Land Cover Mapping in Limited Labels Scenario: A Survey

Rahul Ghosh[1] , **Xiaowei Jia**[2] , **Vipin Kumar**[1]
[1]University of Minnesota
[2]University of Pittsburgh
{ghosh128, kumar001}@umn.edu, xiaowei@pitt.edu

arXiv:2103.02429v2 [cs.LG] 14 Sep 2021

Abstract

Land cover mapping is essential for monitoring global environmental change and managing natural resources. Unfortunately, traditional classification models are plagued by limited training data available in existing land cover products and data heterogeneity over space and time. In this survey, we provide a structured and comprehensive overview of challenges in land cover mapping and machine learning methods used to address these problems. We also discuss the gaps and opportunities that exist for advancing research in this promising direction.

1 Introduction

Global demand for land resources to support human livelihoods and well-being through food, fiber, energy and living space will continue to grow in response to the population expansion and socioeconomic development. This poses a great challenge to the human society, given the increasing competition for land from the need to maintain other essential ecosystem services. Addressing this challenge will require timely information on land use and land cover (LULC) changes, e.g., the conversion of forest to farmland or plantations, the loss of productive cropland due to urbanization, and the degradation of soil due to inappropriate management practices.

Advances in Earth observation technologies have led to the acquisition of vast amounts of accurate, timely, and reliable remote sensing (RS) data that can be used for monitoring changes on the Earth system at a global scale. These rich datasets provide tremendous potential to study the nature and impact of global changes due to natural processes and human activities that shape the physical landscape and environmental quality of our planet.

Supervised classification methods, especially recent deep learning approaches, have achieved significant success in commercial applications in Natural Language Processing (NLP) and Computer Vision (CV) domain, where large training data is available. Given the promise provided by these approaches, there is a growing interest for using these techniques in automated land cover mapping at large scale through the analysis of RS data [Zhu *et al.*, 2017]. However, off-the-shelf supervised classification methods have limited utility due to challenges that are unique to environmental applications. Supervised machine learning algorithms, e.g., advanced deep neural networks, require sufficient labeled training instances which are representative of the test data. Such training data is often scarce in land cover applications given high manual labor and material cost required in manual labeling (e.g., visual inspection) and field study. This is further exacerbated by the high-dimensional nature of spatio-temporal remote sensing data. Moreover, land covers commonly show much variability across space and time, e.g., the same crop can look different in different years and in different regions due to variability in weather conditions and farming practice. Additionally, the availability of multiple RS data sources acquired at different spatial and temporal resolutions, and other heterogeneous data, e.g., elevation, thermal anomalies, and night-time light intensity, provides unique algorithmic challenges that need to be addressed. The goal of this survey is to summarize recent developments in ML methods that are being used to address these challenges in the context of LULC. We focus on scenarios with a broad range of limited labels, e.g., scarce labels or imperfect labels, which also makes it challenging to model the heterogeneity, the large number of classes, and the high-dimensional spatio-temporal data. Additionally, we discuss the gaps and opportunities that exist for advancing research in this promising direction.

We organize the paper as follows. Section 2 describes different challenges that can be achieved using latest ML techniques for learning with limited data. Then, Section 3 discusses ML methods that researchers are using to address these challenges. Table 1 categorizes the work surveyed in this paper by challenge and approach. We hope that this survey will also be valuable for domain scientists who are interested in exploring the use of ML to enhance modeling in mapping target land covers. Although the focus of this survey is largely on land cover application, the ideas carry over to a wide range of spatio-temporal data.

2 Challenges

In this section, we provide an overview of a diverse set of challenges that have been encountered in land cover mapping.

2.1 Heterogeneity in space and time

One major challenge in land cover mapping is the presence of a rich variety in the land cover characteristics in different

regions and at different times, i.e., data heterogeneity. For example, crops can be planted under different soil types, precipitation, and other weather conditions. Farmers also use different land management practices (e.g., conservation tillage, low phosphorous), future crop rotation plan, and local social exchanges. All these factors can result in the heterogeneity over different places and different years.

Traditional ML methods that learn a single model of target land covers over the entire population of training data fail to account for this heterogeneity of the data in space and time, thus resulting in poor prediction performance when applied across diverse regions and times. Additionally, supervised learning algorithms require sufficient availability of labeled instances for training, which are expected to be adequately representative of the test data. However, training labels are often only available for specific regions and years due to the high associated costs of obtaining gold-standard labels. This results in poor prediction performance when the learned model is generalized to a larger region or different years.

2.2 Large number of classes

Land cover products commonly involve a large number of land cover classes. For example, USDA Crop Data Layer [CDL, 2021] has defined over 150 land cover classes in the area of the US. Accurately mapping all these land covers requires the access to training data of all the classes. However, some land cover classes can be very rare. Such scenario is very common in the fine-grained land cover mapping, e.g., crop classification, where we may not have labeled instances of many such rare crops. Moreover, training labels are often only available for specific regions and years. Hence, unseen/novel land cover classes may appear when the trained model is applied to other regions or other years. Traditional ML models are designed to make classification over a fixed set of classes and cannot identify new classes that are not included in the training set.

