ERDC/GRL TR-22-3

Geospatial Research Laboratory

US Army Corps of Engineers®
Engineer Research and Development Center

ERDC
ENGINEER RESEARCH & DEVELOPMENT CENTER

Snow-Covered Region Improvements to a Support Vector Machine-Based Semi-Automated Land Cover Mapping Decision Support Tool

Francis O'Neill, Kristofer Lasko, Elena Sava

October 2022

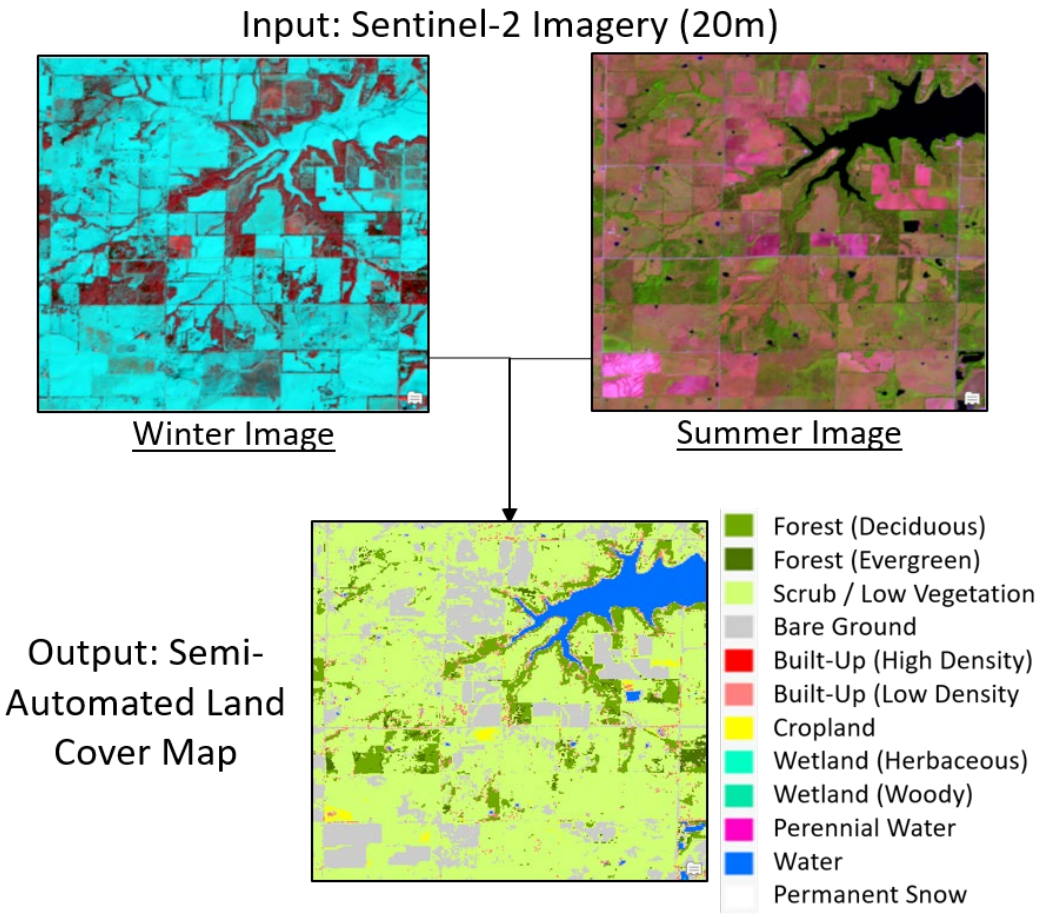

The U.S. Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdc.usace.army.mil.

To search for other technical reports published by ERDC, visit the ERDC online library at https://erdc-library.erdc.dren.mil.

ERDC/GRL TR-22-3
October 2022

Snow-Covered Region Improvements to a Support Vector Machine-Based Semi-Automated Land Cover Mapping Decision Support Tool

Francis O'Neill, Kristofer Lasko, Elena Sava

Geospatial Research Laboratory
U.S. Army Engineer Research and Development Center
7701 Telegraph Road
Alexandria, VA 22315-3864

Final Report

Prepared for Headquarters, U.S. Army Corps of Engineers
Washington, DC 20314-1000

Under PE 633463 , Project AU1 // Tactical Geospatial Information Capabilities / Enhanced Terrain Processing—Demonstration

Abstract

This work builds on the original semi-automated land cover mapping algorithm and quantifies improvements to class accuracy, analyzes the results, and conducts a more in-depth accuracy assessment in conjunction with test sites and the National Land Cover Database (NLCD). This algorithm uses support vector machines trained on data collected across the continental United States to generate a pre-trained model for inclusion into a decision support tool within ArcGIS Pro. Version 2 includes an additional snow cover class and accounts for snow cover effects within the other land cover classes. Overall accuracy across the continental United States for Version 2 is 75% on snow-covered pixels and 69% on snow-free pixels, versus 16% and 66% for Version 1. However, combining the "crop" and "low vegetation" classes improves these values to 86% for snow and 83% for snow-free, compared to 19% and 83% for Version 1. This merging is justified by their spectral similarity, the difference between crop and low vegetation falling closer to land use than land cover. The Version 2 tool is built into a Python-based ArcGIS toolbox, allowing users to leverage the pre-trained model—along with image splitting and parallel processing techniques—for their land cover type map generation needs.

Contents

Figures and Tables

Figures

Tables

Preface

This study was conducted for the US Army Corps of Engineers under PE 633463 , Project AU1, "Tactical Geospatial Information Capabilities / Enhanced Terrain Processing—Demonstration."

The work was performed under the Data and Signature Analysis Branch of the TIG Research Division, U.S. Army Engineer Research and Development Center, Geospatial Research Laboratory (ERDC-GRL). At the time of publication, Ms. Jennifer Smith was Chief, Data Signature and Analysis Branch; Mr. Jeffrey Murphy was Chief, TIG Research Division; and Mr. Austin Davis was Technical Director of ERDC-GRL. The Director of ERDC-GRL was Mr. David R. Hibner.

COL Christian Patterson was Commander of ERDC, and Dr. David W. Pittman was the Director.

Acronyms and Abbreviations

CONUS	Continental United States
ETP	enhanced terrain processing
GB	gigabyte
GRL	Geospatial Research Laboratory
OCONUS	outside continental United States
L2A	Level-2A (surface reflectance)
NIR	Near-infrared
NLCD	National Land Cover Database
RAM	random access memory
ROI	region of interest
SWIR	shortwave infrared
SVM	support vector machine
S2	Sentinel-2
USGS	United States Geological Survey
VisNav	Visual Navigation Dataset

This page intentionally left blank.

1 Introduction

The Semi-Automated Land Cover Tool presented in this report enables the rapid creation of up-to-date land cover maps over any continental United States (CONUS) area of interest (AOI). The tool takes as input two Sentinel-2 (S2) satellite images, one from the winter and one from the summer, and outputs a segmented map that classifies each image pixel as belonging to one of 12 land cover classes. This report focuses on the enhancements made between the original Version 1 (V1, previously described in Lasko and Sava 2021) and a new Version 2 (V2). V2 expands the tool's applicability to AOIs where snow is present.

The ability of the land cover tool to translate regularly updated satellite imagery into useful geospatial data is critical. Land cover is an important variable for a variety of applications, including cross-country mobility and cover and concealment. However, too often soldiers must rely on existing land cover maps with coarse resolution (30 m pixels or larger) and/or high latency (multiple years out of date), such as the GeoCover and VisNav datasets. The Semi-Automated Land Cover Tool addresses both of these issues through its use of S2 data, which offer both higher resolution (20 m pixels) and lower latency (S2 has a revisit time of 10 days).

V1 achieved high accuracy across a wide variety of CONUS land cover classes. However, the classification algorithm suffered performance degradation in AOIs where snow was present in one or both satellite images. The spectral signature at each input sample point is drawn from a pair of two satellite images, one from winter and one from summer; however, the initial training dataset did not include any image pairs in which one or both of the images contained snow cover. This omission poses a serious drawback and decrease in accuracy for users applying the tool to their own imagery in which snow cover occurs in the winter, summer, or both dates. The spectral signature of any land cover type is substantially different when one of the images is covered in snow (e.g., winter image with snow, summer image with low vegetation, compared with low vegetation in both dates). Additional training data collection was therefore required to address this issue and ensure good performance over the entire extent of CONUS, regardless of snow cover conditions. The collection and integration of these additional training data forms the focus of the V2 improvements described in this report.

Of the major land cover types, snow cover is critical and sometimes unaccounted for. Snow cover is important as it plays a major role in surface radiation (e.g., high albedo), energy budgets, and hydrology (Kongoli, Romanov, and Ferraro 2012). Snow and glacier melt play a major role in water supply for many mountainous regions of the world, especially during summer and spring when precipitation may be low and snowmelt is the primary water source (Van Loon 2015). Decrease in snow levels can result in drought in areas that rely heavily on snowmelt, which infiltrates below the root zone of plants more effectively than rainwater and is often a major contributor to groundwater recharge (Biemans et al. 2019; Dierauer et al. 2019). Snow can also be a critical land cover class for terrain understanding in military applications. Snowfall can reduce effectiveness and stability of communications, data efforts, and mobility of personnel or vehicles and may affect concealment. Thus, it is critical to improve the performance of the Semi-Automated Land Cover Tool over AOIs that include snow cover.

1.1 Background

Land cover represents the physical land type, which can include vegetation, water, or manmade features found on the Earth's surface. The term *land cover* is often incorrectly used interchangeably with the term *land use*. However, land use refers to the purpose of the land cover—for example, agriculture, which can include multiple different land cover types such as trees (e.g. plantation), row crops, grassland (e.g. ranching agriculture), or water (aquaculture). These distinctions are important for a holistic understanding of a particular region; however, land use mapping is much more challenging as it can be difficult to identify land use directly via satellite remote sensing. Land cover provides the more fundamental characteristics useful for terrain analysis, as well as navigation through the environment.

Satellite remote sensing enables recurrent, near-global observations of Earth's surface at a range of scales from entire continents to small areas and spatial pixel resolution of kilometers to centimeters. Data from public satellites such as Landsat 8 or S2 provide consistent and high image acquisition angles, high geolocation accuracy, geometric correction, robust atmospheric correction, and surface reflectance products. Because of these qualities, the above sensors are optimal for the creation of land cover products. S2 with its constellation of two satellites provides near-global observations of the earth about every 6 days. The resulting observations of

Earth's surface are an improvement on Landsat as S2 offers additional spectral bands such as red-edge, as well as a higher spatial resolution (10 m for red, green, blue, and near-IR bands, or 20 m for most of the remaining bands) compared to 30 m for Landsat 8 (European Space Agency 2012).

Many different land cover products have been created using S2 and Landsat 8. One of the most popular and robust products is the National Land Cover Database (NLCD), which uses a dense annual time series stack of Landsat 8 imagery to map multiple land cover classes across the United States (Jin et al. 2019). The NLCD is produced at intermittent time periods; it is currently available for the years 2001, 2003, 2006, 2008, 2011, 2013, 2016, and 2019. The most recent NLCD has an overall accuracy of 90.6% at the 30 m scale (Wickham et al. 2021). Despite this high apparent accuracy, Lasko and Sava (2021) find that in some cases, the 2016 NLCD identifies land use rather than land cover. For example, in one instance the trees and grass of a golf course were identified as "Developed" by the 2016 NLCD due to the course's location within an urban area. Other regional and global land cover products also exist, each with their own strengths and limitations. One example is the Corine Land Cover Database, which, although it boasts high fidelity and accuracy, is limited in geographic scope to Europe and has a resolution of 100 m (Büttner 2014). Another example is the North American Land Change Monitoring System, which is available at higher resolution (30 m) but has a low update frequency (3–5 years) (Latifovic et al. 2017). Other land cover products commonly used by soldiers, including the NGA Visual Navigation dataset and the GeoCover dataset, have similar issues (National Geospatial Intelligence Agency 2010; Cunningham et al. 2002). These drawbacks limit the usefulness of the above products since they may omit key terrain details, mislabel image pixels, or fail to capture land cover changes that have occurred since they were last updated.

