EasySDE User Guide

Oscar García
University of Northern British Columbia
garcia@unbc.ca

September 2009

Contents

> The Guide is definitive. Reality is frequently inaccurate.
> (Douglas Adams, *The Restaurant at the End of the Universe*)

1   What

EasySDE is essentially a GUI front-end to some old software for developing forest height growth (or site index) models. Fairly sophisticated and statistically efficient methods are used, making good use of almost any kind of repeated-measurements data. The old software was not particularly user-friendly, which undoubtedly limited its appeal. *EasySDE* greatly simplifies the process. Its ease of use should compare favourably to most other site modelling procedures.

In short, various parametrizations of the Richards equation can be used. For estimation, environmental and measurement error sources are modelled through a stochastic differential equation (SDE). All parameters are estimated simultaneously by maximum likelihood (don't panic, all this will be explained shortly). Data can consist of two or more height-age pairs on any number of sample plots. The number of measurements per plot and the measurement interval lengths can vary. Either permanent sample plot or stem analysis data may be used.

Some methodological background is given in the next section. *EasySDE* usage instructions follow. This should be more than sufficient for model development using the variants that past experience has shown to be the most useful. If you really know what you are doing, and feel adventurous, for added flexibility you might want to run the underlying programs directly. Or even hack into the code, or use some of its components for other purposes. Details are contained in the *Advanced and Programmer Notes* included in this distribution. All sources are freely available under the MIT license (www.opensource.org/licenses).

The models and methods are described formally in García (1983), and Seber and Wild (2003, pages 354–356 and Appendix C), but what follows should be enough to get you started. Extensive experimental results are reported, fo instance, in García (1999, 2005). Related methods were used by Rennolls (1995).

The software runs on Microsoft Windows and on Linux. Everything here applies equally to both unless stated otherwise. The screen pictures were generated on Windows, Linux will look slightly different.

2 Theory

2.1 Growth, height, and site

On a given site, and under "average" weather conditions, the predicted height growth depends on the current height:

$$\frac{\mathrm{d}H}{\mathrm{d}t} = f(H) , \tag{1}$$

where H is top height and t is time (or age). It is usually assumed that, within limits, other stand variables such as density do not significantly affect top height. It might be argued that growth also depends on age. But in fact, age, as elapsed time since birth, does not have a physical presence at any given time — other than perhaps in the number of growth rings. Therefore, it can hardly have a direct causal effect. Only accumulated changes reflected in the current condition (state) of the stand can be causes. In principle, ageing of meristems might have some influence on growth rates. For the most part, however, tree size should be the dominant variable.

Starting from some initial height H_0 at time t_0, accumulation (i.e., integration) of the growth rate in (1) produces the height H_1 at some other time t_1. This can be calculated taking advantage of (1) being a "separable differential equation":

$$\frac{\mathrm{d}H}{f(H)} = \mathrm{d}t$$

$$\int_{H_0}^{H_1} \frac{\mathrm{d}H}{f(H)} = \int_{t_0}^{t_1} \mathrm{d}t = t_1 - t_0 .$$

If we can analytically integrate the left-hand side, and then solve for H_1, we obtain a formula for H_1 as a function of t_1, t_0, and H_0.

A flexible and commonly used model has $f(H) = \eta H^m - \kappa H$. It was proposed in this form by von Bertalanffy (1938, 1949), not just with $m = 2/3$ as it is often said, and was popularized for plant growth by Richards (1959). In forestry the model is sometimes called "Chapman-Richards", following Pienaar and Turnbull (1973). It will be convenient to use a power transformation H^c, with $c = 1 - m$, and write it as

$$\frac{\mathrm{d}H^c}{\mathrm{d}t} = b(a^c - H^c) . \tag{2}$$

The equivalence can be verified by differentiating on the left-hand side, and re-arranging. Integrating and solving for H, as explained above, gives the height-age equation:

$$H = a[1 - (1 - H_0^c/a^c)\mathrm{e}^{-b(t-t_0)}]^{1/c} . \tag{3}$$

I have omitted the subscript 1 to simplify. We shall also refer to t as *age*, although it only needs to be a relative time or nominal age; shifting by a constant does not make any difference.

For a whole growth curve, often H_0 and t_0 are taken as zero in (3), forcing the curve to go through the origin. When using breast-height age, $t_0 = 0$ (sometimes 0.5), and $H_0 = 1.3$ m. This is a typical sigmoid curve. The parameter a is the asymptote, b is a "speed" factor, scaling the time axis, and c is a shape parameter determining the relative height of the inflection point.

2.2   Site index

Top height growth is widely used as an indicator of site productivity. The advantage over more direct measures, such as volume production, is that it is less affected by stand density. It is assumed that height-age curves, e.g. (3), vary across sites. Depending on site quality, the curves lie at different levels, without intersecting each other. The curves may be labelled in various ways, a common one being by the height reached at some specified reference or base age: the *site index*.

