A Markov approach to characterizing the PK-PD relationship of anti-migraine drugs

Maas, H.J.

8

Modelling of the PK-PD relationship of anti-migraine drugs using a Markov approach: Summary, conclusions and perspectives

The previous chapters described the application of Markov models to data from clinical trials of migraine. It has been shown that these models can be used to derive concentration-effect relationships, evaluate the effect of changes in formulation, study the role of demographic variables on efficacy and predict the response to medication both during and outside attacks.

The use of Markov models for analysing migraine data is not entirely new. Ebutt and Hassani were the first to propose a non-homogeneous Markov chain in discrete-time time to describe the course of attacks in patients [1]. The Markov chain acknowledges the highly variable course of headache which, though eventually resolving, exhibits features of randomness. More importantly, the Markov chain recognises a disease as a dynamic process giving rise to a pattern of symptoms. Therapeutic intervention is either directed against these symptoms or against the underlying disease process itself. Often, knowledge about the process is lacking, as are adequate biomarkers. In an effort to understand the process, disease stages or states are then defined on clinical grounds. These states may or may not have suspected biological substrates. The dynamic behaviour of such disease models is realised by assigning progression rates from one state to the next. The

dynamics of Markov chains often resemble those of diseases albeit within a limited period of time. This is remarkable given the strict Markov property, which states that the future state of the process only depends on the current state and not on any previous states.

The studies in this thesis were not only aimed at describing the pain course of migraine attacks, but also to relate pain response to the concentration of triptans, *i.e.* to develop a model that describes the PK-PD relationship of triptans. In the following sections, the development and application of the hidden Markov model for migraine attacks are discussed by chapter.

8.1 Developing a disease-state approach to modelling

Chapter 2 described the development of the hidden Markov model and its application to data of clinical trials of sumatriptan. In the development of this model, several issues were addressed.

In selecting the model structure, the clinical differentiation into a pain free and a pain relief state was considered most important. The number of states (three) reflects this differentiation: a baseline state occupied by moderate and severe headache scores, a pain relief state containing mostly mild pain scores and a pain free state. It can be argued that four states would give the most accurate description of the migraine attack as the pain is assessed using a four-point scale. However, it was mainly because of the large heterogeneity in headache scores that it was decided to only use three states. The heterogeneity can partly be explained by the setup of typical acute anti-migraine trials. Patients enrol into a study once their headache intensity is at least moderate. The interval leading up to the point of enrolment is highly variable and as a result patients are all in different phases of their attacks. Some may have a moderate score and will progress to a severe score, while others with moderate headache will not experience an increase anymore. Because these different phases cannot be distinguished from the data, it was decided to group all patients into a single starting state. Another reason for not assigning different states to moderate and severe scores is related to the first reason. Once headache starts to decrease, patients with severe headache first experience moderate headache and then progress to mild and no pain. In other words, the moderate pain phase is a transient phase. This is in contrast to the severe pain phase, in which these patients reside for a longer time. Patients entering the study with moderate pain severity on the other hand are likely to stay in that phase before eventually progressing to mild and pain free states. Thus, for different patients, the moderate pain phase has dynamically different meanings. Treating all transitions from this state as the same, the variability of this transition will increase and the expected benefits of having four states are not materialised.

Rather than adding a hidden layer to the Markov chain, the moderate and severe pain scores could just be grouped manually into a single category and further analysed using a standard three-state Markov chain. This was not done as it appeared that estimating the distribution of the four scores over the three states provides a useful measure of similarity among data sets. If two data sets were to be analysed that differ appreciably in

the length of the observation interval, the states of their models will contain different score distributions. If for example the observation interval only covers the first 2 hours after dosing, relatively few patients will have attained pain free status. Hence, if a state existed that contained only pain-free scores, not enough information would be available to estimate the associated transition rate to this state. The minimisation algorithm circumvents this by finding another distribution of the four scores over the three states. As a result, the transition rates obtained from two data sets that considerably differ in their observation windows cannot be compared but due to the redistribution of scores, predictions of headache course can still be made. This characteristic property was used to determine an optimal observation window for the data sets. On the one hand, the observation period was short enough to ensure that the Markov property was not violated. On the other hand, the period was long enough to generate a distribution of headache scores over the states that reflected the clinical subdivision into a pain state, pain relief state and a pain free state. A disadvantage of using a HMM for clustering scores into states is that the algorithms are more complex and less efficient than those used for regular Markov chains.