Two contrasting methods are commonly used for the creation of thematic maps from remotely sensed imagery: supervised and unsupervised classification. Unsupervised classification generates clusters of spectrally similar pixels. This means that pixels must be assigned to the clusters by human interpreters after the classification takes place. In contrast, supervised classification uses human-generated training data to define the land cover types of interest that are identified in the output thematic map. Supervised classification is more time-consuming due to the necessity of

generating high-quality training data in order to get meaningful results but often produces more useful output and is the dominant method for land cover mapping today (Mohd et al. 2009).

Standard workflow for supervised classification involves the manual labeling of a large number of pixels within the region of interest (ROI), stratified throughout the various target land cover classes (Gómez et al. 2016). These pixel labels and the corresponding spectral profiles of the satellite imagery are used to train a classification model that can then produce a thematic map of the entire ROI. This classification model can consist of any one of a variety of supervised classification algorithms that are commonly used in land cover mapping studies. These algorithms include support vector machines (SVM), ensembles of weak learner decision tree methods such as random forest, various types of artificial neural networks, probabilistic classifiers such as Naïve Bayes, and many more linear and non-linear models (Pedregos et al. 2011).

1.2 Approach

Many of the global and regional land cover products mentioned above rely on a dense time series stack of imagery (often 12-plus images per location), which can be very data intensive and difficult for users to replicate on their own computers. Moreover, these land cover products are typically created for a single year and thus quickly become out of date, especially in areas experiencing land cover change on an annual basis. (The most recent NLCD map was produced in 2019, released in 2020, and is already 3 years out of date at the time of writing.) To address this issue, Lasko and Sava (2021) previously created an SVM-based land cover classification algorithm that uses only two dates of imagery and integrated it into an ArcGIS Pro decision support tool for end users to create a land cover product over a specified region, thereby resolving the issue of up-to-date maps. However, this study expands on that effort by adding an additional land cover class—permanent snow (snow in winter, snow in summer)—and by accounting for seasonal snow cover in the other land cover classes (e.g., snow in winter, bare ground in summer). Lastly, this pre-trained machine learning classifier approach was validated across several CONUS test sites using stratified random sampling points as described in the methodology section.

The pre-trained machine learning model is created by using expertly labled training data on imagery that is representative of the different land

covers found in CONUS. The training data is then fed into an SVM classifier to fit it and it is ultimately packaged with the decision support tool. The SVM classifier is a widely used non-parametric statistical learning classifier with no assumptions made regarding the underlying data distribution. SVMs generally perform well on sparse data, and they are also less likely to overfit the model than a decision tree classifier such as a random forest. The SVM is versatile in that different kernel functions, such as linear, sigmoid, or polynomial, can be used for model training. While the SVM was originally designed for binary classification problems, it has been shown to be effective in multi-label classification problems with high-dimensional feature spaces. This method typically performs effectively in land cover classification studies using shallow learning techniques (Foody and Mathur 2004; Pal and Mather 2005; Jia et al. 2012; Pal and Foody 2012). The SVM algorithm promises to obtain the optimal hyperplane for a training dataset in terms of generalization error. A detailed description of the SVM algorithm can be found in Suykens et al. (1999). Additional description of our classification method is in the methodology section.

1.3 Scope

For a semi-automated land cover classification tool to produce reliable output across CONUS, it must have the capability to recognize and categorize all the potential spectral inputs within the region of intended use. These various inputs include not only the different types of land cover to be identified but also the possible temporary weather conditions over those land cover types. In particular, when labeling image pixels based on two dates of satellite imagery, it is essential that the classifier be able to recognize each land cover class even when the surface spectral signature of the pixel in question is covered by snow in one of the input images.

Accordingly, this report focuses on building a robust semi-automated land cover classification tool with a geographic focus on CONUS and the inclusion of a snow cover class. The report leverages moderate resolution multispectral imagery from Sentinel-2 at 20 m for the analysis. The study limitations include the typical limitations present when creating pre-trained models. The study is designed to create land cover maps over broad areas quickly and accurately for use in general land cover and land use monitoring applications. One of the main limitations is that the pre-trained model may not perform well if it is given imagery that is significantly different from the training data it was fit on. Lastly, the study

uses 20 m imagery, which is of moderate resolution; however, this may be too coarse of a resolution for some applications such as feature extraction, battle damage assessment, or for fine-scale details such as mapping individual buildings, roads, and so on.

1.4 Objective

Our objective is to expand the functionality of the V1 Semi-Automated Land Cover Classification tool to incorporate the snow cover class, while simultaneously maintaining the accuracy in already-existing classes and expanding the tool's capability to seasonality changes. This upgraded V2 classification model is assessed for accuracy and integrated into an ArcGIS Pro Python toolbox. This packaging enables end users to automatically generate their own land cover maps without the need for their own training data in the same way as its predecessor.

2 Data and Sample Sites

2.1 Sentinel-2 imagery

The ability to have continuous satellite data available over a large spatial extent requires a constellation of satellites that have long-term data records, as well as new satellite constellations that collect data at frequent intervals. For these reasons, as well as global coverage and free data access, S2 was selected as the satellite imagery source for this study.

Currently, two operational S2 satellites are available, with multispectral instruments that acquire data at a moderate temporal and spatial resolution. S2 collects data at a 10 m, 20 m, and 60 m spatial resolution depending on the spectral band. It has a temporal resolution of approximately 5 days, with two satellites in constellation. S2 imagery can be acquired from a variety of US sources or directly from the European Space Agency's public-facing servers available at the open access hub website. Each scene file contains 13 multispectral spectral bands ranging from the visible, near infrared (NIR), and shortwave infrared (SWIR) electromagnetic frequency domain. S2 land surface reflectance scenes at Level 2A (L2A) were downloaded and used for the land cover training data collection and model development. Table 1 summarizes the spatial and spectral characteristics of S2 L2A datasets used to develop the land cover classification tool in this study.

Table 1. Multispectral characteristics of S2 (L2A Surface Reflectance) satellite spectral bands used in this study.

Band—Spectral Region	Central Wavelength (nm)	Resolution (m)
Band 2—Blue	490	20
Band 3—Green	560	20
Band 4—Red	665	20
Band 5—Red Edge	705	20
Band 6—Red Edge	740	20
Band 7—Red Edge	783	20
Band 8A—NIR	865	20
Band 11—SWIR	1,610	20
Band 12—SWIR	2,190	20

2.2 Description of training and testing sites

For V1 of the land cover classification tool, multispectral S2 imagery was collected throughout different sites within CONUS to represent the diverse land cover types and ecosystems. Training locations are shown in Figure 1.

Figure 1. Map showing the locations of S2 data sites (19 total) used for building and testing the semi-automated land cover classification models.

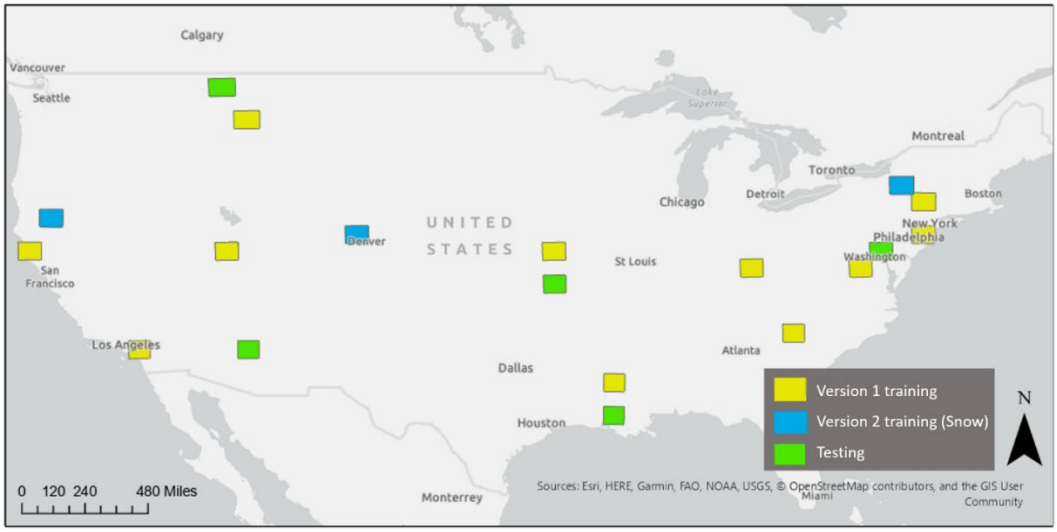

The northern California site is representative of lush evergreen forests found throughout the west coast and secondarily includes a mosaic of low vegetation, water, and small suburban areas. The southern California site includes the dry urban area of Los Angeles as well as arid landscape with less dense forests, bare areas, and croplands. The Utah site is characteristic of arid, rocky, salt-laden landscape but also includes significant areas of cropland and built-up areas. The Mississippi and Louisiana site is characteristic of humid sub-tropical forest, wetlands, irrigated croplands, lush mixed forest, and the Mississippi River Valley landscape. The New York site is representative of temperate continental climate landscape with a mosaic of low vegetation, deciduous and evergreen forest, lakes, and more. The New Jersey site is representative of warm continental climate with wetlands, built-up areas, forests, and more. The Virginia, Maryland, DC scene is representative of warm oceanic climate with a mosaic of densely built-up areas, suburbs, wetlands, rivers, croplands, deciduous forest, and more. The Kentucky and South Carolina scenes are representative of mosaic landscapes in cool continental and humid subtropical environments, respectively. The Kansas scene is representative of large-scale commercial cropland such as soybean and corn. The western Montana site is a mosaic of plains, forest, wetlands, and

croplands. Lastly, the Arizona site is characteristic of semi-arid higher elevation areas.

For V2, additional S2 data were collected at three different CONUS sites (also shown in Figure 1). These new sites were chosen to ensure coverage across a variety of snow-covered land cover classes, as well as the geographically variable spectral signatures within each of those land cover classes. The upstate New York site is characteristic of crop mosaics, the open water of the great lakes, river systems, and both urban and suburban built-up areas. The Colorado site includes large swathes of deciduous and evergreen forests both on the high plateau adjacent to Denver's urban cover and on the slopes of the Rocky Mountains, which also include substantial bare ground and permanent snow. The northern California scene, meanwhile, is representative of managed low-vegetation pasture punctuated by a single area of bare ground (Mt. Shasta), which includes a substantial area of permanent snow.

All of the training land cover classes included in the V1 tool were also represented in these additional training scenes, with three exceptions: the two wetland classes (woody and herbaceous) and "perennial water." The omission of the wetland classes is due to the fact that regions in which wetlands naturally occur are generally too warm to experience regular snowfall. The "perennial water" class, meanwhile, was omitted because the definition of the class makes it impossible to identify in locations where the winter satellite scene is covered in snow. Perennial water is defined as "a wetland, spring, stream, river, pond or lake that only exists for both dates of imagery following precipitation or snowmelt" (Lasko and Sava 2021). At a site where the winter scene is covered in snow, all water features only appear in the summer scene (being rendered invisible with ice and snow in the winter scene), making it impossible to identify which features are the result of snowmelt.