2.3 Paucity of labeled data

Most land cover labels are created by domain experts through visual inspection or field study. Given the substantial human labor and material cost required in the labeling process, it is often difficult to obtain large training data for land cover mapping. This poses a serious challenge for effectively training advanced machine learning models. Recent deep learning models, that have shown high performance in a variety of commercial applications, usually involve a large number of parameters and thus require large labeled training data.

This issue can be further exacerbated by the severe skewness of land covers. In particular, there can be much fewer samples for certain minority land cover classes compared with dominant classes such as forest, grassland and barren land. Although some minority classes have very few amount of training samples, detecting these classes can be critical for understanding the impact of human practice and land cover changes to the environment. For example, only small number of farmers have planted cover crops in the winter time so it is challenging to train a cover crop classifier using such limited data. However, planting cover crops is highly encouraged by the government since they can hep maintain the soil nutrient and prevent water contamination. Hence, the capacity to monitor the extent of cover crops is urgently needed.

2.4 Imperfect labels

Training labels provided in many existing land cover products are noisy due to a number of reasons. First, inconsistent standard used in the annotation process. In a prior study for plantations, [Petersen *et al.*, 2016] divided Landsat images into a grid of 20 × 20 km and visually scanned each gridded structure for multiple types of plantations with the assistance of forest gain and loss information. Such an annotation process results in less than 80% precision according to their independent assessment. Moreover, when applied to map larger regions, this approach may require multiple observers to delineate tree crop plantations, and observers are likely to be inconsistent with one another. Second, many labeling information is present at a much coarser scale, e.g., in the form of bounding boxes (approximate location), image-level labels (whether the input image contains objects), annotated polygons over the region of interest. Furthermore, many land cover products only provide labels for objects of interest, and the negative samples are left unmarked. Lastly, although some proxy labels are available, e.g., active fire for wild fire mapping and night-time light intensity for urban area mapping, they often either over-estimate or under-estimate target land covers.

2.5 High-dimensional spatio-temporal data

Many RS datasets contain reflectance values from multiple spectral bands and thus are high dimensional. Given the power of deep learning models in automatically extracting hidden features, researchers have been using these models in learning new data representations from RS data. Research along this direction is of great importance since it may provide new opportunities to understand the relation between spectral features and land covers. In particular, deep learning models have the promise to extract new vegetation indices from multi-spectral reflectance values. Such indices can be potentially more robust and representative compared with existing vegetation indices, such as Normalized Difference Vegetation Index (NDVI), which are defined based on domain knowledge. Moreover, the learned representation, together with the latest interpretation methods, can provide useful insights about the change of spectral data in response to different land covers.

Besides interpretability, the obtained representations can also be used for a variety of downstream applications. For example, these representations can be used to measure the similarity amongst different land covers. Also, one can gather representations learned from different regions and use their distance to better quantify spatial heterogeneity.

3 ML Methods

3.1 Zero-shot Learning

Zero-shot learning (ZSL) methods have been proposed for classification in a special scenario, where we consider two sets of classes – *seen classes* ($\mathcal{S}$), for which labeled instances

are present, and *unseen classes* ($\mathcal{U}$), for which there are no labeled instances. These two set of classes are disjoint and the aim is to learn a classifier that can classify the instances in the testing set into classes belonging to *unseen classes*. To capture the relationship between seen and unseen classes, ZSL methods create a real-valued semantic space where each class $Y \in \mathcal{S} \cup \mathcal{U}$ (both seen and unseen) has a corresponding vector representation, which is also referred to as the class prototype. To create this semantic space, some form of auxiliary information is often necessary to describe the relation amongst classes in $\mathcal{S} \cup \mathcal{U}$. The class prototype along with the training data are then used to learn the zero-shot classifier. ZSL methods have been used in several land cover studies given their ability to automatically classify rare or unseen classes. Existing ZSL methods can be categorized based on the type of semantic spaces and the methods used to learn the zero-shot classifier [Wang *et al.*, 2019].

Recent literature has demonstrated the capacity of ZSL in levering additional text data to extract semantic space for unseen land cover classes [Li *et al.*, 2017a; Gui *et al.*, 2018; Sumbul *et al.*, 2017; Pradhan *et al.*, 2020]. For example, [Li *et al.*, 2017a] introduced the idea of zero-shot scene classification using high spectral resolution (HSR) images. They first constructed a semantic space for embedding class prototype using the word2vec model [Mikolov *et al.*, 2013] trained on Wikipedia. Then they further created a semantic directed graph over both seen and unseen classes and assign the edge weight as the distance between each pair classes based on their class prototypes. Once establishing the relationship between the seen and unseen, they fine-tune a GoogLeNet model [Szegedy *et al.*, 2015] on the training data. The combination of GoogLeNet and the semantic graph enables producing the probability of a new instance belonging to each of seen or unseen classes. In another work, [Sumbul *et al.*, 2017] extended ZSL methods for fine-grained tree classification by using a combination of label embeddings learned from text data via word2vec, visual embeddings of trees, and hierarchy embeddings based on scientific classification of tree species. Such combination allows creating a semantic space that leverage the benefits of both domain knowledge and learning based methods.

