

The BUGS project: Evolution, critique and future directions

David Lunn^{1,*},[†], David Spiegelhalter¹, Andrew Thomas² and Nicky Best³

¹*Medical Research Council Biostatistics Unit, Institute of Public Health, University Forvie Site, Robinson Way, Cambridge CB2 0SR, U.K.*

²*School of Mathematics and Statistics, Mathematical Institute, North Haugh, St. Andrews, Fife KY16 9SS, U.K.*

³*Department of Epidemiology and Public Health, Imperial College London, St. Mary's Campus, Norfolk Place, London W2 1PG, U.K.*

SUMMARY

BUGS is a software package for Bayesian inference using Gibbs sampling. The software has been instrumental in raising awareness of Bayesian modelling among both academic and commercial communities internationally, and has enjoyed considerable success over its 20-year life span. Despite this, the software has a number of shortcomings and a principal aim of this paper is to provide a balanced critical appraisal, in particular highlighting how various ideas have led to unprecedented flexibility while at the same time producing negative side effects. We also present a historical overview of the BUGS project and some future perspectives. Copyright © 2009 John Wiley & Sons, Ltd.

KEY WORDS: BUGS; WinBUGS; OpenBUGS; Bayesian modelling; graphical models

1. INTRODUCTION

BUGS [1] is a software package for performing Bayesian inference using Gibbs sampling [2, 3]. The BUGS project began at the Medical Research Council Biostatistics Unit in Cambridge in 1989. Since that time the software has become one of the most popular statistical modelling packages, with, at the time of writing, over 30 000 registered users of *WinBUGS* (the Microsoft Windows incarnation of the software) worldwide, and an active on-line community comprising over 8000 members. Typing ‘WinBUGS’ into Google generates over 100 000 hits; in Google Scholar the figure is in excess of 5000; and simply searching Statistics in Medicine on-line gives nearly 100 hits.

BUGS has been just one part of the tremendous growth in the application of Bayesian ideas over the last 20 years. Prior to the widespread introduction of simulation-based methods, Bayesian ideas

*Correspondence to: David Lunn, Medical Research Council Biostatistics Unit, Institute of Public Health, University Forvie Site, Robinson Way, Cambridge CB2 0SR, U.K.

[†]E-mail: david.lunn@mrc-bsu.cam.ac.uk

could only be implemented in circumstances in which solutions could be obtained in closed form in so-called conjugate analyses, or by ingenious but restricted application of numerical integration methods. BUGS has therefore been instrumental in raising awareness of Bayesian modelling among both academic and commercial communities internationally, and has greatly facilitated the routine analysis of many complex statistical problems. Numerous applications of the software can be found in the literature, in a wide array of application areas, including disease mapping [4], pharmacometrics [5], ecology [6], health-economics [7], genetics [8], archaeology [9], psychometrics [10], coastal engineering [11], educational performance [12], behavioural studies [13], econometrics [14], automated music transcription [15], sports modelling [16], fisheries stock assessment [17], and actuarial science [18]. The software is also used widely for the teaching of Bayesian modelling ideas to students and researchers the world over, and several texts use BUGS extensively for illustrating the Bayesian approach [19–26]. The importance of the software has been acknowledged in ‘An International Review of U.K. Research in Mathematics’ (<http://www.cms.ac.uk/irm/irm.pdf>).

Despite this apparent success, we are well aware that BUGS has some significant shortcomings and it is our intention here to provide a balanced critical appraisal of the software in its current form(s). We will also present an overview of the way in which the software has evolved, and a discussion of potential future directions for the BUGS project. We recognize that not all parts of the paper will be of interest to all readers, and so try to provide appropriate guidance.

The structure of this paper is as follows. In Section 2 we explain the intimate and vital connection between BUGS and graphical modelling, while in Section 3 we present an overview of the BUGS language: these sections could be skipped by those familiar to ideas behind the program. In Section 4 we describe how the software has evolved, which goes into more technical detail regarding the ‘guts’ of the program and how the different versions were developed and funded. The technical material is presented in a separate subsection (Section 4.2) and can be skipped by those who are less interested in how BUGS’ capabilities have expanded over the years. Section 5 provides a critical appraisal and discusses features that should be irritatingly familiar to all those who have struggled with the program. A concluding discussion, including some future perspectives, is given in Section 6.

2. A BASIS IN GRAPHICAL MODELLING

The popularity of the software stems from its flexibility, which is due to the exploitation of graphical modelling theory. Consider a simple linear regression problem given by:

$$y_i \sim N(\mu_i, \sigma^2), \quad \mu_i = \alpha + \beta x_i, \quad i = 1, \dots, N$$

for responses y_i and observed values x_i of a single covariate ($i = 1, \dots, N$). As we are working in a Bayesian setting we assign prior distributions to the unknown parameters, α , β and σ^2 (or σ , or $\log \sigma$, say). For example,

$$\alpha \sim N(m_\alpha, v_\alpha), \quad \beta \sim N(m_\beta, v_\beta), \quad \log \sigma \sim U(a, b)$$

for fixed constants m_α , v_α , m_β , v_β , a and b . An alternative representation of this model is the *directed acyclic graph* (DAG; see [27], for example) shown in Figure 1, where each quantity in the model corresponds to a *node* and links between nodes show direct dependence. The graph is *directed* because each link is an arrow; it is *acyclic* because by following the arrows it is not possible to return to a node after leaving it.

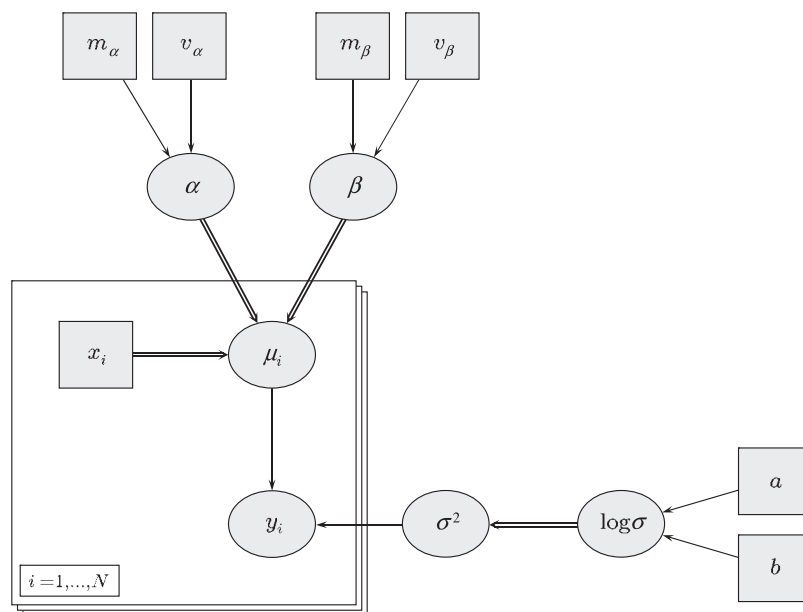


Figure 1. Directed acyclic graph (DAG) corresponding to linear regression example—see text for details.

The notation is defined as follows. Rectangular nodes denote known constants. Elliptical nodes represent either deterministic relationships (i.e. functions) or stochastic quantities, i.e. quantities that require a distributional assumption. Stochastic dependence and functional dependence are denoted by single-edged arrows and double-edged arrows, respectively. Repetitive structures, such as the ‘loop’ from $i = 1$ to N , are represented by ‘plates’, which may be nested if the model is hierarchical.

A node v is said to be a *parent* of node w if an arrow emanating from v points to w ; furthermore, w is then said to be a *child* of v . We are primarily interested in stochastic nodes, i.e. the unknown parameters and the data. When identifying probabilistic relationships between these, deterministic links are collapsed and constants are ignored. Thus, the terms parent and child are usually reserved for the appropriate stochastic quantities. In the above example, the *stochastic parents* of each y_i are α , β and $\log \sigma$, whereas we can refer to μ_i and σ^2 as *direct parents*.

DAGs can be used to describe pictorially a very wide class of statistical models through describing the ‘local’ relationships between quantities. It is when these models become complicated that the benefits become obvious. DAGs communicate the essential structure of the model without recourse to a large set of equations. This is achieved by abstraction: the details of distributional assumptions and deterministic relationships are ‘hidden’. This is conceptually similar to object-oriented programming (OOP) (see [28], for example), which realization has been instrumental in maximizing the flexibility of BUGS, as we will discuss later.

The joint distribution of all nodes V in *any* DAG can be factorized as follows [29]:

$$p(V) = \prod_{v \in V} p(v | \mathcal{P}_v) \quad (1)$$

where \mathcal{P}_v denotes the set of (stochastic) parents of node v . With regard to Bayesian inference, the data D and the unknowns θ , say, together form a partition of V , and so the joint posterior density $p(\theta|D) \propto p(V)$. For the purposes of Gibbs sampling we are interested in the full conditional distributions of each unknown stochastic node conditional on the values of all other nodes in the graph. For two arbitrary sets of nodes A and B , say, let $A \setminus B$ denote ‘all elements of A except those in B ’. Then the full conditional for a specific node w , say, is denoted $p(w|V \setminus w)$. As $\{w, V \setminus w\}$ is also a partition of V , we have that the full conditional is also proportional to the right-hand side of (1). However, $p(w|V \setminus w)$ is a distribution in w only, and so we can ignore any factors not involving w , giving

$$p(w|V \setminus w) \propto p(w|\mathcal{P}_w) \times \prod_{v \in \mathcal{C}_w} p(v|\mathcal{P}_v)$$

where \mathcal{C}_w denotes the children of w . The beauty of the factorization (1) is thus two-fold. First, we can write down the joint posterior for any DAG simply by knowing the relationship between each node and its parents. Second, the full conditional distribution for any node (or set of nodes) is a *local computation* on the graph, involving only the node-parent dependencies for the node of interest itself and its children. Thus, one only ever needs to consider a small part of the model at any given time, without needing to take account of the bigger picture. The BUGS software exploits these facts first by providing a language (the BUGS language) for specifying arbitrary child–parent relationships, and by ‘inverting’ these to determine the set of nodes relevant to each full conditional calculation: these comprise the children, parents and ‘co-parents’ collectively known as the ‘Markov blanket’. The software also provides a mechanism whereby each node can communicate, to the ‘inference engine’, how it is related to its parents. With these facilities in place, Gibbs sampling on arbitrary graphs is conceptually straightforward.

3. THE BUGS LANGUAGE

As noted above, BUGS provides its own language for the textual specification of graphical models. The language is designed to mimic the mathematical specification of the model in terms of parent–child relationships. Stochastic relationships are denoted with ‘ \sim ’ whereas logical/deterministic relationships are denoted with a ‘ \leftarrow ’. Repetitive structures are represented using ‘for-loops’, which may be nested if the model is hierarchical, say. The following code specifies the linear regression graph referred to in the previous section:

```
model {
  for (i in 1:N) {
    y[i] ~ dnorm(mu[i], tau)
    mu[i] <- alpha + beta * x[i]
  }
  alpha ~ dnorm(m.alpha, p.alpha)
  beta ~ dnorm(m.beta, p.beta)
  log.sigma ~ dunif(a, b)
  sigma <- exp(log.sigma)
  sigma.sq <- pow(sigma, 2)
  tau <- 1 / sigma.sq
}
```

```

p.alpha  <- 1 / v.alpha
p.beta   <- 1 / v.beta
}

```

The code should be reasonably self-explanatory but a few notes are required for clarification. First note that normal distributions (`dnorm(. , .)`) are parameterized in terms of precision (equal to $1/\text{variance}$) as opposed to variance. This can be a source of confusion and reflects the fact that the software was originally designed to exploit mathematical conveniences (so-called conjugacy) where possible. Although this is no longer necessary and the need for an intuitive prior typically outweighs the benefit of any such conjugacy, the parameterization remains unchanged. In addition, note that the data, `y[1:N]`, `x[1:N]`, `N`, and the values of `m.alpha`, `v.alpha`, `m.beta`, `v.beta`, `a` and `b` in this case, are loaded into the system separately. Finally, note that the `pow(. , .)` function raises the first argument to the power of the second argument, and so the `sigma.sq` variable represents σ^2 .

