

Aster Models for Life History Analysis

BY CHARLES J. GEYER

*School of Statistics, University of Minnesota, 313 Ford Hall, 224 Church Street S. E.,
Minneapolis, Minnesota, 55455, U. S. A.*

charlie@stat.umn.edu

STUART WAGENIUS

*Institute for Plant Conservation Biology, Chicago Botanic Garden, 1000 Lake Cook Road,
Glencoe, Illinois, 60022, U. S. A.*

wagenius@chicagobotanic.org

AND RUTH G. SHAW

*Department of Ecology, Evolution, and Behavior, University of Minnesota, 100 Ecology
Building, 1987 Upper Buford Circle, St. Paul, Minnesota, 55108, U. S. A.*

rshaw@superb.ecology.umn.edu

Abstract

We present a new class of statistical models, designed for life history analysis of plants and animals, that allow joint analysis of data on survival and reproduction over multiple years, allow for variables having different probability distributions, and correctly account for the dependence of variables on earlier variables. We illustrate their utility with an analysis of data taken from an experimental study of *Echinacea angustifolia* sampled from remnant prairie populations in western Minnesota. These models generalize both generalized linear models and survival analysis. The joint distribution is factored as a product of conditional

distributions, each an exponential family with the conditioning variable being the sample size of the conditional distribution. The model may be heterogeneous, each conditional distribution being a different exponential family. We show that the joint distribution is a flat exponential family and derive its canonical parameters, Fisher information, and other properties. These models are implemented in an R package ‘aster’ available from CRAN.

Some key words: Conditional Exponential Family; Flat Exponential Family; Generalized Linear Model; Graphical Model; Maximum Likelihood.

1. INTRODUCTION

This article introduces a class of statistical models we call *aster models*. They were invented for life history analysis of plants and animals and are best introduced by an example about perennial plants observed over several years. For each individual planted, at each census, we record whether it is alive, whether it flowers, and its number of flower heads. These data are complicated, especially when recorded for several years, but simple conditional models may suffice. We consider mortality status (dead or alive) to be Bernoulli given the preceding mortality status. Similarly, flowering status given mortality status is also Bernoulli. Given flowering, the number of flower heads may have a zero-truncated Poisson distribution (Martin, et al., 2005). Figure 1 shows the graphical model for a single individual.

[Figure 1 about here.]

This aster model generalizes both discrete time Cox regression (Cox, 1972; Breslow, 1972, 1974) and generalized linear models (McCullagh & Nelder, 1989). Aster models apply to any similar conditional modelling. We could, for example, add other variables, such as seed count modelled conditional on flower head count.

A simultaneous analysis that models the joint distribution of all the variables in a life history analysis can answer questions that cannot be addressed through separate analyses of each variable conditional on the others. Joint analysis also deals with structural zeros in the data (e. g., a dead individual remains dead and cannot flower, so in Figure 1 any arrow that leads from a variable that is zero to another variable implies that other variable must also be zero), which present intractable missing data problems in separate analyses of individual variables. Aster models have no problem with structural zeros; likelihood inference automatically handles them correctly.

Aster models are simple graphical models (Lauritzen, 1996, Section 3.2.3) in which the joint density is a product of conditionals as in equation (1) below. No knowledge of graphical model theory is needed to understand aster models. One innovative aspect of aster models is the interplay between two parameterizations described in Sections 1.2 and 1.3 below. The *conditional canonical parameterization* arises when each conditional distribution in the product is an exponential family and we use the canonical parameterization for each. The *unconditional canonical parameterization* arises from observing that the joint model is a full flat exponential family (Barndorff-Nielsen, 1978, Chapter 8) and using the canonical parameters for that family, defined by equation (5) below.

1.1. Factorization and Graphical Model

Variables in an aster model are denoted X_j where j runs over nodes of a graph. A general aster model is a chain graph model (Lauritzen, 1996, pp. 7 and 53) having both *arrows* (directed edges) and *lines* (undirected edges). Figure 1 is special, having only arrows. Arrows go from *parent* to *child*, lines between *neighbors*. Nodes that are not children are *root* nodes. Those that are not parents are *terminal* nodes.

