



MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets

Mark A. Friedl^{a,*}, Damien Sulla-Menashe^a, Bin Tan^b, Annemarie Schneider^c, Navin Ramankutty^d, Adam Sibley^a, Xiaoman Huang^a

^a Department of Geography and Environment, Boston University, 675 Commonwealth Avenue, Boston, MA 02215, USA

^b Earth Resources Technology, Inc., NASA Goddard Space Flight Center, Code 614.5, Greenbelt, MD 20771, USA

^c Center for Sustainability and the Global Environment, University of Wisconsin-Madison, 1710 University Avenue, Room 264, Madison, Wisconsin 53726 USA

^d Department of Geography and Program in Earth System Science, 627 Burnside Hall, 805 Sherbrooke Street W., Montreal, QC, Canada H3A 2K6

ARTICLE INFO

Article history:

Received 23 March 2009

Received in revised form 27 August 2009

Accepted 29 August 2009

Keywords:

Global land cover

MODIS

Classification

ABSTRACT

Information related to land cover is immensely important to global change science. In the past decade, data sources and methodologies for creating global land cover maps from remote sensing have evolved rapidly. Here we describe the datasets and algorithms used to create the Collection 5 MODIS Global Land Cover Type product, which is substantially changed relative to Collection 4. In addition to using updated input data, the algorithm and ancillary datasets used to produce the product have been refined. Most importantly, the Collection 5 product is generated at 500-m spatial resolution, providing a four-fold increase in spatial resolution relative to the previous version. In addition, many components of the classification algorithm have been changed. The training site database has been revised, land surface temperature is now included as an input feature, and ancillary datasets used in post-processing of ensemble decision tree results have been updated. Further, methods used to correct classifier results for bias imposed by training data properties have been refined, techniques used to fuse ancillary data based on spatially varying prior probabilities have been revised, and a variety of methods have been developed to address limitations of the algorithm for the urban, wetland, and deciduous needleleaf classes. Finally, techniques used to stabilize classification results across years have been developed and implemented to reduce year-to-year variation in land cover labels not associated with land cover change. Results from a cross-validation analysis indicate that the overall accuracy of the product is about 75% correctly classified, but that the range in class-specific accuracies is large. Comparison of Collection 5 maps with Collection 4 results show substantial differences arising from increased spatial resolution and changes in the input data and classification algorithm.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Global land cover maps provide thematic characterizations of the Earth's surface that capture biotic and abiotic properties and that are closely tied to the ecological condition of land areas. Because surface properties affect biosphere–atmosphere interaction, accurate land cover information is required to parameterize land surface processes in regional-to-global scale Earth system models (Bonan et al., 2002b; Ek et al., 2003; Running & Coughlan, 1988; Sellers et al., 1997; Sterling & Ducharme, 2008). Further, humans depend heavily on goods and services provided by terrestrial ecosystems (Foley et al., 2005), and the global area of land dominated by humans has expanded rapidly in the last 100 years (Ellis & Ramankutty, 2008; Goldewijk, 2001; Ramankutty & Foley, 1999; Sanderson et al., 2002; Vitousek et al., 1997). As a consequence, land use and land cover modification by humans are among the most important agents of environmental

change at local to global scales and have significant implications for ecosystem health, water quality, and sustainable land management (Foley et al., 2005; Lubchenco, 1998). Reliable information regarding the state of global land cover is therefore essential.

Until about fifteen years ago, global land cover datasets were based on pre-existing maps and atlases compiled from ground surveys, national mapping programs, and highly generalized biogeographic maps (Matthews, 1983; Olson, 1982; Wilson & Henderson-Sellers, 1985). In the 1990's, global datasets derived from the AVHRR made it possible to map large scale land cover for the first time based on land surface properties observed from remote sensing (DeFries et al., 1995; Defries & Townshend, 1994; Hansen et al., 2000; Loveland et al., 2000; Stone et al., 1994; Townshend, 1998). As newer moderate resolution remote sensing data sources have emerged (e.g., MODIS, SPOT VEGETATION, MERIS), substantial effort has been focused on developing improved characterizations of global land cover. The current generation of global land cover products include the GLC2000 product produced from SPOT VEGETATION (Bartholome & Belward, 2005), the MODIS Collection 4 Land Cover Product (Friedl et al., 2002),

* Corresponding author.

E-mail address: friedl@bu.edu (M.A. Friedl).

the MODIS Collection 4 Vegetation Continuous Fields product (Hansen et al., 2002), and most recently, the GlobeCover product produced using data from MERIS (Arino et al., 2008).

As each of these datasets have been produced, new approaches have been developed to solve the unique and substantial challenges associated with global land cover mapping. The GLC2000 and GlobCover products were largely developed using unsupervised classification techniques, while the MODIS land cover product uses a supervised approach. The MODIS Vegetation Continuous Fields product also uses a supervised approach, but maps continuous values of vegetation cover at each pixel instead of discrete classes. In each case, unique technical approaches and solutions have been brought to bear on the problem of mapping land cover at very large scales using remote sensing.

In this paper we describe the MODIS Collection 5 Land Cover Type product, which has recently become available to the scientific community. The reprocessing model adopted by the MODIS science team is invaluable because it allows changes to be implemented to algorithms and input data based on experience gained from previous collections. Collection 5 is the latest version of the Land Cover Type product and includes significant changes relative to the Collection 4 product. Here we describe the methods and datasets used to create the Collection 5 product, focusing on changes that have been made to the algorithm and datasets relative to Collection 4.

2. Overview of Collection 5 algorithm refinements

The MODIS land cover product is designed to support scientific investigations that require information related to the current state and seasonal-to-decadal scale dynamics in global land cover properties. The product consists of two suites of science datasets. MODIS *Land Cover Type* (MCD12Q1; Friedl et al., 2002) includes five main layers in which land cover is mapped using different classification systems. Hereafter, we refer to this as the MLCT product. The MODIS *Land Cover Dynamics* product (MCD12Q2; Zhang et al., 2006) includes seven layers, and has been developed to support studies of seasonal phenology and interannual variation in land surface and ecosystem properties. The Collection 5 land cover dynamics product is described elsewhere (Ganguly et al., in press). Here we discuss the land cover type product only.

The MLCT product consists of five different land cover classifications (Table 1) that are produced for each calendar year. These layers include the 17-class International Geosphere–Biosphere Programme classification (IGBP; Loveland & Belward, 1997); the 14-class University of Maryland classification (UMD; Hansen et al., 2000); a 10-class system used by the MODIS LAI/FPAR algorithm (Lotsch et al., 2003; Myneni et al., 2002); an 8-Biome classification proposed by Running et al. (1995); and a 12-Class plant functional type classification

Table 1
Classifications included in the MOD12Q1 product.

	IGBP	UMD	LAI/FPAR	BGC	PFT
Forests	Evergreen needleleaf forest (1)	Evergreen needleleaf forest	Evergreen needleleaf forests	Evergreen needleleaf vegetation	Evergreen needleleaf tree
	Deciduous needleleaf forest (2)	Deciduous needleleaf forest	Deciduous needleleaf forests	Deciduous needleleaf vegetation	Needleleaf deciduous tree
	Evergreen broadleaf forest (3)	Evergreen broadleaf forest	Evergreen broadleaf forests	Evergreen broadleaf vegetation	Broadleaf evergreen tree
	Deciduous broadleaf forest (4)	Deciduous broadleaf forest	Deciduous broadleaf forests	Deciduous broadleaf vegetation	Broadleaf deciduous tree
	Mixed forests (5)	Mixed forest			
Woodlands	Woody savannas (8)	Woody savannas	Savanna		
	Savannas (9)	Savannas			
Grasses/cereals	Grasslands (10)	Grassland	Grasses/cereal crops	Annual grass vegetation	Grass
Shrublands	Closed shrublands (6)	Closed shrublands	Shrublands		Shrub
	Open shrublands (7)	Open shrublands	Shrublands		Shrub
Croplands and mosaics	Croplands (12)	Croplands	Broadleaf crops	Annual broadleaf vegetation	Cereal crop
	Cropland/natural vegetation mosaics (14)				Broadleaf crop
Seasonally or permanently inundated	Permanent wetlands (11)				
Unvegetated	Urban and built-up land (13)	Urban and built-up land	Urban	Urban	Urban
	Barren or sparsely vegetated (16)	Barren or sparsely vegetated	Unvegetated	Unvegetated	Barren or sparsely vegetated
	Water (17)	Water	Water	Water	Water
	Permanent snow and ice (15)				Snow and ice

Shaded boxes indicate no corresponding class relative to IGBP; numbers in parentheses indicate IGBP class numbers used in this paper. Classification acronyms are defined in Section 2.

described by Bonan et al. (2002a). In addition to these classification layers, the MLCT product provides the most likely IGBP class and a continuous measure of “classification confidence” at each pixel (McIver & Friedl, 2001). A lower spatial resolution climate-modeling grid (MCD12C1) is produced at 0.05° spatial resolution for users who do not require the spatial detail afforded by the 500-m land cover product. The MCD12C1 product provides the dominant land cover type as well as the sub-grid frequency distribution of land cover classes within each 0.05° cell. For practical reasons the discussion here focuses on the IGBP layer.

The MODIS land cover type product is produced using an ensemble supervised classification algorithm. The base algorithm is a decision tree (C4.5; Quinlan, 1993), and ensemble classifications are estimated using boosting (Freund & Schapire, 1996; Quinlan, 1996; Schapire et al., 1998). The use of boosting is central because it allows the algorithm to derive estimates of class-conditional probabilities for each class at each pixel (Friedman et al., 2000; McIver & Friedl, 2001). As in Collection 4, we use an ensemble of ten boosted decision trees to generate the product. Results from the ensemble decision trees are post-processed to correct classification results for biases inherent to the decision tree algorithm caused by specific properties of the training sample, and to exploit extant information related to the geographic distribution of global land cover (Section 4). The basic processing steps of this algorithm are presented schematically in Fig. 1, and details are provided elsewhere (Friedl et al., 1999, 2000a, 2002; McIver & Friedl, 2001, 2002). Here we describe modifications that are unique to Collection 5. Specifically, we focus on: (1) revisions to the MODIS land cover training site database, (2) changes to input features, (3) refinement of ancillary data layers that are merged with ensemble decision tree results to produce the final product, and (4) methods we have developed to tune, refine, and stabilize classification results across different years.

3. Data

3.1. Training data

High quality training data are essential to the MLCT algorithm, and each reprocessing of the MODIS land cover product has provided an invaluable opportunity to revise and augment the MODIS land cover training site database. Because of the importance of this database and because source data can become out-of-date, maintenance of the site database is an important and ongoing process. The ability to update and revise this database afforded by periodic reprocessing has been highly beneficial and has resulted in a mature database of land cover training sites.