The BUGS language is a *declarative* language. This means that it does not matter in which order the various statements are made, which makes it easier for the system to interpret the model specification. The downside of this, however, is that there is no scope for ‘logical’ expressions such as IF–THEN–ELSE statements. There are ways around this, for example, by making use of the `step(.)` function (equal to $1/0$ depending on the positive/negative status of the argument) to switch between modelling terms, but this still represents a significant limitation of the language. Despite this the language has proved enormously successful and capable of expressing a vast array of model-types.

To illustrate the flexibility afforded by the language, let us elaborate the model described above a little. First, suppose that the covariate `x[]` is subjected to measurement error. We could model this, for example, by adding the assumption $x[i] \sim \text{dnorm}(\mu.x[i], \tauau.x)$, where the $\mu.x[i]$ terms are unknown and assigned an appropriate prior, e.g. $\mu.x[i] \sim \text{dnorm}(m.x, p.x)$ with the mean `m.x` and precision `p.x` known, and $\tauau.x$, for the sake of simplicity, also known. Now suppose, for instance, that the responses `y[]` represent growth data on a given individual, collected at times `x[]`. If we had such data on a number of individuals (indexed by $j = 1, \dots, K$) as opposed to just one, then we might wish to assume different individuals to have distinct but similar (exchangeable) parameters via $\alpha_j \sim N(m_\alpha, v_\alpha)$, $\beta_j \sim N(m_\beta, v_\beta)$, $j = 1, \dots, K$, where now m_α , v_α , m_β and v_β are unknown and assigned appropriate priors, e.g.

$$m_\alpha \sim N(0, 100^2), \quad \sqrt{v_\alpha} \sim U(0, 100), \quad m_\beta \sim N(0, 100^2), \quad \sqrt{v_\beta} \sim U(0, 100)$$

To implement this now hierarchical model we simply add lines to represent the priors for m_α , v_α , m_β and v_β , e.g. `m.alpha ~ dnorm(0, 0.0001)`, and nest any individual-specific statements within a loop over individuals (indexed by `j`):

```

for (j in 1:K) {
  for (i in 1:N) {
    y[i,j] ~ dnorm(mu[i,j], tau)
    mu[i,j] <- alpha[j] + beta[j] * x[i]
  }
  alpha[j] ~ dnorm(m.alpha, p.alpha)
  beta[j]  ~ dnorm(m.beta, p.beta)
}

```

Suppose we wish to fit a non-linear growth curve, $\mu_{ij} = \alpha_j - \beta_j \gamma_j^{x_i}$, say, as opposed to a straight line. Then we would rewrite the definition of `mu[i, j]` as follows:

```
mu[i, j] <- alpha[j] - beta[j] * pow(gamma[j], x[i])
```

If the α_j s and β_j s are as before and we further suppose $\gamma_j \sim U(0, 1)$ for $j = 1, \dots, K$, then we simply include the line `gamma[j] ~ dunif(0, 1)` within the `j`-loop (but outside the `i`-loop). Finally, to robustify the model against any outlying observations we might wish to assume the `y[i, j]`s arise from a Student-*t* as opposed to a normal distribution. Then we would modify the assumed relationship between each `y[i, j]` and its parents to `y[i, j] ~ dt(mu[i, j], tau, d)`, where the degrees-of-freedom parameter `d` is either assumed to be known (e.g. `d <- 4`) or assigned an appropriate prior, Poisson, say, `d ~ dpois(...)`.

One of the reasons the BUGS language is so flexible is that it is open ended. It is quite straightforward to introduce new ‘vocabulary’, i.e. new distributions and functions, and indeed this can be done without any part of the existing software being modified or even recompiled. This is due to the object- and component-oriented [30] nature of the software, which is discussed throughout.

4. EVOLUTION OF BUGS

4.1. Inception and early years

In the early 1980s one focus of work in artificial intelligence concerned *expert systems*, which were programs designed to mimic expertise in complex situations. A basic principle was a separation of the *knowledge base* that encapsulated expertise, from the *inference engine* that controlled the response to new evidence and queries: there was agreement that a natural computer implementation of a knowledge base was in a *declarative* rather than a *procedural* form, so that the program would express local relationships between entities which could then be manipulated using the inference engine.

Where there was substantial disagreement and sometimes passionate argument was about how to deal with uncertainty. For example, extremely influential work from Stanford featured rule-based systems in which uncertainty was handled using a system of ‘certainty factors’ attached to rules and manipulated according to somewhat arbitrary principles—this was labeled as ‘*ad hoc* quantitative mumbo-jumbo’ by a prominent statistician (Smith, discussion of [31]). The challenge was to come up with a more rigorous process based on probability theory but that still retained the attractive aspect of local computation. The potential role of graphical models, also to become known as Bayesian networks, was argued both in the U.S. [32] and in the U.K. [33], and exact means of propagating uncertainty in tree structures [34] and general networks [35] were developed.

Pearl [36] noted that simulation methods could also be used for approximate inference in directed graphical models, and Spiegelhalter and Lauritzen [37] showed that graphical models could also include parameters as nodes, hence facilitating general Bayesian inference. It became apparent that, by merging these ideas, the computations required for Bayesian inference were simply an iterative sequence of local calculations on the graph, and that each such calculation, viewed abstractly, was identical. OOP would therefore prove highly effective: not only does it provide a means of describing relationships between entities of different, but related, types, i.e. a *virtual graphical*

model (VGM),[‡] but it also allows abstract computation, vastly simplifying implementation (and maintenance) of the algorithm (see [38] for details). Putting these ideas together led to the concept of a general program that could handle arbitrary models and would use simple Gibbs sampling to simulate from nodes conditional on their neighbours in the graph. A start was made on the BUGS program in 1989 with the appointment of Andrew Thomas to the MRC Biostatistics Unit: it is rather extraordinary that at the same time the classic MCMC work of Gelfand and Smith [3] was being carried out 80 miles away in Nottingham but entirely independently and from a rather different starting point.

The language chosen for implementation was Modula-2 [39]. Actually this is not an object-oriented language but it can be used to a similar effect. In addition it is modular (as one might expect from its name), which facilitates the hiding of detail, an important concept in designing complex architectures, for much the same reasons as it is effective in graphical modelling. An obvious alternative might have been C++; however, as Andrew Thomas points out '*I am unable to program in C++. If I was clever enough to program in C++ I would have a very good job in say banking.*'

A prototype on a PC was demonstrated at the 4th Bayesian Valencia meeting in April 1991 [40], and a Unix version at the Practical Bayesian Statistics meeting in Nottingham in 1992 [1]. Version 0.1 for Unix was released in February 1993 and further refinements followed until a stable Version 0.5 in 1995. Early versions were distributed on disks and a voluntary fee of £25 or \$40 solicited. A large impetus was the INSERM workshop on MCMC methods in 1993 which had a follow-up workshop in Cambridge using BUGS and led to the *MCMC in Practice* book [41] (which, rather remarkably, is still selling well). Throughout this period interest in MCMC was steadily growing, and the portfolio of examples grew, including the standard lip-cancer spatial mapping example which was made to work on the day before the Basel–Amsterdam Bayesian Riverboat conference in 1993.

The name BUGS shows that only simple Gibbs sampling was used at first, and even then only using simple inversion, conjugate updating, or adaptive rejection sampling for log-concave full conditionals. The parameterizations for each distribution were chosen for computational convenience, which is why the normal ended up being parameterized in terms of precision (inverse-variance). After some time a grid-based Metropolis–Hastings algorithm [42, 43] was implemented as a somewhat crude means of performing more general sampling. However, it was not until the project moved (in 1996) to Imperial College, London, that the sampling capabilities of the software were expanded. We describe such technological developments in the following subsection, which can be skipped by uninterested readers. In short, we describe how and why a Windows version was developed, and how various modelling challenges in fields such as pharmacokinetics, disease mapping and genetic epidemiology led to improved algorithms and greater flexibility.

4.2. Technological evolution

4.2.1. A stand-alone (Windows) version.

In the mid-1990s a decision was taken to create a stand-alone version of the software, one which could provide interactive diagnostic and inference tools. This led to a new choice of programming language, for two main reasons. First the Microsoft Windows operating system had been rapidly gaining massive popularity and it seemed an ideal

[‡]Simplistically speaking, a VGM uses *objects* and *pointer variables*, respectively, to represent nodes and directed links in the graph.

environment for the kind of fully interactive software that was planned; it was also thought that we would reach a much wider audience through this medium. Second, the parallels between OOP and graphical modelling were becoming more and more apparent and it was thought that these could be better exploited. In addition, the use of a *component-oriented* framework [30] would facilitate the growth of the BUGS-language-vocabulary, by allowing new distributions and functions to be implemented with relative ease.

The programming language chosen was Component Pascal [44], implemented within a *Rapid Application Development* environment known as *BlackBox Component Builder* (<http://www.oberon.ch/blackbox.html>). Component Pascal follows the tradition of Pascal but is actually incompatible with it. The language encompasses both procedural and object-oriented paradigms and is modular, again making for well-structured architectures and greater ease of maintenance. An in-built *Garbage Collector* [45] ensures reliable memory management and a *run-time linking* facility allows components to be loaded *on demand*, which feature is key in ensuring open-endedness of the BUGS language.

The BlackBox framework supports rapid application development by providing infrastructure for creating the features inherent with interactive software, such as menu commands, dialogue boxes and graphics. It is very powerful but much of the functionality is provided at a very low level. For example, the graphics subsystem is so low-level that an entirely new subsystem has had to be built on top of this for rendering statistical plots in a reasonably intuitive manner. It has been difficult to justify investing too much time in such developments due to the academic nature of BUGS' funding over the years. Hence, the software's graphics have always been a little basic, although, hopefully, they are adequate.

4.2.2. Expansion of capabilities. Shortly after the first stand-alone version (WinBUGS) was functional, work began on expanding its capabilities. The first task was to implement a general purpose Metropolis–Hastings sampler for updating continuous scalar quantities whose full conditional distribution is neither available in closed form nor log-concave. The proposal distribution is a normal distribution centred at the current point with a self-tuning variance adjusted every 100 iterations with a view to (ultimately) achieving an acceptance rate of between 20 and 40 per cent, an intuitive balance between the regularity of movement and the distance travelled in each move. These limits are fairly arbitrary but are based loosely on formal analysis [46].

The self-tuning variance is encapsulated within a Metropolis–Hastings *updater object*, an instance of which is coupled to each node in the graph requiring Metropolis updates. Hence, a distinct and independently adapted proposal variance is used for each such node, so that the sampling of each full conditional is 'optimized' as opposed to using a globally adapted variance which may perform poorly for some nodes. The initial size of the proposal variance is small, e.g. 10^{-4} . This ensures regular movements (high acceptance) towards the target mode(s), which allows reliable location of the vicinity of the target density in a reasonably short time. As the target density is homed-in upon, the proposal variance is gradually increased until it is of a similar, but somewhat larger, size to that of the target distribution.

The self-tuning Metropolis updater was initially developed as a part of the PKBugs project, which aimed to provide an interface (to WinBUGS) for pharmacokinetic modelling. Not only did PKBugs provide a test-bed for serious non-linear modelling, whereby the reliability of the Metropolis updater's adaption algorithm could be honed, but it also served as a proof-of-concept study for fully exploiting the software's various extensibility features. In particular, PKBugs showed that the BUGS language could easily be expanded to include (arbitrarily) complex functions, and

that specialized interfaces for describing the statistical model and for presenting the output in different ways could be implemented in a straightforward manner.