Let F and J denote root and non-root nodes. Aster models have very special chain graph structure determined by a partition \mathcal{G} of J and a function $p : \mathcal{G} \rightarrow J \cup F$. For each

$G \in \mathcal{G}$ there is an arrow from $p(G)$ to each element of G and a line between each pair of elements of G . For any set S , let \mathbf{X}_S denote the vector whose components are X_j , $j \in S$. The graph determines a factorization

$$\Pr(\mathbf{X}_J|\mathbf{X}_F) = \prod_{G \in \mathcal{G}} \Pr\{\mathbf{X}_G|X_{p(G)}\} \quad (1)$$

(compare equation 3.23 in Lauritzen, 1996).

Elements of \mathcal{G} are called *chain components* because they are connectivity components of the chain graph (Lauritzen, 1996, pp. 6–7). Figure 1 having no undirected edges, each node is a chain component by itself. Allowing nontrivial chain components allows the elements of \mathbf{X}_G to be conditionally dependent given $X_{p(G)}$ with merely notational changes to the theory. In our example (Section 4) the graph consists of many copies of Figure 1, one for each individual plant. Individuals have no explicit representation. For any set S , let $p^{-1}(S)$ denote the set of G such that $p(G) \in S$. Then each subgraph consisting of one $G \in p^{-1}(F)$, its descendants (children, children of children, etc.), and arrows and lines connecting them, corresponds to one individual. If we make each such G have a distinct root element $p(G)$, then the set of descendants of each root node corresponds to one individual. Although all individuals in our example have the same subgraph, this is not required.

1.2. Conditional Exponential Families

We take each of the conditional distributions in (1) to be an exponential family having canonical statistic \mathbf{X}_G that is the sum of $X_{p(G)}$ independent and identically distributed random vectors, possibly a different such family for each G . Conditionally, $X_{p(G)} = 0$ implies $\mathbf{X}_G = \mathbf{0}_G$ almost surely. If $j \neq p(G)$ for any G , then the values of X_j are unrestricted. If the distribution of \mathbf{X}_G given $X_{p(G)}$ is infinitely divisible, e. g., Poisson or normal, for each $G \in p^{-1}(\{j\})$, then X_j must be nonnegative-real-valued. Otherwise, X_j must be nonnegative-integer-valued.

The log likelihood for the whole family has the form

$$\sum_{G \in \mathcal{G}} \left[\sum_{j \in G} X_j \theta_j - X_{p(G)} \psi_G(\boldsymbol{\theta}_G) \right] = \sum_{j \in J} X_j \theta_j - \sum_{G \in \mathcal{G}} X_{p(G)} \psi_G(\boldsymbol{\theta}_G), \quad (2)$$

where $\boldsymbol{\theta}_G$ is the canonical parameter vector for the G -th conditional family, having components θ_j , $j \in G$, and ψ_G is the *cumulant function* for that family (Barndorff-Nielsen, 1978, pp. 105, 139, and 150) that satisfies

$$E_{\boldsymbol{\theta}_G} \{\mathbf{X}_G | X_{p(G)}\} = X_{p(G)} \nabla \psi_G(\boldsymbol{\theta}_G) \quad (3)$$

$$\text{var}_{\boldsymbol{\theta}_G} \{\mathbf{X}_G | X_{p(G)}\} = X_{p(G)} \nabla^2 \psi_G(\boldsymbol{\theta}_G) \quad (4)$$

where $\text{var}_{\boldsymbol{\theta}}(\mathbf{X})$ denotes the variance-covariance matrix of \mathbf{X} and $\nabla^2 \psi(\boldsymbol{\theta})$ the matrix of second partial derivatives of ψ (Barndorff-Nielsen, 1978, p. 150).

1.3. Unconditional Exponential Families

Collecting terms with the same X_j in (2), we get

$$\sum_{j \in J} X_j \left[\theta_j - \sum_{G \in p^{-1}(\{j\})} \psi_G(\boldsymbol{\theta}_G) \right] - \sum_{G \in p^{-1}(F)} X_{p(G)} \psi_G(\boldsymbol{\theta}_G)$$

and see that

$$\varphi_j = \theta_j - \sum_{G \in p^{-1}(\{j\})} \psi_G(\boldsymbol{\theta}_G), \quad j \in J \quad (5)$$