Training data for the Collection 5 product includes 1860 sites distributed across the Earth's land areas (Fig. 2, Table 2). To ensure that the database captures a wide range of geographic and ecological variability, the database is periodically intersected with a map of Olsen's ecoregions, which allows under- or unsampled regions to be identified (Friedl et al., 2002). Each site consists of a polygon, delineated on Landsat or higher resolution imagery via manual interpretation, where the land cover is uniform and representative of one IGBP class. The size of sites range from 1 500-m MODIS pixel (~0.2 km²) to 376 pixels (~80 km²), but the distribution is highly skewed towards smaller sites: the median size is 16 pixels and 1741 sites cover fewer than 50 MODIS pixels.

For Collection 5, each site was reviewed using Landsat or higher spatial resolution data. As part of this process, we systematically updated our database of image data and replaced older Landsat images acquired in the 1990's with Landsat7 or orthorectified Geocover 2000 imagery. In recent years, the availability of imagery via GoogleEarth© has been extremely valuable as a supplemental data source. In addition to removing sites with low quality labels, correcting labeling errors, and improving the ecological representation of the site database, substantial effort was devoted to adding

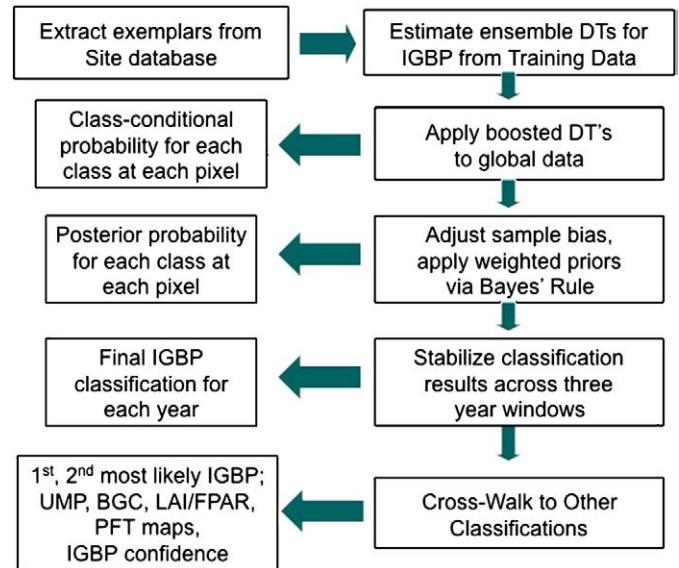


Fig. 1. Flow chart for MOD12Q1 production.

sites in regions where the database had poor representation, and to reducing the size of larger sites. This latter activity was particularly important because spatial correlation within sites leads to significant redundancy in the training data.

Table 2 summarizes the distribution of sites and training data by continent and IGBP land cover type in Collection 5, and Fig. 3 shows the frequency distribution of training sites and pixels in Collection 5 along with differences between each in Collections 4 and 5. Because of its geographic extent and variability, agriculture (class 12) is by far the most heavily sampled class. Wetlands, which are highly diverse at global scales, are also heavily sampled. Deciduous needleleaf forest (class 3), on the other hand, is somewhat under-sampled because it is difficult to identify representative sites for this class at the scale required for site delineation. The most obvious differences between Collection 4 and Collection 5 are decreases in the number of sites or pixels for the evergreen forest, wetlands, agriculture, and agricultural mosaic classes (classes 1, 2, 11, 12, and 14, respectively), and modest increases in the number of sites or pixels for deciduous broadleaf forests and closed shrublands (classes 4, 6). A substantial number of savanna and woody savanna (classes 8, 9) training sites were added, but the total number of training pixels in these two classes decreased. Note that while the number of MODIS training pixels did not substantially change relative to Collection 4, the total area represented by these pixels has decreased four-fold because of the increased spatial resolution used in Collection 5. This change also reflects a focus on using smaller, high quality sites in Collection 5 relative to previous collections.

3.2. Input data and features

Input features used in the MLCT algorithm include spectral and temporal information from MODIS bands 1–7, supplemented by the enhanced vegetation index (EVI; Huete et al., 2002). We also include Collection 5 MODIS Land Surface Temperature (LST; Wan et al., 2002), which was not used in previous Collections. For bands 1–7 and to compute the EVI, we use nadir BRDF-adjusted reflectance (NBAR) data provided by the MODIS BRDF/albedo product (Schaaf et al., 2002). This product provides surface reflectance measurements that are normalized to a consistent nadir view geometry based on BRDF-models of surface anisotropy, thereby minimizing the effect of variable view geometry in surface reflectance data.

Collection 5 NBAR data are produced on a rolling 8-day interval based on 16 days of MODIS surface reflectance data at a spatial

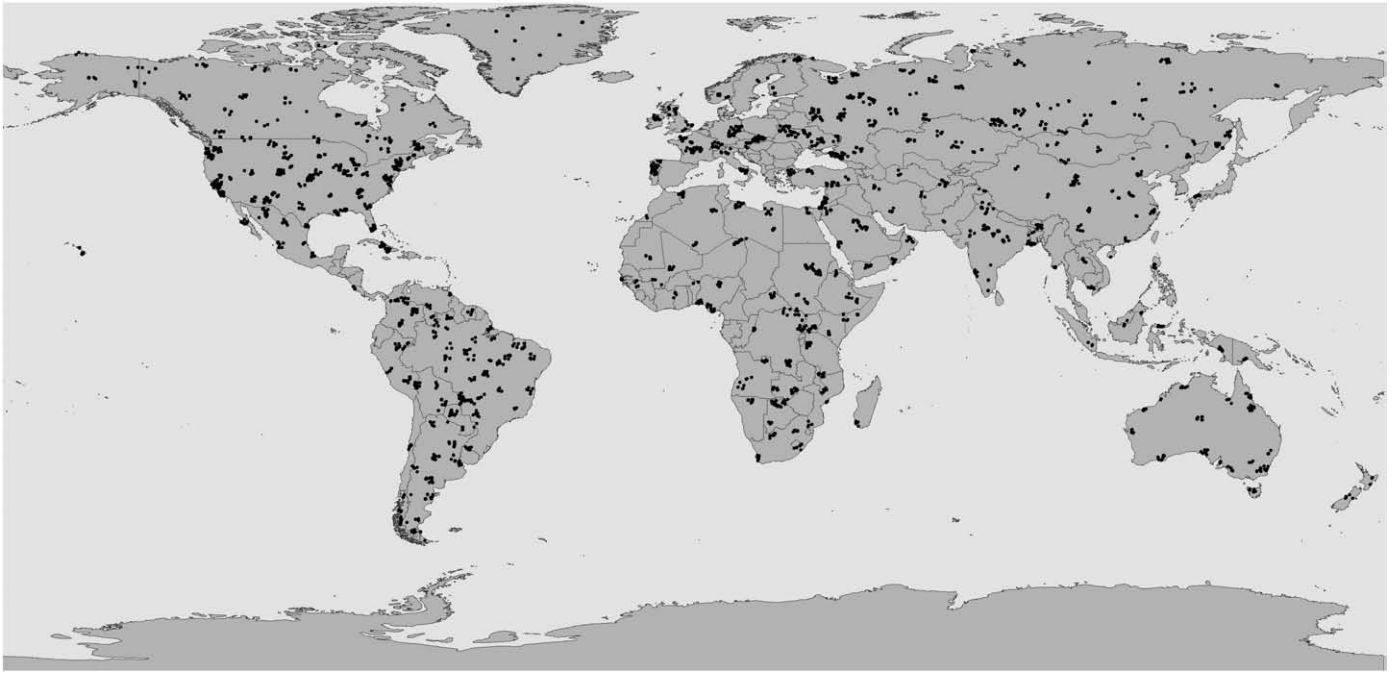


Fig. 2. Map of training sites (identified by dots) used to create the MODIS land cover type product.

resolution of 500-m. The Collection 4 product was produced at 1-km spatial resolution at 16-day intervals (note that the spatial resolution of the MODIS sinusoidal grid is actually 463.313-m and 926.625-m; 1000-m and 500-m are used by convention). This change has two positive implications for the MLCT product. First, the availability of 500-m NBAR data provided the basis for increasing the spatial resolution of the MLCT product to 500-m in Collection 5. Second, because the MLCT algorithm aggregates 8-day values to 32-day averages (to reduce data volumes and using a quality assurance-weighted averaging procedure), fewer missing values caused by clouds and other sources are present in the input features relative to Collection 4.

As in Collection 4, the Collection 5 MLCT product is generated on a calendar year basis. For each calendar year, algorithm inputs include twelve sets of 32-day average NBAR, LST and EVI data. In addition, annual metrics (minimum, maximum and mean values) for the EVI,

LST and NBAR bands are also included as inputs, providing a total of 135 features. In this way, the algorithm is able to exploit information related to the phenology and temporal variability characteristic of land cover types that complements the spectral information provided by MODIS (Friedl et al., 1999; Lloyd, 1990; Loveland et al., 1995; Townshend et al., 1987).

Depending on the location and time of year, MODIS input data can include substantial levels of missing data arising from clouds (especially in the tropics) and low illumination and polar night in the northern high latitudes. C4.5 provides robust algorithms for coping with missing features (Quinlan, 1993). However, if a substantial proportion of the input features are missing, the reliability of classification results degrades. To be conservative, if the number of missing features at a pixel exceeds 84 features, the pixel is not classified. In this situation, the pixel is filled using the most recent Collection 5 label. If

Table 2
Frequency distribution for training site pixels by IGBP class and continent.

IGBP class	NA		SA		AF		EA		AU		Total		Mapped area (km ²)
	Site	Pixel	Site	Pixel	Site	Pixel	Site	Pixel	Site	Pixel	Site	Pixel	
1	43	529	0	0	0	0	69	1036	2	23	114	1588	4,136,838
2	2	41	80	1990	14	236	11	143	24	403	131	2813	13,469,653
3	0	0	0	0	0	0	25	755	0	0	25	755	2,718,250
4	23	222	18	728	12	355	60	1033	0	0	113	2338	2,042,995
5	28	291	3	61	0	0	72	597	0	0	103	949	6,074,390
6	18	286	21	375	15	337	20	261	13	458	87	1717	2,506,826
7	34	530	18	526	27	624	31	1036	5	82	115	2798	20,181,252
8	33	339	13	396	41	791	40	497	11	181	138	2204	13,589,431
9	17	170	27	282	34	700	22	277	10	197	110	1626	8,612,734
10	35	756	26	509	10	274	39	870	8	227	118	2636	15,158,441
11	42	1285	29	562	45	830	40	670	13	274	169	3621	1,651,294
12	76	2094	40	812	29	579	161	4670	21	317	327	8472	12,041,134
13	-	-	-	-	-	-	-	-	-	-	-	-	656,263
14	25	86	27	70	30	295	53	289	4	8	139	748	8,622,136
15	6	205	11	516	0	0	7	154	2	17	26	892	15,725,424
16	4	103	2	35	48	3880	30	866	6	129	90	5013	18,300,756
17	12	392	7	198	8	276	24	982	4	113	55	1961	2,063,628
Total	398	7329	322	7060	313	9177	704	14136	23	2429	1860	40131	147,551,454

NA = North America; SA = South America, AF = Africa; EA = Asia; AU = Australia and Pacific islands. Note that class 13 (urban) was mapped separately and training data is therefore not included in the table.