Finally, five additional two-date S2 image pairs (winter and summer) were collected from across CONUS for model validation. These test sites were chosen to cover the full range of target land cover classes, where possible, in both snow-covered and snow-free regions. The validation imagery locations were chosen as follows: Louisiana, for assessments of woody wetlands and other snow-free land cover types; Arizona, to test performance in arid conditions; Kansas, for snow-covered cropland and deciduous forest; Washington, DC, for its herbaceous bayside wetlands

and other snow-free cover classes; and Montana, for permanent snow classification, snow on bare ground (mountainsides), and evergreen forest.

3 Methodology

3.1 Methodology

The goal of this project is to expand the capability of the extant semi-automated land cover mapping tool described in Lasko and Sava (2021). Specifically, the goal is for improvement on imagery pixels containing snow cover in one or both dates of imagery (because V1 did not account for snow). At its core, this decision support tool is built around a pre-trained machine learning model that enables the user to create up-to-date land cover maps without the collection of additional training data or tedious manual labeling at the pixel level. The pre-trained model in V1 of the tool did not account for snow and thus performed poorly over snow-covered areas. To rectify this shortcoming, a new set of snow-covered pixels was collected as training data and combined with the original training data in order to train a new machine learning classifier capable of producing accurate results over both snowy and snow-free regions. The remainder of the training, evaluation, and ArcGIS Python toolbox integration workflow was replicated from V1 of the tool, as shown in Figure 2.

In this technical report, the focus is on improving the performance of the semi-automated land cover classification tool for regions in which one or both S2 images contain snow-covered pixels without sacrificing the existing accuracy of the V1 tool in regions without the snow cover class.

3.2 Training data

The first step of training a classifier is to collect satellite imagery and label pixels within each of the land cover classes of interest. This process is labor intensive and typically requires many hours of work by an analyst. Although automated methods of training data generation have been demonstrated, manual pixel labeling remains the norm and is the method used in this project (Malinowski et al. 2020).

Figure 2. Depiction of underlying methodology for land cover classification tool.

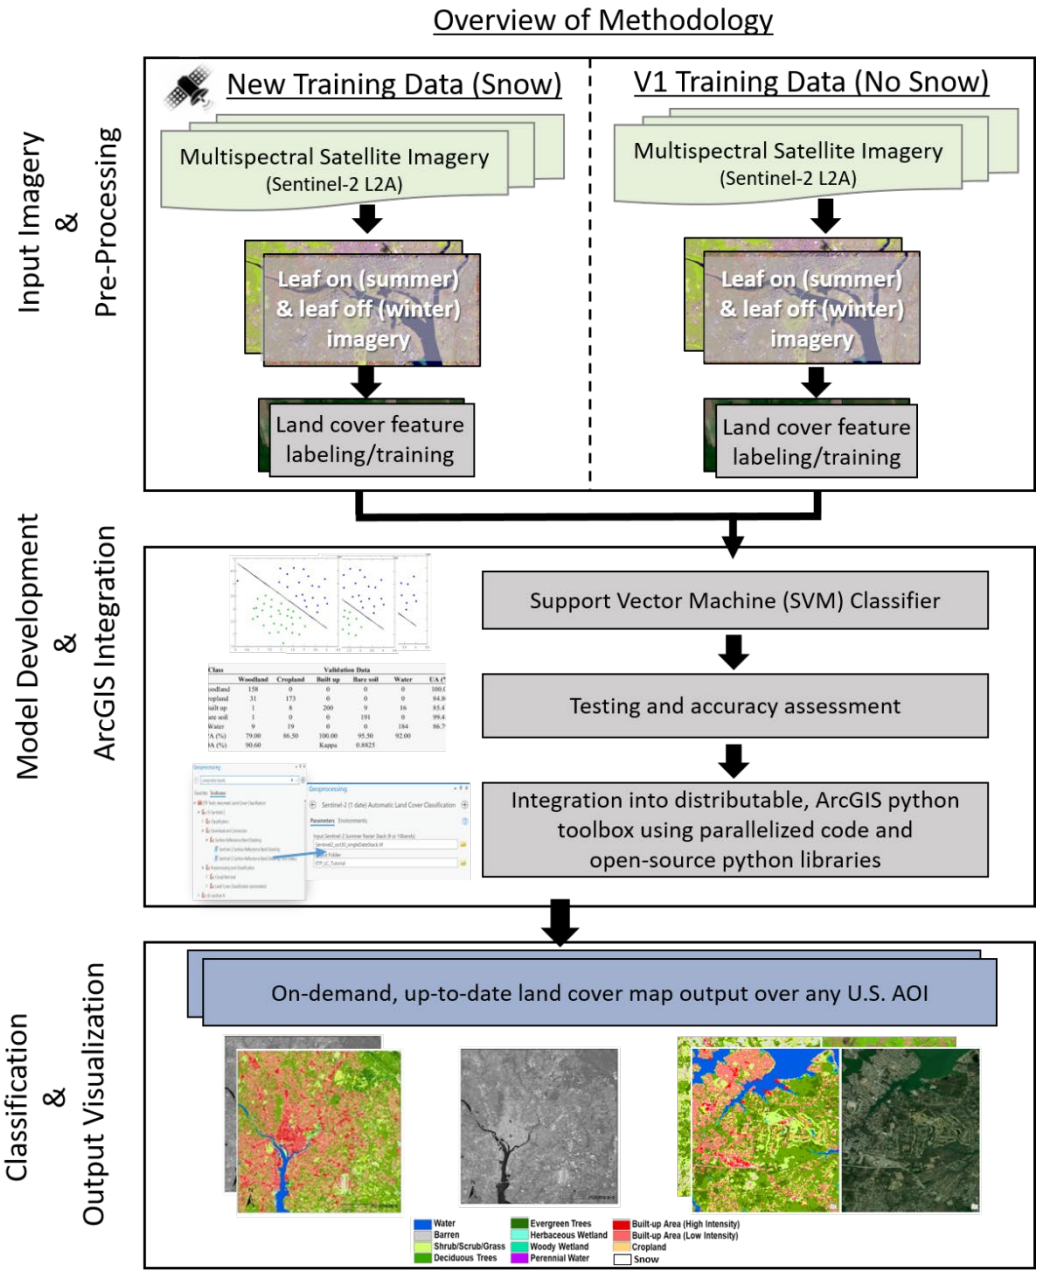

To make an informed decision about the land cover label, multiple information sources are used in addition to the input imagery. Supplementary data sources include higher resolution imagery, time series of key phenological dates, and Google Earth basemaps. However, it is crucial for the training labels to be generated directly on the S2 imagery to avoid any geolocation issues introduced from the complementary datasets (Congalton and Green 2019).

After thousands of training points have been collected, the pixel values for each spectral band along with the associated land cover class labels (e.g., evergreen trees) are combined into a single tabular file, verified, and passed into a machine learning classifier such as an SVM. Often, it is necessary to collect additional training data as initial classifier training may reveal errors.

The core methodological improvement between the V1 and V2 classifiers is the incorporation of additional training data representing snow-covered pixels. Training data are a vital part of any machine learning workflow as the information contained in the set of training examples is used by the model to "learn" how to separate future inputs—in this case, pixels of multispectral and multitemporal imagery—into the desired output classes. For the specific application of land cover mapping, large amounts of training data are needed across a wide variety of biomes and climatological regions. The spectral signature of a single land cover type can vary significantly from place to place even within CONUS (e.g., "deciduous forest" looks quite different in California vs. Maine). Therefore, a machine learning classifier can only learn robust rules for how to segment the complex feature space of the input imagery if it has been trained on sample points that cover many examples of the aforementioned within-class spectral variability.

Crucially, a major driver of such intra-categorical spectral variability is the presence of snow in the winter imagery scene. Figure 3 shows the extent to which snow cover changes the first two principal components of the 18-dimensional spectral signature encoded in two dates' worth of 9-band multispectral S2 imagery. For most of the land cover classes (the exceptions being low vegetation and cropland), the spatial division between snow and no-snow points is so clear that it is possible to draw a curve on the figure that would separate them almost perfectly.

There is one additional quality of the snow-cover-induced spectral variability that is not captured by Figure 3. Even in cases where the spectral signature of snow-covered and snow-free points appear to overlap, it is possible (and, in practice, common—see Section 4 below) for the snow-covered points to overlap more closely with snow-free members of a different class than they do with members of their own class. When a machine learning model that has trained only on snow-free points is

applied to pixels where snow is present, it will therefore mislabel such snow-covered pixels.

Figure 3. Spectral separation of training points with and without snow cover in the winter S2 image. Axes show the first two components of a Principal Component Analysis rotation in order to capture the effect across all 18 input bands in a 2-dimensional figure.

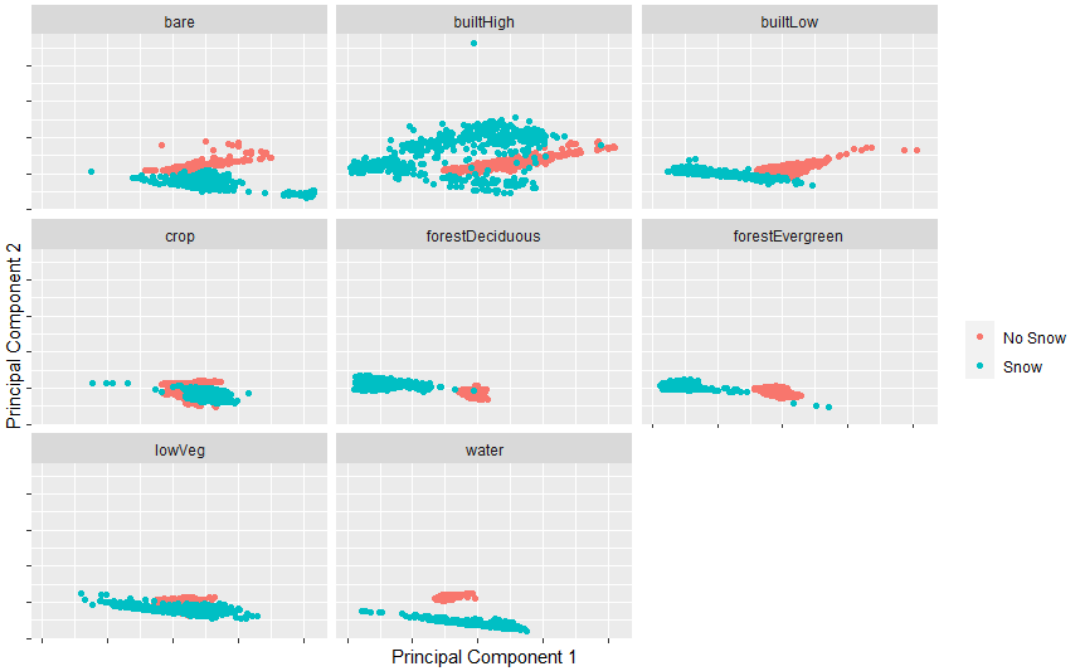

To correct this shortcoming of the V1 tool, 6,466 additional sample points were collected from the snow-covered training sites described in Section 2.2. The spectral profile of the corresponding two-date S2 image stack was extracted at each of these points, the result converted into a machine-learning-ready training dataset. Table 2 shows the breakdown of new training point counts by land cover type.