are the canonical parameters of an unconditional exponential family with canonical statistics X_j . We now write \mathbf{X} instead of \mathbf{X}_J , $\boldsymbol{\varphi}$ instead of $\boldsymbol{\varphi}_J$, and so forth, and let $\langle \mathbf{X}, \boldsymbol{\varphi} \rangle$ denote the inner product $\sum_j X_j \varphi_j$. Then we can write the log likelihood of this unconditional family as

$$l(\boldsymbol{\varphi}) = \langle \mathbf{X}, \boldsymbol{\varphi} \rangle - \psi(\boldsymbol{\varphi}) \quad (6)$$

where the cumulant function of this family is

$$\psi(\boldsymbol{\varphi}) = \sum_{G \in p^{-1}(F)} X_{p(G)} \psi_G(\boldsymbol{\theta}_G). \quad (7)$$

All of the $X_{p(G)}$ in (7) are at root nodes, hence nonrandom, so ψ is a deterministic function. Also the right hand side of (7) is a function of φ by the logic of exponential families (Barndorff-Nielsen, 1978, pp. 105 ff.).

The system of equations (5) can be solved for the θ_j in terms of the φ_j in one pass through the equations in any order that finds θ_j for children before parents. Thus (5) determines an invertible change of parameter.

1.4. Canonical Linear Models

One of the desirable aspects of exponential family canonical parameter linear models defined by reparameterization of the form

$$\varphi = \mathbf{M}\boldsymbol{\beta}, \tag{8}$$

where \mathbf{M} is a known matrix (the *model matrix*), is that because $\langle \mathbf{X}, \mathbf{M}\boldsymbol{\beta} \rangle = \langle \mathbf{M}^T \mathbf{X}, \boldsymbol{\beta} \rangle$ the result is a new exponential family with canonical statistic

$$\mathbf{Y} = \mathbf{M}^T \mathbf{X} \tag{9}$$

and canonical parameter $\boldsymbol{\beta}$. The dimension of this new family will be the dimension of $\boldsymbol{\beta}$, if \mathbf{M} has full rank.

As is well known (Barndorff-Nielsen, 1978, p. 111), the canonical statistic of an exponential family is minimal sufficient. Since we have both conditional and unconditional families in play, we stress that this well-known result is about *unconditional* families. A dimension reduction to a low-dimensional sufficient statistic like (9) does not occur when the conditional canonical parameters $\boldsymbol{\theta}$ are modelled linearly, and this suggests that linear models for the unconditional parameterization may be scientifically more interesting despite their more complicated structure.

1.5. Mean Value Parameters

Conditional mean values

$$\xi_j = E_{\boldsymbol{\theta}_G}\{X_j|X_{p(G)}\} = X_{p(G)} \frac{\partial \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_j}, \quad j \in G \quad (10)$$

are not parameters because they contain random data $X_{p(G)}$, although they play the role of mean value parameters when we condition on $X_{p(G)}$ treating it as constant. By standard exponential family theory (Barndorff-Nielsen, 1978, p. 121), $\nabla \psi_G : \boldsymbol{\theta}_G \mapsto \boldsymbol{\xi}_G$ is an invertible change of parameter.

Unconditional mean value parameters are the unconditional expectations

$$\boldsymbol{\tau} = E_{\boldsymbol{\varphi}}\{\mathbf{X}\} = \nabla \psi(\boldsymbol{\varphi}). \quad (11)$$

By standard theory, $\nabla \psi : \boldsymbol{\varphi} \mapsto \boldsymbol{\tau}$ is an invertible change of parameter. The unconditional expectation in (11) can be calculated using the iterated expectation theorem

$$E_{\boldsymbol{\varphi}}\{X_j\} = E_{\boldsymbol{\varphi}}\{X_{p(G)}\} \frac{\partial \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_j}, \quad j \in G \quad (12)$$

where $\boldsymbol{\theta}$ is determined from $\boldsymbol{\varphi}$ by solving (5). The system of equations (12) can produce the τ_j in one pass through the equations in any order that finds τ_j for parents before children.