Mathematically, this corresponds to a *one-parameter family of curves*. That is, the curve equations for different sites differ only in the value of one scalar

parameter. For instance, in (2) and (3) one of the parameters a, b, or c, may change with site quality, while the others are common to all sites. More generally, this may happen after a re-parametrization: any or all of a, b and c may be assumed to be functions of some other parameter that varies with site. We call a site-dependent parameter *local*, being specific to a sample plot, and the others *global*, pertaining to the model as a whole.

The relationship between the local parameter and the conventional site index S is found by substituting the base age t_s in (3):

$$S = a[1 - (1 - H_0^c/a^c)\mathrm{e}^{-b(t_s - t_0)}]^{1/c} . \tag{4}$$

In simpler cases (4) can be algebraically solved for the local parameter, and this substituted in (3) to obtain a height-age equation indexed by S. A numerical solution may be needed in other instances.

All this is fairly straightforward in a deterministic world, or when dealing with predicted trends as up to here. With actual stand observations we need to be more precise. Implicit in the traditional concept seems to be the idea that site curves and site index represent some sort of average over hypothetical stands that might grow in the site. That is, site index is a property of the site. More recently, some researchers have thought of site index as the actual height reached by a specific stand at the base age: a property of the stand. Failing to distinguish between these two different definitions, which we might dub "site site-index" and "stand site-index", has caused a great deal of confusion and unnecessary controversy (see, for instance, Bailey and Cieszewski, 2000, and references therein).

We take the site site-index view. Specifically, *Site index is the top height at a base age, predicted "at birth", for any hypothetical stands (of a given species, etc.) that might grow on a site*[1].

At least in this view, the site index base age is essentially arbitrary, a device for assigning a number to each curve in the family. Therefore, modelling

---

[1] *Predicted* refers to some kind of point estimate, such as an expected, most likely, or median value. Mean, mode and median converge to the same value as sample size increases, so, presuming an infinite number of hypothetical stands, it makes no difference in the definition. There are differences for estimation, as discussed later. Note that heights predicted by (3) for an existing stand, and in particular the predicted height at base age, differ depending on its current (t_0, H_0). For the same reasons that your life expectancy at birth differs from your life expectancy now.

6

approaches where results change with the chosen reference age are, at best, ugly. Our methods are strictly *base-age invariant*, in the sense of Bailey and Clutter (1974).
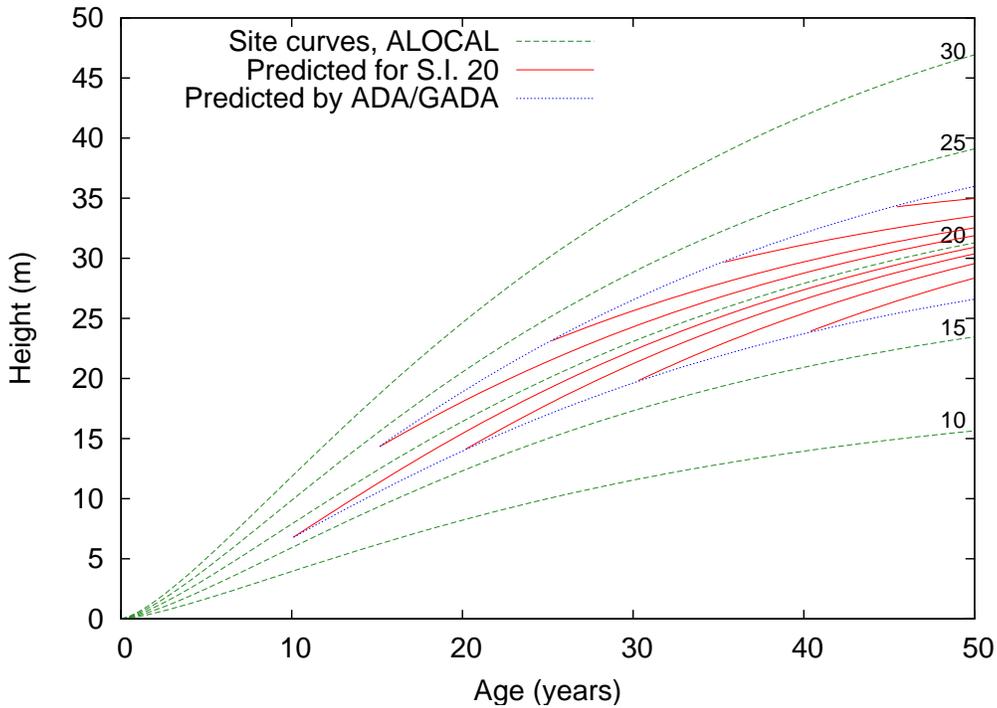

Figure 1: Dashed green and dotted blue: site-index curves. Continuous red: projections for off-curve site-index-20 stands.