The new training data, collected entirely from image pairs containing snow cover on one or both dates, were concatenated and shuffled into the original dataset that was used to train V1 of the semi-automated land cover classification tool. The resulting combined dataset contains 14,315 observations over 12 land cover classes (the 11 classes present in the V1 training dataset, along with a new “Permanent Snow” class).

Table 2. Number of snow-covered training samples by land cover type.

Land Cover Type	Number of Pixels
Water	700
Bare Ground	1,220
Sparse / Low Vegetation	574
Forest (Deciduous)	555
Forest (Evergreen)	546
Built-Up (High Density)	584
Built-Up (Low Density)	592
Cropland	788
Permanent Snow	907

3.3 Model training

Once training data were created, a linear SVM model was built and 85% of the training data was applied. The remaining 15% was withheld for accuracy assessment. The SVM was fit to the training data, enabling the model to classify land cover type. The model was trained using Python 3.6.9 and the support vector machine function in the scikit-learn library. The linear kernel was used in combination with default hyper-parameters; in grid-search testing, modification of hyper-parameters was not found to improve final model performance.

The trained model was then transferred and written into ArcGIS Pro Python toolboxes to enable the prediction of land cover type across different regions based on user input. This trained model can later be updated based on improvements to the model, such as expansion to areas outside the continental United States (OCONUS).

3.4 Python toolbox design

3.4.1 ETP land cover mapping Python toolboxes

ArcGIS Pro contains a suite of commercial, closed-source geospatial processing tools. These tools can be accessed within the ArcGIS Pro GUI through the ArcToolBox functionality. Custom toolboxes and tools can be created using Python programming, which enables new analysis capabilities with third-party Python libraries directly within the ArcGIS Pro GUI.

V1 of the ArcGIS Pro tools included a set of image pre-processing tools, cloud-masking tools, as well as the two-date S2 classification and two-date Landsat 8 classification. V2 updates the toolbox to include a new S2 two-date classification tool that accounts for snow cover, as shown in Figure 4. The classes represented in each version of the two-date S2 classification tool (V1 and V2) are shown in Table 3.

Figure 4. View of the land cover tools as shown integrated into ArcGIS Pro as a Python toolbox. The tools are integrated as a subset of the Enhanced Terrain Processing Toolkit.

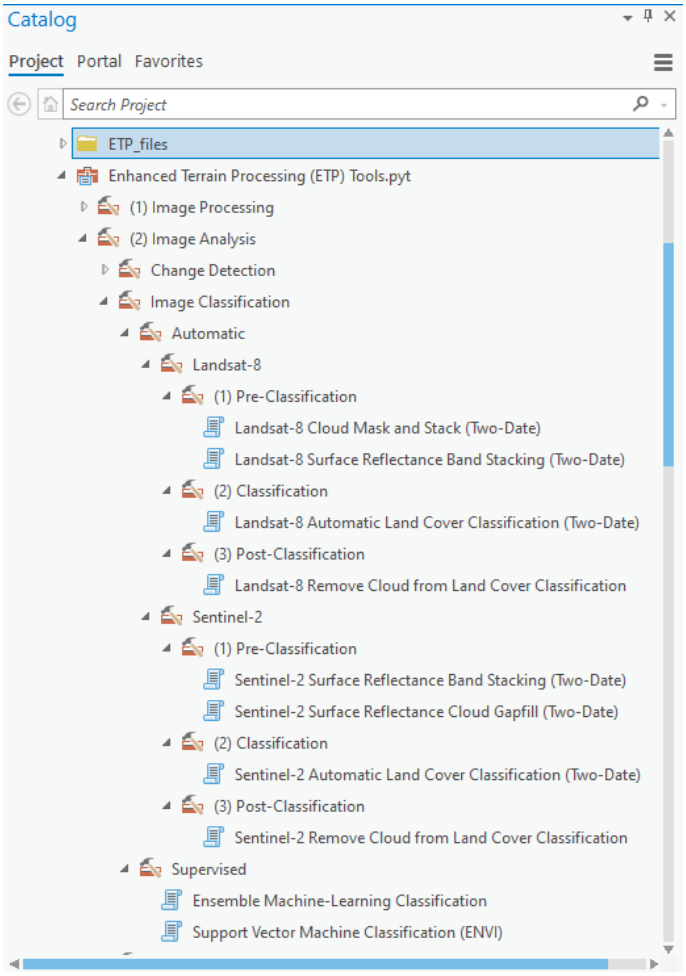

The V2 tool maintains the same image chunking and parallel processing described in Lasko and Sava (2021) for faster execution and compatibility on computers with minimal hardware resources. The output class labels for V2 use the same encoding scheme as the VisNav (Visual Navigation)

dataset and can easily be re-categorized and used interchangeably with other terrain analysis land cover products.

Table 3. Land cover types generated by the two-date classification tool.

V1 Land Cover Types	V2 Land Cover Types
Water	Water
Barren	Barren
Scrub, Shrub, grasses	Scrub, Shrub, grasses
Deciduous Trees	Deciduous Trees
Evergreen Trees	Evergreen Trees
Herbaceous Wetlands	Herbaceous Wetlands
Woody Wetlands	Woody Wetlands
Perennial Water	Perennial Water
Built-up (Low Density)	Built-up (Low Density)
Built-up (High Density)	Built-up (High Density)
Cropland	Cropland
-	Permanent Snow

3.4.2 Installation and usage of ArcGIS Pro tools

The custom ArcGIS Pro Python-based toolbox containing the land cover mapping capabilities is compatible with any Windows computer that has ArcGIS Pro version 2.5 or higher installed and at least 8GB of RAM. Easy-to-follow instructions on installation, usage, and use cases are available in separate tutorial and user guide documents.

4 Example Output

4.1 Comparison of V1 and V2

Performance improvement between V1 and V2 of the semi-automated land cover classification tool is most striking when comparing the output of the tool over a region subset where snow appears in the winter S2 image. Figure 5 and Figure 6 showcase the poor performance of the V1 classifier in such regions and the contrast with the output of the V2classifier. All of the images in this section are drawn from S2 scenes that were not used for model training.

Figure 5. Input and output of tool over Montana Rockies, showing (a) winter satellite imagery (SWIR-NIR-Red), (b) summer satellite imagery (SWIR-NIR-Red), (c) V1 classifier output, (d) V2 classifier output.

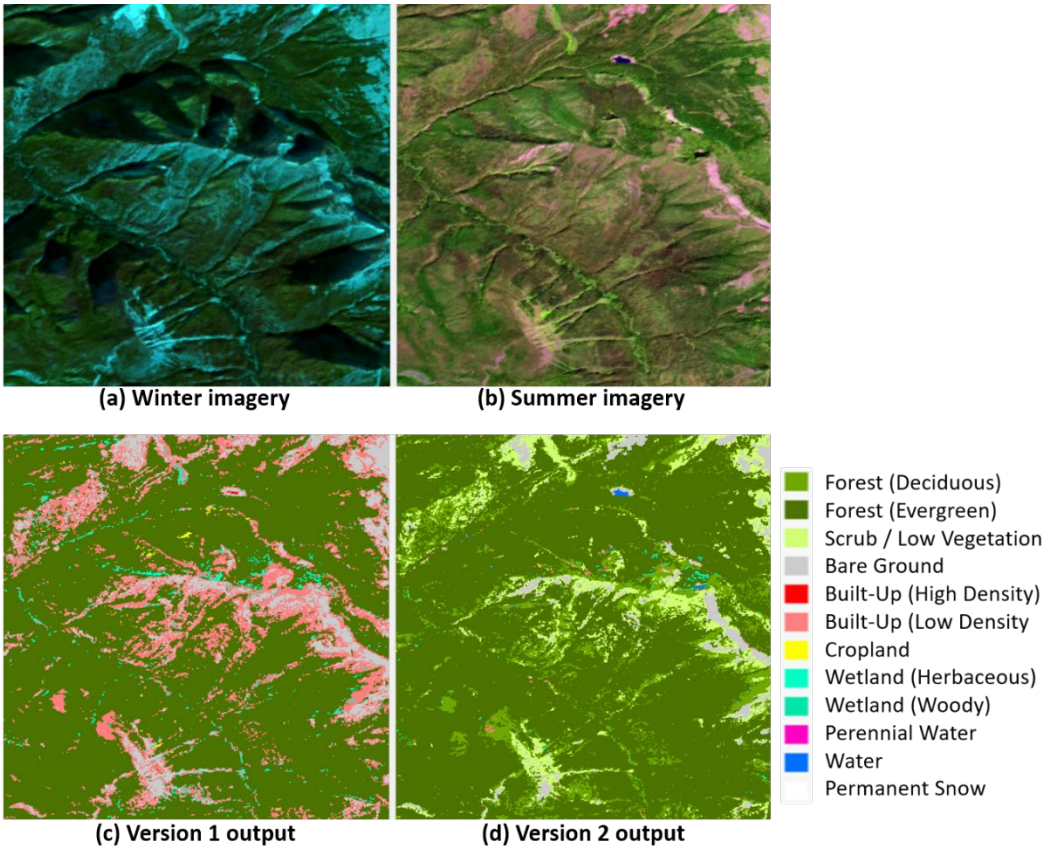

Figure 5 shows the tendency of the V1 classifier to label snowy pixels in mountainous regions as built-up low density rather than their true class value, in this case mainly scrub/low vegetation. The V1 classifier also mistakenly identifies a narrow strip along the ridge of the central mountain as herbaceous wetland, clearly incorrect given the biome in

question. The V2 classifier output, seen in the lower right of Figure 5, corrects the majority of these mislabeled pixels, including recoding the small lake in the upper center from built-up high density to its true class (water).

A similar tendency to misclassify snow-covered vegetation as either built-up low density or wetland is observed in non-mountainous regions, shown in Figure 6. All the Kansas fields in the scene are mislabeled by the V1 classifier. The deciduous forest in the upper third of the image suffers a similar fate, being grouped into the wetland category along with the adjacent fields. The only pixels correctly identified by the V1 tool are a few patches of evergreen forest in the lower center and bottom left corner. In contrast, the V2 output once again shows much closer correspondence to the land cover classes one would expect from visual inspection of the two dates' worth of satellite imagery.

Figure 6. Input and output of tool over Kansas farmland, showing (a) winter satellite imagery (SWIR-NIR-Red), (b) summer satellite imagery (SWIR-NIR-Red), (c) V1 classifier output, (d) V2 classifier output.