2. LIKELIHOOD INFERENCE

2.1. Conditional Models

The score $\nabla l(\boldsymbol{\theta})$ for conditional canonical parameters is particularly simple, having components

$$\frac{\partial l(\boldsymbol{\theta})}{\partial \theta_j} = X_j - X_{p(G)} \frac{\psi_G(\boldsymbol{\theta}_G)}{\partial \theta_j}, \quad j \in G \quad (13)$$

and, if these parameters are modelled linearly as in (8) but with $\boldsymbol{\varphi}$ replaced by $\boldsymbol{\theta}$, then

$$\nabla l(\boldsymbol{\beta}) = \nabla l(\boldsymbol{\theta})\mathbf{M}. \quad (14)$$

Observed Fisher information for $\boldsymbol{\theta}$, which is the matrix $-\nabla^2 l(\boldsymbol{\theta})$, is block diagonal with

$$-\frac{\partial^2 l(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} = X_{p(G)} \frac{\partial^2 \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_i \partial \theta_j}, \quad i, j \in G \quad (15)$$

the only nonzero entries. Expected Fisher information for $\boldsymbol{\theta}$ is the unconditional expectation of observed Fisher information, calculated using (15) and (12). If $I(\boldsymbol{\theta})$ denotes either observed or expected Fisher information for $\boldsymbol{\theta}$ and similarly for other parameters, then

$$I(\boldsymbol{\beta}) = \mathbf{M}^T I(\boldsymbol{\theta}) \mathbf{M}. \quad (16)$$

2.2. Unconditional Models

The score $\nabla l(\boldsymbol{\varphi})$ for unconditional canonical parameters is, as in every (unconditional) exponential family, ‘observed minus expected’

$$\frac{\partial l(\boldsymbol{\varphi})}{\partial \varphi_j} = X_j - E_{\boldsymbol{\varphi}}\{X_j\}$$

the unconditional expectation on the right hand side being evaluated by using (12). If these parameters are modelled linearly as in (8), then we have (14) with $\boldsymbol{\theta}$ replaced by $\boldsymbol{\varphi}$. Note that (13) is not ‘observed minus conditionally expected’ if considered a vector equation because the conditioning would differ amongst components.

Second derivatives with respect to (unconditional) canonical parameters of an exponential family are nonrandom, hence observed and expected Fisher information $I(\boldsymbol{\varphi})$ are equal, given by either of the expressions $\nabla^2 \psi(\boldsymbol{\varphi})$ or $\text{var}_{\boldsymbol{\varphi}}(\mathbf{X})$. Fix $\boldsymbol{\theta}$ and $\boldsymbol{\varphi}$ related by (5). For $i, j \in G$, the iterated covariance formula gives

$$\text{cov}_{\boldsymbol{\varphi}}\{X_i, X_j\} = \frac{\partial^2 \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_i \partial \theta_j} E_{\boldsymbol{\varphi}}\{X_{p(G)}\} + \frac{\partial \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_i} \frac{\partial \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_j} \text{var}_{\boldsymbol{\varphi}}\{X_{p(G)}\}. \quad (17)$$

Otherwise we may assume $j \in G$ and i not a descendant of j so $\text{cov}_{\boldsymbol{\varphi}}\{X_i, X_j | X_{p(G)}\} = 0$ because X_j is conditionally independent given $X_{p(G)}$ of all variables except itself and its descendants. Then the iterated covariance formula gives

$$\text{cov}_{\boldsymbol{\varphi}}\{X_i, X_j\} = \frac{\partial \psi_G(\boldsymbol{\theta}_G)}{\partial \theta_j} \text{cov}_{\boldsymbol{\varphi}}\{X_i, X_{p(G)}\}. \quad (18)$$

Expectations having been calculated using (12), variances, the case $i = j$ in (17), are calculated in one pass through (17) in any order that does parents before children. Then another pass using (17) and (18) gives covariances. The information for β is given by (16) with θ replaced by φ .

2.3. Prediction

By ‘prediction’, we only mean evaluation of a function of estimated parameters, what the `predict` function in R does for generalized linear models. In aster models we have five different parameterizations of interest (β , θ , φ , ξ , and τ). Fisher information for β , already described, handles predictions of β . So this section is about ‘predicting’ the remaining four.

One often predicts for new individuals having different covariate values from the observed individuals. Then the model matrix $\widetilde{\mathbf{M}}$ used for the prediction is different from that used for calculating $\hat{\beta}$ and Fisher information $I(\hat{\beta})$, either observed or expected.