An extension of this idea in the context of land-cover mapping is presented in [Jia *et al.*, 2017]. Instead of using auxiliary text data, authors in this work create the semantic space using sequences of land cover changes collected from other regions in the world. They have shown the effectiveness of this method in identifying deforestation and forest fires without using extra text data. Additionally, [Chen *et al.*, 2019] used the human-defined attributes as semantic information for knowledge transfer between seen classes and unseen classes and used a projection method to learn the zero-shot classifier. Specifically, they projected image features and label attributes to the same space and train a zero-shot classifier, which can then classify the unseen classes by finding the closest class-attribute in the projected space.

Although notable developments of ZSL models have been published for various CV tasks, their applications to RS are still at a nascent stage with the focus on small regions and small sets of target land covers. Moreover, most of existing methods create label-semantic space using the word2vec model, which was originally designed for text embedding and thus are not able to fully explore land cover characteristics. Hence, there is a need for new mechanism to create a semantic space that captures the properties of the classes and also includes the domain knowledge of the hierarchy of these classes. One potential method is the combination-based method of ZSL, where the idea is to construct the classifier for unseen classes by the combination of classifiers for basic elements that are used to constitute the classes. For example, in crop classification, one can create individual models to learn the basic elements during a crop cycle like sowing of seeds, growing season, harvesting period, texture and combining them for building the semantic space guided by the domain knowledge.

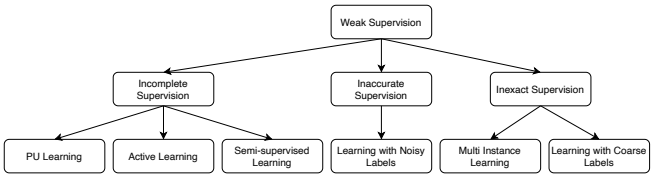

Figure 1: Weak supervision taxonomy.

3.2 Weakly supervised Learning

Weakly supervised learning (WSL) is an area of machine learning covering a variety of studies that attempt to construct predictive models by learning with different forms of weak supervision. In particular, there are three types of weak supervision [Zhou, 2018]: incomplete, inexact and inaccurate. An illustration of these three types along with the learning methodologies used in each of them is shown in Fig. 1. Several research works have demonstrated the benefits of such techniques on RS data for handling the challenges posed by weak supervision.

ML methods that handle incomplete supervision, e.g., only target land covers are labeled, falls under the Positive-Unlabeled (PU) Learning family [Bekker *et al.*, 2020]. [Wei *et al.*, 2020] proposed a scribble based deep learning method for road surface extraction. Specifically, using the scribble based PU road labels from crowd sources such as GPS or Open Street Maps (OSM) [Haklay *et al.*, 2008], negative samples (non-road pixels) are identified using buffers around the scribbles. These non-road pixels along with the scribble based road labels are used to train a semantic segmentation model. Another scenario of incomplete supervision arises when along with a few labeled samples, there is exists large unlabeled data. Semi-supervised learning [Van Engelen *et al.*, 2020] methods are typically used to handle this scenario. [Han *et al.*, 2020b] adopted the Mean Teacher model [Tarvainen *et al.*, 2017] which is a popular semi-supervised technique to train their backbone ResNet [He *et al.*, 2016] model using a few manually annotated samples and unlabeled samples for economic scales prediction.

The availability of large-scale crowd sourced information offers opportunities for building ML models for land cover mapping tasks for which expert-annotated labels are not available. However, crowd sourced data are often plagued by noise and thus require the ability to deal with inaccu-

rate supervision. [Sheehan *et al.*, 2018] showed results on classifying satellite images using weak labels extracted from Wikipedia. Similarly, [Sun *et al.*, 2019] used crowd sourced GPS data for road extraction from aerial imagery. Apart from crowd sourced labels, other satellite products such as night-time lights also form a noisy proxy for several tasks such as urbanization and economic development. In their study, [Han *et al.*, 2020a] showed the usability of night-time lights as a data rich proxy for identifying economic development from satellite imagery.

For pixel-wise land cover mapping tasks, coarse-level labels or image-level labels are often much easier to collect than pixel-level labels. Several methods have been proposed to leverage coarse-level labels for pixel-wise land cover mapping [Zhang *et al.*, 2019; Schmitt *et al.*, 2020; Wang *et al.*, 2020]. In particular, [Schmitt *et al.*, 2020] discussed the challenges and opportunities of using coarse-level and image-level weak labels for land cover mapping. Many of these approaches are inspired by image segmentation approaches using image-level labels [Zhang *et al.*, 2016a; Fu *et al.*, 2018].

Although most of the existing works have focused on single source of weak supervision, combining independent sources of weak supervision in a multi-task framework can force the models to learn from each other. For instance, for mapping urban areas the different sources of weak supervision can be night-time lights (coarse labels) and crowd sourced OSM labels (noisy labels).

3.3 Self-supervised Learning

One common approach for dealing with limited availability of labeled datasets is to pre-train an ML model on existing large labeled datasets for a related problem, and then refine it using a small number of labeled samples for the problem of interest. For example, models for image recognition are first trained using large-scale datasets like ImageNet [Deng *et al.*, 2009] and then are fine-tuned on the limited-size dataset for the downstream task [Huh *et al.*, 2016]. However, such approaches cannot be used for RS due to two key reasons; a) Spectral bands available in RS satellite data are much richer than RGB data in ImageNet (e.g., Sentinel data has 12 bands); b) large scale labeled data for land covers are not available.

