Accounting for training data error in machine learning applied to Earth observations

Arthur Elmes*[1,2], Hamed Alemohammad[3], Ryan Avery[4], Kelly Caylor[4,5], Ronald Eastman[1], Lewis Fishgold[6], Mark A. Friedl[7], Meha Jain[8], Divyani Kohli[9], Juan Carlos Laso Bayas[10], Dalton Lunga[11], Jessica L. McCarty[12], Robert Gilmore Pontius Jr[1], Andrew B. Reinmann[13, 14], John Rogan[1], Lei Song[1], Hristiana Stoynova[13, 14], Su Ye[1], Zhuangfang-Fang Yi[15], Lyndon Estes*[1]

[1] Graduate School of Geography, Clark University, Worcester, MA 01610, USA; reastman@clarku.edu (R.E.); rpontius@clarku.edu (R.G.P.); jrogan@clarku.edu (J.R.); lsong@clarku.edu (L.S.); sye@clarku.edu (S.Y.); lestes@clarku.edu (L.E.)
[2] School for the Environment, University of Massachusetts Boston, Boston, MA 02125, USA; arthur.elmes@umb.edu
[3] Radiant Earth Foundation, San Francisco, CA, 94105, USA; hamed@radiant.earth
[4] Department of Geography, University of California, Santa Barbara, CA 93013, USA; ravery@ucsb.edu
[5] Bren School of Environmental Science and Management, University of California, Santa Barbara, CA 93013, USA; caylor@ucsb.edu
[6] Azavea, Inc., Philadelphia, PA 19123, USA; lfishgold@azavea.com
[7] Department of Earth and Environment, Boston University, Boston, MA 02215; friedl@bu.edu
[8] School for Environment and Sustainability, University of Michigan, 48109, USA; mehajain@umich.edu
[9] Faculty of Geo-Information Science & Earth Observation (ITC), University of Twente, 7514 AE Enschede, The Netherlands; d.kohli@utwente.nl
[10] Center for Earth Observation and Citizen Science, Ecosystems Services and Management Program, International Institute for Applied Systems Analysis (IIASA), Laxenburg, A-2361, Austria; lasobaya@iiasa.ac.at
[11] National Security Emerging Technologies, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA; lungadd@ornl.gov
[12] Department of Geography and Geospatial Analysis Center, Miami University, Oxford, OH 45056, USA; mccartjl@MiamiOH.edu
[13] Environmental Science Initiative, Advanced Science Research Center at the Graduate Center of the City University of New York (CUNY), New York, NY 10065, USA; Andrew.Reinmann@asrc.cuny.edu
[14] Department of Geography and Environmental Science, Hunter College, New York, NY 10065, USA
[15] Development Seed, Washington, DC 20001, USA; geospatialanalystyi@gmail.com
*Correspondence: arthur.elmes@umb.edu, Tel. 1-304-906-7946 (A.E.); lestes@clarku.edu, Tel. 1-202-431-0496 (L.E.)

Abstract

The role of remote sensing in understanding earth systems is growing rapidly, in part due to advances in new machine learning (ML) techniques. These approaches typically rely on large, spatially extensive training datasets to predict categories or continuous quantities. These training data are typically collected by digitizing polygons from high spatial resolution imagery, but also by *in situ* sampling, or by using pre-existing datasets. Training data are usually considered to represent the truth, but in practice almost always have error, stemming from 1) the design of the training data sample, and 2) mistakes made when collecting the sample. The latter is particularly relevant for image-interpreted training data, which is an increasingly common approach, due in part to the rise of citizen science or crowdsourcing campaigns. Training data errors can cause substantial errors in the maps created using ML algorithms, which propagate in derived 'downstream' products. Despite these potential errors and the real-world consequences they can have on map-based decisions, training data error is typically not explicitly accounted for or reported. We review the current practices for collection and treatment of training data for Earth Observation (EO) research, and identify the sources and impacts of training data error. We then illustrate these issues using several case studies that explore various aspects of training data creation and usage across a range of EO applications: infrastructure mapping, global surface flux estimates, and agricultural monitoring. We propose guidelines and recommendations for EO researchers regarding the treatment of training data. To harmonize potentially disparate terminology across disciplines, we recommend the following four terms to describe the data used in model creation and assessment: training data, validation data, training reference data, and map reference data. In general our advice is applicable to all of these classes of data, but our recommendations specifically focus on training data. Researchers should carefully consider the tolerable levels of map error and appropriate error metrics as a first step. Then, training data error should be minimized in research design and during individual data label creation, following procedures such as appropriate legend definition and imagery selection based on spatial resolution and spatio-temporal representativeness, using consensus-based labeling strategies, and by incorporating interactive feedback to interpreters regarding accuracy. We strongly advise that training error is incorporated in model outputs, either directly in bias and variance estimates or at minimum by documenting the sources and implications of error. Training data should be fully documented and should be made openly available via an open training repository, allowing others to replicate and assess its use. To guide researchers in this process, we propose three tiers of training data error accounting standards. Finally, we advise that researchers strive to clearly communicate the magnitude and impacts of training data error on map outputs, with specific consideration to the intended or likely audience and users of the map.

1. Introduction

Recent technological advancements have led to a new era in remote sensing, marked by rapid improvements in our ability to map and measure features on the Earth's surface such as Land Cover and Land Use (LCLU) [e.g. 1,2], vegetation cover and abundance [3], soil moisture [4], infrastructure [5,6], vegetation phenology [7–9], land surface temperature [10,11], and many others. The resulting data are being used by an expanding set of disciplines to gain new insights into socioeconomic and environmental dynamics, such as community-level poverty rates [12], fine-grained changes in surface water [13], and forest cover [14], and carbon accounting [15]. As such, remote sensing is increasingly shaping our understanding of how the world works, and how it is changing.

These breakthroughs are facilitated by several technological advances, particularly the increasing availability of moderate (5-30 m), high (1-5m, High Resolution, HR), and Very High Resolution (<1 m) VHR) imagery, as well as new machine learning (ML) algorithms that increasingly require large, high quality training datasets [e.g. 16,17–20]. The question of training data quality, in addition to training data quantity, is central to determining the ultimate quality of map products generated by ML algorithms, and therefore to downstream products based on those maps [21]. While progress in algorithmic performance continues apace, the standards concerning the collection and use of training data remain uncoordinated across researchers [22]. Additionally, much of the research and development of big data and ML is occurring in industry and in the fields of computer science and (non-spatial) data science, leaving a potential knowledge gap or lag for Earth Observation (EO) scientists [23].

The measurement and communication of map accuracy has been long studied and has become a mature topic in remote sensing and related fields, with a variety of metrics and approaches tailored to different data types, analyses, and user groups [24–32]. However, the impacts of training data error have not been assessed as thoroughly [22], even though such error is almost always present and directly impacts map accuracy [25,33–35]. Of particular concern is the propagation of errors downstream in further quantitative analyses relying on the map products [36]. These derived products are valuable, but potentially magnify errors originating in the initial calibration of the first map, illustrating the pernicious effects of under-reported uncertainty in training data [37]. In general, the problems associated with training data error can be summarized as follows:

1. The 'big data' era vastly increases the demand for human-produced training data.

2. Creation of of map products from EO data rely heavily on human-generated training data, which in most cases contain error, particularly when developed through image interpretation.

3. Uncertainty in these training data is rarely assessed or reported, and training data accuracy is often assumed to be perfect.

4. Errors in the training data may propagate to downstream products in surprising and potentially harmful ways (e.g. leading to bad decisions), and can occur without the map producer and/or map user's knowledge.

These problems spur the need for a comprehensive review of the issues of training data quality. While similar standards and practices exist for the collection and use of map reference data (also commonly called accuracy or validation data), we focus primarily on data used to train and validate ML models internally, prior to final accuracy assessment. We present a summary of current practices in the treatment of training data for categorical and continuous map generation, and identify the most common sources of error and inconsistency. We illustrate the impacts of uncertainty in training data creation within a series of case studies that span a range of typical remote sensing applications, including the mapping of building and road footprints, global surface flux estimates, and the characteristics of agricultural systems. Finally, we propose several recommendations and guidelines for best practices in training data collection and use, regarding definition of acceptable accuracy levels, minimization of error from design and collection standpoints, characterization and incorporation of training error for output maps, and communication of training data error in scientific and public documentation.

1.1 Current Trends in Training Data Collection

A large proportion of remote sensing projects make some use of training data, typically created either using geolocated *in situ* data [33,38], by visually interpreting high and/or very high resolution spatial resolution imagery [39–41], or by interpreting the images to be classified/modeled themselves [e.g. 42,43,44]. Of these collection methods, image interpretation is increasingly common [45], particularly with the rise in crowdsourcing initiatives [20,46]. As such, mapping is strongly constrained by the creation of training data, which are often seen to represent absolute 'truth' [47], if for no other reason than that their accuracy is assumed to be perfect [25,34,48]. However, multiple sources of error are possible and indeed likely in training data, whether collected *in situ* or via image interpretation [38].

The use of large, data-hungry ML algorithms continues to grow in many fields, including remote sensing. Artificial Neural Network (ANN) classification performance generally increases as a function of training sample size, with more complex ANNs, and particularly Convolutional Neural Networks (CNN), producing higher output accuracy [49]. However, such complex neural networks require substantially more training data than traditional statistical models, and like many ML approaches are sensitive to noisy and biased data, producing the logistical difficulty of creating very large, 'clean' training data sets [50–52].

Partially to address this need, a recent trend in large-scale mapping projects is to employ large teams of training data interpreters, often within citizen science campaigns that rely on web-based data creation tools [20,53–55]. Several recent efforts have been entirely devoted to the production of extremely large, open-ended training data sets without a specific mapping goal, but rather to serve as comprehensive benchmarks [56,57].

1.2 Characterizing Training Data Error

Due to different disciplinary lineages, terminology associated with the various datasets used to assess the accuracy of map algorithms is sometimes contradictory or distinct. Here we harmonize terminology by referring to four types of reference data: training, validation, training reference, and map reference. Due to disciplinary differences in terminology and usage, we propose the following definitions for these terms, as they will be used below. Training data, our primary focus, refers to a sample of observations, typically consisting of points or polygons, that relate image pixels to semantic labels. Validation data, typically a semi-independent random subset of training data, are withheld and used to fit ML model parameters and internally evaluate performance. Training reference data are expert-defined exemplar observations used to assess training data error during or after data creation. Map reference data consist of the completely independent set of observations used to assess final map accuracy; while these data may be collected simultaneously with the previous datasets, they must adhere to more stringent collection protocols, and must only be used for accuracy assessment for the final map product, rather than used iteratively in model or map improvement [47]. What we define as map reference data are also referred to as the test set in ML literature [58]; we avoid this term in an attempt to harmonize terminology across EO and ML research.

1.2.1 Map Accuracy Reporting Practices

To understand how training data errors can impact map accuracy, it is necessary to first review current practices and standards for measuring and reporting final map accuracy. While the emphasis of this paper is specifically on training data, as opposed to map reference data, it is necessary to review procedures for accuracy assessment. Sampling protocols for accuracy assessment are more stringent than those for collection of training data [47], but because both training and map reference data are often collected as part of a single campaign or using the same methods (e.g. Xiong et al, 2017), the stricter set of procedures should be followed for both. We therefore summarize several important features and best practices of error analysis.

Error analysis compares a mapped variable to a corresponding reference variable, which is typically derived from training data or some previously created standard product. Map reference data used for accuracy assessment are collected according to sampling and response designs that specify the probabilities of inclusion for each location, and the protocol for creating the labeled reference data, respectively [47,59]. Sample data used for training may also be derived from the

same sample, although this requires strict separation of the reference sample. The protocols for collection of training data specifically may be more purposive and targeted to classes of interest [47]. Critically, the assumption is that both training and map reference data have higher accuracy than the mapped data -- an assumption which is typically unexamined [22,60]. Reference data are often collected via visual image interpretation of high spatial resolution imagery, or by *in situ* field campaigns. Sampling design, whether simple random, stratified random, or systematic is dependent on application and *a priori* knowledge of the study area, and should be probability-based, such that the inclusion probability of each sample relates to the likelihood of that sample unit being included [26,47,61]. If the observations do not have equal probability of selection, then it is essential to convert the sample data to a confusion matrix that reflects an unbiased estimate for the entire population using methods summarized in Stehman and Foody [47].

Map accuracy is typically assessed using a metric or metrics designed to provide information regarding the correspondence of mapped and reference data. The objective of these metrics is to provide insights into the product's expected best use cases and potential shortcomings. Accuracy metrics vary according to whether the mapped variable is categorical or continuous, with each type of variable having its own foundation for error analysis [62–66]. The confusion matrix, i.e. a square contingency table, is the foundation for categorical variables. Conventionally, the table's rows provide mapped categories and the columns show the matching reference categories, with the diagonal entries showing agreement between the two. The confusion matrix is used to calculate user's accuracy (i.e. the complement of commission error intensity), producer's accuracy (i.e. the complement of omission error intensity), and overall accuracy (i.e. the complement of proportion error) [28]. More details on the interpretation of these values and other aspects of the error matrix are provided in several existing publications [24,26,47,62,67–69].

Beyond the aforementioned metrics, a number of other accuracy measures are also calculated from the error matrix. Most prominent among these is the Kappa Index of Agreement [70], which is widely used in the remote sensing and species distribution modelling literature. However, Kappa varies with class prevalence [71] and can be easily misinterpreted, thus its continued use is no longer recommended [27]. More recently, a number of additional metrics have started to be more commonly used for remote sensing accuracy analysis, in part due to contributions from other disciplines such as computer science and information science. Due to differing conventions and objectives within these disciplines, the metrics and terminology relating to error and accuracy are often quite different. To help resolve this confusion, we summarize these metrics and their meanings in Table 1.

A special and increasingly used type of categorical map is derived from object-based image analysis, in which the output map is classified into polygons representing discrete objects [72]. At present there is no commonly accepted standard for reporting the accuracy of such maps in

the remote sensing literature [45], since the optimal set of metrics for polygon accuracy assessment depends on the intended use of the categorical map. For example, edge similarity metrics are useful for assessing the segmentation of individual agricultural fields, whereas area based metrics will fail where multiple objects are frequently mapped as a single object [45]. In the fields of computer vision and deep learning, the mean Average Precision metric, otherwise known as the Jaccard Index, or Intersection over Union, is the primary benchmark for evaluating object based classification, or semantic segmentation, to use the computer vision terminology [73]. This and other area based metrics can be used in a remote sensing context, but for many mapping goals they should be complemented by other metrics that account for shape and edge similarity. Perhaps due to these complexities, many existing studies have assessed the accuracy of object-based maps using per-pixel accuracy assessments, which is problematic because it involves comparing fundamentally different spatial units [45].

The scatter plot, showing the mapped variable on the y-axis and the reference variable on the x-axis, is the foundation of error analysis for continuous variables. Since any point falling off the 1:1 line indicates deviation from a measurement of the true value, a visual assessment of the plot an intuitive first step for assessing error in the mapped variable. Several metrics are commonly used to quantify disagreement between mapped and reference variables, including mean deviation, Root Mean Square Error (RMSE; a.k.a. Root Mean Square Deviation, RMSD), and Mean Absolute Deviation (MAD), although RMSE may be inappropriate, and is frequently misinterpreted as the measurement of average error [74–76]. The Receiver Operating Characteristic (ROC) and the Total Operating Characteristic (TOC) enable analysis of continuous mapped variable relative to a binary reference variable, for example presence or absence [64,77,78]. The area under this curve (AUC) of an ROC/TOC plot is often used as a single measure of overall accuracy that summarizes numerous thresholds for the continuous variable [78].

Most of the metrics reported in Table 1 provide useful information for map users. However, they can only present information concerning final map quality, and do not convey potentially critical information pertaining to the earlier stages of map creation [81]. Critically, information regarding the collection, assessment, and use of training data in model calibration, is usually not communicated, whether by quantitative metrics or qualitative description.

Table 1: Summary of frequently used accuracy metrics.

