COENOSES 4 (3): 175-180, 1989

SIMULATION MODELLING FOR ECOLOGICAL APPLICATIONS [1]

D.W. Goodall, C.S.I.R.O., Division of Wildlife and Ecology, Bag. No. 4 PO, Midland, W.A., Australia

Objectives

More and more, ecologists are called on to provide, or contribute to, predictions of what changes will occur in natural systems, either if left undisturbed or (more often) if they are subject to certain perturbations. For these predictions, the most appropriate way - and the most convincing for the non-ecologist - is often a simulation model.

Simulation models have often been used to contribute to theoretical ecology, but I am here concerned with severely practical applications - the effect of industrial discharges on the algal populations of a lake; population changes in endangered species of birds; the effect of different fire regimes on timber production in forests; an so forth. In such cases, the ecologist is usually trying to meet the demands of a "client" organization or group, usually decision-makers or some sort. The demands most commonly are explicit, and in meeting them the ecologist will need to be in regular consultation with the "clients". Only if the "clients" are convinced that the model incorporates the best scientific knowledge about the system and gives sound predictions will they be prepared to base their decisions upon it. It is par of the ecologist's job to convince the clients that the prediction is sound.

Various modelling approaches may be used. In some cases, existing knowledge is such that one can jump directly from the input variables to the desired predictions, by using *regression* equations in which the desired outputs are the dependent variables and the quantities defining the initial state of the system and relevant inputs are the independent variables. This approach presupposes extensive prior observations on systems analogous to that under study, for durations at least as great as that for which predictions are wanted. Rather similar is the *Markov chain* type of model, where the transition probabilities between different states of the system are known or estimated in advance. Like the regression approach, this assumes extensive prior obervations on analogous systems - as well as the assumption that the transition probabilities are constant, or at least that they change in a definable fashion. Such extensive knowledge of analogous systems is rarely available to the ecological modeller.

Other modelling approaches (often called "*mechanistic*") take account of the chains of causation leading to changes in the state of the system. The rates of change may be expressed as differential equations, including derivates of the variables with respect to time, and the predictions are then obtained by integrating the set of differential equations. This is a *differential model* - common in applied fields based on the physical sciences. In ecology, however, the number of variables is often very large, and the types of equation representing their rates of change are often rather intractable even to numerical integration. Accordingly, the approximation represented by a *difference model* is often preferred - one where each change is calculated over a discrete time-step, during which it is assumed that all variables remain constant - in other words, the time course of each variable in the model is a series of vertical rises linked by horizontal steps, approximating the continuous curve which would usually apply in reality. The following discussion will assume that a difference model is envisaged; but much of what is said will apply whatever modelling approach is adopted.

It follows from the intrinsic relationship between ecologist and client that the first step in developing a predictive simulation model should be jointly to define and clarify the objectives.

The objectives of modelling will, in the first place, consist of a set of output variables to be predicted over a specified period. An important set of decisions concern the aggregation of these output variables - to what extent different variables need to be distinguished, or whether they can be combined in broad groups, and totals or means used in place of the individual values. For instance, it may be necessary to predict the biomass of birds species by species - or even cohort by cohort; or perhaps they can be combined in "guilds" ; or perhaps changes in the total avian biomass are all that is needed. Aggregation can greatly simplify the model; but it may make it more difficult to estimate the parameters needed, and may reduce the value of the predictions. The modelling process itself may indeed require a measure of disaggregation. Even if total avian biomass is the required output, it may only be practicable to model the biomass changes species by species.

The output variables will often need to be tracked throughout the "run" of the model. This will be so if their values influence other relevant variables as they change during the "run" - in other words, if they are *state variables* of the system (state variables are those whose values at any point in time define the state of

[1] An invited lecture presented at the 2nd C.E.T.A. International Workshop on Mathematical Community Ecology, November 19-25, 1989, Gorizia, Italy.

the system at that time). Even if the output variables are not themselves state variables, their values are likely to depend on other variables changing through time which are state variables. In this case, the output variables themselves may need to be calculated only for times when predictions are required, and not to be tracked throughout the "run" of the model. The state variables, on the other hand, will need to be tracked constantly.