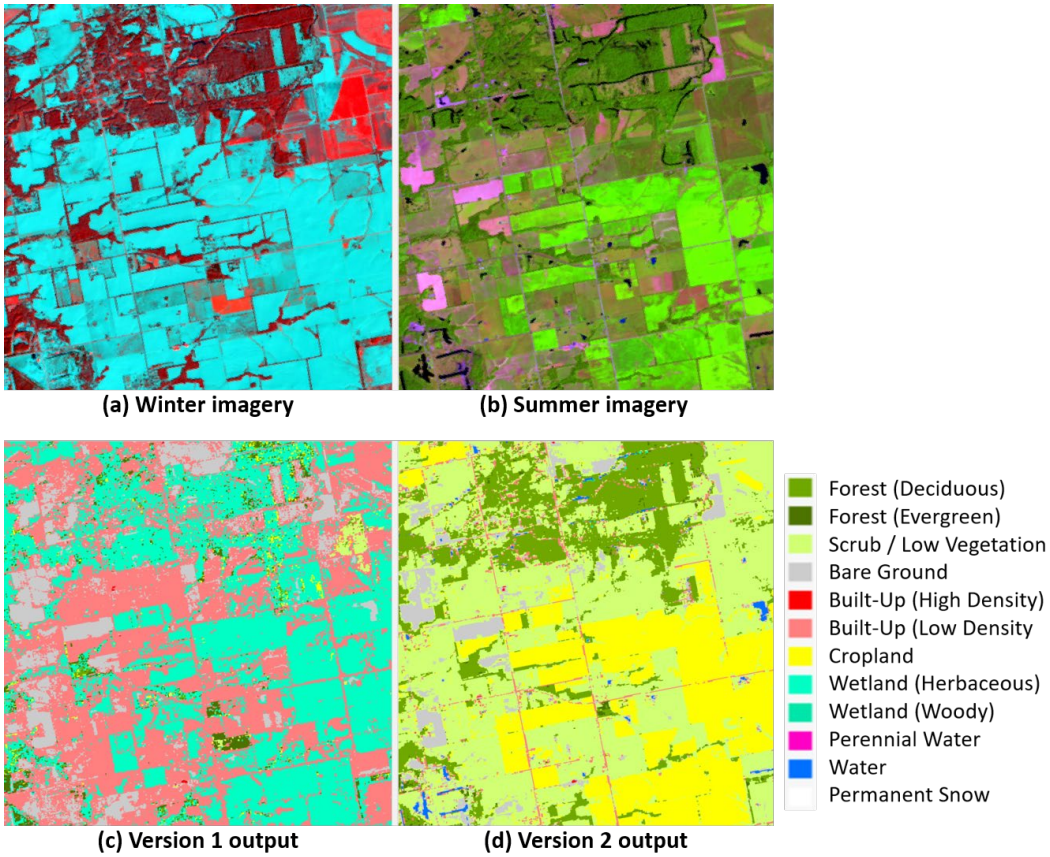

When adding capabilities to an existing machine learning model, there is always the need to address concerns as to whether the expanded

performance in one direction comes at the expense of the model's existing capacity in other domains. In the example of the semi-automated land cover classification tool, one might worry that the addition of snow-cover training data could "confuse" the enhanced model when it is presented to scenes without snow in the winter image. Figure 7 shows that this is not the case; in the Louisiana scene shown, the output of the V1 and V2 classifiers is nearly identical. The new model retains the ability of its predecessor to accurately identify land cover types in snow-free regions. The Mississippi point bar in the upper right quadrant of the scene shows that the model can accurately map herbaceous wetland, one of the most difficult land cover types to distinguish given the relative sparsity of training data.

Figure 7. Input and output of tool over Mississippi River in southern Louisiana, showing (a) winter satellite imagery (SWIR-NIR-Red), (b) summer satellite imagery (SWIR-NIR-Red), (c) V1 classifier output, (d) V2 classifier output.

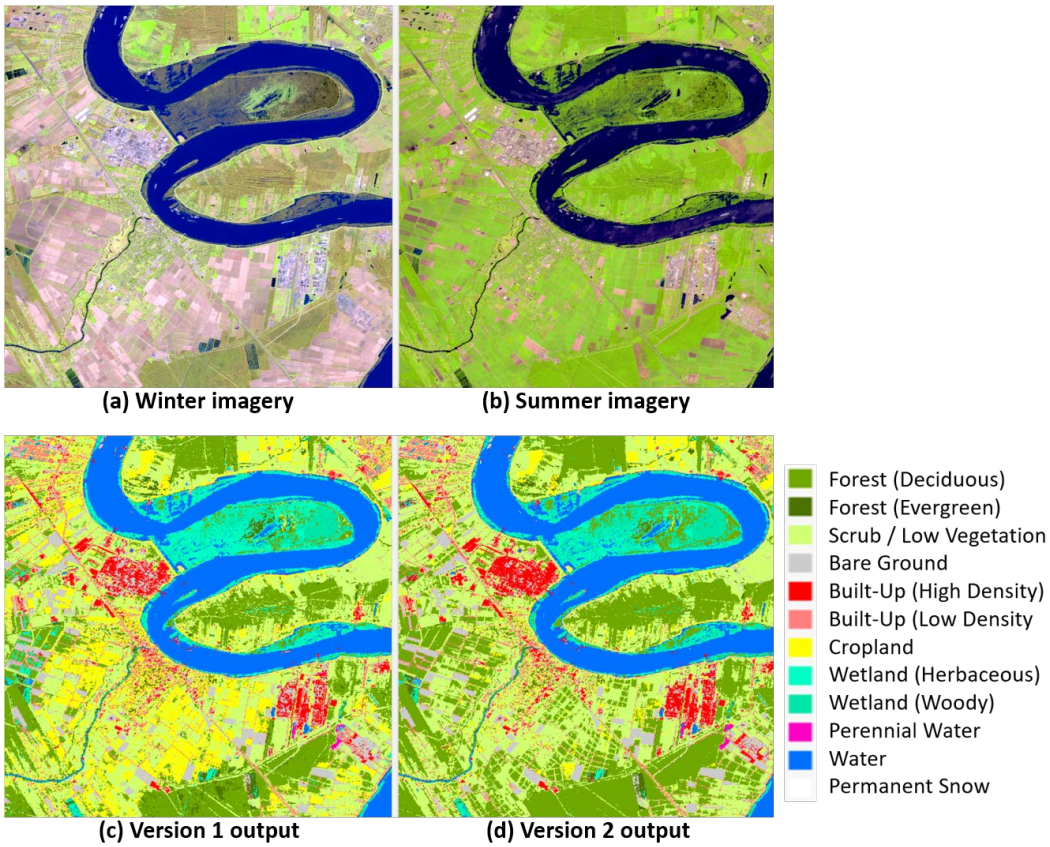

Despite significant performance gains in snow-covered landscapes, however, the V2 classifier is not without its own known flaws. Figure 8 shows a scene in the Montana Rockies (not far from the area shown in Figure 5) where the V2 classifier struggles to distinguish bare ground from

built-up high density. This issue is concentrated in areas of topographic shadow, where the bulk of the mountains themselves cast shadows over the part of their slopes that lie opposite the sun. The relative darkness of the shaded pixels compared to the other members of their land cover class interferes with the model's ability to accurately label them. Note, however, that the V2 output is still substantially more accurate than that of the V1 tool, which, in addition to the errors of its successor, also misclassifies the large swathes of low vegetation in the scene as either built-up low density or herbaceous wetland.

Figure 8. Input and output of tool over another part of Montana Rockies, showing a) winter satellite imagery (SWIR-NIR-Red), (b) summer satellite imagery (SWIR-NIR-Red), (c) V1 classifier output, (d) V2 classifier output.

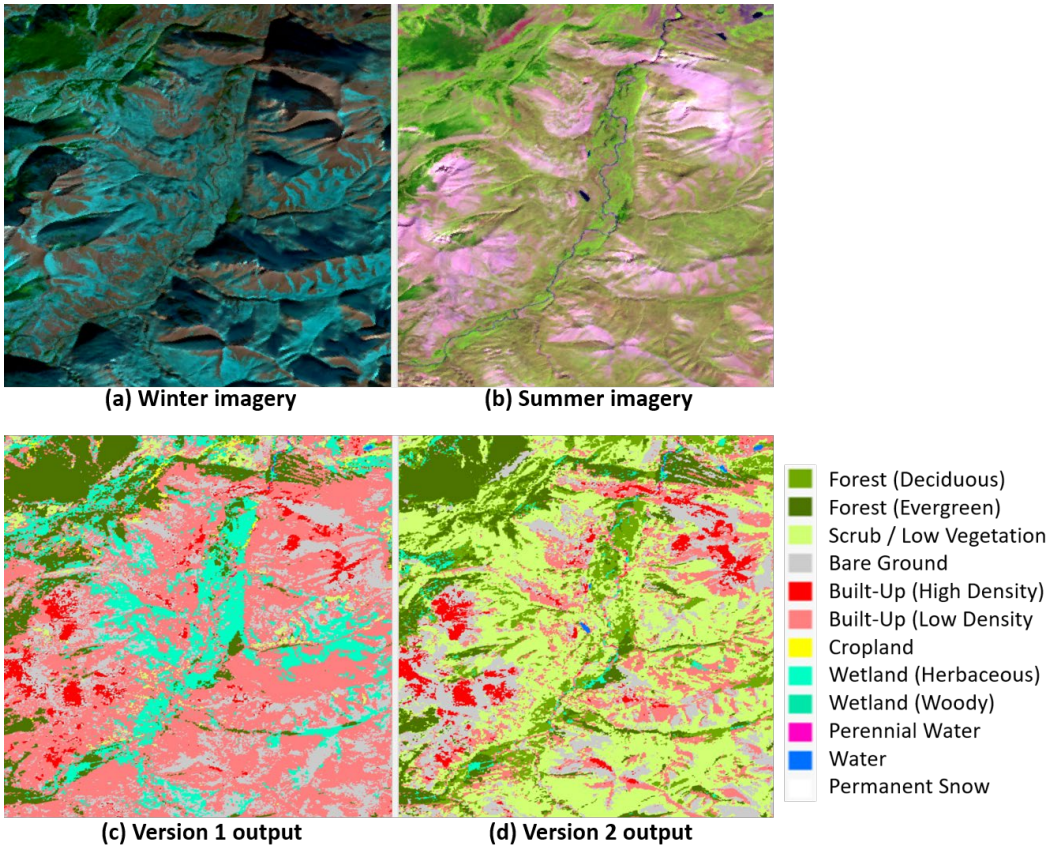

4.2 Comparison to NLCD

Figure 9, Figure 10, and Figure 11 present a comparison between the output of the V2 semi-automated land cover mapping tool and the 2019 NLCD over the same area. All the S2 data in this section are drawn from scenes that were not used for model training. Note that the NLCD is produced at 30 m resolution, while the land cover tool output has 20 m

pixels. The NLCD map was produced for year 2019 using supervised machine learning classification from a dense time series of Landsat 8 satellite observations. This product is freely available from the USGS Earth Explorer website. NLCD land cover labels were re-categorized to match those used in this project, as described in Table 4.

Table 4. Mapping of NLCD land cover classes to semi-automated land cover classifier labels.

NLCD Label	V2 Output Label
Open Water	Water
Perennial Snow/Ice	Snow
Developed, Low Intensity	Built-Up (Low Density)
Developed, Medium Intensity	Built-Up (Low Density)
Developed, High Intensity	Built-Up (High Density)
Barren Land	Bare Ground
Deciduous Forest	Forest (Deciduous)
Evergreen Forest	Forest (Evergreen)
Mixed Forest	Forest (Evergreen)
Shrub/Scrub	Low Vegetation
Herbaceous	Low Vegetation
Hay/Pasture	Low Vegetation
Cultivated Crops	Cropland
Woody Wetlands	Wetlands (Woody)
Emergent Herbaceous Wetlands	Wetlands (Herbaceous)

The area near Nashville, TN, shown in Figure 9 exemplifies the overall agreement between the output of the V2 tool and the NLCD. In particular, the contiguous patches of forest cover match up well between the two products, and the contour of the highway system is visible in both. The V2 tool mislabels the large quarry in the center of the scene as built-up cover, while the NLCD correctly identifies it as bare ground. Most notably, however, there is substantial disagreement between the two products surrounding the cropland class, especially in comparison with the low vegetation and/or bare ground classes. These cover types often suffer from substantial spectral overlap; in fact, cropland is in many ways more of a land use than a land cover. When observed at only one or two temporal points, a given cropped field may appear identical to a patch of low vegetation or barren ground. This issue is discussed in more detail in Section 5.

