Systems Engineering Models of Human-Machine Interaction

Systems Engineering Models of Human–Machine Interaction

Series Volume 6

William B. Rouse
University of Illinois, Urbana, Illinois

NORTH HOLLAND
New York • Oxford

Elsevier North Holland, Inc.
52 Vanderbilt Avenue, New York, New York 10017

Distributors outside the United States and Canada:

Thomond Books
(A Division of Elsevier/North-Holland Scientific Publishers, Ltd.)
P.O. Box 85
Limerick, Ireland

Library of Congress Cataloging in Publication Data

Rouse, William B.
Systems engineering models of human–machine interaction.
(North Holland series in system science and engineering; v. 6)

Bibliography: p.
Includes index.
1. Man–machine systems—Mathematical models. 2. Systems engineering—Mathematical models. I. Title.
TA167.R68 620.8'2 79-25437
ISBN 0-444-00366-5

Desk Editors Michael Gnat; Michael Cantwell
Design Series
Art rendered by Vantage Art, Inc.
Production Manager Joanne Jay
Compositor Comptype
Printer Haddon Craftsmen

Manufactured in the United States of America

Preface

In our technological society, humans are increasingly interacting with machines. The designers of these machines would like to consider this human–machine interaction in the same quantitative manner in which they pursue much of the rest of the design process. To this end, mathematical models of human–machine interaction are needed. A wide variety of such models is presented here.

This book was written with several goals in mind. First of all, one objective was to emphasize the current state of the art rather than provide a thorough historical perspective. Based on this goal, well over 80% of the references cited have publication dates of 1970 or later. However, readers interested in earlier works will find excellent sources on this material noted throughout this book.

A second goal was to provide a treatment of a highly mathematical topic while avoiding calculus, differential equations, Laplace and Fourier transforms, and so on. Thus, the only mathematical prerequisites for this book include basic algebra and probability theory. Although representing all of the models algebraically does result in a few topics (but not too many) being covered rather tersely, it is hoped that the purely algebraic treatment will make the material covered accessible to many more readers. Perhaps those readers who find this book stimulating will go on to study those works that require more extensive mathematical prerequisites.

A third goal was to include basic tutorials on the modeling methodologies of interest and thus avoid requiring the reader to consult other basic sources. Furthermore, along with the tutorials, fairly complete examples of applications are discussed and a breadth of other applications briefly reviewed. The choice of the characteristics noted in this paragraph was motivated by a desire to employ this book as a primary text for a graduate course on human–machine systems that typically includes both engineering and psychology students.

For the most part, this book is based on lecture notes used for this course, which is offered in Industrial Engineering at the University of Illinois, as well as on lecture notes for a short course on modeling human–machine interaction given outside the university. In teaching these courses, I have found that the type of material presented in this book is nicely complemented by having students pursue a series of small design projects in which they have to choose among the various available models, resolve measurement problems, and so on. This approach leads students to realize, from experience, the ways in which models are particularly useful. This end is not served as well by specific exercises in which students primarily learn to manipulate equations. For this reason, a set of such exercises is not included in this book. However, if such material is desired, the numerous texts cited throughout this book are more than adequate sources.

To a great extent, this book is also based on the results of the author's interactions with colleagues at Illinois, elsewhere in the United States, and in other countries as well. I am truly indebted to these individuals, who have contributed greatly to the lines of thought formulated in this book. I am also most grateful to Carolyn Robins for her editorial assistance in preparing the manuscript.

William B. Rouse

Urbana, Illinois

Contents

Chapter 1
Introduction

Humans interact with machines in many ways. Many people drive automobiles. Some people repair automobiles. Others fly airplanes, work in nuclear power plants, or work with computers. In fact, it is not unreasonable to claim that humans, in general, are increasingly interacting with machines.

As designers of machines, engineers should be concerned with the way in which humans interact with machines. This issue is of importance for several reasons. First, one wants to ensure that machines are safe to operate. Second, one would like to know how the human's abilities and limitations affect the performance of the overall human–machine system. With this knowledge, the machine can perhaps be designed so as to complement the human's characteristics. Another reason for being concerned with how humans interact with machines involves the problem of job satisfaction. We should like to design machines so that interacting with them is satisfying, or at least certainly not demeaning. Although job satisfaction and safety are not concerns of this book, we mention these topics here in recognition of their importance.

