Mountains of Confusion

Evaluating Image Enhancement to Improve AI Landscape Classification

Larissa Bron, School of Environmental Studies, 2024

Mountain Legacy Project

University of Victoria | Environmental Studies

Project supervised by Dr. Eric Higgs

This research was supported by the Jamie Cassels Undergraduate Research Awards

Background

The Mountain Legacy Project (MLP) works with the largest systematic collection of mountain photographs in the world. More than 10,000 images have been repeated to quantify landscape change over time.

Mt. Baldy, Kananaskis River Valley

Historic, McArthur, 1889

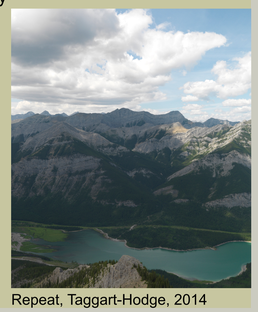

Repeat, Taggart-Hodge, 2014

Quantitative analyses require pixel-by-pixel mapping of ecosystem types in an image. This time-consuming task has considerably held back scaling up research efforts to entire landscapes.

Human created land cover masks

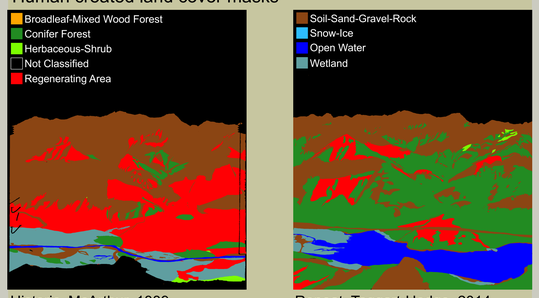

Historic, McArthur, 1889

Repeat, Taggart-Hodge, 2014

The Python Landscape Classification (PyLC) tool was developed to automate ecosystem classification using AI semantic segmentation. The MLP team is working to improve PyLC accuracy to better recognize the diversity of ecosystems in the Rockies.

PyLC generated land cover masks

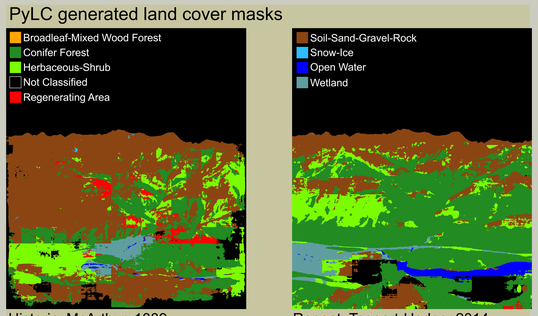

Historic, McArthur, 1889

Repeat, Taggart-Hodge, 2014

Hypothesis

Improving the appearance of ecosystems in images by applying image enhancements will clarify ecosystem characteristics, improving PyLC accuracy for classification.

Methods

1. Updated the original PyLC image set by removing erroneous land cover masks and improving image diversity by adding images from more locations.

Station 3, Wilmore Wilderness Park

Nidd, 1953

Station 443, David Thompson Country

Stephenson, 2011

Forget-Me-Not-Ridge, Kananaskis

Wheeler, 1897

2. Created nine altered versions of the image set using Python functions for three versions of contrast, denoising, and sharpening image enhancements.

Example Python code for a contrast image enhancement

```python
# Apply CLAHE with different combinations and show the results
for i, (clip_limit, tile_grid_size) in enumerate(combinations, start=1):
    # Apply CLAHE with the current combination
    clahe = cv2.createCLAHE(clipLimit=clip_limit, tileGridSize=tile_grid_size)
    enhanced_image = clahe.apply(image)

    # Display the enhanced image with the current parameters
    axs[0, i].imshow(enhanced_image, cmap='gray')
    axs[0, i].set_title(f'Clip Limit: {clip_limit},\nTile Grid Size: {tile_grid_size}')
    axs[0, i].axis('off')
```

3. Ran the baseline and nine sets of enhanced images separately through PyLC grayscale model 2-3, resulting in land cover masks and accuracy metrics read outs.

Accuracy metrics comparing PyLC output with reference masks

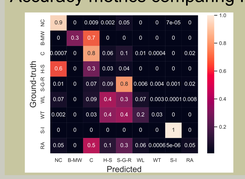

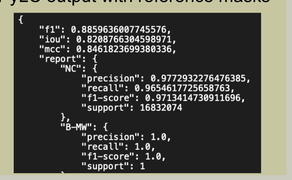

4. Compared the median value of accuracy metrics at two nested levels – F1 scores for land cover categories within each image, and three versions of averaged accuracy metrics.

Further Information

Detailed methods, code, data, images, and references are available on GitHub.