To further the range of non-linear modelling possibilities, other forms of Metropolis proposal distribution for sampling from distributions with restricted support (interval or semi-infinite) were experimented with, until Radford Neal's slice sampler was invented [47]. Like the Metropolis updater this needs adapting over time, by estimating the typical size of a 'slice'. The procedure can be easily automated, however, and, again like the Metropolis updater, the tuning parameter can be encapsulated within an updater object so that each full conditional is adapted to independently.

At this point, WinBUGS seemed capable of fitting pretty much anything we could throw at it, and work was completed on the graphics. This provided a good opportunity to try pushing the system further by developing GeoBUGS, a specialized interface for spatial modelling, which would provide new distributions for spatially correlated quantities and graphical devices for mapping them as well as for facilitating model specification. The majority of such distributions actually correspond to undirected sub-graphs. However, they are easily accommodated within the directed graphical framework of BUGS by simply thinking of them as multivariate quantities represented by a single node in the graph with complex internal structure—if all internal nodes can be updated together, then the internal structure is unimportant for the purposes of performing the encompassing Gibbs scheme.

In another development, an extension of the PKBUGS project led to consideration of how we might implement models described in terms of differential equations, which are important in infectious disease epidemiology and mathematical biology [48, 49], for example, as well as in the motivating pharmacokinetic/toxicokinetic problems that we considered [50]. This work uncovered a significant limitation of the way in which BUGS implemented virtual graphical models (VGMs), an inability to accommodate vector-valued functions. For example, at this point the matrix-inverse function could only provide one element of the inverted matrix for each inversion, requiring $p(p+1)/2$ inversions to obtain the whole matrix. Clearly, such inefficiencies had to be circumvented to make numerical differential equation solving across a grid of time-points feasible. The problem here was that there was no mechanism in place for remembering when calculations had already been performed. One of the most significant advances in the design of the BUGS VGM was to implement such a memory mechanism, whereby logical nodes would know whether any of their stochastic parents had been updated since their last evaluation and thus require reevaluating. This was achieved by all stochastic nodes broadcasting an appropriate message to their descendants whenever their value was updated. This paved the way for implementing more and more complex functions efficiently and also inspired an improvement to the software's Metropolis updaters—a similar mechanism was used to ensure that complex functions saved their current values, in case of rejection, before any relevant Metropolis proposals.

At this point, we had had a great deal of experience in implementing new distributions and functions and realized that much of the work involved, although complicated, was somewhat routine. It became apparent that we could build a module hierarchy for implementing new modelling components whereby only the most basic function-specific or distribution-specific details need be specified directly by the user, and all of the less specific detail would be hidden beneath. This led to the *WinBUGS Development Interface* (WBDev [51]). In many cases, WBDev could offer the user significant (sometimes massive) run-speed improvements, due to 'hard-wiring' of the new distribution/function into the system using compiled code as opposed to specifying the new component via the BUGS language, which is interpreted at run-time. Other advantages included tidier (and less error-prone) BUGS models due to 'packaging' of complex code, and the facility to

exploit a comprehensive computer language for specifying the new component as opposed to the somewhat limited BUGS language, e.g. piecewise functions could be specified using IF–THEN–ELSE constructs rather than via BUGS' `step(.)` function. The WBDev project also served as a proof-of-concept for making the software open source.

4.2.3. Divergence of established and experimental versions. In 2004 Andrew Thomas began work on an open-source version of the software at the University of Helsinki. The OpenBUGS project initially had three main aims. First, decoupling of core functionality and the user interface, to facilitate use of the software from within other environments. The most significant output from this research arm was BRugs, an interface that allows BUGS to be used interactively from within R. Several other packages exist for interfacing with *WinBUGS* through other software, e.g. SAS, Stata, Excel, but the beauty of BRugs is that it is fully interactive and BUGS output can be accessed mid-analysis for exploiting the full power of R in generating diagnostic plots, for instance. A second aim of the OpenBUGS project was to make the software more platform-independent, and to this end a Linux-based version, namely LinBUGS, has been developed; currently, this can be run on any Intel®-chipped machine. Third, and perhaps most importantly, the software was to be taken forwards with an experimental version in which new ideas could be tried without detriment to users requiring an established/stable version (*WinBUGS*). This has led to a version that is much more amenable to extension in terms of its range of updating algorithms. Indeed this development highlights one of the fundamental differences between *WinBUGS* and OpenBUGS. Both versions classify the full conditional distribution of each node to be updated, but whereas *WinBUGS* uses a one-to-one mapping of classification to updating method, OpenBUGS works through a list of available updating methods, typically in order of increasing generality, and allocates methods to all nodes for which the method is appropriate. This latter approach can cope far better with an ever-expanding range of methods, which fits well with an open-source philosophy. Another significant outcome of this research is that infrastructure for parallelizing the software now exists.

Meanwhile, the development of *WinBUGS* at Imperial College was centred more and more around specific applications. For example, reversible jump methods were implemented for a generic class of problems [52], but the main focus of this work became applications in genetic epidemiology [53] and pharmacokinetics [54]. In addition, the differential equation solving capabilities were strengthened to accommodate more complex models in pharmacokinetics and make them more applicable to infectious disease modelling. Hence, the two versions of the software diverged somewhat, each with their own advanced features unavailable in the other. At the time of writing, we are moving to rectify this and facilitate the migration of *WinBUGS*' advanced features over to the OpenBUGS framework, which represents the future of the project, as discussed in Section 6, although *WinBUGS* will always remain available as a stable version for routine use.

4.3. Funding history

In almost two decades, the BUGS project has received a total of ~1.3 million U.K. pounds in direct funding, as well as indirect long-term support for senior staff, including 5 years funding from The Academy of Finland for AT to work on OpenBUGS. As outlined in Table I, the U.K. grant funding has come from three of the U.K.'s research councils (Medical Research Council, Economic and Social Research Council, Engineering and Physical Sciences Research Council) as well as the Wellcome Trust and a pharmaceutical company (AstraZeneca). While most grants have been targeted at a specific development goal, each has contributed generally to the overall

Table I. U.K. funding history of the BUGS project.

Time-scale	Funding body	Main purpose
1989–1993	MRC	Salary for programmer (AT)
1993–1996	ESRC	Analysis of complex longitudinal data
1996–1999	EPSRC	Development of PKBugs (pharmacokinetics)
1997–1999	ESRC	Development of GeoBUGS (spatial modelling)
1999–2004	MRC	Modelling complexity in biomedical research
2003–2004	AstraZeneca	Development of differential equation interface
2003–2006	MRC	Development of tools for handling genetic data
2006–2007	Wellcome Trust	Consolidation of previous work + dissemination

evolution of the software. The project now attracts some core funding from MRC via DS's and DL's work at the Biostatistics Unit in Cambridge, which is currently focussed on building an appropriate infrastructure for the further development of OpenBUGS (see Section 6).

5. CRITICAL APPRAISAL

In this section we attempt to identify the reasons for BUGS' success. We also show how the same positive aspects of its design also lead to some of its main shortcomings. The aim is to provide some insight into the origins of BUGS' various idiosyncrasies.

5.1. Flexibility

It seems that the most fundamental reason behind BUGS' success is its flexibility and generality. The software's Gibbs sampling scheme is designed so that it can be applied to any DAG and, in principle, any directed graph can be specified using the BUGS language, since the language is entirely open-ended and new distributions and functions can thus be added as required. We delve deeper into how these things are possible below but we first highlight the issues that they cause:

- Any model can be specified, regardless of whether it makes any sense. For example, Bernoulli trials cannot be over-dispersed as the variance is determined by the mean, but there is nothing to stop people fitting a model such as

```
y[i] ~ dbern(mu[i])
logit(mu[i]) <- dnorm(m.mu, p.mu)
```

and getting meaningless results. We have seen publications that do this. Similarly, we have frequently seen 'pointless' prior distributions being given, for example

```
beta ~ dnorm(m.beta, p.beta)
m.beta ~ dnorm(0, 0.0001)
p.beta ~ dgamma(0.001, 0.001)
```

instead of directly specifying values for `m.beta` and `p.beta`: this construction does induce some sort of prior on `beta` but adds a completely superfluous layer to the model.

In particular, there is no mechanism in place to ensure that a given model is identifiable. Such mechanisms may seem desirable but any attempt to police use of the language would invariably

restrict freedom, which goes against the philosophy of BUGS somewhat. (The ingenuity of users is such that there will always be models that we could not have anticipated when defining the rules.) It can be quite reasonable, for example, to deliberately over-parameterize hierarchical models, making certain parameters non-identifiable but greatly improving the efficiency of estimation for identifiable contrasts (see [19, Chapter 19], for instance). For example, in a repeated measures model with data $y_{ij} \sim N(\mu + u_i, \sigma^2)$ and $u_i \sim N(0, \phi^2)$, we can rescale the second-stage random effects by an unknown (and non-identifiable) factor, $u_i = \alpha b_i$, say, which can prevent the Gibbs sampler getting stuck.

- Similarly, the BUGS language does not discriminate between models that the software is good at fitting and those for which it may perform poorly. For example, the language easily accommodates autocorrelated time-series models, even those with longitudinal structure on the variances, but the algorithms selected to update the parameters of such models can often struggle. More generally, there is usually very little contextual information to be gleaned from the model specification, and BUGS' MCMC algorithms are required to be reasonably general to cope with the vast array of possible models. These facts together mean that it is very difficult to fully optimize the updating process for a given model: BUGS will almost always run slower than bespoke code.
- The same model can typically be specified in different ways, sometimes leading to (dramatically) different behaviours. A good example of this, which also highlights the previous point, is the autocorrelated model with data $y_i \sim N(\mu_i, \sigma^2)$, $i = 1, \dots, N$, and $\mu_i \sim N(f(\mu_{i-1}), \phi^2)$ for some appropriate function $f(\cdot)$. We can specify this either as written, which gives a *hierarchically centred* model, or, equivalently, via $\mu_i = f(\mu_{i-1}) + u_i$ with $u_i \sim N(0, \phi^2)$. In cases where $f(\cdot)$ contains more than one instance of μ_{i-1} , e.g. $f(\mu_{i-1}) = \beta \mu_{i-1} / (1 - \mu_{i-1})$, then the latter approach, with anything more than a few data ($N > 30$, say), will cause BUGS to 'hang' when building the internal representation of the model. We still do not fully understand why this happens, although BUGS is presumably attempting to construct an excessively complex graph, where the excess complexity grows exponentially with N . (Note that it is usually building complex/large graphs that causes BUGS to hang during compilation of the model.) We have established, however, that both models appear to give the same results for small N , although the hierarchically centred model runs much faster and has far better convergence properties.
- The range of possibilities afforded by the BUGS language means that there are certain modelling scenarios that occur only very rarely, and so bugs in the source code may go unnoticed for long periods of time. For example, one error in updating a Pareto distribution (with uniform likelihood) took over 15 years to discover. Bizarrely, someone else noticed the same error within the same week!

It is also worth noting here that BUGS provides much flexibility with respect to the way in which the simulation is carried out, for example, imposing no limits on the number of Gibbs iterations required and very few on the range of samples that can be used for inference. Hence, the software can be used in a non-rigorous manner for performing exploratory analyses, say. The downside of this, of course, is that the user has to take responsibility for convergence-related decisions, and so, to some extent, they need to have a reasonable understanding of the underlying methodology. BUGS is also flexible in allowing the user to choose which quantities should be monitored during the simulation. This is actually out of necessity since all monitored samples are stored in the computer's RAM for quick access, and to store all values sampled for every variable in a complex

model would soon lead to the computer running out of memory. Hence, the user must again be careful, to monitor only quantities of interest. (Note that more memory-efficient monitors are available for nodes where approximate inference is sufficient.)

