Available online at www.sciencedirect.com
ScienceDirect
Remote Sensing of Environment 109 (2007) 482–499
Remote Sensing of Environment
www.elsevier.com/locate/rse

An empirical InSAR-optical fusion approach to mapping vegetation canopy height

Wayne S. Walker [a,*], Josef M. Kellndorfer [a], Elizabeth LaPoint [b], Michael Hoppus [b], James Westfall [b]

[a] *The Wood Hole Research Center, Falmouth, MA, United States*
[b] *USDA Forest Service, Forest Inventory and Analysis Program, United States*

Received 21 November 2006; received in revised form 1 February 2007; accepted 3 February 2007

Abstract

Exploiting synergies afforded by a host of recently available national-scale data sets derived from interferometric synthetic aperture radar (InSAR) and passive optical remote sensing, this paper describes the development of a novel empirical approach for the provision of regional- to continental-scale estimates of vegetation canopy height. Supported by data from the 2000 Shuttle Radar Topography Mission (SRTM), the National Elevation Dataset (NED), the LANDFIRE project, and the National Land Cover Database (NLCD) 2001, this paper describes a data fusion and modeling strategy for developing the first-ever high-resolution map of canopy height for the conterminous U.S. The approach was tested as part of a prototype study spanning some 62,000 km^2 in central Utah (NLCD mapping zone 16). A mapping strategy based on object-oriented image analysis and tree-based regression techniques is employed. Empirical model development is driven by a database of height metrics obtained from an extensive field plot network administered by the USDA Forest Service–Forest Inventory and Analysis (FIA) program. Based on data from 508 FIA field plots, an average absolute height error of 2.1 m (r=0.88) was achieved for the prototype mapping zone.

Keywords: Vegetation canopy height; Scattering phase center height; InSAR; Radar; Interferometry; Optical; Multi-spectral; SRTM; Landsat ETM+; Forest inventory; FIA; DEM; Object-oriented; Segmentation; Regression trees

1. Introduction

1.1. Motivation

Spatially extensive and accurate maps of vegetation canopy height are of value not only to ecologists and land managers working in diverse fields such as biodiversity conservation, wildfire risk assessment, and timber production, but also to climate change scientists focused on reducing the uncertainty associated with the carbon cycle component of Earth's climate system. High-resolution maps of canopy height have the potential to significantly improve the accuracy of aboveground biomass and carbon stock baselines upon which models of future climate change necessarily depend. Reliable baseline information is also needed for measuring and monitoring carbon fluxes and for verifying emissions reductions in the context of national and international carbon accounting strategies.

Although the forests of the United States and other mid- to high-latitude nations are covered by extensive inventory plot networks, these data are largely inadequate for the provision of high-resolution estimates of aboveground biomass and carbon stocks. Whereas dry biomass, which contains 45 to 50% carbon by weight (Linder & Axelsson, 1982; Reichle et al., 1973), may be well quantified for the localized areas where measurements exist, extrapolation across larger unsampled regions can contribute to considerable estimate uncertainty (Houghton & Goodale, 2004). Consequently, at regional to continental scales, estimates of multi-dimensional forest structural metrics are necessarily acquired through the use of remote sensing technologies in concert with ground-based measurements derived from national forest inventories. The practice of leveraging the combined strengths of forest inventory and

* Corresponding author. Tel.: +1 508 540 9900.
E-mail address: wwalker@whrc.org (W.S. Walker).

0034-4257/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.rse.2007.02.001

satellite image data dates back to the early 1990s in Finland (Tomppo, 1991). More recent examples include applications in northern Europe and the United States (Huang et al., 2002; McRoberts & Liknes, 2005; Reese et al., 2002, 2003; Tomppo et al., 2002).

Numerous approaches have been put forth for the provision of aboveground biomass estimates using the range of available remote sensing technologies including passive optical (e.g., Dong et al., 2003; Myneni et al., 2001), radar (e.g., Dobson et al., 1992; Ranson et al., 1997), and lidar (e.g., Drake et al., 2002; Hyde et al., 2005; Lefsky et al., 1999a,b); however, a technique has yet to be presented that is consistent, reproducible, and applicable across broad geographic extents (Rosenqvist et al., 2003). This is largely due to the fact that biomass is a three-dimensional metric — the accurate estimation of which requires biophysical measures, and therefore remote sensors, that capture both the horizontal (e.g., canopy density/cover) and vertical (e.g., canopy height) structural character of the vegetation (Mette & Hajnsek, 2003; Mette et al., 2004; Treuhaft et al., 2004). While the science of acquiring remotely sensed estimates of horizontal vegetation structure has matured considerably over the past 25 years, only in the last decade have significant advances in instrument development made it possible to obtain consistent and accurate measurements of canopy height and related metrics of vertical vegetation structure (e.g., Lefsky et al., 2002; Treuhaft & Siqueira, 2000). Motivated by these advancements, this research focuses on the three-dimensional structure of forest vegetation in an effort to expand the scientific basis for regional- to continental-scale carbon accounting. Specifically, this research presents an approach to the generation of high-resolution, spatially extensive maps of vegetation canopy height. The approach is the foundation for an ongoing NASA-sponsored project with the ultimate goal of generating the first-ever circa-2000 baseline dataset of vegetation canopy height, aboveground biomass, and carbon stocks for the conterminous U.S. This project is possible, in part, because of the complimentary nature and quasi-synchronous development of several national digital geospatial datasets. The following section provides a brief introduction to these datasets.

1.2. Confluence of national mapping efforts

The last several years have been marked by an unprecedented confluence of high-resolution geospatial data sources and derived products for the conterminous U.S. The first of these datasets was acquired early in 2000 when the NASA-JPL Shuttle Radar Topography Mission (SRTM) used C-band (5.6 cm, 5.3 GHz) interferometric synthetic aperture radar technology (InSAR) to obtain high-resolution (one arc-second) elevation data on a near-global scale for the purpose of generating the most complete digital topographic database of Earth. Rather than reflecting the "bald-earth" surface, an SRTM-derived digital elevation model (DEM) is unique in that it more closely reflects the elevation surface formed by vegetation (e.g., tree canopies) and anthropogenic features (e.g., buildings, towers, etc.). Assuming the elevation of the bald-earth surface is known, an estimate of the interferometric "scattering phase center height" (h_{spc}) can be computed (Brown, 2003; Brown & Sarabandi, 2003; Kellndorfer et al., 2004; Kobayashi et al., 2000; Saich et al., 2001). It follows that the value of h_{spc} is correlated with both the amount and height of vegetation present. Recent research has confirmed the feasibility of using SRTM DEMs together with bald-earth topography data to estimate the height of vegetation canopies (Brown, 2003; Brown & Sarabandi, 2003; Kellndorfer et al., 2004; Walker et al., 2007).

A second dataset with considerable potential to provide information on the horizontal structure of forests is the 2001 National Land Cover Dataset (NLCD; Homer et al., 2004). This multi-layer dataset, currently being developed by the Multi-Resolution Land Characteristics (MRLC) Consortium, uses an ecoregional mapping approach and consists of 1) normalized Tasseled Cap (TC) transformations of Landsat 7 ETM+ imagery from three time periods (early, peak, and late growing season), 2) classified land cover data derived from TC imagery, 3) independent image derivatives of imperviousness and tree canopy density, and 4) independent ancillary data layers including DEM derivatives of slope, aspect and elevation derived from the National Elevation Dataset (NED), which was seamlessly compiled for the entire United Sates for the first time in 1999. All data layers are being released at a grid spacing of 30 m.

A third and final dataset, also under active development, is the multi-partner Landscape Fire and Resource Management Planning Tools Project (LANDFIRE). LANDFIRE is an ecosystem, wildland fire, and wildland fuels mapping project designed to generate a comprehensive suite of spatial data layers describing wildland fuel, existing vegetation composition and structure, historical vegetation conditions, and historical fire regimes. A set of more than 20 national map products is being produced by LANDFIRE using the NLCD ecoregional mapping approach. Specific deliverables include maps of mean fire return interval, percent fire severity, and successional class, as well as existing vegetation type, canopy cover, and canopy height. The canopy height product is currently in development and is slated to be released as a discrete (i.e., five forested height classes) data layer. Aboveground biomass and carbon stocks are not being mapped as part of the LANDFIRE project. Consistent with the NLCD, all LANDFIRE data layers are being released at a grid spacing of 30 m.

The success of a mapping project such as the one proposed here depends largely on the availability of a suitable ground reference database. Complimenting the aforementioned assemblage of national spatial datasets is a national ground reference database available as part of the Forest Inventory and Analysis (FIA) program administered by the USDA Forest Service. In continuous operation since 1930, the FIA program is the only nationwide source of timely, consistent, and reliable forest inventory and monitoring information. The FIA Database (FIADB) contains plot-level forest biometric information collected repeatedly at more than 125,000 locations throughout the United States.

Given the highly complementary nature and quasi-synchronous development of the SRTM, NLCD, and LANDFIRE data sources, an exceptional opportunity exists for exploiting

InSAR/optical synergies. Whereas the SRTM InSAR data provide information pertaining to the vertical structure, i.e., primarily vegetation height, several optically-derived layers provided as part of the NLCD and LANDFIRE projects are suitable for characterizing key aspects of horizontal structure (i.e., vegetation type, canopy cover/density, etc.).

1.3. Objectives

Building on knowledge gained in the context of research conducted by Kellndorfer et al. (2004), Pierce et al. (2006), and Walker et al. (2007), the general objective of this article is to present the results of a proof-of-concept study focused on development of a robust empirical approach for generating a high-resolution, year-2000 baseline estimate of vegetation canopy height for the conterminous U.S. The approach utilizes data fusion, knowledge-based image segmentation, and regression-tree techniques to synergistically exploit the information content of the SRTM interferometric data together with that of data layers obtained from the NED, NLCD and LANDFIRE datasets.

To facilitate development, implementation, and evaluation of the proof-of-concept study, as well as enable future nationwide implementation of the approach, the ecoregional "mapping-zone" concept developed as part of the NLCD 2001 project was adopted for use. The concept, which has also been implemented by the LANDFIRE project, was developed in order to simplify the process of large-scale land cover mapping by stratifying the nation into 66 sub-regions that represent relative homogeneity in terms of biophysical (landform, soil, and vegetation) and spectral characteristics (Homer & Gallant, 2001; Homer et al., 2004). For the purposes of this proof-of-concept study mapping zone 16 (MZ16) was chosen. The zone, which spans over 62,000 km^2 including portions of central Utah, southeastern Idaho, and southwestern Wyoming, was selected because it was the first zone for which all data layers relevant to this research, particularly those currently under production as part of the NLCD and LANDFIRE projects, were available.

2. Mapping zone 16 description

The boundary of MZ16 largely follows that of Ecoregion 19 (Wasatch and Uinta Mountains) of the United States Environmental Protection Agency's Level III Ecoregions of the Conterminous United States (Woods et al., 2001). The zone is composed of a core area of high-elevation, steep, rugged mountains with narrow crests and valleys. This core is flanked in some areas by dissected plateaus and open high mountains (Woods et al., 2001). Elevations within the zone range from 1450 to 4100 m. Over half of the zone is forested, with both vegetation and underlying soils following a pattern of elevational zonation. Low elevations are typically characterized by grasses and a variety of shrubs (often heavily grazed) including sagebrush, chaparral, and mahogany. Low to middle elevations (also grazed) are covered by a range of vegetation types, which include oak and pinyon–juniper woodlands, as well as areas of chaparral, aspen (*Populus tremuloides*), ponderosa pine (*Pinus ponderosa*), and Douglas-fir (*Pseudotsuga menziesii*). Middle to high elevations tend to be covered by large continuous tracts of coniferous forest that include Engelmann spruce (*Picea engelmannii*), subalpine (*Abies lasiocarpa*) and white fir (*Abies concolor*), as well as bristlecone (*Pinus longaeva*), limber (*Pinus flexillis*), and lodgepole pine (*Pinus contorta*). The highest peaks rise well above tree-line and are characterized by alpine vegetation.