This book is concerned with the performance of human–machine systems. Adopting a typical engineering point of view, our goal is to develop methods of analysis that allow one to *predict* performance. This goal should be contrasted with that of trying to *measure* performance. Were measurement our goal, then we would devote this book to the various considerations surrounding the topic of experimental design.

To predict the performance of a human–machine system, we require some representation of the system that allows us to determine how independent variables affect dependent variables. To represent a human–machine system, we shall have to depict both human and machine behavior in compatible terms. For a variety of obvious reasons, it is appropriate to represent human behavior in machine-like terms, as opposed to vice versa. Further, while a good verbal description of the human's tasks, abilities, and limitations is certainly an essential first step towards developing a representation of human behavior, such a description is usually inadequate in that it will only allow *qualitative* statements about the performance characteristics of the human. Our goal is to provide *quantitative* predictions of human performance.

The engineering approach to obtaining quantitative performance predictions for almost any problem is to develop a mathematical model of that problem in terms of how relevant inputs (independent variables) affect interesting outputs (dependent variables). If this approach is successful, and it is not always successful, then the resulting mathematical model can be quite useful in several ways.

The Purposes of Models

In general, there are four major uses of models. First of all, the modeling process is itself beneficial. Developing a mathematical model requires a very organized and thorough pursuit of all the issues surrounding a problem. This is especially true when one gets to the stage of writing a computer program that incorporates the resulting model. At this point, one often finds that various parameters are undefined or perhaps immeasurable. Once the model is sufficiently defined to allow the computer program to produce predictions of performance, then it is not unusual to find the predictions to be ridiculous when compared to the performance of the real system. In this situation, one has obviously overlooked some crucial aspect of the real system. Then, the modeling process iterates and the model is updated. This iterative process of model development and testing provides many insights with respect to the system being studied. These insights are valuable even if the model's predictions are actually never used.

A second use of models is to provide succinct explanations of data. For example, a study of the learning of motor skills may result in a learning curve of rms (root-mean-squared) error versus time. While this tabulation of error versus time summarizes the results of the study, a much more succinct explanation might be found using a single parameter to characterize learning rate within a mathematical model of learning of motor skills. Thus, models can be used to aggregate various plots and tabulations into a few behaviorally meaningful parameters and, in that way, allow much clearer comparisons among tasks and experiments while also providing much simpler "rules of thumb" for systems designers. We shall return to this point repeatedly throughout this book.

Models are also useful for designing experiments. It is not unusual for a model to have many free parameters. Estimating these parameters can present data collection problems. These problems can be lessened if one tests the model to determine the sensitivity of performance predictions to parameter variations. If the model's predictions are fairly insensitive to some of the parameters, then one can assume "reasonable" values for these parameters and avoid investing the effort necessary for estimating them. Similarly, if one is contemplating an empirical study of human performance, one can use the results of a sensitivity analysis with a model to determine which parameters should be varied in the empirical study. This may sound like an unusual procedure. If one had a good model of human performance in a particular task, then why would one run an empirical study of that task? On the other hand, if one did not have a good model, how could one use a model to determine which parameters have the greatest effect on performance? Quite simply, one can often use an approximate model to obtain a feeling for the sensitivity of performance to various parameters. The value of this approach is dependent on how reasonable the model is regardless of the fact that it has not yet been experimentally validated. Using models in this way is part of the art of engineering.

Finally, models are useful in terms of the quantitative predictions that they produce. While this aspect of model usage is often given too much emphasis, relative to the other purposes of models, quantitative predictions are nevertheless useful in the process of designing systems. For example, it is quite important to know how a pilot and aircraft will interact *before* one builds the aircraft. Another example, which we discuss at some length in Chapters 4 and 7, involves embedding predictive models in computer programs in a manner that allows the computer to "understand" the user.

Most of the discussions throughout this book will focus on models in terms of the performance predictions that they produce. However, we want to emphasize here in our early discussions that one should look at such predictions as only one of the benefits of modeling.