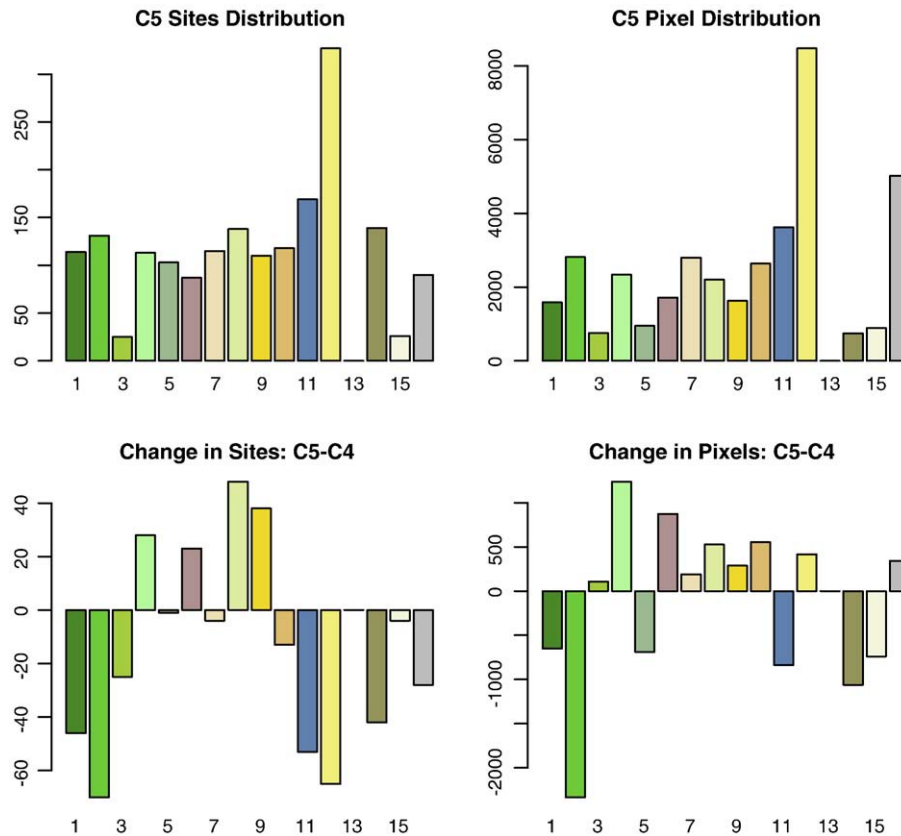


Fig. 3. Barplots showing the frequency distribution of training sites and pixels in Collection 5 along with differences in the number of training sites and pixels between Collections 4 and 5. Numbers on horizontal axis refer to IGBP classes, provided in Table 1.

the problem is persistent (i.e., no Collection 5 values), the pixel is filled using Collection 4 data. This situation is rare, and affects only a very small number of pixels.

4. Post-processing of ensemble decision tree results

4.1. Overview of issues

As we alluded to above, post-processing refinements are applied to the ensemble decision tree output to create the final Collection 5 land cover product (Fig. 1). The specific adjustments we have developed address limitations imposed by: (1) the spectral–temporal information content of MODIS data, and (2) biases that are inherent to tree-based classification models. Fig. 4 presents an example from the south central United States centered on the lower Mississippi valley (MODIS tile h10v05, roughly 1100 km × 1100 km) that provides a step-by-step illustration of the effects and importance of the refinements we describe below.

The limitations identified above arise from two fundamental assumptions of the ensemble decision tree classification algorithm: (1) that the distribution of the training data is representative of the population, and (2) that the features are able to distinguish the classes in the training data. In the present case, neither of these assumptions is strictly valid. The training site database we have compiled is designed to capture geographic and ecological variability, but it is unrealistic to claim or assume that it captures the complete range of variability in global land cover. Similarly, the class frequency distribution of the training data does not reflect the global distribution of land cover classes (Fig. 5). Further, even if both these conditions were met, the land cover class definitions used in the MLCT product were developed in support of science community needs, but not on a thorough understanding of what classes MODIS can consistently identify with high accuracy. As a consequence, the spectral–temporal separability of many classes is

ambiguous (e.g., savanna versus woody savannas versus grasslands), a problem that is compounded by the inclusion of mixture classes (e.g., agricultural mosaic, mixed forests).

For pixels where the training set does not include a good exemplar site or where the spectral–temporal information is equivocal, supervised algorithms over- (under) predict more (less) frequent classes in the training data, leading to bias and errors in classification results (McIver & Friedl, 2002). Below we describe two algorithm refinements we have implemented to address these limitations. In both cases we adjust the class-conditional probabilities produced from the ensemble decision trees using Bayes' rule in association with parameterized prior probabilities.

4.2. Sample bias correction

The first issue, classification bias imposed by the training data, we address via a sample bias correction. This issue is illustrated in Fig. 5, which shows the global frequency distribution for both the training sites and pixels, along with land areas in each class based on our final classification. Clearly, there are substantial differences, the most obvious being agriculture and wetlands. Fig. 4 illustrates how this bias propagates into classification results. Specifically, because they are over-sampled in the site database relative to other classes, agriculture (shown in yellow) and wetlands (shown in dark blue) are over-represented in the ensemble decision tree results. In simple terms, because the classifier is optimized to maximize classification accuracy based on the training data, classification results are biased to over-predict the most common classes in the training data. Conversely, rarer classes in the training data tend to be penalized.

To correct this problem, the class-conditional probabilities estimated by the boosted decision trees are adjusted using Bayes' Rule based on prior probabilities prescribed to be inversely proportional to the number

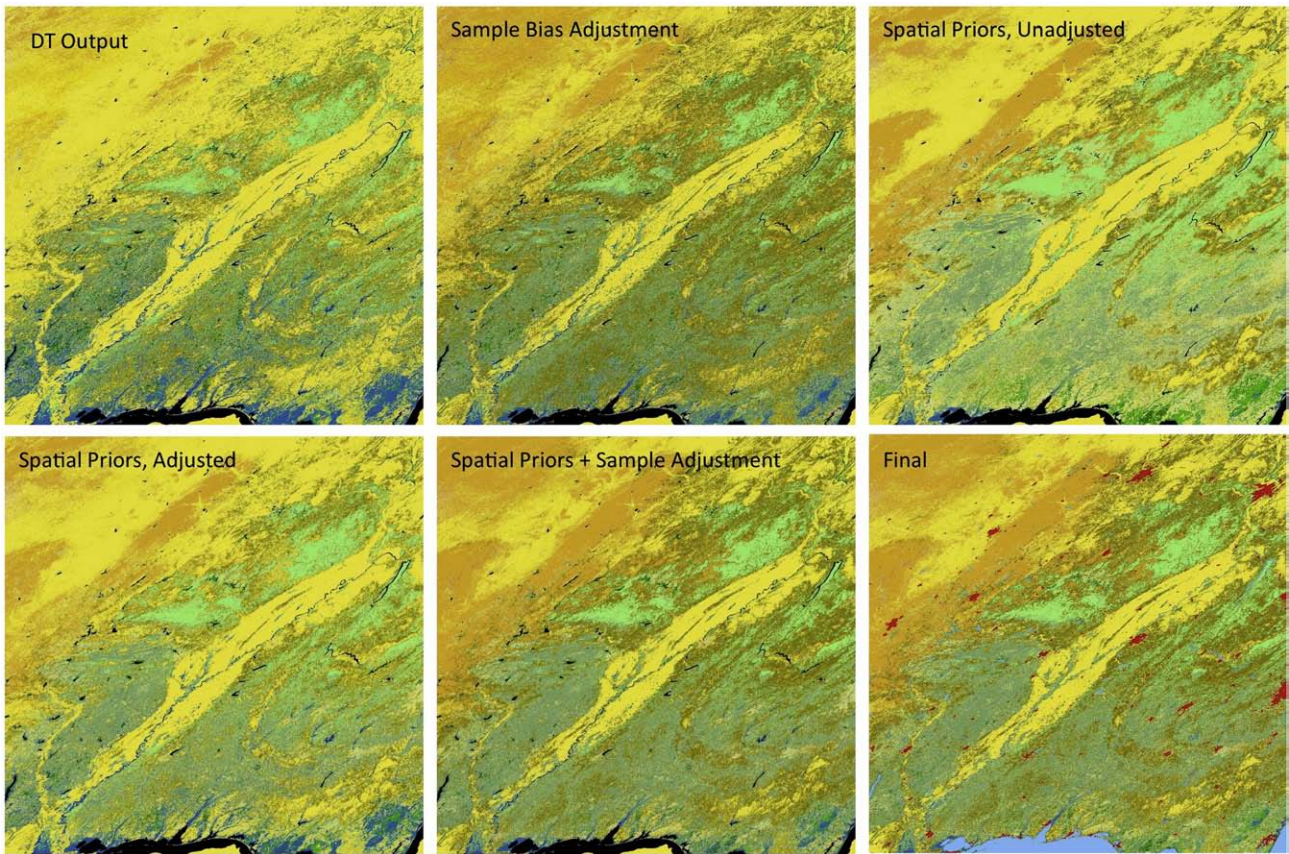


Fig. 4. Image panel for MODIS tile V05H10 (south central United States) showing classification results at each stage of processing. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

of training samples in each class. This has the effect of reducing the posterior probabilities (estimated via Bayes' Rule) for classes that are over-sampled in the training data, and vice versa. Fig. 6 demonstrates how the frequency of over-sampled classes is reduced in the predictions by applying this correction, and vice versa for under-sampled classes. The

Training Data Distribution, Sample Correction, Adjusted Proportions

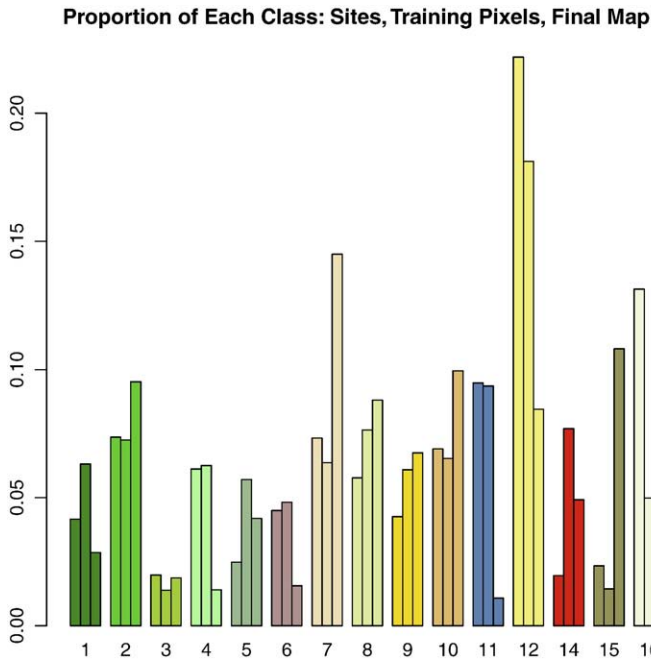


Fig. 5. Barplots showing for each IGBP class the proportion of: training sites (left bar); training pixels (middle bar), and final mapped classes (right bar).