Modelling was performed using an S-Plus library for (hidden) Markov models. This library allowed for the implementation of user-specified functions to link transition rates to covariate values. Because of this, a nonlinear concentration *vs* transition rate relationship could be built into the Markov chain part of the model. This feature allowed for the first time the evaluation of a nonlinear concentration-effect relation within a stochastic dynamic system (the Markov chain). The most obvious advantage is the characterisation of both improvement and worsening of pain over time. Other analytical methods have not been able to describe the migraine attack using these dynamical concepts.

A complicating factor in the development of the concentration-effect relation was the presence of a natural exponent in the coded transition rates. Though the use of these exponents prevented transition rates from becoming negative, it also prohibited describing concentration-dependency of the transition rate using a regular $Emax$ model. Instead, the concentration *vs* transition rate relationship follows the shape of an exponentiated $Emax$ model. Though the shape is roughly equal, the interpretation of drug-related parameters is slightly different.

Another difficulty is related to the identifiability of drug-related parameters ($EC50$ and $Emax$) in the absence of a sufficient number of different doses. Since 50 mg and 100 mg sumatriptan have similar efficacy and because few studies have been performed using the 25 mg oral dose, precise estimations of these parameters could not always be obtained. It was demonstrated during the development of a similar model for sumatriptan that with a broader range of doses, these parameters could indeed be estimated more precisely. The information available in the concentration data was further limited by the fact that only population-averaged predictions of concentration were available for the PK-PD analysis. In most studies, PK measurements and pain assessments were not both present. Instead, PK information was derived mostly from phase I clinical studies. Using these data requires that there are no relevant differences between sumatriptan PK in healthy volunteers and in patients. Though it is often claimed that these differences can be considerable, PK analysis has previously shown that the PK characteristics are

not relevantly different [2]. Yet, it cannot be ruled out that estimates of drug-related parameters would be different had patient PK samples been available.

8.2 Development of diagnostic tools

Goodness-of-fit evaluation of categorical data analysis is notoriously underdeveloped. Except for comparison of likelihoods, no tools were present in the HMM S-Plus library. Therefore, an S-Plus routine was developed for the routine evaluation of goodness-of-fit and for assessing the precision of the predictions (chapter 3). Goodness-of-fit was assessed by comparing observed pain relief and pain free percentages over time with those generated by a recurrent algorithm using the parameter estimates obtained from fitting the data to the hidden Markov model. The recurrent algorithm uses the Kolmogorov differential equations to generate mean score *vs* time profiles. This method is both accurate (*i.e.* it is not an approximation, like simulation methods) and fast. Pain free and pain relief profiles were chosen as a measure of goodness-of-fit rather than the separate pain scores because it was the objective of the Markov modelling to predict these primary and secondary clinical endpoints. Moreover, since moderate and severe pain scores were grouped into a single state, assessing the courses of the separate scores over time would not be helpful as this would lead to consider the use of a 4-state Markov chain, the use of which had been ruled out based on arguments mentioned above.

The uncertainty in the predicted headache response profiles was visualised by constructing 95% confidence intervals around the mean-predicted headache profiles. These intervals were constructed using a modified version of the Kolmogorov algorithm that included the standard errors of the parameter estimates. This calculation used a first-order approximation of the variance. The validity of the approximation was evaluated by performing a series of Monte Carlo simulations. It appeared that a first-order approximation could be used for most of the responses. For the pain free response in the presence of sumatriptan, however, the standard errors were large and the calculated confidence interval differed considerably from the simulated one. The simulated confidence intervals were asymmetrically distributed around the mean. This is due to the exponent in the equations describing the state-to-state transition rates.

The method proposed is fast and works well for small errors. When errors are large, the use of simulation techniques is necessary.