3. Model development database

A prerequisite to the construction of multivariate tree-based regression models relating observed canopy height to SRTM and other remote sensing and ancillary data is the compilation of a model development database (MDDB). The MDDB consists of multiple records corresponding to the number of reference observations (i.e., FIA field plots) available within the mapping zone. Each record contains multiple fields, which correspond to the specific response (derived from the FIADB) and predictor (derived from remote sensing and ancillary data sources) variables on which modeling is to be based. The following section describes the various data acquisition, image processing, and computational steps involved in compilation of the MDDB. A diagrammatic summary of these steps is presented in Fig. 1.

3.1. Data acquisition and preprocessing

3.1.1. SRTM and NED data

For a complete description of both the SRTM and NED digital elevation data, the reader is directed to Kellndorfer et al. (2004). The SRTM C-band and NED DEMs for MZ16 were acquired from the United States Geological Survey (USGS) EROS Data Center (Dean Gesch, *pers. comm.*) in the form of 17 individual raster image tiles each covering an area of one degree by one degree. Tiles from each dataset were mosaiced and an SRTM minus NED difference (SRTMDIFF) image was calculated based on the rationale put forth by Kellndorfer et al. (2004). A topographic slope (SLP) layer was also generated from the NED DEM.

3.1.2. NLCD 2001 data

A detailed summary of the NLCD 2001 data-layer production methods is presented by Huang et al. (2001) and Homer et al. (2004). NLCD 2001 data were acquired from the USGS EROS Data Center (Dean Gesch, *pers. comm.*) and included layers of land cover (LC) and canopy density (CD). The LC layer, acquired primarily for reference purposes, consists of 17 classes that generally approximate the thematic detail represented in the Level II classification of Anderson et al. (1976). Developed independently of the LC layer, the CD layer depicts the spatial distribution of tree canopy density (trees $\geq$5 m tall) as a continuous variable with values ranging from 1 to 100% (Huang et al., 2001).

3.1.3. LANDFIRE data

A single data layer, existing vegetation type (EVT), was acquired from the LANDFIRE project website (www.landfire.gov), which is the primary source for additional information on the LANDFIRE project and individual data-layer production

methods. The EVT classification scheme is based on the terrestrial ecological systems classification developed by NatureServe for the Western Hemisphere (Comer et al., 2003). The scheme was developed to provide a practical, mid-scale classification unit that could be readily mapped from remotely sensed imagery and readily identified by managers in the field. Of the 599 ecological systems (hereafter referred to as classes) identified within the U.S., 60 were observed to occur within MZ16. Of these, 53 were characterized by some form of woody or herbaceous vegetation while the remaining seven were non-vegetated, e.g., water, permanent snow/ice, barren, developed, etc. Of the vegetated classes, 23 were forested.

3.1.4. FIA data

The monitoring component of the FIA program consists of a systematic sample across all public and private lands in the U.S. In the late 1990s, the FIA program adopted a common field plot design consisting of four 1/24th acre fixed-radius (24.0 ft/7.3 m) subplots (FIA, 2004). Field plots are distributed across the landscape with approximately one sample (FIA plot) every

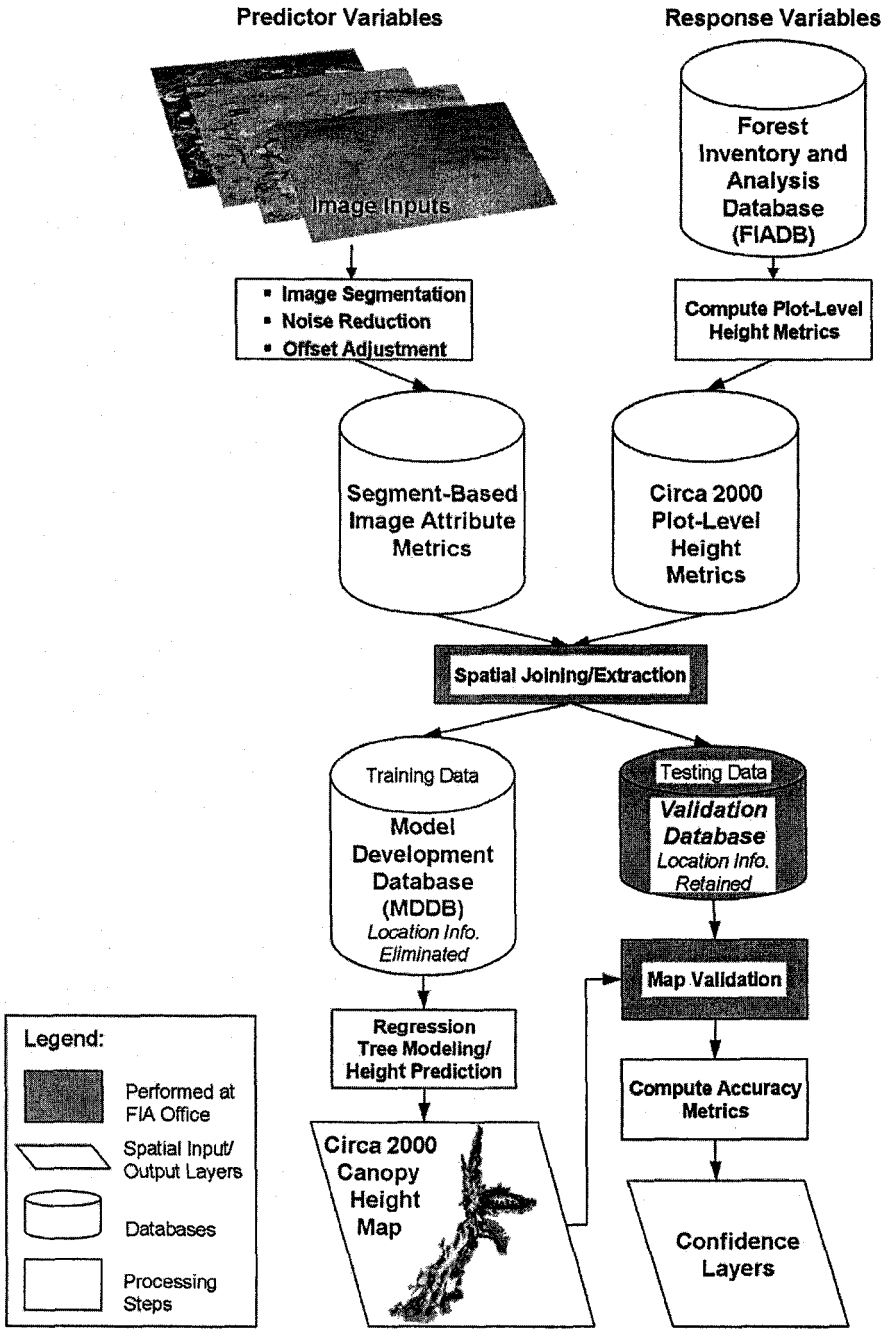

Fig. 1. Process flow diagram depicting the principal steps in the height mapping approach.

6000 acres (≈2428 ha). Each plot (i.e., an assemblage of four subplots) is required to have a tree stocking of at least 10% within a one-acre (0.4 ha) neighborhood. Field crews collect plot-level data on forest type, site attributes, tree species, and tree size including stem diameter and height, and overall tree condition. For further information on the FIA program, the reader is directed to the national FIA website (www.fia.fs.fed.us).

FIA data for MZ16 were downloaded from Version 1.7 of the FIA database (FIADB) (www.ncrs2.fs.fed.us/4801/FIADB). Due to issues involving the timing of data availability, only MZ16 data occurring within the State of Utah (95% of the mapping zone by area) could be accessed. Data from the Utah acquisition were collected during Subcycles 1–4 (i.e., 2000–2003) of Utah's second state-wide inventory. Although the latter Subcycles were surveyed post-2000, i.e., more recently than the year on which this baseline mapping project is focused, inclusion of data from this broader temporal range was deemed acceptable given 1) the relative slow growth of trees in the dry basins and montane landscapes of MZ16 and 2) the need for a sufficient sample size on which to base model development. No attempt was made to adjust (i.e., back cast) the heights of trees measured in the latter Subcycles (post-2000). As a whole, these data represent the most consistent, complete, and accurate source of forest biometric information for the region during the time period of interest.

The FIADB is a relational database consisting of twelve hierarchical tables. Prior to computing plot-level canopy height metrics from the UTAH FIADB, the initial database, containing 3665 plots, was filtered to extract all forested plots, i.e., plots containing trees (stems ≥5.0 in/12.7 cm in diameter) based on entries in the FIADB Tree Table. The filtering procedure produced 1395 plots. Using data contained in the ACTUALHT (i.e., actual height) field of the FIADB TREE Table, three plot-level canopy height metrics were computed including basal-area weighted average height (BAWHT), average height (AVGHT), and maximum height (MAXHT). The BAWHT of each plot was calculated according to:

$$\mathrm{BAWHT} = \frac{1}{\mathrm{BA_{Plot}}} \sum_{i=1}^{n} (\mathrm{BA}_i * \mathrm{ACTUALHT}_i) \tag{3.1}$$

where BA_i is the basal area (m^2) of the ith tree in the plot and is calculated according to:

$$\mathrm{BA} = 0.00007854 * \mathrm{DBH}^2 \tag{3.2}$$

(DBH is the diameter at breast height (cm) from the TREE Table) and $\mathrm{BA_{Plot}}$ is the total basal area (m^2) for the plot and is calculated according to:

$$\mathrm{BA_{Plot}} = \sum_{i=1}^{n} \mathrm{BA}_i \tag{3.3}$$

The AVGHT of each plot was calculated according to:

$$\mathrm{AVGHT} = \frac{1}{n} \sum_{i=1}^{n} \mathrm{ACTUALHT}_i \tag{3.4}$$

A total of seven FIA-based reference variables were added to the MDDB for use in subsequent model development activities. These variables included FLDTPYCD (forest type code), CONDID (condition class number, e.g., stand density, size, origin, etc.), BAWHT, AVGHT, and MAXHT, as well as the aspect (ASP) of the central subplot and a unique plot identification code (PLOTID).

3.2. Extraction of segment-based image attributes

As documented in Kellndorfer et al. (2004) and Walker et al. (2007), the SRTM DEM data contain residual phase noise errors, which left unmitigated, result in erroneous estimates of the h_{spc} and render the data largely unusable as a legitimate source of canopy height information. In Walker et al. (2007), a knowledge-based strategy to phase noise error mitigation, and ultimately h_{spc} calculation, was proposed. The method, developed further in the context of this research, was implemented across MZ16. The following sections provide a description of the approach as implemented (see also Fig. 1).

3.2.1. Segment-based noise mitigation

A segmentation-based approach to sample (i.e., pixel) aggregation and averaging was brought to bear on the problem of SRTM phase noise reduction. Here the term "segmentation" refers to the subdivision of an image or image stack into a number of regions, i.e., polygons or image objects, based on some predefined criteria (Baatz et al., 2004). The software package eCognition provided the computational framework in which image segmentation was accomplished. Unlike block-filtering techniques which impose rectangular averaging schemes, eCognition provides for the automatic and optimal delineation of local homogenous regions, e.g., irregularly-shaped forest tracts, within which sample averaging and consequent noise reduction can be more smartly constrained.

A segmentation strategy was formulated with the general goal of producing image objects that were 1) of sufficient size to provide for adequate sample averaging and noise reduction in forested regions, 2) homogenous in terms of topographic slope, 3) homogenous in terms of vertical forest structure (i.e., canopy height), and 4) homogenous in terms of horizontal forest structure (i.e., canopy density). The realization of this goal was a challenge given the inherent antagonism of object size and object homogeneity. That is to say, all else being equal, as the average size of image objects increases, so does the amount of topographic and structural heterogeneity observed. An average object size of 15–20 pixels was targeted following an evaluation of previous research by Kellndorfer et al. (2004) and Walker et al. (2007). Objects of this size are deemed large enough to provide for adequate noise reduction under most SRTM datatake regimes while at the same time remaining small enough to allow for sufficient within-object homogeneity. Balancing noise reduction with object homogeneity is essential as both are critical to the generation of robust object-based mean estimates of the h_{spc}.

