Applying an Artificial Neural Network to Warfarin Maintenance Dose Prediction

Idit Solomon MD[1,5], Nitsan Maharshak MD[3,5], Gal Chechik MSc[4], Leonard Leibovici MD[2,5], Aharon Lubetsky MD[5,6], Hillel Halkin MD[5,6], David Ezra MD[5,6] and Nachman Ash MD[5,6]

Departments of [1]Ophthalmology and [2]Internal Medicine H, Rabin Medical Center, Petah Tiqva, Israel
[3]Department of Internal Medicine T, Sourasky Tel Aviv Medical Center, Tel Aviv, Israel
[4]School of Mathematical Sciences and [5]Sackler Faculty of Medicine, Tel Aviv University, Ramat Aviv, Israel
[6]Department of Internal Medicine A, Sheba Medical Center, Tel Hashomer, Israel

Key words: neural network, warfarin, drug therapy, dose prediction, computer-assisted therapy

Abstract

Background: Oral anticoagulation with warfarin can lead to life-threatening events as a result of either over-anticoagulation or undertreatment. One of the main contributors to an undesirable warfarin effect is the need to adjust its daily dose for a specific patient. The dose is adjusted empirically based on the experience of the clinician, a method that is often imprecise. There is currently no other well-accepted method for predicting the maintenance dose of warfarin.

Objective: To describe the application of an artificial neural network to the problem of warfarin maintenance dose prediction.

Methods: We designed a neural network that predicts the maintenance dose of warfarin. Data on 148 patients attending a large anticoagulant clinic were collected by file review. Using correlational analysis of the patients' data we selected the best input variables. The network was trained by using the back-propagation algorithm on a subset of our data and the results were validated against the rest of the data. We used a multivariate linear regression to create a comparable model.

Results: The neural network generated reasonable predictions of the maintenance dose (r = 0.823). The results of the linear regression model were similar (r = 0.800).

Conclusion: Neural networks can be applied successfully for warfarin maintenance dose prediction. The results are promising, but further investigation is needed.

IMAJ 2004;6:732–735

For Editorial see page 770

Well-designed clinical trials have shown that warfarin is effective in the prevention of various hypercoagulable states [1]. However, its use is complicated by its narrow therapeutic window, the complex correlation between the dosage and the effect of the drug, and the numerous factors that affect its metabolism – namely, diet, renal or hepatic dysfunction, use of co-drugs, age and gender [2], weight, ethnicity [3] and genotype [4]. As a result, there is only a vague relationship between the maintenance dose or blood level of warfarin and the desired drug effect (prothrombin time or international normalized ratio) [5]. As many patients are given the drug for long periods (months or years), physicians frequently encounter either a suboptimal effect (recurrent thromboembolic events) or bleeding due to overdosage. The venous thromboembolism recurrence rate is approximately 7% at 6 months [6], and the reported frequency of bleeding ranges from 1% to 7% per year, depending on the indication for anticoagulant therapy and the bleeding classification used [7]. These complications pose both a clinical risk, which is sometimes life-threatening, and an economic burden [8].

During the period of our research, warfarin treatment was initiated with a loading phase. The patient was given relatively high doses of the drug for a few days, and the INR was frequently measured. Subsequent doses were then adjusted accordingly. When the INR reached the desired range, the maintenance phase began during which the maintenance dose was estimated empirically by the clinician on the basis of previous INR measurements. This trial-and-error approach has since been found to be inaccurate, with patients being out of their therapeutic range up to 20–50% of the time [9–12].

For many years now, researchers have been testing different methods for predicting the maintenance dose of warfarin. These can be largely divided into two categories: methods based on correlations between the loading phase data and the maintenance dose, and methods based on Bayesian regression-derived pharmacologic models [13]. Some are used in clinical practice [14] and others have been found to be efficient for empiric dosing [15,16], but none has been widely accepted as an alternative for empiric therapy.