8.3 Model consistency across drugs

The concentration-effect relationships of both sumatriptan and naratriptan were derived using the Markov modelling approach. The results confirmed the general clinical observation that, for pain relief, naratriptan is approximately three times more potent than sumatriptan and that the maximum effects are more or less equal.

In the case of sumatriptan, the pharmacodynamic parameters on the transition towards pain free status could not be precisely estimated. However, the naratriptan analy-

ses resulted in estimates with higher precision. This was attributed to the availability of a broader range of doses for the latter drug.

On the same transition, different values were also found for the parameter estimates themselves. Possibly, the lack of sumatriptan dose groups caused unidentifiability and strong correlation among the PD parameters.

Provided that sufficient information is available, the characterisation of concentration-effect relationships using Markov models can give more meaningful parameter estimates than when obtained by other modelling techniques. By considering the fluctuations in headache score over time, a realistic framework for the disease was built that served as the background against which drug action was modelled. Previous PK-PD analyses did not take into account fluctuations and therefore could not describe phenomena such as recurrence. An incomplete characterisation of the disease can lead to biased estimates of the PD parameters. This in turn complicates comparison between drugs and between studies.

The PD estimates were also compared with the values of potency and the maximum effects obtained in *in vitro* and *ex vivo* experiments. The extent of agreement depended on drug and transition, but in general the in vitro values corresponded best with the estimates. This result is promising for the evaluation of the efficacy of new compounds in man. Substitution of in vitro values can give an indication of the efficacy that can be expected in patients.

As for the expected recurrence, no difference was found in the recurrence rates derived for both drugs after correcting for placebo rates. In this respect therefore, the model structure assumes that recurrence is a phenomenon inherent to the patient rather than drug action. This view was previously taken by [3] but is in contrast with findings in [4].

8.4 The effect of covariates on the anti-migraine response

Markov disease models are traditionally used to describe chronic and episodic disorders, covering time spans of months or years. Little is known about the episodic nature of migraine. A combination of endogenous and exogenous triggers seem to illicit attacks, making their frequency highly variable both within and between patients. Stochastic models such as Markov models are the preferred modelling tools under these conditions. In the current analysis, a two-state Markov model with a single layer of variables was used to characterise the attack frequency in the presence and absence of oral sumatriptan treatment. This method was evaluated against simple distribution fitting.

Compared with distribution fitting, the choice of distributions offered in standard Markov models is limited to exponential and Erlang distributions (a concatenation of n exponential distributions, with n an integer). On the other hand, the Markov approach acknowledges the dynamic character of the disease, iterating over the time series of pain data, rather than requiring that the time intervals between attacks be already calculated. for this reason, this approach can also be more easily extended to contain subject-specific

effects.

Preferentially, this type of analysis is performed on time series studies where pain is assessed daily and for a year or longer. The data in the current analysis were derived from an acute study investigating the reproducibility of the drug effect over the course of up to three subsequent attacks within one year. Though this did not impede the estimation of frequencies, if more information had been available, the influence of one or more covariates on the duration of the period between attacks could have been assessed.

The greater flexibility of the distribution analysis was apparent from the accuracy of prediction, which was better than that for the Markov model. Neither model demonstrated a significant difference between the duration of interictal periods in the presence of placebo and active drug. These models can proof especially useful for assessing the effect of prophylactic anti-migraine drugs, where such a difference should be demonstrated.

8.5 The episodic nature of migraine

Many covariates can influence the course of resolution of a migraine attack.

Age, sex, baseline pain score, body weight and type of migraine are all known to positively or negatively impact the pain duration [5, 6]. These factors add to the heterogeneity of the disease and as a result frustrate the search for effective drugs. The current analysis focused on the effect of patient age on both the placebo and the drug effect in oral sumatriptan studies. It has been reported that in contrast to the response to drug, the response to placebo decreases with age [7, 8]. This observation can be explained by assuming a maximum stimulus on the transition rate towards a state of less pain. Once that rate has been reached, other stimuli cannot increase the rate further. In clinical terms, the increased response after placebo in young patients can be attributed to the placebo effect. However, non-treated attacks are also known to be shorter-lasting in children and adolescents. Placing these facts into the above theory, this would mean that either young patients are intrinsically closer to the maximum transition rate or the placebo effect increases the transition rate.