Fig. 6. Barplots showing for each IGBP class: the proportion of training pixels in each class (left bar), the prior probabilities used to implement the sample bias correction (middle bar), and the resulting effective overall likelihood for each class (right bar). The net effect is to reduce the overall likelihood of more heavily sampled classes (and vice versa), thereby reducing the bias imposed by the training sample.




Fig. 7. Parameterization used to prescribe the prior probabilities for classes 12 and 14 (agriculture and agricultural mosaic) based on cropping intensity from Ramankutty et al. (2008).

top middle panel of Fig. 4 illustrates how this adjustment is manifested in the Collection 5 classification results for MODIS tile number h10v05: the reduction in over-sampled classes (agriculture; wetlands) relative to the ensemble decision tree results is clearly evident. In this context, it is important to note that the magnitude of the bias introduced by the sample distribution depends on the degree to which each class is separable in the feature space, and classes that are highly separable are relatively unaffected by the training sample distribution.

4.3. Spatially explicit prior probabilities

The second issue, inadequate class separability compounded by mixture classes, we address by adjusting ensemble decision tree results

based on spatially explicit prior probabilities, which are parameterized using extant information derived from two main sources. The first source of information is the MODIS Collection 4 MLCT product. Using Collection 4 maps, probabilities are estimated at each pixel using a moving window algorithm that computes the regional proportion (i.e., the likelihood) for each class. In Collection 4, we used a 201 by 201 window ($\sim 35,000 \text{ km}^2$) derived from the IGBP DISCover dataset (Loveland et al., 2000) to do this. In Collection 5, we use a 151 by 151 window ($\sim 20,000 \text{ km}^2$). This step assumes that the Collection 4 data capture the regional variability in land cover, and that the regional frequency distribution therefore provides a good basis for parameterizing the prior probability for each IGBP class at each pixel.

As part of this process, the prior probabilities for agriculture (class 12) and agricultural mosaic (class 14) derived from the moving window are replaced with probabilities parameterized using the dataset produced by (Ramankutty et al., 2008), which furnishes estimates of global cropping intensity at 0.05° spatial resolution (roughly 30 km^2 at the equator) for year 2000. Because this dataset uses remote sensing data sources merged with local census data, it provides an excellent basis for parameterizing the prior probabilities for these two classes at much higher spatial resolution than the moving window procedure described above affords. To parameterize the prior probabilities for classes 12 and 14, we use a Gaussian function of cropping intensity centered at 50% to prescribe the local likelihood for class 14 (agricultural mosaic), and a sigmoidal function of cropping intensity for class 12 (agriculture), where the prior probabilities for each class intersect at 60% (Fig. 7). In this way, the probabilities are consistent with the definitions for each class. Also, the maximum prescribed prior probability associated with classes 12 and 14 does not exceed 0.5, thereby reducing the likelihood that the classifier results simply replicate the results of (Ramankutty et al., 2008). After replacing the values for classes 12 and 14 in this fashion, the vector of prior probabilities at each pixel is normalized to sum to 1.0.

Fig. 8 shows the global distribution of the resulting prior probabilities for the agriculture and agricultural mosaic classes, and the net result of applying the merged spatial prior probabilities at a regional scale is shown in the top-right panel of Fig. 4. Like the sample bias adjustment, the main effect is to reduce over-prediction of over-sampled classes in

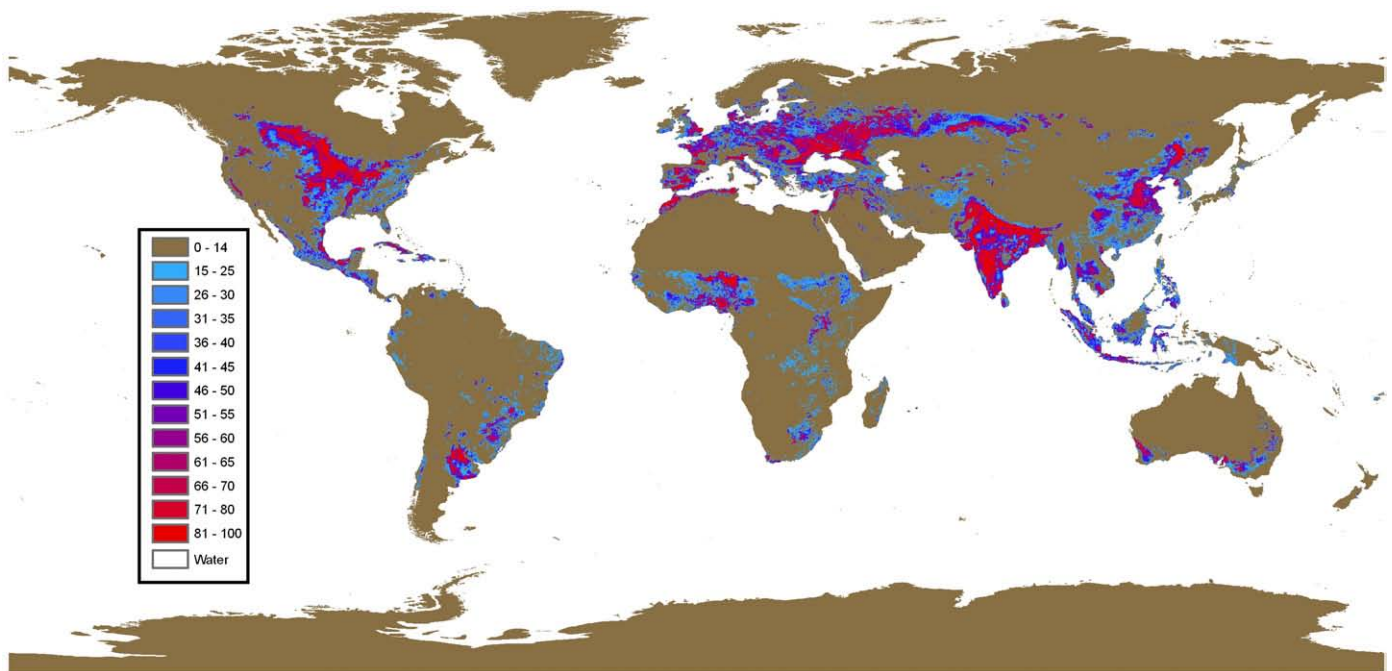


Fig. 8. Map of prior probabilities ($\times 100$) for agriculture (class 12) and agricultural mosaic (class 14) derived from Ramankutty et al. (2008) using the parameterization shown in Fig. 7. Red and pink areas indicate regions with high likelihood of high intensity agriculture. Blue and purple areas indicate areas dominated by less intensive agricultural mosaics. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Effect of Confidence Parameter on Priors




Fig. 9. Schematic showing how the value of the spatial prior probability confidence parameter influences the magnitude of the prior probabilities used to create the final map. The leftmost bar is the original prior probability. The subsequent bars, from left to right, show the adjusted priors for $c = 0.9, 0.5,$ and $0.1,$ respectively. The net effect, as c varies from 1 to 0, is to progressively adjust the priors towards a uniform distribution (i.e., $c = 0$ is equivalent to equal priors).

the training data. Closer inspection also reveals increased grasslands in the northwest quadrant and a smaller proportion of agricultural mosaic throughout the tile. As for the sample bias correction, highly separable classes are relatively unaffected by this correction.

4.4. Tuning classification results

Ideally, we wish to minimize the influence of the spatial prior probabilities and to maximize information from MODIS Collection 5 data. To do this, a tuning parameter is used that controls how heavily the prior probabilities at each pixel are weighted. This “confidence parameter” (c), which ranges from 0–1, is used to apply a linear transformation that weights the spatial priors more or less heavily, depending on the value of c . Specifically, the prior probabilities at each pixel are adjusted using the following expression:

$$P(i) = P(i) + (1 - P(i)) \times (1 - c), \tag{1}$$

where $P(i)$ is the prior probability for class i at any given pixel. After the adjustment has been applied, the vector of $P(i)$'s at each pixel are normalized to sum to 1.

The effect of using different values for the confidence parameter c is illustrated in Fig. 9 for a hypothetical pixel in the northeastern United States. The first bar for each class shows the unadjusted prior probability (i.e., the most likely classes are 1, 4, 5, 12, and 14), and the subsequent bars show the adjusted probabilities for c values of 0.9, 0.5, and 0.1, moving from left to right. The net effect, as c varies from 1 to 0, is to linearly adjust the prior probabilities for each class from their original distribution to a uniform distribution (i.e., equal priors for all classes). For Collection 5, we use a value of 0.25 for c , which is conservative and was determined by extensive trial and error using selected tiles spanning a range of continents and land cover types.

The upper right and bottom left panels of Fig. 4 show the result of fusing the ensemble decision tree and spatial prior probabilities before and after applying the confidence parameter (i.e., $c = 1, 0.25$ respectively). As we indicated above, it is desirable to reduce the weight of the spatial priors as much as possible; i.e., we wish to maximize the information from Collection 5 and minimize the algorithm's dependence on information from Collection 4 and Ramankutty et al. (2008).