In this study we used an artificial neural network to predict the maintenance dose of warfarin using demographic data and loading data. Artificial neural networks are algorithms that can be used to perform non-linear statistical modeling. They provide a novel alternative to logistic regression, the most commonly used method for developing predictive models for dichotomous outcomes in medicine. Neural networks offer a number of advantages: They can implicitly detect complex non-linear relationships between dependent and independent variables, and can detect all possible interactions between predictor variables. In addition, efficiency is improved by the availability of multiple training algorithms. Disadvantages include its "black box" characteristics, great computational burden, proneness to overfitting, and the empiric nature of model development. (For a detailed explanation of neural networks see the Appendix.)

INR = international normalized ratio

The aim of this study was to describe the artificial neural network model we constructed to predict the warfarin maintenance dose.

Patients and Methods

Patients characteristics

The files of 148 patients attending a large anticoagulant clinic in Israel from 1996 to 1997 were reviewed. Written informed consent was obtained from all patients. The study was conducted in accordance with the Helsinki Declaration.

Inclusion criteria were as follows: age above 18 years; warfarin treatment for longer than 1 month, three consecutive INR measurements within the target range no more than 0.5 units apart, complete data on dosage or INR levels in the loading period, warfarin treatment started no earlier than 1994, no change in co-drugs during the loading period, no congestive heart failure and liver disease, and normal INR and partial thromboplastin time before starting warfarin.

Data collection

The variables recorded from the files were selected on the basis of current medical knowledge [5,17], including: age, gender, height, weight, indication for warfarin treatment, smoking and drinking habits, co-drugs, target INR, loading data (warfarin doses, length of loading period, INR at end of loading period), and equilibrium state data (INR and maintenance dose).

Network inputs

The inputs for the model were selected from all possible data items by correlating each of them with the maintenance dose. The correlation was tested by a simple linear regression and Student's *t*-test.

Neural network model

The neural network model was built as a three-layered feed-forward network and was trained by the back-propagation algorithm using Matlab [18]. Data from the first 108 files were used as the training set, and data from the next 40 as the validation set.

Linear regression model

The linear regression model was built with the same inputs using the Statistical Analysis System (SAS) 8 [19].

Comparison of model predictions

The SAS was used to compare the ability of the models to predict the maintenance dose in the validation set. The three sets of maintenance doses (actual, neural network, linear regression model) were each divided into quartiles, and the match between the quartile number of the actual maintenance dose and the quartile number of the predicted subset was determined.

Results

Table 1 presents the characteristics of patients in the training and validation sets. There were no statistically significant between-group differences for any of the items.

The input variables chosen on the basis of the correlational analyses are presented in Table 2. The variables that showed the strongest correlation ($P \leqslant 0.01$) were age, loading ratio (total loading dose divided by the INR at the end of the loading phase, LOADI), and use (or not) of amiodarone as a co-drug (PROCOR).

The linear model, built on the data of the training group, is shown in equation 1. Each coefficient represents the impact of its variable on the maintenance dose.

Equation 1: MD = 3.26 – 0.31 * PROCOR + 0.28*LOADI – 0.032 * AGE ($r = 0.8$, $P = 0.0001$).

The four subsets of maintenance doses (by quartile) are shown in Table 3, and the comparisons of the models' performance are shown in Table 4. When the neural network was used, the predicted maintenance dose was in the same subset as the actual maintenance dose for 30 of the 40 patients in the validation group (75%), and one subset away for the other 10. Overall, the predicted maintenance dose was close to the actual maintenance dose, with a maximum deviation of one subset. In the linear model, the predicted maintenance dose was in the same subset as the actual

Table 1. Characteristics of patients in training and validation sets

Variable	Training set	Validation set
No. of patients	108	40
Male sex	63 (58.33%)	22 (55%)
Age (yrs)*	58.01 ± 15.6 (18–84)	64.4 ± 9.4 (42–80)
Amiodarone as co-drug	15 (13.9%)	4 (10%)
Total loading dose (mg)*	24.9 ± 8.2 (5–50)	24.8 ± 6.7 (13.8–37.5)
No. of loading days*	3.4 ± 0.7	3.3 ± 0.76
INR at end of loading phase*	2.59 ± 1.55 (1.14–10.8)	2.37 ± 1.02 (1.04–6.9)
Maintenance dose (mg)*	4.84 ± 2.26 (1.07–12.5)	5.13 ± 2.57 (1.43–10)

* Mean ± SD (range).