In either case, modelling this scenario requires that two covariates be taken into account: one time-varying (drug concentration) and one constant (age). As both covariates were assumed to affect the transition rates towards states of less pain, they were incorporated in the same equations describing the transition rates in the hidden layer of a three-state HMM.

The range of available age groups was extended by adding a study of migraine in adolescents. On the other hand, the range of time points in this study was limited to times of up to 4 h post dose. This was not expected to severely impact estimation, as the duration of attacks is intrinsically shorter in the young.

The equations describing the transitions to states of less pain were based on an analysis of the interaction of toxic compounds [9]. These equations were embedded into the HMM.

The effect of other continuous covariates can be characterised using a similar model.

Alternatively, covariate effects can also be tested on the open layer of score distributions. For example, the covariate sex can be placed on the distribution of scores in the first state of the HMM, as it was found that female patients are more likely to experience severe headache at the start of the study (difference 5%, unpublished results).

8.6 Perspectives

Hidden/latent states concept The need for a hidden layer in a Markov model is usually based on one of two considerations.

Firstly, the latent states in this layer are used to classify sequences of measurements sharing particular properties. In the analysis of electrocardiogram (ECG) and electroencephalogram (EEG) signals, locally stationary sequences are grouped into separate states [10, 11]. The properties of each state are then described in the open layer by means of statistical distributions. Most biological signals are nonstationary. That is, their statistical characteristics such as mean and variance change over time. The grouping of headache intensity scores measured in the course of a migraine attack is another example of this use of latent states. Each state corresponds to a clinically defined endpoint and may well have a pathophysiological analogue. The number of latent states required to model a nonstationary time sequence depends on degree of nonstationarity in the signal and the signal's information content. The number of states should be sufficiently large to encompass the periods marked by local stationarity. However, their number should be sufficiently small to maintain parameter identifiability.

A second motivation for using latent states is to address the issue of misclassified observations or missing observations. In both cases, these observations are treated as sampled from a distribution that corresponds to one of the latent states. For example, a false-positive outcome of an imperfect diagnostic test will be treated as an observation drawn from a state representing absence of the disease. The correctness of this procedure strongly depends on the validity of the Markov property in a given system.

Hidden Markov models (HMMs) have been successfully applied to model chronically recurring infections, such as herpes [12], papillomavirus [13] and malaria [14]. Misclassification by diagnostic tests is an important concern in these diseases.

Other applications Migraine has been described as a chronic neurological disorder with episodic manifestations [15]. Other disorders in this category include epilepsy and cluster headache. For these diseases aetiology and pathophysiology are incompletely understood. Data collected in a clinical setting usually involves ordinal measures of pain and disability or simply count data. Although individual events may be too short to be effectively characterised by Markov models, the occurrence of events can be the object of Markovian analysis.

Before frequency analysis was introduced in the field of epilepsy, two-state Markov models were used to analyse patterns of seizures [16]. It was conjectured that if the occurrence of an epileptic episode is a random event, it can be described by a Markov

chain. Alternatively, if the events are periodical, a Markov chain cannot characterise the patterns.

The issue of randomness versus periodicity of events is particularly interesting in cluster headache [17]. Cluster periods are mostly periodical (often season-bound). However, the end of every period may constitute the end of the disease itself, rendering the disease dynamics highly non-stationary. Furthermore, every period is characterised by a number of recurrent attacks (up to 8 attacks per 24 h) which may be triggered by rapid eye movement (REM) sleep.

8.7 Extensions to the techniques

Single migraine attacks develop and resolve within a matter of hours, whereas the period between attacks spans multiple days. Applying Markov approaches to analyse the effects of acute anti-migraine drugs on both the length of the attack and the between-attack periods requires that both time frames are combined. This can be achieved by introducing one or more additional levels of progression within a single state of the original process. The resulting structure is referred to as a multi-dimensional or hierarchical Markov model. Applications have arisen in the field of chronic obstructive pulmonary disease (COPD) [18]. This application is also useful when analysing the effects of drugs that can be used both for prophylaxis and abortion of attacks [19].

