FIA BioSum: A User's Guide (DRAFT)

By (tentatively ordered as of 12/31/2018)
        Larry Potts
        Lesley Bross
        Sara Loreno
        Tina Mozelewski
        Jeremy Fried

31 December 2018

NOTE: This Users Guide is in DRAFT status. Changes will continue to be made until it is published as a PNW GTR in late 2019

Table of Contents

Acknowledgements

Foreword

Development of Bioregional Inventory Originated Simulation Under Management (BioSum), an analysis framework supported by field data collected by the USFS Forest Inventory and Analysis Program, began in 2001 with a mission of "forest biomass summarization" as the escalating cost of suppressing large wildland fires inspired a shift in thinking about managing forested landscapes to reduce wildfire hazard. Given that many forests characterized by high fire hazard have an abundance of small-diameter trees traditionally considered unsuitable for conversion into merchantable wood products, concern surfaced as to how to finance fuel treatments in such forests. Much of the removable material consists of "dirty" (i.e., including non-wood components such as bark and foliage) chips generated by chipping entire trees of noncommercial species, and tops and limbs of merchantable trees. A widely touted solution was to further develop the bioenergy industry and encourage appropriate siting of biomass processing facilities to create markets for this otherwise unutilized resource. This led to numerous questions about production economics – how much merchantable wood and dirty chips would be produced by landscape-scale fuel treatment? How much could be feasibly recovered and hauled to processing facilities? Where should bioenergy facilities be built? Increasingly, there were also silvicultural questions – how much area is in need of treatment? What fraction of this could be treated without subsidy? On how much land and for how long would treatments be effective? The list of questions grew quickly, and providing timely responses based on statistically representative data and scientifically valid analytic approaches was challenging. Most questions were sufficiently complex that they went far beyond what any ad-hoc analysis of even the most comprehensive forest inventory data could deliver.

We developed BioSum 0.1 as a framework for analyzing and summarizing woody biomass possibilities, relying on the systematic sample of forests as represented in data collected by the US Forest Service's Forest Inventory and Analysis program. It included off-the-shelf stand simulation, harvest cost, and fire hazard models; a custom-built, geographically explicit transportation cost model; and multiple technical and decision focused assumptions, parameters and heuristics. The system had to be "hand-cranked" to generate useful information, and iterative analysis and experimentation were woefully constrained by the lack of automation in that analysis "system". Initial results for a study area in southwest Oregon became available at the time the Healthy Forests Restoration Act was being debated by the U.S. Congress in 2002, and were received with great interest. Considerable interest was expressed in building BioSum models for other areas of the country. An analysis of a larger study area in western Oregon and northern California soon followed, as well as a bi-state effort covering all of Arizona and New Mexico. Results were used, for example, in strategic planning by the U.S. Forest Service Region 3, the California Department of Forestry and Fire Protection, and a group seeking investment in a bioenergy plant in Lakeview, Oregon.

To enable this kind of analysis to be easily, even routinely, carried out anywhere that the requisite data is available, the PNW Research Station's Focused Science Delivery program provided a seed grant through the National Fire Plan (which was matched and greatly extended by PNW's Forest Inventory and Analysis program) to integrate the tools in the BioSum framework to support analysis of questions about fuel treatment feasibility, costs, and yields, and identifying promising locations for adding wood processing capacity. As development proceeded, new uses were contemplated and the system evolved to allow users to rely on any stand characteristic or set of characteristics as the basis for triggering management activity and/or to define effective management. Moreover, the model evolved from a static, one time evaluation of effectiveness comparing a pre-treatment to a post treatment value to modeling a dynamic, multi-decade sequence of silvicultural activities, costs, revenues and product flows, thanks to tighter integration with the FVS stand projection capability. This advance allowed, for example, evaluation and comparison of treatment longevity among treatments across all kinds of forests. Additional funding to support these developments was provided by the California Energy Commission, the Energy Biosciences Institute, Oregon Department of Forestry, the Joint Fire Sciences Program, and the Pacific Northwest Research Station.

Conceived as an "analyst-friendly" modeling tool (not a "goof-proof" gaming exercise for casual users), this implementation of BioSum is intended for analysts accustomed to using, or willing to master, the Forest Vegetation Simulator (FVS) stand projection model. Familiarity with Microsoft Access is essential for manipulating and querying complex databases, and ability to follow instructions describing a GIS workflow is necessary to compute vehicle travel time estimates for the haul cost component. Users should also understand what kinds of silviculture will both accomplish their objectives (e.g., enhanced resistance to fire, insects and disease; sustainable stand structures; desired habitat elements) and be realistic to implement; they must also be prepared to craft decision rules concerning locally appropriate operational harvest systems for implementing silvicultural prescriptions and what kind of results qualify as evidence of effective management.