Table 2. Correlation of tentative input variables and maintenance doses for the whole set of patients

Variable	Correlation coefficient (r)	*P*
Age (yrs)	-0.2471	0.002
Body surface area	0.1689	0.04
Amiodarone	-0.2433	0.003
Target INR	-0.0059	0.943
Total loading dose	0.4225	<0.001
Loading period	0.193	0.019
INR at end of loading	-0.405	<0.001
Total loading dose/INR	0.7196	<0.001

Table 3. Subsets of maintenance dose by quartiles

	Maintenance dose by quartile			
Dose subset (mg)	1	2	3	4
Actual	MD <3.21	3.21 =MD <4.64	4.64 =MD <7.145	7.145 =MD
NN prediction	MD <3.42	3.42 =MD <4.535	4.535 =MD <5.76	5.760 =MD
Linear model prediction	MD <3.757	3.757 =MD <4.85	4.85 =MD <5.52	5.520 =MD

Table 4. Comparison of model performance

	Predicted subset					
	1		2		3	
Actual subset	NN	LM	NN	LM	NN	LM
1	8 (20)	6 (15)	5 (12.5)	3 (7.5)	0 (0)	1 (2.5)
2	0 (0)	1 (2.5)	6 (15)	5 (12.5)	0 (0)	4 (10)
3	0 (0)	1(2.5)	0 (0)	3 (7.5)	6 (15)	4 (10)
4	0 (0)	0 (0)	0 (0)	0 (0)	5 (12.5)	2 (5)

*Values represent number of patients (percent of the total validation group)

maintenance dose for 23 patients (57.5%), and one subset away for 15 (37.5%). Overall, the predicted maintenance dose was close to the actual maintenance dose in 95% of cases, with a maximum deviation of one subset.

Analysis of the findings in individual patients yielded no statistically significant differences between the actual and predicted maintenance doses (Student's paired *t*-test: $P = 0.275$ for network, $P = 0.216$ for linear model).

Discussion

Artificial neural networks can detect complex non-linear relationships between predictor variables and outcome [20]. In the domain of drug treatment, they have been applied successfully to predict cyclosporine dosage in patients after kidney transplantation [21], tacrolimus blood levels in liver transplant patients [22], and gentamicin serum concentrations from various input data over a range of patient ages and renal function [23]. In this study, we applied a neural network model to warfarin dose prediction. The behavior of warfarin at the maintenance level is difficult to predict in the clinical setting, and this problem is compounded by the drug's complex relationship with possible predictors. We assumed that training the neural network with examples would identify the relationships between the inputs of the model and the outcome. Comparison of the findings with the neural network model and a standard multivariate linear model showed a similar performance ($r = 0.823$ vs. $r = 0.800$).

We could identify only one previous study on the use of a neural network model in patients treated by warfarin [24]; the methods used to select predictor variables are described, but there is no mention of the performance of the model.

The selection of the input variables is a crucial step. We collected data on every factor reported to have a possible effect on warfarin metabolism or clearance. Factors found in published [17] and unpublished (Ash, Leibovici, Ezra) studies to have no correlation to warfarin's biological behavior, such as gender, were eliminated in advance. We then selected those that also had a significant correlation with the maintenance dose in our specific dataset for inclusion in the model. For example, warfarin clearance is known to decrease with age [17], and age had a significant correlation to maintenance dose ($P = 0.002$, Table 2), whereas body surface area had only a borderline correlation to maintenance dose ($P = 0.04$) and was therefore omitted.

Many medications affect warfarin metabolism, most of them by potentiating its action [25]. One of these is amiodarone, a common anti-arrhythmic drug. On the basis of our clinical experience and the strong correlation amiodarone showed with the maintenance dose, we included it too as a predictor variable.

As in previous warfarin dose prediction studies, the loading phase data had the strongest correlation to the maintenance dose. However, in most of the previous studies, one of the loading data items was set as a constant for all patients: either the length of the loading period or the daily dose of warfarin. Such rigidity is not typical of patient management in anticoagulant clinics. Therefore, to make our findings more applicable to the time the study was conducted, we kept the loading data flexible: the length of the loading period was 2, 3 or 4 days, and any warfarin dosage was allowed.