The set of state variables is a most important defining feature of the model, for it is they which will need to be calculated afresh for each time step, or whose rates of change are in a constant dependence one on another, then it may well be possible to aggregate them and hence reduce the complexity of the model. The set of state variables which must all be tracked separately defines a minimal level of complexity.

Apart from the set of state variables which enable the required outputs to be predicted, a second specification of the modelling task is the range of systems to be covered. It is rare for a model to be of interest only as representing a single system. Generally, it should be able to perform predictions over a range of similar systems each of which has its own special features. This range of systems then (the domain of the model, or "universe of discourse", to borrow a philosophical phrase) should also, like the output variables, be defined in consultation between the ecologist and his "client".

A third respect in which the modelling task has to be specified is the various manipulations and modifications to which the system may be subjected. Usually the "clients" are contemplating some form of restraint to be imposed on the system, or some changes whose impact is to be predicted. The range of such modifications should if possible be spelled out in advance, for they may well affect the structure of the model.

Other constraints on the modelling process include the time scale; the precision needed for the predictions; the availability of data, or the possibility of acquiring them. These various constraints may well interact for instance, the longer the time scale modelled, the more difficult it becomes to foresee what manipulations may be desirable or practicable, and the more uncertain are the appropriate weightings for any objective function that may be proposed.

The problem of uncertainty in a model should be faced squarely, and not side-stepped. All model results are uncertain, if only because the values of parameters involved are known only within a penumbra of uncertainty. Provided this uncertainty is clearly stated, and if possible quantified, its existence will not normally worry the "clients"; they will be used to making decisions where a number of relevant factors are uncertain - if the weather, the economic future, natural catastrophes can be taken into account in some way, then likewise uncertainties in ecological prediction can be taken into account. This will be greatly facilitated, however, if the uncertainties can be brought into the oper and quantified.

The objective function

Model development can be greatly facilitated if the "clients" are willing to agree on an objective function. This is a function which embodies all the output variables of interest, and prediction of which constitutes the objective of the modelling activity. Most commonly, the objective function will be a weighted sum of the various output variables; but any single-value function of them will serve.

It should be noted that the objective function should cover not only all the output variables of interest, but also the whole domain over which the model is to be applied. Again, the various parts of the domain may need to be weighted differently - a process which is the responsibility of the "clients". Maybe, for instance, they would like the commonest type of systems to be given most weight, or those types of manipulations which at present seem most likely. These relative weightings will be highly relevant at a later stage of model development, as will be shown below, and may even be taken into account in the early stages. For instance, if the use of fire as a management tool seems unlikely, modelling the effects of fire may be postponed until the rest of the model has been adequately tested.

Even if the "clients" are unwilling to define an objective function, they should at least be urged to give weightings to the different variables predicted, and to the different regions of the model domain. Otherwise useful sensitivity testing (see below) may become impossible.

Model construction

Once the modelling task has been defined, building of the model may begin. First, the state variables are considered one by one, and for each a function expressing its rate of change - or its value at some future time - is defined. This function will include a number of other variables. In so far as these are within the list of state variables already defined, or are quantities to be defined outside the system and not subject to change as a result of system processes, the complexity of the model will not be increased above that already required by the state variables. If other variables subject to change within the system are included in the list of those already enumerated.

The function expressing a rate of change in a state variable will include one or more parameters - constant for the system, or changing in some predefined way irrespective of what happens within the system. These parameters will need to be estimated, and this may prove a severe limitation on the modelling process. The estimate may be derived, for instance, from a regression in which the dependent variables and measured rates of change in similar systems recorded over an interval. Or it may be derived from direct ad hoc experimentation on appropriately simplified subsystems. At the worst, estimates may be used on informed guesswork.