For the most part, BioSum, now at version 5.86, is "analyst-friendly" in that interaction with the software occurs via error-trapped entry screens, push-button and check-box option selection, and drop-down menus. We hope that these features will promote analytic efficiency and ultimately reduce errors that would otherwise detract from the validity of results.

BioSum's developers have considered the possibility that users may ultimately wish to replace some components of the system with better or more locally appropriate versions when these become available. The largely modular program architecture and underlying database platform should facilitate such experimentation by the adventurous analyst seeking a slightly different kind of analysis. BioSum's developers have often exploited these features to extend analysis beyond what the software had been designed to do, and to build a complete replacement for the originally implemented harvest cost model. Some users will

likely be interested in developing back-end reporting tools for summarizing and displaying simulation outputs. Due to limitations of time and funding, as well as the challenge of choosing which outputs to focus on developing, this tool, as of version 5, is not yet designed to produce presentation-ready output. Because all outputs are stored as MS Access database tables, analysts can rely on general purpose office-automation software (e.g., Excel) or statistical analysis tools to format, graph and print simulation results. We expect that issues and foci of inquiry will vary among analysts, so some customization in the summarizing and interpreting of results is inevitable. We are happy to share with other users any back-end tools developed for BioSum at the BioSum web site (http://www.biosum.info), as they become available.

Finally, a word of caution – user-specified assumptions and parameters drive this modeling framework at least as much as the inventory data and transportation cost surface. To ensure that model results will be valid, care must be taken when designing and specifying treatments, defining effectiveness, and setting requirements for including full and partial FIA plots in the population of stands to be treated. Choose your assumptions wisely, and enjoy the powerful insights you are sure to develop using BioSum and FIA data!

Jeremy Fried, Portland, OR 31 December 2018

Chapter 1: Introduction

The Forest Inventory and Analysis (FIA) based BioSum (Bioregional Inventory Originated Simulation Under Management) has evolved as a policy analysis framework and workflow management software into something akin to a Leatherman™—adept at addressing a surprisingly broad spectrum of questions centering on forest management at broad spatial scale. BioSum can address a question as "simple" question as characterizing the distribution of per acre costs and net revenues of applying a particular silvicultural treatment over all the forests for which it might be suited as of today. It can also address more challenging questions that depend on illustrating alternative forest futures based on projecting the outcomes of alternative trajectories defined by the confluence of management, succession and natural disturbance for large, multi-million acre, multi-owner landscapes. BioSum can address complex questions concerning the effectiveness and economic feasibility of managing to improve forest health and resilience, identify propitious sites for facilities that manufacture wood products or generate bioenergy, and quantify the climate benefits produced by enhanced carbon sequestration and substitution potential that result from forest management, in fire-prone landscapes and elsewhere.

All of these questions can be addressed by 1) using FIA plot data to represent initial stand conditions so as to achieve a statistically representative analysis of all forests, or any specified subset, 2) using the Forest Vegetation Simulator (FVS) to implement silvicultural prescriptions, calculate stand metrics that relate to management effectiveness and carbon dynamics, project growth and mortality over time, and simulate disturbances, climate change or any other process for which FVS is parameterized to develop a richly detailed, broad spectrum of potential futures, 3) using modules in the BioSum software to account for costs and revenues of management and evaluate and rank effectiveness of alternative management choices.

BioSum was designed to operate as a "what-if" tool capable of providing objective, statistically valid answers to questions concerning landscape-scale forest treatment. Some examples of the kinds of questions BioSum can address include:

1. Where should a new biomass-to-energy conversion facility be situated?

2. Where is a good place to develop or expand capacity for processing merchantable timber?

3. On how much forest area would treatment X, Y or Z be effective if done today? How would the answer change if we cared about effectiveness over the next 30 years?

4. On how many acres could treatment X pay for itself?

5. How many more acres could be treated if diameter caps were not imposed? What about if $100 per acre was available to subsidize utilization of harvest residues?

6. How would my landscape change, in terms of area by stand structure class or other ecologically-driven aggregate, if fuel treatments were applied wherever they could be effective?