To illustrate some of the conceptual subtleties, assume a model with a as the local parameter, and height zero at age zero ($ALOCAL$ in Section 2.4). The base age is 25 years, $b = 0.04$, and $c = 0.7$. In Figure 1, the dashed green curves are site index curves. That is, heights predicted "at birth" by equation (3), with $t_0 = H_0 = 0$ and various values of the site parameter a. Now, suppose that for whatever reason (e.g., unfavourable weather or heavy weed competition) a 10 year-old stand growing on site 20 happens to be 15% short of the "normal" height for that age and site. See the point at age 10 just below the site-index-20 curve. What should be the future predicted height? The popular ADA/GADA methods (Bailey and Clutter, 1974; Cieszewski and Bailey, 2000) use a transition equation similar to (3), but obtained by eliminating the local parameter in the site-index function.

That greatly simplifies estimation (no local parameters to deal with), but it assumes that, given the current stand condition, growth is independent of site quality. The ADA/GADA projection is the site index curve that passes through the current point, the dotted blue curve. Our approach uses equation (3) with (t_0, H_0) being the current point, and gives the red continuous curve. The stand is still growing on site 20, so that predicted growth should be higher than for a stand of the same age and height but on site 17. Similarly, projections for stands that overshoot their nominal trajectories should be lower than those for stands on higher site qualities that stay the course.

Warning: As with any single-number measure, the concept of site index can fail if pushed too far. It can be a good indicator of relative productivity within biogeoclimatic regions where growth patterns and limiting factors are similar. Otherwise, not just the level, but also curve shape may vary in such a way that crossings occur. An appropriate stratification should be used. Alternatively, multidimensional indicators might be possible, although little or no research in this direction appears to have been done. An experimental model variant with two local parameters is included in this package.

2.3   Estimation

Rational parameter estimation requires some sort of model for the variability of the observations. Experience with regression suggests that moderate deviations from the assumptions — e.g., additive errors, uniform variance, independence, normality, etc. — may not be overly critical for the quality of the estimates. A rough characterization of the error structure is nevertheless desirable.

It is clear, for instance, that the deviations of a stand from the model of equation (3) would tend to increase with time, and that they are not statistically independent. This because the "error", due among other things to weather fluctuations, accumulates over time. Ordinary least-squares is therefore relatively inefficient. Seber and Wild (2003, Ch. 7) review several approaches to modelling this kind of data, and others are found in the forestry literature. Many, such as autoregressive and mixed effects models, trade a somewhat crude representation of error structure for the convenience of using standard statistical packages. We use a more realistic stochastic model that requires

special-purpose software. Hopefully, EasySDE now makes its application not more difficult than other methods.

For each sample plot, we represent the effects of "environmental noise", mostly weather, as a perturbation added to the right-hand side of (2). The perturbation is assumed to be a continuous random process, with values that are independent on non-overlapping intervals. Technically, what is called white noise, the formal derivative of a Wiener or Brownian motion process:

$$\frac{\mathrm{d}H^c}{\mathrm{d}t} = b(a^c - H^c) + \sqrt{2b}\sigma\dot{w}(t) \ . \tag{5}$$

Here $w(t)$ is the standard (mean 0, variance 1) Wiener process, and σ is a new parameter determining the perturbation intensity or variance[2]. This is called a *stochastic differential equation* (SDE). To keep things "simple", we neglect any correlations across plots.

In addition, we recognize that an observed stand top height h_i may differ from the real H_i due to measurement and sampling error. For mathematical convenience we assume

$$h_i^c = H_i^c + \sigma_m\epsilon_i \ , \tag{6}$$

where the ϵ_i are independent standard normal variables, and σ_m is another parameter to be estimated. This makes the error somewhat higher for larger heights (standard error approximately proportional to H_i^{1-c}), which seems reasonable.

Finally, it might happen that growth perturbations are higher at the seedling stage than later on. An optional additional parameter σ_0 may be used to represent extra variability over some time previous to the first measurement.

The general model has then 8 "basic" parameters: $a, b, c, t_0, H_0, \sigma, \sigma_m$, and σ_0. In a specific model version or variant, any of these parameters may be either fixed (e,g. zero), global, or local. They may also be some function of other new global and local parameters — examples later. Actually, a "local parameter" represents many parameters, one for each sample plot. Typically, one needs to estimate hundreds of parameters. Although for most applications only values for the globals among the first five above are used.

---

[2] The factor $\sqrt{2b}$ implies larger effects for higher growth rates when b is local and σ is constant.