Term	Information Content/Typical Usage	Description
Overall Accuracy	Summary metric combining all class accuracies into a single number	Proportion of correctly classified cases divided by the total of all classified cases
User's Accuracy (a.k.a. Precision)	Metric of the intensity of true positives given the classified category in which the true positives were 'found'. The intensity complement of commission error.	Proportion of correctly classified cases relative to the total number of cases classified into the given category
Kappa Index of Agreement	Single metric for overall accuracy	Used to measure the agreement between mapped and reference categories of a dataset while attempting to correct for agreement that occurs by chance.
Producer's Accuracy (a.k.a. Sensitivity, Recall)	Metric indicating the intensity of true positives given the reference category) The intensity complement of omission error.	True positive rate; ratio of correctly classified cases of a given class to the total true cases of that class
Specificity	Metric for commission error; indicates how well the model avoids false positives	True negative rate; ratio of correctly classified negatives to the sum of true negatives and false positives
True Skill Statistic [71]	Metric that combines sensitivity and specificity while accounting for class prevalence	Sensitivity + Specificity - 1
F1 [79,80]	Combined metric of commission and omission error	Equally weighted harmonic mean of precision and recall
Bias (Mean Bias Error)	Quantifies the average difference between predicted and reference variables	The average error, representing the systematic over- or under-prediction of a continuous variable
Root Mean Square Error/Deviation	Measures a combination of the average error and the variability within the distribution of errors	A potentially misleading metric used to measure disagreement between predicted and reference continuous variables
Mean Absolute Deviation	Measures how far points are from Y=X line	Recommended metric to measure disagreement between predicted and reference continuous variables

1.2.2 Accuracy Reporting Practices for Training Data

While substantial effort has been devoted to developing best practices for collecting map reference data for accuracy assessment, a wide diversity of approaches are currently used to create, analyze, and exploit training data [22,45]. Furthermore, while accuracy assessment is a mature topic in the remote sensing literature [26,27,32,47,82], the general topic of training data quality and uncertainty is under-appreciated and requires attention. Specifically, the degree to

which errors in training data impact map error is rarely quantified, and generally not incorporated into final accuracy estimates, despite the fact that response design has an inherent influence on map quality [33]. For example, Swan et al. [83] studied the impact of analyst-generated errors in a building footprint training dataset, finding that increased training data noise substantially reduced model accuracy as measured by recall and precision (see Table 1). Recent research regarding the impact of interpreter uncertainty on map output accuracy has specifically begun to quantify variability in model results arising from training data errors, by assessing the accuracy of image-interpreted training data relative to *in situ* observations [38], and by measuring between-interpreter variability [33,84,85]. McRoberts et al. [33] proposed a method that factors in both sampling and interpreter errors in the training data in assessing overall map accuracy, which produces unbiased variance estimates for mapped variables. In general, scientific literature on training data error recommends that multiple interpreters create training polygons and labels for the same locations, such that aggregation procedures can be used, and allows for comparison to expert defined training reference observations [47]. These methods facilitate a more complete accounting for error variance [22,33,84,85].

In general, however, training data quality, or even a thorough description of training data collection protocols, is typically underreported in peer reviewed publications. For example, Ye et al. [45] found that of 209 journal articles specifically concerning object-based image analysis, roughly one third "failed to give sufficient information concerning accuracy methodology such as sampling scheme and sampling size". Additionally, the authors found that 42% of the studies relied on image interpretation to create training and map reference data, with 10% using existing thematic maps, 26% using field collected data, and 15% using a mixed approach (7% contained no information). Similarly, the results of a keyword-search-based review of 30 recent, top ranked, original research papers on the topic of land cover mapping, conducted on Google Scholar[1], revealed that only 2 make explicit, systematic reviews of training data quality used to perform classification, while 16 made no mention of training data standards at all. Over 75% of these studies used image interpretation, as opposed to *in situ* data, in either training, accuracy assessment, or both. Some of these papers (~25%) used unsupervised classifiers in the processing chain to outline training areas, followed by image interpretation to assign labels to the polygons/pixels. This lack of detail creates substantial uncertainty as to how the training data were determined, reviewed, and incorporated. Taken as a whole, it is evident that current practices regarding the collection and presentation of training data are insufficient for the purposes of systematic, reproducible science.

---

[1] Publication date from 2008 to 2019. Keywords used: remote sensing; land cover, land use, LCLU; original research articles only.

2. Sources of Training Data Error

Training data error falls within two general classes: 1) errors stemming from the design of the training sample, including selection of the data source in the case of image interpretation, and 2) errors made during the process of training data creation, i.e. the process of digitizing and labeling points or polygons, or collection of field measurements. Below we discuss training data errors arising from each of these classes, bearing in mind that the two classes are non-exclusive.

2.1 Design-related errors

Training data error may also arise due to sampling design and data availability, caused by inadequate representativeness, data gaps, and spatiotemporal mis-matches, even with no error from image interpretation *per se*. Data gaps are particularly problematic due to the prevalence of image interpretation as a source of training data collection, and arise when spatially inaccurate or temporally obsolete data are used to generate training polygons and labels [26,39]. An example of how image age can lead to training data error is illustrated in Figure 1, which contrasts smallholder fields that are clearly visible in a satellite base map (Bing maps) with ground data collected in 2018. Center pivot fields were installed after the base map imagery was collected, but before ground data collection, creating a temporal mis-match between the base map and the *in situ* data. Labels generated from the base map, a common practice, in this case would introduce substantial error into an ML algorithm classifying more recent imagery. New HR satellites that have more frequent acquisitions [e.g. PlanetScope, 86] can help minimize such temporal offset problems for projects that are designed to map present-day conditions (e.g. 2018 land cover), but do not solve the problem for maps of periods before such imagery was collected (i.e. before 2016). The same can be said for aerial and Unmanned Aerial Vehicle acquisitions, which are typically limited in geographic and temporal extent [87]. While hardcopy historical copy maps can help supplement data gaps, these data sources come with their own problems, such as errors introduced during scanning and co-registration, and unknown production standards and uncertainties.

Spatial alignment and spatiotemporal coverage can be particularly problematic with HR and VHR commercial satellite imagery, which have narrow swath widths, and are often on-demand, which results in substantially off-nadir view angles, and may leave large gaps in the data record for certain places [39]. Due to large view angles and the lack of adequate digital elevation models, overlapping imagery, or other relevant control points, HR/VHR imagery typically does not meet the same orthorectification and geopositioning standards as coarser resolution, government operated satellites [88–90]. When integrating HR/VHR imagery acquired at different azimuth and elevation angles, features such as building roofs show apparent offset in much the same way as topographical effects; this is particularly problematic when a) mapping

the same type of features over a time series, and/or b) when using an existing vector dataset such as OSM as training data [91–93].

Figure 1. An example of potential training data error that can arise when image interpretation is conducted on older imagery. The underlying imagery is from Bing maps, which shows smallholder agricultural fields near Kulpawn, Ghana. The white polygons were collected by a team of mappers (hired by Meridia) on the ground using handheld GPS in 2018. The smallholder fields were replaced by larger center-pivot irrigation fields sometime after the imagery in the base map was collected.

These artifacts are particularly problematic for large crowdsourcing campaigns in which participants interpret HR/VHR satellite imagery potentially spanning multiple sensors and dates, which is constrained by image availability [39], quality (e.g. cloud cover percentage and atmospheric correction), view/sun angle [94], or image acquisition date. Thus, although much larger and cheaper training datasets are becoming a prized data source for mapping projects, the quality and representativeness issues in such volunteer-generated datasets can be cumbersome to find and clean, and represent important sources of training data error [47,95,96]. In many cases, design-based training data error in crowdsourced data sets may propagate through ML algorithms and map outputs.

Training data are often collected by image interpretation of HR/VHR imagery coregistered with the coarser resolution data used as ML model data. This creates a spatial resolution conflict because often the relationship between image objects and pixel size is fundamentally different, since objects that are clearly delineated as spectrally homogenous areas in HR/VHR imagery are part of mixed pixels in moderate or coarse resolution model imagery. This mis-match, which introduces error from the design standpoint, is similar to the concept of H-resolution versus L-resolution scene models proposed by Strahler et al. [97]; in H-resolution models, the objects of

interest are substantially larger than the pixel size, and vice versa for L-resolution models. The incorporation of mixed pixels may degrade classification model performance, or at least introduce undesired class spectral variability [98–100].

Similar spatial resolution and scaling issues must be dealt with in combining *in situ* measurements with satellite observations for continuous variables. Field collected data often cannot practically cover the entire area of a pixel in the model data, especially for moderate or coarse resolution imagery, and therefore potentially induce scaling errors based on the modifiable areal unit problem [101,102]. Spatial representativeness assessments and interpolation methods are used to limit this problem to the for operational EO science products [103–106], but nevertheless this issue is likely to be a source of training data error for *in situ* collection of training data.

A different design-based problem stems from the practice of iterative algorithm refinement using sample observations as both training and validation data. This problem can best be understood within the context of cross validation, which is an often used and appropriate method for ML parameter selection [e.g. 107]. However, any sample data used as internal validation data, when more than one iteration is performed, have lost their independence and can no longer be treated as map reference data, but must instead are internal validation data. This issue may be less appreciated within the context of approaches such as active learning, in which additional training sites are collected to increase the representativeness of the sample [42,e.g. 108]. Therefore, when discussing truth data terminology, it is important to distinguish between training data, validation data, and map reference data, the latter of which must only be used after algorithm development and map creation have been completed.

2.2 Collection-related errors

There are several common forms of error associated with the collection of training data. The first of these are errors of interpretation, which are mistakes created in the process of manual image interpretation. Manual interpretation of images is a primary source of training data, and often does not provide consistent training labels between interpreters [22,24,84]. Image interpretation may be conducted by people who lack experience in the task, or are unfamiliar with the context of the study area. In an uncommonly thorough analysis of training data error from image interpretation, Powell et al. [84] showed that inter-interpreter agreement ranged from 46% to 92%, depending on land cover, with an overall agreement of 86%. This research, which relied on trained image interpreters, concluded that transitional land cover classes produce substantial disagreement during training data creation; this is particularly problematic since much of the focus of land cover mapping is directed towards change detection. Another image interpretation approach based on an online platform with built-in accuracy assessment relying on pre-determined training reference data [109] shows that average interpreter score in digitizing crop

field boundaries is around ~80%, which holds true whether the interpreter mapped several hundred sites or < 50 (Figure 2). These results suggest that increased image interpretation experience does not eliminate labelling error, even when analysts are highly seasoned as in Powell et al. [84]. These findings underscore the need for improved and widely accepted standards for training data creation and vetting using predefined training reference data or inter-interpreter comparisons [33,38,84,85,110].


Figure 2: Number of sites mapped per worker versus the average score received at reference sites, where workers' maps were compared to reference maps using a built-in accuracy assessment protocol within a crowdsourcing platform for collect cropland data (Estes et al., 2016).

Semantic class definitions can also be inadequate or poorly communicated, leading to labeling error [111,112], particularly during discrimination of land use, as opposed to land cover [113]. This is especially evident in urban environments, which are typically spatially and spectrally heterogeneous even when imaged at high spatial resolution [114]. Critically, such scenes are also semantically vague, even when viewed from ground level. For example, Figure 3 shows a typical example of training data collection for mapping informal settlements (a.k.a slums), in Nairobi, Kenya, in which each of several trained interpreters delineate the same area [115]. Because slums may be defined by sociodemographic factors in addition to spatial and spectral properties, training data creation for such areas is prone to error stemming from semantic issues [112]. Additionally, complex classes such as slums may exhibit high variability between study areas, as local idiosyncrasies link the definition of slums to different physical, remotely observable characteristics. In general, these results illustrate the critical importance of training data error analysis and consensus mapping for HR/VHR image interpretation, particularly for heterogeneous target classes with vague or regionally idiosyncratic semantic definitions. Similar to other target LCLU classes, the generalizability of informal settlement mapping may be limited, even with local knowledge.
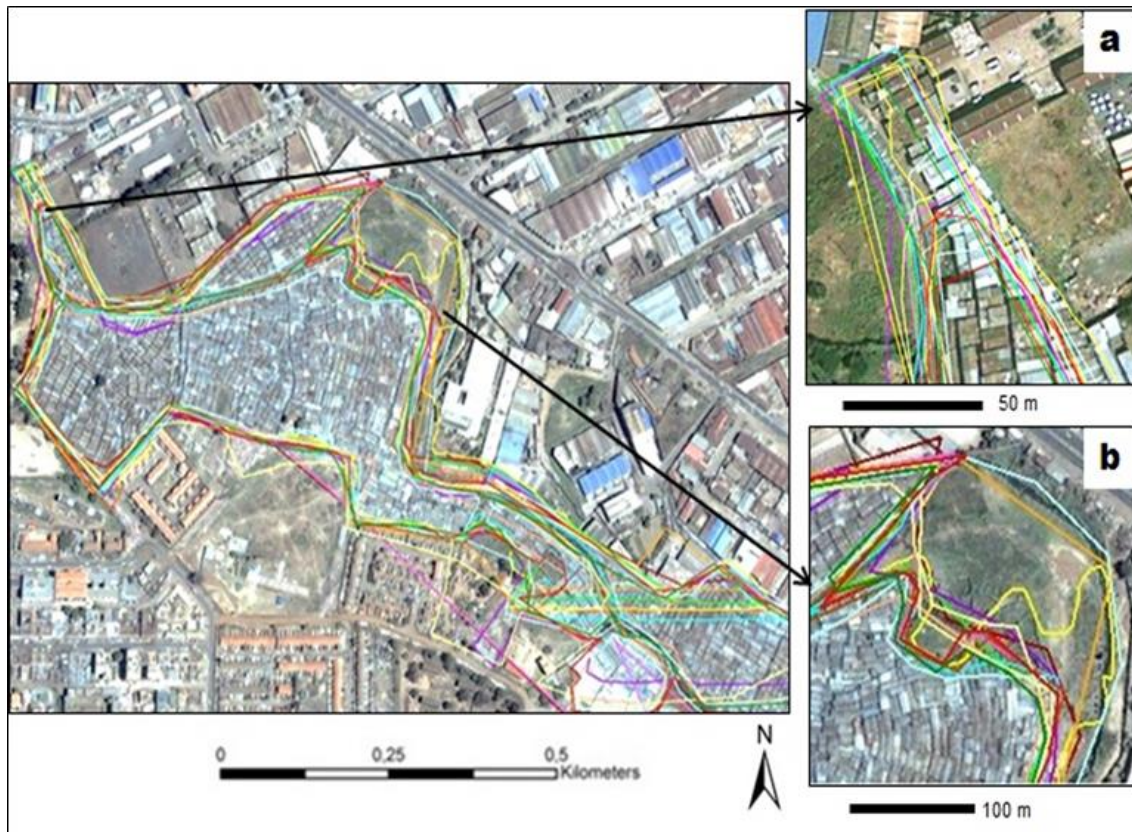
Figure 3: The challenges of mapping slum extent from image interpretation in Nairobi, Kenya. Each colored line indicates a different analyst's delineation of the same slum, illustrating semantic confusion. Adapted with permission from Kohli et al. [115].

Categorical mapping projects typically define a crisp set of non-overlapping categories, rather than a fuzzy set [116,117]. However, many human and natural land covers exhibit continuous gradation between classes, implying that crisp map legends will necessarily cause semantic ambiguity for training data labels, since many pixels represent a mixture of land cover types [118,119]. This problem is particularly evident with moderate and coarse resolution imagery, which often contain pixels comprising mixtures of the classes of interest [41]. When scene objects approximate the spatial dimension of the image resolution, local variance is highest, leading to poor classification accuracies [120]. While substantial research has been devoted to fuzzy set classification and the issue of mixed pixels, crisp categories are still often relied on during the training and testing phases of image classification [121], although Woodcock and Gopal [117] provide a less crisp approach to accuracy assessment. Data errors stemming from unclear class definitions can also be caused by incomplete training or data creation standards, which result in lack of key data or metadata collection in the field or during digitization [113]. Such inadequacies limit the analysis of training data error, and therefore the ability to account for error propagation.

Interpretation training data error also results from the need to collect large training data sets, which is brought on by the expanding extent of remote sensing analyses. Such large-scale training datasets are often collected rapidly, and entail labeling many individual training observations over a short field season or interpretation campaign. This trend is particularly exemplified by recent developments in large-scale collaborative data collections, often referred to as Volunteered Geographic Information (VGI) [55]. These data collection campaigns typically rely on web-based visualization and data contribution. Prominent examples include OpenStreetMap (OSM), DIYLandcover [109], Geo-Wiki[2] [46], Collect Earth[3] [122], and FotoQuest Go[4] [123]. In crowdsourced campaigns as well as smaller-scale data creation projects, the balance between speed and quality is a function of interpreter skill, experience, contextual knowledge, personal interest, and motivations for involvement in the data collection [20]. Interpretation errors can be exacerbated when there is insufficient interpreter education or familiarity with the area being mapped. For example, delineation of different classes of urban land use and socioeconomic status may be extremely difficult based only on the spatial and spectral information provided by an image, and without the benefit of local knowledge [112]. Inadequate or confusing user interfaces may also lead to error [20,112].