5.2. Popularity

The popularity brought about largely by BUGS' generality/flexibility is itself partly responsible for the viral spreading of certain poor practices throughout Bayesian statistics. In particular, such widespread use has led to a degree of ubiquity of potentially inappropriate priors, through, ironically, a lack of exploitation of the software's flexibility. Many people enter the world of Bayesian ideas through WinBUGS, and novice users naturally tend to copy more experienced users' code; hence, 'lazy' priors are easily perpetuated. Examples include:

- The $\text{Gamma}(\varepsilon, \varepsilon)$ prior may be reasonable as a prior for precision parameters for observations (it is approximately the Jeffreys prior which might be better approximated by a uniform prior on the log-variance), but is generally inappropriate for the precision of random effects, particularly for continuous data. It arose in a time when the mathematical convenience that it offered was valuable but, with modern MCMC methods, this is no longer the case. However, it still survives due to early use in our own BUGS examples.
- It is easy to fit huge models to sparse data, and have all the conclusions driven unwittingly by casually made prior assumptions.
- The $\text{I}(\cdot, \cdot)$ mechanism for dealing with censored data can be used for modelling truncated prior distributions *provided they have no unknown parameters*—if there are unknown parameters the inferences will be wrong.

All these errors can be made, and results obtained, without any warning from the program. Sometimes lack of convergence can give a hint that something is wrong, but generally it is the user's (generally untrained) responsibility to identify the problem, which is why the manual is headed by a large public-health warning that '*Gibbs sampling can be dangerous!*'

Another factor contributing to BUGS' popularity is the facility to automatically compute the *Deviance Information Criterion* (DIC), which offers a straightforward means of comparing different models. This can be considered an adaptation of the Akaike Information Criterion to Bayesian models incorporating prior information, whether through fully-specified proper prior distributions or hierarchical models. It features two elements: the posterior mean of the deviance as a measure of fit, and an assessment p_D of the 'effective number of parameters' as a measure of complexity. Adding these components gives the DIC, and the model with the smallest DIC might be considered preferable in the sense of being likely to make better short-term predictions on data of a similar structure to that already observed.

DIC has proved to be very popular, and the original paper [55] is currently the third highest-cited paper in the whole of the mathematical sciences in the last 10 years. But it is not without problems. The method of assessing p_D is not invariant to transformations of the parameters and it can give inappropriate results if there are highly non-normal posterior distributions of the parameters on which prior distributions have been placed. Alternative forms of p_D have been suggested, and the theoretical basis for DIC has been questioned, in particular a suggestion that there is insufficient penalization for complexity [56]. Furthermore, it is still unclear what an appropriate procedure for categorical parameters, as in mixture models [57], might be. We consider DIC as a useful tool, but its implementation could be improved and perhaps guidance for its appropriate use built into the program.

5.3. Use of object-orientation

The software's ability to update any DAG is largely due to its exploitation of OOP. Put very simply, OOP facilitates abstract computation and so the general factorization properties of DAGs can be fully exploited. Even when writing code for the Gibbs scheme we only need to think in abstract terms, of visiting each stochastic node in the graph and navigating to its children, asking each node in turn for its contribution to the relevant full conditional (or its parameters). We need not care here about the nature of each node's distributional or functional form—responsibility for providing the correct information rests with the node objects themselves, which from the point-of-view of the Gibbs scheme are identical, but have much specialized detail hidden within.

Use of OOP and related ideas also makes it possible for the BUGS language to be open ended. Because the software works at an abstract level, it can be designed to make use of modelling components (distributions and functions) without having to know anything about them. Hence, new functions and distributions can be added at any time, and because the software can load new components *on demand*, it does not even need recompiling when new components are added. The same is true with new updating algorithms, which has greatly facilitated the expansion of BUGS' capabilities over the years.

One downside of object-orientation in systems as complex as BUGS is that it can make debugging somewhat difficult. In more traditional programming settings, there would typically be some kind of master program specifying a list of commands to be executed in sequence. A debugger can then be implemented to step through each command in turn and provide current information on parameters of interest. In contrast, objects are generally designed only to make/respond to various requests of/from other objects—they have little or no concept of the order in which external events occur. Consequently, an object-oriented program typically behaves as a complex chain-reaction and the path back to the source of any error may be extremely convoluted. This can make for indecipherable error messages, error messages that have nothing to do with the actual error, and impenetrable 'trap' windows, all of which even the developers of the software sometimes struggle to understand. For example, an 'undefined real result' trap lists the sequence of 'function calls' that led to some impossible calculation, such as the square root of a negative number. The calculation in question may have been attempted by an updater object trying to evaluate some function of interest. To figure out what the updater was trying to do, which node it was dealing with at the time, what the current values of that node and those in its Markov blanket are, and why these might have led to some inconsistency (even if it were that simple), just from a list of function calls and the current values of some 'internal' parameters, is a challenge indeed. As developers of the software we have access to special tools for probing these 'traps' for more information but it can still be difficult sometimes. It has been particularly difficult to fix any of these traps, or replace them with more meaningful error messages, since they represent unanticipated errors—things that were not *meant* to go wrong—it is virtually impossible to get hold of enough contextual information to provide a message that might be of some use to the user.

It should be noted at this point that the object-oriented framework selected for developing BUGS was perhaps, in retrospect, not the best choice. A natural extension of object-orientation is *component*-orientation, whereby the development of distinct software components that can be deployed independently is facilitated; for example, a graphics/equation editor or a spell-checker that can be used from within many other packages, such as a word processor or presentation software. The Component Pascal framework underlying BUGS is component-oriented (as the name

suggests) but was dominated by Java, its nearest competitor, which offered the (vastly) more popular C-style syntax as well as platform independence. However, at the time, Java could not support serious number-crunching and it was thought that with the explosion in personal computing, Unix might become substantially less popular, somewhat lessening the need for platform independence. Component-oriented systems are best exploitable when there is a large user-base and one can make use of other developers' components (reusing code as opposed to having to continually reinvent the wheel). The increased uptake of Java made it even more popular and Component Pascal is now somewhat obsolete (although much loved by its users). This has led to substantial problems in terms of capacity building, i.e. encouraging contributions to BUGS from third-parties as well as recruiting new developers.

5.4. Context-independence of modelling components

Both object- and component-orientation facilitate the development of components that can be used in arbitrary settings. Such context-independence has contributed enormously to the generality/flexibility of BUGS. For example, the majority of distributions in BUGS can be used as priors, likelihood, predictive distributions, and, with the use of the `cut` function, 'distributional constants' [58]; they may also be used at any level of a (hierarchical, say) model. However, the freedom that context-independence offers also creates scope for (unintentional) misuse when components have been designed imperfectly. Examples include:

- Certain distributions can only be used in particular contexts; for example, neither the Wishart nor the Dirichlet distribution, in WinBUGS, should be used as a likelihood, i.e. it is not possible to learn about their parameters (although a trick does exist in the case of the Dirichlet). The user is free to attempt to use such distributions incorrectly, and the software *should* generate a warning, but with the vast array of possibilities there is no guarantee that all such misuses will be caught; this, of course, could lead to unreliable results (without anyone being aware). In a similar vein, the Poisson distribution was once 'upgraded' (for reasons that have been lost in the mists of time) so that it could be used as a prior for continuous quantities, e.g. non-integer-valued data. However, several users have pointed out that in some settings this can lead to non-integer values being sampled for discrete quantities for which a conventional (i.e. discrete) Poisson prior was intended.
- Context independence allows the various algorithms used in BUGS to be switched around, e.g. Metropolis could be used instead of rejection sampling for certain models. This, however, was not fully anticipated in the design of WinBUGS. For the sake of efficiency, many of the samplers used in WinBUGS make assumptions about the context in which they are being used, which do not necessarily hold if applied in alternative settings. For example, a Metropolis updater will assume that any real number is a suitable candidate value for the variable being updated, because the updater was originally designed only for distributions with support on the whole real line. If it is then applied to a node whose full conditional has restricted support, the node may calculate its full conditional density incorrectly (since it assumes, for efficiency, that it will not receive inappropriate input), potentially leading to acceptance of a prohibited value. OpenBUGS, on the other hand, has been designed from a fully component-oriented perspective so that MCMC samplers can be used in arbitrary settings.

6. DISCUSSION AND FUTURE PERSPECTIVES

The current thrust of the BUGS project is towards focussing *all* development efforts on OpenBUGS. At the time of writing there exist only two people, worldwide, with an intimate understanding of the BUGS architecture (AT+DL). This architecture is complex and to simply download the source code and learn how to develop it would be a formidable task. There are numerous researchers and organizations willing to contribute, but to invest the time and resources required just to get to the point of being able to start development work is perhaps not economically viable. To circumvent this problem we are in the process of setting up a charitable foundation to act as a conduit for funding of development work, whereby all funds can be centralized and invested in a small number of developers who can be trained in the internal workings of BUGS and subsequently be responsible for multiple developments. The foundation will have a board of directors responsible for prioritizing proposed developments, which will all be for the common good of users (as opposed to specialized ventures, though we would actively encourage such developments 'externally').

For most of its life BUGS has enjoyed a relatively competition-free environment. MLwiN [59] has provided MCMC facilities for some years now, but only for a restricted class of models. Indeed, for more complex models, MLwiN will actually write a BUGS script. More recently, however, more general-purpose contenders have emerged. In particular, Just Another Gibbs Sampler (JAGS [60]) is a 'BUGS clone', which is implemented in C++ and will therefore run on many platforms as well as being more amenable to user-extension. JAGS has also highlighted and corrected some potential flaws in the BUGS language; for example, JAGS achieves better separation of modelling assumptions from data, with respect to censored observations, for instance. More competition arises from Microsoft's 'Infer.NET' package (<http://research.microsoft.com/en-us/um/cambridge/projects/infernet/>), which is a very general graphical modelling package, supporting many types of graph (directed or undirected) and providing numerous algorithms for inference. However, we anticipate that the focus here will be more on efficient processing of large problems using adequate approximations within a machine-learning context, as opposed to precise statistical inference requiring accurate assessment of second-order (error) properties. A Bayesian MCMC package has also recently been implemented within the SAS software [61].

Meanwhile, there is an ever expanding range of interfaces for using BUGS from other software, including Excel, GenStat, Matlab, R, SAS and Stata. There are also numerous packages to aid in converting BUGS input/output from/to ArcGIS, ArcView, Emacs, Excel, Mathematica, Matlab, R/S and SAS. In addition, there exist packages for automatically generating and running scripts using Perl and Python. And several specialized interfaces for running complex models in WinBUGS are available from www.winbugs-development.org.uk. These include facilities for implementing pharmacokinetic models, reversible jump methods, differential equation models, and for writing your own specialized function/distribution (to be incorporated into the BUGS language).

Despite several shortcomings, such as the lack of IF-THEN-ELSE logic and a somewhat cumbersome mechanism for specifying censoring, the BUGS language has proved enormously successful. It can describe effectively a vast array of different model types, even when the software is not well suited to analysing them. For example, the language can accommodate many time-series (autoregressive) models but Gibbs sampling is notoriously problematic in these settings, with excessive autocorrelation typically apparent in the simulated chains. As BUGS continues to evolve, it is mainly the algorithms implemented in the software that are changing, to offer a more

tailored solution to each individual problem, as opposed to the language itself, which we hope will, with some tweaks, continue long into the future, as a general means of describing graphical models.