The data consists of a number of sample plots, each with a sequence of age-height observation pairs (t_i, h_i). We estimate the parameters by the method of maximum likelihood (ML). It is based on calculating the probability, under the assumed model, of obtaining the observed data. Considered as a function of the parameters, substituting the given data, this is called the likelihood function. The ML estimates are the parameter values that maximize the likelihood function. Apart from an intuitive appeal, ML estimation has a number of nice statistical properties. Two characteristics are particularly attractive for this kind of application. First, no matter how complicated a model, the procedure is well defined: one "simply" obtains the likelihood function and finds a maximum. Second, it is invariant under transformations, that is, the ML estimate for any function of the parameters equals that function of the parameter ML estimates. This is convenient where parametrizations are essentially arbitrary: should we get good estimates for the site index, for its logarithm, for a local parameter, for the height at a certain age? Unlike with other methods, all these are simultaneously possible and compatible.

The likelihood function is obtained through integrating the SDE (5), and incorporating the distributions implied by (6). If you must, see García (1983) or Seber and Wild (2003) for the gory details. A modified full Newton method, using first and second derivatives, is used for efficiency and reliability in minimizing the negative of the log-likelihood over hundreds of variables (parameters). The implementation makes use of the special structure of the problem, with partitioning strategies to take advantage of sparsity in the matrix of second derivatives. Details in García (1980), and in the *Advanced and Programmer Notes* in this package.

It might be worthwhile to clarify a couple of things that have caused some confusion. I probably shot myself in the foot by delving in the *Biometrics* paper into theoretically interesting but practically less important limitations of the model. Some have misunderstood this and dismissed the whole approach as useless, so let us try to keep things in perspective. No model is perfect, nor should it be: striving for perfection has an unacceptable cost in terms of parsimony and estimation efficiency. The first quibble was mentioning that (5) cannot be strictly correct; the additive random perturbation could conceivably make the (transformed) height to go negative. True, but in reality unlikely and inconsequential. The same thing is routinely ignored without a second thought when using regressions with non-negative variables. The

second issue may be a little more troubling, at least from a methodological point of view: Often, although not always, the estimate of either σ or σ_m turns out to be zero. Apparently, sometimes the data does not provide sufficient information to reliably separate the different variance components. Of course, the problem would go away if either one of the two sources of variability is ignored, as has been done in all the other methods I know of. At any rate, this does not seem to have any appreciable impact on the estimates for the parameters of interest (a, b, c, t_0), and might be taken as a warning against the temptation of introducing even more "realistic" error structures.

2.4 Model selection and hypothesis testing

Models with the same number of parameters can be compared and chosen according to the highest value of the computed likelihood, or of its logarithm, the log-likelihood. In comparing models with different number of parameters, extra flexibility needs to be penalized. Various theories suggest adding penalties of one-half to three log-likelihood units per parameter. Akaike's information criterion (AIC), perhaps the most popular, prescribes one unit. Note that the likelihood value by itself is essentially meaningless, it is the relative values, or differences in the logarithms, that are important.

For hypothesis testing, one approach considers a log-likelihood difference of about two units as "significant". More orthodox is a likelihood-ratio test. It is based on the fact that twice the log-likelihood difference is asymptotically distributed as a χ^2 with degrees-of-freedom equal to the difference in number of parameters. For examples see García (1999, 2005).

3 Program usage

3.1 Data format

The data must be arranged in a text file with three columns: plot number, age, and height. Plot numbers must be integers, with up to 9 digits. The other values are free format, and columns may be separated by any number of spaces and/or tabs.

Each row corresponds to one measurement date, with the sequence of plot measurements ordered by increasing age. Different plots must be separated by a blank line. See the examples in the *Work* folder.

Before going any further, it is a good idea to plot the data to check for "funnies". *SitePlot* may be used for this. A detailed description is in Section 4, but for plotting just data it should be self-explanatory.

Example data. Two data sets are included in the *Work* folder:

bcpine.dat Breast-height age (years) and top height (m) from permanent sample plots of lodgepole pine in the Interior of British Columbia. Taken from a file provided with the documentation of VDYP 6:
`http://www.for.gov.bc.ca/hts/vri/ip/software/vdyp`
`/vdyp_manual/vdyp99.html#overlay`

reventlow.dat Total age and heights from stem analysis of 18 beech trees in the First Copenhagen District done by Count Reventlow around the year 1800 (!). From Table 11 of Reventlow (1879). Converted to metric, and stump age calculated by linear extrapolation (a likely underestimate).

Notes on stem analysis. Often heights are interpolated for every year, using methods such as those described by Dyer and Bailey (1987). That is a Bad Idea, the artificial "data" can distort the real trends; use only values for the actual crosscuts. Because in general the tree tip corresponding to the ring count at a crosscut is located some distance above the crosscut, crosscut heights may be adjusted as shown by Dyer and Bailey (1987). However, as suggested by Milner (1992) and by Jim Goudie (quoted by Nigh (1995)), a simpler and probably better method is to use the crosscut height, reducing the ring count by 0.5 years.