7. How much merchantable and submerchantable wood could a million-dollar fuel treatment program generate?

8. What proportion of the acres where fire hazard reduction would be beneficial are within national forests?

9. How would a putative price change in merchantable wood delivered to a mill or harvest residues delivered to a bioenergy or biochar facility affect the area treated and the wood yield? How do these outcomes vary as effectiveness criteria are changed?

10. Which forest restoration treatments have the greatest longevity?

It is also important to mention what BioSum is *not* designed to do. When used for facility siting, for example, the imprecision inherent in the simplifying assumption that all of the acres represented by a plot occur at the plot location, makes it inappropriate for tactical application in pinpointing the truly "best" site. Other considerations, such as power transmission infrastructure (and air quality constraints in the case of a proposed bioenergy facility), as well as access to work force, may be at least as important as minimizing costs of hauling feedstock. Owing to the typically low density of forest inventory plots, BioSum isn't well-suited to analyzing small areas. Anything less than several dozen potentially treatable plots would be cause for concern; robust analyses will build on hundreds of plots covering a broad range of forest stand structures. Moreover, BioSum lacks any spatial representation that would address how treatments relate to reduced burned area, for example, since such questions depend on the specific placement of treatments, something outside the scope of this tool.

Our research group uses BioSum to analyze how landscape-scale forest management could play out over the next few decades, under alternative scenarios focused on forest restoration, hazard reduction, and maximization of climate benefits via GHG emissions avoided. How you use the model is up to you. We're confident that there are many innovative and valid applications that we've never considered—we look forward to hearing about the new ways you find BioSum to be useful.

Required Data and Analyst Skills

Though the example questions outlined above run the gamut, most depend on the same underlying information: inventory tree lists, plot locations, processing site locations, the road network, and a healthy serving of assumptions and parameters brought by the user, including treatment specifications, effectiveness criteria, delivered prices for wood, and some cost parameters. The specific requirements for these assumptions and parameters are documented throughout this User Guide.

Successful BioSum analyses typically require a team covering multiple knowledge domains, including silviculture, fire and fuels management (if management is focused on enhancing resistance to fire), policy and economics, and perhaps most critically, the FVS modeling environment. An experienced FVS analyst can devise and implement appropriate simulations that account for all prescription elements in each silvicultural sequence, including surface fuel treatment; fuel managers and silviculturists also provide, for example, "yard loss" parameters to account for differences in logging slash generated by whole tree harvest systems versus log-length systems that leave tops and limbs where trees are felled. For FVS variants that don't contain an endogenous "establishment model",

local expertise may be required to develop regeneration parameters to be used in the simulation, for example, with the aid of a REPUTE analysis (Vandendrieche 2010).

What is BioSum?