In such campaigns, citizen scientists contribute training data through a web-based platform [46,96,e.g. 108,109]. These approaches can potentially leverage the 'wisdom of the crowd', in which aggregate knowledge from many non-expert participants, post-processed for noise, can generate highly detailed, spatially extensive training data sets [46]. However, in part because of this open-ended nature, such data have inherent errors in both spatial representativeness and image interpretation, primarily stemming from uneven geographic contributions and lack of user experience working with EO data and methods. While some methods may mitigate such noise by leveraging the large number of observations [50–52], the data sets must be treated quite differently than expert-derived training data sets.

Errors also arise in *in situ* training data, caused by measurement error, geolocation inaccuracy, incorrect identification of relevant objects (e.g. vegetation species), and other such mistakes [124]. In addition to these factors, some feature types may actually be more difficult to discern on the ground than from an aerial perspective. Aside from these problems, technologically-induced training data errors arise from many and sundry causes, such as mapping or measurement software or hardware defects, user input error, or measurement error in field measurement devices (e.g. spectroradiometer calibration, GPS hardware). However, accounting for quantitative measurement error is more straightforward than thematic training data creation.

---

[2] [www.geo-wiki.org, 46]

[3] [http://www.openforis.org/tools/collect-earth.htm, 122]

[4] http://fotoquest-go.org/

Textbook tools to quantify measurement error are widely available, and *in situ* data collection procedures often include inter-analyst measurement comparison [125,126].

3. Impacts of Training Data Error

Lack of commonly held standards for creation of training data can lead to inaccurate or non-comparable maps, and thus miscommunication of geographic phenomena. Areal estimates of land cover classes can be substantially biased, which is particularly problematic for assessment of environmentally or economically important LCLU classes, particularly in change analysis [33,60,127]. Similarly, training data error can cause bias and poor variance estimation; for example, McRoberts et al. [33] showed that when interpreter error is not included in accuracy assessments, standard errors can be underestimated by a factor of 2.3. Critically, the choice and size of training data is particularly impactful on the outputs of increasingly used ML algorithms, which are susceptible to overfitting and class imbalance problems [57,107]. Additionally, the assumption of representativeness of training pixels is often overstated, and many training data may in fact not be generalizable to broader scales (discussed by Tuia et al. [128]).

Existing work indicates the potential for errors in the inputs to the map-making process to propagate substantial errors to map outputs [25,34]. While it is possible to quantify error from training data in output maps using the methods of McRoberts et al. [33], the full trajectory of error from initial creation to output map to downstream product has yet to be fully explored. However, recent work from Estes et al. [36] has begun to examine the propagation of error from primary land cover maps to subsequent derived products. This work used a high quality reference cropland map to quantify the errors in 1 km cropland fractional estimates derived from several existing land cover datasets, and quantified how these errors propagated in several types of map-based analyses that required cropland fractions as inputs, including estimates of crop production and yield, carbon densities, evapotranspiration, and food security. The results suggest that downstream errors were in some instances (e.g. carbon stock estimates, Figure 4) several fold larger in magnitude than those in the input cropland maps, whereas in other applications (e.g. evapotranspiration estimates), the errors were muted. In either case, the degree to which the error magnifies or reduces in subsequent maps is hard to anticipate, and the high likelihood that it could have the former effect means that any conclusions based on such land cover-derived maps must be treated with caution if the error propagation is not quantified. Although this analysis did not examine how training data errors impacted the maps they were used to create, it provides some insight into the potential scope of the problem and suggests a method for characterizing how training errors could impact map accuracy.

Figure 4: An examination of how error in pixel-wise cropland fractional estimates (expressed as a percentage, top row) can propagate error (expressed as a percentage) in maps that use land cover data as inputs, such as estimates of carbon density (bottom row). Figure adapted from Estes et al. [36].

Similarly, unassessed errors may exist in many widely used products, such as the National Land Cover Map [NLCD, 129], which, as a well-known and large-area product, is used for a wide variety of purposes, including an input into methods for mapping urban tree canopy biomass. Urban tree play an important but under-examined role in regional carbon cycles [130–132], in addition to other ecosystem services such as latent heat flux and albedo modification [e.g., 132]. With urban lands expected to triple between 2000 and 2030 [133,134], the need to factor them into carbon accounting is becoming more pressing, but urban areas are often omitted from EO studies of carbon dynamics [e.g., MODIS NPP, 135], or else exhibit large uncertainties. The ability to remotely map urban tree cover is limited by a) spatial resolutions that are often too coarse to reconcile the high spatial heterogeneity of urban landscapes and b) training data that is often biased to forested, agricultural, and other rural landscapes. The Landsat-derived NLCD Percent Tree Cover product [136] estimates canopy cover at 30-m resolution across the U.S. for 2011 and 2016, and can be used with allometric equations to map tree canopy cover and biomass dynamics at medium resolution. However, Coulston et al. [136] found that urban areas exhibited the highest level of uncertainty, and Smith et al. [137] found that the NLCD 2001 canopy cover product underestimated canopy cover in Baltimore, Maryland by 11% compared to a high resolution dataset in urban areas. Therefore, to accurately estimate urban biomass, with correctly

specified uncertainty, using NLCD as input for biomass maps will require better accounting of error.

The accuracy of this dataset was further explored in a suburb of Washington, D.C., USA; VHR Wayback World Imagery was used to manually digitize canopy cover within a heterogeneous 225 ha landscape typical of many suburban areas. We found that the NLCD percent canopy cover product underestimates canopy cover by 15.9% (7.6 ha) in this landscape. The largest discrepancy between the manually mapped canopy cover and the NCLD product occurred along forest edges (Figure 5); in these areas, the NLCD underestimated canopy cover by 27%. This discrepancy is particularly important in heterogenous urban landscapes, where forest edges comprise a high proportion of total forest area. Scaling field data from forest plots to the entire forest area on the NIST campus yielded an estimate of 8,164 Mg C stored in aboveground forest biomass, based on our high-resolution map of forest canopy cover, compared to only 5,960 Mg C based on the NLCD canopy cover product. These differences are particularly problematic from a carbon accounting perspective, as temperate forest carbon storage and rates of sequestration are much larger (64% and 89%, respectively) than in forest interiors [138]. Moving forward, it is important to improve the representation of urban areas in large area products such as the NLCD, which serve as inputs for downstream mapping and estimation of regional canopy cover and biomass. Ultimately, specifically assessing the uncertainty of these products will greatly improve the capacity to quantify vegetation and carbon dynamics in urban landscapes.

Figure 5: Spatial variations in canopy cover (A) and uncertainty in canopy cover estimates (B) in forested and non-forested areas of the heterogeneous suburban landscape of the National Institute of Standards and Technology campus in Gaithersburg, Maryland. Percent canopy cover at a 30-m resolution from the commonly used National Land Cover Database (NLCD) Percent Canopy Cover product (and its uncertainty) is superimposed over a high-resolution map of forested areas (hollow outlined polygons) and non-forest trees (e.g., street trees; solid polygons) that were manually mapped using <1-m resolution Wayback World Imagery. Note the lower estimates of percent canopy cover along forest edges (A) and the associated higher levels of uncertainty (B) using the NLCD product.

To represent the growing set of remote sensing applications in which training data error may be encountered, we present a set of case studies that explore both categorical and continuous modeling and mapping. To help lay a common framework, we show a typical methods sequence for a ML-based remote sensing analysis in Figure 6, which also helps clarify the terminology used in this paper. The figure shows the various sources and implications of error in the modeling and mapping process, beginning with issues in the data sources and sample design, and continuing through model training, validation, and ultimately in map accuracy assessment.
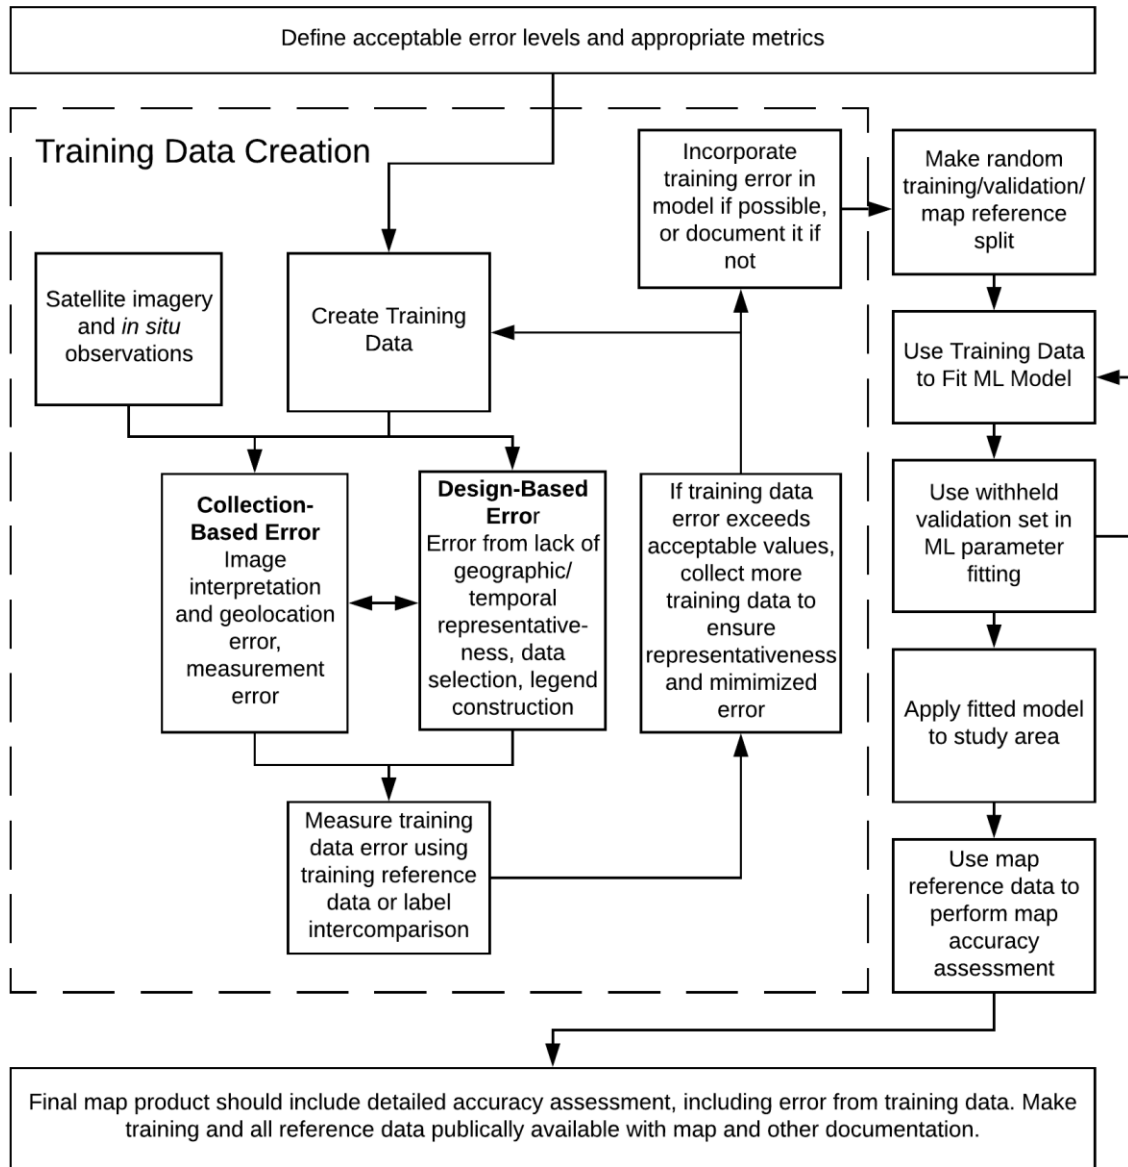
Figure 6: Flow chart of typical workflow for machine learning applications in Earth observation data.

4. Case Studies

The examples presented above show how errors in the land cover maps lead to errors in analyses based on these maps, and are suggestive of the map errors that might result from training data errors. To illustrate this issue more clearly, we provide several case studies across different

mapping applications that represent the broad range of mapping and modeling applications that rely on training data.

4.1 Infrastructure Mapping

Automated building footprint detection is an important but difficult mapping task, potentially benefiting a wide range of applications. The following case study illustrates the use of Raster Vision[5], an open source deep learning framework, to train several models for automated building detection from high resolution imagery[6]. These models perform best when trained on a large number of correctly labeled examples, usually generated by a paid team of professional labelers. An alternative, less costly approach was conducted in which a building segmentation model was trained using labels extracted from OpenStreetMap (OSM). However, the labeled training polygons generated from OSM contain errors: some buildings are missing, and others are poorly aligned with the imagery or missing details. This provides a good test case for experimentation on how noise in the labels affects the accuracy of the resulting model.

In this case study to measure the relationship between label noise and model accuracy, the amount of label noise was varied while holding all other variables constant. To do this, an off-the-shelf dataset was used in place of OSM, into which label errors were systematically introduced. The SpaceNet Vegas buildings data set was used as a source of training data in these experiments. To this relatively large data set (~30,000 labeled buildings)[7], missing and imprecisely drawn building errors were systematically introduced, and the resulting model accuracy was measured. The experimental design consisted of two series of six datasets each, with random deletion or shift of buildings at increasing probabilities and magnitudes, respectively. For each dataset, a UNet semantic segmentation model with a ResNet18 backbone was trained using the fastai/PyTorch plugin for Raster Vision[8]. These experiments, including data preparation and visualization, can be replicated using code at[9].

Figure 7 shows the ground truth and predictions for a variety of scenes and noise levels, showing that the quality of the predictions decreases with the noise level. Also, the background and central portions of buildings tend to be predicted correctly, whereas the outer periphery of buildings presents a greater challenge.

---

[5] https://rastervision.io/

[6] Additional detail available at: https://www.azavea.com/blog/2019/08/05/noisy-labels-deep-learning/

[7] https://spacenetchallenge.github.io/datasets/spacenetBuildings-V2summary.html

[8] https://github.com/azavea/raster-vision-fastai-plugin

[9] https://github.com/azavea/raster-vision-experiments/tree/master/noisy_buildings_semseg

Figure 7: Predictions of the model trained on different noisy datasets. Each row shows a single scene over different noise levels. The top two rows show noisy drops, while the bottom two rows show noisy shifts. The ground truth is outlined in light blue, and the predictions are filled in orange.

4.1.1 Incorporating Noisy Training Label Data

These results are quantified in Figure 8, which shows F1, precision, and recall values for each of the noise levels below (see Table 1 for terminology description). The precision falls more slowly than recall (and even increases for noisy drops), which is consistent with the pattern of errors observed in the prediction plots. Pixels that are predicted as building tend to be in the central portion of buildings, leading to high precision.


Figure 8: The precision, recall, and F1 scores across different noise levels are shown for the cases in which labels are randomly dropped (A) or randomly shifted (B). The F1 score of the building predictions is plotted as a function of the F1 score of the label (C), for the cases when labels were either randomly dropped or shifted.

In panels (A) and (B) of Figure 8, the x-axis shows the noise from randomly dropped and randomly shifted labels, respectively. Panel (C) combines the effects of noisy deletions and noisy shifts on accuracy in a single graph, using F1 to measure error. The F1 score of the noisy versus ground truth labels is a function of the pixel-wise errors; this metric has the benefit of measuring the effect of noise on error in a way that is comparable across datasets and object classes. For instance, a noisy shift of 10 in a dataset with large buildings might result in a different proportion of erroneous label pixels than in another dataset with small buildings. From this, it is shown that some of the shifted datasets have a greater level of noise, but that the prediction F1 scores are similar between the two series when the noise level is similar.