The Modeling Process

The first step in the modeling process involves defining the problem. This includes a statement of the phenomena of interest as well as a choice of performance measures. Thus, for example, one might choose to study the abilities of air traffic controllers to detect potential midair collisions. One might designate probability of detection and time until detection as performance measures. As another example, one might decide to investigate the abilities of maintenance personnel in troubleshooting tasks. Appropriate performance measures for this task would be probability of correct diagnosis, time needed for diagnosis, and cost of diagnosis.

While this process of defining the problem may seem quite straightforward, this is far from true. It is not unusual to realize suddenly that one is working on the wrong problem. Also, it is not uncommon to discover that the performance measures initially chosen are inappropriate. This latter difficulty can frequently be attributed to the fact that performance measures are often only indirect indices of desirable system characteristics. For example, vehicle ride quality is difficult to quantify unless one uses indirect measures such as rms amplitude of vibrations, and so on. Using such indirect measures increases the likelihood of choosing inappropriately. Thus, one should not view problem definition as a necessarily easy portion of the modeling process. Indeed, it is probably the most troublesome step in that it is difficult to obtain training (e.g., by reading books such as this) that will give one the ability to define problems insightfully and appropriately. Experience appears to be the key ingredient in this step of the modeling process.

With the problem defined and candidate performance measures chosen, the next step of the modeling process involves representing the problem. What are the system's inputs? What are its outputs? How do inputs affect outputs? How are the performance measures of interest affected by the system's variables? Representing the problem involves providing answers to these questions.

Up to this point, we have not been concerned with formalizing our conception of human–machine systems. However, some formalization is now necessary if we are to avoid confusing terminology. Figure 1.1 depicts a basic human–machine system. With this figure, we want to make three important clarifications. First of all, unless otherwise specified, the term "system" will refer to the overall combination of human and machine. Second, the term "machine" will refer to everything other than the human. Third, and finally, this figure clearly illustrates the notion that the machine's outputs are the human's inputs and the human's outputs are the machine's inputs.

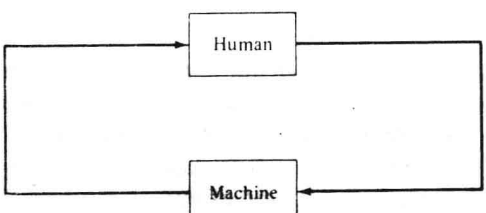

Figure 1.1. Basic human–machine system.

This last idea is central to much of this book. The output of the human's perception/decision/action process may, for example, be the rotation of an automobile's steering wheel. The angle of rotation is the input to the automobile. The machine (i.e., the automobile) converts this angle into a change in the vehicle's path. The human perceives this change and then updates the earlier decision/action regarding the rotation of the steering wheel, and so on.

In terms of representing this problem, one could employ various engineering tools to model how steering wheel inputs affect automobile outputs. However, this book is *not* concerned with this aspect of modeling. Instead, we are concerned with how the path that an automobile is following affects the steering wheel rotations that the human produces. Thus, throughout this book, we shall assume that a model of the machine is available and our discussions will be aimed at developing a model of the human. Of course, as we shall later discuss at some length, the human's behavior is highly influenced by the particular machine involved in the interaction; thus one cannot view the models of human and machine as independent of each other.

Succinctly then, we can view problem representation as the process of developing a relationship between the inputs and outputs of the human. As discussions throughout this book will indicate, the appropriate type of representation depends on a variety of factors. In Chapter 7, we shall consider the problem domains appropriate to the various representation methodologies discussed in this book.

Once a representation of the problem has been formulated, the next step of the modeling process is calculation of performance. Consider the following example. Suppose the input to the human x results in the human producing output $y = a + bx$. Further, assume that the machine transforms its input y into output $x = c + dy$. Finally, the performance measure might be x^2 where the goal is minimization of x^2. Calculation of performance involves determining the simultaneous solution of $y = a + bx$ and $x = c + dy$ and then calculating x^2.

Although this very oversimplified example illustrates the concept of calculating performance, it by no means depicts the usual level of difficulty of such calculations. In fact, it is not unusual to be unable to solve analytically the set of equations embodying the models of human and machine. In such situations, one has to resort to either approximations or simulation.

