

Quantifying the impact of non-adherence to drug therapy: a technical note concerning an application of a branch and bound algorithm

Running title: Non-adherence to therapy: a branch and bound algorithm.

Daniel Smith 222 Hill Lane, Southampton, SO15 7NR, UK

Christina Pagel Clinical Operational Research Unit, University College London

Martin Utley Clinical Operational Research Unit, University College London

Steve Gallivan Clinical Operational Research Unit, University College London

Corresponding Author

Martin Utley (m.utley@ucl.ac.uk)

Clinical Operation Research Unit,

University College London,

Gower Street,

London, WC1E 6BT

UK.

Contact to be cited for readers

Christina Pagel (c.pagel@ucl.ac.uk), Clinical Operational Research Unit, University College London, Gower Street, London WC1E 6BT, UK.

Acknowledgement

Abstract

Pharmacokinetic models typically rely on a key assumption that patients take their medication as prescribed, whereas this is often not the case. We present a branch and bound algorithm that can be used to estimate the time-varying probability that, given a specified pattern of non-adherence to a prescribed regimen, a patient receives no therapeutic benefit from treatment. Use of this algorithm is a much faster method for obtaining this probability than exhaustive computation of the relevant probability distribution. The use of this algorithm to assess, in quantitative terms, the impact of non-adherence on the effectiveness of treatment provides a rational basis for evaluating the potential harm to patients.

Key words

Pharmacy; Stochastic modelling; Computational algorithms; Branch and bound methods

1. Introduction

Pharmacokinetic models are almost invariably based on an assumption that medicines are taken as prescribed, whereas it is widely acknowledged that deviations from the prescribed regimen (termed non-adherence) are common [1,2]. An analytical framework for incorporating some aspects of non-adherence into pharmacokinetic models is presented elsewhere in this journal [3]. This provides a means of calculating the mean and standard deviation of the concentration of a drug in the serum given a particular level of adherence to a specified drug regimen, which is valuable information if one wants to assess the potential for patient harm. In this note, we present an algorithm that can be used to quantify this potential harm in terms of the proportion of time that the concentration of drug in the serum of a patient is so low that no clinical benefit is conferred. This approach is adopted since the probability distributions concerned cannot accurately be approximated by standard parameterised distributions.

In section 2 we discuss some basic concepts of drug therapy and introduce the notation used in our analysis of a patient prescribed a course of drugs. In section 3, we describe a computational algorithm based on a branch and bound technique which can be used to calculate the probability that the concentration of the drug in the serum is below a pre-specified level at a specified point in time.

2. Incorporating non-adherence into a model of pharmacokinetics

A key concept in therapy using medicines is that of a lower therapeutic limit for the concentration of a drug in the serum. If the drug concentration is outside this and in the sub-

therapeutic range, the therapy is judged to be clinically ineffective. When prescribing a drug for a chronic condition, say, the expectation is that the medicine will be taken regularly over a long period of time. The drug dose and frequency are carefully chosen so that, if the medicine is taken as prescribed, the drug concentration should not fall in the sub-therapeutic range, an aim which can be disrupted if there is non-adherence. For a broader discussion of these notions see [4].

The analysis here concerns the concentration of a drug in the serum at some time T following a scheduled sequence of N administrations of the drug, given the possibility of non-adherence. We consider a simple model of non-adherence, assuming that for each scheduled administration there is a probability that the drug is not taken, each scheduled administration being modelled as an independent Bernoulli trial.

For $1 \leq w$ let the random variable X_w have the value 1 if the w -th scheduled administration of the drug occurs and the value 0 otherwise, r_w denoting the probability that the administration does not occur.

Our analysis concerns cases where the combined impact of several drug administrations on the serum concentration of the drug can be calculated as the sum of the individual impacts of each administration had it occurred alone. For $1 \leq w \leq N$, let $C_w(T)$ denote the intended contribution of the w -th scheduled administration to the concentration of the drug in the serum at time T . The exact form of $C_w(T)$ depends on the drug, the schedule of doses, the specified administration times and the underlying pharmacokinetic model (see [3] for one example or [4] for a more general discussion). For the purposes of this paper it is enough to state that these quantities can be calculated.

Let the time varying random function $Z(T)$ denote the serum concentration of the drug at time T due to the combined effect of the drug administrations up to that time given by:

$$Z(T) = \sum_{w=1}^N X_w C_w(T). \quad (1)$$

In principle, the entire probability distribution for the values taken by $Z(T)$ can be calculated by exhaustive enumeration of possibilities, although this becomes computationally more demanding and then infeasible as the number of scheduled administrations increases. In addition, the distribution of $Z(T)$ cannot reliably be approximated by any standard parameterised distribution. For these reasons we have used numerical methods to calculate the probability that $Z(T)$ corresponds to a serum concentration that is within the sub-therapeutic range.