BioSum consists of several components, including a Microsoft Windows™ computer program (BioSum Manager) that manages the workflow of a BioSum analysis (Figure 1), the FIA data on which it feeds, the protocols and guidelines outlined in this Users' Guide, and the problem definition, including assumptions and sideboards, brought by the user. BioSum Manager is a project management "shell", containing several task-oriented modules to manage analysis workflow, that requires that the 32 bit version of Microsoft Office be installed; all other required software is included in the BioSum distribution setup file. BioSum guides analysis workflow through four of BioSum's five modules, while protocols guide the user through travel time calculations, as follows. Starting with the FIA plot data for the user's area of interest, the Database module (1) loads that data as a representative sample of the forests for which management options are to be explored. The FVS module (2) is for defining management alternatives, formulated as a set of silvicultural sequences in FVS, implemented over time, that may differ among initial conditions and may include varying triggering criteria, covering the range of possibilities for which understanding is sought. A no action and/or business as usual sequence is typically an important alternative against which others can be evaluated. BioSum Manager generates FVS readable databases for exchanging information between BioSum and FVS. Travel times (3) are calculated between all plots and all processing facilities for use in computing cost per acre for hauling merchantable wood and utilized harvest residues a.k.a. "chips" or "energy wood". In the Processor module (4), all sequences modeled in FVS are processed through as many alternative Processor scenarios as desired—these may vary, for example, by wood valuation, harvest system used to implement treatments, or the set of available facilities to which wood can be delivered. The Treatment Optimizer (5) guides the user through developing one or more Optimizer scenarios specifying what effective management looks like and setting criteria for selecting the best of multiple effective silvicultural sequences, then summarizes some key landscape metrics, such as area treated by owner class, area feeding each processing facility, and quantities of wood and revenue produced, and cost incurred. It also summarizes outputs for each stand, and because each stand accounts for a statistically determined area of the forest, stand results can be aggregated however desired to describe averages or totals for fractions of the forested landscape that are of particular interest. Many more tables of policy relevant outputs can be produced by querying the tables that result from Treatment Optimizer and from intermediate modules such as Processor.
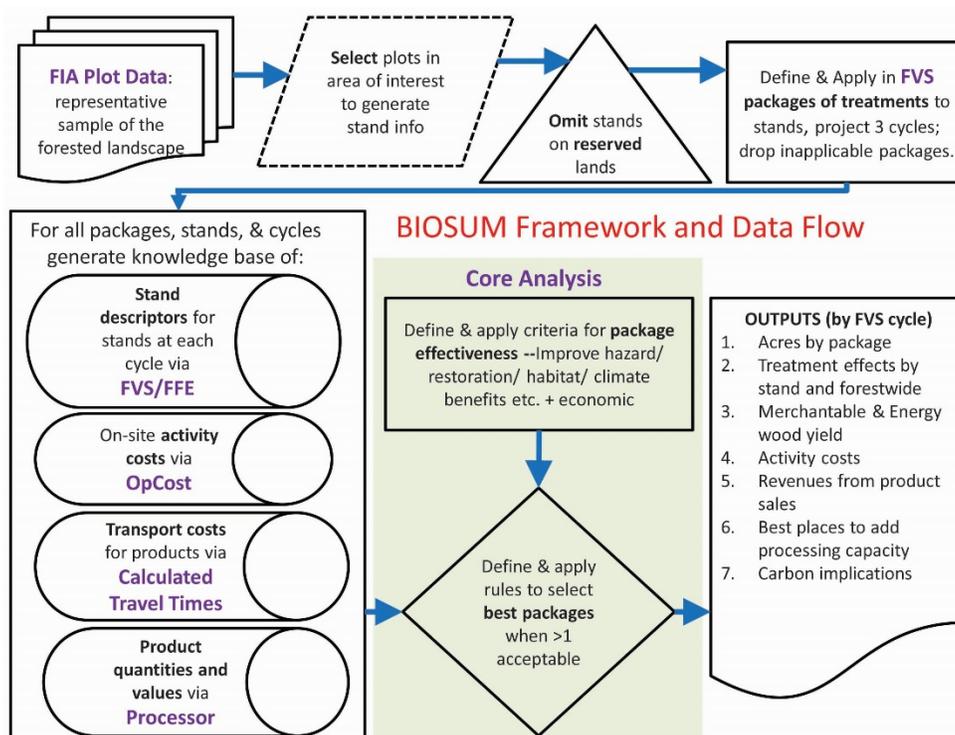
New at Version 5

An infusion of funding in FY18 provided by the PNW-RMA Program's Vegetation Monitoring Science and Analysis team BioSum enabled a temporary expansion of software development and analysis staff that delivered significant dividends in usability and capability compared to earlier versions of BioSum Manager:

- All FIA data is now translated into FVS input databases rather than text files, significantly easing the burden of readying data for analysis in FVS and enabling several new options for conveying additional, plot-specific information to FVS for use in the simulations (see below);

- Users can elect to import FIA down wood, duff and litter and field crew estimated Scott and Burgan surface fuel model into the BioSum project database and the FVS input database for use in simulations rather than rely solely on FVS default assumptions concerning surface fuels;

- Users can elect to import diameter and height measurements of trees collected at previous visits in order to calculate diameter and height growth for use in calibrating FVS projections; this will likely considerably improve the accuracy of FVS projections;

- Treatment Optimizer (formerly Core Analysis) now has access to any measured or calculated stand attribute for defining effective management based on its value at any simulation cycle or on a weighted average or total across all simulation cycles;

- Treatment Optimizer now has the capacity to identify the best management based on both stand condition and economic considerations calculated at any combination of simulation cycles;

- Most modules now include extensive, documented data consistency and integrity checks for analytical quality assurance;

- The harvest cost module, OpCost, was completely revamped and redesigned with an eye towards opening up the assumptions about suitable logging cost equations to the user by placing all harvest system definitions, and equations for the machines that comprise those systems, in Access database tables where they can be easily inspected, modified, disabled or supplemented. This enhancement allows control of which equations are used, enables adding new equations, and greatly facilitates adding new harvest systems, for example, to account for technology change;

- OpCost now includes a brand new set of equations that represent tethered harvest systems; we have found that tethered systems are frequently the lowest cost way to implement restoration management, especially on steep slopes.