Self-supervised learning (SSL) [Jing *et al.*, 2020] is an alternative approach that learns feature representation from unlabeled images. Specifically, SSL methods aim to pre-train ML models using "pre-text" learning tasks related to target classes before fine-tuning ML models using small amount of observations. The central idea is to propose various pretext tasks for the network to solve, in the hope that the network will learn important feature representations related to specific problems, such as image classification, segmentation, and object detection. Some popular self-supervised methods in CV include inpainting patches [Pathak *et al.*, 2016], image colorization [Zhang *et al.*, 2016b], and solving image-jigsaw [Noroozi *et al.*, 2016].

Recently, researchers have started using SSL for land cover mapping. One popular approach is to pre-train the ML model by optimizing the reconstruction performance on unlabeled data. For example, [Kemker *et al.*, 2018] introduced the reconstruction loss for each hidden layer of a stacked convolutional autoencoder for unsupervised training of the model. Specifically, they used a symmetric encoder-decoder framework where the feed-forward encoder segments the satellite images and the decoder network reconstructs the original input back from the compressed representations of the encoder. This method has shown improved predictive performance on RS datasets that have small annotated training data.

Many existing works have also adapted pre-text tasks that are widely used in CV to RS applications and have shown success in improving the performance. For example, [Singh *et al.*, 2018] adopted the inpainting pretext task and showed its superiority over other autoencoder-based self-supervised methods in learning representation for segmenting land covers in several applications. In addition to the inpainting task, the authors propose to use an adversarial training scheme, where an adversarial network, called the coach network, is tasked to increase the difficulty of the inpainting task gradually so that the semantic inpainting model do not overfit to a single type of corruption. In [Tao *et al.*, 2020], the authors compared image inpainting, image-jigsaw and contrastive learning [Chen *et al.*, 2020] based pretext tasks for the RS land cover classification and present the results on three datasets. Similarly, [Vincenzi *et al.*, 2020] use an autoencoder to predict the RGB channels given the other channels of the data as input. During the target classification tasks, the model has two independent branches which takes in the spectral bands and the RGB channels, respectively.

Another line of thought for pretext task, which is commonly used in the NLP domain, is representation learning based on context-similarity. The central idea is that words in similar contexts should have similar representations [Mikolov *et al.*, 2013]. By redefining context as spatial neighborhoods, Tile2Vec [Jean *et al.*, 2019] used this idea in RS domain where they promote the tiles that are close together to have similar representations than the tiles that are far apart.

SSL can also be used to extract shared information across multiple data sources. For example, [Swope *et al.*, 2019] proposed a method that creates different views of a given area and uses the InfoNCE loss to enforce consistency between the two views. To create two views of a multispectral image, standard data-augmentation techniques such as channel dropout, random flips, rotation, cropping and jittering the brightness and contrast are applied randomly. The consistency is enforced between the representations from different levels of a ResNet model. The authors show that the representations learned in this manner outperform the fully supervised ImageNet weights. Moreover, this multiview framework allows the fusion of multiple sources of data for joint representation learning using multi-modal data sources.

While much of the current research has focused on adopting ideas that have been successful in CV and NLP, there is also huge promise for designing new pretext learning tasks specific to land cover characteristics. One promising direction is to use domain knowledge, e.g., crop phenology, to define pre-training objectives. For example, neural networks can be pre-trained to predict several phenology-related values, including the slope of the NDVI curve, the maximum NDVI value, etc.

Table 1: Table of literature classified by challenges(rows) and method(columns).

	ZSL	Weakly-supervised Learning	SSL	TL	Meta Learning	Feature Extraction
Heterogeneity in space and time	[Jia *et al.*, 2017]	[Tong *et al.*, 2020]		[Jia *et al.*, 2019a][Tong *et al.*, 2020][Islam *et al.*, 2019]		
Large number of classes	[Chen *et al.*, 2019] [Pradhan *et al.*, 2020] [Li *et al.*, 2017a] [Gui *et al.*, 2018] [Sumbul *et al.*, 2017]			[Deng *et al.*, 2020]		[Deng *et al.*, 2020]
Paucity of labeled data			[Xie *et al.*, 2018] [Tao *et al.*, 2020] [Vincenzi *et al.*, 2020] [Singh *et al.*, 2018] [Kemker *et al.*, 2018] [Romero *et al.*, 2015]	[Gadiraju *et al.*, 2020][Marmanis *et al.*, 2015][Scott *et al.*, 2017][Rostami *et al.*, 2019][Chen *et al.*, 2018] [Ayush *et al.*, 2020][Deng *et al.*, 2020]	[Zhang *et al.*, 2021][Rußwurm *et al.*, 2020][Li *et al.*, 2020]	[Xie *et al.*, 2018][Gadiraju *et al.*, 2020][Ayush *et al.*, 2020] [Deng *et al.*, 2020][Lu *et al.*, 2017][Romero *et al.*, 2015]
Imperfect labels		[Sun *et al.*, 2019][Wang *et al.*, 2020][Wei *et al.*, 2020][Qiao *et al.*, 2020][Zhang *et al.*, 2016a] [Zhang *et al.*, 2019][Schmitt *et al.*, 2020][Fu *et al.*, 2018] [Han *et al.*, 2020a]		[Xie *et al.*, 2016][Najjar *et al.*, 2017]		
High-dimensional spatio-temporal data	[Demertzis *et al.*, 2020]	[Han *et al.*, 2020b][Sheehan *et al.*, 2018]	[Xie *et al.*, 2018][Swope *et al.*, 2019][Jean *et al.*, 2019]	[Hu *et al.*, 2015][Gargees *et al.*, 2019][Jia *et al.*, 2019b]	[Demertzis *et al.*, 2020]	[Xie *et al.*, 2018][Gargees *et al.*, 2019][Demertzis *et al.*, 2020] [Albert *et al.*, 2017][Liu *et al.*, 2019][Yang *et al.*, 2020]