3. An algorithm for calculating the probability of the drug concentration being in the sub-therapeutic range.

We have developed a branch and bound algorithm (see [5, 6] for general introductions to branch and bound techniques) to find the probability that the serum concentration of a drug is in the sub-therapeutic range at time T . Full details of the algorithm are available on-line (www.ucl.ac.uk/operational-research/downloads); here an overview is given.

It is useful to think of the outcomes of the scheduled drug administrations in terms of a descending directed binary tree with $(N + 1)$ horizontal levels of nodes, the first N levels corresponding to scheduled drug administrations. These levels are not in temporal order, but in decreasing order of the intended contribution of the administrations to the serum concentration at time T . At the nodes on these levels the tree branches to the right or left to distinct nodes on the next level according to whether or not the drug is taken. The nodes on the $(N+1)$ -th level have no further branching. This tree has $(2^{N+1} - 1)$ nodes and 2^N maximal paths leading from the highest level to different nodes on the lowest level. The probability that a particular path is followed depends on the path's constituent Bernoulli probabilities.

We distinguish two categories of nodes. A node is said to be *Type A* (mnemonically, **A**bove) if the drug concentration at time T would be above the sub-therapeutic range whatever subsequent path is followed in the tree. *Type B* (mnemonically, **B**elow) nodes are those where the concentration at time T would be in the sub-therapeutic range whatever subsequent path is followed. Although in principle all nodes could be explicitly categorised in this way, we use a branch and bound algorithm structured to avoid the computational burden that this would entail.

The algorithm involves successively updating three sets of nodes: sets A and B containing nodes known to be Type A or Type B , respectively, and a set U of nodes under consideration. The algorithm alternates between a branching process, increasing the number of nodes in U , and a bounding process, re-assigning nodes from U to the sets A or B if they are Type A or Type B respectively. More details of the algorithm are given in the Appendix.

4. Example application of the algorithm

In this example, we consider adherence to taking mexilitene which is used to treat ventricular arrhythmias. Consider a patient who has been prescribed 250mg doses of mexilitene to be taken at 8 hour intervals [7] but forgets to take the drug with probability 0.2. The lower therapeutic limit for mexilitene is 0.5 mg/L. Figure 1 shows the probability over time that the concentration of drug in the serum is within the sub-therapeutic range. After the first scheduled administration the probability of the drug concentration being within the sub-therapeutic range is 1 since a single drug administration is insufficient to bring the drug concentration higher. However, after a few weeks, a cyclic pattern is established, with the probability falling at each scheduled administration and then rising again. Once such a pattern is established, the expected proportion of time that the drug concentration is within the sub-therapeutic range can be calculated by taking the average of this probability, as obtained using the algorithm described above, over one cycle. In the example given, the expected proportion of time that the drug concentration is within the sub-therapeutic range is just under 10%.

INSERT FIGURE 1 HERE

5. Discussion

We have developed a branch and bound algorithm to solve a complex problem in an application of stochastic pharmacokinetics. Use of this algorithm enables one to calculate rapidly the probability that the concentration of a drug in a patient's serum is within the sub-therapeutic range, given a specified level of stochastic non-adherence to treatment. The algorithm can be adapted to calculate prediction intervals in this context, albeit with additional complexity.

The algorithm presented in this note was developed in response to the specific clinical problem discussed here and in [3] but we consider the approach adopted to have much wider applicability. Essentially the algorithm is used to calculate cumulative probabilities associated with the probability distribution of a sum of independent, differently distributed, discrete random variables in the special case in which each such variable can take just two values, one of these two values being zero. One possible extension of our work is to modify the algorithm such that it can be used in the more general case of a sum of independent, discrete random variables.

References:

1. L. Orsterberg, T. Blaschke, Adherence to Medication. *New England Journal of Medicine*, **353**(2005): 487-497.
2. B. S. Bloom, Daily regimen and adherence with treatment, *British Medical Journal*, **323**(2001):647
3. S. Gallivan, C. Pagel, M. Utley, B. Dean Franklyn, K. Taxis, N. Barber, A technical note concerning non-adherence to drug therapy: exact expressions for the mean and variance of drug concentration, **submitted alongside this paper** to Health Care Management Science 2007.
4. L. Shargel, S. Wu-Pong, A.B.C. Yu, *Applied Biopharmaceutics & Pharmacokinetics*, 5th edition, McGraw Hill, London, 2005.
5. M. Brusco and S. Stahl. *Branch-and-bound applications in combinatorial data analysis*, Springer, New York 2005
6. W.L. Winston, *Operations research : applications and algorithm*, Thomson Brooks, California, 2004. Chapter 9
7. British National Formulary, March 2003, British Medical Association.