Let η be the linear predictor ($\eta = \theta$ for conditional models and $\eta = \varphi$ for unconditional models), ζ any one of θ , φ , ξ , and τ , and $f : \eta \mapsto \zeta$. Suppose we wish to predict

$$g(\beta) = h(\zeta) = h(f(\widetilde{\mathbf{M}}\beta)). \quad (19)$$

Then by the chain rule (19) has derivative

$$\nabla g(\beta) = \nabla h(\zeta) \cdot \nabla f(\eta) \cdot \widetilde{\mathbf{M}}, \quad (20)$$

and, by the ‘usual’ asymptotics of maximum likelihood and the delta method, the asymptotic distribution of the prediction $h(\hat{\zeta}) = g(\hat{\beta})$ is

$$\text{Normal}(g(\beta), [\nabla g(\hat{\beta})]I(\hat{\beta})^{-1}[\nabla g(\hat{\beta})]^T),$$

where $\nabla g(\hat{\beta})$ is given by (20) with $\hat{\eta} = \widetilde{\mathbf{M}}\hat{\beta}$ plugged in for η and $\hat{\zeta} = f(\hat{\eta})$ plugged in for ζ .

We write ‘predictions’ in this complicated form to separate the parts of the specification, the

functions h and ∇h and the model matrix $\widetilde{\mathbf{M}}$, that change from application to application from the part ∇f that does not change and can be done by computer (see Appendix A of the technical report at <http://www.stat.umn.edu/geyer/aster> for details).

To predict mean value parameters one must specify new ‘response’ data X_j as well as new ‘covariate’ data in $\widetilde{\mathbf{M}}$. Unconditional mean value parameters $\boldsymbol{\tau}$ depend on X_j , $j \in F$, whereas conditional mean value parameters $\boldsymbol{\xi}$ depend on X_j , $j \in J \cup F$. It is often interesting, however, to predict $\boldsymbol{\xi}$ for hypothetical individuals with $X_j = 1$ for all j , thus obtaining conditional mean values per unit sample size.

3. SOFTWARE

We have released an R (R Development Core Team, 2004) package `aster` that fits, tests, predicts, and simulates aster models. It uses the R formula mini-language, originally developed for GENSTAT and S (Wilkinson & Rogers, 1973; Chambers & Hastie, 1992) so that model fitting is much like that for linear or generalized linear models. Similarly, R functions `summary.aster`, `anova.aster`, and `predict.aster` provide regression coefficients with standard errors, z statistics, and p -values, likelihood ratio tests for model comparison, and the predictions with standard errors described in Section 2.3. The R function `raster` generates random aster model data for simulation studies or parametric bootstrap calculations. The package is available from CRAN (<http://www.cran.r-project.org>) and is open source.

The current version of the package limits the general model described in this article in several ways. In predictions, only linear h are allowed in (19), but this can be worked around. For general h , observe that $h(\hat{\boldsymbol{\zeta}})$ and $\mathbf{A}^T \hat{\boldsymbol{\zeta}}$, where $\mathbf{A} = \nabla h(\boldsymbol{\zeta})$, have the same standard errors. Thus obtain $h(\hat{\boldsymbol{\zeta}})$ by one call to `predict.aster` and the standard errors for $\mathbf{A}^T \hat{\boldsymbol{\zeta}}$, where $\mathbf{A} = \nabla h(\hat{\boldsymbol{\zeta}})$ by a second call. In models, the only conditional families

currently implemented are Bernoulli, Poisson, and zero-truncated Poisson, but adding another one-parameter exponential family would require only implementation of its ψ , ψ' , and ψ'' functions. Multiparameter conditional families (chain components) are not yet implemented. Allowing terminal nodes that are two-parameter normal or allowing child nodes that are multinomial given a common parent would require more extensive changes to the package.

4. EXAMPLE

Data were collected on 570 individuals of *Echinacea angustifolia*, each having the data structure shown in Figure 1. These plants were sampled as seeds from seven remnant populations that are surviving fragments of the tall-grass prairie that a century ago covered western Minnesota and other parts of the Great Plains of North America. The plants were experimentally randomized at the time of planting into a field within 6.5 km of all populations of origin. The data set contains three predictor variables: `ewloc` and `nsloc` give east-west and north-south positions of individuals and `pop` gives their remnant population of origin. To use the R formula mini-language we need to create more variables: `resp` is a vector comprising the response variables (the M_j , F_j , and H_j), `level` is categorical naming the type of response variable (values M , F , and H), `year` is categorical giving the year, `varb` is the interaction of `level` and `year`, and `hdct` is short for `level = H`.