As far as possible, balance was achieved using a hierarchical (i.e., nested) segmentation approach. Three image layers, including the NED-derived slope (SLP), the SRTM minus NED difference (SRTMDIFF), and the NLCD canopy

density (CD)[1], served as inputs to the segmentation process. The rationale behind this strategy can be summarized as follows. First, a relatively coarse segmentation was generated using slope (SLP) as the primary input. This initial segmentation represented an attempt to incorporate slope information as a way to control for slope-related effects that may have influenced the InSAR response from the terrain. Second, while holding the boundaries of this initial segmentation constant, a meso-scale segmentation was generated in which the SRTMDIFF image was used to further regionalize the existing segments. At this level, the resulting segmentation defined regions in terms of relative homogeneity of both slope and canopy height (i.e., h_{spc}). Third, while holding the boundaries of this second segmentation constant, a relatively fine-scale segmentation was generated in which the CD image was used as the final driver of the regionalization. The decision to progress hierarchically from canopy height to canopy density in the sequence of inputs rather than vice-versa was made in order to ensure that vertical structure was considered at a broader segmentation scale than horizontal structure and, therefore, was the principal driver of segmentation size and subsequently the amount of averaging imposed. Prior to initiating the segmentation procedure, the CD image was selected to provide an analysis mask. In doing so, the segmentation process was constrained to proceed in only those regions where canopy cover was present, i.e., CD values greater than zero.

The segmentation process resulted in a vector layer comprised of over 4.5 million image–object polygons. This vector layer was used to compute a suite of object-based statistical metrics. The first, and perhaps most significant, of these metrics was the mean h_{spc} (MH_{spc}) value, calculated within each of the 4.5 million polygons using data extracted from the SRTMDIFF image. As a result of the object-based averaging procedure, this metric represents a noise-reduced value of the h_{spc} within each polygon. The object-based standard deviation of h_{spc} (SD_{spc}) was calculated as well. In addition to the h_{spc} metrics, object-based means for the SLP, CD and NED images and standard deviations for the CD and NED images were similarly computed. The area (AREA) of each image object was also calculated.

3.2.2. Vertical offset evaluation

As observed by Brown et al. (2005) and Walker et al. (2007), as well as discussed by Kellndorfer et al. (2004) and others, it is possible for both the NED and SRTM DEM data to contain vertical biases of varying magnitudes (centimeters to several meters) (Hensley et al., 2000; Gesch et al., 2002). These biases, typically a result of the DEM production history, are of little concern when the observed offset is a constant value across a region of interest. However, when a nonlinear offset (i.e., a spatial trend) is identified, it must be removed or, at the very least minimized, in order to avoid propagation of this nonlinear error. Nonlinear offsets are commonly manifested as banding or striping in the SRTM–NED difference image.

Following methods described by both Kellndorfer et al. (2004) and Walker et al. (2007), barren and otherwise nonvegetated regions were evaluated to determine if a vertical offset existed between the NED and SRTM DEMs. Although a mean difference (i.e., constant value) of approximately 2.0 m was observed between the DEMs, no significant nonlinear offset or trend was identified.

3.3. Spatial database joining

Although the FIADB, including all plot measurements, is part of the public domain, federal law prohibits the USDA Forest Service from releasing the exact coordinates of FIA plot locations. In an effort to provide access to FIA plot coordinates while maintaining privacy protection and long-term plot integrity, the FIA program established the FIA National Spatial Data Services (NSDS) unit where FIA plot locations can be linked spatially with data acquired from remote sensing/ancillary sources. Because coordinate locations must not leave NSDS computers, all research and development activities involving plot coordinates must be carried out on-site and under the supervision of NSDS staff.

Final compilation of the MDDB required that a spatial join be established between the 1395 FIA plots (and associated reference variables) described in Section 3.1.4 and the eight object-based metrics discussed in the previous section (Fig. 1). The objective of the joining procedure was to extract image–object data from beneath the 1395 FIA plot locations. All data processing was performed at the USDA Forest Service Northern

Table 1
Variables included in the final model development database (MDDB)

Variable	Variable description	FIA-based reference variables	Object-based image variables	Response (R) or predictor (P) variable
PLOTID	Plot Identification Code	X		–
CONDID	Condition Class Number	X		–
FLDTYPCD	Forest Type Code	X		P
BAWHT	Basal-Area Weighted Average Height	X		R
AVGHT	Average Height	X		R
MAXHT	Maximum Height	X		R
ASP	Aspect of Subplot 1	X		P
MH_{SPC}	Mean Scattering Phase Center Height (SRTM)		X	P
SDH_{SPC}	S.D. Scattering Phase Center Height (SRTM)		X	P
MSLP	Mean Slope (NED)		X	P
MCD	Mean Canopy Density (NLCD)		X	P
SDCD	S.D. Canopy Density (NLCD)		X	P
MNED	Mean Elevation (NED)		X	P
SDNED	S.D. Elevation (NED)		X	P
AREA	Image Object Area (eCognition)		X	P

[1] Neither NLCD or LANDFIRE land cover layers were included in the segmentation process in order to avoid the introduction of errors attributable to misclassification. Land cover information was instead incorporated during the process of model development. This topic is addressed further in Section 5.

Research Station, Durham, New Hampshire. Following data extraction, eight object-based image variables were added to the MDDB, bringing the total number of variables to 15. A complete list of variables in the final MDDB is included in Table 1.

Following the spatial joining, the MDDB was filtered further using several criteria. First, filtering was carried out on the mean canopy density (MCD) field to ensure that only cases with a mean canopy density greater than zero were retained. Second, the database was filtered hierarchically on condition class number (CONDID) and forest type code (FLDTYPCD) such that only cases identified by a single condition class (across all subplots) and a single forest type were retained. This was done to ensure the homogeneity of field plots in terms of horizontal and vertical forest structure. Finally, filtering was carried out on the AREA field using a threshold of 1.5 ha (i.e., $\approx$17 pixels). Only cases having an area greater than or equal to 1.5 ha were retained. This final filter was applied to ensure that each FIA plot was associated with an image–object polygon large enough to provide for adequate noise reduction. The filtering procedure resulted in a final MDDB containing a total of 393 FIA plots (i.e., cases; Fig. 2).

4. Model development

4.1. Choice of modeling framework

The availability of regional- to continental-scale datasets derived from remotely sensed and other ancillary sources has motivated recent research into the suitability of available empirical statistical modeling techniques for broad-scale prediction and mapping of forest structural attributes (Frescino et al., 2001; Moisen & Frescino, 2002; Moisen et al., 2003). Working throughout the Interior Western United States, Moisen and Frescino (2002) compared five approaches for modeling six different FIA-derived biophysical response variables and a suite of satellite-derived predictor variables. The approaches included linear models, generalized additive models (GAMS; Guisan et al., 2002), classification and regression trees (CART; Breiman et al., 1984), multivariate adaptive regression splines (MARS; Friedman, 1991), and artificial neural networks. Although the MARS and GAMS models performed marginally better than the others, little appreciable difference among the techniques was observed when applied to real data.

In response to advances in tree-based (i.e., CART) techniques, including boosting (Freund & Shapire, 1996), bagging (Breiman, 1996), and hybrid approaches, Moisen et al. (2003) conducted another comparison as part of prototype mapping activities conducted under the LANDFIRE project. The research has particular relevance in the context of the current mapping project because 1) MZ16 served as the prototype mapping zone for the study and 2) FIA-derived BAWHT served as one of two continuous response variables tested. The BAWHT investigation involved a comparison of three separate modeling techniques including 1) a simple CART model, 2) a MARS model, and 3) a hybrid model combining tree-based (i.e., CART) methods with recursive linear regression implemented as part of the Cubist (www.rulequest.com) package. Based on an evaluation of global performance measures and residual plots, the MARS and Cubist models were judged to have performed similarly (i.e., average errors equal to 2.95 and 2.81 m, respectively) while both outperformed the CART model.

The overall performance and user-friendly characteristics of tree-based modeling strategies have led to the formal adoption of the Cubist approach by both the NLCD and LANDFIRE mapping projects (Homer et al., 2004). Based on insights provided by Moisen et al. (2003) as well as the widespread acceptance of tree-based methods by the mapping applications community, it was determined that regression trees would also

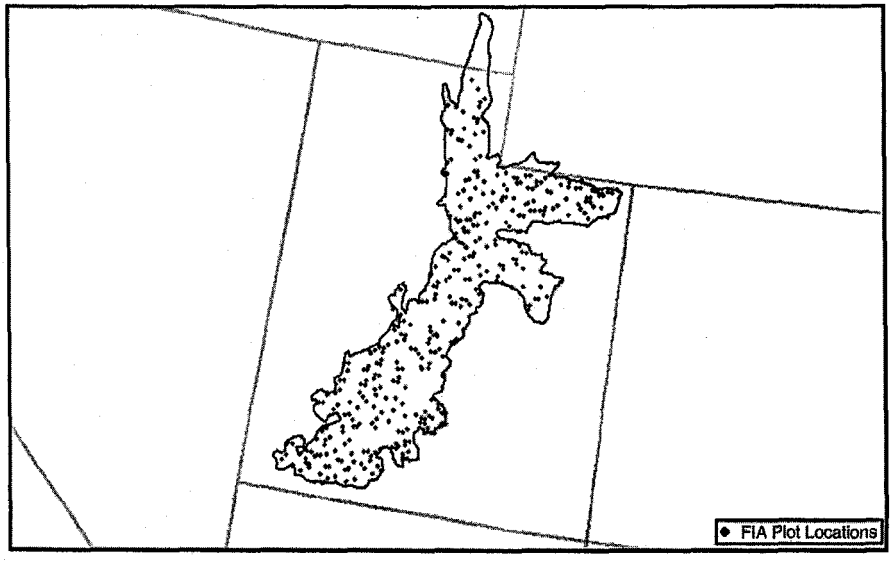

Fig. 2. Approximate locations of 393 FIA field plots within MZ16 (central Utah). Data from these plots were retained as part of the final model development database (MDDB). Figure provided courtesy of FIA National Spatial Data Services.

provide an appropriate modeling framework for the current study. Two tree-based approaches, Cubist and randomForest (Breiman, 2001), were subsequently selected for comparative testing.

4.2. Tree-based regression

4.2.1. General background

Popularized by Breiman et al. (1984), regression trees have evolved considerably in recent years as algorithm development has focused on overcoming known deficiencies and enhancing predictive power (Moisen et al., 2003). Important advancements include the development of bagging and boosting techniques as well as the maturation of hybrid tree-based methods. Bagging (Breiman, 1996) and boosting (Freund & Shapire, 1996) fall into the category of ensemble learning methods where the goal is to construct a "forest" (i.e., ensemble) of expert trees and combine them through a voting scheme (i.e., simple averaging) for the purpose of improving predictive accuracy (Bauer & Kohavi, 1999; Dietterich, 2000). Bagging, a term derived from *b*ootstrap *agg*regation, produces replicate training sets, and hence trees, by sampling with replacement from the training cases. Boosting uses all training cases to construct each tree, but successive trees place extra weight on cases that proved difficult to predict in earlier trees. In bagging, all members of the ensemble have an equal vote and a simple average is used to compute final predictions (Chan et al., 2001; Quinlan, 1996). Conversely, in boosting, different voting strengths are assigned to ensemble members based on their accuracy. The Cubist and randomForest packages are perhaps the most well known examples of regression-tree approaches with ensemble learning enhancements.