These experiments present a small step toward answering the question: how much accuracy is sacrificed by using training data from OSM? Preliminary results indicate, as expected, that accuracy decreases as noise increases, and that the model becomes more conservative as the noise level increases, only predicting central portions of buildings. Furthermore, the noisy shift experiments suggest that the relationship between noise level and accuracy is nonlinear. Future work will quantify the functional form of this relationship, and how it varies with the size of the training set. Some preliminary work toward this goal has been described in Rolnick et al. [139], which focuses on image classification of Imagenet-style images.

One limitation of these results is that the magnitude of error in OSM for most areas is unknown, making it difficult to predict the effect of using OSM labels to train models in a generalized, global sense. 'Noisy' error in OSM can be estimated by measuring the disparity between OSM labels to clean labels, such as the SpaceNet labels used in this case, providing a local estimate of OSM noise. A more general but less rigorous approach is to roughly estimate the noise level by visually inspecting the labels in OSM, and comparing to Figure 7, which shows examples of the labels at different noise levels.

4.1.2 Detecting Roads from Satellite Imagery

Road networks constitute a critical geographical data layer used to assist national decision makers in resource allocation, infrastructure planning, vaccination campaigns, and disaster response, among others. However, accurate and up-to-date road networks are not available in many developing countries. High resolution satellite imagery, paired with deep learning methods, provides the capacity to detect and map roads at large spatial scales. This important goal, however, is dependent on availability of local high quality training data.

To evaluate the impact of local training data availability on predicted road network accuracy, a study was carried out in Kumasi, Ghana [140]. Two datasets were used to train ML models: 1) the SpaceNet[10] Dataset [141] in Khartoum, Sudan, and Las Vegas, USA, and 2) OSM data in Kumasi, Ghana. The SpaceNet Dataset includes high quality road labels with human expert validation, but unfortunately was not available in Kumasi, Ghana. Therefore, the latter study site relied on OSM data, consisting of crowdsourced labels with no accuracy assessment or expert validation. A series of experiments were carried out to assess the feasibility of using transfer learning, using the Raster Vision library in Python for training and evaluation. For all MobileNet V2 models introduced in the following list, the image chip size was set to 300 x 300 pixels, and the training/validation split was 80/20.

The Las Vegas Model was trained and validated on SpaceNet data in Las Vegas, and produced very high accuracy predictions. However, when this model was used in Kumasi, it predicted very few roads, with only scattered road segments. The Khartoum Model was also trained using SpaceNet data in Khartoum,. The Kumasi Model used Maxar (previously DigitalGlobe) WorldView-3 imagery and labels from OSM as input. OSM was used to test the quality of crowdsourced labels in training a road detection model. The Khartoum Model was then fine-tuned on OSM labels in Kumasi for three different steps of 100K, 50K and 10K. All models used the same hyperparameters, to isolate the role of training data on model performances.

To validate the models' performance using an independent dataset, a set of expert labels were generated over a small part of Kumasi. Figure 9 shows the region with human expert data

---

[10] https://spacenetchallenge.github.io/

23

vetting, along with the three model predictions. The Las Vegas model is excluded from this figure as it does not have any meaningful prediction in Kumasi. Quantitative performance metrics were calculated using the human expert labels, which the models had been blind to during training. The results indicate that, as clearly indicated by Figure 9, the F1 score for roads was substantially higher for the Kumasi Model (0.6458) than when using the Khartoum model (0.3780). However, by retraining and fine-tuning the Khartoum model, the F1 score for roads increased to 0.6135. The full accuracy results for this experiment are presented in full in Table A1.
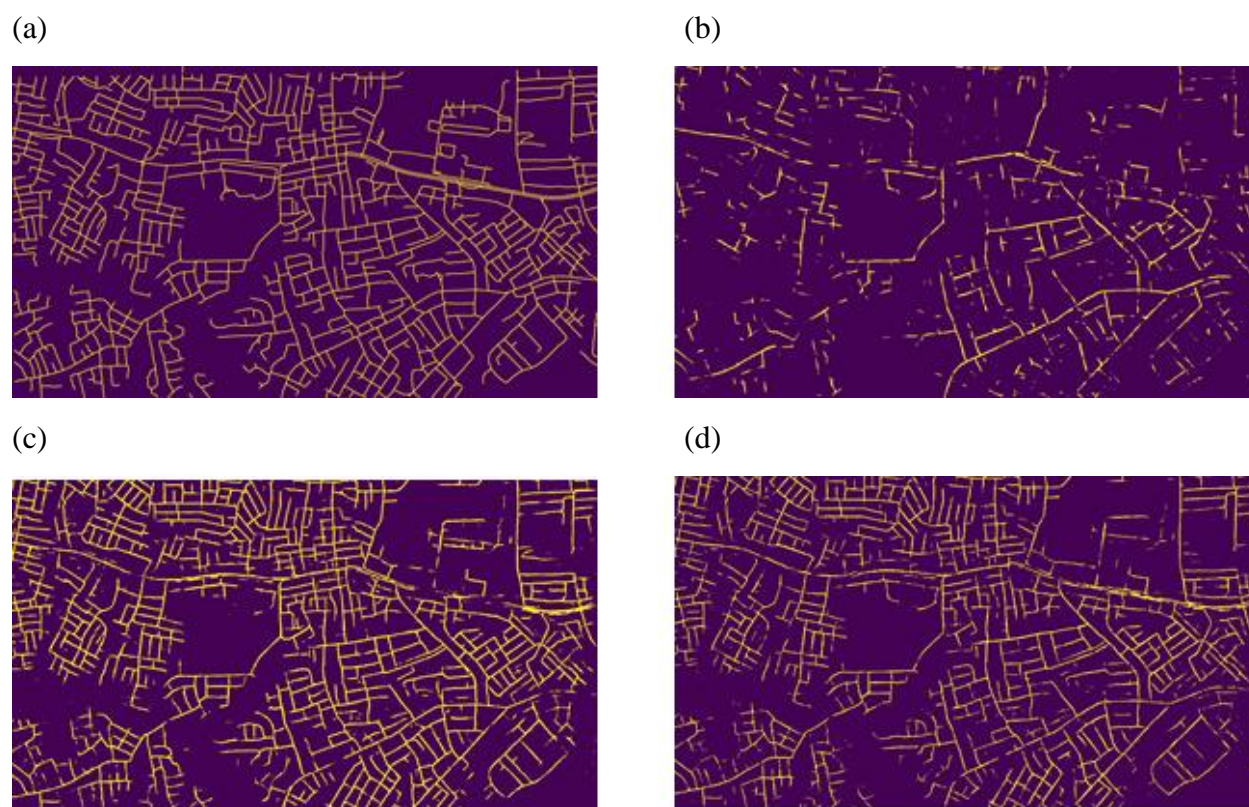
(a)

(b)

(c)

(d)

Figure 9: (A) Labels generated by experts for validation. (B) Predictions from the Khartoum Model. (C) Predictions from Kumasi Model. (D) Predictions from Khartoum Model retrained in Kumasi with 10K steps.

Based on these results, it is concluded that: 1) lack of diverse training data significantly limits the geographic applicability of models, as the types, surfaces, and arrangements of roads various substantially between regions; 2) regional training datasets are essential for the model to learn the feature of roads in that region; and 3) transfer learning from a reasonably similar geography can help models train better.

4.2 Global Surface Flux Estimates

Fluxes at the land-atmosphere boundary play a key role in regulating water, carbon and energy cycles. These fluxes include latent heat flux (LE), sensible heat flux (H), and gross primary production (GPP). While these fluxes cannot be measured directly from remote sensing observations, other remotely sensed variables can be used to estimate these fluxes. Moreover, these three fluxes are highly coupled, and therefore a coupled model is ideal.

A fully connected neural network model was developed for this purpose [142], named Water, Energy, and Carbon with Artificial Neural Networks (WECANN). Inputs to WECANN are remotely sensed estimates of precipitation, soil moisture, net radiation, snow water equivalent, air temperature and solar induced fluorescence (SIF). The target variables for training the model were derived from outputs of global models. However, this presents the difficulty that the target variables are model outputs that can have substantial error, which will propagate in the WECANN model. To mitigate this problem, three independent estimates of each of the three fluxes (LE, H and GPP) were retrieved from the global models. Then a novel statistical approach, named Triple Collocation (TC, Figure A1, equation A1), was used to combine those estimates to a new dataset for training WECANN model.

Triple collocation (TC) is a technique for estimating the unknown error (measured with standard deviations or RMSEs) of three mutually independent measurement systems, without treating any one system as zero-error 'truth' [143]. The three measurements systems estimate a variable collocated in space and time, hence called Triple Collocation. Using these probabilities, at each pixel and at each time one of the three estimates of the target variable is randomly selected to generate the training data.

The results of WECANN model outputs were evaluated against ground measurements from global FLUXNET towers from 2007 to 2015 (Figure 10), using both the coefficient of determination and Root-Mean-Squared-Error (RMSE) to evaluate accuracy. These show that WECANN's correlation was on average 17% higher (range 8-51%) than that of any one of the three individual inputs, while the RMSE was 21% lower (range 4-54%). These differences provide a partial quantification of the error inherent in any one of these training inputs, and show that by combining them using the TC technique, we can reduce error in an ML model for predicting the fluxes at global scale. This case study illustrates a means of assessing and accounting for error in training data, for cases in which these data are not created specifically for

Figure 10: R^2 and RMSE of the WECANN model output against ground measurements from FLUXNET towers in comparison to the three datasets used to generate the target training data for LE (a, b), H (c, d) and GPP (e, f).

the project, but rather are pre-existing data products with potentially quite different characteristics and potentially unknown error.

4.3 Agricultural monitoring

Two agricultural cases illustrate how training data error can impact both categorical and quantitative remotely sensed measures. The first relates to cropland mapping, and is drawn from an ongoing study focused on mapping smallholder agricultural fields at high resolution (3-4 m) in Ghana. The mapping method is based on "active learning", in which a RandomForest-based

[144] ML algorithm is iteratively trained and validated by a crowdsourcing platform, which enlists human trainers to visually interpret and digitize field boundaries visible within the imagery (PlanetScope visual and near-infrared surface reflectance [86]) being classified [108,109,145]. The crowdsourcing platform incorporates a protocol for assessing the accuracy of training labels, in which each worker is periodically directed to a training reference site where the boundaries are already known, but are not visible to the worker. Using these training reference sites, their maps are then scored using a multi-dimensional accuracy assessment algorithm (Estes et al, in prep), resulting in an average training data accuracy score for each worker that ranges between 0 (complete disagreement with reference) and 1 (perfect agreement). Each label site is mapped by at least five workers, and the resulting worker-specific accuracy scores are used within a Bayesian merging algorithm to combine the five sets of labels into a single consensus label, which is then used to train the RandomForest classifier. Here we use the worker-specific training accuracy scores to assess the impact of label quality on map accuracy, by assessing three variants of two RandomForest-generated maps, one over Central (~3,400 km^2) and one over Northern Ghana (~3,100 km^2). The first two maps were trained using labels generated by the least accurate worker to map each training site, the second two by the most accurate worker to map each site, and the third using the consensus labels. The accuracy of each pair of maps was then assessed against an independent training reference set (reserved consensus labels) using the True Skill Statistic [71] (sensitivity + specificity - 1, with scores ranging from - 1 to 1). The results show a substantial difference in accuracy between the maps trained with the least and most accurate workers' labels (Figure 11A), with the former having 7-9% more skill than the latter, while maps based on consensus labels have ~3% more skill than those of the most accurate workers' labels.

The second case relates to remotely sensed crop estimates of wheat yields collected in 48 smallholder fields in Bihar, India in 2016-17 [146]. Yield data were collected via eight 2x1 m^2 crop cuts within each field, and PlanetScope-derived green chlorophyll vegetation indices (GCVI) were calculated over each field from imagery collected over four dates during the growing season (January 13, February 25, March 12, and April 14, 2017). A RandomForest regression was trained on the yield measured for each field, using the four dates of GCVI values as predictors. To test the effect of training data error on the resulting yield predictions, three types of noise were artificially introduced into the yield data used for training: 1) a systematic 0.5 ton/ha overestimate, and randomly distributed errors sampled from a normal distribution with a mean of 0 ton/ha and 2) standard deviations of 0.5 ton/ha and 3) 1 ton/ha. A baseline model fit to unperturbed data was also developed. Each model was trained on three separate randomly selected subsets of 32 perturbed observations, and the predictions were made for the remaining 16 held-out (independent) yield observations, which were not perturbed. This three-fold cross validation process was repeated 50 times, with each permutation using a different random seed to construct the folds, in order to achieve stable error metrics. The model performance was assessed

by calculating the averages of the mean absolute error (MAE) of the prediction, and the R^2 of regressions fit between prediction and observed values (Figure 11B).
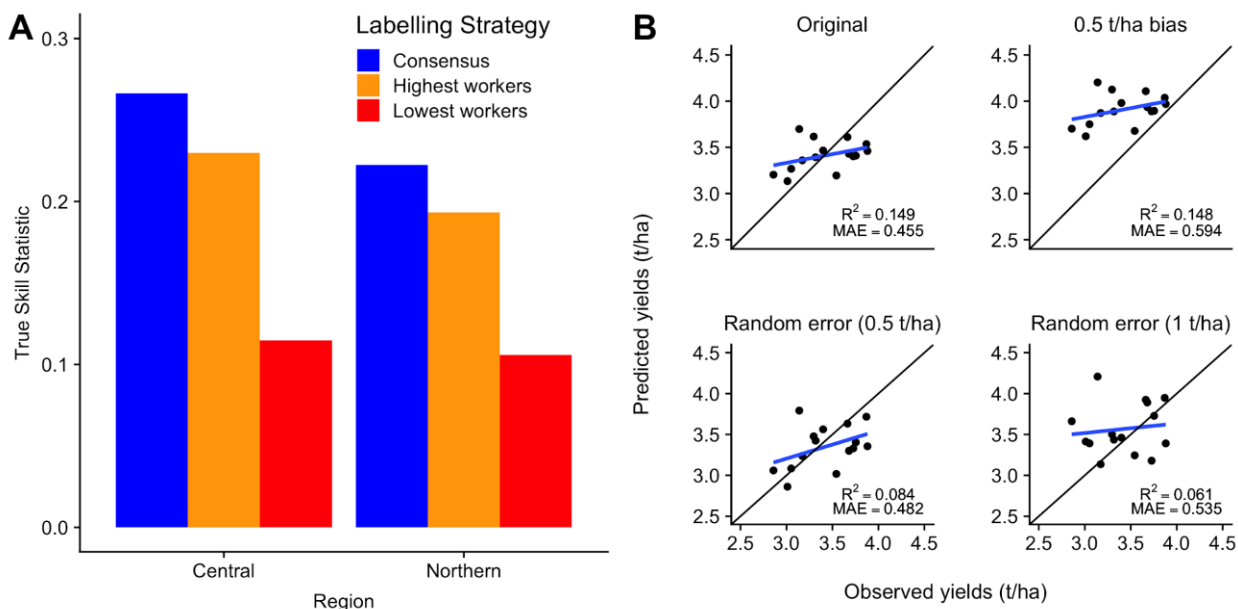

Figure 11. A comparison (A) of the accuracy (based on the True Skill Statistic) of cropland maps over two areas of Ghana when generated by labels of different levels of quality (red = least accurate workers' labels; orange = most accurate workers' labels; blue = "consensus" labels made by merging all workers' labels). (B) Results from a RandomForest model of wheat yields trained on satellite-derived vegetation indices, showing the relationship between predicted yield and independent observed yields, in terms of the fit against the 1:1 line and the regression slope of the relationship (points and regression line represent the mean of a single randomly selected model permutation). The average mean absolute error (MAE) and average regression R^2s calculated across all permutations are shown for each model.

The results show that four models, including the baseline, compressed the range of yields, as seen in the shallow slope between observed versus predicted values, but prediction error was 18-31% higher when training yields had either the high level of random or systematic error within them. The smaller amount of random noise only added about 6% error to the predictions, suggesting that RandomForest is tolerant to some training error. Note that the average R^2 of the observed-predicted regression fit was nearly the same for the systematic error case as the baseline, which shows that this metric can be an unreliable measure of performance for quantitative measures, and that it is important to assess fit against the 1:1 line and using a metric such as mean absolute error.

5. Guidelines and Recommendations

Although our review and case studies show that impacts of training data error on EO applications can vary, and that the existing literature on this problem is scarce, several best practices and guidelines emerge from this work. Below we outline a series of suggested steps for minimizing and accounting for training data error, within the context of undertaking and assessing the accuracy of a typical ML-based mapping project.