Figure 9. Tool output vs. NLCD over Tennessee, showing (a) winter satellite imagery, (b) summer satellite imagery, (c) 2019 NLCD, (d) V2 classifier output.

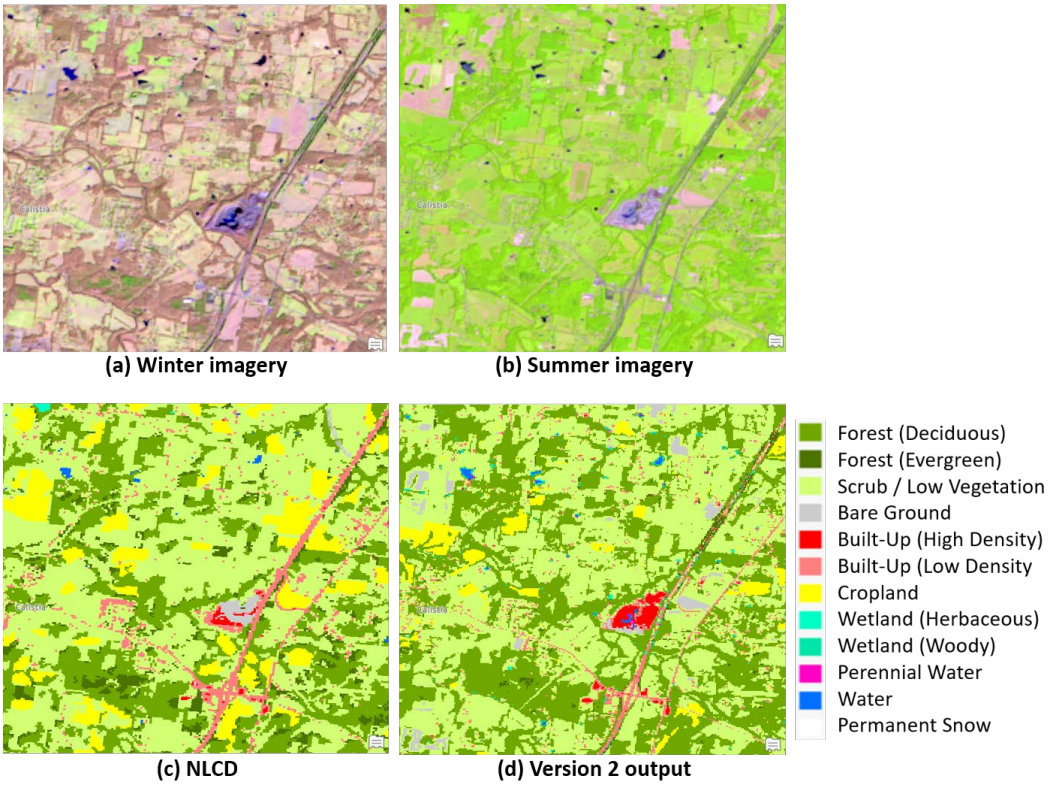

A similar case, in an area with snow cover in the winter scene, is shown in Figure 10. This image subset comes from a different part of the same Kansas scene shown in Figure 6 in Section 4.1. Similarly to Figure 9, here the water and forest classes show good agreement between the NLCD and V2 output, while the cropland, low vegetation, and bare ground categories differ substantially. As above, these varying labels are due to the spectral similarity between these cover types. In this case, with snow on the ground in the winter scene, one must also note that the V2 tool is forced to use only the summer scene for much of its labeling purposes. The NLCD, on the other hand, has a dense time series of an entire year at its disposal. Note, for example, the large area in the upper center of the scene, labeled "bare" by the V2 classifier but "low vegetation" by the NLCD. While these areas may indeed be vegetated for some part of the year, in the two S2 images used by the V2 classifier, no vegetation is visible. This temporal variability in land cover type accounts for a large part of the disagreement between these two thematic maps.

Figure 10. Tool output vs. NLCD over Kansas, showing (a) winter satellite imagery, (b) summer satellite imagery, (c) 2019 NLCD, (d) V2 classifier output.

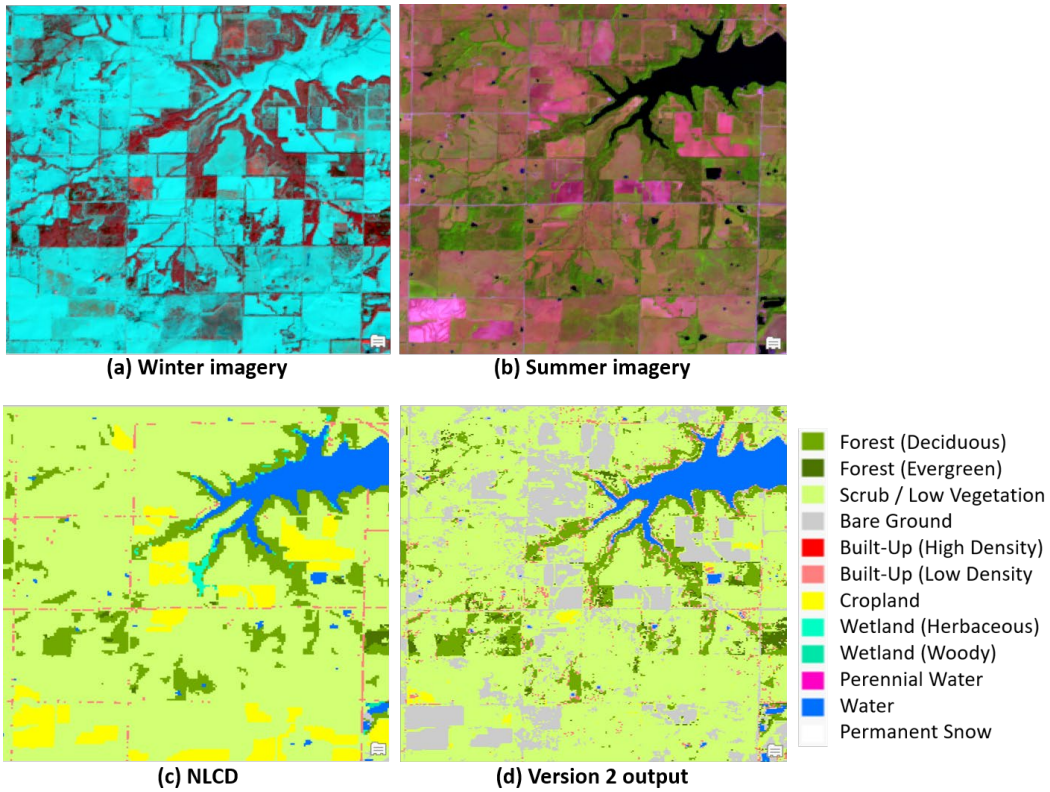

The previous example showed a case in which the input imagery seems to support the label chosen by the V2 classifier rather than the NLCD. Nowhere is this more apparent than in Figure 11, which shows the labels of both classes over a part of the Sierra Nevada Mountains in California. The NLCD shows the majority of the scene as covered in low vegetation, while the V2 classifier output labels those same pixels as bare ground. An examination of the input S2 imagery, however, provides far more evidence for the bare ground label than for low vegetation. Any plant life is conspicuously absent in both the winter scene (when the mountains are covered in snow) and in the summer scene. If vegetation is present, it is too scattered and sparse to be readily seen at 20 m resolution. This might be another example of temporal variability in vegetation cover, previously seen in Figure 10. If this is the case, however, it is impossible to determine from the two S2 images that form the input for the V2 classifier. It may also be that the land cover type has changed between 2019 (when the latest NLCD was created) and 2020 (when the summer S2 image was captured)—note the importance of an up-to-date land cover product.

Figure 11. Tool output vs. NLCD over the Sierra Nevada Mountains, showing (a) winter satellite imagery, (b) summer satellite imagery, (c) 2019 NLCD, (d) V2 classifier output.

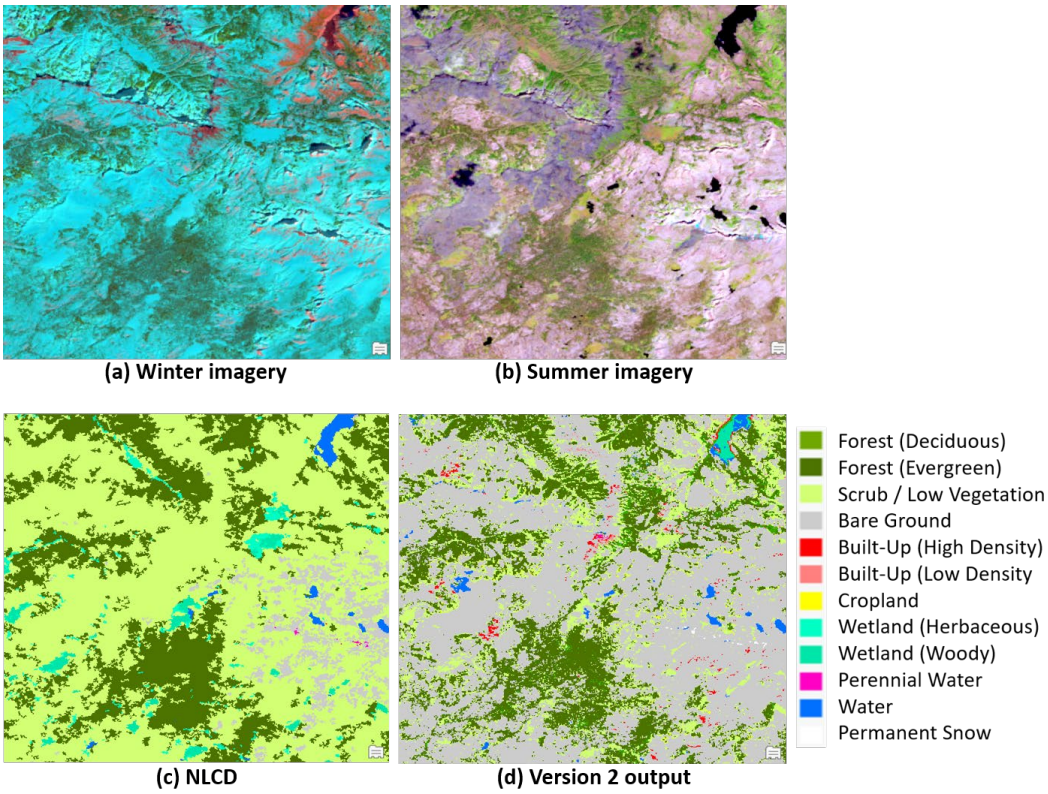

In fact, the possibility of land cover change should always be kept in mind when comparing the static 2019 NLCD to the V2 classifier output, which is generated on more recent satellite imagery. Cropland, the subject of much discussion around the first two figures in this section, may be planted in one year and lie fallow in the next (thus causing it to appear either as bare or sparsely vegetated in a satellite image). A key advantage of the land cover tool presented here is its ability to generate up-to-date maps using the latest earth observation data to capture these year-to-year changes.

5 Accuracy Assessment

Accuracy assessment is a critical part of any image classification project. To measure the accuracy of both classification algorithms, we used two methods. First, 15% of the training dataset was initially withheld from model training (i.e., not used to train machine learning algorithm) and was then used for model evaluation. Then, once the model was fully trained, additional sample points were collected on S2 scenes outside the training dataset extent. Both versions of the tool were then evaluated again using these new points.