REFERENCES

1. Gilks WR, Thomas A, Spiegelhalter DJ. A language and program for complex Bayesian modelling. *The Statistician* 1994; **43**:169–178.
2. Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis* 1984; **6**:721–741.
3. Gelfand AE, Smith AFM. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 1990; **85**:398–409.
4. Lawson AB, Browne WJ, Vidal Rodeiro CL. *Disease Mapping with WinBUGS and MLwiN*. Statistics in Practice. Wiley: Chichester, 2004.
5. Mu S, Ludden TM. Estimation of population pharmacokinetic parameters in the presence of non-compliance. *Journal of Pharmacokinetics and Pharmacodynamics* 2003; **30**:53–81.
6. McCarthy MA, Parris KM. Clarifying the effect of toe clipping on frogs with Bayesian statistics. *Journal of Applied Ecology* 2004; **41**:780–786.
7. O'Hagan A, Stevens JW, Montmartin J. Bayesian cost-effectiveness analysis from clinical trial data. *Statistics in Medicine* 2001; **20**:733–753.
8. Scurrah KL, Palmer LJ, Burton PR. Variance components analysis for pedigree-based censored survival data using generalized linear mixed models (GLMMs) and Gibbs sampling in BUGS. *Genetic Epidemiology* 2000; **19**:127–148.
9. Millard A, Gowland R. A Bayesian approach to the estimation of age of humans from toothwear. *Archeologia e Calcolatori* 2002; **13**:197–210.
10. Eaves L, Erkanli A, Silberg J, Angold A, Maes HH, Foley D. Application of Bayesian inference using Gibbs sampling to item-response theory modeling of multi-symptom genetic data. *Behavior Genetics* 2005; **35**:765–780.
11. Milheiro-Oliveira P. Bayesian statistical methods for modeling and prediction of major landslides in coastal cliffs. *Coastal Engineering Journal* 2007; **49**:45–61.
12. Ayers E, Junker B. IRT modeling of tutor performance to predict end-of-year exam scores. *Educational and Psychological Measurement* 2008; **68**:972–987.
13. Chevrolat J-P, Golmard J-L, Ammar S, Jouvent R, Boisvieux J-F. Modelling behavioral syndromes using Bayesian networks. *Artificial Intelligence in Medicine* 1998; **14**:259–277.
14. Meyer R, Yu J. BUGS for a Bayesian analysis of stochastic volatility models. *Econometrics Journal* 2000; **3**:198–215.
15. Weihs C, Ligges U. Parameter optimization in automatic transcription of music. In *From Data and Information Analysis to Knowledge Engineering: Proceedings of the 29th Annual Conference of the Gesellschaft für Klassifikation e.V. University of Magdeburg*, 9–11 March 2005, Spiliopoulou M, Kruse R, Borgelt C, Nürnberger A, Gaul W (eds), Studies in Classification, Data Analysis, and Knowledge Organization. Springer: Berlin, 2006; 740–747.
16. Swartz TB, Gill PS, Beaudoin D, deSilva BM. Optimal batting orders in one-day cricket. *Computers and Operations Research* 2006; **33**:1939–1950.
17. Wyatt RJ. Mapping the abundance of riverine fish populations: integrating hierarchical Bayesian models with a geographic information system (GIS). *Canadian Journal of Fisheries and Aquatic Sciences* 2003; **60**:997–1006.
18. Katsis A, Ntzoufras I. Bayesian hypothesis testing for the distribution of insurance claim counts using the Gibbs sampler. *Journal of Computational Methods in Science and Engineering* 2005; **5**:201–214.
19. Gelman A, Hill J. *Data Analysis using Regression and Multilevel/Hierarchical Models*. Cambridge University Press: Cambridge, 2007.
20. Carlin BP, Louis TA. *Bayesian Methods for Data Analysis* (3rd edn). CRC Press: Boca Raton, 2008.
21. Congdon P. *Bayesian Statistical Modelling*. Wiley: Chichester, 2001.
22. Lawson AB. *Bayesian Disease Mapping: Hierarchical Modeling in Spatial Epidemiology*. CRC Press: Boca Raton, 2009.
23. Woodworth GG. *Biostatistics: A Bayesian Introduction*. Wiley: New York, 2004.
24. Broemeling LD. *Bayesian Biostatistics and Diagnostic Medicine*. CRC Press: Boca Raton, 2007.

25. Ntzoufras I. *Bayesian Modeling using WinBUGS*. Wiley: New York, 2009.
26. Gill J. *Bayesian Methods: A Social and Behavioral Sciences Approach*. Chapman & Hall, CRC: London, Boca Raton, 2002.
27. Spiegelhalter DJ. Bayesian graphical modelling: a case-study in monitoring health outcomes. *Applied Statistics* 1998; **47**:115–133.
28. Stefik M, Bobrow DG. Object-oriented programming: themes and variations. *The AI Magazine* 1985; **6**(4):40–62.
29. Lauritzen SL, Dawid AP, Larsen BN, Leimer HG. Independence properties of directed Markov fields. *Networks* 1990; **20**:491–505.
30. Szyperski C. Component-oriented programming: a refined variation of object-oriented programming. *The Oberon Tribune* 1995; **1**:1–5.
31. Spiegelhalter DJ, Knill-Jones RP. Statistical and knowledge-based approaches to clinical decision-support systems, with an application in gastroenterology (with Discussion). *Journal of the Royal Statistical Society, Series A* 1984; **147**:35–77.
32. Pearl J. Reverend Bayes on inference engines: a distributed hierarchical approach. *Proceedings of the American Association for Artificial Intelligence National Conference on AI*, Pittsburgh, 1982; 133–136.
33. Spiegelhalter DJ. Probabilistic reasoning in predictive expert systems. In *Uncertainty in Artificial Intelligence*, Kanal LN, Lemmer J (eds). North-Holland: Amsterdam, 1986; 47–68.
34. Kim JH, Pearl J. A computational model for combined causal and diagnostic reasoning in inference systems. *Proceedings of the IJCAI-83*, Karlsruhe, Germany, 1983; 190–193.
35. Lauritzen SL, Spiegelhalter DJ. Local computations with probabilities on graphical structures and their application to expert systems (with Discussion). *Journal of the Royal Statistical Society, Series B* 1988; **50**:157–224.
36. Pearl J. Evidential reasoning using stochastic simulation. *Artificial Intelligence* 1987; **32**:245–257.
37. Spiegelhalter DJ, Lauritzen SL. Sequential updating of conditional probabilities on directed graphical structures. 1990; 20:579–605.
38. Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing* 2000; **10**:325–337.
39. Wirth N. *Programming in Modula-2*. Springer: Santa Clara, 1983.
40. Thomas A, Spiegelhalter DJ, Gilks WR. BUGS: a program to perform Bayesian inference using Gibbs sampling. In *Bayesian Statistics*, Bernardo JM, Berger JO, Dawid AP, Smith AFM (eds), vol. 4. Oxford University Press: Oxford, U.K., 1992; 837–842.
41. Gilks WR, Richardson S, Spiegelhalter DJ. *Markov Chain Monte Carlo in Practice*. Chapman & Hall: London, 1996.
42. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equations of state calculations by fast computing machines. *Journal of Chemical Physics* 1953; **21**:1087–1091.
43. Hastings WK. Monte Carlo sampling-based methods using Markov chains and their applications. *Biometrika* 1970; **57**:97–109.
44. Oberon Microsystems Inc. Component Pascal Language Report. *Technical Report*, Oberon Microsystems Inc., Zurich, 2001.
45. Jones R, Lins R. *Garbage Collection: Algorithms for Automatic Dynamic Memory Management*. Wiley: New York, 1996.
46. Gelman A, Gilks WR, Roberts GO. Weak convergence and optimal scaling of random walk Metropolis algorithms. *Annals of Applied Probability* 1997; **7**:110–120.
47. Neal RM. Markov chain Monte Carlo methods based on ‘slicing’ the density function. *Technical Report 9722*, Department of Statistics, University of Toronto, 1997.
48. Hethcote HW. The mathematics of infectious diseases. *SIAM Review* 2000; **42**:599–653.
49. Jones DS, Sleeman BD. *Differential Equations and Mathematical Biology*. Chapman & Hall: Boca Raton, 2003.
50. Lunn DJ. Bayesian analysis of population pharmacokinetic/pharmacodynamic models. In *Probabilistic Modeling in Bioinformatics and Medical Informatics*, Husmeier D, Dybowski R, Roberts S (eds). Springer: London, 2005; 351–370.
51. Lunn DJ. WinBUGS Development Interface (WBDev). *ISBA Bulletin* 2003; **10**(3):10–11.
52. Lunn DJ, Best N, Whittaker JC. Generic reversible jump MCMC using graphical models. *Statistics and Computing* 2008; DOI: 10.1007/s11222-008-9100-0.
53. Lunn DJ, Whittaker JC, Best N. A Bayesian toolkit for genetic association studies. *Genetic Epidemiology* 2006; **30**:231–247.
54. Lunn DJ. Automated covariate selection and Bayesian model averaging in population PK/PD models. *Journal of Pharmacokinetics and Pharmacodynamics* 2008; **35**:85–100.

55. Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A. Bayesian measures of model complexity and fit (with Discussion). *Journal of the Royal Statistical Society, Series B* 2002; **64**:583–639.
56. Plummer M. Penalized loss functions for Bayesian model comparison. *Biostatistics* 2008; **9**:523–539.
57. Celeux G. Mixture models for classification. In *Advances in Data Analysis*, Decker R, Lenz HJ (eds). Springer: Berlin, 2007; 3–14.
58. Lunn D, Best N, Spiegelhalter D, Graham G, Neuenschwander B. Combining MCMC with ‘sequential’ PKPD modelling. *Journal of Pharmacokinetics and Pharmacodynamics* 2009; DOI: 10.1007/s10928-008-9109-1.
59. Rasbash J, Browne W, Goldstein H, Yang M, Plewis I, Healy M, Woodhouse G, Draper D, Langford I, Lewis T. *A User's Guide to MLwiN* (2nd edn). Institute of Education: London, 2000.
60. Plummer M. *JAGS Version 1.0.3 Manual*. IARC: Lyon, France, 2008.
61. SAS Institute Inc. *SAS/STAT 9.2 User's Guide*. SAS Institute Inc.: Cary, NC, 2008.

COMMENTARY

Comments on ‘The BUGS project: Evolution, critique and future directions’

Mary Kathryn Cowles*,[†],[‡]

Department of Statistics and Actuarial Science, The University of Iowa, Iowa City, IA, U.S.A.

In this age of completely computer-based statistical analysis, a statistical method becomes accepted only when accessible software to implement it has stood the test of time. WinBUGS has been a primary vehicle to bring Bayesian modeling into the mainstream in business, industry, and government, as well as academic research.

The paper tracks the evolution of BUGS through two decades characterized by rapid-fire changes in hardware and operating systems, in MCMC algorithms, and in both the magnitude of data sets to be analyzed and the complexity of Bayesian models to be fit to them. Beginning with the original, command-based version for the Unix and DOS operating systems, BUGS then became WinBUGS, with a more interactive user interface, new updating algorithms enabling the fitting of much wider classes of models, and built-in graphical capabilities.

The authors straightforwardly discuss both the advantages and the down-sides of choices that were made in the transition to WinBUGS, particularly the object-oriented design. While affording flexibility and extensibility, this characteristic is also responsible for the infamous WinBUGS ‘traps’—unintelligible error messages that give the user next to no hint of where the problem actually lies.

The paper does not emphasize a WinBUGS offering that partially offsets the poor error messages, namely the excellent user manual accessible within WinBUGS from a drop-down menu. This manual documents all aspects of WinBUGS use and provides a worked example in a tutorial framework.

In a welcome move, in 2004, members of the WinBUGS team began the development of OpenBUGS, an open-source version of WinBUGS. The first aim of this effort—to make it easy to interface OpenBUGS with other statistical software, for both data input and output postprocessing—has been satisfied neatly in the Windows environment. In particular, the OpenBUGS team developed the R [1] package BRugs [2], which enables the user to prepare data and initial values in R, pass them and the location of a model-specification file to OpenBUGS, and collect the MCMC samples produced for graphical and numeric analysis in R.