Hybrid tree-based methods attempt to exploit the strengths of both standard regression-tree algorithms and local modeling techniques to enhance the predictive abilities of final models (Moisen et al., 2003). A number of hybridization strategies have been proposed in the literature, and the reader is referred to Torgo (1999) for a detailed summary.

4.2.2. Cubist

The Cubist package (www.rulequest.com) implements a hybrid tree-based approach that combines a regression-tree algorithm with local modeling using a proprietary variant of linear least squares regression (R. Quinlan, *pers. comm.*). Options available in Cubist include *composite* and *committee* models. Whereas composite models combine regression trees with instance-based or nearest-neighbor models (Quinlan, 1993), committee models provide ensemble learning capabilities similar to that of boosting. Additionally, Cubist provides for *f-fold cross validation*, which is a common method for obtaining more reliable estimates of predictive accuracy, particularly when working with datasets that are not large enough to support separate training and testing populations.

4.2.3. randomForest

First proposed by Breiman (2001), the concept of random forests (RF) adds an additional layer of randomness to the bagging strategy described above (Breiman, 1996; Liaw & Wiener, 2002). In addition to constructing each standard regression tree in the ensemble with a different bootstrapped sample of the data, the RF algorithm incorporates a unique approach to node splitting. Whereas node splitting is typically accomplished using the best (i.e., optimal) split among all predictors, RF node splitting is achieved using the best split among a random subset of predictors chosen at each node (Breiman, 2001; Liaw & Wiener, 2002). The RF bootstrap samples used to construct each tree omit approximately 1/3rd of the cases. These hold-out cases are referred to as out-of-bag (OOB) (Breiman, 2001, 2006).

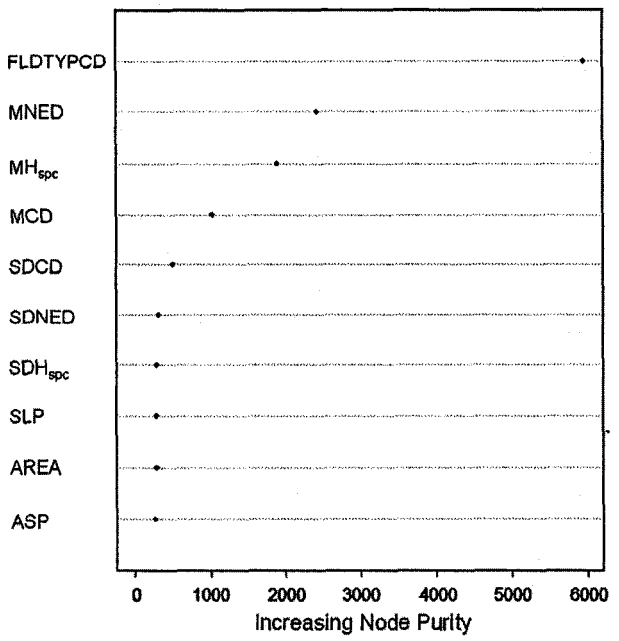

Fig. 3. Variable importance plot generated in randomForest indicating the relative importance of the first four-predictor variables over that of the last six in the prediction of basal-area weighted average height (BAWHT).

The randomForest package provides an R interface to the original Fortran programs written by Breiman and Cutler (available at www.stat.berkeley.edu/ users/breiman) (Liaw & Wiener, 2002). R is a programming environment for statistical computing and graphics (R Core Development Team, 2005; www.r-project.org), and is available as Free Software under the terms of the Free Software Foundation's GNU General Public License. The R implementation offers several options for fine tuning and analyzing the RF model as well as a number of instructive text and graphical outputs. Among these, perhaps the most useful is the *variable importance plot* (VIP). To produce a VIP plot, the RF algorithm estimates the importance of each predictor by computing how much the error increases for a given tree when OOB data for each predictor are randomly permuted while all other predictors are left unchanged (Liaw & Wiener, 2002).

5. Implementation

One of the principal goals of the regression-tree analysis was to test and evaluate the Cubist and randomForest packages as part of an operational framework for broad-scale map

generation. This meant not only an assessment based on predictive accuracy, but also one that considered practical issues such as accessibility, flexibility, and interoperability. Although Cubist and randomForest share much of the same regression-tree functionality, they differ somewhat in terms of the specific parameter options offered and the terminology used to describe them. Nevertheless, a one-to-one parameterization of the models was effectively achieved.

The first step in model parameterization focused on predictor-variable selection. Toward this end, three randomForest VIP plots were generated, i.e., one for each response variable, to determine which of the ten predictor variables warranted inclusion in the modeling comparison. The plots were produced using all cases in the MDDB (i.e., 393) and the default randomForest settings which consist of a 500-tree ensemble, a random sample of three predictors per node, and a minimum node size of five cases. Fig. 3 includes the VIP plot produced for BAWHT, which is highly representative of the other height metrics (i.e., AVGHT and MAXHT). Inspection of the plot indicated that forest type code (FLDTYPCD), mean elevation (MNED), mean scattering phase center height (MH_{spc}), and mean canopy density (MCD) possessed considerable explanatory value in the regression-tree model. The remaining variables had importance values that were very similar to one another and not particularly high when compared to the other four. Taking into consideration both variable economy as well as an intuitive sense of agreement with the VIP plot, the decision was made to proceed with the Cubist–randomForest comparison using only the first four predictors in Fig. 3. It is important to note that FLDTYPCD was retained as a predictor variable in spite of being derived from the FIADB and not a continuous raster data layer. This decision was made because FLDTYPCD is a ground-observed variable and, therefore, represents the most accurate source of forest cover type information available from the perspective of model development. The replacement of FLDTYPCD with a spatially-continuous surrogate layer in the context of model prediction and map generation is addressed in Section 6.2.

For the purposes of the primary Cubist–randomForest comparison, all possible combinations (15 total) of the four-predictor variables were used to construct 45 different ensemble regression trees predicting each of the three height response variables (i.e., BAWHT, AVGHT, MAXHT). The specific objectives of this comparison were to 1) determine which suite of predictor variables possesses the greatest explanatory value

Table 2
Performance measures for each of 45 Cubist and randomForest tree-based regression models used to predict basal area weighted height (BAWHT)

	Cubist			randomForest		
	Basal Area Weighted Height (BAWHT)			Basal Area Weighted Height (BAWHT)		
	r_{train} $n=294$	r_{test} $n=99$	Average error (m)	r_{train} $n=294$	r_{test} $n=99$	Average error (m)
One variable models						
MH_{spc}	0.55	0.73	3.5	0.88	0.55	3.9
MNED	0.70	0.64	3.2	0.91	0.51	3.9
MCD	0.49	0.54	4.1	0.88	0.42	4.4
FLDTYPCD	0.88	0.87	2.0	0.89	0.87	2.2
Two variable models						
MH_{spc}, MNED	0.83	0.75	2.6	0.95	0.75	2.6
MH_{spc}, MCD	0.61	0.78	3.0	0.93	0.73	3.1
MH_{spc}, FLDTYPCD	0.91	0.90	1.9	0.95	0.89	2.0
MNED, MCD	0.78	0.74	2.7	0.95	0.74	2.7
MNED, FLDTYPCD	0.89	0.87	2.1	0.93	0.86	2.2
MCD, FLDTYPCD	0.88	0.88	2.0	0.94	0.87	2.0
Three variable models						
MH_{spc}, MNED, MCD	0.83	0.79	2.5	0.96	0.82	2.4
MH_{spc}, MNED, FLDTYPCD	0.91	0.90	1.9	0.96	0.89	1.9
MH_{spc}, MCD, FLDTYPCD	0.91	0.89	1.9	0.97	0.89	2.0
MNED, MCD, FLDTYPCD	0.88	0.88	2.0	0.95	0.88	2.0
Four variable models						
MH_{spc}, MNED, MCD, FLDTYPCD	0.91	0.90	1.8	0.97	0.89	1.9

The performance of all possible combinations of the four-predictor variables (i.e., MH_{spc}, MNED, MCD, and FLDTYPCD) is reported.

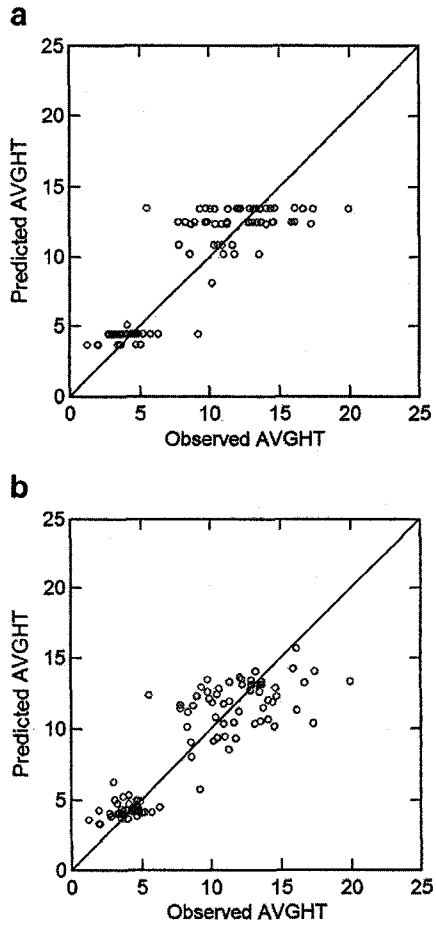

Fig. 4. Plots of observed versus predicted AVGHT (m) derived from Cubist tree-based models in which a) a single predictor variable (i.e., FLDTYPCD) and b) all four-predictor variables (i.e., MH_{spc}, MNED, MCD, and FLDTYPCD) were used in model development.

Table 3
Performance measures for six cross-validated randomForest tree-based regression models used to predict each of three height response variables (i.e., BAWHT, AVGHT, MAXHT)

Cross-validation results	Basal Area Weighted Height (BAWHT)		Average Height (AVGHT)		Maximum Height (MAXHT)	
	$r_{cross\ validation}$ (n=393)	Average error (m)	$r_{cross\ validation}$ (n=393)	Average error (m)	$r_{cross\ validation}$ (n=393)	Average error (m)
Cubist						
MH_{spc}, MNED, MCD, FLDTYPCD	0.89	1.9	0.90	1.5	0.87	3.0
randomForest						
MH_{spc}, MNED, MCD, FLDTYPCD	0.89	2.0	0.90	1.5	0.87	3.1

Models were constructed using all four-predictor variables (i.e., MH_{spc}, MNED, MCD, and FLDTYPCD).

in terms of canopy height prediction, 2) determine which of the three response variables is predicted most accurately, and 3) determine how accurately canopy height can be estimated, i.e., what level of error can be expected. The comparison was conducted using randomly selected independent training and testing datasets of 294 and 99 cases, respectively (Fig. 1). In both Cubist and randomForest, 500-tree ensembles (called *committees* in Cubist and *forests* in randomForest) were grown using the training dataset, and models constructed from the training dataset were then validated against (i.e., used to predict) the unseen cases in the testing dataset.

The performance of both the Cubist and randomForest models was evaluated using the correlation coefficient (r) and the average error (AE). The correlation coefficient was reported for both the training and testing datasets as a measure of linear agreement between the observed and predicted values. The average error was based on data from the testing dataset only and represents the average of the absolute differences between observed and predicted values. To simulate the bootstrapping procedure implemented in randomForest, all Cubist regression-tree models were constructed using a 3-fold cross validation.

A secondary Cubist–randomForest comparison was conducted wherein all (i.e., 393) cases and all (i.e., 4) predictor variables were used to construct three different ensemble regression trees — one for each response variable. The primary objective of this analysis was to compare the accuracy of selected Cubist–randomForest models developed under a cross-validation strategy to that of models developed under the independent validation strategy described above.

6. Results and discussion

6.1. Model development and validation

The results of the primary Cubist–randomForest comparison are presented in Table 2. The table includes model performance measures for 15 different tree ensembles representing all possible predictor-variable combinations corresponding to the response variable basal area weighted height (BAWHT). The results from the average height (AVGHT) and maximum height (MAXHT) variables were not included as they did not differ significantly from those of BAWHT. Based on an evaluation of these results, a number of observations can be made. First, inspection of the average errors suggests that predictive accuracy does not differ significantly between the Cubist and randomForest models. This generally holds true regardless of the predictor-variable combination.