Overview of BioSum Modules and Potential Analyses

Database- Chapter 2

A BioSum analysis begins with identifying an area of interest, typically in excess of 500,000 acres, to select plots to be included. This can include all the FIA inventory plots within a geographic area, or just a subset that meets user-specified criteria relating to, for example, ownership class, forest type, or wildland urban interface context. Because FIA data covers all forested lands, it is usually important to identify (and remove from consideration) plots within protected areas such as wilderness, parks and designated roadless study areas, for example. Some criteria for plot inclusion can be assessed from the inventory database (e.g., wilderness conditions will be attributed as reservcd=1 in the COND table); others can be derived via overlay with GIS layers. Some plots have more than one condition present – for example, two age classes of forest, or a forest and a non-forest condition, or forest controlled by members of two different landowner groups (e.g., a National Forest and a private holding) – so conditions are the basic analysis unit on which treatments are simulated. Conditions become stands when loaded into FVS.

Because the FIA inventory is a statistically representative sample of the whole forest, BioSum calculates an "expansion factor" for each condition that accounts for the number of acres in the forested landscape represented by the condition. These expansion factors are crucial for generating BioSum results that can be applied to the greater landscape. Data in FIADB formatted CSV files can be downloaded from the national FIA website (https://apps.fs.usda.gov/fia/datamart/datamart.html) and imported into an Access database to make it ready for use by BioSum. The database module imports this data into a "project" directory structure (a set of nested directories or folders) under which all information for a BioSum analysis is stored and processed, including inputs, intermediate databases and final outputs.

Travel Time- Chapter 3

There are two important calculations that require spatial analysis: 1) calculating the haul time from the plot to each potential processing site (this is used, with hourly truck operation cost and capacity, to develop unit transportation costs per acre for each stand), and 2) calculating the distance from the plot to the nearest road for use as a forwarding/skidding/yarding distance when modeling harvest costs. The timing within the overall BioSum workflow for these calculations is flexible, as long as it follows completion of the Database module and precedes the Treatment Optimizer. The analysis relies on a two point layers (plots and potential processing sites), and a line layer (roads) attributed with travel time per unit distance. This workflow can be carried out in ArcGIS or R, following the procedures outlined in this user's guide. The outputs are a table of truck travel times in ==units== between the publicly-available approximate locations of each FIA plot and a user-supplied list of existing and potential processing facilities that accept merchantable wood or harvest residues.

FVS- Chapter 4

Silvicultural prescriptions to be modelled are described and labeled by the user in this module, then assembled into sequences of prescriptions (referenced as "packages" in BioSum parlance) that are applied over time. Multiple treatments can and should be specified so that in the later stages of analysis, the Treatment Optimizer module will have a range of choices from which to choose for each stand. It is generally the case that no single prescription is best for all stands, no matter how "best" is determined.

The actual prescriptions in the form of FVS keywords and parameters are developed separately within the FVS software environment in SUPPOSE or off-line. Prescriptions designed for forest restoration are typically specified in terms of which trees to prioritize for removal and/or retention (for example, based on size and species), desired residual stand density, and limits on the size of trees allowed to be harvested, although prescription design is really only limited by the analyst's imagination and expertise in translating their ideas into a keyword file.

An unlimited number of silivicultural sequences can be specified, each containing a prescription or grow-only (no action) choice applied at the outset of each of four growth cycles of five or ten years each. Harvesting systems can be specified separately for each prescription, if desired, or specified as applicable to all prescriptions at a later stage of the analysis when developing Processor scenarios. BioSum performs a number of translations on and reformats the FIADB data to ready it for projection and simulation in FVS (e.g., supplying site index values compatible with the site index equations used by FVS and mapping FIA species codes to their FVS counterparts), then generates one FVS readable Access database per FVS variant that can also be populated with FIA collected information on down wood and other surface fuels or Scott and Burgan surface fuel models, when available.

The last workflow phase in the FVS module imports FVS output into the BioSum project environment so that, for every tree in the FVS CUTLIST tables (which contain records of trees harvested by the prescriptions applied in each silvicultural sequence at each growth/treatment cycle), biomass and volume can be computed using the same species

and region-specific FIA equation systems used in standard FIA compilation and reporting systems.

