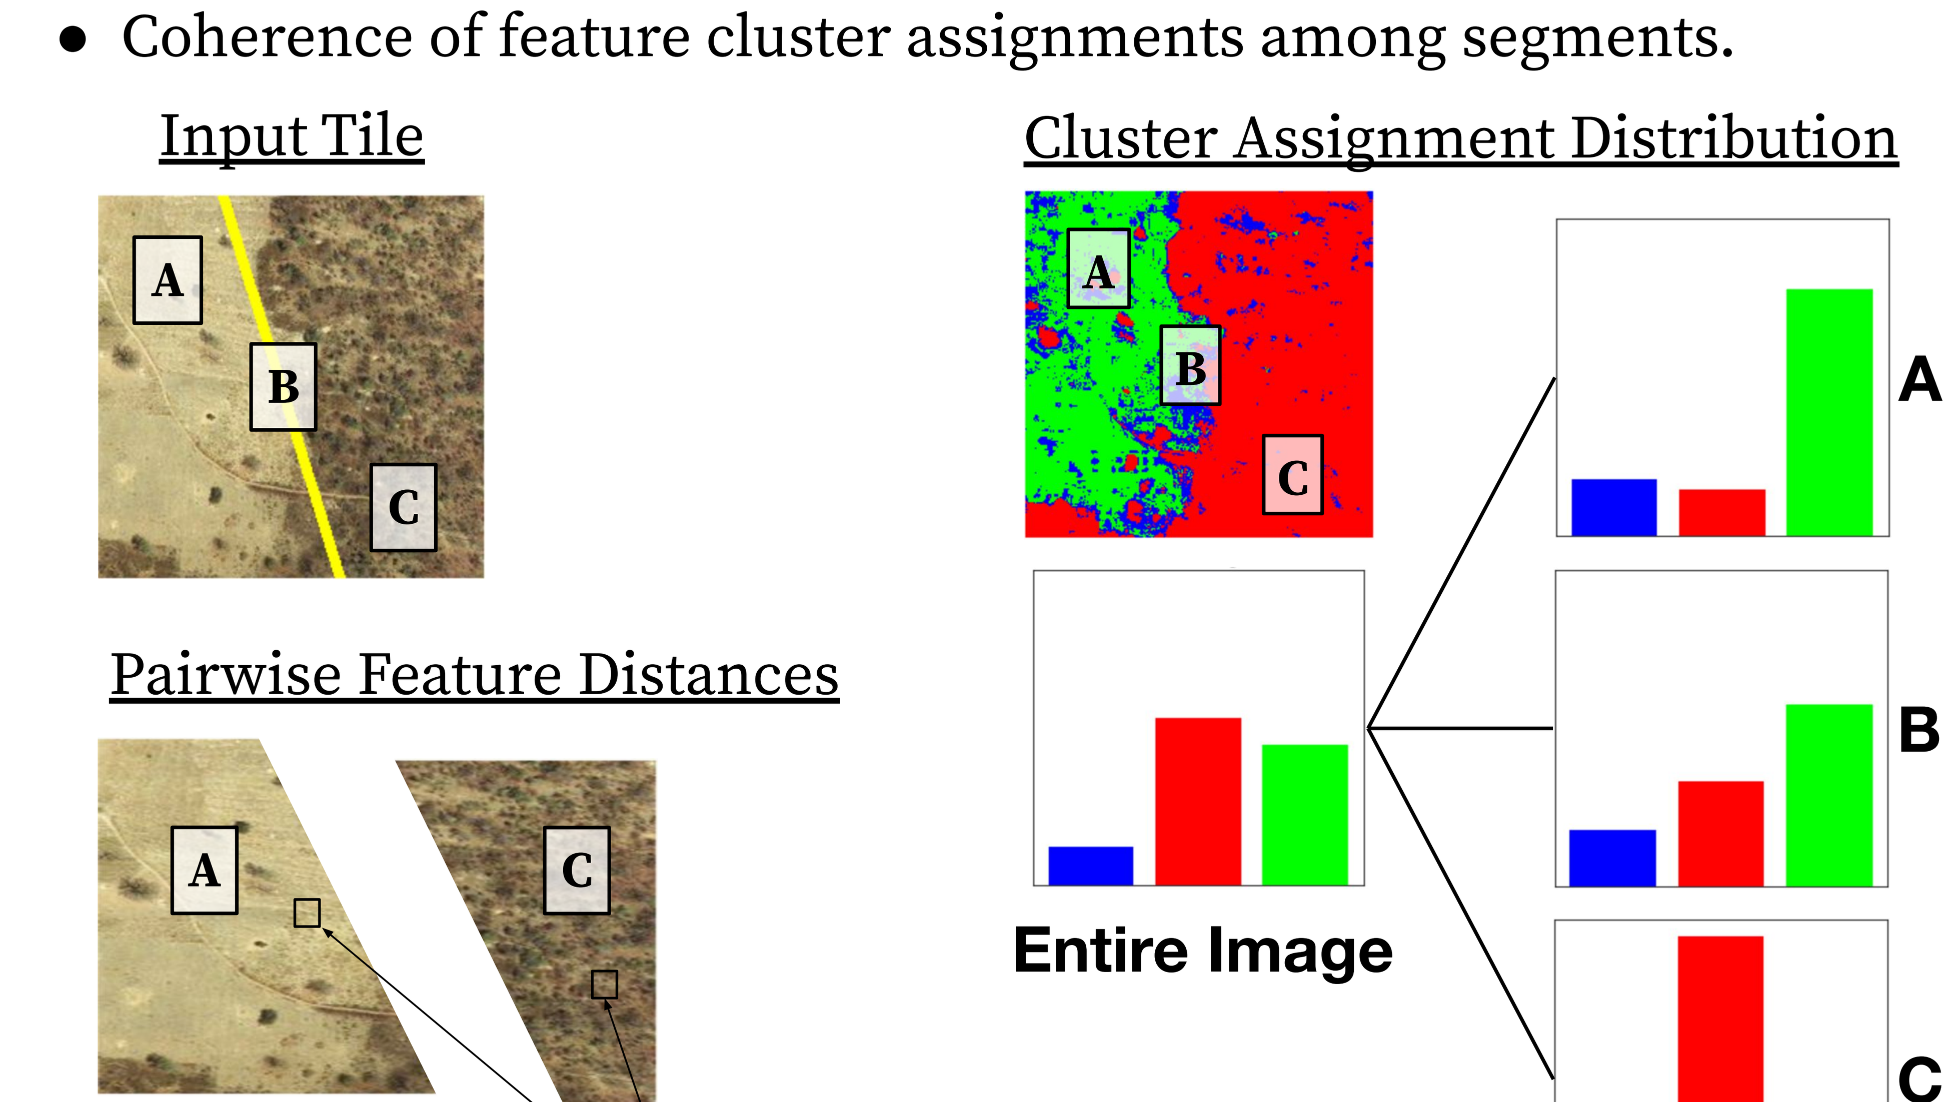

$$D(A, C) = \mathbb{E}_{a \in A, c \in C}[d(a, c)]$$

Results

Results from our best performing methods:

Method	Accuracy↑	τ↑	Footrule↓
BorderCut	**65.85 ± 1.6**	0.145 ± 0.02	283.18 ± 8.8
Clustering	58.80	**0.449**	**186.68**

Accuracy measures rate of agreement with pairwise human judgments.
τ and **Footrule** measure agreement between rankings derived from ground truth vs. predictions.

Results:

- BorderCut does best on raw accuracy of pairwise annotations.
- Clustering does best on global ranking metrics.

Conclusions:

- At 65% accuracy, the task is far from solved
- More data and better techniques are needed

Project page with code and data: https://fw.cs.wwu.edu/~wehrwes/BorderLegibility/