Second, predictive accuracy generally improves with the number of predictor variables in the model. Although the presence of such a trend is not surprising from a statistical point of view, it is consistent with what is expected given the variable importance information presented in Fig. 3. The existence of this trend appears to be package-independent; however, it is noticeably stronger in the randomForest case, i.e., the average error is almost

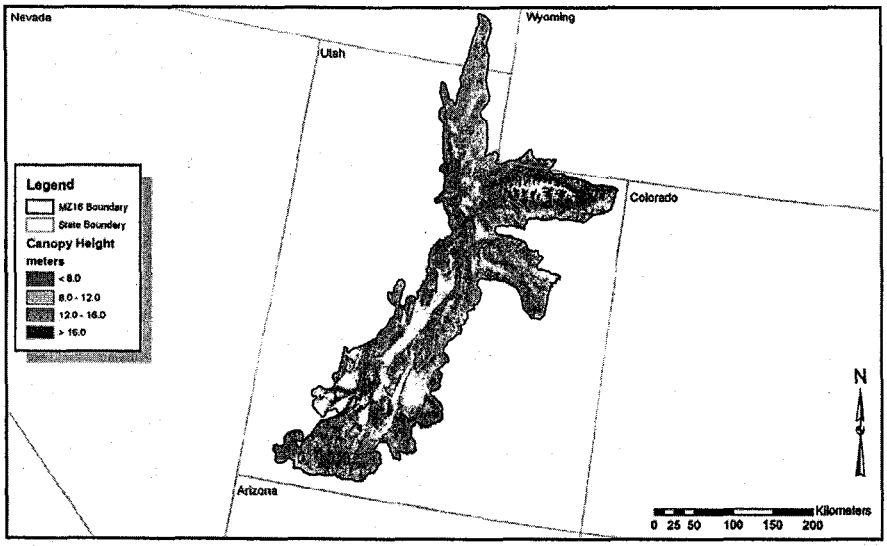

Fig. 5. Map of basal-area weighted average height (BAWHT) (MZ16 — central Utah). Legend reflects BAWHT in meters.

Table 4
Tabular cross-walk for converting between FIA forest type and LANDFIRE existing vegetation type classes

No. of plots	FIA Forest Cover Type (FLDTYPCD)		LANDFIRE Existing Vegetation Type (EVT)
125	Pinyon–Juniper — 180	⇦	Colorado Plateau Pinyon–Juniper — 2016
		⇦	Great Basin Pinyon–Juniper — 2019
		⇦	Rocky Mountain Foothill Limber pine–Juniper — 2049
		⇦	Inter-Mountain Basins Juniper — 2115
22	Douglas-fir — 201	⇦	Rocky Mountain Montane Dry-Mesic Mixed Conifer — 2051
		⇦	Rocky Mountain Montane Mesic Mixed Conifer — 2052
15	Ponderosa Pine — 221	⇦	Southern Rocky Mountain Ponderosa Pine —2054
		⇦	Rocky Mountain Ponderosa Pine — 2117
46	Engelmann Spruce/ Subalpine Fir — 260	⇦	Rocky Mountain Subalpine Dry-Mesic Spruce-Fir — 2055
		⇦	Rocky Mountain Subalpine Mesic Spruce-Fir — 2056
18	White Fir — 261	⇦	White Fir — 2208
28	Lodgepole Pine — 281	⇦	Rocky Mountain Lodgepole Pine — 2050
4	Foxtail/Limber/ Bristlecone Pine — 360	⇦	Inter-Mountain Basins Subalpine Limber-Bristlecone Pine — 2020
		⇦	Rocky Mountain Subalpine Montane Limber-Bristlecone Pine — 2057
61	Aspen — 901	⇦	Rocky Mountain Aspen — 2011
		⇦	Inter-Mountain Basins Aspen-Mixed Conifer — 2061
55	Deciduous Oak — 925	⇦	Shrub Live Oak — 2215
		⇦	Gambel Oak — 2217
13	Mountain Mahogany — 953	⇦	Inter-Mountain Basins Mountain Mahogany — 2062
6	Intermountain Maple — 954	⇦	Rocky Mountain Bigtooth Maple — 2012
393			

always minimized when all four-predictor variables are included. Nevertheless, it is important to note that the difference in error reduction observed following the inclusion of additional, i.e., third or fourth, variables is not always significant. For example,

Table 5
Summary of BAWHT map accuracy based on performance measures computed from three different sets of FIA field plots and associated map values

Polygon size range	No. of plots	Percent of total plots	Correlation coefficient (r)	Average error (m)
<1.5 ha (17 pixels)	133	25	0.81	2.6
≥1.5 ha (17 pixels)	375	75	0.90	1.9
All Polygons	508	100	0.88	2.1

BAWHT can be estimated by randomForest using either three (e.g., MH_{spc}, MNED, and FLDTYPCD) or four predictors (e.g., MH_{spc}, MNED, MCD, and FLDTYPCD) with an average error of 1.9 m (r_{test}=0.89) in both cases (Table 2). In this example, the addition of MCD appears to have negligible explanatory value for the prediction of BAWHT. This particular result is not unexpected given that MCD exhibits the lowest importance value of the four-predictor variables (Fig. 3).

The previous example underscores the need to consider not only performance measures but also variable importance when evaluating tree-based models. At the same time, it is perhaps equally important to consider the distributional characteristics of the predictions, and the implications these characteristics might have for the final map product. For example, AVGHT can be estimated by randomForest with an average error of 1.6 m (r_{test}=0.89) using all four predictors. Alternatively, the same estimation can be made with nearly the same level of performance (i.e., 1.6 m average error; r_{test}=0.88) using only a single variable, forest cover type (i.e., FLDTYPCD), as a predictor. Statistically speaking, a model based on FLDTYPCD alone would seem to afford both model simplicity and predictive power. However, as a categorical predictor, FLDTYPCD lacks the necessary variance to produce map products with acceptable local accuracy. This lack of variance is manifested clearly in Fig. 4a, wherein a plot of actual versus predicted AVGHT is observed to have a pronounced horizontal banding pattern. In contrast, Fig. 4b illustrates the relatively even distribution that results when all four predictors are included in the model. The consequence of basing model development on a single categorical predictor is that the hybrid randomForest (or Cubist) algorithm is effectively reduced to a

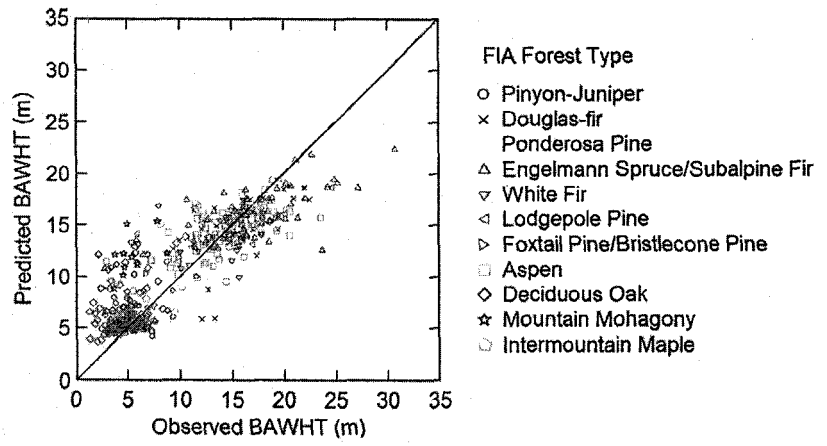

Fig. 6. Plot of observed versus predicted BAWHT based on map extractions from beneath 508 FIA sample locations.

Table 6
Summary of BAWHT map accuracy based on performance measures computed for each of the 11 FIA forest types observed within MZ16

No. of plots (n = 508)	FIA Forest Cover Type (FLDTYPCD)	Observed average BAWHT (m)	Observed S.D. BAWHT (m)	Predicted average BAWHT (m)	Predicted S.D. BAWHT (m)	Average error (m)	Error as a percentage of BAWHT (%)
67	Deciduous Oak – 925	5.0	2.8	6.9	2.73	2.4	48.8
18	Mountain Mahogany – 953	5.0	1.4	9.2	3.62	4.2	83.1
170	Pinyon-Juniper – 180	5.1	1.3	6.0	1.90	1.3	26.1
9	Intermountain Maple – 954	7.0	2.0	7.9	2.95	2.2	31.2
2	Foxtail/Limber/Bristlecone Pine – 360	8.1	1.8	9.5	1.22	2.2	26.8
18	Ponderosa Pine – 221	14.0	4.2	12.8	3.9	2.3	16.7
80	Aspen – 901	14.5	3.4	14.4	2.1	2.1	14.7
21	White Fir – 261	14.6	2.5	14.2	2.2	1.8	12.4
32	Lodgepole Pine – 281	14.8	3.3	14.9	1.7	1.7	11.6
29	Douglas-fir – 201	15.7	3.8	14.8	3.4	2.8	17.5
62	Engelmann Spruce/Subalpine Fir – 260	17.0	4.8	16.1	2.4	2.8	16.2

Forest types are presented in order of increasing observed average BAWHT (see third column). Dashed line separates short-statured types (above) from all others (below).

standard regression tree. That is to say, AVGHT is predicted not with a series of local regression models, but rather with a series of local constants, i.e., mean values (Fig. 4a). As a result, although the statistical performance of FLDTYPCD is noteworthy (Table 2), this variable does not demonstrate sufficient predictive ability in a mapping context.

6.2. Map generation and accuracy assessment

An effort was undertaken to determine how the tree-based modeling framework might be most efficiently incorporated into a streamlined workflow for model-based prediction and map generation (Fig. 1). As part of this effort, a second Cubist–randomForest comparison was conducted wherein all cases (i.e., 393) and all predictor variables (i.e., 4) were used to construct three different ensemble regression trees, one for each response variable (Table 3). The primary objective of this analysis was to compare the accuracy of selected Cubist–randomForest models developed under a cross-validation strategy to that of models developed under the independent validation strategy summarized in the previous section. The analysis was constrained to include only models constructed from all four predictors because, among the 15 different predictor-variable combinations evaluated, the four-predictor models are consistently among the most accurate (Table 3).

Generally speaking, model performance is very consistent with that reported under the independent training and testing scenario (Table 2). For example, Cubist produces average errors of 1.9, 1.5, and 3.0 m for BAWHT, AVGHT, and MAXHT, respectively, under the cross-validation strategy. Conversely, average errors of 1.8, 1.5, and 3.0 m are achieved for BAWHT, AVGHT, and MAXHT, respectively, under the independent validation strategy (Table 2). In addition, the observed correlation coefficients are very consistent among the models. Overall, these results suggest that reliable estimates of predictive accuracy are indeed obtainable through cross-validation as implemented in both Cubist and randomForest. This is particularly advantageous when the number of available cases is insufficient to form suitable training and testing populations.

The ultimate goal of this proof-of-concept study was to produce a high-resolution map of vegetation canopy height for MZ16. The model results presented in Tables 2 and 3 suggest that a number of different height maps could legitimately be produced, each having different strengths and weaknesses depending on the application. For example, a map of average height (AVGHT) might be produced as a general purpose layer to inform analyses related to fire modeling, habitat management, or biodiversity conservation. Alternatively, a map of maximum height (MAXHT) might be useful for identifying tracts of old growth forest or forested areas prone to windthrow. In the context of the current research, which is motivated by the need for accurate baseline estimates of aboveground biomass and carbon stocks, it is hypothesized that a map of basal-area weighted average height (BAWHT) will be most useful. That is to say, BAWHT is likely to be a more robust predictor of aboveground biomass and carbon stocks than other height metrics because it takes into consideration not only stem height, but stem diameter as well.