3.4 Transfer Learning

Transfer learning (TL) is a class of methods which aim to improve a ML model on a target domain by transferring information from a related source domain. When there is limited labeled training data for the target domain, the existence of large scale datasets in a separate but related source domain can also provides the opportunity to transfer the knowledge gained from the source domain to the target task. TL has been extensively used in ML domain where models trained on large-scale datasets such as ImageNet have been successfully used as a starting point for solving the target task.

Several studies have also shown the benefit of transferring the knowledge from large-scale CV datasets to RS data. In particular, many of these methods aim to deal with the paucity of labeled data [Marmanis *et al.*, 2015; Hu *et al.*, 2015; Najjar *et al.*, 2017; Chen *et al.*, 2018; Scott *et al.*, 2017]. For example, [Hu *et al.*, 2015] present a comprehensive comparison of several network models pre-trained on large-scale datasets like ImageNet and Places [Zhou *et al.*, 2014], and then fine-tuned on multiple RS benchmark datasets. Several other studies have shown the benefit of pre-training on ImageNet datasets in target tasks like road safety mapping [Najjar *et al.*, 2017], airplane detection [Chen *et al.*, 2018] and other land cover mapping datasets [Scott *et al.*, 2017].

Another approach to TL is to pretrain the ML model using data rich proxy labels. This approach has several benefits: 1. Proxy labels can be very specific to the target class labels of interest (e.g., night time lights can be a good proxy for infrastructure available or level of poverty); 2. Transfer can be done meaningfully as they are more closely related to the target task than ImageNet. For instance, [Xie *et al.*, 2016] pretrain a ResNet model first using ImageNet dataset followed by night-time lights. The authors show that this method can achieve good performance when fine-tuned on limited labels of poverty prediction. As another example, [Ayush *et al.*, 2020] used the publicly available xView [Lam *et al.*, 2018] dataset, which is one of the largest and most diverse publicly available overhead imagery datasets for object detection. They also showed the success of transferring the models learned using xView as the source dataset to the target task of poverty prediction.

Several studies have also explored domain adaptation techniques for knowledge transfer. Domain adaptation aims to extract feature representation that is invariant across source and target domains and thus can potentially alleviate data heterogeneity. [Rostami *et al.*, 2019] used a domain adaptation method to transfer the knowledge from Electro-optical(EO) domain, which has large number of labels available, to the synthetic aperture radar (SAR) image domain, where labelling can be difficult. They used two encoder networks to extract a shared invariant cross-domain embedding space such that the distributions of encoded EO and SAR data is minimized in this latent space. [Jia *et al.*, 2019a] utilized a cyclic generative adversarial network (GAN) model to perform the domain adaptation between source and target domains which enables them to address the heterogeneity in croplands and grasslands over different regions and different years. In similar lines, [Islam *et al.*, 2019] proposed a learning based domain adaptation technique using an adversarial loss to align marginal data distribution between a source region and a target region. Then a few labelled samples from the target domain is used to align class specific data distributions between the two domains, based on the contrastive semantic alignment loss.

Additionally, TL has been used to extract meaningful feature representations which can then be used for unsupervised grouping of land cover scenes. [Gargees *et al.*, 2019] used ImageNet and RSI-CB256 [Li *et al.*, 2017b] benchmark dataset to train a land cover classification model. This model is then used as a feature extractor and the extracted features from RS images can be used for clustering.

3.5 Meta Learning

Meta learning is a strategy that automatically learns how to adapt models quickly to different tasks. In meta learning, a machine learning model gains experience from a distribution

of related tasks and uses this experience to improve its future learning performance. Meta learning is performed in two levels: a) inner level where the model tries to solve a single task, and b) outer level where the meta model is updated using the experience gained by solving different tasks in the inner level. This enables the model to quickly learn a new task.

Several works have shown the benefit of using meta learning methods for handling the paucity of labeled data in land cover mapping [Zhang *et al.*, 2021; Rußwurm *et al.*, 2020]. [Rußwurm *et al.*, 2020] used model agnostic meta learning (MAML) for the problem of inductive transfer-learning, where the generalization is induced by a few labeled examples in the target domain. They show results on the Sen12MS dataset [Schmitt *et al.*, 2019] and the DeepGlobe challenge. [Li *et al.*, 2020] proposed a novel meta learning framework for few-shot RS scene classification. Specifically, they generated different tasks sampled from a task family and learn a metric space that is tuned to perform well on new tasks. [Zhang *et al.*, 2021] proposed a three stage meta learning based strategy for few-shot classification of aerial scene images. First, a feature extractor is trained using the base set and a standard cross-entropy loss. Second, the meta learning classifier is trained over a set of episodes in the meta-training stage while comparing the query features with the mean of support features. Finally, the meta learning classifier is tested on a set of novel set referred to as the meta-test stage. [Demertzis *et al.*, 2020] proposed Model-Agnostic Meta-Ensemble Zero-shot Learning (MAME-ZSL) which exploits the MAML++ algorithm for the detection of unseen classes in land cover mapping.