Step 1: Define acceptable level of accuracy and choose appropriate metric

As a starting point, analysts and map-makers should first determine the level of error that is tolerable for their particular application, and then determine the metrics that are best suited for assessing map accuracy within the context of the research questions or hypotheses [147]. For example, if the goal is to estimate the number of hectares of a category, then commission and omission errors of that category cancel each other out, and therefore the gross error for the category is less relevant than the net error for the category [27]. In the case of a continuous variable, in which the absolute accuracy of the mapped value is of greatest importance (e.g. for predicting grain yields in different areas), assessing the fit of the relationship between predicted and observed values along the 1:1 line is important, together with measures such as mean absolute error and mean error [75,76].

It is important to emphasize that the ability to determine tolerable map error is limited by the extent to which the amount of error in the map *reference data* is known. The error in these data determine the upper limit of achievable map accuracy, since a model's predictions cannot be more accurate than its map reference data (or if they are, then that cannot be known). Put more simply, if the map reference data are only 90% accurate, then the map can be at most 90% accurate. While this may seem intuitive, it is often not sufficiently considered, since reference data accuracy is often assumed to be perfect [25,34,48], when in fact, as with training data, there is always likely to be some error in map reference data. Map reference data error has two important ramifications. First, if the map reference data error is greater than the tolerable error level, then there is little point in attempting to improve the map algorithm, unless a different reference dataset with tolerable error can be obtained. Second, if the error in the map reference data is unknown, then it is impossible to properly assess map accuracy, including how that accuracy is impacted by training data error.

Although the aforementioned considerations relate to map reference data, they are highly relevant to and intertwined with training data. Training and reference data are often affected by the same sources of error, particularly when both datasets are collected as part of the same methods, which is common practice [42,108]. The procedures for minimizing and accounting for

errors in both datasets are thus often the same. Our subsequent recommendations therefore cover both training and map reference datasets, except where we indicate necessary distinctions. Furthermore, the method we recommend for assessing labelling uncertainty may provide the most practical and effective method for assessing the degree of truth in many reference data (see section 1.2.2).

Step 2: Minimize design-related errors

The next logical step in a mapping project is to design a strategy for collecting the training and map reference samples. Although there are numerous factors to consider, there are several general aspects of design that can help minimize potential training errors.

Sample design

The first of these relates to sampling design itself, i.e. where, when, how many, and what type of samples are placed. With respect to the training data, to a certain extent this depends on the requirements of the selected ML algorithm, which can, for example, have differing requirements with respect to class balance [e.g. 107]. However, one critical design aspect relates to cases where training and reference samples are collected simultaneously. In these cases, it is essential that considerations of the training data sample not undermine the standards required for an independent, probabilistic map reference sample [sensu 48]. These standards are not necessarily violated when both training and reference samples are collected by the same people [47]. Instead, the potential for error arises if iterative refinement of the algorithm is done against the map reference dataset. This problem can best be understood within the context of cross validation, which is appropriate for ML parameter selection [e.g. 107], but when the number of folds exceed one (as in our yield estimation case study; Figure 11B) then the portions excluded from training lose statistical independence and can no longer serve as the map reference [58]. However, the risk of losing map reference data independence may also occur when training sites are selected iteratively, in order to increase their representativeness and improve ML performance [42,e.g. 108]. If the gain due to new training sites is assessed against the map reference, then it will also lose independence after the first iteration. Moreover, any error in the map reference sample will be integrated into the final map. Xiong et al [42] avoided this problem by assessing performance gains from additional training data through visual assessment of their ML-generated cropland maps against other similar land cover products. A more quantitative approach is to divide an initial sample into three splits: one for training, the second for validation for algorithm improvements, including those related to the addition of new training sites, and the third as the map reference, used only for final accuracy assessment. This partition approach was used for the cropland mapping case study [Figure 11A, 148].

Geographic representativeness and the degree to which they capture the variability in the feature of interest is an important sample design consideration [109,149]. The road mapping case study

shows the errors that can result when maps are trained with samples that do not adequately represent the features in a particular region. Training data can in practice be highly localized or relevant or accurate for a limited spatial extent or temporal period [112,140]. This problem may become more relevant given the increase in stock or benchmark training libraries, and attempts to transfer pre-trained models to other regions, time periods, or scales of observation [57,150]. While such training libraries can present an immense benefit to large extent EO research, if these are to be relied on for training, their representativeness of the features of interest should be assessed, and augmented as needed, as in Khartoum model case study (Figure 9D). However, we suggest that the best practice is to train using data specifically collected within a bounded region for the feature being mapped [e.g. within a particular agroecoregion, 42,151], and avoid over-generalizing or transferring models to other regions [152]. When using citizen science or crowdsourcing approaches to collect these data, representativeness is ensured by directing labellers to the selected training data reference sites [e.g. 109], rather than having them select regions to map.

Samples should also be temporally representative of the imagery that is being classified. That is, relative to the imagery being classified, the training (and map reference) sample should fall within a window of time that matches the characteristic rate of change of the feature being mapped. This temporal interval can be estimated by assessing the temporal autocorrelation in the feature of interest [153]. For rapidly changing features, such as the timing of deforestation events, snow/ice melt, and vegetation coverage during phenological transition, the sample may need to be captured within a few days or weeks of the acquisition of the imagery being classified, whereas for slower-moving features a sample collected within a few years may be sufficient.

Training Data Sources

The requirements for temporal representativeness makes the source of training imagery a critical consideration for projects that rely on image interpretation. The use of basemap imagery is not recommended for training when mapping dynamic features, given the broad range and uneven distribution of image ages in these basemaps [39], unless the age of the imagery being classified can be matched to that of the training sample. Otherwise, there is substantial potential for introducing error into the algorithm (e.g. Figure 1), and its impact may be hard to assess, particularly if the reference sample was also collected from the basemap. The goal of temporal representativeness must be balanced with the need to have a sufficiently high spatial resolution for accurate image interpretation, which helps minimize errors during the collection of the sample (see Step 3). Beyond matters of cost, this tradeoff is presumably one reason why HR/VHR basemaps are widely used [39]. New commercial imagery, such as PlanetScope [86], which are collected at high frequency (near-daily) with a sufficient spatial resolution (3-4 m) for many visual interpretation tasks, may be a preferable source of training imagery for developing maps that represent the post-2016 period.

It is also important to consider and account for additional sensor and imagery characteristics that shape interpreter data entry, such as spatial resolution, atmospheric quality (e.g. clouds, haze), sensor view angle, sun angle, spectral band selection, and image rendering contrast stretches. These characteristics should be taken into consideration in the sample design stage, although this information may not be available to obtain *a priori* [53].

Legend design

For thematic maps, legend design merits special consideration as it relates to training data, particularly for multi-temporal and/or large area projects that rely on multiple image datasets [39]. As discussed above in section 2 above, objects of interest, including land cover types, should be at least twice as large as the pixel resolution of the imagery used in the classification algorithm [97,120,154]. When image spatial resolution is too coarse relative to the scene elements of interest, image interpretation errors are likely due to mixed pixels [98–100]. This implies that in designing a legend, researchers should select classes that can be effectively mapped using the coarsest resolution imagery that will be incorporated in the model, and avoid the problem of collecting training samples having mixed pixels. This is especially relevant since HR/VHR imagery is often used to create training data polygons and labels, while moderate or coarse resolution imagery is most often used in model processing [e.g. 42,155–157].

Continuous training data, particularly those collected *in situ*, are often point samples, and therefore a sampling protocol should be used to match field measurements and pixel dimensions in order to avoid scaling problems associated with the modifiable areal unit problem [101,102]. Spatial representativeness should be considered as a limiting factor for legend design, and to the extent possible, researchers should attempt to use categories that are supported by both the spatial resolution of the model data and the field sampling protocols to be used [103–106].

Step 3: Minimize collection-related errors

There are numerous ways to collect training data for categorical and continuous mapping projects, each with their own sources of error. There are therefore many potential approaches for minimizing associated collection errors, which in many cases may be relevant only to a particular variable [e.g. for agricultural area estimates 158]. Practices for controlling measurement errors in continuous variables are often well understood and part of conventional field data collection practices. We thus focus here primarily on strategies to minimize error in image interpretation, and increasingly common practice used for training categorical mapping algorithms. We also touch on the specific case of model-derived training data.

McRoberts et al. [33] showed that an increased number of interpreters reduces bias in the estimated variable. Since this is the only readily controlled factor during image interpretation campaigns, these authors recommend at least three image interpreters be used, such that at

minimum a majority vote interpretation rule be possible. For more unequal stratum (i.e. land cover category) proportional size, estimated bias was shown to be larger, and thus as many as seven interpreters are recommended in those cases. In the cropland mapping case study, five interpreters were used per site.

Beyond the design considerations mentioned above, there are a number of steps that can be taken to minimize error during the collection of the training sample. Whenever possible, we recommend using built-in accuracy assessment protocols, particularly for image interpretation [109]. For example, active feedback during training label creation can help reduce errors on a rolling basis, by providing interpreters information regarding their performance [159]. This strategy relies on predefined training reference data for comparison with interpreter-generated labels. Training reference datasets can be limited in size compared to the ultimate sample size, provided that training reference locations are randomly presented to interpreters during the data creation campaign.

Since absolute truth is an impossible standard, consensus among domain expert interpreters can be used as the best and most practical measure of 'truth' for the training and map reference data [22,47]. As illustrated by the agricultural and surface flux case studies, methods involving multiple independent measurements at the same location provide a mechanism for assessing training data uncertainty (Figure 12), in addition to minimizing error during creation. In the case of image interpretation, this involves multiple interpreters assessing the same site, similar to *in situ* methods in which multiple analysts independently make the same measurement of a continuous variable.

In any case, image interpreters should be given thorough instruction regarding remote sensing principles as well as local or regional contextual information. Analysts with local domain expertise are particularly helpful in consistent identification of idiosyncratic land covers [115]. Our recommendations regarding interpreter education are particularly relevant for crowdsourcing or citizen science data collection campaigns, in which it cannot be assumed that participants have any formal experience in image interpretation.

As described above, image interpretation is inadvisable in cases where the available imagery does not support the legend categories, or the similar but potentially more hazardous case that HR/VHR imagery is used to create training samples that are then used with coarser resolution imagery when ingested into the ML model [see 97]. Assuming that researchers correctly specify their data selection and legend design, image interpretation errors due to insufficient resolution should be minimized; however, special care should be given to borderline classes, or classes exhibiting a high degree of spatial and/or spectral variability due to land cover mixtures within the pixel [98–100]. In such cases, we recommend that training protocols include the practice of leaving a buffer around training polygons whose pixels may be mixed over relatively small distances, as described by Xiong et al. [42].

A special case of training collection for continuous variable mapping is approaches in which a simulated data, typically a process-based model, is used to generate the training data. A prominent example of this approach is seen with crop yield mapping, where field-estimates of yield are hard to obtain and can have substantial error [160]. One particular example is the Scalable Yield Mapping (SCYM) method, in which a crop simulator is used to generate a large number of realizations of yield under various environmental and management conditions, and the outputs are used to train an empirical model (typically ML) with remotely retrievable predictors that is used for yield mapping [161,162]. Training data errors in such cases can be minimized by rigorously calibrating models (itself a challenging task) according to best practices from the relevant modeling literature, such as those recommended for crop simulation [e.g. 163]. Alternatively, if modeled training data are necessary but careful calibration is not possible (e.g. because the data are pre-existing), then a merging approach such as Triple Collocation (Section 4.2) may be employed to reduce training error.

Step 4. Assess training data error

The best way to assess training data error is to measure it directly. For continuous variables, calculating measurement error should be possible in many cases, even for model-generated training data, in which the variance can be calculated from simulation treatments [e.g. 163]. For categorical mapping, training error can be measured using an internal accuracy assessment protocol that makes use of predefined training training reference data (e.g. Estes et al., [109]. During training data creation, whether *in situ* or via image interpretation, we recommend calculating quality metrics for the data creators, relating to speed, precision, and consistency. This recommendation is based on experience in crowdsourced data creation [96,109], but it is applicable to any type of data collection, and could greatly bolster the understanding and quantification of error propagation.

However, it can be challenging to produce training and map reference data [25,26,33,69], and indeed in some cases the true category is not clear whether looking at an image or standing on the ground in the same location. When target land cover classes are vague or difficult to delineate, local knowledge can help compensate for difficulties in image interpretation (Figure 3). At the minimum, we recommend reporting the percentage of the reference observations for which the category is ambiguous.

If an accuracy protocol including pre-defined training reference observations is not possible, then we recommend that researchers calculate uncertainty estimates based on repeated measures approaches, as described above and shown in Figure 12; this is useful for both training and map reference data. If no quantification of training data error is possible, then researchers should at the very least clearly document the data creation methods, and detail likely sources of error and potential uncertainties.

Table 2: Template and procedure for documenting training data.

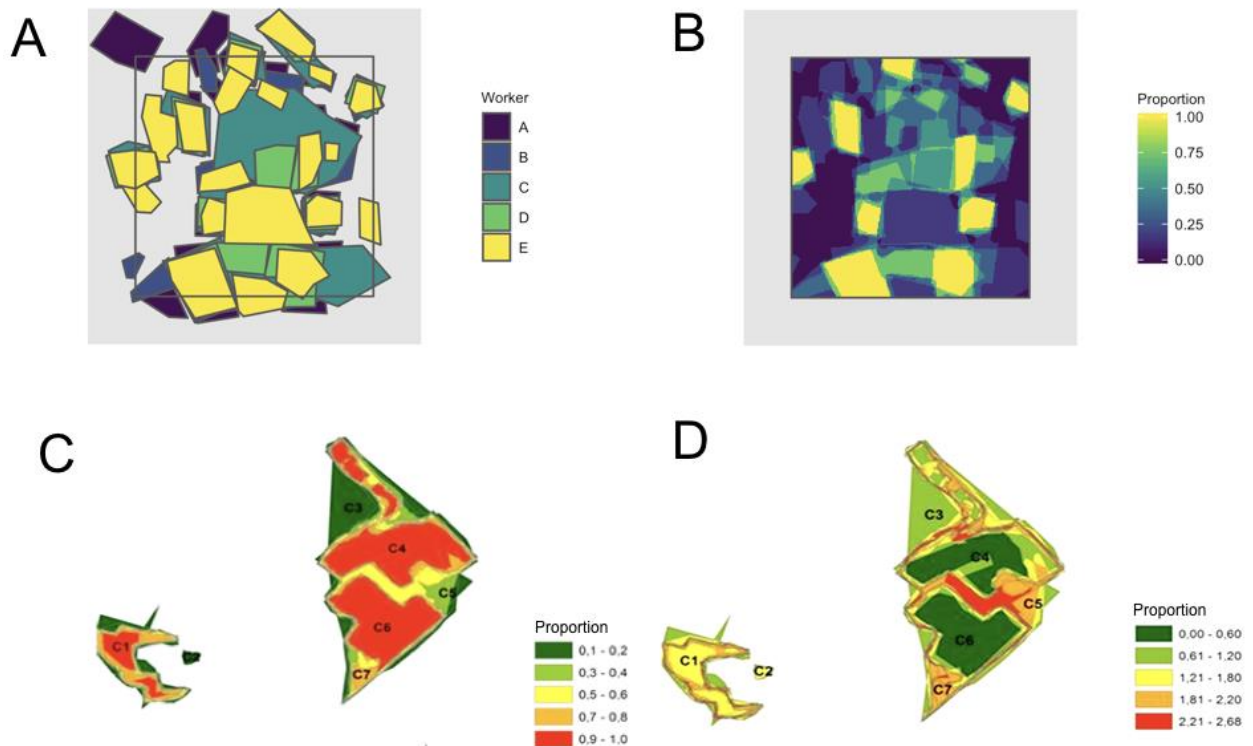|	
Training data set name	
How data were created (technical details)	
Funding source	
Purpose	
LULC definitions	
Time period	
Spatial extent	
How it was collected (samples of field sheets)	
Spatial resolution (image, field, quadrat, point location)	
Image ID	
Date and time of creation	

Figure 12: Two examples of consensus-based mapping approaches and their potential use for assessing training (or reference) data uncertainty. Panel A shows a collection of crop field boundary polygons drawn by five independent workers around crop fields visible in PlanetScope imagery collected over Ghana. These labels can be converted into a heat map (B) showing the overall agreement, the inverse of uncertainty. Similarly, 19 independent experts were asked to delineate slum settlements in image subset from Cape Town, South Africa. The polygons are converted into overall agreement and the uncertainty is modeled using random sets - (C) shows the covering function, which is then used to calculate standard deviation of random set (D). Both these metrics indicate the variability as well as stability in boundaries delineated by different experts. Adapted with permission from Kohli et al. (2016).