*Correspondence to: Mary Kathryn Cowles, Department of Statistics and Actuarial Science, The University of Iowa, Iowa City, IA, U.S.A.

[†]E-mail: kcowles@stat.uiowa.edu

[‡]Associate Professor.

In the short-term future, I would hope to see more attention given to the second-stated aim of OpenBUGS—to make the software platform-independent. The Linux-based research computing environment is very popular in academic statistical research, in large part because R [1] provides excellent statistical computing and graphics and interfaces easily with LaTeX [3] for document preparation. Furthermore, the current economic climate is making Linux attractive to businesses as a lower-cost computing environment. Thus, a fully functional and fully supported version of OpenBUGS for Linux would substantially expand the use of OpenBUGS among academic statisticians, and perhaps in business as well. Indeed a version of BUGS for Linux—linbugs—does exist, and I am successfully running it under the Ubuntu 9.04 distribution of Linux. Linbugs is an implementation for Linux of ClassicBUGS, a command line version of BUGS that has no graphical or convergence diagnostic capabilities. Linbugs seems to be able to fit the same range of models that WinBUGS can, including spatial models. As of this writing, the OpenBUGS developers have not yet gotten the BRugs R package to work with linbugs. However, linbugs can write the MCMC samples it produces out to text files that can then easily be read by the R package coda [4], which provides a wide range of tools for graphical and numeric convergence assessment and posterior analysis. I hope to see the OpenBUGS team begin to provide comprehensive support and development for linbugs, beginning with compilation instructions for various Linux distributions.

The third stated aim of the OpenBUGS project is to provide an ‘experimental version in which new ideas could be tried without detriment to users requiring an established/stable version (WinBUGS)’. An exciting prospect mentioned in this regard is the potential for parallelizing OpenBUGS, a development that should make the use of OpenBUGS for larger data sets, for example in spatial and environmental statistics, more feasible.

I hope the BUGS team will also consider creating a web-based version of OpenBUGS, perhaps along the lines of Rweb [5] for R.

Many thanks to the indefatigable BUGS team for 20 years of work in providing software for Bayesian research and production data analysis, and best wishes for their charitable foundation and continuing development of BUGS.

REFERENCES

1. R Development Core Team. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, 2008. Available from: <http://www.R-project.org>.
2. Thomas A, O’Hara B, Ligges U, Sturtz S. Making BUGS open. *R News* 2006; **6**(1):12–17.
3. Lamport L. *LaTeX: A Document Preparation System*. Addison-Wesley: Reading, MA, 1986.
4. Plummer M, Best N, Cowles MK, Vines K. CODA: output analysis and diagnostics for MCMC. R package version 0.13-3, 2008.
5. Banfield J. Statistical analysis on the Web. Available from: <http://www.math.montana.edu/Rweb/>.

DISCLAIMER

I was one of the initial developers of the coda package back in the mid-1990s. Others have since improved and expanded it, and ported it to R.

COMMENTARY

Comments on ‘The BUGS project: Evolution, critique and future directions’[‡]

Andrew Gelman^{*,†}

Department of Statistics and Department of Political Science, Columbia University, New York, NY, U.S.A.

I first saw BUGS in a demonstration version at a conference in 1991, but I didn’t take it seriously until over a decade later, when I found that some of my Ph.D. students in political science were using Bugs to fit their models. It turned out that Bugs’s modeling language was ideal for students who wanted to fit complex models but didn’t have a full grasp of the mathematics of likelihood functions, let alone Bayesian inference and integration. I also learned that the modular structure of BUGS was a great way for students, and researchers in general, to think more about modeling and less about deciding which conventional structure should be fit to data.

Since then, my enthusiasm for BUGS has waxed and waned, depending on what sorts of problems I was working on. For example, in our study of income and voting in U.S. states [1], my colleagues fit all our models in BUGS. Meanwhile we kept running into difficulty when we tried to expand our model in different ways, most notably when going from varying-intercept multilevel regressions, to varying-intercept, varying-slope regressions, to models with more than two varying coefficients per group. Around this time I discovered lmer [2], a function in R which fits multilevel linear and generalized linear models allowing for varying intercepts and slopes. The lmer function can have convergence problems and does not account for uncertainty in the variance parameters, but it is faster than Bugs and in many cases more reliable—so much so that Jennifer Hill and I retooled our book on multilevel models to foreground lmer and de-emphasize Bugs, using the latter more as a way of illustrating models than as a practical tool.

What does BUGS do best and what does it do worst? In short, BUGS excels with complicated models for small data sets. For example, we fit a fairly elaborate discrete choice model to data from a laboratory economics experiment with 30 trials on each of about 100 participants [3]. The sample size was small enough that we wanted to use a hierarchical model to obtain stable inference for all the people, so we programmed something up in BUGS and it worked right away.

Well, not right away. We first had to devise a work-around because BUGS was crashing on a logistic model. We added a couple lines in the model to bound the probabilities between 0.01 and

*Correspondence to: Andrew Gelman, Department of Statistics and Department of Political Science, Columbia University, New York, NY, U.S.A.

†E-mail: gelman@stat.columbia.edu

‡Comments on ‘The BUGS project: evolution, critique and future directions,’ by David Lunn, David Spiegelhalter, Andrew Thomas, and Nicky Best, for *Statistics in Medicine*.

0.99. As with many situations in which a specification is set up for computational reasons [4], it turned out that the model made statistical sense as well: bounding the probabilities allowed for the occasional outlier, an issue we had never previously thought about in the context of binary data (see [5] for a cleaner solution to this problem using the *t* distribution in place of the logistic). Anyway, the point is that in this example, BUGS worked well, it was the direct solution to a good answer, and even its problems led to a new solution which would actually have been difficult to implement in non-BUGS software.

When does BUGS not work so well? With large data sets, multivariate structures, and regression coefficients. For one thing, it can be difficult to keep track of regression coefficients in the BUGS modeling language, which allows no subroutines or macros.

For example, instead of:

```
for (i in 1:n){
  y[i] ~ dnorm (y.hat[i], tau.y)
  y.hat[i] <- a[county[i]] + b[county[i]]*x[i]
  e.y[i] <- y[i] - y.hat[i]
}
tau.y <- pow(sigma.y, -2)
sigma.y ~ dunif (0, 1000)
```

we want something like:

```
y ~ norm (a[county] + b[county]*x, sigma.y)
```

This is all algorithmic and could be programmed, but Bugs is not set up so as to allow this.

Convergence can be slow and, in fact, each iteration can take a long time to run. When generalizing our model of income and voting to allow four coefficients for each state (corresponding to individual income, religious attendance, their interaction, and a constant term), BUGS basically ground to a halt, to the extent that I wouldn't have trusted its results, even had I decided to be patient and run it all night.

Winston Churchill said that sometimes the truth is so precious, it must be attended by a bodyguard of lies. Similarly, for a model to be believed, it must, except in the simplest of cases, be accompanied by similar models that either give similar results or, if they differ, do so in a way that can be understood. In the example of voting by income and religion, we simply took a step back and fit the model using *lmer*. This didn't always work either, but when *lmer* had problems we stripped down the model in various ways until it worked. It ran fast enough that we were able to experiment.

What are the great things about BUGS?

1. It really works! I can use it in my applied research. (For software to be useful to me in this way, it doesn't have to work all the time, it only has to solve some problems that can't easily be solved in other ways, and to break down in recognizable ways, so that it doesn't pretend to work when it doesn't.)
2. BUGS is easy to use and to teach, with intuitive syntax.
3. It is free.
4. It can be called directly from R (see [6] for many illustrations of the advantages of running BUGS in this way, most notably for preprocessing of data and postprocessing of inferences).

What are some problems with BUGS?

1. It often needs lots of hand-holding adjustment of the model to work.
2. Efficiently programmed models can get really long, ugly, and bug-prone, as I have learned in implementing redundant parameterizations in hierarchical models [7].
3. You can't debug it by running interactively, as you can in R, Matlab, C, Fortran, etc.
4. There's no easy way to go inside and improve the sampling, either by writing shortcuts in the BUGS language or by inserting your own jumping rules inside the inference engine. Attempts to 'trick' BUGS into using efficient updates can sometimes backfire.
5. As far as I can tell, BUGS does not understand hierarchical models but rather thinks of scalar parameters as individual entities. This may be partly a historical issue—before the mid-1990s, hierarchical modeling was generally considered a special topic rather than central to concepts of Bayesian inference and prior information—but, whatever the reason, the current implementation of BUGS does not allow easy setup or fast computation with hierarchical models.

There is also a problem—not really the fault of BUGS itself—that its users typically focus on inference as the only goal, ignoring the other two stages of Bayesian data analysis: model building and model checking. Often it can take awhile to get the model to converge, and once that happens there's a tendency to just stop, relax, and present the results. I believe that an important future development should be automatic implementations of model checking using posterior simulations. This can be done—for example, by creating a 'shadow' replicated variable for each variable in the model, and then allowing user-specified or default graphical comparisons—but, like everything else, it requires work.

I shouldn't really complain—after all, BUGS is free and, as noted in the article under discussion, has had an incredibly beneficial influence, especially considering that it was put together by just a few dedicated people. I hope that the next 20 years of BUGS are as successful as the first.

REFERENCES

1. Gelman A, Shor B, Bafumi J, Park D. Rich state, poor state, red state, blue state: what's the matter with Connecticut? *Quarterly Journal of Political Science* 2007; **2**:345–367.
2. Bates D. Fitting linear models in R using the lme4 package. *R News* 2005; **5**(1):27–30.
3. Casella A, Gelman A, Palfrey TR. An experimental study of storable votes. *Games and Economic Behavior* 2005; **57**:123–154.
4. Gelman A. Parameterization and Bayesian modeling. *Journal of the American Statistical Association* 2004; **99**:537–545.
5. Liu C. Robit regression: a simple robust alternative to logistic and probit regression. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, Gelman A, Meng XL (eds). Wiley: London, 2004; 227–238.
6. Gelman A, Hill J. *Data Analysis using Regression and Multilevel/Hierarchical Models*. Cambridge University Press: Cambridge, 2007.
7. Gelman A, Huang Z, van Dyk D, Boscardin WJ. Using redundant parameters to fit hierarchical models. *Journal of Computational and Graphical Statistics* 2007; **17**:95–122.

COMMENTARY

Comments on ‘The BUGS project: Evolution, critique and future directions’

Martyn Plummer^{*,†}

International Agency for Research on Cancer, 150 Cours Albert Thomas, 69372 Lyon Cedex 08, France

BUGS has had a substantial impact on statistical practice in many application areas. This timely review could have appeared in any number of journals. But it is no coincidence that it appears in *Statistics in Medicine*. The driving force behind the development of BUGS throughout most of its 20-year history has been biostatistics. This is largely due to the people behind the BUGS project, but I also believe that the BUGS approach to modelling is highly suited to medical research.

Medical studies are constrained by relatively small sample sizes. At the same time they have many uncontrolled sources of variation that need to be accounted for. Standardized approaches may discard or hide some of the extra complexity in the data. BUGS invites us to take the opposite approach: instead of coercing the data into an existing model, the model is expanded to fit the data.

From this perspective, the original model becomes a *design motif* embedded in a larger graphical model. The success of the BUGS approach relies on the ability of the inference engine to recognize common design motifs and match them with appropriate methods. As noted by Lunn and colleagues, BUGS does not work so well when the engine has to fall back on generic sampling algorithms. Much of the reputation of Bayesian modelling for inefficiency stems from over-reliance on such methods.

It may be worth thinking about graphical models in terms of *patterns*, an idea that was first introduced in architecture [1], but has also found widespread use in software engineering [2]. A pattern describes a reusable solution to a common design problem. In graphical modelling, a pattern would be a solution to analyzing part of a large graphical model, which forms a recognizable design motif. A simple example of a pattern, which exploits small-scale features of the graph, is a conjugate updater for a single node. One of the key differences between WinBUGS and OpenBUGS is that the latter is capable of recognizing larger design motifs, such as a generalized linear models, and applying an appropriate block-updater.