It will be seen that simplicity in the model is highly desirable - in limiting both the number of state variables to be modelled and the number of parameters to be estimated. If it is possible to develop two independent models, both of which predict the same variables with similar precision, then the simpler model is to be preferred. But the comparison would need to be performed not for a single system only, but for a sample representing the whole range of systems to be covered - a matter which will be discussed later.

If the "clients" can specify beforehand the precision needed for the model, it may be possible initially to choose a simple type of model which may be expected to meet this requirement - especially if this type has already been used elsewhere with known results. In general, however, it is better to proceed from a more complex model, and to simplify it progressively as some parts of complexities are shown in the course of sensitivity testing (see below) to have little influence on the outcome.

It is desirable to involve the "clients", as far as possible, in all stages of model-building. If they have agreed that the way in which the rates and interdependence of the processes are modelled is reasonable, and in accord with current knowledge, and that simplifications introduced do not go too far, then they will *ipso facto* share responsibility for subsequent stages. It is accordingly desirable that work on the model should be presented stage to the "clients" in forms which they can understand and either approve on criticise. Any implied assumptions should be made explicit, so that they can be explicitly accepted, or rejected and replaced.

Understanding of the model will be facilitated if functional relations can be expressed in the form of diagrams, as well as the formal quasi-mathematical expressions needed for computer coding. Tables may also be a valuable means of communication between ecologist and "client"; they are almost as readily comprehensible to the "client" as diagrams, and can even be used directly in the model, by categorization or interpolation, if no simple mathematical expression seems appropriate.

When the structure of the model has been decided and agreed in this way, computer coding can start. Some languages intended specifically for simulation are available, and can often be used for ecological models, depending on the way they have been constructed. But any general-purpose programming language can serve the same purpose.

Meanwhile, parameter estimation should begin. As indicated, this can take various forms - regression; *ad hoc* experimentation; or even informed guesses. The last should not be disdained. "Clients" with practical experience of the systems in question are likely to have semi-quantitative knowledge of the processes, often only semi-conscious, which can provide a valuable starting-point if no formal estimate of the parameter is available.

Any parameter estimate should be associated with a measure of uncertainty. In principle, this should extend over the whole universe of discourse - the range of systems and manipulations to be covered. Depending on the structure of the model, different parameter estimates may be needed for these different situations, the estimate of uncertainty should cover this range. The reproduction rate of a particular animal species, for instance, may be known with high precision for a particular forest type, and this precision would be the appropriate figure if use of the model were confined to this forest type. But the parameter is also known to vary in other types of vegetation. If a common estimate is to be used in the model for all vegetation types, it is the latter variation which is relevant. This is analogous to the situation in agricultural field trials, where response estimates and their errors will apply to a particular trial only; if an estimate of the variation of response over a landscape is needed, a series of trials must be set up, and the variation of responses between trials is the relevant figure. The same applies, *mutatis mutandis*, to a model.

Uncertainty measures should also be sought for estimates based on informed guesses. To obtain these will generally be no more difficult than to obtain the guesses themselves - the "client" can usually just as readily indicate the range within which a parameter lies as a particular most likely value.

The treatment of uncertainty in a model should be distinguished from stochastic modelling. There are some types of model where the sthochastic element may be all-important. Where the model's purpose is to predict the probability of extinction of a particular species, this will often apply. Models for this purpose might indeed predict the probability distribution of population sizes, the parameters of the distribution being absolute values, though the outcome would itself be stochastic. More commonly, the model may include a number of different processes, each individually stochastic - birth, death, predation, infection, and so forth; and the model outcome depends on the combination of these events to give a result which is not analytically predictable. In such cases, the choice may be that the model should treat each of these processes as random, and develop a probability distribution for the outcome as an empirical result of repeated "runs" of the model, using different "seeds" for a pseudo-random generator. This type of stochastic model is usually appropriate only where the number of individual elements (individuals) tracked separately through time is quite small.