Majority of the research works above have been on the MAML direction of meta learning. The use of other meta learning algorithms such as metric-based meta learning is still rare, where the central idea is to automatically learn a similarity measure between two instances. If the similarity between the instances of target and source tasks can be successfully captured, it should lead in effective solving of the target tasks given its similarity to a particular source task. For instance, considering the identification of cashew plantations, corn crops and palm oil plantations as the (target, source_1 and source_2) task, the target task shares more similarity with source_2(both being plantations) than source_1. Thus, while identifying cashew plantations, the source task of identifying of palm oil plantation will be more beneficial. Another popular form of metric-based meta learning is prototypical networks [Snell *et al.*, 2017], where a prototype feature vector is defined for every class using instance embeddings. Here, there is an opportunity to incorporate domain knowledge about the classes to create such prototypical vectors.

3.6 Feature extraction methods

RS products often contain reflectance values from multiple spectral bands, e.g. MODIS has 7 spectral bands, Sentinel 2A has 13 spectral bands, and some hyper-spectral products have even more. The reflectance value on each spectral band can be sensitive to different atmosphere and ground conditions, e.g., clouds, leaf. Environmental scientists have previously created multiple indices, NDVI, Urban Index (UI), Normalized Difference Built-up Index (NDBI), based on their domain knowledge. These indices are essentially combinations of multiple spectral bands and have shown effectiveness in identifying specific land covers. However, the performance can be significantly degraded for distinguishing multiple similar land covers, e.g., multiple crops can have similar NDVI patterns if they are planted and harvested on close dates.

The advent of new data-driven feature extraction methods have provided opportunities to extract more robust features from multiple spectral bands. Moreover, the state-of-the-art deep learning structures, such as convolutional neural network and recurrent neural network, also enable incorporating spatial and temporal context into the data representation. The extracted feature representation using these advanced techniques have shown improved performance in several applications. For example, prior work [Deng *et al.*, 2020] has used a meta feature extractor based on Convolutional Neural Networks (CNNs) that extracts representative feature representations and enables learning with much fewer amounts of training samples. Moreover, such feature extraction can be also conducted using unlabeled data [Romero *et al.*, 2015].

Land cover changes are commonly driven by natural processes or human interventions, which have specific temporal and spatial patterns, e.g., deforestation continues to happen in nearby areas. This requires new feature extraction methods that incorporate both spatial and temporal relationships that are consistent to the change of weather conditions, human activities and their interactions. Additionally, feature extraction methods that can interpret such data dependencies will be highly valuable for better understanding land cover changes.

4 Discussions, Future Directions and Concluding Remarks

In this survey, we reviewed key challenges faced in land cover mapping and machine learning methods for addressing them. Given the advances in machine learning and deep learning, we anticipate these techniques to play an important role in future works on monitoring land cover changes. In Table 1, we summarized relevant literature on different challenges using different methods. Also note that many boxes in the Table 1 have no and very few entries, many of which represent opportunities for future work. For example, SSL has been widely used for pre-training ML models so that they need much fewer training samples for fine-tuning. There is also a great promise for using SSL for extracting new representations that can be used for downstream applications, such as clustering spatial regions.

Another promising direction is to combine RS with other data sources such as street-level images to jointly study land covers. This would require new multi-view and representation learning methods to handle data from different views and in different spatial and temporal resolutions [Jia *et al.*, 2019b]. Additionally, many environmental processes, e.g., the growth of crops in response to fertilizer, soil and water conditions, cannot be directly observed through RS data. Hence, it is important to combine RS data with underlying physical/bio-chemical processes. For example, the emerging research on physics-guided machine learning [Willard *et al.*, 2020] has provided huge potential on this direction.

5 Acknowledgement

This work was funded by the NSF awards 1838159 and 1739191. Rahul Ghosh is supported by UMII MNDrive Graduate Fellowship.

References

[Albert *et al.*, 2017] Adrian Albert et al. Using convolutional networks and satellite imagery to identify patterns in urban environments at a large scale. In *SIGKDD*, 2017.

[Ayush *et al.*, 2020] Kumar Ayush et al. Generating interpretable poverty maps using object detection in satellite images. *arXiv:2002.01612*, 2020.

[Bekker *et al.*, 2020] Jessa Bekker and Jesse Davis. Learning from positive and unlabeled data: A survey. *Machine Learning*, 2020.

[CDL, 2021] USDA cropland data layer, 2021. https://www.nass.usda.gov/Research_and_Science/Cropland/SARS1a.php.

[Chen *et al.*, 2018] Zhong Chen et al. End-to-end airplane detection using transfer learning in remote sensing images. *Remote Sensing*, 2018.