3.2   Model specification

On running *EasySDE*, the window in Figure 2 appears. The items in it are as follows:
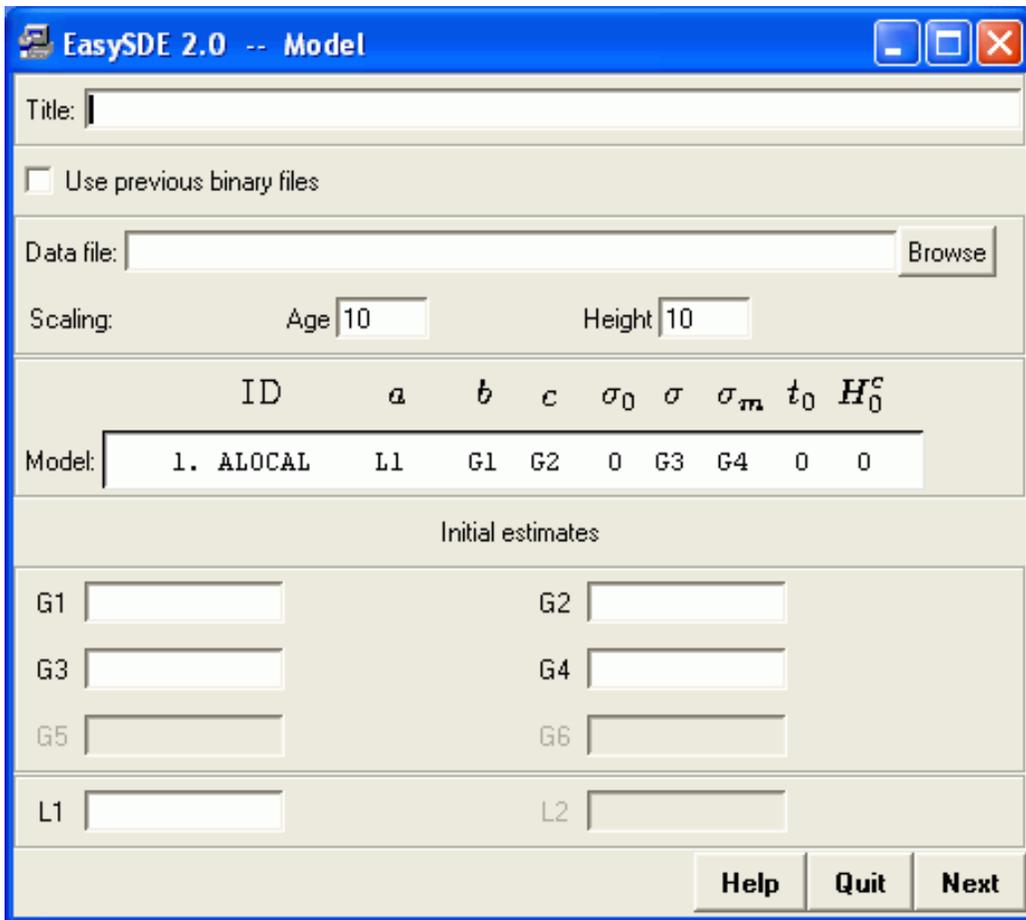
Figure 2: Stage 1: Model specification window

Title. One line of optional descriptive information to be printed at the top of the program output.

Binary files. If checked, the data and local parameter estimates from the previous run are used. The model version in the new run must use the same local parameter(s).

File. Name of the input data file.

Scaling. The ages and heights in the data file are divided by these values. Convergence is usually faster and more reliable if the scaled ages and heights are not too far from unity. If the scale factors are different from one, the initial estimates and the output correspond to the scaled model. Values like 10 or 100 make things easier, simplyfying the calculations when converting back.

Model. Click to select a model version. The `Gi` stand for global parameters, and the `Li` for local parameters. A number of models are available in EasySDE:

`ALOCAL` has a as local parameter, and is forced to pass through the origin. It produces anamorphic site index curves. `A-T0` to `A-T0S0` include non-zero origin parameters, and/or the additional initial variance σ_0.

`BLOCAL` to `B-T0S0` are similar, but taking b as the local parameter. The resulting curves are "polymorphic". Actually, proportional along the t axis.

`LINEAR` and `POWER` are more general models, where curve shape varies in a more flexible way with site. At the cost of an extra parameter. Unfortunately, only a numerical solution of equation (3) for the local as a function of site is possible, so they are not as convenient for many uses[3]. They can be useful for assessing the adequacy of the simpler models, and as a last resort if the others are not good enough. These models were used in García (1996).

`ABLOCALS` has both a and b as local parameters, so it does not produce traditional site index curves. It may be useful for hypothesis testing and experimentation. Graphing the relationship between the estimated a and b might suggest appropriate re-parametrizations.

`ALOCBH` to `ABLOCBH` are similar to the above, but using breast-height age. They force a height of 1.3 m at age 0.5 (the standard used in BC). If your breast-height definitions are different, the easiest may be to transform your data temporarily and adjust the results later.

These are the models that generally have been found most useful in practice. Other variants can be implemented with just a little more effort; see the

---

[3] Conjecture: there are no "truly" polymorphic models with a base-age invariant explicit expression for the indexing parameter in terms of site index.

