Journal of Pharmacokinetics and Biopharmaceutics, Vol. 20, No. 4, 1992

Commentary on "Neural Networks in Pharmacodynamic Modeling. Is Current Modeling Practice of Complex Kinetic Systems at a Dead End?"

Ronald A. Siegel[1]

Received May 4, 1992

INTRODUCTION

The emergence of neural networks in the past decade as an alternative means of pattern recognition and prediction has wide implications, as Peter Veng-Pedersen and Nishit Modi discuss, concerning the nature of scientific and engineering endeavors. On the one hand, it is evident that a wide variety of patterns and relations can be recognized readily by a neural net, but would take inordinate amounts of time to understand by the conventional scientific method. Yet, the following question can be asked: "If I can train a neural net to make correct decisions or estimations for a particular set of patterns or phenomena, have I learned anything that I could apply to other, related phenomena?" In this commentary I first discuss these issues from a variety of angles, without drawing any specific conclusions regarding the superiority or inferiority of the neural net approach compared to the more conventionally accepted methods of scientific inquiry. I then raise and discuss some points regarding Veng-Pedersen and Modi's specific application.

Before proceeding, however, I wish to constructively criticize the premise that an article, such as was written by Veng-Pedersen and Modi, belongs at this juncture in the "Perspectives in Pharmacokinetics" section of this Journal. It is clear that neural nets are interesting and should be investigated as one way of approaching pharmacodynamics. However, a

[1]Department of Pharmacy and Pharmaceutical Chemistry, University of California, San Francisco, California 94143-0446.

0090-466X/92/0800-0413$06.50/0

perspective is usually attained after considerable experience. There is as yet no published peer-reviewed example applying neural nets to pharmacodynamics, in which the whole methodolgy is laid out for study and criticism. While the present article describes neural nets to some degree and shows the performance of a particular construction with a particular data set, it is not complete as a research article. I would have preferred to see several research articles appear before a Perspectives article was written, because at that time more experience would have been accrued, and comparisons could be made with more traditional pharmacodynamic approaches. In this light, it seems presumptuous in the title to ask: "Is Current Modeling Practice of Complex Kinetic Systems at a Dead End?".

SOME THOUGHTS ABOUT NEURAL NETS

1. Humans, as well as many animals, are not "hard-wired" to perform all tasks they will encounter. Real neural systems are plastic in the sense that they can alter their topology and strength of their connections in response to training, as in artificial neural nets. Without this ability, there could be no learning, and every task would have to be approached naively each time it is required.

Humans and other animals are remarkably adept at performing many complex tasks. No human makes conscious trajectory calculations when picking up a pencil, playing a musical instrument, or throwing a football. Yet it is clear that some type of computation is being performed. (The computational skill of a quarterback throwing a "bomb" downfield in the presence of a defense is astounding.)

There is nothing intrinsically wrong with training a neural net to execute complex tasks providing, of course, that it works. If a black box can be successful, then that's fine. After all, the trained human being performs in many ways as a black box. Since we do not require understanding of precisely how a human performs a certain set of tasks, there may be no need to demand such an understanding of a machine.

2. Neural nets are by no means the only example of a computational technique whose application can generate correct results, but without yielding insight to the investigator. As another example, consider the numerical solution of partial differential equations (PDEs). Textbooks abound which contain elegant theories and solutions of linear PDEs over simple geometric domains. With some effort, one can obtain an intuitive understanding of the behavior of systems governed by the PDEs in the simple geometries, such that predictions can be made for new situations such as a change in

boundary conditions. Often, real problems can be approximated by these more simple, solvable systems. (Hence the physicists' joke about the "spherical cow.") However, when the governing PDEs are highly nonlinear, or the geometry cannot be well approximated by a simpler case, one must resort to numerical techniques. While it is true that techniques such as boundary layer theory and local linearization can provide insight into local behavior, they do not provide global predictive capability for new geometries or boundary conditions. Hence, scientists and engineers already have experience with techniques which provide, in common with neural nets, useful results without yielding much insight.

3. Although the previous thoughts argue that neural net techniques have a place in science, it should not be inferred that neural nets should replace more accepted scientific *modi operandi.* Let us imagine that high-speed computers and neural net techniques were available in the late 19th century. At the time much effort was being spent trying to understand the temperature dependence of the spectrum of blackbody radiation. The results could not be predicted by classical physical theories. A scientist back then could have decided that the process was too complex to comprehend, and instead could have trained a neural net on the blackbody data, and yielded an accurate predictor of radiation intensity at any temperature and frequency.

This success would have been empty, however, compared to what actually happened. As is well known, Planck invented the quantum description of radiation to provide an explanation for the experimental observations. The development of quantum theory that followed from Planck's insight, along with the technological spinoffs (including semiconductors), might never have occurred if physicists were satisfied with a neural net's predictions. (In this fable we ignore the fact that the semiconductors which are integral to today's high-speed computers would have been needed to implement the neural net for the blackbody studies.)

4. A conclusion that might be drawn is the following: Neural nets may be very useful for developing relations between inputs and outputs for short-term needs. Nevertheless, they should not be seen as replacing theoretical science. Neural nets, in their present form, are unable to develop theories such as quantum mechanics or link-model descriptions relating kinetics of drug concentration and drug effect. These theories provide great simplification and insights which can be used to design future experiments. Neural net outputs merely provide answers to specific questions. Therefore, I propose that neural nets are "tactical" components, whereas more traditional theoretical methods are "strategic" components in the armamentarium of science.

TECHNICAL COMMENTS

Careful inspection of Fig. 4 of Veng-Pedersen and Modi's paper shows that the neural net predictions in the posttraining phase (210–320 min) are somewhat biased. Most of the observed heart rate data lies below the predicted curve. This may seem a minor effect, but it is consistent. In addition, the data level off finally around a baseline value of 250 beats/min, while the neural net predicts a continuing increase in heart rate. The latter effect might be explained by noting that the net has no experience with a leveling off, i.e., in the training phase, drug levels never decrease enough for the effect to return to baseline. Thus, the net extrapolates in the posttraining phase, but in an incorrect manner.

A related issue concerns the meaning of the *relative prediction accuracy* (RPA). It is stated in the paper that an average RPA of about 78%, as obtained for the data sets studied, may be close to what is achievable in practice. I see no reason why expected RPA values of 100% could not be attained. The RPA seems to me to be a measure of *bias*. If the predicted curve goes in the "middle" of the scattered data, and no detectable serial correlations in the error exist, then the neural net predictor should probably be considered adequate. If the statistics of the scatter in the training and posttraining phases are the same, then such a predictor would have an expected RPA of 100%. In this case, one should expect that a diminished RPA value would reflect either bias or predictor inconsistency.

The development of useful measures of goodness-of-fit for neural net predictions should be a fruitful area of research. In particular, the statistical properties of measures such as the RPA could be investigated. Offhand, the RPA should have an F-type distribution. The difficulty in making a definite statement, however, will lie in specifying the number of parameters involved in the modeling process since, by definition, the model is transparent to the analyst.

CONCLUSIONS

Veng-Pedersen and Modi have begun some very interesting investigations. The neural net technique as applied to pharmacodynamics appears to be a promising area of research. However, the results are far from conclusive at present. I believe it is too early to state that the neural net techniques will supplant the more traditional modeling techniques, particularly for short-term goals. Moreover, I hope that the former do not *completely* replace the latter, as the latter are what ultimately lead to the insights which are the main event in scientific discovery and progress.