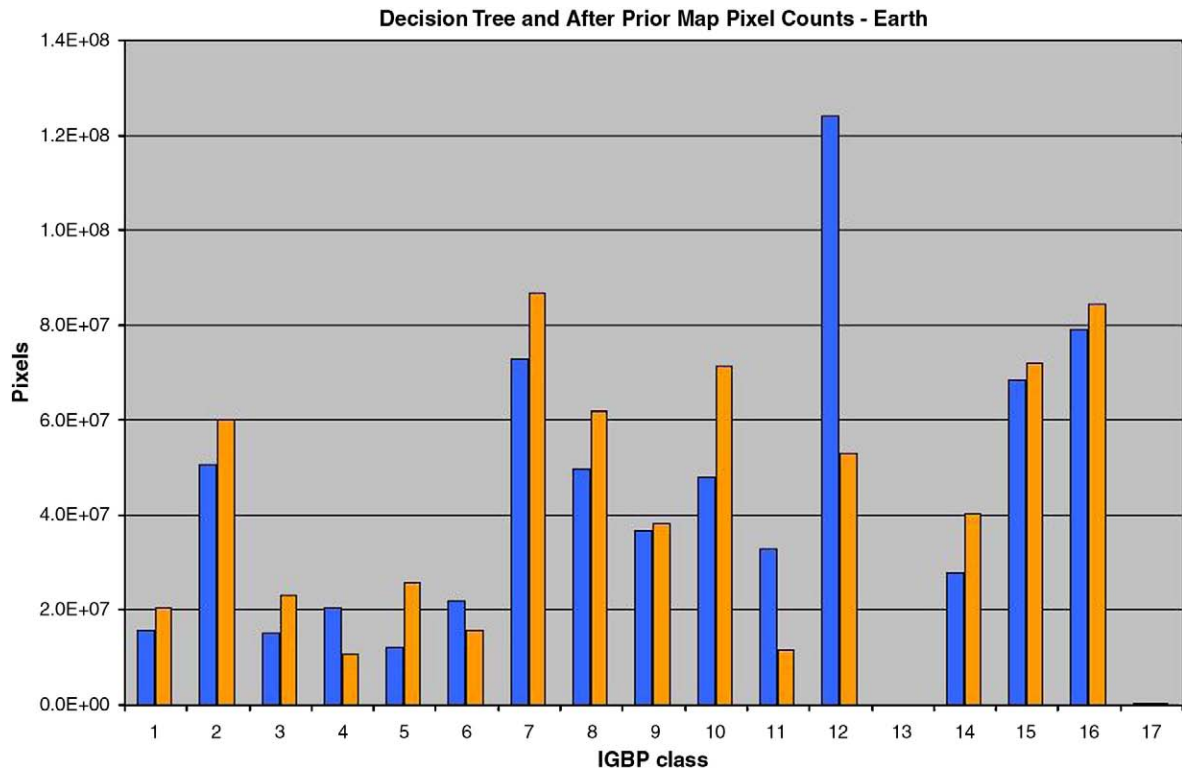


Fig. 10. Differences in the number of pixels mapped in each class at global scale by the ensemble decision trees (blue) and after applying the post-processing adjustments (orange) described in Section 4. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Figs. 4 and 9 reveal that by using a relatively low value for c , the net effect is quite modest. The final step is to merge the results from the sample bias correction and the spatial prior adjustment, which is shown in the bottom middle panel of Fig. 4.

Fig. 10 shows the global class frequency distribution before and after applying the sample bias and spatially explicit prior probability adjustments. Overall, the adjustments change 6.5% of pixels. However, these adjustments are not distributed uniformly. Agriculture is the most common class mapped by the ensemble decision trees, but is reduced by over 50% once the adjustments have been applied. The mapped areas for deciduous broadleaf forests, closed shrublands, and wetlands are also reduced as a result of the sample bias and spatial prior adjustments. These changes are balanced by increases in the area mapped for all other classes, with the largest increases in grasslands, savannas, open shrublands, mixed forests, and agricultural mosaic.

5. Special cases: urban land use, wetlands, and deciduous needleleaf forests

In addition to the post-processing described above, experience from previous MODIS collections has demonstrated that several classes are particularly problematic and difficult to map. In particular, wetlands and deciduous needleleaf forests tend to be over-represented, even after the adjustments described above are applied. To correct this, we have identified thresholds (lower boundaries) for posterior probabilities that are required for a pixel to be labeled as wetland or deciduous needleleaf forest. If this threshold is not met, the label is replaced with the next most likely class.

For deciduous needleleaf forests we applied a threshold of $P > 0.7$, which was determined based on extensive trial and error. Applying this threshold substantially reduced errors of commission associated with this class. Wetlands, on the other hand, presented substantial errors of omission and commission. To reduce errors of commission, we required a minimum posterior probability threshold of $P > 0.75$. To reduce errors of omission, the algorithm examines the decision tree results (i.e., prior to applying the adjustments described in Sections 4.2 and 4.3) and retains pixels with very high class-conditional probabilities for wetlands ($P > 0.9$). This is required because wetlands are relatively rare and are frequently small in extent. As a consequence, the wetlands class tends to have low prior probability outside of large wetlands complexes. Further, Fig. 3 shows that wetlands (class 11) are relatively heavily sampled in the site database. The net effect of both the sample bias and spatial prior adjustments is therefore to impose low prior probabilities for this class, leading to errors of omission, particularly in regions where wetlands are not spatially extensive. Using the strategy described above, we eliminate




Fig. 12. Proportion of pixels changing from year-to-year from 2001–2005, before and after applying stabilization.

many spurious wetland pixels arising from oversampling in the training set, but retain high confidence wetland pixels predicted by the classifier prior to post-processing adjustments. In cases where the criteria described above are not met, the maximum likelihood class is replaced with the second most likely class. Extensive inspection of the resulting maps in association with high-resolution imagery indicates that these strategies provide the best qualitative compromise between errors of omission and commission associated with the wetlands and deciduous needleleaf classes.

We have implemented a similar set of adjustments to improve representation of water in inland areas, particularly along coastlines. These adjustments affect a small proportion of pixels. Specific details are beyond the scope of this paper, but visual inspection of results clearly reveals the benefit from these adjustments. Future versions of the product will use a newly created land-water mask that should resolve much of this problem.

Urban land areas present a particularly difficult case. Urban cores tend to be sparsely vegetated and are often difficult to distinguish from barren and sparsely vegetated land areas. Conversely, suburban land areas are easily confused with natural vegetation classes. Further, the density and form of urban areas at global scales vary widely with climate and socio-economic factors. In Collection 4, we addressed these issues by mapping urban land areas as a separate class using prior probabilities based on a combination of gridded population data and the DMSP nighttime lights dataset (Schneider et al., 2003). In Collection 5, we use a different approach. Specifically, global urban




Fig. 11. Urban land cover (in yellow) overlaid on a Landsat scene for Guangzhou, China: (a) Landsat-based classification, (b) MODIS Collection 5, (c) MODIS Collection 4. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 3
User and producer accuracies, standard errors, and 95% confidence intervals for Collection 5 IGBP classes based on cross-validation.

IGBP land cover class	Producer's accuracy (%)				User's accuracy (%)			
	PA	Std. err.	CI–	CI+	UA	Std. err.	CI–	CI+
1.	89.8	2.3	85.3	94.4	78.0	5.3	67.5	88.6
2.	92.6	2.4	88.0	97.2	83.1	3.2	76.8	89.5
3.	67.3	10.9	45.8	88.7	90.4	4.6	81.4	99.4
4.	68.9	6.2	56.7	81.0	75.9	5.3	65.6	86.3
5.	76.2	5.7	65.1	87.3	53.1	6.1	41.1	65.1
6.	63.4	5.9	51.9	74.9	47.0	5.5	36.1	57.8
7.	48.3	6.2	36.1	60.5	74.1	5.2	63.8	84.4
8.	45.2	4.1	37.2	53.3	34.3	4.5	25.4	43.2
9.	22.6	4.4	13.9	31.3	39.0	6.0	27.2	50.8
10.	73.6	4.1	65.7	81.6	55.9	4.2	47.6	64.2
11.	70.6	4.2	62.4	78.7	96.4	1.8	92.7	99.9
12.	83.3	2.0	79.4	87.1	92.8	1.5	89.8	95.8
14.	60.5	5.7	49.3	71.7	27.5	3.6	20.5	34.6
15.	75.6	10.9	54.4	96.9	96.8	2.3	92.2	100.0
16.	95.8	1.4	93.1	98.4	92.7	2.1	88.5	96.8
17.	96.6	1.9	92.8	100.0	99.3	0.4	98.6	100.0

land areas were mapped using an ecoregion-based stratification with eighteen strata, where training data and supervised classifications were developed and tuned to each stratum. Because of the fragmented form of many urban areas, the higher spatial resolution used in Collection 5 provides a significantly improved representation of urban land use. This is visually evident from inspection of the map product (Fig. 11), and is confirmed by a validation based on Landsat-derived maps of urban land use for 135 cities. Full details are provided in (Schneider et al., in press).

6. Stabilization of results across years and cross-walking classification schemes

The final component of the MLCT algorithm includes two elements: reducing year-to-year variability in classification results and creating the additional layers to the IGBP classification. Reducing the amount of interannual change in classifications is a particularly difficult challenge because classification results in heterogeneous areas and ecotones are unstable and tend to toggle year-to-year between similar classes. As we described above, this problem arises because many landscapes include mixtures of classes at 500-m spatial resolution and because the spectral–temporal signature of some land cover classes is not easily separable in MODIS data. Further, year-to-year variability in phenology and disturbances such as fire, drought, and

insect infestations present a highly variable system that is difficult to consistently characterize globally at annual time scales. As a consequence, classification results at global scales can vary substantially from one year to the next (Fig. 12). While most of these changes involve classes that are ecologically proximate and arise from poor spectral–temporal separability in MODIS data (e.g., mixed forest and deciduous broadleaf forest; grassland and open shrublands), it is desirable to reduce the amount of spurious year-to-year change in the maps.

To address this, the MLCT algorithm imposes constraints on year-to-year variation in classification results at each pixel. To do this, we use the posterior probability associated with the primary label in each year. If the classifier predicts a different class from the preceding year, the class-label is changed only if the posterior probability associated with the new label is higher than the probability associated with the previous label. To avoid propagating incorrect or out-of-date labels in areas of change, we apply this procedure using three-year windows. In this way, we perpetuate high quality labels, account for the possibility of land cover change, and reduce the amount of interannual variation in labels to about 10% (Fig. 12). Note, however, that this level of change is still well above the amount of actual global land cover change. Thus, land cover change should not be inferred by differencing the MLCT product across years.

The final step in creating the MLCT product is generating the UMD, LAI/FPAR, 8-Biome, and Plant Functional Type layers. This is accomplished by cross-walking the IGBP layer to each classification scheme using the associated posterior probabilities in combination with global classifications for leaf type (broadleaf, needleleaf), phenology (deciduous, evergreen) and crop type (cereal, broadleaf). The logic behind this process provides internal consistency among classes across the different classification systems.

7. Cross-validation analysis of accuracy

To provide a quantitative assessment of map accuracy, we performed a 10-fold cross-validation analysis using the training site database. This method is widely used to assess classification accuracy when independent validation data are not available. To do this, the training site database was stratified into 10 unique subsets, each with 186 randomly selected sites. To avoid spatial correlation in training and test data, sites (not individual pixels) were used as the sampling unit (Friedl et al., 2000b). Using this approach, ten classifications were performed using the MLCT algorithm, each based on a unique combination of 9 subsets to train the data, and using the remaining subset as a “test” set. In this way, every pixel in the database was classified based on an independent

Table 4
Confusion matrix for Collection 5 IGBP classes based on cross-validation.