Traditional machine learning best practices rely heavily on the first method of accuracy assessment: holding out a random subset of the training data for use as a validation dataset. However, when working in geospatial applications it is important to remember the so-called First Law of Geography: "Everything is related to everything else, but near things are more related than distant things" (Tobler 1970). In the case of classification accuracy, this means that holding out validation points from within the training dataset will give an inflated measure of the model's performance. Even though the classifier did not see the validation points during training, random sampling means that it likely did train on at least a few points that are spatially near to each of the validation points. Thus, since within-class spectral signatures are spatially autocorrelated, the overall effect is similar to validating on the training dataset itself (a known faux-pas in machine learning).

The second method of accuracy assessment—choosing additional test points in new geographic regions—seeks to mitigate the above issues. The additional testing points used in this method were chosen over sites not used for the collection of training data, and thus the accuracy of the classifier when labeling them is more representative of its performance in general.

The calculation of within-class accuracy for a given model and test dataset is fairly simple: one need only count the number of correctly labeled points in that class and then divide by the total number of points in that class. Aggregating those class-wise accuracy values into an overall accuracy measure, however, is more complicated. Simply taking the mean of the class accuracies is not sufficient; such a method gives equal weight to each class, when in fact some land cover types (such as deciduous forest) are

much more common than others (such as herbaceous wetland). The overall accuracy of a classifier should represent the likelihood of correctly labeling an arbitrary pixel within the applicable region (in this case, CONUS). Thus, the class-wise accuracies must be weighted by an outside measure of the prevalence of each class.

We used the 2019 NLCD to calculate these class weights. The NLCD land cover classes were re-mapped to match the labels used in this project, as described in Table 4 in Section 4.2. Pixel counts of each cover type across CONUS were normalized to create a measure of each class's prevalence, shown in Table 5. These values were used as weights to calculate the weighted mean of class-wise accuracy for each of the two assessment methods. This weighted mean represents the overall classifier accuracy.

Table 5. Relative prevalence of land cover types in CONUS, calculated from 2019 NLCD.

Land Cover Class	Normalized Prevalence
Water	0.053
Bare Ground	0.010
Forest (Deciduous)	0.127
Forest (Evergreen)	0.115
Low Vegetation	0.441
Cropland	0.163
Built-Up Low Density	0.027
Built-Up High Density	0.004
Herbaceous Wetland	0.015
Woody Wetland	0.045

Given that the cropland class blurs the line between land cover and land use, along with its spectral similarity to scrub/low vegetation (the two are often distinguishable to the human eye only by the geometric shapes of cropped fields, information not available to a pixel-wise classifier), we perform each measure of accuracy two times: once with cropland and low vegetation as two separate classes and once with their values considered equal. As expected, aggregating cropland into the low-vegetation class substantially improves overall accuracy since confusion between the two labels is common. Results from both methods are reported in each of the sections below.

5.1 Data held out from training

When training both versions of the semi-automated land cover classifier, 15% of the training dataset is held in reserve for purposes of validation. This amounts to 1,441 samples from the V1 training dataset and 2,147 samples from the V2 training dataset. The trained classification models were used to predict land cover type from the spectral signature of each validation sample, and the resulting predictions were then compared to the original label. A confusion matrix of the results is shown in Figure 12.

Figure 12. Confusion matrices of validation data accuracy, faceted by classifier version. Value shown is percentage of target pixels that receive a given predicted label; thus columns sum to 1 but rows may not.

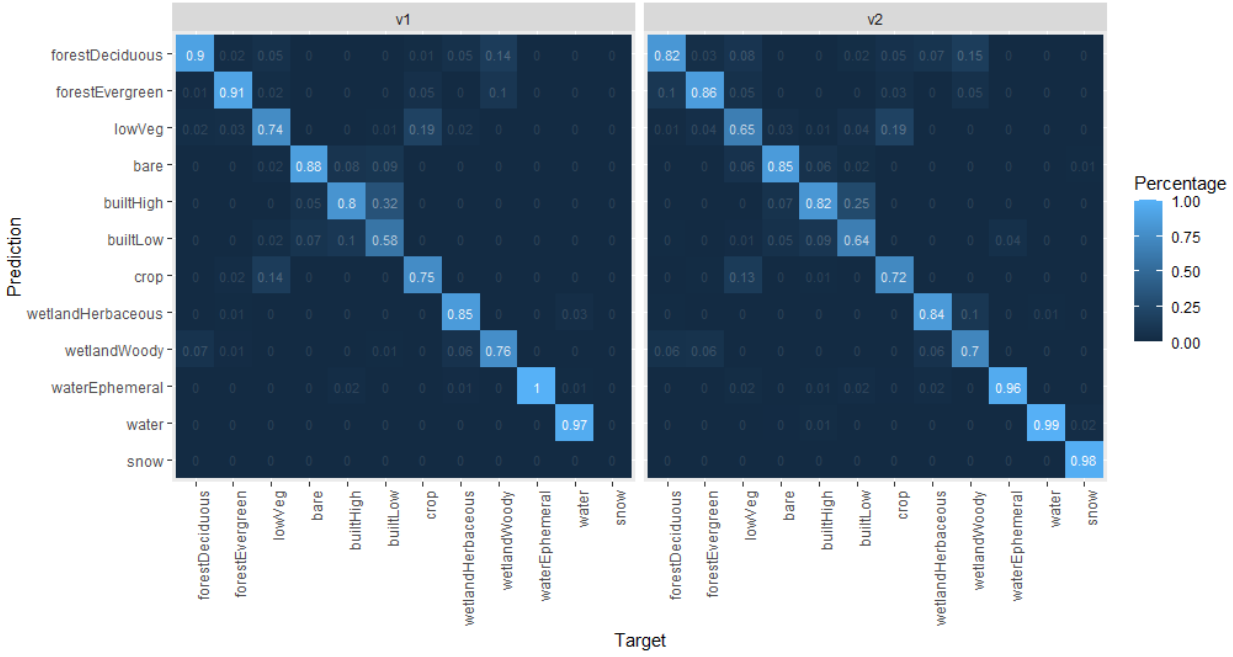

Classification accuracy as calculated in this way is not divided by snow cover (the analysis in Section 5.2, however, does include such a distinction). The training samples are randomly shuffled before selection of the validation subset, and information regarding whether a given sample comes from a snow-covered scene is obfuscated in the process. Taken in isolation, this analysis appears to indicate that the V2 classifier is less accurate than its predecessor in most categories (the exceptions being bare ground and water). The more detailed and representative assessment in Section 5.2, however, will show that this loss of accuracy is an illusion. Note also that the validation dataset used for this analysis of the V1 classifier is significantly smaller than that used to assess the V2 classifier,

as by definition the V1 training dataset does not include any of the snow-covered sample points.

For both versions of the classifier, the least accurate cover type in this analysis was Built-Up Low Density. A majority of misclassification in that category stemmed from mislabeling as Built-Up High Density, a similar cover type whose definition has a somewhat fuzzy overlap with Built-Up Low Density. Both classifier versions also show confusion between the crop and low vegetation classes, a known issue addressed in the previous section.

Overall validation accuracy of the two classifier versions is shown in Table 6. When cropland and scrub/low vegetation are kept as separate land cover classes, neither version of the tool reaches the 80% benchmark dictated by the mission requirements. Combining those two classes into one, however, results in both versions exceeding the 80% threshold.

Table 6. Overall classifier accuracy over validation samples held out from training dataset.

-	Crop and Low Vegetation Separate	Crop and Low Vegetation Consolidated
V1	0.798	0.908
V2	0.754	0.868

5.2 Additional test data

To form a more unbiased picture of the land cover classifier's performance on imagery that was not used to generate the training dataset, an additional assessment was performed using five testing scenes of two-date S2 imagery, none of which contain any points used for model training. Details on the selection of test scene selection and placement can be found in Section 2.2.

Sample points were manually placed and labeled over corresponding pixels in each of the test imagery scenes. Table 7 shows the distribution of test pixels by land cover type.

Table 7. Number of test samples by land cover type.

Land Cover Type	Number of Samples
Water	542
Bare Ground	576
Sparse / Low Vegetation	475
Forest (Deciduous)	408
Forest (Evergreen)	488
Wetland (Herbaceous)	88
Wetland (Woody)	145
Ephemeral Water	137
Built-Up (High Density)	451
Built-Up (Low Density)	468
Cropland	446
Permanent Snow	79

Both versions (V1 and V2) of the semi-automated land cover classification tool were then used to create output classified maps from the five validation image pairs. The predicted land cover class in those outputs were extracted at the sample point locations and compared to the target labels determined prior to the application of the tool. Figure 13 shows a confusion matrix of the results.

In scenes without snow cover, accuracy for all classes has either remained unchanged or improved from V1 to V2, with the exception of woody wetland (which saw a 0.14 percentage point decrease). These improvements are minor across the board, in the range of two or three percentage points.

In scenes with snow cover, meanwhile, the performance improvement between versions is more dramatic. Using the predicted labels of the V1 classifier, only four out of the eight classes with snow-covered representation achieve accuracies above 50%. In fact, for most snow-covered classes the V1 classifier performs worse than chance. With the labels from the V2 classifier, in contrast, all eight snow-covered classes reach above 50% correct labeling, with 100% of permanent snow and snow-covered cropland points being correctly identified by the V2 algorithm.

Figure 13. Confusion matrices of test data accuracy. Faceted by classifier version (horizontal) and snow vs. no-snow (vertical). Value shown is percentage of target pixels that receive a given predicted label; thus columns sum to 1 but rows may not.