PyLC Responds to Image Enhancements

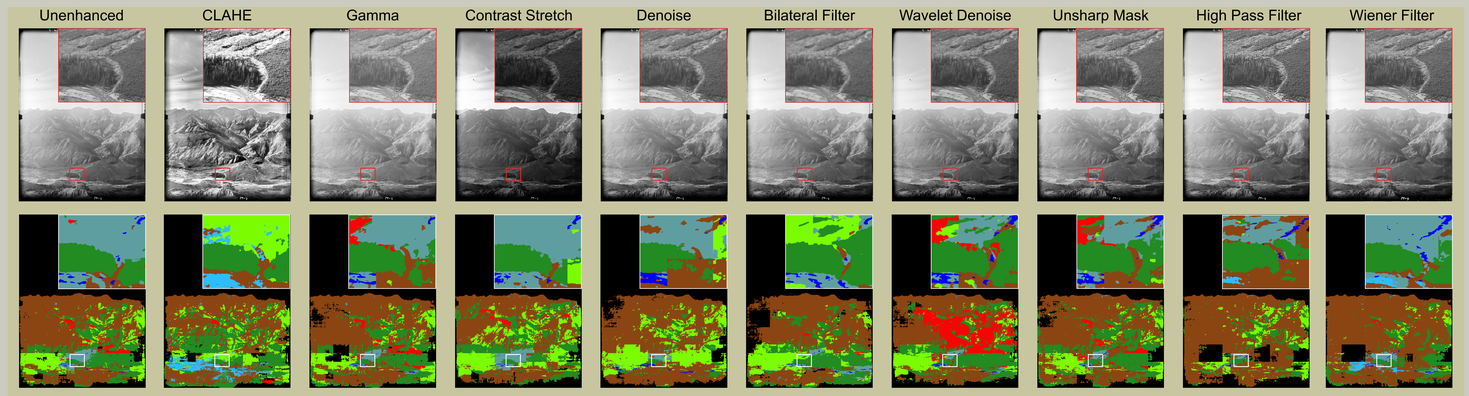

Each successive enhancement changed the resulting PyLC mask when compared with the unenhanced baseline.

Comparing Accuracy Between Image Enhancements

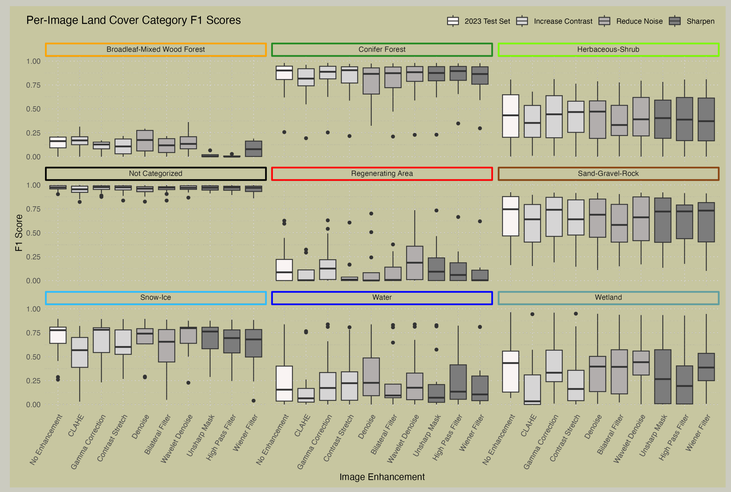

No image enhancement consistently improved the F1 score for any land cover category, instead the result was typically no difference or a decrease in F1. The goal is for all categories to achieve F1 scores comparable to Not Categorized and Conifer Forest.

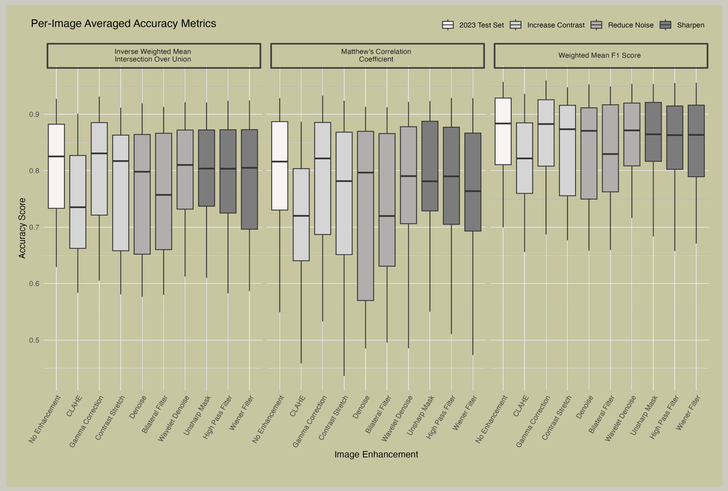

The averaged accuracy metrics each use different weighting to quantify PyLC accuracy for all pixels. Despite this difference, each show a similar trend that image enhancements either don't affect or decrease the accuracy for land cover classification.

Discussion

Image enhancements tended to decrease accuracy, though each changed the resulting PyLC mask. This lends support toward limited model generalization and a framework for data augmentation.

Changing the way images appear may hinder PyLC which was trained with only 95 images [1]. Efforts to improve accuracy may benefit more from increasing the number of training images [2].

Intervening in an AI pipeline using image enhancements may be better suited to data augmentation [3], increasing model generalizability while decreasing the number of manual training masks needed.

Acknowledgements

I am grateful for the opportunity to carry out research on WSANEC and Lekwungen lands. Thank you to my supervisor, Dr. Eric Higgs, and the MLP lab team for the support, as well as the incredible librarians, researchers, and students at UVic for great conversations.