Training site label	Classification output label																
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1	1426	13	40	16	74	1	3	225	0	0	31	0	0	0	0	1	0
2	19	2550	0	90	52	18	0	46	0	0	276	6	0	11	0	0	0
3	2	0	508	2	19	0	3	10	10	0	7	0	0	0	0	0	0
4	0	23	11	1611	64	24	0	217	30	3	100	14	0	23	0	0	0
5	103	24	143	250	723	3	0	21	5	2	68	5	0	8	2	0	3
6	0	21	11	0	1	1086	519	149	98	225	70	116	0	5	8	7	0
7	0	0	1	0	0	179	1351	9	23	116	1	32	0	5	0	104	0
8	35	44	22	291	14	131	64	997	691	77	202	246	0	82	0	16	0
9	0	0	11	5	0	53	27	276	367	95	31	54	0	20	0	0	0
10	0	3	0	0	0	140	430	66	213	1938	40	414	0	26	111	77	0
11	2	48	7	0	2	0	0	2	4	5	2406	19	0	0	0	0	0
12	0	0	0	0	0	9	64	52	25	118	60	6963	0	84	66	2	61
13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14	1	28	1	73	0	47	22	134	130	22	105	498	0	402	0	0	0
15	0	0	0	0	0	0	1	0	0	13	0	0	0	0	580	4	0
16	0	0	0	0	0	20	314	0	27	14	0	4	0	0	0	4802	1
17	0	0	0	0	0	0	0	0	0	0	12	0	0	0	0	0	1831

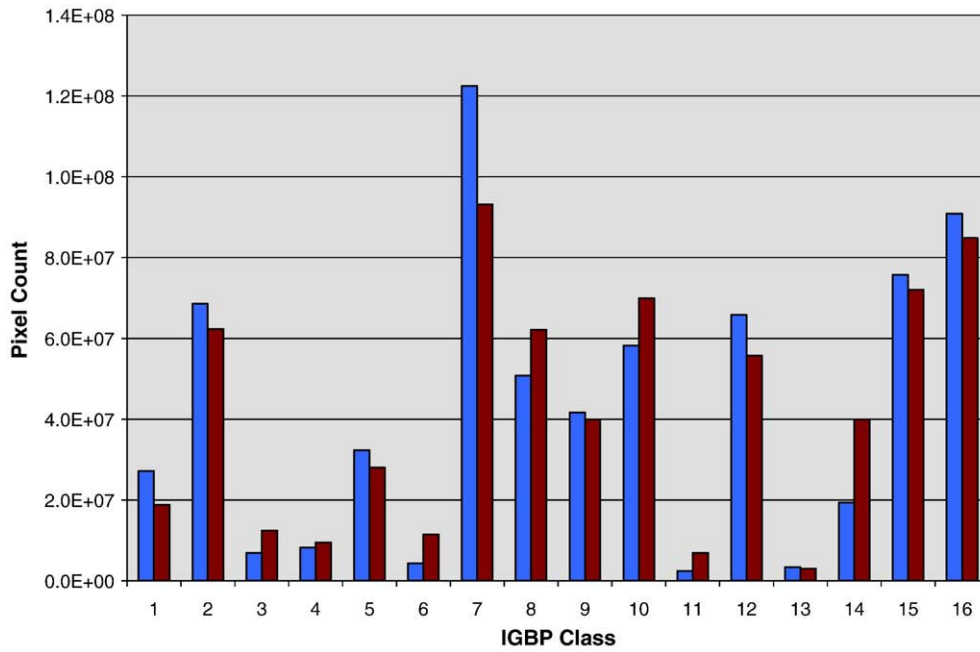


Fig. 13. Barplots showing the number of pixels in each class in Collection 4 (blue) and Collection 5 (red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

training set. In this context, it is important to note that because the cross-validation runs are based on 90-percent random samples of the training data, they are likely to have modestly lower predictive accuracy than the classifications used in operational production of the product. Thus, the classification accuracies and standard errors reported here should be conservative.

Because the analysis used sites as the basic sampling unit, estimates for standard errors reported below are based on cluster sampling (Cochran, 1977; Stehman, 1997). Confusion matrices and other quantities related to classification accuracy are well-described in the liter-

ature and will not be discussed here, except to indicate that we followed well-established community protocols used for this type of analysis (Foody, 2002; Strahler et al., 2006). Space does not allow a detailed presentation and we are preparing a separate paper that presents a more complete characterization of results from this analysis. Here we present an overview of results from 2005, which are representative of results from the multi-year record.

The overall accuracy across all classes in the 2005 map is 74.8%. The error variance on this estimate is 1.3%, yielding a 95% confidence interval of 72.3–77.4%. User and producer accuracies (Table 3) are

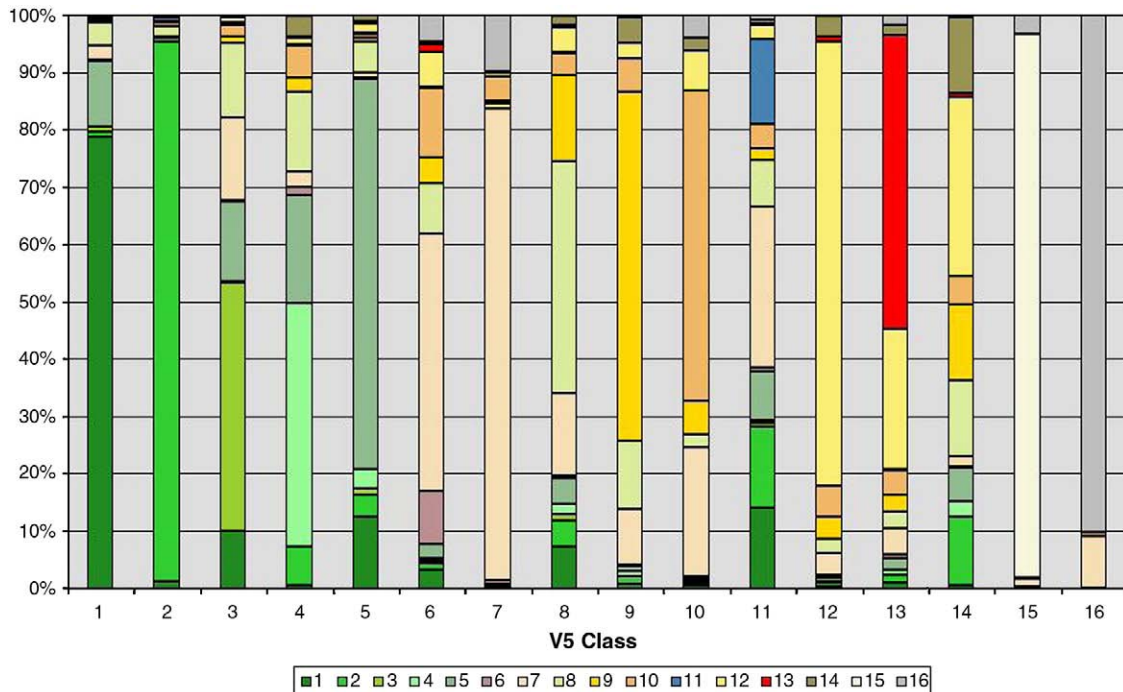


Fig. 14. Proportion of pixels derived from each class in Collection 4 for each class in Collection 5.