Advanced and Programmer Notes, Section 4.

It is recommended to start with the simplest model in each class (e.g., `ALOCAL` or `BLOCAL`), and use the estimated parameters as starting point for more general variants. The *Binary files* option can be useful there.

Initial estimates. These are the starting points for the likelihood optimization iterations. They must correspond to the specified scaling, if any. For the local(s), a common value is initially used for all plots (unless the *Binary files* option is specified). Tab and Shift-Tab may be used to move among the entry fields.

Convergence can be slow or fail if initial estimates are not good enough. When possible, try stepping through models of increasing complexity, using previous estimates as starting points. Using *Binary files* takes advantage of the previous individual-plot locals.

It is also a good idea to try several starting points, to (partially) guard against local optima.

Some suggestions when starting from scratch: For a, guess a reasonable height upper bound. Parameter c usually lies between 0.5 and 1; try 0.7. Try the reciprocal of the (scaled) base age for b. The σ's are a little trickier; with decent scaling something around 0.05 should be OK.

Help. Displays an abbreviated version of this section.

Quit. As it says.

Next. Advances to the next stage.

3.3 Pre-processing

On clicking `Next`, the information is checked, and prepared for use by the main estimation program. If there are no missing or incorrect entries in the
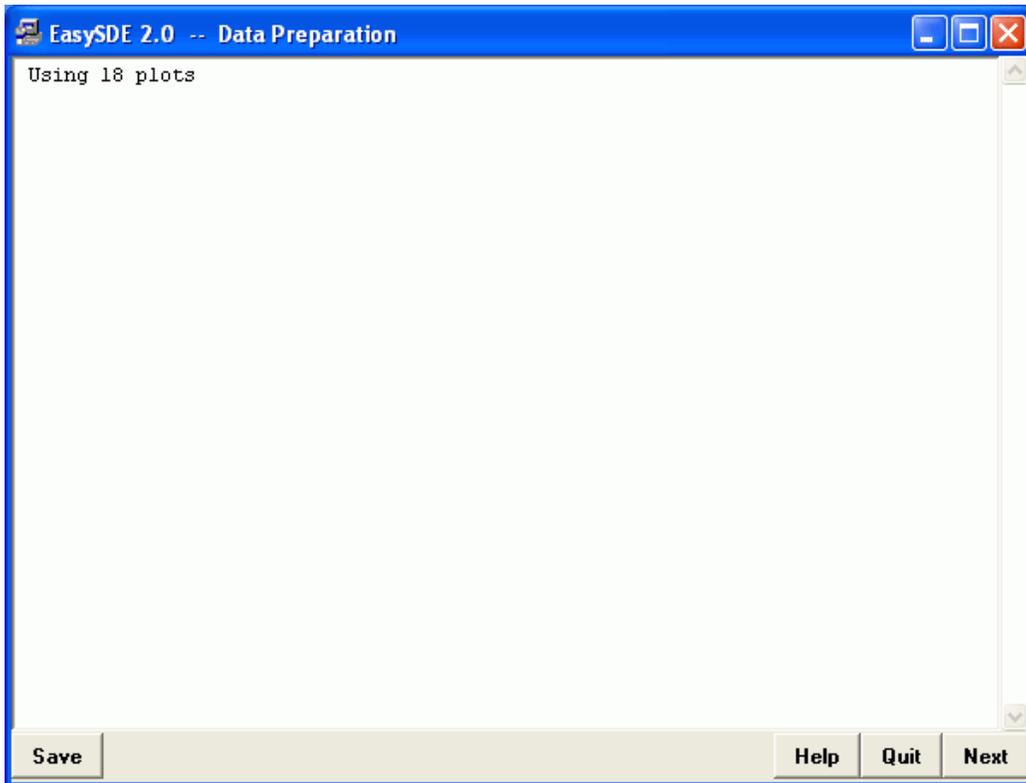
Model dialog, the window in Figure 3 appears.

Figure 3: Stage 2: Data pre-processing window.

No user input is required. The input data file is checked for inconsistent plot numbers and decreasing ages. If there are no error messages, the number of plots is displayed, and we can proceed to the next stage by clicking Next. The contents of the message window can be saved to a file through the Save button.

3.4 Run

The same window is then used to display the iteration log as the optimization proceeds. To monitor progress, values for the parameters and for the function being minimized (which is -2 times the log-likelihood) are shown for each

iteration, together with some other more esoteric information. Most relevant, the contribution to the objective function and the local parameter estimates for a few plots are shown, followed by iteration number, function, and global parameter values. This was really more useful in the good old days of overnight runs, to see if the program had not crashed.

It should be possible to abort the run by clicking the `Quit` button (but it isn't!). Otherwise, the program runs for up to a maximum of 200 iterations. If this happens try a better starting point and/or better scaling.