Processor- Chapter 5

BioSum's PROCESSOR module accomplishes two essential tasks to characterize, for each combination of stand and silvicultural sequence and for each growth or treatment cycle, wood production and associated value and costs: 1) computing per acre volume, weight and value of a) harvested merchantable wood, by user-specified species group and tree diameter class, and b) "dirty chips" (so-called because these chips derive from harvest residues delivered to the landing, and contain non-woody materials such as bark and foliage that make them unsuited for producing pulp, but potentially salable as fuel for bioenergy generation, depending on markets); and 2) estimating per acre treatment costs for harvest operations and any related activities such as treatment of surface fuels. Most of task 2 is accomplished using the OpCost software (Bell et al. 2017a), developed in R, with logic similar to Dykstra et al.'s (2009) Fuel Reduction Cost Simulator (FRCS; Fight, Hartsough and Noordjik 2006) which OpCost replaced. OpCost differs from FRCS in that it relies on parameters from more recent cost studies for the machines used in some of the supported harvest systems and is backed by a validation study (Bell et al. 2017b).

OpCost relies on machine production rates from the peer-reviewed harvest operations literature, and based primarily on elemental time analysis studies, to estimate harvest cost for all simulated forest operations activities in BioSum. The key inputs are similar to those in FRCS: harvest system, field observed slope and yarding distance (relying on distance from plot center to the nearest road as a proxy) at the stand level and, for four tree size classes, average volume per tree, number of trees harvested per acre, wood density, hardwood fraction, and residue fraction. OpCost estimates machine usage times for all machines in a harvest system, and translates these to cost estimates based on conventional cost control processes (D. M. Matthews 1942). OpCost also estimates "move-in" costs for getting equipment to a treatment site and ready to operate. These are ultimately combined to generate a total cost to harvest and process wood volume on a per acre basis.

Treatment Optimizer- Chapter 6

Treatment Optimizer integrates intermediate outputs from all the analyses generated via earlier modules in the workflow into user-crafted scenarios that address the landscape scale policy questions for which BioSum was designed. In Treatment Optimizer, analysts can interactively specify effectiveness criteria based on any of the FIA or FVS attributes at any simulation cycle or combination of cycles (e.g., weighted average values or totals over all cycles) to define what should be considered effective management. For example, effectiveness could depend on the post-treatment crown fire potential at cycle 1 being 20% lower than the pre-treatment crown fire potential, as indicated by one or more stand attributes. Alternatively, effectiveness could require that mean mortality volume over all four cycles under moderate fire conditions be lower for silvicultural sequence with active management than the mean mortality volume in a grow-only sequence. With the right choice of COMPUTE variables specified in the FVS prescription keyword file to allow tracking this attribute, one might deem sequences that move 50% of the stand basal area into trees larger than 25 inches dbh by the end of cycle 4 to be effective if the goal is to

accelerate the development of old growth characteristics. Up to four different stand attributes can be used in defining management effectiveness, and complex logic is supported, including making effectiveness contingent on improvements in more than one attribute, or limiting the allowable disimprovement in one attribute that accompanies the improvement in another.

Stand-sequence combinations that don't achieve the user-defined effectiveness threshold drop out of consideration, and those that remain are subject to heuristic optimization, where a user's overarching goal *after ensuring management effectiveness*, drives selection of the overall best treatment for each stand and the acres it represents. An optimization criterion could, for example, be minimizing crown fire potential, maximizing quadratic mean diameter in the residual (post-treatment) stand, maximizing mean canopy cover over 40 years, or minimizing net management costs (or maximizing net revenue). Users also specify attributes or treatment rankings to break ties that might occur in the optimization so that, for a given scenario, there will be a single best sequence identified for each stand that has one or more effective silvicultural sequences.

Treatment Optimizer can then summarize outcomes of the scenario for the whole forest landscape—for example, how many acres can be effectively managed, how frequently does each silvicultural sequence considered prove best, how much merchantable wood and collected harvest residues would be generated, and net revenue or cost, per cycle, for implementing the management program on all eligible acres. These kinds of summaries can be aggregated and reported, for example, by ownership, product species group, product tree diameter class, forest type, stand density class, treatment style, or cycle (decade). Aggregation of outputs is also possible by wood processing facility. Sequences and acres can be filtered by whether net revenues exceed any specified threshold, including negative values that indicate allowance for some degree of subsidy. Some of these tabular outputs are reported automatically in database tables; others can be easily generated via queries of the Treatment Optimizer table output.