A series of computer programs was implemented in R to generate a raster map of BAWHT for MZ16 at a resolution of 30 m (Fig. 5). The map was produced using the four-predictor,

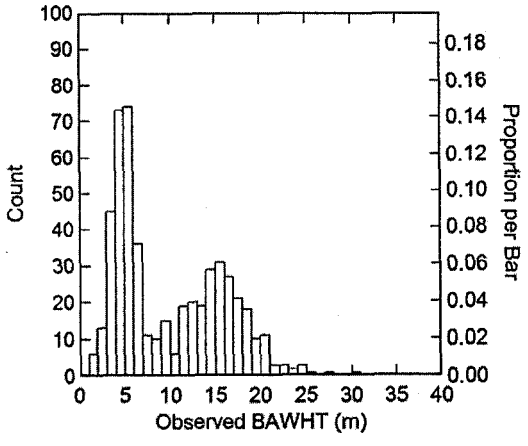

Fig. 7. Histogram of observed BAWHT within MZ16 illustrating the bimodal distribution of height values.

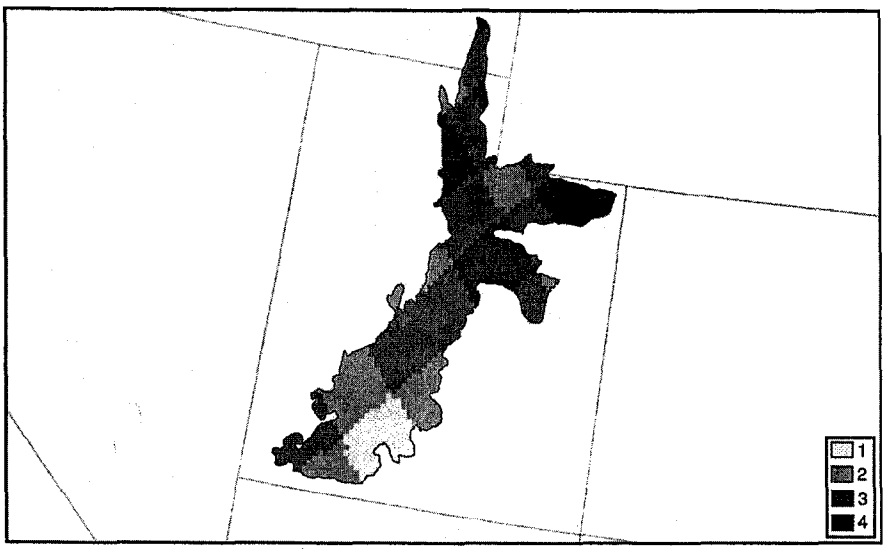

Fig. 8. Map illustrating the location and extent of different SRTM datatake regimes within MZ16 (central Utah). Legend reflects the number of SRTM datatakes, which generally increases with distance north.

cross-validated randomForest model reported in Table 3. Although the Cubist model was observed to have a slight edge in terms of overall predictive performance (Table 3), the randomForest framework was ultimately selected as the source of tree-based models for map generation. This decision was motivated in large measure by the greater accessibility and flexibility offered by randomForest within the open-source R environment, including the ability to implement parallel processing on Linux-based computing clusters. The map was constructed from pixel-by-pixel predictions based on inputs from the segment-based mean values associated with the MH_{spc} predictor variable together with data from the original NED, CD, and FLDTYPCD raster layers. The segment-based mean values associated with the MNED and MCD variables were not used in map generation because noise reduction was not an issue where these layers were concerned and given that the pixel-to-pixel variation present in the original raster layers represented a potentially important source of explanatory information. As a surrogate for the FIADB-derived FLDTYPCD variable, the LANDFIRE EVT layer (see Section 3.1.3) was used as a spatially explicit predictor variable. This was accomplished by developing a cross-walk between FIA forest type and LANDFIRE existing vegetation type (Table 4).

A plot of observed versus predicted BAWHT is shown in Fig. 6. This plot illustrates the predictive performance of the map itself and was constructed using predicted BAWHT values extracted from the map following generation of the raster layer (Fig. 5). Map extractions were performed by FIA NSDS; predicted values were extracted from individual 30-m pixels associated with the centers of 508 FIA field plots. Of these plots, 375 were part of the final MDDB[2]. The remaining 133 plots were originally excluded from model development because they occurred within image–object polygons that were deemed too small (i.e.,<1.5 ha/17 pixels) to provide for adequate noise reduction (see Section 3.3). Performance measures for each of the aforementioned plot subsets (133 and 375), as well as for the total set (508), are reported in Table 5. The 133-plot subset produced an average error of 2.6 m (r=0.81; Table 5). This subset holds perhaps the greatest value for the evaluation of map accuracy because it provides not only for an independent validation of the BAWHT map (i.e., due to exclusion from the MDDB), but also for a worst-case estimate of average map error (i.e., due to insufficient averaging). As expected, the average error of 1.9 m (r=0.90) attributed to the 375-plot subset is 1) markedly lower (by approximately 0.7 m) than that of the subset based on 133 plots and 2) generally consistent with that reported in Table 3 for the 293-plot, cross-validated model result (average error=2.0 m). It should be noted that this result is likely biased to a small degree by the fact that 18 of the original 393 MDDB plots were not included in the 375-plot subset (for further details, see Footnote 2). In summary, when the accuracy of the map is evaluated as a whole, i.e., when the two plot subsets are combined, an average error of 2.1 m (r=0.88) is observed (Table 5).

It is also possible to evaluate the accuracy of the BAWHT map in terms of how well individual FIA forest type classes and groups of classes are predicted (Table 6). A cursory inspection of Fig. 6 reveals an obvious clustering of data values, which is an artifact of the strongly bimodal height distribution observed within MZ16 (Fig. 7). Over half of the 508 FIA field plots (i.e., 266; 52%) are associated with forested types (e.g., woodlands) of inherently short stature (i.e., average BAWHT ≤ 8.5 m), including Deciduous Oak, Mountain Mahogany, Pinyon–Juniper, Intermountain Maple, and Foxtail/Limber/Bristlecone Pine (Table 6). The remaining plots (i.e., 242; 48%) comprise the typical mid- to high-elevation forest types of the Intermountain West including Ponderosa Pine, Lodgepole

[2] The final MDDB included a total of 393 plots. Of these, 18 plots (i.e., 393-375) were not associated with map predictions due to misclassification errors (e.g., non-forest classes assigned to forested pixels) in the LANDFIRE EVT layer. As a result, these plots (i.e., 5% of the MDDB) are not represented in Fig. 6.

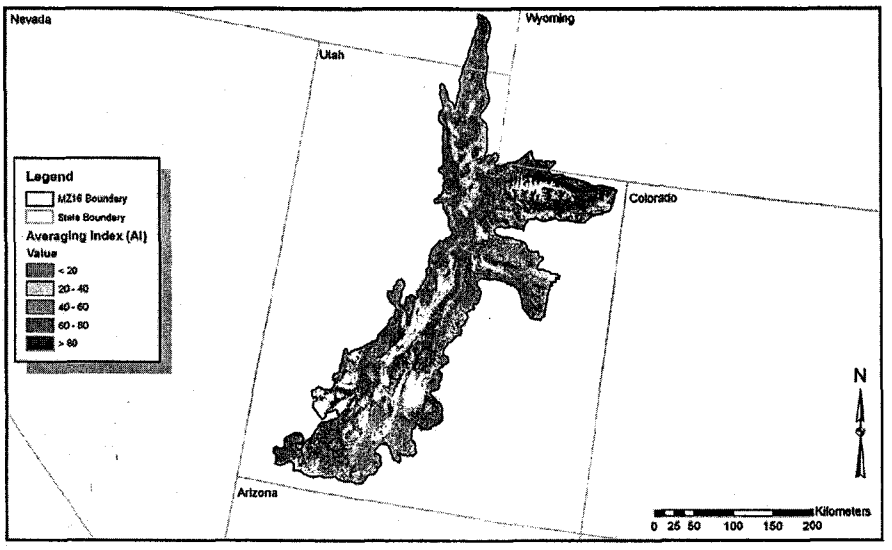

Fig. 9. Averaging index (AI) map (MZ16 — central Utah). Legend reflects the number of SRTM C-band datatakes multiplied by image–object polygon size (in pixels). Whereas the number of acquired datatakes ranges from 1 to 4 (Fig. 8), polygon size ranges from 1 to 551 pixels. Thus, the map assumes values between 1 (i.e., one datatake multiplied by a polygon consisting of one pixel) and 2204, with the observed and theoretical maximums being equal.

Pine, Douglas-fir, and Trembling Aspen among others. Within this group, the average BAWHT is greater than 13.5 m (Table 6).

Also included in Table 6 are accuracy statistics associated with each of the 11 FIA forest type classes for which BAWHT values were predicted and mapped. In general, a reasonably close correspondence is revealed between the observed and predicted values reported for both the average and standard deviation of BAWHT. Error statistics are reported both as an average and as a percentage of average BAWHT. Not surprisingly, error constitutes the greatest percentage of average BAWHT among the five short-statured types (Table 6). Although the average error among these classes is less than 3.0 m with the exception of the Mountain Mahogany class (i.e., 4.2 m), the error percentage of average BAWHT remains above 25%. Conversely, among the remaining classes, where average error is also less than 3.0 m, the error percentage of average BAWHT is consistently below 20%. It is important to point out that this trend is exacerbated to some degree by the variance attenuation associated with the BAWHT predictions (Fig. 6; see Cohen et al., 2003). In particular, observed BAWHT values less than the mean (i.e., short-statured trees) tend to be somewhat overpredicted, resulting in markedly higher error percentages.

6.3. Map confidence estimation

The presence of random phase noise in the SRTM DEMs has been shown to be a potentially significant source of error in SRTM-derived estimates of canopy height (Kellndorfer et al., 2004; Walker et al., 2007). It has also been shown that sample averaging, accomplished using multiple SRTM datatakes and/or aggregations of SRTM pixels, represents a viable strategy for reducing phase noise error. In the context of the current

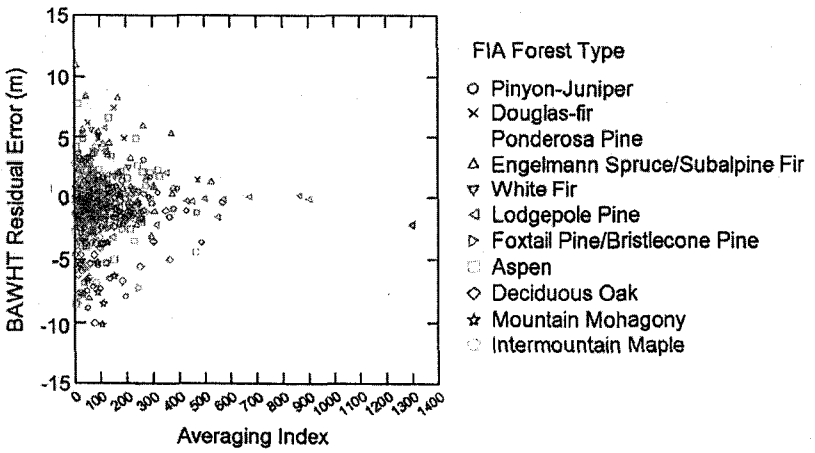

Fig. 10. Plot illustrating the relationship between the averaging index (AI) and BAWHT residual error. The range of the residuals, reflected in the vertical spread of points, is observed to decrease with increasing values of AI.