In most multi-state applications it is assumed that transition rates do not vary between patients. That is, every transition is characterised by a distribution with fixed parameters. In reality, any parameter that determines disease progression displays natural variability and as a result, can be characterised by a statistical distribution. Application of mixed effects models offers an appealing approach to account for inter-patient variability [20, 21].

As the structures of Markov models become increasingly complex, parameter estimation becomes ever more difficult. Bayesian estimation methods can aid in two ways. Firstly, by using Markov Chain Monte Carlo (MCMC) algorithms which are readily available in software for Bayesian estimation, the problem of finding (approximations of) the likelihood function can be avoided. A disadvantage is that MCMC is generally less robust than analytic statistical methods. Secondly, if the data available are insufficient to estimate all parameters in the mixed model, it is still possible to obtain estimates using prior information. Other data sources can be exploited

When a disease is detected at an early stage, it may be more amenable to treatment. Systematic screening of a population can therefore be an effective way of reducing mortality from a disease.

Establishing an adequate screening policy for the disease requires knowledge of its natural history. The type of individuals to screen, and the time of the screen, should be chosen according to the risk of onset of the disease, whereas the intervals between successive screens should be chosen according to the risk of progression. Both risks may vary with characteristics of the patients and the risk of progression may vary with the current stage of the disease.

References

[1] H. Hassani and A. Ebbutt. Use of a stochastic model for repeated binary assessment. 1996, *Stat.Med.*, 15, 2617–2627.

[2] V.F. Cosson and E. Fuseau. Mixed effect modeling of sumatriptan pharmacokinetics during drug development: II. From healthy subjects to phase 2 dose ranging in patients. 1999, *J.Pharmacokinet.Biopharm.*, 27, 149–171.

[3] W.H. Visser, J. Burggraaf, L.M. Muller, R.C. Schoemaker, P.A. Fowler, A.F. Cohen, and M.D. Ferrari. Pharmacokinetic and pharmacodynamic profiles of sumatriptan in migraine patients with headache recurrence or no response. 1996, *Clin.Pharmacol.Ther.*, 60, 452–460.

[4] G. Geraud, C. Keywood, and J.M. Senard. Migraine headache recurrence: relationship to clinical, pharmacological and pharmacokinetic properties of triptans. 2003, *Headache*, 43, 376–388.

[5] G. Geraud, J. Olesen, V. Pfaffenrath, P. Tfelt-Hansen, R. Zupping, H.C. Diener, and R. Sweet. Comparison of the efficacy of zolmitriptan and sumatriptan: issues in migraine trial design. 2000, *Cephalalgia*, 20, 30–38.

[6] R.B. Lipton. Methodologic issues in acute migraine clinical trials. 2000, *Neurology*, 55, S3–S7.

[7] D.W. Lewis, P. Winner, and W. Wasiewski. The placebo responder rate in children and adolescents. 2005, *Headache*, 45, 232–239.

[8] P. Winner. Triptans for migraine management in adolescents. 2002, *Headache*, 42, 675–679.

[9] C.N. Barton, R.C. Braunberg, and L. Friedman. Nonlinear statistical models for the joint action of toxins. 1993, *Biometrics*, 49, 95–105.

[10] W.D. Penny and S.J. Roberts. Dynamic models for nonstationary signal segmentation. 1999, *Comput.Biomed.Res.*, 32, 483–502.

[11] J. Pardey, S. Roberts, and L. Tarassenko. A review of parametric modelling techniques for EEG analysis. 1996, *Med. Eng Phys.*, 18, 2–11.

[12] C.M. Crespi, W.G. Cumberland, and S. Blower. A queueing model for chronic recurrent conditions under panel observation. 2005, *Biometrics*, 61, 193–198.

[13] A. Bureau, J.P. Hughes, and S.C. Shiboski. An S-Plus implementation of Hidden Markov Models in Continuous Time. 2000, *J.Comp.Graph.Stat.*, 9, 621–632.

[14] T. Smith and P. Vounatsou. Estimation of infection and recovery rates for highly polymorphic parasites when detectability is imperfect, using hidden Markov models. 2003, *Stat.Med.*, 22, 1709–1724.