Advanced Ad-hoc Analyses

For some questions, linking Treatment Optimizer results to attributes in the plot and condition tables, some of which may have been added to these tables via GIS overlay or other model application, can address a broader set of questions that link to location. For example, to name a few, how much of the wood from fuel treatments would be generated within the wildland urban interface? How much of the area requiring treatment is on steep slopes? How does site class impact treatment effectiveness and longevity? How does management success and cost relate to owner group? Some analysts may wish to explore post-treatment forest structure and fire resilience, compare post to corresponding pre-treatment values, compare habitat quality under a management sequence with what a grow-only simulation delivers, or track carbon pools over time, in managed and unmanaged sequences, with and without disturbance at different frequencies and intensities. All of the FVS table outputs that the analyst chooses to specify in their KCP files (that control FVS operation), including carbon, trees, potential fire and more are available for this kind of analysis by analysts comfortable developing database queries to stitch these different kinds of results together with the BioSum outputs. Moreover, because each stand

is linked to a plot with a defined location in space, some users may wish to explore results geographically, for example, to identify regions where treatment effectiveness is generally high, economic feasibility is high, or pre-treatment hazard is especially high. It can also be useful to plot stands attributed (i.e., shaded) by product destination to convey the concept of potential woodsheds under a given configuration of processing sites.

An earlier version of BioSum was used to jointly optimize treatments and wood processing facility locations (Daugherty et al. 2008), and this may be of interest to some users. In that application, analysts extracted the data developed from the modules described in chapters 2-5, as input to a mixed-integer optimization algorithm to arrive at a joint optimization solution to the problem of which treatments to apply and which processing sites to build (and to what capacity). We used the same kinds of objectives for treatment effectiveness as can now be applied as heuristics in the Treatment Optimizer module.

Work Flow Considerations

With the exception of Geoprocessing, for which there is some flexibility as to the "when", the modules outlined above must be executed sequentially because each successive step relies on analysis completed in the previous step. When performing iterative simulations, it is sometimes necessary to revisit earlier steps to implement alternative assumptions. For example, to model the effects of different assumptions about merchantable wood prices, the analyst would need to rerun the Processor module before running a new Treatment Optimizer Scenario (and linking to the newly created Processor scenario). To understand the effect of a different network of processing sites, the analyst would revisit geoprocessing operations before re-running Treatment Optimizer.

The time required to complete each module varies. Database can typically be completed quite quickly – in as little as an hour for a single state, depending on whether down wood and calibration data are also requested. Geoprocessing for a large area can requires several hours to complete, and has some iterative elements that may require multiple analyses to get a complete travel times table. However, this workflow is usually only done once. After silvicultural prescriptions have been generated as KCP files (the challenging part), the push of a few buttons is all that is required to initiate and complete the tasks in the FVS module, but some of them take several hours for a statewide project with many prescriptions specified. Running the prescriptions in FVS requires the most time; on some statewide projects, this has consumed a week of computer time, though that particular task can be divided among multiple computers and batched up so that it can run largely unattended.

Processor is easy to setup, and runs unattended for several hours when there are many prescriptions (e.g., 30) and multiple FVS variants for the project area. Treatment Optimizer runs quickly once the effectiveness and optimization criteria have been established by the user, usually completing in 2 to 5 minutes, though depending on how complex and ambitious, defining those criteria could take as little as 5 minutes or as long as an hour, the first time through. That rapid execution turnaround time makes it easy to modify and rerun scenarios to see the effect of different assumptions on the questions of interest.

A word about simulation

Simulation models like BioSum are designed to capture the essence of a "system" (in this case, the application of landscape-scale forest management) in an area, volume, value, cost and effectiveness accounting model intended to mimic the functioning of the system. If sufficiently realistic, simulation models can help identify solutions to critical problems, and robust models may even be useful in addressing problems not anticipated when they were developed. One motivation for developing this simulator was concern that much of current decision-making about fuel treatments, and forest management in general, is driven more by conjecture and belief, or by distortions introduced by simplistic performance benchmarks and the budgeting process, than by forest and fire science.

One of the under-appreciated aspects of simulation models is their ability to enhance our understanding of a system by quantifying the relationships within the system being modeled. In the case of using BioSum to evaluate forest restoration potential, we are forced to quantify achievement of fuel hazard and other restoration goals, and to be explicit about the silvicutural treatments we would consider implementing to them. We are also forced to examine closely the costs associated with implementing prescriptions and with hauling removed material to places where it can be utilized, providing some revenue to offset management costs.