Another type of model where a stochastic element is clearly relevant is one in which the output is dependent on seasonal weather conditions, and the "clients" wish to know the probability distribution of different output values, given existing climatic uncertainty. A deterministical model may be build, and then "run" with different weather sequences derived as random samples from existing weather records. The results can then take the form of a joint probability distribution for the various output variables, in any particular part of the model domain.

Sensitivity testing

Once the model is coded and running, the measures of uncertainty mentioned above will form a basis for sensitivity testing. For this, a sample of systems and scenarios should be set up, covering the domain to which the simulation is to apply. Simulations will then be performed for each of the situations in the sample, with parameter values varied one by one to the upper and lower limits of uncertainty, while the rest of the parameters are held at the most likely value. If a number of output variables are to be predicted, they will show varying sensitivity to modifications in the values of the different parameters. Sensitivity of each of the output variables to each of the parameters will also vary within the universe of discourse. This is where the weightings referred to earlier become important.

If the different output variables, systems and manipulations have been combined into a single objective function, with the arbitraty weights allotted, the sensitivity of the objective function to variation in the different parameter estimates will show the overall importance which should be attached to them. If the uncertainty in a particular parameter has only a negligible effect on the objective function, then improvement in that parameter estimate is unimportant. If variation in another parameter, within the bounds set for its uncertainty, has a substantial effect on values of the objective function (because variables in which this parameter plays an important part, or systems or manipulations in which it is highly influential, have been given heavy weights), then clearly one should seek methods for improving the estimate of this parameter.

At this point, simplification of the model can be considered. If, for instance, the output variables prove to be highly insensitive to a particular parameter, then it may be possible to change the formulation of that part of the model so that this unimportant parameter is eliminated - a quadratic expression converted into a linear one, for instance. If a state variable is not directly involved in the calculation of the output variables, it may be possible to eliminate it, and all the processes specifically related to it, or to treat it as constant. This possibility may be tested by sensitivity "runs" of the model with and without these changes. If the effects, over the whole universe of discourse, are negligible, the simplification may be adopted without hesitation. Similarly, it may be possible to remove a whole module if its relevance to changes in the output variables proves to be minimal. Another type of simplification which can be tested in this way is aggregation. If disaggregated variables are not specifically required as output, their replacement by one or a few aggregated variables (guilds or life-forms in place of species, for instance; or combining different soil types) may result in a very welcome reduction in the variables themselves, in the parameters to be estimated, and in the complexity of the model generally.

Validation

As has been said, the "clients" expect the model to provide best estimates for an objective function, or for a set of specified variables, under the various conditions of interest, including often one or more proposed manipulations of the system. But they would also like estimates of the variance associated with each of these estimates - particularly the estimates of changes resulting from manipulations. These pose a special challenge to the modelling process. If the fiducial distribution of each parameter estimate were known, a Monte Carlo process could be used to estimate the variances required. But this situation would be very rare. A partial approach would make use of the uncertainty quantified for each parameter estimate, whether the uncertainty is formally calculated or merely guesswork. But even if this *a priori* estimation were possible, a series of validation tests in each of which the model results could be compared with corresponding real-world observations would be more satisfactory. These would have the advantage of taking into account elements of the system which had not been included in the simulation.

What is required for validation is observations on systems within the universe of discourse in which initial values of all state variables are known, together with any inputs needed by the model. The values of the variables required as model output are observed after a time interval within the range for which predictions are required. The differences between the observed output variables and those predicted by the model using the same inputs then constitute the best possible test of the validity of the model. Clearly, each such validation will apply only to a particular point within the universe of discourse. A series of such tests covering the whole universe of discourse will be required to give confidence in applying the model to a new situation -which will generally be intermediate between the situation included in the validation tests.