Simulation solution involves developing an analog and/or digital computer program that emulates the set of equations of interest and allows one to calculate various properties of the resulting solutions. Since human–machine systems problems frequently have probabilistic aspects, this calculation is often of a statistical nature. Thus, one must replicate the simulation solutions many times to obtain a sufficient number of samples of performance measures to be able to estimate average performance within reasonable confidence limits. Although simulation approaches to calculating performances certainly work, one usually attempts to avoid simulation unless analytical solutions, perhaps with appropriate approximations, are not possible. This is simply due to the fact that simulation solutions are typically more time-consuming and costly than analytical solutions. Nevertheless, as will be discussed at various points throughout later chapters, many realistic and important problems have to be solved using simulation.

The next step in the modeling process involves experimental validation of the model. There are two parts to this process. First, since most (but not all) models typically have several free parameters, usually reflecting some behavioral assumptions (e.g., reaction time), one may have to experimentally determine the values of these parameters. This involves adjusting the parameters until the performance of the model matches actual system performance as closely as possible.

There are many misconceptions about this performance-matching process. Some researchers (i.e., nonmodelers) have claimed that with two or three free parameters they can make model performance match virtually any empirical results. This is patently absurd. For example, if the actual system is such that $y = x^2$ and one chooses to use $y = a + bx$ as a model, no amount of manipulation of the constants a and b is going to match this model to the real system over a reasonable range of x. The important point to note here is that the structure of the model will preclude certain types of behavior regardless of the values of its parameters. Thus, one should be careful in deriving a model's structure and, for the most part, not apologize about having to estimate the model's parameters.

On the other hand, some researchers (i.e., modelers) tend to let parameterization get out of hand. Occasionally someone will use 10, 15, or 20 free parameters to match a single performance measure (e.g., rms error). This clearly presents methodological problems. However, overparameterization also presents difficulties in terms of interpreting the meaning of parameter variations. This subverts some of the purposes of modeling, namely, providing succinct explanations of data and providing assistance in designing experiments.

The second part of the experimentation process involves using the model to predict performance and then empirically determining how close the predictions are to actual occurrences. Of course, if the problem of interest has probabilistic aspects, then one does not expect a model's predictions to be perfect. The essential question is this: Does the model provide more predictive ability than one could obtain by simply using the mean of past performance as a predictor of future performance? The answer to this question can be quantified in terms of the percent of the variation about the mean that is explained by the model. One useful rule of thumb here is the following: The percent of the variance that a model will be able to explain (predict) is inversely related to the robustness of the human–machine system environment being studied. Thus, for example, one can probably model variations of detector time in a pattern recognition task to a greater extent than one can model variations of task time in problem solving situations (Rouse and Rouse 1979).

The results of experimentation lead to the comparison step of the modeling process. Although it is difficult to separate experimentation from comparison, as evidenced by the issues discussed in the last few paragraphs, we nevertheless want to single out some very important aspects of the comparison process. First of all, the distinction between behavior and performance should be noted. Within this book, the term "behavior" will be used to refer to *what* the human does whereas the term "performance" will refer to *how well* it is done. In a tracking task, for example, the human's behavior is the specific time history of control movements produced whereas the performance is the rms tracking error that results.

Since a variety of patterns of behavior might result in the same performance, it is very much easier to develop models to predict performance than it is to develop models to predict behavior. For many engineering applications, performance predictions are all that is necessary. When such is the case, then comparisons of the performance of the model with empirical measurements of performance are sufficient to validate the model. However, such validation does *not* allow one to infer that the model's behavior matches human behavior. If one is interested in behavior as well as performance, then one should make validating comparisons between the behavior of the model and the behavior of the human. Although this principle does not have to be followed religiously, it should not be flagrantly violated.

Since a model that can accurately predict behavior will also be able to accurately predict performance (but not vice versa), a behavioral model is much "stronger" in the sense that it more completely describes the human as he performs the task of interest. This point has been convincingly argued by Gregg and Simon (1967). They contrast process (behavioral) models and statistical (performance) models and defend in detail the notions briefly outlined in these last two paragraphs.