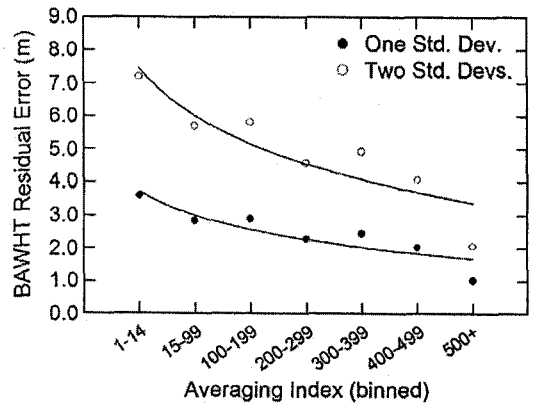

Fig. 11. Curves illustrating the level of error expected for 68% (one standard deviation) and 95% (two standard deviations) of the BAWHT map.

research, a segmentation-based approach to phase noise reduction was implemented and subsequent model development was carried out using only data associated with polygons greater than or equal to 1.5 ha (≈ 17 pixels) in size. In addition to reducing phase noise error in the SRTMDIFF layer, it is reasonable to assume that this averaging strategy contributed to error reduction in both the CD and NED layers as well.

Given the prominent role played by sample averaging in the generation of the BAWHT map, and in the interest of exploring further the results reported in Table 6, an effort was undertaken to investigate in detail the relationship between sample averaging and map accuracy. It was hypothesized that this relationship should manifest itself as a functional trend in which the magnitude of residual errors decreases as the level of averaging increases. Support for this hypothesis would suggest that more confidence be placed in BAWHT height predictions, i.e., individual map pixels, associated with larger polygons and vice-versa.

To test the hypothesis, it was necessary to quantify the amount of sample averaging applied to each 30-m pixel in the BAWHT map. A simple averaging index (AI) was developed in which the index value computed for each pixel reflected the number of SRTM C-band datatakes associated with the pixel (Fig. 8) multiplied by the size (measured in pixels) of the particular image–object polygon to which the pixel belonged. An AI map of MZ16 is shown in Fig. 9. From this map, individual pixel values were extracted from beneath the center of each of the 508 FIA field plots used in the previous accuracy assessment. Whereas the AI map values range from 1 to 2204, values in the extracted subset ranged from 2 to 1296.

A plot of AI versus residual error is presented in Fig. 10. Inspection of the plot reveals a reasonably clear relationship between the two variables. As hypothesized, the magnitude of the residual values, as evidenced by their vertical spread, is observed to decrease as the values of the AI decrease. Based on the information contained in Fig. 10, it was possible to generate a pair of confidence curves relating predictive accuracy, as manifested in the residual values, to the degree of averaging applied, as reflected in the AI (Fig. 11). The curves were generated from the standard deviations of residual values occurring within seven separate AI bins. The primary curve (closed circles) represents one standard deviation from the mean residual value in each bin. Assuming the sample size in each bin is large enough to be representative of the map as a whole, then this curve defines the maximum error in each bin that can be expected for 68% of map pixels. Similarly, the secondary curve (open circles), representing two standard deviations from the mean, defines the maximum level of error in each bin that can be expected for 95% of map pixels. Overall, the curves reveal that 68% of pixels are expected to have errors less than ±3.7 m (Fig. 11).

A more complete picture of error dynamics within the BAWHT map is provided by Table 7, which illustrates the distribution of AI values among the seven bins used to create the confidence curves referenced in Fig. 11. For example, the table reveals that roughly 87% of pixels in the BAWHT map have AI values of 15 or greater. Thus, it can be said with nearly 70% confidence that 87% of the map area will have errors less than ± 3.0 m.

7. General discussion

The proof-of-concept study presented here is similar, in many respects, to work recently conducted by Moisen et al. (2003) in the context of LANDFIRE prototype mapping activities. In general, both efforts focus on the mapping of BAWHT in MZ16 using hybrid tree-based regression techniques. However, the specifics of the two approaches differ to a large degree. The principal difference involves the use by Moisen et al. (2003) of multiple predictor variables derived from passive optical data sources including multi-date, multi-band Landsat data and Landsat-derived tasseled cap indices. Additionally, Moisen et al. (2003) based model development activities on a larger complement of FIA field plots (2052) acquired over a broader time frame (1993–2001). Given these and other methodological differences between the two efforts, it is not possible to make any firm judgments about the merits of either approach based on the information available. Nevertheless, a number of inferences can be made in the context of a performance-based comparison. In terms of overall accuracy, the InSAR-optical fusion approach reported on here compares quite favorably with that of Moisen et al. (2003). Whereas, the Cubist-based committee model[3] implemented by Moisen et al. (2003) produced an average error of 2.81 m (9.23 ft) and a correlation coefficient of 0.75, virtually the same model implemented here resulted in an average error of 2.06 m (6.76 ft) and a correlation coefficient of 0.88. Moisen et al. (2003) modeled BAWHT as a function of 40 predictor variables; only four were considered in the context of the current modeling effort. In general, it is impossible to ascribe the observed improvement in predictive accuracy (0.75 m) afforded by the current approach to any particular element of either approach. Nevertheless, the strategy presented here possesses

[3] The randomForest package was not included in the comparative analysis of modeling approaches conducted by Moisen et al. (2003; see Section 4.2.1 for further details). As a result, it was not possible to present here the results of a randomForest comparison.

Table 7
Distribution of averaging index (AI) values within the mapped region of MZ16

AI values	Distribution (%)
1	0.2
2–9	7.2
10–14	6.0
15–99	47.1
100–199	23.0
200–299	8.9
300–399	3.8
400–499	1.8
500+	2.0
	100.0

characteristics that make it uniquely appealing. Principal among these is the inclusion of a continuous predictor variable with demonstrated sensitivity to forest structure in the vertical dimension. Regardless of ecoregional-specific influences (e.g., topography), the SRTM-based predictor is expected to remain a significant explanatory contributor. Perhaps equally appealing is the level of variable economy afforded by the approach. Because model development was predicated on a physical understanding of the predictors and their role in explaining vegetation canopy height, few variables were needed to achieve a relatively high level of accuracy. Given the overall performance and intuitive appeal of the approach, the potential exists to contribute in a meaningful way to the improvement of the LANDFIRE canopy height product.

While the overall accuracy of the BAWHT map is quite encouraging, important caveats associated with various aspects of the mapping approach warrant acknowledgement. First, although the FIA program is currently the only nationwide source of consistent and reliable forest inventory data, the FIADB was never intended for use as a reference database for remote sensing applications. Generally speaking, the design, size, and spatial frequency of FIA field plots is not well suited to broad-scale, high-resolution mapping of forest structural attributes. Nevertheless, this proof-of-concept study provides yet further evidence to support the use of the FIADB in regional- to continental-scale mapping efforts. Simply put, projects of this sort would be wholly impossible without the long-standing commitment of the U.S. Congress and USDA Forest Service to support and maintain the FIA network. Given the certainty of advancements in high-resolution, stand-level remote sensing technologies such as lidar, the FIA network will likely need to evolve and keep pace with such advancements if it is going to remain relevant to the needs of an ever-expanding and technically-inclined user community.

A second caveat involves the use of derivative map products such as the CD and LANDFIRE EVT layers as predictor variables in model development. Because errors in derivative products are not uncommon, and because they can be difficult to track given the general lack of reliable and/or published accuracy assessments, derivatives are best avoided in projects where unwanted error is likely to accumulate in a final map product. Both the derivative CD and LANDFIRE EVT are well removed from the raw Landsat ETM+ data used in their production. Although the value of both products, and particularly the LANDFIRE EVT, has been demonstrated clearly in the current context, their replacement by either Landsat ETM+ at-satellite reflectance data or tasseled-cap indices is worthy of future testing.

Finally, it is important to acknowledge that the value of the SRTM MH_{spc} variable for canopy height prediction is certain to vary considerably among the 66 mapping zones that comprise the conterminous U.S. For example, in MZ16 and throughout much of the western U.S., pronounced elevational gradients govern precipitation and solar radiation to pattern the distribution and associated vertical structure of vegetation (Barnes et al., 1998). As a result, the relative explanatory power of the SRTM MH_{spc} predictor is less than that of either the FLDTYPCD or MNED variables (Fig. 3). Additionally, it is well known that InSAR DEM errors tend to increase in regions of diverse topography where layover, shadow, and concomitant decorrelation can greatly complicate the process of phase unwrapping (Hanssen, 2001). For these reasons, MZ16 may not provide the most appropriate test for an InSAR-optical fusion approach to canopy height mapping. At the same time, success of the approach in the western U.S. might be viewed as an encouraging indicator of performance throughout much of the East where predictors like FLDTYPCD and MNED correlate much less with canopy height.

8. Conclusions

The proof-of-concept study presented here reveals that production of the first InSAR/optical fusion-based continuous-surface map of vegetation canopy height for the conterminous U.S. is an ambitious goal, but one that is certainly achievable. The innovative segmentation-based approach brings advanced processing and analysis techniques to bear on some of the best available spatial data for the purpose of filling a significant void in our capacity to quantify trends in vertical forest structure across broad spatial scales. The approach provides both the theoretical and operational framework for future work, focused not only on completion of the national map of canopy height, but also on subsequent generation of the first-ever circa-2000 baseline dataset of aboveground biomass and carbon stocks for the conterminous U.S. These layers are being generated as part of the National Biomass and Carbon Dataset 2000 (NBCD 2000), which is scheduled for completion in early 2009. Maps of vegetation canopy height, aboveground dry biomass, and carbon, together with spatial accuracy metrics, will be accessible at 30 m postings via the U.S. Geological Survey Seamless Data Distribution System.

Acknowledgements

The authors would like thank Michelle West, Burton Barnes, Leland Pierce, and Jesse Bishop for their contributions to the preparation of an earlier draft of this manuscript. The development of NBCD 2000 is funded by the NASA Terrestrial Ecology Program under grant number NNG05GI27G. Generation of the canopy height layer is supported by the United States Geological Survey under the LANDFIRE program. PCI Geomatics supports the project through the donation of Geomatica software licenses. Definiens Imaging co-sponsored one license of eCognition.

Collin Homer (NLCD 2001), Dean Gesch (NED), Zhi-Liang Zhu (LANDFIRE), and Birgit Peterson (LANDFIRE) at the USGS EROS Data Center as well as the USDA Forest Service Forest Inventory and Analysis program are all acknowledged for their collaborative contributions.

References

Anderson, J. F., Hardy, E. E., Roach, J. T., & Witmer, R. E. (1976). *A land use and land cover classification system for use with remote sensor data* (pp. 28). U.S. Geological Survey.

Baatz, M., Benz, U., Dehghani, S., Heynen, M., Höltje, A., Hoffmann, P., et al. (2004). *eCognition: User Guide 4: Definiens Imaging GmbH.*

Barnes, B. V., Zak, D. R., Denton, S. R., & Spurr, S. H. (1998). *Forest ecology* (4th ed). New York: John Wiley and Sons.

Bauer, E., & Kohavi, R. (1999). An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine Learning, 36*, 105–139.

Breiman, L. (1996). Bagging predictors. *Machine Learning, 26*, 123–140.

Breiman, L. (2001). Random forests. *Machine Learning, 45*, 5–23.

Breiman, L. (2006). *Manual on setting up, using, and understanding random forests v4.0.* In (p. 33): ftp://ftp.stat.berkeley.edu/pub/users/breiman/Using_random_forests_v4.0.pdf

Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. G. (1984). *Classification and regression trees.* Belmont, CA: Wadsworth International Group.

Brown, C. G. (2003). Tree height estimation using Shuttle Radar Topography Mission and ancillary data. *Electrical Engineering and Computer Science* (pp. 219). Ann Arbor: The University of Michigan.

Brown, C. G., & Sarabandi, K. (2003). Estimation of red pine tree height using Shuttle Radar Topography Mission and ancillary data. *Proceedings of the International Geoscience and Remote Sensing Symposium* (pp. 2850–2852). Toronto, Canada: IEEE.

Brown, C. G., Sarabandi, K., & Pierce, L. E. (2005). Validation of the Shuttle Radar Topography Mission height data. *IEEE Transactions on Geoscience and Remote Sensing, 43*, 1707–1715.