We fit many models. See Appendix D of the technical report at <http://www.stat.umn.edu/geyer/aster> for details. Scientific interest focuses on the model comparison shown in Table 1.

[Table 1 about here.]

The models are nested. The linear predictor for Model 4 can be written

$$\varphi_j = \mu_{L_j, Y_j} + \alpha_{L_j} U_j + \beta_{L_j} V_j + \gamma_{L_j, P_j} \quad (21)$$

where L_j , Y_j , U_j , V_j , and P_j are `level`, `year`, `ewloc`, `nsloc`, and `pop`, respectively, and the alphas, betas, and gammas are regression coefficients. Model 3 is the submodel of Model 4 that imposes the constraint $\gamma_{M,P} = \gamma_{F,P}$, for all populations P . Model 2 is the submodel of Model 3 that imposes the constraint $\gamma_{M,P} = \gamma_{F,P} = 0$, for all P . Model 1 is the submodel of Model 2 that imposes the constraint $\gamma_{L,P} = 0$, for all L and P .

All models contain the graph node effect (`varb` or μ_{L_j, Y_j}) and the quantitative spatial effect (`level:(nsloc + ewloc)` or $\alpha_{L_j} U_j + \beta_{L_j} V_j$), which was chosen by comparing many models (details in the technical report). We explain here only differences amongst models, which involve only categorical predictors. In an unconditional aster model, which these are, such terms require the maximum likelihood mean value parameters for each category (summed over all individuals in the category) to match their observed values (‘observed equals expected’). Model 4 makes observed equal expected for *total head count* $\sum_j H_j$, for *total flowering* $\sum_j F_j$, and for *total survival* $\sum_j M_j$ within each population. Model 3 makes observed equal expected for *total head count* $\sum_j H_j$ and for *total non-head count* $\sum_j (M_j + F_j)$ within each population. Model 2 makes observed equal expected for *total head count* $\sum_j H_j$ within each population.

From purely statistical considerations, Model 3 is the best of these four nested models. Model 4 does not fit significantly better. Model 2 fits significantly worse. Model 3 is difficult to interpret scientifically, because ‘non-head count’ $\sum_j (M_j + F_j)$ is unnatural, in effect scoring 0 for dead, 1 for alive without flowers, or 2 for alive with flowers.

Model 2 is the model of primary scientific interest. Evolutionary biologists are fundamentally interested in fitness. For our purposes here the fitness of an individual may be defined as its contribution over its lifespan in descendants to the next generation (see Beatty, 1992; Keller, 1992; Paul, 1992, for further discussion). Fitness is notoriously difficult to measure. For these data the most direct surrogate measure for fitness is *total head*

count $\sum_j H_j$. The currently available data represent a small fraction of this plant's life-span. To obtain more complete measures of fitness, we are continuing these experiments and collecting these data for successive years.

Biologists call all our measured variables (the M_j , F_j , and H_j) *components of fitness*. Since M_j and F_j contribute to fitness (descendants) only through H_j , in an aster model the unconditional expectation of H_j (its unconditional mean value parameter) completely accounts for the contributions of M_j and F_j . Strictly speaking, this is not quite true, since we do not have H_j measured over the whole life span, so the last M_j contains the information about future reproduction, but it becomes truer as more data are collected in future years. Moreover, we have no data about life span and do not wish to inject subjective opinion about future flower head count into the analysis.

The statistical point of this is that the M_j and F_j are in the model only to produce the correct stochastic structure. If we could directly model the marginal distribution of the H_j (but we can't), we would not need the other variables. They are 'nuisance variables' that must be in the model but are of no interest in this particular analysis (the mean value parameters for those variables are nuisance parameters). Model 3 is the best according to the likelihood ratio test, but does it fit the variables of interest better than Model 2? We do not know of an established methodology addressing this issue, so we propose looking at confidence intervals for the mean value parameters for total flower head count shown in Figure 2.