[Chen *et al.*, 2019] Hong Chen et al. Generalized zero-shot vehicle detection in remote sensing imagery via coarse-to-fine framework. In *IJCAI*, 2019.

[Chen *et al.*, 2020] Ting Chen et al. A simple framework for contrastive learning of visual representations. In *ICML*. PMLR, 2020.

[Demertzis *et al.*, 2020] Konstantinos Demertzis and Lazaros Iliadis. Geoai: A model-agnostic meta-ensemble zero-shot learning method for hyperspectral image analysis and classification. *Algorithms*, 2020.

[Deng *et al.*, 2009] Jia Deng et al. Imagenet: A large-scale hierarchical image database. In *CVPR*. Ieee, 2009.

[Deng *et al.*, 2020] Jingyu Deng et al. Few-shot object detection on remote sensing images. *arXiv:2006.07826*, 2020.

[Fu *et al.*, 2018] Kun Fu et al. Wsf-net: Weakly supervised feature-fusion network for binary segmentation in remote sensing image. *Remote Sensing*, 2018.

[Gadiraju *et al.*, 2020] Krishna Karthik Gadiraju et al. Comparative analysis of deep transfer learning performance on crop classification. In *SIGSPATIAL Workshops*, 2020.

[Gargees *et al.*, 2019] Rasha S Gargees and Grant J Scott. Deep feature clustering for remote sensing imagery land cover analysis. *GRSL*, 2019.

[Gui *et al.*, 2018] Rong Gui et al. A generalized zero-shot learning framework for polsar land cover classification. *Remote Sensing*, 2018.

[Haklay *et al.*, 2008] Mordechai Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE Pervasive computing*, 2008.

[Han *et al.*, 2020a] Sungwon Han et al. Learning to score economic development from satellite imagery. In *SIGKDD*, 2020.

[Han *et al.*, 2020b] Sungwon Han et al. Lightweight and robust representation of economic scales from satellite imagery. In *AAAI*, 2020.

[He *et al.*, 2016] Kaiming He et al. Deep residual learning for image recognition. In *CPVR*, 2016.

[Hu *et al.*, 2015] Fan Hu et al. Transferring deep convolutional neural networks for the scene classification of high-resolution remote sensing imagery. *Remote Sensing*, 2015.

[Huh *et al.*, 2016] Minyoung Huh et al. What makes imagenet good for transfer learning? *arXiv:1608.08614*, 2016.

[Islam *et al.*, 2019] Kazi Aminul Islam et al. Semi-supervised adversarial domain adaptation for seagrass detection in multispectral images. In *ICDM*. IEEE, 2019.

[Jean *et al.*, 2019] Neal Jean et al. Tile2vec: Unsupervised representation learning for spatially distributed data. In *AAAI*, 2019.

[Jia *et al.*, 2017] Xiaowei Jia et al. Incremental dual-memory lstm in land cover prediction. In *SIGKDD*, 2017.

[Jia *et al.*, 2019a] Xiaowei Jia et al. Classifying heterogeneous sequential data by cyclic domain adaptation: An application in land cover detection. In *SDM*. SIAM, 2019.

[Jia *et al.*, 2019b] Xiaowei Jia et al. Recurrent generative networks for multi-resolution satellite data: An application in cropland monitoring. In *IJCAI*, 2019.

[Jing *et al.*, 2020] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *TPAMI*, 2020.

[Kemker *et al.*, 2018] Ronald Kemker et al. Low-shot learning for the semantic segmentation of remote sensing imagery. *TGRS*, 2018.

[Lam *et al.*, 2018] Darius Lam et al. xview: Objects in context in overhead imagery. *arXiv:1802.07856*, 2018.

[Li *et al.*, 2017a] Aoxue Li et al. Zero-shot scene classification for high spatial resolution remote sensing images. *TGRS*, 2017.

[Li *et al.*, 2017b] Haifeng Li et al. Rsi-cb: A large scale remote sensing image classification benchmark via crowdsource data. *arXiv:1705.10450*, 2017.

[Li *et al.*, 2020] Haifeng Li et al. Rs-metanet: Deep meta metric learning for few-shot remote sensing scene classification. *arXiv:2009.13364*, 2020.

[Liu *et al.*, 2019] Xuning Liu et al. Siamese convolutional neural networks for remote sensing scene classification. *GRSL*, 2019.

[Lu *et al.*, 2017] Xiaoqiang Lu et al. Jm-net and cluster-svm for aerial scene classification. In *IJCAI*, 2017.

[Marmanis *et al.*, 2015] Dimitrios Marmanis et al. Deep learning earth observation classification using imagenet pretrained networks. *GRSL*, 2015.

[Mikolov *et al.*, 2013] Tomas Mikolov et al. Distributed representations of words and phrases and their compositionality. *NeurIPS*, 2013.

[Najjar *et al.*, 2017] Alameen Najjar et al. Combining satellite imagery and open data to map road safety. In *AAAI*, 2017.

[Noroozi *et al.*, 2016] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*. Springer, 2016.

[Pathak *et al.*, 2016] Deepak Pathak et al. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.

[Petersen *et al.*, 2016] Rachael Petersen et al. Mapping tree plantations with multispectral imagery: preliminary results for seven tropical countries. *WRI*, 2016.