Chan, J. C. -W., Huang, C., & DeFries, R. S. (2001). Enhanced algorithm performance for land cover classification from remotely sensed data using bagging and boosting. *IEEE Transactions on Geoscience and Remote Sensing, 39*, 693–695.

Cohen, W. B., Maiersperger, T. K., Gower, S. T., & Turner, D. P. (2003). An improved strategy for regression of biophysical variables and Landsat ETM+ data. *Remote Sensing of Environment, 84*, 561–571.

Comer, P., Faber-Langendoen, D., Evans, R., Gawler, S., Josse, C., Kittel, G., et. al. (2003). Ecological systems of the United States: A working classification of U.S. terrestrial systems. In NatureServe (Ed.) (p. 75). Arlington, Virginia: NatureServe.

Dietterich, T. G. (2000). An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Machine Learning, 40*, 139–157.

Dobson, M. C., Ulaby, F. T., LeToan, T., Beaudoin, A., Kasischke, E. S., & Christensen, N. (1992). Dependence of radar backscatter on coniferous forest biomass. *IEEE Transactions on Geoscience and Remote Sensing, 30*, 412–415.

Dong, J., Kaufmann, R. K., Myneni, R. B., Tucker, C. J., Kauppi, P. E., Liski, J., et al. (2003). Remote sensing estimates of boreal and temperate forest woody biomass: carbon pools, sources, and sinks. *Remote Sensing of Environment, 84*, 393–410.

Drake, J., Dubayah, R., Knox, R., Clark, D., & Blair, J. B. (2002). Sensitivity of large-footprint lidar to canopy structure and biomass in a neotropical rainforest. *Remote Sensing of Environment, 81*, 378–392.

FIA. (2004). Forest Inventory and Analysis (FIA) National Core Field Guide. Volume 1: Field data collection procedures for Phase 2 plots (Version 2.0). In (p. 208): U.S. Department of Agriculture Forest Service.

Frescino, T. S., Edwards Jr., T.C., & Moisen, G. G. (2001). Modeling spatially explicit forest structural attributes using generalized additive models. *Journal of Vegetation Science, 12*, 15–26.

Freund, Y., & Shapire, R. (1996). Experiments with a new boosting algorithm. *Proceedings of the Thirteenth International Conference on Machine Learning* (pp. 148–156). San Francisco: IEEE.

Friedman, J. H. (1991). Multivariate adaptive regression splines. *Annals of Statistics, 19*, 1–141.

Gesch, D., Oimoen, M., Greenlee, S., Nelson, C., Steuck, M., & Tyler, D. (2002). The National Elevation Dataset. *Photogrammetric Engineering and Remote Sensing, 68*, 5–11.

Guisan, A., Edwards Jr., T. C., & Hastie, T. (2002). Generalized linear and generalized additive models in studies of species distributions: Setting the scene. *Ecological Modeling, 157*, 89–100.

Hanssen, R. (2001). *Radar interferometry: Data interpretation and error analysis.* The Netherlands: Kluwer Academic Publishers.

Hensley, S., Rosen, P., & Gurrola, E. (2000). Topographic map generation for the Shuttle Radar Topography Mission C-band SCANSAR interferometry. *Proceedings of SPIE* (pp. 179–189).

Homer, C., & Gallant, A. (2001). Partitioning the conterminous United States into mapping zones for Landsat TM land cover mapping. In (p. 7). Sioux Falls: Department of Interior, United States Geological Survey.

Homer, C., Huang, C., Yang, L., Wylie, B., & Coan, M. (2004). Development of a 2001 national landcover database for the United States. *Photogrammetric Engineering and Remote Sensing, 70*, 829–840.

Houghton, R. A., & Goodale, C. L. (2004). Effects of land-use change on the carbon balance of terrestrial ecosystems. In R. S. DeFries, G. P. Asner, & R. A. Houghton (Eds.), *Ecosystems and land use change* (pp. 85–98). Washington, D.C.: American Geophysical Union.

Huang, C., Yang, L., Homer, C., Coan, M., Rykhus, R., Zhang, Z., et al. (2002). Synergistic use of FIA plot data and Landsat 7 ETM+ images for large area forest mapping. *Proceedings of the Third Annual Forest Inventory and Analysis Symposium, Gen; Tech. Rep. NC, Vol. 230* (pp. 50–55). St. Paul, MN: U.S. Department of Agriculture, Forest Service, North Central Research Station.

Huang, C., Yang, L., Wylie, B., & Homer, C. (2001). A strategy for estimating tree canopy density using Landsat 7 ETM+ and high resolution images over large areas. *Third International Conference on Geospatial Information in Agriculture and Forestry.*

Hyde, P., Dubayah, R., Peterson, B., Blair, J. B., Hofton, M., Hunsaker, C., et al. (2005). Mapping forest structure for wildlife habitat analysis using waveform lidar: Validation of montane ecosystems. *Remote Sensing of Environment, 96*, 427–437.

Kellndorfer, J. M., Walker, W. S., Pierce, L. E., Dobson, M. C., Fites, J. A., Hunsaker, C., et al. (2004). Vegetation height estimation from Shuttle Radar Topography Mission and National Elevation Datasets. *Remote Sensing of Environment, 93*, 339–358.

Kobayashi, Y., Sarabandi, K., Pierce, L., & Dobson, M. C. (2000). An evaluation of the JPL TOPSAR for extracting tree heights. *IEEE Transactions on Geoscience and Remote Sensing, 38*, 2446–2454.

Lefsky, M. A., Cohen, W. B., Acker, S. A., Parker, G. G., Spies, T. A., & Harding, D. (1999). Lidar remote sensing of the canopy structure and biophysical properties of Douglas-fir western hemlock forests. *Remote Sensing of Environment, 70*, 339–361.

Lefsky, M. A., Cohen, W. B., Parker, G. G., & Harding, D. J. (2002). Lidar remote sensing for ecosystem studies. *Bioscience, 52*, 19–30.

Lefsky, M. A., Harding, D., Cohen, W. B., Parker, G., & Shugart, H. H. (1999). Surface lidar remote sensing of basal area and biomass in deciduous forests of Eastern Maryland, USA. *Remote Sensing of Environment, 67*, 83–98.

Liaw, A., & Wiener, M. (2002). Classification and regression by randomForest. *R News, 2/3*, 18–22.

Linder, S., & Axelsson, B. (1982). Changes in C uptake and allocation patterns as a result of irrigation and fertilization in a young Pinus sylvestris stand. In R. H. Waring (Ed.), *C Uptake and Allocation in Subalpine Ecosystems as a Key to Management. Corvallis, OR: IUFRO Workshop, 2–3 Aug. 1982, Oregon State University.*

McRoberts, R. E., & Liknes, G. C. (2005). Assessing the effects of forest fragmentation using satellite imagery and forest inventory data. *Proceedings of the fourth annual forest inventory and analysis symposium, Gen; Tech. Rep. NC, Vol. 252* (pp. 117–120). St. Paul, MN: U.S. Department of Agriculture, Forest Service, North Central Research Station.

Mette, T., Hajnsek, I., & Papathanassiou, K. (2003). Height–biomass allometry in temperate forests. *IEEE International Geoscience and Remote Sensing Symposium, 3*, 1942–1944.

Mette, T., Papathanassiou, K., Hajnsek, I., Pretzsch, H., & Biber, P. (2004). Applying a common allometric equation to convert forest height from Pol-InSAR data to forest biomass. *IEEE International Geoscience and Remote Sensing Symposium, 1*, 269–272.

Moisen, G. G., & Frescino, T. S. (2002). Comparing five modeling techniques for predicting forest characteristics. *Ecological Modeling, 157*, 209–225.

Moisen, G. G., Frescino, T. S., Huang, C., Vogelmann, J. E., & Zhu, Z. (2003). Predictive modeling of forest cover and tree canopy height in the central Rocky Mountains of Utah. *Proceedings of the 2003 Meeting of the American Society of Photogrammetry and Remote Sensing* (ASPRS) (pp. 11). Anchorage, AK: ASPRS.

Myneni, R. B., Dong, J., Tucker, C. J., Kaufmann, R. K., Kauppi, P. E., Liski, J., et al. (2001). A large carbon sink in the woody biomass of northern forests. *Proceedings of the National Academy of Sciences, 98*, 14784–14789.

Pierce, L., Kellndorfer, J., Walker, W., & Barros, O. (2006). Evaluation of the horizontal resolution of SRTM elevation data. *Photogrammetric Engineering and Remote Sensing, 72*, 1235–1244.

Quinlan, J. R. (1993). *C4.5: Programs for machine learning.* Los Altos, CA: Morgan Kaufmann.

Quinlan, J. R. (1996). Bagging, boosting, and C4.5. *Proceedings of the Thirteenth National Conference on Artificial Intelligence* (pp. 725–730). Cambridge, MA: AAAI Press/MIT Press.

R_Core_Development_Team. (2005). *R: A language and environment for statistical computing.* Vienna: R Foundation for Statistical Computing.

Ranson, K. J., Sun, G., Lang, R. H., Chauhan, N. S., Cacciola, R. J., & Kilic, O. (1997). Mapping of boreal forest biomass from spaceborne synthetic aperture radar. *Journal of Geophysical Research, 102*, 29,599–529,610.

Reese, H., Nilsson, M., Granqvist Pahlén, T., Hagner, O., Joyce, S., Tingelöf, U., et al. (2003). Countrywide estimates of forest variables using satellite data and field data from the National Forest Inventory. *Ambio, 32*, 542–548.

Reese, H., Nilsson, M., Sandström, P., & Olsson, H. (2002). Applications using estimates of forest parameters derived from satellite and forest inventory data. *Computers and Electronics in Agriculture, 37*, 37–55.

Reichle, D. E., Dinger, B. E., Edwards, N. T., Harris, W. F., & Sollins, P. (1973). Carbon flow and storage in a forest ecosystem. In G. M. Woodwell & E.V. Pecan (Eds.), *Carbon and Biosphere* (pp. 345–365). Springfield, VA: U.S. Atomic Energy Commission.

Rosenqvist, A., Milne, A., Lucas, R., Imhoff, M., & Dobson, C. (2003). A review of remote sensing technology in support of the Kyoto Protocol. *Environmental Science and Policy, 6*, 441–455.

Saich, P., Balzter, H., Luckman, A., & Skinner, L. (2001). Deriving forest characteristics using polarimetric InSAR measurements and models. *Proceedings of the International Geoscience and Remote Sensing Symposium* (pp. 3096–3098). Sydney, Australia: IEEE.

Tomppo, E. (1991). Satellite image-based national forest inventory of Finland. *International Archives of Photogrammetry and Remote Sensing, 28*, 419–424.

Tomppo, E., Nilsson, M., Rosengren, M., Aalto, P., & Kennedy, P. (2002). Simultaneous use of Landsat-TM and IRS-1C WiFS data in estimating large area tree stem volume and aboveground biomass. *Remote Sensing of Environment, 82*, 156–171.

Torgo, L. (1999). Inductive learning of tree-based regression models. *Department of Computer Science* (pp. 243). Porto, Portugal: University of Porto.

Treuhaft, R. N., Law, B. E., & Asner, G. P. (2004). Forest attributes from radar interferometric structure and its fusion with optical remote sensing. *BioScience, 54*, 561–571.

Treuhaft, R. N., & Siqueira, P. R. (2000). Forest structure of vegetated land surfaces from interferometric and polarimetric radar. *Radio Science, 35*, 141–177.

Walker, W. S., Kellndorfer, J. M., & Pierce, L. E. (2007). Quality assessment of SRTM C- and X-band interferometric data: Implications for the retrieval of vegetation canopy height. *Remote Sensing of Environment, 106*, 428–448.

Woods, A. J., Lammers, D. A., Bryce, S. A., Omernik, J. M., Denton, R. L., Domeier, M., et al. (2001). *Ecoregions of Utah (2 sided color poster with map) (scale 1:1,175,000)* Reston, VA: U.S. Geological Survey.