[Figure 2 about here.]

Although we have no formal test to propose, we claim it is obvious that Model 3 is no better than Model 2 at 'predicting' the best surrogate of expected fitness. We take this as justification for using Model 2 in scientific discussion and infer from it significant differences among the populations in flower head count and, thus, fitness.

Model 3 being difficult to interpret scientifically, Model 4 is the next larger readily interpretable model. Model 4 does fit significantly better than Model 2 ($P = 0.00016$), which implies differences among populations in mortality and flowering (the M_j and F_j) that may be of scientific interest even though they make no direct contribution to fitness (since Model 2 already fully accounts for their contributions through H_j).

Note that we would have gotten very different results had we used a conditional model (not shown, see Section D.3.2 of the technical report). The parameters of interest are *unconditional* expectations of total flower head count. This alone suggests an unconditional model. Furthermore, we see in (5) that unconditional aster models ‘mix levels’ passing information up from children to parents. This is why Model 2 in our example was successful in predicting total head count while only modelling `pop` effects at head count nodes. Because a *conditional* aster model does not mix levels in this way, it *must* model all levels and so usually needs many more parameters than an unconditional model.

5. DISCUSSION

The key idea of aster models (as we see it) is the usefulness of what we have called unconditional aster models, which have low-dimensional sufficient statistics (9). Following Geyer and Thompson (1992), who argued in favour of exponential family models with sufficient statistics chosen to be scientifically interpretable, an idea they attributed to Jaynes (1978), we emphasize the value of these models in analyses of life histories and overall fitness.

We don’t insist, though. The R package `aster` is even-handed with respect to conditional and unconditional models and conditional and unconditional parameters. Users may use whatever seems best to them. Any joint analysis is better than any separate analyses of different variables. We have, however, one warning. In an unconditional aster model, as in all exponential family models, the map from canonical parameters to mean value parameters

is monotone; with sufficient statistic vector \mathbf{Y} given by (9) having components Y_k we have

$$-\frac{\partial^2 l(\boldsymbol{\beta})}{\partial \beta_k^2} = \frac{\partial E_{\boldsymbol{\beta}}\{Y_k\}}{\partial \beta_k} > 0. \quad (22)$$

This gives regression coefficients their simple interpretation: an increase in β_k causes an increase in $E_{\boldsymbol{\beta}}\{Y_k\}$, other betas being held constant. The analog of (22) for a conditional model is

$$-\frac{\partial^2 l(\boldsymbol{\beta})}{\partial \beta_k^2} = \frac{\partial}{\partial \beta_k} \sum_{G \in \mathcal{G}} \sum_{j \in G} E_{\boldsymbol{\beta}}\{X_j | X_{p(G)}\} m_{jk} > 0, \quad (23)$$

where m_{jk} are components of \mathbf{M} . Because we cannot move the sums in (23) inside the conditional expectation, there is no corresponding simple interpretation of regression coefficients. So conditional aster models are algebraically simple but statistically complicated. Unconditional aster models are algebraically complicated but statistically simple. They can be simply explained as flat exponential families having the desired sufficient statistics. They are the aster models that behave according to the intuitions derived from linear and generalized linear models.

We saw in our example that aster models allowed us to successfully model fitness. In humans, Darwinian fitness is not the *summum bonum* that it is for other organisms, but neither is mere survival! Survival analysis, by its very name, rules out any consideration of quality of life. Aster models do permit such consideration.

ACKNOWLEDGEMENTS

Our key equation (5) was figured out about 1980 when R. G. S. was a graduate student and C. J. G. not yet a graduate student. R. G. S.'s Ph. D. adviser Janis Antonovics provided some funding for development (from a now forgotten source, but we thank it), but we did not publish then, and other work intervened. High quality life history data, generated recently by S. W., Julie Etterson, David Heiser, and Stacey Halpern, and work of Helen Hangelbroek (funded by NSF DMS-0083468) about what in hindsight we would call a conditional aster

model analysis of mortality data, motivated implementation of these ideas. Were it not for Janis's enthusiastic encouragement back then, we might not have been able to do this now. We should also acknowledge the R team. Without R we would never have made software as usable and powerful as the `aster` package. We also thank Robert Gentleman for pointing out the medical implications.