[Pradhan *et al.*, 2020] Biswajeet Pradhan et al. Unseen land cover classification from high-resolution orthophotos using integration of zero-shot learning and convolutional neural networks. *Remote Sensing*, 2020.

[Qiao *et al.*, 2020] Rui Qiao et al. Simple weakly supervised deep learning pipeline for detecting individual red-attacked trees in vhr remote sensing images. *Remote Sensing Letters*, 2020.

[Romero *et al.*, 2015] Adriana Romero et al. Unsupervised deep feature extraction for remote sensing image classification. *TGRS*, 2015.

[Rostami *et al.*, 2019] Mohammad Rostami et al. Deep transfer learning for few-shot sar image classification. *Remote Sensing*, 2019.

[Rußwurm *et al.*, 2020] Marc Rußwurm et al. Meta-learning for few-shot land cover classification. In *CVPR Workshops*, 2020.

[Schmitt *et al.*, 2019] Michael Schmitt et al. Sen12ms–a curated dataset of georeferenced multi-spectral sentinel-1/2 imagery for deep learning and data fusion. *arXiv:1906.07789*, 2019.

[Schmitt *et al.*, 2020] Michael Schmitt et al. Weakly supervised semantic segmentation of satellite images for land cover mapping–challenges and opportunities. *arXiv:2002.08254*, 2020.

[Scott *et al.*, 2017] Grant J Scott et al. Training deep convolutional neural networks for land–cover classification of high-resolution imagery. *GRSL*, 2017.

[Sheehan *et al.*, 2018] Evan Sheehan et al. Learning to interpret satellite images using wikipedia. *arXiv*, 2018.

[Singh *et al.*, 2018] Suriya Singh et al. Self-supervised feature learning for semantic segmentation of overhead imagery. In *BMVC*, 2018.

[Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv:1703.05175*, 2017.

[Sumbul *et al.*, 2017] Gencer Sumbul et al. Fine-grained object recognition and zero-shot learning in remote sensing imagery. *TGRS*, 2017.

[Sun *et al.*, 2019] Tao Sun et al. Leveraging crowdsourced gps data for road extraction from aerial imagery. In *CVPR*, 2019.

[Swope *et al.*, 2019] Aidan M Swope et al. Representation learning for remote sensing: An unsupervised sensor fusion approach. 2019.

[Szegedy *et al.*, 2015] Christian Szegedy et al. Going deeper with convolutions. In *CVPR*, 2015.

[Tao *et al.*, 2020] Chao Tao et al. Remote sensing image scene classification with self-supervised paradigm under limited labeled samples. *GRSL*, 2020.

[Tarvainen *et al.*, 2017] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv:1703.01780*, 2017.

[Tong *et al.*, 2020] Xin-Yi Tong et al. Land-cover classification with high-resolution remote sensing images using transferable deep models. *RSE*, 2020.

[Van Engelen *et al.*, 2020] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 2020.

[Vincenzi *et al.*, 2020] Stefano Vincenzi et al. The color out of space: learning self-supervised representations for earth observation imagery. *arXiv:2006.12119*, 2020.

[Wang *et al.*, 2019] Wei Wang et al. A survey of zero-shot learning: Settings, methods, and applications. *TIST*, 2019.

[Wang *et al.*, 2020] Sherrie Wang et al. Weakly supervised deep learning for segmentation of remote sensing imagery. *Remote Sensing*, 2020.

[Wei *et al.*, 2020] Yao Wei et al. Scribble-based weakly supervised deep learning for road surface extraction from remote sensing images. *arXiv:2010.13106*, 2020.

[Willard *et al.*, 2020] Jared Willard et al. Integrating physics-based modeling with machine learning: A survey. *arXiv:2003.04919*, 2020.

[Xie *et al.*, 2016] Michael Xie et al. Transfer learning from deep features for remote sensing and poverty mapping. In *AAAI*, 2016.

[Xie *et al.*, 2018] Yiqun Xie et al. An unsupervised augmentation framework for deep learning based geospatial object detection: a summary of results. In *SIGSPATIAL*, 2018.

[Yang *et al.*, 2020] Xinghao Yang et al. Tensor canonical correlation analysis networks for multi-view remote sensing scene recognition. *TKDE*, 2020.

[Zhang *et al.*, 2016a] Fan Zhang et al. Weakly supervised learning based on coupled convolutional neural networks for aircraft detection. *TGRS*, 2016.

[Zhang *et al.*, 2016b] Richard Zhang et al. Colorful image colorization. In *ECCV*. Springer, 2016.

[Zhang *et al.*, 2019] Libao Zhang et al. Hierarchical weakly supervised learning for residential area semantic segmentation in remote sensing images. *GRSL*, 2019.

[Zhang *et al.*, 2021] Pei Zhang et al. Few-shot classification of aerial scene images via meta-learning. *Remote Sensing*, 2021.

[Zhou *et al.*, 2014] Bolei Zhou et al. Learning deep features for scene recognition using places database. 2014.

[Zhou, 2018] Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National science review*, 2018.

[Zhu *et al.*, 2017] Xiao Xiang Zhu et al. Deep learning in remote sensing: A comprehensive review and list of resources. *GRSM*, 2017.