Future development possibilities for BUGS are not limited to refining the inference engine. There is plenty of scope for using BUGS as a platform for exploring ideas in Bayesian modelling.

*Correspondence to: Martyn Plummer, International Agency for Research on Cancer, 150 Cours Albert Thomas, 69372 Lyon Cedex 08, France.

†E-mail: plummer@iarc.fr

The BUGS authors have already done this with the deviance information criterion and the cut function. Conversely, popular software packages can become vehicles for spreading practices of dubious rigour. Lunn and colleagues give some examples of ill-defined models that are all too familiar from my own experience as a reviewer. Clearly, the user base has expanded past the point where the current *caveat emptor* approach has any real impact. It would be enormously helpful if BUGS could provide some 'red flags' in such cases to alert the user to potential problems in the model definition. More generally, it is currently difficult to track the flow of information in a graphical model from the data to the parameters. I am sure that this is a rich area for future research which will eventually lead to improved understanding of the way complex models behave.

There are some technological challenges that are easy to identify, because the changes in computing practice are taking place now. Multi-core PCs are now standard on the desktop. The ability to exploit their parallel processing capabilities should be a priority. A second challenge is the move to 64-bit operating systems. These are currently used only in high-performance computing environments, but will become more widespread when demand for memory exceeds the limit imposed by the 32-bit operating systems, which limit the per-process user address space to 2 GB.

Future technological developments, on the other hand, are notoriously hard to predict. In the mid-1990s when the BUGS project focused exclusively on Windows, few people would have predicted that the Unix operating system would get a new lease of life in the forms of GNU/Linux and Mac OS X (although, to be fair, some did). The choice of Component Pascal and the Black Box component framework at that time was, I suppose, guided by aesthetic criteria, and Andrew Thomas's self-professed inability to program in C++ is a reaction to its lack of elegance and simplicity. From the perspective of 2009, it seems that these choices have had a detrimental lock-in effect but it is too early to be sure.

Technological factors are only part of the story. The human factor is much more important, especially now that BUGS has been established as an open source project. Just as the current success of BUGS is a product of its small but influential development team; hence, so its future depends on attracting stakeholders who are not just users of the software but are actively involved in its development. The influence of BUGS is undeniable, so this should be an attractive prospect for anyone wishing to improve the current practice in Bayesian modelling.

REFERENCES

1. Alexander C. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press: New York, 1977.
2. Gamma E, Helm R, Johnson R, Vlissides J. *Design Patterns: Elements of Reusable Object-oriented Software*. Addison-Wesley: Boston, MA, 1995.

COMMENTARY

Comments on ‘The BUGS project: Evolution, critique, and future directions’

Armando Teixeira-Pinto^{1,*},[†] and Sharon-Lise T. Normand^{2,3}

¹*Faculty of Medicine, Department of Biostatistics and Medical Informatics, CINTESIS, University of Porto, Porto, Portugal*

²*Harvard Medical School, Department of Health Care Policy, Boston, U.S.A.*

³*Harvard School of Public Health, Department of Biostatistics, Boston, U.S.A.*

KEY WORDS: graphical models; causal inference; prediction

1. INTRODUCTION

Lunn, Spiegelhalter, Thomas, and Best (henceforth Lunn *et al.*) [1] provide a comprehensive review of the BUGS software package in this issue. The authors describe how computational algorithms developed for graphical models were exploited for the purposes of Bayesian inference.

They discuss the software evolution from its early version for UNIX platforms to the development of WinBUGS, a stand-alone application running on Microsoft Windows and, more recently, the release of the opensource version OpenBUGS, which is described as the future direction for BUGS development. Finally, the authors provide a contemporary appraisal and comment on the impact of BUGS.

2. GRAPHICAL MODELS AS A BASIS FOR COMPUTATION

In their public introduction of BUGS at the Fifth Valencia International Meeting during June 1994 in Spain, Spiegelhalter, Thomas, and Best [2] described the use of graphical models to express conditional independence assumptions for the joint distribution of observed and unobserved variables. With its release, BUGS eliminated many problems associated with the implementation of

*Correspondence to: Armando Teixeira-Pinto, Faculty of Medicine, Department of Biostatistics and Medical Informatics, CINTESIS, University of Porto, Al. Prof. Hernani Monteiro, 4200-319 Porto, Portugal.

[†]E-mail: tpinto@med.up.pt, tpinto@post.harvard.edu

Contract/grant sponsor: National Institute of Mental Health; contract/grant number: R01-MH61434

Bayesian methods. Statistical researchers could actually apply Bayesian methods to real problems and, importantly, work collaboratively with other scientists using these methods.

2.1. *Early commentaries*

Several researchers provided early commentaries regarding the new software. Issues related to the comprehensive and extendability of the software were raised [3]. The software's lack of inclusion of multivariate nodes and restriction to log-concave distributions were criticized but these have since been remedied. Normand [3] also suggested using predictive distributions to assess model fit, referencing Dempster's idea [4] of calculating the predictive log-likelihood with respect to the posterior distribution of the model parameters. The deviance information criterion has become one of the main model summary statistics in Bayesian computation.

With the initial release of BUGS, many were concerned with model convergence [5] and this worry persists. Gelfand [6] specifically identified problems associated with specification of the hyper-parameter distributions for the precision parameters for random effects in particular. A related concern is the risk that the set of conditional distributions may be inconsistent [7]. These issues are as relevant today as they were more than two decades ago.

Finally, Fienberg [8] noted that BUGS ignored the causal nature of the inference associated with graphical model. The software designers replied that the use of the graph was to represent assumptions of conditional independence in order to facilitate the derivation of full conditional distributions. With increasing emphasis on causal inference by statisticians, creatively extending the software to address causal questions seems to be a particularly important and exciting new direction.

2.2. *Bugs 20 years later*

Lunn *et al.* [1] describe the impressive technological and platform changes made to BUGS. These enhancements have undoubtedly increased access to and use of Bayesian methods. The authors provide a fair assessment of the advantages and the limitations of BUGS. We had a few additional observations.

First, nothing is described about innovations in graphical modeling theory and how BUGS has capitalized (or intends to capitalize) on these innovations. With increased analyses of large and complex databases that are likely to include genetic information, demographic information, social factors, and so on, issues relating to model selection will become critical. Methods proposed by Drton and Perlman [9] provide model selection for Gaussian graphs—do the authors see any role for similar algorithms for graphs consisting of mixed nodes?

Second, a recurrent problem that is frequently reported by BUGS users and also addressed by the authors is the error reporting system. In fact, when something goes wrong during the Markov Chain Monte Carlo (MCMC) computation it is virtually impossible to identify the source of the error. The suggestions presented in the user's manual include different choice of initial values, more informative priors, and re-parameterization of the model. However, this trial-and-error approach quickly becomes untenable when the model has a certain level of complexity and the number of parameters is large. Users would greatly benefit from a comprehensive error reporting system.

Third, the dependence on Windows platform is another disadvantage. Some efforts are being made to permit Linux and MAC users to run WinBUGS on their computers but at the moment this is only possible through WINE, a freeware application that permits a Unix-like operating system to execute programs written for Microsoft Windows [10].

Table I. Some sample macros and packages that communicate with WinBUGS and OpenBUGS.

Software	Macro/package	Reference
SAS	RASmacro	Sparapani, R. Some SAS macros for BUGS data. Retrieved June 7, 2009, from http://www2.mcw.edu/pcor/bugs/tr049.pdf . Zhiyong Z. A modified version of Rodney Sparapani's SAS Macros. Retrieved June 7, 2009, from http://www.psychstat.org/us/article.php/61.htm
	WinBUGSio	Smith, MK, Richardson H. WinBUGSio: A SAS Macro for the Remote Execution of WinBUGS. <i>Journal of Statistical Software</i> 2007; 23 (9):1–10
	per cent WinBUGS	Zhang L. Bayesian data analysis using per cent WinBUGS. Retrieved June 7, 2009, from http://www.lexjansen.com/pharmasug/2008/sp/sp12.pdf
STATA	winbugsfromstata	Thompson J, Palmer T, Moreno S. Bayesian Analysis in Stata using WinBUGS. <i>The Stata Journal</i> 2006; 6 (4):530–549.
R&S - Plus	R2WinBUGS	Sturtz S, Ligges U and Gelman A. R2WinBUGS: A Package for running WinBUGS from R. <i>Journal of Statistical Software</i> 2005; 3 (12):1–16
	CODA	Plummer M, Best N, Cowles K, Vines K. CODA: Convergence diagnosis and output analysis for MCMC. <i>R news</i> 2006; 6 (1):7–11
	BOA	Smith B. boa: An R package for MCMC output convergence assessment and posterior inference. <i>Journal of Statistical Software</i> 2007; 21 (11):1–37

The lack of tools for data management and limited ability to read data in different formats makes the use of additional software for data preparation and exploration of the results inevitable. Several packages are available that allow other statistical software to communicate with BUGS in a transparent way to the user. These include, for example, SAS, R, S-Plus, and STATA (Table I). In our view, this synergy is beneficial for WinBUGS development because it not only allows more focus on improving the MCMC engine but also provides easy access to users of other statistical software, who, otherwise, would be reluctant to learn how to use one more program.

An important feature of WinBUGS not to be undervalued is the documentation built into the software. This documentation can be accessed through the Help menu and it is very comprehensive. The developers have included three volumes of over 40 worked examples based on real data problems properly documented and explained. These examples facilitate an understanding of the software and its versatility, and provide a valuable resource for teaching Bayesian methods in general.

3. CONCLUDING REMARKS

BUGS opened the door for applied Bayesian analyses in an unprecedented manner. In a recent review of the history of Bayesian statistics in Medicine, Ashby [11] referred to BUGS or WinBUGS more than 10 times in the article. Lunn *et al.* [1] indicated a Web search using Google with the word 'WinBUGS' generated over than 100 000 hits. In fact, by adding the word 'OpenBUGS' the number increases to 250 000 hits, which is an impressive number especially if compared with a competitor that gets 35 000 hits.

We thank Lunn *et al.* [1] for the opportunity to comment on their article, and more importantly, to express gratitude for making it possible to implement Bayesian analyses. An ancient Greek [12] commented ‘Not everyone can sail to Corinth’—Lunn *et al.* [1] have made this possible!

ACKNOWLEDGEMENTS

Dr Normand was supported in part by grant R01-MH61434 from the National Institute of Mental Health.

REFERENCES

1. Lunn D, Spiegelhalter D, Thomas A, Best N. The BUGS project: evolution, critique, and future directions. *Statistics in Medicine* 2009.
2. Spiegelhalter D, Thomas A, Best N. Computation on graphical models (with Discussion). In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 407–425.
3. Normand S-L. Discussion of ‘computation on Bayesian graphical models’. In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 420–421.
4. Dempster AP. The direct use of likelihood for significance testing. In *Proceedings of the Conference on Foundational Questions in Statistical Inference*, Barndorff-Nielsen O, Blaeslid P, Shou G (eds), University of Aarhus, 1974; 335–352.
5. Carlin B. Discussion of ‘computation on Bayesian graphical models’. In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 421–422.
6. Gelfand AE. Discussion of ‘computation on Bayesian graphical models’. In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 422–423.
7. Dawid AP. Discussion of ‘computation on Bayesian graphical models’. In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 422.
8. Fienberg SE. Discussion of ‘computation on Bayesian graphical models’. In *Bayesian Statistics 5*, Bernardo JM, Berger JO, David AP, Smith AFM (eds). Oxford University Press: Oxford, 1996; 422.
9. Drton M, Perlman MD. Model selection for Gaussian concentration graphs. *Biometrika* 2004; **91**(3):591–602.
10. Wine User Guide. Retrived June 7, 2009. Available from: <http://www.winehq.org/docs/wineusr-guide/index>.
11. Ashby D. Bayesian statistics in medicine: a 25 year review. *Statistics in Medicine* 2006; **25**:3589–3631.
12. Attributed to *Anonymous*.