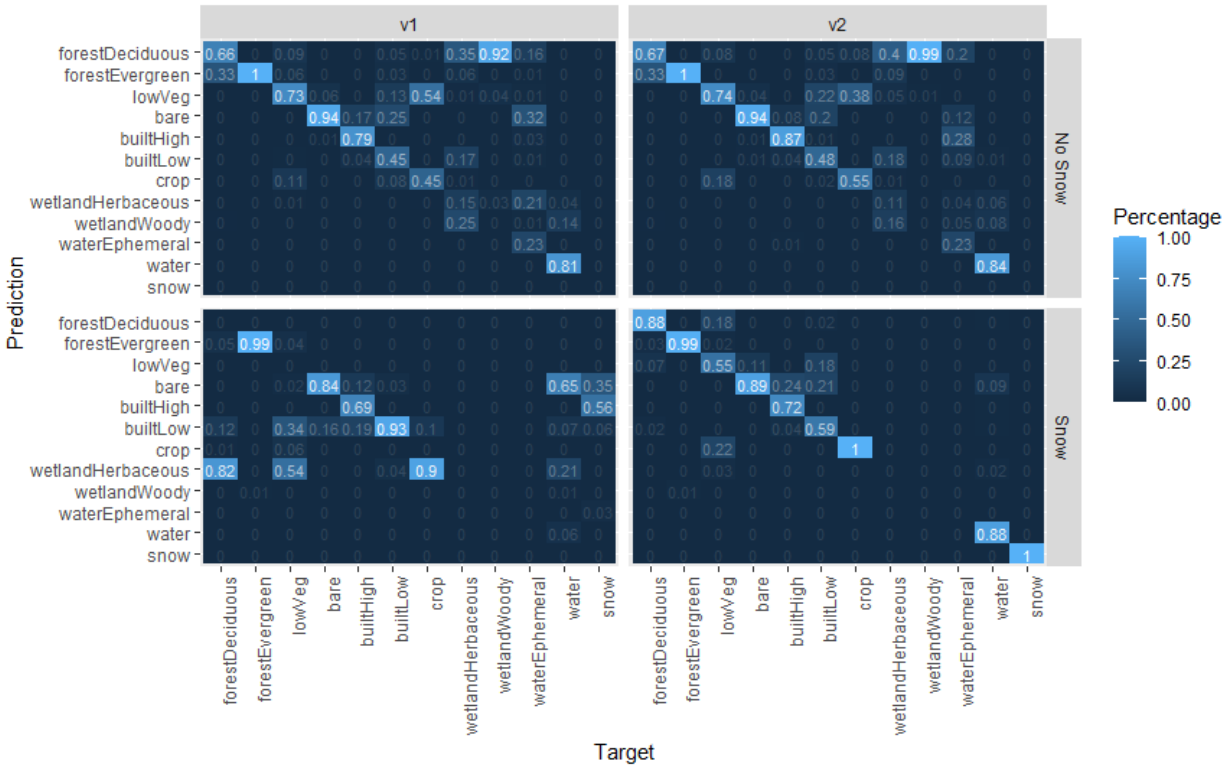

In imagery without winter snow, the V2 classifier continues to struggle with built-up low density pixels, cropland, both wetland classes (woody and herbaceous), and the perennial water class. Three of these five cases can be explained by the overlap of spectral signatures between each of these classes and another output class. Built-up low density cover represents a mixture of impervious surface and vegetation; as a result, many pixels in this class are misidentified as belonging to the low vegetation class. Woody wetland, from above, has a canopy similar to that of deciduous forest, which is indeed where the classifier incorrectly categorizes nearly all such pixels. Pixels from the cropland class, meanwhile, are often (38%) misclassified as low vegetation; in fact, the distinction between these two output classes has more to do with land use than land cover.

The confusion around herbaceous wetland and perennial water do not admit quite the same spectral justification as the other three. The misclassification of these two labels is spread around three or more other classes. It is notable that both herbaceous wetland and perennial water are in the lowest three classes in terms of number of training points (639 and

162, respectively). Although the SVM weights training points in order to avoid bias in favor of one class over another, the dearth of training input means that the variety of spectral signatures learned by the SVM for these two classes is more limited than for the other cover types.

In imagery where snow is present in the winter scene, some of the same classification errors remain. When separating cropland from low vegetation, the issue seen in snow-free imagery is reversed: cropland is correctly labeled, while low vegetation pixels are sometimes (22%) misclassified as crop. Built-up low density pixels suffer from nearly identical confusion to that seen in snow-free images, while there is some additional slippage (24%) from built-up high density into the bare ground class. This may be due to the fact that when snow is present in the winter scene, both bare ground and high density impervious surfaces are usually covered in a layer of snow. Since this snow cover has a nearly identical spectral signature regardless of whether the ground below is bare or built up, the classifier must rely solely on the summer image to distinguish the two categories, resulting in reduced accuracy.

The improvements in test accuracy described above stand in stark contrast to the decline seen in the validation accuracy between V1 and V2 described in Section 5.1. The pattern of improvement remains the same when considering the overall accuracy across all categories (weighted by NLCD prevalence), which is shown in Table 8. Note that, while absolute test accuracy numbers are lower than those seen in the validation accuracy (as expected), the results when cropland and low vegetation are consolidated into a single class label still exceed the 80% threshold dictated by the mission requirements. Moreover, the shockingly low accuracy of V1 on snow-covered pixels makes clear the necessity of the transition to V2.

Table 8. Overall classifier accuracy over additional test points.

Version and Snow Coverage	Crop and Low Vegetation Separate	Crop and Low Vegetation Consolidated
V1 with No Snow	0.664	0.828
V1 with Snow	0.163	0.189
V2 with No Snow	0.686	0.827
V2 with Snow	0.753	0.869

6 Conclusion

This project led to the development of a new version of the semi-automated land cover classification tool, which uses a pre-trained machine learning model to create land cover type maps for use in CONUS. New sample points were collected from sites across CONUS to incorporate snow-covered terrain into the machine learning training dataset. Using the combined dataset, a pre-trained machine learning model was created for use with two dates of S2 multispectral imagery. The satellite imagery model was then integrated into ArcGIS Pro GUI using Python toolboxes with parallel processing and image chunking to serve as an easily accessible decision support tool for USACE customers.

The decision support tools enable a user to create a land cover type map without the need to generate training data or to rely on costly or outdated static maps. Further, the ArcGIS Pro tools allow a map to be generated in a matter of minutes with minimal user input. The land cover type class scheme for V2 is equivalent to that used in VisNav and can easily be re-categorized for interchangeability with other land cover maps used in terrain analysis applications.

The estimated overall accuracy of the V2 classifier is 83% over snow-covered pixels and 87% over snow-free pixels. Compared to V1 of the tool, V2 achieves superior performance in cases when snow cover is present in one or both of the input images (accuracy 3.6 times higher) while retaining the capabilities of V1 in snow-free scenes. After testing the imagery across varied ecoregions of CONUS, the study found future areas of improvement to the models should include better separability of wetland cover classes, built-up low-density areas, perennial water, and topographic shadow. These errors are attributed both to spectral similarities found across the geographically-distributed training imagery and to the low prevalence and high within-class spectral variability of several of the above-land cover types, particularly wetland and perennial water.

Ultimately, the semi-automated land cover mapping tool serves as a foundation on which to build more advanced land cover mapping tools. Future improvements under this project will expand the geographic coverage of the model outside of CONUS, leverage advanced deep learning frameworks for more robust classification, and incorporate higher spatial resolution sensors.

References

Biemans, H., C. Siderius, A. Lutz, S. Nepal, B. Ahmad, T. Hassan, W. von Bloh, et al. 2019. "Importance of Snow and Glacier Meltwater for Agriculture on the Indo-Gangetic Plain." *Nature Sustainability* 2 (7): 594–601.

Büttner, G. 2014. "CORINE Land Cover and Land Cover Change Products." In I. Manakos and M. Braun, *Land Use and Land Cover Mapping in Europe*. Dordrecht: Springer.

Congalton, R. G., and K. Green. 2019. *Assessing the Accuracy of Remotely Sensed Data: Principles and Practices*. Boca Raton: CRC Press.

Cunningham, D. J., J. E. Melican, E. Wemmelmann, and T. B. Jones. 2002. "GeoCover LC—A Moderate Resolution Global Landcover Database." *ESRI Proceedings*.

Dierauer, J. R., D. M. Allen, and P. H. Whitfield. 2019. "Snow Drought Risk and Susceptibility in the Western United States and Southwestern Canada." *Water Resources Research* 55 (4): 3076–3091.

European Space Agency. 2012. *Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services*. Noordwijk: ESA Communications. Retrieved from https://sentinel.esa.int/documents/247904/349490/s2_sp-1322_2.pdf.

Gómez, C., J. C. White, and M. A. Wulder. 2016. "Optical Remotely Sensed Time Series Data for Land Cover Classification: A Review." *ISPRS Journal of Photogrammetry and Remote Sensing* 116: 55–72.

Jin, S., C. Homer, J. Dewitz, P. Danielson, and D. Howard. 2019. "National Land Cover Database (NLCD) 2016 Science Research Products." In *AGU Fall Meeting Abstracts* (pp. B11I–2301). Retrieved from https://ui.adsabs.harvard.edu/abs/2019AGUFM.B11I2301J.

Kongoli, C., P. Romanov, and R. Ferraro. 2012. "Snow Cover Monitoring from Remote Sensing Satellites: Possibilities for Drought Assessment." In B. D. Wardlow and M. C. Anderson, *Remote Sensing of Drought: Innovative Monitoring Approaches*. Boca Raton, FL: CRC Press, 359–383.

Lasko, K., and E. Sava. 2021. *Semi-Automated Land Cover Mapping Using an Ensemble of Support Vector Machines with Moderate Resolution Imagery Integrated into a Custom Decision Support Tool*. Alexandria, VA: United States Army Corps of Engineers Engineering Research and Development Center.

Latifovic, R., D. Pouliot, and I. Olthof. 2017. "Circa 2010 Land Cover of Canada: Local Optimization Methodology and Product Development." *Remote Sensing* 9: 1098.

Malinowski, R., S. Lewinski,M. Rybicki, E. Gromny, M. Jenerowicz, M. Krupinski, A. Nowakowski, et al. 2020. "Automated Production of a Land Cover/Land Use Map of Europe Based on Sentinel-2 Imagery." *Remote Sensing* 12 (21).

Mohd, H. I., H. Pakhriazad, and M. Shahrin. 2009. "Evaluating Supervised and Unsupervised Techniques for Land Cover Mapping Using Remote Sensing Data." *Geografia : Malaysian Journal of Society and Space* 5 (1): 1–10.

National Geospatial-Intelligence Agency. 2010. "Visual Navigation Land Cover / Land Use Dataset." Retrieved from https://data.amerigeoss.org/dataset/land-cover-30m-nga.

Pedregos, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12 (85): 2825–2830.

Tobler, W. 1970. "A Computer Movie Simulating Urban Growth in the Detroit Region." *Economic Geography* 46 (1): 234–240.

van Loon, A. F. 2015. "Hydrological Drought Explained." *WIREs Water* 2 (4): 359–392.

Wickham, J., S. Stehman, D. G. Sorenson, L. Gass, and J. A. Dewitz. 2021. "Thematic Accuracy Assessment of the NLCD 2016 Land Cover for the Conterminous United States." *Remote Sensing of Environment* 257 (112357).

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
10-2022	Final Technical Report	

4. TITLE AND SUBTITLE

Snow-Covered Region Improvements to a Support Vector Machine-Based Semi-Automated Land Cover Mapping Decision Support Tool

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
633463

6. AUTHOR(S)

Francis O'Neill, Kristofer Lasko, Elena Sava

5d. PROJECT NUMBER
AU1

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

ERDC Geospatial Research Lab
7701 Telegraph Rd
Alexandria, VA 22315

8. PERFORMING ORGANIZATION REPORT NUMBER

ERDC/GRL-TR-22-3

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Headquarters
U.S. Army Corps of Engineers
Washington, DC 20314-1000

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This work builds on the original semi-automated land cover mapping algorithm and quantifies improvements to class accuracy, analyzes the results, and conducts a more in-depth accuracy assessment in conjunction with test sites and the National Land Cover Database (NLCD). This algorithm uses support vector machines trained on data collected across the continental United States to generate a pre-trained model for inclusion into a decision support tool within ArcGIS Pro. Version 2 includes an additional snow cover class and accounts for snow cover effects within the other land cover classes. Overall accuracy across the continental United States for Version 2 is 75% on snow-covered pixels and 69% on snow-free pixels, versus 16% and 66% for Version 1. However, combining the "crop" and "low vegetation" classes improves these values to 86% for snow and 83% for snow-free, compared to 19% and 83% for Version 1. This merging is justified by their spectral similarity, the difference between crop and low vegetation falling closer to land use than land cover. The Version 2 tool is built into a Python-based ArcGIS toolbox, allowing users to leverage the pre-trained model—along with image splitting and parallel processing techniques—for their land cover type map generation needs.

15. SUBJECT TERMS
Land cover--Remote sensing; Remote-sensing images; Cold regions; Geospatial data; Machine learning; algorithms; Geographic information systems

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)
Unclassified	Unclassified	Unclassified	None	46	