As shown in Table 2, the INR at the end of the loading phase was strongly correlated to the maintenance dose ($r = -0.405$, $P < 0.001$), as was the total loading dose ($r = 0.4225$, $P < 0.001$). The strongest correlation to the maintenance dose, however, was shown by their quotient LOADI, i.e., the total loading dose divided by the INR at the end of the loading phase ($r = 0.7196$, $P = <0.001$). LOADI is the index of patient resistance to warfarin and reflects all the information needed about the loading period. We used this value as an input variable, together with age and co-treatment with amiodarone.

Our results are encouraging, but they are not sufficiently accurate, especially at dosage extremes, to recommend the use of a neural network for predicting the maintenance dose of warfarin in a clinical setting. In addition, machine learning warrants further study because existing pharmacokinetic models for warfarin dose prediction are not ideal. To improve the performance of our model, we may need to integrate more biological information. Specifically, our model is missing two important data items that are known to affect warfarin dosage: the amount of vitamin K in the diet, and genotyping, which seems to be a promising method for identifying or predicting warfarin-sensitive patients.

Another drawback of our study was the relatively small sample size (n=148). Much larger series are usually used to train a neural network, although the number of patients included here is comparable to that used in other applied neural network models in medicine. Too few patients affect the ability of the model to generalize and perform accurately on completely different datasets [22]. Accordingly, like other warfarin prediction models, ours is not applicable to every patient, as the inclusion criteria limit its ability for extrapolation beyond the specific series studied. For example, the model will not perform well for patients with hepatic failure or congestive heart disease.

Finally, the everyday costs of applying a neural network model are low, and the mean output is for the software only. We conclude that the use of a neural network to predict the maintenance dose of warfarin is promising but requires more thorough investigation. Future efforts should be directed at identifying additional patient characteristics, such as the amount of vitamin K in the diet, ethnicity [3], genotype [4], and even pharmacokinetic calculations that can improve the model.