Besides deciding whether comparisons will be based on behavior or performance, one also must choose the range over which validation is to be attempted. Quite often the particular application motivating the modeling effort will dictate the range over which one should vary the conditions for which the model is being tested. One should choose the range of testing conditions so that later predictions will be interpolations rather than extrapolations. If unforeseen situations later cause one to have to use extrapolations, then one is much better off if a behavioral model is available, especially if one is only interested in performance predictions for the extrapolated conditions. The reason for this is obvious. The more fine-grained the validation process is, the more reasonable extrapolations will be, particularly if these extrapolations are not as fine-grained as the validation.

Summarizing our discussions of the modeling process thus far, the following steps of the process are important to note:

1. Definition
2. Representation
3. Calculation
4. Experimentation
5. Comparison
6. Iteration

With the exception of iteration, we have discussed all of these steps of the modeling process. The iteration step has been added at this point to emphasize that the modeling process is really not as straightforward as this introductory chapter may have led one to believe. One typically goes through a seemingly never-ending process of refinement whereby model inadequacies are eliminated and the range of a model validity is extended. Occasionally as one is trying to eliminate inadequacies and extend the range of validity, one finds that a whole new representation methodology is needed. Several transitions of this type will be considered throughout this book. In general, iteration results in the modeling process never terminating, basically because the modeling process is synonymous with the process of knowledge acquisition and organization whose inherent goal is growth and change.

The Use of Analogies

One of the most powerful problem solving methods of science and technology is the use of analogies. Basically, this involves viewing a new problem as if it were an old problem for which one is likely to know the solution or, at least, possess considerable insight. For example, one might view the central nervous system as an electrical circuit and then employ various circuit analysis methods in an attempt to understand the central nervous system. This analogy might prove quite satisfactory until one must deal with the chemical nature of central nervous system activity. Then, the electrical circuit analogy might have to be modified or perhaps replaced.

The iterative process of adopting, modifying, and replacing analogies is central to science and engineering. When an analogy within a particular area of research gathers a sufficient number of adherents, it is often termed a *paradigm* (Kuhn 1962). The emergence of paradigms is a focal point of study for those who research the history of science and engineering. One especially important notion that can be gleaned from this research is that the emergence, and especially the maintenance, of a paradigm can usually partially be attributed to a social consensus among researchers rather than to purely technical considerations (Ziman 1968). This phenomenon results in tremendous inertia in the sense that the need for a new paradigm often has to be overwhelming before an old paradigm is replaced.

This book considers a variety of analogies that are of use when analyzing human–machine systems. Of all the analogies that we will consider, perhaps the only one that has achieved the status of a paradigm is the servomechanism analogy. The basic idea is that the human acts as an error-nulling device (e.g., a thermostat) when driving an automobile, flying an airplane, or doing just about anything. Because of Norbert Wiener's efforts in this area (Wiener 1948), the servomechanism analogy is often referred to as the cybernetic paradigm. However, as we shall discuss in Chapter 7, it is perhaps unfortunate that cybernetics often is viewed so narrowly.

The servomechanism analogy or paradigm has proven to be quite useful because of the great number of control theory methods available for analysis of systems that can be represented as servomechanisms. Chapter 3, in particular, is devoted to discussing some of these methods. However, while the servomechanism paradigm is still predominant in the human–machine systems area, several new analogies have emerged that are also quite useful.

The need for new analogies has been precipitated by the increasing use of automation. The increasing use of computers to perform control tasks has resulted in the human's role becoming more like that of a monitor and supervisor. In such a role, the human can have responsibility for more tasks. Furthermore, as a backup for the computer, the human has to help in detection and diagnosis of system failures. Viewing the human as a monitor, supervisor, and diagnostician leads to three new analogies: ideal observer, time-shared computer, and logical problem solver. The ideal observer analogy, discussed in Chapter 2, is not really that new. However, the recent emphasis on the human's role as a monitor has increased the potential usefulness of this analogy. The usefulness of the time-shared computer analogy, considered in Chapter 4, is related to the desire to represent the human's abilities to supervise multiple tasks. This analogy has continually been embellished as the technology of time-shared computing has developed. The logical problem solver analogy, examined in Chapter 5, is relatively new. Its development reflects a recognition of the importance of understanding the human's role in systems where the ultimate responsibility for system failure is the human's.