Clicking `Next` takes us to:

3.5 Results

A similar window displays the final results report. The report starts with the run identification, and the termination condition: "`CONVERGED`" for a successful run.

For each plot, the following is shown: Plot number. Number of measurements. Contribution to -2 times the log-likelihood (large values flag outliers). Local parameter(s) estimate(s). Their approximate standard error. An estimated site index for each plot can be obtained from the local parameter using equation (4).

General info follows: Number of plots. Total number of measurements. Log-likelihood.

Then, info on the global parameters: Estimates. Approximate standard errors. Matrix of approximate correlations between them[4].

Finally, for the local parameter(s): Average of the estimates. Mean standard error. Mean correlations among local parameters. Mean correlations between local and global parameters.

Note that all results involve the specified scaling factors, if any. Back scaling can be figured-out by substituting in (5) the age and height divided by the

---

[4] As mentioned at the end of Section 2.3, sometimes one of the variance estimates turns out to be zero. Take with a pinch of salt.

respective factors. It may also be obtained from the *SitePlot* program, see below. It can be a good idea to verify through a run without scaling, starting from the final parameter values.

Before exiting, the user is prompted to optionally save the report, if not already saved.

3.6   Behind the scenes

If you must know, the front-end runs `DataPrep.exe`, which reads the input data file and generates intermediate files with the data (`DATA.tmp`, binary), the locals (`PARS.tmp`, binary), and other info (`CTRL.tmp`, text). If everything goes well, these are then passed on to `SDEfit.exe`, which does the hard work. The final local estimates and other information are left in `PARS.tmp`, and results in `REPORT.tmp` (a text file). There is a little more flexibility when running these programs manually; a list of plot-specific local initial estimates can then be used, see the advanced notes.

4   Graphing

SitePlot aids in plotting the data and site index curves. Running SitePlot brings up the window in Figure 4. It generates commands for *gnuplot*, a free graphing program (included in the Windows installation).

The first part is similar to EasySDE. The **Title** appears at the top of the graph. Not all models are supported. **Scaling** factors and scaled parameter values can be used. It is possibly to plot only data, only curves, or both.

Once the parameters are given, the site index equation can be displayed with the **Show equation** button. This gives top height as a function of age, site index, and index age (Section 2.2). The original units are used (no scaling). The equation is in gnuplot format, which is fairly standard except for using `**` for exponentiation.

On the lower part of the window, labels and ranges for the Age and Height axes can be given or changed. Any strings can be used for the labels. If a
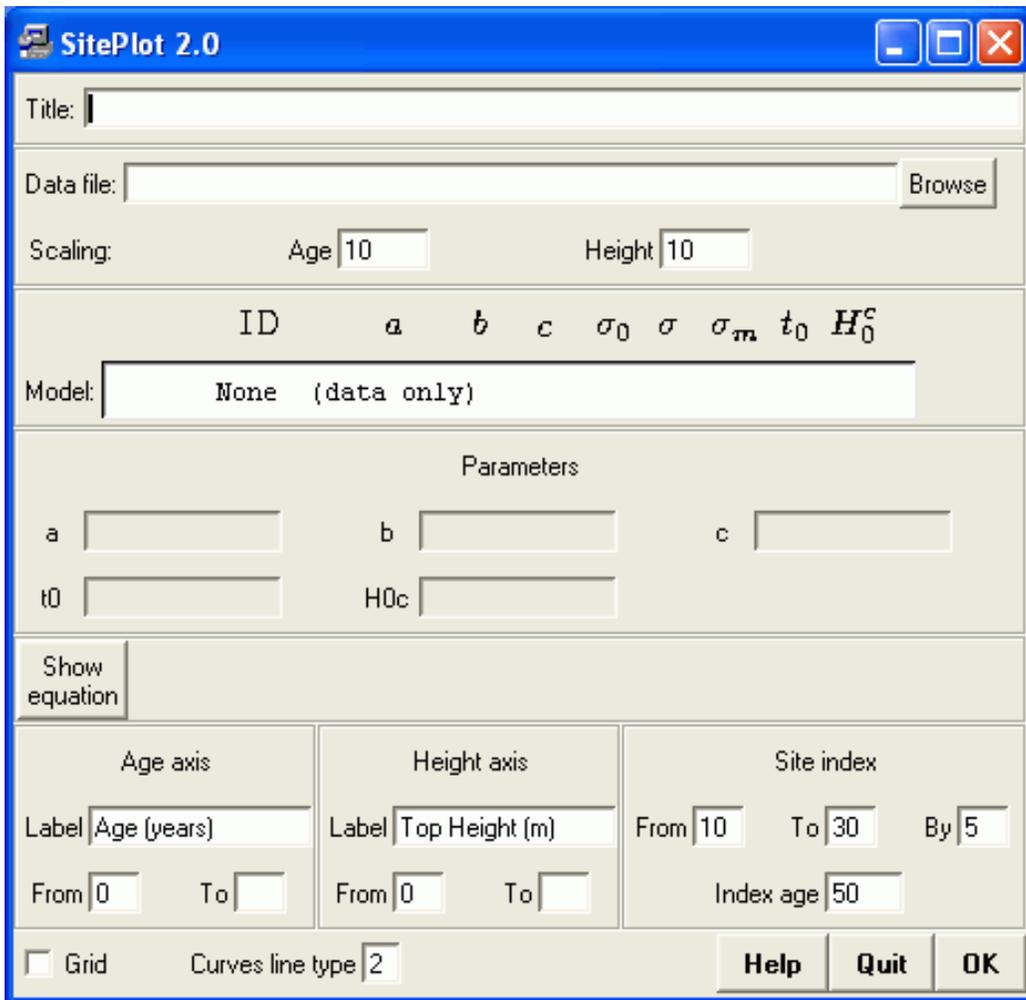
Figure 4: Graphing with SitePlot.