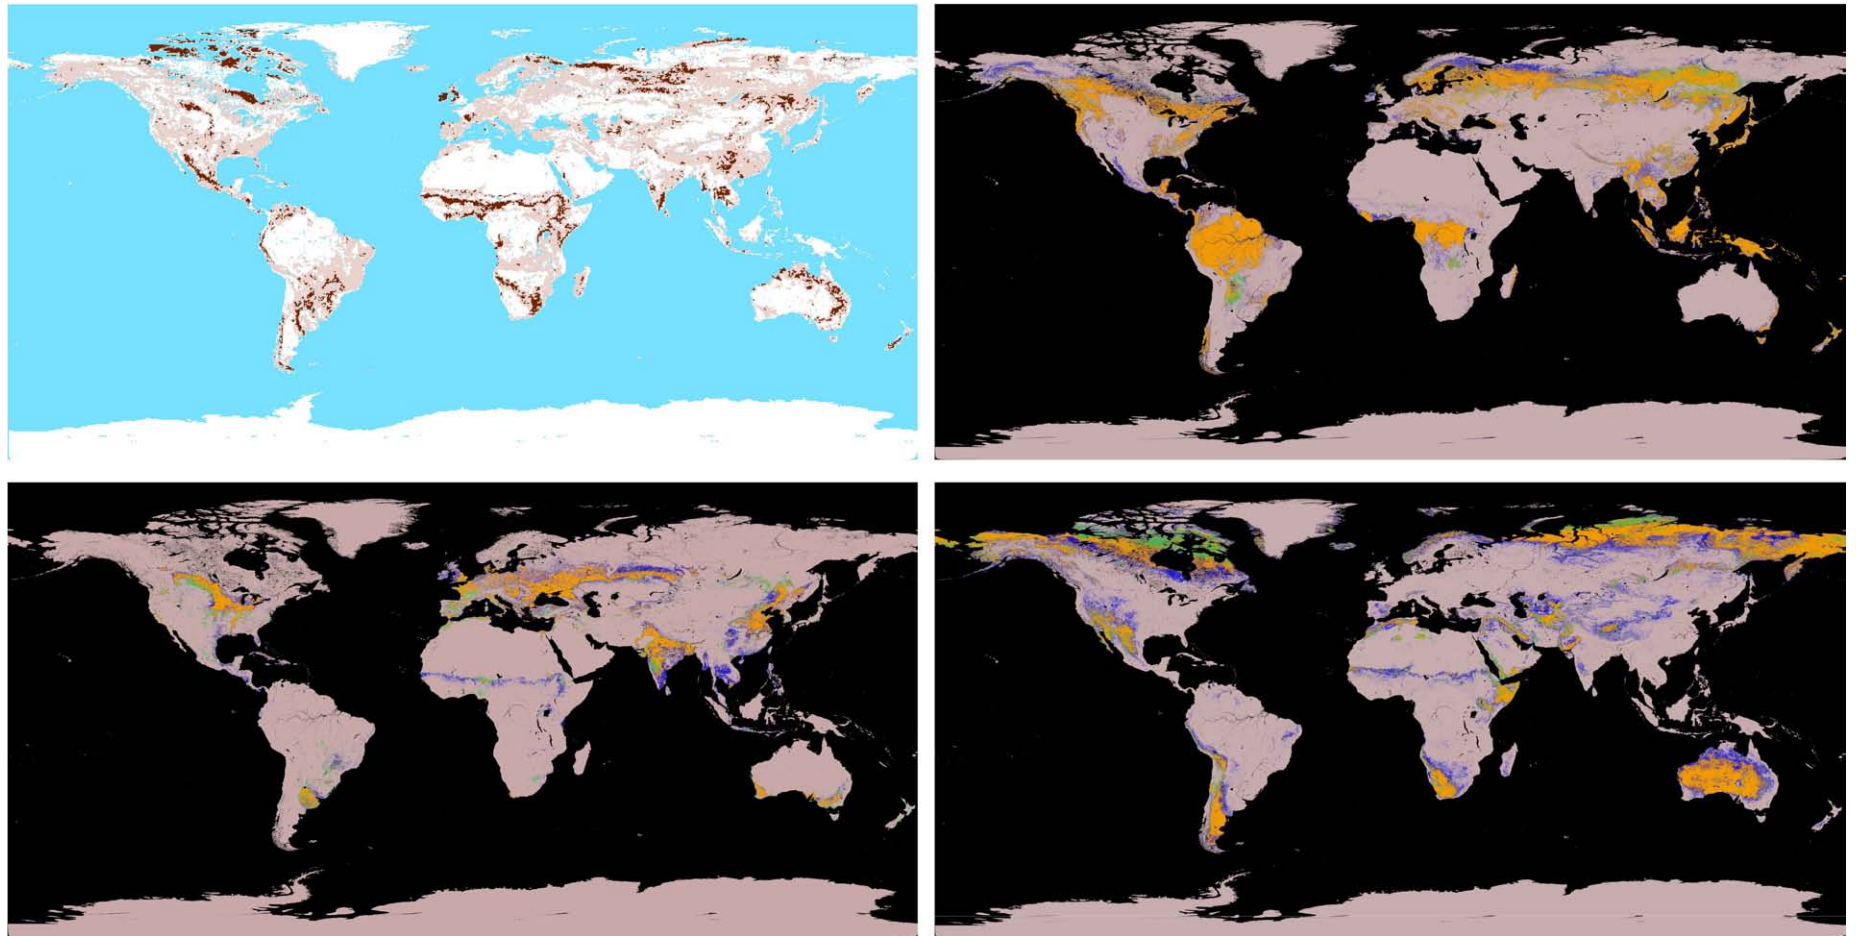


Fig. 15. Maps showing the geographic distribution of major differences in Collection 4 versus Collection 5 for key classes. Upper left: Percentage total change in 50 × 50km cells: brown = 70–100%; beige = 30–70%; white = 0–30%. Upper right: Forests (classes 1–5); Lower left: Agriculture. Lower right: Shrublands. Color key: orange = same in Collections 4 and 5; blue = Collection 4 only; green = Collection 5 only.

generally greater than 70%, but some classes show especially low accuracies. In particular, open shrublands, woody savannas and savannas exhibit low producer accuracies, while mixed forests, closed shrublands, savannas and woody savannas, grasslands and agricultural mosaic show low user accuracies. On a more positive note, the forest classes (1–5) show generally good accuracies, as did agriculture. Not surprisingly, the water, snow and ice, and barren and sparsely vegetated classes showed high user and producer accuracies. Note that the urban class was not included in this analysis. A separate analysis using a large sample of independent validation sites indicates that the accuracy of this class is about 93% (Schneider et al., in press).

Table 4 presents the confusion matrix produced by the cross-validation analysis and helps to explain many of the patterns observed in Table 3. In particular, it is clear from Table 4 that the source of low user and producer accuracies, and by extension most of the error in the map, arises from confusion among a subset of ecologically similar classes. Confusion between savannas and woody savannas (classes 8, 9) is substantial. Woody savannas are also confused with forest classes 1 and 4, and agricultural mosaic (class 14), and open shrublands are confused with the closed shrublands, grasslands, and barren and sparsely vegetated classes. These patterns demonstrate that classification errors are largely concentrated among classes that encompasses ecological and biophysical gradients, and that are quite similar both functionally and in terms of their spectral-temporal properties. They also suggest that depending on user needs, it may make sense to improve map quality by aggregating classes (e.g., classes 8 and 9).

8. Comparison with Collection 4 land cover type

To conclude our analysis, we present an overview of major changes in the mapped distribution of land cover classes in Collection 5 relative to Collection 4. We summarize the between and within-class differences, along with geographic patterns in these differences. To make this comparison, we compare data for 2004 (the last year the Collection 4 data were produced) using Collection 4 data resampled to 500-m to allow direct comparison with Collection 5 results.

Collection 5 presents a substantially different representation of global land cover relative to Collection 4. Overall, about 31% of the Earth's land surface is labeled differently. These differences can be attributed to a number of different sources. First, changes to the algorithm, input features, and training data lead to different results. Second, Collection 5 NBAR data provide a four-fold increase in spatial resolution. Third, many of the differences between Collections 4 and 5 are among classes that are geographically, ecologically, and spectrally similar. This last point is closely related to the first two: higher spatial resolution and refined training and input data provide a better basis in Collection 5 for distinguishing among classes in heterogeneous landscapes relative to Collection 4. Conversely, areas with large expanses of uniform cover (e.g., sub-tropical deserts, undisturbed tropical forests, etc) are relatively unchanged.

Fig. 13 presents a bar chart showing the number of 500-m pixels classified in each of the IGBP classes in Collections 4 and 5. The area occupied by most classes changed by roughly 5–15% (note that the differences do not reflect the total change, but only changes in the total area occupied by each class). Open shrublands, which was by far the largest class in Collection 4, decreased substantially. Conversely, agricultural mosaic doubled relative to Collection 4. Wetlands, a small but important class, tripled in area.

Fig. 14 presents a more detailed characterization of the differences. For each IGBP class in Collection 5, this figure shows the proportion of pixels derived from each IGBP class in Collection 4. Some classes in Collection 5, for example classes 1, 2, 7, 15, and 16, are relatively consistent with Collection 4. The remaining classes, on the other hand, show substantial diversity in their Collection 4 IGBP class labels.

However, close inspection shows that many (but not all) of the differences reflect changes among similar classes, as we described above. For example, most of the “changed” pixels in class 5 (mixed forests) were labeled as other forest classes in Collection 4. Similarly, many pixels labeled as class 8 (woody savannas) in Collection 5 were previously labeled as one of the forest classes, woody savannas, or grasslands in Collection 4, which together encompass a gradient in tree cover and climate regimes.

By mapping the locations where Collections 4 and 5 differ in their class labels, geographic patterns in the results presented in Figs. 13 and 14 become clearly evident (Fig. 15). Space does not allow a detailed presentation of these patterns and here we focus on three main cover types: agriculture, forests, and wetlands. The area mapped as Forests (classes 1–5) in Collection 5 generally decreased, both in the high latitudes and in the tropics. In both regions, forests in Collection 4 are widely replaced in Collection 5 with woodland classes (savannas, woody savannas) that are representative of high latitude and tropical woodland and forest ecosystems. Extensive regions of open shrublands in subarctic and sub-tropical regions in Collection 4 were replaced by the grassland and barren and sparsely vegetated classes, except in the high arctic, where open shrublands expanded. The area occupied by agriculture decreased in central Eurasia, parts of tropical Asia, and the Sahel in Collection 5, but increased in the central United States, Europe, and China. In all three cases, much of the difference was associated with changes to and from the agricultural mosaic class.

9. Discussion and conclusions

The Collection 4 MODIS land cover type product was publicly released in 2004. In the intervening years, the algorithm and datasets used to produce this product have been substantially revised. As a result, the Collection 5 product, which was released in late 2008, is considerably different from the Collection 4 product. The goal of this paper is to document the main changes that have been made to the algorithm and to provide a general characterization of how the Collection 5 product differs from Collection 4.

The most important change in Collection 5 is that the MODIS Land Cover Type Product is being produced at 500-m spatial resolution. This alone is a major refinement and substantially modifies the product relative to Collection 4. In addition, as we have documented in this paper, the MLCT algorithm and input data have been heavily revised. Training data and ancillary datasets used to produce the product have been updated, input features have been modified, and the methods used to post-process the ensemble decision tree results (sample bias and spatial prior probability adjustments) have been refined. The Collection 5 product includes a new urban layer that was produced independent of the main algorithm using methods specifically developed for global mapping of urban land cover, and class-specific fixes designed to improve mapping of wetlands and deciduous needleleaf forests were implemented. Finally, a method was developed to help stabilize classification results and reduce the level of spurious year-to-year change in the datasets. Cross-validation accuracy assessment indicates an overall accuracy of 75%, with substantial variability in class-specific accuracies.

The net result is a substantially different and revised representation of global land cover relative to Collection 4. The higher resolution afforded by 500-m MODIS data in combination with changes to the algorithm produce considerable differences between Collections 4 and 5. The most prominent differences are that the area mapped as forests and open shrublands decreased, while the area mapped as grassland, savannas and agricultural mosaic classes increased. More generally, differences between Collections 4 and 5 are ubiquitous outside of the forested tropics, sub-tropical deserts, and polar ice sheets, regions where the Earth's land surface is characterized by large tracts of uniform land cover.

Moving forward, we plan to address a number of important challenges in Collection 6. First, we plan to migrate the classification system used by MODIS to a system that is consistent with the FAO Land Cover Classification System (Ahlqvist, 2008; Jansen & Di Gregorio, 2002). This will provide a classification that conforms to international community standards, and which provides a more effective distinction between land cover and land use. In addition, we plan to reduce the need for extensive tuning, adjustments, and class-specific solutions that are currently part of the algorithm. In the final analysis, it will probably never be possible to create high quality global land cover maps in a fully automated fashion without some manual adjustment and refinement. However, greater automation and repeatability is a desirable goal and the Collection 5 version of the MODIS land cover product represents a significant step forward in this regard. Indeed, with each new collection of MODIS data we move closer to this objective.

The land cover remote sensing community is rapidly moving towards higher spatial resolution products based on Landsat and other medium resolution sensors at continental and larger scales. Ten years ago, the prospect of global operational processing of 500-m MODIS data was formidable. As the MODIS land cover and other products have demonstrated, robust, repeatable, and semi-automated mapping of global land cover and other land surface variables from remote sensing is both feasible and useful to the science community. Moving forward, significant challenges exist in merging high frequency moderate resolution observations from sensors like MODIS with lower frequency but higher spatial resolution sensors such as Landsat. Experience gained from producing the MODIS global land cover type product should support and inform this next generation of higher resolution products.

Acknowledgements

The research described in this paper was supported by NASA Cooperative Agreement Number NNX08AE61A. The authors also thank the many individuals who have contributed to this effort over the years. We also thank Steve Stehman for informal and helpful discussions related to accuracy assessment.