REFERENCES

- BARNDORFF-NIELSEN, O. E. (1978). *Information and Exponential Families*. Chichester: John Wiley.
- BEATTY, J. (1992). Fitness: Theoretical contexts. In *Keywords in Evolutionary Biology*, Ed. E. F. Keller and E. A. Lloyd, pp. 115–119. Cambridge, Mass: Harvard University Press.
- BRESLOW, N. (1972). Discussion of the paper by Cox (1972). *J. R. Statist. Soc. B* **34**, 216–17.
- BRESLOW, N. (1974). Covariance analysis of censored survival data. *Biometrics*, **30**, 89-99.
- CHAMBERS, J. M. & HASTIE, T. J. (1992). Statistical models. In *Statistical Models in S*, J. M. Chambers and T. J. Hastie, eds. Pacific Grove, Calif.: Wadsworth & Brooks/Cole.
- COX, D. R. (1972). Regression models and life-tables (with discussion). *J. R. Statist. Soc. B* **34**, 187–220.
- GEYER, C. J. & THOMPSON, E. A. (1992). Constrained Monte Carlo maximum likelihood for dependent data, (with discussion). *J. R. Statist. Soc. B* **54** 657–99.
- JAYNES, E. T. (1978). Where do we stand on maximum entropy? In *The Maximum*

- Entropy Formalism*, Ed. R. D. Levine and M. Tribus, pp. 15-118. Cambridge, Mass: MIT Press.
- KELLER, E. F. (1992). Fitness: Reproductive ambiguities. In *Keywords in Evolutionary Biology*, Ed. E. F. Keller and E. A. Lloyd, pp. 120–121. Cambridge, Mass: Harvard University Press.
- LAURITZEN, S. L. (1996). *Graphical Models*. New York: Oxford University Press.
- MARTIN, T. G., WINTLE, B. A., RHODES, J. R., KUHNERT, P. M., FIELD, S. A., LOW-CHOY, S. J., TYRE, A. J., and POSSINGHAM, H. P. (2005). Zero tolerance ecology: improving ecological inference by modelling the source of zero observations. *Ecology Letters* **8** 1235–1246.
- MCCULLAGH, P. & NELDER, J. A. (1989). *Generalized Linear Models*, 2nd ed. London: Chapman & Hall.
- PAUL, D. (1992). Fitness: Historical perspective. In *Keywords in Evolutionary Biology*, Ed. E. F. Keller and E. A. Lloyd, pp. 112–114. Cambridge, Mass: Harvard University Press.
- R DEVELOPMENT CORE TEAM (2004). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org>.
- WILKINSON, G. N. & ROGERS, C. E. (1973). Symbolic description of factorial models for analysis of variance. *Applied Statistics* **22**, 392–99.




Figure 1: Graph for *Echinacea* Aster Data. Arrows go from parent nodes to child nodes. Nodes are labelled by their associated variables. The only root node is associated with the constant variable 1. M_j is the mortality status in year $2001 + j$. F_j is the flowering status in year $2001 + j$. H_j is the flower head count in year $2001 + j$. The M_j and F_j are Bernoulli conditional on their parent variables being one (and zero otherwise). The H_j are zero-truncated Poisson conditional on their parent variables being one (and zero otherwise).




Figure 2: Confidence Intervals for Total Head Count. 95% (non-simultaneous) confidence intervals for the unconditional expectation of total flower head count (all three years) for individuals from different populations and central spatial location. Solid bars based on Model 2 in Table 1. Dashed bars based on Model 3 in Table 1.

Table 1: Tests for Model Comparison. The model formulae are given above and the analysis of deviance below (deviance is twice log likelihood).

- 1: `resp ~ varb + level:(nsloc + ewloc)`
- 2: `resp ~ varb + level:(nsloc + ewloc) + hdct * pop - pop`
- 3: `resp ~ varb + level:(nsloc + ewloc) + hdct * pop`
- 4: `resp ~ varb + level:(nsloc + ewloc) + level * pop`

Model Number	Model d. f.	Model Deviance	Test d. f.	Test Deviance	Test <i>p</i> -value
1	15	2728.72			
2	21	2712.54	6	16.18	0.013
3	27	2684.86	6	27.67	0.00011
4	33	2674.70	6	10.17	0.12