This kind of simulation can be very useful for assessing the impacts of changes in both the policy environment (e.g., what kinds of fuel treatments are acceptable, how much funding is available to subsidize fuel treatments) and the economic environment (e.g., investment in new processing capacity, availability of new, low-cost techniques for removal of small diameter woody material). Simulation can be automated so as to produce a wide variety of possible scenarios and associated outcomes, without regard to current treatment plans, in the hope that a near optimal solution will be identified. Alternatively, it can begin from a "base case" that represents current plans and conditions, and be guided by an analyst familiar with the system and with the range of scenarios that might be worth considering.

Given the time and expense required to build and run a simulator like BioSum, some will ask, "Why bother?" One answer is that the complexity of managing large, multi-owner, forested landscapes is such that good decisions are difficult to make. Another is that bad siting decisions could have catastrophic consequences for biomass processing investors, the ecological integrity of the forested landscape, or both. While an experienced procurement forester may recognize and understand some of the simpler qualitative relationships in the system, predicting economic feasibility is far more difficult. Complex interactions can be explored in-depth using simulation, and intuitive expectations with respect to simpler relationships can be rigorously tested and refined.

In any case, simulation offers a safer means of experimenting with a system than implementing changes on a trial basis. No prudent land manager would embark on a fuel treatment program without some confidence that there would be markets for the materials removed (or that funding would be available to cover treatment costs) and that the treatments would be effective. No investor would build a bioenergy facility without confidence that a consistent supply of feedstock would be available to sustain it over a reasonable period of time. This kind of tool offers all involved an opportunity to conduct

due diligence and to arrive at an analytically supportable basis for management decisions that commit funds and/or future actions. In this context, a simulation approach can lead to better management decisions.

Simulation modeling is both a science and an art because of the inescapable trade-off between model realism and data requirements. Relatively simple models such as "my fuel treatment planner" can help us understand treatment options, costs and effectiveness at a stand level, but are not of much help in answering questions about where to build processing capacity or forecasting how much funding would be needed to have an appreciable impact on fuel hazard at landscape scale. On the other hand, extremely complex models, capturing every nuance and possible permutation of the system, can be burdensome to parameterize and difficult to interpret. It is generally agreed that the best models are those which capture the elements of a system that is of critical concern to the clients needing the model, without being needlessly complex. BioSum's design is the result of a conscious effort to balance complexity with realism, and to allow the user to make a number of decisions concerning the appropriate level of complexity. Simulation parameters for everything from effectiveness criteria to prescription sets to harvest and haul costs can be adjusted to enhance model realism. Sensitivity analysis will be an excellent way to examine the impacts of the assumptions specified by model users and to provide information about which may merit more intensive effort to perfect.

In developing BioSum, we tried to tailor the model to the needs of both land managers and those interested in building processing capacity. We hope that these users will find the end product of our collective effort to be worthwhile.

Learning BioSum

This analysis system is robust, comprehensive and flexible; however, it also has the potential to feel overwhelming from the perspective of a new user. Analysts should allow themselves time to come up to speed, and carefully read the documentation for each module to ensure that the analyses they conduct will be valid. While the treatment optimizer module is enticing and eagerly anticipated by most users, the results will only be valid if care has been taken in the preceding workflow.

New users are advised to start small—pick a small area and develop just a couple of prescriptions to build a test dataset to explore the kinds of analysis that are possible using Treatment Optimizer. Then, if all looks promising, collect and assemble the data, assumptions and parameters required for all modules to support an analysis of the target study area. For cases where this documentation does not address questions you have about how to run the model, BioSum experts at PNW-FIA can be contacted for advice. CONTACT INFO

Literature Cited

Barbour, R. James, Jeremy S. Fried, Peter J. Daugherty, Glenn Christensen, and Roger Fight. 2008. "Potential biomass and logs from fire-hazard-reduction treatments in Southwest Oregon and Northern California." Forest policy and economics 10(6): 400-407.

Bell, Conor K., Robert F. Keefe, and Jeremy S. Fried. 2017a. "OpCost: an open-source system for estimating costs of stand-level forest operations." Gen. Tech. Rep. PNW-GTR-960. Portland, OR: US Department of Agriculture, Forest Service, Pacific Northwest Research Station 23 p.

Bell, Conor K., Robert F. Keefe, and Jeremy S. Fried. 2017b. "Validation of the OpCost logging cost model using contractor surveys." International journal of forest engineering 28(2):73-84.