The precision needed in a model can sometimes be specified in advance, but the precision attainable can rarely be stated before the model is built. If an existing model is to be applied to a new situation, previous validation tests will provide a good indication, so long as the tests have covered that part of the domain within which the new situation is located - that, for instance, one is not extending the use of the model to a new geographical region or new types of manipulation. Otherwise, one can only guess whether the degree of agreement between model and real-life results already found can be expected in the new situation.

For a new model not yet validated, all that can be said about its precision is that it cannot be greater than that imposed by the uncertainty of the parameter estimates. If Monte Carlo tests of the effects of varying parameter values show that the resulting values of the objective function may differ by, say, 20%, then one may be reasonably confident that validation tests might deviate by at least this much from the modelling results.

Validation tests are often used in a hit-and-miss process of model improvement. Discrepancies between observed and modelled values are noted, these discrepancies are ascribed - largely intuitively - to a particular module, process or parameters of the model, and special efforts are directed to correcting these perceived sources of discrepancy. But this procedure of using validation results to guide model improvement is of doubtful value. Usually many parts of the model are involved in the calculation of each change, and intuition is a poor guide to their respective contributions, unless supported by *ad hoc* sensitivity tests. Where a parameter estimate has been based on guesswork, attempts are sometimes made to improve it later by trial and error during the validation phase, small adjustments being made until the real-life and model results are sufficiently close; but this is practicable only if the number of parameter estimates to be improved in this way is very small, and can never really take into account a series of validation tests covering the whole model domain.

Often a model its intended to perform predictions within a long time-span - longer than is practicable for field trials and observations. In such a case, one may have to content oneself with partial validations over much shorter terms. Since the behaviour of a system is likely to change considerably during the course of development - a forest proceeding from establishment to harvest, or a lake from oligotrophy to eutrophy - such partial validations would need to be applied to systems at different stages of development in order that model predictions over the whole time-span should inspire confidence.

In any model intended for repeated use in different circumstances, model-building and validation should be continuing processes. Any observations on a system within the domain modelled can serve for validation, provided all the input needs are met - initial records of the value of each state variable, all inputs needed during the "run" of the model, and the required output variables at some later date. One may "run" the model with the same initial values and inputs, performing a Monte Carlo test with parameter values sampled from the zone of uncertainty associated with each. If the observed values of the output variables fall within the penumbra of uncertainty of the repeated model predictions, then the validation is satisfactory. If not, the discrepancy should be investigated - which may require some extra observations on the real-life system. If the system is being observed over an extended period, and "runs" of the model during the early part of this period have revealed these discrepancies, the extra observations may be possible for the latter part of the real-life test, so that model improvement can proceed concurrently with it.

How should validation results be used? Should they constitute a formal test of the accuracy of the predictions made by the model? Clearly, exact agreement between observed results and model predictions is not to be expected, for all the parameter estimates used in the model are surrounded by uncertainty - apart from known or unknown discrepancies between the structure of the model and that of the system. If there are enough validation tests, then the difference between the observed and modelled values for the objective function could in principle be used as a dependent variable, and a regression of this on a set of independent variables describing the initial state of the system and its location within the domain of the model could be used to correct any new "run" of the model. But once such an extensive set of field results is available, the choice may well be to use regression equations rather than a mechanistic model. It is at a much earlier stage, when complete validation tests are few or perhaps non-existent, that the model is valuable. Thus it should probably be recognized that the value of successful validation tests - "successful" in the sense of supporting the model rather than casting doubt upon it - is rather for building confidence in the model than for incorporation in any formal calculations. A model is worth building in situations, or for periods, where the field trials which it is replacing (and which would be used for validation) are impracticable in the number and on the scale (in space and time) which are indicated. Modelling is a *pis aller*; but ecology (unlike some other sciences) very often has to deal with large scales of space and time, and highly variable subject-matter. These factors greatly limit the use of experimental approaches in ecology, so that the *pis aller* of modelling can be a most valuable substitute.