Appendix. Description of the branch and bound algorithm.

We use the vertex set $V = \{v_i\}$ to denote the set of nodes of the directed tree, which are arbitrarily indexed except that the top node is denoted v_1 . For $1 \leq i \leq 2^{N+1} - 1$, we denote the probability that the outcomes of Bernoulli trials would result in a path passing through the node v_i by $p(v_i)$.

In the summary of the branch and bound algorithm below, the symbol j is used to denote the index of the node within U that is selected for branching. The algorithm is summarised as follows:

Step 1 INITIALISATION

*Set A and B as empty and U as containing just one node v_1 , the top node in the tree.
Set $p(v_1) = 1$, since all possible paths emanate from v_1 . Compute the maximum value of drug concentration at time T associated with this node if all the scheduled drug administrations occurred and assign v_1 to B if possible. END if U is now empty.*

Step 2 BRANCHING

Update U , replacing the node selected for branching by the two nodes it leads to in the next level of the tree. Compute the probabilities $p(\cdot)$ and the minimum and maximum values of drug concentration at time T associated with these nodes.

Step 3 BOUNDING

If possible, reduce U , by re-assigning the two new nodes to the sets A or B if possible.

Step 4 SELECTING THE NEXT BRANCH NODE

IF the set U is empty, sum $p()$ for nodes in B to give the solution THEN END.

This is a stopping condition indicating that the search for the probability that the drug concentration at time T is in the sub-therapeutic range is complete.

ELSE select as the branching node that v_j in U which leads to the highest value of $Z(T)$.

GO TO Step 2

To illustrate the operation of the algorithm, consider a sequence of six scheduled administrations. For one such sequence, Figure A1 shows the nodes of the tree that belonged to the set U at some point during the operation of the algorithm.

The nodes which are shown as black represent nodes of the tree assigned to the set A by the time the algorithm finished. Those shown as white are those included in the set B used to calculate the probability of the final drug concentration being in the sub-therapeutic range. This is the total probability associated with all paths passing through one of the white nodes. The nodes shaded grey are those that were members of U , were branched on and removed from U .

INSERT FIGURE A1 HERE

Figure Caption

Figure 1 The probability of drug concentration being in the sub-therapeutic range for mexiletene if the probability of each administration not occurring is 0.2.

Figure A1 Example nodes evaluated during a run of the algorithm for 6 scheduled administrations.

Nodes shown as black are those in the set A at the final step of the algorithm. Those shown in white are in set B at the final step of the algorithm and are included in the calculation of the cumulative probability that the drug concentration is in the sub-therapeutic range.

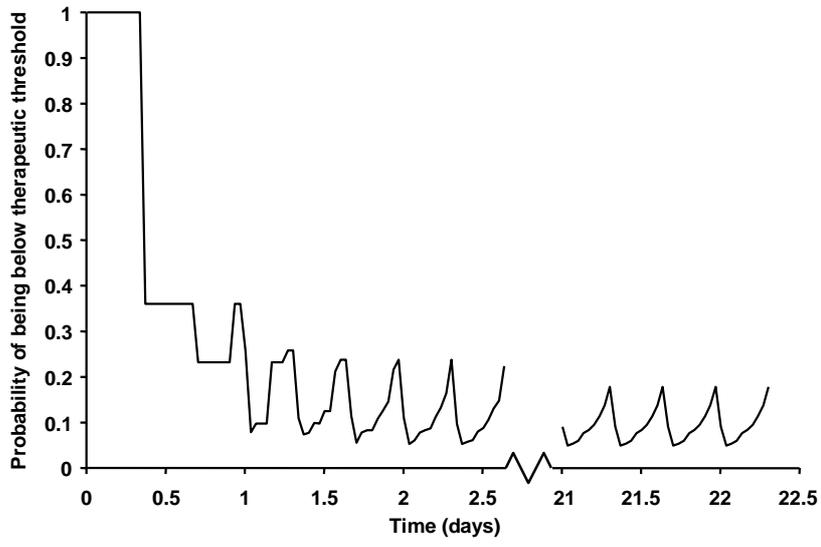


Figure 1