References

1. Hirsh J, Dalen JE, Deykin D, Poller L, Bussey H. Oral anticoagulants. Mechanism of action, clinical effectiveness, and optimal therapeutic range. *Chest* 1995;108:S213–46.
2. Oates A, Jackson PR, Austin CA, Cahnner KS. A new regimen for starting warfarin therapy in out-patients. *Br J Clin Pharmacol* 1998;46:157–61.
3. Absher RK, Moore ME, Parker MH. Patient specific factors predictive of warfarin dosage requirements. *Ann Pharmacother* 2002;36:1512–17.
4. Higashi MK, Veenstra DL, Kondo LM, et al. Association between CYP2C9 genetic variants and anticoagulation-related outcomes during warfarin therapy. *JAMA* 2002;287:1690–8.
5. Routledge PA, Chapman PH, Davies DM, Rawlins MD. Pharmacokinetics and pharmacodynamics of warfarin at steady state. *Br J Clin Pharmacol* 1979;8:243–7.
6. White RH. The epidemiology of venous thromboembolism. *Circulation* 2003;107(23 Suppl 1):I4–18.
7. Beyth RJ. Management of hemorrhagic complications associated with oral anticoagulant treatment. *Expert Opin Drug Saf* 2002;1:129–36.
8. Goldberg Y, Berliner S, Seligsohn U. Clinical and economic advantages of anticoagulant treatment in specialized clinics. *Harefuah* 2002;141:830–2 (Hebrew).
9. Duxbury BM. Therapeutic control of anticoagulant treatment. *Br Med J (Clin Res Ed)* 1982;284:702–4.
10. Copplestone A, Roath S. Assessment of therapeutic control of anticoagulation. *Acta Haematol* 1984;71:376–80.
11. Harries AD, Birtwell AJ, Jones DB. Anticoagulant control [Letter]. *Lancet* 1981;i:1320.
12. Doble N, Baron JH. Anticoagulation control with warfarin by junior hospital doctors. *J R Soc Med* 1987;80:627.
13. Lee C, Coleman RW, Mungall DR. Effect of using warfarin plasma concentrations in Bayesian forecasting of prothrombin-time response. *Clin Pharm* 1987;6:406–12.
14. Ryan PJ, Gilbert M, Rose PE. Computer control of anticoagulant dose for therapeutic management. *Br Med J* 1989;299:1207–9.
15. Poller L, Wright D, Rowlands M. Prospective comparative study of computer programs used for management of warfarin. *J Clin Pathol* 1993;46:299–303.
16. White RH, Mungall D. Outpatient management of warfarin therapy: comparison of computer-predicted dosage adjustment to skilled professional care. *Ther Drug Monit* 1991;13:46–50.
17. Mungall DR, Ludden TM, Marshall J, Hawkins DW, Talbert RL, Crawford MH. Population pharmacokinetics of racemic warfarin in adult patients. *J Pharmacokinet Biopharm* 1985;13:213–27.
18. Demuth H, Beale M. Neural Network Toolbox for Use with MATLAB-Userguide. Natick, MA: Mathworks Inc., 1993.
19. SAS Users Guide: Statistics. Cary, NC: SAS Institute Inc., 1982:119–37, 217–21.
20. Tu J. Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes. *J Clin Epidemiol* 1996;49:1225–31.
21. Camps-Valls G, Porta-Oltra B, Soria-Olivas E, et al. Prediction of cyclosporine dosage in patients after kidney transplantation using neural networks. *IEEE Trans Biomed Eng* 2003;50:442–8.
22. Chen HY, Chen TC, Min DI, Fischer GW, Wu YM. Prediction of tacrolimus blood levels by using the neural network with genetic algorithm in liver transplantation patients. *Ther Drug Monit* 1999;21:50–6.
23. Corrigan BW, Mayo PR, Jamali F. Application of a neural network for gentamicin concentration prediction in a general hospital population. *Ther Drug Monit* 1997;19:25–8.
24. Narayanan M, Lucas S. A genetic algorithm to improve a neural network to predict a patient's response to warfarin. *Methods Inf Med* 1993;32:55–8.
25. Wells PS, Holbrook AM, Crowther NR, Hirsh J. Interactions of warfarin with drugs and food. *Ann Intern Med* 1994;121:676–83.

Correspondence: Dr. N. Ash, Dept. of Internal Medicine A, Sheba Medical Center, Tel Hashomer 52621, Israel.
Phone: (972-3) 530-2636
Fax: (972-3) 535-2789
email: nachman.ash@sheba.health.gov.il

Appendix

A typical neural network consists of a series of nodes ("neurons") arranged in three layers: input, hidden, output. Each of the hidden and output nodes contains an "activation function." The hidden nodes allow the network to model complex non-linear mathematical relationships between the input (predictor) variables and a corresponding output (outcome) variable. Each node is connected to every node in the next layer by a "connection weight" – equivalent to the beta-coefficient in a regression model – and contains the "knowledge" acquired by the neural network after training. The connection weights in an untrained neural network are assigned random values. The major purpose of training a neural network is to estimate the optimal values of these connection weights: the weights are gradually adjusted to minimize the difference between the predicted output of the network and the known value of the outcome variable. Minimizing the error is the essence of the back-propagation algorithm, a popular learning algorithm used in the present study.

Capsule

Obesity, endoplasmic reticulum stress, and diabetes

Obesity is a major risk factor for type 2 diabetes, but how the two conditions are linked is not fully understood. Scientists now identify endoplasmic reticulum (ER) stress as a possible connection. The ER is a network of intracellular membranes involved in membrane and secretory protein synthesis and processing, and its function is sensitive to pathologic stress, such as changes in nutrient levels. Obesity triggers ER stress in liver and fat cells and this stress, in turn, disrupts insulin signaling. When placed on a high fat diet, genetically manipulated mice with elevated levels of ER stress developed peripheral insulin resistance at a higher rate than controls. The ER stress signaling pathway may be a useful target for the development of new therapeutics for type 2 diabetes.

Science 2004;306:457

E. Israeli