range limit is left blank it is calculated by gnuplot automatically from the data and/or function values. The age limits must be given when not plotting data.

For plotting **site index** curves, one needs the index age, the range from lowest to highest site index, and the site index interval between curves.

Finally, we may choose to draw a grid or not, and different line types for the curves can be selected (enter `test` in gnuplot to see the types available).

19

If everything goes well, on clicking **OK** gnuplot is executed, and the graph is displayed. The user is left in gnuplot, where the graph can be embellished, printed (on Windows, right-click on the graph), saved in various formats, etc. Enter `set terminal` and `help terminal` to see the output formats available. For instance, entering `set output 'filename.eps'`, followed by `set term post eps 20` and `replot`, produces a good quality Encapsulated PostScript file.

Gnuplot is an excellent general-purpose graphing program, well worth learning. See www.gnuplot.info. There is a good brief tutorial at www.duke.edu/~hpgavin/gnuplot.html.

The gnuplot commands are left in the file *SitePlot.plt*. This can be loaded from gnuplot, possibly after manual editing. On Windows the gnuplot executable is *wgnuplot.exe*.

References

Bailey, R. L., Cieszewski, C. J., 2000. Development of a well-behaved site-index equation: jack pine in north-central Ontario: Comment. Canadian Journal of Forest Research 30, 1667–1668.

Bailey, R. L., Clutter, J. L., 1974. Base-age invariant polymorphic site curves. Forest Science 20, 155–159.

Cieszewski, C. J., Bailey, R. L., 2000. Generalized algebraic difference approach: Theory based derivation of dynamic site equations with polymorphism and variable asymptotes. Forest Science 46 (1), 116–126.

Dyer, M. E., Bailey, R. L., 1987. A test of six methods for estimating true heights from stem analysis data. Forest Science 33, 3–13.

García, O., 1980. A stochastic differential equation model for the height growth of forest stands. Presented at the 5th Australian Statistical Conference, Sydney. http://web.unbc.ca/~garcia/unpub/oz80.pdf.

García, O., 1983. A stochastic differential equation model for the height growth of forest stands. Biometrics 39, 1059–1072.

García, O., 1996. Toward new site index curves for Douglas-fir in the Netherlands. Working paper, Royal Veterinary and Agricultural University, Unit of Forestry, Fredericksberg, Denmark, `http://web.unbc.ca/~garcia/unbub/dfsite.pdf`.

García, O., 1999. Height growth of *Pinus radiata* in New Zealand. New Zealand Journal of Forestry Science 29 (1), 131–145.

García, O., 2005. Comparing and combining stem analysis and permanent sample plot data in site index models. Forest Science 51 (4), 277–283.

Milner, K. S., 1992. Site index and height growth curves for ponderosa pine, Western larch, lodgepole pine, and Douglas-fir in Western Montana. Western Journal of Applied Forestry 7 (1), 9–14.

Nigh, G. D., 1995. Compatibility improvements and bias reduction in height-age models. Research Report 03, British Columbia Ministry of Forests, Research Branch.

Pienaar, L. V., Turnbull, K. J., 1973. The Chapman-Richards generalization of von Bertalanffy's growth model for basal area growth and yield in even-aged stands. Forest Science 19, 2–22.

Rennolls, K., 1995. Forest height growth modelling. Forest Ecology and Management 71 (3), 217–225.

Reventlow, C. D. F., 1879. A Treatise of Forestry. Society of Forest History, Horshølm, Denmark, (English translation, 1960).

Richards, F. J., 1959. A flexible growth function for empirical use. Journal of Experimental Botany 10, 290–300.

Seber, G. A. F., Wild, C. J., 2003. Nonlinear Regression. Wiley Interscience, (also Wiley 1989).

von Bertalanffy, L., 1938. A quantitative theory of organic growth (Inquiries on growth laws. II). Human Biology 10, 181–213.

von Bertalanffy, L., 1949. Problems of organic growth. Nature 163, 156–158.