COMMENTARY

Comments on ‘The BUGS project: Evolution, critique and future directions’

Jon Wakefield*,†

Departments of Statistics and Biostatistics, University of Washington, Seattle, WA, U.S.A.

It was a pleasure to read the paper of Lunn *et al.*, and I congratulate the authors on their achievement; by providing a means for the routine implementation of Bayesian methods the BUGS program has revolutionized statistics. My own small part in the history of BUGS was to work with the authors on the PKBUGS project, [1]. My own PhD, and those of Dave Lunn and Nicky Best all concerned Bayesian population pharmacokinetic modeling [2–4], and so it was a natural step for us to use the BUGS machinery to carry out the required hierarchical nonlinear modeling.

With respect to Section 5.2 of the paper, I have for a number of years [5] been troubled by the specification of inverse gamma priors for the variance of random effects. A major difficulty is that beyond the linear model, it is not straightforward to think about the random effects due to the scale on which they act, particularly for binomial data. I will briefly describe a method that I have lately experimented with [6]. For concreteness we will consider a generic log odds ratio in cluster i , $\delta_i | \sigma^2 \sim_{\text{iid}} N(0, \sigma^2)$ with $\sigma^{-2} \sim \text{Ga}(a, b)$. Integration over σ^2 gives the marginal distribution for δ_i , and is a t distribution with $d = 2a$ degrees of freedom, location zero, and scale $\Sigma = b/a$. To construct a prior distribution we require a careful interpretation of δ_i , or more informatively, $\exp(\delta_i)$. The latter is the deviation in cluster i of the odds, from the median of the distribution of the odds, across all clusters. Hence we may refer to $\exp(\delta_i)$ as residual odds, since it is relative to the median odds across areas. To pick a and b we give a range for $\exp(\delta_i)$. In particular, for a range $(1/R, R)$ we use the relationship $\pm t_{0.025}^d \sqrt{\Sigma} = \pm \log R$, where t_r^d is the $100 \times r$ th quantile of a Student t random variable with d degrees of freedom, to give $a = d/2$, $b = (\log R)^2 d / 2 (t_{1-(1-q)/2}^d)^2$. We (somewhat arbitrarily) choose $d = 1$, to give a log Cauchy marginal distribution for $\exp(\delta_i)$. As an example, for a 95% range of $[0.1, 10]$ we obtain $a = 0.5$, $b = 0.0164$.

Given the ubiquity of MCMC, I recently enjoyed reading the article of Rue *et al.* [7], which describes the integrated nested Laplace approximation (INLA) approach to Bayesian computation, a method that cleverly combines Laplace and numerical integration. I would be interested in the authors’ view of this work. I have found it particularly useful for carrying out simulation studies in real time to compare Bayesian procedures with competitors. Of course, this new found ability opens up a number of difficult questions such as how to choose priors ‘fairly’ in a simulation setting.

*Correspondence to: Jon Wakefield, Departments of Statistics and Biostatistics, University of Washington, Seattle, WA, U.S.A.

†E-mail: jonno@u.washington.edu

Finally, I wonder whether some of the flexibility of allowing, for example, Poisson distributions that are defined for non-integer responses could be used creatively for providing Bayesian analogues of useful frequentist techniques such as quasi-likelihood or sandwich estimation. I have recently [8, 9] been performing Bayesian inference using asymptotic normal likelihoods combined with normal priors (in a genome-wide association studies context, in which potentially millions of analyses are required). An advantage of this framework (which has been previously used by a number of authors, including [10] is that asymptotic distributions based on quasi-likelihood, sandwich variance estimates (including their use in GEE), and other non-standard sampling situations such as partial likelihood, two-phase sampling, etc., can be incorporated in the analysis. Many of these procedures have no formal Bayesian justification, though from a practical standpoint they are all (at least to me!) clearly a good idea. While BUGS is not required for the normal/normal framework just described, perhaps some of the flexibility of BUGS could be used to combine the practical attractiveness of the frequentist approaches listed above, with prior distributions that smooth the empirical estimators. The latter may be particularly useful in small sample situations in which, for example, sandwich estimates may be unstable.

REFERENCES

1. Lunn DJ, Best N, Thomas A, Wakefield J, Spiegelhalter D. Bayesian analysis of population PK/PD models: general concepts and software. *Journal of Pharmacokinetics and Pharmacodynamics* 2002; **29**:271–307.
2. Wakefield JC. The Bayesian analysis of pharmacokinetic models. *Ph.D. Thesis*, University of Nottingham, 1992.
3. Best NG. Dynamic models for post-transplant patient monitoring. *Ph.D. Thesis*, University of Cambridge, 1993.
4. Lunn D. The application of Markov chain Monte Carlo techniques to the study of population pharmacokinetics. *Ph.D. Thesis*, University of Manchester, 1995.
5. Kelsall JE, Wakefield JC. Discussion of ‘Bayesian models for spatially correlated disease and exposure data’, by Best *et al.* In *Bayesian Statistics*, vol. 6, Bernardo JM, Berger J, Dawid AP, Smith AFM (eds). Oxford University Press: Oxford, 1991; 151.
6. Wakefield JC. Multi-level modelling, the ecologic fallacy, and hybrid study designs. *International Journal of Epidemiology* 2009; **38**:330–336.
7. Rue H, Martino S, Chopin N. Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with Discussion). *Journal of the Royal Statistical Society, Series B* 2009; **71**:319–392.
8. Wakefield J. A Bayesian measure of the probability of false discovery in genetic epidemiology studies. *American Journal of Human Genetics* 2007; **81**:208–227.
9. Wakefield J. Bayes factors for genome-wide association studies: comparison with *P*-values. *Genetic Epidemiology* 2009; **33**:79–86.
10. Spiegelhalter DJ, Freedman LS, Parmar MKB. Bayesian approaches to randomized trials (with Discussion). *Journal of the Royal Statistical Society, Series A* 1994; **157**:357–416.

REJOINDER

Rejoinder to commentaries on ‘The BUGS project: Evolution, critique and future directions’

We are very grateful to the discussants for their contributions, which express the mixture of gratitude and frustration that seems to characterize many BUGS users.

Some common themes can be identified. Both Andrew Gelman and Martyn Plummer suggest that the program could be made more ‘intelligent’, both in automatically spotting when more efficient sampling algorithms could be used, such as block updating in hierarchical models, and in trying to help users in their modelling. The latter aim could require the development of a basic ‘expert system’ that, in Plummer’s words, puts out ‘red flags’ when, for example, an apparently innocuous prior is being very influential, or there is clear conflict between one part of the data and either the prior or another part of the data. In both cases the key idea is probably refining the design of the VGM to provide more contextual information. This has already been identified as a long-term strategy for recognizing larger design motifs (in Plummer’s words), as well as for facilitating the operation of alternative inference engines on the BUGS VGM.

Martyn Plummer and Kate Cowles both comment on the need to exploit parallel processing opportunities. Indeed, this seems an obvious step given that multi-chain Gibbs samplers represent examples of *embarrassingly parallel* problems (where there is no need for communication between processors), although there is great scope for implementing more efficient parallelization schemes. As mentioned in the paper, some of the infrastructure required for parallelizing BUGS has been incorporated into the design of OpenBUGS, but it is unlikely that this will fully come to fruition before issues with the Linux version (see below) have been addressed. Actually, it may interest readers to know that a parallel version of WinBUGS apparently already exists in the form of *GridBUGS* [1], developed by a team at Johnson and Johnson for distributing applications across their computer grid.

One of the main sources of frustration, as pointed out by both Kate Cowles and Teixeira-Pinto and Normand, is the error reporting system. We are sure we could do better here, but it is difficult to say by how much. The problem might be due in part to the modular nature of BUGS. Each modelling component can perform basic checks on itself and its neighbours when initialized, but such checks are context-independent and localized, as well as potentially not comprehensive enough. Hence, context-driven and long-range conflicts between modelling components may always be difficult to catch. As noted above, extracting more contextual information from the virtual graph is a long-term objective. Perhaps this contextual information can also be used to improve error handling.

Both Andrew Gelman and Teixeira-Pinto and Normand suggest the use of predictive distributions in order to carry out model checking—this is not difficult to do within BUGS but does currently require special effort. Progress towards a flexible and stable open-source version will, we hope, make it much easier for such features to be automated.

Andrew Gelman rightly makes comparisons with the R-function *lmer*, and his excellent book with Jennifer Hill shows the relative strengths of each package. We have found a common practice is to use BUGS as a ‘gold standard’, but to then seek faster non-MCMC implementations as approximations that allow rapid exploratory analysis.

Jon Wakefield focuses on some tricky modelling issues within Bayesian analysis. Our own preference is to avoid inverse-gamma priors altogether for variances of random effects, and follow Gelman [2] in using half-normal or half-Cauchy priors for random-effect standard deviations, since these place finite but non-zero weight on the variance component being precisely zero, but decline steadily in order to down-weight the possibility of complete independence. His suggestion regarding the partial use of non-model-based likelihoods is intriguing, and follows our use of the 'cut' function [3] to allow inferences that are not based on a full probability model: however, each case threatens nocturnal visits from the Bayesian thought-police.

We fully support Kate Cowles' and Teixeira-Pinto and Normand's pleas for less Windows-dependence, and again point to the opportunities that the formation of a 'BUGS Foundation' offers to support such developments in the future. In particular, at the time of writing, we were gearing up towards a major new release of OpenBUGS, which we hope will attract much of the existing WinBUGS community. One of our fundamental prerequisites for release is a fully functioning and comprehensively documented (including compilation instructions) Linux implementation.

Finally, Teixeira-Pinto and Normand allude to the possibility of learning about the structure of the graph as part of the analysis. We have experimented to some extent with such ideas in implementing the *Jump* interface [4]. A more general approach might make use of a dynamic VGM, one in which links between nodes can easily be switched on and off. We are not sure if the data structures currently used for representing graphs are well suited to such endeavours, however. Perhaps this is an area for the next generation of graphical modelling software, which the 'BUGS Foundation' will aim to support.

DAVID LUNN

*MRC Biostatistics Unit, Institute of Public Health
University Forvie Site, Robinson Way, Cambridge CB2 0SR, U.K.
E-mail: david.lunn@mrc-bsu.cam.ac.uk*

DAVID SPIEGELHALTER

*Institute of Public Health, University of Forvie Site
Robinson Way, Cambridge CB2 0SR, U.K.
E-mail: david.spiegelhalter@mrc-bsu.cam.ac.uk*

ANDREW THOMAS

*School of Mathematics and Statistics
University of St Andrews, St Andrews, U.K.
E-mail: andrewt@mcs.st-andrews.ac.uk*

NICKY BEST

*Division of Epidemiology, Public Health and Primary Care
Imperial College, South Kensington Campus, London SW7 2AZ, U.K.
E-mail: n.best@imperial.ac.uk*

REFERENCES

1. Girgis IG, Schaible T, Nandy P, De Ridder F, Mathers J, Mohanty S. Parallel Bayesian methodology for population analysis. 2007. Available at: www.page-meeting.org/?abstract=1105.
2. Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* 2006; **1**:515–533.
3. Lunn D, Best N, Spiegelhalter D, Graham G, Neuenschwander B. Combining MCMC with 'sequential' PKPD modelling. *Journal of Pharmacokinetics and Pharmacodynamics* 2009; **36**:19–38.
4. Lunn DJ, Best N, Whittaker JC. Generic reversible jump MCMC using graphical models. *Statistics and Computing* 2008; DOI: 10.1007/s11222-008-9100-0.