Step 5. Evaluate and communicate the impact of training data error

Due to the wide range of remote sensing research currently underway, a wide variety of training data and classification algorithms are in use. Therefore, it is not possible to specify a single protocol for treatment of training data error. Instead, we outline three tiers that represent different levels of accounting for the impact of training data errors on resulting map products.

Tier 1

The optimal training data accuracy assessment, termed Tier 1, involves comparison of training data to gold standard training reference data, representing truth. While the term 'truth' may entail

an essentially unattainable standard of universal veracity, we suggest that expert consensus, as discussed above, provides a practical benchmark, as well as defining an upper bound on the knowable map accuracy. Using such training reference data, researchers can quantify training data error directly, and thus incorporate it in estimates of output map variance and bias, for example using methods presented by McRoberts et al. [33]. For quantitative variables in which there are substantial training data errors, we recommend the use of type 2 regression, which considers error in both the remotely sensed measurement and the training data, as opposed to type 1 regression, which ignores potential error in the training data. Alternatively, absent such an error metric that explicitly factors training error into its formula, the impact of these errors can be quantified by repeatedly training the model with training datasets adjusted to represent the distributions of measured errors within the training sample, and quantify the resulting differences in the final map accuracy metrics (as in the cropland case study, section 4.3). This approach is likely the most useful in cases where training data are taken from large, pre-existing stock or benchmark libraries.

Tier 2

If it is not possible to directly measure and quantify error in training data, the next best course of action is to introduce a plausible range of simulated error into training data and evaluate its impact on the accuracy of maps trained with these perturbed datasets. If multiple workers are tasked with collecting training data for the same site, then the variance in their data can be used to estimate the uncertainty bounds (e.g. Figure 12). This approach is illustrated in the building mapping case study (section 4.1.1), which illustrates the sensitivity of key accuracy metrics to two different kinds of simulated labelling errors. The wheat yield case study (see section 4.3) provide an example of this approach for a continuous variable. Benchmark datasets may help facilitate this type of training data error impacts analysis, by providing a common set of training labels with known characteristics that can be systematically modified.

Tier 3

The bare minimum practice for training error accounting should be for researchers to release their training dataset during publication of their dataset, assuming it is not an existing published one, as has been recommended (Stehman and Foody, 2019) and sometimes done for reference data [e.g., 42]. These data should be documented with standard metadata, as shown in Table 2, and should also contain a description of the potential uncertainties associated with it. For example, the SpatioTemporal Asset Catalog (STAC)[11] provides a framework for standardization of metadata for EO data, and is increasingly seen as an international standard for geospatial data[12]. Additionally, we recommend that researchers use and present the relevant accuracy measure, determined by the study design and goals, and describe clearly that the reported

---

[11] https://stacspec.org/

[12] https://github.com/radiantearth/stac-spec/tree/dev/extensions/label

accuracy is itself uncertain, due to possible errors in both training and reference data. These steps allow for scientific replication of research, and would also help build an extensive training data repository.

Communication of training data error

Finally, uncertainty that has propagated through the analysis from training data should be faithfully reported in any reporting documentation and maps. We advise that researchers consider the broader communication of consequences of training data error and map accuracy, in a time when EO information products are increasingly used by the public and policy domains. Incomplete error reporting serves to limit to the scientific validity and usefulness of these products [47].

It is critical to consider the map's intended audience in order to present uncertainty as clearly and usefully as possible. While we advise to always provide full accounting for error sources along with the final accuracy metrics and qualification, it is worth considering carefully how this critical information could best be crafted and presented such that non-specialist audiences will be able to understand the map and its limitations. In general, we recommend including the error on or as close to the actual map whenever possible, whether by means of metrics, the error matrix, and/or by using cartographic techniques for representing uncertainty. Examples of effective cartographic techniques for conveying uncertainty include selection of appropriate, intuitive, and color-blind friendly color schemes for classes and symbols, varying color value and saturation and font/line weight to indicate levels of uncertainty, use of crisp versus blurred boundaries and symbols to indicate the range of uncertainty, or display of consensus maps or side-by-side juxtaposition in cases of multiple, mutually exclusive predictions for the same place and time have been made (e.g. representing differently specified models) [29,30]. Maps of consensus of training polygons and output classes or segments can provide valuable uncertainty information to users, such as shown in Figure 12A-B.

Towards an Open Training Data Repository

For the scientific community, the ideal standard of openness and replicability is to provide a complete description of training data collection practices, appropriate accuracy metrics, and perhaps most importantly of all, the raw data. Ideally, we recommend the creation of a centralized, open source database of all available and relevant training data, using the details collected in the proposed template (Table 2), and drawing inspiration from projects such as STAC. This type of open repository, taking inspiration from similar large-scale databases for computer vision [ImageNet, 164,SIFT10M Dataset, 165], and remote sensing [DeepSat, 166,UC Merced Land Use Dataset, 167], should contain full training metadata, citations to the peer-reviewed literature, as well as links to downloadable versions of training data collection protocols. Following the philosophy of free and open source software, we strongly recommend

that researchers embrace open source data, which is the only way by which a study can be truly reproduced.

6. Conclusions

The current treatment of training data in remote sensing research is insufficient for the purposes of reproducibility, uncertainty estimation, and communication of results. This is particularly problematic for error propagation to higher level information products, in which error originating in supposedly gold standard training data may have difficult to predict effects. The purpose of this paper is to call attention to this issue and promote the recommendations outlined above.

It is important to distinguish between the types of 'truth' data that may be used during an EO based mapping project. Here, we refer to four categories of such data: training, validation, training reference and map reference data, of which the latter has the most stringent sampling design requirements [25,82]. The independence of map reference data must be preserved, meaning that whether or not it is collected simultaneously with training data, it must be reserved for final map accuracy analysis. Any data used to enhance or develop ML model parameters (i.e. both training and validation data), must be treated separately from map reference data. In general our advice is applicable to both map reference and training data, but our recommendations specifically focus on the latter. For a comprehensive review of protocols for reference data collection, we refer the reader to Stehman and Foody [47].

To account for error in training data, we recommend the following steps. Researchers should carefully consider the tolerable levels of error in the desired map before collecting or creating training or reference data. This exercise involves translating the needs of the map users into quantifiably comparable metrics, and will guide the selection of these metrics. Next, we recommend that researchers strive to minimize error originating from both design- and collection-related errors; these include appropriate legend definition and imagery selection based on spatial resolution, spatiotemporal representativeness, consensus-based labeling strategies, and interactive feedback to interpreters regarding accuracy. Because it is not possible to completely eliminate training data error, we strongly advise that such error is incorporated in model outputs, either directly in bias and variance estimates or, if such incorporation is not possible, by documenting the sources and implications of error. As a minimum standard, training data should be fully documented and should be made openly available, allowing others to replicate and assess its use. To guide researchers in this process, we have proposed three tiers of training data error accounting standards, which we feel may provide a common basis for comparison of the ability to account for training data error. Finally, we advise that researchers strive to clearly communicate the magnitude and impacts of training data error on map outputs, with specific consideration to the intended or likely audience and users of the map.

Acknowledgements

References

1. Chen, J.; Chen, J.; Liao, A.; Cao, X.; Chen, L.; Chen, X.; He, C.; Han, G.; Peng, S.; Lu, M.; et al. Global Land Cover Mapping at 30 m Resolution: A POK-Based Operational Approach. *ISPRS J. Photogramm. Remote Sens.* **2015**, *103*, 7–27.
2. Friedl, M.A.; Sulla-Menashe, D.; Tan, B.; Schneider, A.; Ramankutty, N.; Sibley, A.; Huang, X. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets. *Remote Sens. Environ.* **2010**, *114*, 168–182.
3. Song, X.-P.; Hansen, M.C.; Stehman, S.V.; Potapov, P.V.; Tyukavina, A.; Vermote, E.F.; Townshend, J.R. Global land change from 1982 to 2016. *Nature* **2018**, *560*, 639–643.
4. Mohanty, B.P.; Cosh, M.H.; Lakshmi, V.; Montzka, C. Soil Moisture Remote Sensing: State-of-the-Science. *Vadose Zone J.* **2017**, *16*.
5. Daudt, R.C.; Le Saux, B.; Boulch, A.; Gousseau, Y. Guided Anisotropic Diffusion and Iterative Learning for Weakly Supervised Change Detection. *arXiv [cs.CV]* 2019.
6. Hecht, R.; Meinel, G.; Buchroithner, M. Automatic identification of building types based on topographic databases – a comparison of different data sources. *International Journal of Cartography* **2015**, *1*, 18–31.
7. Zhang, X.; Jayavelu, S.; Liu, L.; Friedl, M.A.; Henebry, G.M.; Liu, Y.; Schaaf, C.B.; Richardson, A.D.; Gray, J. Evaluation of land surface phenology from VIIRS data using time series of PhenoCam imagery. *Agric. For. Meteorol.* **2018**, *256-257*, 137–149.
8. Tan, B.; Morisette, J.T.; Wolfe, R.E.; Gao, F.; Ederer, G.A.; Nightingale, J.; Pedelty, J.A. An Enhanced TIMESAT Algorithm for Estimating Vegetation Phenology Metrics From MODIS Data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2011**, *4*, 361–371.
9. Zhang, X.; Friedl, M.A.; Schaaf, C.B. Global vegetation phenology from Moderate Resolution Imaging Spectroradiometer (MODIS): Evaluation of global patterns and comparison with in situ measurements: GLOBAL PHENOLOGY FROM MODIS. *J. Geophys. Res.* **2006**, *111*, 981.
10. Wan, Z. New refinements and validation of the MODIS Land-Surface Temperature/Emissivity products. *Remote Sens. Environ.* **2008**, *112*, 59–74.
11. Jiménez-Muñoz, J.C.; Sobrino, J.A.; Skoković, D.; Mattar, C.; Cristóbal, J. Land Surface Temperature Retrieval Methods From Landsat-8 Thermal Infrared Sensor Data. *IEEE Geoscience and Remote Sensing Letters* **2014**, *11*, 1840–1843.
12. Jean, N.; Burke, M.; Xie, M.; Davis, W.M.; Lobell, D.B.; Ermon, S. Combining satellite imagery and machine learning to predict poverty. *Science* **2016**, *353*, 790–794.
13. Pekel, J.-F.; Cottam, A.; Gorelick, N.; Belward, A.S. High-resolution mapping of global surface water and its long-term changes. *Nature* **2016**, *540*, 418–422.
14. Hansen, M.C.; Potapov, P.; Tyukavina, A. Comment on "Tropical forests are a net carbon source based on aboveground measurements of gain and loss." *Science* **2019**, *363*.
15. Gutierrez-Velez, V.H.; Pontius, R.G. Influence of carbon mapping and land change modelling on the prediction of carbon emissions from deforestation. *Environ. Conserv.* **2012**, *39*, 325–336.
16. Loveland, T.R.; Reed, B.C.; Brown, J.F.; Ohlen, D.O.; Zhu, Z.; Yang, L.; Merchant, J.W. Development of a global land cover characteristics database and IGBP DISCover from 1 km AVHRR data. *Int. J. Remote Sens.* **2000**, *21*, 1303–1330.
17. Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE conference on computer vision and pattern recognition; Ieee, 2009; pp. 248–255.
18. Helber, P.; Bischke, B.; Dengel, A.; Borth, D. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2019**, 1–10.

19. Liu, Q.; Hang, R.; Song, H.; Li, Z. Learning Multiscale Deep Features for High-Resolution Satellite Image Scene Classification. *IEEE Trans. Geosci. Remote Sens.* **2018**, *56*, 117–126.
20. Laso Bayas, J.C.; Lesiv, M.; Waldner, F.; Schucknecht, A.; Duerauer, M.; See, L.; Fritz, S.; Fraisl, D.; Moorthy, I.; McCallum, I.; et al. A global reference database of crowdsourced cropland data collected using the Geo-Wiki platform. *Sci Data* **2017**, *4*, 170136.
21. Fortier, J.; Rogan, J.; Woodcock, C.E.; Runfola, D.M. Utilizing Temporally Invariant Calibration Sites to Classify Multiple Dates and Types of Satellite Imagery. *Photogrammetric Engineering & Remote Sensing* **2011**, *77*, 181–189.
22. Pengra, B.W.; Stehman, S.V.; Horton, J.A.; Dockter, D.J.; Schroeder, T.A.; Yang, Z.; Cohen, W.B.; Healey, S.P.; Loveland, T.R. Quality control and assessment of interpreter consistency of annual land cover reference data in an operational national monitoring program. *Remote Sens. Environ.* **2019**, 111261.
23. Zhu, X.X.; Tuia, D.; Mou, L.; Xia, G.; Zhang, L.; Xu, F.; Fraundorfer, F. Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources. *IEEE Geoscience and Remote Sensing Magazine* **2017**, *5*, 8–36.
24. Foody, G.M. Status of land cover classification accuracy assessment. *Remote Sens. Environ.* **2002**, *80*, 185–201.
25. Foody, G.M. Assessing the accuracy of land cover change with imperfect ground reference data. *Remote Sensing of Environment* 2010, *114*, 2271–2285.
26. Olofsson, P.; Foody, G.M.; Herold, M.; Stehman, S.V.; Woodcock, C.E.; Wulder, M.A. Good practices for estimating area and assessing accuracy of land change. *Remote Sens. Environ.* **2014**, *148*, 42–57.
27. Pontius, R.G.; Millones, M. Death to Kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *Int. J. Remote Sens.* **2011**, *32*, 4407–4429.
28. Congalton, R.G.; Green, K. *Assessing the accuracy of remotely sensed data: principles and practices*; CRC press, 2008;.
29. Monmonier, M. Cartography: uncertainty, interventions, and dynamic display. *Prog. Hum. Geogr.* **2006**, *30*, 373–381.
30. MacEachren, A.M. Visualizing Uncertain Information. *1* **1992**, 10–19.
31. Goodchild, M.F.; Gopal, S. *The Accuracy Of Spatial Databases*; CRC Press, 1989; ISBN 9780203490235.
32. Congalton, R.G. A review of assessing the accuracy of classifications of remotely sensed data. *Remote Sens. Environ.* **1991**, *37*, 35–46.
33. McRoberts, R.E.; Stehman, S.V.; Liknes, G.C.; Næsset, E.; Sannier, C.; Walters, B.F. The effects of imperfect reference data on remote sensing-assisted estimators of land cover class proportions. *ISPRS J. Photogramm. Remote Sens.* **2018**, *142*, 292–300.
34. Carlotto, M.J. Effect of errors in ground truth on classification accuracy. *Int. J. Remote Sens.* **2009**, *30*, 4831–4849.
35. Franklin, J.; Simons, D.K.; Beardsley, D.; Rogan, J.M.; Gordon, H. Evaluating Errors in a Digital Vegetation Map with Forest Inventory Data and Accuracy Assessment Using Fuzzy Sets. *Trans. GIS* **2001**, *5*, 285–304.
36. Estes, L.; Chen, P.; Debats, S.; Evans, T.; Ferreira, S.; Kuemmerle, T.; Ragazzo, G.; Sheffield, J.; Wolf, A.; Wood, E.; et al. A large-area, spatially continuous assessment of land cover map error and its impact on downstream analyses. *Glob. Chang. Biol.* **2018**, *24*, 322–337.
37. Fritz, S.; See, L.; McCallum, I.; Schill, C.; Obersteiner, M.; van der Velde, M.; Boettcher, H.; Havlík, P.; Achard, F. Highlighting continued uncertainty in global land cover maps for the user community. *Environmental Research Letters* 2011, *6*, 044005.
38. Copass, C.; Antonova, N.; Kennedy, R. Comparison of Office and Field Techniques for Validating Landscape Change Classification in Pacific Northwest National Parks. *Remote Sensing* **2018**, *11*, 3.

39. Lesiv, M.; See, L.; Laso Bayas, J.C.; Sturn, T.; Schepaschenko, D.; Karner, M.; Moorthy, I.; McCallum, I.; Fritz, S. Characterizing the Spatial and Temporal Availability of Very High Resolution Satellite Imagery in Google Earth and Microsoft Bing Maps as a Source of Reference Data. *Land* **2018**, *7*, 118.