Manuscript received: March 1989

The precision needed in a model can sometimes be specified in advance, but the precision attainable can rarely be stated before the model is built. If an existing model is to be applied to a new situation, previous validation tests will provide a good indication, so long as the tests have covered that part of the domain within which the new situation is located - that, for instance, one is not extending the use of the model to a new geographical region or new types of manipulation. Otherwise, one can only guess whether the degree of agreement between model and real-life results already found can be expected in the new situation.

For a new model not yet validated, all that can be said about its precision is that it cannot be greater than that imposed by the uncertainty of the parameter estimates. If Monte Carlo tests of the effects of varying parameter values show that the resulting values of the objective function may differ by, say, 20%, then one may be reasonably confident that validation tests might deviate by at least this much from the modelling results.

Validation tests are often used in a hit-and-miss process of model improvement. Discrepancies between observed and modelled values are noted, these discrepancies are ascribed - largely intuitively - to a particular module, process or parameters of the model, and special efforts are directed to correcting these perceived sources of discrepancy. But this procedure of using validation results to guide model improvement is of doubtful value. Usually manu parts of the model are involved in the calculation of each change, and intuition is a poor guide to their respective contributions, unless supported by *ad hoc* sensitivity tests. Where a parameter estimate has been based on guesswork, attempts are sometimes made to improve it later by trial and error during the validation phase, small adjustments being made until the real-life and model results are sufficiently close; but this is practicable only if the number of parameter estimates to be improved in this way is very small, and can never really take into account a series of validation tests covering the whole model domain.

Often a model its intended to perform predictions within a long time-span - longer than is practicable for field trials and observations. In such a case, one may have to content oneself with partial validations over much shorter terms. Since the behaviour of a system is likely to change considerably during the course of development - a forest proceeding from establishment to harvest, or a lake from oligotrophy to eutrophy - such partial validations would need to be applied to systems at different stages of development in order that model predictions over the whole time-span should inspire confidence.

In any model intended for repeated use in different circumstances, model-building and validation should be continuing processes. Any observations on a system within the domain modelled can serve for validation, provided all the input needs are met - initial records of the value of each state variable, all inputs needed during the "run" of the model, and the required output variables at some later date. One may "run" the model with the same initial values and inputs, performing a Monte Carlo test with parameter values sampled from the zone of uncertainty associated with each. If the observed values of the output variables fall within the penumbra of uncertainty of the repeated model predictions, then the validation is satisfactory. If not, the discrepancy should be investigated - which may require some extra observations on the real-life system. If the system is being observed over an extended period, and "runs" of the model during the early part of this period have revealed these discrepancies, the extra observations may be possible for the latter part of the real-life test, so that model improvement can proceed concurrently with it.

How should validation results be used? Should they constitute a formal test of the accuracy of the predictions made by the model? Clearly, exact agreement between observed results and model predictions is not to be expected, for all the parameter estimates used in the model are surrounded by uncertainty - apart from known or unknown discrepancies between the structure of the model and that of the system. If there are enough validation tests, then the difference between the observed and modelled values for the objective function could in principle be used a dependent variable, and a regression of this on a set of independent variables describing the initial state of the system and its location within the domain of the model could be used to correct any new "run" of the model. But once such an extensive set of field results is available, the choice may well be to use regression equations rather than a mechanistic model. It is at a much earlier stage, when complete validation tests are few or perhaps nonexistent, that the model is valuable. Thus it should probably be recognized that the value of successful validation tests - "successful" in the sense of supporting the model rather than casting doubt upon it - is rather for building confidence in the model than for incorporation in any formal calculations. A model is worth building in situations, or for periods, where the field trials which it is replacing (and which would be used for validation) are impracticable in the number and on the scale (in space and time) which are indicated. Modelling is a *pis aller*; but ecology (unlike some other sciences) very often has to deal with large scales of space and time, and highly variable subject-matter. These factors greatly limit the use of experimental approaches in ecology, so that the *pis aller* of modelling can be a most valuable substitute.

Manuscript received: March 1989