References

- Ahlqvist, O. (2008). In search of classification that supports the dynamics of science: the FAO Land Cover Classification System and proposed modifications. *Environment and Planning, B – Planning and Design*, 35, 169–186.
- Arino, O., Bicheron, P., Achard, F., Latham, J., Witt, R., & Weber, J. L. (2008). GLOBCOVER – The most detailed portrait of Earth. *ESA Bulletin-European Space Agency*, 24–31.
- Bartholome, E., & Belward, A. S. (2005). GLC2000: A new approach to global land cover mapping from Earth observation data. *International Journal of Remote Sensing*, 26, 1959–1977.
- Bonan, G. B., Levis, S., Kergoat, L., & Oleson, K. W. (2002). Landscapes as patches of plant functional types: An integrating concept for climate and ecosystem models. *Global Biogeochemical Cycles*, 16.
- Bonan, G. B., Oleson, K. W., Vertenstein, M., Levis, S., Zeng, X. B., Dai, Y. J., et al. (2002). The land surface climatology of the community land model coupled to the NCAR community climate model. *Journal of Climate*, 15, 3123–3149.
- Cochran, W. G. (1977). *Sampling Techniques*, 3rd Edition. Wiley, New York, 428 pp.
- DeFries, R., Hansen, M., & Townshend, J. (1995). Global discrimination of land cover types from metrics derived from AVHRR Pathfinder data. *Remote Sensing of Environment*, 54, 209–222.
- DeFries, R. S., & Townshend, J. R. G. (1994). NDVI-derived land-cover classifications at a global-scale. *International Journal of Remote Sensing*, 15, 3567–3586.
- Ek, M. B., Mitchell, K. E., Lin, Y., Rogers, E., Grunmann, P., Koren, V., et al. (2003). Implementation of Noah land surface model advances in the National Centers for Environmental Prediction operational mesoscale Eta model. *Journal of Geophysical Research-Atmospheres*, 108.
- Ellis, E. C., & Ramankutty, N. (2008). Putting people in the map: Anthropogenic biomes of the world. *Frontiers in Ecology and the Environment*, 6, 439–447.
- Foley, J. A., DeFries, R., Asner, G. P., Barford, C., Bonan, G., Carpenter, S. R., et al. (2005). Global consequences of land use. *Science*, 309, 570–574.
- Foody, G. M. (2002). Status of land cover classification accuracy assessment. *Remote Sensing of Environment*, 80, 185–201.
- Freund, Y., & Schapire, R. E. (1996). A decision-theoretic generalization of on-line learning and an application to boosting. *Computational Learning Theory: Second European Conference, EuroCOLT'95* (pp. 23–27).
- Friedl, M. A., Brodley, C. E., & Strahler, A. H. (1999). Maximizing land cover classification accuracies produced by decision trees at continental to global scales. *IEEE Transactions on Geoscience and Remote Sensing*, 37, 969–977.
- Friedl, M. A., McIver, D. K., Hodges, J. C. F., Zhang, X. Y., Muchoney, D., Strahler, A. H., et al. (2002). Global land cover mapping from MODIS: Algorithms and early results. *Remote Sensing of Environment*, 83, 287–302.
- Friedl, M. A., Muchoney, D., McIver, D., Gao, F., Hodges, J. C. F., & Strahler, A. H. (2000). Characterization of North American land cover from NOAA-AVHRR data using the EOS MODIS land cover classification algorithm. *Geophysical Research Letters*, 27, 977–980.
- Friedl, M. A., Woodcock, C., Gopal, S., Muchoney, D., Strahler, A. H., & Barker-Schaaf, C. (2000). A note on procedures used for accuracy assessment in land cover maps derived from AVHRR data. *International Journal of Remote Sensing*, 21, 1073–1077.
- Friedman, J., Hastie, T., & Tibshirani, R. (2000). Additive logistic regression: A statistical view of boosting. *Annals of Statistics*, 28, 337–374.
- Ganguly, S., Friedl, M. A., Tan, B., Zhang, X., & Verma, M. (in press). Vegetation Phenology from MODIS: Characterization of the Collection 5 Global Land Cover Dynamics Product, submitted, *Remote Sensing of Environment*.
- Goldewijk, K. K. (2001). Estimating global land use change over the past 300 years: The HYDE Database. *Global Biogeochemical Cycles*, 15, 417–433.
- Hansen, M. C., Defries, R. S., Townshend, J. R. G., & Sohlberg, R. (2000). Global land cover classification at 1 km spatial resolution using a classification tree approach. *International Journal of Remote Sensing*, 21, 1331–1364.
- Hansen, M. C., Defries, R. S., Townshend, J. R. G., Sohlberg, R., Dimiceli, C., & Carroll, M. (2002). Towards an operational MODIS continuous field of percent tree cover algorithm: Examples using AVHRR and MODIS data. *Remote Sensing of Environment*, 83, 303–319.
- Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sensing of Environment*, 83, 195–213.
- Jansen, L. J. M., & Di Gregorio, A. (2002). Parametric land cover and land-use classifications as tools for environmental change detection. *Agriculture Ecosystems & Environment*, 91, 89–100.
- Lloyd, D. (1990). A phenological classification of terrestrial vegetation cover using shortwave vegetation index imagery. *International Journal of Remote Sensing*, 11, 2269–2279.
- Lotsch, A., Tian, Y., Friedl, M. A., & Myneni, R. B. (2003). Land cover mapping in support of LAI and FPAR retrievals from EOS-MODIS and MISR: Classification methods and sensitivities to errors. *International Journal of Remote Sensing*, 24, 1997–2016.
- Loveland, T. R., & Belward, A. S. (1997). The IGBP-DIS global 1 km land cover data set, DISCover: First results. *International Journal of Remote Sensing*, 18, 3291–3295.
- Loveland, T. R., Merchant, J. W., Brown, J. F., Ohlen, D. O., Reed, B. C., Olson, P., et al. (1995). Seasonal land-cover regions of the United States. *Annals of the Association of American Geographers*, 85, 339–355.
- Loveland, T. R., Reed, B. C., Brown, J. F., Ohlen, D. O., Zhu, Z., Yang, L., et al. (2000). Development of a global land cover characteristics database and IGBP DISCover from 1 km AVHRR data. *International Journal of Remote Sensing*, 21, 1303–1330.
- Lubchenco, J. (1998). Entering the century of the environment: A new social contract for science. *Science*, 279, 491–497.
- Matthews, E. (1983). Global vegetation and land-use – new high-resolution data-bases for climate studies. *Journal of Climate and Applied Meteorology*, 22, 474–487.
- McIver, D. K., & Friedl, M. A. (2001). Estimating pixel-scale land cover classification confidence using nonparametric machine learning methods. *IEEE Transactions on Geoscience and Remote Sensing*, 39, 1959–1968.
- McIver, D. K., & Friedl, M. A. (2002). Using prior probabilities in decision-tree classification of remotely sensed data. *Remote Sensing of Environment*, 81, 253–261.
- Myneni, R. B., Hoffman, S., Knyazikhin, Y., Privette, J. L., Glassy, J., Tian, Y., et al. (2002). Global products of vegetation leaf area and fraction absorbed PAR from year one of MODIS data. *Remote Sensing of Environment*, 83, 214–231.
- Olson, J. S. (1982). Earth's vegetation and atmospheric carbon dioxide. In W. C. Clark (Ed.), *Carbon Dioxide Review: 1982* (pp. 388–399). New York: Oxford University Press.
- Quinlan, J. R. (1993). *[C]4.5: Programs for Machine Learning*. San Mateo, CA: Morgan Kaufmann.
- Quinlan, J. R. (1996). Bagging, boosting, and C4.5. *Proceedings of the Thirteenth National Conference on Artificial Intelligence, Portland, OR* (pp. 725–730).
- Ramankutty, N., Evan, A. T., Monfreda, C., & Foley, J. A. (2008). Farming the planet: 1. Geographic distribution of global agricultural lands in the year 2000. *Global Biogeochemical Cycles*, 22.
- Ramankutty, N., & Foley, J. A. (1999). Estimating historical changes in global land cover: Croplands from 1700 to 1992. *Global Biogeochemical Cycles*, 13, 997–1027.
- Running, S. W., & Coughlan, J. C. (1988). A general-model of forest ecosystem processes for regional applications. 1. hydrologic balance, canopy gas-exchange and primary production processes. *Ecological Modelling*, 42, 125–154.
- Running, S. W., Loveland, T. R., Pierce, L. L., Nemani, R., & Hunt, E. R. (1995). A remote-sensing based vegetation classification logic for global land-cover analysis. *Remote Sensing of Environment*, 51, 39–48.
- Sanderson, E. W., Jaitheh, M., Levy, M. A., Redford, K. H., Wannebo, A. V., & Woolmer, G. (2002). The human footprint and the last of the wild. *Bioscience*, 52, 891–904.
- Schaaf, C. B., Gao, F., Strahler, A. H., Lucht, W., Li, X. W., Tsang, T., et al. (2002). First operational BRDF, albedo nadir reflectance products from MODIS. *Remote Sensing of Environment*, 83, 135–148.
- Schapire, R. E., Freund, Y., Bartlett, P., & Lee, W. S. (1998). Boosting the margin: A new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26, 1651–1686.
- Schneider, A., Friedl, M. A., McIver, D. K., & Woodcock, C. E. (2003). Mapping urban areas by fusing multiple sources of coarse resolution remotely sensed data. *Photogrammetric Engineering and Remote Sensing*, 69, 1377–1386.

- Schneider, A., Friedl, M. A., & Potere, D. (in press). A new map of urban extent from MODIS satellite data. *Environmental Research Letters*.
- Sellers, P. J., Dickinson, R. E., Randall, D. A., Betts, A. K., Hall, F. G., Berry, J. A., et al. (1997). Modeling the exchanges of energy, water, and carbon between continents and the atmosphere. *Science*, 275, 502–509.
- Stehman, S. V. (1997). Estimating standard errors of accuracy assessment under cluster sampling. *Remote Sensing of Environment*, 60, 258–269.
- Sterling, S., & Ducharne, A. (2008). Comprehensive data set of global land cover change for land surface model applications. *Global Biogeochemical Cycles*, 22.
- Stone, T. A., Schlesinger, P., Houghton, R. A., & Woodwell, G. M. (1994). A map of the vegetation of south-america based on satellite imagery. *Photogrammetric Engineering and Remote Sensing*, 60, 541–551.
- Strahler, A. H., Boschetti, L., Foody, G. M., Friedl, M. A., Hansen, M. C., Herold, M., et al. (2006). *Global land cover validation: Recommendations for evaluation and accuracy assessment of global land cover maps*. Luxembourg: European Commission – DG Joint Research Centre, Institute for Environment and Sustainability EUR 22156 EN-DG, 48 pp.
- Townshend, J. R. G. (1998). Global data sets for land applications from the advanced very high resolution radiometer: An introduction. *International Journal of Remote Sensing*, 15(17), 3319–3332.
- Townshend, J. R. G., Justice, C. O., & Kalb, V. (1987). Characterization and classification of South American land cover types using satellite data. *International Journal of Remote Sensing* (pp. 1189–1207).
- Vitousek, P. M., Mooney, H. A., Lubchenco, J., & Melillo, J. M. (1997). Human domination of Earth's ecosystems. *Science*, 277, 494–499.
- Wan, Z. M., Zhang, Y. L., Zhang, Q. C., & Li, Z. L. (2002). Validation of the land-surface temperature products retrieved from Terra Moderate Resolution Imaging Spectroradiometer data. *Remote Sensing of Environment*, 83, 163–180.
- Wilson, M., & Henderson-Sellers, A. (1985). A global archive of land cover and soils data for use in general circulation climate models. *Journal of Climatology*, 5, 119–143.
- Zhang, X. Y., Friedl, M. A., & Schaaf, C. B. (2006). Global vegetation phenology from Moderate Resolution Imaging Spectroradiometer (MODIS): Evaluation of global patterns and comparison with in situ measurements. *Journal of Geophysical Research-Biogeosciences*, 111, G04017. doi:10.1029/2006JG000217.