40. Biradar, C.M.; Thenkabail, P.S.; Noojipady, P.; Li, Y.; Dheeravath, V.; Turral, H.; Velpuri, M.; Gumma, M.K.; Gangalakunta, O.R.P.; Cai, X.L.; et al. A global map of rainfed cropland areas (GMRCA) at the end of last millennium using remote sensing. *Int. J. Appl. Earth Obs. Geoinf.* **2009**, *11*, 114–129.

41. Sulla-Menashe, D.; Gray, J.M.; Abercrombie, S.P.; Friedl, M.A. Hierarchical mapping of annual global land cover 2001 to present: The MODIS Collection 6 Land Cover product. *Remote Sens. Environ.* **2019**, *222*, 183–194.

42. Xiong, J.; Thenkabail, P.S.; Tilton, J.C.; Gumma, M.K.; Teluguntla, P.; Oliphant, A.; Congalton, R.G.; Yadav, K.; Gorelick, N. Nominal 30-m Cropland Extent Map of Continental Africa by Integrating Pixel-Based and Object-Based Algorithms Using Sentinel-2 and Landsat-8 Data on Google Earth Engine. *Remote Sensing* **2017**, *9*, 1065.

43. Mallinis, G.; Emmanoloudis, D.; Giannakopoulos, V.; Maris, F.; Koutsias, N. Mapping and interpreting historical land cover/land use changes in a Natura 2000 site using earth observational data: The case of Nestos delta, Greece. *Appl. Geogr.* **2011**, *31*, 312–320.

44. Jawak, S.D.; Luis, A.J. Improved land cover mapping using high resolution multiangle 8-band WorldView-2 satellite remote sensing data. *JARS* **2013**, *7*, 073573.

45. Ye, S.; Pontius, R.G.; Rakshit, R. A review of accuracy assessment for object-based image analysis: From per-pixel to per-polygon approaches. *ISPRS J. Photogramm. Remote Sens.* **2018**, *141*, 137–147.

46. Fritz, S.; See, L.; Perger, C.; McCallum, I.; Schill, C.; Schepaschenko, D.; Duerauer, M.; Karner, M.; Dresel, C.; Laso-Bayas, J.-C.; et al. A global dataset of crowdsourced land cover and land use reference data. *Sci Data* **2017**, *4*, 170075.

47. Stehman, S.V.; Foody, G.M. Key issues in rigorous accuracy assessment of land cover products. *Remote Sens. Environ.* **2019**, *231*, 111199.

48. Stehman, S.V. Sampling designs for accuracy assessment of land cover. *Int. J. Remote Sens.* **2009**, *30*, 5243–5272.

49. Brodrick, P.G.; Davies, A.B.; Asner, G.P. Uncovering Ecological Patterns with Convolutional Neural Networks. *Trends Ecol. Evol.* **2019**, *34*, 734–745.

50. Xiao, T.; Xia, T.; Yang, Y.; Huang, C.; Wang, X. Learning from massive noisy labeled data for image classification. In Proceedings of the Proceedings of the IEEE conference on computer vision and pattern recognition; cv-foundation.org, 2015; pp. 2691–2699.

51. Frénay, B.; Verleysen, M. Classification in the presence of label noise: a survey. *IEEE Trans Neural Netw Learn Syst* **2014**, *25*, 845–869.

52. Brodley, C.E.; Friedl, M.A. Identifying Mislabeled Training Data. *1* **1999**, *11*, 131–167.

53. Lesiv, M.; Laso Bayas, J.C.; See, L.; Duerauer, M.; Dahlia, D.; Durando, N.; Hazarika, R.; Kumar Sahariah, P.; Vakolyuk, M. 'yana; Blyshchyk, V.; et al. Estimating the global distribution of field size using crowdsourcing. *Glob. Chang. Biol.* **2019**, *25*, 174–186.

54. Fritz, S.; McCallum, I.; Schill, C.; Perger, C.; See, L.; Schepaschenko, D.; van der Velde, M.; Kraxner, F.; Obersteiner, M. Geo-Wiki: An Online Platform for Improving Global Land Cover. *Environmental Modelling & Software* **2012**, *31*, 110–123.

55. Goodchild, M.F. Citizens as sensors: the world of volunteered geography. *GeoJournal* **2007**, *69*, 211–221.

56. Van Etten, A.; Lindenbaum, D.; Bacastow, T.M. SpaceNet: A Remote Sensing Dataset and Challenge Series. *arXiv [cs.CV]* 2018.

57. Sumbul, G.; Charfuelan, M.; Demir, B.; Markl, V. BigEarthNet: A Large-Scale Benchmark Archive

For Remote Sensing Image Understanding. *arXiv [cs.CV]* 2019.

58. Kohavi, R.; Others A study of cross-validation and bootstrap for accuracy estimation and model selection. In Proceedings of the Ijcai; Montreal, Canada, 1995; Vol. 14, pp. 1137–1145.
59. Stehman, S.V.; Czaplewski, R.L. Design and Analysis for Thematic Map Accuracy Assessment: Fundamental Principles. *Remote Sens. Environ.* **1998**, *64*, 331–344.
60. Foody, G.M. Ground reference data error and the mis-estimation of the area of land cover change as a function of its abundance. *Remote Sens. Lett.* **2013**, *4*, 783–792.
61. Stehman, S.V. Practical Implications of Design-Based Sampling Inference for Thematic Map Accuracy Assessment. *Remote Sens. Environ.* **2000**, *72*, 35–45.
62. Kuzera, K.; Pontius, R.G. Importance of Matrix Construction for Multiple-Resolution Categorical Map Comparison. *GISci. Remote Sens.* **2008**, *45*, 249–274.
63. Pontius, R.G.; Thontteh, O.; Chen, H. Components of information for multiple resolution comparison between maps that share a real variable. *Environ. Ecol. Stat.* **2008**, *15*, 111–142.
64. Pontius, R.G.; Parmentier, B. Recommendations for using the relative operating characteristic (ROC). *Landsc. Ecol.* **2014**, *29*, 367–382.
65. Pontius, R.G. Component intensities to relate difference by category with difference overall. *Int. J. Appl. Earth Obs. Geoinf.* **2019**, *77*, 94–99.
66. Pontius, R.G., Jr.; Connors, J. Range of Categorical Associations for Comparison of Maps with Mixed Pixels. *Photogrammetric Engineering & Remote Sensing* **2009**, *75*, 963–969.
67. Aldwaik, S.Z.; Pontius, R.G., Jr. Intensity analysis to unify measurements of size and stationarity of land changes by interval, category, and transition. *Landsc. Urban Plan.* **2012**, *106*, 103–114.
68. Pontius, R.G.; Gao, Y.; Giner, N.M.; Kohyama, T.; Osaki, M.; Hirose, K. Design and Interpretation of Intensity Analysis Illustrated by Land Change in Central Kalimantan, Indonesia. *Land* **2013**, *2*, 351–369.
69. Foody, G.M. Harshness in image classification accuracy assessment. *Int. J. Remote Sens.* **2008**, *29*, 3137–3158.
70. Cohen, J. A Coefficient of Agreement for Nominal Scales. *Educ. Psychol. Meas.* **1960**, *20*, 37–46.
71. Allouche, O.; Tsoar, A.; Kadmon, R. Assessing the Accuracy of Species Distribution Models: Prevalence, Kappa and the True Skill Statistic (TSS). *J. Appl. Ecol.* **2006**, *43*, 1223–1232.
72. Blaschke, T. Object based image analysis for remote sensing. *ISPRS J. Photogramm. Remote Sens.* **2010**, *65*, 2–16.
73. Garcia-Garcia, A.; Orts-Escolano, S.; Oprea, S.; Villena-Martinez, V.; Garcia-Rodriguez, J. A Review on Deep Learning Techniques Applied to Semantic Segmentation. *arXiv [cs.CV]* 2017.
74. Willmott, C.J.; Matsuura, K. On the use of dimensioned measures of error to evaluate the performance of spatial interpolators. *Int. J. Geogr. Inf. Sci.* **2006**, *20*, 89–102.
75. Willmott, C.J.; Matsuura, K.; Robeson, S.M. Ambiguities inherent in sums-of-squares-based error statistics. *Atmos. Environ.* **2009**, *43*, 749–752.
76. Willmott, C.J.; Matsuura, K. Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Clim. Res.* **2005**, *30*, 79–82.
77. Pontius, R.G., Jr; Si, K. The total operating characteristic to measure diagnostic ability for multiple thresholds. *Int. J. Geogr. Inf. Sci.* **2014**, *28*, 570–583.
78. Fielding, A.H.; Bell, J.F. A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environ. Conserv.* **1997**, *24*, 38–49.
79. Weaver, J.; Moore, B.; Reith, A.; McKee, J.; Lunga, D. A Comparison of Machine Learning Techniques to Extract Human Settlements from High Resolution Imagery. In Proceedings of the IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium; ieeexplore.ieee.org, 2018; pp. 6412–6415.
80. Stehman, S.V. Statistical rigor and practical utility in thematic map accuracy assessment. *Photogramm. Eng. Remote Sens.* **2001**, *67*, 727–734.

81. Rogan, J.; Mietkiewicz, N. Land cover change detection. *Land Resources Monitoring, Modeling, and Mapping with Remote Sensing; Thenkabail, PS, Ed* **2015**, 579–603.

82. Olofsson, P.; Foody, G.M.; Stehman, S.V.; Woodcock, C.E. Making better use of accuracy data in land change studies: Estimating accuracy and area and quantifying uncertainty using stratified estimation. *Remote Sens. Environ.* **2013**, *129*, 122–131.

83. Swan, B.; Laverdiere, M.; Yang, H.L. How Good is Good Enough?: Quantifying the Effects of Training Set Quality. In Proceedings of the Proceedings of the 2Nd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery; ACM: New York, NY, USA, 2018; pp. 47–51.

84. Powell, R.L.; Matzke, N.; de Souza, C.; Clark, M.; Numata, I.; Hess, L.L.; Roberts, D.A. Sources of error in accuracy assessment of thematic land-cover maps in the Brazilian Amazon. *Remote Sens. Environ.* **2004**, *90*, 221–234.

85. Clark, M.L.; Aide, T.M.; Riner, G. Land change for all municipalities in Latin America and the Caribbean assessed from 250-m MODIS imagery (2001–2010). *Remote Sens. Environ.* **2012**, *126*, 84–103.

86. Planet Team Planet Application Program Interface: In Space for Life on Earth. San Francisco, CA 2017.

87. Manfreda, S.; McCabe, M.F.; Miller, P.E.; Lucas, R.; Pajuelo Madrigal, V.; Mallinis, G.; Ben Dor, E.; Helman, D.; Estes, L.; Ciraolo, G.; et al. On the Use of Unmanned Aerial Systems for Environmental Monitoring. *Remote Sensing* **2018**, *10*, 641.

88. Toutin, T. Geometric processing of IKONOS Geo images with DEM. In Proceedings of the Proceedings of ISPRS Joint Workshop "High Resolution Mapping from Space" 2001; pdfs.semanticscholar.org, 2001; pp. 19–21.

89. Reinartz, P.; Müller, R.; Schwind, P.; Suri, S.; Bamler, R. Orthorectification of VHR optical satellite data exploiting the geometric accuracy of TerraSAR-X data. *ISPRS J. Photogramm. Remote Sens.* **2011**, *66*, 124–132.

90. Aguilar, M.A.; Saldaña, M. del M.; Aguilar, F.J. Assessing geometric accuracy of the orthorectification process from GeoEye-1 and WorldView-2 panchromatic images. *Int. J. Appl. Earth Obs. Geoinf.* **2013**, *21*, 427–435.

91. Chen, J.; Zipf, A. DeepVGI: Deep Learning with Volunteered Geographic Information. In Proceedings of the Proceedings of the 26th International Conference on World Wide Web Companion; International World Wide Web Conferences Steering Committee: Republic and Canton of Geneva, Switzerland, 2017; pp. 771–772.

92. Kaiser, P.; Wegner, J.D.; Lucchi, A.; Jaggi, M.; Hofmann, T.; Schindler, K. Learning Aerial Image Segmentation From Online Maps. *IEEE Trans. Geosci. Remote Sens.* **2017**, *55*, 6054–6068.

93. Audebert, N.; Le Saux, B.; Lefèvre, S. Joint learning from earth observation and openstreetmap data to get faster better semantic maps. In Proceedings of the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops; 2017; pp. 67–75.

94. Tewkesbury, A.P.; Comber, A.J.; Tate, N.J.; Lamb, A.; Fisher, P.F. A critical synthesis of remotely sensed optical image change detection techniques. *Remote Sens. Environ.* **2015**, *160*, 1–14.

95. Stehman, S.V.; Fonte, C.C.; Foody, G.M.; See, L. Using volunteered geographic information (VGI) in design-based statistical inference for area estimation and accuracy assessment of land cover. *Remote Sens. Environ.* **2018**, *212*, 47–59.

96. Waldner, F.; Schucknecht, A.; Lesiv, M.; Gallego, J.; See, L.; Pérez-Hoyos, A.; d'Andrimont, R.; de Maet, T.; Bayas, J.C.L.; Fritz, S.; et al. Conflation of expert and crowd reference data to validate global binary thematic maps. *Remote Sens. Environ.* **2019**, *221*, 235–246.

97. Strahler, A.H.; Woodcock, C.E.; Smith, J.A. On the nature of models in remote sensing. *Remote Sens. Environ.* **1986**, *20*, 121–139.

98. Shao, Y.; Lunetta, R.S. Comparison of support vector machine, neural network, and CART

algorithms for the land-cover classification using limited training data points. *ISPRS J. Photogramm. Remote Sens.* **2012**, *70*, 78–87.

99. Foody, G.M. Relating the land-cover composition of mixed pixels to artificial neural network classification output. *Photogramm. Eng. Remote Sens.* **1996**, *62*, 491–498.

100. Moody, A.; Gopal, S.; Strahler, A.H. Artificial neural network response to mixed pixels in coarse-resolution satellite data. *Remote Sens. Environ.* **1996**, *58*, 329–343.

101. Oppenshaw, S.; Taylor, P. A million or so correlation coefficients. *Statistical methods in the spatial sciences. Pion, London* **1979**.

102. Jelinski, D.E.; Wu, J. The modifiable areal unit problem and implications for landscape ecology. *Landsc. Ecol.* **1996**, *11*, 129–140.

103. Weiss, M.; de Beaufort, L.; Baret, F.; Allard, D.; Bruguier, N.; Marloie, O. Mapping leaf area index measurements at different scales for the validation of large swath satellite sensors: first results of the VALERI project. In Proceedings of the 8th International symposium in physical measurements and remote sensing, Aussois (France); w3.avignon.inra.fr, 2001; pp. 125–130.

104. Tian, Y.; Woodcock, C.E.; Wang, Y.; Privette, J.L.; Shabanov, N.V.; Zhou, L.; Zhang, Y.; Buermann, W.; Dong, J.; Veikkanen, B.; et al. Multiscale analysis and validation of the MODIS LAI product: I. Uncertainty assessment. *Remote Sens. Environ.* **2002**, *83*, 414–430.

105. Masuoka, E.; Roy, D.; Wolfe, R.; Morisette, J.; Sinno, S.; Teague, M.; Saleous, N.; Devadiga, S.; Justice, C.O.; Nickeson, J. MODIS Land Data Products: Generation, Quality Assurance and Validation. In *Land Remote Sensing and Global Environmental Change: NASA's Earth Observing System and the Science of ASTER and MODIS*; Ramachandran, B., Justice, C.O., Abrams, M.J., Eds.; Springer New York: New York, NY, 2011; pp. 509–531 ISBN 9781441967497.

106. Cohen, W.B.; Justice, C.O. Validating MODIS terrestrial ecology products: linking in situ and satellite measurements. *Remote Sens. Environ.* **1999**, *70*, 1–3.

107. Maxwell, A.E.; Warner, T.A.; Fang, F. Implementation of machine-learning classification in remote sensing: an applied review. *Int. J. Remote Sens.* **2018**, *39*, 2784–2817.

108. Debats, S.R.; Estes, L.D.; Thompson, D.R.; Caylor, K.K. *Integrating active learning and crowdsourcing into large-scale supervised landcover mapping algorithms*; PeerJ Preprints, 2017;.

109. Estes, L.D.; McRitchie, D.; Choi, J.; Debats, S.; Evans, T.; Guthe, W.; Luo, D.; Ragazzo, G.; Zempleni, R.; Caylor, K.K. A Platform for Crowdsourcing the Creation of Representative, Accurate Landcover Maps. *Environmental Modelling & Software* **2016**, *80*, 41–53.

110. Neigh, C.S.R.; Carroll, M.L.; Wooten, M.R.; McCarty, J.L.; Powell, B.F.; Husak, G.J.; Enenkel, M.; Hain, C.R. Smallholder crop area mapped with wall-to-wall WorldView sub-meter panchromatic image texture: A test case for Tigray, Ethiopia. *Remote Sens. Environ.* **2018**, *212*, 8–20.

111. Comber, A.; Fisher, P. What is land cover? *Environment and Planning* **2005**.

112. Kohli, D.; Sliuzas, R.; Kerle, N.; Stein, A. An ontology of slums for image-based classification. *Comput. Environ. Urban Syst.* **2012**, *36*, 154–163.

113. Verburg, P.H.; Neumann, K.; Nol, L. Challenges in using land use and land cover data for global change studies. *Glob. Chang. Biol.* **2011**, *17*, 974–989.

114. Weng, Q. Remote sensing of impervious surfaces in the urban areas: Requirements, methods, and trends. *Remote Sens. Environ.* **2012**, *117*, 34–49.

115. Kohli, D.; Stein, A.; Sliuzas, R. Uncertainty analysis for image interpretations of urban slums. *Comput. Environ. Urban Syst.* **2016**, *60*, 37–49.

116. Rocchini, D. While Boolean sets non-gently rip: A theoretical framework on fuzzy sets for mapping landscape patterns. *Ecol. Complex.* **2010**, *7*, 125–129.

117. Woodcock, C.E.; Gopal, S. Fuzzy set theory and thematic maps: accuracy assessment and area estimation. *Int. J. Geogr. Inf. Sci.* **2000**, *14*, 153–172.

118. Rocchini, D.; Foody, G.M.; Nagendra, H.; Ricotta, C.; Anand, M.; He, K.S.; Amici, V.; Kleinschmit, B.; Förster, M.; Schmidtlein, S.; et al. Uncertainty in ecosystem mapping by remote

sensing. *Comput. Geosci.* **2013**, *50*, 128–135.

119.   Zhang, J.; Foody, G.M. A fuzzy classification of sub-urban land cover from remotely sensed imagery. *Int. J. Remote Sens.* **1998**, *19*, 2721–2738.

120.   Woodcock, C.E.; Strahler, A.H. The factor of scale in remote sensing. *Remote Sens. Environ.* **1987**, *21*, 311–332.

121.   Foody, G.M. The continuum of classification fuzziness in thematic mapping. *Photogramm. Eng. Remote Sens.* **1999**, *65*, 443–452.

122.   Bey, A.; Sánchez-Paus Díaz, A.; Maniatis, D.; Marchi, G.; Mollicone, D.; Ricci, S.; Bastin, J.-F.; Moore, R.; Federici, S.; Rezende, M.; et al. Collect Earth: Land Use and Land Cover Assessment through Augmented Visual Interpretation. *Remote Sensing* **2016**, *8*, 807.

123.   Fritz, S.; Sturn, T.; Karner, M.; Moorthy, I.; See, L.; Laso Bayas, J.C.; Fraisl, D. FotoQuest Go: A Citizen Science Approach to the Collection of In-Situ Land Cover and Land Use Data for Calibration and Validation.; pure.iiasa.ac.at, 2019.

124.   Thompson, I.D.; Maher, S.C.; Rouillard, D.P.; Fryxell, J.M.; Baker, J.A. Accuracy of forest inventory mapping: Some implications for boreal forest management. *For. Ecol. Manage.* **2007**, *252*, 208–221.

125.   Bland, M.J.; Altman, D.G. Statistics notes: Measurement error. *BMJ* **1996**, *312*, 1654.

126.   Martin, D. An Introduction to "The Guide to the Expression of Uncertainty in Measurement." In *Evaluation of measurement data -- Guide to the expression of uncertainty in measurement*; JCGM, 2008; pp. 1–10.

127.   Foody, G.M. The impact of imperfect ground reference data on the accuracy of land cover change estimation. *Int. J. Remote Sens.* **2009**, *30*, 3275–3281.

128.   Tuia, D.; Pasolli, E.; Emery, W.J. Using active learning to adapt remote sensing image classifiers. *Remote Sensing of Environment* 2011, *115*, 2232–2242.

129.   Multi-Resolution Land Characteristics Consortium (U.S.) National land cover dataset (NLCD).

130.   Menon, S.; Akbari, H.; Mahanama, S.; Sednev, I.; Levinson, R. Radiative forcing and temperature response to changes in urban albedos and associated CO2 offsets. *Environ. Res. Lett.* **2010**, *5*, 014005.

131.   Hutyra, L.R.; Yoon, B.; Hepinstall-Cymerman, J.; Alberti, M. Carbon consequences of land cover change and expansion of urban lands: A case study in the Seattle metropolitan region. *Landsc. Urban Plan.* **2011**, *103*, 83–93.

132.   Reinmann, A.B.; Hutyra, L.R.; Trlica, A.; Olofsson, P. Assessing the global warming potential of human settlement expansion in a mesic temperate landscape from 2005 to 2050. *Sci. Total Environ.* **2016**, *545-546*, 512–524.

133.   Seto, K.C.; Güneralp, B.; Hutyra, L.R. Global forecasts of urban expansion to 2030 and direct impacts on biodiversity and carbon pools. *Proc. Natl. Acad. Sci. U. S. A.* **2012**, *109*, 16083–16088.

134.   Angel, S.; Parent, J.; Civco, D.L.; Blei, A.; Potere, D. The dimensions of global urban expansion: Estimates and projections for all countries, 2000–2050. *Prog. Plann.* **2011**, *75*, 53–107.

135.   Hardiman, B.S.; Wang, J.A.; Hutyra, L.R.; Gately, C.K.; Getson, J.M.; Friedl, M.A. Accounting for urban biogenic fluxes in regional carbon budgets. *Sci. Total Environ.* **2017**, *592*, 366–372.

136.   Coulston, J.W.; Moisen, G.G.; Wilson, B.T.; Finco, M.V.; Cohen, W.B.; Brewer, C.K. Modeling percent tree canopy cover: a pilot study. *Photogrammetric Engineering & Remote Sensing 78 (7): 715--727* **2012**, *78*, 715–727.

137.   Smith, M.L.; Zhou, W.; Cadenasso, M.; Grove, M.; Band, L.E. Evaluation of the national land cover database for hydrologic applications in urban and suburban Baltimore, Maryland 1. *JAWRA Journal of the American Water Resources Association* **2010**, *46*, 429–442.

138.   Reinmann, A.B.; Hutyra, L.R. Edge effects enhance carbon uptake and its vulnerability to climate change in temperate broadleaf forests. *Proc. Natl. Acad. Sci. U. S. A.* **2017**, *114*, 107–112.

139.   Rolnick, D.; Veit, A.; Belongie, S.; Shavit, N. Deep Learning is Robust to Massive Label Noise.

arXiv [cs.LG] 2017.

140. Nachmany, Y.; Alemohammad, H. Detecting Roads from Satellite Imagery in the Developing World. In Proceedings of the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops; openaccess.thecvf.com, 2019; pp. 83–89.

141. The SpaceNet Catalog SpaceNet on Amazon Web Services (AWS). "Datasets." 2018.

142. Alemohammad, S.H.; Fang, B.; Konings, A.G.; Aires, F.; Green, J.K.; Kolassa, J.; Miralles, D.; Prigent, C.; Gentine, P. Water, Energy, and Carbon with Artificial Neural Networks (WECANN): A statistically-based estimate of global surface turbulent fluxes and gross primary productivity using solar-induced fluorescence. *Biogeosciences* **2017**, *14*, 4101–4124.

143. McColl, K.A.; Vogelzang, J.; Konings, A.G.; Entekhabi, D.; Piles, M.; Stoffelen, A. Extended triple collocation: Estimating errors and correlation coefficients with respect to an unknown target. *Geophys. Res. Lett.* **2014**, *41*, 6229–6236.

144. Breiman, L. Random Forests. *Mach. Learn.* **2001**, *45*, 5–32.

145. Debats, S.R.; Luo, D.; Estes, L.D.; Fuchs, T.J.; Caylor, K.K. A Generalized Computer Vision Approach to Mapping Crop Fields in Heterogeneous Agricultural Landscapes. *Remote Sens. Environ.* **2016**, *179*, 210–221.

146. Jain, M., B. Singh, P. Rao, A.K. Srivastava, S. Poonia, G. Azzari, A.J. McDonald, D.B. Lobell Yield Gains from Sustainable Intensification can be Evaluated and Improved Using Satellite Data.

147. Pontius, R.G. Criteria to Confirm Models that Simulate Deforestation and Carbon Disturbance. *Land* **2018**, *7*, 105.

148. Lyndon D. Estes, Su Ye, Lei Song, Ron Eastman, Sitian Xiong, Tammy Woodard, Boka Luo, Dennis McRitchie, Ryan Avery, Kelly Caylor, Stephanie, Debats. Improving cropland maps through tight integration of human and machine intelligence. *In preparation*.

149. Waldner, F.; De Abelleyra, D.; Verón, S.R.; Zhang, M.; Wu, B.; Plotnikov, D.; Bartalev, S.; Lavreniuk, M.; Skakun, S.; Kussul, N.; et al. Towards a set of agrosystem-specific cropland mapping methods to address the global cropland diversity. *Int. J. Remote Sens.* **2016**, *37*, 3196–3231.

150. Castelluccio, M.; Poggi, G.; Sansone, C.; Verdoliva, L. Land Use Classification in Remote Sensing Images by Convolutional Neural Networks. *arXiv [cs.CV]* 2015.

151. Azevedo, T., Sr.; Souza, C.M., Jr.; Shimbo, J.; Alencar, A. MapBiomas initiative: Mapping annual land cover and land use changes in Brazil from 1985 to 2017.; adsabs.harvard.edu, 2018; Vol. 2018.

152. Brown, J.F.; Tollerud, H.J.; Barber, C.P.; Zhou, Q.; Dwyer, J.L.; Vogelmann, J.E.; Loveland, T.R.; Woodcock, C.E.; Stehman, S.V.; Zhu, Z.; et al. Lessons learned implementing an operational continuous United States national land change monitoring capability: The Land Change Monitoring, Assessment, and Projection (LCMAP) approach. *Remote Sens. Environ.* **2019**, 111356.

153. Estes, L.; Elsen, P.R.; Treuer, T.; Ahmed, L.; Caylor, K.; Chang, J.; Choi, J.J.; Ellis, E.C. The spatial and temporal domains of modern ecology. *Nat Ecol Evol* **2018**, *2*, 819–826.

154. Jensen, J.R.; Cowen, D.C. Remote sensing of urban/suburban infrastructure and socio-economic attributes. *Photogramm. Eng. Remote Sens.* **1999**, *65*, 611–622.

155. Rozenstein, O.; Karnieli, A. Comparison of methods for land-use classification incorporating remote sensing and GIS inputs. *Appl. Geogr.* **2011**, *31*, 533–544.

156. Dorais, A.; Cardille, J. Strategies for Incorporating High-Resolution Google Earth Databases to Guide and Validate Classifications: Understanding Deforestation in Borneo. *Remote Sensing* **2011**, *3*, 1157–1176.

157. Sexton, J.O.; Urban, D.L.; Donohue, M.J.; Song, C. Long-term land cover dynamics by multi-temporal classification across the Landsat-5 record. *Remote Sens. Environ.* **2013**, *128*, 246–258.

158. Carletto, C.; Gourlay, S.; Winters, P. From Guesstimates to GPStimates: Land Area Measurement and Implications for Agricultural Analysis. *J. Afr. Econ.* **2015**, *24*, 593–628.

159. Laso Bayas, J.C.; See, L.; Fritz, S.; Sturn, T.; Perger, C.; Dürauer, M.; Karner, M.; Moorthy, I.;

Schepaschenko, D.; Domian, D.; et al. Crowdsourcing In-Situ Data on Land Cover and Land Use Using Gamification and Mobile Technology. *Remote Sensing* **2016**, *8*, 905.

160. Burke, M.; Lobell, D.B. Satellite-based assessment of yield variation and its determinants in smallholder African systems. *Proc. Natl. Acad. Sci. U. S. A.* **2017**.

161. Jin, Z.; Azzari, G.; You, C.; Di Tommaso, S.; Aston, S.; Burke, M.; Lobell, D.B. Smallholder maize area and yield mapping at national scales with Google Earth Engine. *Remote Sens. Environ.* **2019**, *228*, 115–128.

162. Lobell, D.B.; Thau, D.; Seifert, C.; Engle, E.; Little, B. A Scalable Satellite-Based Crop Yield Mapper. *Remote Sens. Environ.* **2015**, *164*, 324–333.

163. Grassini, P.; van Bussel, L.G.J.; Van Wart, J.; Wolf, J.; Claessens, L.; Yang, H.; Boogaard, H.; de Groot, H.; van Ittersum, M.K.; Cassman, K.G. How Good Is Good Enough? Data Requirements for Reliable Crop Yield Simulations and Yield-Gap Analysis. *Field Crops Res.* **2015**, *177*, 49–63.

164. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* **2015**, *115*, 211–252.

165. Fu, X.; McCane, B.; Mills, S.; Albert, M. NOKMeans: Non-Orthogonal K-means Hashing. In *Computer Vision -- ACCV 2014*; Cremers, D., Reid, I., Saito, H., Yang, M.-H., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, 2015; Vol. 9003, pp. 162–177 ISBN 9783319168647.

166. Basu, S.; Ganguly, S.; Mukhopadhyay, S.; DiBiano, R.; Karki, M.; Nemani, R. DeepSat: A Learning Framework for Satellite Imagery. In Proceedings of the Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems; ACM: New York, NY, USA, 2015; pp. 37:1–37:10.

167. Yang, Y.; Newsam, S. Bag-of-visual-words and spatial extensions for land-use classification. *Proceedings of the 18th SIGSPATIAL international* **2010**.

Appendix

In this problem, TC-based RMSE estimates at each pixel were used to compute a priori probability (P_i) of selecting a particular dataset:

$$P_i = \frac{\dfrac{1}{\sigma_{\varepsilon_i}^2}}{\sum_{i=1}^{3} \dfrac{1}{\sigma_{\varepsilon_i}^2}}$$

(A1)

P_i is the probability of selecting measurement system i, $\sigma_{\varepsilon i}$ is the standard deviation of the random error in measurement system i.

Figure A1 depicts how X_T (the training time series for a pixel) is formed by sampling from X_1, X_2, and X_3 over time.


Figure A1: Schematic of product selection using the Triple Collocation approach.

Table A1 presents quantitative results of comparing each of the three models trained for the road detection case in Kumasi, Ghana to the validation labels. This region (shown in Figure 9) included 5,406,942 road pixels and 50,627,010 background pixels.

Table A1: Performance metrics of different models in Kumasi, Ghana

	F1	IOU	Precision	Recall
Khartoum Model				
Average	0.6659	0.5723	0.7758	0.6267
Road	0.3780	0.2330	0.6250	0.2709
Background	0.9538	0.9116	0.9266	0.9862
Kumasi Model				
Average	0.8004	0.6955	0.7693	0.8450
Road	0.6458	0.4769	0.5662	0.7513
Background	0.9552	0.9142	0.9725	0.9386
Khartoum Model retrained in Kumasi				
Average	0.7869	0.6830	0.7965	0.7780
Road	0.6135	0.4425	0.6363	0.5921
Background	0.9603	0.9236	0.9568	0.9639

Figure A2 shows a qualitative comparison of different model outputs along with the validation labels over a sample area of Figure 4.

(a)

(b)

(c)

(d)

(e)

Figure A2: Sample prediction results in Kumasi, Ghana. (a) Input imagery. (b) Predictions from the Las Vegas model. (c) Predictions from the Khartoum model. (d) Prediction from the Kumasi model. (e) Predictions from the Khartoum Model retrained in Kumasi. [in panel b-e model predictions are in shaded color overlaid with validation labels in red on top of imagery]